

DANTE-W: Diffuse Albedo Neural Texturing in the Wild

Guangyu Wang^{1*}, Tianheng Lu^{1*}, Ruqi Huang^{1✉}, and Lu Fang^{1✉}

¹Tsinghua University, Beijing 100084, China

Abstract. Classical mesh texturing techniques blend captured multi-view images directly, which inevitably suffer from baked-in shading and casted shadows that compromise visual fidelity during relighting. To circumvent this issue, we present a neural texturing framework, namely DANTE-W, to enable high-fidelity diffuse albedo texture recovery from unstructured image collections for large-scale, in-the-wild scenes, which integrates seamlessly with traditional 3D reconstruction pipelines. Given a reconstructed mesh and its surface parameterization, our method fuses view-space generative albedo priors into a coherent texture space via an expressive neural representation, while substantially enhancing fine-grained textural details through physically principled neural rendering. To comprehensively evaluate our method, we curate a benchmark dataset featuring diverse, fine-grained textures, comprising both real-world in-the-wild scenes and synthetic objects. Extensive experiments verify the effectiveness of our approach in reconstructing accurate albedo textures and boosting relighting fidelity. Project page: dante-wild.github.io.

Keywords: Texture mapping · Neural representation · Diffuse albedo estimation

1 Introduction

Reconstructing high-quality textured assets of real-world large-scale scenes plays a crucial role in domains of virtual reality, filmmaking, and gaming. However, delivering lifelike experiences with these in-the-wild digitizations remains challenging, as it necessitates high-fidelity and 3D-consistent rendering of fine-grained details under arbitrary lighting conditions. In particular, we identify and address the critical limitations of existing methods as follows.

First of all, existing 3D reconstruction paradigms lack generalizable intrinsic decomposition capability. While traditional image-based 3D reconstruction techniques [21, 41, 45, 57, 61] scale well to hundreds of unstructured input images and large-scale scenes, they fail to reason about intrinsic material properties beyond raw pixel observations – specifically, to disentangle surface reflectance (e.g., diffuse albedo) from illumination effects (i.e., irradiance) – leading to baked-in

[✉]Correspondence to: Lu Fang (fanglu@tsinghua.edu.cn, luvision.net), Ruqi Huang (ruqihuang@sz.tsinghua.edu.cn, rqhuang88.github.io).

*Authors contributed equally to this work.




Fig. 1: We present DANTE-W, a neural texturing framework for high-fidelity diffuse albedo recovery in the wild. Compared to vanilla mesh texturing with baked-in lighting effects (e.g., noon-time shading and strong roof-edge shadowing on this pavilion), our method effectively disentangles a 3D-consistent diffuse albedo texture with exceptional photorealism. Leveraging physically principled neural rendering, DANTE-W faithfully reconstructs fine-grained albedo details, enabling hyper-realistic free-view relighting.

shading and shadows that compromise the visual fidelity when varying the illumination. In pursuit of intrinsic decomposition, a wealth of neural inverse rendering approaches [10, 25, 29, 31, 37, 39, 50, 67] integrate physically-based rendering [12, 32] with neural representations [33, 47, 49, 63], which jointly reconstruct shape, illumination, and material by imposing learned priors [11, 42, 43, 68] or empirical regularizations [5, 15, 19, 22, 23, 40, 44, 56, 59, 65]. However, they fall short in generalizing to challenging in-the-wild scenarios featuring complex illumination and highly detailed texture, due to oversimplified rendering formulations and the inherent lighting-material ambiguity.

On the other end of the spectrum, recent 2D generation-based methods [38, 66] demonstrate emergent generalizable albedo reasoning capabilities, by inheriting strong priors from pretrained diffusion models [1, 2, 9, 55] and finetuning on high-quality, task-specific synthetic datasets [18, 54, 70]. However, these methods operate on a per-image or per-video-clip basis, leading to unstable predictions under drastic viewpoint variations while struggling to accurately represent fine details – primarily due to the lack of spatial awareness, the probabilistic nature of generative modelling, and the information loss within VAE. Moreover, these diffusion-based renderers remain incompatible with modern graphics pipelines and are inefficient for rendering. These limitations underscore the need to bridge these coarse diffusion clues with expressive 3D representations for robust, in-the-wild intrinsic decomposition.

In light of the above observations, we present DANTE-W, a neural texturing framework seamlessly integrating with traditional 3D mesh reconstruction pipelines [21, 45, 57] for high-fidelity diffuse albedo recovery, with an emphasis on challenging in-the-wild scenarios characterized by: 1) large-scale, near-

Lambertian scenes with intricate, fine-grained textural details; 2) strong lighting variations involving complex shading and shadowing; 3) highly unstructured viewpoints spanning diverse observation distances.

To combine the best of both worlds, we fuse generalizable, view-space albedo priors – obtained from state-of-the-art diffusion models [38, 66] – into an expressive, 3D-consistent neural texture, thereby aggregating the view-dependent diffusion outputs to suppress flickering hallucinations. It is worth noting that, though, naively lifting per-view prior is prone to over-smoothed results, due to the inaccurate and view-inconsistent diffusion predictions at fine details.

To address this issue, we propose to learn accurate details from the original images and resolve the ambiguity through the lens of frequency. Our observation is that in-the-wild lighting effects tend to be smoother than fine-grained albedo variations, which is also supported by prior findings [53] that irradiance can be accurately approximated within 1% error using only nine spherical harmonic coefficients corresponding to the lowest-frequency modes. To this end, we disentangle high-frequency albedo components from irradiance by explicitly differentiating the frequency bands of their respective neural representations, thereby encouraging high-frequency information to flow into the albedo representation when neural rendering the raw imagery.

Given a mesh reconstructed and parameterized using off-the-shelf tools [45], our neural texturing framework operates efficiently – requiring only a few minutes for large-scale scenes with thousands of unconstrained multi-view images. The resulting neural texture can be directly baked into a standard UV-map and readily imported into any modern graphics engine (e.g., Blender [20]) for physically-based relighting, real-time rendering, and flexible editing.

