Neural Radiance Fields:
Fundamentals to Applications

CS380 Talk 1.
Youngju Na
M.S. Student @ SGVR Lab

KAIST SGVR Lab

yjna2907@kaist.ac.kr

Background: Novel View Synthesis

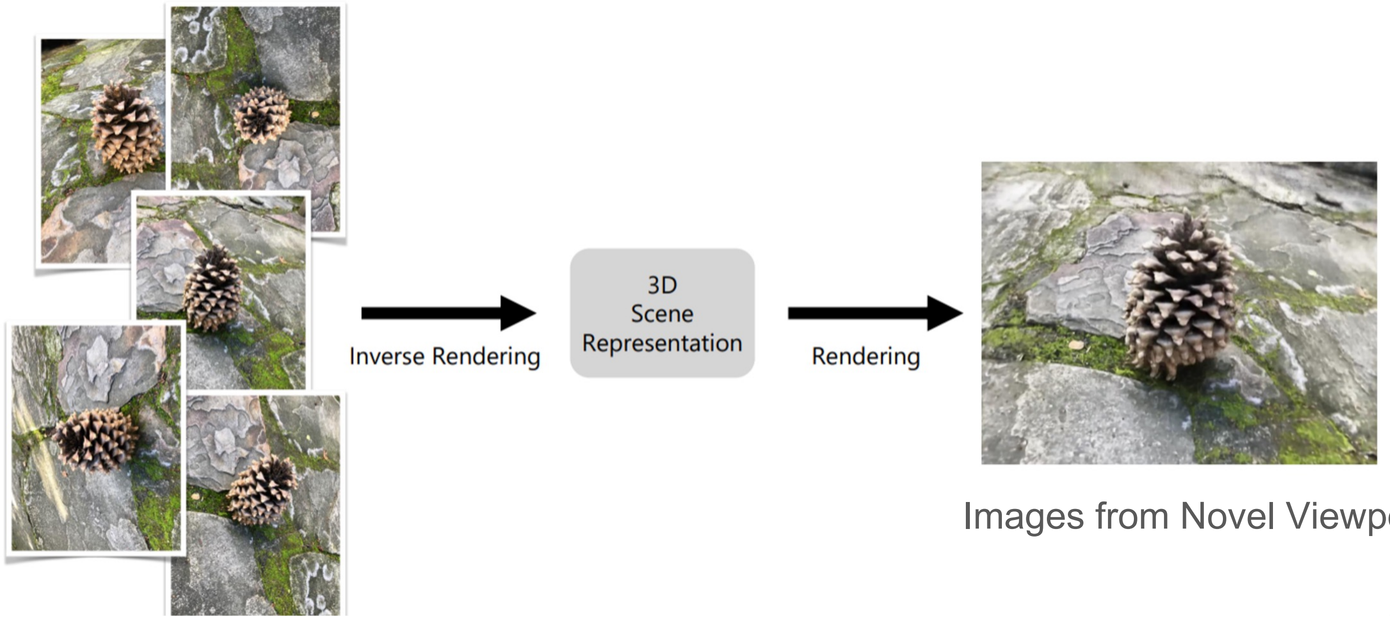

Images from Novel Viewpoints

Images from multiple camera viewpoints

Neural Radiance Fields ECCV 2020 Oral - Best Paper Honorable Mention

Input: images from various camera viewpoints

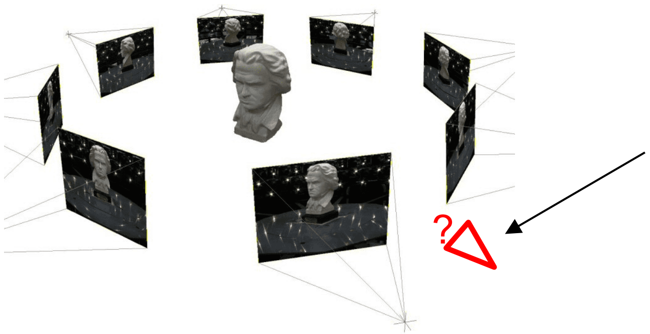

Output: images from novel camera viewpoints

Source: https://theaisummer.com/nerf/

Examples (synthesized from novel views)

Videos: https://www.youtube.com/watch?v=JuH79E8rdKc&t=191s

https://www.matthewtancik.com/nerf

Implicit Representation

$f(\cdot)$ is a parameterized 2D/3D scalar field

x: coordinate

$f(x) = \|x^2\| - 1$

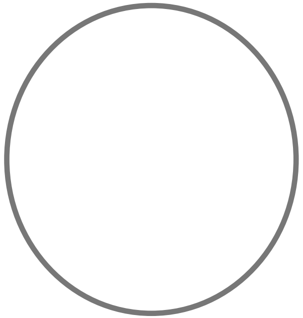

x

$f(x) = ?$

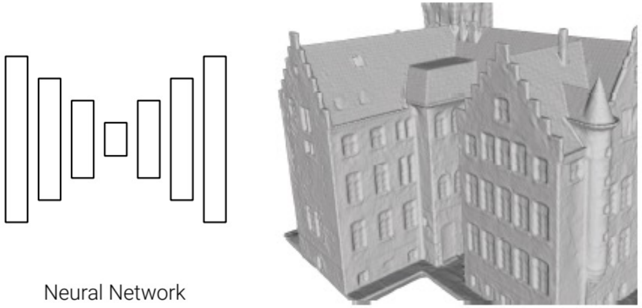

Represent 3D Scene as Continuous functions

Signed Distance Function (SDF) or Occupancy Fields

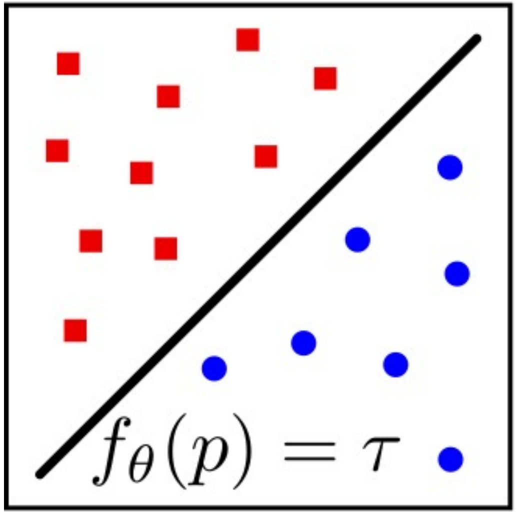

Occupancy networks, Mescheder et al.

NeRF 3D Representations

Neural Network as a continuous shape representaiton.

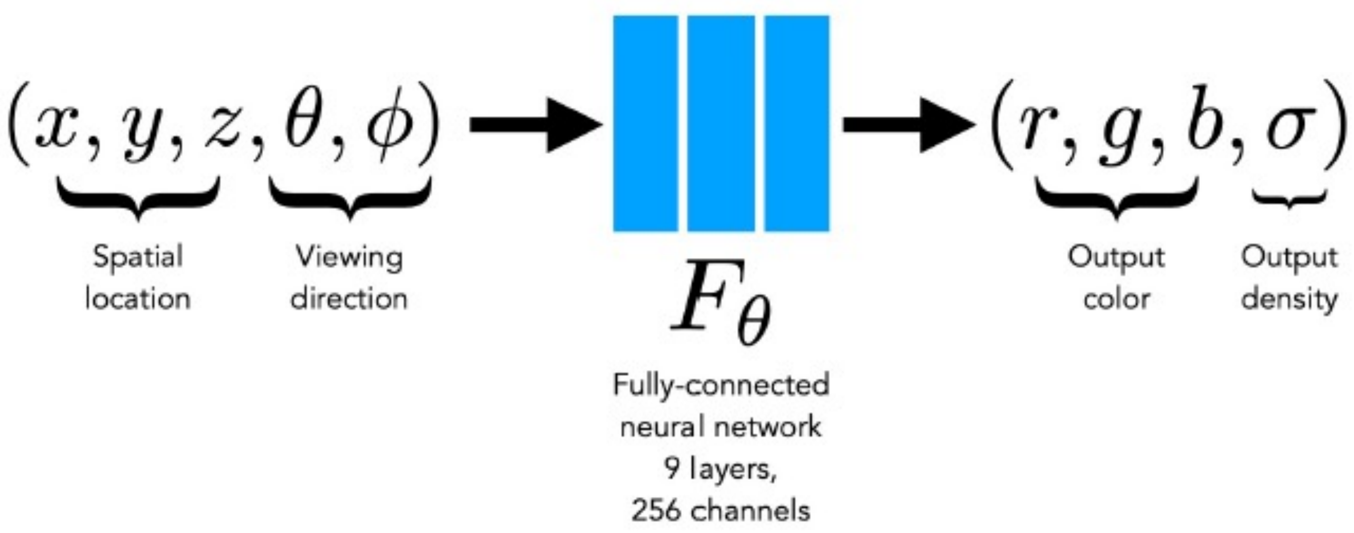

How do we learn 3D representations from 2D images?

Method Overview

Cast Rays => Estimate 3D Representations => **Volume Rendering** => 2D Photometric Loss

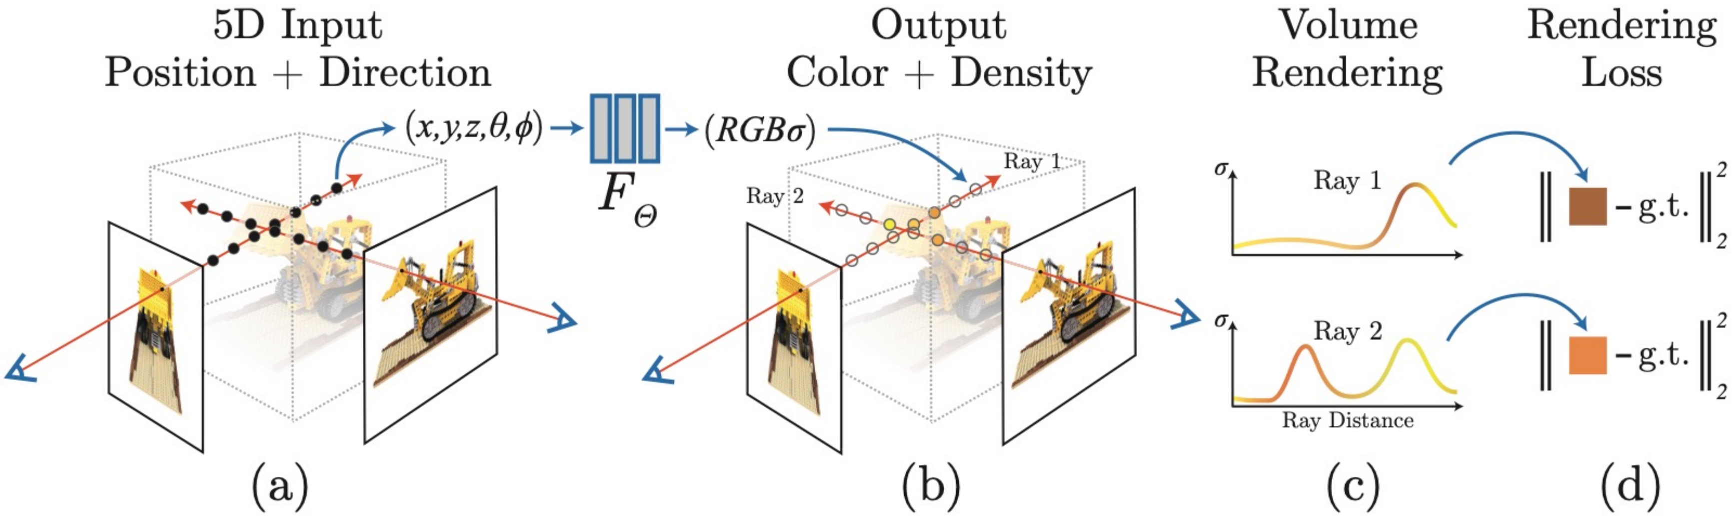

https://www.matthewtancik.com/nerf

Neural Volumetric Rendering

Neural Volumetric **Rendering**

computing color along rays
through 3D space

What color is this pixel?

https://sites.google.com/berkeley.edu/nerf-tutorial/home

Cameras and rays

- We need the mathematical mapping from (*camera*, *pixel*) → *ray*
- Then can abstract underlying problem as learning the function *ray* → *color* (the "plenoptic function")

