Encoding Shapes and their Differences with Functional Maps

MS9 - Shape processing,
Curves and Surfaces, July 3rd, 2018

Maks Ovsjanikov

Joint with: *E. Corman, O. Azencot, R. Rustamov, M. Ben-Chen, R. Huang, L. Guibas…*

Shape Comparison

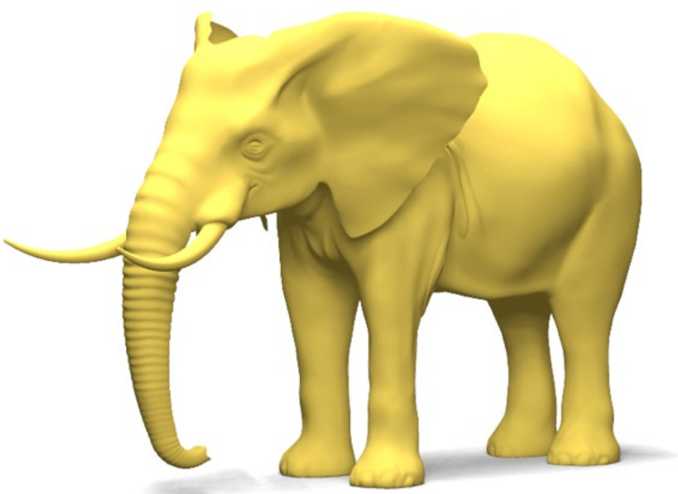

$\stackrel{?}{=}$

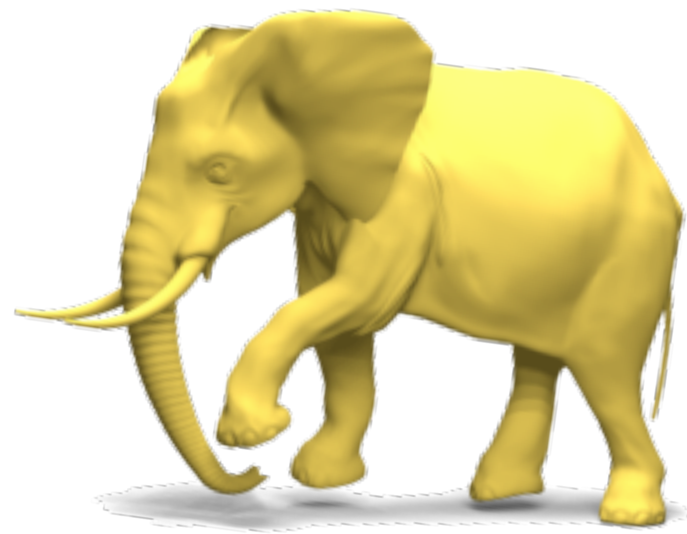

Given a pair of 3D shapes, quantify if they are *similar*.

Shape **Comparison**

Given *a set* of 3D shapes, quantify if they are similar.
If not, find regions of dissimiliarity.

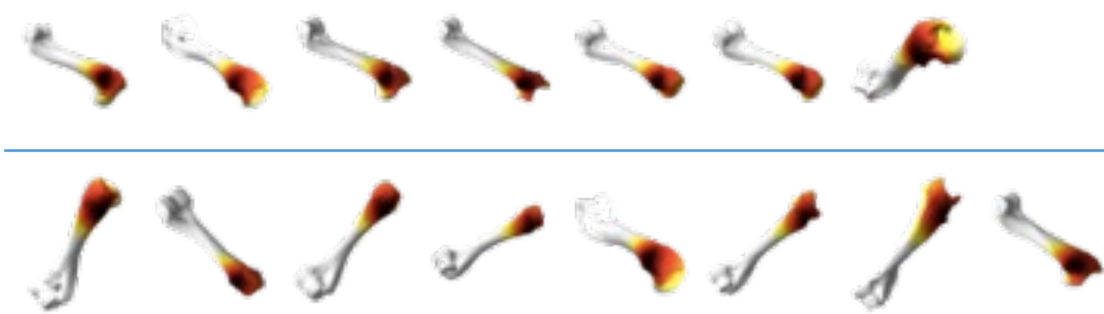

Sanglier
Chambord

Sanglier
Compiègne

Data: Anthony Herrel, Raphaël Cornette,
Muséum National d'Histoire Naturelle, CNRS

What is a Shape?

- Continuous: a surface embedded in 3D.
- Discrete: a graph embedded in 3D (triangle mesh).

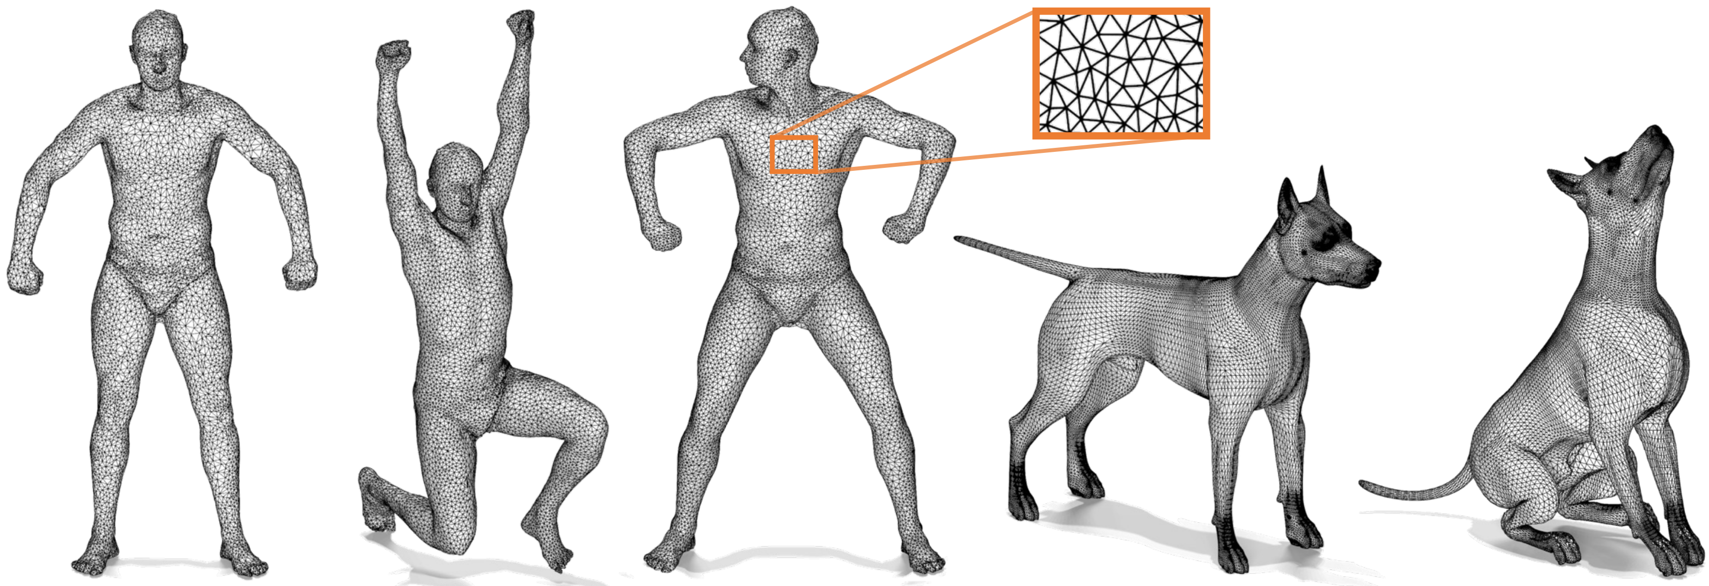

5k – 200k triangles

Shapes from the SCAPE, TOSCA and FAUST datasets

Today

- Encoding differences between shapes.
- Recovering shapes from functional operators.
- Unbiased (base shape-free) shape comparison.

Main observation:

- Can fully encode *shape spaces* with functional maps without assuming fixed connectivity.

Background: Functional Maps

Rather than comparing *points* on objects it is often easier to compare *real-valued functions* defined on them[1,2].

[1] *Functional Maps: A Flexible Representation of Maps Between Shapes,* O., Ben-Chen, Solomon, Butscher. Guibas, SIGGRAPH 2012

[2] *Computing and Processing Correspondences with Functional Maps,* O. et al., SIGGRAPH Courses 2017

Background: Functional Maps

Rather than comparing *points* on objects it is often easier to compare *real-valued functions* defined on them. Such maps can be represented as matrices.

[1] *Functional Maps: A Flexible Representation of Maps Between Shapes,* O., Ben-Chen, Solomon, Butscher. Guibas, SIGGRAPH 2012

[2] *Computing and Processing Correspondences with Functional Maps,* O. et al., SIGGRAPH Courses 2017

Background: Functional Maps

Computing functional maps is often *much* easier (reduces to least squares) than point-to-point maps.

Can think of a functional map as an matrix of size $n_{V_2} \times n_{V_1}$, or of size, $k_2 \times k_1$, in a reduced basis.

Computing and Processing Correspondences with Functional Maps, O. et al., SIGGRAPH Courses 2017

Problem Setup

- Given a pair of shapes and a *functional* map between them, detect similarities and *differences* (distortion) across them.

