ReLightGS: Relighting Existing 3D Gaussian Splatting Scenes Without Retraining

Kaan Altıntaş[1] Responsible Professor: Michael Weinmann [1], Examiner: Emir Demirović [1]

[1]EEMCS, Delft University of Technology

Project Page

1. Introduction

Motivation:

- Existing 3D Gaussian Splatting (3DGS) [2] models bake the capture lighting into the learned Gaussian representation.
- Relightable 3DGS methods usually need the original images and training optimization.
- Impractical when only a trained 3DGS scene is available.

Our Contribution:

- We relight already trained 3DGS scenes directly from the fixed Gaussian representation.
- From the fixed Gaussian representation, we estimate depth, normals, materials, and visibility, then combine these cues in a deferred lighting pass.

2. Background & Related Work

Background:

- Novel-view synthesis methods such as NeRF [1]and 3D Gaussian Splatting (3DGS) [2] learn scene representations from captured images to render the scene from new viewpoints.
- 3DGS is useful for interactive graphics because it represents the scene as explicit anisotropic Gaussians and renders them with rasterization at real-time frame rates.

Related Work:

- Existing relightable NeRF and 3DGS methods methods such as NeRV [3], GS-IR [5], improve this by estimating reflectance, normals, materials, and visibility.
- However, they require the original multi-view images and additional optimization during training.

In contrast, our method works in a more constrained setting: only an already trained, exported .ply Gaussian scene is available and without any retraining performed.

4. Quantitative Results

- Compared against inverse-rendering baselines
- Ours method ranks second
- **Exceeds GS-IR** on all metrics (MAE, PSNR, SSIM, LPIPS)

	Method	Normal MAE↓	Relight PSNR↑	Relight SSIM↑	Relight LPIPS↓
Neural field based rendering	NeRFactor [24]	6.314	23.383	0.908	0.131
	InvRender [25]	5.074	23.973	0.901	0.101
	NVDiffrec [16]	6.078	19.880	0.879	0.104
	TensoIR [10]	4.100	28.580	0.944	0.081
Gaussian splat based rendering	GS-IR [13]	4.948	24.374	0.885	0.096
	Ours	4.780	25.920	0.906	0.095

Variant	Normal MAE↓	Relight PSNR↑	Relight SSIM↑	Relight LPIPS↓
w/o material labels	4.78	24.58	0.8836	0.1106
w/ material labels	4.78	25.92	0.9060	0.0950
Improvement	–	+1.34	+0.0224	−0.0156

