

---

# Neural Gaffer: Relighting Any Object via Diffusion

---

Haian Jin<sup>1</sup> Yuan Li<sup>2</sup> Fujun Luan<sup>3</sup> Yuanbo Xiangli<sup>1</sup> Sai Bi<sup>3</sup>  
Kai Zhang<sup>3</sup> Zexiang Xu<sup>3</sup> Jin Sun<sup>4</sup> Noah Snavely<sup>1</sup>

<sup>1</sup>Cornell Tech, Cornell University <sup>2</sup>Zhejiang University  
<sup>3</sup>Adobe Research <sup>4</sup>University of Georgia

Project Website: <https://neural-gaffer.github.io/>

## Abstract

Single-image relighting is a challenging task that involves reasoning about the complex interplay between geometry, materials, and lighting. Many prior methods either support only specific categories of images, such as portraits, or require special capture conditions, like using a flashlight. Alternatively, some methods explicitly decompose a scene into intrinsic components, such as normals and BRDFs, which can be inaccurate or under-expressive. In this work, we propose a novel end-to-end 2D relighting diffusion model, called *Neural Gaffer*<sup>1</sup>, that takes a single image of any object and can synthesize an accurate, high-quality relit image under any novel environmental lighting condition, simply by conditioning an image generator on a target environment map, without an explicit scene decomposition. Our method builds on a pre-trained diffusion model, and fine-tunes it on a synthetic relighting dataset, revealing and harnessing the inherent understanding of lighting present in the diffusion model. We evaluate our model on both synthetic and in-the-wild Internet imagery and demonstrate its advantages in terms of generalization and accuracy. Moreover, by combining with other generative methods, our model enables many downstream 2D tasks, such as text-based relighting and object insertion. Our model can also operate as a strong relighting prior for 3D tasks, such as relighting a radiance field.

## 1 Introduction

Lighting plays a key role in our visual world, serving as one of the fundamental elements that shape our interpretation and interaction with 3D space. The interplay of light and shadows can highlight textures, reveal contours, and enhance the perception of shape and form. In many cases, there is a desire to relight an image—that is, to modify it as if it were captured under different lighting conditions. Such a relighting capability enables photo enhancement (e.g., relighting portraits [18; 67]), facilitates consistent lighting across various scenes for filmmaking [59; 22], and supports the seamless integration of virtual objects into real-world environments [38; 19].

However, relighting single images is a challenging task because it involves the complex interplay between geometry, materials, and illumination. Hence, many classical model-based inverse rendering approaches aim to explicitly recover shape, material properties, and lighting from input images [2; 75; 47; 83; 87], so that they can then modify these components and rerender the image. These methods often require multi-view inputs and can suffer from: 1) *model limitations* that prevent faithful estimation of complex light, materials, and shape in real-world scenes, and 2) *ambiguities* in determining these factors without strong data-driven priors.

To circumvent these issues, image-based relighting techniques [18; 72; 76] avoid explicit scene reconstruction and focus on high-quality image synthesis. However, some image-based methods

---

<sup>1</sup>In film and television crews, the gaffer is responsible for managing lighting, including associated resources such as labor, lighting instruments, and electrical equipment

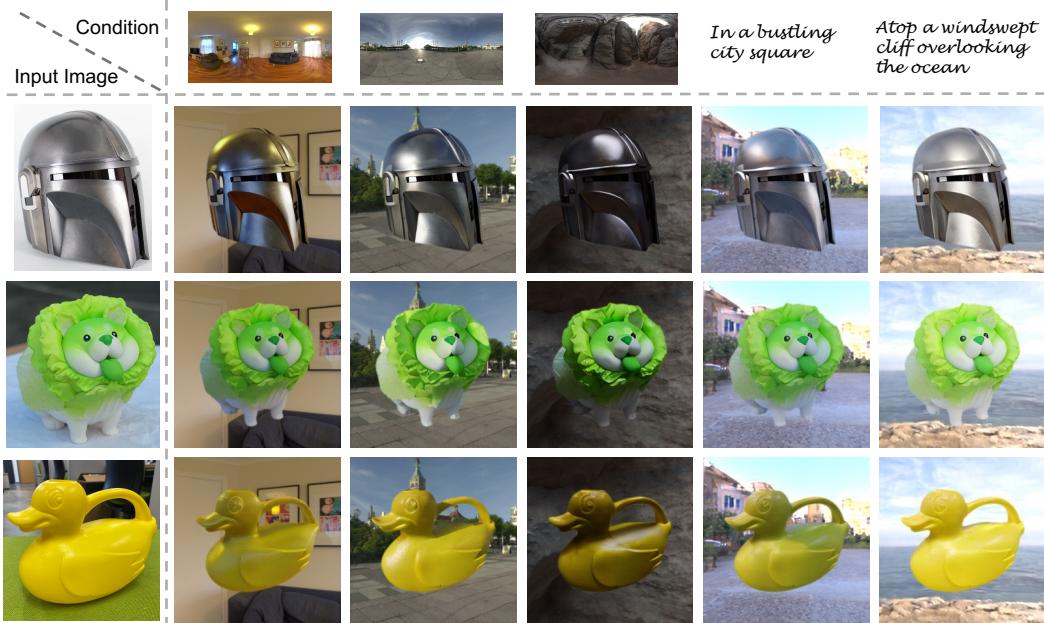


Figure 1: **Single-image relighting results on real data.** Neural Gaffer supports single-image relighting for various input images under diverse lighting conditions, using either image-conditioned input (i.e., an environment map) or text-conditioned input (i.e., a description of the target lighting). These results demonstrate our model’s capability to adapt to diverse lighting scenarios while preserving the visual fidelity of the original objects. Our relighting results remain consistent with the lighting rotation. Please see the supplementary webpage for additional video results produced from real input images.

[18; 58; 76; 44; 3] require inputs under multiple lighting conditions or even sophisticated capture setups (e.g., a light stage [18]), while recent deep learning-based techniques [67; 51; 89] enable single-view relighting but often work only on specific categories like portraits.

In this paper, we propose a novel category-agnostic approach for single-view relighting that tackles the challenges mentioned above. We introduce an end-to-end 2D relighting diffusion model, named *Neural Gaffer*, that takes a single image of any object as input and synthesizes an accurate and high-quality relit image under any environmental lighting condition, specified as a high-dynamic range (HDR) environment map. In contrast to previous learning-based single-view models that are trained for specific categories, we leverage powerful diffusion models, trained on diverse data, and enable relighting of objects from arbitrary categories. Unlike prior model-based approaches, our diffusion-based data-driven framework enables the model to learn physical priors from an object-centric synthetic dataset featuring physically-based materials and High Dynamic Range (HDR) environment maps. This facilitates more accurate lighting effects compared to traditional methods.

Our model exhibits superior generalization and accuracy on both synthetic and real-world images (See Fig. 1). It also seamlessly integrates with other generative methods for various 2D image editing tasks, such as object insertion. Neural Gaffer also can serve as a robust relighting prior for neural radiance fields, facilitating a novel two-stage 3D relighting pipeline. Hence, our diffusion-based solution can relight any object under any environmental lighting condition for both 2D and 3D tasks. By harnessing the power of diffusion models, our approach generalizes effectively across diverse scenes and lighting conditions, overcoming the limitations of prior model-based and image-based methods.

## 2 Related Work

**Diffusion models.** Recently, diffusion-based image generation models [26; 49] have dominated visual content generation tasks. Stable Diffusion [60] not only generates impressive results in the text-to-image setting, but also enables a wide range of image applications. ControlNet [84] adds trainable neural network modules to the U-Net structure of diffusion models to generate images

conditioned on edges, depth, human pose, etc. Diffusion models can also change other aspects of an input image [45; 9], such as image style, via image-to-image models [52]. ControlCom [80] and AnyDoor [14] are image composition models that can blend a foreground object into a target image. Huang et al. provide a comprehensive review of diffusion-based image editing [28].

Recent work has also shown that, despite being trained on 2D images, other kinds of information can be unlocked from diffusion models. A notable example is Zero-1-to-3 [42], which uncovers 3D reasoning capabilities in Stable Diffusion by fine-tuning it on a view synthesis task using synthetic data, and conditioning the model on viewpoint. Similar to the spirit of Zero-1-to-3 and its follow-up work (such as [63; 27]), our work unlocks the latent capability of diffusion models for relighting.

**Single-image relighting.** Relighting imagery is a difficult task that requires an (explicit or implicit) understanding of geometry, materials, lighting, and their interaction through light transport. Single-image relighting is especially challenging due to its ill-posedness; there are infinite explanations of a single image in terms of underlying factors. Classic single-image methods, like SIRFS [2], focus on explicit estimation of scene properties, via optimization with hand-crafted priors. More recent learning-based methods can learn such priors from data, but either focus on specific image types, such as outdoor scenes [73; 69; 78; 40; 21], portraits [53; 71; 64; 67; 48; 51; 23; 56], or human bodies [31; 37; 29], or require special capture conditions like camera flash [39; 61]. Single-image relighting methods for general objects, with realistic appearance, materials, and lighting, remain elusive. In addition, most recent single-image relighting methods [51; 39; 61; 12] still explicitly decompose an image into components, then reconstruct it with modified lighting. This strategy allows the components to be explicitly controlled, but also limits their realism based on the representational ability of the chosen components, such as the Disney BSDF [11] or Phong reflection model [10]. In contrast, similar to [67; 79], we achieve object relighting without an explicit scene decomposition. By training an image-based and lighting-conditioned diffusion model, we can relight any given object from a single photo.

