IllumiNeRF: 3D Relighting Without Inverse Rendering

Xiaoming Zhao^{1,3*} Pratul P. Srinivasan² Dor Verbin²
Keunhong Park¹ Ricardo Martin-Brualla¹ Philipp Henzler¹
Google Research ²Google DeepMind ³University of Illinois Urbana-Champaign

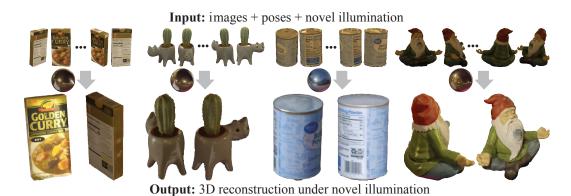


Figure 1: Given a set of posed input images under an *unknown* lighting (four exemplar images from the set are shown on top), IllumiNeRF produces high-quality novel views (bottom) relit under a target lighting (illustrated as chrome balls). Inputs obtained from the Stanford-ORB dataset [27].

Abstract

Existing methods for relightable view synthesis — using a set of images of an object under unknown lighting to recover a 3D representation that can be rendered from novel viewpoints under a target illumination — are based on inverse rendering, and attempt to disentangle the object geometry, materials, and lighting that explain the input images. Furthermore, this typically involves optimization through differentiable Monte Carlo rendering, which is brittle and computationally-expensive. In this work, we propose a simpler approach: we first relight each input image using an image diffusion model conditioned on target environment lighting and estimated object geometry. We then reconstruct a Neural Radiance Field (NeRF) with these relit images, from which we render novel views under the target lighting. We demonstrate that this strategy is surprisingly competitive and achieves state-of-the-art results on multiple relighting benchmarks. Please see our project page at illuminerf.github.io.

1 Introduction

Capturing an object's appearance so that it can be accurately rendered in novel environments is a central problem in computer vision whose solution would democratize 3D content creation for augmented and virtual reality, photography, filmmaking, and game development. Recent advances in view synthesis [36] have made impressive progress in reconstructing a 3D representation that can be rendered from novel viewpoints, using just a set of observed images. However, those methods

38th Conference on Neural Information Processing Systems (NeurIPS 2024).

^{*}Work done as an intern at Google.

typically only recover the appearance of the object under the captured illumination, and *relightable view synthesis* — rendering novel views of the captured object under arbitrary target environments — remains challenging.

Recent methods for recovering relightable 3D representations treat this task as *inverse rendering*, and attempt to estimate the geometry, materials, and illumination that jointly explain the input images using physically-based rendering methods. These approaches typically involve gradient-based optimization through differentiable Monte Carlo rendering procedures, which are noisy and computationally-expensive. Moreover, the inverse rendering optimization problem is brittle and inherently ambiguous; many potential sets of geometry, materials, and lighting can explain the input images, but many of these incorrect explanations produce obviously implausible renderings when rendered under novel unobserved illumination.

We propose a different approach that avoids inverse rendering and instead leverages a generative image model fine-tuned for the task of relighting. Given a set of images viewing an object and a desired target illumination, we use a single-image 2D Relighting Diffusion Model that outputs relit images of the object under the target illumination. Due to the ambiguous nature of the problem, each sample of the generative model encodes a different explanation of the object's materials, geometry and the input illumination. However, as opposed to optimization-based inverse rendering, such samples are all *plausible* relit images since they are the output of the trained diffusion model.

Instead of attempting to recover a single explanation of the underlying object's appearance, we sample multiple plausible relit images for each observed viewpoint, and treat the underlying explanations as samples of unobserved latent variables. To recover a final consistent 3D representation of the relit object, we use the full set of sampled relit images from all viewpoints to train a "latent NeRF" that reconciles all the samples into a single 3D representation, which can be rendered to produce plausible relit images from novel viewpoints.

The key contribution of our work is a new paradigm for relightable 3D reconstruction that replaces 3D inverse rendering with: generating samples with a single-image 2D Relighting Diffusion Model followed by distilling these samples into a 3D latent NeRF representation. We demonstrate that this strategy is surprisingly competitive and outperforms existing most 3D inverse rendering baselines on the TensoIR [23] and Stanford-ORB [27] relighting and view synthesis benchmarks.

2 Related Work

Our work addresses the task of relightable 3D reconstruction by using a lighting-conditioned diffusion model as a generative prior for single-image relighting. It is closely related to prior work in relightable 3D reconstruction, inverse rendering, and single-image relighting. Below, we review these lines of work and discuss how they relate to our proposed approach.

Relightable 3D Reconstruction The goal of relightable 3D reconstruction is to reconstruct a 3D representation of an object that can be relit by novel illumination conditions and rendered from novel camera poses. In scenarios where an object is observed under multiple lighting conditions [12], it is trivial to render its appearance under novel illumination that is a linear combination of the observed lighting conditions, due to the linear behavior of light. This approach is generally limited to laboratory capture scenarios where it is possible to observe an object under a lighting basis.

In more casual capture scenarios, the object is observed under just a single or a small handful of lighting conditions. Existing works typically address this setting using methods based on inverse rendering that explicitly factor an object's appearance into the underlying 3D geometry, object material properties, and lighting that jointly explain the observed images. State-of-the-art approaches to 3D inverse rendering [9, 10, 17, 23, 26, 33, 38, 46, 47] generally utilize the following strategy: they start with a neural field representation of 3D geometry (typically volume density as in NeRF [36], hybrid volume-surface representations as in NeuS [57] and VolSDF [59], or meshes extracted from neural field representations) from the input images, equip the model with a representation of surface materials (*e.g.* spatially-varying BRDF parameters) and lighting, and jointly optimize these factors through a differentiable physics-based rendering procedure [40]. While methods may differ in their choice of geometry, material, and lighting representations, and employ different techniques to accelerate the evaluation of the rendering integral, they generally all follow this same high-level inverse rendering strategy. Unfortunately, even if the geometry is known, inverse rendering

is a notoriously ambiguous problem [43, 52] and many combinations of materials and lighting can explain an object's appearance. However, not all of these combinations are plausible, and incorrect factorizations that explain observed images under one lighting condition may produce glaring artifacts when rendered under different lighting. Furthermore, differentiable physics-based rendering is computationally-expensive as thousands of samples are needed for Monte Carlo estimates of the rendering integral, typically requires custom implementations [2, 3, 22, 28, 32, 35, 54], and the resulting inverse rendering loss landscape is non-smooth and difficult to optimize effectively with gradient descent [14].

Single Image Relighting Instead of using inverse rendering to recover object material parameters which can be relit with physically-based rendering techniques, we train a diffusion model that can directly sample from the distribution of relit images conditioned on a target lighting condition. This diffusion model is essentially a generative single-image relighting model. Early single image relighting techniques employed optimization-based inverse rendering [4]. Subsequent methods trained deep convolutional neural networks to output image geometry, materials, and lighting [29, 30], or in some cases, to directly output relit images [48, 7, 8].

Most related to our method are a few recent works that have trained diffusion models for single image relighting. LightIt [25] trains a model similar to ControlNet [63] to relight outdoor images under arbitrary sun positions conditioned on input normals and shading. DiffusionLight [41] estimates the lighting of an image by using a ControlNet to inpaint the color pixels of a chrome ball in the middle of the scene, from which an environment map can be recovered.

Most similar to our work is the concurrent method of DiLightNet [61] that focuses on single image relighting. DiLightNet uses a ControlNet-based [63] approach to condition a single-image relighting diffusion model on a target environment map. DiLightNet uses a set of "radiance cues" [15] — renderings of the object's geometry (obtained from an off-the-shelf monocular depth network) with various roughness levels under the target environment illumination — as conditioning. Our method instead focuses on 3D relighting, where multiple of images of an object are available. It uses a similar single-image relighting diffusion model conditioned on radiance cues. Unlike DiLightNet which uses geometry from monocular depth estimation to render radiance cues, we use geometry estimated from the input views using a state-of-the-art surface reconstruction method [56]. This allows our model to better model complex light transport effects such as interreflections caused by occluded geometry.