To comprehensively benchmark our method, we curate a dataset namely *GigaLit*, containing both in-the-wild scenes and synthetic objects, each characterized by complex, high-resolution textures under diverse lighting conditions. Extensive evaluations demonstrate the state-of-the-art performance of DANTE-W – our method outperforms diffusion-based renderers [38, 66] by +2.63/3.85dB in PSNR and achieves $\sim 25\%$ lower LPIPS for diffuse albedo recovery, and by +6.13/7.15dB in PSNR and over 40% reduction in LPIPS for relighting, underscoring a substantial advancement towards unprecedented robustness and fidelity for in-the-wild relighting applications. To summarize, our main contributions are as follows:

- We present DANTE-W to recover high-fidelity diffuse albedo for in-the-wild scenes, by lifting view-space diffusion priors to a consistent neural texture representation.
- We propose a physically principled neural rendering framework to disentangle 3D-consistent, fine-grained albedo details from the original imagery via a frequency decomposition perspective.
- We introduce *GigaLit*, a benchmark dataset for in-the-wild diffuse albedo recovery and relighting, featuring unprecedented textural details, complex illumination variations, and unstructured viewpoints, where our method demonstrates significant improvements over prior works.

2 Related work

2.1 Classical mesh texturing

Traditional mesh texturing methods [8, 26, 41, 45, 61] take posed multi-view images and a parameterized mesh from multi-view stereo [21, 57] as inputs, and optimize a Markov random field (MRF) energy, where the data term encodes view selection strategies [26, 41, 61] and the smoothness term enforces texture seam consistency via various color adjustment schemes [35, 58]. Among them, Bi et al. [8] introduces a global patch-based optimization method to enable high-quality, aligned texture mapping from misaligned input images and noisy geometry. Ling et al. [41] uses parallelized Loopy Belief Propagation (LBP) to enable per-face view selection, which effectively mitigates non-linear illumination differences with a hierarchical blending scheme. However, these vanilla texturing methods directly operate on pixel colour, thus leaving complex shading and shadowing effects baked into the resulting texture. This limitation substantially degrades visual fidelity in downstream relighting applications.

2.2 Intrinsic decomposition

Intrinsic decomposition aims to separate lighting and underlying surface reflectance from a single image [7, 13, 14, 66] or multi-view image sequence [10, 29, 37, 38, 50]. The majority of methods leverage physically-based rendering [22, 25, 29, 39, 44, 50], generative material estimation [37, 43, 64], or a combination of both [17, 42] to jointly recover shape and reflectance for 3D relighting. Some recent works perform diffusion-based generative relighting [3, 28, 52, 60, 69] directly in the image domain without explicit intrinsic reasoning. Despite their effectiveness on small-scale, object-centric scenarios [18, 34] with highly structured capture setting, they fail to generalize to complex, large-scale scenes with unconstrained image collections. Another line of works [19, 31, 40] extends inverse rendering to in-the-wild scenarios by imposing tailored priors on sky [5, 23] or shadow and sunlight [5, 15, 56, 59]. However, they struggle to accurately estimate diffuse albedo that is fully disentangled from irradiance, as the imposed priors are insufficient to resolve the inherent lighting-material ambiguity.

On the other hand, recent advances [38, 66] demonstrate compelling generalization capabilities by finetuning general purpose image [55] or video generation models [1, 2] on high-quality synthetic data [18, 54, 70] for in-the-wild intrinsic decomposition. However, their diffusion outputs suffer from the loss of details and 3D-inconsistent hallucinations, thus failing to empower high-fidelity, free-viewpoint rendering. In this work, we propose to aggregate view-space albedo priors from these state-of-the-art diffusion-based renderers into a 3D-consistent neural texture, while leveraging physically-principled neural rendering to further disentangle high-frequency albedo details from the original imagery. Our method effectively eliminates the ambiguity and demonstrates superior expressivity for challenging in-the-wild scenarios. Unlike existing frameworks, we explicitly formulate diffuse albedo recovery within the mesh texturing phase to ensure seamless compatibility with well-established 3D reconstruction and graphics pipelines.

3 Method

In this section, we introduce DANTE-W, aiming to recover high-fidelity diffuse albedo texture for in-the-wild scenes. Given a collection of unstructured multi-view images and a mesh reconstructed and parameterized using off-the-shelf 3D reconstruction tools, we lift view-dependent, diffusion-based albedo priors to a consistent and expressive 2D neural field defined on the UV-texture space. To enhance fine-grained, 3D-consistent albedo details, we present a physically principled neural rendering framework that disentangles albedo from irradiance by differentiating the frequency bands of their neural representations while jointly optimizing for neural rendering loss and albedo regularization. The resulting neural texture can then be baked into standard texture map and seamlessly integrated with modern graphics engines for relighting.

In the following, we first briefly introduce our shading formulation and neural encoding module (Sec. 3.1). We then elaborate on our neural texture representation for diffuse albedo (Sec. 3.2) and our neural rendering framework to recover fine albedo details (Sec. 3.3).

3.1 Preliminaries

Diffuse shading model. In this work, we focus on in-the-wild scenes with approximately Lambertian surfaces, consistent with the assumptions commonly adopted in traditional 3D reconstruction literature [21, 57, 61]. Our goal, however, is to take a step further by decomposing the observed colour $\mathbf{c} \in \mathbb{R}^3$ into its diffuse albedo $\mathbf{a}_d \in \mathbb{R}^3$ and diffuse irradiance $\mathbf{s}_d \in \mathbb{R}^3$ components, according to the image formation model:

$$\mathbf{c} = \mathbf{a}_d \odot \mathbf{s}_d, \quad (1)$$

where the diffuse albedo $\mathbf{a}_d$ describes the intrinsic base color of dielectric opaque surfaces, and the diffuse irradiance $\mathbf{s}_d$ specifies the incident lighting at a surface position integrated over the upper cosine-weighted hemisphere [30], which contributes to diverse shading and shadowing effects. Therefore, texturing with $\mathbf{a}_d$ instead of $\mathbf{c}$ significantly improves visual fidelity under novel lighting conditions.

