

Diffusion-Based Material Regularization for Physics-Based Inverse Rendering

Jingwang Ling¹, Lifan Wu², Feng Xu³, and Shuang Zhao¹(✉)

¹ University of Illinois Urbana-Champaign ² NVIDIA


³ BNRist and School of Software, Tsinghua University

{ling23,shzhao}@illinois.edu, lifanw@nvidia.com, xufeng2003@gmail.com

Stanford-ORB cup_scene006



Stanford-ORB pepsi_scene004



Input Relight-GT Ours Neural-PBIR MaterialFusion

Fig. 1: Qualitative comparison of relighting on the Stanford-ORB dataset [21] between our method, Neural-PBIR [38], and MaterialFusion [26]. With our new material clustering regularizer, we avoid baked-in shadows while accurately modeling spatially varying materials (top). Our method is also more robust to strong highlights on glossy metallic surfaces, producing more accurate reflections (bottom).

Abstract. Reconstructing physics-based 3D assets—geometry, materials, and illumination—from multi-view images is a core problem in computer graphics and vision, and a prerequisite for realistic relighting and editing. Physics-based inverse rendering offers an accurate image-formation model, but is severely underconstrained: without strong priors, illumination is baked into materials, and reconstructions generalize poorly to novel views and lighting. Data-driven diffusion models, in contrast, predict visually plausible materials, yet their predictions rarely satisfy the rendering equation and are not directly usable for physics-based rendering. We bridge these two paradigms rather than replacing either. Our key idea is to treat the predictions of a state-of-the-art diffusion model not as target material values but as a similarity kernel for optimization: we introduce a regularization loss that penalizes deviations

in the optimized material over surface regions where the diffusion predictions are near-constant, while leaving the optimization free to match the input images. Built on this regularizer, our end-to-end pipeline jointly reconstructs geometry, materials, and illumination, yielding high-quality assets that drop into standard rendering pipelines and relight faithfully. On the Synthetic4Relight, Stanford-ORB, and DTC-Synthetic datasets, our method significantly outperforms state-of-the-art baselines in both reconstruction accuracy and relighting quality.

Keywords: Inverse rendering · Material estimation · Diffusion models · Relighting

1 Introduction

Physics-based 3D assets, including object geometry, material properties, and illumination conditions, are essential to a wide array of applications in virtual/augmented reality, e-commerce, and entertainment. Consequently, reconstructing these assets from captured images is a core problem in computer graphics and vision.

To this end, a common approach is *inverse rendering*, also known as *analysis by synthesis*. By formulating reconstruction as an optimization problem, these techniques iteratively refine asset parameters, typically via some variant of stochastic gradient descent, to minimize the discrepancy between input images and their corresponding renderings. Unfortunately, inverse-rendering optimizations are generally underconstrained, especially for material reconstruction. Without proper priors or regularization, the reconstructed assets often generalize poorly to novel viewing or lighting conditions (e.g., suffering from severe “baking” artifacts), significantly limiting their usefulness.

Recently, *data-driven* methods have emerged as a powerful alternative. Trained on large quantities of real and synthetic image data, these methods directly infer physics-based parameters without the need for optimization. However, despite being visually plausible and often free of baking artifacts, the reconstructions produced by these methods typically lack the physical accuracy required to generate renderings that faithfully match the input images.

In this paper, we introduce a technique that regularizes inverse-rendering-based material optimization using a state-of-the-art data-driven model: DiffusionRenderer [24]. Instead of simply initializing inverse-rendering optimizations with diffusion predictions, we exploit the spatial consistency encoded in the model’s predictions. Specifically, we devise a novel regularization loss that penalizes deviations in optimized material parameters over surface regions where the diffusion predictions are near-constant. Building on this regularization scheme, we develop an end-to-end inverse-rendering pipeline that reconstructs geometry, material, and illumination from multi-view images of an object.

Our method is designed to bridge the gap between data-driven material prediction and physics-based inverse rendering, rather than replacing either component alone. Diffusion-based predictors such as DiffusionRenderer [24] and

RGB $\leftrightarrow$ X [45] can produce perceptually convincing intrinsic maps, but these predictions are not guaranteed to satisfy the rendering equation or reproduce held-out relighting results. Conversely, vanilla physics-based optimization in Mitsuba [17] provides an accurate rendering model, but remains underconstrained under sparse views and unknown illumination, making it prone to material-lighting ambiguity and baking artifacts. The key distinction of our approach is to use diffusion-predicted G-buffers as a similarity kernel for physics-based optimization, rather than as fixed target values or a generic smoothness prior.

To demonstrate the effectiveness of our method, we show that it outperforms several state-of-the-art baselines [26, 38] on the Synthetic4Relight [48], Stanford-ORB [21], and DTC-Synthetic [11] datasets.

2 Related Work

Physics-based inverse rendering. Inverse rendering reconstructs scene parameters—geometry, material, and lighting—by matching renderings to observed images. Recent advances in physics-based differentiable rendering [49] enable gradient-based optimization of these parameters via Monte Carlo light transport, an approach known as analysis by synthesis. The problem is fundamentally ill-posed because of inherent ambiguities among these parameters [10, 33], so many systems add regularization to constrain the solution space [29, 36, 51]. Neural representations such as NeRFs [31] and Gaussian Splatting [20] achieve impressive reconstruction and novel view synthesis [19, 38, 39, 54], but extracting 3D assets compatible with standard graphics pipelines remains challenging. Flash Cache [3] and IRGS [14] improve reconstruction quality through more accurate light transport, using a fast radiance cache and inter-reflection-aware 2D Gaussian ray tracing, respectively. These advances are orthogonal to our focus: they target light-transport accuracy, whereas our regularization targets material-lighting ambiguity. Our work addresses surface-based inverse rendering, producing standard triangle meshes and physically based materials (e.g., the Disney BRDF) that are directly usable in traditional rendering engines.

Data-driven priors for inverse rendering. Data-driven methods have advanced material estimation and intrinsic decomposition. For material estimation, models such as RGB $\leftrightarrow$ X [45] and DiffusionRenderer [24] predict intrinsic G-buffers (albedo, roughness, metallic, normal) from single images or videos. For lighting, DiffusionLight [8, 35] and LuxDiT [25] use diffusion models to estimate environmental illumination. A growing line of work integrates such priors directly into the inverse-rendering optimization loop. MaterialFusion [26] and IntrinsicAnything [7] use diffusion priors to guide reconstruction, while VideoMat [32] extracts PBR materials with video diffusion models. MatMart [43] adopts a two-stage diffusion framework: it first predicts materials from observed views, then completes unobserved regions through prior-guided generation, using view-material cross-attention for multi-view consistency. However, relying solely on per-view predictions can introduce inconsistencies and bias. Rather than treating these predictions as ground truth, we use them as structural guidance to

regularize our physics-based optimization, making it robust to prediction errors from the data-driven models.

