

Latent Feature-Guided Diffusion Models for Shadow Removal

Kangfu Mei^{*1,2}, Luis Figueroa², Zhe Lin², Zhihong Ding²,
Scott Cohen², Vishal M. Patel^{†1}

¹ Johns Hopkins University, ² Adobe Research

<https://kfmei.page/shadow-diffusion/>

Abstract

Recovering textures under shadows has remained a challenging problem due to the difficulty of inferring shadow-free scenes from shadow images. In this paper, we propose the use of diffusion models as they offer a promising approach to gradually refine the details of shadow regions during the diffusion process. Our method improves this process by conditioning on a learned latent feature space that inherits the characteristics of shadow-free images, thus avoiding the limitation of conventional methods that condition on degraded images only. Additionally, we propose to alleviate potential local optima during training by fusing noise features with the diffusion network. We demonstrate the effectiveness of our approach which outperforms the previous best method by 13% in terms of RMSE on the AISTD dataset. Further, we explore instance-level shadow removal, where our model outperforms the previous best method by 82% in terms of RMSE on the DESOBA dataset.

1. Introduction

Images captured in natural illumination often contain shadows caused by objects blocking the light from the illumination source. Shadows can degrade the performance of many computer vision algorithms, such as detection, segmentation, and recognition [31, 51]. Furthermore, removing shadows is essential for photo-editing applications such as distractor removal [49] and relighting [20]; which may rely on instance-level shadow removal. Therefore, it is critical to develop methods that can automatically remove shadows from captured images as works explored in literature.

Recently, diffusion models [40] with hierarchical denoising autoencoders [17] have shown to achieve impressive synthesis performance in terms of sample quality and diversity. The conditional generation ability further allows for

iterative refinement and fine-grained control according to certain conditions. Motivated by the success of diffusion-based image restoration models [36, 39], we adapt diffusion models for the task of shadow removal by conditioning on the input shadow image and corresponding shadow mask as a baseline approach to generate shadow-free images. However, preserving and generating high-fidelity textures and colors in the shadow region after removal is non-trivial. The baseline model appears to favor borrowing textures from the surrounding non-shadow areas rather than focusing on restoring the original details underneath the shadow, which results in incorrect color mixtures and loss of detail in the shadow region. In Fig. 2, we show one of the representative issues of image-mask conditioning, *i.e.*, the model synthesizes results containing an incorrect color mixture.

Intuitively, the intensity drop in shadow regions means that diffusion models are typically guided more strongly by the surrounding non-shadow areas. However, this guidance can harm the fidelity of the result if the texture and color under the shadow differs significantly from the surrounding areas. In addition, the multi-head attention module [43] used in diffusion models can exacerbate this issue by extracting global information. This motivates us to consider guiding the conditioned diffusion models with an additional latent feature space that captures external perceptual shadow-free information as the shadow removal priors.

Our proposed method differs from latent diffusion models (LDMs) [36] in that we incorporate a learnable feature encoder to discover a latent feature space. To optimize the latent feature encoder, we minimize the difference between the feature space of shadow images and that of shadow-free images, using it as the loss function. Through experimentation, we have found that optimizing the encoder together with the diffusion models leads to a compact and perceptual latent feature space. Additionally, we demonstrate that pretraining the diffusion model on shadow-free images simplifies the optimization process and is crucial for achieving high-fidelity synthesis. By guiding the diffusion models on the latent feature space, instead of just conditioning on the shadow image and mask, we observe significant improve-

^{*}This work was done during an internship at Adobe Research.

[†]Vishal M. Patel was supported by NSF CAREER award 2045489.




Figure 1. Given a shadow mask, our method effectively removes shadows and recovers the underlying details for shadows at the general level (top two rows) or instance level (bottom two rows). From left to right, we show the input image, shadow mask, SG-ShadowNet [44] result, our method result, and shadow-free images for comparisons.

ments in shadow removal capability.

In addition to the proposed latent feature space guidance, we propose an improved diffusion network that addresses the issue of *posterior collapse* [5, 16, 52], which refers to the local optima of diffusion models. We identify the local optimum as the degrading effect of the noise variable and introduce a Dense Latent Variable Fusion (DLVF) module that includes dense skip connections between the embedding of the noise and diffusion network. DLVF significantly improves shadow removal results without introducing additional parameters or running complexity. In summary, this paper makes the following contributions:

- A new shadow removal model that addresses the challenging task of general and instance-level shadow removal. This is the first work, to the best of our knowledge, to demonstrate the applicability of diffusion models for instance shadow removal.
- We show that it is possible to acquire compact and perceptual guidance in a learned feature space that is optimized together with the diffusion models, without relying on handcrafted features or physical quantities.
- We identify the local optimum of diffusion models that degrades the model results and introduce a dense latent variable fusion module to alleviate it, leading to significant performance improvement.

2. Related Work

Shadow Removal. The major challenge of modern learning-based shadow removal approaches comes from the large diversity of real-world shadow scenes. The performance of recent shadow removal methods degrades significantly on out-of-distribution scenes [1]. Various approaches




Figure 2. Our baseline method, which conditions diffusion models solely on shadow and mask images, produces incorrect results such as color mixing in highlight areas. In contrast, our proposed method generates results with consistent and reasonable colors that match the surrounding area.

have been explored for addressing this issue, such as using physical illumination models, handcrafted priors, and image gradients [8, 9, 15]. The recent trend has been in developing learning-based methods that can predict shadow-free scenes [2, 11, 21, 23, 27] or intermediate factors [25, 26] for restoration. These methods have improved from previous methods in learning data [27], shadow effects [23], network architecture [2, 13, 45], and learning target decomposition [25, 26, 50]. In particular, generative models have gained some traction for shadow removal. ARGAN [6] removes the effect of shadow in a progressive manner determined by a discriminator. Nevertheless, these end-to-end GAN-based methods lack generalizability on the out-of-distribution shadow images without significant modifications. Recent diffusion models have shown promising performance in general image restoration tasks but are rarely explored in shadow removal [30, 32, 37]. In this work, we first propose to apply diffusion models for removing shadows, to leverage their impressive capacity of perceptual synthesis, which is shown to be capable of gradually preserving details in denoising sequences.

