#### Neural Rendering & Neural Radiance Fields(NeRFs)

Arslan Artykov, PhD Candidate @ ENPC



#### **Novel View Synthesis**

• Step 1: use the images to calibrate cameras (sparse scene geometry)



#### **Novel View Synthesis**

• Step 2: use the images to build a dense 3D scene representation



Kerbl, Kopanas et al. "3D Gaussian Splatting for Real-Time Radiance Field Rendering" SIGGRAPH 2023

### **Forward Graphics (Rendering)**

• Given a 3D scene specification and Cameras, yield images



#### **Inverse Graphics**



#### Put All Together...



#### Neural Radiance Fields (NeRF) as an approach to inverse rendering



## **Deep learning for 3D reconstruction**

 Previously: we reconstruct geometry by running stereo or multi-view stereo on a set of images

"Classical" approach

- How can we leverage powerful tools of deep learning?
  - Deep neural networks
  - GPU-accelerated stochastic gradient descent

## NeRF and related methods – Key ideas

- We need to create a loss function and a scene representation that we can optimize using gradient descent to reconstruct the scene
- Differentiable rendering

## **Rendering (Signature)**

- Given
  - Scene: {lights, materials, geometry, ...}



- Camera Ray (origin + direction)



- Generate
  - color of the ray/pixel



#### **Differentiable Rendering**



Given an observable variable (pixel colors), we will build a differentiable forward model that we then use to estimate unobserved (latent) variables (geometry, appearance)

## NeRF == Differentiable Rendering with a **Neural Volumetric Representation**

# Neural Volumetric Rendering

querying the radiance value along rays through 3D space



# Neural Volumetric Rendering

continuous, differentiable rendering model without concrete ray/surface intersections



# **Neural** Volumetric Rendering

using a neural network as a scene representation, rather than a voxel grid of data



## NeRF: Representing Scenes as Neural Radiance **Fields for View Synthesis** ECCV 2020



#### Ben Mildenhall\*



**UC Berkeley** 



Pratul Srinivasan\*



**UC Berkeley** 



Matt Tancik\*



**UC Berkeley** 



Jon Barron



Google Research





Ravi Ramamoorthi

UC San Diego



Ren Ng



**UC Berkeley** 





#### **NeRF** Overview

- Volumetric rendering
- Neural networks as representations for spatial data
- Neural Radiance Fields (NeRF)

#### **NeRF** Overview

#### Volumetric rendering

- Neural networks as representations for spatial data
- Neural Radiance Fields (NeRF)

#### **Volume Rendering**

Scene is a cloud of tiny colored particles



#### Volume Rendering Digest (for NeRF)

Andrea Tagliasacchi<sup>1,2</sup> Ben Mildenhall<sup>1</sup> <sup>1</sup>Google Research <sup>2</sup>Simon Fraser University

Neural Radiance Fields [3] employ simple volume rendering as a way to overcome the challenges of differentiating through ray-triangle intersections by leveraging a probabilistic notion of visibility. This is achieved by assuming the scene is composed by a cloud of light-emitting particles whose density changes in space (in the terminology of physically-based rendering, this would be described as a volume with absorption and emission but no scattering [4, Sec 11.1]. In what follows, for the sake of exposition simplicity, and without loss of generality, we assume the emitted light *does not* change as a function of view-direction. This technical report is a condensed version of previous reports [1, 2], but rewritten in the context of NeRF, and adopting its commonly used notation<sup>1</sup>.

**Transmittance.** Let the density field  $\sigma(\mathbf{x})$ , with  $\mathbf{x} \in \mathbb{R}^3$  indicate the differential likelihood of a ray bitting a particle (i.e. the probability of fitting a particle white traveling an infinitesimal distance). We reparameterize the density along a given ray  $\mathbf{r}=(\mathbf{0}, \mathbf{d})$  as a scalar function  $\sigma(t)$ , since any point  $\mathbf{x}$  along the ray can be written as  $\mathbf{r}(t)=\mathbf{o}+t\mathbf{d}$ . Density is closely tied to the transmittance function  $\mathcal{T}(t)$ , which indicates the probability of a ray traveling over the interval [0, t) without hitting any particles. Then the probability  $\mathcal{T}(t+dt)$  of not hitting a particle when taking a differential step dt is equal to  $\mathcal{T}(t)$ , the likelihood of the ray reaching t, times  $(1 - dt - \sigma(t))$ , the probability of not hitting anything during the step:

