RT-Splatting: Joint Reflection-Transmission Modeling with Gaussian Splatting

Supplementary Material

In the supplementary material, we provide additional implementation details of our method (Sec. A). We also present an ablation study that examines the sensitivity of Specular-Aware Gradient Gating to the gating strength k (Sec. B). Finally, we report a per-scene breakdown of quantitative metrics and include further qualitative results (Sec. C).

A. Implementation Details

We strictly follow the evaluation protocols established in 3DGS-DR [43] and Ref-GS [50]. Consistent with these methods, we adopt the spherical domain strategy to define the foreground region of interest. Furthermore, we adopt their data preprocessing standards by applying consistent image downsampling factors across all experiments to ensure a fair comparison. All experiments are conducted on a single NVIDIA RTX 4090 GPU.

A.1. Density Control Strategy

Our density control strategy adapts the standard 2DGS [14] scheme to align with our proposed occupancy-opacity factorization. While standard 2DGS relies on opacity resets to regulate the number of primitives and prevent floaters, our decoupled representation requires managing the Gaussian primitives based on both geometric occupancy and optical opacity. We therefore introduce an interleaved reset schedule: instead of the standard opacity reset every 3,000 iterations, we perform resets every 1,500 iterations, alternating between geometric occupancy σ and optical opacity α. Furthermore, we perform pruning based on geometric occupancy σ rather than optical opacity α to prevent the erroneous removal of semi-transparent structures.

A.2. Specular-Aware Gradient Gating

For the local variance calculation in Eq. (5), we use a 3×3 window size. The hyperparameter k is set to 4 based on our sensitivity analysis in Sec. B. In practice, the gradient modulation described in Eq. (6) is implemented via a partial stop-gradient mechanism. Let $\mathtt{sg}(\cdot)$ denote the stop-gradient operator. The gated transmission color $\tilde{\mathbf{C}}_{\text{trans}}$ used for final composition is defined as:

$$\tilde{\mathbf{C}}_{\text{trans}} = (1 - g) \cdot \mathtt{sg}(\mathbf{C}_{\text{trans}}) + g \cdot \mathbf{C}_{\text{trans}} \,. \tag{8}$$

This formulation ensures $\tilde{\mathbf{C}}_{\text{trans}} \equiv \mathbf{C}_{\text{trans}}$ during the forward pass, while scaling the backward gradients by g.

A.3. Losses

Following 2DGS [14], we minimize the normal consistency loss $\mathcal{L}_{\text{n}}$ to enforce geometric alignment between rendered normals and depth gradients. For appearance supervision $\mathcal{L}_{\text{img}}$, we augment the standard $\mathcal{L}_1$ and D-SSIM [39] losses with the perceptual loss $\mathcal{L}_{\text{perc}}$ from EnvGS [41], with $\lambda = 0.2$ and $\lambda_{\text{perc}} = 0.01$:

$$\mathcal{L}_{\text{img}} = (1 - \lambda)\mathcal{L}_1 + \lambda\mathcal{L}_{\text{D-SSIM}} + \lambda_{\text{perc}}\mathcal{L}_{\text{perc}} \,. \tag{9}$$

Combining these with our transparent mask regularization $\mathcal{L}_{\text{mask}}$ in Eq. (7), the total loss is given as:

$$\mathcal{L} = \mathcal{L}_{\text{img}} + \lambda_{\text{n}}\mathcal{L}_{\text{n}} + \lambda_{\text{mask}}\mathcal{L}_{\text{mask}} \,. \tag{10}$$

We empirically set $\lambda_{\text{n}} = 0.05$ and $\lambda_{\text{mask}} = 0.01$.

Table 4. Sensitivity analysis of the gating strength hyperparameter k, averaged over transparent regions of all scenes.

k	PSNR↑	SSIM↑	LPIPS↓
0	38.574	0.9861	0.0179
1	38.436	0.9860	0.0185
2	38.631	0.9862	0.0178
4	**38.696**	**0.9865**	**0.0175**
8	38.608	0.9863	0.0176
16	38.523	0.9864	0.0179
32	38.417	0.9863	0.0180

B. Additional Ablation Study

To evaluate the impact of our Specular-Aware Gradient Gating mechanism, we conduct a sensitivity analysis on the gating strength hyperparameter k defined in Eq. (5). The quantitative results, which are averaged over the transparent regions across all eight scenes, are summarized in Tab. 4. The performance peaks at $k = 4$, which offers the optimal trade-off between suppressing specular artifacts and preserving transmission details.

