

GS³: Efficient Relighting with Triple Gaussian Splatting

ZOUBIN BI, YIXIN ZENG, CHONG ZENG, FAN PEI, XIANG FENG, KUN ZHOU, and HONGZHI WU, State Key Lab of CAD&CG, Zhejiang University, China



Fig. 1. From 500-2,000 multi-view one-light-at-a-time (OLAT) input photographs, we train a representation consisting of spatial and angular Gaussians, for real-time, high-quality novel lighting-and-view synthesis. The reconstructions of a number of challenging objects with complex geometry and appearance are shown above. Please refer to the accompanying video for animated results with varying lighting and view conditions. We achieve a training time of 40-70 minutes and a rendering speed of 90 frames per second on a single commodity GPU.

We present a spatial and angular Gaussian based representation and a triple splatting process, for real-time, high-quality novel lighting-and-view synthesis from multi-view point-lit input images. To describe complex appearance, we employ a Lambertian plus a mixture of angular Gaussians as an effective reflectance function for each spatial Gaussian. To generate self-shadow, we splat all spatial Gaussians towards the light source to obtain shadow values, which are further refined by a small multi-layer perceptron. To compensate for other effects like global illumination, another network is trained to compute and add a per-spatial-Gaussian RGB tuple. The effectiveness of our representation is demonstrated on 30 samples with a wide variation in geometry (from solid to fluffy) and appearance (from translucent to anisotropic), as well as using different forms of input data, including rendered images of synthetic/reconstructed objects, photographs captured with a handheld camera and a flash, or from a professional lightstage. We achieve a training time of 40-70 minutes and a rendering speed of 90 fps on a single commodity GPU. Our results compare favorably with state-of-the-art techniques in terms of quality/performance.

CCS Concepts: • **Computing methodologies** → **Reflectance modeling; Image-based rendering.**

Additional Key Words and Phrases: 3D Gaussian splatting, reflectance, geometry

*Corresponding authors: Kun Zhou & Hongzhi Wu.

**The first two authors contributed equally.

Authors' Contact Information: Zoubin Bi; Yixin Zeng; Chong Zeng; Fan Pei; Xiang Feng; Kun Zhou; Hongzhi Wu, hww@acm.org, State Key Lab of CAD&CG, Zhejiang University, Hangzhou, China.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

SA Conference Papers '24, December 3–6, 2024, Tokyo, Japan

© 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 979-8-4007-1131-2/24/12

<https://doi.org/10.1145/3680528.3687576>

ACM Reference Format:

Zoubin Bi, Yixin Zeng, Chong Zeng, Fan Pei, Xiang Feng, Kun Zhou, and Hongzhi Wu. 2024. GS³: Efficient Relighting with Triple Gaussian Splatting. In *SIGGRAPH Asia 2024 Conference Papers (SA Conference Papers '24)*, December 3–6, 2024, Tokyo, Japan. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3680528.3687576>

1 INTRODUCTION

Realistically reproducing the look of a physical object at different view and lighting conditions in the virtual world has been a long-standing problem in computer graphics and computer vision. It is critical in various applications, including cultural heritage, e-commerce and visual effects.

Digital representation of shape and appearance plays a key role in this task. Traditional representations like 3D surface mesh and parametric spatially-varying bidirectional reflectance distribution function (SVBRDF) are widely used in both academia and industry [Dorsey et al. 2010]. However, they are inherently difficult to jointly optimize with respect to input photographs, therefore often leading to suboptimal results. In the past five years, implicit representations, such as Neural Radiance Fields (NeRF) [Mildenhall et al. 2020], demonstrate extraordinary ability in high-quality novel view synthesis, and even relighting [Jin et al. 2023; Lyu et al. 2022; Zeng et al. 2023]. But these techniques often suffer from expensive training computation and/or slow rendering speed, limiting their applications in practice.

Recently, 3D Gaussian Splatting (GS) [Kerbl et al. 2023] gains tremendous popularity in high-quality and efficient reconstruction of Lambertian-dominant objects/scenes baked with static lighting, by essentially upgrading to a differentiable tile-based splatting method. Considerable research efforts [Gao et al. 2023; Jiang et al. 2023; Liang et al. 2023b; Saito et al. 2023] are made to extend GS towards novel lighting-and-view synthesis. However, high-quality relighting remains challenging, as complex appearance like

anisotropic reflectance is not modeled, and the shading computation is usually confined to *surface geometry* only.

In this paper, we present a novel representation based on spatial and angular Gaussians along with a triple splatting process, for real-time, high-quality novel lighting-and-view synthesis from around 500-2,000 multi-view input images, lit with one point light at a time. To describe complex appearance, we replace the spherical harmonics (SH) associated with each vanilla spatial Gaussian with a Lambertian and a mixture of angular Gaussians (a differentiable anisotropic spherical Gaussian modified from [Saito et al. 2023; Xu et al. 2013]), essentially representing a microfacet normal distribution (i.e., 1st splatting). To efficiently support self-shadow, we splat all spatial Gaussians toward the light source, by reusing the same high-performance pipeline as the original screen-space splatting (i.e., 2nd splatting). To compensate for other effects like global illumination, we employ an additional multi-layer perceptron (MLP) to compute a RGB tuple for each spatial Gaussian. The above three factors are splatted to the camera and mixed to produce an image (i.e., 3rd splatting), whose difference with a corresponding input photograph drives the optimization of our representation in an end-to-end, fully differentiable manner.

The effectiveness of our representation is demonstrated on samples with a wide variation in geometry and appearance. With a modest increase in footprint and training/runtime computation compared with GS, we obtain high-performance and -quality synthesis results under novel lighting and view conditions. These results compare favorably with state-of-the-art techniques in terms of quality/performance. Our representation can handle a wide spectrum of input data, including rendered images of synthetic/reconstructed objects, as well as photographs captured with a smartphone and a flash, or from a professional lightstage. Our code and data are publicly available at <https://GSrelight.github.io/>.

2 RELATED WORK

Below we review the most relevant work mainly in chronological order. While some existing papers require additional lighting estimation/decoupling, we would like to emphasize that this paper focuses on a general relightable representation only. We assume that the lighting is known or calibrated, and can work well with a wide spectrum of input data, from synthetic images to photographs captured with a low-end camera or a high-end lightstage. Interested readers are referred to excellent recent surveys for a broader view of the topic [Fei et al. 2024; Tewari et al. 2022; Wu et al. 2024].

2.1 Traditional Relighting

While widely deployed in practice, traditional representations, such as 3D surface mesh and parametric SVBRDF which varies with location, view and lighting directions, are challenging to optimize jointly. The majority of existing work performs separate estimations of shape and appearance, the latter of which is typically represented as attributes defined on a known 3D geometry. Dense lights are used to remove adversarial effects like strong specular reflections to enable geometry reconstruction with multi-view stereo, prior to reflectance estimation [Kang et al. 2019; Tunwattanapong et al. 2013]. Zhou et al. [2013] recover a 3D shape from multi-view photometric

cues, and then compute isotropic surface reflectance. Structured illumination is adopted to recover highly precise surface geometry, after which the appearance is computed [Holroyd et al. 2010; Xu et al. 2023]. Despite training an image-space neural renderer [Gao et al. 2020; Philip et al. 2021], both methods learn to relight using buffers rendered with fixed, non-optimizable geometry. Due to the difficulty in performing an end-to-end, joint optimization of shape and appearance, the result quality of the above work is limited: once computed, errors in geometric estimation cannot be easily fixed, and may contaminate the subsequent appearance reconstruction.

