Spec-Gloss Surfels and Normal-Diffuse Priors for Relightable Glossy Objects

Supplementary Material

6. Indirect Illumination

We adopt the indirect lighting formulation of Ref-Gaussian [33]. The specular IBL is split into direct and indirect components, modulated by visibility $V \in \{0, 1\}$ along the reflected direction $\boldsymbol{\omega}_r = 2(\boldsymbol{\omega}_o \cdot \mathbf{n})\mathbf{n} - \boldsymbol{\omega}_o$. The visibility is computed via bounding volume hierarchy (BVH) for accelerated ray tracing on an extracted TSDF mesh, which is updated periodically (i.e. every 3k iterations) for efficiency. The visible part uses the standard prefiltered environment lookup, while the occluded part is modeled per Gaussian with low-order spherical harmonics $S_i(\cdot)$ evaluated at the reflection direction $\boldsymbol{\omega}_r$ and alpha-blended in screen space:

$$L_{\text{ind}}(\boldsymbol{\omega}_r) = \sum_{i=1}^{N} l_{\text{i}}(\boldsymbol{\omega}_r)\, \alpha_i \prod_{j=1}^{i-1}(1 - \alpha_j), \tag{9}$$

where $l_{\text{i}}(\boldsymbol{\omega}_r) = S_i(\boldsymbol{\omega}_r)$. This formulation captures interreflections while retaining real-time performance.

7. Material Parameterizations in the Literature

In the following equations, we present the computation of the outgoing color for the various material parameterizations. The Disney BRDF with the metallic-roughness parameterization estimates L_o as

$$L_o^{\text{MR}} = (1 - m)bL_d + (F_0(m, b)\beta_1 + \beta_2)L_s, \tag{10}$$

which entangles the base color b to the metallic m through F_0. GaussianShader [9] simplifies the rendering equation to solve the entanglement issue by learning F_0 and using it for energy conservation on the diffuse appearance as well, departing, however, from principled models in literature or industry. The simplified outgoing radiance equation is defined as:

$$L_o^{\text{Simple}} = \sigma(b - ln3) + (1 - k_s)L_d + (k_s\beta_1 + \beta_2)L_s, \tag{11}$$

where σ is the sigmoid function. More recently, Ref-Gaussian [33] tried a similar approach that solved the entanglement issue by replacing base color b in F_0 with a specular tint term k_s

$$L_o^{\text{MRS}} = b + (F_0(m, k_s)\beta_1 + \beta_2)L_s, \tag{12}$$

where the Fresnel coefficients β_1, β_2 are given from a lookup table indexed with roughness r and the nv product of the surface normal $\mathbf{n}$ and viewing ray direction $\mathbf{v}$. Finally, L_d is the irradiance from the environment, convolved with a diffuse kernel, and L_s is the prefiltered specular environment. Nevertheless, it faces the same portability issues as GaussianShader, and while both methods perform well on the NVS task, performance deteriorates in relighting.

8. Additional Results

8.1. Scene Editing

Tab. 5 and Fig. 12 compare our relighting results with Gaussian Splatting baselines, showing that our method produces more plausible relighting than Ref-Gaussian [33], 3DGS-DR [12], and GaussianShader [9]. Fig. 13 compares the three material parameterizations (SG, MR, MRS) and validates our choice, as SG consistently outperforms the MR and MRS variants. Fig. 14 further illustrates the material editing capabilities of our method.

8.2. Recovered Environment Map Results

Figures 15-17 present additional results on environment map recovery for the Shiny Synthetic [30], Glossy Synthetic [21], and Shiny Real [30] scenes.

8.3. Reconstruction Results

In Figures 18 - 20, we present qualitative results for all examined scenes. The corresponding quantitative results are presented in Tab. 6. We observe that our method is competitive and often outperforms Gaussian splatting and implicit baselines. The gap in performance in real scenes is mostly due to the limitation of 2DGS in representing dense, thin structures and sparsely seen background features.

8.4. Material Decomposition

In Fig. 21 and Fig. 22, we present the material decomposition of our method, the diffuse and specular components, the surface normals, as well as visibility, direct and indirect light terms for all examined synthetic scenes.

8.5. Diffuse Color and Surface Normal Priors

Fig. 23 and Fig. 24 present sample predictions of StableDelight [25] and StableNormal [34] for all examined scenes from the Shiny Synthetic [30], Glossy Synthetic [21] and Shiny Real [30] scenes. StableDelight manages to remove specular reflections in most cases, but fails in cases of mirror-like surfaces (e.g. toaster and glossy teapot), confusing the reflections as part of the diffuse appearance of the scene. StableNormal, on the other hand, fails to recognize the surface of the liquid in the coffee scene and oversmooths complex geometries (e.g. cat, luyu) and occasionally confuses interreflections as separate geometries (e.g. tbell).

Table 5. Per-scene relighting and NVS comparison on the Glossy Synthetic [21] scenes. We highlight the first, second, and third best results.

