Modelling Open Curves with Neural Implicit SDF

31-08-2021 Lu Meng

Content

Introduction

Introduction: Signed distance function (SDF)

- Definition:
  - The signed distance function determines the **distance** of a given point x from the boundary of a curve Ω.
  - The function is **negative** inside Ω and **positive** outside.
- The zero-value contour is the **boundary** of the curve.
- We can derive any curve from its SDF.

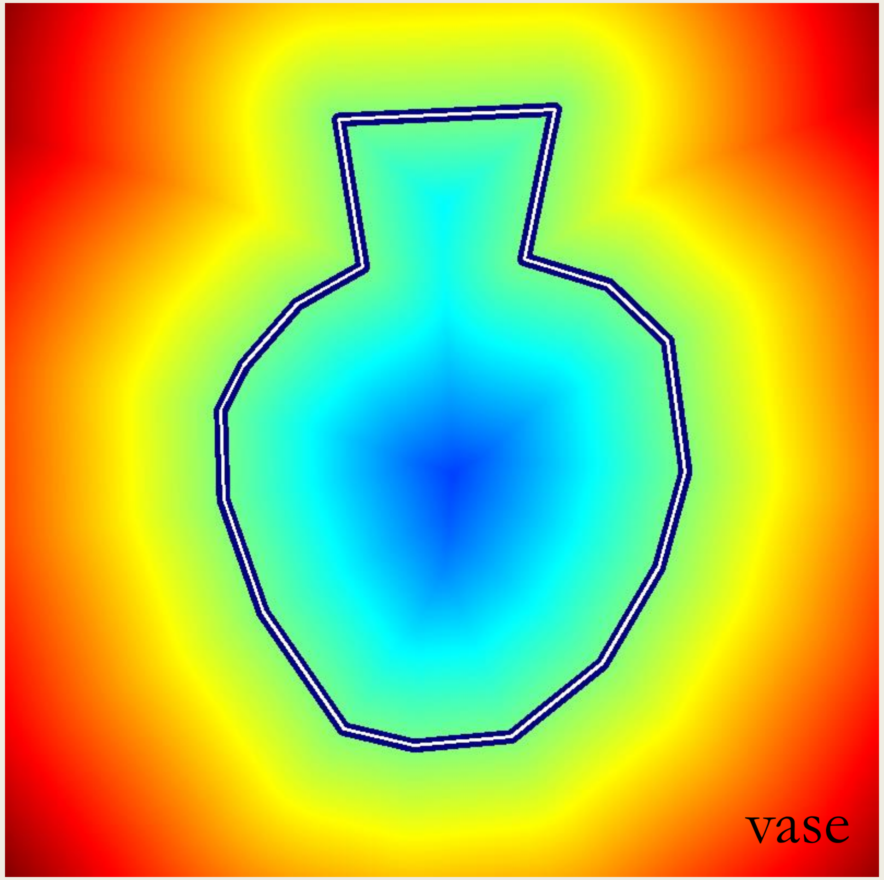

Red area: **positive** SDF
Blue area: **negative** SDF
Black line: **zero-value** SDF contour
White line: true curve

Introduction: Truncated signed distance function (TSDF)

$$TSDF = clamp(SDF, threshold)$$

- The zero-value contour is the **boundary** of the curve.

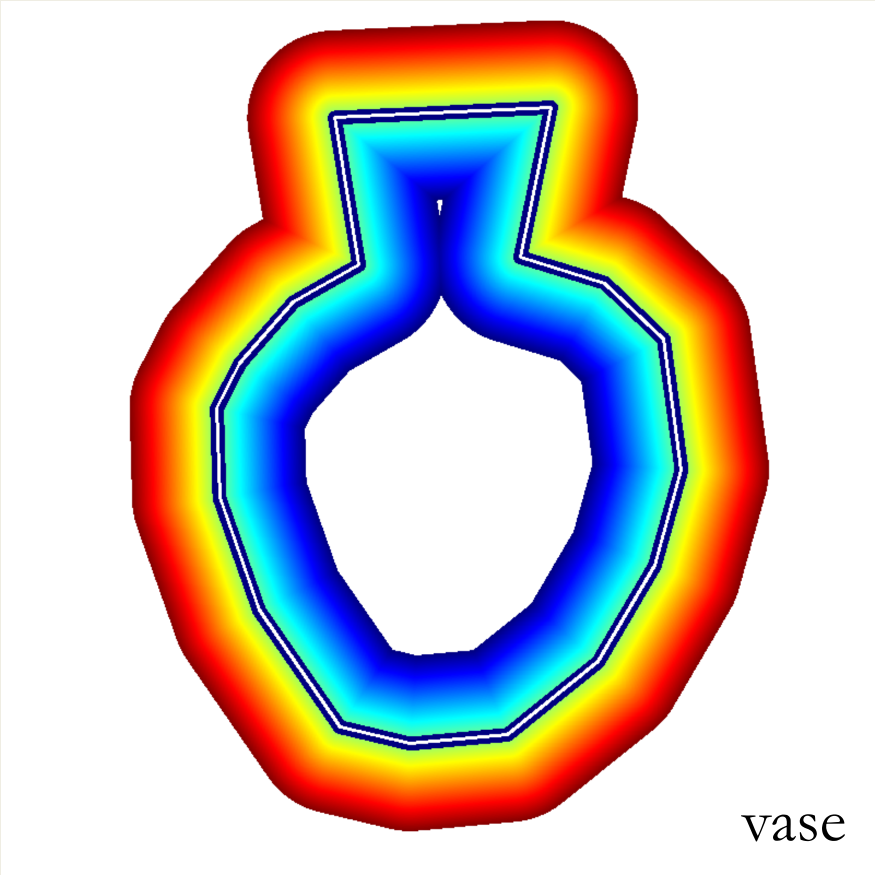

Red area: **positive** SDF
Blue area: **negative** SDF
Black line: **zero-value** SDF contour
White line: true curve

Introduction: TSDF of open curves

- The open curve is assigned a **direction**.
- The **sign** is determined by whether a point x is on the left or right of the curve.

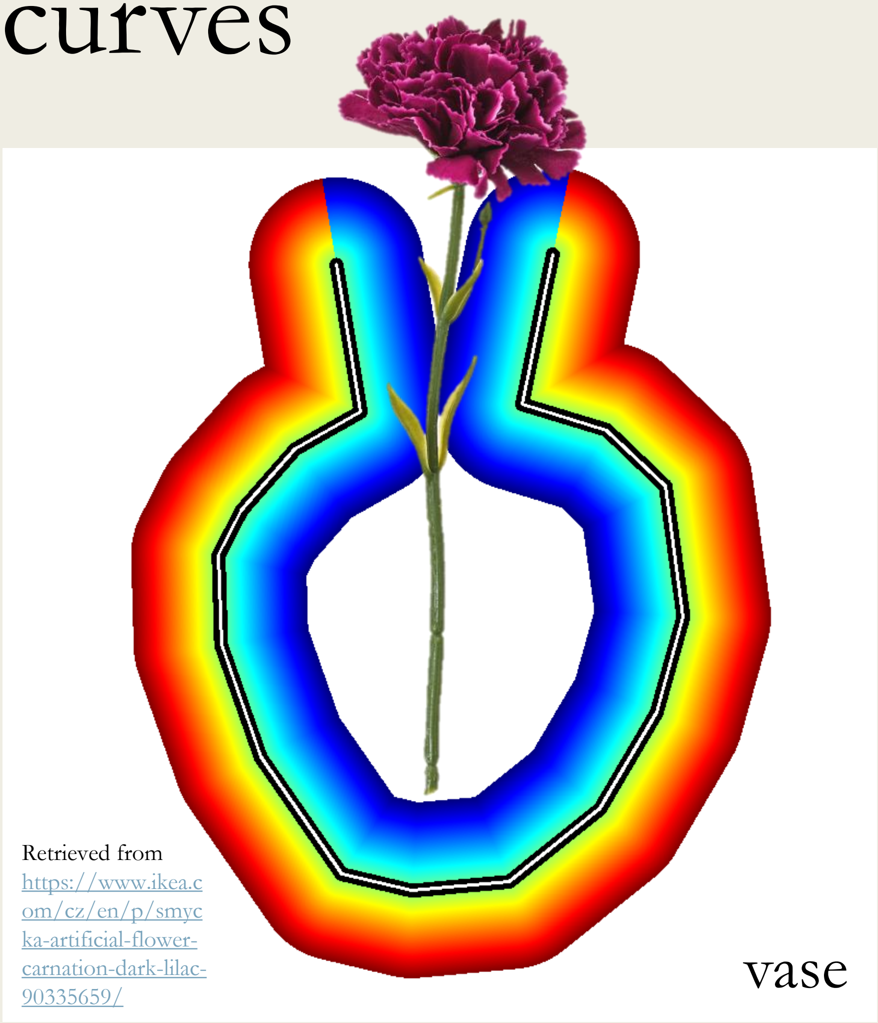

Red area: **positive** SDF
Blue area: **negative** SDF
Black line: **zero-value** SDF contour
White line: true curve

Introduction: TSDF of open curves

- The open curve is assigned a **direction**.
- The **sign** is determined by whether a point x is on the left or right of the curve.

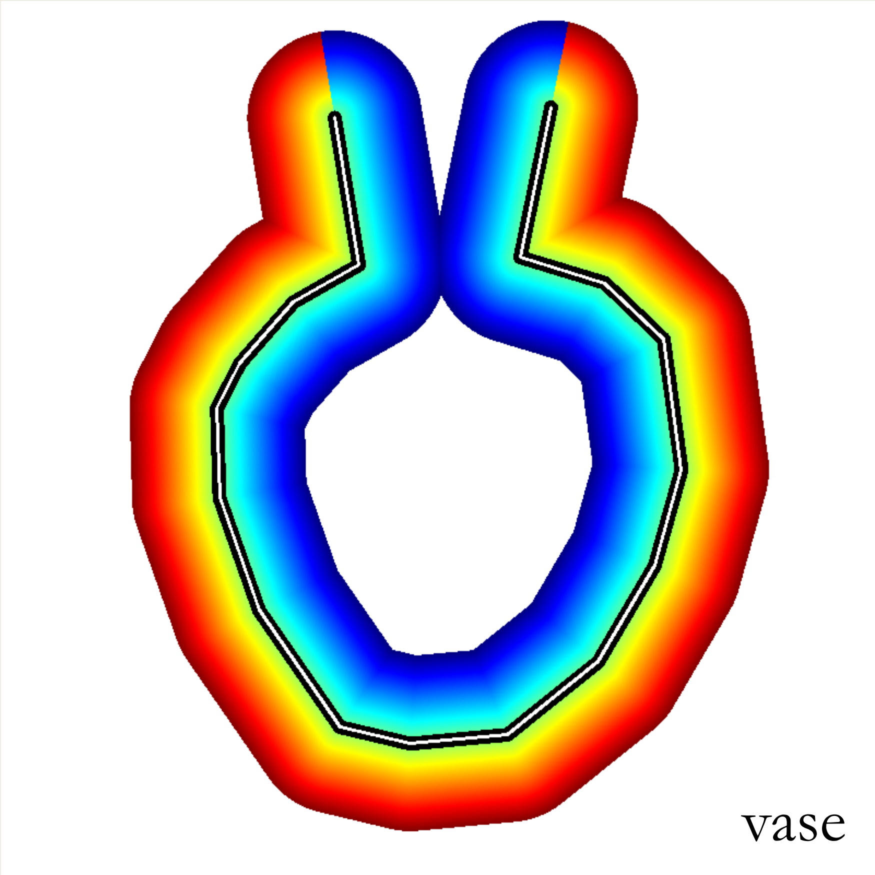

Red area: **positive** SDF
Blue area: **negative** SDF
Black line: **zero-value** SDF contour
White line: true curve

Network

Network: Idea

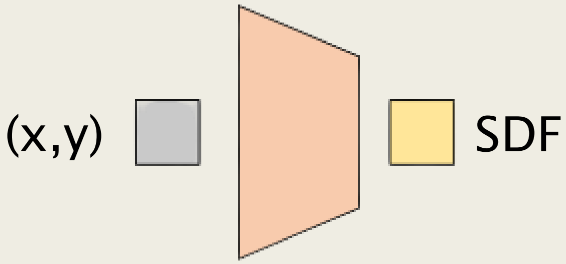

$$f_\theta(\boldsymbol{x}) \approx TSDF(\boldsymbol{x}), \forall \boldsymbol{x} \in \Omega$$

- Use a **neural network** to model the TSDF of a certain curve.
- The **curve information** is contained in the weight $\boldsymbol{\theta}$ of the network.

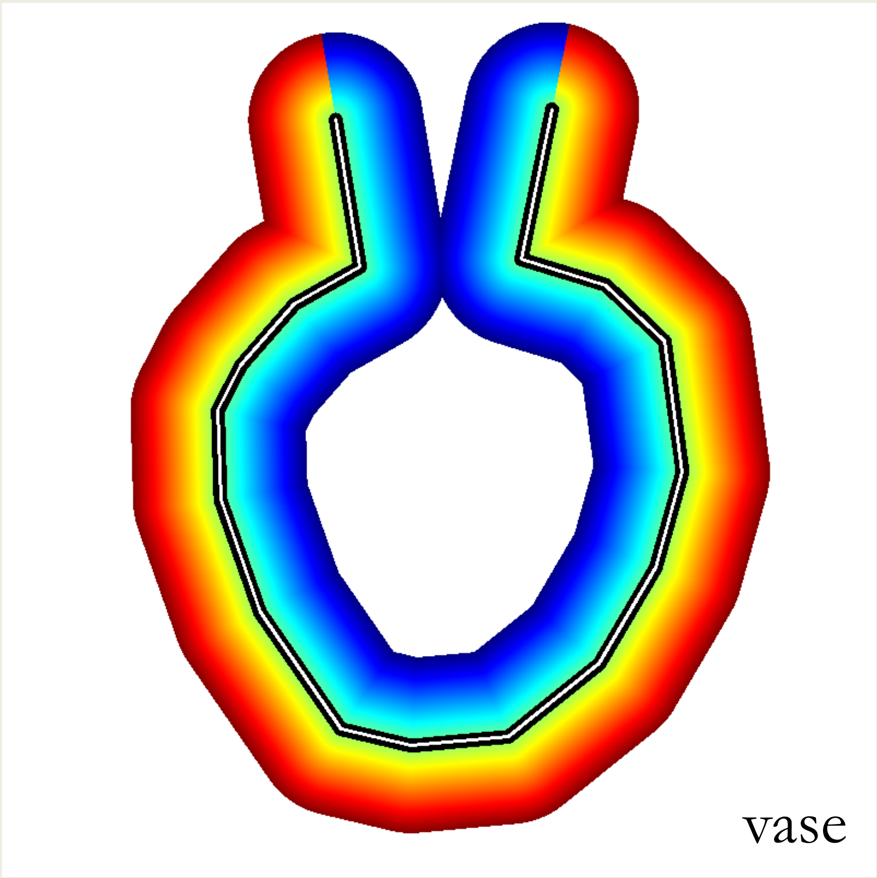

Red area: **positive** SDF
Blue area: **negative** SDF
Black line: **zero-value** SDF contour
White line: true curve

Network: Data preparation

$X := \{(\boldsymbol{x}, s) : SDF(\boldsymbol{x}) = s\}$

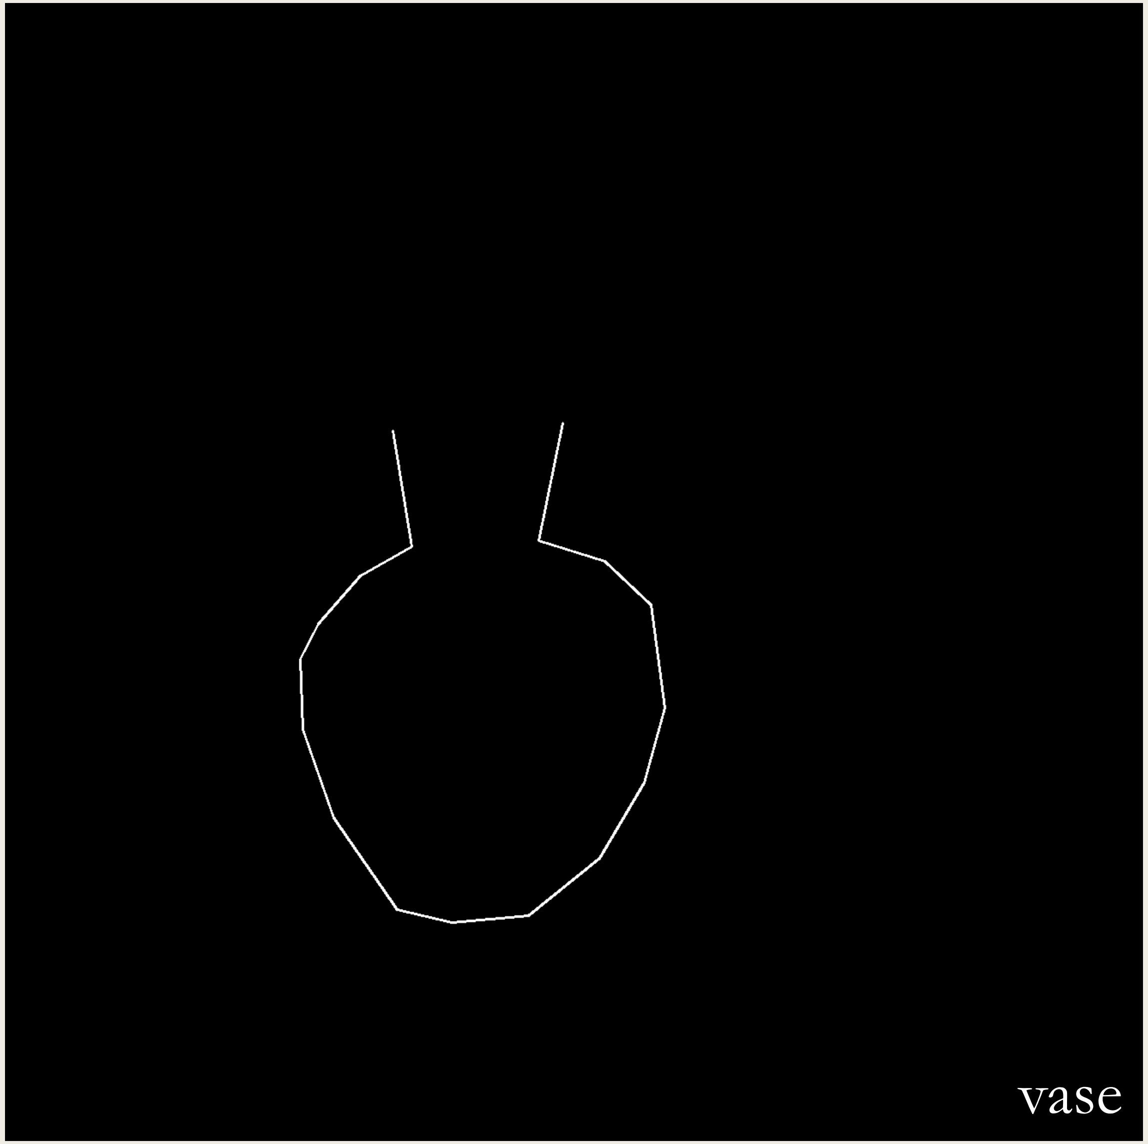

Network: Data preparation

$$X := \{(\boldsymbol{x}, s) : SDF(\boldsymbol{x}) = s\}$$

- **Normalize** the curve so that it is bounded by a unit circle.