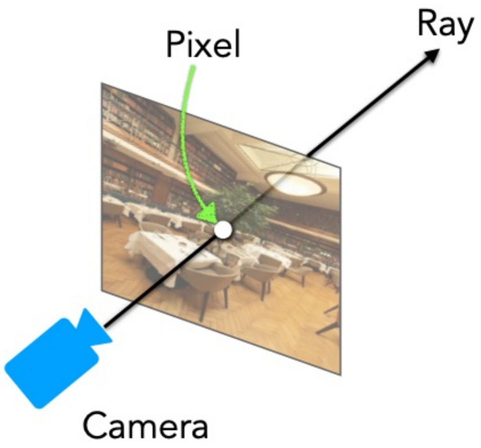

https://sites.google.com/berkeley.edu/nerf-tutorial/home

Coordinate frames + Transforms: world-to-camera

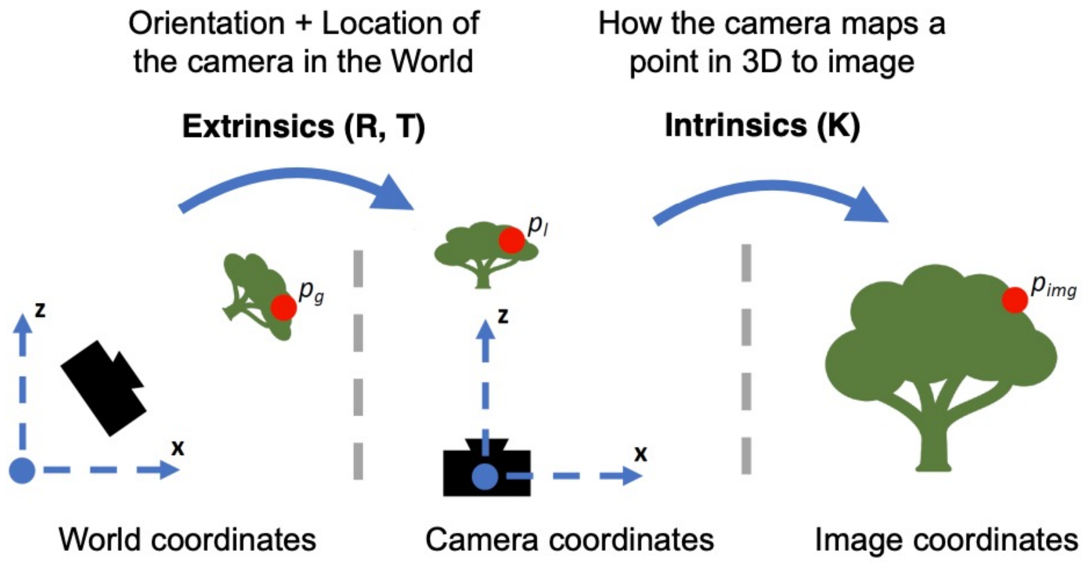

Figure credit: Peter Hedman

https://sites.google.com/berkeley.edu/nerf-tutorial/home

Calculating points along a ray

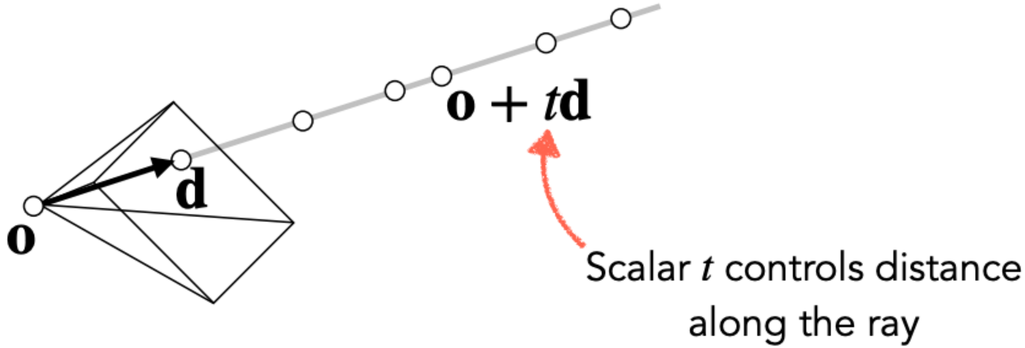

https://sites.google.com/berkeley.edu/nerf-tutorial/home

Neural **Volumetric** Rendering

continuous, differentiable
rendering model without
concrete ray/surface intersections

KAIST

https://sites.google.com/berkeley.edu/nerf-tutorial/home

Surface vs. volume rendering

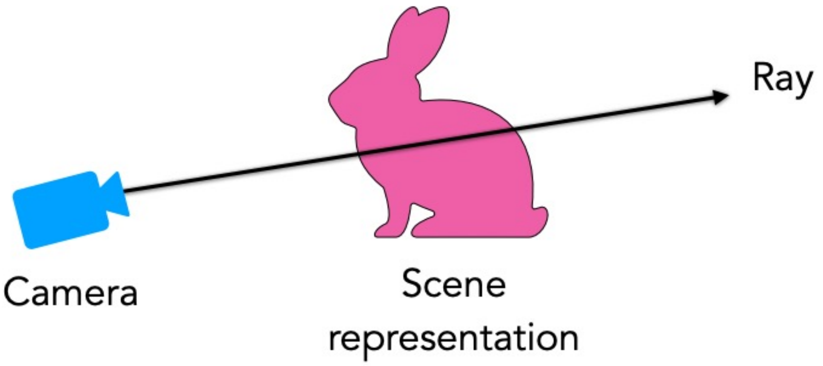

Want to know how ray interacts with scene

KAIST

https://sites.google.com/berkeley.edu/nerf-tutorial/home

Surface vs. volume rendering

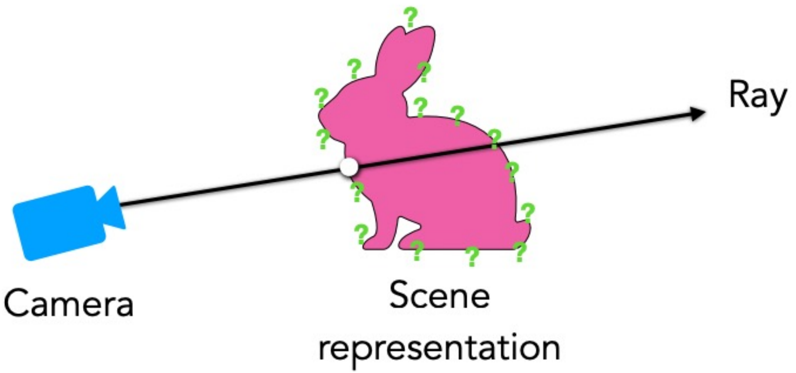

Surface rendering — loop over geometry, check for ray hits

https://sites.google.com/berkeley.edu/nerf-tutorial/home

Surface vs. volume rendering

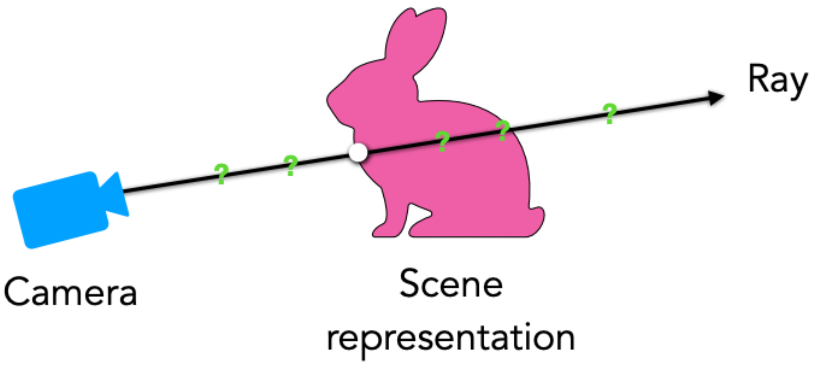

Volume rendering — loop over ray points, query geometry

Volumetric formulation for NeRF

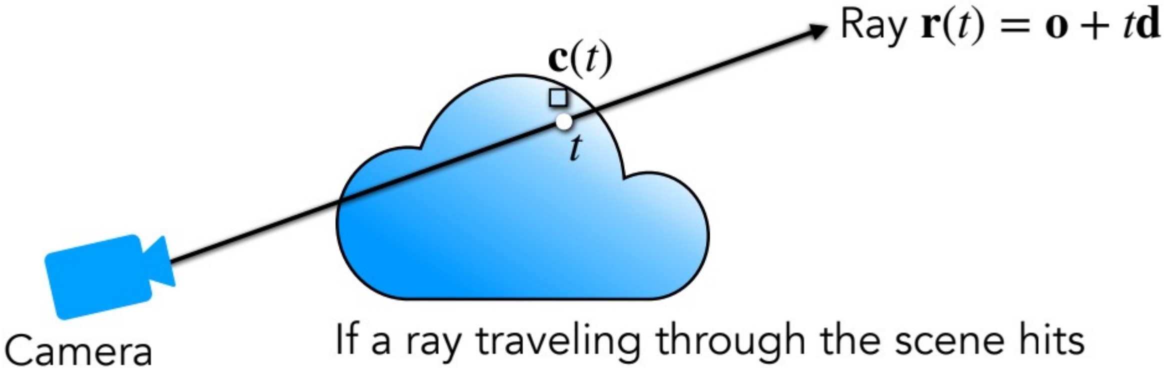

If a ray traveling through the scene hits a particle at distance t along the ray, we return its color $\mathbf{c}(t)$

https://sites.google.com/berkeley.edu/nerf-tutorial/home

What does it mean for a ray to "hit" the volume?

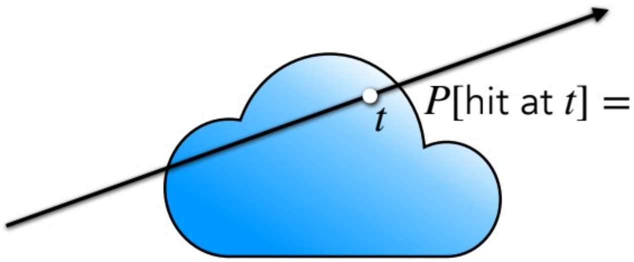

$P[\text{hit at } t] = \sigma(t)\, dt$

This notion is *probabilistic:* chance that ray hits a particle in a small interval around t is $\sigma(t)\, dt$.

σ is called the "volume density"

Probabilistic interpretation

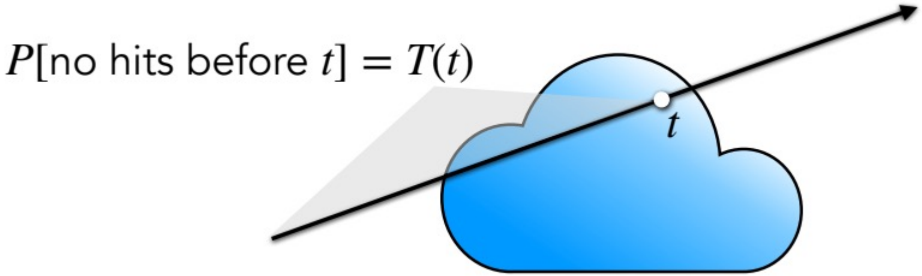

To determine if t is the *first* hit along the ray, need to know $T(t)$: the probability that the ray makes it through the volume up to t.
$T(t)$ is called "transmittance"

https://sites.google.com/berkeley.edu/nerf-tutorial/home

PDF for ray termination

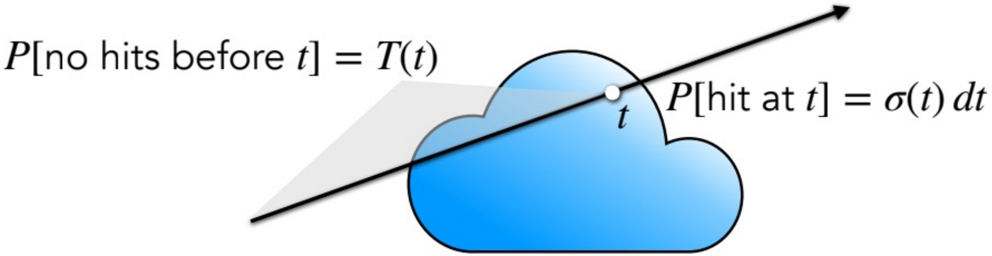

Finally, we can write the probability that a ray terminates at t as a function of only sigma

$$P[\text{first hit at } t] = P[\text{no hit before } t] \times P[\text{hit at } t]$$

$$= T(t)\sigma(t)dt$$

$$= \exp\left(-\int_{t_0}^{t} \sigma(s)\,ds\right)\sigma(t)\,dt$$

KAIST

https://sites.google.com/berkeley.edu/nerf-tutorial/home

Expected value of color along ray

This means the expected color returned by the ray will be

$$\int_{t_0}^{t_1} T(t)\sigma(t)\mathbf{c}(t)\,dt$$

Note the nested integral!

