IRGS: Inter-Reflective Gaussian Splatting with 2D Gaussian Ray Tracing

Supplementary Material

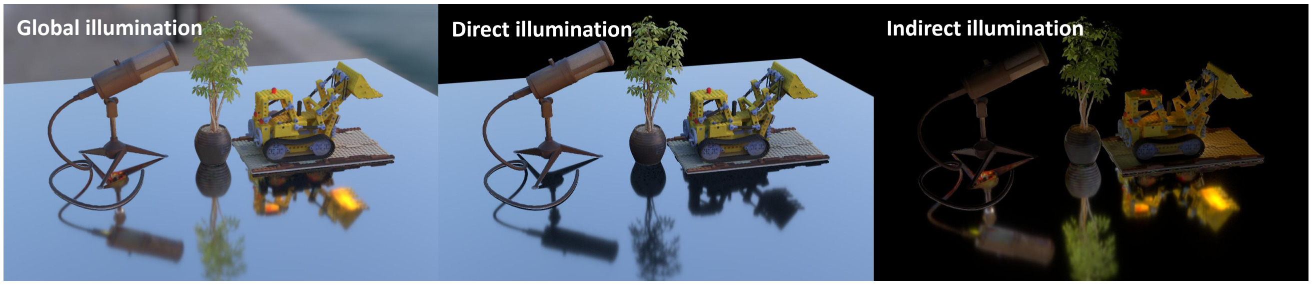

Figure 9. Global, direct, and indirect illumination in a Gaussian-based scene using our *IRGS*.

6. Relighting details

Given a sampled incident direction, we conduct 2D Gaussian ray tracing to obtain the intersected Gaussians along the ray, and aggregate the albedo, roughness, and normal by alpha-blending: $\{\mathcal{A}_r, \mathcal{R}_r, \mathcal{N}_r\} = \sum_i \omega_i \{\boldsymbol{a}_i, r_i, \boldsymbol{n}_i\}$, where $\omega_i = \frac{T_\mathrm{i}\alpha_\mathrm{i}}{\sum_\mathrm{i} T_\mathrm{i}\alpha_\mathrm{i}}$. Then, to avoid the extensive Monte Carlo sampling, we pre-integrate the cubemap, which allows us to obtain the diffuse L_d and specular term L_s for incident direction using only a single query. Specifically, we apply split-sum approximation for the specular term:

$$L_s \approx \int_\Omega f_s(\boldsymbol{\omega}_i, \boldsymbol{\omega}_o)(\boldsymbol{\omega}_i \cdot \mathcal{N}_r) d\boldsymbol{\omega}_i \cdot \int_\Omega L_i(\boldsymbol{\omega}_i) D(\boldsymbol{\omega}_i, \boldsymbol{\omega}_o)(\boldsymbol{\omega}_i \cdot \mathcal{N}_r) d\boldsymbol{\omega}_i, \quad (18)$$

where the left term depends solely on $(\omega_i \cdot \mathcal{N}_r)$ and roughness R_r, this allows the results to be pre-integrated and stored in a 2D lookup texture map. The right term represents the integral of incident radiance, which can also be pre-integrated before rendering. Consequently, the indirect radiance of the incident ray is given as: $L_\mathrm{ind} = L_\mathrm{d} + L_\mathrm{s}$.

7. Composited scene details

In Fig. 1 and Fig. 9, we relight a scene composed of four Gaussian-based objects: "Mic," "Ficus," "Lego," and "Ground." The objects "Mic," "Ficus," and "Lego" are reconstructed using the proposed *IRGS* framework, while "Ground" is manually designed using a set of parallel 2D Gaussians. To better illustrate the inter-reflective properties of *IRGS*, we assign an additional metallic property m to each Gaussian. Specifically, we set $m = 0$ for "Mic", "Ficus", and "Lego", and $m = 1$ for "Ground". For efficient rendering, we employ importance sampling with 512 rays distributed using cosine-weighted sampling for the diffuse term and 256 rays distributed using GGX sampling for the specular term. In Fig. 9, "Direct Illumination" considers only the direct incident radiance, "Indirect Illumination" accounts for only the indirect radiance, and "Global Illumination" combines both direct and indirect radiance for full rendering.