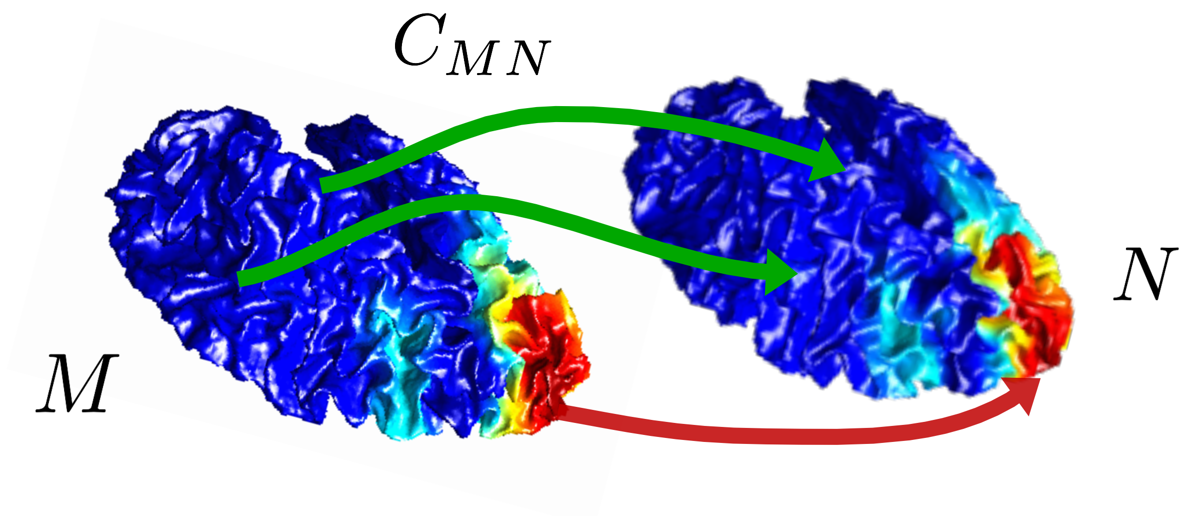

- Do it in a *multi-scale* way (not be sensitive to *local* changes).
- Accommodate approximate *soft* (functional) maps

Map-Based Exploration of Intrinsic Shape Differences and Variability, *Rustamov, O., Azencot, Ben-Chen, Chazal, Guibas*, SIGGRAPH 2013

Shape Differences Definition

Given a functional map $C_{MN} : \mathcal{F}(M) \to \mathcal{F}(N)$
and a choice of functional inner products: $\langle \cdot, \cdot \rangle_M, \langle \cdot, \cdot \rangle_N$

Define *a shape difference operator* as linear operator D, s.t.

$$\langle f, D(g) \rangle_M = \langle C_{MN}(f), C_{MN}(g) \rangle_N \;\; \forall f, g$$

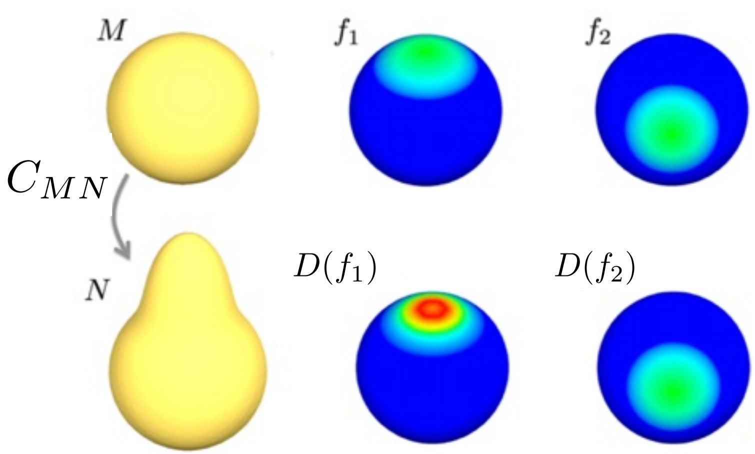

Shape Differences Definition

Given a functional map $C_{MN} : \mathcal{F}(M) \to \mathcal{F}(N)$
and a choice of functional inner products: $\langle \cdot, \cdot \rangle_M, \langle \cdot, \cdot \rangle_N$

Define *a shape difference operator* as linear operator *D*, s.t.

$$\langle f, D(g) \rangle_M = \langle C_{MN}(f), C_{MN}(g) \rangle_N \;\; \forall f, g$$

Existence and uniqueness of *D* is guaranteed by the Riesz representation theorem.

Map-Based Exploration of Intrinsic Shape Differences and Variability, *Rustamov, O., Azencot, Ben-Chen, Chazal, Guibas,* SIGGRAPH 2013

Shape Differences Definition

Given a functional map $C_{MN} : \mathcal{F}(M) \to \mathcal{F}(N)$
and a choice of functional inner products: $\langle \cdot, \cdot \rangle_M , \langle \cdot, \cdot \rangle_N$

Define *a shape difference operator* as linear operator *D*, s.t.

$$\langle f, D(g) \rangle_M = \langle C_{MN}(f), C_{MN}(g) \rangle_N \;\; \forall f, g$$

We let *V* and *R*, be operators associated with L_2 and H_1 inner products:

$$V :< f, g >_{L_2} = \int f(x) g(x) d\mu$$

$$R :< f, g >_{H_1} = \int \langle \nabla f(x), \nabla g(x) \rangle \, d\mu$$

Shape Differences Definition

Given a functional map $C_{MN} : \mathcal{F}(M) \to \mathcal{F}(N)$
and a choice of functional inner products: $\langle \cdot, \cdot \rangle_M, \langle \cdot, \cdot \rangle_N$

Define *a shape difference operator* as linear operator D, s.t.

$$\langle f, D(g) \rangle_M = \langle C_{MN}(f), C_{MN}(g) \rangle_N \;\; \forall f, g$$

We let V and R, be operators associated with L_2 and H_1 inner products:

$$V :< f, g >_{L_2} = \int f(x) g(x) d\mu$$

$$R :< f, g >_{H_1} = \int \langle \nabla f(x), \nabla g(x) \rangle \, d\mu = < f, \Delta g >_{L_2}$$

Shape Differences Definition

Given a functional map $C_{MN} : \mathcal{F}(M) \to \mathcal{F}(N)$
and a choice of functional inner products: $\langle \cdot, \cdot \rangle_M, \langle \cdot, \cdot \rangle_N$

Define *a shape difference operator* as linear operator D, s.t.

$$\langle f, D(g) \rangle_M = \langle C_{MN}(f), C_{MN}(g) \rangle_N \;\; \forall f, g$$

We let V and R, be operators associated with L_2 and H_1 inner products. In the discrete setting, reduces to simply matrix transposes and inverses:

$$< f, g >_{L_2} = f^T A g$$

$$< f, g >_{H_1} = f^T L g$$

Shape Differences Properties

Theorem:

If C_{MN} comes from a point to point map, then:

1) $V = Id$ if and only if the map is *area-preserving*.

2) $R = Id$ if and only if the map is *conformal*.

1) $\langle f, g \rangle_{L_2(M)} = \langle C_{MN}(f), C_{MN}(g) \rangle_{L_2(N)} \quad \forall f, g$

2) $\langle f, g \rangle_{H_1(M)} = \langle C_{MN}(f), C_{MN}(g) \rangle_{H_1(N)} \quad \forall f, g$

Shape Differences in Collections

- Since shape differences $D_{M,N1}, D_{M,N2}$ are operators with the same domain/range, we can *compare distortion* on multiple shapes.

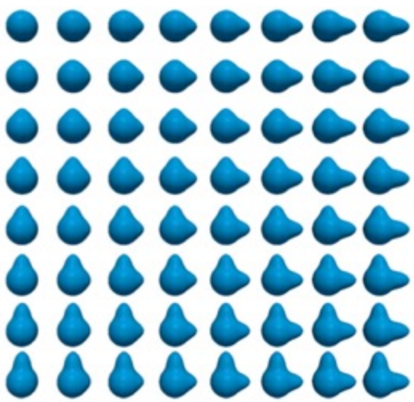

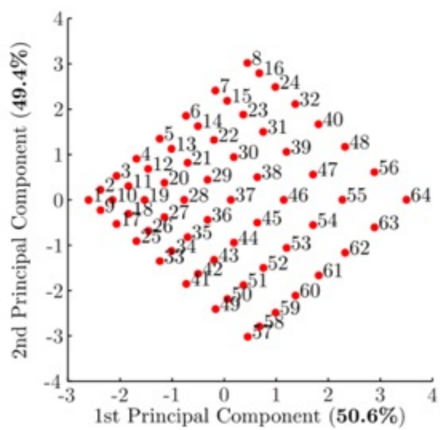

PCA on area-based shape differences

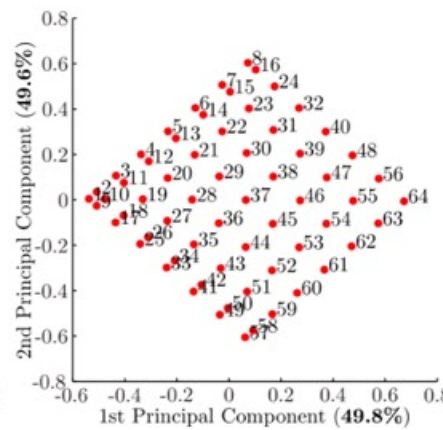

PCA on conformal shape differences

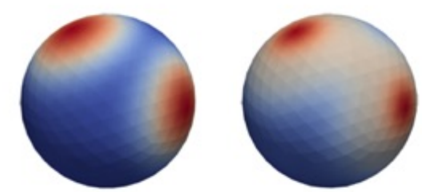

Variability localization for area

Variability localization for conformal

Map-Based Exploration of Intrinsic Shape Differences and Variability, *Rustamov, O., Azencot, Ben-Chen, Chazal, Guibas*, SIGGRAPH 2013

Shape Differences in Collections

- Since shape differences $D_{M,N1}, D_{M,N2}$ are operators with the same domain/range, we can *compare distortion* on multiple shapes.

