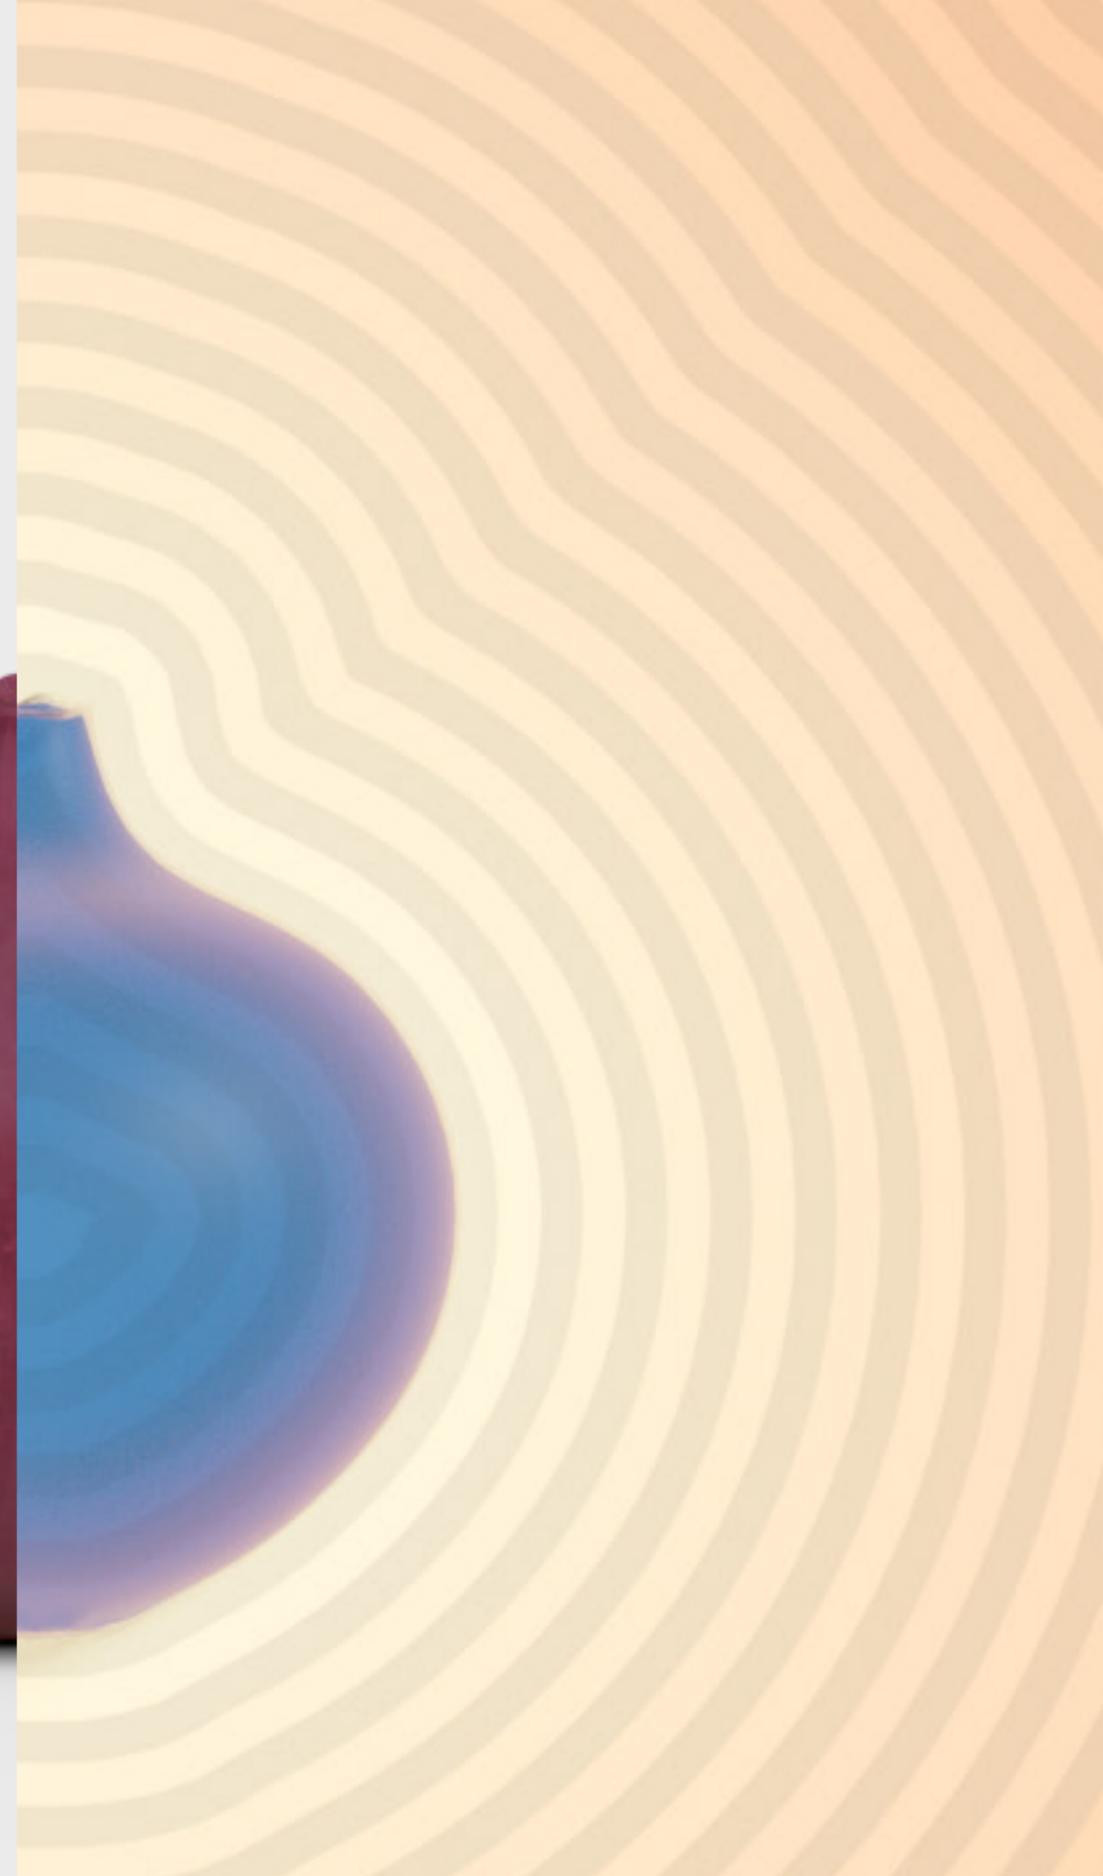
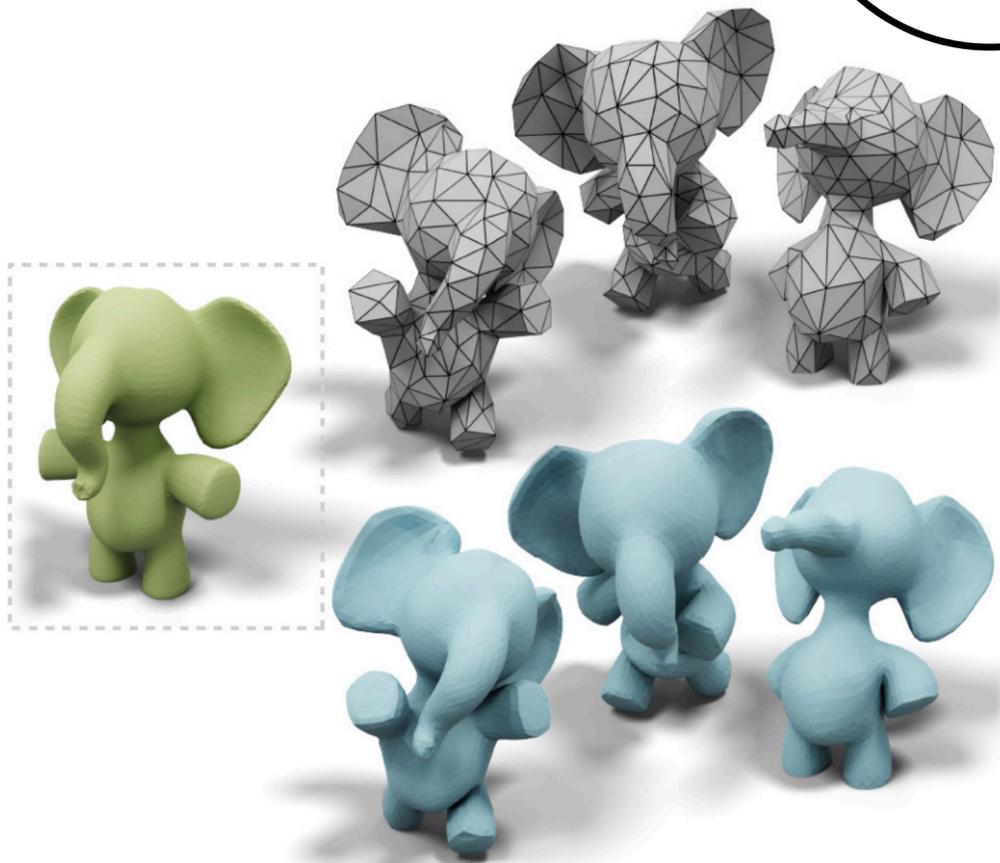


# Neural Implicit Shape Representations

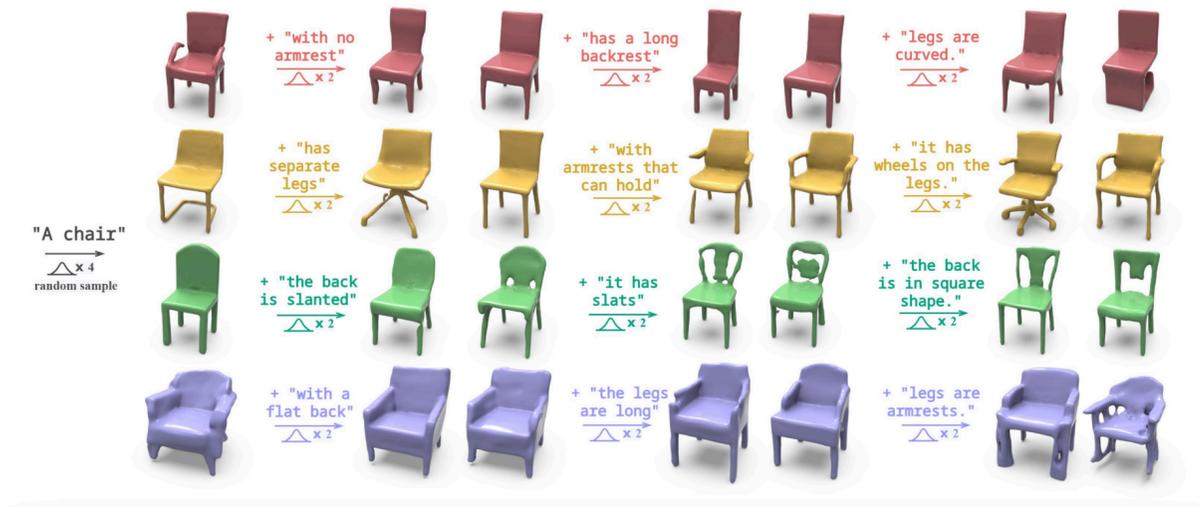
*Deep Learning for Geometry Processing*  
Daniel Ritchie and Zoë Marschner



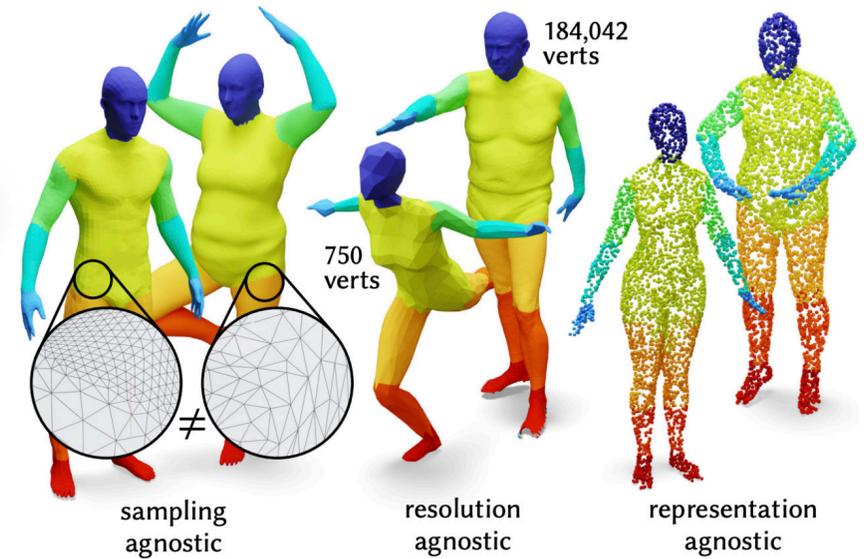
**Editing**  
(e.g., Neural Subdivision [Liu et al., 2020])

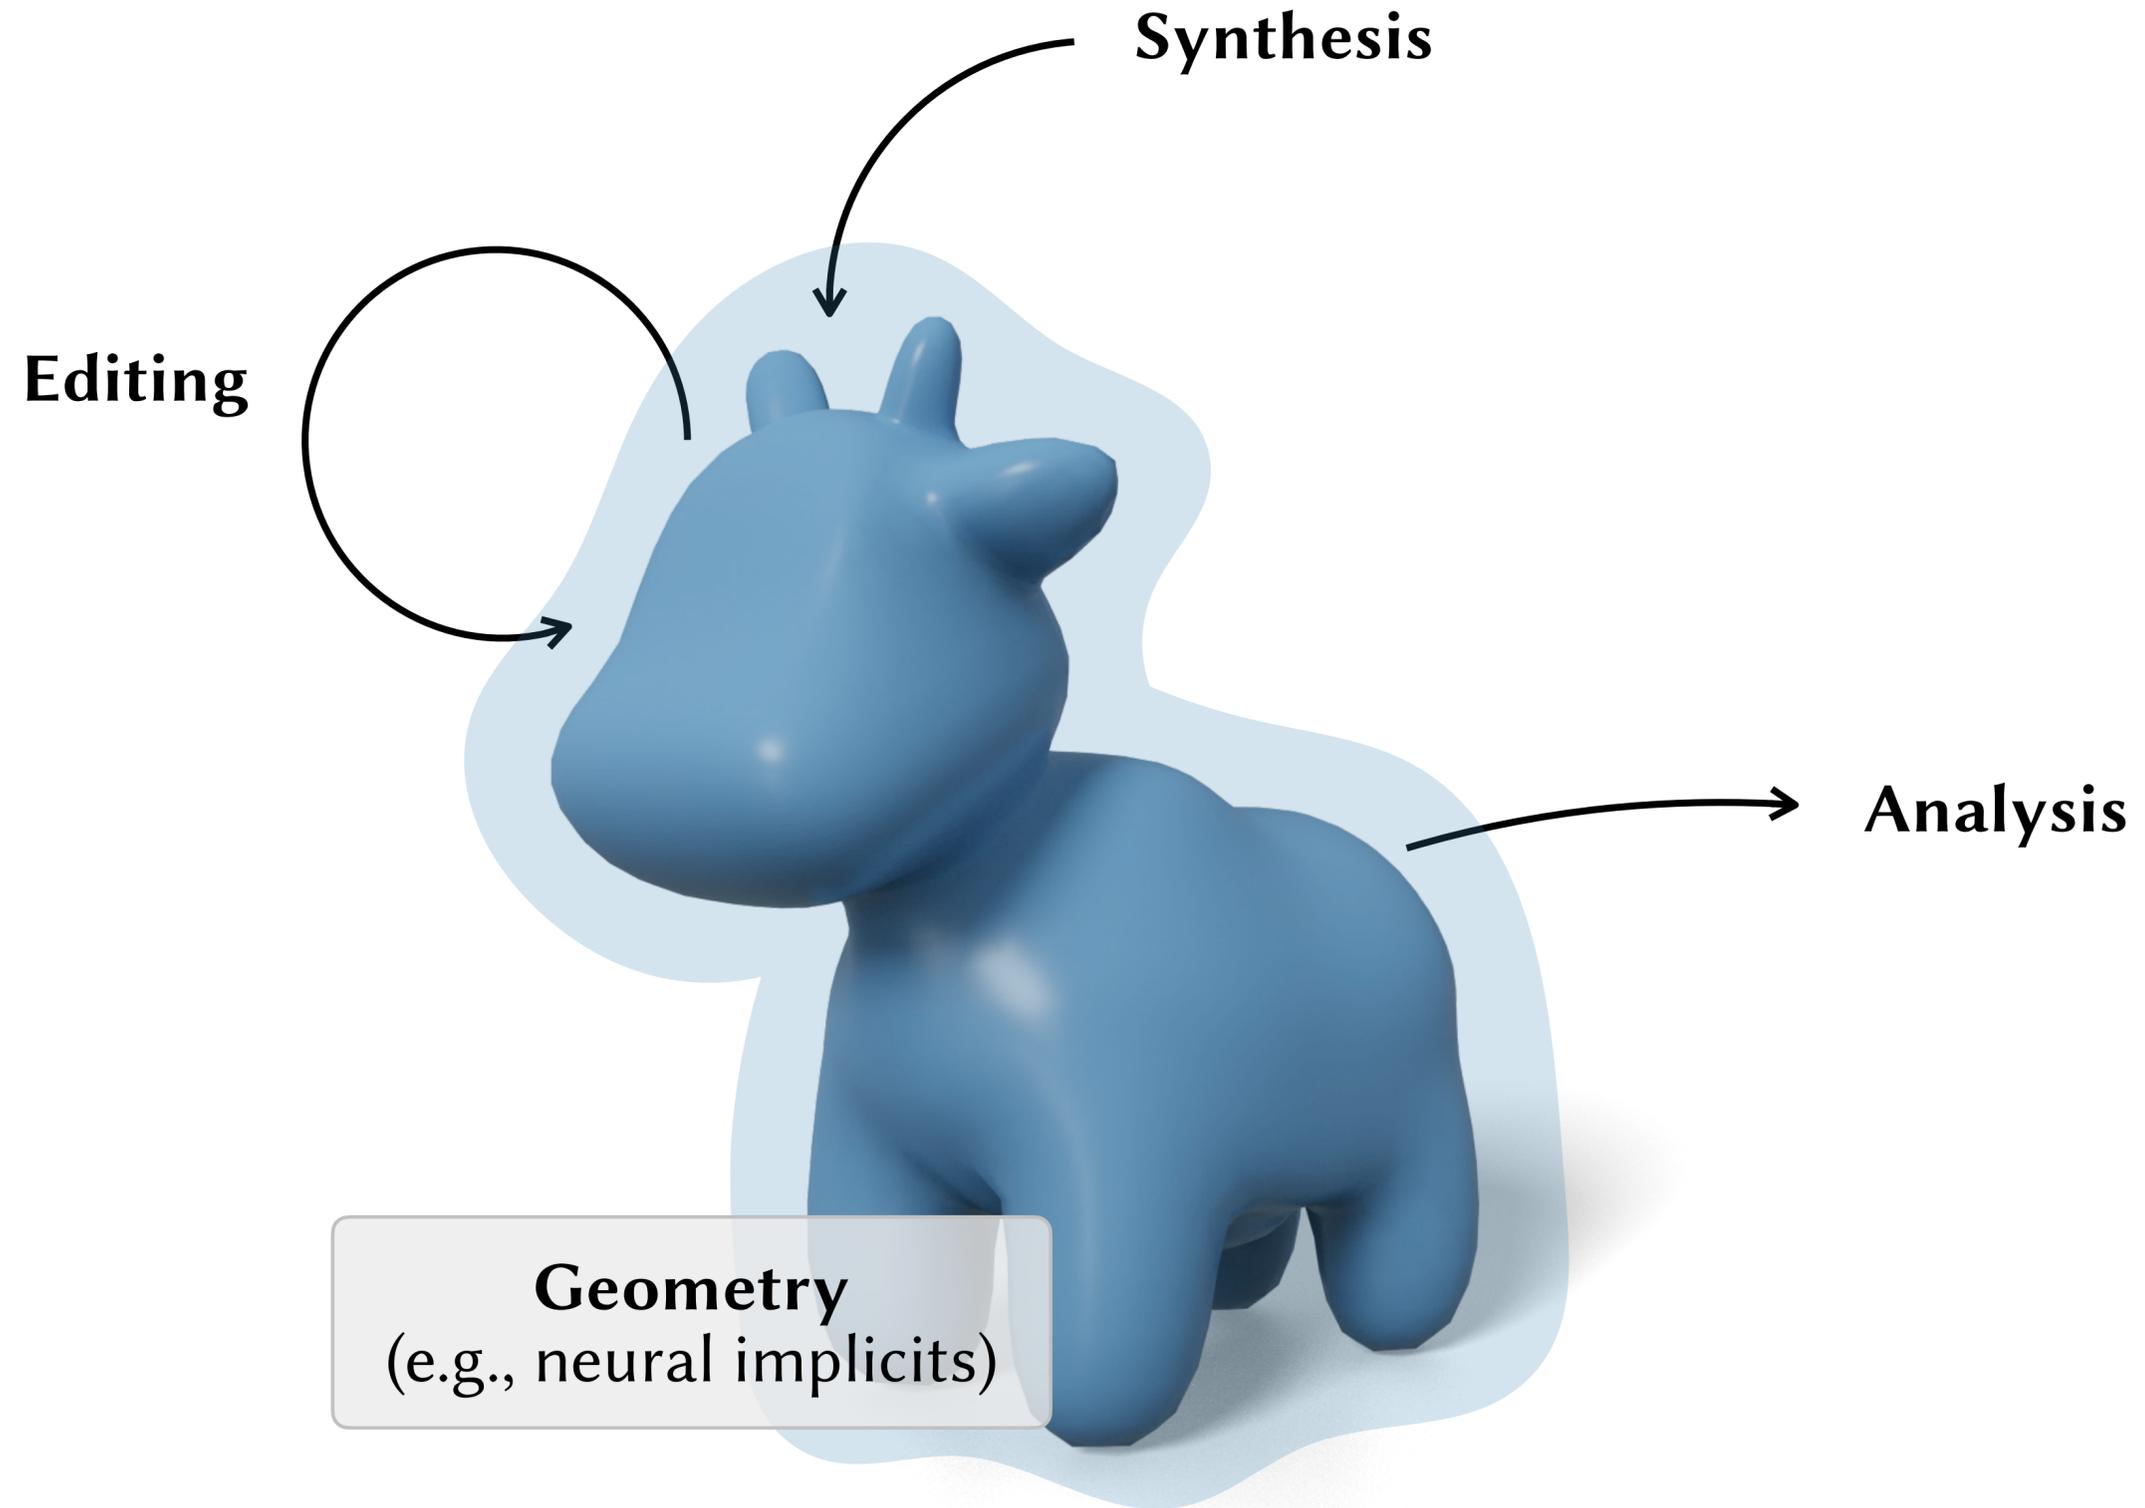


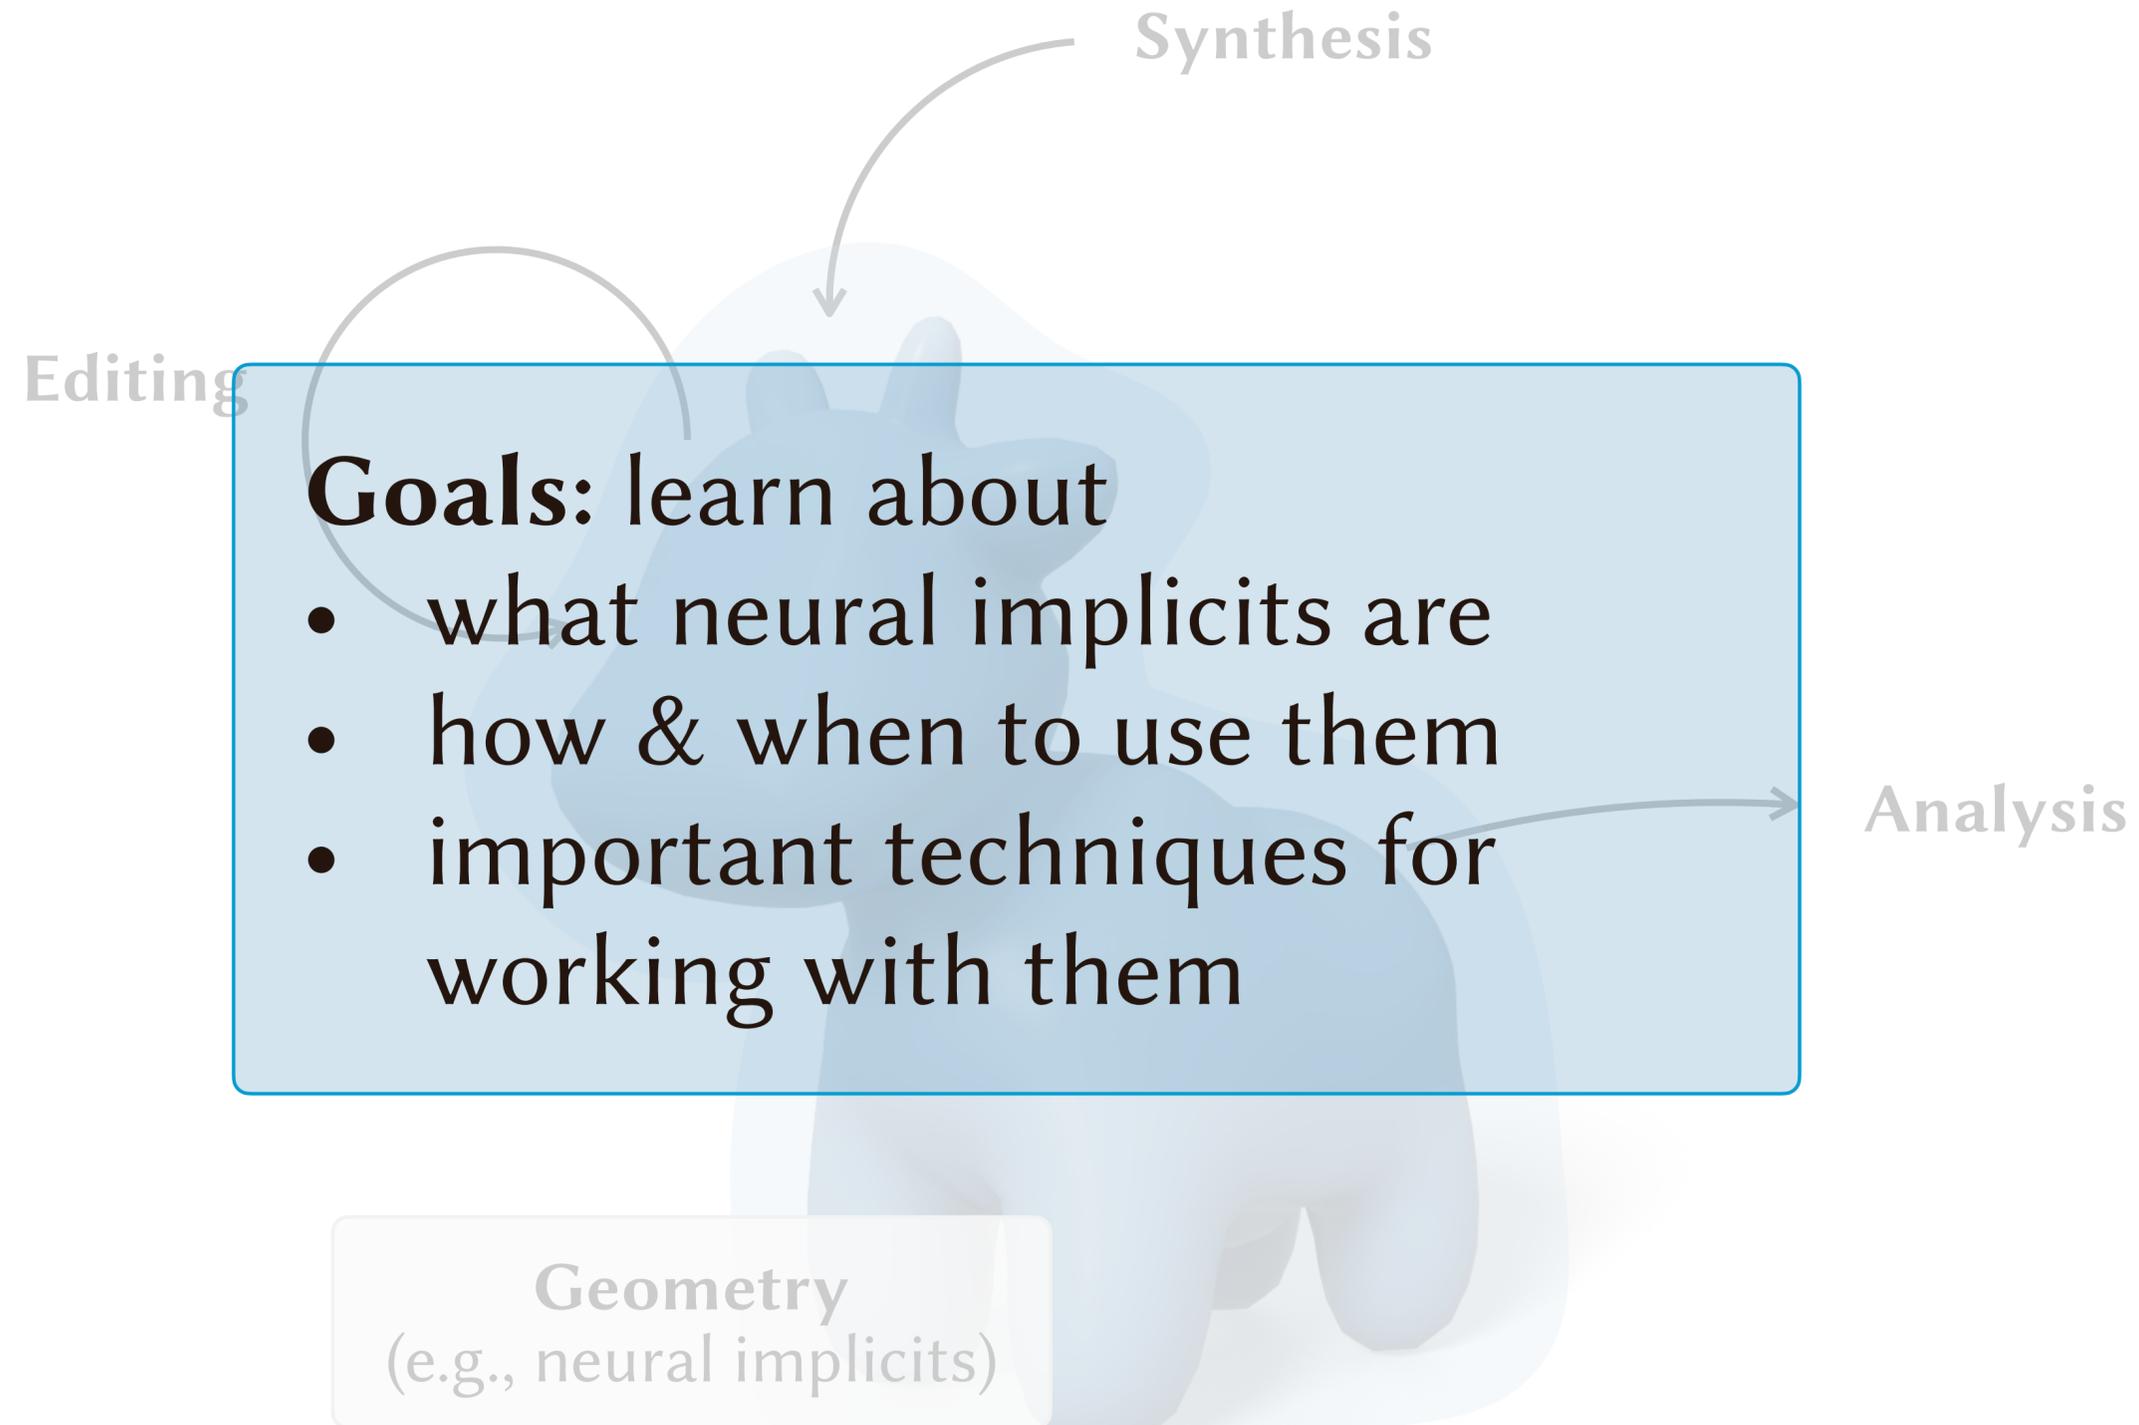
**Synthesis**  
(e.g., ShapeCrafter [Fu et al., 2023])

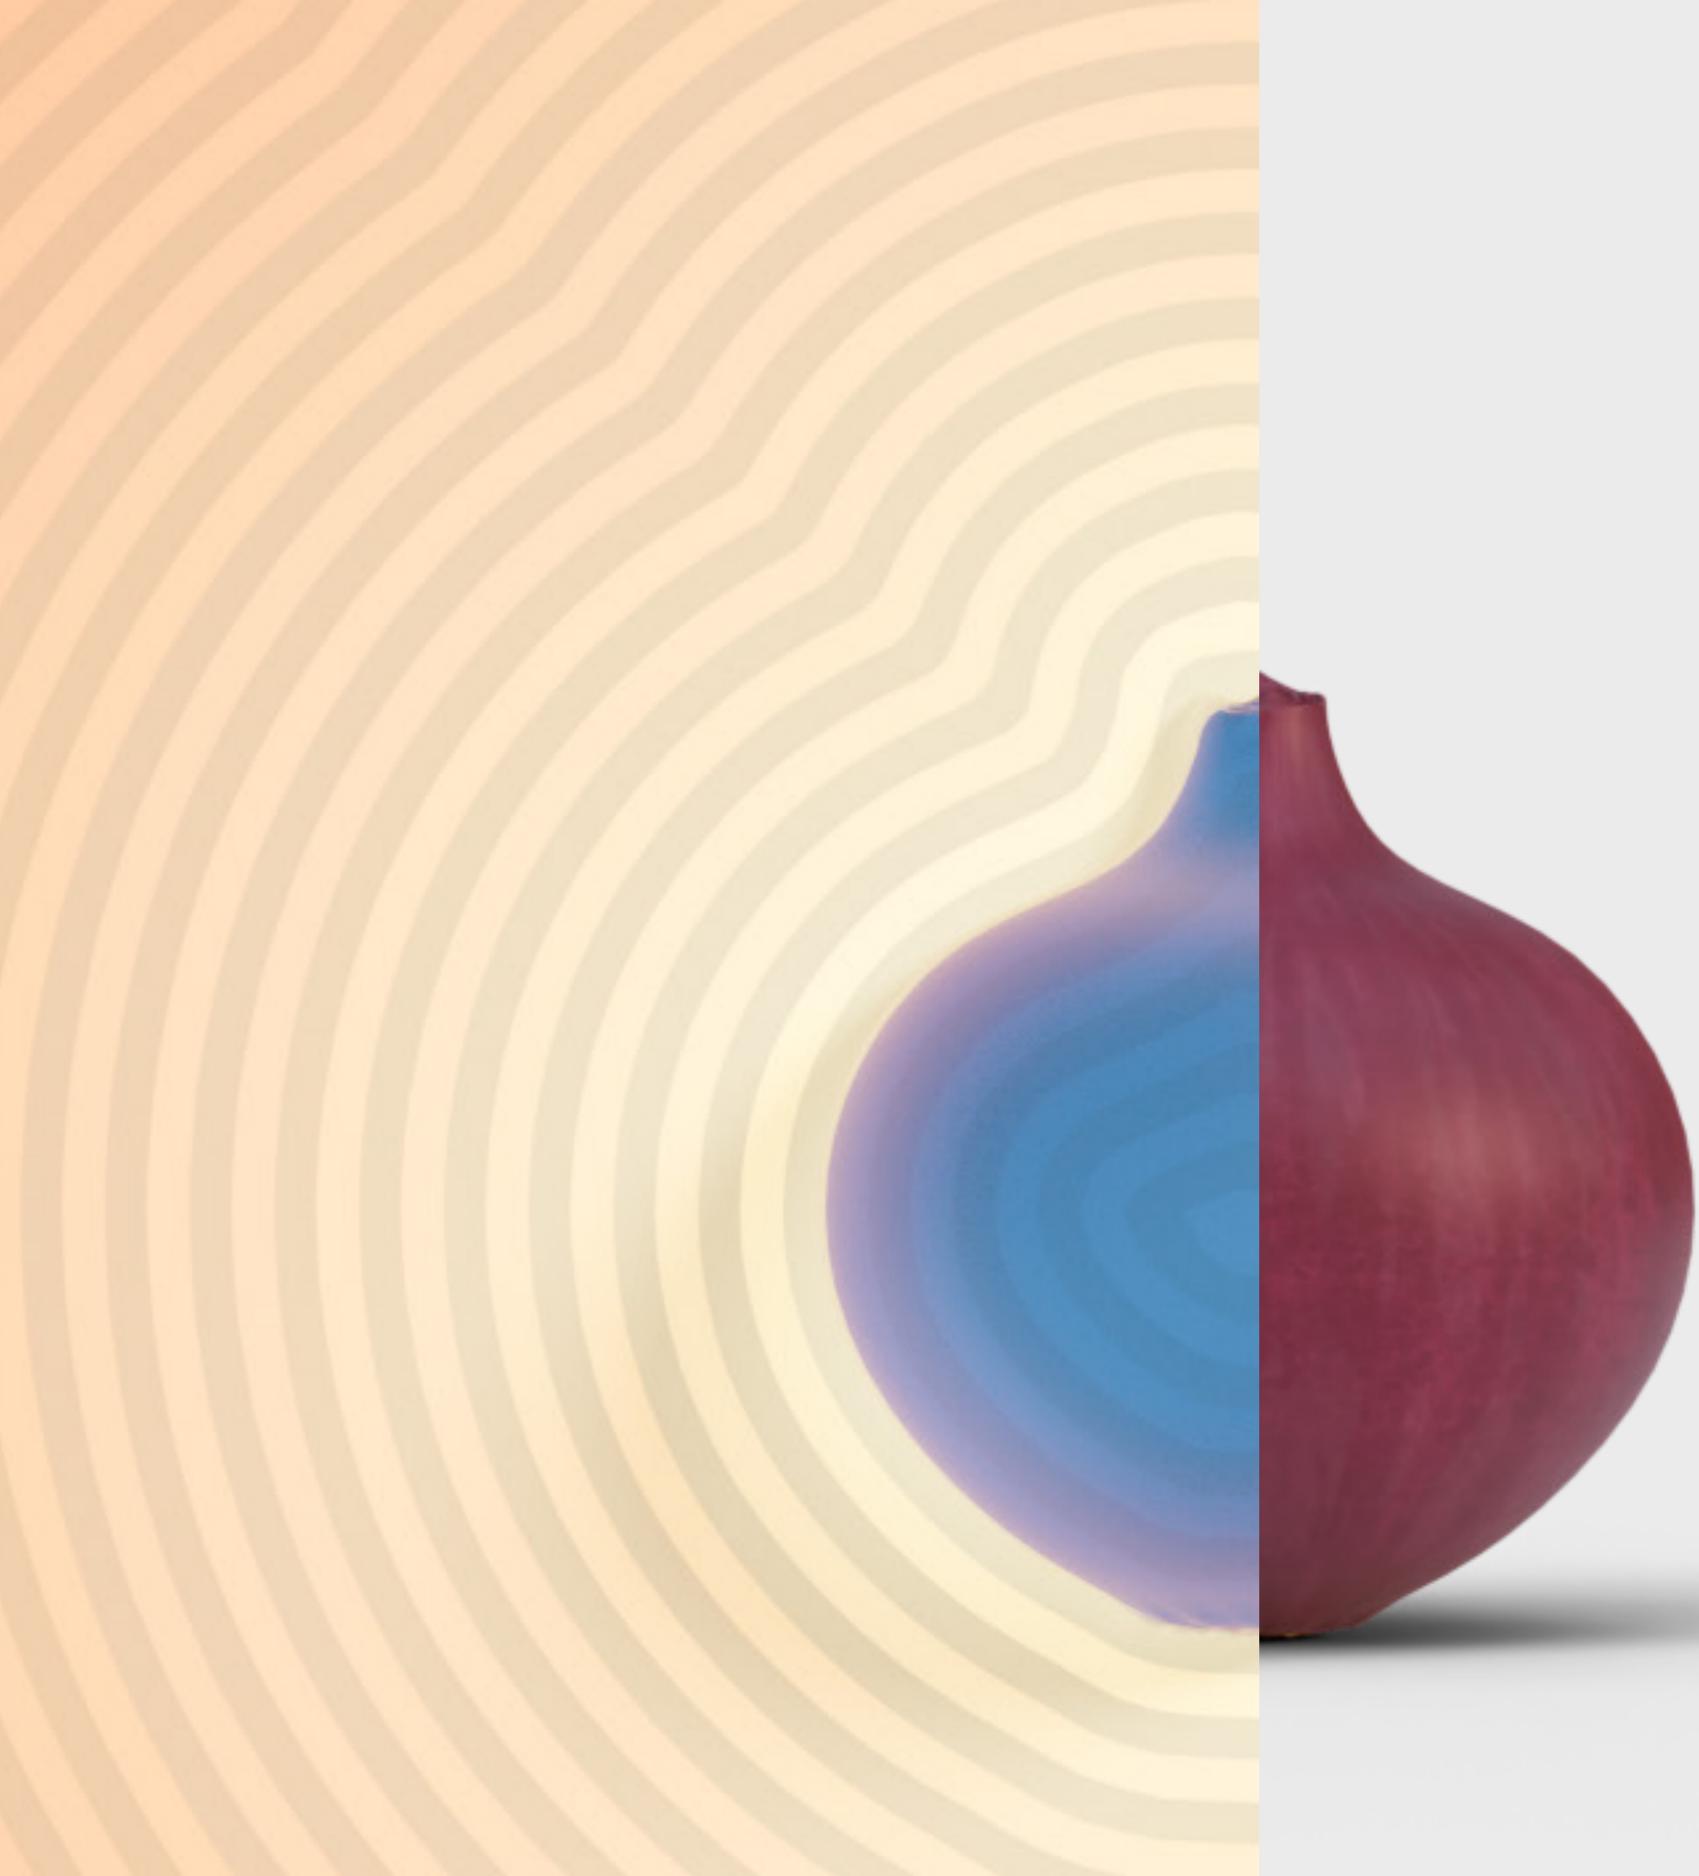


**Analysis**  
(e.g., DiffusionNet [Sharp et al., 2022])









## OVERVIEW

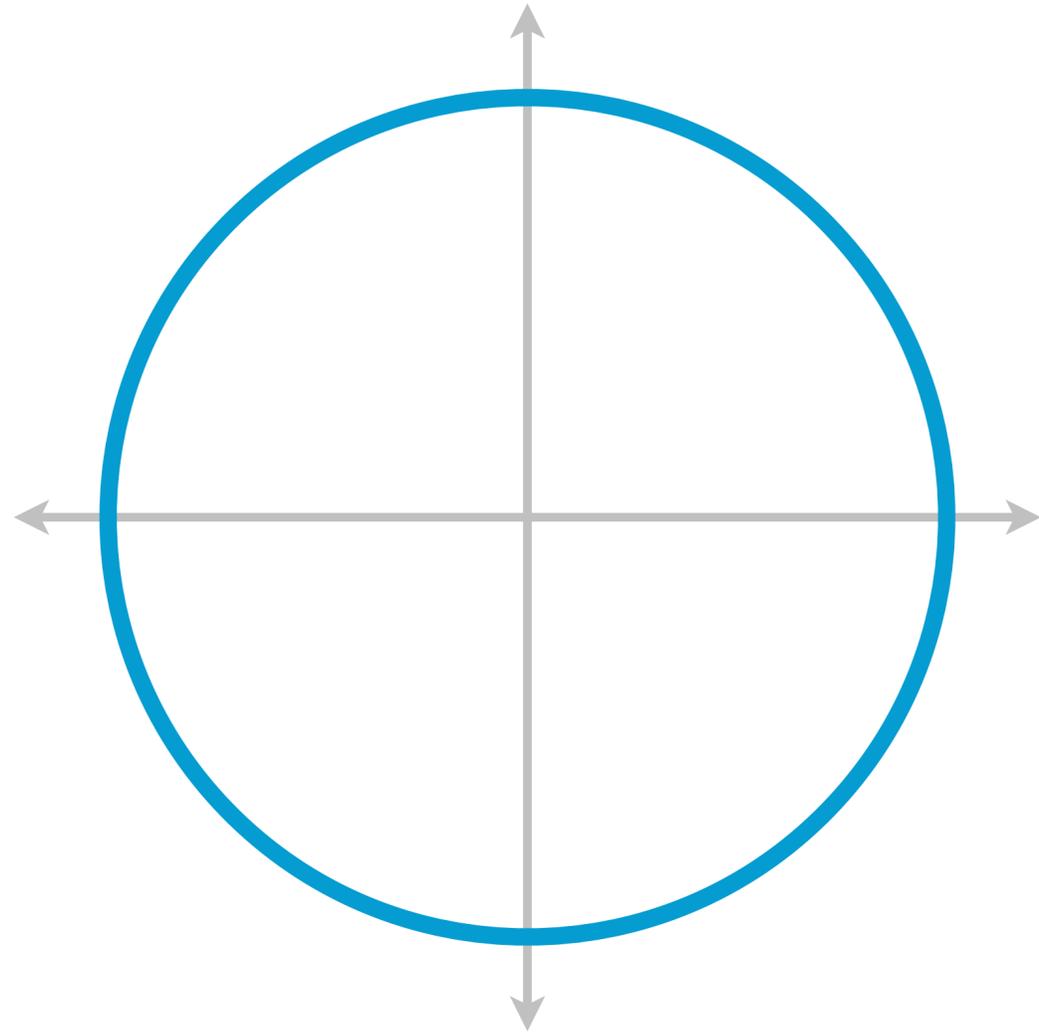
### IMPLICIT SURFACES

1. Background & History
2. Digital Representations

### NEURAL IMPLICIT

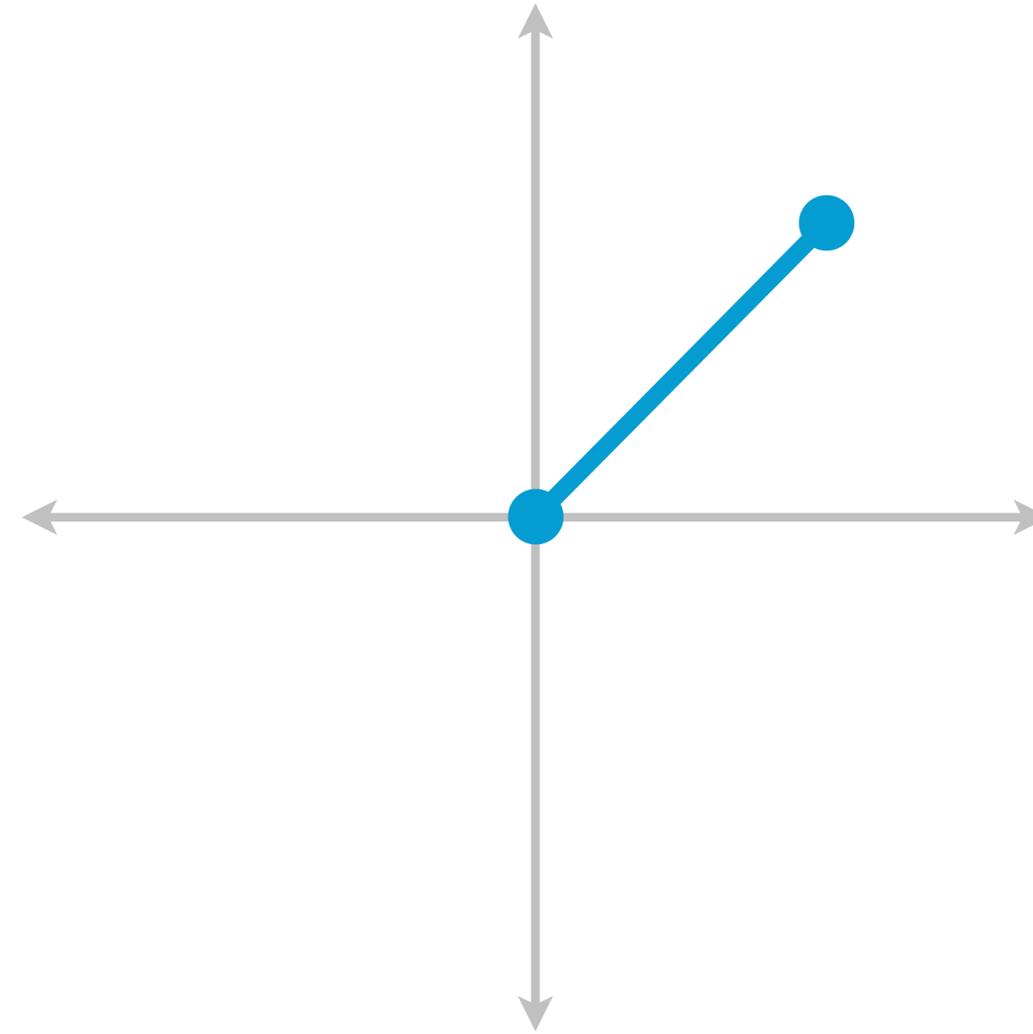
3. The Basics
4. Why & Why Not Use Them?
5. Working out the Details
6. Geometric Operations

### CONCLUSION



IMPLICIT

$$x^2 + y^2 = r^2$$



EXPLICIT

$$f(t) = (1,1)t \quad t \in [0,1]$$

# Implicit functions: definition

IMPLICIT

$$f(x, y) : \mathbb{R}^2 \rightarrow \mathbb{R}$$

$$f(x, y) = x^2 + y^2 - r^2$$



# Implicit functions: definition

## $c$ -Level Sets

- surface defined by set of points where  $f$  is equal to  $c$

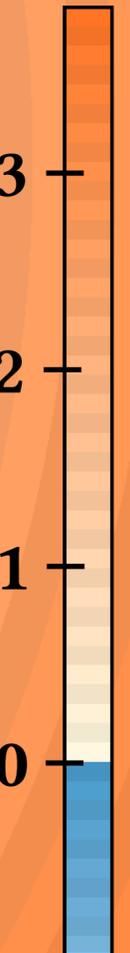
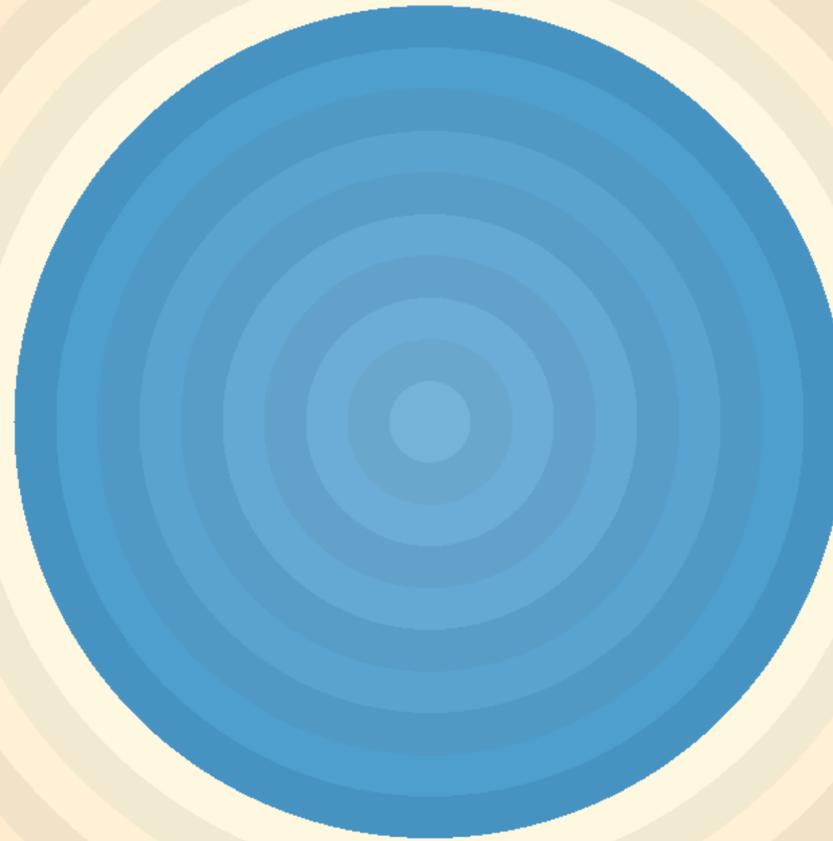
$$f(x, y) = c$$

$$c = \overset{-1}{\bullet} \text{---} \overset{2}{\bullet}$$



## ***c*-Level Sets**

- surface defined by set of points where  $f$  is equal to  $c$
- typically, *zero-level set* is treated as the surface

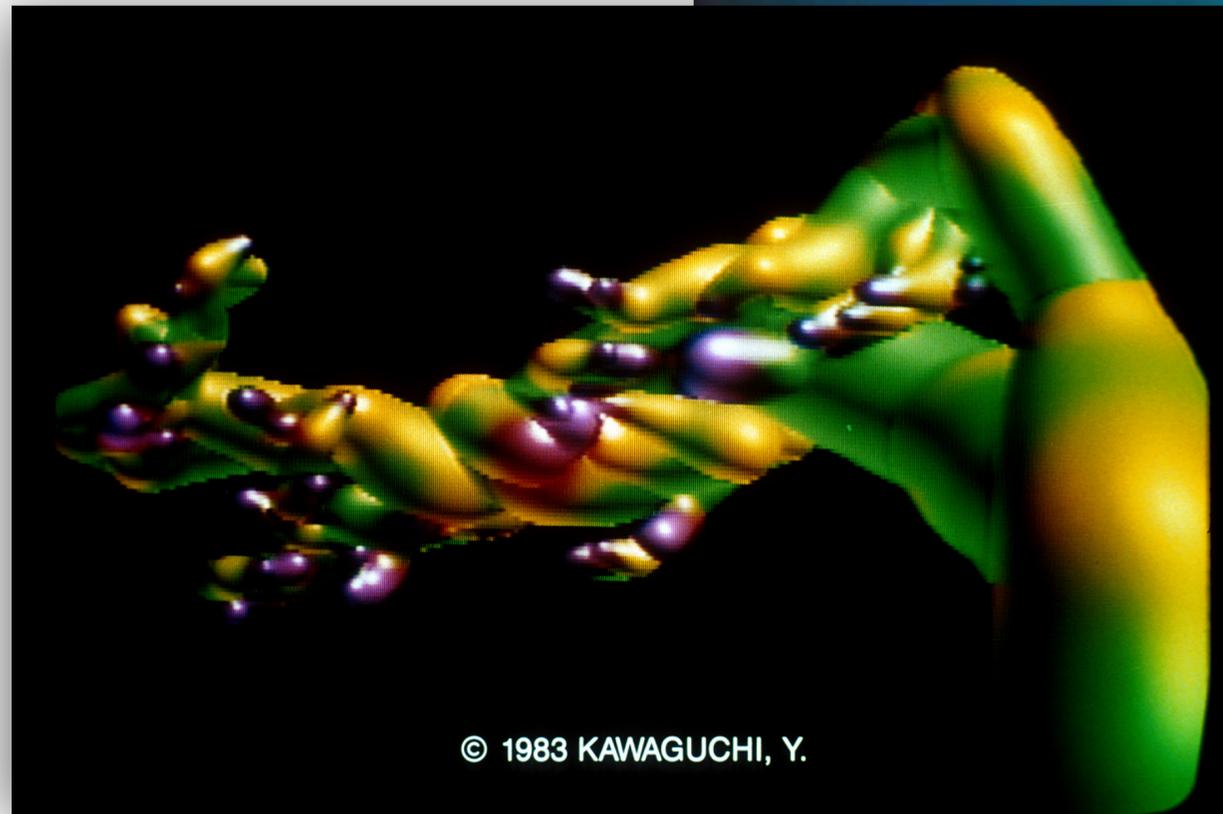




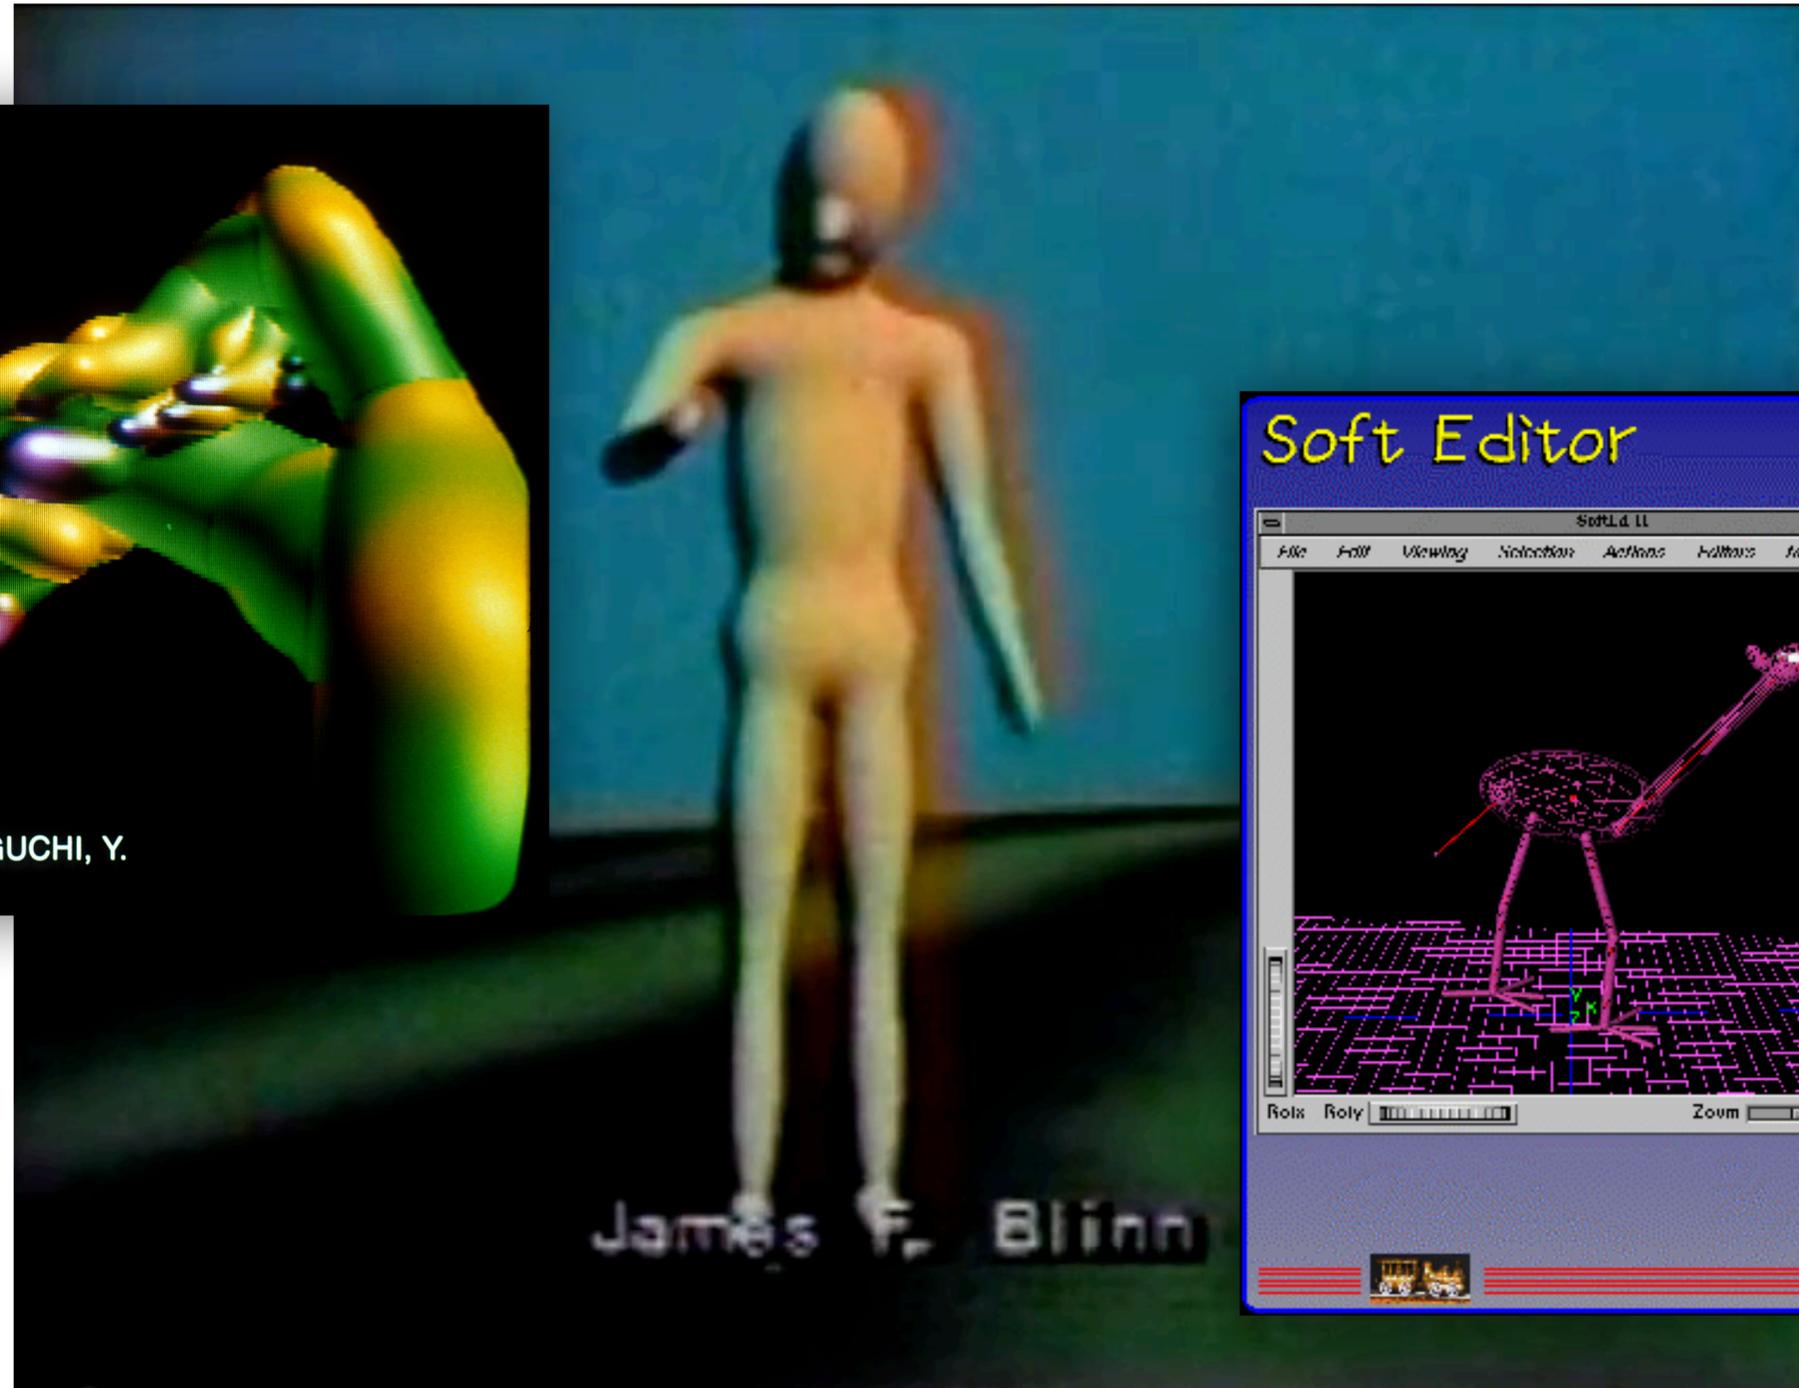


Jim Blinn's "Blobby molecules" (1982)

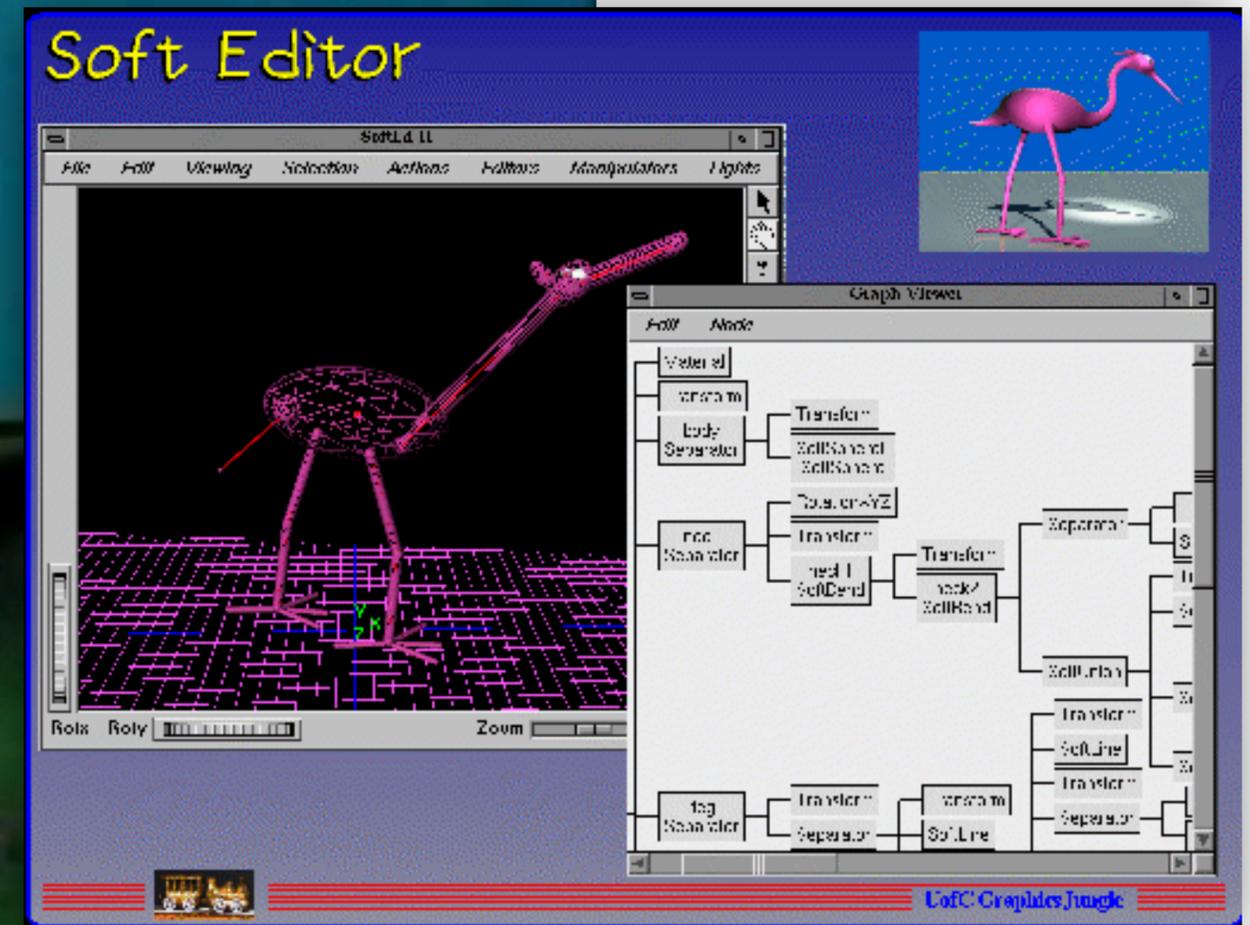
# Implicits in early graphics history



“metaballs”



“soft objects”



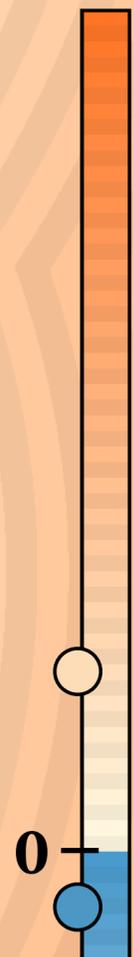
📖 For more background on implicits, see Chapter 21 of *Fundamentals of Computer Graphics*

## **IMPLICIT SURFACES: BIG IDEAS**

1. Constructive Solid Geometry
2. Taxonomy of implicits
3. Rendering
4. Surface Extraction

“outside!”

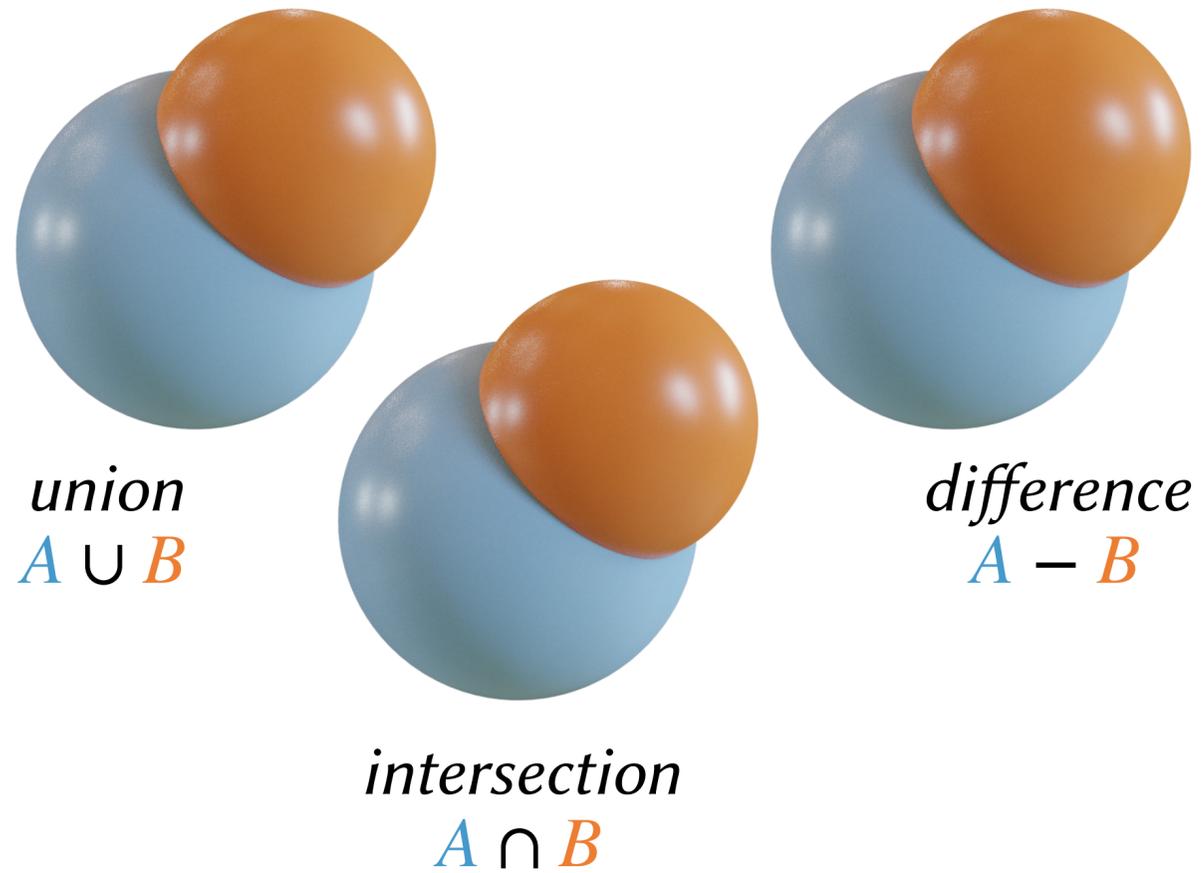
“inside!”



