# **UniRelight: Learning Joint Decomposition and Synthesis for Video Relighting**

Kai  $He^{1,2,3}$  Ruofan Liang<sup>1,2,3</sup> Jacob Munkberg<sup>1</sup> Jon Hasselgren<sup>1</sup>

Nandita Vijaykumar<sup>2,3</sup> Alexander Keller<sup>1</sup> Sanja Fidler<sup>1,2,3</sup>

Igor Gilitschenski<sup>2,3†</sup> Zan Gojcic<sup>1†</sup> Zian Wang<sup>1,2,3†</sup>

<sup>1</sup>NVIDIA <sup>2</sup>University of Toronto <sup>3</sup>Vector Institute

#### **Abstract**

We address the challenge of relighting a single image or video, a task that demands precise scene intrinsic understanding and high-quality light transport synthesis. Existing end-to-end relighting models are often limited by the scarcity of paired multi-illumination data, restricting their ability to generalize across diverse scenes. Conversely, two-stage pipelines that combine inverse and forward rendering can mitigate data requirements but are susceptible to error accumulation and often fail to produce realistic outputs under complex lighting conditions or with sophisticated materials. In this work, we introduce a general-purpose approach that jointly estimates albedo and synthesizes relit outputs in a single pass, harnessing the generative capabilities of video diffusion models. This joint formulation enhances implicit scene comprehension and facilitates the creation of realistic lighting effects and intricate material interactions, such as shadows, reflections, and transparency. Trained on synthetic multi-illumination data and extensive automatically labeled real-world videos, our model demonstrates strong generalization across diverse domains and surpasses previous methods in both visual fidelity and temporal consistency. Our project page is https://research.nvidia.com/labs/toronto-ai/UniRelight/.

#### 1 Introduction

Lighting is crucial in defining a scene's visual appearance, influencing both aesthetics and the perception of objects and materials. Modifying illumination while preserving scene content enables a wide range of applications, including creative editing, simulation, and robust vision systems. Realistic relighting necessitates an accurate understanding of scene geometry and materials, as well as faithful light transport modeling. Recent methods aim to learn these complex relationships from data, but a significant challenge is the scarcity of multi-illumination datasets capturing the same scene under varied lighting conditions. Acquiring such data in the real world requires controlled environments, calibrated equipment, and repeated captures, making it costly and impractical. As a result, most existing approaches rely on synthetic data or focus on narrow domains like portraiture, limiting their ability to generalize to real-world scenes.

To address the lack of multi-illumination data, recent works [69, 39] have decomposed the relighting task into two stages: an inverse rendering step that estimates scene attributes such as albedo, normal, and depth (i.e., G-buffers), followed by a forward rendering step that synthesizes relit images conditioned on these estimates. This pipeline avoids the need for multi-illumination supervision, allowing each stage to be trained separately using single-illumination data. However, inverse and forward rendering are inherently coupled, and modeling them separately with discrete intermediate representations introduces significant limitations. The forward renderer is highly sensitive to errors in

<sup>†</sup> Joint Advising



Figure 1: Given an input image (top left) or video, our method *jointly* estimates albedo (bottom left) and synthesizes relit videos with novel lighting conditions using provided HDR probes. Notably, our estimated albedo maps effectively demodulate shadows and specular highlights, while the relit images exhibit plausible shadows and specular highlights.

the estimated G-buffers and struggles to capture complex material properties not encoded by them, such as transparency and subsurface scattering.

Learning an end-to-end relighting model in low data regime is challenging. While prior methods for intrinsic estimation (e.g., geometry [28, 19] and albedo [32, 39]) trained on synthetic data often generalize well to real-world scenes, we observe that directly training relighting models as video-to-video translation leads to poor generalization to unseen domains. We argue that robust relighting requires an explicit understanding of illumination and scene properties.

Hence, we propose a relighting framework that *jointly* models the distribution of scene intrinsics and illumination. Inspired by VideoJAM [12], we train a video generative model that *jointly* denoises latent space for relighting and albedo demodulation in a single pass. Practically, we concatenate both latent representations and treat them as a single video clip. This design is motivated by the hypothesis that demodulation provides a strong prior for the relighting task, such as removing shadows. This *joint* formulation encourages the model to learn an internal representation of scene structure, leading to improved generalization across diverse and unseen domains. Our approach contrasts with two-stage pipelines that rely on explicit G-buffer estimation. Instead, we implicitly reason about intrinsic scene representation, enabling better representation learning, reducing error accumulation, and modeling of complex visual effects such as specular highlights, transparency, and subsurface scattering.

Our formulation also offers flexibility to learn from diverse sources of supervision. We leverage high-quality synthetic data for full supervision and complement it with real-world single illumination videos that can be auto-labeled at scale. This hybrid training strategy enables the model to handle complex lighting effects while significantly improving generalization to unseen domains, see Figure 1. UNIRELIGHT enables high-quality relighting and intrinsic decomposition from a single input image or video, producing temporally consistent shadows, reflections, and transparency, and outperforms state-of-the-art methods.

# 2 Related Work

**Inverse rendering** estimates intrinsic scene properties such as geometry, materials, and lighting from input images. Traditional approaches use hand-crafted priors within an optimization framework [34, 3, 21, 11, 77, 2] to handle low-order effects like diffuse shading and reflectance. Recently, the field has been revitalized by applying supervised and self-supervised learning [4, 33, 35, 57, 66, 36, 37, 10, 61, 60, 63, 6] to inverse rendering tasks. These methods typically require large, domain-specific datasets, and struggle to generalize outside the training domain.

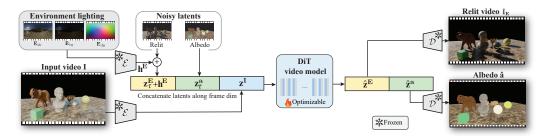


Figure 2: **Method overview.** Given an input video  $\mathbf{I}$  and a target lighting configuration  $(\mathbf{E}_{\text{ldr}}, \mathbf{E}_{\text{log}}, \mathbf{E}_{\text{dir}})$ , our method jointly predicts a relit video  $\hat{\mathbf{I}}_E$  and its corresponding albedo  $\hat{\mathbf{a}}$ . We use a pretrained VAE encoder-decoder pair  $(\mathcal{E}, \mathcal{D})$  to map input and output videos to a latent space. The latents for the target relit video and albedo are concatenated along the temporal (frame) dimension with the encoded input video. Lighting features  $\mathbf{h}^E$ , derived from the environment maps, are concatenated along the channel dimension with the relit video latent. A finetuned DiT video model denoises the joint latent according to Equation 1, enabling consistent generation of both relit appearance and intrinsic decomposition.

Most related to our approach are methods that leverage large generative image and video models for inverse rendering [18, 32, 54, 40, 30, 14, 47, 41]. Notably, RGB $\leftrightarrow$ X [69] demonstrates both intrinsic decomposition and neural rendering from intrinsics using finetuned image diffusion models. The DiffusionRenderer [39] extends RGB $\leftrightarrow$ X to video and also supports relighting.

**Relighting** is the task of modifying the lighting conditions in images or videos. Many methods first reconstruct 3D scenes from multi-view images [46, 22, 13, 9, 71, 76, 70, 56, 62, 38, 64, 42, 58, 25]. Previously, surface reflectance has been captured using calibrated light stages [15]. Now, relighting is supported by material properties recovered from inverse rendering. These methods typically require per-scene optimization and are limited to static, object-centric scenes with single illumination. Training across multiple scenes has led to the exploration of latent feature learning [78, 43, 72, 7], often incorporating neural rendering modules with intrinsic buffers as priors [53, 20, 50, 31, 65].

Recent approaches [55, 31, 68, 26, 39, 5, 75, 45] leverage diffusion models for relighting tasks. These methods are often domain-specific, such as portraits, single objects, and outdoor scenes. Despite promising results, a significant challenge remains the need for multi-illumination datasets [48], which are difficult to capture at scale.

**Joint generative modeling** approaches enable diffusion models to predict multiple modalities. Matrix3D [44] predicts pose estimation, depth, and novel view synthesis using a single DiT [52] model. VideoJAM [12] extends this by predicting both generated pixels and their corresponding motion from a single DiT. We leverage this approach to jointly predict a relit image and albedo.