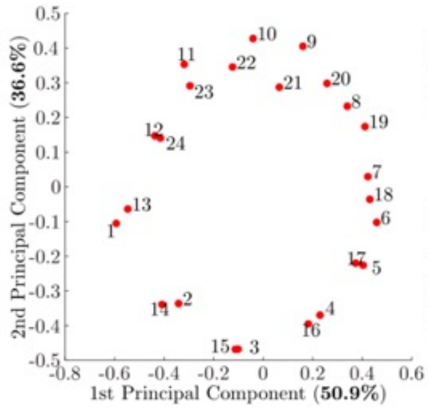

PCA on area-based shape differences

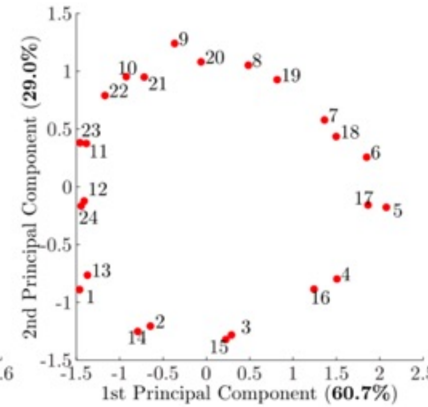

PCA on conformal shape differences

Variability localization for area

Variability localization for conformal

Map-Based Exploration of Intrinsic Shape Differences and Variability, *Rustamov, O., Azencot, Ben-Chen, Chazal, Guibas*, SIGGRAPH 2013

Can we recover the metric?

Given a base shape M and two shape difference operators, can we recover the target shape?

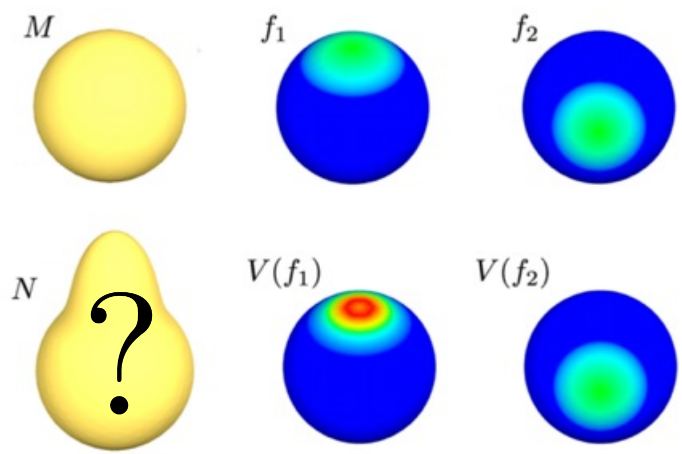

Functional Characterization of Intrinsic and Extrinsic Geometry
Corman, Solomon, Ben-Chen, Guibas, O. TOG 2017

Can we recover the metric?

Given a base shape M and two shape difference operators, can we recover the target shape?

Possible limitation:

Shape difference operators are blind to *isometric deformations*.

$$\langle f, D(g) \rangle_M = \langle C_{MN}(f), C_{MN}(g) \rangle_N \;\; \forall f, g$$

If C_{MN} preserves inner products, then $D = Id$.

Functional Characterization of Intrinsic and Extrinsic Geometry
Corman, Solomon, Ben-Chen, Guibas, O. TOG 2017

Can we recover the metric?

Given a base shape M and two shape difference operators, can we recover the target shape?

Possible limitation:

Shape difference operators are blind to isometric deformations.

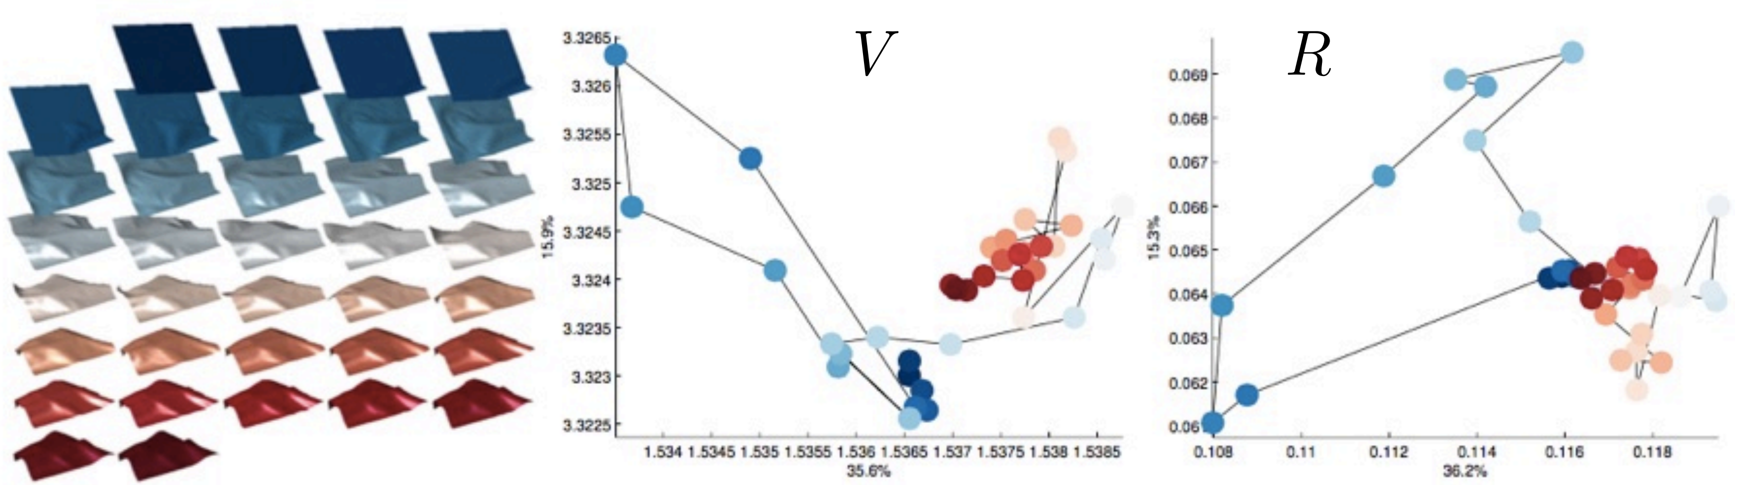

Can we recover the metric?

Given a base shape M and two shape difference operators, can we recover the target shape?

Possible limitation:
Shape difference operators are blind to isometric deformations.

Best hope:
Recover the metric and solve for the pose.

Functional Characterization of Intrinsic and Extrinsic Geometry
Corman, Solomon, Ben-Chen, Guibas, O. TOG 2017

A metric on the triangle mesh

From inner products to the metric on a triangle mesh:

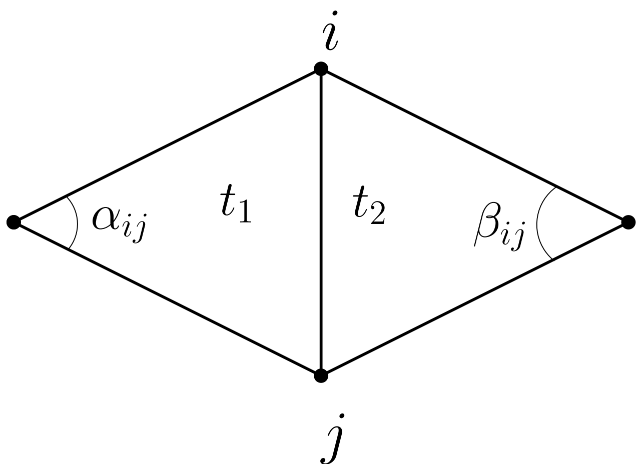

$$L_{ij} = < \nabla e_i, \nabla e_j >$$
$$= \frac{1}{2}\cot(\alpha_{ij}) + \frac{1}{2}\cot(\beta_{ij})$$

Given the inner product between every pair of functions can we recover the metric? **Probably[1,2]**

When the information is exact

[1] Zeng et al. *Discrete heat kernel determines discrete Riemannian metric*. Graph. Models , 2012
[2] De Goes et al. *Weighted triangulations for geometry processing*, TOG, 2014

A metric on the triangle mesh

From metric to inner products on a triangle mesh:

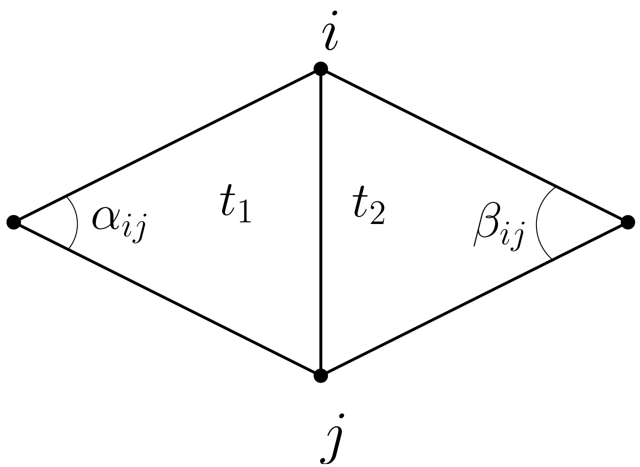

$$L_{ij} = < \nabla e_i, \nabla e_j >$$
$$= \frac{1}{2}\cot(\alpha_{ij}) + \frac{1}{2}\cot(\beta_{ij})$$

Given the Laplacian of a shape can we recover the metric?

- *What if it is known approximately?*
- *Using Shape Difference Operators?*

Zeng et al. *Discrete heat kernel determines discrete Riemannian metric*. Graph. Models , 2012
De Goes et al. *Weighted triangulations for geometry processing*, TOG, 2014

Recovering the metric

From metric to inner products on a triangle mesh:

Theorem:

Given the two shape difference operators, the discrete metric can be recovered by solving 2 linear systems that are ``almost always" full-rank.

Propose convex regularization, for noisy/underconstrained systems.

Functional Characterization of Intrinsic and Extrinsic Geometry
Corman, Solomon, Ben-Chen, Guibas, O. TOG 2017

Recovering the shape

With only the edge-lengths, there are multiple near-isometries. Recovering the exact pose is hard.

Functional Characterization of Intrinsic and Extrinsic Geometry
Corman, Solomon, Ben-Chen, Guibas, O. TOG 2017

Extrinsic Information

Can we add additional extrinsic information? Encode the *second fundamental form*?