3 Method

3.1 Problem Formulation

Given a dataset of images of an object and corresponding camera poses $\mathcal{D} = \{(I_i, \pi_i)\}_{i=1}^N$, the general goal of relightable 3D reconstruction is to estimate a model with parameters θ that when rendered, produces relit versions of the dataset under unobserved target illumination L^T . This can be expressed as:

$$\theta^* = \operatorname*{argmax}_{\theta} p(\mathcal{D}_{\theta}^T | \mathcal{D}), \tag{1}$$

where $\mathcal{D}_{\theta}^T \triangleq \left\{ \left(\mathrm{relight}(\mathcal{D}, L^T, \pi_i, \theta), \pi_i \right) \right\}_{i=1}^N$ is a relit version of the original dataset under target illumination L^T using model θ . Note that Eq. (1) only maximizes the likelihood of the original given poses after relighting. However, by using view synthesis, we can then turn the collection of relit images into a 3D representation which can be rendered from arbitrary poses. For brevity, we therefore omit the implicit dependence of \mathcal{D}^T in θ .

This relighting problem has traditionally been solved by using inverse rendering. Inverse rendering techniques do not maximize the probability of the relit renderings, but instead recover a single point estimate of the most likely scene geometry G, materials M, and lighting L (note that this is the "source" lighting condition for the observed images) that together explain the input dataset, and then use physically-based rendering to relight this factorized explanation under the target lighting. Inverse rendering seeks to recover $\theta^{\rm IR} = (G^\star, M^\star)$, where:

$$G^{\star}, M^{\star}, L^{\star} = \underset{G, M, L}{\operatorname{argmax}} p(G, M, L | \mathcal{D}) = \underset{G, M, L}{\operatorname{argmax}} p(\mathcal{D} | G, M, L) p(G, M, L). \tag{2}$$

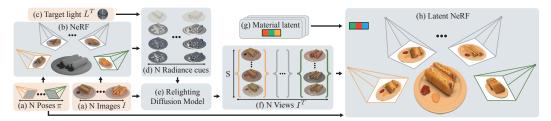


Figure 2: **Overview.** Given a set of images I and camera poses π in (a), we run NeRF to extract the 3D geometry as in (b). Based on this geometry and a target light shown in (c), we create radiance cues for each given input view as in (d). Next, we independently relight each input image using a single-image Relighting Diffusion Model illustrated in (e) and sample S possible solutions for each given view displayed in (f). Finally, we distill the relit set of images into a 3D representation through a Latent NeRF optimization as in (g) and (h).

The first data likelihood term is computed by physics-based rendering of the estimated model and the second prior term is often factorized into separate handcrafted priors on geometry, materials, and lighting [23, 33, 43].

A relighting approach based on inverse rendering then renders each image I in \mathcal{D} corresponding to camera pose π using the recovered geometry and materials, illuminated by the target lighting L^T , resulting in relight($\mathcal{D}, L^T, \pi, \theta^{\mathrm{IR}}$). This approach has three main issues. First, the differentiable rendering procedures used to compute the gradient of the likelihood term are computationally-expensive. Second, it requires careful modeling of light transport which is cumbersome and existing differentiable renderers do not account for many types of lighting and material effects seen in the real world. Third, there are often ambiguities between M and L, meaning that any errors in their decomposition may be apparent in the relit data. It is quite difficult to design effective handcrafted priors on geometry, materials, and lighting, so inverse rendering procedures frequently recover explanations that have a high data likelihood (are able to render the observed data) but produce clearly incorrect results when re-rendered under different illumination.

3.2 Model Overview

We propose an approach that attempts to maximize the probability of relit images in Eq. (1) without using an explicit physically-based model of the object's lighting or materials. First, let us introduce a latent variable Z that can be thought of as implicitly representing the input images' lighting along with the object's material and geometry parameters. We can write the likelihood of the relit data as:

$$p(\mathcal{D}^T|\mathcal{D}) = \int p(\mathcal{D}^T, Z|\mathcal{D}) dZ = \int p(\mathcal{D}^T|Z, \mathcal{D}) p(Z|\mathcal{D}) dZ.$$
 (3)

Introducing these latent variables lets us consider all relit renderings in the dataset, $\mathcal{D}_i^T \triangleq (I_i^T, \pi_i)$, as conditionally independent, since the rendering under the target lighting L^T is deterministic given the object's geometry and materials. This enables writing the likelihood as:

$$p(\mathcal{D}^T|\mathcal{D}) = \int \left[\prod_{i=1}^N p(\mathcal{D}_i^T|Z_i, \mathcal{D}_i) \right] \underbrace{p(Z|\mathcal{D}) \, dZ}_{\text{latent Prior}}.$$
 (4)

We propose to model this with a latent NeRF model, as used by Martin-Brualla *et al.* [34] that is able to render novel views under the target illumination for any sampled latent vector. We describe this model in Sec. 3.3. We train this NeRF model by generating a large quantity of sampled relit images with the same target lighting but with different (unknown) latent vectors using a *Relighting Diffusion Model* which we will describe in Sec. 3.4. In this way, the latent NeRF model effectively distills a large dataset of relit images sampled by the diffusion model into a single 3D representation that can render novel views of the object under the target lighting for any sampled latent.

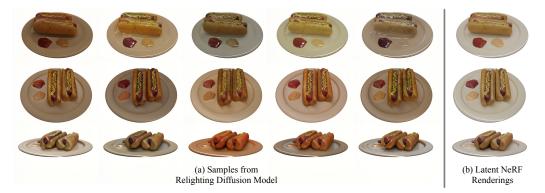


Figure 3: **Relit samples** *vs.* **latent NeRF.** (a) Samples of the Relighting Diffusion Model (Sec. 3.4) for the same target environment map, and (b) renderings from the optimized Latent NeRF (Sec. 3.3) for a fixed value of the latent. The diffusion samples correspond to different latent explanations of the scene and our latent NeRF optimization is able to effectively optimize these latent variables along with the NeRF model's parameters to produce consistent renderings for each latent explanation.

3.3 Latent NeRF Model

We wish to model the distribution in Eq. (4) in a manner that lets us render images that correspond to relit views of the object for any sampled latent Z. We choose to model this with a latent code NeRF 3D representation, inspired by prior works that condition NeRFs on latent codes to represent sources of variation such as the time of day during capture [34]. This latent NeRF optimizes a set of latent codes that are used to condition the view-dependent color function represented by the NeRF, enabling it to render novel views of the relit object under the target illumination for any sampled latent code. In our implementation, the latent NeRF's geometry does not depend on the latent code, so the latent code may be interpreted as only representing the object's material properties.