# 3 Preliminaries: Video Diffusion Models

Diffusion models approximate a data distribution  $p_{\text{data}}(\mathbf{I})$  by learning to iteratively denoise samples corrupted by Gaussian noise [59, 23, 17]. For efficiency, most video diffusion models (VDMs) operate in a lower-dimensional latent space [8, 1]. Given an RGB video  $\mathbf{I} \in \mathbb{R}^{L \times H \times W \times 3}$  consisting of L frames at resolution  $H \times W$ , a pre-trained VAE encoder  $\mathcal E$  encodes the video into a latent tensor  $\mathbf{z} = \mathcal E(\mathbf{I}) \in \mathbb{R}^{l \times h \times w \times C}$ . Then, the final video  $\hat{\mathbf{I}}$  is reconstructed by decoding  $\mathbf{z}$  with a pre-trained VAE decoder  $\mathcal D$ . Both the training and inference stages of the VDM are conducted in this latent space.

In this work, we fine-tune a recent Diffusion Transformer (DiT) video model, Cosmos-Predict1 [1]. Encoding and decoding to and from latent space are performed by the pre-trained VAE Cosmos-1.0-Tokenizer-CV8x8x8, which compresses the video by a factor of eight along the spatial and temporal dimensions:  $l=\frac{L}{8}, C=16, h=\frac{H}{8},$  and  $w=\frac{W}{8}$ . The base model supports text- and image-guided video generation at a resolution of  $704\times1280$  pixels.

To train the VDM, noisy versions  $\mathbf{z}_{\tau} = \alpha_{\tau}\mathbf{z}_0 + \sigma_{\tau}\epsilon$  are constructed by adding Gaussian noise  $\epsilon$ , with the noise schedule provided by  $\alpha_{\tau}$  and  $\sigma_{\tau}$  following the EDM [27]. The diffusion model parameters  $\theta$  of the denoising function  $\mathbf{f}_{\theta}$  are optimized using the denoising score matching objective [27]. Once trained, iteratively applying  $\mathbf{f}_{\theta}$  to a sample of Gaussian noise will produce a sample of  $p_{\text{data}}(\mathbf{I})$ .

# 4 Method

We propose a generative relighting framework by fine-tuning a DiT-based video diffusion model [52, 1], which jointly predicts the albedo and relit appearance from an input image or video under a target lighting condition.

At the core of our method is a joint denoising architecture: the latent representations for albedo and relit video are concatenated at the *token* level and denoised in a single pass using the DiT model. This formulation enables cross-modal interaction via self-attention, allowing the model to capture shared scene structure and improve generalization and temporal consistency.

We leverage a combination of synthetic datasets, multi-illumination data, and auto-labeled real-world data to train our model, ensuring robust generalization and high-quality output. In the following sections, we detail the model architecture, data strategy, and training objectives.

# 4.1 Model Design

Given an input video  $\mathbf{I}$  and a temporally-varying target lighting condition  $\mathbf{E} \in \mathbb{R}^{L \times H \times W \times 3}$  of the same dimensions, our goal is to train a model  $\mathbf{f}_{\theta}$  that *jointly* denoises the albedo  $\mathbf{a}$  of the input video and a relit video  $\mathbf{I}_{\mathbf{E}}$  under the target illumination  $\mathbf{E}$ .

As illustrated in Figure 2, the model comprises a VAE encoder-decoder pair, denoted  $(\mathcal{E}, \mathcal{D})$ , and a transformer-based denoising function,  $\mathbf{f}_{\theta}$ . We use the VAE encoder  $\mathcal{E}$  to separately encode the input video  $\mathbf{I}_{\mathbf{E}}$ , producing the corresponding latent tensors  $(\mathbf{z}^{\mathbf{I}}, \mathbf{z}^{\mathbf{a}}, \mathbf{z}^{\mathbf{E}})$ .

Unlike VideoJAM [12], which compresses a pair of latents  $(\mathbf{z}^x, \mathbf{z}^y)$  using a linear layer, we find this approach does not yield acceptable quality for our applications. We adopt a simple yet effective strategy: concatenating the latents  $\mathbf{z}^I$ ,  $\mathbf{z}^a$ , and  $\mathbf{z}^E$  along the temporal (frame) dimension. This formulation enables the DiT model to apply full self-attention across input, albedo, and relit frames, facilitating cross-modal information exchange.

**Token embeddings.** To distinguish between the three modalities in the concatenated sequence, we combine standard RoPE positional embeddings with dedicated type embeddings. We first apply the same positional embedding to each of the three video clips, encoding the spatial-temporal position of tokens within each clip. To indicate modality, we introduce a learnable *type embedding*  $\mathbf{c}_{\text{emb}} \in \mathbb{R}^{K_{\text{emb}} \times C_{\text{emb}}}$ , where  $K_{\text{emb}} = 3$  denotes the number of video types. Each type embedding is broadcast across spatial-temporal dimensions and concatenated along the channel dimension with its corresponding latent representation:  $\mathbf{c}_{\text{emb}}^0$  is appended to  $\mathbf{z}^\mathbf{E}$ ,  $\mathbf{c}_{\text{emb}}^1$  to  $\mathbf{z}^\mathbf{a}$ , and  $\mathbf{c}_{\text{emb}}^2$  to  $\mathbf{z}^\mathbf{I}$ .

**Condition mask.** We attach a binary mask to each frame indicating whether it should be treated as a condition (e.g., input source video) or a denoising target (e.g., relit video). During training, we randomly vary which modalities are used as inputs, allowing the model to generalize under partial supervision. For instance, we may randomly treat a as either a conditioning input or a prediction target (see Section 4.3 for details).

**Encoding HDR lighting.** Environment maps contain high dynamic range (HDR) values that often exceed the intensity range of standard images, while the VAEs used in latent diffusion models are trained on low dynamic range (LDR) inputs and cannot directly encode these high-intensity signals.

To address this, we follow prior work [26, 39] and represent the lighting condition using three complementary buffers: (1) an LDR panorama  $\mathbf{E}_{ldr}$  obtained via standard Reinhard tonemapping; (2) a normalized log-intensity map  $\mathbf{E}_{log} = \log(\mathbf{E}+1)/E_{max}$ , where  $E_{max}$  is the maximum log intensity; and (3) a directional encoding  $\mathbf{E}_{dir} \in \mathbb{R}^{L \times H \times W \times 3}$ , where each pixel stores a unit vector indicating its direction in the camera coordinate system. These representations are passed through the VAE encoder, yielding  $\mathbf{h^E} = (\mathcal{E}(\mathbf{E}_{ldr}), \mathcal{E}(\mathbf{E}_{log}), \mathcal{E}(\mathbf{E}_{dir}))$  which is concatenated along the channel dimension with  $\mathbf{z^E}$ . We also append a binary condition mask to indicate whether lighting features are present. For input video and albedo tokens, we use zero-padded placeholders to maintain shape consistency.

#### 4.2 Data Strategy

Similar to prior work [39], we combine a large-scale synthetic dataset with smaller real-world datasets. The synthetic dataset offers a large volume of high-quality data, which, when combined with

powerful diffusion models, shows impressive generalization to unseen domains [28, 19]. However, using synthetic target images biases the model output, resulting in a rendered look. Therefore, we also utilize smaller and lower-quality datasets of automatically labeled real-world data.

**Synthetic data curation.** Our synthetic dataset consists of rendered video clips that are tuples  $(\mathbf{I}, \mathbf{I_E}, \mathbf{a}, \mathbf{E})$  of input video, relit result, albedo, and environment map. To create a large dataset with complex and diverse scenes and lighting conditions, we follow the methodology of the Diffusion-Renderer [39] and procedurally generate simple scenes with randomized environment map lighting. We curate a collection of 36,500 3D objects from the Objaverse LVIS subset, 4,260 PBR materials, and 766 HDRI environment maps, gathered from publicly available resources. Each scene contains a ground plane with a randomly selected material, and up to three randomly placed 3D objects. We perform collision detection to avoid intersecting objects. Additionally, we place up to three primitive objects (cube, sphere, and cylinder) with randomized shapes and materials, to increase variety. A randomly selected HDR environment map illuminates each scene. We generate videos with random motions including camera orbits, camera oscillation, lighting rotation, object rotation and translation. Each scene is rendered with two different illuminations under the same motion.

We rendered the dataset using a custom path tracer based on OptiX [51] at high sample counts with path length three to capture global illumination effects, producing a total of 108k videos with 57 frames per video at a resolution of  $704 \times 1280$  pixels.

MIT multi-illumination labeling. The multi-illumination dataset of Murmann et al. [48] consists of 985 real-world indoor scenes for training and 30 scenes for testing, each lit under 25 different conditions. Each image contains a reflective chrome sphere that has been isolated and exported as an HDR image, which can be easily remapped to the latitude-longitude representation expected by our model. We create a dataset by randomly selecting pairs from the 25 lighting conditions and extracting  $(\mathbf{I}, \mathbf{E}, \mathbf{I_E})$ .

