NeuroMorph: Unsupervised Shape Interpolation and Correspondence in One Go

Marvin Eisenberger∗,† , David Novotny∗ , Gael Kerchenbaum∗ , Patrick Labatut∗ , Natalia Neverova∗ , Daniel Cremers† , Andrea Vedaldi∗ Facebook AI Research∗ , Technical University of Munich†

Seminar presentation by Askar Kolushev

Problem description

Relate two 3D shapes

Source

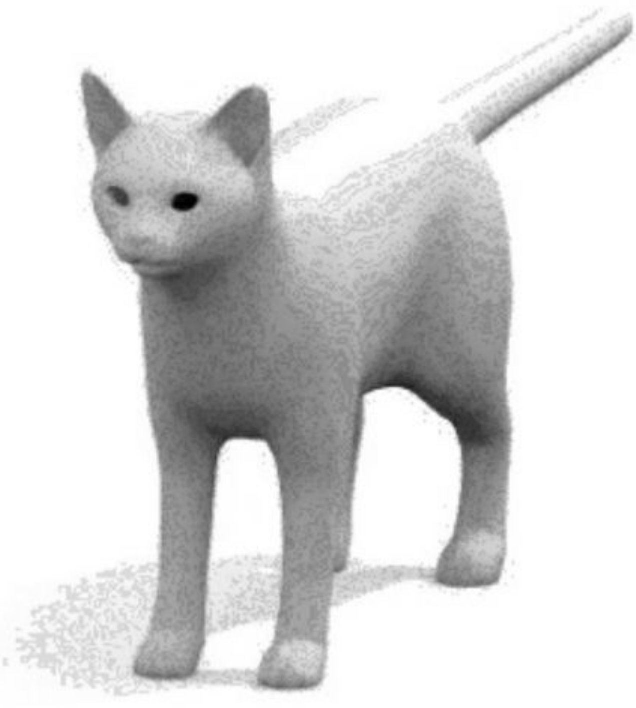

Target

Shape interpolation

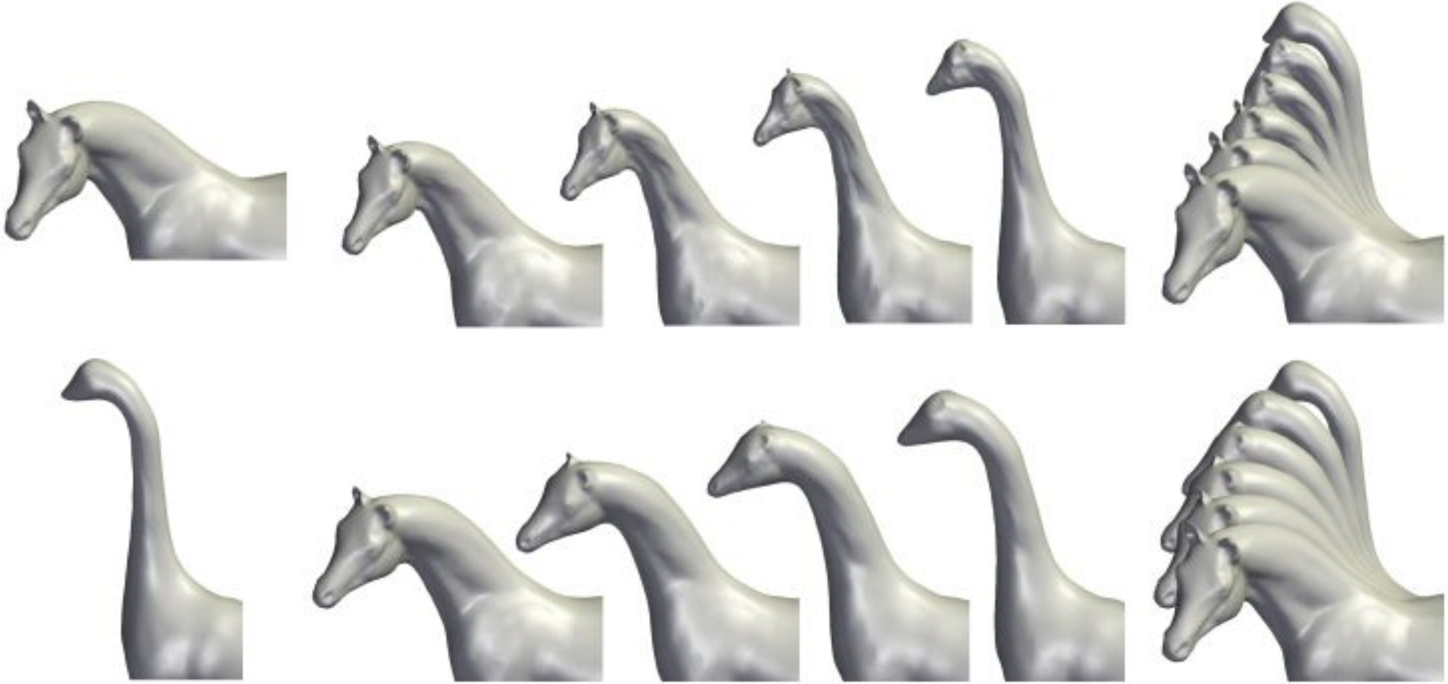

Xu, D., Zhang, H., Wang, Q., & Bao, H. (2005). Poisson shape interpolation. *Graph. Model.*, 68, 268-281.

