Image-Based Rendering

CSE 274, Lecture 10: Beyond NeRFs

Ravi Ramamoorthi

http://www.cs.ucsd.edu/~ravir

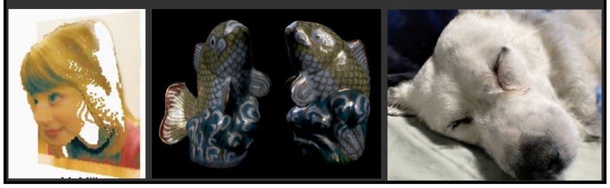

1

To Do

2

Motivation

- Neural Radiance Fields key for view synthesis
- Tons of follow-on work, best papers at almost every subsequent vision/graphics/learning conference
- We talk at a high level about a few key developments
  - Antialiasing (MIPNerf etc.)
  - Feature Grids: Triplanes (EG3D), TensoRF
  - InstantNGP
  - Gaussian Splatting
  - Single Image methods and GenAI
  - Stochastic Geometry Fields
- Only small subset, thousands of papers each year

3

Neural Volumes

- Neural volumes, differentiable ray marching
- Volume explicit voxel grid, using CNN
  - Considered MLPs like NeRF but w/o positional encoding

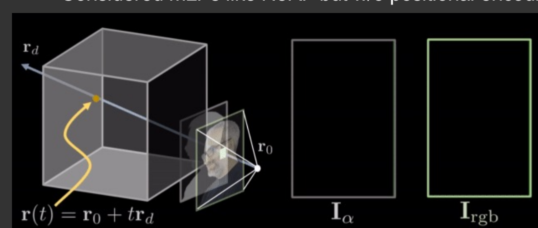

Lombardi et al. SIGGRAPH 2019

Seminal paper volume rendering, points to NeRFs, doesn't get due recognition always

4

Neural Volumes

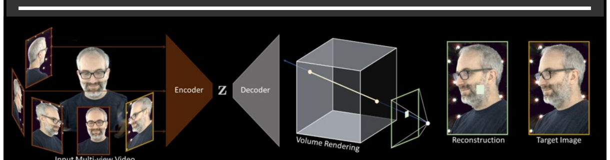

- Many Key Ideas
  - *Volumetric not surface representation (as regular grid)*
  - Neural network representation of volume
  - Directly optimize loss with respect to reconstruction
  - Differentiable Ray Marching
- Subsequent work retains notions above (note: gaussian splatting does not use a neural network)

5

Neural Radiance Fields (NeRF)

Mildenhall, Srinivasan, Tancik, Barron, Ramamoorthi, Ng 20

6

Toy problem: storing 2D image data

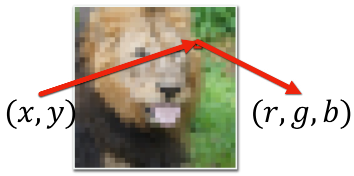

Usually we store an image as a 2D grid of RGB color values

7

Toy problem: storing 2D image data

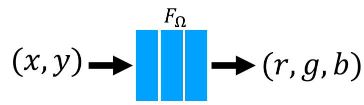

What if we train a simple fully-connected network (MLP) to do this instead?

8

Naive approach fails!

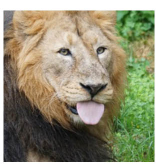

Ground truth image

Neural network output

9

Problem:

"Standard" coordinate-based MLPs cannot represent high frequency functions

10

Solution:

Pass input coordinates through a high frequency mapping first

11

Example mapping: "positional encoding"

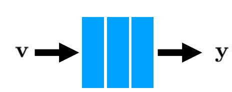

$$\begin{pmatrix} \sin(\mathbf{v}), \cos(\mathbf{v}) \\ \sin(2\mathbf{v}), \cos(2\mathbf{v}) \\ \sin(4\mathbf{v}), \cos(4\mathbf{v}) \\ \dots \\ \sin(2^{L-1}\mathbf{v}), \cos(2^{L-1}\mathbf{v}) \end{pmatrix}$$

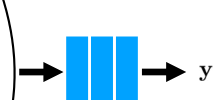

12

Problem solved… but why?

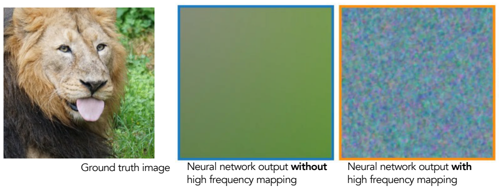

13

Fourier Features Let Networks Learn High Frequency Functions in Low Dimensional Domains