A complementary line of work learns to relight images or 3D representations directly, bypassing explicit PBR material estimation. Single-image methods [18, 44] and NeRF-distillation approaches [40, 50] produce compelling relighting but do not yield exportable PBR material maps, limiting downstream use in standard rendering pipelines. Single-image methods such as DiLightNet [44] further operate on a fixed input view and do not synthesize the held-out novel views required by multi-view relighting benchmarks such as Stanford-ORB. Feed-forward methods such as RelitLRM [47] offer fast inference but rely on fixed learned priors that may not generalize to all real-world appearances. LightSwitch [27] casts multi-view-consistent relighting as a denoising task, and Alzayer et al. [2] handle in-the-wild captures with large illumination variation. Like this line of work, our method leverages data-driven priors, but it combines them with physics-based optimization to improve the physical accuracy of the reconstructed material asset rather than directly synthesizing relit images.

Regularization of material parameters. Optimizing spatially varying materials from sparse views requires effective regularization to prevent overfitting and baking artifacts. A classic strategy is reflectance sharing [55], which assumes the scene contains a limited set of materials and shares observations across surface points with similar appearance. For example, Lensch et al. [23] use hard cluster assignments derived from the current optimization state, which can be unreliable early on when materials still contain baked artifacts. This idea has been formalized through basis BRDFs [9, 52] that enforce piecewise simplicity. Other approaches rely on explicit material segmentation [12, 37], which is itself difficult to obtain accurately. Sharma et al. [37] additionally learn a perceptual similarity metric for materials, but expose it only as point-wise similarity maps or binary masks, neither of which is readily applicable to inverse rendering. Our approach is inspired by cross-channel regularization techniques such as joint bilateral filtering (JBF) [34], which uses guidance buffers to refine a target image and can be implemented efficiently for high-dimensional data via permutohedral lattices [1]. In inverse rendering, JBF-like terms have been used to regularize specular parameters using albedo [29, 36], and Wiersma et al. [42] use a related similarity kernel to estimate uncertainty over BRDF parameters when geometry and illumination are known. We instead tackle the joint reconstruction of geometry, materials, and lighting under unknown illumination by constructing a stable similarity kernel from externally predicted G-buffers. This encourages regions with similar predicted G-buffers to converge to similar optimized parameters, constraining the solution space while still allowing physical corrections.

3 Method

Taking as input N multi-view images $\{I_i\}_{i=1}^N$ of a static object under unknown illumination, our goal is to reconstruct physics-based representations of the object’s shape, materials, and illumination, all directly compatible with standard




Fig. 2: Overview of our pipeline. From N multi-view images $\{I_i\}$ under unknown illumination, we (1) predict per-view intrinsic G-buffers $\mathbf{G} = [\mathbf{A}, \mathbf{R}, \mathbf{M}, \mathbf{N}]$ with a conditional diffusion model; (2) reconstruct a voxel-grid SDF by neural volume rendering, supervised by the predicted normals $\mathbf{N}$ ($\mathcal{L}_{\text{shape}}$); and (3) jointly optimize shape, spatially varying material, and an environment map by differentiable rendering, minimizing the photometric loss $\mathcal{L}_{\text{img}}$ and our material clustering regularizer $\mathcal{L}_{\text{mat}}$, with albedo regularized through the scale-agnostic transform ψ . The result is a renderer-ready PBR asset that relights faithfully under novel illumination. *Bottom:* the diffusion-predicted G-buffer $\mathbf{g} = [\mathbf{A}, \mathbf{R}, \mathbf{M}]$ guides joint bilateral filtering of the rendered G-buffer $\hat{\mathbf{g}}$. The G-buffers shown are from an early training iteration for illustration purposes and do not reflect the final reconstruction quality.

rendering pipelines. Specifically, we seek asset parameters that produce renderings closely resembling the input images. When multiple solutions fit the images equally well, we aim to obtain physically plausible and visually coherent reconstructions that minimize baked-in artifacts (when viewed under novel conditions).

To this end, we introduce a technique that operates in three stages: *preprocessing*, *neural shape reconstruction*, and *physics-based inverse rendering* (PBIR). Fig. 2 provides an overview. In the following, we detail each of these stages.

Preprocessing. For each view i , our technique calibrates the camera parameters and leverages a conditional diffusion model [24, 45] to predict intrinsic G-buffers $\mathbf{G}_i$ comprising the base color $\mathbf{A}_i$, roughness $\mathbf{R}_i$, metallic $\mathbf{M}_i$, and surface normal $\mathbf{N}_i$. These buffers will be used in the following stages to supervise the reconstruction of object shape and materials.

Neural surface reconstruction. With the per-view G-buffers obtained, we then reconstruct the shape of the object represented as a voxel-grid signed distance function (SDF) using neural volume rendering [31, 38, 41]. To further improve reconstruction quality, we incorporate a normal-supervision loss that encourages

the rendered normal buffer $\hat{\mathbf{N}}_i$ to match the diffusion-predicted normals $\mathbf{N}_i$:

$$\mathcal{L}_{\text{shape}} = \sum_{i=1}^N \sum_p H_\delta \left(1 - \hat{\mathbf{N}}_{i,p}^\top \mathbf{N}_{i,p} \right), \quad (1)$$

where $\hat{\mathbf{N}}_{i,p}$ and $\mathbf{N}_{i,p}$ denote, respectively, the p -th pixel of $\hat{\mathbf{N}}_i$ and $\mathbf{N}_i$. In addition, H_δ is the Huber penalty—which improves robustness to imperfect normal predictions—given by

$$H_\delta(t) = \begin{cases} \frac{1}{2}t^2, & t \leq \delta, \\ \delta(t - \frac{1}{2}\delta), & \text{otherwise,} \end{cases} \quad (2)$$

where $t \geq 0$, and δ is set to 0.03 to downweight angular differences larger than 15° .

In practice, the normal-supervision loss in Eq. (1) enhances surface details and helps mitigate concave shape artifacts—especially on glossy surfaces (see Fig. 7). With the SDF reconstructed, we extract a mesh using Marching Cubes [28] as the initial shape of the object.

We note that, given the calibrated camera parameters and a reconstructed object geometry, the diffusion-predicted G-buffers $\{\mathbf{G}_i\}$ could be directly back-projected onto the object surface to complete the material reconstruction. However, as we will demonstrate in Sec. 4.4, materials reconstructed using this strategy often fail to produce renderings that match the ground truth. Instead, we will use these G-buffers as priors to greatly improve the robustness of our PBIR stage.

