3D Reconstruction Using Gaussian Splatting

Daniel Perazzo

Orientadores: Luiz Velho, João Paulo Lima and Tiago Novello

75 anos impa
Instituto de Matemática Pura e Aplicada

Universidade Federal de Pernambuco

voxar labs

Centro de Informática UFPE

1. Introduction

1. Introduction

- 3D Reconstruction is a important problem!
- One of the most important in Computer Vision

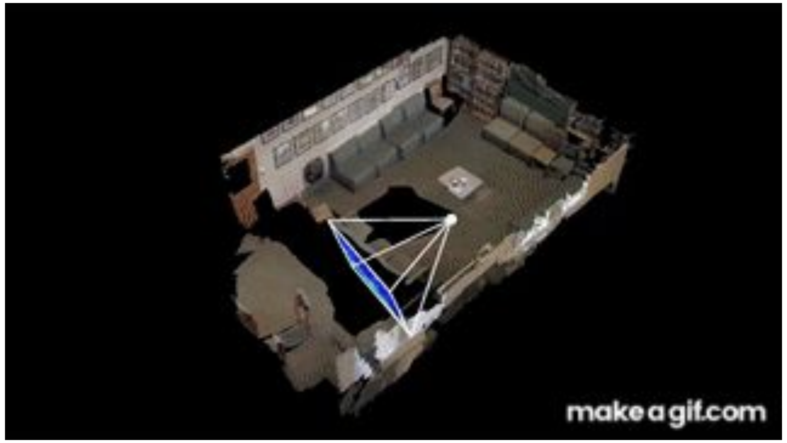

3D Reconstruction

1. Introduction :: Applications

- Intersections with other fields
- Cool Applications of 3D Reconstruction
  - Entertainment
  - Archeological Preservation
  - Medical Imaging
  - And more!

1. Introduction :: Basics

- Represent images as matrices
    - 3 channel matrices (RGB)
- We specify camera as matrices
    - Rotation and Translation for extrinsics
    - Intrinsic matrices

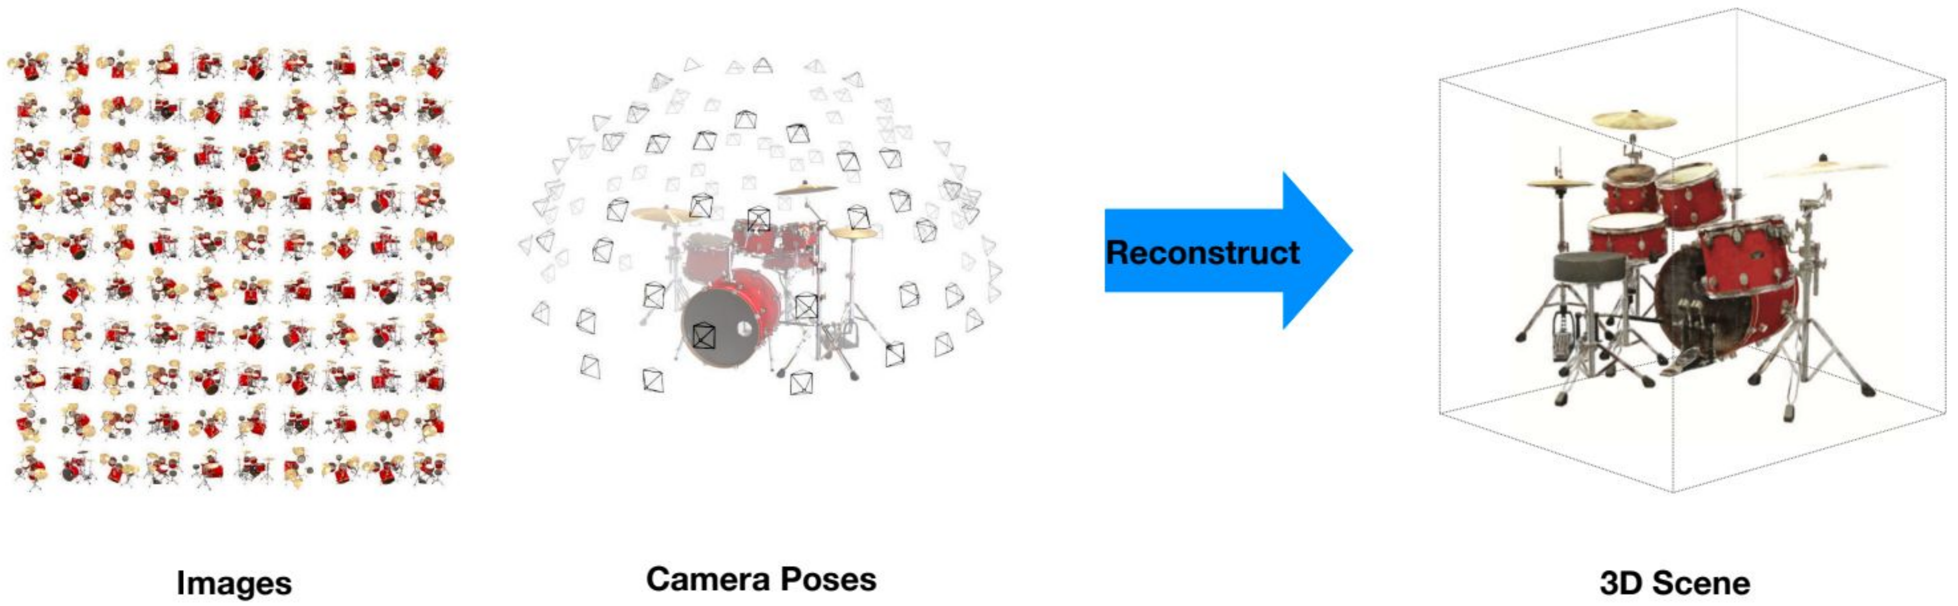

1. Introduction :: Basics

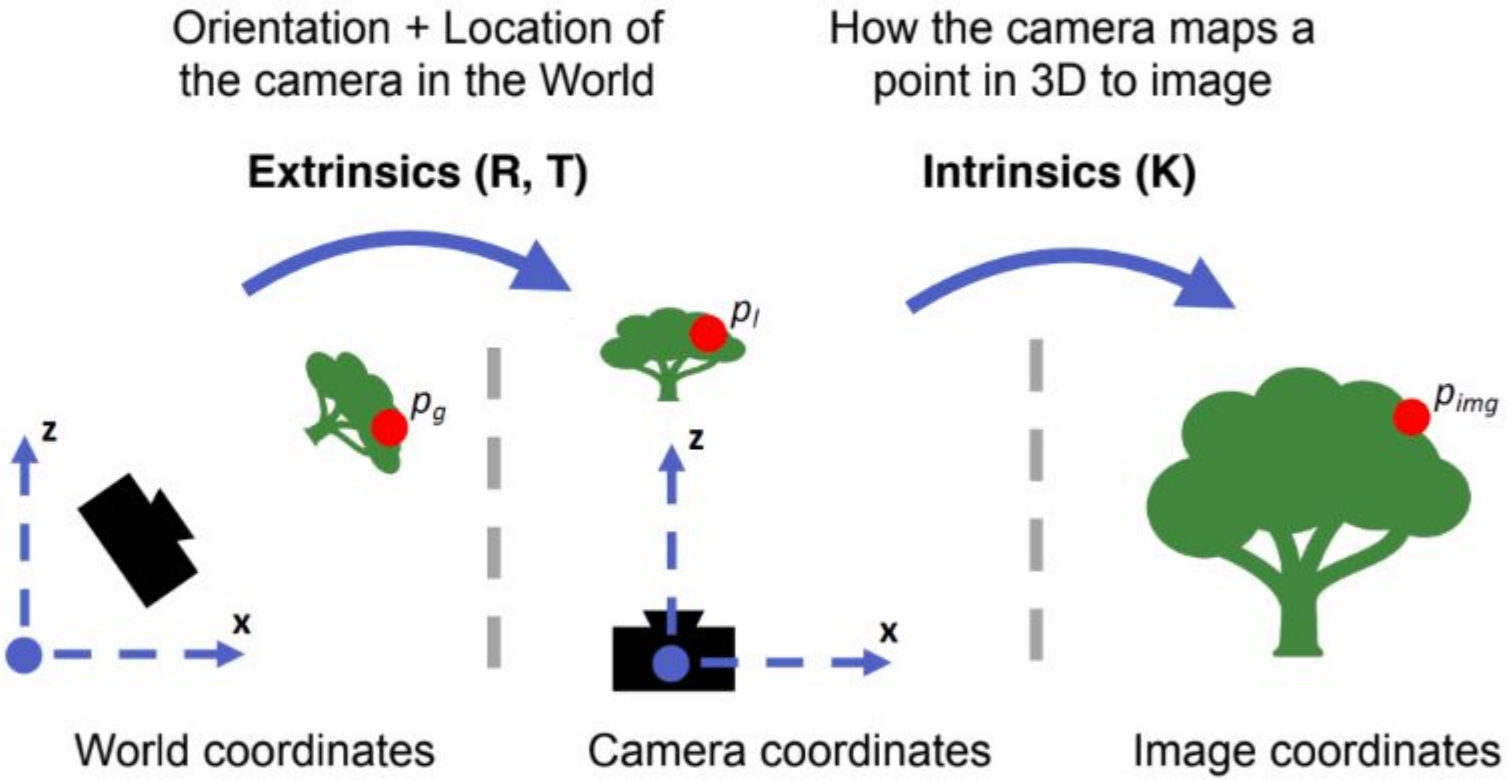

Figure credit: Peter Hedman

1. Introduction :: Examples

- What is the format of the object?
  - Point Cloud (RGB)
  - Mesh
  - Etc.

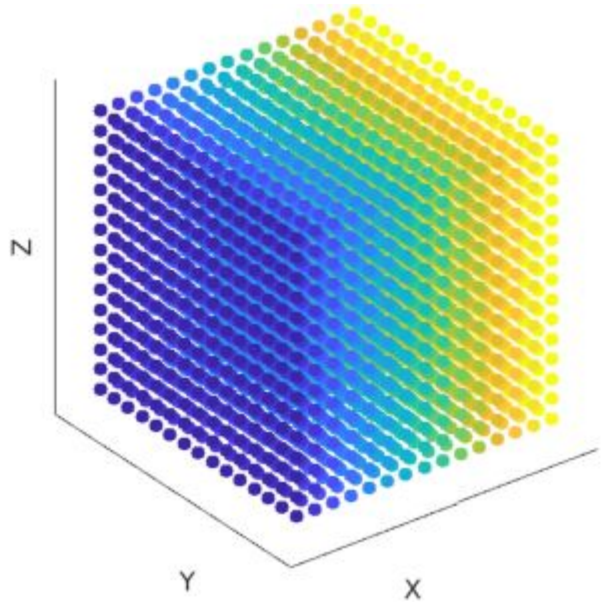

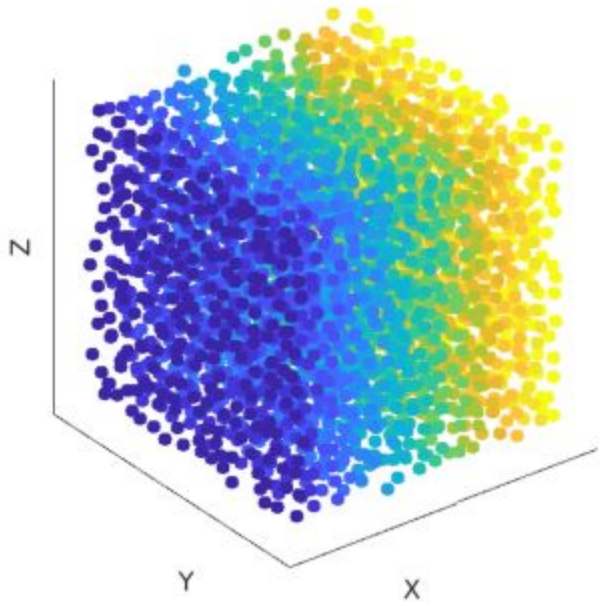

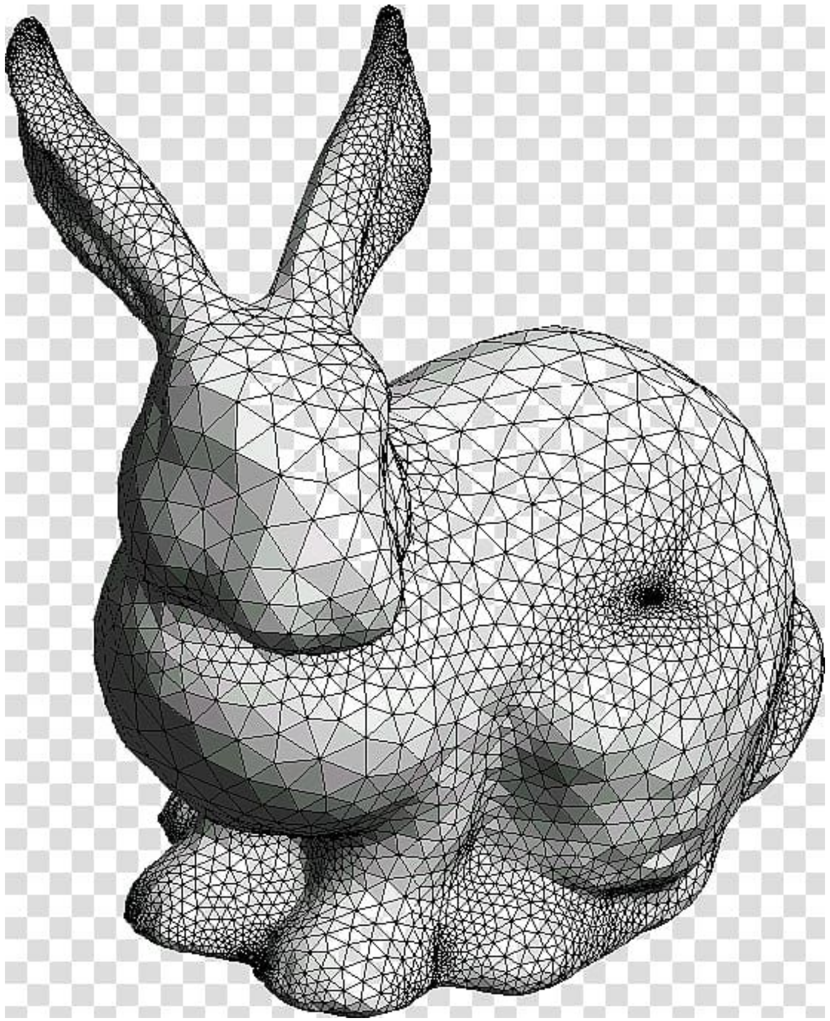

1. Introduction :: Examples

- What is the format of the object?
  - Point Cloud (RGB)
  - Mesh
  - Etc.