Multi-resolution hash encoding. Recent advances in neural fields [6, 36, 49] adopt a hybrid representation that combines multi-resolution hash encoding [49] with a tiny multilayer perceptron neural network (MLP) decoder to faithfully capture fine details. Specifically, hash encoding allocates a hierarchy of independent feature grids across multiple spatial resolutions, where each grid vertex is mapped to a hash entry storing a learnable high-dimensional feature vector. Let $\{V_\ell\}_{\ell=1}^L$ be the set of spatial grid resolutions, and given an input position $\mathbf{x}$ and a resolution V_ℓ , we denote by $\psi_\ell(\mathbf{x}) \in \mathbb{R}^Z$ the Z -dimensional hash feature obtained via linear interpolation of hash entries at the vertices of the grid cell enclosing $\mathbf{x}$. The hash features across all resolutions are concatenated as:

$$\psi(\mathbf{x}) = (\psi_1(\mathbf{x}), \dots, \psi_L(\mathbf{x})) \in \mathbb{R}^{LZ}, \quad (2)$$

which are subsequently fed into the lightweight MLP to decode the desired scene attributes.




Fig. 2: An overview of our physically principled neural rendering framework. Given a reconstructed mesh of the scene with surface parameterization, we represent diffuse albedo texture $\mathbf{a}_d$ using a high-resolution 2D hash encoding $\psi_a^{2D}(\cdot)$ and irradiance $\mathbf{s}_d$ using a low-resolution 3D hash encoding $\psi_s^{3D}(\cdot)$. This explicit frequency-band discrepancy (with the maximal grid resolutions satisfying $V_{L_a}^{(a)} \gg V_{L_s}^{(s)}$) effectively facilitates the disentanglement between the two intrinsic components. We guide the low-frequency component of diffuse albedo with view-space diffusion priors (Eq. (4)) and recover fine-grained albedo details by neural rendering the raw observations (Eq. (6)).

3.2 Albedo representation

Hash-encoded neural texture. Given a reconstructed mesh with its surface parameterization, we represent the diffuse albedo texture in terms of an expressive neural field defined in the UV-coordinate system, as shown in Fig. 3.

Specifically, we apply multi-resolution hash encoding [49] in the 2D texture space with a hierarchy of spatial grid resolutions $\{V_\ell^{(a)}\}_{\ell=1}^{L_a}$. We then use a lightweight MLP, $\mathcal{F}_{\Theta_a}$, to interpret the hash features and decode the diffuse albedo value. Let $\hat{x} \in \mathbb{R}^2$ be the parametric UV-coordinate of a surface position, we denote by $\psi_a^{2D}(\hat{x}) \in \mathbb{R}^{L_a Z_a}$ the concatenated hash features across all L_a resolutions queried at $\hat{x}$, then the diffuse albedo $\mathbf{a}_d \in \mathbb{R}^3$ is modelled as:

$$\mathbf{a}_d = \mathcal{F}_{\Theta_a}(\psi_a^{2D}(\hat{x})). \quad (3)$$




Fig. 3: An illustration of the proposed hash-encoded neural texture representation.

Benefited from the expressiveness of multi-resolution hash encoding, our neural albedo texture effectively represents high-resolution, intricate details in a memory-efficient manner. Moreover, by organizing a coarse-to-fine hierarchy of spatial grid resolutions, our neural texture inherently mitigates texture seams without the need for any tailored designs [35].

Lifting view-space albedo priors. Since disentangling diffuse albedo from raw observations is a highly ill-posed problem, we leverage state-of-the-art intrinsic decomposition diffusion models (e.g., RGB $\leftrightarrow$ X [66] or DiffusionRenderer [38]) to impose a strong guidance for albedo estimation. To achieve this, we first run the pretrained diffusion model on input images $\{\mathcal{I}_k\}$ to generate the albedo buffer

$\{\mathcal{A}'_k\}$. We then distill our neural texture using these view-space albedo priors via rasterization and neural rendering:

$$\mathcal{L}_a = \sum_k \sum_{\mathbf{r}} \|\mathbf{a}_d - \mathcal{A}'_k(\mathbf{r})\|. \quad (4)$$

To elaborate, for each pixel ray $\mathbf{r}$ from the k -th image, we leverage rasterization to obtain the UV-coordinate $\hat{\mathbf{x}}$ of the corresponding ray-surface intersection and calculate the diffuse albedo value $\mathbf{a}_d$ using Eq. (3). We jointly optimize the learnable hash entries and the MLP decoder by minimizing the discrepancies between our albedo texture $\mathbf{a}_d$ and the corresponding screen-space albedo value $\mathcal{A}'_k(\mathbf{r})$ generated by the pretrained diffusion model.

Remaining issues. We find naively lifting view-dependent diffusion outputs (Eq. (4)) leads to suboptimal results, with the distilled albedo texture exhibiting excessive blurriness, as shown in Fig. 4, 6, 8, 7 (w/o PR). We attribute this degradation to two primary factors: 1) the 2D generative framework fails to maintain 3D consistency, particularly for highly intricate details; 2) the VAE struggles to reconstruct fine-grained details accurately due to the inherent information loss of the compressed latent space. These limitations motivate us to further recover high-fidelity albedo details via physically-based neural rendering.

3.3 Recovering fine-grained albedo texture

Leveraging diffusion priors as strong guidance for the low-frequency component of diffuse albedo, we propose to further learn 3D consistent, high-frequency albedo details from the raw input images through a physically principled neural rendering framework. Observing that fine albedo details generally exhibit higher spatial frequencies than irradiance for in-the-wild scenes, we can mitigate the albedo-irradiance ambiguity by exploiting their frequency discrepancy within the underlying neural representations.

Lighting representation. Unlike diffuse albedo, which is defined only on the scene surface and can be effectively represented in the UV texture space, light interacts intricately with the scene geometry in 3D space through multiple bounces, yielding shading and shadows shaped by the full set of participating scene structures within the light transport. Therefore, we model global illumination as a volumetric neural field to enable implicit light transport reasoning over the complete 3D scene geometry.