One Option:

Use dihedral angles to represent encode principal curvatures.

Difficulty:

Angle-based values are both unstable and difficult to recover in the presence of noise.

Second Fundamental Form is a *quadratic form,* not an angle.

Extrinsic Information

Can we add additional extrinsic information? Encode the second fundamental form?

Main idea : offset surfaces.

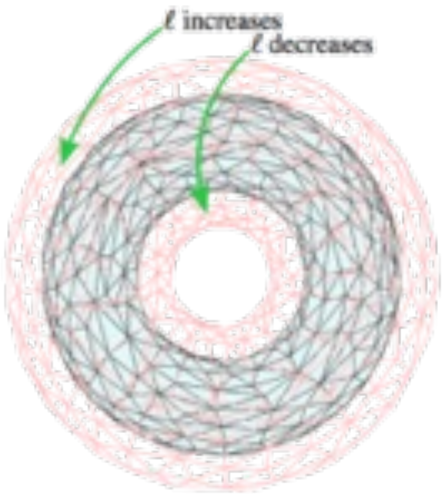

Edge-lengths change according to curvature of the offset surface.

Given a family of immersions, where each point follows the outward normal direction:

$$\left.\frac{\partial g}{\partial t}\right|_{t=0} = 2h|_{t=0} \text{ and } \left.\frac{\partial \mu}{\partial t}\right|_{t=0} = H\mu,$$

g : Metric (first fundamental form)

h : Second fundamental form

μ : Local area

H : Mean curvature

Shape Differences Based on Offset Surfaces

Given two shapes, compute four difference operators: two between the shapes, and two between their offsets.

V_{M^o,N^o}

R_{M^o,N^o}

$V_{M,N}$

$R_{M,N}$

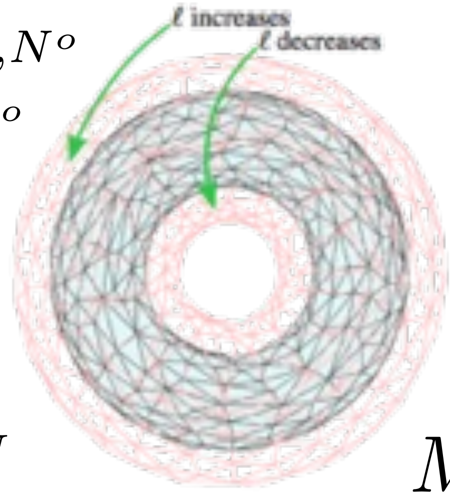

M

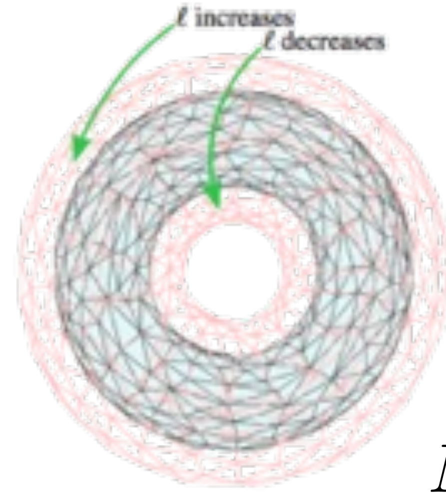

N

$V_{M,N}, R_{M,N}$ encode change in metric,
V_{M^o,N^o}, R_{M^o,N^o} encode change in curvature

Exploring shapes with extrinsic information

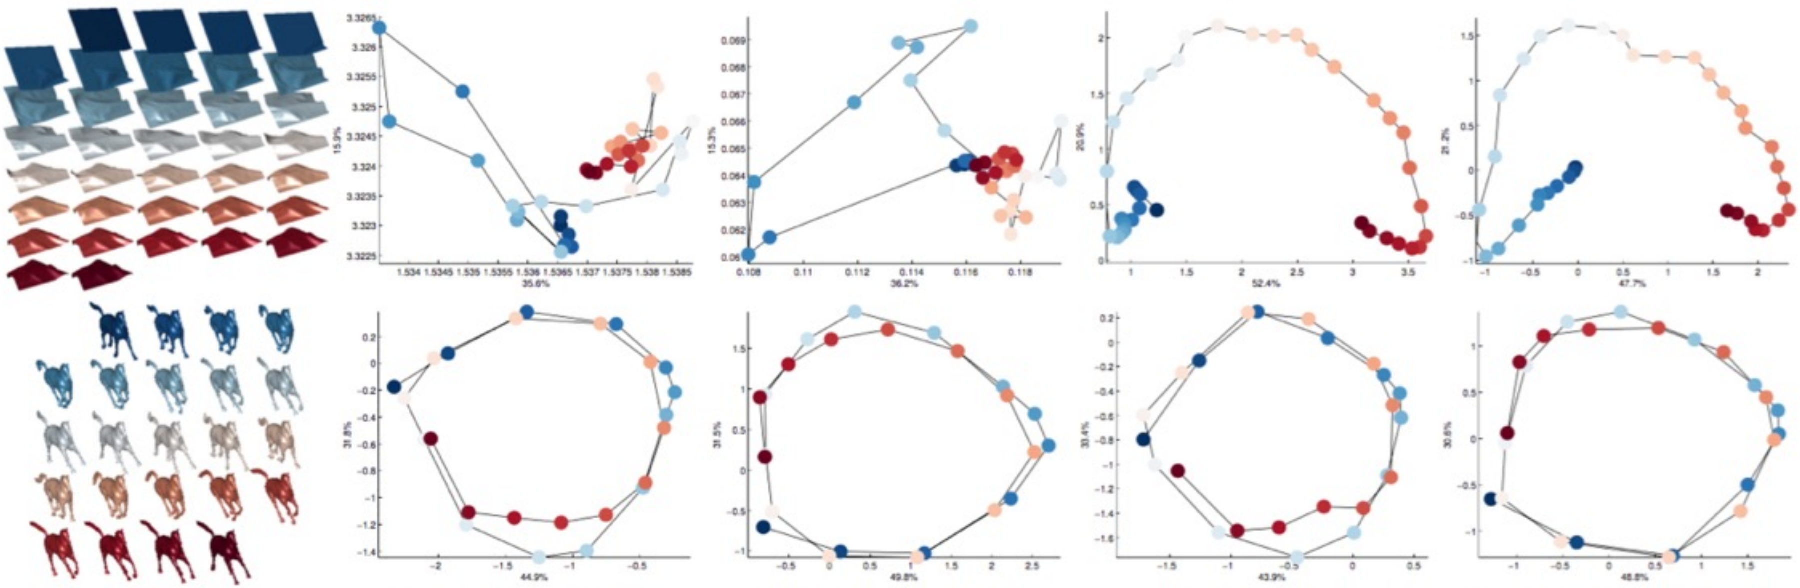

Shapes | Area-based (no offset) | Conformal (no offset) | Area-based (with offset) | Conformal (with offset)

PCA of various shape difference operators

Reconstruction from shape differences

Consequence:

Given the *four* shape difference operators, the shape can be recovered by solving 4 linear systems of equations.

Shape reconstruction can be phrased as reconstruction based on lengths of tetrahedra.

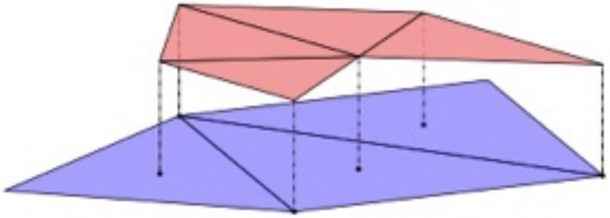

Mesh (blue) and offset (red)

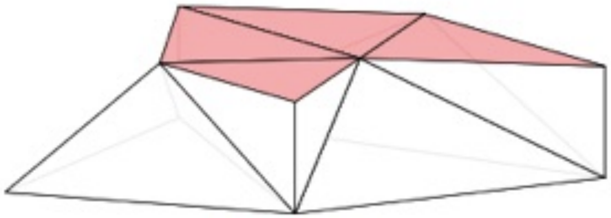

Thickening

Reconstruction from shape differences

Consequence:

An operator view:

The shape is fully encoded by two operators for the first and two for the second fundamental forms.

A coherent, parallel theory in the continuous and discrete case.

Functional Characterization of Intrinsic and Extrinsic Geometry
Corman, Solomon, Ben-Chen, Guibas, O. TOG 2017

Shape Recovery from operators

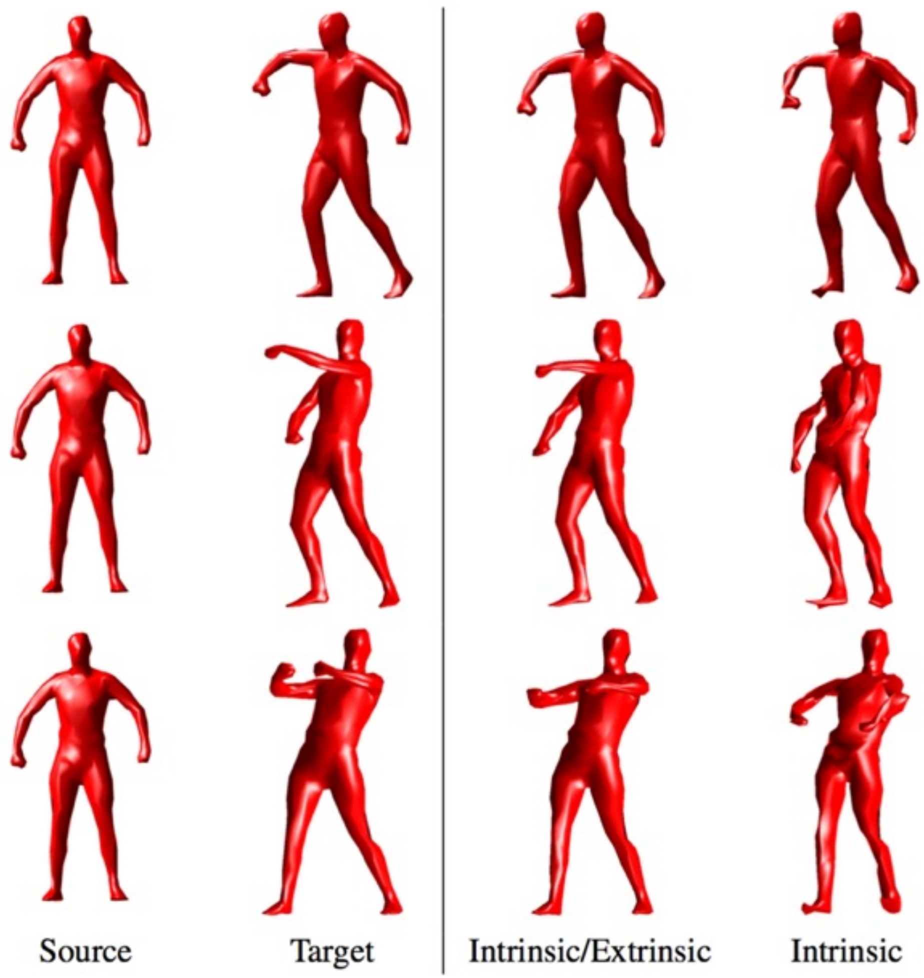

Shape Recovery from operators

Can use the pipeline for interpolation/extrapolation, even with different connectivity.

