Supplementary Material for CoR-GS

Overview

We organize the material as follows. We first provide more results of the point disagreement and rendering disagreement in Sec. A. Then, the hyperparameters and the design of our method are discussed in Sec. B. We provide more visualization results in Sec. C. Finally, we discuss the future work of our method in Sec. D.

A Additional Results of Disagreement

A.1 Behaviors of Disagreement

We provide the point disagreement and the rendering disagreement recorded in more scenes. We also provide the two disagreements of our CoR-GS. The results are shown in Figs. I and II. We observe that the two 3D Gaussian radiance fields exhibit different behaviors in various scenes and the disagreements increase significantly during densification. Integrating co-pruning and pseudo-view co-regularization, CoR-GS effectively suppresses the point disagreement and the rendering disagreement of vanilla 3DGS.

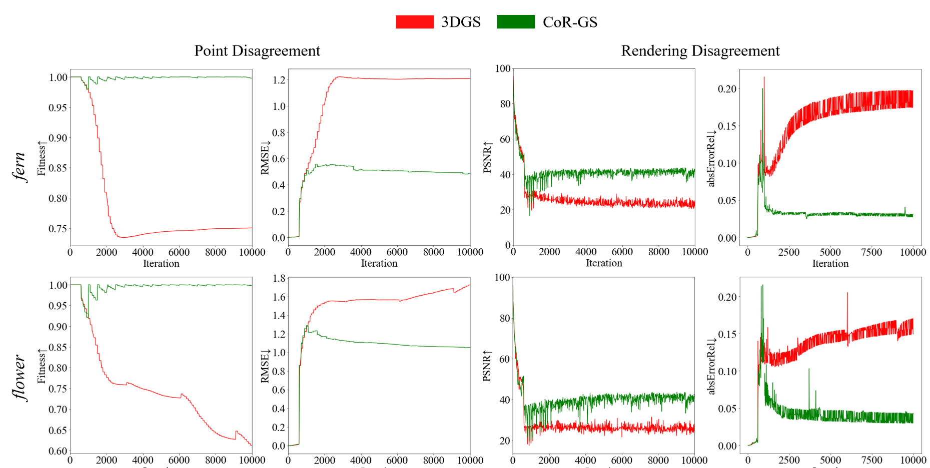

Fig. I: The recorded different behaviors of two 3d Gaussian radiance fields on the 3-view LLFF dataset during training.