To obtain pseudo-groundtruth albedo labels a, we re-implement the inverse renderer from Diffusion-Renderer [39] based on the Cosmos-1.0-Diffusion-7BVideo2World model [49], fine-tuned on our synthetic dataset. The resulting model produces high-quality albedo estimates across diverse images and videos. For each scene, we estimate albedo under all 25 lighting conditions and average the results to obtain a stable albedo map.

**Real-world auto-labeling.** While multi-illumination datasets provide effective supervision for relighting, they are costly to acquire and limited in scale. In contrast, the Internet offers an abundance of high-quality videos captured under single illumination. To leverage this resource, we curate 150k real-world video clips, each with 57 frames at a resolution of  $704\times1280$ , and automatically annotate them with albedo maps to create paired data  $(\mathbf{I}, \mathbf{a})$  using the reproduced inverse rendering model. These auto-labeled RGB–albedo pairs significantly increase the diversity and realism of our training set, enabling improved generalization to real-world scenes.

#### 4.3 Training

Our model is trained on a combination of the synthetic video dataset, the MIT multi-illumination dataset, and real-world auto-labeled data. For image datasets, we treat images as single-frame videos. For both the synthetic video dataset and the MIT multi-illumination dataset, each data sample consists of an input video  $\mathbf{I}_{\mathbf{E}}$  its corresponding albedo  $\mathbf{a}$ , a new environment map  $\mathbf{E}$ , and the target relit video  $\mathbf{I}_{\mathbf{E}}$  under this new illumination. The target latent variable for these datasets is constructed by concatenating the latent of the relit video,  $\mathbf{z}_0^{\mathbf{E}}$ , and the corresponding albedo,  $\mathbf{z}_0^{\mathbf{a}}$ , along the temporal dimension. Noise is introduced independently to both  $\mathbf{z}_0^{\mathbf{E}}$  and  $\mathbf{z}_0^{\mathbf{a}}$  to produce  $\mathbf{z}_{\tau}^{\mathbf{E}}$  and  $\mathbf{z}_{\tau}^{\mathbf{a}}$ . The model parameters are optimized by minimizing the objective function:

$$\hat{\mathbf{z}}^{\mathbf{E}}(\theta), \hat{\mathbf{z}}^{\mathbf{a}}(\theta) = \mathbf{f}_{\theta}([\mathbf{z}_{\tau}^{\mathbf{E}} + \mathbf{h}^{\mathbf{E}}, \mathbf{z}_{\tau}^{\mathbf{a}}, \mathbf{z}^{\mathbf{I}}]; \mathbf{c}_{\text{emb}}, \tau)$$
(1)

$$\mathcal{L}(\theta) = \mathbb{E}_{(\mathbf{z}_0^{\mathbf{E}}, \mathbf{z}_0^{\mathbf{a}}) \sim p_{\text{data}}, \epsilon \sim \mathcal{N}(0, \sigma^2 I)} \left[ \left\| \hat{\mathbf{z}}^{\mathbf{E}}(\theta) - \mathbf{z}_0^{\mathbf{E}} \right\|_2^2 + \lambda_{\mathbf{a}} \left\| \hat{\mathbf{z}}^{\mathbf{a}}(\theta) - \mathbf{z}_0^{\mathbf{a}} \right\|_2^2 \right], \quad (2)$$

where  $[\cdot]$  denotes concatenation in the temporal dimension, and  $\lambda_{\mathbf{a}}=0.1$  is a scalar weight for albedo.

**Training strategy.** Our training process incorporates specific conditioning strategies tailored to the data source. For synthetic data and the MIT multi-illumination dataset, we apply three different



Figure 3: Qualitative comparison on the synthetic dataset and MIT multi-illumination dataset. Our method produces high-quality inter-reflections and shadows in synthetic scenes (top rows). Crucially, on the MIT multi-illumination dataset (bottom rows), it delivers relighting results with higher accuracy than baselines, which fail when faced with complex materials.

conditioning strategies. In 12% of training steps, the input video is dropped (i.e.,  $\mathbf{z^I} = \mathbf{0}$ ), and the ground truth albedo serves as the sole conditioning signal. The denoising function is then  $\mathbf{f}_{\theta}([\mathbf{z}_{\tau}^{\mathbf{E}} + \mathbf{h^E}, \mathbf{z_0^a}, \mathbf{0}]; \mathbf{c}_{\text{emb}}, \tau)$ . In another 18% of the training steps, the model is conditioned on both the input video  $\mathbf{z^I}$  and the ground truth albedo, resulting in the function  $\mathbf{f}_{\theta}([\mathbf{z_{\tau}^E} + \mathbf{h^E}, \mathbf{z_0^a}, \mathbf{z^I}]; \mathbf{c}_{\text{emb}}, \tau)$ . When using these two specific conditions, the conditioning latents of the input video and albedo have their corresponding condition masks set to 1. For the remaining 70% of the training steps, we use the default denoising function where we take input video tokens as conditions and predict the denoised relit video and albedo.

Our real-world auto-labeled data consists of RGB videos and corresponding albedos but lacks pairwise RGB videos and environment maps. The provided albedo  $\mathbf{a}$  is used as a condition, and the RGB video is treated as the target relit video  $\mathbf{I_E}$ . Since the the original input video and the environment map information are unavailable for this dataset, their respective representations are set to  $\mathbf{0}$ : The environment map embedding  $\mathbf{h^E}$  within the primary input array becomes  $\mathbf{0}$ , and the conditional input video representation  $\mathbf{z^I}$  is dropped. The denoising function is thus  $\mathbf{f_{\theta}}([\mathbf{z_T^E} + \mathbf{0}, \mathbf{z_0^a}, \mathbf{0}]; \mathbf{c_{emb}}, \tau)$ . We adopt a 10% probability of dropping conditions to enable classifier-free guidance [24].

# 5 Results

Before evaluating our method, we introduce baselines, metrics, and datasets. Then we conduct ablation studies on the principle of joint modeling and real-world auto-labeling of data. We refer to model implementation details in the Supplement.

**Baselines.** For the relighting task, we compare with the 2D methods DiLightNet [68], Neural-Gaffer [26], and DiffusionRenderer [39]. As the original DiffusionRenderer is built on Stable Video Diffusion [8], to isolate the effect of the base model and enable a fair algorithmic comparison, we re-implement DiffusionRenderer using the same Cosmos [49] backbone as our method, and adopt a data curation strategy closely following the original paper. We refer to this enhanced variant as DiffusionRenderer (Cosmos). For albedo estimation, we compare with IntrinsicImageDiffusion [32] and DiffusionRenderer [39].

**Metrics.** We apply the standard image metrics PSNR, SSIM, and LPIPS [74] to evaluate relighting and albedo estimation results. For relighting, we also conduct a user study to evaluate the perceptual quality of the produced results. Following prior works [32, 26, 39], we counteract scale ambiguity in both relighting and albedo estimation by a three-channel scaling factor using least-squares error minimization before computing the metrics.



Figure 4: **Qualitative comparison on in-the-wild data.** Our method generates more plausible results than the baselines, with higher quality and more realistic appearance.

Table 1: Quantitative evaluation of relighting, including a user study, where "Ours preferred" indicates the preference over the baselines. A preference over the >50% indicates Ours outperforming baselines.

		theticSce SSIM↑		PSNR ↑		ılti-illumiı LPIPS ↓	nation Ours preferred
DiLightNet [68]	19.13	0.584	0.319	15.79	0.539	0.368	$92\% \pm 8\%$
Neural Gaffer [26]	19.17	0.638	0.263	17.87	0.683	0.241	$84\% \pm 2\%$
DiffusionRenderer [39]	24.17	0.768	0.217	16.99	0.582	0.342	$96\% \pm 4\%$
DiffusionRenderer (Cosmos)	26.61	0.841	0.222	17.29	0.622	0.355	$88\% \pm 8\%$
Ours	26.97	$\overline{0.847}$	0.190	20.76	0.749	0.251	n/a

**Datasets.** We have curated a high-quality synthetic dataset named *SyntheticScenes* for quantitative and qualitative evaluation of both relighting and albedo estimation tasks. This dataset includes 3D assets from PolyHaven [67] and Objaverse [16], ensuring that no assets were used in the training of our method or any of the baseline methods. *SyntheticScenes* features 40 scenes, each constructed with a plane textured with random physically-based materials. Each scene is rendered into a 57-frame video sequence under four different lighting conditions, incorporating orbiting camera motions and rotating environment lighting. We cycle through the lighting conditions, selecting one video as the input, and the following as the relighting target, resulting in four different relighting tasks per scene. Additionally, we utilize the test set of *MIT multi-illumination* benchmark [48], which includes 30 scenes under 25 different illuminations. In this dataset, images are captured sequentially under adjacent flashlight illuminations. For our relighting tasks, the image captured under the *i*-th lighting condition serves as input, while the image captured under the (i+12)-th lighting condition is utilized as the target illumination.