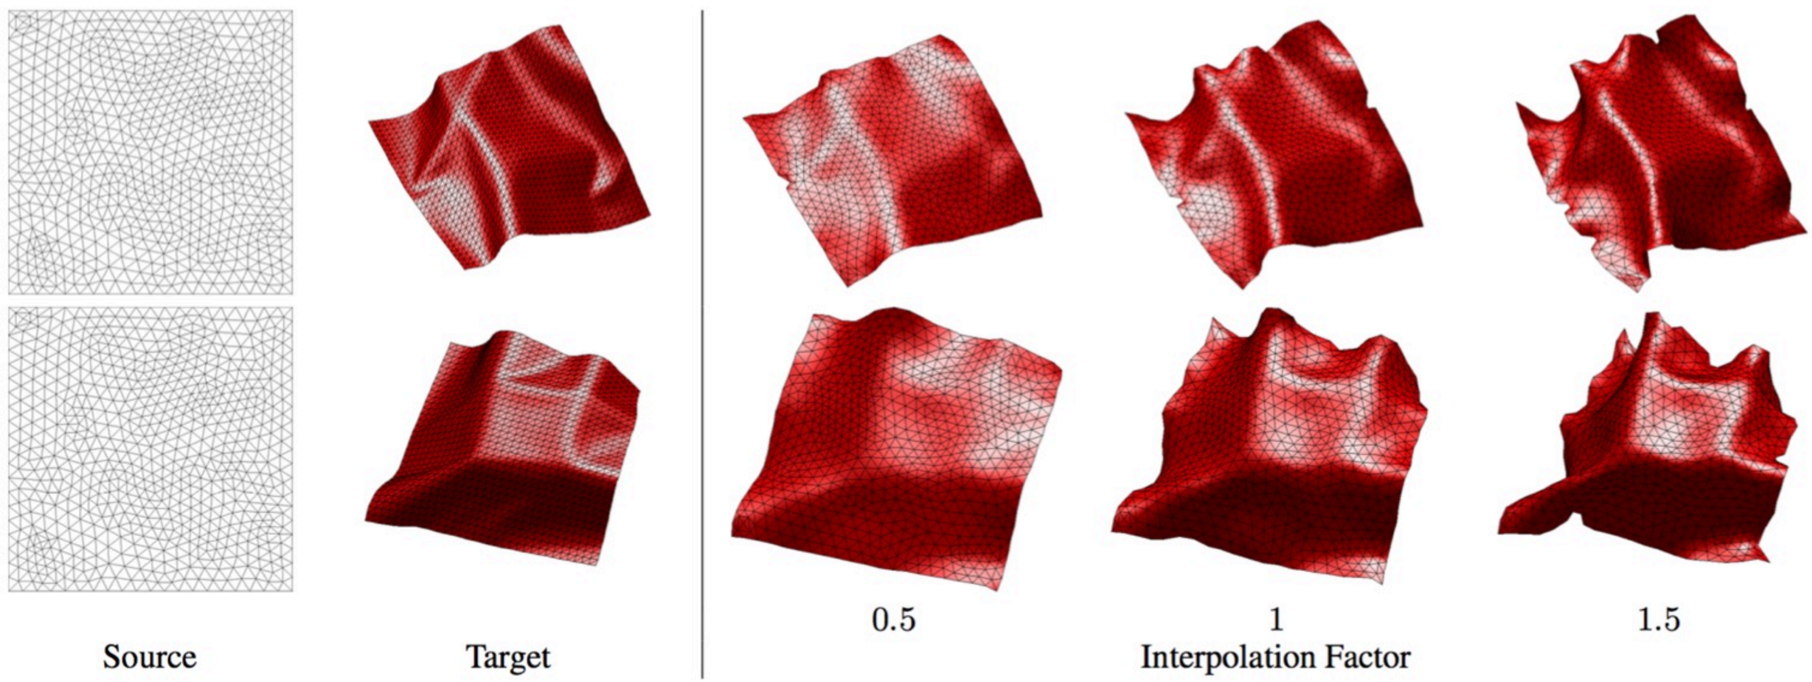

Shape Recovery from operators

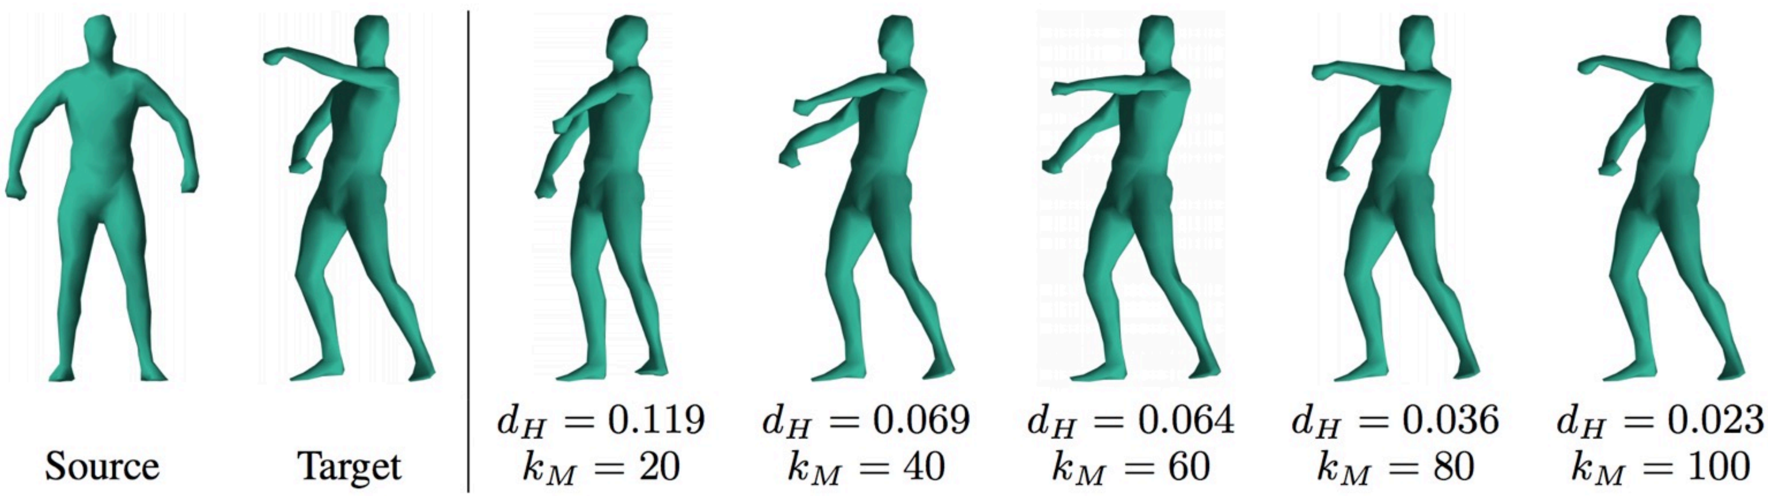

Functional Characterization of Intrinsic and Extrinsic Geometry
Corman, Solomon, Ben-Chen, Guibas, O. TOG 2017

Unbiased Shape Differences

Basic shape differences require a *star-shaped* graph.

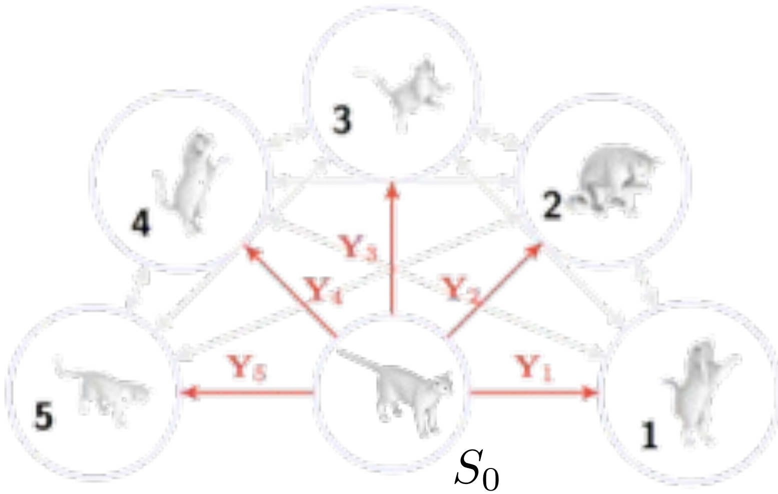

S_0

$$S_i \Longleftrightarrow D_{S_0,S_i}$$

Map-Based Exploration of Intrinsic Shape Differences and Variability, *Rustamov, O., Azencot, Ben-Chen, Chazal, Guibas*, SIGGRAPH 2013

Unbiased Shape Differences

Basic shape differences require a *star-shaped* graph.

