	PSNR		
No. Views	3	6	100
Mip-NeRF	17.73	23.87	36.65
Rendered	20.32	25.90	35.09
URF	20.09	26.12	36.36
Ours	23.81	27.91	36.96

Table 5.1.: Comparison of model performance. While depth supervised methods (Rendered depth, Urban Radiance Fields [Rem+22]) improve results over the base model for few views, they adversely affect reconstruction quality when supervision is dense. Our loss by contrast consistently improves reconstruction results.

We hypothesize that our method encourages correctly initialized densities in regions that are multiview consistent, while simultaneously discouraging floaters and density in general in regions that we know to be empty. The key advantage over URF's formulation is the consideration of bounded intervals that let the network retain expansive volumes that are expressive and multiview consistent, rather than constricting entire volumes, supervised by discrete pixel-wise quantities, to predict δ-shaped densities.

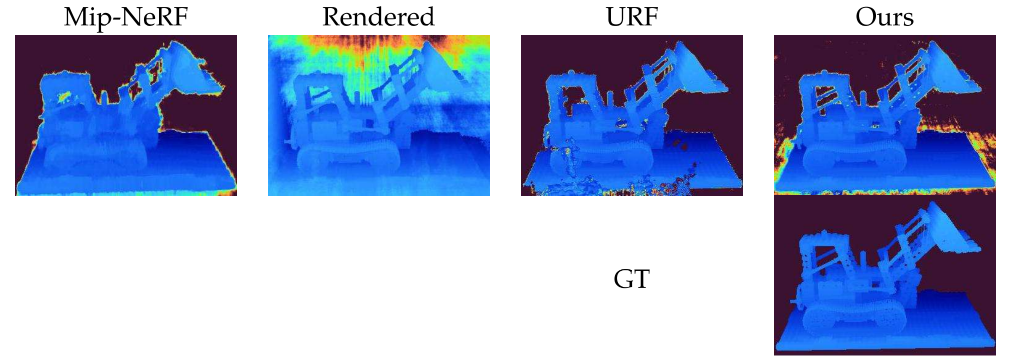

Figure 5.2.: Qualitative inspection of predicted depth maps during 3-view training shows how different densities are enforced. Mip-NeRF assigns irregular densities that purely minimize reconstruction loss from 3 views, and therefore don't have to be constrained to surfaces. Rendered Loss assigns semi-transparent and white densities everywhere to ensure the weighted training rays sum to their respective depth values. URF overfits to training views and enforces δ-shaped densities, while ours models very fine geometry (see holes in the Lego model) with minimal supervision.