



Volume Rendering of Neural Implicit Surfaces

Lior Yariv

Jiatao Gu

Yoni Kasten Yaron Lipman

Weizmann Institute of Science

Facebook AI Research

INTRODUCTION

Previous works and
paper contribution

01

RELATED WORK

Neural Scene Representation
Multi-view 3D Reconstruction

02

METHODS

SDF as a geometric function
Volume rendering
Bounding the error
Sampling Algorithm
Training

03

TABLE OF CONTENTS

04

EXPERMENTS

Qualitative and
quantitative results

05

CONCLUSION

Wrapping it up

06

SUPPLEMENTALS

Extra reference material

01


INTRODUCTION

Background and Related Works






INTRODUCTION

- Volume Rendering
 - Neural Radiance Field (NeRF)
 - In nutshell
 - Drawbacks?
 - VolSDF Contribution
 - New density representation and reconstruction
 - Moving from generic to geometric density function
- 



RELATED WORKS

- Neural Scene Representation & Rendering
 - Combining neural implicit functions with volume rendering
 - + Expressive representation power
 - + low memory foot-print
 - Recovered geometry
 - Opacity function approximation
- 

RELATED WORKS

- Multi-view 3D Reconstruction
 - Depth-based approaches
 - Complex pipeline
 - Voxel-based approaches
 - + Directly model objects in a volume
 - Limited to low resolution
 - Require accurate object masks





03

METHODS

A New Geometric Model Density


(1/5) Density As Transformed SDF


- What is SDF?
- Drawback of current model
- Improvement of previous model
- Advantages



(1/5) Density As Transformed SDF

- What is SDF?
 - A signed distance function is an n-dimensional implicit function, which associates a scalar value with each point of its n-dimensional domain.





volumetric density



signed distance function




produced rendering



(1/5) Density As Transformed SDF

- What is SDF?
- Drawback of current model



(1/5) Density As Transformed SDF

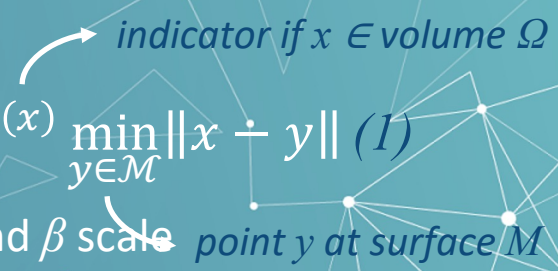
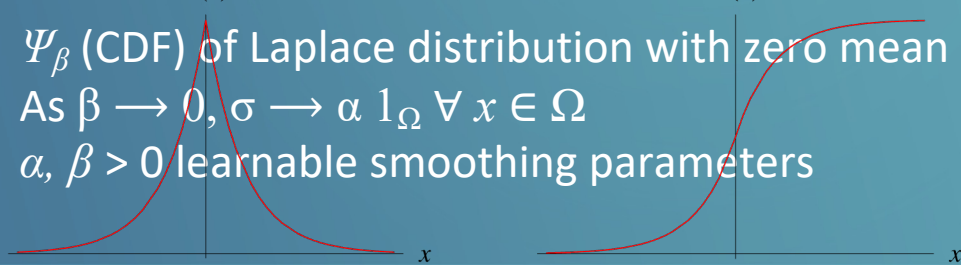
- What is SDF?
- Drawback of current model



- Geometric volume density at point x

$$\sigma(x) = \alpha \Psi_{\beta}(-d_{\Omega}(x)) \quad (2) \quad d_{\Omega}(x) = (-1)^{1_{\Omega}(x)} \min_{y \in \mathcal{M}} \|x - y\| \quad (1)$$

Ψ_{β} (CDF) of Laplace distribution with zero mean and β scale
 As $\beta \rightarrow 0, \sigma \rightarrow \alpha 1_{\Omega} \forall x \in \Omega$
 $\alpha, \beta > 0$ learnable smoothing parameters



(1/5) Density As Transformed SDF

- What is SDF?
- Drawback of current model



- Geometric volume density

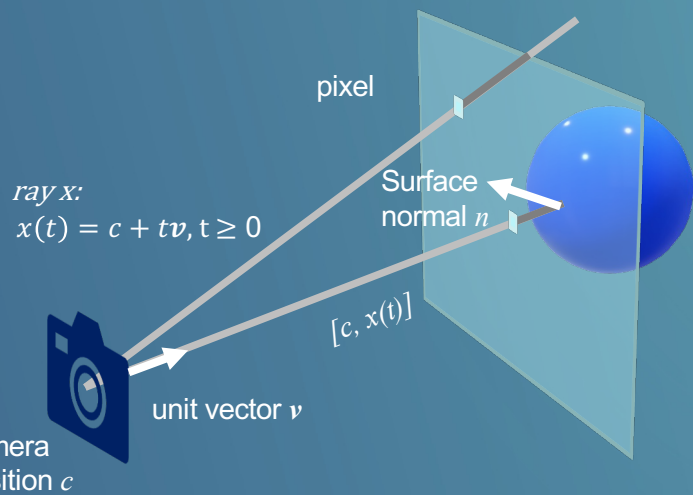
$$\sigma(x) = \alpha \Psi_{\beta}(-d_{\Omega}(x)) \quad (2)$$

