

Appendix

A. Voronoi vs naïve decomposition – Table 3

To test that our learned decomposition actually contributes to an improvement in accuracy, we test against a version of DeRF with a naïve decomposition – where cells are fixed into a grid. We find that the learned decomposition results in improvement in all metrics.

B. Additional experiments on real data

Per-scene results – Figures 10–17. In this section we provide further comprehensive experiments on the real capture scenes showing how DeRFs with up to 16 heads perform in terms of reconstruction quality and inference cost.

We find that DeRF models provide the best quality-computation trade-off, for all cases in terms of LPIPS, and almost all cases for PSNR and SSIM. Moreover, for the strong majority of scenes, highly decomposed DeRF models give the best results, especially for the perceptual and structural metrics. We also find the advantage of DeRF is most consistent for scenes where the absolute error is higher, i.e. where gains in performance are needed the most. All experiments in this section use the same sample counts as [13].

Variance Test – Table 4. Additionally, to help explain the increasingly chaotic nature of the graphs as quality increases, we perform an experiment in which we train eight identical DeRF models on the scene with the highest reconstruction accuracy (the “room” scene). In this test we find a significant difference in results depending only on initialization,

especially for PSNR. This suggests that, while the current experiments already provide enough evidence for the efficacy of DeRF, extracting robust statistics via multiple runs may deliver a more conclusive answer at the expense of compute. We note however, our results already demonstrate that decomposing the scene is almost always better, especially in terms of perceptual metrics, and when less computation is used for inference.

Video supplementary. We further direct the interested readers to the video supplementary that demonstrates the rendering quality and the decomposition in 3D.

	PSNR↑	SSIM↑	LPIPS↓
DeRF	28.55	0.89	0.25
DeRF-Grid	28.07	0.88	0.28

Table 3. **Voronoi vs grid decomposition –** Quantitative results for two equal-capacity 64-head, 32-unit models trained on the “room” scene. The DeRF model uses our Voronoi decomposition and the DeRF-Grid uses an untrained decomposition where the cells are arranged in a $4 \times 4 \times 4$ grid. Our Voronoi decomposition brings significant benefit in terms of rendering quality.

	PSNR↑	SSIM↑	LPIPS↓
Minimum	29.15	0.92	0.15
Maximum	29.86	0.93	0.16

Table 4. **Variation across initialization –** We report the variation in rendering quality resulting from different initializations by showing the minimum and maximum metrics, achieved by a set of 16-head, 128-unit DeRF models trained on the “room” scene.

