



Differentiable rendering For Multi-View 3D Reconstruction

January 23rd, 2023

Multi-View 3D Reconstruction



I_1, \dots, I_n





P_1, \dots, P_n

$$P_i = K_i[R_i|t_i]$$




Dense 3D reconstruction
(Mesh, Voxels, Implicit)

Classical methods




NeRF – Volume rendering




$\sigma = 1$

$\sigma = 50$

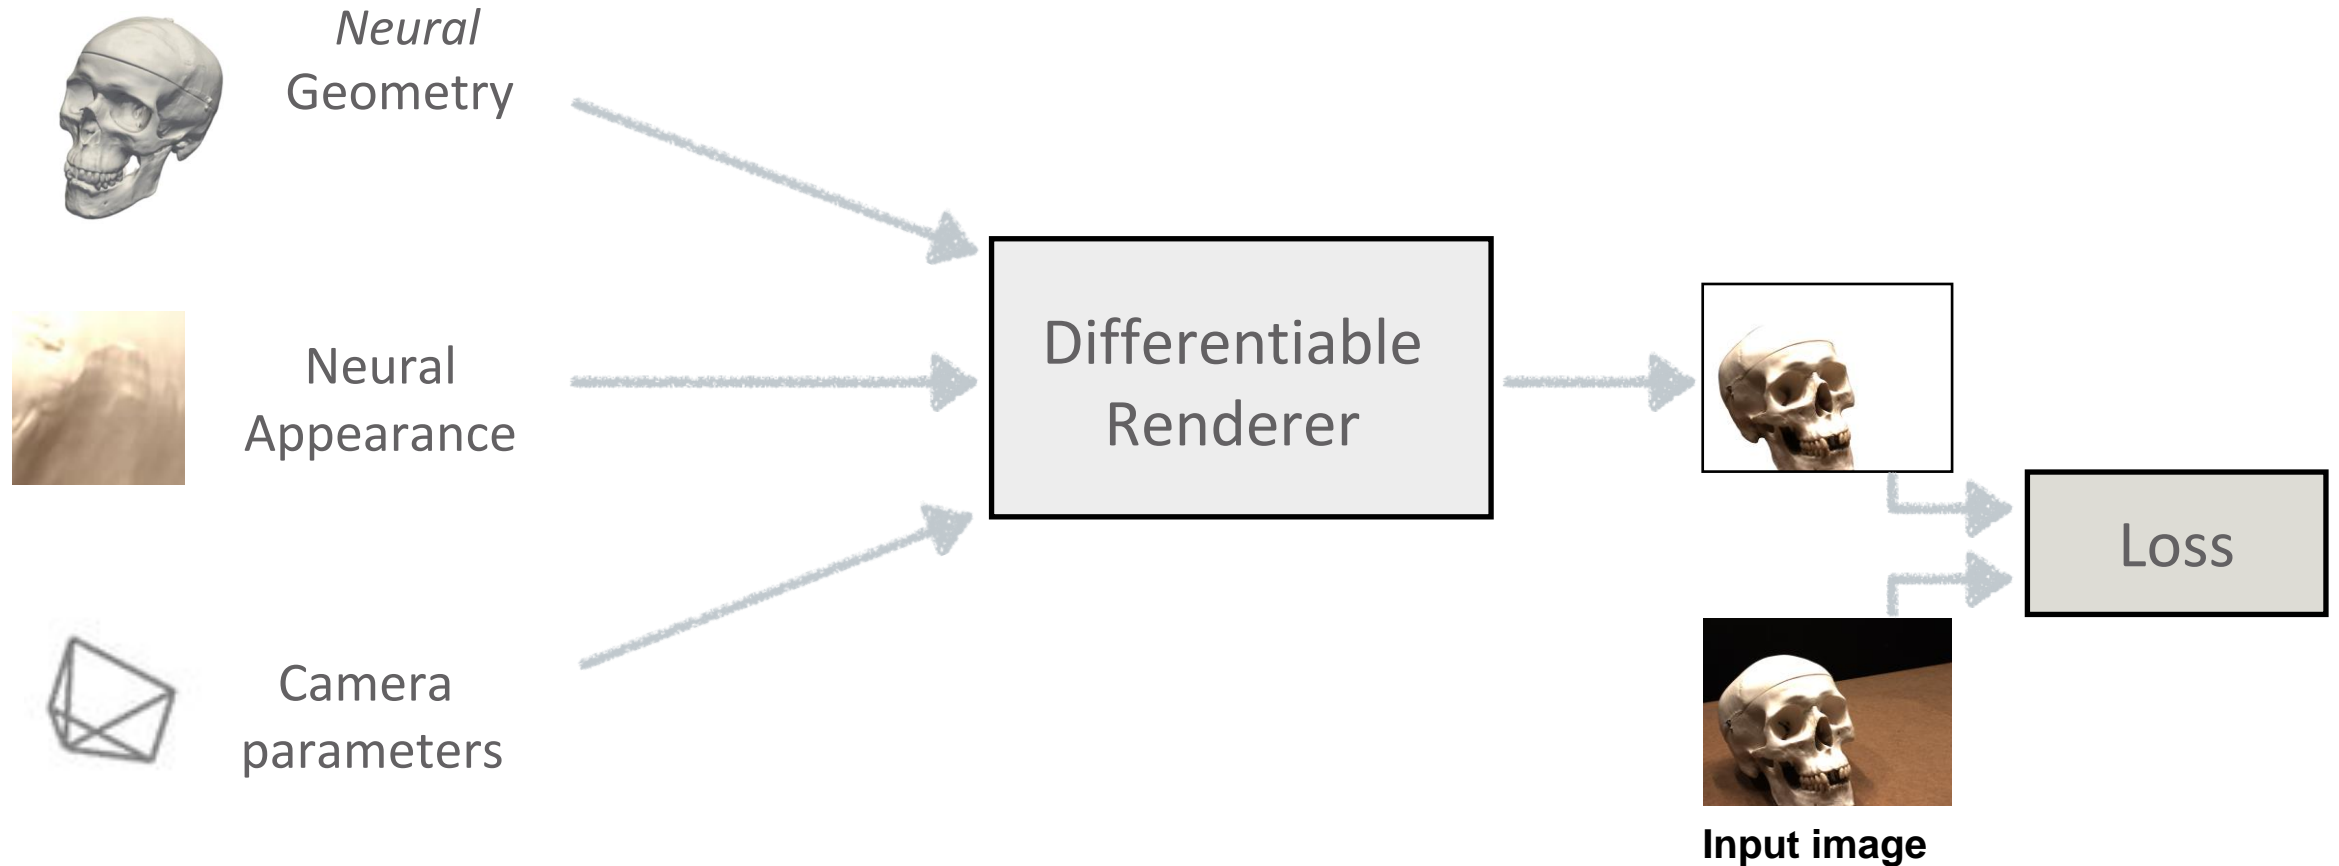


$\sigma = 500$

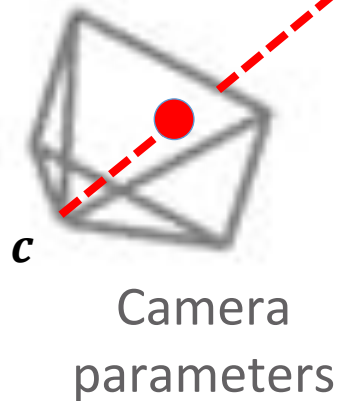
Differentiable **surface** rendering



Implicit Differentiable Renderer (IDR)