On the other hand, few exception papers try to conduct a highly involved optimization that alternates between solving for shape and reflectance [Nam et al. 2018; Wu et al. 2015; Xia et al. 2016]; the latter two even solve for unknown environment lighting. However, due to the non-differentiable nature of directly optimizing common traditional representations, approximations/tricks have to be applied from one place to another. Therefore, the result quality is still not satisfactory. It is not even clear if the optimization converges.

2.2 Neural Relighting

With the advances in deep learning, neural implicit representations and/or modern large-scale optimization tools make it possible to jointly solve for geometry and appearance in a fully differentiable, end-to-end fashion. Compared with traditional relighting, direct optimization with respect to input photographs leads to higher quality results. Implicit representations like NeRF [Mildenhall et al. 2020] demonstrate unprecedented quality in novel view synthesis. And considerable research efforts are made to extend the idea to relighting [Bi et al. 2020; Munkberg et al. 2022; Sun et al. 2021; Zhang et al. 2021b]. Due to the space limit, below we briefly review representative approaches.

One class of existing work takes images under *unknown* environment lighting(s) as input, and has to deal with the fundamental lighting-material ambiguity. These methods typically integrate approximate physical-based rendering (PBR), along with various regularization to better condition the optimization. Boss et al. [2021] assume spatial coherence and jointly optimize a compressed latent BRDF space. Zhang et al. [2021a] employ a homogeneous specular appearance. Both methods ignore occlusion and indirect illumination. Extending from surface BRDFs, a MLP-predicted microflake volume is proposed in [Zhang et al. 2023]. Jin et al. [2023] calculate visibility from the volume transmittance in a Siamese radiance field, and consider second-bounce illumination. A pre-trained neural renderer is proposed in [Liang et al. 2023a], as a neural approximation of the explicitly PBR rendering equation. Lyu et al. [2022] incorporate PBR prior by bootstrapping light transport modeling with synthesized OLAT images as training data, and refine the result with captured photographs. The shape and appearance are not optimized in tandem.

Another class of work directly takes photographs captured with *known/calibrated* lighting conditions as input. Srinivasan et al. [2021] train MLPs to predict fields of volumetric density, surface normal, material parameters, intersection, and visibility, which are jointly optimized via inverse rendering. Yu et al. [2023] employ a neural

scattering function that approximates radiance transfer from a distant light, with OLAT input images. Recently, Zeng et al. [2023] improve over [Gao et al. 2020] with a neural implicit radiance representation, and add shadow and highlight hints to help a network to model high frequency light transport effects. A hybrid point-volumetric representation is proposed in concurrent work [Chung et al. 2024] for efficient inverse rendering. Due to hard visibility thresholding, transparent/furry objects are not supported.

While most state-of-the-art neural techniques can produce high-quality results, both the training and rendering costs are substantially more expensive than, e.g., GS-based methods. And it is non-trivial to directly apply the ideas here to GS, due to the considerable differences between the representations.

2.3 Gaussian-Splatting-Based Relighting

Recently, 3D Gaussian Splatting [Kerbl et al. 2023] introduces a highly efficient differentiable rasterization pipeline for a Gaussian-based representation, substantially improving the training time and runtime performance. Low-order SH is employed in each vanilla Gaussian to represent Lambertian-dominant appearance variations under a fixed environment lighting. Several approaches replace SH with higher-frequency functions to improve the view-dependent synthesis quality, which, however, cannot support lighting change [Malarz et al. 2023; Yang et al. 2024; Ye et al. 2024].

Towards the goal of relighting with more complex, general appearance, a number of techniques have been proposed. Similar to neural relighting, the majority of related work here takes images under an unknown environment lighting as input [Gao et al. 2023; Jiang et al. 2023; Liang et al. 2023b; Shi et al. 2023]. The basic idea is to model the appearance for each 3D Gaussian as an isotropic parametric BRDF, precompute or ray-trace the visibility, sample indirect illumination and store as low-frequency SH, and perform inverse rendering. All these methods require well defined *surface* normals to properly regularize their optimizations, which limits the applicable geometry to opaque ones with clear boundaries.

Another line of work takes images captured with varying light sources as input. Saito et al. [2023] propose a relightable head avatar. For each 3D Gaussian, the specular reflectance is modeled as a single learnable isotropic spherical Gaussian, and the light visibility is computed from a neural network.

All the above work for general objects/scenes does not handle challenging appearance such as anisotropic reflections, with the exception of specialized models (e.g., [Luo et al. 2024] for hair). Their relighting quality is limited, when compared with latest neural relighting approaches (e.g., [Zeng et al. 2023]). In comparison, we present the first general GS-based relightable representation for complex geometry and appearance. Unlike the aforementioned work, we do not rely on any regularizations/strong priors in the optimization (e.g., we do not require well defined surface normals), which can be fragile in handling complex cases. Our quality is comparable to or higher than state-of-the-art neural relighting, while our computation is substantially more efficient, by exploiting the differentiable rasterization pipeline of GS.

3 PRELIMINARIES

Our pipeline builds upon the highly efficient GS [Kerbl et al. 2023]. Similarly, we represent the geometry with anisotropic 3D Gaussians (or *spatial Gaussians* in this paper), whose density at a 3D location \mathbf{p} is defined as:

$$G_{\text{spa}}(\mathbf{p}) = \exp\left(-\frac{1}{2}(\mathbf{p} - \boldsymbol{\mu})^\top \Sigma^{-1}(\mathbf{p} - \boldsymbol{\mu})\right). \quad (1)$$

Here $\boldsymbol{\mu}$ is the 3D center of the Gaussian, and Σ is a covariance matrix $\Sigma = RSS^\top R^\top$, where S is a scaling matrix and R is a rotation matrix.

Each spatial Gaussian is associated with an opacity γ_j and a color \mathbf{c}_j . To generate an image, the spatial Gaussians are projected to the screen as 2D Gaussian splats. Then, for each pixel, its intersecting Gaussian splats are sorted and alpha-blended as follows:

$$\boldsymbol{\zeta} = \sum_j \mathbf{c}_j \beta_j \gamma_j T_j, \quad (2)$$

$$T_j = \prod_{k=0}^{j-1} (1 - \beta_k \gamma_k). \quad (3)$$

Here $\boldsymbol{\zeta}$ is the final color of the current pixel, T_j is the cumulative opacity for the top-most j Gaussian splats, and β_k is the density at the current pixel center of the k -th Gaussian splat.

4 OUR APPROACH

We take photographs of an object/scene from different calibrated views, lit with one point light at a time as input, and output a set of spatial Gaussians to represent the geometry, each of which is associated with an opacity and an appearance function, mainly represented as a linear combination of angular Gaussians.