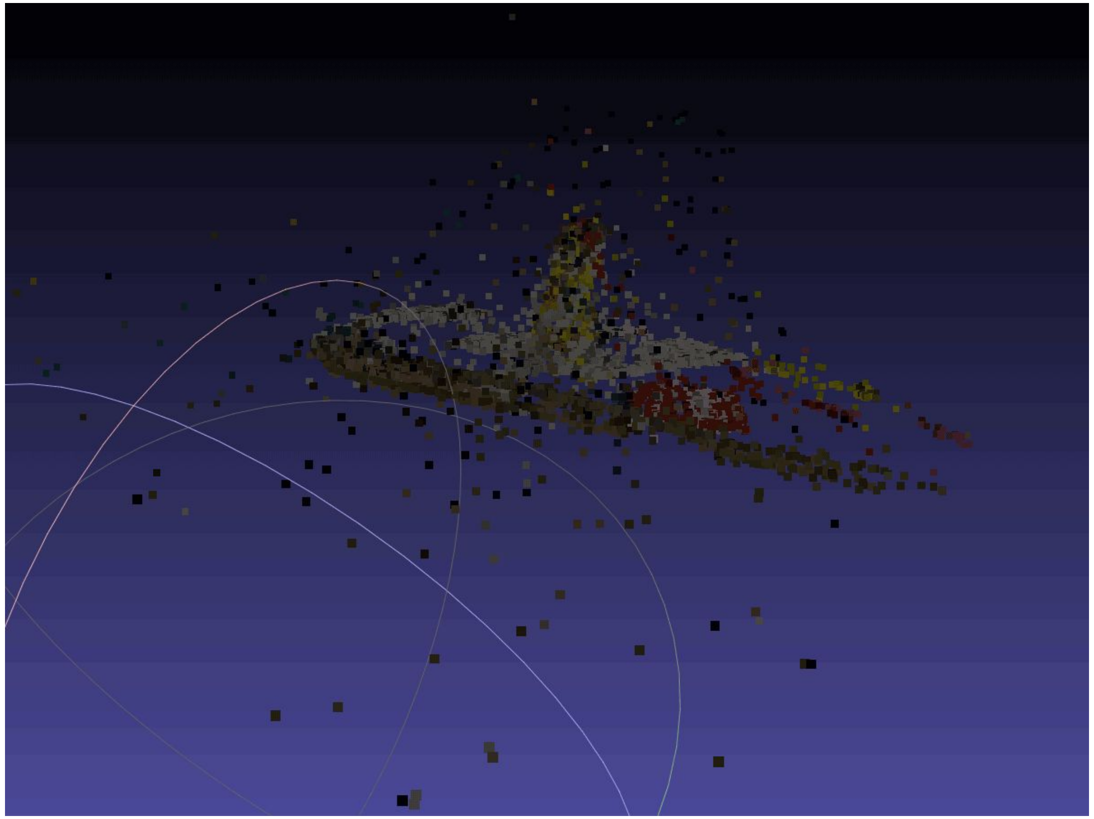

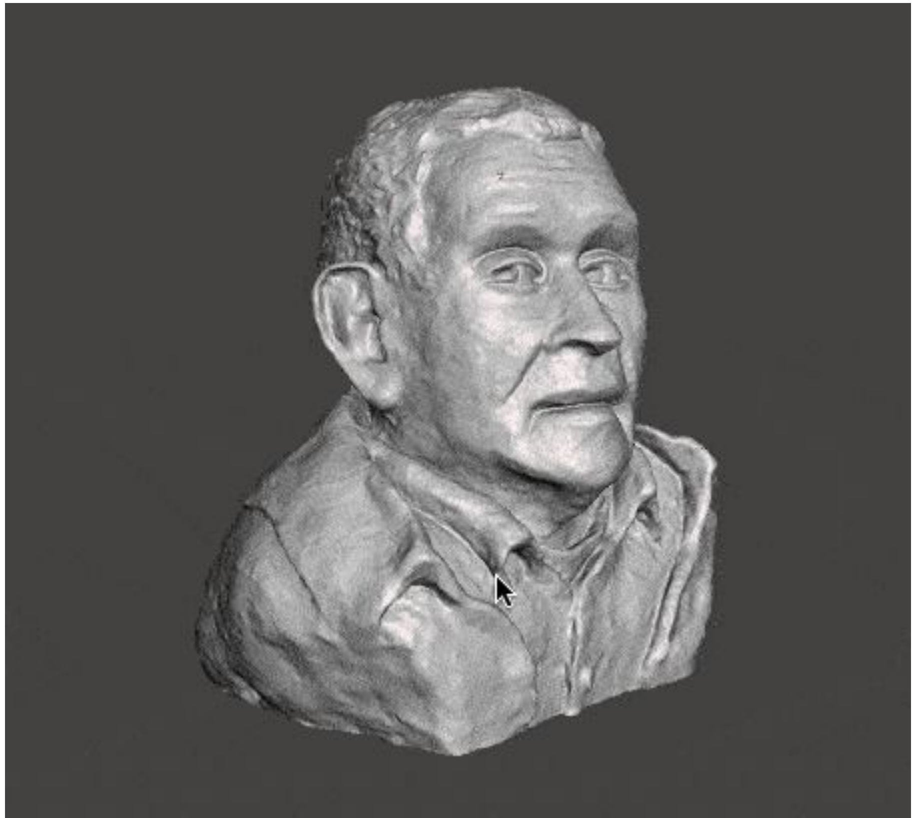

1. Introduction :: Classic Techniques

- There are some techniques for extraction
- Most famous is probably COLMAP

1. Introduction :: Modern Methods

- However, many new techniques for 3D Reconstruction now exist!
- New representations!
    - NeRFs
    - Gaussian Splatting

1. Introduction :: Schedule

- Give a presentation of this methods
- Give an intuition on how they work
- And what current works are being done
- What kind of research we are doing (Visgraf)

2. Neural Radiance Fields

2. Neural Radiance Fields

- NeRFs were the first-version of this current wave
- First, developed for “view-synthesis”
- Represents the object as a neural-network
- Uses volume rendering to extract novel views
- Trained on a set of images with camera parameters
    - But what are all this?

Mildenhall, Ben, et al. "Nerf: Representing scenes as neural radiance fields for view synthesis." European conference on computer vision. Springer, Cham, 2020.

2. Neural Radiance Fields :: Volume Rendering

- A form to create images from “clouds”
  - Used in games, movies, VFX, etc.
  - Create smoke effects
  - Solves a integral
  - Differentiable
    - Unlike other types of rendering

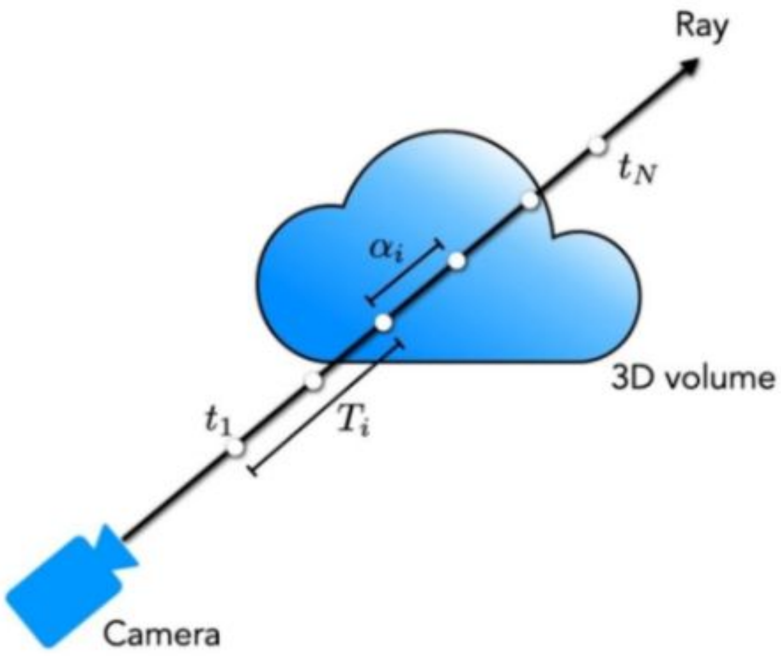

Max, Nelson, and Min Chen. Local and global illumination in the volume rendering integral. No. UCRL-PROC-216495. Lawrence Livermore National Lab.(LLNL), Livermore, CA (United States), 2005.

2. Neural Radiance Fields :: Volume Rendering

- The equation that gives the color for a ray is:

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

- Where:

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

- We can analyze this integral numerically

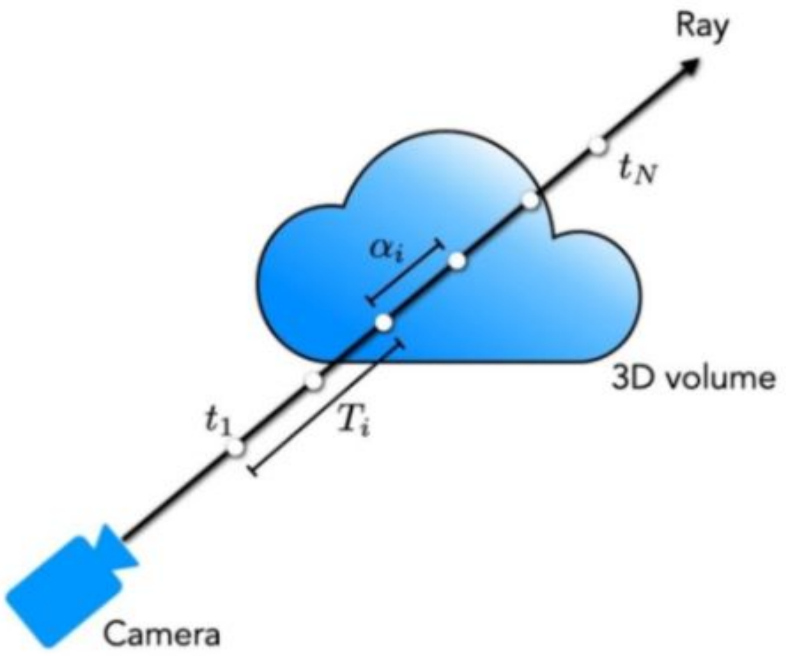

Max, Nelson, and Min Chen. Local and global illumination in the volume rendering integral. No. UCRL-PROC-216495. Lawrence Livermore National Lab.(LLNL), Livermore, CA (United States), 2005.