To optimize the parameters θ of the latent NeRF model, we maximize the log-likelihood, which by using Eq. (4), can be written as the following maximization problem:

$$\theta^* = \operatorname*{argmax}_{\theta} \log p(\mathcal{D}_{\theta}^T | \mathcal{D}) = \operatorname*{argmax}_{\theta} \log \int \left[\prod_{i=1}^N p(\mathcal{D}_i^T | Z_i, \mathcal{D}_i) \right] p(Z | \mathcal{D}) dZ. \tag{5}$$

Because integrating over all possible latents Z is intractable, we use a heuristic inference strategy and replace the integral with the maximum a posteriori (MAP) estimate of Z:

$$\theta^* \approx \operatorname*{argmax}_{\theta} \max_{Z} \left\{ \sum_{i=1}^{N} \log p(\mathcal{D}_i^T | Z_i, \mathcal{D}_i) + \log p(Z | \mathcal{D}) \right\}.$$
 (6)

By assuming a Gaussian model over the data given the materials, the first term in Eq. (6) is a reconstruction loss over the images. However, since we do not have access to the true latent vector Z, we assume a uniform prior over them, turning the second term in Eq. (6) into a constant. In practice, similar to prior work on NeRFs optimized to generate new views given a dataset containing images with varying appearance, we rely on the NeRF model to resolve any mismatches in the appearance of different images [34]. See Fig. 3 for illustrations. The minimization of the negative log-likelihood can then be written as:

$$\theta^* = \arg\min_{\theta} \min_{Z} \sum_{i=1}^{N} \|\mathcal{D}_i^T - \text{latent-NeRF}(\theta, Z_i, \pi_i)\|^2.$$
 (7)

3.4 Relighting Diffusion Model

In order to train the latent NeRF model described in Sec. 3.3, we use a Relighting Diffusion Model (RDM) to generate S samples for each viewpoint from $p(\mathcal{D}_i^T|\mathcal{D}_i)$. In other words, given an input image and target lighting L^T , the single-image RDM samples S images corresponding to relit



Figure 4: Example radiance cues for a view of the 'hotdog' scene.

versions of D_i that have a high likelihood given the new target light L^T . We then associate each sample $s \in \{1, ..., S\}$ with its own latent code $Z_{i,s}$ and sum over all samples when training the latent NeRF (Eq. (7)).

Our RDM is implemented as an image denoising diffusion model that is conditioned by the input image and target lighting. To encode the target lighting, we use image-space radiance cues [15, 44, 61], visualized in Fig. 4. These radiance cues are generated by using a simple shading model to render a handful of images of the object's estimated geometry under the target lighting. This procedure is designed to provide information about the effects of specularities, shadows, and global illumination, without requiring the diffusion network to learn these effects from scratch. In our experiments, we use four different pre-defined materials to render radiance cues: one diffuse material with a pure white albedo, and three purely-specular materials with roughness values $\{0.05, 0.13, 0.34\}$. We use GGX [55] as the shading model. For more details, please refer to Sec. A.2.

The RDM architecture consists of a pretrained latent image diffusion model, similar to StableDiffusion [45], and uses a ControlNet [63] based approach to condition on the radiance cues. Please refer to Sec. A.3 for more architecture details.

4 Experiments

4.1 Experimental Setup

Relighting Dataset We render objects from Objaverse [13] under varying poses and illuminations. For each object, we randomly sample 4 poses, and render each under 4 different lighting conditions. We represent the lighting as HDR environment maps, and randomly sample from a dataset of 509 environment maps from Polyhaven [60]. For more details, see Sec. A.4.

Evaluation Datasets We evaluate our method on two datasets: TensoIR [23], a synthetic benchmark, and Stanford-ORB [27], a real-world benchmark. TensoIR contains renderings of four synthetic objects rendered under six lighting conditions. Following [23], we use the training split of 100 renderings with "sunset" lighting as input $\{I_i\}$. We then evaluate on 200 poses, each of which has renderings under five different environment maps, *i.e.*, "bridge", "city", "fireplace", "forest", and "night", for a total of 4000 renderings. Stanford-ORB is a real-world benchmark for inverse rendering on data captured in the wild. It contains 14 objects with various materials and captures each object under three different lighting settings, resulting in 42 (object, lighting) pairs. For the task of relighting, we are given images of an object under a single lighting condition and follow the benchmark protocol to evaluate relit images of the object under the two target lighting settings.

Baselines We compare our method to several existing inverse rendering approaches. On both benchmarks, we compare to NeRFactor [65] and InvRender [66]. On the synthetic benchmark, we additionally compare to TensoIR [23], the current top-performing approach on that benchmark. For the Stanford-ORB benchmark, we additionally compare to PhySG [62], NVDiffRec [38], NeRD [10], NVDiffRecMC [17], and Neural-PBIR [47].

Our Model Inference At inference time, the ideal embedding vector Z that best corresponds to the actual material is unknown. One approach to find this vector is to optimize Z to match a subset of the test set images (as in [34]). However, to ensure a fair comparison, we avoid this optimization. Instead, we set Z=0 for all views when rendering test images.

Table 1: **TensoIR benchmark** [23]. We evaluate four objects. Each object has five target lightings, each of which is associated with 200 poses, resulting in evaluating 4000 renderings in total. Running time for baselines are copied from [23]. Our time is A (geometry optimization on GPU) + B (diffusion sampling on TPU) + C (latent NeRF optimization on GPU). Best and 2nd-best are highlighted.

	PSNR↑	SSIM↑	LPIPS↓	Wall-clock Time↓	Device
NeRFactor [65]	23.383	0.908	0.131	> 100 h	a RTX 2080 Ti
InvRender [66]	23.973	0.901	0.101	15 h	a RTX 2080 Ti
TensoIR [23]	28.580	0.944	0.081	5 h	a RTX 2080 Ti
Ours	29.709	0.947	0.072	0.75 h + 1 h + 0.75 h	16 A100 40GB + a TPUv5
Ours (single GPU)	29.245	0.946	0.073	2 h + 1 h + 2 h	a A100 40GB + a TPUv5

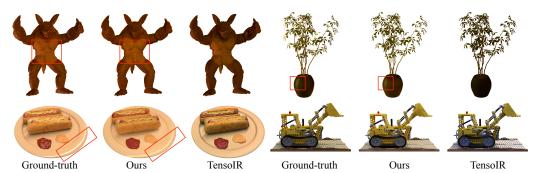


Figure 5: **Qualitative results on TensoIR.** Renderings from all approaches have been rescaled with respect to the ground-truth as mentioned in Eq. (4.1). Unlike TensoIR, our method faithfully recovers specular highlights and colors as indicated in red.

Evaluation Metrics For both benchmarks, we evaluate the quality of 3D relighting by reporting image metrics for rendered images. We report PSNR, SSIM [58], and LPIPS-VGG [64] on low dynamic range (LDR) images. Additionally, we report PSNR on high dynamic range (HDR) images on Stanford-ORB following the benchmark protocol, denoted as PSNR-H while the PSNR on LDR images is marked as PSNR-L. For approaches that do not produce HDR renderings, including ours, we convert the LDR renderings to linear values by using the inverse of the sRGB tone mapping curve. Due to the inherent ambiguities for the relighting task, we follow prior works [23, 27] and apply a channel-wise scale factor to RGB channels to match the ground truth image before computing metrics. Following established evaluation practices on Stanford-ORB, we compute the scale per output image individually whereas for TensoIR we compute a global scale factor that is used for all output images.²

4.2 Benchmarking

Unless otherwise specified, all results are produced using S=16 samples (see Sec. 3.4) and make use of 16 A100 40GB GPUs (batch size of 2^{14} rays for NeRF optimization). We also provide results on a single A100 40GB GPU (batch size of 2^{13} for NeRF optimization).

We report quantitative results on the TensoIR benchmark in Tab. 1, and show qualitative examples in Fig. 5. We significantly outperform all competitors quantitatively on all metrics with comparable or improved wall-clock time. Visually our method is capable of recovering specular highlights whereas prior methods struggle to model these.

Similarly, we report results on Stanford-ORB in Tab. 2 and Fig. 6. Our proposed approach quantitatively improves upon all baselines, except those of Neural-PBIR [47], indicating the effectiveness of IllumiNeRF in real world scenarios. Note that although Neural-PBIR achieves better metrics than us, Fig. 6 shows that their relighting results are mostly diffuse, even for highly-glossy objects, and that they lack many of the strong specular highlights that our method is able to recover. This behavior of their model may explain their better metrics despite worse qualitative performance for specular highlights, because the illumination maps provided by Stanford-ORB do not correspond to the

²Please refer to https://github.com/StanfordORB/Stanford-ORB/blob/962ea6d2cc/scripts/test.py#L36 and https://github.com/Haian-Jin/TensoIR/blob/2a7a4d00/renderer.py#L12.