To efficiently render an image under a point light, a deferred shading approach is adopted. First, we color each spatial Gaussian by evaluating its appearance function, and splat them into a *shading image* (Sec. 4.1). Next, for each spatial Gaussian, we compute a shadow value by splatting all of them towards the light (which we call *shadow splatting*), and refine it with an MLP. We color each spatial Gaussian with its own shadow value, and splat them into a *shadow image* (Sec. 4.2). Finally, we color each spatial Gaussian with another MLP that represents unhandled effects like global illumination, and splat them into a *residual image* (Sec. 4.3). The final rendering result is computed by multiplying the shading image with the shadow one, and adding the residual image, on a per-pixel basis. Please refer to Fig. 2 for a graphical illustration.

For training, an end-to-end, joint optimization of spatial Gaussians and corresponding appearance functions are performed. Similar to vanilla GS, it minimizes the differences between the input photographs and our rendering. Adaptive density control in GS is also applied.

4.1 Shading

To accurately represent complex appearance (e.g., anisotropic reflections), we replace for each spatial Gaussian the low-order SH as defined in vanilla GS with a view- and lighting-dependent function f as:

$$f(\boldsymbol{\omega}'_i, \boldsymbol{\omega}'_o) = \rho_d f_d(\boldsymbol{\omega}'_i) + \rho_s f_s(\boldsymbol{\omega}'_i, \boldsymbol{\omega}'_o). \quad (4)$$

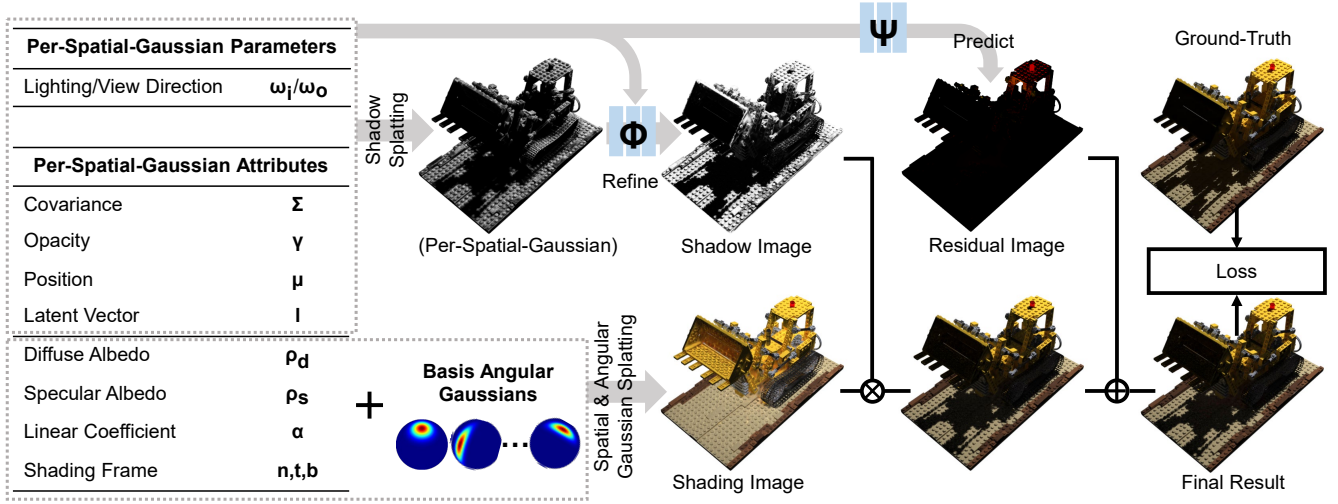


Fig. 2. Our deferred-shading-based pipeline. First, we color each spatial Gaussian by evaluating its appearance function, defined as a Lambertian plus a linear combination of basis angular Gaussians, and splat into a shading image. Next, for each spatial Gaussian, we compute a shadow value by splatting all of them towards the light (i.e., shadow splatting), and refine it with an MLP. We color each spatial Gaussian with its shadow value, and splat them into a shadow image. Finally, we color each spatial Gaussian with another MLP that represents unhandled effects like global illumination, and splat them into a residual image. The final rendering result is computed by multiplying the shading image with the shadow one, and adding the residual image, on a per-pixel basis.

Here ρ_d/ρ_s is the diffuse/specular albedo, f_d/f_s is the diffuse/specular appearance function to be defined in the remainder of this subsection, and ω'_i/ω'_o is the local lighting/view direction, respectively. The learnable shading frame at a spatial Gaussian is defined as $[\mathbf{n}, \mathbf{t}, \mathbf{b}]$, which, as the name suggests, consists of normal, tangent and binormal. In practice, we do not explicit store the three vectors. Instead, we use a unit quaternion to express an equivalent rotation transform from the world space to the shading frame of a particular Gaussian. Note that the shading frame is totally independent from the 3 axes of the spatial Gaussian (represented in Σ).

To generate a shading image, we color each spatial Gaussian by evaluating f , and then perform splatting. Below we describe the details about f_d and f_s .

Diffuse Appearance. We modify a common cosine-weighted Lambertian function to the following function f_d :

$$f_d(\omega'_i) = \frac{\text{ELU}(\mathbf{n}' \cdot \omega'_i) + \varepsilon(1 - \frac{1}{e})}{(1 + \varepsilon(1 - \frac{1}{e}))\pi}. \quad (5)$$

Here \mathbf{n}' is the normal in the shading frame, and we set the hyperparameter of ELU and ε as 0.01. Note that f_d is slightly different from the standard definition, and its gradient is non-zero for any ω'_i , which is amenable for differentiable optimization. In comparison, the original cosine-weighted Lambertian has a zero gradient over the lower hemisphere, which would form a "dead zone" once the optimization gets stuck there.

Specular Appearance. To represent complex all-frequency specular appearance, we model f_s as a mixture of modified anisotropic spherical Gaussians (denoted as *angular Gaussians* in this paper) below:

$$f_s(\omega'_i, \omega'_o) = \sum_j \alpha_j G_{\text{ang},j}(\mathbf{h}'), \quad (6)$$

where α_j is a weight, and \mathbf{h}' is the half vector computed as $\mathbf{h}' = \frac{\omega'_i + \omega'_o}{\|\omega'_i + \omega'_o\|}$. $G_{\text{ang},j}$ is an angular Gaussian, defined as:

$$G_{\text{ang}}(\mathbf{h}') = \frac{1}{\sigma_z} \exp\left(-\frac{1}{2} \left(\frac{\arccos(\mathbf{h}' \cdot \mathbf{z}) \sqrt{\left(\frac{s'_x}{\sigma_x}\right)^2 + \left(\frac{s'_y}{\sigma_y}\right)^2}}{\sigma_z} \right)^2\right). \quad (7)$$

Here $[\mathbf{x}, \mathbf{y}, \mathbf{z}]$ is the local frame of the angular Gaussian, s' is the normalized result of the projection of \mathbf{h}' onto the x-y plane, and $\sigma_x/\sigma_y/\sigma_z$ are the standard deviations in three dimensions. Note that we extend the isotropic definition in [Saito et al. 2023] with [Xu et al. 2013] to support anisotropy. Directly employing the definition from [Xu et al. 2013] often cannot model highly specular reflections well, and its smooth term is unfriendly to differentiable optimization, according to [Saito et al. 2023] and our pilot study. We call the evaluation of Eq. 6 as *angular Gaussian splatting*, as it involves the mixing of multiple Gaussians.

