

# Lecture: Neural Fields Part 1

# How does Computer Vision & Graphics work together?



Current Image

## Explicit 3D Reconstruction with Neural Fields.



3D Intrinsic Components

Explicit: Reconstruct 3D  
(Introduction to Graphics Lectures)

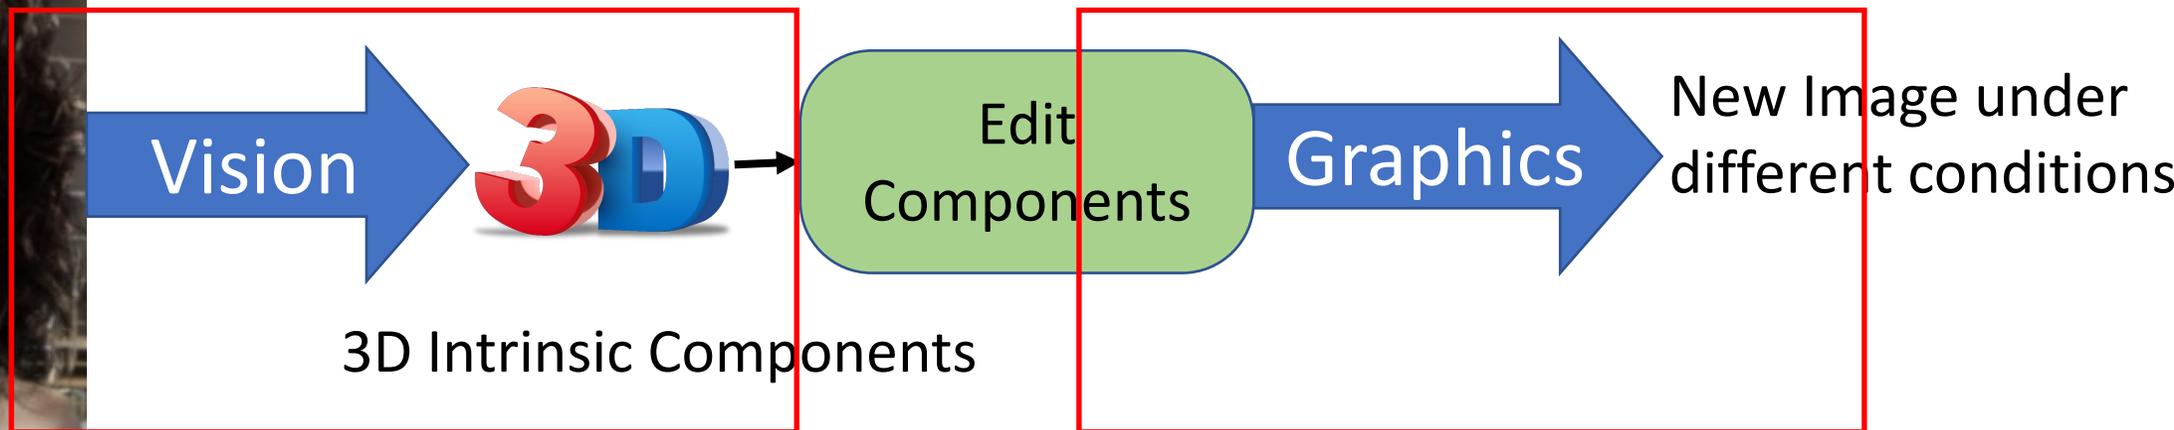
Implicit: Neural Representation  
(Generative Models Lectures)

Change:

- Viewpoint
- Lighting
- Reflectance
- Background
- Attributes
- Many others...

# Outline

- What is Neural Fields & why it got so much attention?
- The Prelude: Neural Implicit Surfaces
- Introduction to Volume Rendering

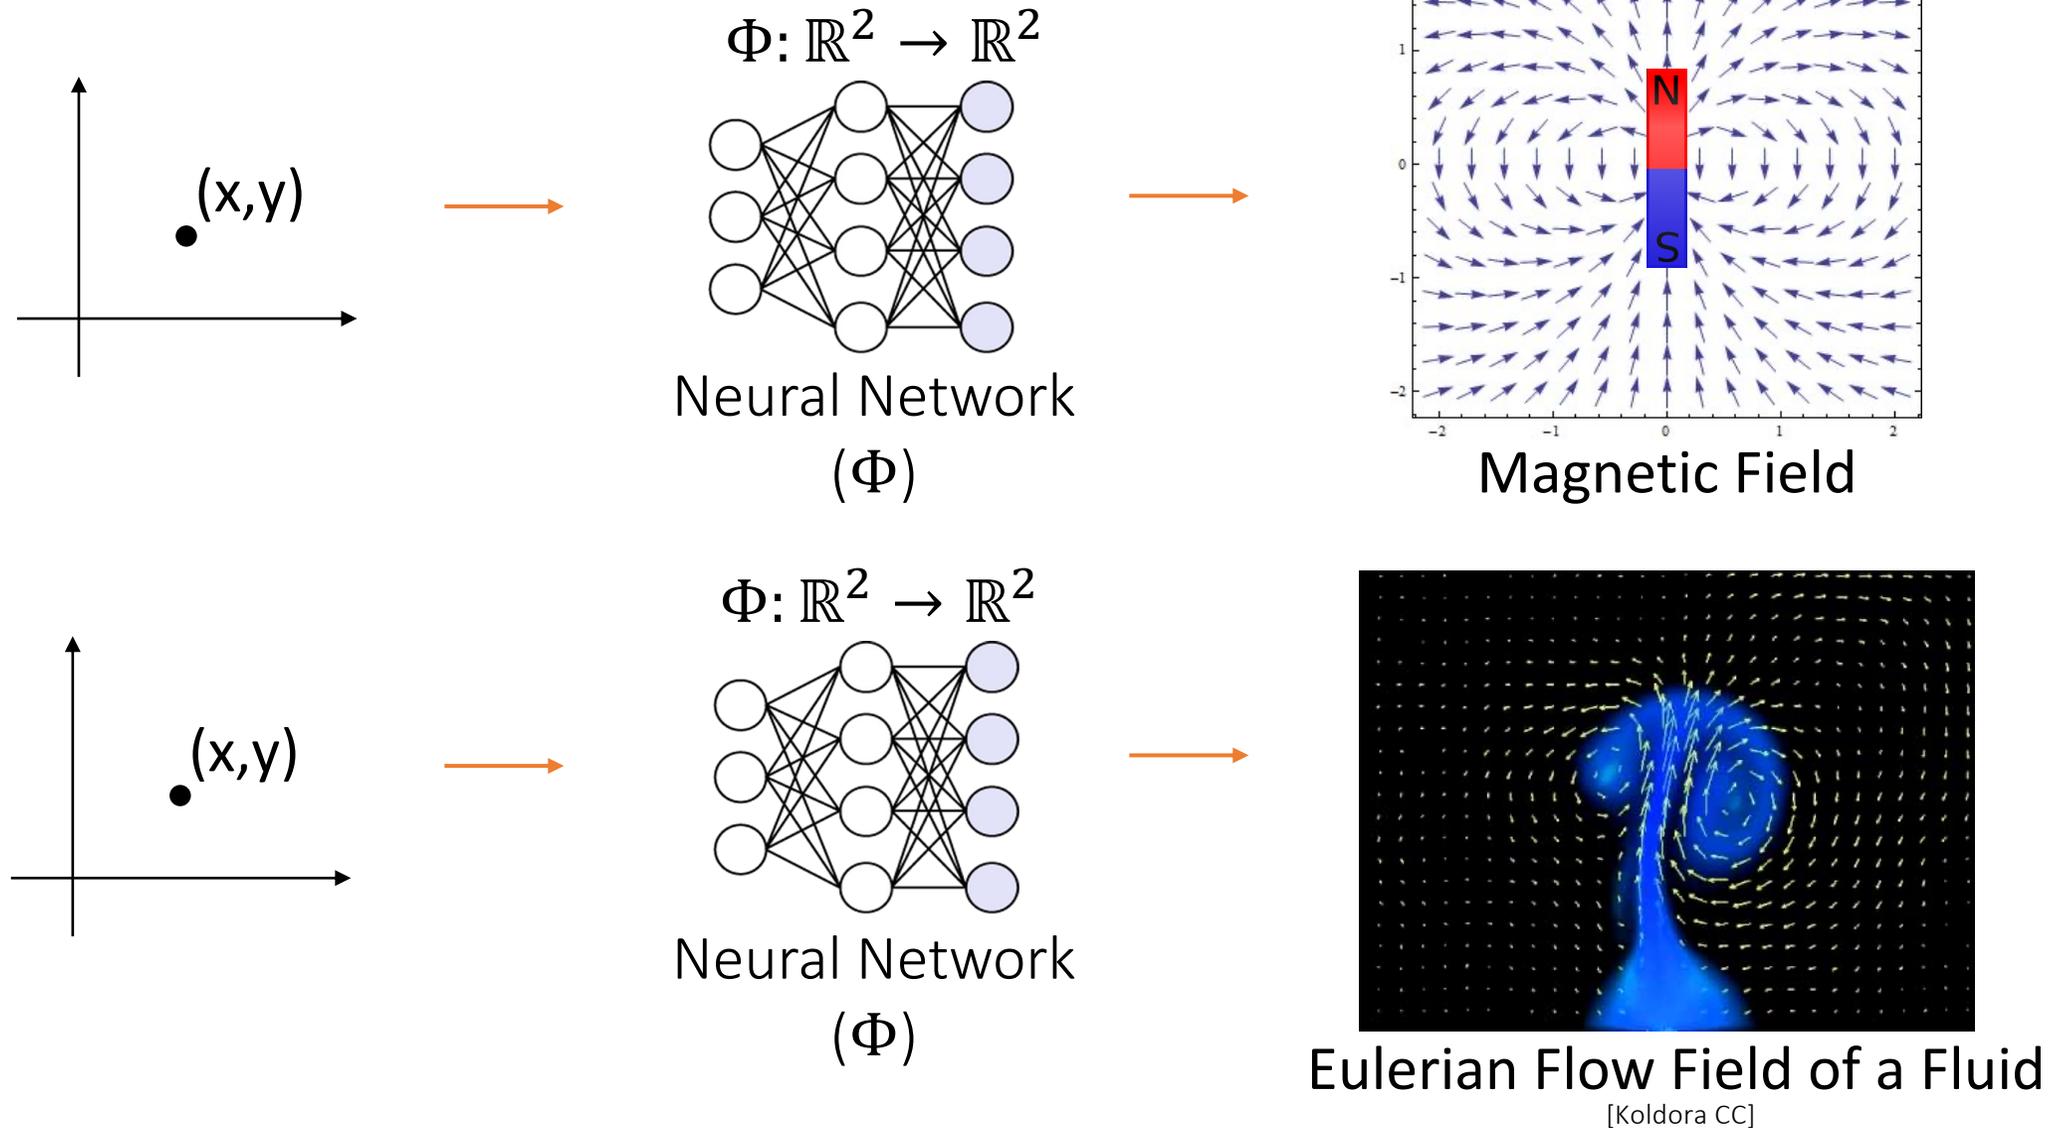


Current Image

# Outline

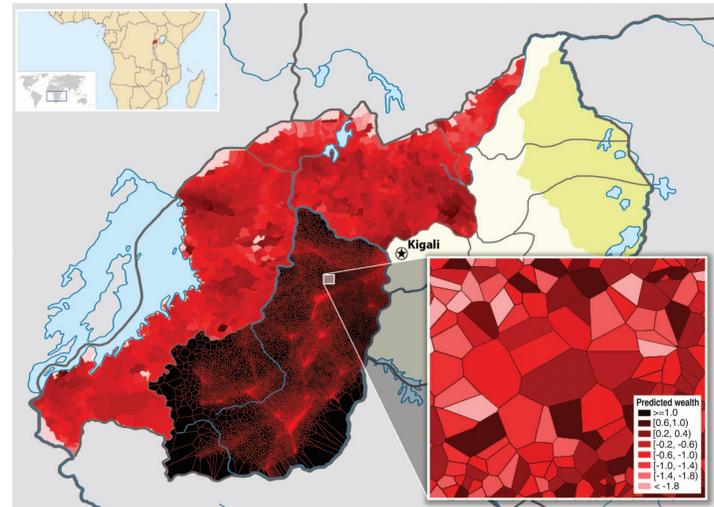
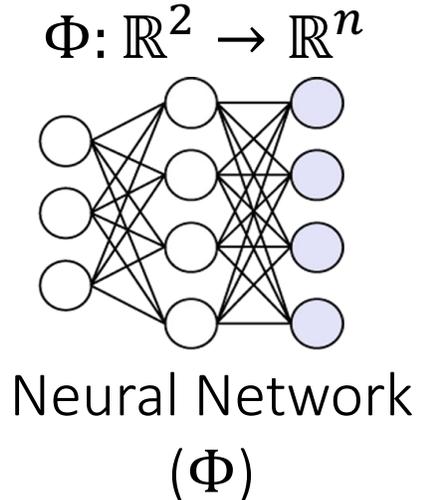
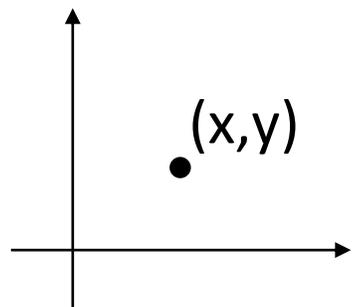
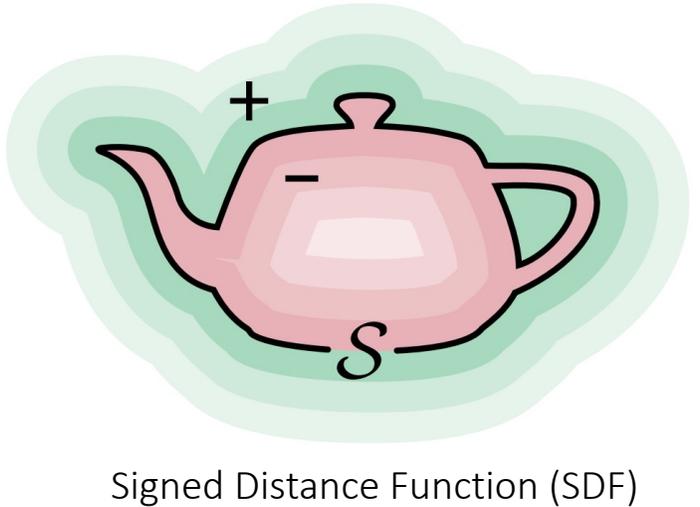
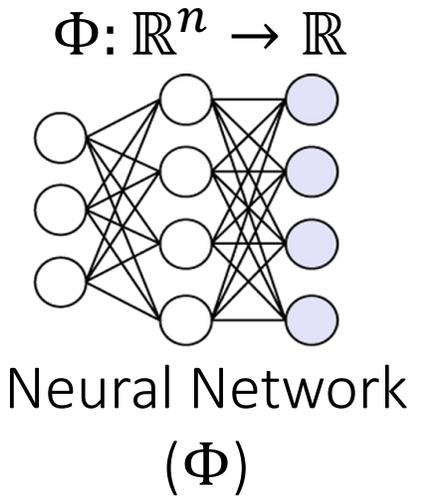
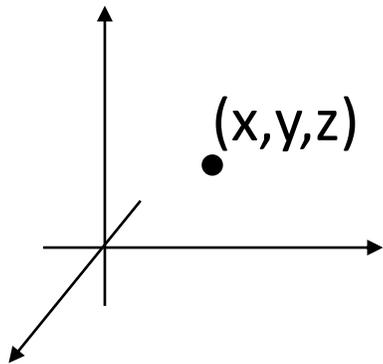
- What is Neural Fields & why it got so much attention?
- The Prelude: Neural Implicit Surfaces
- Introduction to Volume Rendering

# What are neural fields?



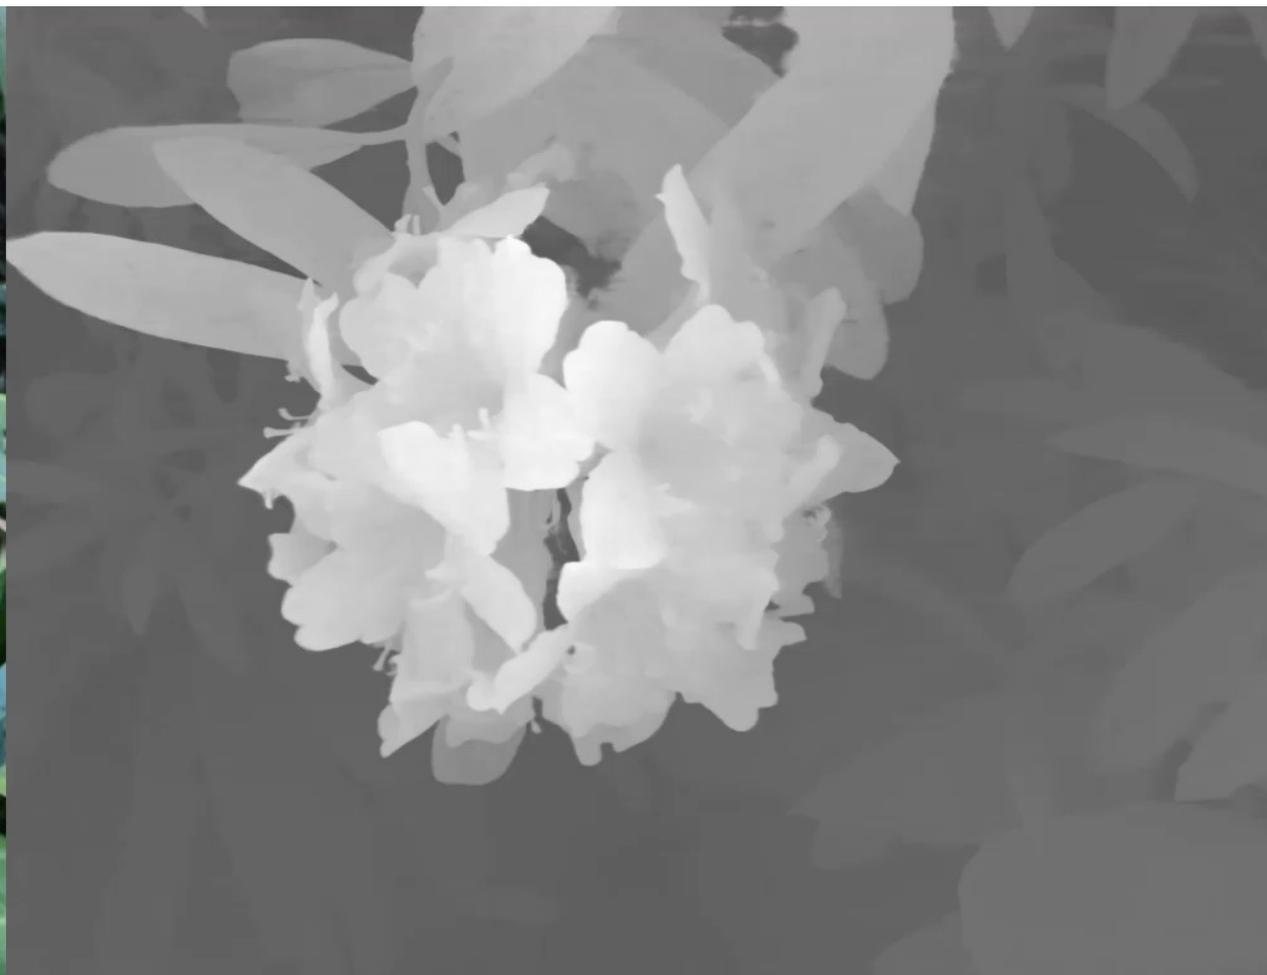
[Koldora CC]

# What are neural fields?



NeRF (Neural Radiance Field) has revolutionized  
Computer Vision & Graphics in past 2 years!

Let's look at some of the stunning results it produced!



**NeRF: Representing Scenes as Neural Radiance Fields for View Synthesis,**  
Ben Mildenhall, *Pratul Srinivasan*, Matthew Tancik\*, Jonathan Barron, Ravi Ramamoorthi, Ren Ng, ECCV 2020.

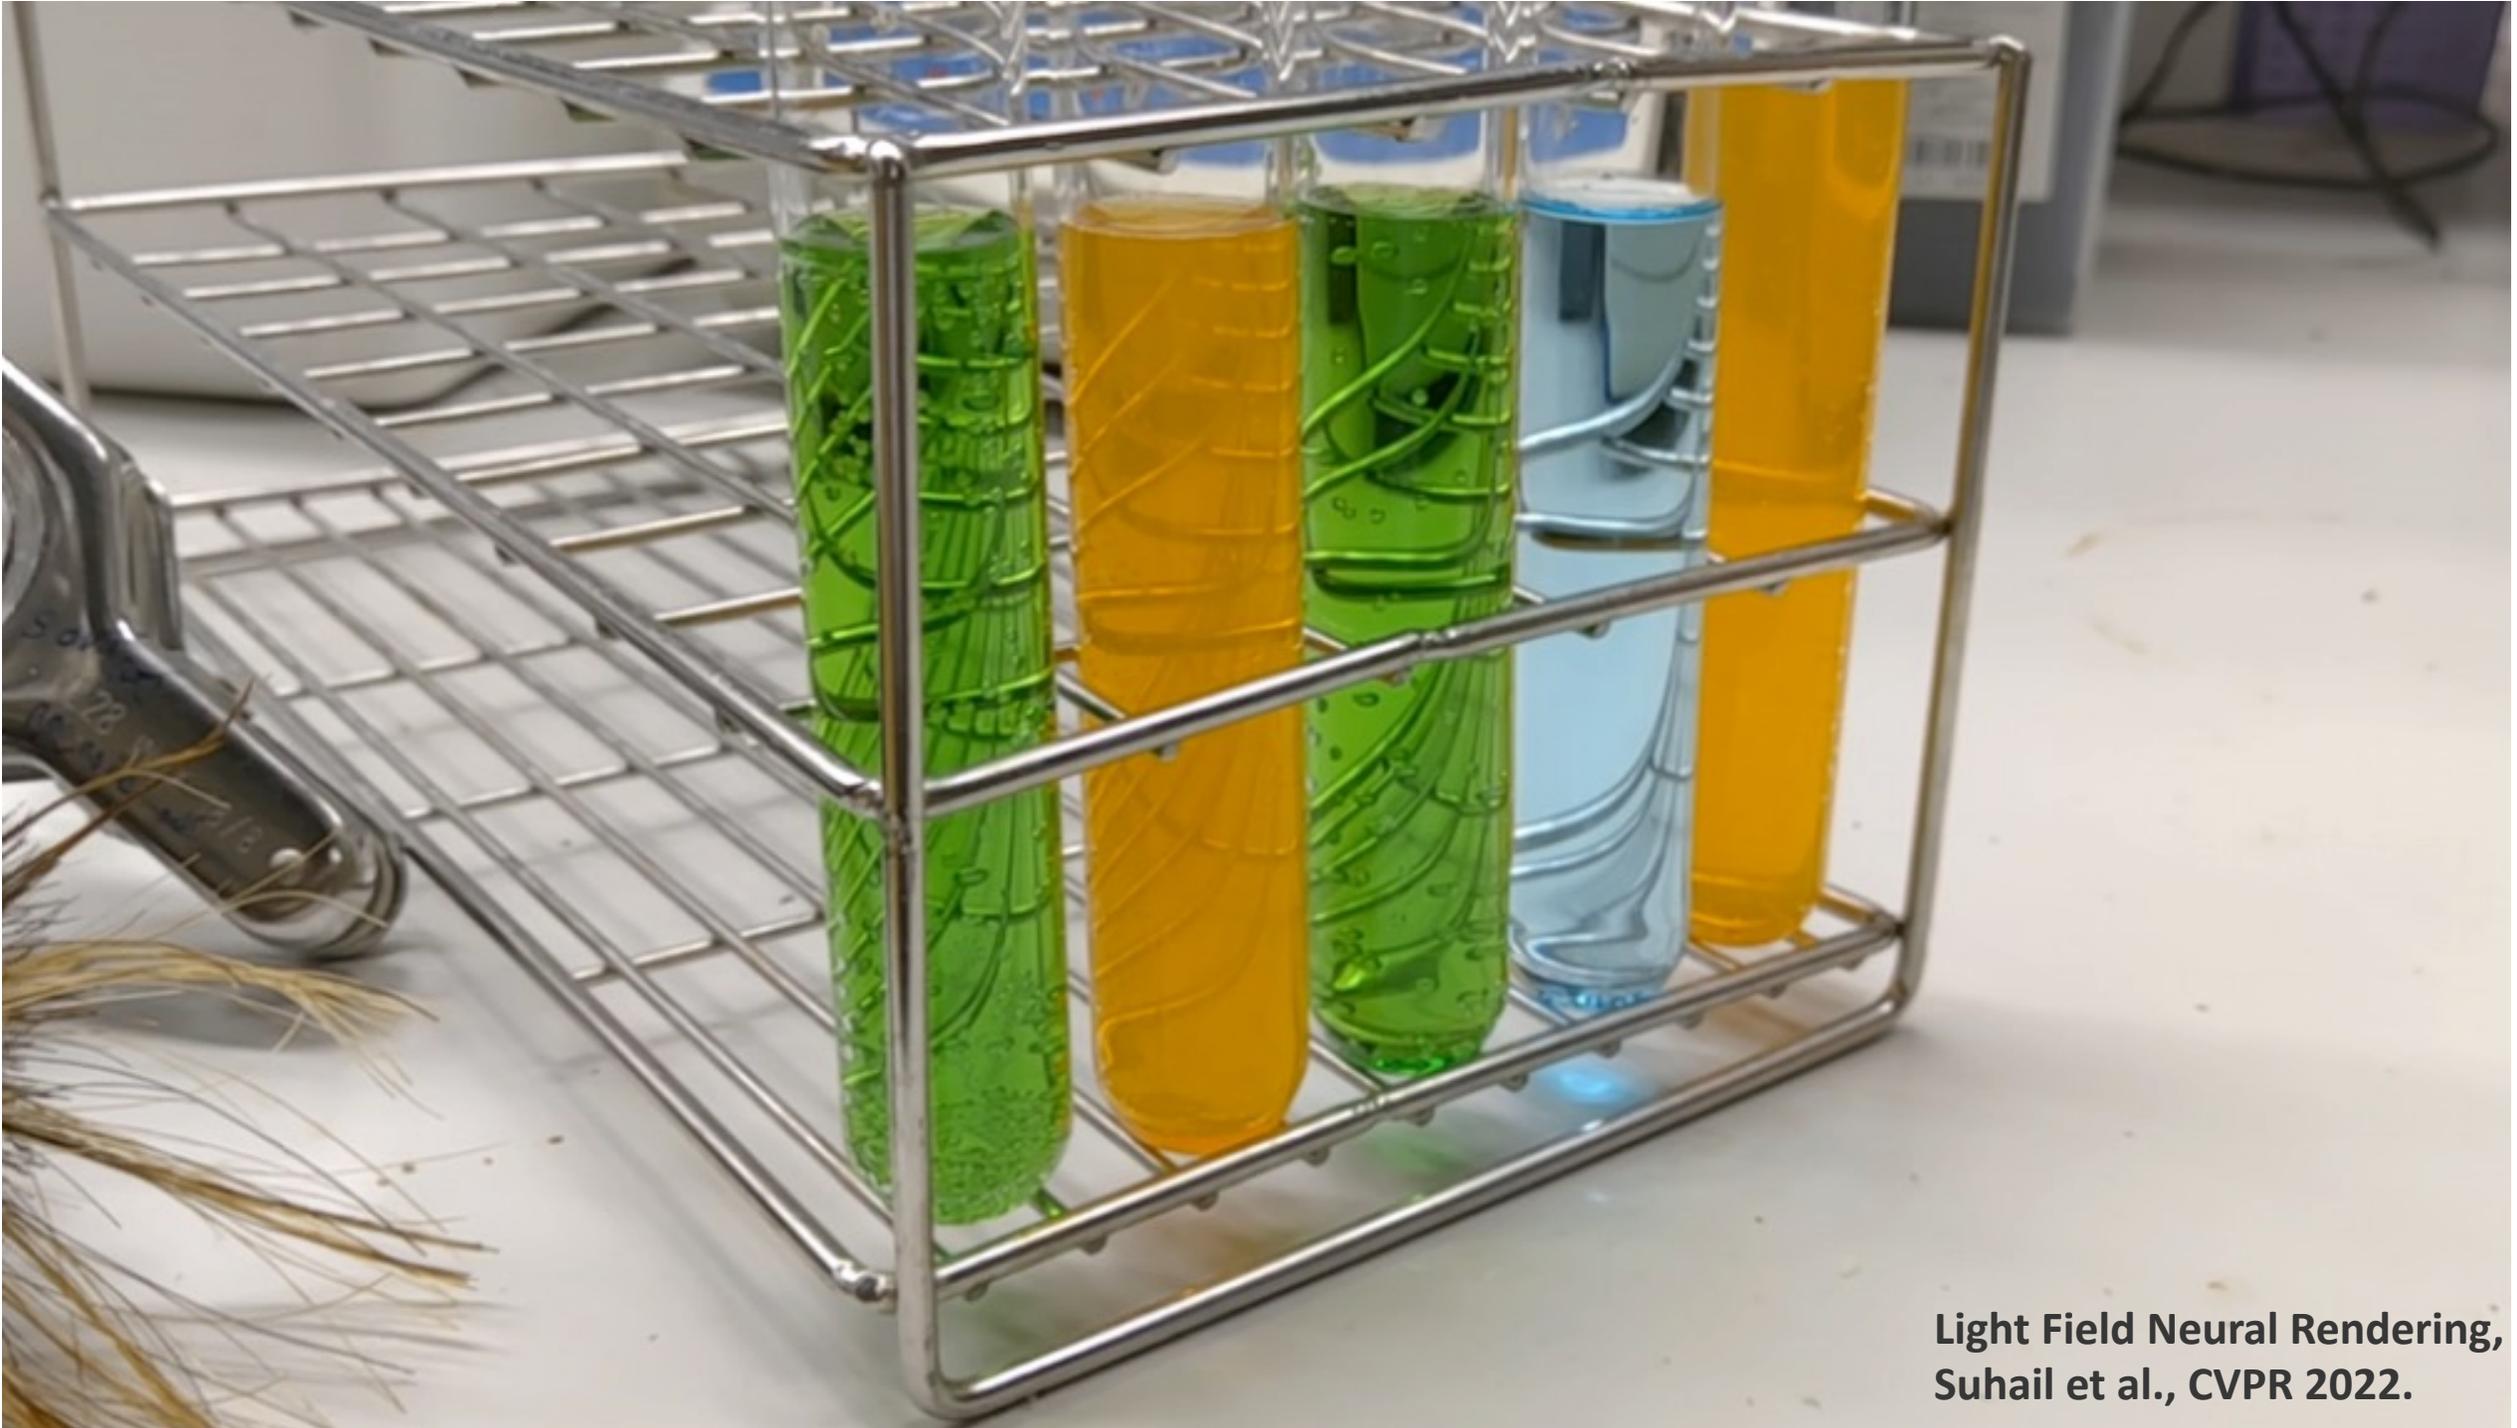


**NeRF: Representing Scenes as Neural Radiance Fields for View Synthesis,**  
Ben Mildenhall, *Pratul Srinivasan*, Matthew Tancik\*, Jonathan Barron, Ravi Ramamoorthi, Ren Ng, ECCV 2020.

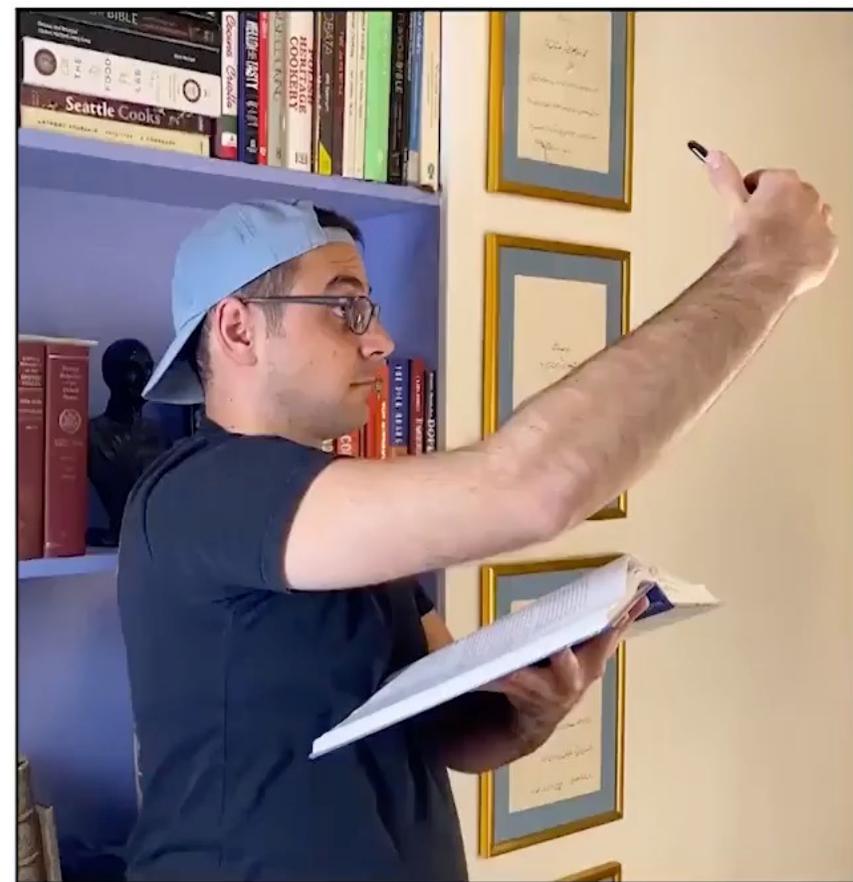




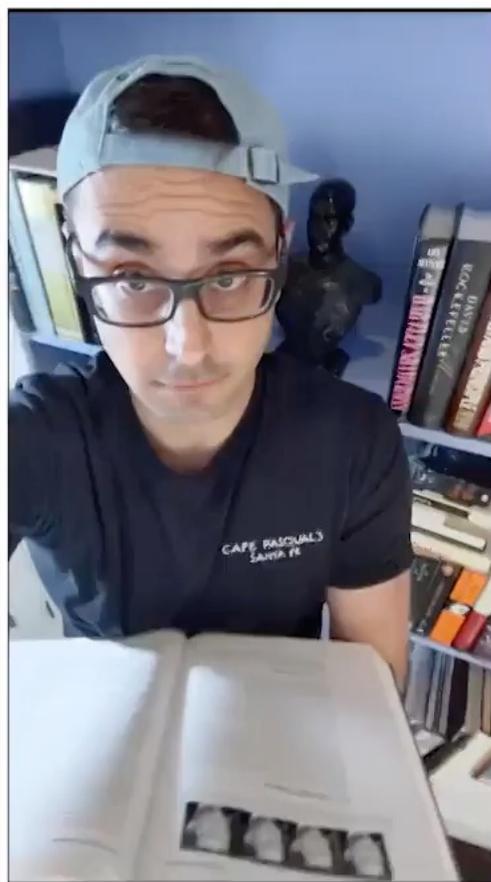
Block-NeRF: Scalable Large Scene Neural View Synthesis, CVPR 2022.



Light Field Neural Rendering,  
Suhail et al., CVPR 2022.



(a) Capture Process



(b) Input



(c) Nerfie



(d) Nerfie Depth

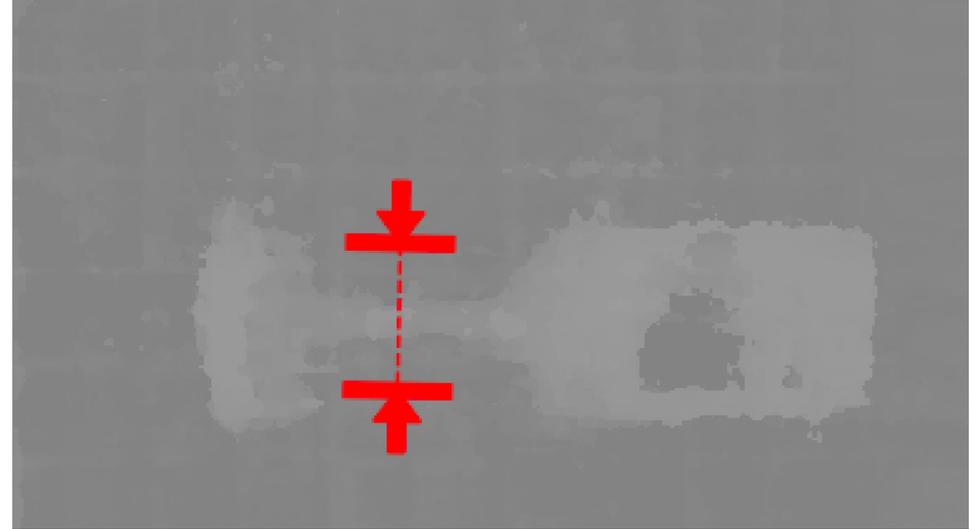
**NeRFies: Deformable Neural Radiance Fields**, Keunhong Park et al., ICCV 2021.



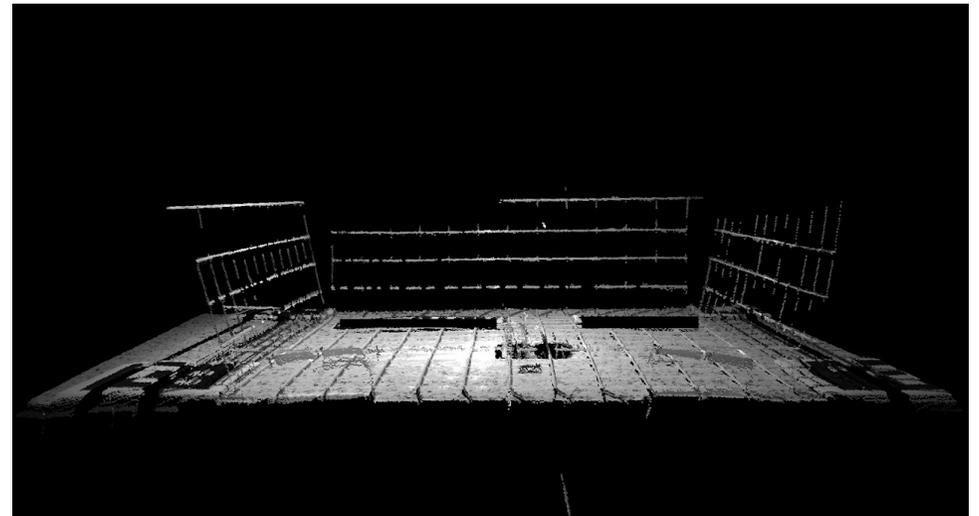
Neural 3D Video Synthesis  
from Multi-view Video,  
Li et al., CVPR 2022



“GRAM: Generative Radiance Manifolds for 3D-Aware Image Generation”, Deng et al., CVPR 2022.

