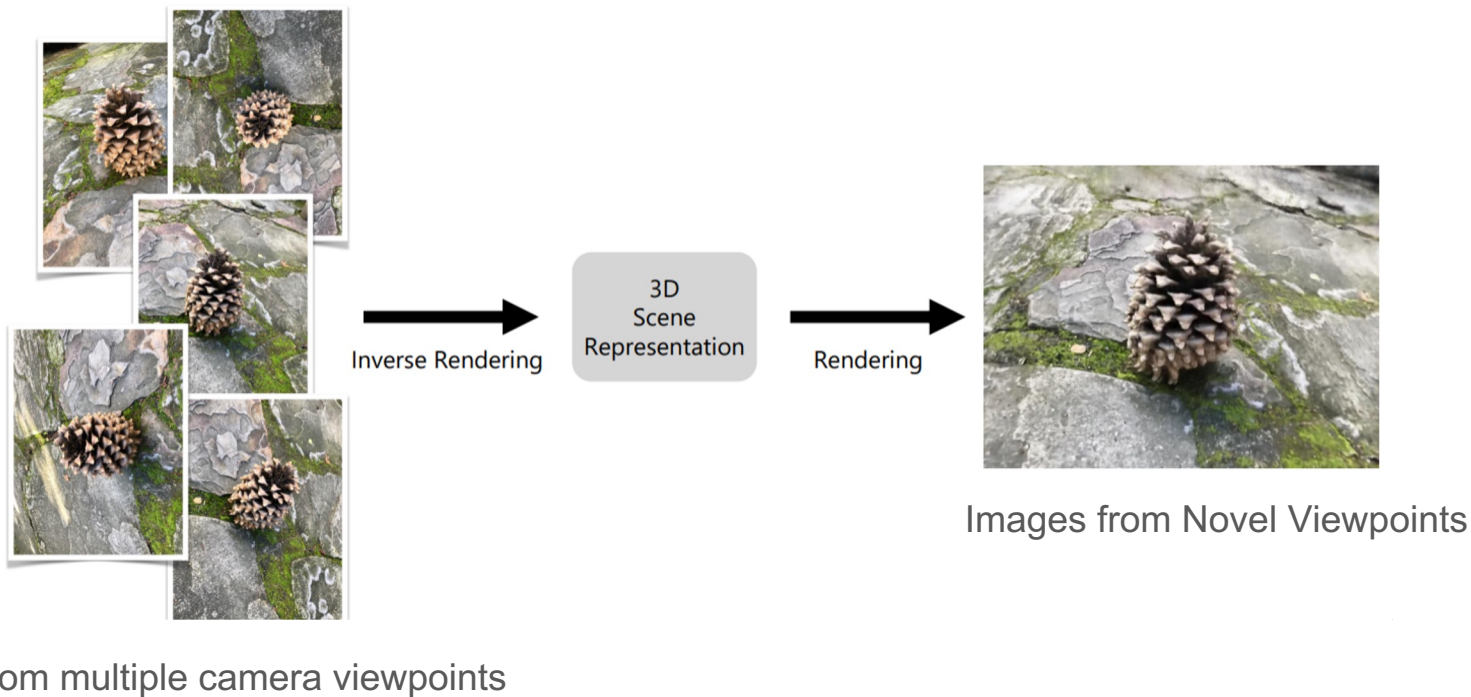


Radiance Fields: NeRF to 3D Gaussian Splatting

Youngju Na
M.S. Student @ SGVR Lab




Background: Novel View Synthesis



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)