$$T(t + dt) = T(t) \cdot (1 - dt \cdot \sigma(t)) \qquad (1)$$

$$\frac{\mathcal{T}(t+dt) - \mathcal{T}(t)}{dt} \equiv \mathcal{T}'(t) = -\mathcal{T}(t) \cdot \sigma(t)$$
<sup>(2)</sup>

This is a classical differential equation that can be solved as follows:

$$\mathcal{T}'(t) = -\mathcal{T}(t) \cdot \sigma(t) \tag{3}$$

$$\frac{T'(t)}{T(t)} = -\sigma(t) \qquad (4)$$

$$\int_{a}^{b} \frac{\mathcal{T}'(t)}{\mathcal{T}(t)} dt = -\int_{a}^{b} \sigma(t) dt$$
(5)

$$\log \mathcal{T}(t)|_{a}^{b} = -\int_{a}^{o} \sigma(t) dt$$

$$\mathcal{T}(t) = \left( \int_{a}^{b} \sigma(t) dt \right)$$
(6)

$$\mathcal{T}(a \to b) \equiv \frac{\mathcal{T}(b)}{\mathcal{T}(a)} = \exp\left(-\int_{a}^{b} \sigma(t) \, dt\right) \tag{7}$$

where we define  $T(a \rightarrow b)$  as the probability that the ray travels from distance a to b along the ray without hitting a particle, which is related to the previous notation by  $T(t) = T(0 \rightarrow t)$ .

#### **Volumetric Formulation for NeRF**



#### **Volumetric Density**

• Probability that ray stops in a small interval around t is  $\sigma(t) dt$ 

 $\sigma(t)$  is called Volume[tric] Density



#### **Scene Representation**

• Our 3D scene representation is therefore a field:

$$\Phi: \mathbb{R}^3 \to \mathbb{R}^3 \times \mathbb{R}^+; \quad \Phi(\mathbf{x}) = (\sigma, \mathbf{c})$$



#### **Occlusion Modelling**

- Should we retrieve the color  $\mathbf{c}(t)$ ? ...only if position  $\mathbf{r}(t)$  is not occluded
  - Transmittance T(t) is the probability of no particles hit in [0, t) range



Relating  $\sigma(t)$  to T(t)

• Hit probabilities are statistically independent along ray

 $P[\text{no hit before } t + dt] = P[\text{no hit before } t] \cdot P[\text{no hit at } t]$  $T(t + dt) = T(t) \cdot (1 - \sigma(t) dt)$ 



#### Transmittance

$$T(0 \to t) := T(t) = \exp\left(-\int_0^t \sigma(u) \, du\right)$$

No hits before t is equal to integral over density up until t.

(1-
$$T(t)$$
): Opacity (1- $T(t)$ )' =  $T(t) \cdot \sigma(t)$ 

#### **Volume Rendering**

$$C = \int_0^\infty T(t) \cdot \sigma(t) \cdot \mathbf{c}(t) \, dt \quad \text{where} \quad T(t) = \exp\left(-\int_0^t \sigma(u) \, du\right)$$

- How do we solve this?
- Discretize the nested integral!

#### **Approximating the Integral**

- Split the ray up into N segments with endpoints  $\{t_1, t_2, ..., t_{N+1}\}$
- Segment length is  $\delta_i = t_{i+1} t_i$ 
  - warning: non-necessarily uniform!





#### **Approximating the Nested Integral**

- Assume volume density/color are constant within interval (i.e. Reimann)
- Warning: constant density ≠ constant transmittance!

$$C = \int_0^\infty T(t) \cdot \sigma(t) \cdot \mathbf{c}(t) \, dt \approx \sum_{i=1}^n \int_{t_i}^{t_{i+1}} T(t) \cdot \sigma_i \cdot \mathbf{c}_i \, dt$$



#### **Forward Model - NeRF**

$$C = \int_0^\infty T(t) \cdot \sigma(t) \cdot \mathbf{c}(t) \, dt \approx \sum_{i=1}^n \int_{t_i}^{t_{i+1}} T(t) \cdot \sigma_i \cdot \mathbf{c}_i \, dt$$
  
get rid of the inner  
integral...