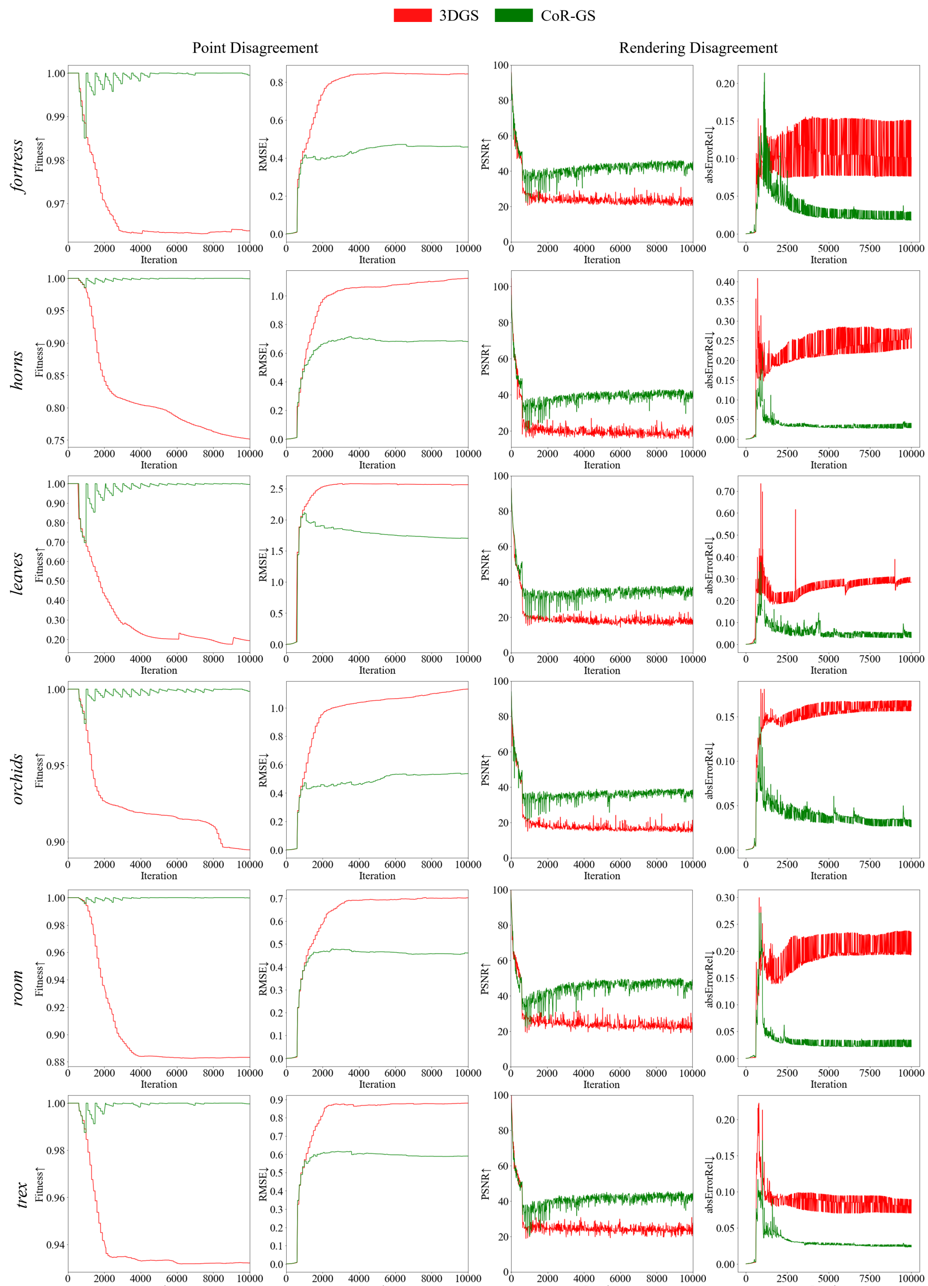

Fig. II: The recorded different behaviors of two 3d Gaussian radiance fields on the 3-view LLFF dataset during training.