Neural Geometry



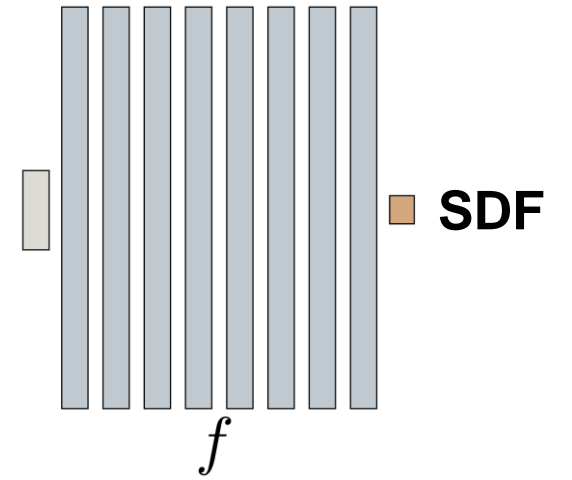
v

Neural Geometry



\hat{x}


Implicit Neural Representation



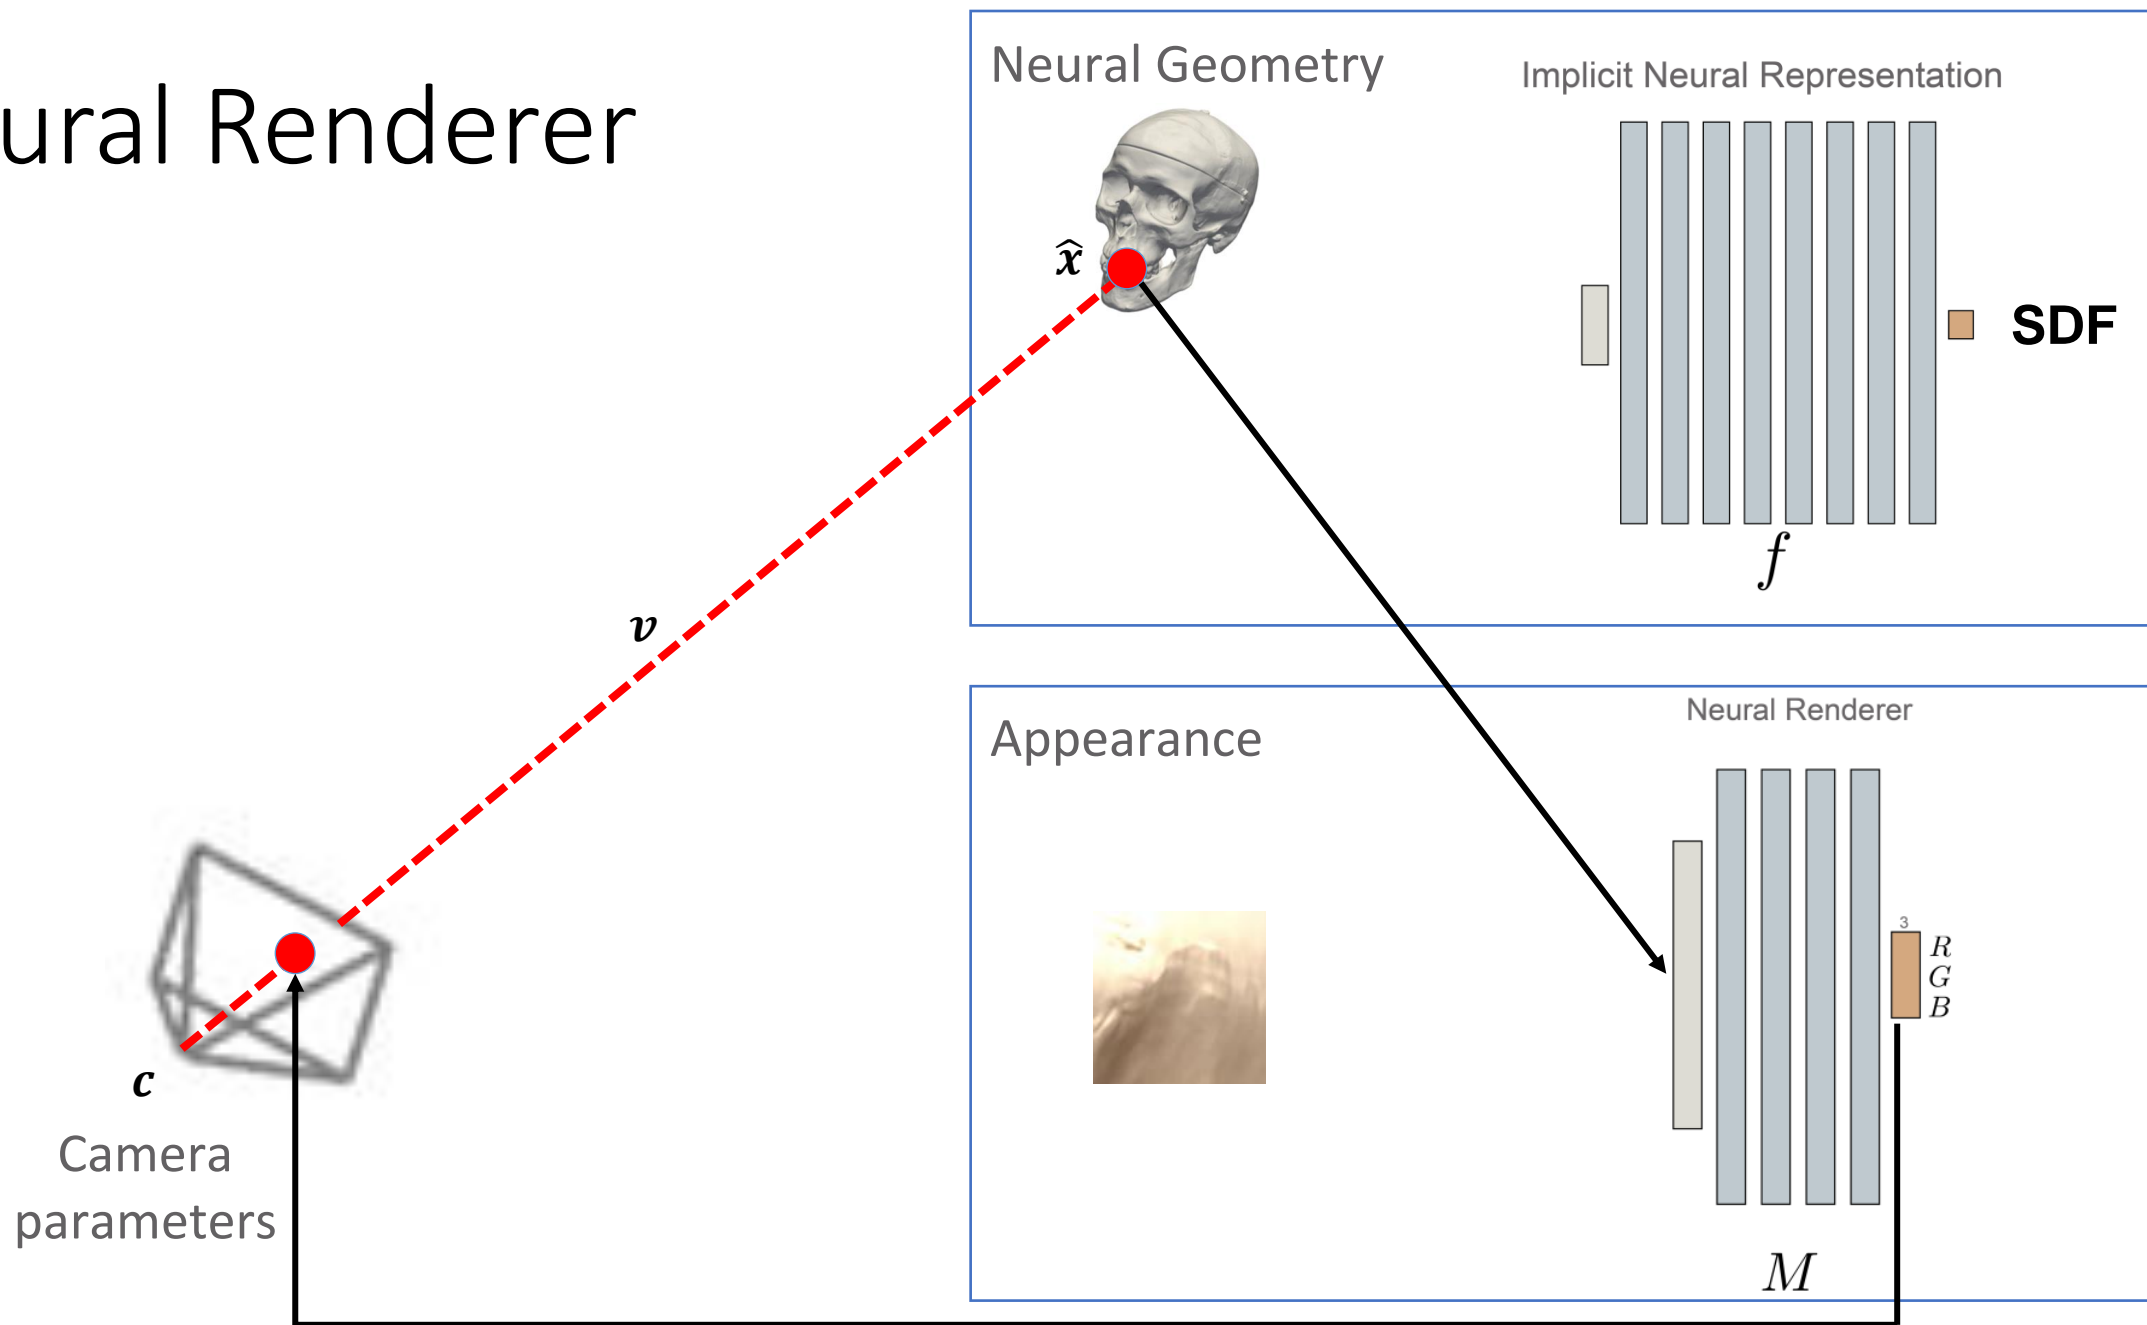
$$\mathcal{S}_\theta = \{x \in \mathbb{R}^3 \mid f(x; \theta) = 0\}$$