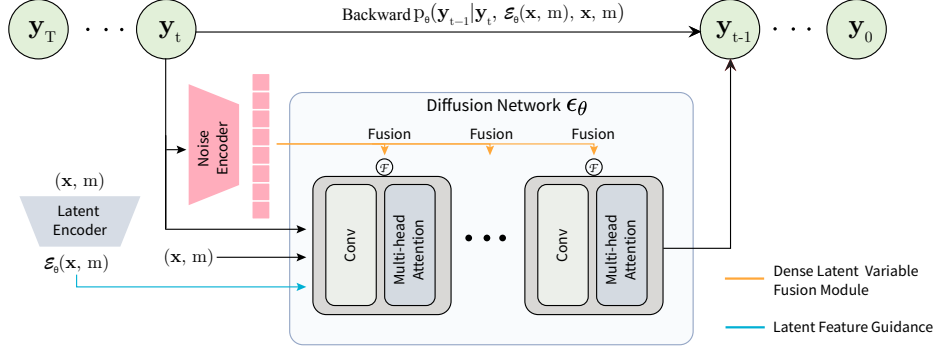


Figure 3. Our diffusion model architecture is illustrated in this backward diffusion diagram. The latent feature encoder $\mathcal{E}_\theta(\cdot)$ takes the shadow image $\mathbf{x} \in \mathbb{R}^{3 \times H \times W}$ and shadow mask $m \in \mathbb{R}^{1 \times H \times W}$ as input, with a resolution of $H \times W$, and acquires the latent feature in a compressed dimension of $1 \times H \times W$. The diffusion network $\epsilon_\theta(\cdot)$ conditioned on (\mathbf{x}, m) takes the latent feature concatenated with the noisy image $\mathbf{y}_t \in \mathbb{R}^{3 \times H \times W}$ as input, and estimates the noiseless image $\mathbf{y}_{t-1} \in \mathbb{R}^{3 \times H \times W}$ at each diffusion process $p_\theta(\cdot)$. In this process, the noise encoder takes the noise image \mathbf{y}_t as input and acquires a 1-D vector as the noise embedding, which is fused with the diffusion network features by modulation for escaping the local optima.

Latent Feature Space Guidance. Guidance has become an essential component of diffusion models and powers spectacular image generation results in recent works. Typical guidance for diffusion models includes class information [4], text description [35, 38], and even gradients [18]. Nevertheless, these features cannot be easily adopted in shadow removal to provide more guidance than images. In literature, physical quantities and handcrafted features have been heavily explored for guiding the restoration network. Zhu et al. [53] propose to guide the network with an estimated shadow-invariant color map, and Wan et al. [44] propose to guide the network with coarse de-shadowed images. Illumination invariant representations [7, 12, 41] are another related approach that aims to decompose intrinsic images by finding quantities invariant to color, density, or shading. In our approach, we define a new latent feature space for guiding diffusion models. By maximizing the similarity between the shadow and shadow-free latents, we empirically demonstrate that it better guides the diffusion model to remove shadows by encapsulating essential perceptual information as a shadow-free prior.

Posterior Collapse. The problem of posterior collapse refers to undesirable local optima first observed in the training of VAE models [24]. Efforts to address it have included aggressive optimization of the inference network proposed by He et al. [16], weakening the generator by Fu et al. [10], and changing the objective function by Tolstikhin et al. [42]. In this work, we show that although this issue has primarily been investigated in VAE models, conditional diffusion models can also suffer from similar issues. Specifically, the conditions used in diffusion models usually provide stronger guidance compared to the latent noise variable. Inspired by previous efforts to address the issue, we propose a new Dense Latent Variable Fusion (DLVF) module for dif-

fusion models and experimentally demonstrate that this design improvement improves shadow removal results without introducing additional costs or modifications. Different from the other latent-based diffusion methods [33, 36], ours uses simpler pixel space and models shadow-free image distribution.

3. Proposed Method

3.1. Conditional Diffusion Models

Diffusion Forward Process. The denoising diffusion models have been shown to be effective for modeling complex data distributions by reversing a gradual noising process. For the shadow-free image distribution, we define the forward diffusion process that destroys a shadow-free image $\mathbf{y} \sim q(\mathbf{y})$ with T successive standard noises:

$$q(\mathbf{y}_t | \mathbf{y}_{t-1}) = \mathcal{N}(\mathbf{y}_t; \sqrt{\beta_t} \mathbf{y}_0, (1 - \beta_t) \mathbf{I}). \quad (1)$$

Alternatively, we can use the reparameterization trick [24] to express this as:

$$\begin{aligned} q(\mathbf{y}_t | \mathbf{y}_0) &= \mathcal{N}(\mathbf{y}_t; \sqrt{\bar{\alpha}_t} \mathbf{y}_0, (1 - \bar{\alpha}_t) \mathbf{I}) \\ &= \sqrt{\bar{\alpha}_t} \mathbf{y}_0 + \epsilon \sqrt{1 - \bar{\alpha}_t}, \epsilon \sim \mathcal{N}(0, \mathbf{I}), \end{aligned} \quad (2)$$

where the variance schedule $\{\beta_1, \dots, \beta_T\}$ linearly increases and has a closed form $\alpha_t = 1 - \beta_t$ and $\bar{\alpha}_t := \prod_{s=1}^t \alpha_s$.

Diffusion Backward Process. The reversion of $q(\mathbf{y}_t | \mathbf{y}_{t-1})$ is tractable by conditioning on image \mathbf{y}_0 , and it results in sampling arbitrary shadow-free images from noise $\mathbf{y}_T \sim \mathcal{N}(0, \mathbf{I})$ for removal as:

$$q(\mathbf{y}_{t-1} | \mathbf{y}_t, \mathbf{y}_0) = \mathcal{N}(\mathbf{y}_{t-1}; \tilde{\mu}(\mathbf{y}_t, \mathbf{y}_0), \tilde{\beta}_t \mathbf{I}). \quad (3)$$

According to Bayes' rule and Eq. (2), we represent $\tilde{\mu}_t$ as:

$$\tilde{\mu}_t(\mathbf{y}_t, \mathbf{y}_0) := \frac{1}{\sqrt{\alpha_t}} (\mathbf{y}_t - \frac{1 - \alpha_t}{\sqrt{1 - \bar{\alpha}_t}} \epsilon_t). \quad (4)$$




Figure 4. The diagram illustrates the two-stage learning approach used in our proposed method. In the pretraining stage (top row), the diffusion network is trained on shadow-free images to learn a latent feature space that captures informative shadow-free priors as guidance. In the finetuning stage (bottom row), we initialize the diffusion network with the pretraining weights from (a) for shadow removal under the latent feature guidance.

