A Implementation Details

Progressive Training. Our experiments follow the reflectance *Composition D* and the progressive *Schedule I* described in Sec. 5.3. For fair comparison, we train each scene for 100K iterations on the NRHints and GS^3 datasets, and for 60K iterations on the SSS-GS dataset to match the original settings. The overall progressive training process is illustrated in Fig. 6 and Fig. 7.

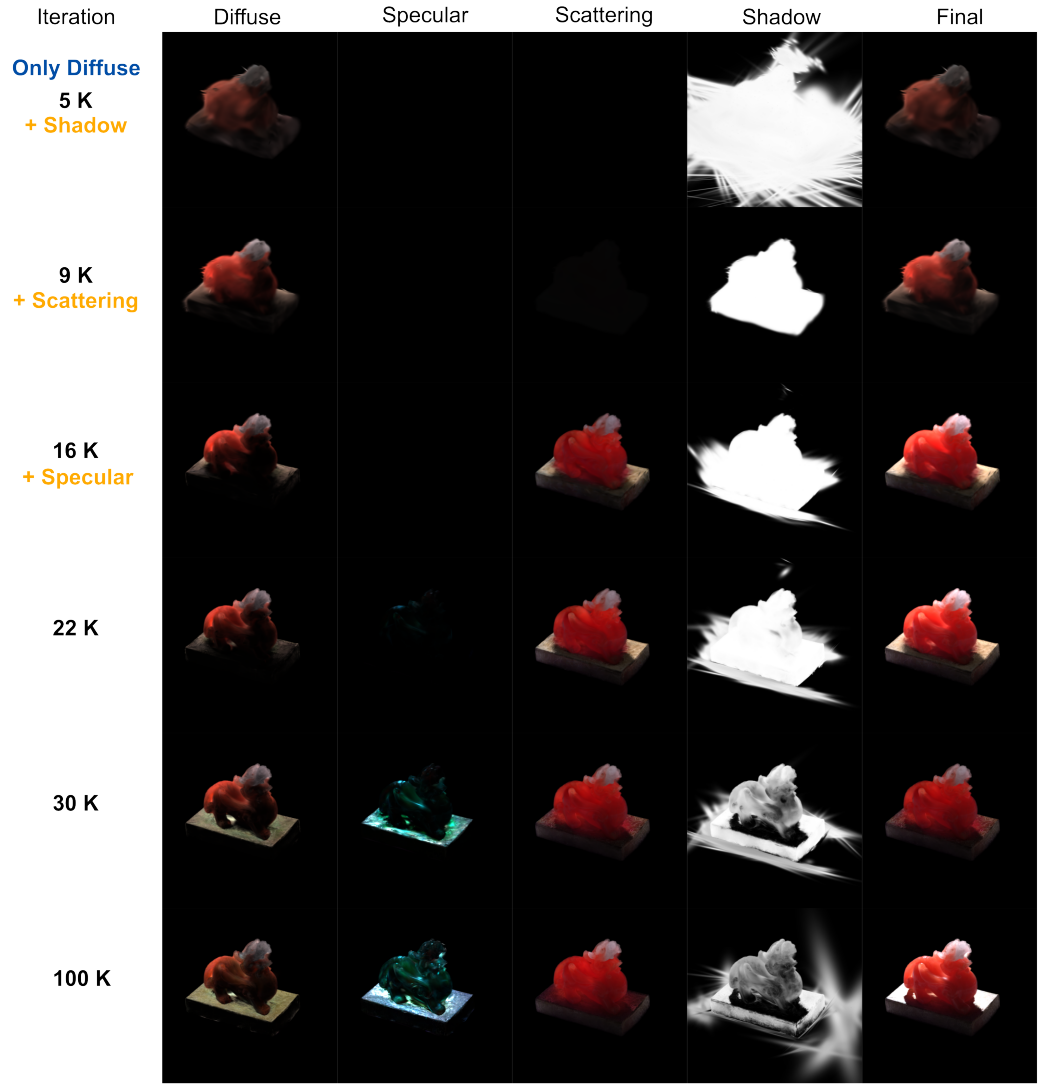


Figure 6: Illustration of the progressive training schedule on the *pixiu* scene from the NRHints real dataset (Zeng et al., 2023). The model is trained for a total of 100K iterations.

Specifically, during the initial 5K iterations, only the diffuse term contributes to shading. This warm-up phase stabilizes coarse geometry and appearance, since all Gaussians are randomly initialized (10K points uniformly distributed on the unit sphere) from NeRF JSON inputs. These early diffuse-only iterations are therefore critical for forming a reliable Gaussian structure.

From 5K iterations onward, we introduce the shadow term, which provides a first approximation of light visibility and substantially improves lighting initialization. At 9K iterations, the subsurface-scattering term is activated to model low-frequency, multiple-scattering effects in translucent materials. Between 9K and 16K iterations, shadow gradients are temporarily held fixed so that the scattering term can converge without interference from competing gradients. This schedule is motivated by the longer convergence needs of scattering, which would otherwise remain underfitted.

After 16K iterations, we introduce the specular term to capture high-frequency details and view-dependent highlights. Optimizing specular too early tends to dominate smoother reflectance components, especially subsurface scattering, since the optimizer naturally prioritizes sharper signals (see Sched. H in Fig. 5 and Fig. 12). To address this, we temporarily freeze gradients of the scattering term between 13K and 20K iterations, allowing specular learning to progress without suppressing low-frequency transport. Additionally, from 16K to 20K iterations we suspend updates to the ASG lobe parameters (scale and rotation), so the specular module first focuses on Fresnel intensity before refining anisotropic lobe orientations.

In summary, the progressive schedule gradually disentangles low- and high-frequency reflectance phenomena, balancing the convergence speed of each component and yielding a stable and physically consistent optimization.

