Implicit Neural Representations in Geometry Processing

POPILSS

Julie Digne

19/06/2025

Joint work with **Camille Buonomo**, **Mattéo Clémot**, Raphaëlle Chaine

Geometry Processing

Geometry Processing

Turning real-world object surfaces into virtual representations and processing them.

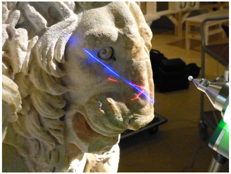

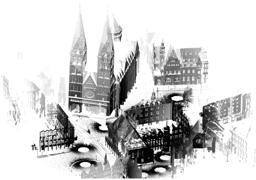

Bremen University

CNRS-MAP

Shape representation

- Point sets: sparse but no watertight PointNet [Qi 2017], KPConv [Thomas 19]...
- Meshes: efficient but needs to be computed EdgeConv [Hanocka 2019], DeltaConv [Wiersma 2022]
- Implicit Representation: hard to use for analysis purpose

Issues for Deep Learning on surfaces

Need for an equivariant convolution on the surface.

Shape Database

Lack of data variety (geometry and topology)

ShapeNet

Implicit Representation in Geometry Processing

- Representing a shape through a set of distances to a set of primitives [Bloomenthal 90]
- Mesh reconstruction (Marching Cubes [Lorensen 87]) or Direct Rendering (Sphere tracing [Hart 96])
- For surface reconstruction: from a point set build a signed distance field [Hoppe 92]
- Poisson Surface Reconstruction [Kazhdan 2006], [Alexa 2003] Moving Least Squares Surfaces

A long standing idea

Find a good function basis for representing the signed distance function.

Implicit Neural Representation in Geometry Processing

INR

Train a neural network to represent a shape. (Deep SDF [Park 2019], Occupancy Network [Mescheder 2018]).

- Optimization per shape: no database.

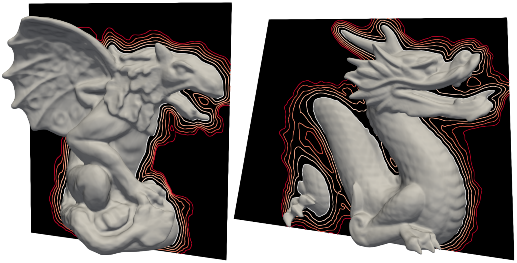

[Lipman 2019]

Implicit Neural Representation in Geometry Processing

INR

Train a neural network to represent a shape. (Deep SDF [Park 2019], Occupancy Network [Mescheder 2018]).

- Optimization per shape: no database.

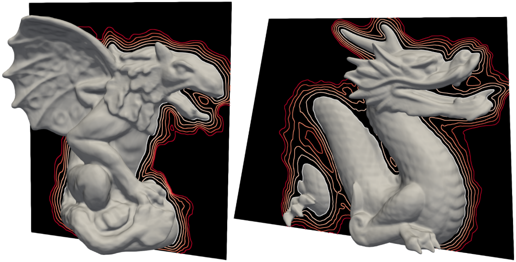

[Lipman 2019]

Focus on two problems:

- Shape analysis: extract the topology of a shape
- Shape synthesis: interpolate between two shapes

Signed distance fields

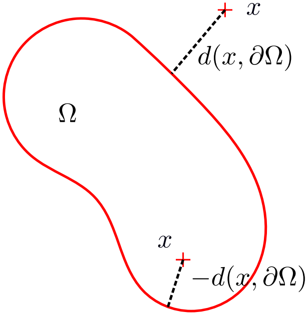

Signed distance field are useful

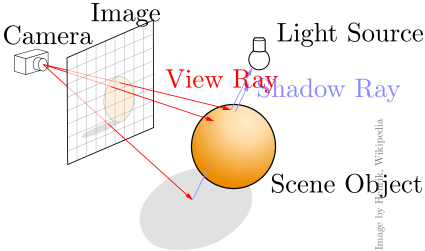

- Requires to compute ray/surface intersection.
- Direct intersection with explicit representations (Meshes/Geometric primitives)

Sphere tracing [Hart 1996]

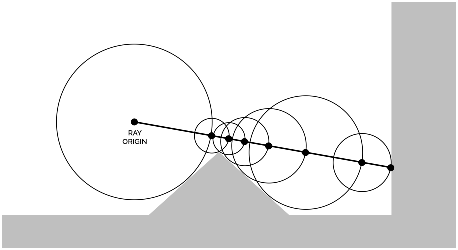

1. Input: a point x and direction $\mathbf{v}$, a signed distance field u.

2. Initialize $t = 0$

3. While $t < D$

  3.1 $x_t = x + t\mathbf{v}$

  3.2 $d = u(x_t)$

  3.3 If $d < \varepsilon$ Return x_t

  3.4 Else Increment $t = t + d$

Neural Skeleton: Regularizing INR away from the surface

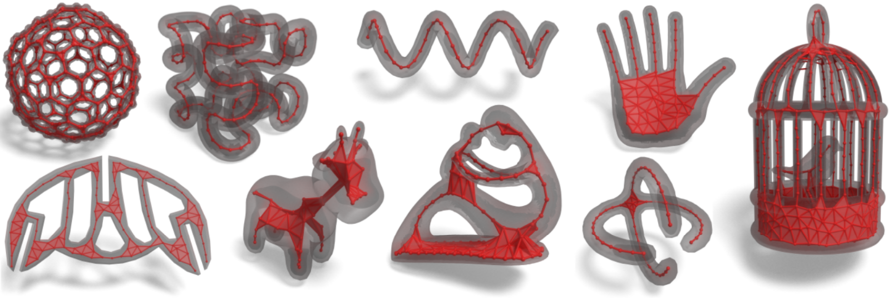

- *Neural skeleton: implicit neural representation away from the surface*, Mattéo Clémot and Julie Digne 2023

Extract the medial axis of a shape based on an INR

- Extract topological data from shapes (genus, medial axis) even with noise, missing data and outliers.
- All the topological information is included within the signed distance field.
- How do we represent this signed distance field? Implicit Neural Representation

Medial Axis

Definition

A point p belongs to the medial axis of a compact shape if it has at least two distinct nearest neighbors on the shape surface.

