Supplementary Materials for GaRe: Relightable 3D Gaussian Splatting for Outdoor Scenes from Unconstrained Photo Collections

Haiyang Bai[1] Jiaqi Zhu[1] Songru Jiang[1] Wei Huang[1] Tao Lu[2]
Yuanqi Li[1] Jie Guo[1] Runze Fu[3] Yanwen Guo[1] Lijun Chen[1,*]

[1]Nanjing University [2]Brown University [3]JSTI Group

https://baihyyut.github.io/GaRe/

1. Dataset Preprocess

Similar to the original Gaussian model, our method utilizes a dataset of pose-calibrated outdoor images and sparse point clouds generated via Structure-from-Motion (SfM). Additionally, we enhance the preprocessing step by extracting sky masks $\mathcal{M}_{sky}$ and coarse sun visibility $\mathcal{V}_c$ for all training images, which are used in Sec. 3.2 and Sec. 3.3. The detailed preprocessing steps are outlined below.

1.1. Sky Mask

The sky mask is a binary mask that represents the sky and non-sky regions. To generate an accurate sky mask, we leverage the pre-trained Depth Anything V2 [4] model to estimate the depth maps for all photo collections, with predictions encoded as relative depth disparities. Given the empirical assumption that the sky lies at an infinite distance, it is characterized by near-zero disparity values. Thus, we adopt a disparity threshold of $\tau_D = 0.1$, categorizing pixels as sky (1) or non-sky (0).

1.2. Coarse Sun Visibility

Sun visibility is represented as a binary mask indicating whether sunlight can directly reach the surface of the scene. To extract sun visibility, we first pre-classify all photo collections into sunny and cloudy weather conditions based on [5]. Sun visibility extraction is performed only on sunny images, while for cloudy images, sun visibility is manually set to zero, indicating complete occlusion of the sun. For the sunny images, we convert the RGB images $\mathcal{I}$ into the HSV color space, where the V channel $\mathcal{I}_v$ represents brightness. Sunlit regions exhibit higher brightness values, while shaded areas have lower values. To enhance the contrast between sunlit and non-sunlit regions, we apply a gamma transformation as follows:

$$\mathcal{I}'_V = (\beta \cdot \mathcal{I}_V - \epsilon)^{\gamma}. \tag{1}$$

Here, β, ϵ, and γ are the scaling coefficient, offset coefficient, and gamma transformation exponent, respectively, with values set to $\frac{1}{255}$, 0.1 and 1.5. Subsequently, we apply a brightness threshold of $\tau_V = 0.3$ on $\mathcal{I}'_v$ to classify the regions into two categories: areas with higher brightness values are labeled as sun-visible, while those with lower values are marked as sun-invisible.

2. Sky Semantics

The two-step filtering strategy effectively mitigates the misleading influence of floating Gaussians in the sky on surface visibility. However, it introduces inaccuracies in the visibility of the sky region. To resolve this, we define a sky semantic $\boldsymbol{o}_k \in \mathbb{R}^+$ for each Gaussian and transform the optimization as a binary classification task. This is done by by minimizing the cross-entropy loss (CEL) between the rendered sky mask $\hat{\mathcal{M}}_{sky}$, generated by splatting the sky semantic, and the ground truth sky mask $\mathcal{M}_{sky}$, obtained through Sec. 1.1. The loss function is:

$$\mathcal{L}_{sem} = \text{CEL}(\mathcal{M}_{sky}, \hat{\mathcal{M}}_{sky}). \tag{2}$$

Once the sky semantics are optimized, the framework can render the sky mask $\hat{\mathcal{M}}_{sky}$ from any arbitrary viewpoint. This capability facilitates continuous guidance for defining the sky region that remains consistently visible to the sun.

3. Cloudy Day

In our experimental dataset, there are several instances representing scenarios without direct sunlight, $i.e.$, cloudy days. In these cases, the effect of sunlight is disregarded, and the global illumination for each viewpoint is decomposed into sky shading and indirect shading, as shown in Fig. 1.

* Corresponding author.