With labels improves relighting quality ↑ PSNR & SSIM, ↓ LPIPS

3. Methodology

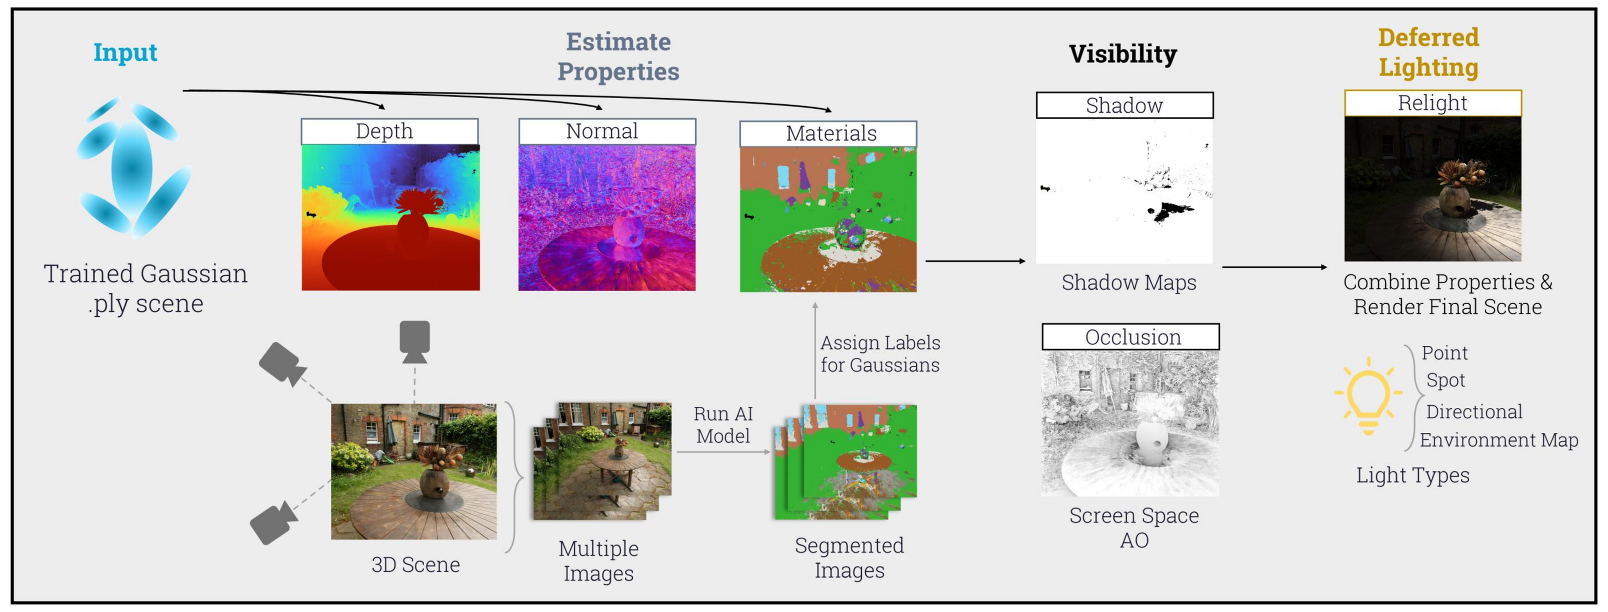

3.1 Depth & Normals

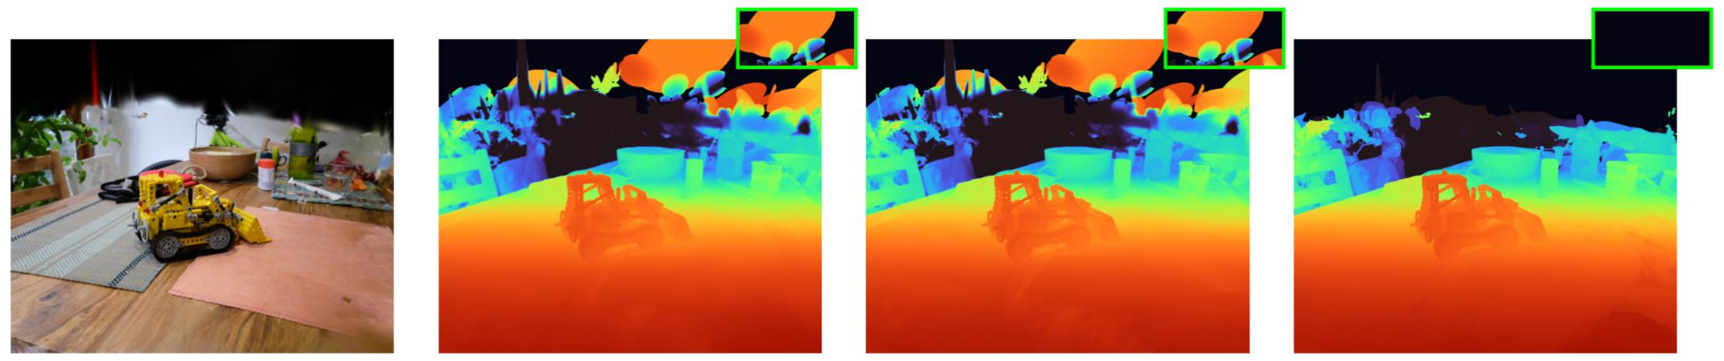

(a) RGB view (b) Weighted (c) Projected-plane (d) Iso-opacity

- **Weighted depth:** Averages Gaussians Depths per pixel.
- **Projected-plane depth:** Treats each splat as a local plane.
- **Iso-opacity depth:** Picks the depth where accumulated opacity crosses a threshold.