2. Neural Radiance Fields :: Volume Rendering

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

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

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

colors

weights

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:

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

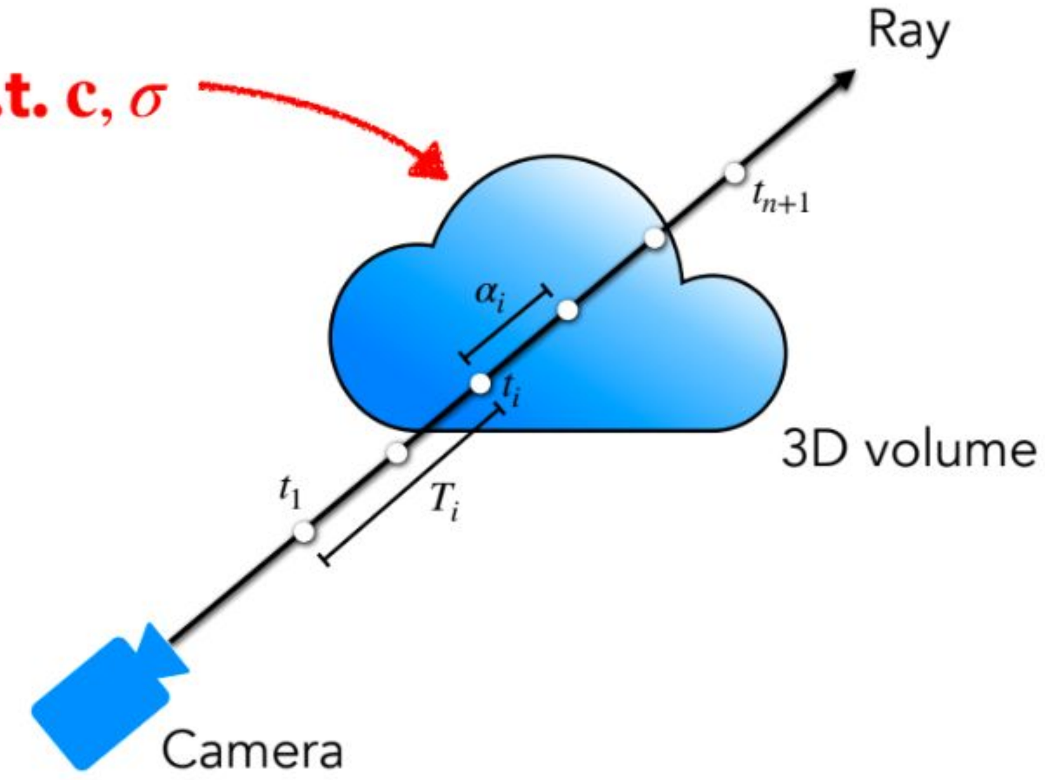

3. Neural Radiance Fields :: Pipeline

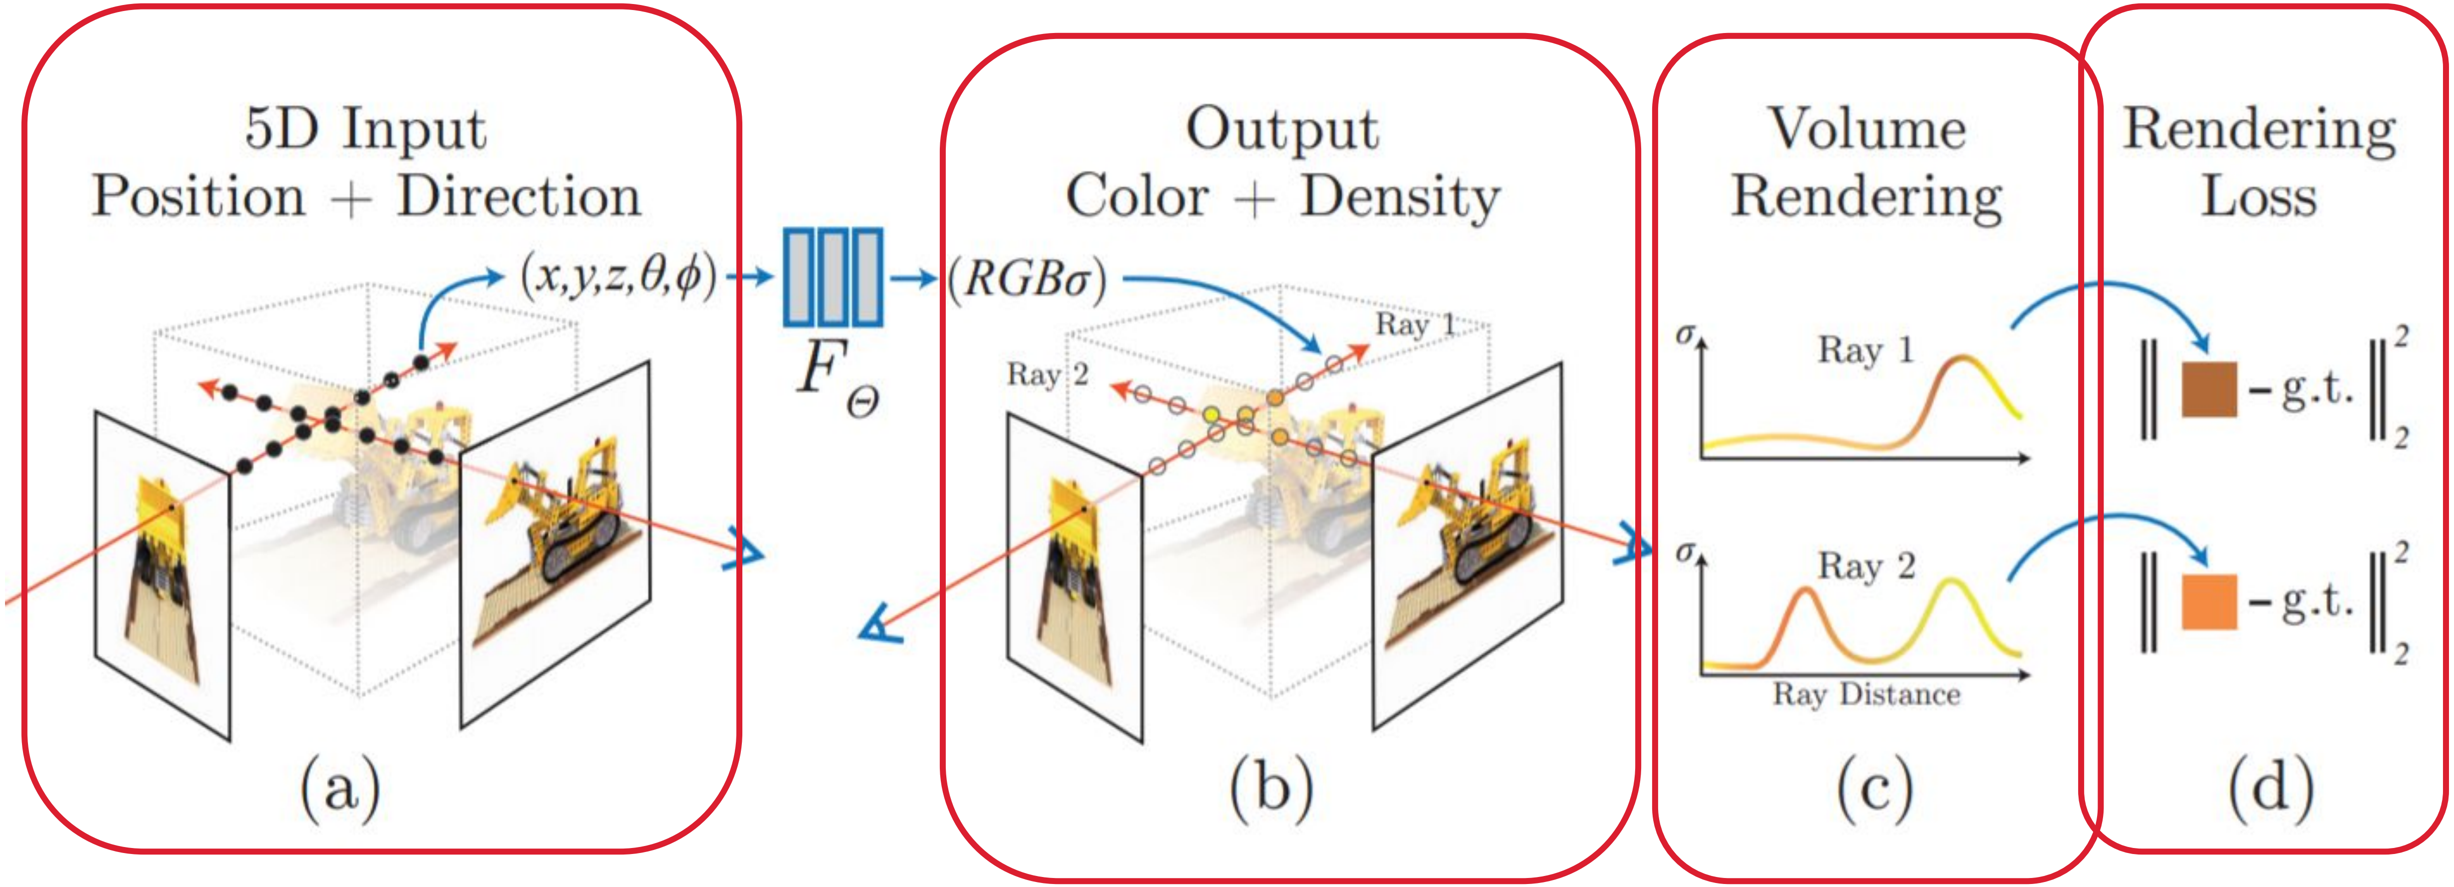

2. Neural Radiance Fields :: Results

- Creates many photorealistic effects
- Improved on many subsequent works
    - Instant-NGP
    - Plenoxels

Max, Nelson, and Min Chen. Local and global illumination in the volume rendering integral. No. UCRL-PROC-216495. Lawrence Livermore National Lab.(LLNL), Livermore, CA (United States), 2005.

2. Neural Radiance Fields :: Results

- Creates many photorealistic effects
- Improved on many subsequent works
  - Instant-NGP
  - Plenoxels
  - Directvoxgo++...

Perazzo, Daniel, et al. "DirectVoxGO++: Grid-based fast object reconstruction using radiance fields." Computers & Graphics 114 (2023): 96-104.

2. Neural Radiance Fields :: Drawbacks

- Originally, took 9 hours to train a single scene
  - Time was reduced to 30 minutes
- However, at the core, their method is expensive

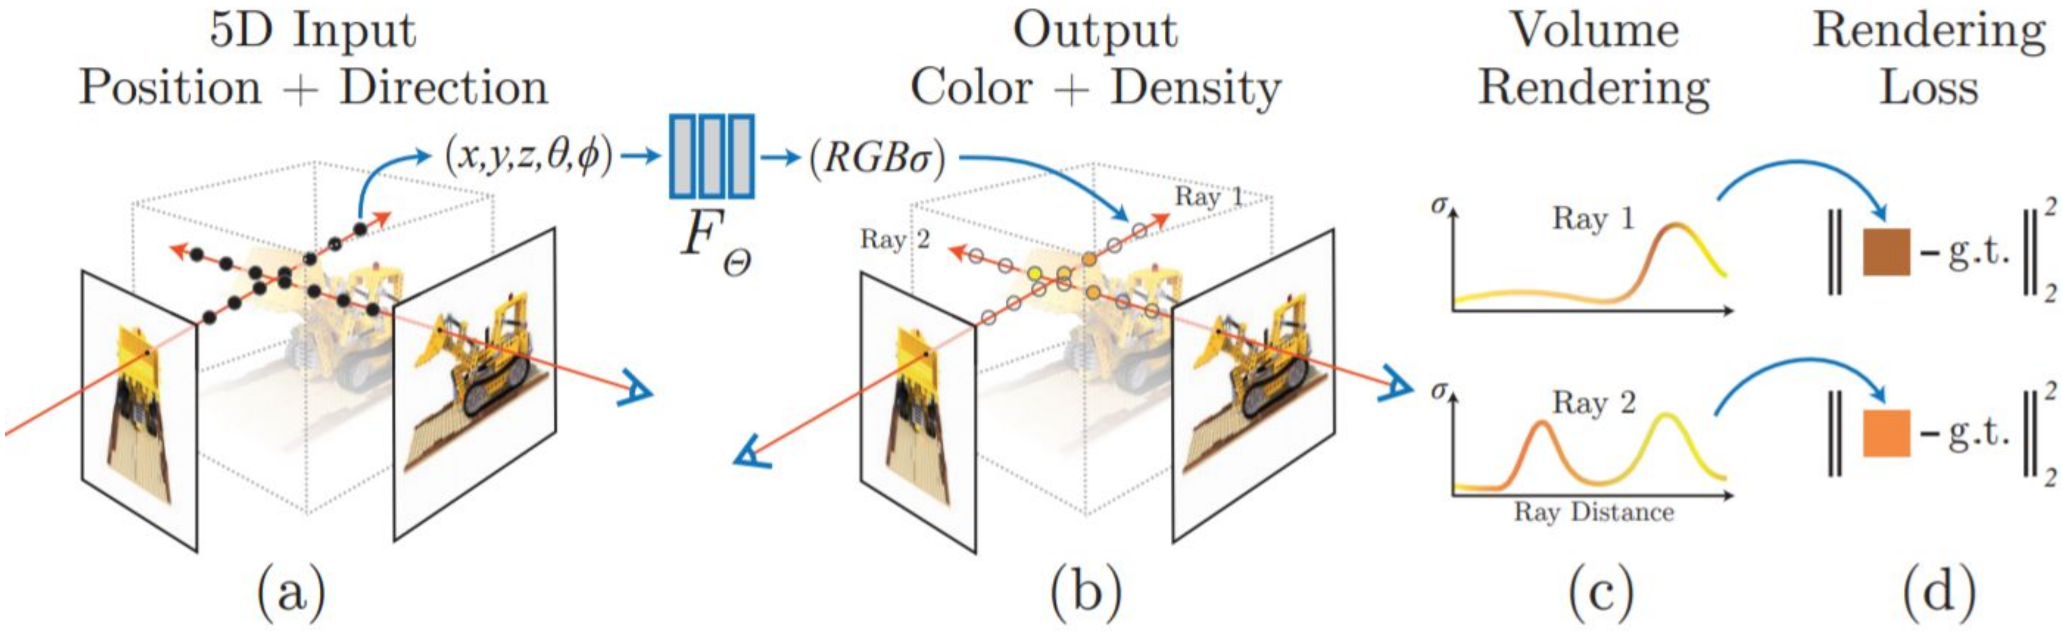

2. Neural Radiance Fields :: Drawbacks

- Originally, took 9 hours to train a single scene
  - Time was reduced to 30 minutes
- However, at the core, their method is expensive
  - Needs other ideas…

