RNG: Relightable Neural Gaussians

Supplementary Material

We propose RNG, a novel relightable asset with neural Gaussians. Without assumptions in the shading model and geometry types, we enable the relighting for both fluffy objects and surface-like scenes. In this supplementary material, we provide extra quality validation and metrics in Sec. 6, and discuss the choice in hybrid optimization and the network size in Sec. 7.

6. Additional validation

In Fig. 9, we visualize the obtained depth maps and shadow cues of our model of both real and synthetic objects. The depth map and shadow cues match well with our renderings and the ground truth. Our forward-deferred optimization strategy provides us with high-quality geometries, and together with reasonable shadow information, our model predicts close results to the reference.

In Table 3, we compare our neural appearance model to the vanilla 3DGS with SHs. We run both methods on datasets rendered with environment lighting and compare the NVS quality. Since the shadow cue and depth refinement MLP are disabled in our method, we only run forward shading for our method. Overall, our neural radiance representation provides more capacity and power in various scenes.

In Fig. 13, we demonstrate the effectiveness of our shadow cue by showcasing an example where the shadow quality dominates. We observe a significant quality improvement in the shadow by adding shadow cues into our model.

In Fig. 10, we move the light source towards and away from the object, showing the different lighting effects. Thanks to the shadow cue, our model shows robustness under different light conditions and can produce reasonable light effects.

In Fig. 15, we show the relighting results of RNG under moving point lights. Each column in the figure shows a different light direction. Our model can render scenes under novel lights with realistic appearance and high-quality details, and can properly model the self-shadowing effects. We suggest the reader refer to the supplementary video for more validation.

7. Additional discussion

The choice in hybrid optimization. To improve the shadow quality and avoid blurry artifacts, we suggest a deferred shading process, regularizing the appearance of shadows in the image space. As shown in Fig. 11, the appearance of shadows is not obtained by blending Gaus-

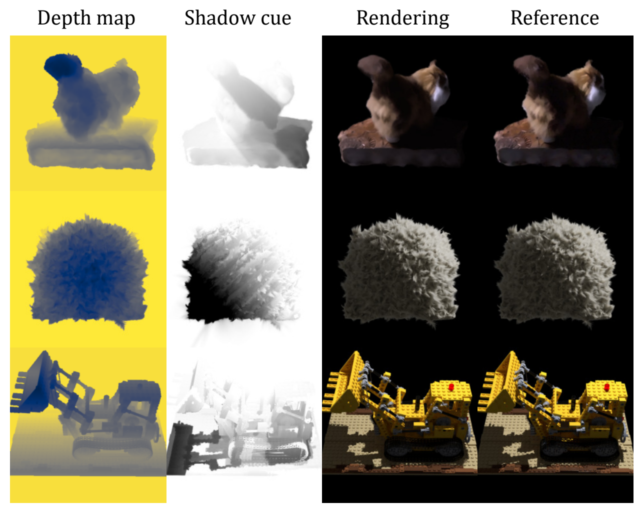

Figure 9. The visualization of depth maps and shadow cue maps of our model for different objects. The two-stage hybrid optimization strategy provides clear and accurate geometries, and the shadow cue also correctly reflects the visibility information.

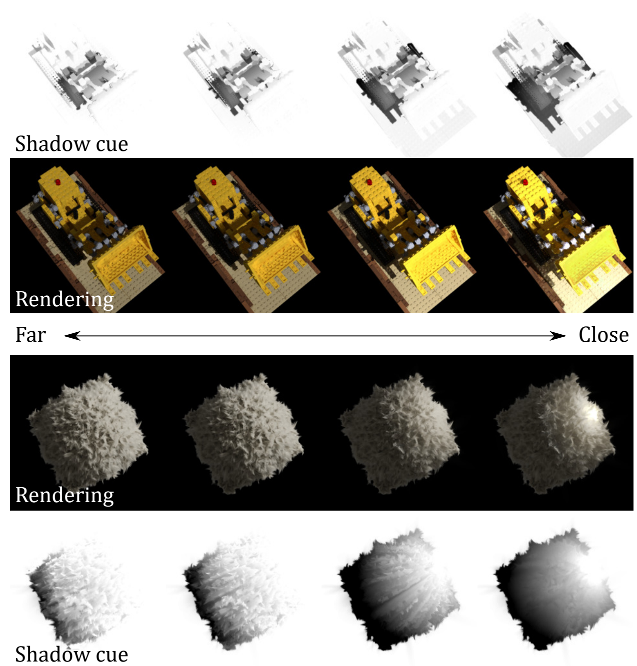

Figure 10. The renderings and visualization of shadow cues when moving the point light towards and away in the scene. The shadow cues correctly reflect the movement of the light source, and our model produces plausible renderings.

