

DeS3: Adaptive Attention-Driven Self and Soft Shadow Removal Using ViT Similarity

Yeying Jin¹, Wei Ye², Wenhan Yang³, Yuan Yuan², Robby T. Tan¹

¹National University of Singapore

²Huawei International Pte Ltd

³Peng Cheng Laboratory

e0178303@u.nus.edu, yewei10@huawei.com, yangwh@pcl.ac.cn, yuanyuan10@huawei.com, robbytan@nus.edu.sg

Abstract

Removing soft and self shadows that lack clear boundaries from a single image is still challenging. Self shadows are shadows that are cast on the object itself. Most existing methods rely on binary shadow masks, without considering the ambiguous boundaries of soft and self shadows. In this paper, we present DeS3, a method that removes hard, soft and self shadows based on adaptive attention and ViT similarity. Our novel ViT similarity loss utilizes features extracted from a pre-trained Vision Transformer. This loss helps guide the reverse sampling towards recovering scene structures. Our adaptive attention is able to differentiate shadow regions from the underlying objects, as well as shadow regions from the object casting the shadow. This capability enables DeS3 to better recover the structures of objects even when they are partially occluded by shadows. Different from existing methods that rely on constraints during the training phase, we incorporate the ViT similarity during the sampling stage. Our method outperforms state-of-the-art methods on the SRD, AISTD, LRSS, USR and UIUC datasets, removing hard, soft, and self shadows robustly. Specifically, our method outperforms the SOTA method by 16% of the RMSE of the whole image on the LRSS dataset.

Introduction

Shadows can be categorized into hard shadows, soft shadows, and self shadows (Salvador, Cavallaro, and Ebrahimi 2004; Huang and Chen 2009). When an object blocks beams of light, shadows are formed on surfaces or nearby objects, which are referred to as hard and soft shadows. Hard shadows have sharp boundaries, while soft shadows gradually transition from shadow to non-shadow regions without any distinct boundaries. Self shadows occur when a portion of an object obstructs beams of light, resulting in shadows being cast on the object itself. Removing all types of shadows is intractable due to ambiguities between shadow and non-shadow regions. Soft and self shadows are particularly challenging to identify and remove compared to hard shadows.

To remove shadows, most of the methods use an off-the-shelf shadow detection method (Zhu et al. 2018) or user interaction (Gryka, Terry, and Brostow 2015) to obtain binary shadow masks. However, these binary shadow masks are




Figure 1: The results of SOTA supervised method (Wan et al. 2022) and weakly-supervised method (Liu et al. 2021b) in removing (a) self shadow, (b) soft shadow, and (c) hard shadow. Our DeS3 can preserve meaningful objects (duck, paper, bollard, etc.) during the reverse sampling, and achieve better shadow removal results.

problematic to acquire, particularly for soft and self shadows as shown in Fig. 3. Moreover, current shadow removal methods fail to handle self shadow images since obtaining the ground truth for outdoor self shadows is intractable.

The weakly-supervised methods (Le and Samaras 2020; Liu et al. 2021b) require binary shadow masks to distinguish shadow and shadow-free patches. These patches from the same images can reduce the domain gap of shadow and shadow-free domains. However, cropping patches is time-consuming (Liu et al. 2021b) and synthesizing shadows on the shadow-free regions cannot represent real-world shadow images. Moreover, these weakly-supervised methods require shadows to be homogeneous, which is often not the case, particularly for self shadows and soft shadows.

Existing unsupervised methods (Hu et al. 2019b; Liu et al. 2021a; Jin, Sharma, and Tan 2021) are GAN-based (Zhu et al. 2017), which require unpaired shadow and shadow-free images for training. They largely rely on the statistical similarity between images from the two domains (Le and Samaras 2020) (shadow and shadow-free domains). Unfortunately, once the two domains are statistically different, these methods produce hallucination/fake contents (Zhu et al. 2018) and suffer from unstable training.

Most SOTA methods rely on masks, but obtaining masks for self and soft shadows is intractable. Fixed attention in (Jin, Sharma, and Tan 2021) for self-shadows are unreliable. Hence, we propose an adaptive attention to progressively focus on self-shadows. For this reason, we employ the diffusion. Unlike other generative models, diffusion works progressively.

In this paper, we introduce *DeS3*¹, a diffusion-based method that removes hard, soft, and self shadows from a single image using adaptive attention and ViT similarity. “De” means “DeShadow”, “S3” means three types of shadows. Each step in the diffusion process brings about changes in our adaptive attention, as shown in Fig. 4, enabling DeS3 to effectively eliminate self and soft shadows that lack clear boundaries. Moreover, using the ViT similarity, our DeS3 can preserve object structures even when the object is partially occluded by self and soft shadows. To guide the reverse sampling in estimating the structural features, our basic idea is to inject the pretrained features from DINO-ViT (Caron et al. 2021). Specifically, since the keys’ self-similarity (Shechtman and Irani 2007) contains object structures (Tumanyan et al. 2022), they are utilized to compute the ViT similarity loss between the denoised samples and the input shadow condition. As shown in Fig. 1, our method can remove shadows and retain various objects (duck, paper, bollard, etc.).

In summary, we make the following contributions:

1. We introduce DeS3, the first shadow removal network, that performs shadow removal robustly on hard, soft and self shadows from a single image.
2. Our DeS3 does not require shadow masks from the dataset or a shadow detector. Through the adaptive attention in the progressive diffusion process, DeS3 learns to adjust to all types of shadow regions, particularly self-shadow regions.
3. To maintain object structure features, we integrate the ViT similarity loss into the reverse sampling. ViT’s features are independent from shadows and thus can extract structures more robustly.