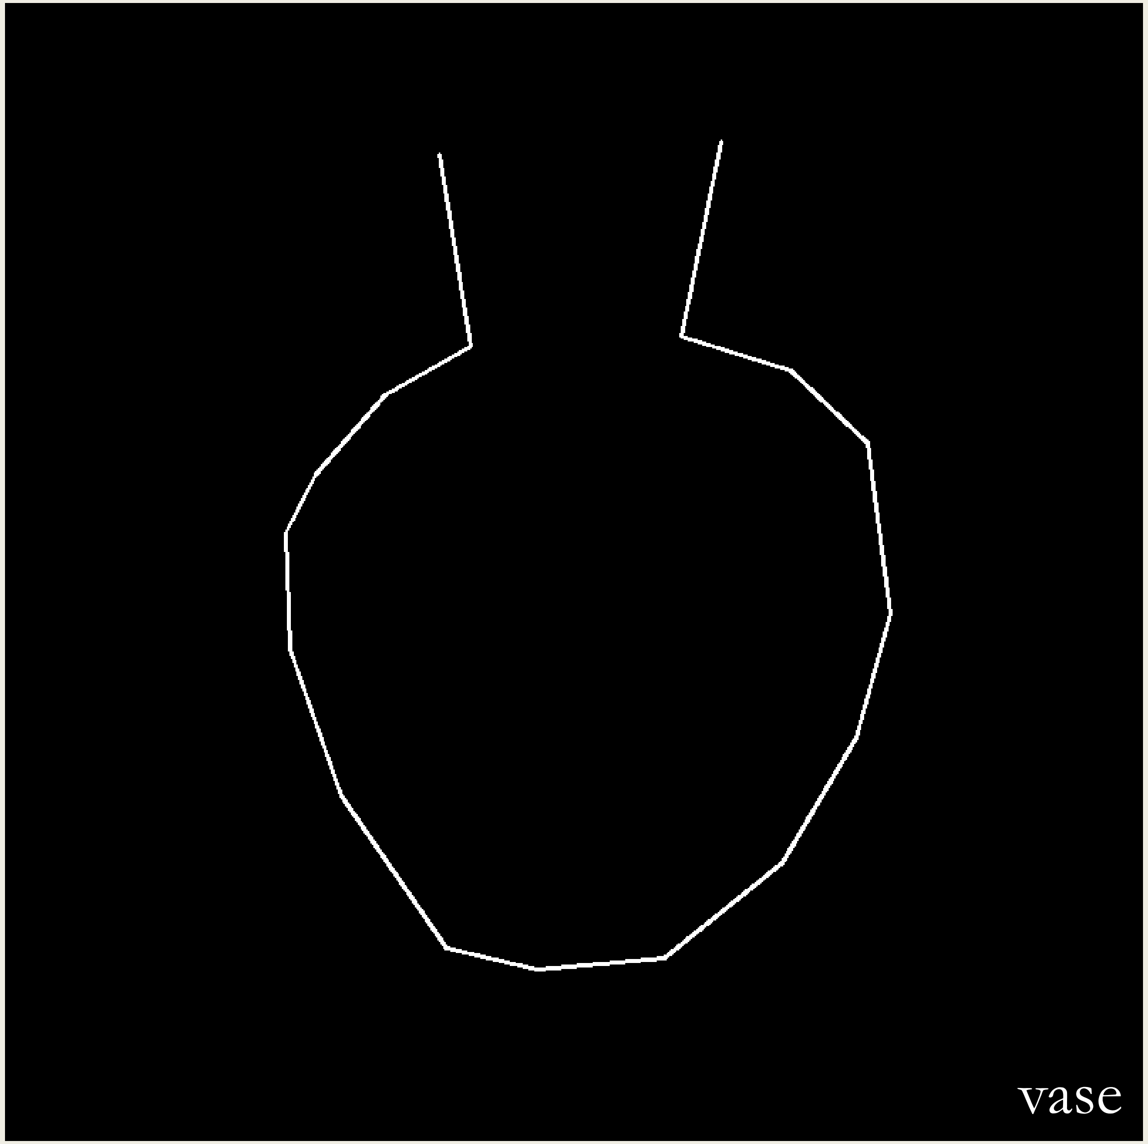

Network: Data preparation

$$X := \{(\boldsymbol{x}, s) : SDF(\boldsymbol{x}) = s\}$$

- **Normalize** the curve so that it is bounded by a unit circle.
- Generate 2,000 **uniform** sampling points.

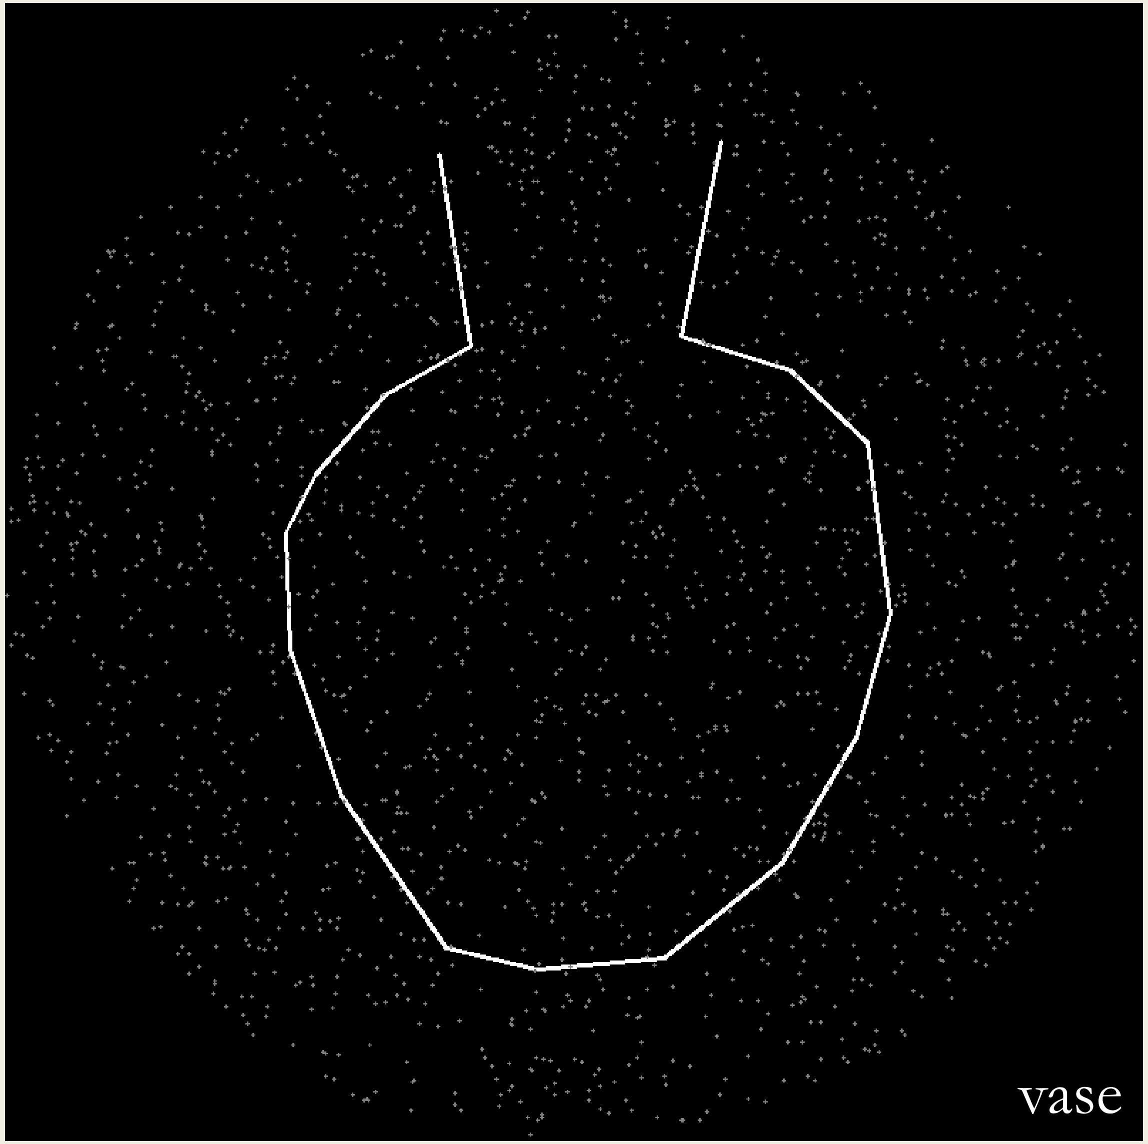

Network: Data preparation

$$X := \{(\boldsymbol{x}, s) : SDF(\boldsymbol{x}) = s\}$$

- **Normalize** the curve so that it is bounded by a unit circle.
- Generate 2,000 **uniform** sampling points.
- Generate 12,000 **Gaussian** sampling points.
- Obtain more information about the **zero-value set**.

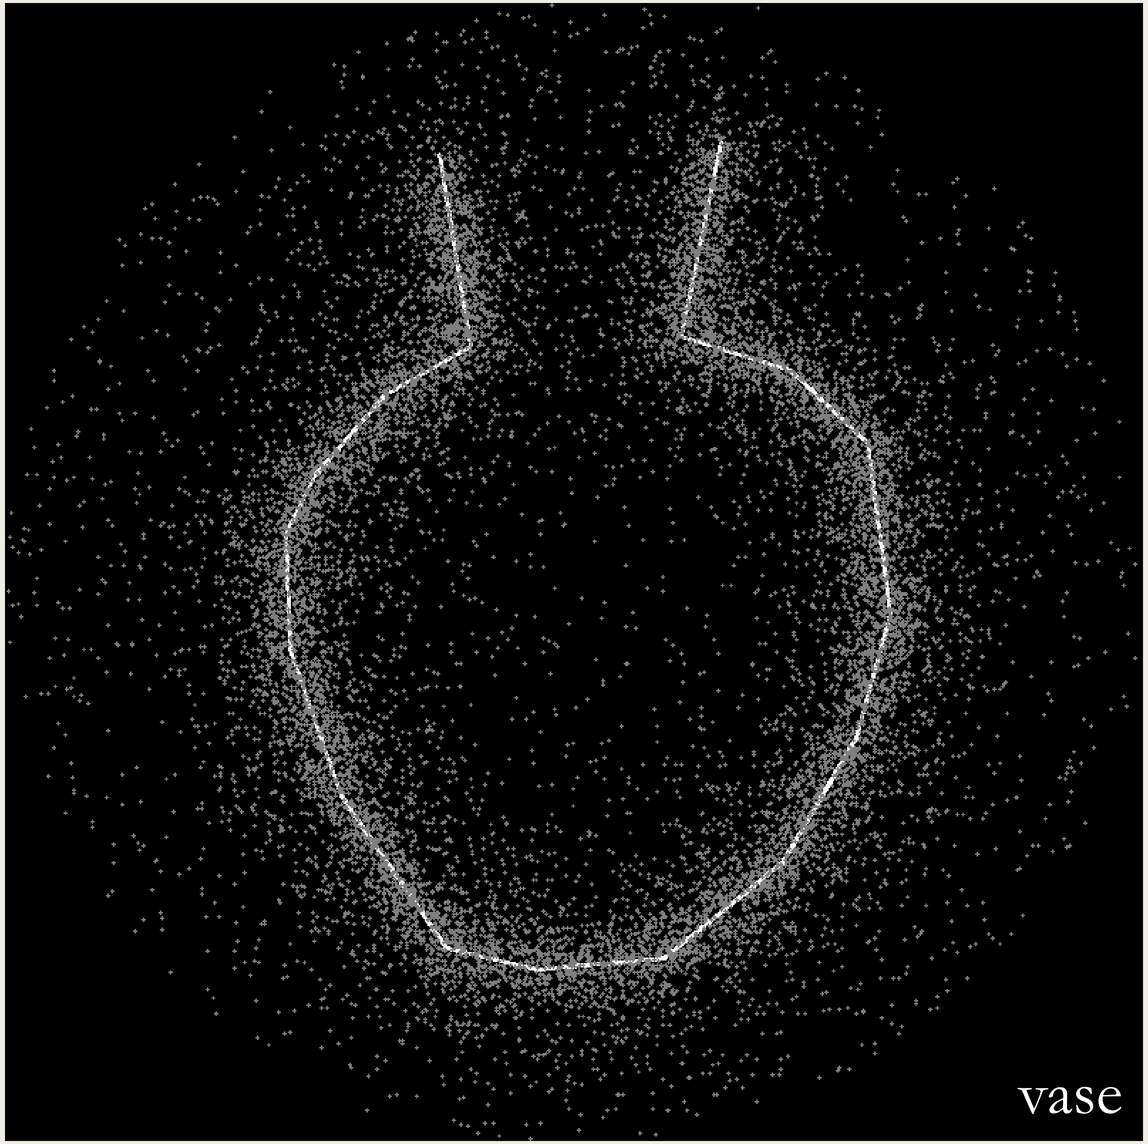

Network: Data preparation

$$X := \{(\boldsymbol{x}, s) : SDF(\boldsymbol{x}) = s\}$$

- Calculate the corresponding **ground truth SDF**.
  - Find the **minimum value** among the distances from the point to each segment.
  - Determine the **sign** according to which side of the curve the point is on.