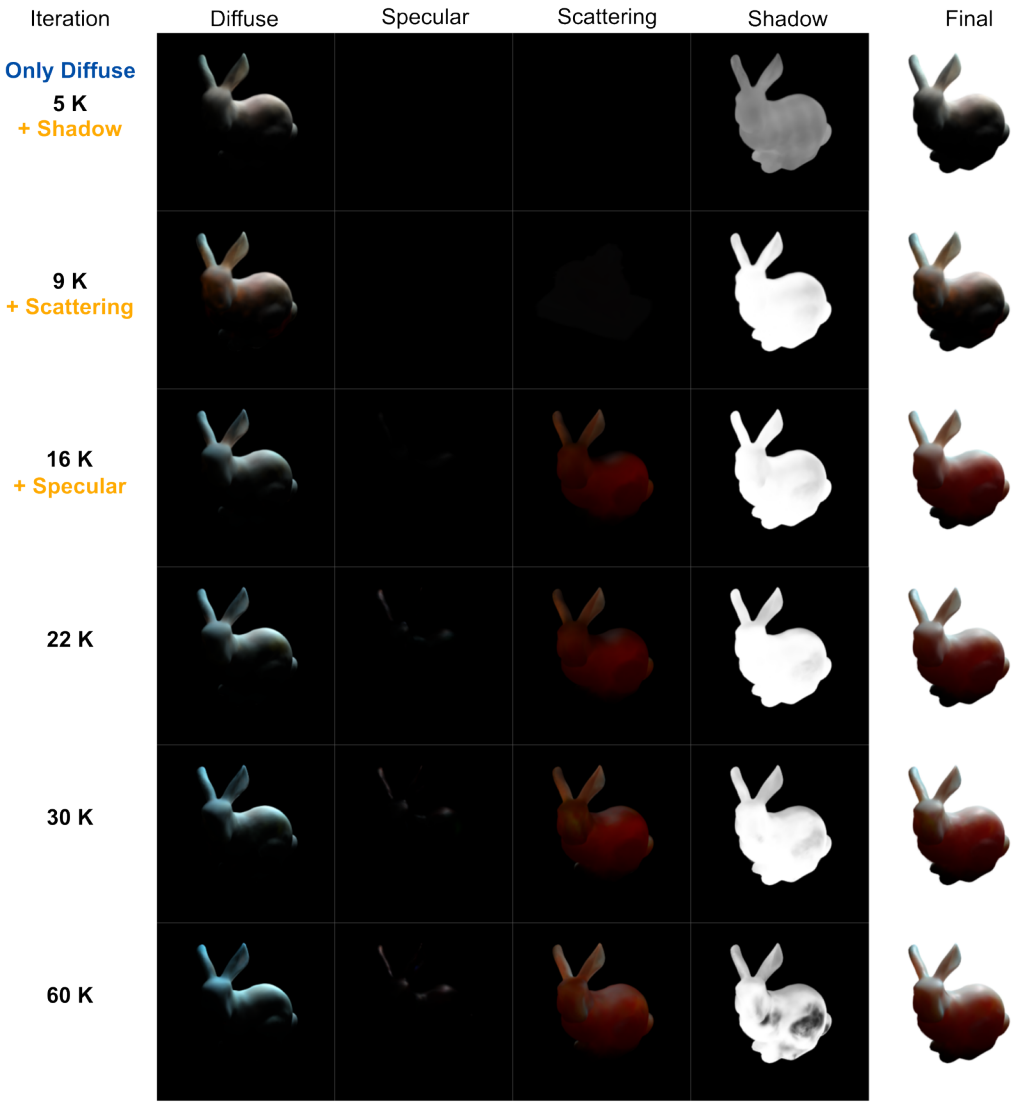


Figure 7: Illustration of the progressive training schedule on the *bunny* scene from the SSS-GS synthetic dataset (Dihlmann et al., 2025). The model is trained for a total of 60K iterations.

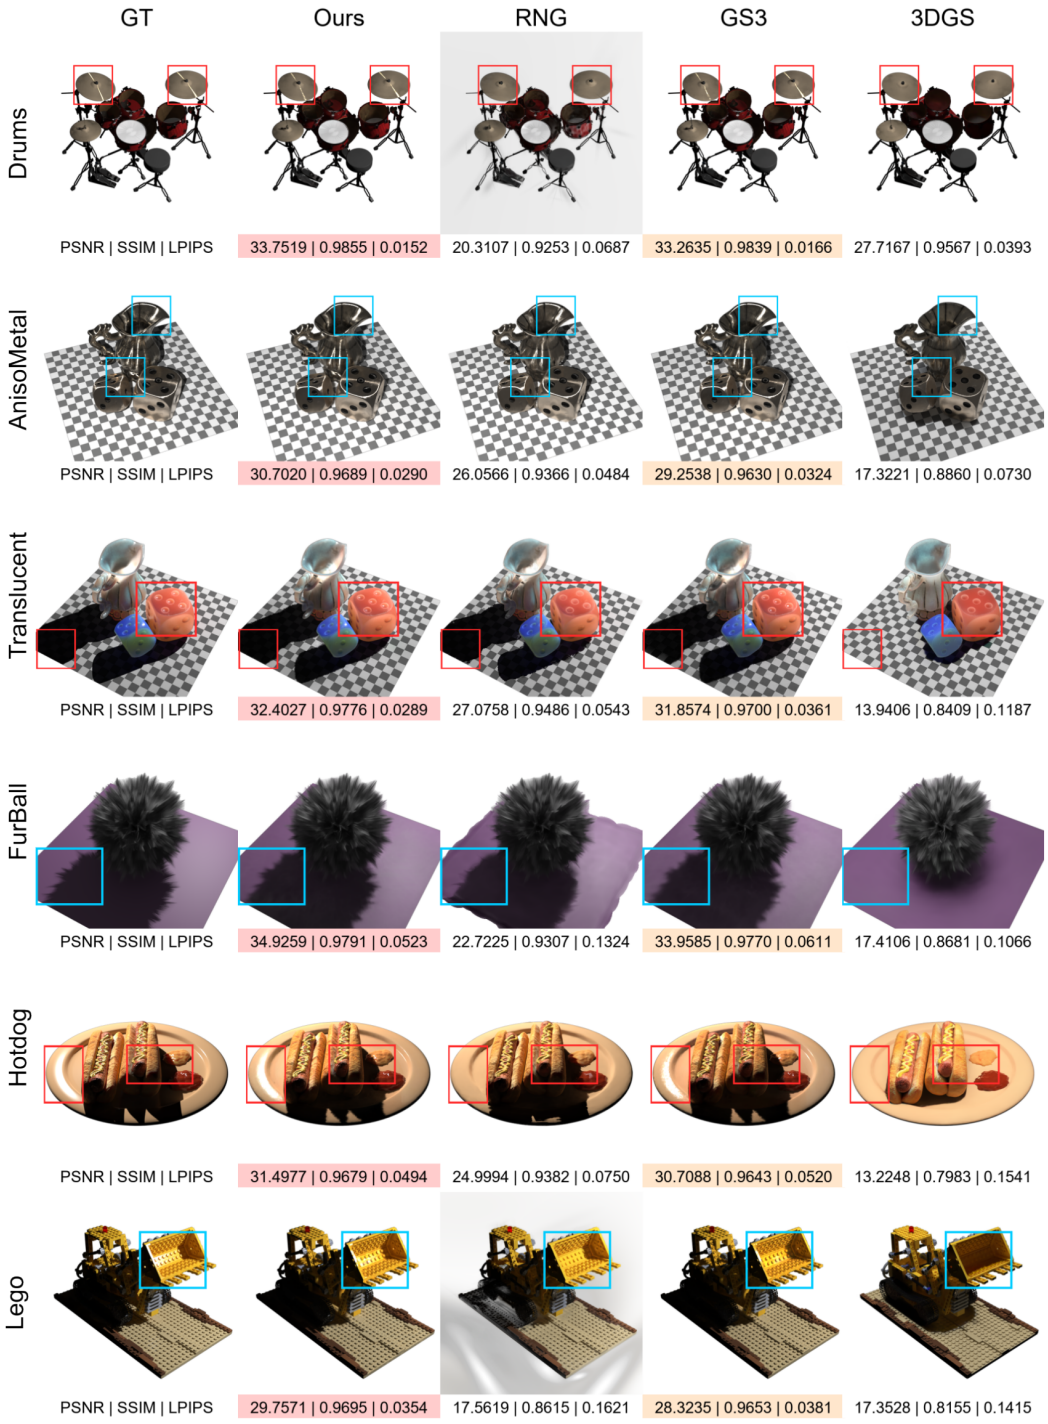


Figure 8: Qualitative comparison of relighting results on novel test-time lighting from synthetic datasets provided by GS^3 (Bi et al., 2024). Note that the rendered view of RNG's results on the *Drums* and *Lego* scenes are obtained using their official code trained using a set of white-background images.

