

Co3D-10-Category			Noise Ablation			2D Inpainting		
	PSNR↑ / LPIPS↓	FID↓		PSNR↑ / LPIPS↓	FID↓		PSNR↑ / LPIPS↓	FID↓
pixelNeRF	<b>17.96</b> / 0.479	158.50	With noise	17.24 / 0.40	92.23	Determ.	<b>21.35</b> / 0.11	9.18
SparseFusion	11.76 / 0.770	257.63	Ours	<b>18.19</b> / <b>0.34</b>	<b>56.64</b>	Ours	20.18 / <b>0.09</b>	<b>4.25</b>
Ours	17.62 / <b>0.368</b>	<b>66.81</b>						

Table 1: (left) Evaluation for 10 categories of Co3D. (middle) Ablation with and without noise added to camera poses. (right) 2D inpainting results.



Figure 1: Ablation experiment. Using noisy camera poses at training time leads to lower-quality renderings.




Figure 2: Results for a general model trained on 10 Co3D categories.



Figure 3: Results on the Objaverse dataset.




Figure 4: Inpainting experiment. Our training dataset has images like the input, with random patches missing. At test time, we can learn to complete the partial observations. Deterministic baseline learns a blurry completion.




Figure 5: Extracted point clouds demonstrate 3D consistent reconstructions, also see point cloud in Fig. 6

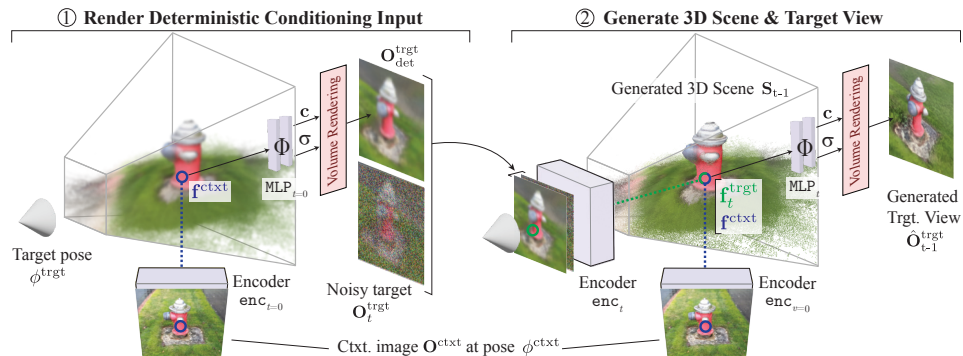


Figure 6: Updated Pipeline figure that depicts our inverse graphics pipeline.