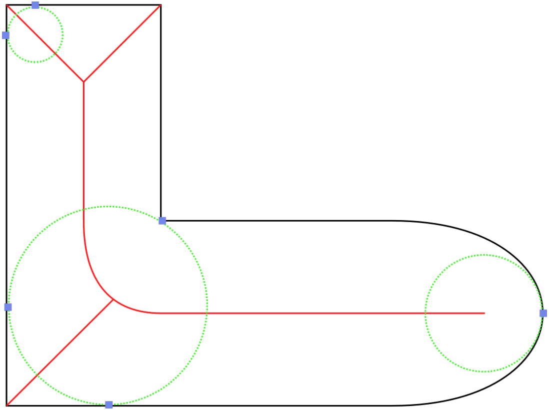

Medial Axis: classical methods

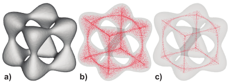

from [Tagliasacchi et al. 2016]

from [Tagliasacchi et al. 2016]

- Curve Skeleton (Mesh Contraction [Au 2008], Mean Curvature Flow [Tagliasacchi 2012], visual hull [Livesu 2012], local symmetries [Tagliasacchi 2009, Huang 2012])
- Computational Geometry: Voronoi subcomplex [Dey 2002], power crust [Amenta 2001]
- Signed distance field and voxelization: VoxelCores [Yan 2018]

Result

Often needs to be *compressed*. [Dou 2022]

Optimization Process

- Input data a set of points $(x_i, \mathbf{n}_i), i \in I$
- Look for u continuous and a.e. $\mathcal{C}^1$ such that:

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

Optimization Process

- Input data a set of points $(x_i, \mathbf{n}_i), i \in I$
- Look for u continuous and a.e. $\mathcal{C}^1$ such that:

$$\begin{cases} \|\nabla u\| & = 1 \\ u_{|\partial \mathcal{S}} & = 0 \\ \nabla u_{|\partial \mathcal{S}} & = \mathbf{n} \end{cases} \quad (1)$$

- Loss [Gropp 2020]

$$l(\theta) = \int_{x \in \mathcal{S}} (|u_\theta(x)|^2 + \tau \|\nabla u_\theta(x) - \mathbf{n}(x)\|^2) dx + \lambda \int_{x \in \mathbb{R}^3} (\|\nabla u_\theta(x)\| - 1)^2 dx$$

Optimization Process

- Input data a set of points $(x_i, \mathbf{n}_i), i \in I$
- Look for u continuous and a.e. $\mathcal{C}^1$ such that:

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

- Loss [Gropp 2020]

$$l(\theta) = \int_{x \in \mathcal{S}} (|u_\theta(x)|^2 + \tau \|\nabla u_\theta(x) - \mathbf{n}(x)\|^2) dx + \lambda \int_{x \in \mathbb{R}^3} (\|\nabla u_\theta(x)\| - 1)^2 dx$$

- Approximation by Monte Carlo

Eikonal Equation

- Infinite number of solutions
- Viscosity solution theory: allows to select the right solution
- Use smooth eikonal equation (not practical [Lipman 2019])

$$\|\nabla u\| - \varepsilon \Delta u = 1$$

- Consequence: blobs appear

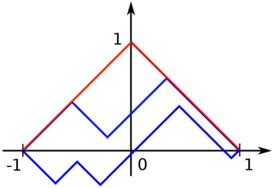

[Camilli 2014]

Infinite number of solutions

Not an issue close to the surface – but far away?

Which neural network?

- MLP (6 layers, 128-256 neurons/layer) with ReLU activation functions
- ReLU: piecewise constant gradient.
- Sitzman (2021) replaces ReLU with sine activation function: smooth function

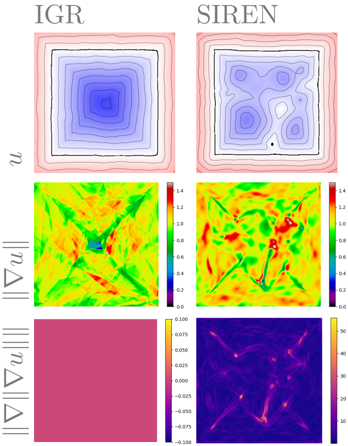

TV regularization

- Look for a smooth surrogate for the signed distance function
- Medial axis: zeros of the gradient
- Add the TV of the gradient norm.

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

- Rationale: minimize the measure of the zeros of the gradient set (counter-example!)

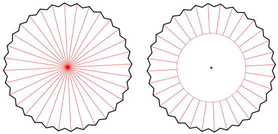

TV regularization - some analysis

- The TV term favors that u has no second order differential content along the gradient lines

Since $\nabla u = (u_x, u_y, u_z)$, it follows:

$$\begin{aligned}
\nabla \|\nabla u\| &= \nabla \sqrt{u_x^2 + u_y^2 + u_z^2} \\
&= \frac{1}{2\|\nabla u\|} \begin{pmatrix} 2u_x u_{xx} + 2u_y u_{xy} + 2u_z u_{xz} \\ 2u_x u_{xy} + 2u_y u_{yy} + 2u_z u_{yz} \\ 2u_x u_{zx} + 2u_y u_{zy} + 2u_z u_{zz} \end{pmatrix} \\
&= H_u \frac{\nabla u}{\|\nabla u\|}
\end{aligned}$$

Total loss

- Eikonal loss:

$$\mathcal{L}_{eikonal} = \int_{\mathbb{R}^3} (1 - \|\nabla u(p)\|)^2 \, dp \tag{2}$$

- Surface loss:

$$\mathcal{L}_{\text{surface}} = \int_{\partial\Omega} u(p)^2 dp + \int_{\partial\Omega} 1 - \frac{\mathbf{n}(p) \cdot \nabla u(p)}{\|\mathbf{n}(p)\| \, \|\nabla u(p)\|} dp \tag{3}$$

- Learning point loss

$$\mathcal{L}_{\text{learning}} = \sum_{p \in \mathcal{P}} (u(p) - d(p))^2 + \sum_{p \in \mathcal{P}} 1 - \frac{\nabla u(p) \cdot \nabla d(p)}{\|\nabla u(p)\| \, \|\nabla d(p)\|} \tag{4}$$

- + TV loss

Loss

$$\mathcal{L} = \lambda_e \mathcal{L}_{eikonal} + \lambda_s \mathcal{L}_{surface} + \lambda_l \mathcal{L}_{learning} + \lambda_{TV} \mathcal{L}_{TV} \tag{5}$$

