RegNeRF: Regularizing Neural Radiance Fields for View Synthesis from Sparse Inputs

Michael Niemeyer Jonathan T. Barron Ben Mildenhall
Mehdi S. M. Sajjadi Andreas Geiger Noha Radwan

Google

MAX PLANCK INSTITUTE FOR INTELLIGENT SYSTEMS

EBERHARD KARLS UNIVERSITÄT TÜBINGEN

Motivation

Neural Radiance Fields (NeRFs) achieve impressive view synthesis results when **many input images** are given:

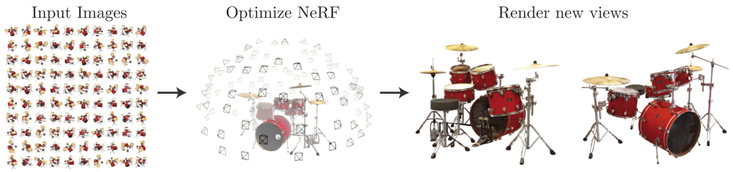

However, in real world applications, the input is often **sparse** and obtaining a dense coverage of the scene can be costly.

The Challenge

We observe that for **sparse inputs**, the view synthesis **results** for NeRF are **degenerate**:

Sparse Set of 3 Input Views

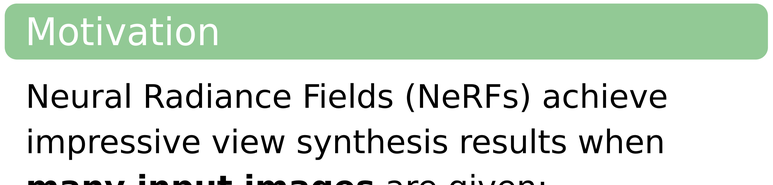

Novel Views Generated by NeRF

We propose RegNeRF which enables **realistic view synthesis** for sparse inputs:

Same Novel Views Generated by RegNeRF

How do we achieve this?

Our Method

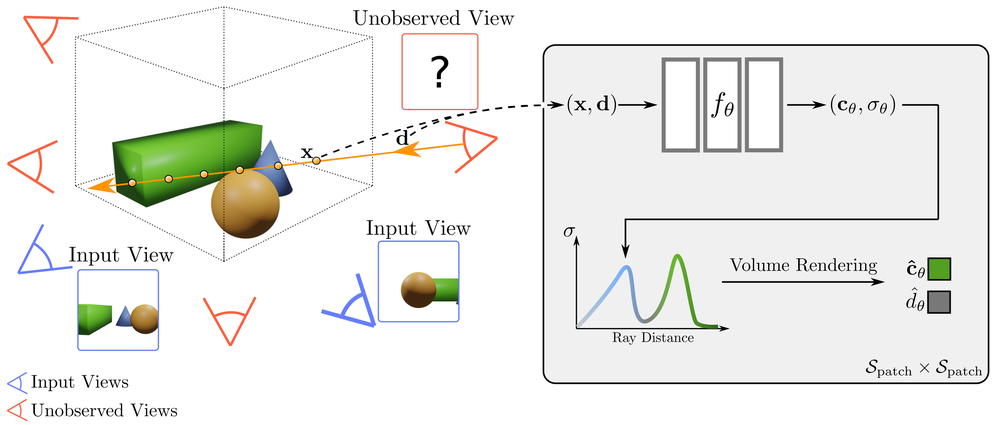

Next to the reconstruction loss on the input views, we perform the following for unobserved views:

1. Appearance Regularization:

$$\hat{P} = \{\hat{\mathbf{c}}_\theta(\mathbf{r}_{ij})\}_{ij} \longrightarrow \mathcal{L}_{\text{NLL}}(\theta) = -\log p_Z\left(\phi\left(\hat{P}\right)\right)$$

2. Geometry Regularization:

$$\left\{\hat{d}_\theta(\mathbf{r}_{ij})\right\}_{ij} \longrightarrow \mathcal{L}_{\text{DS}}(\theta) = \sum_{i,j}\left(\hat{d}_\theta(\mathbf{r}_{ij}) - \hat{d}_\theta(\mathbf{r}_{i+1j})\right)^2 + \left(\hat{d}_\theta(\mathbf{r}_{ij}) - \hat{d}_\theta(\mathbf{r}_{ij+1})\right)^2$$

