Neural skeleton: implicit neural representation away from the surface

Shape Modeling International 2023

Mattéo Clémot, Julie Digne

LIRIS Lyon 1

July 14, 2023

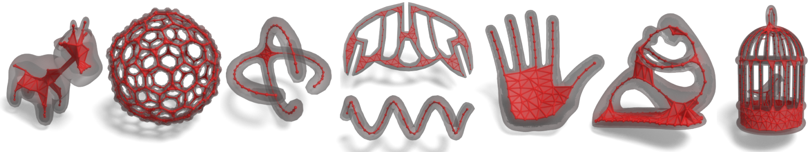

Skeletonization

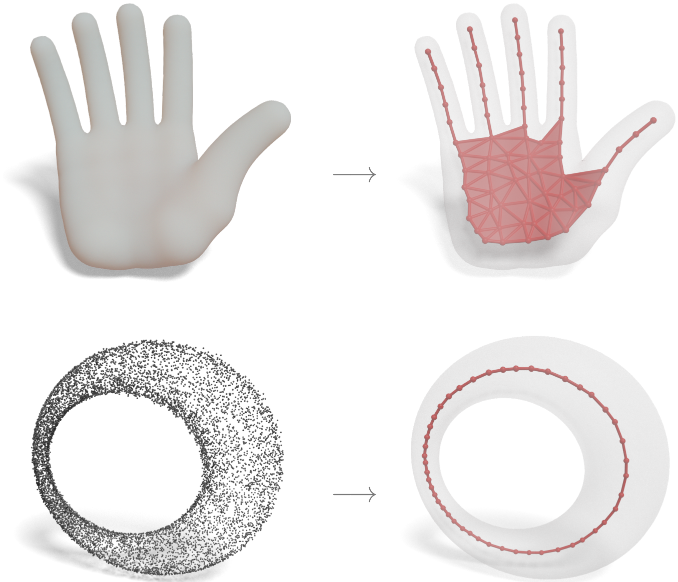

Applications

Shape segmentation, shape matching...

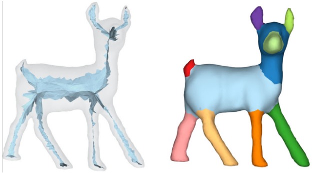

[Lin 2020]

A long standing problem

- Mean Curvature Skeleton [Tagliasacchi 2012]

[Tagliasacchi 2016]

A long standing problem

- Mean Curvature Skeleton [Tagliasacchi 2012]
- L_1-Medial Skeleton [Huang 2013]

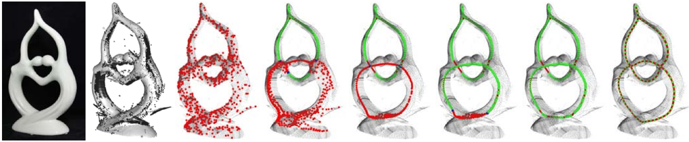

[Huang 2013]

A long standing problem

- Mean Curvature Skeleton [Tagliasacchi 2012]
- L_1-Medial Skeleton [Huang 2013]
- Voxel Cores [Yan 2018]

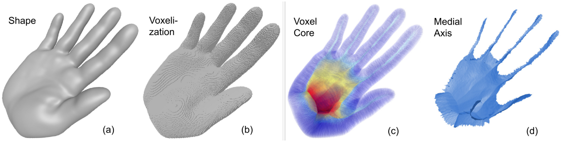

[Yan 2018]

A long standing problem

- Mean Curvature Skeleton [Tagliasacchi 2012]
- L_1-Medial Skeleton [Huang 2013]
- Voxel Cores [Yan 2018]
- Coverage Axis [Dou 2022]

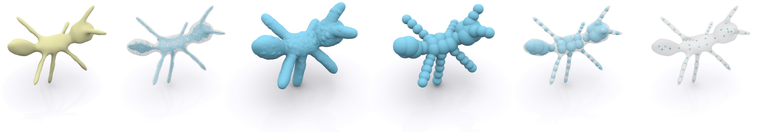

[Dou 2022]

A long standing problem

- Mean Curvature Skeleton [Tagliasacchi 2012]
- L_1-Medial Skeleton [Huang 2013]
- Voxel Cores [Yan 2018]
- Coverage Axis [Dou 2022]
- Point2Skeleton [Lin 2021] (needs a database; beyond our scope)

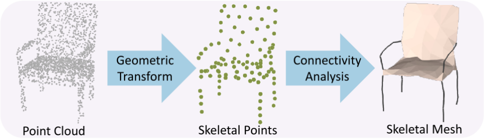

[Lin 2021]

Medial axis

Medial axis $\text{med}(\Omega)$: points $\mathbf{x}$ of $\mathbb{R}^d$ such that $d(\mathbf{x}, \partial\Omega)$ is reached at least two times.