Physics-based inverse rendering. With the initial mesh obtained, the last stage of our technique jointly optimizes the object’s shape, materials, and illumination to minimize:

$$\mathcal{L} = \mathcal{L}_{\text{img}} + \lambda_{\text{mat}} \mathcal{L}_{\text{mat}}, \quad (3)$$

where $\mathcal{L}_{\text{mat}}$ regularizes the material using the diffusion-predicted G-buffers $\{\mathbf{G}_i\}$ and λ_{mat} is the corresponding scalar weight. In addition, $\mathcal{L}_{\text{img}}$ is the image rendering loss measured in relative MSE [22, 30]:

$$\mathcal{L}_{\text{img}} = \sum_{i=1}^N \left\| \frac{\hat{I}_i - I_i}{\text{sg}(\hat{I}_i) + \epsilon} \right\|_2^2, \quad (4)$$

where $\hat{I}_i$ denotes the image rendered under the i -th view, $\text{sg}(\cdot)$ is the stop-gradient operator, and $\epsilon > 0$ is a small constant for stable normalization. Because our datasets contain HDR images, the relative loss prevents a small number of bright pixels from dominating the optimization.

In the following, we detail our material regularization strategy in Sec. 3.1 and describe our scale-agnostic albedo transformation in Sec. 3.2.

3.1 Implicit Material Clustering Regularization

Directly optimizing spatially varying materials from sparse multi-view observations under unknown illumination is severely underconstrained. Each surface location is observed from only a limited set of views and lighting conditions, often causing the optimization to bake illumination effects into the material. Recent conditional diffusion models can infer plausible per-view intrinsic G-buffers from RGB inputs [24, 45], producing results that are often artifact-free and semantically coherent. However, these predictions are not constrained by the rendering equation and thus lack physical accuracy, leading to suboptimal visual quality when used directly for rendering.

Recent methods have integrated diffusion-based material priors into physics-based inverse rendering [26, 32]. However, directly regularizing toward the predicted G-buffers creates a tension with photometric fitting; approaches based on score distillation with annealed diffusion noise [26, 53] or global scale-invariant losses [7, 32] do not fully address material-dependent biases in a scene with multiple materials. We instead leverage the implicit grouping of same-material regions found in diffusion predictions, and regularize our solution to remain consistent within these regions while allowing for per-region deviations. Akin to reflectance sharing [55], this effectively reduces the degrees of freedom by sharing observations across similar material regions.

A practical challenge is to identify which image regions correspond to the same material without committing to an explicit segmentation. One could fit a small set of BRDF bases with spatially sparse mixing (so that each region selects a single basis) [9, 52]. However, this introduces model selection issues (e.g., number of bases) and sensitivity to region boundaries. We instead use diffusion-predicted G-buffers to define an implicit, soft notion of material similarity across pixels as follows. For a given view, let $\mathbf{g} = [\mathbf{A}, \mathbf{R}, \mathbf{M}]$ denote the predicted material G-buffer, where each vector $\mathbf{g}_p$ is the concatenation of channels at pixel p . We define the similarity kernel between pixels p and q as

$$k_{p,q}(\mathbf{g}) = \exp\left(-\frac{\|\mathbf{g}_p - \mathbf{g}_q\|_2^2}{2\sigma_g^2}\right), \quad (5)$$

where σ_g controls the decay of similarity with respect to G-buffer distance.

Let $\hat{\mathbf{g}} = [\hat{\mathbf{A}}, \hat{\mathbf{R}}, \hat{\mathbf{M}}]$ be the differentially rendered counterpart of $\mathbf{g}$ for the same view. We encourage pixels that are similar in the predicted G-buffers to share similar optimized material by computing a kernel-weighted average via a joint bilateral filter:

$$\text{JBF}(\hat{\mathbf{g}}; \mathbf{g})_p = \frac{\sum_q k_{p,q}(\mathbf{g}) \hat{\mathbf{g}}_q}{\sum_q k_{p,q}(\mathbf{g})}, \quad (6)$$

where $\text{JBF}(\hat{\mathbf{g}}; \mathbf{g})$ denotes the filtered buffer, and subscript p indicates its value at pixel p . This filter can be implemented efficiently and differentially using permutohedral lattices [1]. The bottom panel of Fig. 2 illustrates these buffers during a training iteration.

In Eqs. (5) and (6), we use one kernel $k_{p,q}$ obtained using the concatenated $\mathbf{g} = [\mathbf{A}, \mathbf{R}, \mathbf{M}]$ to regularize all (i.e., albedo, roughness, and metallic) channels. This ensures the kernel considers differences in all channels and, therefore, reduces the risk of over-regularization (e.g., when some predicted G-buffer channels are overly smooth).

Finally, we define the per-view material regularization term as the L1 distance between the rendered and filtered G-buffers:

$$\mathcal{L}_{\text{mat}} = \|\hat{\mathbf{g}} - \text{JBF}(\hat{\mathbf{g}}; \mathbf{g})\|_1. \quad (7)$$

3.2 Scale-Agnostic Albedo Transformation

Applying material regularization directly to albedo under unknown illumination can bias optimization: the optimizer might scale down the albedo to reduce the regularization loss, while the lighting compensates to match the rendered appearance, resulting in overly intense light sources. To address this issue, we regularize albedo in a scale-agnostic space by transforming the rendered albedo as

$$\psi(\hat{\mathbf{A}}) = \text{sg}(\hat{\mathbf{A}}) \odot \log([\hat{\mathbf{A}}]_\epsilon), \quad (8)$$

where $\odot$ denotes element-wise multiplication, $\text{sg}(\cdot)$ is the stop-gradient operator, and $[a]_\epsilon = \max(a, \epsilon)$ is used for numerical stability with a small $\epsilon > 0$. The log mapping makes the region-wise average scale-agnostic: within a region, a multiplicative rescaling of the albedo becomes an additive offset in log space, which cancels out in the difference $\hat{\mathbf{g}} - \text{JBF}(\hat{\mathbf{g}}; \mathbf{g})$. However, the log function also rescales gradients by a factor proportional to $1/\hat{\mathbf{A}}$. Multiplying by the detached $\text{sg}(\hat{\mathbf{A}})$ cancels this gradient scaling and appropriately downweights dark regions. Ultimately, we replace $\hat{\mathbf{A}}$ in $\hat{\mathbf{g}}$ with $\psi(\hat{\mathbf{A}})$ when evaluating $\mathcal{L}_{\text{mat}}$.

Current diffusion models predict per-view G-buffers that may be inconsistent across views. Our regularization is relatively robust to such inconsistencies because they often manifest as per-region shifts or scalings, which are largely addressed by our design.

4 Experiments

After describing implementation details and datasets, we compare against prior inverse-rendering and relighting methods on Stanford-ORB, Synthetic4Relight, and DTC-Synthetic, explore alternative regularization choices, ablate key components, and provide additional relighting comparisons. The supplementary material provides additional kernel-design ablations, implementation details, the RGB $\leftrightarrow$ X upstream-model experiment, experiments applying our material regularizer to a Gaussian Splatting-based inverse-rendering pipeline [14], and extensive intrinsic G-buffer visualizations. The supplementary video shows dynamic relighting and visualizations difficult to assess from static figures.