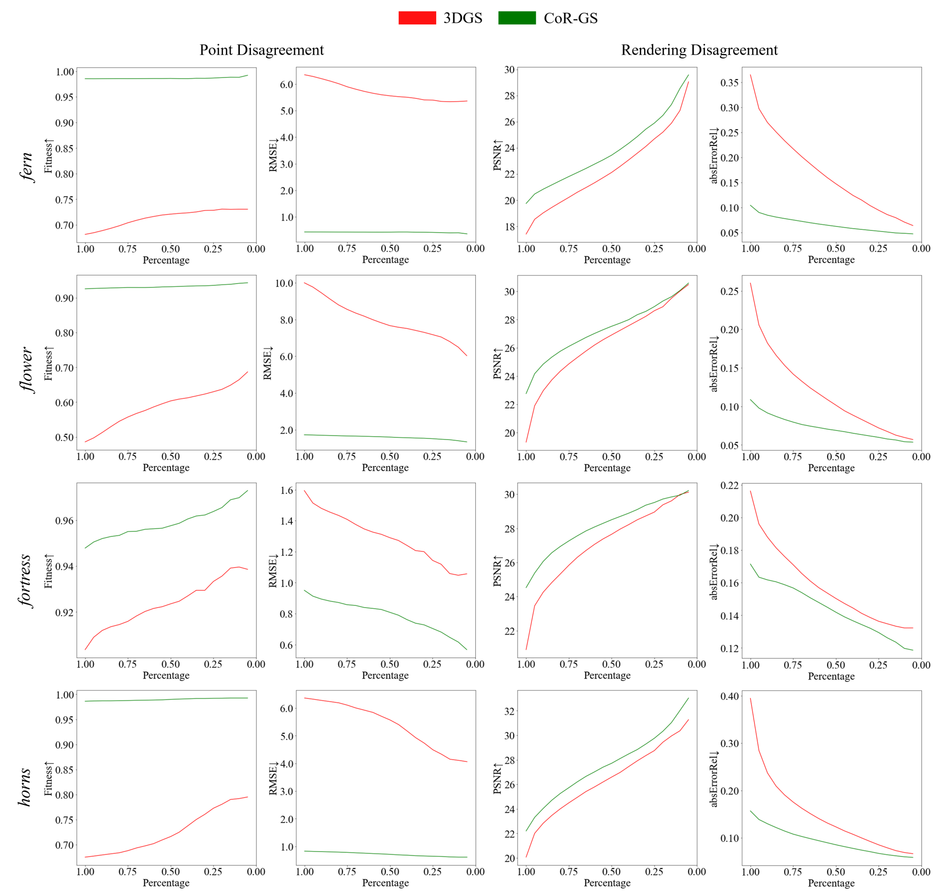

Fig. III: The correlation between the disagreements and reconstruction quality on the 3-view LLFF dataset. With the reduction of regions with higher disagreement scores, the reconstruction quality averaging the remaining regions continuously improves.

A.2 Disagreement and Accurate Reconstruction

In Figs. III and IV, we provide results on more scenes to demonstrate the negative correlation between the disagreement and accurate reconstruction. By suppressing two disagreements, CoR-GS has an overall improvement in reconstruction quality to vanilla 3DGS, which is reflected in areas with different disagreement scores. We further provide visualization of the reconstructed Gaussian points in Fig. V.

B Additional Ablation Results

We provide more ablation results of the distance threshold in co-pruning and the regularization term in pseudo-view co-regularization in Tab. I.

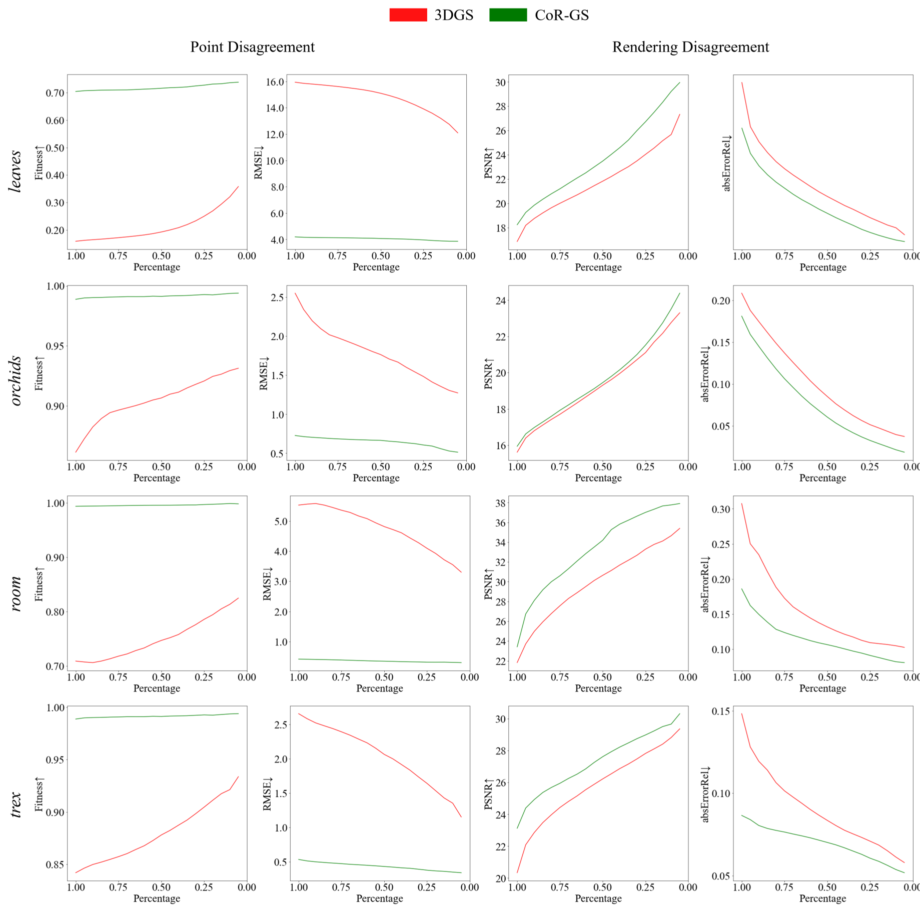

Fig. IV: The correlation between the disagreements and reconstruction quality on the 3-view LLFF dataset. With the reduction of regions with higher disagreement scores, the reconstruction quality averaging the remaining regions continuously improves.