Furthermore, for a particular object/scene, a set of the basis angular Gaussians are shared across all spatial Gaussians, essentially exploiting the spatial coherence to better condition the optimization, as common in related work [Chen et al. 2014; Lensch et al. 2003; Nam et al. 2018]. Consequently, for each spatial Gaussian, the complete set of learnable parameters to represent an f consist of $[\mathbf{n}, \mathbf{t}, \mathbf{b}]$, $[\mathbf{x}, \mathbf{y}, \mathbf{z}]$, $[\sigma_x, \sigma_y, \sigma_z]$, ρ_d, ρ_s , and the set of weights $\{\alpha_j\}$ to linearly combine the shared basis angular Gaussians, whose total number is 8 in main experiments. Note that when the input appearance information is sufficient to condition our optimization [Ma et al. 2021; Tunwattanapong et al. 2013], it is possible to use a separate set of angular Gaussians for each spatial Gaussian, instead of sharing them, to further improve the result quality.

4.2 Shadowing

The shading image does not account for any shadowing effects. To support them, a naïve but expensive method would trace a shadow ray from a scene point to the point light, and perform a line integral of opacity along the way as the final visibility result with respect to the light. To develop an efficient algorithm for shadow computation on Gaussians, we observe that traditional shadow mapping [Williams 1978] reuses the high-performance rasterization pipeline for the view from the light instead of the camera. Here we apply a similar idea, by performing high-performance Gaussian splatting towards the point light. Please see Fig. 3 for an illustration.




Fig. 3. Shadow computation. We first splat all spatial Gaussians towards the light source, and compute the accumulated opacity T_m and the density β_m at each intersection with each shadow ray (left). For each spatial Gaussian, all of its T_m with respect to different shadow rays are weighted-averaged by corresponding β_m to obtain a shadow value T (center). Next, this shadow value is refined with a small MLP. Finally, Gaussians with refined shadow values (T') are splatted towards the camera to produce the shadow image (right). The above Gaussians may appear incorrectly, when viewed in a program other than Adobe Acrobat.

Specifically, we first splat all spatial Gaussians towards the point light, by setting up a perspective camera whose center is the light, similar to how a shadow map is computed (i.e., shadow splatting). Here we use the same image resolution as an input image. For a shadow ray (with an index of m), a number of spatial Gaussians whose 2D splats intersecting the ray are sorted, according to the distances to the light. Next, we compute the cumulative opacity T_m , according to Eq. 3, for an intersecting spatial Gaussian as its shadow value.

Note that the 2D splat of a spatial Gaussian is likely to intersect a number of shadow rays, with multiple shadow values computed at each ray. In this case, we compute an average $T = \frac{\sum_m \beta_m T_m}{\sum_m \beta_m}$ as the result, weighted by β_m , the projected 2D Gaussian density at an intersection. Similar to shadow mapping, we apply a shadow bias of 0.015 to alleviate "z-fighting".

To further improve the shadow quality (e.g., denoise), we refine for each spatial Gaussian its shadow value with an MLP, Φ , as follows:

$$T' = \Phi(T, \omega_i; \mu, \mathbf{l}). \quad (8)$$

Here T/T' is the shadow value before/after neural refinement, respectively. In addition, Φ is a small 3-layer MLP with {32, 32, 32}. Each hidden layers is followed by a leaky ReLU activation, while the output layer a sigmoid one. The parameter μ (i.e., the spatial Gaussian center) is designed to make the MLP spatial aware, and the learnable \mathbf{l} is a per-spatial-Gaussian 6D latent vector. We apply 4-band positional encoding to μ and ω_i before sending to Φ . Finally,

we color each spatial Gaussian with its refined shadow value, and splat into a shadow image.

4.3 Other Effects

To consider other light transport not modeled in Sec. 4.1 & 4.2 (e.g., global illumination), we employ another MLP, $\Psi(\omega_o; \mu, \mathbf{l})$, to predict the impact of these effects for each spatial Gaussian. Specifically, Ψ is a 3-layer MLP with {128, 128, 128}. Each hidden layer is followed by a leaky ReLU activation, while the output layer a sigmoid one. The MLP is a function of ω_o only, which is a common parameterization for representing indirect illumination in real-time rendering [Akenine-Mller et al. 2018]. In addition, the parameter μ is the spatial Gaussian center, and the learnable \mathbf{l} is a shared latent vector defined in Sec. 4.2. We apply 4-band positional encoding to μ and ω_o . Finally, we evaluate Ψ to color each spatial Gaussian, and splat into a residual image.

4.4 Training

Loss. We use the loss function from [Kerbl et al. 2023], defined as follows:

$$\mathcal{L} = (1 - \lambda)\mathcal{L}_1 + \lambda\mathcal{L}_{D-SSIM}. \quad (9)$$

Here \mathcal{L}_1 is the L1 image loss, \mathcal{L}_{D-SSIM} is the sum of SSIM [Wang et al. 2004] of each difference image between an input photograph and our corresponding rendering, and $\lambda = 0.2$ in all experiments. Note that we do not impose any regularization on intermediate results/components. We find it sufficient and elegant to use the above end-to-end image loss for high-quality relighting.

Initialization. We initialize geometry properties of spatial Gaussians, following vanilla GS. For an angular Gaussian, we randomly sample σ_z from [0.13, 0.69], and set $\sigma_x = 0.5$ and $\sigma_y = 1.0$. Both ρ_d and ρ_s are initialized as (1, 1, 1), and each α is set to 0.5. The local/shading frame of each spatial/angular Gaussian is initialized to align with the axes in the world space.

Training Strategy & Details. To reduce the probability of getting stuck with an undesired local minimum and to increase robustness in practice, we use a two-stage training strategy to gradually increase the degrees of freedom. First, we limit the appearance function to be the Lambertian term only, and train for 15K iterations. We find that the shading frames converge stably in this stage. Next, we use the full appearance function with specular reflections, and train for 100K iterations. In all experiments, we employ the Adam optimizer with a momentum of 0.9. The learning rate varies with different parameters, similar to [Kerbl et al. 2023]. For ρ_d and ρ_s , the learning rate is set to 0.01. For angular Gaussians, we use a fixed rate of 0.01 before 40K iterations, and exponentially decay it to 0.0001 at 90K, and fix it afterwards.

4.5 Rendering

For a given view and a point light, the rendering process with our representation is described in the beginning of Sec. 4. To support a directional light, we switch from perspective projection to orthographic one in the shadow splatting process. Furthermore, rendering with an environment light is implemented as a linear combination of the results under a number of sampled directional lights.

5 RESULTS & DISCUSSIONS

All experiments are conducted on a workstation with dual AMD EPYC 7763 CPUs, 768GB DDR4 memory and an NVIDIA GeForce RTX 4090 GPU. It takes 40-70 minutes to train our representation (120K-750K spatial Gaussians and 8 basis angular Gaussians), whose rendering speed is over 90fps.