**IMPLICIT SURFACES: BIG IDEAS**

1. **CSG**
2. Taxonomy of implicits
3. Rendering
4. Surface Extraction

## Constructive solid geometry



### IMPLICIT SURFACES: BIG IDEAS

1. **CSG**
2. Taxonomy of implicits
3. Rendering
4. Surface Extraction

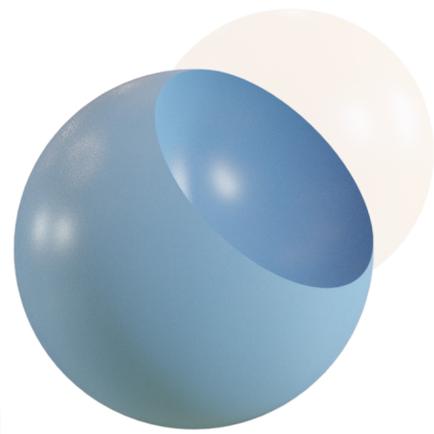
## Constructive solid geometry



*union*  
 $A \cup B$



*intersection*  
 $A \cap B$



*difference*  
 $A - B$

### IMPLICIT SURFACES: BIG IDEAS

1. **CSG**
2. Taxonomy of implicits
3. Rendering
4. Surface Extraction

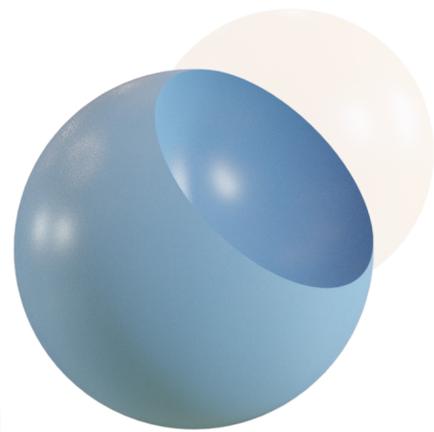
### Constructive solid geometry



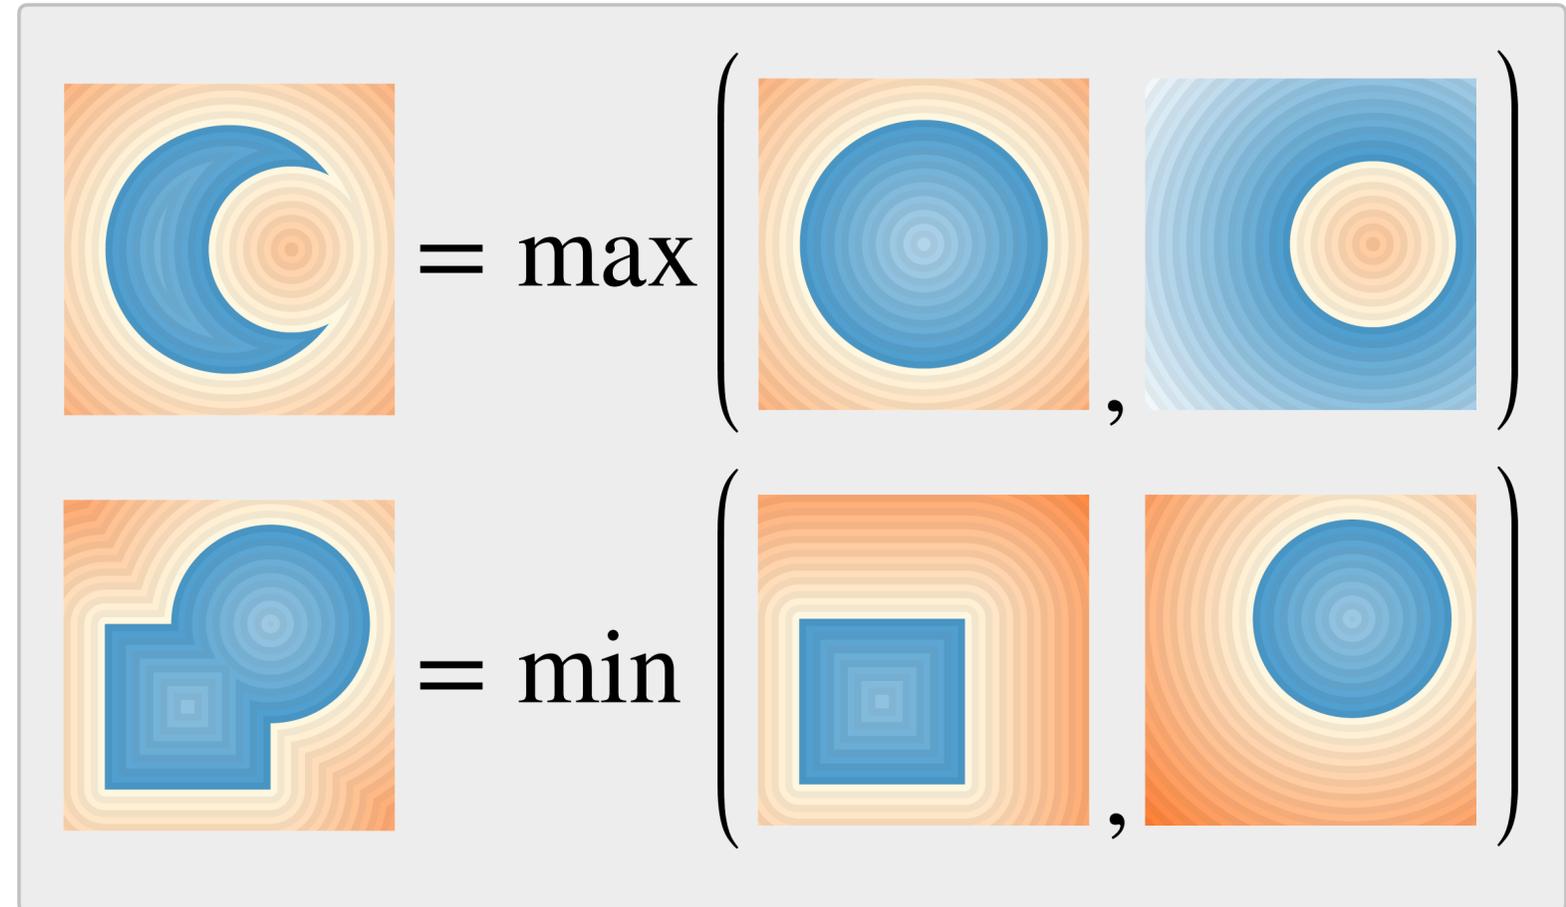
union  
 $A \cup B$



intersection  
 $A \cap B$



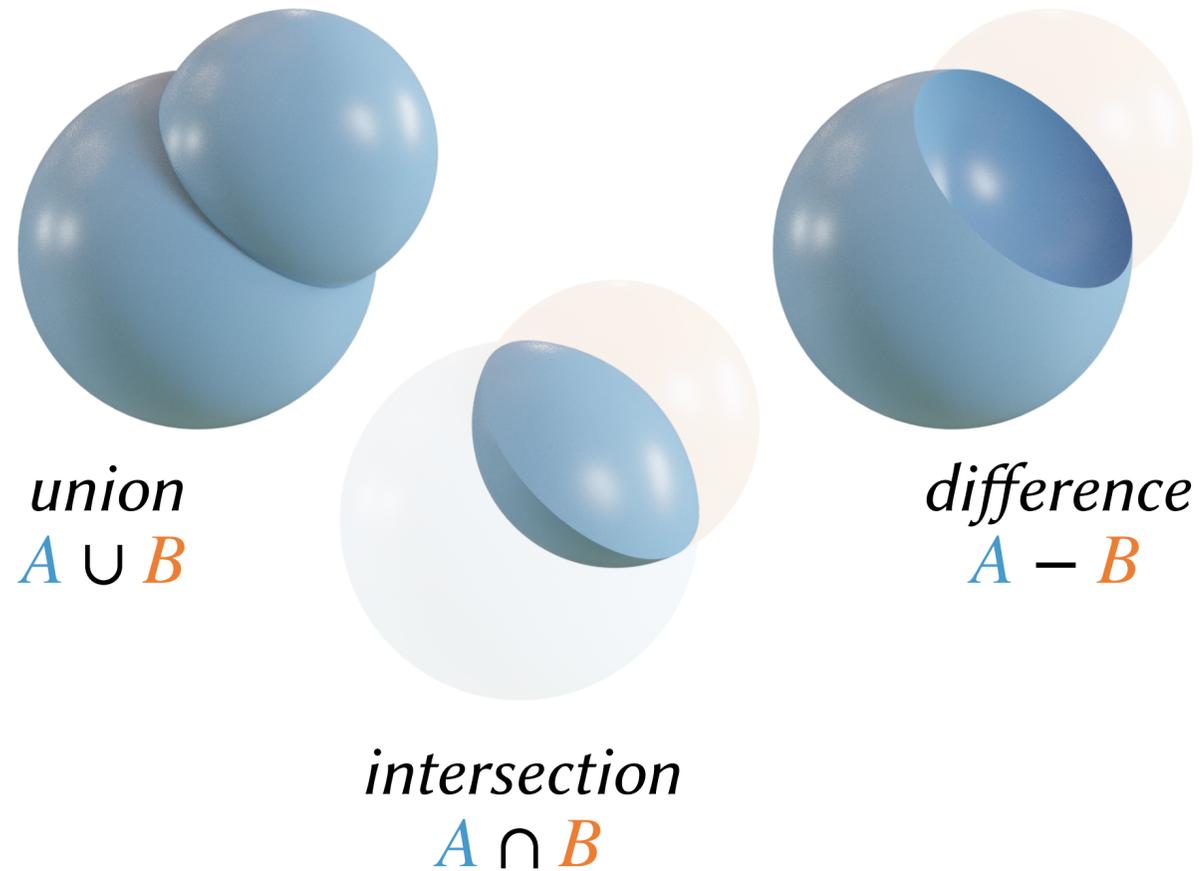
difference  
 $A - B$



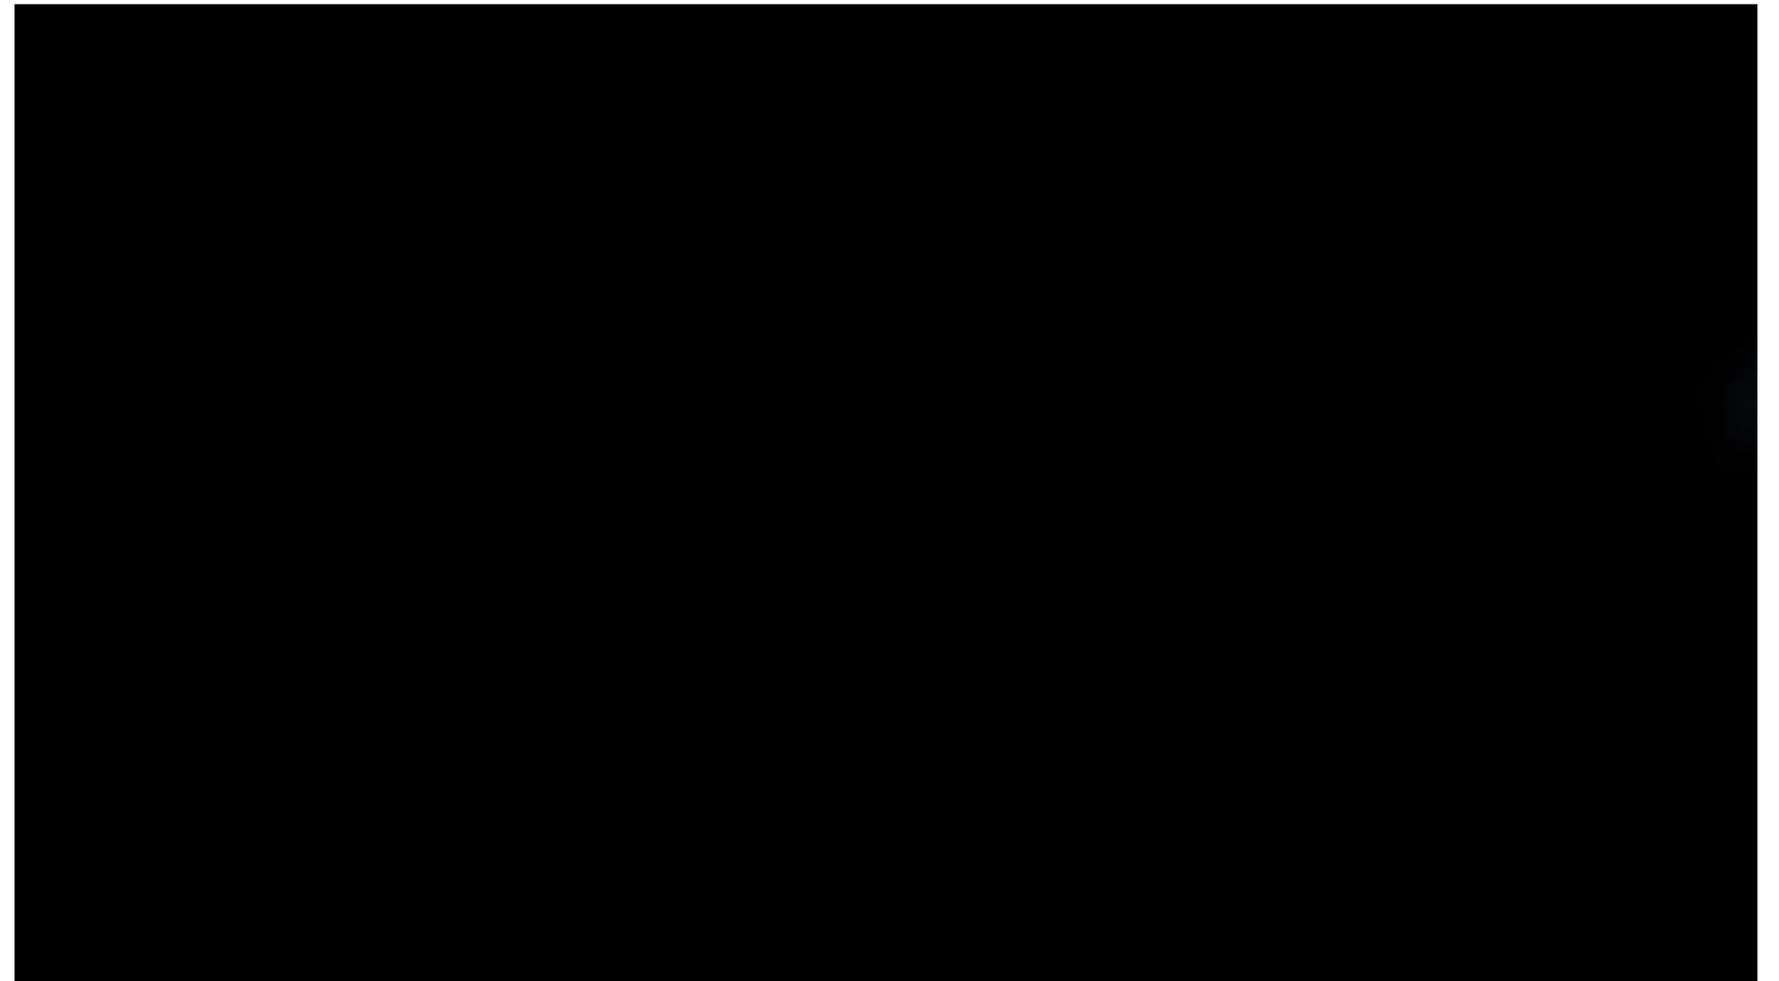
### IMPLICIT SURFACES: BIG IDEAS

1. **CSG**
2. Taxonomy of implicits
3. Rendering
4. Surface Extraction

## Constructive solid geometry



## Modeling with implicits



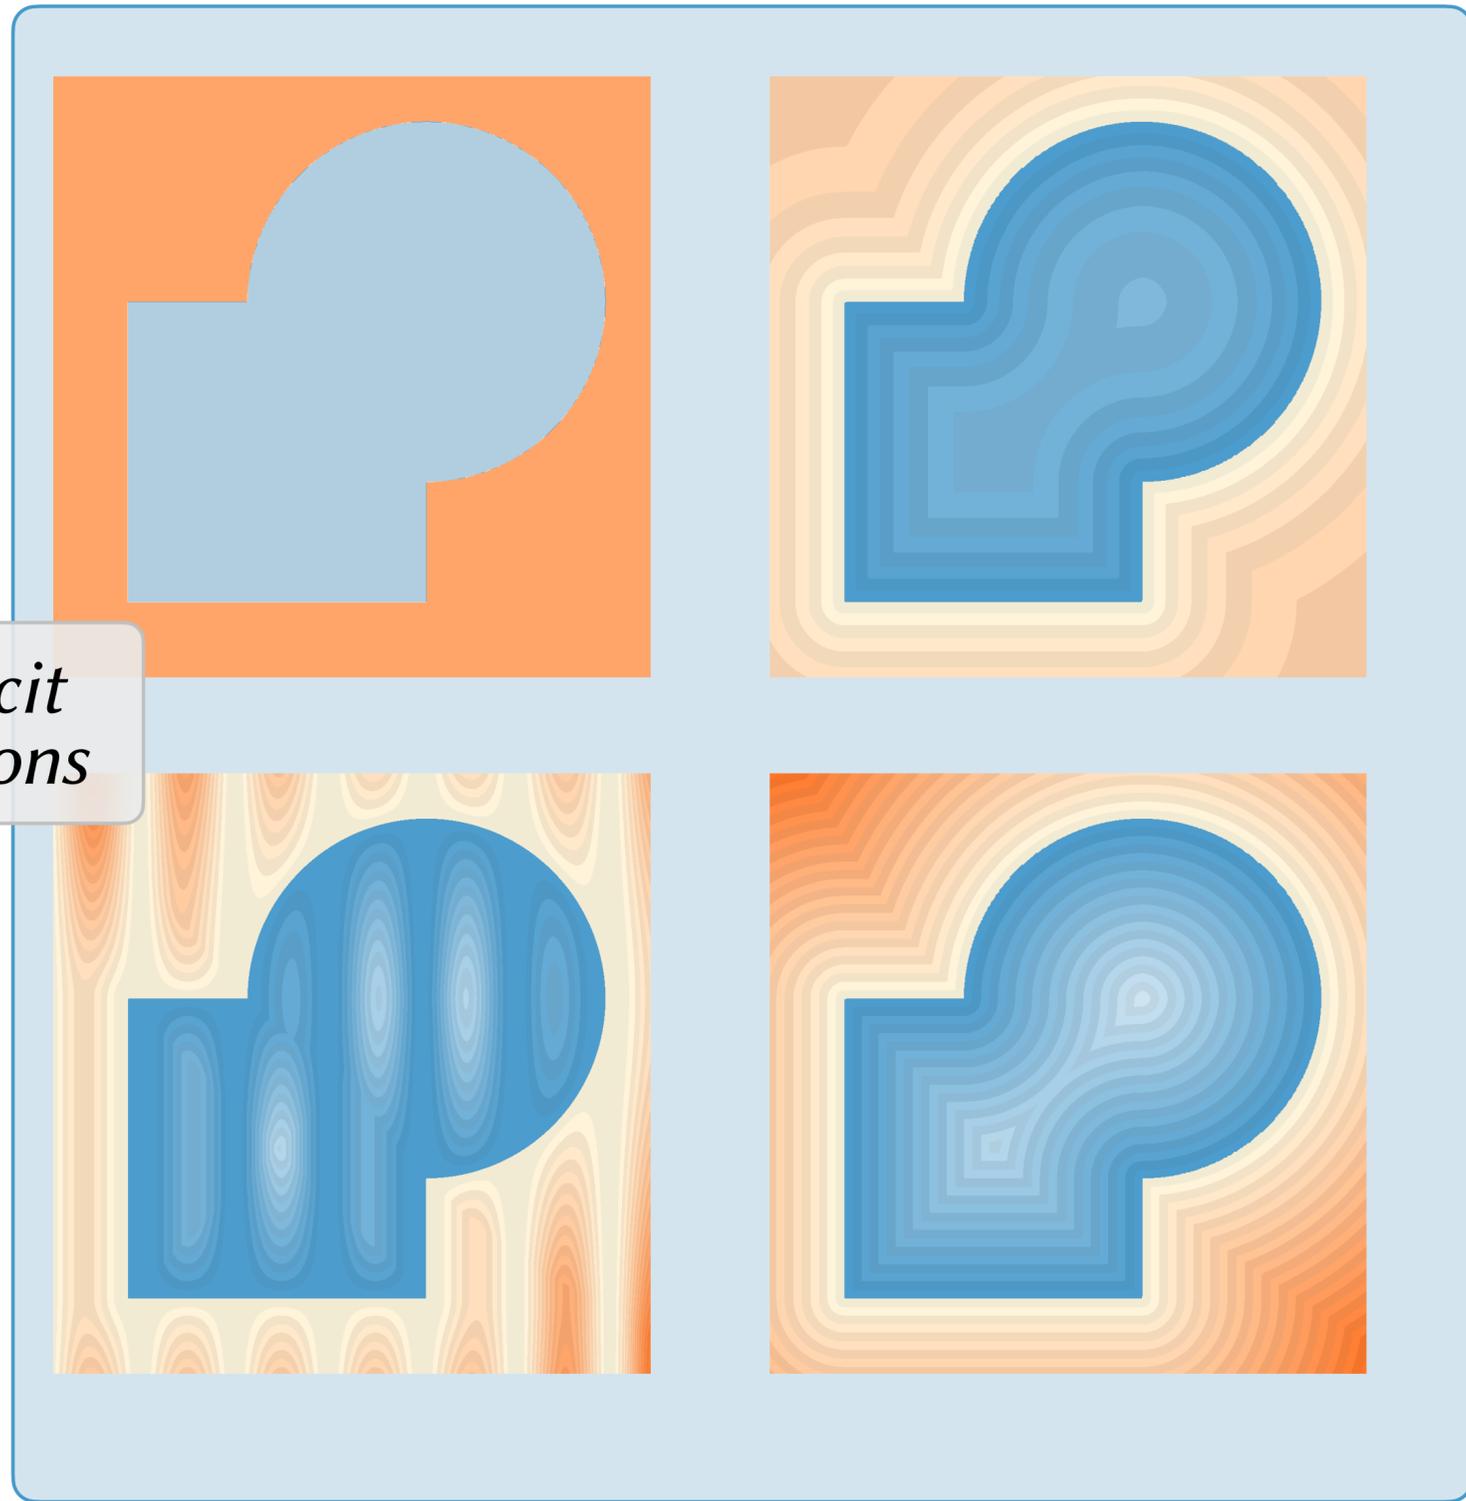
*Complex shapes from simple primitives!*

### IMPLICIT SURFACES: BIG IDEAS

1. **CSG**
2. Taxonomy of implicits
3. Rendering
4. Surface Extraction

📖 For more amazing implicit function art, see [Inigo Quilez's blog](#)

*implicit functions*



**IMPLICIT SURFACES: BIG IDEAS**

1. CSG
2. **Taxonomy of implicits**
3. Rendering
4. Surface Extraction

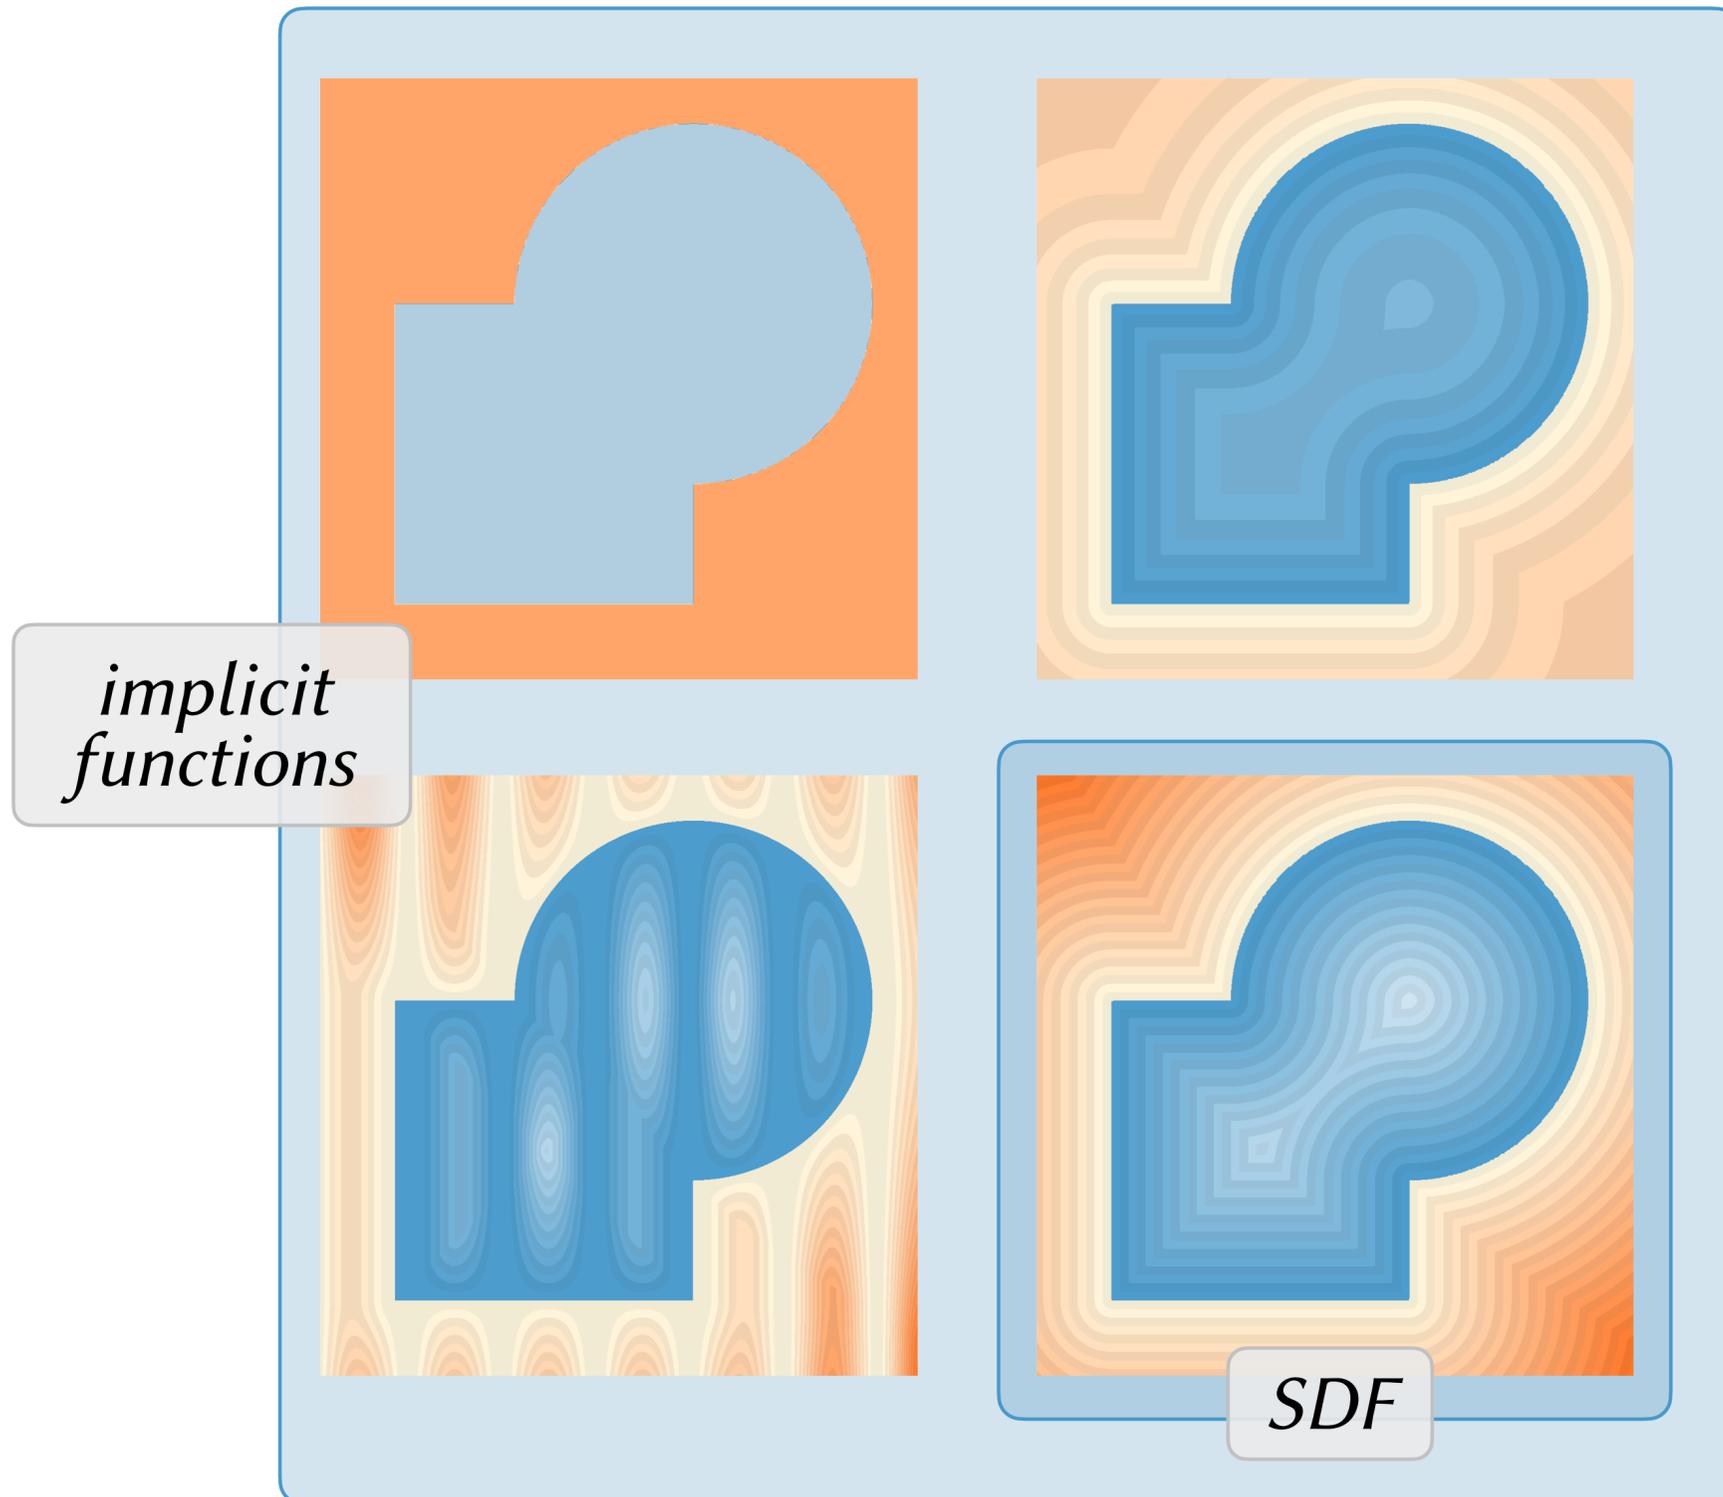
### Indicator Function

$$f(\mathbf{x}) = \begin{cases} -1 & \mathbf{x} \in \text{shape} \\ 1 & \mathbf{x} \notin \text{shape} \end{cases}$$



### IMPLICIT SURFACES: BIG IDEAS

1. CSG
2. **Taxonomy of implicits**
3. Rendering
4. Surface Extraction

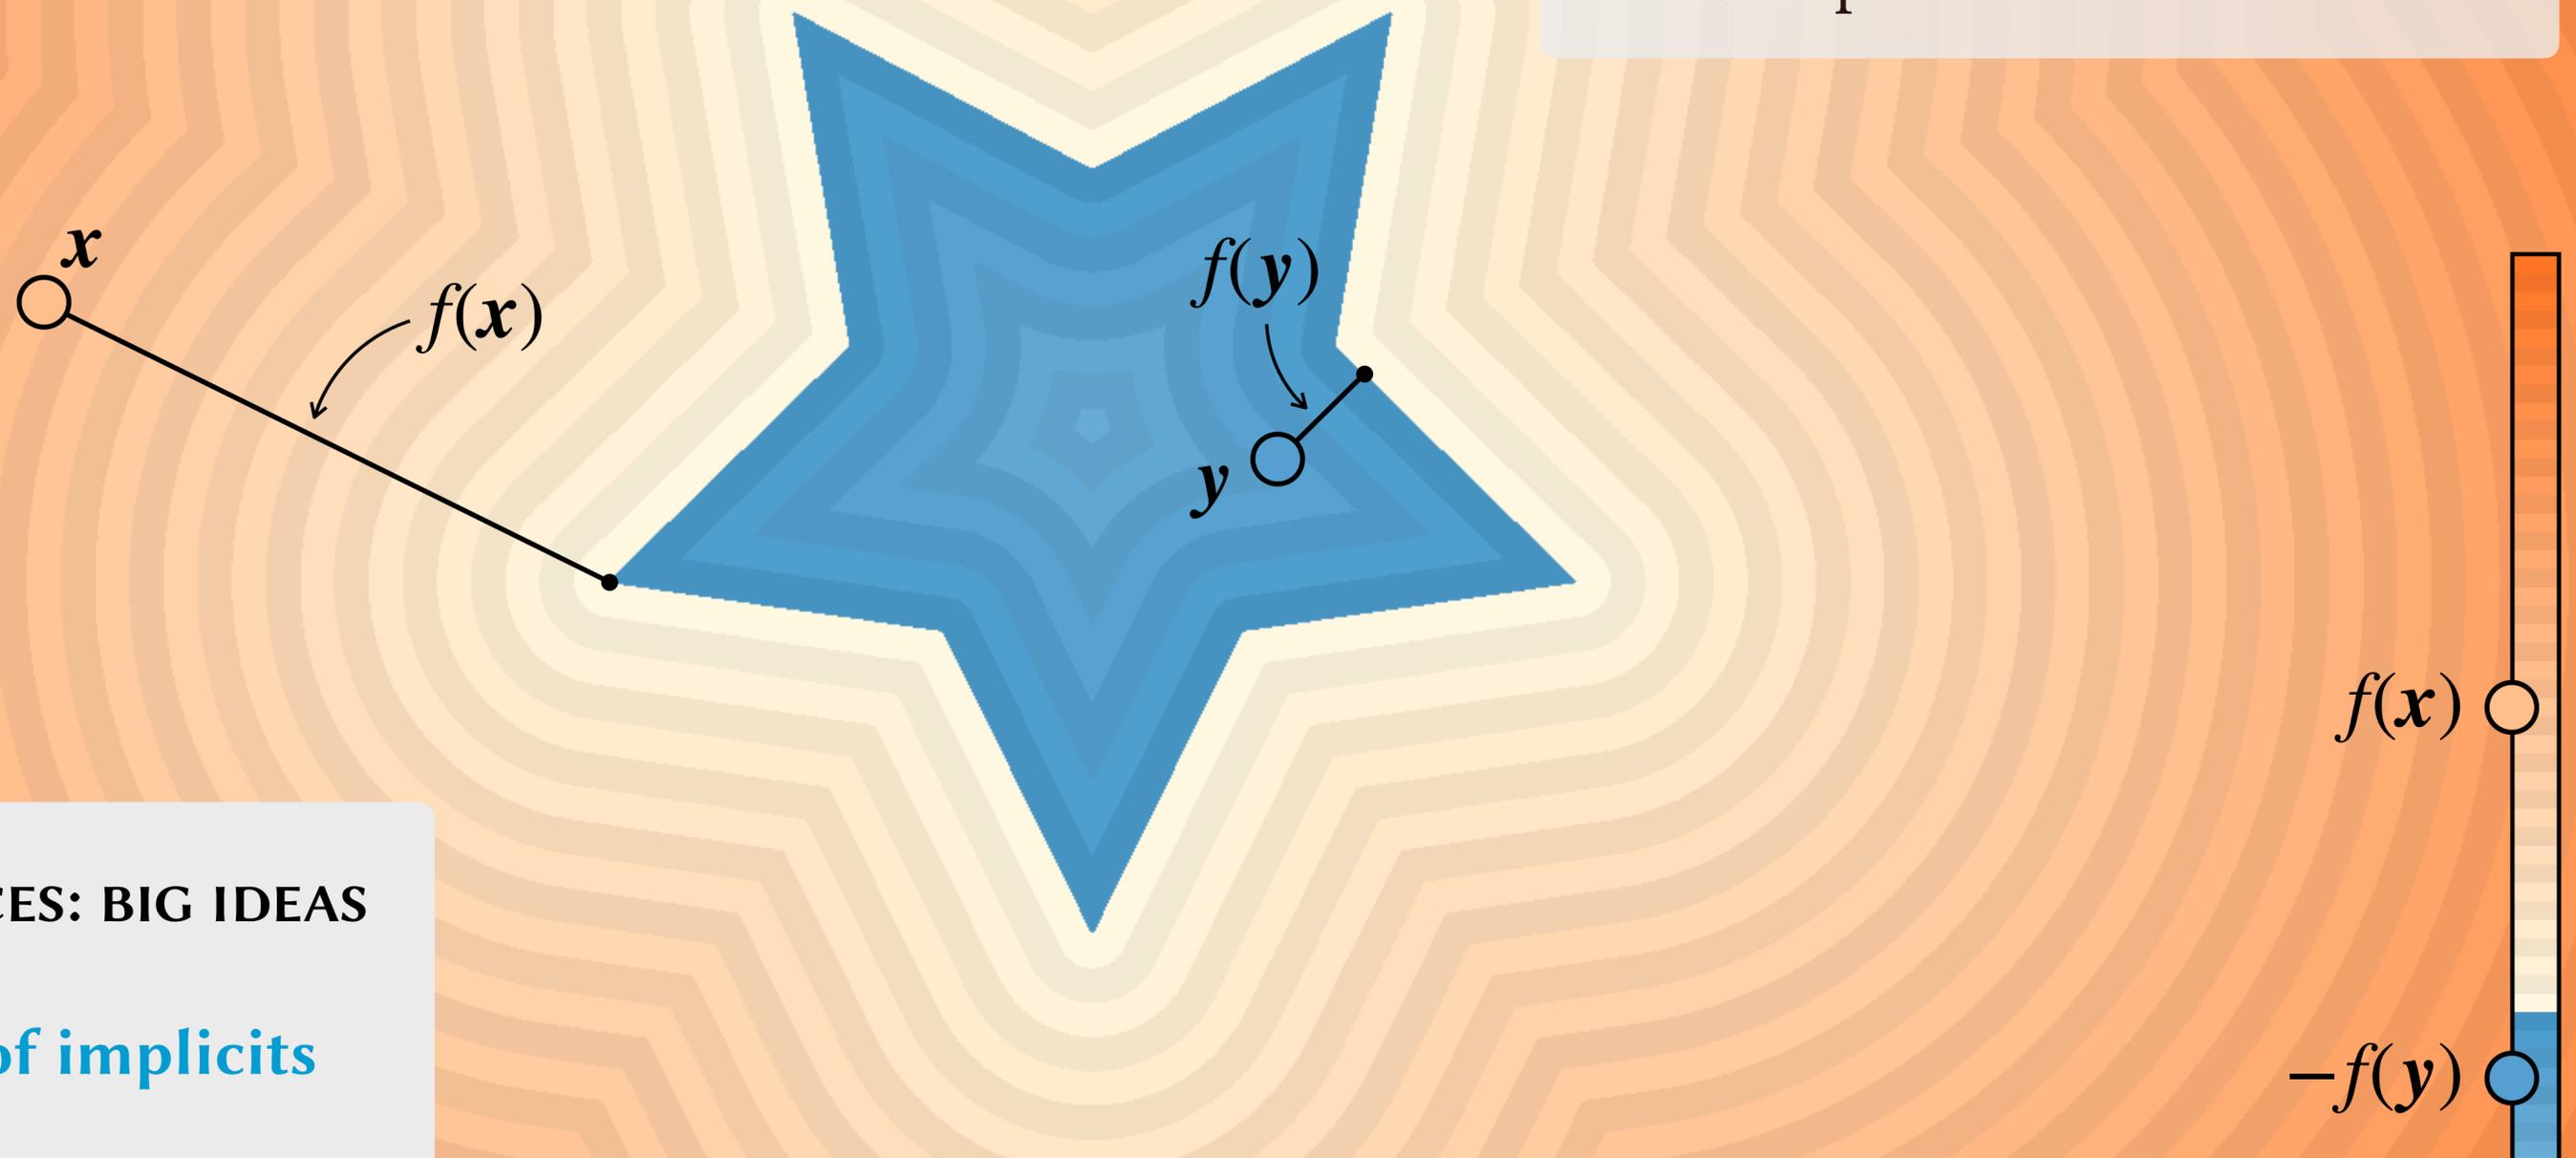


## IMPLICIT SURFACES: BIG IDEAS

1. CSG
2. **Taxonomy of implicits**
3. Rendering
4. Surface Extraction

## Signed Distance Function

- $|f(\mathbf{x})|$  encodes distance to closest point on surface

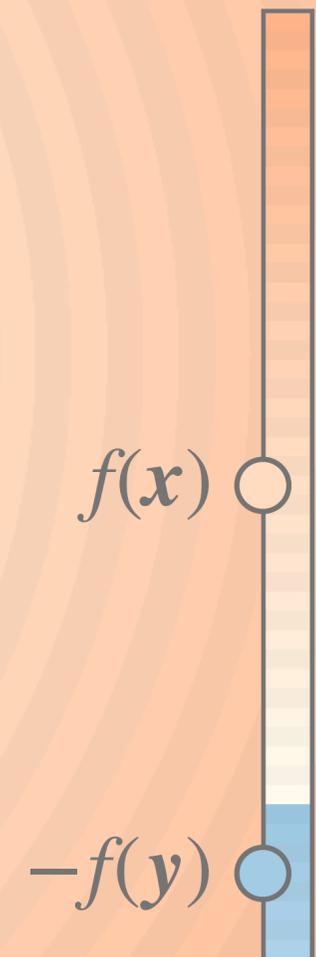
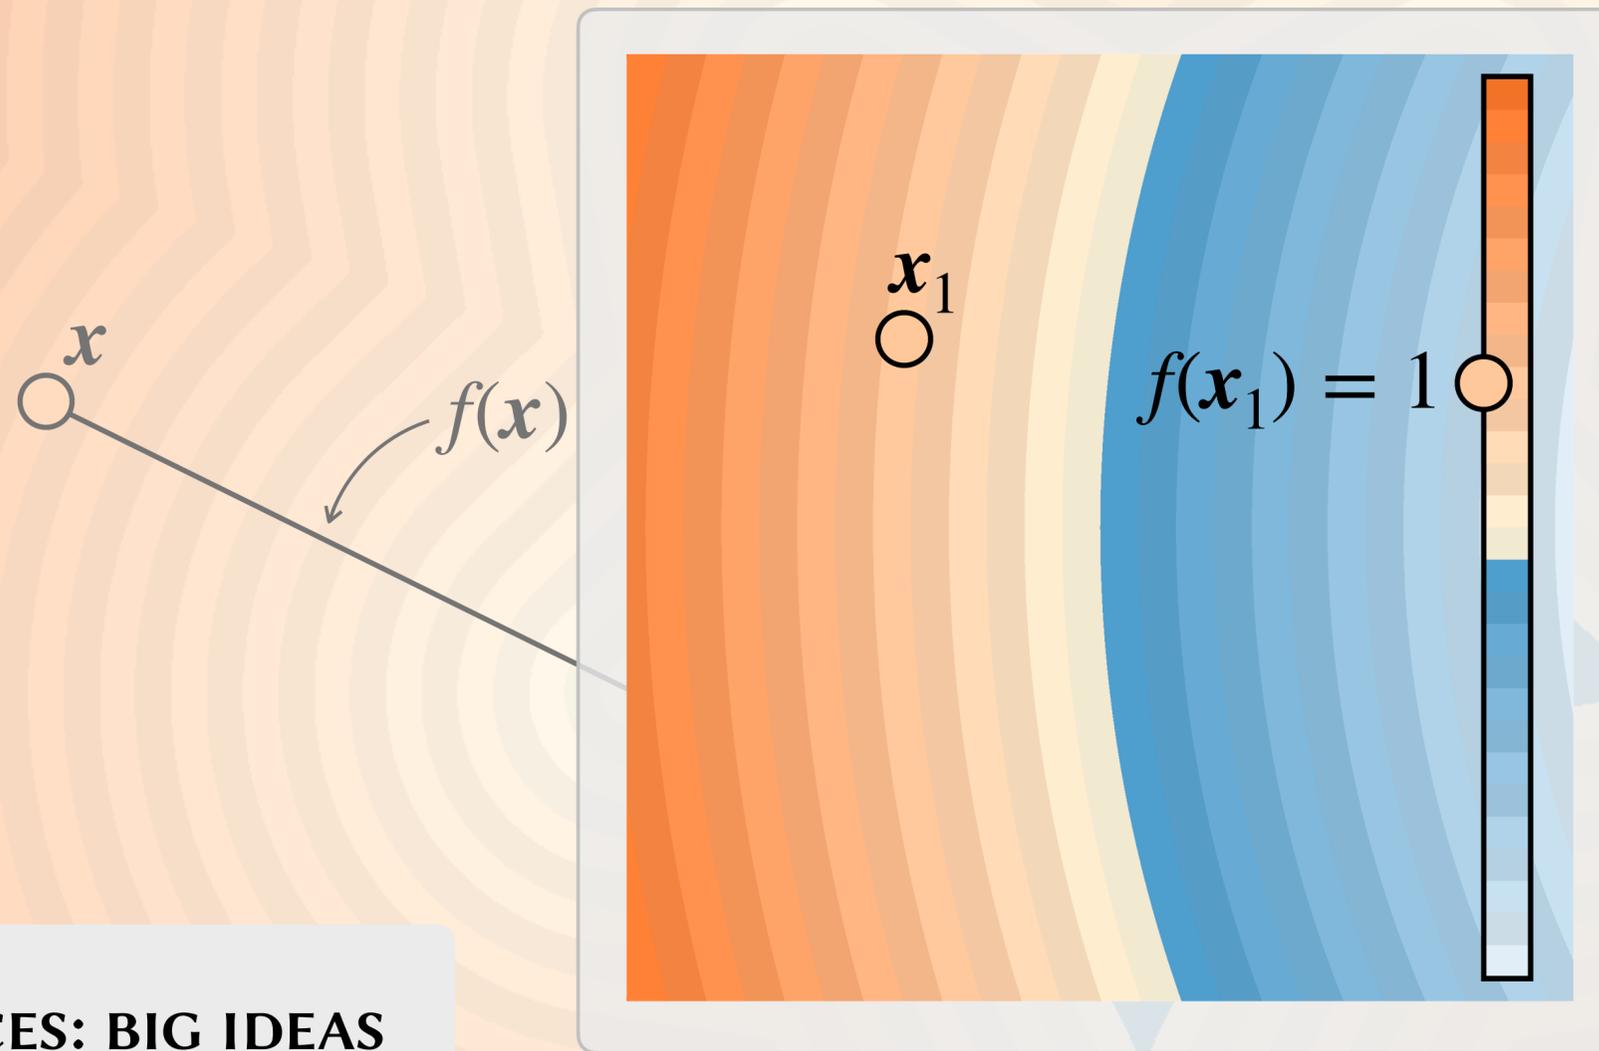


### IMPLICIT SURFACES: BIG IDEAS

1. CSG
2. Taxonomy of implicits
3. Rendering
4. Surface Extraction

## Signed Distance Function

- $|f(\mathbf{x})|$  encodes distance to closest point on surface

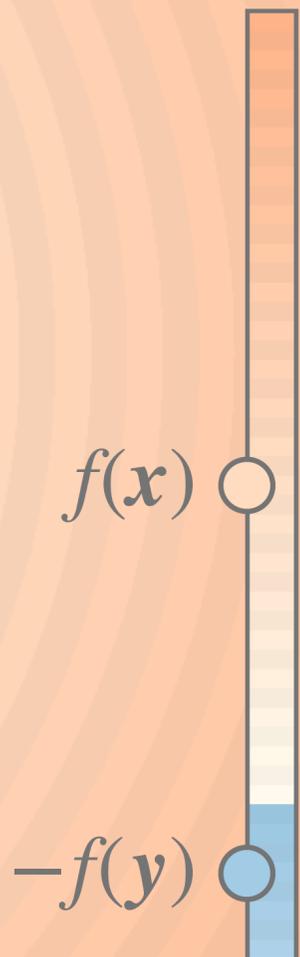
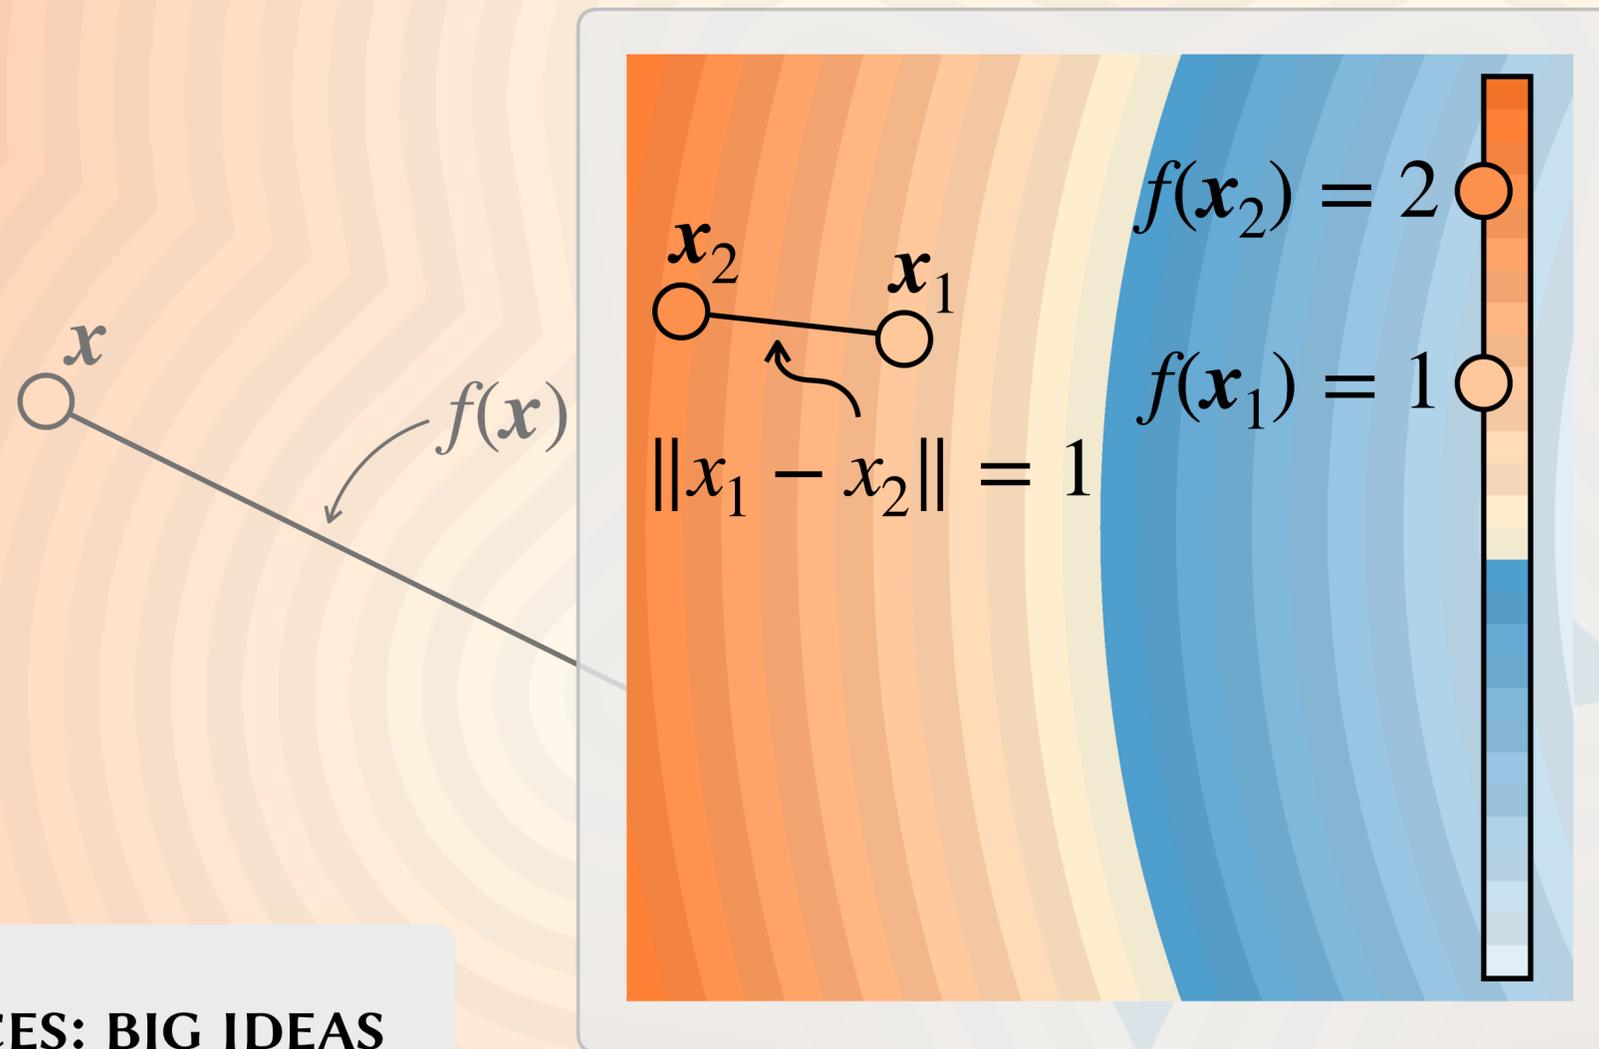


### IMPLICIT SURFACES: BIG IDEAS

1. CSG
2. Taxonomy of implicits
3. Rendering
4. Surface Extraction

## Signed Distance Function

- $|f(\mathbf{x})|$  encodes distance to closest point on surface

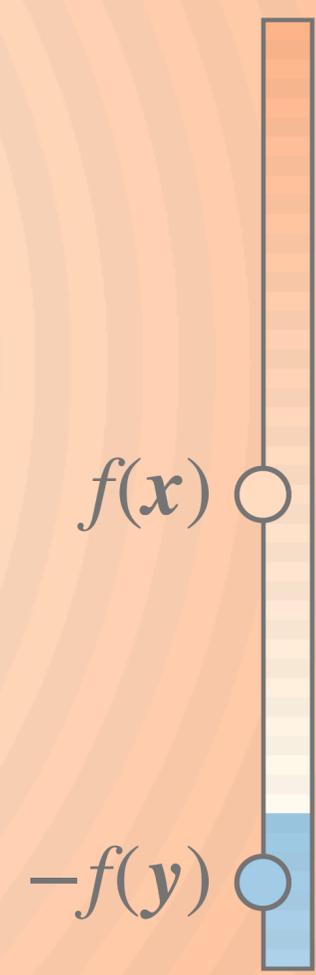
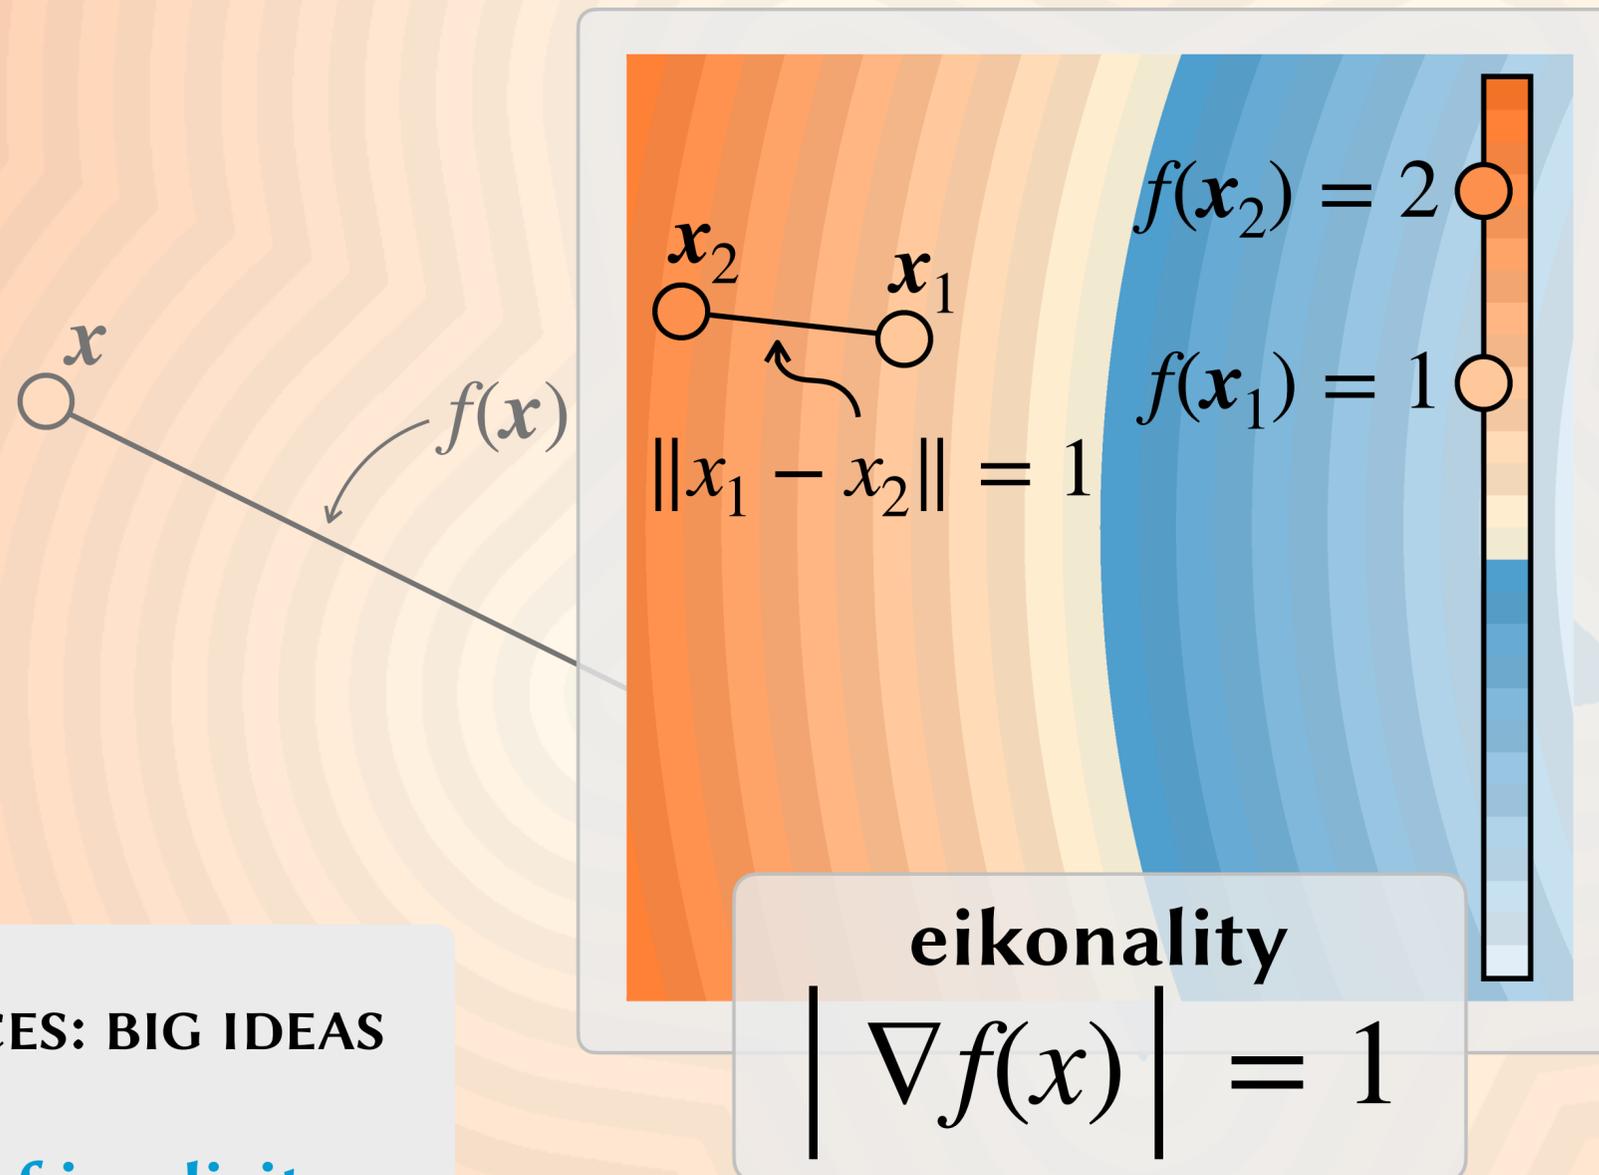


## IMPLICIT SURFACES: BIG IDEAS

1. CSG
2. Taxonomy of implicits
3. Rendering
4. Surface Extraction

### Signed Distance Function

- $|f(x)|$  encodes distance to closest point on surface

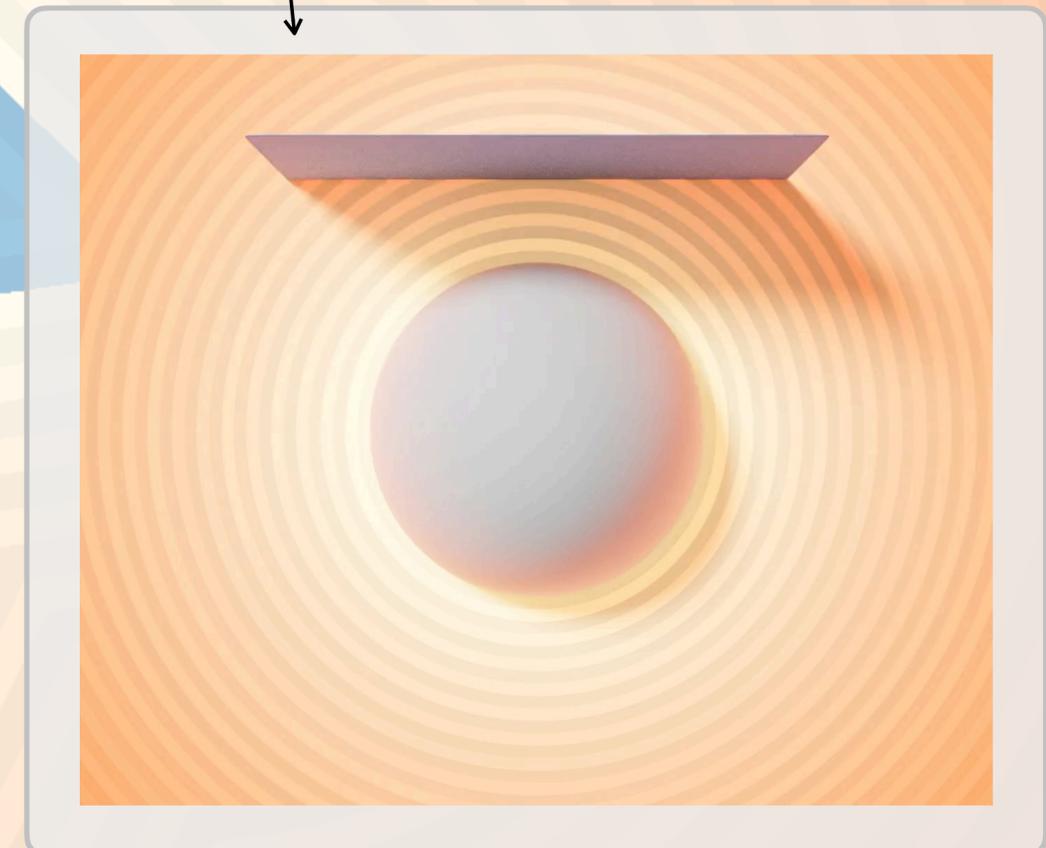
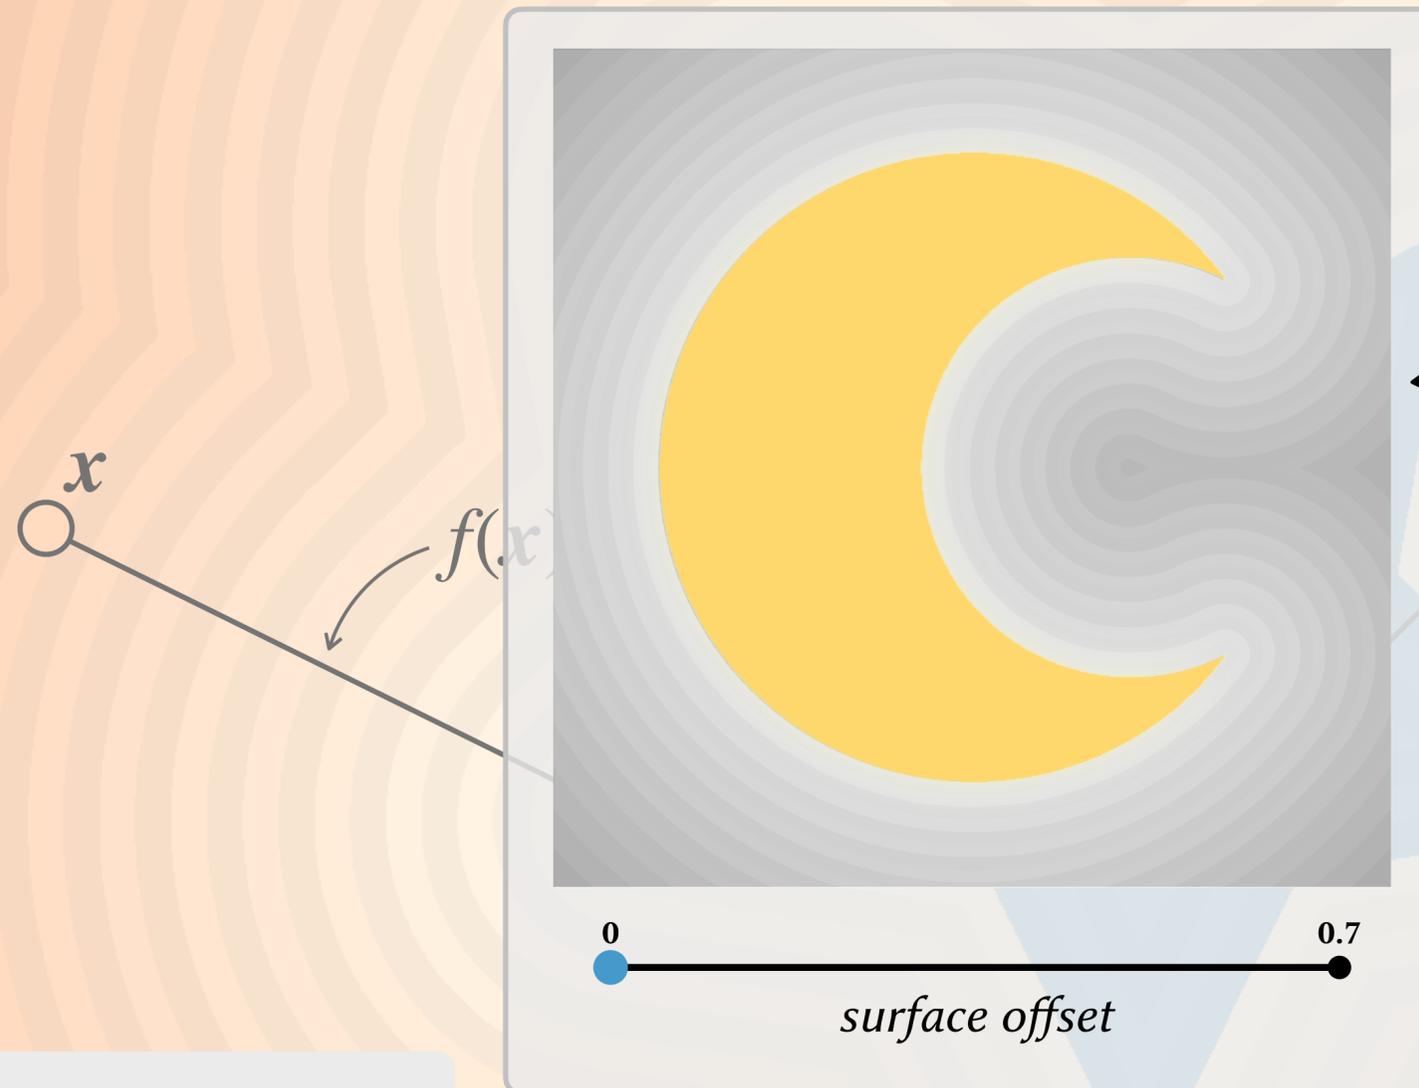


### IMPLICIT SURFACES: BIG IDEAS

1. CSG
2. Taxonomy of implicits
3. Rendering
4. Surface Extraction

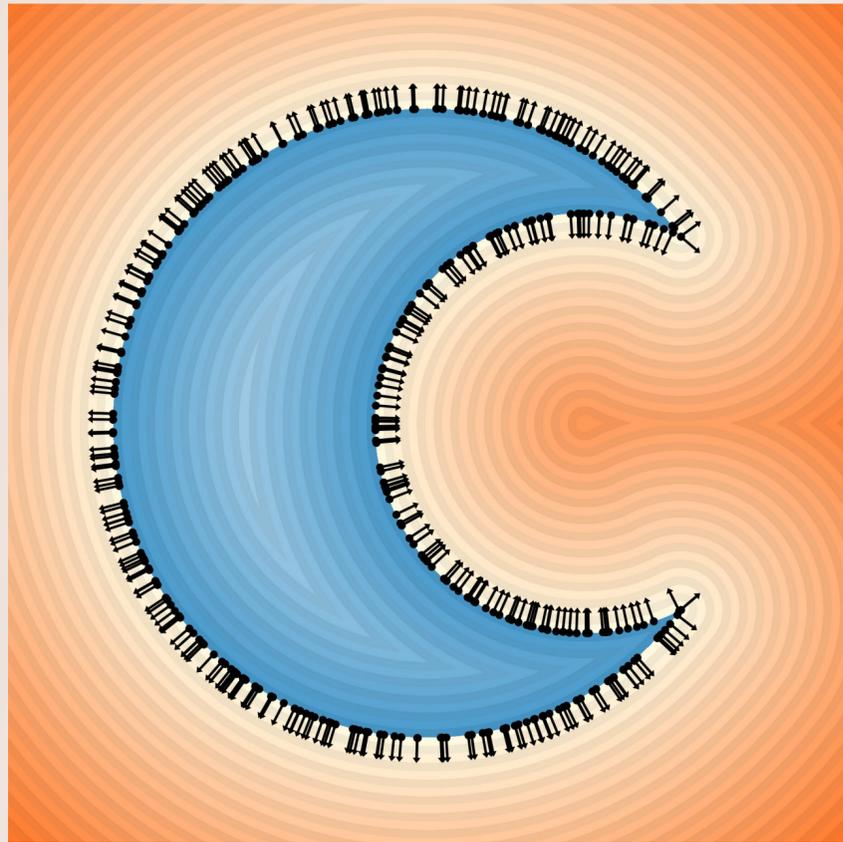
## Signed Distance Function

- $|f(\mathbf{x})|$  encodes distance to closest point on surface
- makes many tasks *way easier* 😊
- e.g., *offset surfaces, distance checks for simulation, etc.*

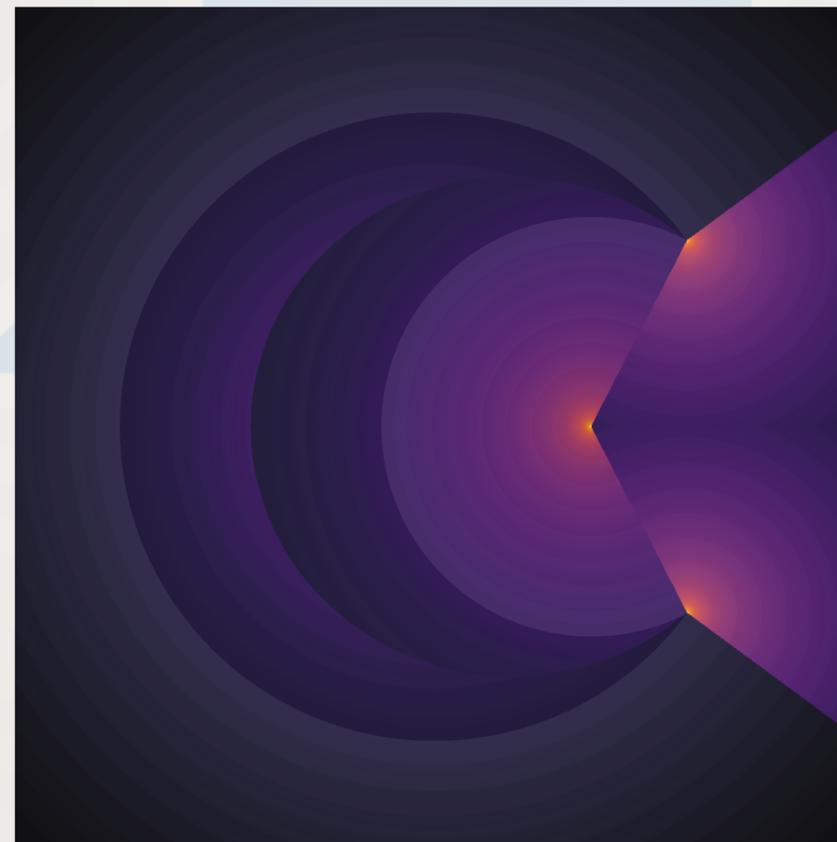


## IMPLICIT SURFACES: BIG IDEAS

1. CSG
2. **Taxonomy of implicits**
3. Rendering
4. Surface Extraction



Surface normals ( $\nabla f$ )



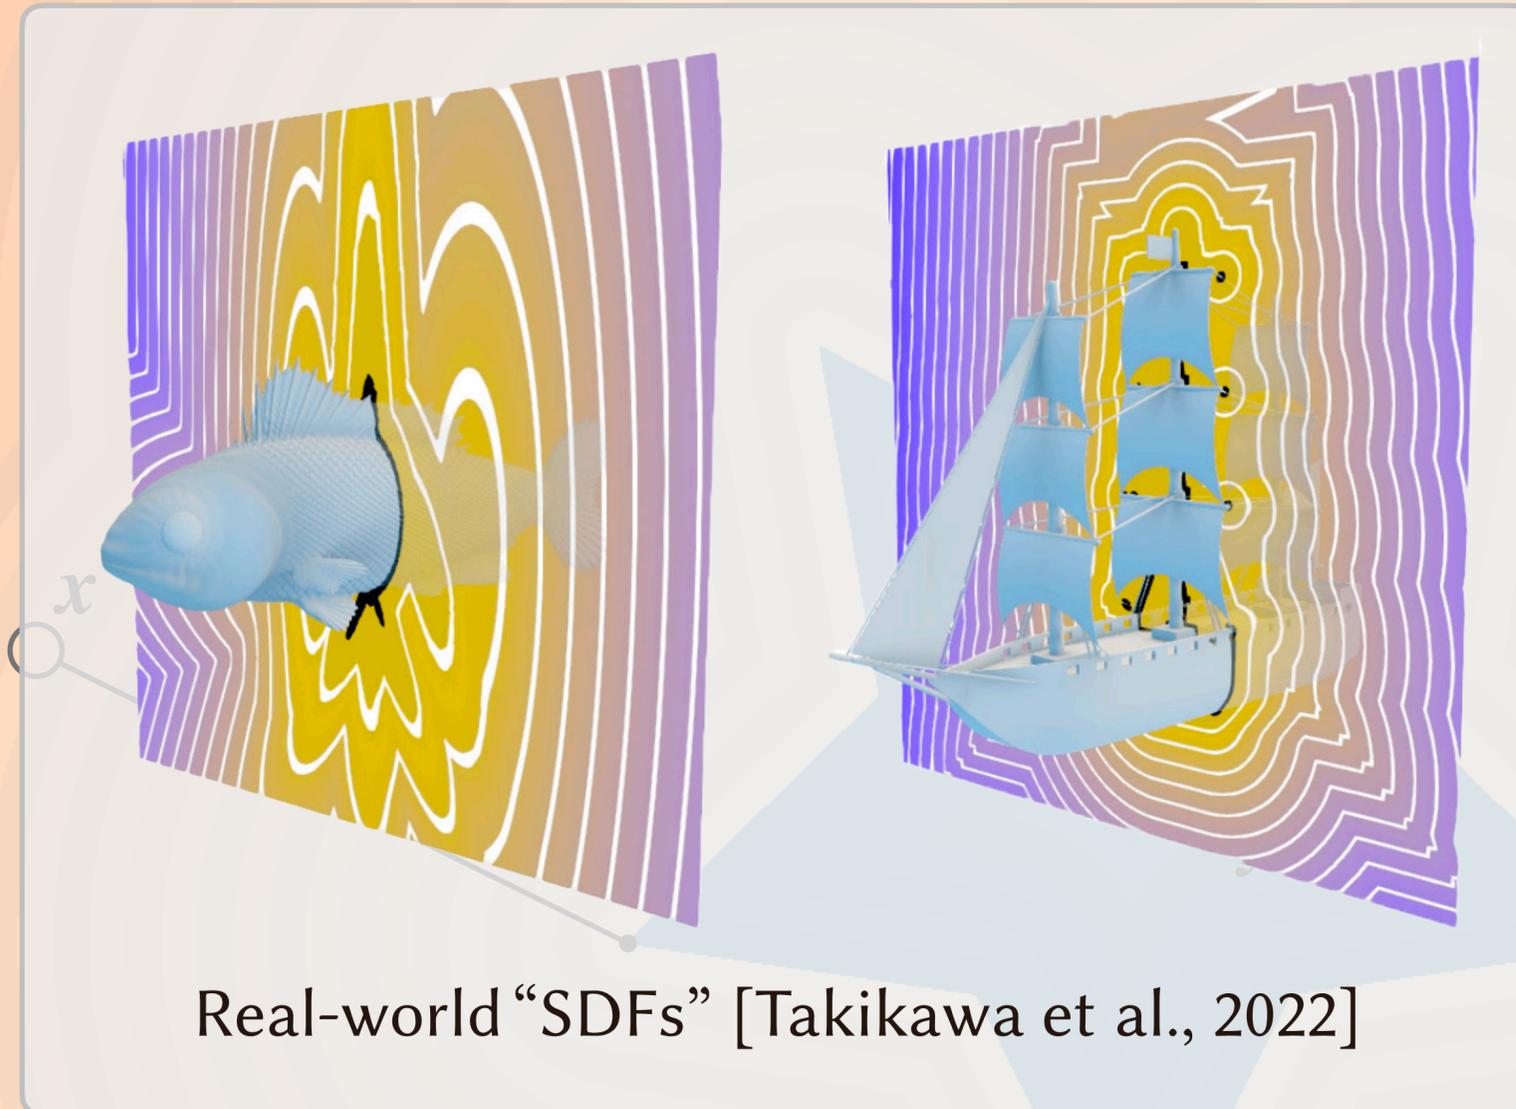
Mean curvature ( $\Delta f$ )

## Signed Distance Function

- $|f(\mathbf{x})|$  encodes distance to closest point on surface
- makes many tasks *way easier* 😊
  - e.g., *offset surfaces, distance checks for simulation, etc.*
  - e.g., *geometric quantities*

## IMPLICIT SURFACES: BIG IDEAS

1. CSG
2. **Taxonomy of implicits**
3. Rendering
4. Surface Extraction



Real-world “SDFs” [Takikawa et al., 2022]

## Signed Distance Function

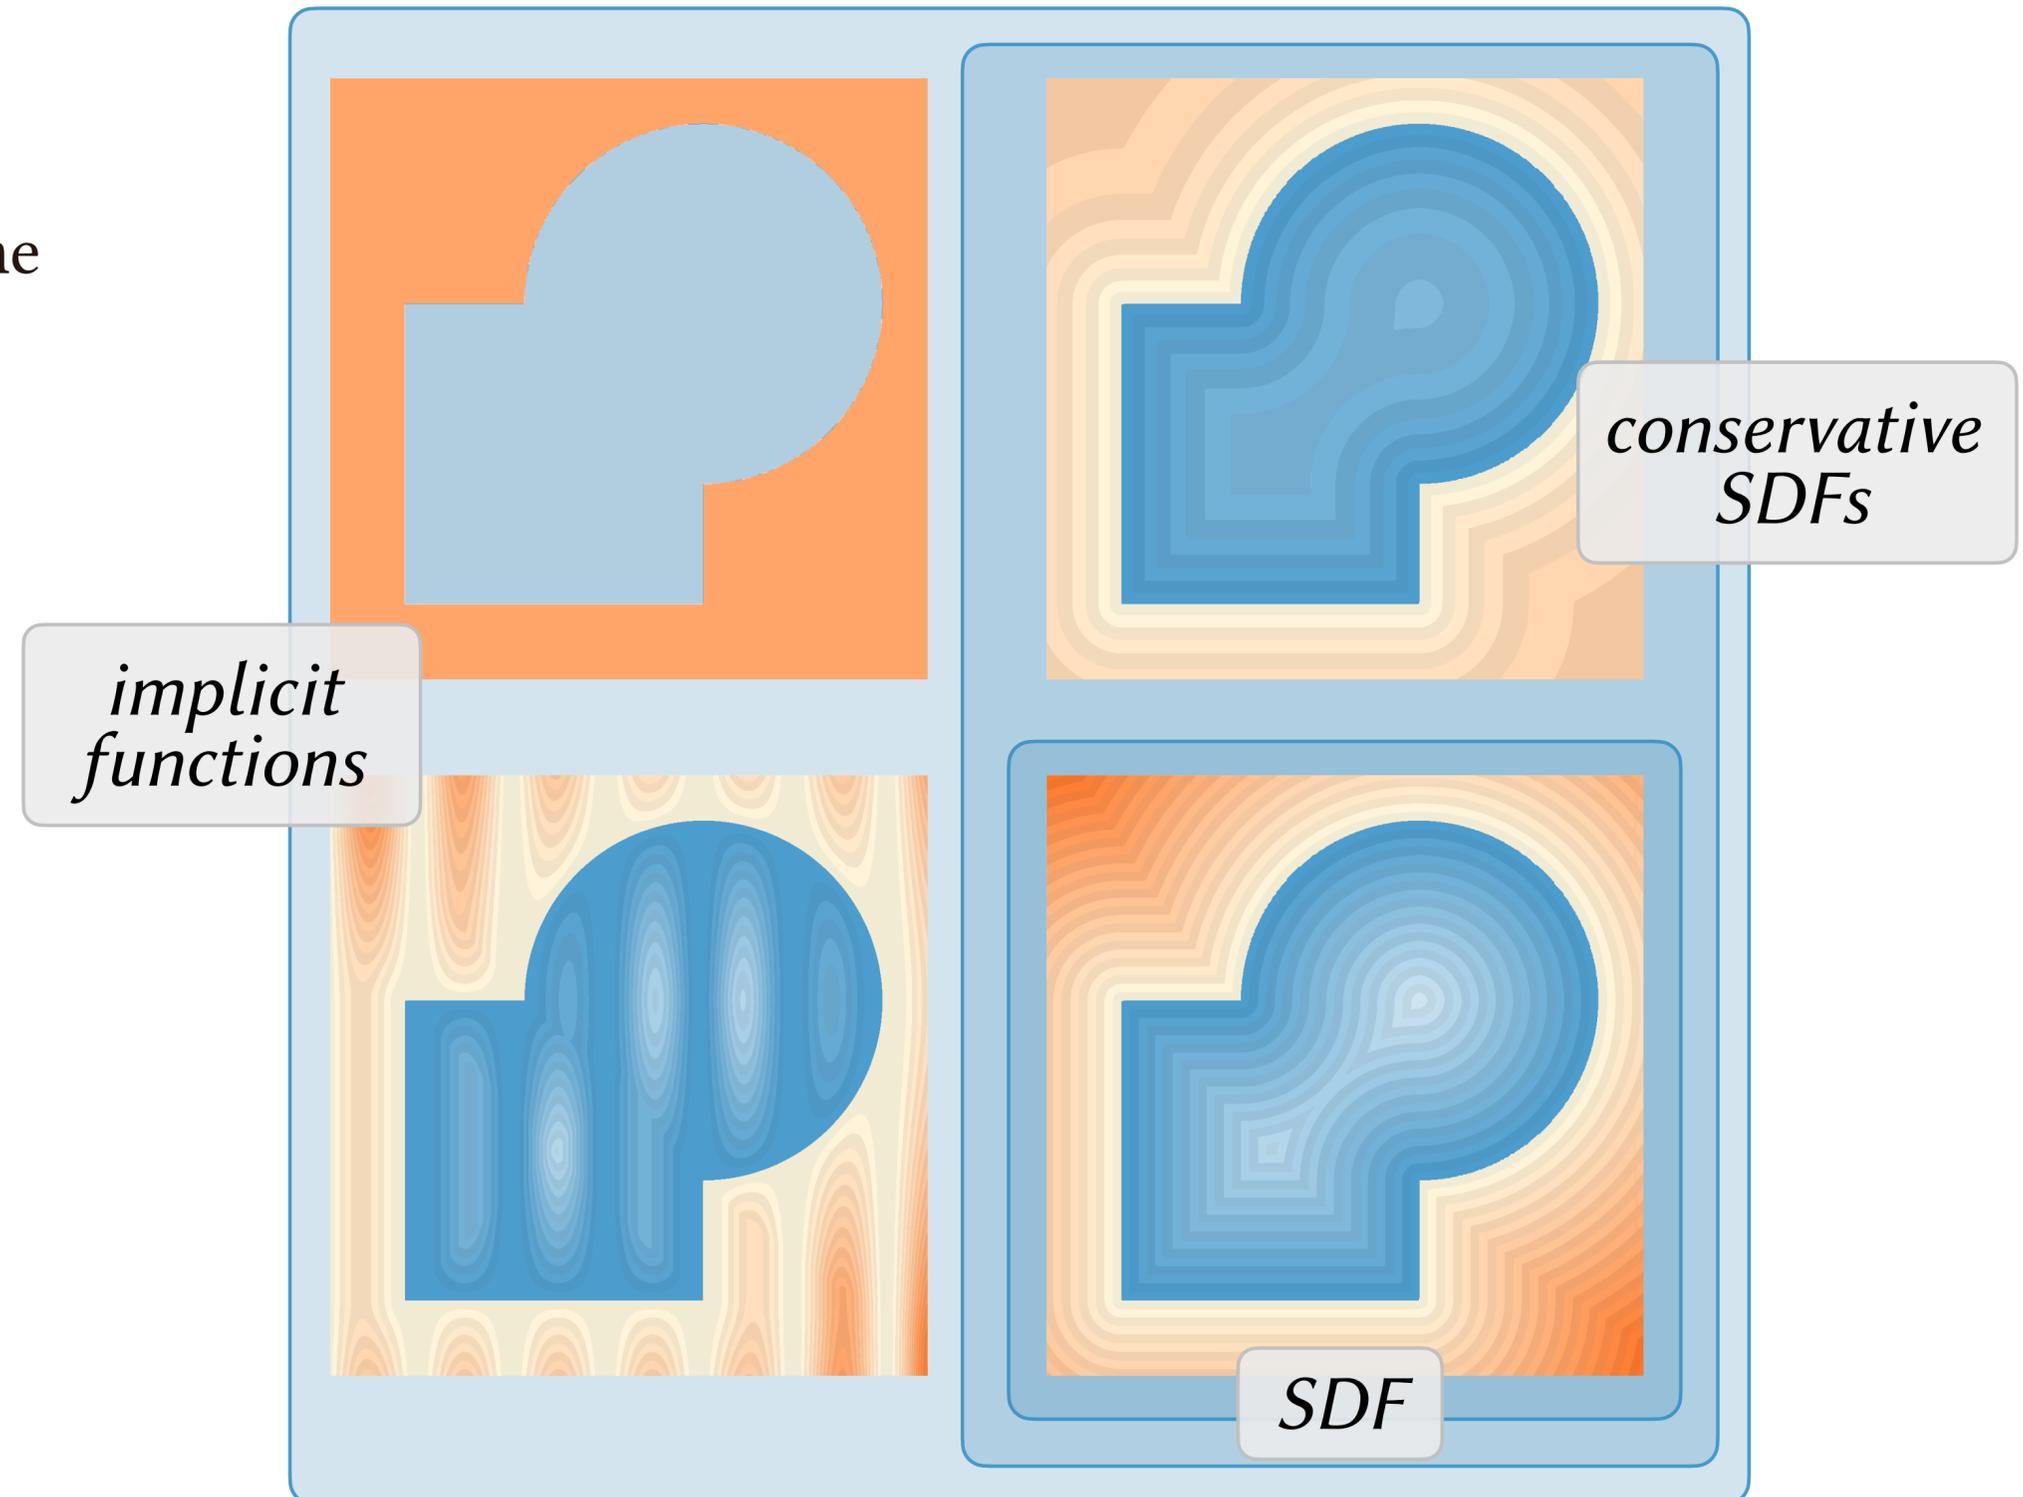
- $|f(\mathbf{x})|$  encodes distance to closest point on surface
- makes many tasks *way* easier 😊
- distance property is hard to maintain ☹️
  - *e.g., not preserved by most editing operations*

## IMPLICIT SURFACES: BIG IDEAS

1. CSG
2. **Taxonomy of implicits**
3. Rendering
4. Surface Extraction

## Conservative SDFs

- (aka approximate SDFs)
- $|f(\mathbf{x})|$  is less than the distance to the closest point on the surface