For all views, we perform:

3. Sample Space Annealing:

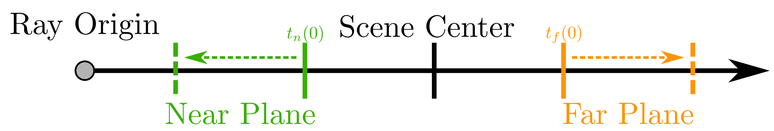

$$t_n(i) = t_m + (t_n - t_m)\eta(i)$$
$$t_f(i) = t_m + (t_f - t_m)\eta(i)$$
$$\eta(i) = \min\left(\max\left(\frac{i}{N_t}, p_s\right), 1\right)$$

Results

View Synthesis

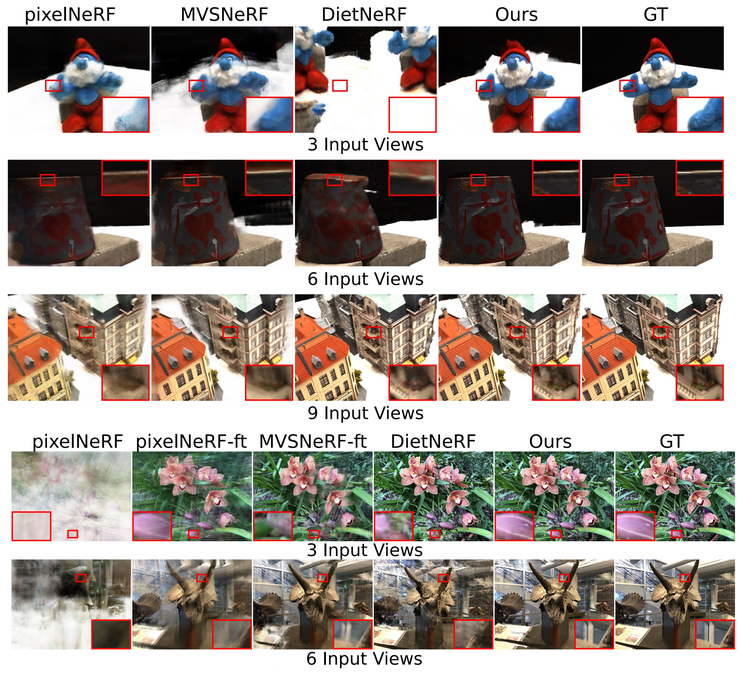

Geometry Comparison

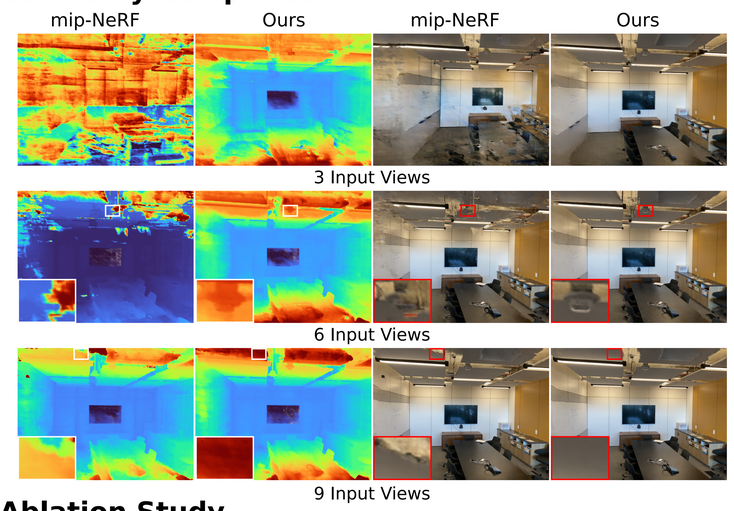

Ablation Study

michael.niemeyer@tue.mpg.de