#### 5.1 Evaluation

Qualitative comparison. In Figure 3 and Figure 4, we present a comparative analysis of our method against recent state-of-the-art relighting techniques: DiLightNet [68], NeuralGaffer [26], and DiffusionRenderer [39]. Our approach demonstrates superior performance, particularly in handling intricate shadows and inter-reflections. As illustrated in Figure 3, while DiffusionRenderer performs well on the synthetic dataset, it struggles to accurately represent complex materials—such as anisotropic surfaces, glass, and transparent objects—when utilizing G-buffers as an intermediate state, leading to suboptimal results. NeuralGaffer and DiLightNet do not produce accurate shadows and reflections, leading to poor relighting effects. Specifically, NeuralGaffer often makes very subtle, or no changes to the input video.

Figure 4 compares relighting results on in-the-wild data. For the first two object-centric scenes, we provided DiLightNet and NeuralGaffer with masks (we used DiLightNet's auto-masking, which is based on Segment Anything [29]) to produce more reasonable results. Notably, DiffusionRenderer struggles with fur, human skin, and car windows in the first three cases, as these are hard to represent

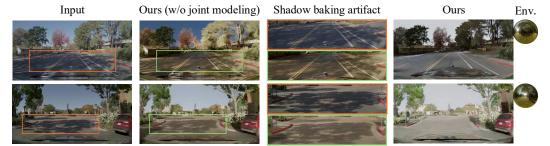


Figure 5: **Ablation on joint modeling.** Relighting results on urban street scenes. The orange and green crops highlight regions where the pure relighting model (w/o joint modeling) clearly bakes shadows from the input image into the relit result. Our joint model correctly demodulates the shadows.

with simple G-buffers. It also misestimates the material of the plastic basin, leading to metal-like appearances. Overall, our relighting model is more general-purpose and can be used more effectively in diverse scenes, yielding comparatively more accurate and high-fidelity relighting results.

**Quantitative comparison.** In Table 1, we present a comparative analysis of our method against all baselines on both *SyntheticScenes* and the MIT multi-illumination benchmark [48]. DiffusionRenderer works well on *SyntheticScenes*, and the Cosmos version—a stronger base model—leads to improved fidelity. However, on the MIT multi-illumination dataset, DiffusionRenderer fails to model the more challenging materials using only the G-buffers, and even the Cosmos version cannot bridge this gap.

An outlier is observed in the LPIPS score for NeuralGaffer on the MIT multi-illumination dataset. As mentioned before, NeuralGaffer makes minimal changes to the input and performs poorly, but still achieves the lowest LPIPS score. We speculate that LPIPS considers the images sufficiently similar due to the relatively small deviation in light direction for this dataset.

**User study.** We conducted a user study to evaluate the perceptual quality of our relighting method compared to the baseline methods on the MIT multi-illumination dataset. Participants were shown a ground-truth relit image alongside two relighting results—one generated by our method and the other by a baseline model, with the order randomly shuffled. They were asked to select the result that more closely resembled the ground truth, considering aspects such as transparency, shadows, and reflections. For each sample pair, 11 users made a binary selection, and majority voting determined the preferred method for each comparison. We repeated the study three times, involving a total of 33 users, and report the average percentage of samples where our method is preferred over baselines in Table 1, along with the standard deviation for the three experiments. Our results are consistently preferred, strongly outperforming the other baselines.

**Albedo estimation.** While our primary focus is on relighting, we also evaluate the quality of our albedo estimates compared to the baselines in Table 2. While the Cosmos version of DiffusionRenderer has a slight edge, our method performs or on par with previous work in all metrics.

Table 2: Quantitative evaluation of albedo estimation on *SyntheticScenes*.

	SyntheticScenes		
	PSNR ↑	SSIM ↑	LPIPS $\downarrow$
IntrinsicImageDiffusion [32]	16.41	0.543	0.395
DiffusionRenderer [39]	26.22	0.837	0.166
DiffusionRenderer (Cosmos)	28.56	0.911	0.131
Ours	28.07	0.877	0.167

# 5.2 Ablation Study

Ablation on joint modeling. We ablate our joint modeling approach with quantitative and qualitative results in Table 3 and Figure 5. We compare our model to a pure relighting model (without joint modeling), whose denoising function is defined as  $\mathbf{f}_{\theta}([\mathbf{z}_{\tau}^{E} + \mathbf{h}^{E}]; \mathbf{z}^{I}, \tau)$ . Both models are trained exclusively on synthetic data. The quantitative results in Table 3 reveal that our joint model slightly improves relighting quality compared to the ablated model. The table also shows that providing ground truth albedo as input to our joint model further enhances quality, highlighting the model's capacity to leverage albedo information effectively. For a qualitative assessment, Figure 5 compares the results on urban street scenes outside the training distribution, where urban street scenes typically exhibit strong shadows under sunlight. The joint modeling approach demonstrates superior generalization, characterized by reduced false shadowing, as clearly shown in the figures, whereas the model without joint modeling largely bakes the shadow from the input.



Figure 6: Ablation on real-world auto-labeled data. Although the dataset is sparsely labeled, it helps the model generalize to natural images.



Figure 7: **Application for data augmentation.** The top left image with green outline is the input image.

Table 3: Ablations on joint modeling designs evaluated on the *SyntheticScenes* dataset.

Ablated Versions	PSNR ↑	SSIM ↑	LPIPS ↓
Ours (w/o joint modeling)	26.42	0.842	0.191
Ours	26.97	0.847	0.190
Ours (w/ GT albedo)	27.15	0.857	0.181

Table 4: User study of relighting on <i>StreetScenes</i> . Each row compares
an ablated variant against the base version (Ours w/o auto-labeled data),
reporting percentage of samples where users prefer the base version.

Ablated Versions	Ours (w/o auto-labeled data) preferred		
Ours (w/o auto-labeled data, w/o joint modeling)	68% ± 14%		
Ours (w/ auto-labeled data) Ours (w/o auto-labeled data, w/ inference-time albedo)	$45\% \pm 8\%$ $54\% \pm 13\%$		

Ablation on real-world automatically labeled data. We ablate the usefulness of our automatically labeled real-world data in Figure 6. The model trained on solely the synthetic and MIT multi-illumination datasets (center column) exhibits noticeable artifacts for out-of-focus regions and outdoor scenes. This is expected, as both datasets have very limited depth of field and primarily feature indoor scenes or lack sky-like backgrounds. As shown in the figure, even though the real-world auto-labeled dataset contains only a subset  $(\mathbf{I}, \mathbf{a})$  of the video labels, it significantly helps the model generalize to out-of-domain data, thereby improving image quality.

**User study.** We conducted a user study to evaluate the perceptual impact of our design choices on relighting quality, using 19 urban street scenes. For each scene, participants were shown a reference image and two relit videos produced by ablated versions of our method (randomized order). They were asked to select the video with more realistic lighting, focusing on aspects such as shadows and reflections. As before, each comparison was evaluated by 11 participants through binary selection, and the preferred result was determined by majority vote. We repeated the study three times with a total of 33 participants and report the average preference rate and the standard deviation in Table 4.

When trained only on multi-illumination data (w/o auto-labeled data), the joint modeling variant is strongly preferred over direct relighting. Adding real-world auto-labeled data further improves perceptual quality, consistent with our qualitative and quantitative findings. When comparing models with and without estimated albedo at inference time, the larger standard deviation in user preference suggests the perceptual quality is comparable. This indicates our method does not strictly rely on estimated albedo at test time.

#### 5.3 Application: Illumination Augmentation

Our model's strong generalization capability enables effective data augmentation for scenarios such as driving and robotics scenes. We test our model under conditions without environment maps and with different random seeds. Figure 7 shows five random illuminations of one input scene. Our model generates diverse data, including nighttime and dusk scenes, demonstrating that it accurately models the illumination distribution and can sample realistic relighting results under varying lighting conditions.