B TECHNICAL DETAILS

Diffuse Term. We adopt a Lambertian BRDF to model the diffuse component, assuming view-independent cosine-weighted reflection:

$$f_d = \max(0, \mathbf{n} \cdot \omega_i), \tag{9}$$

where $\mathbf{n}$ is the surface normal and ω_i is the incident light direction. This simple formulation stabilizes the reconstruction of low-frequency reflectance and provides a physically interpretable baseline for shading.

Specular Term. The specular component is modeled as a Fresnel factor multiplied by an anisotropic spherical Gaussian (ASG) reflectance function:

$$f_s = F(\omega_o, \mathbf{h}) \cdot D_{\text{ASG}}(\mathbf{h}), \tag{10}$$

where $\mathbf{h}$ is the half-vector between light and view directions, and $F(\omega_o, \mathbf{h})$ denotes the Fresnel reflectance term, approximated with Schlick's formulation (Schlick, 1994).

The ASG reflectance function is expressed as a weighted sum of N global ASG bases:

$$D_{\text{ASG}}(\mathbf{h}) = \sum_{j=1}^{N} G_j(\mathbf{h}) \cdot \alpha_j, \tag{11}$$

where each basis G_j takes the geometric form (Xu et al., 2013):

$$G_j(\mathbf{h}) = \exp\left(-\lambda_j(\mathbf{h} \cdot \mathbf{x}_j)^2 - \mu_j(\mathbf{h} \cdot \mathbf{y}_j)^2\right), \tag{12}$$

with $(\mathbf{x}_j, \mathbf{y}_j)$ defining the local anisotropy axes. By leveraging a compact set of global ASG bases with learnable weights α_j, we achieve expressive, view-dependent reflections without the need for per-Gaussian specular storage, maintaining efficiency while preserving rendering fidelity.

Table 3: Quantitative comparison with SSS-GS (Dihlmann et al., 2025) and KiloOSF (Yu et al., 2022) on the synthetic datasets provided by SSS-GS, trained for 60K iterations. The best and second-best results (based on PSNR) are highlighted in red and orange, respectively.

Method \ Dataset	Bunny		Candle		Dragon		Soap		Statue		Average	
	Train	Test	Train	Test	Train	Test	Train	Test	Train	Test	Train	Test
PSNR ↑												
KiloOSF	-	-	-	-	-	-	-	-	-	-	-	25.91 ± 1.88
SSS-GS	-	-	-	-	-	-	-	-	-	-	-	35.01 ± 1.01
Ours (w/o Opt)	40.7672	37.2270	40.0682	38.3662	39.3646	36.6325	45.1439	40.4385	38.1997	34.5404	40.7087	37.4409
Ours (w/ Opt)	40.8704	37.2960	43.2426	41.1038	40.8462	37.3363	45.2659	40.6914	39.1271	35.3434	41.8705	38.3542
SSIM ↑												
KiloOSF	-	-	-	-	-	-	-	-	-	-	-	0.93 ± 0.02
SSS-GS	-	-	-	-	-	-	-	-	-	-	-	0.972 ± 0.01
Ours (w/o Opt)	0.9922	0.9859	0.9910	0.9875	0.9874	0.9789	0.9950	0.9908	0.9879	0.9782	0.9907	0.9843
Ours (w/ Opt)	0.9920	0.9861	0.9947	0.9921	0.9900	0.9812	0.9951	0.9912	0.9901	0.9812	0.9924	0.9863
LPIPS ↓												
KiloOSF	-	-	-	-	-	-	-	-	-	-	-	0.83 ± 0.09
SSS-GS	-	-	-	-	-	-	-	-	-	-	-	0.040 ± 0.01
Ours (w/o Opt)	0.0113	0.0179	0.0134	0.0172	0.0157	0.0240	0.0060	0.0104	0.0153	0.0235	0.0123	0.0186
Ours (w/ Opt)	0.0111	0.0173	0.0072	0.0099	0.0121	0.0209	0.0056	0.0099	0.0135	0.0212	0.0099	0.0158

C ADDITIONAL RESULTS ON THE SYNTHETIC DATASET

GS^3 Synthetic Dataset. We compare our method against four representative baselines: the original 3D Gaussian Splatting (Kerbl et al., 2023), GI-GS (Chen et al., 2024), GS^3 (Bi et al., 2024), and RNG (Fan et al., 2025). These methods are evaluated on reconstruction of training views from the synthetic datasets released by GS^3, which contain diverse reflectance properties and serve as a comprehensive benchmark for relightable rendering. Quantitative results are reported in Tab. 1b of the main paper, and qualitative comparisons are shown in Fig. 8. We additionally compare against several representative relighting baselines, as presented in Appendix F.

To ensure a fair comparison, all methods were trained for 100K iterations on the same training sets, followed by rendering on both the training and test splits. Reconstruction quality is evaluated on the