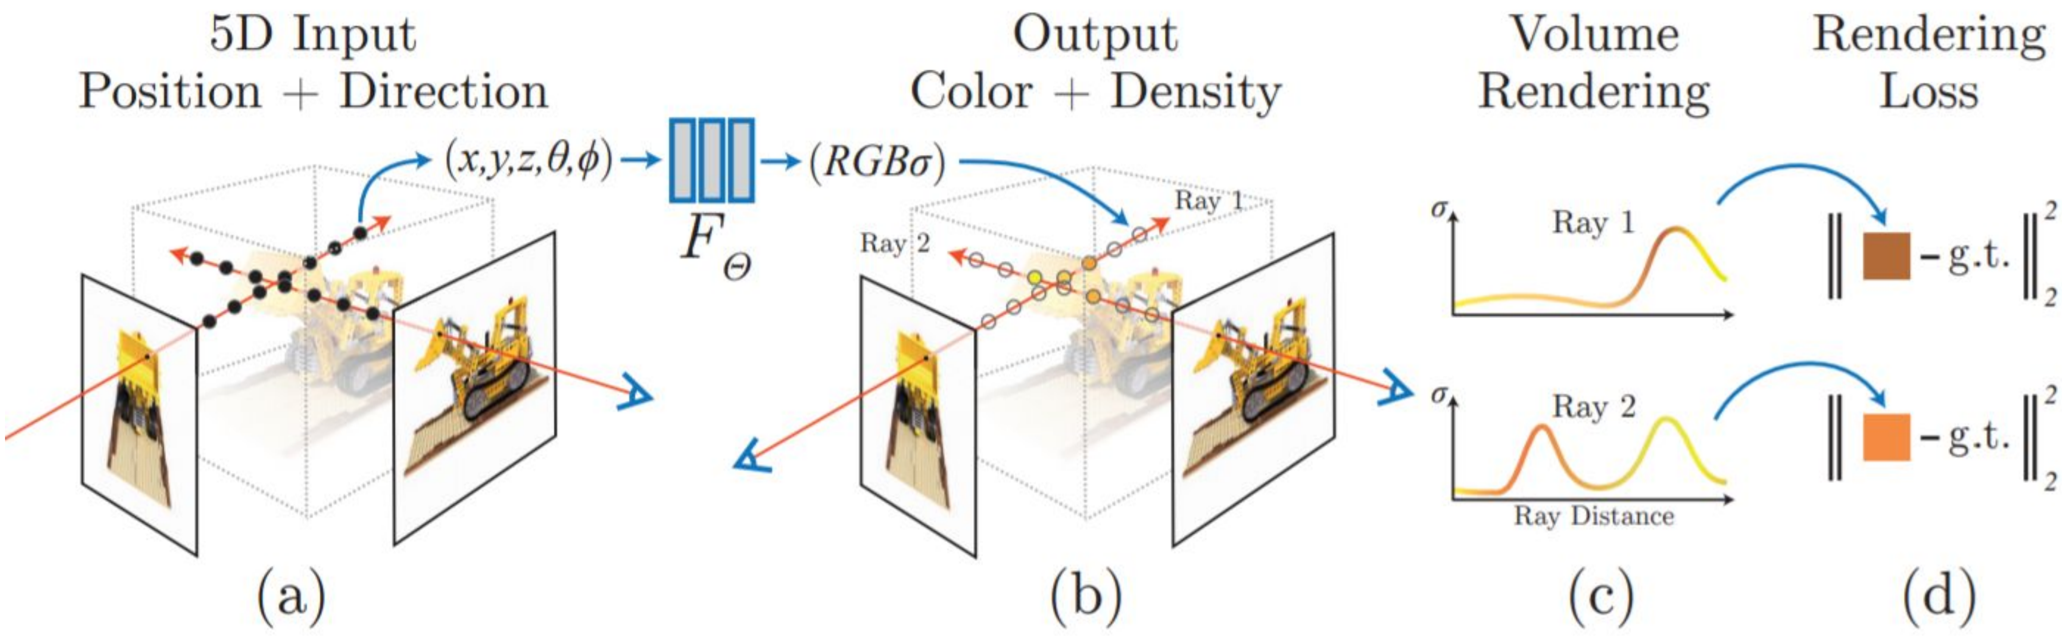

4. Gaussian Splatting

- One of the hottest papers in computer vision
    - Solves the main NeRF drawback!
        - Speed

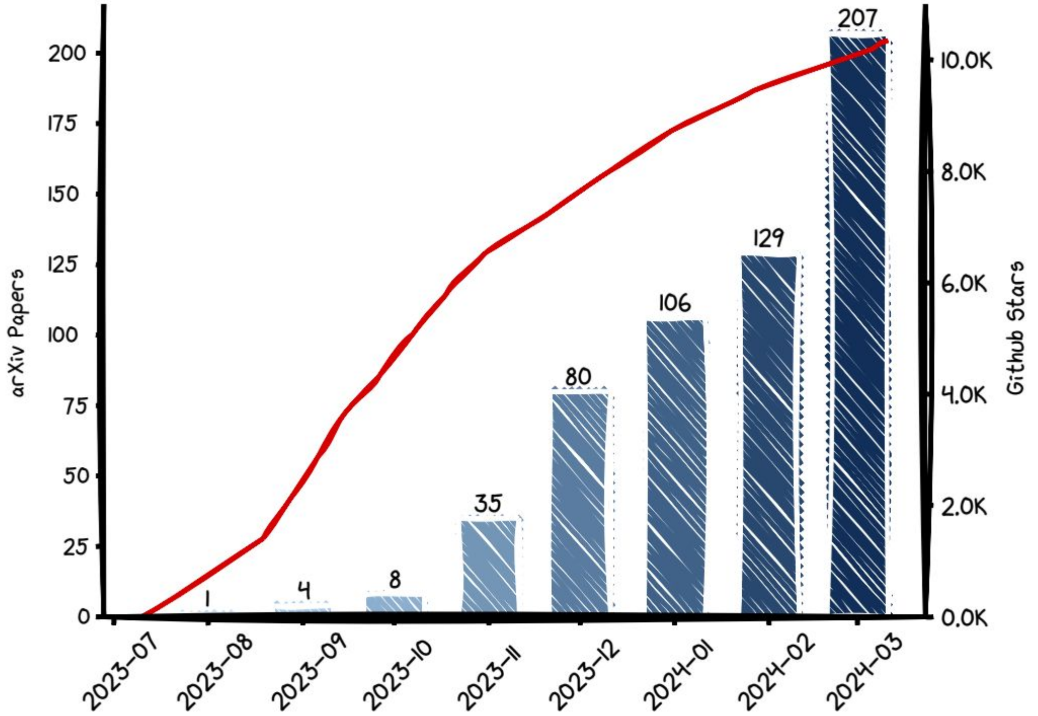

Kerbl, Bernhard, et al. "3d gaussian splatting for real-time radiance field rendering." ACM Transactions on Graphics 42.4 (2023): 1-14.

4. Gaussian Splatting :: Basics

- Actually a "Classical" paper in disguise
  - No Deep Learning
- Implements classical ideas into 3D Reconstruction

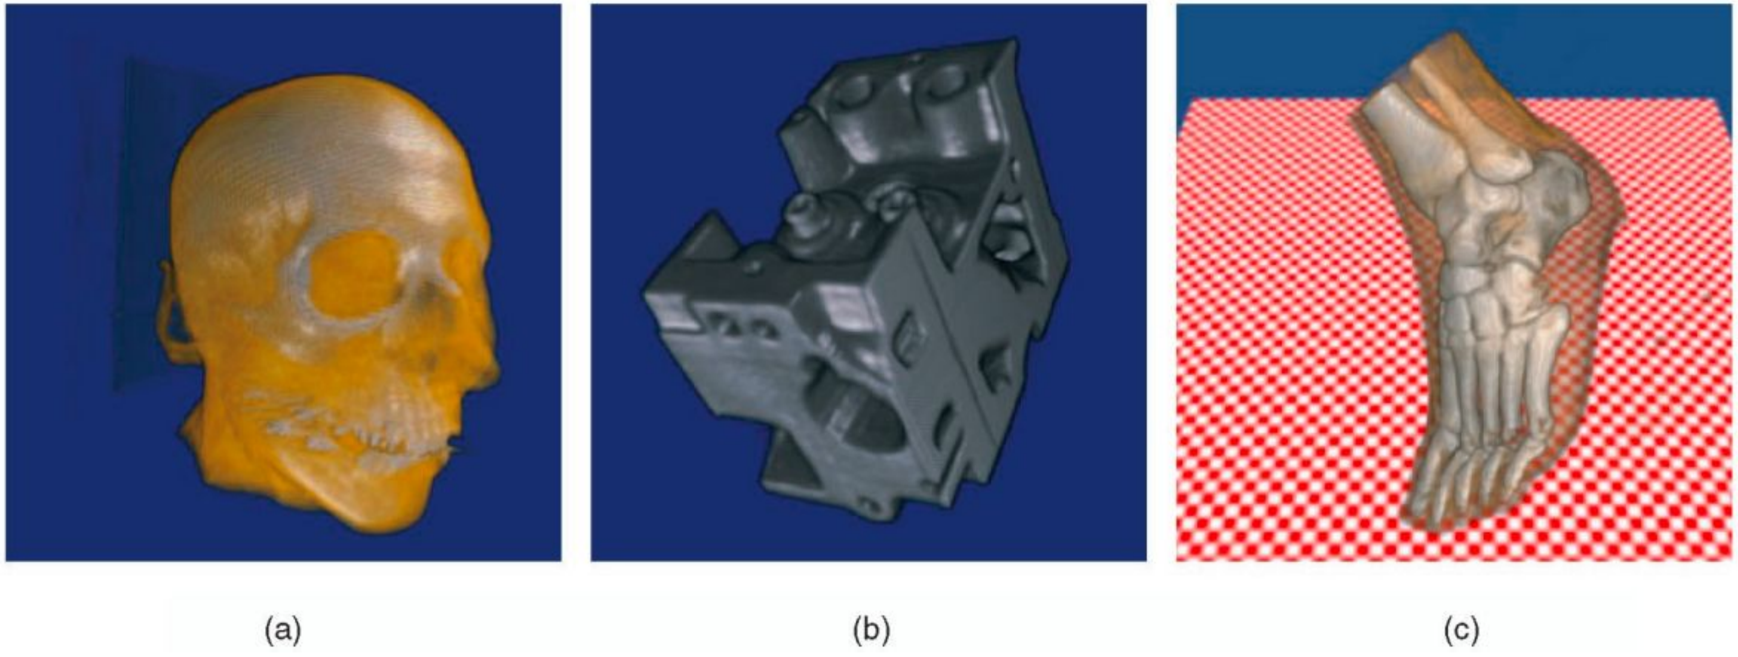

Zwicker, Matthias, et al. "EWA splatting." IEEE Transactions on Visualization and Computer Graphics 8.3 (2002): 223-238.

4. Gaussian Splatting :: Basics

- Impressive results!
- Really fast
  - Absurdly fast!
- But, how?

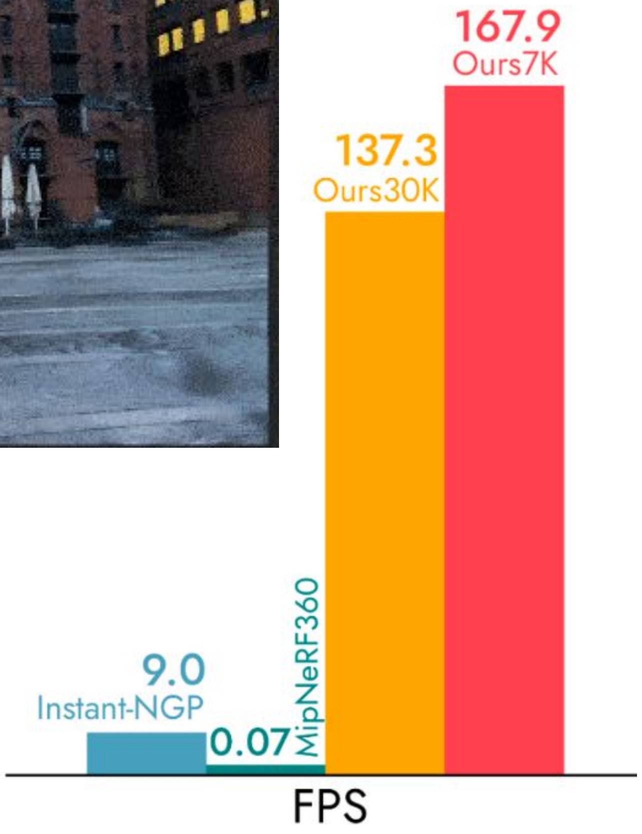

Kerbl, Bernhard, et al. "3d gaussian splatting for real-time radiance field rendering." ACM Transactions on Graphics 42.4 (2023): 1-14.

4. Gaussian Splatting :: Basics

- Impressive results!
- Really fast
  - Absurdly fast!
- But, how?

Kerbl, Bernhard, et al. "3d gaussian splatting for real-time radiance field rendering." ACM Transactions on Graphics 42.4 (2023): 1-14.