	Relighting							Novel View Synthesis						
	bell	**cat**	**luyu**	**potion**	**tbell**	**teapot**	***avg***	**bell**	**cat**	**luyu**	**potion**	**tbell**	**teapot**	***avg***
PSNR ↑														
GShader [9]	20.34	15.92	15.83	14.04	18.47	19.25	17.31	28.07	31.81	27.18	30.09	24.48	23.58	27.55
3DGS-DR [12]	20.59	20.36	21.34	20.01	19.08	19.94	20.22	31.65	33.86	28.71	32.79	28.94	25.36	30.14
RefGauss [33]	19.87	21.58	20.45	20.20	21.35	21.45	20.82	32.86	33.01	30.04	33.07	29.84	26.68	30.92
Ours (MR)	17.03	21.73	21.29	27.57	15.81	18.62	20.34	28.61	32.20	27.95	33.31	27.59	24.05	28.95
Ours (MRS)	21.66	26.14	21.80	27.35	22.94	20.82	23.45	30.71	33.11	29.46	33.38	30.01	25.76	30.40
Ours (SG)	25.40	26.82	22.73	27.13	23.77	22.72	24.76	33.40	33.54	29.73	33.57	30.12	26.96	31.22
SSIM ↑														
GShader [9]	0.900	0.868	0.828	0.783	0.875	0.883	0.856	0.919	0.961	0.914	0.938	0.898	0.901	0.921
3DGS-DR [12]	0.886	0.894	0.875	0.862	0.883	0.906	0.884	0.962	0.976	0.936	0.957	0.952	0.936	0.953
RefGauss [33]	0.888	0.919	0.877	0.870	0.912	0.894	0.893	0.969	0.973	0.952	0.963	0.962	0.947	0.961
Ours (MR)	0.866	0.851	0.902	0.856	0.937	0.787	0.865	0.944	0.970	0.934	0.967	0.948	0.929	0.949
Ours (MRS)	0.914	0.899	0.937	0.894	0.938	0.912	0.904	0.962	0.976	0.950	0.968	0.966	0.946	0.961
Ours (SG)	0.942	0.946	0.905	0.939	0.930	0.928	0.932	0.975	0.978	0.953	0.968	0.967	0.956	0.966
LPIPS ↓														
GShader [9]	0.114	0.104	0.098	0.156	0.138	0.115	0.121	0.098	0.056	0.064	0.088	0.122	0.091	0.086
3DGS-DR [12]	0.103	0.097	0.083	0.115	0.110	0.080	0.098	0.064	0.040	0.053	0.075	0.067	0.067	0.058
RefGauss [33]	0.104	0.070	0.079	0.110	0.087	0.087	0.090	0.040	0.040	0.043	0.064	0.058	0.058	0.051
Ours (MR)	0.142	0.091	0.099	0.079	0.133	0.099	0.107	0.072	0.047	0.055	0.054	0.072	0.069	0.062
Ours (MRS)	0.091	0.068	0.070	0.078	0.080	0.079	0.078	0.045	0.038	0.043	0.054	0.049	0.055	0.047
Ours (SG)	0.062	0.064	0.065	0.077	0.074	0.067	0.068	0.032	0.035	0.042	0.053	0.050	0.046	0.043

Figure 12. Qualitative relighting comparison under three environment maps (corridor, golf, neon) on the Glossy Synthetic dataset. Our method typically yields more faithful relighting than competing Gaussian Splatting baselines.

Figure 13. Qualitative relighting comparison under three environment maps (corridor, golf, neon) on the Glossy Synthetic dataset. We compare our method against three alternative material parameterizations, and the results support our choice of the specular–glossiness (SG) parameterization.

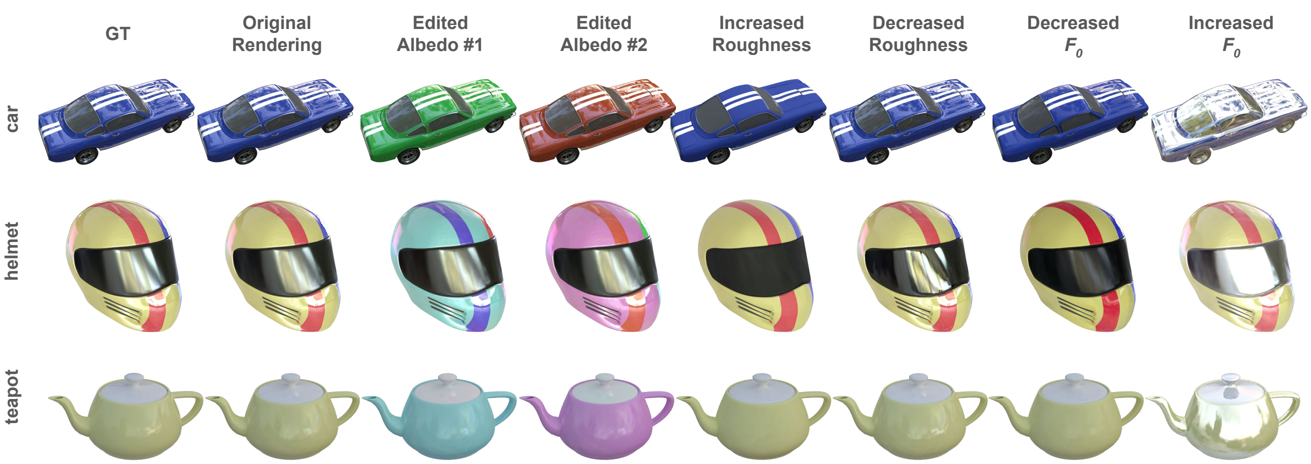

Figure 14. Demonstration of the material-editing capabilities of our method. After training, we modify the reconstructed materials and re-render three scenes from the Shiny Synthetic dataset [30] with increased or decreased albedo, roughness, or F_0.