8. More results

8.1. Results on Synthetic4Relight

We further provide a qualitative comparison on three additional scenes from the Synthetic4Relight dataset [43], including "air ballons" (Fig. 10), "hotdog" (Fig. 11), and "jugs" (Fig. 12). It is evident that our estimated material properties, environment lighting, and relighted images are the most realistic compared to the competitors, GS-IR [19] and R3DG [10]. In Fig. 13, we further visualize the different estimated components in novel view, illustrating *IRGS*'s ability to capture accurate inter-reflection. Additionally, in Fig. 14, we compare the rendered normal maps of each scene with those from GS-IR [19] and R3DG [10]. Thanks to the accurate geometry provided by 2D Gaussian primitives, our normal maps exhibit superior fidelity.

8.2. Results on TensoIR

We present a qualitative comparison of the rendered normal maps and estimated albedo maps in Fig. 15 and Fig. 16, respectively. *IRGS* produces smooth normal maps, and clean albedo maps with minimal shadow artifacts, attributed to the accurate modeling of incident radiance through 2D Gaussian ray tracing.

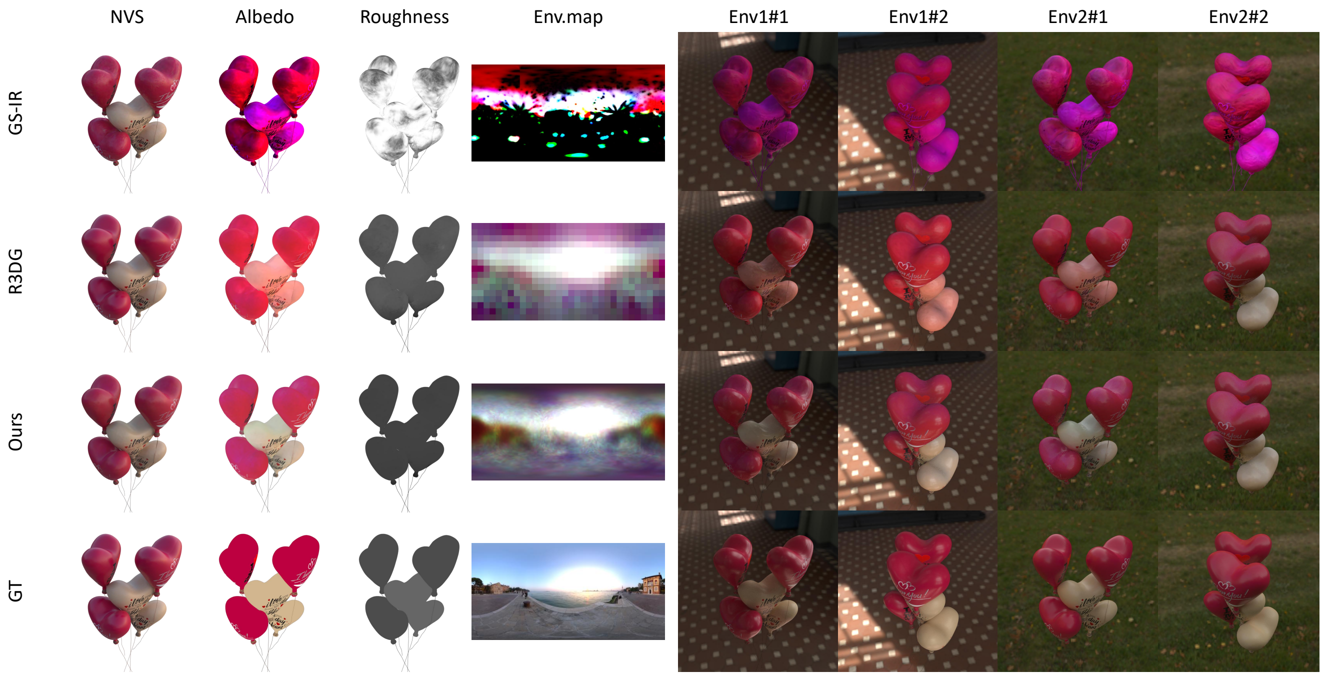

Figure 10. Qualitative comparison of NVS, material and lighting estimation, and relighting results on the Synthetic4Relight dataset [43].