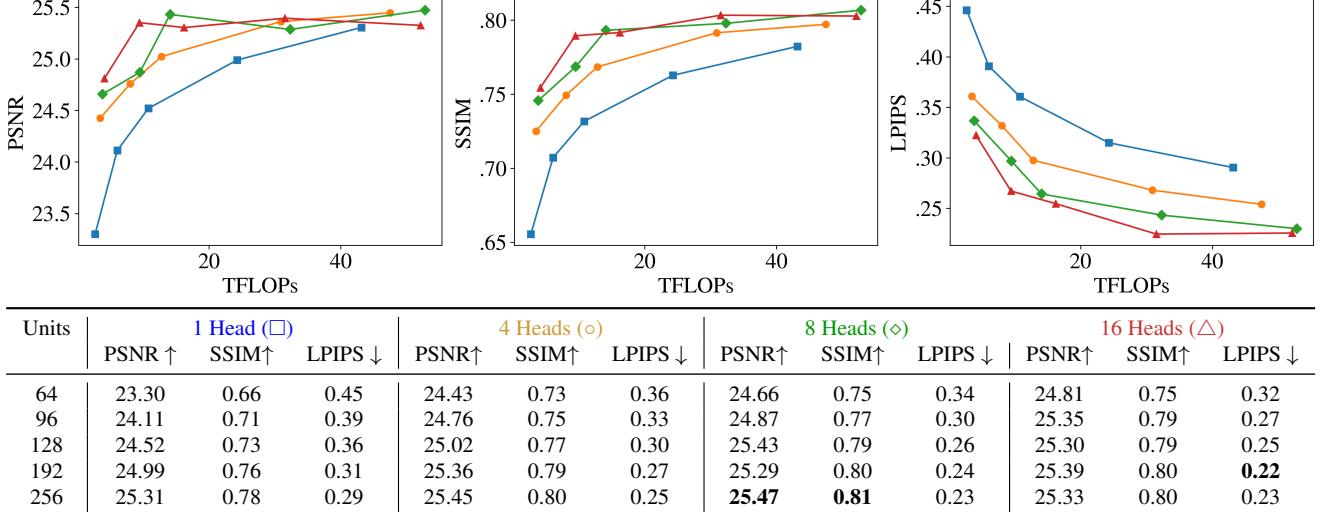


Figure 10. Quantitative results for the “fern” scene.

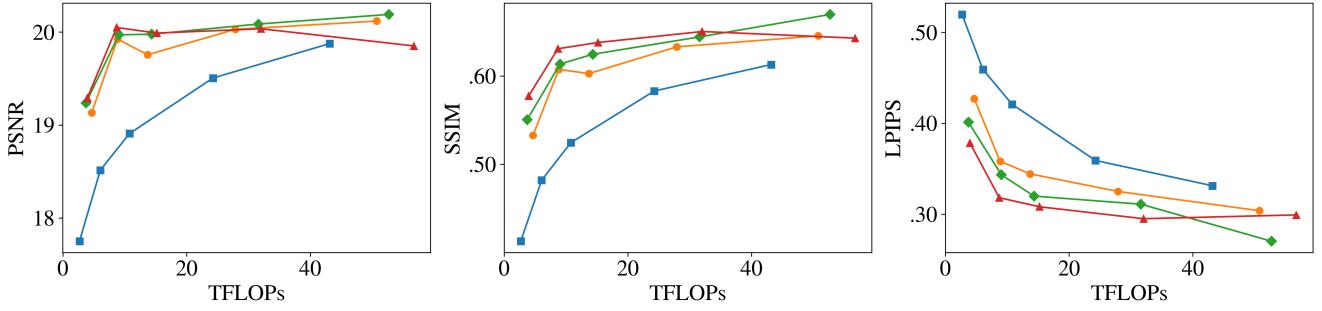


Figure 11. Quantitative results for the “orchids” scene.

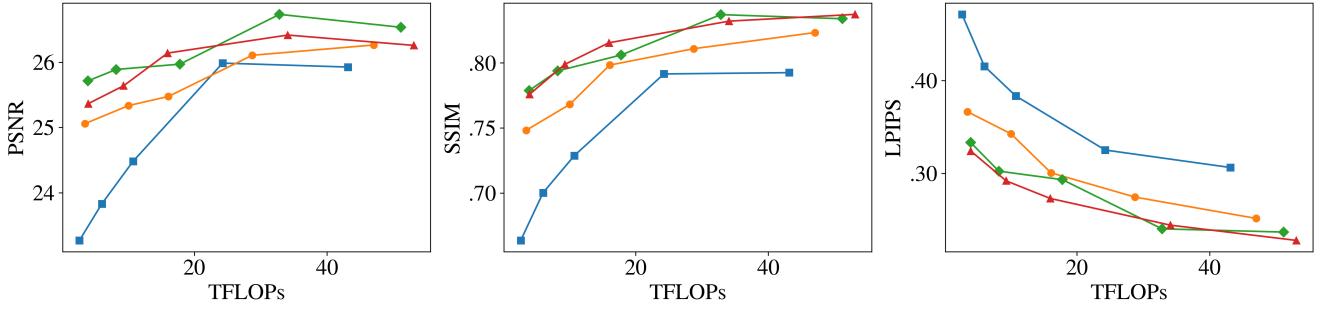


Figure 12. Quantitative results for the “horns” scene.