Comprehensive experiments on the SRD, AISTD, LRSS, UIUC and USR datasets demonstrate that DeS3 outperforms the state-of-the-art methods, particularly on self and soft shadows. Our DeS3 outperforms ShadowDiffusion (Guo et al. 2023b) by 16% error reduction (from 3.49 to 3.01) in terms of RMSE on the overall areas of the LRSS dataset.

Related Work

Different image priors have been explored for single image shadow removal, e.g., modeling of illumination and color (Finlayson et al. 2005; Jin et al. 2023a), image regions (Guo, Dai, and Hoiem 2012; Vicente, Hoai, and Samaras 2017; Jin et al. 2022), image gradients (Gryka, Terry, and Brostow 2015; Jin et al. 2023b). However, traditional shadow removal methods may produce unsatisfactory results. Recently, supervised learning-based shadow removal methods (Chen et al. 2021; Fu et al. 2021; Liu et al. 2023a,b)

have shown promising performance. DeshadowNet (Qu et al. 2017) removes shadows in an end-to-end manner. DSC (Hu et al. 2019a) captures global and context information from the direction-aware spatial attention module. SP+M-Net (Le and Samaras 2019) and SP+M+I-Net (Le and Samaras 2021) remove shadow using image decomposition. CANet (Chen et al. 2021) is a two-stage context-aware network. BMNet (Zhu et al. 2022a) removes shadows using invertible neural networks. ST-CGAN (Wang, Li, and Yang 2018) jointly detects and removes shadows. ARGAN (Ding et al. 2019) uses LSTM attention to detect shadows. DHAN (Cun, Pun, and Shi 2020) employs the SMGAN to generate shadow matting. RIS-GAN (Zhang et al. 2020) explores the relationship of the residual images. SG-ShadowNet (Wan et al. 2022) treats shadow removal as intra-image style transfer. These methods use Conditional GAN, StyleGAN (Karras, Laine, and Aila 2019) or Uformer (Guo et al. 2023a; Wang et al. 2022) as their architecture. However, supervised methods fail to handle self-shadows, as there are no ground truths (creating ones is intractable).

The weakly-supervised shadow removal methods (Le and Samaras 2020; Liu et al. 2021a) require shadow masks to distinguish shadow and non-shadow patches or to produce pseudo shadow pairs, making them unsuitable for handling soft or self-shadows. Unsupervised shadow removal methods (Hu et al. 2019b; Liu et al. 2021a; Jin, Sharma, and Tan 2021) need non-shadow reference images, which are hard to obtain if the input has self-shadows. They rely on CycleGAN (Zhu et al. 2017; Jin, Yang, and Tan 2022) and learn from unpaired data.

Denoising Diffusion Probabilistic Models (DDPM) (Ho, Jain, and Abbeel 2020) and Denoising Diffusion Implicit Models (DDIM) (Song, Meng, and Ermon 2021) have recently demonstrated promising generative ability (Li et al. 2023). Palette (Saharia et al. 2022a), Unit-DDPM (Sasaki, Willcocks, and Breckon 2021), RDDM (Liu et al. 2023c) were proposed for image-to-image translation. WeatherDiffusion (Özdenizci and Legenstein 2023) is a diffusion method presented for weather removal. However, they are not designed to remove shadows and lack object structures (Preechakul et al. 2022). Shadowdiffusion (Guo et al. 2023b) uses shadow degradation as a prior and unrolling diffusion to remove shadows. However, the degradation model’s accuracy depends on the mask, and hence this method may not work well for soft and self-shadow removal.

Proposed Method: DeS3

Fig. 2 shows our pipeline, comprising a forward diffusion (dashed line) and reverse sampling (solid line). The main goal of DeS3 is to remove hard, soft and self shadows from a single image using an end-to-end network (without masks).

Conditional DDIM We use Denoising Diffusion Implicit Models (DDIM) (Song, Meng, and Ermon 2021) as our generative model. DDIM is trained on a clean image $\mathbf{x}_0 \sim q(\mathbf{x}_0)$ via a forward diffusion process $q(\mathbf{x}_t | \mathbf{x}_{t-1})$ that sequentially adds the Gaussian noise at every time steps t , *i.e.*, $q(\mathbf{x}_t | \mathbf{x}_{t-1}) = \mathcal{N}(\mathbf{x}_t; \sqrt{1 - \beta_t} \mathbf{x}_{t-1}, \beta_t \mathbf{I})$, where $\{\beta_t\}_{t=0}^T$ is

¹Our data and code is available at: https://github.com/jinyeying/DeS3_Deshadow.