We use *iso-opacity depth* → sharper surfaces, fewer floaters

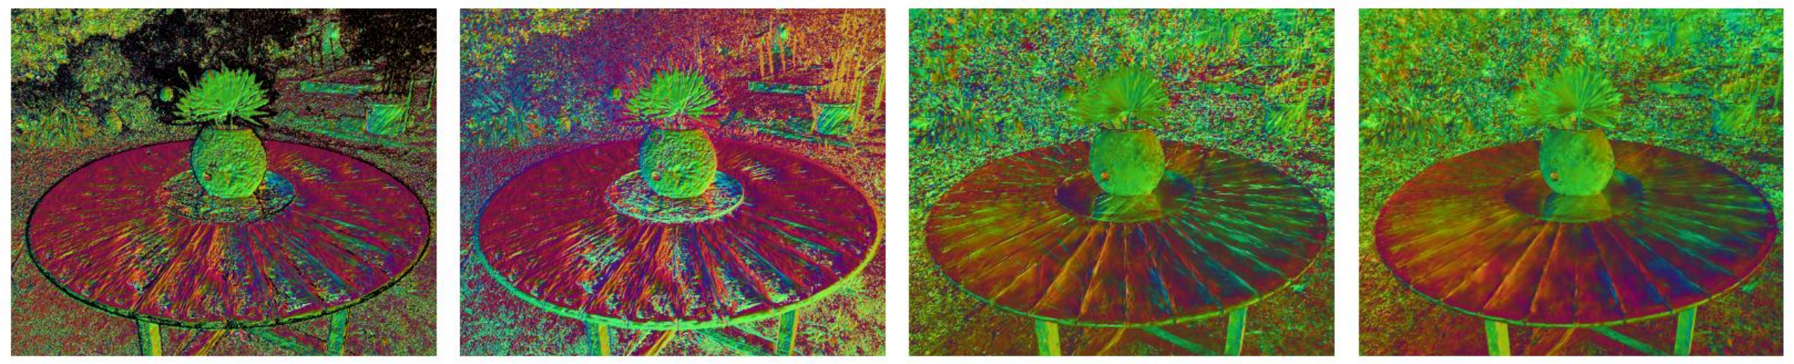

(a) Cross (b) Plane-fit (c) Projected-plane (d) Max density

- **Cross-product:** Normals from neighboring depth pixels; noisy at edges
- **Plane-fit:** Fits a small plane to nearby depth points.; smoother
- **Projected-plane:** Per-Gaussian depth plane; some artifacts
- **Max-density plane:** Computed analytically from each Gaussian's covariance