- Advantages



(2/5) Volume Rendering of σ

- Volume rendering integral
 - approximating the integrated (i.e., summed) light radiance along this ray reaching the camera



The probability a light particle succeeds traversing the segment $[c, x(t)]$ without bouncing off

$$T(t) = \exp\left(-\int_0^t \sigma(x(s)) ds\right) \quad (4)$$


The volume rendering equation is the expected light along the ray,

$$I(c, v) = \int_0^\infty L(x(t), n(t), v) \tau(t) dt \quad (7)$$

where $L(x, n, v)$ is the radiance field

(2/5) Volume Rendering of σ

- Opacity as CDF:



$$T(t) = \exp\left(-\int_0^t \sigma(x(s)) ds\right) \quad (4)$$

$$\tau(t) = \sigma(x(t)) T(t) = \frac{dO}{dt}(t) \quad \text{PDF}$$



(2/5) Volume Rendering of σ



- Sampling τ
 - PDF τ is typically extremely concentrated near the object's boundary
 - the choice of the sample points S has a crucial effect on the approximation quality
 - NeRF: second, coarse network was trained specifically for the approximation of the opacity
 - VolSDF: sampling S is computed by a sampling algorithm based on an error bound for the opacity approximation



(2/5) Volume Rendering of σ

VoISFD

NeRF



(2/5) Volume Rendering of σ

VoISFD

NeRF



(3/5) Bound on The Opacity Approximation Error

- Transparency for ray x with sample point s :

$$T(t) = \exp\left(-\int_0^t \sigma(x(s)) ds\right)$$

$$O(t) = 1 - T(t)$$

$$O(t) = 1 - \exp\left(-\int_0^t \sigma(x(s)) ds\right)$$

$$O(t) = 1 - \exp(-\hat{R}(t))$$

$$\hat{R}(t) = \sum_{i=1}^{k-1} \delta_i \sigma_i + (t - t_k) \sigma_k$$



(3/5) Bound on The Opacity Approximation Error


- What about the error in summation?

$$\sigma(x) = \alpha \Psi_{\beta}(-d_{\Omega}(x))$$

$$\Psi_{\beta} = \begin{cases} \frac{1}{2} \exp\left(\frac{s}{\beta}\right), s \leq 0 \\ 1 - \frac{1}{2} \exp\left(-\frac{s}{\beta}\right), s > 0 \end{cases}$$

$$\frac{d}{dt} \Psi_{\beta} = \frac{1}{2\beta} \exp\left(-\frac{|s|}{\beta}\right)$$

$$\left| \frac{d}{ds} \sigma(x(s)) \right| \leq \frac{\alpha}{2\beta} \exp\left(-\frac{d_i^*}{\beta}\right), \text{ where } d_i^* = \min_{\substack{s \in [t_i, t_{i+1}] \\ y \notin B_i \cup B_{i+1}}} \|x(s) - y\|$$



(3/5) Bound on The Opacity Approximation Error

- This bound can be used to derive an error bound for \hat{R}

$$\left| \frac{d}{ds} \sigma(x(s)) \right| \leq \frac{\alpha}{2\beta} \exp\left(-\frac{d_i^*}{\beta}\right), \text{ where } d_i^* = \min_{\substack{s \in [t_i, t_{i+1}] \\ y \notin B_i \cup B_{i+1}}} \|x(s) - y\|$$

$$\hat{R}(t) = \sum_{i=1}^{k-1} \delta_i \sigma_i + (t - t_k) \sigma_k$$

$$|E(t)| \leq \hat{E}(t) = \frac{\alpha}{4\beta} \left(\sum_{i=1}^{k-1} \delta_i^2 \exp\left(-\frac{d_i^*}{\beta}\right) + (t - t_k)^2 \exp\left(-\frac{d_k^*}{\beta}\right) \right)$$

- So the error of the appr. opacity \hat{O} can be bounded as

$$|O(t) - \hat{O}(t)| \leq \exp(-\hat{R}(t)) \left(\exp(\hat{E}(t)) - 1 \right)$$


(3/5) Bound on The Opacity Approximation Error

- Taking the maximum over all intervals furnishes a bound $B_{\mathcal{T},\beta}$ as a function of \mathcal{T} and β

$$|O(t) - \hat{O}(t)| \leq \exp(-\hat{R}(t)) \left(\exp(\hat{E}(t)) - 1 \right)$$

$$\max_{t \in [0, M]} |O(t) - \hat{O}(t)| \leq B_{\mathcal{T},\beta} = \max_{k \in [n-1]} \left\{ \exp(-\hat{R}(t_k)) \left(\exp(\hat{E}(t_{k+1})) - 1 \right) \right\}$$

- Where is \mathcal{T} a set of samples
 $\mathcal{T} = \{t_i\}_{i=1}^n, 0 = t_1 < \dots < t_n = M$



(4/5) Sampling Algorithm

- Using the bound to compute sampling:
- $I(\mathbf{c}, \mathbf{v}) = \int_0^\infty L(x(t), n(t), \mathbf{v})\tau(t)dt$
- $I(\mathbf{c}, \mathbf{v}) \approx \hat{I}_S(\mathbf{c}, \mathbf{v}) = \sum_{i=1}^{m-1} \hat{t}_i L_i$

Algorithm 1: Sampling algorithm.