Sphere tracing

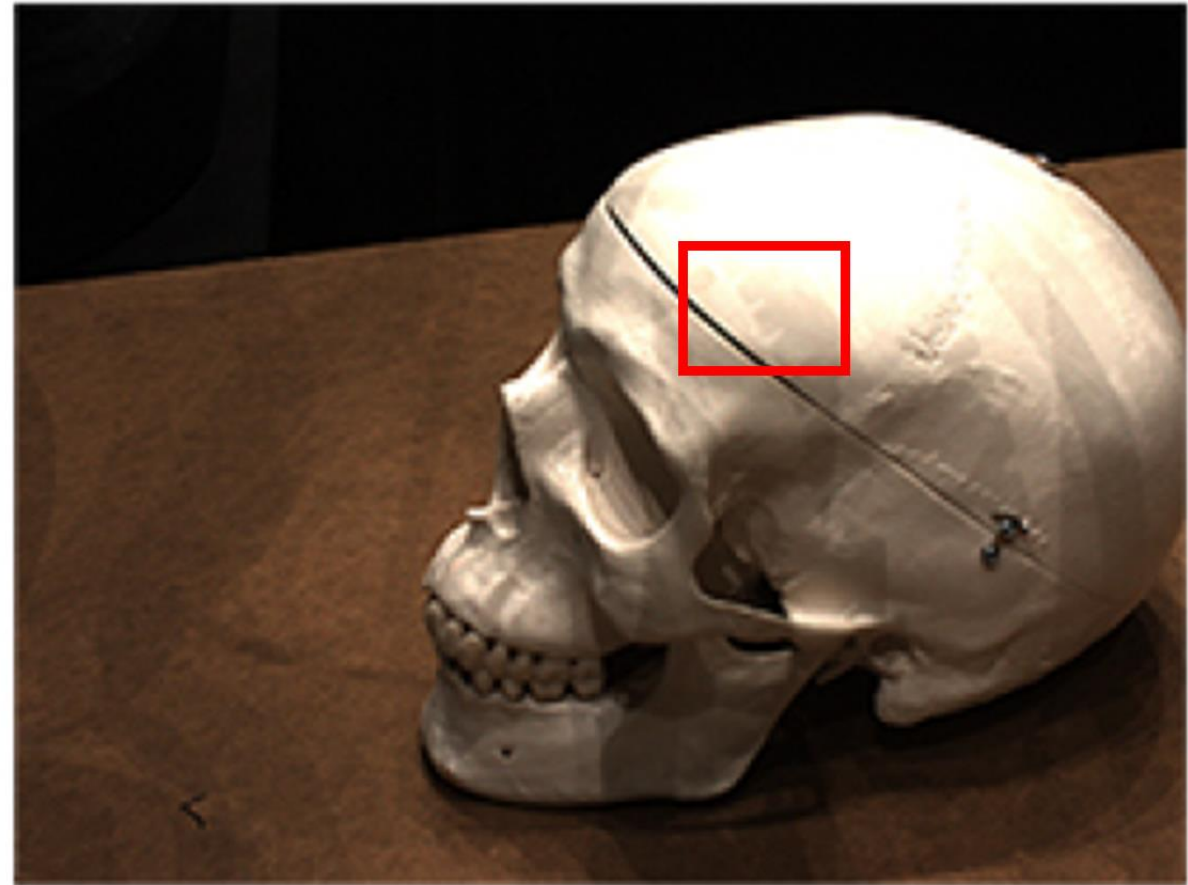

- Finding intersection point



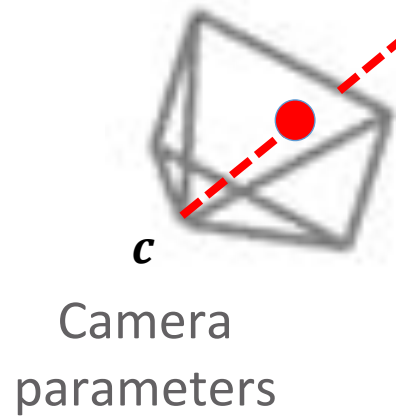