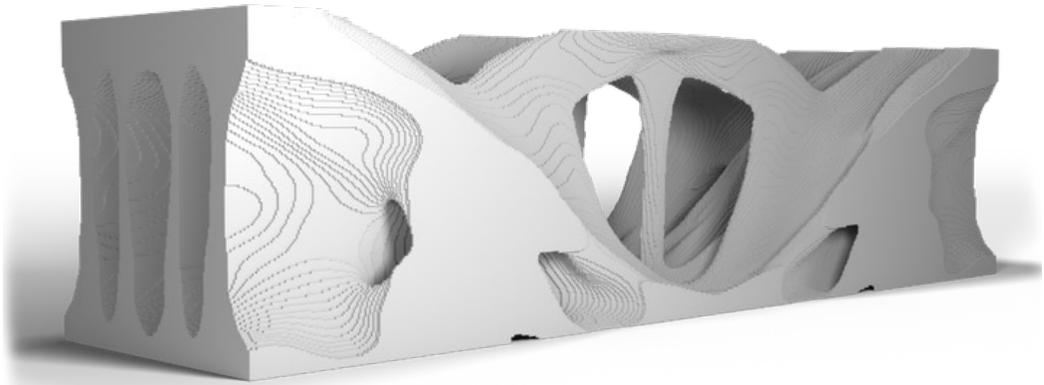


Grasp shown in red

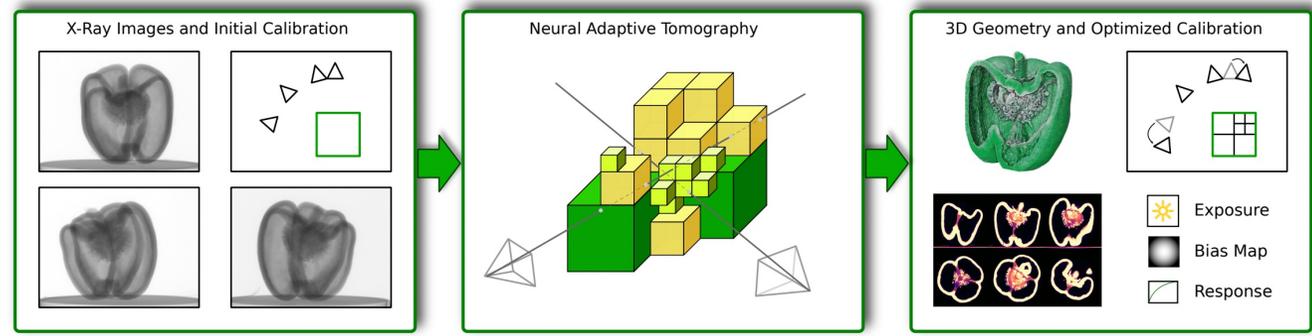


Point cloud from sensor

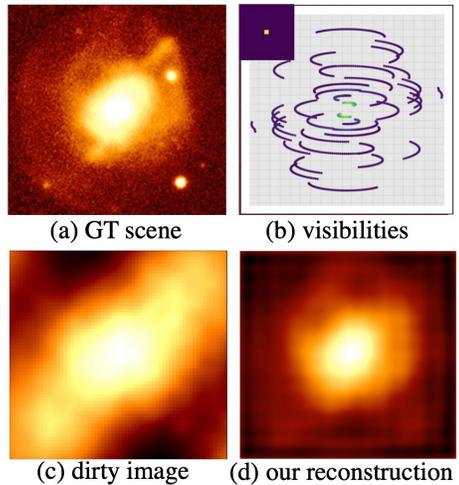
# Neural Fields for Science and Engineering



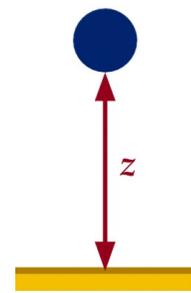
Topology Optimization [Doosti et al. 2021]



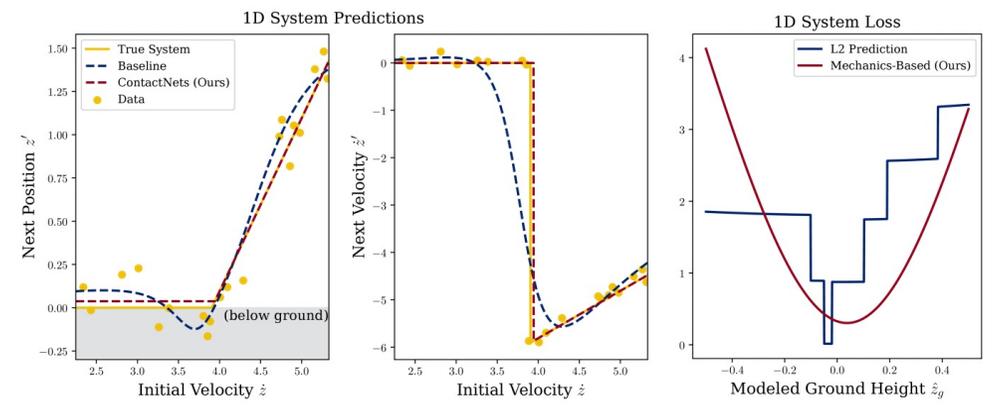
Tomographic Reconstruction [Ruckert et al. 2022]



Astronomical Interferometry [Wu et al. 2021]



(a) 1D System

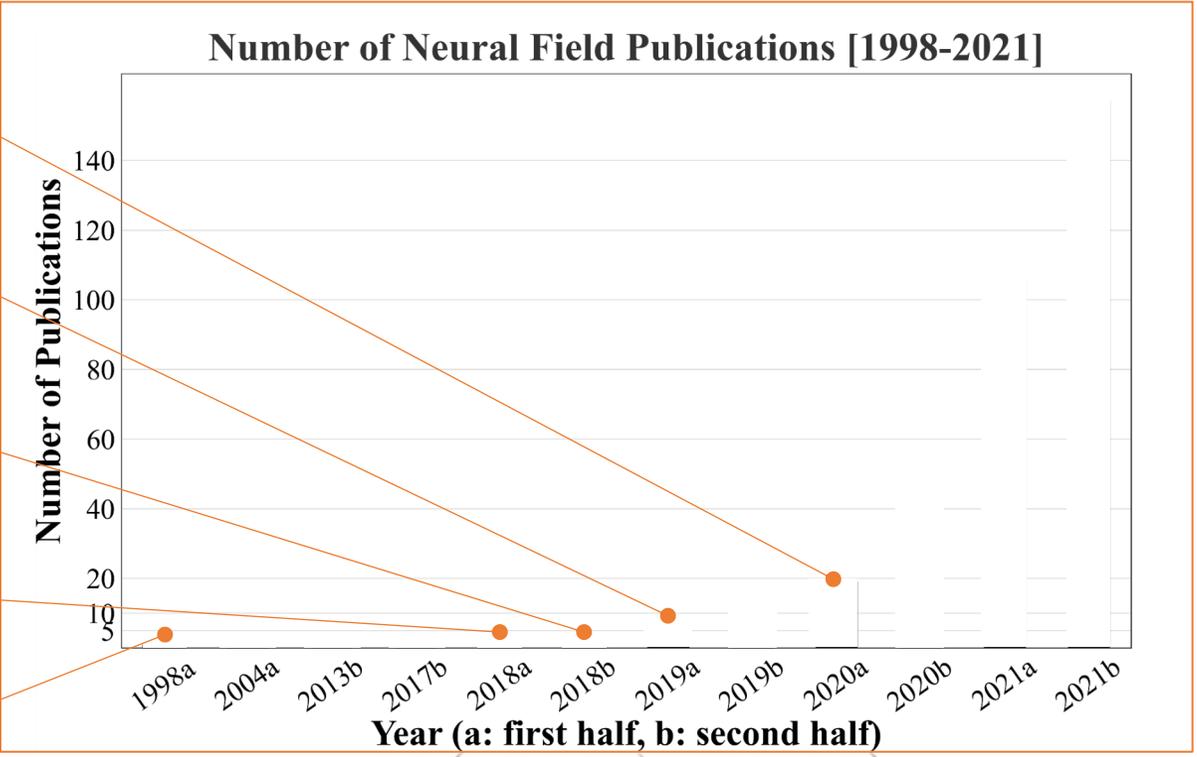
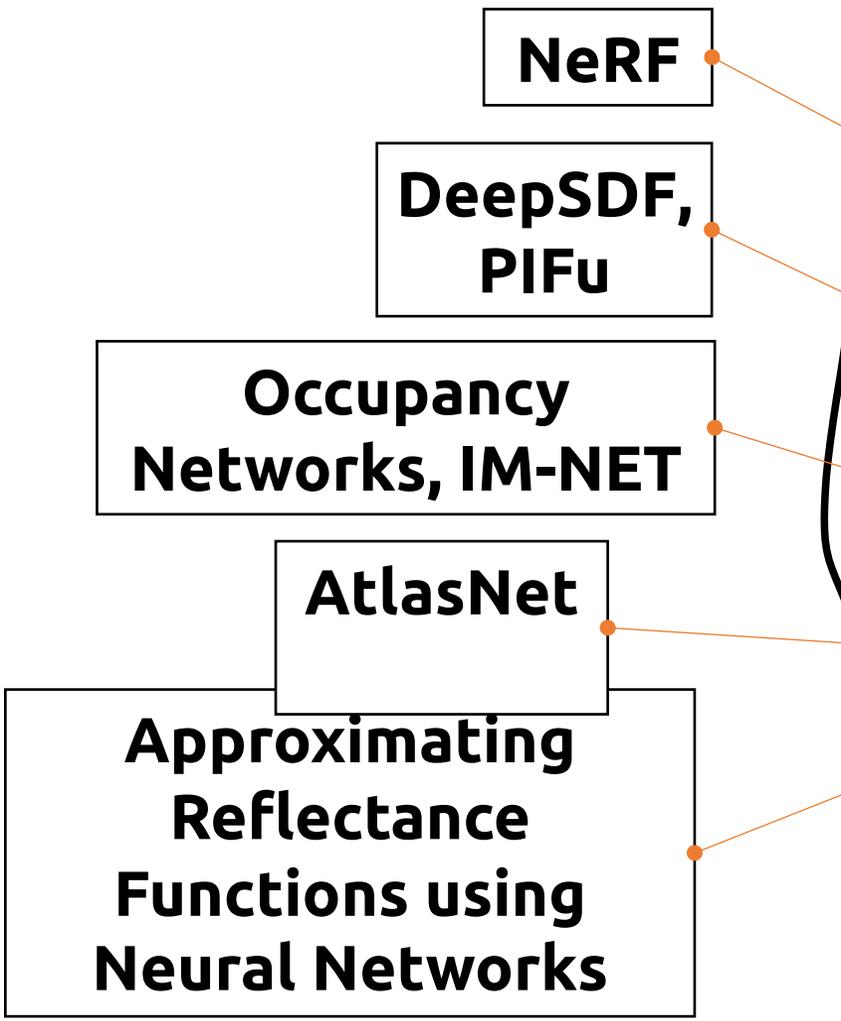


(b) Model Predictions

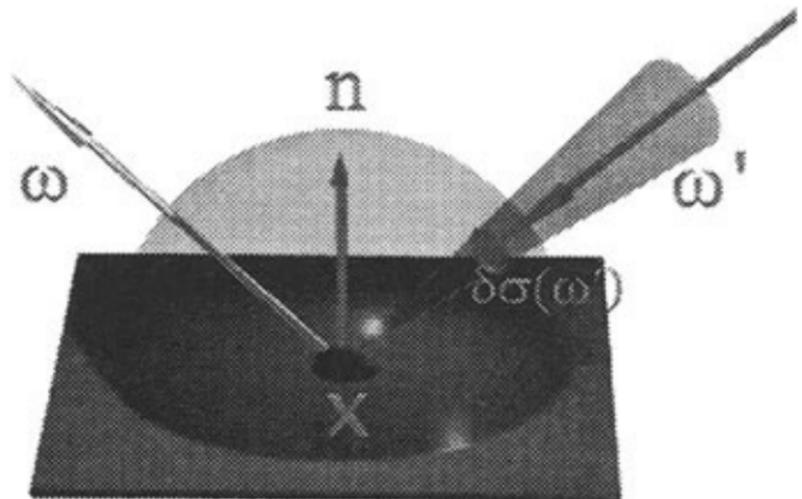
(c) Loss Landscape

Contact Dynamics [Pfrommer, Halm et al. 2022]

# The “Cambrian Explosion” of Neural Fields

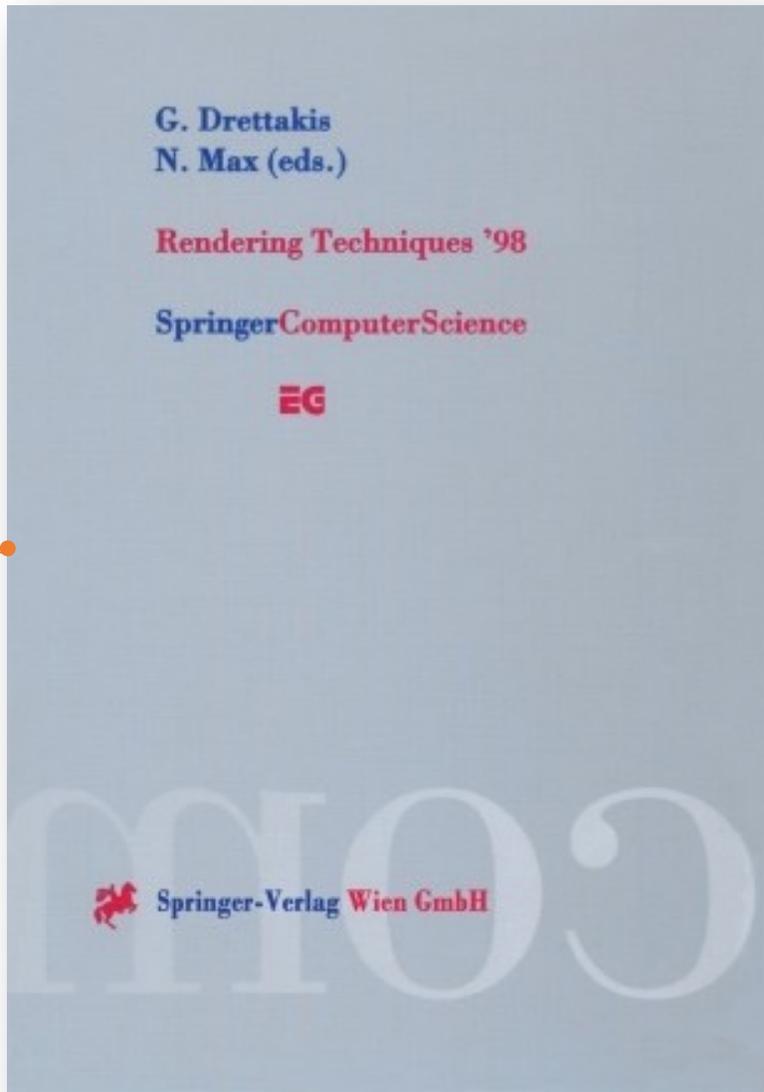


# The “Cambrian Explosion” of Neural Fields



[Gargan and Neelamkavil 1998]

**Approximating  
Reflectance  
Functions using  
Neural Networks**





# Outline

- What is Neural Fields & why it got so much attention?
- **The Prelude: Neural Implicit Surfaces**
- Introduction to Volume Rendering

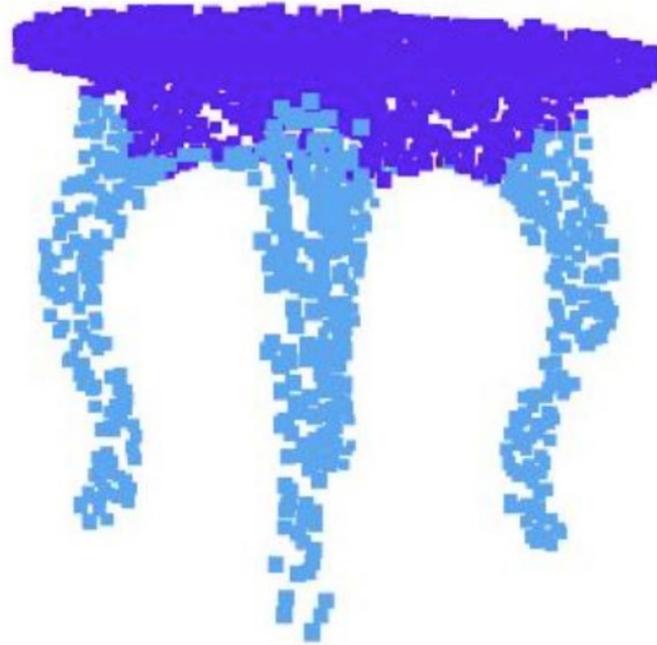
# Representations for 3D Deep Learning

Voxel



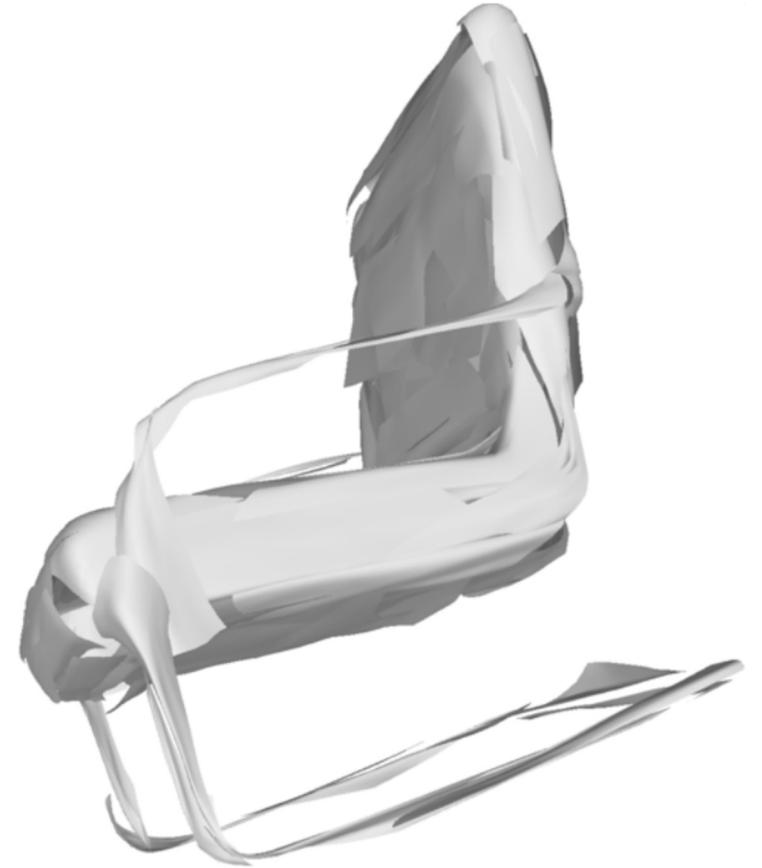
Wu et al. 2016

Points



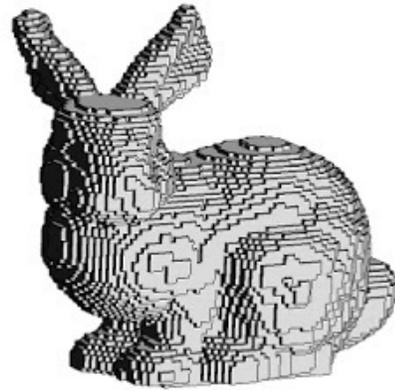
Qi et al. 2017

Meshes

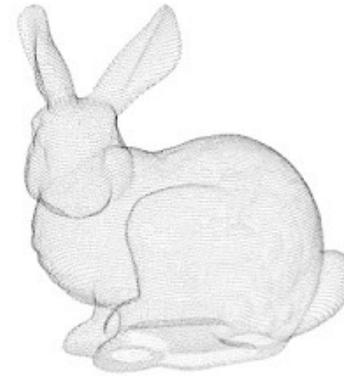


Groueix et al. 2018

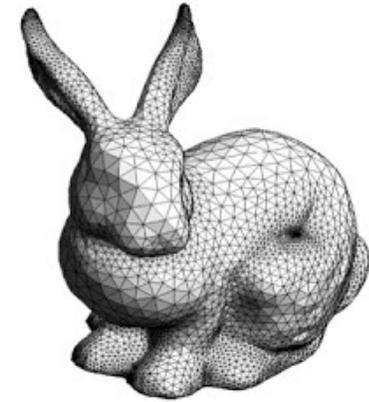
# 3D Representations (Explicit)



**Voxel**



**Point cloud**



**Polygon mesh**

Memory efficiency

Poor

Not good

Good

Textures

Not good

No

Yes

For neural networks

Easy

Not easy

Not easy

We adopt **polygon mesh** for its high potential

Images are from

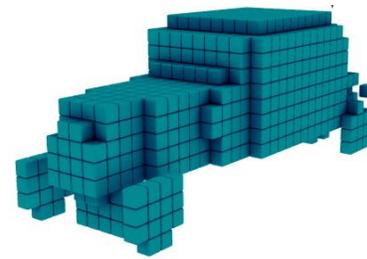
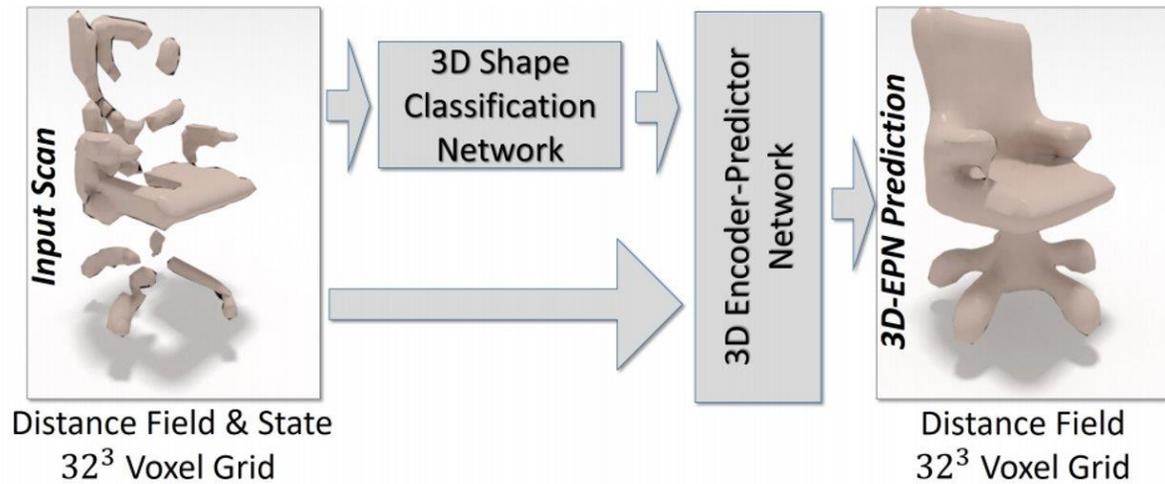
<http://cse.iitkgp.ac.in/~pb/research/3dpoly/3dpoly.html>

<http://waldyrrious.net/learning-holography/pb-cgh-formulas.xhtml>

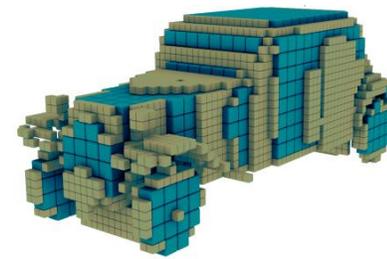
<http://www.cs.mun.ca/~omeruvia/philosophy/images/BunnyWire.gif>

# Voxel Representation

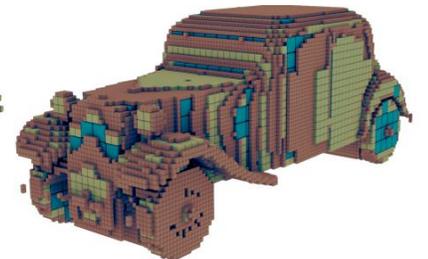
- Memory Intensive, Computationally Expensive ( $N^3$ )



$32^3$



$64^3$

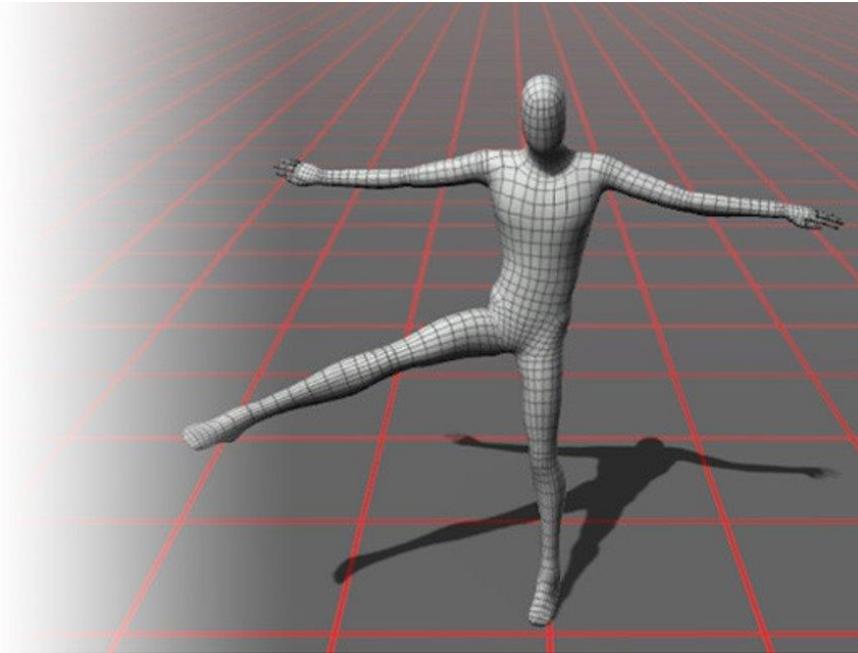


$128^3$



# Mesh Representation

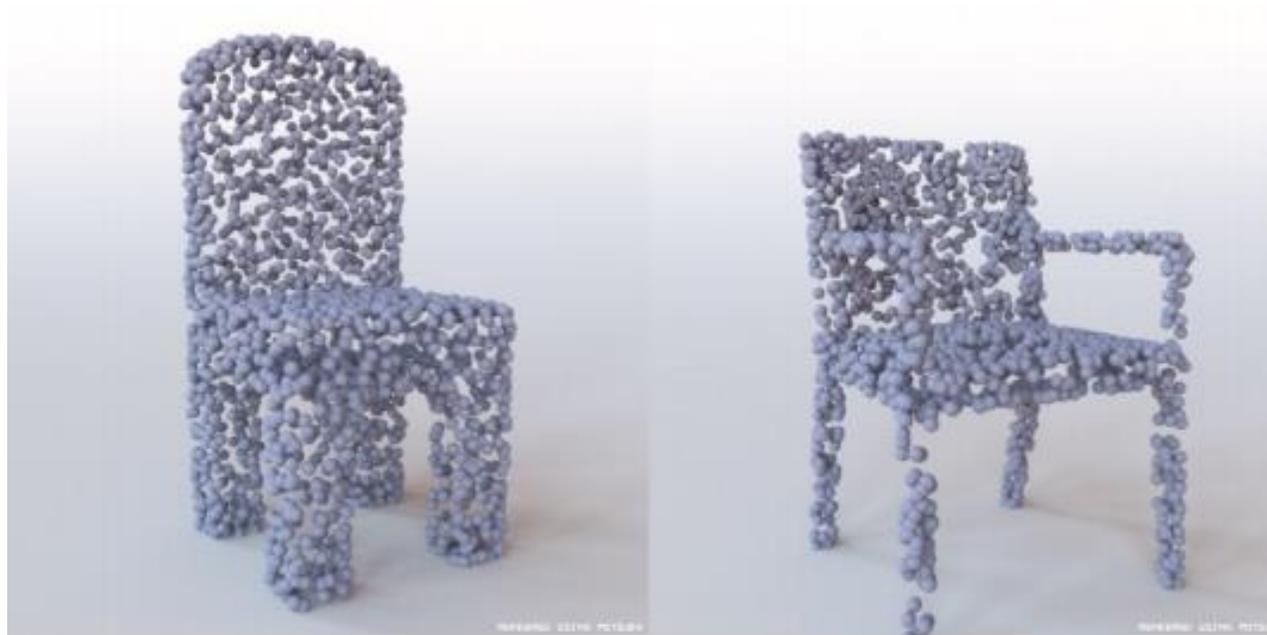
- Fixed Topologies (relies on Templates)
- Combinatorial Problem → Discrete Vertices and Connections



Groueix et al. 2018

# Point Cloud Representation

- Does not Define a Surface
- Not suitable for Visualization, Texturing, etc

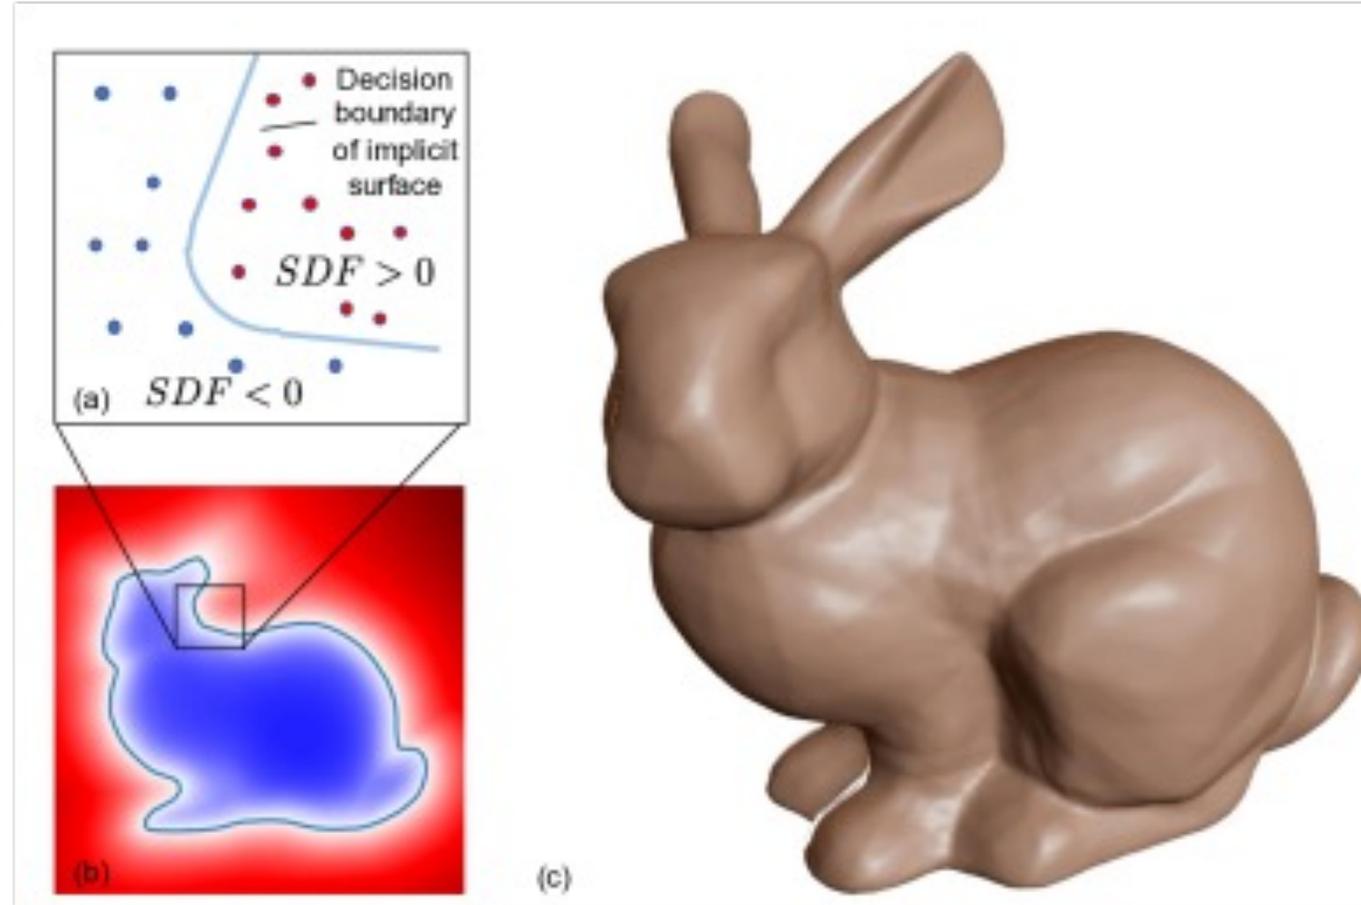


Achlioptas et al.

# Surface Representation: Signed Distance Function (SDF) - implicit representation via level set

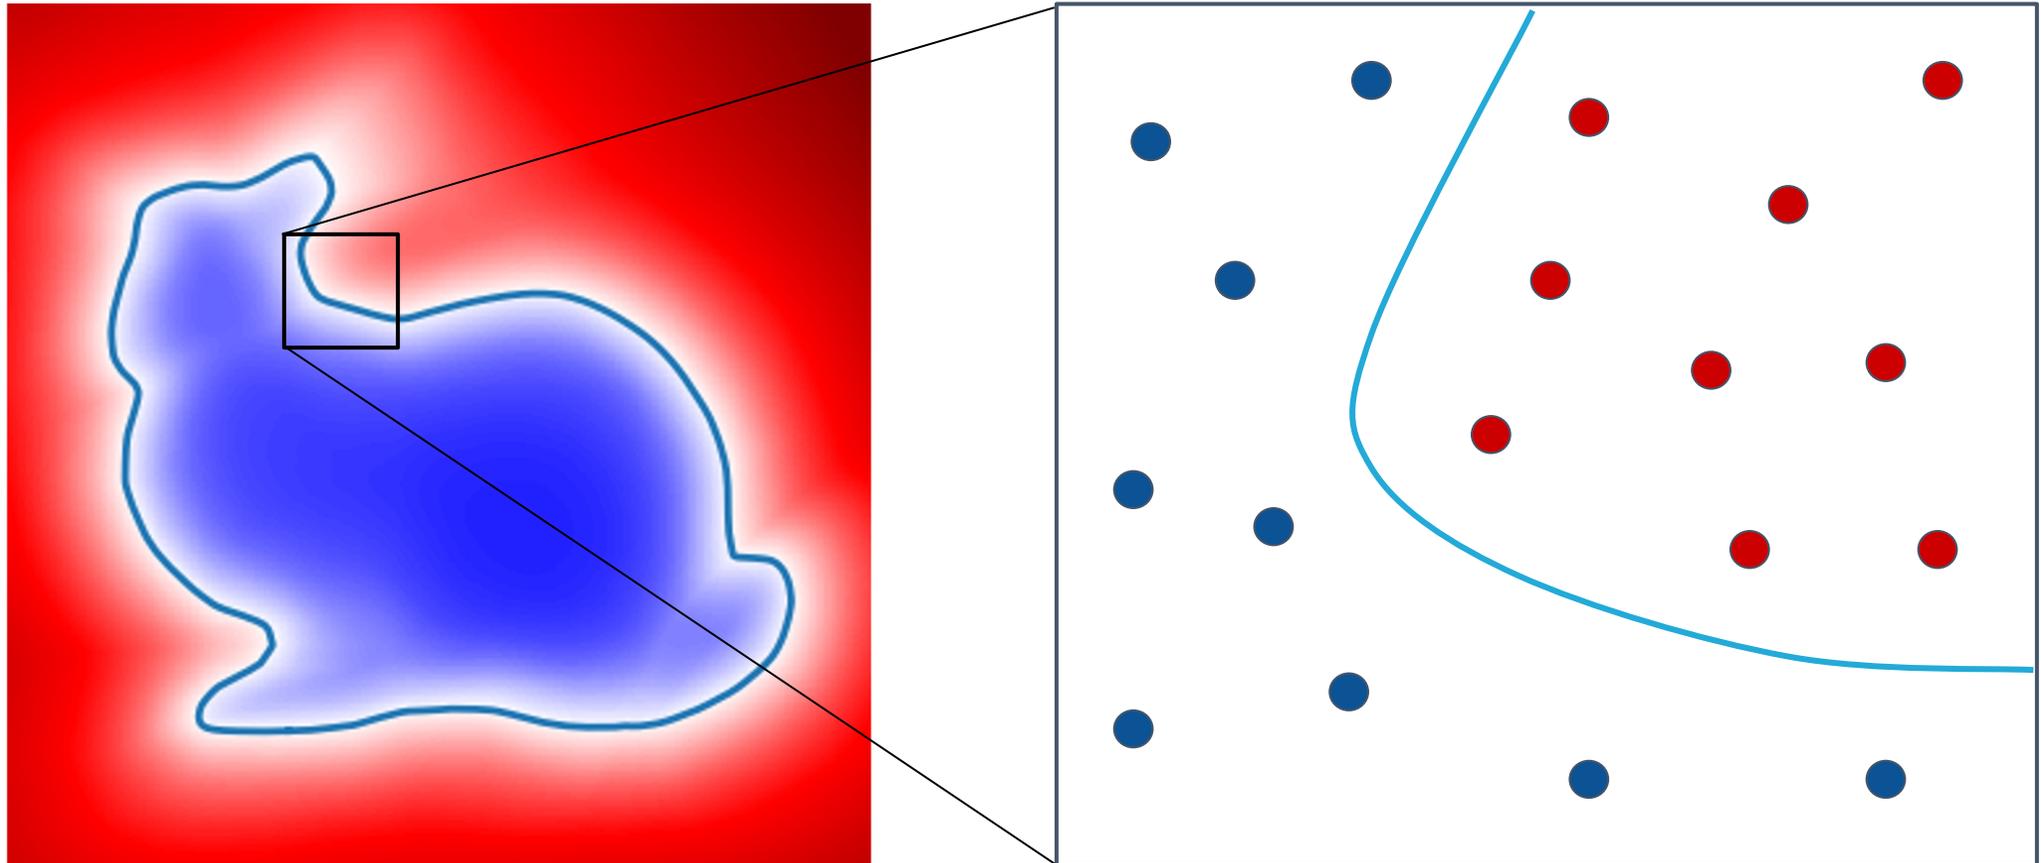
$SDF(X) = 0$ , when  $X$  is on the surface.  
 $SDF(X) > 0$ , when  $X$  is outside the surface  
 $SDF(X) < 0$ , when  $X$  is inside the surface

Note: SDF is an implicit representation!  
Suitable for neural networks but hard to  
import inside existing graphics software.