Shape interpolation

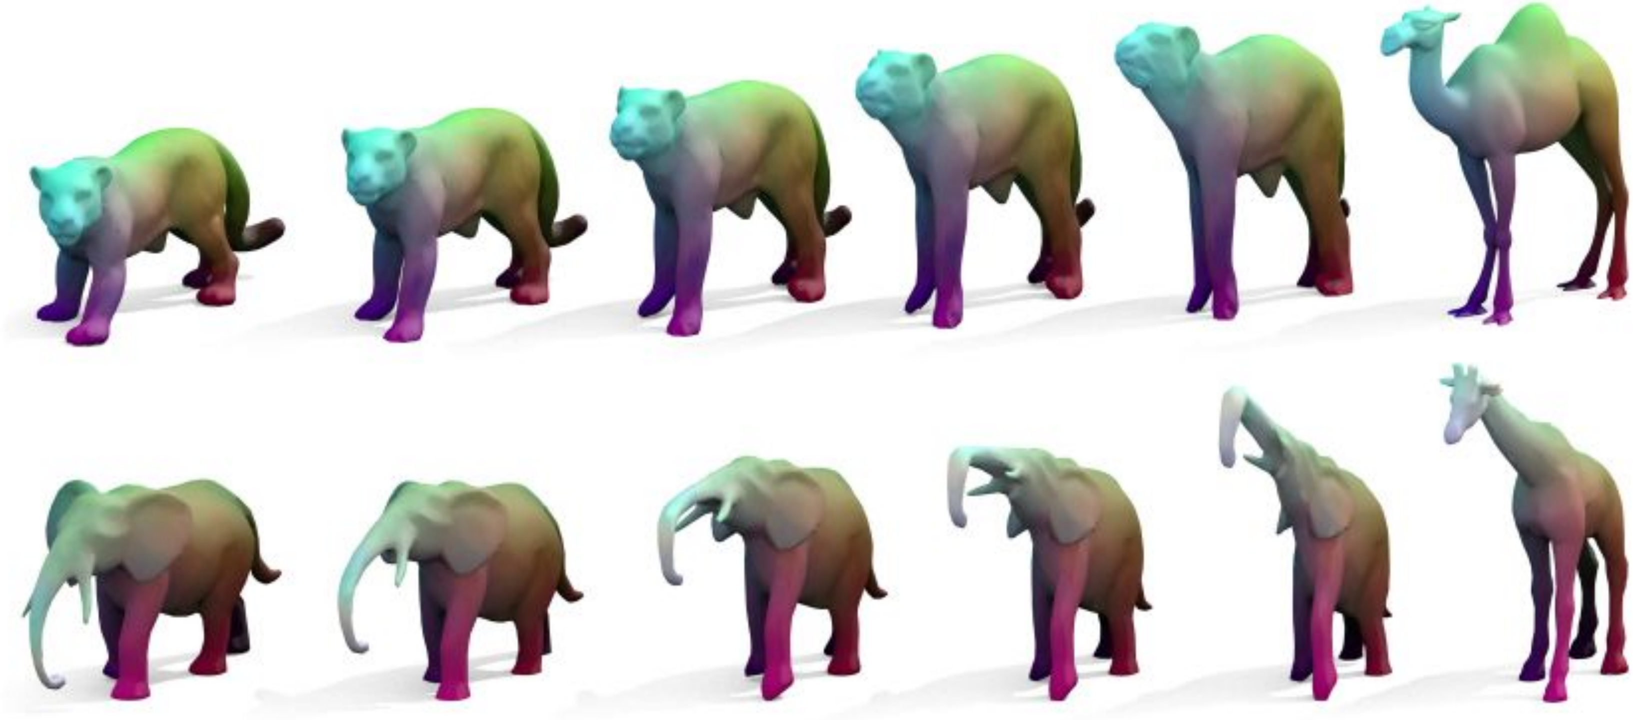

Shape interpolation

Geometric characterisation:

- Represent with low-dimensional manifolds;
- Interpolate shapes directly.

In both cases, find geodesic paths between the corresponding points - paths with minimum number of edges.

Statistical characterization (generative models):

- Occupancy probabilities on 3D voxel grid;
- Decode point clouds or 3D meshes;
- Implicit representation with a neural network.

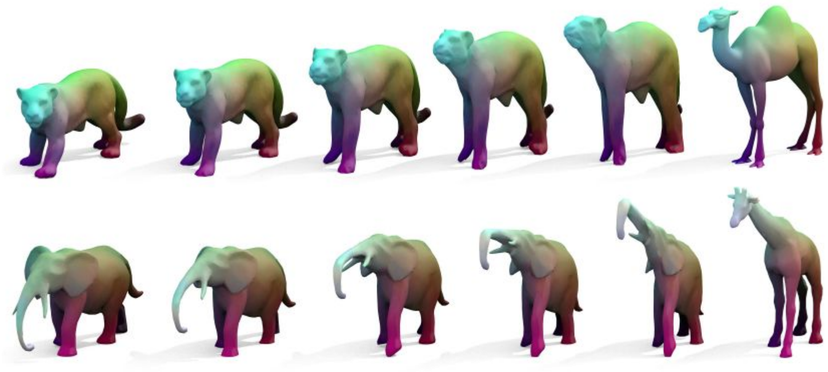

Geometrical characterization of 3D shapes

Low-dimensional manifold in a high-dimensional space

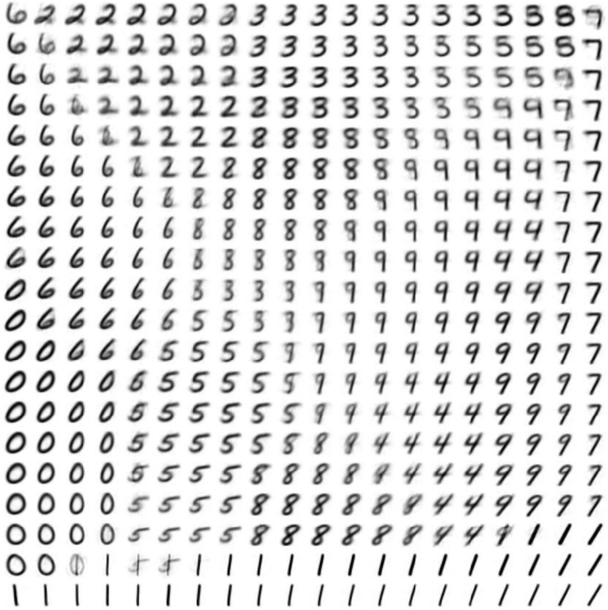

Fyffe, Graham. (2019). Closed Form Variances for Variational Auto-Encoders.

Shape interpolation

Geometric characterisation:

- Represent with low-dimensional manifolds;
- Interpolate shapes directly.

In both cases, find geodesic paths between the corresponding points - paths with minimum number of edges.

Statistical characterization (generative models):

- Occupancy probabilities on 3D voxel grid;
- Decode point clouds or 3D meshes;
- Implicit representation with a neural network.

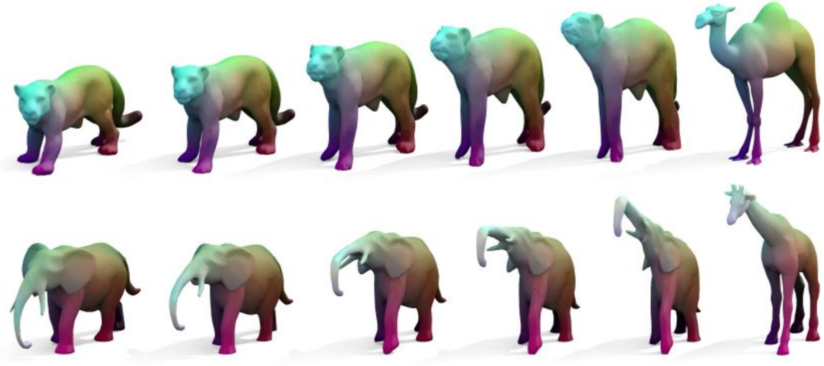

3D shape representation

3D voxel grid

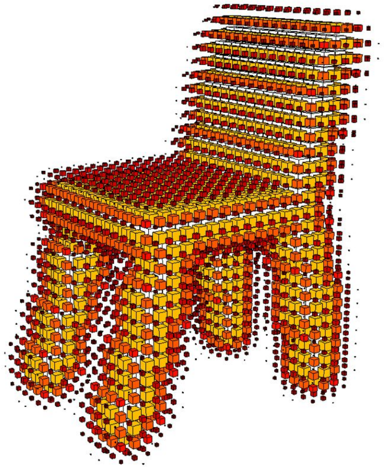

Kakillioglu, Burak & Ren, Ao & Wang, Yanzhi & Velipasalar, Senem. (2020). 3D Capsule Networks for Object Classification With Weight Pruning. IEEE Access. PP. 1-1. 10.1109/ACCESS.2020.2971950.

