Rebuttal for LRM-Zero: Training Large Reconstruction Models with Synthesized Data

Anonymous Author(s)
Affiliation
Address
email

Table 1: *LRM-Zero* vs. *LRM-Zero-Obja* vs. GS-LRM at the 8-input-view, 256 resolution setting. Z means Zeroverse and O means Objaverse. The *LRM-Zero* (first row) and GS-LRM (second row) results are from experiment 9 in Tab. 8 and experiment 2 in Tab. 9 in main text. The *LRM-Zero-Obja* result (third row) is obtained by training on 800K *Zeroverse* data and 800K Objaverse data. While *LRM-Zero-Obja* outperforms *LRM-Zero*, it underperforms GS-LRM.

data	scaling			GSO			ABO		
	Training Steps def. 1x, 80K	Model Size def. 1x, 300M	Data Size def. 1x, 400K	PSNR ↑	SSIM ↑	LPIPS ↓	PSNR ↑	SSIM ↑	LPIPS ↓
Z	2x	1x	4x	31.15	0.960	0.034	29.02	0.935	0.064
O	2x	1x	2x	**33.12**	**0.973**	**0.024**	**31.75**	**0.957**	**0.047**
Z&O	2x	1x	4x	32.11	0.968	0.027	30.70	0.950	0.052

Table 2: Scaling down GS-LRM's training data size. When training on only 200K instead of 800K Objaverse data, GS-LRM's performance drops by only 0.1 PSNR on GSO.

id	scaling			GSO			ABO		
	Training Steps def. 1x, 80K	Model Size def. 1x, 300M	Data Size def. 2x, 800K	PSNR ↑	SSIM ↑	LPIPS ↓	PSNR ↑	SSIM ↑	LPIPS ↓
1	1x	1x	2x	**29.59**	**0.944**	**0.050**	**28.92**	**0.926**	**0.074**
2	1x	1x	0.5x	29.42	0.942	0.052	28.75	0.924	0.075

Table 3: Extending training stability experiments. When using around only 10% boolean difference augmentation, *LRM-Zero* can train stably even with the GS-LRM's default 0.5 perceptual loss weight. This shows that when we limit the ratio of boolean difference augmentation, we do not need to change any training hyperparameters from GS-LRM to stabilize training on *Zeroverse*.

	dataset			training			result
id	hf-only	boolean	wireframe	perceptual loss weight (default 0.5)	Gaussian scale clipping (default -1.2)	view angle threshold (default 60)	GSO PSNR, if finished
11	92%	8%	0%	0.2	default	default	30.86
13	85%	10%	5%	default	default	default	30.62

Submitted to 38th Conference on Neural Information Processing Systems (NeurIPS 2024). Do not distribute.