We reconstruct objects/scenes with a wide variation in geometry (from solid to fluffy) and appearance (from translucent to anisotropic) from multi-view point-lit images. Four forms of input data are tested: (1) rendered images of synthetic NeRF [Mildenhall et al. 2020]; (2) rendered images of captured results from [Kang et al. 2019] and OpenSVBRDF [Ma et al. 2023]; (3) captured point-lit photographs from NRHints [Zeng et al. 2023]; and (4) multi-view photometric images captured by a professional lightstage. The input images for (1) and (2) have a spatial resolution of 512^2 , while the resolution is 512^2 or 1024^2 for (3) and (4).

For comparison experiments, we take the official code of each method and retrain with the same set of 500-2,000 input images. For quantitative assessments of reconstruction quality, we compute PSNR, SSIM, and LPIPS [Zhang et al. 2018] averaged over all test images. Please also refer to the accompanying video for animated results with varying view and lighting. For ablation experiments, please refer to the supplemental material, as the space of the main paper is limited.

5.1 Results

Synthetic. Fig. 11 and Fig. 9 show our reconstruction results from rendered images of synthetic/captured subjects. For each subject, we use 2,000 training images and 400 test ones, with randomly sampled view and point light. TRANSLUCENT exhibits complex subsurface scattering; ANISOMETAL and DRUMS contain strong anisotropic appearance; FURBALL has a fuzzy geometry; and LEGO contains complex occlusions and shadows. In addition, MATERIALBALLS, TOWER, FABRIC, CUP and EGG show complex, spatially-varying specular reflections. In all cases, we successfully reconstruct a wide variety of challenging shapes and appearance with our representation.

Captured. In Fig. 6 and Fig. 9, we generate realistic novel lighting-and-view synthesis results from the data of [Zeng et al. 2023], acquired with a handheld camera and a smartphone with a flash. Exactly the same training and test data from their paper are used. CUP-FABRIC consists of translucent materials (cup), and isotropic (balls) and anisotropic reflections (fabric); PIXIU shows strong subsurface scattering and self-occlusions; FISH, CLUTTERED and CAT contain intricate details, like shadow and glint (ground-plane) and complex appearance (fur). PIKACHU includes glossy highlights and considerable self-occlusions.

In Fig. 5, we reconstruct from multi-view photometric photographs, captured with a professional lightstage with 24,576 LEDs and 2 cameras, similar to [Kang et al. 2019]. For each subject, we use 2,000 training images and 400 test images. ZHAOJUN and BOOT contain furry geometry and complex glinty/anisotropic appearance; Fox has highly complex occlusions and reflections on grass-like ground and the hair; LI'LONES is a doll with challenging parts like long hair, which often requires special, separate handling in existing work [Saito et al. 2023]. CONTAINER and NEFERTITI include highly

specular reflections with complex spatial variations. For all these highly challenging cases, we demonstrate high-quality reconstructions with a unified representation.




Fig. 4. Visualization of the basis angular Gaussians and the spatial distributions of corresponding weights. Each image shows the spatial distribution of the weight for a particular basis angular Gaussian, which is color-coded in the bottom-left corner. A brighter pixel indicates a larger weight.

Visualizations. In Fig. 8, we visualize the components of the pipeline, for a better understanding of what is going on under the hood. Note that the intermediate results demonstrate decent quality/decoupling from each other, despite that we do *not* impose regularization/constraint on any of them, unlike the majority of GS-based work. This shows the elegance of our simple, end-to-end image loss. In addition, Fig. 4 visualizes the basis angular Gaussians and the spatial distributions of corresponding weights (via splatting) of the same scene.

5.2 Comparisons

In Fig. 10, we compare with Neural Radiance Transfer Fields [Lyu et al. 2022]. Both approach are trained with the same set of 2,000 directional-lit images. Our results show higher-quality shadows and specular reflections than theirs. Moreover, due to the surface-based appearance representation, they cannot work well on fuzzy geometry like FURBALL.

Fig. 9 compares with NRHints [Zeng et al. 2023], the state-of-the-art relightable implicit representation. We can reconstruct the anisotropic highlights in DRUMS and the specular reflections on the floor of FISH, which their approach struggles with. For FUR, LEGO and CAT, they fail to reproduce many of the original spatial details, although they sometimes achieve a higher score than ours with the blurry reconstructions. Compared with NRHints, we obtain higher-quality or comparable results, and more than *an order of magnitude* higher performance both in training (40-70min vs. 15 hrs) and rendering (90fps vs. <1fps) on the same workstation.

In Fig. 7, we compare with OSF [Yu et al. 2023], one state-of-the-art method for reconstructing sub-surface scattering appearance, on TRANSLUCENT. While not explicitly modeled, scattering effects are faithfully recovered with our representation. In comparison, the lack of self-occlusion handling and accurate surface reflectance in OSF leads to a lower-quality result.

In Fig. 12, we compare with GS based/NeRF-like relighting methods, which take images lit with unknown environment lighting as

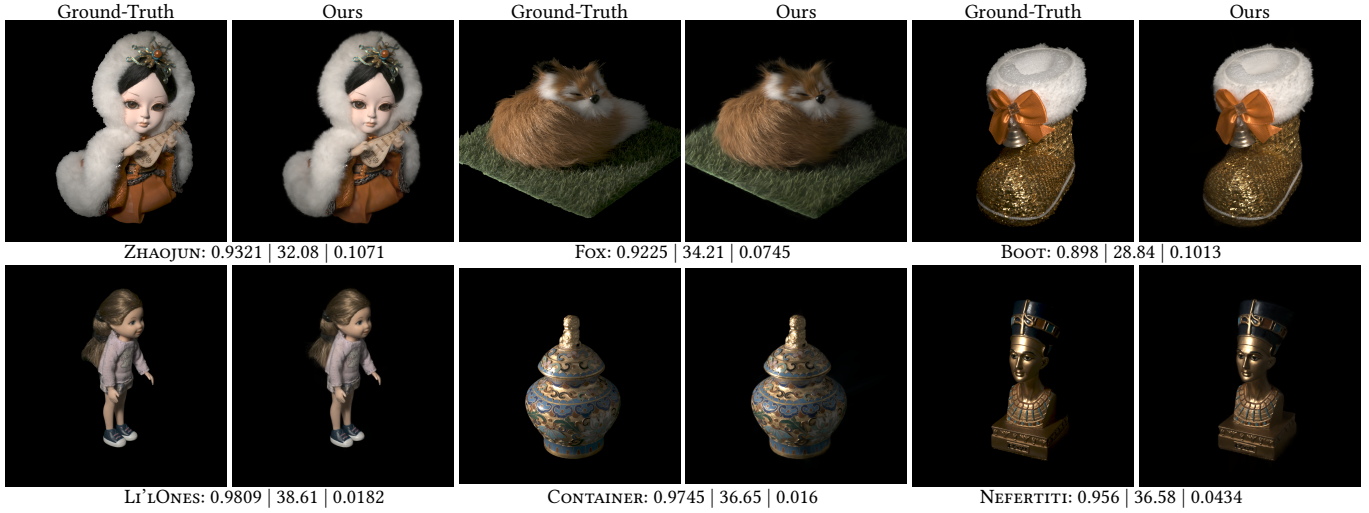


Fig. 5. Our relighting results on captured data from a professional lightstage. For each pair of images, the left one is the ground-truth photograph, and the right is our result. Average errors in SSIM, PSNR and LPIPS are reported at the bottom.