Lemma

The **signed-distance function** $u_\Omega(\mathbf{x}) = \begin{cases} -d(\mathbf{x}, \partial\Omega) & \text{if } \mathbf{x} \in \Omega \\ d(\mathbf{x}, \partial\Omega) & \text{if } \mathbf{x} \in \overline{\Omega}^C \end{cases}$ is differentiable almost everywhere, verifies the eikonal equation $\|\nabla u_\Omega\| = 1$ where this is the case, and the medial axis $\text{med}(\Omega)$ is exactly the points of non-differentiability.

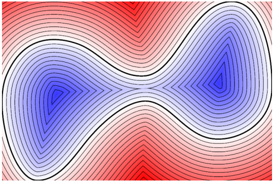

SDF u_Ω

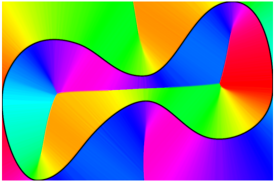

Direction of ∇u_Ω

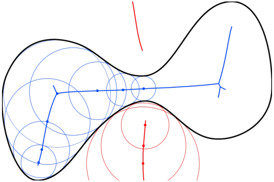

Medial axis $\text{med}(\Omega)$

Medial axis properties

- Ω and $\text{med}(\Omega)$ have the same homotopy type [Lieutier2004]

Medial axis properties

- Ω and $\text{med}(\Omega)$ have the same homotopy type [Lieutier2004]
- $\text{med}(\Omega)$ is unstable

About implicit neural representations

INR

Train a neural network to encode a shape into its parameters.

$$u_\theta : \begin{array}{rcl} \mathbb{R}^3 & \rightarrow & \mathbb{R} \\ (x, y, z) & \mapsto & \mathrm{SDF}_\Omega(x, y, z) \end{array}$$

- DeepSDF [Park 2019], Occupancy Networks [Mescheder 2019]...
- Optimization per shape / on a database
- Focus on surface reconstruction and visualization

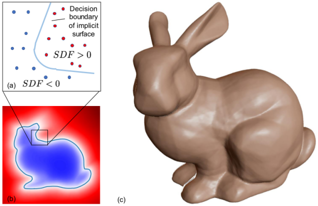

[Park 2019]

Can we use INRs to extract a skeleton?
$\rightarrow$ Leverage neural priors to get robustness.

Overview

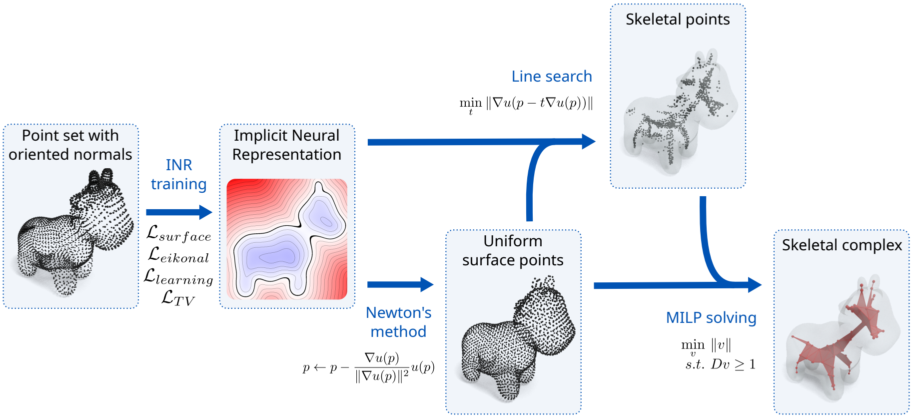

Point set with oriented normals
INR training
$\mathcal{L}_{surface}$
$\mathcal{L}_{eikonal}$
$\mathcal{L}_{learning}$
$\mathcal{L}_{TV}$
Implicit Neural Representation
Line search
$\min_t \|\nabla u(p - t\nabla u(p))\|$
Skeletal points
Newton's method
$p \leftarrow p - \frac{\nabla u(p)}{\|\nabla u(p)\|^2} u(p)$
Uniform surface points
MILP solving
$\min_v \|v\|$ $s.t.\ Dv \geq 1$
Skeletal complex

INR general principle

- Input: point cloud with normals $(\mathbf{x}_i, \mathbf{n}_i)$
- Look for u such that:

$$\begin{cases} \|\nabla u\| &= 1 \\ u_{|\partial\Omega} &= 0 \\ \nabla u_{|\partial\Omega} &= \mathbf{n} \end{cases}$$

- Loss function from [Gropp 2020]:

$$\ell(\theta) = \frac{1}{|I|} \sum_{i \in I} \left(|u_\theta(\mathbf{x}_i)| + \tau \|\nabla u_\theta(\mathbf{x}_i) - \mathbf{n}_i\| \right) + \lambda \mathbb{E}_{\mathbf{x}}[(\|\nabla u_\theta(\mathbf{x})\| - 1)^2]$$

Which neural network architecture, which activation function?