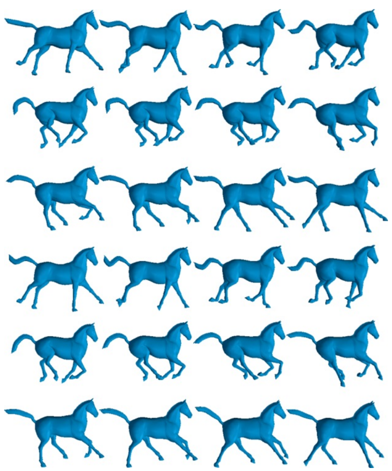

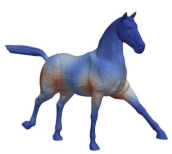

S_0

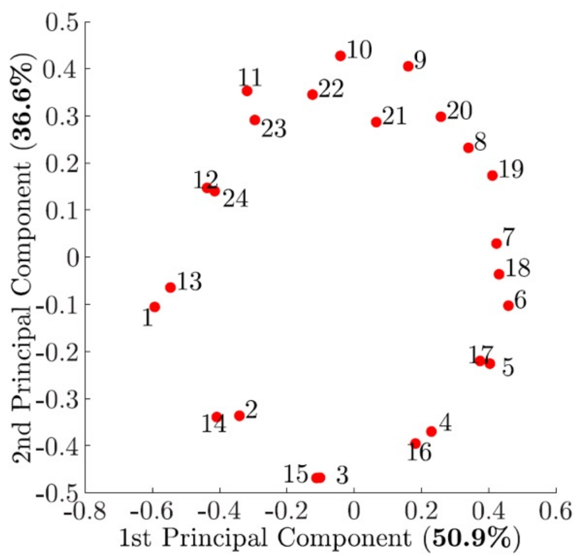

$$S_i \Longleftrightarrow D_{S_0,S_i}$$

Map-Based Exploration of Intrinsic Shape Differences and Variability, *Rustamov, O., Azencot, Ben-Chen, Chazal, Guibas*, SIGGRAPH 2013

Unbiased Shape Differences

What happens if there is no single *base shape?*

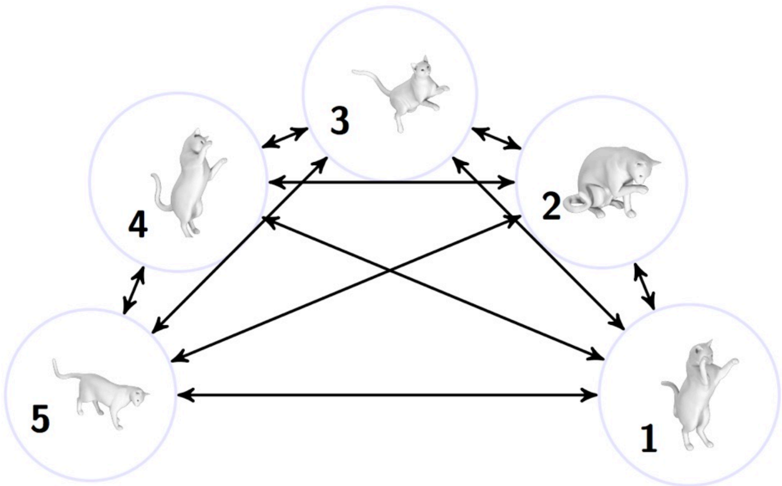

Latent Space Representation for Shape Analysis and Learning, Huang, Achlioptas, Guibas, O. 2018 (arXiv:1806.03967)

Unbiased Shape Differences

Can define and compute a *latent shape* with well-defined geometric structure.

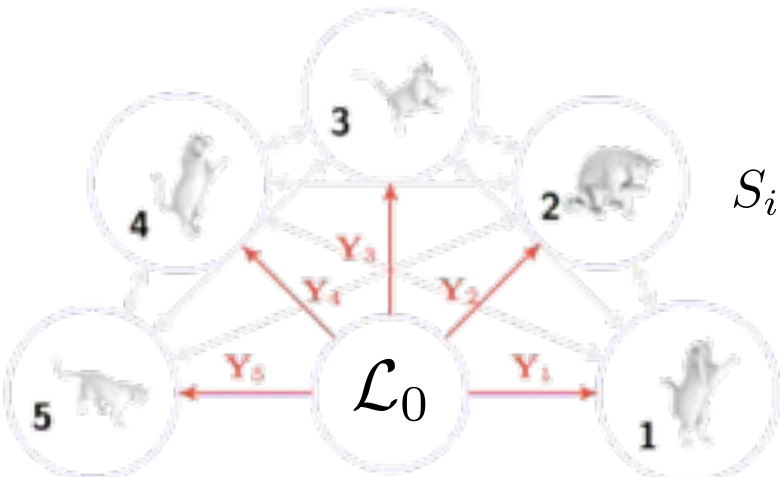

$$S_i \Longleftrightarrow D_{\mathcal{L}_0,S_i} =$$

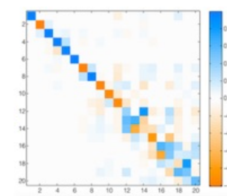

Every shape in the collection is represented as a small-sized matrix, **independent of a base shape**!

Latent Space Representation for Shape Analysis and Learning, Huang, Achlioptas, Guibas, O. 2018 (arXiv:1806.03967)

Latent Shape Spaces

- Given a *functional map nework,* enforcing loop closure by creating a "latent" shape:

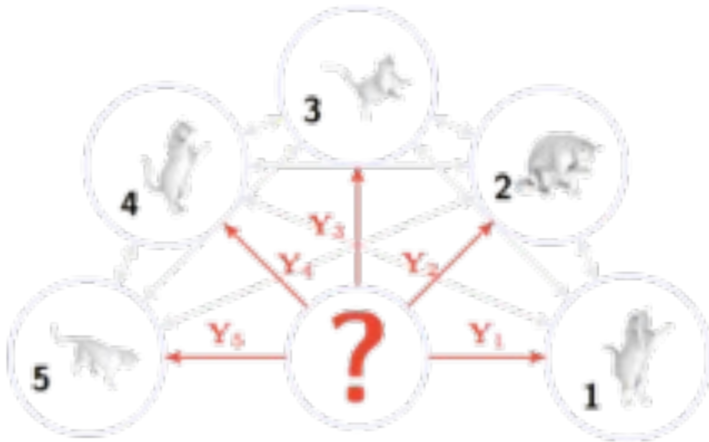

$$C_{ij} = Y_j Y_i^{-1}$$

- Find the optimal latent space by solving: $\min_{\mathbf{Y}} \|C_{ij} Y_i - Y_j\|_F^2$
If $\mathbf{Y}^T \mathbf{Y} = Id$ this reduces to an eigenvalue problem.

- Given Y_i's, solve for C_{ij} to enforce consistency. Restart.

Image Co-Segmentation via Consistent Functional Maps Wang, Huang, Guibas, CVPR 2013

Functional map networks for analyzing and exploring large shape collections, Huang, Wang, Guibas, SIGGRAPH 2014

Latent Shape Spaces

- Given a *functional map nework,* enforcing loop closure by creating a “latent” shape:

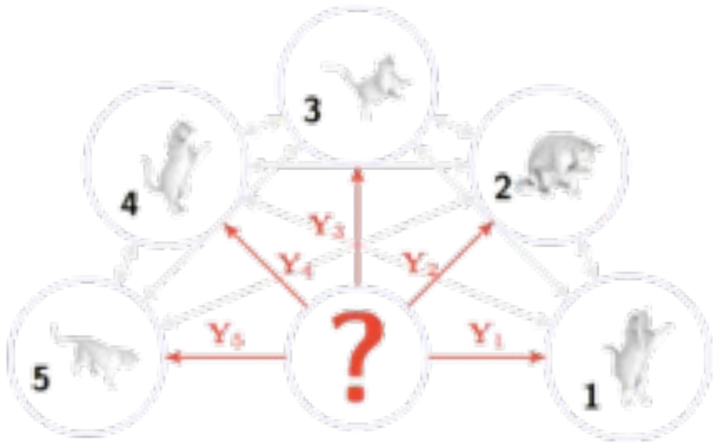

$$C_{ij} = Y_j Y_i^{-1}$$

- Main observation: the latent shape can be endowed with metric and measure structure (although is not embeddable).