Approximating the nested integral

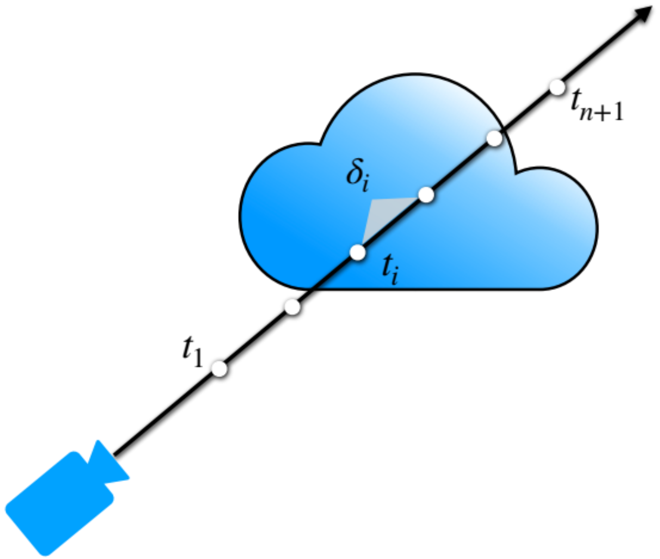

We use quadrature to approximate the nested integral,
splitting the ray up into n segments with endpoints $\{t_1, t_2, \ldots, t_{n+1}\}$
with lengths $\delta_i = t_{i+1} - t_i$

KAIST

https://sites.google.com/berkeley.edu/nerf-tutorial/home

Approximating the nested integral

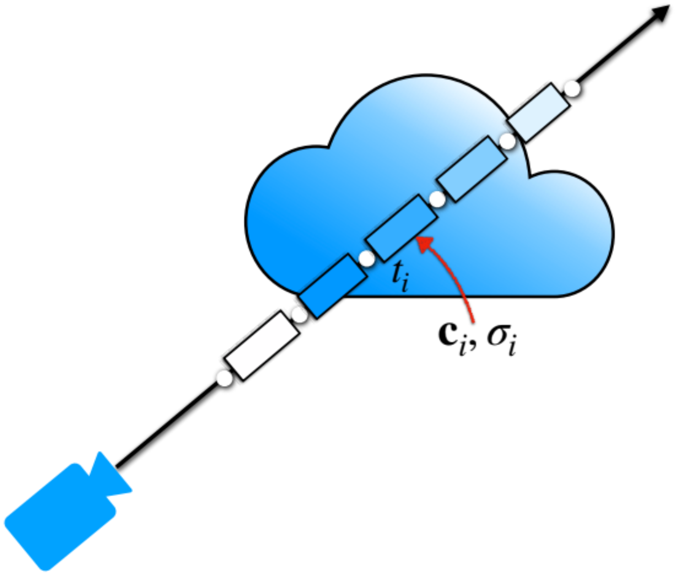

We assume volume density and color are roughly constant within each interval

https://sites.google.com/berkeley.edu/nerf-tutorial/home

Summary: volume rendering integral estimate

Rendering model for ray $\mathbf{r}(t) = \mathbf{o} + t\mathbf{d}$:

$$\mathbf{c} \approx \sum_{i=1}^{n} T_i \alpha_i \mathbf{c}_i$$

weights colors

How much light is blocked earlier along ray:

$$T_i = \prod_{j=1}^{i-1} (1 - \alpha_j)$$

How much light is contributed by ray segment i:

$$\alpha_i = 1 - \exp(-\sigma_i \delta_i)$$

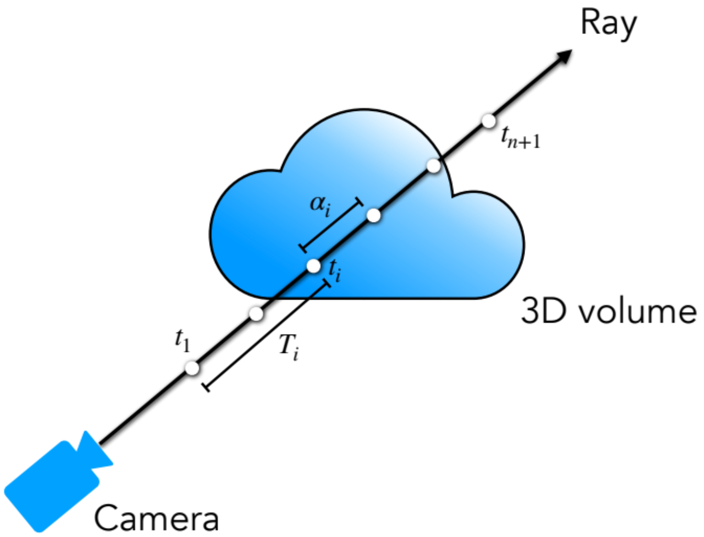

Detailed derivation

Volume rendering is trivially differentiable

Rendering model for ray $\mathbf{r}(t) = \mathbf{o} + t\mathbf{d}$:

$$\boxed{\mathbf{c} \approx \sum_{i=1}^{n} \underbrace{T_i \alpha_i}_{\text{weights}} \underbrace{\mathbf{c}_i}_{\text{colors}}}$$

differentiable w.r.t. $\mathbf{c}, \sigma$

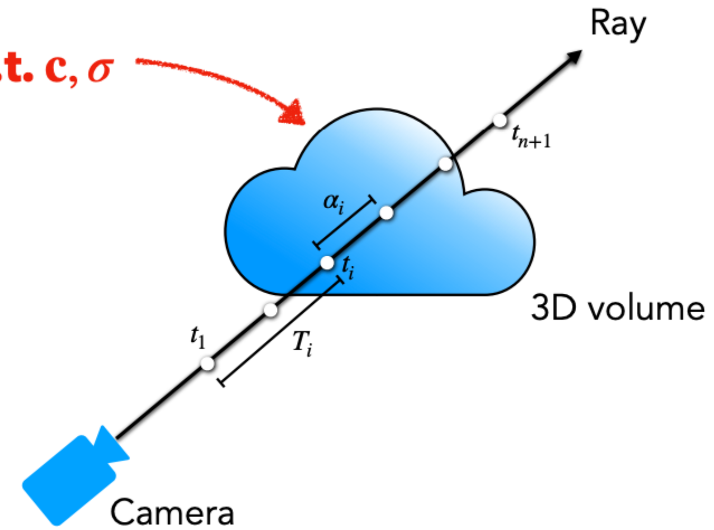

How much light is blocked earlier along ray:

$$T_i = \prod_{j=1}^{i-1} (1 - \alpha_j)$$

How much light is contributed by ray segment i:

$$\alpha_i = 1 - \exp(-\sigma_i \delta_i)$$

https://sites.google.com/berkeley.edu/nerf-tutorial/home

Video

Mildenhall et al., NeRF: Representing Scenes as Neural Radiance Fields for View Synthesis, ECCV 2020

Novel View Synthesis & View Dependency

https://www.matthewtancik.com/nerf

Resources

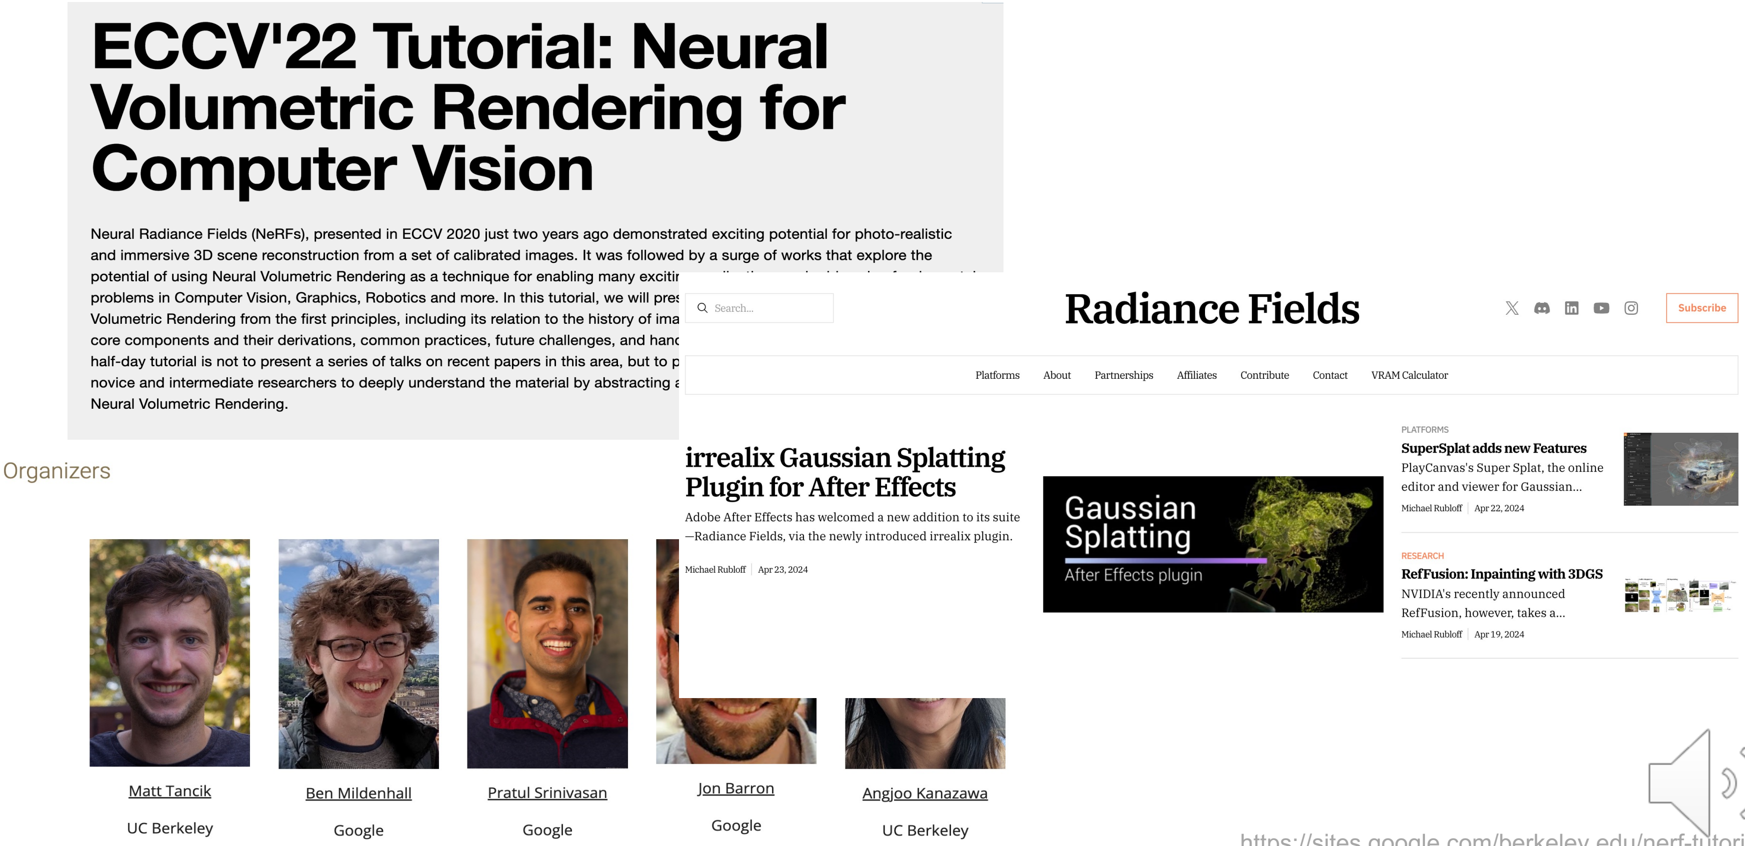

ECCV'22 Tutorial: Neural Volumetric Rendering for Computer Vision

Neural Radiance Fields (NeRFs), presented in ECCV 2020 just two years ago demonstrated exciting potential for photo-realistic and immersive 3D scene reconstruction from a set of calibrated images. It was followed by a surge of works that explore the potential of using Neural Volumetric Rendering as a technique for enabling many excitir
problems in Computer Vision, Graphics, Robotics and more. In this tutorial, we will pres
Volumetric Rendering from the first principles, including its relation to the history of ima
core components and their derivations, common practices, future challenges, and hanc
half-day tutorial is not to present a series of talks on recent papers in this area, but to p
novice and intermediate researchers to deeply understand the material by abstracting a
Neural Volumetric Rendering.