# 6 Discussion

To overcome data scarcity and the limitations of decoupled two-stage methods in realistic relighting, we present a novel framework UNIRELIGHT, that *jointly* models scene intrinsics and illumination. This approach enhances generalization, reduces error accumulation, and more effectively captures complex visual effects by implicitly reasoning about scene properties rather than relying on explicit G-buffers. Furthermore, when trained on real-world auto-labeled data, our model achieves state-of-the-art, high-quality, and realistic relighting results.

Limitations and future work. While our method performs well on diverse data, our model cannot handle emitting objects, such as toggling lights within scenes. Our design, which is focused on environmental lighting, does not extend to emittance effects, which remains an area of future research. Moreover, incorporating text-based relighting could significantly enhance usability. Future work could investigate the conditioning of lighting with natural language by incorporating text cross-attention in our DiT model, a feature currently not implemented. While our method can generate reasonable shadows, integrating our generative approach with explicit 3D information—such as traced shadow buffers—represents a promising direction for improving physical accuracy. Besides, in UNIRELIGHT, we incorporate the albedo to improve the model's generalization, and predicting additional properties could further enhance the model's understanding of scene structure and materials. However, memory requirements scale with the number of properties, making this approach expensive for high-resolution videos; it is worth exploring this direction further.

# Acknowledgments and Disclosure of Funding

The authors thank Tianshi Cao and Huan Ling for their insightful discussions that contributed to this project. This work was conducted at and supported by NVIDIA.

#### References

- [1] Niket Agarwal, Arslan Ali, Maciej Bala, Yogesh Balaji, Erik Barker, Tiffany Cai, Prithvijit Chattopadhyay, Yongxin Chen, Yin Cui, Yifan Ding, et al. Cosmos world foundation model platform for physical AI. arXiv preprint arXiv:2501.03575, 2025. 3, 4
- [2] Jonathan T. Barron and Jitendra Malik. Shape, illumination, and reflectance from shading. <u>TPAMI</u>, 2014.
- [3] Harry Barrow, J. M. Tenenbaum, A. Hanson, and E. Riseman. Recovering intrinsic scene characteristics. Comput. Vis. Syst, 1978. 2
- [4] Sean Bell, Kavita Bala, and Noah Snavely. Intrinsic images in the wild. ACM ToG, 2014. 2
- [5] Shrisha Bharadwaj, Haiwen Feng, Victoria Abrevaya, and Michael J. Black. GenLit: Reformulating Single-Image Relighting as Video Generation. arXiv preprint arXiv:2412.11224, 2024. 3
- [6] Anand Bhattad, Daniel McKee, Derek Hoiem, and D. A. Forsyth. StyleGAN knows normal, depth, albedo, and more. arXiv preprint arXiv:2306.00987, 2023. 2
- [7] Anand Bhattad, James Soole, and D.A. Forsyth. Stylitgan: Image-based relighting via latent control. In CVPR, 2024. 3
- [8] Andreas Blattmann, Tim Dockhorn, Sumith Kulal, Daniel Mendelevitch, Maciej Kilian, Dominik Lorenz, Yam Levi, Zion English, Vikram Voleti, Adam Letts, et al. Stable video diffusion: Scaling latent video diffusion models to large datasets. arXiv preprint arXiv:2311.15127, 2023. 3, 6
- [9] Mark Boss, Raphael Braun, Varun Jampani, Jonathan T. Barron, Ce Liu, and Hendrik P.A. Lensch. NeRD: neural reflectance decomposition from image collections. In ICCV, 2021. 3
- [10] Mark Boss, Varun Jampani, Kihwan Kim, Hendrik P.A. Lensch, and Jan Kautz. Two-shot spatially-varying BRDF and shape estimation. In CVPR, 2020. 2
- [11] Adrien Bousseau, Sylvain Paris, and Frédo Durand. User-assisted intrinsic images. ACM ToG, 2009. 2
- [12] Hila Chefer, Uriel Singer, Amit Zohar, Yuval Kirstain, Adam Polyak, Yaniv Taigman, Lior Wolf, and Shelly Sheynin. VideoJAM: Joint appearance-motion representations for enhanced motion generation in video models. arXiv preprint arXiv: 2502.02492, 2025. 2, 3, 4
- [13] Wenzheng Chen, Joey Litalien, Jun Gao, Zian Wang, Clement Fuji Tsang, Sameh Khalis, Or Litany, and Sanja Fidler. DIB-R++: Learning to predict lighting and material with a hybrid differentiable renderer. In NeurIPS, 2021. 3
- [14] Zhifei Chen, Tianshuo Xu, Wenhang Ge, Leyi Wu, Dongyu Yan, Jing He, Luozhou Wang, Lu Zeng, Shunsi Zhang, and Yingcong Chen. Uni-Renderer: Unifying Rendering and Inverse Rendering Via Dual Stream Diffusion. arXiv preprint arXiv:2412.15050, 2025. 3

- [15] Paul Debevec, Tim Hawkins, Chris Tchou, Haarm-Pieter Duiker, Westley Sarokin, and Mark Sagar. Acquiring the reflectance field of a human face. In SIGGRAPH, 2000. 3
- [16] 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. In CVPR, 2023. 7
- [17] Prafulla Dhariwal and Alexander Quinn Nichol. Diffusion models beat GANs on image synthesis. In NeurIPS, 2021. 3
- [18] Xiaodan Du, Nicholas Kolkin, Greg Shakhnarovich, and Anand Bhattad. Generative models: What do they know? Do they know things? Let's find out! arXiv preprint arXiv:2311.17137, 2024. 3
- [19] Xiao Fu, Wei Yin, Mu Hu, Kaixuan Wang, Yuexin Ma, Ping Tan, Shaojie Shen, Dahua Lin, and Xiaoxiao Long. GeoWizard: unleashing the diffusion priors for 3D geometry estimation from a single image. In ECCV, 2024. 2, 5
- [20] David Griffiths, Tobias Ritschel, and Julien Philip. Outcast: Outdoor single-image relighting with cast shadows. Computer Graphics Forum, 2022. 3
- [21] Roger Grosse, Micah K. Johnson, Edward H. Adelson, and William T. Freeman. Ground truth dataset and baseline evaluations for intrinsic image algorithms. In <u>ICCV</u>, pages 2335–2342. IEEE, 2009. 2
- [22] Jon Hasselgren, Nikolai Hofmann, and Jacob Munkberg. Shape, light, and material decomposition from images using Monte Carlo rendering and denoising. arXiv preprint arXiv:2206.03380, 2022. 3
- [23] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In NeurIPS, 2020. 3
- [24] Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. arXiv preprint arXiv:2207.12598, 2022.
- [25] Yingwenqi Jiang, Jiadong Tu, Yuan Liu, Xifeng Gao, Xiaoxiao Long, Wenping Wang, and Yuexin Ma. GaussianShader: 3D Gaussian splatting with shading functions for reflective surfaces. In CVPR, 2024. 3
- [26] Haian Jin, Yuan Li, Fujun Luan, Yuanbo Xiangli, Sai Bi, Kai Zhang, Zexiang Xu, Jin Sun, and Noah Snavely. Neural Gaffer: Relighting any object via diffusion. In NeurIPS, 2024. 3, 4, 6, 7
- [27] Tero Karras, Miika Aittala, Timo Aila, and Samuli Laine. Elucidating the design space of diffusion-based generative models. In NeurIPS, 2022. 3
- [28] Bingxin Ke, Anton Obukhov, Shengyu Huang, Nando Metzger, Rodrigo Caye Daudt, and Konrad Schindler. Repurposing diffusion-based image generators for monocular depth estimation. In CVPR, 2024. 2, 5
- [29] 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. arXiv preprint arXiv:2304.02643, 2023. 7
- [30] Peter Kocsis, Lukas Höllein, and Matthias Nießner. IntrinsiX: High-Quality PBR Generation using Image Priors. arXiv preprint arXiv:2504.01008, 2025. 3
- [31] Peter Kocsis, Julien Philip, Kalyan Sunkavalli, Matthias Nießner, and Yannick Hold-Geoffroy. LightIt: illumination modeling and control for diffusion models. In CVPR, 2024. 3
- [32] Peter Kocsis, Vincent Sitzmann, and Matthias Nießner. Intrinsic image diffusion for indoor single-view material estimation. In CVPR, 2024. 2, 3, 6, 8
- [33] Balazs Kovacs, Sean Bell, Noah Snavely, and Kavita Bala. Shading annotations in the wild. In <u>CVPR</u>, 2017. 2
- [34] Edwin H. Land and John J. McCann. Lightness and retinex theory. J. Opt. Soc. Am., 1971. 2
- [35] Zhengqi Li and Noah Snavely. CGintrinsics: Better intrinsic image decomposition through physically-based rendering. In ECCV, 2018. 2
- [36] Zhengqin Li, Mohammad Shafiei, Ravi Ramamoorthi, Kalyan Sunkavalli, and Manmohan Chandraker. Inverse rendering for complex indoor scenes: Shape, spatially-varying lighting and svbrdf from a single image. In <u>CVPR</u>, 2020. 2