Implicit Representation


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

x : coordinate

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



x




Neural Network

$f(x) = ?$



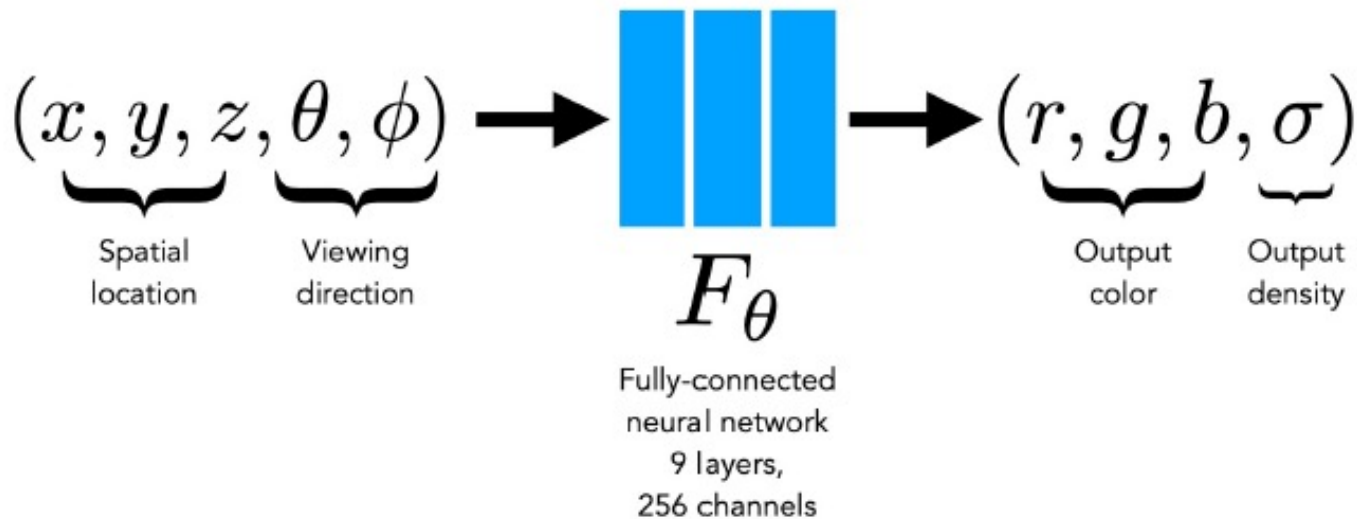
Represent 3D Scene as Continuous functions

Signed Distance Function (SDF) or Occupancy Fields



NeRF 3D Representations

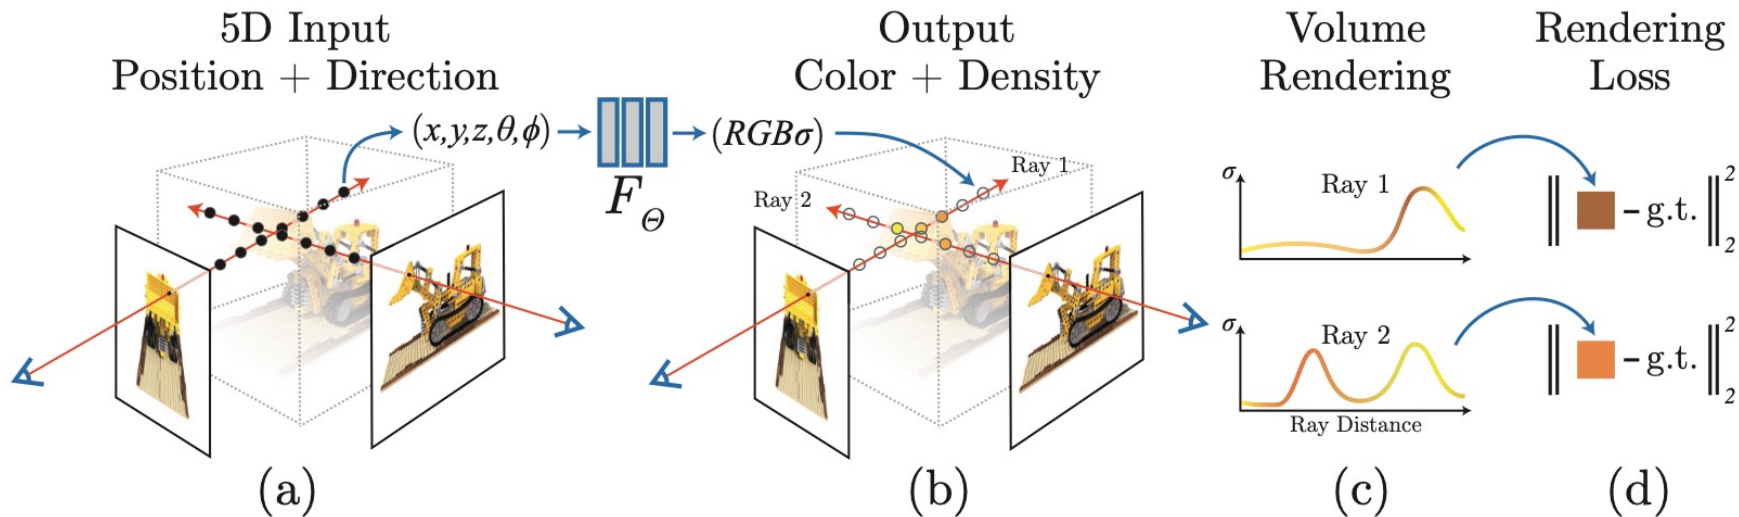
Neural Network as a continuous shape representation.