Resulting Fields

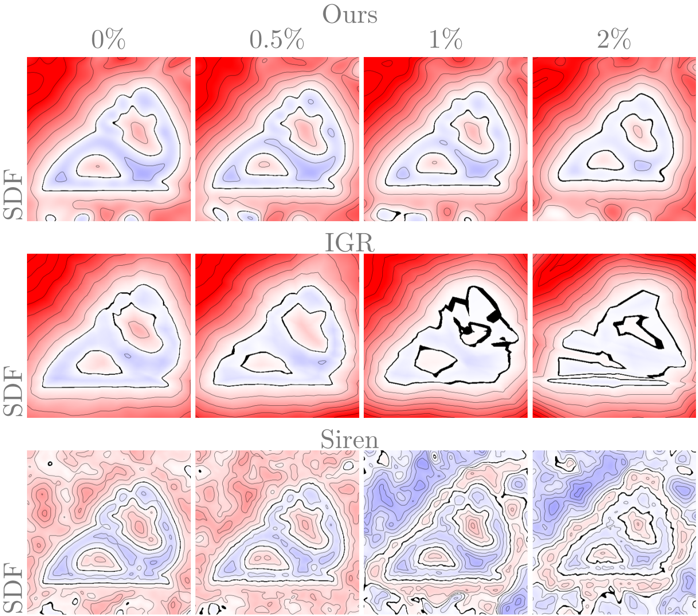

$\|\nabla u\|$

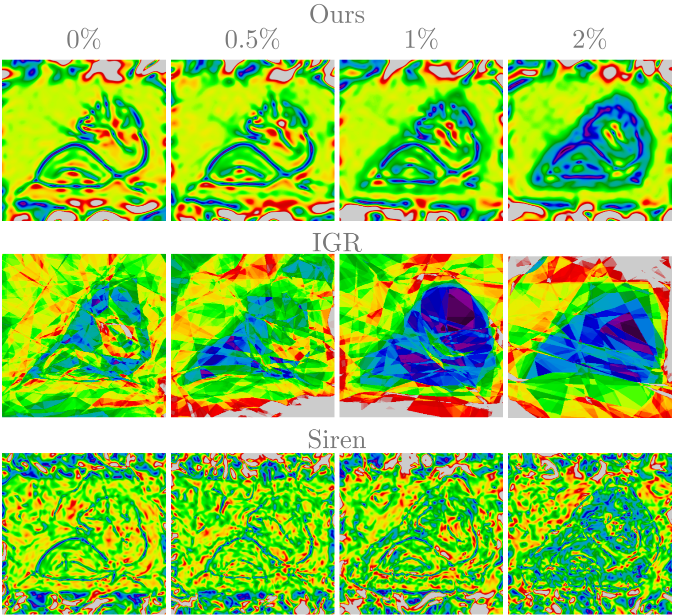

$\nabla \|\nabla u\|$

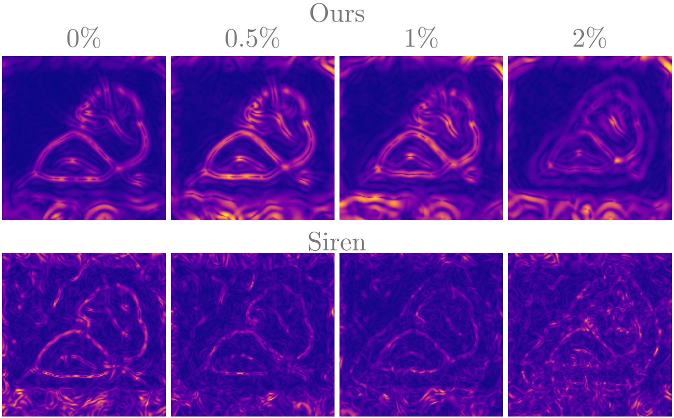

Overview

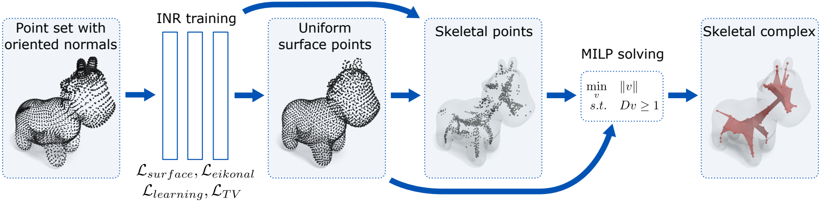
Point set with oriented normals
INR training
$\mathcal{L}_{surface}, \mathcal{L}_{eikonal}$
$\mathcal{L}_{learning}, \mathcal{L}_{TV}$
Uniform surface points
Skeletal points
MILP solving
$\min_v \|v\|$ $s.t.$ $Dv \geq 1$
Skeletal complex

Uniform surface sampling

- Why? used later for skeletal compression
- Sample N points in the ambient space and iterate Newton steps

$$p \leftarrow p - \frac{\nabla u(p)}{\|\nabla u(p)\|^2} u(p) \tag{6}$$

- Regularization by repulsion on the tangent plane
- Reprojection on the surface

GPU skeleton tracing

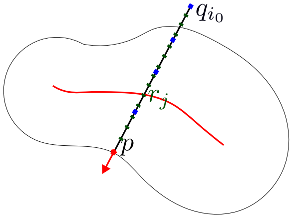

- Sample N points (p_i) on the surface using Newton's method
- For $i = 1 \cdots N$
  - Sample n points $q_j = p_i - t\frac{h}{n}\frac{\nabla u(p_i)}{\|\nabla u(p_i)\|}$ $(t = 1 \cdots n)$
  - Find i_0 the smallest index such that $u(q_{i_0}) > 0$
  - Sample n points r_j between p and q_{i_0}
  - Find r_j with lowest $\|\nabla u(r_j)\|$ value

Simplicial complex extraction [Dou 2022]

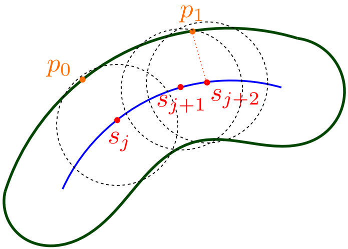

- N points x_i, M skeletal points s_i with distance r_i.
- Coverage matrix: D $(N \times M)$

$$D_{ij} = 1 \text{ if } \|p_i - s_j\| - r_j \leq \delta \text{ and } 0 \text{ otherwise}$$

- Mixed Integer Linear Problem:

$$\begin{array}{cl} \min & \|v\|_2 \\ \text{s.t.} & Dv \succeq 1 \end{array} \tag{7}$$

- Link selected points (weighted triangulation)