Tancik et al., NeurIPS 2020

14

Use neural network to replace large N-d array

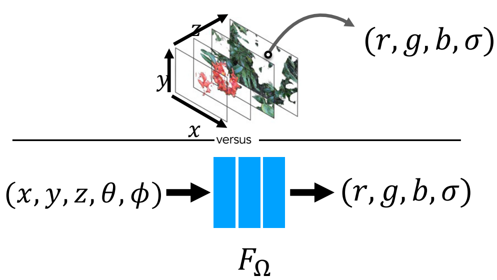

15

Mip-NeRF

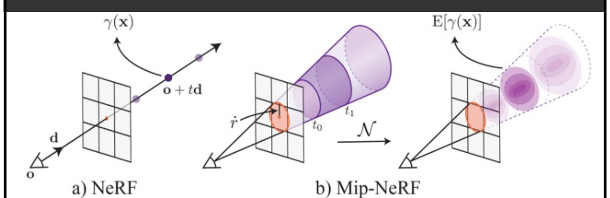

- https://www.youtube.com/watch?v=EpH175PY1A0&t=11s
- Subsequent work: MipNeRF360 (CVPR 22);
  Tri-MipRF, ZipNeRF (ICCV 23); Rip-NeRF (SIG 24)

Barron et al. ICCV 2021

16

Feature Fields

- Neural Volumes and NeRFs are two extremes
  - Purely explicit voxel grid vs compact neural (large) MLP
- Hybrid representations: Feature Grids
  - Coarse voxel grid of abstract features
  - Smaller (often just 2-layer) MLP to decode features to standard volume density and color as in NeRF
  - No positional encoding needed any more, only features
- Memory-Time tradeoff: more memory for faster computation time and training
- Explosion in types of feature grids (and some revisiting PRT and early light transport)
- Gaussian splatting uses gaussians, no MLP at all

17

TriPlanes EG3D

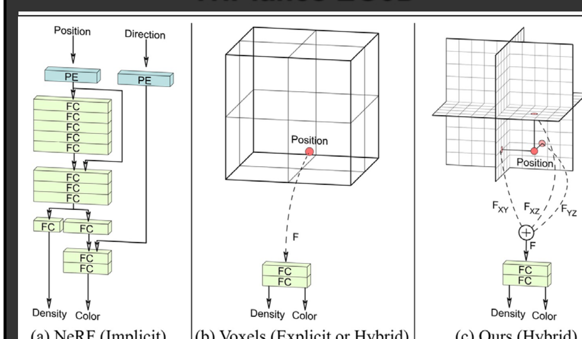

Chan et al. CVPR 22. https://nvlabs.github.io/eg3d/

18

Tensorial Radiance Fields

- Tensor factorization of feature grid into lower-dimensional 2D components (Triplanes is also a 2D factorization)

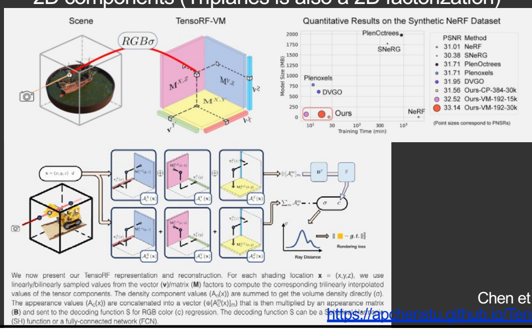

Chen et al. 22
https://apchenstu.github.io/TensoRF/

19

2022: NVIDIA Instant NGP – Real Time NeRFs

TIME Magazine Named NVIDIA Instant NeRF a Best Invention of 2022

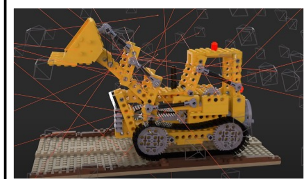

The original neural radiance fields paper from 2020 took many hours to train a model and rendered at 0.03fps

With Instant NGP, NeRFs train in seconds, and render at real time rates (30 fps or faster)

Thomas Muller, Alex Evans, Christopher Schied, and Alex Keller 22

20

Multiresolution Hash Encoding

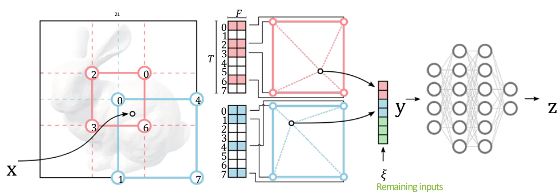

21

Reducing compute-per-sample: learned hash grids (Instant NGP)

Müller, et al 2022, *Instant Neural Graphics Primitives with a Multiresolution Hash Encoding*

22

3D Gaussian Splatting for Radiance Fields

Kerbl et al. 23 https://www.youtube.com/watch?v=T_kXY43VZnk

23

NeRFs for Digital Twins

24

NeRFs for Landscapes

Instant NeRF (@jonstephens85 on Twitter); slide courtesy Thomas Muller

25

NeRFs for Furry Volumes

Instant NeRF (@vibrantnebula on Twitter); slide courtesy Thomas Muller

26

NeRFs applied to almost every vision/graphics++ problem

Over 1400 research papers related to NeRFs in 2022 alone

Decomposing NeRF for Editing via Feature Field Distillation
Kobayashi et al NeurIPS 2022

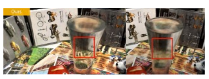

Eikonal Fields for Refractive Novel-View Synthesis
Bemana et al SIGGRAPH 2022

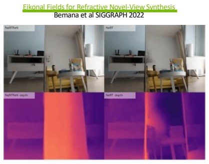

NeRFReN Neural Radiance Fields with Reflections
Guo et al CVPR 2022

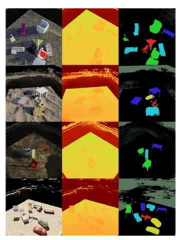

NeSF: Neural Semantic Fields for Generalizable Semantic Segmentation of 3D Scenes
Vora, Radwan, et al TMLR 2022

27 NVIDIA

27

NeRFs + Language Models

- Text-based editing with large language models (Haque et al.)