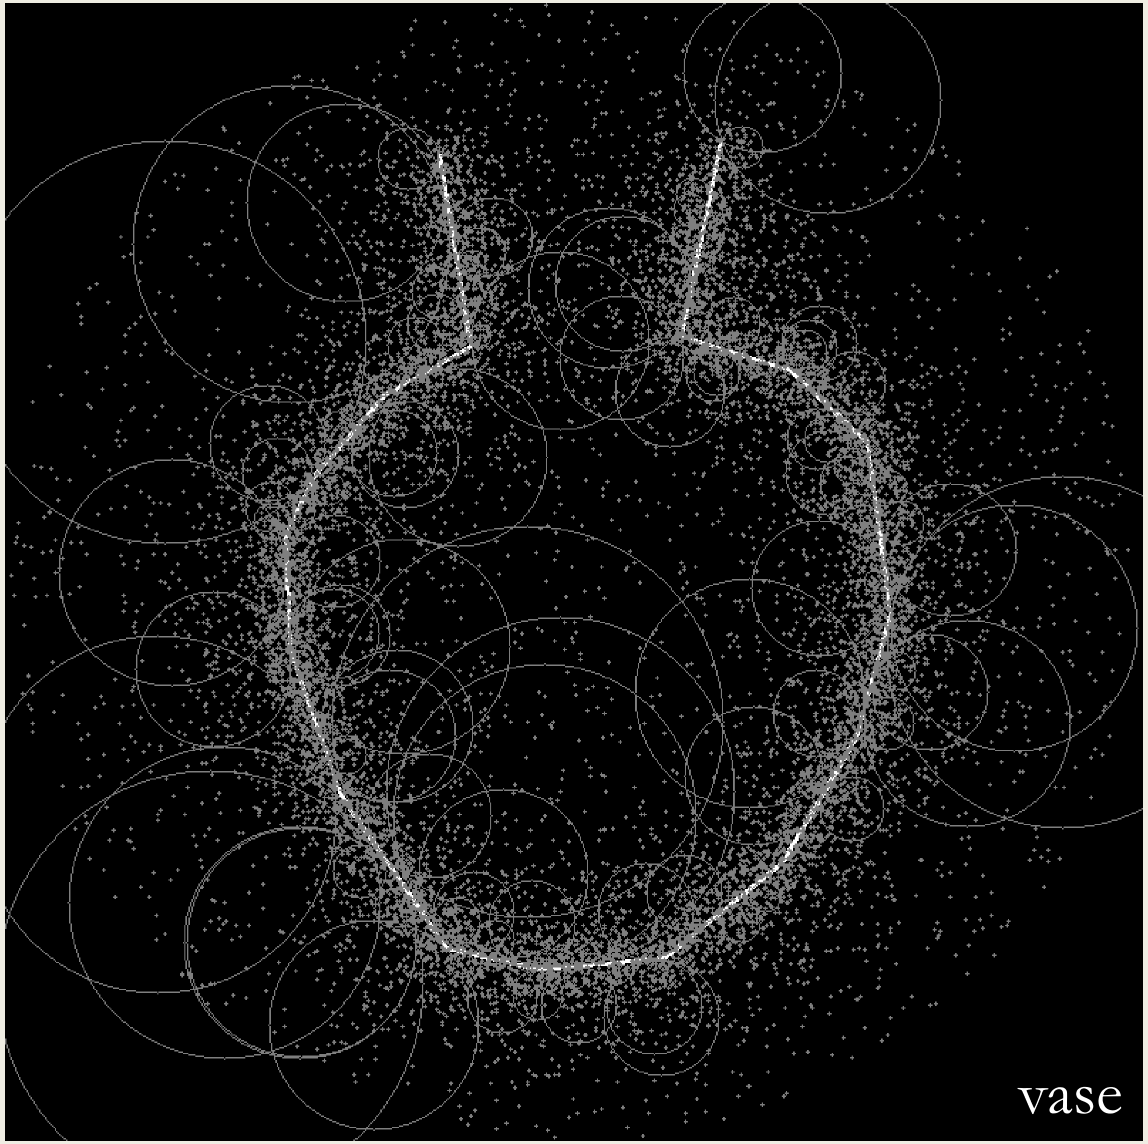

Network: Structure

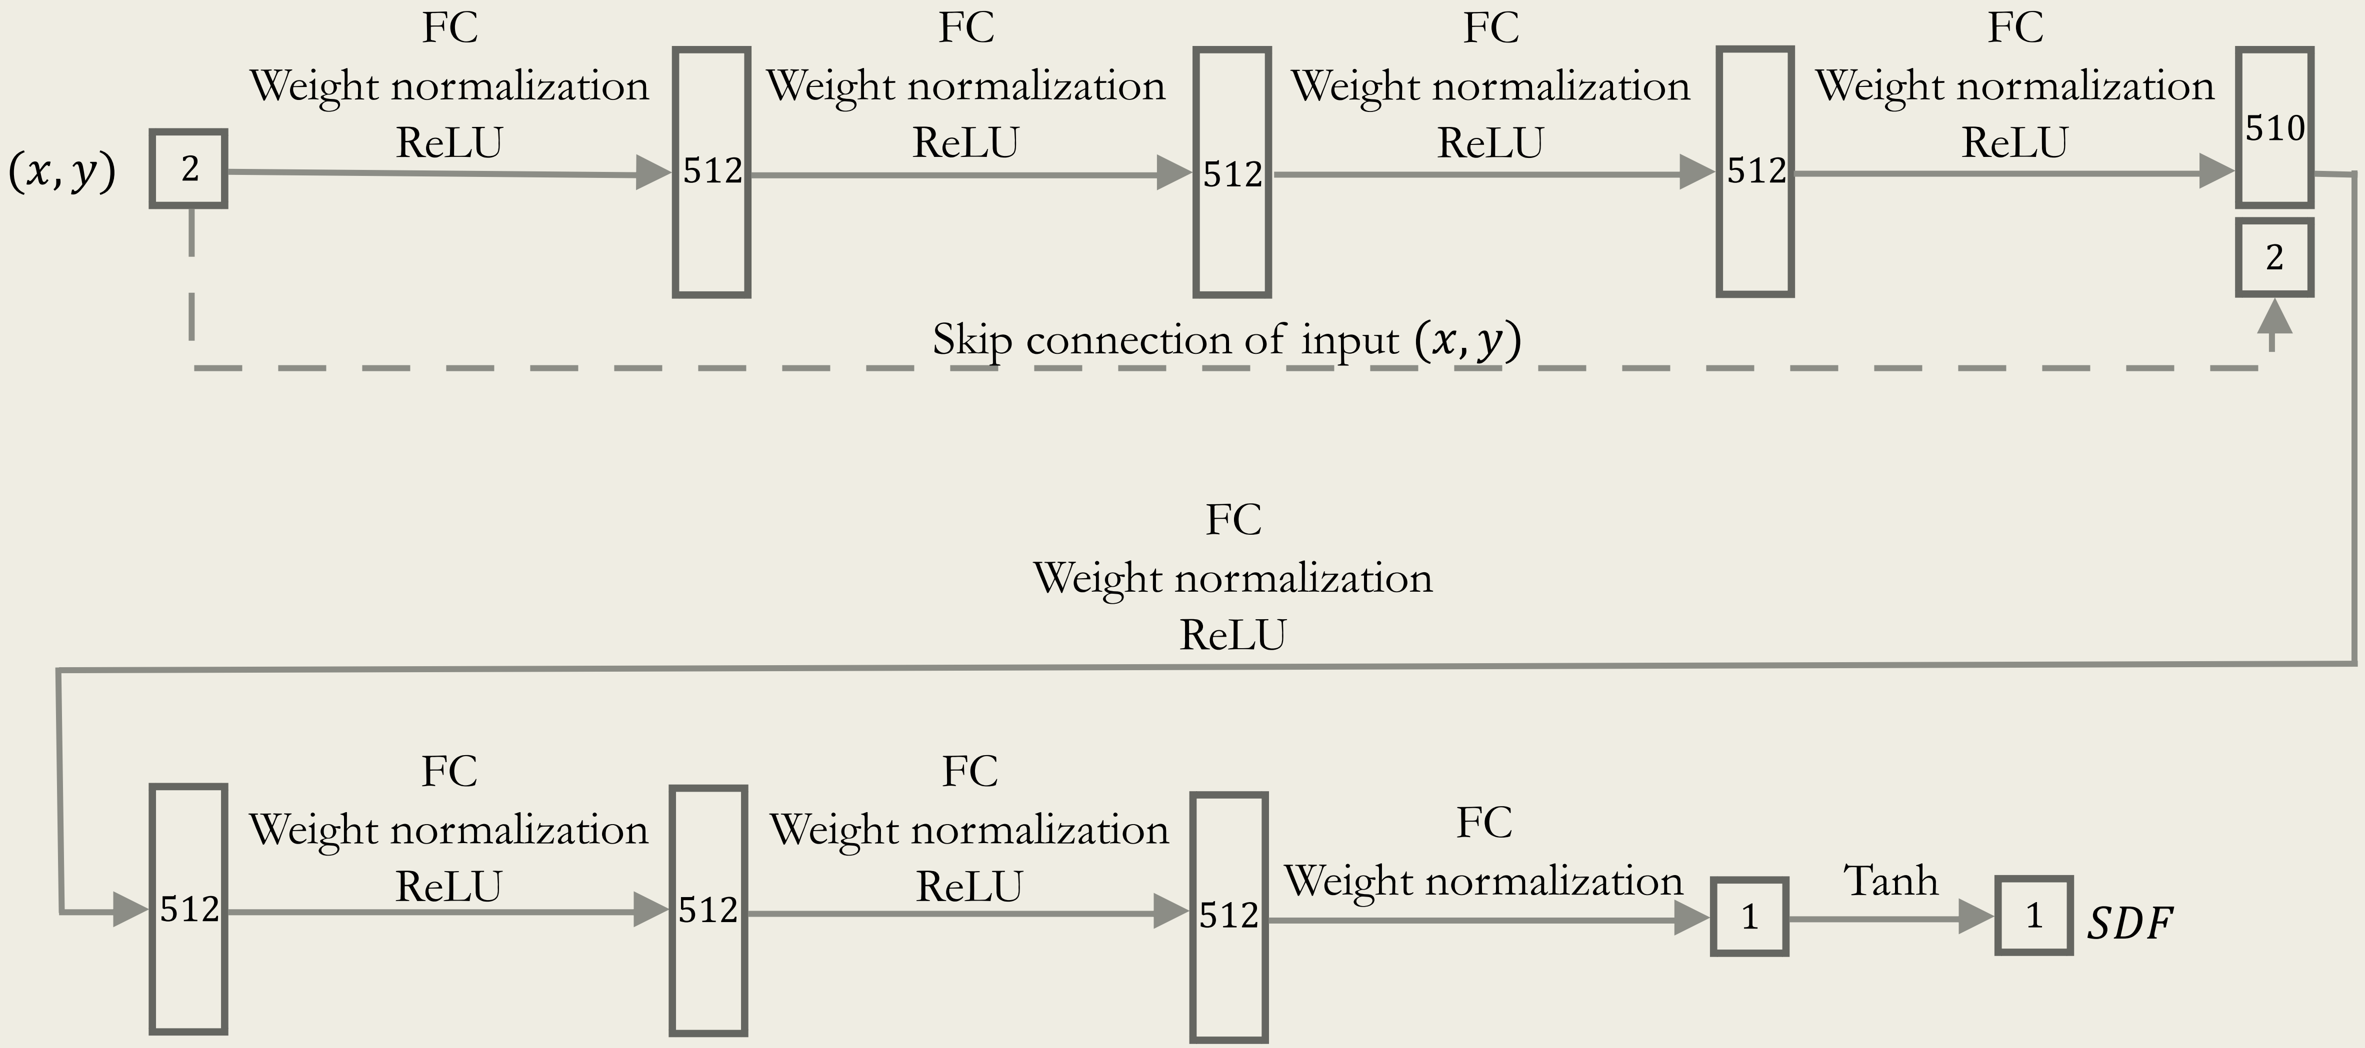

Network: Training details

- Loss function:
  - $\mathcal{L}(f_\theta(\boldsymbol{x}), s) = |\mathrm{clamp}(f_\theta(\boldsymbol{x}), \delta) - \mathrm{clamp}(s, \delta)|$
- Optimizer:
  - Adam
- Configs:
  - Learning rate: 1e-5
  - Epochs: 1,000
  - Batch size: 100

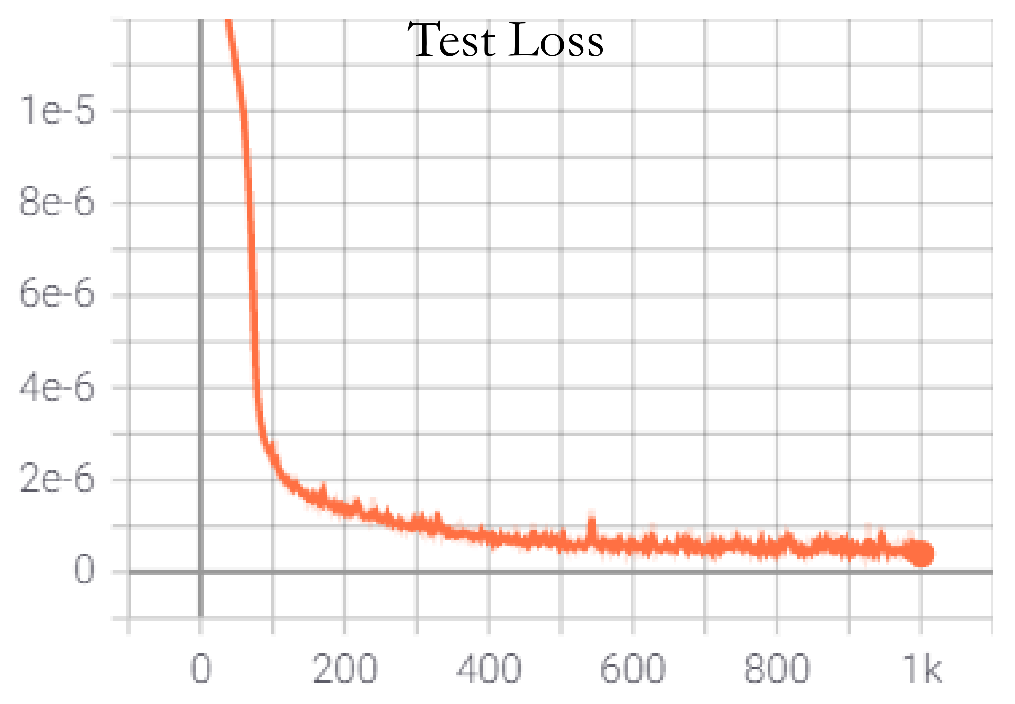

Problems & Solutions

Problems: Curves

Normalized Curves

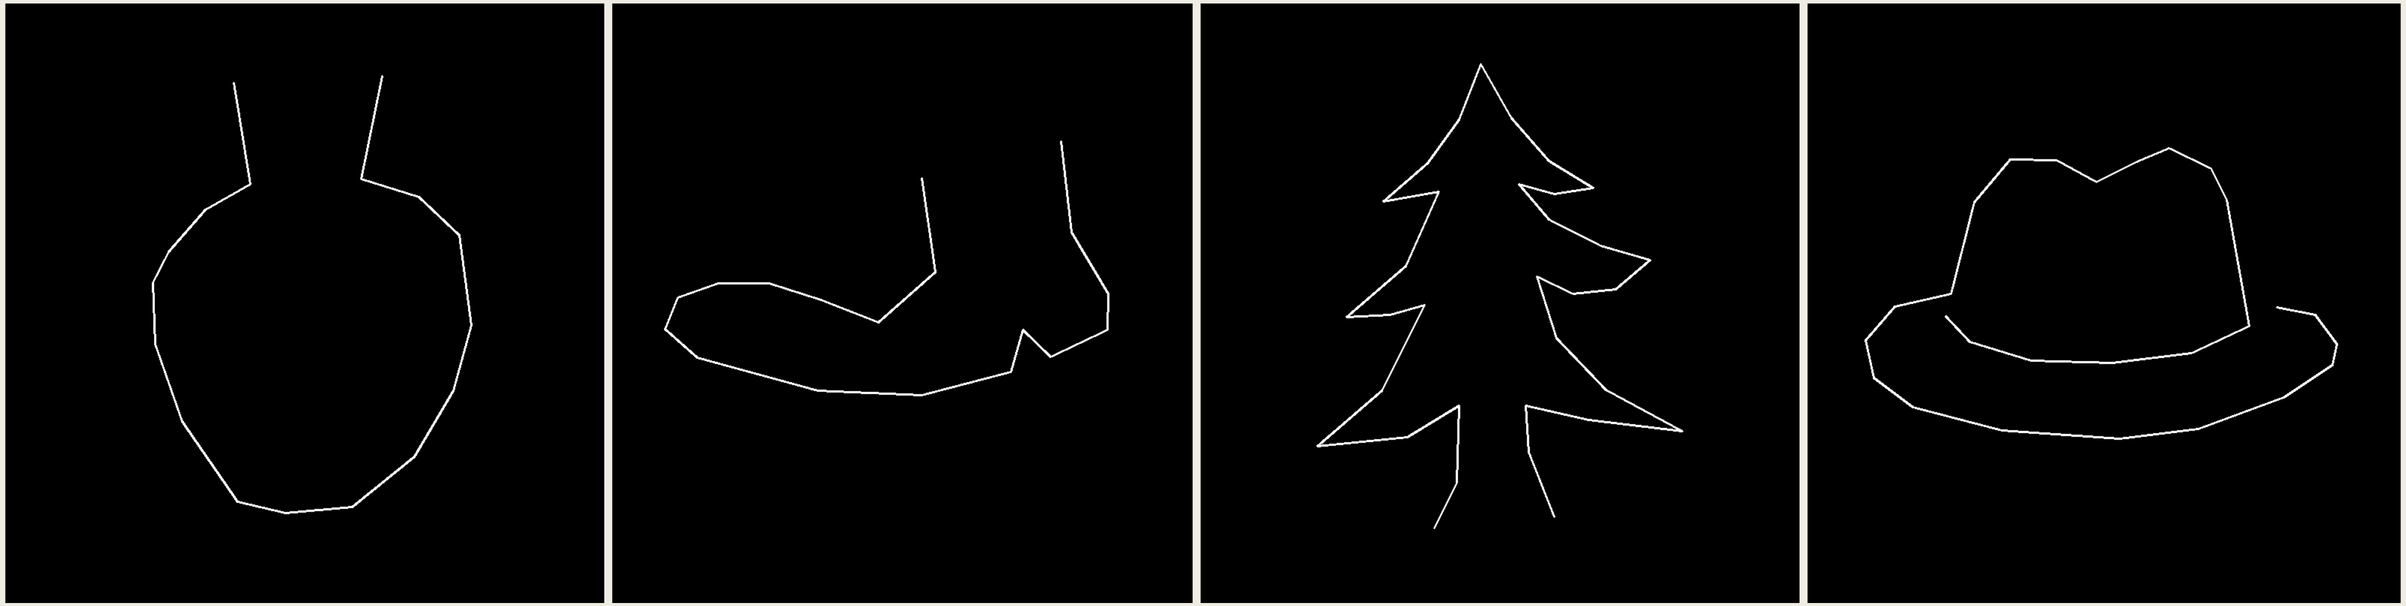

vase shoe tree hat

Problems: Outliers

Red area: **positive** SDF
Blue area: **negative** SDF
Black line: **zero-value** SDF contour
White line: true curve

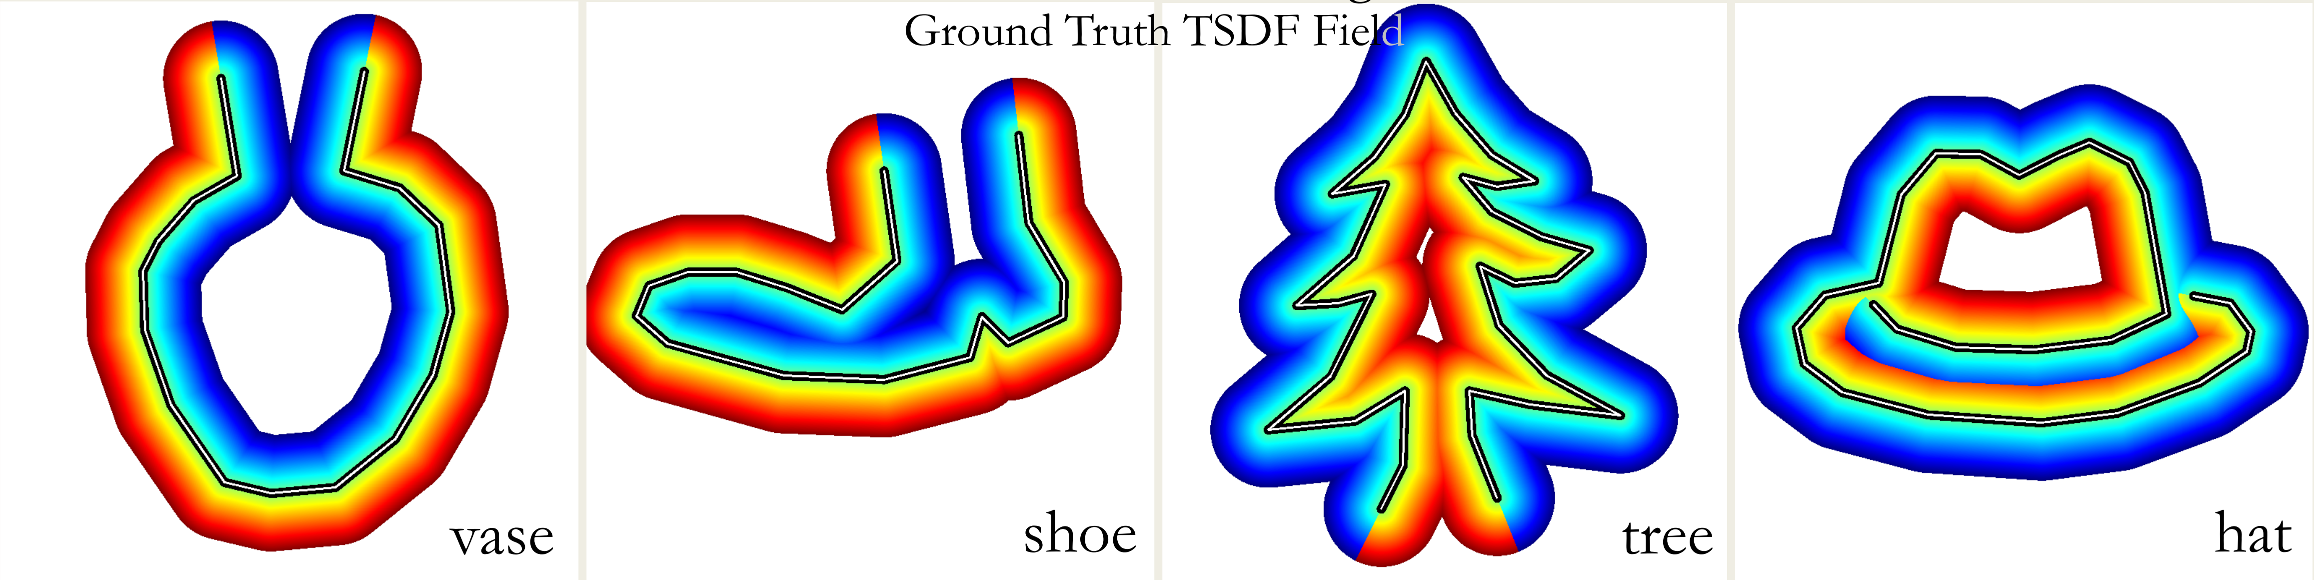

	vase	shoe	tree	hat
Avg. L1 error:	$2.364e-3$	$1.445e-3$	$1.585e-3$	$8.365e-4$
Max L1 error:	$1.309e-1$	$1.127e-1$	$1.196e-1$	$1.340e-1$