Neural Renderer




View dependent color



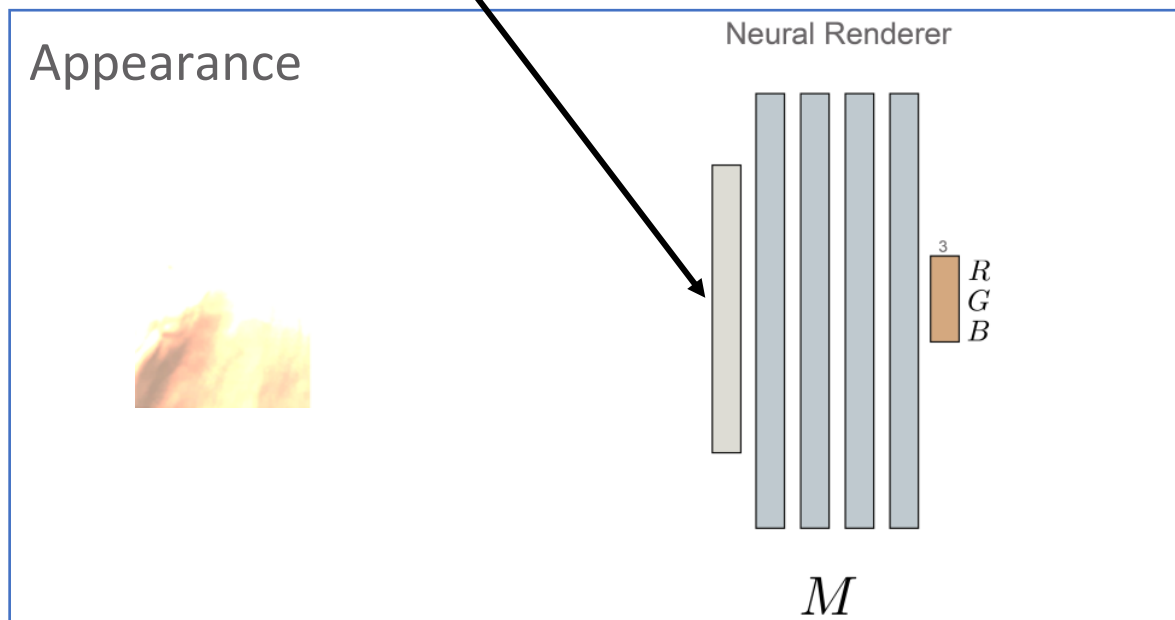
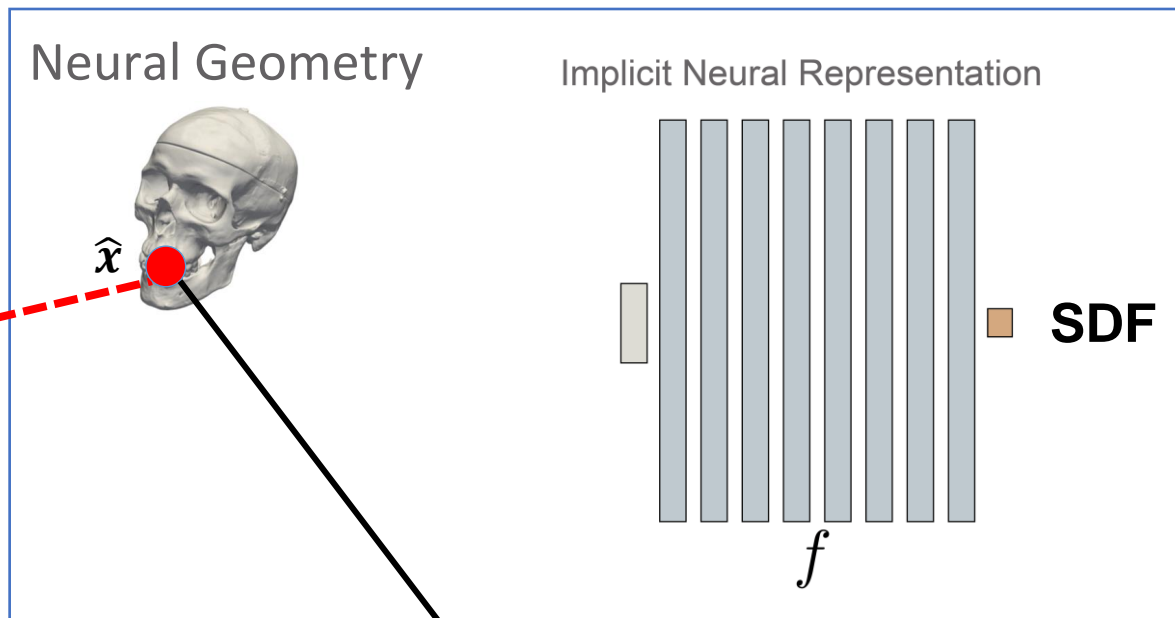