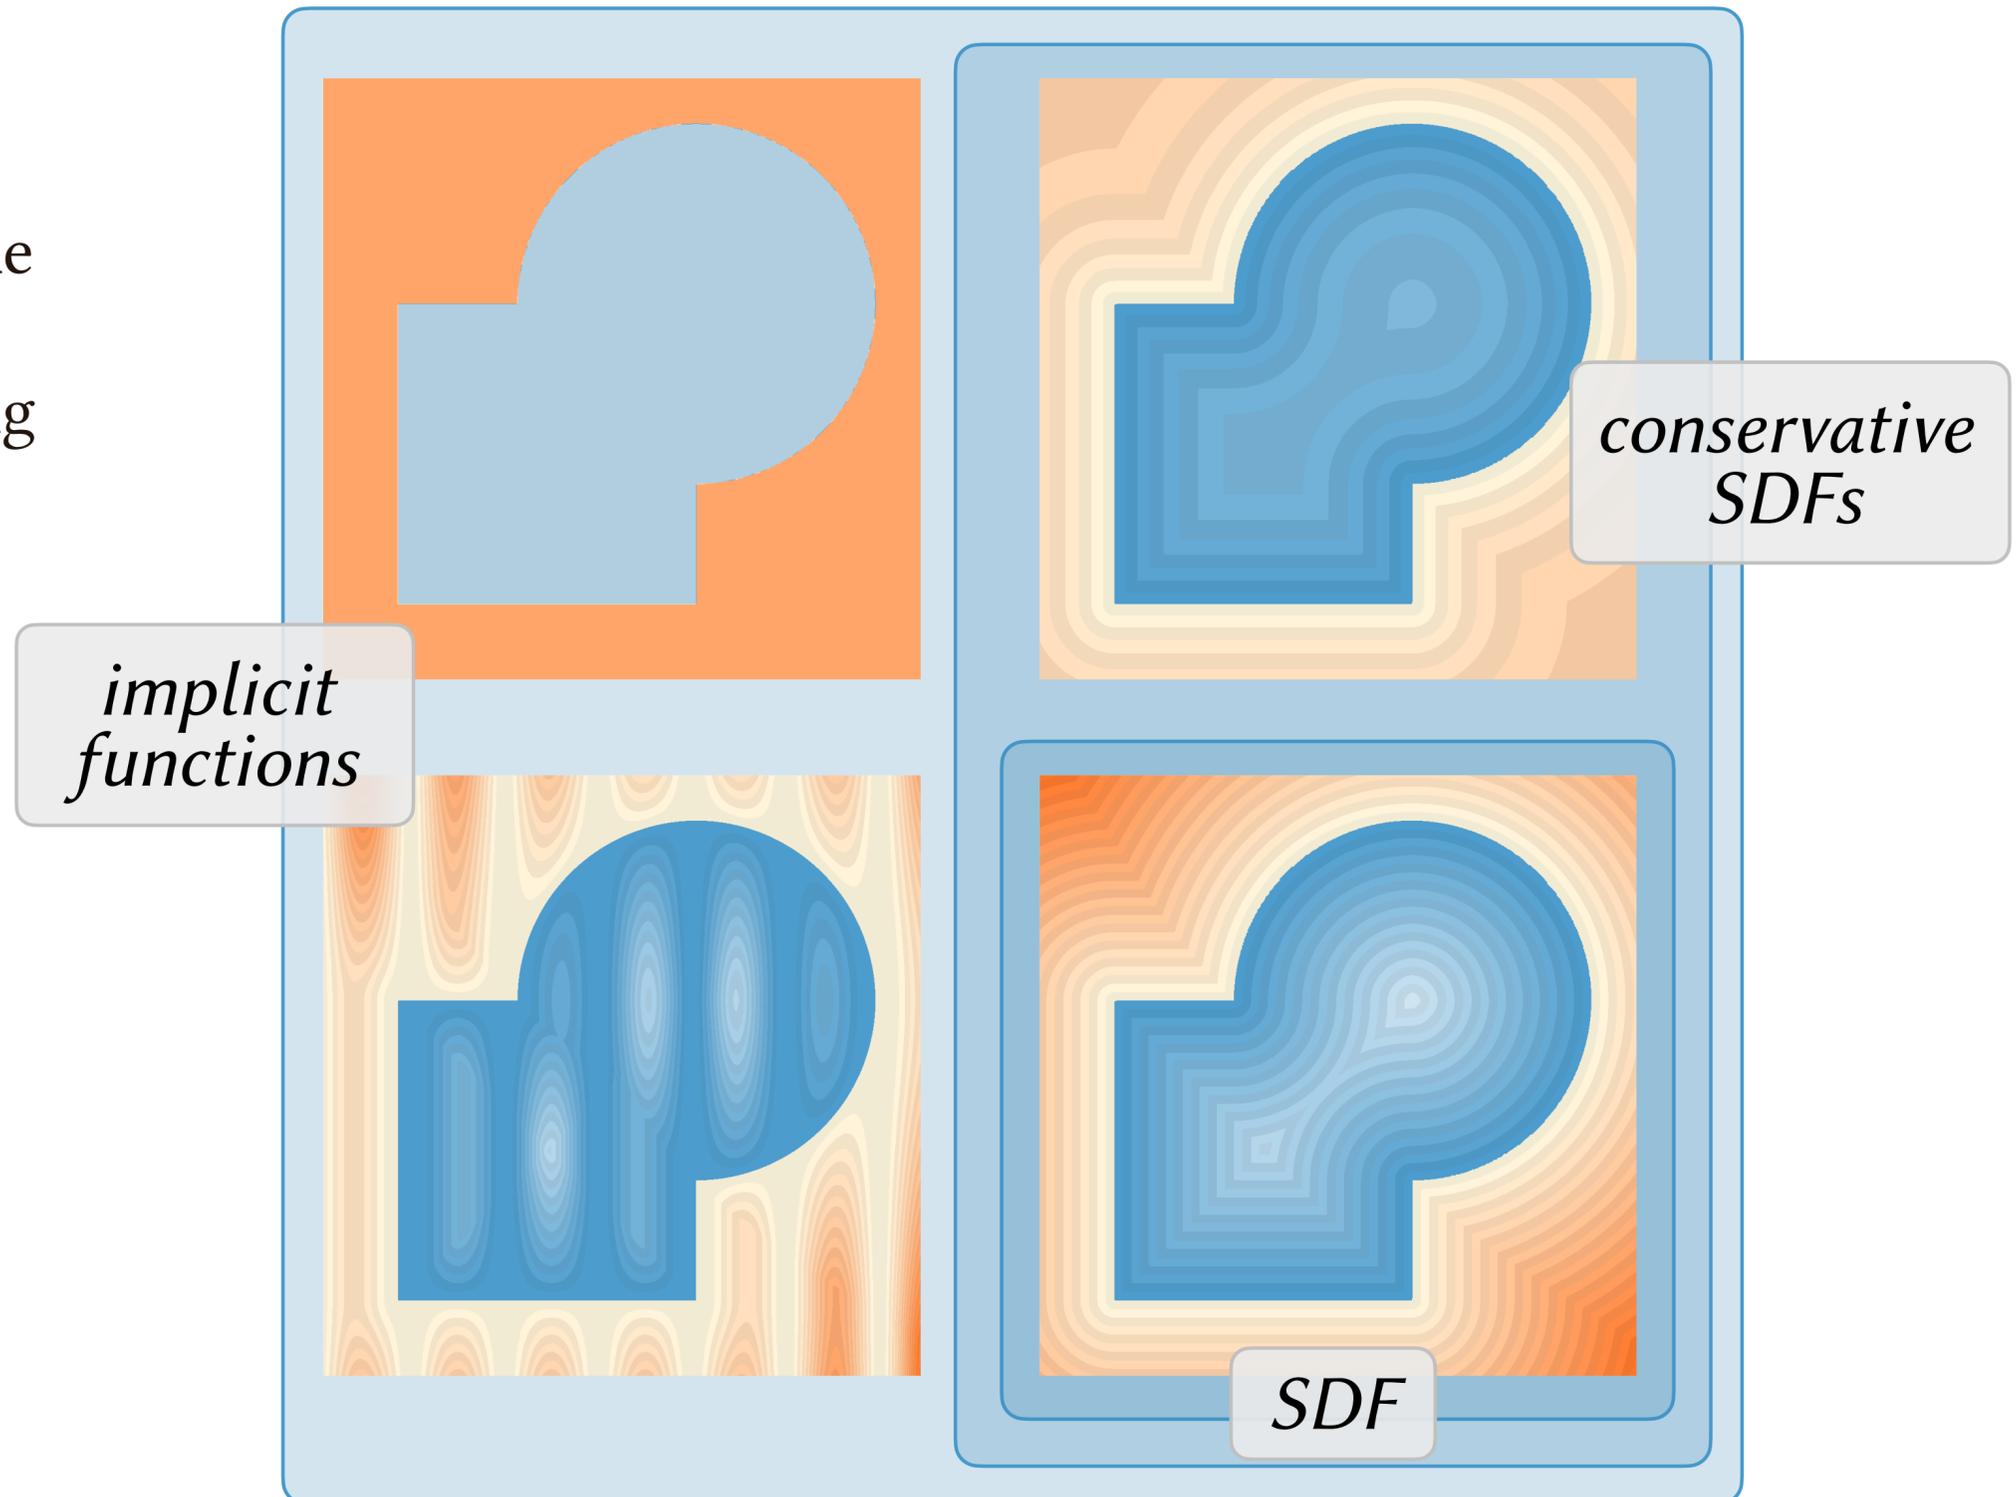


### IMPLICIT SURFACES: BIG IDEAS

1. CSG
2. **Taxonomy of implicits**
3. Rendering
4. Surface Extraction

## Conservative SDFs

- (aka approximate SDFs)
- $|f(\mathbf{x})|$  is less than the distance to the closest point on the surface
- tradeoff between ease of maintaining & having useful properties

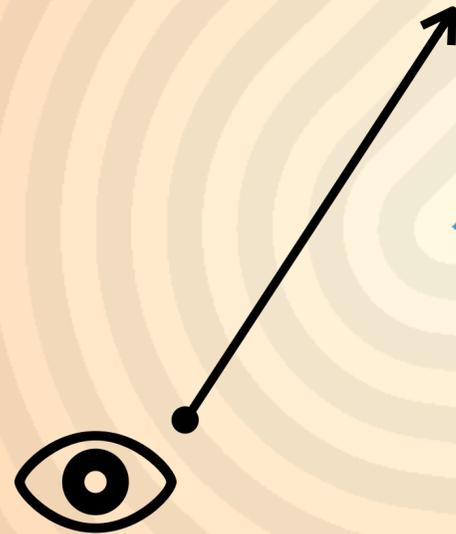


### IMPLICIT SURFACES: BIG IDEAS

1. CSG
2. **Taxonomy of implicits**
3. Rendering
4. Surface Extraction

## Sphere Tracing

- introduced by [Hart, 1996]
- technique to ray trace conservative SDFs



### IMPLICIT SURFACES: BIG IDEAS

1. CSG
2. Taxonomy of implicits
3. **Rendering**
4. Surface Extraction

## Sphere Tracing

- introduced by [Hart, 1996]
- technique to ray trace conservative SDFs
- $|f(x)|$  is distance we can *definitely* travel without crossing the surface

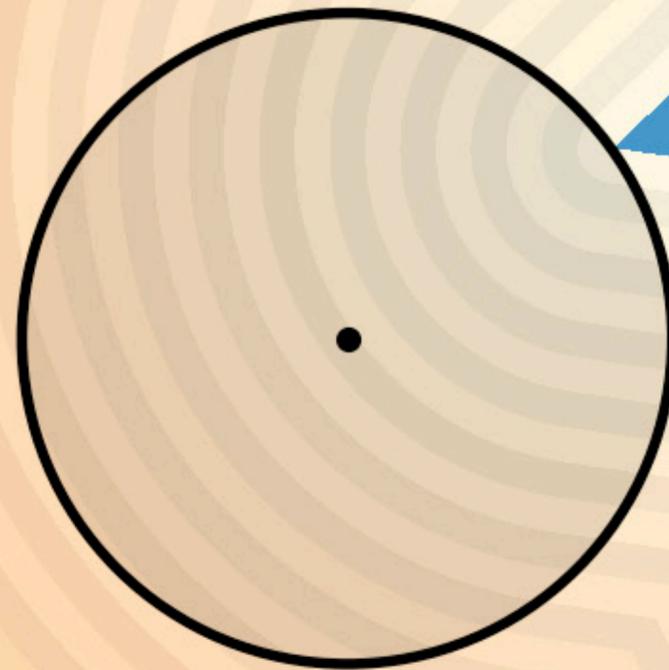


### IMPLICIT SURFACES: BIG IDEAS

1. CSG
2. Taxonomy of implicits
3. **Rendering**
4. Surface Extraction

## Sphere Tracing

- introduced by [Hart, 1996]
- technique to ray trace conservative SDFs
- $|f(x)|$  is distance we can *definitely* travel without crossing the surface



### IMPLICIT SURFACES: BIG IDEAS

1. CSG
2. Taxonomy of implicits
3. **Rendering**
4. Surface Extraction

## Marching cubes

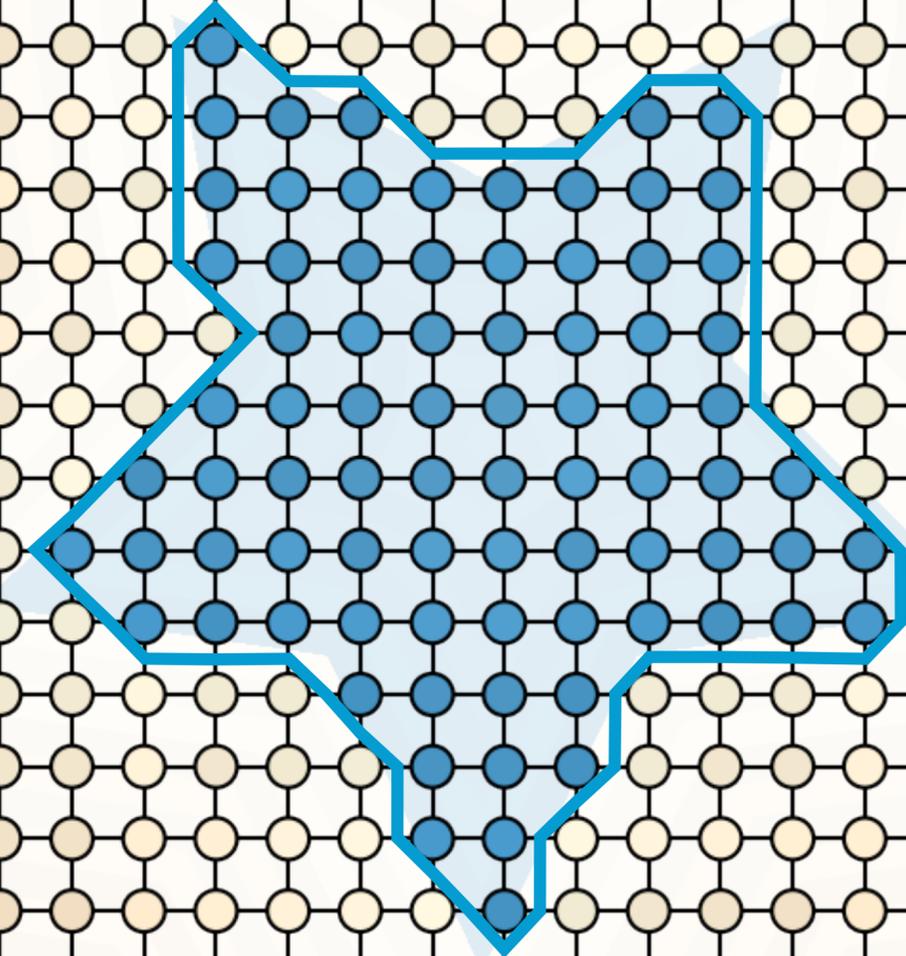
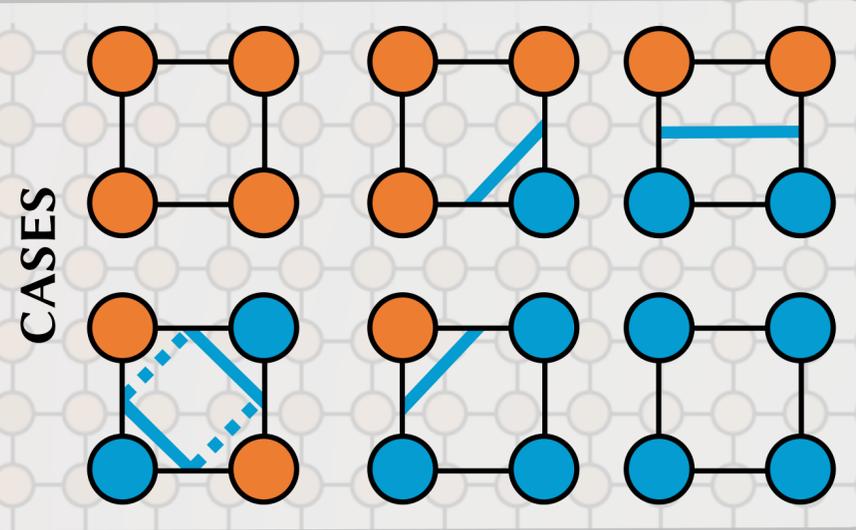
- introduced by [Lorensen & Cline, 1987]
- identify voxels with surface present, draw surface in each of these voxels

### IMPLICIT SURFACES: BIG IDEAS

1. CSG
2. Taxonomy of implicits
3. Rendering
4. **Surface Extraction**

## Marching cubes

- introduced by [Lorensen & Cline, 1987]
- identify voxels with surface present, draw surface in each of these voxels

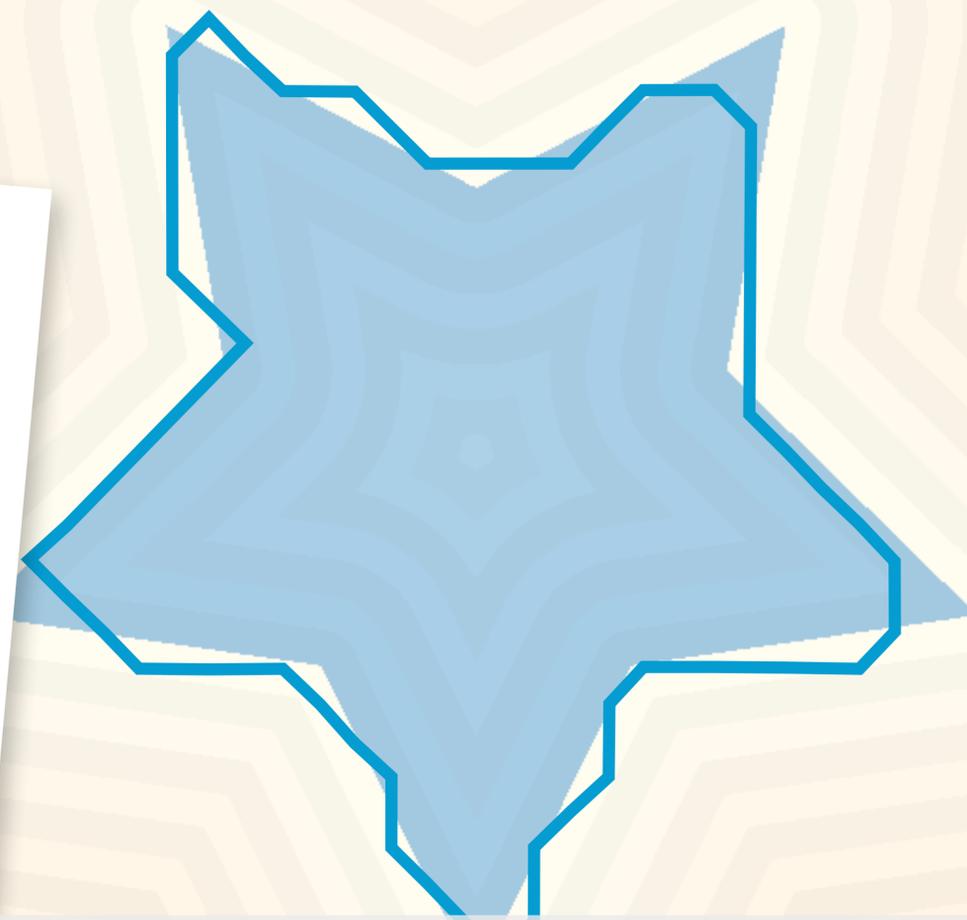


### IMPLICIT SURFACES: BIG IDEAS

1. CSG
2. Taxonomy of implicits
3. Rendering
4. **Surface Extraction**

# Marching cubes

- introduced by [Lorensen & Cline, 1987]
- identify voxels with surface present, draw surface in each of these voxels

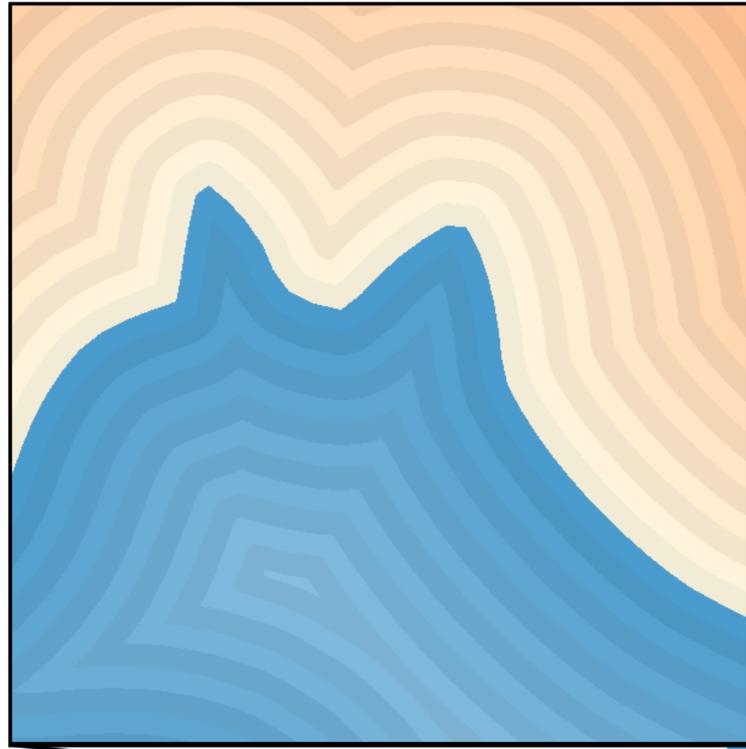


## IMPLICIT SURFACE

1. CSG
2. Taxonomy of implicits
3. Rendering
4. Surface Extraction

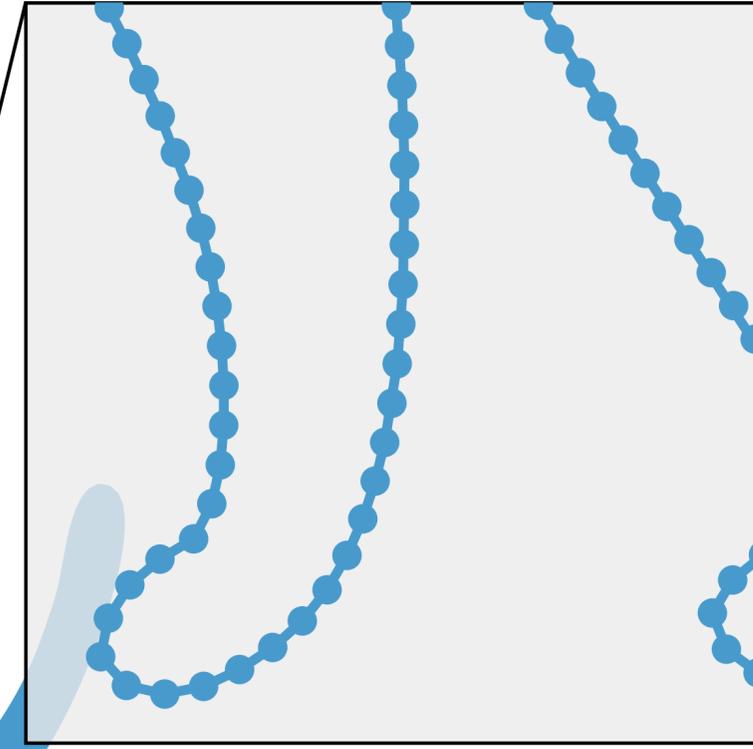
- recent work on improved surface extraction, e.g.,
  - using distance info for SDF case [Sellán et al., 2023]
  - using machine learning & data [Chen et al. 2022]

# Implicit vs. explicit representations



## IMPLICIT

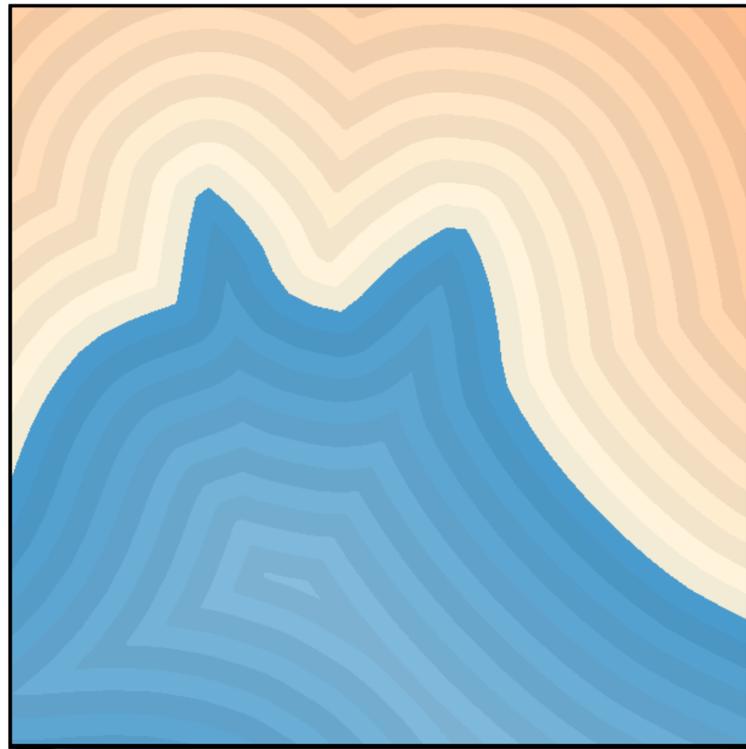
- answer questions about points in space
  - e.g., is this point inside the surface?



## EXPLICIT

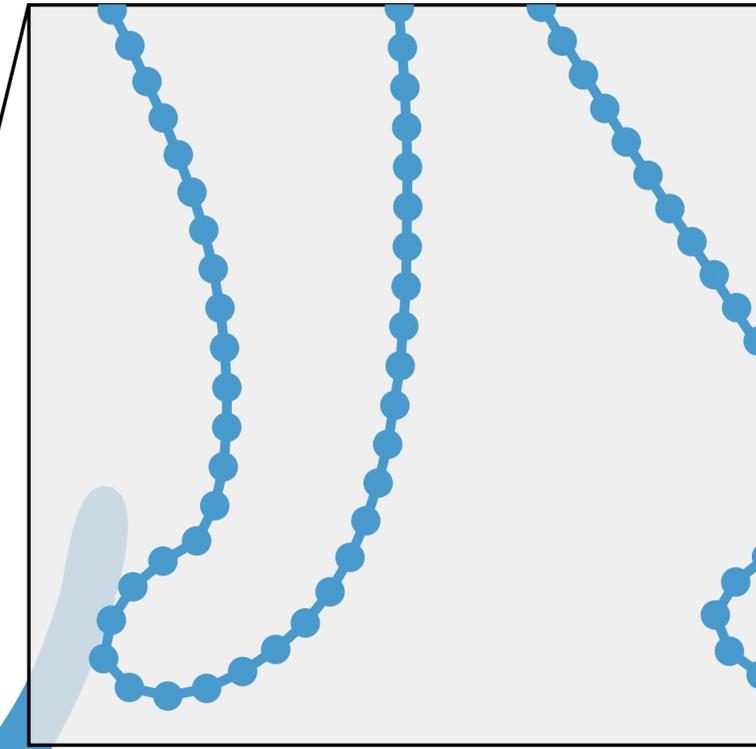
- answer questions about surface
  - e.g., what is the total surface area

# Implicit vs. explicit representations



## IMPLICIT

- answer questions about points in space
  - e.g., is this point inside the surface?
- watertight, easy CSG, ...



## EXPLICIT

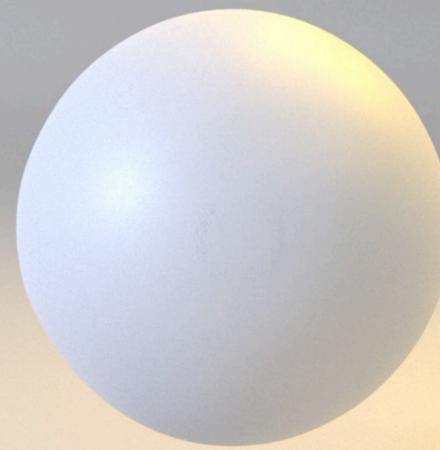
- answer questions about surface
  - e.g., what is the total surface area
- often faster, often want to operate directly on surface, ...

# Implicit vs. explicit representations

in *implicit* representations, topology changes are continuous (very useful for inverse tasks!)

## Optimization

Optimizing shape, albedo  
& roughness

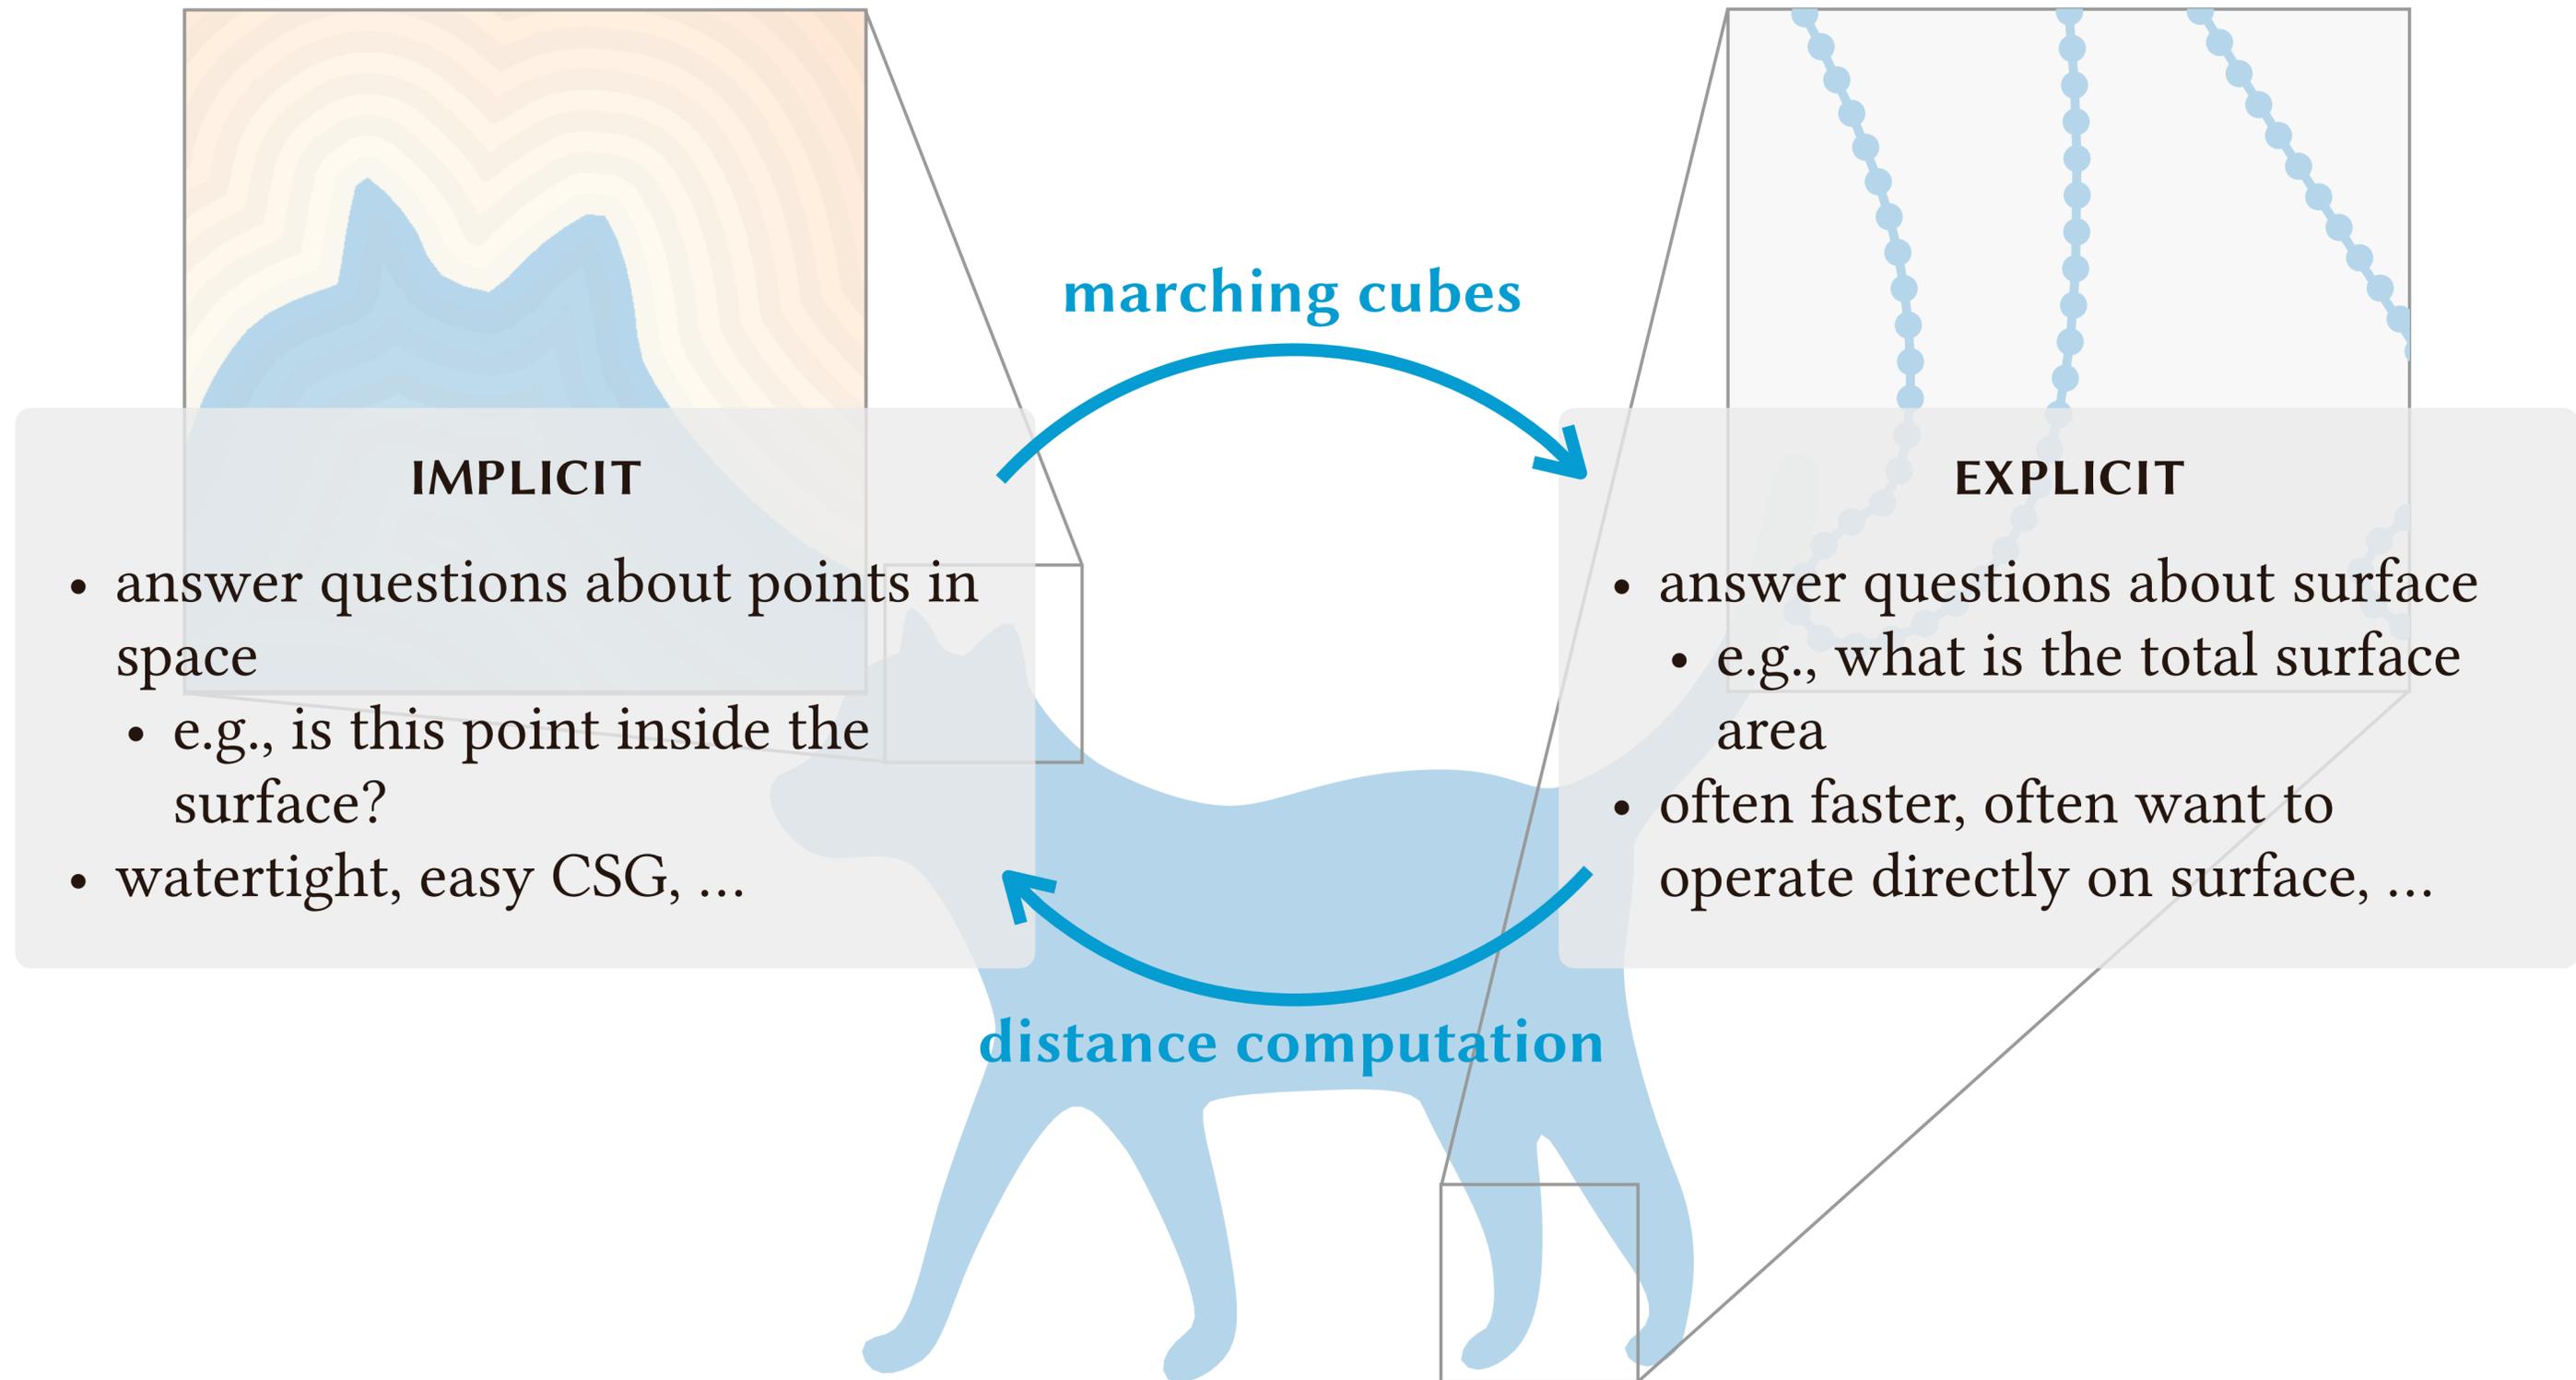


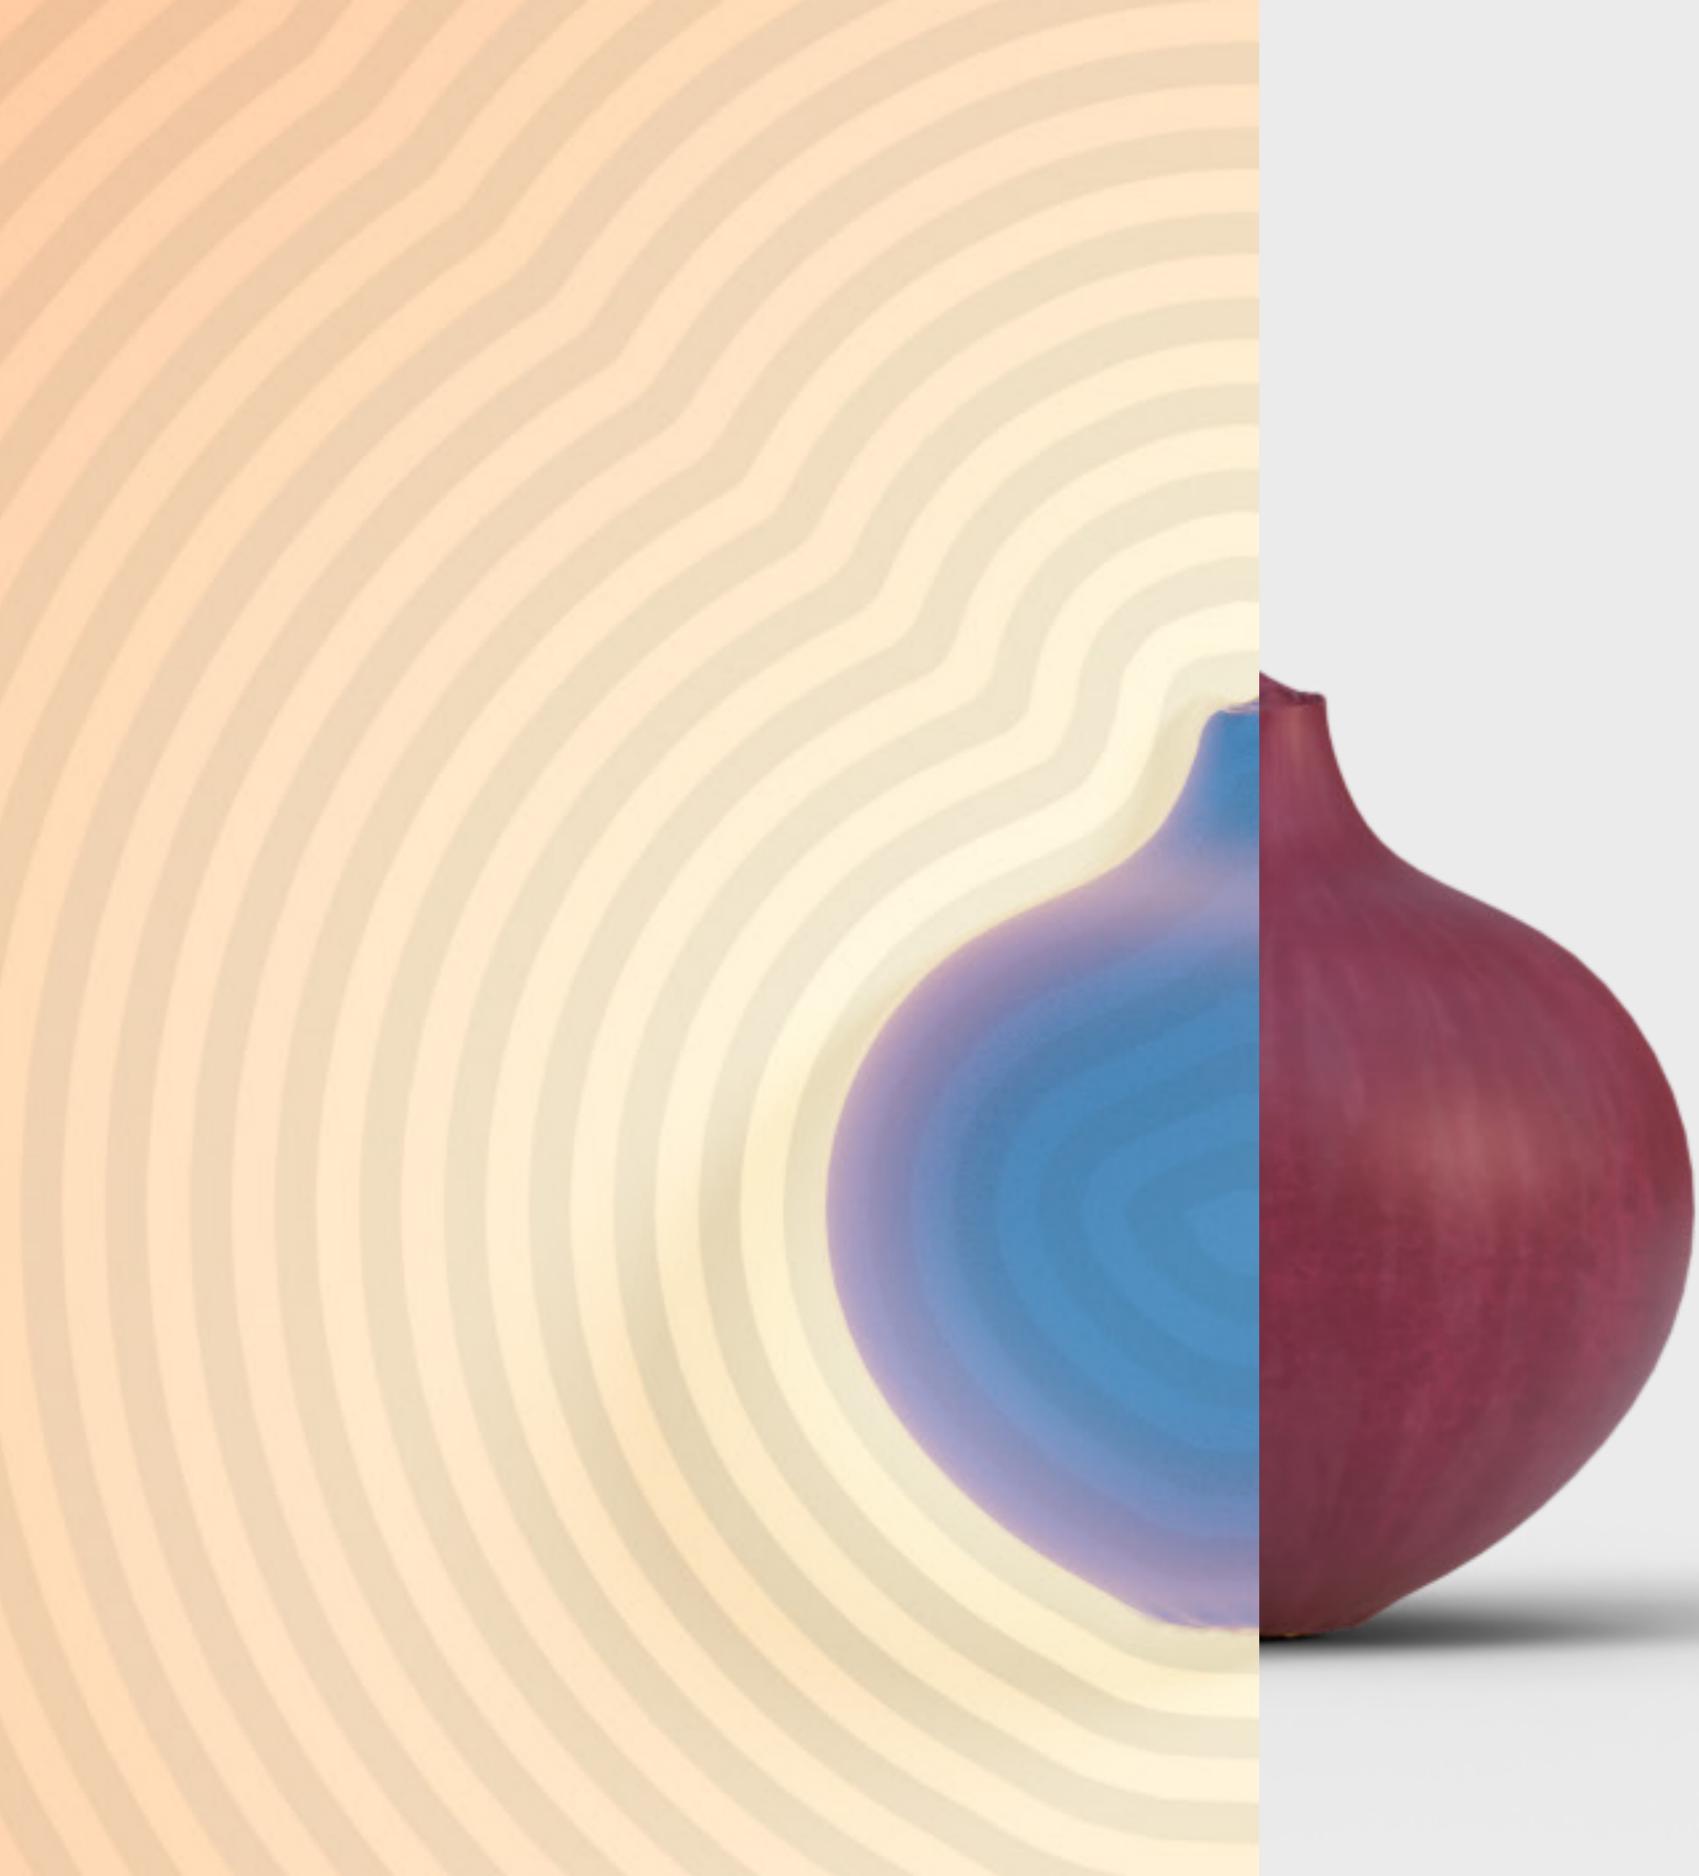
Iteration 0

- answer questions about surface space
  - e.g., is this surface?
- watertight, ea

about surface  
total surface  
want to  
n surface, ...

# Implicit vs. explicit representations





## OVERVIEW

### IMPLICIT SURFACES

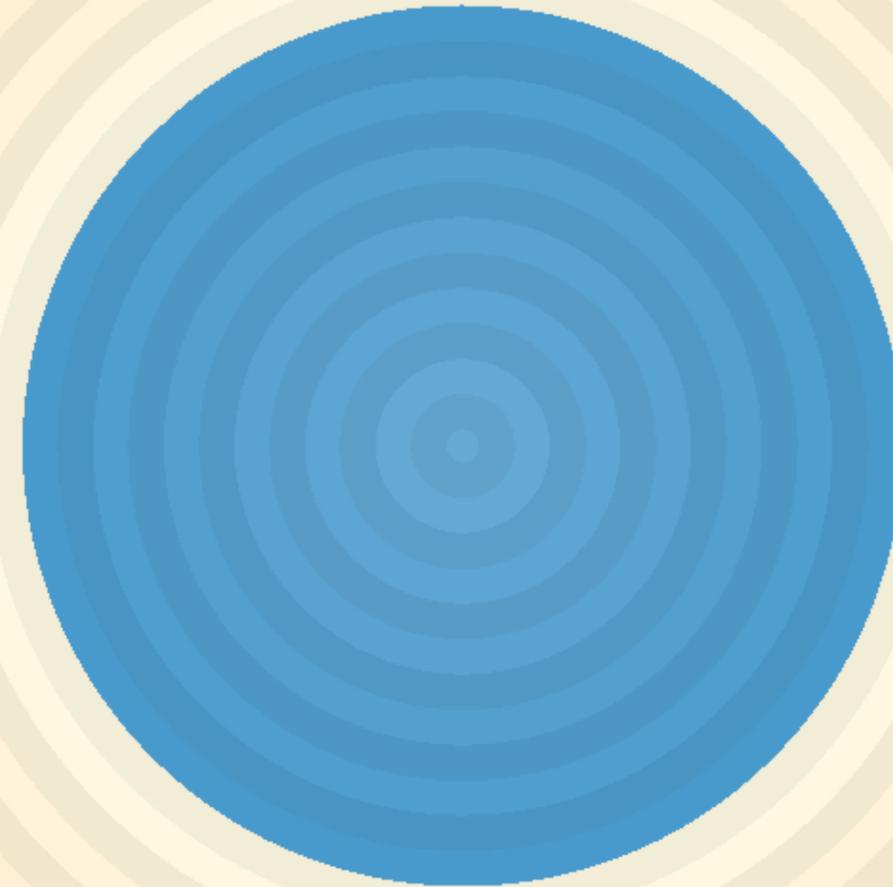
1. Background & History
2. Digital Representations

### NEURAL IMPLICIT

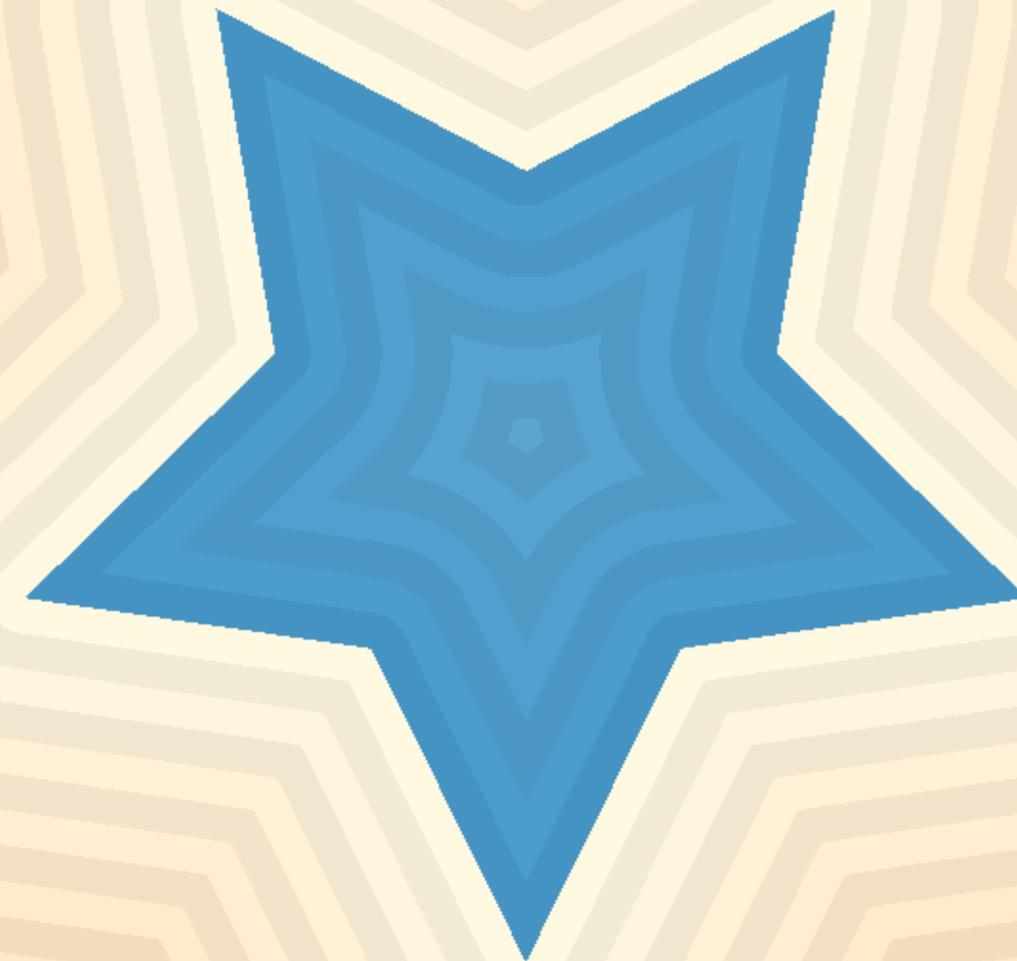
3. The Basics
4. Why & Why Not Use Them?
5. Working out the Details
6. Geometric Operations

### CONCLUSION

$$f(x, y) = x^2 + y^2 - r^2$$



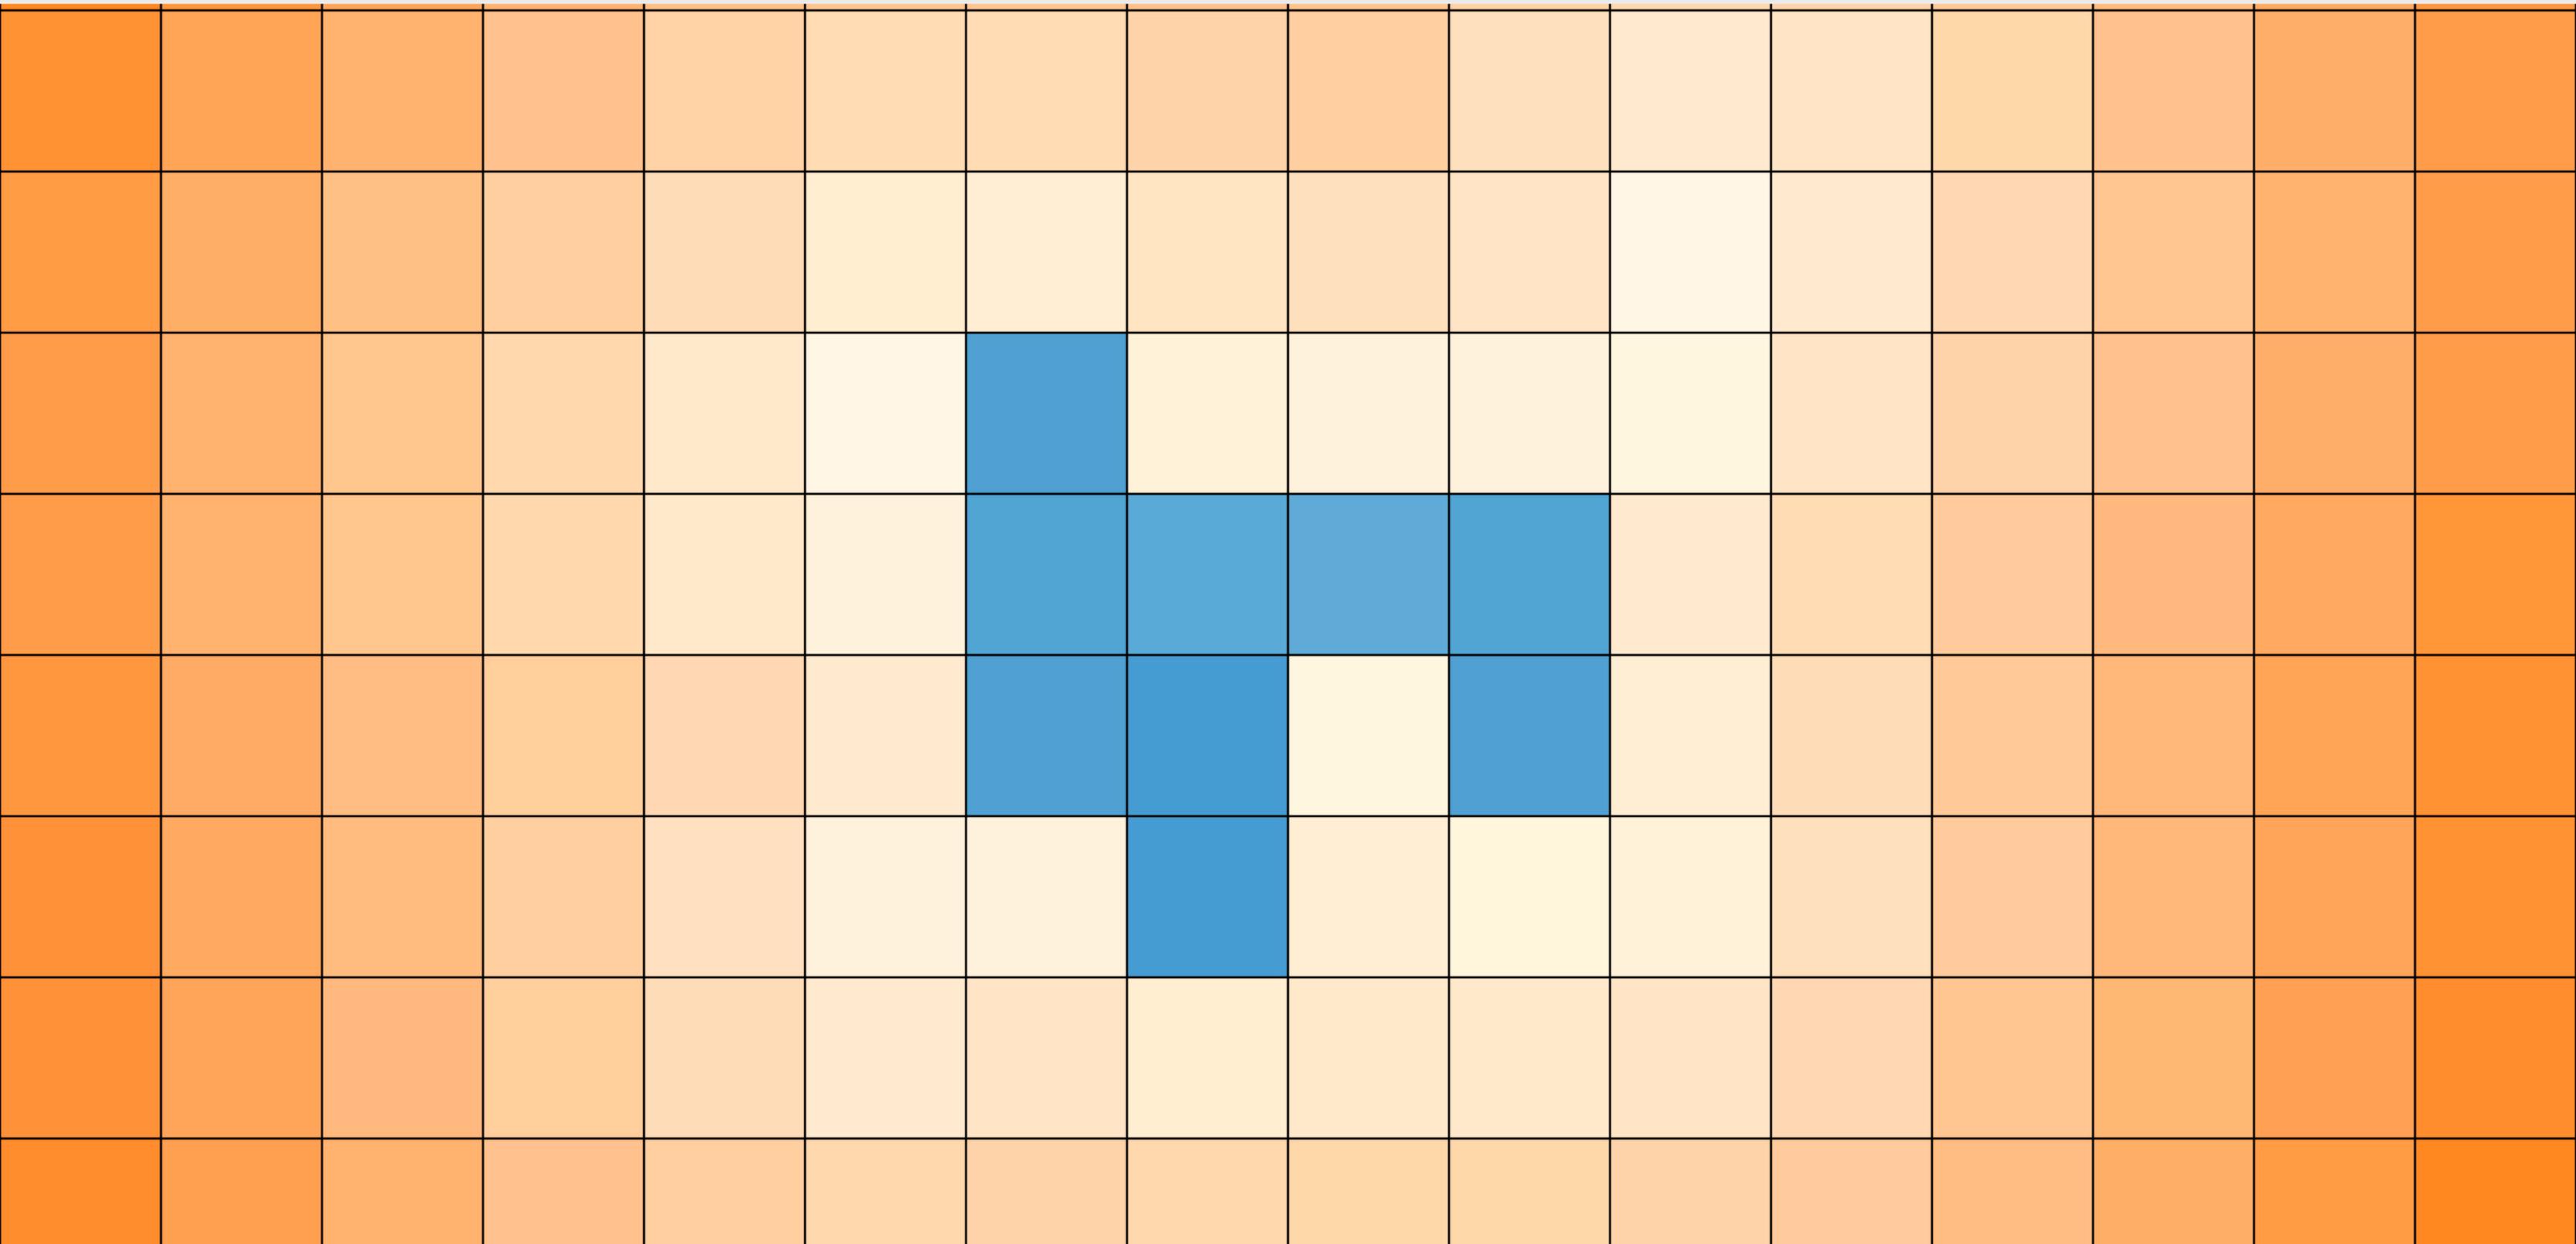
```
float sdStar5(in vec2 p, in float r, in float rf)
{
    const vec2 k1 = vec2(0.809016994375, -0.587785252292);
    const vec2 k2 = vec2(-k1.x,k1.y);
    p.x = abs(p.x);
    p -= 2.0*max(dot(k1,p),0.0)*k1;
    p -= 2.0*max(dot(k2,p),0.0)*k2;
    p.x = abs(p.x);
    p.y -= r;
    vec2 ba = rf*vec2(-k1.y,k1.x) - vec2(0,1);
    float h = clamp( dot(p,ba)/dot(ba,ba), 0.0, r );
    return length(p-ba*h) * sign(p.y*ba.x-p.x*ba.y);
}
```



- slow to compute complex formulas
- can't always find analytic representation

📖 For analytic definitions of SDFs, see “Distance Functions” page on Inigo Quillez’s blog

# Grid-based representations



# Grid-based representations: just a function!

$w_{0,0}$	$w_{1,0}$	$w_{2,0}$	$w_{3,0}$
$w_{0,1}$	$w_{1,1}$	$w_{2,1}$	$w_{3,1}$
$w_{0,2}$	$w_{1,2}$	$w_{2,2}$	$w_{3,2}$
$w_{0,3}$	$w_{1,3}$	$w_{2,3}$	$w_{3,3}$

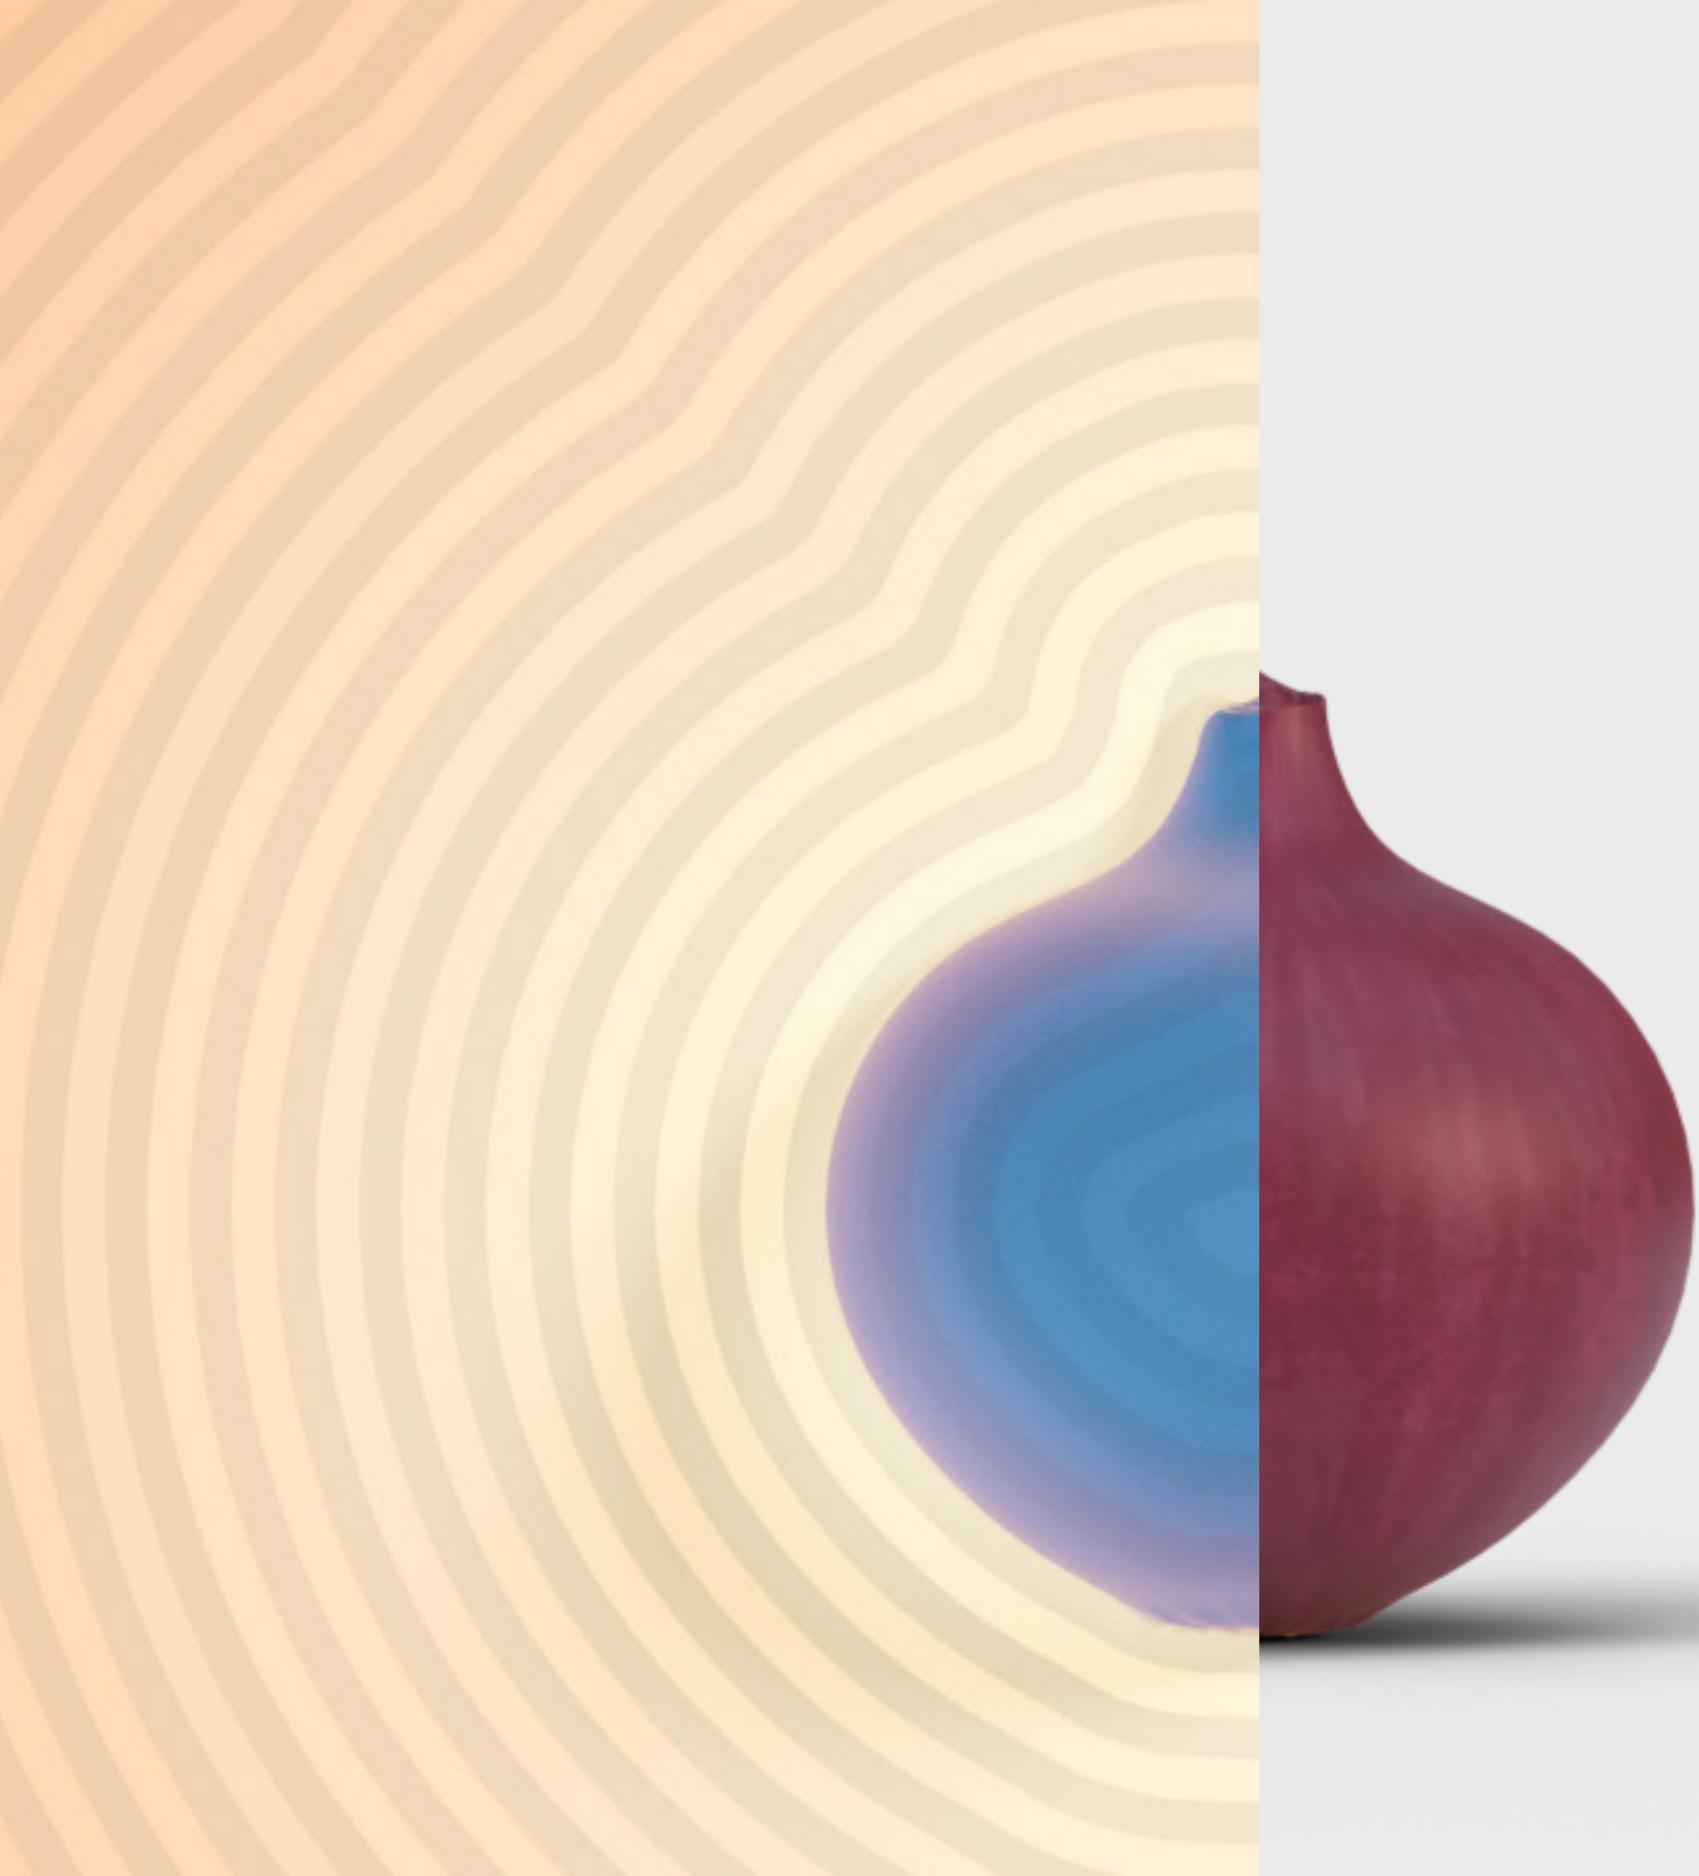
$$f(x, y) = ([x] = 0 \ \& \ [y] = 0) w_{0,0} + ([x] = 1 \ \& \ [y] = 0) w_{1,0} + \dots$$

# Generalization: *any* function space

Degree 0



$$f(x, y) = w_{n,0}x^n + w_{n-1,1}x^{n-1}y + \dots + w_{0,1}y + w_{0,0}$$



## OVERVIEW

### IMPLICIT SURFACES

1. Background & History
2. Digital Representations

### NEURAL IMPLICIT

3. **The Basics**
4. Why & Why Not Use Them?
5. Working out the Details
6. Geometric Operations

### CONCLUSION

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

$$f_{\theta}(x) = A_{\theta}^N \varphi(A_{\theta}^1 \varphi(A_{\theta}^0 x + b_{\theta}^0) + b_{\theta}^1 + \dots) + b_{\theta}^N$$

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

$$f_{\theta}(x) = A_{\theta}^N \varphi(A_{\theta}^1 \varphi(A_{\theta}^0 x + b_{\theta}^0) + b_{\theta}^1 + \dots) + b_{\theta}^N$$

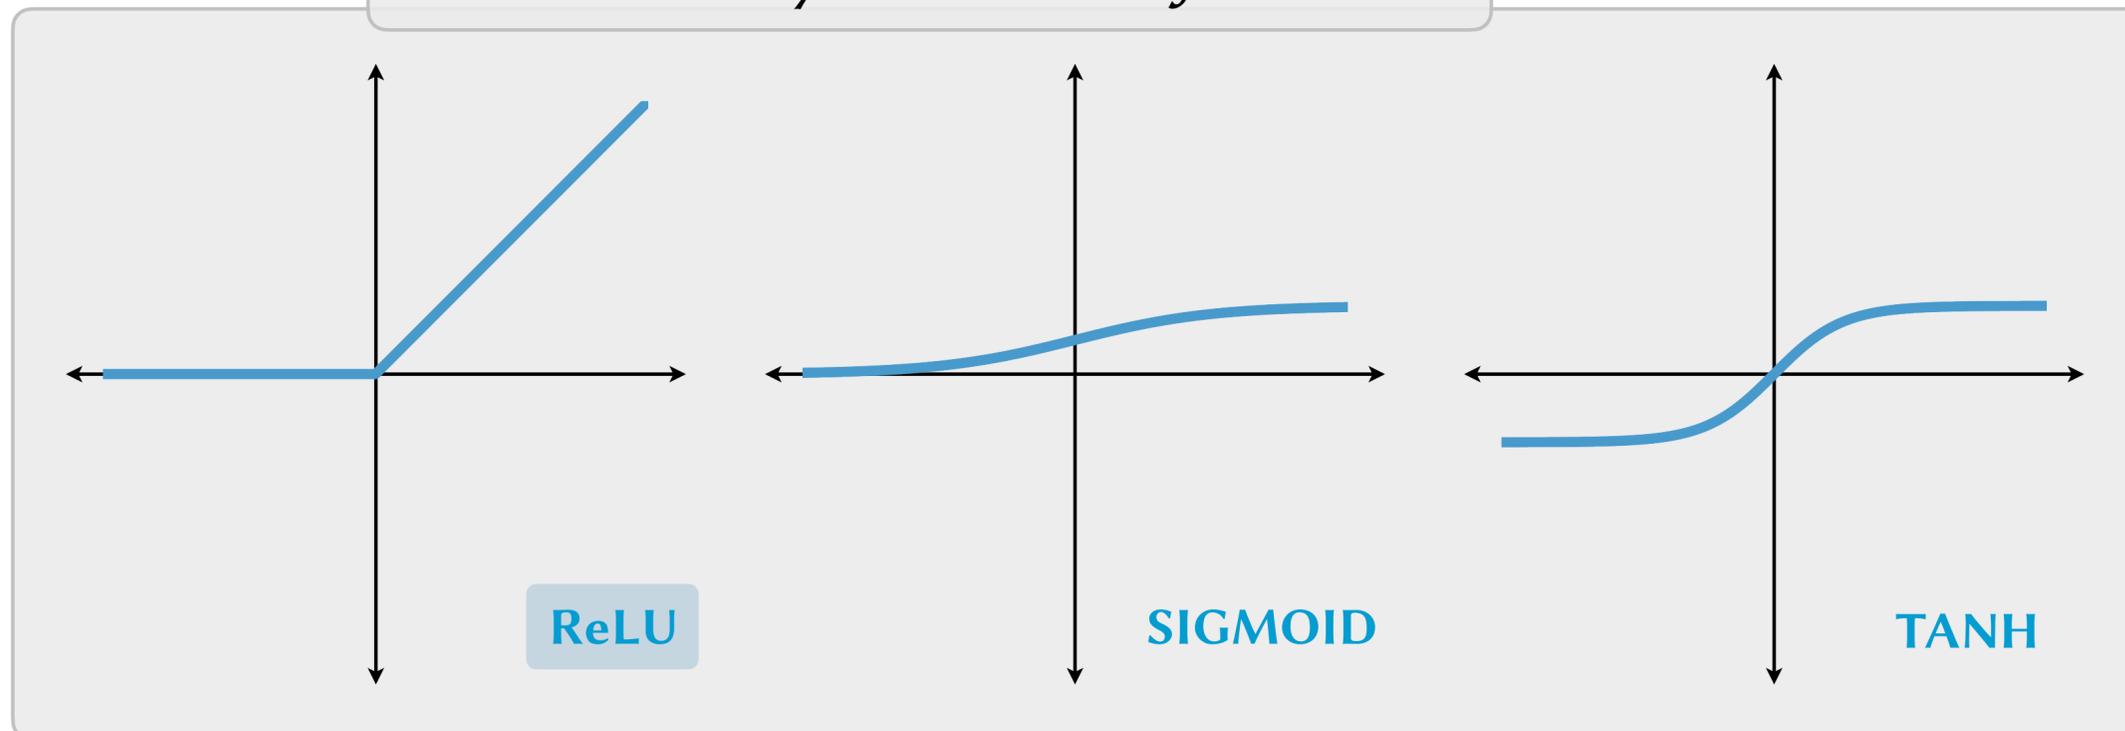
*linear layer*

# Neural networks are a function space

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

$$f_{\theta}(x) = A_{\theta}^N \varphi(A_{\theta}^1 \varphi(A_{\theta}^0 x + b_{\theta}^0) + b_{\theta}^1 + \dots) + b_{\theta}^N$$

*non-linearity/activation function*



📖 For an intro to Neural Networks, see the course “[6.036 Intro to ML](#)” on OCW

# Neural networks are a function space

*MLP*

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

*hidden layers*

$$f_{\theta}(x) = A_{\theta}^N \varphi(A_{\theta}^1 \varphi(A_{\theta}^0 x + b_{\theta}^0) + b_{\theta}^1 + \dots) + b_{\theta}^N$$

*output layer*

*input layer*

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

$$f_{\theta}(x) = A_{\theta}^N \varphi(A_{\theta}^1 \varphi(A_{\theta}^0 x + b_{\theta}^0) + b_{\theta}^1 + \dots) + b_{\theta}^N$$

**Key idea:** neural networks provide a space of functions with parameters  $\theta$

This CVPR paper is the Open Access version, provided by the Computer Vision Foundation. Except for this watermark, it is identical to the accepted version; the final published version of the proceedings is available on IEEE Xplore.