Ho et al. [17] suggests modeling the process with p_θ by optimizing the *variational lower bound* (L_{VLB}) as:

$$L_{VLB} = L_T + L_{T-1} + \dots + L_0, \quad (5)$$

which is defined with Kullback–Leibler (KL) divergence as:

$$\begin{aligned} L_T &= D_{\text{KL}}(q(\mathbf{y}_T | \mathbf{y}_0) \parallel p_\theta(\mathbf{y}_T)), \\ L_t &= D_{\text{KL}}(q(\mathbf{y}_t | \mathbf{y}_{t+1}, \mathbf{y}_0) \parallel p_\theta(\mathbf{y}_t | \mathbf{y}_{t+1})), \\ L_0 &= -\log p_\theta(\mathbf{y}_0 | \mathbf{y}_1), \end{aligned} \quad (6)$$

and the effective simplification of L_t is

$$L_{\text{simple}} := E \left[\|\epsilon - \epsilon_\theta(\mathbf{y}_t, t)\|^2 \right]. \quad (7)$$

Diffusion Conditioning. A straightforward approach to producing shadow-free results is to condition diffusion models on the shadow image \mathbf{x} and shadow mask m by concatenating them with noise \mathbf{y}_t along the channel dimension:

$$p_\theta(\mathbf{y}_{t-1} | \mathbf{y}_t) := p_\theta(\mathbf{y}_{t-1} | \mathbf{y}_t, \mathbf{x}, m, t). \quad (8)$$

In the following sections, we will discuss our improvement based on the baseline following Eq. (8) that takes image \mathbf{x} , \mathbf{y}_t , and mask m as input and predicts the shadow-free noise \mathbf{y}_{t-1} for effective shadow removal as $p_\theta(\mathbf{y}_{t-1} | \mathbf{y}_t, t)$.

3.2. Latent Feature Guidance

The proposed latent feature encoder $\mathcal{E}_\theta(\cdot)$ uses the same network architecture as the diffusion network $\epsilon_\theta(\cdot)$ with the exception of a timestep embedding and predicts a single-channel feature map that has the same spatial dimension as the shadow image x . It guides the diffusion process Eq. (8) by concatenating the guidance with conditions:

$$p_\theta(\mathbf{y}_{t-1} | \mathbf{y}_t) := p_\theta(\mathbf{y}_{t-1} | \mathbf{y}_t, \mathcal{E}_\theta(\mathbf{x}, m), t). \quad (9)$$




Figure 5. Visual comparisons of different guidance strategies in shadow removal literature. (a) to (d): shadow image, invariant color map [53], coarse deshadowed image [44], and our learned latent feature. Our approach provides more perceptual information than (b) and contains fewer shadow features than (c), which still retains a shadow boundary.

We propose to learn to extract shadow-free priors using the latent feature space by minimizing the invariant loss between the encoded shadow-free images and shadow images with shadow masks as:

$$\arg \min_{\theta} \|\mathcal{E}_\theta(\mathbf{y}_0, \mathbf{0}) - \mathcal{E}_\theta(\mathbf{x}, m)\|^2. \quad (10)$$

In order to extract a compact and perceptual feature space to guide the diffusion model, we optimize the encoder together with the whole network during training based on Eq. (7):

$$\begin{aligned} L_{\text{simple}} := E \left[\|\epsilon - \epsilon_\theta(\mathbf{y}_t, \mathcal{E}_\theta(\mathbf{x}, m), \mathbf{x}, m, t)\|^2 \right] \\ + \|\mathcal{E}_\theta(\mathbf{y}_0, \mathbf{0}) - \mathcal{E}_\theta(\mathbf{x}, m)\|^2. \end{aligned} \quad (11)$$

Moreover, we empirically find that pretraining the diffusion model $\epsilon_\theta(\cdot)$ and then finetuning it accelerates the optimization of Eq. (11). Intuitively, the pretraining strategy provides a good starting point to finetune the diffusion model, such that the encoder has already learned to model the important characteristics of shadow-free images such as shadow-free textures and colors. This feature space provides strong guidance during finetuning for minimizing shadow features with the invariant loss, allowing the model to achieve higher-quality results.




Figure 6. We visualize the mean space of variables to show the collapse and our effects. The horizontal and vertical axis represent the mean of predicted \mathbf{y}_t and \mathbf{y}_{t-1} , respectively. The dashed diagonal line represents when the approximate noise is relevant. By projecting denoised samples, the results show the network with our DLVF (third) successfully moves points onto the diagonal line and away from collapses compared to without it (second).

Subsequently, we propose a two-stage learning approach

for guiding the diffusion models including pretraining and finetuning as shown in Fig. 4 as:

- Optimize the diffusion network ϵ_θ together with the latent encoder \mathcal{E}_θ for modeling the characteristics of shadow-free images by minimizing the loss:

$$E \left[\|\epsilon - \epsilon_\theta(\mathbf{y}_t, \mathcal{E}_\theta(\mathbf{y}_0, 0), \mathbf{y}_0, m, t)\|^2 \right]. \quad (12)$$

- Finetune the encoder \mathcal{E}_θ and diffusion network ϵ_θ by optimizing Eq. (11) to effectively remove shadows and preserve the underlying texture.

