

# GR3EN: Generative Relighting for 3D Environments

Xiaoyan Xing<sup>1,3</sup> Philipp Henzler<sup>2</sup> Junhwa Hur<sup>1</sup> Runze Li<sup>1</sup>  
Jonathan T. Barron<sup>1</sup> Pratul P. Srinivasan<sup>1</sup> Dor Verbin<sup>1</sup>

<sup>1</sup>Google DeepMind <sup>2</sup>Google Research <sup>3</sup>University of Amsterdam

<https://GR3EN-relight.github.io>




Figure 1. Our method provides fine-grained control over the lighting of a 3D scene. Given (a) multi-view input images, we first (b) reconstruct a 3D representation. Given target lighting conditions (here turning off the left overhead light and turning on the lamp on the dresser), our method outputs a (c) relit 3D scene, which can be used to render (d) novel views of the scene under the target lighting.

## Abstract

*We present a method for relighting 3D reconstructions of large room-scale environments. Existing solutions for 3D scene relighting often require solving under-determined or ill-conditioned inverse rendering problems, and are as such unable to produce high-quality results on complex real-world scenes. Though recent progress in using generative image and video diffusion models for relighting has been promising, these techniques are either limited to 2D image and video relighting or 3D relighting of individual objects. Our approach enables controllable 3D relighting of room-scale scenes by distilling the outputs of a video-to-video relighting diffusion model into a 3D reconstruction. This side-steps the need to solve a difficult inverse rendering problem, and results in a flexible system that can relight 3D reconstructions of complex real-world scenes. We validate our approach on both synthetic and real-world datasets to show that it can faithfully render novel views of scenes under new lighting conditions.*

## 1. Introduction

Relighting 3D scene representations is a long-standing challenge that is crucial to enable the use of captured real-world scenes in computer graphics applications. Traditionally, this problem has been addressed by first using inverse rendering techniques to decompose a scene’s appearance into materials, geometry, and illumination, and then re-rendering that scene under modified illumination conditions.

Since inverse rendering is inherently ill-posed, solutions often rely on strong priors or simplifying assumptions, such as Lambertian reflectance, global directional lighting, or convex geometry. Consequently, these simplifying assumptions compromise photorealism and restrict generalization to complex environments. Recent methods leverage large generative image and video models for relighting, effectively bypassing the complexities of inverse rendering. However, existing approaches are limited to 2D relighting of images, or 3D relighting of isolated 3D objects illuminated by an environment map (see a summary of these methods in Table 1) — there exist no effective solutions for relighting 3D scene reconstructions, a task which is partic-

Table 1. **Summary of prior relighting methods.** We introduce the first generative relighting framework for room-scale 3D environments. Our method enables novel view synthesis (NVS) of the relit scene, is conditioned only on the target lighting, and does not require access to any additional G-buffer data such as depth, surface normals, or material properties.

| Method             | Conditioning                  | Relit Modality    | NVS |
|--------------------|-------------------------------|-------------------|-----|
| LightLab [29]      | depth map, light source image | image             | ×   |
| Light-A-Video [57] | text                          | video             | ×   |
| DRenderer [26]     | G-buffer, environment map     | video             | ×   |
| UniRelight [16]    | albedo, environment map       | video             | ×   |
| ROGR [43]          | environment map               | 3D object         | ✓   |
| Ours               | light source video            | 3D representation | ✓   |

ularly difficult due to the complex lighting occlusion and interreflection effects that are present in even the simplest room-scale captures.

We introduce a relighting approach that leverages generative video models to relight 3D reconstructions of real-world scenes. First, we present a method for fine-tuning a pre-trained video diffusion model to perform consistent video-to-video relighting given a user-specified target lighting condition and source video. We then use this video-to-video model to relight video sequences rendered from a reconstructed 3D scene representation. Finally, we introduce a technique for distilling these relit 2D videos into a high-fidelity 3D scene representation that looks like the input scene, but relit under new lighting conditions.

Our system is flexible and enables users to perform intuitive 3D scene lighting edits such as turning existing lights on and off, changing their color, or removing outdoor lighting entering through a window. We quantitatively validate our system on synthetic scenes to demonstrate that it can render high-fidelity novel views with different illumination conditions, and we qualitatively demonstrate that it enables rendering photorealistic novel views of real-world 3D scene reconstructions under novel user-controlled lighting.

## 2. Related Work

**Inverse rendering.** Early solutions for scene relighting were built upon inverse rendering systems, which decompose a scene’s appearance into its constituent factors of materials, illumination, and geometry (or some subset of these, assuming the others are known) [6, 38, 39, 50]. These methods either directly invert a simple image formation model [19], or iteratively optimize these factors [4, 9, 10, 24, 25, 50, 52], typically using simplified models for reflectance and lighting. Early attempts at inverse rendering parameterize scene geometry using meshes or depth maps, while the rise of radiance fields [22, 33] enabled effective approaches for inverse rendering that parameterize a scene using the weights of a neural network [3, 21, 30, 40, 55, 58] or a set of 3D Gaussians [12, 14]. While this has substantially improved inverse rendering results by enabling

gradient-based optimization through a differentiable image formation process, this line of work is still limited by the central challenge of inverse rendering: that it is highly ill-posed. As such, these methods rely on simplifying assumptions of the scene or brittle priors over scene content, which severely limits photorealism, or rely on an abundance of input data (*e.g.*, images of both novel views *and* novel lighting conditions), which limits the settings in which these models can be used.

**2D image relighting.** Many recent approaches have attempted to circumvent the challenging inverse rendering decomposition problem by directly learning to relight images in a fully data-driven manner. These approaches have trained feed-forward neural networks [15, 32, 35, 41, 54], or generative image models such as Generative Adversarial Networks (GANs) [7, 34, 42, 45] and more recently diffusion models [11, 17, 20, 23, 29, 44, 47, 48, 51, 53] to generate plausible relit images, conditioned on a desired target lighting. While these methods have been effective at producing plausible 2D results, they do not produce 3D outputs that can be used for view synthesis or other graphics applications. A concurrent work, LuxRemix [27], proposes a relighting framework based on image diffusion models that first decomposes a scene into its constituent light sources, followed by a trained multi-view harmonization model to propagate lighting across different views for scene relighting. In contrast, our method directly achieves consistent room-scale relighting without requiring explicit light source decomposition or view-wise harmonization.

**Video relighting.** Building on recent advancements in video diffusion models, a number of works have addressed the challenge of 2D video relighting by using diffusion models [8, 13, 16, 26, 31, 57]. However, these approaches provide limited control over lighting, only allowing editing coarse control such as modifying directional lights or using text prompts to describe the target light.