Figure 6: **Qualitative results on Stanford-ORB.** Renderings from all approaches have been rescaled with respect to the ground-truth as mentioned in Sec. 4.1. Areas where our approach performs well are highlighted. Our approach produces high-quality renderings with plausible specular reflections.

incident illumination at the object's location, since they were captured using a light probe which was moved for each image in the dataset [27]. This means that even given perfect materials and geometry, the images relit by *any* method cannot match with the true captured images, which is most noticeable in specular highlights. This mismatch penalizes methods like ours, which recover such specularities, over ones that recover mostly diffuse appearance with no apparent specular highlights [52]. For a more detailed discussion see Sec. B.

We also provide qualitative results for different latent codes in Fig. 7. These results demonstrate that the optimized latent codes effectively capture various plausible explanations of the materials.

Table 2: **Stanford-ORB benchmark** [27]. We evaluate 14 objects, each of which was captured under three different lightings. For each (object, lighting) pair, we evaluate renderings of the same object under the other two lightings, resulting in evaluating 836 renderings. †denotes models trained with the ground-truth 3D scans and pseudo materials optimized from light-box captures. Best and 2nd-best are highlighted.

	PSNR-H↑	PSNR-L↑	SSIM↑	LPIPS↓
NVDiffRecMC [17]†	25.08	32.28	0.974	0.027
NVDiffRec [38]†	24.93	32.42	0.975	0.027
PhySG [62]	21.81	28.11	0.960	0.055
NVDiffRec [38]	22.91	29.72	0.963	0.039
NeRD [10]	23.29	29.65	0.957	0.059
NeRFactor [65]	23.54	30.38	0.969	0.048
InvRender [66]	23.76	30.83	0.970	0.046
NVDiffRecMC [17]	24.43	31.60	0.972	0.036
Neural-PBIR [47]	26.01	33.26	0.979	0.023
Ours	25.42	32.62	0.976	0.027
Ours (single GPU)	25.56	32.74	0.976	0.027

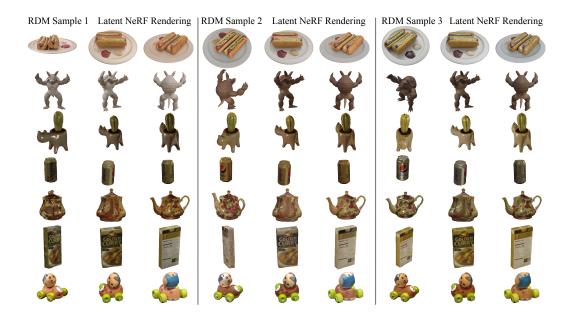


Figure 7: **Renderings from various latents.** Each column shows 1) a Relighting Diffusion Model (RDM) sample and 2) two latent NeRF renderings using the sample's latent code. The diffusion samples are selected uniformly from all N (#views) \times S (#samples per view) diffusion generations. Each row shows results from the same object and lighting with latent codes capturing various plausible explanations of the materials.

4.3 Ablations

We evaluate ablations of our model on TensoIR's hotdog scene in Tab. 3, and visualize them in Fig. 8. We reach the following conclusions: 1) The latent NeRF model is essential: optimizing a standard NeRF cannot reconcile variations across views, even if we only generate a single sample per viewpoint for optimization (S=1). 2) More diffusion samples help: by increasing S, the number of samples from the RDM per viewpoint, we observe consistent improvements across almost all metrics. This corroborates our intuition that using an increased number of samples helps the latent NeRF effectively fit the target distribution (Eq. (4)) in a more stable way.

Table 3: **Ablations.** We conduct ablation studies on the Hotdog scene from TensoIR [23]. We evaluate renderings of 200 novel test camera poses, each under five target environment map lighting conditions, resulting in evaluating 1000 renderings in total. Best is highlighted.

S	Latent	PSNR↑	SSIM↑	LPIPS↓
1	Х	24.957	0.921	0.099
1	✓	26.321	0.925	0.097
4	✓	27.409	0.936	0.087
16	√	27.950	0.939	0.082

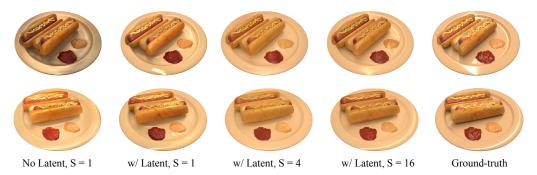


Figure 8: Using a standard NeRF instead of a latent NeRF model is unable to reconcile training samples with different underlying latent explanations. Using a latent NeRF model significantly increases the accuracy of rendered specular appearance, and increasing the number of samples S from the RDM used to train the latent NeRF model further increases the quality of the output renderings.

4.4 Limitations

Our model relies on high quality geometry estimated by UniSDF [56] (see Sec. A.1) to provide sufficiently good radiance cues for conditioning the RDM (Sec. 3.4). Any missing structure will lead our model to miss specular reflections, as seen on the top left of the salt can result in Fig. 6's second column. Errors in geometry also affect the quality of synthesized novel views, *e.g.* the missing thin branches from the plant in Fig. 5 or fine details of the cactus (column 4) in Fig. 6. Note that our RDM, trained on high-quality synthetic geometry, will inherently improve with future advances in geometry reconstruction. Our approach is not suited for real-time relighting, as it requires generating new samples with the RDM and optimizing a NeRF for any new target lighting condition.

5 Conclusion

We have proposed a new paradigm for the task of relightable 3D reconstruction. Instead of decomposing an object's appearance into lighting and material factors and then relighting the object with physically-based rendering, we use a single-image Relighting Diffusion Model (RDM) to sample a varied collection of proposed relit images given a target illumination, and distill these samples into a single consistent 3D latent NeRF representation. This 3D representation can be rendered to synthesize novel views of the object under the target lighting. Perhaps surprisingly, this paradigm consistently outperforms existing inverse rendering methods on synthetic and real-world object relighting benchmarks. This new paradigm's success is likely due to the RDM's ability to generate a large number of proposals for the new relit images. This is in contrast to prior works based on inverse rendering, which first estimates a single material model and then uses it for relighting, since errors in material estimation may propagate to the relit images. We believe that this paradigm may be used to improve data capture, material and lighting estimation, and that it may be used to do so robustly on real-world data.

Acknowledgements

We would like to thank Ben Poole and Ruiqi Gao for insightful discussions. We thank Yunzhi Zhang and Zhengfei Kuang for providing their qualitative results for the Stanford-ORB [27] baseline, and Haian Jin for the TensoIR [23] baseline results. We are also grateful to Abhijit Kundu and Henna Nandwani for their infrastructure support.