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

Method Overview

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



Neural Volumetric Rendering

Neural Volumetric **Rendering**


computing color along rays
through 3D space



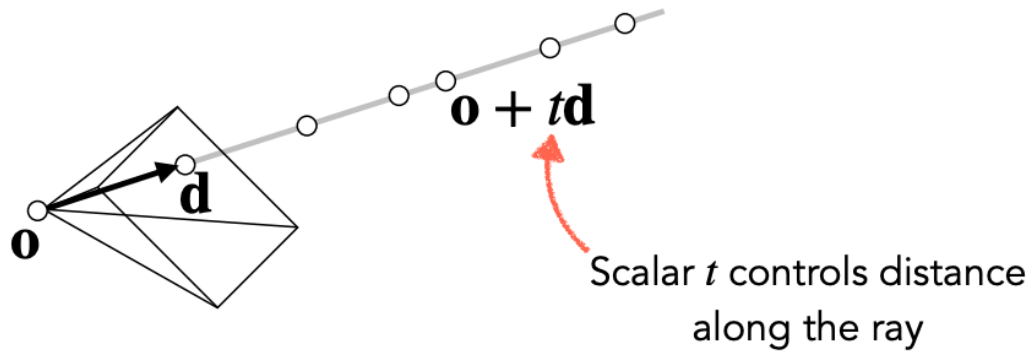
What color is this pixel?

Cameras and rays

- We need the mathematical mapping from $(camera, pixel) \rightarrow ray$
- Then can abstract underlying problem as learning the function $ray \rightarrow color$ (the “plenoptic function”)



Calculating points along a ray




Neural **Volumetric** Rendering

continuous, differentiable
rendering model without
concrete ray/surface intersections




Surface vs. volume rendering




Surface rendering — loop over geometry, check for ray hits

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 $\mathbf{c}(t)$

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




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”

PDF for ray termination



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


$$\begin{aligned} P[\text{first hit at } t] &= P[\text{no hit before } t] \times P[\text{hit at } t] \\ &= T(t)\sigma(t)dt \\ &= \exp\left(-\int_{t_0}^t \sigma(s) ds\right) \sigma(t) dt \end{aligned}$$

Expected value of color along ray

This means the expected color returned by the ray will be


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

color




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 $\{t_1, t_2, \dots, t_{n+1}\}$ with lengths $\delta_i = t_{i+1} - t_i$

Approximating the nested integral



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

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 $\mathbf{r}(t) = \mathbf{o} + t\mathbf{d}$:

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

weights → T_i and α_i → colors

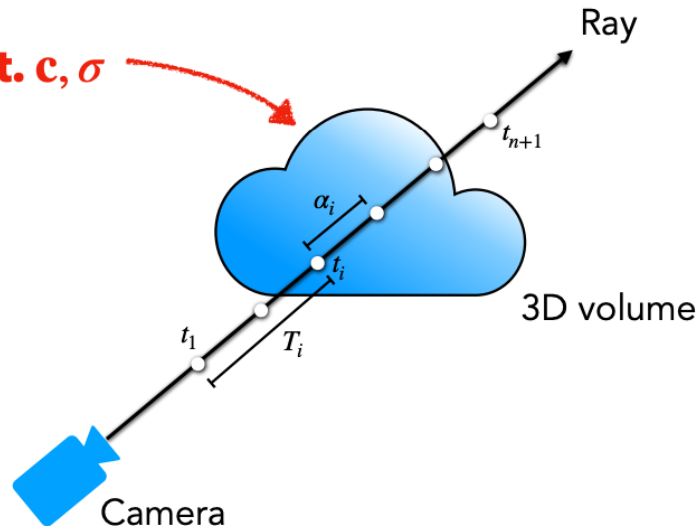
differentiable w.r.t. \mathbf{c}, σ

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 :

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



Video

Novel View Synthesis & View Dependency




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 exciting applications in Computer Vision, Graphics, Robotics and more. In this tutorial, we will present the fundamentals of Neural Volumetric Rendering from the first principles, including its relation to the history of image core components and their derivations, common practices, future challenges, and hands-on. This half-day tutorial is not to present a series of talks on recent papers in this area, but to provide a hands-on experience for novice and intermediate researchers to deeply understand the material by abstracting the concepts of Neural Volumetric Rendering.

Organizers




[Matt Tancik](#)
UC Berkeley




[Ben Mildenhall](#)
Google



[Pratul Srinivasan](#)
Google



[Jon Barron](#)
Google



[Angjoo Kanazawa](#)
UC Berkeley

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

Radiance Fields



Subscribe

[Platforms](#) [About](#) [Partnerships](#) [Affiliates](#) [Contribute](#) [Contact](#) [VRAM Calculator](#)

PLATFORMS

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



RESEARCH

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




Resources

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



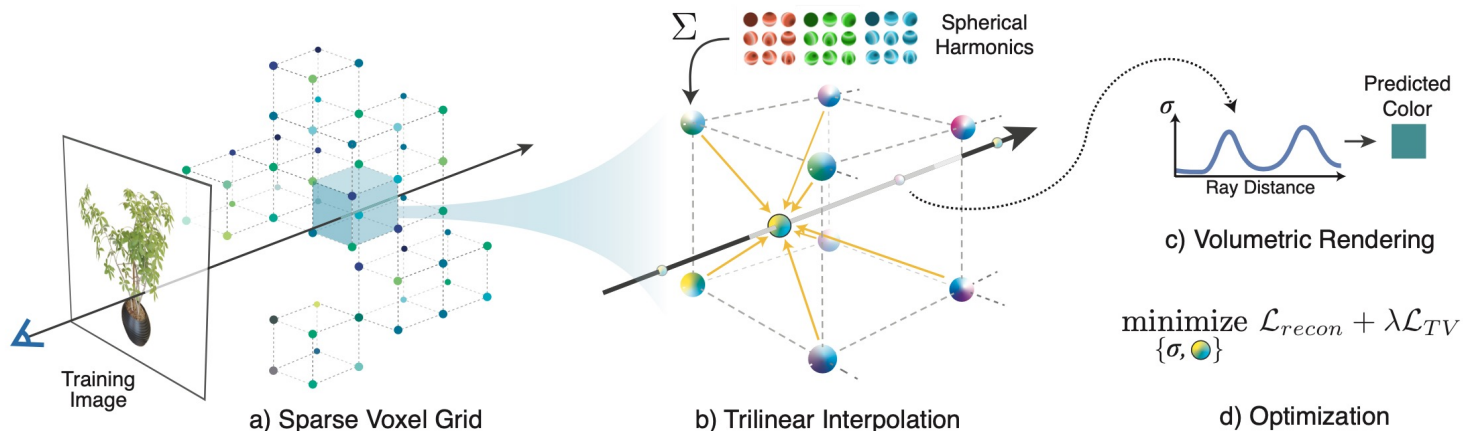
A collaboration friendly studio for NeRFs



3D Gaussian Splatting

SIGGRAPH 2023 best paper award

Plenoxel [CVPR'22]: Fast Optimization / Rendering




What does it optimize?