Results

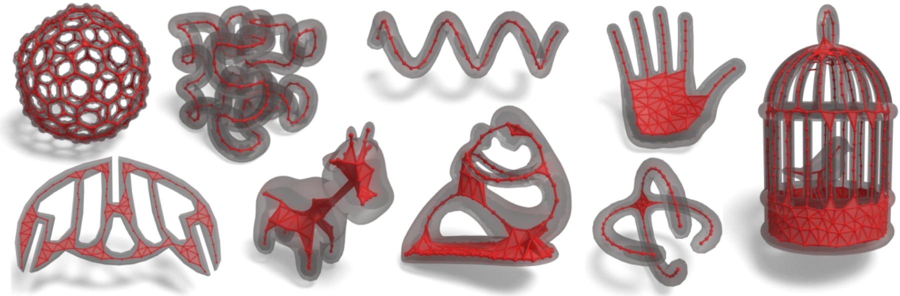

Results

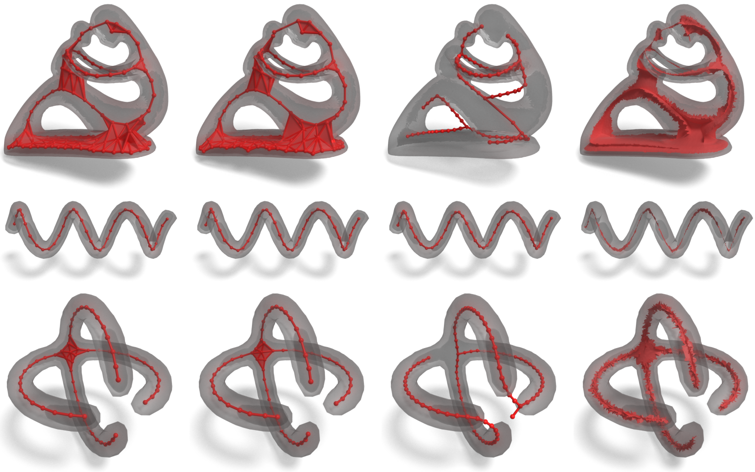

Ours

Coverage Axis

L_1-medial skeleton

VoxelCores

Results

Ours Coverage Axis VoxelCores

results

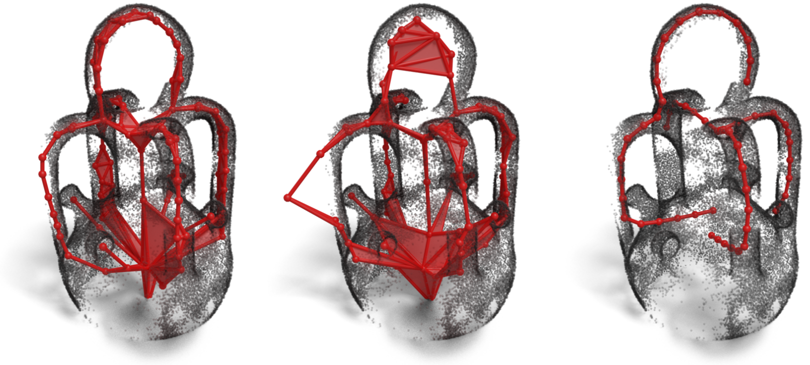

Ours | Coverage Axis | L_1-medial skeleton

Results

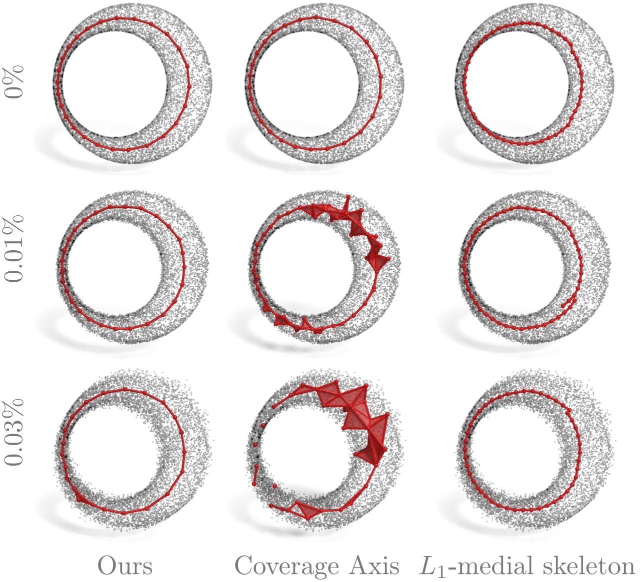

Results

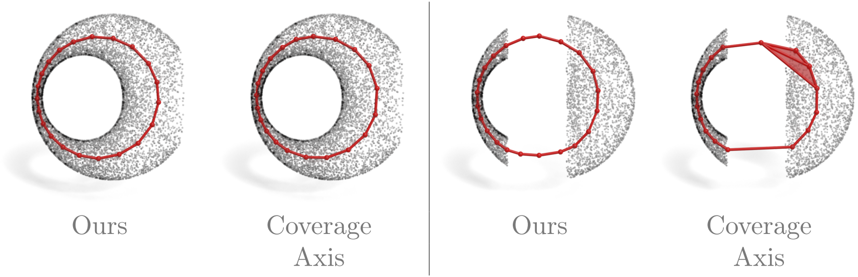

Ours | Coverage Axis | Ours | Coverage Axis

With noise

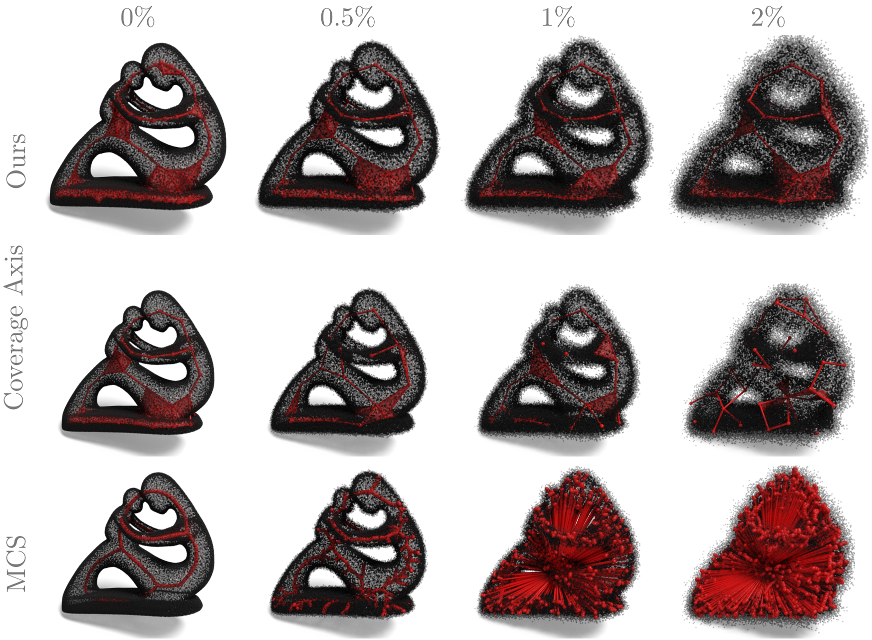

With noise

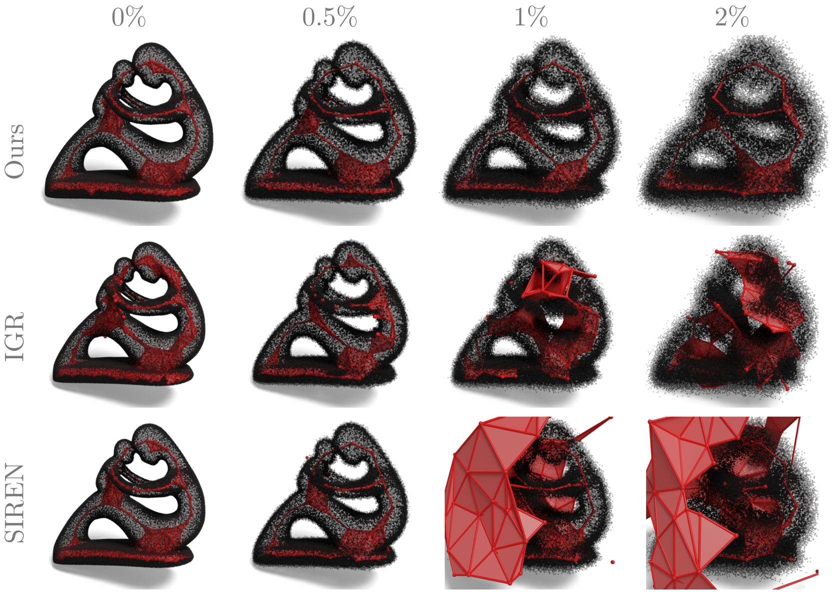

Beyond analysis: shape synthesis

Problem statement

Given two shapes $\mathcal{S}_0$ and $\mathcal{S}_1$ find a continuum of shapes $\mathcal{S}_t$, $(0 \leq t \leq 1)$ interpolating between $\mathcal{S}_0$ and $\mathcal{S}_1$