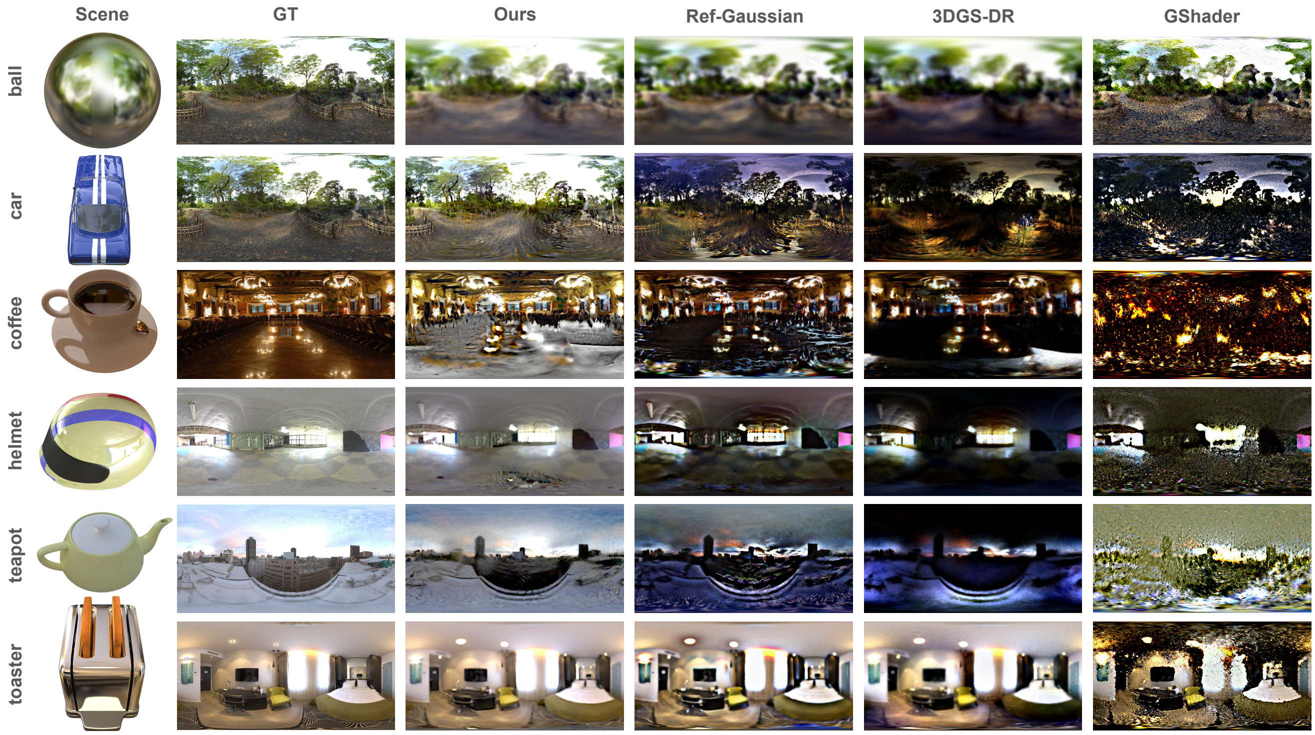

Figure 15. Comparison of recovered environment maps on the scenes of the Shiny Synthetic dataset [30].

Figure 16. Comparison of recovered environment maps from the scenes of the Glossy Synthetic dataset [21].

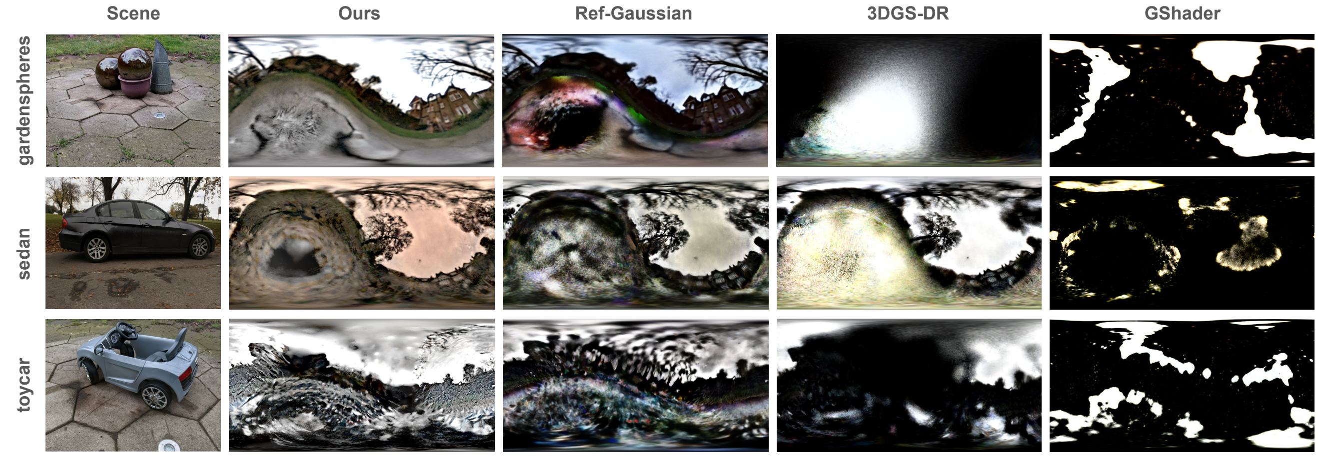

Figure 17. Comparison of recovered environment maps on the scenes of the Shiny Real dataset [30].

Table 6. Quantitative comparison of novel view synthesis averaged over scenes within each dataset. Methods labeled with ✓in the Rel. column support relighting while those labeled with ✗do not. For RefGauss [33], although higher scores were reported on Shiny Real [30], we were unable to reproduce them. We therefore report results obtained using the authors' released code. We highlight the first, second, and third best results.

