How Much Geometry Lies in The Laplacian?

Encoding and recovering the discrete metric on triangle meshes

Distance Geometry Workshop in Bad Honnef,
November 23, 2017

Maks Ovsjanikov

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

Laboratoire d'Informatique de l'École polytechnique

l'X ÉCOLE POLYTECHNIQUE UNIVERSITÉ PARIS-SACLAY

cnrs

What is Geometry Processing

Broad Goals:

To create mathematical models and practical tools for digital representation, manipulation and analysis of 3D shapes.




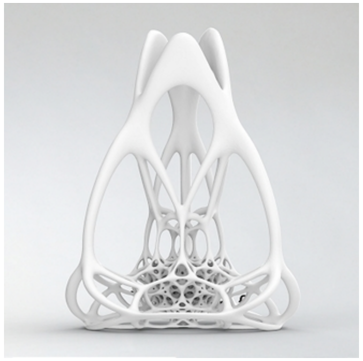
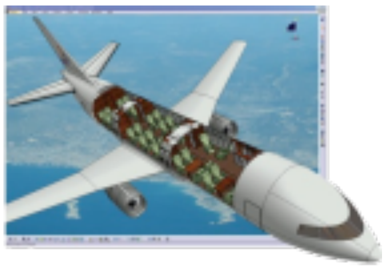
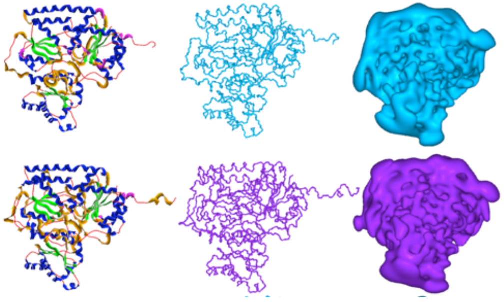

What is a Shape?

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

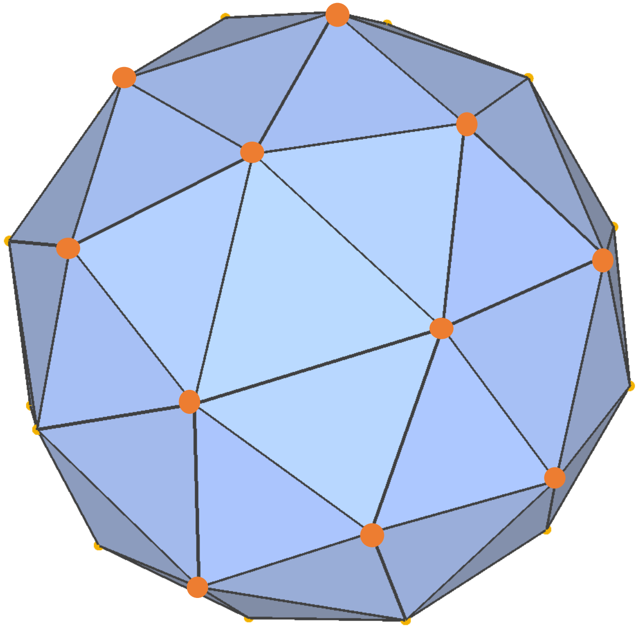

Common assumptions:

- Connected.
- Manifold.
- Without Boundary.

Why triangle meshes

- Functions are piecewise linear inside triangles.
- Can compute gradients.
- Edge lengths correspond to (2x2) matrices inside triangles

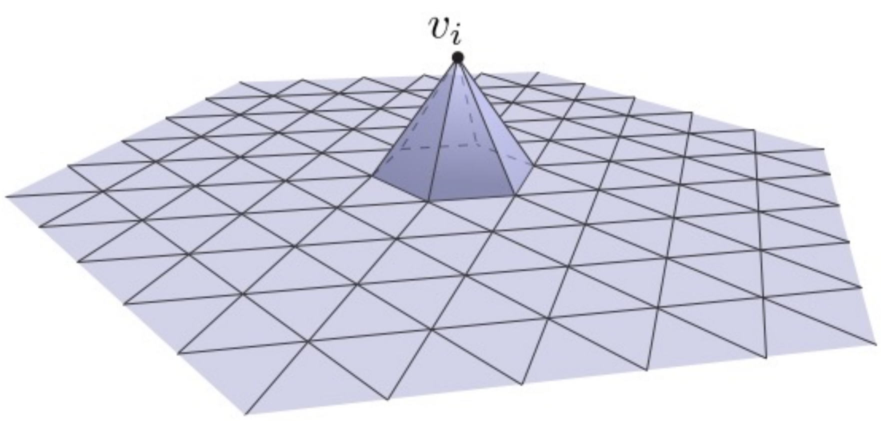

Piecewise-linear functions

$f : \mathcal{V} \rightarrow \mathbb{R}$

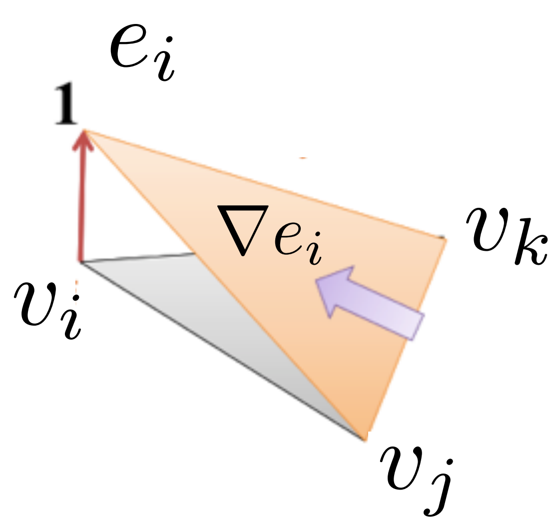

K. Crane and Botsch et al.

What is a Shape?