To demonstrate the effectiveness of our proposed latent feature guidance, Fig. 5 compares it with existing guidance strategies in shadow removal literature, including [44], which conditions on estimated coarse de-shadowed images, and [53], which conditions on estimated invariant color maps for restoration. Our approach preserves more shadow-free perceptual details compared to the estimated invariant color map, which only consists of large color blocks. Similarly, our approach retains fewer shadow features compared to the estimated coarse de-shadowed image, which still retains shadow boundaries that may lead to incorrect results.

3.3. Dense Latent Variable Fusion Module

The phenomenon known as posterior collapse occurs when the training procedure of generative models falls into a trivial local optimum of L_{VLB} , causing the model to ignore the latent variable and collapse the model posterior to the prior, which has only been discussed in VAE [16]. Given the intrinsic similarity between diffusion models and VAE, we first determine the collapse issue of diffusion models under guidance and then address it with a new module.

In our proposed diffusion models, we parameterize the variational distribution $p_\theta(\mathbf{y}_{t-1} | \mathbf{y}_t, \mathbf{y}_0)$ with the latent variable \mathbf{y}_t under the guidance $\mathcal{E}_\theta(\mathbf{x}, m)$ in Eq. (9). In this case, the local optima are characterized by:

$$\begin{aligned} p_\theta(\mathbf{y}_{t-1}) &= p_\theta(\mathbf{y}_{t-1} | \mathbf{y}_t, \mathcal{E}_\theta(\mathbf{x}, m), t) \\ &= p_\theta(\mathbf{y}_{t-1} | \mathbf{y}_t) p_\theta(\mathbf{y}_{t-1} | \mathcal{E}_\theta(\mathbf{x}, m), t) \quad (13) \\ &:= p_\theta(\mathbf{y}_{t-1} | \mathcal{E}_\theta(\mathbf{x}, m), t). \end{aligned}$$

This is undesirable since a crucial goal of diffusion models is to produce diverse outputs. This is particularly important for shadow removal, where complex shadow distributions exist that cannot be easily represented by guidance alone.

Much attention has been devoted to remedying the posterior collapse of VAE models. However, some of these methods weaken the encoder or modeling capability of posterior-related components, as observed in [10, 16]. Other approaches, such as those proposed in [42, 48], significantly complicate the optimization.

In this work, we introduce a new Dense Latent Variable Fusion (DLVF) module that works in tandem with the diffusion network to establish strong links between the latent variable and the generated results. To elaborate, for each

block $\mathcal{G}(\cdot)^n$ of the diffusion network [4] at level n , the feature h^{n-1} and the embedding of $\text{emb}(\mathbf{y}_t)$ generated by a three-layer MLP are both inputted as:

$$h^n = \mathcal{G}(h^{n-1}, \text{emb}(\mathbf{y}_t) \downarrow_{2n}, t)^n, \quad (14)$$

where \downarrow_{2n} denotes the pooling operation with a scaling factor of $2n$ to match the dimension of the features h^{n-1} . To achieve a larger receptive field for the latent variable embedding, we employ fully-connected layers as an additional encoder before inputting them into the network and use adaptive pooling operations to transform the noise into vectors:

$$\mathbf{y}'_t = \mathcal{P}_{pooling}(\mathcal{E}_{noise}(\mathbf{y}_t)), \mathbf{y}'_t \in \mathbb{R}^{1 \times N}, \quad (15)$$

where N is the size of the vector noise. The connection between the embedding $\text{emb}(\mathbf{y}_t) \downarrow_{2n}$ and the network features h^{n-1} is conducted by point-wise summing.

To demonstrate the collapse and effectiveness of our method, we visualize the correspondence between the latent variable \mathbf{y}_t and approximated \mathbf{y}_{t-1} , which are randomly selected from T denoising processes, shown in Fig. 6. In comparison to the baseline without our fusion strategies, our method shows a stronger correspondence between the two variables, indicating a better optimum in the training dynamics. This ultimately results in more effective removal.

4. Experiments

We provide further implementation details, including the settings of the network and optimizer, in the supplemental.

Shadow Removal Benchmarks. We conduct both quantitative and qualitative comparisons on three benchmarks: ISTD [46], AISTD [25], and SRD [34]. The ISTD dataset is a real-world shadow-removal benchmark that consists of 1,330 image triplets for training and 540 image triplets for testing. The image triplet includes the shadow image, shadow mask, and the corresponding shadow-free image. The shadow mask is extracted from the binary difference between the shadow image and the shadow-free image. The AISTD dataset uses the same scene as the ISTD dataset but avoids inconsistent color between the shadow and shadow-free image for accurate comparisons. SRD contains different scenes and consists of 2,680 image pairs for training and 408 image pairs for testing. Since SRD does not contain binary masks for the shadow regions, we follow the common practice and use the masks generated by Cu et al. [3]. For data processing, we empirically dilate all shadow masks in a kernel size of $k = 21$ to address incomplete shadow masks.

Instance-level Shadow Removal Benchmark. We conduct various experiments with visual comparisons on shadow images collected from the internet. The major difference between the above benchmarks and instance-shadow images is the number of shadows in the image, whereas the latter usually has more than one shadow instances. The major collections of our instance-shadow images come from the shadow object association (SOBA)

Table 1. **Quantitative result comparisons of our methods and the state-of-the-art methods on AISTD.** The best and second-best performance is indicated with **bold** and *italic* respectively. We use \uparrow and \downarrow to suggest better high/low score.