Point cloud

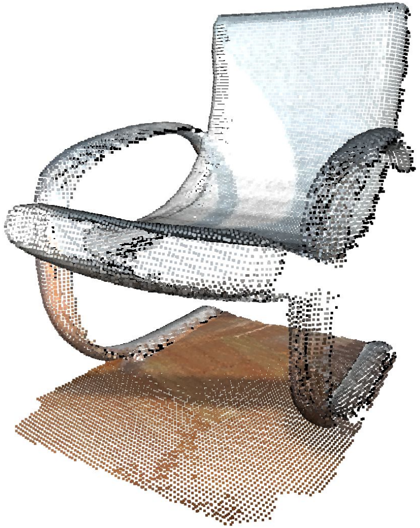

www.open3d.org

3D mesh

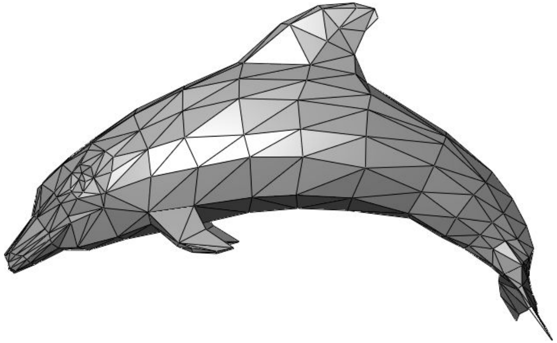

Implicit representation

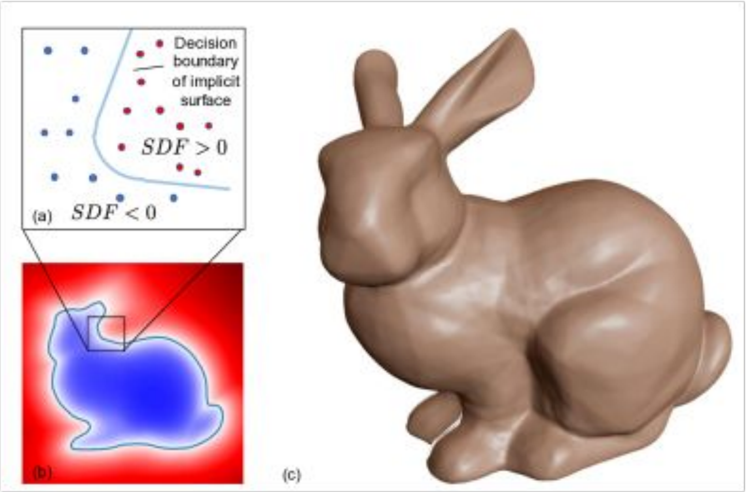

Jeong Joon Park, Peter Florence, Julian Straub, Richard Newcombe, Steven Lovegrove: “DeepSDF: Learning Continuous Signed Distance Functions for Shape Representation”, 2019

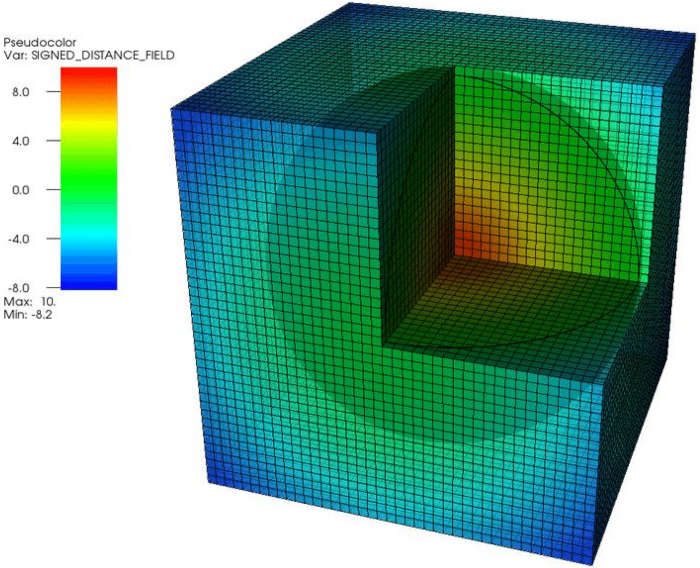

Shriwise, Patrick & Davis, Andrew & Jacobson, Lucas & Wilson, Paul. (2017). Particle Tracking Acceleration via Signed Distance Fields in DAGMC. Nuclear Engineering and Technology.