## Occupancy Networks: Learning 3D Reconstruction in Function Space

Lars Mescheder<sup>1</sup> Michael Oechsle<sup>1,2</sup> Michael Niemeyer<sup>1</sup> Sebastian Nowozin<sup>3†</sup> Andreas Geiger<sup>1</sup>  
<sup>1</sup>Autonomous Vision Group, MPI for Intelligent Systems and University of Tübingen  
<sup>2</sup>ETAS GmbH, Stuttgart  
<sup>3</sup>Google AI Berlin

{firstname.lastname}@tue.mpg.de nowozin@gmail.com

### Abstract

With the advent of deep neural networks, learning-based approaches for 3D reconstruction have gained popularity. However, unlike for images, in 3D there is no canonical representation which is both computationally and memory efficient yet allows for representing high-resolution geometry of arbitrary topology. Many of the state-of-the-art learning-based 3D reconstruction approaches can hence only represent very coarse 3D geometry or are limited to a restricted domain. In this paper, we propose *Occupancy Networks*, a new representation for learning-based 3D reconstruction methods. *Occupancy networks* implicitly represent the 3D surface as the continuous decision boundary of a deep neural network classifier. In contrast to existing approaches, our representation encodes a description of the 3D output at infinite resolution without excessive memory footprint. We validate that our representation can efficiently encode 3D structure and can be inferred from various kinds of input. Our experiments demonstrate competitive results, both qualitatively and quantitatively, for the challenging tasks of 3D reconstruction from single images, noisy point clouds and coarse discrete voxel grids. We believe that occupancy networks will become a useful tool in a wide variety of learning-based 3D tasks.

### 1. Introduction

Recently, learning-based approaches for 3D reconstruction have gained popularity [4, 9, 23, 58, 75, 77]. In contrast to traditional multi-view stereo algorithms, learned models are able to encode rich prior information about the space of 3D shapes which helps to resolve ambiguities in the input. While generative models have recently achieved remarkable successes in generating realistic high resolution images [36, 47, 72], this success has not yet been replicated in the 3D domain. In contrast to the 2D domain, the com-

<sup>†</sup>Part of this work was done while at MSR Cambridge.

4460

## DeepSDF: Learning Continuous Signed Distance Functions for Shape Representation

Jeong Joon Park<sup>1,2†</sup> Peter Florence<sup>2,3†</sup> Julian Straub<sup>3</sup> Richard Newcombe<sup>3</sup> Steven Lovegrove<sup>3</sup>  
<sup>1</sup>University of Washington <sup>2</sup>Massachusetts Institute of Technology <sup>3</sup>Facebook Reality Labs



### Abstract

Computer graphics, 3D computer vision and robotics communities have produced multiple approaches to representing 3D geometry for rendering and reconstruction. These provide trade-offs across fidelity, efficiency and compression capabilities. In this work, we introduce *DeepSDF*, a learned continuous Signed Distance Function (SDF) representation of a class of shapes that enables high quality shape representation, interpolation and completion from partial and noisy 3D input data. *DeepSDF*, like its classical counterpart, represents a shape's surface by a continuous volumetric field: the magnitude of a point in the field represents the distance to the surface boundary and the sign indicates whether the region is inside (-) or outside (+) of the shape, hence our representation implicitly encodes a shape's boundary as the zero-level-set of the learned function while explicitly representing the classification of space as being part of the shapes interior or not. While classical SDF's both in analytical or discretized voxel form typically represent the surface of a single shape, *DeepSDF* can represent an entire class of shapes. Furthermore, we show state-of-the-art performance for learned 3D shape representation and completion while reducing the model size by an order of magnitude compared with previous work.

### 1. Introduction

Deep convolutional networks which are a mainstay of image-based approaches grow quickly in space and time complexity when directly generalized to the 3D spatial dimension, and more classical and compact surface representations such as triangle or quad meshes pose problems in training since we may need to deal with an unknown number of vertices and arbitrary topology. These challenges have limited the quality, flexibility and fidelity of deep learning approaches when attempting to either input 3D data for processing or produce 3D inferences for object segmentation and reconstruction. In this work, we present a novel representation and approach for generative 3D modeling that is efficient, expressive, and fully continuous. Our approach uses the concept of a SDF, but unlike common surface reconstruction techniques which discretize this SDF into a regular grid for evaluation and measurement denoising, we instead learn a generative model to produce such a continuous field. The proposed continuous representation may be intuitively understood as a learned shape-conditioned classifier for which the decision boundary is the surface of the shape itself, as shown in Fig. 2. Our approach shares the generative aspect of other works seeking to map a latent space to a distribution of complex shapes in 3D [54], but critically differs in the central representation. While the notion of an

<sup>†</sup>Work performed during internship at Facebook Reality Labs.

arXiv:1901.05103v1 [cs.CV] 16 Jan 2019

## Learning Implicit Fields for Generative Shape Modeling

Zhiqin Chen  
Simon Fraser University  
zhiqinc@sfu.ca

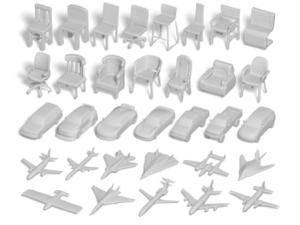
Hao Zhang  
Simon Fraser University  
haoz@sfu.ca

### Abstract

We advocate the use of implicit fields for learning generative models of shapes and introduce an implicit field decoder, called IM-NET, for shape generation, aimed at improving the visual quality of the generated shapes. An implicit field assigns a value to each point in 3D space, so that a shape can be extracted as an iso-surface. IM-NET is trained to perform this assignment by means of a binary classifier. Specifically, it takes a point coordinate, along with a feature vector encoding a shape, and outputs a value which indicates whether the point is outside the shape or not. By replacing conventional decoders by our implicit decoder for representation learning (via IM-AE) and shape generation (via IM-GAN), we demonstrate superior results for tasks such as generative shape modeling, interpolation, and single-view 3D reconstruction, particularly in terms of visual quality. Code and supplementary material are available at <https://github.com/czq142857/implicit-decoder>.

### 1. Introduction

Unlike images and video, 3D shapes are not confined to one standard representation. Up to date, deep neural networks for 3D shape analysis and synthesis have been developed for voxel grids [19, 48], multi-view images [42], point clouds [1, 35], and integrated surface patches [17]. Specific to generative modeling of 3D shapes, despite the many progresses made, the shapes produced by state-of-the-art methods still fall far short in terms of visual quality. This is reflected by a combination of issues including low-resolution outputs, overly smoothed or discontinuous surfaces, as well as a variety of topological noise and irregularities. In this paper, we explore the use of implicit fields for learning deep models of shapes and introduce an implicit field decoder for shape generation, aimed at improving the visual quality of the generated models, as shown in Figure 1. An implicit field assigns a value to each point  $(x, y, z)$ . A shape is represented by all points assigned to a specific value and is typically rendered via iso-surface extraction such as Marching Cubes. Our implicit field decoder, or simply implicit decoder, is trained to perform this assignment task, by means of a binary classifier, and it has a very simple architecture; see Figure 2. Specifically, it takes a point coordinate  $(x, y, z)$ , along with a feature vector encoding a shape, and outputs a value which indicates whether the point is outside the shape or not. In a typical application setup, our decoder, which is coined *IM-NET*, would follow an encoder which outputs the shape feature vectors and then return an implicit field to define an output shape. Several novel features of IM-NET impact the visual quality of the generated shapes. First, the decoder output can be sampled at any resolution and is not limited by the resolution of the training shapes; see Figure 1. More importantly, we concatenate point coordinates with shape features, feeding both as input to our implicit decoder, which learns the inside/outside status of any point relative to a shape. In contrast, a classical convolution/deconvolution-based neural network (CNN) operating on voxelized shapes is typically trained to predict voxels relative to the extent of the bounding volume of a shape. Such a network learns



arXiv:1812.02822v5 [cs.GR] 16 Sep 2019

## Scene Representation Networks: Continuous 3D-Structure-Aware Neural Scene Representations

Vincent Sitzmann<sup>1</sup> Michael Zollhöfer<sup>1</sup> Gordon Wetzstein<sup>1</sup>  
{sitzmann, zollhoefer}@cs.stanford.edu, gordon.wetzstein@stanford.edu  
Stanford University  
vsitzmann.github.io/srnns/

### Abstract

Unsupervised learning with generative models has the potential of discovering rich representations of 3D scenes. While geometric deep learning has explored 3D-structure-aware representations of scene geometry, these models typically require explicit 3D supervision. Emerging neural scene representations can be trained only with posed 2D images, but existing methods ignore the three-dimensional structure of scenes. We propose *Scene Representation Networks (SRNs)*, a continuous, 3D-structure-aware scene representation that encodes both geometry and appearance. SRNs represent scenes as continuous functions that map world coordinates to a feature representation of local scene properties. By formulating the image formation as a differentiable ray-marching algorithm, SRNs can be trained end-to-end from only 2D images and their camera poses, without access to depth or shape. This formulation naturally generalizes across scenes, learning powerful geometry and appearance priors in the process. We demonstrate the potential of SRNs by evaluating them for novel view synthesis, few-shot reconstruction, joint shape and appearance interpolation, and unsupervised discovery of a non-rigid face model.<sup>1</sup>

### 1 Introduction

A major driver behind recent work on generative models has been the promise of unsupervised discovery of powerful neural scene representations, enabling downstream tasks ranging from robotic manipulation and few-shot 3D reconstruction to navigation. A key aspect of solving these tasks is understanding the three-dimensional structure of an environment. However, prior work on neural scene representations either does not or only weakly enforces 3D structure [1–4]. Multi-view geometry and projection operations are performed by a black-box neural renderer, which is expected to learn these operations from data. As a result, such approaches fail to discover 3D structure under limited training data (see Sec. 4), lack guarantees on multi-view consistency of the rendered images, and learned representations are generally not interpretable. Furthermore, these approaches lack an intuitive interface to multi-view and projective geometry important in computer graphics, and cannot easily generalize to camera intrinsic matrices and transformations that were completely unseen at training time. In geometric deep learning, many classic 3D scene representations, such as voxel grids [5–10], point clouds [11–14], or meshes [15] have been integrated with end-to-end deep learning models and have led to significant progress in 3D scene understanding. However, these scene representations are discrete, limiting achievable spatial resolution, only sparsely sampling the underlying smooth surfaces of a scene, and often require explicit 3D supervision.

<sup>1</sup>Please see supplemental video for additional results.

33rd Conference on Neural Information Processing Systems (NeurIPS 2019), Vancouver, Canada.

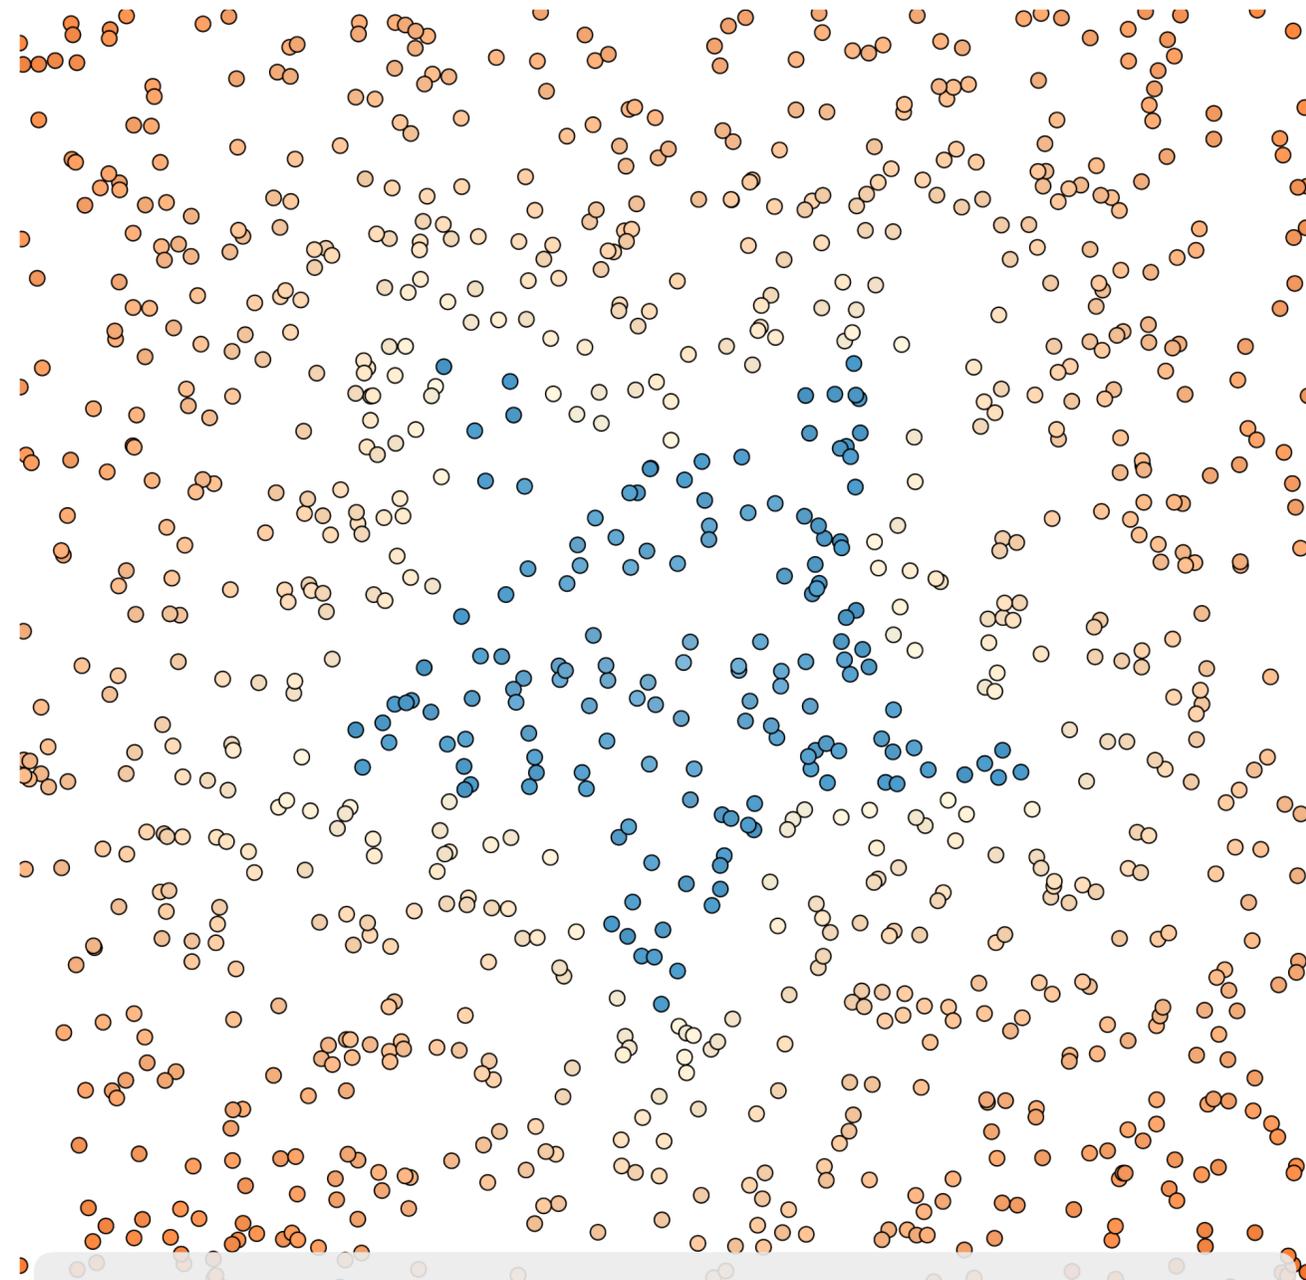
arXiv:1906.01618v2 [cs.CV] 28 Jan 2020

Occupancy Networks  
Mescheder et al., 2019

DeepSDF  
Park et al., 2019

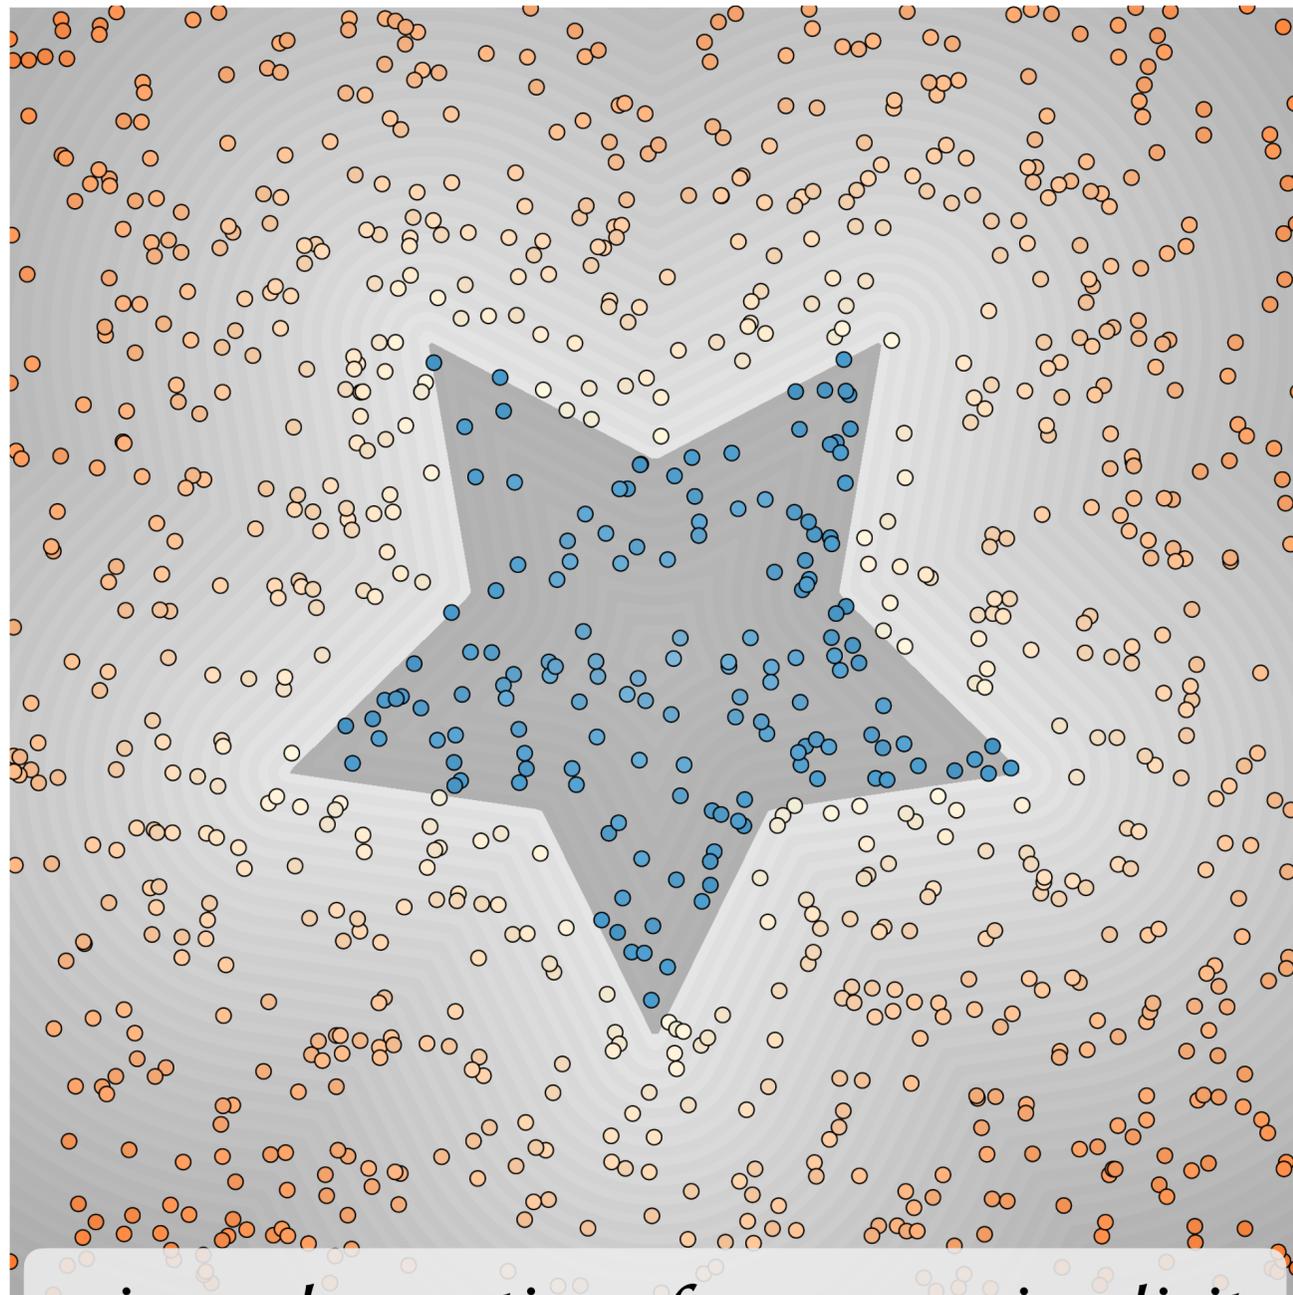
IM-NET  
Chen et al., 2019

Scene Rep. Nets.  
Sitzmann et al., 2019



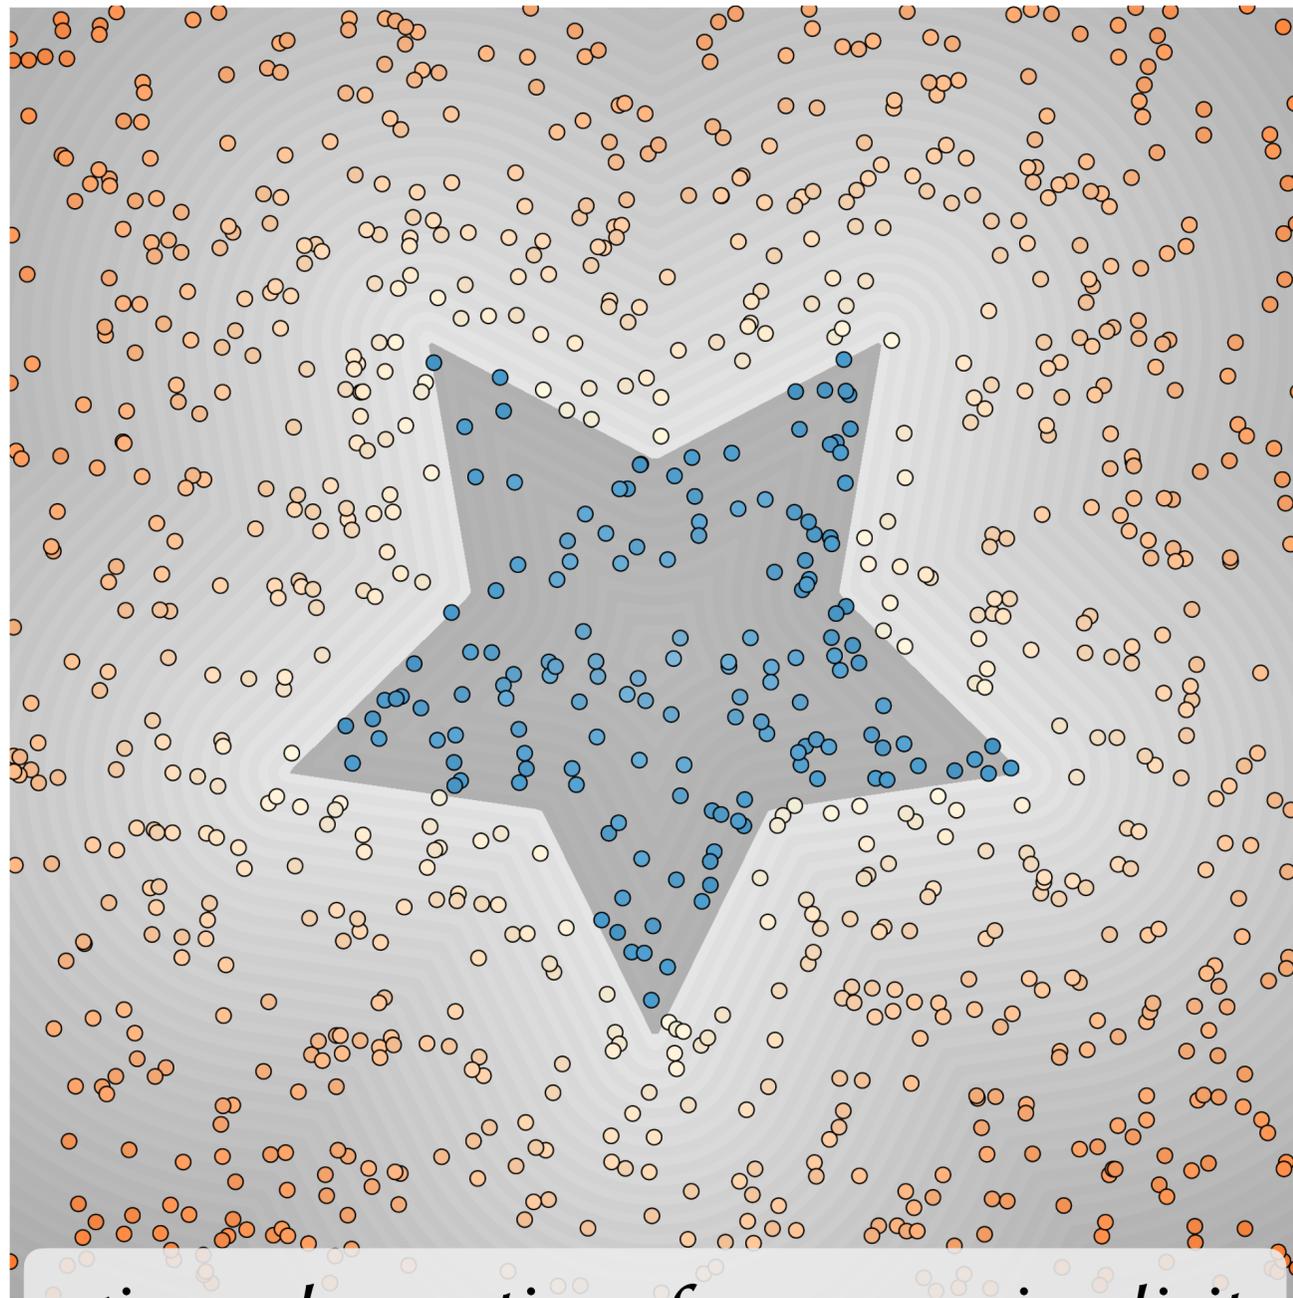
*given observations from some implicit*

given:  $(\mathcal{X}, \mathcal{V})$  with  $f(x) = v$



*given observations from some implicit*

given:  $(\mathcal{X}, \mathcal{V})$  with  $f(x) = v$



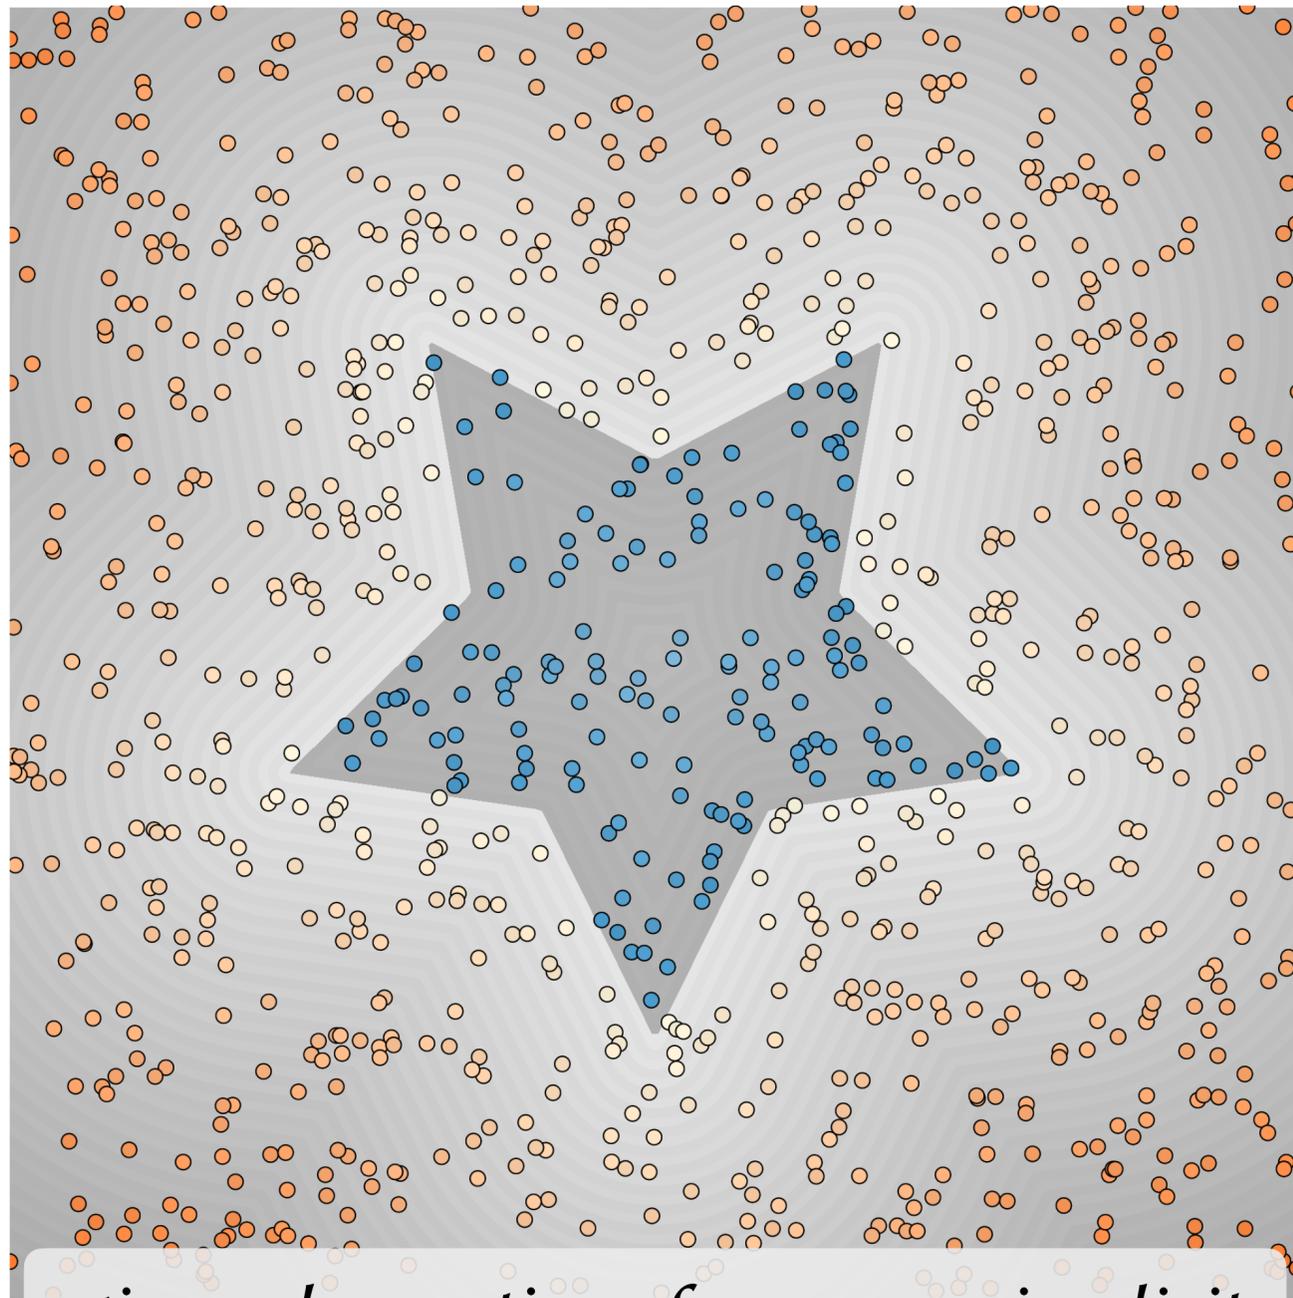
*given observations from some implicit*

given:  $(\mathcal{X}, \mathcal{V})$  with  $f(x) = v$

$$\min_{\theta} \sum_{x \in \mathcal{X}} (f_{\theta}(x) - v)^2$$

— “loss function”

*find best fitting  $f$  in func. space*



*given observations from some implicit*

given:  $(\mathcal{X}, \mathcal{V})$  with  $f(x) = v$

$$\min_{\theta} \sum_{x \in \mathcal{X}} (f_{\theta}(x) - v)^2$$

*optimized using stochastic  
gradient descent*

*find best fitting  $f$  in func. space*

# In practice: SDF fit with pytorch

```
class SDF_NN(nn.Module):
    def __init__(self, num_layers, hidden_size):
        super().__init__()

        self.layers = nn.ModuleList([nn.Linear(2, hidden_size), nn.ReLU()])

        for _ in range(num_layers-2):
            self.layers.append(nn.Linear(hidden_size, hidden_size))
            self.layers.append(nn.ReLU())

        self.layers.append(nn.Linear(hidden_size, 1))

    def forward(self, inp):
        out = inp

        for layer in self.layers:
            out = layer(out)

        return out
```

```
def train(model, data_gen, loss_fn, steps=10000, step_size=0.001):
    optim = torch.optim.Adam(model.parameters(), lr=step_size)

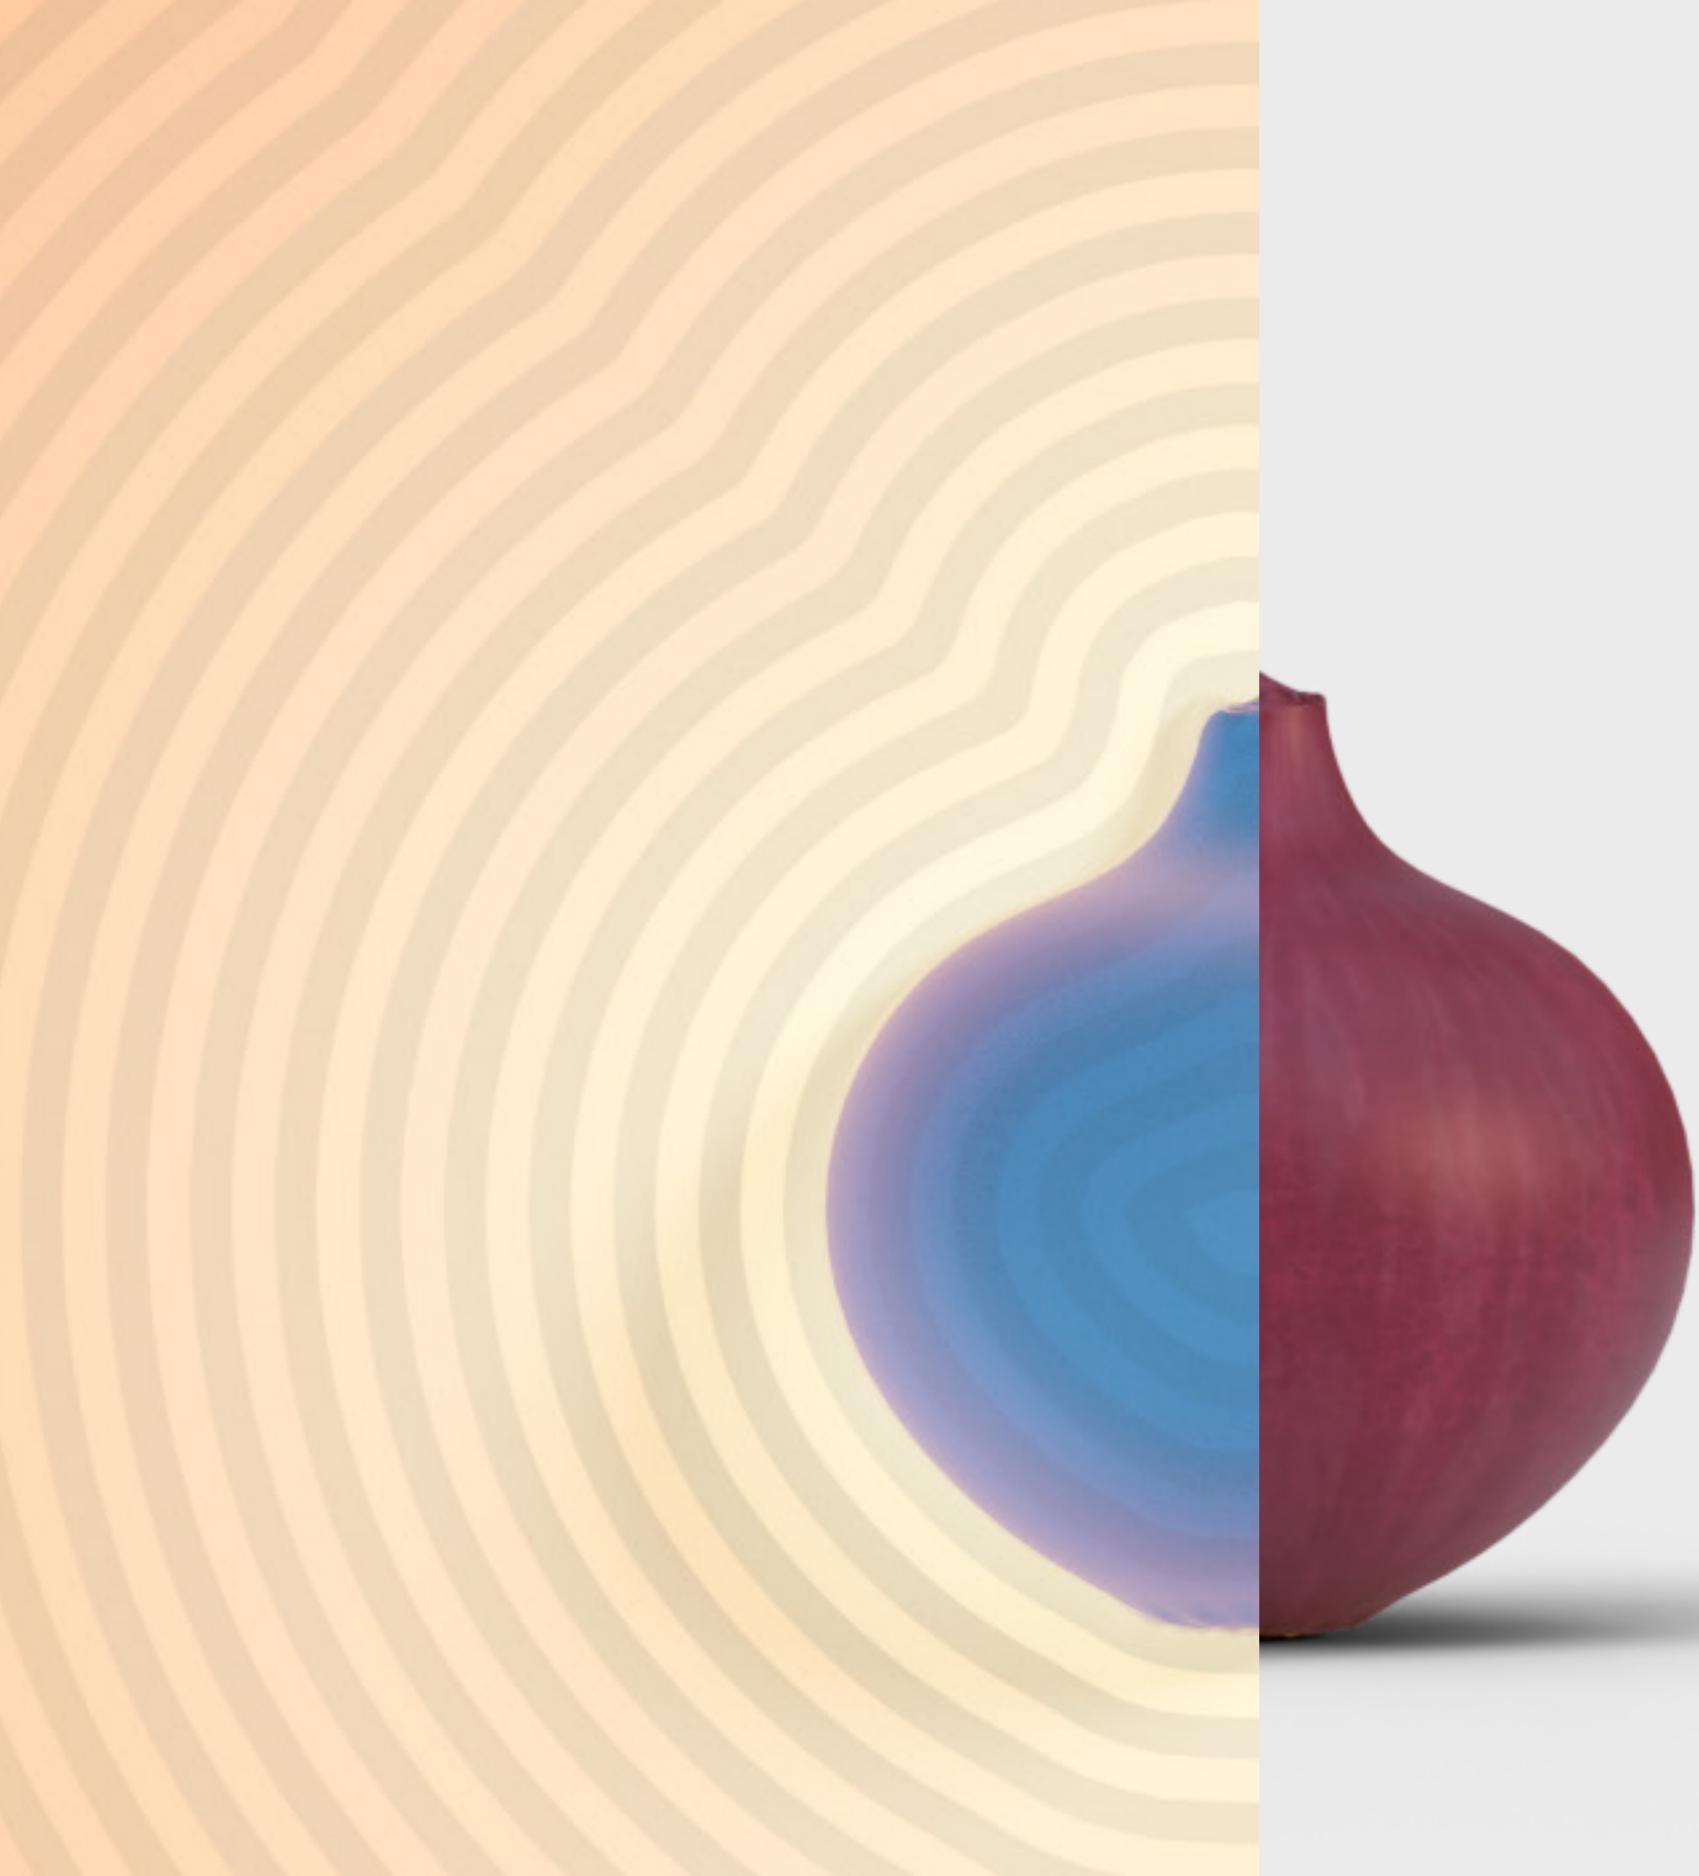
    for i in tqdm(range(steps)):
        pts, gts = data_gen()

        pred = model(pts)
        loss = loss_fn(pred, gts)

        loss.backward()
        optim.step()
        optim.zero_grad()

        if i % 100 == 0:
            tqdm.write(f'Step: {i}, loss: {loss}')

    torch.save(model.state_dict(), './trained_model.pt')
```



## OVERVIEW

### IMPLICIT SURFACES

1. Background & History
2. Digital Representations

### NEURAL IMPLICIT

3. The Basics
4. **Why & Why Not Use Them?**
5. Working out the Details
6. Geometric Operations

### CONCLUSION

**Question:** is this a good function space?

## WHY?

1. **Representative ability**
2. Adaptability
3. Trivially differentiable

**Question:** is this a good function space?

## Universal Approximation Theorem [Hornik et al., 1989]

A MLP with one hidden layer of infinite width and ReLU nonlinearities can represent any continuous function



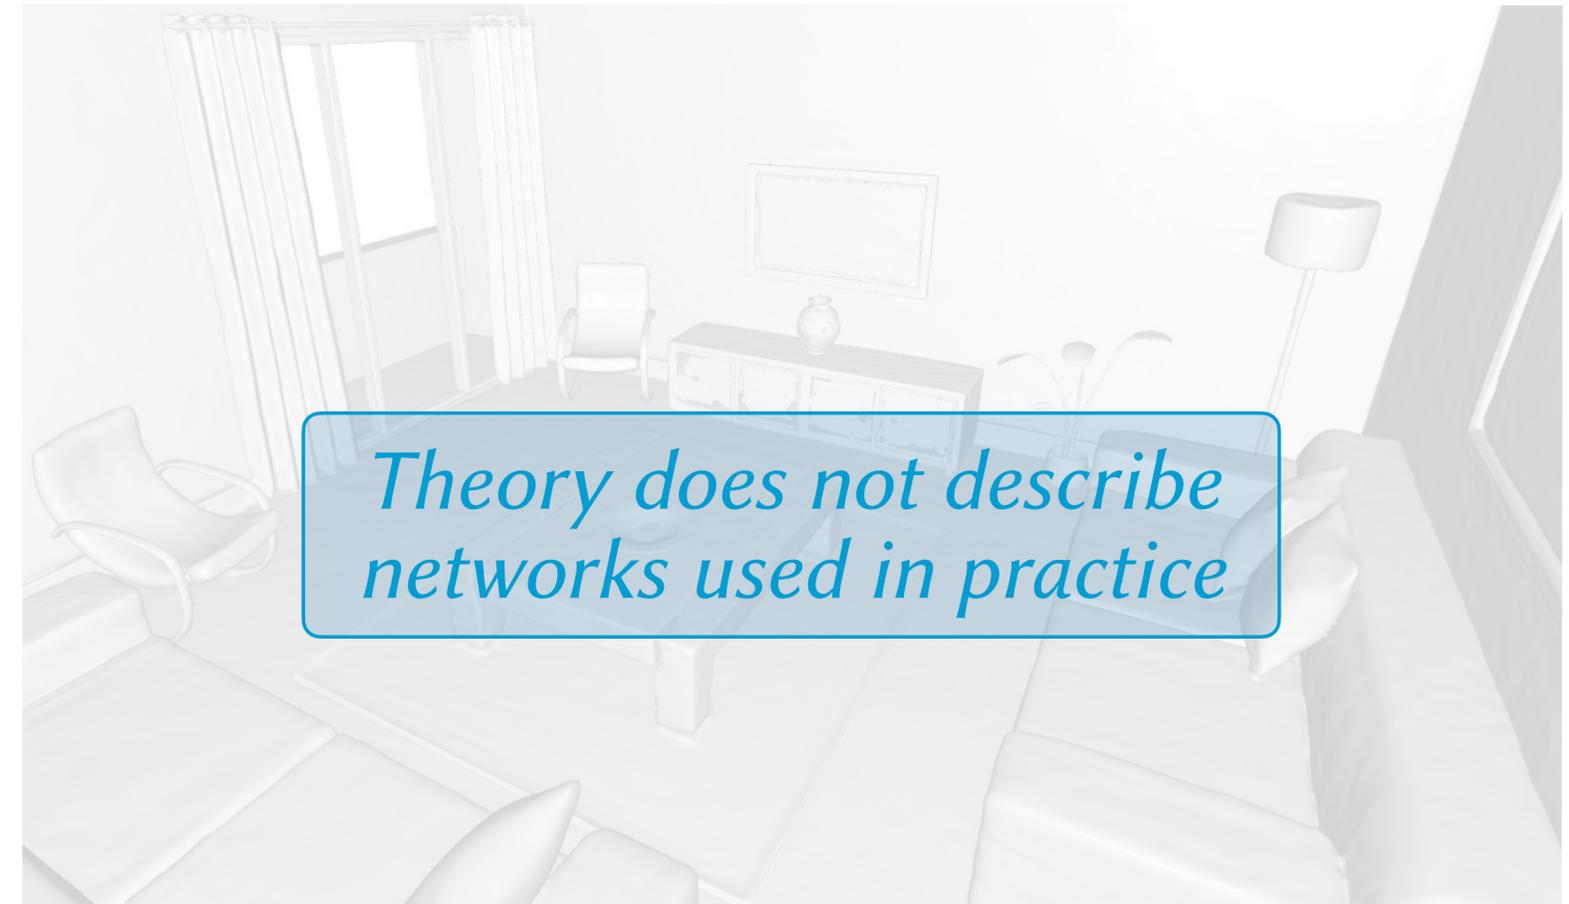
### WHY?

1. **Representative ability**
2. Adaptability
3. Trivially differentiable

**Question:** is this a good function space?

## Universal Approximation Theorem [Hornik et al., 1989]

A MLP with one hidden layer of **infinite** width and ReLU nonlinearities can represent any continuous function

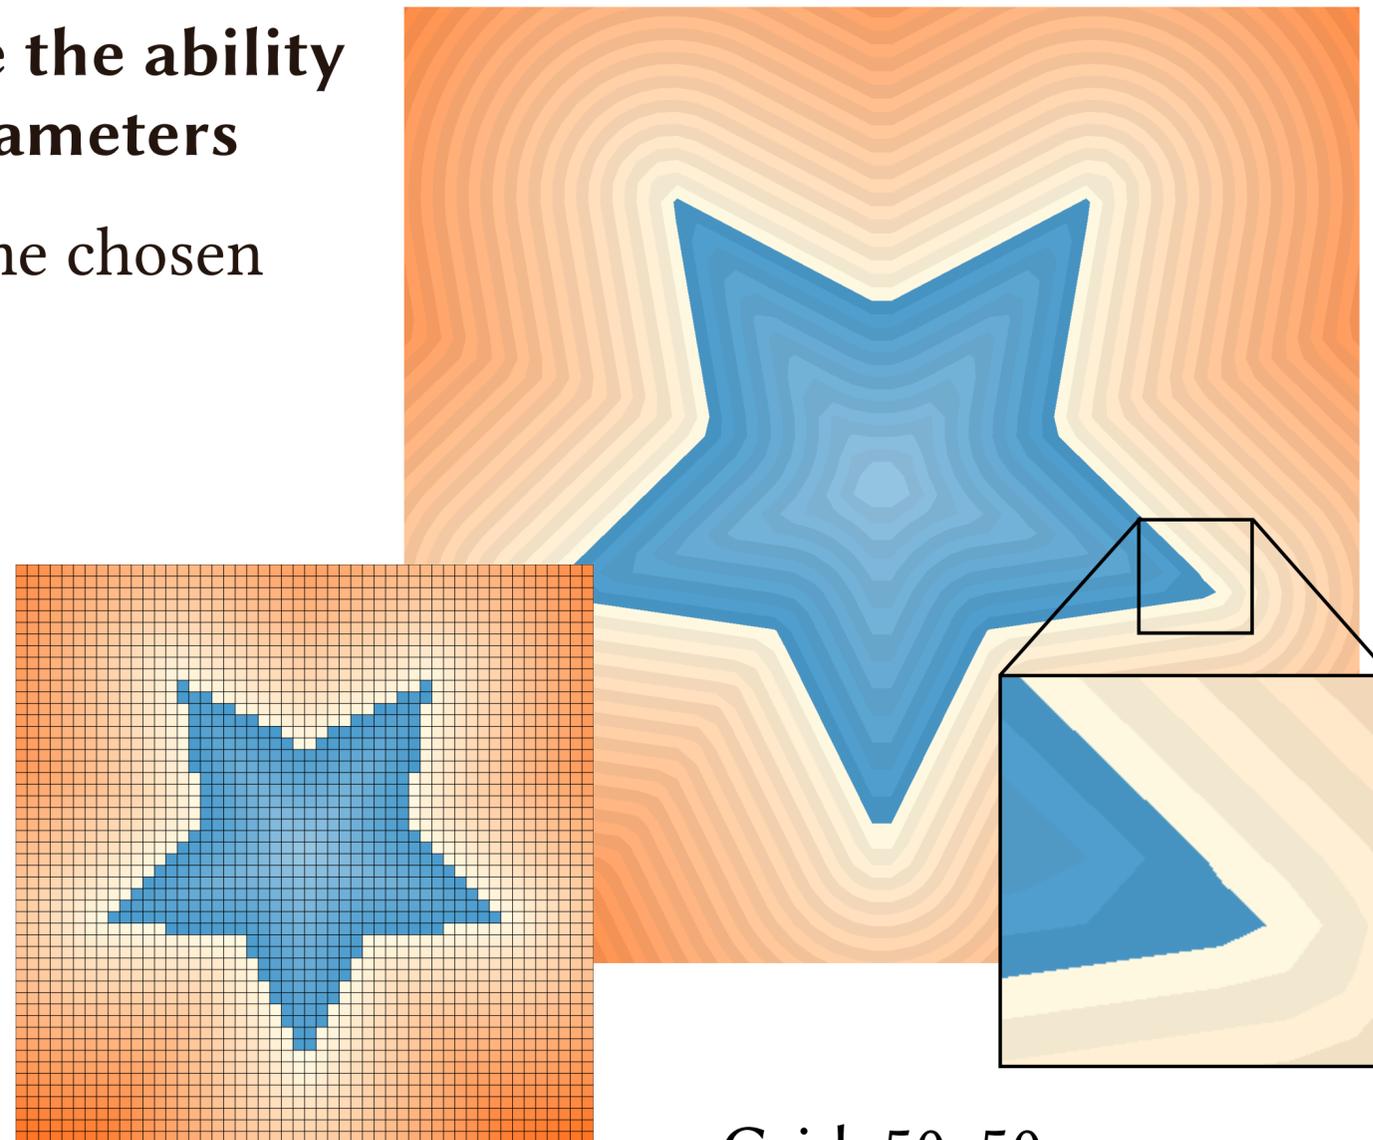


### WHY?

1. **Representative ability**
2. Adaptability
3. Trivially differentiable

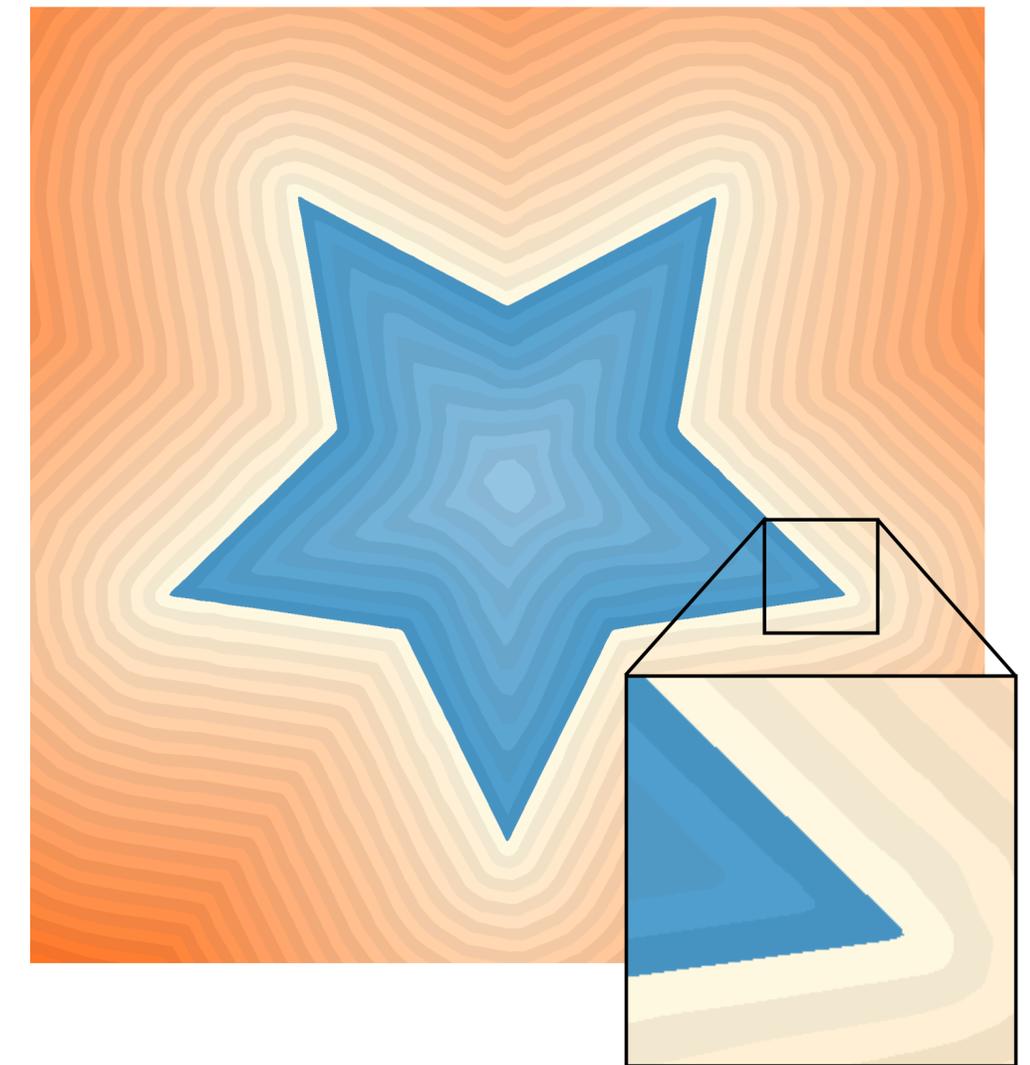
## Neural implicits have the ability to adaptively use parameters

- but this depends on the chosen loss function!



Grid: 50x50

**2500 parameters**



Neural Network: 4 layers, 34  
hidden size

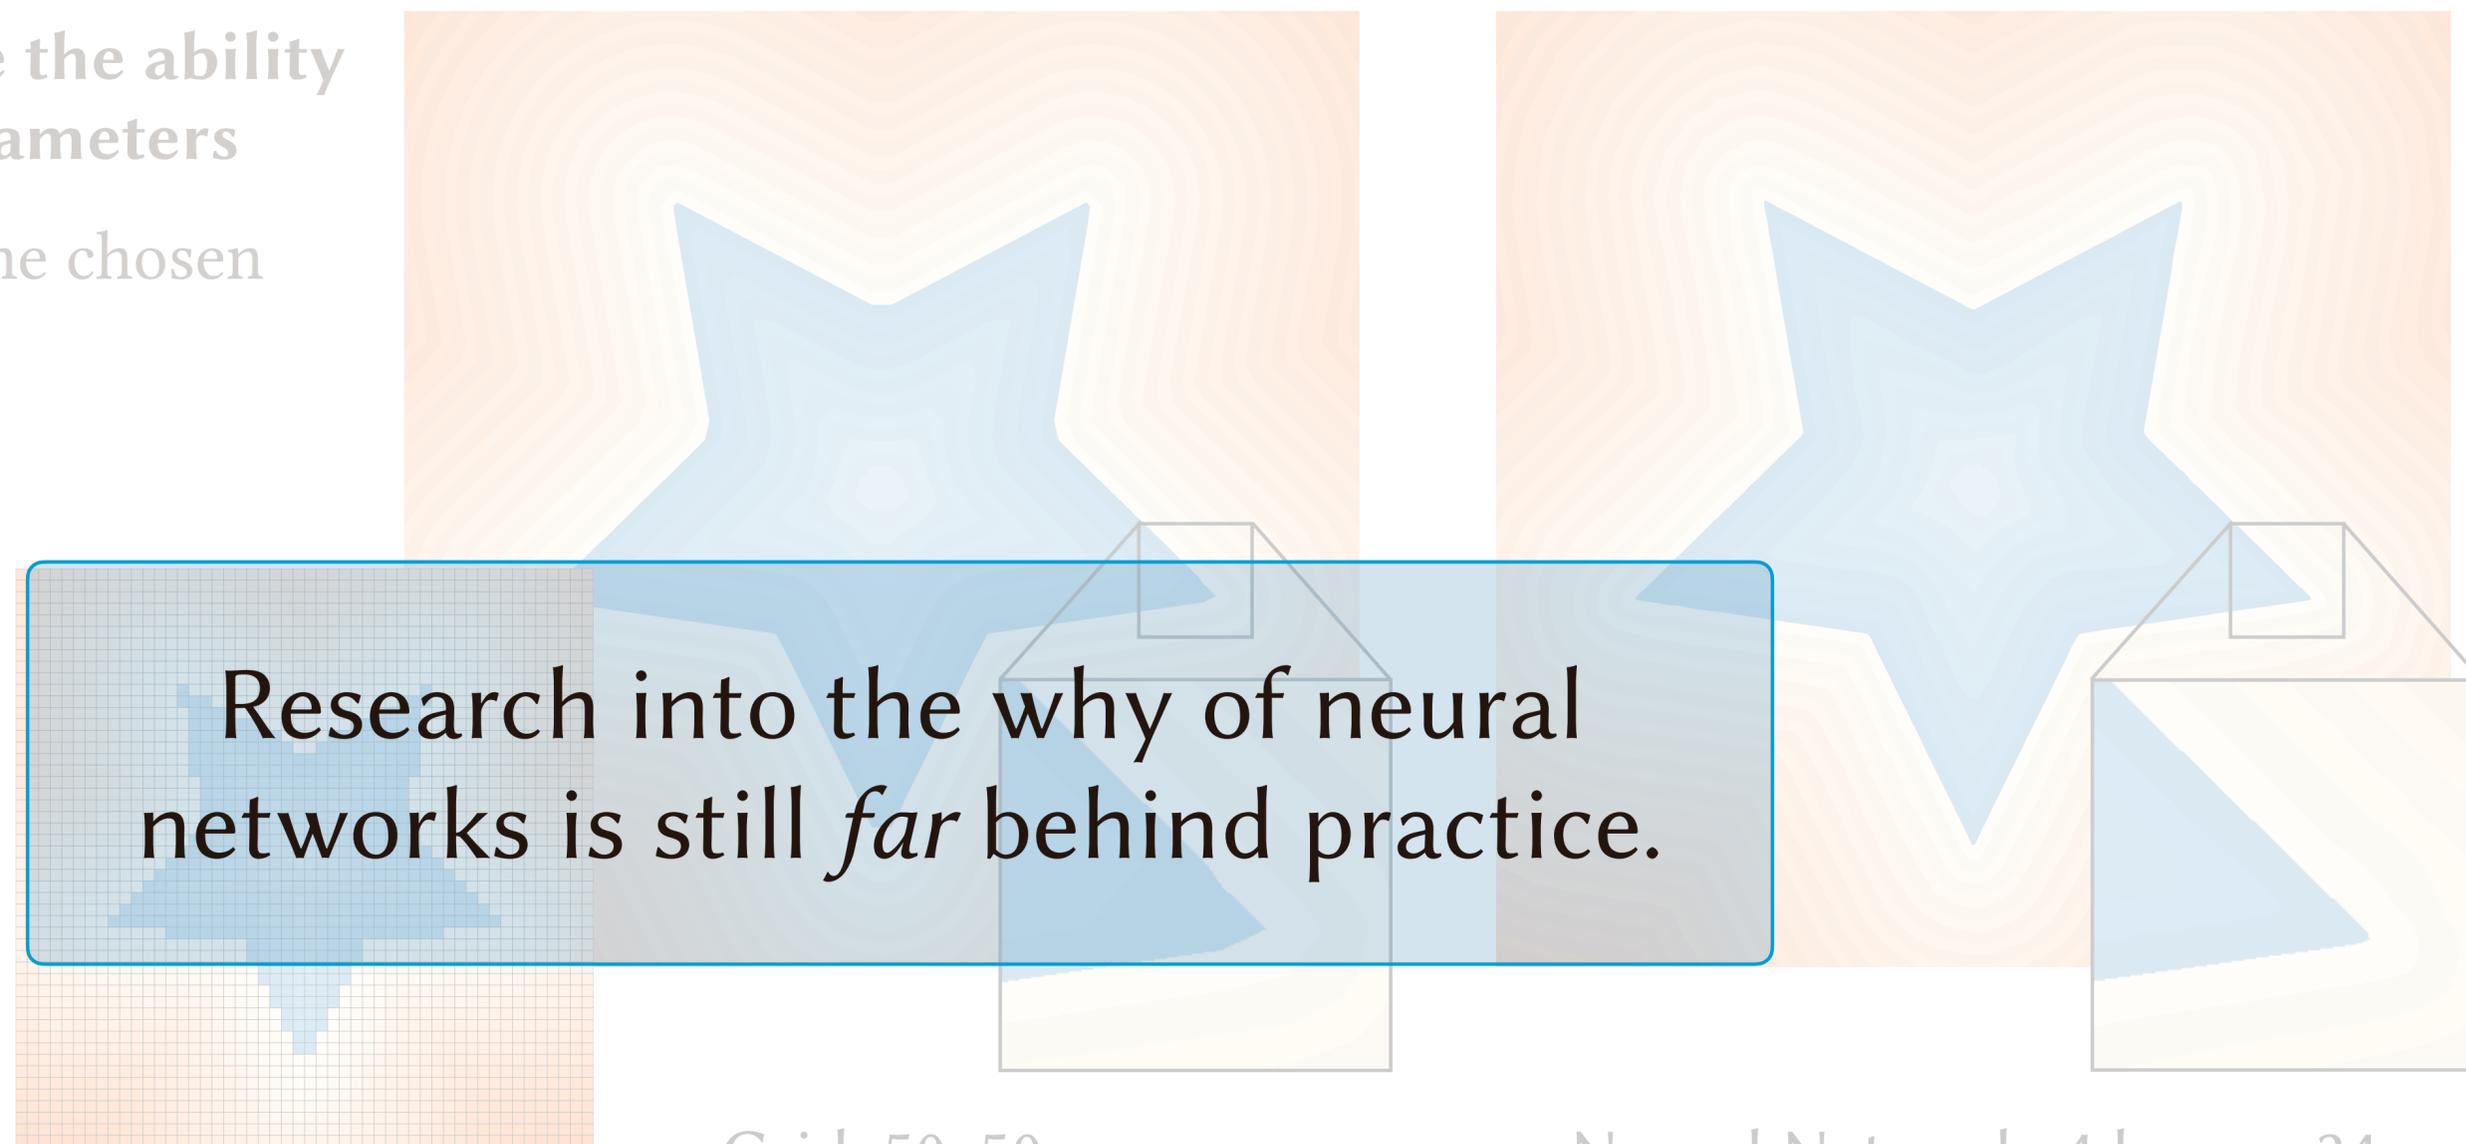
**2414 parameters**

## WHY?

1. Representative ability
2. **Adaptability**
3. Trivially differentiable

Neural implicit have the ability to adaptively use parameters

- but this depends on the chosen loss function!



Research into the why of neural networks is still *far* behind practice.

Grid: 50x50

2500 parameters

Neural Network: 4 layers, 34 hidden size

2414 parameters

WHY?

1. Representative ability
2. **Adaptability**
3. Trivially differentiable

**Neural implicit can easily be used to for many problems requiring a differentiable geometry representation**

## **WHY?**

1. Representative ability
2. Adaptability
3. ✨ **Trivially differentiable** ✨

Neural implicits can easily be used to for many problems requiring a differentiable geometry representation

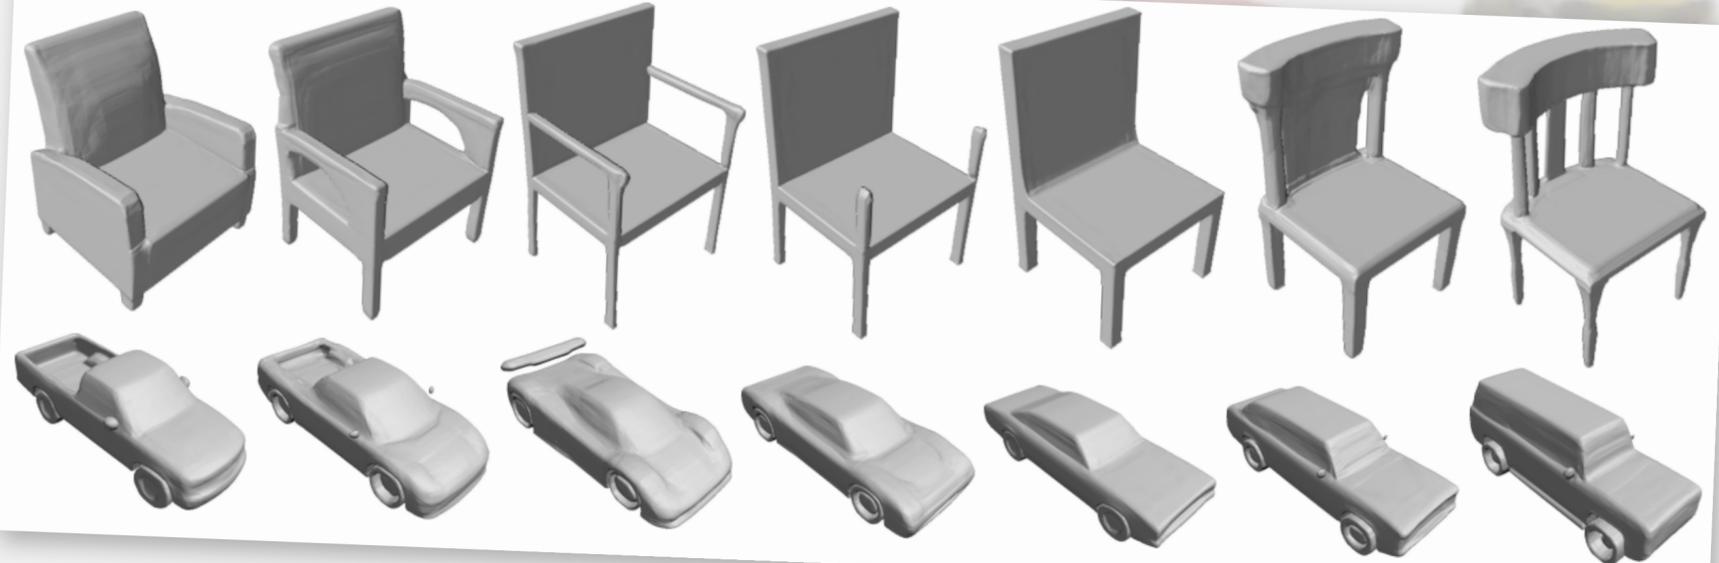


Single shot reconstruction [Sitzmann et al., 2019]

## WHY?

1. Representative ability
2. Adaptability
3. ✨ **Trivially differentiable** ✨

Neural implicits can easily be used to for many problems requiring a differentiable geometry representation

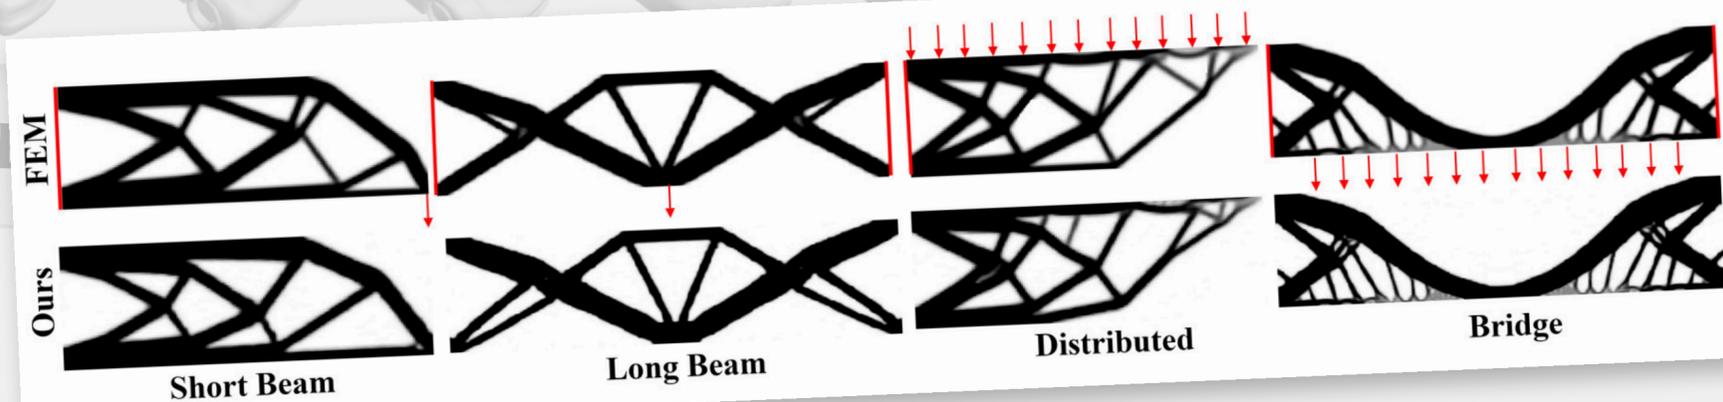
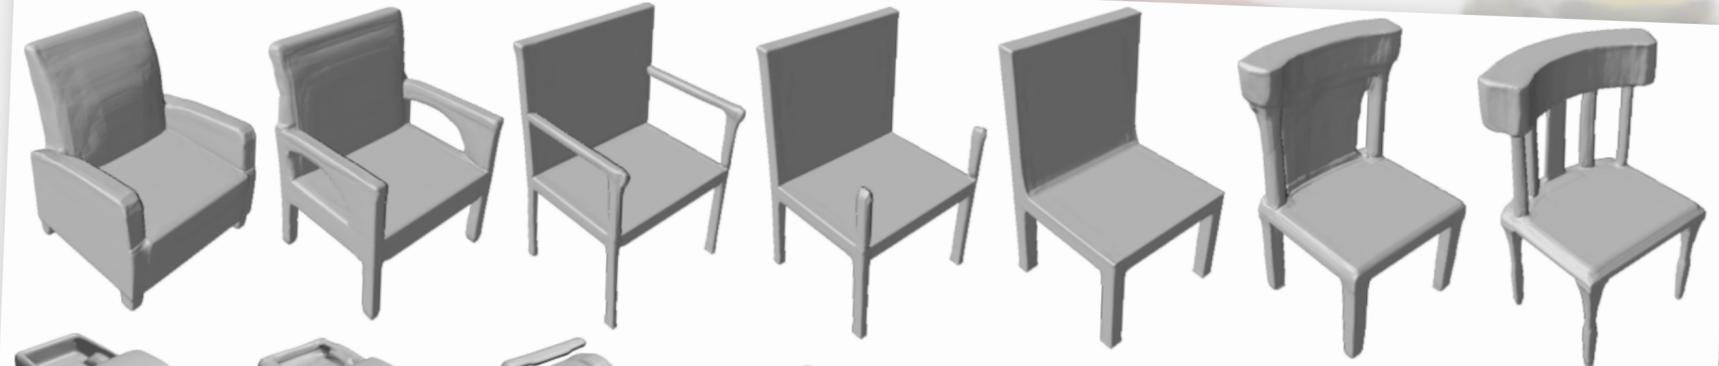


Latent space interpolation [Park et al., 2019]

## WHY?

1. Representative ability
2. Adaptability
3. ✨ Trivially differentiable ✨

Neural implicit can easily be used to for many problems requiring a differentiable geometry representation



Topology Optimization [Zehnder et al., 2021]

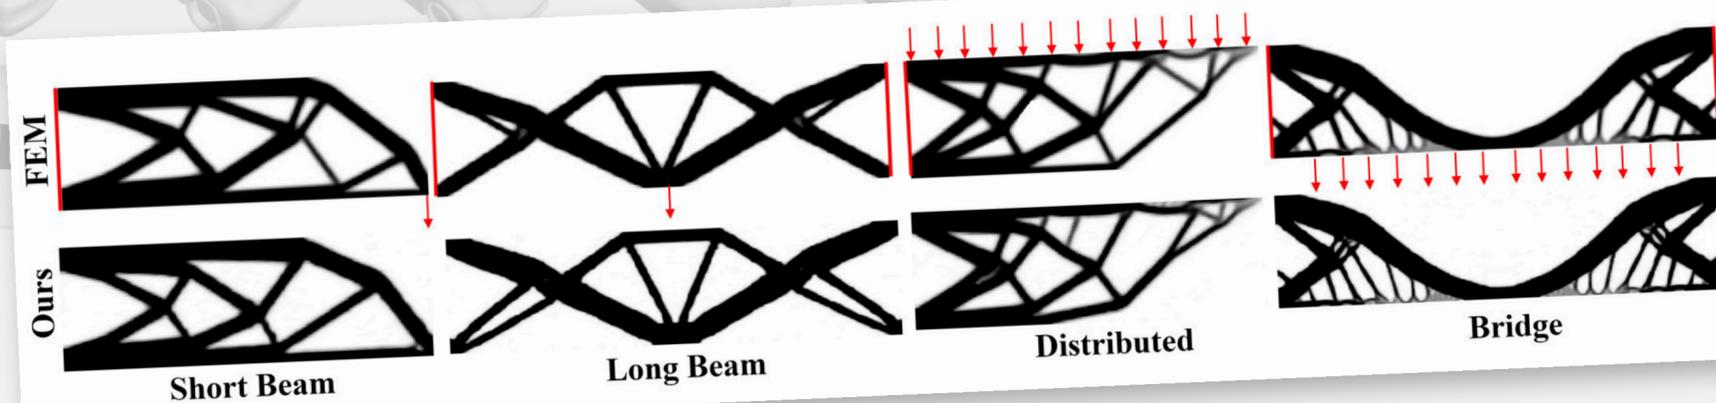
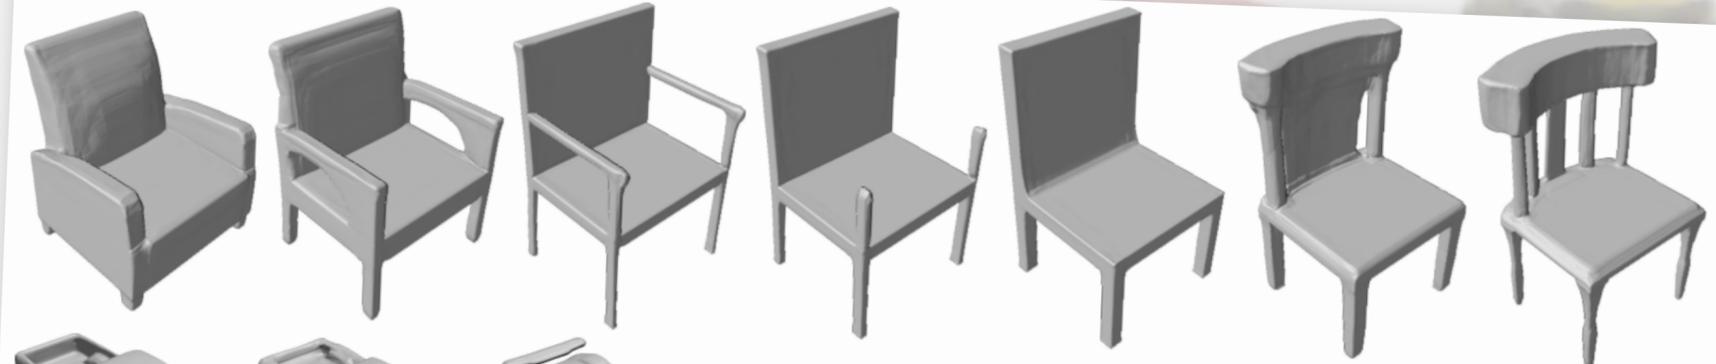
## WHY?

1. Representative ability
2. Adaptability
3. ✨ Trivially differentiable ✨

Neural implicit can easily be used to for many problems requiring a differentiable geometry representation

### In general:

- inverse tasks
- learning from data
- integrate with ML toolbox



Topology Optimization [Zehnder et al., 2021]

### WHY?

1. Representative ability
2. Adaptability
3. ✨ Trivially differentiable ✨

## Properties hold only through *soft losses*

**Loss function:** describes problem to be solved, e.g.,

- SDF property
- reconstruction loss
- ...

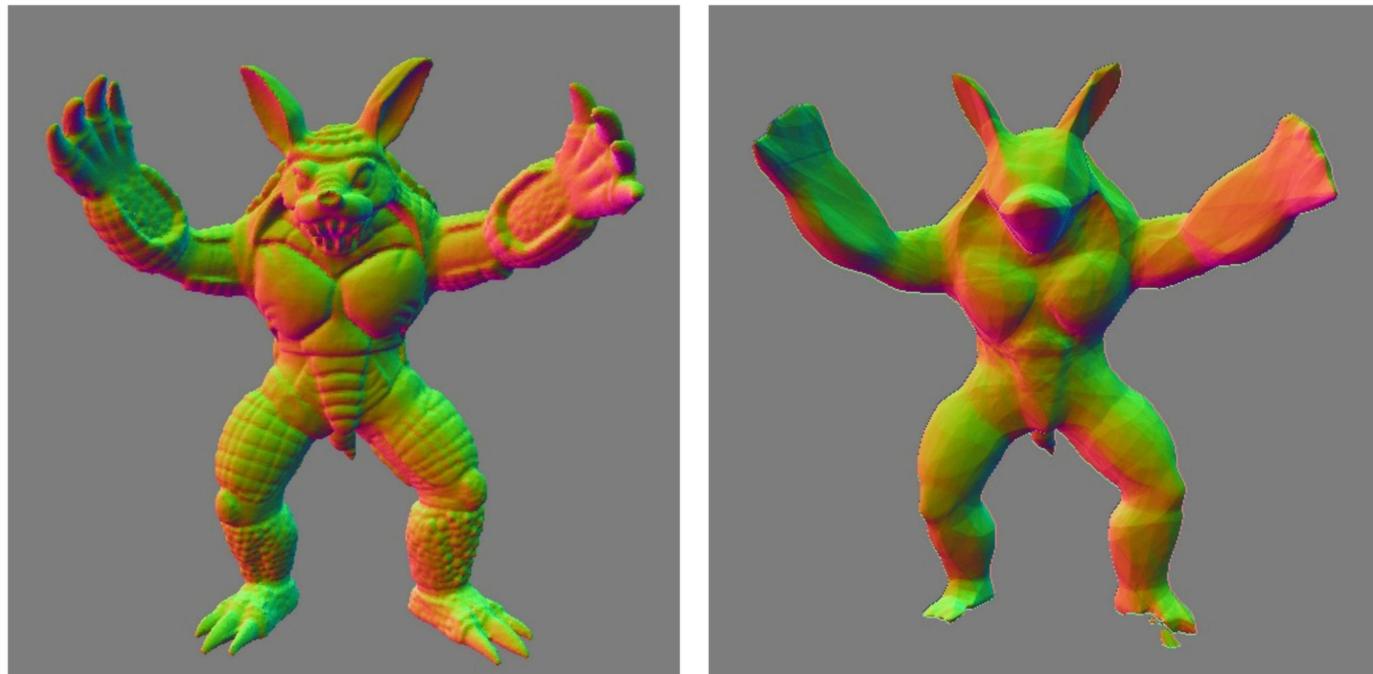
$$\min_{\theta} \mathcal{L}(f_{\theta})$$

### WHY NOT?

1. **Imprecision**
2. Editing (etc.) is difficult
3. The wrong representation

## Properties hold only through *soft losses*

- can be difficult to quantify amount of error
- issue for application which require guarantees!