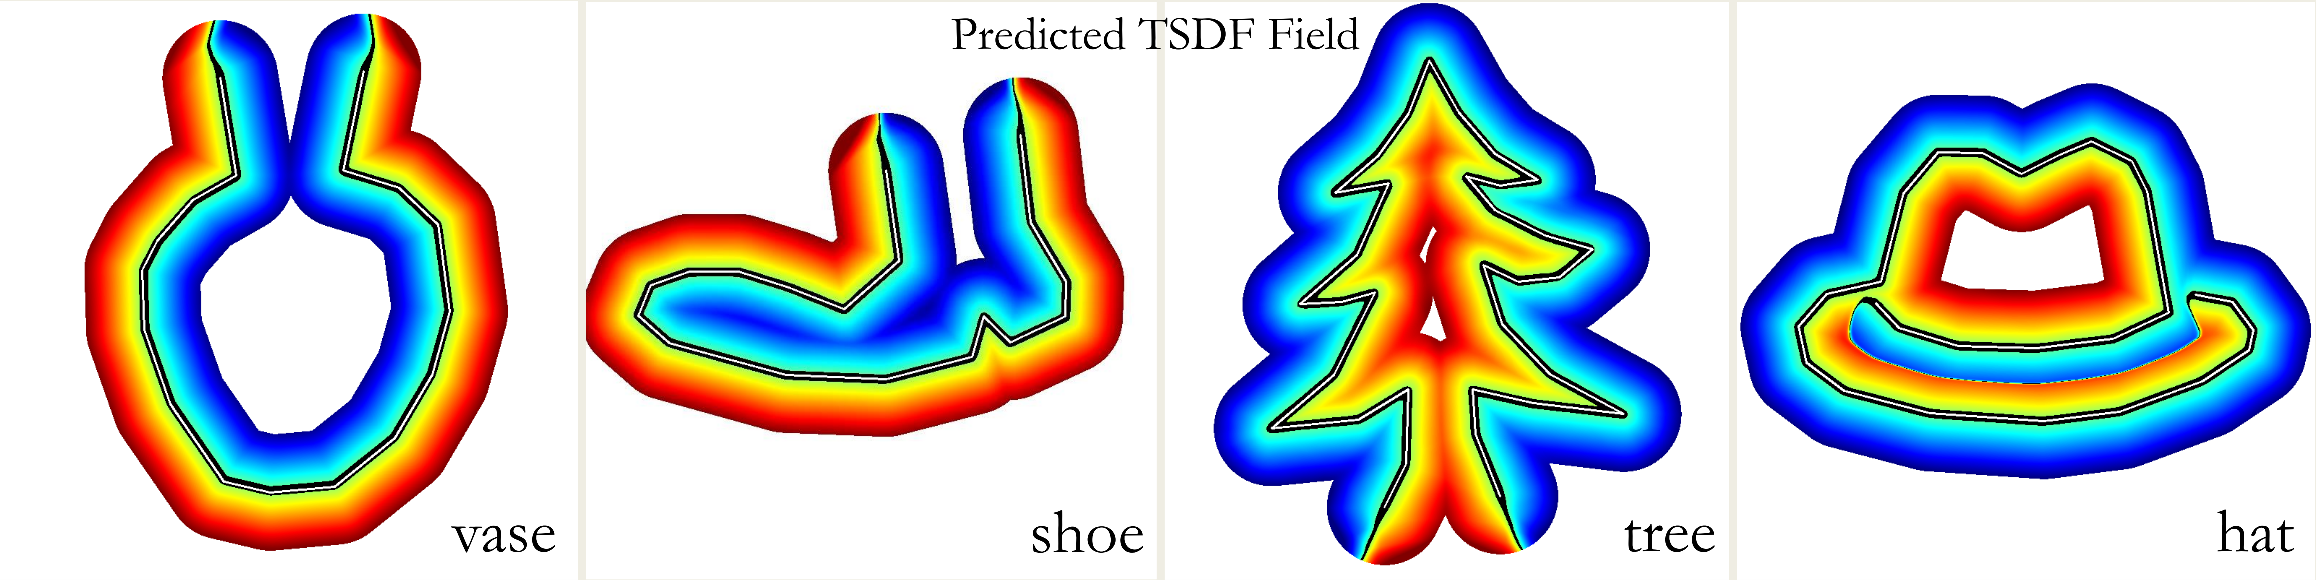

Problems: Outliers

Red area: **positive** SDF
Blue area: **negative** SDF
Black line: **zero-value** SDF contour
White line: true curve

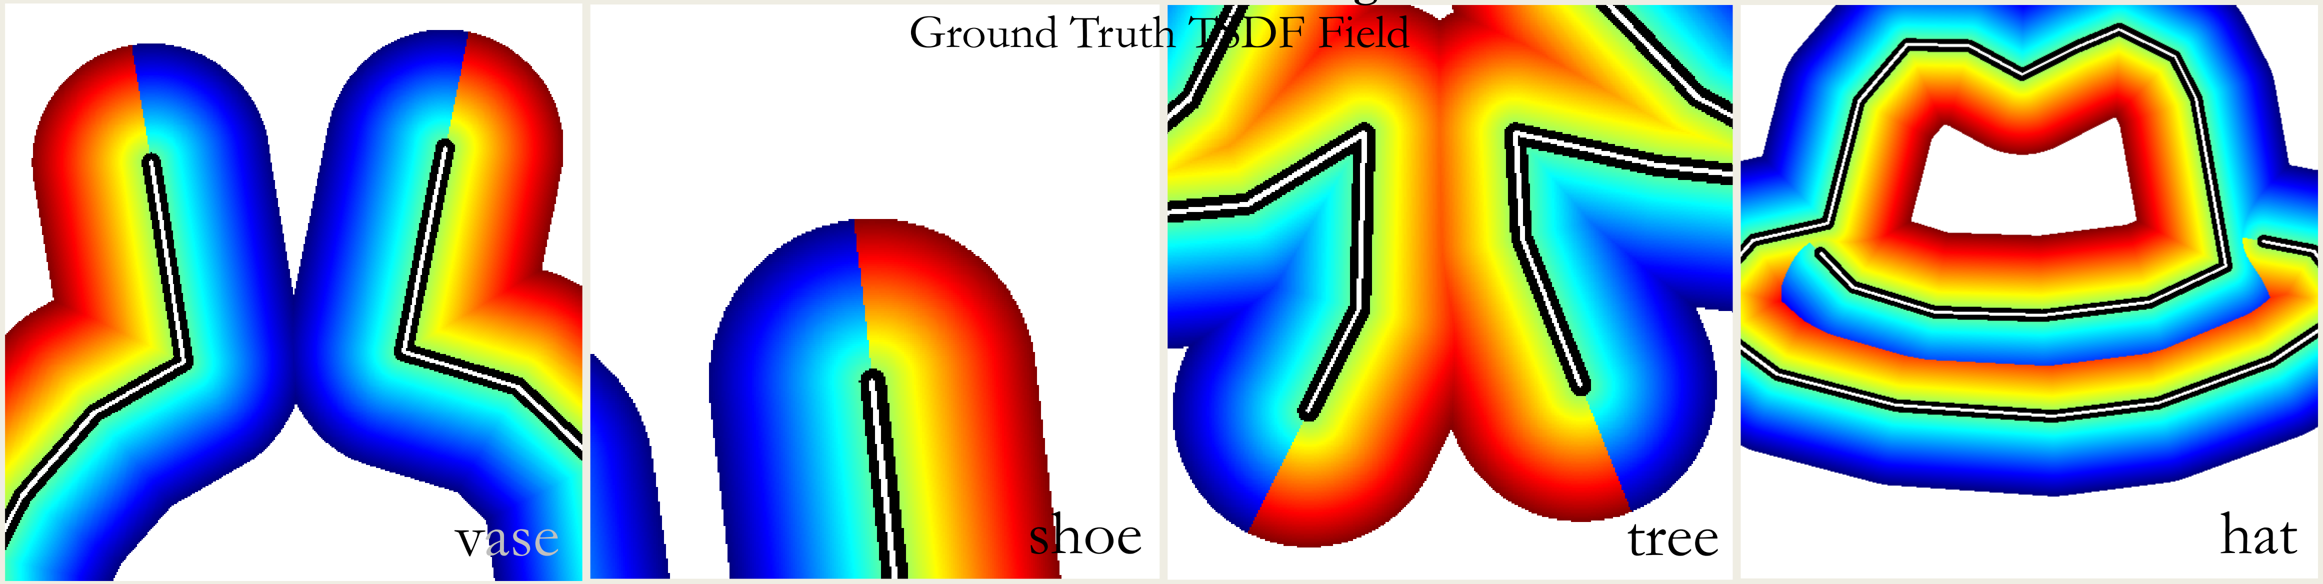

	vase	shoe	tree	hat
Avg. L1 error:	$2.364e-3$	$1.445e-3$	$1.585e-3$	$8.365e-4$
Max L1 error:	$1.309e-1$	$1.127e-1$	$1.196e-1$	$1.340e-1$

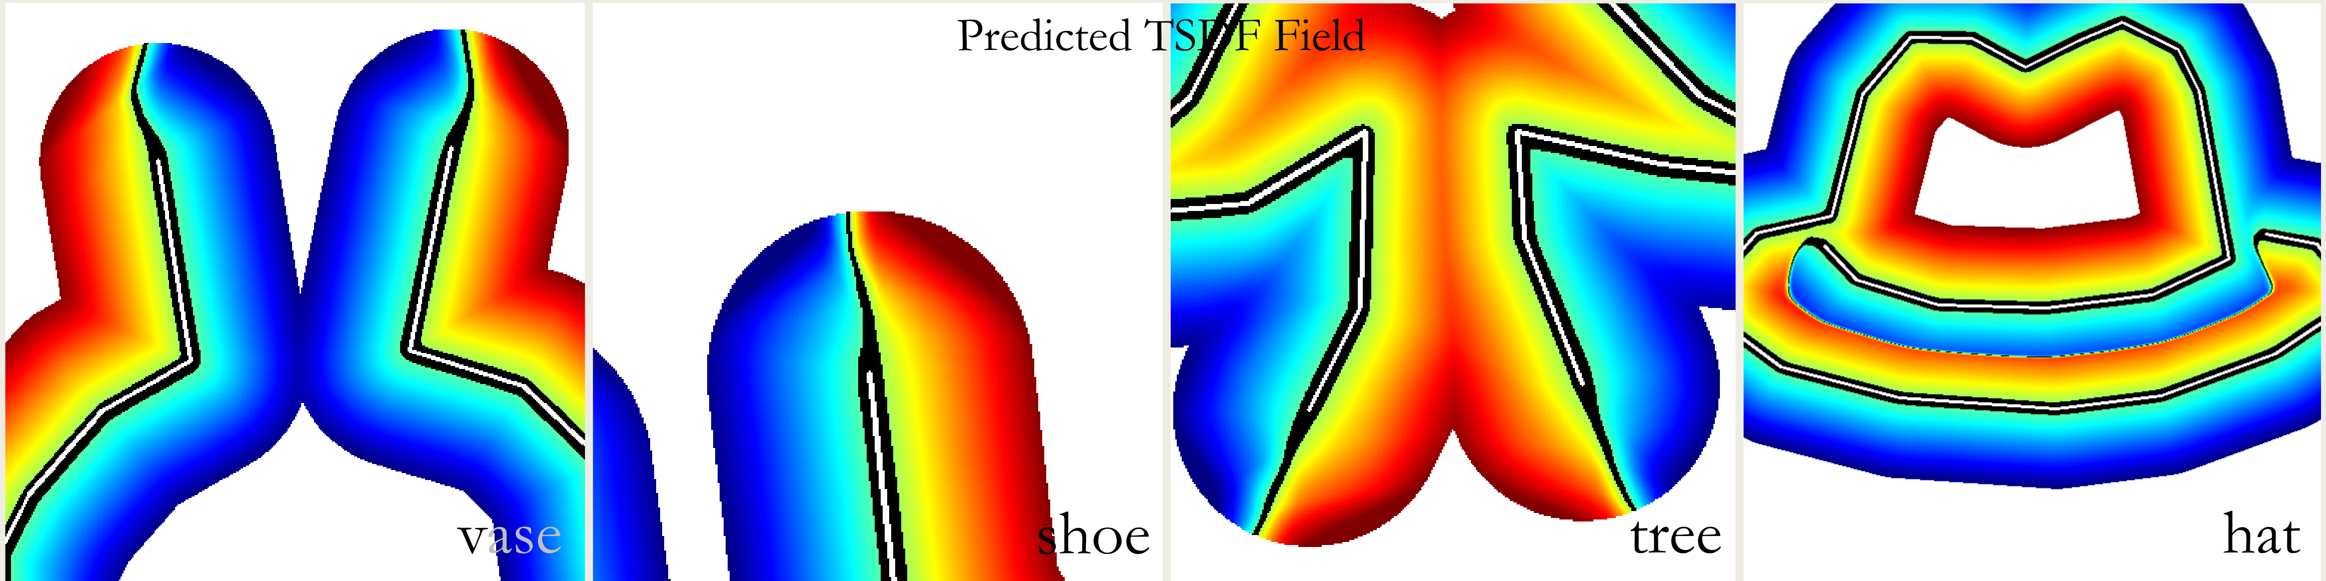

Problems: Outliers around endpoints (Type I)

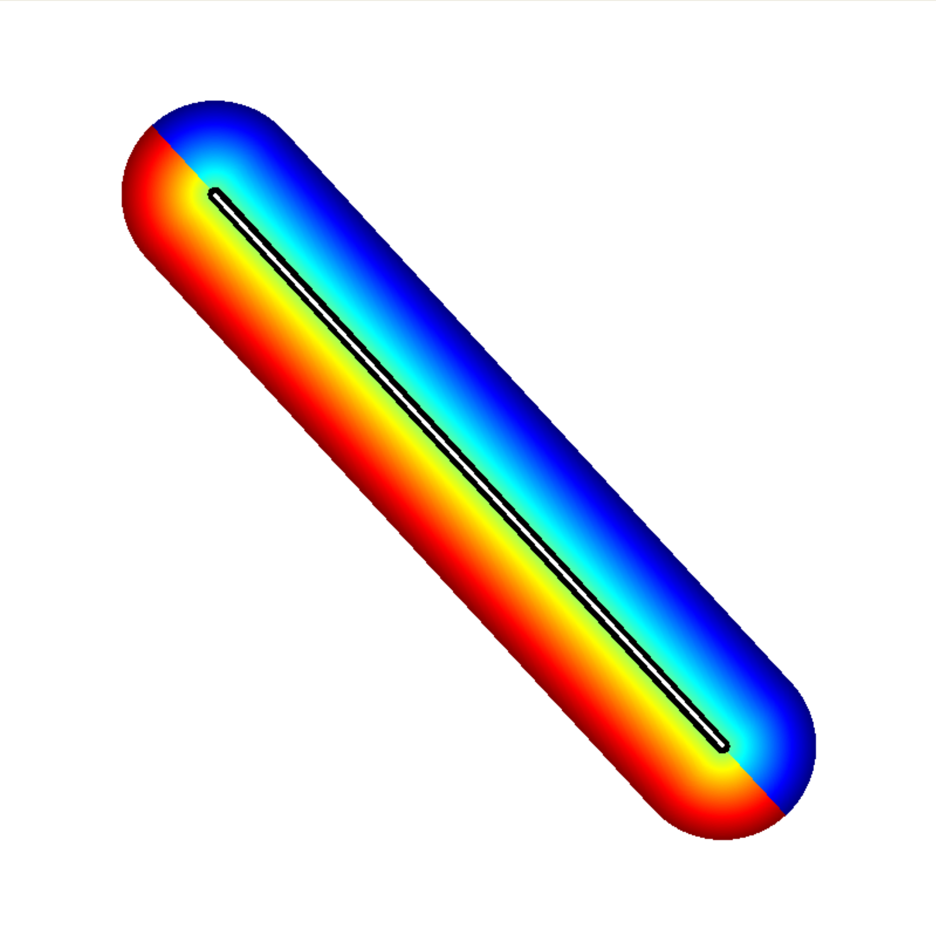

Ground Truth TSDF Field

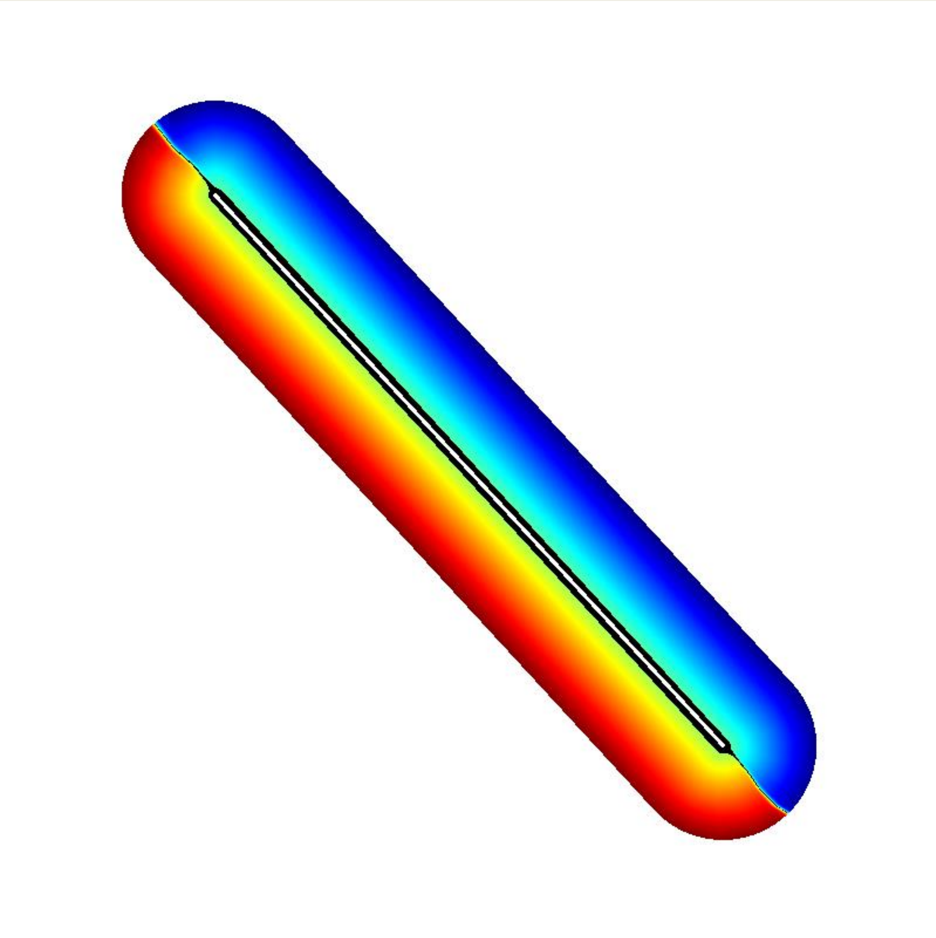

Predicted TSDF Field

Red area: **positive** SDF

Blue area: **negative** SDF

Black line: **zero-value** SDF contour

White line: true curve

Problems: Outliers around endpoints (Type I)