Latent Space Representation for Shape Analysis and Learning, Huang, Achlioptas, Guibas, O. 2018 (arXiv:1806.03967)

Unbiased Shape Differences

Given a *chain-shaped* functional map graph:

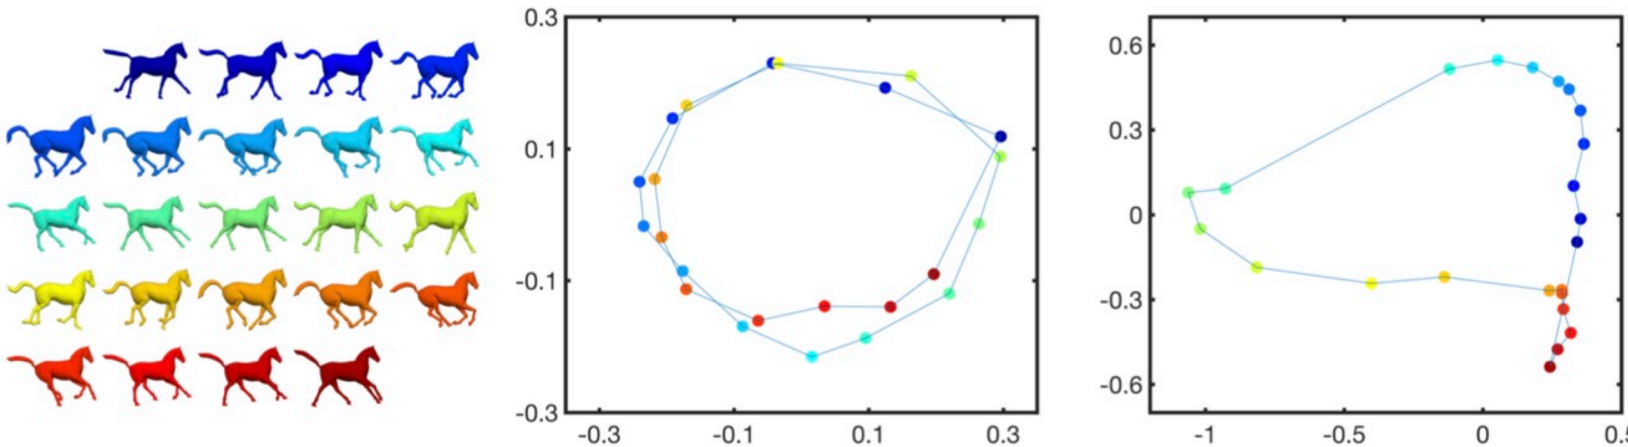

(a) input shape collection **(b) PCA on latent shape differences** **(c) PCA on shape differences to a fixed base shape**

Latent Space Representation for Shape Analysis and Learning, Huang, Achlioptas, Guibas, O. 2018 (arXiv:1806.03967)

Functional Shape Differences + Learning

Each shape is represented as a small-sized matrix. Can use deep-learning (CNN-based) techniques!

$$S_i \Longleftrightarrow D_{\mathcal{L}_0, S_i} =$$

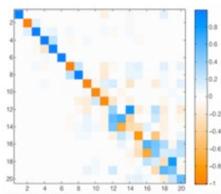

Shape reconstruction with convolutional neural networks:

Latent Space Representation for Shape Analysis and Learning, Huang, Achlioptas, Guibas, O. 2018 (arXiv:1806.03967)

Functional Shape Differences + Learning

Algebraic operations on the difference matrices:

$$D_{S_i,S_j} = D_{S_i}^{-1} D_{S_j}$$

Useful for deformation and style transfer.

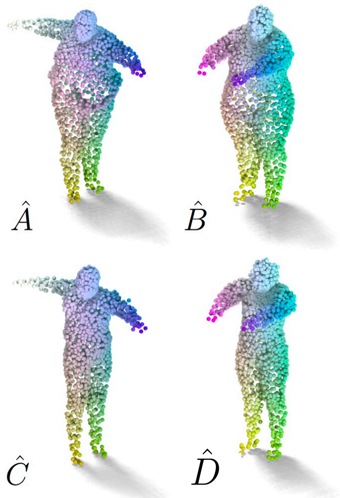

Pose Transfer

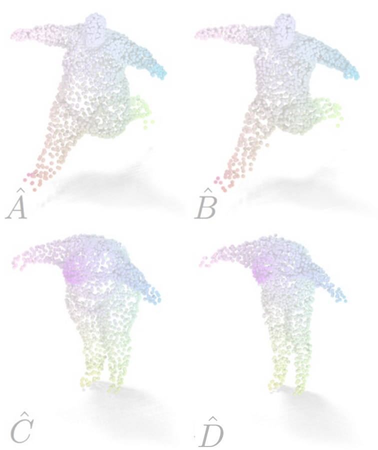

Style Transfer

Latent Space Representation for Shape Analysis and Learning, Huang, Achlioptas, Guibas, O. 2018 (arXiv:1806.03967)

Functional Shape Differences + Learning

Each shape is represented as a small-sized matrix. Can use deep-learning (CNN-based) techniques!

$$S_i \Longleftrightarrow D_{\mathcal{L}_0, S_i} =$$

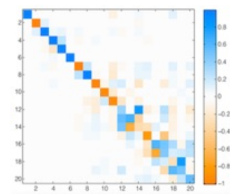

Shape *analysis* via deep learning:

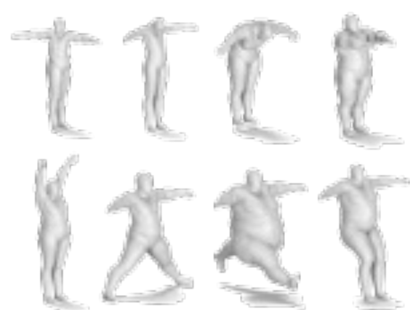

Shape Collection

Estimation accuracy:

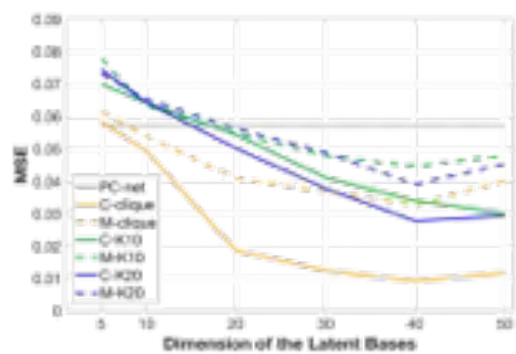

Latent Space Representation for Shape Analysis and Learning, Huang, Achlioptas, Guibas, O. 2018 (arXiv:1806.03967)

Conclusion

- Shape differences allow to encode the shapes as linear operators.
- Can recover the metric from a inner products (shape differences or Laplacian) even in a noisy/approximate case.
- Define unbiased shape differences, by considering *latent* shapes.

Thank you!

Questions?

Exploring shapes with extrinsic information

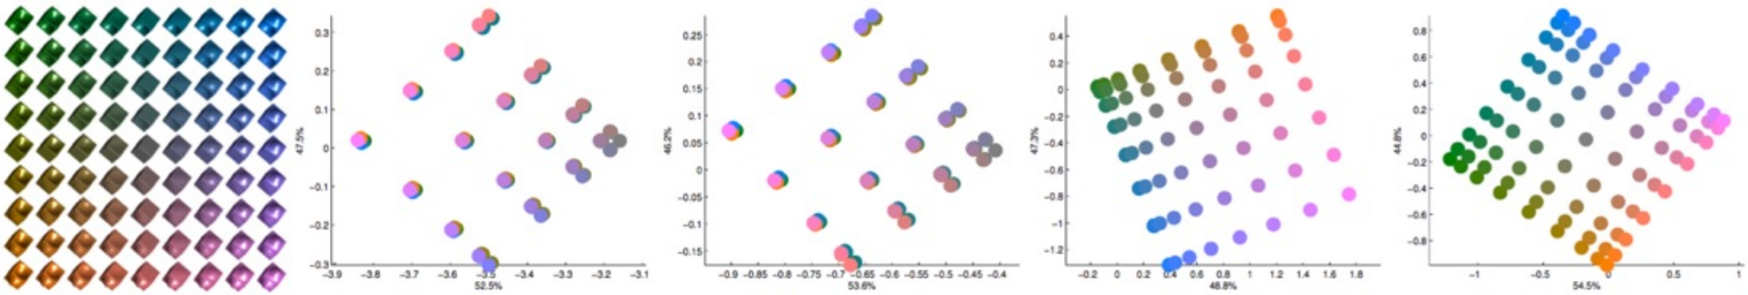

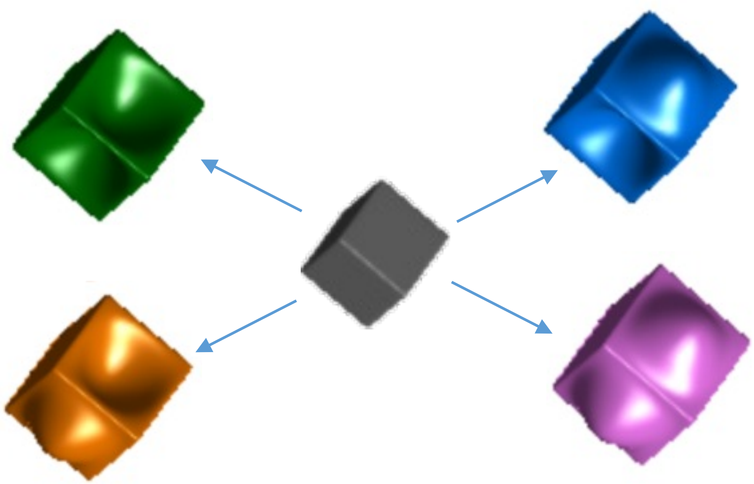

Extrinsic shape differences can distinguish between “inward” and “outward” deformations.

Shape Differences

- Fully characterize intrinsic (metric) distortion using two linear functional operators.
- Can compute areas of maximal distortion through eigen-decomposition.
- Can *compare* distortion of different pairs A->B, vs C->D.