**3D relighting and inverse rendering.** Inverse rendering methods involve estimating 3D models that can be relit under different lighting conditions, typically from multiple images. Recent advances in 3D, such as NeRF [46] and 3D Gaussian Splatting [32], have shown remarkable novel view synthesis results. However, these representations are not themselves relightable as the physical light transport is baked into the radiance field. Efforts have been made toward creating relightable 3D representations, such as PhySG [83], NeRFactor [87], TensoIR [30] and others [6; 8; 7; 5; 4; 82; 88; 35; 66; 77; 81]. These methods often rely on pure optimization to minimize a rendering loss, sometimes incorporating hand-crafted material and/or lighting regularization terms (e.g., smoothness or sparsity) to prevent overfitting. However, these approaches often struggle with specular highlights and glossy materials. We show that our proposed Neural Gaffer method can be applied in an inverse rendering framework, serving as a powerful prior and regularizer. This is because our model has been trained on a large dataset encompassing diverse lighting and material conditions, effectively learning these priors from high-quality data. Our model improves 3D relighting quality, while baselines tend to produce artifacts due to the baking of lighting information into the color field.

### 3 Method

Neural Gaffer is a generalizable image-based 2D diffusion model for relighting general objects; we finetune a pretrained image-conditioned diffusion model [42] with a synthetic relighting dataset to fulfill this goal. In Sec. 3.1, we describe how we construct a large-scale high-quality relighting dataset from Objaverse [20]. Then, in Sec. 3.2, we discuss the design and training of a 2D diffusion model to support zero-shot single-image relighting. Finally, we showcase applications of our model in downstream tasks, including 2D object insertion and 3D radiance field relighting tasks.

#### 3.1 RelitObjaverse Dataset

Capturing ground-truth lighting for real object images is both time-consuming and challenging. Existing real-world relighting datasets [41; 36; 70] are constrained to a very limited number of objects. To address this limitation, we opt to render a large-scale synthetic object relighting dataset.

We use Objaverse [20] as our data source, which comprises about 800K synthetic 3D object models of varying quality. To avoid 3D object models with poor geometry or texture, we developed a series of

filtering mechanisms. Specifically, we first utilize a classifier trained on a manually annotated subset to categorize and filter out the 3D objects with low geometry and texture quality. In addition, we also design a set of filtering rules based on 3D object attributes, such as geometry bounding boxes, BRDF attributes, and file size, eventually yielding a subset (of  $\sim 90K$  objects) of the original Objaverse with high quality. To ensure abundant lighting variations, we collected 1,870 HDR environment maps from the Internet, and augmented them with horizontal flipping and rotation, resulting in  $\sim 30K$  different environment maps in total.

We render images with the Cycles renderer from Blender [17]. During rendering, we randomly sample 12 camera poses for each object. For each viewpoint, we render 16 images under different environment maps. In contrast to existing datasets that only render with environment maps [16; 62], we additionally render an image under a randomly set area light source to increase lighting diversity. In total, for each object, there are  $12 \times (16 + 1) = 204$  images rendered under different lighting conditions and from different viewpoints. Our final synthetic relighting dataset has about 18.4M  $512 \times 512$  rendered images with ground-truth lighting. We will release our full relighting dataset, named the RelitObjaverse dataset, upon acceptance to facilitate reproducibility.

### 3.2 Lighting-conditioned 2D Relighting Diffusion Model

Given a single object image  $x \in \mathbb{R}^{H \times W \times 3}$  and a new lighting condition  $E \in \mathbb{R}^{H' \times W' \times 3}$  (represented as an HDR environment map), our goal is to learn a model  $f$  that can synthesize a new image  $\hat{x}_E$  with correct appearance changes under the new lighting, while preserving the object's identity:

$$\hat{x}_E = f(x, E). \quad (1)$$

We fine-tune an image-to-image diffusion model using our rendered data. In particular, we adopt the pre-trained Zero-1-to-3 model [42] and discard its components and weights relevant to novel-view generation, repurposing the model for relighting. To incorporate lighting conditions, we extend the input latents of the diffusion model, concatenating the encoded input image and the environment map with the denoising latent. We introduce two design decisions that are critical for the effective integration of lighting into the latent diffusion model.

**Environment map rotation.** In general, each pixel of an environment map  $E$  corresponds to a 3D direction in the absolute world coordinate frame, whereas our relighting is performed in image space, that is, locked to the camera's coordinate frame. Therefore, we need a way for a user to specify the relationship between the environment map and camera coordinate frames. One option is to concatenate the relative 3D rotation to the embedding, similar to how Zero-1-to-3 [42] encodes the relative camera pose for the novel view synthesis task. However, we observe that such a design makes it harder for the network to learn the relighting task, perhaps because it needs to interpret the complex relationship between 3D rotation parameters and the desired rotated lighting. Our solution is instead to rotate the target environment map to align with the target camera coordinate frame, prior to feeding it to the model. After rotation, each pixel of the rotated environment map corresponds to a fixed lighting direction in the target camera coordinate frame.

**HDR-LDR conditioning.** Pixel values in the HDR environment map have a theoretical range of  $[0, +\infty)$ , and large pixel values can lead to unstable diffusion training. To rescale it to  $[0, 1.0]$ , a naive idea is to divide it by its maximum value, but this leads to a loss of total energy information. Also, if the map has a very bright region (e.g., the sun), then most of the other values will approach 0 and the environment map will be too dark after division, causing a loss of lighting detail. To solve this, we split the HDR environment map  $E$  into two maps:  $E_L$  is an LDR representation of  $E$ , computed via standard tone-mapping [57; 24], and  $E_H$  is a nonlinearly normalized representation of  $E$ , computed by a logarithmic mapping and divided by the new maximum value.  $E_L$  encodes lighting detail in low-intensity regions, which is useful for generating relighting details (such as reflection details).  $E_H$  maintains the lighting information across the initial spectrum. Inputting  $E_L$

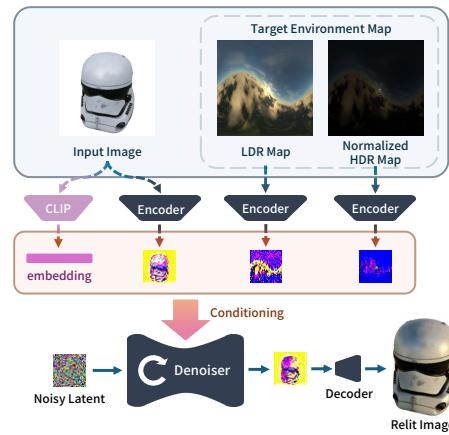


Figure 2: **Model architecture.** Neural Gaffer is an img2img latent diffusion model conditioned on the input image and rotated lighting maps.

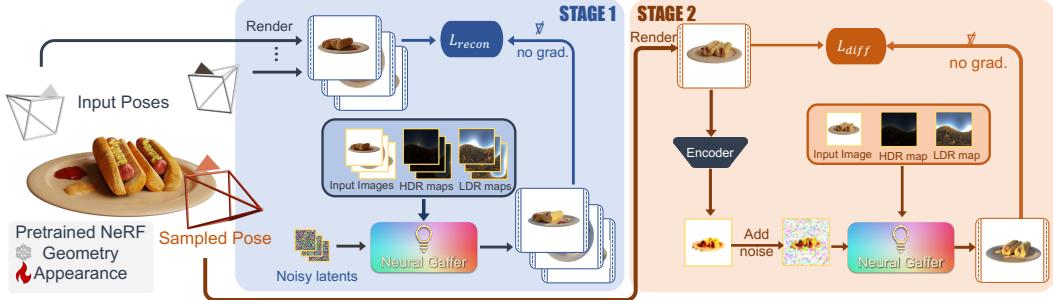


Figure 3: **Relighting a 3D neural radiance field.** Given an input NeRF and a target environmental lighting, in **Stage 1**, we use Neural Gaffer to predict relit images at each predefined camera viewpoint. We then tune the appearance field to overfit the multi-view relighting predictions with a reconstruction loss. In **Stage 2**, we further refine the appearance of the coarsely relit radiance field via the diffusion guidance loss. Using this pipeline, we can relight a NeRF model in minutes.

and  $E_H$  together to the diffusion model means that the network can reason effectively about the full original energy spectrum. Additionally, our design helps to achieve a balanced exposure in the generated output, preventing areas from appearing too dark or too washed out after normalization due to extreme brightness.