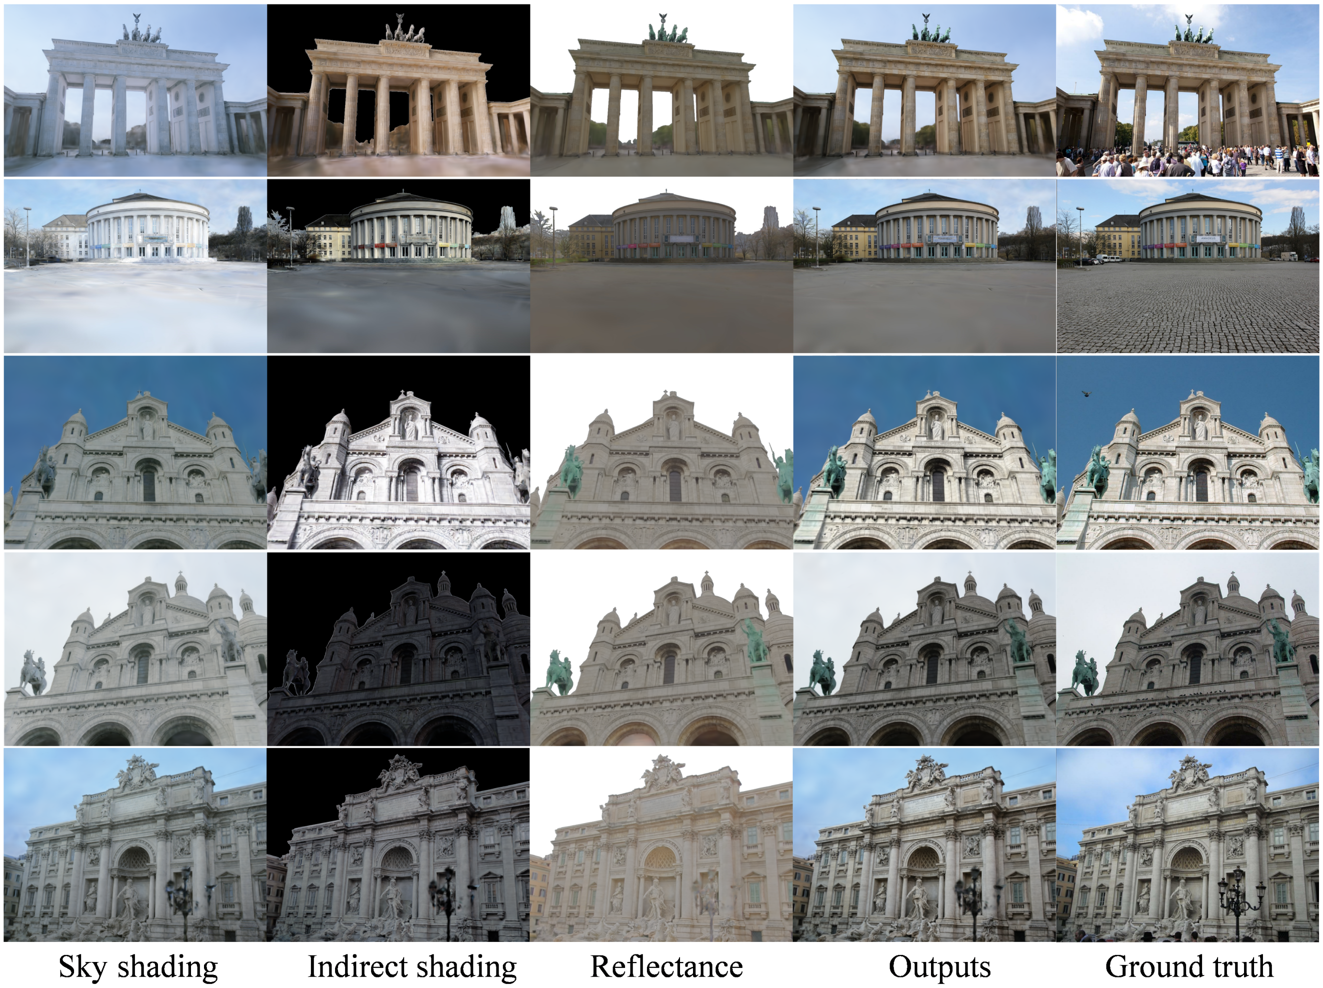

Figure 1. Visualization of the illumination decomposition results under cloudy conditions, where direct sunlight is absent. The scene is divided into sky shading, indirect shading, and reflectance. The sky shading models diffuse radiance from the overcast sky, while indirect shading captures global light transport within the scene. The reflectance preserves material properties independent of lighting. Together, these components enable a physically consistent and accurate reconstruction of the overall illumination in cloudy conditions.

4. Limitations and Future Work

Our method is not without limitations. Firstly, due to the phased training process in our framework, the overall training duration is relatively longer compared to state-of-the-art methods such as 3DGS [2] and SWAG [1]. However, it remains superior to Wildgaussians [3] and GS-W [6], as demonstrated in Table 1. Second, achieving realistic shadow effects is highly dependent on accurate scene geometry. Although our approach effectively filters out Gaussians with ambiguous occlusion relationships in the sky and removes redundant Gaussians caused by transient objects, challenges remain in constructing precise scene surfaces. This is particularly true for areas like the ground, where the lack of texture leads to sparse Gaussian distributions, complicating the accurate extraction of sharp shadow contours. Lastly, our current framework is unable to transfer illumination from unseen images, which is a promising direction for future research.

References

[1] Hiba Dahmani, Moussab Bennehar, Nathan Piasco, Luis Roldao, and Dzmitry Tsishkou. Swag: Splatting in the wild images with appearance-conditioned gaussians. In *ECCV*, pages 325–340. Springer, 2025. 2

[2] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splatting for real-time radiance field rendering. *ACM TOG*, 42(4):139–1, 2023. 2

[3] Jonas Kulhanek, Songyou Peng, Zuzana Kukelova, Marc Pollefeys, and Torsten Sattler. Wildgaussians: 3d gaussian splatting in the wild. *arXiv preprint arXiv:2407.08447*, 2024. 2

[4] Lihe Yang, Bingyi Kang, Zilong Huang, Zhen Zhao, Xiaogang Xu, Jiashi Feng, and Hengshuang Zhao. Depth anything v2. *arXiv preprint arXiv:2406.09414*, 2024. 1

[5] Piaopiao Yu, Jie Guo, Fan Huang, Cheng Zhou, Hongwei Che, Xiao Ling, and Yanwen Guo. Hierarchical disentangled representation learning for outdoor illumination estimation and editing. In *ICCV*, pages 15313–15322, 2021. 1

[6] Dongbin Zhang, Chuming Wang, Weitao Wang, Peihao Li, Minghan Qin, and Haoqian Wang. Gaussian in the wild: 3d gaussian splatting for unconstrained image collections. In *ECCV*, pages 341–359. Springer, 2025. 2