B.1 Distance Threshold

The distance threshold controls how co-pruning considers Gaussians as no-matching. We see the threshold $\tau = 5$ and $\tau = 10$ get the best results. The threshold $\tau = 3$ performs worse than the final implementation $\tau = 5$, which prunes Gaussians strictly. The threshold $\tau = 30$ imposes very loose constraints on Gaussians, which perform similarly to solely with pseudo-view co-regularization.

B.2 Color and Depth Co-regularization

We also ablate the regularization term with rendered color and depth. Following the depth-regularized 3DGS [1, 3], we compute the depth loss with the Pearson

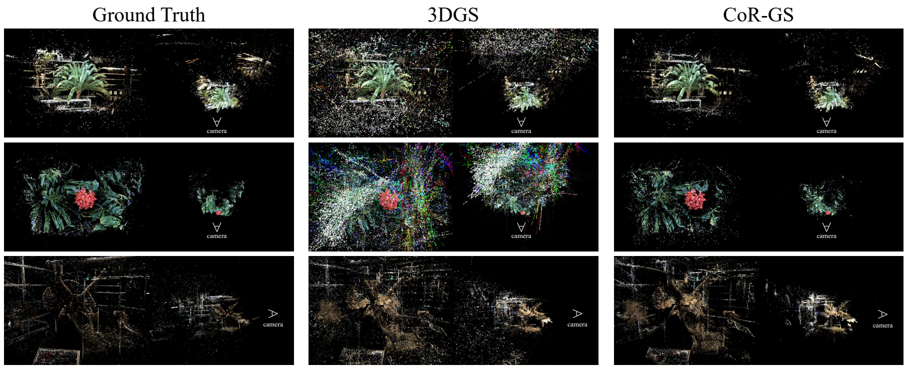

Fig. V: Visulazation of point clouds of 3D Gaussian radiance fields. We obtain the ground truth by training radiance fields with dense views. CoR-GS reconstructs more compact representations, and are more similar to dense-view representations.

Table I: Ablation study of CoR-GS on the 3-view LLFF setting.

Distance Threshold τ	PSNR ↑	SSIM ↑	LPIPS ↓
3	20.36	0.709	0.198
5	20.45	**0.712**	**0.196**
10	**20.46**	0.711	**0.196**
30	20.30	0.707	0.198

Color Co-reg.	Depth Co-reg.	PSNR ↑	SSIM ↑	LPIPS ↓
✓		**20.45**	**0.712**	0.196
	✓	20.01	0.685	0.206
✓	✓	**20.45**	0.711	**0.195**

Number of Radiance Fields	PSNR ↑	SSIM ↑	LPIPS ↓
2	20.45	0.712	0.196
3	20.58	0.721	**0.190**
4	**20.61**	**0.723**	**0.190**

3DGS Baseline	PSNR ↑	SSIM ↑	LPIPS ↓
FSGS [3]	20.43	0.682	0.248
CoR-GS	20.45	0.712	0.196
CoR-FSGS	**20.93**	**0.730**	**0.194**

correlation. We find that although depth co-regularization alone helps, the impact becomes very weak when used in conjunction with color co-regularization. This is because color co-regularization also imposes constraints on depth information due to the sorting process implied for rendering. Therefore, we do not perform co-regularization on depth maps in this case.