To elaborate, we leverage hash encoding to represent lighting, with an explicit control over its frequency band by setting the maximal grid resolution. Let $\mathbf{x} \in \mathbb{R}^3$ be the 3D coordinate of a surface position, we denote by $\{V_\ell^{(s)}\}_{\ell=1}^{L_s}$ the set of 3D feature grid resolutions and by $\psi_s^{3D}(\mathbf{x}) \in \mathbb{R}^{L_s Z_s}$ the concatenated lighting features across all L_s grid resolutions. To account for unconstrained lighting variations, we introduce a per-image learnable embedding $\mathbf{e}_k \in \mathbb{R}^E$ and concatenate it with the multi-resolution lighting features [46, 62]. Let $\mathcal{F}_{\Theta_s}$ be a lightweight MLP head, the irradiance $\mathbf{s}_d \in \mathbb{R}^3$ is calculated as:

$$\mathbf{s}_d = \mathcal{F}_{\Theta_s}(\psi_s^{3D}(\mathbf{x}), \mathbf{e}_k). \quad (5)$$

Physically principled neural rendering. An overview of our full pipeline is illustrated in Fig. 2, where we represent diffuse albedo $\mathbf{a}_d$ and irradiance $\mathbf{s}_d$ respectively using Eq. (3) and Eq. (5), and we compose the final RGB colour according to Eq. (1). Crucially, we disambiguate albedo and irradiance by 1) imposing diffusion-based albedo guidance (Eq. (4)) and 2) differentiating their maximal grid resolutions, i.e., $V_{L_a}^{(a)} \gg V_{L_s}^{(s)}$. This design enables fine-grained details to flow into the diffuse albedo component when recreating the raw input image:

$$\mathcal{L}_c = \sum_k \sum_{\mathbf{r}} \|\gamma(\mathbf{c}) - \mathcal{I}_k(\mathbf{r})\|, \quad (6)$$

where we denote by $\mathcal{I}_k(\mathbf{r})$ the raw RGB value at pixel location $\mathbf{r}$ of the k -th input image, and by $\gamma(\cdot)$ a fixed tone mapping function that converts linear color to sRGB space [4].

Our full loss function involves the neural rendering loss (Eq. (6)) and the diffusion-based albedo regularization term (Eq. (4)), which are jointly optimized with respect to the learned hash entries and the MLP parameters:

$$\mathcal{L} = \lambda \mathcal{L}_a + \mathcal{L}_c, \quad (7)$$

where λ denotes the relative weight of the albedo regularization and we empirically set it to 1.0 for all scenes in our implementation.

3.4 Implementation details

Exposure compensation. To compensate for unconstrained camera response, we introduce a pair of learnable scalars $w_k, b_k \in \mathbb{R}$ for each input image $\mathcal{I}_k$ to rectify the computed RGB color $\mathbf{c}$ [16].

Texture map baking. Our hash-encoded neural texture can be readily transformed to standard texture map, making our framework integrate seamlessly with well-established graphics pipelines. To generate the texture map, we first initialize a UV-map at the desired texture resolution, with each pixel storing its 2D UV-coordinate $\hat{\mathbf{x}}$. This UV-map is then passed through our neural albedo renderer (Eq. (3)) to produce the diffuse albedo map.

Ray batching. We cache the rasterization buffer for all input views prior to training, which stores, for each pixel, the 3D and UV coordinates of the ray-surface intersection and a binary foreground indicator. During training, we randomly sample a set of foreground pixel rays across all views for each gradient descent step. We empirically find this strategy leads to slightly better performance and faster convergence.

4 Experiments

4.1 GigaLit benchmark

Existing benchmarks [18, 29, 34] for intrinsic decomposition and relighting are limited to highly constrained, object-centric scenarios characterized by: 1) simple

albedo textures with minimal fine-grained details; 2) static illumination represented by an environment map lacking complex variations in shading and shadowing; 3) evenly distributed viewpoints centered around the object at a fixed observation scale.

In light of these limitations, we introduce *GigaLit*, a new benchmark dataset designed to characterize the complexity and richness of in-the-wild scenes for diffuse albedo recovery and relighting. Specifically, GigaLit contains 6 real-world outdoor scenes for qualitative evaluation, featuring intricate details, strong shading, and complex shadowing effects. For each scene, we capture hundreds to thousands of unstructured, high-quality images across multiple observation distances and different times of the day using a DSLR camera. We then run Metashape Pro [45] for structure-from-motion, multi-view stereo, surface reconstruction and parameterization.

To support quantitative evaluation, we also meticulously curate 10 synthetic objects with fine albedo details and 4K resolution texture from BlenderKit [51]. We leverage Blender [20] to simulate challenging lighting variations by altering environment maps collected from PolyHeaven [24] and directional spot lights. To mimic the unstructured nature of in-the-wild data, we render each synthetic object along an upper hemispherical trajectory with varying camera-to-object distances. We refer readers to the supplement for additional dataset details.

Baselines. We test our method with two state-of-the-art intrinsic decomposition models – RGB $\leftrightarrow$ X [66] and Cosmos-DiffusionRenderer (DR) [38] – both supporting forward and inverse rendering via diffusion-based inference. We compare against their raw diffusion outputs, vanilla mesh texturing methods [8,45,61], and LightSwitch [43], which performs material estimation and generative relighting via multi-view diffusion. We also compare with the ablated variants of our framework without physically principled neural rendering ([Ours+RGB $\leftrightarrow$ X/DR] (w/o PR)) throughout all experiments. For mesh-based methods, we represent diffuse albedo as standard UV texture map and perform relighting using Blender’s Cycles engine [27].

Experimental setup. For real-captured, in-the-wild scenes, we qualitatively compare all baselines by recovering diffuse albedo texture under the varied illuminations during data collection and relighting with novel environment maps from PolyHeaven. For synthetic data, we pair each object with 5 different environment maps to simulate the unconstrained lighting for albedo recovery, where each input view of the object is rendered using a randomly selected one. We hold out another 5 environment maps shared across objects for relighting evaluation. We use PSNR, SSIM and LPIPS as quantitative metrics, and report the performance as the average over 10 objects, 250 views, and 5 lighting conditions.

4.2 Comparative results

Results on in-the-wild scenes. The visual comparisons on GigaLit in-the-wild scenes are presented in Fig. 4. As observed, vanilla mesh texturing methods like Metashape Pro [45] fail to disentangle shading and shadowing from the underlying diffuse albedo, leaving these irradiance components baked into the UV




Fig. 4: Visual comparison of diffuse albedo recovery on in-the-wild scenes. By lifting view-space diffusion priors upon a coherent neural texture via physically principled neural rendering (PR), our method robustly eliminates baked-in shading and shadowing effects that remain challenging for vanilla mesh texturing (Metashape Pro [45]), while refining raw diffusion outputs of intrinsic decomposition models (RGB $\leftrightarrow$ X [66] and Cosmos-DiffusionRenderer (DR) [38]) with faithful, fine-grained, 3D-consistent details.