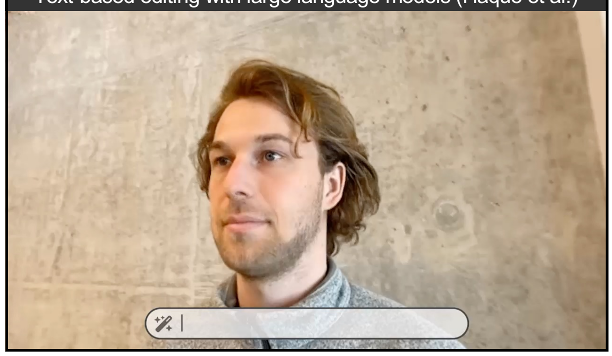

28

NeRFs in Production (Google, Luma)

Google Street View

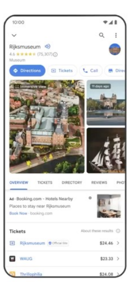

Luma AI iphone app

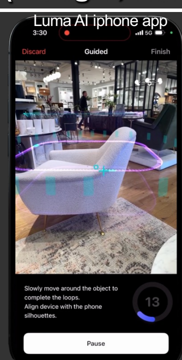

29

Real-Time Radiance Fields (1 image)

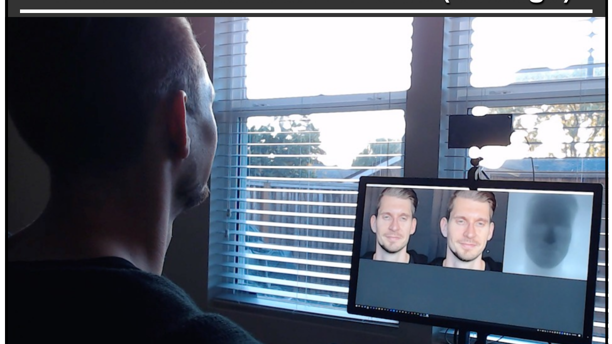

Trevithick et al. 23

30

Lifting Text-Based Avatars to 3D

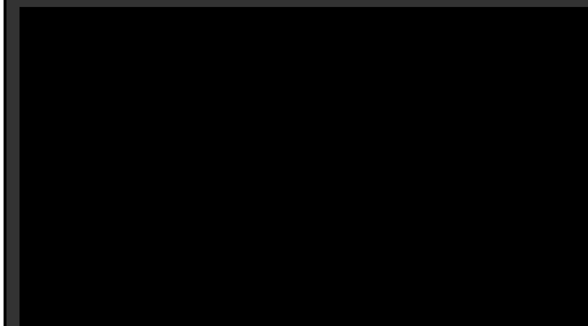

31

3D Videoconferencing

SIGGRAPH 23 Emerging Technologies

32

Modern Text-Conditioned Diffusion (One-2-3-45++)

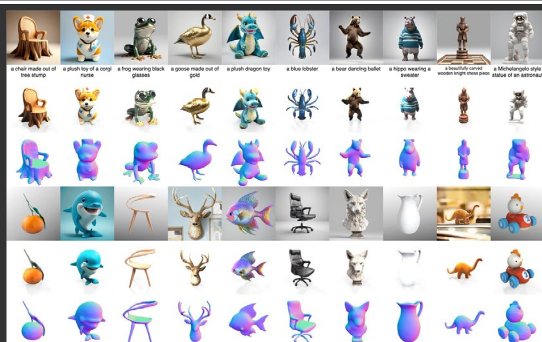

Liu et al. 24

33

CAT3D: Create Anything in 3D

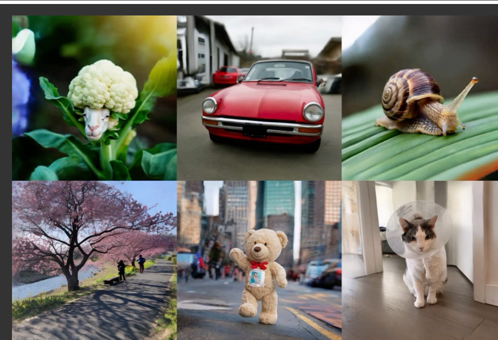

Gao et al. 24. See also https://www.worldlabs.ai/blog

34

Stochastic Geometry: Objects as Volumes

- Volumetric rendering is great, but what if you want an actual 3D representation?
- Possible with stochastic view of geometry, unifies objects with hard surfaces and volumes (Miller et al. CVPR 24)
- https://www.youtube.com/watch?v=ZKUcowfFDPk

35