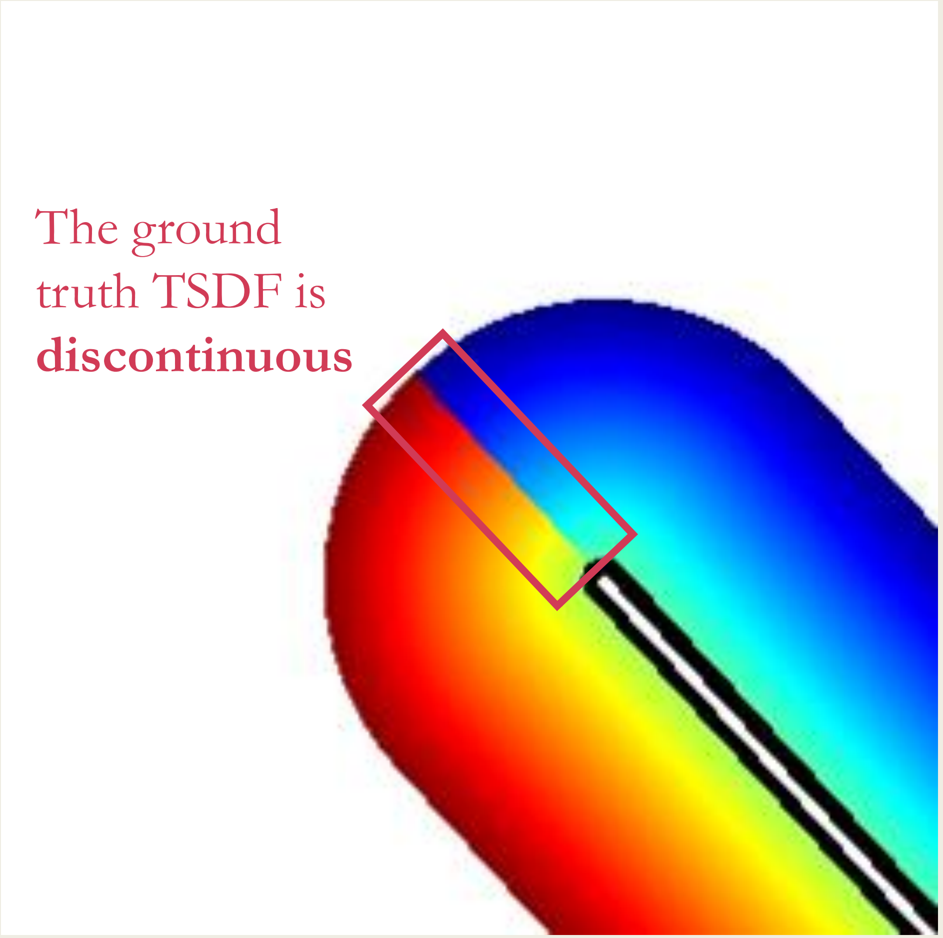

Ground Truth TSDF Field

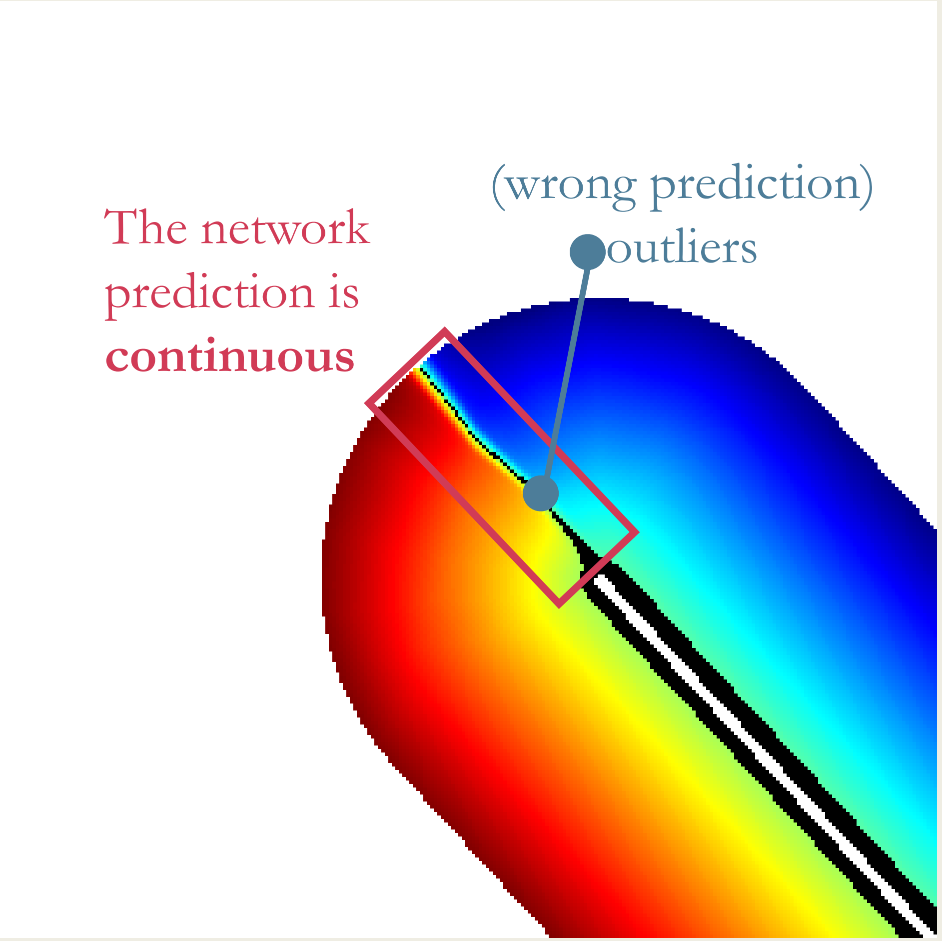

Predicted TSDF Field

Red area: **positive** SDF
Blue area: **negative** SDF
Black line: **zero-value** SDF contour
White line: true curve

Solutions: Cut-off at endpoints

Predicted TSDF before Cut-off

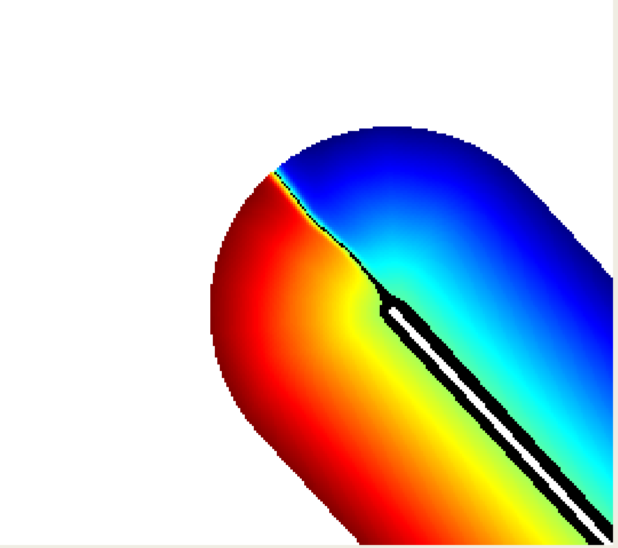

Predicted TSDF after Cut-off

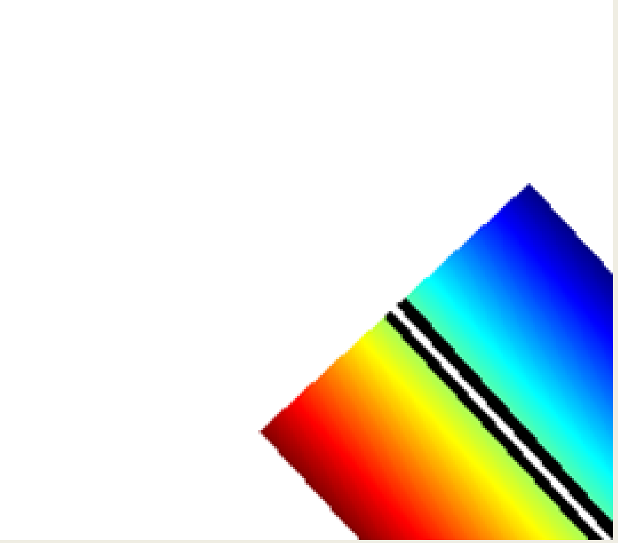

- **Type I outliers**: Outliers near endpoints
  - The **positive-negative conjunction** is caused by an endpoint.
- **Idea**:
  - Cut off the **half circle** at endpoints.
- **Methodology**:
  - Filter out points whose **nearest point on the curve** is one of the endpoints.

Solutions: Cut-off at endpoints

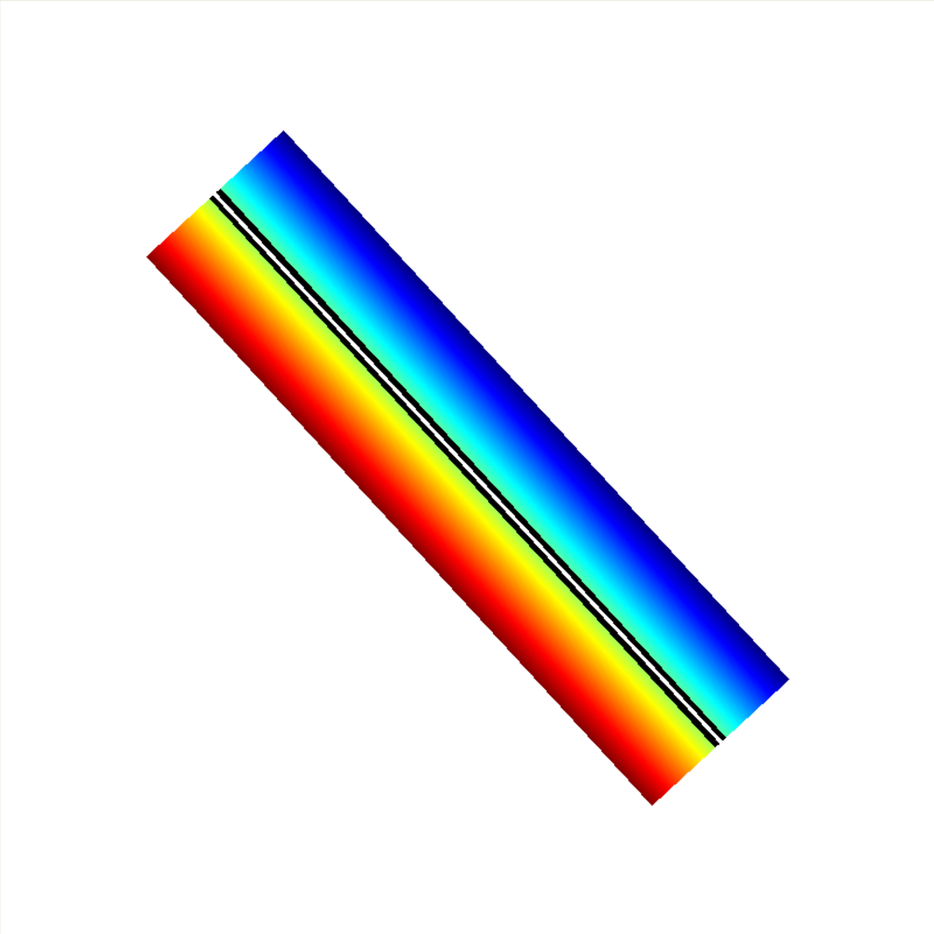

Ground Truth TSDF Field

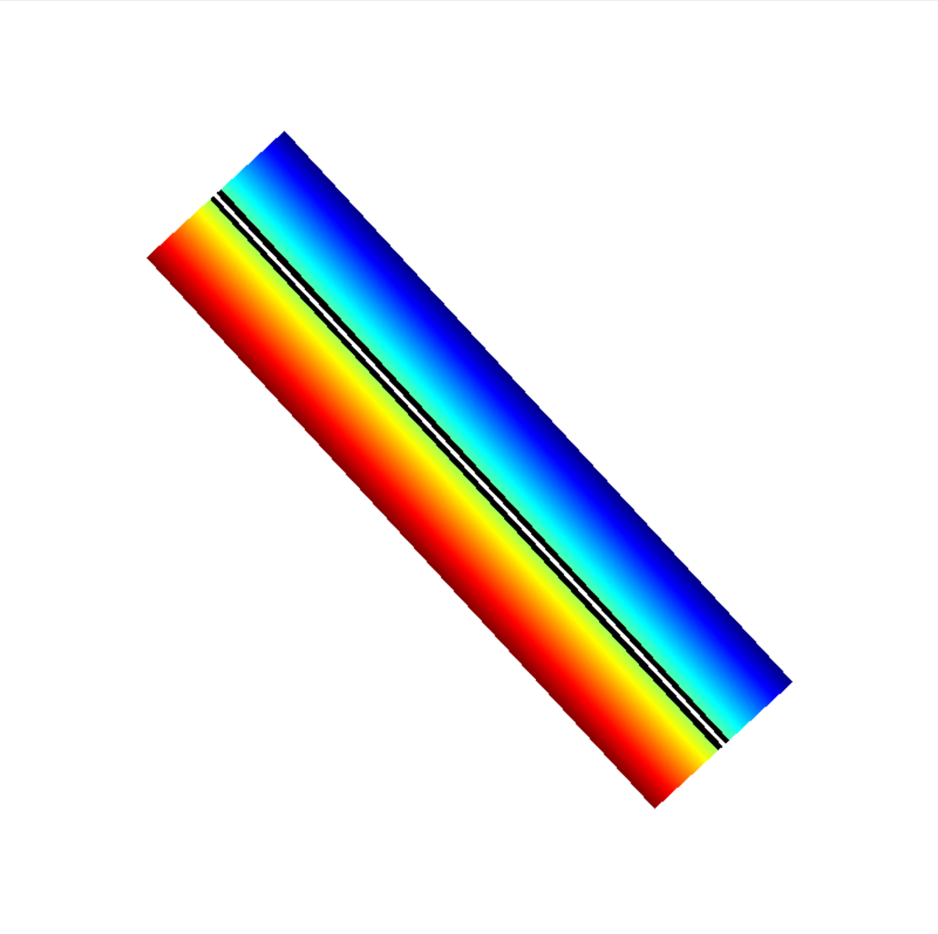

Predicted TSDF Field

Red area: **positive** SDF
Blue area: **negative** SDF
Black line: **zero-value** SDF contour
White line: true curve

Problems: Outliers between two segments (Type II)

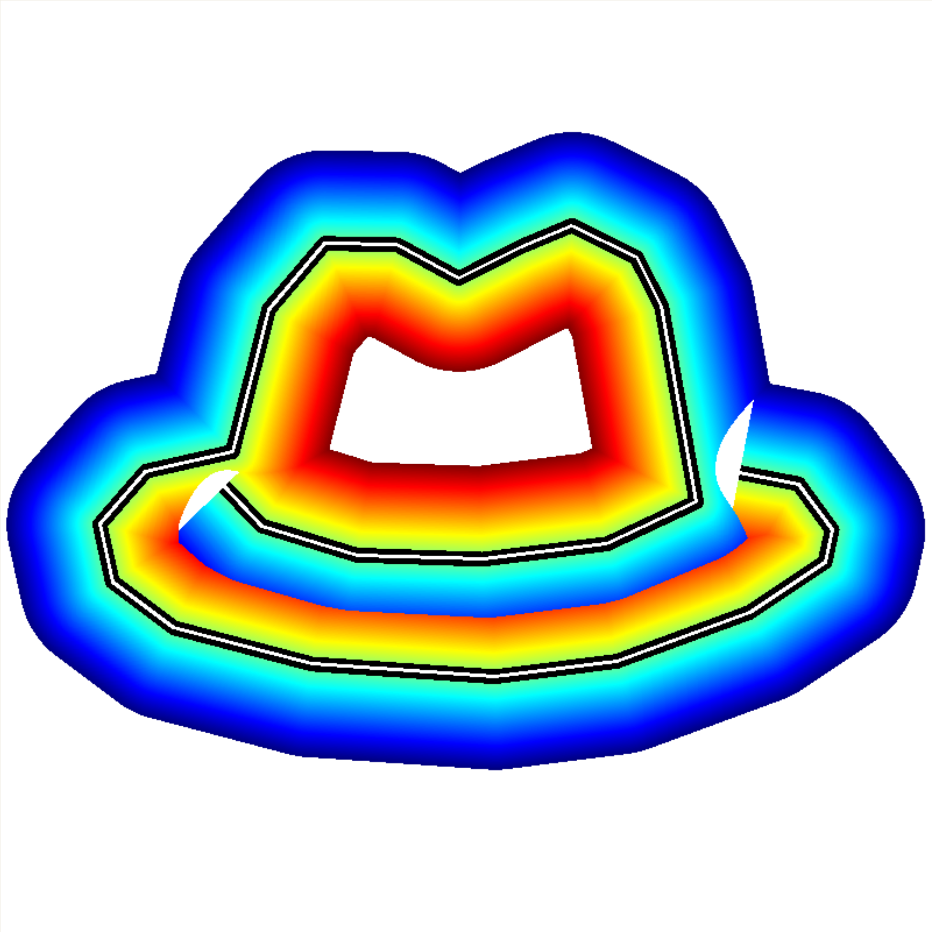

Ground Truth TSDF Field

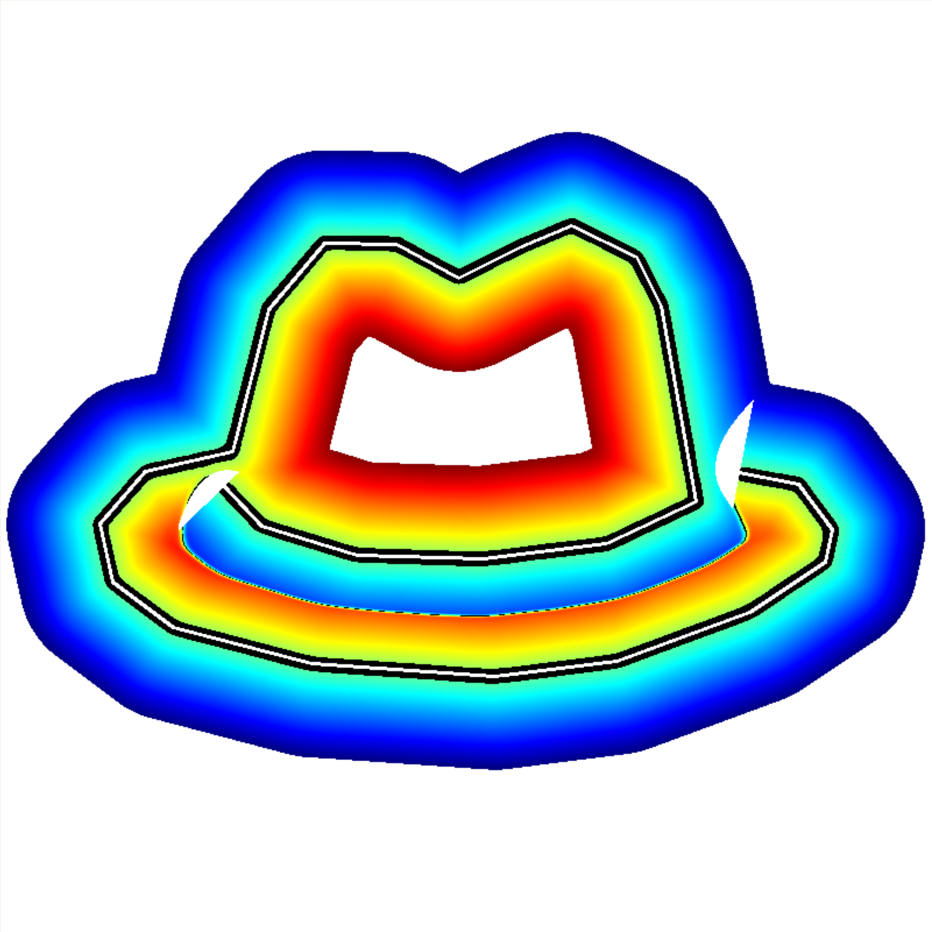

Predicted TSDF Field

Red area: **positive** SDF
Blue area: **negative** SDF
Black line: **zero-value** SDF contour
White line: true curve

Problems: Outliers between two segments (Type II)