Input: error threshold $\epsilon > 0$; β

- 1 Initialize $\mathcal{T} = \mathcal{T}_0$
 - 2 Initialize β_+ such that $B_{\mathcal{T}, \beta_+} \leq \epsilon$
 - 3 **while** $B_{\mathcal{T}, \beta} > \epsilon$ and not *max_iter* **do**
 - 4 upsample \mathcal{T}
 - 5 **if** $B_{\mathcal{T}, \beta_+} < \epsilon$ **then**
 - 6 Find $\beta_\star \in (\beta, \beta_+)$ so that
 $B_{\mathcal{T}, \beta_\star} = \epsilon$
 - 7 Update $\beta_+ \leftarrow \beta_\star$
 - 8 **end**
 - 9 **end**
 - 10 Estimate \hat{O} using \mathcal{T} and β_+
 - 11 $\mathcal{S} \leftarrow$ get fresh m samples using \hat{O}^{-1}
 - 12 **return** \mathcal{S}
-

(4/5) Sampling Algorithm

Setting β as:

$$\beta \geq \frac{\alpha M}{4(n-1) \log(1+\epsilon)}$$

For $n > 0$, $\epsilon > 0$ and $B_{\mathcal{T}, \beta} \leq \epsilon$

Here $n = 128$ was used.

Algorithm 1: Sampling algorithm.

Input: error threshold $\epsilon > 0$; β

- 1 Initialize $\mathcal{T} = \mathcal{T}_0$
 - 2 Initialize β_+ such that $B_{\mathcal{T}, \beta_+} \leq \epsilon$
 - 3 **while** $B_{\mathcal{T}, \beta} > \epsilon$ *and not max_iter* **do**
 - 4 upsample \mathcal{T}
 - 5 **if** $B_{\mathcal{T}, \beta_+} < \epsilon$ **then**
 - 6 Find $\beta_\star \in (\beta, \beta_+)$ so that
 $B_{\mathcal{T}, \beta_\star} = \epsilon$
 - 7 Update $\beta_+ \leftarrow \beta_\star$
 - 8 **end**
 - 9 **end**
 - 10 Estimate \hat{O} using \mathcal{T} and β_+
 - 11 $\mathcal{S} \leftarrow$ get fresh m samples using \hat{O}^{-1}
 - 12 **return** \mathcal{S}
-

(4/5) Sampling Algorithm

We initialize \mathcal{T} with uniform sampling \mathcal{T}_0 \longrightarrow

Algorithm 1: Sampling algorithm.

Input: error threshold $\epsilon > 0$; β

- 1 Initialize $\mathcal{T} = \mathcal{T}_0$
 - 2 Initialize β_+ such that $B_{\mathcal{T}, \beta_+} \leq \epsilon$
 - 3 **while** $B_{\mathcal{T}, \beta} > \epsilon$ *and not max_iter* **do**
 - 4 upsample \mathcal{T}
 - 5 **if** $B_{\mathcal{T}, \beta_+} < \epsilon$ **then**
 - 6 Find $\beta_\star \in (\beta, \beta_+)$ so that
 $B_{\mathcal{T}, \beta_\star} = \epsilon$
 - 7 Update $\beta_+ \leftarrow \beta_\star$
 - 8 **end**
 - 9 **end**
 - 10 Estimate \hat{O} using \mathcal{T} and β_+
 - 11 $\mathcal{S} \leftarrow$ get fresh m samples using \hat{O}^{-1}
 - 12 **return** \mathcal{S}
-

(4/5) Sampling Algorithm

pick $\beta_+ > \beta$ so that the error bound satisfies the required ϵ bound



Algorithm 1: Sampling algorithm.

Input: error threshold $\epsilon > 0$; β

- 1 Initialize $\mathcal{T} = \mathcal{T}_0$
 - 2 Initialize β_+ such that $B_{\mathcal{T}, \beta_+} \leq \epsilon$
 - 3 **while** $B_{\mathcal{T}, \beta} > \epsilon$ *and not max_iter* **do**
 - 4 upsample \mathcal{T}
 - 5 **if** $B_{\mathcal{T}, \beta_+} < \epsilon$ **then**
 - 6 Find $\beta_\star \in (\beta, \beta_+)$ so that
 $B_{\mathcal{T}, \beta_\star} = \epsilon$
 - 7 Update $\beta_+ \leftarrow \beta_\star$
 - 8 **end**
 - 9 **end**
 - 10 Estimate \hat{O} using \mathcal{T} and β_+
 - 11 $\mathcal{S} \leftarrow$ get fresh m samples using \hat{O}^{-1}
 - 12 **return** \mathcal{S}
-

(4/5) Sampling Algorithm

n samples are added to \mathcal{T} to reduce β_+
while keeping $B_{\mathcal{T},\beta}$ within error bound




Algorithm 1: Sampling algorithm.