**Diffusion model architecture and training.** Our latent diffusion model comprises an encoder  $\mathcal{E}$ , a denoiser U-Net  $\epsilon_\theta$ , and a decoder  $\mathcal{D}$ . During each training iteration, given the paired input-target images  $x$  and  $\hat{x}_E$ , and the target lighting condition  $E$ , we (1) rotate  $E$  into the target camera coordinate frame based on the specified lighting direction, forming a rotated map  $\hat{E}$ ; (2) transform  $\hat{E}$  into an LDR environment map  $\hat{E}_L$  and a normalized HDR environment map  $\hat{E}_H$ ; (3) and finally, encode  $\hat{E}_L$ ,  $\hat{E}_H$  and the input image  $x$  with the encoder  $\mathcal{E}$ . The concatenation of the encoded latent maps are used to condition the denoising U-Net  $\epsilon_\theta$ . Fig. 2 depicts our 2D relighting system. We fine-tune the latent diffusion model with the following objective:

$$\min_{\theta} \mathbb{E}_{z \sim \mathcal{E}(x), t, \epsilon \sim \mathcal{N}(0, 1)} \|\epsilon - \epsilon_\theta(z_t, t, \hat{\mathcal{E}}, c(x))\|_2^2, \quad (2)$$

where  $t \sim [1, 1000]$  is the diffusion time step,  $\hat{\mathcal{E}} = \{\mathcal{E}(x), \mathcal{E}(\hat{E}_L), \mathcal{E}(\hat{E}_H)\}$  represents the union of encoded latents for  $x$ ,  $\hat{E}_L$  and  $\hat{E}_H$ , and  $c(x)$  represents the CLIP embedding of the input image  $x$ . Example 2D relighting results from our trained model, applied to real images, are shown in Fig. 1.

### 3.3 Relighting a 3D radiance field with diffusion prior guidance

We can utilize Neural Gaffer as a prior for 3D relighting. In particular, given a 3D neural radiance field consisting of a density field  $\mathcal{G}_\sigma$  and a view-dependent appearance field  $\mathcal{G}_a$  [46], we consider the task of relighting it using our 2D relighting diffusion model as a prior. We find that simply relighting the input NeRF with the popular Score Distillation Sampling (SDS) method [55] leads to over-saturated and unnatural results, as shown in Fig. 8. We instead propose a two-stage pipeline (Fig. 3) to achieve high-quality 3D relighting in minutes. In the first stage, we use our 2D relighting diffusion model to relight input views and overfit the original NeRF’s appearance field, giving us coarsely relit renderings at sampled viewpoints. In the second stage, our model is used to further refine the appearance details with a diffusion loss and an exponential annealed multi-step scheduling.

**Stage 1: Coarse relighting with reconstruction loss.** Given a NeRF with parameters  $\psi$ , the multi-view input images  $x_{\text{input}} = \{x_i\}_{i=1}^N$  and the target lighting  $E_{\text{target}}$ , we first use our diffusion model  $f$  to generate multi-view relighting prediction images  $\hat{x}_{\text{relit}} = \{\hat{x}_i\}_{i=1}^N$ , where  $\hat{x}_i = f(x_i, E_{\text{target}})$ .

Then, we freeze the density field  $\mathcal{G}_\sigma$  of the NeRF, and only optimize the appearance field of the input NeRF with reconstruction loss  $L_{\text{recon}}$ , aiming to minimize the reconstruction error between a rendered image  $x = x(\psi, \pi_{\text{input}})$  and the relit image  $x_{\text{relit}}$  at pose  $\pi_{\text{input}}$ :

$$\mathcal{L}_{\text{recon}}(\psi) = \|(x(\psi, \pi_{\text{input}}) - x_{\text{relit}})\|_2^2. \quad (3)$$

After this stage, we get an initial blurry and noisy relighting result.

**Stage 2: Detail refinement with diffusion guidance loss.** The results from the first stage can be blurry and can exhibit noise because the multi-view relighting predictions results may be inconsistent.

We further refine the relit NeRF’s appearance in Stage 2. We found that directly using the SDS loss [55] for refinement leads to noisy and over-saturated results (see Fig. 8). Inspired by [68; 74; 45], we treat the blurry rendered images from Stage 1 as an intermediate denoising result. We first render an image  $\hat{x}$  under the sampled camera pose  $\pi_{\text{sampled}}$ . We encode and perturb it to a noisy latent  $\hat{z}_t(\hat{x}, t^*)$  with an intermediate noise level  $t^*$ , and initiate the multi-step DDIM [65] denoising process for  $\hat{z}_t(\hat{x}, t^*)$  from  $t^*$ , yielding a latent sample  $\hat{z}_0$ . This latent is decoded to an image  $\hat{x}_{\text{relit}}(t^*) = D(\hat{z}_0)$ , which we use to supervise the rendering with the following diffusion guidance loss:

$$\mathcal{L}_{\text{diff}}(\psi) = w_1 \|\hat{x} - \hat{x}_{\text{relit}}(t^*)\|_1 + w_2 L_p(\hat{x}, \hat{x}_{\text{relit}}(t^*)), \quad (4)$$

where  $L_p$  is the perceptual distance LPIPS [86] and  $w_1$  and  $w_2$  are scale factors to balance the two terms. We exponentially decrease  $t^*$  during optimization.

## 4 Experiments

### 4.1 Implementation details

**Diffusion model training.** We fine-tune our model starting from Zero123’s [42] checkpoint and discard its original linear projection layer for image embedding and pose. We only fine-tune the UNet of the diffusion model and freeze other parts. During fine-tuning, we use a reduced image size of  $256 \times 256$  and a total batch size of 1024. Both the LDR and normalized HDR environment maps are resized to  $256 \times 256$ . We use AdamW [43] and set the learning rate to  $10^{-4}$  for training. We fine-tune our model for 80K iterations on 8 A6000 GPUs for 5 days.

**3D relighting.** Given a neural radiance field reconstructed from multi-view inputs (we use TensoRF [13]), we first freeze its density field and then optimize the input NeRF’s appearance field for 2,500 iterations in Stage 1, with a ray batch size of 4096. In Stage 2, we set the initial timestep  $t^*$  to decrease from 0.4 to 0.05 exponentially.  $w_1$  and  $w_2$  in Eqn. 4 are set to be 0.5 for all of our test cases. We optimize the results from Stage 1 for 500 iterations, and for each iteration, we sample one view to compute the diffusion guidance loss in Eqn. 4. Stage 1 takes about 5 minutes and Stage 2 takes about 3 minutes on a single A6000 GPU.

**Downstream 2D tasks.** Our model supports downstream 2D image editing tasks, such as text-based relighting and object insertion. For text-based relighting, given a text prompt, we first use Text2Light [15] to generate an HDR panorama, and, using that as an environment map, relight images using our diffusion model. Fig. 1 (right two columns) shows example results; more are on our project webpage. For object insertion, given an object image and a target (perspective) scene image, we segment the foreground object with SAM [34], estimate lighting from the target background image with DiffusionLight [54], perform relighting, and finally generate a shadow attached to the object using [50]. Our workflow can preserve the object’s identity and achieve high-quality visual quality.

### 4.2 Baselines and metrics

We use common image metrics including PNSR, SSIM, and LPIPS [86] to evaluate relighting results quantitatively. In addition, since there exists an ambiguity between lighting intensity and object color, we also compute channel-aligned results as in [87; 88], where each channel of the relighting prediction is rescaled to align the average intensity with the ground-truth. We compute metrics for both the raw and the channel-aligned results for quantitative evaluation.

For the single-image relighting task, there are very few prior works that can handle general objects under natural illumination. We compare to a very recent work DiLightNet [79], which trains a lighting-conditioned ControlNet [84] to relight images, which is achieved by conditioning its ControlNet with a set of rendered images using predefined materials, estimated geometry of the input image, and target lighting. Per our request, the authors ran their model on our validation data. We select 48 high-quality objects from Objaverse as validation objects, which are unseen during training. We render each object under 4 different camera poses. For each camera, we randomly sample 12 unseen environment maps to render the target relit images, and one additional environment map to render the input. We also compare with IC-Light [85]. Since our lighting conditioning is different from its setting (they use background images), we only provide qualitative comparisons in our appendix and project webpage. For relighting 3D radiance fields, we compare with two 3D inverse rendering methods: TensoIR [30] and NVDIFREC-MC [25]. We select 6 objects from the NeRF-Synthetic

dataset [46] and Objaverse [20] for testing. For each object, 100 camera poses are sampled to generate training images and 20 camera poses are sampled for testing images. Each pose is rendered under four distinct, previously unseen environmental lighting conditions. One lighting condition is designated for training images, while the remaining three are reserved for relighting testing.

### 4.3 Results analysis

**Comparing single-image relighting with DiLightNet [79].** As shown in Fig. 4, we evaluate across a diverse set of unseen objects and unseen target lighting conditions. Our approach showcases superior fidelity to the intended lighting settings, maintaining consistent color accuracy and preserving intricate details. Compared with DiLightNet [79], our model can generate more accurate highlights, shadows, and high-frequent reflections. These results demonstrate our model’s ability to accurately understand the underlying physical properties of the object in the input image, such as geometry and materials, and faithfully reproduce the desired relighting appearance. Quantitative results are in Tab. 1.

**Single-image relighting on real data.** We evaluate our single-image relighting model on a set of real-world input photos (Fig. 1), under either image-conditioned (environment map) or text-conditioned (text description of the desired target lighting) target illumination. As shown in the figure, our model effectively adapts to diverse lighting conditions and produces realistic, high-fidelity relighting results under novel illumination. In addition, our relighting results are consistent and stable with the lighting condition switching. Please refer to our webpage to get more video results of more real data.