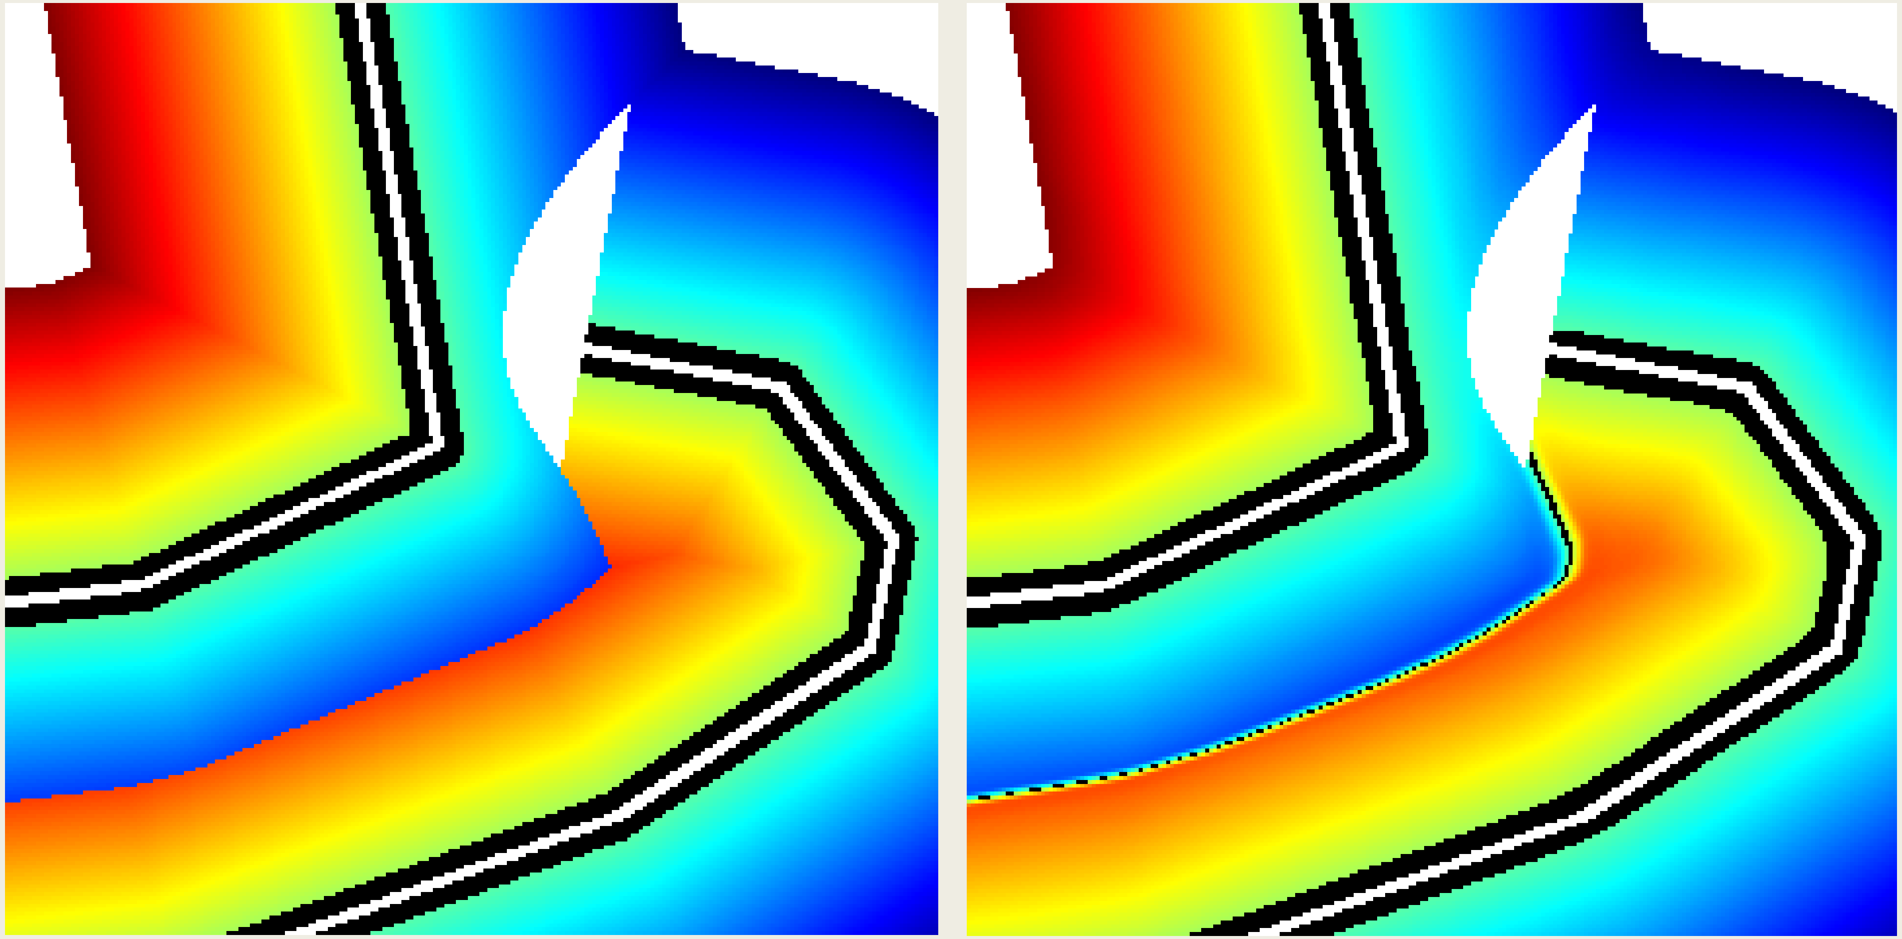

Ground Truth TSDF Field

Predicted TSDF Field

Red area: **positive** SDF
Blue area: **negative** SDF
Black line: **zero-value** SDF contour
White line: true curve

Solutions: Adaptive truncation distance

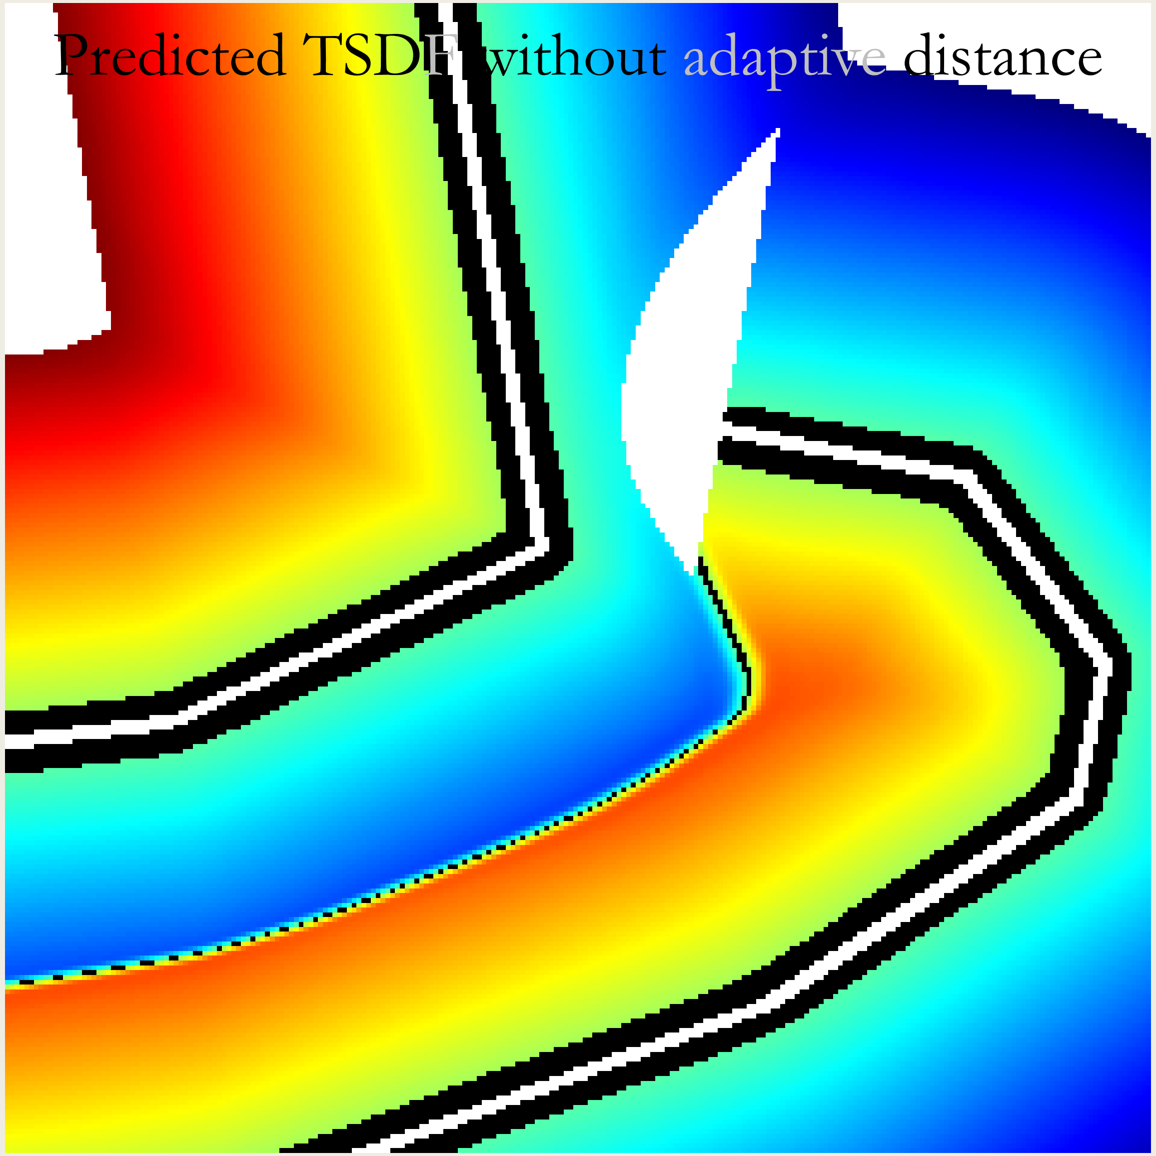

- **Type II outliers**: Outliers between two segments
  - The positive-negative conjunction is caused by **two parts of the curves.**
- **Idea**:
  - Implement an **adaptive truncation distance**.

Solutions: Adaptive truncation distance

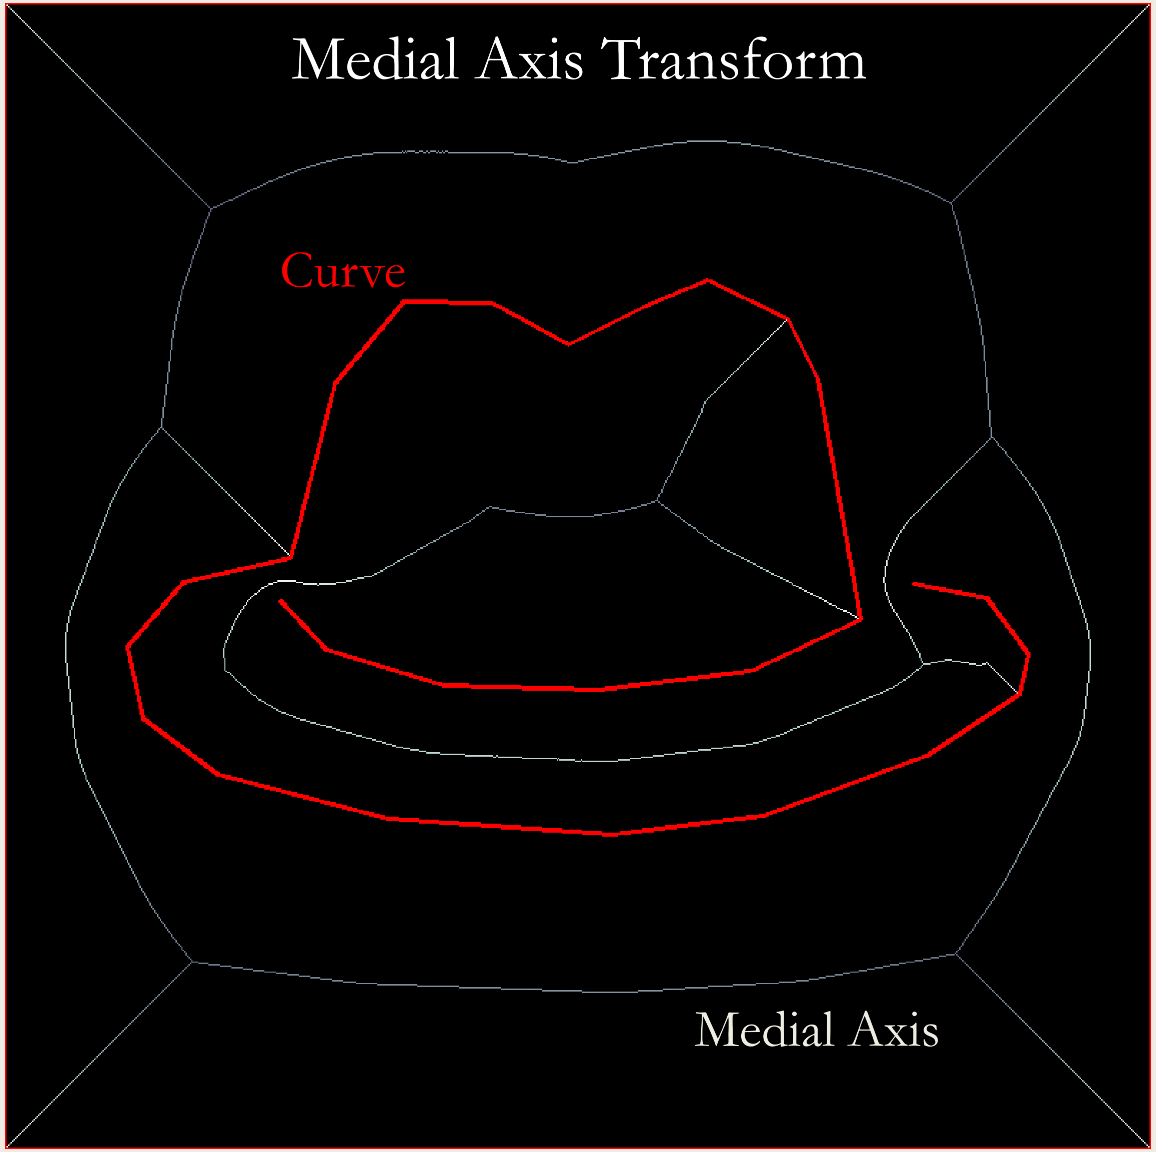

- **Medial Axis Transform**:
  - The **medial axis** is the set of all points **having more than one closest point** on the curve's boundary.
- **Methodology**:
  - For every point on the **medial axis**, if the point is near the **positive-negative junction**, we filter out **nearby points**.

Solutions: Adaptive truncation distance

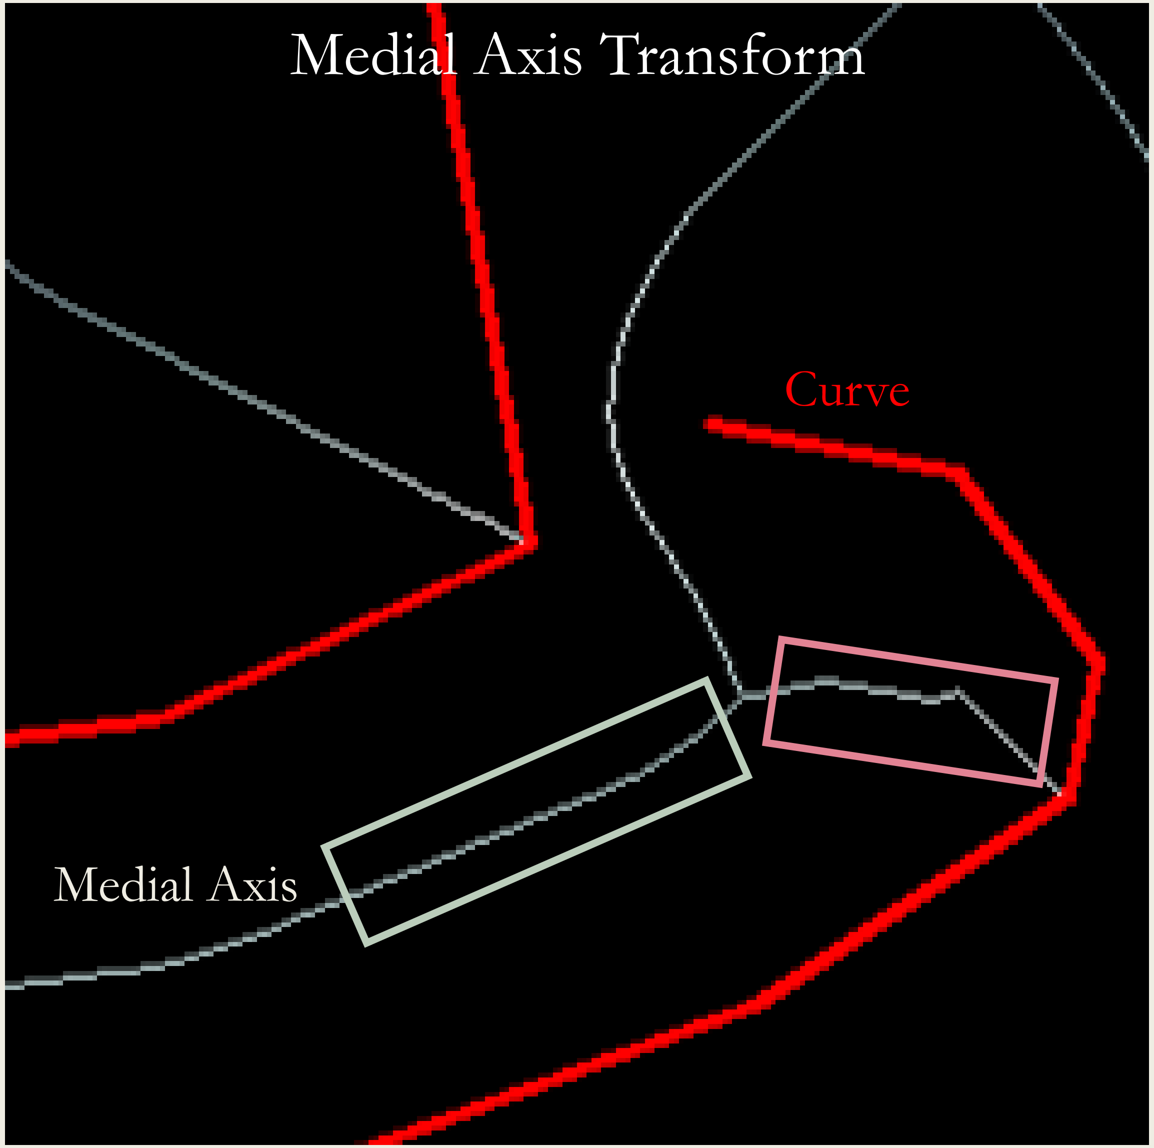

- **Medial Axis Transform**:
  - The **medial axis** is the set of all points **having more than one closest point** on the curve's boundary.
- **Methodology**:
  - For every point on the **medial axis**, if the point is near the **positive-negative junction**, we filter out **nearby points**.