4. Gaussian Splatting :: Workings

- Basic pipeline for 3DGS
- We will detail each of these

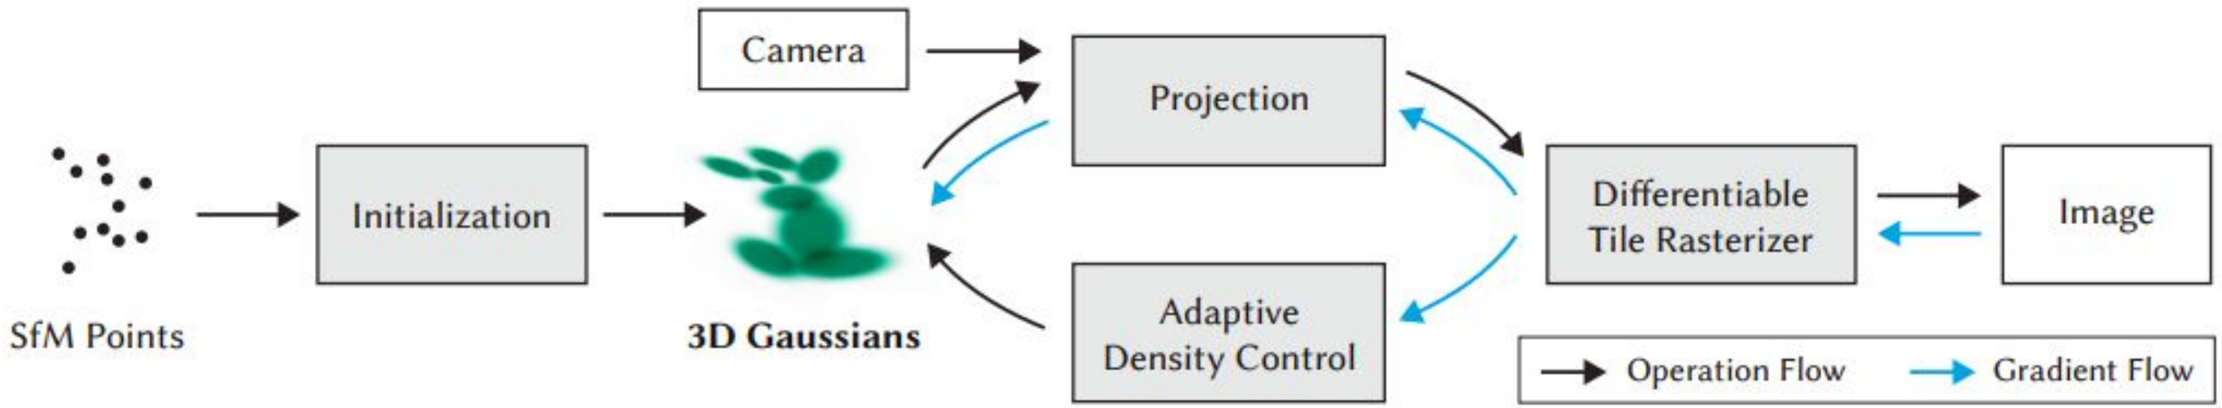

Fig. 2. Optimization starts with the sparse SfM point cloud and creates a set of 3D Gaussians. We then optimize and adaptively control the density of this set of Gaussians. During optimization we use our fast tile-based renderer, allowing competitive training times compared to SOTA fast radiance field methods. Once trained, our renderer allows real-time navigation for a wide variety of scenes.

4. Gaussian Splatting :: Workings

- Imagine 3D Gaussians floating in space
  - Each with their “density”
  - But how can we render (take a picture)

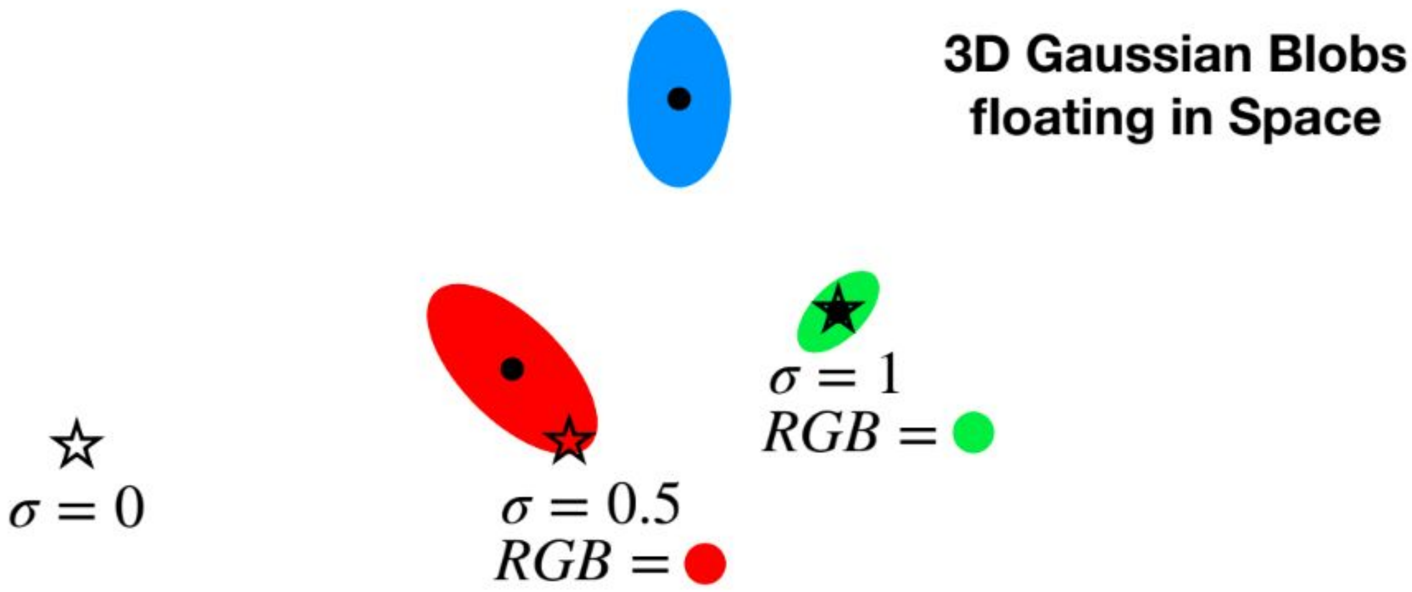

4. Gaussian Splatting :: Workings

- We can perform rendering like in NeRFs
  - Perform Volume Rendering

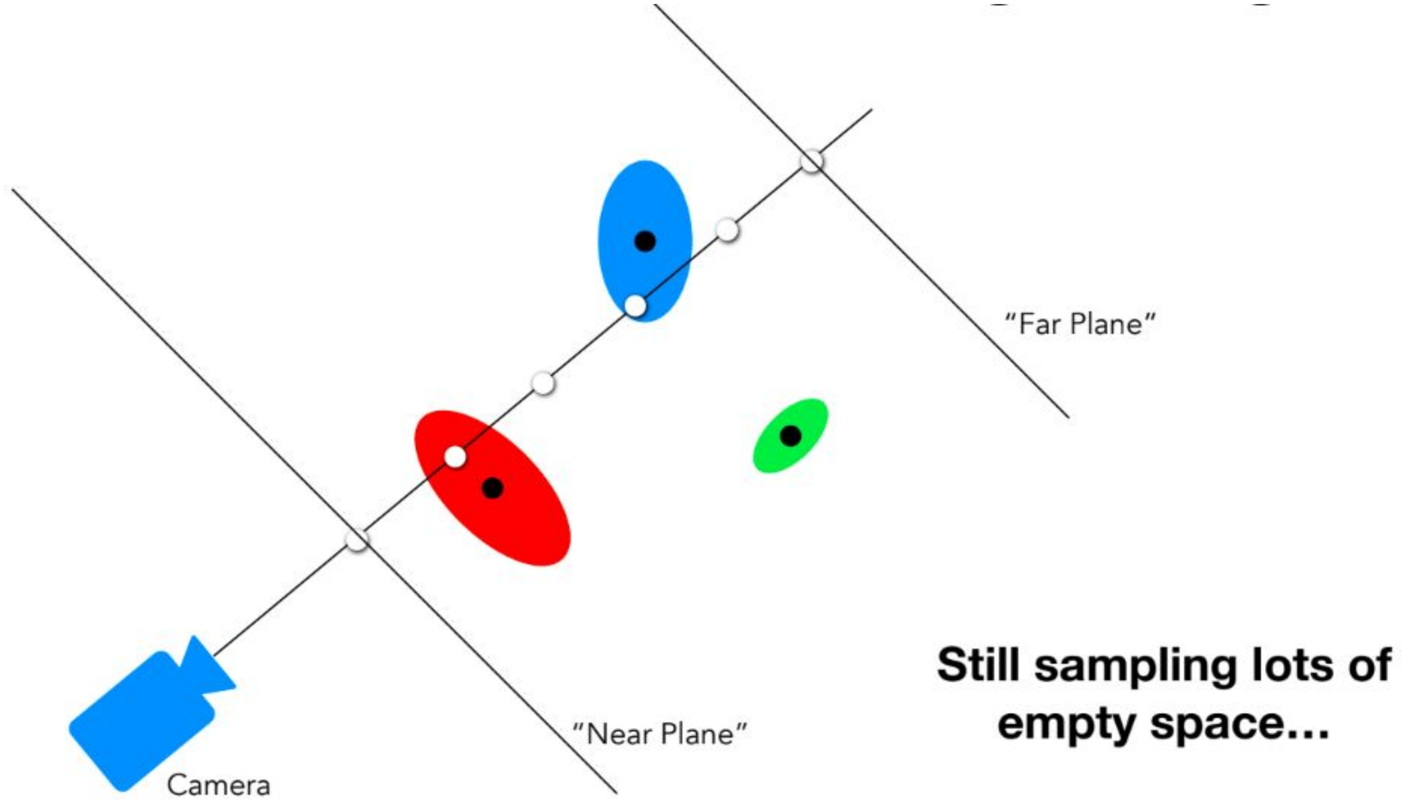

4. Gaussian Splatting :: Workings

- However, we can skip this sampling
  - Compute a integral analytically
  - The idea is to “Splat the gaussians” on the camera
- Project the Gaussian on the scene and compose then
  - Rasterization
- Much faster than sampling

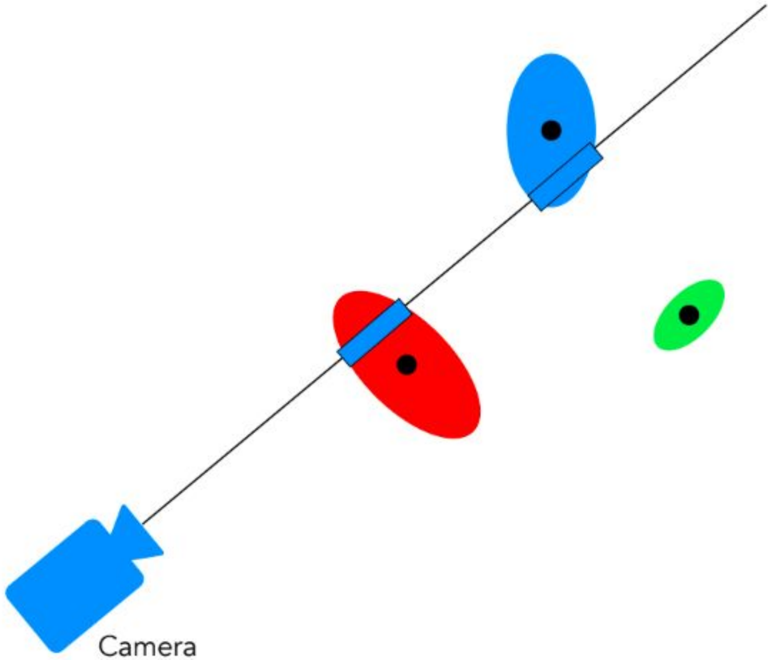

4. Gaussian Splatting :: Workings

- Once the Gaussians are already splatted on the camera
  - Order the gaussians based on the distance
  - Perform a composition on the color

$$C(p) = \sum_{i \in N} c_i f_i^{2D}(p) \prod_{j=1}^{i-1} (1 - f_j^{2D}(p))$$

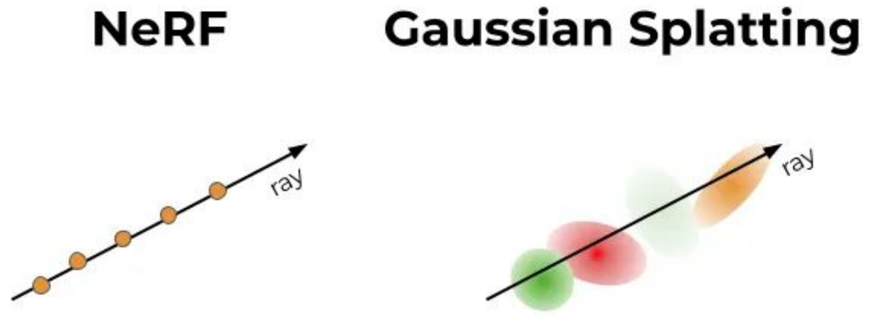

- Equation really similar to NeRF!
- Implemented rasterizer is fast!
  - Implemented directly on the GPU
  - Uses many fast methods from CUDA

4. Gaussian Splatting :: Representation

- And how do we represent the Gaussians?
  - Not really Gaussians…
- Parameterize the Gaussians
  - Mean
  - Covariance
  - Opacity
  - Color
- We can diagonalize covariance:

$$\Sigma = RSS^T R^T$$

- Optimize Rotation and Scales

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

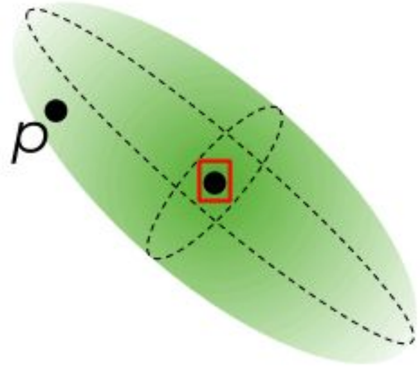

- How do we parameterize the color?
    - View-dependent
- Spherical harmonics
    - Parameters for the Gaussian