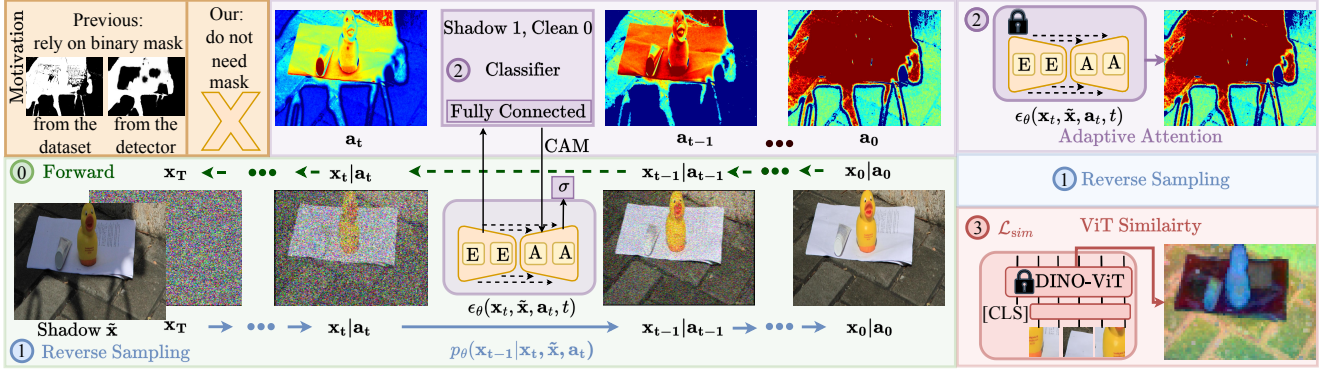


Figure 2: The architecture and the motivation of our DeS3. (1) The forward diffusion is shown in green. The reverse sampling starts from the noise map \mathbf{x}_T concatenated with the conditional shadow inputs $\tilde{\mathbf{x}}$. Our DeS3 samples image \mathbf{x}_t at each time step t . (2) We inject a classifier into the noise prediction network $\epsilon_\theta(\mathbf{x}_t, \tilde{\mathbf{x}}, \mathbf{a}_t, t)$. Adaptive attention \mathbf{a}_t is progressively improved at each time step t . (3) To guide the reverse sampling to output the object structure features, we have the ViT similarity loss \mathcal{L}_{sim} , extracted keys from the pre-trained DINO-ViT.




Figure 3: Shadows can be categorized into hard, soft, and self shadows. Left: Soft and self shadows. Right: Wrong binary masks, due to the ambiguous boundaries (the red boxes for the incorrect and blue boxes for mis-detected masks).




Figure 5: The results of adaptive attention that enables our reverse sampling to focus on hard, soft and self shadows.




Figure 4: Adaptive attention is refined during inference.

a variance schedule. The forward diffusion of length T expressed as:

$$q(\mathbf{x}_{1:T} | \mathbf{x}_0) = \prod_{t=1}^T q(\mathbf{x}_t | \mathbf{x}_{t-1}). \quad (1)$$

Diffusion learns to reverse the process in Eq.(1) with the Gaussian transitions, *i.e.*, $p_\theta(\mathbf{x}_{t-1} | \mathbf{x}_t) = \mathcal{N}(\mathbf{x}_{t-1}; \boldsymbol{\mu}_\theta(\mathbf{x}_t, t), \boldsymbol{\Sigma}_\theta(\mathbf{x}_t, t))$, at each time step.

The reverse denoising process is parameterized by a trainable network (e.g., U-Net), which estimates the mean




Figure 6: d is DC (2021)'s result, and f is ours. We integrate adaptive attention with diffusion to tackle self-shadows.

$\boldsymbol{\mu}_\theta(\mathbf{x}_t, t)$ and variance $\boldsymbol{\Sigma}_\theta(\mathbf{x}_t, t)$ with parameter θ . This reverse process starts from a standard normal distribution $p(\mathbf{x}_T) = \mathcal{N}(\mathbf{x}_T; \mathbf{0}, \mathbf{I})$ and follows: $p_\theta(\mathbf{x}_{0:T}) = p(\mathbf{x}_T) \prod_{t=1}^T p_\theta(\mathbf{x}_{t-1} | \mathbf{x}_t)$. The noise prediction network $\epsilon_\theta(\mathbf{x}_t, t)$ can be optimized by this objective function: $L_\theta(\text{DM}) = \mathbb{E}_{\mathbf{x}_0, \epsilon, t} [\|\epsilon - \epsilon_\theta(\mathbf{x}_t, t)\|^2]$. After the optimization, we can sample from the learned parameterized Gaussian transitions $p_\theta(\mathbf{x}_{t-1} | \mathbf{x}_t)$ by: $\mathbf{x}_{t-1} = \boldsymbol{\mu}_\theta(\mathbf{x}_t, t) + \boldsymbol{\Sigma}_\theta^{1/2}(\mathbf{x}_t, t)\epsilon$, $\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$. Our DeS3 use conditional DDIM to output deterministic and consistent samples.

We reverse process $p_\theta(\mathbf{x}_{0:T} | \tilde{\mathbf{x}})$, without changing the diffusion process $q(\mathbf{x}_{1:T} | \mathbf{x}_0)$ in Eq.(1). We replace the

noise prediction network $\epsilon_\theta(\mathbf{x}_t, t)$ with $\epsilon_\theta(\mathbf{x}_t, \tilde{\mathbf{x}}, t)$, where $\tilde{\mathbf{x}}$ denotes the shadow image. Here, we use the conditional DDIM (Song, Meng, and Ermon 2021) and input shadow images $\tilde{\mathbf{x}}$ and shadow-free \mathbf{x}_0 pairs ($\tilde{\mathbf{x}}, \mathbf{x}_0$) to output shadow-free image. We concatenate $\tilde{\mathbf{x}}$ and \mathbf{x}_t channel-wisely, and input them to the deterministic reverse process. Therefore, we treat shadow removal as a reverse process with the loss function $L_\theta(\text{CDM})$ and parameterize a noise prediction network $\epsilon_\theta(\mathbf{x}_t, \tilde{\mathbf{x}}, t)$:

$$L_\theta(\text{CDM}) = \mathbb{E}_{\mathbf{x}_0, \epsilon, t} \left[\|\epsilon - \epsilon_\theta(\mathbf{x}_t, \tilde{\mathbf{x}}, t)\|^2 \right], \quad (2)$$

$$\hat{\mathbf{x}}_0(\mathbf{x}_t) = \frac{1}{\sqrt{\alpha_t}} \left(\mathbf{x}_t - \sqrt{1 - \alpha_t} \cdot \epsilon_\theta(\mathbf{x}_t, \tilde{\mathbf{x}}, t) \right),$$

where $\hat{\mathbf{x}}_0(\mathbf{x}_t)$ refers to the estimated clean image from the sample \mathbf{x}_t . For brevity, we short $\hat{\mathbf{x}}_0(\mathbf{x}_t)$ as \mathbf{x} .