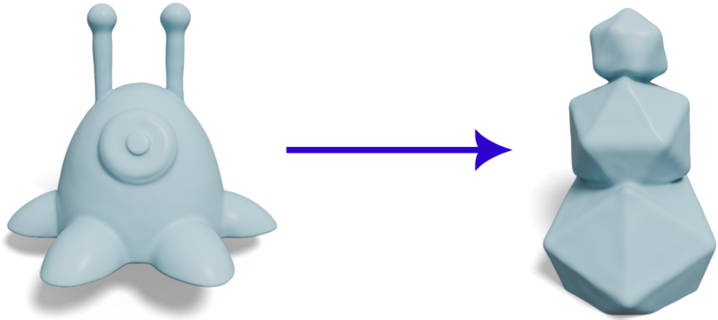

Beyond analysis: shape synthesis

Problem statement

Given two shapes $\mathcal{S}_0$ and $\mathcal{S}_1$ find a continuum of shapes $\mathcal{S}_t$, $(0 \leq t \leq 1)$ interpolating between $\mathcal{S}_0$ and $\mathcal{S}_1$

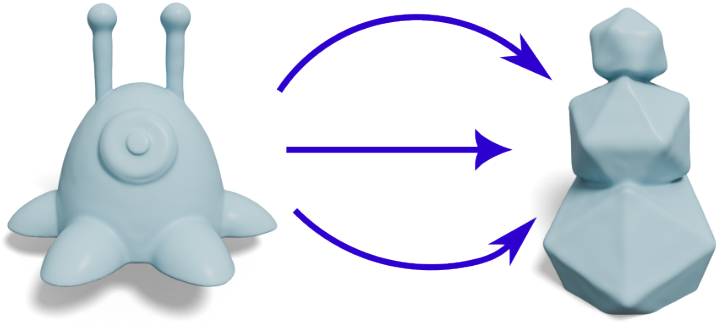

Shape interpolation: an ill-posed problem

Needs further hypotheses:

- As rigid as possible

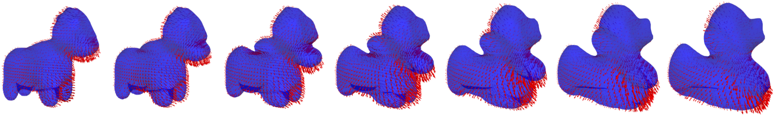

[Tao 2016]

Shape interpolation: an ill-posed problem

Needs further hypotheses:

- As rigid as possible
- Landmark correspondences

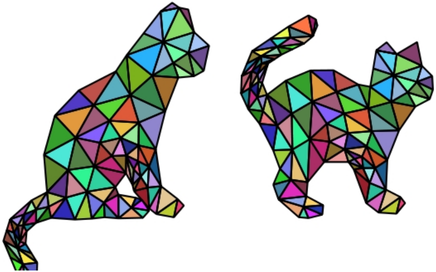

[Baxter 2011]

Shape interpolation: an ill-posed problem

Needs further hypotheses:

- As rigid as possible
- Landmark correspondences
- Least displacement

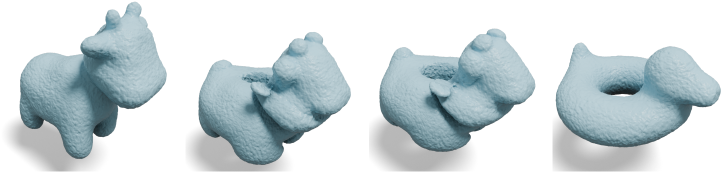

GeomLoss [Feydy 17,19]

Shape interpolation: an ill-posed problem

Needs further hypotheses:

- As rigid as possible
- Landmark correspondences
- Least displacement
- Volume preserving

[Eisenberger 2018]

Volume Preserving shape morphing

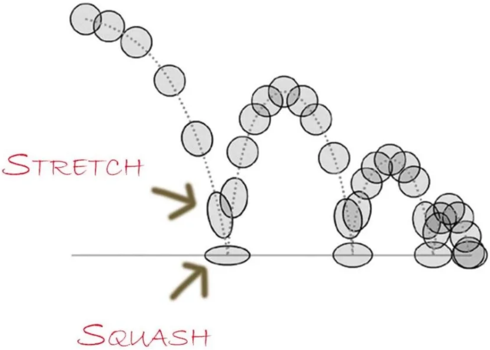

[Disney's 12 Principles of Animation]

Our setting

Find a volume preserving shape morphing without explicit correspondences.

Deformation vector field

Advection

Look for a vector field $V(x,t)$ deforming a $\mathcal{S}_0$ into $\mathcal{S}_1$

ODE for the motion of a point p:

$$\dot{p}(t) = V(p(t), t) \tag{8}$$

Deformation vector field

Advection

Look for a vector field $V(x, t)$ deforming a $\mathcal{S}_0$ into $\mathcal{S}_1$

ODE for the motion of a point p:

$$\dot{p}(t) = V(p(t), t) \tag{8}$$

How do we link V with the implicit representation?

Level-set equation (LSE)

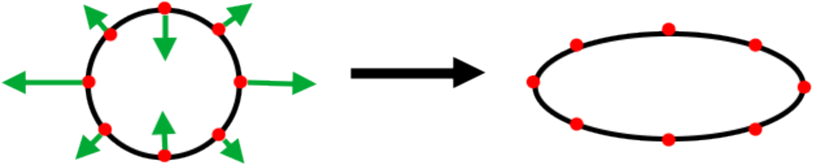

Motion of $p \in \mathcal{S}_0$:

$$\dot{p}(t) = \mathbf{V}(p(t), t) \text{ with } p(0) = p$$

$$f(x,t) \text{ verifies : } f(p(t),t) = 0$$

$$\frac{df(p(t),t)}{dt} = \frac{\partial f}{\partial t}(p(t),t) + \langle \nabla f(p(t),t), \dot{p}(t) \rangle = 0$$

See e.g. [Osher 2000]

Level-set equation (LSE)

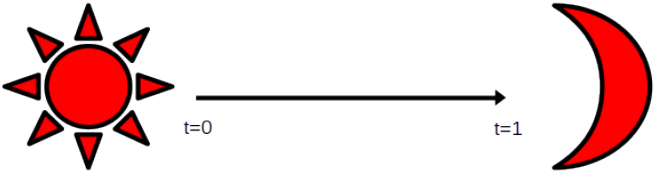

$\mathcal{S}_0$ $\mathcal{S}_1$

$$\begin{cases} \frac{\partial f}{\partial t} + \langle \nabla f, \mathbf{V} \rangle & = 0 \text{ on } \mathbb{R}^d \times [0,1] \\ f(x,0) & = g_0(x) \ \forall x \in \mathbb{R}^d \\ f(x,1) & = g_1(x) \ \forall x \in \mathbb{R}^d \end{cases}$$