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

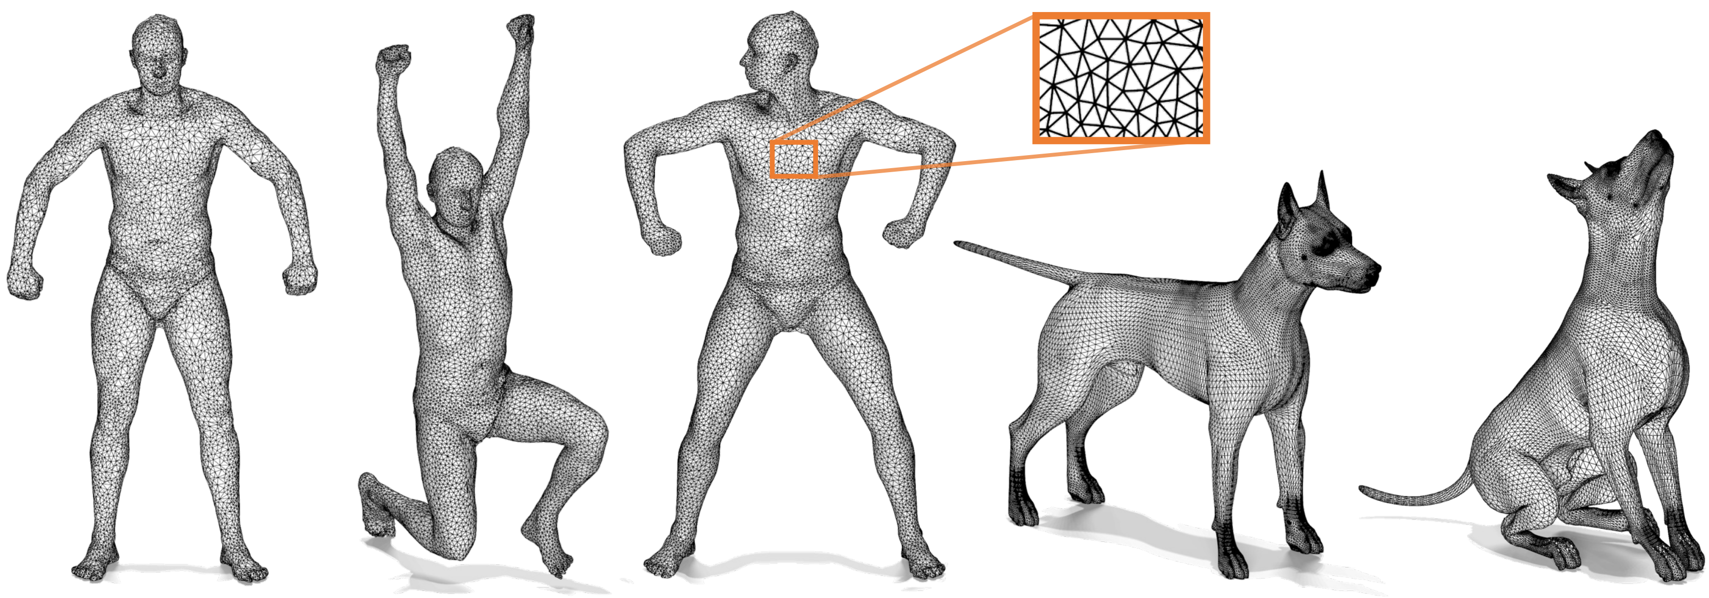

5k – 200k triangles

Shapes from the SCAPE, TOSCA and FAUST datasets

Shape Comparison

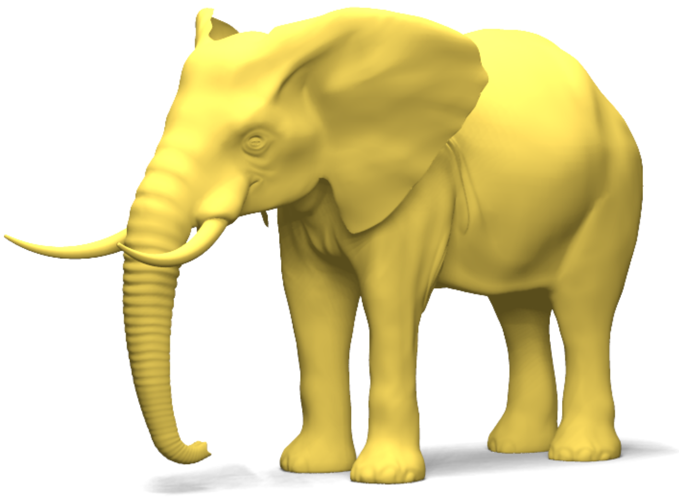

$$\overset{?}{=}$$

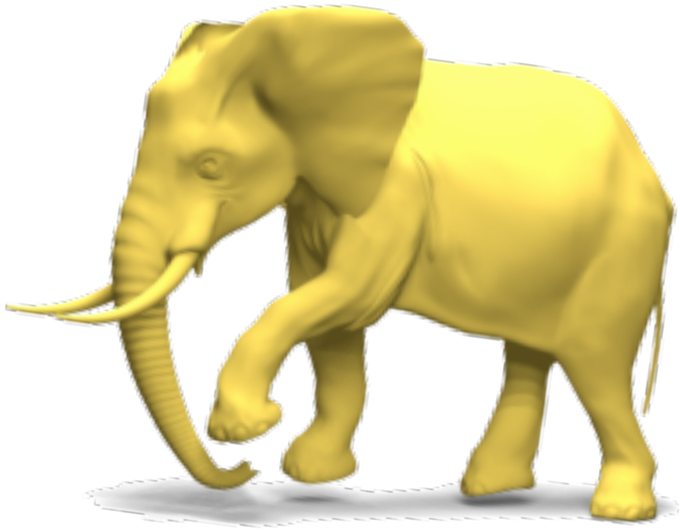

Given two 3D shapes, quantify if they are *similar.*

Shape Matching

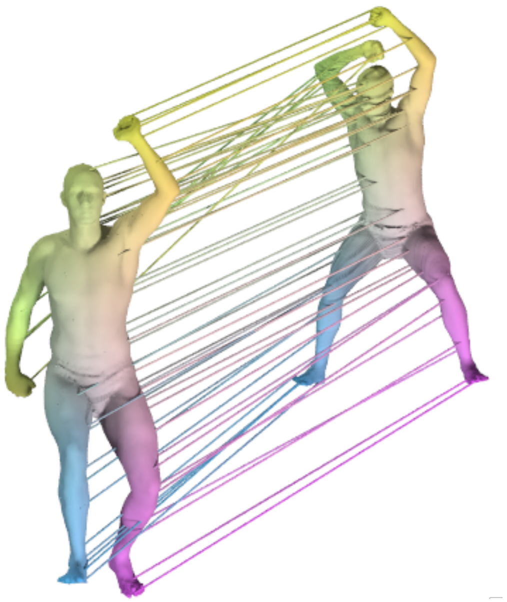

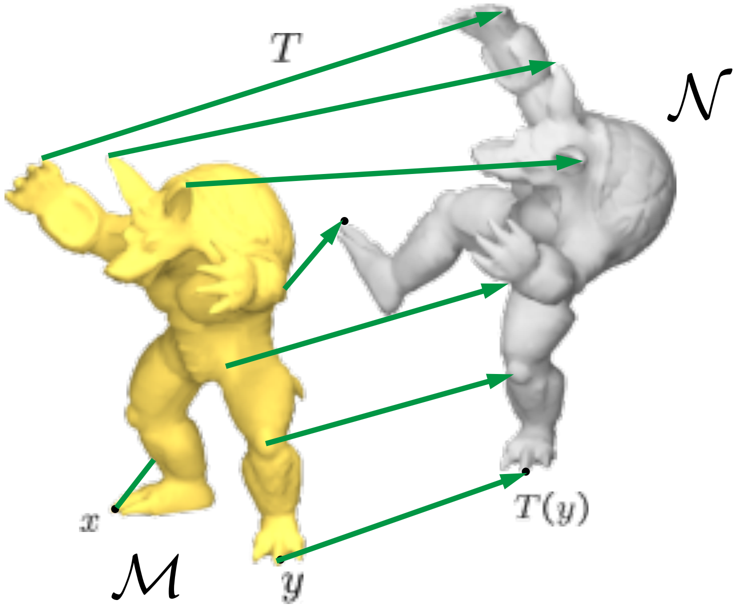

Given two 3D shapes, find corresponding points.