		<i>shadow region</i>			<i>non-shadow region</i>			<i>all image</i>		
Method		RMSE \downarrow	PSNR \uparrow	SSIM \uparrow	RMSE \downarrow	PSNR \uparrow	SSIM \uparrow	RMSE \downarrow	PSNR \uparrow	SSIM \uparrow
SP+M-Net [25]	ICCV-19	5.93	37.96	0.990	3.05	35.77	0.973	3.51	32.90	0.957
DHAN [3]	AAAI-20	11.38	33.18	0.987	7.15	27.10	0.972	7.81	25.65	0.954
Param+M+D-Net [26]	ECCV-20	9.67	33.46	0.985	2.91	34.85	0.974	3.98	30.13	0.945
G2R-ShadowNet [27]	CVPR-21	7.38	36.24	0.988	3.00	35.26	0.975	3.69	31.90	0.953
Auto-Exposure [11]	CVPR-21	6.57	36.30	0.976	3.83	31.10	0.874	4.27	29.44	0.838
DC-ShadowNet [23]	ICCV-21	10.57	32.15	0.976	3.82	34.99	0.969	4.80	28.75	0.925
EMDN [54]	AAAI-22	7.94	36.44	0.986	4.78	31.80	0.962	5.28	29.98	0.940
SG-ShadowNet [44]	ECCV-22	5.93	37.25	0.989	3.00	35.27	0.975	3.46	32.42	0.956
BMN [53]	CVPR-22	<i>5.69</i>	<i>38.00</i>	0.991	<i>2.52</i>	<i>37.35</i>	0.981	<i>3.02</i>	<i>33.93</i>	<i>0.966</i>
Palette Diffusion [37]	SIGGRAPH-22	15.40	-	-	7.82	-	-	6.41	-	-
Repaint Diffusion [30]	CVPR-22	12.90	-	-	10.66	-	-	24.90	-	-
LFG-Diffusion	Ours	5.15	39.36	0.991	2.47	37.69	0.981	2.90	34.69	0.968
<i>shadow image</i>		39.72	20.87	0.944	2.51	36.63	0.980	8.38	20.45	0.908

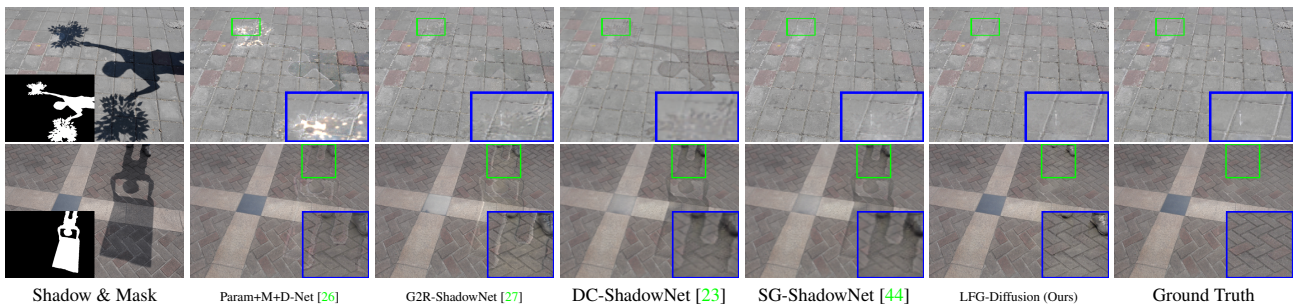


Figure 7. **Visual comparisons of the representative hard shadow removal results on AISTD dataset.** Here we highlight the details of shadow regions that are marked with green box in the blue box area, where ours best perseveres details and removes shadow effects. Please see the supplement for additional visual results.

dataset [47]. We use the manually manipulated shadow-free images of the DESOBA dataset [19] as ground truths for removing shadows at the instance level. For the network training, we synthesize shadow image triplets following the method proposed by Inoue et al. [22]. Please see the supplement for a deep analysis of the synthesized data.

4.1. Performance Evaluation

We evaluate our proposed algorithm against state-of-the-art shadow-removal methods, including SP+M-Net [25], DHAN [3], Param+M+D-Net [26], G2R-ShadowNet [27], Auto-Exposure [11], DC-ShadowNet [23], EMDN [54], BMN [53], and SG-ShadowNet [44], as well as two representative image restoration diffusion models, Palette Diffusion [30], and Repaint Diffusion [37]. The evaluation metrics include the Root Mean Square Error (RMSE) between the shadow-free results and the ground truth in the LAB color space as well as the Peak Signal-to-Noise Ratio (PSNR) and structural similarity (SSIM) in the RGB space. We also provide the metrics measured on the whole image and non-shadow region for reference. Following previous methods [11, 26, 27], we interpolate the results with a resolution of 256×256 for evaluation. We also present the metrics evaluated on the shadow images for reference.

The first row in Tab. 1 shows the quantitative results on the AISTD dataset. Compared with the representative end-to-end learning-based methods, including EMDN [54], Auto-Exposure [11] and DHAN [3], ours significantly outperforms them in all regions. The performance gap between them and ours in the *non-shadow* and *full* regions further indicates the superiority of our model in generating high-quality textures of backgrounds. As expected, the comparison between ours and the other generative methods, including BMN [53], DC-ShadowNet [23], and G2R-ShadowNet [27] demonstrate that our method achieves equal performance improvement in different regions. In contrast, the other methods fail in regions with specific textures. The difference suggests the guidance effectiveness of our modeled latent feature, which is capable of balancing the unbalanced guidance from the surrounding non-shadow areas and shadow regions via the invariant loss function to aid the model in preserving texture and color. The results shown in Tab. 2 on the SRD and ISTD datasets further demonstrate the superiority of our method over the others.

Visual comparison from the AISTD dataset in Fig. 7 and SRD dataset in Fig. 8 further validates the effectiveness of our method. As shown in Fig. 7, our method demonstrates robustness to imperfect shadow mask inputs and preserves

Table 2. **Quantitative comparison results of our methods and the state-of-the-art methods on the *ISTD* dataset and *SRD* dataset.** We want to remark on a slight performance drop in the non-shadow region of our method. The reason is that the two benchmarks are unadjusted, which means the shadow and shadow-free image pairs were captured at different lighting environments. The color inconsistency would result in inaccurate non-shadow region and all image measurement.

