Lagrangian Hashing for Compressed Neural Field Representations

Supplementary Material

We visualize the Lagrangian representation across $\tilde{L}(=2)$ levels, the effect of learnable Gaussian standard deviation, other comparisons with PointNeRF [2] and PAPR [3] and more detailed quantitative results for NeRF reconstruction task. Furthermore, we provide more rendering results in https://theialab.github.io/laghashes.

6 Limitations

Our method is currently restricted to object-centric scenes. Our approach does not achieve an effective Eulerian-Lagrangian hybrid representation for unbounded scenes when trained with the scene contraction function. Due to our emphasis on learning a compact representation of 3D objects, we left further exploration of Lagrangian representations for unbounded scenes to future work.

7 Lagrangian representation across $\tilde{L}(=2)$ levels

We show our lagrangian representation across $\tilde{L}$ LoDs. As shown in Fig. 9, both scales learn the complete point representation and the fine LoD scale tends to capture more high-frequency details.

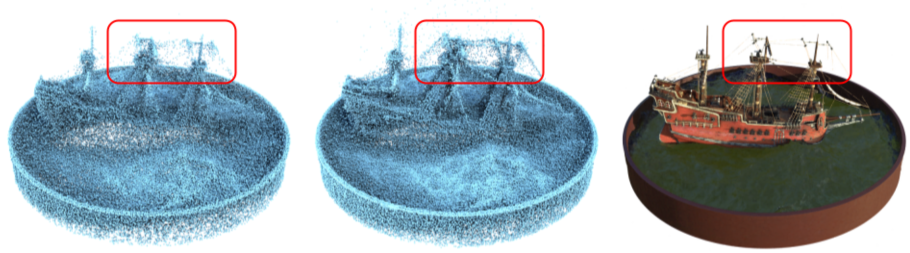

Pre-final LoD | Final LoD | Rendered RGB image

Fig. 9: Notice the mast of the ship(highlighted region), where we see that the final LoD represents high-frequency details better than the pre-final LoD.

8 Learnable vs fixed standard deviation

We quantitatively compare the learnable Gaussian standard deviation(σ) vs fixed Gaussian standard deviation(σ). We show that the fixed standard deviation, sim-

Table 5: Learnable vs Fixed standard deviation: the fixed standard deviation (preferred in our method) is even slightly better than the learnable ones.

Method	# Params	Truck	Barn	Family	Caterpillar	Avg.
Fixed $\sigma(B = 2^{14})$	0.92M	26.77	26.75	32.30	25.94	27.92
Trainable $\sigma(B = 2^{14})$	1.02M	26.67	26.68	32.22	25.88	27.86

pler and with a smaller number of parameters, achieves a slightly better performance. Note this is unsurprising because the fixed standard deviation adaptive to the grid size has the guaranteed spatial coverage whereas the learnable standard deviation might cause degenerate cases (Gaussians that are too big or too small).

Table 6: Comparisons with other point-based representations, PointNeRF [2], PAPR [3], and 3DGS [1]: Our method achieves better/equivalent performance (PSNR↑) to other representation, even though we don't use point initialization as in [2] and don't use 2D U-Net to remove artifacts from the screen space as in [3].

Method	# Params	Lego	Mic	Materials	Chair	Hotdog	Ficus	Drums	Ship	Avg.
InstantNGP($B = 2^{19}$)	12.10M	35.67	36.85	29.60	35.71	37.37	33.95	25.44	30.29	33.11
3DGS(#G=210k) [1]	12.35M	35.89	36.71	30.48	35.37	38.05	35.48	26.24	31.64	33.73
PointNeRF [2]	5.00M	32.65	35.54	26.97	35.09	35.49	33.24	25.01	30.18	31.77
PAPR [3]	6.80M	32.62	35.64	29.54	33.59	36.40	36.50	25.35	26.92	32.07
3DGS(#G=110k) [1]	6.83M	35.28	36.54	30.46	34.80	37.77	35.48	26.19	31.48	33.50
Ours($B = 2^{17}$)	6.68M	35.60	36.45	29.63	35.61	37.23	33.89	25.67	30.84	33.12

Table 7: Comparisons with 3DGS [1]: When the parameter count is lowered, our method achieves better performance (PSNR ↑) while we find that 3DGS experiences a sharp decline in novel view synthesis quality.