- [37] Zhengqin Li, Ting-Wei Yu, Shen Sang, Sarah Wang, Sai Bi, Zexiang Xu, Hong-Xing Yu, Kalyan Sunkavalli, Miloš Hašan, Ravi Ramamoorthi, et al. OpenRooms: an end-to-end open framework for photorealistic indoor scene datasets. arXiv preprint arXiv:2007.12868, 2020. 2
- [38] Ruofan Liang, Huiting Chen, Chunlin Li, Fan Chen, Selvakumar Panneer, and Nandita Vijaykumar. ENVIDR: Implicit differentiable renderer with neural environment lighting. In ICCV, 2023. 3
- [39] Ruofan Liang, Zan Gojcic, Huan Ling, Jacob Munkberg, Jon Hasselgren, Zhi-Hao Lin, Jun Gao, Alexander Keller, Nandita Vijaykumar, Sanja Fidler, and Zian Wang. Diffusionrenderer: Neural inverse and forward rendering with video diffusion models. In CVPR, 2025. 1, 2, 3, 4, 5, 6, 7, 8, 22, 23
- [40] Ruofan Liang, Zan Gojcic, Merlin Nimier-David, David Acuna, Nandita Vijaykumar, Sanja Fidler, and Zian Wang. Photorealistic object insertion with diffusion-guided inverse rendering. In ECCV, 2024. 3
- [41] Ruofan Liang, Kai He, Zan Gojcic, Igor Gilitschenski, Sanja Fidler, Nandita Vijaykumar, and Zian Wang. Luxdit: Lighting estimation with video diffusion transformer, 2025. 3
- [42] Zhihao Liang, Qi Zhang, Ying Feng, Ying Shan, and Kui Jia. GS-IR: 3D Gaussian splatting for inverse rendering. arXiv preprint arXiv:2311.16473, 2023. 3
- [43] Andrew Liu, Shiry Ginosar, Tinghui Zhou, Alexei A. Efros, and Noah Snavely. Learning to factorize and relight a city. In ECCV, 2020. 3
- [44] Yuanxun Lu, Jingyang Zhang, Tian Fang, Jean-Daniel Nahmias, Yanghai Tsin, Long Quan, Xun Cao, Yao Yao, and Shiwei Li. Matrix3D: Large Photogrammetry Model All-in-One. <u>arXiv preprint</u> arXiv:2502.07685, 2025. 3
- [45] Yiqun Mei, Mingming He, Li Ma, Julien Philip, Wenqi Xian, David M George, Xueming Yu, Gabriel Dedic, Ahmet Levent Taşel, Ning Yu, et al. Lux post facto: Learning portrait performance relighting with conditional video diffusion and a hybrid dataset. In CVPR, 2025. 3
- [46] Jacob Munkberg, Jon Hasselgren, Tianchang Shen, Jun Gao, Wenzheng Chen, Alex Evans, Thomas Müller, and Sanja Fidler. Extracting triangular 3D models, materials, and lighting from images. <u>arXiv preprint</u> arXiv:2111.12503, 2021. 3
- [47] Jacob Munkberg, Zian Wang, Ruofan Liang, Tianchang Shen, and Jon Hasselgren. VideoMat: Extracting PBR Materials from Video Diffusion Models. In EGSR - CGF Track, 2025. 3
- [48] Lukas Murmann, Michael Gharbi, Miika Aittala, and Fredo Durand. A multi-illumination dataset of indoor object appearance. In <u>ICCV</u>, 2019. 3, 5, 7, 8
- [49] NVIDIA. Cosmos world foundation model platform for physical AI. arXiv preprint arXiv:2501.03575, 2025. 5, 6, 21
- [50] 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 ToG, 2021. 3
- [51] Steven G. Parker, James Bigler, Andreas Dietrich, Heiko Friedrich, Jared Hoberock, David Luebke, David McAllister, Morgan McGuire, Keith Morley, Austin Robison, and Martin Stich. OptiX: a general purpose ray tracing engine. ACM ToG, 2010. 5
- [52] William Peebles and Saining Xie. Scalable diffusion models with transformers. arXiv:2212.09748, 2022. 3, 4
- [53] Julien Philip, Michaël Gharbi, Tinghui Zhou, Alexei A Efros, and George Drettakis. Multi-view relighting using a geometry-aware network. ACM ToG, 2019. 3
- [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. arXiv preprint arxiv:2312.09168, 2023. 3
- [55] Yohan Poirier-Ginter, Alban Gauthier, Julien Philip, Jean-François Lalonde, and George Drettakis. A Diffusion Approach to Radiance Field Relighting using Multi-Illumination Synthesis. <u>Computer Graphics</u> Forum, 2024. 3
- [56] Viktor Rudnev, Mohamed Elgharib, William Smith, Lingjie Liu, Vladislav Golyanik, and Christian Theobalt. NeRF for outdoor scene relighting. In ECCV, 2022. 3

- [57] Soumyadip Sengupta, Jinwei Gu, Kihwan Kim, Guilin Liu, David W. Jacobs, and Jan Kautz. Neural inverse rendering of an indoor scene from a single image. In ICCV, 2019. 2
- [58] Yahao Shi, Yanmin Wu, Chenming Wu, Xing Liu, Chen Zhao, Haocheng Feng, Jingtuo Liu, Liangjun Zhang, Jian Zhang, Bin Zhou, et al. GIR: 3D Gaussian inverse rendering for relightable scene factorization. arXiv preprint arXiv:2312.05133, 2023. 3
- [59] Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In ICML, 2015. 3
- [60] Zian Wang, Wenzheng Chen, David Acuna, Jan Kautz, and Sanja Fidler. Neural light field estimation for street scenes with differentiable virtual object insertion. In <u>ECCV</u>, 2022. 2
- [61] Zian Wang, Jonah Philion, Sanja Fidler, and Jan Kautz. Learning indoor inverse rendering with 3D spatially-varying lighting. In ICCV, 2021. 2
- [62] Zian Wang, Tianchang Shen, Jun Gao, Shengyu Huang, Jacob Munkberg, Jon Hasselgren, Zan Gojcic, Wenzheng Chen, and Sanja Fidler. Neural fields meet explicit geometric representations for inverse rendering of urban scenes. In CVPR, June 2023. 3
- [63] Felix Wimbauer, Shangzhe Wu, and Christian Rupprecht. De-rendering 3D objects in the wild. In <u>CVPR</u>, 2022. 2
- [64] Chen Xi, Peng Sida, Yang Dongchen, Liu Yuan, Pan Bowen, Lv Chengfei, and Zhou. Xiaowei. IntrinsicAnything: learning diffusion priors for inverse rendering under unknown illumination. <u>arXiv preprint</u> arXiv: 2404.11593, 2024. 3
- [65] Xiaoyan Xing, Konrad Groh, Sezer Karaoglu, Theo Gevers, and Anand Bhattad. Luminet: Latent intrinsics meets diffusion models for indoor scene relighting. arXiv preprint arXiv:2412.00177, 2024. 3
- [66] Ye Yu and William A. P. Smith. InverseRenderNet: learning single image inverse rendering. In <u>CVPR</u>, 2019. 2
- [67] Greg Zaal and et al. Poly Haven The Public 3D Asset Library, 2024. 7
- [68] Chong Zeng, Yue Dong, Pieter Peers, Youkang Kong, Hongzhi Wu, and Xin Tong. DiLightNet: fine-grained lighting control for diffusion-based image generation. In <u>SIGGRAPH</u>, 2024. 3, 6, 7, 21
- [69] Zheng Zeng, Valentin Deschaintre, Iliyan Georgiev, Yannick Hold-Geoffroy, Yiwei Hu, Fujun Luan, Ling-Qi Yan, and Miloš Hašan. RGB↔X: image decomposition and synthesis using material-and lighting-aware diffusion models. In SIGGRAPH, 2024. 1, 3
- [70] Kai Zhang, Fujun Luan, Zhengqi Li, and Noah Snavely. IRON: inverse rendering by optimizing neural SDFs and materials from photometric images. In CVPR, 2022. 3
- [71] Kai Zhang, Fujun Luan, Qianqian Wang, Kavita Bala, and Noah Snavely. PhySG: Inverse rendering with spherical Gaussians for physics-based material editing and relighting. In CVPR, 2021. 3
- [72] Longwen Zhang, Qixuan Zhang, Minye Wu, Jingyi Yu, and Lan Xu. Neural video portrait relighting in real-time via consistency modeling. In <u>ICCV</u>, 2021. 3
- [73] Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. Scaling in-the-wild training for diffusion-based illumination harmonization and editing by imposing consistent light transport. In ICLR, 2025. 22, 23
- [74] Richard Zhang, Phillip Isola, Alexei A. Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In CVPR, 2018. 6
- [75] Tianyuan Zhang, Zhengfei Kuang, Haian Jin, Zexiang Xu, Sai Bi, Hao Tan, He Zhang, Yiwei Hu, Milos Hasan, William T. Freeman, Kai Zhang, and Fujun Luan. RelitLRM: Generative relightable radiance for large reconstruction models. In ICLR, 2025. 3
- [76] 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. <u>ACM</u> ToG, 2021. 3
- [77] Qi Zhao, Ping Tan, Qiang Dai, Li Shen, Enhua Wu, and Stephen Lin. A closed-form solution to retinex with nonlocal texture constraints. TPAMI, 34(7):1437–1444, 2012.
- [78] Hao Zhou, Sunil Hadap, Kalyan Sunkavalli, and David W. Jacobs. Deep single-image portrait relighting. In ICCV, October 2019. 3