Shape **Matching**

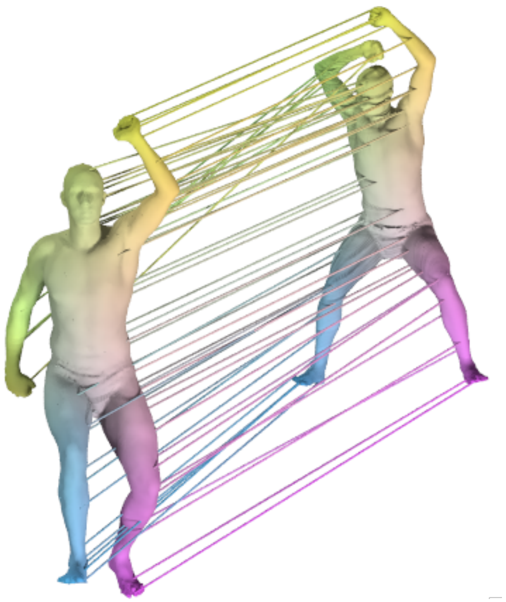

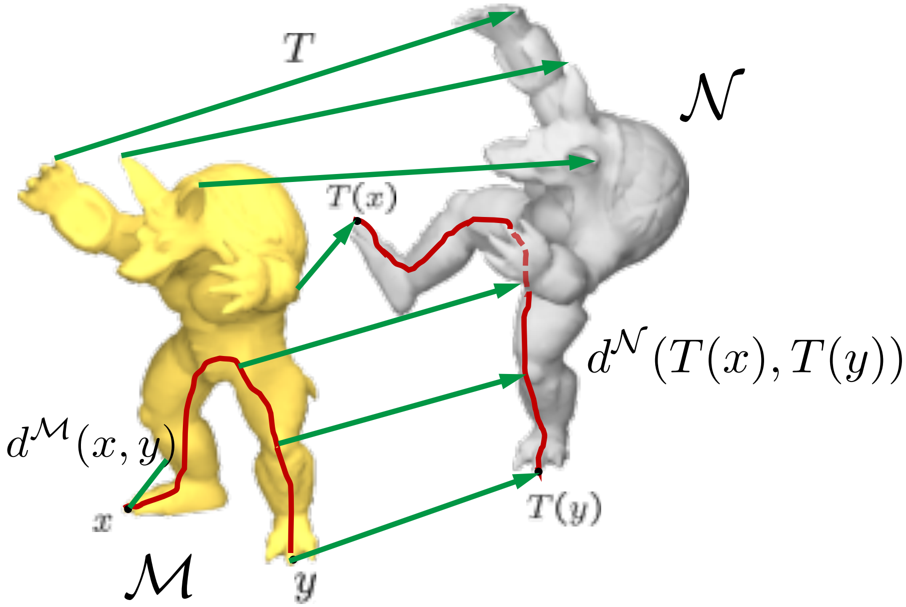

Given two 3D shapes, find intrinsically *isometric* correspondences.

Why Shape Matching

Given a correspondence, we also can **detect and measure the areas of change:**

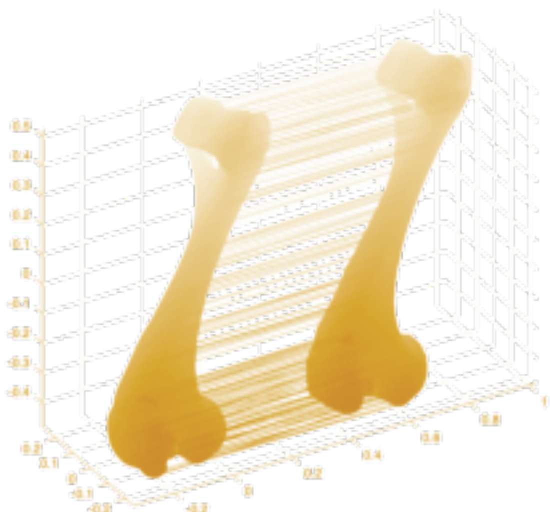

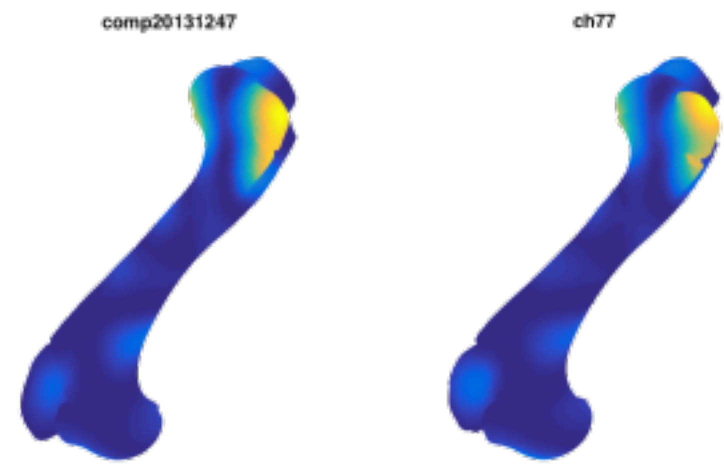

Data from: FunEvol group (CNRS, MNHN)

Today

- Encoding shape changes

- Recovering the shape from the Laplacian-based quantities.

Main observation:

- Many tasks can be formulated through manipulation of linear operators defined on (L2) function spaces.

- Can recover the metric even from noisy data.