**Loss function:** describes problem to be solved, e.g.,

- SDF property
- reconstruction loss
- ...

$$\min_{\theta} \mathcal{L}(f_{\theta})$$

## WHY NOT?

1. **Imprecision**
2. Editing (etc.) is difficult
3. The wrong representation

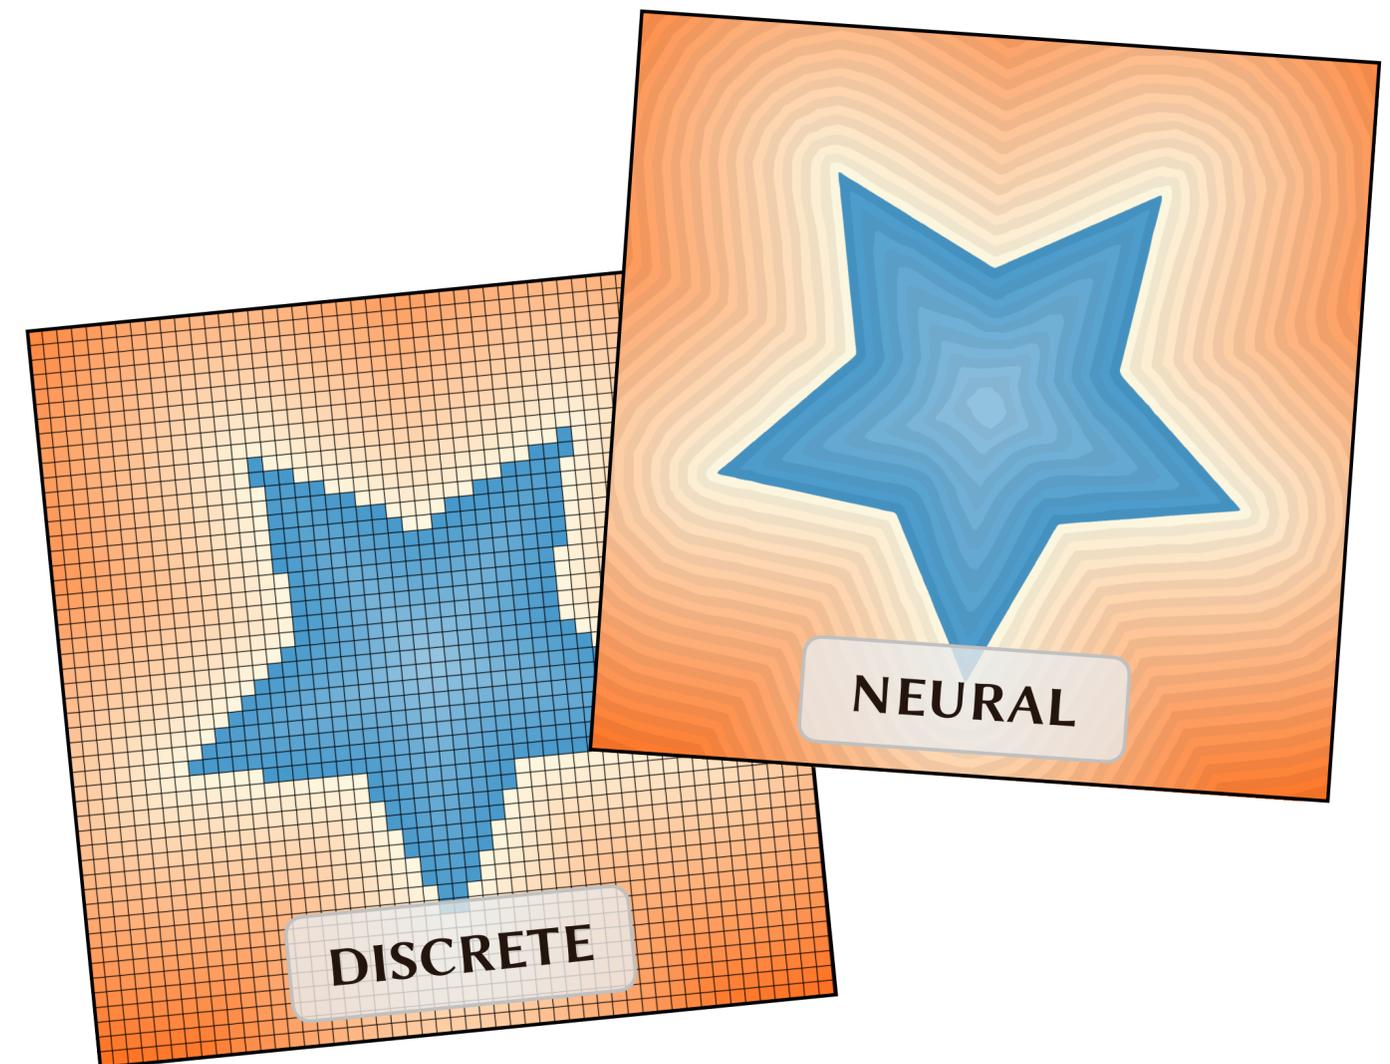
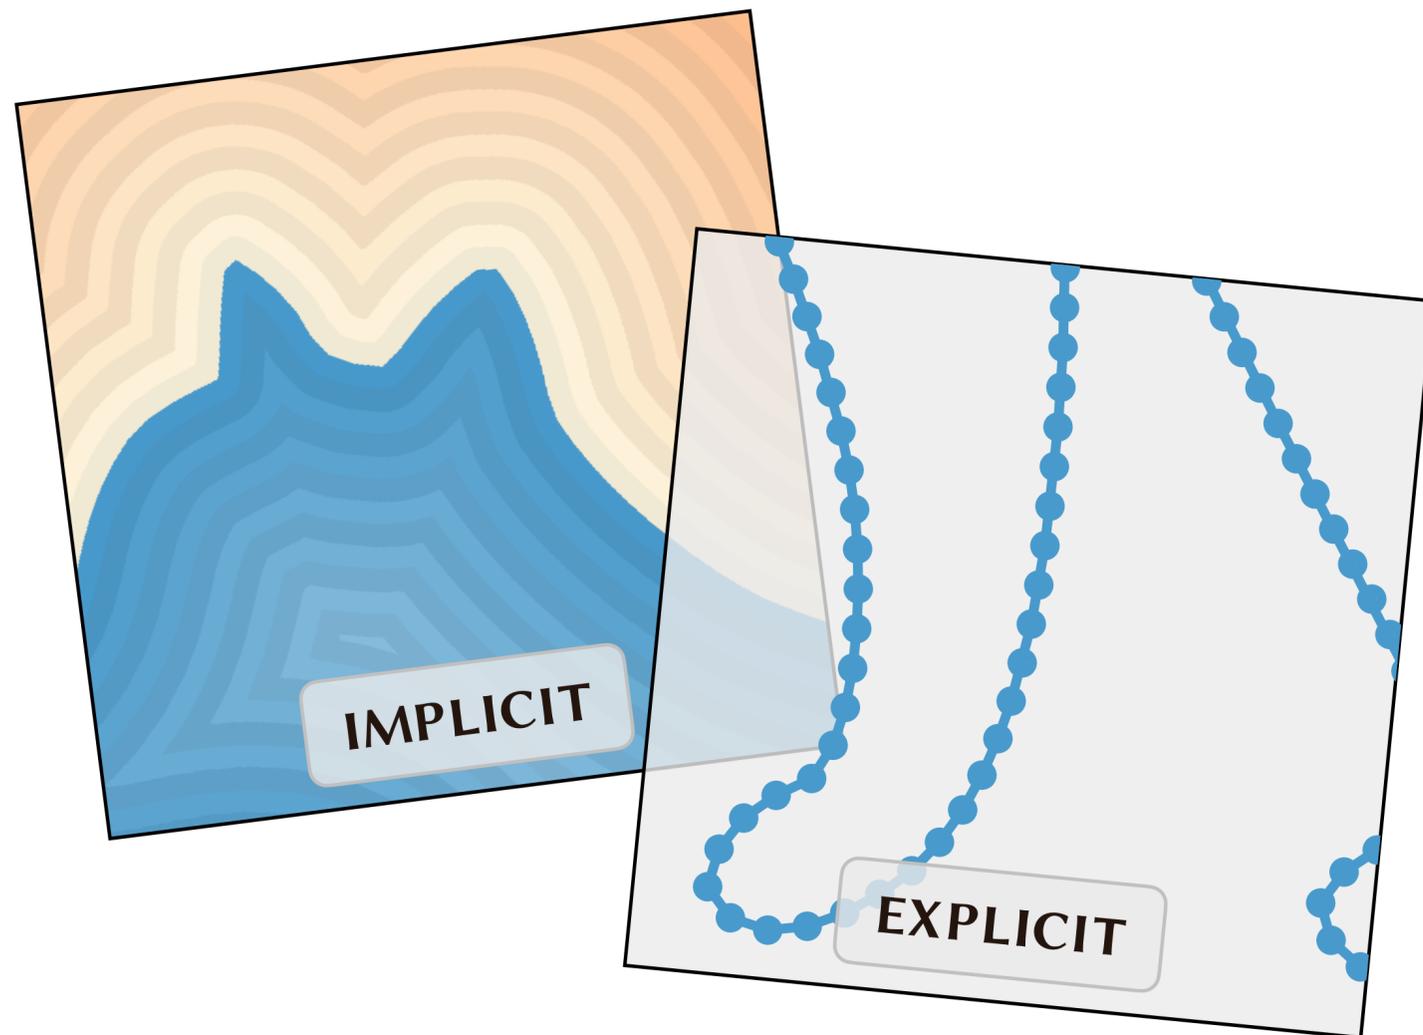
- geometric data stored in weights of neural network → downstream operations can be difficult
  - e.g., editing



## WHY NOT?

1. Imprecision
2. **Editing (etc.) is difficult**
3. The wrong representation

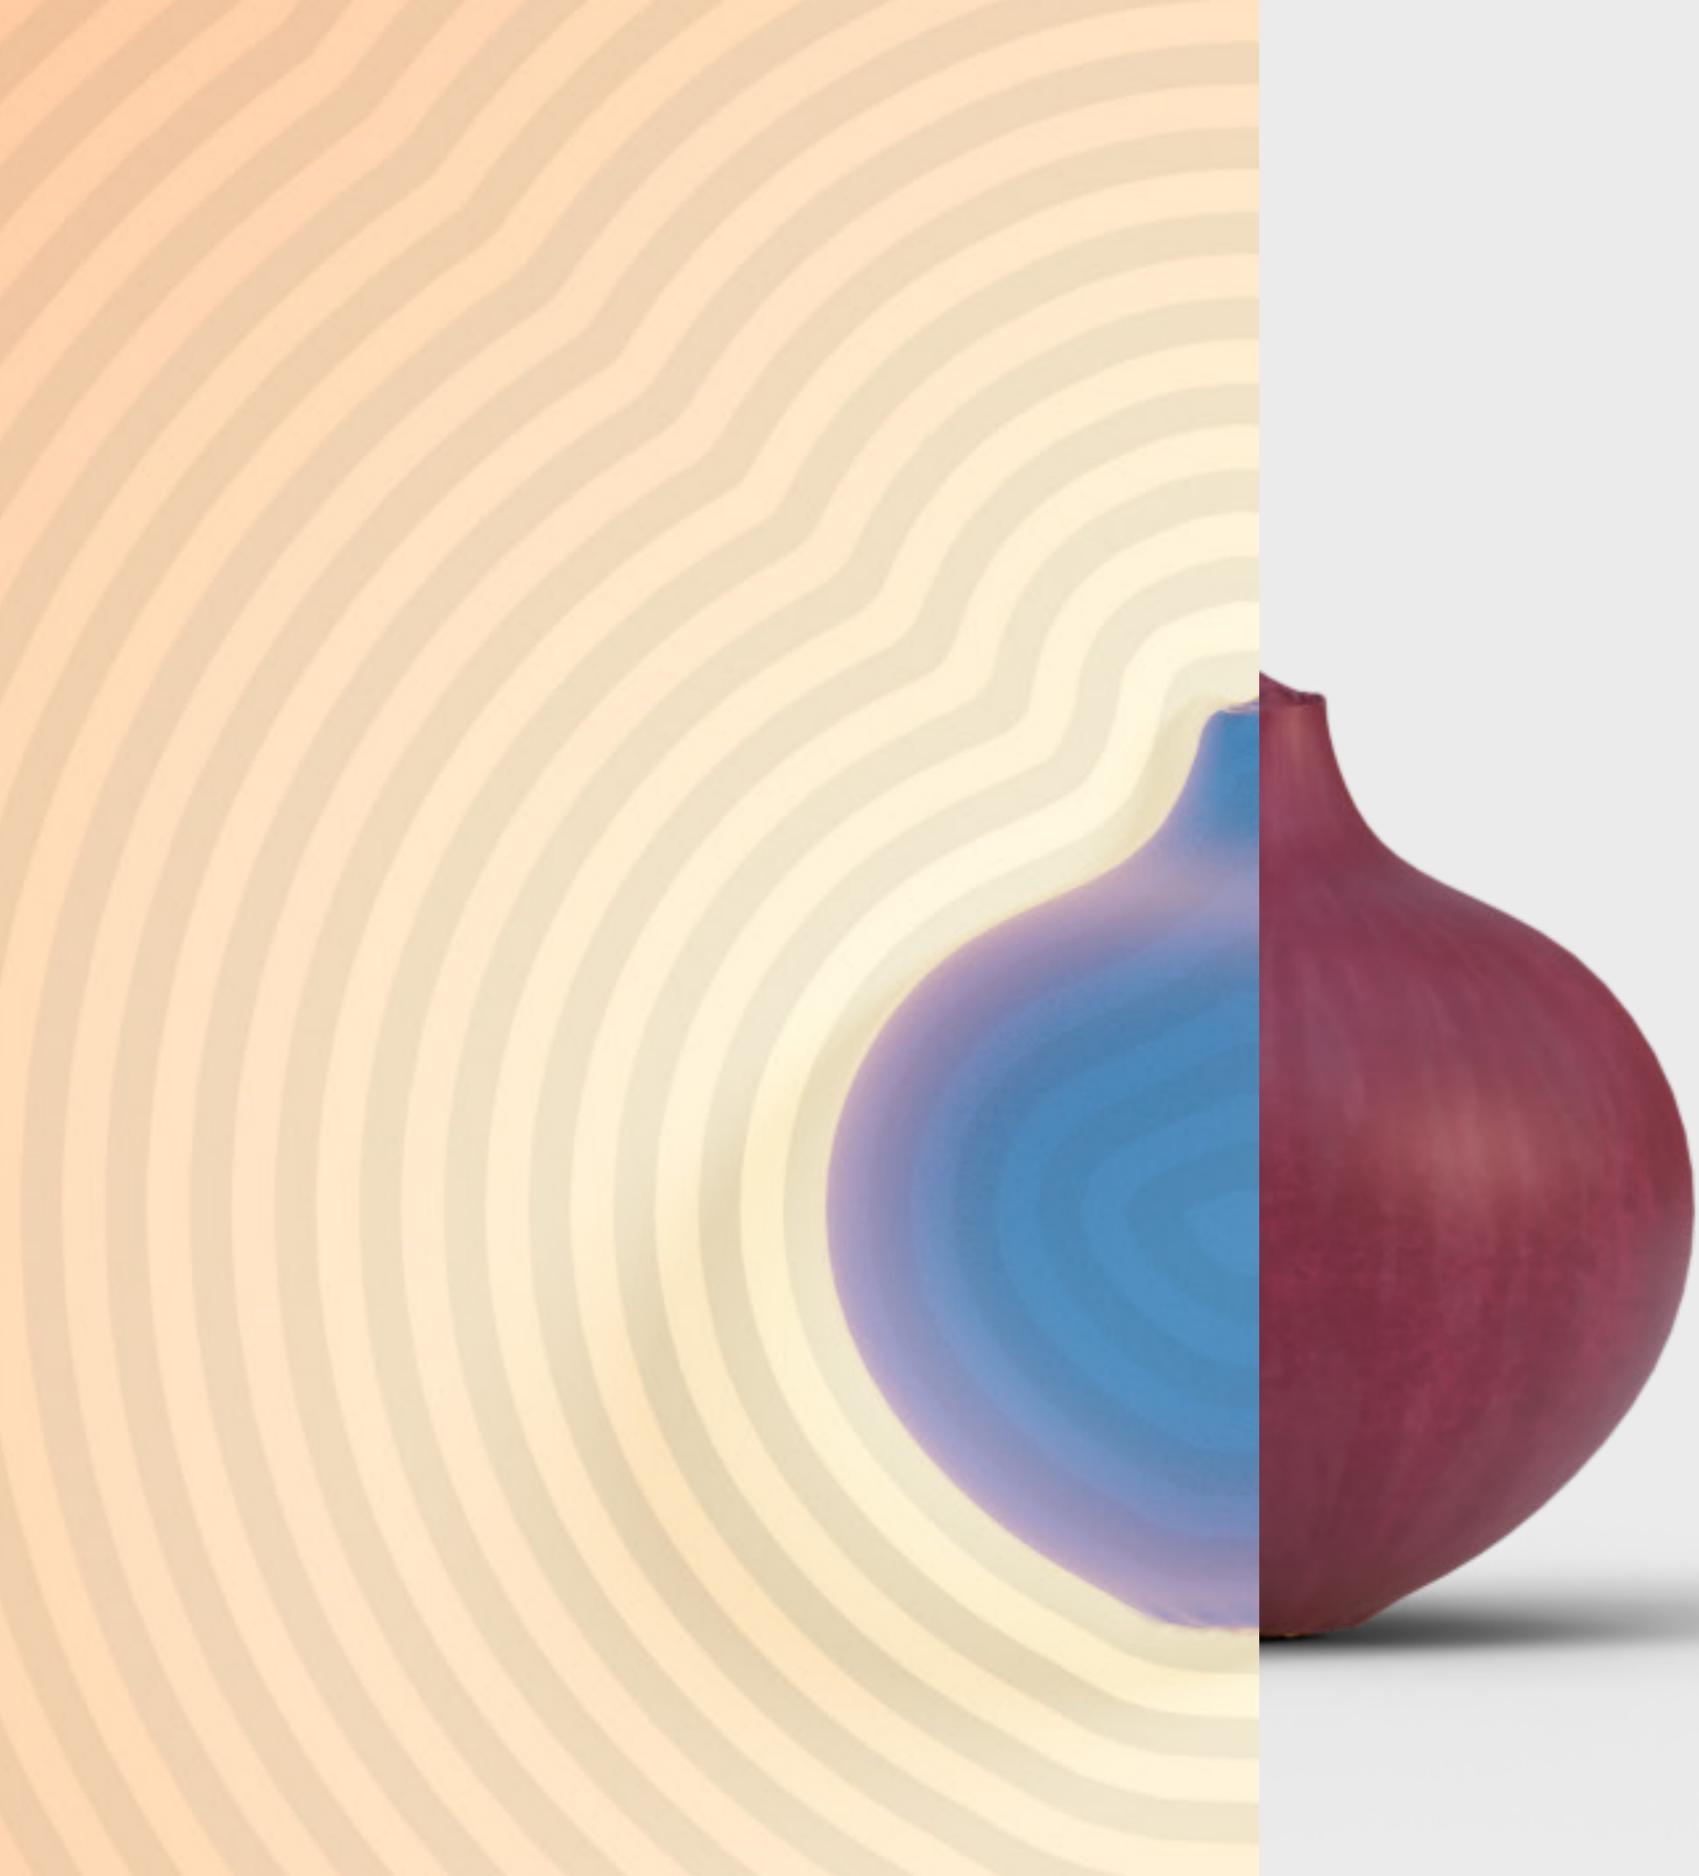
● ●———●  
*layer 6,  $A_{50,50}$*



## WHY NOT?

1. Imprecision
2. Editing (etc.) is difficult
3. ✨ **The wrong representation** ✨

**Pick the right representation for your problem!**



## OVERVIEW

### IMPLICIT SURFACES

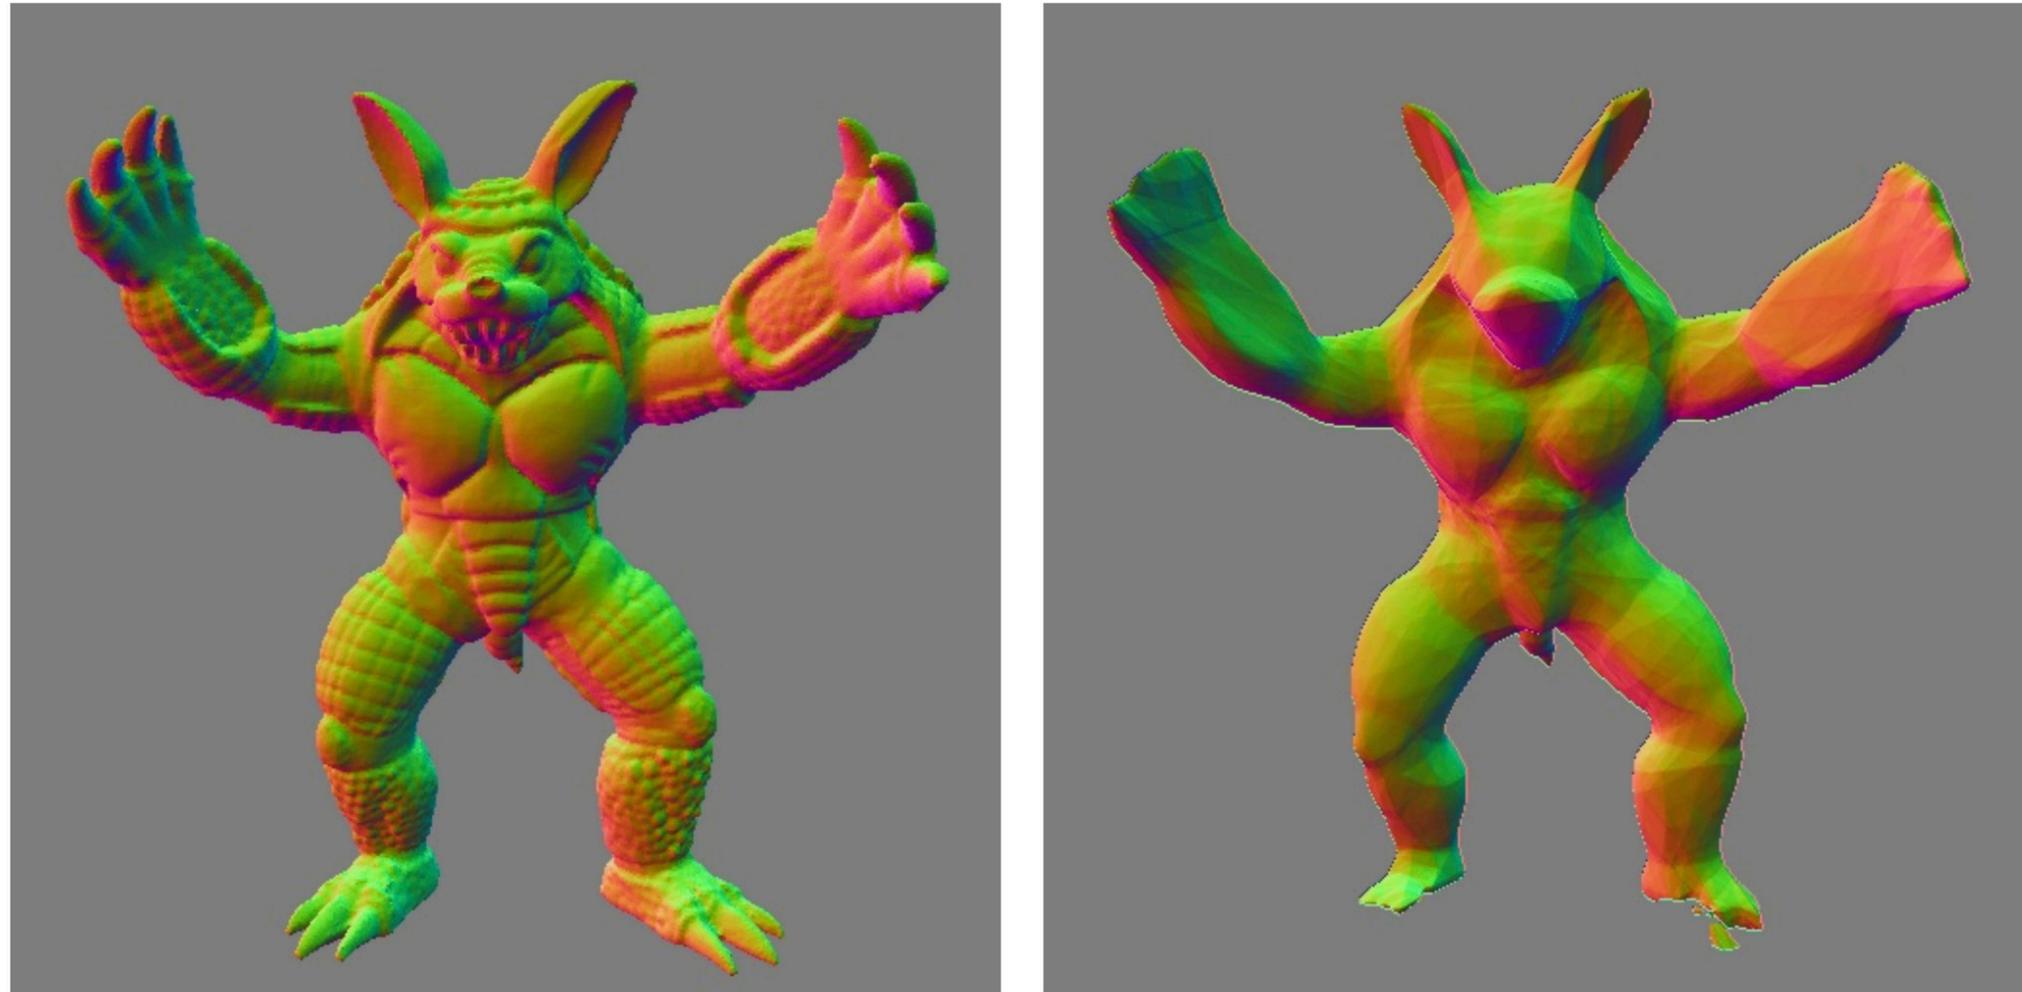
1. Background & History
2. Digital Representations

### NEURAL IMPLICIT

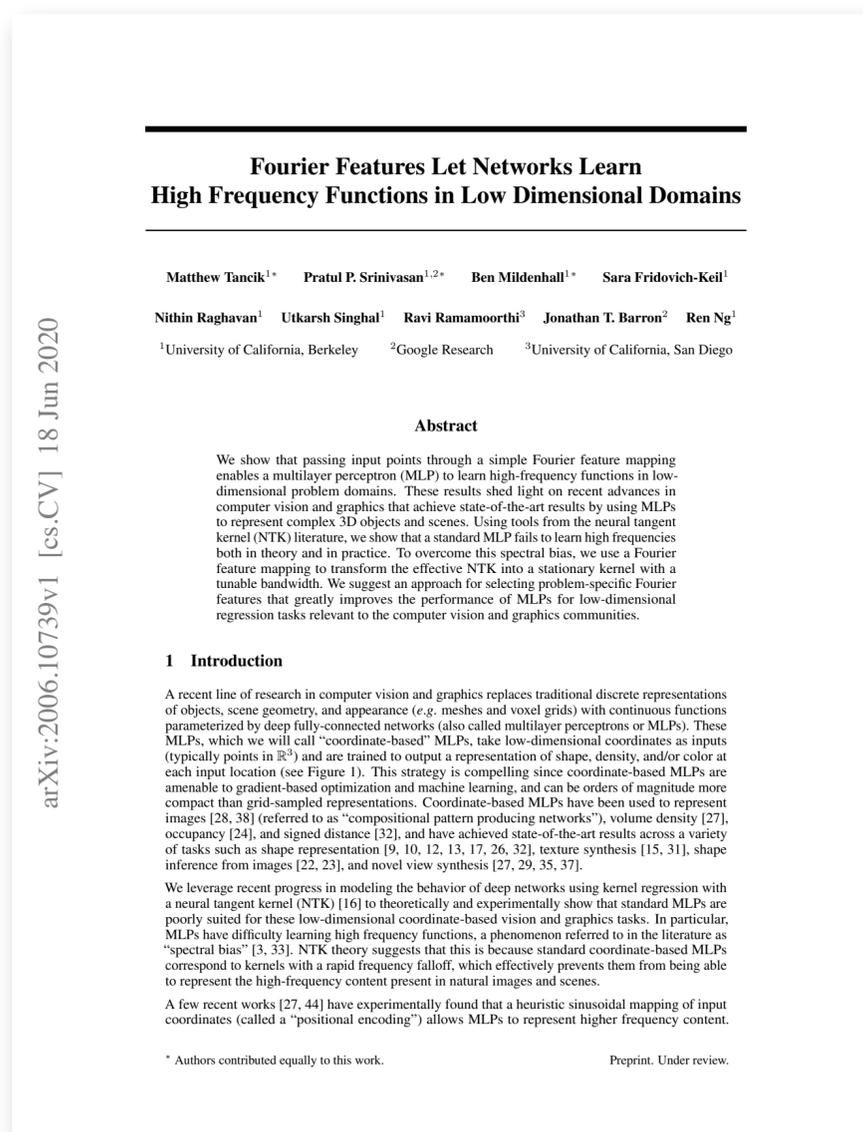
3. The Basics
4. Why & Why Not Use Them?
5. [Working out the Details](#)
6. Geometric Operations

### CONCLUSION

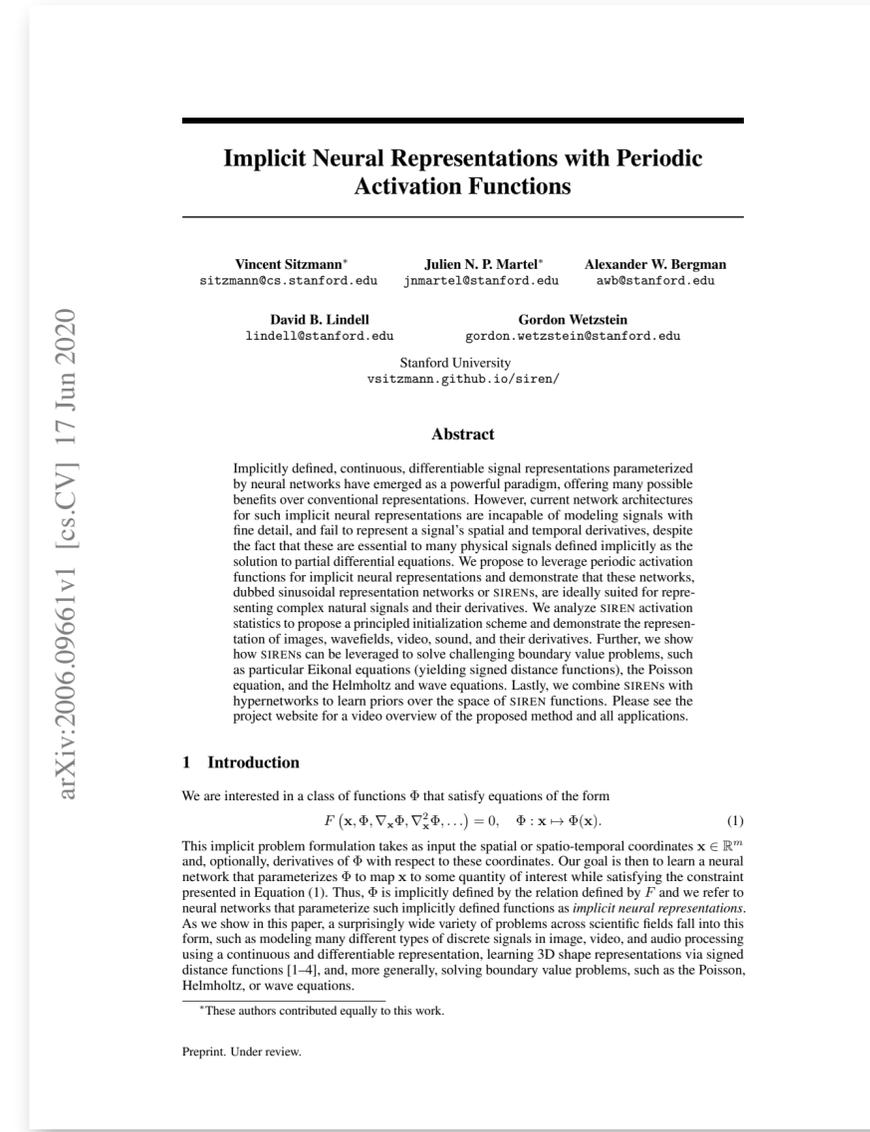
# Sharp features in neural implicits



Sharp features are hard for standard techniques to capture!



Fourier Features [Tancik et al., 2020]

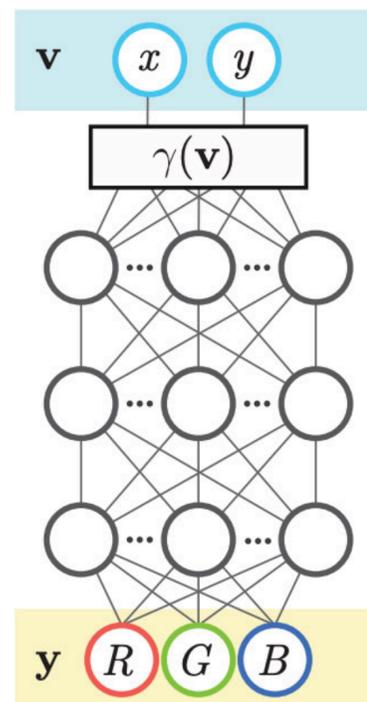


SIREN [Sitzmann et al., 2020]

# Sharp features in neural implicits

augment input points with “Fourier features”

$$\gamma(\mathbf{v}) = [a_1 \cos(2\pi \mathbf{b}_1^T \mathbf{v}), a_1 \sin(2\pi \mathbf{b}_1^T \mathbf{v}), \dots, a_m \cos(2\pi \mathbf{b}_m^T \mathbf{v}), a_m \sin(2\pi \mathbf{b}_m^T \mathbf{v})]$$



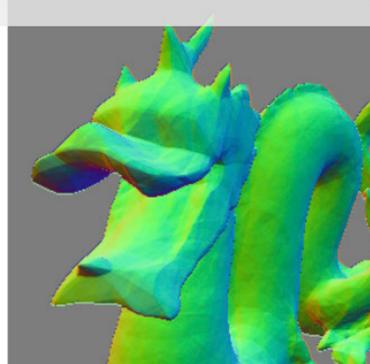
No Fourier features  
 $\gamma(\mathbf{v}) = \mathbf{v}$

With Fourier features  
 $\gamma(\mathbf{v}) = \mathbf{FF}(\mathbf{v})$

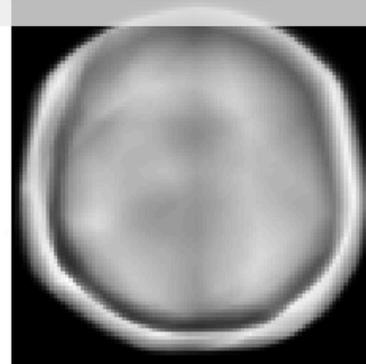
(a) Coordinate-based MLP



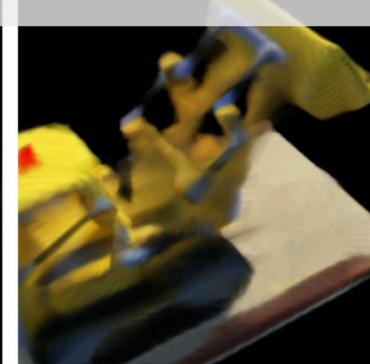
(b) Image regression  
 $(x, y) \rightarrow \text{RGB}$



(c) 3D shape regression  
 $(x, y, z) \rightarrow \text{occupancy}$



(d) MRI reconstruction  
 $(x, y, z) \rightarrow \text{density}$



(e) Inverse rendering  
 $(x, y, z) \rightarrow \text{RGB, density}$

Representations with Periodic Activation Functions

Julien N. P. Martel\*    Alexander W. Bergman  
jnmartel@stanford.edu    awb@stanford.edu

Gordon Wetzstein  
gordon.wetzstein@stanford.edu

Stanford University  
wetzstein.github.io/siren/

Abstract

Continuous, differentiable signal representations parameterized by neural networks have emerged as a powerful paradigm, offering many possible applications. However, current network architectures using such representations are incapable of modeling signals with sharp features, despite the fact that many physical signals are defined implicitly as solutions to partial differential equations. We propose to leverage periodic activation functions in neural representations and demonstrate that these networks, when used as implicit representations or SIRENs, are ideally suited for representing signals and their derivatives. We analyze SIREN activation functions and propose a principled initialization scheme and demonstrate the representations' ability to solve challenging boundary value problems, such as the Poisson equation (yielding signed distance functions), the Helmholtz equation, and the wave equation. Lastly, we combine SIRENs with implicit representations to solve a wide variety of problems across scientific fields. Please see the full paper for an overview of the proposed method and all applications.

Equations  $\Phi$  that satisfy equations of the form

$$\nabla_x \Phi, \nabla_x^2 \Phi, \dots = 0, \quad \Phi: \mathbf{x} \mapsto \Phi(\mathbf{x}). \quad (1)$$

$\mathbf{x}$  takes as input the spatial or spatio-temporal coordinates  $\mathbf{x} \in \mathbb{R}^m$  with respect to these coordinates. Our goal is then to learn a neural map  $\Phi$  to some quantity of interest while satisfying the constraint that  $\Phi$  is implicitly defined by the relation defined by  $F$  and we refer to such implicitly defined functions as *implicit neural representations*. A wide variety of problems across scientific fields fall into this category, including discrete signals in image, video, and audio processing. In this paper, we focus on learning 3D shape representations via signed distance functions, solving boundary value problems, such as the Poisson equation, and solving boundary value problems, such as the Poisson equation.

Fourier Features [Tancik et al., 2020]

SIREN [Sitzmann et al., 2020]



arXiv:2003.04618v2 [cs.CV] 1 Aug 2020

## Convolutional Occupancy

Songyou Peng<sup>1,2</sup> Michael Niemeyer<sup>2,3</sup>  
Marc Pollefeys<sup>1,5</sup> Andreas

<sup>1</sup>ETH Zurich <sup>2</sup>Max Planck Institute for Intel  
<sup>3</sup>University of Tübingen <sup>4</sup>Amazon, Tü

**Abstract.** Recently, implicit neural representations for learning-based 3D reconstruction. While results, most implicit approaches are limited to geometry of single objects and do not scale to more complex scenes. The key limiting factor of implicit methods is the lack of a connected network architecture which does not capture information in the observations or incorporate translational equivariance. In this paper, we propose Convolutional Occupancy Networks, a more flexible implicit representation for 3D reconstruction of objects and 3D scenes. By using encoders with implicit occupancy decoders, our method enables structured reasoning in 3D space, improving the effectiveness of the proposed representation by geometry from noisy point clouds and low-resolution data. We empirically find that our method enables 3D reconstruction of single objects, scales to generalizes well from synthetic to real data.

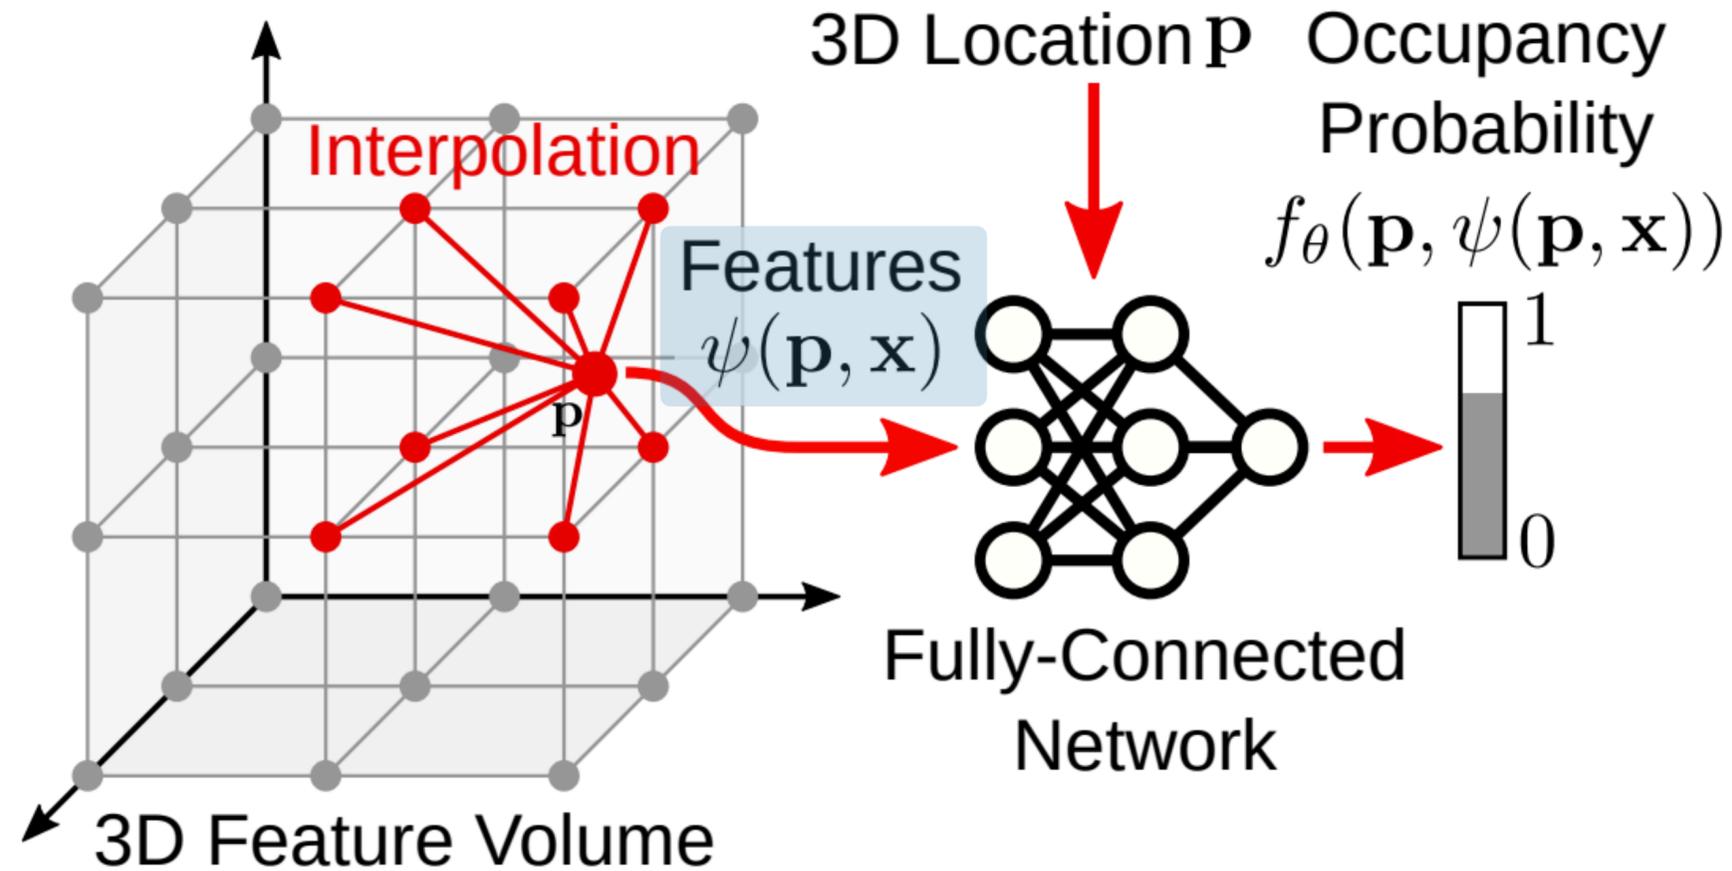
### 1 Introduction

3D reconstruction is a fundamental problem in computer vision applications. An ideal representation of 3D geometry should have the following properties: a) encode complex geometries and large scenes, b) encapsulate local and global information in terms of memory and computation.

Unfortunately, current representations for 3D geometry do not satisfy all of these requirements. Volumetric representations like voxels are low-resolution due to their large memory requirements. Point cloud representations discard topological relationships. Neural networks [13] are often hard to predict using neural networks.

Recently, several works [3, 26, 27, 31] have introduced representations which represent 3D structures using learned occupancy or signed distance functions. In contrast to explicit representations, implicit methods do not discretize 3D space during training, thus resulting in continuous representations of 3D geometry without topology restrictions. While inspiring many follow-up

\* This work was done prior to joining Amazon.



## Convolutional Occupancy Network [Peng et al., 2020]

# Hybrid discrete/continuous fields

## Instant Neural Graphics Primitives with a Multiresolution Hash Encoding

THOMAS MÜLLER, NVIDIA, Switzerland  
ALEX EVANS, NVIDIA, United Kingdom  
CHRISTOPH SCHIED, NVIDIA, USA  
ALEXANDER KELLER, NVIDIA, Germany

<https://nvlabs.github.io/instant-ngp>

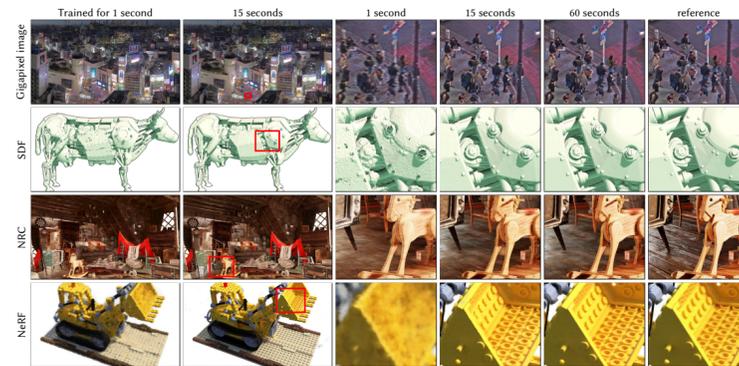


Fig. 1. We demonstrate instant training of neural graphics primitives on a single GPU for multiple tasks. In *Gigapixel image* we represent a gigapixel image by a neural network. *SDF* learns a signed distance function in 3D space whose zero level-set represents a 2D surface. *Neural radiance caching (NRC)* [Müller et al. 2021] employs a neural network that is trained in real-time to cache costly lighting calculations. Lastly, *NeRF* [Mildenhall et al. 2020] uses 2D images and their camera poses to reconstruct a volumetric radiance-and-density field that is visualized using ray marching. In all tasks, our encoding and its efficient implementation provide clear benefits: rapid training, high quality, and simplicity. Our encoding is task-agnostic: we use the same implementation and hyperparameters across all tasks and only vary the hash table size which trades off quality and performance. Tokyo gigapixel photograph ©Trevor Dobson (CC BY-NC-ND 2.0), Lego bulldozer 3D model ©Håvard Dalen (CC BY-NC 2.0)

Neural graphics primitives, parameterized by fully connected neural networks, can be costly to train and evaluate. We reduce this cost with a versatile new input encoding that permits the use of a smaller network without sacrificing quality, thus significantly reducing the number of floating point and memory access operations: a small neural network is augmented by a multiresolution hash table of trainable feature vectors whose values are optimized through stochastic gradient descent. The multiresolution structure allows the network to disambiguate hash collisions, making for a simple

Authors' addresses: Thomas Müller, NVIDIA, Zürich, Switzerland, [tmueller@nvidia.com](mailto:tmueller@nvidia.com); Alex Evans, NVIDIA, London, United Kingdom, [alex@nvidia.com](mailto:alex@nvidia.com); Christoph Schied, NVIDIA, Seattle, USA, [cschied@nvidia.com](mailto:cschied@nvidia.com); Alexander Keller, NVIDIA, Berlin, Germany, [akeller@nvidia.com](mailto:akeller@nvidia.com).

© 2022 Copyright held by the owner/author(s). Publication rights licensed to ACM. This is the author's version of the work. It is posted here for your personal use. Not for redistribution. The definitive Version of Record was published in *ACM Transactions on Graphics*, <https://doi.org/10.1145/3528223.3530127>.

architecture that is trivial to parallelize on modern GPUs. We leverage this parallelism by implementing the whole system using fully-fused CUDA kernels with a focus on minimizing wasted bandwidth and compute operations. We achieve a combined speedup of several orders of magnitude, enabling training of high-quality neural graphics primitives in a matter of seconds, and rendering in tens of milliseconds at a resolution of 1920×1080.

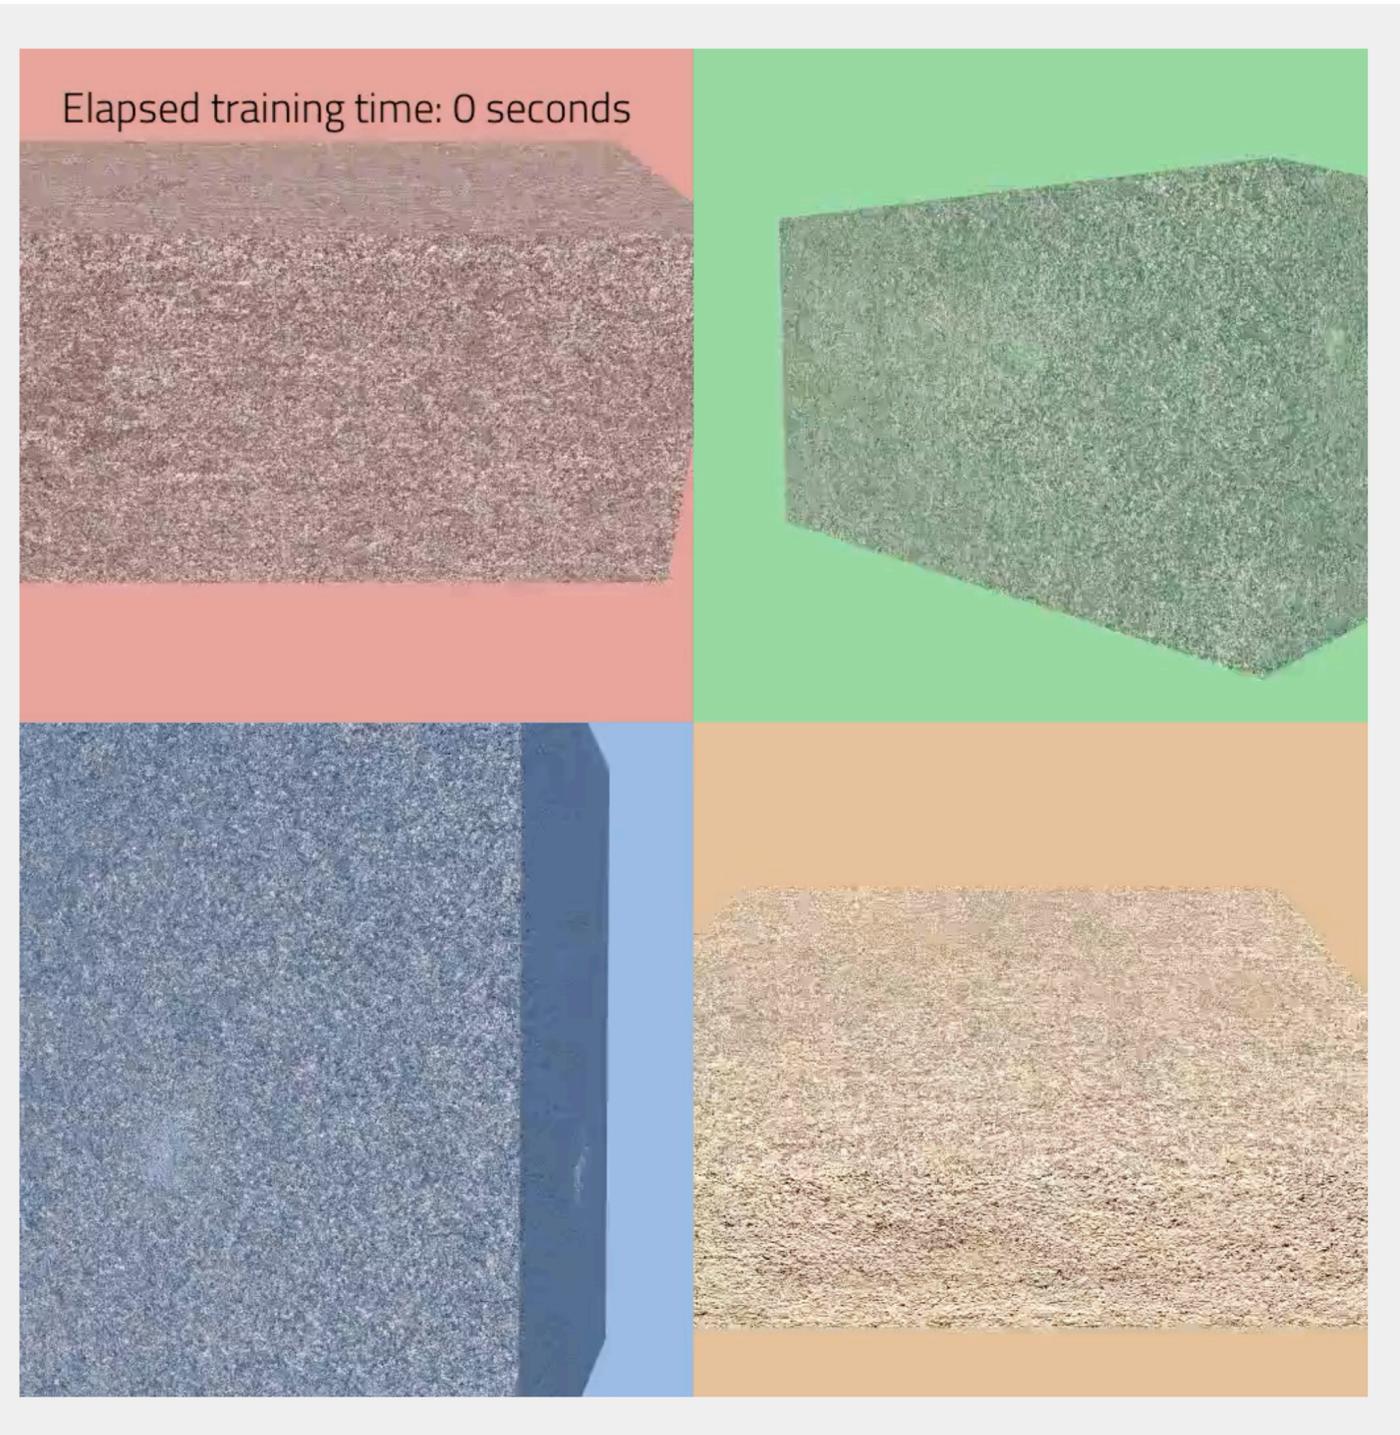
CCS Concepts: • **Computing methodologies** → *Massively parallel algorithms; Vector/streaming algorithms; Neural networks.*

Additional Key Words and Phrases: Image Synthesis, Neural Networks, Encodings, Hashing, GPUs, Parallel Computation, Function Approximation.

### ACM Reference Format:

Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. 2022. Instant Neural Graphics Primitives with a Multiresolution Hash Encoding. *ACM Trans. Graph.* 41, 4, Article 102 (July 2022), 15 pages. <https://doi.org/10.1145/3528223.3530127>

ACM Trans. Graph., Vol. 41, No. 4, Article 102. Publication date: July 2022.



## Instant NGP [Müller et al., 2022]

# Training neural implicits: sampling points

arXiv:2009.09808v3 [cs.GR] 17 Jan 2021

## ON THE EFFECTIVENESS OF WEIGHT-ENCODED NEURAL IMPLICIT 3D SHAPES

Thomas Davies<sup>1</sup>, Derek Nowrouzezahrai<sup>2</sup> & Alec Jacobson<sup>1</sup>  
<sup>1</sup>Department of Computer Science, University of Toronto  
<sup>2</sup>Centre for Intelligent Machines, McGill University

### ABSTRACT

A neural implicit outputs a number indicating whether the given query point in space is inside, outside, or on a surface. Many prior works have focused on *latent-encoded* neural implicits, where a latent vector encoding of a specific shape is also fed as input. While affording latent-space interpolation, this comes at the cost of reconstruction accuracy for any *single* shape. Training a specific network for each 3D shape, a *weight-encoded* neural implicit may forgo the latent vector and focus reconstruction accuracy on the details of a single shape. While previously considered as an intermediary representation for 3D scanning tasks or as a toy-problem leading up to latent-encoding tasks, weight-encoded neural implicits have not yet been taken seriously as a 3D shape representation. In this paper, we establish that weight-encoded neural implicits meet the criteria of a first-class 3D shape representation. We introduce a suite of technical contributions to improve reconstruction accuracy, convergence, and robustness when learning the signed distance field induced by a polygonal mesh — the *de facto* standard representation. Viewed as a lossy compression, our conversion outperforms standard techniques from geometry processing. Compared to previous latent- and weight-encoded neural implicits we demonstrate superior robustness, scalability, and performance.

### 1 INTRODUCTION

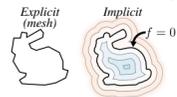
While 3D surface representation has been a foundational topic of study in the computer graphics community for over four decades, recent developments in machine learning have highlighted the potential that neural networks can play as effective parameterizations of solid shapes.

The success of neural approaches to shape representations has been evidenced both through their ability of representing complex geometries as well as their utility in end-to-end 3D shape learning, reconstruction, and understanding and tasks. These approaches also make use of the growing availability of user generated 3D content and high-fidelity 3D capture devices, e.g., point cloud scanners.

For these 3D tasks, one powerful configuration is to represent a 3D surface  $\mathcal{S}$  as the set containing any point  $\vec{x} \in \mathbb{R}^3$  for which an implicit function (i.e., a neural network) evaluates to zero:

$$\mathcal{S} := \{ \vec{x} \in \mathbb{R}^3 \mid f_{\theta}(\vec{x}; \vec{z}) = 0 \}, \quad (1)$$

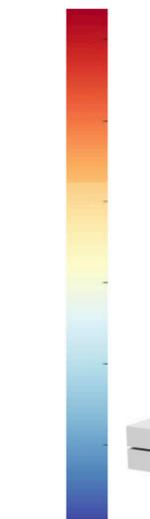
where  $\theta \in \mathbb{R}^m$  are the network weights and  $\vec{z} \in \mathbb{R}^k$  is an input latent vector encoding a particular shape. In contrast to the *de facto* standard polygonal mesh representation which *explicitly* discretizes a surface's geometry, the function  $f$  *implicitly* defines the shape  $\mathcal{S}$  encoded in  $\vec{z}$ . We refer to the representation in Eq. (1) as a *latent-encoded neural implicit*.



Park et al. (2019) propose to optimize the weights  $\theta$  so each shape  $\mathcal{S}_i \in \mathcal{D}$  in a dataset or shape distribution  $\mathcal{D}$  is encoded into a corresponding latent vector  $\vec{z}_i$ . If successfully trained, the weights  $\theta$  of their DEEPSDF implicit function  $f_{\theta}$  can be said to generalize across the “shape space” of  $\mathcal{D}$ . As always with supervision, reducing the training set from  $\mathcal{D}$  will affect  $f$ 's ability to generalize and can lead to overfitting. Doing so may seem, at first, to be an ill-fated and uninteresting idea.

Our work considers an extreme case — when the training set is reduced to a single shape  $\mathcal{S}_i$ . We can draw a simple but powerful conclusion: in this setting, one can completely forgo the latent vector

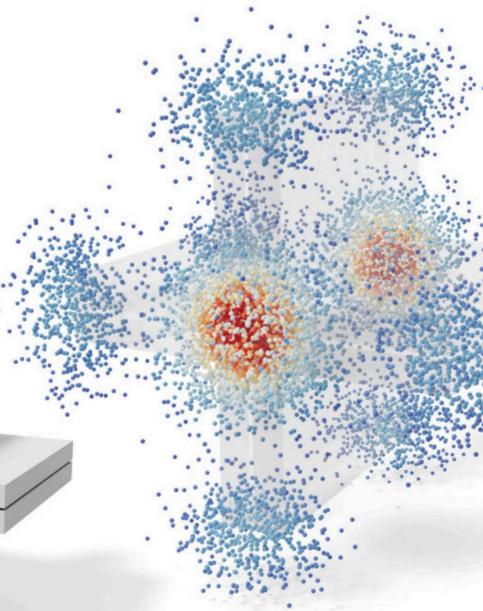
high  
density



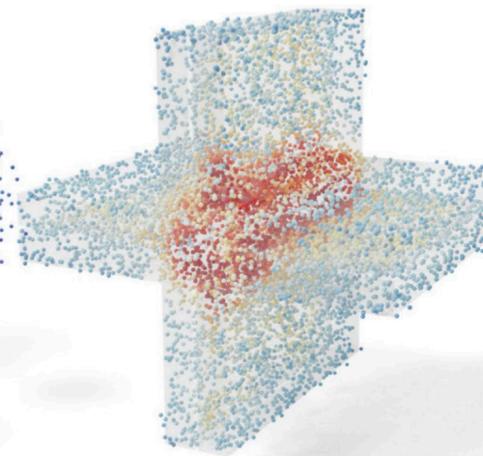
low  
density



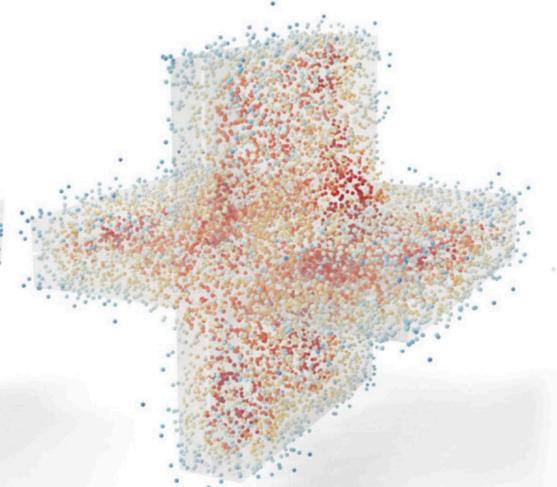
Vertex



Surface



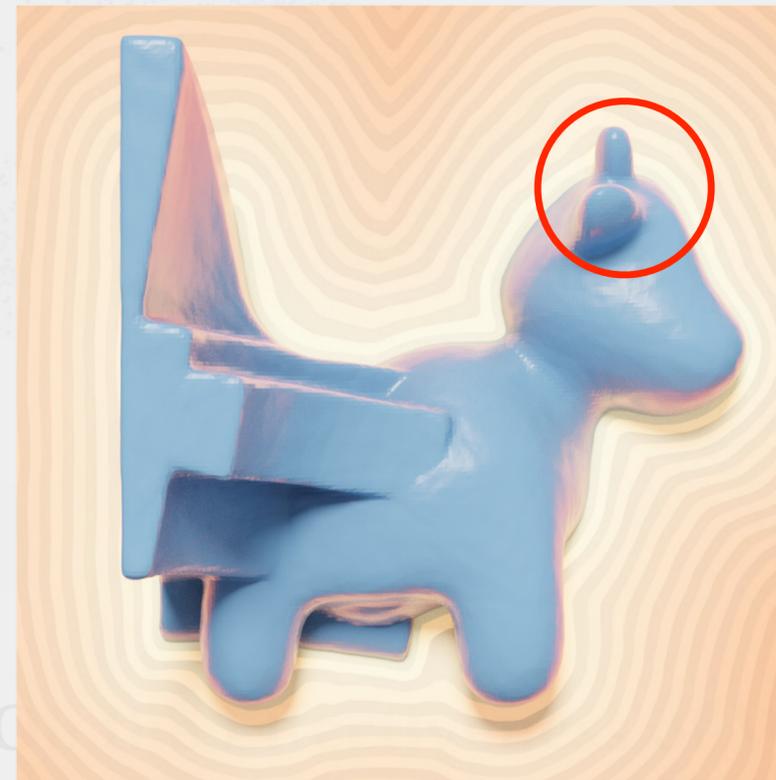
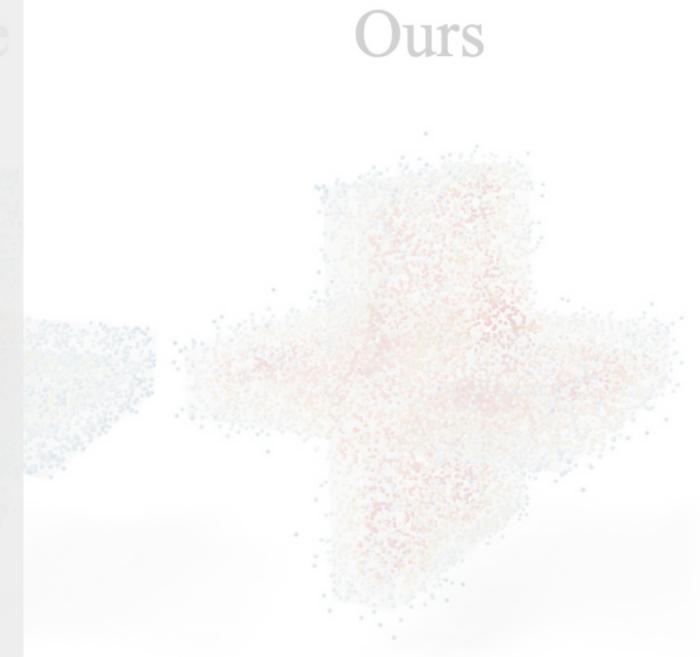
Ours



from **On the Effectiveness of Weight-Encoded  
Neural Implicit 3D Shapes** [Davies et al., 2021]

# Training neural implicits: sampling points

care more about reconstruction quality near surface  $\Rightarrow$   
bias loss, achieved via importance sampling near surface



Weight-Encoded  
Neural Implicit 3D Shapes [Davies et al., 2021]

ON THE EFFECTIVENESS OF  
WEIGHT-ENCODED NEURAL IMPLICIT 3D SHAPES

Thomas Davies<sup>1</sup>, Derek Nowrouzezahrai<sup>2</sup> & Alec Jacobson<sup>1</sup>  
<sup>1</sup>Department of Computer Science, University of Toronto  
<sup>2</sup>Centre for Intelligent Machines, McGill University

## ABSTRACT

A neural implicit outputs a number indicating whether the given query point in space is inside, outside, or on a surface. Many prior works have focused on *latent-encoded* neural implicits, where a latent vector encoding of a specific shape is also used as input. While affording latent-space interpolation, this comes at the cost of reconstruction accuracy for any *single* shape. Training a specific network for each 3D shape, a *weight-encoded* neural implicit may forgo the latent vector and focus reconstruction accuracy on the details of a single shape. While previously considered as an intermediary representation for 3D scanning tasks or as a toy-problem leading up to latent-encoding tasks, weight-encoded neural implicits have not yet been taken seriously as a 3D shape representation. In this paper, we establish that weight-encoded neural implicits meet the criteria of a first-class 3D shape representation. We introduce a suite of technical contributions to improve reconstruction accuracy, convergence, and robustness when learning the signed distance field induced by a polygonal mesh — the *de facto* standard representation. Viewed as a lossy compression, our conversion outperforms standard techniques from geometry processing. Compared to previous latent- and weight-encoded neural implicits we demonstrate superior robustness, scalability, and performance.

## 1 INTRODUCTION

While 3D surface representation has been a foundational topic of study in the computer graphics community for over four decades, recent developments in machine learning have highlighted potential that neural networks can play as effective parameterizations of solid shapes.

The success of neural approaches to shape representations has been evidenced both through the ability of representing complex geometries as well as their utility in end-to-end 3D shape learning, reconstruction, and understanding and tasks. These approaches also make use of the growing availability of user generated 3D content and high-fidelity 3D capture devices, e.g., point cloud scans.

For these 3D tasks, one powerful configuration is to represent a 3D surface  $S$  as the set containing any point  $\vec{x} \in \mathbb{R}^3$  for which an implicit function (i.e., a neural network) evaluates to zero:

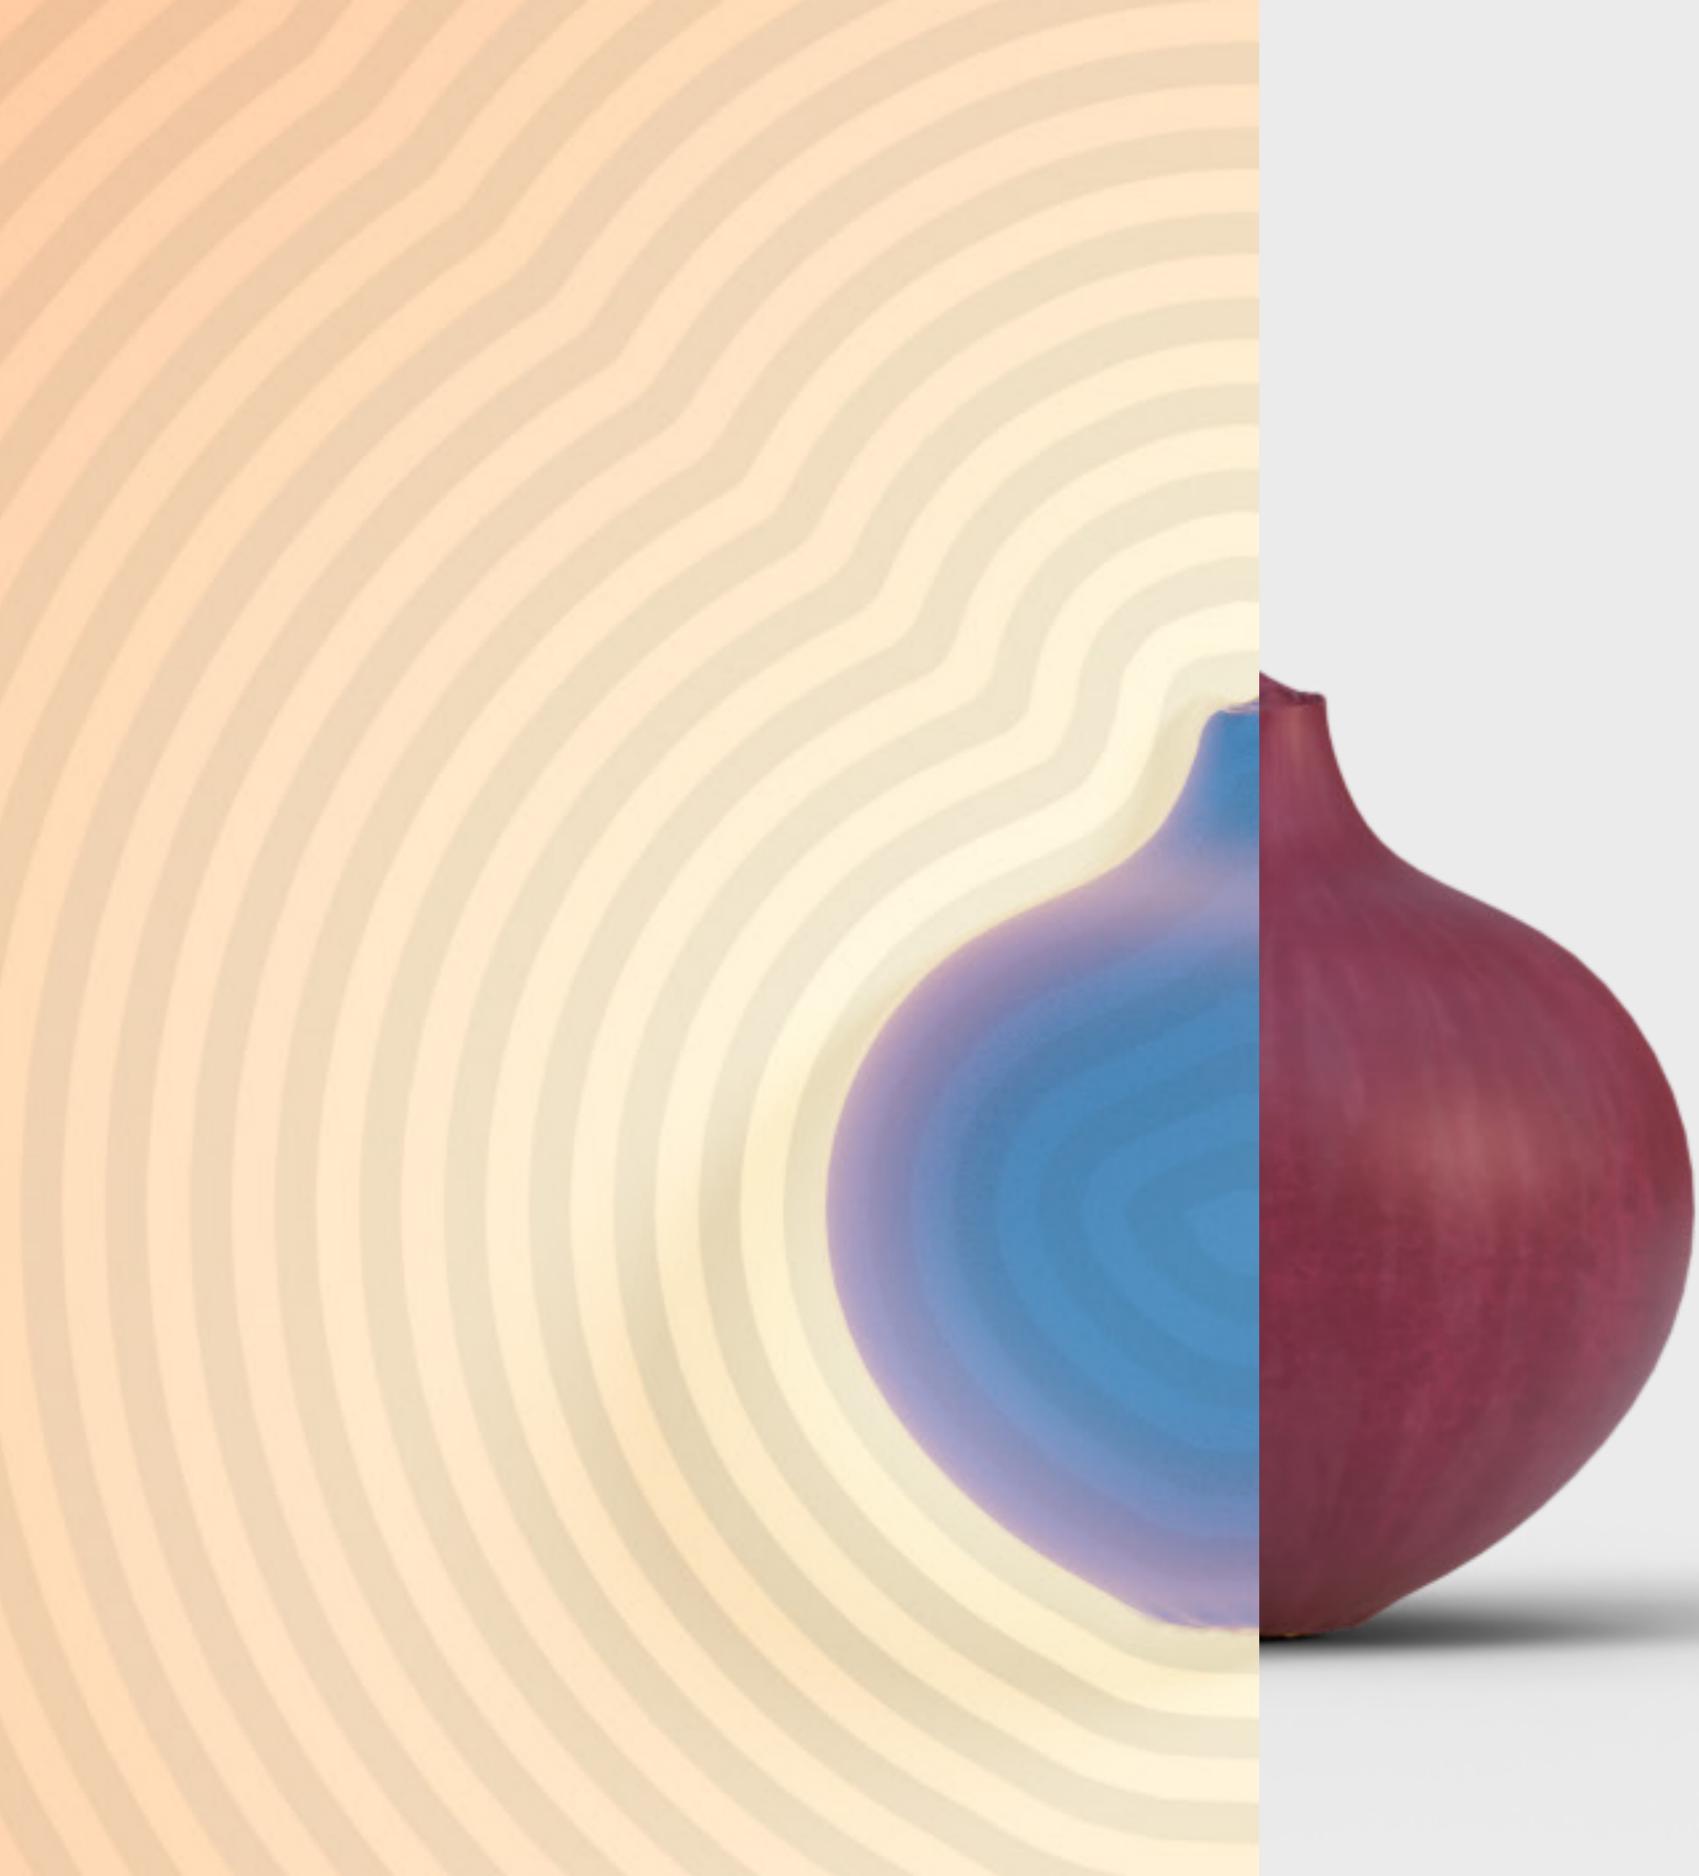
$$S := \{ \vec{x} \in \mathbb{R}^3 \mid f_\theta(\vec{x}; \vec{z}) = 0 \},$$

where  $\theta \in \mathbb{R}^m$  are the network weights and  $\vec{z} \in \mathbb{R}^k$  is an input latent vector encoding a particular shape. In contrast to the *de facto* standard polygonal mesh representation which *explicitly* discretizes a surface's geometry, the function  $f$  *implicitly* defines the shape  $S$  encoded in  $\vec{z}$ . We refer to the representation in Eq. (1) as a *latent-encoded neural implicit*.

Park et al. (2019) propose to optimize the weights  $\theta$  so each shape  $S_i \in \mathcal{D}$  in a dataset or shape distribution  $\mathcal{D}$  is encoded into a corresponding latent vector  $\vec{z}_i$ . If successfully trained, the weights  $\theta$  of their DEEPSDF implicit function  $f_\theta$  can be said to generalize across the “shape space” of  $\mathcal{D}$ . As always with supervision, reducing the training set from  $\mathcal{D}$  will affect  $f$ 's ability to generalize and can lead to overfitting. Doing so may seem, at first, to be an ill-fated and uninteresting idea.