The shadow removal starts from the noise map, $\mathbf{x}_T \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$, with the condition shadow input $\tilde{\mathbf{x}}$, and applies the diffusion towards the target non-shadow image \mathbf{x}_0 . The sampling from the reverse process $\mathbf{x}_{t-1} \sim p_\theta(\mathbf{x}_{t-1} | \mathbf{x}_t, \tilde{\mathbf{x}})$ employs: $\mathbf{x}_{t-1} = \sqrt{\alpha_{t-1}} \left(\frac{\mathbf{x}_t - \sqrt{1 - \alpha_t} \cdot \epsilon_\theta(\mathbf{x}_t, \tilde{\mathbf{x}}, t)}{\sqrt{\alpha_t}} \right) + \sqrt{1 - \alpha_{t-1}} \cdot \epsilon_\theta(\mathbf{x}_t, \tilde{\mathbf{x}}, t)$.

Adaptive Classifier-driven Attention

Our novelty lies in adaptive attention, which can highlight the regions of all three types of shadows, particularly self-shadows. Figs. 6e to 6f show our results using adaptive attention, and additional results can be found in Fig. 5. Our adaptive attention is progressively refined throughout the reverse process, as illustrated in Fig. 4. In contrast, (Jin, Sharma, and Tan 2021) relies solely on fixed CAMs, which tend to struggle with handling self-shadows, as demonstrated in Figs. 6c to 6d.

Since a plain diffusion model is not aware of the shadow regions, we design an adaptive attention to allow the reverse sampling to focus on the shadow regions. Our soft attention map differs from existing binary masks shown in Fig. 3, which are hard to obtain for self and soft shadows. We fuse the attention into diffusion by injecting the classifier in the U-Net. Specifically, by labeling the clear image \mathbf{x}_0 as 0, and the shadow condition $\tilde{\mathbf{x}}$ as 1, and using binary classification with Class Activation Map (CAM) (Zhou et al. 2016). Given an image $x \in \{\tilde{\mathbf{x}}, \mathbf{x}_0\}$, $L_{\text{cam}} = -(\mathbb{E}_{x \sim \tilde{\mathbf{x}}}[\log(C(x))] + \mathbb{E}_{x \sim \mathbf{x}_0}[\log(1 - C(x))])$, where C is the classifier. By leveraging the information from the classifier, the CAM attention is learned from the data and can focus on hard, soft, and self-shadow regions, as shown in Fig. 5.

Unlike the fixed attention in (Jin, Sharma, and Tan 2021), our attention is adaptive and progressively refines throughout the diffusion process. Besides the CAM attention, we employ a residual map refinement. By utilizing shadow and shadow-free pairs during training, we can employ a residual map to refine classifier-driven attention progressively. That is to say, the noise estimation network $\epsilon_\theta(\mathbf{x}_t, \tilde{\mathbf{x}}, \mathbf{a}_t, t)$ has two tasks: estimating the noise and progressively refining the attention \mathbf{a}_t . We add one Sigmoid Conv layer after the last layer of the noise estimation network, and obtain the residual map \mathbf{m}_{res} by computing the difference map between shadow




Figure 7: We visualize the deep DINO-ViT features. These object structure ViT features help DeS3 preserve the object/scene structures (e.g., fire hydrant in self shadow example, duck, bag, bear in soft shadow examples).

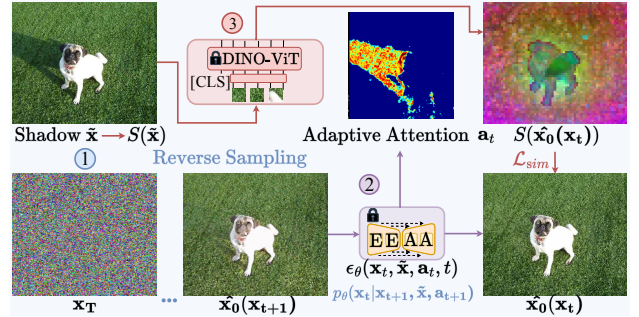


Figure 8: In our reverse sampling process, we use the loss function to guide the output to follow the ViT similarity. When \mathcal{L}_{sim} is high, continue the reverse sampling till 25 steps. When \mathcal{L}_{sim} is low, the reverse sampling can be stopped early.

and shadow-free pair. This difference map is then used to guide the refinement of \mathbf{a}_t via the following loss:

$$L_{\text{att}} = \mathbb{E}_{t \sim [1, T]} \left[\|\mathbf{a}_t - \mathbf{m}_{\text{res}}\|^2 \right], \quad (3)$$

where $\mathbf{m}_{\text{res}} = \sigma(\mathbf{x}_0 - \tilde{\mathbf{x}})$, $\sigma(\cdot)$ is the Sigmoid Conv layer.

Object Structures in Reverse Sampling

Unlike the VGG loss (Jin, Sharma, and Tan 2021) used in training, we propose to use a DINO-ViT loss as a stopping criterion in the reverse sampling (inference) stage. DINO-ViT demonstrates greater robustness in preserving object structures across all types of shadows when compared to VGG, as shown in Fig. 13. Moreover, in contrast to the losses in (Jin, Sharma, and Tan 2021), our DINO-ViT loss is exclusively employed during inference, serving as a stopping criterion, as depicted in Fig. 8. If the total loss L_t is smaller than that in the subsequent diffusion step, we halt the reverse process earlier.