Fig. 5: Generalization on sharp shadows and strong specularities using GigaLit scenes.

texture. For recent diffusion-based renderers such as Cosmos-DiffusionRenderer (DR) [38] and RGB $\leftrightarrow$ X [66], their raw diffusion outputs primarily lack accuracy, stability and 3D consistency across varying viewpoints, particularly at fine details. By contrast, our method enables strictly 3D-consistent and more faithful diffuse albedo recovery. As shown in Fig. 5, our method also generalizes well to more complex scenarios, e.g., sharp shadows and strong reflections, due to the robustness of meshing and our 3D-consistent neural texture. We refer readers to the supplement for more qualitative results.

Results on synthetic data. In Tab. 1, we report the mean metrics on the GigaLit synthetic subset. Our method significantly outperforms all baselines in terms of albedo recovery, yielding +4.49dB PSNR over Metashape Pro [45], +3.85dB PSNR and $\sim 25\%$ reduction in LPIPS relative to RGB $\leftrightarrow$ X [66], and +2.63dB PSNR and $\sim 24\%$ lower LPIPS relative to Cosmos-DiffusionRenderer (DR) [38]. For the task of relighting, our method also demonstrate strong performance, achieving +3.94dB PSNR over Metashape Pro [45], +7.15dB PSNR and $\sim 43\%$ lower LPIPS relative to RGB $\leftrightarrow$ X, and +6.13dB PSNR and $\sim 45\%$ reduction in LPIPS relative to DiffusionRenderer. Note that LightSwitch [43] struggles to generalize to complex, varied illumination and unstructured viewpoints, yielding significant hallucinations in both albedo and relighting. The qualitative comparisons are presented in Fig. 6. Under diverse, strong lighting conditions, diffusion-based renderers [38, 66] struggle in physically-accurate relighting, whereas our method demonstrates exceptional robustness by recovering faithful and detailed albedo texture and relighting with physically-based ray tracing [27] and well-established graphics engines [20].

Results on Stanford-ORB dataset. We also conduct experiments on the public Stanford-ORB dataset [34], and the visual results of diffuse albedo recovery are presented in Fig. 7. Compared to raw diffusion outputs from DiffusionRenderer (DR) [38] and RGB $\leftrightarrow$ X [66], our method demonstrates significant improvements in recovering 3D-consistent, fine-grained details while maintaining the robustness in irradiance removal – zoom in to see our high-quality details of the rendered text on the object surface. We also remark that the ground-truth albedo annotations provided by the Stanford-ORB dataset suffer from low capture quality and significant approximation errors, which exhibits pronounced noise and considerable degradation of fine textural details – see the last column ‘Dataset GT’ of Fig. 7. Since there are no precise diffuse albedo measurements

Table 1: Quantitative evaluations on GigaLit synthetic dataset.

	Diffuse Albedo			Relighting		
	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$
Waechter et al. [61]	21.43	0.889	0.103	23.78	0.915	0.092
Bi et al. [8]	21.48	0.885	0.111	23.81	0.902	0.097
Metashape Pro [45]	21.50	0.892	0.103	23.83	0.917	0.089
LightSwitch [43]	16.89	0.807	0.185	17.79	0.808	0.189
RGB $\leftrightarrow$ X [66]	21.65	0.860	0.125	20.30	0.849	0.134
DiffusionRenderer (DR) [38]	23.36	0.879	0.120	21.64	0.873	0.130
[Ours+RGB $\leftrightarrow$ X] (w/o PR)	25.32	0.890	0.120	27.27	0.918	0.099
[Ours+RGB$\leftrightarrow$X] (Full)	25.50	0.911	0.093	27.45	0.933	0.076
[Ours+DR] (w/o PR)	25.69	0.907	0.115	27.45	0.928	0.092
[Ours+DR] (Full)	25.99	0.921	0.091	27.77	0.938	0.071



Fig. 6: Qualitative results of diffuse albedo recovery and relighting on GigaLit synthetic objects, compared with vanilla texturing (Metashape Pro [45]), raw outputs of RGB $\leftrightarrow$ X [66] and Cosmos-DiffusionRenderer (DR) [38], and the ablated variant of our method in terms of physically principled neural rendering (PR). The albedo is recovered under the *Original Lit* and subsequently relighted under the *Novel Lit*. Our method effectively removes strong lighting effects and faithfully preserves fine details, while achieving superior relighting fidelity by capitalizing on the standard UV texture and physically-based ray tracing [27].



Fig. 7: Qualitative comparisons on diffuse albedo recovery using Stanford-ORB dataset [34]. The raw diffusion outputs from both DiffusionRenderer (DR) [38] and RGB $\leftrightarrow$ X [66] tend to exhibit hallucinated artifacts caused by the probabilistic nature of generative modelling and the information loss of the VAE latent. Naively lifting view-space priors without our physically principled neural rendering (Ours+DR/RGB $\leftrightarrow$ X (w/o PR)) also leads to severe blurriness, since the diffusion model fails to reason about 3D consistency and yields view-dependent flickering artifacts. By contrast, our method (Ours+DR/RGB $\leftrightarrow$ X) accurately recovers finer details while robustly eliminates strong irradiance effects. Note that the albedo ground truth in Stanford-ORB (Dataset GT) contains noisy capture artifacts that significantly deteriorate fine-grained textures, making it unsuitable for quantitative benchmarking of high-fidelity albedo recovery. Please zoom in to see our high-quality details.




Fig. 8: Visualizations of diffuse albedo recovery and relighting on a large-scale, in-the-wild scene – The Pavilion of Prince Teng – reconstructed using over 2,500 unstructured photographs. Our method learns a consistent diffuse albedo texture with significantly finer details compared to the naive lifting of diffusion priors (w/o PR). We also demonstrate superior fidelity compared to the vanilla texturing approach, Metashape Pro [45], which is prone to baked-in irradiance. Please zoom in to see details and lighting effects.




Fig. 9: Ablation study on frequency discrepancy. Differentiating the maximal grid resolution facilitates the disentanglement.