Fig. 6. Our relighting results on captured data from [Zeng et al. 2023]. For each pair of images, the left one is the ground-truth photograph, and the right is our rendering result. Average errors in SSIM, PSNR and LPIPS are reported at the bottom.




Fig. 7. Comparison with [Yu et al. 2023]. From the left to right, the ground-truth, the results of our approach and [Yu et al. 2023], respectively. Average errors in SSIM, PSNR and LPIPS are reported at the bottom of related images.

input, including GaussianShader [Jiang et al. 2023], GS-IR [Liang et al. 2023b], Relightable 3D Gaussian [Gao et al. 2023] and Tenso-IR [Jin et al. 2023]. For these methods, we train with the same set of 2,000 images rendered with an environment map, and test under a different one. To be as fair as possible, the same ground-truth environment map is supplied to all alternative methods as input during training; our approach uses the same number of point-lit input images; Another test environment map is used for relighting. Our

quality clearly surpasses alternative approaches in all cases. This is not surprising: state-of-the-art relighting techniques (e.g., [Zeng et al. 2023]) take images with *known* lighting as input, rather than with unknown environment lighting, to produce high-quality results.

6 LIMITATIONS & FUTURE WORK

Our work is subject to a number of limitations. First, we do not consider transparent materials, such as glass or gems. It will be interesting to replace the MLP Ψ with an explicitly modeling of other light transports like refraction or internal reflection. Moreover, under certain lighting/view conditions, our shadows are not as crisp as, e.g., using a mesh-based representation for geometry. Also for extremely high-frequency anisotropic appearance, the reconstructed highlights might blink. The main cause for both cases is the insufficient granularity of spatial Gaussians. We are intrigued to develop more advanced density control method, as well as establish additional direct gradient pathway(s), to solve this problem.

In the future, it is promising to search for the optimal acquisition conditions (i.e., view/lighting) for our representation to reduce the number of input images and improve the reconstruction quality at

the same time, by exploiting, e.g., the highly efficient learned illumination multiplexing [Kang et al. 2021]. While the images we capture with a professional lightstage is already useful as a benchmark for future research, it is desirable to build a large-scale database in the spirit of [Ma et al. 2023] to facilitate generative tasks like [Poirier-Ginter et al. 2024; Zeng et al. 2024] based on our representation.

ACKNOWLEDGMENTS

We would like to express our gratitude to Yue Dong, Guojun Chen, Xianmin Xu and Yaowen Chen for their generous help to this project. This work is partially supported by NSF China (62332015, 62227806 & 62421003), the Fundamental Research Funds for the Central Universities (226-2023-00145), the XPLOER PRIZE, and Information Technology Center and State Key Lab of CAD&CG, Zhejiang University.