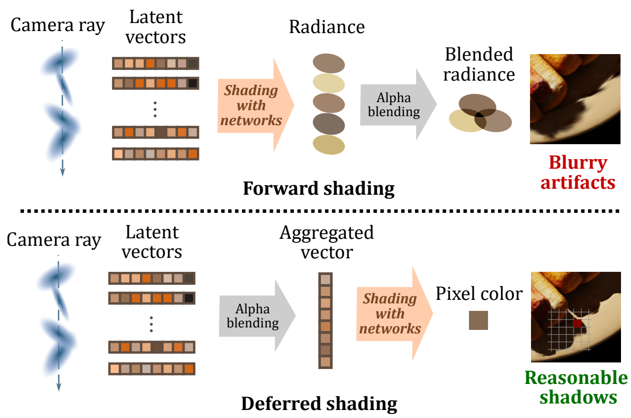

Figure 11. The difference in producing shadows between forward shading and deferred shading. For comparison, alpha blending usually leads to wrong and blurry shadows, while deferred shading provides more capability and flexibility to produce plausible colors for the blended image space features.

sians, but directly decided in the image space, avoiding the artifacts brought by the blending operation. However, according to our observation, forward shading produces better geometry, while deferred shading leads to outliers and floaters. We show such observations in Fig. 12. Therefore, we suggest a two-stage hybrid optimization strategy in the end, preserving both the geometry and shadow qualities.

Network sizes. In Fig. 14, we show the prediction accuracy with varying sizes of the neural Gaussian decoder. We use FURBALL for example, since this scene includes both complex appearance and shadows. All metrics are normalized so that higher values indicate better performance for ease of comparison. The variants are tagged by the number of hidden units and hidden layers, and our choice is $(256, 4)$. Our choice yields the best results among all variants, balancing between quality and computational complexity.

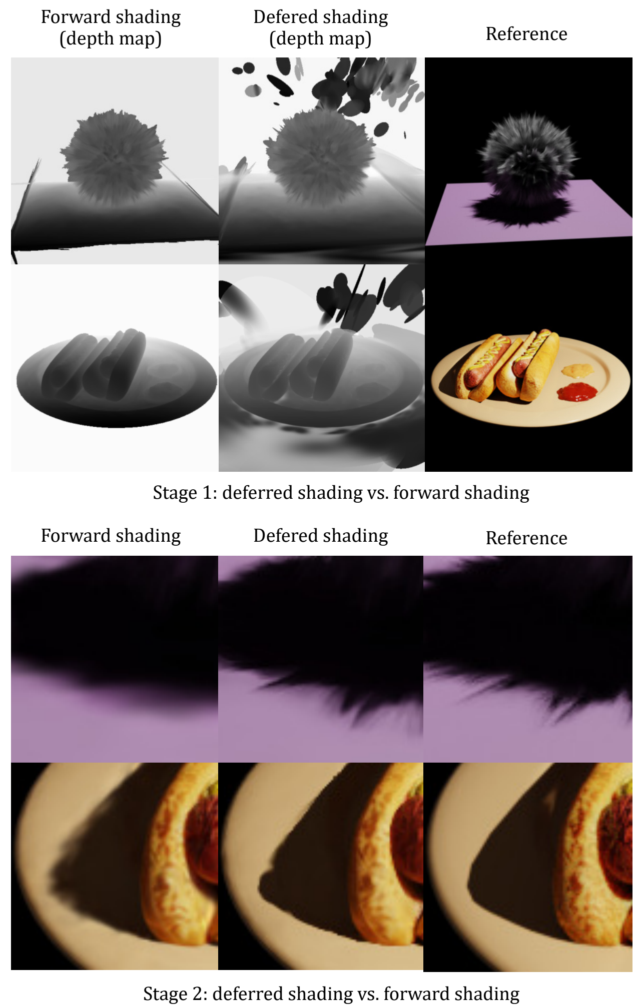

Figure 12. The depth/shadow cue visualization and rendering comparison between forward and deferred shading. Forward shading produces better geometry, while deferred shading produces better shadows.

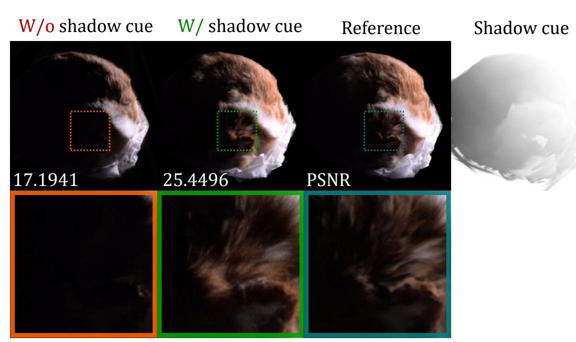

Figure 13. The quality difference at shadow regions with/without the shadow cues. With the shadow cues applied, the network can further improve the quality of dark pixels, as well as provide a clearer boundary for the shadow region.