**Generative 3D relighting.** Relighting objects and scenes is a core task in many 3D content creation workflows, and several recent works have focused on relighting reconstructed 3D representations. Neural Gaffer [20] and IllumiNeRF [56] trained 2D diffusion models to relight a set of images, conditioned on a target environment lighting, and used these relit images to reconstruct NeRFs of individual objects. Extensions of this approach to multiview diffusion models improves the consistency of the results by jointly relighting multiple images of an object captured under varying illumination conditions [1, 43]. Other approaches use generative models to first estimate material properties and then relight the objects [28]. While these techniques produce impressive relighting results for isolated objects, they all rely on environment lighting representations that cannot be applied to larger scenes.






Figure 2. **Our pipeline for 3D scene relighting.** Starting with a 3D reconstruction (top right), we first render a video of the scene along a chosen camera path (top left). We then use our relighting diffusion model to relight the input video given a target lighting, which is provided as a video whose values specify the color and intensity of the target light sources. Finally, the output relit video is distilled back into the 3D representation (bottom right), resulting in a relit 3D scene that can be rendered from novel viewpoints.

### 3. Method

Given a set of posed multiview images and user-provided target lighting, our method produces a 3D representation of the relit scene. Our pipeline addresses this task in three stages. First, we reconstruct a 3D scene representation from the input multiview images. Second, we render a video sequence from this 3D representation and pass it along with the target lighting through our video-to-video relighting diffusion model, to generate a relit video from the same camera path. Finally, we “distill” the relit appearance from the relit video back into the 3D representation, to get a final 3D representation that enables rendering the relit scene from any camera pose. This pipeline (not including the initial 3D reconstruction) is illustrated in Figure 2.

In Sections 3.1 and 3.2 we describe the relighting model and its training process, and in Section 3.3 we describe the 3D reconstruction and distillation pipeline.

#### 3.1. Relighting Model

The main component of our method is a video-to-video relighting diffusion model. We start with the Wan 2.2 TI2V-

5B base model [46], and we fine-tune it for video relighting. Our model takes an input video and a specification of the target lighting. The input video is rendered from our 3D reconstruction, and it provides the majority of the content of the output video, such as geometry, materials, and camera path. The goal of the relighting model is to maintain the content of the input video while only changing the lighting.

Our lighting conditioning must model all light sources in indoor environments, which consists of directly observed light sources (*e.g.*, ceiling lights or standing lamps), and from unobserved sources such as the sun shining through a window.

For observed indoor light sources, we provide the relighting model with a per-pixel annotation of the target lighting, containing its intensity and its hue. For unobserved light sources outside the scene, we assume that external illumination is dominated by the sun, so we provide a single scalar for the sun’s intensity.

**Visible lighting conditioning.** Given an RGB video  $\mathbf{I} \in \mathbb{R}^{T \times H \times W \times 3}$  consisting of  $T$  frames and spatial dimensions  $H \times W$ , we define the visible lighting conditioning as a

video  $\mathbf{M} \in \mathbb{R}^{T \times H \times W \times 3}$ .

The frames of  $\mathbf{M}$  correspond to the same cameras and content as  $\mathbf{I}$ , but their values control the light sources. A pixel with RGB value  $\mathbf{c} \in \mathbb{R}^3$  indicates that the light source observed at that pixel should have intensity  $\mathbf{c}$ , whereas a pixel with value  $(-1, -1, -1)$  indicates that the light source should not be edited. Having a special value of  $-1$  for not editing is important because otherwise we would have to set  $\mathbf{c}$  to exactly the right value at every pixel. We provide examples of our lighting conditioning videos  $\mathbf{M}$  in Figure 3 and the supplement.

Prior work on relighting using generative models such as LightLab [29] concatenates the input image and the lighting condition along the image’s channel dimension. This uses the fact that the content in the input and output images should be aligned. However, in our experiments this approach did not yield high quality video relighting results. Other approaches like UniRelight [16] instead concatenate the output and conditioning videos along the temporal dimension, along with positional embedding that is identical for all videos (*i.e.*, the target lighting videos, input video, and output video all use the same positional embedding). Instead, we found another method to work better than these existing approaches: we also concatenate the videos along the temporal dimension, but we use rotary positional encodings [18]. This design allows full self-attention and provides different positional encodings to different videos. Empirically we show that this approach results in superior performance compared to other conditioning strategies.

**Directional lighting conditioning.** Our representation of unobserved light is extremely simple, consisting of a single scalar for the sun’s intensity. Similar to LightLab [29], the intensity gets lifted to a high-dimensional space using Fourier features, passed through a shallow MLP to get the directional lighting embedding  $\mathbf{d}$  and fed into the model’s cross-attention layers.

### 3.2. Relighting Model Training

Each batch for training our diffusion model consists of an input video  $\mathbf{I}$ , the lighting condition video  $\mathbf{M}$ , and the directional lighting embedding  $\mathbf{d}$ . A pretrained, frozen autoencoder [46] is used to map the video sequences to spatial-temporal latents, one for the input video  $\mathbf{z}_I$ , and one for the lighting condition video  $\mathbf{z}_M$ . Those two latents are concatenated along the temporal dimension with the output video latent (also compressed by the same shared autoencoder) which is being denoised  $\mathbf{z}$ . A flow-based noise scheduler  $f_{\text{noise}}(\mathbf{z}, t)$  introduces noise to  $\mathbf{z}$  at a given timestep  $t$ . The parameters  $\theta$  of the relighting model  $f_\theta$  are optimized by minimizing the objective function at time step  $t$ :

$$\mathcal{L}(\theta) = \|\hat{\mathbf{z}}_t - f_{\text{noise}}(\mathbf{z}, t)\|_2^2, \quad (1)$$

where the noise added to the latent is computed as:

$$\hat{\mathbf{z}}_t = f_\theta(f_{\text{noise}}(\mathbf{z}_t, t), \mathbf{z}_I, \mathbf{z}_M, \mathbf{d}, \mathbf{e}; t). \quad (2)$$

After denoising, we use the pretrained autoencoder’s decoder to obtain the relit output video.

During training, we sample the timestep  $t$  with a strong bias towards higher noise magnitudes, by drawing it from the following distribution:

$$p(t) = \begin{cases} \frac{\rho}{\tau} & \text{if } 0 \leq t \leq \tau, \\ \frac{1-\rho}{1-\tau} & \text{if } \tau < t \leq 1, \end{cases} \quad (3)$$

where we set  $\tau = 0.4$  and  $\rho = 0.85$ , meaning that 85% of training steps are done on 40% of the noisiest timesteps.

