3D Reconstruction Algorithm

Brandan Herron

Columbia University and Amazon Summer Undergraduate Research Experience Program, Columbia University, Amazon

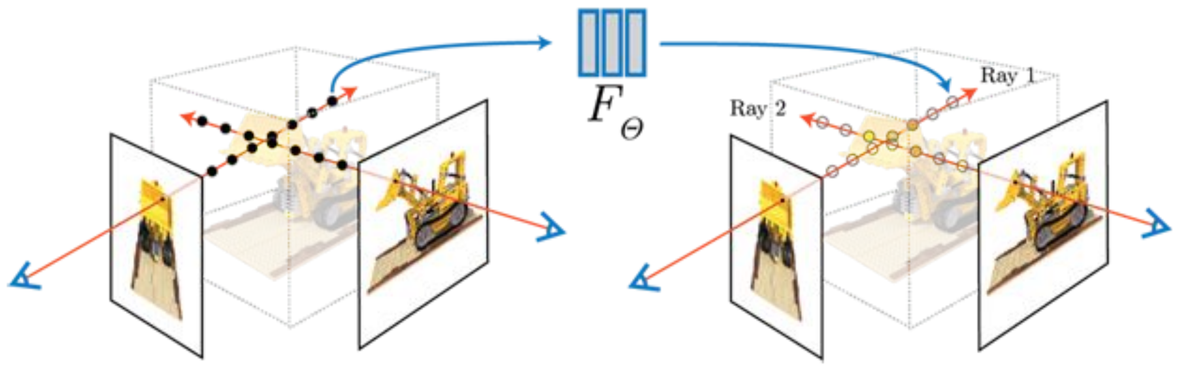

Ben et al. (2022)

Introduction

- Neural Radiance Fields (NeRF) is a fully-connected neural network that can generate novel views of complex 3D scenes, based on a set of 2D images.
- NeRFs contain a neural network which is an implicit function that exists everywhere as a continuous representation.
- Input data is a 3D location (x, y, z) and a 2D viewing direction (θ, Φ) and the output is an emitted color (r, g, b) and volume density (α).
- NeRF takes 12 hours to train compared to 10 minutes to train when using Local Light Field Fusion (LLFF), however, NeRF only requires 5MB, compared to 15GB for LLFF (3000x more than NeRF.

Methods

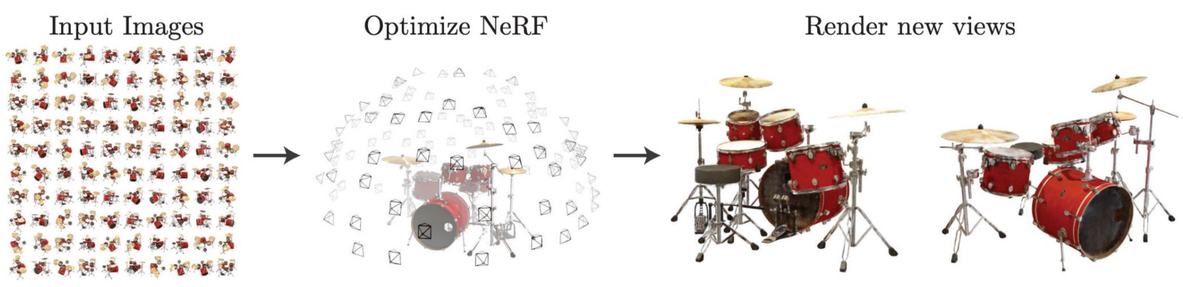

Ben et al. (2022)

- We sample the 5D coordinates (location and viewing direction) along camera rays.
- Feeding those locations into an MLP to produce a color and volume density.
- The volume rendering technique is used to composite these values into a 2D image.

Skills Gained

PyTorch

Google Developers

- Understanding neural radiance field scene representation, and volume rendering with radiance fields algorithms.
- Utilizing Pytorch for deep learning models.
- Using Colab to write and execute code.
- How to measure peak signal to noise ratio.
- Understand what Local Light Field Fusion (LLFF) is and why NeRF is more efficient than LLFF.

Results

- For our test, we used a set of 177 images.
- Rather than use the cloud, we thought a linux computer would be efficient enough to run the experiment.
- Over a few hours, we can see “I heart NY” and detail on the wooden patterns on the table starts to appear.

Reference

Ben M., Pratul S., Matthew T., Jonathan B., Ravi R., Ren N. (2022). NeRF: Representing Scenes as Neural Radiance Fields for View Synthesis. *https://arxiv.org/abs/2003.08934*