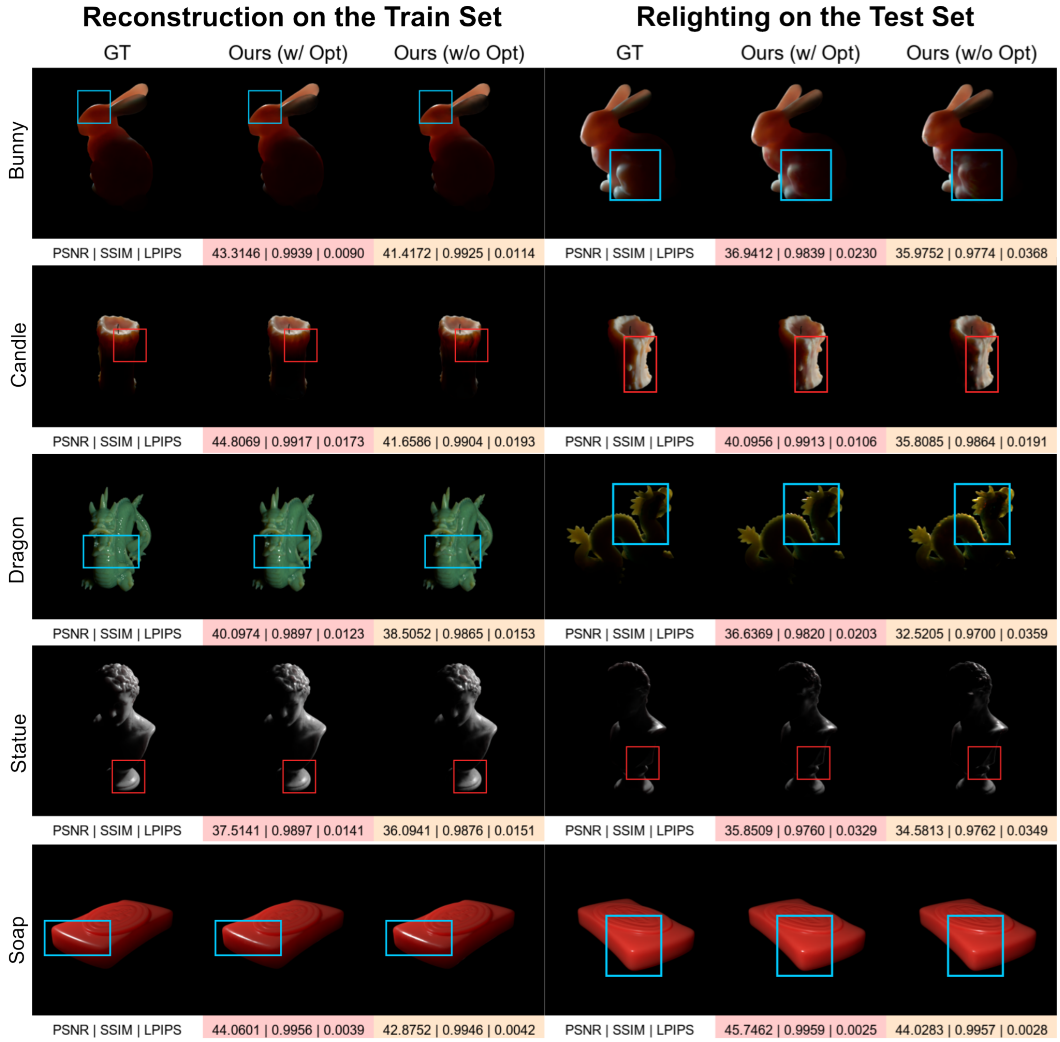


Figure 9: Qualitative comparison on synthetic datasets from SSS-GS (Dihlmann et al., 2025). The left column shows reconstruction results on training views, while the right column presents relighting results on novel test-time lighting. The best and second-best results (based on PSNR) are highlighted in red and orange, respectively.

training set, while relighting performance is assessed on the test set, which contains unseen lighting conditions. Our method achieves superior reconstruction of specular highlights across most scenes, such as *Drums* and *Hotdog*, and demonstrates more accurate shadow reconstruction in complex cases like *FurBall*.

Additionally, we observed that GS^3 occasionally produces shadow artifacts at certain viewpoints—especially in scenes like *FurBall*—and we intentionally avoided including such anomalous views in the comparison to maintain fairness.

Since the GS^3 synthetic datasets contain limited subsurface scattering effects, we further validate the effectiveness of our model on the recently released synthetic dataset from SSS-GS, which features more prominent subsurface scattering phenomena.

SSS-GS Synthetic Dataset. We further evaluate our method on the recently released synthetic dataset from SSS-GS (Dihlmann et al., 2025). While both scale-down and full-resolution versions of the synthetic dataset exist, only the scale-down version (500 images per split at 256×256 resolution)

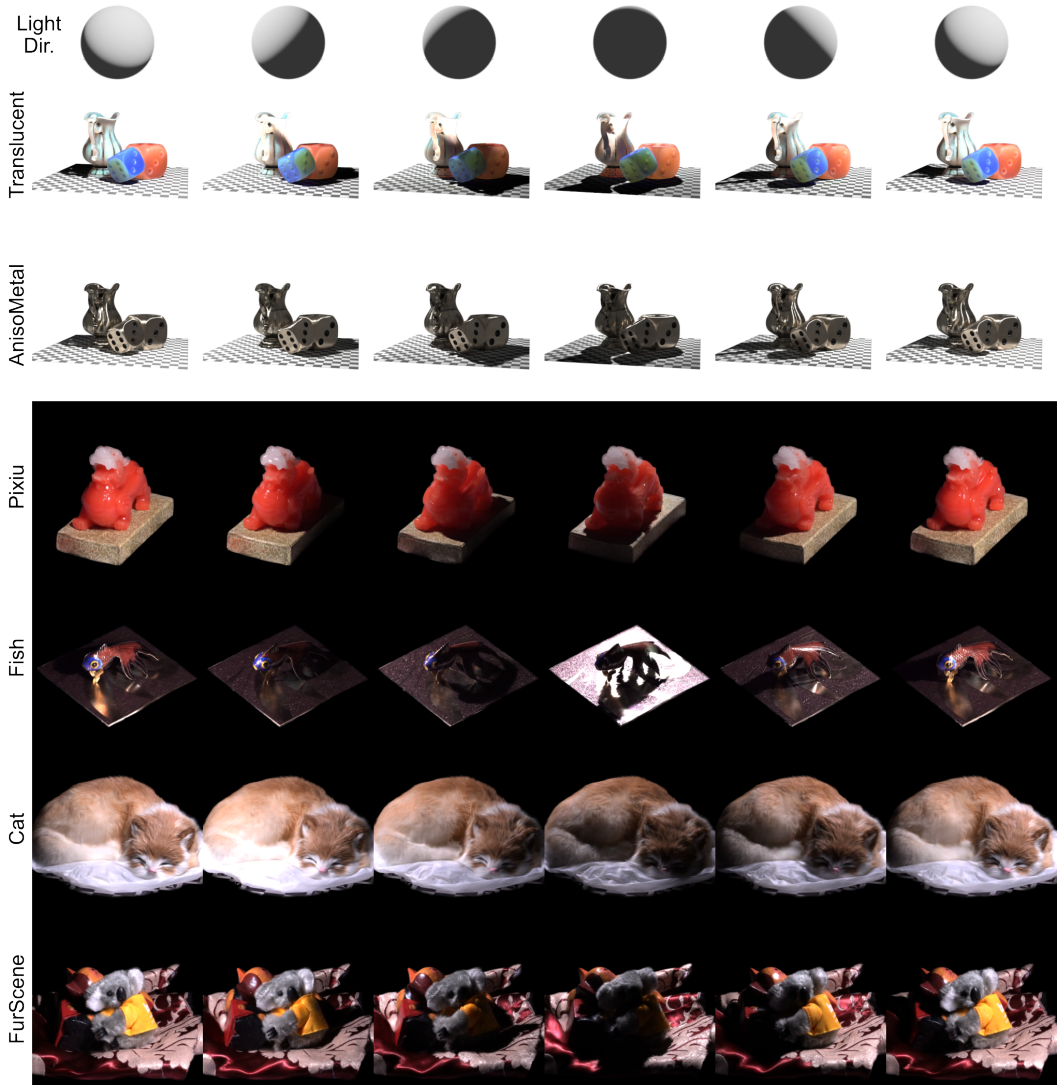


Figure 10: Qualitative relighting results on both real and synthetic datasets. Each scene is rendered from the same camera viewpoint under six novel lighting conditions.

is publicly available. In contrast, the real datasets are currently released only in full resolution (13,193 images per split at 800×649). Due to the limited time, we conducted our experiments only on the publicly available synthetic subset. This includes five scenes: *bunny*, *candle*, *dragon*, *soap*, and *statue*.