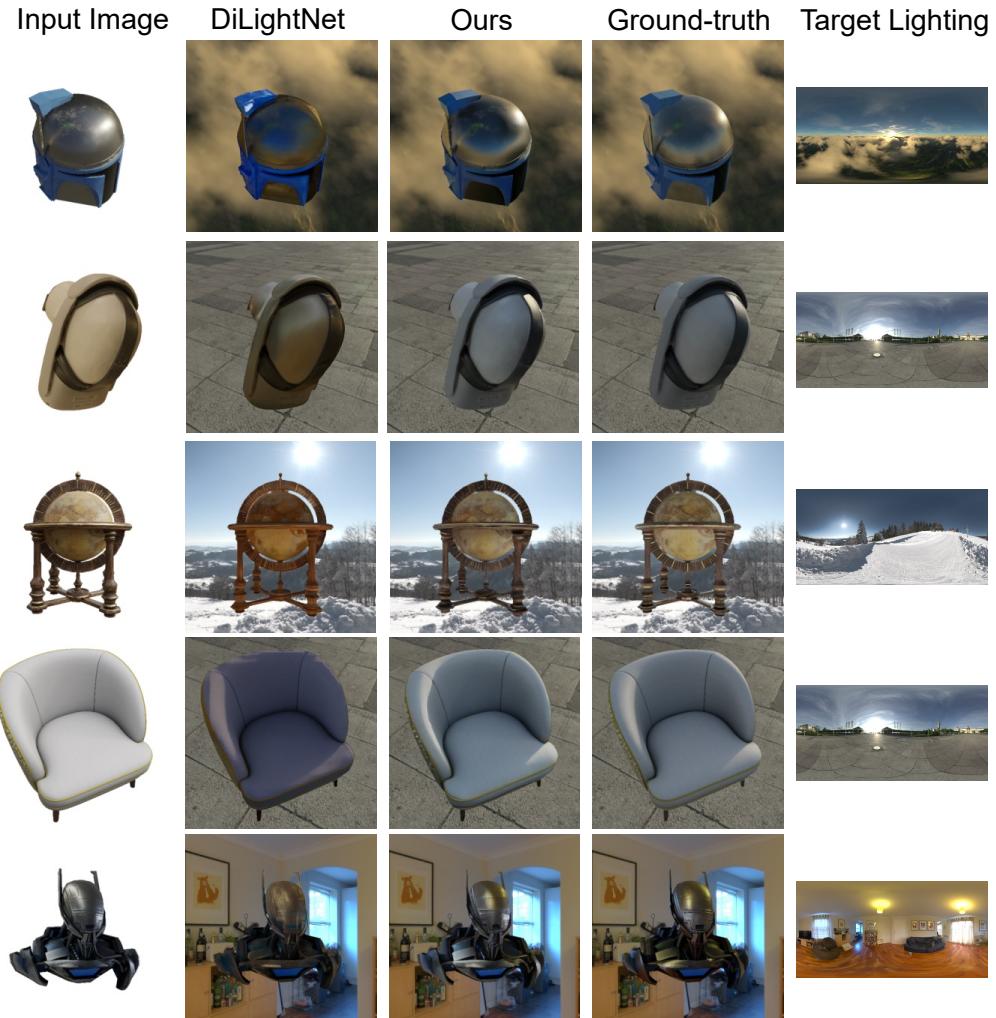


Figure 4: **Single-image relighting comparison with DiLightNet [79] under diverse lighting.** Our method demonstrates superior fidelity to the target lighting, maintains more consistent color and detail, and can generate more accurate highlights, shadows, and high-frequent reflections.

Table 1: Single-image relighting quantitative comparison.

Method	Relighting			Channel-aligned Relighting		
	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$
Ours	<b>26.706</b>	<b>0.927</b>	<b>0.041</b>	<b>29.829</b>	<b>0.939</b>	<b>0.028</b>

Table 2: 3D relighting quantitative comparison.

Method	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$
Ours	<b>29.11</b>	<b>0.930</b>	<b>0.039</b>
TensoIR [30]	25.80	0.878	0.106

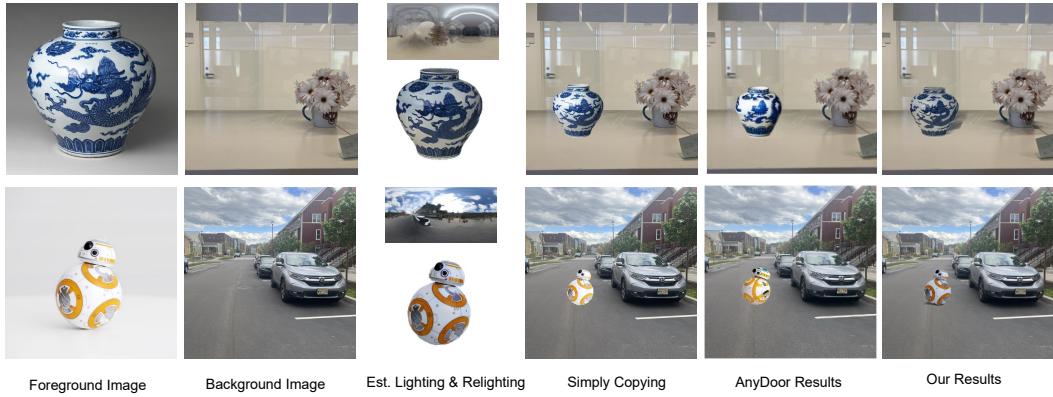


Figure 5: **Object insertion.** Our diffusion model can be applied to object insertion. Compared with AnyDoor [14], our method better preserves the identity of the inserted object and achieves higher-quality results.

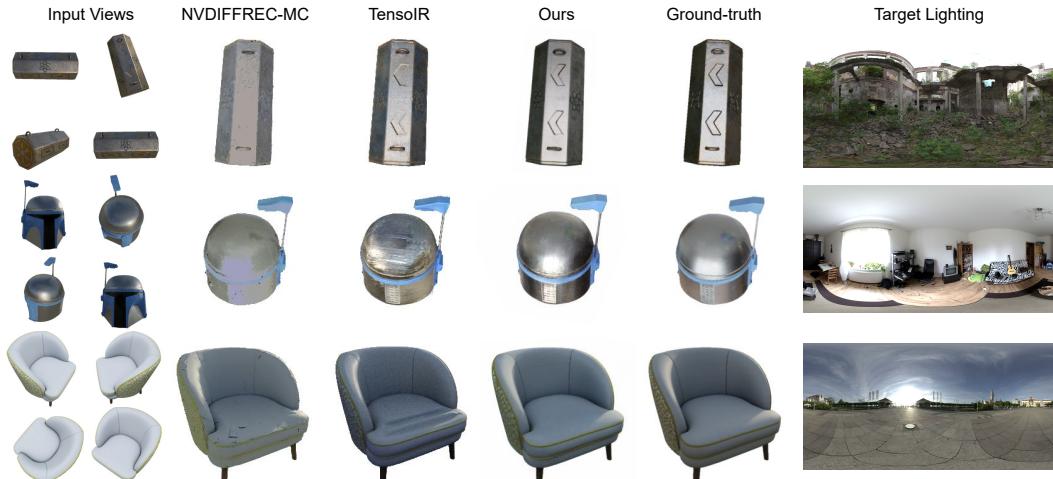


Figure 6: **Relighting 3D objects.** We take multi-view images and a target lighting condition as input and compare our method’s relighting results to NVDIFFREC-MC [25] and TensoIR [30]. (The shown images are channel-aligned results.) Lacking strong priors, the baselines tend to miss specular highlights (NVDIFFREC-MC), bake the original highlights (TensoIR), or produce artifacts (NVDIFFREC-MC and TensoIR). In contrast, our method implicitly learns to model the complex interplay of light and materials from the training process of our relighting diffusion model, resulting in more accurate relighting results with fewer artifacts. Video comparisons are available on our project webpage.

**Relighting 3D objects.** Fig. 6 shows the results of our two-stage method for relighting 3D objects. Given an input radiance field(NeRF) representing a 3D object, we first tune NeRF for coarse relighting in Stage 1, followed by a finer relighting result in Stage 2 that improves fidelity. We also compared to conventional inverse rendering methods such as NVDIFFREC-MC [25] and TensoIR [30], reporting quantitative results in Tab. 2 and showing qualitative comparisons in Fig. 6 and Fig. 13.



Figure 7: **Ablations on the lighting conditioning designs** for our single-image relighting diffusion model.

#### Ablation studies for alternate lighting conditioning designs for our diffusion model.

As ablations, we evaluated different lighting conditioning designs for our diffusion model: 1) only inputting one environment map rather than both the rotated LDR and normalized HDR map, and 2) controlling the lighting direction by adding the relative 3D rotation between the default environment map coordinate frame and the target camera frame to the U-Net CLIP embedding branch without rotating the environment maps explicitly. Due to limited computational resources, when performing ablations for the model design, we train the different models with 2 A6000 GPUs for 48K iterations. We provide a quantitative comparison between them in Tab. 3 and a qualitative comparison in Fig. 7, showcasing the effectiveness of our design. As shown in Fig. 7, controlling the lighting direction by only adding the relative 3D lighting rotation to the U-Net CLIP embedding branch can't get the relighting results with correct lighting directions for all testing cases. Without inputting the LDR map, the model fails to generate the correct reflection for the first test case. Without the HDR map, it cannot produce the correct specular highlights for the third test case. Only when both the LDR and HDR maps are provided can the model

Table 3: Ablations on diffusion model designs.