We find that this bias promotes much faster convergence and improved relighting fidelity. This is because plausible relighting requires global reasoning to ensure illumination consistency across different parts of the scene, and correctness with respect to the target lighting. This low-frequency structure is primarily established during high noise levels. Low noise levels are responsible for refining high-frequency details, which is a relatively trivial task in our setting, since the scene’s content and texture are already provided in the input video.

### 3.3. Reconstruction and 3D Distillation

The first step in our pipeline is obtaining a 3D reconstruction of a scene. We can reconstruct a 3D scene from a collection of multi-view images using Neural Radiance Fields [33] or 3D Gaussian Splatting [22]. However, because our method only relies on videos rendered from the 3D model (and not the multi-view images used to reconstruct it), our method can also be applied to existing 3D reconstructions.

Once a 3D reconstruction is available, we render it into a video to be passed to our relighting video-to-video model. We render an elliptical camera path using a fisheye camera with a wide field of view. This is because the video diffusion model can only process a limited number of frames at once, so increasing the field of view of the camera improves the coverage of the scene.

After passing this video through our relighting model, we distill the relit video by fine-tuning the 3D model to match the relit video. The relit 3D representation then enables rendering images and videos from arbitrary new viewpoints.

Throughout the main paper, our results are obtained by using Zip-NeRF for reconstruction and 3D distillation due to its high quality [5]. However, nothing about our approach is specific to any particular radiance field technique, and one could instead use a real-time (but slightly lower quality) approach such as 3D Gaussian Splatting [22].



Figure 3. **A visualization of our “one-light-at-a-time” data synthesis pipeline.** We (a) separately render the scene illuminated by each of the light sources, (b) multiply each one by its target light color, and (c) combine all light sources according to Equation 4 to get the target illumination.

## 4. Training Data

Training a generative model for relighting requires a large-scale dataset of various scenes that are rendered under diverse, controllable lighting conditions. To this end, we use Infinigen [36, 37] to construct a large-scale synthetic dataset of photorealistic room-scale environments. Our dataset comprises over 300 indoor scenes spanning a variety of room types, including dining rooms, kitchens, and bedrooms. Each scene is rendered under 30 distinct lighting conditions, resulting in more than 9,000 unique scene-lighting combinations. In total, the dataset consists of 108,000 unique videos of 81 frames each. We render all images using Blender’s Cycles engine at a resolution of  $720 \times 1280$  pixels. We include samples from our dataset in the supplemental material.

Given a training scene, for each training camera  $\pi$  we render a set of  $L + 1$  “one-light-at-a-time” (OLAT) images  $I^\pi = \{I_\ell^\pi\}_{\ell=0}^L$ , where  $L$  denotes the number of artificial light sources, and the additional light corresponds to the unobserved outside light sources. The  $\ell$ th image is rendered by turning off all lights, and only turning on the  $\ell$ th light source. We render and store these images in high dynamic range linear color space, enabling the synthesis of arbitrary target lighting conditions during training via linear combinations of the basis images:

$$I_{\text{target}}^\pi = \gamma \left( e \cdot \sum_{\ell=0}^L \mathbf{c}_\ell I_\ell^\pi \right). \quad (4)$$

Here,  $\mathbf{c}_\ell \in [0, 1]^3$  represents the intensity of the  $\ell$ th light source,  $e$  is a selected exposure value, and  $\gamma$  is a simple sRGB tone mapping function [2]. The OLAT data synthesis pipeline is illustrated in Figure 3.

During training, the intensities  $\{\mathbf{c}_\ell\}_{\ell=0}^L$  are selected randomly. Every light source is independently set to be on or off ( $\mathbf{c}_\ell=0$ ) with equal probabilities, with at least one light being on (*i.e.*, we reject samples where all lights are off).

For lights that are on,  $\mathbf{c}_\ell$  is chosen to reflect typical indoor scene illuminants: for 80% of iterations the color values are set to the color of blackbody radiation with temperature uniformly chosen between 2500K and 10500K, and for the remaining 20% of iterations it is assigned a uniformly random color in HSV space. This sampling strategy encourages the model to learn a robust representation of lighting without overfitting to specific lighting patterns.

**Camera settings.** Our training camera paths are designed to capture large portions of the scene. We use equisolid fisheye cameras with a  $180^\circ$  field of view. Each elliptical rendering path consists of 81 evenly-spaced cameras.

## 5. Experiments

We evaluate our method on two data sources to thoroughly assess its performance, robustness, and generalization capabilities. Our first set of experiments shows qualitative results on real-world indoor scenes from the Eyeful Tower dataset [49]. We show that our method is successful at handling the domain gap between synthetic and real data: even though our relighting model is trained solely on synthetic data, it is able to relight real scenes. Additionally, while our training data is diverse, it contains a limited number of lighting fixtures. However, we demonstrate that our approach is able to generalize to unseen light shapes and render plausible relit videos which can be distilled into 3D representations. We conduct a large-scale user study to evaluate the performance of our method against LightLab [29], and the improved Cosmos version of DiffusionRenderer [26].

Our second dataset consists of held-out synthetic scenes rendered via Infinigen, that follows the same distribution as our training data. We use this dataset for qualitatively and quantitatively testing our method, and we perform an ablation study on this data to validate our design of the diffusion model and its training.

### 5.1. Implementation Details

**Relighting model.** Our video diffusion model, based on Wan 2.2 [46], is finetuned from the standard weights for 11,200 iterations on our relighting dataset. We use a batch size of 64, processing video clips of 81 frames at a resolution of  $352 \times 640$ . Training takes  $\sim 36$  hours on 64 NVIDIA H100 GPUs.

**3D distillation.** To distill the relit video into the final relit Zip-NeRF model, we optimize it using the standard training pipeline for 2000 iterations using the Adam optimizer with a learning rate of  $10^{-3}$ . This takes  $\sim 10$  minutes on 8 NVIDIA H100 GPUs.

### 5.2. Baselines

To the best of our knowledge, our method is the first to relight full 3D scenes with fine-grained, per-source control






Figure 4. **Real-world relighting results on the Eyeful Tower dataset.** The top row shows a frame from the input video (and an inset on the right), with the target lighting condition displayed in the bottom-right corner of the input frame. The target light requires turning off the right half of the overhead lights as well as the light coming from outside the room. The remaining rows show the relit results obtained by Nano Banana, LightLab, and our method. Note that our method produces realistic edits with convincing global illumination effects like the shadow cast by the object on the table. Refer appendix for results of other views.

over intensity and color. Prior 2D image editors such as LightLab [29] provide similar controls but operate on one image at a time and thus do not maintain multi-view or temporal consistency. Conversely, scene editors like DiffusionRenderer ensure temporal consistency but offer limited control and cannot target specific light sources. Text-driven video relighting such as Light-A-Video [57] can modify lighting using text prompts, but only for certain video types and without explicit per-light control. Because no existing approach offers controllable 3D scene relighting, we introduce a new benchmark for this task and compare against the closest available baselines from these adjacent settings.