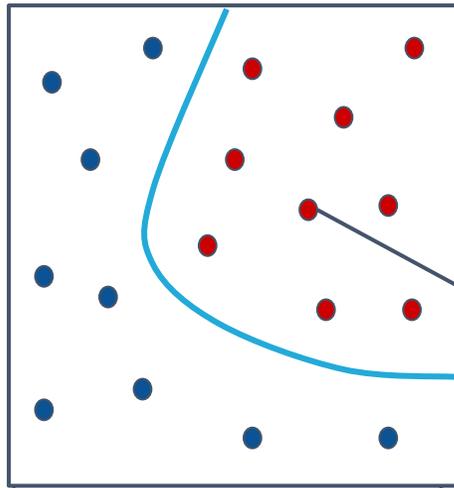


**Deep SDF: Use a neural network (co-ordinate based MLP) to represent the SDF function.**

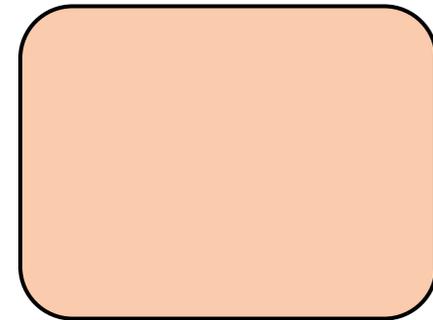
# Surface as Decision Boundary



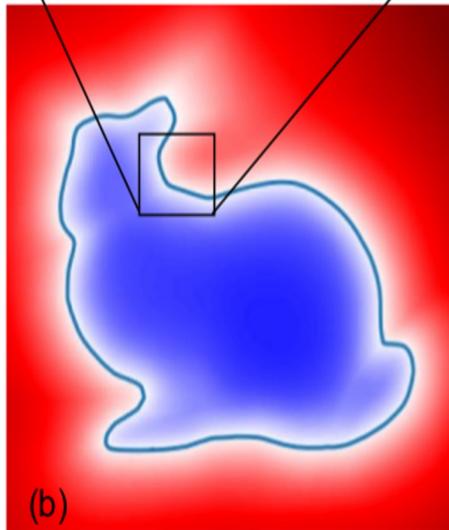
# Regression of Continuous SDF



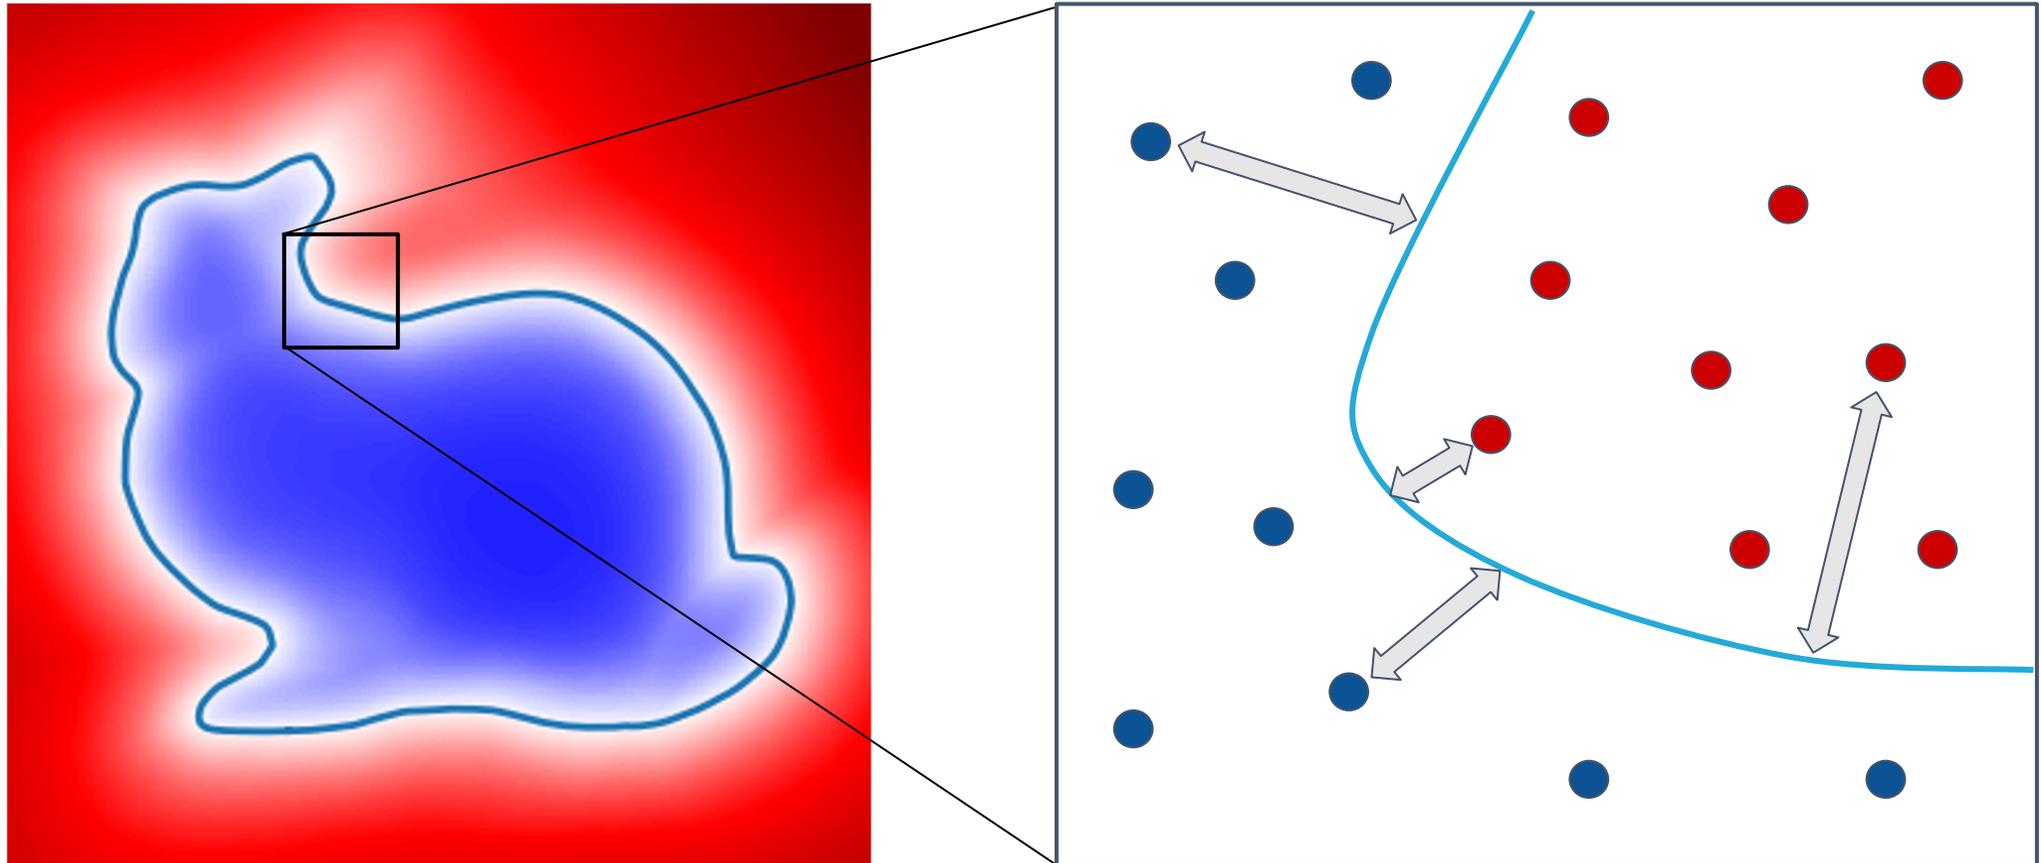
$(x, y, z)$



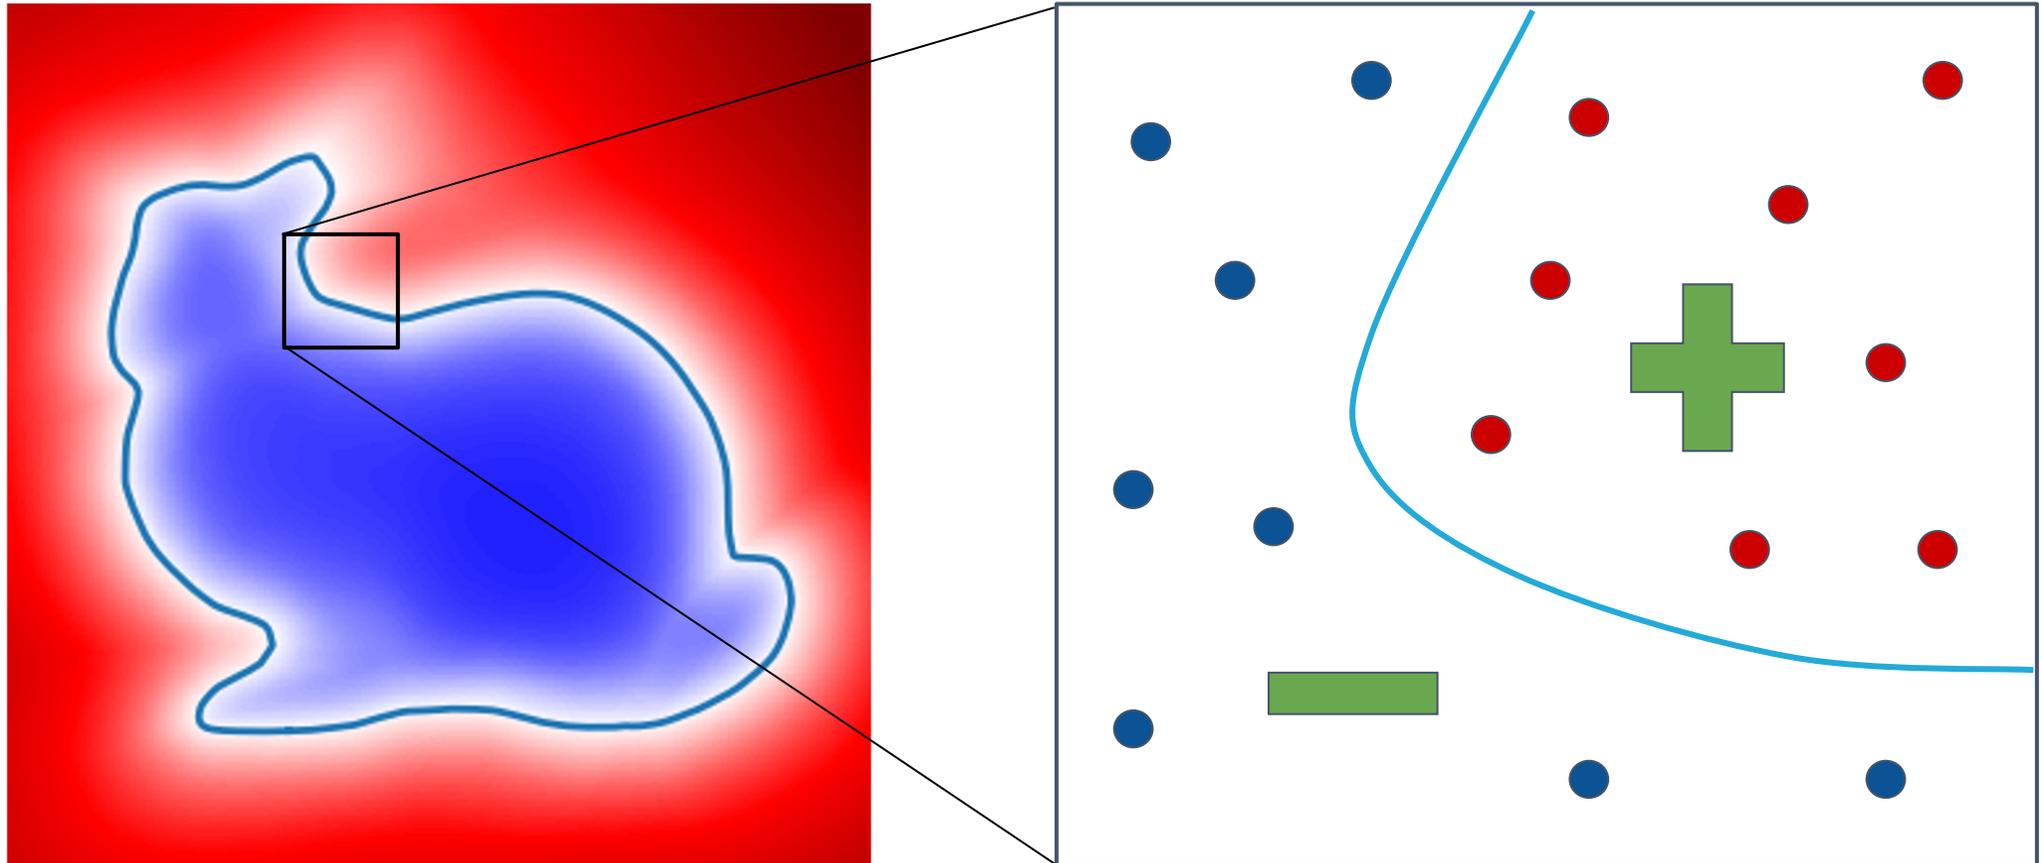
SDF



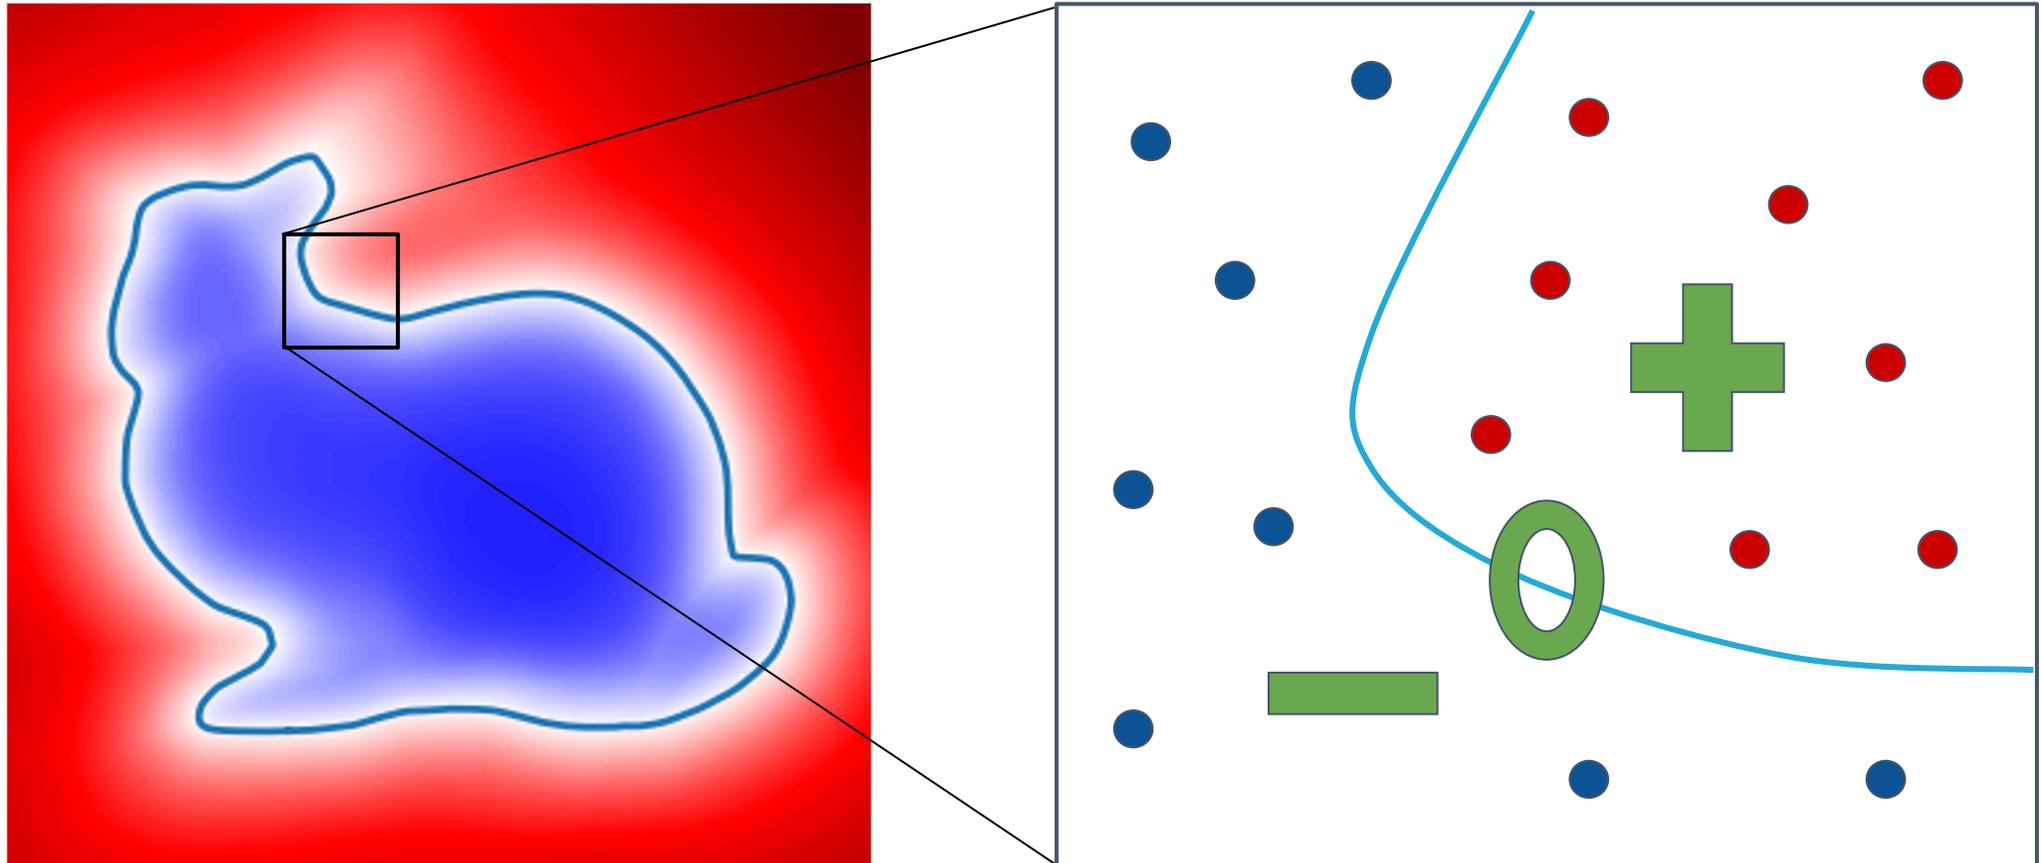
# Signed Distance Function



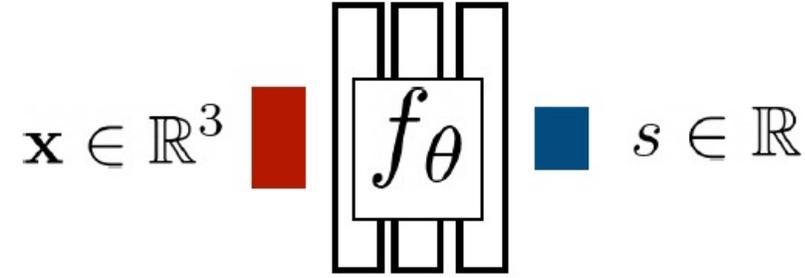
# Signed Distance Function



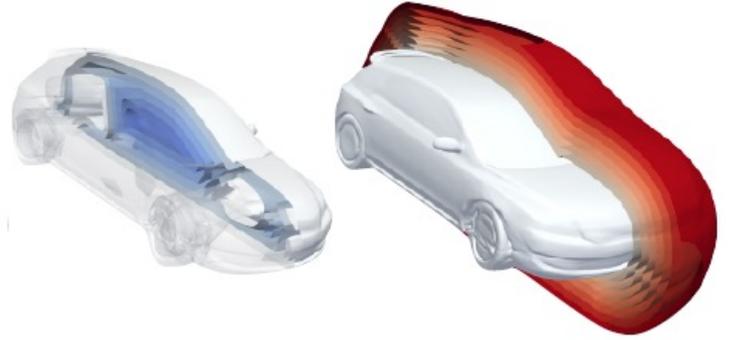
# Signed Distance Function



# Instance-specific SDFs



$$f_\theta : \mathbb{R}^3 \rightarrow \mathbb{R}$$



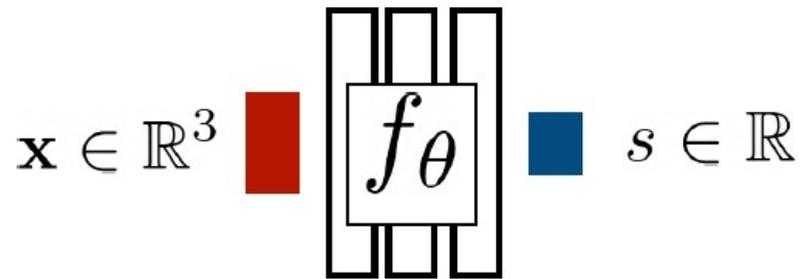
**Signed Distance Field (for a single instance):**

(position)  $\rightarrow$  (distance)

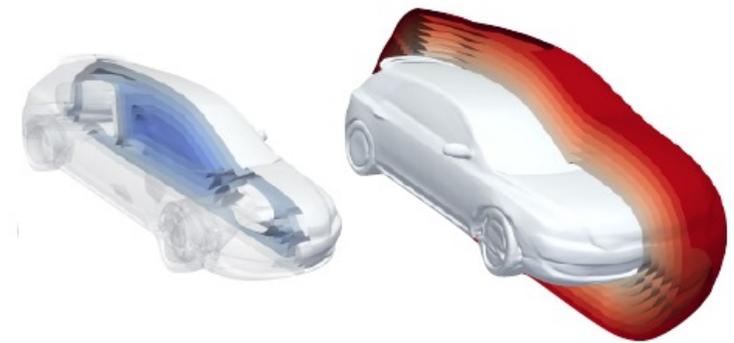
if 6 layer network with 1000-dim feature space, about 6M parameters per instance!

# Questions we want to answer

- How do we create a 3D mesh from a SDF function? (SDF rendering)
- How do we generalize this to any objects? (Training Deep SDF)
- How do we use this during inference? (Testing DeepSDF)
- How do DeepSDF concept extends to NeRF and other methods?



$$f_\theta : \mathbb{R}^3 \rightarrow \mathbb{R}$$

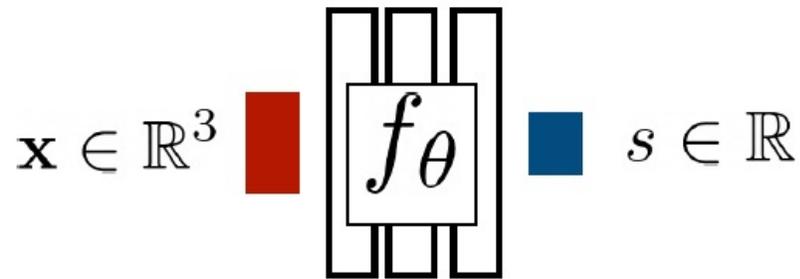


**Signed Distance Field (for a single instance):**

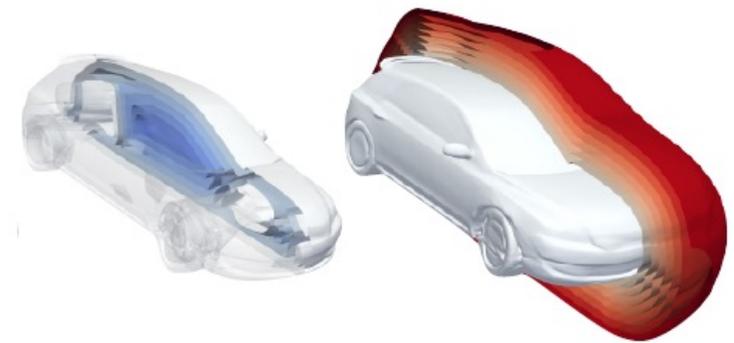
(position)  $\rightarrow$  (distance)

# Questions we want to answer

- How do we create a 3D mesh from a SDF function? (SDF rendering)
- How do we generalize this to any objects? (Training Deep SDF)
- How do we use this during inference? (Testing DeepSDF)
- How do DeepSDF concept extends to NeRF and other methods?



$$f_\theta : \mathbb{R}^3 \rightarrow \mathbb{R}$$



**Signed Distance Field (for a single instance):**

(position)  $\rightarrow$  (distance)

# Neural Implicit Surfaces

$$\{\mathbf{p} \mid f_{\theta}(\mathbf{p})\} = 0$$

Neural network with parameters  $\theta$

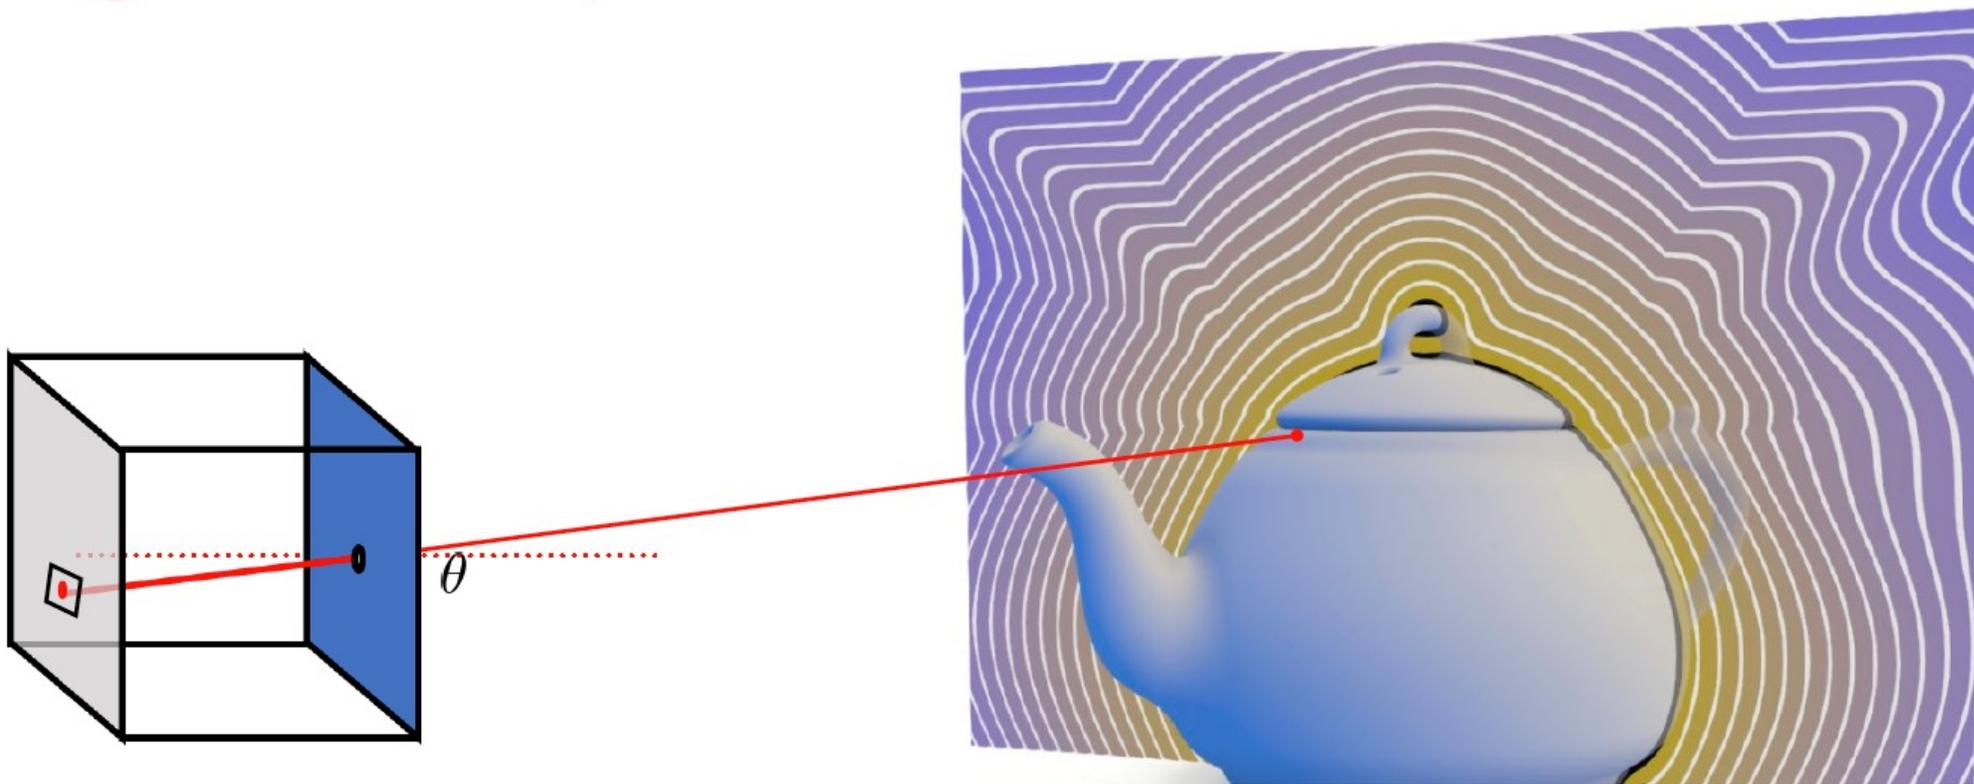
Allows representing complex geometry

Q: What additional constraint is required for a signed distance field?

$$\|\nabla f\| = 1$$



# Rendering Neural Implicit Surfaces

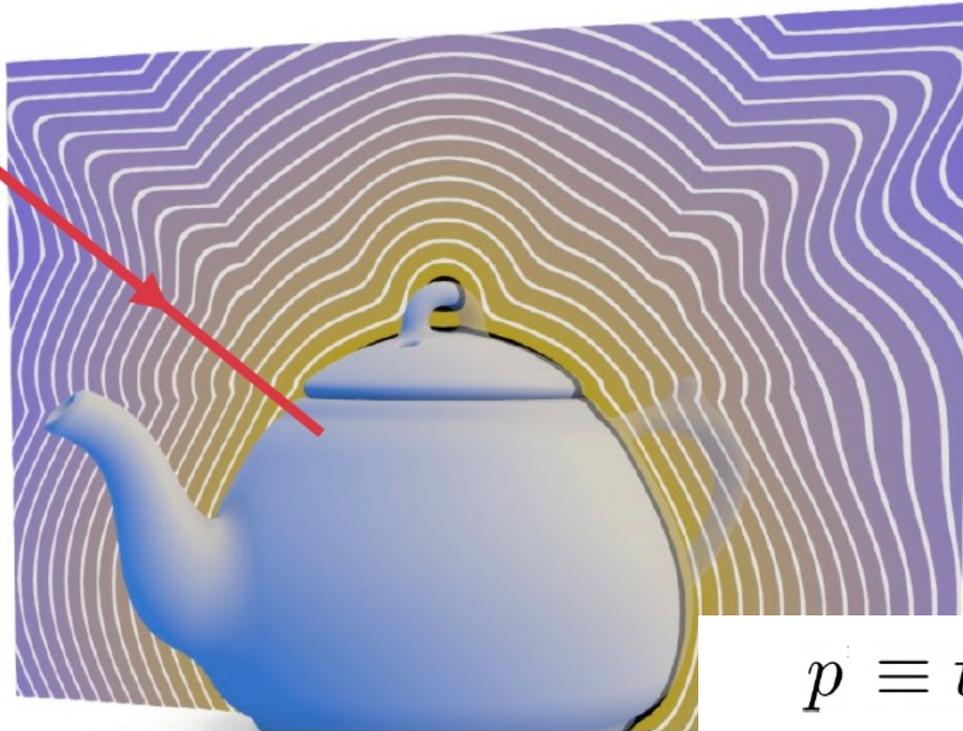


The appearance at each pixel is determined by a unique surface point

**Where does a ray intersect an implicit surface?**

**How to model appearance?**

# Intersections of Rays with Implicit Surfaces



## Input:

Implicit Representation  $f_{\theta}(\cdot)$

Query Ray  $r : x_0 + td$

## Output:

Intersection point:  $p$

How to solve this?

- Use root-finding methods (e.g. secant method)

$$p \equiv t^*d + x_0$$

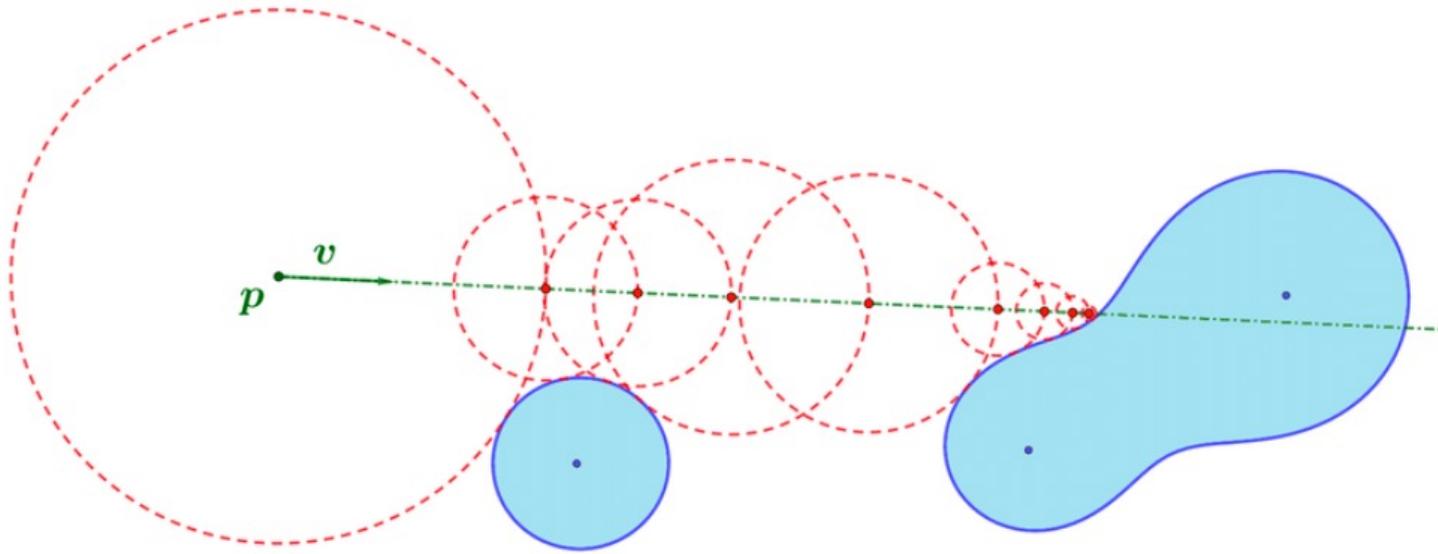
$$\text{s.t. } f_{\theta}(p) = 0$$

$$\bar{f}_{\theta}(t) \equiv f_{\theta}(x_0 + td)$$

**a 1-d search problem!**

$d$  = unit vector denoting the direction of the ray.  
 $t$  = scalar distance between any point  $p$  on the ray and the origin  $x_0$ .

# Intersections of Rays with SDFs: Sphere Tracing



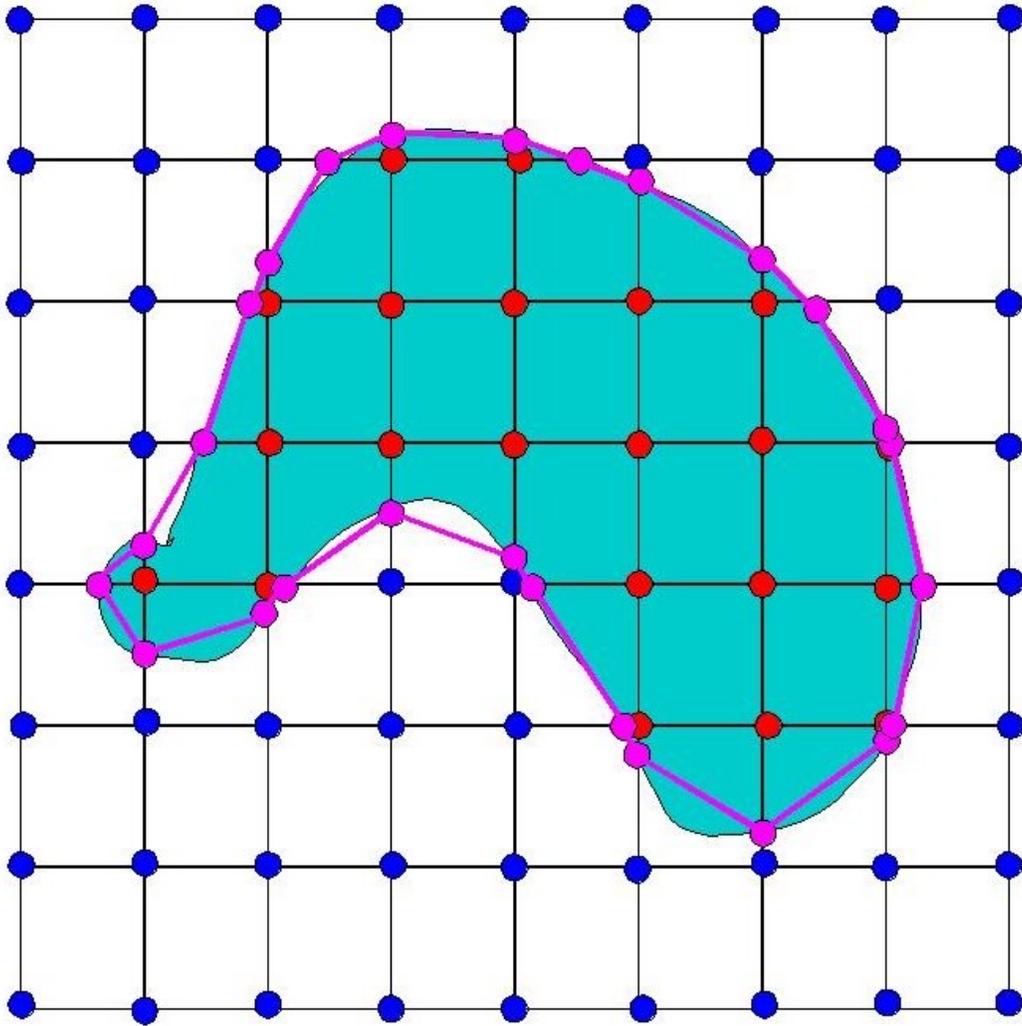
While  $f(p) > \epsilon$  :

$$t = t + f(p)$$

$$p = x_0 + td$$

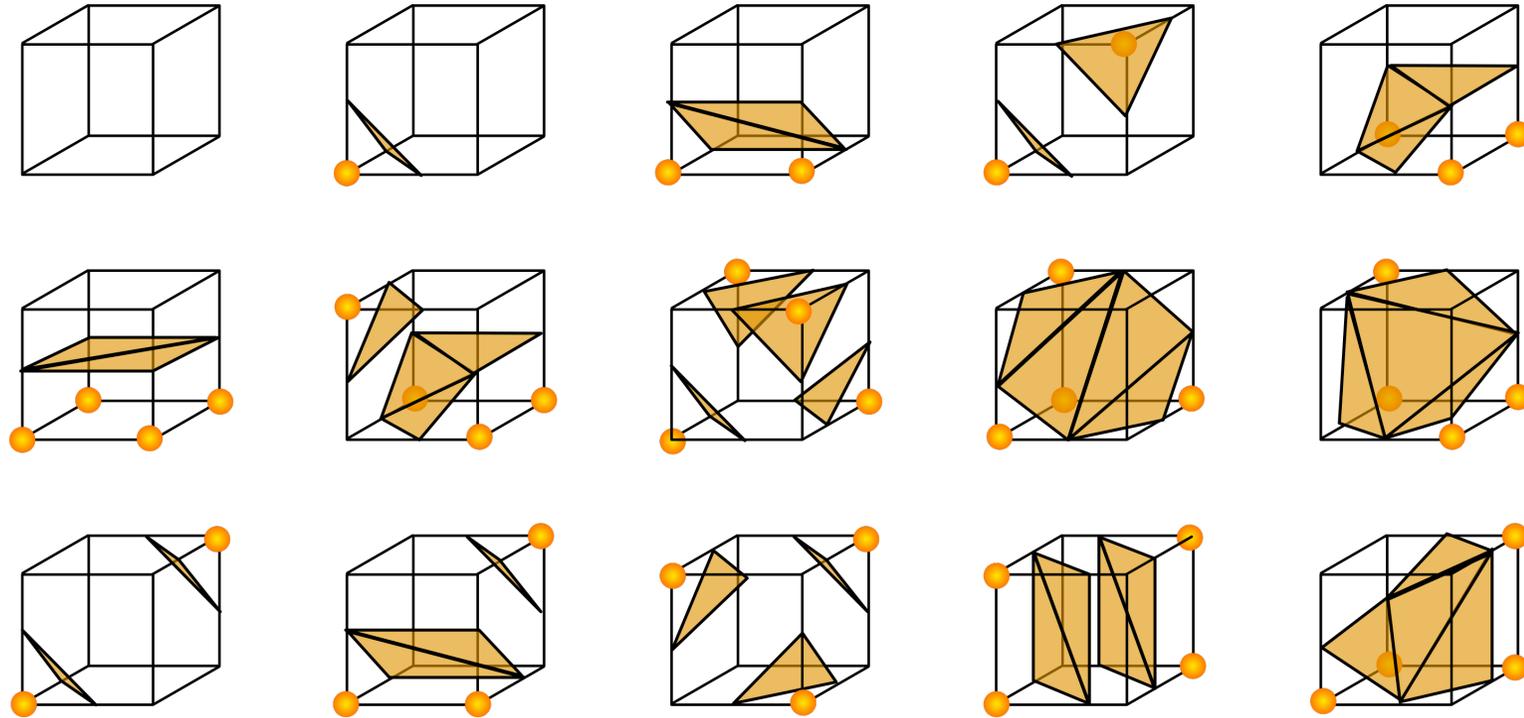
Basically, walk along the ray until close to a surface (or we exceed max-steps)

# Marching Cubes in 2D



- Each vertex of a cell can be labelled +ve or -ve, so total  $2^4$  possibilities.
- After labeling each vertex draw boundary between +ve and -ve.
- Refine the boundary.