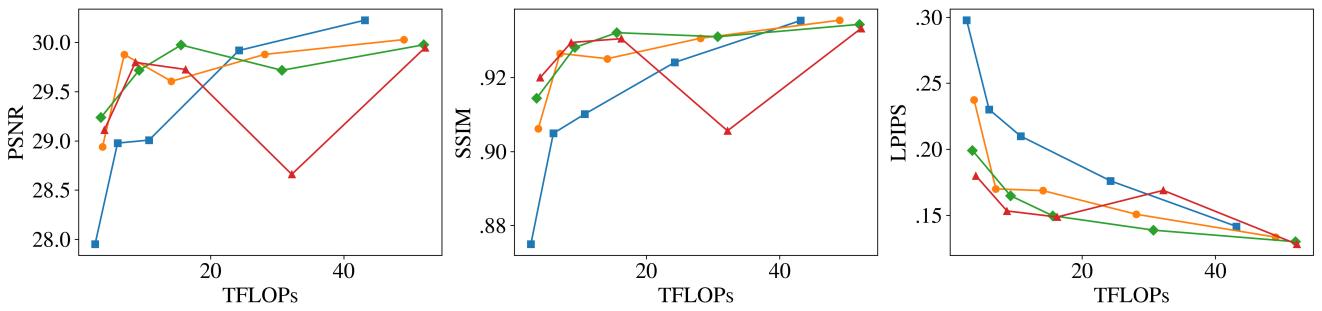
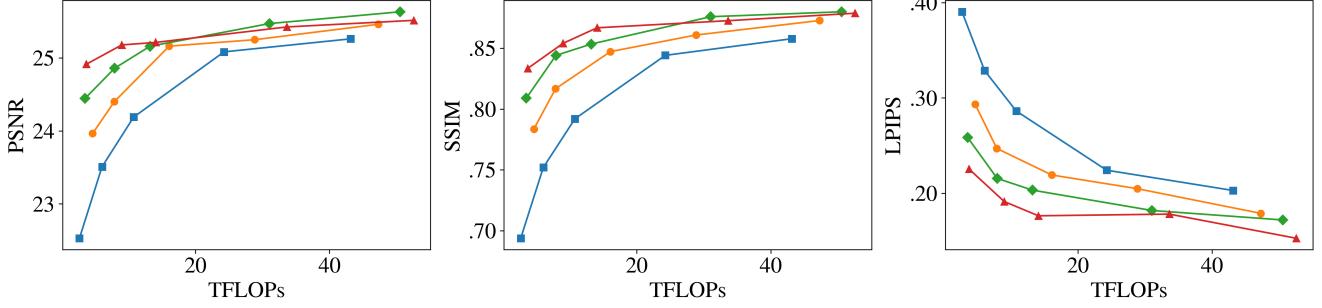
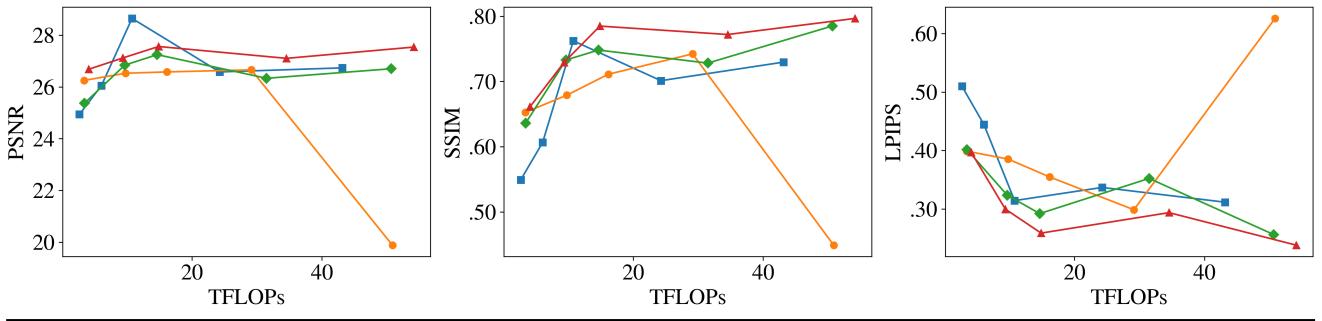


Figure 13. Quantitative results for the “room” scene.