**Conditioning for LightLab.** LightLab only supports editing a single light source. Therefore in order to modify multiple light sources we apply their model multiple times, each



Figure 5. **3D light insertion.** We add a new out-of-distribution light source to a real reconstructed 3D scene. Despite the training data not containing this type of light, and despite the relighting model being trained only using synthetic data, our method still obtains realistic results. Note in particular the consistent global illumination effects observed for example in the shadows of the objects on the table.




Figure 6. **Comparison of our relit novel views with those obtained by LightLab.** We show novel view renderings of the original scene (left), compared with the same image relit by LightLab (center), and the corresponding viewpoint rendered from our 3D relit scene (right). Note that LightLab fails to relight the scene whether the target light is visible (top row), as well as when it is not directly visible (bottom row).

only adjusting the color or intensity of one light source.

**Conditioning for DiffusionRenderer.** DiffusionRenderer relights scenes based on an environment map. Therefore for each target illumination we use our relit NeRF to render a target environment map. Since DiffusionRenderer supports at most 57 frames, we evaluate it on the first 57 frames of our camera path and report its results on this subset.

**Conditioning for Light-A-Video.** Light-A-Video is conditioned on text. For each sequence, we use prompts that first describe the unobserved illumination, and then specify the color of each observed light and whether it is on or off. We employ its WAN 2.1-based variation, as it handles longer videos and higher resolution.



Table 2. **User Study on Real-World Scene Video Relighting.** We conduct a human preference study evaluating three critical aspects: frame-to-frame consistency (F2F-C), object identity preservation (O-Pres), and lighting similarity to the provided lighting description (L-Sim). We report the win rates (%) of our method against each competing approach.

| Opponent Method        | F2F-C. $\uparrow$ | O-Pres. $\uparrow$ | L-Sim. $\uparrow$ |
|------------------------|-------------------|--------------------|-------------------|
| Input                  | 61.1              | 55.6               | 88.9              |
| LightLab [29]          | 86.1              | 80.6               | 86.1              |
| DiffusionRenderer [26] | 80.6              | 83.3               | 88.9              |
| Ours Overall           | 75.9              | 73.1               | 88.0              |

### 5.3. Real-World Scene Relighting Results

We evaluate our method on captured scenes from the Eyeful Tower dataset [49]. This dataset consists of high-quality captures of indoor environments. We show qualitative results of our entire pipeline (reconstruction, relighting, and 3D distillation) in Figures 1, 2, 4, and 5. Perhaps surprisingly, our model was never trained on light fixture shapes present in this data, such as the large rectangular ceiling light sources in the top half of Figure 4, but it successfully generalizes to out-of-distribution lighting conditions.

Figure 6 shows our relit 3D model renders under different lighting conditions, demonstrating the ability of our full pipeline to render the scene under different illumination conditions, compared with LightLab [29], a single-image relighting method, which fails relighting this image, and therefore cannot be used for 3D relighting.

Figures 8 and 11 show the results of our relighting model, demonstrating that our diffusion model provides fine-grained control over the illumination of the scene, rendering realistic global illumination effects like cast shadows and reflections.

Finally, our project page contains many more video results and comparisons with other methods.

**User study.** We conducted a user study comparing our method with LightLab [29], DiffusionRenderer [26], and the original input video (which has high quality but does not respect the target lighting). Users were shown pairs of videos with identical camera paths and a text description of the target lighting condition, then asked to evaluate frame-to-frame consistency, object identity preservation, and lighting similarity.

Table 2 presents the results. Our method achieves an overall win rate of 79.0%, with strong performance across all criteria: 88.0% for lighting similarity, 75.9% for temporal consistency, and 73.1% for identity preservation. We outperform both LightLab and DiffusionRenderer with 84.3% win rates, and achieve 68.5% against the artifact-free input, demonstrating that users prefer our improved lighting fidelity. Details are provided in the supplementary material.

Table 3. **Quantitative evaluation of per-frame relighting.** Our method outperforms DiffusionRenderer, which is a method capable of relighting 3D scenes, but does not allow as much control as our model. The “input video” entry denotes the metrics computed between the input video and the relit video, for reference.

| Method                 | PSNR $\uparrow$ | SSIM $\uparrow$ | LPIPS $\downarrow$ |
|------------------------|-----------------|-----------------|--------------------|
| Input video            | 13.11           | 0.679           | 0.252              |
| DiffusionRenderer [26] | 15.17           | 0.553           | 0.668              |
| LightLab [29]          | 11.95           | 0.699           | 0.281              |
| Light-A-Video [57]     | 10.88           | 0.525           | 0.663              |
| <b>Ours</b>            | <b>24.77</b>    | <b>0.741</b>    | <b>0.130</b>       |

### 5.4. Synthetic Scene Relighting Results

In Table 3 we report our relighting results compared with DiffusionRenderer (Cosmos version) [26], LightLab [29] and Light-A-Video (WAN 2.1 version) [57]. Although DiffusionRenderer can theoretically relight 3D scenes, its control over light sources is extremely limited. As shown in Figure 12, DiffusionRenderer fails to accurately modify visible light sources as well as exterior light sources that are not directly visible. While LightLab can edit the color of the lights, its control is limited, and it fails to turn on exterior light sources.

### 5.5. Ablation Studies

We perform an ablation study to validate two of our major design choices in our relighting model’s architecture and training process. The results are reported in Table 4. First, we replace our conditioning approach described in Section 3.1 with concatenation along the channel dimension, as done in LightLab [29], which results in a significant degradation in relighting quality. Second, we replace the sampling strategy described in Section 3.2 with standard uniform samples  $t \sim U([0, 1])$ , which also leads to a degradation in the accuracy of the relit results.

### 5.6. Additional Applications

**3D light insertion.** Our method can also be used to insert new light sources with arbitrary shapes into a 3D scene. Figure 5 shows an example where we insert into the scene a new “neon” light. This is done by projecting the target light shape onto the planar wall, and using it as the conditioning lights. Despite being trained on a limited set of lighting conditions, which does not include this type of light, our method produces plausible-looking results.

**Video relighting.** Although our full method is designed for relighting full 3D scenes, the video relighting component can be used on its own, enabling video relighting without requiring 3D reconstruction. We show results of our video relighting model in Figure 7. Our relighting model is capable of removing exterior lights and their associated



Figure 7. **Real-world video relighting.** Our model is capable of relighting real-world videos (top row). In the first example (middle row), the lamp is turned on and its color set to yellow; note the cast shadow of the ball and the reflectance on the surface of the ball. In the second example (bottom row), our model generalizes and successfully transforms the ball into a light source, thereby producing plausible global illumination effects.