- Spherical Harmonics (SH) Coefficients
- Volume Density

=> No Neural Network

Spherical Harmonics (SH)

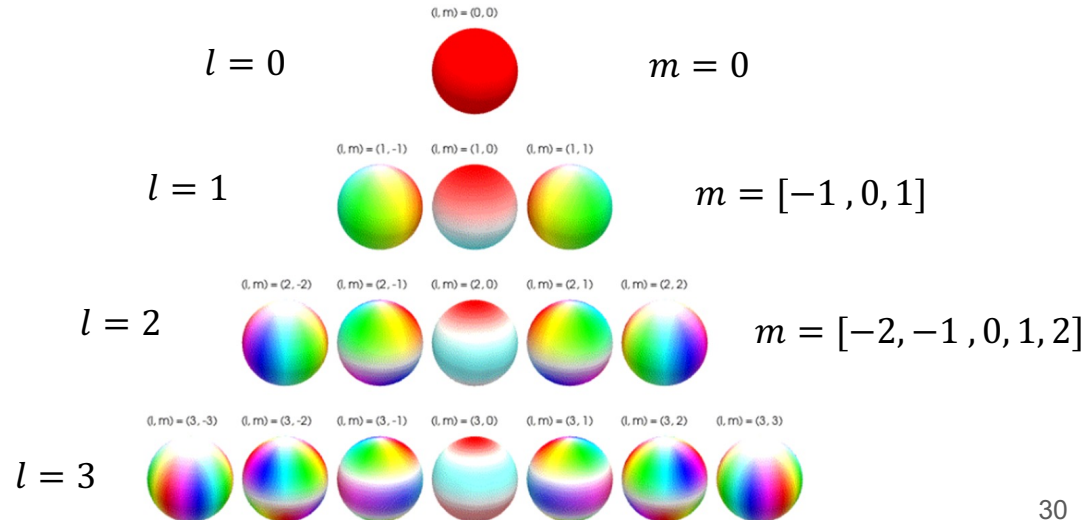


(θ : polar angle, ϕ : azimuthal angle)


If we solve Laplace equation in surface point,

$$Y_l^m(\theta, \phi) = \sqrt{\frac{(2l+1)(l-|m|)!}{4\pi(l+|m|)!}} P_l^{|m|}(\cos\theta) e^{im\phi}$$

$$P_\ell^{(|m|)}(\cos\theta) = (-1)^m \frac{(\ell+|m|)!}{(\ell-|m|)!} P_\ell^{(-|m|)}(\cos\theta)$$



Fast Optimization / Rendering : Plenoxel [CVPR'22]

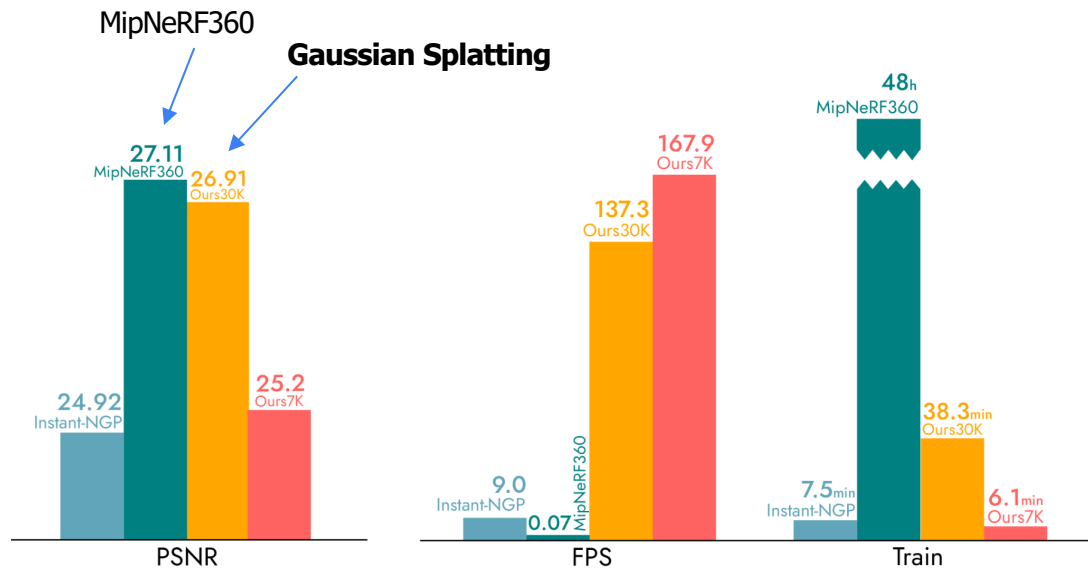
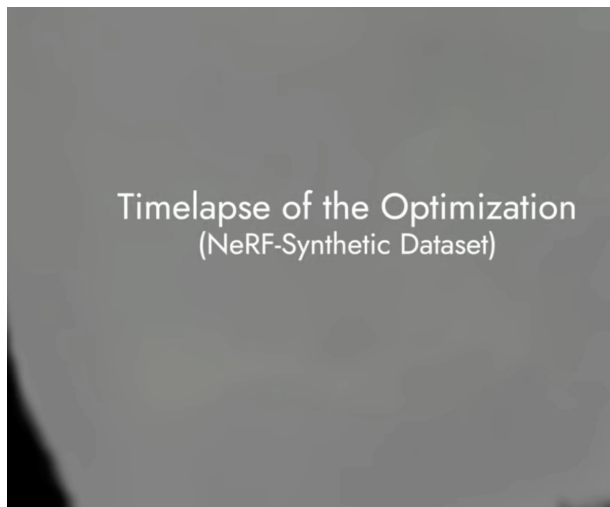