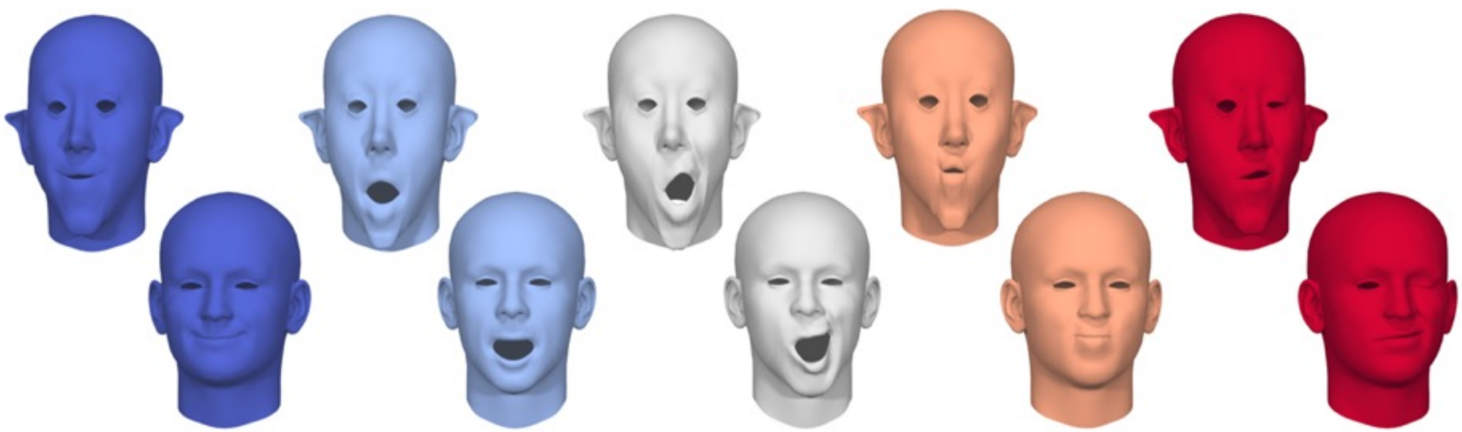

Map-Based Exploration of Intrinsic Shape Differences and Variability, *Rustamov, O., Azencot, Ben-Chen, Chazal, Guibas*, SIGGRAPH 2013

Functional Approach to Mappings

Given two shapes and a pointwise bijection $T : \mathcal{N} \to \mathcal{M}$

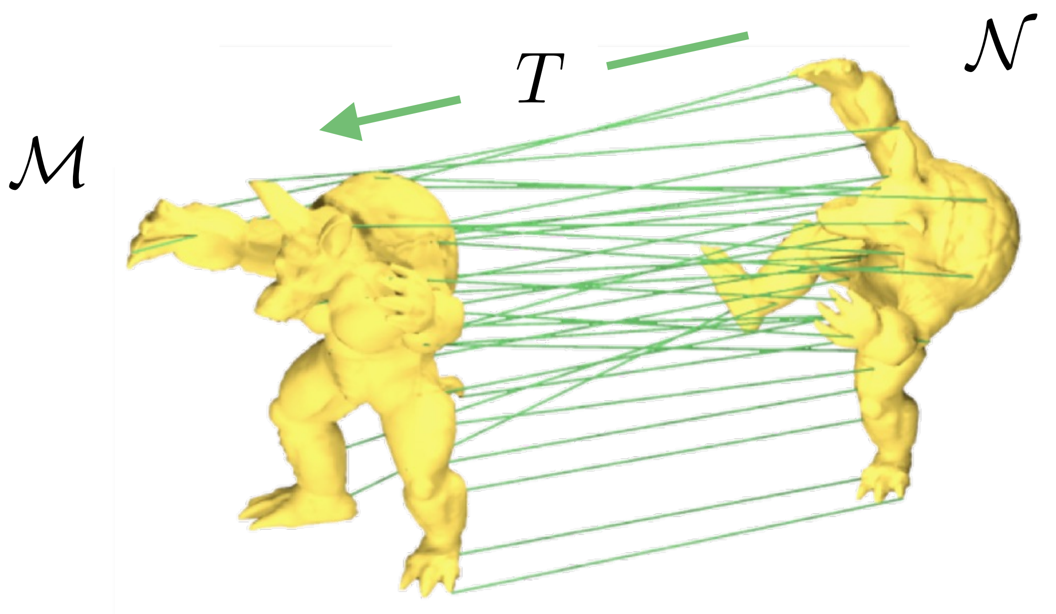

The map T induces a functional correspondence:

$$T_F(f) = g, \text{ where } g = f \circ T$$

Functional maps: a flexible representation of maps between shapes, *O., Ben-Chen, Solomon, Butscher, Guibas*, SIGGRAPH 2012

Functional Approach to Mappings

Given two shapes and a pointwise map $T : \mathcal{N} \to \mathcal{M}$

$$T_F(f) = g : \mathcal{N} \to \mathbb{R}$$

$$T_F$$

$$f : \mathcal{M} \to \mathbb{R}$$

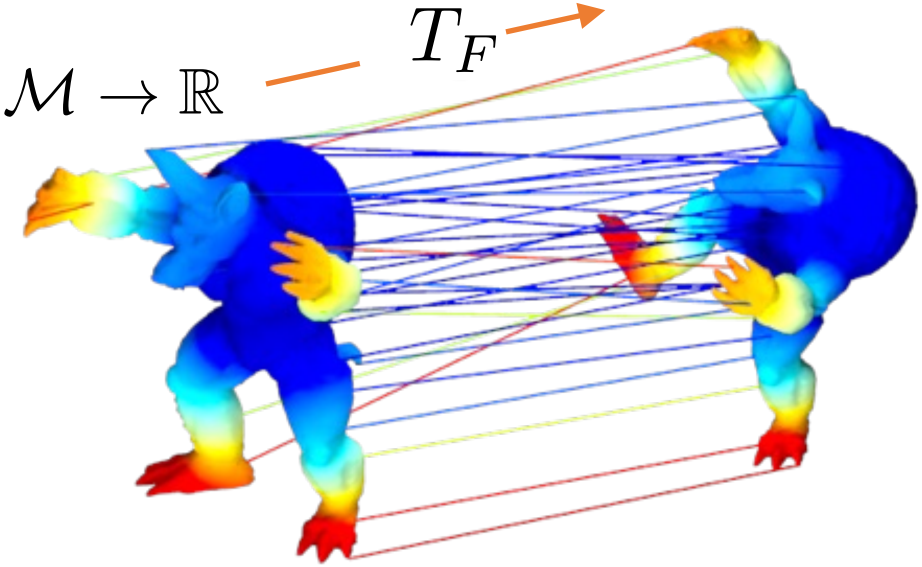

The map T induces a functional correspondence:

$$T_F(f) = g, \text{ where } g = f \circ T$$

Functional maps: a flexible representation of maps between shapes, *O., Ben-Chen, Solomon, Butscher, Guibas*, SIGGRAPH 2012

Functional Approach to Mappings

Given two shapes and a pointwise map $T : \mathcal{N} \to \mathcal{M}$

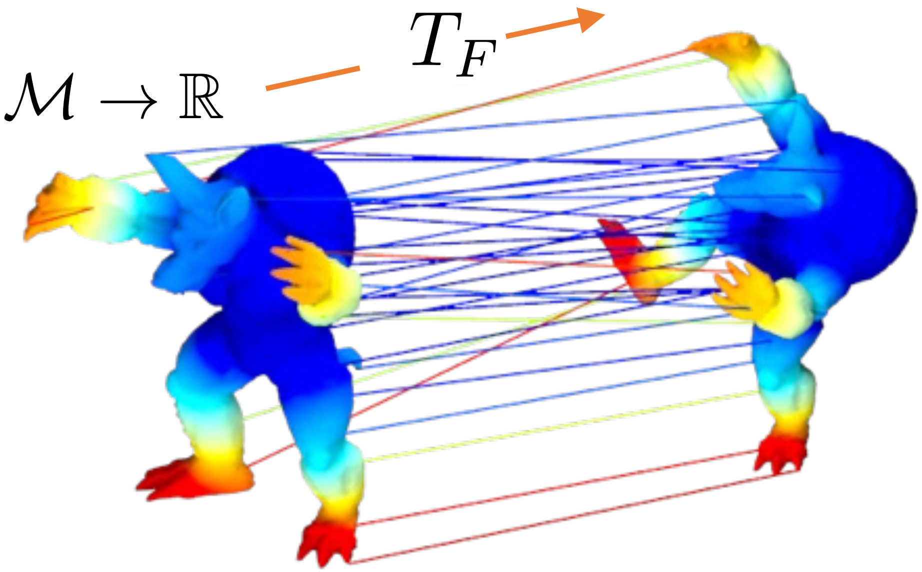

The induced functional correspondence is linear:

$$T_F(\alpha_1 f_1 + \alpha_2 f_2) = \alpha_1 T_F(f_1) + \alpha_2 T_F(f_2)$$

Observation

Assume that both: $f \in \mathcal{L}_2(\mathcal{M}), g \in \mathcal{L}_2(\mathcal{N})$

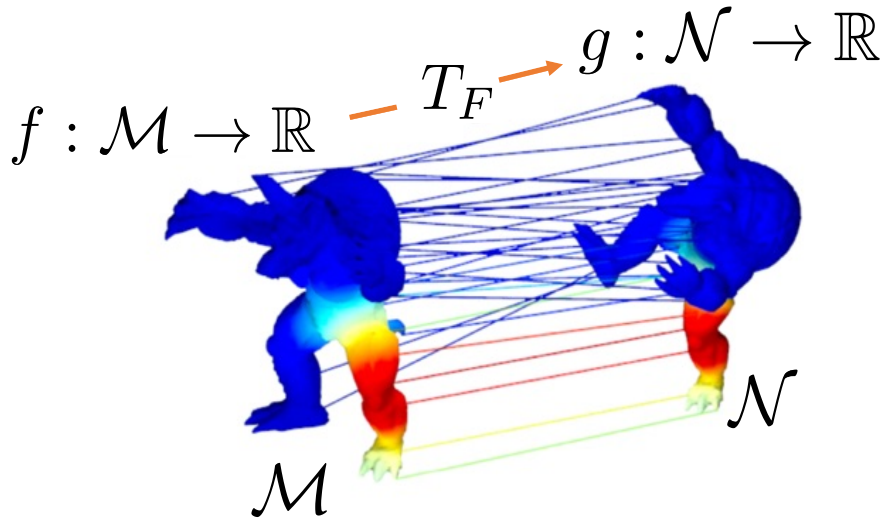

Express both f and $T_F(f)$ in terms of *basis functions:*

$$f = \sum_i a_i \phi_i^{\mathcal{M}} \qquad g = T_F(f) = \sum_j b_j \phi_j^{\mathcal{N}}$$

Since T_F is linear, there is a linear transformation from $\{a_i\}$ to $\{b_j\}$.