In contrast to the existing shadow removal methods that use the supervised losses (requiring ground truth shadow-

Method	Train	RMSE ↓			PSNR ↑			SSIM ↑		
		S	NS	ALL	S	NS	ALL	S	NS	ALL
DHAN (2020)	P+M	8.94	4.80	5.67	33.67	34.79	30.51	0.978	0.979	0.949
Auto (2021)	P+M	8.56	5.75	6.51	32.26	31.87	28.40	0.966	0.945	0.893
EM (2022b)	P+M	10.00	6.04	7.20	29.44	26.67	24.16	0.937	0.879	0.779
BM (2022a)	P+M	6.61	3.61	4.46	35.05	36.02	31.69	0.981	0.982	0.956
SG (2022)	P+M	7.53	3.08	4.33	33.88	36.43	31.38	0.981	0.987	0.959
DC (2021)	UP	7.70	3.65	4.66	34.00	35.53	31.53	0.975	0.981	0.955
SR3 (2022b)	P	-	-	-	35.44	34.35	31.29	0.980	0.970	0.946
W.Dif. (2023)	P	-	-	-	33.38	31.15	28.45	0.981	0.972	0.951
DSC (2019a)	P	8.81	4.41	5.71	30.65	31.94	27.76	0.960	0.965	0.903
Ours	P	5.88	2.83	3.72	37.45	38.12	34.11	0.984	0.988	0.968

Table 1: Quantitative results on the SRD dataset. S, NS and ALL represent shadow, non-shadow and entire regions. M shows ground truth shadow masks are also used in training. “P” and “UP” stand for “Paired” and “Unpaired”.




Figure 9: Self shadow removal results on UIUC dataset. (a) Input images. (b) Our results. (c)~(g) Results of the state-of-the-art methods. Note that our method does not need additional masks in training and testing.

free) in training, we use the self-tuned losses in the reverse sampling (inference). Preserving object/scene structures is critical for shadow removal (Qu et al. 2017). Therefore, we introduce the ViT similarity loss to guide the reverse sampling. To preserve the object structures, our key idea is to exploit a self-similarity prior (Shechtman and Irani 2007). For the soft shadow example in Fig. 7, the left and right parts of the bag contain different lighting conditions and appearances, but they do share similar structures

(i.e., self-similarity). A pre-trained DINO-ViT (Caron et al. 2021) can provide better object structure features, compared to the CNN backbone, transformer-based backbones shown in (Karmali et al. 2022). Motivated by this, we use a pre-trained DINO-ViT as our feature extractor, enabling us to capture deep features (Amir et al. 2021), and providing useful feature supervision to our DeS3.

DINO-ViT Keys We extract keys (deep features) from the pre-trained DINO-ViT. In the multi-head self-attention




Figure 10: Hard shadow removal results on the USR (top), AISTD (second) and SRD (bottom four rows) datasets.

layer, the keys contain the object structure information (Tumanyan et al. 2022). In Fig. 7, we show the Principal Component Analysis (PCA) visualization of the keys’ self-similarity and demonstrate the top three components as RGB at the deepest layer of DINO-ViT. As one can observe, the feature representations capture the object structures, helping our DeS3 preserve and recover meaningful features.

We propose the ViT similarity loss between the intermediate keys of the input image and the estimated denoised output image during the reverse sampling given a noise estimation network $\epsilon_\theta(\mathbf{x}_t, \tilde{\mathbf{x}}, \mathbf{a}_t, t)$ and ViT model:

$$\mathcal{L}_{sim}(\tilde{\mathbf{x}}, \mathbf{x}) = \|S^l(\tilde{\mathbf{x}}) - S^l(\mathbf{x})\|_F, \quad (4)$$

where S is the self-similarity descriptor, $S^l(x) \in \mathbb{R}^{(n) \times (n)}$, with $n \times n$ dimension, where n is the number of patches. In DINO-ViT, the input image is split into n non-overlapping patches. $l = 11$ is the deepest ViT layer. $\|\cdot\|_F$ is the Frobenius norm. The self-similarity descriptor is defined as:

$$S^l(\mathbf{x})_{ij} = \cos(k_i^l(\mathbf{x}), k_j^l(\mathbf{x})) = 1 - \frac{k_i^l(\mathbf{x}) \cdot k_j^l(\mathbf{x})}{\|k_i^l(\mathbf{x})\| \cdot \|k_j^l(\mathbf{x})\|},$$

where $\cos(\cdot)$ is the cosine similarity between spatial keys k_i and k_j .

Experiments

Implementation To ensure fair comparisons, all the baselines, including ours are trained and tested on the same datasets. We trained our DeS3 on each dataset and tested on the corresponding dataset, e.g., for SRD (Qu et al. 2017), we used 2680 SRD and 408 SRD images for training and testing, respectively. We use 1000 steps for training and 25 steps for inference, noise schedule β_t linear from 0.0001 to 0.02. α and β are empirically set to 0.5 in training.




Figure 11: Shadow removal results on the soft shadow dataset LRSS (top three) and SRD (bottom three rows) datasets. Note that our method does not need additional masks in training and testing.

$\mathcal{L}_t(\tilde{\mathbf{x}}, \mathbf{x}) = \lambda_{sim} \mathcal{L}_{sim}$, which is the total reverse sampling loss, where $\lambda_{sim} = 2$, we balance the self-tuned ViT loss to select the good sampling output.

Shadow Removal on Hard Shadows The SRD dataset consists of 2680 training and 408 testing pairs of shadow and shadow-free images without shadow masks. It contains hard shadows and soft shadows. For the baseline methods that require shadow masks, we additionally use the processing results of DHAN (Cun, Pun, and Shi 2020) for SRD shadow masks.