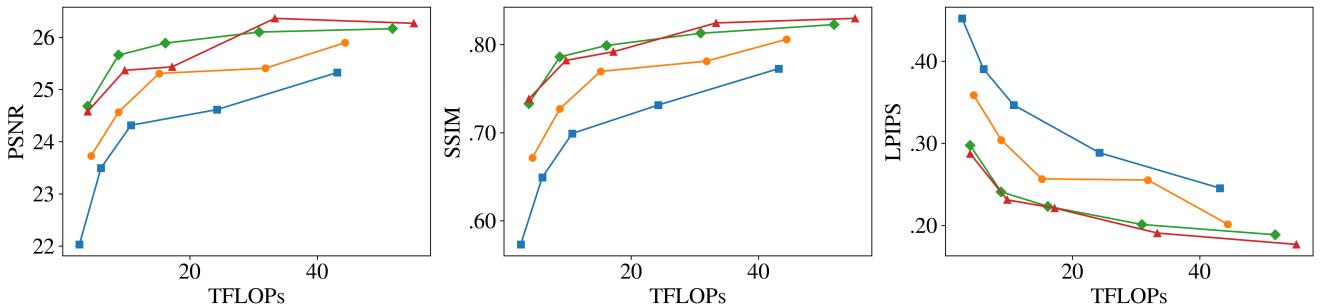
Units	1 Head (□)			4 Heads (○)			8 Heads (◇)			16 Heads (△)		
	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓
64	22.53	0.69	0.39	23.96	0.78	0.29	24.44	0.81	0.26	24.91	0.83	0.23
96	23.51	0.75	0.33	24.40	0.82	0.25	24.86	0.84	0.22	25.18	0.85	0.19
128	24.19	0.79	0.29	25.16	0.85	0.22	25.16	0.85	0.20	25.21	0.87	0.18
192	25.08	0.84	0.22	25.25	0.86	0.20	25.47	0.88	0.18	25.42	0.87	0.18
256	25.26	0.86	0.20	25.46	0.87	0.18	25.63	0.88	0.17	25.51	0.88	0.15

Figure 14. Quantitative results for the “trex” scene.



Units	1 Head (□)			4 Heads (○)			8 Heads (◇)			16 Heads (△)		
	PSNR↑	SSIM↑	LPIPS ↓	PSNR↑	SSIM↑	LPIPS ↓	PSNR↑	SSIM↑	LPIPS ↓	PSNR↑	SSIM↑	LPIPS ↓
64	24.94	0.55	0.51	26.25	0.65	0.40	25.37	0.64	0.40	26.68	0.66	0.40
96	26.05	0.61	0.44	26.53	0.68	0.39	26.84	0.73	0.32	27.12	0.73	0.30
128	28.65	0.76	0.31	26.58	0.71	0.35	27.25	0.75	0.29	27.56	0.78	0.26
192	26.58	0.70	0.34	26.66	0.74	0.30	26.33	0.73	0.35	27.10	0.77	0.29
256	26.74	0.73	0.31	19.88	0.45	0.63	26.70	0.79	0.26	27.54	0.80	0.24

Figure 15. Quantitative results for the “fortress” scene.



Units	1 Head (□)			4 Heads (○)			8 Heads (◇)			16 Heads (△)		
	PSNR↑	SSIM↑	LPIPS ↓	PSNR↑	SSIM↑	LPIPS ↓	PSNR↑	SSIM↑	LPIPS ↓	PSNR↑	SSIM↑	LPIPS ↓
64	22.03	0.57	0.45	23.73	0.67	0.36	24.68	0.73	0.30	24.57	0.74	0.29
96	23.49	0.65	0.39	24.57	0.73	0.30	25.66	0.79	0.24	25.36	0.78	0.23
128	24.31	0.70	0.35	25.30	0.77	0.26	25.88	0.80	0.22	25.43	0.79	0.22
192	24.61	0.73	0.29	25.40	0.78	0.26	26.10	0.81	0.20	26.36	0.82	0.19
256	25.32	0.77	0.25	25.89	0.81	0.20	26.16	0.82	0.19	26.26	0.83	0.18

Figure 16. Quantitative results for the “flower” scene.