Organizers

Matt Tancik
UC Berkeley

Ben Mildenhall
Google

Pratul Srinivasan
Google

Jon Barron
Google

Angjoo Kanazawa
UC Berkeley

Search...

Radiance Fields

Subscribe

Platforms About Partnerships Affiliates Contribute Contact VRAM Calculator

irrealix Gaussian Splatting Plugin for After Effects

Adobe After Effects has welcomed a new addition to its suite —Radiance Fields, via the newly introduced irrealix plugin.

Michael Rubloff | Apr 23, 2024

Gaussian Splatting
After Effects plugin

PLATFORMS
SuperSplat adds new Features
PlayCanvas's Super Splat, the online editor and viewer for Gaussian...
Michael Rubloff | Apr 22, 2024

RESEARCH
RefFusion: Inpainting with 3DGS
NVIDIA's recently announced RefFusion, however, takes a...
Michael Rubloff | Apr 19, 2024

https://sites.google.com/berkeley.edu/nerf-tutorial/home
https://radiancefields.com/

Resources

docs passing | pypi package 1.0.3 | Core Tests. passing | License Apache 2.0

nerfstudio

A collaboration friendly studio for NeRFs

Documentation | Viewer | Colab

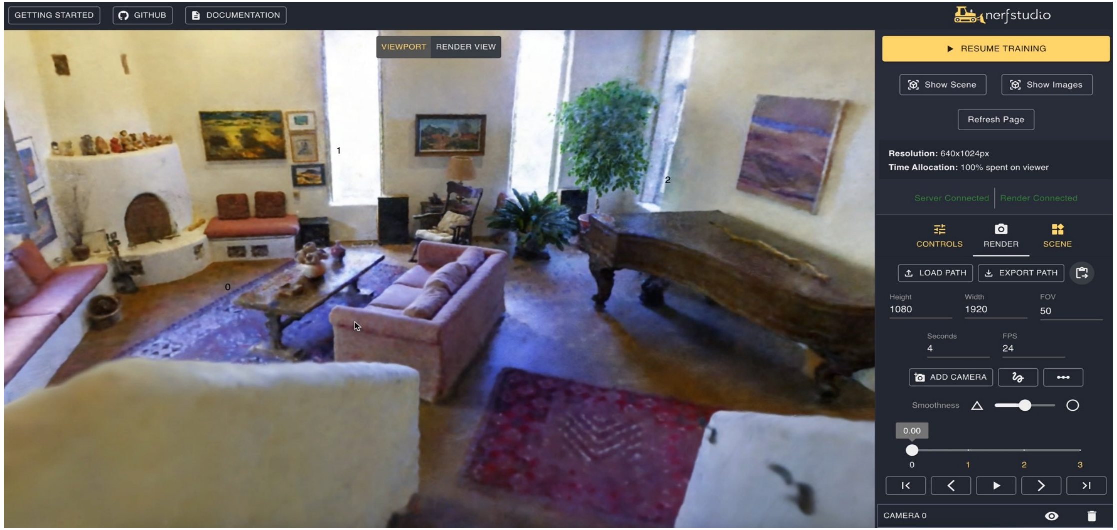

https://docs.nerf.studio/

Limitations / Applications

NeRF's Limitations and Applications

1. Slow rendering / optimization time => Fast Rendering

2. Assume static scene => Dynamic NeRF

3. Per-Scene Optimization => Generalizable Methods

4. Not a mesh => Surface Reconstruciton

List goes on and on…!
NeRF has been cited 6800+

KAIST

Fast Optimization / Rendering : Plenoxel [CVPR'22]

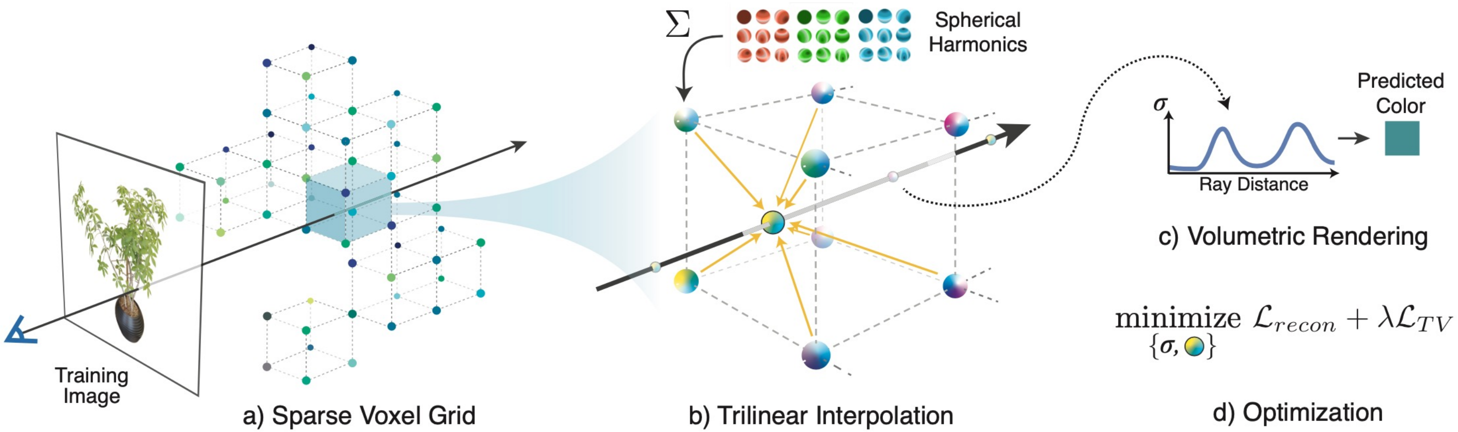

https://alexyu.net/plenoxels/

Fast Optimization / Rendering : Plenoxel [CVPR'22]

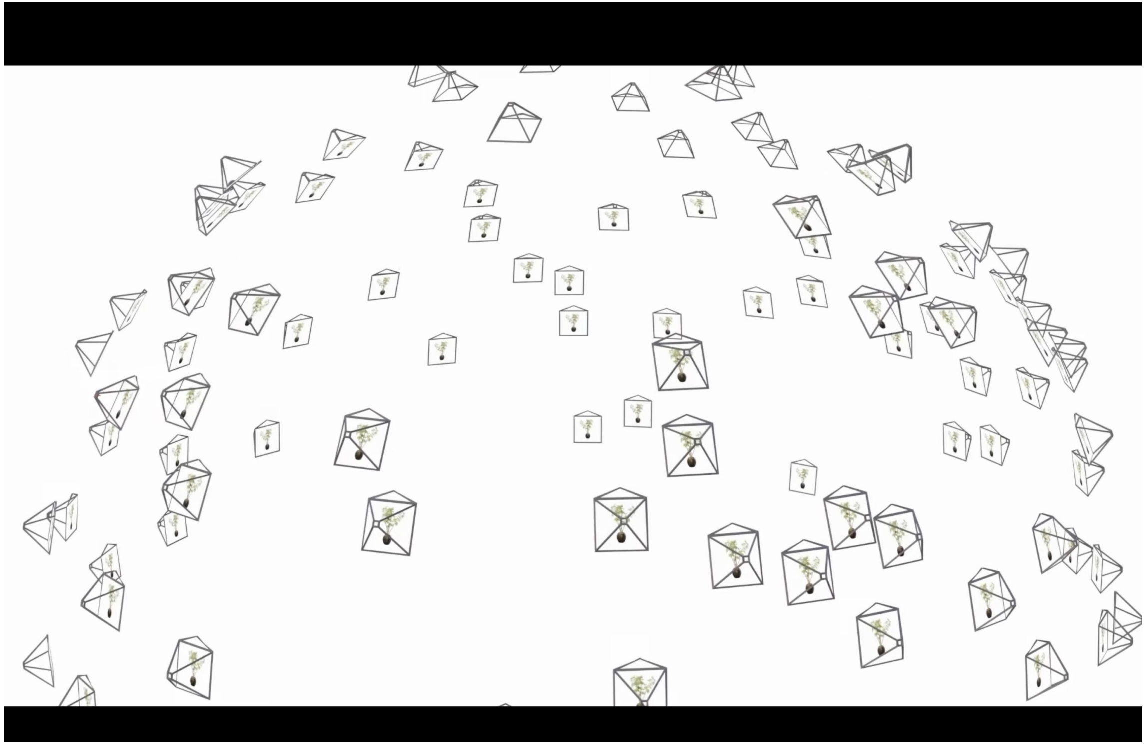

https://alexyu.net/plenoxels/

The world we capture is usually Dynamic / Deformable

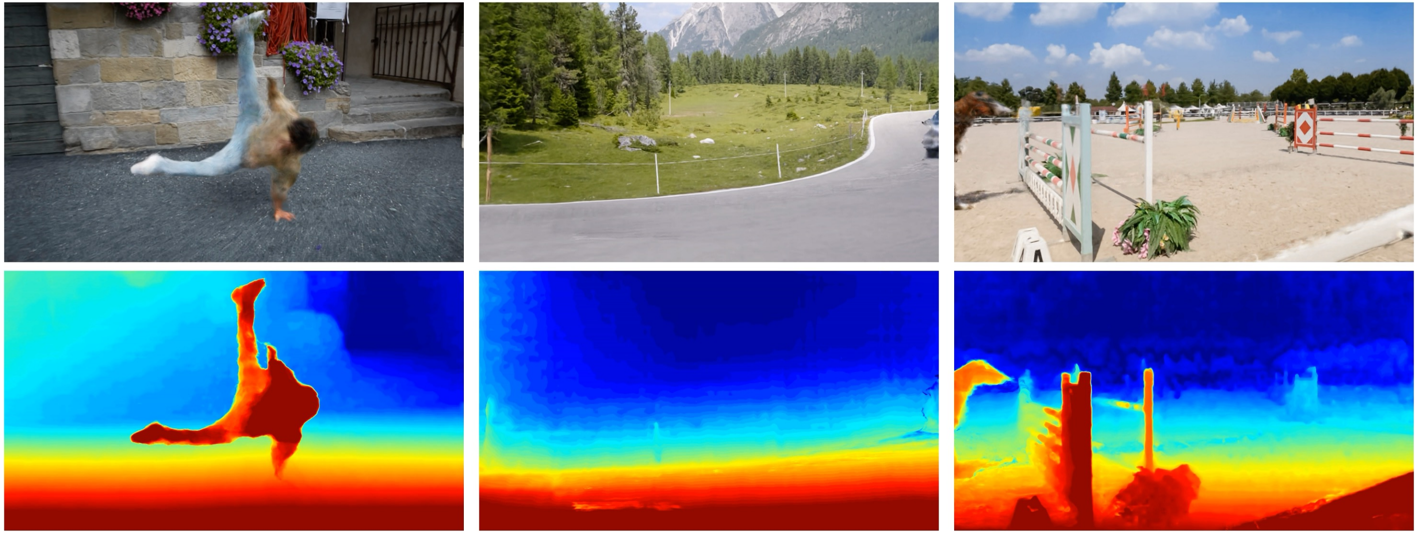

Simple baseline for adding time

$$(x, y, z, \theta, \phi, t) \rightarrow F_{\Omega} \rightarrow (r, g, b, \sigma)$$

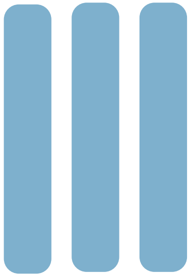

F_{Ω}

Hard without simultaneous multiple view!