Fig. 10: Run time comparison with Metashape Pro [45] in terms of mesh texturing. Our method demonstrates comparable visual quality while achieving over $7\times$ acceleration.

for real-world scenes, the ground-truth albedo in Stanford-ORB serves only as a coarse reference, failing to provide a principled and rigorous quantitative evaluation for the task of high-fidelity diffuse albedo recovery, particularly for highly intricate textural details. By contrast, our method effectively learns accurate, fine-grained details that exhibit comparable visual fidelity to those in the original input images, significantly outperforming the quality of the ground truth annotations. This limitation motivates us to curate high-quality synthetic data in our GigaLit dataset for more reasonable quantitative evaluations.

Ablation studies. We ablate the proposed physically principled neural rendering (PR) framework and the comparative results (denoted as (w/o PR)) are shown in Fig. 4, 6, 7, 8, and Tab. 1. This design effectively compensates for the loss of high-frequency albedo details when lifting inconsistent screen-space albedo priors. We also ablate the frequency discrepancy design in terms of our albedo and lighting representation, and the visual results are presented in Fig. 9. Without explicitly differentiating the frequency bands of the underlying neural representations, the neural rendering paradigm suffers from ambiguity that tends to attribute fine details to the irradiance component.

Efficiency report. Leveraging advanced neural encoding implementations [48], our diffuse albedo neural texturing framework is highly efficient, with the entire optimization generally taking 5~8 minutes per scene on a single Nvidia RTX 3090 GPU – over $7\times$ faster than the multi-view mesh texturing process of Metashape Pro [45], while maintaining the comparable visual quality, as shown in Fig. 10. In terms of storage, our method consumes 0.8GB of GPU memory for object-centric synthetic scenes (maximal hash grid resolution $V_{L_a}^{(a)} = 2048$), and 3.3GB for in-the-wild scenes ($V_{L_a}^{(a)} = 8192$).

5 Conclusion

We introduce DANTE-W, the first high-fidelity diffuse albedo neural texturing approach for complex, in-the-wild scenes, taking as input unconstrained image collections. The core idea is to lift inconsistent view-space diffusion priors into a coherent and expressive neural texture representation, while incentivizing 3D-consistent, highly-intricate albedo details by exploiting frequency-band discrepancy within a physically principled neural rendering framework. Extensive

experiments demonstrate the superiority of our method in robustly alleviating strong irradiance effects and achieving unparalleled relighting fidelity.

Limitation & Future Work. Despite the effectiveness in lifting view-dependent, diffusion-based albedo priors with accurate, 3D-consistent details, our current method is designed as a mesh texturing step within traditional 3D reconstruction pipeline and therefore does not account for non-Lambertian effects. One future research direction is to exploit expressive neural texture in advanced scene representations (e.g., Gaussian Splatting [33]) and include specular materials into the proposed framework.

Acknowledgements

This work is supported in part by Natural Science Foundation of China (NSFC) under contract No. 62125106 and 62427804, in part by Tsinghua University Dushi Program (No.20251080107), in part by the Beijing Outstanding Young Scientist Program under contract No. JWZQ20240101009, in part by the XPLOER PRIZE.

References

1. Agarwal, N., Ali, A., Bala, M., Balaji, Y., Barker, E., Cai, T., Chattopadhyay, P., Chen, Y., Cui, Y., Ding, Y., et al.: Cosmos world foundation model platform for physical ai. arXiv preprint arXiv:2501.03575 (2025)
2. Alhaija, H.A., Alvarez, J., Bala, M., Cai, T., Cao, T., Cha, L., Chen, J., Chen, M., Ferroni, F., Fidler, S., et al.: Cosmos-transfer1: Conditional world generation with adaptive multimodal control. arXiv preprint arXiv:2503.14492 (2025)
3. Alzayer, H., Henzler, P., Barron, J.T., Huang, J.B., Srinivasan, P.P., Verbin, D.: Generative multiview relighting for 3d reconstruction under extreme illumination variation. In: Proceedings of the Computer Vision and Pattern Recognition Conference. pp. 10933–10942 (2025)
4. Anderson, M., Motta, R., Chandrasekar, S., Stokes, M.: Proposal for a standard default color space for the internet—srgb. In: Color and imaging conference. vol. 4, pp. 238–245. Society of Imaging Science and Technology (1996)
5. Bai, H., Zhu, J., Jiang, S., Huang, W., Lu, T., Li, Y., Guo, J., Fu, R., Guo, Y., Chen, L.: Gare: Relightable 3d gaussian splatting for outdoor scenes from unconstrained photo collections. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 26456–26465 (2025)
6. Barron, J.T., Mildenhall, B., Verbin, D., Srinivasan, P.P., Hedman, P.: Zip-nerf: Anti-aliased grid-based neural radiance fields. ICCV (2023)
7. Bell, S., Bala, K., Snavely, N.: Intrinsic images in the wild. ACM Transactions on Graphics (TOG) **33**(4), 1–12 (2014)
8. Bi, S., Kalantari, N.K., Ramamoorthi, R.: Patch-based optimization for image-based texture mapping. ACM Trans. Graph. **36**(4), 106–1 (2017)
9. Blattmann, A., Dockhorn, T., Kulal, S., Mendelevitch, D., Kilian, M., Lorenz, D., Levi, Y., English, Z., Voleti, V., Letts, A., et al.: Stable video diffusion: Scaling latent video diffusion models to large datasets. arXiv preprint arXiv:2311.15127 (2023)