References

- [1] H. Aanæs, R. R. Jensen, G. Vogiatzis, E. Tola, and A. B. Dahl. Large-Scale Data for Multiple-View Stereopsis. *IJCV*, 2016. 14
- [2] S. Bangaru, M. Gharbi, T.-M. Li, F. Luan, K. Sunkavalli, M. Hasan, S. Bi, Z. Xu, G. Bernstein, and F. Durand. Differentiable Rendering of Neural SDFs through Reparameterization. In *SIGGRAPH Asia*, 2022. 3
- [3] S. Bangaru, L. Wu, T.-M. Li, J. Munkberg, G. Bernstein, J. Ragan-Kelley, F. Durand, A. Lefohn, and Y. He. SLANG.D: Fast, Modular and Differentiable Shader Programming. ACM TOG, 2023. 3
- [4] J. T. Barron and J. Malik. Shape, Illumination, and Reflectance from Shading. TPAMI, 2014. 3
- [5] J. T. Barron, B. Mildenhall, D. Verbin, P. P. Srinivasan, and P. Hedman. Mip-NeRF 360: Unbounded Anti-Aliased Neural Radiance Fields. In CVPR, 2021. 14
- [6] J. T. Barron, B. Mildenhall, D. Verbin, P. P. Srinivasan, and P. Hedman. Zip-NeRF: Anti-Aliased Grid-Based Neural Radiance Fields. In *ICCV*, 2023. 14
- [7] A. Bhattad and D. A. Forsyth. Cut-and-Paste Object Insertion by Enabling Deep Image Prior for Reshading. In 3DV, 2022.
- [8] A. Bhattad, J. Soole, and D. Forsyth. StyLitGAN: Image-Based Relighting via Latent Control. In CVPR, 2024. 3
- [9] S. Bi, Z. Xu, P. Srinivasan, B. Mildenhall, K. Sunkavalli, M. Hašan, Y. Hold-Geoffroy, D. Kriegman, and R. Ramamoorthi. Neural Reflectance Fields for Appearance Acquisition. *ArXiv*, 2020. 2
- [10] M. Boss, R. Braun, V. Jampani, J. T. Barron, C. Liu, and H. P. A. Lensch. NeRD: Neural Reflectance Decomposition from Image Collections. In *ICCV*, 2021. 2, 6, 9, 17
- [11] J. Bradbury, R. Frostig, P. Hawkins, M. J. Johnson, C. Leary, D. Maclaurin, G. Necula, A. Paszke, J. VanderPlas, S. Wanderman-Milne, and Q. Zhang. JAX: Composable Transformations of Python+NumPy Programs, 2018. URL http://github.com/google/jax. 14, 15
- [12] P. Debevec, T. Hawkins, C. Tchou, H.-P. Duiker, W. Sarokin, and M. Sagar. Acquiring the Reflectance Field of a Human Face. In *ACM CGIT*, 2000. 2
- [13] M. Deitke, D. Schwenk, J. Salvador, L. Weihs, O. Michel, E. VanderBilt, L. Schmidt, K. Ehsani, A. Kemb-havi, and A. Farhadi. Objaverse: A Universe of Annotated 3D Objects. In CVPR, 2023. 6, 16
- [14] M. Fischer and T. Ritschel. Plateau-Reduced Differentiable Path Tracing. In CVPR, 2023. 3
- [15] D. Gao, G. Chen, Y. Dong, P. Peers, K. Xu, and X. Tong. Deferred Neural Lighting: Free-viewpoint Relighting from Unstructured Photographs. ACM TOG, 2020. 3, 6
- [16] K. Greff, F. Belletti, L. Beyer, C. Doersch, Y. Du, D. Duckworth, D. J. Fleet, D. Gnanapragasam, F. Golemo, C. Herrmann, T. Kipf, A. Kundu, D. Lagun, I. Laradji, H.-T. D. Liu, H. Meyer, Y. Miao, D. Nowrouzezahrai, C. Oztireli, E. Pot, N. Radwan, D. Rebain, S. Sabour, M. S. M. Sajjadi, M. Sela, V. Sitzmann, A. Stone, D. Sun, S. Vora, Z. Wang, T. Wu, K. M. Yi, F. Zhong, and A. Tagliasacchi. Kubric: a scalable dataset generator. In CVPR, 2022. 14
- [17] J. Hasselgren, N. Hofmann, and J. Munkberg. Shape, Light, and Material Decomposition from Images using Monte Carlo Rendering and Denoising. In *NeurIPS*, 2022. 2, 6, 9, 17
- [18] J. Heek, A. Levskaya, A. Oliver, M. Ritter, B. Rondepierre, A. Steiner, and M. van Zee. Flax: A Neural Network Library and Ecosystem for JAX, 2023. URL http://github.com/google/flax. 16
- [19] D. Hendrycks and K. Gimpel. Gaussian Error Linear Units (GELUs). ArXiv, 2016. 16

- [20] J. Ho and T. Salimans. Classifier-Free Diffusion Guidance. ArXiv, 2022. 16
- [21] J. Ho, A. Jain, and P. Abbeel. Denoising Diffusion Probabilistic Models. In NeuIPS, 2020. 16
- [22] W. Jakob, S. Speierer, N. Roussel, and D. Vicini. Dr.Jit: A Just-In-Time Compiler for Differentiable Rendering. *ACM TOG*, 2022. 3
- [23] H. Jin, I. Liu, P. Xu, X. Zhang, S. Han, S. Bi, X. Zhou, Z. Xu, and H. Su. TensoIR: Tensorial Inverse Rendering. In CVPR, 2023. 2, 4, 6, 7, 10, 11
- [24] D. P. Kingma and J. Ba. Adam: A Method for Stochastic Optimization. ArXiv, 2014. 14
- [25] P. Kocsis, J. Philip, K. Sunkavalli, M. Nießner, and Y. Hold-Geoffroy. LightIt: Illumination Modeling and Control for Diffusion Models. In CVPR, 2024. 3
- [26] Z. Kuang, K. Olszewski, M. Chai, Z. Huang, P. Achlioptas, and S. Tulyakov. NeROIC: Neural Rendering of Objects from Online Image Collections. ACM TOG, 2022.
- [27] Z. Kuang, Y. Zhang, H.-X. Yu, S. Agarwala, E. Wu, J. Wu, et al. Stanford-ORB: A Real-World 3D Object Inverse Rendering Benchmark. In *NeurIPS*, 2023. 1, 2, 6, 7, 8, 9, 11, 17
- [28] T.-M. Li, M. Aittala, F. Durand, and J. Lehtinen. Differentiable Monte Carlo Ray Tracing through Edge Sampling. ACM TOG, 2018. 3
- [29] Z. Li, Z. Xu, R. Ramamoorthi, K. Sunkavalli, and M. Chandraker. Learning to reconstruct shape and spatially-varying reflectance from a single image. *ACM TOG*, 2018. 3
- [30] Z. Li, M. Shafiei, R. Ramamoorthi, K. Sunkavalli, and M. Chandraker. Inverse Rendering for Complex Indoor Ccenes: Shape, Spatially-varying Lighting and SVBRDF from a Single Image. In *CVPR*, 2020. 3
- [31] W. E. Lorensen and H. E. Cline. Marching Cubes: A High Resolution 3D Surface Construction Algorithm. In ACM CGIT, 1987. 14
- [32] G. Loubet, N. Holzschuch, and W. Jakob. Reparameterizing Discontinuous Integrands for Differentiable Rendering. *ACM TOG*, 2019. 3
- [33] A. Mai, D. Verbin, F. Kuester, and S. Fridovich-Keil. Neural Microfacet Fields for Inverse Rendering. In ICCV, 2023. 2, 4
- [34] R. Martin-Brualla, N. Radwan, M. S. M. Sajjadi, J. T. Barron, A. Dosovitskiy, and D. Duckworth. NeRF in the Wild: Neural Radiance Fields for Unconstrained Photo Collections. In *CVPR*, 2021. 4, 5, 6
- [35] Merlin Nimier-David and Sébastien Speierer and Benoît Ruiz and Wenzel Jakob. Radiative backpropagation: An adjoint method for lightning-fast differentiable rendering. ACM TOG, 2020.
- [36] B. Mildenhall, P. P. Srinivasan, M. Tancik, J. T. Barron, R. Ramamoorthi, and R. Ng. NeRF: Representing Scenes as Neural Radiance Fields for View Synthesis. In ECCV, 2020. 1, 2
- [37] T. Müller, A. Evans, C. Schied, and A. Keller. Instant neural graphics primitives with a multiresolution hash encoding. *ACM TOG*, 2022. 14
- [38] J. Munkberg, J. Hasselgren, T. Shen, J. Gao, W. Chen, A. Evans, T. Müller, and S. Fidler. Extracting Triangular 3D Models, Materials, and Lighting From Images. In CVPR, 2022. 2, 6, 9, 17
- [39] M. Oechsle, S. Peng, and A. Geiger. UNISURF: Unifying Neural Implicit Surfaces and Radiance Fields for Multi-View Reconstruction. In ICCV, 2021. 14
- [40] M. Pharr, W. Jakob, and G. Humphreys. Physically Based Rendering: From Theory to Implementation. MIT Press, 2023. 2
- [41] P. Phongthawee, W. Chinchuthakun, N. Sinsunthithet, A. Raj, V. Jampani, P. Khungurn, and S. Suwa-janakorn. Diffusionlight: Light probes for free by painting a chrome ball. In CVPR, 2024. 3
- [42] A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin, J. Clark, G. Krueger, and I. Sutskever. Learning Transferable Visual Models From Natural Language Supervision. In *ICML*, 2021. 15, 16
- [43] R. Ramamoorthi and P. Hanrahan. A Signal-Processing Framework for Inverse Rendering. In ACM CGIT, 2001. 3, 4