The INR we are using

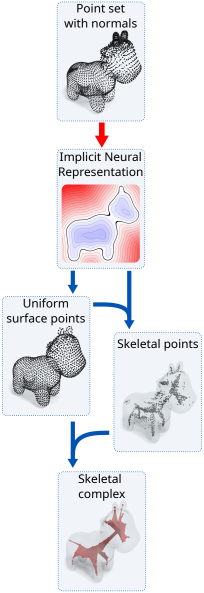

Architecture based on SIREN [Sitzmann 2020]

- MLP (6 layers, 64 neurons per layer, pretrained on a sphere SDF)
- Periodic activation function: $\sigma = \sin$

Now:

- Looking for a smooth approximation of the SDF
- Non-differentiabilities $\longleftrightarrow$ low gradient's norm

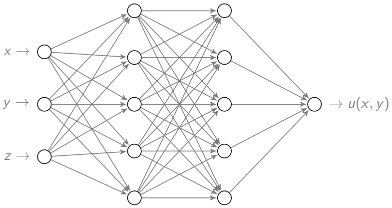

Far from the surface

- Infinite number of a.e. differentiable solutions to $\begin{cases} \|\nabla u\| & = 1 \\ u_{|\partial\Omega} & = 0 \\ \nabla u_{|\partial\Omega} & = \mathbf{n} \end{cases}$
- Blobs can appear

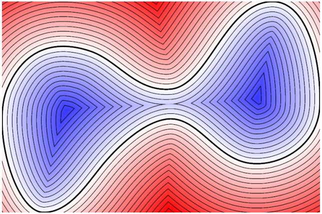

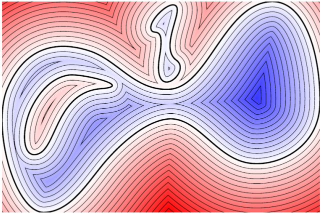

- Viscosity solution theory can help theoretically, but not practical

$$\|\nabla u_\varepsilon\| + \varepsilon \Delta u_\varepsilon = 1 \qquad u = \lim_{\varepsilon \to 0^+} u_\varepsilon$$

TV regularization

- Add a gradient's norm total variation regularization term in the loss function:

$$\mathcal{L}_{\mathrm{TV}} = \int_{\mathbb{R}^3} \|\nabla \|\nabla u\|(p)\| \mathrm{d}p$$

- Initial idea: minimize the measure of discontinuities / where $\|\nabla u\| < 1$ ...but wrong

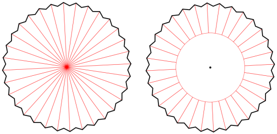

- Works with anchor points in the ambient space with coarse distance estimate

Loss function

- Surface loss:
$$\mathcal{L}_{\mathrm{s}}(\theta) = \int_{\partial\Omega} u_\theta(p)^2 \mathrm{d}p + \tau \int_{\partial\Omega} 1 - \frac{\mathbf{n}(p) \cdot \nabla u_\theta(p)}{\|\mathbf{n}(p)\| \, \|\nabla u_\theta(p)\|} \mathrm{d}p$$
- Eikonal loss:
$$\mathcal{L}_{\mathrm{e}}(\theta) = \int_{\mathbb{R}^3} \left(1 - \|\nabla u_\theta(p)\|\right)^2 \mathrm{d}p$$
- Learning points loss:
$$\mathcal{L}_{\mathrm{l}}(\theta) = \frac{1}{|\mathcal{P}|} \sum_{p \in \mathcal{P}} (u_\theta(p) - d(p))^2$$
- TV regularization loss:
$$\mathcal{L}_{\mathrm{TV}}(\theta) = \int_{\mathbb{R}^3} \|\nabla \|\nabla u_\theta\|(p)\| \mathrm{d}p$$

Final loss function

$$\mathcal{L} = \lambda_{\mathrm{e}} \mathcal{L}_{\mathrm{e}}(\theta) + \lambda_{\mathrm{s}} \mathcal{L}_{\mathrm{s}}(\theta) + \lambda_{\mathrm{l}} \mathcal{L}_{\mathrm{l}}(\theta) + \lambda_{\mathrm{TV}} \mathcal{L}_{\mathrm{TV}}(\theta)$$