# Marching Cubes in 3D



In 3D coloring vertex +ve/-ve has  $2^8$  possibilities. Brute force search is expensive.

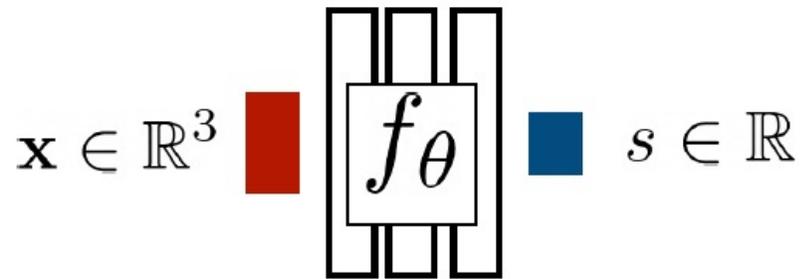
However, we can obtain only 15 unique possibilities from  $2^8$

# How do we render SDF into a mesh?

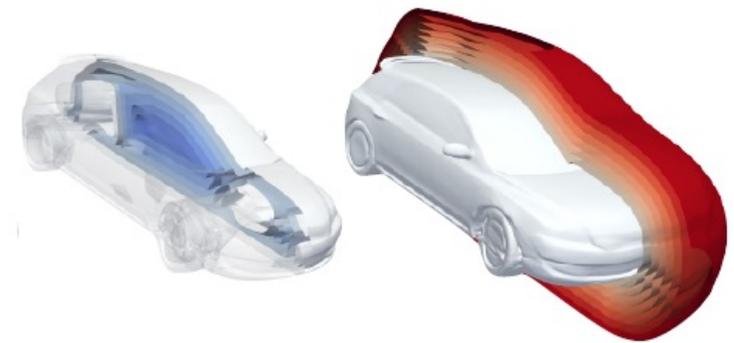
- Secant Method (Find roots of a 1-d search problem along each ray)
- Sphere Tracing
- Marching Cubes

# Questions we want to answer

- How do we create a 3D mesh from a SDF function? (SDF rendering)
- How do we generalize this to any objects? (Training Deep SDF)
- How do we use this during inference? (Testing DeepSDF)
- How do DeepSDF concept extends to NeRF and other methods?



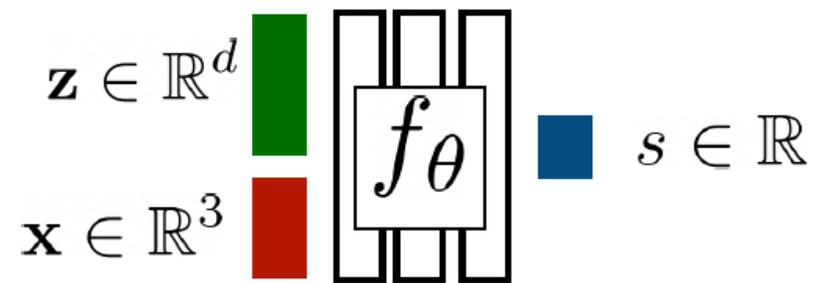
$$f_\theta : \mathbb{R}^3 \rightarrow \mathbb{R}$$



**Signed Distance Field (for a single instance):**

(position)  $\rightarrow$  (distance)

# Latent Conditioning based SDFs



$$f_\theta : \mathbb{R}^3 \times \mathbb{R}^d \rightarrow \mathbb{R}$$

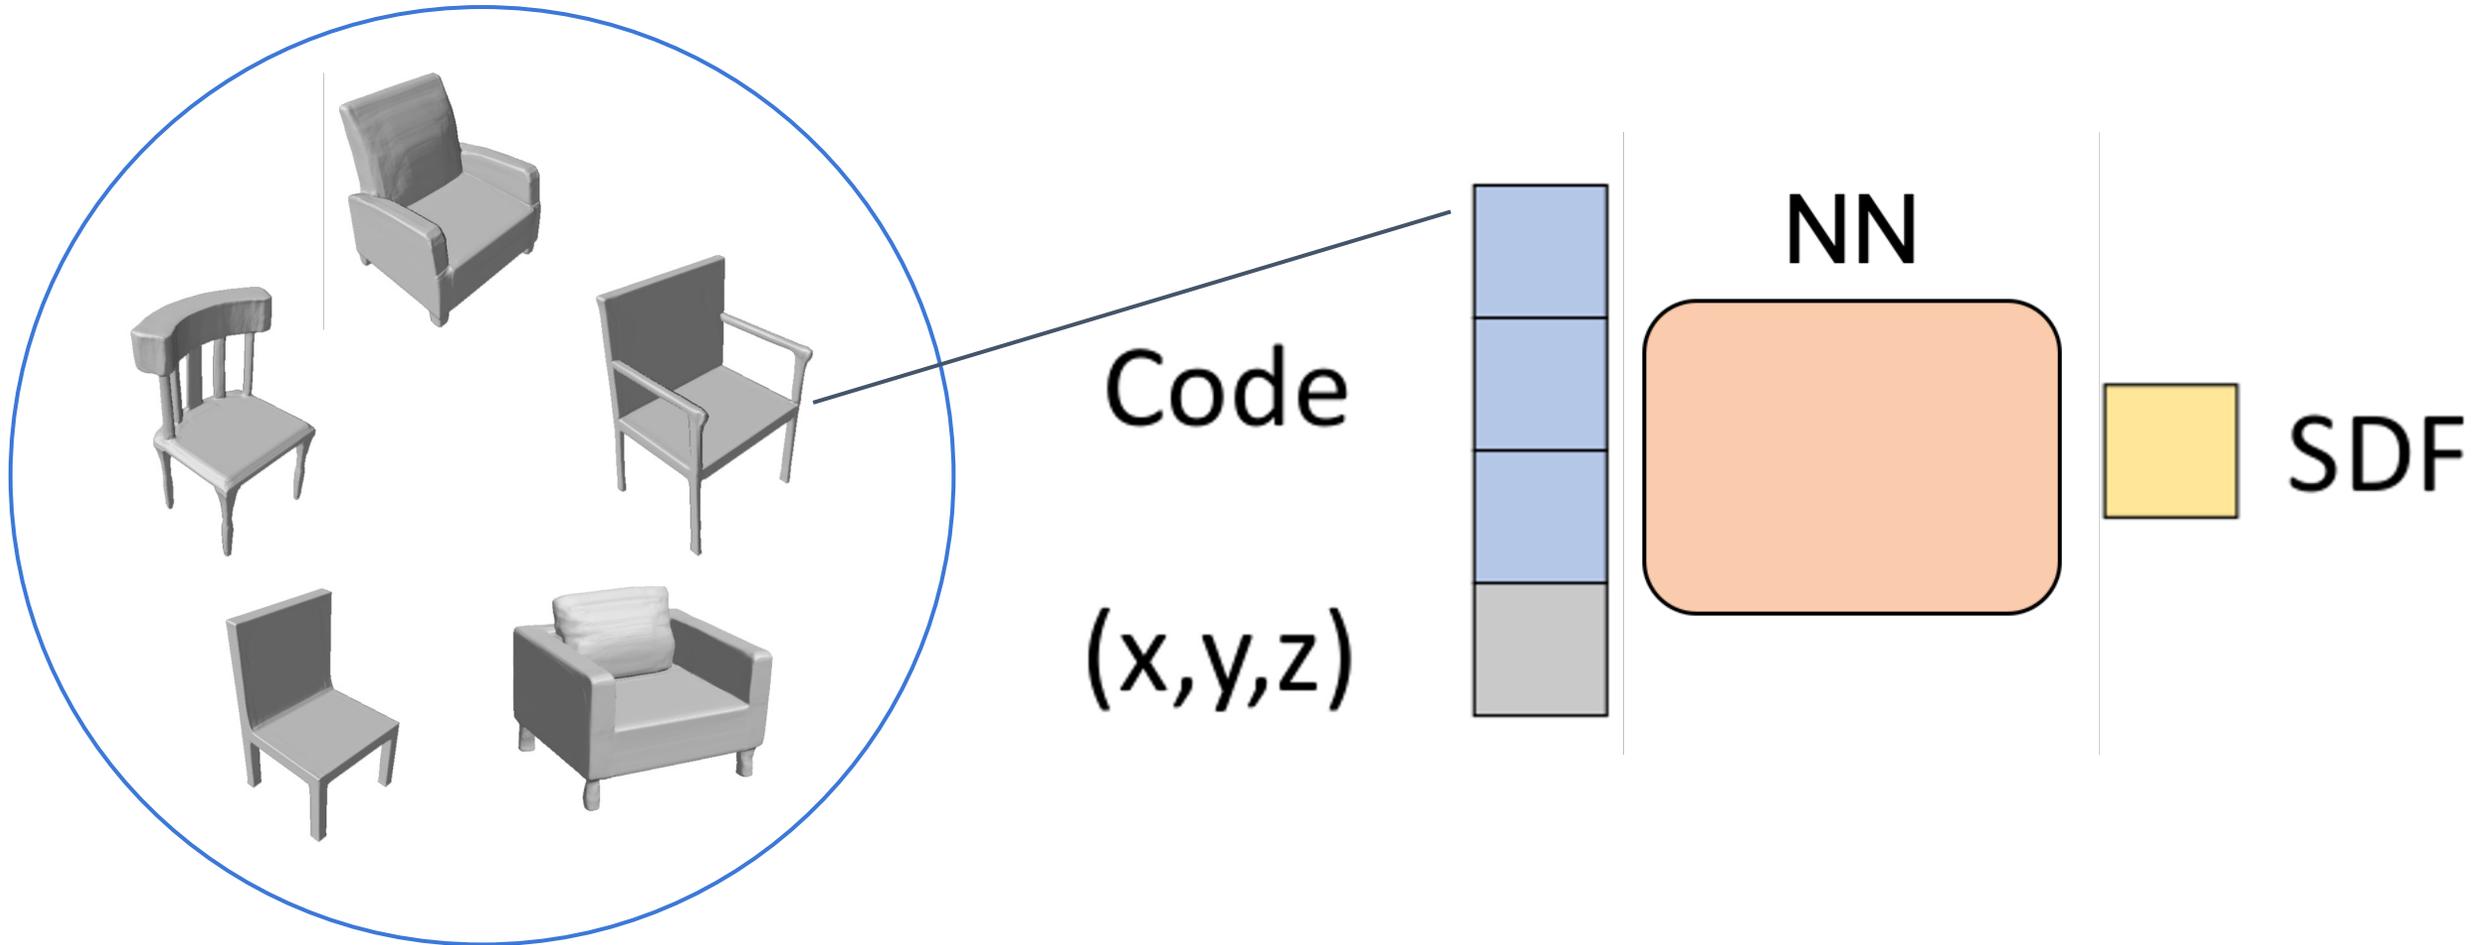
**Generalizable Signed Distance Field:**

(latent code, position)  $\rightarrow$  (distance)

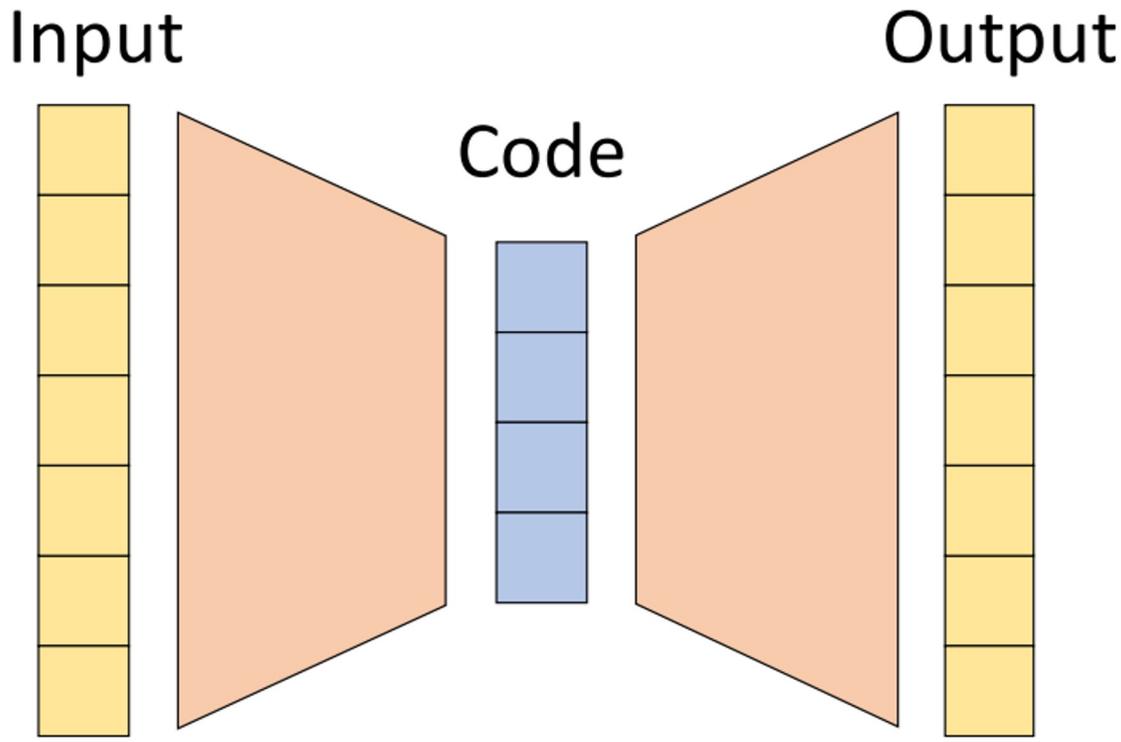
Each object is represented by a corresponding latent code (only  $d$  parameters per instance)