Figure 8. **Control of light intensity.** Our method enables fine control of the target intensity. The leftmost three columns display the relit scenes under the same target lights but different target light intensities. The regions primarily lit by the modified target light are marked in blue boxes and those dominated by light that is not modified are in orange. The middle three columns and the rightmost three columns are insets corresponding to the blue and orange boxed regions, respectively. Note the dramatic changes in local appearance within the target-lit regions, contrasted with the nearly constant illumination in regions primarily lit by unmodified light sources.

Table 4. **Ablation study of model design.** We observe that replacing our conditioning approach with simple channel-wise concatenation, or removing the high-noise training bias, significantly degrades relighting performance.

|                       | PSNR $\uparrow$ | SSIM $\uparrow$ | LPIPS $\downarrow$ |
|-----------------------|-----------------|-----------------|--------------------|
| Channel concatenation | 17.06           | 0.648           | 0.172              |
| w.o. high noise bias  | 23.82           | 0.637           | 0.192              |
| <b>Ours</b>           | <b>24.77</b>    | <b>0.741</b>    | <b>0.130</b>       |

lighting effects such as reflections and cast shadows.

## 6. Conclusion

We have presented a system to relight 3D reconstructions of entire scenes. Our technique renders a video from the 3D scene representation, relights it using a video-to-video diffusion model conditioned on user-specified target lighting,

and distills this relit video back into the 3D reconstruction to generate a relit 3D scene representation.

Our experiments have demonstrated that our approach produces photorealistic relit 3D representations of real-world scenes, faithfully representing global illumination effects such as cast shadows and specular highlights, which prior methods often struggle with. This opens up opportunities for editing and using 3D scene reconstructions as controllable graphics assets.

**Limitations.** Although our relighting model generalizes well across diverse light sources, its performance is limited by the diversity of light source assets used during training. For example, in Figure 5, the light source is not entirely emissive. This could be mitigated by incorporating a broader range of light source assets during training.

## References

- [1] Hadi Alzayer, Philipp Henzler, Jonathan T. Barron, Jia-Bin Huang, Pratul P. Srinivasan, and Dor Verbin. 2025. Generative Multiview Relighting for 3D Reconstruction under Extreme Illumination Variation. *CVPR* (2025). 2
- [2] Matthew Anderson, Ricardo Motta, Srinivasan Chandrasekar, and Michael Stokes. 1996. Proposal for a standard default color space for the internet—sRGB. *CIC* (1996). 5
- [3] Benjamin Attal, Dor Verbin, Ben Mildenhall, Peter Hedman, Jonathan T. Barron, Matthew O’Toole, and Pratul P. Srinivasan. 2024. Flash Cache: Reducing Bias in Radiance Cache Based Inverse Rendering. *ECCV* (2024). 2
- [4] Jonathan T. Barron and Jitendra Malik. 2014. Shape, illumination, and reflectance from shading. *IEEE TPAMI* (2014). 2
- [5] Jonathan T. Barron, Ben Mildenhall, Dor Verbin, Pratul P. Srinivasan, and Peter Hedman. 2023. Zip-NeRF: Anti-Aliased Grid-Based Neural Radiance Fields. *ICCV* (2023). 4, 12
- [6] Harry Barrow, J Tenenbaum, A Hanson, and E Riseman. 1978. Recovering intrinsic scene characteristics. *Comput. Vis. Syst* 2 (1978), 3–26. 2
- [7] Anand Bhattad, James Soole, and David A. Forsyth. 2024. StylitGAN: Image-Based Relighting via Latent Control. *CVPR* (2024). 2
- [8] Ziqi Cai, Kaiwen Jiang, Shu-Yu Chen, Yu-Kun Lai, Hongbo Fu, Boxin Shi, and Lin Gao. 2024. Real-time 3D-aware Portrait Video Relighting. *CVPR* (2024). 2
- [9] Chris Careaga and Yağız Aksoy. 2025. Physically Controllable Relighting of Photographs. *SIGGRAPH* (2025). 2
- [10] Tim Caselitz, Michael Krawez, Jugesh Sundram, Mark Van Loock, and Wolfram Burgard. 2020. Camera Tracking in Lighting Adaptable Maps of Indoor Environments. *ICRA* (2020). 2
- [11] Jun Myeong Choi, Annie Wang, Pieter Peers, Anand Bhattad, and Roni Sengupta. 2025. ScribbleLight: Single Image Indoor Relighting with Scribbles. *CVPR* (2025). 2
- [12] Jiahui Fan, Fujun Luan, Jian Yang, Milos Hasan, and Beibei Wang. 2025. RNG: Relightable Neural Gaussians. *CVPR* (2025). 2
- [13] Ye Fang, Zeyi Sun, Shangzhan Zhang, Tong Wu, Yinghao Xu, Pan Zhang, Jiaqi Wang, Gordon Wetzstein, and Dahua Lin. 2025. RelightVid: Temporal-Consistent Diffusion Model for Video Relighting. *arXiv:2501.16330* (2025). 2
- [14] Jian Gao, Chun Gu, Youtian Lin, Hao Zhu, Xun Cao, Li Zhang, and Yao Yao. 2024. Relightable 3D gaussian: Real-time point cloud relighting with BRDF decomposition and ray tracing. *ECCV* (2024). 2
- [15] David Griffiths, Tobias Ritschel, and Julien Philip. 2022. OutCast: Single Image Relighting with Cast Shadows. *CGF* (2022). 2
- [16] Kai He, Ruofan Liang, Jacob Munkberg, Jon Hasselgren, Nandita Vijaykumar, Alexander Keller, Sanja Fidler, Igor Gilitschenski, Zan Gojcic, and Zian Wang. 2025. UniRelight: Learning Joint Decomposition and Synthesis for Video Relighting. *NeurIPS* (2025). 2, 4
- [17] Mingming He, Pascal Clausen, Ahmet Levent Taşel, Li Ma, Oliver Pilarski, Wenqi Xian, Laszlo Rikker, Xueming Yu, Ryan Burgert, Ning Yu, et al. 2024. DiffRelight: Diffusion-Based Facial Performance Relighting. *SIGGRAPH Asia* (2024). 2
- [18] Byeongho Heo, Song Park, Dongyoon Han, and Sangdoo Yun. 2024. Rotary position embedding for vision transformer. *ECCV* (2024). 4
- [19] Berthold KP Horn. 1974. Determining Lightness from an Image. *Computer Graphics and Image Processing* (1974). 2
- [20] Haian Jin, Yuan Li, Fujun Luan, Yuanbo Xiangli, Sai Bi, Kai Zhang, Zexiang Xu, Jin Sun, and Noah Snively. 2024. Neural Gaffer: Relighting Any Object via Diffusion. *NeurIPS* (2024). 2
- [21] Haian Jin, Isabella Liu, Peijia Xu, Xiaoshuai Zhang, Songfang Han, Sai Bi, Xiaowei Zhou, Zexiang Xu, and Hao Su. 2023. TensorIR: Tensorial Inverse Rendering. *CVPR* (2023). 2
- [22] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 2023. 3D Gaussian splatting for real-time radiance field rendering. *ACM TOG* (2023). 2, 4
- [23] Peter Kocsis, Julien Philip, Kalyan Sunkavalli, Matthias Nießner, and Yannick Hold-Geoffroy. 2024. LightIt: Illumination modeling and control for diffusion models. *CVPR* (2024). 2
- [24] Zhengqin Li, Jia Shi, Sai Bi, Rui Zhu, Kalyan Sunkavalli, Miloš Hašan, Zexiang Xu, Ravi Ramamoorthi, and Manmohan Chandraker. 2022. Physically-Based Editing of Indoor Scene Lighting from a Single Image. *ECCV* (2022). 2
- [25] Zhen Li, Lingli Wang, Mofang Cheng, Cihui Pan, and Jiaqi Yang. 2023. Multi-view Inverse Rendering for Large-scale Real-world Indoor Scenes. *CVPR* (2023). 2
- [26] Ruofan Liang, Zan Gojcic, Huan Ling, Jacob Munkberg, Jon Hasselgren, Zhi-Hao Lin, Jun Gao, Alexander Keller, Nandita Vijaykumar, Sanja Fidler, and Zian Wang. 2025. DiffusionRenderer: Neural Inverse and Forward Rendering with Video Diffusion Models. *CVPR* (2025). 2, 5, 7, 12
- [27] Ruofan Liang, Norman Müller, Ethan Weber, Duncan Zauss, Nandita Vijaykumar, Peter Kontschieder, and Christian Richardt. 2026. LuxRemix: Lighting Decomposition and Remixing for Indoor Scenes. *arXiv preprint arXiv:2601.15283* (2026). 2
- [28] Yehonathan Litman, Fernando De la Torre, and Shubham Tulsiani. 2025. LightSwitch: Multi-view Relighting with Material-guided Diffusion. *ICCV* (2025). 2
- [29] Nadav Magar, Amir Hertz, Eric Tabellion, Yael Pritch, Alex Rav-Acha, Ariel Shamir, and Yedid Hoshen. 2025. LightLab: Controlling Light Sources in Images with Diffusion Models. *SIGGRAPH* (2025). 2, 4, 5, 6, 7, 12
- [30] Alexander Mai, Dor Verbin, Falko Kuester, and Sara Fridovich-Keil. 2023. Neural microfacet fields for inverse rendering. *ICCV* (2023). 2
- [31] Yiqun Mei, Mingming He, Li Ma, Julien Philip, Wenqi Xian, David M George, Xueming Yu, Gabriel Dedic, Ahmet Levent Taşel, Ning Yu, et al. 2025. Lux post facto: Learning portrait performance relighting with conditional video diffusion and a hybrid dataset. *CVPR* (2025). 2