#### **NeRF as Alpha Blending**

$$C = \sum_{n=1}^{N} T_n \cdot \underbrace{(1 - \exp(-\sigma_n \delta_n))}_{\alpha_n \equiv \text{ opacity}} \cdot \mathbf{c}_n \quad \text{where} \quad T_n = \exp\left(\sum_{k=1}^{n-1} - \sigma_k \delta_k\right)$$

$$C = \sum_{n=1}^{N} T_n \cdot \alpha_n \cdot \mathbf{c}_n, \text{ where}$$

$$T_n = \prod_{k=1}^{n-1} (1 - \alpha_k) \dots \text{ segment occlusion}$$

# Volume rendering estimation: integrating color along a ray

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



How much light is blocked earlier along ray:

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

Computing the color for a set of rays through the pixels of an image yields a rendered image





# Volume rendering estimation: integrating color along a ray

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

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

How much light is blocked earlier along ray:

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

Computing the color for a set of rays through the pixels of an image yields a rendered image





#### **NeRF** Overview

- ► Volumetric rendering
- Neural networks as representations for spatial data
- Neural Radiance Fields (NeRF)

#### Toy problem: storing 2D image data



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

#### Toy problem: storing 2D image data



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

#### Naive approach fails!



Ground truth image



Neural network output fit with gradient descent

#### Problem:

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

**Spectral Bias:** Neural networks are biased to fit lower frequency signals.

#### Solution:

# Pass input coordinates through a high frequency mapping first



#### Problem solved!



Ground truth image



Neural network output without high frequency mapping

Neural network output with high frequency mapping

#### **NeRF** Overview

- ► Volumetric rendering
- Neural networks as representations for spatial data
- Neural Radiance Fields (NeRF)

# NeRF = volume rendering + coordinate-based network

#### How do we store the values of $\mathbf{c}, \sigma$ at each point in space



#### Extension: view-dependent field



#### Putting it all together



Train network using gradient descent to reproduce all input views of scene



#### **NeRF Requires(incomplete list):**

• Limited to either bounded or forward facing scenes

• Pixel-perfect camera calibration(e.g. COLMAP)

• Scene to be completely static

- Hundreds of images per scene(overfitting)
- Long training time(~1 day/scene)

• Real-time rendering is almost impossible(volume integration requires multiple samples at each ray)

# Results

# NeRF encodes convincing view-dependent effects using directional dependence



#### NeRF encodes detailed scene geometry with occlusion effects



## Summary

- Represent the scene as volumetric colored "fog"
- Store the fog color and density at each point as an MLP mapping 3D position (x, y, z) to color c and density σ
- Render image by shooting a ray through the fog for each pixel
- Optimize MLP parameters by rendering to a set of known viewpoints and comparing to ground truth images

#### 3D Gaussian Splatting for Real-Time Radiance Field Rendering SIGGRAPH 2023

#### 3D Gaussian Splatting for Real-Time Radiance Field Rendering

BERNHARD KERBL<sup>\*</sup>, Inria, Université Côte d'Azur, France GEORGIOS KOPANAS<sup>\*</sup>, Inria, Université Côte d'Azur, France THOMAS LEIMKÜHLER, Max-Planck-Institut für Informatik, Germany GEORGE DRETTAKIS, Inria, Université Côte d'Azur, France



Fig. 1. Our method achieves real-time rendering of radiance fields with quality that equals the previous method with the best quality [Barron et al. 2022], while only requiring optimization times competitive with the fastest previous methods [Fridovich-Keil and Yu et al. 2022]. Kuller et al. 2022]. Key to this performance is a novel 3D Gaussian scene representation coupled with a real-time differentiable renderer, which offers significant speedup to both scene optimization and novel view synthesis. Note that for comparable training times to instantNGP [Müller et al. 2022], we achieve similar quality to theirs; while this is the maximum quality they reach, by training for 5 min we achieve state-oft-he-art quality, even slightly better than Mip-NeRF306 [Barron et al. 2022].