10. Boss, M., Braun, R., Jampani, V., Barron, J.T., Liu, C., Lensch, H.: Nerd: Neural reflectance decomposition from image collections. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 12684–12694 (2021)
11. Boss, M., Jampani, V., Braun, R., Liu, C., Barron, J., Lensch, H.: Neural-pil: Neural pre-integrated lighting for reflectance decomposition. *Advances in Neural Information Processing Systems* **34**, 10691–10704 (2021)
12. Burley, B., Studios, W.D.A.: Physically-based shading at disney. In: *Acm siggraph*. vol. 2012, pp. 1–7. vol. 2012 (2012)
13. Careaga, C., Aksoy, Y.: Colorful diffuse intrinsic image decomposition in the wild. *ACM Transactions on Graphics (TOG)* **43**(6), 1–12 (2024)
14. Careaga, C., Aksoy, Y.: Intrinsic image decomposition via ordinal shading. *ACM Trans. Graph.* **43**(1) (2023)
15. Chang, Y., Kim, Y., Seo, S., Yi, J., Kwak, N.: Fast sun-aligned outdoor scene relighting based on tensorf. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. pp. 3626–3636 (2024)
16. Chen, D., Li, H., Ye, W., Wang, Y., Xie, W., Zhai, S., Wang, N., Liu, H., Bao, H., Zhang, G.: Pgsr: Planar-based gaussian splatting for efficient and high-fidelity surface reconstruction. *IEEE Transactions on Visualization and Computer Graphics* (2024)
17. Chen, X., Peng, S., Yang, D., Liu, Y., Pan, B., Lv, C., Zhou, X.: Intrinsicanything: Learning diffusion priors for inverse rendering under unknown illumination. In: *European Conference on Computer Vision*. pp. 450–467. Springer (2024)
18. Deitke, M., Schwenk, D., Salvador, J., Weihs, L., Michel, O., VanderBilt, E., Schmidt, L., Ehsani, K., Kembhavi, A., Farhadi, A.: Objaverse: A universe of annotated 3d objects. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 13142–13153 (2023)
19. Feng, W., Ye, K., Zhang, Q., Zhang, Q., Li, N.: 2d gaussian splatting for outdoor scene decomposition and relighting. In: *Proceedings of the Thirty-Fourth International Joint Conference on Artificial Intelligence*. pp. 990–998 (2025)
20. Foundation, T.B.: Blender 3.6 (2023)
21. Furukawa, Y., Hernández, C., et al.: Multi-view stereo: A tutorial. *Foundations and trends® in Computer Graphics and Vision* **9**(1-2), 1–148 (2015)
22. Gao, J., Gu, C., Lin, Y., Li, Z., Zhu, H., Cao, X., Zhang, L., Yao, Y.: Relightable 3d gaussians: Realistic point cloud relighting with brdf decomposition and ray tracing. In: *European Conference on Computer Vision*. pp. 73–89. Springer (2024)
23. Gardner, J.A., Kashin, E., Egger, B., Smith, W.A.: The sky’s the limit: Relightable outdoor scenes via a sky-pixel constrained illumination prior and outside-in visibility. In: *European conference on computer vision*. pp. 126–143. Springer (2024)
24. Greg, Z., Rob, T., Rico, C., James, R.C., Dario, B., Jenelle, v.H., et al.: Polyhaven: a curated public asset library for visual effects artists and game designers (2021)
25. Hasselgren, J., Hofmann, N., Munkberg, J.: Shape, light, and material decomposition from images using monte carlo rendering and denoising. *Advances in Neural Information Processing Systems* **35**, 22856–22869 (2022)
26. Huang, X., Zhu, Q., Jiang, W.: Gpvc: Graphics pipeline-based visibility classification for texture reconstruction. *Remote Sensing* **10**(11), 1725 (2018)
27. Iraci, B.: Blender cycles: lighting and rendering cookbook. Packt Publishing Ltd (2013)
28. Jin, H., Li, Y., Luan, F., Xiangli, Y., Bi, S., Zhang, K., Xu, Z., Sun, J., Snavely, N.: Neural gaffer: Relighting any object via diffusion. *Advances in Neural Information Processing Systems* **37**, 141129–141152 (2024)

29. Jin, H., Liu, I., Xu, P., Zhang, X., Han, S., Bi, S., Zhou, X., Xu, Z., Su, H.: Tensor: Tensorial inverse rendering. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 165–174 (2023)
30. Kajiya, J.T.: The rendering equation. In: Proceedings of the 13th annual conference on Computer graphics and interactive techniques. pp. 143–150 (1986)
31. Kaleta, J., Kania, K., Trzciński, T., Kowalski, M.: Lumigauss: relightable gaussian splatting in the wild. In: 2025 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV). pp. 1–10. IEEE (2025)
32. Karis, B., Games, E.: Real shading in unreal engine 4. Proc. Physically Based Shading Theory Practice **4**(3), 1 (2013)
33. Kerbl, B., Kopanas, G., Leimkühler, T., Drettakis, G.: 3d gaussian splatting for real-time radiance field rendering. ACM Trans. Graph. **42**(4), 139–1 (2023)
34. Kuang, Z., Zhang, Y., Yu, H.X., Agarwala, S., Wu, E., Wu, J., et al.: Stanford-orb: a real-world 3d object inverse rendering benchmark. In: Advances in Neural Information Processing Systems Datasets and Benchmarks Track (2023)
35. Lempitsky, V., Ivanov, D.: Seamless mosaicing of image-based texture maps. In: 2007 IEEE conference on computer vision and pattern recognition. pp. 1–6. IEEE (2007)
36. Li, Z., Müller, T., Evans, A., Taylor, R.H., Unberath, M., Liu, M.Y., Lin, C.H.: Neuralangelo: High-fidelity neural surface reconstruction. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 8456–8465 (2023)
37. Li, Z., Wu, T., Tan, J., Zhang, M., Wang, J., Lin, D.: Idarb: Intrinsic decomposition for arbitrary number of input views and illuminations. arXiv preprint arXiv:2412.12083 (2024)
38. Liang, R., Gojcic, Z., Ling, H., Munkberg, J., Hasselgren, J., Lin, Z.H., Gao, J., Keller, A., Vijaykumar, N., Fidler, S., Wang, Z.: Diffusionrenderer: Neural inverse and forward rendering with video diffusion models. In: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (June 2025)
39. Liang, Z., Zhang, Q., Feng, Y., Shan, Y., Jia, K.: Gs-ir: 3d gaussian splatting for inverse rendering. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 21644–21653 (2024)
40. Liao, L., Zhang, C., Wu, T., Lv, H., Deng, B., Gao, L.: Rosgs: Relightable outdoor scenes with gaussian splatting. arXiv preprint arXiv:2509.11275 (2025)
41. Ling, X., Qin, R.: Large-scale and efficient texture mapping algorithm via loop belief propagation. IEEE transactions on geoscience and remote sensing **61**, 1–11 (2023)
42. Litman, Y., Patashnik, O., Deng, K., Agrawal, A., Zawar, R., De la Torre, F., Tulsiani, S.: Materialfusion: Enhancing inverse rendering with material diffusion priors. In: 2025 International Conference on 3D Vision (3DV). pp. 802–812. IEEE (2025)
43. Litman, Y., De la Torre, F., Tulsiani, S.: Lightswitch: Multi-view relighting with material-guided diffusion. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 27750–27759 (2025)
44. Liu, Y., Wang, P., Lin, C., Long, X., Wang, J., Liu, L., Komura, T., Wang, W.: Nero: Neural geometry and brdf reconstruction of reflective objects from multiview images. ACM Transactions on Graphics (ToG) **42**(4), 1–22 (2023)
45. LLC, A.: Agisoft metashape professional (2023)
46. Martin-Brualla, R., Radwan, N., Sajjadi, M.S., Barron, J.T., Dosovitskiy, A., Duckworth, D.: Nerf in the wild: Neural radiance fields for unconstrained photo collec-