In the neural framework

- Given $\mathbf{V}$

LSE loss:

$$l_{LSE} = \int_{\mathbb{R}^d \times [0,1]} \mid \frac{\partial f_\theta}{\partial t} - \langle \nabla f_\theta, \mathbf{V} \rangle \mid \, \mathrm{d}p\mathrm{d}t$$

Boundary condition :

$$l_{Dirichlet} = \sum_{i=0,1} \int_{\mathbb{R}^d} \mid g_i - f_\theta(.,i) \mid \, \mathrm{d}p + \sum_{i=0,1} \int_{x \in \partial \mathcal{S}_i} |f_\theta(x,i)| dx$$

$$l_{Neumann} = \sum_{i=0,1} \int_{\mathcal{S}_i} \mid 1 - \langle \nabla f_\theta(.,i), \vec{N}_{\mathcal{S}_i} \rangle \mid \, \mathrm{d}p$$

Eikonal loss :

$$l_{Eik} = \int_{\mathbb{R}^d \times [0,1]} \mid 1 - \mid \nabla f_\theta(p,t) \mid\mid \, \mathrm{d}p\mathrm{d}t$$

Network: fully connected networks (6 hidden layers, 128 neurons per layer, and Sine activations).

LSE with handcrafted V (Nise) [Novello et al. 2023]

for interpolation:

$$\mathbf{V}(p,t) = -(g_1(p) - f_\theta(p,t))\frac{\nabla f_\theta(t,p)}{\|\nabla f_\theta(t,p)\|}$$

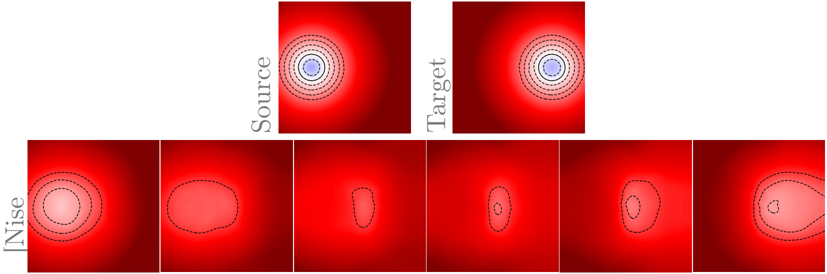

Joint learning of f_θ and $\mathbf{V}_\theta$

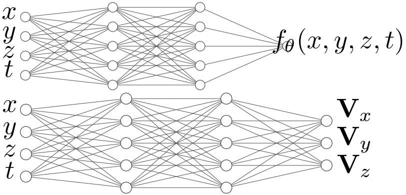

- Estimate both f_θ and $\mathbf{V}$ through two neural networks trained end-to-end

$$l_{LSE} = \int_{\mathbb{R}^d \times [0,1]} \frac{\partial f_\theta}{\partial t} - \langle \nabla f_\theta, \mathbf{V}_\theta \rangle \mid \mathrm{d}p\mathrm{d}t$$

Joint learning of f_θ and $\mathbf{V}_\theta$

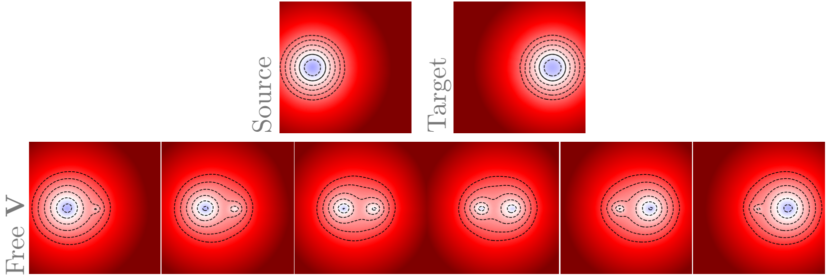

Joint learning of f_θ and $\mathbf{V}_\theta$

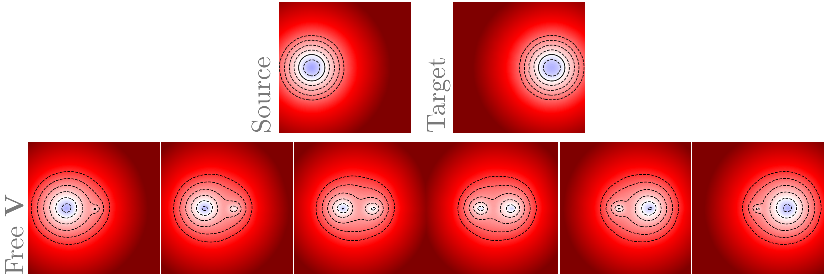

Fails with a simple translation.

Volume preservation

Sufficient condition

If $\operatorname{div} \mathbf{V} = 0$, then the shape advected by $\mathbf{V}$ has constant volume.

See also [Richter-Powell 2022].

Volume preservation

Sufficient condition

If $\operatorname{div} \mathbf{V} = 0$, then the shape advected by $\mathbf{V}$ has constant volume.

- To build a divergence free vector field $\mathbf{V}$, estimate $\mathbf{D}$ and set $\mathbf{V} = \operatorname{curl} \mathbf{D}$.

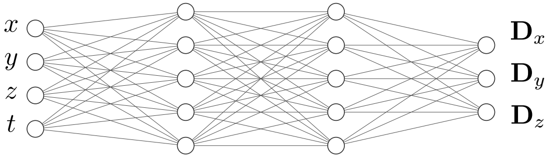

See also [Richter-Powell 2022].

Result

Div.-free

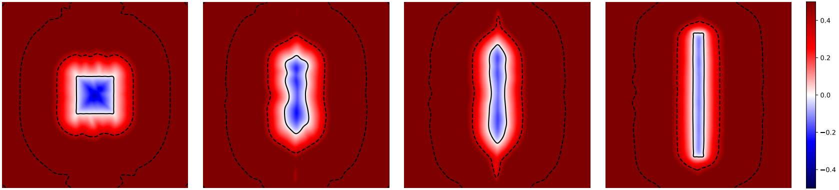

Eikonal incompatibility

Remark

If $\mathrm{div}(\mathbf{V}) = 0$, $\mathbf{V}$ preserves the volume of *every level set.*

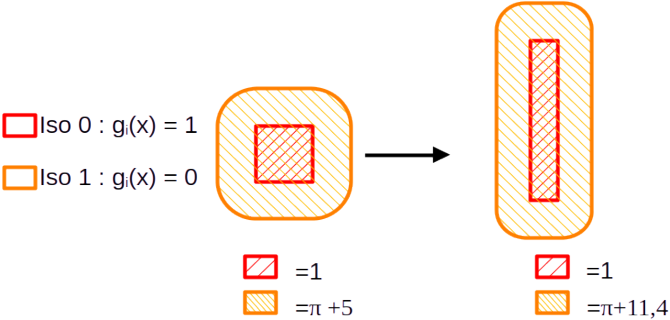

Only the conservation of the 0-levelset is necessary.

Adaptive-divergence

Idea

Augment a divergence free vector field by a quantity that vanishes along the integral lines $\partial \mathcal{S}_0$.