Method	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$
Ours	<b>26.052</b>	<b>0.923</b>	<b>0.035</b>
Ours w/o LDR map	25.503	0.920	0.045
Ours w/o HDR map	25.822	0.922	0.036
Ours w/o rotation	24.455	0.915	0.051

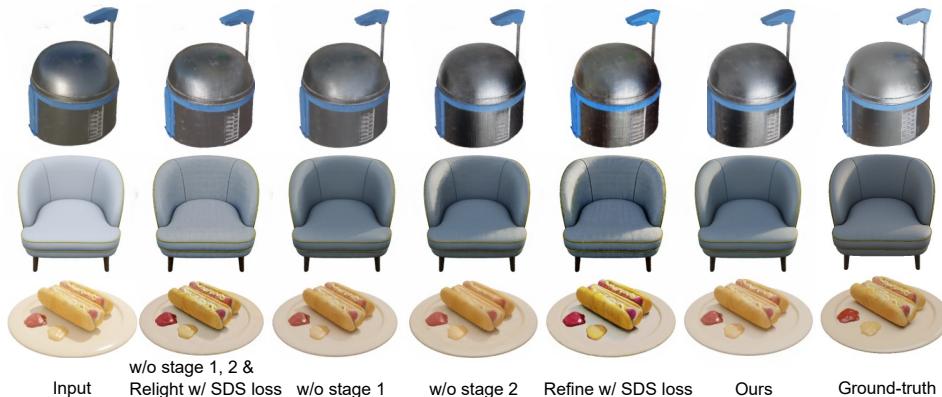


Figure 8: **Ablations on methods for relighting 3D radiance fields.** Our full model achieves the most accurate and visually realistic relighting results among all variants. Please zoom in to see detailed differences across ablation variants.

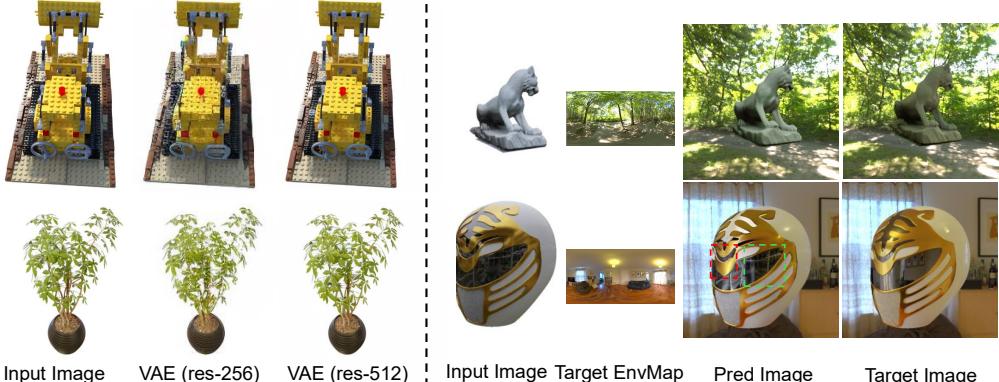


Figure 9: **Limitation Analysis.** The left side shows the encoding-decoding results using the VAE [33] in our base diffusion model [1; 42]. These results indicate that the VAE struggles to preserve identity for objects with fine details at low resolution ( $256 \times 256$ ), which in turn affects the performance of our model on such objects. On the right side, we present several failed cases of our method in Objaverse instances.

second sofa scene, or remove shadows in the third hotdog scene. Stage 1 ensures overall relighting accuracy, while Stage 2 helps reduce noise and leads to more detailed results. Without Stage 1, we cannot generate the correct specular highlight for the first scene, and the details of the shadows on the sofa scene’s backrest become noisy and less accurate. Without Stage 2, the appearance of the relit scenes becomes noisier and blurrier, especially in the sofa scene. Replacing our diffusion guidance loss (Eqn. 4) in Stage 2 with an SDS loss results in unnatural, noisy, and over-saturated relighting. Our full model consistently produces the most accurate and visually realistic results across all test variants. Please zoom in to see detailed differences across ablation variants.

## 5 Limitations

While our model can achieve overall consistent relighting results under rotating or changing lighting conditions, our results can still exhibit minor inconsistencies since the model is generative.

Given the high resource demands of data preprocessing (specifically, rotating the HDR environment map) and model training, and considering our limited university resources, we trained the model at a lower image resolution of  $256 \times 256$ . In Fig. 9, we use VAE [33] in our base diffusion model [1; 42] to encode input images into latent maps, and then directly decode them. We found that VAE struggles to preserve identity for objects with fine details even from latent maps encoded from the input images at this resolution, which in turn results in many relighting failure cases at this resolution. Finetuning our model at a higher resolution will greatly help solve this issue.

We observe two key failure cases when relighting Objaverse Instances, as shown on the right side of Fig. 9. First, our model struggles with color ambiguity in diffuse objects, leading to inaccuracies in relit colors (first row). This issue arises from inherent ambiguities in the input image, where color could be attributed to either material or lighting. Including the background of the input image may help the model better infer the lighting conditions of input images and address this challenge. We give more discussions in the Appendix. C. Second, while the model generates low-frequency highlights well, it often misses high-frequency reflections in relit images (second row).

In addition, we discuss our limitations regarding portrait relighting in the Appendix. B.

## 6 Conclusion

We proposed an end-to-end 2D relighting diffusion model, Neural Gaffer, that addresses the long-standing challenge of single-image relighting through a data-driven framework. Our model synthesizes accurate and high-quality relighted images from a single input image under any environmental lighting condition, demonstrating superior generalization and accuracy on both synthetic and real-world datasets. Neural Gaffer not only enhances various 2D tasks, such as text-based relighting and object insertion but also serves as a strong relighting prior for 3D tasks, enabling a simplified two-stage 3D relighting pipeline. By overcoming the limitations and ambiguities of traditional model-based inverse rendering approaches, our diffusion-based solution extends the applicability of relighting techniques to a broad range of practical applications.

**Acknowledgments.** We thank Minghua Liu, Chao Xu, et al from Prof. Hao Su’s group at UCSD for assisting with the Objaverse filtering process. We thank Xinrui Liu at Cornell for helping with data capture and shadow generation for the object insertion application. We thank Chong Zeng at Zhejiang University for helping run his model (DiLightNet) on our validation data

This work was done while Haian Jin was a full-time student at Cornell. The selection of data and the generation of all figures and results was led by Cornell University. This work was funded in part by the National Science Foundation (IIS-2211259). Jin Sun is partly supported by a gift from Google.

## References

- [1] Stable diffusion image variations - a hugging face space by lambdalabs.
- [2] Jonathan T. Barron and Jitendra Malik. Shape, illumination, and reflectance from shading. *Trans. Pattern Analysis and Machine Intelligence*, 2015.
- [3] Sai Bi, Stephen Lombardi, Shunsuke Saito, Tomas Simon, Shih-En Wei, Kevyn Mcphail, Ravi Ramamoorthi, Yaser Sheikh, and Jason Saragih. Deep relightable appearance models for animatable faces. *ACM Transactions on Graphics (TOG)*, 40(4):1–15, 2021.
- [4] Sai Bi, Zexiang Xu, Pratul Srinivasan, Ben Mildenhall, Kalyan Sunkavalli, Miloš Hašan, Yannick Hold-Geoffroy, David Kriegman, and Ravi Ramamoorthi. Neural reflectance fields for appearance acquisition. *arXiv preprint arXiv:2008.03824*, 2020.
- [5] Sai Bi, Zexiang Xu, Kalyan Sunkavalli, Miloš Hašan, Yannick Hold-Geoffroy, David Kriegman, and Ravi Ramamoorthi. Deep reflectance volumes: Relightable reconstructions from multi-view photometric images. In *Proc. European Conf. on Computer Vision (ECCV)*, pages 294–311. Springer, 2020.
- [6] Mark Boss, Raphael Braun, Varun Jampani, Jonathan T Barron, Ce Liu, and Hendrik Lensch. Nerd: Neural reflectance decomposition from image collections. In *Proc. Int. Conf. on Computer Vision (ICCV)*, pages 12684–12694, 2021.
- [7] Mark Boss, Andreas Engelhardt, Abhishek Kar, Yuanzhen Li, Deqing Sun, Jonathan Barron, Hendrik Lensch, and Varun Jampani. Samurai: Shape and material from unconstrained real-world arbitrary image collections. *Neural Information Processing Systems*, 35:26389–26403, 2022.
- [8] Mark Boss, Varun Jampani, Raphael Braun, Ce Liu, Jonathan Barron, and Hendrik Lensch. Neural-pil: Neural pre-integrated lighting for reflectance decomposition. *Neural Information Processing Systems*, 34:10691–10704, 2021.
- [9] Tim Brooks, Aleksander Holynski, and Alexei A Efros. Instructpix2pix: Learning to follow image editing instructions. In *Proc. Computer Vision and Pattern Recognition (CVPR)*, pages 18392–18402, 2023.
- [10] Phong Bui-Tuong. Illumination for computer generated pictures. *CACM*, 1975.
- [11] Brent Burley and Walt Disney Animation Studios. Physically-based shading at disney. In *Acm Siggraph*, volume 2012, pages 1–7. vol. 2012, 2012.
- [12] Yuanhao Cai, Hao Bian, Jing Lin, Haoqian Wang, Radu Timofte, and Yulun Zhang. Retinex-former: One-stage retinex-based transformer for low-light image enhancement. In *ICCV*, 2023.
- [13] Anpei Chen, Zexiang Xu, Andreas Geiger, Jingyi Yu, and Hao Su. Tensorf: Tensorial radiance fields. In *European Conference on Computer Vision (ECCV)*, 2022.
- [14] Xi Chen, Lianghua Huang, Yu Liu, Yujun Shen, Deli Zhao, and Hengshuang Zhao. Anydoor: Zero-shot object-level image customization, 2024.
- [15] Zhaoxi Chen, Guangcong Wang, and Ziwei Liu. Text2light: Zero-shot text-driven hdr panorama generation, 2023.
- [16] Jasmine Collins, Shubham Goel, Kenan Deng, Achleshwar Luthra, Leon Xu, Erhan Gundogdu, Xi Zhang, Tomas F. Yago Vicente, Thomas Dideriksen, Himanshu Arora, Matthieu Guillaumin, and Jitendra Malik. Abo: Dataset and benchmarks for real-world 3d object understanding, 2022.
- [17] Blender Online Community. *Blender - a 3D modelling and rendering package*. Blender Foundation, Stichting Blender Foundation, Amsterdam, 2018.
- [18] Paul Debevec, Tim Hawkins, Chris Tchou, Haarm-Pieter Duiker, Westley Sarokin, and Mark Sagar. Acquiring the reflectance field of a human face. In *Proceedings of the 27th annual conference on Computer graphics and interactive techniques*, pages 145–156, 2000.
- [19] Paul E. Debevec. Rendering synthetic objects into real scenes: bridging traditional and image-based graphics with global illumination and high dynamic range photography. *Proceedings of the 25th annual conference on Computer graphics and interactive techniques*, 1998.