# **NeurIPS Paper Checklist**

The checklist is designed to encourage best practices for responsible machine learning research, addressing issues of reproducibility, transparency, research ethics, and societal impact. Do not remove the checklist: **The papers not including the checklist will be desk rejected.** The checklist should follow the references and follow the (optional) supplemental material. The checklist does NOT count towards the page limit.

Please read the checklist guidelines carefully for information on how to answer these questions. For each question in the checklist:

- You should answer [Yes], [No], or [NA].
- [NA] means either that the question is Not Applicable for that particular paper or the relevant information is Not Available.
- Please provide a short (1–2 sentence) justification right after your answer (even for NA).

The checklist answers are an integral part of your paper submission. They are visible to the reviewers, area chairs, senior area chairs, and ethics reviewers. You will be asked to also include it (after eventual revisions) with the final version of your paper, and its final version will be published with the paper.

The reviewers of your paper will be asked to use the checklist as one of the factors in their evaluation. While "[Yes]" is generally preferable to "[No]", it is perfectly acceptable to answer "[No]" provided a proper justification is given (e.g., "error bars are not reported because it would be too computationally expensive" or "we were unable to find the license for the dataset we used"). In general, answering "[No]" or "[NA]" is not grounds for rejection. While the questions are phrased in a binary way, we acknowledge that the true answer is often more nuanced, so please just use your best judgment and write a justification to elaborate. All supporting evidence can appear either in the main paper or the supplemental material, provided in appendix. If you answer [Yes] to a question, in the justification please point to the section(s) where related material for the question can be found.

# IMPORTANT, please:

- Delete this instruction block, but keep the section heading "NeurIPS Paper Checklist",
- · Keep the checklist subsection headings, questions/answers and guidelines below.
- Do not modify the questions and only use the provided macros for your answers.

#### 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 abstract and introduction clearly state our contributions, which are fully supported by the experiments and analysis presented.

#### 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: We discuss limitations in Section 6.

#### 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: The paper does not contain theoretical results or formal proofs.

#### 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 crossreferenced.
- 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 include all necessary details for reproduction in the main paper and supplementary materials.

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: [No]

Justification: We plan to release the code and data upon acceptance. The internal guidelines of our institution prevent us from releasing code at this stage.

# 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, hyperparameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: All relevant details are provided in the paper and supplementary materials.

#### 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: Due to the high computational cost of training large-scale video diffusion models, we did not perform multiple runs to report error bars.

#### 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: We include sufficient information in supplementary materials.

# 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: Our work adheres to the NeurIPS Code of Ethics.

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: [Yes]

Justification: We discussed societal impacts of our work in the 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: Our work does not involve models or data with high risk of misuse.

#### 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: Third-party assets are properly credited, and their licenses are respected.

#### 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 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: Our paper does not release new assets.

# 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: [Yes]

Justification: We include the details of all crowdsourcing experiments in supplementary materials.

#### 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: [Yes]

Justification: An approval is obtained for all of our experiments.

#### 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.

# 16. Declaration of LLM usage

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does not impact the core methodology, scientific rigorousness, or originality of the research, declaration is not required.

Answer: [NA]

Justification: LLMs were not used in the core methods of this research.

#### Guidelines:

- The answer NA means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
- Please refer to our LLM policy (https://neurips.cc/Conferences/2025/LLM) for what should or should not be described.

In this Appendix, we first discuss the broader impact of our project (Appendix A). We then provide additional implementation details of our model and experiments (Appendix B), followed by further results and analysis (Appendix C). Please refer to the accompanying video for additional qualitative results and comparisons.

# A Broader Impact

We present UniRelight, a generative framework that jointly estimates albedo and synthesizes relit videos from a single input, enabling diverse lighting manipulation across both synthetic and real-world scenes. This capability can support a range of applications, including creative content generation, visual effects, virtual production, and potentially data augmentation for training more robust computer vision models in domains such as robotics and autonomous driving.

As with all generative video models, UniRelight may reflect biases present in its training data. Such biases could lead to relighting results that fail to generalize to underrepresented scene types or lighting conditions. Furthermore, tools for lighting manipulation carry the risk of misuse, such as altering or misrepresenting visual content in sensitive contexts like surveillance or media.

We discourage the use of UniRelight in applications where relighting may contribute to misinformation, misattribution, or privacy violations. In human-centric use cases, we recommend careful dataset curation to ensure fair representation across skin tones, races, and gender identities. Practitioners are encouraged to critically assess and de-bias training data to mitigate unintended harms where appropriate.

# **B** Experimental Details

#### **B.1** Implementation Details

We fine-tune our models based on Cosmos-Predict1-7B-Video2World [49], a pre-trained DiT video diffusion model.

The encoded latents  $\mathbf{z^I}$ ,  $\mathbf{z^a}$ ,  $\mathbf{z^E}$ ,  $\mathcal{E}(\mathbf{E}_{ldr})$ ,  $\mathcal{E}(\mathbf{E}_{log})$ , and  $\mathcal{E}(\mathbf{E}_{dir})$  are all in  $\mathbb{R}^{l \times h \times w \times C}$ , where C=16. We use  $C_{emb}=3$  as the dimension of the type embedding. Thus, the concatenated tokens have a channel dimension of  $16+1+(16+1)\times 3+3=71$ , where each latent has an associated binary condition mask (added as an extra channel), and the lighting features ( $\mathbf{E}_{ldr}$ ,  $\mathbf{E}_{log}$ ,  $\mathbf{E}_{dir}$ ) each include a condition mask to indicate whether they are provided.

We adopt an image-video co-training strategy and train the model in two stages. Firstly, we train with only synthetic data by mixing the image data (sampling one frame from the video data) and the video data in a ratio of 1:1 for 15,000 iterations. Then we train with all data with random sampling, including synthetic video data, synthetic image data, real-world auto-labeled data, and MIT multi-illumination data, in a ratio of 8:1:3:2 for 12,000 iterations. We augment the real-world auto-labeled data with random flipping.

All training is done with a batch size of 64, using the AdamW optimizer with a learning rate of  $2\times 10^{-5}$ , with mixed-precision (BF16) training at a resolution of  $480\times 848$  pixels. The AdamW optimizer was employed with a weight decay of 0.1. The exponential decay rates for the moment estimates  $\beta$  are set to 0.9 for the first moment and 0.99 for the second moment, with  $\epsilon$  at  $1\times 10^{-10}$ . The total training of two stages takes around 4 days on 32 A100 GPUs.

During inference, we use 35 denoising steps. We do not apply classifier-free guidance (CFG), as we empirically found that inference without CFG yields more accurate and visually consistent results.

**Baseline configurations.** Since DiLightNet [68] requires a text prompt per example, we use meta/llama-3.2-90b-vision-instruct<sup>0</sup> to generate a short prompt for each example in the datasets based on the first image in each clip with the instruction "What is in this image? Describe the materials. Be concise and produce an answer with a few sentences, no more than 50 words."

