A View-consistent Sampling Method for Regularized Training of Neural Radiance Fields

Overview

- Ray sampling is a critical component to training neural radiance fields for artifact-free 3D scene representation.

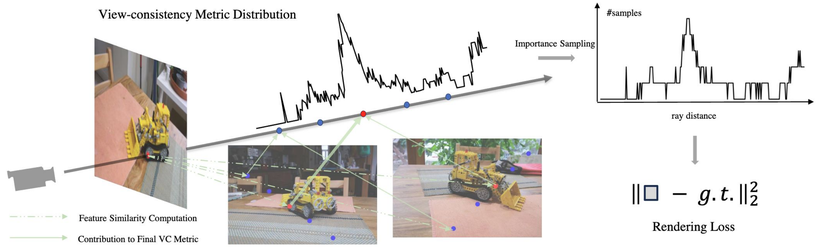

- ✔ We propose **view-consistent sampling**: Our central idea is to pre-compute a view-consistency distribution along rays and to perform importance sampling according to this distribution.

  As a result, the sampling will concentrate around surface points instead of random points in the capture volume.

- The question is: How can we compute a good view-consistency distribution in practice?

- ✔ For each sampled point, we will evaluate its view-consistent score using a tailored feature representation that retains geometric information.

Methodology

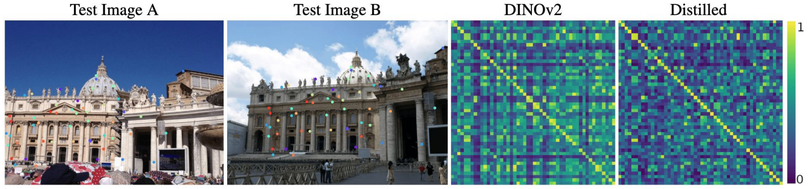

- ✔ We first **distill geometric information** from DINOv2: We train our distillation network by enforcing distilled features to be as close as possible at matching locations, and as far as possible at non-matching locations.

  We visualize the feature similarities across the two test views, the distilled features retains geometric information and thus create near-identity matrix.

- ✔ **View-consistency metric:** For each sampled point, we extract distilled features from its projection locations at all views, the metric is then simply computed by comparing the reference feature with projection features.

- ✔ **Depth-pushing loss:** We also employ a depth-pushing loss to favor distant sampled points to prevent background collapse.

Results

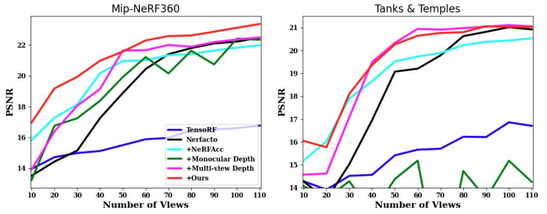

Find code and paper here!

Our sampling method has proven effective in NeRF training and outperforms regularizations using monocular and multi-view depth estimations.