C. Additional Results

Tables 5 and 6 report per-scene quantitative metrics for entire images and transparent regions, respectively. In addition, Fig. 6 presents further qualitative results to demonstrate the decomposition capabilities of our method. In contrast, baseline methods often fail to simultaneously reconstruct surface reflections and the transmission behind the surface, recovering only one component at the expense of the other. We also recommend viewing the supplementary video to best appreciate the temporal consistency and visual quality of our results.

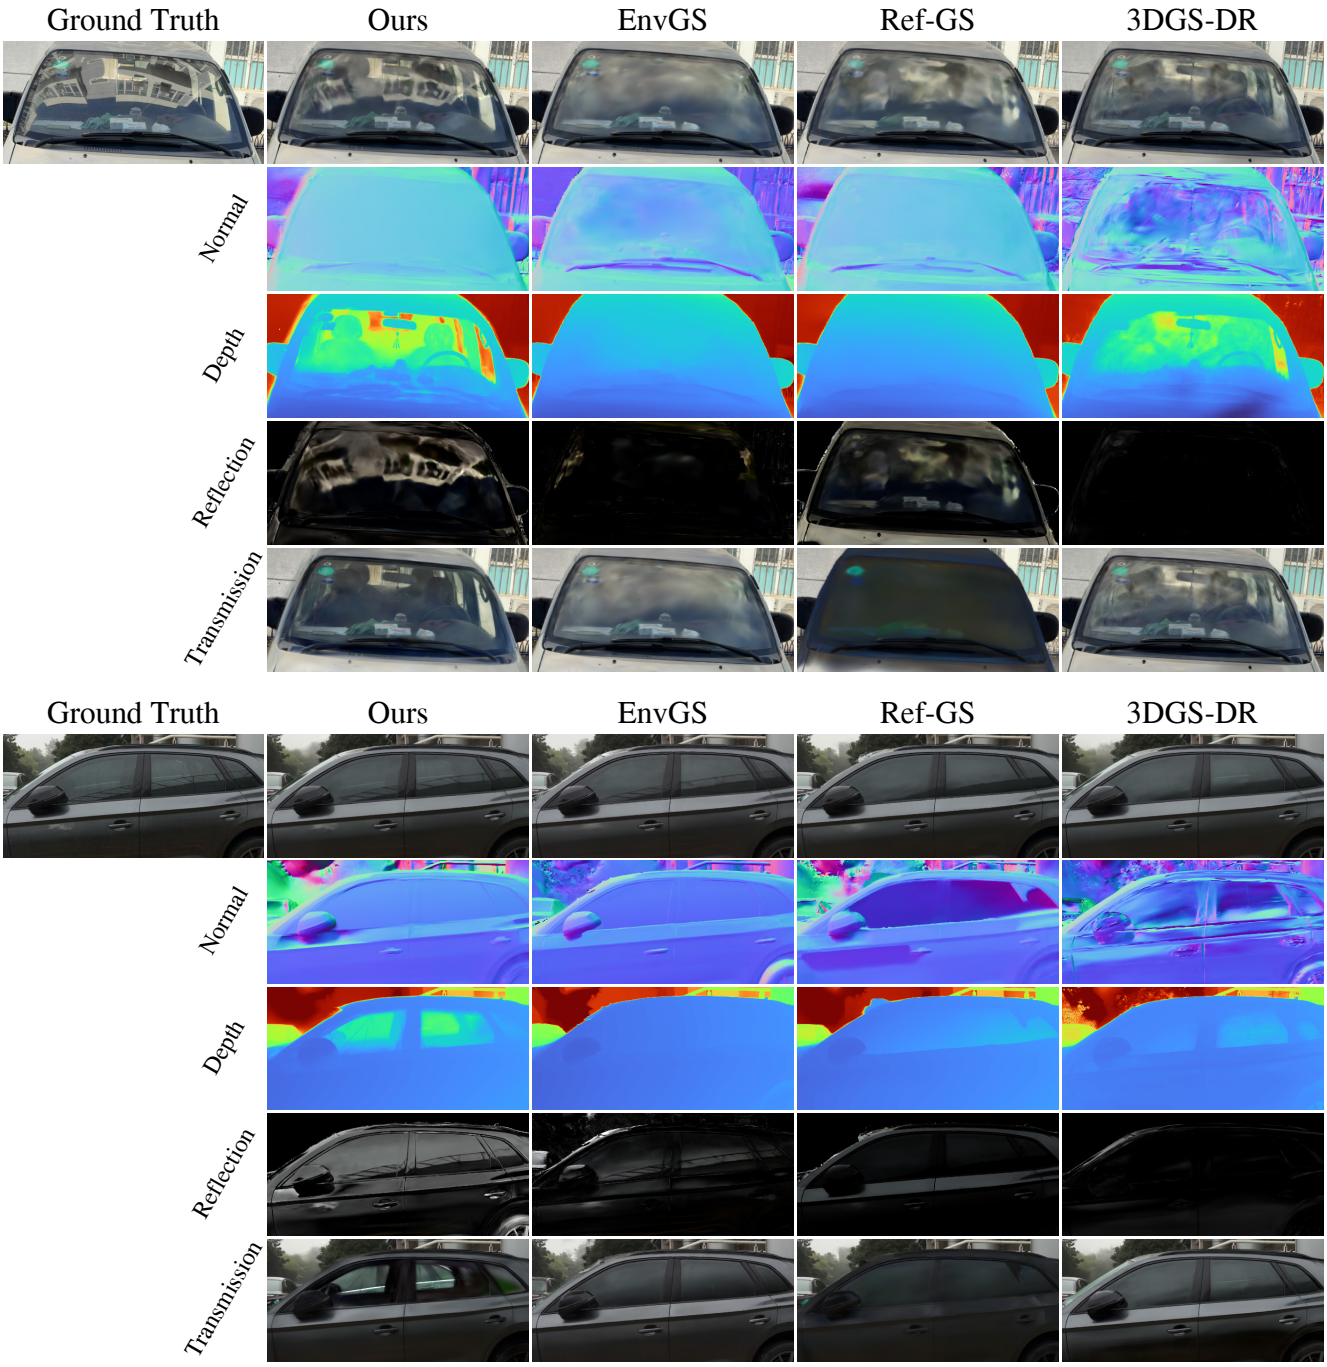

Figure 6. **Visual decomposition on real-world scenes.** We visualize the decomposed components of our method compared to baselines. For our method, *Normal* captures the surface geometry, while *Depth* corresponds to the volumetric accumulation. Since baseline methods do not explicitly model semi-transparent transmission, we visualize their diffuse component in the *Transmission* row. Our method achieves high-fidelity separation of the *Reflection* and *Transmission* layers.

Methods	PSNR ↑							
	Sedan	*Toycar*	*Compact*	*Hatchback*	*Audi*	*Truck*	*Van*	*Swab*
3DGS [18]	26.038	23.779	29.506	26.556	27.683	25.394	23.436	31.578
2DGS [14]	25.467	24.012	29.602	26.623	27.487	25.113	23.317	30.032
GShader [17]	25.709	23.888	28.542	25.880	26.913	23.738	22.189	20.077
3DGS-DR [43]	26.117	24.123	29.232	27.088	27.778	25.246	23.385	28.882
Ref-GS [50]	26.397	24.176	29.706	26.556	27.858	24.902	22.351	30.251
EnvGS [41]	26.790	24.637	29.703	27.457	28.686	25.571	23.043	30.651
Ours	27.071	24.794	30.598	27.619	29.011	25.848	23.744	33.817