Sources for the talk

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

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

Background: Functional Maps

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

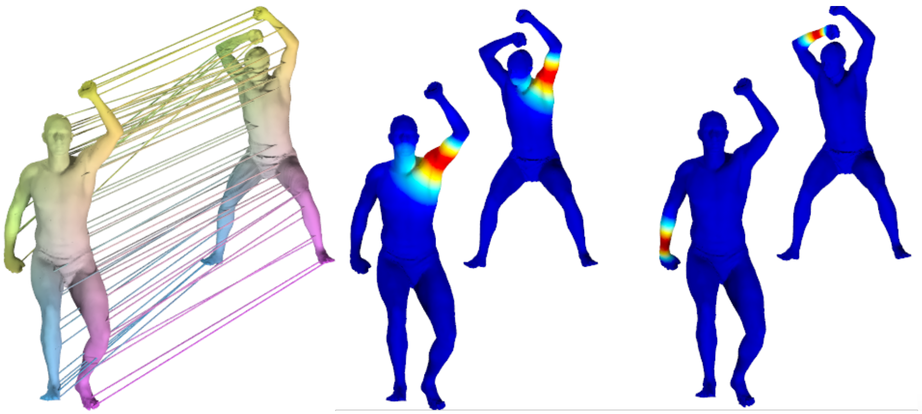

[2] *Functional Maps: A Flexible Representation of Maps Between Shapes,* O., Ben-Chen, Solomon, Butscher. Guibas, SIGGRAPH 2012

[3] *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.

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

[3] *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.

In practice, can think of a functional map as an matrix of size $n_{V_2} \times n_{V_1}$.

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

Motivation

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

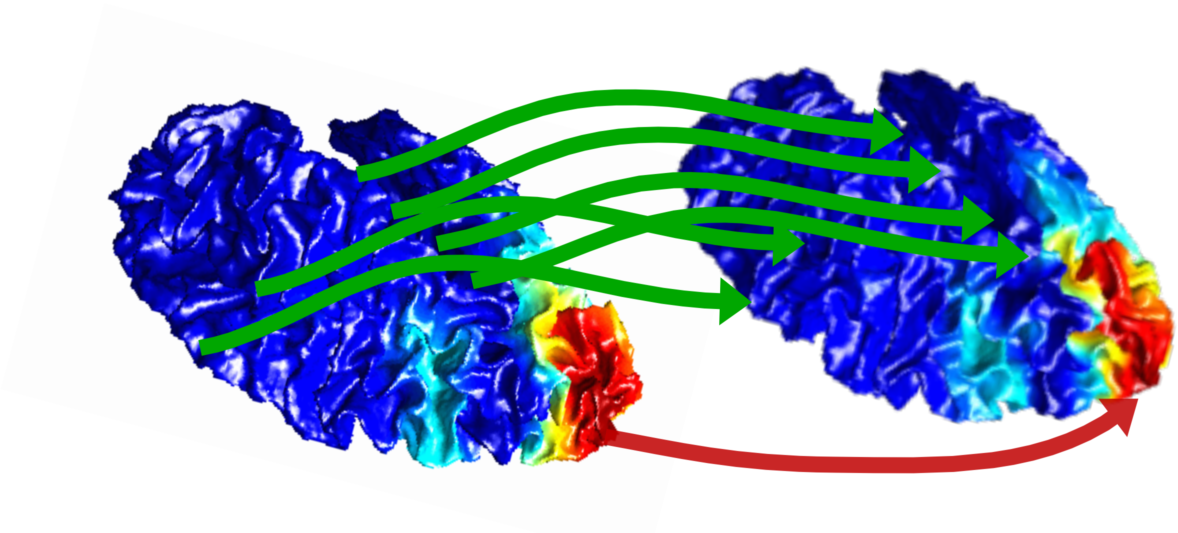

- Do it in a *multi-scale* way (not be sensitive to *local* changes).
- Accommodate approximate *soft* (functional) maps

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 an inner product norm: $\|f\|_M^2 =< f, f >_M$

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

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

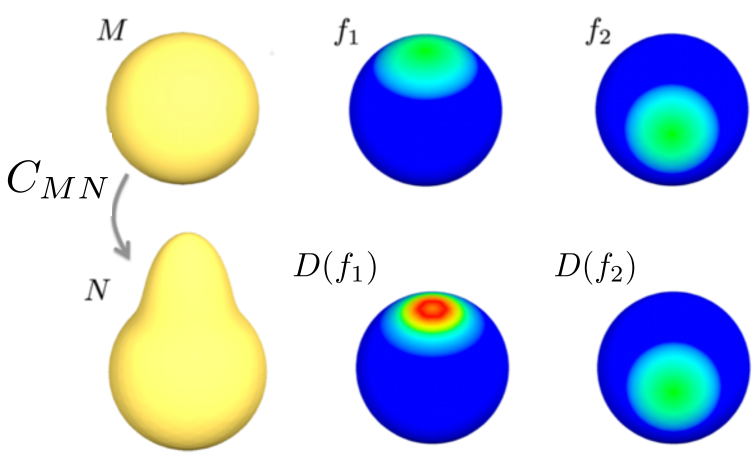

Shape Differences Definition

Given a functional map $C_{MN} : \mathcal{F}(M) \rightarrow \mathcal{F}(N)$
and an inner product norm: $\|f\|_M^2 =< f, f >_M$

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 an inner product norm: $\|f\|_M^2 =< f, f >_M$

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 an inner product norm: $\|f\|_M^2 =< f, f >_M$

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 $C_{MN} : \mathcal{F}(M) \rightarrow \mathcal{F}(N)$
and an inner product norm: $\|f\|_M^2 =< f, f >_M$

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.

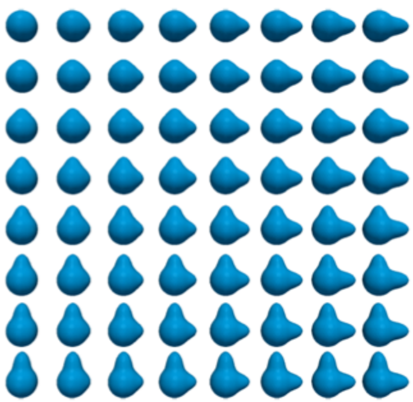

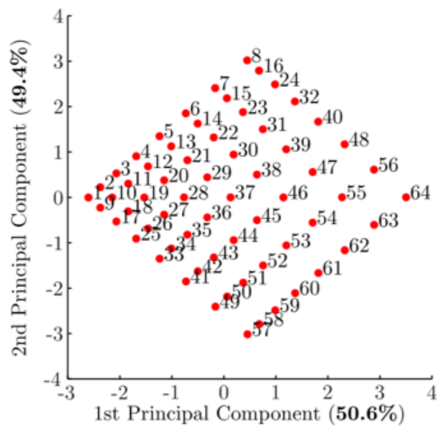

PCA on area-based shape differences

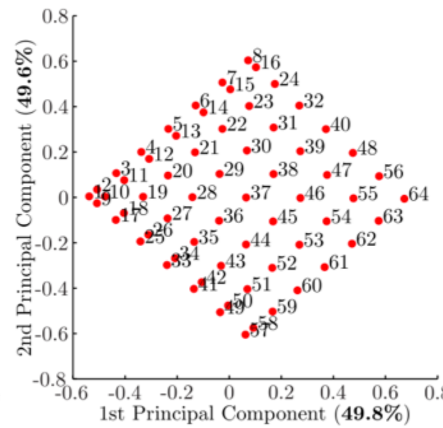

PCA on conformal shape differences

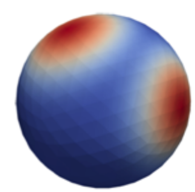

Variability localization for area

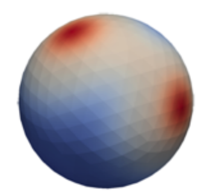

Variability localization for conformal

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.