The same neural net parameters across **all** objects

# Coding Multiple Shapes

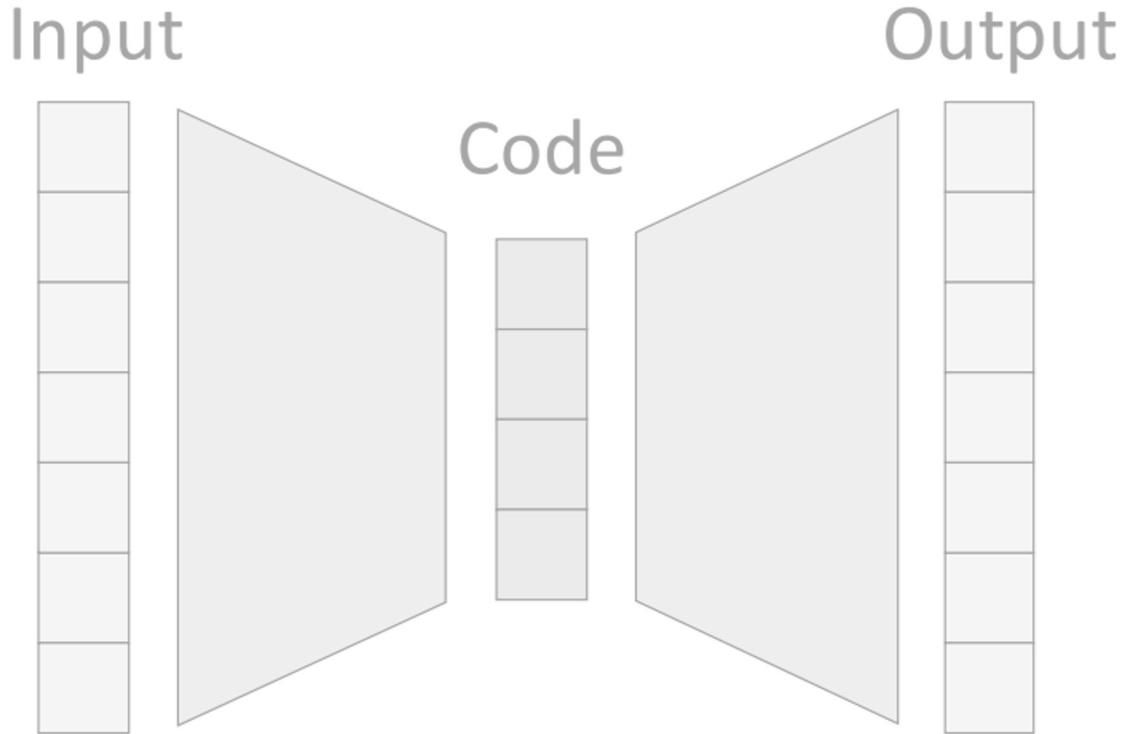


# Auto-Encoder

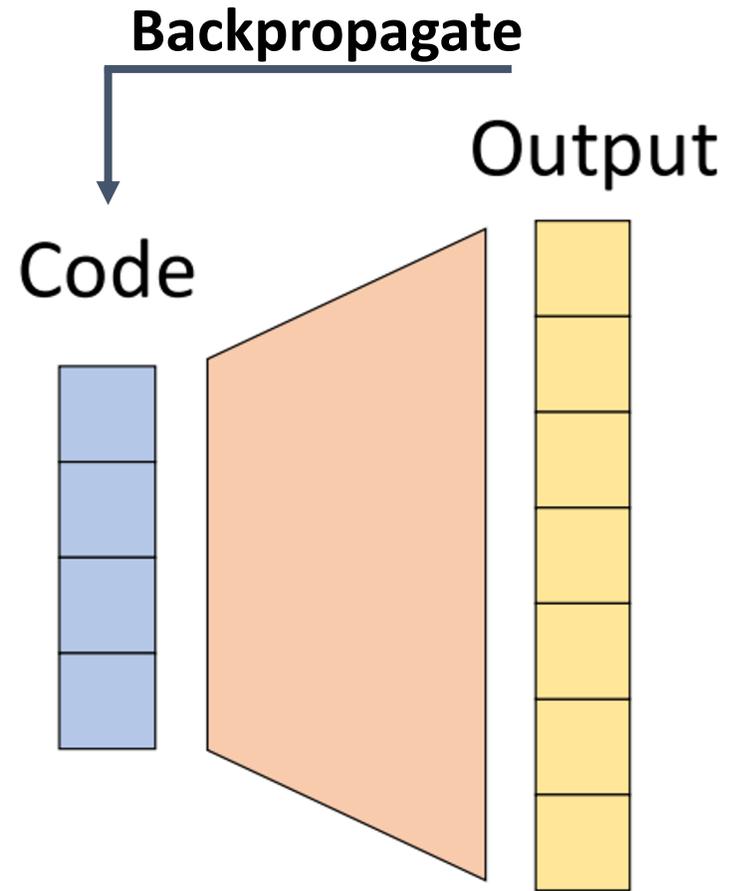


**Auto-Encoder**

# Auto-Decoder

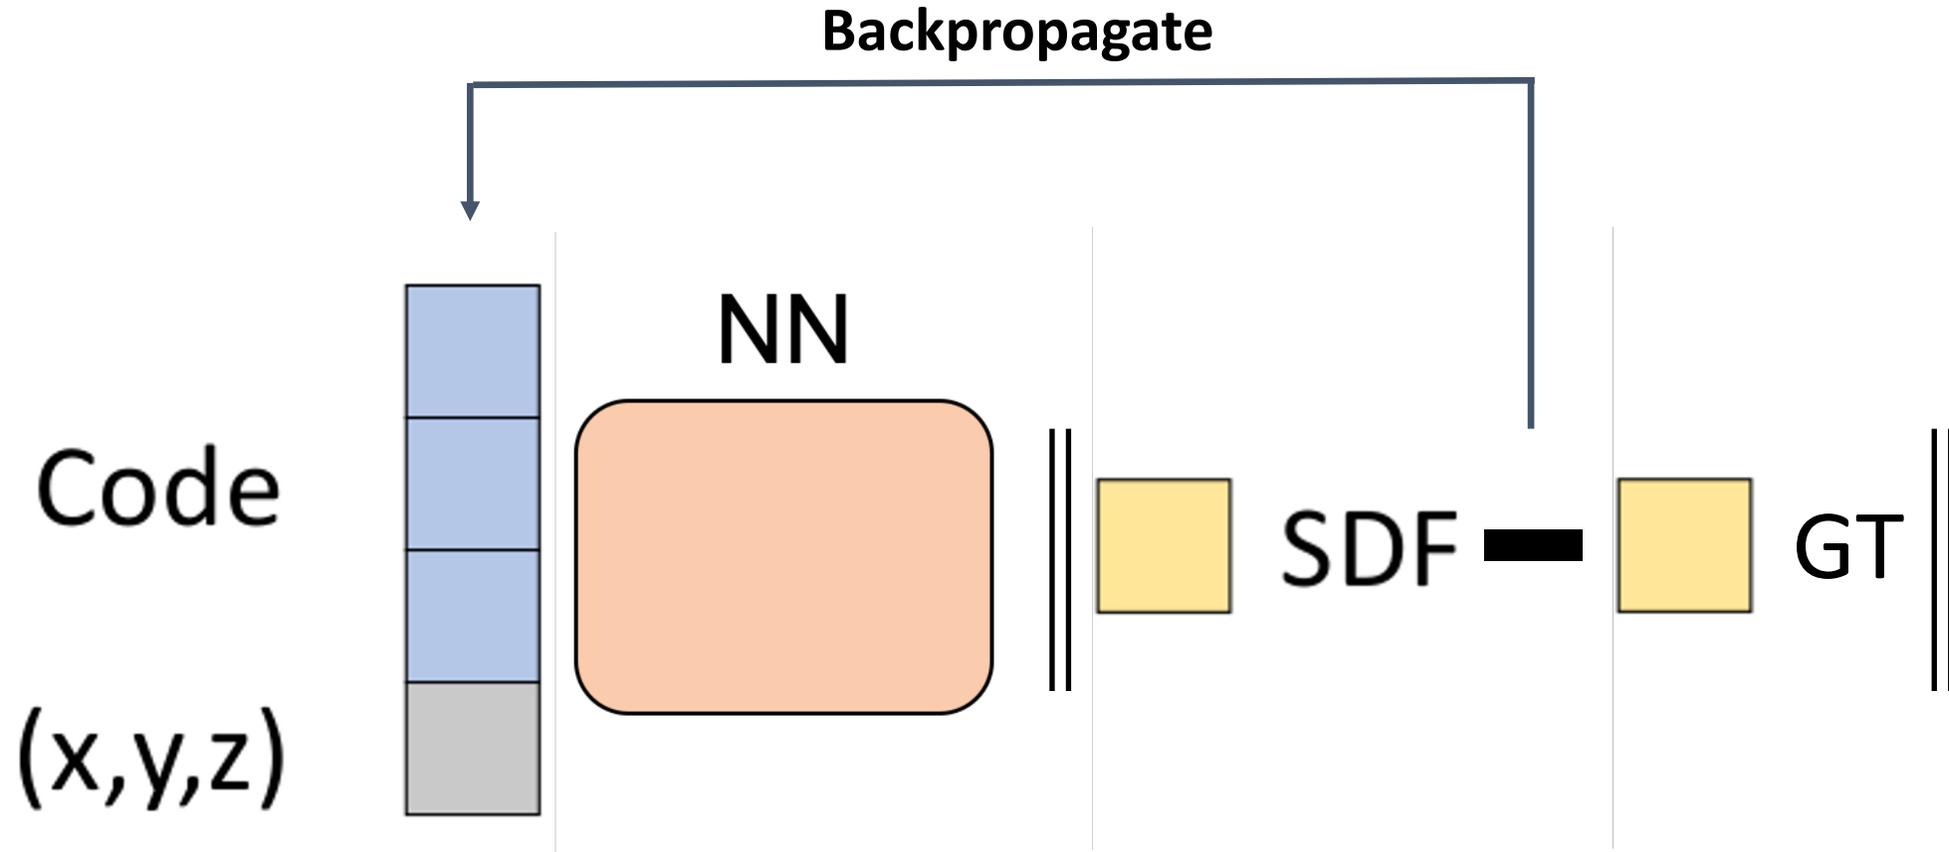


**Auto-Encoder**



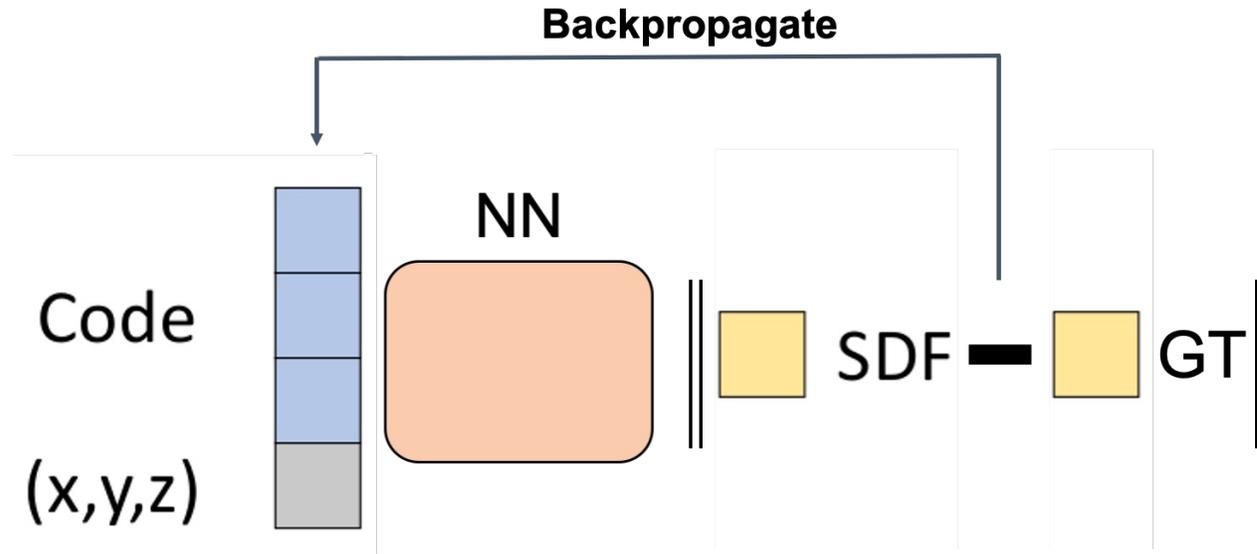
**Auto-Decoder**

# Auto-Decoder



During Training: Optimize for both NN parameters and Code

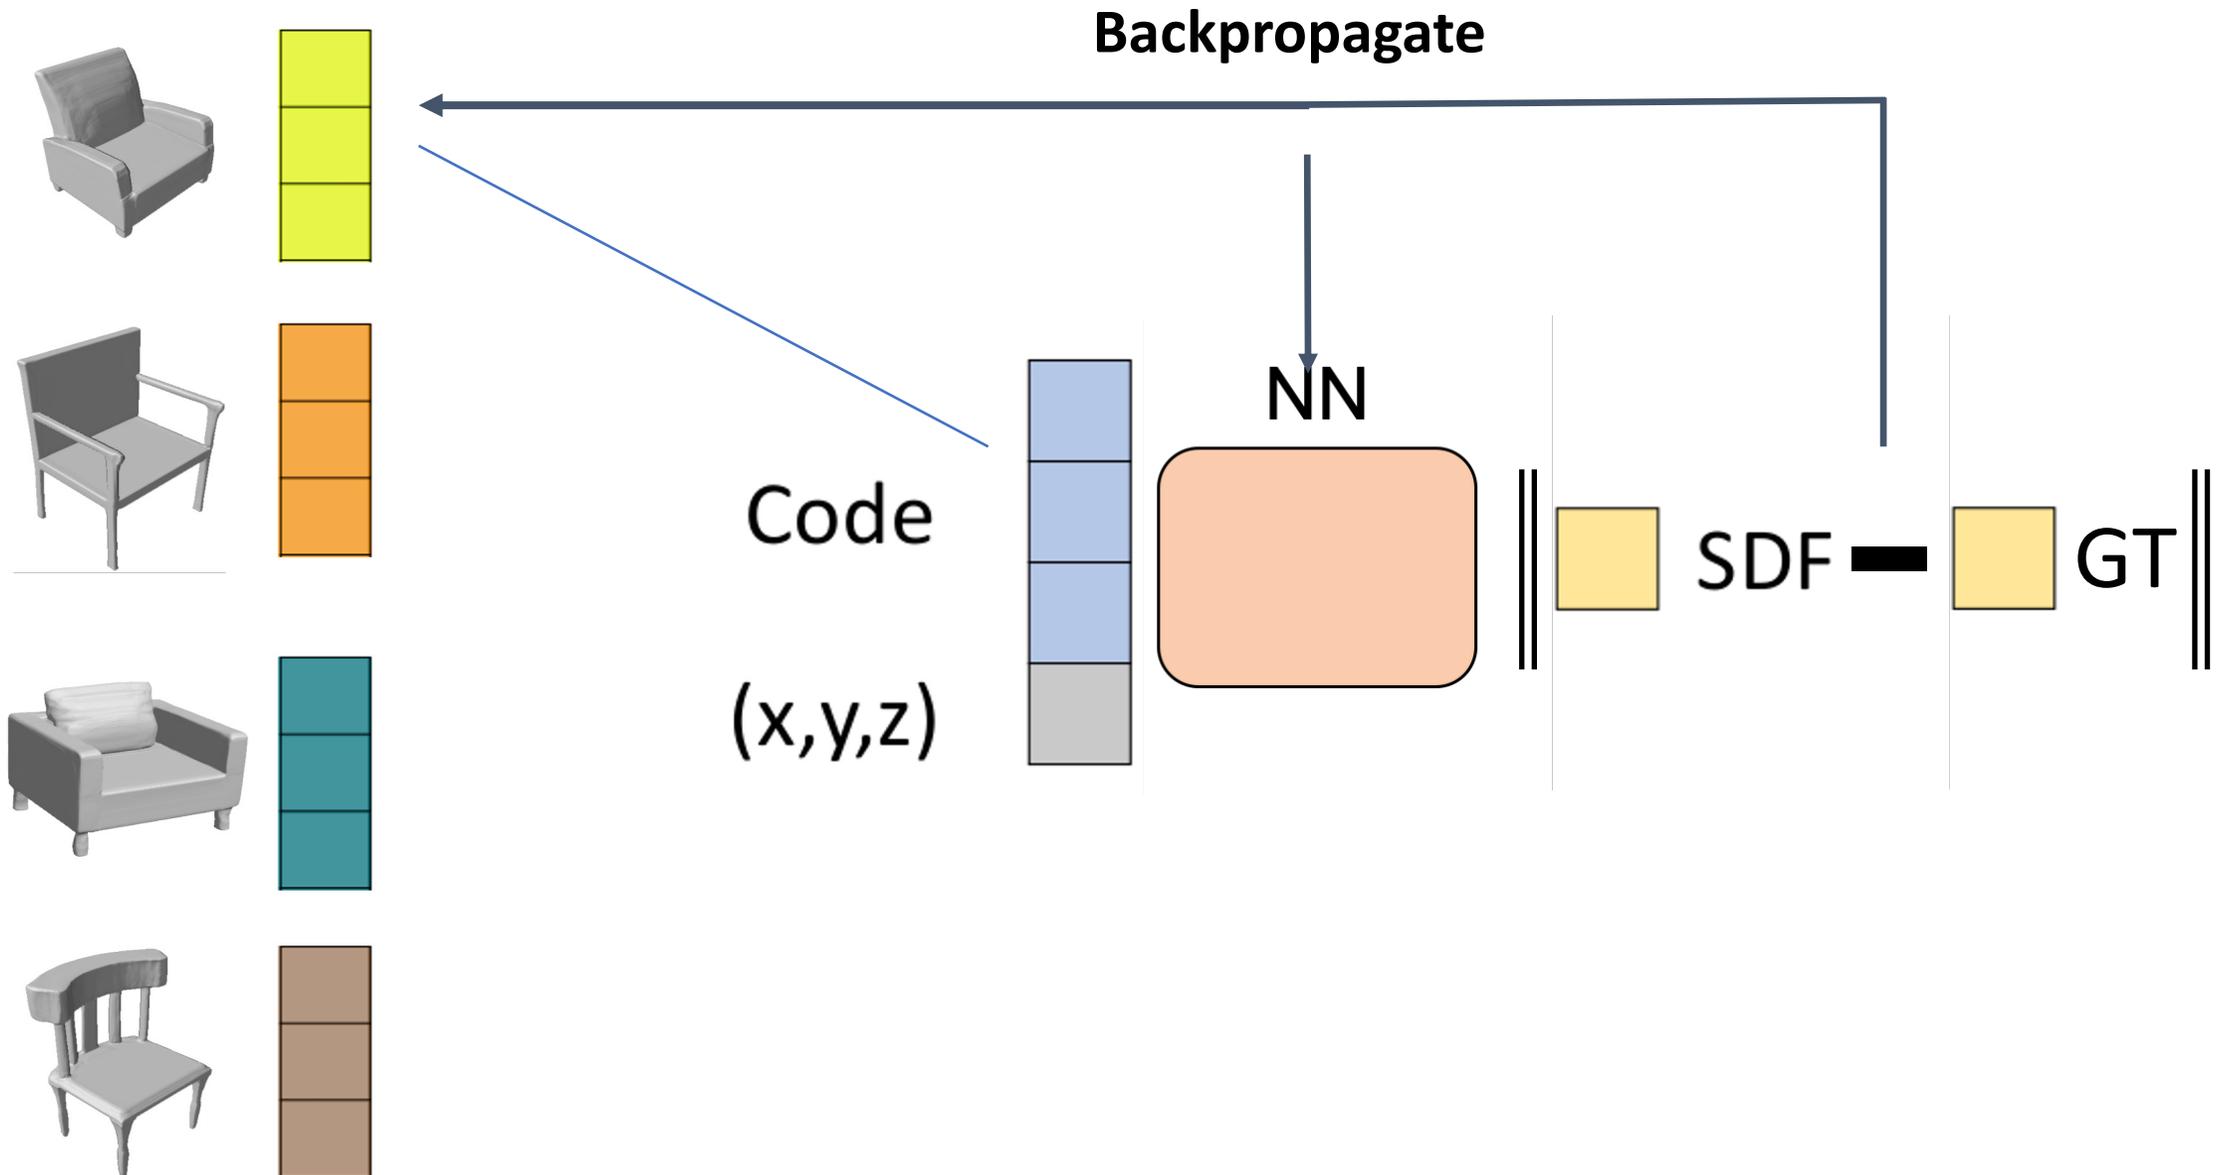
# Benefits of Auto-Decoder



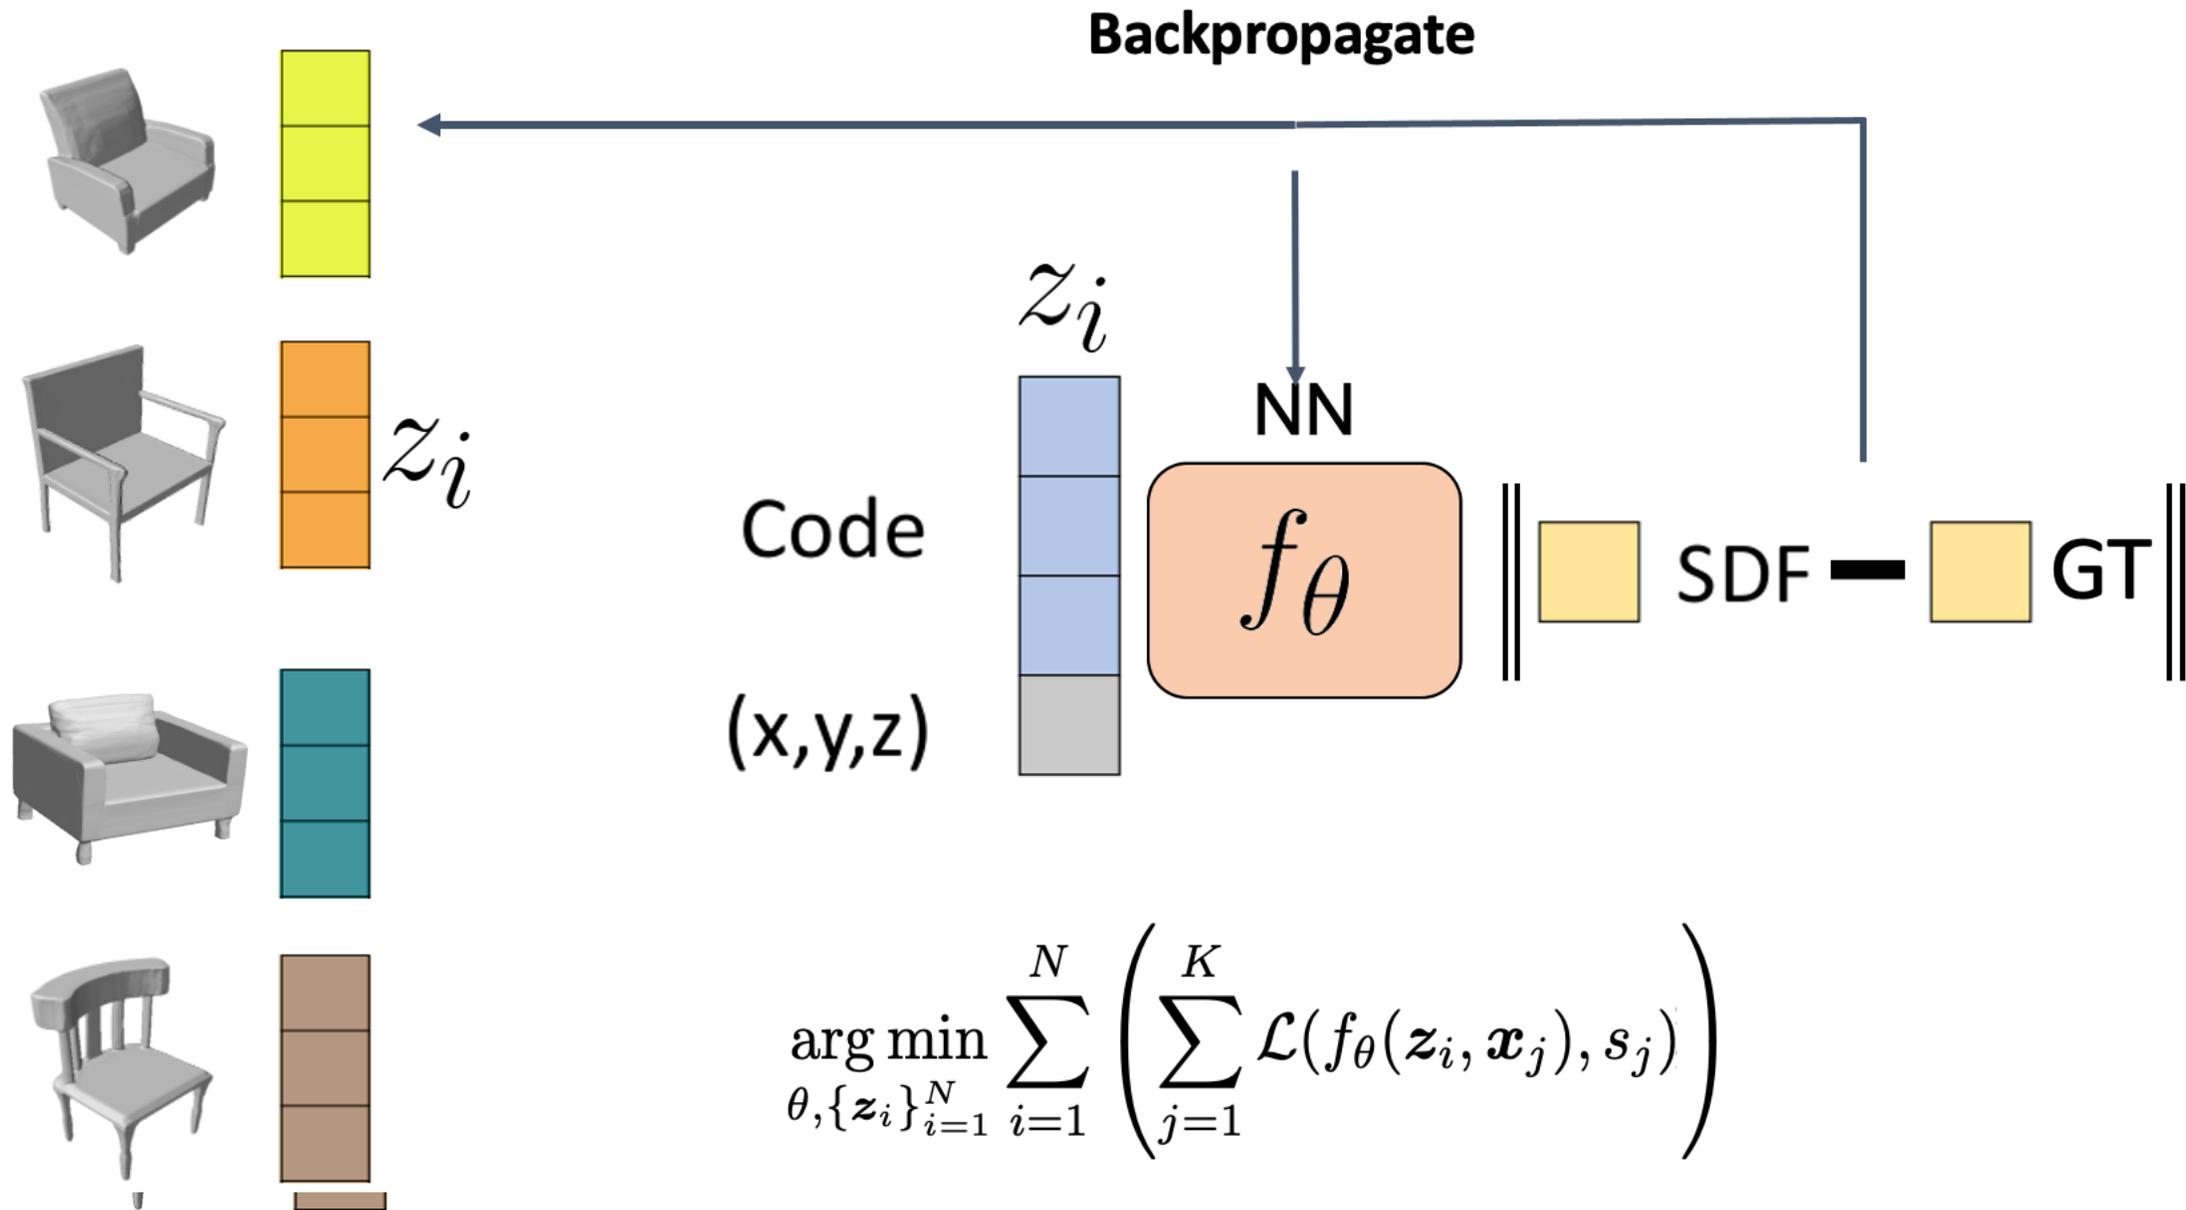
## Benefits during Inference

1. Any Number of Observations – Partial
2. More Controlled Inference – e.g. Accuracy, Priors

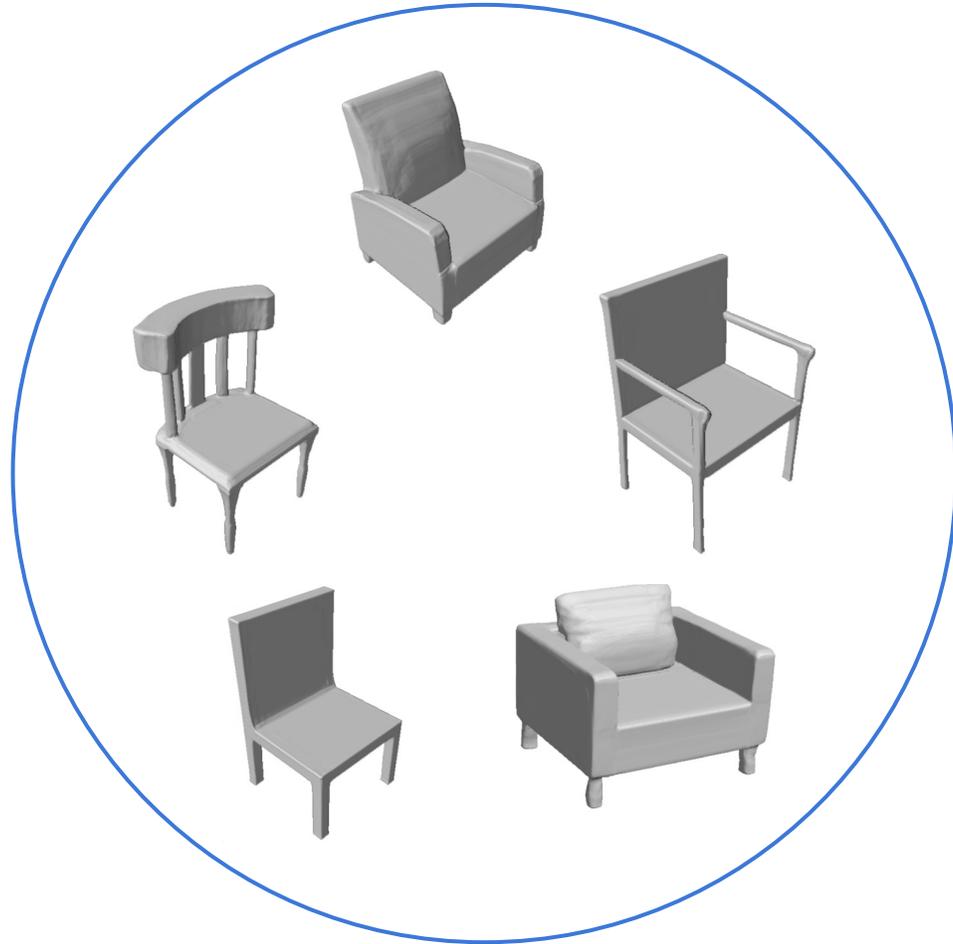
# Auto-Decoder Training



# Auto-Decoder Training

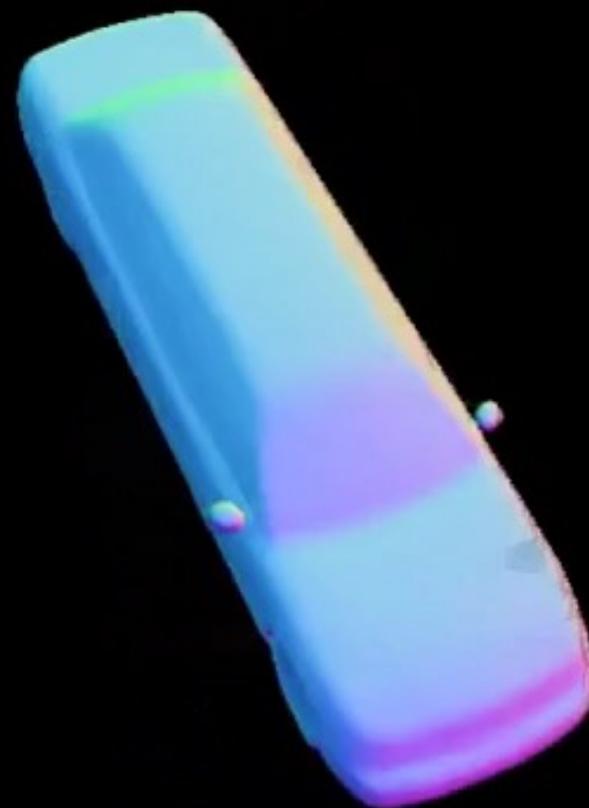


# Latent Space of Shapes



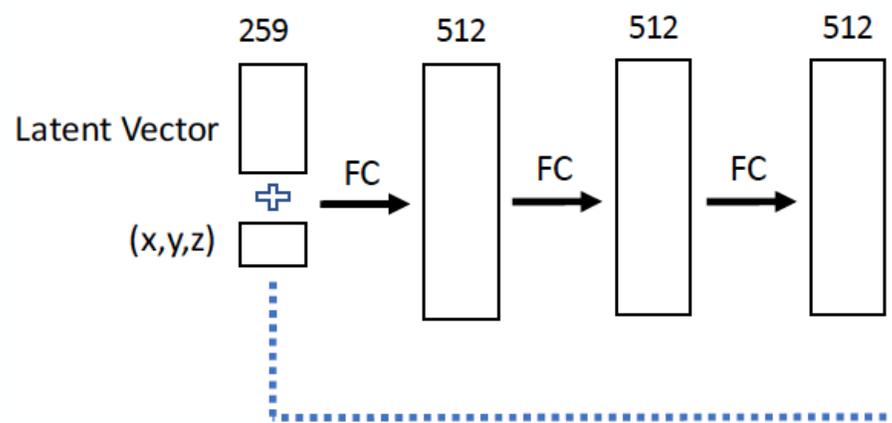


Learned Chair Shape Space



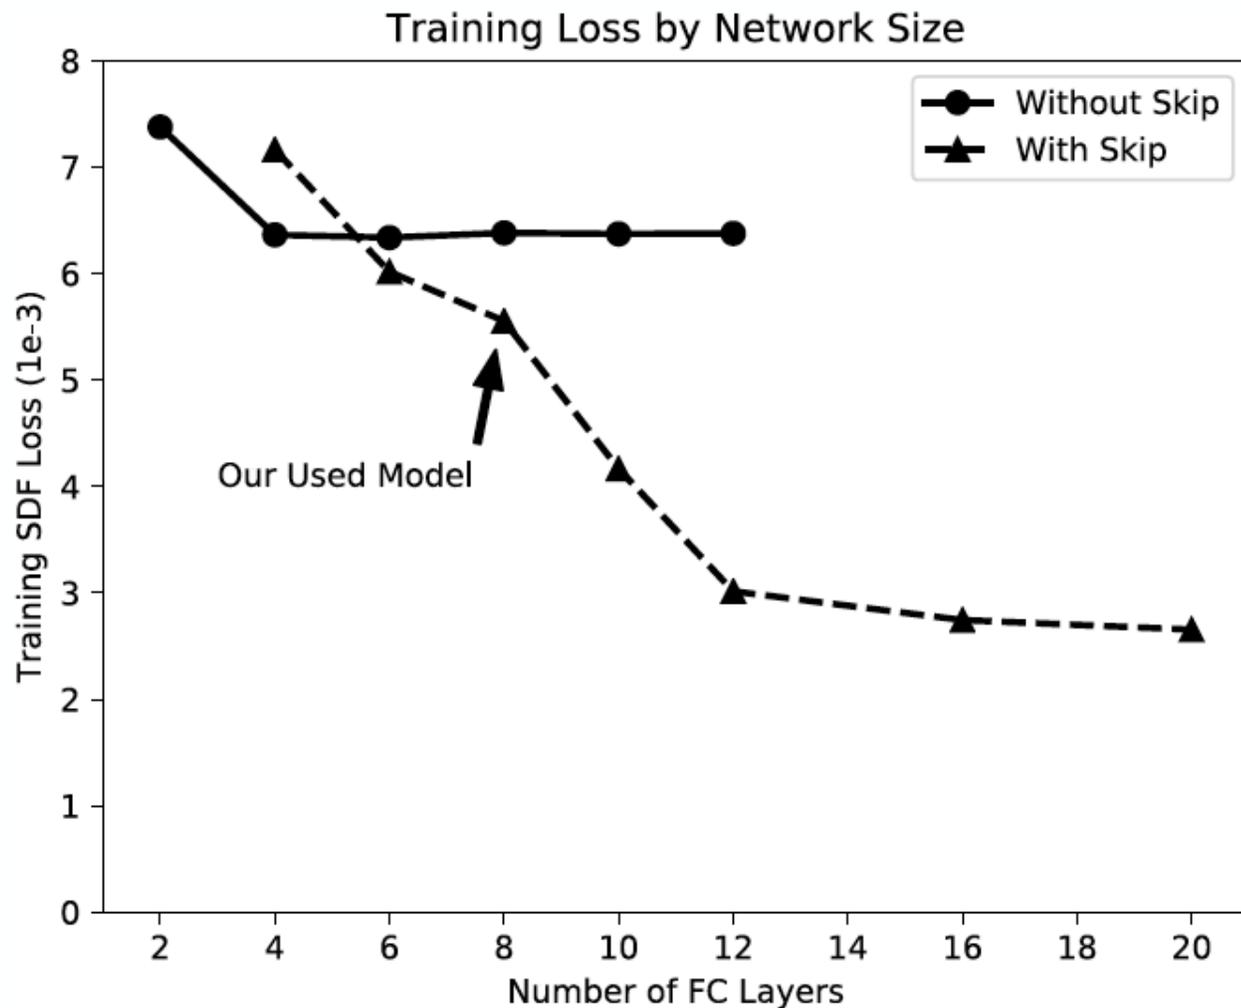
Learned Car Shape Space

# Conditional Neural SDFs



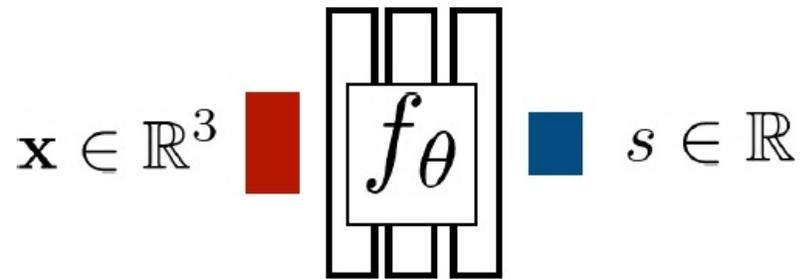
Simplest approach — just concat

Skip connecti

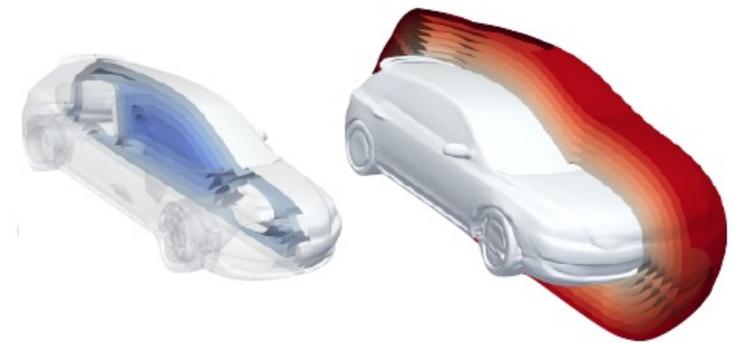


# Questions we want to answer

- How do we create a 3D mesh from a SDF function? (SDF rendering)
- How do we generalize this to any objects? (Training Deep SDF)
- How do we use this during inference? (Testing DeepSDF)
- How do DeepSDF concept extends to NeRF and other methods?



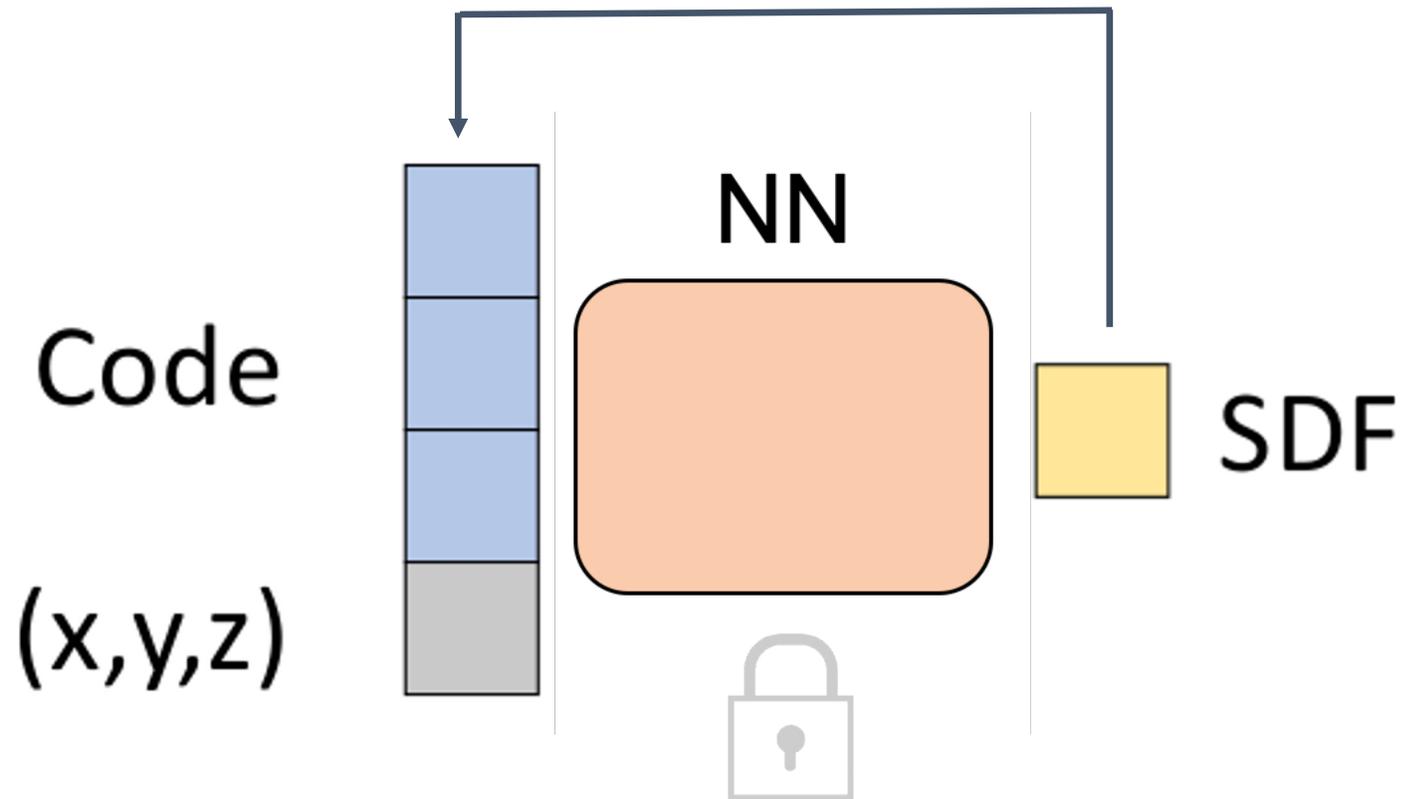
$$f_\theta : \mathbb{R}^3 \rightarrow \mathbb{R}$$



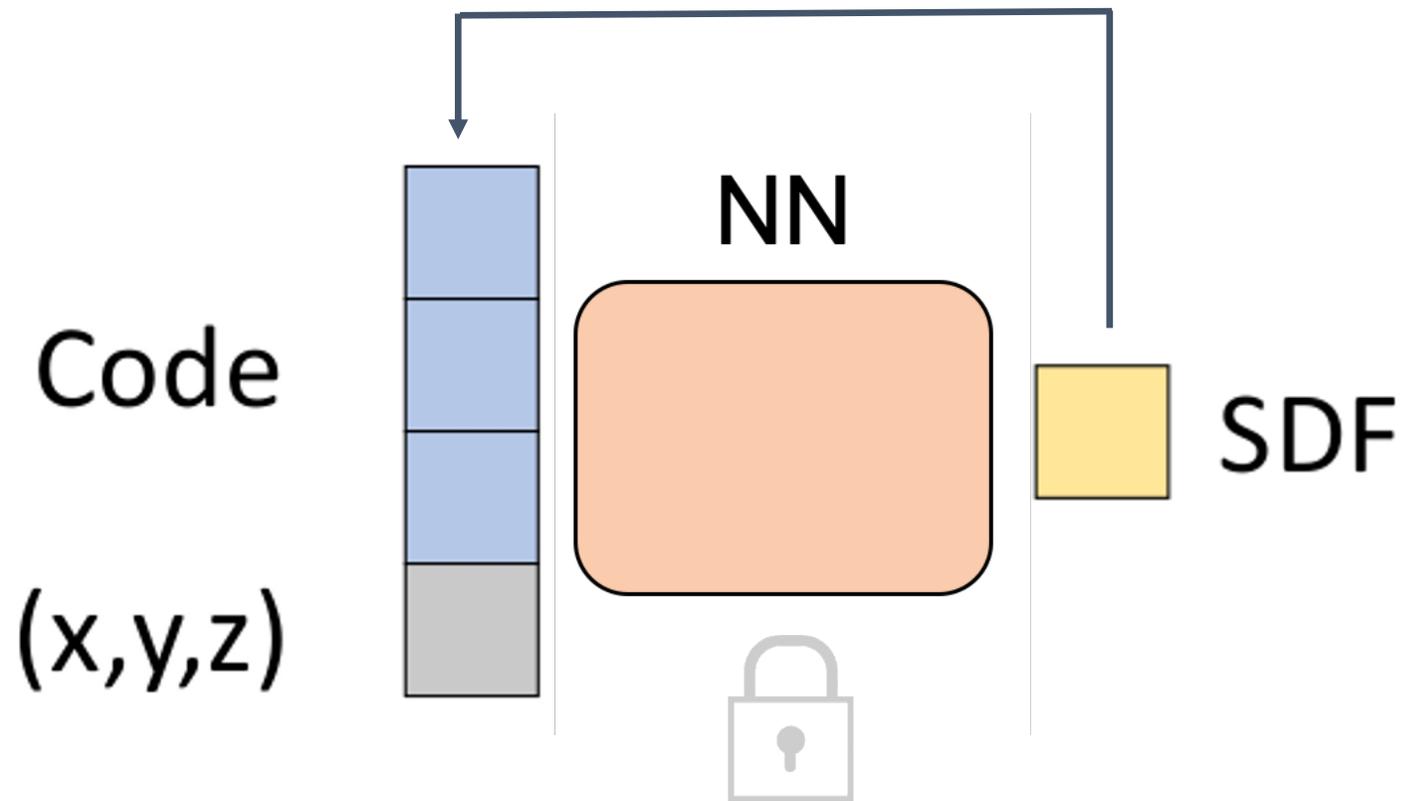
**Signed Distance Field (for a single instance):**

(position)  $\rightarrow$  (distance)

# Auto-Decoder Inference



# Auto-Decoder Inference



Optimize the Code until the SDF matches the Test Shape



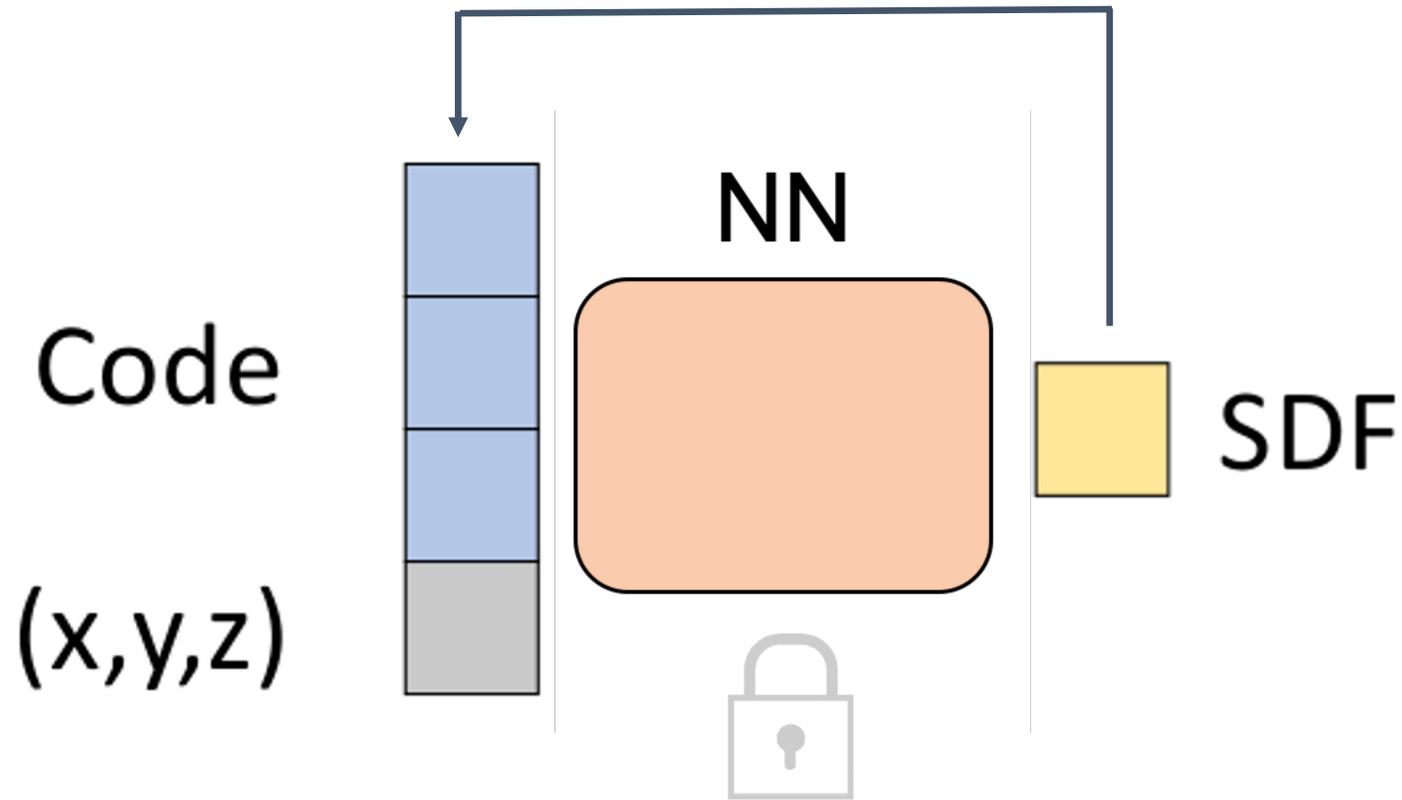
**Test Shape**



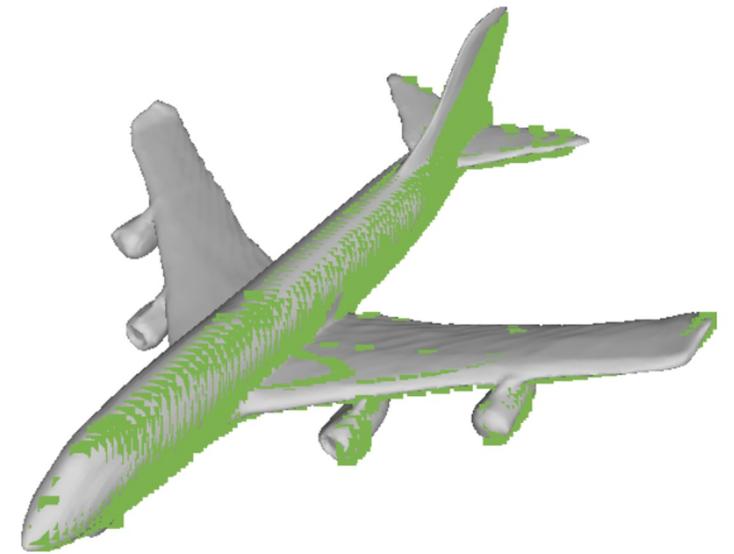
**Reconstruction**

# Auto-Decoder Inference

Input = Partial Observation



Input

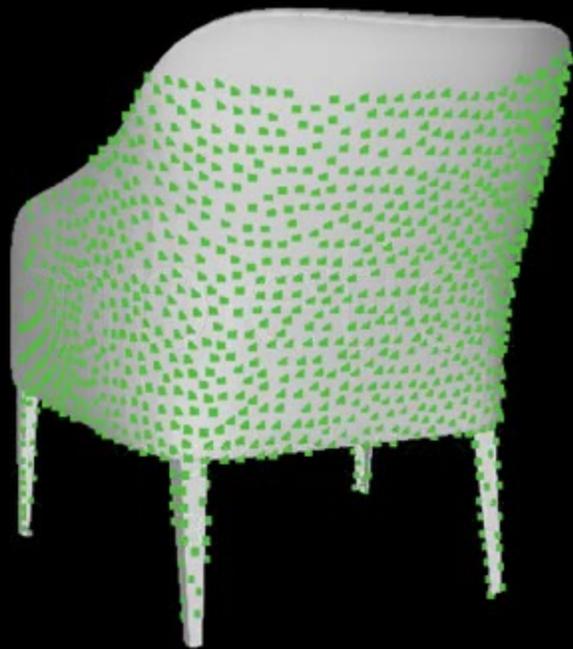


Reconstruction

# DeepSDF: Learning Continuous Signed Distance Functions for Shape Representation



GT Shape



Sampled Depth Map  
(Green Input)

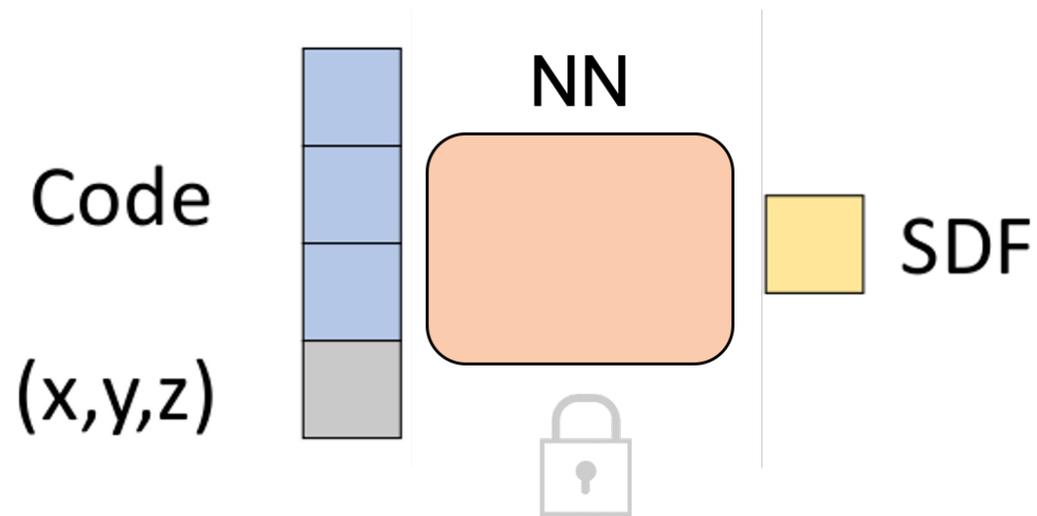


DeepSDF Decoding  
Output

Example evolution of shape completion for a partial depth map (ADAM iteration: 1-400)

# Adding Priors to Inference

$$\hat{\mathbf{z}} = \arg \min_{\mathbf{z}} \sum_{(\mathbf{x}_j, \mathbf{s}_j) \in X} \mathcal{L}(f_{\theta}(\mathbf{z}, \mathbf{x}_j), \mathbf{s}_j)$$



# Adding Priors to Inference

$$\hat{\mathbf{z}} = \arg \min_{\mathbf{z}} \sum_{(\mathbf{x}_j, \mathbf{s}_j) \in X} \mathcal{L}(f_{\theta}(\mathbf{z}, \mathbf{x}_j), \mathbf{s}_j)$$

Distribution Prior:  $\frac{1}{\sigma^2} \|\mathbf{z}\|_2^2$

SDF Regularization:  $(\|\nabla_{\mathbf{x}} f(\mathbf{x}; \theta)\| - 1)^2$  (Matan et al. 2020)

Normal Regularization:  $\|\nabla_{\mathbf{x}} f(\mathbf{x}_i; \theta) - \mathbf{n}_i\|$

# Results

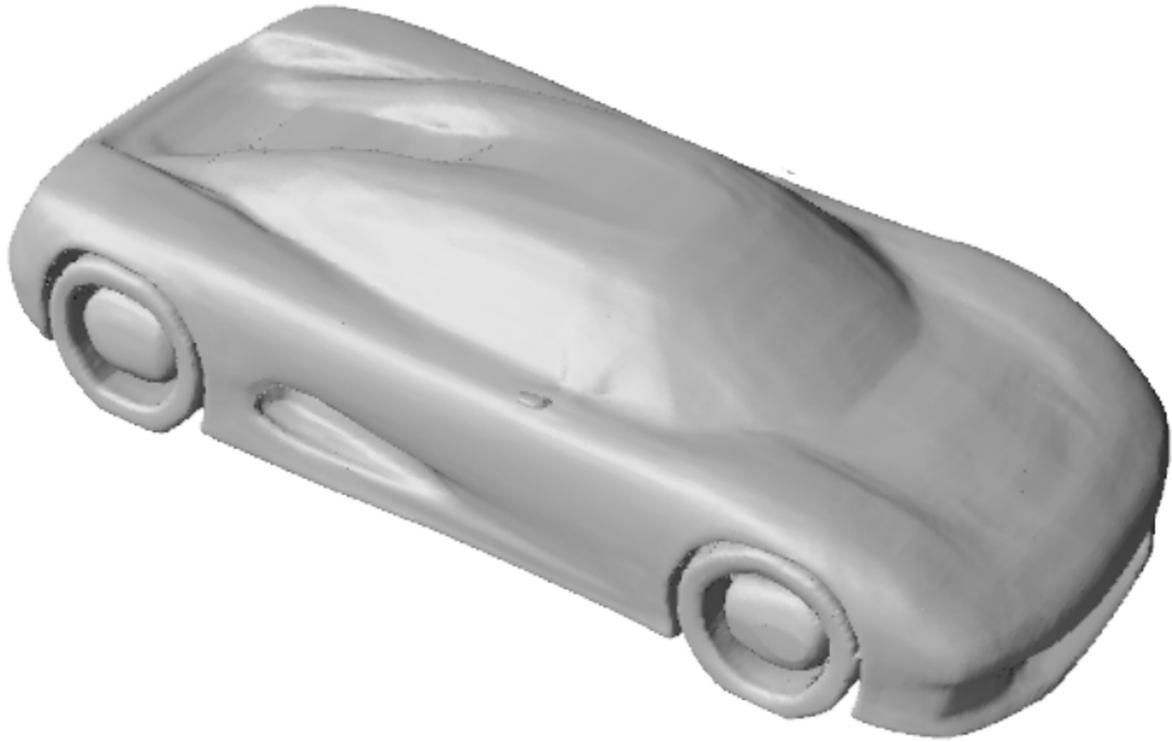
## Auto-encoding unknown shapes

CD, median					
AtlasNet-Sph.	0.511	0.079	0.389	2.180	0.330
AtlasNet-25	0.276	0.065	0.195	0.993	0.311
DeepSDF	<b>0.072</b>	<b>0.036</b>	<b>0.068</b>	<b>0.219</b>	<b>0.088</b>