Fig. 10 (bottom four rows) and Table 1 show the hard shadow removal results on the SRD dataset. Fig. 11 (bottom three rows) show the soft shadow removal results on the SRD dataset, which show our DeS3 outperforms the SOTA methods. The best values for each metric are highlighted in **bold**. Moreover, DeS3 achieves the highest PSNR and SSIM values in shadow, non-shadow and overall areas. PSNR and SSIM are widely used in image generation tasks (2023c; 2023b; 2023a). Fig. 10 (second row) show results on the AISTD dataset, which demonstrate the superiority of our DeS3 compared to baseline results. The USR dataset (Hu et al. 2019b) is an unpaired dataset, including hard, soft, and self shadows. We evaluate on USR test shadow images and show results in Fig. 10 (first row).

Shadow Removal on Soft and Self Shadows The LRSS (Gryka, Terry, and Brostow 2015) is a soft shadow dataset with 134 shadow images²; we followed (Jin, Sharma, and Tan 2021; Gryka, Terry, and Brostow 2015), using the

²LRSS Dataset are obtained from their project website: <http://visual.cs.ucl.ac.uk/pubs/softshadows/>

Aspects	Ours	S.Diff.	SG	G2R	SP+M+I	DC	SP+M
1.DeS \uparrow	8.6\pm1.3	5.8 \pm 3.9	<u>6.2\pm2.7</u>	5.6 \pm 3.8	6.0 \pm 4.0	5.4 \pm 1.8	5.6 \pm 2.3
2.Real \uparrow	9.0\pm0.9	6.9 \pm 2.8	6.9 \pm 2.0	<u>7.2\pm3.2</u>	7.0 \pm 1.5	6.0 \pm 0.5	6.6 \pm 1.8

Table 2: User study on the self shadow removal of the UIUC dataset (Guo, Dai, and Hoiem 2011), our method obtained the highest mean (the max score is 10), showing our method is effective in shadow removal (DeS) and visually realistic (Real). The best results are in bold, the second best are in underline.

Method	Guo	Gryka	DHAN	SP+M	MS.GAN	DC	Guo (at)	S.Diff.	Ours
RMSE \downarrow	6.02	4.38	7.92	7.48	7.13	<u>3.48</u>	5.87	3.49	3.01
PSNR \uparrow	27.88	29.25	25.57	23.93	25.12	<u>31.01</u>	28.02	28.86	33.95
Train	P+M	P+M+S	P+M+S	P+M	UP	UP	P	P+M	P

Table 3: Results on the LRSS (soft shadow) dataset. M and S represent ground truth shadow masks and synthetic data are used in training. ‘‘P’’ and ‘‘UP’’ stand for ‘‘Paired’’ and ‘‘Unpaired’’. Our DeS3 does not need shadow masks.



Figure 12: Ablation studies on the ViT similarity loss.




Figure 13: Ablation study on ViT vs VGG, the output of ViT is better than VGG.

same 34 LRSS images with their corresponding shadow-free images for evaluation shown in Table 3. Our DeS3 achieves the highest PSNR and lowest RMSE.

Fig. 11 (top three rows) show results on the LRSS dataset. Besides the SOTA baselines, we compared our method with the traditional hard and soft shadow removal method (Guo, Dai, and Hoiem 2012). Since the majority of the SOTA baselines need the shadow mask for evaluation, we use the shadow mask provided by LRSS dataset (Gryka, Terry, and Brostow 2015). However, the provided binary shadow mask does not consider the ambiguous boundaries of soft and self shadows, thus leading to remaining shadows in the ambiguous boundaries.

UCF (Zhu et al. 2010) and UIUC dataset (Guo, Dai, and Hoiem 2011), contain 245 and 108 images. UIUC provides 30 images where shadows are caused by objects in the scene, which are used for self shadow evaluation, the results are shown in Fig. 9. For the baseline methods that require shadow masks, we additionally use the detection results of BDRAR (Zhu et al. 2018).

Method	S \downarrow	NS \downarrow	Overall \downarrow
w/o Adaptive Classifier Atten	6.35	3.02	4.46
w/o \mathcal{L}_{sim}	6.21	2.99	3.94
Ours (complete model)	5.88	2.83	3.72

Table 4: Ablation experiments of our method using the SRD dataset. The numbers represent RMSE, lower is better.

Ablation Studies Fig. 12 and Table. 4 show the effectiveness of the ViT similarity loss used in DeS3. From Fig. 12 (second row), the ViT similarity loss help in maintaining the structures of the mountain, owing to the features (the self-similarity keys) of the pre-trained DINO-ViT. We compare our method with two diffusion baselines: WeatherDiffusion (Özdenizci and Legenstein 2023) and SR3 (Saharia et al. 2022b) shown in Tab.1. More results are in the supplementary material.

Conclusion

We have proposed DeS3, the first shadow removal network that performs shadow removal robustly on hard, soft and self shadows. Unlike existing methods, our DeS3 does not rely on masks during training and testing. To guide our reverse sampling process to preserve object/scene structure information, we propose the self-tuned ViT similarity loss. Unlike existing methods, our DeS3 performs shadow removal robustly on outdoor self shadows, which is intractable to have ground truths.

Acknowledgments

This research is supported by the National Research Foundation, Singapore under its AI Singapore Programme (AISG Award No: AISG2-PhD/2022-01-037(T)). Wenhan Yang’s research is supported in part by the Basic and Frontier Research Project of PCL, and the Major Key Project of PCL.