Neural Renderer




v



Neural Renderer




Neural Renderer

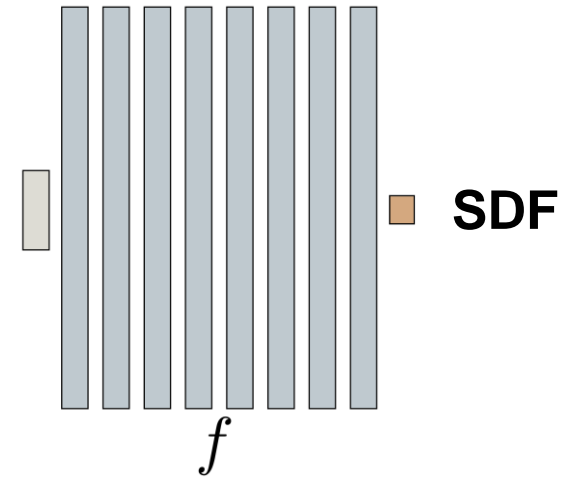


v

Neural Geometry



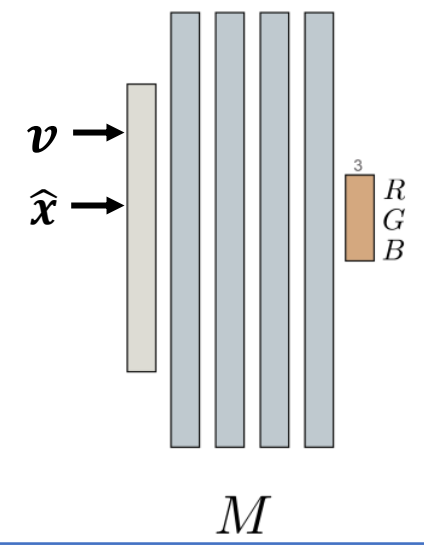
Implicit Neural Representation



Appearance

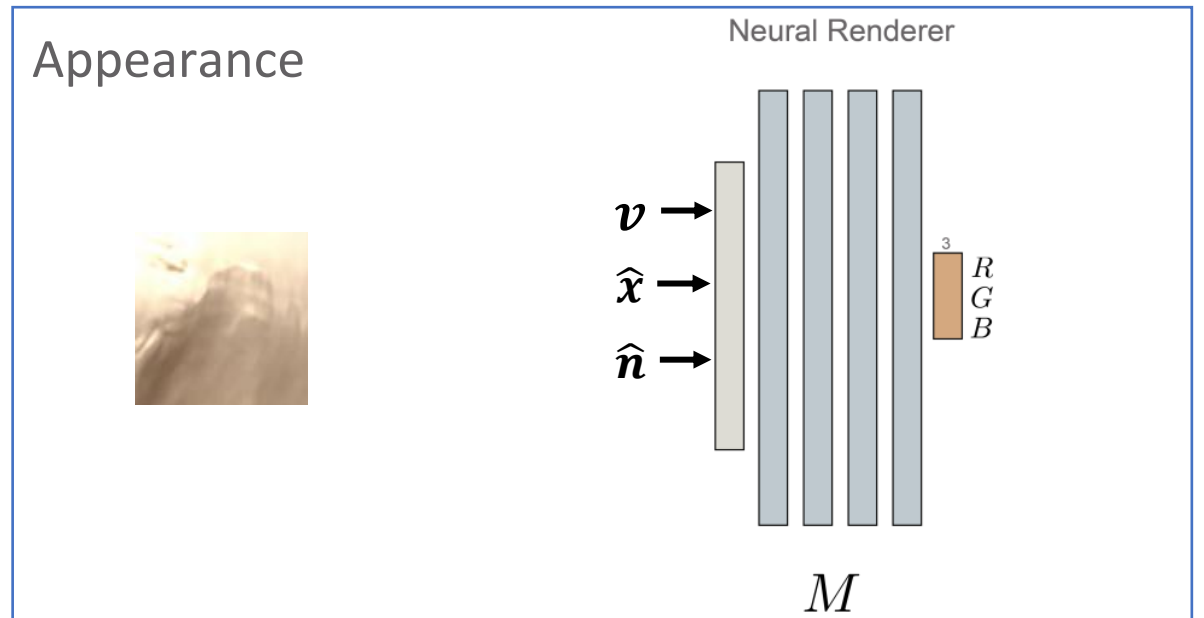
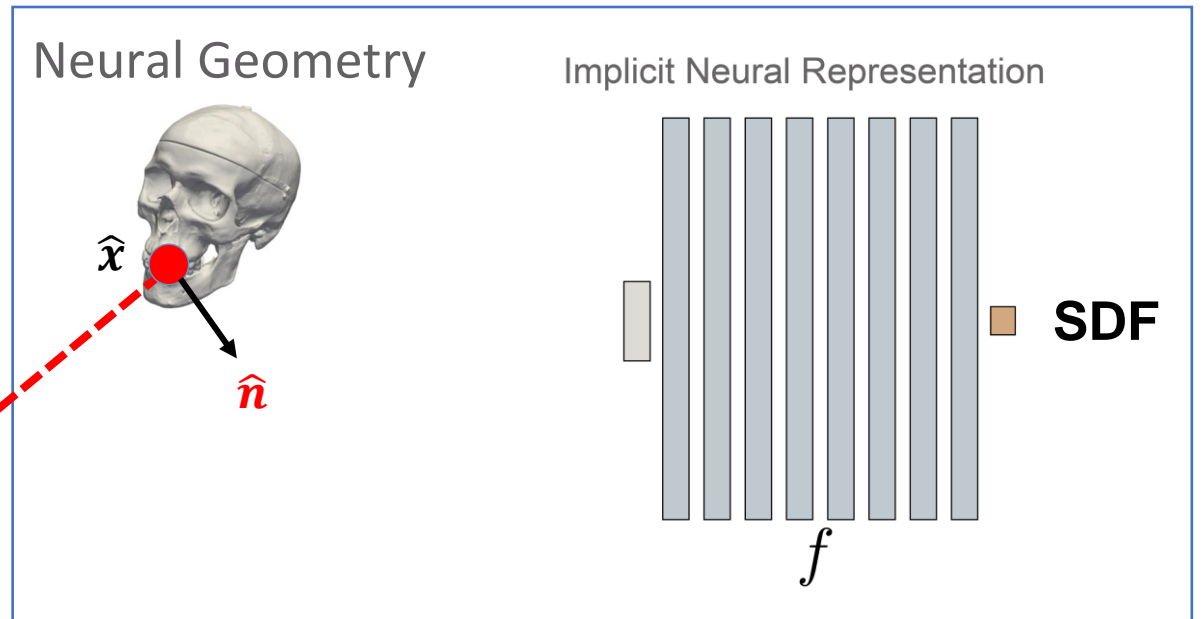



Neural Renderer




Neural Renderer

- Can we render a different geometry with the same renderer?
- What kind of input can “encourage” the renderer to generalize?