- [20] Matt Deitke, Dustin Schwenk, Jordi Salvador, Luca Weihs, Oscar Michel, Eli VanderBilt, Ludwig Schmidt, Kiana Ehsani, Aniruddha Kembhavi, and Ali Farhadi. Objaverse: A universe of annotated 3d objects. *arXiv preprint arXiv:2212.08051*, 2022.
- [21] David Griffiths, Tobias Ritschel, and Julien Philip. Outcast: Outdoor single-image relighting with cast shadows. *Computer Graphics Forum*, 41(2):179–193, 2022.
- [22] Kaiwen Guo, Peter Lincoln, Philip L. Davidson, Jay Busch, Xueming Yu, Matt Whalen, Geoff Harvey, Sergio Orts, Rohit Pandey, Jason Dourgarian, Danhang Tang, Anastasia Tkach, Adarsh Kowdle, Emily Cooper, Mingsong Dou, S. Fanello, Graham Fyffe, Christoph Rhemann, Jonathan Taylor, Paul E. Debevec, and Shahram Izadi. The relightables. *ACM Transactions on Graphics (TOG)*, 38:1 – 19, 2019.
- [23] Yuxuan Han, Zhibo Wang, and Feng Xu. Learning a 3d morphable face reflectance model from low-cost data, 2023.
- [24] Samuel W. Hasinoff, Dillon Sharlet, Ryan Geiss, Andrew Adams, Jonathan T. Barron, Florian Kainz, Jiawen Chen, and Marc Levoy. Burst photography for high dynamic range and low-light imaging on mobile cameras. *ACM Trans. Graph.*, 35(6), dec 2016.
- [25] Jon Hasselgren, Nikolai Hofmann, and Jacob Munkberg. Shape, Light, and Material Decomposition from Images using Monte Carlo Rendering and Denoising. *arXiv:2206.03380*, 2022.
- [26] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Neural Information Processing Systems*, 33:6840–6851, 2020.
- [27] Hanzhe Hu, Zhizhuo Zhou, Varun Jampani, and Shubham Tulsiani. Mvd-fusion: Single-view 3d via depth-consistent multi-view generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9698–9707, 2024.
- [28] Yi Huang, Jiancheng Huang, Yifan Liu, Mingfu Yan, Jiaxi Lv, Jianzhuang Liu, Wei Xiong, He Zhang, Shifeng Chen, and Liangliang Cao. Diffusion model-based image editing: A survey. *arXiv preprint arXiv:2402.17525*, 2024.
- [29] Chaonan Ji, Tao Yu, Kaiwen Guo, Jingxin Liu, and Yebin Liu. Geometry-aware single-image full-body human relighting. In *ECCV*, page 388–405, 2022.
- [30] Haian Jin, Isabella Liu, Peijia Xu, Xiaoshuai Zhang, Songfang Han, Sai Bi, Xiaowei Zhou, Zexiang Xu, and Hao Su. Tensoir: Tensorial inverse rendering. In *Proc. Computer Vision and Pattern Recognition (CVPR)*, pages 165–174, 2023.
- [31] Yoshihiro Kanamori and Yuki Endo. Relighting humans: occlusion-aware inverse rendering for full-body human images. *ACM Trans. Graph.*, 37(6), 2018.
- [32] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splatting for real-time radiance field rendering. *ACM Transactions on Graphics*, 42(4):1–14, 2023.
- [33] Diederik P Kingma and Max Welling. Auto-encoding variational bayes, 2022.
- [34] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C. Berg, Wan-Yen Lo, Piotr Dollár, and Ross Girshick. Segment anything, 2023.
- [35] Zhengfei Kuang, Kyle Olszewski, Menglei Chai, Zeng Huang, Panos Achlioptas, and Sergey Tulyakov. Neroic: Neural rendering of objects from online image collections. *ACM Transactions on Graphics (TOG)*, 41(4):1–12, 2022.
- [36] Zhengfei Kuang, Yunzhi Zhang, Hong-Xing Yu, Samir Agarwala, Shangzhe Wu, and Jiajun Wu. Stanford-orb: A real-world 3d object inverse rendering benchmark, 2024.
- [37] Manuel Lagunas, Xin Sun, Jimei Yang, Ruben Villegas, Jianming Zhang, Zhixin Shu, Belen Masia, and Diego Gutierrez. Single-image full-body human relighting. *Eurographics Symposium on Rendering - DL-only Track*, 2021.
- [38] Zhengqin Li, Jia Shi, Sai Bi, Rui Zhu, Kalyan Sunkavalli, Milovs Havsan, Zexiang Xu, Ravi Ramamoorthi, and Manmohan Chandraker. Physically-based editing of indoor scene lighting from a single image. *ArXiv*, abs/2205.09343, 2022.
- [39] Zhengqin Li, Zexiang Xu, Ravi Ramamoorthi, Kalyan Sunkavalli, and Manmohan Chandraker. Learning to reconstruct shape and spatially-varying reflectance from a single image. *ACM Transactions on Graphics (TOG)*, 37:1 – 11, 2018.
- [40] Andrew Liu, Shiry Ginosar, Tinghui Zhou, Alexei A. Efros, and Noah Snavely. Learning to factorize and relight a city, 2020.
- [41] Isabella Liu, Linghao Chen, Ziyang Fu, Liwen Wu, Haian Jin, Zhong Li, Chin Ming Ryan Wong, Yi Xu, Ravi Ramamoorthi, Zexiang Xu, and Hao Su. Openillumination: A multi-illumination dataset for inverse rendering evaluation on real objects, 2024.
- [42] Ruoshi Liu, Rundi Wu, Basile Van Hoorick, Pavel Tokmakov, Sergey Zakharov, and Carl Vondrick. Zero-1-to-3: Zero-shot one image to 3d object, 2023.