Mip-NeRF360 vs Plenoxels

Train	1.6 days	~ 30 mins
FPS	0.06s	6.8s
Memory	8.6MB	2.1GB

Gaussian Splatting: Fast 3D Reconstruction and Rendering (3DGS)


- Gaussian Splatting is a fast training and real-time rendering framework.
- Takes \sim 1hour training and achieves **>120 FPS**.



Gaussian Splatting: Fast 3D Reconstruction and Rendering


- Representation: 3D Gaussians

1D Gaussians




$$g(x) = \frac{1}{\sigma\sqrt{2\pi}} \exp\left(-\frac{(x - \mu)^2}{2\sigma^2}\right)$$

2D Gaussians



3D Gaussians




$$g(x) = \exp\left(-\frac{1}{2}(x - \mu)^T \Sigma^{-1}(x - \mu)\right)$$

Generalized multivariate gaussian distribution (without normalization)

Gaussian Splatting: Fast 3D Reconstruction and Rendering

- Representation: 3D Gaussians


$$f_i(p) = \sigma(\alpha_i) \exp\left(-\frac{1}{2}(p - \mu_i)\Sigma_i^{-1}(p - \mu_i)\right)$$



Each 3D Gaussian is parametrized by:

- **Mean μ :** 3D position (x, y, z)
- **Covariance $\Sigma = RSS^T R^T$:** (Scale S, Rotation R)
- **Opacity: $\sigma(\alpha)$**
- **Color parameters:** spherical harmonics (SH) coefficients.


Input Views



3D Gaussian




Novel Views




3D Gaussian Splatting (3D-GS)

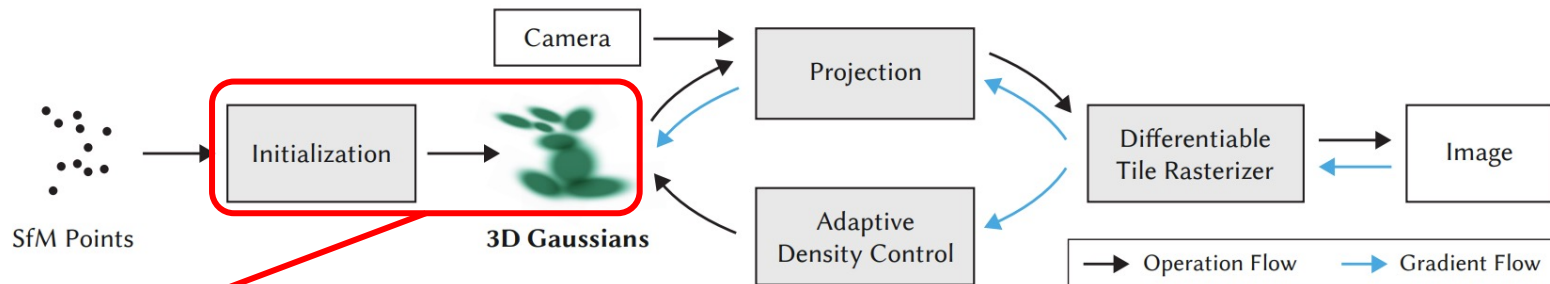
3D Gaussian Splatting for Real-Time Radiance Field Rendering [SIGGRAPH'23 Best Paper Award]