Solutions: Adaptive truncation distance

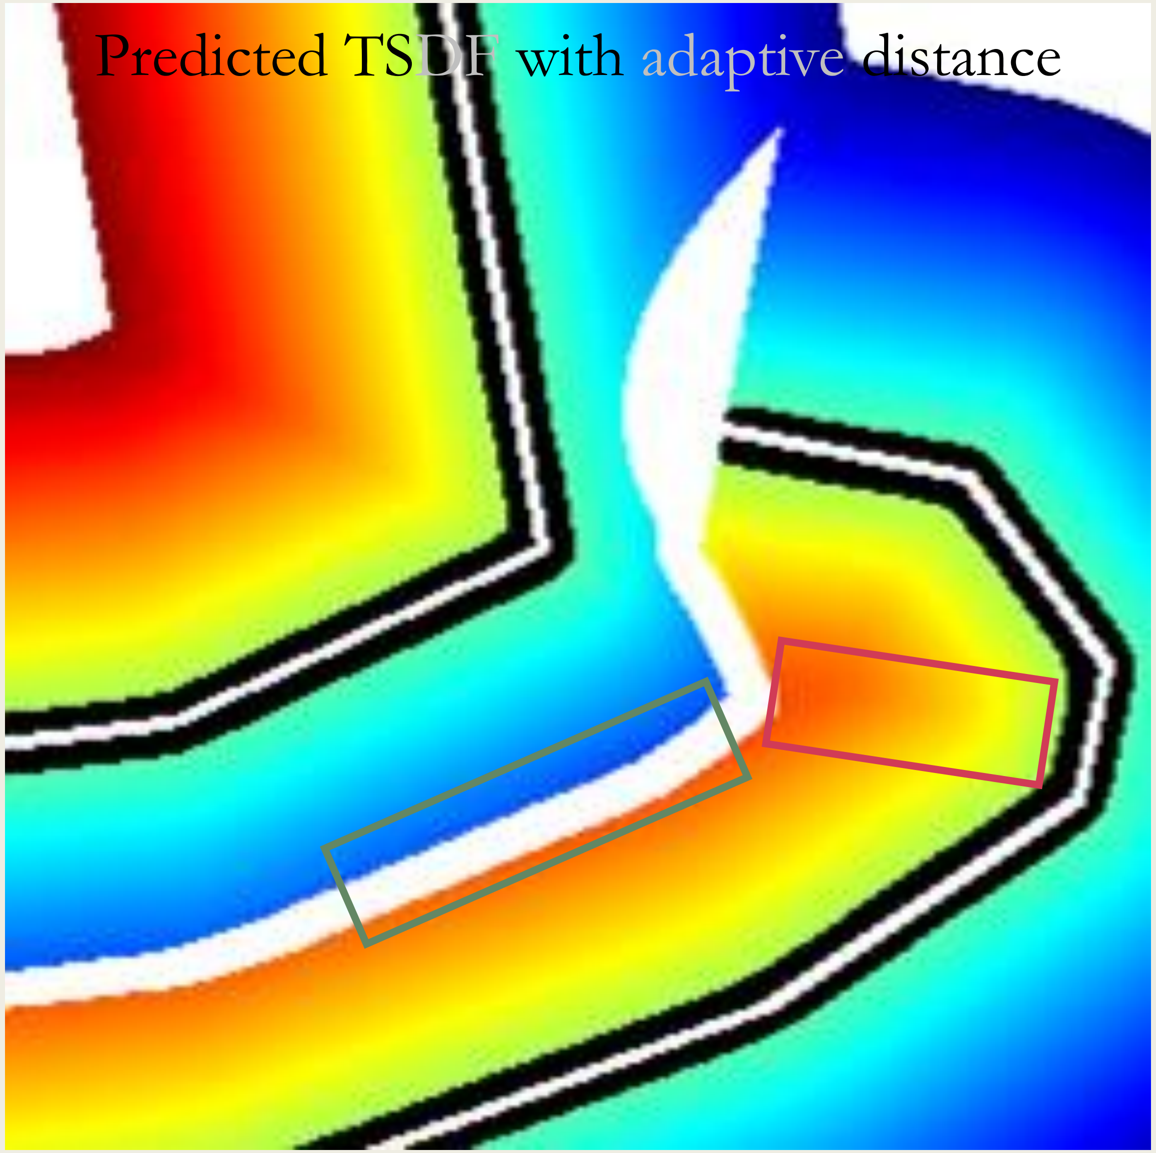

- **Medial Axis Transform:**
  - The **medial axis** is the set of all points **having more than one closest point** on the curve's boundary.
- **Methodology:**
  - For every point on the **medial axis**, if the point is near the **positive-negative junction**, we filter out **nearby points.**

Solutions: Adaptive truncation distance

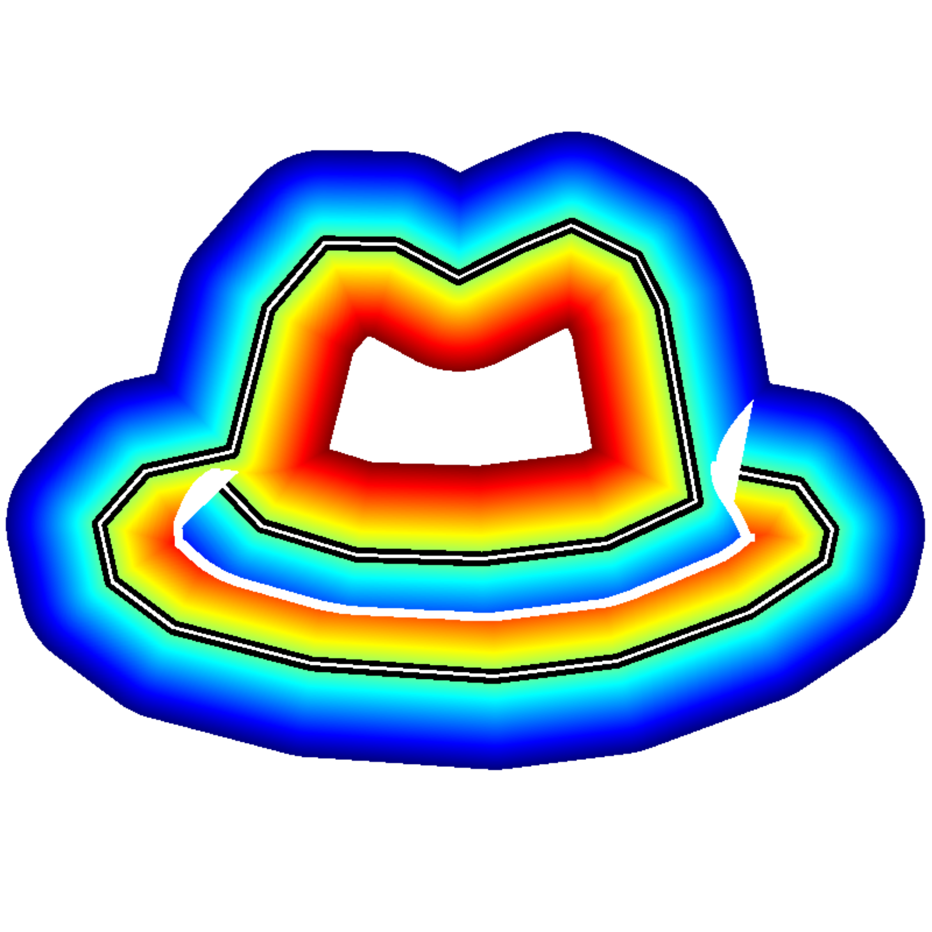

Ground Truth TSDF Field

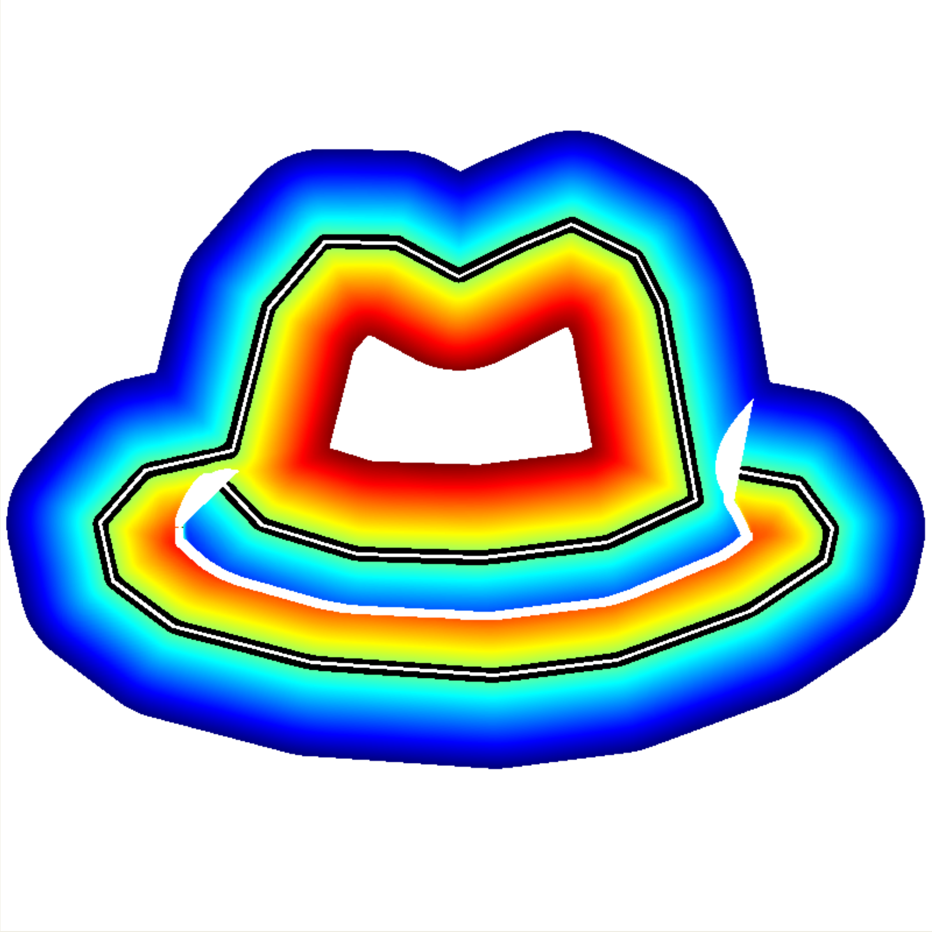

Predicted TSDF Field

Red area: **positive** SDF
Blue area: **negative** SDF
Black line: **zero-value** SDF contour
White line: true curve

Results

Results: No outliers

Red area: **positive** SDF
Blue area: **negative** SDF
Black line: **zero-value** SDF contour
White line: true curve

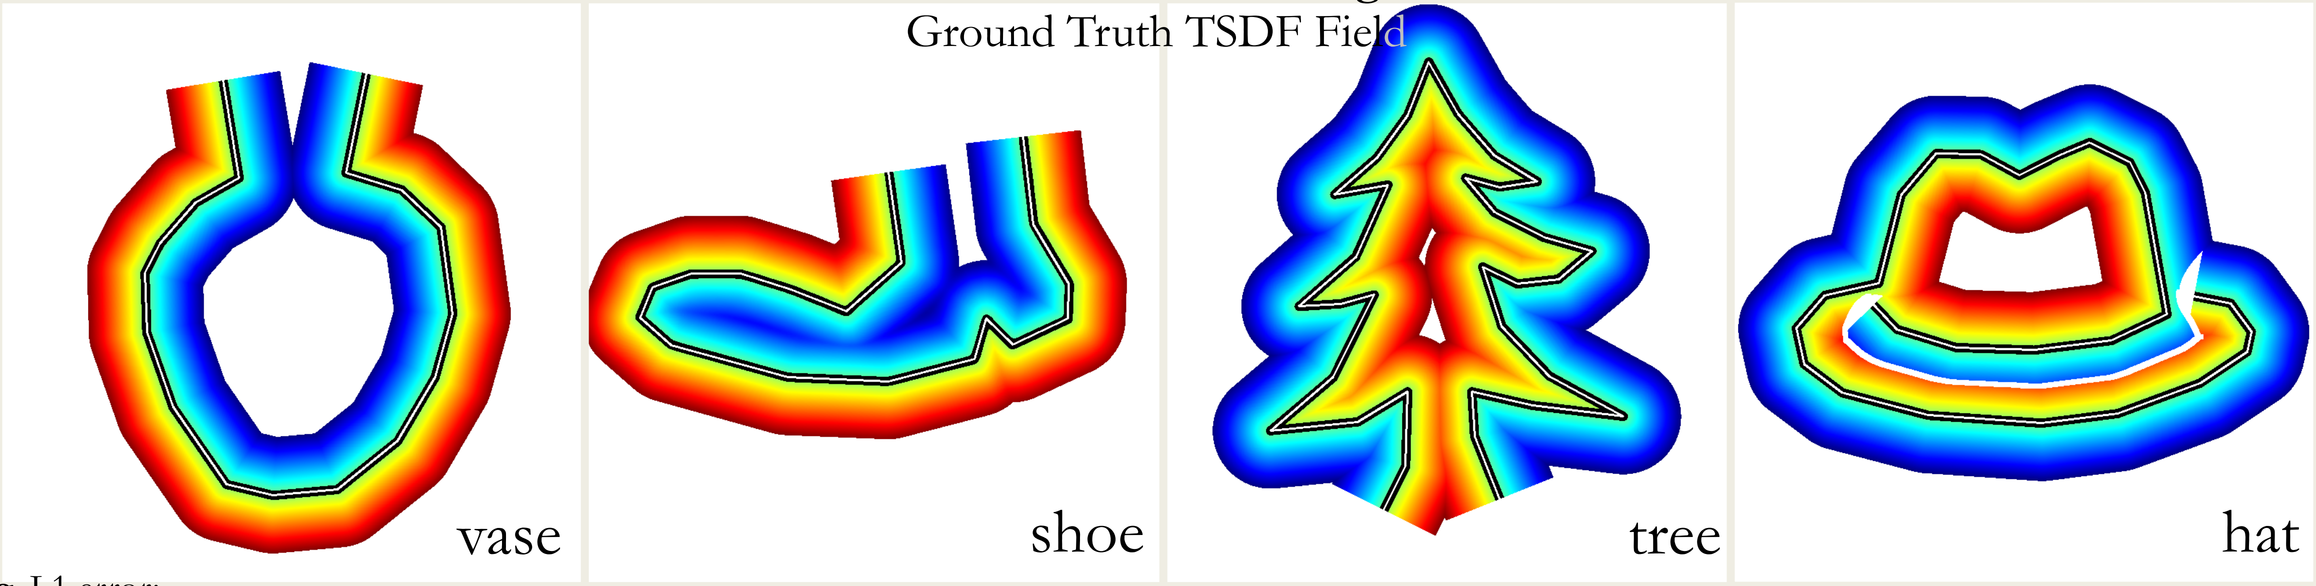

	vase	shoe	tree	hat
Avg. L1 error:	$2.364e-3 \rightarrow 1.666e-3$	$1.445e-3 \rightarrow 5.893e-4$	$1.585e-3 \rightarrow 1.041e-3$	$8.365e-4 \rightarrow 4.520e-4$
Max L1 error:	$1.309e-1 \rightarrow 8.080e-3$	$1.127e-1 \rightarrow 5.428e-2$	$1.196e-1 \rightarrow 1.195e-2$	$1.340e-1 \rightarrow 7.647e-2$

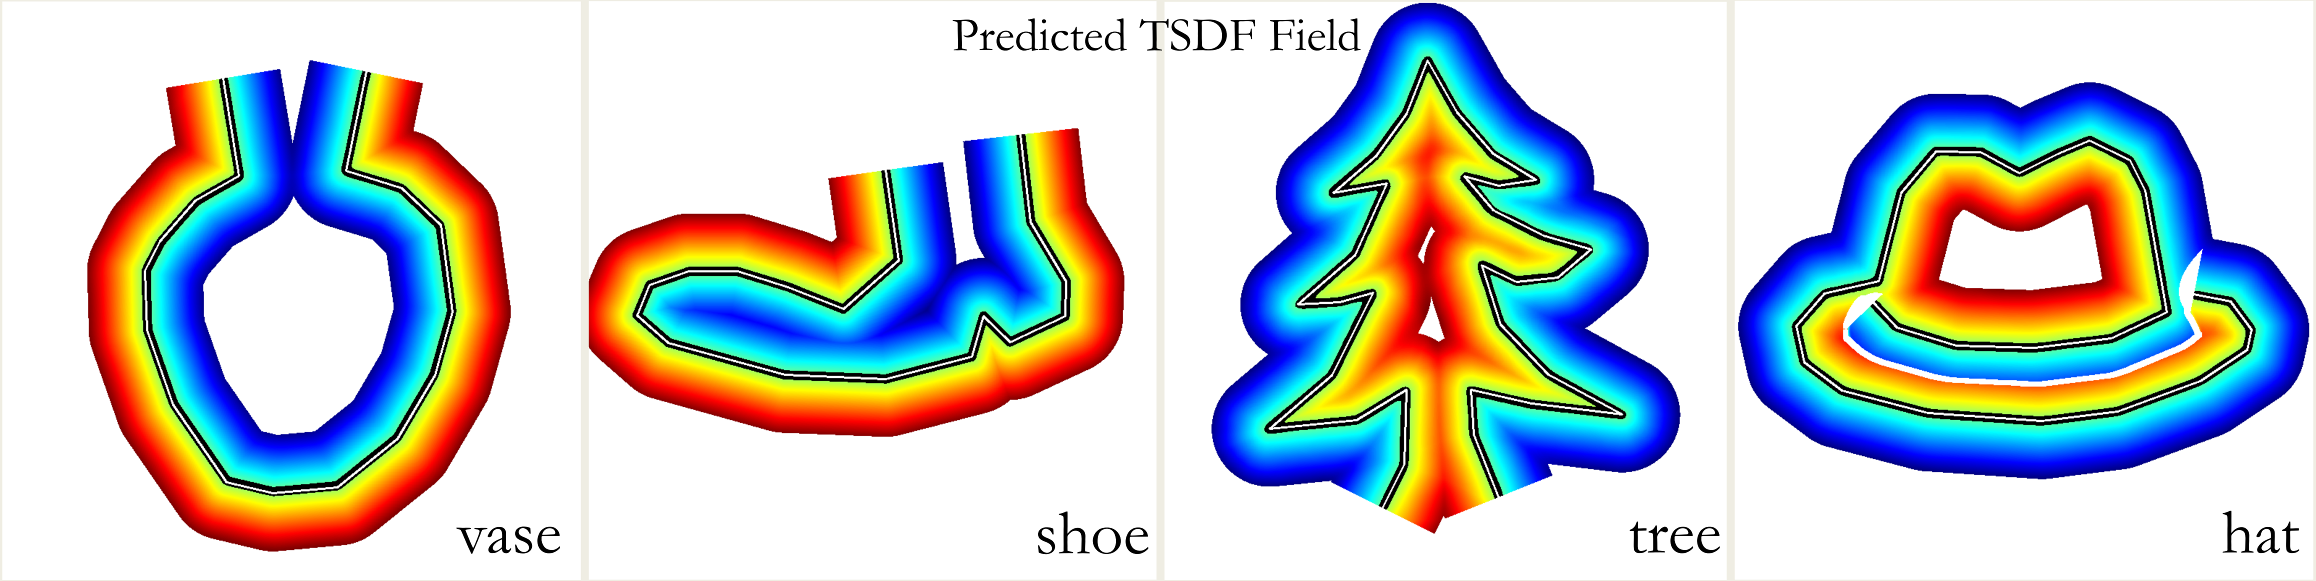

Results: No outliers

Red area: **positive** SDF
Blue area: **negative** SDF
Black line: **zero-value** SDF contour
White line: true curve