Comparison

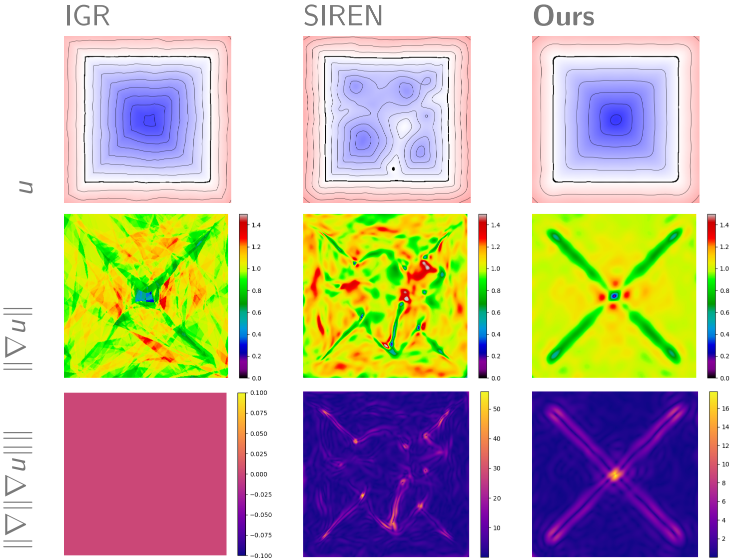
IGR
SIREN
Ours
u
$\|\nabla u\|$
$\|\nabla\|\nabla u\|\|$
1.4
1.2
1.0
0.8
0.6
0.4
0.2
0.0
0.100
0.075
0.050
0.025
0.000
−0.025
−0.050
−0.075
−0.100
50
40
30
20
10
16
14
12
10
8
6
4
2

Comparison – SDF slices

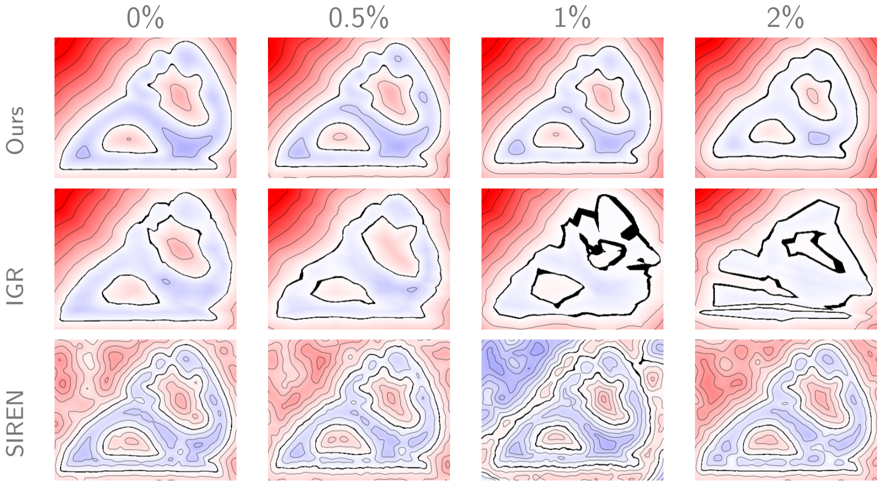

Comparison – $\|\nabla u\|$ slices

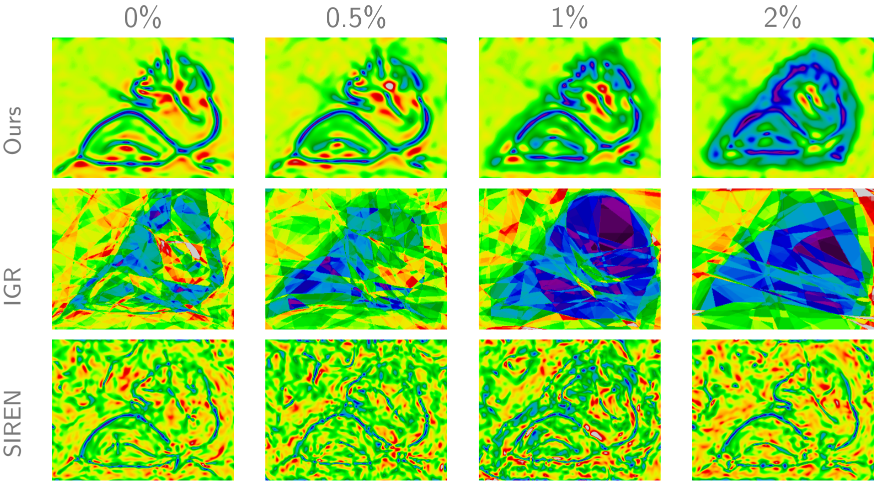

Uniform surface sampling

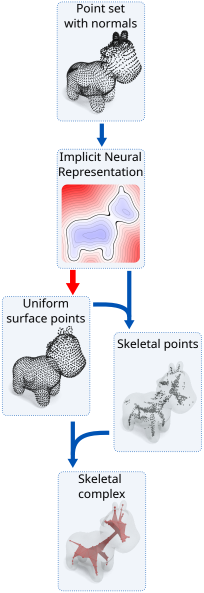

Uniform surface sampling method from [Yifan 2021]