Our work considers an extreme case — when the training set is reduced to a single shape  $S_1$ . We draw a simple but powerful conclusion: in this setting, one can completely forgo the latent vector





## OVERVIEW

### IMPLICIT SURFACES

1. Background & History
2. Digital Representations

### NEURAL IMPLICIT

3. The Basics
4. Why & Why Not Use Them?
5. Working out the Details
6. **Geometric Operations**

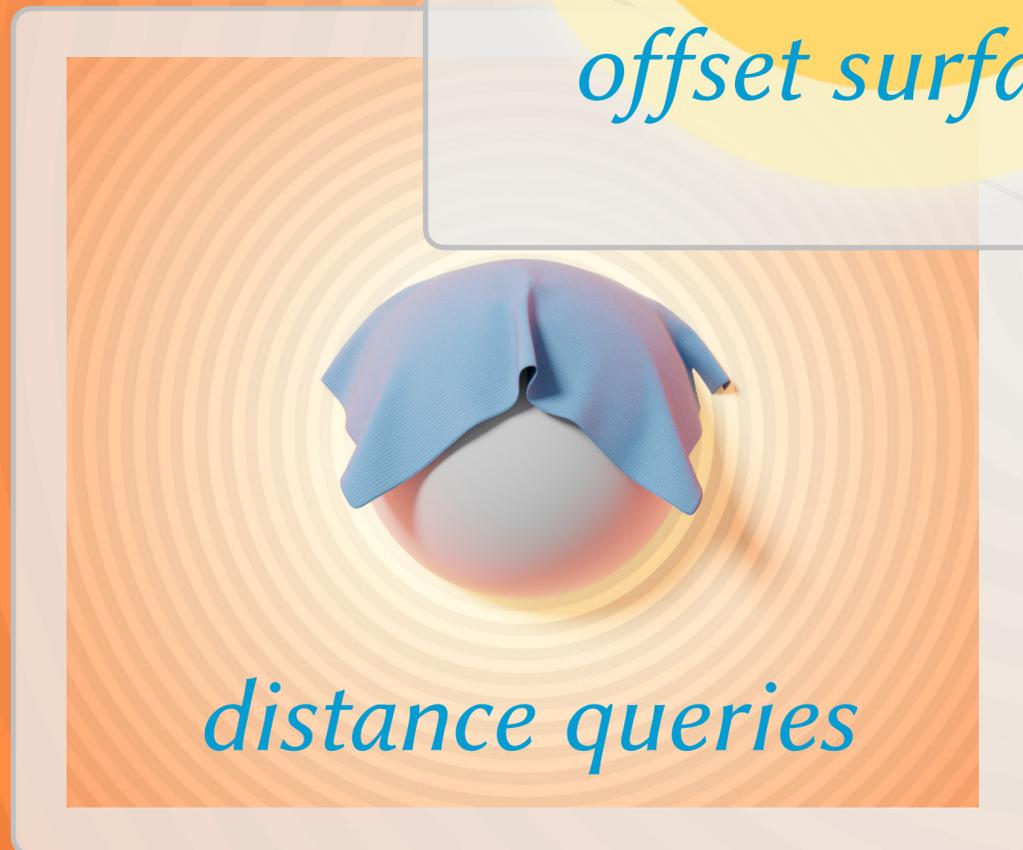
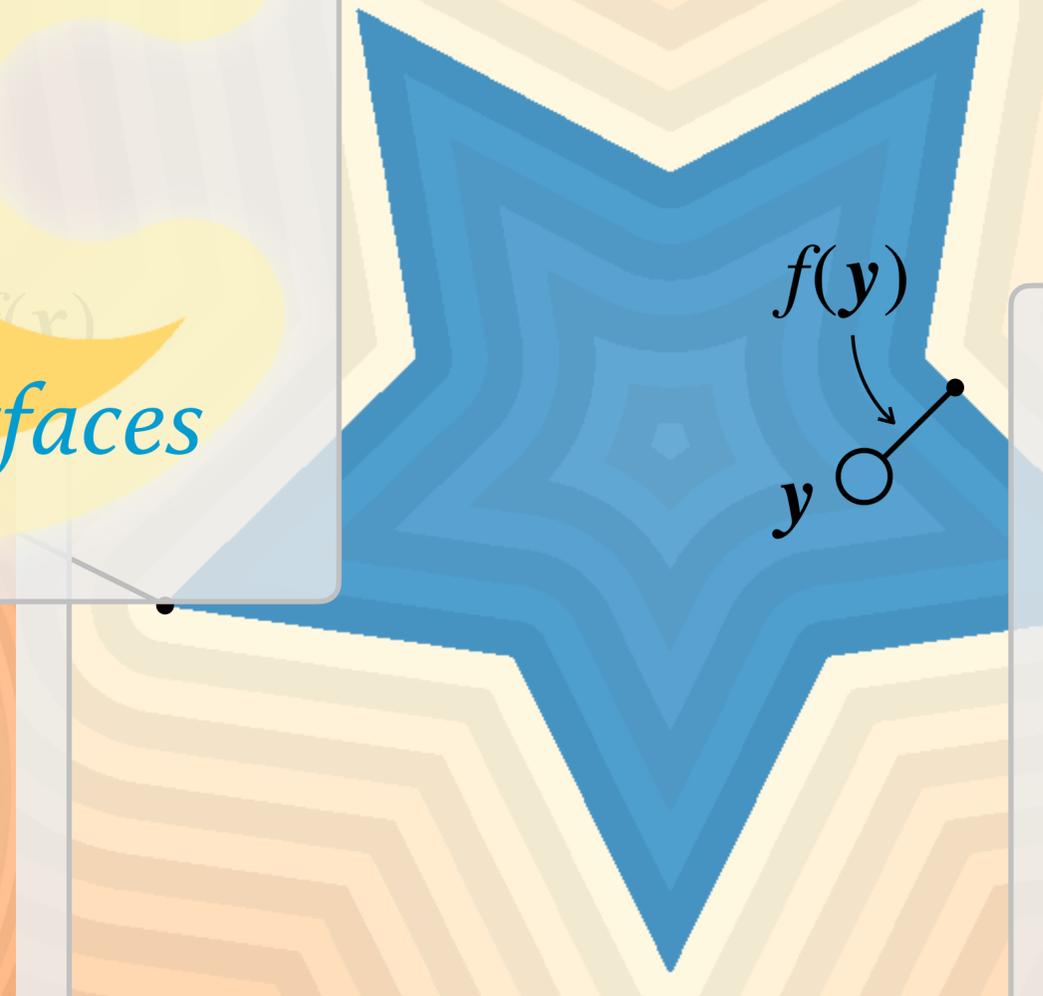
### CONCLUSION

## Signed Distance Function

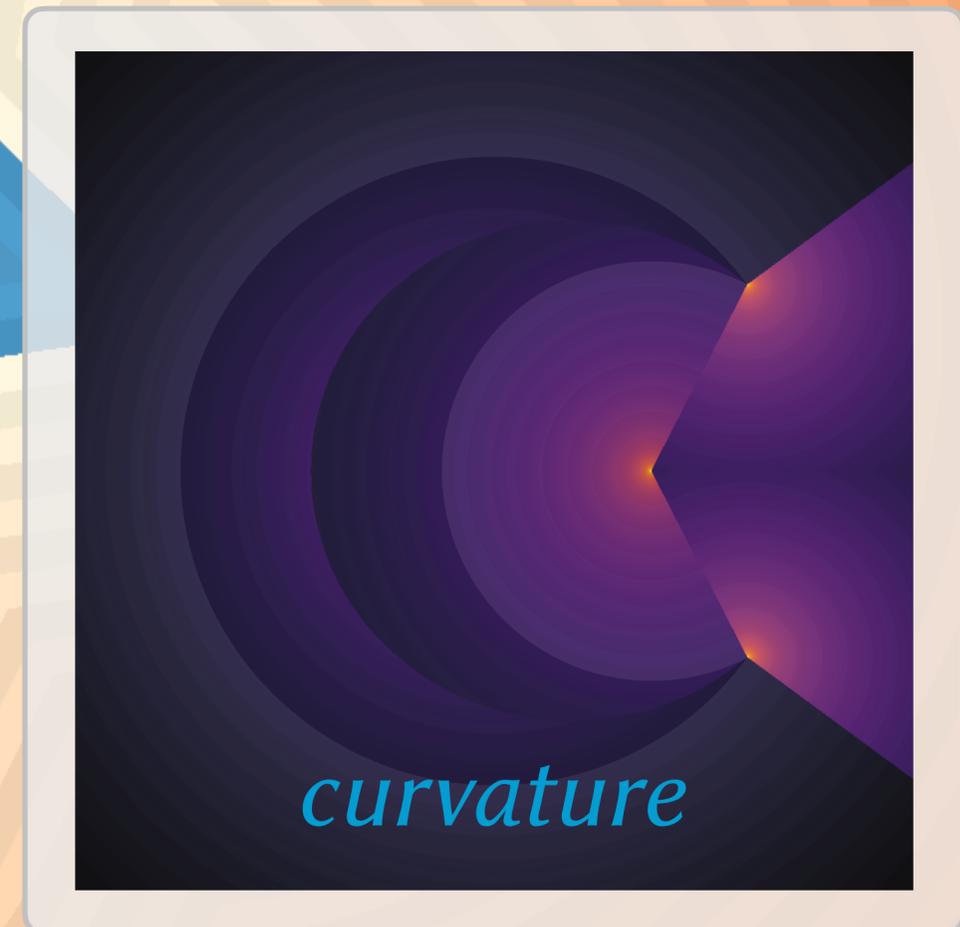
- $|f(\mathbf{x})|$  encodes distance to closest point on surface
- many geometric queries easy



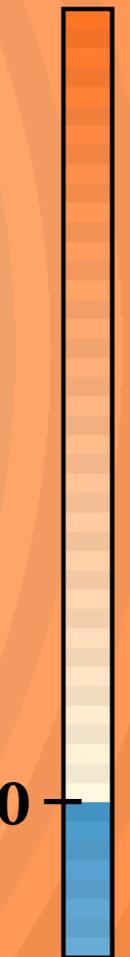
*offset surfaces*



*distance queries*



*curvature*



How do we make a neural implicit an SDF?

How do we make a neural implicit an SDF?

$$\mathcal{L}(x) = \mathcal{L}_{\text{org}}(x) + \mathcal{L}_{\text{SDF}}(x)$$

*add regularization term to loss function*

arXiv:2002.10099v2 [cs.LG] 9 Jul 2020

**Implicit Geometric Regularization for Learning Shapes**

Amos Gropp<sup>1</sup> Lior Yariv<sup>1</sup> Niv Haim<sup>1</sup> Matan Atzmon<sup>1</sup> Yaron Lipman<sup>1</sup>

**Abstract**

Representing shapes as level sets of neural networks has been recently proved to be useful for different shape analysis and reconstruction tasks. So far, such representations were computed using either: (i) pre-computed implicit shape representations; or (ii) loss functions explicitly defined over the neural level sets.

In this paper we offer a new paradigm for computing high fidelity implicit neural representations directly from raw data (i.e., point clouds, with or without normal information). We observe that a rather simple loss function, encouraging the neural network to vanish on the input point cloud and to have a unit norm gradient, possesses an implicit geometric regularization property that favors smooth and natural zero level set surfaces, avoiding bad zero-loss solutions.

We provide a theoretical analysis of this property for the linear case, and show that, in practice, our method leads to state of the art implicit neural representations with higher level-of-details and fidelity compared to previous methods.

**1. Introduction**

Recently, level sets of neural networks have been used to represent 3D shapes (Park et al., 2019; Atzmon et al., 2019; Chen & Zhang, 2019; Mescheder et al., 2019), i.e.,

$$\mathcal{M} = \{x \in \mathbb{R}^3 \mid f(x; \theta) = 0\}, \quad (1)$$

where  $f: \mathbb{R}^3 \times \mathbb{R}^m \rightarrow \mathbb{R}$  is a multilayer perceptron (MLP); we call such representations *implicit neural representations*. Compared to the more traditional way of representing surfaces via implicit functions defined on volumetric grids (Wu et al., 2016; Chen et al., 2016; Dai et al., 2017; Stutz & Geiger, 2018), implicit representations enjoy the benefit of relating to pieces of a continuum (i.e., parameters) of the model rather than to the fixed discretization of the grid. So far, most previous works using implicit neural representations computed  $f$  with 3D supervision; that is, by comparing  $f$  to a known (or pre-computed) implicit representation of some shape. (Park et al., 2019) use a regression loss to approximate a pre-computed signed distance functions of shapes; (Chen & Zhang, 2019; Mescheder et al., 2019) use classification loss with pre-computed occupancy function.

In this work we are interested in working directly with raw data: Given an input point cloud  $\mathcal{X} = \{x_i\}_{i \in I} \subset \mathbb{R}^3$ , with or without normal data,  $\mathcal{N} = \{n_i\}_{i \in I} \subset \mathbb{R}^3$ , our goal is to compute  $\theta$  such that  $f(x; \theta)$  is approximately the signed distance function to a plausible surface  $\mathcal{M}$  defined by the point data  $\mathcal{X}$  and normals  $\mathcal{N}$ .

Some previous works are constructing implicit neural representations from raw data. In (Atzmon et al., 2019) no 3D supervision is required and the loss is formulated directly on the zero level set  $\mathcal{M}$ ; iterative sampling of the zero level set is required for formulating the loss. In a more recent work, (Atzmon & Lipman, 2020) use unsigned regression to introduce good local minima that produces useful implicit neural representations, with no 3D supervision and no zero level set



Figure 1. Learning curves from 2D point clouds (white disks) using our method; black lines depict the zero level sets of the trained neural networks,  $\mathcal{M}$ . The implicit geometric regularization during the optimization reaches a plausible explanation of the data.

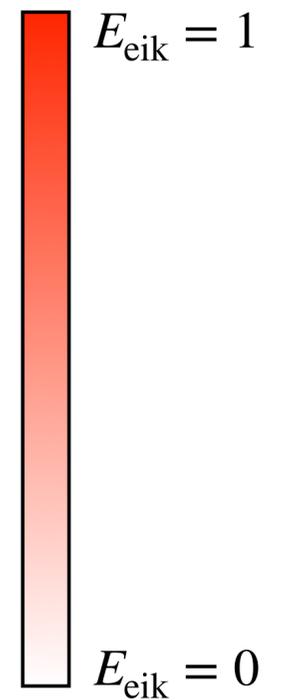
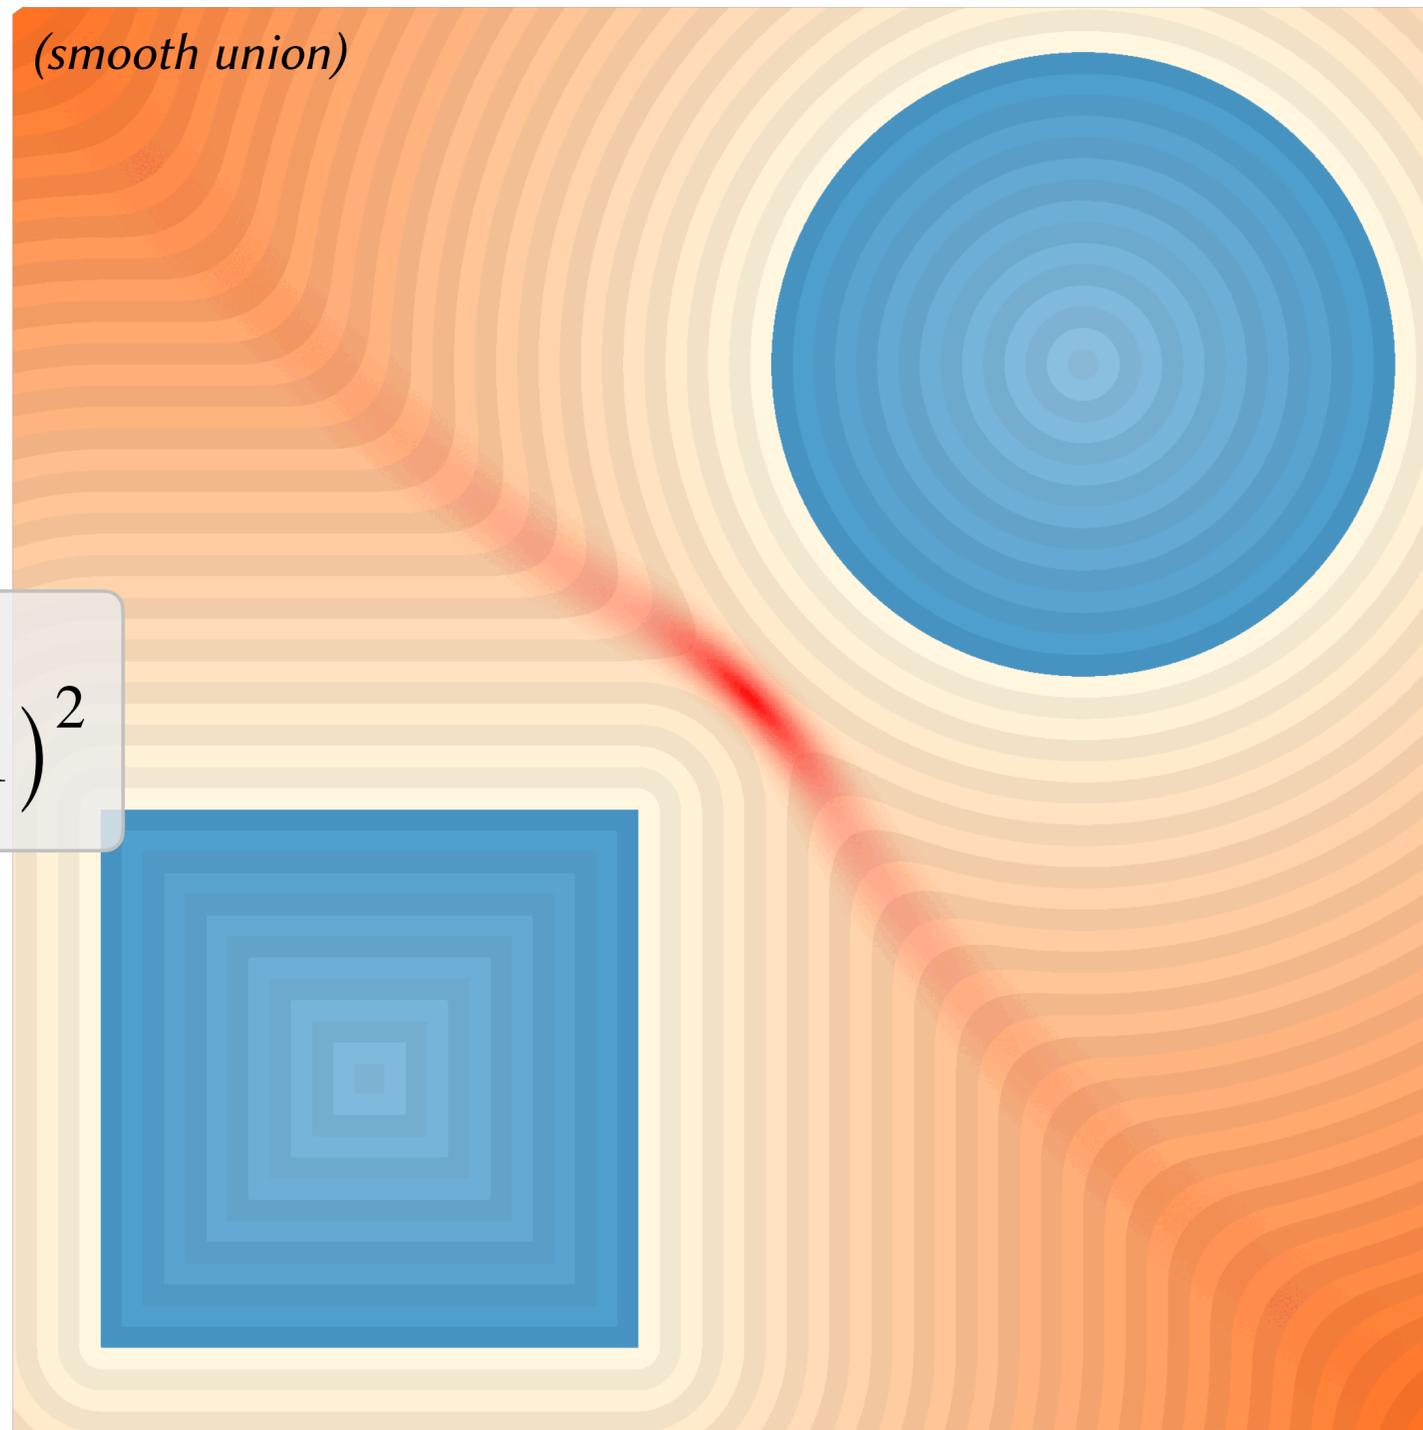
et al., 2016; Chen et al., 2016; Dai et al., 2017; Stutz & Geiger, 2018); implicit representations enjoy the benefit of relating to pieces of a continuum (i.e., parameters) of the model rather than to the fixed discretization of the grid. So far, most previous works using implicit neural representations computed  $f$  with 3D supervision; that is, by comparing  $f$  to a known (or pre-computed) implicit representation of some shape. (Park et al., 2019) use a regression loss to approximate a pre-computed signed distance functions of shapes; (Chen & Zhang, 2019; Mescheder et al., 2019) use classification loss with pre-computed occupancy function.

In this work we are interested in working directly with raw data: Given an input point cloud  $\mathcal{X} = \{x_i\}_{i \in I} \subset \mathbb{R}^3$ , with or without normal data,  $\mathcal{N} = \{n_i\}_{i \in I} \subset \mathbb{R}^3$ , our goal is to compute  $\theta$  such that  $f(x; \theta)$  is approximately the signed distance function to a plausible surface  $\mathcal{M}$  defined by the point data  $\mathcal{X}$  and normals  $\mathcal{N}$ .

Some previous works are constructing implicit neural representations from raw data. In (Atzmon et al., 2019) no 3D supervision is required and the loss is formulated directly on the zero level set  $\mathcal{M}$ ; iterative sampling of the zero level set is required for formulating the loss. In a more recent work, (Atzmon & Lipman, 2020) use unsigned regression to introduce good local minima that produces useful implicit neural representations, with no 3D supervision and no zero level set

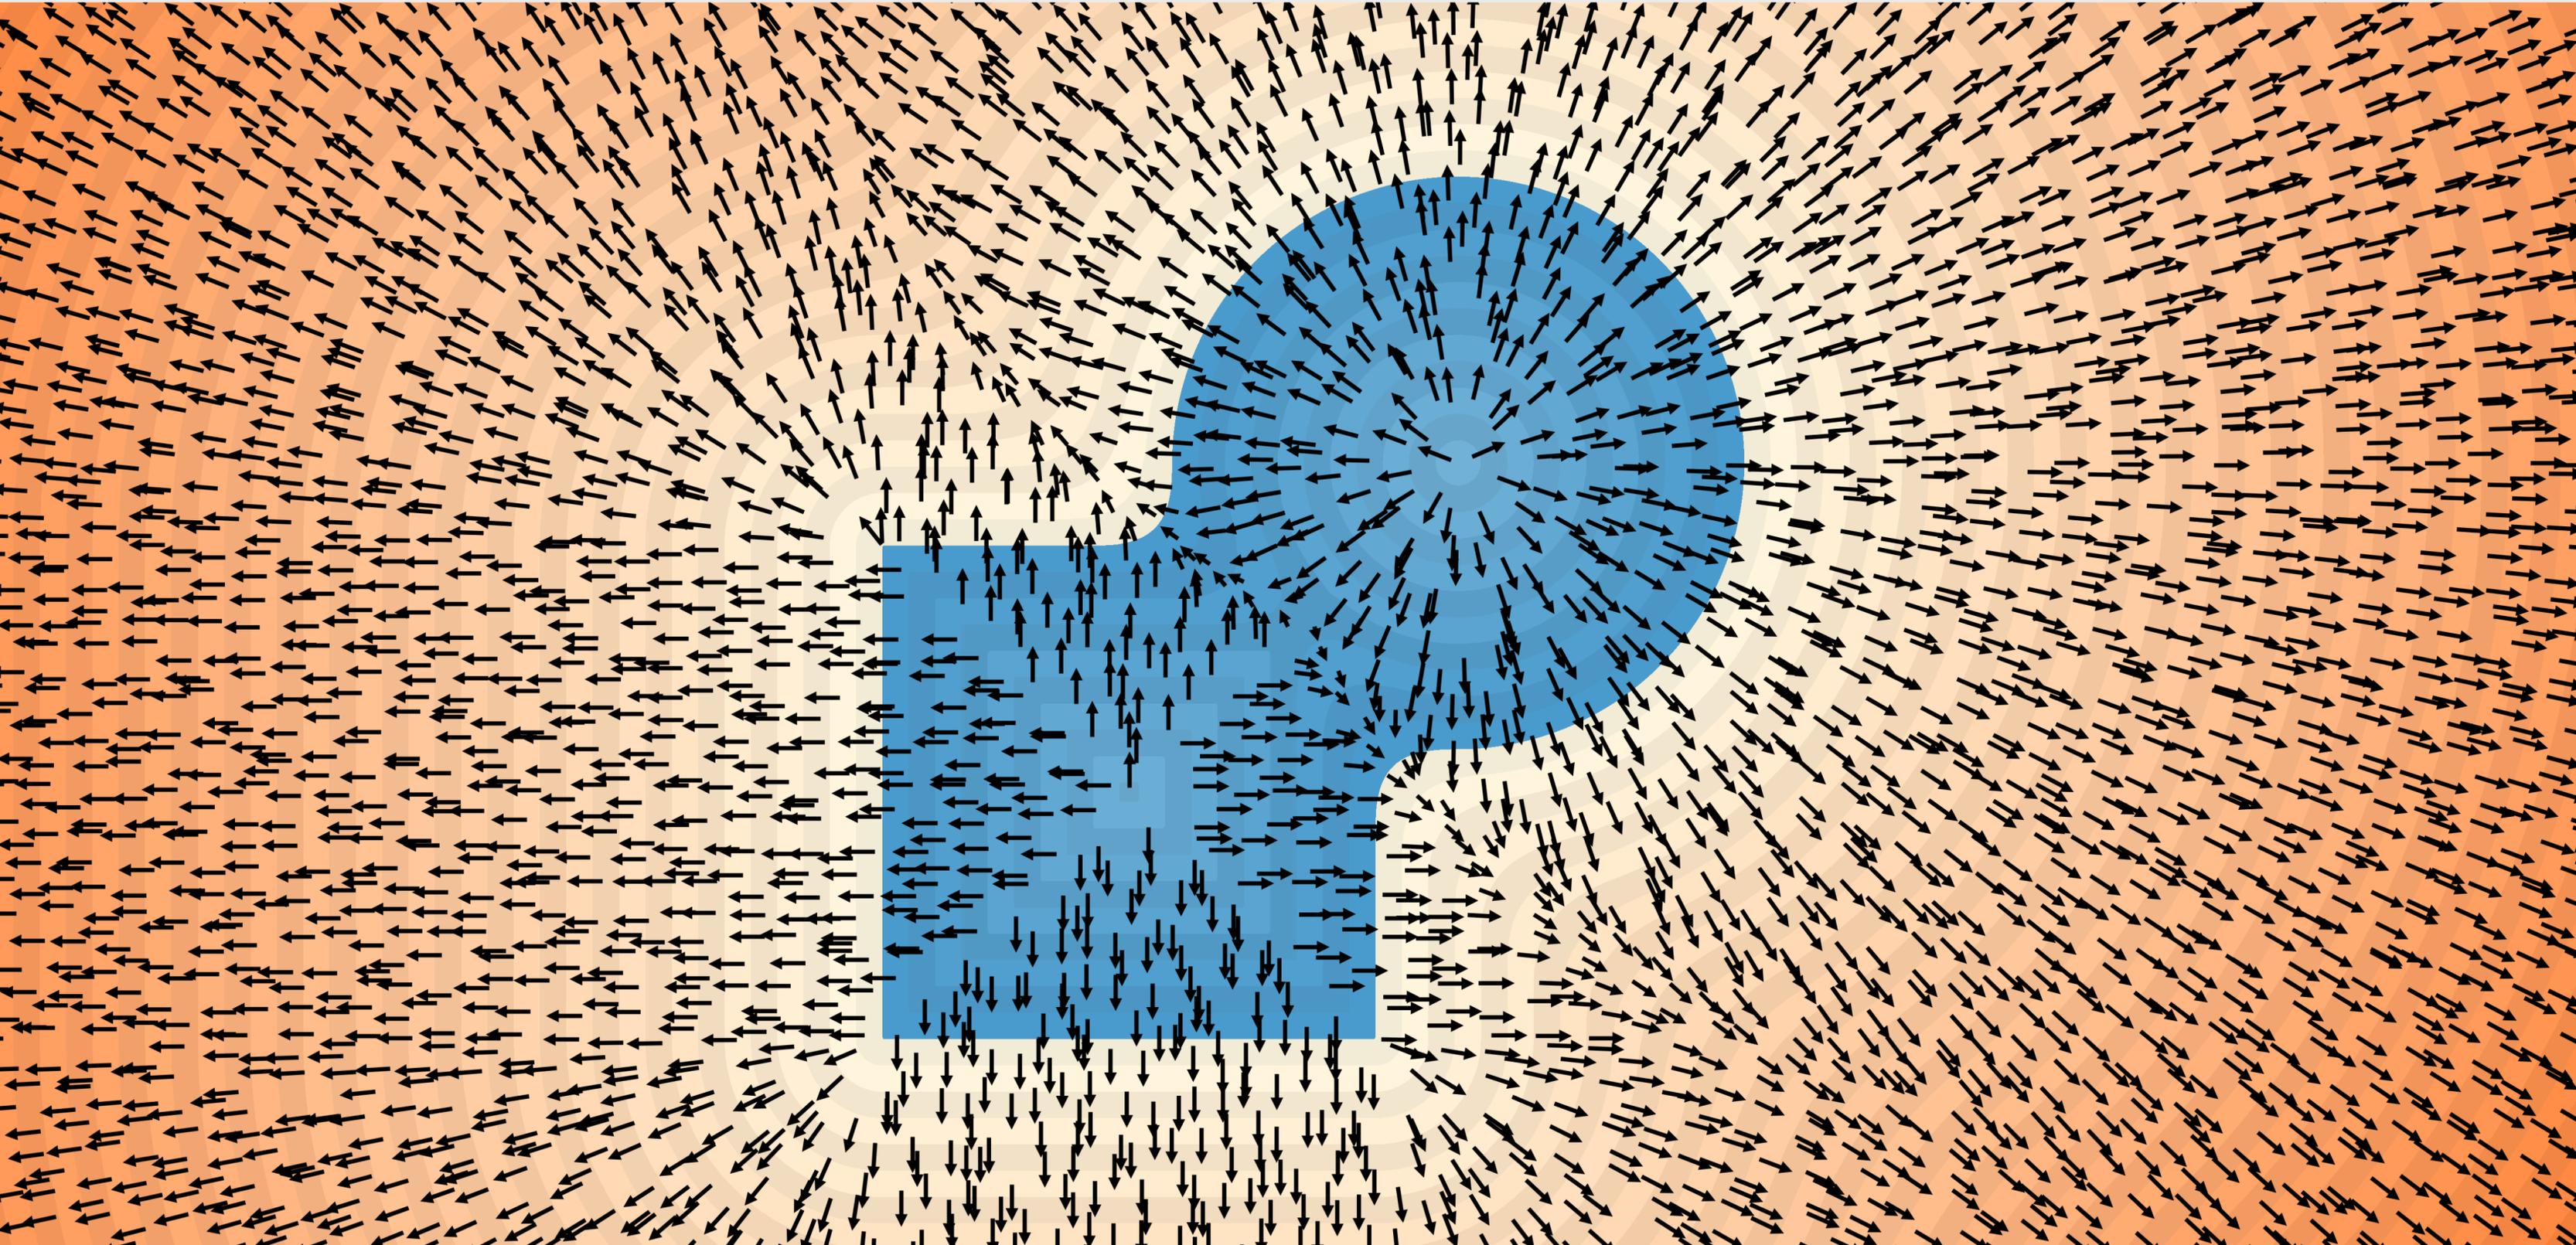
*Eikonal Loss*

$$E_{\text{eik}} = (\|\nabla f_{\theta}\| - 1)^2$$

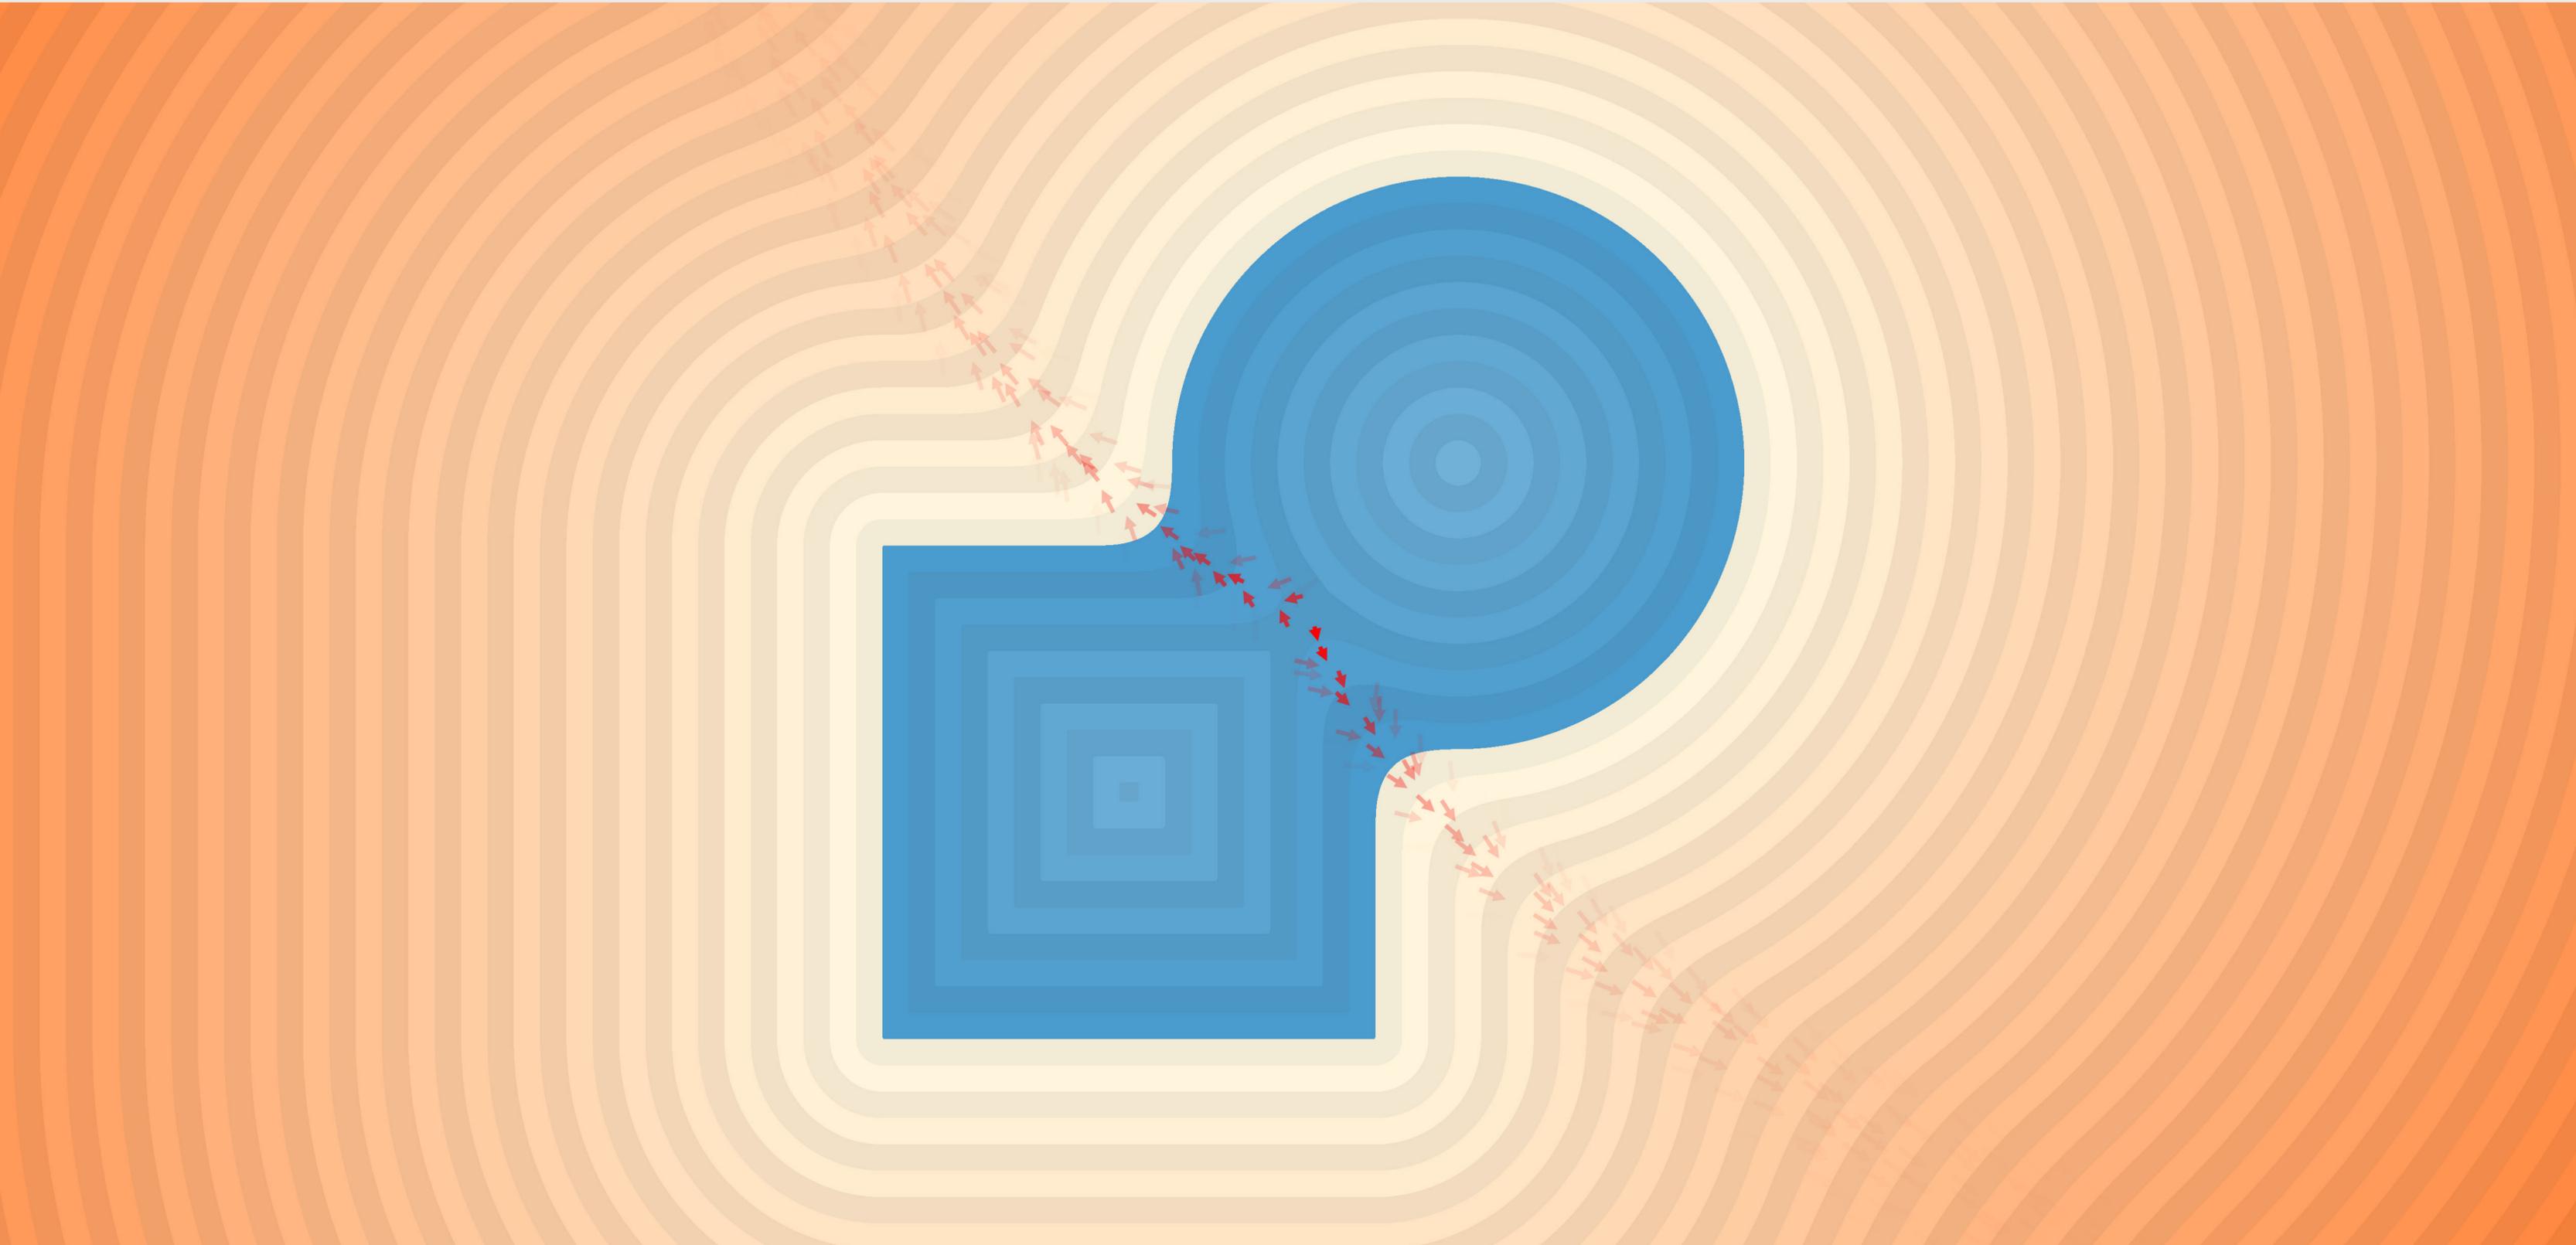


[Gropp et. al, 2020]

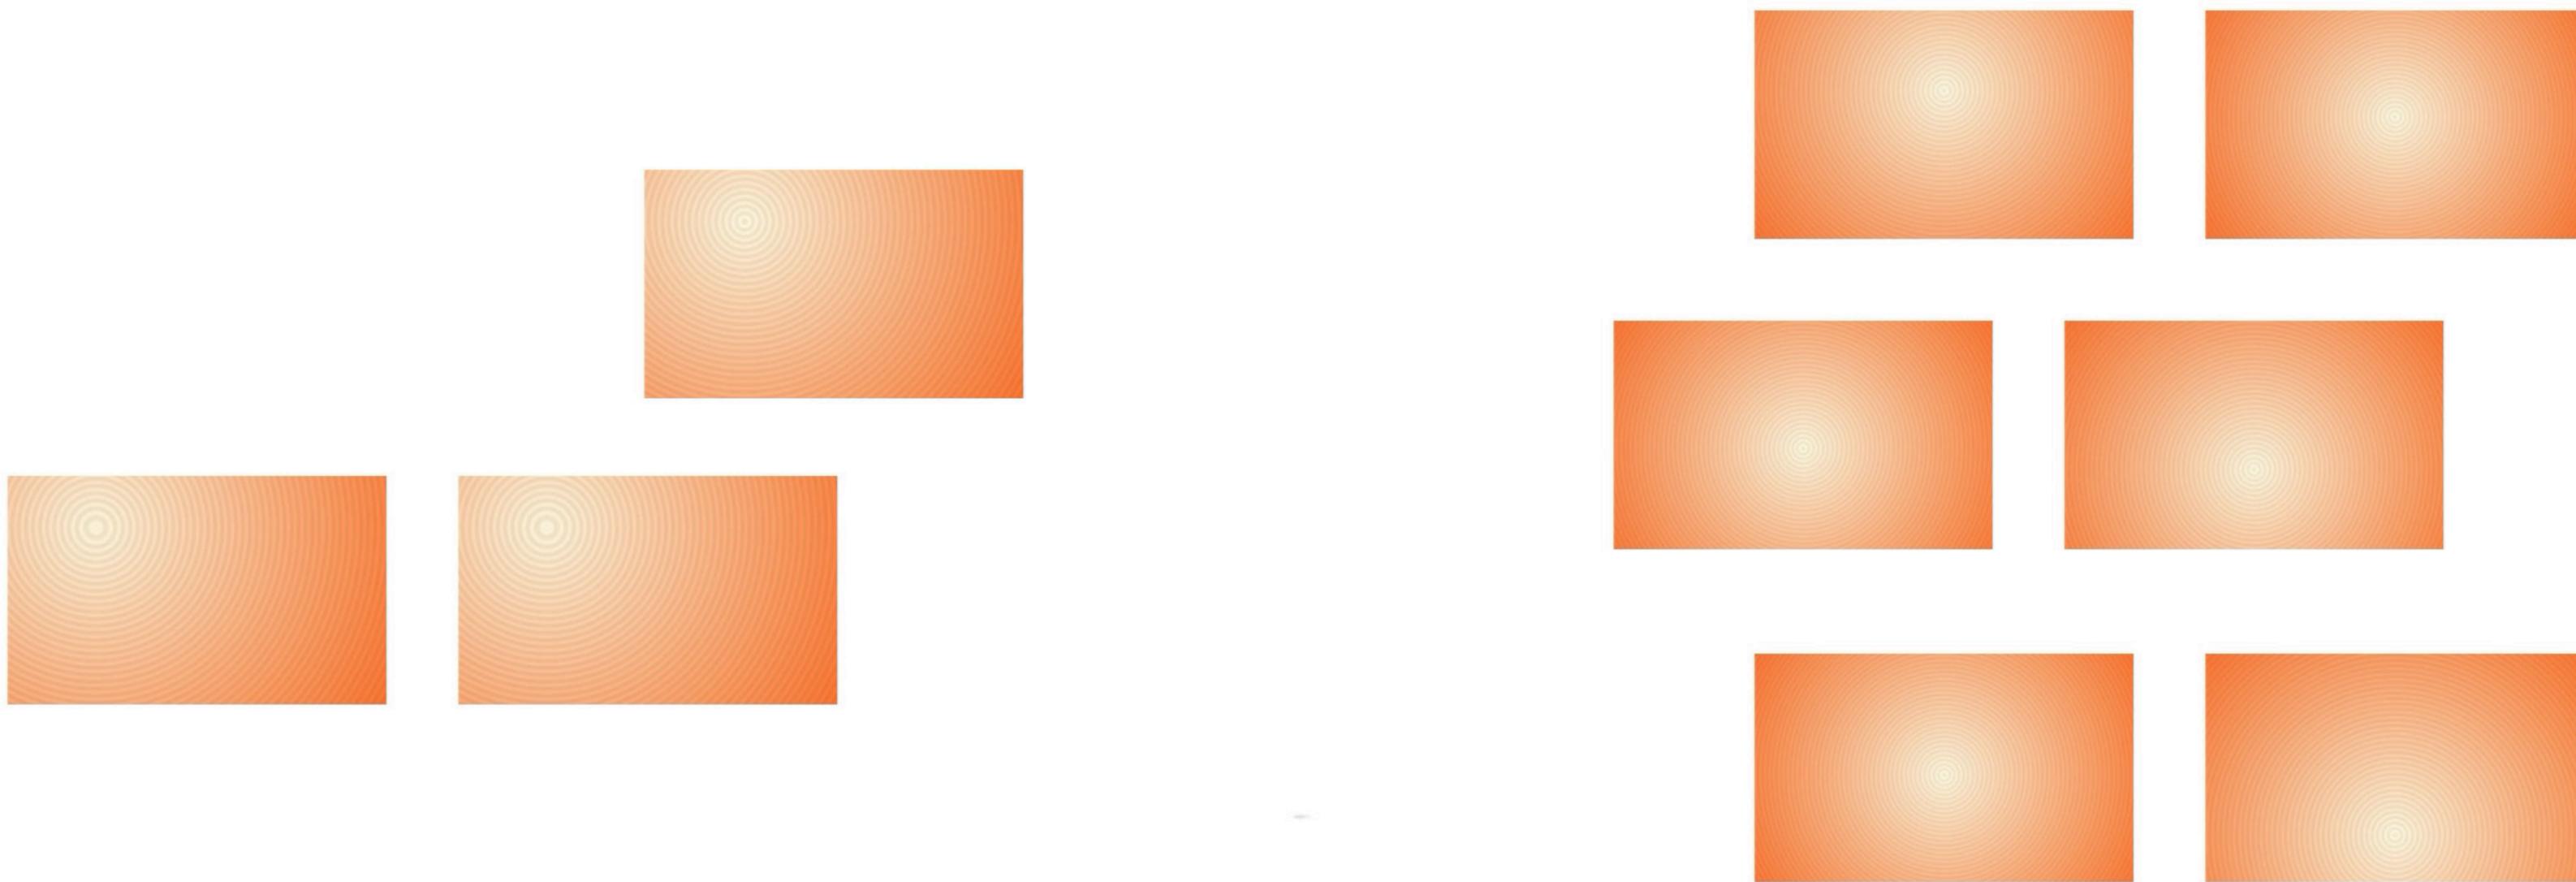
# Eikonal loss: walkthrough



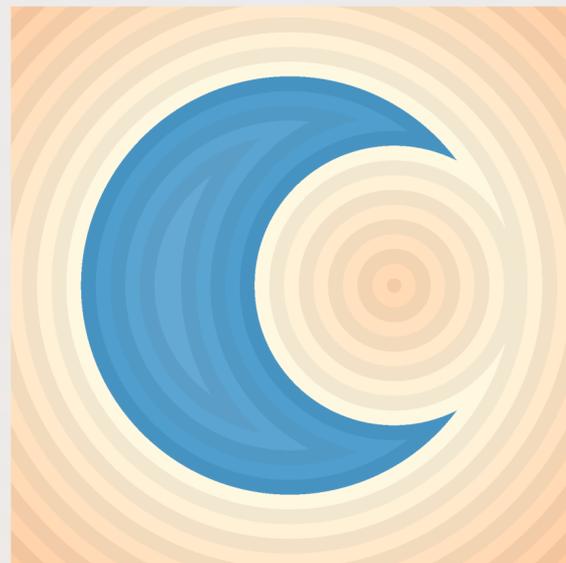
# Eikonal loss: walkthrough



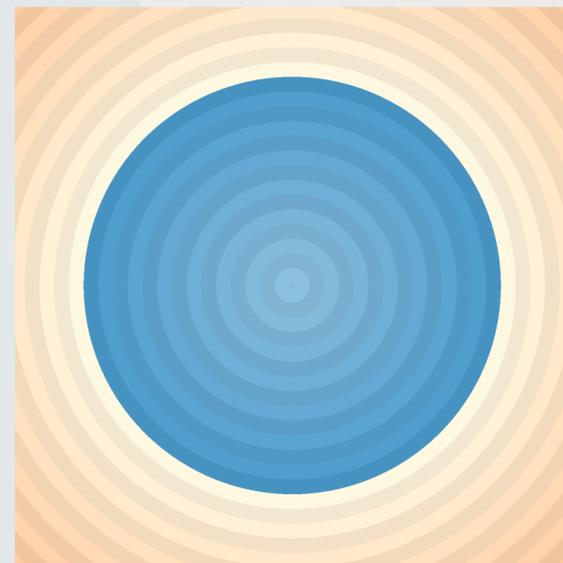
# Eikonal but not SDF



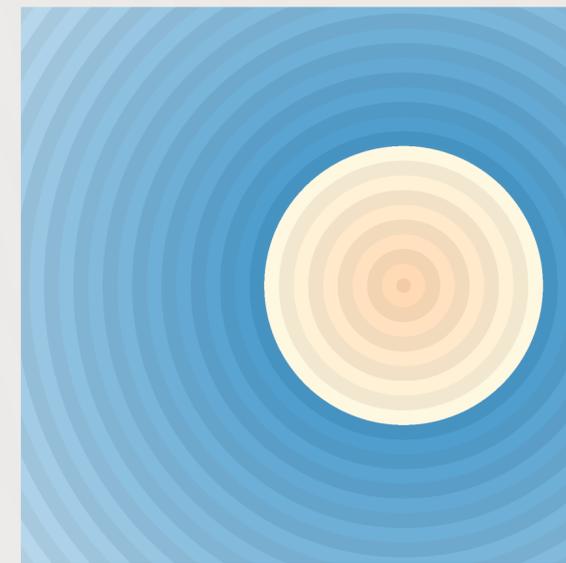
# Eikonal but not SDF



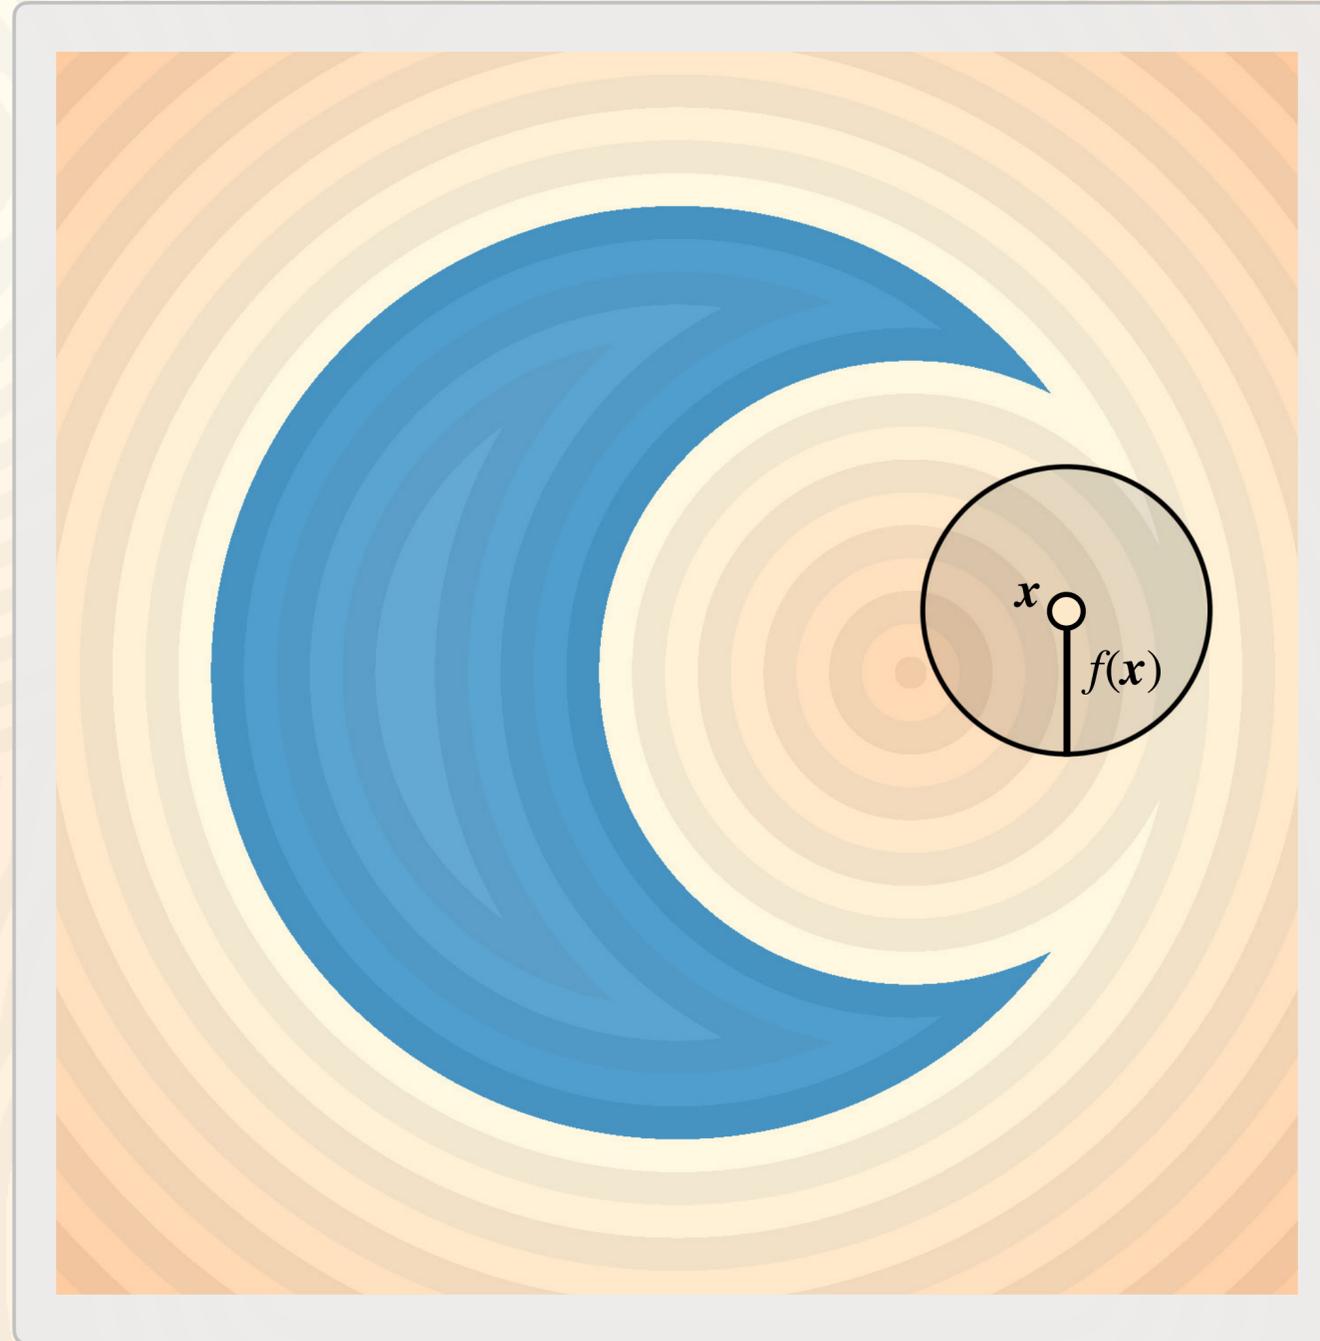
$= \max$



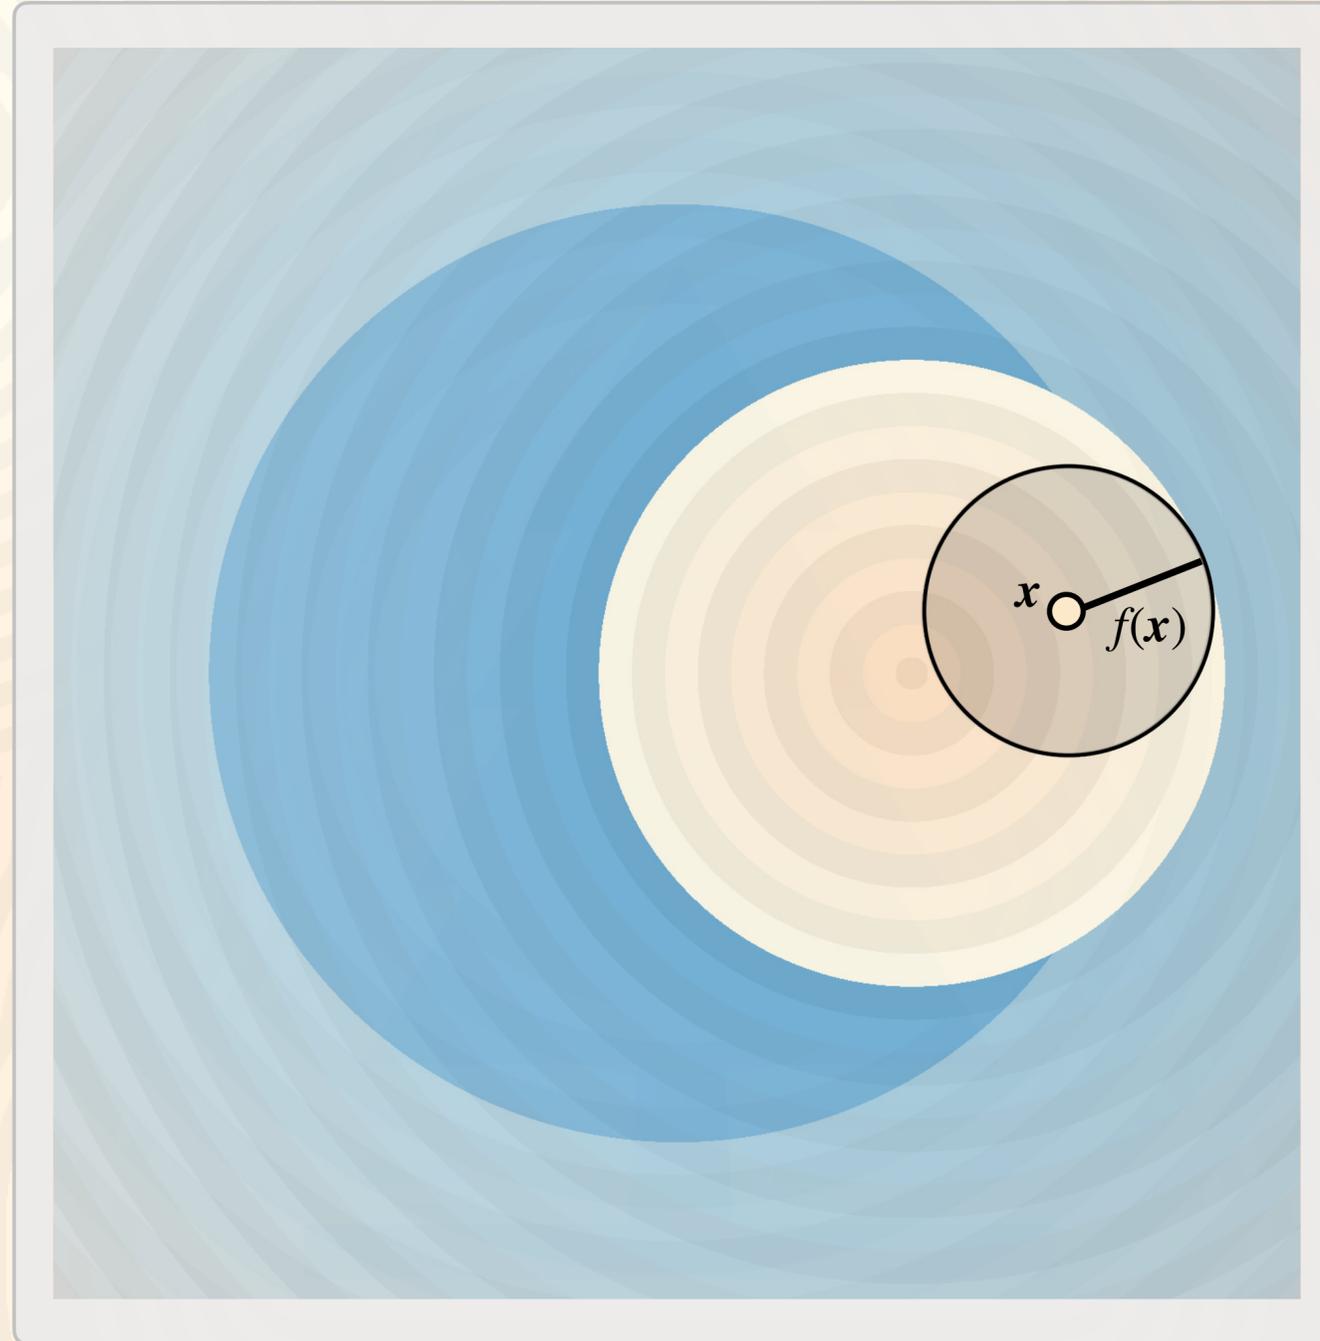
,



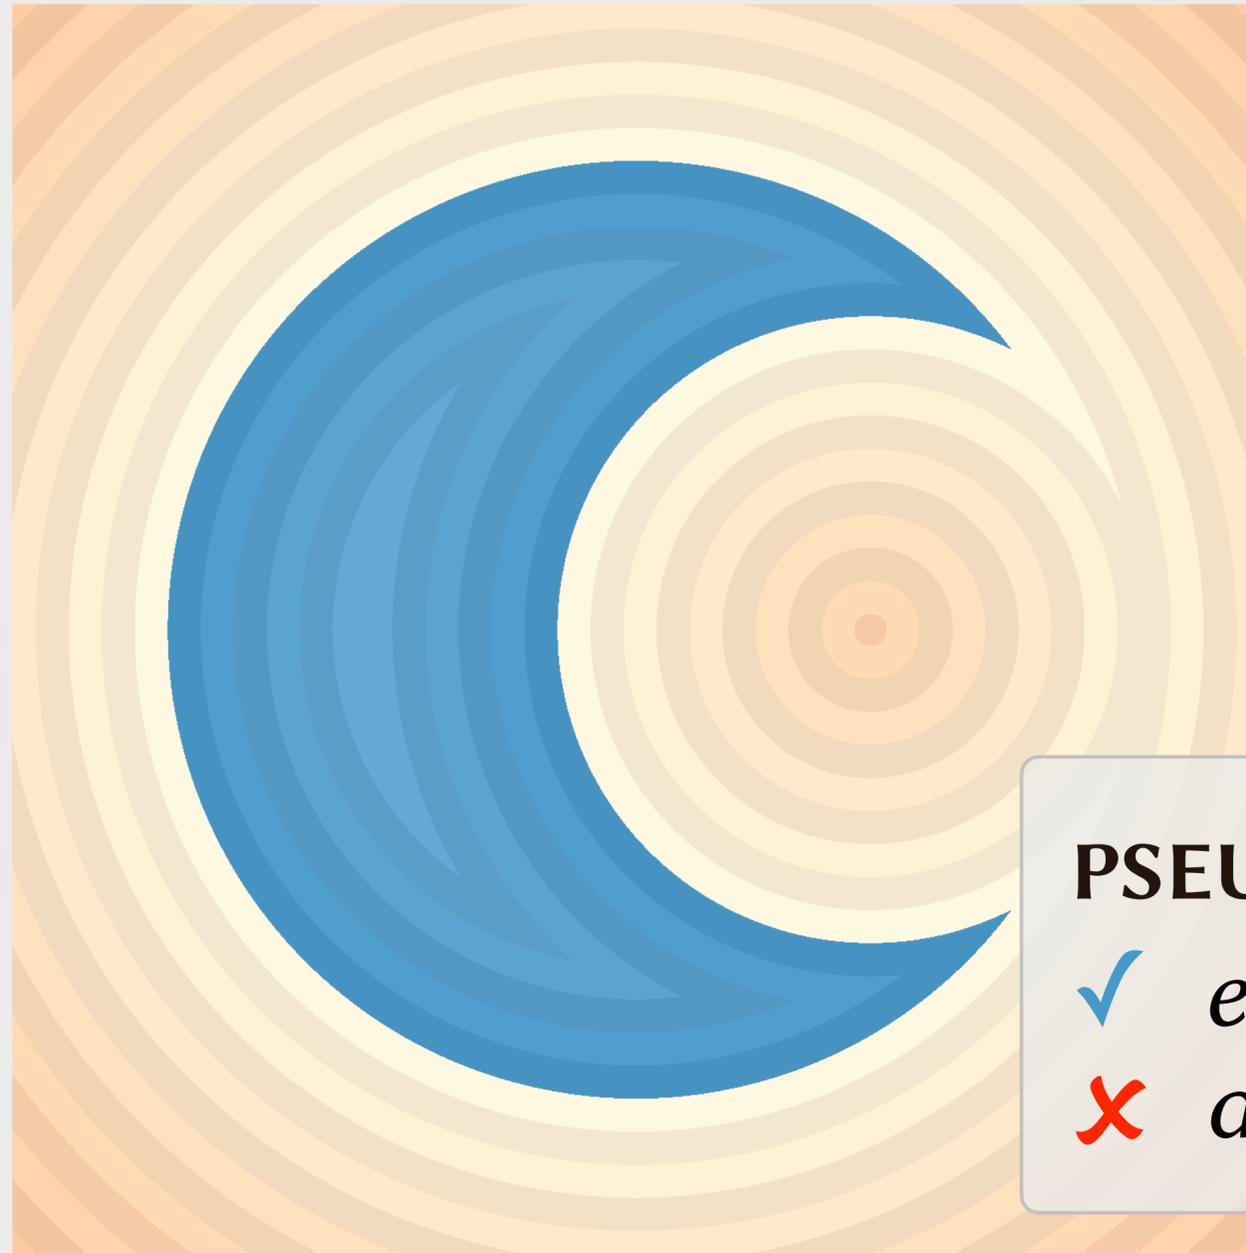
# Eikonal but not SDF



# Eikonal but not SDF



# Eikonal but not SDF



## PSEUDO-SDF

- ✓ *eikonality*
- ✗ *distance property*

# Eikonal loss



# Eikonal loss



Eikonal Loss  
avg. value =  $7.1 \times 10^{-5}$

**Constructive Solid Geometry on Neural Signed Distance Fields**

Zoë Marschner  
zoem@cmu.edu  
Massachusetts Institute of Technology  
Carnegie Mellon University  
United States

Silvia Sellán  
sgsellan@cs.toronto.edu  
University of Toronto  
Canada

Hsueh-Ti Derek Liu  
hsuehtil@gmail.com  
Roblox Research  
University of Toronto  
Canada

Alec Jacobson  
jacobson@cs.toronto.edu  
University of Toronto  
Adobe Research  
Canada



**Figure 1:** Our method allows for the computation of exact neural SDFs of CSG operations. Here, we train one network to learn the swept volume of a stellated dodecahedron shape, parametric over the control points of the cubic Bézier path it is swept along. Specific swept volumes within this parameter space are then unioned together and with cylinders, resulting in a neural implicit which thanks to our regularization term forms an exact SDF of the word “SDF”.

**ABSTRACT**  
Signed Distance Fields (SDFs) parameterized by neural networks have recently gained popularity as a fundamental geometric representation. However, editing the shape encoded by a neural SDF remains an open challenge. A tempting approach is to leverage common geometric operators (e.g., boolean operations), but such edits often lead to incorrect non-SDF outputs (which we call Pseudo-SDFs), preventing them from being used for downstream tasks. In this paper, we characterize the space of Pseudo-SDFs, which are eikonal yet not true distance functions, and derive the *closest point loss*, a novel regularizer that encourages the output to be an exact SDF. We demonstrate the applicability of our regularization to many operations in which traditional methods cause a Pseudo-SDF to

arise, such as CSG and swept volumes, and produce a true (neural) SDF for the result of these operations.

**CCS CONCEPTS**  
• Computing methodologies → Shape modeling; Shape representations.

**KEYWORDS**  
signed distance field, neural implicit, CSG, swept volumes

**ACM Reference Format:**  
Zoë Marschner, Silvia Sellán, Hsueh-Ti Derek Liu, and Alec Jacobson. 2023. Constructive Solid Geometry on Neural Signed Distance Fields. In *SIGGRAPH Asia 2023 Conference Papers (SA Conference Papers '23)*, December 12–15, 2023, Sydney, NSW, Australia. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3610548.3618170>

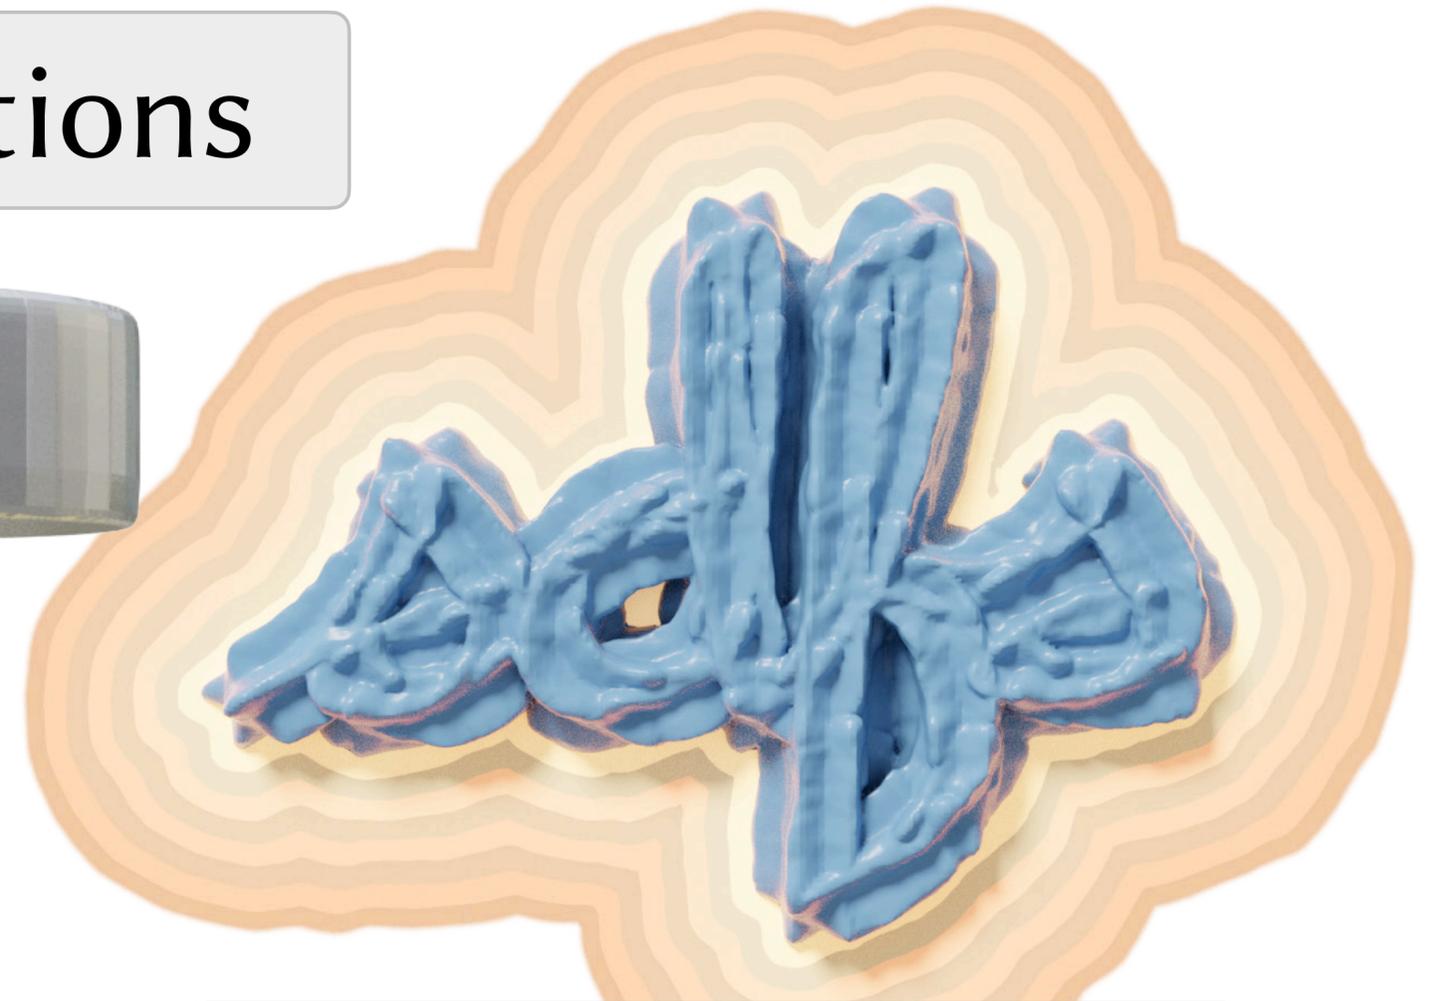
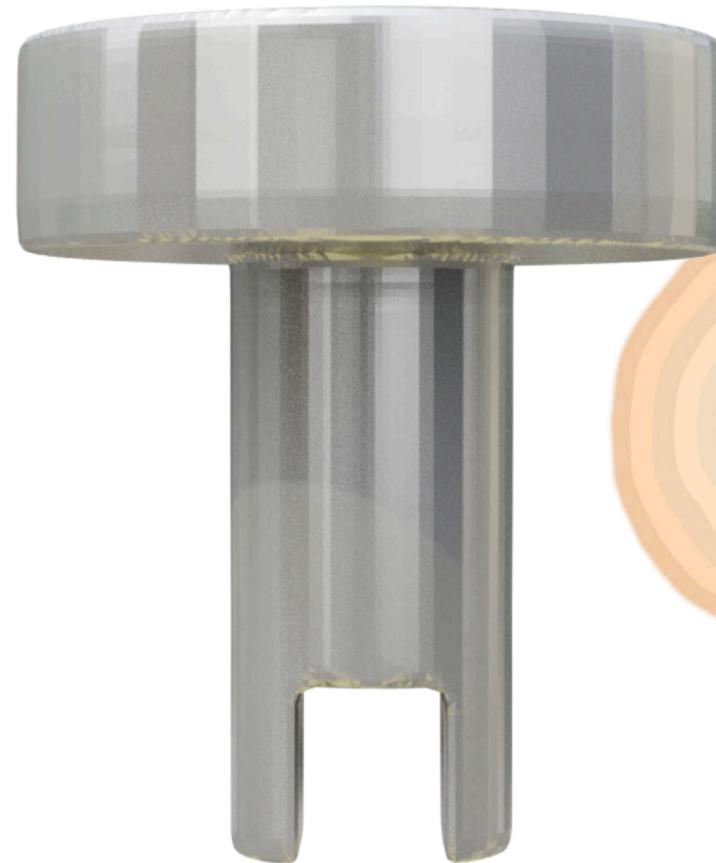
**1 INTRODUCTION**  
*Neural implicit functions* have gained attention as a fundamental representation of 3D objects due to their state-of-the-art performance in tasks like compression and reconstruction, as well as their generative power. They describe the boundary of a solid shape as



This work is licensed under a Creative Commons Attribution International 4.0 License.

