Nerfbusters: Removing Ghostly Artifacts from Casually Captured NeRFs

Frederik Warburg*[1], Ethan Weber*[1], Matthew Tancik[1], Aleksander Hołyński[1,2], Angjoo Kanazawa[1]

UC Berkeley[1], Google Research[2]

BERKELEY ARTIFICIAL INTELLIGENCE RESEARCH

Research at Google

SCAN ME

TLDR: We propose an evaluation procedure for in-the-wild NeRFs, and we present a method that uses a 3D diffusion prior to clean NeRFs

Motivation

- Casually captured NeRFs have artifacts such as floaters of flawed geometry
- Current NeRF evaluation fails to address the in-the-wild setting with novel views far from train

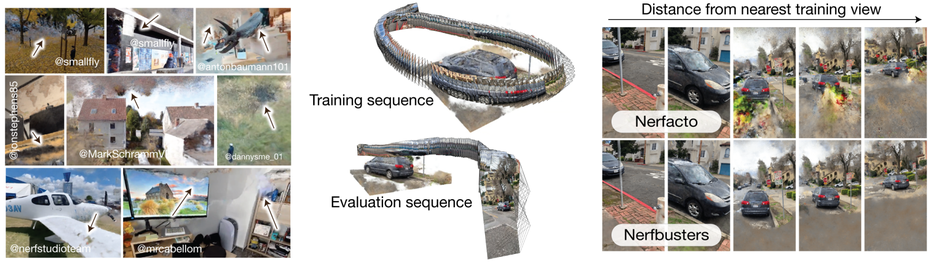

Main Contributions

- We propose an evaluation setting with two separate videos for training and evaluation
- We create and release the Nerfbusters Dataset which consists of 12 captures, 2 videos each
- We train a local 3D diffusion prior and apply a Density Score Distillation Sampling (DSDS) loss

Evaluation Procedure

- We train a NeRF on both train and evaluation videos to create pseudo GT
- We evaluate on visible regions and report coverage

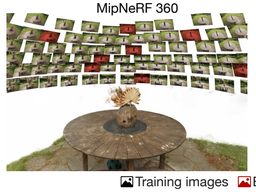

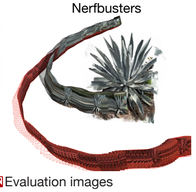

Training images Evaluation images

	Translation to NN	Rotation to NN
Nerfbusters	0.62	28.51
Synthetic [20]	0.29	18.12
LLFF [19]	0.24	2.07
Phototourism [11]	0.01	9.47
MipNeRF-360 [2]	0.07	10.49

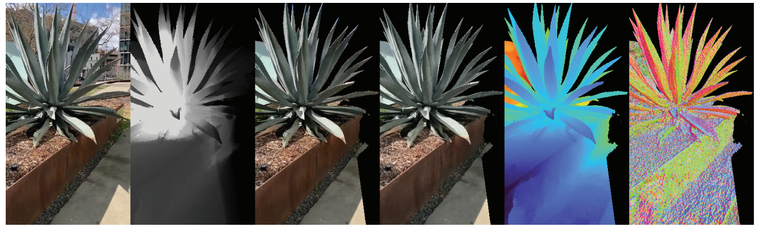

3D Diffusion Prior

- We train a local 3D diffusion prior on small cubes on ShapeNet
- We sample NeRF densities and apply the prior with a Density Score Distillation Sampling (DSDS) loss
- Our method is applied as a post-processing step
- We importance sample cube locations based on NeRF rendering densities

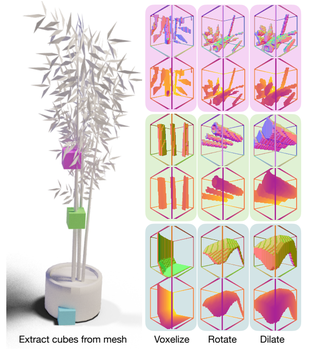

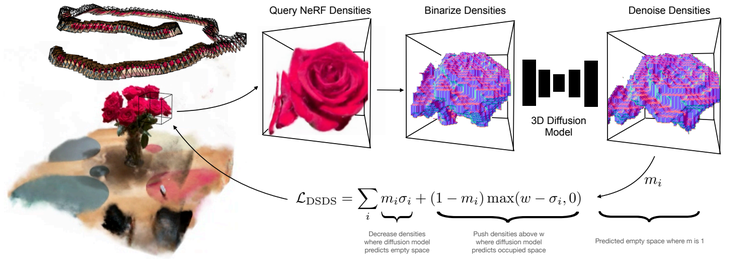

Visibility Loss

- We penalize densities if they are not visible by any training views
- The visibility loss is easy to implement with camera frustum checks

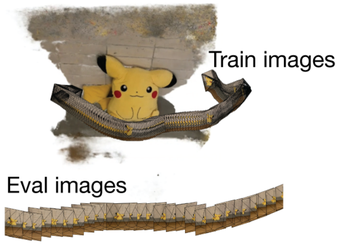

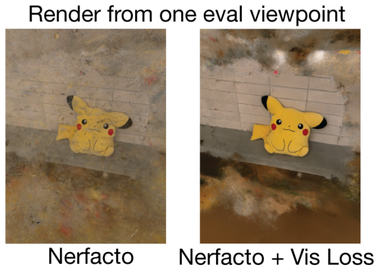

Results

- Our learned 3D prior favorably reduces floaters compared to hand-crafted methods such as sparsity loss, total-variation loss, or RegNeRF

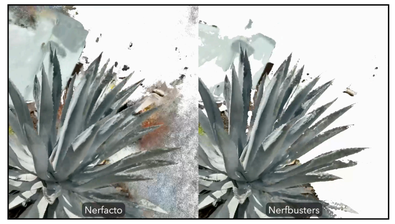

	PSNR ↑	SSIM ↑	LPIPS ↓	Depth ↓	Disp. ↓	Mean ° ↓	Median ° ↓	% 30° ↑	Coverage ↑
Nerfacto Pseudo GT	25.98	0.8591	0.1019	0.0	0.0	0.0	0.0	1.0	0.893
Nerfacto	17.00	0.5267	0.3800	126.277	1.510	60.63	54.638	0.254	**0.896**
+ Visibility Loss	17.81	0.5538	0.3432	100.057	1.041	57.73	51.335	0.280	0.854
+ Vis + Sparsity [42]	17.81	0.5536	0.3445	92.168	1.145	57.77	51.399	0.280	0.854
+ Vis + TV [7]	17.84	0.5617	0.3409	74.015	0.382	61.93	56.164	0.242	0.843
+ Vis + RegNeRF [25]	17.49	0.5396	0.3585	182.447	1.200	59.39	53.267	0.268	0.858
+ Vis + DSDS (Ours)	**17.99**	**0.6060**	**0.2496**	**54.453**	**0.114**	**54.77**	**47.981**	**0.295**	0.630

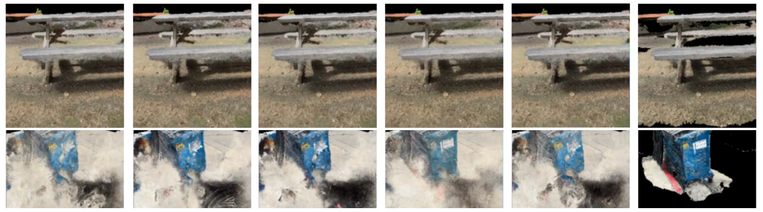

Limitations

- Our method deletes geometry but cannot inpaint color
- Our method does not scale to large scenes due to spatial aliasing