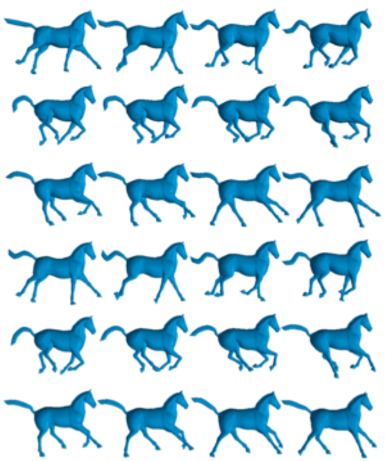

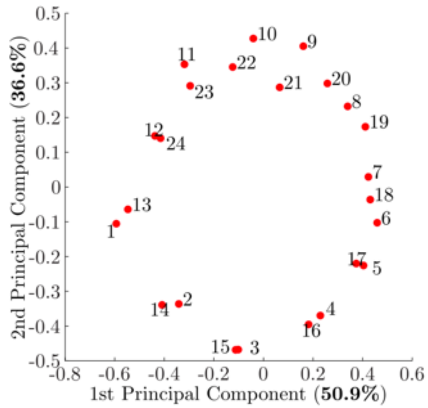

PCA on area-based shape differences

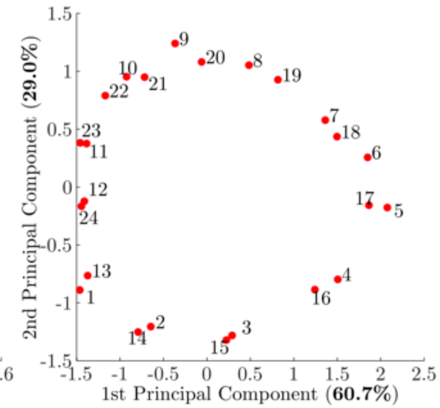

PCA on conformal shape differences

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

Comparing Shape Differences

Find a shape D_i, such that the difference between shapes B and D_i is as close as possible to the difference between A and C_i.

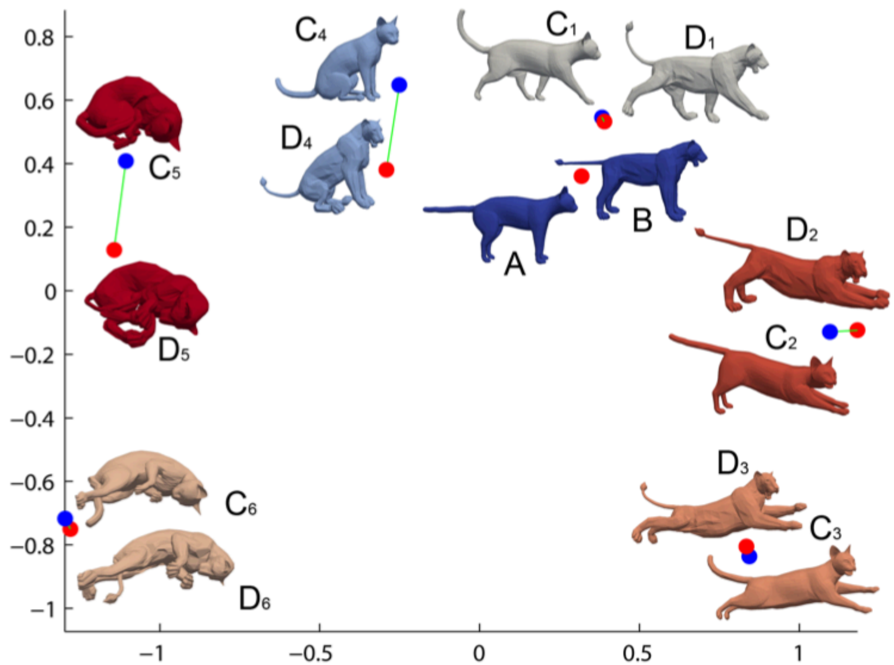

Map-Based Exploration of Intrinsic Shape Differences and Variability, *Rustamov, O., Azencot, Ben-Chen, Chazal, Guibas*, SIGGRAPH 2013

Recap

Shape differences represent the distortion as a pair of linear operators, defined via:

$$< f,\ D(g) >_M = < F(f), F(g) >_N \quad \forall\ f, g$$

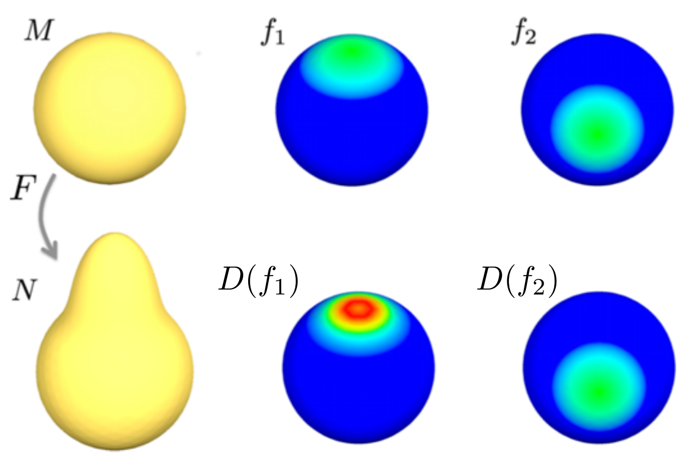

Question

How much information is contained in these operators?

Theorem:

If F comes from a point map:

$$R = Id, \text{ and } V = Id$$

If and only if the map is an intrinsic isometry.

$$V \ : \int_M fg d\mu^M = \int_N fg d\mu^N$$

$$R \ : \int_M < \nabla f, \nabla g > \mu^M = \int_N < \nabla f, \nabla g > \mu^N$$

Can we recover the metric?