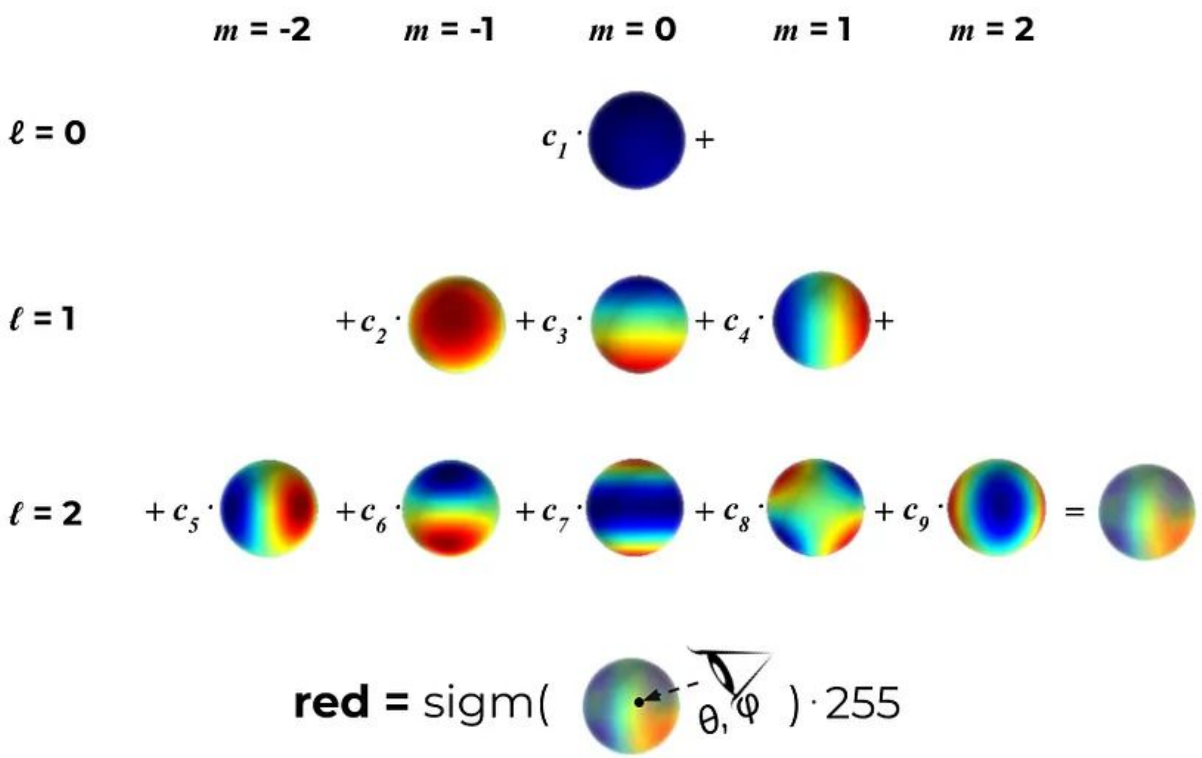

4. Gaussian Splatting :: Optimization

- Optimization with traditional Deep Learning optimizers (Adam)
- However, there is a problem,
  - Subject to local minima
- So, we start with a Sparse point cloud!!
  - COLMAP

4. Gaussian Splatting :: Optimization

- Optimization with traditional Deep Learning optimizers (Adam)
- However, there is a problem,
  - Subject to local minima
- So, we start with a Sparse point cloud!!
  - COLMAP

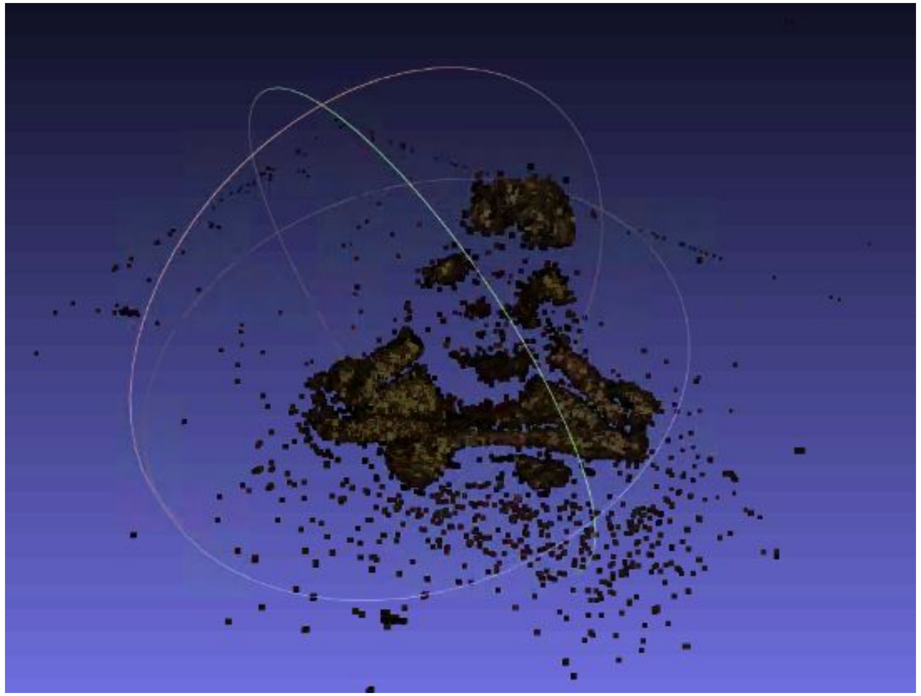

4. Gaussian Splatting :: Optimization

- Optimization with traditional Deep Learning optimizers
- However, there is a problem,
  - Subject to local minima
- So, we start with a Sparse point cloud!!
  - COLMAP

4. Gaussian Splatting :: Optimization

- Finally the authors devised a method to grow the Gaussians
  - Clone and split the gaussians
  - According if the Gaussian is “Under” or “Over” Reconstructed
- Increases the performance of the technique
- Uses L1 and D-SSIM loss for optimization

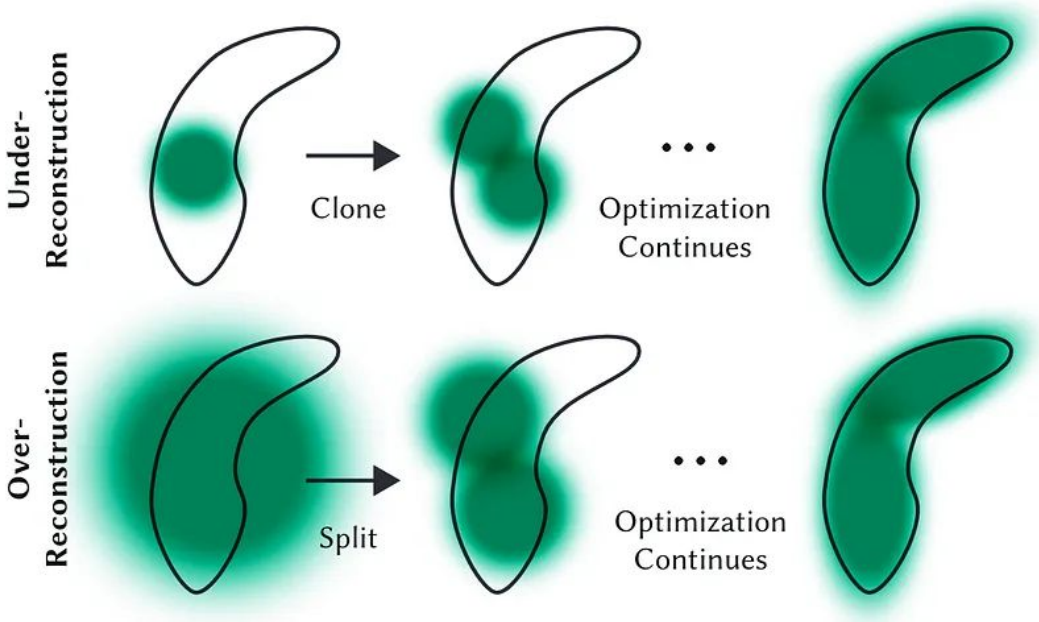

4. Gaussian Splatting :: Results

- Great Results!!
  - Really fast to render

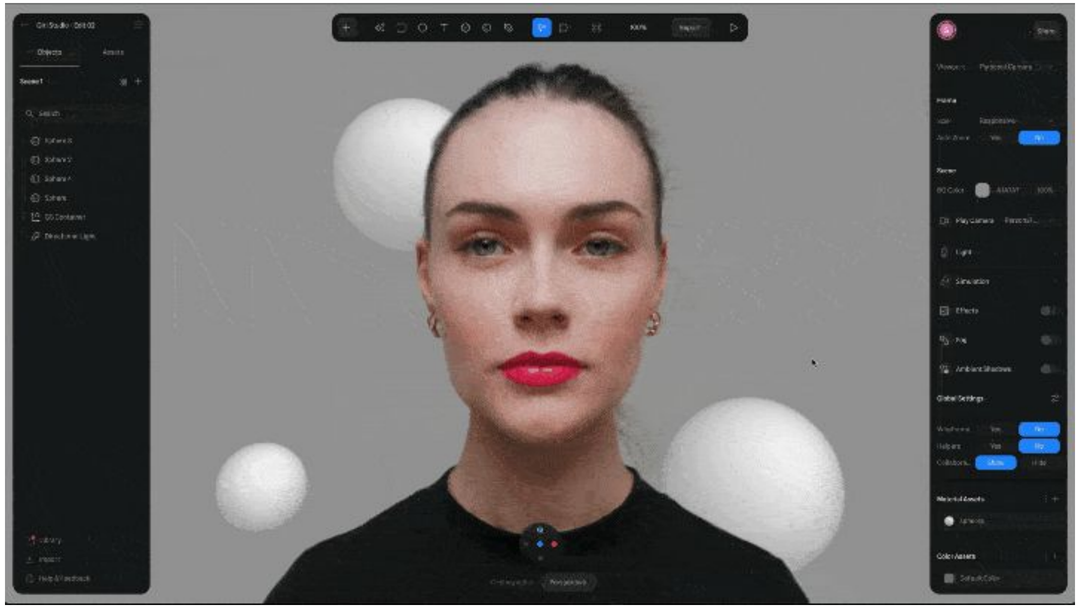

4. Gaussian Splatting :: Drawbacks

- There are many advantages to Gaussian Splatting
- However, there can be some flaws
  - The geometry can be bad (mesh)
  - Not its focus

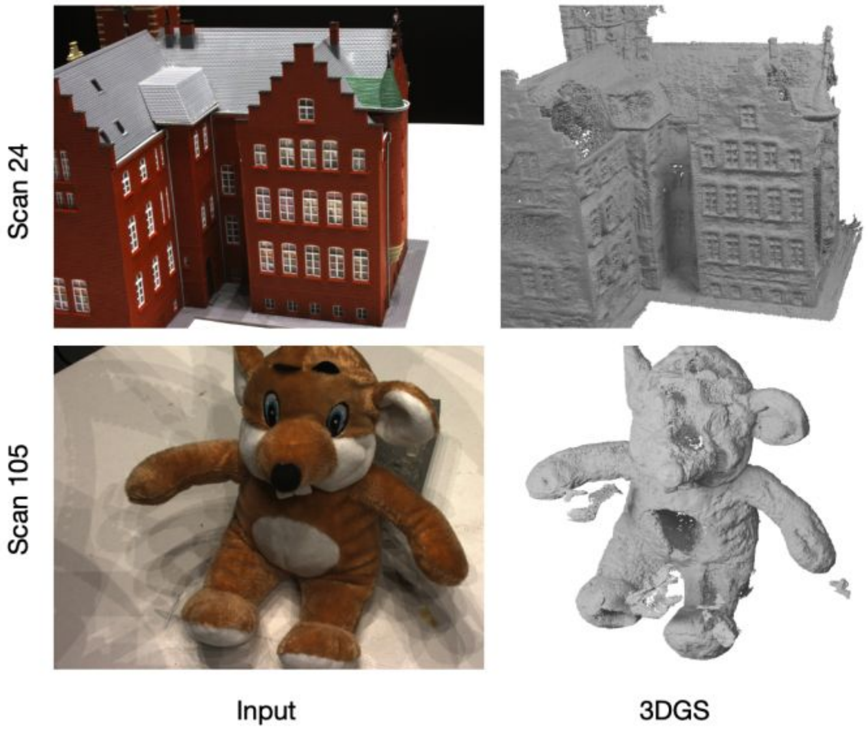

5. Geometry + Gaussian Splatting

5. Geometry + Gaussian Splatting :: Examples

- Techniques which aim to create gaussianswith better geometry
- What kinds of works and ideas do they employ
  - We will focus on two
    - SUGAR
    - 2DGS

5. Geometry + Gaussian Splatting :: SUGAR

- First paper on extracting meshes from Gaussian Splats
    - Adapts the Gaussians so they are more easier to convert
    - Extract a mesh using Poisson Surface Reconstruction
    - Optimizes Gaussians using mesh

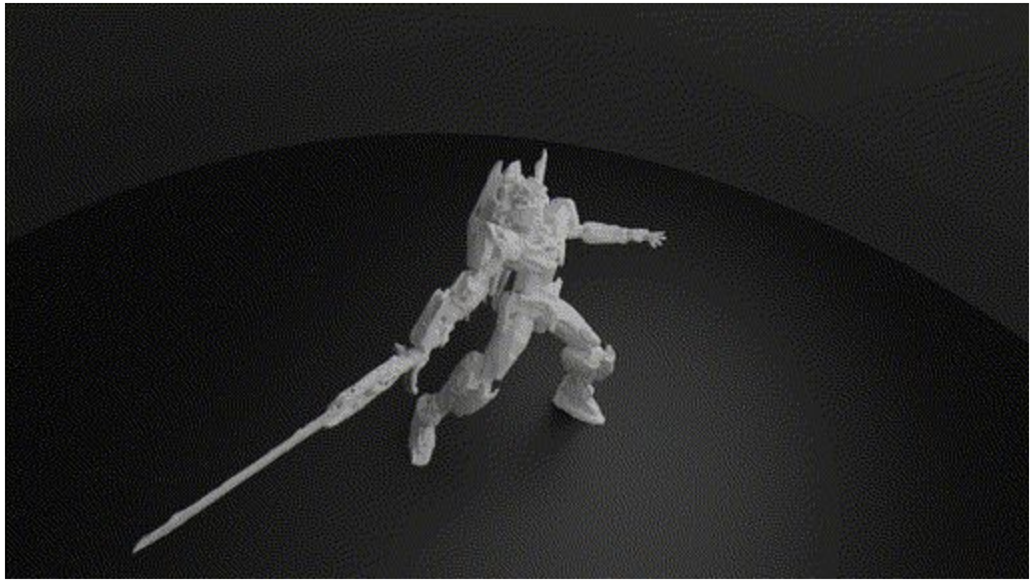

Guédon, Antoine, and Vincent Lepetit. "Sugar: Surface-aligned gaussian splatting for efficient 3d mesh reconstruction and high-quality mesh rendering." arXiv preprint arXiv:2311.12775 (2023).