(a) <i>ISTD</i> dataset results.										(b) <i>SRD</i> dataset result.			
Method	shadow region			non-shadow region			all image			Method	shadow region		
	RMSE ↓	PSNR ↑	SSIM ↑	RMSE ↓	PSNR ↑	SSIM ↑	RMSE ↓	PSNR ↑	SSIM ↑		RMSE ↓	PSNR ↑	SSIM
DHAN [3]	7.53	35.82	<u>0.989</u>	5.33	30.95	0.971	5.68	29.09	0.953	DHAN [3]	8.94	33.67	0.978
G2R-ShadowNet [27]	10.72	31.63	0.975	-	-	-	-	-	-	G2R-ShadowNet [27]	-	-	-
Auto-Exposure [11]	7.82	34.94	0.973	5.59	28.57	0.862	5.94	27.19	0.824	Auto-exposure [11]	8.86	34.93	0.963
DC-ShadowNet [23]	11.43	31.69	0.976	5.86	28.92	0.956	6.62	26.38	0.917	DC-ShadowNet [23]	7.86	36.34	0.970
EMDN [54]	7.94	<u>36.44</u>	0.986	4.78	31.80	0.962	5.28	29.98	0.940	EMDN [54]	9.83	32.48	0.928
SG-ShadowNet [44]	-	-	-	-	-	-	-	-	-	SG-ShadowNet [44]	8.00	35.53	0.974
BMN [53]	<u>7.44</u>	35.73	<u>0.989</u>	4.61	32.73	<u>0.976</u>	<u>5.06</u>	<u>30.26</u>	<u>0.957</u>	BMN [53]	<u>7.40</u>	<u>36.81</u>	0.979
LFG-Diffusion (Ours)	6.41	37.19	0.990	<u>4.65</u>	<u>32.60</u>	0.977	4.93	30.64	0.963	LFG-Diffusion (Ours)	6.81	37.42	0.979
shadow image	32.67	22.43	0.953	6.77	27.27	0.974	10.86	20.56	0.908	shadow image	46.24	19.95	0.889




Figure 8. **Visual comparisons of the representative hard shadow cases from the *SRD* dataset.**

the textures as well as removing other subtle shadow effects.

4.2. Instance Shadow Removal Evaluation

For real-world applications, shadows cast by objects in the scene are usually instance-level; thus, preserving the other shadows while accurately removing the target instance shadow is crucial. Here, we compare our method with the most recent shadow removal work SG-ShadowNet [44] to demonstrate the generalizability of our method, where we finetuned it with the same dataset synthesized for our experiments. Sample results are shown in Fig. 9. Compared with the SG-ShadowNet, ours thoroughly removes the shadow from the images. As far as we know, this is the first work to demonstrate the applicability of instance shadow removal.

4.3. Ablation Study and Analysis

Effects of the DLVF module. In Tab. 3, we investigate the effectiveness of the proposed DLVF module. We use two alternative methods for comparison: a lagged posterior approach [16] for addressing the posterior collapse, which aggressively optimizes the diffusion network before optimizing the latent feature encoder, and the baseline approach that uses a diffusion network without the fusion strategy. The results show that lagged posterior is less effective, with only a slight improvement margin over the baseline, which could be due to the large complexity of diffusion

Table 3. **Effects of different types of strategies for addressing the posterior collapse in diffusion models.** We only show shadow region results that are distinguishing.

Settings	Region	AISTD dataset		
		RMSE ↓	PSNR ↑	SSIM ↑
w/o. fusion	shadow	6.75	36.34	0.990
lagged posterior	shadow	<u>6.65</u>	<u>36.27</u>	0.990
dense fusion (Ours)	shadow	5.92	37.70	0.990

models and difficulty in training. In contrast, our proposed dense fusion scheme outperforms the baseline by a margin of 0.83 RMSE. Moreover, we visually demonstrate its effectiveness by showing the correspondence of the denoised results y_{t-1} and y_t in Fig. 6. These results validate the idea proposed in our DLVF, *i.e.*, fusing more noise features into each block of the diffusion network is a promising approach for alleviating the local optima of training diffusion models.

Effects of Latent Feature Space Guidance. Tab. 4 compares different types of diffusion model guidance for removing shadows, including (a) estimated invariant color map, (b) estimated coarse de-shadowed image, (c) learned latent feature space without invariant loss, and (d) learn latent feature space with our two-stage learning. Our proposed setting achieves a significantly better numerical




Figure 9. Visual comparisons of the real instance shadow removal results on the DeSOBA dataset.

performance compared to the others. Interestingly, the guidance (*i.e.* coarse deshadowed) that provides the most pixel information performs worse than the guidance (*i.e.* invariant color map) that only provides a simple color map. After deeply looking at their visualization in Fig. 5, we observe that even coarse de-shadowed image still contains shadow boundary that may mislead the diffusion models, while the color map omits most shadow features, which demonstrates that only encapsulating shadow-free features is crucial for improving the performance. Correspondingly, our latent feature is acquired by minimizing the difference between the encoded features of shadow and shadow-free images, which implicitly omits shadow features, and it contains more perceptual features because we optimize it together with diffusion models for learning denoising. Therefore, it guides diffusion models with more shadow-free features and outperforms the compared methods.

Table 4. Effects of different types of diffusion model guidance that provides shadow-free priors.

Settings	Region	AISTD dataset		
		RMSE ↓	PSNR ↑	SSIM ↑
invariant color map [53]	shadow	7.72	36.24	0.986
coarse deshadowed [44]	shadow	8.03	35.74	0.988
$\bar{\mathcal{E}}_{\theta}(\mathbf{x}, m)$	shadow	7.59	36.65	0.984
$\mathcal{E}_{\theta}(\mathbf{x}, m)$ (Ours)	shadow	5.92	37.70	0.990

Table 5. Complexity comparisons of our distilled lighter model with the accelerated diffusion solver.

Method	params	time	AISTD dataset (RMSE)		
			shadow	non-shadow	all
BMN	0.4M	1.69s	5.69	2.52	3.02
G2R-ShadowNet	22.8M	0.36s	7.38	3.00	3.69
LFG-Diffusion (Ours)	82.6M	2.76s	5.15	2.47	2.90
LFG-Diffusion (Distilled)	25.5M	0.24s	5.21	2.34	2.94

Model Complexity Analysis. Our work focuses on adapting diffusion models to address shadow removal, and therefore we prioritize exploration over analysis of model complexity and inference time. However, we demonstrate the feasibility of our approach in terms of model complexity and inference time using advanced technologies such as

those proposed in [28] for reducing model parameters without sacrificing performance, and [29] for accelerating diffusion sampling in Tab. 5. We find that even with similar settings, our lighter model outperforms compared methods with better restoration performance and is also faster.