Neural Renderer

- Adding a global feature to allow secondary lighting effects and self shadows




v

Neural Geometry



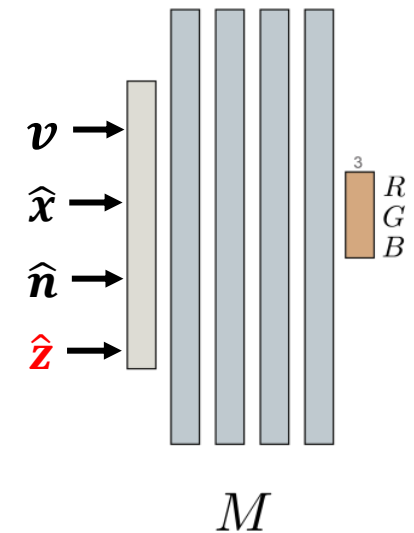
Implicit Neural Representation




Appearance



Neural Renderer




Training




v

Neural Geometry




Implicit Neural Representation




Appearance



Neural Renderer



Positional encoding



No PE, 5000 epochs




2000 epochs

Results



Results - Comparison



Ablation study



$$M(\hat{x}, \hat{n}, v, z)$$



$$M(\hat{x}, \hat{n}, v, z)$$




$$M(\hat{x}, \hat{n}, v, z)$$



$$M(\hat{x}, \hat{n}, v, z)$$

Disentangling geometry and appearance



Input images

Geometry

Rendering


IDR

Pros:

- Simultaneously learns geometry and appearance
- Highly accurate 3D reconstruction

Cons:

- Requires initialization
- Requires object masks



Random Init.
No Masks




Sphere Init.
No Masks¹



Sphere Init.
With Masks

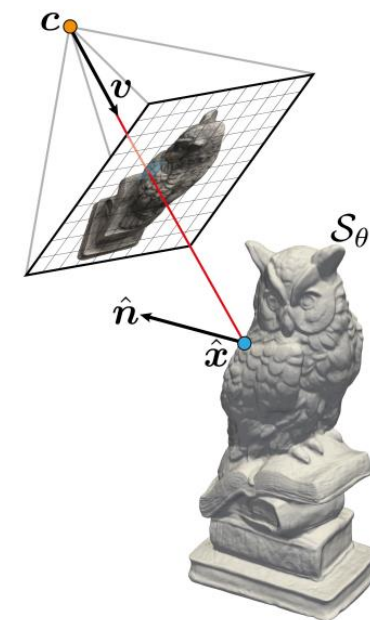
Nerf – Volume Rendering

- Poor geometry



IDR – Surface Rendering

- Requires masks
- Requires initialization

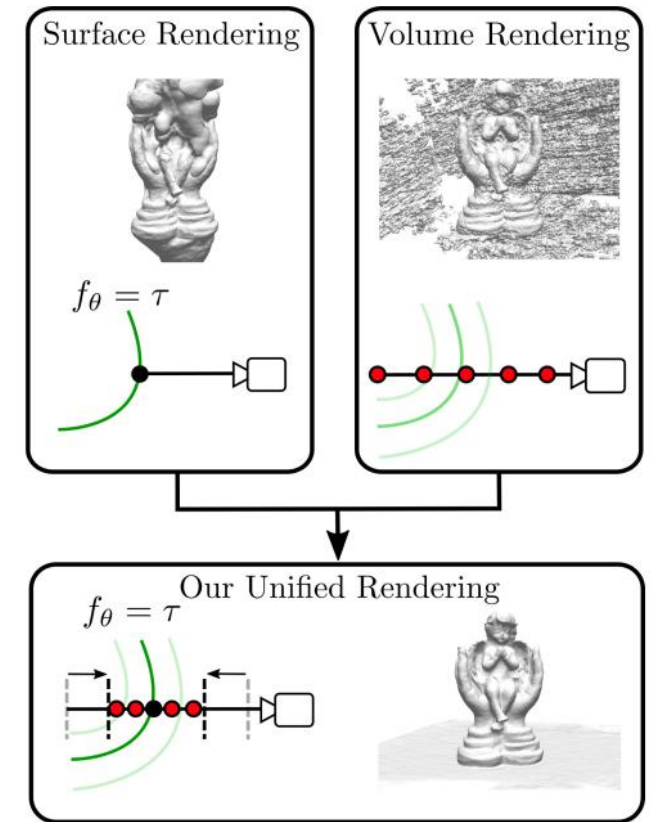


Density **Vs** Geometry

Sampling **Vs** Intersection point

UNISURF: Unifying Neural Implicit Surfaces and Radiance Fields for Multi-View Reconstruction

- Unified framework for implicit surfaces and radiance fields
- Reconstructing **solid** objects from a set of RGB images without masks