Method	Rel.	Shiny Synthetic [30]							Glossy Synthetic [21]							Shiny Real [30]			
		ball	car	coffee	helmet	teapot	toaster	*avg*	bell	cat	luyu	potion	tbell	teapot	*avg*	garden	sedan	toycar	*avg*
PSNR ↑																			
Ref-NeRF	✗	33.16	30.44	33.99	29.94	45.12	26.12	33.13	30.02	29.76	25.42	30.11	26.91	22.77	27.50	22.01	25.21	23.65	23.62
3DGS	✗	27.65	27.26	32.30	28.22	45.71	20.99	30.36	25.11	31.36	26.97	30.16	23.88	21.51	26.50	21.75	26.03	23.78	23.85
GShader	✓	30.99	27.96	32.39	28.32	45.86	26.28	31.97	28.07	31.81	27.18	30.09	24.48	23.58	27.54	21.74	24.89	23.76	23.46
ENVIDR	✓	41.02	27.81	30.57	32.71	42.62	26.03	33.46	30.88	31.04	28.03	32.11	28.64	26.77	29.58	21.47	24.61	22.92	23.00
3DGS-DR	✓	33.66	30.39	34.65	31.69	47.12	27.02	34.09	31.65	33.86	28.71	32.79	28.94	25.36	30.22	21.82	26.32	23.83	23.99
Ref-GS	✗	36.10	30.94	34.38	33.40	46.69	27.28	34.80	31.70	33.15	29.46	32.64	30.08	26.47	30.58	22.48	26.63	24.20	24.44
RefGauss	✓	37.01	31.04	34.63	32.32	47.16	28.05	35.04	32.86	33.01	30.04	33.07	29.84	26.68	30.92	22.79	25.13	24.01	23.98
Ours	✓	38.14	31.92	34.72	33.72	46.77	27.69	35.50	33.40	33.54	29.73	33.57	30.12	26.96	31.22	22.69	26.17	24.19	24.35
SSIM ↑																			
Ref-NeRF	✗	0.971	0.950	0.972	0.954	0.995	0.921	0.961	0.941	0.944	0.901	0.933	0.947	0.897	0.927	0.584	0.720	0.633	0.646
3DGS	✗	0.937	0.931	0.972	0.951	0.996	0.894	0.947	0.892	0.959	0.916	0.938	0.908	0.881	0.916	0.571	0.771	0.637	0.660
GShader	✓	0.966	0.932	0.971	0.951	0.996	0.929	0.958	0.919	0.961	0.914	0.938	0.898	0.901	0.922	0.576	0.728	0.637	0.647
ENVIDR	✓	0.997	0.943	0.962	0.987	0.995	0.990	0.979	0.954	0.965	0.931	0.960	0.947	0.957	0.952	0.561	0.707	0.549	0.606
3DGS-DR	✓	0.979	0.962	0.976	0.971	0.997	0.943	0.971	0.962	0.976	0.936	0.957	0.952	0.936	0.953	0.581	0.773	0.639	0.664
Ref-GS	✗	0.981	0.961	0.973	0.975	0.997	0.950	0.973	0.965	0.973	0.946	0.957	0.956	0.944	0.957	0.507	0.783	0.682	0.657
RefGauss	✓	0.981	0.964	0.976	0.959	0.997	0.942	0.970	0.969	0.973	0.952	0.963	0.962	0.947	0.964	0.616	0.731	0.642	0.663
Ours	✓	0.989	0.974	0.977	0.978	0.997	0.949	0.978	0.975	0.978	0.953	0.968	0.967	0.956	0.966	0.615	0.761	0.661	0.679
LPIPS ↓																			
Ref-NeRF	✗	0.166	0.050	0.082	0.086	0.012	0.083	0.080	0.102	0.104	0.098	0.084	0.114	0.098	0.100	0.251	0.234	0.231	0.239
3DGS	✗	0.162	0.047	0.079	0.081	0.008	0.125	0.084	0.104	0.062	0.064	0.093	0.125	0.102	0.092	0.248	0.206	0.237	0.230
GShader	✓	0.121	0.044	0.078	0.074	0.007	0.079	0.067	0.098	0.056	0.064	0.088	0.122	0.091	0.087	0.274	0.259	0.239	0.257
ENVIDR	✓	0.020	0.046	0.083	0.036	0.009	0.081	0.046	0.054	0.049	0.059	0.072	0.069	0.041	0.057	0.263	0.387	0.345	0.332
3DGS-DR	✓	0.098	0.033	0.076	0.049	0.005	0.081	0.057	0.064	0.040	0.053	0.075	0.067	0.067	0.061	0.247	0.208	0.231	0.229
Ref-GS	✗	0.098	0.034	0.082	0.045	0.006	0.070	0.056	0.049	0.041	0.046	0.076	0.073	0.064	0.058	0.242	0.196	0.236	0.225
RefGauss	✓	0.098	0.033	0.076	0.050	0.006	0.074	0.056	0.040	0.040	0.043	0.064	0.058	0.058	0.047	0.278	0.277	0.279	0.278
Ours	✓	0.073	0.027	0.078	0.038	0.006	0.073	0.049	0.032	0.035	0.042	0.053	0.050	0.046	0.043	0.291	0.222	0.264	0.259