Input: error threshold $\epsilon > 0$; β

- 1 Initialize $\mathcal{T} = \mathcal{T}_0$
 - 2 Initialize β_+ such that $B_{\mathcal{T},\beta_+} \leq \epsilon$
 - 3 **while** $B_{\mathcal{T},\beta} > \epsilon$ and not *max_iter* **do**
 - 4 upsample \mathcal{T}
 - 5 **if** $B_{\mathcal{T},\beta_+} < \epsilon$ **then**
 - 6 Find $\beta_\star \in (\beta, \beta_+)$ so that
 $B_{\mathcal{T},\beta_\star} = \epsilon$
 - 7 Update $\beta_+ \leftarrow \beta_\star$
 - 8 **end**
 - 9 **end**
 - 10 Estimate \hat{O} using \mathcal{T} and β_+
 - 11 $\mathcal{S} \leftarrow$ get fresh m samples using \hat{O}^{-1}
 - 12 **return** \mathcal{S}
-

(4/5) Sampling Algorithm

We use the bisection method (10 max iterations) to search for β_* and update β_+



Algorithm 1: Sampling algorithm.

Input: error threshold $\epsilon > 0$; β

- 1 Initialize $\mathcal{T} = \mathcal{T}_0$
 - 2 Initialize β_+ such that $B_{\mathcal{T}, \beta_+} \leq \epsilon$
 - 3 **while** $B_{\mathcal{T}, \beta} > \epsilon$ and not *max_iter* **do**
 - 4 upsample \mathcal{T}
 - 5 **if** $B_{\mathcal{T}, \beta_+} < \epsilon$ **then**
 - 6 Find $\beta_* \in (\beta, \beta_+)$ so that
 $B_{\mathcal{T}, \beta_*} = \epsilon$
 - 7 Update $\beta_+ \leftarrow \beta_*$
 - 8 **end**
 - 9 **end**
 - 10 Estimate \hat{O} using \mathcal{T} and β_+
 - 11 $\mathcal{S} \leftarrow$ get fresh m samples using \hat{O}^{-1}
 - 12 **return** \mathcal{S}
-

(4/5) Sampling Algorithm

Run iteratively until $B_{\mathcal{T},\beta} \leq \epsilon$ (5 max iter)



Algorithm 1: Sampling algorithm.

Input: error threshold $\epsilon > 0$; β

- 1 Initialize $\mathcal{T} = \mathcal{T}_0$
 - 2 Initialize β_+ such that $B_{\mathcal{T},\beta_+} \leq \epsilon$
 - 3 **while** $B_{\mathcal{T},\beta} > \epsilon$ and not *max_iter* **do**
 - 4 upsample \mathcal{T}
 - 5 **if** $B_{\mathcal{T},\beta_+} < \epsilon$ **then**
 - 6 Find $\beta_\star \in (\beta, \beta_+)$ so that
 $B_{\mathcal{T},\beta_\star} = \epsilon$
 - 7 Update $\beta_+ \leftarrow \beta_\star$
 - 8 **end**
 - 9 **end**
 - 10 Estimate \hat{O} using \mathcal{T} and β_+
 - 11 $\mathcal{S} \leftarrow$ get fresh m samples using \hat{O}^{-1}
 - 12 **return** \mathcal{S}
-