Method	# Params	Lego	Mic	Materials	Chair	Hotdog	Ficus	Drums	Ship	Avg.
3DGS(#G=210k) [1]	12.35M	35.89	36.71	30.48	35.37	38.05	35.48	26.24	31.64	33.73
3DGS(#G=3.5k) [1]	0.18M	16.40	23.27	15.98	19.66	18.59	28.05	15.60	21.74	19.91
Ours($B = 2^{11}$)	0.18M	31.15	32.65	28.52	32.44	35.67	31.98	25.07	28.26	30.72

9 Other comparisons: PointNeRF, PAPR, and 3DGS

We further show quantitative comparison of our method with other point-based representations, PointNeRF [2], PAPR [3]. As shown in Table 6, our method quantitatively outperforms PointNeRF even though we do not use a COLMAP-based initialization for our point representation. We also quantitatively outperform PAPR while we learn 3D consistent point representation and do not use the 2D U-Net architecture that PAPR proposed to remove artifacts from the screen space. While our method performs similarly to 3DGS at high parameter settings, we observe that 3DGS experiences a significant drop in novel view synthesis quality at lower parameter counts, as shown in Table 7, whereas our method maintains comparatively high quality.

10 More Detailed Quantitative results

We provide the metric scores broken down by scene on both datasets. Table 8 and 9 shows the per-scene scores for the NeRF Synthetic dataset and Table 10 and 11 shows the scores for the Tanks & Temples dataset.

Table 8: NeRF synthetic dataset – SSIM scores

Method	Lego	Mic	Materials	Chair	Hotdog	Ficus	Drums	Ship	Avg.
InstantNGP($B = 2^{19}$)	0.977	0.989	0.945	0.985	0.980	0.981	0.934	0.859	0.956
Ours($B = 2^{17}$)	0.978	0.991	0.947	0.984	0.981	0.981	0.934	0.892	0.961

Table 9: NeRF synthetic dataset – LPIPS scores

Method	Lego	Mic	Materials	Chair	Hotdog	Ficus	Drums	Ship	Avg.
InstantNGP($B = 2^{19}$)	0.027	0.017	0.072	0.028	0.037	0.038	0.086	0.180	0.065
Ours($B = 2^{17}$)	0.027	0.015	0.070	0.024	0.036	0.049	0.083	0.139	0.055

Table 10: Tanks & Temples - SSIM scores

Method	Truck	Barn	Family	Caterpillar	Avg.
InstantNGP($B = 2^{19}$)	0.913	0.837	0.955	0.912	0.904
Ours($B = 2^{17}$)	0.910	0.848	0.954	0.910	0.906

Table 11: Tanks & Temples - LPIPS scores

Method	Truck	Barn	Family	Caterpillar	Avg.
InstantNGP($B = 2^{19}$)	0.138	0.289	0.076	0.152	0.164
Ours($B = 2^{17}$)	0.144	0.280	0.079	0.153	0.164

References

1. Kerbl, B., Kopanas, G., Leimkühler, T., Drettakis, G.: 3d gaussian splatting for real-time radiance field rendering. TOG (Proc. of SIGGRAPH) (2023)
2. Xu, Q., Xu, Z., Philip, J., Bi, S., Shu, Z., Sunkavalli, K., Neumann, U.: Point-nerf: Point-based neural radiance fields. In: CVPR (2022)
3. Zhang, Y., Peng, S., Moazeni, A., Li, K.: Papr: Proximity attention point rendering. NeurIPS (2024)