ShadowDiffusion Comparison. Given the similar between the recent ShadowDiffusion [14] (SD) and our method, which both characterize the shadow-free image distribution by conditioning diffusion models, ours further explores shadow removal at the instance level without any modifications to the model. The other difference is majorly in the method complexity, *i.e.*, tackling challenges such as color-mixing and collapse, often arising from direct conditioning on shadow images. SD integrates a pre-trained shadow removal network. In contrast, ours models the shadow-free priors through two-stage learning and mitigates collapse using dense fusion modules. Tab. 6 demonstrates our efficiency in the shadow region of AISTD. (‡We use their evaluation settings that give different numbers.)

Table 6. Quantitative comparison with ShadowDiffusion.

Model	params(↓)	time(↓)	RMSE(↓)
ShadowDiffusion [14]	602.6M*	7.54s*	4.9‡
LFG-Diffusion (Ours)	82.6M	2.76s	5.0‡

5. Conclusion

In this work, we introduced a novel class of diffusion models that significantly outperform existing shadow removal methods at the general and instance level. By incorporating a latent feature space that captures perceptual shadow-free priors, we have shown that this guidance can mitigate the unbalanced guidance issue between shadow and non-shadow areas during restoration. Furthermore, we have proposed the DLVF module, which strengthens the connections between latent variable of noise and the diffusion network to prevent local optimum. Our comprehensive evaluations and analyses have demonstrated the superior effectiveness of our method compared to existing state-of-the-art shadow removal methods. We believe that our proposed diffusion model-based technique has the potential to be applied to other similar ill-posed low-level problems.

References

- [1] Eli Arbel and Hagit Hel-Or. Shadow removal using intensity surfaces and texture anchor points. *IEEE TPAMI*, 2010. 2
- [2] Zipei Chen, Chengjiang Long, Ling Zhang, and Chunxia Xiao. Canet: A context-aware network for shadow removal. In *ICCV*, 2021. 2
- [3] Xiaodong Cun, Chi-Man Pun, and Cheng Shi. Towards ghost-free shadow removal via dual hierarchical aggregation network and shadow matting GAN. In *AAAI*, 2020. 5, 6, 7
- [4] Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. In *NeurIPS*, 2021. 3, 5
- [5] Adji B Dieng, Yoon Kim, Alexander M Rush, and David M Blei. Avoiding latent variable collapse with generative skip models. In *International Conference on Artificial Intelligence and Statistics*, 2019. 2
- [6] Bin Ding, Chengjiang Long, Ling Zhang, and Chunxia Xiao. Argan: Attentive recurrent generative adversarial network for shadow detection and removal. In *ICCV*, 2019. 2
- [7] Graham D Finlayson, Subho S Chatterjee, and Brian V Funt. Color angular indexing. In *ECCV*, 1996. 3
- [8] Graham D Finlayson, Mark S Drew, and Cheng Lu. Entropy minimization for shadow removal. *IJCV*, 2009. 2
- [9] Graham D Finlayson, Steven D Hordley, Cheng Lu, and Mark S Drew. On the removal of shadows from images. *IEEE TPAMI*, 2005. 2
- [10] Hao Fu, Chunyuan Li, Xiaodong Liu, Jianfeng Gao, Asli Celikyilmaz, and Lawrence Carin. Cyclical annealing schedule: A simple approach to mitigating kl vanishing. *arXiv preprint arXiv:1903.10145*, 2019. 3, 5
- [11] Lan Fu, Changqing Zhou, Qing Guo, Felix Juefei-Xu, Hongkai Yu, Wei Feng, Yang Liu, and Song Wang. Auto-exposure fusion for single-image shadow removal. In *CVPR*, 2021. 2, 6, 7
- [12] Brian V. Funt and Graham D. Finlayson. Color constant color indexing. *IEEE TPAMI*, 1995. 3
- [13] Lanqing Guo, Siyu Huang, Ding Liu, Hao Cheng, and Bihan Wen. Shadowformer: Global context helps image shadow removal. *arXiv preprint arXiv:2302.01650*, 2023. 2
- [14] Lanqing Guo, Chong Wang, Wenhan Yang, Siyu Huang, Yufei Wang, Hanspeter Pfister, and Bihan Wen. Shadowdiffusion: When degradation prior meets diffusion model for shadow removal. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14049–14058, 2023. 8
- [15] Ruiqi Guo, Qieyun Dai, and Derek Hoiem. Paired regions for shadow detection and removal. *IEEE TPAMI*, 2012. 2
- [16] Junxian He, Daniel Spokoyny, Graham Neubig, and Taylor Berg-Kirkpatrick. Lagging inference networks and posterior collapse in variational autoencoders. *arXiv preprint arXiv:1901.05534*, 2019. 2, 3, 5, 7
- [17] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *NeurIPS*, 2020. 1, 4
- [18] Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. *arXiv preprint arXiv:2207.12598*, 2022. 3
- [19] Yan Hong, Li Niu, and Jianfu Zhang. Shadow generation for composite image in real-world scenes. In *AAAI*, 2022. 6
- [20] Andrew Hou, Michel Sarkis, Ning Bi, Yiyong Tong, and Xiaoming Liu. Face relighting with geometrically consistent shadows. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4217–4226, 2022. 1
- [21] Xiaowei Hu, Yitong Jiang, Chi-Wing Fu, and Pheng-Ann Heng. Mask-shadowgan: Learning to remove shadows from unpaired data. In *ICCV*, 2019. 2
- [22] Naoto Inoue and Toshihiko Yamasaki. Learning from synthetic shadows for shadow detection and removal. *IEEE TCSVT*, 2020. 6
- [23] Yeying Jin, Aashish Sharma, and Robby T Tan. Dshadownet: Single-image hard and soft shadow removal using unsupervised domain-classifier guided network. In *ICCV*, 2021. 2, 6, 7
- [24] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013. 3
- [25] Hieu Le and Dimitris Samaras. Shadow removal via shadow image decomposition. In *ICCV*, 2019. 2, 5, 6
- [26] Hieu Le and Dimitris Samaras. From shadow segmentation to shadow removal. In *ECCV*, 2020. 2, 6
- [27] Zhihao Liu, Hui Yin, Xinyi Wu, Zhenyao Wu, Yang Mi, and Song Wang. From shadow generation to shadow removal. In *CVPR*, 2021. 2, 6, 7
- [28] Chengqiang Lu, Jianwei Zhang, Yunfei Chu, Zhengyu Chen, Jingren Zhou, Fei Wu, Haiqing Chen, and Hongxia Yang. Knowledge distillation of transformer-based language models revisited. *arXiv preprint arXiv:2206.14366*, 2022. 8
- [29] Cheng Lu, Yuhao Zhou, Fan Bao, Jianfei Chen, Chongxuan Li, and Jun Zhu. Dpm-solver: A fast ode solver for diffusion probabilistic model sampling in around 10 steps. *arXiv preprint arXiv:2206.00927*, 2022. 8
- [30] Andreas Lugmayr, Martin Danelljan, Andres Romero, Fisher Yu, Radu Timofte, and Luc Van Gool. Repaint: Inpainting using denoising diffusion probabilistic models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11461–11471, 2022. 2, 6
- [31] Haijian Ma, Qiming Qin, and Xinyi Shen. Shadow segmentation and compensation in high resolution satellite images. In *IGARSS 2008-2008 IEEE International Geoscience and Remote Sensing Symposium*, volume 2, pages II–1036, 2008. 1
- [32] Chenlin Meng, Yutong He, Yang Song, Jiaming Song, Jiajun Wu, Jun-Yan Zhu, and Stefano Ermon. Sdedit: Guided image synthesis and editing with stochastic differential equations. In *International Conference on Learning Representations*, 2021. 2
- [33] Konpat Preechakul, Nattanat Chatthee, Suttisak Widadwongsa, and Supasorn Suwajanakorn. Diffusion autoencoders: Toward a meaningful and decodable representation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10619–10629, 2022. 3
- [34] Liangqiong Qu, Jiandong Tian, Shengfeng He, Yandong Tang, and Rynson WH Lau. Dshadownet: A multi-context embedding deep network for shadow removal. In *CVPR*, 2017. 5