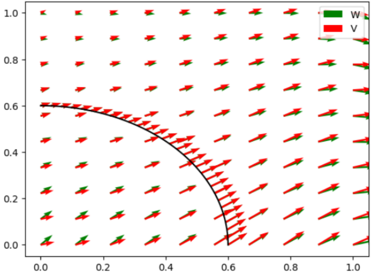

Adaptive divergence : Formal definition

Definition

$\mathbf{V} \in \mathbb{R}^d$ has adaptive divergence w.r.t. $\mathcal{S}_0$ if there exists a divergence-free $\mathbf{W}$ with associated flow $\phi_{\mathbf{W}}$ and a vector field $\mathbf{F}$ s.t. $\forall t \in [0,1]$:

$$\forall x \in \mathbb{R}^d \; \mathbf{V}(x,t) = \mathbf{W}(x,t) + \mathbf{F}(x,t)$$

$$\forall x \in \partial\mathcal{S}_0 \; \mathbf{F}(\phi_{\mathbf{W}}(x,t),t) = 0$$

$$\forall x \in \partial\mathcal{S}_0 \; \operatorname{div}(\mathbf{F})(\phi_{\mathbf{W}}(x,t),t) = 0 \text{ so } \operatorname{div}(\mathbf{V})(\phi_{\mathbf{W}}(x,t),t) = 0$$

Volume preservation with *Adadiv*

Theorem (Volume Preservation)

If $\mathbf{V}$ *has adaptive divergence w.r.t.* $\mathcal{S}$ *and corresponding flow* $\phi_{\mathbf{V}}$*:*

$$\mathrm{d}_t Vol(\phi_{\mathbf{V}}(\mathcal{S}, t)) = 0$$

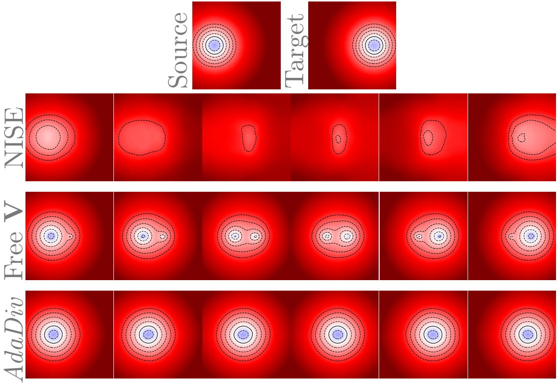

Adaptive Divergence - Numerical construction

$$\mathbf{V}_\theta = \mathrm{curl}(\mathbf{D}_\theta) + \beta(f_\theta)\nabla u_\theta$$

$$\beta(x) : \mathbb{R} \longrightarrow \mathbb{R}$$

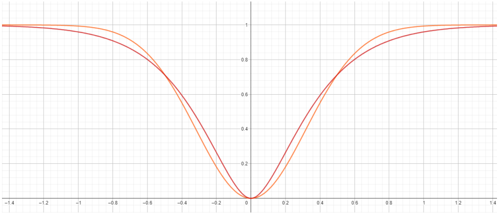

$1 - e^{-\alpha|x|}(1 + \alpha|x|)$ $\quad$ $1 - e^{-\alpha x 2}$

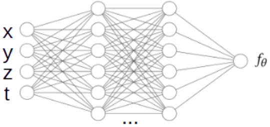

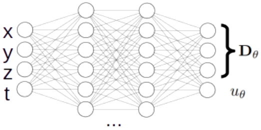

Result in 2D

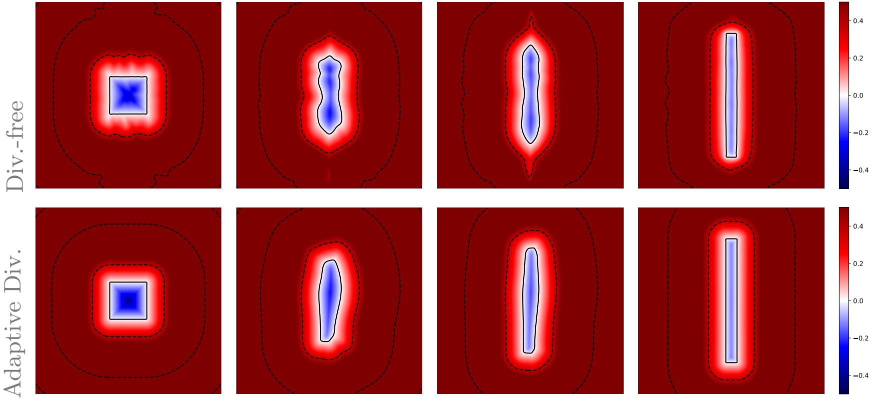

Result in 2D

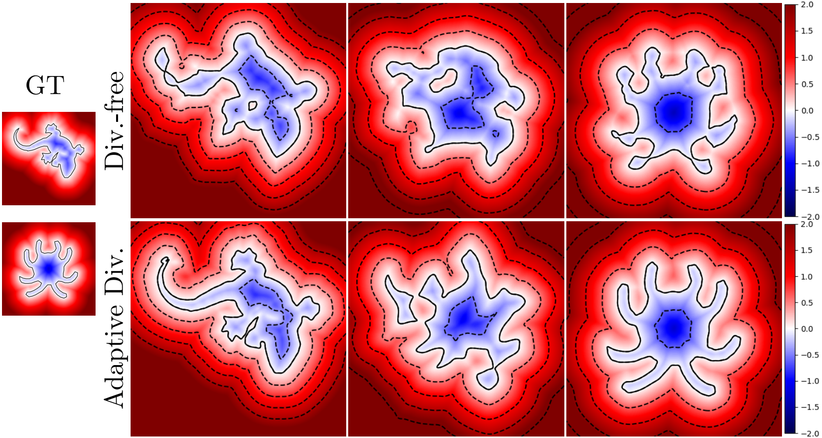

Results

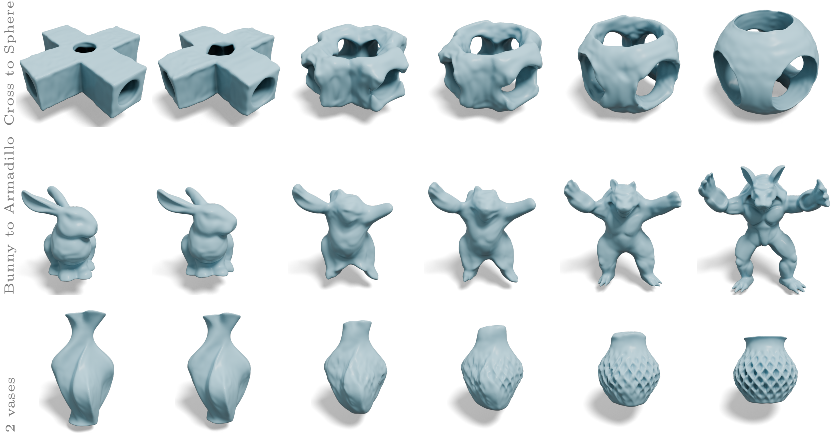

Topological changes

Weak link between f_θ and $\mathbf{V}_\theta$

Optimization problem :

$$l_{LSE} = \int_{\mathbb{R}^d \times [0,1]} \frac{\partial f_\theta}{\partial t} - \langle \nabla f_\theta, \mathbf{V}_\theta \rangle \mid \mathrm{d}p\mathrm{d}t$$

$$\partial_t f_\theta + \langle \nabla f_\theta, \mathbf{V}_\theta \rangle \neq 0$$