1. Projection on the surface with Newton's method

$$p \leftarrow p - \frac{\nabla u(p)}{\|\nabla u(p)\|^2} u(p)$$

2. Uniformization with repulsion steps in the tangent plane using the k-nearest neighbors

Skeleton sampling

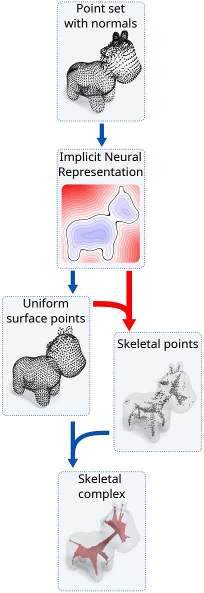

Lemma

Let $x \in \partial\Omega$. There exists $t > 0$ such that

$$x - t\nabla u_\Omega(x) \in \mathrm{med}(\Omega).$$

1. compute rays from the surface sample in the directions $-\nabla u$

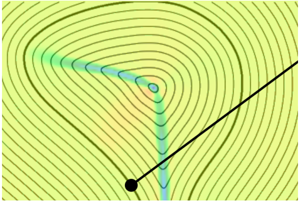

Skeleton sampling

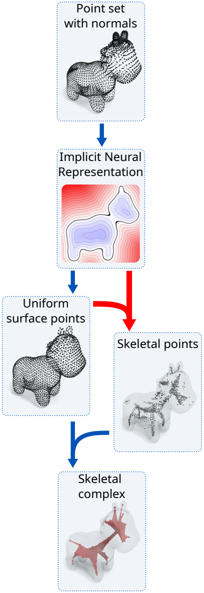

Lemma

Let $x \in \partial\Omega$. There exists $t > 0$ such that

$$x - t\nabla u_\Omega(x) \in \mathrm{med}(\Omega).$$

1. compute rays from the surface sample in the directions $-\nabla u$
2. find where they intersect the surface

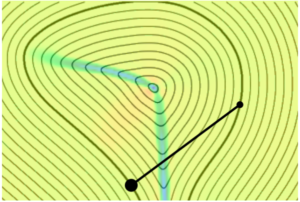

Skeleton sampling

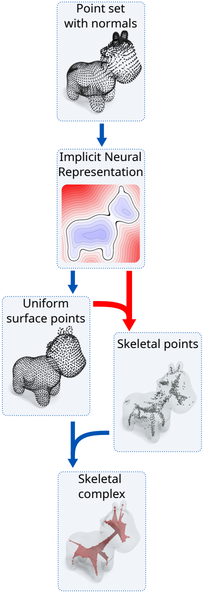

Lemma

Let $x \in \partial\Omega$. There exists $t > 0$ such that

$$x - t\nabla u_\Omega(x) \in \mathrm{med}(\Omega).$$

1. compute rays from the surface sample in the directions $-\nabla u$
2. find where they intersect the surface
3. sample them and find the smallest $\|\nabla u\|$

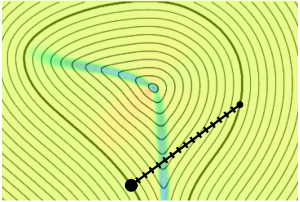

Skeleton sampling

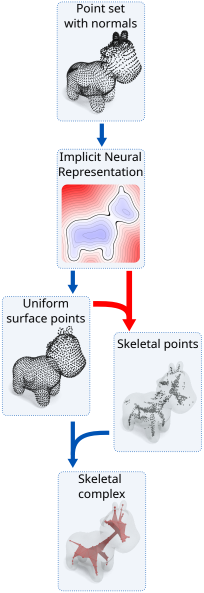

Lemma

Let $x \in \partial\Omega$. There exists $t > 0$ such that

$$x - t\nabla u_\Omega(x) \in \mathrm{med}(\Omega).$$

1. compute rays from the surface sample in the directions $-\nabla u$
2. find where they intersect the surface
3. sample them and find the smallest $\|\nabla u\|$

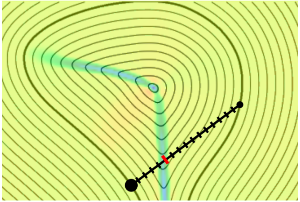

Skeletal points selection

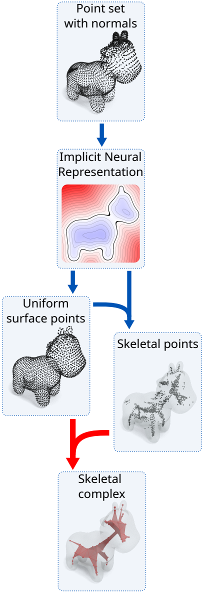

Set cover formulation from Coverage Axis [Dou 2022]