Signed distance field vs Velocity field

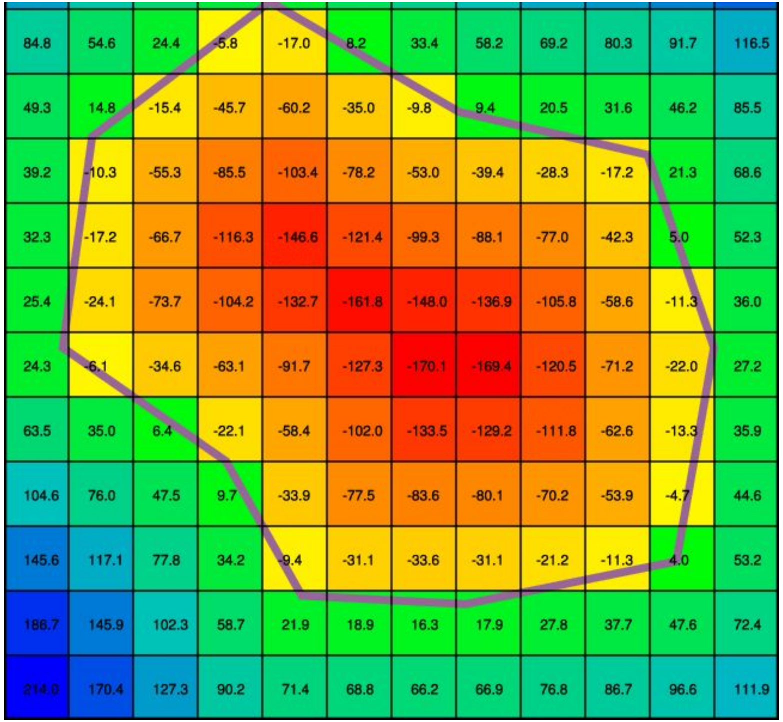

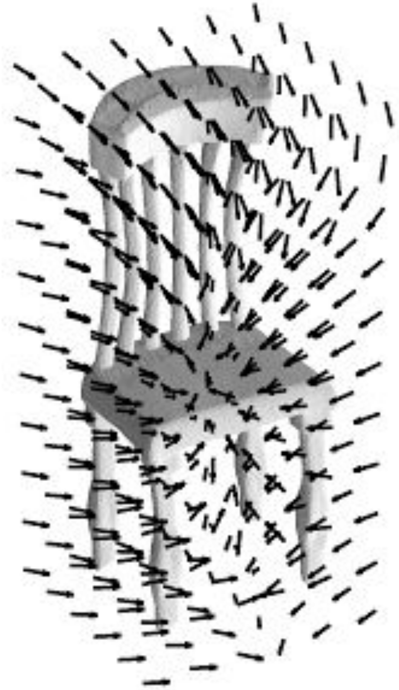

Point correspondence

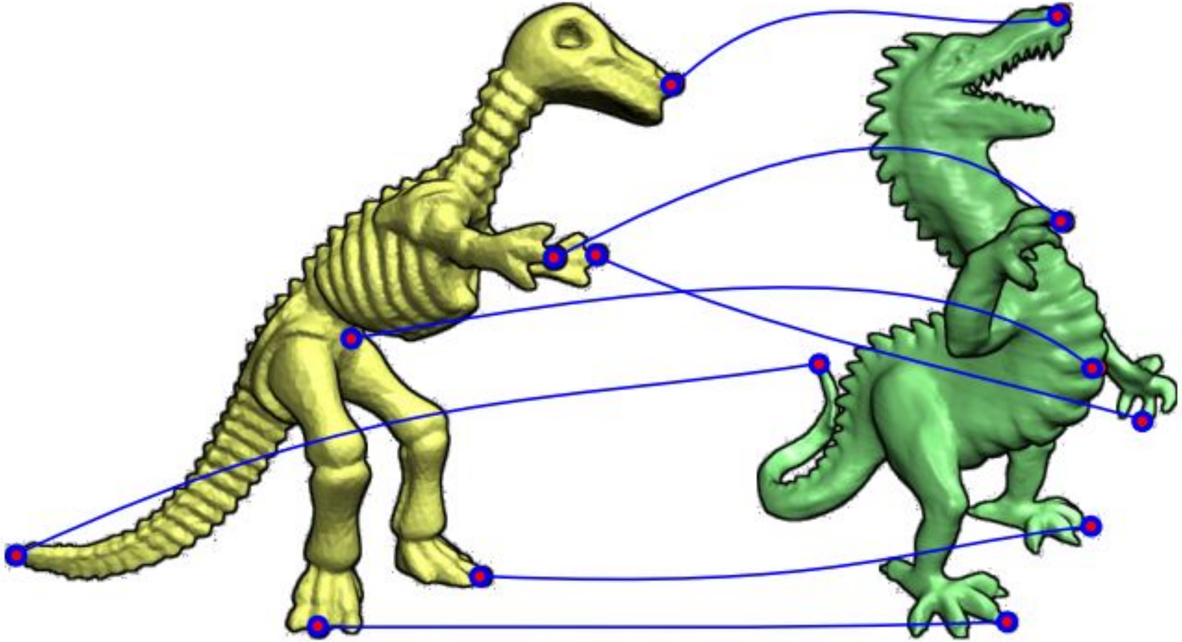

Oliver van Kaick, Hao Zhang, Ghassan Hamarneh, and Daniel Cohen-Or. A survey on shape correspondence. *Computer Graphics Forum*, 30(6):1681–1707, 2011.

Point correspondence

Point correspondence

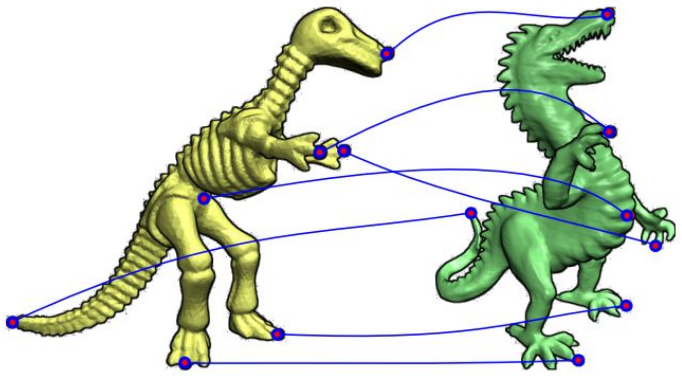

- Axiomatic
- Machine learning-based
- Manual

Limitations of previous works

- Point correspondence is solved separately
- Othen done manually
- => time-consuming data preparation
- => not enough training data
- => no inter-class interpolation

NeuroMorph

Relate two 3D shapes

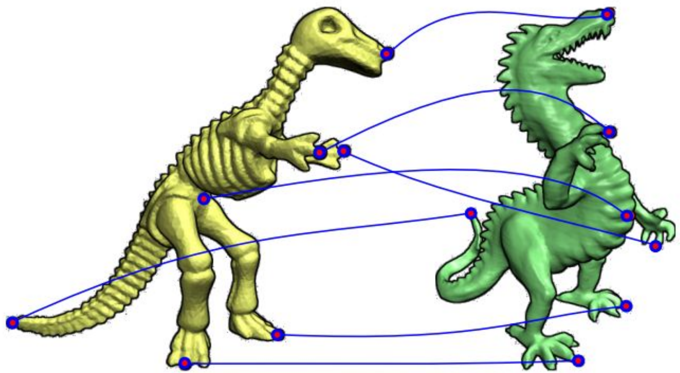

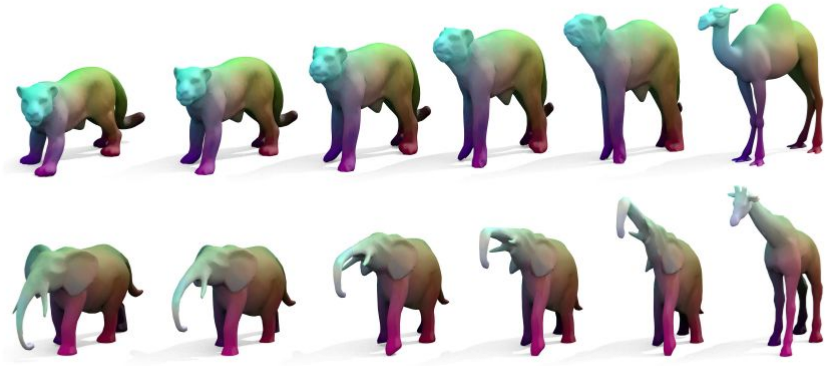

Relate two 3D shapes

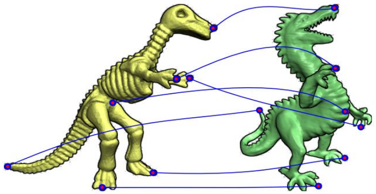

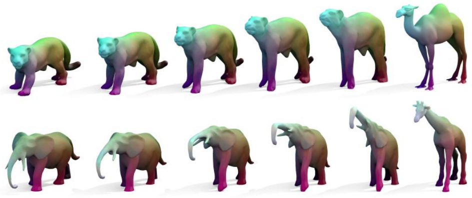

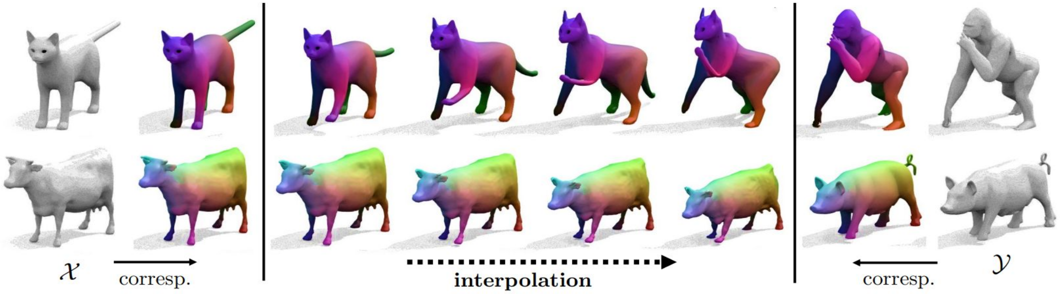

Single feed-forward pass and unsupervised

Relate two 3D shapes

Shape interpolation

Use geometric representation:

- Represent with low-dimensional manifolds;
- Interpolate shapes directly.

In both cases, find geodesic paths between the corresponding points - paths with minimum number of edges.

Statistical characterization (generative models):

- Occupancy probabilities on 3D voxel grid;
- Decode point clouds or 3D meshes;
- Implicit representation with a neural network.