Radiance Field methods have recently revolutionized novel-view synthesis of scenes captured with multiple photos or videos. However, achieving high visual quality still requires neural networks that are costly to train and render, while recent faster methods inevitably trade off speed for quality. For unbounded and complete scenes (rather than isolated objects) and 1080p resolution rendering, no current method can achieve real-time display rates. We introduce three key elements that allow us to achieve state-of-the-art visual quality while maintaining competitive training times and importantly allow high-quality real-time (> 30 fps) novel-view synthesis at 1080p resolution. First, starting from sparse points produced during camera calibration, we represent the scene with 3D Gaussians that preserve desirable properties of continuous volumetric radiance fields for scene optimization while avoiding unnecessary computation in empty space; Second, we perform interleaved optimization/density control of the 3D Gaussians, notably optimizing anisotropic covariance to achieve an accurate representation of the scene: Third, we develop a fast visibility-aware rendering algorithm that supports anisotropic splatting and both accelerates training and allows realtime rendering. We demonstrate state-of-the-art visual quality and real-time rendering on several established datasets.

#### CCS Concepts: • Computing methodologies → Rendering; Point-based models; Rasterization; Machine learning approaches.

"Both authors contributed equally to the paper.

Authors' addresses: Bernhard Kerbh bernhard kerbh@inriafr, Inria, Université Côte d'Azur, France; Georgios Kopanas, georgios kopanas@inriafr, Inria, Université Côte d'Azur, France; Thomas Leinkühller, thomas leinkuchler@mpi-inf.mgg.de, Max-Planck-Institut für Informatik, Germany; George Drettakis, george.drettakis@inria.fr, Inria, Université Côte d'Azu, France.

Publication rights licensed to ACM. ACM acknowledges that this contribution was authored or co-authored by an employee, contractor or affiliate of an attoined government. As such, the Government retains a nonexclusive, royally-free right to publish or reproduce this article, or to allow others to do so, for Government purposes only. © 2018 Copyright held by the owner-author(s), Publication rights licensed to ACM. 0730-0301/2018/0-ART0 315.00 https://doi.org/X0X0XXXXXXX Additional Key Words and Phrases: novel view synthesis, radiance fields, 3D gaussians, real-time rendering

#### ACM Reference Format:

Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 2018. 3D Gaussian Splatting for Real-Time Radiance Field Rendering. ACM Trans. Graph. 0, 0, Article 0 (2018), 14 pages. https://doi.org/XXXXXXX XXXXXXX

#### 1 INTRODUCTION

Meshes and points are the most common 3D scene representations because they are explicit and are a good fit for fast GPU/CUDA-based rasterization. In contrast, recent Neural Radiance Field (NeRF) methods build on continuous scene representations, typically optimizing a Multi-Laver Perceptron (MLP) using volumetric rav-marching for novel-view synthesis of captured scenes. Similarly, the most efficient radiance field solutions to date build on continuous representations by interpolating values stored in, e.g., voxel [Fridovich-Keil and Yu et al. 2022] or hash [Müller et al. 2022] grids or points [Xu et al. 2022]. While the continuous nature of these methods helps optimization, the stochastic sampling required for rendering is costly and can result in noise. We introduce a new approach that combines the best of both worlds: our 3D Gaussian representation allows optimization with state-of-the-art (SOTA) visual quality and competitive training times, while our tile-based splatting solution ensures real-time rendering at SOTA quality for 1080p resolution on several previously published datasets [Barron et al. 2022; Hedman et al. 2018; Knapitsch et al. 2017] (see Fig. 1).

Our goal is to allow real-time rendering for scenes captured with multiple photos, and create the representations with optimization times as fast as the most efficient previous methods for typical real scenes. Recent methods achieve fast training [Fridovich-Kell

#### **Volume Splatting**

• 3D scene is represented as 3D Gaussians(blobs).