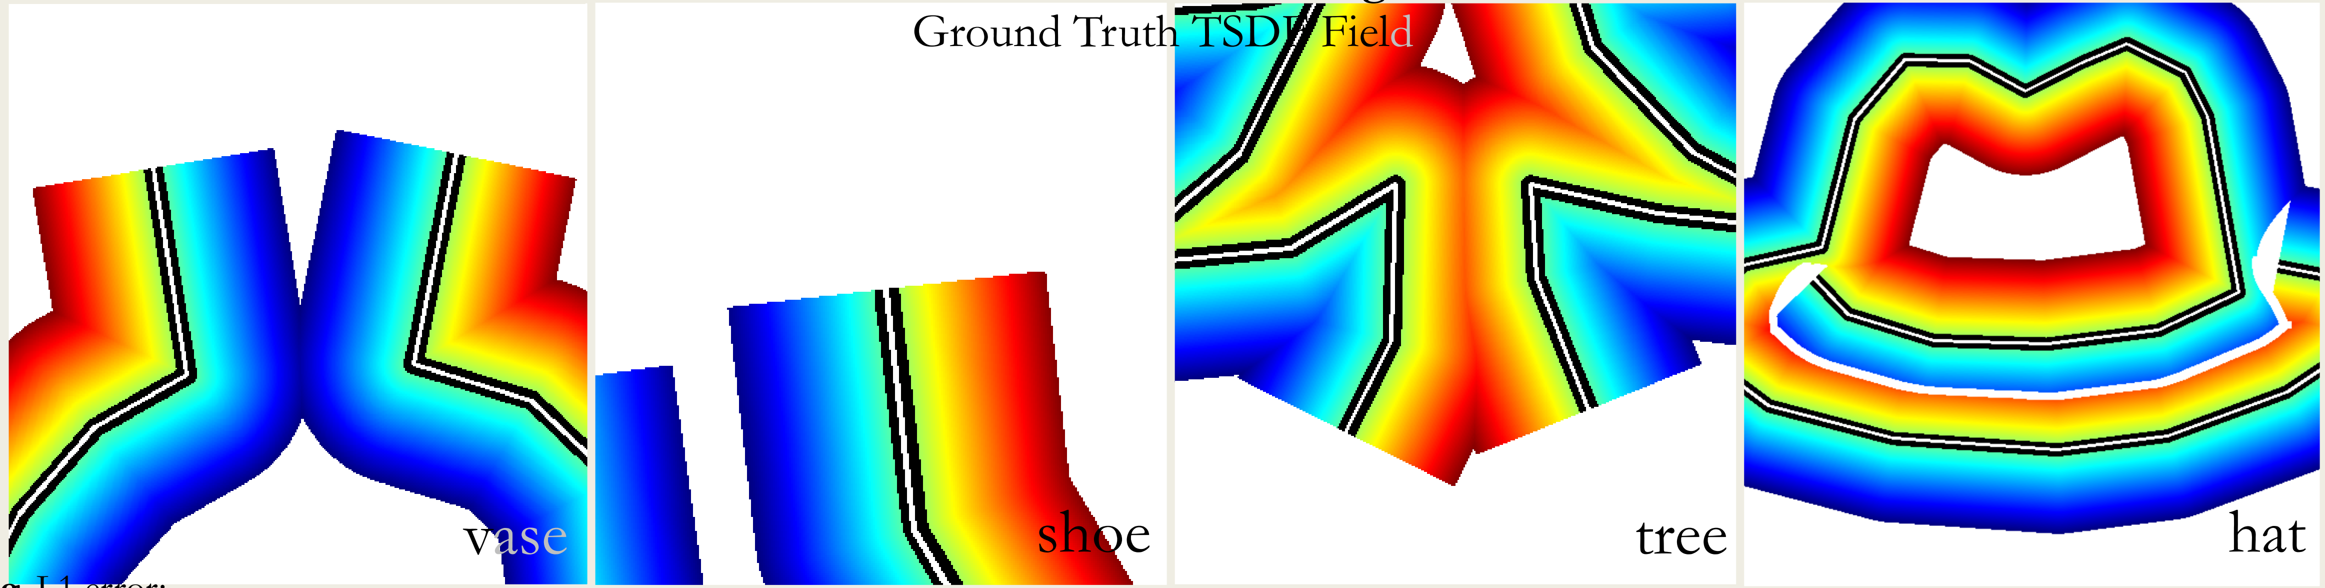

Avg. L1 error:

vase	shoe	tree	hat
$2.364e-3 \rightarrow 1.666e-3$	$1.445e-3 \rightarrow 5.893e-4$	$1.585e-3 \rightarrow 1.041e-3$	$8.365e-4 \rightarrow 4.520e-4$

Max L1 error:

vase	shoe	tree	hat
$1.309e-1 \rightarrow 8.080e-3$	$1.127e-1 \rightarrow 5.428e-2$	$1.196e-1 \rightarrow 1.195e-2$	$1.340e-1 \rightarrow 7.647e-2$

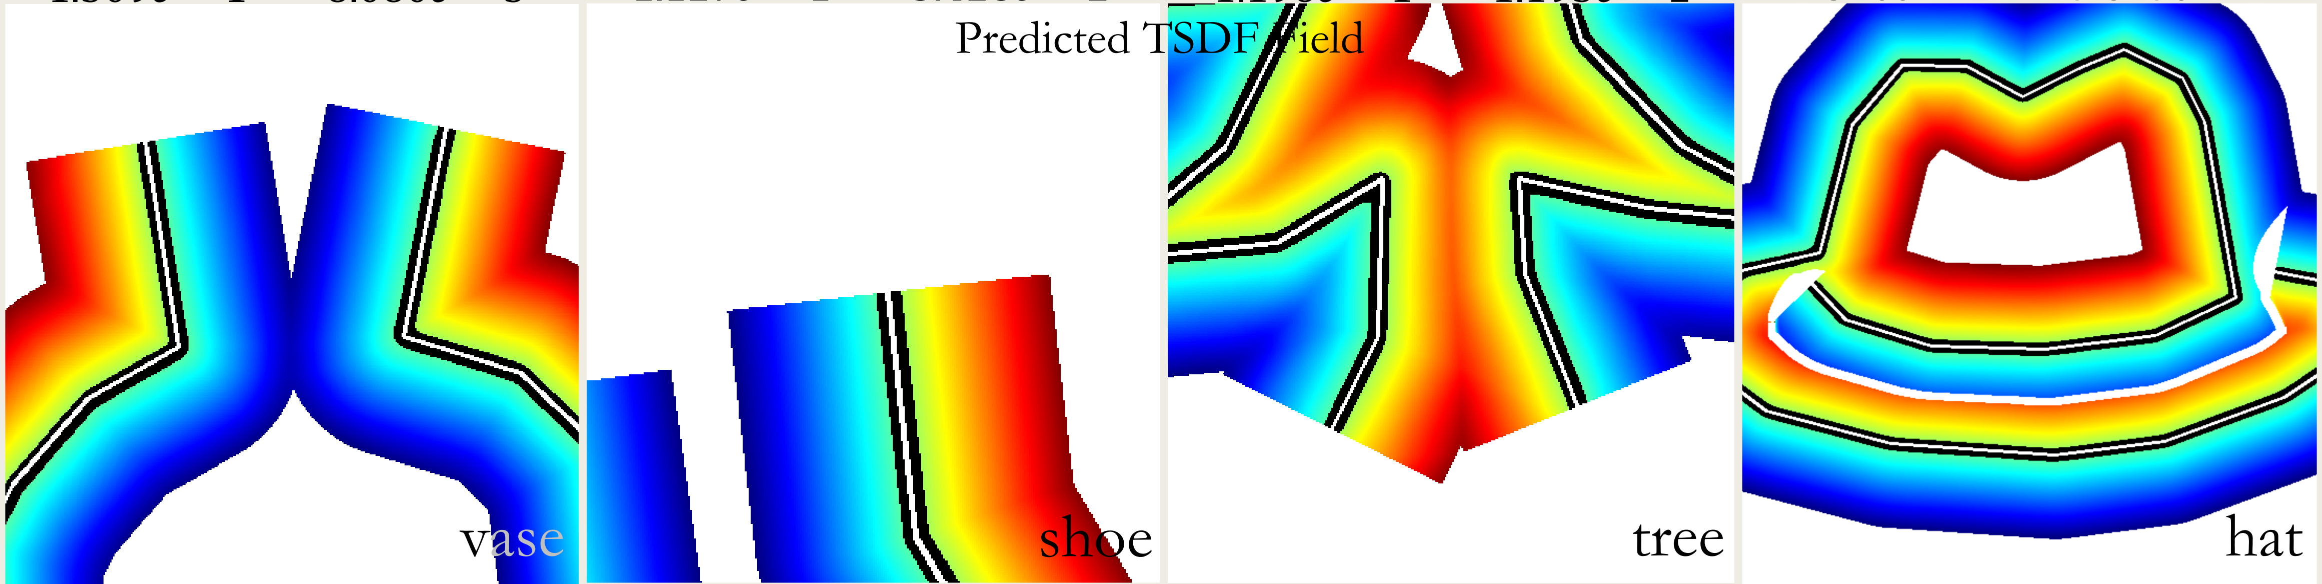

Results: Comparison

Red area: **positive** SDF
Blue area: **negative** SDF
Black line: **zero-value** SDF contour
White line: true curve

Chibane, J., Mir, A., & Pons-Moll, G. (2020). Neural unsigned distance fields for implicit function learning.

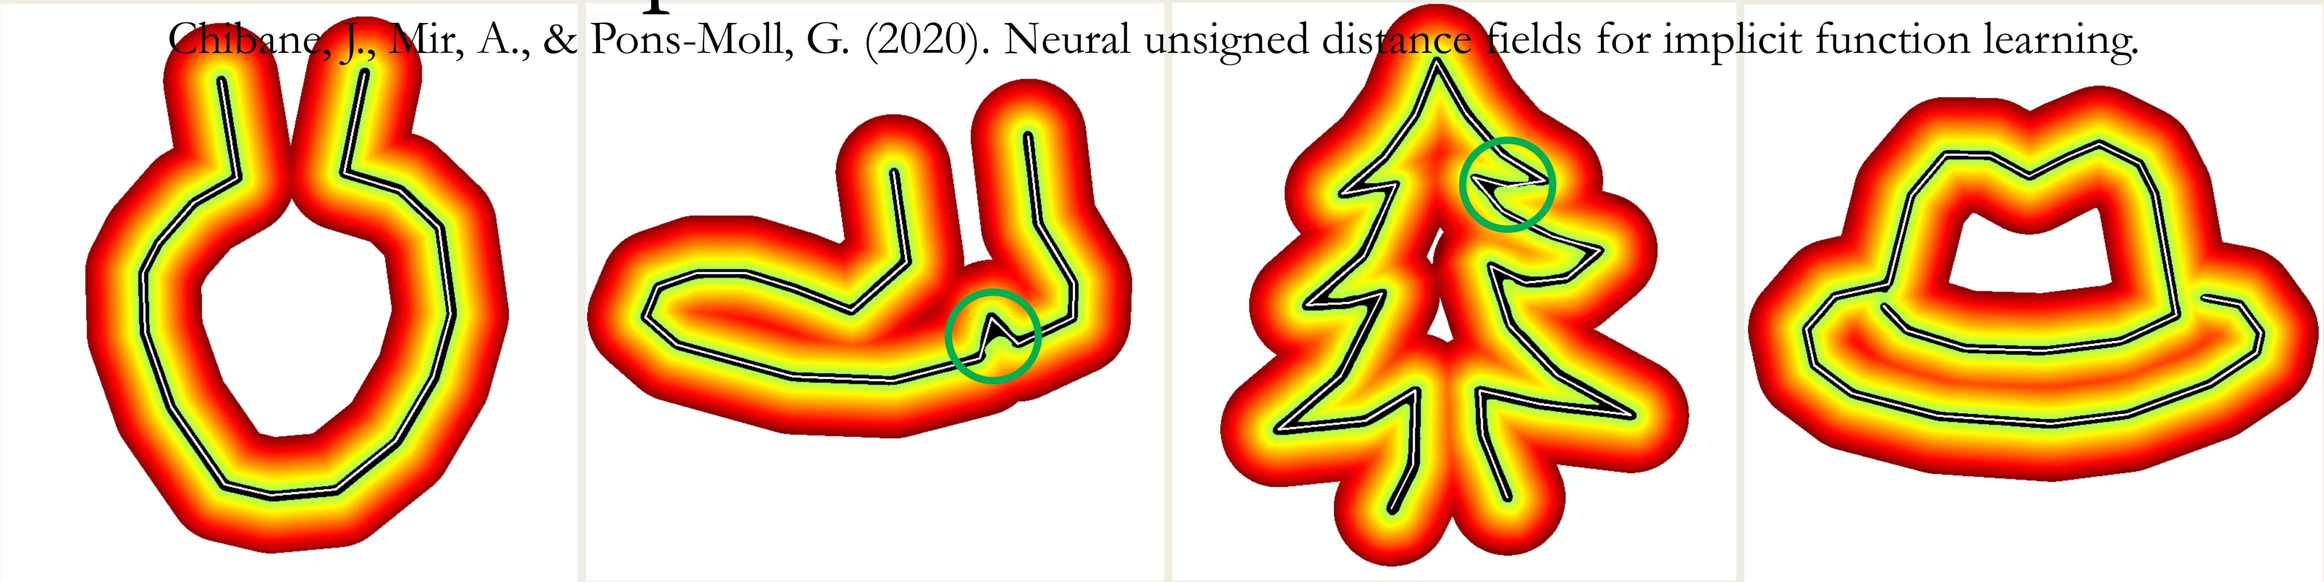

Avg. L1 error:	$1.465e-3$	$7.052e-4$	$1.416e-3$	$3.639e-4$
	$1.666e-3$	$5.893e-4$	$1.041e-3$	$4.520e-4$

My Results

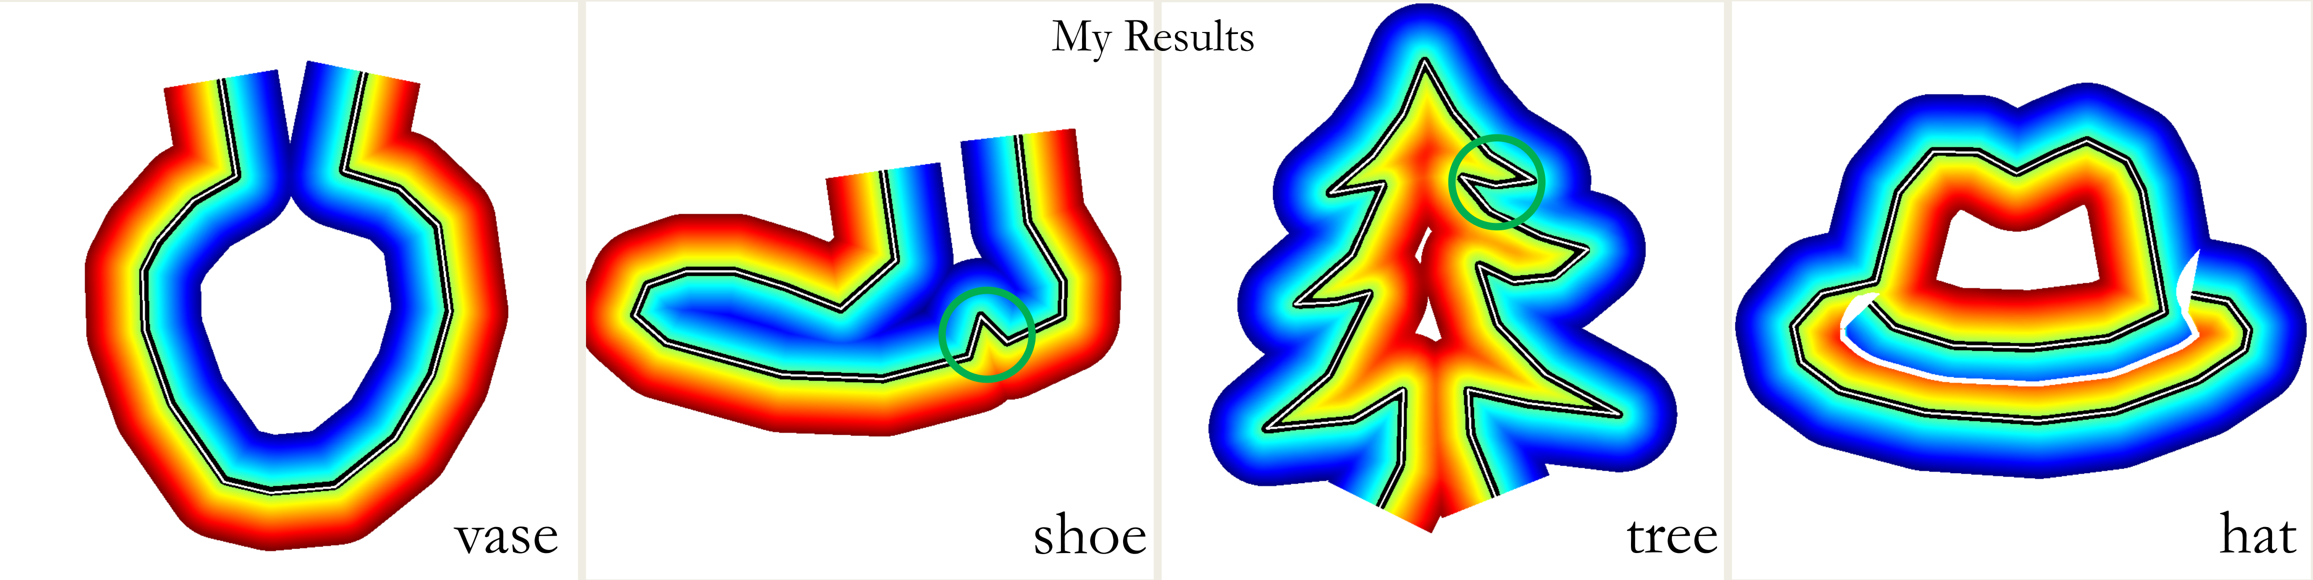

Future work

Future work: Applications & Extensions

- **Applications:**
  - Collision detection using **neural representations**.
- **Extensions:**
  - Extension to **3D open surfaces**.
  - Endpoints of open curves → **Boundary** of open surfaces
  - Medial axis transform in 2D → **Medial axis transform** in 3D

The end

Bibliography

- Chibane, J., Mir, A., & Pons-Moll, G. (2020). Neural unsigned distance fields for implicit function learning. arXiv preprint arXiv:2010.13938.
- Park, J. J., Florence, P., Straub, J., Newcombe, R., & Lovegrove, S. (2019). Deepsdf: Learning continuous signed distance functions for shape representation. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (pp. 165-174).