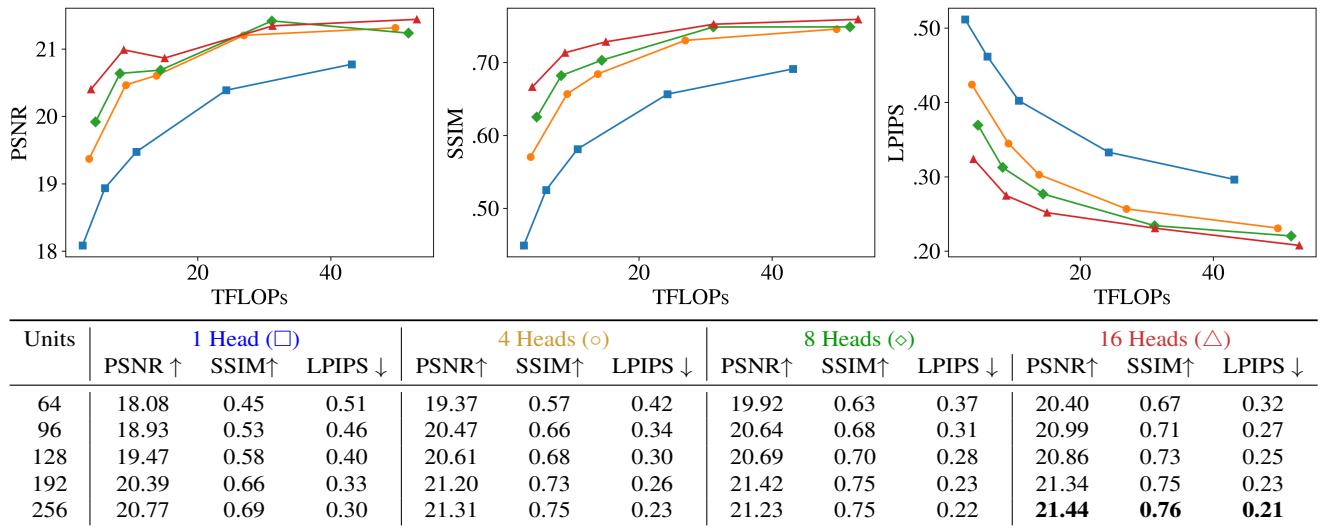


Figure 17. Quantitative results for the “leaves” scene.

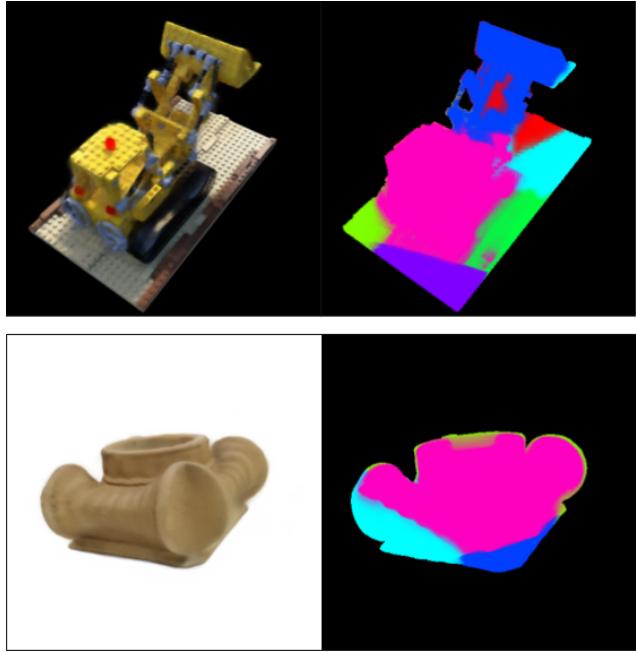


Figure 18. **Synthetic Data** – Renders and decompositions for the NeRF [13] “lego” scene (top) and DeepVoxels [18] “greek” scene (bottom).

C. Synthetic data – Figure 18

While we are mainly interested in photorealistic scenes, our method can also learn decomposed models from synthetic data. To show this we provide renders and visualize the decompositions for DeRF models trained on two synthetic scenes. Each model uses 192 units and 8 heads.