References

- Amir, S.; Gandelsman, Y.; Bagon, S.; and Dekel, T. 2021. Deep vit features as dense visual descriptors. *arXiv preprint arXiv:2112.05814*.
- Caron, M.; Touvron, H.; Misra, I.; Jégou, H.; Mairal, J.; Bojanowski, P.; and Joulin, A. 2021. Emerging properties in self-supervised vision transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 9650–9660.
- Chen, Z.; Long, C.; Zhang, L.; and Xiao, C. 2021. CANet: A Context-Aware Network for Shadow Removal. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 4743–4752.
- Cun, X.; Pun, C.-M.; and Shi, C. 2020. Towards ghost-free shadow removal via dual hierarchical aggregation network and shadow matting gan. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, 10680–10687.
- Ding, B.; Long, C.; Zhang, L.; and Xiao, C. 2019. Argan: Attentive recurrent generative adversarial network for shadow detection and removal. In *Proceedings of the IEEE/CVF international conference on computer vision*, 10213–10222.
- Finlayson, G. D.; Hordley, S. D.; Lu, C.; and Drew, M. S. 2005. On the removal of shadows from images. *IEEE transactions on pattern analysis and machine intelligence*, 28(1): 59–68.
- Fu, L.; Zhou, C.; Guo, Q.; Juefei-Xu, F.; Yu, H.; Feng, W.; Liu, Y.; and Wang, S. 2021. Auto-exposure fusion for single-image shadow removal. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 10571–10580.
- Gryka, M.; Terry, M.; and Brostow, G. J. 2015. Learning to remove soft shadows. *ACM Transactions on Graphics (TOG)*, 34(5): 1–15.
- Guo, L.; Huang, S.; Liu, D.; Cheng, H.; and Wen, B. 2023a. ShadowFormer: Global Context Helps Image Shadow Removal. *arXiv preprint arXiv:2302.01650*.
- Guo, L.; Wang, C.; Yang, W.; Huang, S.; Wang, Y.; Pfister, H.; and Wen, B. 2023b. Shadowdiffusion: When degradation prior meets diffusion model for shadow removal. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 14049–14058.
- Guo, R.; Dai, Q.; and Hoiem, D. 2011. Single-image shadow detection and removal using paired regions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2033–2040.
- Guo, R.; Dai, Q.; and Hoiem, D. 2012. Paired regions for shadow detection and removal. *IEEE transactions on pattern analysis and machine intelligence*, 35(12): 2956–2967.
- Ho, J.; Jain, A.; and Abbeel, P. 2020. Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems*, 33: 6840–6851.
- Hu, X.; Fu, C.-W.; Zhu, L.; Qin, J.; and Heng, P.-A. 2019a. Direction-Aware Spatial Context Features for Shadow Detection and Removal. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 42(11): 2795–2808.
- Hu, X.; Jiang, Y.; Fu, C.-W.; and Heng, P.-A. 2019b. Mask-ShadowGAN: Learning to remove shadows from unpaired data. In *Proceedings of the IEEE International Conference on Computer Vision*, 2472–2481.
- Huang, J.-B.; and Chen, C.-S. 2009. Moving cast shadow detection using physics-based features. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, 2310–2317. IEEE.
- Jin, Y.; Li, R.; Yang, W.; and Tan, R. T. 2023a. Estimating reflectance layer from a single image: Integrating reflectance guidance and shadow/specular aware learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, 1069–1077.
- Jin, Y.; Lin, B.; Yan, W.; Yuan, Y.; Ye, W.; and Tan, R. T. 2023b. Enhancing visibility in nighttime haze images using guided apsf and gradient adaptive convolution. In *Proceedings of the 31st ACM International Conference on Multimedia*, 2446–2457.
- Jin, Y.; Sharma, A.; and Tan, R. T. 2021. DC-ShadowNet: Single-Image Hard and Soft Shadow Removal Using Unsupervised Domain-Classifer Guided Network. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 5027–5036.
- Jin, Y.; Yan, W.; Yang, W.; and Tan, R. T. 2022. Structure Representation Network and Uncertainty Feedback Learning for Dense Non-Uniform Fog Removal. In *Proceedings of the Asian Conference on Computer Vision*, 2041–2058.
- Jin, Y.; Yang, W.; and Tan, R. T. 2022. Unsupervised night image enhancement: When layer decomposition meets light-effects suppression. In *European Conference on Computer Vision*, 404–421. Springer.
- Karmali, T.; Parihar, R.; Agrawal, S.; Rangwani, H.; Jampani, V.; Singh, M.; and Babu, R. V. 2022. Hierarchical Semantic Regularization of Latent Spaces in StyleGANs. In *European Conference on Computer Vision*, 443–459. Springer.
- Karras, T.; Laine, S.; and Aila, T. 2019. A style-based generator architecture for generative adversarial networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 4401–4410.
- Le, H.; and Samaras, D. 2019. Shadow removal via shadow image decomposition. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 8578–8587.
- Le, H.; and Samaras, D. 2020. From shadow segmentation to shadow removal. In *European Conference on Computer Vision*, 264–281. Springer.
- Le, H.; and Samaras, D. 2021. Physics-based shadow image decomposition for shadow removal. *IEEE Transactions on Pattern Analysis & Machine Intelligence*, (01): 1–1.
- Li, X.; Ren, Y.; Jin, X.; Lan, C.; Wang, X.; Zeng, W.; Wang, X.; and Chen, Z. 2023. Diffusion Models for Image Restoration and Enhancement—A Comprehensive Survey. *arXiv preprint arXiv:2308.09388*.
- Liu, J.; Wang, Q.; Fan, H.; Li, W.; Qu, L.; and Tang, Y. 2023a. A Decoupled Multi-Task Network for Shadow Removal. *IEEE Transactions on Multimedia*.