- [35] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*, 2022. 3
- [36] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *CVPR*, 2022. 1, 3
- [37] Chitwan Saharia, William Chan, Huiwen Chang, Chris Lee, Jonathan Ho, Tim Salimans, David Fleet, and Mohammad Norouzi. Palette: Image-to-image diffusion models. In *ACM SIGGRAPH 2022 Conference Proceedings*, pages 1–10, 2022. 2, 6
- [38] Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily Denton, Seyed Kamyar Seyed Ghasemipour, Burcu Karagol Ayan, S Sara Mahdavi, Rapha Gontijo Lopes, et al. Photorealistic text-to-image diffusion models with deep language understanding. *arXiv preprint arXiv:2205.11487*, 2022. 3
- [39] Chitwan Saharia, Jonathan Ho, William Chan, Tim Salimans, David J Fleet, and Mohammad Norouzi. Image super-resolution via iterative refinement. *IEEE TPAMI*, 2022. 1
- [40] Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *ICML*, 2015. 1
- [41] Markus Andreas Stricker and Markus Orengo. Similarity of color images. In *Storage and retrieval for image and video databases III*, volume 2420, pages 381–392. SPIE, 1995. 3
- [42] Ilya Tolstikhin, Olivier Bousquet, Sylvain Gelly, and Bernhard Schoelkopf. Wasserstein auto-encoders. *arXiv preprint arXiv:1711.01558*, 2017. 3, 5
- [43] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NeurIPS*, 2017. 1
- [44] Jin Wan, Hui Yin, Zhenyao Wu, Xinyi Wu, Yanting Liu, and Song Wang. Style-guided shadow removal. In *ECCV*, 2022. 2, 3, 4, 5, 6, 7, 8
- [45] Jin Wan, Hui Yin, Zhenyao Wu, Xinyi Wu, Zhihao Liu, and Song Wang. Crformer: A cross-region transformer for shadow removal. *arXiv preprint arXiv:2207.01600*, 2022. 2
- [46] Jifeng Wang, Xiang Li, and Jian Yang. Stacked conditional generative adversarial networks for jointly learning shadow detection and shadow removal. In *CVPR*, 2018. 5
- [47] Tianyu Wang, Xiaowei Hu, Qiong Wang, Pheng-Ann Heng, and Chi-Wing Fu. Instance shadow detection. In *CVPR*, 2020. 6
- [48] Zichao Yang, Zhiting Hu, Ruslan Salakhutdinov, and Taylor Berg-Kirkpatrick. Improved variational autoencoders for text modeling using dilated convolutions. In *International conference on machine learning*, pages 3881–3890. PMLR, 2017. 5
- [49] Edward Zhang, Ricardo Martin-Brualla, Janne Kontkanen, and Brian L Curless. No shadow left behind: Removing objects and their shadows using approximate lighting and geometry. In *CVPR*, 2021. 1
- [50] Ling Zhang, Chengjiang Long, Xiaolong Zhang, and Chunxia Xiao. Ris-gan: Explore residual and illumination with generative adversarial networks for shadow removal. In *AAAI*, 2020. 2
- [51] Wuming Zhang, Xi Zhao, Jean-Marie Morvan, and Liming Chen. Improving shadow suppression for illumination robust face recognition. *IEEE TPAMI*, 2018. 1
- [52] Shengjia Zhao, Jiaming Song, and Stefano Ermon. Infovae: Information maximizing variational autoencoders. *arXiv preprint arXiv:1706.02262*, 2017. 2
- [53] Yurui Zhu, Jie Huang, Xueyang Fu, Feng Zhao, Qibin Sun, and Zheng-Jun Zha. Bijective Mapping Network for Shadow Removal. In *CVPR*, 2022. 3, 4, 5, 6, 7, 8
- [54] Yurui Zhu, Zeyu Xiao, Yanchi Fang, Xueyang Fu, Zhiwei Xiong, and Zheng-Jun Zha. Efficient model-driven network for shadow removal. In *AAAI*, 2022. 6, 7