REFERENCES

- Tomas Akenine-Miller, Eric Haines, and Naty Hoffman. 2018. *Real-Time Rendering, Fourth Edition* (4th ed.). A. K. Peters, Ltd., USA.
- Sai Bi, Zexiang Xu, Pratul Srinivasan, Ben Mildenhall, Kalyan Sunkavalli, Miloš Hašan, Yannick Hold-Geoffroy, David Kriegman, and Ravi Ramamoorthi. 2020. Neural Reflectance Fields for Appearance Acquisition. *arXiv:2008.03824* [cs.CV]
- Mark Boss, Raphael Braun, Varun Jampani, Jonathan T Barron, Ce Liu, and Hendrik Lensch. 2021. Nerd: Neural reflectance decomposition from image collections. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 12684–12694.
- Guojun Chen, Yue Dong, Pieter Peers, Jiawan Zhang, and Xin Tong. 2014. Reflectance scanning: Estimating shading frame and BRDF with generalized linear light sources. *ACM Transactions on Graphics (TOG)* 33, 4 (2014), 1–11.
- Hoon-Gyu Chung, Seokjun Choi, and Seung-Hwan Baek. 2024. Differentiable Point-based Inverse Rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 4399–4409.
- Julie Dorsey, Holly Rushmeier, and François Sillion. 2010. *Digital modeling of material appearance*. Elsevier.
- Ben Fei, Jingyi Xu, Rui Zhang, Qingyuan Zhou, Weidong Yang, and Ying He. 2024. 3D Gaussian Splatting as New Era: A Survey. *IEEE Transactions on Visualization and Computer Graphics* (2024), 1–20. <https://doi.org/10.1109/TVCG.2024.3397828>
- Duan Gao, Guojun Chen, Yue Dong, Pieter Peers, Kun Xu, and Xin Tong. 2020. Deferred neural lighting: free-viewpoint relighting from unstructured photographs. *ACM Transactions on Graphics (TOG)* 39, 6 (December 2020), 258.
- Jian Gao, Chun Gu, Youtian Lin, Hao Zhu, Xun Cao, Li Zhang, and Yao Yao. 2023. Relightable 3D Gaussian: Real-time Point Cloud Relighting with BRDF Decomposition and Ray Tracing. *arXiv:2311.16043* [cs.CV]
- Michael Holroyd, Jason Lawrence, and Todd Zickler. 2010. A coaxial optical scanner for synchronous acquisition of 3D geometry and surface reflectance. *ACM Trans. Graph.* 29, 4, Article 99 (jul 2010), 12 pages. <https://doi.org/10.1145/1778765.1778836>
- Yingwenqi Jiang, Jiadong Tu, Yuan Liu, Xifeng Gao, Xiaoxiao Long, Wenping Wang, and Yuexin Ma. 2023. GaussianShader: 3D Gaussian Splatting with Shading Functions for Reflective Surfaces. *arXiv:2311.17977* [cs.CV]
- Haian Jin, Isabella Liu, Peijia Xu, Xiaoshuai Zhang, Songfang Han, Sai Bi, Xiaowei Zhou, Zexiang Xu, and Hao Su. 2023. Tensor: Tensorial inverse rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 165–174.
- Kaizhang Kang, Minyi Gu, Cihui Xie, Xuanda Yang, Hongzhi Wu, and Kun Zhou. 2021. Neural Reflectance Capture in the View-Illumination Domain. *IEEE Transactions on Visualization and Computer Graphics* 29, 2 (2021), 1450–1462.
- Kaizhang Kang, Cihui Xie, Chengan He, Mingqi Yi, Minyi Gu, Zimin Chen, Kun Zhou, and Hongzhi Wu. 2019. Learning efficient illumination multiplexing for joint capture of reflectance and shape. *ACM Trans. Graph.* 38, 6 (2019), 165–1.
- Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 2023. 3D Gaussian Splatting for Real-Time Radiance Field Rendering. *ACM Transactions on Graphics* 42, 4 (July 2023). <https://repo-sam.inria.fr/fungraph/3d-gaussian-splatting/>
- Hendrik PA Lensch, Jan Kautz, Michael Goesele, Wolfgang Heidrich, and Hans-Peter Seidel. 2003. Image-based reconstruction of spatial appearance and geometric detail. *ACM Transactions on Graphics (TOG)* 22, 2 (2003), 234–257.
- Ruofan Liang, Huiting Chen, Chunlin Li, Fan Chen, Selvakumar Panneer, and Nandita Vijaykumar. 2023a. Envidr: Implicit differentiable renderer with neural environment lighting. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 79–89.
- Zhihao Liang, Qi Zhang, Ying Feng, Ying Shan, and Kui Jia. 2023b. GS-IR: 3D Gaussian Splatting for Inverse Rendering. *arXiv:2311.16473* [cs.CV]
- Haimin Luo, Min Ouyang, Zijun Zhao, Suyi Jiang, Longwen Zhang, Qixuan Zhang, Wei Yang, Lan Xu, and Jingyi Yu. 2024. GaussianHair: Hair Modeling and Rendering with Light-aware Gaussians. *arXiv preprint arXiv:2402.10483* (2024).
- Linjie Lyu, Ayush Tewari, Thomas Leimkuehler, Marc Habermann, and Christian Theobalt. 2022. Neural Radiance Transfer Fields for Relightable Novel-view Synthesis with Global Illumination. *arXiv:2207.13607* [cs.CV]
- Xiaohe Ma, Kaizhang Kang, Ruisheng Zhu, Hongzhi Wu, and Kun Zhou. 2021. Free-form scanning of non-planar appearance with neural trace photography. *ACM Transactions on Graphics (TOG)* 40, 4 (2021), 1–13.
- Xiaohe Ma, Xianmin Xu, Leyao Zhang, Kun Zhou, and Hongzhi Wu. 2023. OpenSVBRDF: A Database of Measured Spatially-Varying Reflectance. *ACM Trans. Graph.* 42, 6, Article 254 (dec 2023), 14 pages. <https://doi.org/10.1145/3618358>
- Dawid Malarz, Weronika Smolak, Jacek Tabor, Slawomir Tadeja, and Przemysław Spurek. 2023. Gaussian Splatting with NeRF-based Color and Opacity. *arXiv:2312.13729* [cs.CV]
- Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, Jonathan T. Barron, Ravi Ramamoorthi, and Ren Ng. 2020. NeRF: Representing Scenes as Neural Radiance Fields for View Synthesis. In *ECCV*.
- Jacob Munkberg, Jon Hasselgren, Tianchang Shen, Jun Gao, Wenzheng Chen, Alex Evans, Thomas Müller, and Sanja Fidler. 2022. Extracting Triangular 3D Models, Materials, and Lighting From Images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 8280–8290.
- Giljoo Nam, Joo Ho Lee, Diego Gutierrez, and Min H. Kim. 2018. Practical SVBRDF Acquisition of 3D Objects with Unstructured Flash Photography. *ACM Trans. Graph.* 37, 6, Article 267 (dec 2018), 12 pages.
- Julien Philip, Sébastien Morgenthaler, Michaël Gharbi, and George Drettakis. 2021. Free-viewpoint indoor neural relighting from multi-view stereo. *ACM Transactions on Graphics (TOG)* 40, 5 (2021), 1–18.
- Yohan Poirier-Ginter, Alban Gauthier, Julien Philip, Jean-François Lalonde, and George Drettakis. 2024. A Diffusion Approach to Radiance Field Relighting using Multi-Illumination Synthesis. *Computer Graphics Forum* (2024). <https://doi.org/10.1111/cgfg.15147>
- Shunsuke Saito, Gabriel Schwartz, Tomas Simon, Junxuan Li, and Giljoo Nam. 2023. Relightable Gaussian Codec Avatars. *arXiv:2312.03704* [cs.GR]
- Yahao Shi, Yanmin Wu, Chenming Wu, Xing Liu, Chen Zhao, Haocheng Feng, Jingtuo Liu, Liangjun Zhang, Jian Zhang, Bin Zhou, Errui Ding, and Jingdong Wang. 2023. GIR: 3D Gaussian Inverse Rendering for Relightable Scene Factorization. *arXiv:2312.05133* [cs.CV]
- Pratul P Srinivasan, Boyang Deng, Xiuming Zhang, Matthew Tancik, Ben Mildenhall, and Jonathan T Barron. 2021. Nerv: Neural reflectance and visibility fields for relighting and view synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 7495–7504.
- Tiancheng Sun, Kai-En Lin, Sai Bi, Zexiang Xu, and Ravi Ramamoorthi. 2021. Nelf: Neural light-transport field for portrait view synthesis and relighting. *arXiv preprint arXiv:2107.12351* (2021).
- Ayush Tewari, Justus Thies, Ben Mildenhall, Pratul Srinivasan, Edgar Tretschk, Wang Yifan, Christoph Lassner, Vincent Sitzmann, Ricardo Martin-Brualla, Stephen Lombardi, et al. 2022. Advances in neural rendering. In *Computer Graphics Forum*, Vol. 41. Wiley Online Library, 703–735.
- Borom Tunwattananong, Graham Fyffe, Paul Graham, Jay Busch, Xueming Yu, Abhijeet Ghosh, and Paul Debevec. 2013. Acquiring reflectance and shape from continuous spherical harmonic illumination. *ACM Transactions on graphics (TOG)* 32, 4 (2013), 1–12.
- Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. 2004. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing* 13, 4 (2004), 600–612.
- Lance Williams. 1978. Casting curved shadows on curved surfaces. In *Proceedings of the 5th annual conference on Computer graphics and interactive techniques*. 270–274.
- Hongzhi Wu, Zhaotian Wang, and Kun Zhou. 2015. Simultaneous localization and appearance estimation with a consumer RGB-D camera. *IEEE transactions on visualization and computer graphics* 22, 8 (2015), 2012–2023.
- Tong Wu, Yu-Jie Yuan, Ling-Xiao Zhang, Jie Yang, Yan-Pei Cao, Ling-Qi Yan, and Lin Gao. 2024. Recent Advances in 3D Gaussian Splatting. *arXiv:2403.11134* [cs.CV]
- Rui Xia, Yue Dong, Pieter Peers, and Xin Tong. 2016. Recovering shape and spatially-varying surface reflectance under unknown illumination. *ACM Trans. Graph.* 35, 6, Article 187 (dec 2016), 12 pages. <https://doi.org/10.1145/2980179.2980248>
- Kun Xu, Wei-Lun Sun, Zhao Dong, Dan-Yong Zhao, Run-Dong Wu, and Shi-Min Hu. 2013. Anisotropic spherical Gaussians. *ACM Trans. Graph.* 32, 6, Article 209 (nov 2013), 11 pages. <https://doi.org/10.1145/2508363.2508386>
- Xianmin Xu, Yuxin Lin, Haoyang Zhou, Chong Zeng, Yaxin Yu, Kun Zhou, and Hongzhi Wu. 2023. A Unified Spatial-Angular Structured Light for Single-View Acquisition of Shape and Reflectance. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 206–215.
- Ziyi Yang, Xinyu Gao, Yangtian Sun, Yihua Huang, Xiaoyang Lyu, Wen Zhou, Shaohui Jiao, Xiaojuan Qi, and Xiaogang Jin. 2024. Spec-Gaussian: Anisotropic View-Dependent Appearance for 3D Gaussian Splatting. *arXiv:2402.15870* [cs.CV]
- Keyang Ye, Qiming Hou, and Kun Zhou. 2024. 3D Gaussian Splatting with Deferred Reflection. (2024).
- Hong-Xing Yu, Michelle Guo, Alireza Fathi, Yen-Yu Chang, Eric Ryan Chan, Ruohan Gao, Thomas Funkhouser, and Jiajun Wu. 2023. Learning Object-centric Neural Scattering Functions for Free-viewpoint Relighting and Scene Composition. *Transactions on Machine Learning Research* (2023).
- Chong Zeng, Guojun Chen, Yue Dong, Pieter Peers, Hongzhi Wu, and Xin Tong. 2023. Relighting Neural Radiance Fields with Shadow and Highlight Hints. In *ACM SIGGRAPH 2023 Conference Proceedings*.
- Chong Zeng, Yue Dong, Pieter Peers, Youkang Kong, Hongzhi Wu, and Xin Tong. 2024. DiLightNet: Fine-grained Lighting Control for Diffusion-based Image Generation. In *ACM SIGGRAPH 2024 Conference Papers*.
- Kai Zhang, Fujun Luan, Qianqian Wang, Kavita Bala, and Noah Snavely. 2021a. 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*. 5453–5462.
- Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. 2018. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 586–595.
- Xiuming Zhang, Pratul P Srinivasan, Boyang Deng, Paul Debevec, William T Freeman, and Jonathan T Barron. 2021b. Nerfactor: Neural factorization of shape and reflectance under an unknown illumination. *ACM Transactions on Graphics (ToG)*