- [44] P. Ren, J. Wang, J. M. Snyder, X. Tong, and B. Guo. Pocket Reflectometry. ACM TOG, 2011. 6
- [45] R. Rombach, A. Blattmann, D. Lorenz, P. Esser, and B. Ommer. High-Resolution Image Synthesis with Latent Diffusion Models. In CVPR, 2022. 6, 15
- [46] P. P. Srinivasan, B. Deng, X. Zhang, M. Tancik, B. Mildenhall, and J. T. Barron. Nerv: Neural reflectance and visibility fields for relighting and view synthesis. In CVPR, 2021. 2
- [47] C. Sun, G. Cai, Z. Li, K. Yan, C. Zhang, C. S. Marshall, J.-B. Huang, S. Zhao, and Z. Dong. Neural-PBIR Reconstruction of Shape, Material, and Illumination. In *ICCV*, 2023. 2, 6, 7, 9
- [48] T. Sun, J. T. Barron, Y.-T. Tsai, Z. Xu, X. Yu, G. Fyffe, C. Rhemann, J. Busch, P. Debevec, and R. Ramamoorthi. Single Image Portrait Relighting. ACM TOG, 2019. 3
- [49] J. Tang, Z. Chen, X. Chen, T. Wang, G. Zeng, and Z. Liu. LGM: Large Multi-View Gaussian Model for High-Resolution 3D Content Creation. ArXiv, 2024. 16
- [50] The Blender Foundation. Blender 2.93. URL https://www.blender.org/. 14
- [51] D. Verbin, P. Hedman, B. Mildenhall, T. E. Zickler, J. T. Barron, and P. P. Srinivasan. Ref-NeRF: Structured View-Dependent Appearance for Neural Radiance Fields. In CVPR, 2022. 14
- [52] D. Verbin, B. Mildenhall, P. Hedman, J. T. Barron, T. Zickler, and P. P. Srinivasan. Eclipse: Disambiguating Illumination and Materials using Unintended Shadows. In CVPR, 2024. 3, 8
- [53] D. Verbin, P. P. Srinivasan, P. Hedman, B. Mildenhall, B. Attal, R. Szeliski, and J. T. Barron. NeRF-Casting: Improved View-Dependent Appearance with Consistent Reflections. ArXiv, 2024. 14
- [54] D. Vicini, S. Speierer, and W. Jakob. Differentiable Signed Distance Function Rendering. ACM TOG, 2022. 3
- [55] B. Walter, S. Marschner, H. Li, and K. E. Torrance. Microfacet Models for Refraction through Rough Surfaces. In *Rendering Techniques*, 2007. 6, 14
- [56] F. Wang, M.-J. Rakotosaona, M. Niemeyer, R. Szeliski, M. Pollefeys, and F. Tombari. UniSDF: Unifying Neural Representations for High-Fidelity 3D Reconstruction of Complex Scenes with Reflections. In NeurIPS, 2024. 3, 10, 14
- [57] P. Wang, L. Liu, Y. Liu, C. Theobalt, T. Komura, and W. Wang. NeuS: Learning Neural Implicit Surfaces by Volume Rendering for Multi-view Reconstruction. In *NeurIPS*, 2021. 2
- [58] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli. Image quality assessment: from error visibility to structural similarity. TIP, 2004. 7
- [59] L. Yariv, J. Gu, Y. Kasten, and Y. Lipman. Volume Rendering of Neural Implicit Surfaces. In *NeurIPS*, 2021. 2
- [60] G. Zaal, R. Tuytel, R. Cilliers, J. R. Cock, A. Mischok, S. Majboroda, D. Savva, and J. Burger. Polyhaven: a Curated Public Asset Library for Visual Effects Artists and Game Designers, 2021. 6, 17
- [61] C. Zeng, Y. Dong, P. Peers, Y. Kong, H. Wu, and X. Tong. DiLightNet: Fine-grained Lighting Control for Diffusion-based Image Generation. In SIGGRAPH, 2024. 3, 6, 16
- [62] K. Zhang, F. Luan, Q. Wang, K. Bala, and N. Snavely. PhySG: Inverse Rendering with Spherical Gaussians for Physics-based Material Editing and Relighting. In CVPR, 2021. 6, 9, 17
- [63] L. Zhang, A. Rao, and M. Agrawala. Adding Conditional Control to Text-to-Image Diffusion Models. In ICCV, 2023. 3, 6, 15, 16
- [64] R. Zhang, P. Isola, A. A. Efros, E. Shechtman, and O. Wang. The Unreasonable Effectiveness of Deep Features as a Perceptual Metric. In CVPR, 2018.
- [65] X. Zhang, P. P. Srinivasan, B. Deng, P. E. Debevec, W. T. Freeman, and J. T. Barron. Nerfactor. In ACM TOG, 2021. 6, 7, 9, 14, 17
- [66] Y. Zhang, J. Sun, X. H. He, H. Fu, R. Jia, and X. Zhou. Modeling Indirect Illumination for Inverse Rendering. In CVPR, 2022. 6, 7, 9, 17

Appendix – IllumiNeRF: 3D Relighting Without Inverse Rendering

This appendix is organized as follows:

- 1. Sec. A provides more implementation details;
- 2. Sec. B details the inconsistent illumination issue on Stanford-ORB.

A Additional Implementation Details

A.1 Latent NeRF Model and Geometry Estimator

We use JAX [11] to implement both the geometry estimator and Latent NeRF model as UniSDF [56], a state-of-the-art volume rendering approach based on a signed distance function (SDF). The advantage of using UniSDF is that it enables easily extracting a mesh from the SDF, which we can then import into a standard rendering engine such as Blender [50] in order to compute radiance cues. Additionally, UniSDF decouples geometry from appearance, allowing us to fix the weights related to geometry and only optimize for weights that model the appearance. Note that future NeRF/SDF approaches with improved geometric reconstruction can be seamlessly integrated in our method.