Fig. 3: Qualitative comparison of relighting on the Stanford-ORB dataset [21] and DTC-Synthetic dataset [11] between our method, Neural-PBIR [38], and MaterialFusion [26].

4.1 Implementation Details

For the SDF stage, we largely follow [38] and add the normal regularization term with weight 10^{-5} . For the PBIR stage, we represent the shape as a mesh, the materials with the Disney principled BRDF [4] (optimizing base color, metallic, and roughness), and the lighting as an environment map. We parameterize the spatially varying BRDF parameters and the environment map using Dictionary Fields [5, 6], a neural field that shares a set of basis functions across locations (we use their 2D configuration). The environment map uses an exponential ac-



Relight-GT Ours N.-PBIR M. Fusion Relight-GT Ours N.-PBIR M. Fusion

Fig. 4: Qualitative comparison of relighting on the Synthetic4Relight dataset [48] between our method, Neural-PBIR [38], and MaterialFusion [26].

tivation. BRDF outputs use no activation; we clamp them to $[0, 1]$ and apply an L_1 penalty to out-of-range values with a weight of 10^{-2} .

We optimize using an initial learning rate of 3×10^{-2} with cosine annealing to 10^{-3} . We set $\lambda_{\text{mat}} = 0.1$, $\sigma_g = 0.02$ and $\epsilon = 10^{-2}$, and run 900 iterations. We initialize the BRDF and environment-map parameters to uniform values near 0.5. Our pipeline is implemented in Mitsuba 3 [17] and PyTorch. We use 256 samples per pixel for primal rendering and 64 samples per pixel for backward-mode automatic differentiation; each iteration renders 6 randomly sampled views. On a single RTX 5090 GPU, the SDF stage takes around 10 minutes, DiffusionRenderer preprocessing around 15 minutes, and the PBIR stage around 7 minutes per case.

4.2 Datasets and Metrics

We evaluate our method and baselines on a real-world benchmark, Stanford-ORB [21], and two synthetic datasets, Synthetic4Relight [48] and DTC-Synthetic. DTC-Synthetic contains seven scenes that we render using objects from DigitalTwinCatalog [11], designed to more thoroughly test reconstruction of glossy objects and strong cast shadows.

Stanford-ORB contains 42 scenes with captured relighting ground truth; we use the official open-source benchmark code for evaluation. For Synthetic4Relight, we follow the evaluation protocol of [48]. For DTC-Synthetic, we render 60 relit images under three novel environment maps from [48] (20 images per lighting condition) and report relighting metrics in the low dynamic range (LDR) domain. We report PSNR, SSIM, and LPIPS [46].

4.3 Comparison

We compare our method against the following approaches: *Neural-PBIR* [38], as a representative of purely analysis-by-synthesis inverse rendering that optimizes

Table 1: Quantitative results of our method and baselines on the Stanford-ORB dataset.

	Geometry			Novel Scene Relighting				Novel View Synthesis			
	Depth	Normal	Shape	PSNR-H	PSNR-L	SSIM	LPIPS	PSNR-H	PSNR-L	SSIM	LPIPS
	↓	↓	↓	↑	↑	↑	↓	↑	↑	↑	↓
Neural-PBIR	0.32	0.061	<u>0.42</u>	<u>26.07</u>	<u>33.39</u>	<u>0.980</u>	<u>0.023</u>	<u>28.90</u>	<u>36.92</u>	<u>0.987</u>	<u>0.019</u>
MaterialFusion	0.30	0.044	1.08	23.52	31.18	0.968	0.037	26.17	34.67	0.975	0.032
Ours	0.25	0.014	0.30	27.22	34.98	0.981	0.021	29.58	38.46	0.987	0.017

Table 2: Quantitative results of our method and baselines on the Synthetic4Relight and DTC-Synthetic datasets.

Method	Synthetic4Relight										DTC-Synthetic		
	Relighting			Aligned albedo			Albedo			Rough.	Relighting		
	PSNR	SSIM	LPIPS	PSNR	SSIM	LPIPS	PSNR	SSIM	LPIPS	MSE	PSNR	SSIM	LPIPS
	↑	↑	↓	↑	↑	↓	↑	↑	↓	↓	↑	↑	↓
Neural-PBIR	27.83	<u>0.974</u>	0.058	25.72	0.951	0.074	22.67	0.951	0.091	0.016	39.18	0.9928	0.0100
MaterialFusion	20.20	0.944	0.091	24.28	0.952	0.070	15.65	0.931	0.094	<u>0.015</u>	28.63	0.9746	0.0237
Ours	32.02	0.975	0.049	27.83	0.960	0.045	27.04	0.960	0.048	0.013	43.21	0.9961	0.0063
Ours (w/o reg.)	<u>28.94</u>	0.973	<u>0.055</u>	<u>26.63</u>	<u>0.956</u>	<u>0.065</u>	<u>23.79</u>	<u>0.954</u>	<u>0.084</u>	0.021	<u>40.23</u>	<u>0.9954</u>	<u>0.0065</u>

shape, spatially varying materials, and illumination without learned priors, and *MaterialFusion* [26], as a representative of diffusion-prior-aided inverse rendering via score distillation [53]. We use the official implementations and follow the recommended hyperparameters. We also test several previous regularization techniques on our pipeline; please see Sec. 4.4.

We show a qualitative comparison of novel-scene relighting on the Stanford-ORB and DTC-Synthetic datasets in Figs. 1 and 3.

Our method is robust to strong directional lighting in the input images, where hard cast shadows induce high-contrast appearance variation across an object. In this setting, purely analysis-by-synthesis optimization (e.g., Neural-PBIR) can absorb shadowing into the estimated materials, yielding baked artifacts that persist under relighting. By grouping observations within likely same-material regions, our regularization effectively reduces the degrees of freedom in spatially varying materials and discourages fitting cast shadows as texture; as a result, we remove baked shadows on the side of `car_scene002`, on the red cube in `Block_RedBlue`, and in the cast-shadow pattern around the cup handle in `cup_scene006`, which are typically visible in the Neural-PBIR results. This same-material grouping also improves material estimation. We recover more accurate roughness, as reflected by specular highlights that better match the reference relighting images in `ball_scene003`, `blocks_scene005`, on the flower pot of `cactus_scene005`, and on the top of `chips_scene003`. We also recover more accurate metallic parameters, reproducing the shiny metallic appearance during relighting on top of `baking_scene003`, in `pitcher_scene001` and `pepsi_scene004`, and on top of `TeaPot_EmeraldGoldTop`. With the help of our regularization on normals, our method further improves surface shape and mitigates concave-shape artifacts.