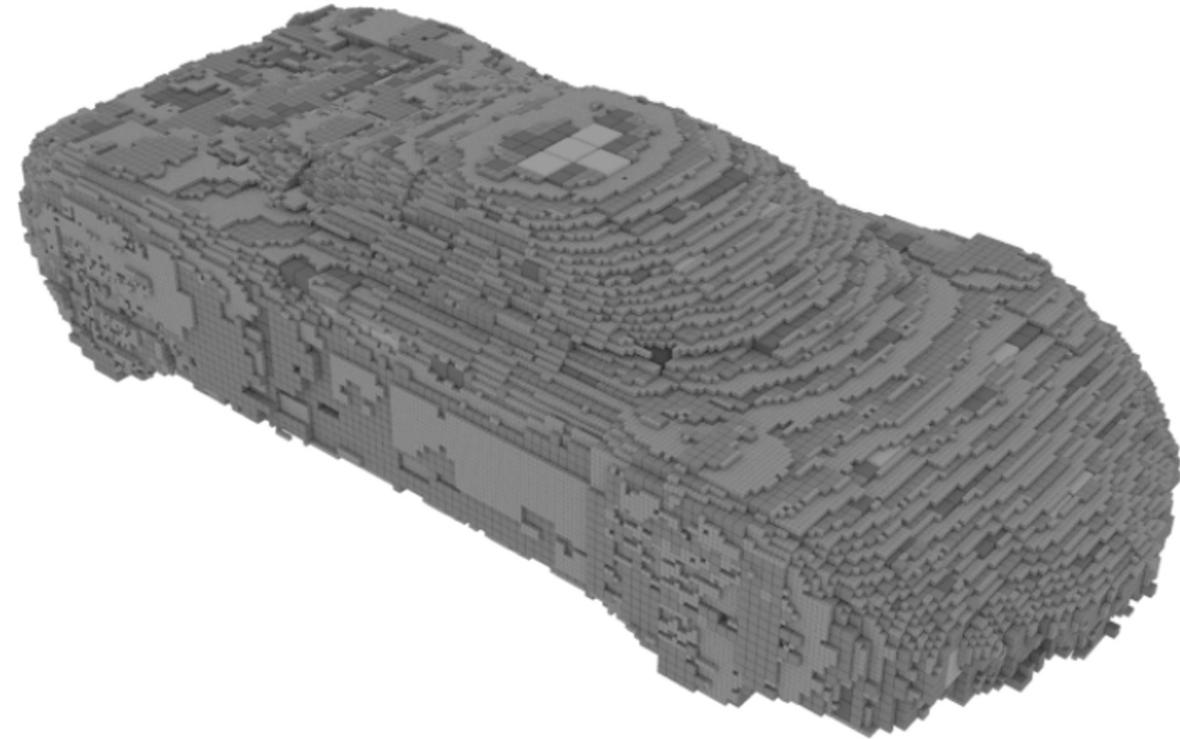
## Shape completion

Method \Metric	CD, med.	<i>lower is better</i>			<i>higher is better</i>	
		CD, mean	EMD	Mesh acc.	Mesh comp.	Cos sim.
chair						
3D-EPN	2.25	2.83	0.084	0.059	0.209	0.752
DeepSDF	<b>1.28</b>	<b>2.11</b>	<b>0.071</b>	<b>0.049</b>	<b>0.500</b>	<b>0.766</b>
plane						
3D-EPN	1.63	2.19	0.063	0.040	0.165	0.710
DeepSDF	<b>0.37</b>	<b>1.16</b>	<b>0.049</b>	<b>0.032</b>	<b>0.722</b>	<b>0.823</b>

# Results: Comparison with Octree-Based



Our  
Reconstruction



Octree Based

# Results: Comparisons with Mesh-Based



Ground Truth



Our Reconstruction

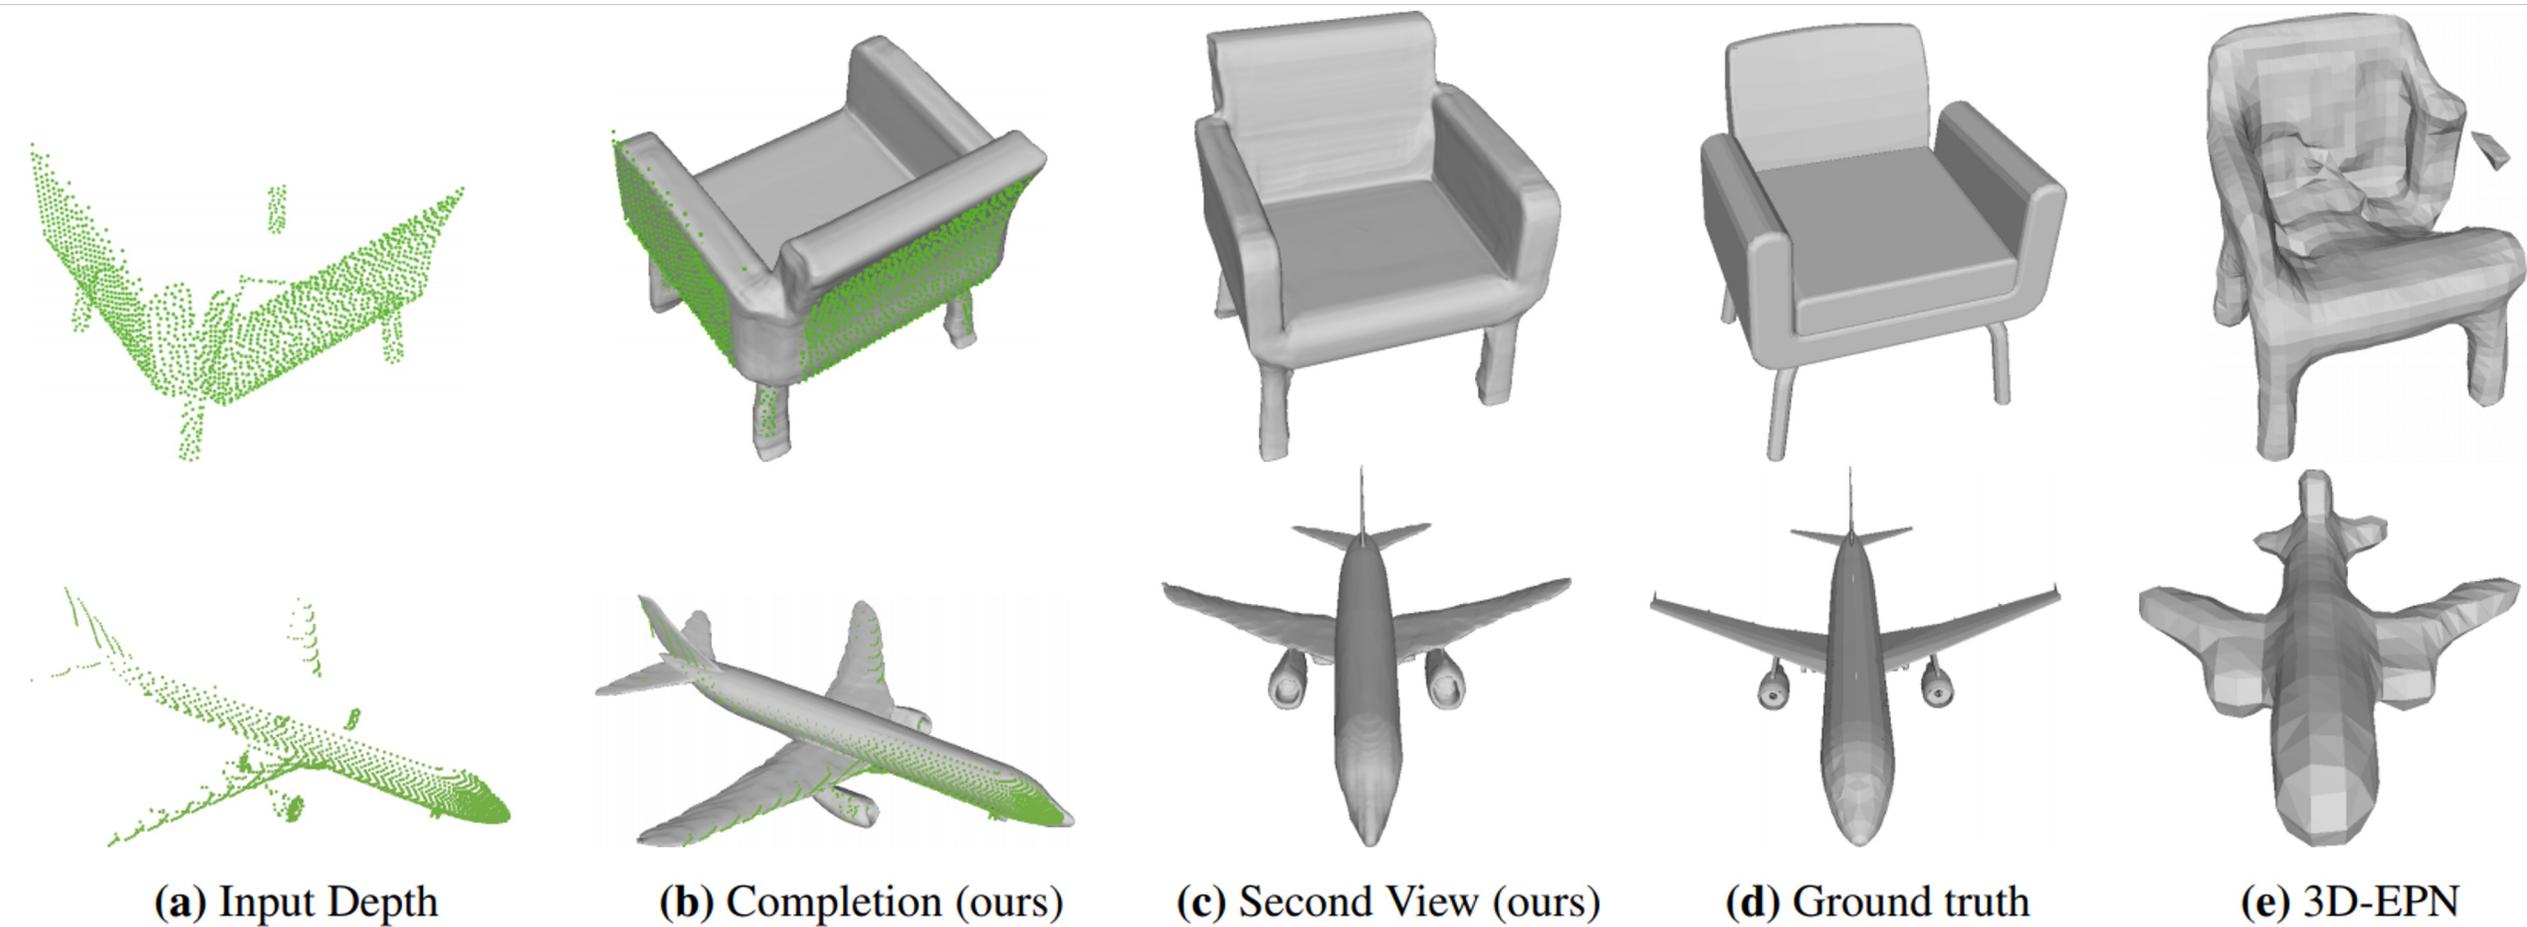


Atlasnet (25 Patches)



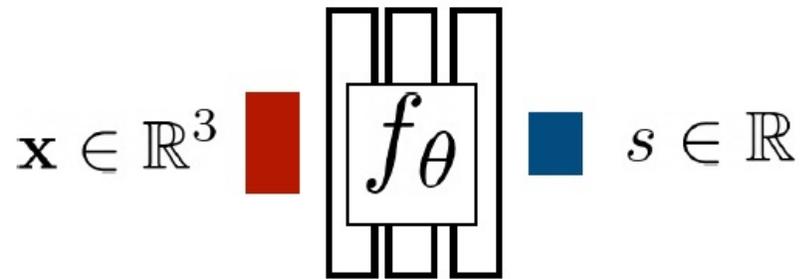
Atlasnet (1 Patch)

# Shape Completion

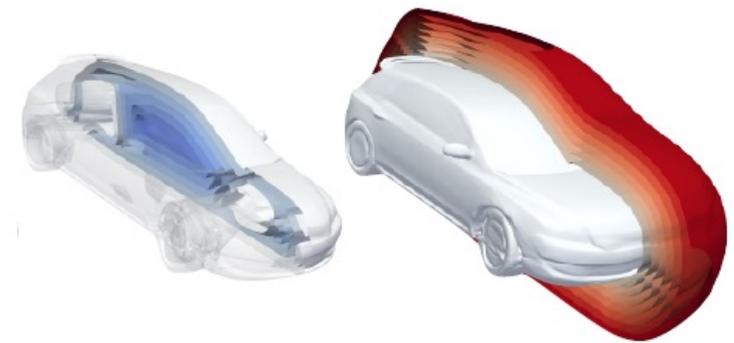


# Questions we want to answer

- How do we create a 3D mesh from a SDF function? (SDF rendering)
- How do we generalize this to any objects? (Training Deep SDF)
- How do we use this during inference? (Testing DeepSDF)
- How do DeepSDF concept extends to NeRF and other methods?

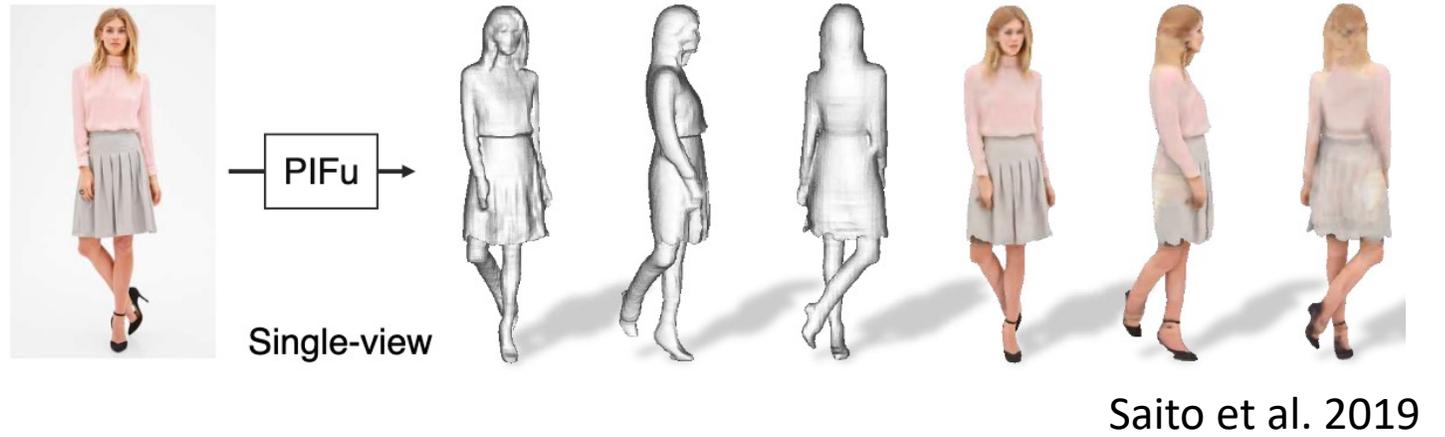
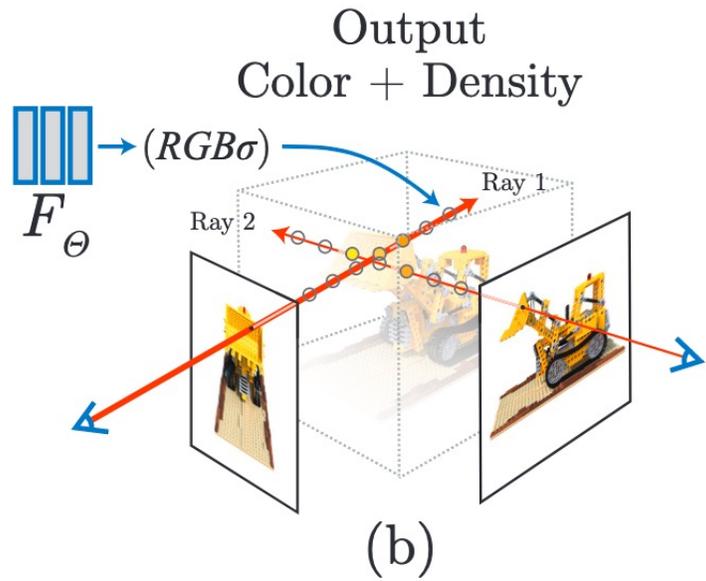
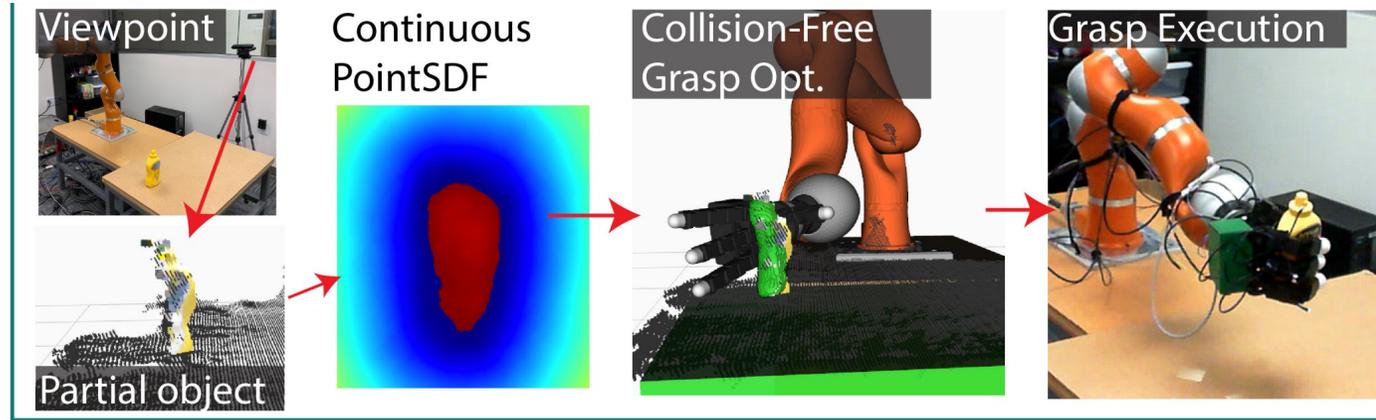
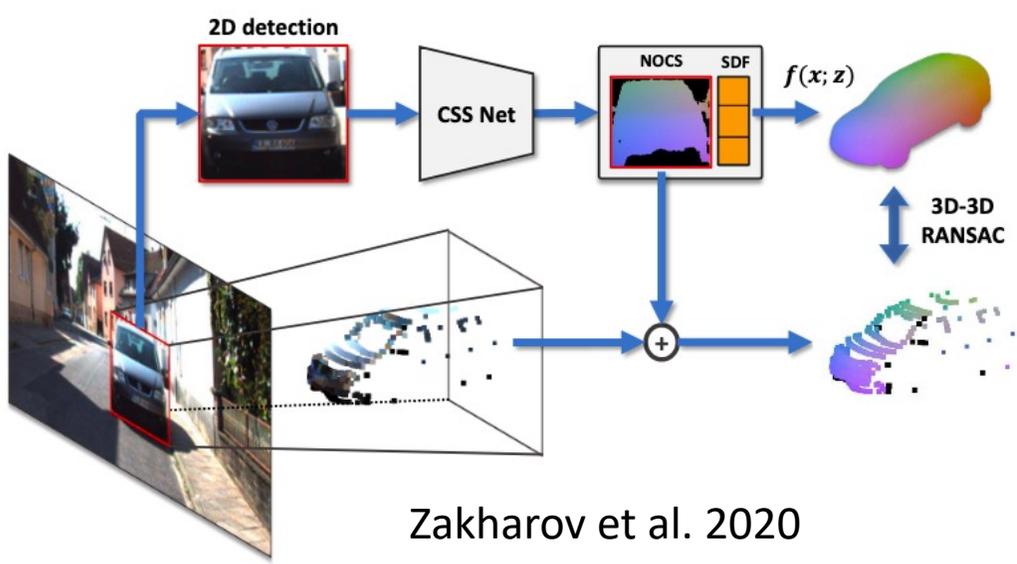


$$f_\theta : \mathbb{R}^3 \rightarrow \mathbb{R}$$

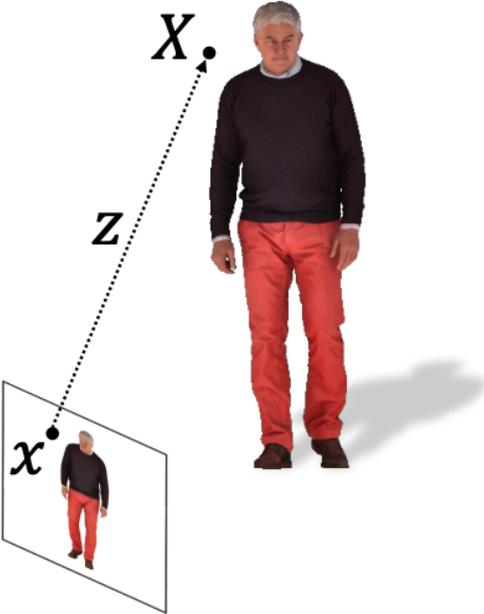
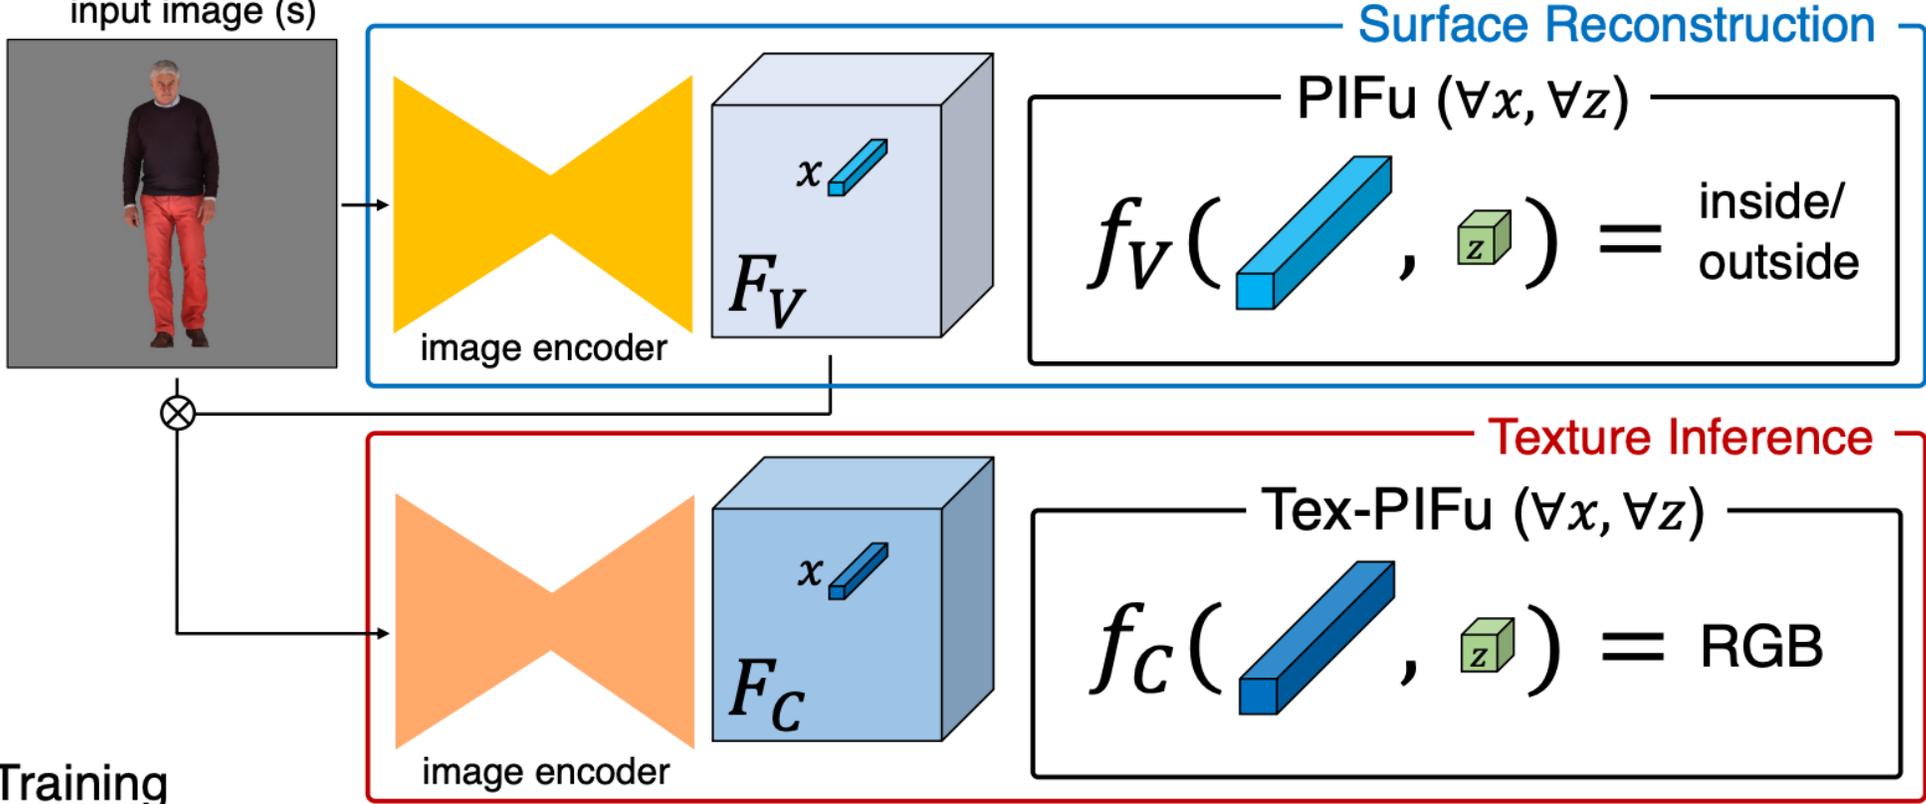


**Signed Distance Field (for a single instance):**

(position)  $\rightarrow$  (distance)