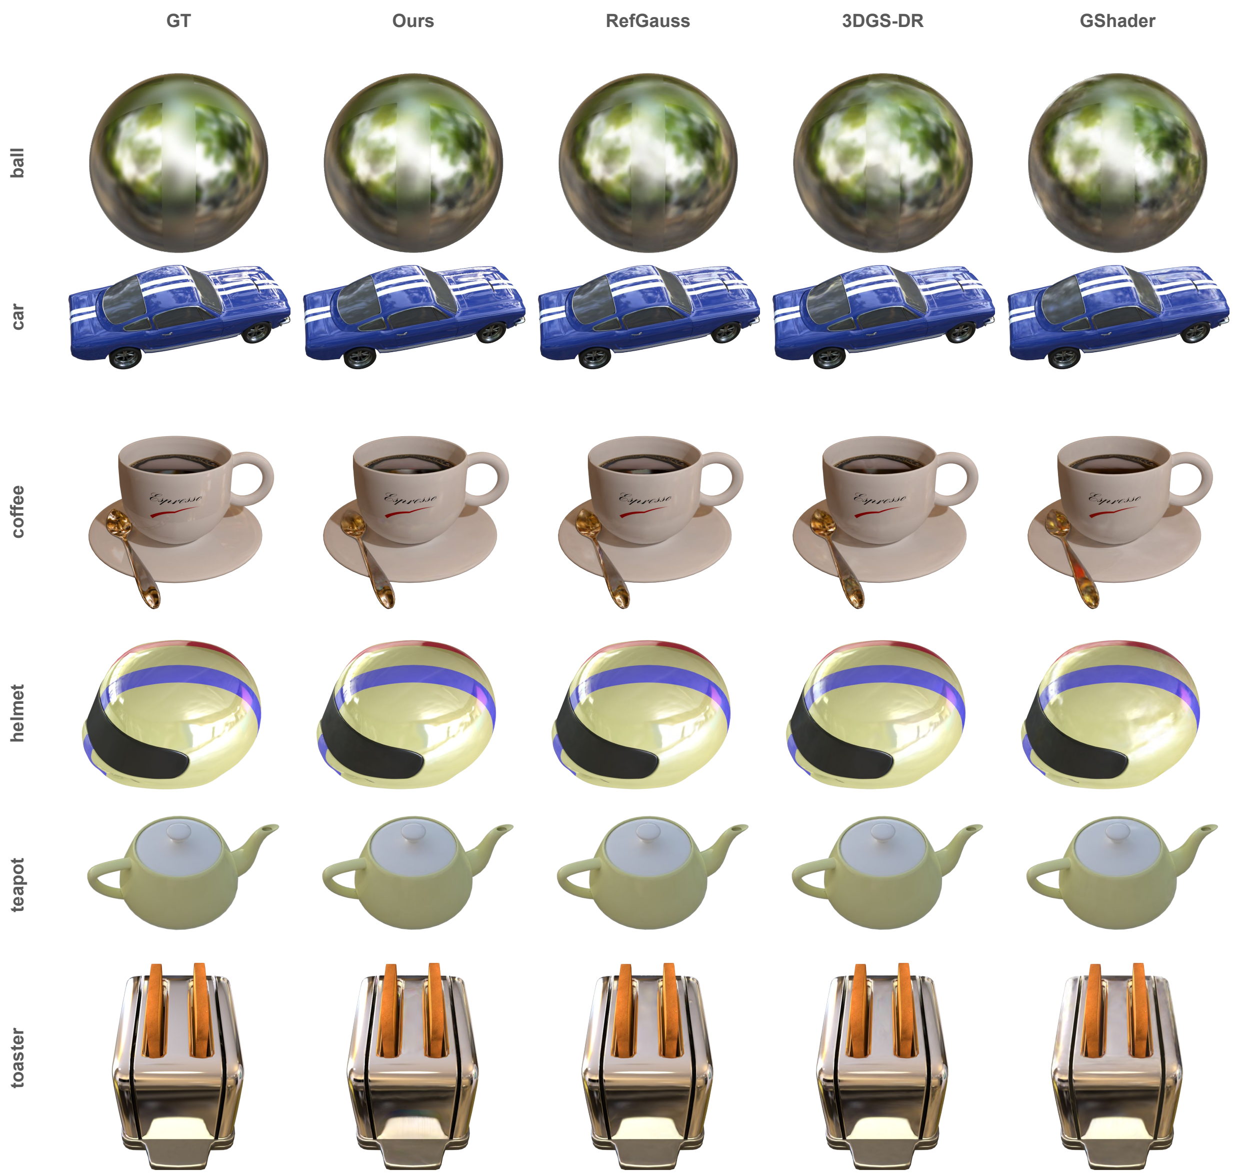

Figure 18. Comparison of rendering quality on the Shiny Synthetic dataset [30].

Figure 19. Comparison of rendering quality on the Glossy Synthetic dataset [21].