5. Geometry + Gaussian Splatting :: SUGAR

- But how does it adapts the Gaussians?
  - Flatten them into disks
  - Easier to get the normal

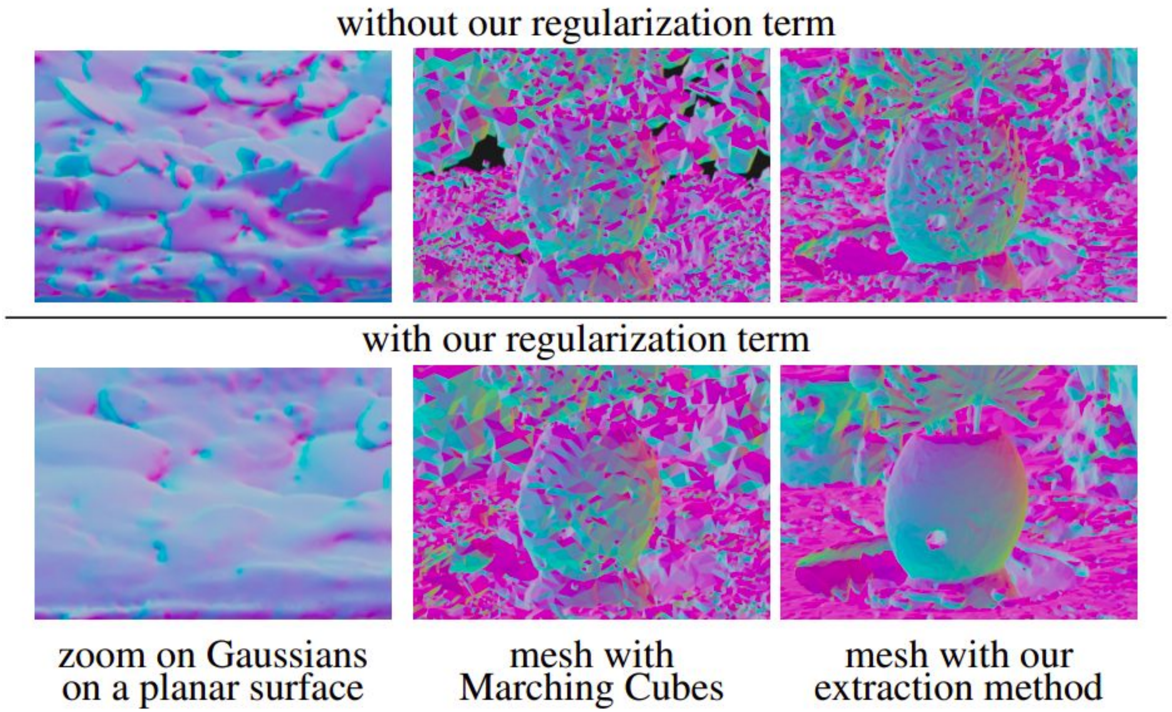

Guédon, Antoine, and Vincent Lepetit. "Sugar: Surface-aligned gaussian splatting for efficient 3d mesh reconstruction and high-quality mesh rendering." arXiv preprint arXiv:2311.12775 (2023).

5. Geometry + Gaussian Splatting :: SUGAR

- However, results are not perfect
- Mesh can be noisy

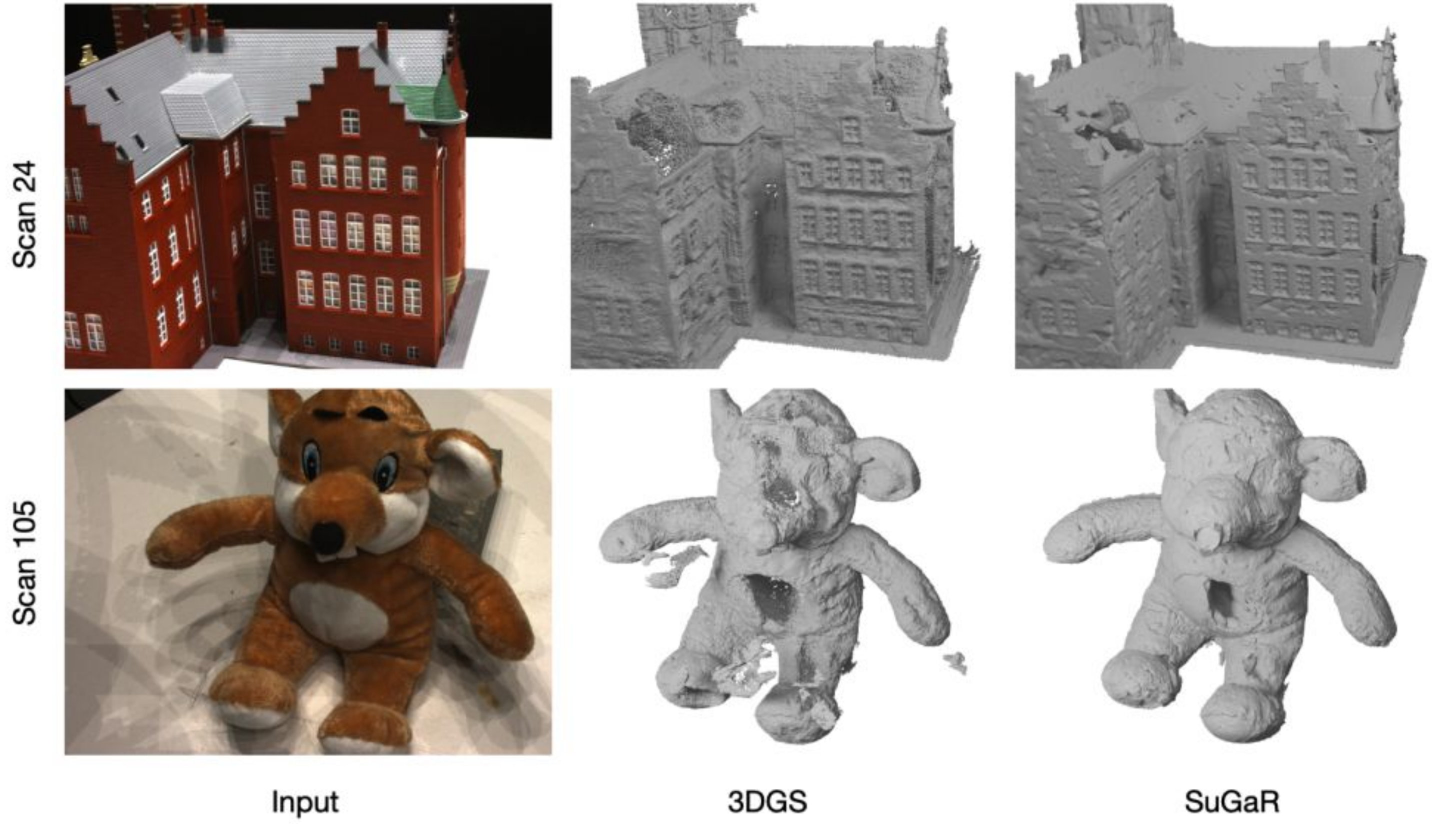

Guédon, Antoine, and Vincent Lepetit. "Sugar: Surface-aligned gaussian splatting for efficient 3d mesh reconstruction and high-quality mesh rendering." arXiv preprint arXiv:2311.12775 (2023).

5. Geometry + Gaussian Splatting :: 2DGS

- Key Idea: Use 2D Gaussians instead of 3D Ones
  - Approximates better the geometry (mean will be on surface)
  - We can use well-defined normals (at least the direction)

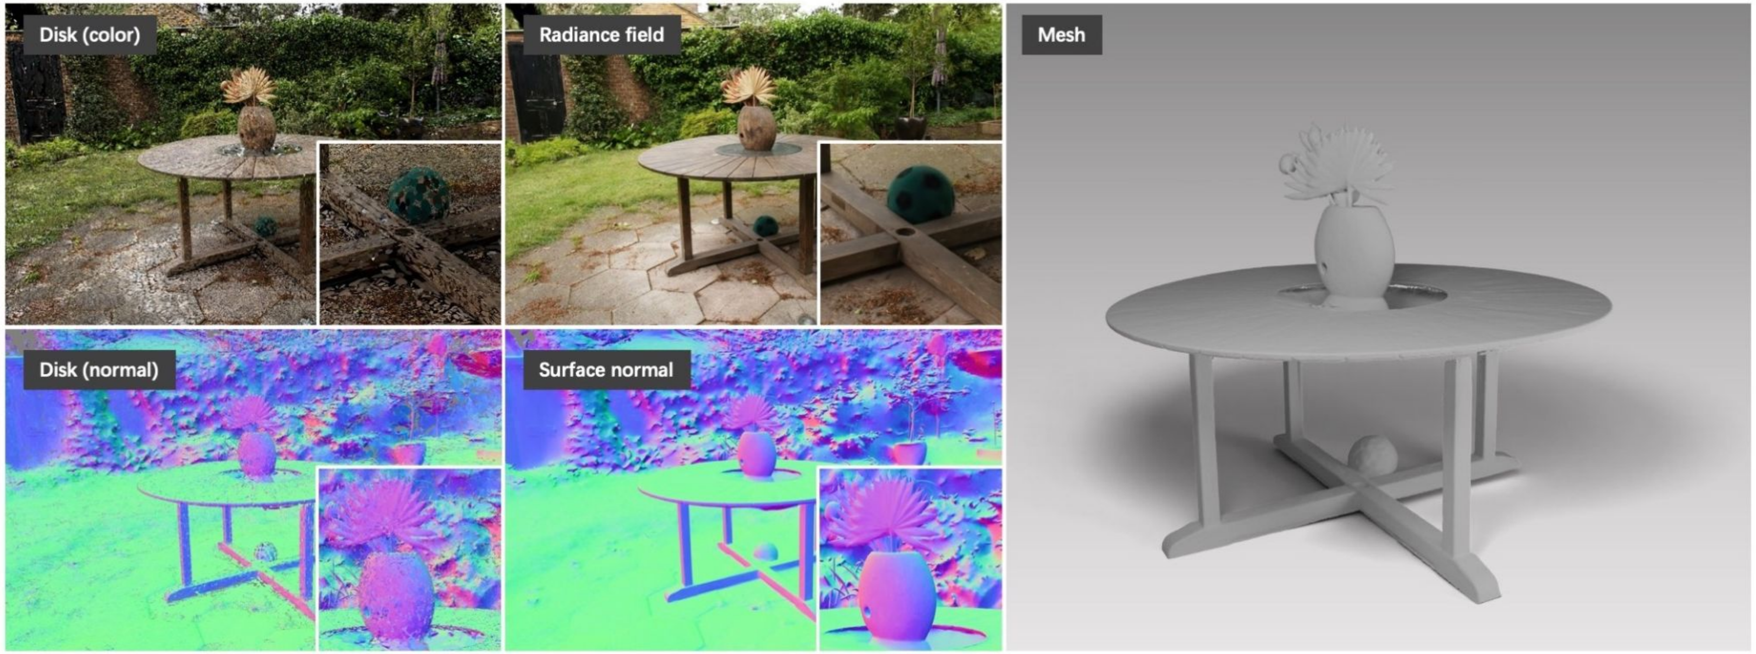

(a) 2D disks as surface elements (b) 2D Gaussian splatting (c) Meshing

Huang, Binbin, et al. "2d gaussian splatting for geometrically accurate radiance fields." arXiv preprint arXiv:2403.17888 (2024).

