

LumiGauss: Relightable Gaussian Splatting in the Wild

Supplementary Material

Joanna Kaleta^{1,2*} Kacper Kania¹ Tomasz Trzcinski^{1,4,5} Marek Kowalski³

¹Warsaw University of Technology ²Sano Centre for Computational Medicine ³Microsoft ⁴IDEAS NCBR ⁵Tooploox

A. Dataset Processing

Occluders. To exclude occluders from training images we use masks provided with OSR dataset [4].

Test set. We test our approach on 5 viewpoints for each scene, as it was originally proposed in [4]. For testing, we use test masks provided by [4] and we strictly follow their evaluation protocol. For SSIM, we report the average value over the segmentation mask, utilizing the scikit-image implementation with a window size of 5 and eroding the segmentation mask by the same window size to exclude the influence of pixels outside the mask on the metric value.

Testing with ground truth environment map. The authors of [2] made an effort to recover steps for environment map preprocessing and alignment. The preprocessing step is available in their repository, accessible at: https://github.com/JADGardner/neusky/blob/main/notebooks/nerfosr_envmaps.ipynb. The detailed discussion on SOL-NeRF [5] approach to environment map alignment is included in the NeuSky main paper [2] and also confirmed with SOL-NeRF authors.

B. Implementation details

The appearance embedding vector is set to a size of 24 dimensions. For predicting the environment map, we use MLP with 3 fully-connected layers of size 64. We trained all models for 40000 iterations, the first training stage is set to 20000 iterations. The learning rate for MLP and embedding is set to 0.002, which after first training stage is reduced to 0.0002. We train gaussian spherical harmonics with a learning rate of 0.002. We set the loss function weights as follows: for ℓ_{0-1} $\lambda_1 = 0.001$, for ℓ_+ $\lambda_2 = 0.05$, for $\ell_{\bullet \leftrightarrow \circ}$ $\lambda_3 \in \{1.0, 10.0\}$, for ℓ_{\bullet} $\lambda_4 = 10.0$. In the second training stage we set $\lambda_{\bullet} = 0.001$.

We adhere to the original Gaussian splatting densification and pruning protocols, with a densification interval of 500 iterations and an opacity reset interval of 3000 iterations. We apply regularizations to align Gaussians with surfaces, as originally described in [3]. Additionally, we utilize

the dual visibility concept proposed in [3]. This ensures that the Gaussians are always correctly oriented towards the camera. Dual visibility effectively produces consistent world normals, with visible normals being consistent and non-visible ones contributing minimally to the rendering. Regularization of Spherical Harmonics \mathbf{d}_k is dependent on gaussian normals. Since normals are rotated to always face the camera, to maintain alignment between each Gaussian’s normal and its associated \mathbf{d}_k , we also rotate \mathbf{d}_k accordingly.

We run all experiments using a single NVIDIA A100 80GB or RTX 2080 Ti 128 GB.

C. Relighting - additional results

Please reach for additional results to the attached videos.

D. Qualitative comparison - additional results

In Fig. 1 we show the qualitative comparison of our method, NeRF-OSR, and SR-TensoRF. We show the landmark relit with ground truth environment map for NeRF-OSR and LumiGauss. SR-TensoRF reconstructs ground truth using only daytime (timestamp).

In Fig. 3, we show the qualitative comparison of our method, NeRF-OSR, and SR-TensoRF. We use the *default synthetic* environment map provided by [4]. This environment map was used for visualisation purposes in [1]. We use it to ensure a fair comparison and consistency with results from concurrent works. We also present albedo and normals extracted from the reconstructed scene. Please note that our model produces much cleaner results. Compared to the baselines, it reconstructs sharp features in small elements of the buildings, which is also reflected in the quantitative results ???. LumiGauss also gracefully smooths out the elements of scenes that are variable across the images, such as trees and clouds. On the other hand, NeRF-OSR and SR-TensoRF produce artifacts that negatively impact the output reconstructions.

In Fig. 2 we present additional comparison with concurrent works. We focus on normal and albedo quality.

In Fig. 4 we present additional results of novel view syn-

*Corresponding authors: joanna.kaleta.dokt@pw.edu.pl

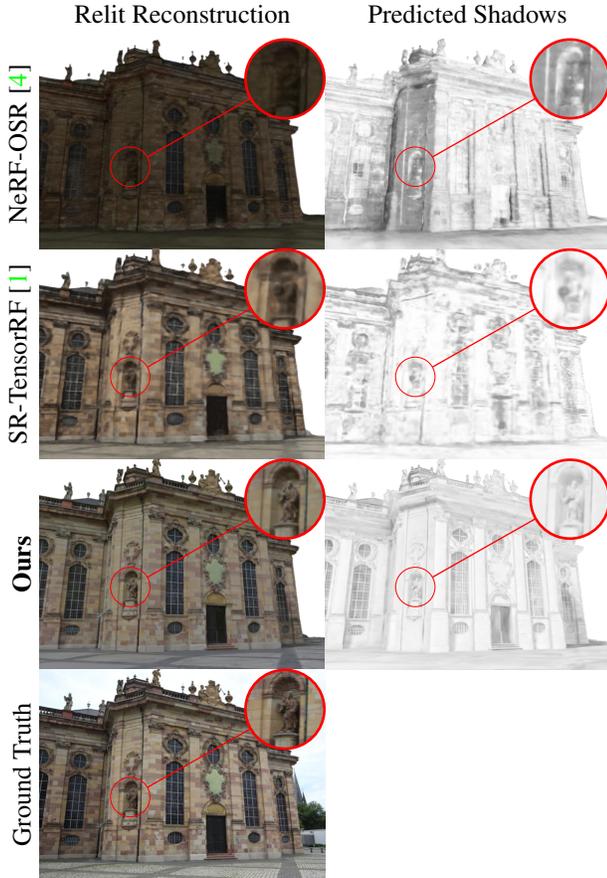


Figure 1. **Qualitative comparison of scene reconstruction for the selected photo session.** Results for NeRF-OSR [4] and LumiGauss were generated using ground truth environment maps for selected photo session, while ST-TensorRF [1] used extracted timestamp. Results for NeRF-OSR and SR-TensorF reported originally in [1].