Fig. 5: Qualitative comparison of relighting between our method and alternative material regularization strategies, including directly back-projecting DiffusionRenderer predictions (Diffusion-BP), optimizing from this initialization without regularization (w/o reg.), and using a non-data-driven albedo-specular correlation regularizer (d-s corr.).




Fig. 6: Qualitative relighting comparison between our method and the global scale-invariant loss from VideoMat [32].



MaterialFusion is based on a different shape reconstruction pipeline [15], so shape quality is not directly comparable. It often eliminates baked-in shadows but tends to underfit the input images, failing to reconstruct spatially varying color patterns in `cup_scene006` and on the side of `baking_scene003`, yielding over-smoothed materials, less accurate relighting, and occasional color drift. We attribute these failures to the score-distillation target not always aligning with the input scene and the lack of an explicit mechanism to correct per-region inaccuracies, which can cause over-regularization.

We show a qualitative comparison of relighting on the Synthetic4Relight dataset in Fig. 4. Compared to Neural-PBIR, our method avoids baking on the `chair` back and the `hotdog` plate. Our method also achieves a more accurate overall color tone and produces more accurate reflections on `air_balloons` and the inside wall of `jugs`. MaterialFusion tends to cause color shifts in relighting on `air_balloons` and `jugs`.

We present the quantitative comparison results on Stanford-ORB in Tab. 1 and on Synthetic4Relight and DTC-Synthetic in Tab. 2.

Table 3: Stanford-ORB relighting comparison: (a) alternative regularization choices and (b) previous-method baselines.

Method	PSNR-H $\uparrow$	PSNR-L $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$
<i>(a) Alternative regularization choices</i>				
Diffusion-BP	26.06	33.09	0.976	0.0333
w/o reg.	26.11	33.78	0.978	0.0223
d-s corr.	26.37	34.10	0.979	0.0215
scale inv.	26.38	34.21	0.980	0.0225
Ours	27.22	34.98	0.981	0.0213
<i>(b) Comparison with previous methods</i>				
NVDiffRecMC [15]	24.43	31.60	0.972	0.0360
R3DG [13]	23.40	30.96	0.968	0.0348
Material Anything [16]	24.46	31.40	0.968	0.0387
Neural Gaffer [18]	24.32	31.46	0.963	0.0447
IllumiNeRF [50]	25.56	32.74	0.976	0.0270
Vanilla Mitsuba [17]	25.62	33.16	0.977	0.0232



Our method outperforms the baselines on relighting accuracy as well as albedo/roughness estimation and shape reconstruction. As Neural-PBIR is the current state of the art on Stanford-ORB, surpassing it sets a new state of the art on this dataset by a notable margin.

4.4 Exploring Alternative Choices

We evaluate alternative material-regularization designs against our implicit material clustering regularizer. As a first baseline, we directly back-project our predicted material maps onto our reconstructed shape. The resulting materials are visually plausible but do not match the relighting ground truth (Fig. 5), motivating an analysis-by-synthesis stage that enforces consistency between rendered and observed images. However, simply initializing the optimization with these maps and optimizing without our regularization quickly reintroduces baking artifacts, such as the cast-shadow pattern in `Block_RedBlue` and dotted specular highlights in `grogue_scene003` (the w/o reg case in Fig. 5), indicating that the initialization remains far from convergence and the optimization can fall into local minima that trade off image fit against material coherence.

We also compare against a global scale-invariant loss, in the spirit of VideoMat [32] and IntrinsicAnything [7], implemented following VideoMat with our predicted material maps as guidance (Fig. 6). A global adjustment is less effective at suppressing baked-in shadows (e.g., `cup_scene007`) and can underfit localized material details (e.g., the metallic foil lettering in `curry_scene001`), since it cannot accommodate region-specific corrections.

Finally, we evaluate a non-data-driven diffuse-specular self-correlation regularizer [29]. We construct the correlation signal from the current optimized base color and use it to regularize the metallic and roughness buffers. However, when the base color already contains baked artifacts, the regularizer can prop-




Fig. 9: Additional qualitative Stanford-ORB relighting results grouped by case. Each row shows the shared input and relighting target once, followed by previous-method baselines and our result.

agate these errors and become less effective (Fig. 5, denoted as d-s corr.), as evidenced by cast shadows in *Block_RedBlue* (top row) and dotted artifacts in *grogu_scene003* (bottom row).

The same trends hold quantitatively on Stanford-ORB: all alternative choices in Tab. 3(a) underperform our implicit material clustering regularizer in relighting accuracy.

4.5 Ablation Studies

We evaluate several components of our method in this section. In Fig. 7, we compare reconstructions with and without our regularization on normals. Without it, the reconstructed surface can exhibit concave artifacts, especially on glossy surfaces, which in turn introduce severe highlight-related artifacts. Adding the normal regularization yields a more faithful shape that better matches the reference. We ablate our scale-agnostic albedo transformation in Fig. 8. In the first row, we show the relighting reference and results, and in the second row, we show the reference and reconstructed environment map of the original lighting condition. Without the transformation, the optimization tends to scale down base color to reduce the regularization, which is compensated by a brighter and more spatially spread-out lighting estimate (second row), deviating from the reference lighting (left) and degrading reconstruction quality.

Finally, Tab. 2 reports a quantitative ablation on Synthetic4Relight and DTC-Synthetic. Overall, the improvement of our full system comes from three parts: our improved inverse-rendering pipeline, our regularization on normals, and our material regularization. The ablation isolates the last part: removing our material regularization consistently reduces relighting accuracy across both datasets.

4.6 Additional Relighting Comparisons

We further compare with additional relighting and inverse-rendering baselines on Stanford-ORB: NVDiffRecMC [15], R3DG [13], Material Anything [16], Neural Gaffer [18], IllumiNeRF [50], and a vanilla Mitsuba inverse-rendering baseline [17]. Unlike the main comparison with Neural-PBIR and MaterialFusion, these methods do not all reconstruct comparable shape and PBR assets; some directly perform relighting or use representations outside the mesh-based setting, so we report only held-out relighting quality.

As shown in Tab. 3(b) and Fig. 9, our method achieves the best relighting accuracy. The vanilla Mitsuba baseline shows that a physically based renderer alone is insufficient under sparse-view inverse rendering, where material-lighting ambiguities cause baked-in appearance. Unlike image-based relighting methods such as Neural Gaffer and IllumiNeRF, our method also produces relightable PBR assets while improving held-out relighting metrics.

5 Discussion and Conclusion

Limitations and future work. Our approach assumes DiffusionRenderer predicts material maps that are locally consistent within same-material regions. When this breaks down—e.g., for rare appearances outside the training distribution—the prior becomes unreliable and can over-regularize, blurring materials. DiffusionRenderer also operates at limited resolution, so matching its predictions leaves high-frequency details less constrained; higher-resolution diffusion models would address this. Finally, we derive our similarity kernel from a concatenation of material channels; alternative combinations may prove preferable as diffusion models evolve. Residual artifacts also remain near shadow boundaries, which we attribute to imperfect geometry; future work could refine shape using signals from our regularization to better align it with cast-shadow boundaries.