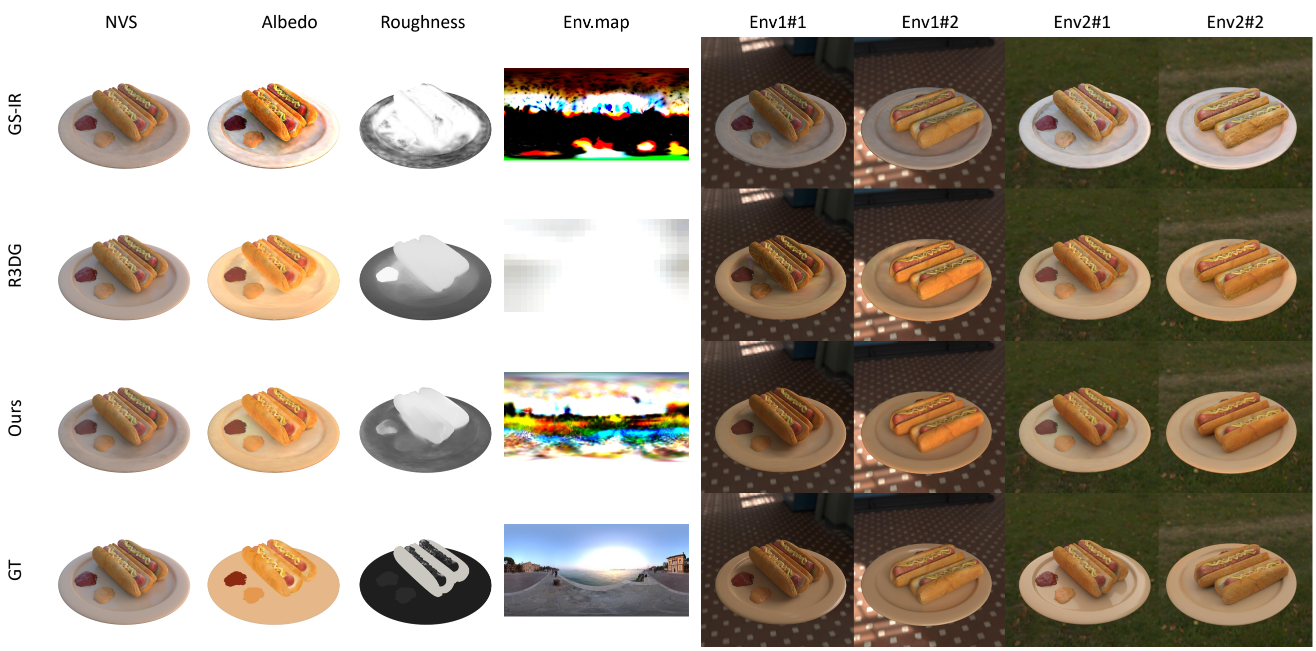

Figure 11. Qualitative comparison of NVS, material and lighting estimation, and relighting results on the Synthetic4Relight dataset [43].