Methods	SSIM ↑							
	Sedan	*Toycar*	*Compact*	*Hatchback*	*Audi*	*Truck*	*Van*	*Swab*
3DGS [18]	0.768	0.637	0.895	0.845	0.874	0.878	0.778	0.949
2DGS [14]	0.770	0.653	0.888	0.848	0.875	0.870	0.768	0.945
GShader [17]	0.760	0.649	0.885	0.840	0.860	0.844	0.747	0.855
3DGS-DR [43]	0.768	0.655	0.883	0.846	0.869	0.873	0.773	0.943
Ref-GS [50]	0.778	0.658	0.898	0.845	0.864	0.871	0.758	0.943
EnvGS [41]	0.777	0.667	0.878	0.848	0.881	0.876	0.753	0.942
Ours	0.790	0.684	0.893	0.851	0.887	0.881	0.785	0.958

Methods	LPIPS ↓							
	Sedan	*Toycar*	*Compact*	*Hatchback*	*Audi*	*Truck*	*Van*	*Swab*
3DGS [18]	0.205	0.237	0.147	0.196	0.150	0.148	0.233	0.193
2DGS [14]	0.234	0.238	0.175	0.198	0.164	0.174	0.266	0.219
GShader [17]	0.226	0.258	0.167	0.195	0.177	0.192	0.282	0.453
3DGS-DR [43]	0.209	0.240	0.163	0.204	0.160	0.166	0.255	0.203
Ref-GS [50]	0.205	0.236	0.149	0.196	0.183	0.156	0.263	0.217
EnvGS [41]	0.207	0.233	0.167	0.183	0.139	0.163	0.267	0.253
Ours	0.188	0.227	0.141	0.177	0.137	0.129	0.215	0.179

Table 5. **Per-scene metrics on entire images.** We report PSNR, SSIM, and LPIPS on scenes from Ref-Real [35], NeRF-Casting [36], EnvGS [41], T&T [21], and our self-captured scenes.

Methods	PSNR ↑							
	Sedan	*Toycar*	*Compact*	*Hatchback*	*Audi*	*Truck*	*Van*	*Swab*
3DGS [18]	35.514	40.462	37.212	37.543	38.742	36.567	30.652	34.481
2DGS [14]	35.315	40.379	37.276	36.555	37.932	36.541	30.720	33.126
GShader [17]	35.394	40.545	35.194	36.017	38.300	35.334	28.836	25.155
3DGS-DR [43]	36.489	40.137	37.301	37.241	39.518	36.655	30.547	33.480
Ref-GS [50]	37.037	40.338	37.212	37.543	38.540	35.894	29.711	33.581
EnvGS [41]	36.729	40.638	36.777	37.609	39.475	36.492	29.917	33.535
Ours	37.727	40.899	40.785	38.927	42.015	38.238	32.429	38.551

Methods	SSIM ↑							
	Sedan	*Toycar*	*Compact*	*Hatchback*	*Audi*	*Truck*	*Van*	*Swab*
3DGS [18]	0.984	0.995	0.988	0.991	0.990	0.991	0.949	0.979
2DGS [14]	0.985	0.995	0.988	0.991	0.989	0.991	0.951	0.978
GShader [17]	0.984	0.995	0.984	0.990	0.988	0.988	0.945	0.944
3DGS-DR [43]	0.986	0.995	0.988	0.991	0.989	0.990	0.949	0.978
Ref-GS [50]	0.987	0.994	0.988	0.991	0.988	0.989	0.947	0.976
EnvGS [41]	0.986	0.995	0.987	0.991	0.990	0.989	0.949	0.977
Ours	0.988	0.995	0.992	0.993	0.992	0.992	0.954	0.986

Methods	LPIPS ↓							
	Sedan	*Toycar*	*Compact*	*Hatchback*	*Audi*	*Truck*	*Van*	*Swab*
3DGS [18]	0.016	0.005	0.015	0.010	0.018	0.008	0.053	0.042
2DGS [14]	0.017	0.005	0.014	0.010	0.020	0.008	0.056	0.049
GShader [17]	0.017	0.006	0.019	0.011	0.020	0.010	0.061	0.107
3DGS-DR [43]	0.015	0.005	0.015	0.010	0.019	0.009	0.054	0.044
Ref-GS [50]	0.014	0.006	0.015	0.010	0.022	0.010	0.060	0.052
EnvGS [41]	0.015	0.006	0.016	0.009	0.018	0.010	0.058	0.055
Ours	0.012	0.005	0.010	0.008	0.015	0.007	0.049	0.034

Table 6. **Per-scene metrics on transparent regions.**