- [43] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization, 2019.
- [44] Abhimitra Meka, Christian Haene, Rohit Pandey, Michael Zollhöfer, Sean Fanello, Graham Fyffe, Adarsh Kowdle, Xueming Yu, Jay Busch, Jason Dourgarian, et al. Deep reflectance fields: high-quality facial reflectance field inference from color gradient illumination. *ACM Transactions on Graphics (TOG)*, 38(4):1–12, 2019.
- [45] 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.
- [46] Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, Jonathan T. Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. In *ECCV*, 2020.
- [47] Giljoo Nam, Joo Ho Lee, Diego Gutierrez, and Min H Kim. Practical svbrdf acquisition of 3d objects with unstructured flash photography. *ACM Transactions on Graphics (TOG)*, 37(6):1–12, 2018.
- [48] Thomas Nestmeyer, Jean-François Lalonde, Iain Matthews, and Andreas Lehrmann. Learning physics-guided face relighting under directional light. In *Proc. Computer Vision and Pattern Recognition (CVPR)*, pages 5124–5133, 2020.
- [49] Alexander Quinn Nichol and Prafulla Dhariwal. Improved denoising diffusion probabilistic models. In *International conference on machine learning*, pages 8162–8171. PMLR, 2021.
- [50] Li Niu, Wenyang Cong, Liu Liu, Yan Hong, Bo Zhang, Jing Liang, and Liqing Zhang. Making images real again: A comprehensive survey on deep image composition. *arXiv preprint arXiv:2106.14490*, 2021.
- [51] Rohit Pandey, Sergio Orts Escolano, Chloe Legendre, Christian Haene, Sofien Bouaziz, Christoph Rhemann, Paul Debevec, and Sean Fanello. Total relighting: learning to relight portraits for background replacement. *ACM Trans. Graph.*, 40(4), 2021.
- [52] Gaurav Parmar, Krishna Kumar Singh, Richard Zhang, Yijun Li, Jingwan Lu, and Jun-Yan Zhu. Zero-shot image-to-image translation. In *ACM SIGGRAPH 2023 Conference Proceedings*, pages 1–11, 2023.
- [53] Pieter Peers, Naoki Tamura, Wojciech Matusik, and Paul Debevec. Post-production facial performance relighting using reflectance transfer. *ACM Trans. Graph.*, 26(3), 2007.
- [54] Pakkapon Phongthawee, Worameth Chinchuthakun, Nontaphat Sinsunthithet, Amit Raj, Varun Jampani, Pramook Khungurn, and Supasorn Suwajanakorn. Diffusionlight: Light probes for free by painting a chrome ball. In *ArXiv*, 2023.
- [55] Ben Poole, Ajay Jain, Jonathan T. Barron, and Ben Mildenhall. Dreamfusion: Text-to-3d using 2d diffusion. *arXiv*, 2022.
- [56] Anurag Ranjan, Kwang Moo Yi, Jen-Hao Rick Chang, and Oncel Tuzel. Facelit: Neural 3d relightable faces. In *Proc. Computer Vision and Pattern Recognition (CVPR)*, pages 8619–8628, 2023.
- [57] Erik Reinhard, Greg Ward, Sumanta Pattanaik, and Paul Debevec. *High Dynamic Range Imaging: Acquisition, Display, and Image-Based Lighting (The Morgan Kaufmann Series in Computer Graphics)*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 2005.
- [58] Peiran Ren, Yue Dong, Stephen Lin, Xin Tong, and Baining Guo. Image based relighting using neural networks. *ACM Transactions on Graphics (ToG)*, 34(4):1–12, 2015.
- [59] Christian Richardt, Carsten Stoll, Neil A. Dodgson, Hans-Peter Seidel, and Christian Theobalt. Coherent spatiotemporal filtering, upsampling and rendering of RGBZ videos. *Computer Graphics Forum (Proceedings of Eurographics)*, 31(2), May 2012.
- [60] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proc. Computer Vision and Pattern Recognition (CVPR)*, pages 10684–10695, 2022.
- [61] Shen Sang and M. Chandraker. Single-shot neural relighting and svbrdf estimation. In *ECCV*, 2020.
- [62] Jian Shi, Yue Dong, Hao Su, and Stella X. Yu. Learning non-lambertian object intrinsics across shapenet categories, 2016.
- [63] Ruoxi Shi, Hansheng Chen, Zhuoyang Zhang, Minghua Liu, Chao Xu, Xinyue Wei, Linghao Chen, Chong Zeng, and Hao Su. Zero123++: a single image to consistent multi-view diffusion base model, 2023.
- [64] Zhixin Shu, Ersin Yumer, Sunil Hadap, Kalyan Sunkavalli, Eli Shechtman, and Dimitris Samaras. Neural face editing with intrinsic image disentangling. In *Proc. Computer Vision and Pattern Recognition (CVPR)*, pages 5541–5550, 2017.
- [65] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models, 2022.

- [66] Pratul P Srinivasan, Boyang Deng, Xiuming Zhang, Matthew Tancik, Ben Mildenhall, and Jonathan T Barron. Nerv: Neural reflectance and visibility fields for relighting and view synthesis. In *Proc. Computer Vision and Pattern Recognition (CVPR)*, pages 7495–7504, 2021.
- [67] Tiancheng Sun, Jonathan T Barron, Yun-Ta Tsai, Zexiang Xu, Xueming Yu, Graham Fyffe, Christoph Rhemann, Jay Busch, Paul Debevec, and Ravi Ramamoorthi. Single image portrait relighting. *ACM Trans. Graph.*, 38(4), 2019.
- [68] Jiaxiang Tang, Jiawei Ren, Hang Zhou, Ziwei Liu, and Gang Zeng. Dreamgaussian: Generative gaussian splatting for efficient 3d content creation. *arXiv preprint arXiv:2309.16653*, 2023.
- [69] Murat Türe, Mustafa Ege Çiklabakkal, Aykut Erdem, Erkut Erdem, Pınar Satılmış, and Ahmet Oguz Akyüz. From noon to sunset: Interactive rendering, relighting, and recolouring of landscape photographs by modifying solar position. In *Computer Graphics Forum*, pages 500–515. Wiley Online Library, 2021.
- [70] Benjamin Ummenhofer, Sanskar Agrawal, Rene Sepulveda, Yixing Lao, Kai Zhang, Tianhang Cheng, Stephan Richter, Shenlong Wang, and German Ros. Objects with lighting: A real-world dataset for evaluating reconstruction and rendering for object relighting, 2024.
- [71] Yang Wang, Lei Zhang, Zicheng Liu, Gang Hua, Zhen Wen, Zhengyou Zhang, and Dimitris Samaras. Face relighting from a single image under arbitrary unknown lighting conditions. *Trans. Pattern Analysis and Machine Intelligence*, 31(11):1968–1984, 2008.
- [72] Andreas Wenger, Andrew Gardner, Chris Tchou, Jonas Unger, Tim Hawkins, and Paul Debevec. Performance relighting and reflectance transformation with time-multiplexed illumination. *ACM Transactions on Graphics (TOG)*, 24(3):756–764, 2005.
- [73] Jung-Hsuan Wu and Suguru Saito. Interactive relighting in single low-dynamic range images. *ACM Trans. Graph.*, 36(2), 2017.
- [74] Rundi Wu, Ben Mildenhall, Philipp Henzler, Keunhong Park, Ruiqi Gao, Daniel Watson, Pratul P. Srinivasan, Dor Verbin, Jonathan T. Barron, Ben Poole, and Aleksander Holynski. Reconfusion: 3d reconstruction with diffusion priors. *arXiv*, 2023.
- [75] Rui Xia, Yue Dong, Pieter Peers, and Xin Tong. Recovering shape and spatially-varying surface reflectance under unknown illumination. *ACM Transactions on Graphics (TOG)*, 35(6):1–12, 2016.
- [76] Zexiang Xu, Kalyan Sunkavalli, Sunil Hadap, and Ravi Ramamoorthi. Deep image-based relighting from optimal sparse samples. *ACM Transactions on Graphics (ToG)*, 37(4):1–13, 2018.
- [77] Yao Yao, Jingyang Zhang, Jingbo Liu, Yihang Qu, Tian Fang, David McKinnon, Yanghai Tsin, and Long Quan. Neifl: Neural incident light field for physically-based material estimation. In *European Conference on Computer Vision*, pages 700–716. Springer, 2022.
- [78] Ye Yu, Abhimitra Meka, Mohamed Elgarib, Hans-Peter Seidel, Christian Theobalt, and William AP Smith. Self-supervised outdoor scene relighting. In *ECCV*, pages 84–101, 2020.
- [79] Chong Zeng, Yue Dong, Pieter Peers, Youkang Kong, Hongzhi Wu, and Xin Tong. Dilightnet: Fine-grained lighting control for diffusion-based image generation, 2024.
- [80] Bo Zhang, Yuxuan Duan, Jun Lan, Yan Hong, Huijia Zhu, Weiqiang Wang, and Li Niu. Control-com: Controllable image composition using diffusion model. *arXiv preprint arXiv:2308.10040*, 2023.
- [81] Jingyang Zhang, Yao Yao, Shiwei Li, Jingbo Liu, Tian Fang, David McKinnon, Yanghai Tsin, and Long Quan. Neifl++: Inter-reflectable light fields for geometry and material estimation. In *Proc. Int. Conf. on Computer Vision (ICCV)*, pages 3601–3610, 2023.
- [82] Kai Zhang, Fujun Luan, Zhengqi Li, and Noah Snavely. Iron: Inverse rendering by optimizing neural sdofs and materials from photometric images. In *Proc. Computer Vision and Pattern Recognition (CVPR)*, pages 5565–5574, 2022.
- [83] Kai Zhang, Fujun Luan, Qianqian Wang, Kavita Bala, and Noah Snavely. Physgc: Inverse rendering with spherical gaussians for physics-based material editing and relighting. In *Proc. Computer Vision and Pattern Recognition (CVPR)*, pages 5453–5462, 2021.
- [84] Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. Adding conditional control to text-to-image diffusion models, 2023.
- [85] Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. Ic-light github page, 2024.
- [86] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proc. Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [87] Xiuming Zhang, Pratul P Srinivasan, Boyang Deng, Paul Debevec, William T Freeman, and Jonathan T Barron. Nerfactor: Neural factorization of shape and reflectance under an unknown illumination. *ACM Transactions on Graphics (ToG)*, 40(6):1–18, 2021.

- [88] Yuanqing Zhang, Jiaming Sun, Xingyi He, Huan Fu, Rongfei Jia, and Xiaowei Zhou. Modeling indirect illumination for inverse rendering. In *Proc. Computer Vision and Pattern Recognition (CVPR)*, 2022.
- [89] Hao Zhou, Sunil Hadap, Kalyan Sunkavalli, and David W Jacobs. Deep single-image portrait relighting. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 7194–7202, 2019.