Conclusion. We presented an end-to-end pipeline that jointly recovers geometry, spatially varying materials, and illumination from multi-view images. Rather than replacing physics-based optimization with a data-driven predictor, we bridge the two: we treat DiffusionRenderer’s predictions not as target material values but as a similarity kernel, penalizing deviations in the optimized material over regions where the predictions are near-constant while still fitting the input images. This suppresses baking artifacts and yields relightable assets, significantly outperforming state-of-the-art baselines on the Synthetic4Relight, Stanford-ORB, and DTC-Synthetic datasets.

Acknowledgments

We thank the anonymous reviewers for their feedback. This work was partially supported by NSF grant 2553564. This work used NCSA Delta [award OAC 2005572] through ACCESS allocation CIS260003, supported by NSF grants #2138259, #2138286, #2138307, #2137603, and #2138296.

References

1. Adams, A., Baek, J., Davis, M.A.: Fast high-dimensional filtering using the permutohedral lattice. *Comput. Graph. Forum* **29**(2), 753–762 (2010). <https://doi.org/10.1111/J.1467-8659.2009.01645.X>
2. Alzayer, H., Henzler, P., Barron, J.T., Huang, J., Srinivasan, P.P., Verbin, D.: Generative multiview relighting for 3d reconstruction under extreme illumination variation. In: 2025 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2025, Nashville, TN, USA, June 11–15, 2025. pp. 10933–10942. Computer Vision Foundation / IEEE (2025). <https://doi.org/10.1109/CVPR52734.2025.01021>
3. Attal, B., Verbin, D., Mildenhall, B., Hedman, P., Barron, J.T., O’Toole, M., Srinivasan, P.P.: Flash cache: Reducing bias in radiance cache based inverse rendering. In: Computer Vision - ECCV 2024 - 18th European Conference, Milan, Italy, September 29–October 4, 2024, Proceedings, Part XXVIII. Lecture Notes in Computer Science, vol. 15086, pp. 20–36. Springer (2024). https://doi.org/10.1007/978-3-031-73390-1_2
4. Burley, B.: Physically-based shading at disney (2012)
5. Chen, A., Xu, Z., Wei, X., Tang, S., Su, H., Geiger, A.: Dictionary fields: Learning a neural basis decomposition. *ACM Trans. Graph.* **42**(4), 156:1–156:12 (2023). <https://doi.org/10.1145/3592135>
6. Chen, A., Xu, Z., Wei, X., Tang, S., Su, H., Geiger, A.: Factor fields: A unified framework for neural fields and beyond. *CoRR* **abs/2302.01226** (2023). <https://doi.org/10.48550/ARXIV.2302.01226>
7. 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: Leonardis, A., Ricci, E., Roth, S., Russakovsky, O., Sattler, T., Varol, G. (eds.) Computer Vision - ECCV 2024 - 18th European Conference, Milan, Italy, September 29–October 4, 2024, Proceedings, Part LX. Lecture Notes in Computer Science, vol. 15118, pp. 450–467. Springer (2024). https://doi.org/10.1007/978-3-031-73027-6_26
8. Chinchuthakun, W., Phongthawee, P., Raj, A., Jampani, V., Khungurn, P., Suwajanakorn, S.: Diffusionlight-turbo: Accelerated light probes for free via single-pass chrome ball inpainting. *CoRR* **abs/2507.01305** (2025). <https://doi.org/10.48550/ARXIV.2507.01305>
9. Chung, H., Choi, S., Baek, S.: Differentiable inverse rendering with interpretable basis brdfs. In: 2025 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2025, Nashville, TN, USA, June 11–15, 2025. pp. 475–484. Computer Vision Foundation / IEEE (2025). <https://doi.org/10.1109/CVPR52734.2025.00053>
10. Chung, H., Kim, J., Kim, S., Ye, J.C.: Parallel diffusion models of operator and image for blind inverse problems. In: 2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2023, Vancouver, BC, Canada, June 17–24, 2023. pp. 6059–6069. IEEE (2023). <https://doi.org/10.1109/CVPR52729.2023.00587>
11. Dong, Z., Chen, K., Lv, Z., Yu, H., Zhang, Y., Zhang, C., Zhu, Y., Tian, S., Li, Z., Moffatt, G., Christofferson, S., Fort, J., Pan, X., Yan, M., Wu, J., Ren, C.Y., Newcombe, R.A.: Digital twin catalog: A large-scale photorealistic 3d object digital twin dataset. In: 2025 IEEE/CVF Conference on Computer Vision and

- Pattern Recognition, CVPR 2025, Nashville, TN, USA, June 11-15, 2025. pp. 753–763. Computer Vision Foundation / IEEE (2025). <https://doi.org/10.1109/CVPR52734.2025.00079>
12. Fischer, M., Georgiev, I., Groueix, T., Kim, V.G., Ritschel, T., Deschaintre, V.: Sama: Material-aware 3d selection and segmentation. CoRR **abs/2411.19322** (2024). <https://doi.org/10.48550/ARXIV.2411.19322>
 13. 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: Leonardis, A., Ricci, E., Roth, S., Russakovsky, O., Sattler, T., Varol, G. (eds.) Computer Vision - ECCV 2024 - 18th European Conference, Milan, Italy, September 29-October 4, 2024, Proceedings, Part XLV. Lecture Notes in Computer Science, vol. 15103, pp. 73–89. Springer (2024). https://doi.org/10.1007/978-3-031-72995-9_5
 14. Gu, C., Wei, X., Zeng, Z., Yao, Y., Zhang, L.: IRGS: inter-reflective gaussian splatting with 2d gaussian ray tracing. In: 2025 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2025, Nashville, TN, USA, June 11-15, 2025. pp. 10943–10952. Computer Vision Foundation / IEEE (2025). <https://doi.org/10.1109/CVPR52734.2025.01022>
 15. Hasselgren, J., Hofmann, N., Munkberg, J.: Shape, light, and material decomposition from images using monte carlo rendering and denoising. In: Koyejo, S., Mohamed, S., Agarwal, A., Belgrave, D., Cho, K., Oh, A. (eds.) Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2022, NeurIPS 2022, New Orleans, LA, USA, November 28-December 9, 2022 (2022)
 16. Huang, X., Wang, T., Liu, Z., Wang, Q.: Material anything: Generating materials for any 3d object via diffusion. In: 2025 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2025, Nashville, TN, USA, June 11-15, 2025. pp. 26556–26565. Computer Vision Foundation / IEEE (2025). <https://doi.org/10.1109/CVPR52734.2025.02473>
 17. Jakob, W., Speierer, S., Roussel, N., Vicini, D.: DR.JIT: a just-in-time compiler for differentiable rendering. ACM Trans. Graph. **41**(4), 124:1–124:19 (2022). <https://doi.org/10.1145/3528223.3530099>
 18. 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. In: Advances in Neural Information Processing Systems 38: Annual Conference on Neural Information Processing Systems 2024, NeurIPS 2024, Vancouver, BC, Canada, December 10-15, 2024 (2024)
 19. Jin, H., Liu, I., Xu, P., Zhang, X., Han, S., Bi, S., Zhou, X., Xu, Z., Su, H.: Tensor: Tensorial inverse rendering. In: 2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2023, Vancouver, BC, Canada, June 17-24, 2023. pp. 165–174. IEEE (2023). <https://doi.org/10.1109/CVPR52729.2023.00024>
 20. Kerbl, B., Kopanas, G., Leimkühler, T., Drettakis, G.: 3d gaussian splatting for real-time radiance field rendering. ACM Trans. Graph. **42**(4), 139:1–139:14 (2023). <https://doi.org/10.1145/3592433>
 21. Kuang, Z., Zhang, Y., Yu, H., Agarwala, S., Wu, S., Wu, J.: Stanford-orb: A real-world 3d object inverse rendering benchmark. In: Oh, A., Naumann, T., Globerson, A., Saenko, K., Hardt, M., Levine, S. (eds.) Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10-16, 2023 (2023)