Through a deformation network

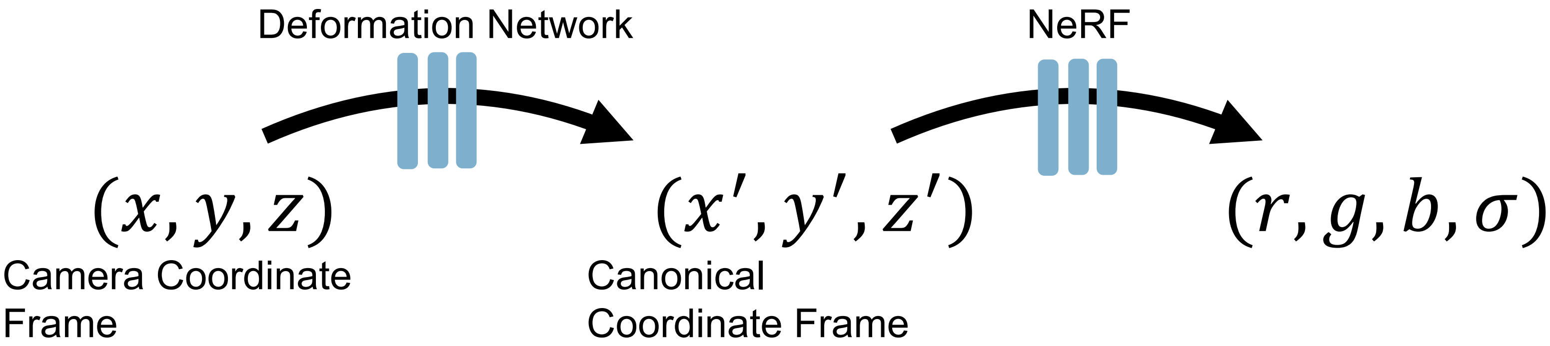

Still very under constrained

D-NeRF [Pumarola et al. CVPR 2021], Nerfies [Park et al. ICCV 2021], HyperNeRF, NR-NeRF, etc..

D-NeRF: Neural Radiance Fields for Dynamic Scenes

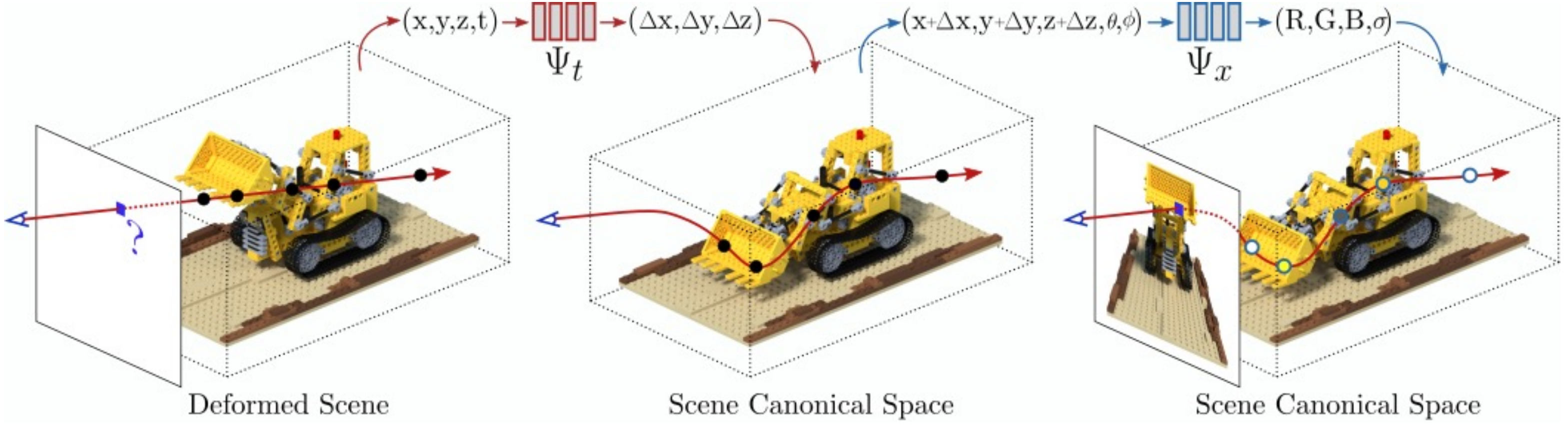

https://www.albertpumarola.com/research/D-NeRF

Dynamic NeRFs

RoDynRF, Liu et al. CVPR'23

DynIBaR, Li et al. CVPR'23

https://www.albertpumarola.com/research/D-NeRF

NeRF requires Per-Scene Optimization
Generalizable Methods with Prior Knowledge

NeRF requires Per-Scene Optimization with Dense Views

1. Scene-specific representation

2. Sparse input camera viewpionts

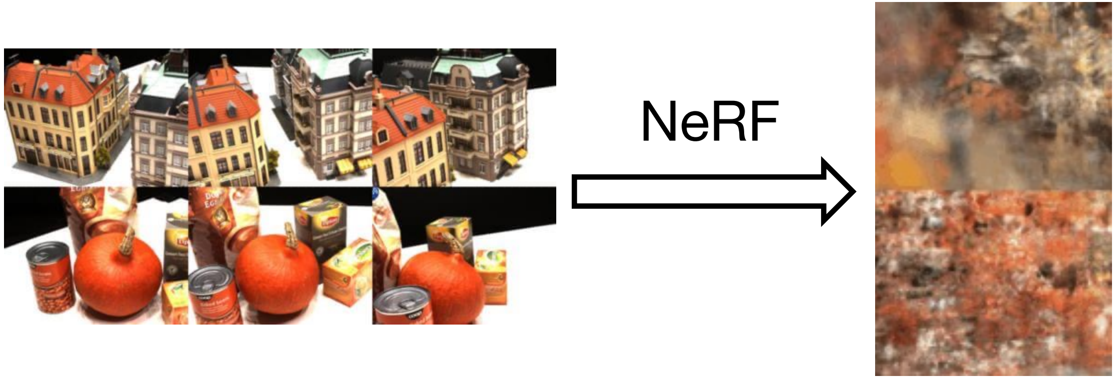

Not Generalizable

Cannot share representations across scenes or views

Few-Shot / One-Shot NeRF

- One-Shot NeRF (pixelNeRF [Yu et al. CVPR'21])

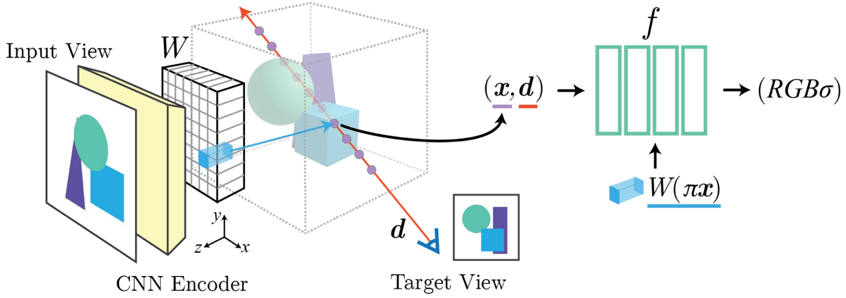

PixelNeRF NeRF

Few-shot (3~10 views): pixelNeRF, IBRNet [Wang et al. CVPR'21], MVSNeRF [Chen et al. ICCV'21], etc…

Challenging for predicting completely unseen large scenes

Surface Extraction from NeRF

Volume Density Fields thresholding $\sigma > c$

NeRF's Volume Density

$\sigma = 1$ $\sigma = 10$ $\sigma = 50$ $\sigma = 100$ $\sigma = 500$

- No explicit definition of surface
- Surface not satisfactory

NeRF + Signed Distance Function (SDF)
Learn SDF field as a scene representation

NeRFs with Signed Distance Function (SDF)

Minimum distance to the closest surface with sign (positive, negative).

Signed Distance Function

$$f(x) = \begin{cases} -d(x, \partial\Omega) & x \in \Omega \\ d(x, \partial\Omega) & x \notin \Omega \end{cases}$$

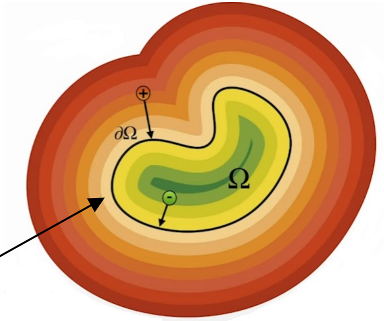

Surface

$$S = \{x \in R^3 | f(x) = 0\}$$

Zero-level set

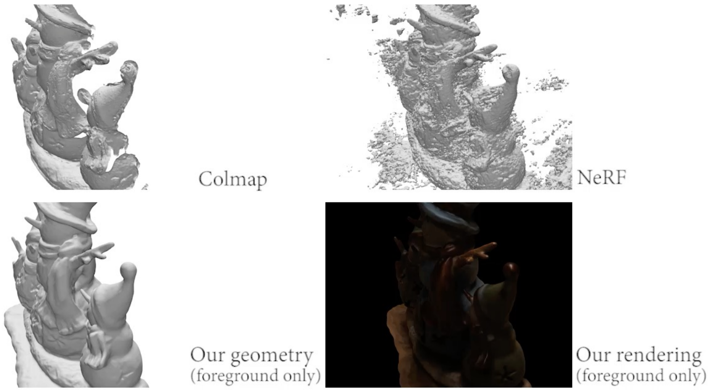

NeuS, VolSDF, Neuralangelo, etc.

$$(x, d) \in R^5$$

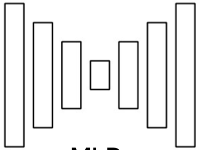

MLP

Volume Density -> SDF

$$(\boldsymbol{c_\theta}, \boxed{\boldsymbol{s_\theta}}) \in \boldsymbol{R^4}$$

UFORecon: Generalizable Sparse-View Surface Reconstruction from Arbitrary and UnFavOrable Sets

CVPR 2024

Youngju Na, Woo Jae Kim, Kyu Beom Han, Suhyeon Ha, Sung-eui Yoon

KAIST

arXiv Code

Generalizable Surface Reconstruction

- Reconstruct from unseen objects or scene
- Few-Shot (3-5 images)

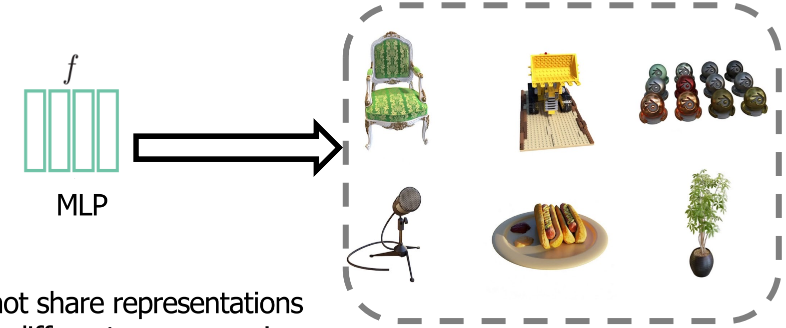

Cannot share representations
across different scenes or views

Scene generalizability

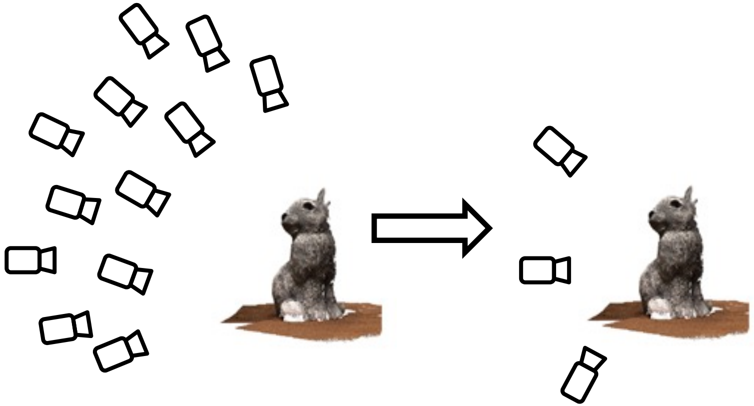

Dense-view Reconstruction
(~100 images)

Sparse-view Reconstruction

Viewpoints Generalizability

KAIST

Assumption: Optimal View Set Assumption

Only consider Predefined **Optimal Camera View Set** as inputs both in train and test time.
Optimal Camera view is defined by view-selection score [1] or nearest neighbors.