- tions. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 7210–7219 (2021)
47. Mildenhall, B., Srinivasan, P.P., Tancik, M., Barron, J.T., Ramamoorthi, R., Ng, R.: Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM* **65**(1), 99–106 (2021)
 48. Müller, T.: tiny-cuda-nn (4 2021)
 49. Müller, T., Evans, A., Schied, C., Keller, A.: Instant neural graphics primitives with a multiresolution hash encoding. *ACM Transactions on Graphics (ToG)* **41**(4), 1–15 (2022)
 50. Munkberg, J., Hasselgren, J., Shen, T., Gao, J., Chen, W., Evans, A., Müller, T., Fidler, S.: Extracting Triangular 3D Models, Materials, and Lighting From Images. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 8280–8290 (June 2022)
 51. Petr, D., Vilém, D., Karolína, H., Monika, R., Adam, K., et al.: Blenderkit: Evergrowing 3d asset library integrated into blender (2018)
 52. Poirier-Ginter, Y., Gauthier, A., Phillip, J., Lalonde, J.F., Drettakis, G.: A diffusion approach to radiance field relighting using multi-illumination synthesis. In: *Computer Graphics Forum*. vol. 43, p. e15147. Wiley Online Library (2024)
 53. Ramamoorthi, R., Hanrahan, P.: An efficient representation for irradiance environment maps. In: Proceedings of the 28th annual conference on Computer graphics and interactive techniques. pp. 497–500 (2001)
 54. Roberts, M., Ramapuram, J., Ranjan, A., Kumar, A., Bautista, M.A., Paczan, N., Webb, R., Susskind, J.M.: Hypersim: A photorealistic synthetic dataset for holistic indoor scene understanding. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 10912–10922 (2021)
 55. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 10684–10695 (2022)
 56. Rudnev, V., Elgharib, M., Smith, W., Liu, L., Golyanik, V., Theobalt, C.: Nerf for outdoor scene relighting. In: *European Conference on Computer Vision (ECCV)* (2022)
 57. Schönberger, J.L., Zheng, E., Frahm, J.M., Pollefeys, M.: Pixelwise view selection for unstructured multi-view stereo. In: *Computer Vision–ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part III 14*. pp. 501–518. Springer (2016)
 58. Shen, T., Wang, J., Fang, T., Zhu, S., Quan, L.: Color correction for image-based modeling in the large. In: *Asian Conference on Computer Vision*. pp. 392–407. Springer (2016)
 59. Sun, J.M., Wu, T., Yang, Y.L., Lai, Y.K., Gao, L.: Sol-nerf: Sunlight modeling for outdoor scene decomposition and relighting. In: *SIGGRAPH Asia 2023 Conference Papers (SA Conference Papers '23)* (2023)
 60. Tang, J., Levine, M.J., Verbin, D., Garbin, S.J., Nießner, M., Brualla, R.M., Srinivasan, P.P., Henzler, P.: Rogr: Relightable 3d objects using generative relighting. In: *The Thirty-ninth Annual Conference on Neural Information Processing Systems* (2025)
 61. Waechter, M., Moehrle, N., Goesele, M.: Let there be color! large-scale texturing of 3d reconstructions. In: *European conference on computer vision*. pp. 836–850. Springer (2014)
 62. Wang, G., Zhang, J., Wang, F., Huang, R., Fang, L.: Xscale-nvs: Cross-scale novel view synthesis with hash featurized manifold. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 21029–21039 (2024)

63. Wang, P., Liu, L., Liu, Y., Theobalt, C., Komura, T., Wang, W.: Neus: learning neural implicit surfaces by volume rendering for multi-view reconstruction. In: Proceedings of the 35th International Conference on Neural Information Processing Systems. pp. 27171–27183 (2021)
64. Wu, X., Zhu, P., Lyu, J., Liu, X., Guo, J., Guo, Y., Xu, W., Lyu, C.: Stableintrinsic: Detail-preserving one-step diffusion model for multi-view material estimation. arXiv preprint arXiv:2508.19789 (2025)
65. Yao, Y., Zhang, J., Liu, J., Qu, Y., Fang, T., McKinnon, D., Tsin, Y., Quan, L.: Neif: Neural incident light field for physically-based material estimation. In: European conference on computer vision. pp. 700–716. Springer (2022)
66. Zeng, Z., Deschaintre, V., Georgiev, I., Hold-Geoffroy, Y., Hu, Y., Luan, F., Yan, L.Q., Havs.an, M.: Rgbx: Image decomposition and synthesis using material- and lighting-aware diffusion models. ACM SIGGRAPH 2024 Conference Papers (2024), <https://api.semanticscholar.org/CorpusID:269484178>
67. Zhang, K., Luan, F., Wang, Q., Bala, K., Snavely, N.: Physg: Inverse rendering with spherical gaussians for physics-based material editing and relighting. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 5453–5462 (2021)
68. Zhang, X., Srinivasan, P.P., Deng, B., Debevec, P., Freeman, W.T., Barron, J.T.: Nerfactor: Neural factorization of shape and reflectance under an unknown illumination. ACM Transactions on Graphics (ToG) **40**(6), 1–18 (2021)
69. Zhao, X., Srinivasan, P., Verbin, D., Park, K., Martin Brualla, R., Henzler, P.: Illuminerf: 3d relighting without inverse rendering. Advances in Neural Information Processing Systems **37**, 42593–42617 (2024)
70. Zhu, J., Luan, F., Huo, Y., Lin, Z., Zhong, Z., Xi, D., Wang, R., Bao, H., Zheng, J., Tang, R.: Learning-based inverse rendering of complex indoor scenes with differentiable monte carlo raytracing. In: SIGGRAPH Asia 2022 Conference Papers. pp. 1–8 (2022)