Following the experimental setup in the SSS-GS paper (Dihlmann et al., 2025), our method was trained for 60K iterations and rendered on a black background to ensure a fair comparison. We compare our method against SSS-GS and KiloOSF (Yu et al., 2022) using the quantitative results reported in the SSS-GS paper. Quantitative results are summarized in Tab. 1c, which reports overall metrics including average reconstruction quality and training cost. Per-scene quantitative results are further detailed in Tab. 3, allowing a finer-grained comparison across individual scenes. Qualitative comparisons are provided in Fig. 9.

We include two variants of our method in this comparison: Ours (w/o Opt) refers to our approach without camera and lighting optimization, while Ours (w/ Opt) includes joint optimization of both camera poses and light directions. This setting was briefly introduced in Sec. 4.5 of the main paper, where we progressively refine camera poses and lighting positions during training. The inclusion of camera and lighting optimization (w/ Opt) leads to noticeable improvements in both camera and

light estimates, resulting in an average increase of approximately 1dB in PSNR on both training and test sets. Despite the input images being downscaled to 256×256, subtle lighting differences remain discernible, particularly in the relighting results on test views of scenes such as *bunny* and *dragon* (right columns).

D Analysis of Relighting

To complement the OLAT test set, which provides ground-truth images for quantitative evaluation, we construct a validation setup for relighting with synthetic light–camera trajectories stored in JSON format. In this setting, each scene is illuminated by a single point light source, and the incident light direction is derived from the relative position between the light source and the center of each Gaussian, ensuring shading consistency at the Gaussian level. Unlike the test set, the validation setup does not include ground-truth images and is instead used to assess generalization under novel light and view configurations. Specifically, the light source is placed on a circular path around the object, sweeping 360° in azimuth with a step of 2.4°. Next, the camera is rotated by 180° to the back side of the object with a step of 2°, after which the light completes another full 360° sweep around the back view. Finally, the camera is rotated 180° back to the original frontal view.

This configuration provides dense sampling of both lighting and viewing conditions, enabling a comprehensive assessment of relighting fidelity. We perform relighting under this setup for both real and synthetic datasets (see Fig. 10). The results demonstrate that our method maintains consistency across varying illumination while faithfully preserving view-dependent effects.

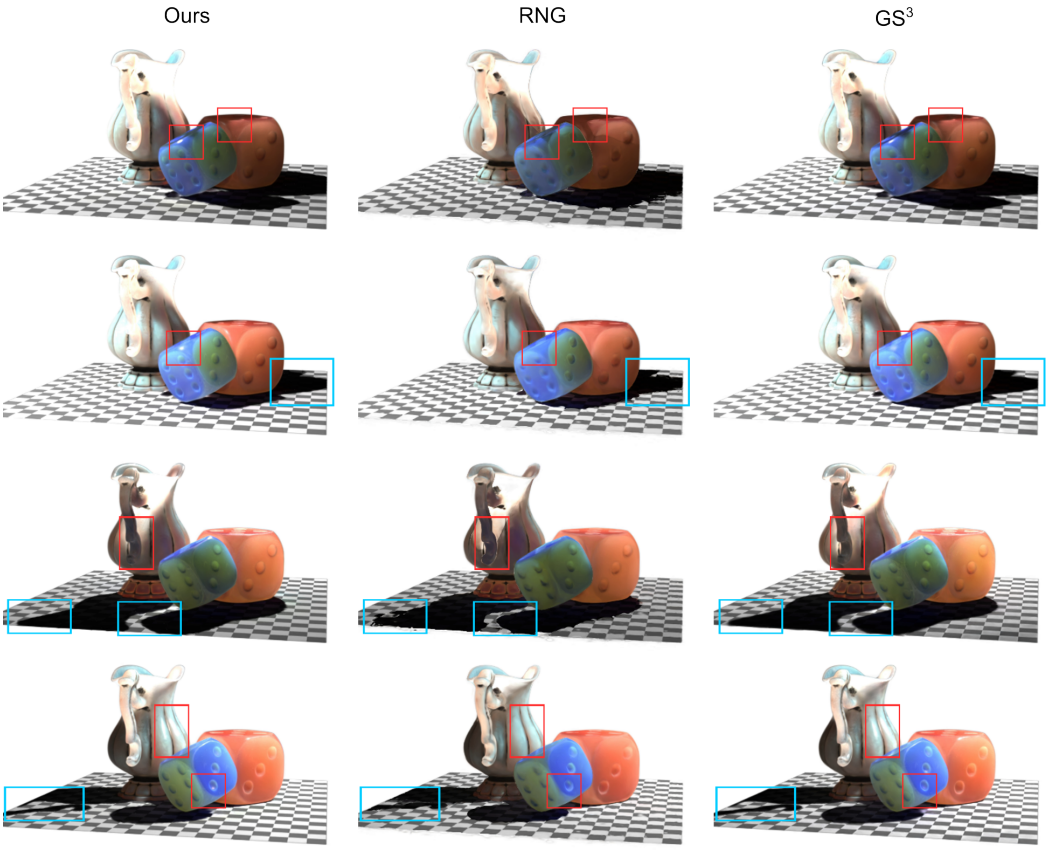


Figure 11: Qualitative relighting comparison on the *Translucent* scene. Each row shows renderings from the same camera viewpoint under four novel lighting conditions. Our method is compared against RNG (Fan et al., 2025) and GS^3 (Bi et al., 2024), demonstrating superior fidelity in reproducing light–material interactions.

Furthermore, we compare relighting results on the *Translucent* scene with two OLAT-based Gaussian relighting methods, RNG (Fan et al., 2025) and GS^3 (Bi et al., 2024), as shown in Fig. 11. The results demonstrate the superior relighting quality of our method, particularly in handling complex light–material interactions. For additional qualitative relighting results, please refer to the supplementary video.

In summary, our relighting analysis not only verifies the robustness of our approach under diverse lighting and viewing conditions, but also lays a fundamental basis for future research on controllable relighting, with broad applications in appearance editing, material-aware reconstruction, and immersive content creation.

Table 4: Ablation study results on both real and synthetic datasets. The best/second-best results are colored in red / orange.