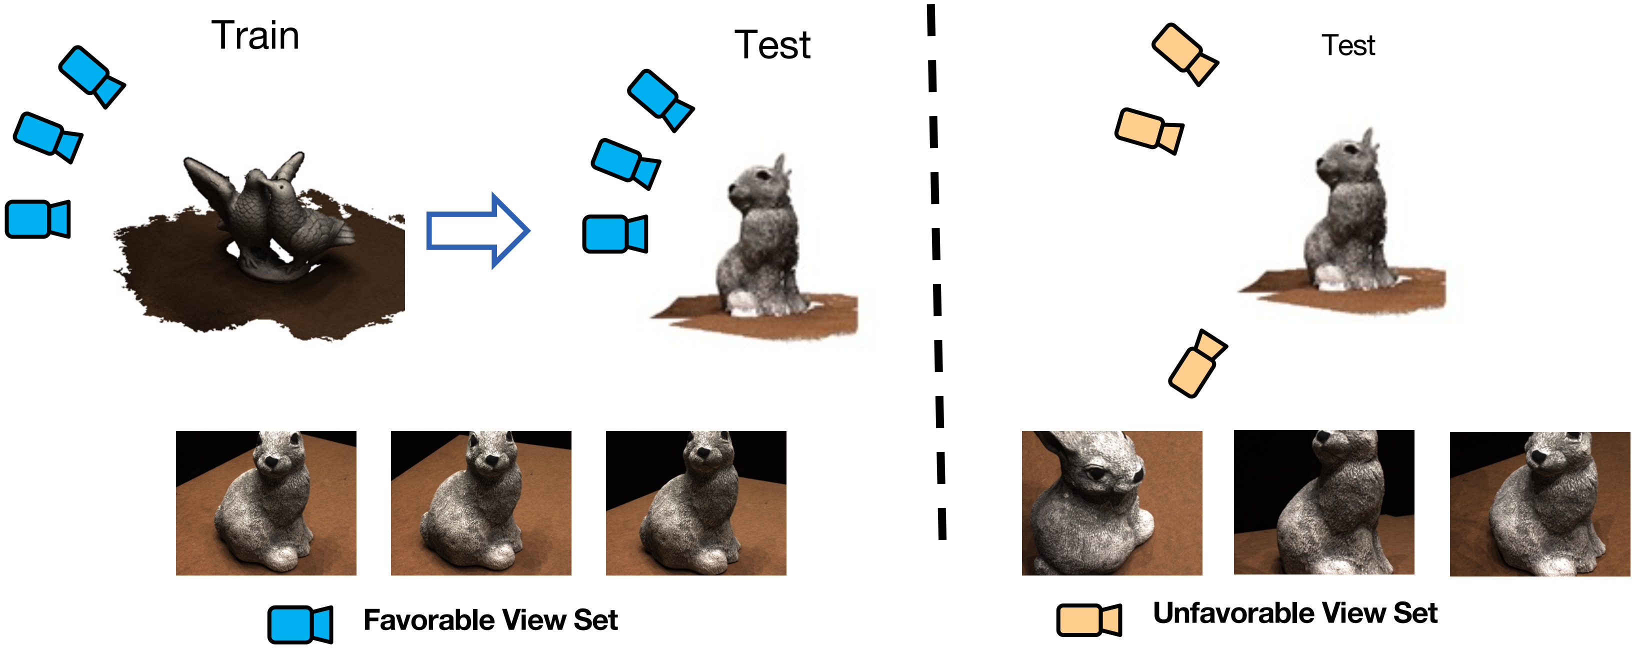

[1]: MVSNet: Depth Inference for Unstructured Multi-view Stereo

Observation: Degenerate Solution for Unfavorable Sets

Test with **VolRecon [CVPR'23]** shows that unfavorable pair outputs degenerate solution.

Best View Set (common Protocol)

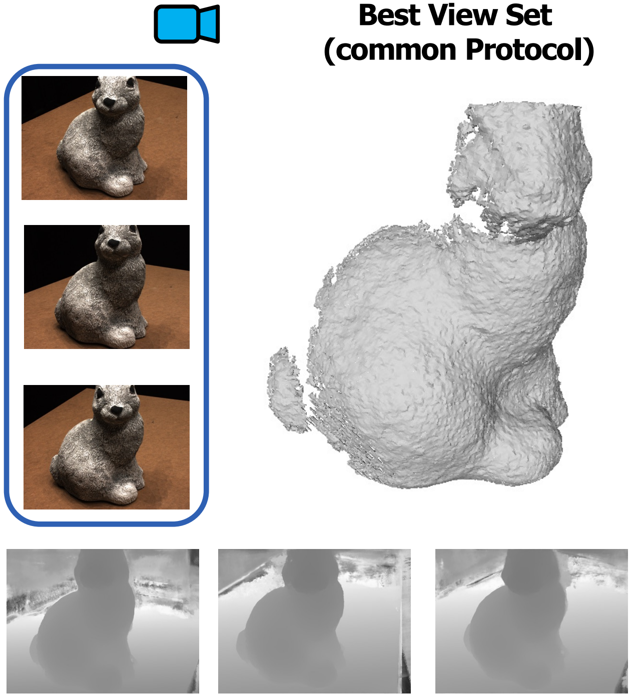

Estimated Depth

Unfavorable View Set

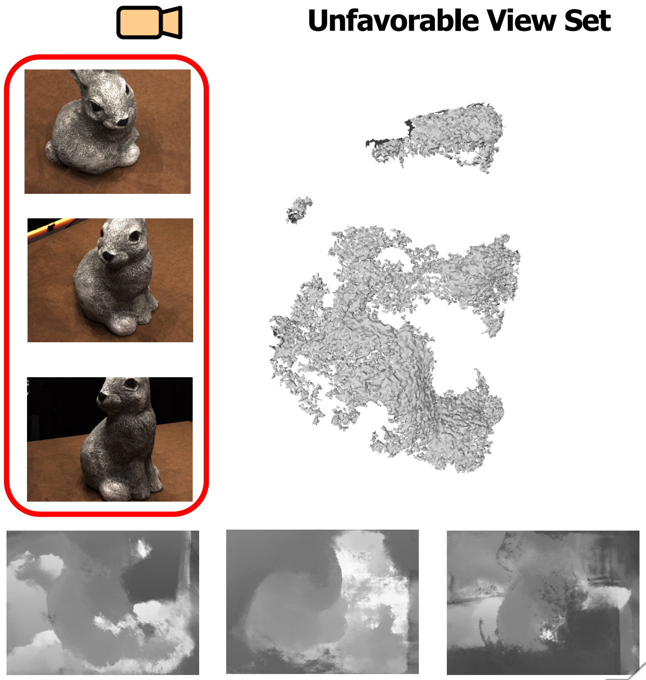

Estimated Depth

Not View-Combination Generalizable

Rethinking Inference Scenario

- In practice, we can’t always guarantee **Optimal View Sets**
- View Sets in train time ≠ View Sets in inference time