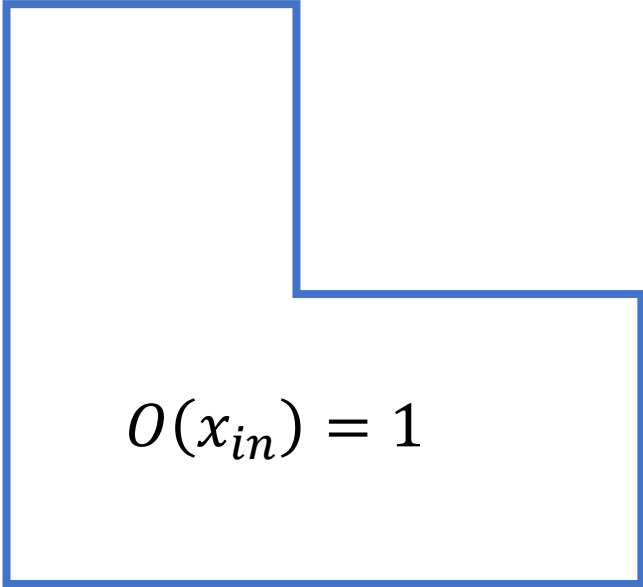


Density to occupancy

- Occupancy:

$$O: \mathbb{R}^3 \rightarrow \{0,1\}$$

$$O(x_{out}) = 0$$



$$O(x_{in}) = 1$$

Density to occupancy

- Occupancy:

$$O_{\theta}: \mathbb{R}^3 \rightarrow [0,1]$$

$$O(x_{out}) = 0$$



$$O(x_{surf}) = 0.5$$

$$O(x_{in}) = 1$$

Density to occupancy

- NeRF

$$\hat{C}(\mathbf{r}) = \sum_{i=1}^N T_i (1 - \exp(-\sigma_{\theta}(\mathbf{x}_i) \delta_i)) c_{\theta}(\mathbf{x}_i, \mathbf{d})$$

$$T_i = \exp\left(-\sum_{j<i} \sigma_{\theta}(\mathbf{x}_j) \delta_j\right)$$

$\sigma(x): R^3 \rightarrow R^+$ is the density function.
 $c(x, d)$ is the color function.

- Can be written as:

$$\hat{C}(\mathbf{r}) = \sum_{i=1}^N \alpha_i(\mathbf{x}_i) \prod_{j<i} (1 - \alpha_j(\mathbf{x}_j)) c(\mathbf{x}_i, \mathbf{d})$$

$$\alpha_i(\mathbf{x}) = 1 - \exp(-\sigma(\mathbf{x}) \delta_i)$$

$\alpha \in [0,1]$ become an occupancy indicator

Density to occupancy

- NeRF

$$\hat{C}(\mathbf{r}) = \sum_{i=1}^N T_i (1 - \exp(-\sigma_{\theta}(\mathbf{x}_i) \delta_i)) c_{\theta}(\mathbf{x}_i, \mathbf{d})$$

$$T_i = \exp\left(-\sum_{j<i} \sigma_{\theta}(\mathbf{x}_j) \delta_j\right)$$

$\sigma(x): R^3 \rightarrow R^+$ is the density function.
 $c(x, d)$ is the color function.

- Can be written as:

$$\hat{C}(\mathbf{r}) = \sum_{i=1}^N o(\mathbf{x}_i) \prod_{j<i} (1 - o(\mathbf{x}_j)) c(\mathbf{x}_i, \mathbf{d})$$

$o(x): R^3 \rightarrow [0,1]$ is the occupancy function

Rendering formula

- Volume rendering

$$\hat{C}_v(\mathbf{r}) = \sum_{i=1}^N o_\theta(\mathbf{x}_i) \prod_{j<i} (1 - o_\theta(\mathbf{x}_j)) c_\theta(\mathbf{x}_i, \mathbf{n}_i, \mathbf{h}_i, \mathbf{d})$$

x is a 3D points
 n is the surface normal
 h is a feature vector
 d is the viewing direction

- Surface rendering

$$\hat{C}_s(\mathbf{r}) = c_\theta(\mathbf{x}_s, \mathbf{n}_s, \mathbf{h}_s, \mathbf{d})$$




Sampling Vs surface point

1. Find surface point


2. Define intervals:

$$\Delta_k = \max(\Delta_{\max} \exp(-k \beta), \Delta_{\min})$$

3. Sample ray








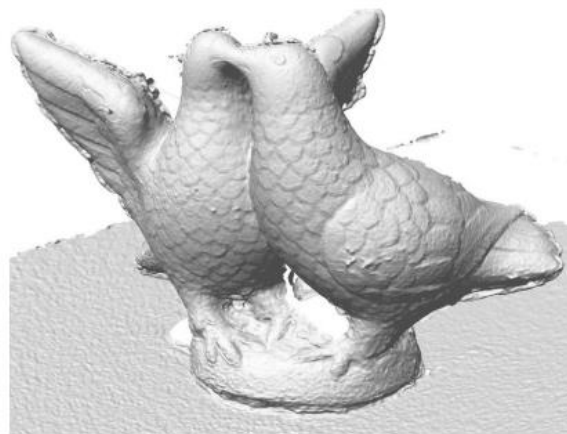
Results



Results



GT view



COLMAP [48]

NeRF [34]

Ours

Neural 3D Reconstruction in the Wild



Summary

- The underlying geometry, obtained by volume renderers is usually non-smooth and contains artifacts.
- Differentiable surface renderers produce highly accurate 3D reconstructions. However, they depend on object masks.
- Unifying surface and volume rendering is possible!

Questions?