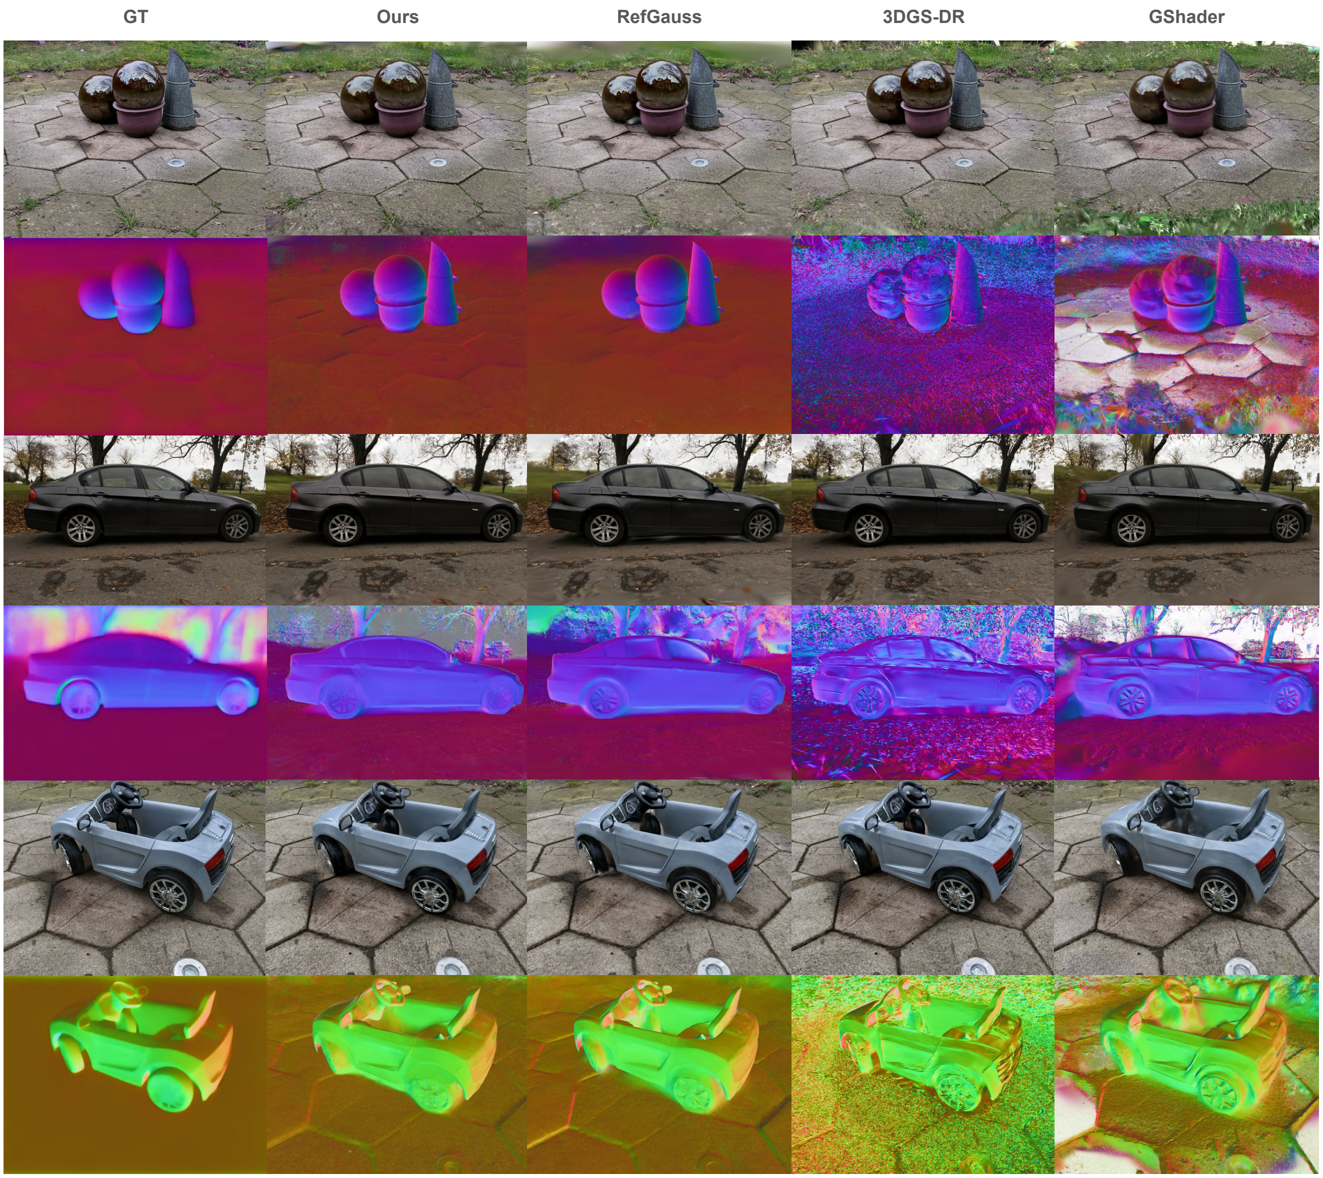

Figure 20. Comparison of rendering quality on the Shiny Real dataset [30]. In the leftmost column, we show the GT RGB test images and the pseudo-GT surface normals, generated by StableNormal [34], and used to supervise the predicted normals.