View-Combination Generalizability

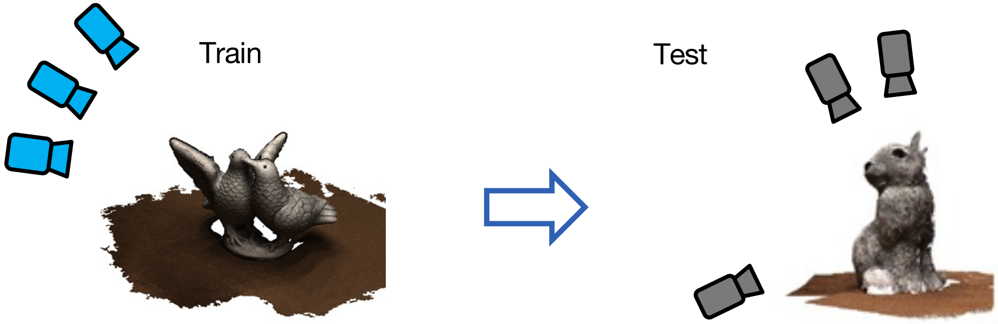

Generalizability**: Scene + Viewpoints + View Combination**

Modeling Correlation between input Views

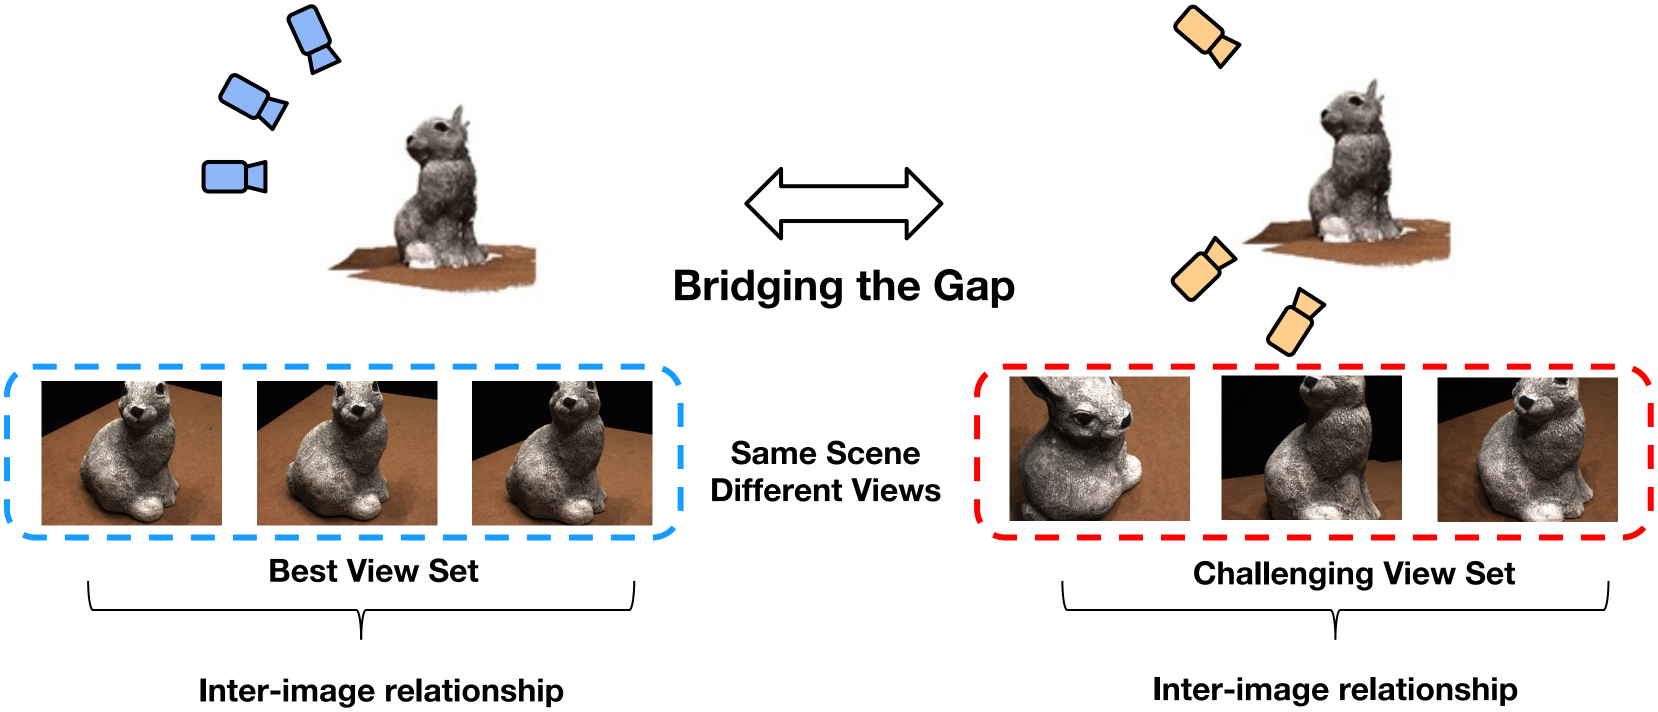

Utilizing inter-image relationship as robust prior

Pair-wise Cross Transformer

Extract feature considering the relationship across the images

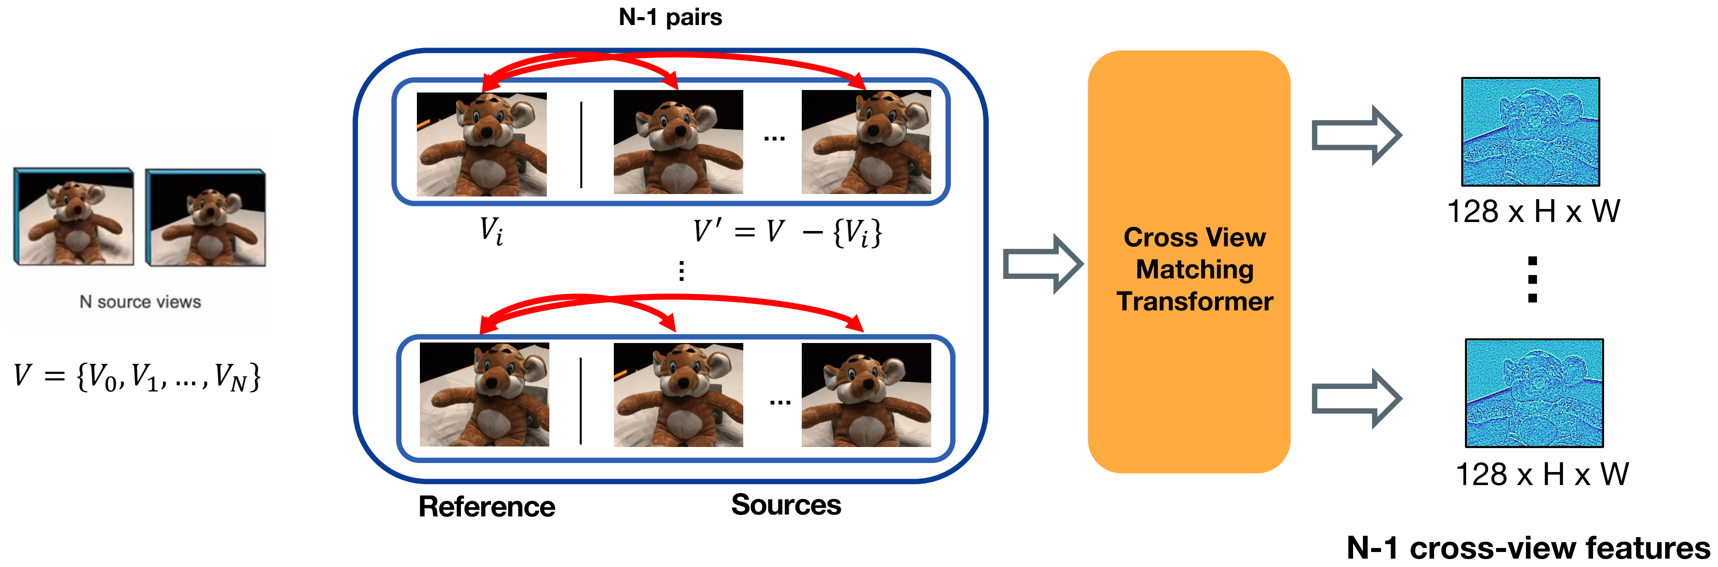

Global Correlation Frustums

Building Correlation Volume from cross-view features.

Learn global correlation among all source images.

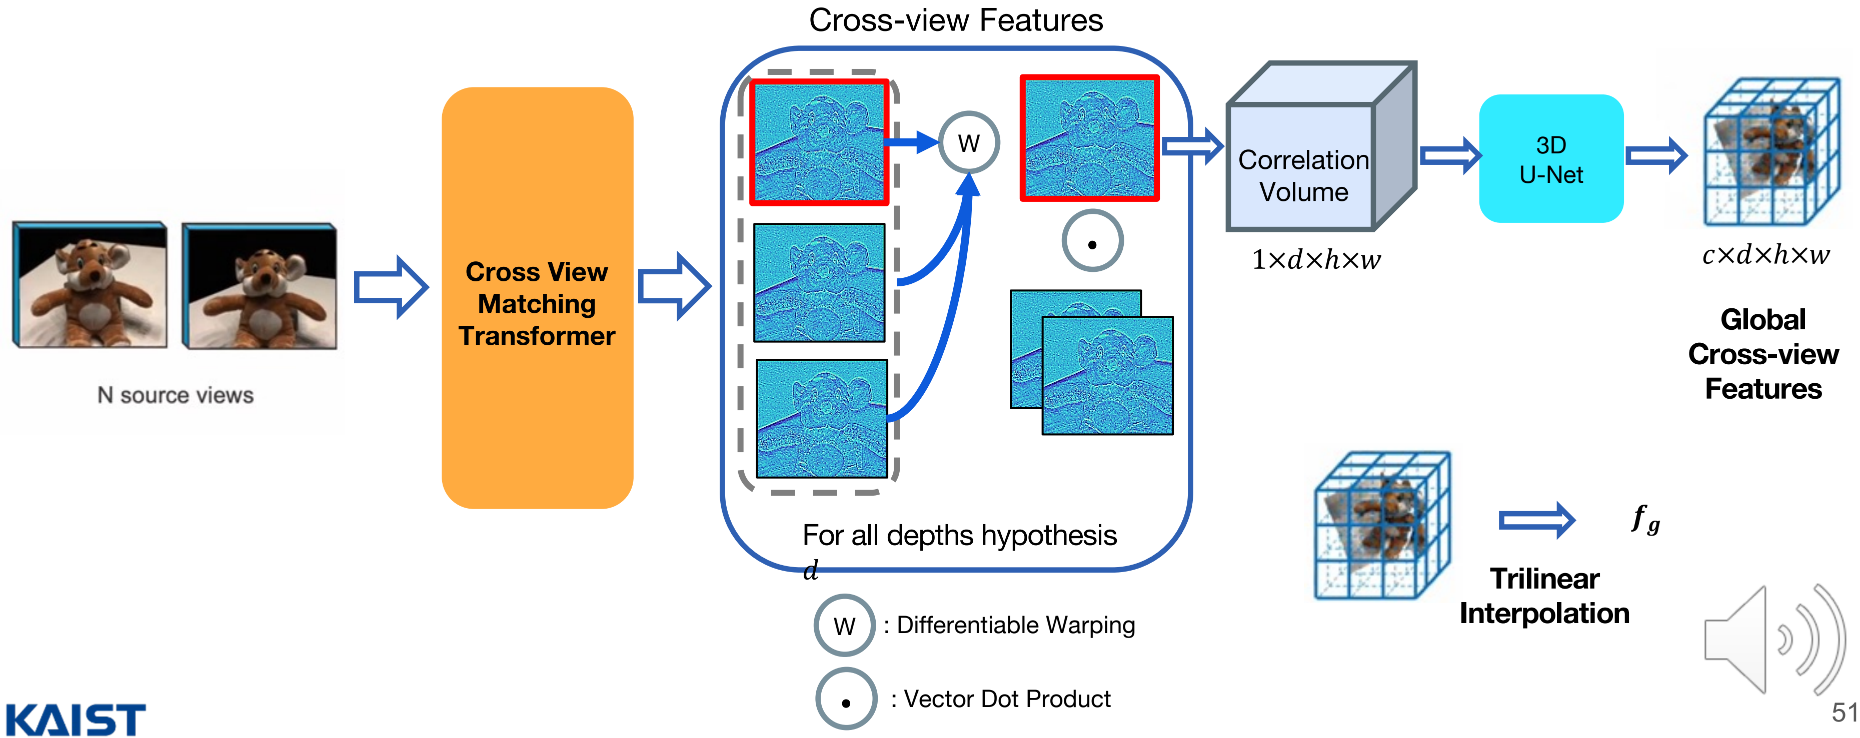

Overall Pipeline of UFORecon

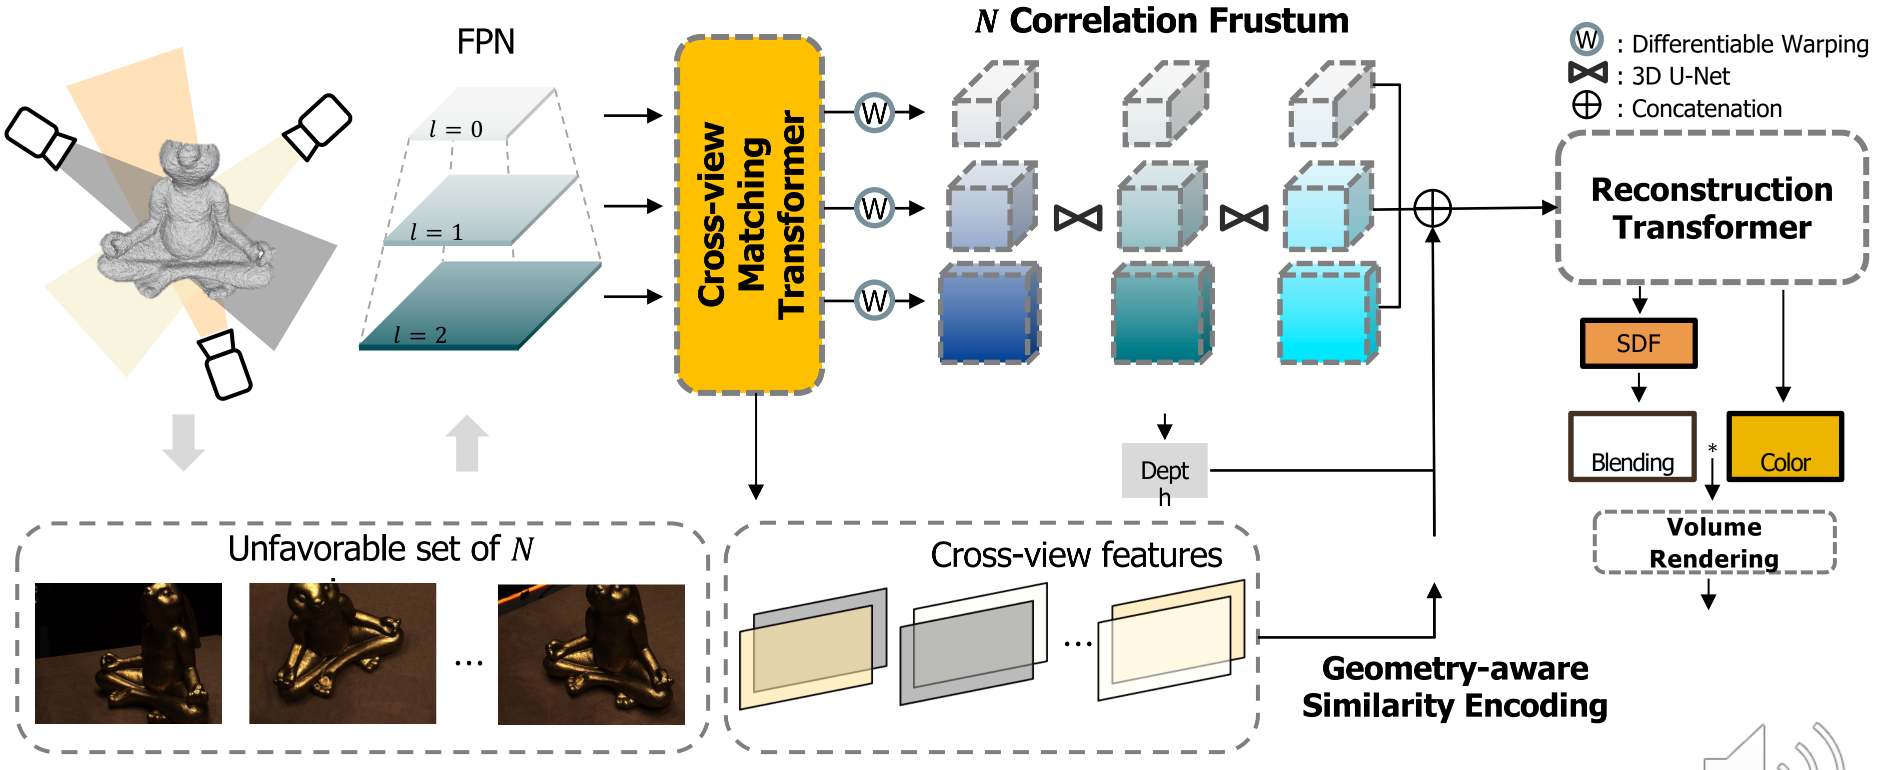

KAIST

Results

Unfavorable Source Views

Results

Volrecon

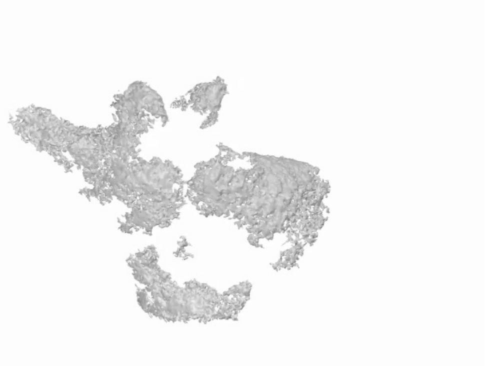

Ours

Ours (Random Set Training)