thesis and comparison with concurrent works. Similarly to NeRF-OSR, we relight our scenes with the **default synthetic** map provided by NeRF-OSR for visualization purposes. This environment map does not correspond to GT images.

E. Ablation study - additional results

In Fig. 5 we present renders from training without selected regularization terms.

References

- [1] Yeonjin Chang, Yearim Kim, Seunghyeon Seo, Jung Yi, and Nojun Kwak. Fast sun-aligned outdoor scene relighting based on tensorf. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 3626–3636, 2024. 1, 2, 3, 4
- [2] James AD Gardner, Evgenii Kashin, Bernhard Egger, and William AP Smith. The sky’s the limit: Relightable outdoor scenes via a sky-pixel constrained illumination prior and outside-in visibility. In *European Conference on Computer Vision*, pages 126–143. Springer, 2025. 1, 3
- [3] Binbin Huang, Zehao Yu, Anpei Chen, Andreas Geiger, and Shenghua Gao. 2d gaussian splatting for geometrically accurate radiance fields. In *ACM SIGGRAPH 2024 Conference Papers*, pages 1–11, 2024. 1
- [4] Viktor Rudnev, Mohamed Elgharib, William Smith, Lingjie Liu, Vladislav Golyanik, and Christian Theobalt. Nerf for outdoor scene relighting. In *European Conference on Computer Vision*, pages 615–631. Springer, 2022. 1, 2, 3
- [5] Jia-Mu Sun, Tong Wu, Yong-Liang Yang, Yu-Kun Lai, and Lin Gao. SOL-NeRF: Sunlight Modeling for Outdoor Scene Decomposition and Relighting. In *SIGGRAPH Asia 2023 Conference Papers*, SA ’23, New York, NY, USA, 2023. Association for Computing Machinery. 1



Figure 2. **Qualitative comparison of predicted albedo and rendered normals.** Results for NERF-OSR, FEGR, SOL-NeRF, NeuSky reported originally in [2].

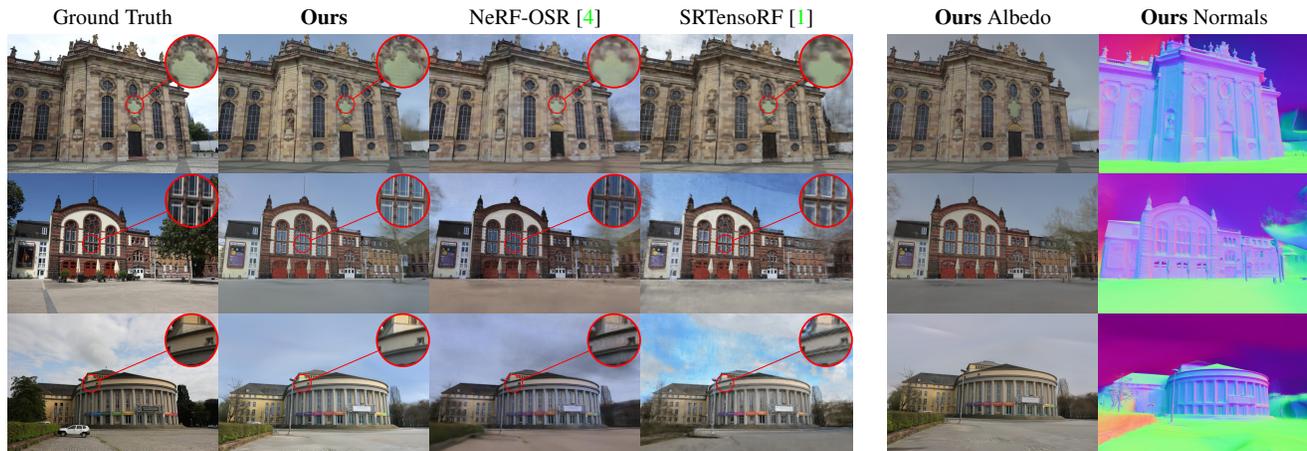


Figure 3. **Qualitative results** – Showcase of novel view synthesis using shadowed radiance transfer. We present albedo and normals produced by our method. Our method generates much sharper renderings. Please see zoom-ins to see details on the quality difference, such as surface smoothness and edge sharpness of small building elements. We use visual results for SR-TenSoRF and NeRF-OSR presented originally in [1]. Please note that, in this comparison the environment map used to create renders for NeRF-OSR and LumiGauss **does not match** the illumination in ground truth. LumiGauss and NeRF-OSR employ the **default** environment map provided by NeRF-OSR **for clear visualisation purpose only**. SR-TenSoRF do not rely on any environment map, instead it utilizes daytime information.





Figure 5. **Ablation study for relighting with external environment map.** The full model results in the clearest render. The strongest quality drop is observed when components restricting D_k are omitted.