Point correspondence

- Manual
- Axiomatic
- Machine learning-based

Goal

$$\mathcal{X} - source$$
$$\mathcal{Y} - target$$

$$f : (\mathcal{X}, \mathcal{Y}) \longmapsto (\Pi, \Delta)$$

$$\Pi - correspondence\ matrix$$
$$\Delta(t) - interpolation\ flow$$

Goal

$$\Pi \in [0, 1]^{n \times m}$$

$$\Pi_{i,j} = \mathbb{P}(p_i \leftrightarrow p_j)$$

$$\Delta(t) \in \mathbb{R}^{n \times 3}, t \in [0, 1]$$

$$\mathbf{X}(t) = \mathbf{X} + \Delta(t)$$

$$\mathbf{X}(0) = \mathbf{X}$$

$$\mathbf{X}(1) = \Pi\mathbf{Y}$$

Point Correspondence - Feature Extraction

$$\tilde{\mathbf{X}} = (\mathbf{X}, \mathbf{N})$$

$$\tilde{\mathbf{X}} \rightarrow \tilde{\mathbf{X}}' \rightarrow \tilde{\mathbf{X}}'' \rightarrow \ \ldots$$

$$\tilde{\mathbf{x}}_i' := \max_{j:(i,j)\in\mathcal{E}} h_\phi(\tilde{\mathbf{x}}_i, \tilde{\mathbf{x}}_j - \tilde{\mathbf{x}}_i)$$

$$\tilde{\mathbf{x}}_i'' := (\tilde{\mathbf{x}}_i', \max_{1\leq i\leq n} \tilde{\mathbf{x}}_i)$$

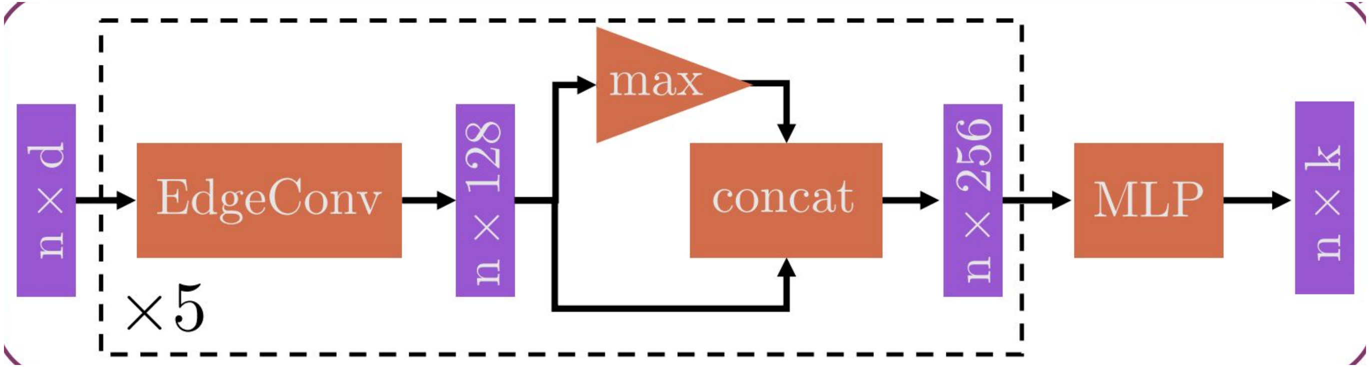

Point Correspondence - Pairwise Feature Comparison

$$\Pi \in [0,1]^{n \times m}$$

$$\Pi_{i,j} = \mathbb{P}(p_i \leftrightarrow p_j)$$

$$\tilde{\mathbf{X}} = \Phi(\mathcal{X}) \in \mathbb{R}^{n \times d}$$

$$\tilde{\mathbf{Y}} = \Phi(\mathcal{Y}) \in \mathbb{R}^{n \times d}$$

$$\Pi_{ij} := \frac{exp(\sigma s_{ij})}{\sum_{k=1}^{m} exp(\sigma s_{ik})} \quad s_{ij} := \frac{\langle \tilde{\mathbf{x}}_i, \tilde{\mathbf{y}}_j \rangle}{\|\tilde{\mathbf{x}}_i\|_2 \|\tilde{\mathbf{y}}_j\|_2}$$

Interpolation

$$\mathbf{X}(t) = \mathbf{X} + \Delta(t)$$
$$\mathbf{Z} := (\mathbf{X}, \Pi\mathbf{Y} - \mathbf{X}, \mathbf{1}t)$$
$$\mathbf{Z} \to \mathbf{Z}' \to \mathbf{Z}'' \ldots \to \mathbf{V} \in \mathbb{R}^{n \times 3}$$
$$\Delta(t) = t\mathbf{V}(t)$$