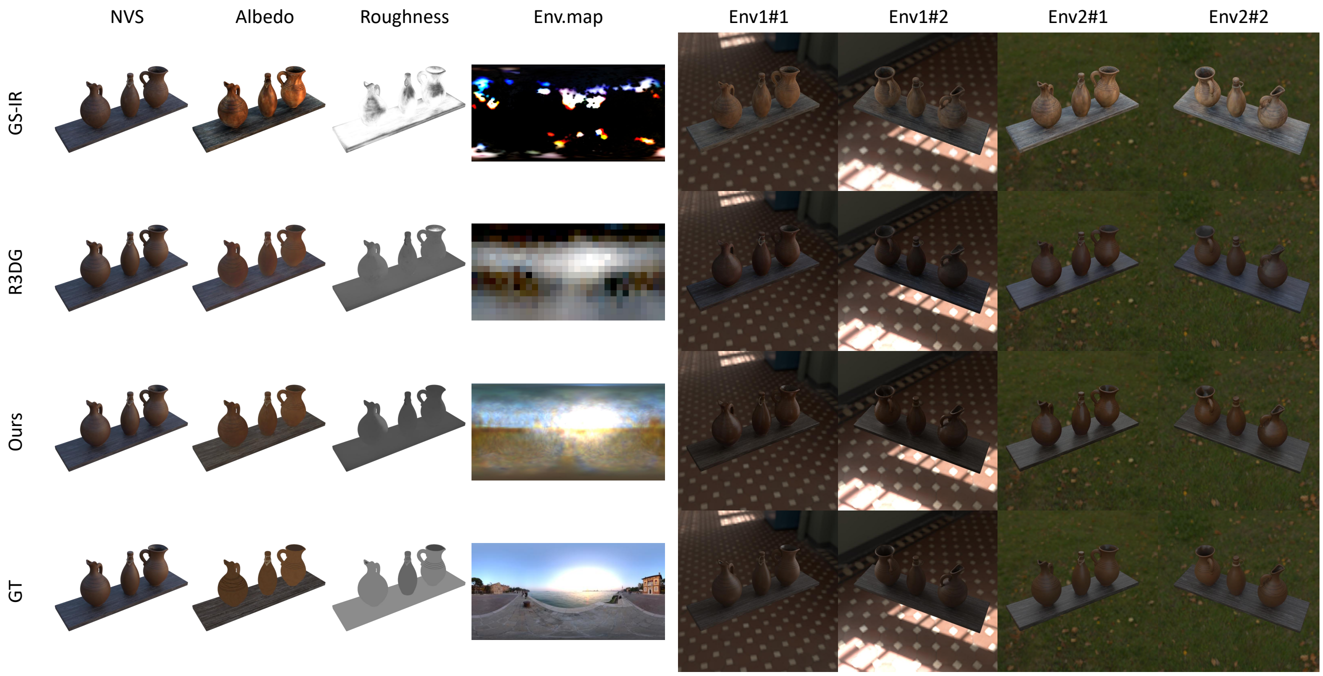

Figure 12. Qualitative comparison of NVS, material and lighting estimation, and relighting results on the Synthetic4Relight dataset [43].

Figure 13. Visualization of estimated components in novel view, including the averaged direct radiance L_{dir}, averaged indirect radiance L_{ind}, averaged visibility V (ambient occlusion), averaged incident radiance L_{i}, diffuse, specular, and final PBR color.

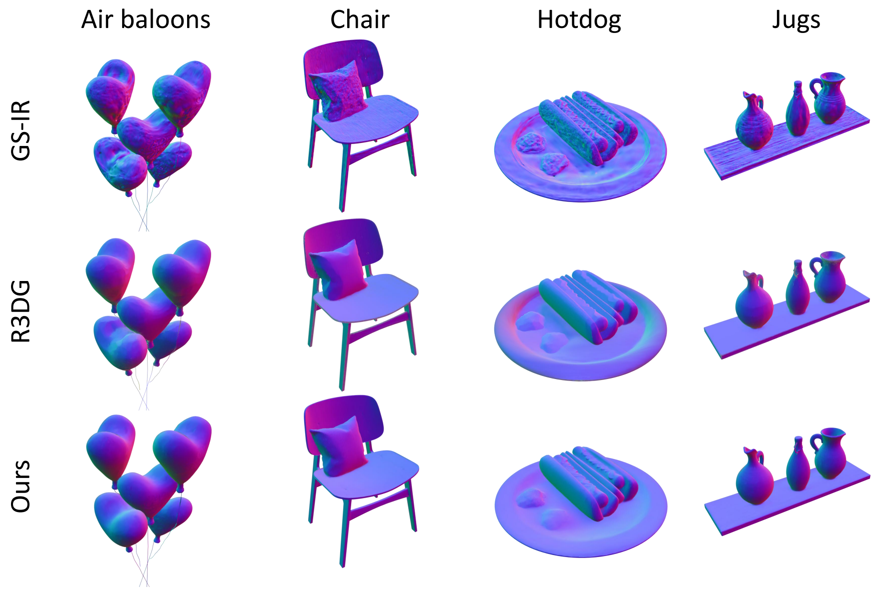

Figure 14. Qualitative comparison of rendered normal maps on the Synthetic4Relight dataset [43]. Note that, Synthetic4Relight dataset does not provide ground truth normal maps.