- N surface points (p_i), M skeletal points (s_j), $N \times M$ coverage matrix $\mathbf{D}$:

$$\mathbf{D}_{ij} = \begin{cases} 1 & \text{if } \|p_i - s_j\| \leq r_j + \delta \\ 0 & \text{otherwise} \end{cases}$$

- Set cover formulation (mixed-integer linear problem):

$$\begin{aligned} \min \quad & \|\mathbf{v}\|_1 \\ \text{s.t.} \quad & \mathbf{D}\mathbf{v} \succcurlyeq 1 \\ & \mathbf{v} \in \{0,1\}^M \end{aligned}$$

Skeletal points meshing

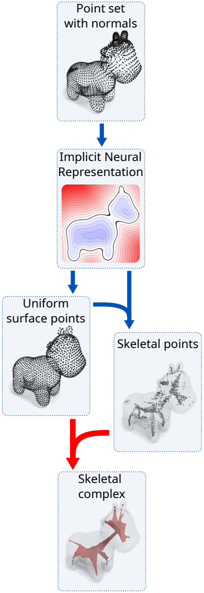

Final step: mesh the selected skeletal points

- Weighted Delaunay triangulation of selected skeletal points and the surface samples

$$\mathbf{RT}(\{(s_j, r_j), s_j \in S \mid v_j = 1\} \cup \{(p_i, \delta), p_i \in P\})$$

- Keep the edges and triangles between selected skeletal points appearing in this triangulation