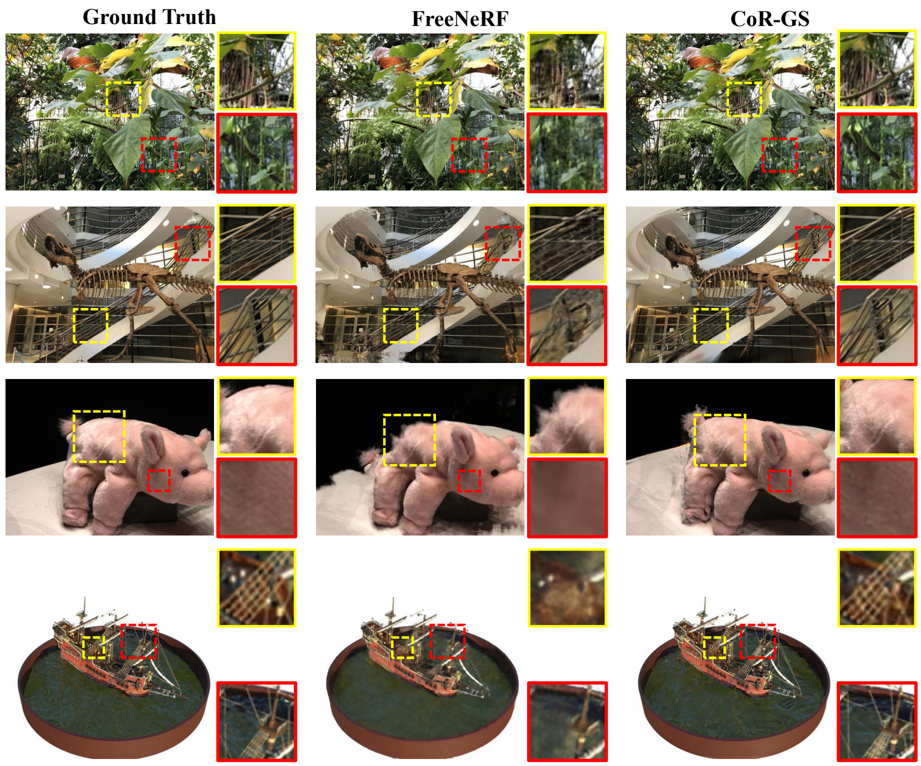

Fig. VI: Qualitative comparison of FreeNeRF and CoR-GS of 3-view LLFF, 3-view DTU and 8-view Blender.

B.3 Number of Radiance Fields

The co-regularization between two 3D Gaussian radiance fields can be naturally extended to utilizing more radiance fields. In this implementation, co-pruning prunes Gaussians exhibit high point disagreement concerning any of the other radiance fields. Pseudo-view co-regularization suppresses the rendering disagreement concerning each of the other radiance fields. With more radiance fields, the disagreement reflects the inaccurate reconstruction more accurately, resulting in further improvements than using two fields.

B.4 Incorporating with Depth Regularized 3DGS

Equipped with our co-regularization method, 3DGS can render high-quality novel views with sparse training views. Considering depth regularization has been proven effective in the sparse view setting, we apply our co-regularization method to the depth regularized method FSGS [3]. The ablation result demonstrates that our method can work well with depth regularization, and especially exhibits an advantage on the image structural SSIM score.

C Additional Visualization Results

C.1 Comparison with NeRFs

In Fig. VI, we provide the visualization comparison between state-of-the-art NeRF-based method FreeNeRF [2] and our CoR-GS. Compared to FreeNeRF, our method demonstrates advantages in reconstructing realistic high-frequency details and better geometry of thin structures.

C.2 CoR-GS Visualizations

This section provides more visualization results of CoR-GS. Figs. VII to X provide the novel view synthesis of CoR-GS on DTU, LLFF, Mip-NeRF360 and Blender datasets.

D Discussion and Future Work

In this paper, we propose to regularize sparse-view 3DGS from a co-regularization perspective and validate its wide effectiveness in various scenarios, with different numbers of input views and with the depth regularization method. This paper demonstrates a negative correlation between disagreed behaviors of 3D Gaussian radiance fields and the reconstruction quality in the sparse-view setting. We will investigate utilizing the disagreed behaviors in more 3DGS applications, such as video reconstruction with spatially and temporally sparse input.

References

1. Xiong, H., Muttukuru, S.: Sparsegs: Real-time 360° sparse view synthesis using gaussian splatting (10 2023) 4
2. Yang, J., Pavone, M., Wang, Y.: Freenerf: Improving few-shot neural rendering with free frequency regularization. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 8254–8263 (2023) 7
3. Zhu, Z., Fan, Z., Jiang, Y., Wang, Z.: Fsgs: Real-time few-shot view synthesis using gaussian splatting (2023) 4, 5, 6