- Liu, J.; Wang, Q.; Fan, H.; Tian, J.; and Tang, Y. 2023b. A Shadow Imaging Bilinear Model and Three-Branch Residual Network for Shadow Removal. *IEEE Transactions on Neural Networks and Learning Systems*.
- Liu, J.; Wang, Q.; Fan, H.; Wang, Y.; Tang, Y.; and Qu, L. 2023c. Residual denoising diffusion models. *arXiv preprint arXiv:2308.13712*.
- Liu, Z.; Yin, H.; Mi, Y.; Pu, M.; and Wang, S. 2021a. Shadow removal by a lightness-guided network with training on unpaired data. *IEEE Transactions on Image Processing*, 30: 1853–1865.
- Liu, Z.; Yin, H.; Wu, X.; Wu, Z.; Mi, Y.; and Wang, S. 2021b. From Shadow Generation to Shadow Removal. In *CVPR*.
- Özdenizci, O.; and Legenstein, R. 2023. Restoring vision in adverse weather conditions with patch-based denoising diffusion models. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- Preechakul, K.; Chatthee, N.; Wizadwongsa, S.; and Suwajanakorn, S. 2022. Diffusion autoencoders: Toward a meaningful and decodable representation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 10619–10629.
- Qu, L.; Tian, J.; He, S.; Tang, Y.; and Lau, R. W. 2017. Dshadownet: A multi-context embedding deep network for shadow removal. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 4067–4075.
- Saharia, C.; Chan, W.; Chang, H.; Lee, C.; Ho, J.; Salimans, T.; Fleet, D.; and Norouzi, M. 2022a. Palette: Image-to-image diffusion models. In *ACM SIGGRAPH 2022 Conference Proceedings*, 1–10.
- Saharia, C.; Ho, J.; Chan, W.; Salimans, T.; Fleet, D. J.; and Norouzi, M. 2022b. Image super-resolution via iterative refinement. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- Salvador, E.; Cavallaro, A.; and Ebrahimi, T. 2004. Cast shadow segmentation using invariant color features. *Computer vision and image understanding*, 95(2): 238–259.
- Sasaki, H.; Willcocks, C. G.; and Breckon, T. P. 2021. Unit-ddpm: Unpaired image translation with denoising diffusion probabilistic models. *arXiv preprint arXiv:2104.05358*.
- Shechtman, E.; and Irani, M. 2007. Matching local self-similarities across images and videos. In *2007 IEEE Conference on Computer Vision and Pattern Recognition*, 1–8. IEEE.
- Song, J.; Meng, C.; and Ermon, S. 2021. Denoising diffusion implicit models. In *International Conference on Learning Representations*.
- Tumanyan, N.; Bar-Tal, O.; Bagon, S.; and Dekel, T. 2022. Splicing ViT Features for Semantic Appearance Transfer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 10748–10757.
- Vicente, T. F. Y.; Hoai, M.; and Samarasinghe, D. 2017. Leave-one-out kernel optimization for shadow detection and removal. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 40(3): 682–695.
- Wan, J.; Yin, H.; Wu, Z.; Wu, X.; Liu, Y.; and Wang, S. 2022. Style-Guided Shadow Removal. In *European Conference on Computer Vision*, 361–378. Springer.
- Wang, C.; Pan, J.; Lin, W.; Dong, J.; and Wu, X.-M. 2023a. SelfPromer: Self-Prompt Dehazing Transformers with Depth-Consistency. *arXiv preprint arXiv:2303.07033*.
- Wang, C.; Pan, J.; Wang, W.; Dong, J.; Wang, M.; Ju, Y.; and Chen, J. 2023b. PromptRestorer: A Prompting Image Restoration Method with Degradation Perception. In *Thirty-seventh Conference on Neural Information Processing Systems*.
- Wang, J.; Li, X.; and Yang, J. 2018. 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*, 1788–1797.
- Wang, J.; Qian, X.; Zhang, M.; Tan, R. T.; and Li, H. 2023c. Seeing What You Said: Talking Face Generation Guided by a Lip Reading Expert. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 14653–14662.
- Wang, Z.; Cun, X.; Bao, J.; Zhou, W.; Liu, J.; and Li, H. 2022. Uformer: A general u-shaped transformer for image restoration. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 17683–17693.
- Zhang, L.; Long, C.; Zhang, X.; and Xiao, C. 2020. Risgan: Explore residual and illumination with generative adversarial networks for shadow removal. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, 12829–12836.
- Zhou, B.; Khosla, A.; Lapedriza, A.; Oliva, A.; and Torralba, A. 2016. Learning Deep Features for Discriminative Localization. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Zhu, J.; Samuel, K. G.; Masood, S. Z.; and Tappen, M. F. 2010. Learning to recognize shadows in monochromatic natural images. In *2010 IEEE Computer Society conference on computer vision and pattern recognition*, 223–230. IEEE.
- Zhu, J.-Y.; Park, T.; Isola, P.; and Efros, A. A. 2017. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Proceedings of the IEEE International Conference on Computer Vision*, 2223–2232.
- Zhu, L.; Deng, Z.; Hu, X.; Fu, C.-W.; Xu, X.; Qin, J.; and Heng, P.-A. 2018. Bidirectional feature pyramid network with recurrent attention residual modules for shadow detection. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 121–136.
- Zhu, Y.; Huang, J.; Fu, X.; Zhao, F.; Sun, Q.; and Zha, Z.-J. 2022a. Bijective Mapping Network for Shadow Removal. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 5627–5636.
- Zhu, Y.; Xiao, Z.; Fang, Y.; Fu, X.; Xiong, Z.; and Zha, Z.-J. 2022b. Efficient Model-Driven Network for Shadow Removal. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, 3635–3643.