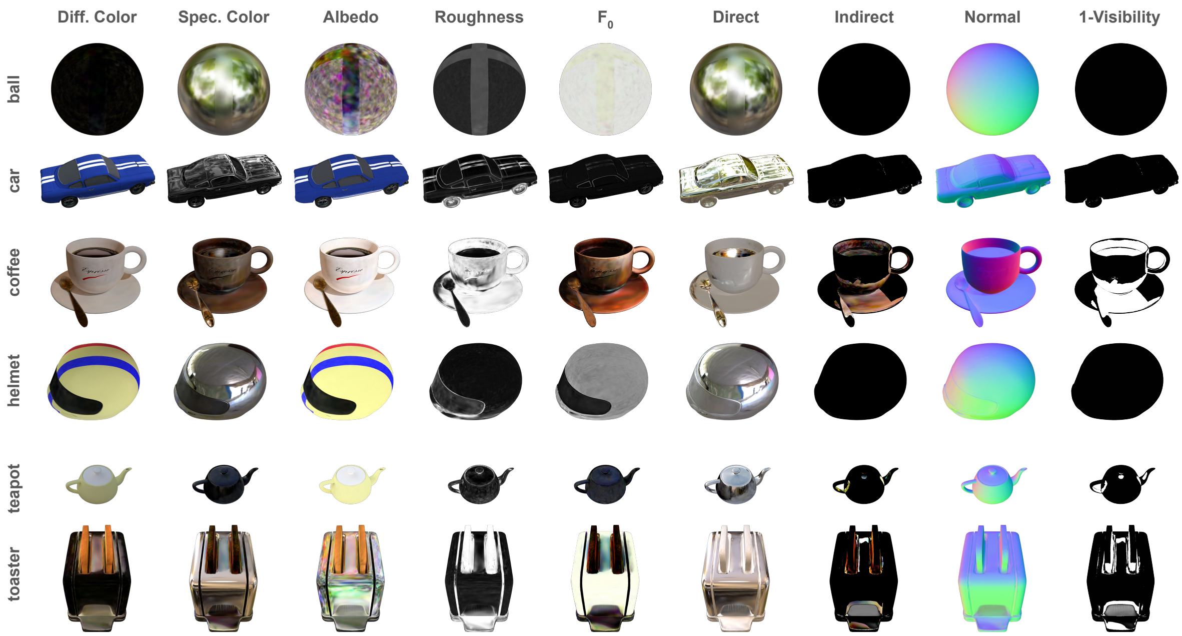

Figure 21. Material decomposition and outputs of our method on the Shiny Synthetic dataset [30]

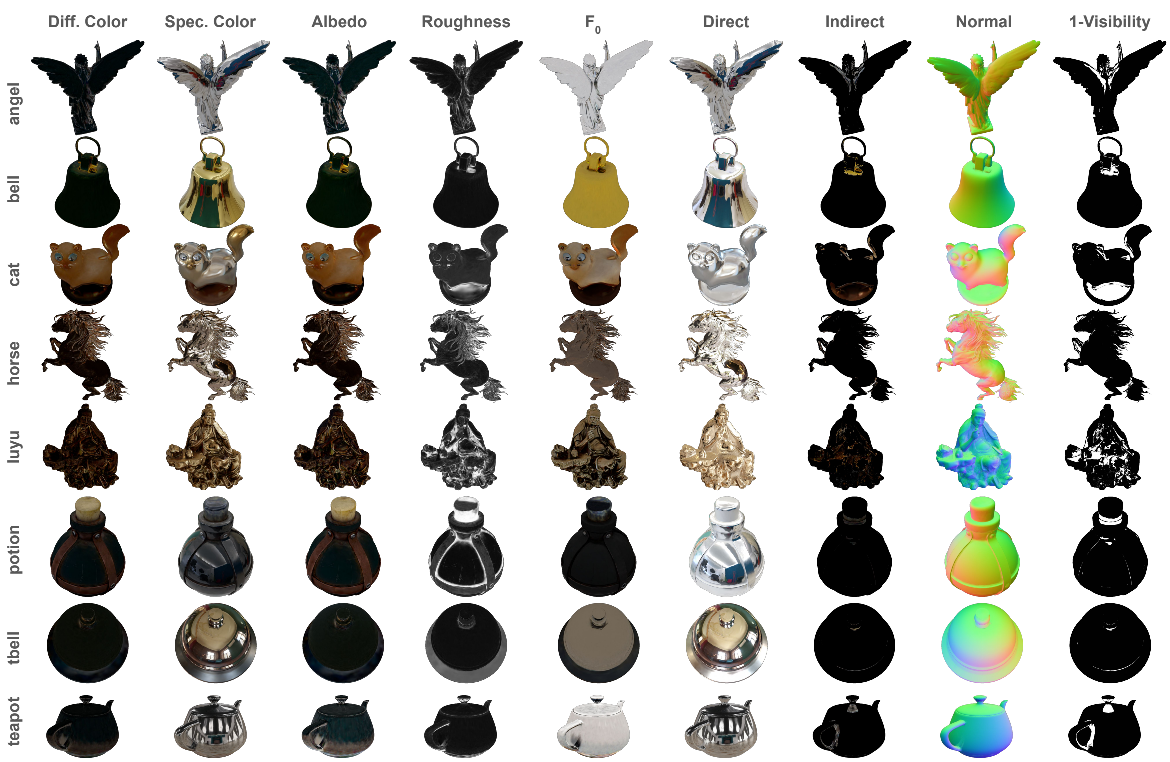

Figure 22. Material decomposition and outputs of our method on the Glossy Synthetic dataset [21]

Figure 23. Example predictions of StableDelight [25] on scenes from the Shiny Synthetic [30], Glossy Synthetic [21], and Shiny Real [30] datasets.