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

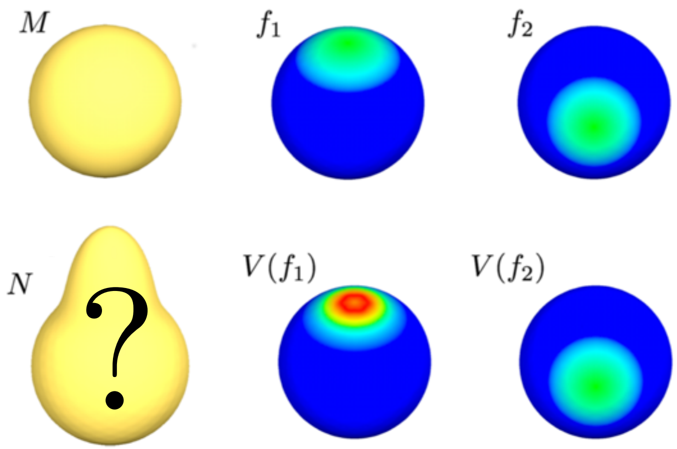

Functional Characterization of Intrinsic and Extrinsic Geometry
Corman, Solomon, Ben-Chen, Guibas, O. TOG 2017

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.

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

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

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.

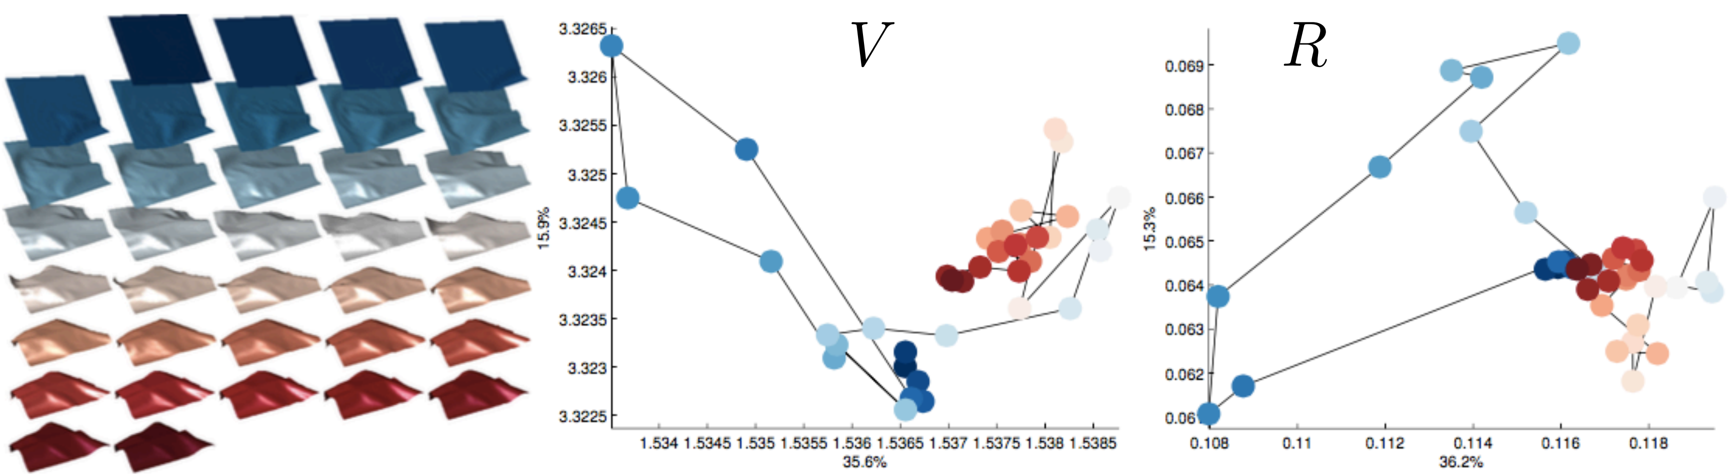

Can we recover the metric?

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 metric to inner products on a triangle mesh:

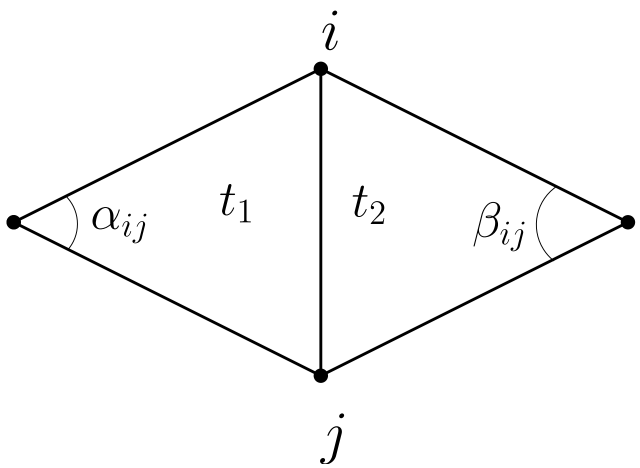

$$L_{ij} = < \nabla e_i, \nabla e_j >$$
$$= \frac{1}{2}\cot(\alpha_{ij}) + \frac{1}{2}\cot(\beta_{ij})$$

Given the 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:

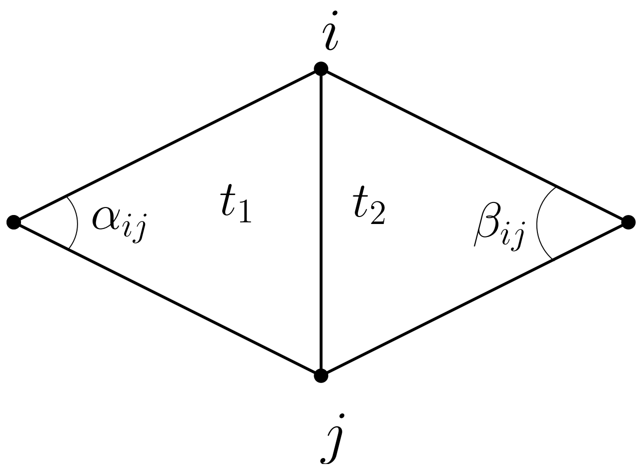

$$L_{ij} = < \nabla e_i, \nabla e_j >$$
$$= \frac{1}{2}\cot(\alpha_{ij}) + \frac{1}{2}\cot(\beta_{ij})$$

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.