Dataset		NRHints		GS^3				SSS-GS		Average
Scenes		Pixiu	Fish	Translucent	FurBall	Lego	Hotdog	Bunny	Dragon	
PSNR ↑										
A: Diff	Train	20.2869	23.7003	16.6152	18.5678	20.6157	18.4717	21.3548	27.7851	20.9247
	Test	20.1878	24.6146	15.3146	17.4010	17.6430	16.4279	21.3013	27.7071	20.0747
B: D + S	Train	20.7321	24.8949	16.6980	18.3871	20.6891	18.5180	21.4677	27.9589	21.1682
	Test	20.5692	25.3627	15.3200	17.1781	17.6290	16.4429	21.2236	27.6713	20.1746
C: D + S + SSS	Train	25.1545	24.8573	26.4267	18.3845	25.6734	18.4628	22.7074	28.8883	23.8194
	Test	24.8100	25.3275	25.6081	17.1655	22.8370	16.3730	22.3572	28.5486	22.8784
D: Full (Ours)	Train	33.6065	32.0748	32.6058	35.4793	31.1434	32.4901	40.7672	39.3646	34.6915
	Test	31.1213	31.1646	32.3919	35.1639	30.4664	32.1330	37.2270	36.6325	33.2876
E: Full - S	Train	32.3487	28.0442	30.9386	33.7976	30.0839	30.0792	34.8697	36.3361	32.0623
	Test	30.5952	28.2837	30.6761	34.0481	29.2294	30.4379	33.9633	35.8654	31.6374
F: Full - SSS	Train	31.5332	31.9130	29.8005	32.6455	30.0656	31.5471	35.7335	37.5666	32.6006
	Test	30.4095	31.1763	30.6253	33.8019	29.0048	32.0991	33.6715	35.9968	32.0982
H: Joint	Train	32.5452	31.4431	29.3508	32.5538	29.8266	30.5236	34.0744	37.0531	32.1713
	Test	31.0880	30.4943	30.6756	33.6832	29.1154	31.7083	31.7589	35.2621	31.7232
I: Prog. Phys (Ours)	Train	33.6065	32.0748	32.6058	35.4793	31.1434	32.4901	40.7672	39.3646	34.6915
	Test	31.1213	31.1646	32.3919	35.1639	30.4664	32.1330	37.2270	36.6325	33.2876
J: Prog. NonPhys	Train	32.5606	31.3321	29.6999	32.8201	30.7004	30.9712	35.9882	37.6112	32.7105
	Test	31.0973	30.2406	30.5955	33.9654	29.4718	31.2575	34.0090	36.0222	32.0824
K: Prog. Merge	Train	33.3438	32.0247	31.7351	32.6704	30.2933	31.4903	39.1104	37.5251	33.5241
	Test	31.0486	31.2827	31.8883	33.9799	29.2607	32.0399	36.6183	35.8016	32.7400

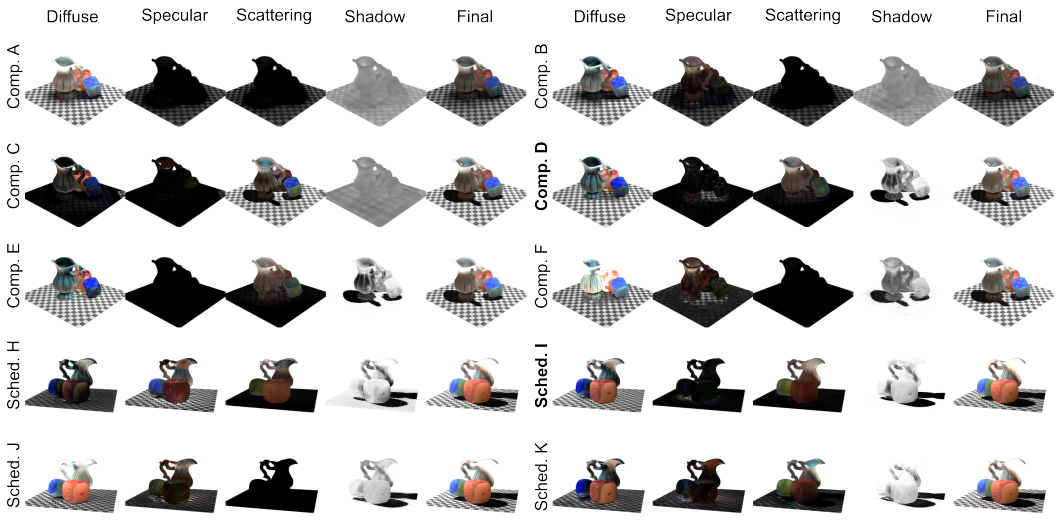


Figure 12: Visualization of reconstructed components under different reflectance decompositions and training schedules on the *Translucent* scene from the GS^3 synthetic dataset. Top: six reflectance compositions (Comp. A-F); Bottom: four training schedules (Sched. H-K).

E Additional Ablation Study

Reflectance Components and Training Schedule. In addition to the Pixiu example presented in the main paper (Fig. 5 and Tab. 2), we provide expanded ablation results across more real and synthetic scenes to illustrate the generality of our observations. For the reflectance components, we report per-scene comparisons for diffuse, specular, subsurface scattering (SSS), and shadow terms, highlighting consistent trends in directional visibility, highlight formation, and translucent appearance. We also include additional per-scene evaluations, following our progressive optimization strategy, for the training schedule. The supplementary Tab. 4 and visualizations Fig.12 confirm that the behaviors observed in the main paper hold robustly across diverse materials and lighting conditions.

Table 5: Quantitative comparison when applying the shadow term to the SSS component.

Dataset / Scenes		NRHints			GS^3				SSS-GS		Average
		Pixiu	Fish	FurScene	Translucent	FurBall	Lego	Hotdog	Bunny	Dragon	
PSNR ↑											
Shadow-on-SSS	Train	32.1850	31.7531	31.6750	30.8066	35.4219	31.1002	32.1094	35.9068	37.3346	33.1436
	Test	30.9756	31.1206	30.6575	30.0740	34.9420	30.3487	31.5272	33.6508	35.8476	32.1271
Ours	Train	33.6065	32.0748	31.7846	32.6058	35.4793	31.1434	32.4901	40.7672	39.3646	34.3685
	Test	31.1213	31.1646	30.7349	32.3919	35.1639	30.4664	32.1330	37.2270	36.6325	33.0039
SSIM ↑											
Shadow-on-SSS	Train	0.9484	0.9337	0.9570	0.9761	0.9768	0.9701	0.9755	0.9841	0.9815	0.9670
	Test	0.9432	0.9252	0.9510	0.9751	0.9732	0.9562	0.9721	0.9762	0.9760	0.9609
Ours	Train	0.9524	0.9363	0.9576	0.9835	0.9776	0.9706	0.9776	0.9922	0.9874	0.9705
	Test	0.9452	0.9260	0.9518	0.9823	0.9733	0.9570	0.9743	0.9859	0.9789	0.9638
LPIPS ↓											
Shadow-on-SSS	Train	0.0797	0.0811	0.0695	0.0278	0.0502	0.0336	0.0305	0.0285	0.0236	0.0472
	Test	0.0825	0.0880	0.0726	0.0277	0.0450	0.0424	0.0327	0.0360	0.0285	0.0506
Ours	Train	0.0751	0.0779	0.0690	0.0201	0.0482	0.0331	0.0293	0.0113	0.0157	0.0422
	Test	0.0791	0.0855	0.0724	0.0200	0.0442	0.0419	0.0326	0.0179	0.0240	0.0464