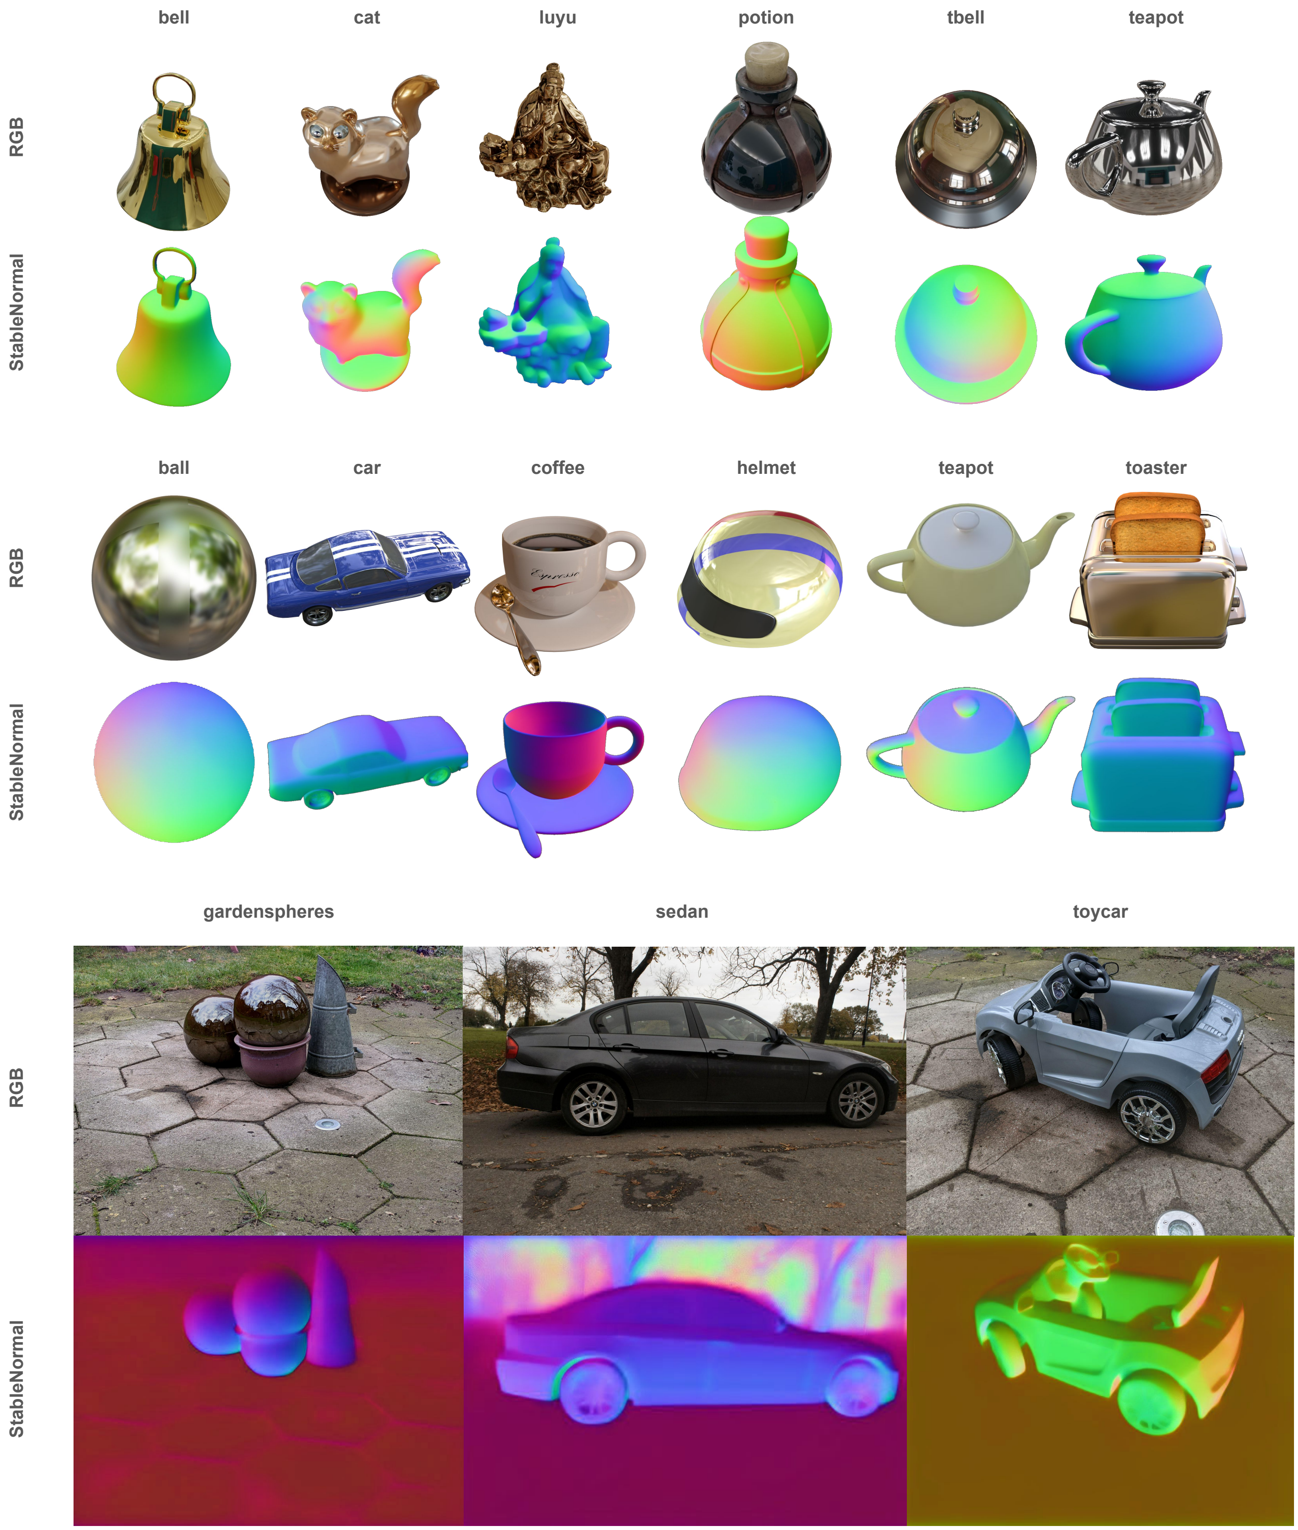

Figure 24. Example predictions of StableNormal [34] on scenes from the Shiny Synthetic [30], Glossy Synthetic [21], and Shiny Real [30] datasets.