Step 1. Initialize points via Structure from Motion (SfM)



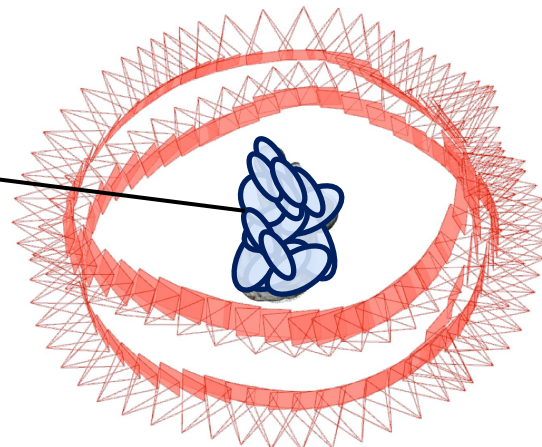
3D Gaussian Splatting (3D-GS)



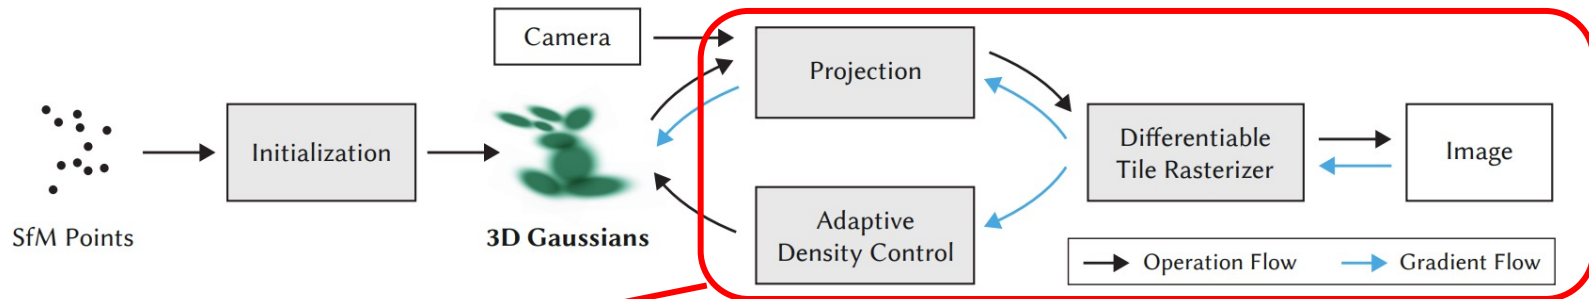
Step 2. Represent points with multivariate Gaussians and assign parameters




μ : Position
 Σ : Covariance
 S : Color
 α : alpha



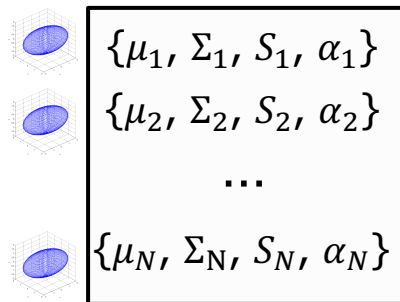
3D Gaussian Splatting (3D-GS)



Step 3. Convert Gaussian primitives into an Image (Rasterize)