(4/5) Sampling Algorithm

Use final \mathcal{T} and β_+ to est. opacity \hat{O}

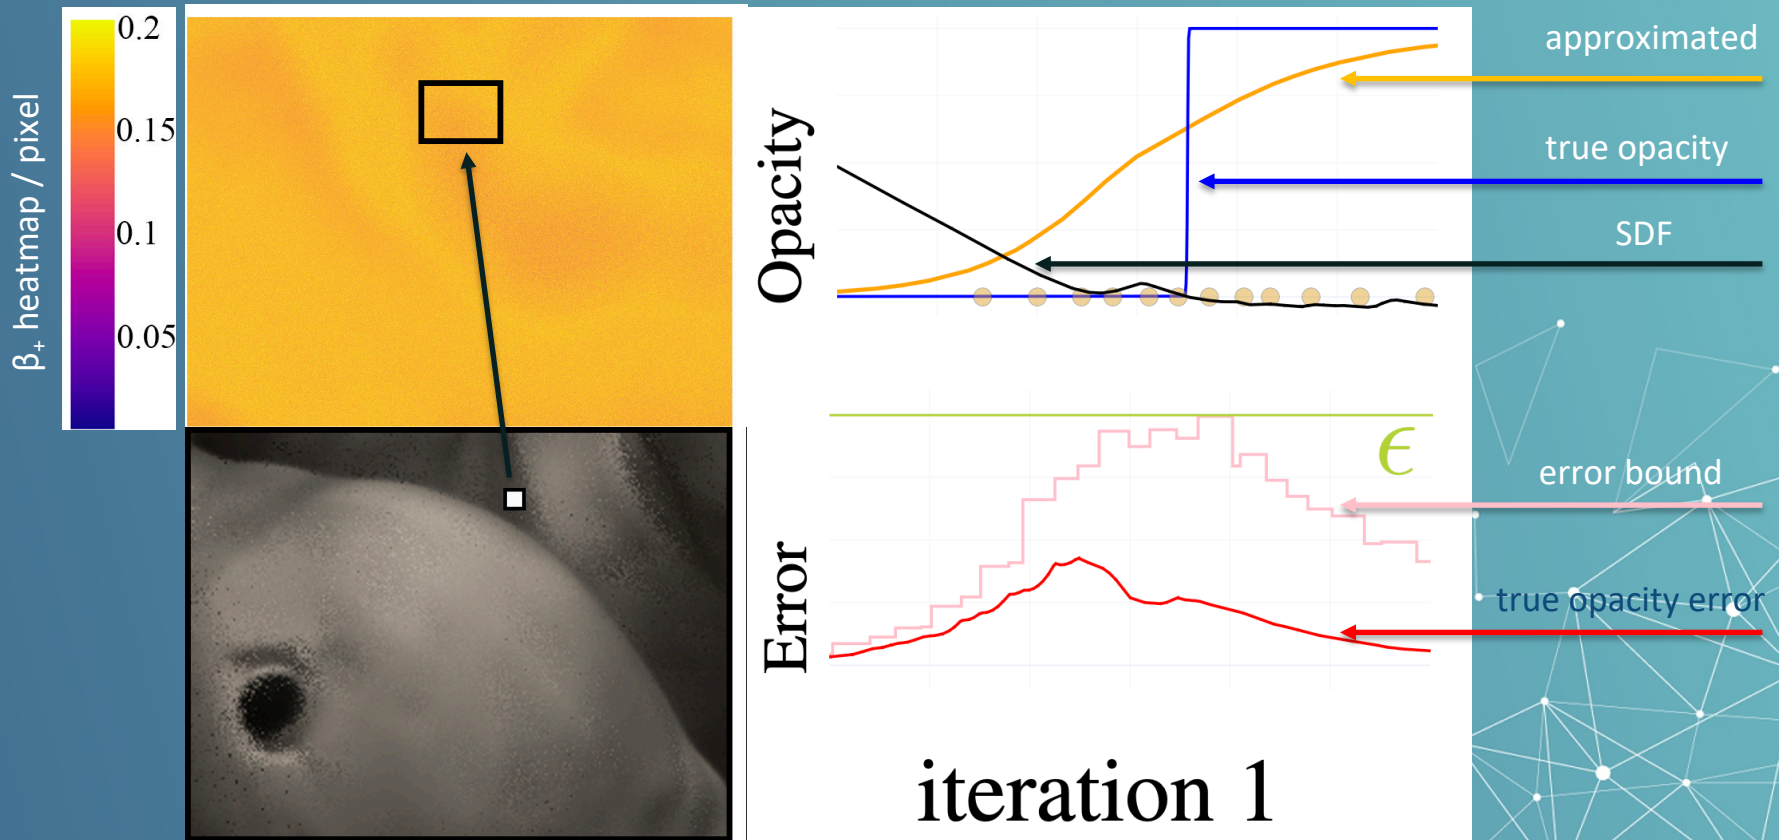


Algorithm 1: Sampling algorithm.



Input: error threshold $\epsilon > 0$; β

- 1 Initialize $\mathcal{T} = \mathcal{T}_0$
 - 2 Initialize β_+ such that $B_{\mathcal{T}, \beta_+} \leq \epsilon$
 - 3 **while** $B_{\mathcal{T}, \beta} > \epsilon$ and not *max_iter* **do**
 - 4 upsample \mathcal{T}
 - 5 **if** $B_{\mathcal{T}, \beta_+} < \epsilon$ **then**
 - 6 Find $\beta_\star \in (\beta, \beta_+)$ so that
 $B_{\mathcal{T}, \beta_\star} = \epsilon$
 - 7 Update $\beta_+ \leftarrow \beta_\star$
 - 8 **end**
 - 9 **end**
 - 10 Estimate \hat{O} using \mathcal{T} and β_+
 - 11 $\mathcal{S} \leftarrow$ get fresh m samples using \hat{O}^{-1}
 - 12 **return** \mathcal{S}
-