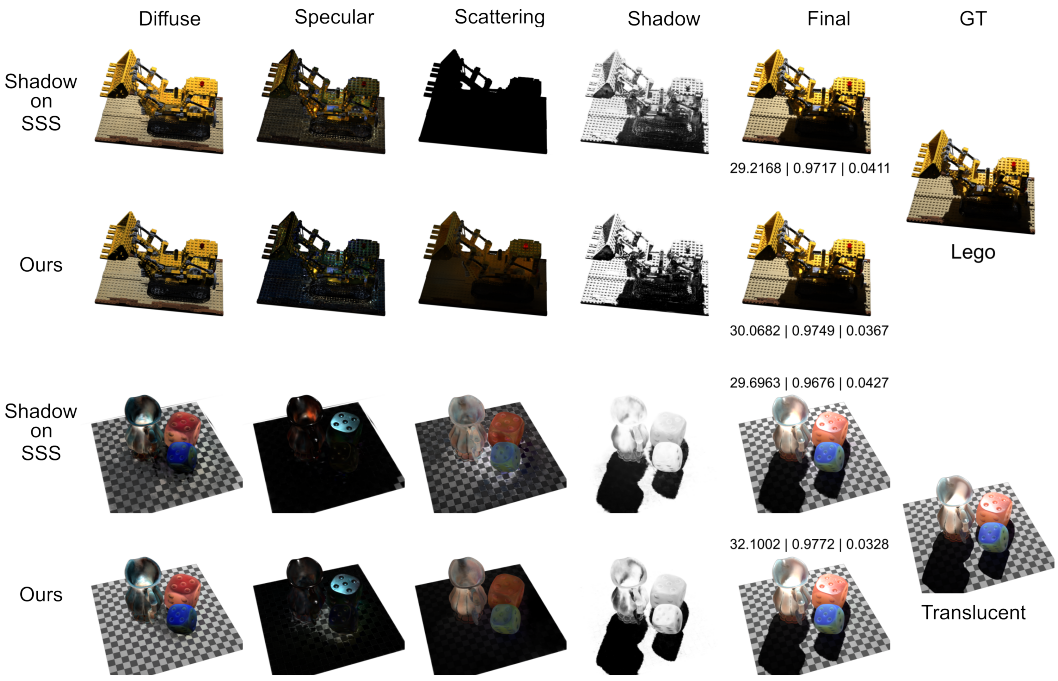


Figure 13: Qualitative comparison for the Shadow–SSS interaction ablation study. Applying the shadow term to the SSS component leads to over-darkened scattering, loss of back-lit translucency, and reduced soft shading. Our formulation preserves translucent appearance while maintaining consistent shadow behavior.

Shadow–SSS Interaction. To examine how the shadow term $S(\mathbf{x})$ interacts with subsurface scattering (SSS), we evaluate an alternative shading variant in which the shadow term is also applied to the SSS component. For clarity, we restate the shading model used in the main paper (Eq. 3):

$$\text{Ours: } (c_d f_d + c_s f_s) \cdot S(\mathbf{x}) + c_{sss} f_{sss} \tag{13}$$

where the shadow term modulates only the diffuse and specular contributions. In the ablation, we instead apply the shadow term to all components, resulting in

$$\text{Shadow-on-SSS: } (c_d f_d + c_s f_s + c_{sss} f_{sss}) \cdot S(\mathbf{x}) \tag{14}$$

This design highlights the distinction between surface occlusion and volumetric diffusion. The SSS module is directly supervised by multi-view images, which already account for visibility and attenuation effects. Introducing $S(\mathbf{x})$ to the SSS component results in a secondary attenuation that reduces back-lit translucency and causes overly dark scattering. Our ablation study confirms that this approach leads to lower reconstruction quality across translucent, semi-translucent, and opaque scenes.

Tab. 5 reports per-scene metrics. The "Shadow-on-SSS" variant performs worse in all categories, with the largest differences observed in scenes that contain strong subsurface transport. Fig. 13 shows representative qualitative results that demonstrate the loss of translucency and soft scattering when the shadow term is applied to the SSS component.

F Additional Relighting Baselines

This section provides additional comparisons on the Synthetic OLAT dataset and further clarifies the modeling differences among recent relighting approaches. Existing relighting methods based on NeRF or 3D Gaussian Splatting generally differ in how illumination is parameterized and reconstructed, which directly determines their behavior under point-light relighting.

Relighting under unknown illumination. The first class of datasets contains scenes captured under *unknown* and often complex illumination, such as outdoor environments or indoor scenes dominated by global illumination. Methods in this category, including TensoIR (Jin et al., 2023), R3DG (Gao et al., 2024), IRGS (Gu et al., 2025), and GI-GS (Chen et al., 2024), must jointly infer surface reflectance, geometry, and an environment map from the observed radiance. Because the incoming light distribution is not provided, these approaches rely on explicit or residual global-illumination modeling (e.g., multi-bounce shading, occlusion volumes, or deferred visibility terms) to explain indirect energy that cannot be deduced from direct lighting alone. The recovered illumination is typically represented as a low-frequency environment map, making these methods effective for ambient relighting but fundamentally limited in reproducing the high-frequency, spatially localized behavior characteristic of point-light transport.

Relighting under known illumination. The second class of datasets follows a controlled One-Light-at-a-Time (OLAT) protocol, where each training view is illuminated by a *known* single point light with calibrated position. Approaches such as GS^3 (Bi et al., 2024), RNG (Fan et al., 2025), and ours leverage this setting, which provides explicit per-light supervision and cleanly separates geometry, BRDF, and illumination. Unlike unknown-light datasets, OLAT observations directly reveal high-frequency shading cues, including directional visibility, sharp-to-soft shadow transitions, and localized subsurface transport, so explicit global-illumination terms become unnecessary. In this setting, introducing residual multi-bounce components often leads to ambiguity by entangling reflectance and transport. Instead, OLAT-oriented methods focus on accurately modeling direct point-light transport; our approach follows this paradigm using a physically structured Gaussian-domain shading model that captures both fine-scale direct effects and the small, naturally occurring low-frequency residual energy in controlled environments.