Our parameterization of the UniSDF model is similar to the one used in the original paper for the DTU dataset [1], with four key changes. First, we reduce the number of rounds of proposal sampling (as introduced by mip-NeRF 360 [5]) from two to one, using 64 proposal samples. Second, we use the asymmetric predicted normal loss from NeRF-Casting [53]:

$$\mathcal{L}_{p} = \sum_{i} \left(\lambda_{1} \omega_{i} \| \nabla \mathbf{n}_{i} - \nabla \mathbf{n}_{i}' \|^{2} + \lambda_{2} \nabla \omega_{i} \| \mathbf{n}_{i} - \nabla \mathbf{n}_{i}' \|^{2} + \lambda_{3} \nabla \omega_{i} \| \nabla \mathbf{n}_{i} - \mathbf{n}_{i}' \|^{2} \right), \quad (S1)$$

where ω_i is the volume rendering weight of the *i*-th sample, \mathbb{Z} denotes the stop-gradient operator, \mathbf{n}_i and \mathbf{n}_i' are the *i*-th sample's density normals and predicted normals respectively (see [51]), and we set $\lambda_1 = \lambda_2 = 10^{-3}$, $\lambda_3 = 10^{-2}$. Third, like NeRF-Casting [53], we use an additional hash grid encoding [37] with 15 scales between a resolution of 32 and 4096, used only for outputting predicted normals. Fourth, we further encourage the local smoothness of the predicted normals \mathbf{n}' by using a smoothness loss similar to [39, 65]:

$$\mathcal{L}_s = \lambda_4 \sum_i \omega_i \|\mathbf{n}'(\mathbf{x}_i + \boldsymbol{\varepsilon}) - \mathbf{n}'(\mathbf{x}_i)\|^2,$$
 (S2)

where \mathbf{x}_i is the 3D position of the *i*-th sample, and $\varepsilon \sim \mathcal{N}(0, \sigma^2 I)$ is an isotropic Gaussian random variable used to perturb the sample locations. We set $\lambda_4 = 0.1$ and $\sigma = 0.01$.

We find that these modifications result in better and smoother geometry necessary for our model's ability to relight objects with specular highlights.

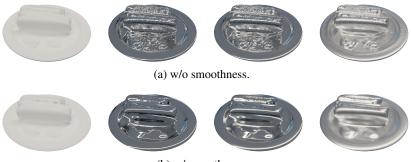
Finally, to incorporate the GLO embeddings, we utilize an MLP to predict an element-wise scale and shift value to be applied to the 'bottleneck' feature of UniSDF, similar to AffineGLO in Zip-NeRF [6].

For both geometry estimation and latent NeRF optimization, we utilize the Adam [24] optimizer with $\beta_1=0.9,\,\beta_2=0.99,\,$ and $\varepsilon=1\times10^{-15}.$ We decay our learning rate logarithmically from 5×10^{-3} to 5×10^{-4} over 25k training iterations with cosine-scheduled warmup in the first 500 steps.

A.2 Radiance Cues

Geometry To extract radiance cues we first optimize UniSDF [56] on the input images. After optimization, we convert the SDF representation to a mesh using marching cubes [31] with threshold set to be zero.

Rendering We use Blender Cycles [50], a physically-based path-tracer to render the radiance cues. We run Blender via the Kubric python wrapper [16], and we use the estimated geometry with the predefined materials based on the GGX material model [55], as described in Sec. 3.4.



(b) w/ smoothness.

Figure S1: Effects of shading normal smoothing function.

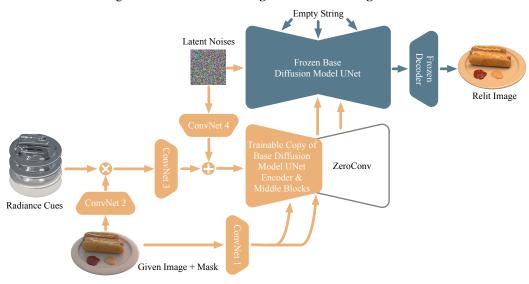


Figure S2: Schematics of our ControlNet-based diffusion model.

Shading Normals In order to produce smoothly-varying specular highlights which look realistic, we need the normals used for shading to be smooth. By default, Blender computes normals for shading based on the input geometry, which may be noisy. To mitigate this, we can feed the predicted normals \mathbf{n}' described in Sec. A.1 to Blender and enable its shading normal smoothing function which applies to the predicted normals, and uses them for shading. However, over-smoothness may harm the photorealism of the rendered shadows. See Fig. S1 for qualitative comparison on radiance cues rendered without enabling the shading normal smoothing (Fig. S1a) and with the feature enabled (Fig. S1b). In our implementation, we exploit a hybrid strategy: we utilize radiance cues without smoothness for the diffuse material and use radiance cues with smoothness for the specular materials. Concretely, our final radiance cues are composed of the first rendering in Fig. S1a and the right three ones in Fig. S1b.

A.3 Relighting Diffusion Model

We implement our relighting diffusion model in JAX [11]. We illustrate the architecture of the model for inference in Fig. S2. We build upon a text-to-image latent diffusion model which is similar to the model of Rombach *et al.* [45]. It denoises gaussian noise of size $64 \times 64 \times 8$ and decodes the output latent features into a relit image of size $512 \times 512 \times 3$. The model was not conditioned on text input, receiving only empty strings via a CLIP text encoder [42]. During training the base model is frozen.

Following ControlNet [63], we create a trainable copy of the base diffusion model's UNet encoder and middle blocks and append them with a ZeroConv-based blocks to the frozen base model. The

Table S1: Fig. S2's ConvNet 1 Structure. Convolution layer's definition is represented as (kernel size, stride, padding). We use SiLU [19] as the activation function between layers. Layer 8 uses zero initialization while the other layers use Flax's [18] default initialization³. In our implementation, we have H = W = 512.

Index	Layer	Output Shape
0 (input)	-	$H \times W \times 4$
1	(3, 1, 1)	$H \times W \times 16$
2-1	(3, 1, 1)	$H \times W \times 16$
2-2	(3, 2, 1)	$H/2 \times W/2 \times 32$
3-1	(3, 1, 1)	$H/2 \times W/2 \times 32$
3-2	(3, 2, 1)	$H/4 \times W/4 \times 64$
4-1	(3, 1, 1)	$H/4 \times W/4 \times 64$
4-2	(3, 2, 1)	$H/8 \times W/8 \times 128$
5-1	(3, 1, 1)	$H/8 \times W/8 \times 128$
5-2	(3, 2, 1)	$H/16 \times W/16 \times 256$
6-1	(3, 1, 1)	$H/16 \times W/16 \times 256$
6-2	(3, 2, 1)	$H/32 \times W/32 \times 512$
7-1	(3, 1, 1)	$H/32 \times W/32 \times 512$
7-2	(3, 2, 1)	$H/64 \times W/64 \times 512$
8	(3, 1, 1)	$H/64 \times W/64 \times 1024$
9	flatten	$(H/64 \times W/64) \times 1024$

Table S2: **Fig. S2's ConvNet 3 Structure.** Convolution layer's definition is represented as (kernel size, stride, padding). We use SiLU [19] as the activation function between layers. Layer 5 uses zero initialization while the other layers uses Flax [18] default initialization³. In our implementation, we have H = W = 512.

Index	Layer	Output Shape
0 (input)	-	$H \times W \times 12$
1	(3, 1, 1)	$H \times W \times 16$
2-1	(3, 1, 1)	$H \times W \times 16$
2-2	(3, 2, 1)	$H/2 \times W/2 \times 32$
3-1	(3, 1, 1)	$H/2 \times W/2 \times 32$
3-2	(3, 2, 1)	$H/4 \times W/4 \times 96$
4-1	(3, 1, 1)	$H/4 \times W/4 \times 96$
4-2	(3, 2, 1)	$H/8 \times W/8 \times 256$
5	(3, 1, 1)	$H/8 \times W/8 \times 320$

given masked image and radiance cues are first fed through ConvNet 2 (see Fig. 4 in [61] for details) and ConvNet 3 (see Tab. S2). The resulting output is added to the output of the latent noise, which is fed through ConvNet 4. ConvNet 4 consists of a single convolution layer with kernel size 3, stride 1, padding 1, and 320 output channels. Given that the trainable copy was designed for tokenized text input, the masked image is first fed through ConvNet 1 (see Tab. S1) to generate representative embeddings. To ensure compatibility between the output of ConvNet 1 (size 64) and the CLIP [42] encoder's text output shape, zero-valued tensors are appended, increasing the size to 77.