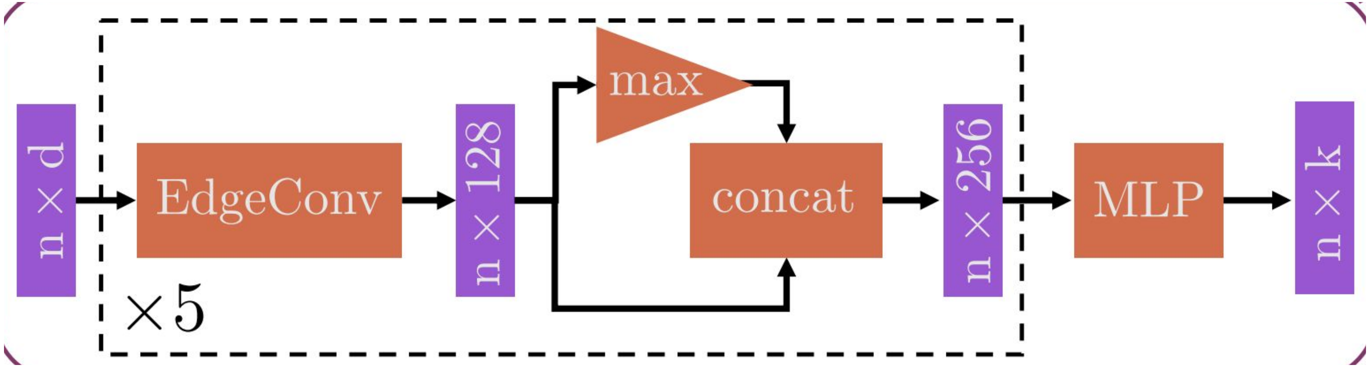

NeuroMorph - Full Architecture

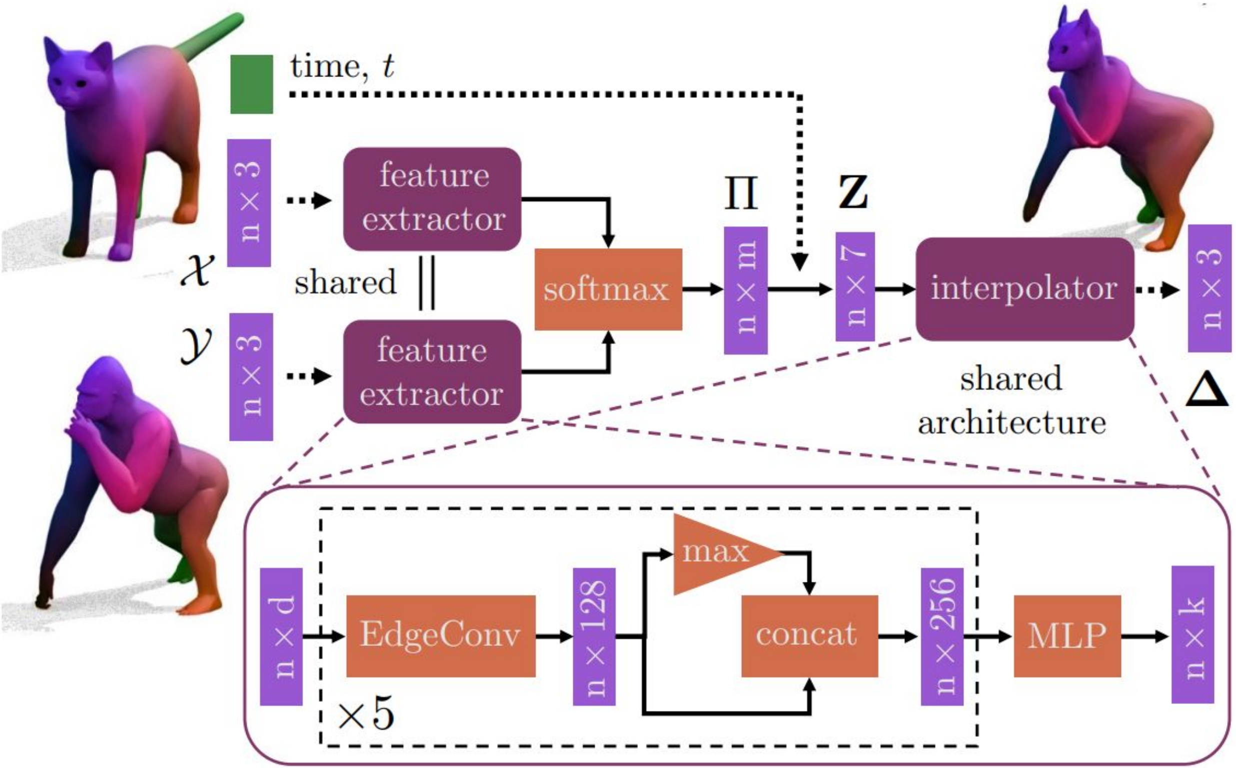

Interpolation - Trivial Solution

$$\mathbf{X}(0) = \mathbf{X} + 0 \cdot \mathbf{V}(0) = \mathbf{X}$$
$$\mathbf{X}(1) = \mathbf{X} + 1 \cdot (\Pi\mathbf{Y} - \mathbf{X}) = \Pi\mathbf{Y}$$

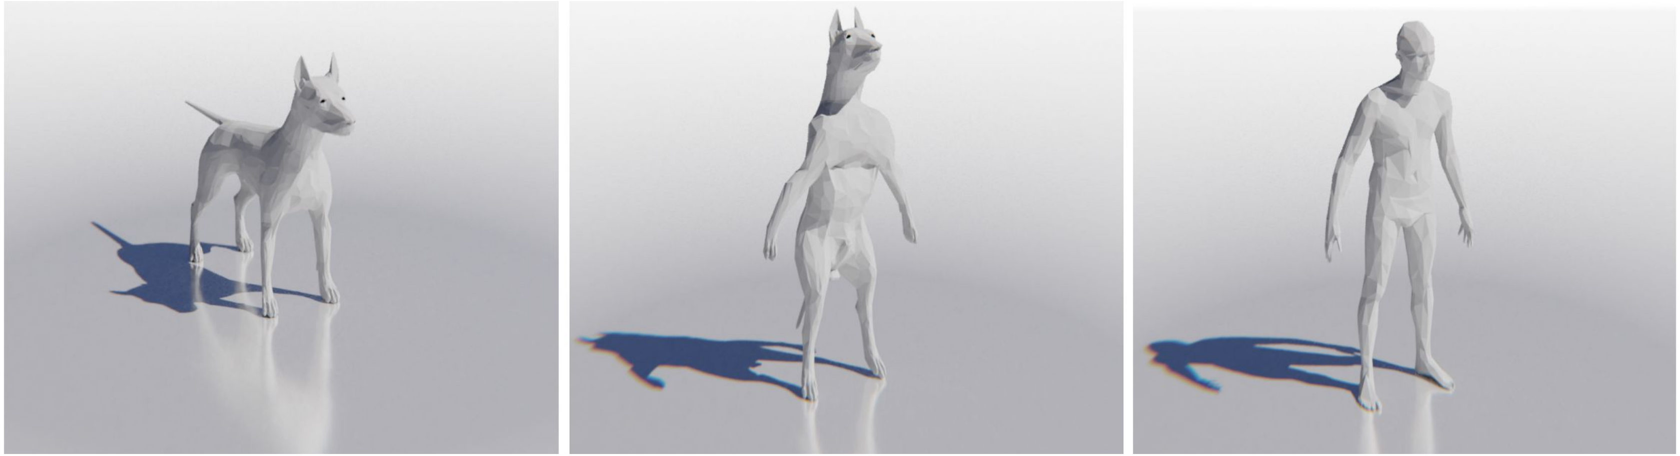

Learning

1. Correctly correspond and interpolate the Source to the Target
2. Keep intermediate models geometrically plausible

$$\ell := \lambda_{reg}\ell_{reg} + \lambda_{arap}\ell_{arap} + \lambda_{geo}\ell_{geo}$$

1. Registration loss
2. As-rigid-as-possible loss
3. Geodesic distance preservation loss

Learning - Registration Loss

$$\ell_{reg}(\mathbf{X}_T, \mathbf{T}, \Pi) := \|\Pi\mathbf{Y} - \mathbf{X}_T\|_2^2$$

$$\Delta(0) = 0$$

$$\Delta(1) = \Pi\mathbf{Y} - \mathbf{X}$$

Learning - As-Rigid-As-Possible Loss

$$E_{arap}(\mathbf{X}_k, \mathbf{X}_{k+1}) := \frac{1}{2} \min_{\substack{\mathbf{R}_i \in SO(3) \\ i=1,\dots,n}} \sum_{(i,j)\in\mathcal{E}} \|\mathbf{R}_i(\mathbf{X}_{k,j} - \mathbf{X}_{k,i}) - (\mathbf{X}_{k+1,j} - \mathbf{X}_{k+1,i})\|_2^2$$

$$\ell_{arap}(\mathbf{X}_0 \dots, \mathbf{X}_T) := \sum_{k=0}^{T-1} E_{arap}(\mathbf{X}_k, \mathbf{X}_{k+1}) + E_{arap}(\mathbf{X}_{k+1}, \mathbf{X}_k)$$

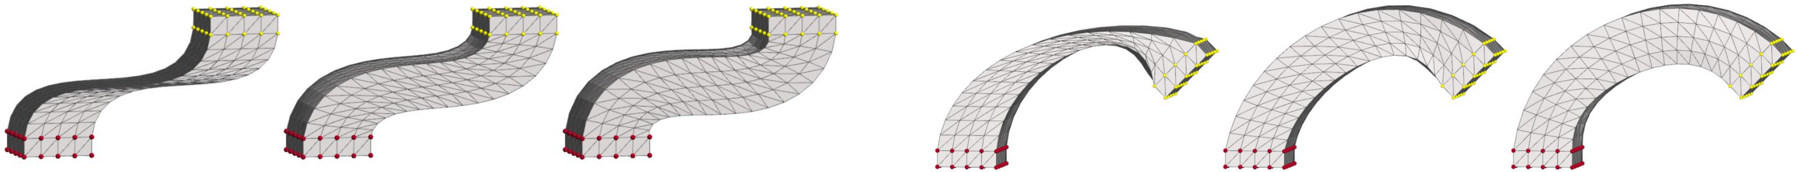

Olga Sorkine and Marc Alexa. 2007. As-rigid-as-possible surface modeling. In Proceedings of the fifth Eurographics symposium on Geometry processing (SGP '07). Eurographics Association, Goslar, DEU, 109–116.