Table 3. NVS Comparison of our neural appearance model and the vanilla SH-based 3DGS under static environment lighting. We provide (from left to right) PSNR(↑), SSIM (↑) and LPIPS (↓) for comparison, and the prevailing results are marked as **bold**. Our neural radiance representation is more flexible and powerful than the SHs in terms of NVS quality. Note that in this case, the shadow cue and depth refinement MLP are disabled in our method.

Scene	Vanilla 3DGS			Ours (forward shading only)		
Armadillo	45.8581	0.9959	0.0023	**49.2875**	**0.9976**	**0.0009**
CupPlane	43.1947	0.9957	0.0021	**47.3267**	**0.9974**	**0.0010**
Ficus	36.9734	0.9937	0.0038	**39.8921**	**0.9964**	**0.0019**
Flowers	36.0157	0.9918	0.0049	**37.2329**	**0.9941**	**0.0036**
HairBlue	38.5114	0.9766	0.0197	**39.5742**	**0.9811**	**0.0146**
Hotdog	35.3799	0.9941	0.0047	**42.6534**	**0.9972**	**0.0015**
Lego	42.0764	0.9962	0.0021	**44.9111**	**0.9981**	**0.0009**
Average	39.7157	0.9920	0.0057	**42.9826**	**0.9946**	**0.0035**

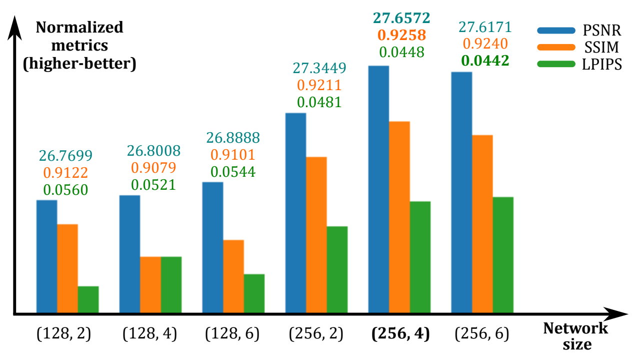

Figure 14. The comparison of variants with different network sizes. We test on FURBALL dataset. All metrics are normalized to be higher-better, and the best values of each metric are marked as **bold**. Our chosen configuration $(256, 4)$ achieves the balance between quality and complexity.

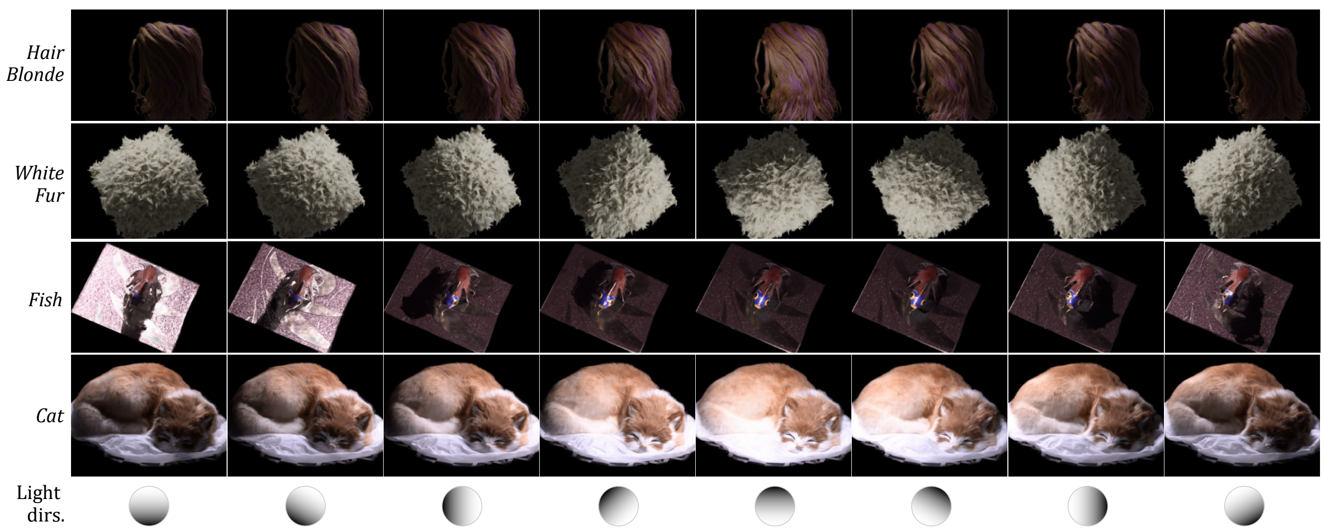

Figure 15. The relighting results of RNG. Each column shows a different point light direction. Our model can render scenes in novel views and lights with realistic appearance with high-quality details, and can properly model the self-shadowing effects.