#### How to Render? Sort/Splat/Blend

- Sort: based on depth
- Splat: compute the shape of the Gaussian after projection  $(3D \rightarrow 2D)$
- Blend: alpha composite the gaussians (in 2D / screen space)

$$c = \sum_{i \in \mathcal{N}} c_i w_i \prod_{j=1}^{i-1} (1-w_j)$$



$$p = \begin{bmatrix} p_x \\ p_y \\ p_z \end{bmatrix} \quad \Sigma = \begin{bmatrix} \sigma_x & \sigma_{xy} & \sigma_{xz} \\ & \sigma_y & \sigma_{yz} \\ & & \sigma_z \end{bmatrix} \qquad o \qquad SH$$

position shape opacity color

### Splat

- Positions (3D→2D)
  - Just apply the classic matrices from graphics:  $x' = P \bullet V \bullet x$
  - Perspective effect
    - divide by z-component of homogeneous coordinate
- Covariance  $(3D \rightarrow 2D)$ 
  - World to view transformation is linear (Rotation+Translation)
  - Linear transformations apply a simple change to covariance matrix
  - Perspective effect

$$\label{eq:sigma_def} \begin{split} \textbf{L} \quad \boldsymbol{\Sigma}_{2d} = \textbf{J}_{proj} \ \textbf{R} \ \boldsymbol{\Sigma}_{3d} \ \textbf{R}^{T} \ \textbf{J}_{proj}^{T} \end{split}$$

J: Jacobian of the affine approx. of projective matrix
R: 3D Gaussian orientation
∑: Covariance matrix

#### Blend

- How to implement this in a shader? (for serving in browser)
- Alpha Blending:



• For 2 Gaussians:

$$c = c_0 w_0 + c_1 w_1 (1 - w_0)$$

• For 3 Gaussians:

$$c = c_0 w_0 + c_1 w_1 (1 - w_0) + c_2 w_2 (1 - w_1) (1 - w_0)$$

#### Optimization

• How to ensure that Gaussians remain Gaussians?

$$egin{aligned} \mathcal{L}(oldsymbol{ heta}) &= \mathbb{E}_{\mathbf{C} \sim \{\mathbf{C}_i\}} \ \mathbb{E}_{\mathbf{r} \sim \mathbf{C}} \ ||\mathbf{C}(\mathbf{r};oldsymbol{ heta}) - \mathbf{C}_i(\mathbf{r})||_2^2 \ oldsymbol{ heta} &= \{(\mathbf{p}_k, o_k, \mathbf{\Sigma}_k, \mathbf{h}_k)\} \end{aligned}$$

- Covariance matrices ought to be symmetric and positive semi-definite
  - all eigenvalues should be positive (or zero)
- Solution: re-parameterize the problem

$$\begin{split} \boldsymbol{\Sigma} &= \mathbf{R} \mathbf{S} \mathbf{S}^{\top} \mathbf{R}^{\top} \\ \mathbf{R} &= \text{quat2mat}(\mathbf{q}), \quad \|\mathbf{q}\|_2 = 1 \\ \mathbf{S} &= \text{diag}([s_x, s_y, s_z]) \end{split}$$

#### Can we just train a 3DGS now?

• No... you need to implement the "adaptive density control" heuristics





Kerbl, Kopanas et al. "3D Gaussian Splatting for Real-Time Radiance Field Rendering" SIGGRAPH 2023

heuristics

#### What are these heuristics?

- Densification
  - increase the number of points whenever under-reconstruction is detected, as measured by positional gradients
- Pruning
  - of very large Gaussians
- Opacity reset (every N iterations)
  - to remove floaters in the representation
- Require a good initialization (e.g. SFM)
  - as all the above are local heuristics





Kerbl, Kopanas et al. "3D Gaussian Splatting for Real-Time Radiance Field Rendering" SIGGRAPH 2023

#### Acknowledgements

• Slides modified from:

- Fundamentals of (neural) Inverse Rendering (Lecture talk at ICVSS '24)
   Assoc. Prof. Andrea Tagliasacchi
- Cornell Uni. CS5670: Computer Vision

## Questions?