As each of the baselines generates videos in different resolutions, for UNet-based baselines, we run inference on the model with the video first resized to  $486 \times 864$  and then center-cropped to a

<sup>0</sup>https://www.llama.com/

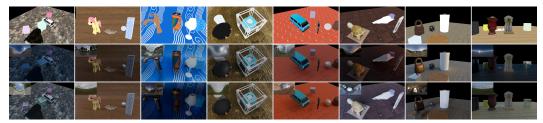


Figure 8: **Synthetic data visualization.** Randomly sampled example images are shown from our synthetic rendering data. The top-most images show albedo maps, while the bottom two rows display rendered scenes under two different illuminations with the corresponding environment maps.

Table 5: Quantitative comparison with IC-Light.

	PSNR↑	SSIM↑	$LPIPS \downarrow$
IC-Light [73] Ours	18.08 <b>23.19</b>	0.834 <b>0.901</b>	0.096 <b>0.079</b>

Table 6: Evaluation of inference runtime cost.

Runtime cost (seco		
DiffusionRenderer [39]	566.6	
DiffusionRenderer (Cosmos)	780.0	
Ours	<b>445.5</b>	

resolution of  $448 \times 832$ ; for our DiT-based model, we run inference on the model with the resized video with resolution of  $486 \times 864$  and then center-cropped to  $448 \times 832$  to align the results.

**Quantitative comparison configurations.** For quantitative evaluation, we apply background masks to the synthetic dataset to focus on foreground appearance. For the MIT multi-illumination dataset, we follow the dataset protocol and mask out light probes in all outputs before computing metrics.

#### **B.2** User Study Details

We conducted two user studies on Amazon Mechanical Turk to evaluate the perceptual quality of relighting results.

**MIT multi-illumination dataset** is a public benchmark with ground truth relighting. Participants were shown three images: a ground truth relit image and two relighting results—one generated by our method and one by a baseline model. Their task was to choose the result that more closely resembled the ground truth, considering attributes such as transparency, shadows, and reflections.

The exact instructions shown to participants were as follows:

Carefully compare Image A, the Reference Image, and Image B. Your task is to determine which image (A or B) is more similar to the Reference Image.

To make an informed decision, you may zoom in to examine the details. Pay close attention to aspects such as lighting, reflections, and shadows, as these can affect how natural the image appears.

Once you have compared the images, select the one that best matches the Reference Image.

- ☐ Image A
- ☐ Image B

We evaluated 30 scenes from the test set, comparing our method against four baselines. The study was repeated three times with 11 unique participants in each run. In total, this resulted in  $30 \times 4 \times 11 \times 3 = 3960$  individual comparisons.

**StreetScenes Dataset.** This dataset contains 19 urban street scenes without ground-truth relighting. Participants were shown a reference image along with two relit videos generated by different ablated versions of our method.

The instructions presented to participants were as follows:

In this study, you will be shown a Reference Image and two videos – Video A and Video B – that changes the lighting of the scene. Your task is to watch both



Figure 9: **Qualitative comparison with IC-Light [73].** We provide the environmental background used for IC-Light conditioning, with the reference environment ball on the left. Our method produces higher-quality and more accurate relighting results.

videos and choose which one (A or B) you think has more realistic shadows and reflections. To make an informed decision, you may zoom in to examine the details. Pay close attention to aspects such as lighting, reflections, and shadows, as these can affect how natural the image appears.

Once you have compared the videos, select the one that has more realistic lighting effects.

- ☐ Video A
- □ Video B

We evaluated 19 scenes, comparing a base version against two ablated versions of our method. The study was repeated three times with 11 unique participants in each run. In total, this resulted in  $19 \times 2 \times 11 \times 3 = 1254$  individual comparisons.

#### **B.3** Synthetic data visualization

We show a synthetic data visualization in Figure 8. Each scene contains albedo videos, two environment maps, and pairwise videos rendered under the environment maps.

# C Additional Results

**Runtime cost.** We evaluate the inference runtime of our model on a 57-frame video at a resolution of  $480 \times 848$ . The overall inference time for performing 35 denoising steps, including VAE encoding and decoding, is 445.5 seconds, measured on a single A100 GPU.

To contextualize this cost, we compare our method against two baselines: DiffusionRenderer [39] and its Cosmos-based variant. All methods use 35 denoising steps for consistency. DiffusionRenderer is run at  $448 \times 832$  resolution (slightly smaller than ours but divisible by 32 to fit its architecture) while the Cosmos variant is run at our native resolution.

For baseline methods, the total runtime is the sum of the inverse rendering and forward rendering durations. Notably, DiffusionRenderer requires five separate inverse rendering passes and one forward rendering pass per video, resulting in significantly higher computational cost. In contrast, our approach performs joint relighting and albedo estimation in a single pass and is correspondingly faster. Full timing results are shown in Table 6.

**Comparison with IC-Light [73].** We compare our method with the single-image relighting approach IC-Light [73] on object-centric synthetic data, as shown in Table 5 and Figure 9. Note that the

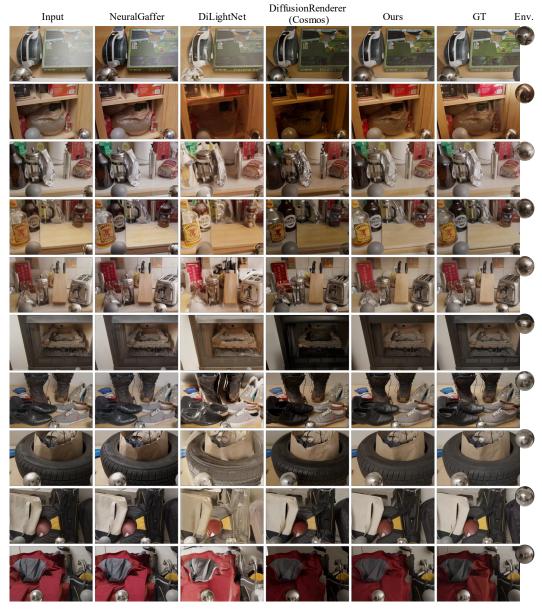


Figure 10: **Additional qualitative comparison on MIT multi-illumination dataset.** Our method consistently achieves more accurate relighting results than all baselines on the MIT multi-illumination dataset, demonstrating strong capability in relighting complex materials.

two methods follow different relighting formulations: IC-Light is designed for object relighting using background context as the primary cue—without access to an explicit environment map, while our method is conditioned on full HDR illumination, but is not specifically tuned for object-centric data.

Our method shows improved quantitative and qualitative performance. Since IC-Light relies on background appearance as its primary cue and has less information about the surrounding lighting, it may retain input-specific effects in its outputs—such as specular highlights and shadows, which can limit accuracy under novel lighting conditions. In contrast, our method produces more faithful relighting results, with sharper specular highlights, more realistic shadows, and improved visual fidelity.

Additional qualitative comparison on the MIT multi-illumination dataset. We provide additional qualitative comparisons on the MIT Multi-Illumination dataset in Figure 10. To ensure a fair comparison, we include results from our re-implemented version of DiffusionRenderer using the Cosmos backbone, which achieves higher visual fidelity than the original implementation. Our method



Figure 11: **Additional qualitative results under point-light illumination.** The bottom right of each column indicates the target lighting conditions. Our results show strong robustness of our method under point-light illumination.

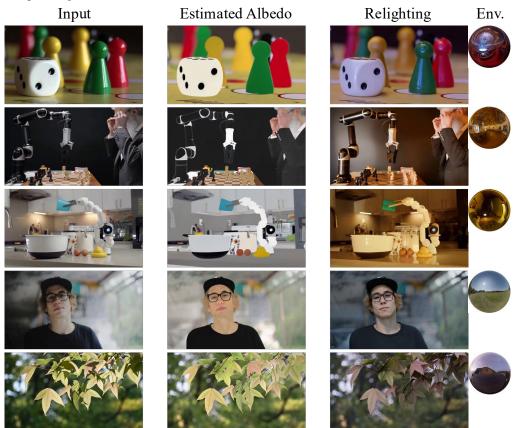


Figure 12: **Additional qualitative results on real scenes.** Our method provides high-quality albedo estimation and realistic relighting results.

consistently produces more accurate transparency, specular highlights, and shadows across scenes, demonstrating strong capability in handling complex materials and outperforming all baselines in visual quality.

**Additional qualitative results under point-light illumination.** We further evaluate the robustness of our method in extreme cases, such as point-light illuminations, which do not exist in our training data. As shown in Figure 11, our method produces high-quality relighting results, demonstrating the strong robustness and generalization capability of our method.

**Additional qualitative results on real scenes.** We present additional results on real scenes in Figure 12. Our method produces high-quality albedo and relighting results with realistic specular highlights and shadows under target lighting conditions.