- [32] Yiqun Mei, He Zhang, Xuaner Zhang, Jianming Zhang, Zhixin Shu, Yilin Wang, Zijun Wei, Shi Yan, HyunJoon Jung, and Vishal M. Patel. 2023. LightPainter: Interactive Portrait Relighting with Freehand Scribble. *CVPR* (2023). 2
- [33] Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, Jonathan T. Barron, Ravi Ramamoorthi, and Ren Ng. 2020. NeRF: Representing Scenes as Neural Radiance Fields for View Synthesis. *ECCV* (2020). 2, 4
- [34] Xingang Pan, Xudong Xu, Chen Change Loy, Christian Theobalt, and Bo Dai. 2021. A Shading-Guided Generative Implicit Model for Shape-Accurate 3D-Aware Image Synthesis. *NeurIPS* (2021). 2
- [35] Rohit Pandey, Sergio Orts-Escolano, Chloe Legendre, Christian Haene, Sofien Bouaziz, Christoph Rhemann, Paul E. Debevec, and Sean Ryan Fanello. 2021. Total Relighting: Learning to Relight Portraits for Background Replacement. *SIGGRAPH* (2021). 2
- [36] Alexander Raistrick, Lahav Lipson, Zeyu Ma, Lingjie Mei, Mingzhe Wang, Yiming Zuo, Karhan Kayan, Hongyu Wen, Beining Han, Yihan Wang, et al. 2023. Infinite Photorealistic Worlds using Procedural Generation. *CVPR* (2023). 5, 12
- [37] Alexander Raistrick, Lingjie Mei, Karhan Kayan, David Yan, Yiming Zuo, Beining Han, Hongyu Wen, Meenal Parakh, Stamatis Alexandropoulos, Lahav Lipson, et al. 2024. Infinigen Indoors: Photorealistic Indoor Scenes using Procedural Generation. *CVPR* (2024). 5, 12
- [38] Ravi Ramamoorthi and Pat Hanrahan. 2001. A signal-processing framework for inverse rendering. *SIGGRAPH* (2001). 2
- [39] Yoichi Sato, Mark D. Wheeler, and Katsushi Ikeuchi. 1997. Object shape and reflectance modeling from observation. *SIGGRAPH* (1997). 2
- [40] Pratul P. Srinivasan, Boyang Deng, Xiuming Zhang, Matthew Tancik, Ben Mildenhall, and Jonathan T. Barron. 2021. NeRV: Neural Reflectance and Visibility Fields for Relighting and View Synthesis. *CVPR* (2021). 2
- [41] Tiancheng Sun, Jonathan T Barron, Yun-Ta Tsai, Zexiang Xu, Xueming Yu, Graham Fyffe, Christoph Rhemann, Jay Busch, Paul Debevec, and Ravi Ramamoorthi. 2019. Single image portrait relighting. *ACM TOG* (2019). 2
- [42] Feitong Tan, Sean Fanello, Abhimitra Meka, Sergio Orts-Escolano, Danhang Tang, Rohit Pandey, Jonathan Taylor, Ping Tan, and Yinda Zhang. 2022. VoLux-GAN: A Generative Model for 3D Face Synthesis with HDRI Relighting. *SIGGRAPH* (2022). 2
- [43] Jiapeng Tang, Matthew Levine, Dor Verbin, Stephan J. Garbin, Matthias Niessner, Ricardo Martin-Brualla, Pratul P. Srinivasan, and Philipp Henzler. 2025. ROGR: Relightable 3D Objects using Generative Relighting. *NeurIPS* (2025). 2
- [44] LuminaBrush Team. 2024. LuminaBrush: Illumination Drawing Tools for Text-to-Image Diffusion Models. 2
- [45] Ayush Tewari, Mohamed Elgharib, Gaurav Bharaj, Florian Bernard, Hans-Peter Seidel, Patrick Pérez, Michael Zollhofer, and Christian Theobalt. 2020. StyleRig: Rigging StyleGAN for 3D Control over Portrait Images. *CVPR* (2020). 2
- [46] Ang Wang, Baole Ai, Bin Wen, Chaojie Mao, Chen-Wei Xie, Di Chen, Feiwu Yu, Haiming Zhao, Jianxiao Yang, Jianyuan Zeng, Jiayu Wang, Jingfeng Zhang, Jingren Zhou, Jinkai Wang, Jixuan Chen, Kai Zhu, Kang Zhao, Keyu Yan, Lianghua Huang, Mengyang Feng, Ningyi Zhang, Pandeng Li, Pingyu Wu, Ruihang Chu, Ruili Feng, Shiwei Zhang, Siyang Sun, Tao Fang, Tianxing Wang, Tianyi Gui, Tingyu Weng, Tong Shen, Wei Lin, Wei Wang, Wei Wang, Wenmeng Zhou, Wenten Wang, Wenting Shen, Wenyuan Yu, Xi-anzhong Shi, Xiaoming Huang, Xin Xu, Yan Kou, Yangyu Lv, Yifei Li, Yijing Liu, Yiming Wang, Yingya Zhang, Yitong Huang, Yong Li, You Wu, Yu Liu, Yulin Pan, Yun Zheng, Yuntao Hong, Yupeng Shi, Yutong Feng, Zeyinzi Jiang, Zhen Han, Zhi-Fan Wu, and Ziyu Liu. 2025. Wan: Open and Advanced Large-Scale Video Generative Models. *arXiv:2503.20314* (2025). 3, 4, 5
- [47] Xiaoyan Xing, Konrad Groh, Sezer Karaoglu, Theo Gevers, and Anand Bhattad. 2025. LumiNet: Latent Intrinsic Meets Diffusion Models for Indoor Scene Relighting. *CVPR* (2025). 2
- [48] Xiaoyan Xing, Vincent Tao Hu, Jan Hendrik Metzen, Konrad Groh, Sezer Karaoglu, and Theo Gevers. 2024. Retinex-diffusion: On controlling illumination conditions in diffusion models via retinex theory. *arXiv preprint arXiv:2407.20785* (2024). 2
- [49] Lining Xu, Vasu Agrawal, William Laney, Tony Garcia, Aayush Bansal, Changil Kim, Samuel Rota Bulò, Lorenzo Porzi, Peter Kontschieder, Aljaž Božič, et al. 2023. VR-NeRF: High-fidelity virtualized walkable spaces. *SIGGRAPH Asia* (2023). 5, 7, 12
- [50] Yizhou Yu, Paul Debevec, Jitendra Malik, and Tim Hawkins. 1999. Inverse global illumination: recovering reflectance models of real scenes from photographs. *SIGGRAPH* (1999). 2
- [51] Chong Zeng, Yue Dong, Pieter Peers, Youkang Kong, Hongzhi Wu, and Xin Tong. 2024. DiLightNet: Fine-grained lighting control for diffusion-based image generation. *SIGGRAPH* (2024). 2
- [52] Edward Zhang, Michael F. Cohen, and Brian Curless. 2016. Emptying, Refurnishing, and Relighting Indoor Spaces. *ACM TOG* (2016). 2
- [53] Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. 2025. Scaling In-the-Wild Training for Diffusion-based Illumination Harmonization and Editing by Imposing Consistent Light Transport. *ICLR* (2025). 2
- [54] Xiao Zhang, William Gao, Seemantdar Jain, Michael Maire, David Forsyth, and Anand Bhattad. 2024. Latent Intrinsic Emerge from Training to Relight. *NeurIPS* (2024). 2
- [55] Xiuming Zhang, Pratul P Srinivasan, Boyang Deng, Paul Debevec, William T. Freeman, and Jonathan T. Barron. 2021. Nerfactor: Neural factorization of shape and reflectance under an unknown illumination. *ACM TOG* (2021). 2
- [56] Xiaoming Zhao, Pratul P. Srinivasan, Dor Verbin, Keunhong Park, Ricardo Martin Brualla, and Philipp Henzler. 2024. IllumiNeRF: 3D Relighting Without Inverse Rendering. *NeurIPS* (2024). 2
- [57] Yujie Zhou, Jiazi Bu, Pengyang Ling, Pan Zhang, Tong Wu, Qidong Huang, Jinsong Li, Xiaoyi Dong, Yuhang Zang, Yuhang Cao, Anyi Rao, Jiaqi Wang, and Li Niu. 2025. Light-