Learning - Geodesic Distance Preservation Loss

$$\ell_{geo}(\Pi) := \|\Pi \mathbf{D}_{\mathcal{Y}} \Pi^T - \mathbf{D}_{\mathcal{X}}\|_2^2$$

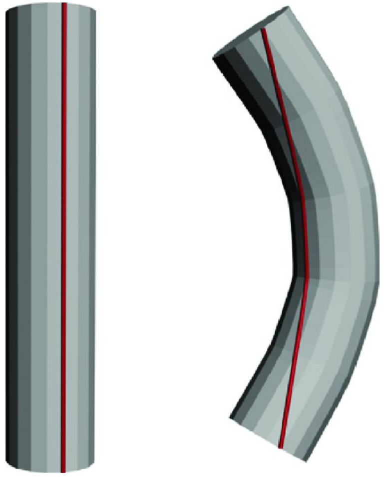

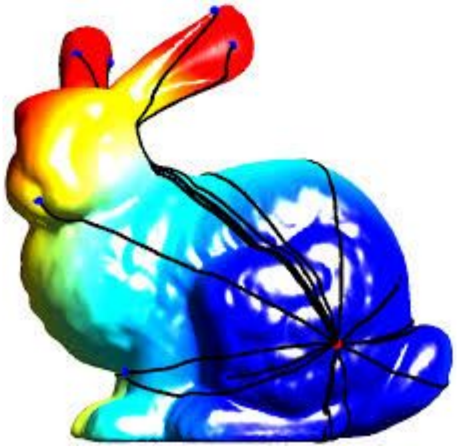

Mykhalchuk, Vasyl & Cordier, Frederic & Seo, Hyewon. (2013). Landmark transfer with minimal graph. Computers & Graphics. 37. 539–552. 10.1016/j.cag.2013.04.005.

Gabriel Peyré, Laurent D. Cohen. Geodesic Methods for Shape and Surface Processing. Tavares, João Manuel R.S.; Jorge, R.M. Natal. Advances in Computational Vision and Medical Image Processing: Methods and Applications, Springer Verlag, pp.29-56, 2009, Computational Methods in Applied Sciences, Vol. 13, ff10.1007/978-1-4020-9086-8ff. ffhal-00365899

Experiments and Results

Experiments - Point Correspondence

Datasets:

1. FAUST (remeshed)
2. SCHREC20
3. G-S-H (Galgo, Sphynx, Human)

Metric:

Princeton benchmark protocol (geodesic distance normalized by square root area of the mesh).

Competitors:

1. BCICP
2. ZoomOut
3. Smooth Shells
4. 3D-CODED
5. FMNet
6. GeoFMNet
7. SurFMNet
8. Unsupervised FMNet
9. Weakly supervised FMNet
10. Deep shells

Experiments - Point Correspondence - Results

		err	p.p.	w/o p.p.
Axiomatic	*BCICP*	*6.4*	*-*	*-*
	ZoomOut	*6.1*	*-*	*-*
	Smooth Shells	*2.5*	*-*	*-*
Supervised	*3D-CODED*	*2.5*	*-*	*-*
	FMNet	*5.9*	*PMF*	*11*
	GeoFMNet	*1.9*	*ZO*	*3.1*
Unsupervised	*SurFMet*	*7.4*	*ICP*	*15*
	Unzip FMNet	*5.7*	*PMF*	*10*
	Weakly sup. FMNet	*1.9*	*ZO*	*3.3*
	Deep shells	*1.7*	*-*	*-*
	NeuroMorph	***1.5***	***SL***	***2.3***

FAUST

Experiments - Point Correspondence - Results

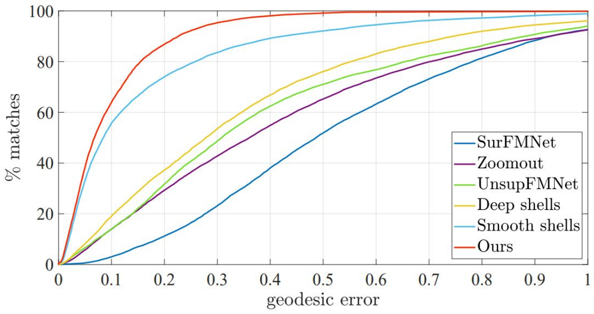

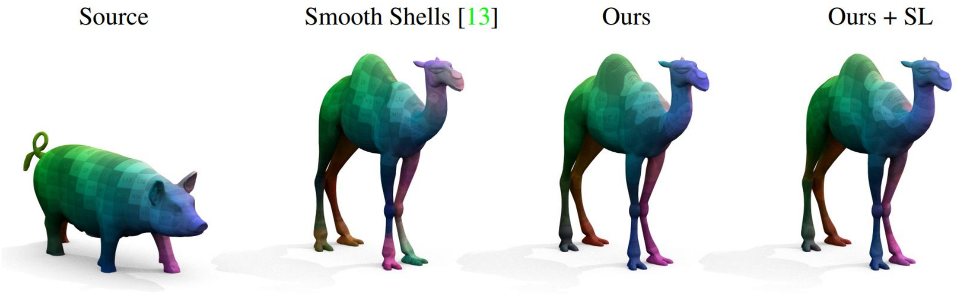

SHREC20

Experiments - Point Correspondence - Results

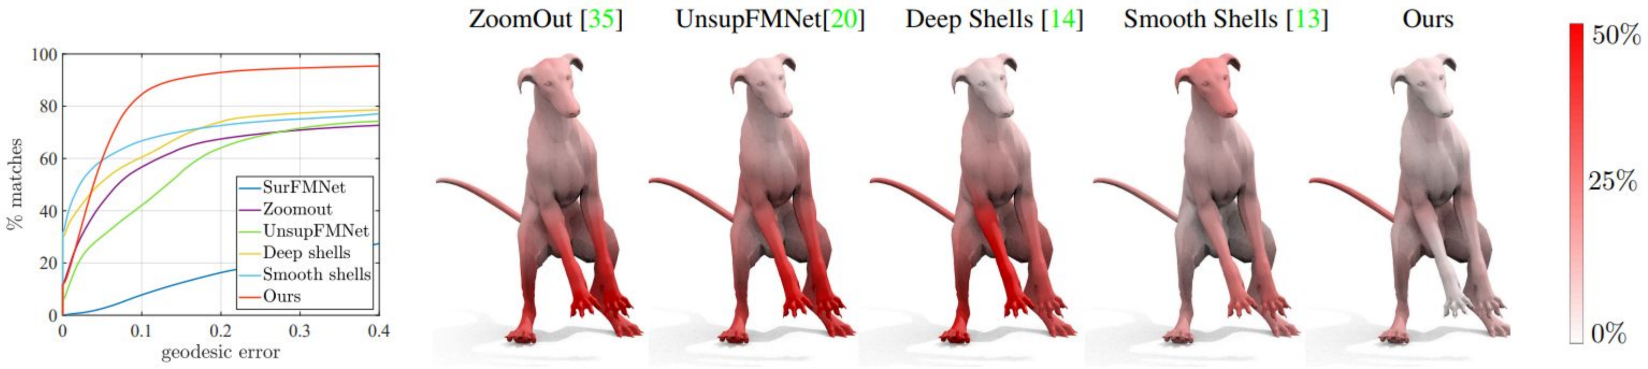

G-S-H

Experiments - Shape Interpolation

Datasets:

1. FAUST
2. MANO

Metric:

1. Conformal distortion
2. Chamfer distance

Competitors:

1. ShapeFlow
2. LIMP
3. Hamiltonian interpolation

Experiments - Shape Interpolation - Metrics

$$F(\mathcal{X}) = A\mathcal{X} + b$$

$$\kappa_F(A) = \frac{trace(A^T A)}{\det A}$$

Conformal distortion

$$CD(S_1, S_2) = \frac{1}{S_1} \sum_{x \in S_1} \min_{y \in S_2} \|x - y\|_2^2 + \frac{1}{S_2} \sum_{y \in S_2} \min_{x \in S_1} \|y - x\|_2^2,$$

Chamfer distance

Experiments - Shape Interpolation - Results

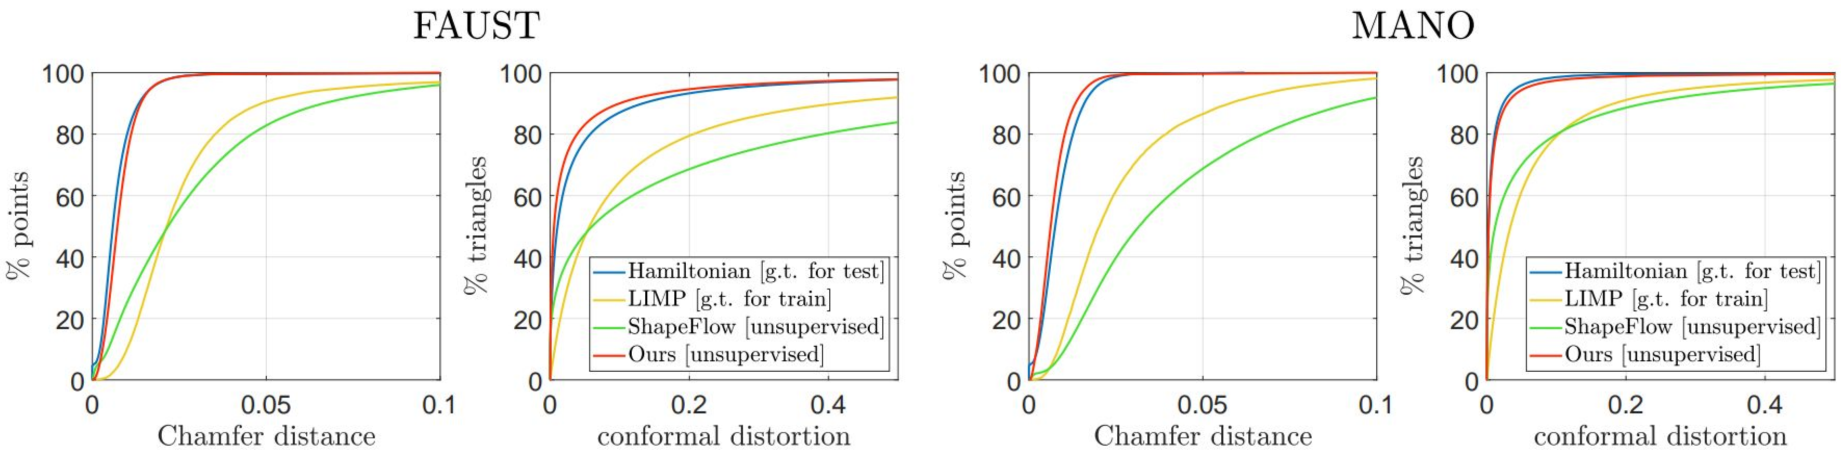

Experiments - Shape Interpolation - Results

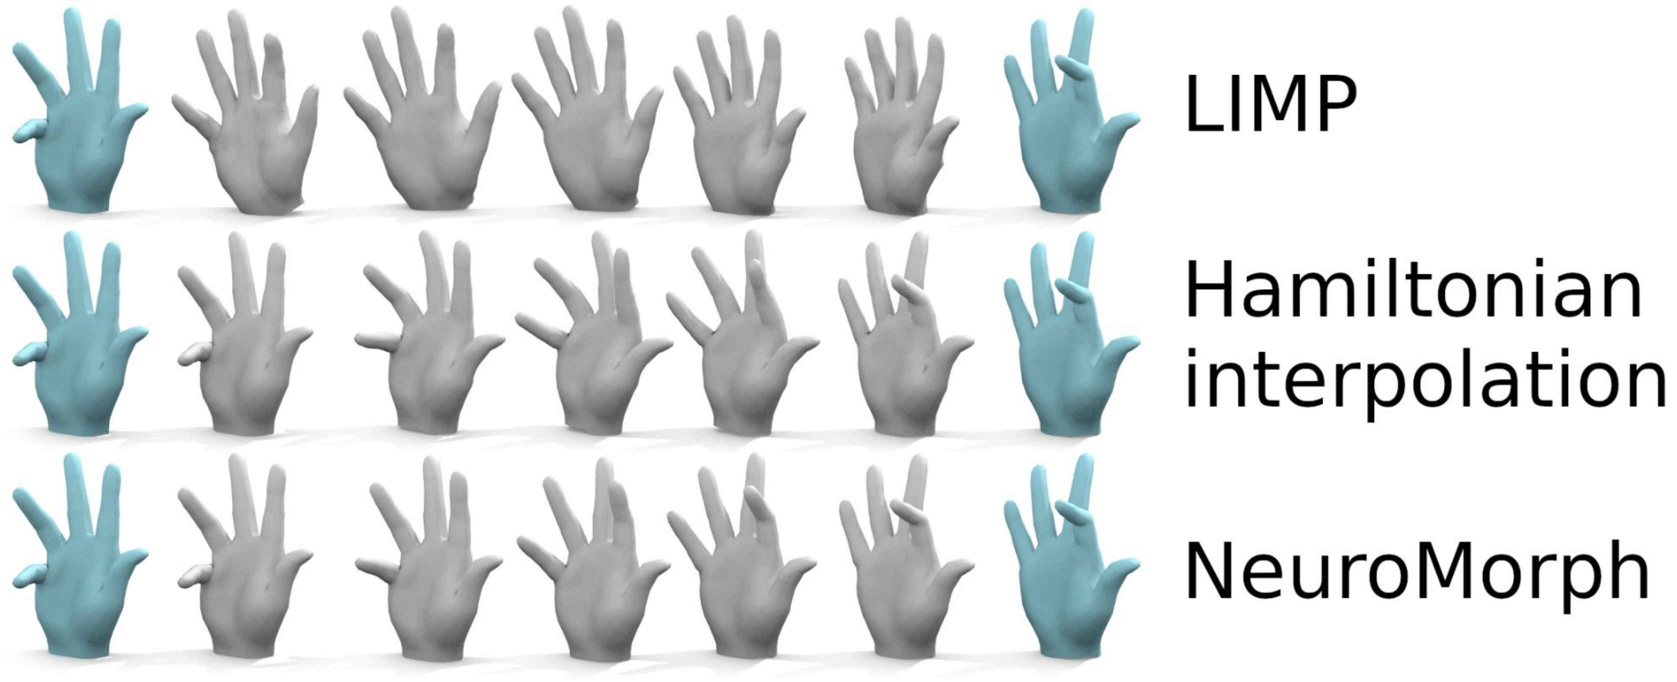

Application - Data Augmentation