5. Geometry + Gaussian Splatting :: 2DGS

- Similar idea to 3DGS
  - Splat the 2D Gaussians to the image frame
  - Same rendering as in 3DGS

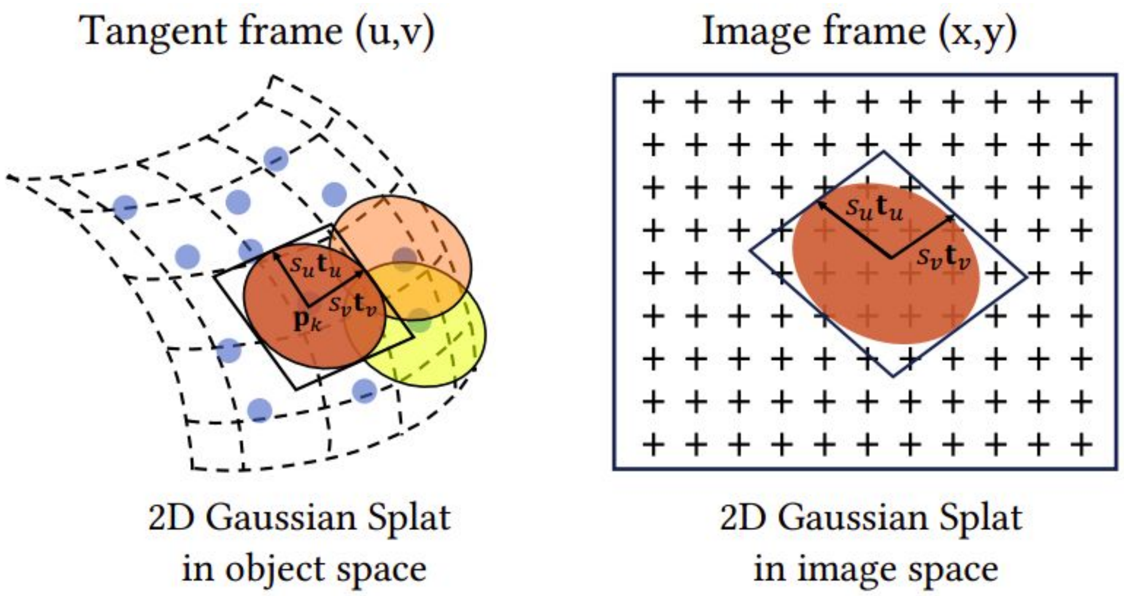

Huang, Binbin, et al. "2d gaussian splatting for geometrically accurate radiance fields." arXiv preprint arXiv:2403.17888 (2024).

5. Geometry + Gaussian Splatting :: 2DGS :: Results

- Performance much better than both previous works!
- Allows a much better reconstruction of geometry

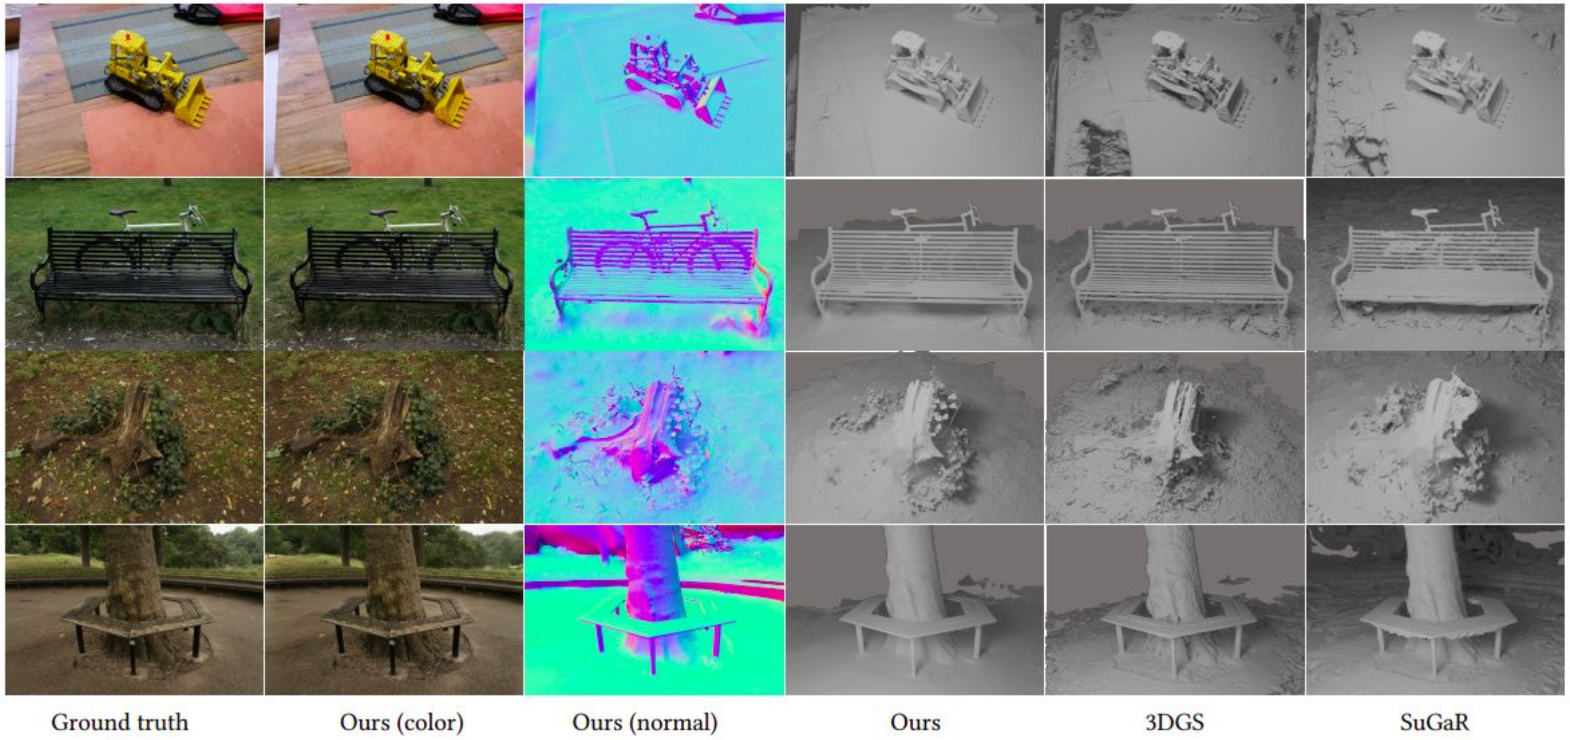

Huang, Binbin, et al. "2d gaussian splatting for geometrically accurate radiance fields." arXiv preprint arXiv:2403.17888 (2024).

5. Geometry + Gaussian Splatting :: 2DGS :: Results

- Performance much better than both previous works!
- Allows a much better reconstruction of geometry

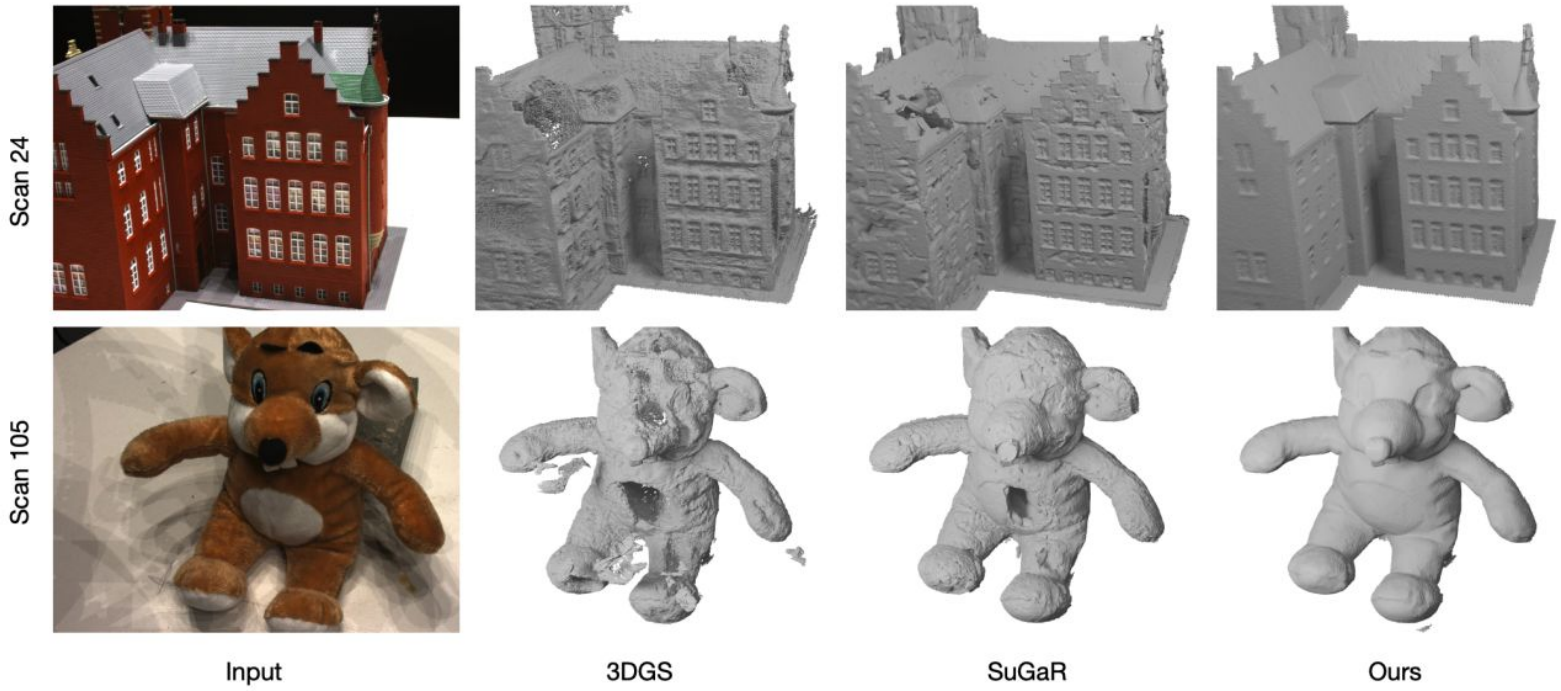

Huang, Binbin, et al. "2d gaussian splatting for geometrically accurate radiance fields." arXiv preprint arXiv:2403.17888 (2024).

5. Geometry + Gaussian Splatting :: 2DGS

- Given all this, what we can do about it?
- Work on the geometry of the objects
  - Use ideas from geometry which were explored by the group

5. Geometry + Gaussian Splatting :: 2DGS

- Currently working on understanding 2DGS and modifying it
  - Exploring how we can use the normals

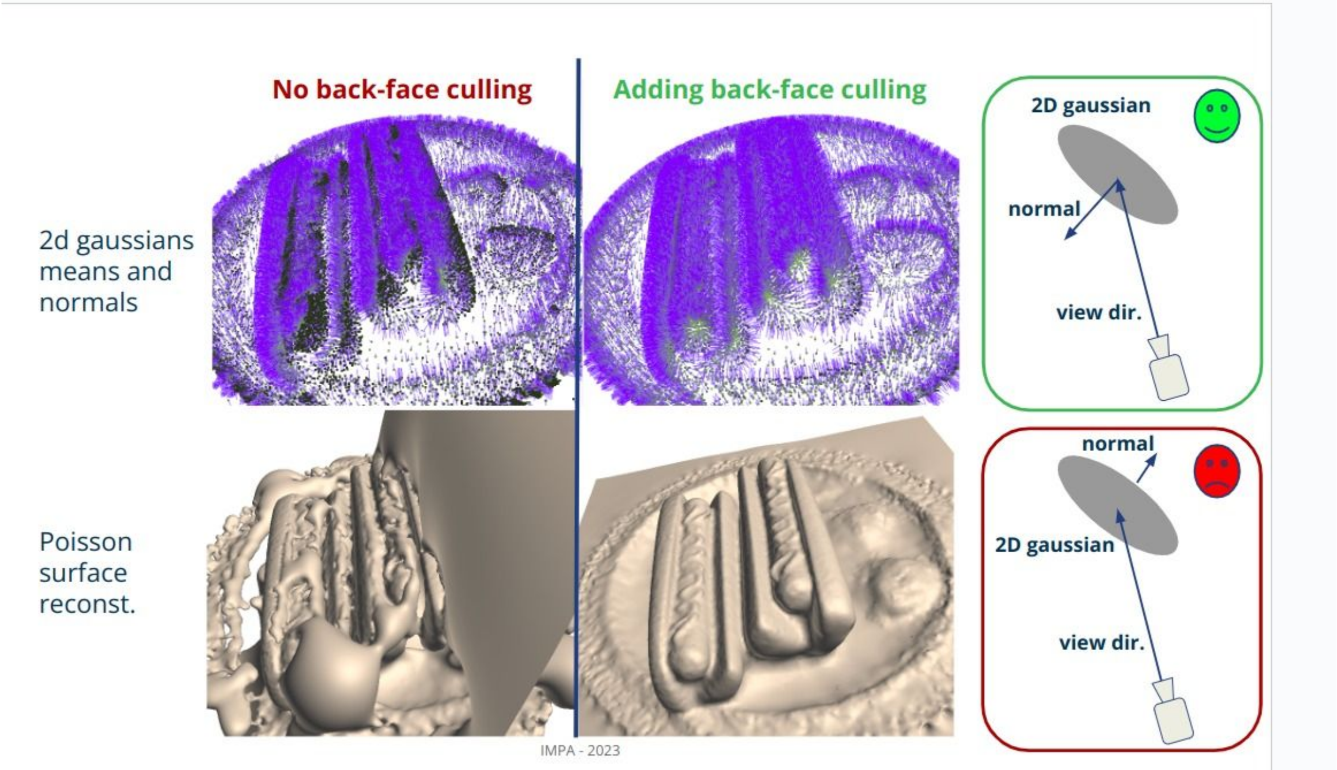

Huang, Binbin, et al. "2d gaussian splatting for geometrically accurate radiance fields." arXiv preprint arXiv:2403.17888 (2024).

5. Conclusion

5. Conclusion

- Gaussian Splatting is a really important and great technique for 3D Reconstruction
- More recent developments such as 2D Gaussian Splatting allows us to better extract geometric attributes such as normals
- For more information, there will be a course next semester!

3D Reconstruction Using Gaussian Splatting

Daniel Perazzo

Orientadores: Luiz Velho, João Paulo Lima and Tiago Novello

75 anos impa
Instituto de Matemática Pura e Aplicada

Universidade Federal de Pernambuco

voxar labs

Centro de Informática UFPE