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

A metric on the triangle mesh

Alternative expression for the cotangent weights:

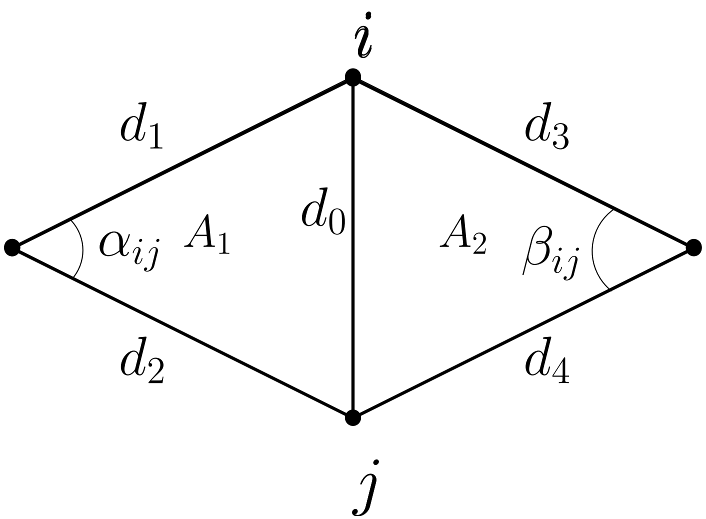

$$< \nabla e_i, \nabla e_j >= \frac{1}{2}\cot(\alpha_{ij}) + \frac{1}{2}\cot(\beta_{ij})$$

$$= \frac{1}{8A_1}(d_0^2 - d_1^2 - d_2^2)$$

$$+ \frac{1}{8A_2}(d_0^2 - d_3^2 - d_4^2)$$

Re-write the weights in terms of edge lengths.

Boscaini et al. *Shape-from-operator: Recovering shapes from intrinsic operators,* CGF, 2015

Corman et al. *Functional Characterization of Intrinsic and Extrinsic Geometry,* TOG 2017

Recovering the metric

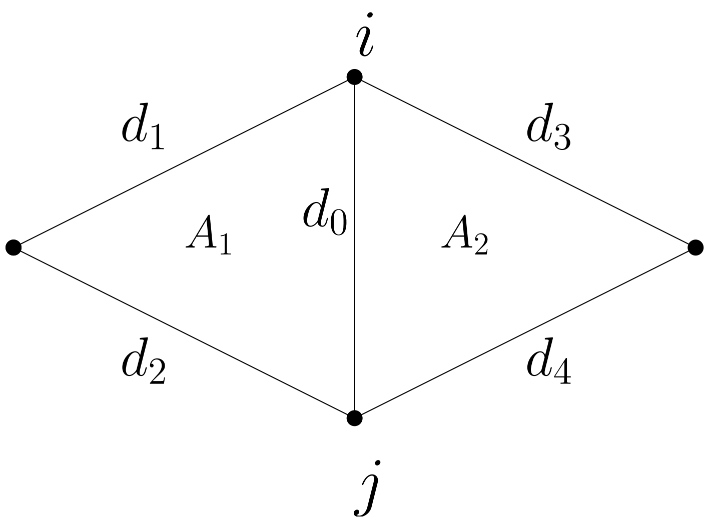

$$< e_i, e_j > = \frac{1}{12}(A_1 + A_2)$$

$$\begin{aligned} < \nabla e_i, \nabla e_j > &= \frac{1}{8A_1}(d_0^2 - d_1^2 - d_2^2) \\ &+ \frac{1}{8\,A_2}(d_0^2 - d_3^2 - d_4^2) \end{aligned}$$

- The areas are linear in the L2 inner product and for fixed areas, the *squared edge lengths* are linear in the H1 inner product.
- The resulting linear systems are *generically* invertible.

Recovering the metric

From Laplacian to the metric:

Theorem:

The edge lengths can be recovered via two linear systems from two matrices of inner products (functions and gradients = cotangent weights), Both are *generically* invertible.

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

Recovering the metric

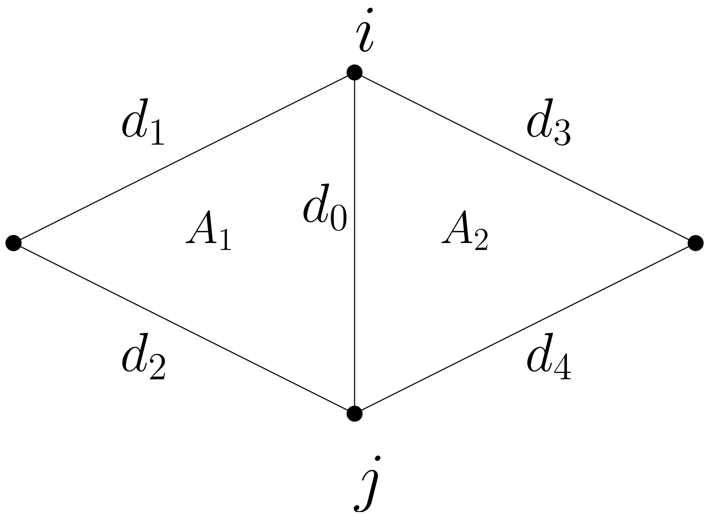

$$< e_i, e_j > = \frac{1}{12}(A_1 + A_2)$$

$$< \nabla e_i, \nabla e_j > = \frac{1}{8A_1}(d_0^2 - d_1^2 - d_2^2) + \frac{1}{8\,A_2}(d_0^2 - d_3^2 - d_4^2)$$

- The areas are linear in the L2 inner product and for fixed areas, the *squared edge lenghts* are linear in the H1 inner product.
- The resulting linear systems are *generically* invertible.
- Can be phrased as a least squares problem even if matrices are noisy/functions are in a different basis.

Recovering the metric

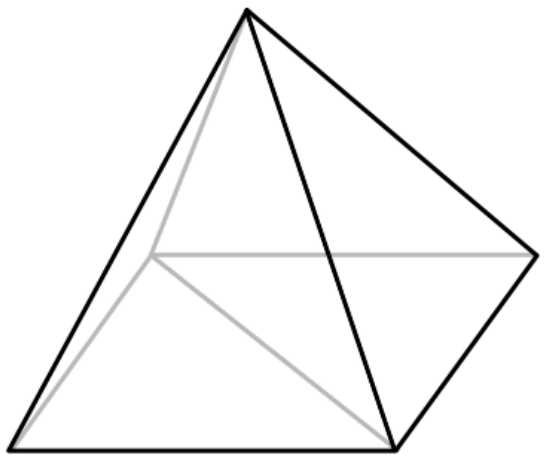

$$< \nabla e_i, \nabla e_j > \; = \frac{1}{8A_1}(d_0^2 - d_1^2 - d_2^2) + \frac{1}{8\,A_2}(d_0^2 - d_3^2 - d_4^2)$$

A mesh for which $C(\ell^2;\mu)$ is not invertible when $\mu = \mathbf{1}$.

- The areas are linear in the L2 inner product and for fixed areas, the *squared edge lenghts* are linear in the H1 inner product.
- The resulting linear systems are *generically* invertible.