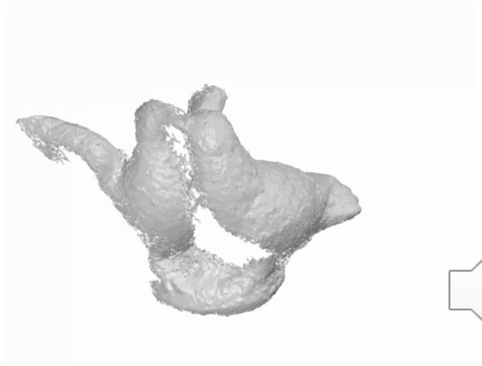

Results

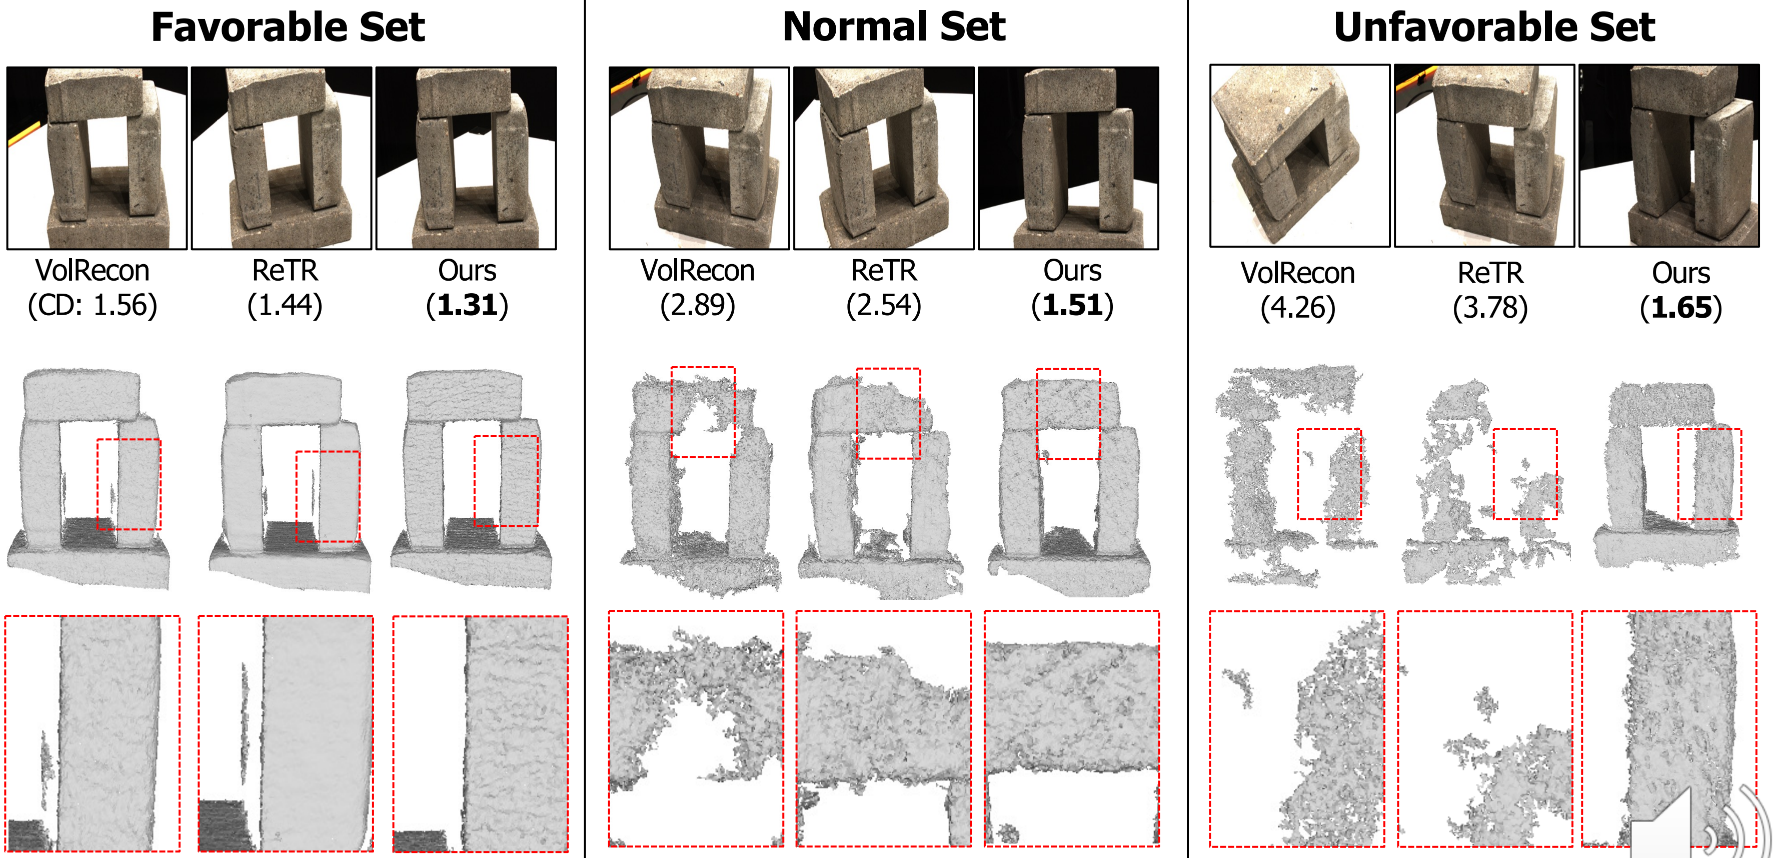

KAIST

Radiance Fields with Generative Models

Generating NeRFs from 2D Generative Models

DreamFusion [Poole et al. arXiv 2022]

Enabling specific edits

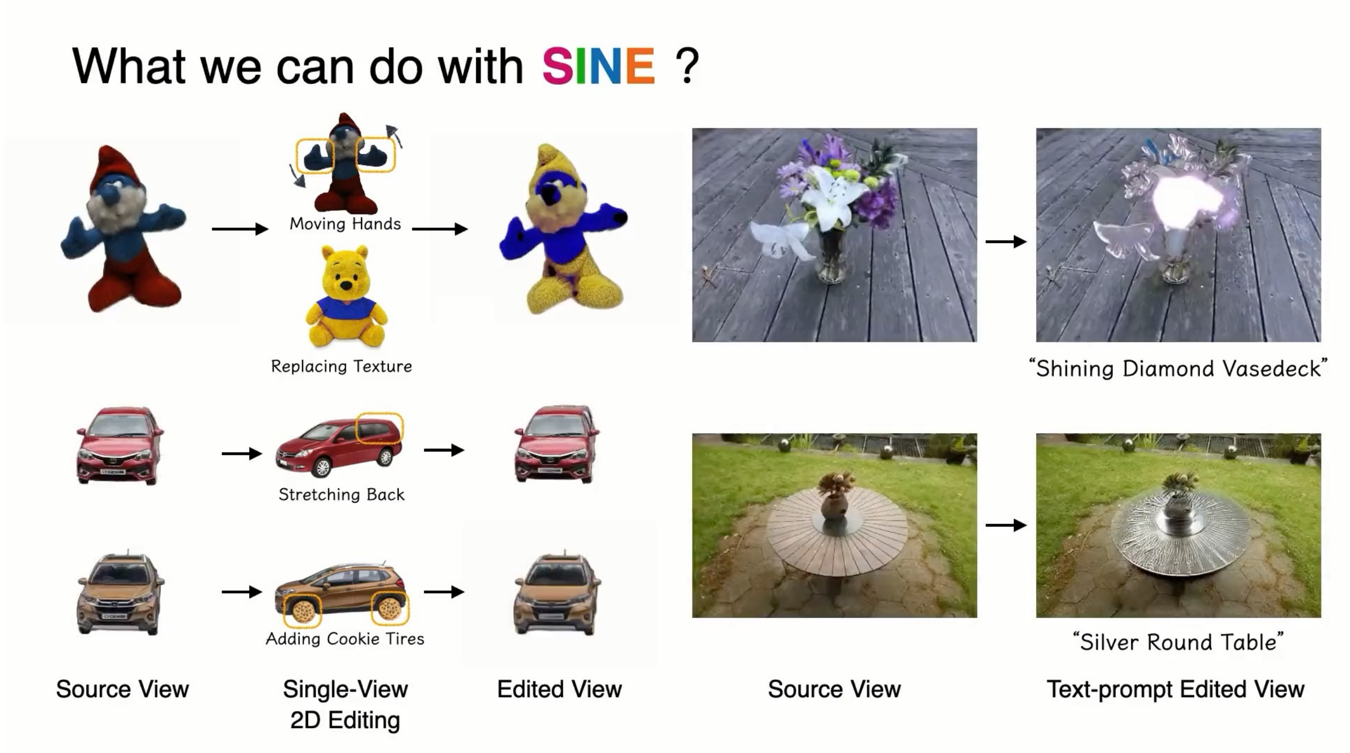

https://zju3dv.github.io/sine/

Semantic Editing

ClipNeRF Wang et al. CVPR 2022, Feature Field Distillation - Kobayashi et al. NeurIPS 2022 … Many more papers here!!

Manipulating captured scenes

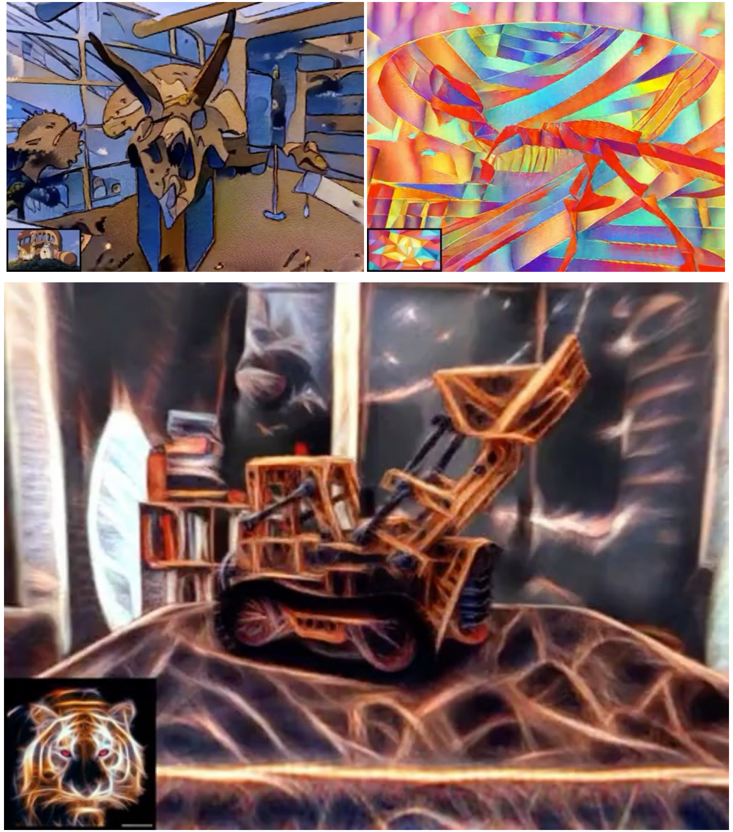

Artistic Radiance Fields Zhang et al. ECCV 2022

Radiance Fields in 2024

NeRF showed various possibilities and potentials but rendering is SLOW..

Next Representations?

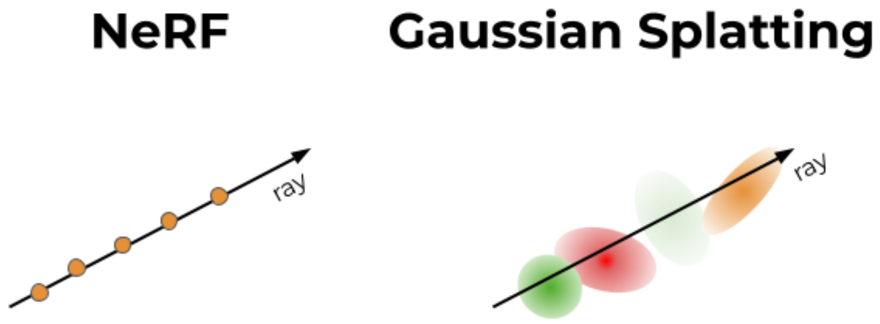

[Short Video]

[Long Video]

Check out yourself! https://github.com/MrNeRF/awesome-3D-gaussian-splatting

Thank you

yjna2907@kaist.ac.kr