(4/5) Sampling Algorithm - Qualitative

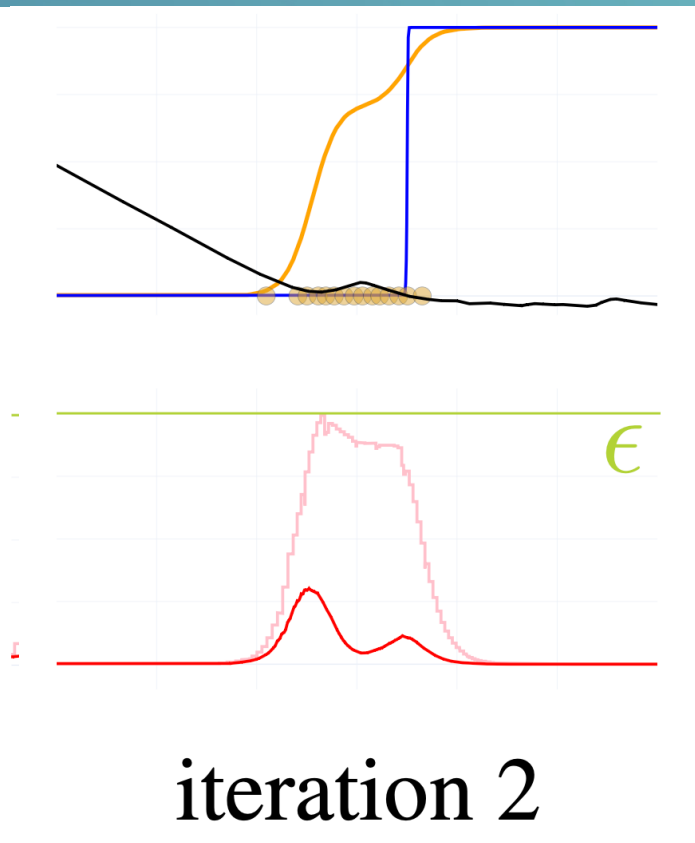


(4/5) Sampling Algorithm - Qualitative



Opacity

Error



approximated

true opacity

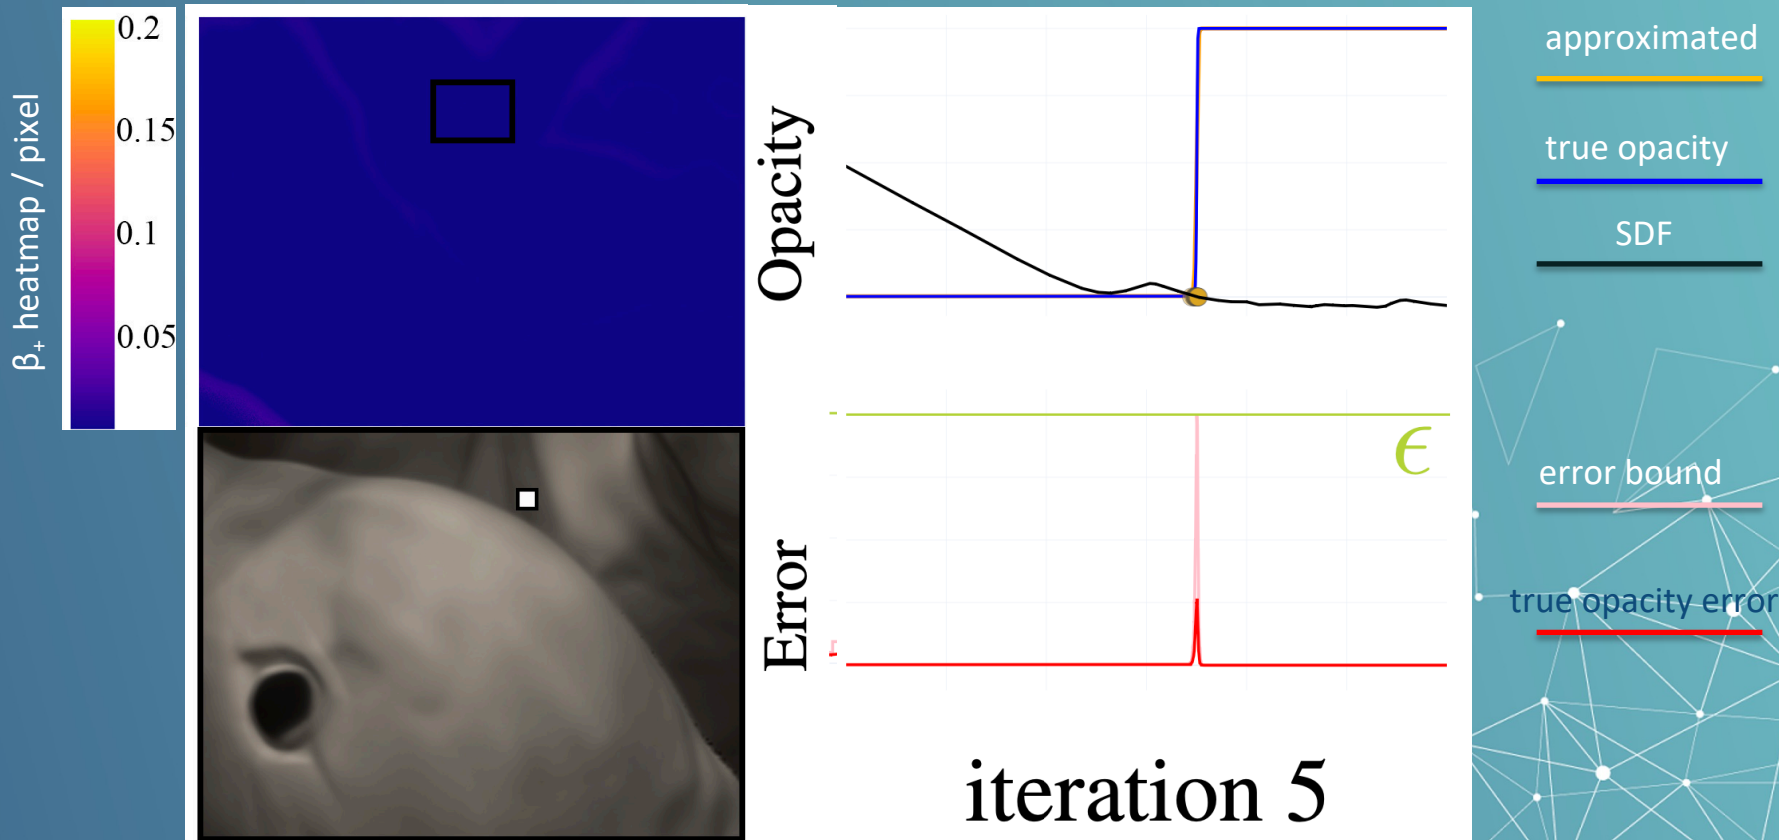
SDF

error bound

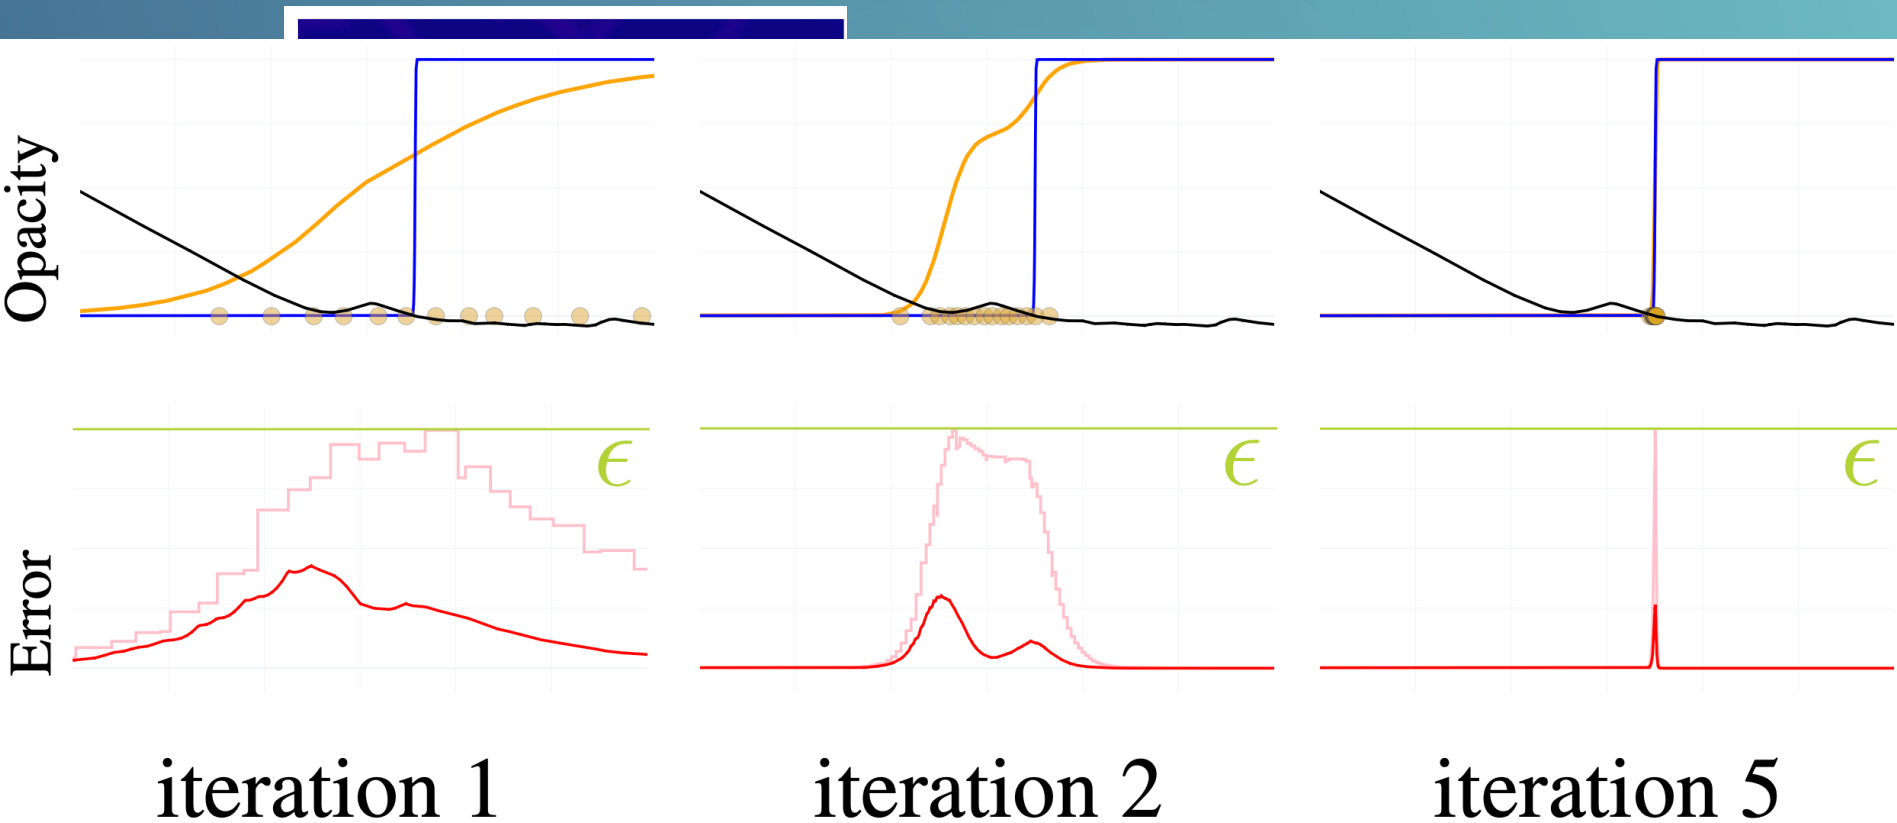
true opacity error

ϵ


(4/5) Sampling Algorithm - Qualitative



(4/5) Sampling Algorithm - Qualitative



(4/5) Sampling Algorithm - Qualitative




(5/5) Training

- 2x MLP:
 - Approximating the SDF of the learned geometry, and global geometry feature z of dimension 256 :
 - $f_\phi(x) = (d(x), z(x)) \in \mathbb{R}^{1+256}$
 - Presenting the scene's radiance field with learnable parameters ψ :
 - $L_\psi(x, n, v, z) \in \mathbb{R}^3$
 - Two scalar learnable parameters
 - $\alpha, \beta \in \mathbb{R}$, with $\alpha = \beta^{-1}$
 - Positional encoding for x and v , same as NeRF.



(5/5) Training

- For each pixel p a triplet (I_p, c_p, v_p)
 - $I_p \in \mathbb{R}^3$ is its intensity (RGB color)
 - $c_p \in \mathbb{R}^3$ is its camera location
 - $v_p \in \mathbb{R}^3$ is the viewing direction (camera to pixel)
- Training loss:
 - $\mathcal{L}(\theta) = \mathcal{L}_{\text{RGB}}(\theta) + \lambda \mathcal{L}_{\text{SDF}}(\varphi)$ (17)
 - $\mathcal{L}_{\text{RGB}}(\theta) = \mathbb{E}_p \left\| I_p - \hat{I}_S(c_p, v_p) \right\|_1$ (18) color loss