- 40, 6 (2021), 1–18.
- Youjia Zhang, Teng Xu, Junqing Yu, Yuteng Ye, Yanqing Jing, Junle Wang, Jingyi Yu, and Wei Yang. 2023. Nemf: Inverse volume rendering with neural microflake field. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 22919–22929.
- Zhenglong Zhou, Zhe Wu, and Ping Tan. 2013. Multi-view photometric stereo with spatially varying isotropic materials. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 1482–1489.




Fig. 8. Visualization of various components of our representation. From the left image to right, the final rendering result, diffuse albedo ρ_d , specular albedo ρ_s , normal \mathbf{n} , shadow function Φ and other effects Ψ ($\times 2.5$ for a better visualization). An image gamma of 1 is used for this figure.

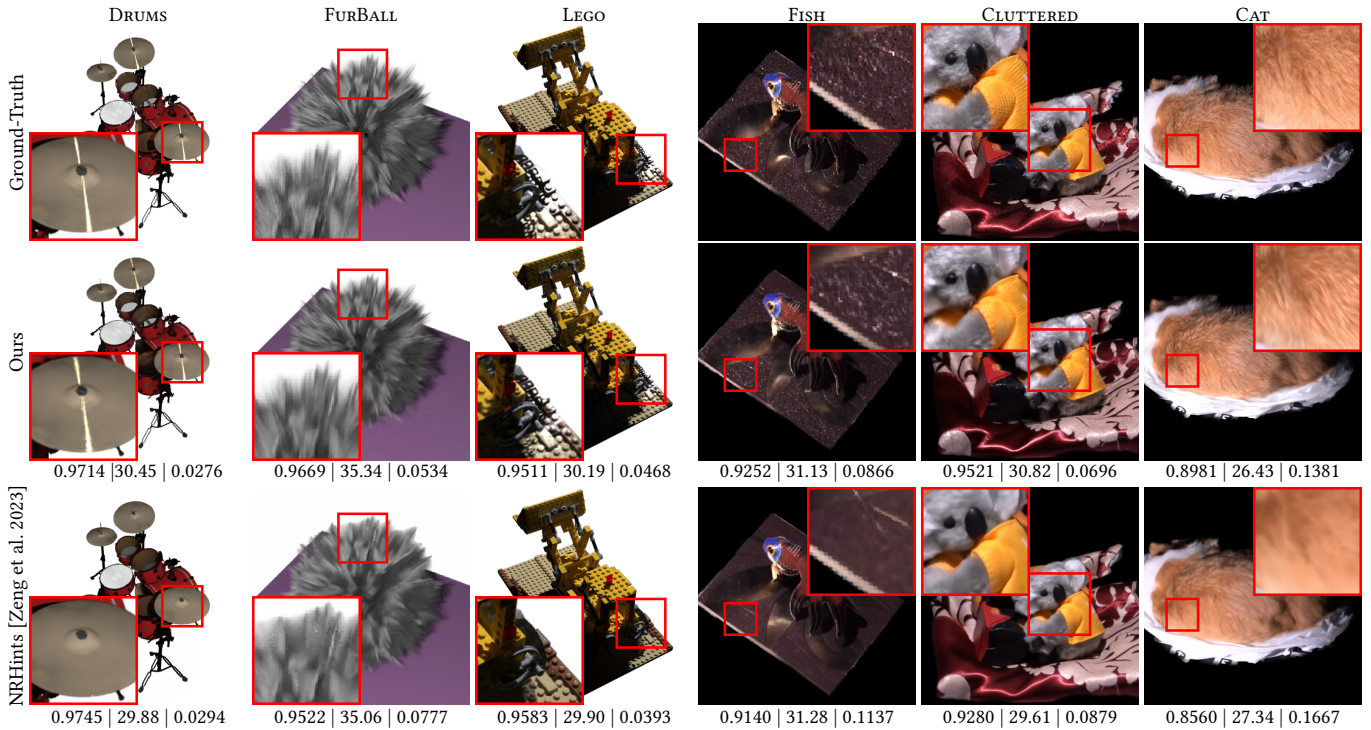


Fig. 9. Comparisons to [Zeng et al. 2023]. From the top row to bottom, the ground-truth, results of our approach and [Zeng et al. 2023], respectively. From the 1st column to 3rd, synthetic data from [Mildenhall et al. 2020]; from the 4th column to the last, captured data from [Zeng et al. 2023]. Average errors in SSIM, PSNR and LPIPS are reported at the bottom of each related image.




Fig. 10. Comparisons to [Lyu et al. 2022]. From every 3 consecutive images, the ground-truth, result with our approach and [Lyu et al. 2022], respectively. Average errors in SSIM, PSNR and LPIPS are reported at the bottom of each related image.




Fig. 11. Our relighting results on synthetic data/rendered images of captured data [Kang et al. 2019; Ma et al. 2023; Zeng et al. 2023]. For each pair of images, the left one is the ground-truth, and the right is our result. Average errors in SSIM, PSNR and LPIPS are reported at the bottom.




Fig. 12. Comparisons with approaches using environment-lit input images. For images from the left column to right in each row: the ground-truth, rendering results of our approach, GaussianShader [Jiang et al. 2023], GS-IR [Liang et al. 2023b], Relightable 3D Gaussian [Gao et al. 2023] and TensorIR [Jin et al. 2023], respectively. Average errors in SSIM, PSNR and LPIPS are reported at the bottom of each related image.