We use *max-density plane normals* → coherent, stable

5. Qualitative Results

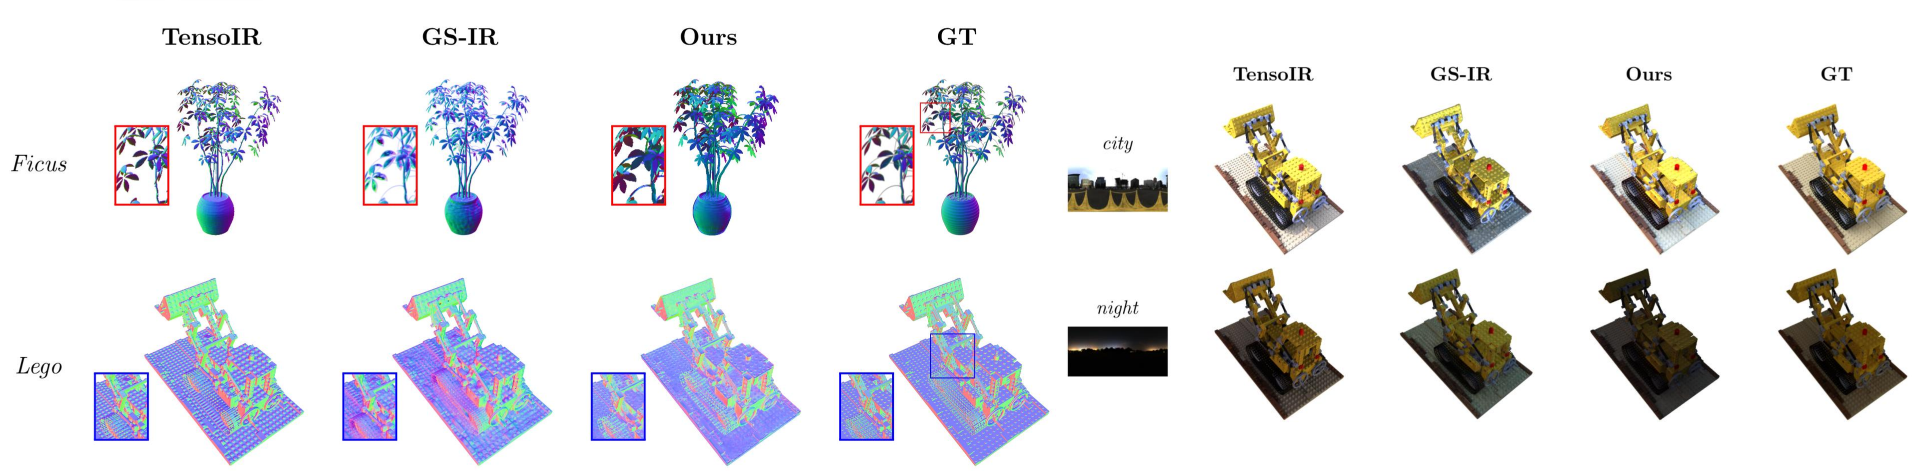

- GS-IR shows artifacts and normal-direction errors
- Ours keeps surface orientation more consistent than GS-IR
- Ours recovered normals slightly thick around fine details

3.2 Material Segmentation Model

(a) Input View (b) Labeled View

- **Dense Material Segmentation (DMS):** predicts material labels from an image
- From multiple segmented image, label assigned to Gaussians
- Mapped to Physically Based Rendering (PBR) params: *diffuse, specular, roughness, metallic*

3.3 Visibility and Occlusion

(a) Input render (b) AO visibility

- Shadows are computed with light-view depth maps, while SSAO uses reconstructed depth and normals to darken contact regions and concavities.

3.4 Deferred Lighting

(a) Input render (b) Relit result

- **Combine all properties:** depth, normals, color, materials, shadows, ambient occlusion
- **Lighting:** point, spot, directional and environment lighting

6. Conclusions & Future Work

Key Achievements

- Relighting with no retraining from original images
- Plausible visual quality
- Interactive lighting → fast visual experimentation

Limitations

- Original lighting partly baked into Gaussian colors
- Material labels cause local errors and bleeding

Future Work

- Better estimation of albedo (true surface color) to remove baked lighting
- Boundary-aware material segmentation to reduce local errors

7. References

[1] B. Mildenhall, P. P. Srinivasan, M. Tancik, J. T. Barron, R. Ramamoorthi, and R. Ng. NeRF: Representing Scenes as Neural Radiance Fields for View Synthesis. ECCV, 2020.
[2] B. Kerbl, G. Kopanas, T. Leimkühler, and G. Drettakis. 3D Gaussian Splatting for Real-Time Radiance Field Rendering. ACM TOG, 2023.
[3] P. P. Srinivasan, B. Deng, X. Zhang, M. Tancik, B. Mildenhall, and J. T. Barron. NeRV: Neural Reflectance and Visibility Fields for Relighting and View Synthesis. CVPR, 2021.
[4] J. Gao, C. Gu, Y. Lin, Z. Li, H. Zhu, X. Cao, L. Zhang, and Y. Yao. Relightable 3D Gaussians: Realistic Point Cloud Relighting with BRDF Decomposition and Ray Tracing. ECCV, 2024.
[5] Z. Liang, Q. Zhang, Y. Feng, Y. Shan, and K. Jia. GS-IR: 3D Gaussian Splatting for Inverse Rendering. CVPR, 2024.
[6] T. Wu, J.-M. Sun, Y.-K. Lai, Y. Ma, L. Kobbelt, and L. Gao. DeferredGS: Decoupled and Editable Gaussian Splatting with Deferred Shading. 2024.
[7] Y. Jiang, J. Tu, Y. Liu, X. Gao, X. Long, W. Wang, and Y. Ma. GaussianShader: 3D Gaussian Splatting with Shading Functions for Reflective Surfaces. 2023.
[8] P. Upchurch and R. Niu. A Dense Material Segmentation Dataset for Indoor and Outdoor Scene Parsing. ECCV, 2022.
[9] B. Zhang, C. Fang, R. Shrestha, Y. Liang, X. Long, and P. Tan. RaDe-GS: Rasterizing Depth in Gaussian Splatting. ACM TOG, 2026.
[10] NVIDIA. Vulkan Gaussian Splatting: Lighting and Shadows. nvpro-samples documentation, 2026.
[11] S. Kheradmand, D. Vicini, G. Kopanas, D. Lagun, K. M. Yi, M. Matthews, and A. Tagliasacchi. StochasticSplats: Stochastic Rasterization for Sorting-Free 3D Gaussian Splatting. ICCV, 2025.

kaltiintas@student.tudelft.nl