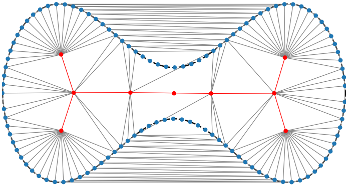

Results

Results

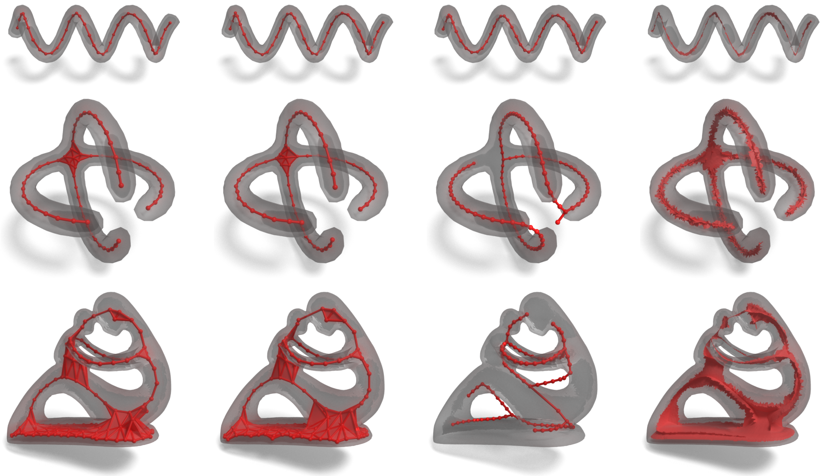

Ours | Coverage Axis | L_1-medial skeleton | Voxel Cores

Results

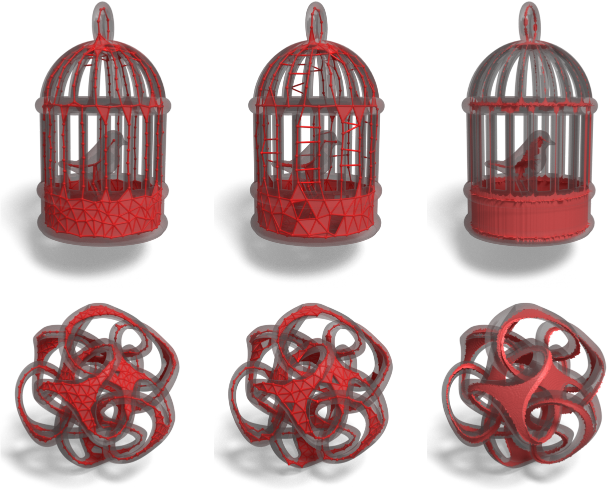

Ours | Coverage Axis | Voxel Cores

Noise robustness

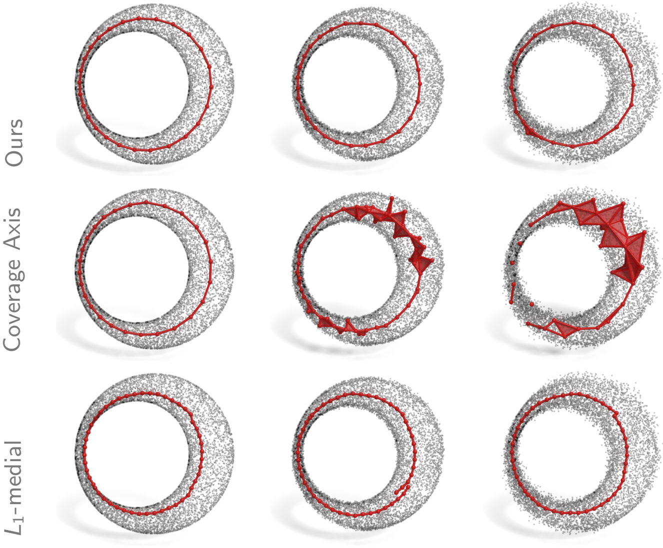

Noise robustness

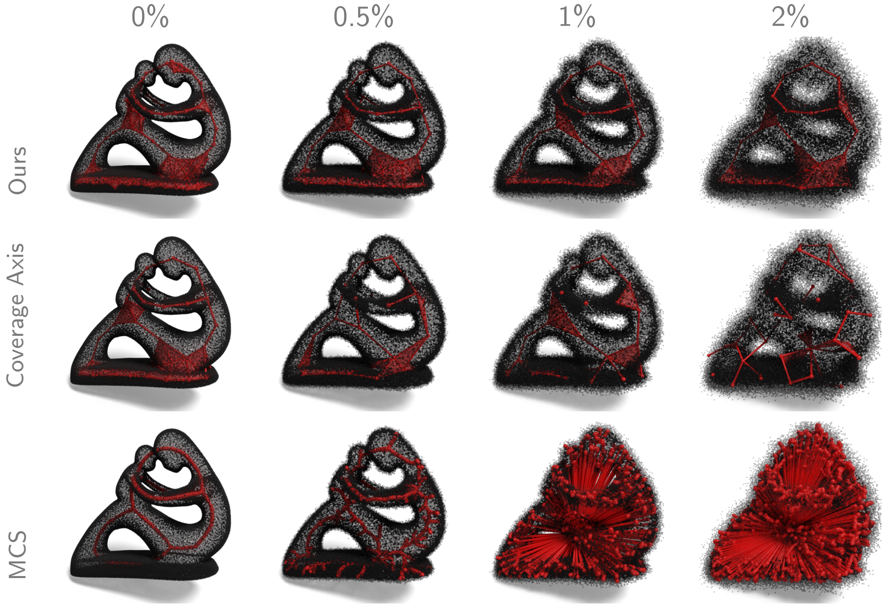

Noise robustness

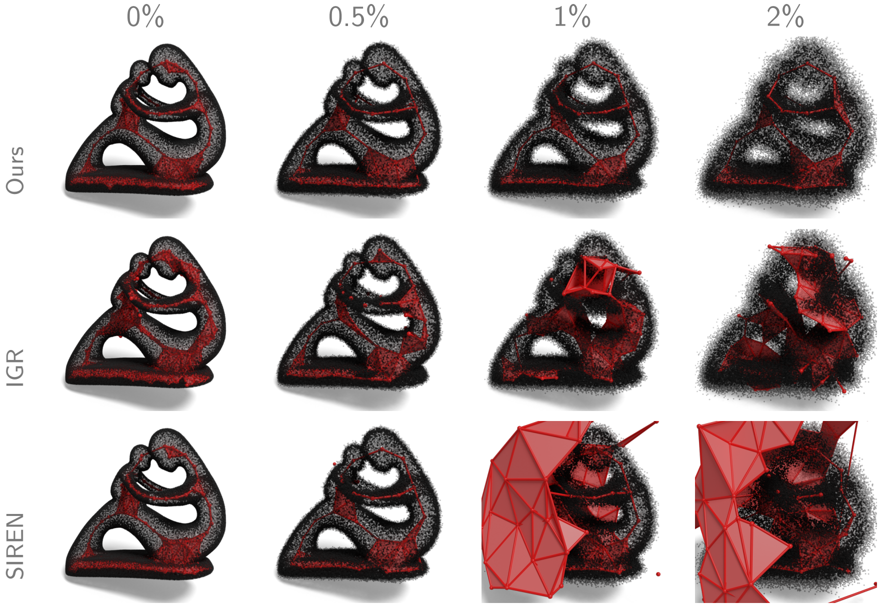

Missing data robustness

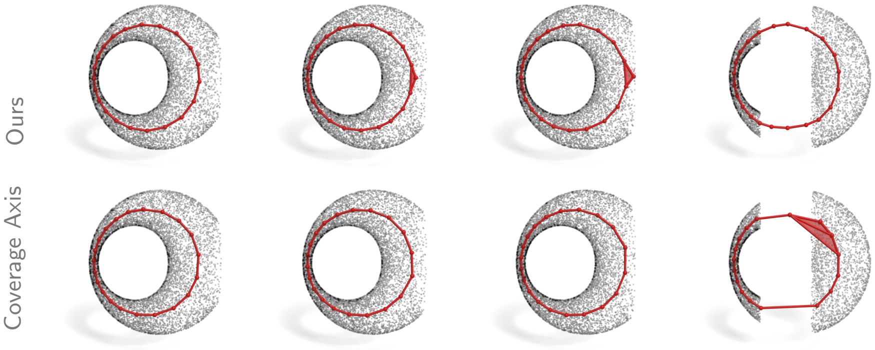

Ablation study table

Shape	Ours	No TV	ReLU No TV	SoftPlus	SoftPlus No TV	No uniform resampling	No learning loss
Noise 0.003	0.011	0.017	0.24	0.038	0.038	0.01	0.79
Noise 0.005	0.015	0.019	0.24	0.028	0.037	0.014	0.70
Noise 0.01	0.021	0.18	0.25	0.035	0.045	0.021	0.79
Noise 0.03	0.25	0.27	0.28	0.27	0.095	0.25	0.72
Truncated 1	0.13	0.30	0.28	0.15	0.15	0.29	0.72
Truncated 2	0.11	0.38	0.41	0.13	0.13	0.12	0.71
Truncated 3	0.12	0.27	0.27	0.18	0.14	0.12	0.72

Table: Ablation study on a torus with added noise and cropped parts (Hausdorff distance).

Limitations & future work

- Time: 2 minutes (laptop with Nvidia RTX 3050)

Limitations & future work

- Time: 2 minutes (laptop with Nvidia RTX 3050)
- Structure of the output

Limitations & future work

- Time: 2 minutes (laptop with Nvidia RTX 3050)