A-Video: Training-free Video Relighting via Progressive Light Fusion. *ICCV* (2025). [2](#), [6](#), [7](#)

- [58] Jingsen Zhu, Yuchi Huo, Qi Ye, Fujun Luan, Jifan Li, Dianbing Xi, Lisha Wang, Rui Tang, Wei Hua, Hujun Bao, and Rui Wang. 2023. I<sup>2</sup>-SDF: Intrinsic Indoor Scene Reconstruction and Editing via Raytracing in Neural SDFs. *CVPR* (2023). [2](#)

# Appendix

## A. Implementation Details

### A.1. Relighting Diffusion Model

**Training.** As described in the main paper, we fine-tune the WAN 2.2 T12V-5B model using the AdamW optimizer with a learning rate of  $10^{-6}$ . The additional MLP modules introduced for our method are trained jointly with the diffusion model, using a learning rate of  $10^{-5}$ . We use Fully Sharded Data Parallel (FSDP) to distribute the model across multiple hosts for efficient large-scale training.

**Inference.** At inference time, we generate an initial random noise tensor matching the spatial and temporal dimensions of the input video latent, and apply 50 denoising steps using the same flow-matching scheduler employed during training. The full inference process requires approximately  $\sim 100$  seconds for generating an 81-frame video at a resolution of  $640 \times 352$  on 8 H100 GPUs.

### A.2. Reconstruction and 3D Distillation

As mentioned in the main paper, our 3D reconstruction is done using Zip-NeRF [5]. Once the Zip-NeRF model is trained, we render an elliptical camera path and feed it through our relighting model.