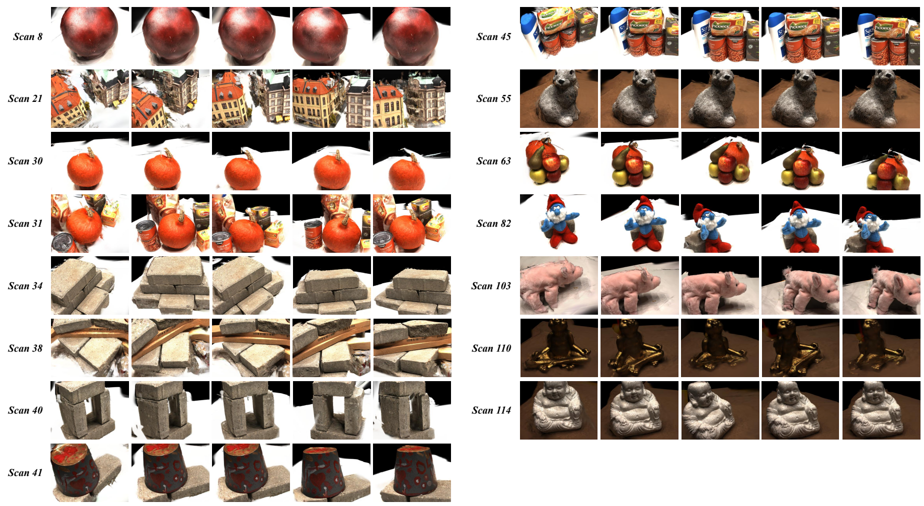

Fig. VII: Qualitative Visualizations of the 3-view DTU setting.

Fig. VIII: Qualitative Visualizations of the 3-view LLFF setting.

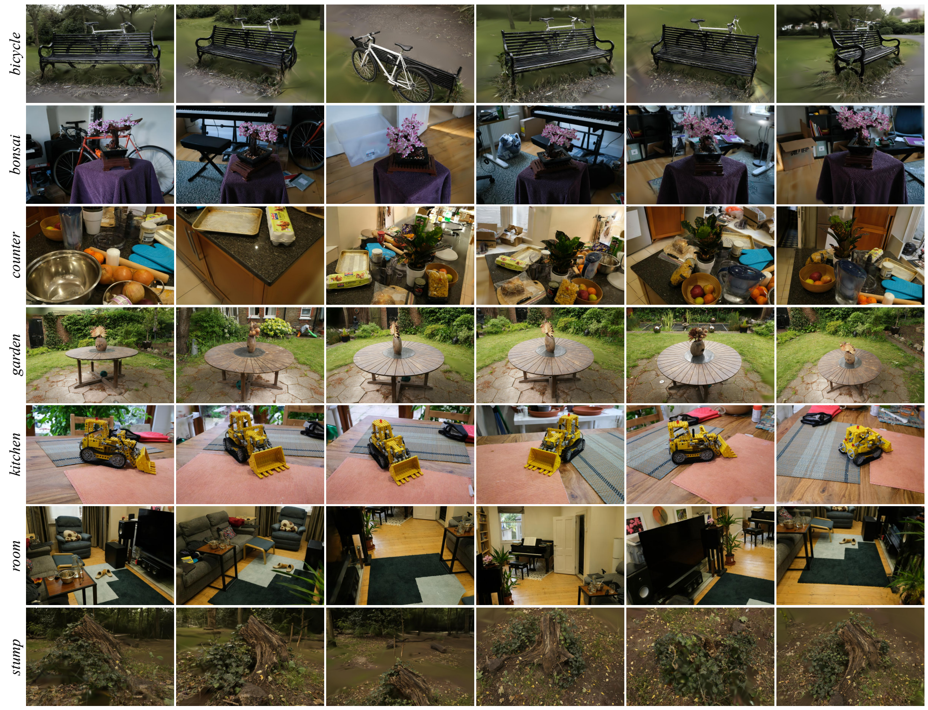

Fig. IX: Qualitative Visualizations of the 24-view Mip-NeRF360 setting.

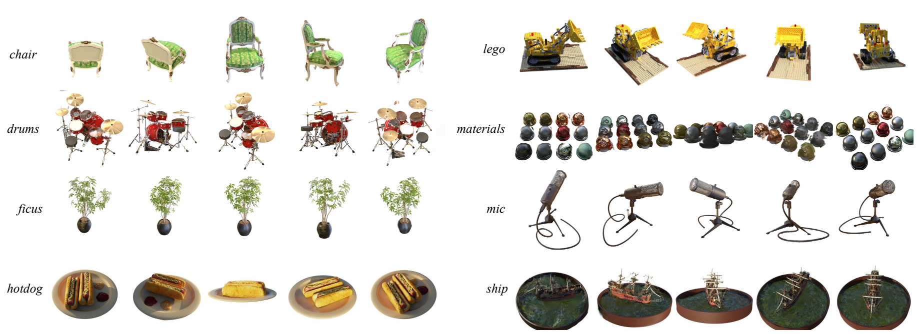

Fig. X: Qualitative Visualizations of the 8-view Blender setting.