Expanded baselines. To broaden the relighting comparison, we additionally evaluate TensoIR and R3DG on the GS^3 Synthetic OLAT dataset. Both methods are re-trained using their official implementations. Since these approaches rely on environment-map estimation, they must explain point-light observations using low-frequency illumination representations. This mismatch introduces ambiguity and typically manifests as blurred shadows, reduced directional contrast, and attenuated high-frequency shading. Quantitative results in Tab. 1b and Tab. 6 show that OLAT-targeted

methods consistently outperform unknown-lighting relighting models, and our method achieves the highest accuracy across all scenes and metrics.

Table 6: Quantitative comparison results with expanded relighting baselines: TensoIR and R3DG. The best results are colored in red.

Method \ Dataset		Translucent	AnisoMetal	Drums	FurBall	Hotdog	Lego	Average
		PSNR ↑						
TensoIR	Train	16.9800	18.1800	26.4400	20.1600	17.1900	17.5200	19.4117
	Test	15.9997	16.8545	24.8974	20.0509	17.1500	17.0317	18.6640
R3DG	Train	17.1815	18.1650	27.3024	21.6486	19.5616	19.4608	20.5533
	Test	16.4664	17.1060	25.0378	20.1443	17.1896	16.2879	18.7053
Ours	Train	32.6058	31.1077	34.2448	35.4793	32.4901	31.1434	32.8452
	Test	32.3919	30.0448	33.5514	35.1639	32.1330	30.4664	32.2919

These expanded comparisons demonstrate that relighting methods designed for unknown or environment-map illumination do not perform well in the point-light OLAT scenario, where high-frequency directional cues are essential. In contrast, known-light approaches—particularly those that incorporate physically structured transport modeling—achieve significantly more accurate and consistent results. Our method attains state-of-the-art performance on the GS^3 Synthetic OLAT dataset.

G Analysis of Shadow

Shadow Pipeline. This section analyzes the shadow term and visualizes the behavior of our visibility formulation. Although a point light theoretically produces a sharp umbra boundary, our continuous transmittance model yields smooth and geometry-aware transitions. The per-ray transmittance v_i accumulates attenuation along each shadow ray, and its spatial variation across neighboring rays naturally induces soft penumbra regions. Aggregating these ray-wise values into a per-Gaussian coarse visibility $\hat{v}_g$ further smooths local discontinuities and captures how each Gaussian contributes to shadowing. The refined shadow $S(x)$ then maps these visibility cues to pixel-space shadow intensities, suppressing residual artifacts and producing stable, physically interpretable soft shadows. As shown in Fig. 3, the progression from v_i to $\hat{v}_g$ and finally to $S(x)$ illustrates how continuous visibility modeling produces coherent, geometry-consistent soft shadows under point-light illumination.

Comparison with Screen-Space Shadow Baselines. We additionally compare our formulation with screen-space opacity-accumulation strategies used in methods such as GS^3 (Bi et al., 2024) and RNG (Fan et al., 2025), as illustrated in Fig. 14. Since these approaches compute shadowing after projection, they are highly sensitive to depth ordering and often exhibit unstable or overly sharp shadow boundaries, particularly around thin structures or regions with multi-layer occlusion. Screen-space accumulation also struggles to maintain consistency under viewpoint changes, as small perturbations in projected splat order can produce flickering or discontinuities.

In contrast, our volumetric visibility formulation integrates attenuation along the light ray in 3D, independent of screen-space ordering. This yields smoother and more geometry-consistent transitions, stable penumbra behavior, and improved handling of dense or overlapping Gaussians and concave geometry. These comparisons highlight the advantages of continuous transmittance and emphasize the importance of modeling visibility at both the ray and Gaussian levels rather than relying solely on post-projection image-space accumulation.

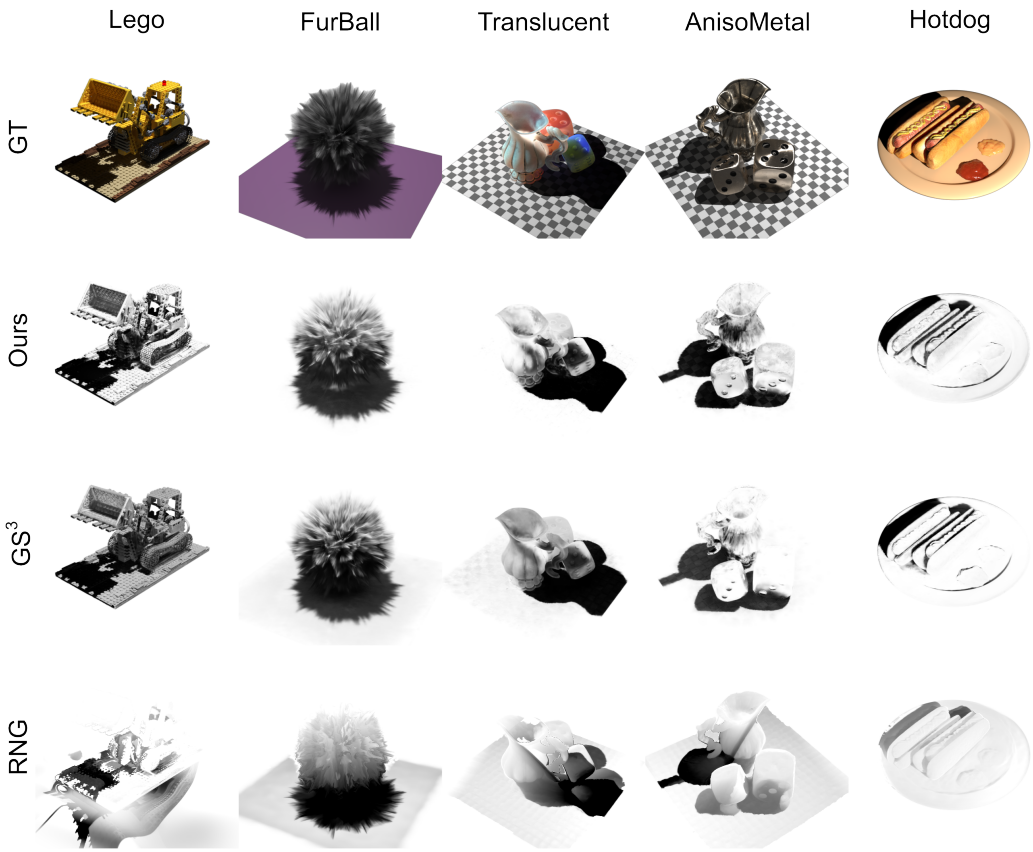


Figure 14: Comparison with screen-space shadow baselines. Top row: reference renderings. Second row: shadows produced by our method. Third and fourth rows: shadows from GS^3 (Bi et al., 2024) and RNG (Fan et al., 2025), which both rely on screen-space opacity accumulation.

H LLM USAGE

Throughout this study, we used LLMs only to assist with writing—correcting grammar and refining phrasing to improve clarity.

We **did not** use LLMs to search for or identify related works; all literature was found by the authors.

LLMs **did not** contribute to the intellectual development of the research.