22. Lehtinen, J., Munkberg, J., Hasselgren, J., Laine, S., Karras, T., Aittala, M., Aila, T.: Noise2noise: Learning image restoration without clean data. In: Dy, J.G., Krause, A. (eds.) *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholmsmässan, Stockholm, Sweden, July 10-15, 2018*. *Proceedings of Machine Learning Research*, vol. 80, pp. 2971–2980. PMLR (2018)
23. Lensch, H.P.A., Kautz, J., Goesele, M., Heidrich, W., Seidel, H.: Image-based reconstruction of spatial appearance and geometric detail. *ACM Trans. Graph.* **22**(2), 234–257 (2003). <https://doi.org/10.1145/636886.636891>
24. Liang, R., Gojcic, Z., Ling, H., Munkberg, J., Hasselgren, J., Lin, Z., Gao, J., Keller, A., Vijaykumar, N., Fidler, S., Wang, Z.: Diffusionrenderer: Neural inverse and forward rendering with video diffusion models. *CoRR* **abs/2501.18590** (2025). <https://doi.org/10.48550/ARXIV.2501.18590>
25. Liang, R., He, K., Gojcic, Z., Gilitzenski, I., Fidler, S., Vijaykumar, N., Wang, Z.: Luxdit: Lighting estimation with video diffusion transformer. *CoRR* **abs/2509.03680** (2025). <https://doi.org/10.48550/ARXIV.2509.03680>
26. Litman, Y., Patashnik, O., Deng, K., Agrawal, A., Zawar, R., la Torre, F.D., Tulsiani, S.: Materialfusion: Enhancing inverse rendering with material diffusion priors. In: *International Conference on 3D Vision, 3DV 2025, Singapore, March 25-28, 2025*. pp. 802–812. IEEE (2025). <https://doi.org/10.1109/3DV66043.2025.00079>
27. Litman, Y., la Torre, F.D., Tulsiani, S.: Lightswitch: Multi-view relighting with material-guided diffusion. *CoRR* **abs/2508.06494** (2025). <https://doi.org/10.48550/ARXIV.2508.06494>
28. Lorensen, W.E., Cline, H.E.: Marching cubes: A high resolution 3d surface construction algorithm. In: Stone, M.C. (ed.) *Proceedings of the 14th Annual Conference on Computer Graphics and Interactive Techniques, SIGGRAPH '87, Anaheim, California, USA, July 27-31, 1987*. pp. 163–169. ACM (1987). <https://doi.org/10.1145/37401.37422>
29. Luan, F., Zhao, S., Bala, K., Dong, Z.: Unified shape and SVBRDF recovery using differentiable monte carlo rendering. *Comput. Graph. Forum* **40**(4), 101–113 (2021). <https://doi.org/10.1111/CGF.14344>
30. Mildenhall, B., Hedman, P., Martin-Brualla, R., Srinivasan, P.P., Barron, J.T.: Nerf in the dark: High dynamic range view synthesis from noisy raw images. In: *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2022, New Orleans, LA, USA, June 18-24, 2022*. pp. 16169–16178. IEEE (2022). <https://doi.org/10.1109/CVPR52688.2022.01571>
31. Mildenhall, B., Srinivasan, P.P., Tancik, M., Barron, J.T., Ramamoorthi, R., Ng, R.: Nerf: Representing scenes as neural radiance fields for view synthesis. In: Vedaldi, A., Bischof, H., Brox, T., Frahm, J. (eds.) *Computer Vision - ECCV 2020 - 16th European Conference, Glasgow, UK, August 23-28, 2020, Proceedings, Part I. Lecture Notes in Computer Science*, vol. 12346, pp. 405–421. Springer (2020). https://doi.org/10.1007/978-3-030-58452-8_24
32. Munkberg, J., Wang, Z., Liang, R., Shen, T., Hasselgren, J.: Videomat: Extracting PBR materials from video diffusion models. *Comput. Graph. Forum* **44**(4) (2025). <https://doi.org/10.1111/CGF.70180>
33. Oscanoa, J., Alkan, C., Abraham, D., Nurdinova, A., Ennis, D., Vasanaawala, S., Mardani, M., Pauly, J.M.: Variational diffusion models for MRI blind inverse problems. In: *NeurIPS 2023 Workshop on Deep Learning and Inverse Problems (2023)*
34. Petschnigg, G., Szeliski, R., Agrawala, M., Cohen, M.F., Hoppe, H., Toyama, K.: Digital photography with flash and no-flash image pairs. *ACM Trans. Graph.* **23**(3), 664–672 (2004). <https://doi.org/10.1145/1015706.1015777>