## B. Real-world 3D Environment Relighting

We conduct experiments on four scenes from the Eye-ful Tower dataset [49]: `officela`, `office_view2`, `kitchen`, and `seating_area`. For each scene, we first mark the most prominent light sources in 3D. This is done by selecting image pixels in 2D and back-projecting them to 3D, and modifying the color of the 3D points corresponding to the light sources to 1 and other points to 0. The resulting 3D light masks are then rendered along the same elliptical camera path to obtain per-frame 2D light masks that are aligned with the input video frames. Figure 10 shows the comparison on more views.

### B.1. Details of User Study

In our user study, participants were presented with two videos generated by one of four sources: LightLab [29], DiffusionRenderer [26] (Cosmos), our method, and the input video. We include the input video as a baseline since, although it does not respect the target lighting condition, it is entirely free of relighting artifacts and therefore often exhibits very high quality. The two videos presented to each user share the same camera path. We provide users with a text description of the target lighting condition to reduce ambiguity, specifying the on/off state, and color of each indoor light source, as well as the exterior lighting condition.

Figure 9. **Illustration of our user study.** The user is asked to provide feedback of the questions based on the explained criteria.

Figure 9 shows the interface used in our user study. Each of 12 users is asked to make 18 comparisons, making up a total of 216 responses. Users were asked to compare the two videos and determine which better satisfies three criteria: frame-to-frame consistency, object identity preservation, and lighting similarity to the target condition. Each comparison was presented in randomized order to mitigate position bias, and responses were aggregated across all raters. The lower win rate against Input (68.5%) is expected, as the original video naturally preserves temporal consistency and object identity perfectly—yet users still preferred our method, indicating they value the improved lighting fidelity our approach provides.

### B.2. Qualitative Comparison

We present some of the qualitative results in Figures 13, and 14 (Light-A-Video, DiffusionRenderer (Cosmos), LightLab and Ours). Please refer to the [webpage](#) for more video examples.

### B.3. Fine-grained Lighting Control

We show the fine-grained individual light source control in Figures 15 and 16. Please refer to the [webpage](#) to see more examples.

## C. Data Generation Pipeline Visualization

Figure 17 illustrates the training data used in our approach. Built upon the Infinigen framework [36, 37], our training set comprises diverse indoor scenes featuring various floor plans and materials. Our OLAT (One-Light-at-a-Time) data generation pipeline simulates a wide range of lighting conditions within each scene. Notably, both interior lighting (including color, intensity, and on/off states) and exterior lighting are systematically varied across different instances.

**Images are shown on the following pages.**





Figure 10. **Real-world relighting results (Extension of Figure 4).** Each pair of rows shows a video frame and a corresponding inset of relighting results obtained using Nano Banana, DiffusionRenderer (Cosmos version), LightLab, and our method. The light near the window is turned off, so no illumination should enter through it. In the top frame, our method removes the TV reflection caused by the light that has been turned off, while the baselines fail to remove it. In the middle frame, our model correctly maintains the correct reflection while removing the one that has been turned off, while also correctly removing the external lights. In contrast, Nano Banana and LightLab make the removed reflections overly dark and fail to turn off the external lights. In the bottom frame, only our method is able to remove the wall reflection without corrupting textures. DiffusionRenderer struggles to control lighting conditions effectively, as the environment map lacks sufficient information for accurate light source conditioning. Additionally, the rendered results tend to be overly diffuse due to the inherent difficulty in precise material property estimation. See our webpage for interactive comparisons.




Figure 11. **Additional real-world results.** Here we present qualitative comparisons of our method with LightLab, Light-A-Video, and DiffusionRenderer on the kitchen scene from the Eyeful Tower dataset. LightLab, which relights each video frame independently, exhibits noticeable temporal flickering artifacts. Light-A-Video offers limited lighting control and produces low-quality outputs. DiffusionRenderer yields temporally consistent results but suffers from blurriness and poor visual quality. In contrast, our method achieves realistic relighting that adheres closely to the target lighting configuration. Notably, while other methods struggle with scenes involving multiple light sources, our approach effectively controls the lighting solely based on the provided target lights.



Figure 12. **Results on a held-out synthetic scene.** Here we show qualitative comparisons between our method and existing approaches—including LightLab, Light-A-Video, and DiffusionRenderer. These are conducted on a test scene from our Infinigen dataset. LightLab exhibits reduced flickering due to the relatively stable exposure in synthetic scenes; however, it fails to turn on the exterior lighting and does not reproduce accurate color tones. Light-A-Video successfully turns on the exterior light but struggles to recover the correct indoor lighting. DiffusionRenderer, while achieving better temporal consistency than in real-world scenarios, also fails to reproduce accurate color. In contrast, our method not only turns on both the exterior light and the indoor lamp but also achieves accurate color rendering for both, demonstrating superior visual fidelity and lighting control.






Figure 13. **Additional real-world results.** Qualitative comparisons of our method with LightLab, Light-A-Video, and DiffusionRenderer on the `officela` scene from the Eyeful Tower dataset. LightLab independently relights each frame, resulting in noticeable temporal flickering artifacts. Light-A-Video offers limited lighting control and produces low-quality results. DiffusionRenderer achieves temporal consistency but suffers from blurriness and low visual fidelity. In contrast, our method realistically relights the scene based on the target lighting conditions. Notably, it accurately generates cast shadows for the object on the table and effectively removes undesired reflectance caused by lighting.




Figure 14. **Additional real-world results.** Qualitative comparisons of our method with LightLab, Light-A-Video, and DiffusionRenderer on the `seating_area` scene from the Eyeful Tower dataset. Since LightLab separately relights each frame of the video it suffers from flickering temporal artifacts. Light-A-Video provides very limited control and low-quality results. DiffusionRenderer provides consistent but low-quality and blurry results. Our approach is able to realistically relight the scene according to the provided target lights.



Figure 15. **Fine-grained color control on a real-world scene.** Top row: input video frames. Middle row: target lighting conditions indicating the ceiling lamps are emitting red and green respectively. Bottom row: our relit results. Our method is able to plausibly relight the scene even when the target lighting colors are non-ordinary.




Figure 16. **Controllable lighting in 3D scenes.** Our method enables fine-grained control over individual light sources in reconstructed environments. Users can selectively toggle indoor and exterior lights (a→b), activate specific ceiling lamp groups (b→c), and modify emission colors (c→d), demonstrating flexible and realistic relighting capabilities.





Figure 17. **An example of a scene used to train our method.** Each column corresponds to a different lighting condition, and each row to a different view. We train our relighting model by feeding into it a video of the scene under one lighting condition and predicting the same scene under a target lighting condition.



Figure 17. **(continued.)** Another scene from the dataset we use to train our method.