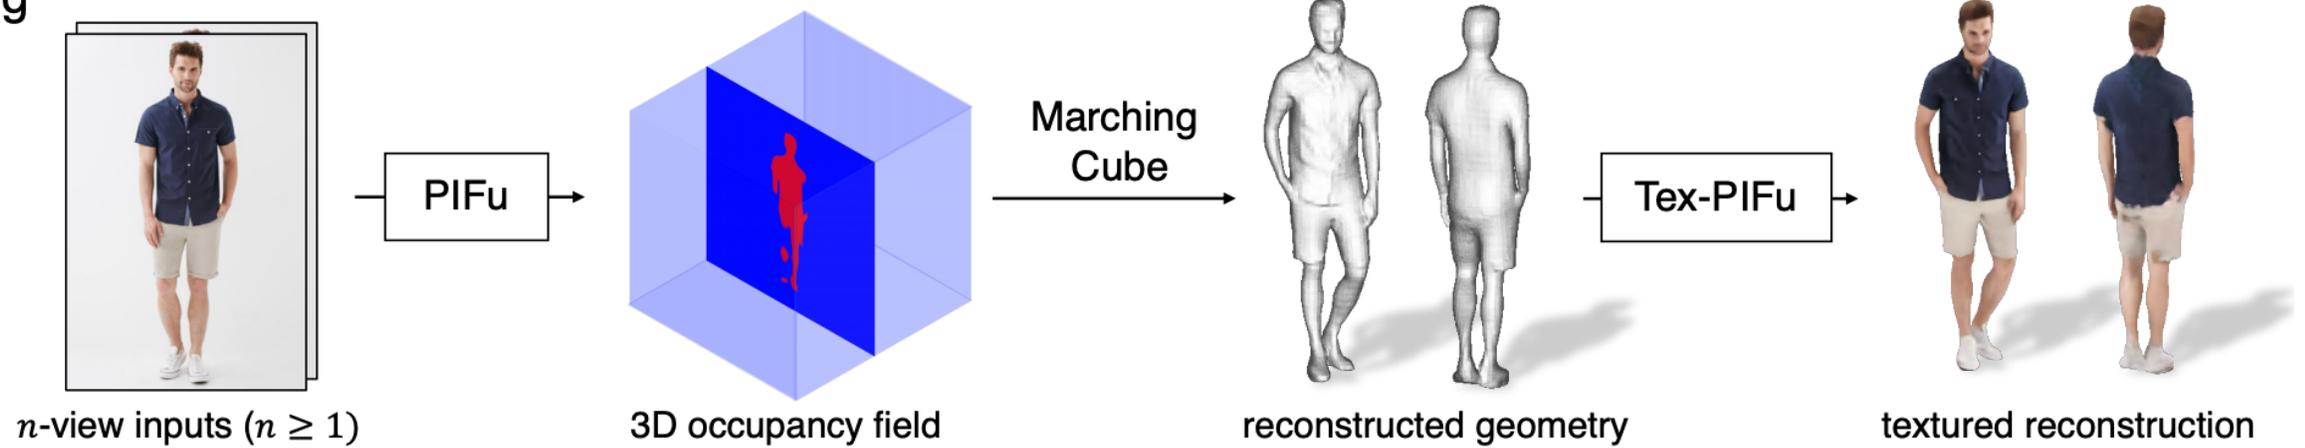


# DeepSDF Extensions: PiFU

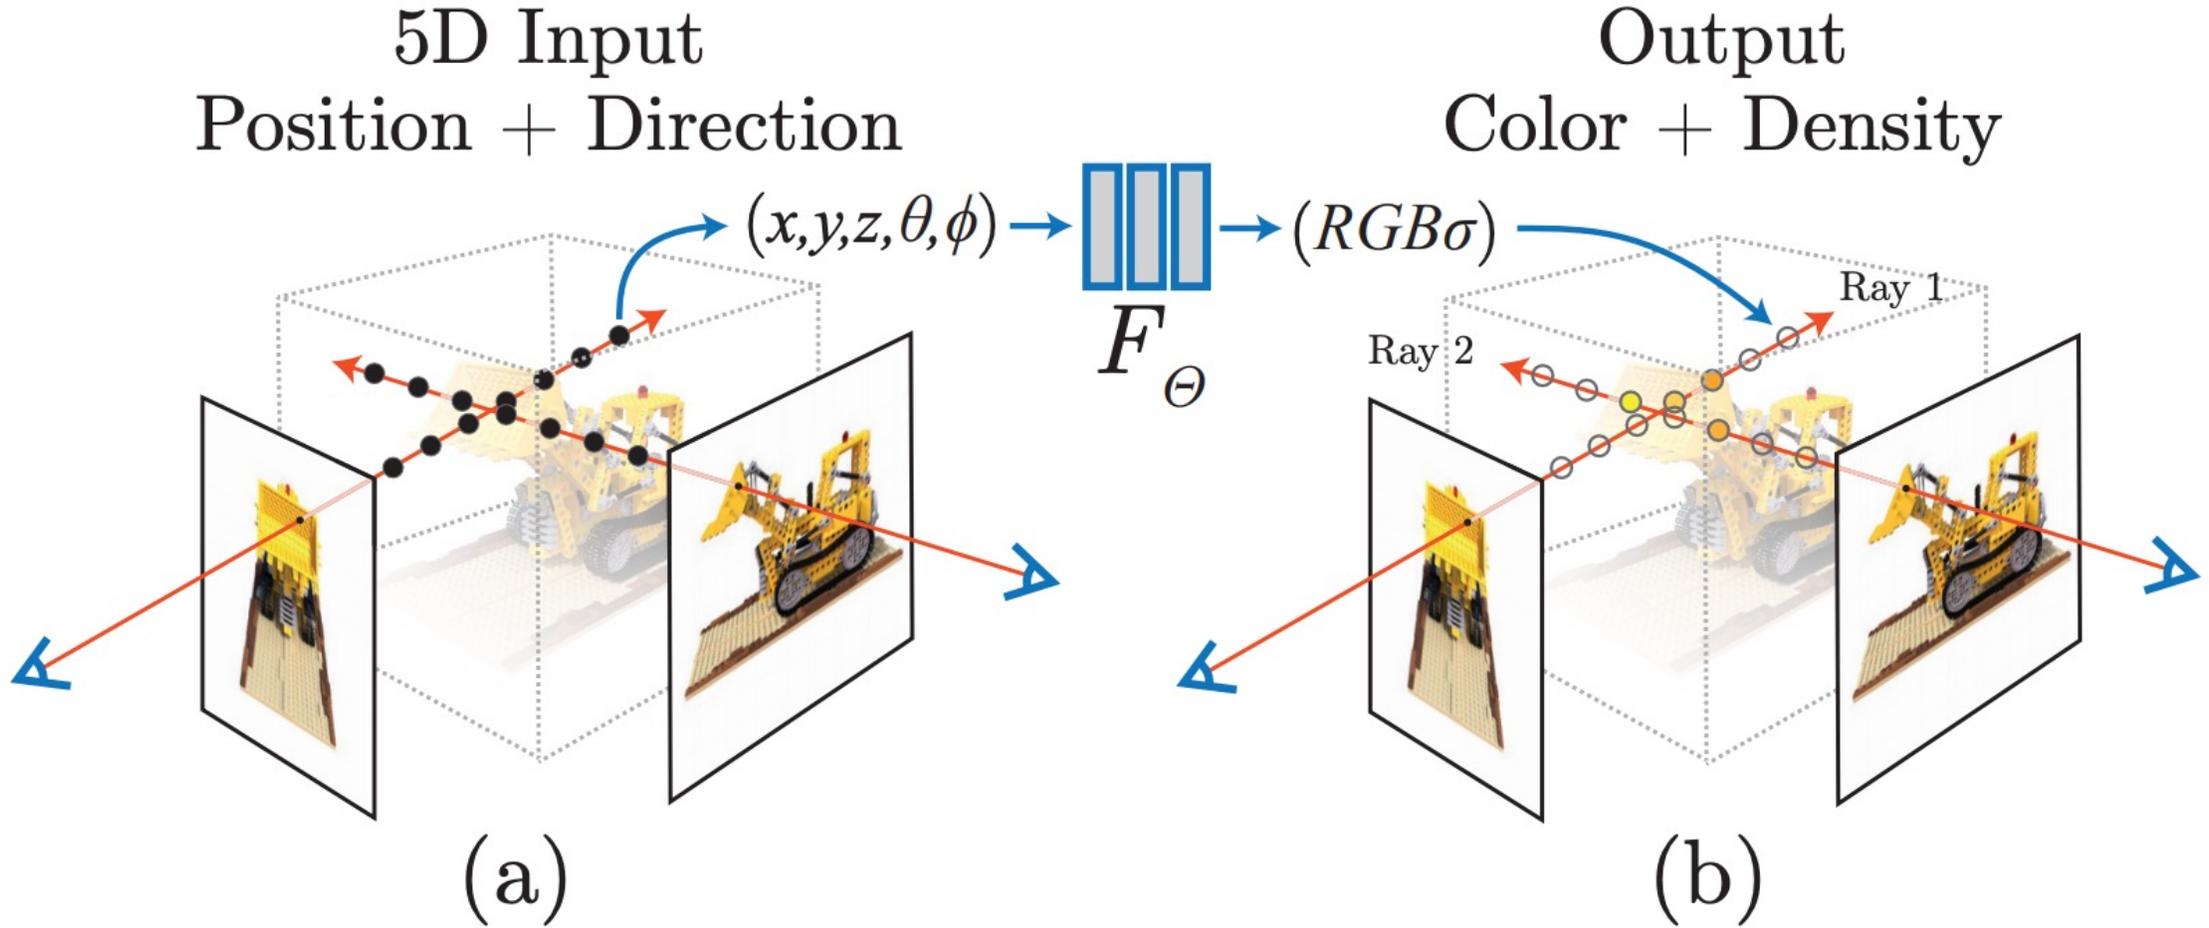


# DeepSDF Extensions: PiFu

Testing



# DeepSDF Extensions: NeRF

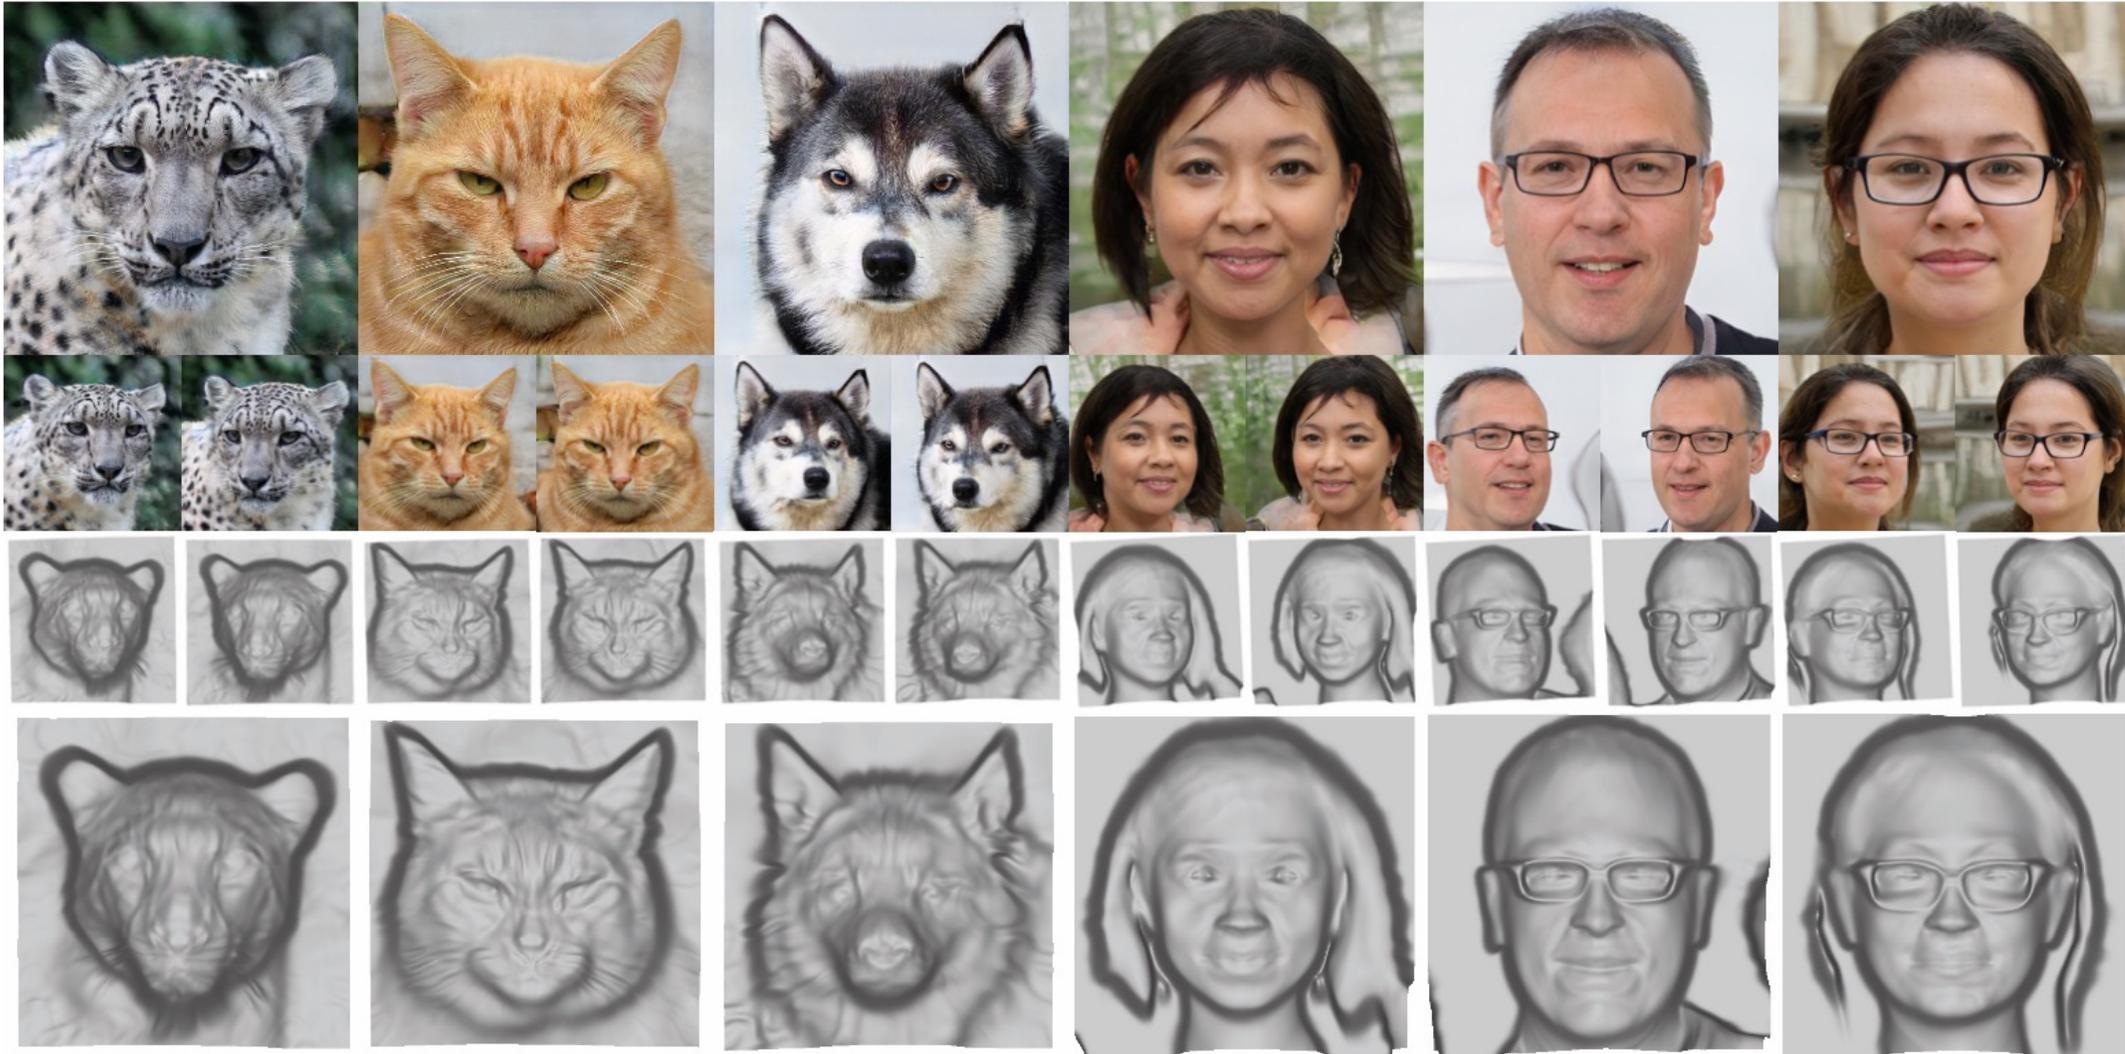


# DeepSDF Extensions: NeRF



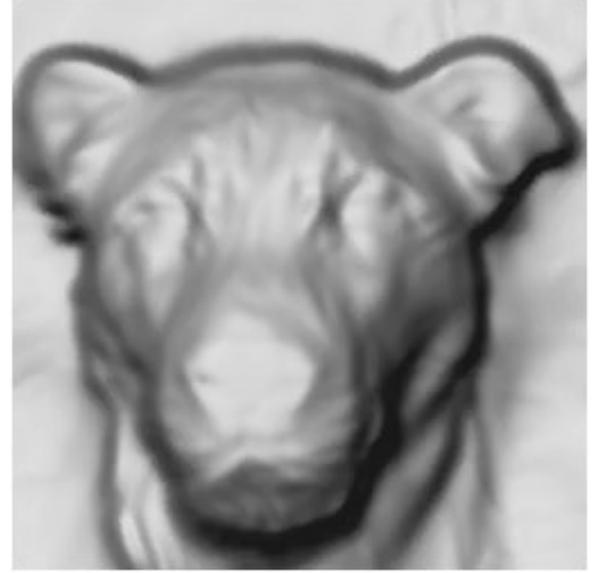
# DeepSDF Extension: StyleSDF

- A 3D GAN using DeepSDF + NeRF modeling

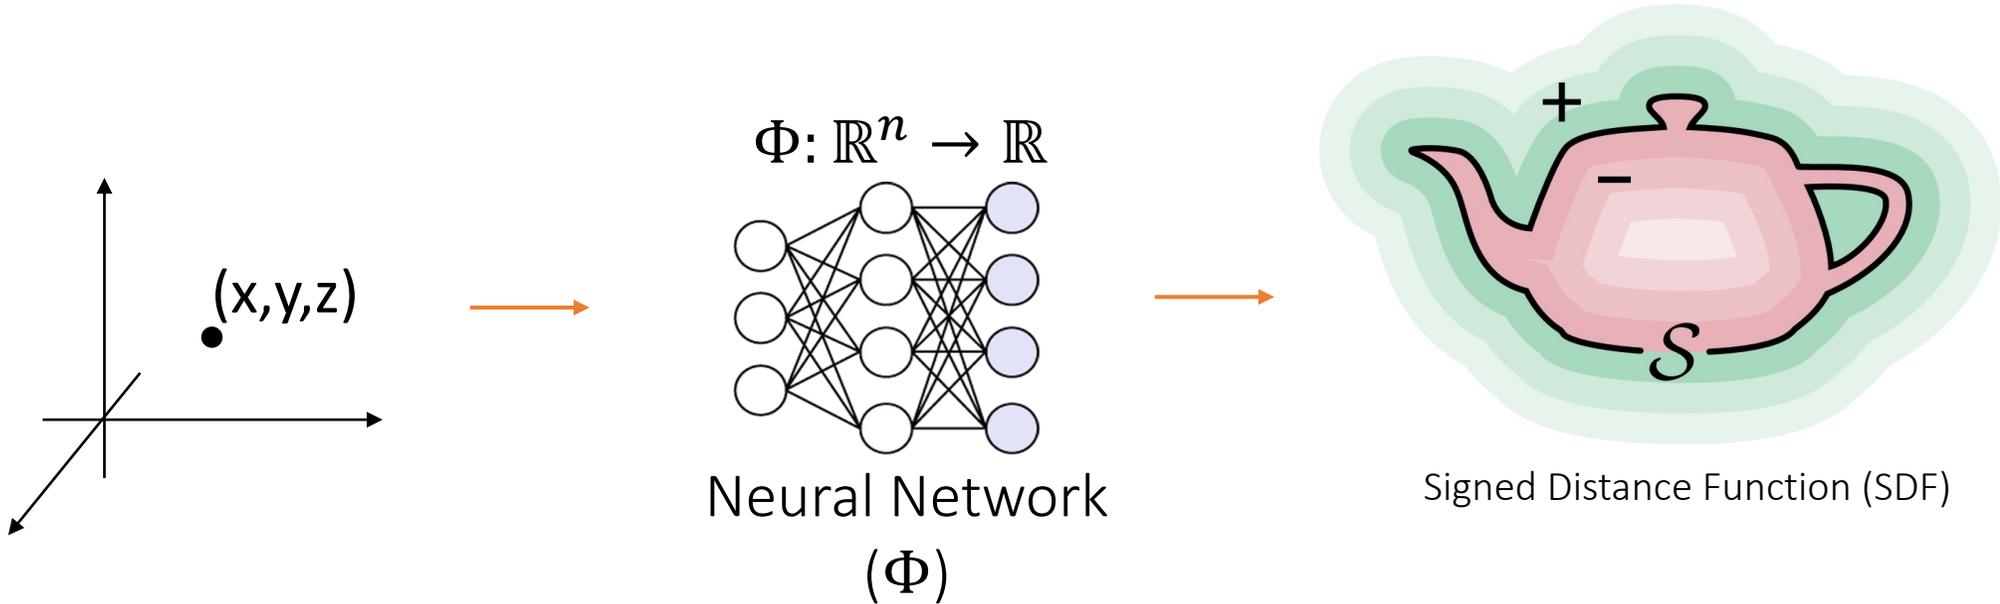


Or-El et al. 2021

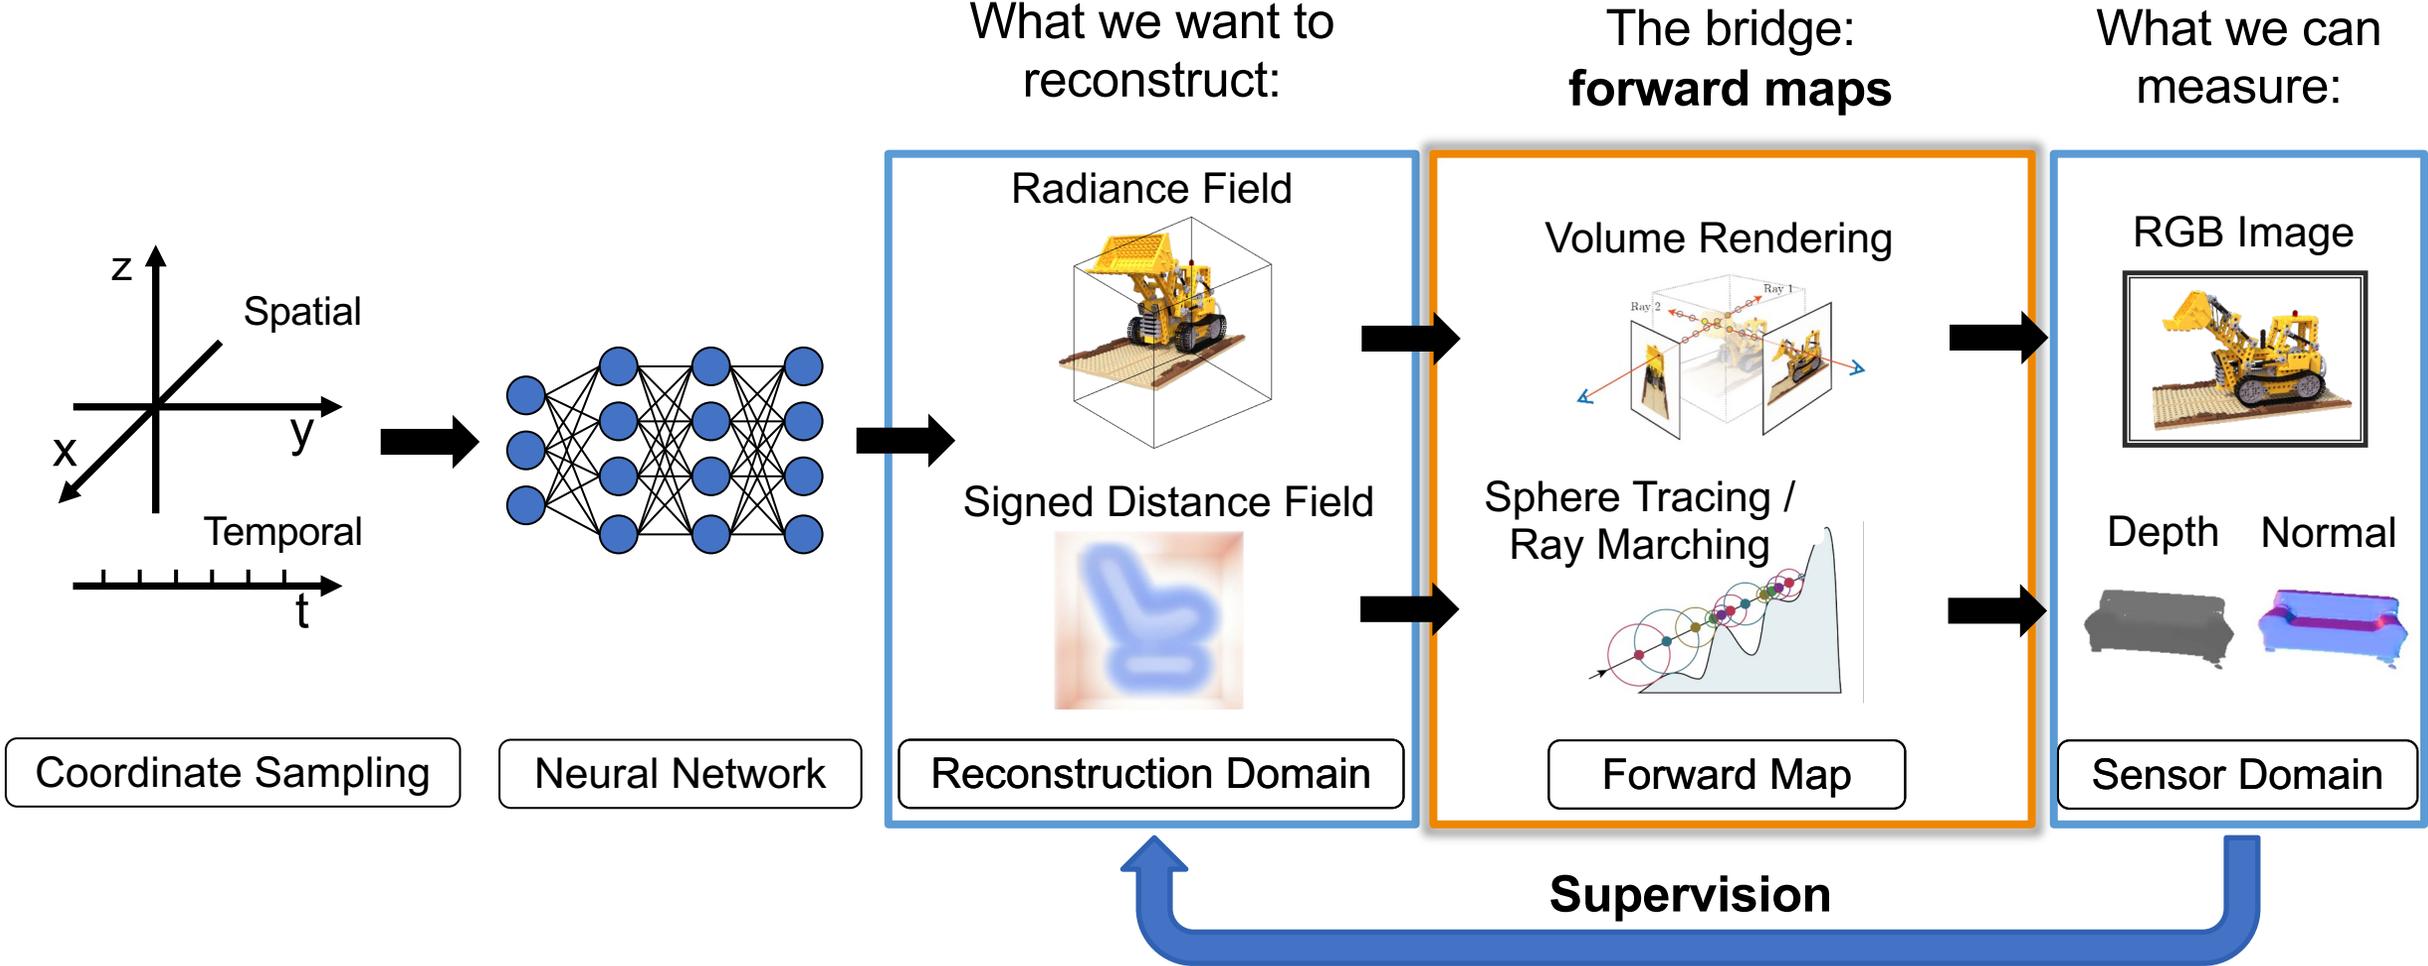
# DeepSDF Extension: StyleSDF



# What are neural fields?



# Neural Field General Framework

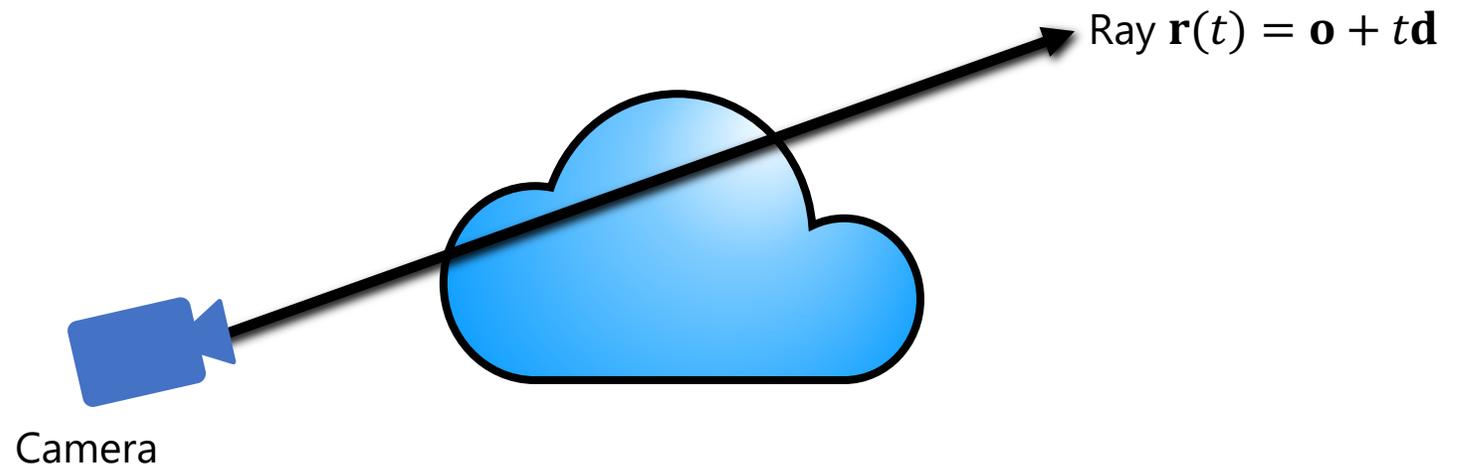


# Outline

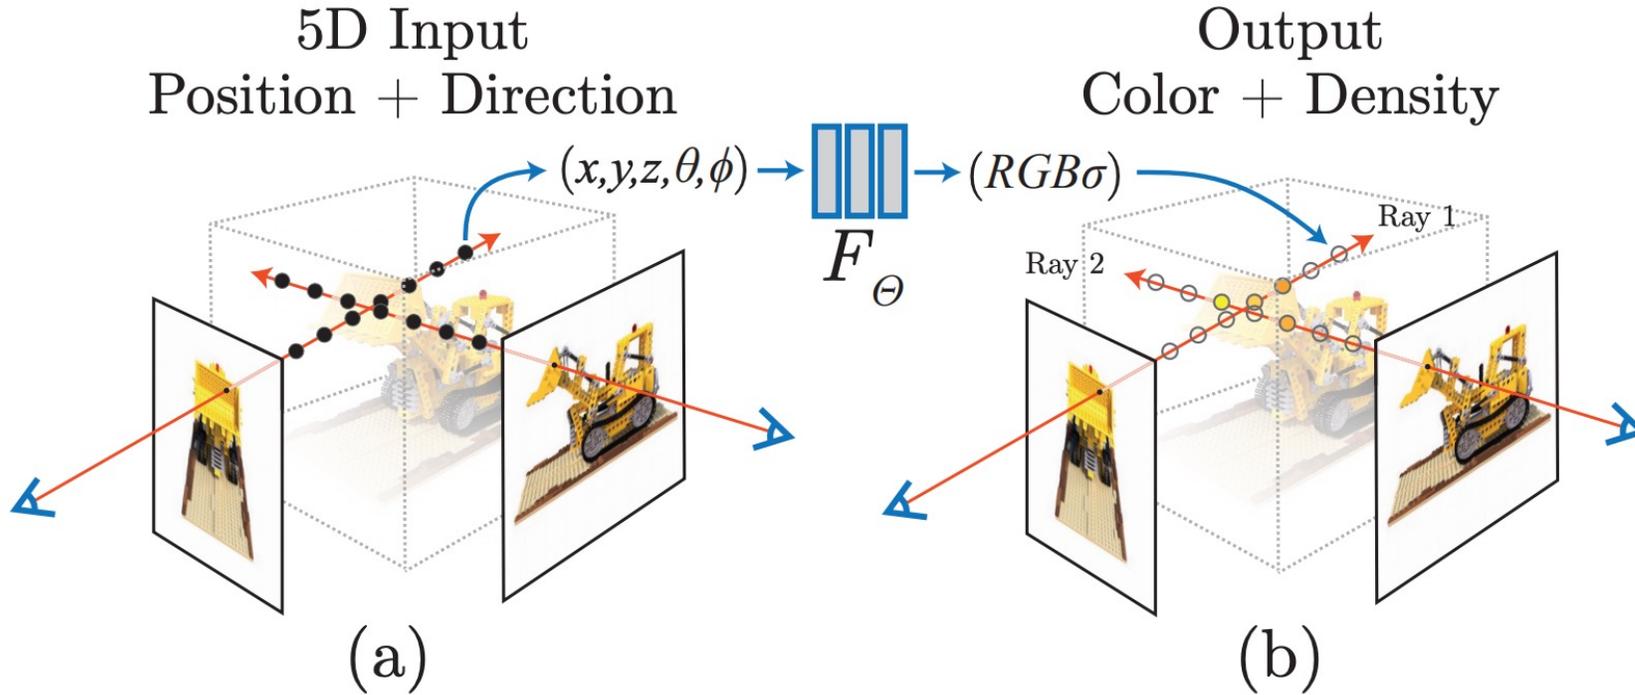
- What is Neural Fields & why it got so much attention?
- The Prelude: Neural Implicit Surfaces
- **Introduction to Volume Rendering**

# What is Volume Rendering?

- Assume a cloud of tiny colored particles in 3D. Each particle has a RGB color and a density.
- Take a pixel on image plane, and shoot a ray from the camera center, through the pixel and into the 'cloud of tiny colored particles'
- What should be the color for that pixel?



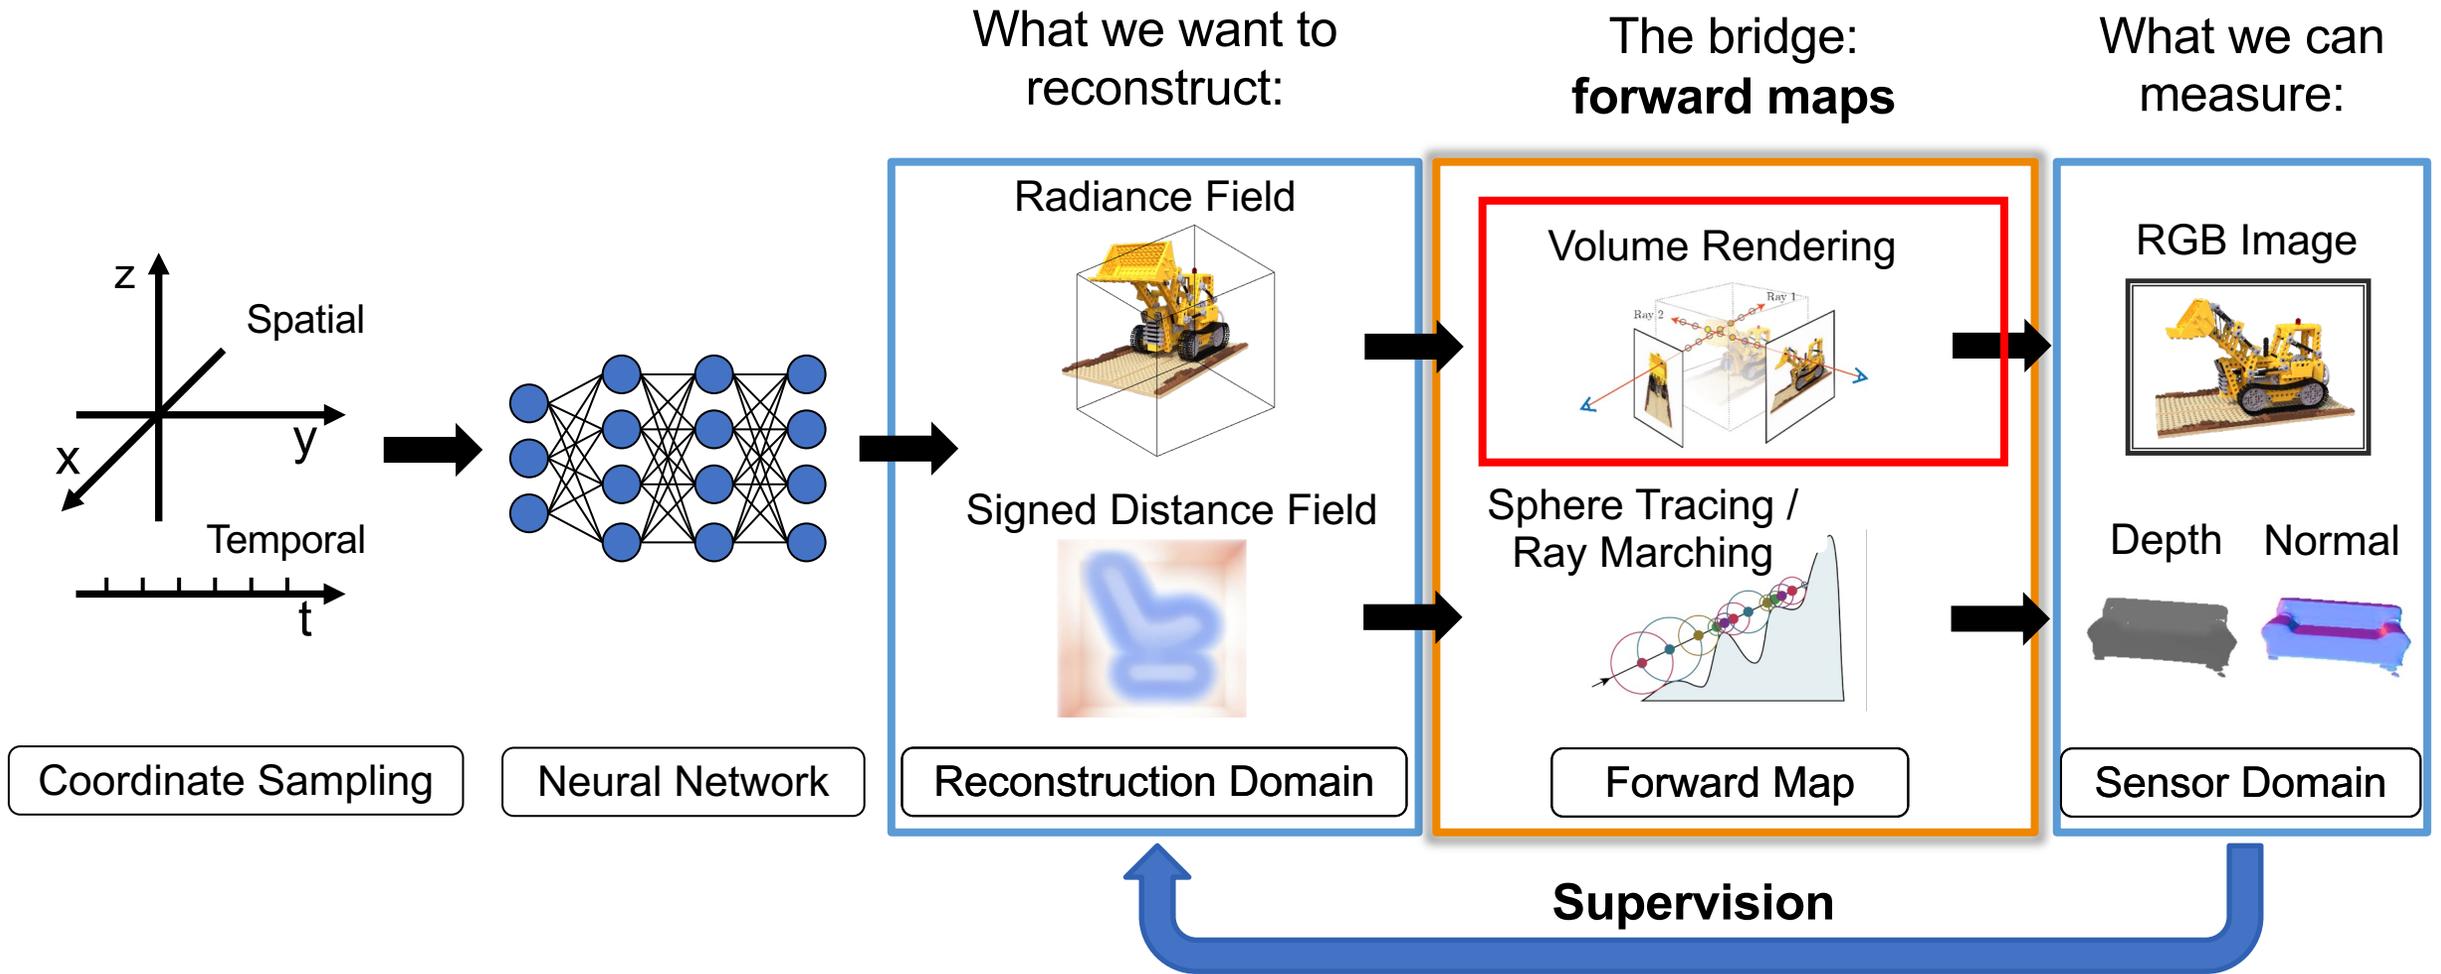
# Volume Rendering in NeRF



In NeRF, we model a 3D scene with a 'cloud of tiny colored particles'.

# Neural Field General Framework

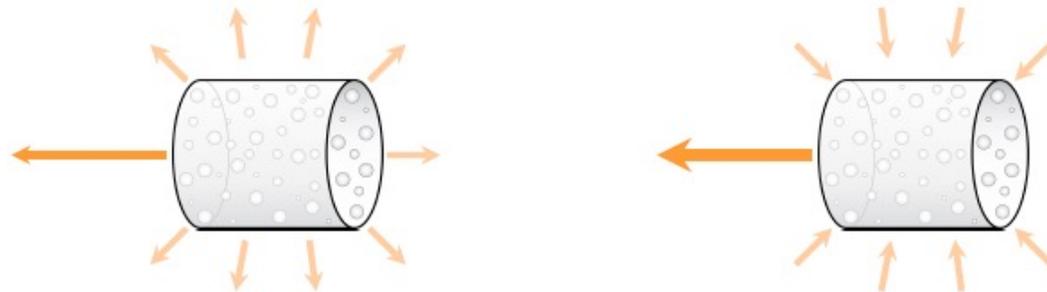
Volume Rendering tells us how to take 'cloud of tiny colored particles' and create an image.



# Radiative Transfer Equation



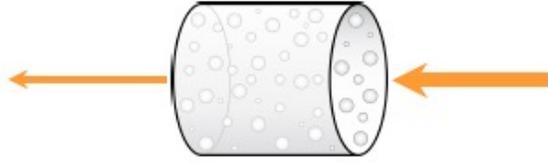
$$dL(\mathbf{x}, \vec{\omega}) = \underbrace{-\sigma_a(\mathbf{x})L(\mathbf{x}, \vec{\omega})dz - \sigma_s(\mathbf{x})L(\mathbf{x}, \vec{\omega})dz}_{\text{Losses}} + \underbrace{+\sigma_a(\mathbf{x})L_e(\mathbf{x}, \vec{\omega})dz + \sigma_s(\mathbf{x})L_s(\mathbf{x}, \vec{\omega})dz}_{\text{Gains}}$$



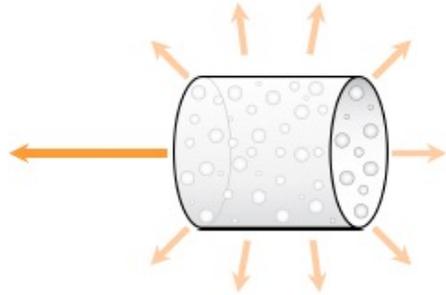
Emission

In-scattering

# Emission-Absorption Model (Ignoring Scattering)



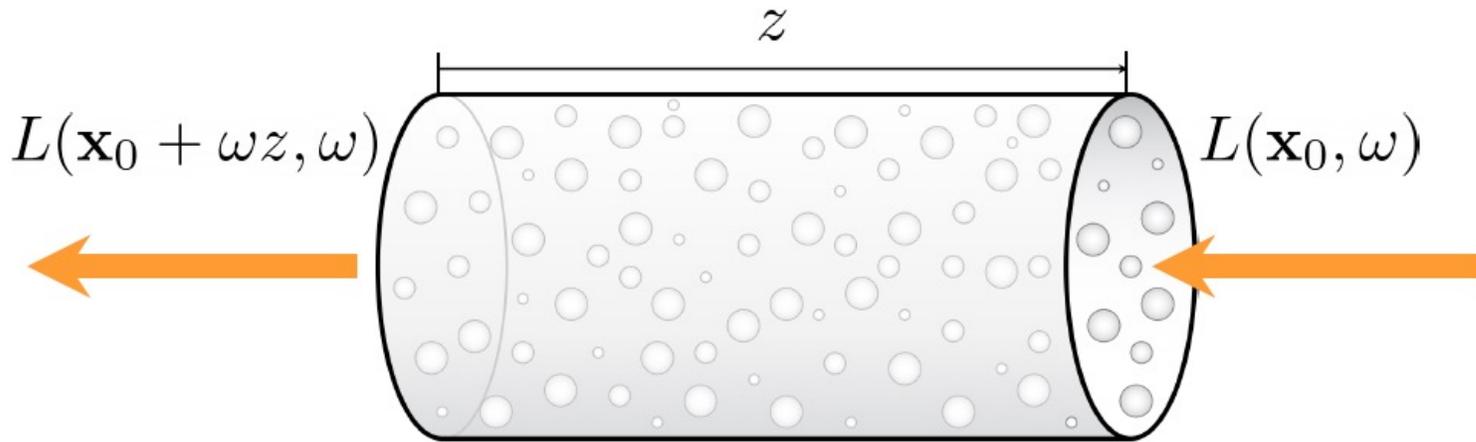
$$dL(\mathbf{x}, \vec{\omega}) = -\sigma_a(\mathbf{x})L(\mathbf{x}, \vec{\omega})dz - \sigma_s(\mathbf{x})L(\mathbf{x}, \vec{\omega})dz \\ + \sigma_a(\mathbf{x})L_e(\mathbf{x}, \vec{\omega})dz + \sigma_s(\mathbf{x})L_s(\mathbf{x}, \vec{\omega})dz$$



Will drop the subscript moving forward

$$\sigma_a \equiv \sigma$$

# Absorption-only Volume Rendering



**Q: What if we have a homogenous medium?** (uniform coefficient)

$$dL(\mathbf{x}, \omega) = -\sigma L(\mathbf{x}, \omega) dz$$

**Can you prove this?**

$$L(\mathbf{x}_0 + \omega z, \omega) = e^{-\sigma z} L(\mathbf{x}_0, \omega)$$

What if we have a non-homogenous medium?  
In non-homogenous medium, coefficient of absorption  $\sigma$  varies with location

$$T(\mathbf{x}, \mathbf{y}) = e^{-\int_{t=0}^z \sigma(\mathbf{x} + \omega \mathbf{t}) dt} L(\mathbf{x}_0, \omega)$$

**Transmittance**

# Transmittance

$$T(\mathbf{x}, \mathbf{y}) = e^{-\int_{t=0}^z \sigma(\mathbf{x} + \omega \mathbf{t}) dt} L(\mathbf{x}_0, \omega)$$

**Transmittance**

$$T(\mathbf{x}, \mathbf{y})$$

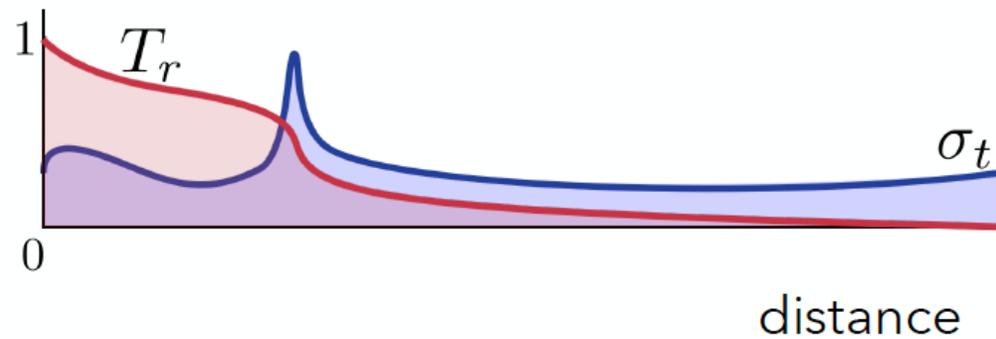
What fraction of radiance at  $\mathbf{x}$  in direction of  $\mathbf{y}$ , reaches  $\mathbf{y}$ ?  
(along a straight line under absorption-only model)

**Homogenous Medium:**

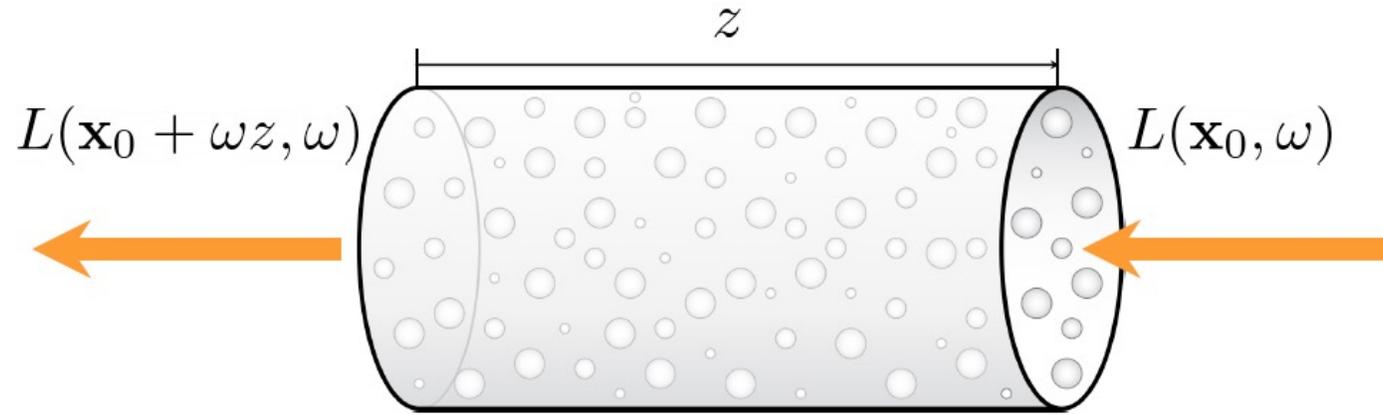
$$e^{-\sigma \|\mathbf{x} - \mathbf{y}\|}$$

**Non-Homogenous Medium:**

$$e^{-\int_{t=0}^{\|\mathbf{x} - \mathbf{y}\|} \sigma(\mathbf{x} + \omega \mathbf{t}) dt}$$