35. Phongthawee, P., Chinchuthakun, W., Sinsunthithet, N., Jampani, V., Raj, A., Khungurn, P., Suwajanakorn, S.: Diffusionlight: Light probes for free by painting a chrome ball. In: 2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2024, Seattle, WA, USA, June 16-22, 2024. pp. 98–108. IEEE (2024). <https://doi.org/10.1109/CVPR52733.2024.00018>
36. Schmitt, C., Donné, S., Riegler, G., Koltun, V., Geiger, A.: On joint estimation of pose, geometry and svbrdf from a handheld scanner. In: 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13-19, 2020. pp. 3490–3500. Computer Vision Foundation / IEEE (2020). <https://doi.org/10.1109/CVPR42600.2020.00355>
37. Sharma, P., Philip, J., Gharbi, M., Freeman, B., Durand, F., Deschaintre, V.: Materialistic: Selecting similar materials in images. *ACM Trans. Graph.* **42**(4), 154:1–154:14 (2023). <https://doi.org/10.1145/3592390>
38. Sun, C., Cai, G., Li, Z., Yan, K., Zhang, C., Marshall, C.S., Huang, J., Zhao, S., Dong, Z.: Neural-pbir reconstruction of shape, material, and illumination. In: IEEE/CVF International Conference on Computer Vision, ICCV 2023, Paris, France, October 1-6, 2023. pp. 18000–18010. IEEE (2023). <https://doi.org/10.1109/ICCV51070.2023.01654>
39. Sun, H., Gao, Y., Xie, J., Yang, J., Wang, B.: SVG-IR: spatially-varying gaussian splatting for inverse rendering. In: 2025 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2025, Nashville, TN, USA, June 11-15, 2025. pp. 16143–16152. Computer Vision Foundation / IEEE (2025). <https://doi.org/10.1109/CVPR52734.2025.01505>
40. Tang, J., Lavine, M., Verbin, D., Garbin, S.J., Nießner, M., Martin-Brualla, R., Srinivasan, P.P., Henzler, P.: ROGR: relightable 3d objects using generative relighting. *CoRR* **abs/2510.03163** (2025). <https://doi.org/10.48550/ARXIV.2510.03163>
41. 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: Ranzato, M., Beygelzimer, A., Dauphin, Y.N., Liang, P., Vaughan, J.W. (eds.) *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual*. pp. 27171–27183 (2021)
42. Wiersma, R., Philip, J., Hasan, M., Mullia, K., Luan, F., Eisemann, E., Deschaintre, V.: Uncertainty for SVBRDF acquisition using frequency analysis. In: *Proceedings of the Special Interest Group on Computer Graphics and Interactive Techniques Conference Papers, SIGGRAPH Conference Papers '25, Vancouver, BC, Canada, August 10-14, 2025*. pp. 169:1–169:12. ACM (2025). <https://doi.org/10.1145/3721238.3730592>
43. Wu, X., Zhu, P., Lyu, J., Liu, X., Guo, J., Guo, Y., Xu, W., Lyu, C.: Matmart: Material reconstruction of 3d objects via diffusion. *CoRR* **abs/2511.18900** (2025). <https://doi.org/10.48550/ARXIV.2511.18900>
44. Zeng, C., Dong, Y., Peers, P., Kong, Y., Wu, H., Tong, X.: Dilightnet: Fine-grained lighting control for diffusion-based image generation. In: *ACM SIGGRAPH 2024 Conference Papers, SIGGRAPH 2024, Denver, CO, USA, 27 July 2024-1 August 2024*. p. 73. ACM (2024). <https://doi.org/10.1145/3641519.3657396>
45. Zeng, Z., Deschaintre, V., Georgiev, I., Hold-Geoffroy, Y., Hu, Y., Luan, F., Yan, L., Hasan, M.: Rgb $\leftrightarrow$ x: Image decomposition and synthesis using material- and lighting-aware diffusion models. In: *Burbano, A., Zorin, D., Jarosz, W. (eds.) ACM SIGGRAPH 2024 Conference Papers, SIGGRAPH 2024, Denver, CO, USA, 27 July*

- 2024-1 August 2024. p. 75. ACM (2024). <https://doi.org/10.1145/3641519.3657445>
46. Zhang, R., Isola, P., Efros, A.A., Shechtman, E., Wang, O.: The unreasonable effectiveness of deep features as a perceptual metric. CoRR **abs/1801.03924** (2018)
 47. Zhang, T., Kuang, Z., Jin, H., Xu, Z., Bi, S., Tan, H., Zhang, H., Hu, Y., Hasan, M., Freeman, W.T., Zhang, K., Luan, F.: Relitlm: Generative relightable radiance for large reconstruction models. In: The Thirteenth International Conference on Learning Representations, ICLR 2025, Singapore, April 24-28, 2025. OpenReview.net (2025)
 48. Zhang, Y., Sun, J., He, X., Fu, H., Jia, R., Zhou, X.: Modeling indirect illumination for inverse rendering. In: 2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2022, New Orleans, LA, USA, June 18-24, 2022. pp. 18622–18631. IEEE (2022). <https://doi.org/10.1109/CVPR52688.2022.01809>
 49. Zhao, S., Jakob, W., Li, T.M.: Physics-based differentiable rendering: A comprehensive introduction. In: ACM SIGGRAPH 2020 Courses. pp. 14:1–14:30 (2020)
 50. Zhao, X., Srinivasan, P.P., Verbin, D., Park, K., Martin-Brualla, R., Henzler, P.: Illuminerf: 3d relighting without inverse rendering. In: Advances in Neural Information Processing Systems 38: Annual Conference on Neural Information Processing Systems 2024, NeurIPS 2024, Vancouver, BC, Canada, December 10-15, 2024 (2024)
 51. Zheng, H., Chu, W., Zhang, B., Wu, Z., Wang, A., Feng, B., Zou, C., Sun, Y., Kovachki, N.B., Ross, Z.E., Bouman, K.L., Yue, Y.: Inversebench: Benchmarking plug-and-play diffusion priors for inverse problems in physical sciences. In: The Thirteenth International Conference on Learning Representations, ICLR 2025, Singapore, April 24-28, 2025. OpenReview.net (2025)
 52. Zhou, Z., Chen, G., Dong, Y., Wipf, D.P., Yu, Y., Snyder, J.M., Tong, X.: Sparse-as-possible SVBRDF acquisition. ACM Trans. Graph. **35**(6), 189:1–189:12 (2016). <https://doi.org/10.1145/2980179.2980247>
 53. Zhu, J., Zhuang, P., Koyejo, S.: HIFA: high-fidelity text-to-3d generation with advanced diffusion guidance. In: The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024. OpenReview.net (2024)
 54. Zhu, Z., Wang, B., Yang, J.: GS-ROR2: bidirectional-guided 3dgs and SDF for reflective object relighting and reconstruction. ACM Trans. Graph. **45**(1), 4:1–4:19 (2026). <https://doi.org/10.1145/3759248>
 55. Zickler, T.E., Enrique, S., Ramamoorthi, R., Belhumeur, P.N.: Reflectance sharing: Image-based rendering from a sparse set of images. In: Deussen, O., Keller, A., Bala, K., Dutré, P., Fellner, D.W., Spencer, S.N. (eds.) Proceedings of the Eurographics Symposium on Rendering Techniques, Konstanz, Germany, June 29-July 1, 2005. pp. 253–264. Eurographics Association (2005). <https://doi.org/10.2312/EGWR/EGSR05/253-264>