Multivariate Gaussians



1. Project into 2d according to the camera
2. Sort by depths
3. Aggregate Gaussians (alpha-blending)

Rasterization



Training (No Neural Network)




optimize



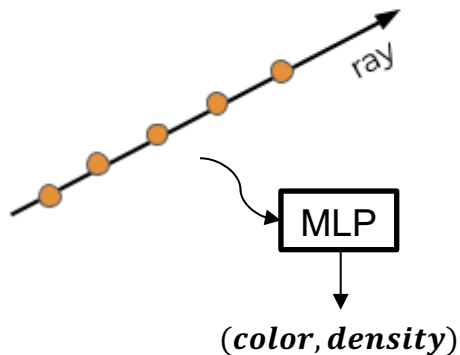
Experimental Results

3DGS has shown high quality rendering with 130+ FPS (real-time).



Dataset Method Metric	Mip-NeRF360					
	SSIM [↑]	PSNR [↑]	LPIPS [↓]	Train	FPS	Mem
Plenoxels	0.626	23.08	0.463	25m49s	6.79	2.1GB
INGP-Base	0.671	25.30	0.371	5m37s	11.7	13MB
INGP-Big	0.699	25.59	0.331	7m30s	9.43	48MB
M-NeRF360	0.792 [†]	27.69 [†]	0.237 [†]	48h	0.06	8.6MB
Ours-7K	0.770	25.60	0.279	6m25s	160	523MB
Ours-30K	0.815	27.21	0.214	41m33s	134	734MB

NeRF




Implicit representation (Neural networks)

Volume Rendering

Slow rendering, low memory

Gaussian Splatting



Explicit representation (3D Gaussians)

Rasterization

Real-time rendering, higher memory




Applications

NeRF/3DGS Applications

1. Assume static scene => Dynamic Scene
 2. Generative Models (Text-to-3D, Image-to-3D, etc.)
 3. Relighting / Light Modeling
 4. Navigation / Autonomous Driving
- Etc.

List goes on and on...!
NeRF has been cited 6800+
3DGS, 1400+

The world we capture is usually Dynamic / Deformable



Dynamic NeRFs / 3DGS



RoDynRF, Liu et al. CVPR'23



DynIBaR, Li et al. CVPR'23




Bae et al. ECCV'24

NeRF/3DGS requires Per-Scene Optimization

Generalizable Methods with Prior Knowledge

NeRF/3DGS requires Per-Scene Optimization with Dense Views

1. Scene-specific representation



2. Sparse input camera viewpoints





Not Generalizable

Cannot share representations across scenes or views

Generalizable NeRF / 3DGS



- Note 1. No Per Scene Optimization ❌, Generalizable ✅
- Note 2. No Dense Views ❌, Only 2-3 images ✅

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




• One-Shot 3DGS (PixelSplat [Charatan et al. CVPR'24])

Input Views



Novel Views



Radiance Fields with Generative Models

Text-to-3D, Image-to-3D

LGM, ECCV'24



“motorcycle”



“mech suit”



“ghost lantern”



“furry fox head”



“dresser”



“swivel chair”



“astronaut”



“mushroom house”

DreamScenes, ECCV'24



A DSLR photo of a living room



DSLR photo of a cyberpunk style bedroom, cyberpunk style




A minecraft cubes world with lake and mountains in the far distance and grass cubes in the near distance

SLAM

Localization and Mapping

GaussNav, ECCV'24

Goal Image



RGB Observation



Semantic Gaussians



Other applications

Table of contents

- [Seminal Paper introducing 3D Gaussian Splatting](#)
- [3D Object Detection](#)
- [Autonomous Driving](#)
- [Avatars](#)
- [Classic work](#)
- [Compression](#)
- [Diffusion](#)
- [Dynamics and Deformation](#)
- [Editing](#)
- [Language Embedding](#)
- [Mesh Extraction and Physics](#)
- [Misc](#)
- [Regularization and Optimization](#)
- [Rendering](#)
- [Reviews](#)
- [SLAM](#)
- [Sparse](#)
- [Navigation and Autonomous Driving](#)
- [Poses](#)
- [Large-Scale](#)

Explore other applications that
might interest you

<https://github.com/MrNeRF/awesome-3D-gaussian-splatting>

Thank you

yjna2907@kaist.ac.kr