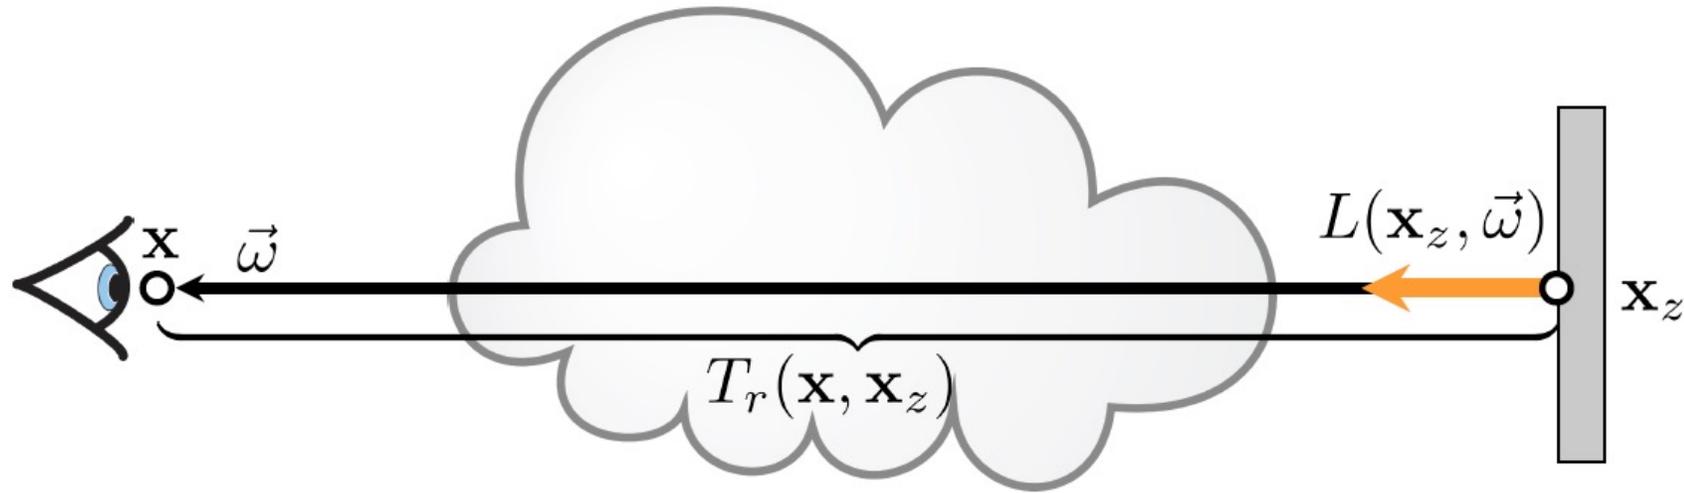


# Absorption-only Volume Rendering



$$L(\mathbf{x}_0 + \omega z, \omega) = T(\mathbf{x}_0, \mathbf{x}_0 + \omega z)L(\mathbf{x}_0, \omega)$$

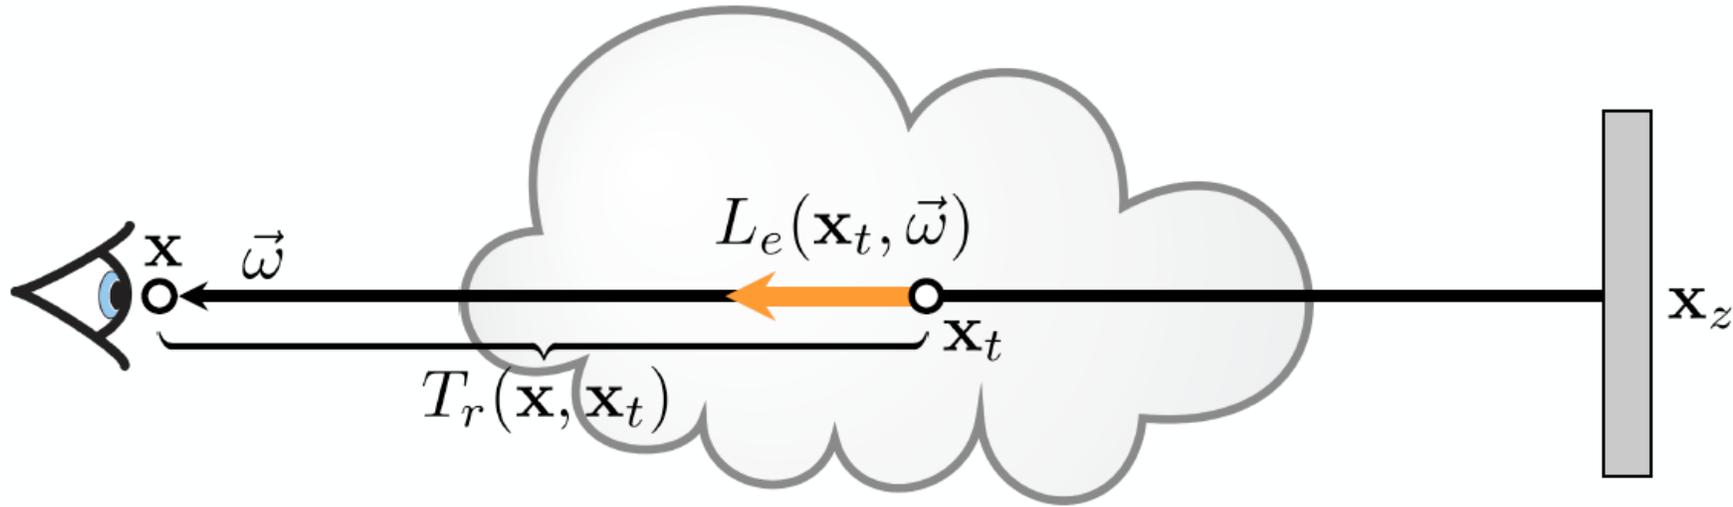
# Absorption-only Volume Rendering



$$L(\mathbf{x}, \omega) = T(\mathbf{x}, \mathbf{x}_z)L(\mathbf{x}_z, \omega)$$

**Radiance from 'outside' the medium**

# Emission-Absorption Volume Rendering



$$L(\mathbf{x}, \omega) = T(\mathbf{x}, \mathbf{x}_z)L(\mathbf{x}_z, \omega)$$

$$+ \int_0^z T(\mathbf{x}, \mathbf{x}_t)\sigma(\mathbf{x}_t)L_e(\mathbf{x}_t, \omega)dt$$

**Accumulated Emitted Radiance from inside**

# Emission-Absorption Volume Rendering

## Special Case: Homogenous **non-emitting** medium

$$T(\mathbf{x}, \mathbf{y}) = e^{-\int_{t=0}^z \sigma(\mathbf{x} + \omega \mathbf{t}) dt} L(\mathbf{x}_0, \omega)$$

**Transmittance**

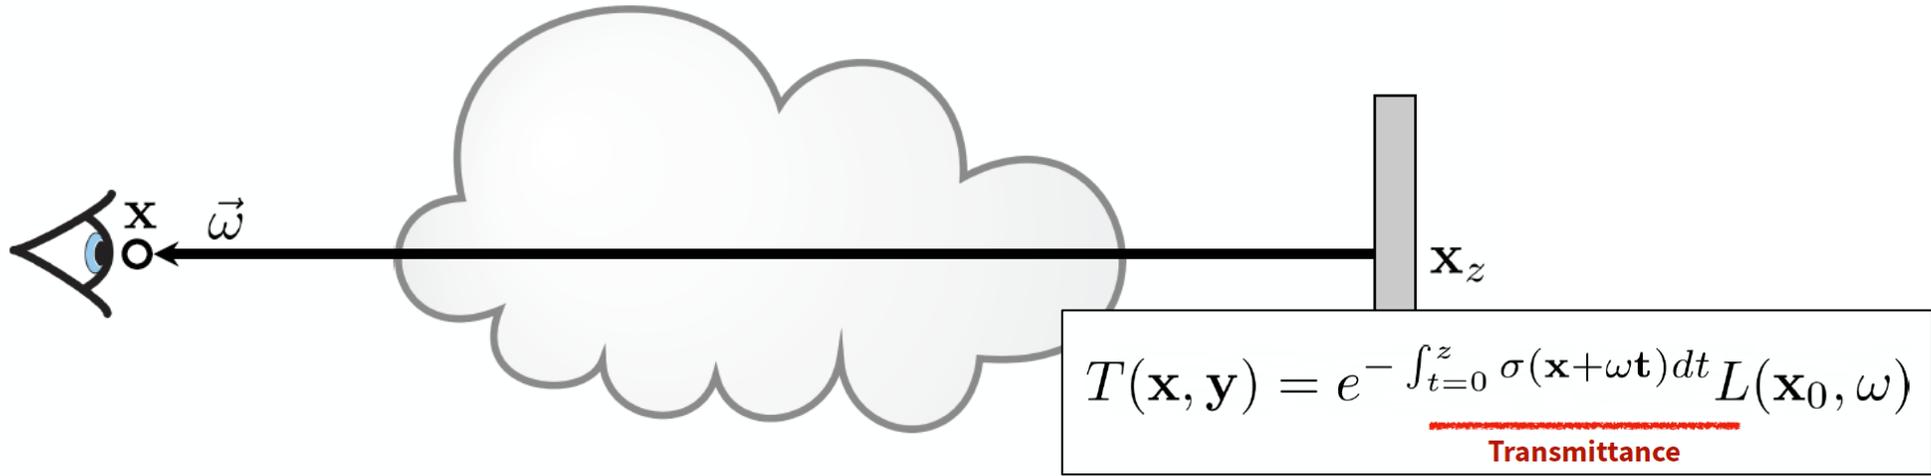


$$L(\mathbf{x}, \omega) = T(\mathbf{x}, \mathbf{x}_z) L(\mathbf{x}_z, \omega)$$

$$T(\mathbf{x}, \mathbf{x}_z) = e^{-\sigma \Delta}$$

# Emission-Absorption Volume Rendering

## Special Case: Homogenous **emitting** medium



$$L(\mathbf{x}, \omega) = e^{-\sigma \Delta} L(\mathbf{x}_z, \omega)$$

$$+ \int_0^z T(\mathbf{x}, \mathbf{x}_t) \sigma(\mathbf{x}_t) L_e(\mathbf{x}_t, \omega) dt$$

$$L_e(\mathbf{x}, \omega) \equiv C$$



$$\int \lambda e^{-\lambda x} dx = -e^{-\lambda x} + c$$

$$L(\mathbf{x}, \omega) = e^{-\sigma \Delta} L(\mathbf{x}_z, \omega)$$

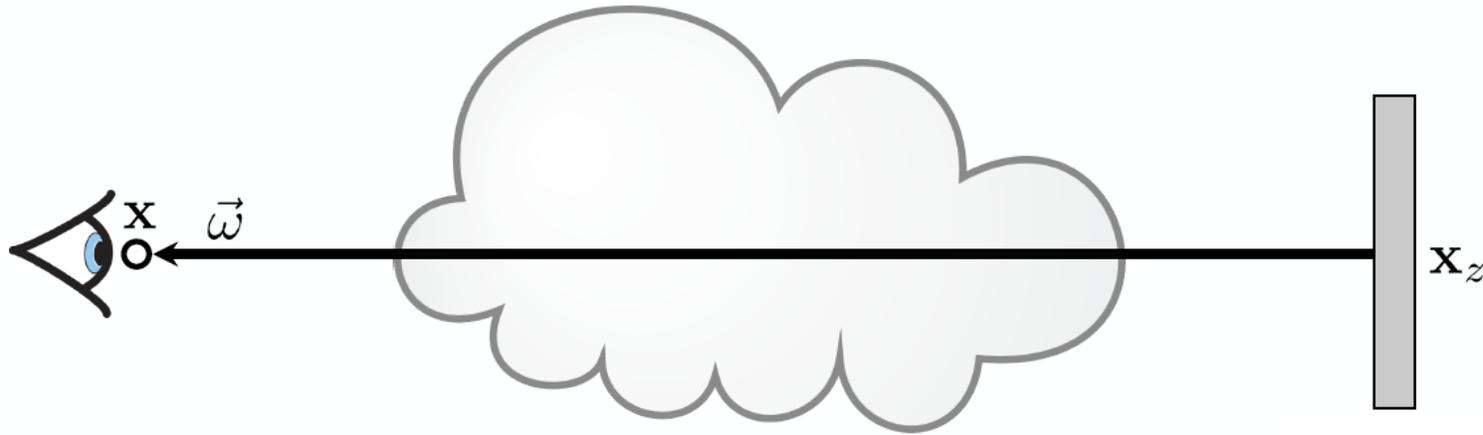
$$+ \int_{\delta=0}^{\Delta} e^{-\sigma \delta} \sigma C d\delta$$



$$L(\mathbf{x}, \omega) = e^{-\sigma \Delta} L(\mathbf{x}_z, \omega) + (1 - e^{-\sigma \Delta}) C$$

# Emission-Absorption Volume Rendering

## Special Case: Homogenous **emitting (only)** medium



$$L(\mathbf{x}, \omega) = \int_0^z T(\mathbf{x}, \mathbf{x}_t) \sigma(\mathbf{x}_t) L_e(\mathbf{x}_t, \omega) dt$$

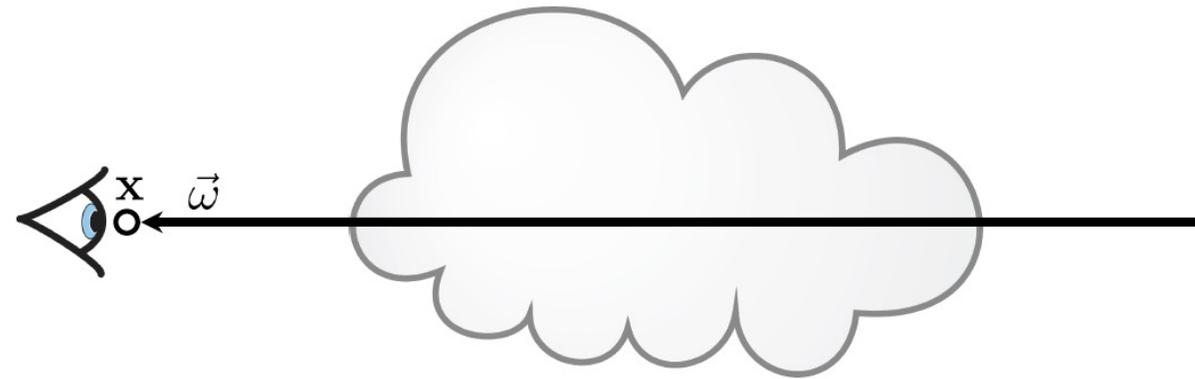
**Let's ignore light from outside medium (for now)**

$$L(\mathbf{x}, \omega) = e^{-\sigma \Delta} L(\mathbf{x}_z, \omega)$$

$$+ \int_0^z T(\mathbf{x}, \mathbf{x}_t) \sigma(\mathbf{x}_t) L_e(\mathbf{x}_t, \omega) dt$$



$$L(\mathbf{x}, \omega) = e^{-\sigma \Delta} L(\mathbf{x}_z, \omega) + (1 - e^{-\sigma \Delta}) C$$



$$L(\mathbf{x}, \omega) = (1 - e^{-\sigma \Delta}) C$$

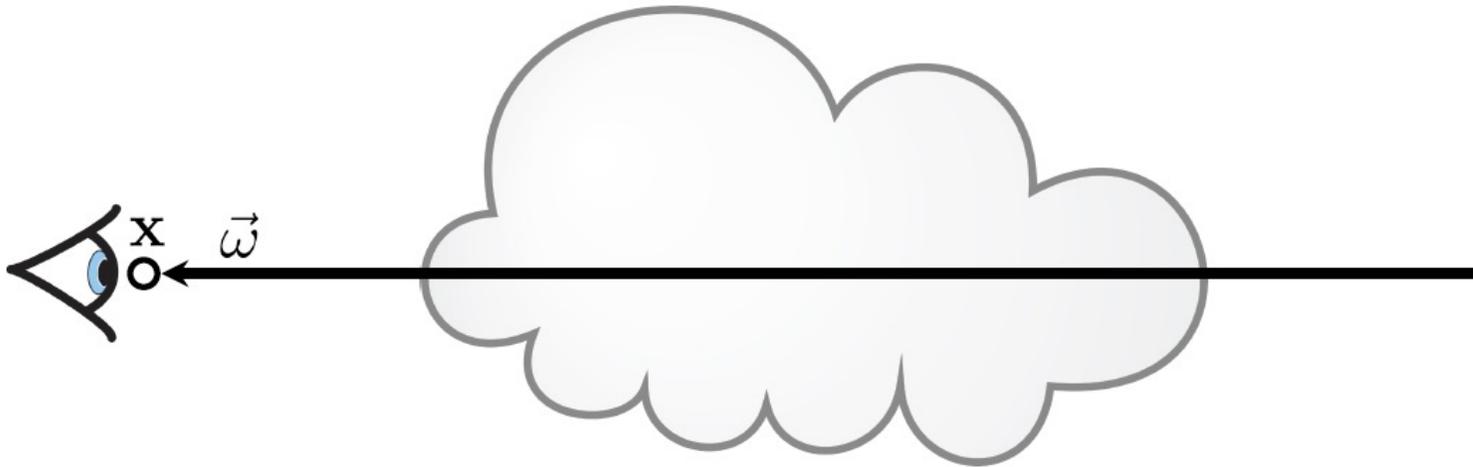
**Emission from a homogenous segment of length  $\Delta$**

# Emission-Absorption Volume Rendering

## Special Case: Homogenous **emitting (only)** medium

Volume Rendering Model to be used in NeRF

Assumption: Ignore light from outside medium.

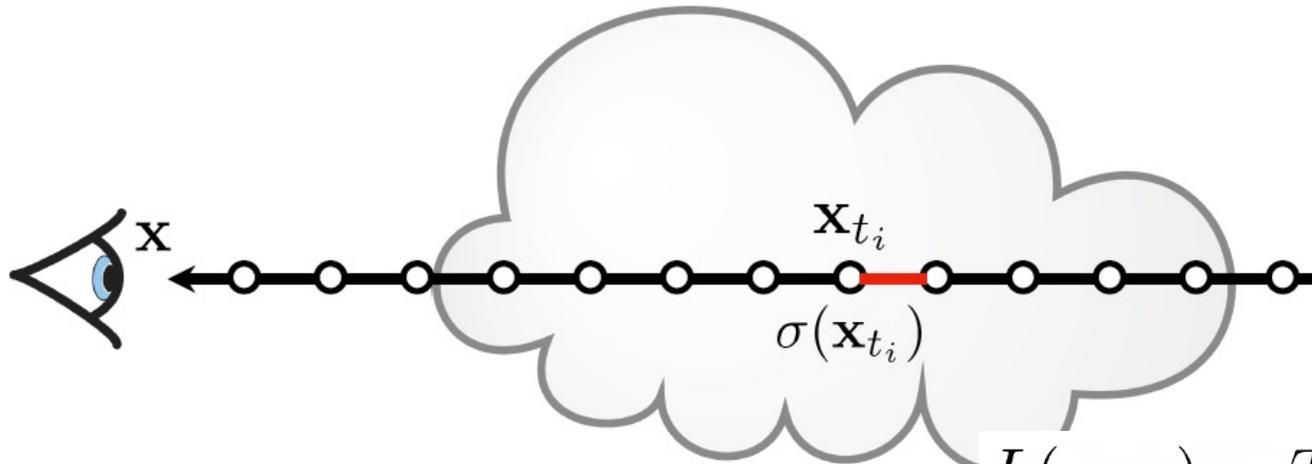


For NeRF this means: the object you are trying to render only emits light, There is no lighting being emitted by the background.

$$L(\mathbf{x}, \omega) = (1 - e^{-\sigma \Delta})C$$

Emission from a homogenous segment of length  $\Delta$

# Computational Volume Rendering: Ray Marching



$$L(\mathbf{x}, \omega) = T(\mathbf{x}, \mathbf{x}_z) L(\mathbf{x}_z, \omega)$$

$$L(\mathbf{x}, \omega) = \sum_{i=1}^N (\text{contribution from } i^{\text{th}} \text{ segment}) \rightarrow L(\mathbf{x}, \omega) = \sum_{i=1}^N T(\mathbf{x}, \mathbf{x}_{t_i}) \cdot \text{emission from } i^{\text{th}} \text{ segment}$$

**Approximate with a discrete sum**

$\mathbf{x}_{t_i}$  :  $i^{\text{th}}$  sample along ray at depth  $t_i$

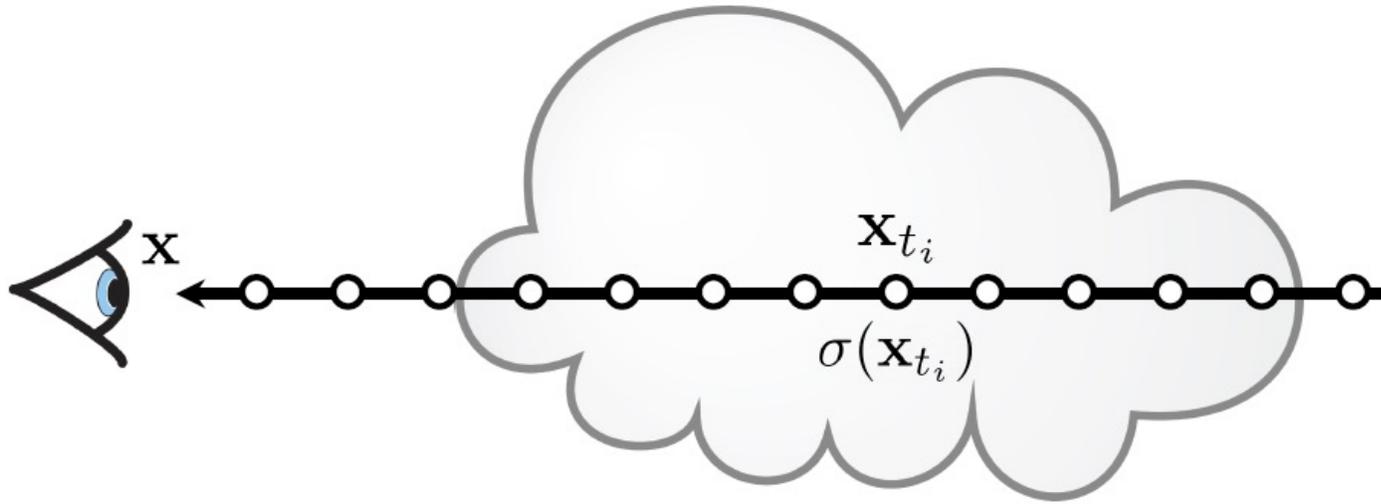
$\Delta t$  : distance between successive samples



$$L(\mathbf{x}, \omega) = (1 - e^{-\sigma \Delta}) C$$

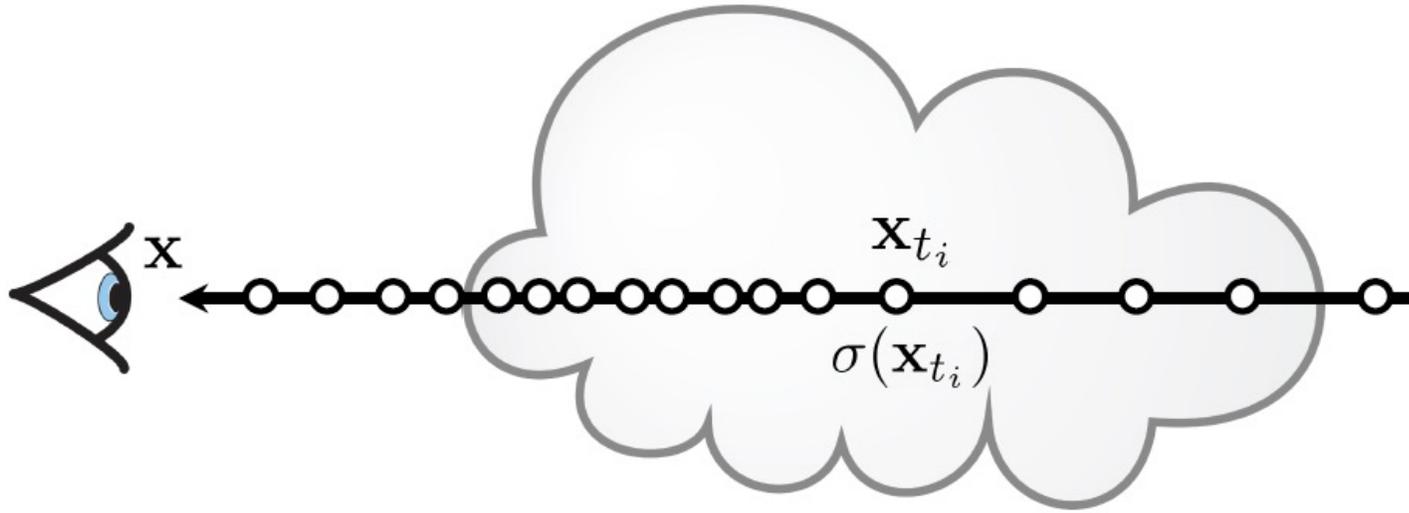
$$L(\mathbf{x}, \omega) = \sum_{i=1}^N T(\mathbf{x}, \mathbf{x}_{t_i}) \cdot (1 - e^{-\sigma_{t_i} \Delta t}) L_e(\mathbf{x}_{t_i}, \omega)$$

# Computational Volume Rendering: Ray Marching



1. Draw uniform samples along a ray (N segments, or N+1 points)
2. Compute transmittance between camera and each sample
3. Aggregate contributions across segments to get overall radiance (color)

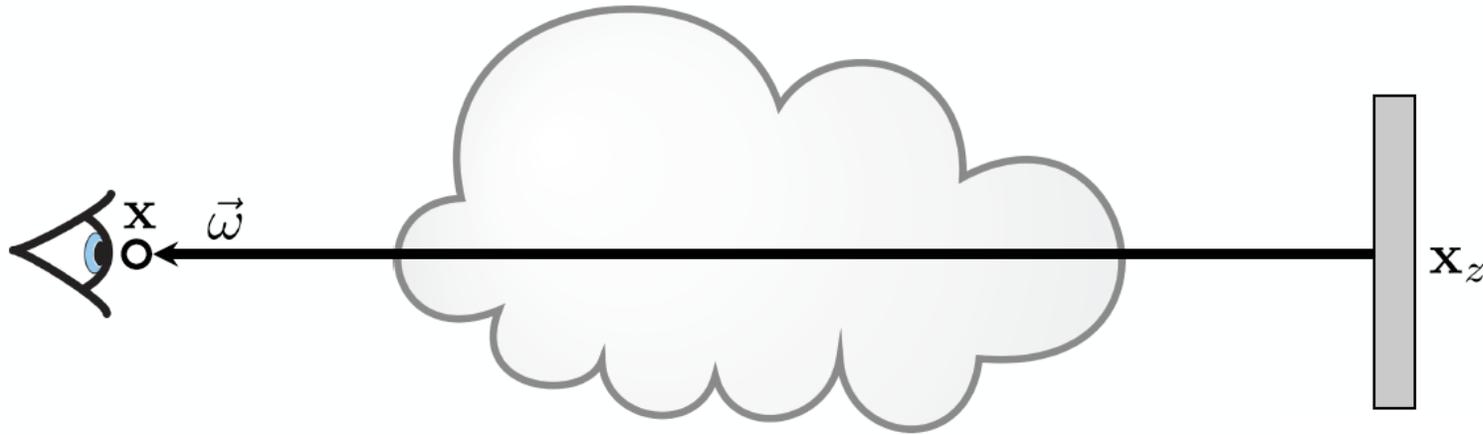
# Computational Volume Rendering: Ray Marching



1. Draw **non-uniform** samples along a ray
2. Compute transmittance between camera and each sample
3. Aggregate contributions across segments to get overall radiance (color)

# Emission-Absorption Volume Rendering

## Special Case: Homogenous **emitting (only)** medium



$$L(\mathbf{x}, \omega) = \int_0^z T(\mathbf{x}, \mathbf{x}_t) \sigma(\mathbf{x}_t) L_e(\mathbf{x}_t, \omega) dt$$

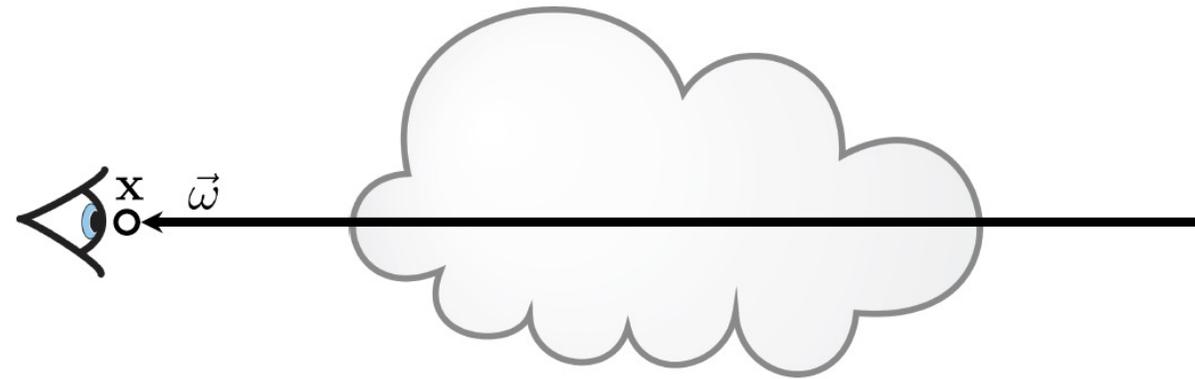
**Let's ignore light from outside medium (for now)**

$$L(\mathbf{x}, \omega) = e^{-\sigma \Delta} L(\mathbf{x}_z, \omega)$$

$$+ \int_0^z T(\mathbf{x}, \mathbf{x}_t) \sigma(\mathbf{x}_t) L_e(\mathbf{x}_t, \omega) dt$$



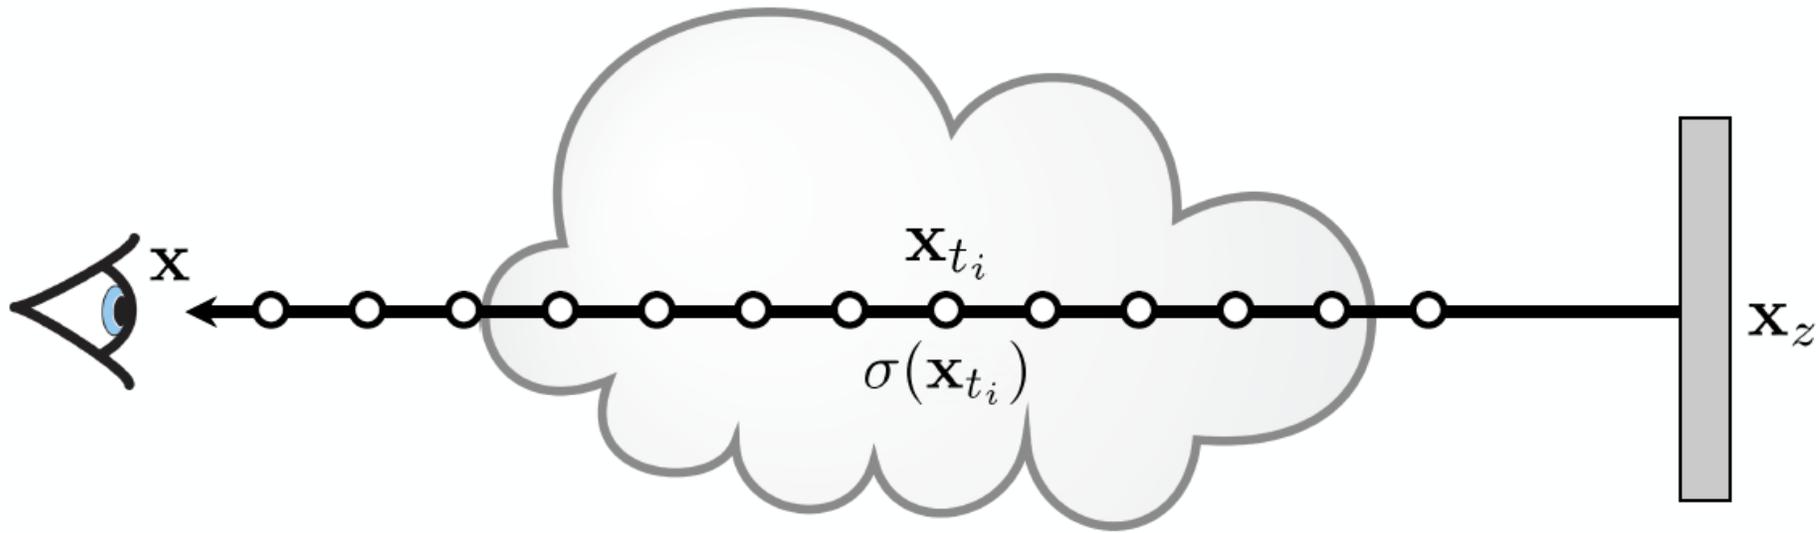
$$L(\mathbf{x}, \omega) = e^{-\sigma \Delta} L(\mathbf{x}_z, \omega) + (1 - e^{-\sigma \Delta}) C$$



$$L(\mathbf{x}, \omega) = (1 - e^{-\sigma \Delta}) C$$

**Emission from a homogenous segment of length  $\Delta$**

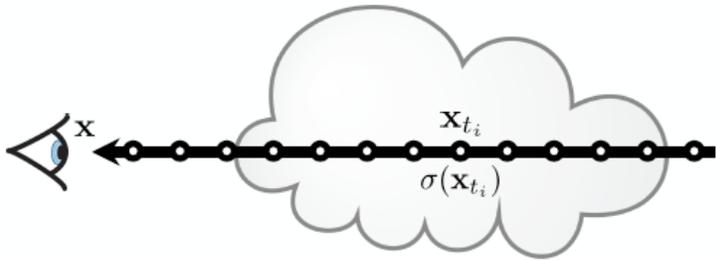
# Enabling Background Radiance



$$L(\mathbf{x}, \omega) = \sum_{i=1}^N T(\mathbf{x}, \mathbf{x}_{t_i})(1 - e^{-\sigma_{t_i} \Delta t}) L_e(\mathbf{x}_{t_i}, \omega)$$

$$T(\mathbf{x}, \mathbf{x}_{t_i}) = T(\mathbf{x}, \mathbf{x}_{t_{i-1}}) e^{-\sigma_{t_{i-1}} \Delta t}$$

# Computational Volume Rendering: A summary



$$\begin{array}{l} \sigma_{t_i} \equiv \sigma(\mathbf{x}_{t_i}) \\ L_e(\mathbf{x}_{t_i}, \omega) \end{array} \longrightarrow L(\mathbf{x}, \omega)$$

$$L(\mathbf{x}, \omega) = \sum_{i=1}^N T(\mathbf{x}, \mathbf{x}_{t_i}) \cdot (1 - e^{-\sigma_{t_i} \Delta t}) \underline{L_e(\mathbf{x}_{t_i}, \omega)}$$

If we can compute:

a) (per-point) density

b) (per-point, direction) emitted light,

we can render **any** ray through the medium

Equivalently, we can render an image from any camera viewpoint (using H\*W rays)

Note: **Differentiable** process w.r.t. the **density, emitted light**

and also camera parameters if density, emission are differentiable functions of position, direction

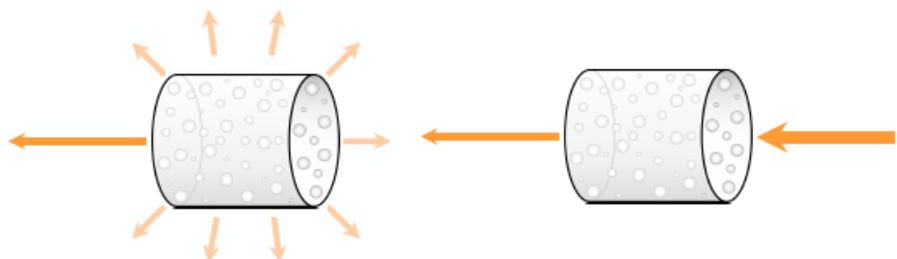
# Volume Rendering in NeRF

$$L(\mathbf{x}, \omega) = \sum_{i=1}^N T(\mathbf{x}, \mathbf{x}_{t_i}) \underbrace{(1 - e^{-\sigma_{t_i} \Delta t})}_{\text{opacity}} L_e(\mathbf{x}_{t_i}, \omega)$$

$L_e(\cdot)$  = RGB color of cloud of tiny particles.

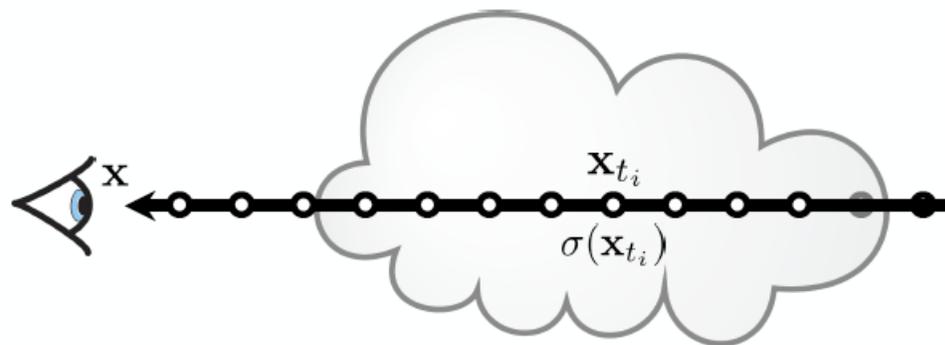
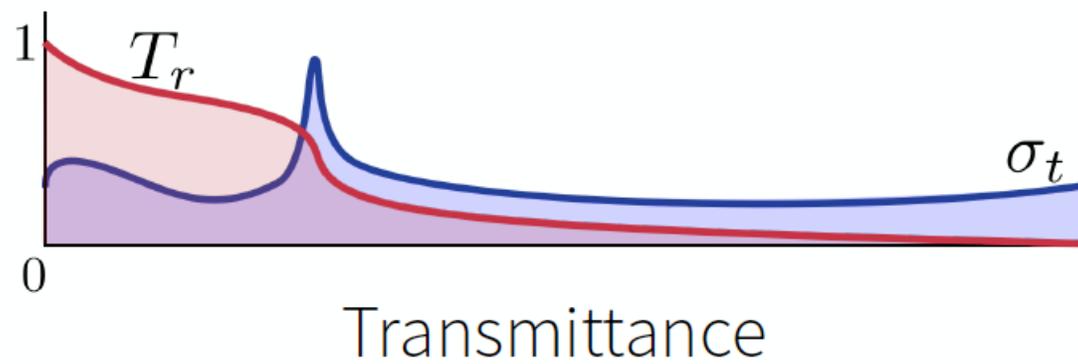
$\sigma$  = density of tiny colored particles

# Summary



$$dL(\mathbf{x}, \vec{\omega}) = -\sigma_a(\mathbf{x})L(\mathbf{x}, \vec{\omega})dz + \sigma_a(\mathbf{x})L_e(\mathbf{x}, \vec{\omega})dz$$

Emission-Absorption Model



$$L(\mathbf{x}, \omega) = \sum_{i=1}^N T(\mathbf{x}, \mathbf{x}_{t_i}) \cdot (1 - e^{-\sigma_{t_i} \Delta t}) L_e(\mathbf{x}_{t_i}, \omega)$$
$$T(\mathbf{x}, \mathbf{x}_{t_i}) = T(\mathbf{x}, \mathbf{x}_{t_{i-1}}) e^{-\sigma_{t_i} \Delta t}$$

A computational algorithm

# Slide Credits

- [“Neural Fields in Computer Vision”](#), CVRP 2022 Tutorial.
- Shubham Tulsiani, “Learning for 3D Vision”, Spring 2022, CMU
- Leo Guibas, JJ Park, “Neural Models for 3D geometry”, Spring 2022, Stanford.