- Structure of the output
- No topological guarantees

Limitations & future work

- Time: 2 minutes (laptop with Nvidia RTX 3050)
- Structure of the output
- No topological guarantees
- Adaptation to latent shape space (DeepSDF [Park 2019])

Conclusion

- TV regularization term to enable skeleton extraction from an INR
- Code: `https://github.com/MClemot/SkeletonLearning` (Replicability Stamp)
- Funding: Agence Nationale de la Recherche, grant ANR-19-CE45-0015 (TOPACS)

Thank you!

Ablation

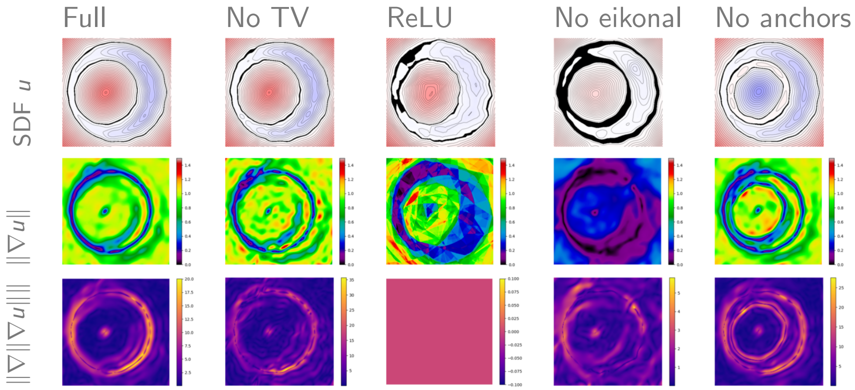
Full
No TV
ReLU
No eikonal
No anchors
SDF u
$||\nabla u||$
$||\nabla||\nabla u|||$

Comparison table

Shape	Ours	SIREN	IGR	MCS	Voxel Cores
clean	0.42	7.9	1.2	2.4	0.41
crop1	1.04	1.1	1.4	2.5	1.3
crop2	1.9	2.0	1.5	2.6	2.0
crop3	0.77	7.9	1.2	2.6	1.15
crop4	0.46	1.5	2.7	2.5	1.5
sub 25%	0.35	8.3	0.86	2.6	0.42
sub 50%	0.38	8.1	1.2	2.5	0.37
var 0.05%	0.46	8.3	1.3	2.5	0.40
var 0.1%	0.45	7.9	1.1	2.6	0.39
var 1%	0.49	0.79	1.9	2	0.67
var 2%	0.57	0.97	3	0.84	1.3

Table: Quantitative comparisons on a synthetic sphere-mesh shape, cropped or degraded with increasing noise (Hausdorff distance). Percentage values for the noise correspond to the noise variance (percentage of the diagonal).

Comparison – $\|\nabla\|\nabla u\|\|$ slices