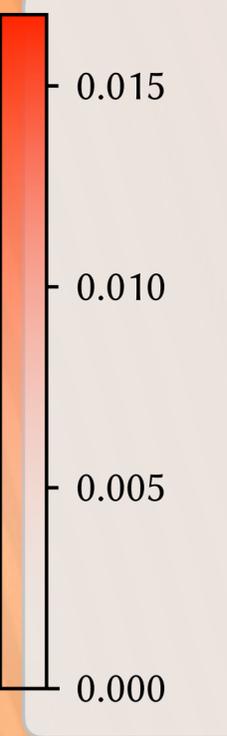
SA Conference Papers '23, December 12–15, 2023, Sydney, NSW, Australia  
© 2023 Copyright held by the owner/author(s).  
ACM ISBN 979-8-4007-0315-7/23/12.  
<https://doi.org/10.1145/3610548.3618170>

## CSG Operations



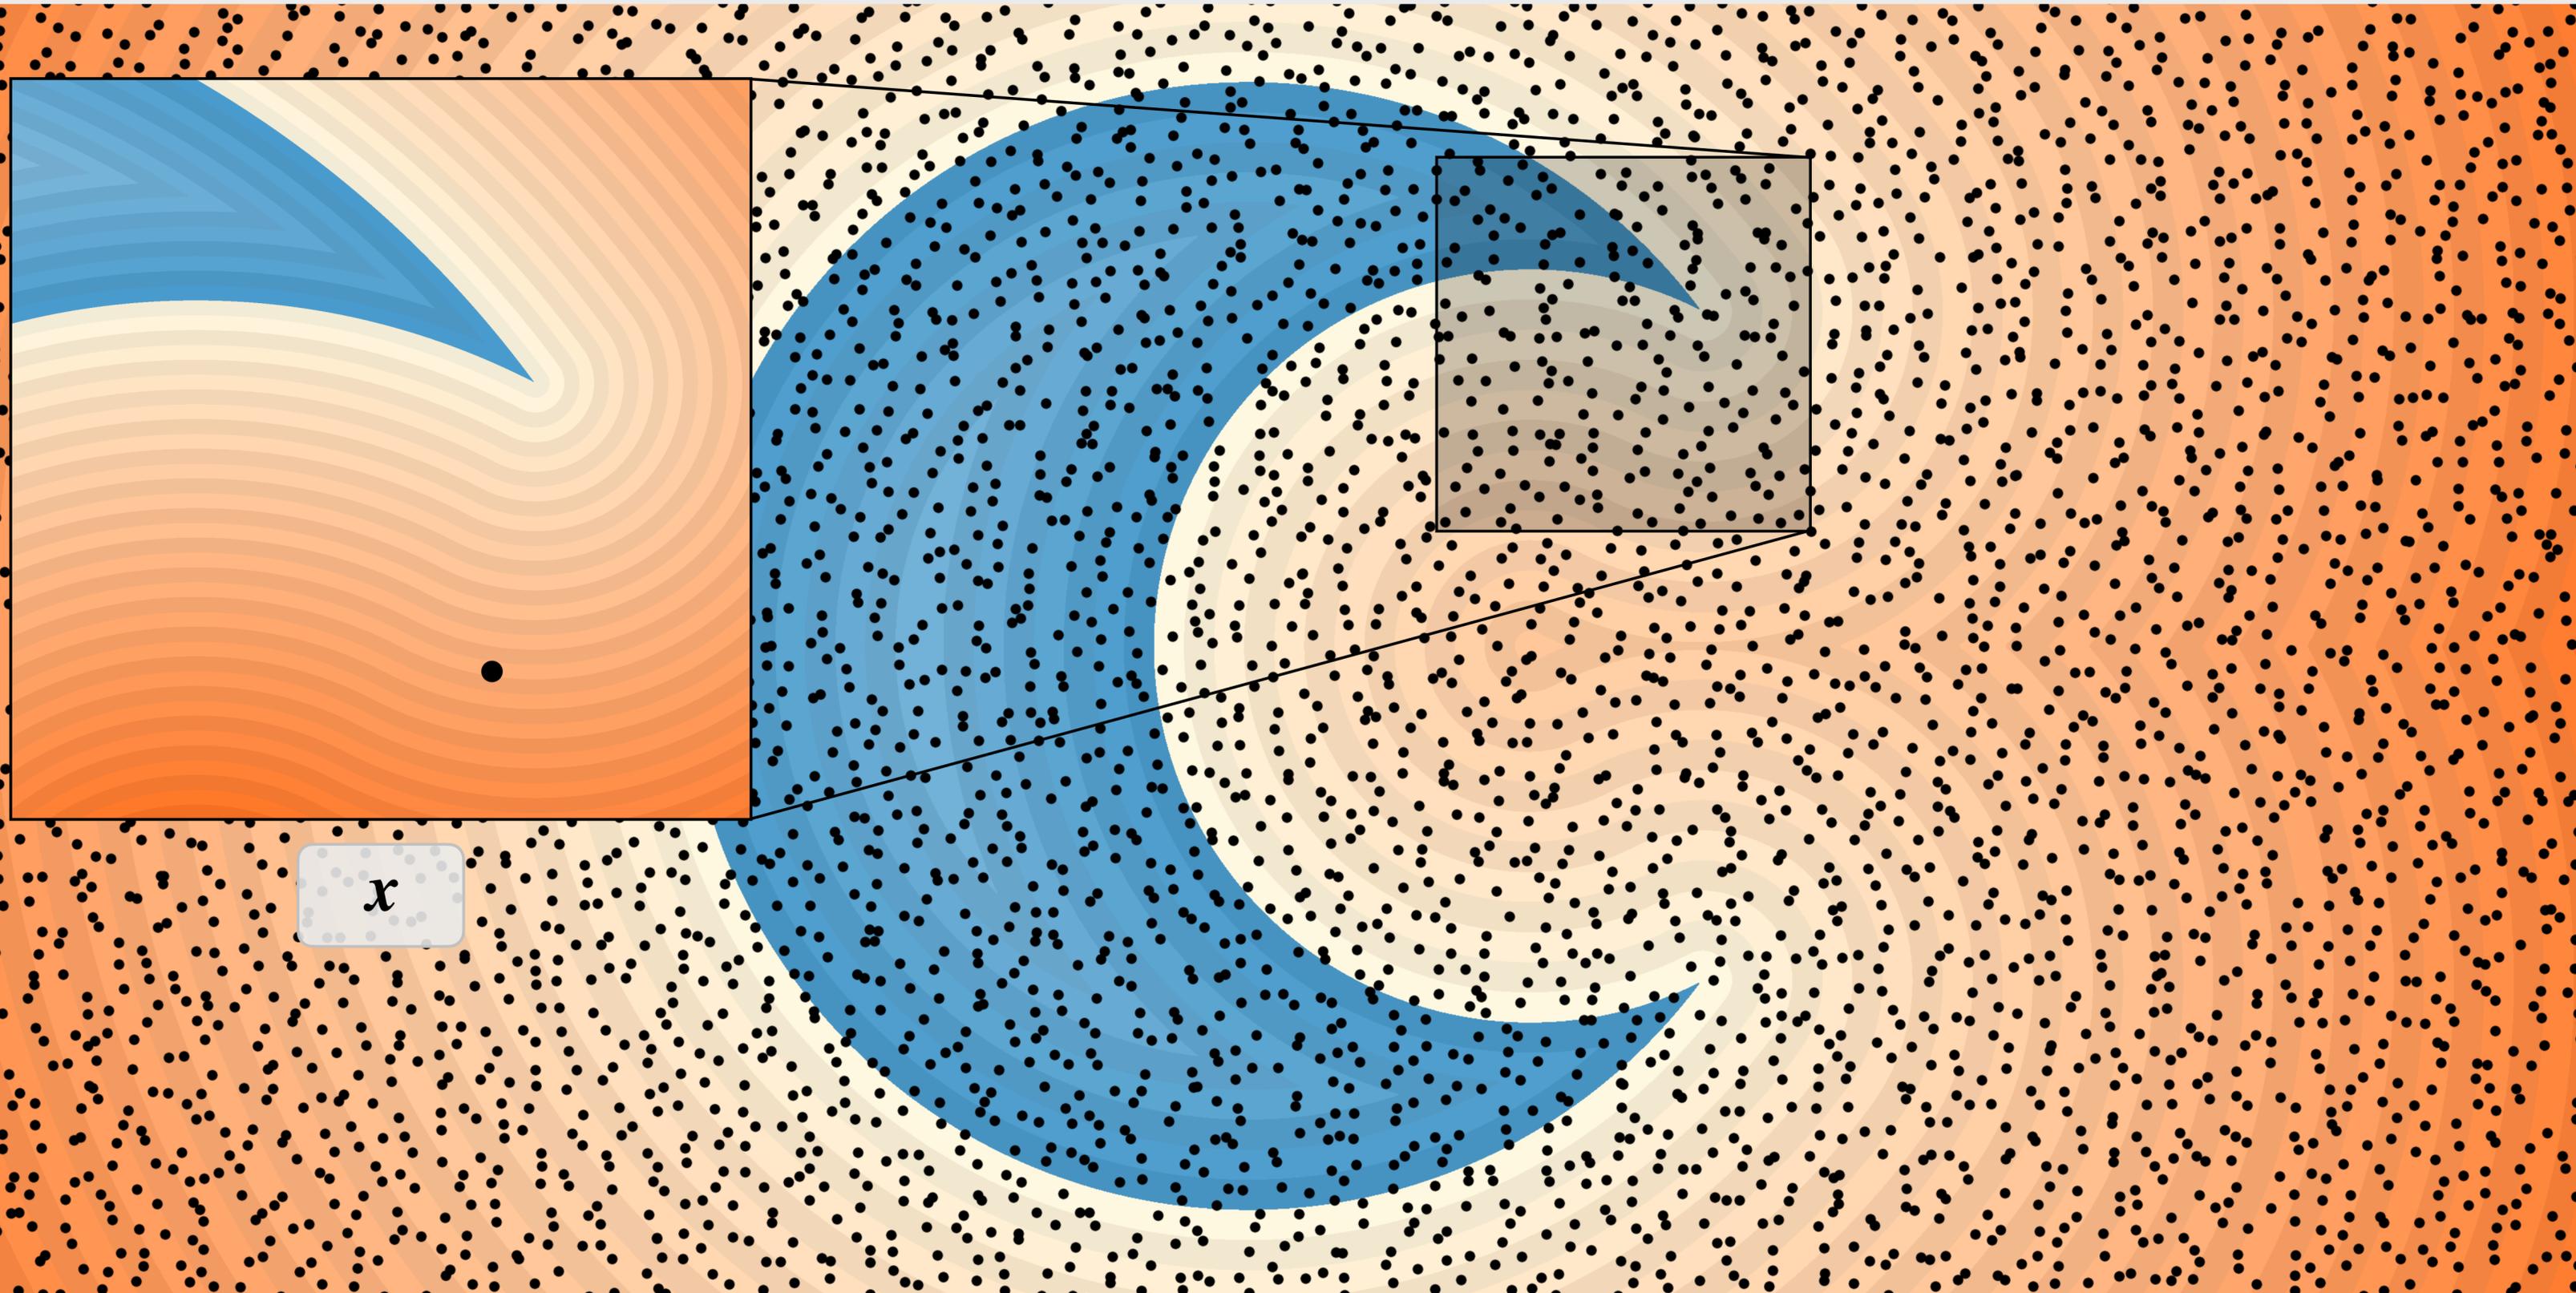
## Swept Volumes

**CSG on Neural SDFs**  
[Marschner et al., 2023]

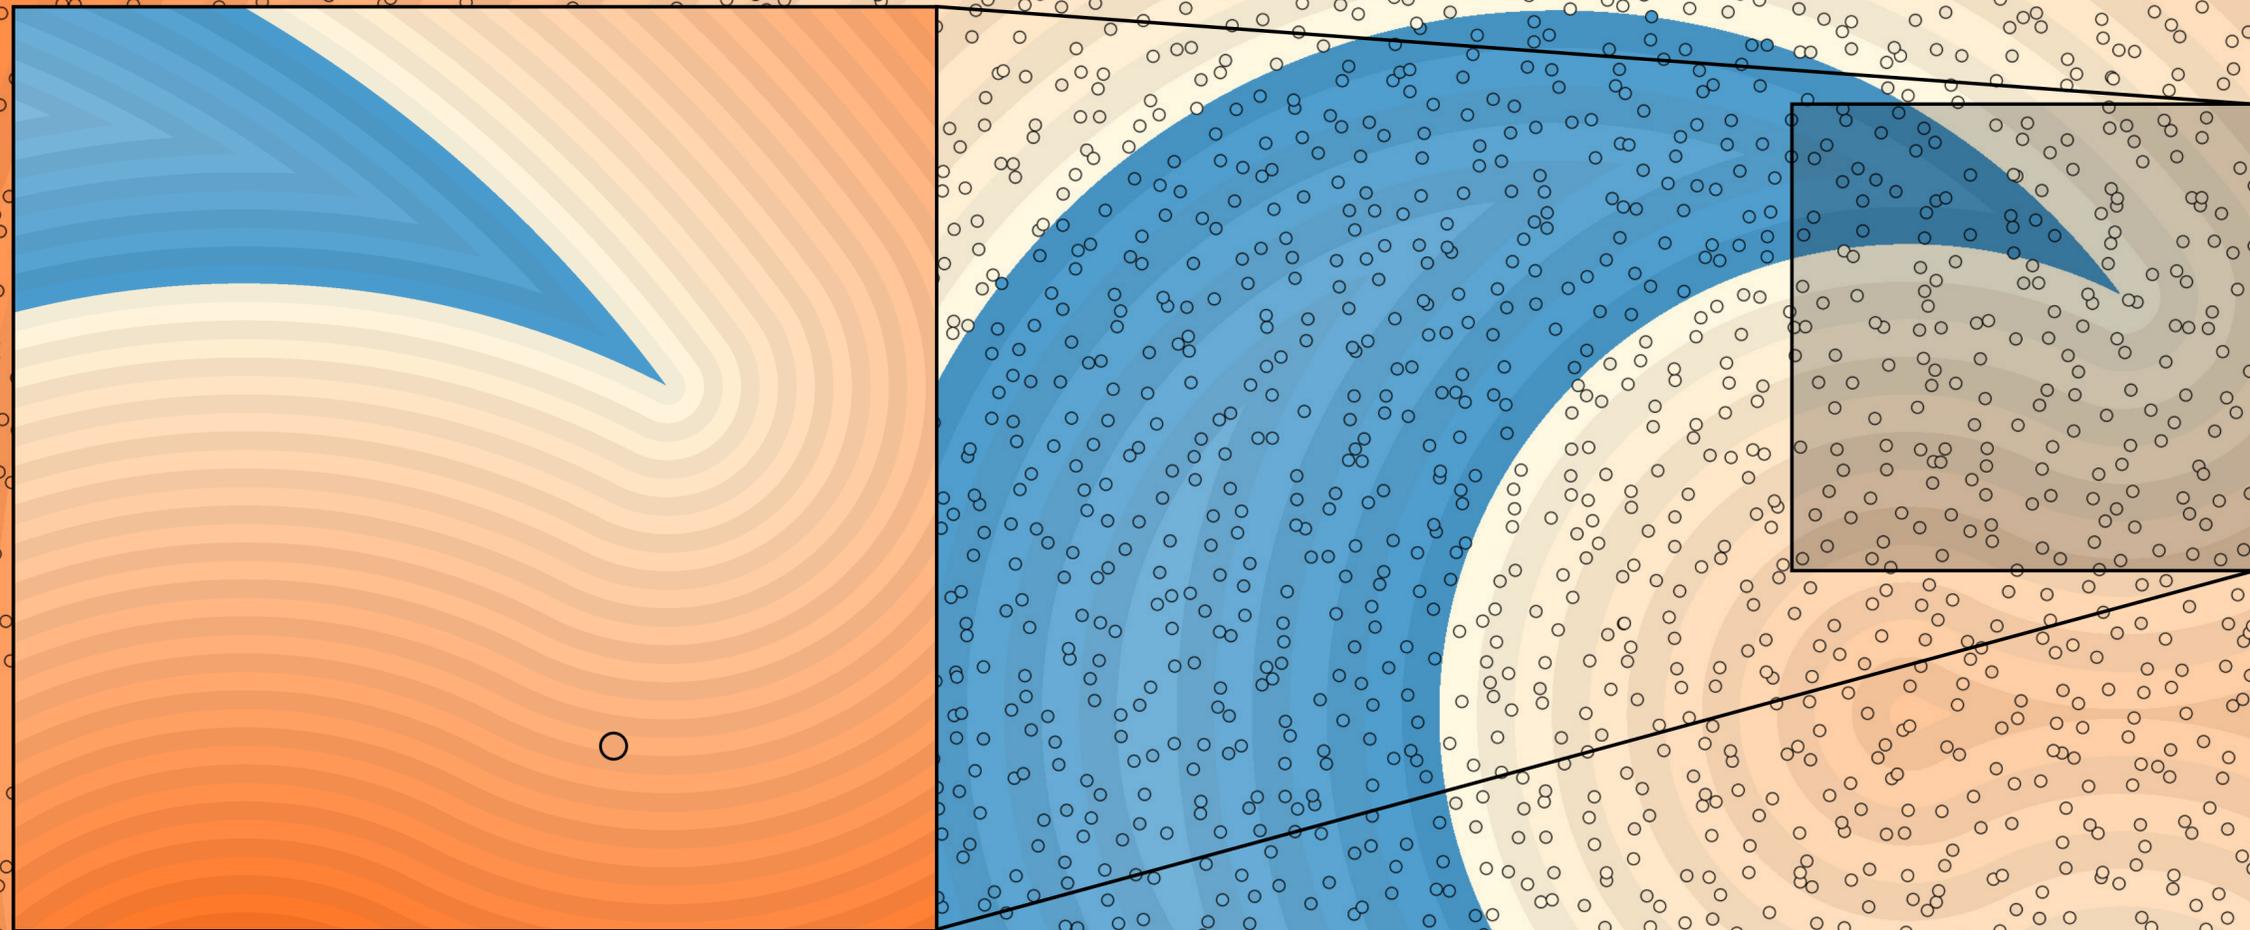


Eikonal Loss  
avg. value =  $7.1 \times 10^{-5}$

# Closest point loss: walkthrough

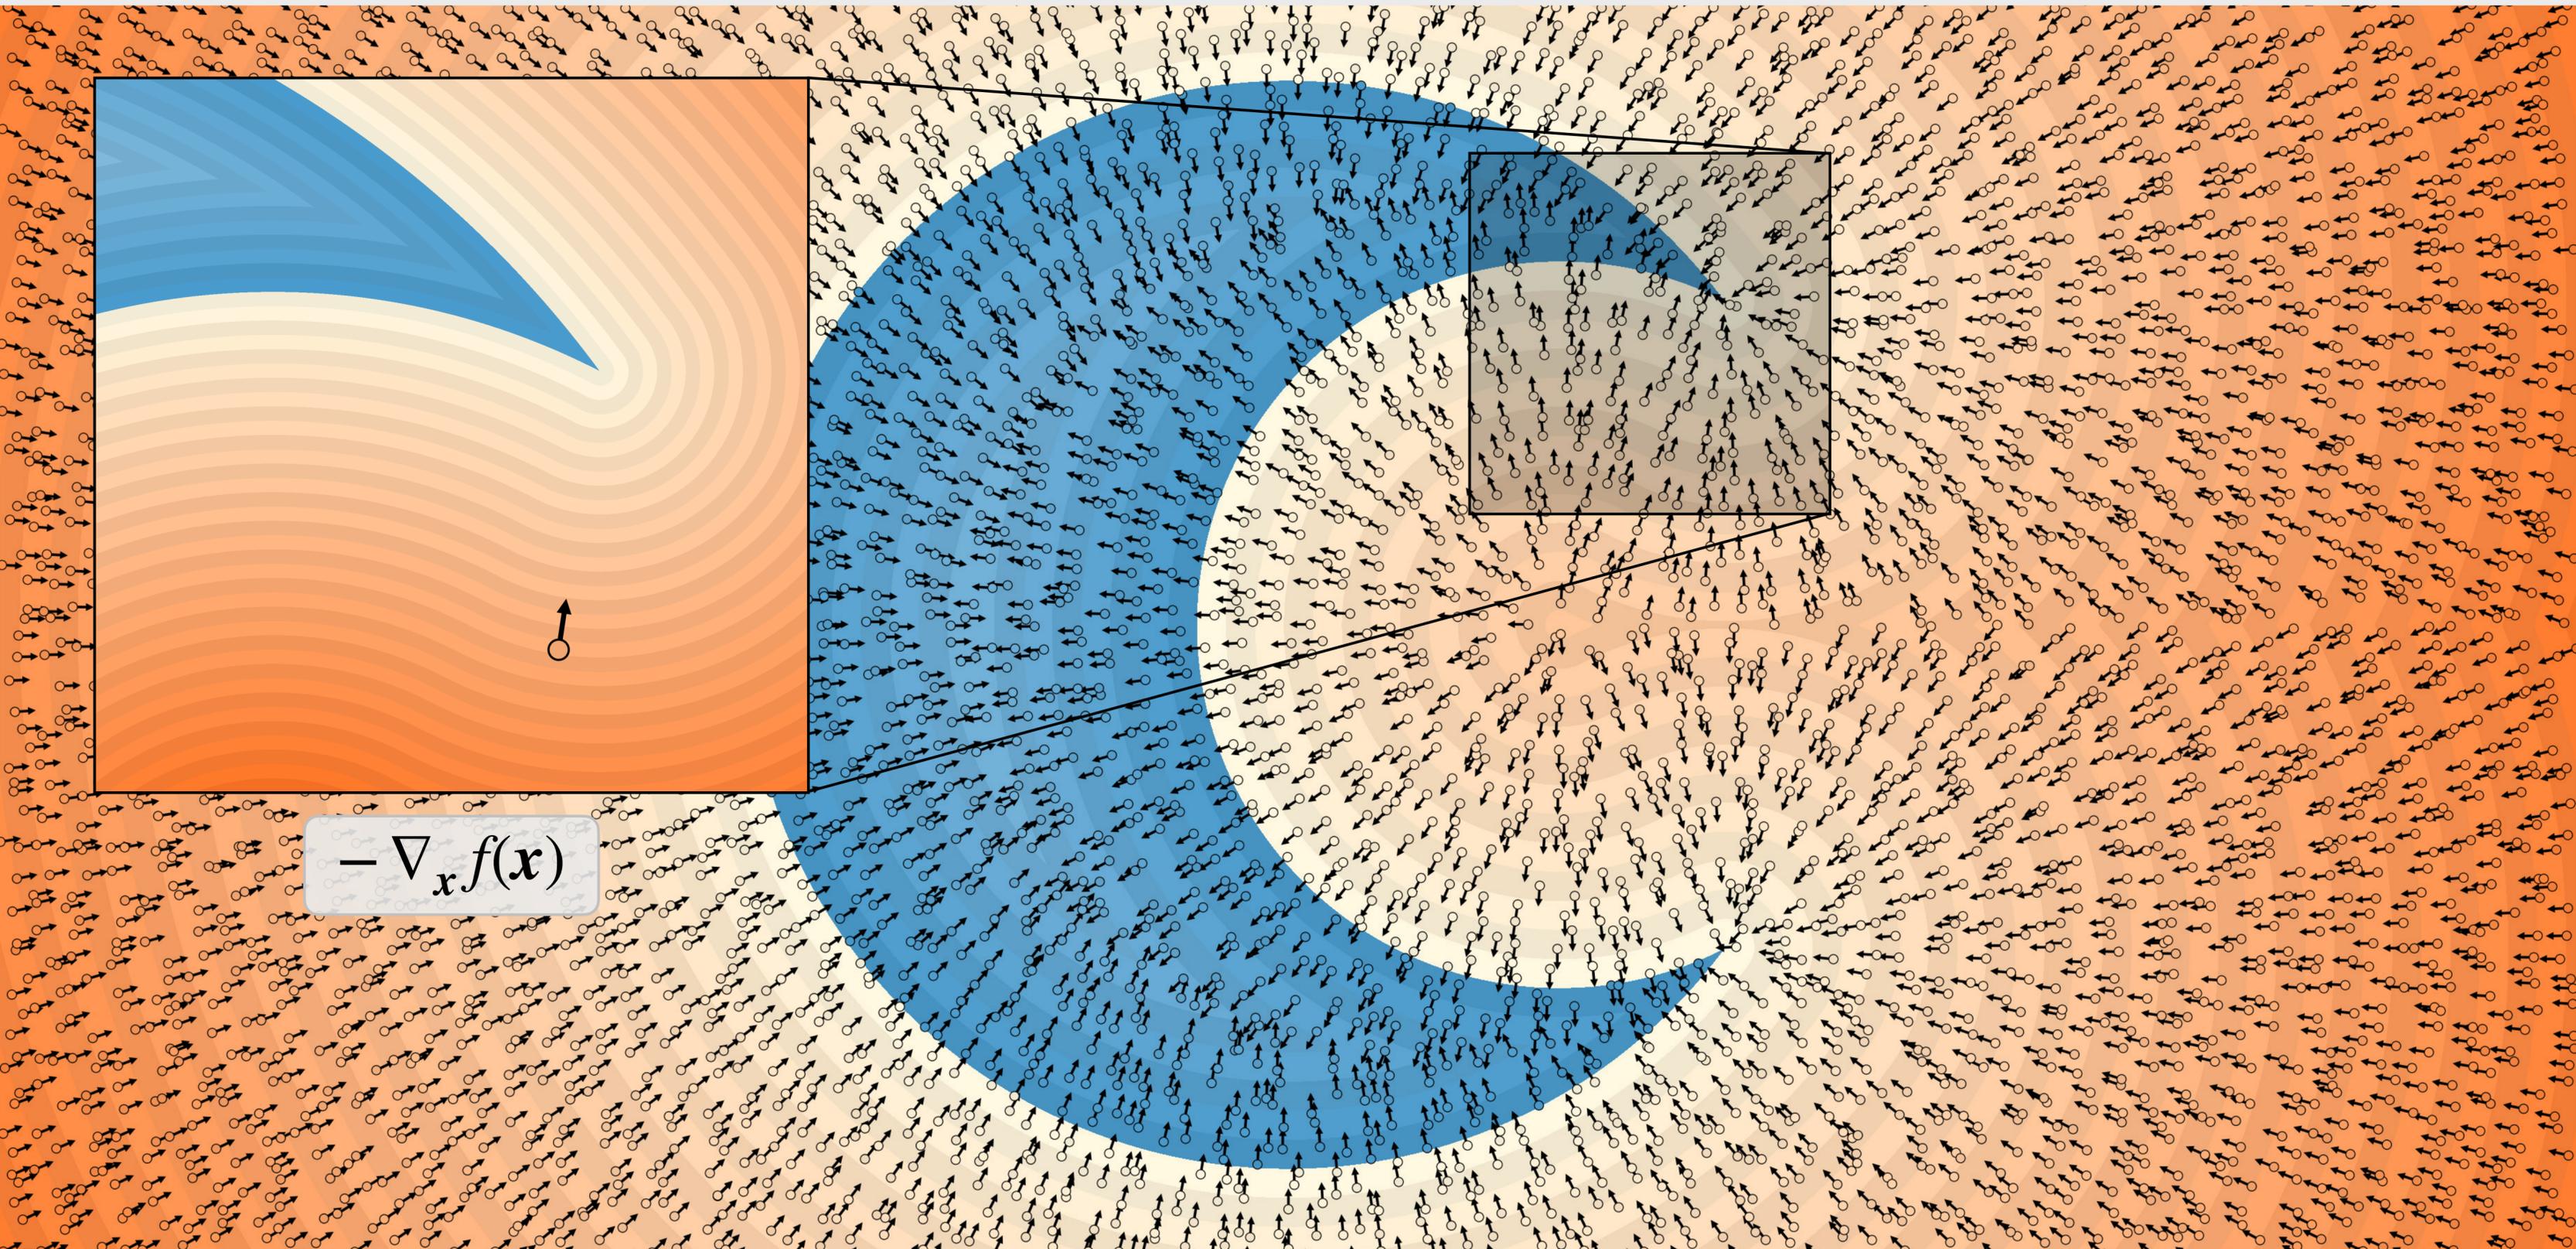


# Closest point loss: walkthrough

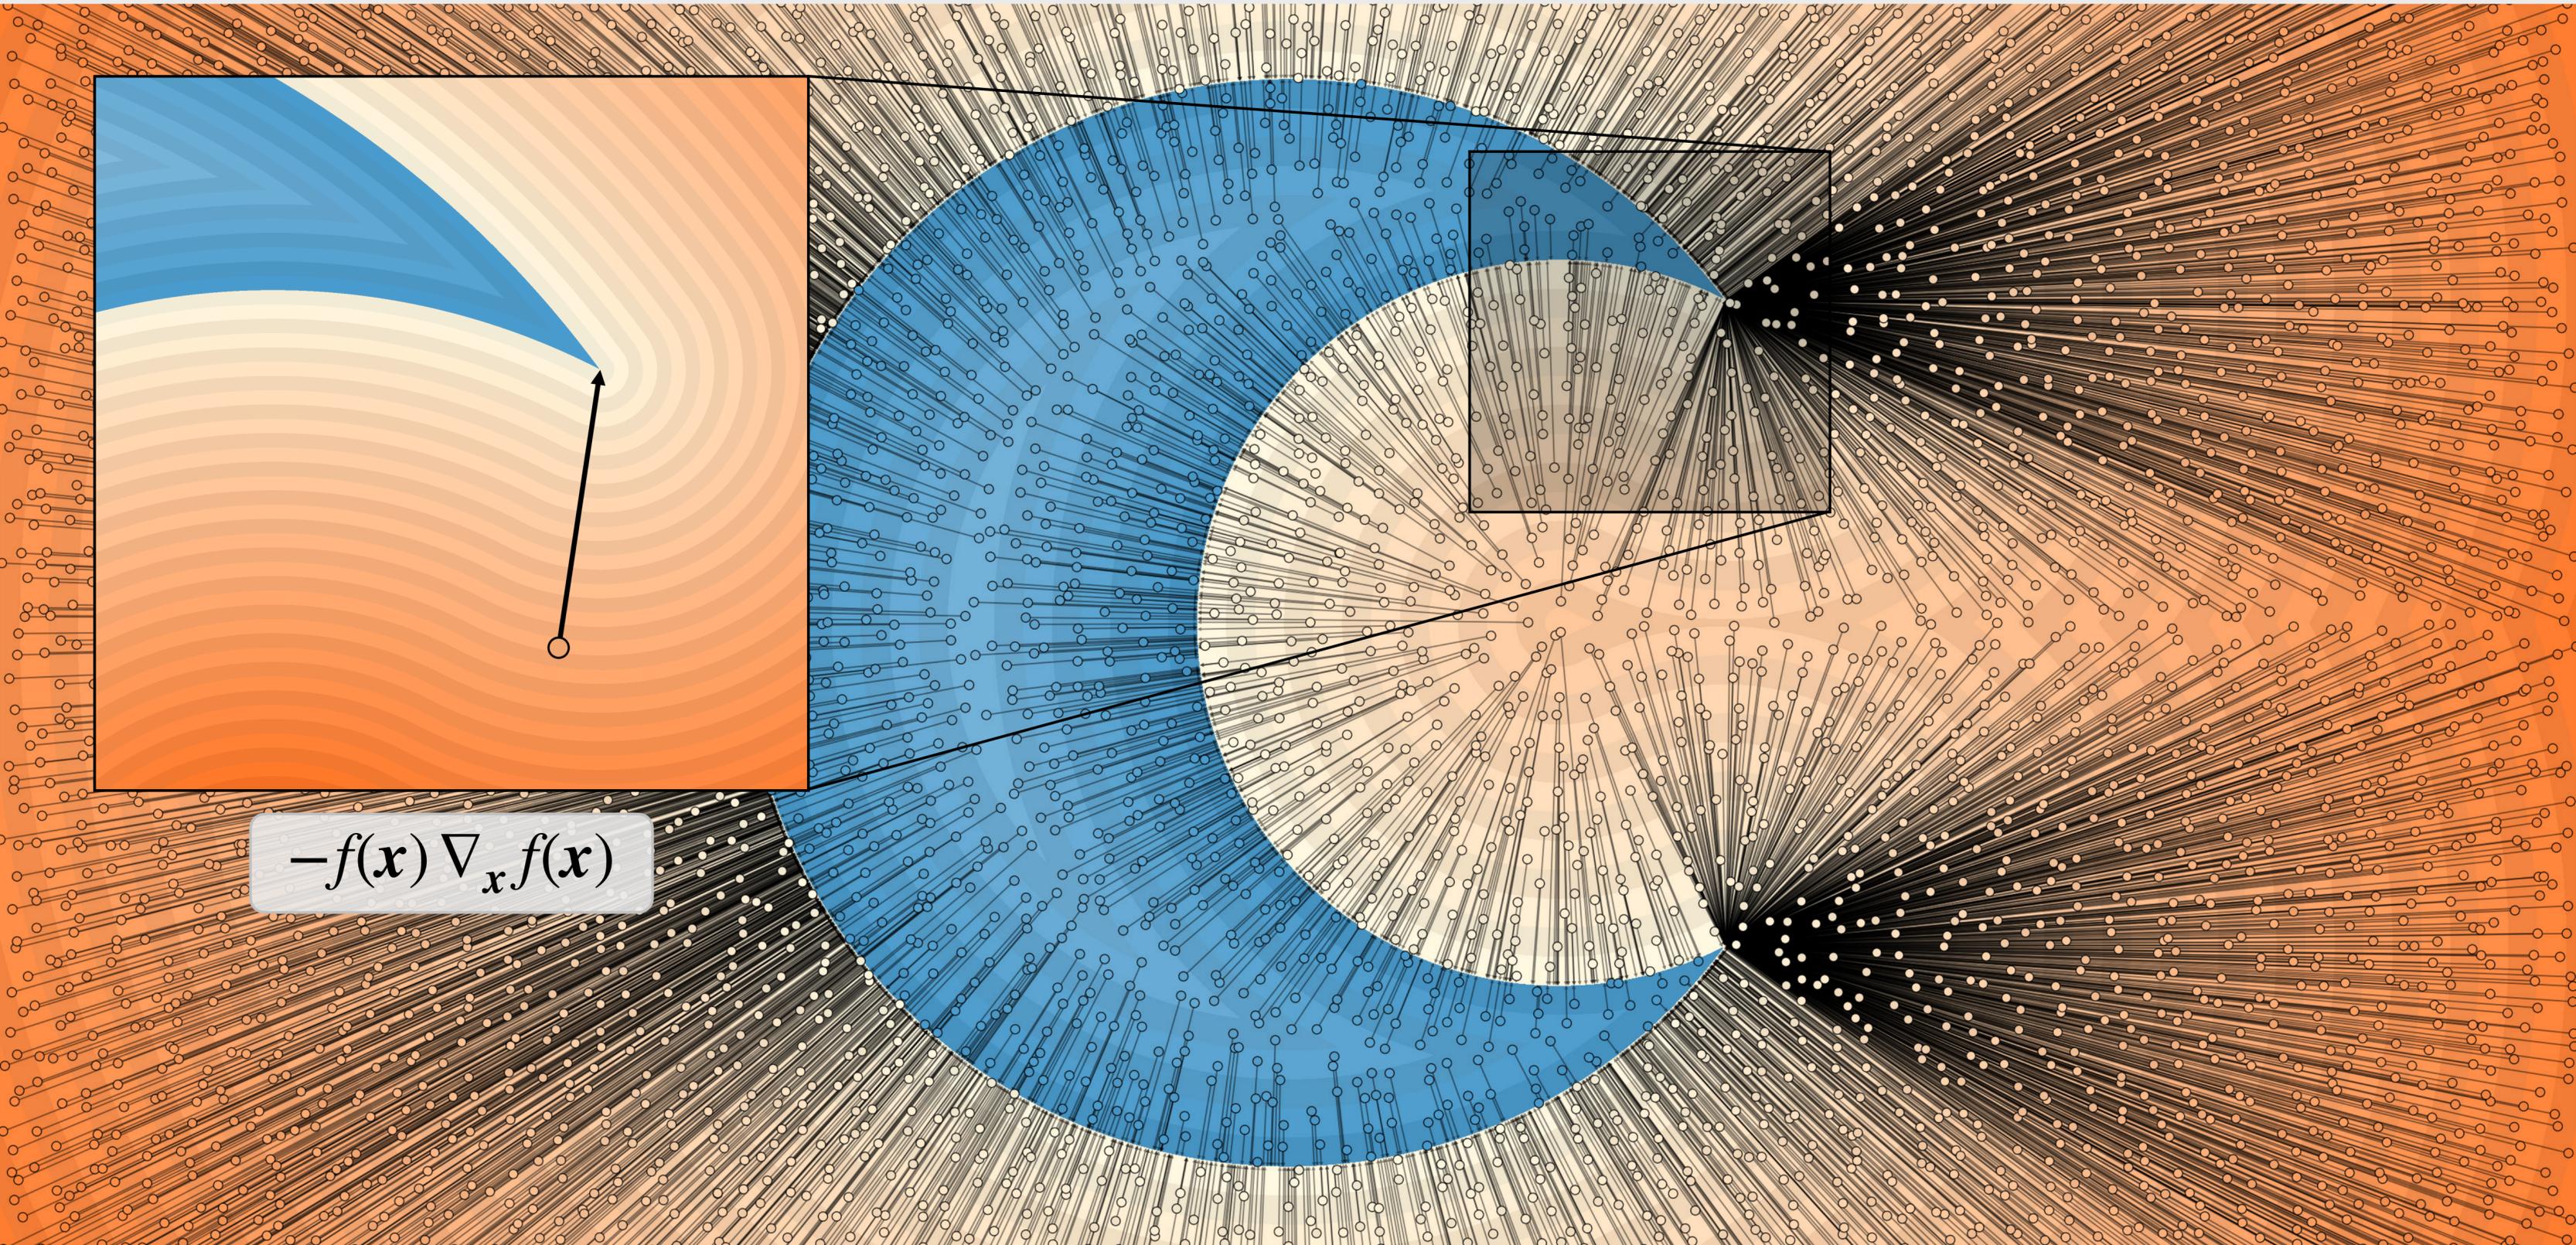


$f(x)$

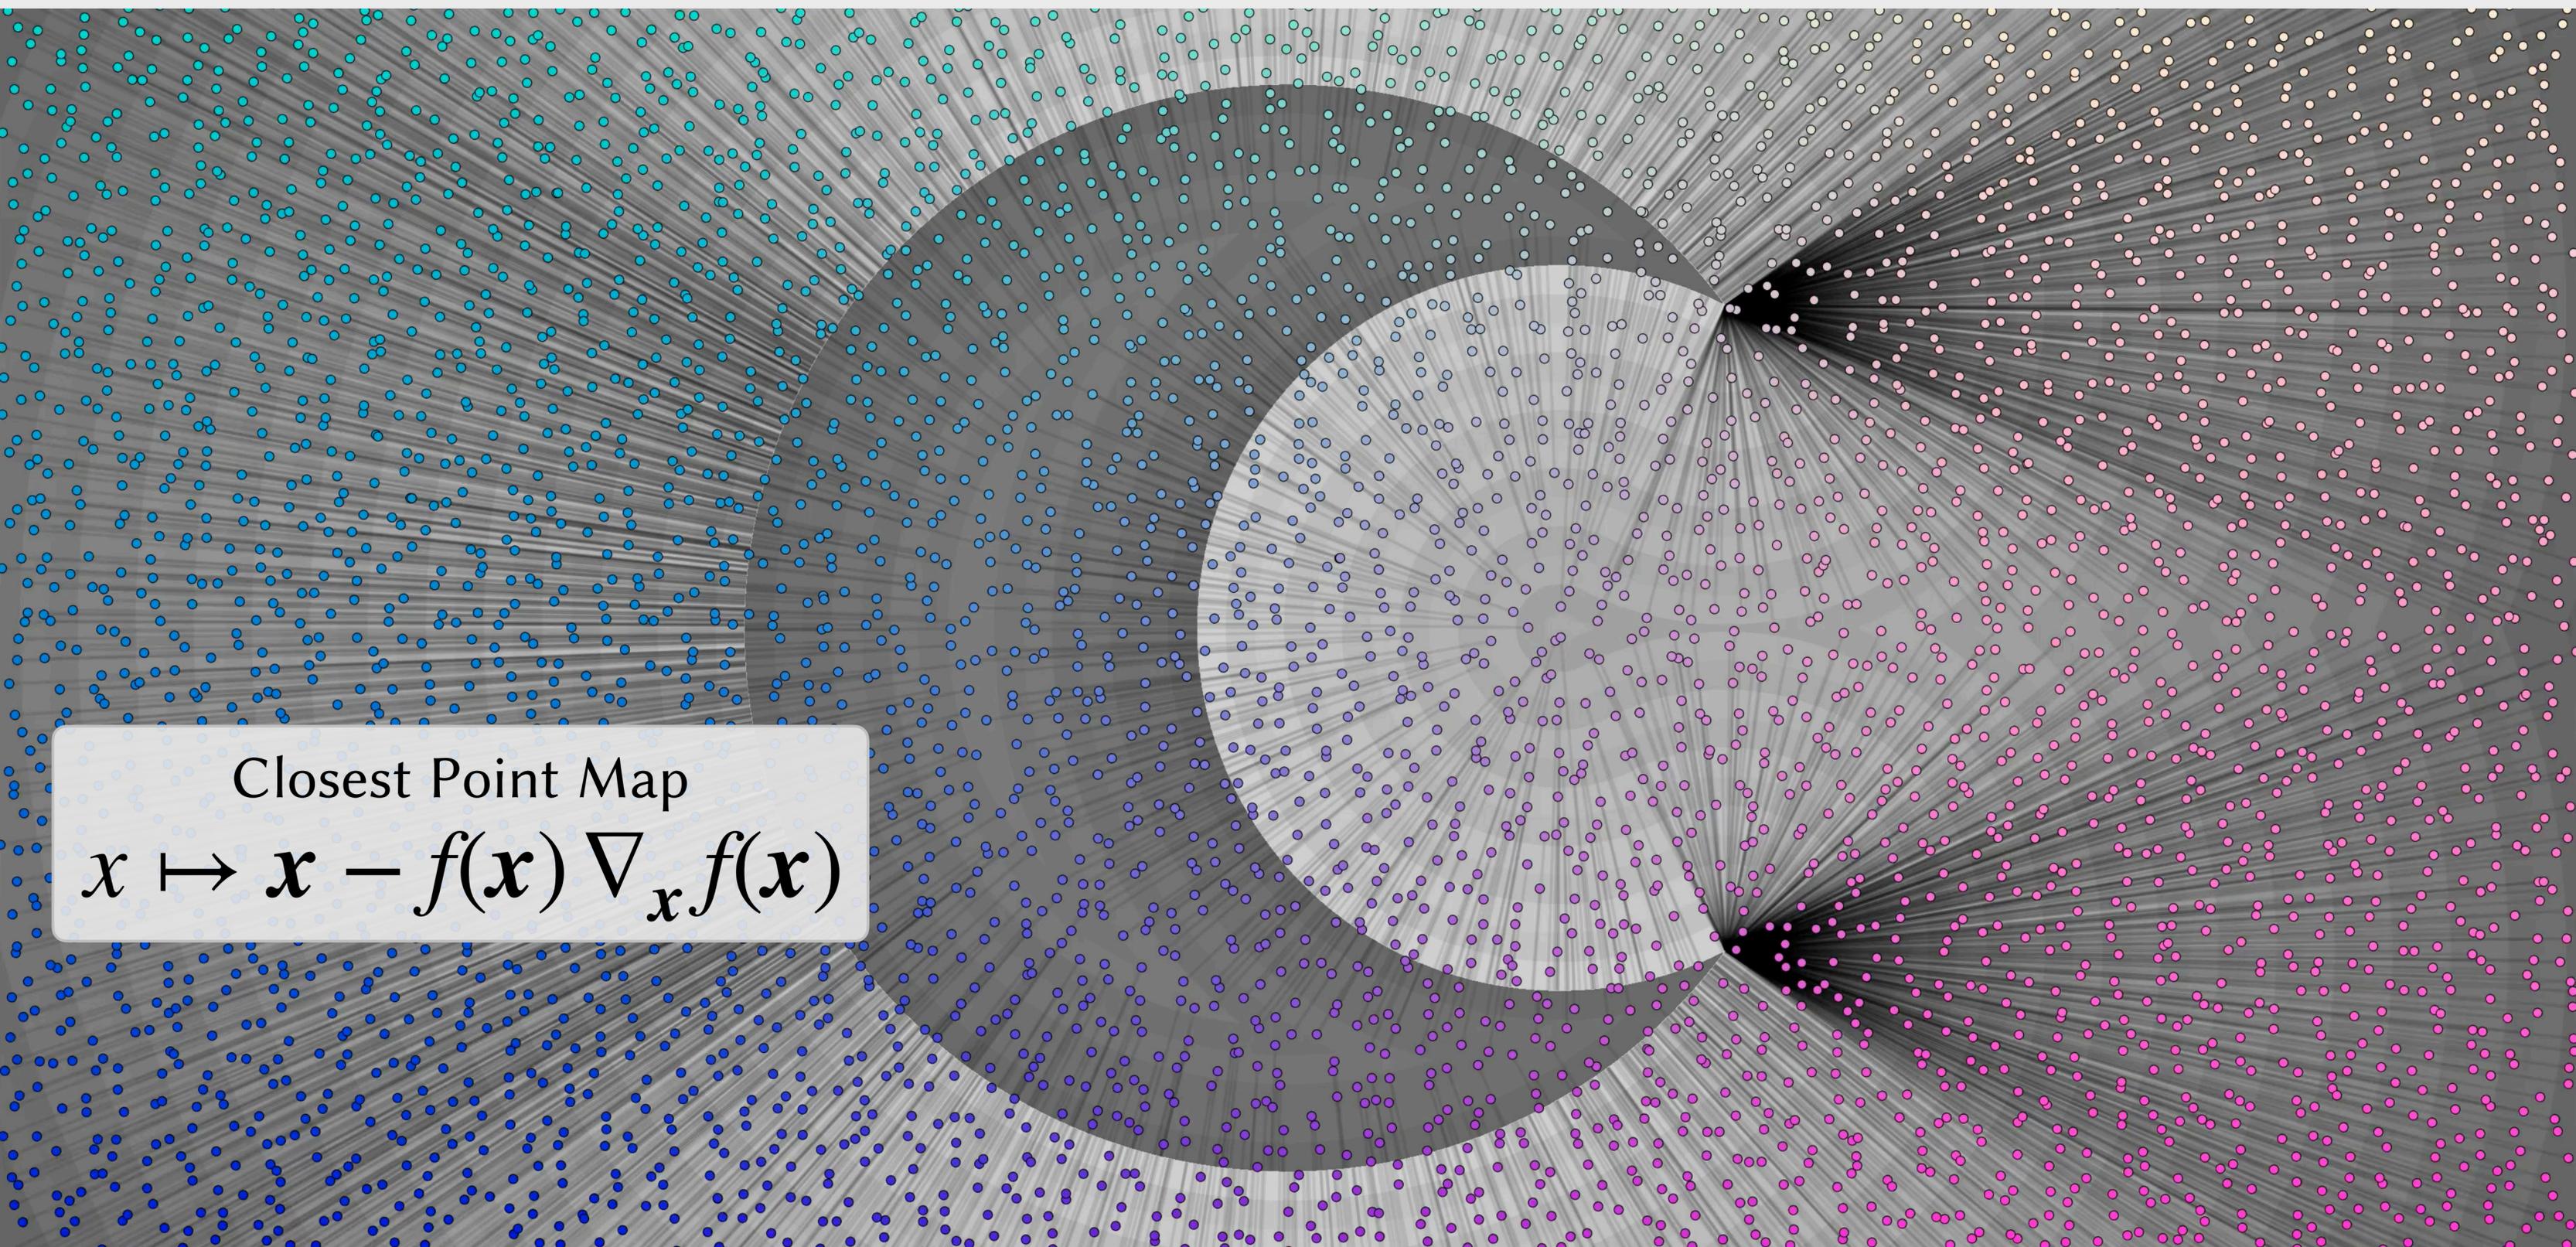
# Closest point loss: walkthrough



# Closest point loss: walkthrough



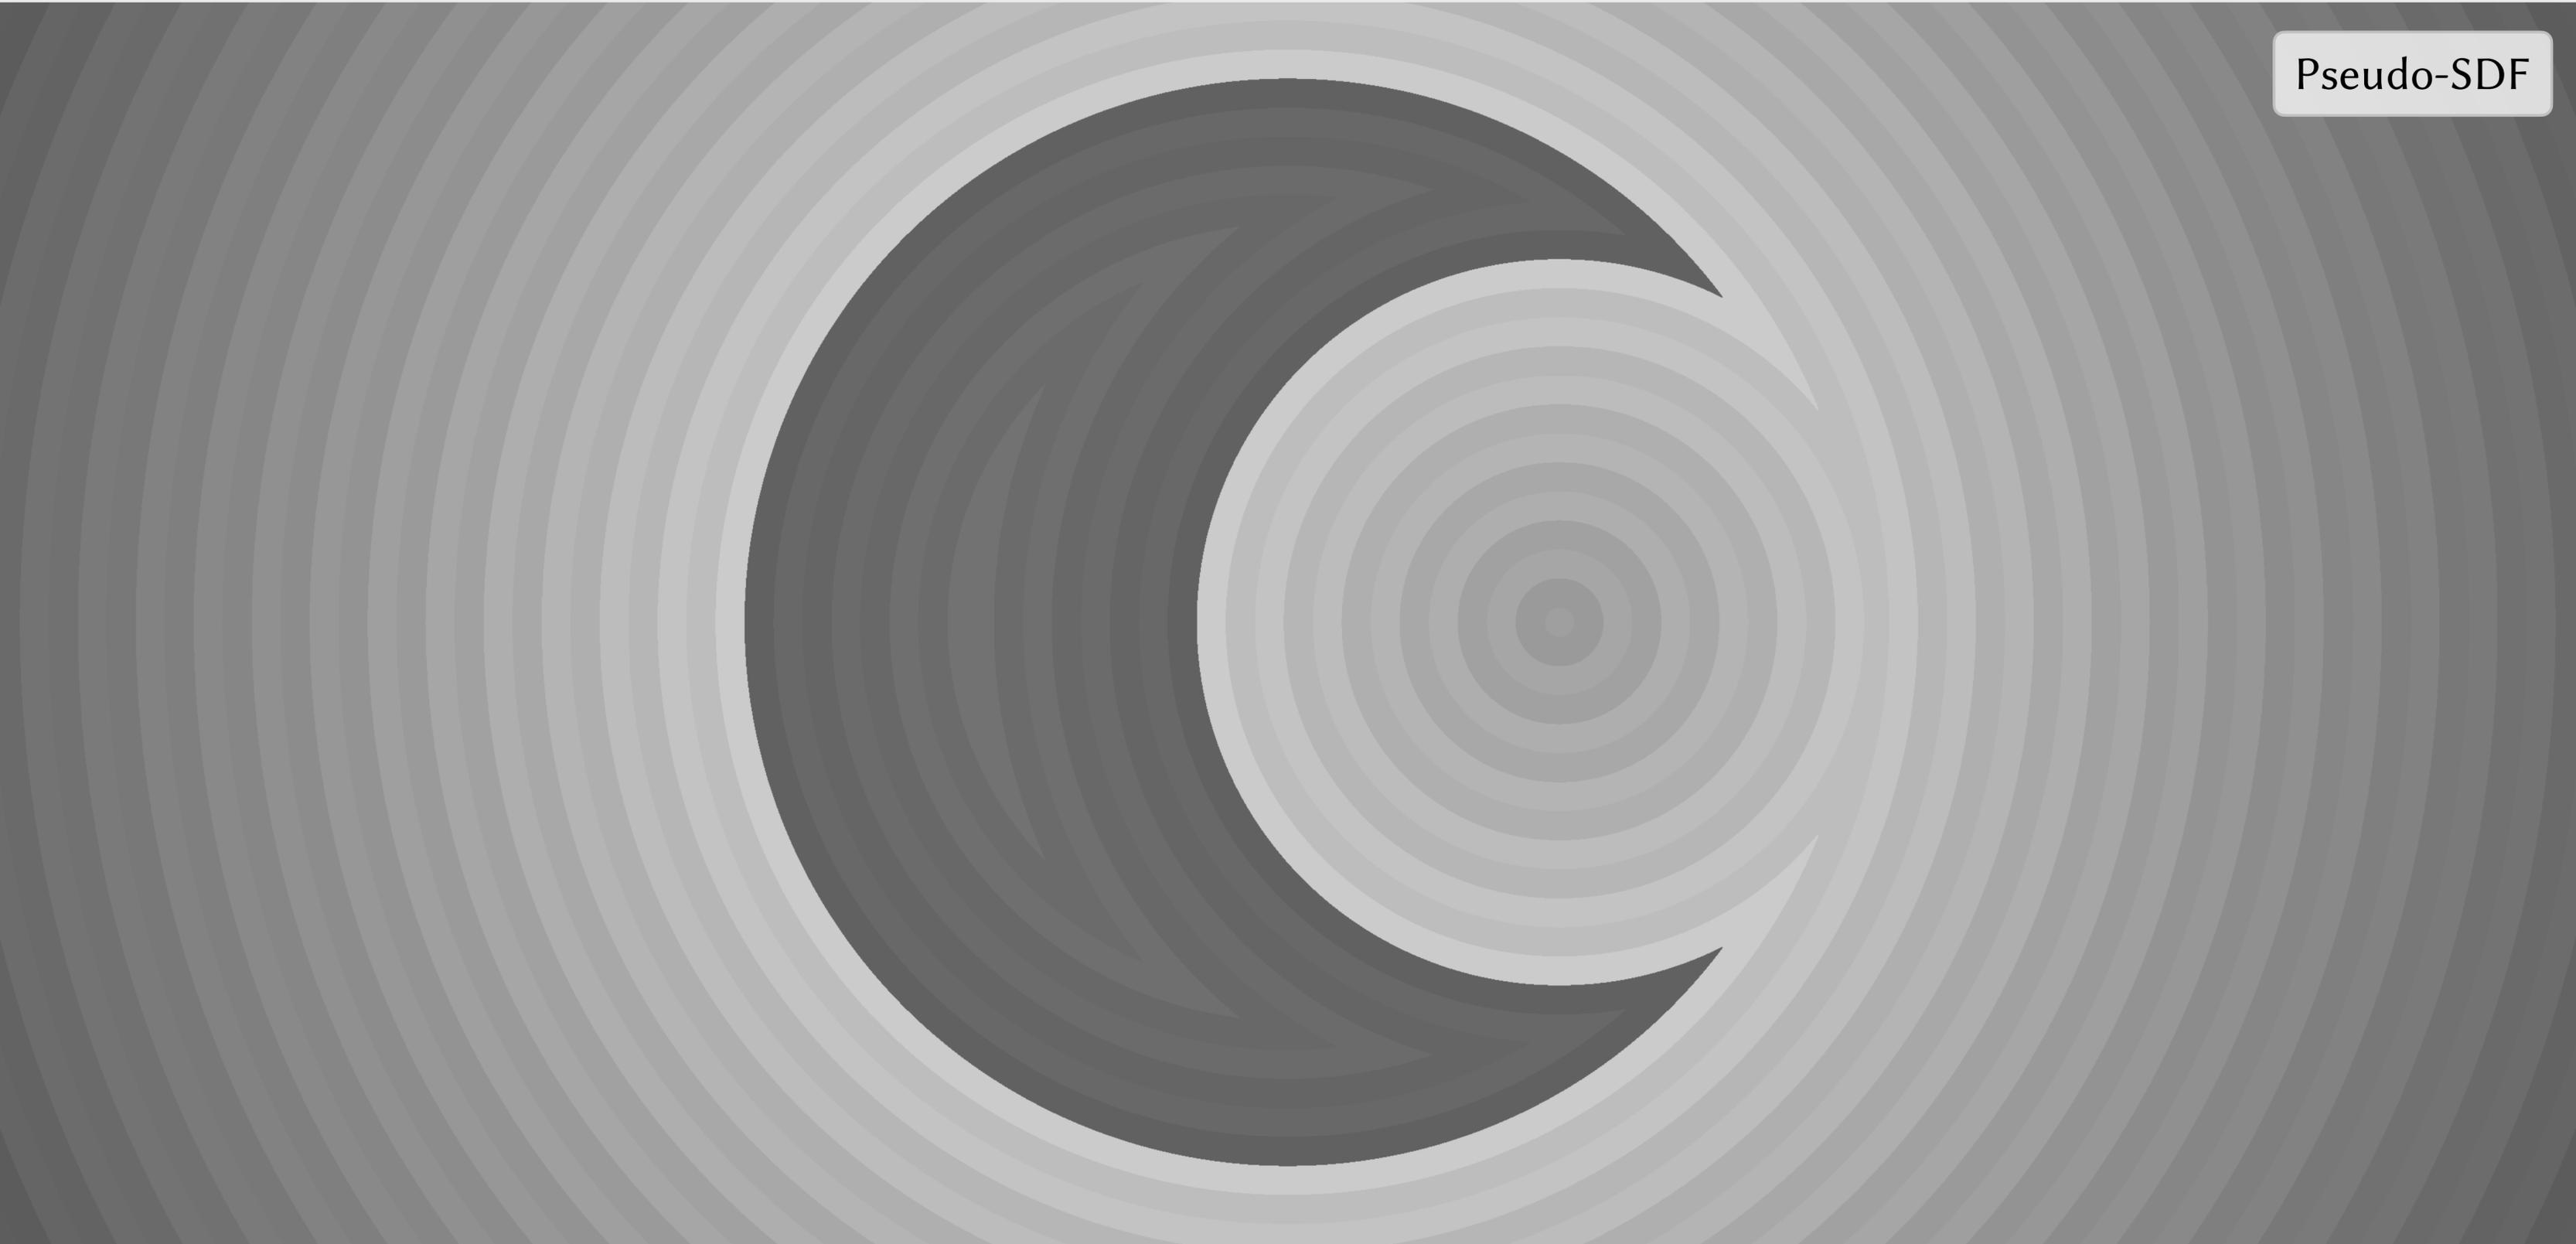
# Closest point loss: walkthrough



Closest Point Map  
 $x \mapsto \mathbf{x} - f(\mathbf{x}) \nabla_{\mathbf{x}} f(\mathbf{x})$

# Closest point loss: walkthrough

Pseudo-SDF



# Closest point loss: walkthrough

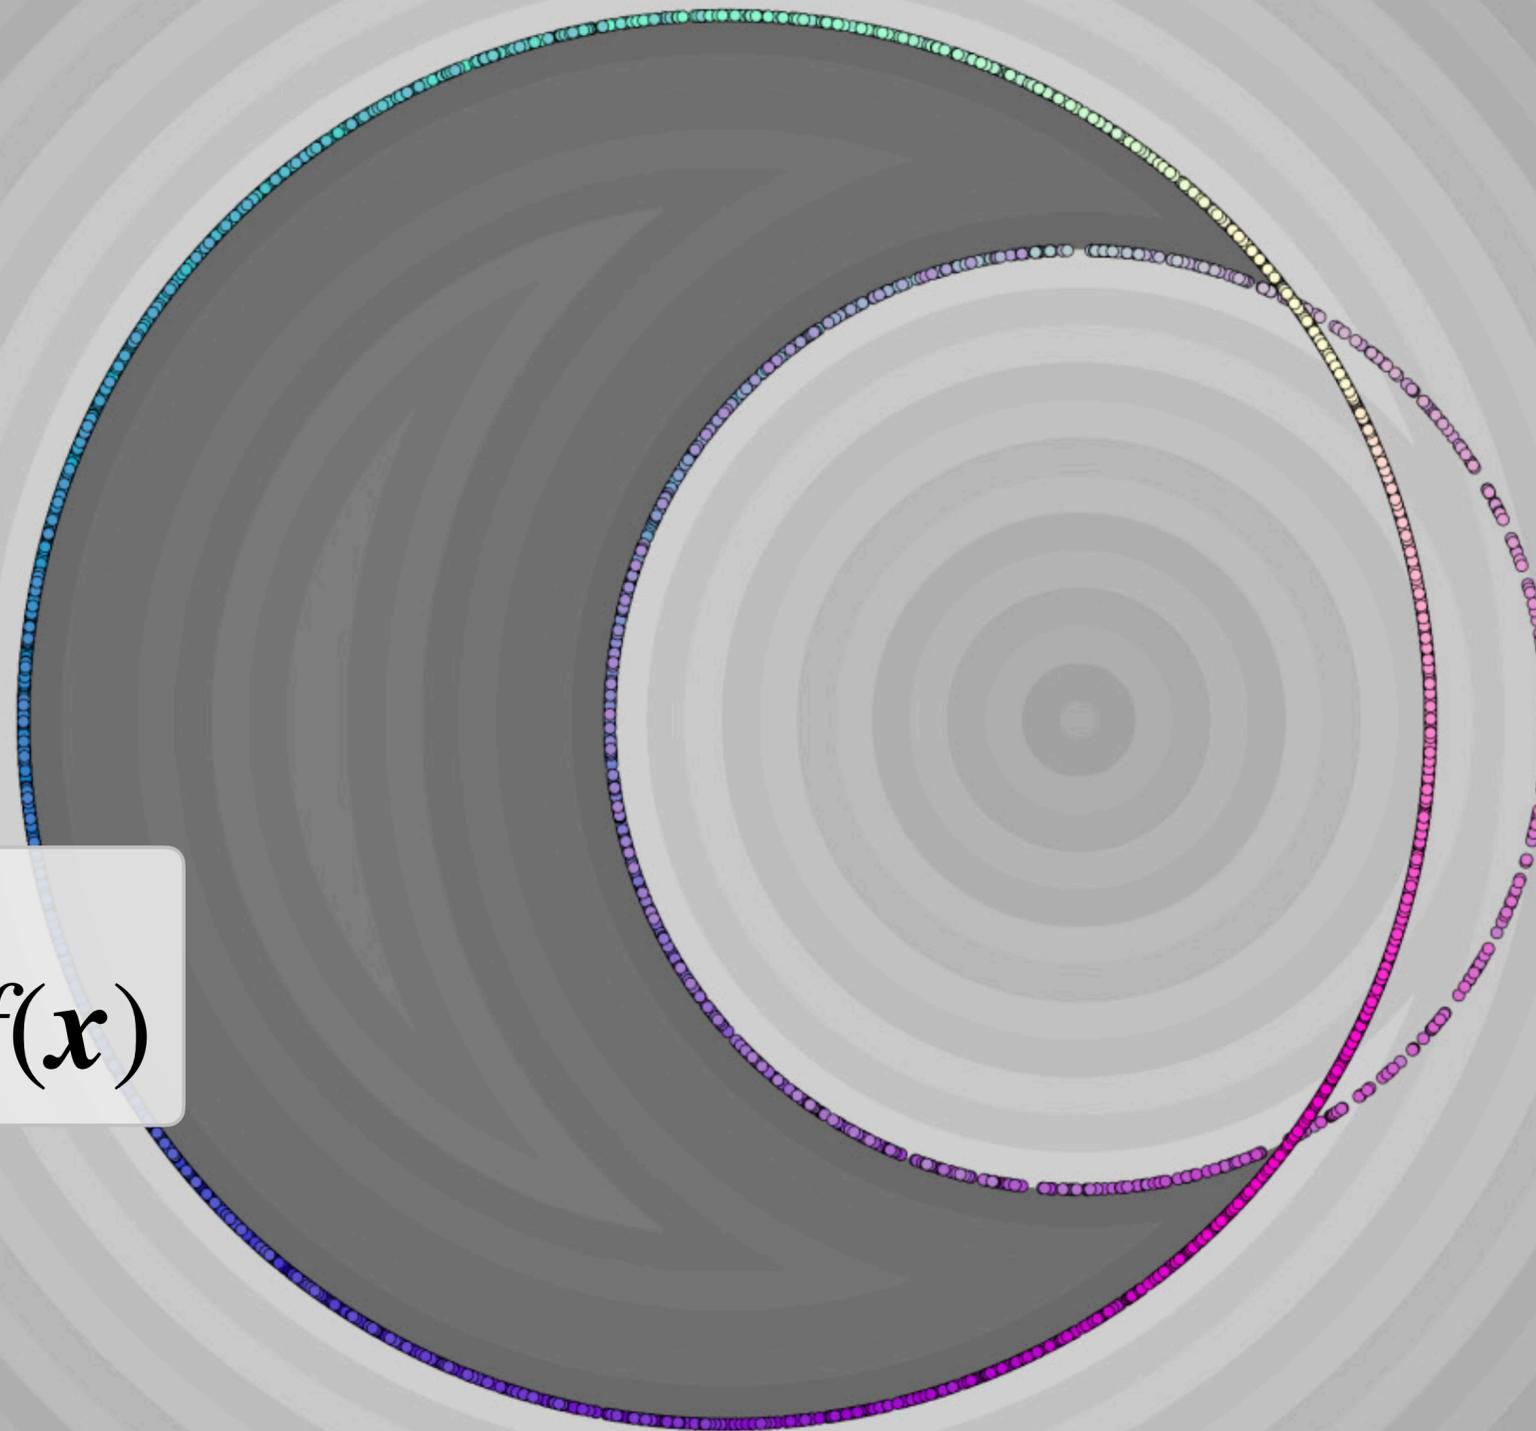
Pseudo-SDF

Closest Point Map

$$\mathbf{x} \mapsto \mathbf{x} - f(\mathbf{x}) \nabla_{\mathbf{x}} f(\mathbf{x})$$

# Closest point loss: walkthrough

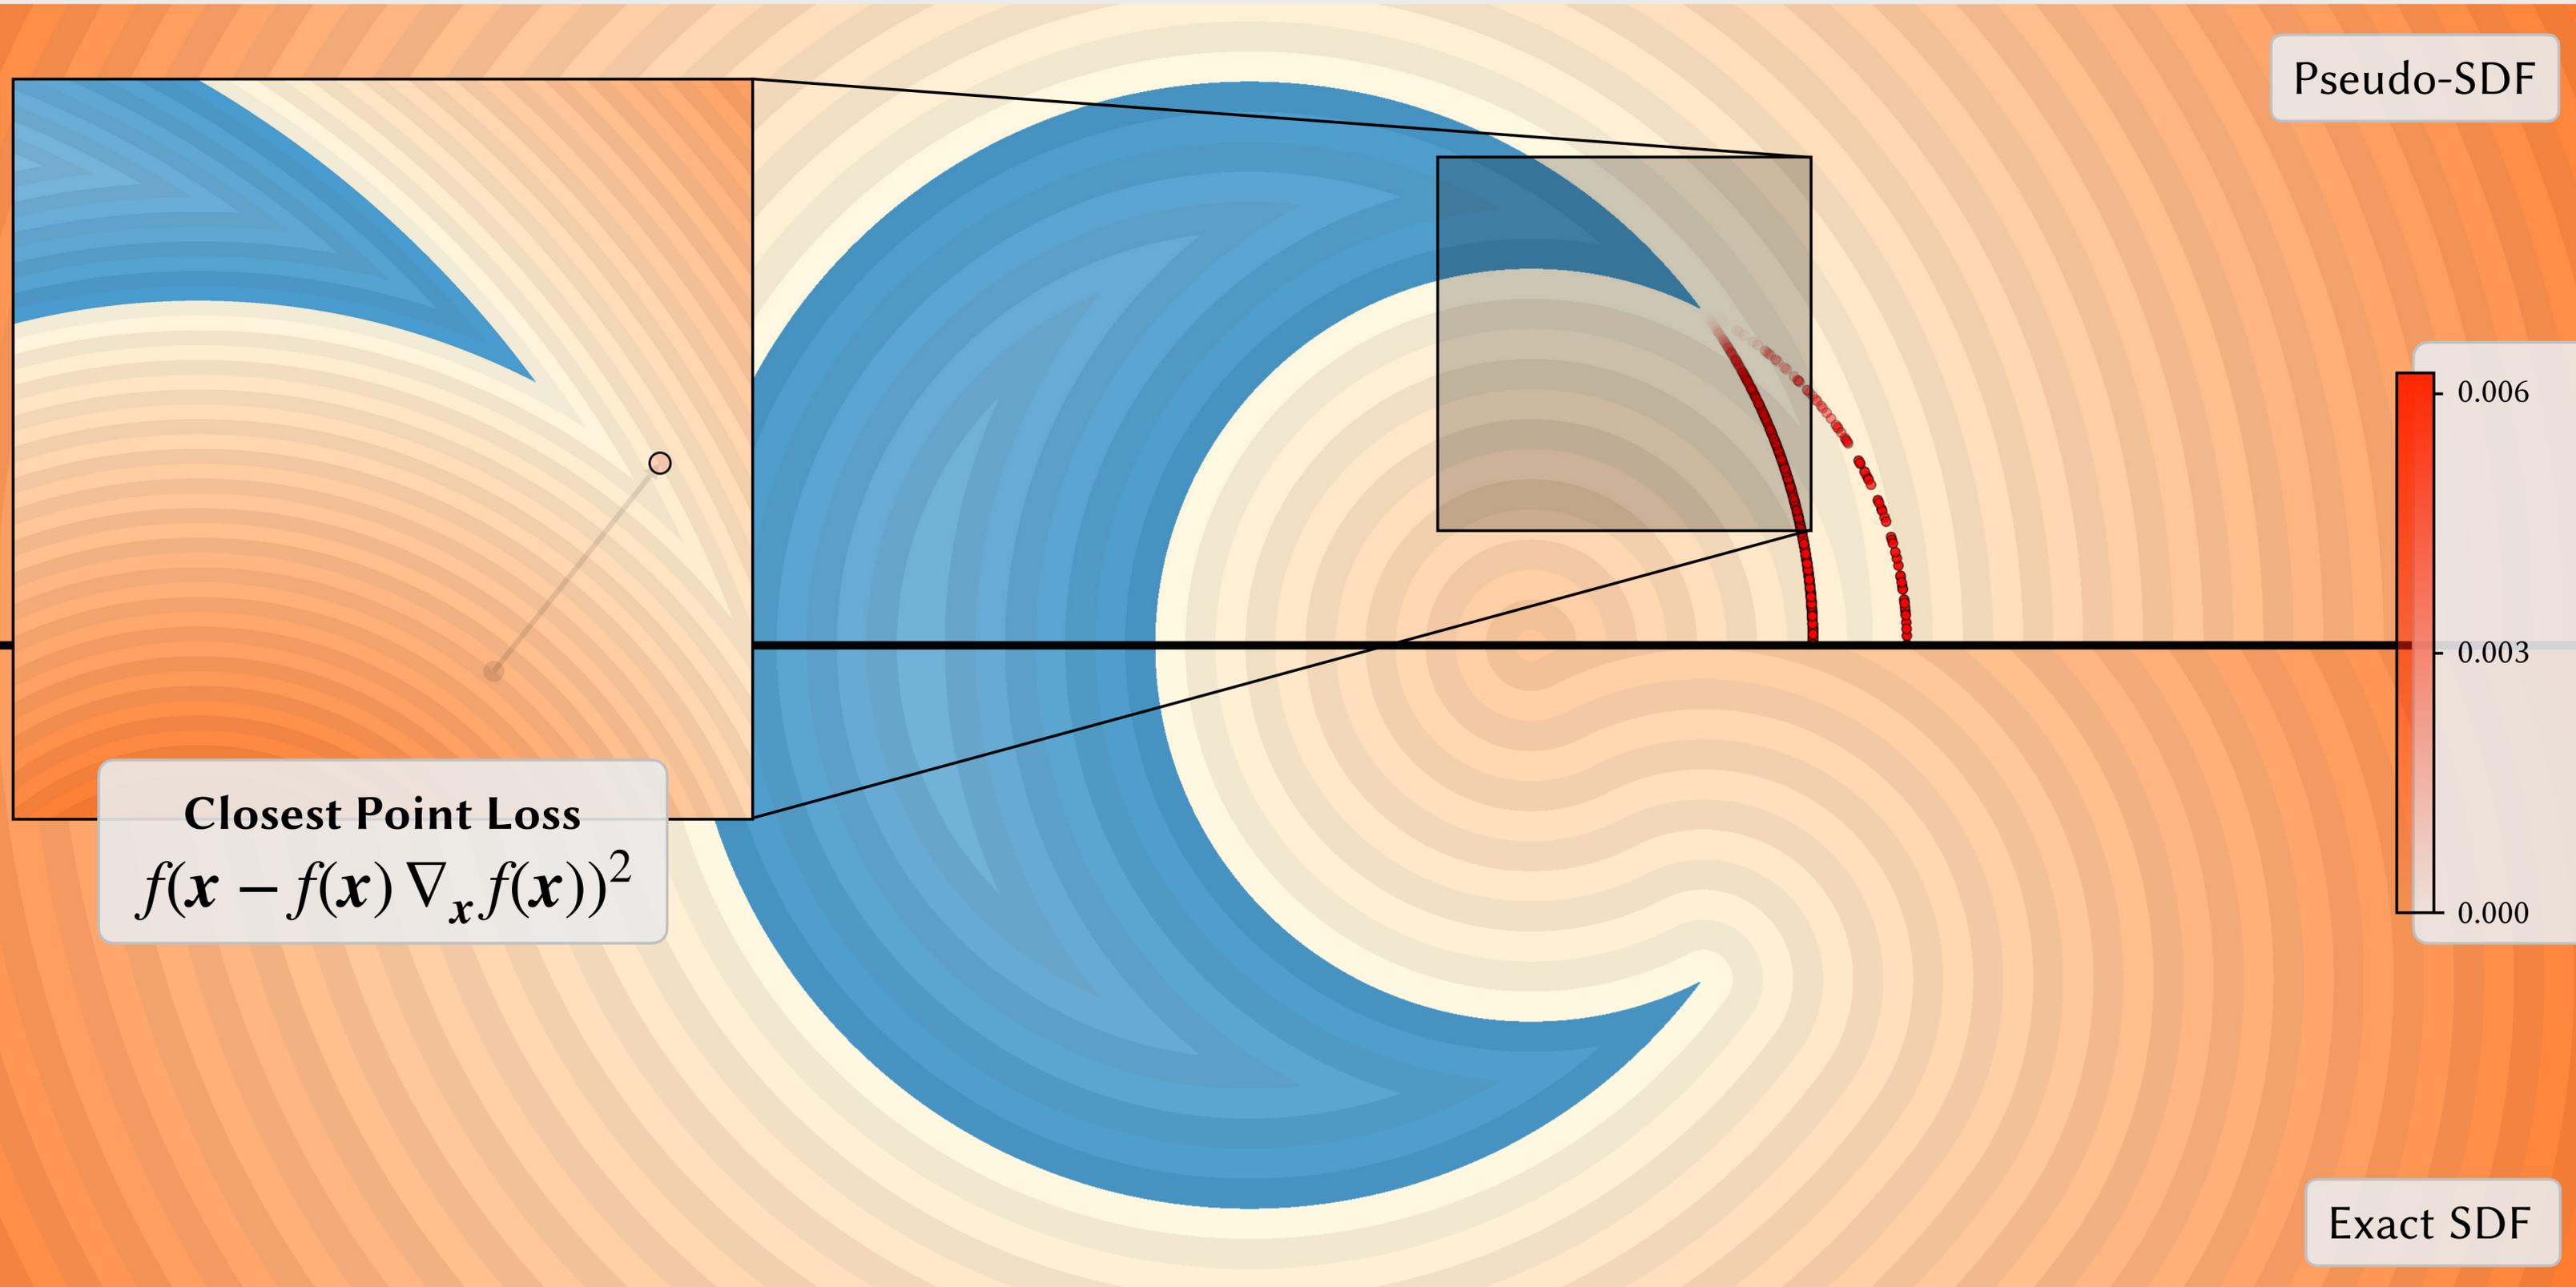
Pseudo-SDF



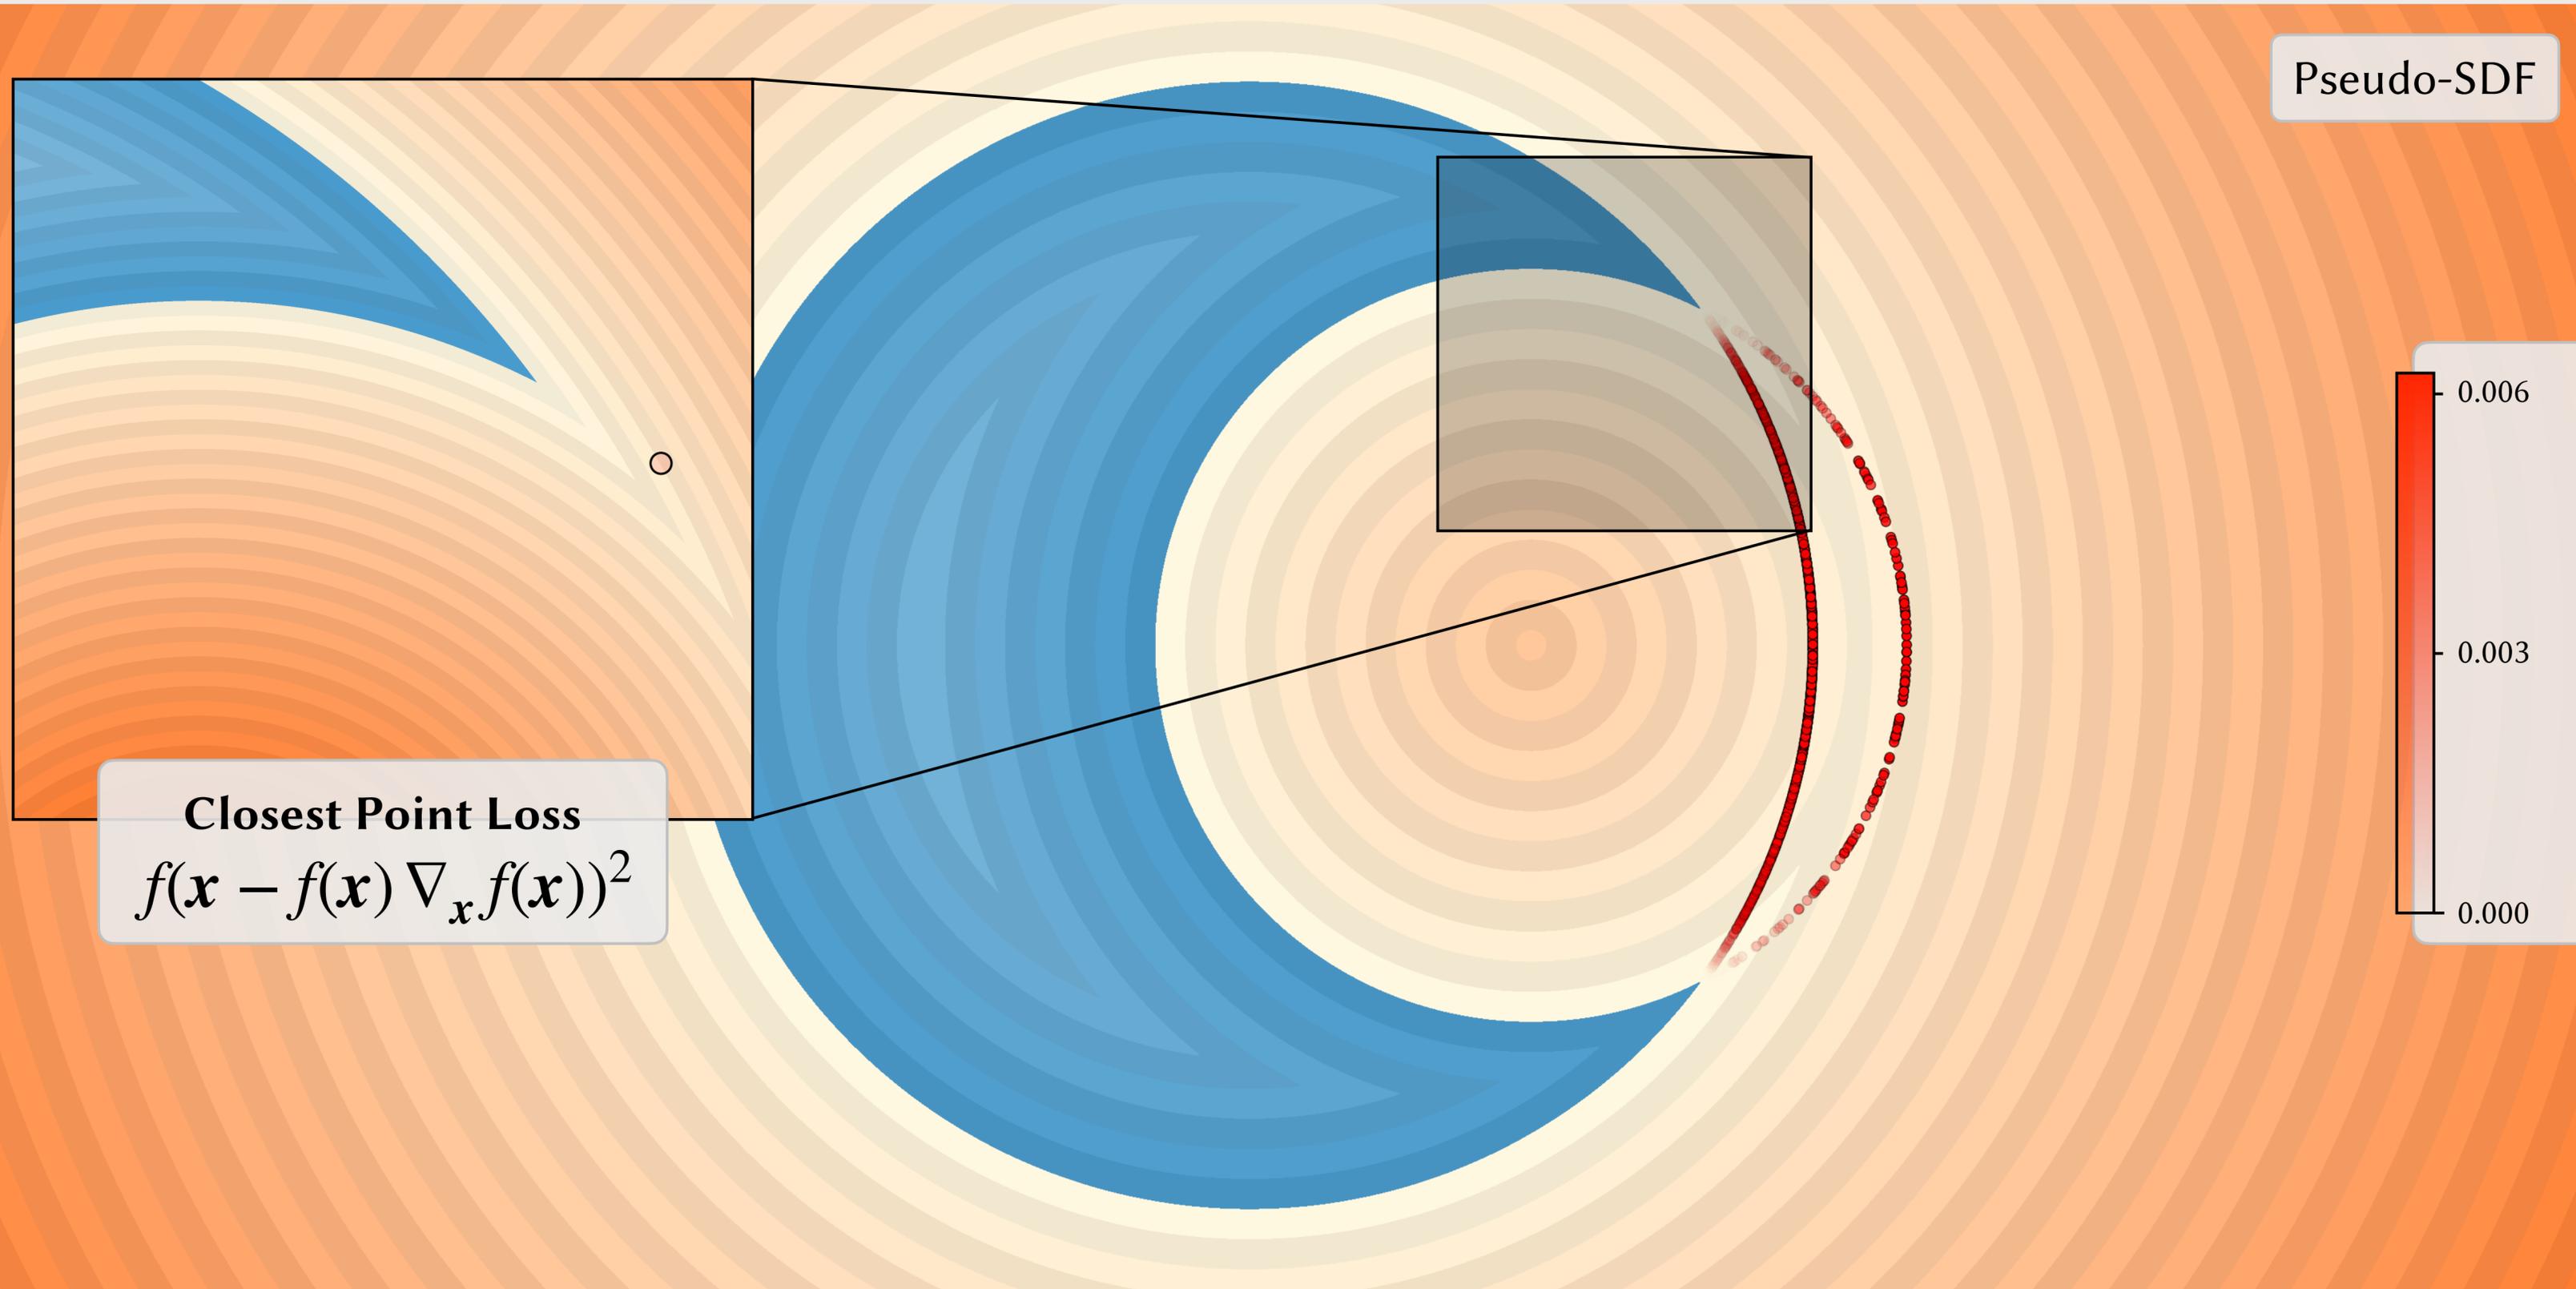
Closest Point Map

$$\mathbf{x} \mapsto \mathbf{x} - f(\mathbf{x}) \nabla_{\mathbf{x}} f(\mathbf{x})$$