 - $\mathcal{L}_{\text{SDF}}(\varphi) = \mathbb{E}_z (\| \nabla d(z) \| - 1)^2$ (18) Eikonal loss










04

EXPERIMENTS

Method evaluation on the challenging task
of multiview 3D surface reconstruction

Multi-view 3D reconstruction

- Quantitative results for the **DTU dataset**
- DTU multi-view image; different objects; fixed camera and lighting parameters




Multi-view 3D reconstruction

- Quantitative results for the **DTU dataset**
- DTU multi-view image; different objects; fixed camera and lighting parameters




	Scan	24	37	40	55	63	65	69	83	97	105	106	110	114	118	122	Mean
Chamfer Distance	IDR	1.63	1.87	0.63	0.48	1.04	0.79	0.77	1.33	1.16	0.76	0.67	0.90	0.42	0.51	0.53	0.90
	colmap₇	0.45	0.91	0.37	0.37	0.90	1.00	0.54	1.22	1.08	0.64	0.48	0.59	0.32	0.45	0.43	0.65
	colmap₀	0.81	2.05	0.73	1.22	1.79	1.58	1.02	3.05	1.40	2.05	1.00	1.32	0.49	0.78	1.17	1.36
	NeRF	1.92	1.73	1.92	0.80	3.41	1.39	1.51	5.44	2.04	1.10	1.01	2.88	0.91	1.00	0.79	1.89
	VolSDF	1.14	1.26	0.81	0.49	1.25	0.70	0.72	1.29	1.18	0.70	0.66	1.08	0.42	0.61	0.55	0.86
PSNR	NeRF	26.24	25.74	26.79	27.57	31.96	31.50	29.58	32.78	28.35	32.08	33.49	31.54	31.0	35.59	35.51	30.65
	VolSDF	26.28	25.61	26.55	26.76	31.57	31.5	29.38	33.23	28.03	32.13	33.16	31.49	30.33	34.9	34.75	30.38

Multi-view 3D reconstruction




- Qualitative results for the DTU dataset





Colmap



NeRF



VoISDF



IDR




Multi-view 3D reconstruction

- Quantitative results for the **BlendedMVS** dataset