We train the diffusion model using an approach similar to ControlNet [63], with a large dataset of synthetic objects rendered under multiple lighting conditions. Each training example for fine-tuning consists of a pair of images that view the same object with the same camera parameters, illuminated by two different environment map (see Sec. A.4). We fine-tune the diffusion model to predict one of these two images, given the other image as well as the corresponding radiance cues rendered using the synthetic object's geometry. Note that for synthetic objects, we do not need to estimate the geometry G nor to enable the Blender normal smoothing function to compute the radiance cues since we already have the ground-truth meshes and the normals from synthetic objects are smooth enough. We fine-tune the base model for 150k steps using batch size of 512 examples and a learning rate of 10^{-4} , which is linearly warmed up from 0 over the first 1k steps. The fine-tuning takes around 2 days on 32 TPUv5 chips. Besides, we always use the empty string as the text input to effectively make the fine-tuned model image-based.

At inference time, we use the DPPM scheduler [21] without classifier-free guidance [20] to produce samples at 512×512 resolution.

A.4 Training Data Processing

We use Objaverse [13] as the synthetic dataset. To filter out low-quality objects, we use the list from [49] to get our initial set of 156,330 ones.⁴ By additionally removing (semi-)transparent ones, we have a final set of 152,649 objects. If the object only contains geometry, we manually assign a homogeneous texture (ShaderNodeBsdfDiffuse) with a color uniformly sampled from $[0,1]^3$. Further, if the object does not have the material information, we assign it a Blender Glossy BSDF

³https://github.com/google/flax/blob/144486b5fa7b3dfb/flax/core/nn/linear.py#L27

⁴https://github.com/ashawkey/objaverse_filter/tree/dc9e7cd0df8626f30df02bb



Figure S3: Stanford-ORB's per-image illuminations from teapot_scene002: inconsistent sun and tripod shape/location.

Table S3: Issues on illuminations of Stanford-ORB [27]. We create two sets of reference renderings with the ground-truth geometry and material: 1) using a fixed illumination to render all evaluation images for a same (object, lighting) pair; and 2) using the **per-image** illumination provided by the benchmark. We then evaluate each approach's renderings with respect to the two sets of reference renderings respectively. We also list each approach's illumination selection in the second column. For each row, better performance between the two evaluations is highlighted. Apparently and consistently, the numerical results favor matched illumination selection.

	Illumination	Reference w/ Fixed Illumination				Reference w/ Per-image Illumination			
	Selection	PSNR-H↑	PSNR-L↑	SSIM↑	LPIPS↓	PSNR-H↑	PSNR-L↑	SSIM↑	LPIPS↓
PhySG [62]	Per-image	22.55	28.05	0.959	0.056	22.71	28.19	0.959	0.055
NVDiffRec [38]	Per-image	23.47	29.35	0.960	0.037	23.71	29.60	0.960	0.037
NeRD [10]	Per-image	24.05	30.20	0.968	0.053	24.22	30.36	0.969	0.053
NeRFactor [65]	Per-image	24.38	30.70	0.970	0.049	24.55	30.85	0.970	0.048
InvRender [66]	Per-image	24.50	30.75	0.970	0.047	24.68	30.93	0.971	0.046
NVDiffRecMC [17]	Per-image	25.17	31.19	0.970	0.037	25.45	31.53	0.970	0.036
Ours	Fixed	26.29	32.45	0.973	0.029	26.00	32.11	0.973	0.029
Ours (single GPU)	Fixed	26.34	32.53	0.974	0.029	26.05	32.17	0.973	0.029
Real Images	_	26.25	32.69	0.975	0.024	26.73	33.27	0.977	0.023

material (ShaderNodeBsdfGlossy), whose roughness value is uniformly sampled from [0.02, 0.5] and base color is set to be the same as the homogeneous texture. The mixing factor between the specular and diffuse materials (ShaderNodeMixShader) is uniformally sampled from [0,1].

As we discussed in Sec. A.3, our diffusion training requires image pairs under different lightings. For this, we select 509 equirectangular environment maps from [60]. For each object, we sample four camera poses on a sphere centered around it. For each camera, we randomly sample two environment maps and augment them with random horizontal shift, vertical flip, and RGB channel shuffle. We then use Blender's Cycle path tracer to render an image of resolution 512×512 with 512 samples per pixel for each environment map using a camera whose focal length is set to be 512.

B Stanford-ORB Illumination Issues

Stanford-ORB provides estimated **per-image** illumination via moving a light probe for each image, see Fig. S3 for an example. Ideally, fixed illumination per object would match reality, but aligning the object and light probe is challenging. This limitation in the Stanford-ORB benchmark can significantly affect results, especially in areas with specular highlights as demonstrated in Tab. S3. Consequently, there is no "correct" way to do relighting in Stanford-ORB: our results use fixed illumination, while competitors use per-image illumination.

We rendered the ground truth geometry and materials, obtained by Stanford-ORB in a controlled studio environment (see [27]'s Sec.3.2.2). Each view was rendered under **fixed** illumination (consistent environment map) and **per-image** illumination (unique environment map per view). Fig. S4 shows both renderings alongside the corresponding real image. Note the significant variations, especially in areas with specular highlights (see marked regions and PSNR).



Figure S4: For each object, from left to right, we show 1) reference rendering w/ **fixed** illumination; 2) reference rendering w/ **per-image** illumination" (see Tab. S3 caption for details); and 3) the real captured image. We show PSNR-L between 1) vs. 3) and 2) vs. 3) respectively. Different illumination settings vary significantly.

We computed metrics for each method using both reference renderings as ground truth whose quantative results are in Tab. S3. Our method excels under fixed illumination but performs worse with per-image illumination. Competitors show the opposite trend, doing better with per-image illumination. Neural-PBIR did not release code or complete results, thus it is missing from the table.

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 precede 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]

The claims about performance are verified by an extensive quantitative and qualitative analysis as described in the experiments section.

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 the limitations of our method in the limitations section. Most notably our method bakes in material and target lighting into the NeRF, *i.e.* for each target light our method needs to optimize a new NeRF model.

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: We do not have any theoretical results.

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 attempted to provide all necessary information to reproduce the results. We provide detailed insights into dataset and radiance cue generation, model architectures, hyperparameters and model training in the supplemental.

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 have not made the code or model weights available online, however, the Objaverse dataset is publicly available as well as the datasets required for the Stanford-ORB and TensoIR benchmarks.

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: We provide information about model training, hyperparameters, optimizer and learning rates in the appendix. For the evaluation we use the data splits defined by the benchmarks.

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: We evaluate our method against the most popular benchmarks which do not 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: For both diffusion model training and NeRF optimization we provide the compute requirements in the appendix.

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: To the best of our knowledge we follow 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: [No]

Justification: We propose a novel paradigm for an existing task — 3D relighting. If reviewers believe our method could cause potential harm, we are happy to include a statement.

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: [No]

Justification: We do not plan to release our trained models.

Guidelines:

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

12. Licenses for existing assets

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

Answer: [Yes]

Justification: We cite the works that provide the assets used for data set generation, as well as the use of pre-trained model weights in the appendix.

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: [Yes]

Justification: Yes, we explain how we render 3D assets in the appendix section.

Guidelines:

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

14. Crowdsourcing and Research with Human Subjects

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

Answer: [NA]

Justification: We did not conduct research with humans or crowdsourced any tasks.

Guidelines:

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

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

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

Answer: [NA]

Justification: No studies were undertaken.

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.