# Closest point loss: walkthrough

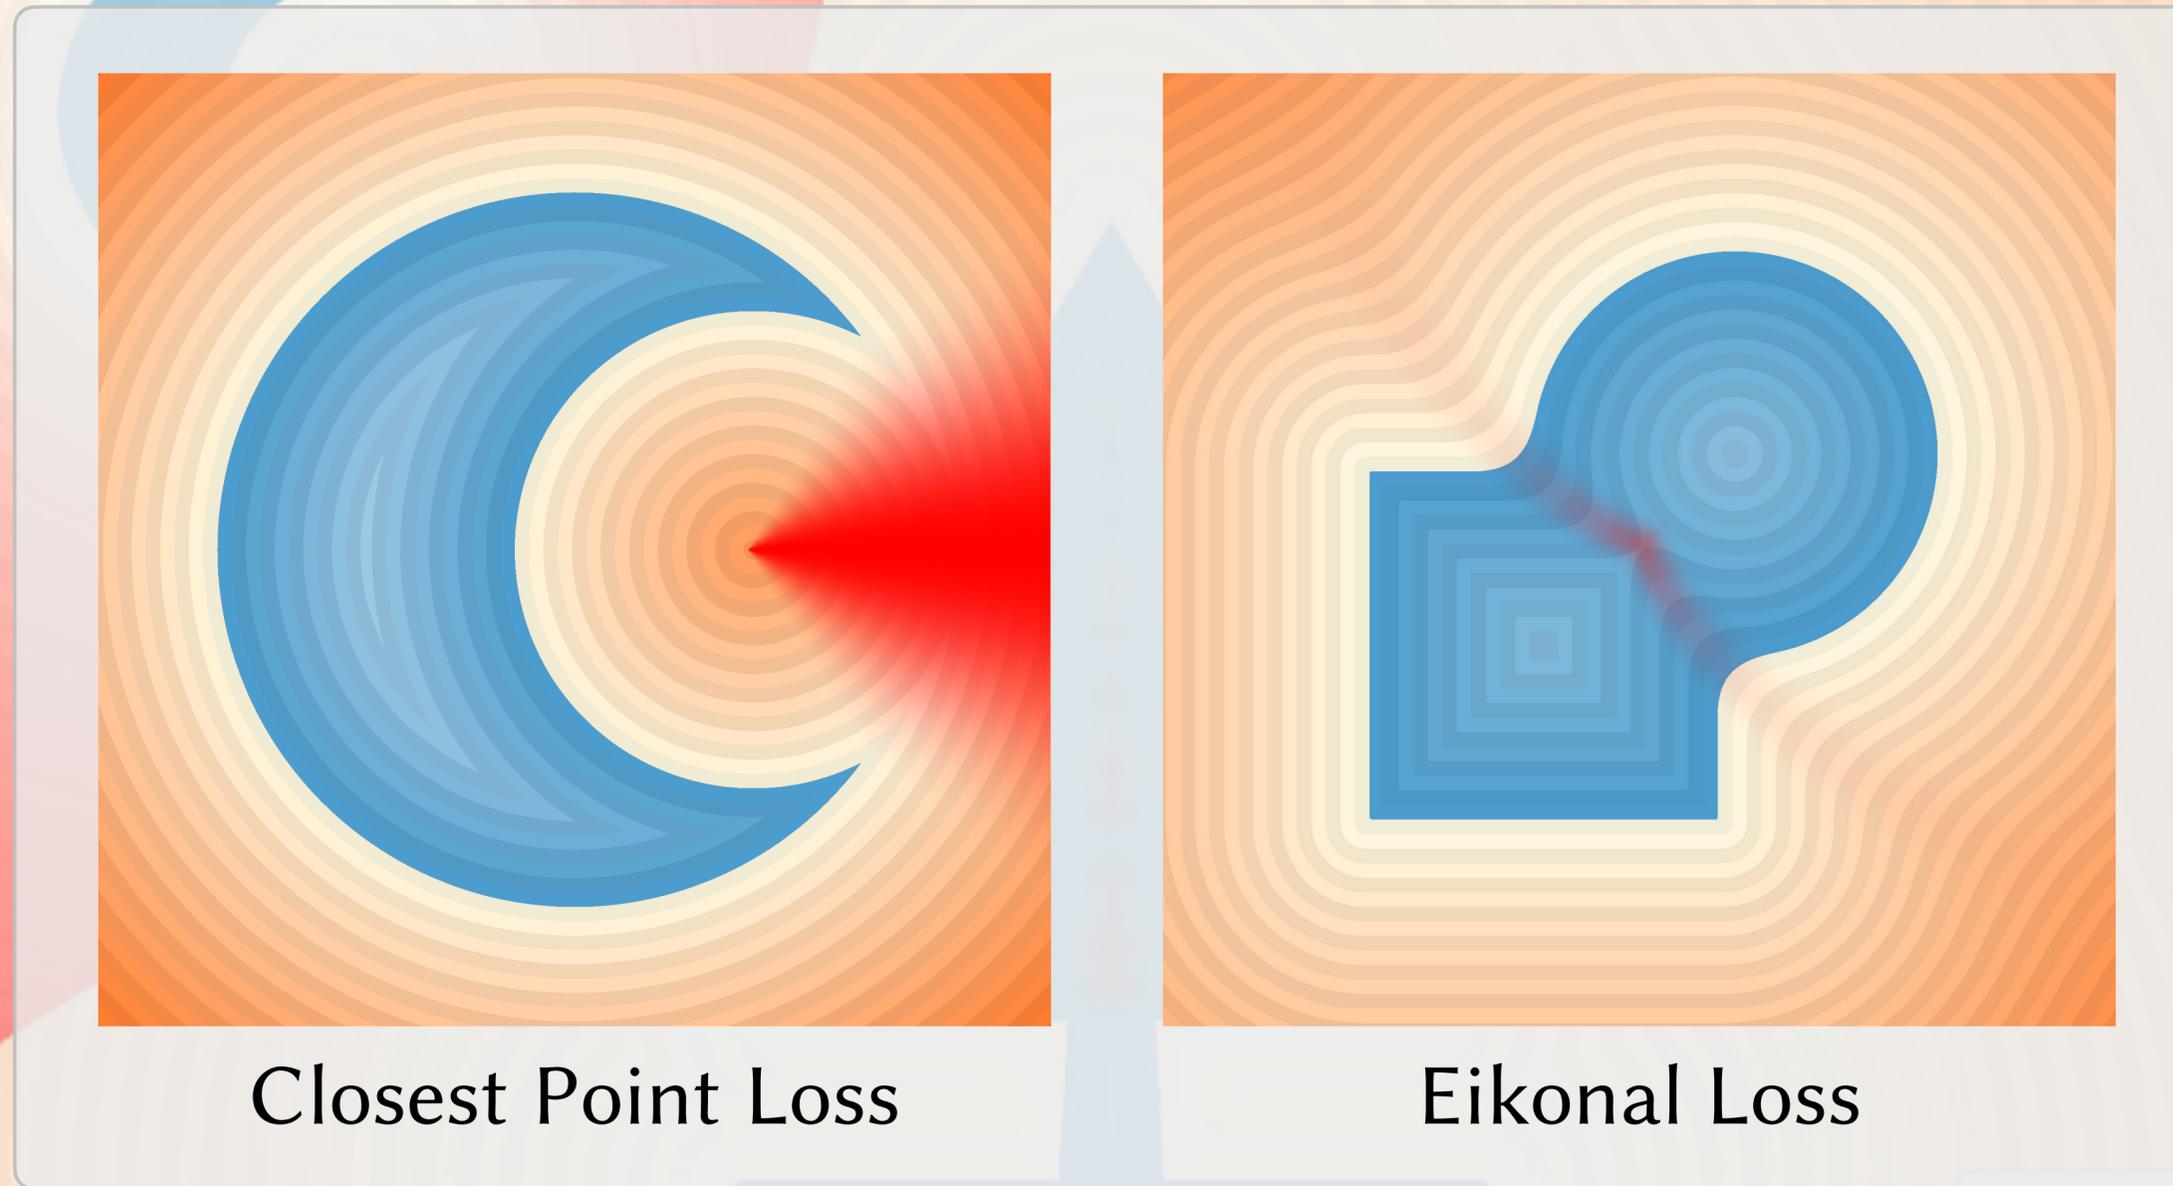


# Closest point loss: walkthrough



**Closest Point Loss**

$$E_{\text{CP}} = f \left( \mathbf{x} - f(\mathbf{x}) \nabla_{\mathbf{x}} f(\mathbf{x}) \right)^2$$



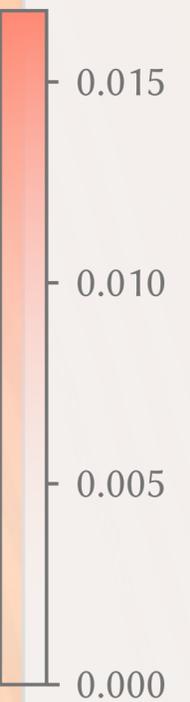
Closest Point Loss

Eikonal Loss

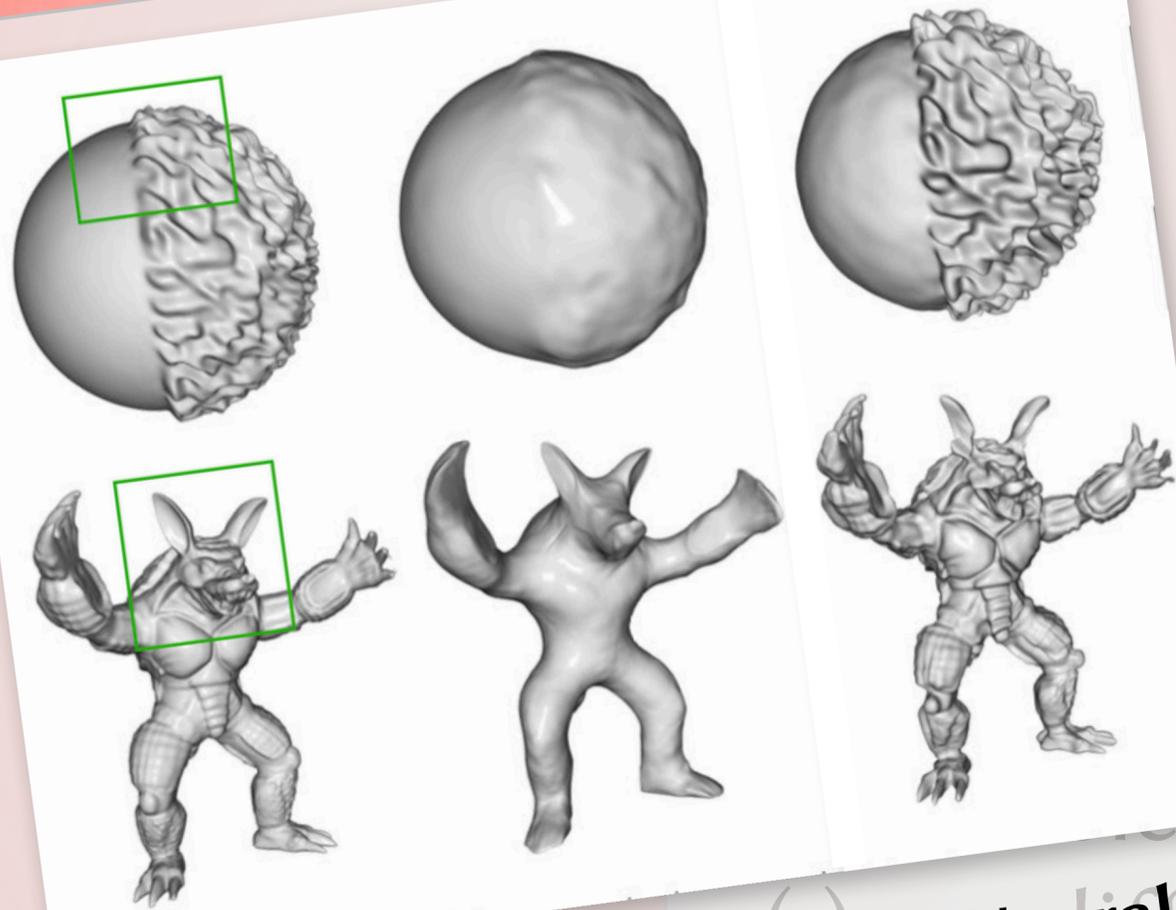
**Closest Point Loss (ours)**  
avg. value = 0.003

## Neural signed distance fields...

- (+) enable many geometric operations
- (-) are difficult to maintain

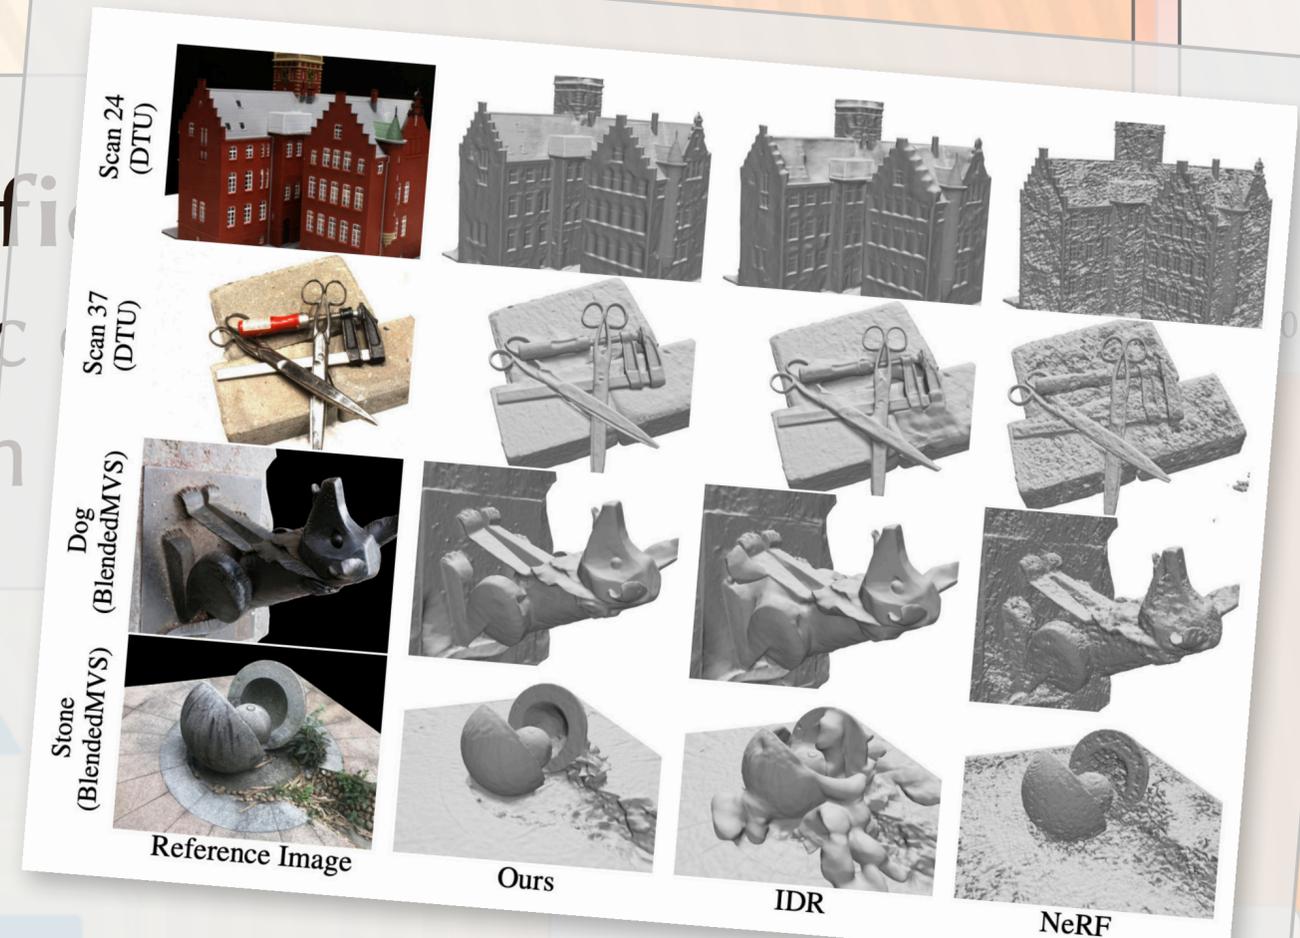


**Closest Point Loss (ours)**  
avg. value = 0.003



Geometry Processing with Neural Fields [Yang et al., 2021]

signed distance fields  
many geometric  
difficult to maintain



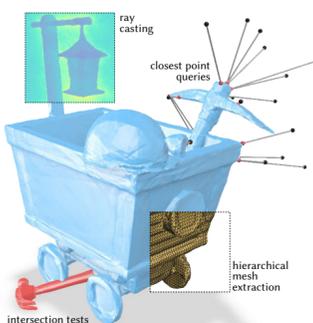
NeuS [Wang et al., 2021]

Closest Point Loss (ours)  
avg. value: 0.003

## Spelunking the Deep: Guaranteed Queries on General Neural Implicit Surfaces via Range Analysis

NICHOLAS SHARP, University of Toronto, Canada  
ALEX JACOBSON, University of Toronto, Adobe Research, Canada

Neural implicit representations, which encode a surface as the level set of a neural network applied to spatial coordinates, have proven to be remarkably effective for optimizing, compressing, and generating 3D geometry. Although these representations are easy to fit, it is not clear how to best evaluate geometric queries on the shape, such as intersecting against a ray or finding a closest point. The predominant approach is to encourage the network to have a signed distance property. However, this property typically holds only approximately, leading to robustness issues, and holds only at the conclusion of training, inhibiting the use of queries in loss functions. Instead, this work presents a new approach to perform queries directly on general neural implicit functions for a wide range of existing architectures. Our key tool is the application of range analysis to neural networks, using automatic arithmetic rules to bound the output of a network over a region; we conduct a study of range analysis on neural networks, and identify variants of affine arithmetic which are highly effective. We use the resulting bounds to develop geometric queries including ray casting, intersection testing, constructing spatial hierarchies, fast mesh extraction, closest-point evaluation, evaluating bulk properties, and more. Our queries can be efficiently evaluated on GPUs, and offer concrete accuracy guarantees even on randomly-initialized networks, enabling their use in training objectives and beyond. We also show a preliminary application to inverse rendering.



CCS Concepts: • Computing methodologies → Shape analysis; Shape representations; • Mathematics of computing → Interval arithmetic.  
Additional Key Words and Phrases: implicit surfaces, neural networks, range analysis, geometry processing

### 1 INTRODUCTION

Representing shapes presents a fundamental dilemma across visual and scientific computing: point clouds and voxel grids are easy to process efficiently, but lack explicit connectivity information; meshes offer a concise and precise description of a surface, but may require difficult unstructured computation, etc. Recently, neural implicit representations have emerged as a promising alternative for a variety of important tasks—the basic idea is to encode a surface as a level set of a neural network applied to spatial coordinates. These neural implicit surfaces inherit many of the strengths which have made neural networks ubiquitous across visual computing, including effective gradient-based optimization, integration with data-driven priors and objectives, and straightforward parallelization on modern hardware.

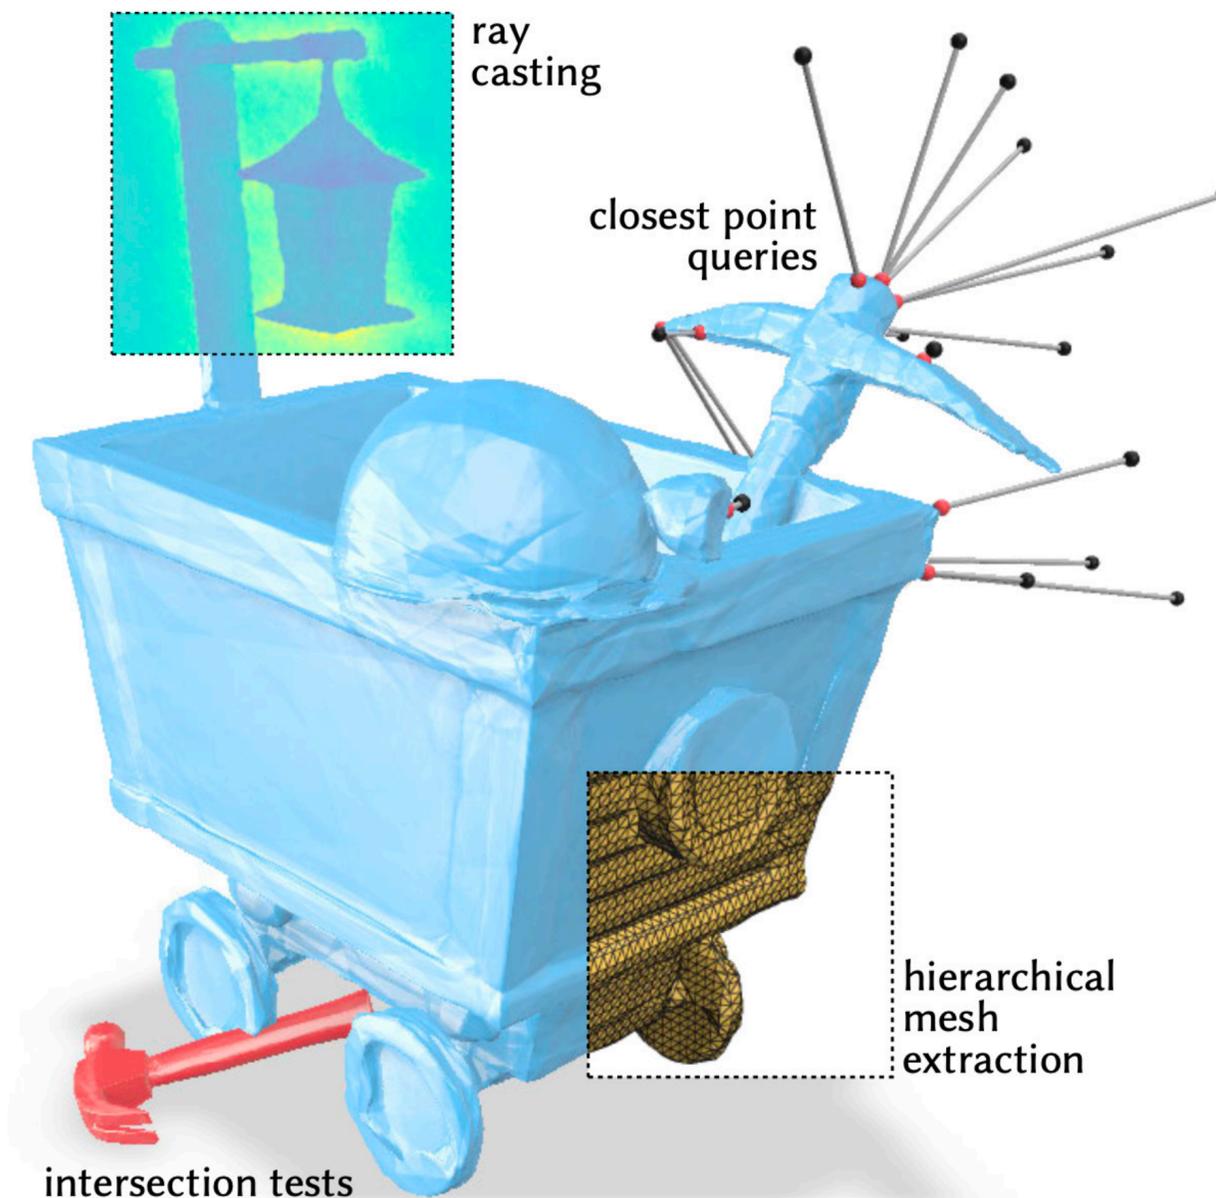
However, there is a price to pay in return for these strong properties: there is no clear strategy for evaluating even the most basic geometric queries against a neural implicit surface, such as intersecting a ray with the surface, or finding a closest point. It would

Authors' addresses: Nicholas Sharp, University of Toronto, Canada, nsharp@cs.toronto.edu; Alex Jacobson, University of Toronto, Adobe Research, Canada, jacobson@cs.toronto.edu.

Fig. 1. Our method enables geometric queries on neural implicit surfaces, without relying on fitting a signed distance function. Several queries are shown here on a neural implicit occupancy function encoding a mine cart. These operations open up new explorations of deep implicit surfaces.

seem that the *only* thing we can do with such a function is to sample it at a point. In a sense, the powerful generality of neural networks is exactly what makes them difficult to query—because they can approximate arbitrary functions with adaptive spatial resolution, it is very difficult to characterize the geometry of their level sets.

One popular recourse is to attempt to fit implicit functions which not only encode a surface via their zero level set, but furthermore have a signed-distance function (SDF) property away from the level set: the magnitude of the function gives the distance to the surface. Although exact SDFs are well-suited for many queries in geometry processing, approximate neural SDFs leave much to be desired. First, such networks are only *approximately* SDFs, and may overestimate the distance to the surface, causing queries to fail unpredictably. More importantly, the SDF property only applies *after* a network has been successfully fitted; thus we cannot make use of geometric queries in the early stages of training, e.g., to define geometric loss functions. Even more broadly, relaxing the expectation that a network fits an SDF opens up a broader class of neural network formulations and objectives, such as those based on occupancy (e.g., as in Section 5).



## Guaranteed Queries on General Neural Implicit Surfaces via Range Analysis [Sharp et al., 2022]

## Spelunking the Deep: Guaranteed Queries on General Neural Implicit Surfaces via Range Analysis

NICHOLAS SHARP, University of Toronto, Canada  
ALEX JACOBSON, University of Toronto, Adobe Research, Canada

Neural implicit representations, which encode a surface as the level set of a neural network applied to spatial coordinates, have proven to be remarkably effective for optimizing, compressing, and generating 3D geometry. Although these representations are easy to fit, it is not clear how to best evaluate geometric queries on the shape, such as intersecting against a ray or finding a closest point. The predominant approach is to encourage the network to have a signed distance property. However, this property typically holds only approximately, leading to robustness issues, and holds only at the conclusion of training, inhibiting the use of queries in loss functions. Instead, this work presents a new approach to perform queries directly on general neural implicit functions for a wide range of existing architectures. Our key tool is the application of range analysis to neural networks, using automatic arithmetic rules to bound the output of a network over a region; we conduct a study of range analysis on neural networks, and identify variants of affine arithmetic which are highly effective. We use the resulting bounds to develop geometric queries including ray casting, intersection testing, constructing spatial hierarchies, fast mesh extraction, closest-point evaluation, evaluating bulk properties, and more. Our queries can be efficiently evaluated on GPUs, and offer concrete accuracy guarantees even on randomly-initialized networks, enabling their use in training objectives and beyond. We also show a preliminary application to inverse rendering.

CCS Concepts: • Computing methodologies → Shape analysis; Shape representations; Mathematics of computing → Interval arithmetic.  
Additional Key Words and Phrases: implicit surfaces, neural networks, range analysis, geometry processing

### 1 INTRODUCTION

Representing shapes presents a fundamental dilemma across visual and scientific computing: point clouds and voxel grids are easy to process efficiently, but lack explicit connectivity information; meshes offer a concise and precise description of a surface, but may require difficult unstructured computation, etc. Recently, neural implicit representations have emerged as a promising alternative for a variety of important tasks—the basic idea is to encode a surface as a level set of a neural network applied to spatial coordinates. These neural implicit surfaces inherit many of the strengths which have made neural networks ubiquitous across visual computing, including effective gradient-based optimization, integration with data-driven priors and objectives, and straightforward parallelization on modern hardware.

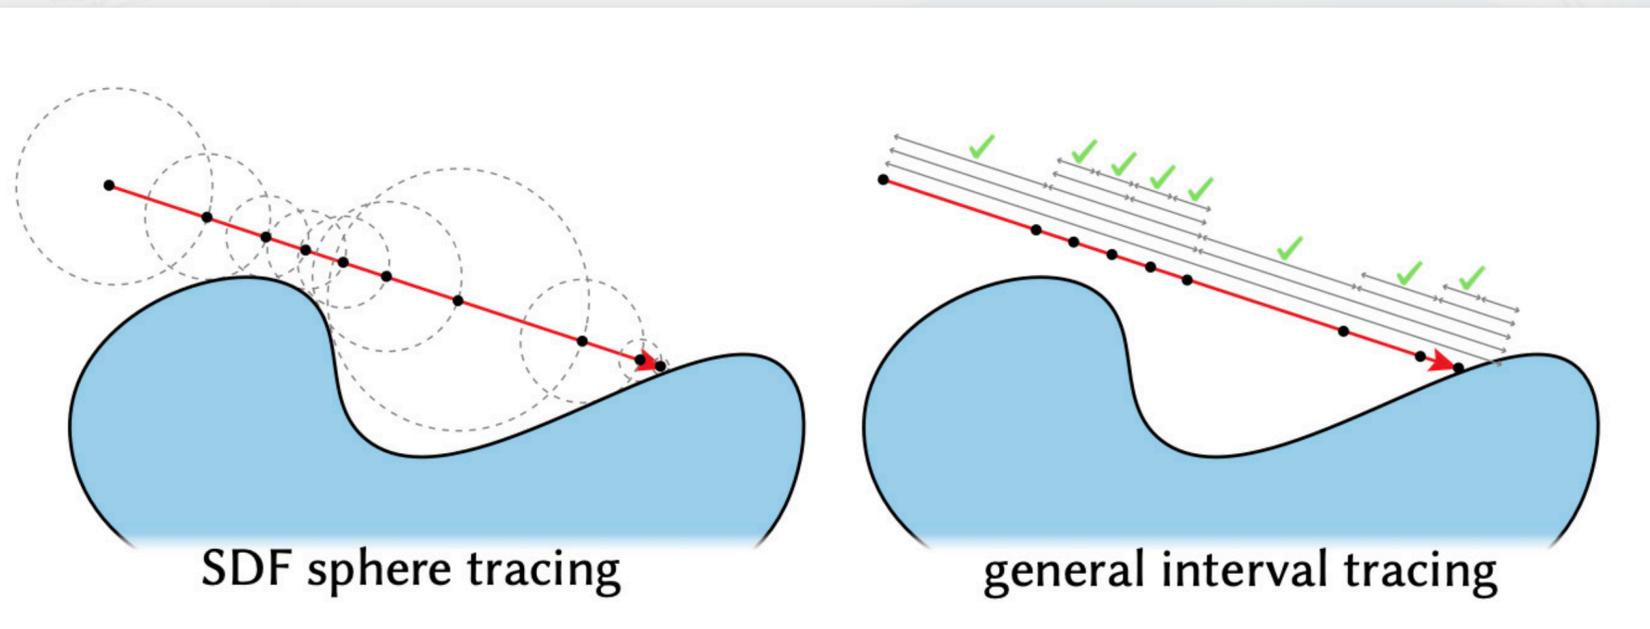
However, there is a price to pay in return for these strong properties: there is no clear strategy for evaluating even the most basic geometric queries against a neural implicit surface, such as intersecting a ray with the surface, or finding a closest point. It would

Authors' addresses: Nicholas Sharp, University of Toronto, Canada, nsharp@cs.toronto.edu; Alex Jacobson, University of Toronto, Adobe Research, Canada, jacobson@cs.toronto.edu

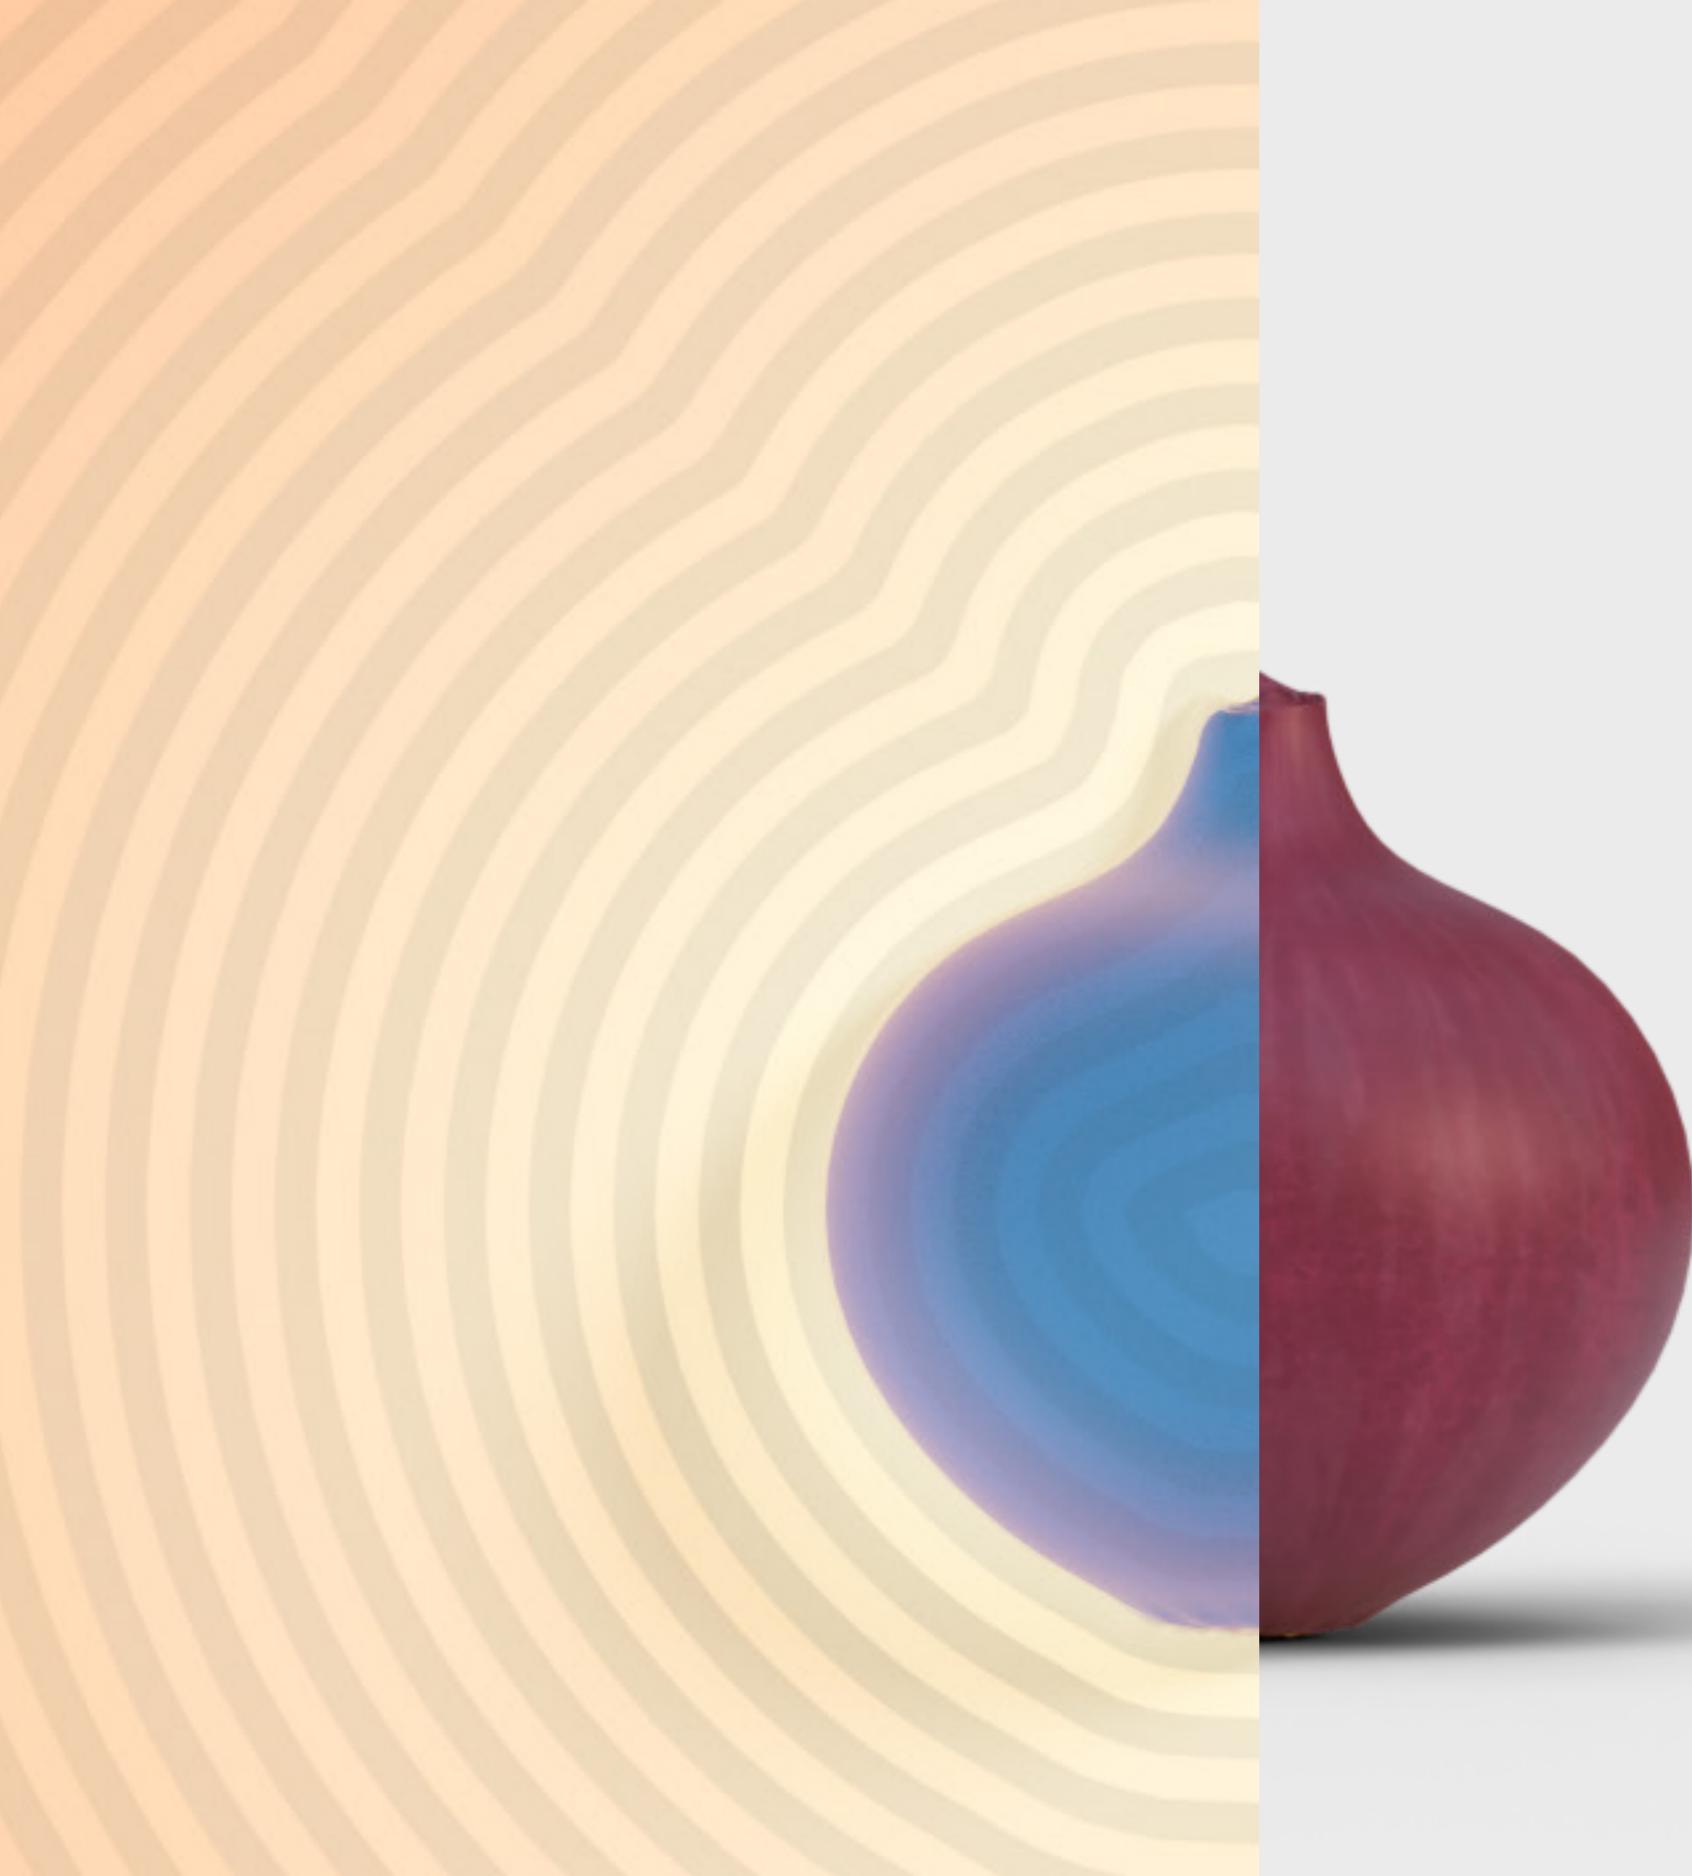
Fig. 1. Our method enables queries without relying on a signed distance function. Shown here on a neural implicit surface. These operations are performed on a GPU.

seem that the only way to find a point on a surface is to sample it at a point. In a sense, this is exactly what makes implicit surfaces so appealing: they are approximate and precise at the same time, but it is very difficult to characterize them. One popular approach is to encourage the network to have a signed distance function property. However, this property typically holds only approximately, leading to robustness issues, and holds only at the conclusion of training, inhibiting the use of queries in loss functions. Instead, this work presents a new approach to perform queries directly on general neural implicit functions for a wide range of existing architectures. Our key tool is the application of range analysis to neural networks, using automatic arithmetic rules to bound the output of a network over a region; we conduct a study of range analysis on neural networks, and identify variants of affine arithmetic which are highly effective. We use the resulting bounds to develop geometric queries including ray casting, intersection testing, constructing spatial hierarchies, fast mesh extraction, closest-point evaluation, evaluating bulk properties, and more. Our queries can be efficiently evaluated on GPUs, and offer concrete accuracy guarantees even on randomly-initialized networks, enabling their use in training objectives and beyond. We also show a preliminary application to inverse rendering.

## Technique: interval analysis on neural network



## Guaranteed Queries on General Neural Implicit Surfaces via Range Analysis [Sharp et al., 2022]



## OVERVIEW

### IMPLICIT SURFACES

1. Background & History
2. Digital Representations

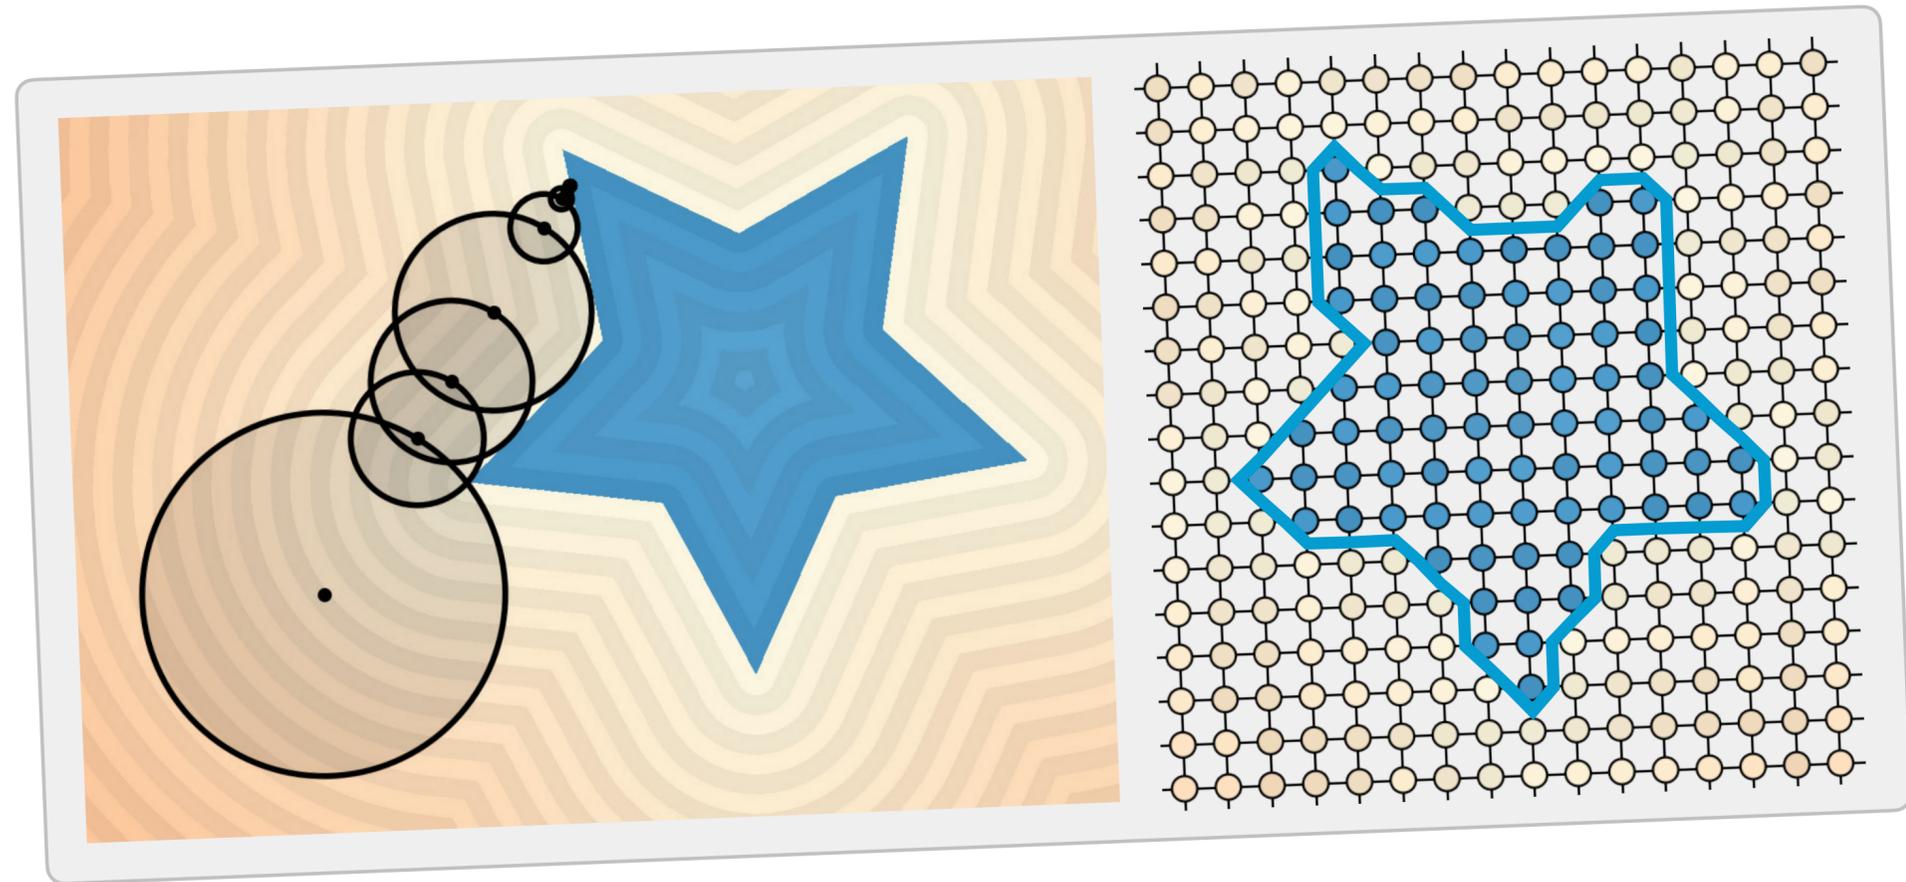
### NEURAL IMPLICIT

3. The Basics
4. Why & Why Not Use Them?
5. Working out the Details
6. Geometric Operations

## CONCLUSION

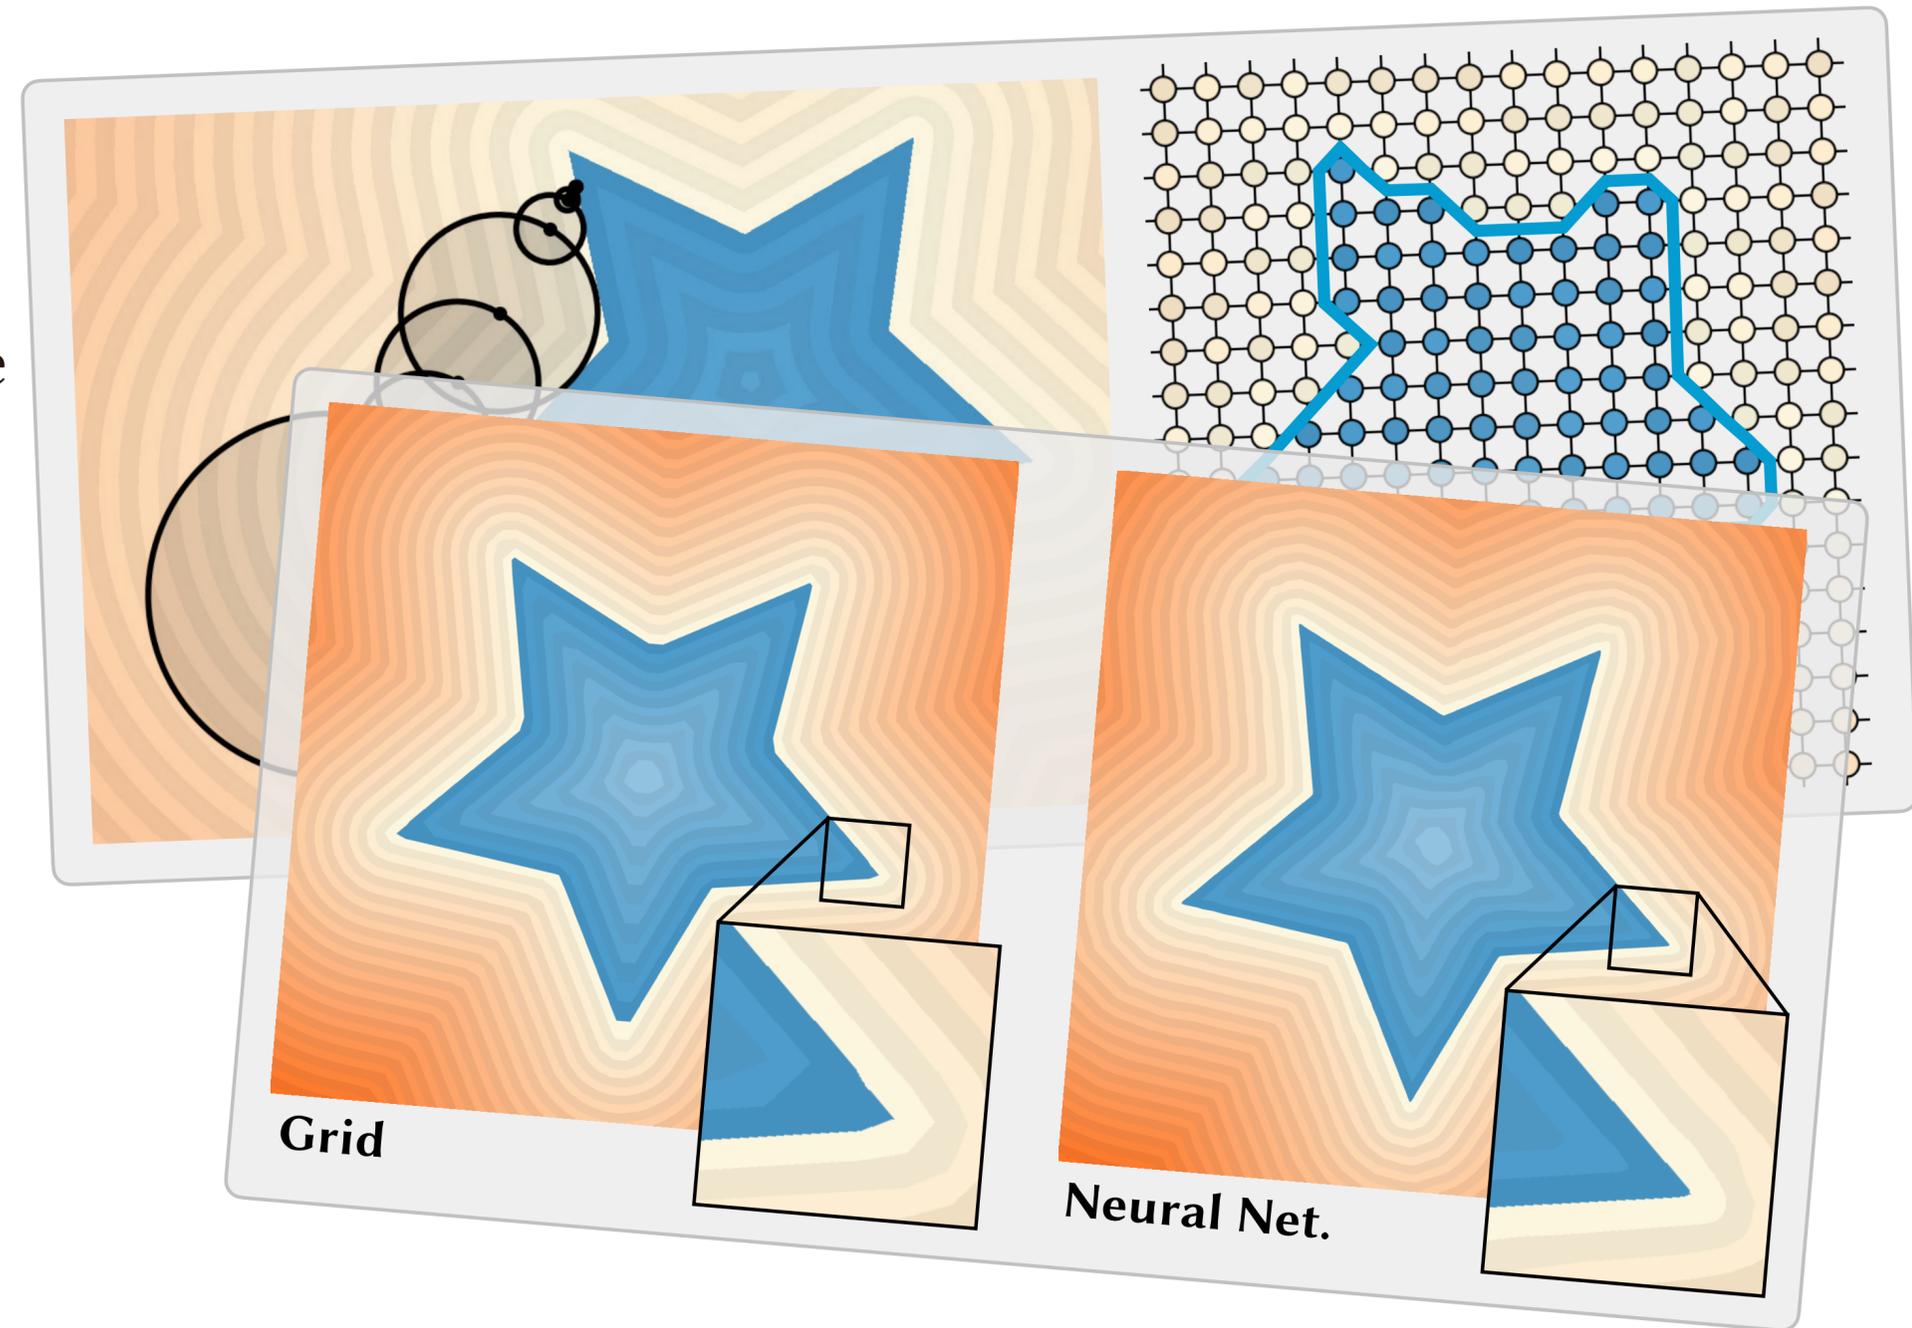
## We saw...

- how to work with a classic representation (implicit functions)



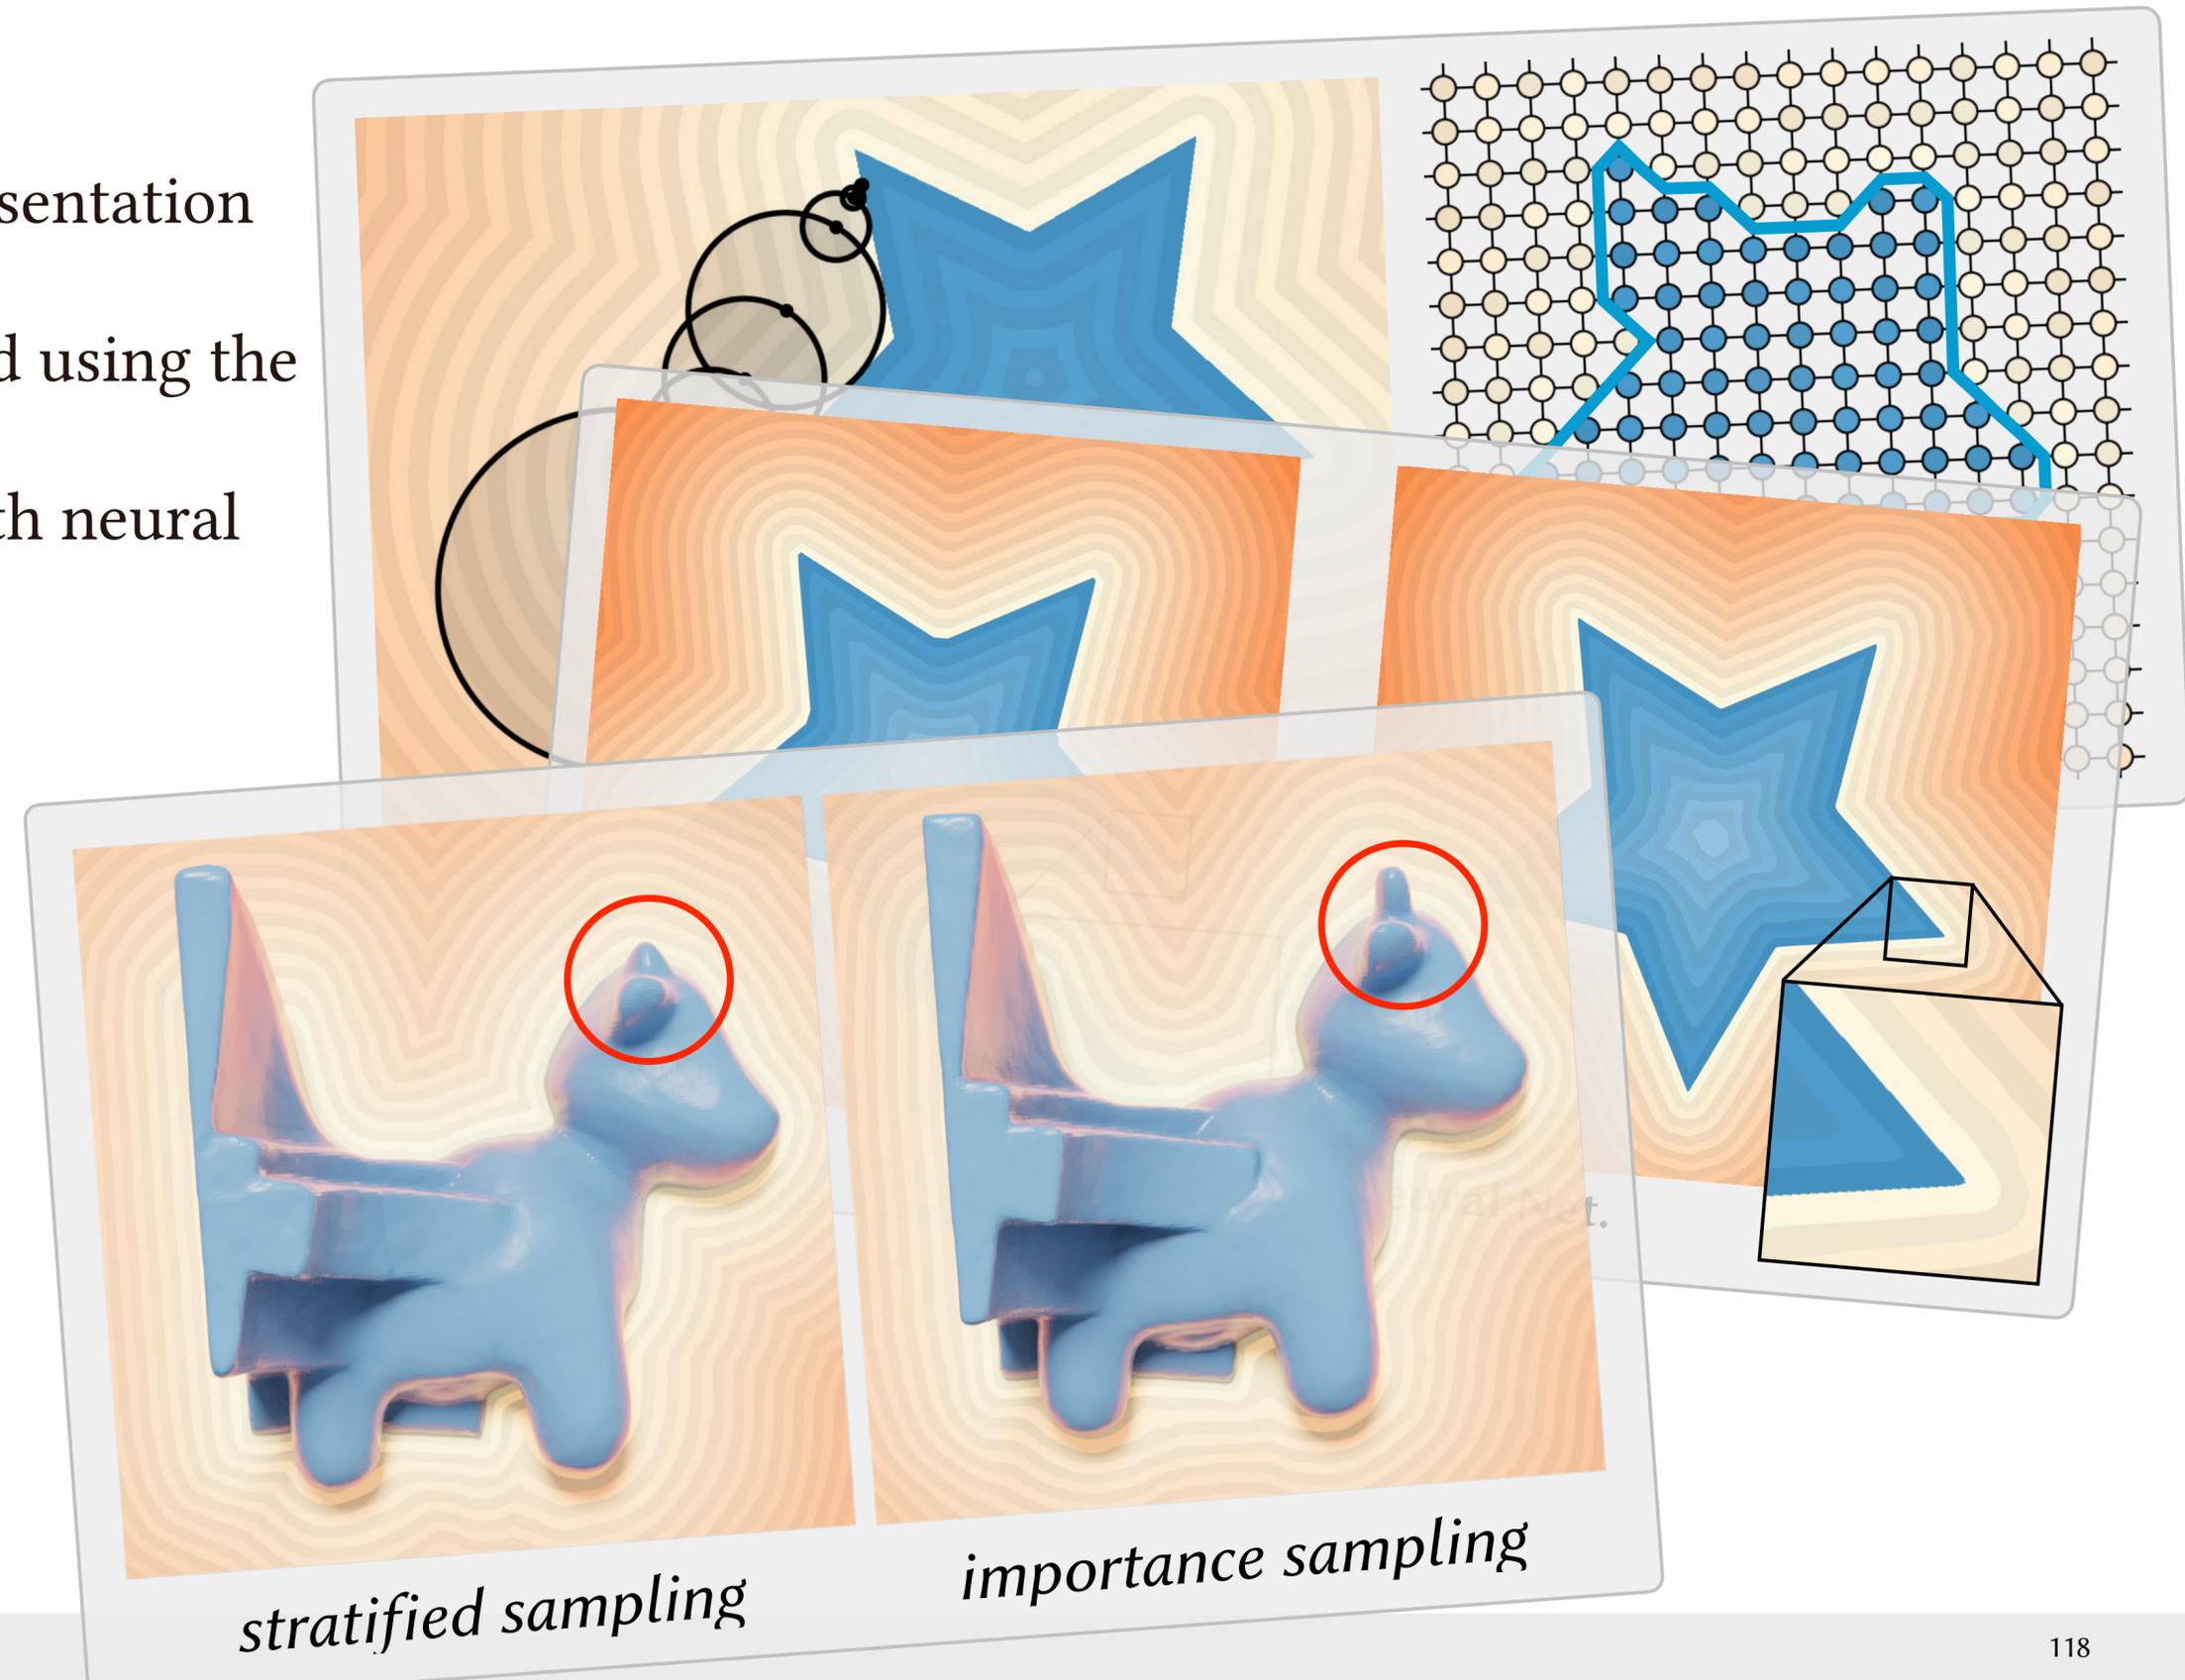
## We saw...

- how to work with a classic representation (implicit functions)
- implicit functions can be encoded using the neural network function space



## We saw...

- how to work with a classic representation (implicit functions)
- implicit functions can be encoded using the neural network function space
- some techniques for working with neural implicits

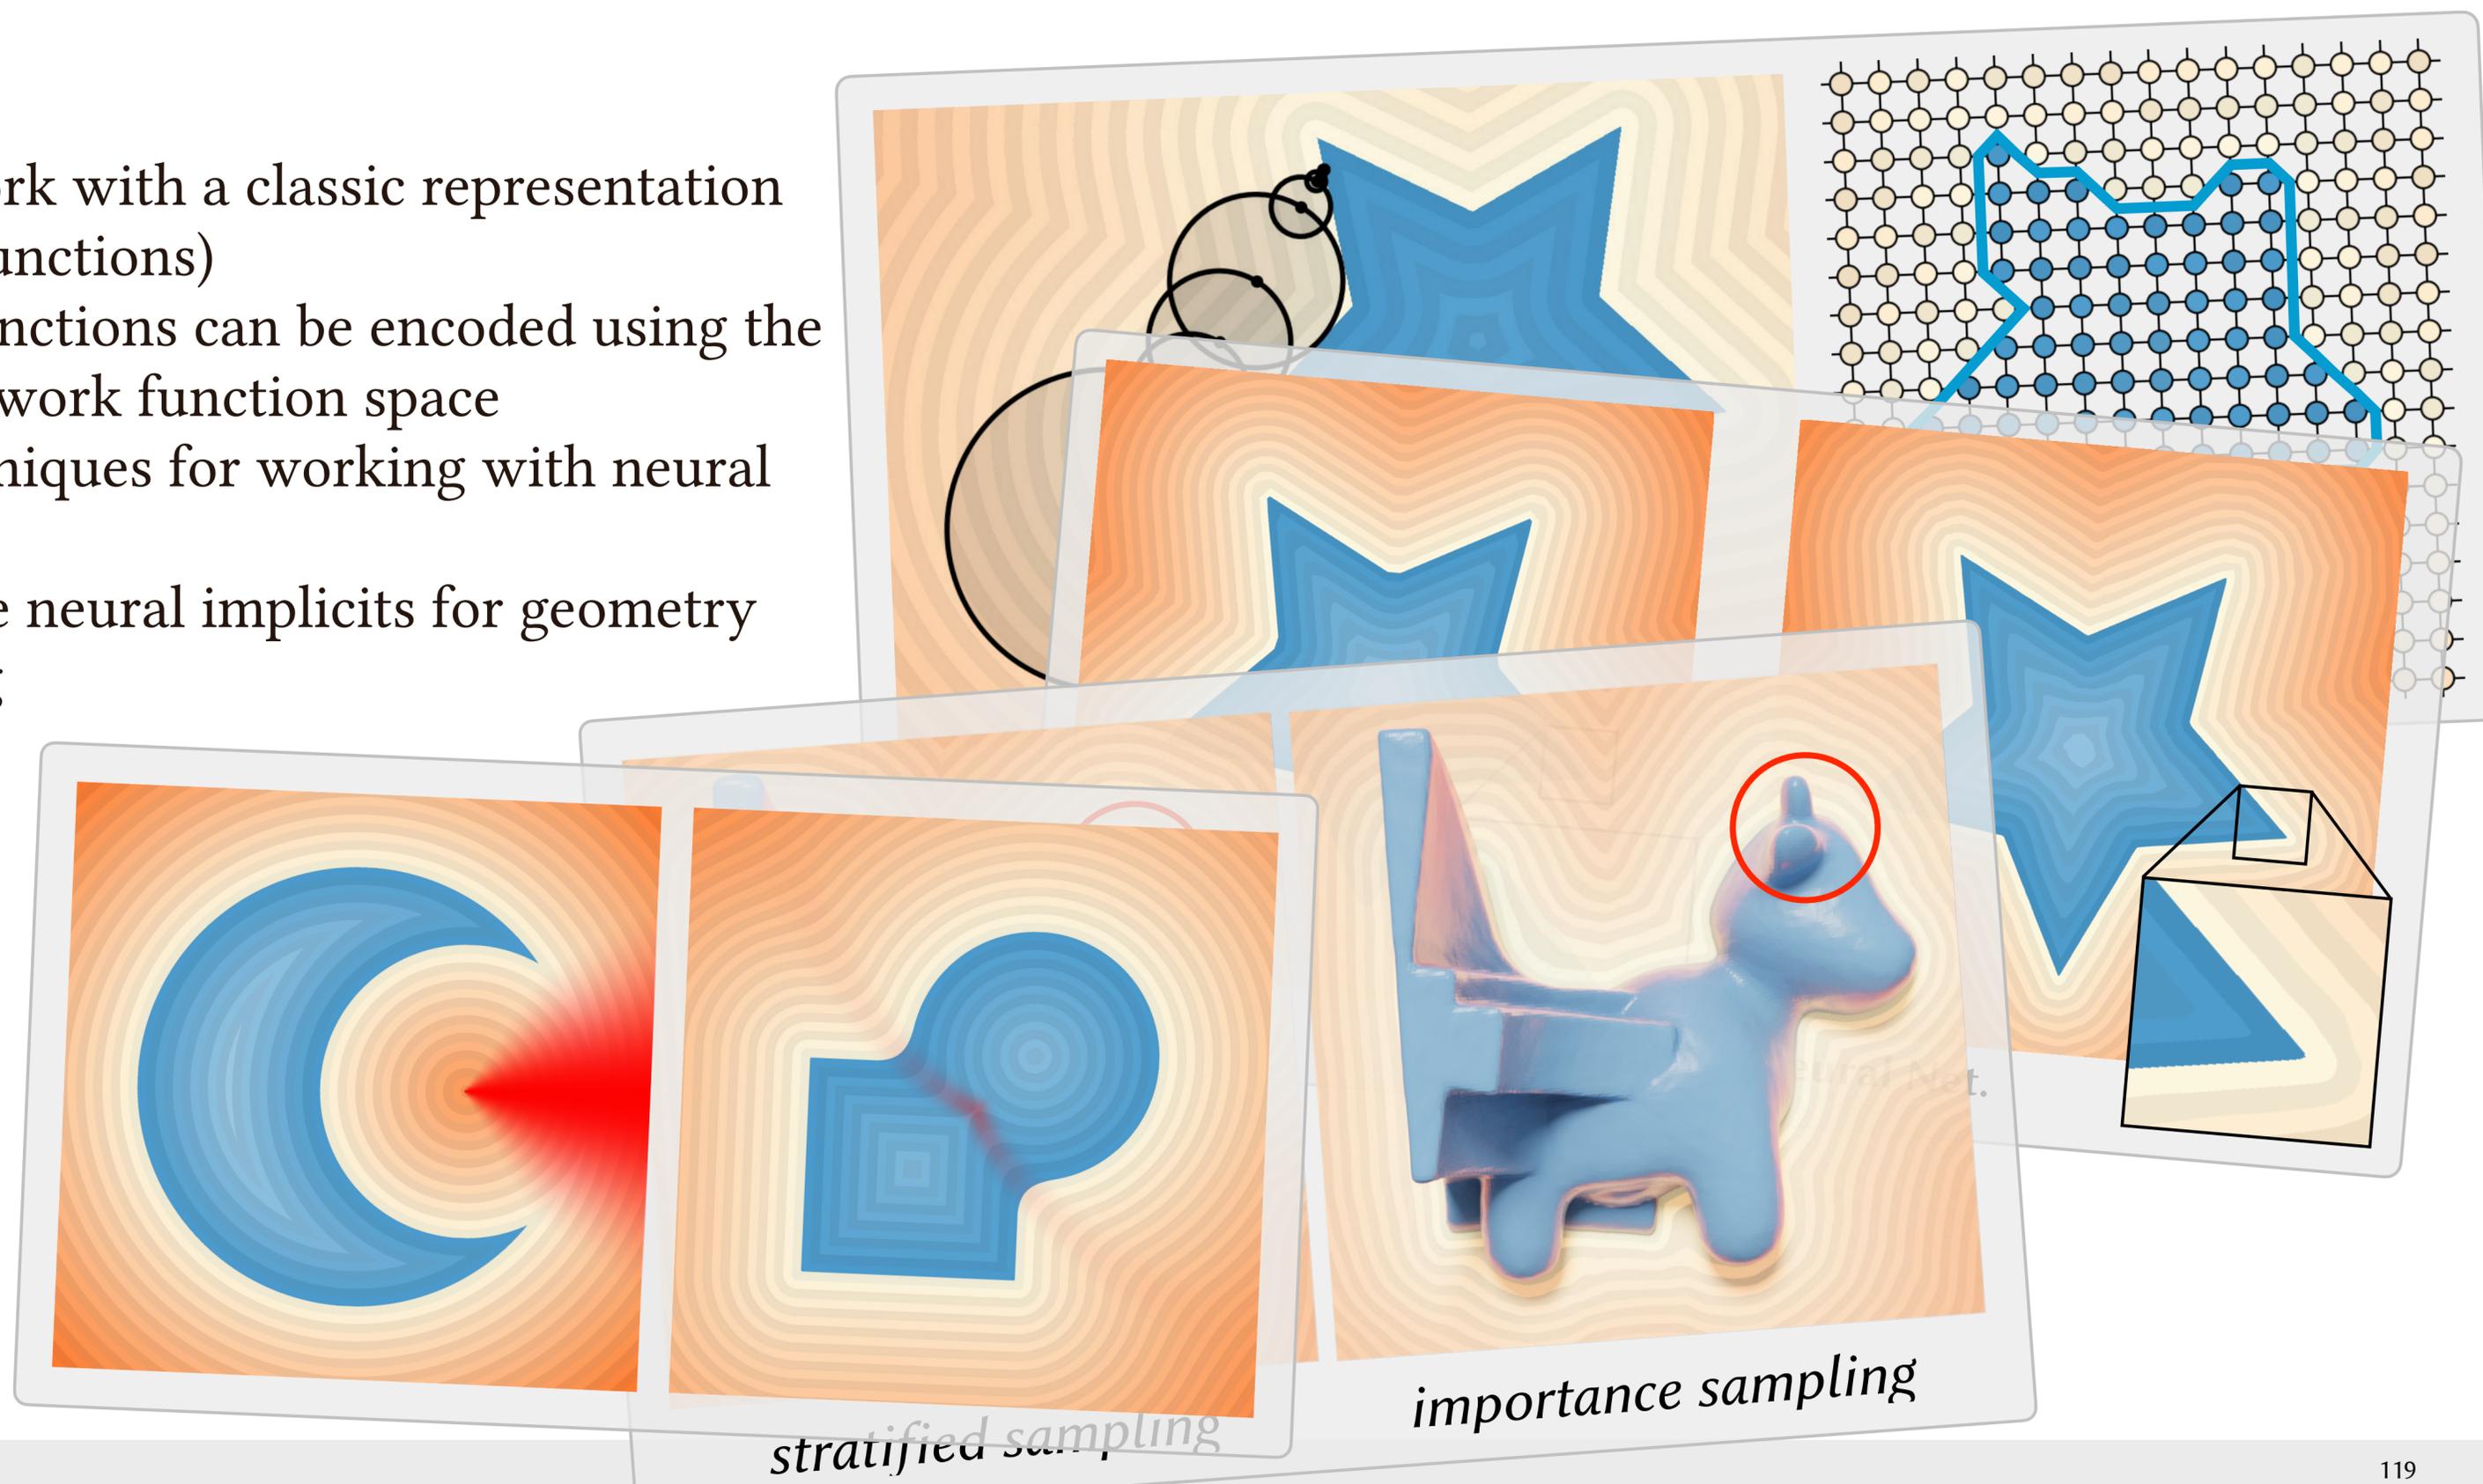


*stratified sampling*

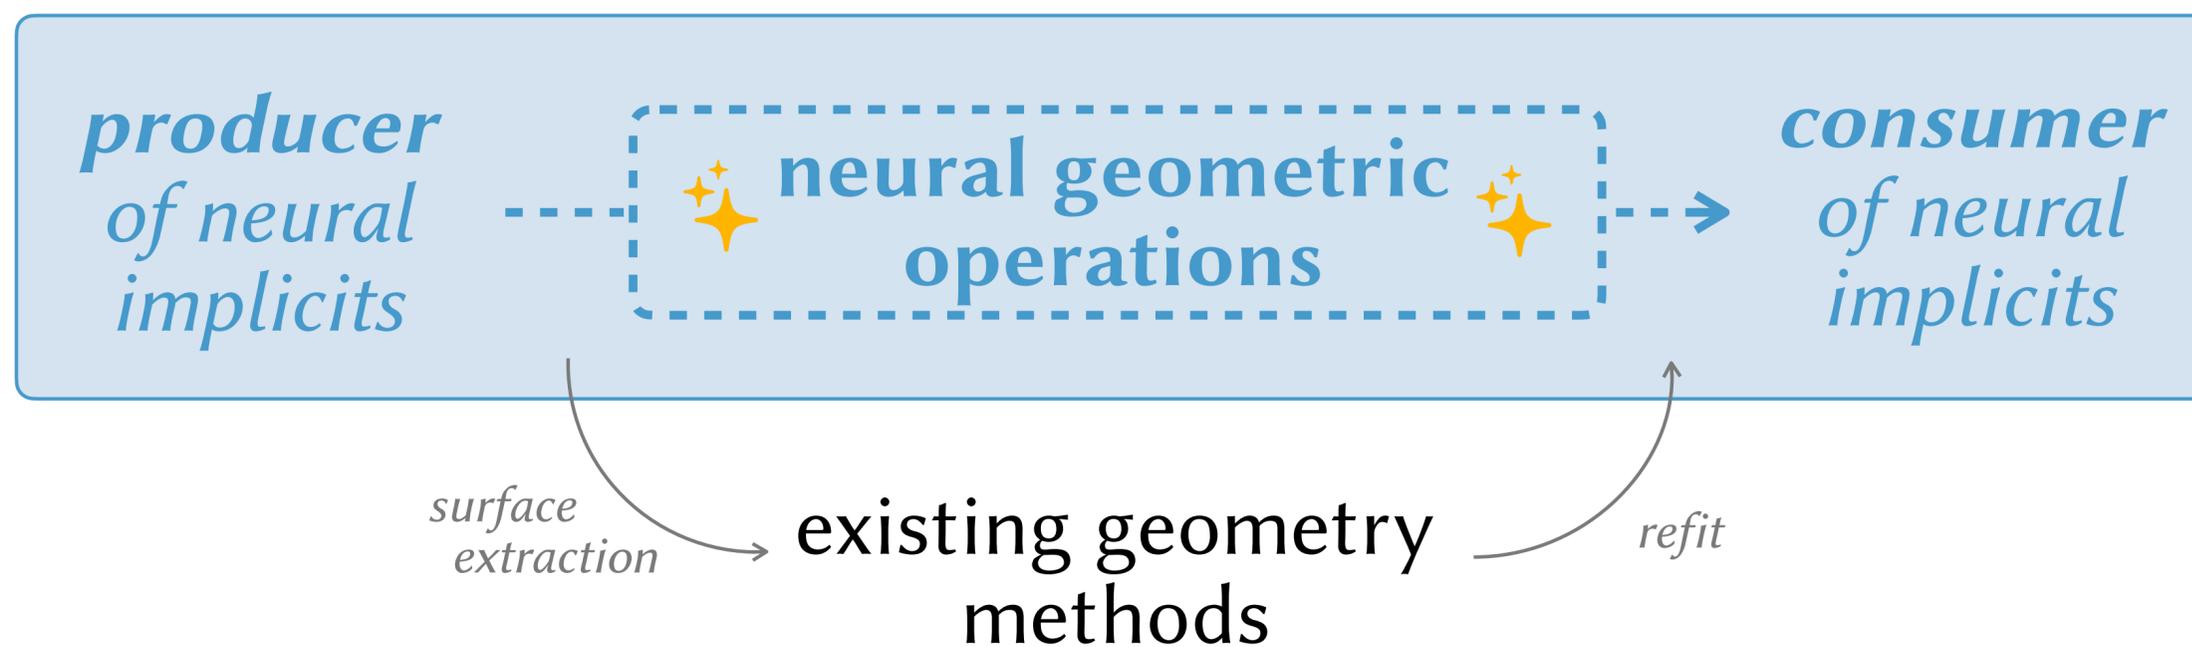
*importance sampling*

## We saw...

- how to work with a classic representation (implicit functions)
- implicit functions can be encoded using the neural network function space
- some techniques for working with neural implicits
- how to use neural implicits for geometry processing



## Neural World





Zoë Marschner [zoem@cmu.edu](mailto:zoem@cmu.edu)



Daniel Ritchie [daniel\\_ritchie@brown.edu](mailto:daniel_ritchie@brown.edu)