$\Rightarrow$ No exact volume conservation, even if $\mathrm{div}(\mathbf{V}_\theta) = 0$

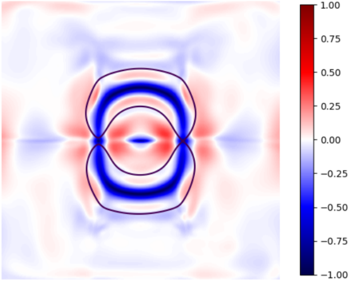

LSE loss

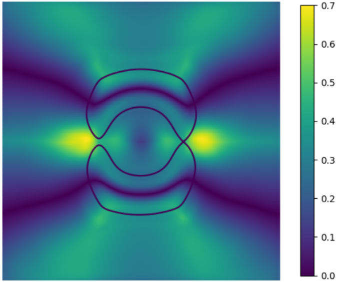

Eikonal loss

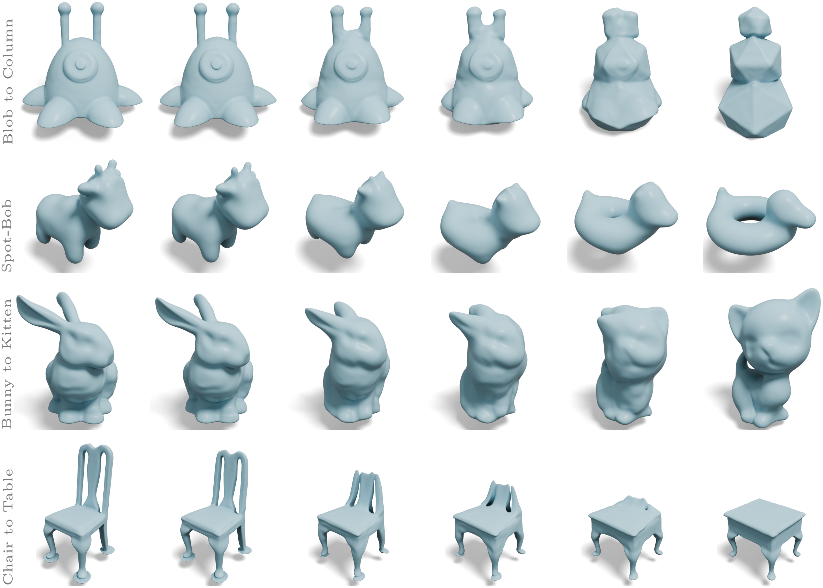
Blob to Column
Spot-Bob
Bunny to Kitten
Chair to Table

Results

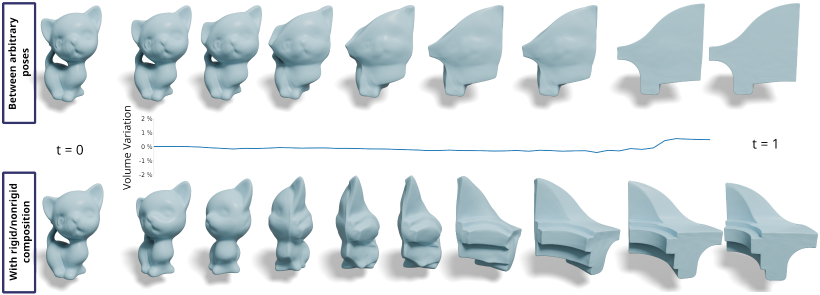
Between arbitrary poses
t = 0
Volume Variation
2 %
1 %
0 %
-1 %
-2 %
t = 1
With rigid/nonrigid composition

Volume evolution comparison

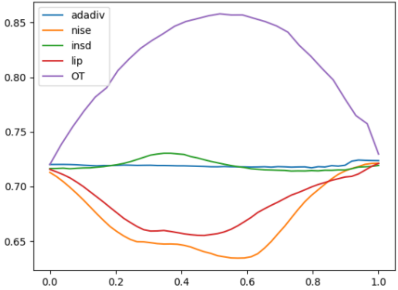

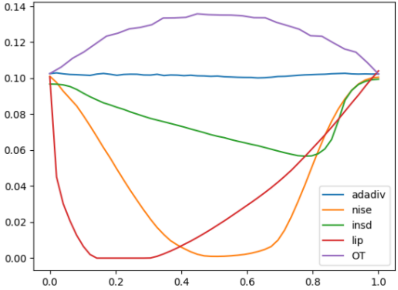

- Adaptive divergence - Our method
- Neural implicit surface evolution [Novello et al. 2023]
- landmark free INSD [Sang et al. 2025]
- Lipschitz-MLP[Liu et al. 2022]
- Optimal transport with Sinkhorn loss [Feydy et al. 2018]

Video

Robustness to noise

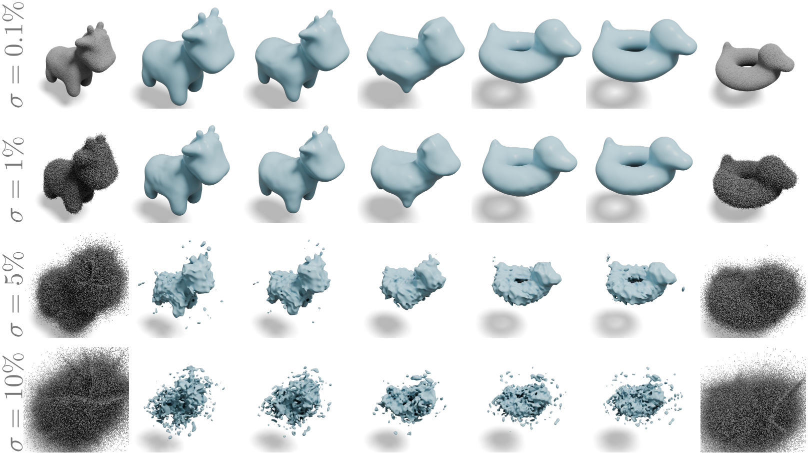

Conclusion

- INR biased toward smoothness (noise robustness vs loss of details)
- Useful for shape analysis
- Relaxation of the divergence free constraint
- Consistent intermediate steps with volume preservation
- Better with 1-Lipschitz guaranteed neural networks?

- *Volume Preserving Neural Shape Morphing*, C. Buonomo, J. Digne, R. Chaine, Computer Graphics Forum, Symposium on Geometry Processing, 2025
- *Neural skeleton: implicit neural representation away from the surface*, M. Clémot and J. Digne, Shape Modeling International 2023 (best paper award)