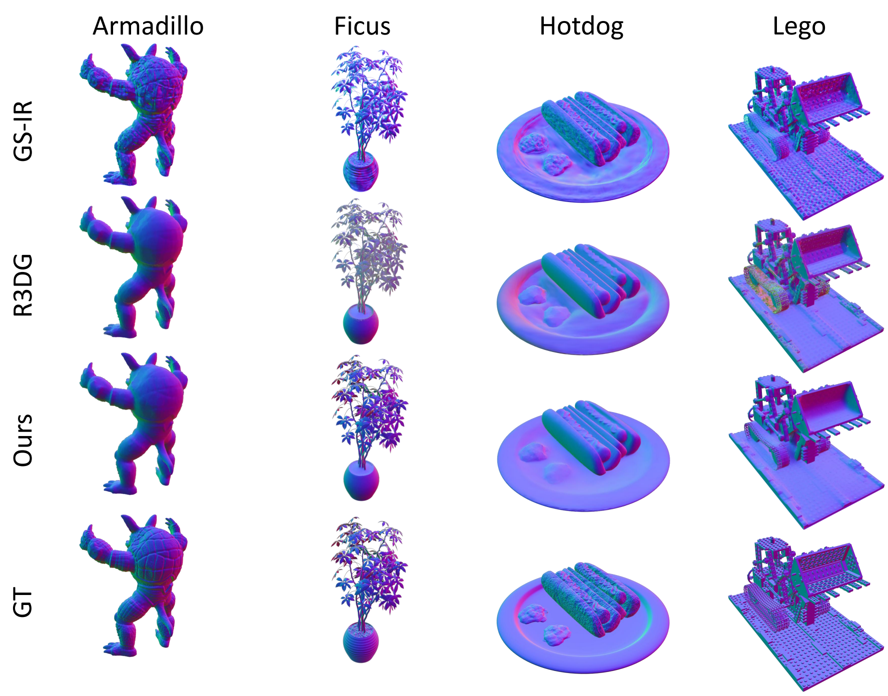

Figure 15. Qualitative comparison of rendered normal maps on the TensoIR dataset [15].

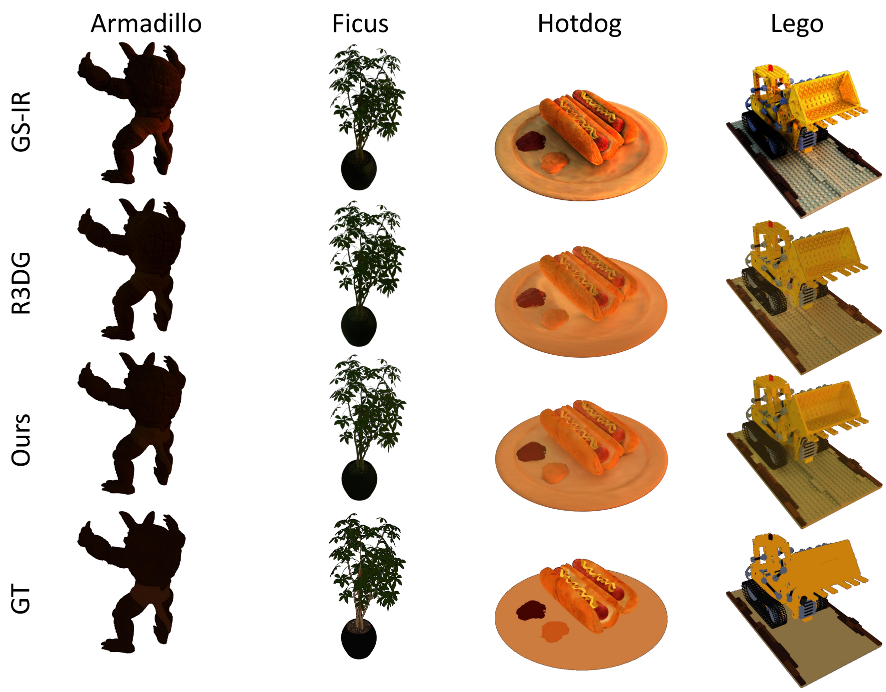

Figure 16. Qualitative comparison of estimated albedo maps on the TensoIR dataset [15]. Note that, we scale each RGB channel by a global scalar.