## A Comparison with IC-Light

In this section, we conduct a comparison with a recent single-image relighting method, IC-Light [85]. In Fig. 10, we present a comparative analysis of relighting results, focusing on the consistency of specular highlights in comparison to IC-Light. Our method consistently adjusts the specular highlights in accordance with the rotation of the lighting, ensuring a coherent relighting effect. In contrast, IC-Light demonstrates inconsistent movement of highlights, with certain regions showing minimal changes. For additional video comparisons, please visit our supplemental webpage.

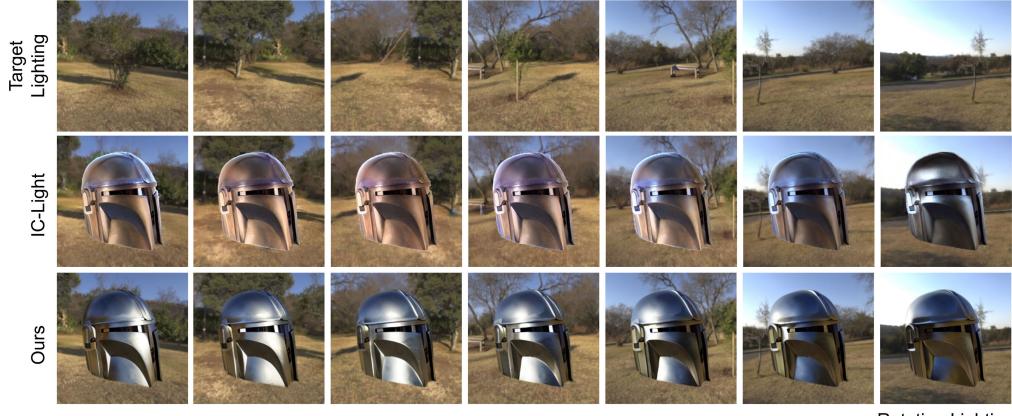


Figure 10: **Comparison of single-image relighting with IC-Light [85].** In this experiment, we compare relighting results and evaluate the consistency of specular highlights against IC-Light [85]. Each column represents a uniformly rotated version of the same environment map image, progressing from left to right. Our method consistently adjusts highlights, faithfully reproducing the object under novel lighting conditions. In contrast, IC-Light [85] exhibits inconsistent highlight movement, with certain regions changing just slightly. Further comparisons are available in videos provided on the supplemental webpage.

## B Portrait Relighting Results

Our model is fine-tuned on object data, therefore, while it is category-agnostic, it might not achieve the highest quality results on specific domains—for instance, for the case of portrait relighting, methods that are specifically trained for portrait photos may produce higher quality results. We show our results on the portrait in Fig. 11 of the appendix.

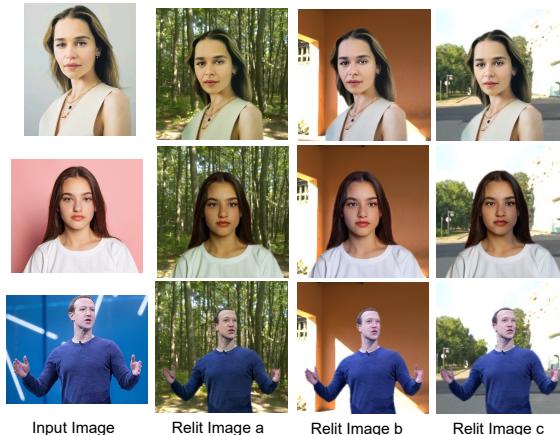


Figure 11: **Human portrait relighting results.** Overall, in most cases, our method shows good generalization to human portraits (see the second row of the figure above). However, there are instances where it produces suboptimal results. For example, in the first row, the model fails to remove the shadow on the neck in the input image. In the failure case shown in the last row, the model does not preserve critical portrait details, such as the shape of the mouth. This issue may stem from the low-resolution limitations discussed in Sec. 5. Further fine-tuning with additional portrait data and at a higher resolution could enhance the model’s performance on portrait images.

### C Inherent color ambiguity



Figure 12: **Color ambiguity analysis.** We found that our data-driven model can handle the inherent color ambiguity to some extent, especially when the objects’ materials in the input image are relatively specular. Because the diffusion model can infer the lighting of the input image through the reflections or specular highlights

Color ambiguity is an inherent issue in the single-image relighting task. That said, we found that our data-driven model can handle this problem to some extent, especially when the objects’ materials in the input image are relatively specular. Because the diffusion model can infer the lighting of the input image through the reflections or specular highlights. In Fig. 12, we show two such examples: we render each object with two different environment maps and relight them with the same environment map. In both cases, the input images have different colors and we want to relight them under the same target lighting. It turned out that both input images can be relit to the results that match the ground truth results, which indicates our model has some ability to handle the color ambiguity.

As for the general diffuse object, we think our model could learn an implicit prior over common colors that show up in object albedos, as well as common colors that occur in lighting, and use that to help resolve this ambiguity. However, this may fail sometimes, as we discussed in Sec.5. One potential solution could be also inputting the background of the input image to the diffusion model since the diffusion model can learn to infer the lighting of the input image from the background. We leave this as future work.

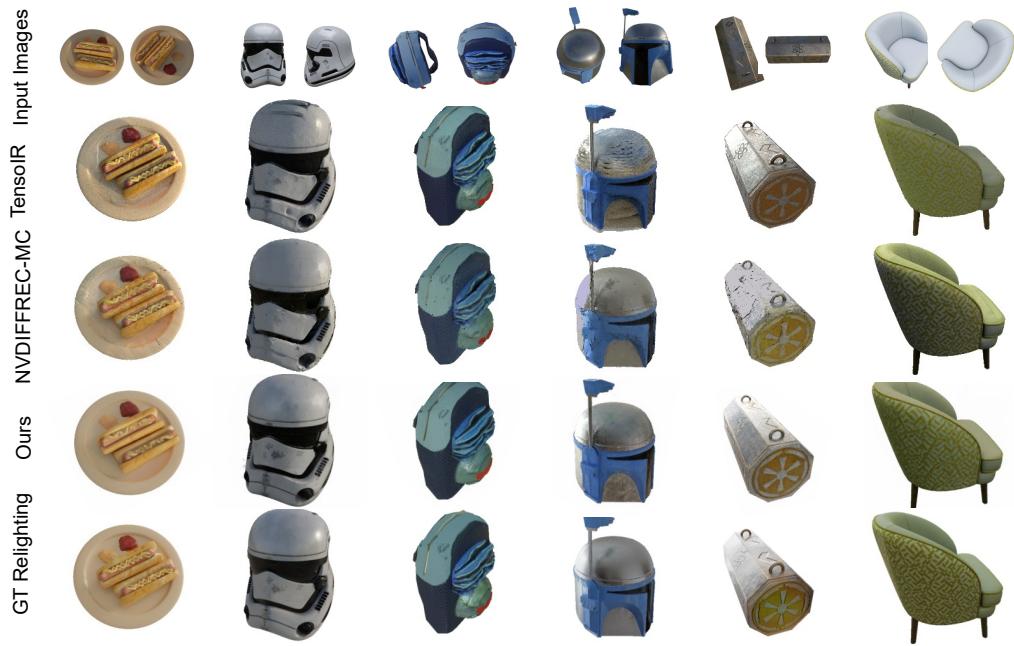


Figure 13: **More results of relighting 3D objects.** We present additional 3D relighting comparison results, which include all of our testing objects. The results show that our methods achieve better qualitative relighting outcomes for all the tested objects

## NeurIPS Paper Checklist

### 1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [\[Yes\]](#)

Justification: The main focus of this paper is to train a lighting-aware image-conditioned diffusion model; it aligns with our title.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

### 2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [\[Yes\]](#)

Justification: Please see the Limitation Section.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

### 3. Theory Assumptions and Proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [\[NA\]](#)

Justification: [NA]

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

#### 4. Experimental Result Reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: We have included enough details to reproduce our results.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
  - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
  - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
  - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
  - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

#### 5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: **[Yes]**

Justification: We will release all code and models once upon acceptance.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

## 6. Experimental Setting/Details

Question: Does the paper specify all the training and test details (e.g., data splits, hyper-parameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: **[Yes]**

Justification: we have described in details our experimental setups in the experiments section.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

## 7. Experiment Statistical Significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: **[No]**

Justification: As each experiment takes significant amount of compute resource, we don't have enough computing resources to complete the significance test.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).

- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

## 8. Experiments Compute Resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [\[Yes\]](#)

Justification: Yes, provide both data creation and model training.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

## 9. Code Of Ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines>?

Answer: [\[Yes\]](#)

Justification: The paper is written on our own and follows the honest codes.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

## 10. Broader Impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [\[No\]](#)

Justification: Please refer to the main paper.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.

- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

## 11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: [NA]

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

## 12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: We have cited all previous papers that lead to our current paper's presentation.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, [paperswithcode.com/datasets](https://paperswithcode.com/datasets) has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.

- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

### 13. New Assets

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [NA]

Justification: No new assets were introduced.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

### 14. Crowdsourcing and Research with Human Subjects

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: [NA]

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

### 15. Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: [NA]

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.