- Large collection of 113 scenes. High quality GT.
- 9 different scenes were selected


	Scene	Doll	Egg	Head	Angel	Bull	Robot	Dog	Bread	Camera	Mean
Chamfer l_1	Our Improvement (%)	54.0	91.2	24.3	75.1	60.7	27.2	47.7	34.6	51.8	51.8
	NeRF++	26.95	27.34	27.23	30.06	26.65	26.73	27.90	31.68	23.44	27.55
PSNR	VolSDF	25.49	27.18	26.36	29.79	26.01	26.03	28.65	31.24	22.97	27.08

Multi-view 3D reconstruction



- Qualitative results for the BlendedMVS dataset




GT











NeRF++



VoISDF





Multi-view 3D reconstruction

- Qualitative results for the BlendedMVS dataset



05

CONCLUSION

You can enter here a subtitle if you need it



CONCLUSIONS AND REMARKS

- The paper does not have a proof of correctness for the sampling algorithm.
- Representing non-watertight manifolds and/or manifolds with boundaries, such as zero thickness surfaces, is not possible with an SDF
- Assumption of homogeneous density; extending it to more general density models would allow representing a broader class of geometries



CONCLUSIONS AND REMARKS


- High quality geometries can be learned in an unsupervised manner.
- Accurate geometry reconstruction from images can be used for malice purposes.



LIMITATIONS



GT




Rendering




Geometry

(a)



(b)



Rendering



Geometry

(c)

06

SUPPLEMENTALS

You can enter here a subtitle if you need it







Thank you.