Enforcing the Triangle Inequality

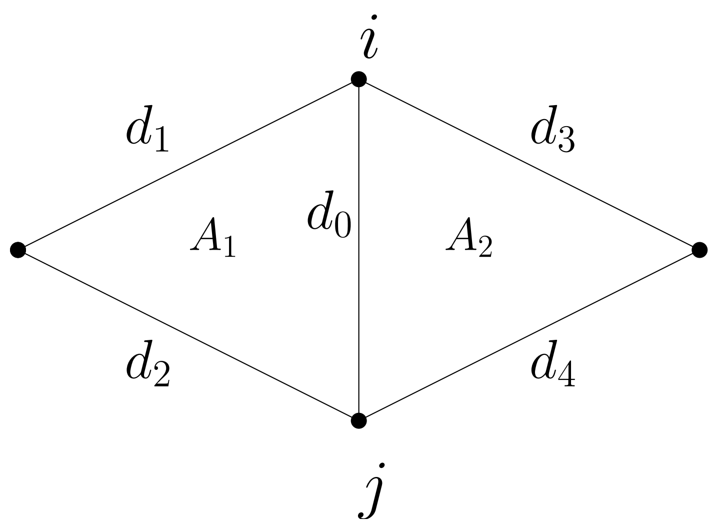

$$< \nabla e_i, \nabla e_j > \; = \frac{1}{8A_1}(d_0^2 - d_1^2 - d_2^2) + \frac{1}{8\,A_2}(d_0^2 - d_3^2 - d_4^2)$$

- Regularization, for noisy/incomplete linear systems:

$$\mathbf{E} = \frac{1}{2}\begin{pmatrix} 2x_1 & x_3 - x_1 - x_2 & x_2 - x_1 - x_3 \\ x_3 - x_1 - x_2 & 2x_2 & x_1 - x_2 - x_3 \\ x_2 - x_1 - x_3 & x_1 - x_2 - x_3 & 2x_3 \end{pmatrix}$$

Is positive semi-definite if and only if x_1, x_2, x_3 are non-negative and their square roots satisfy the triangle inequality.

From Metric to Geometry

Problem:

Given a triangle mesh with approximate edge lengths
Recover the embedding.

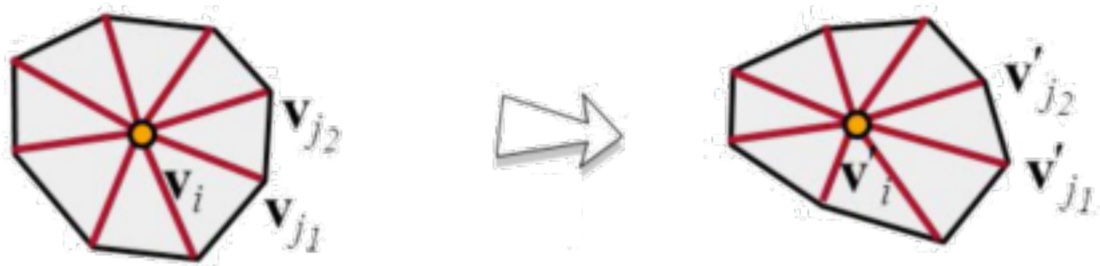

Main idea: deform the triangles to match the target metric.

$$\mathcal{E}(\mathbf{p}') = \sum_{t \in \mathcal{M}} \min_{\mathbf{Q}_t \in SO(3)} A_t \left\| \mathbf{J}_t(\mathbf{p}') - \mathbf{Q}_t \tilde{\mathbf{W}}_t^{-1} \right\|_F^2$$

Iterate between computing $\mathbf{p}'$ and $\mathbf{Q}_t$.

Panozzo et al., *Frame Fields: Anisotropic and Non-Orthogonal Cross Fields,* SIGGRAPH 2014

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
Solomon, Corman, Ben-Chen, Guibas, O. Conditionally accepted at TOG 2016

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.

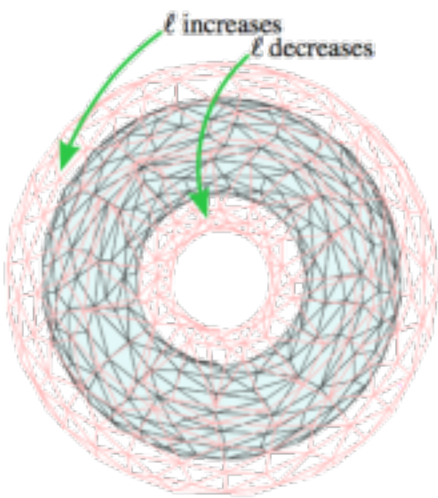

Edge-lengths change according to curvature of the offset surface.

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.

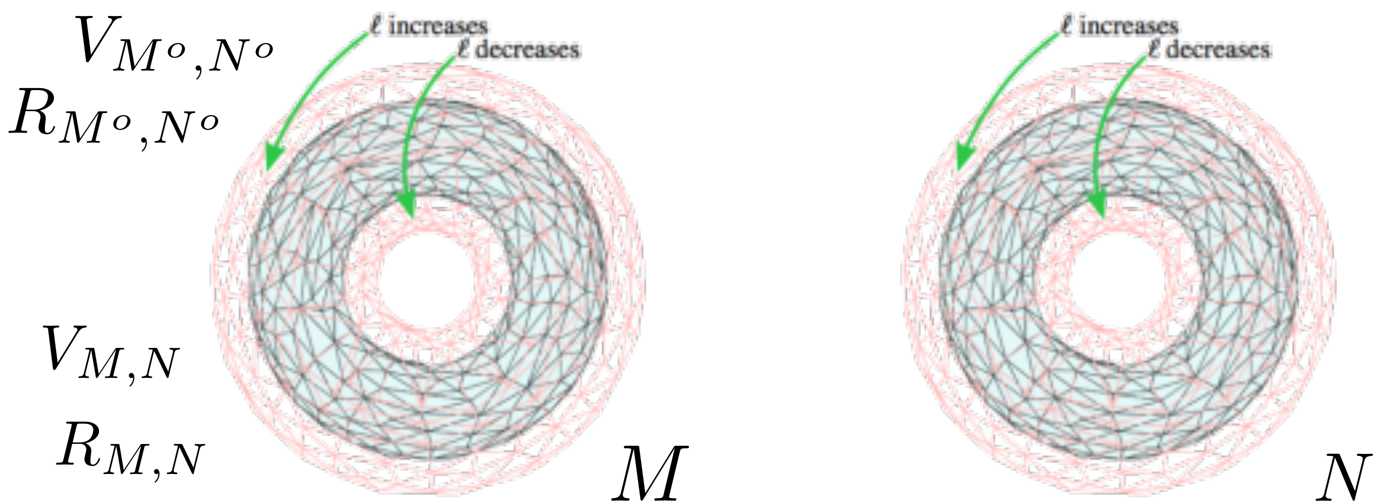

$V_{M,N}, R_{M,N}$ encode change in metric,
V_{M^o,N^o}, R_{M^o,N^o} encode change in curvature

Exploring shapes with extrinsic information

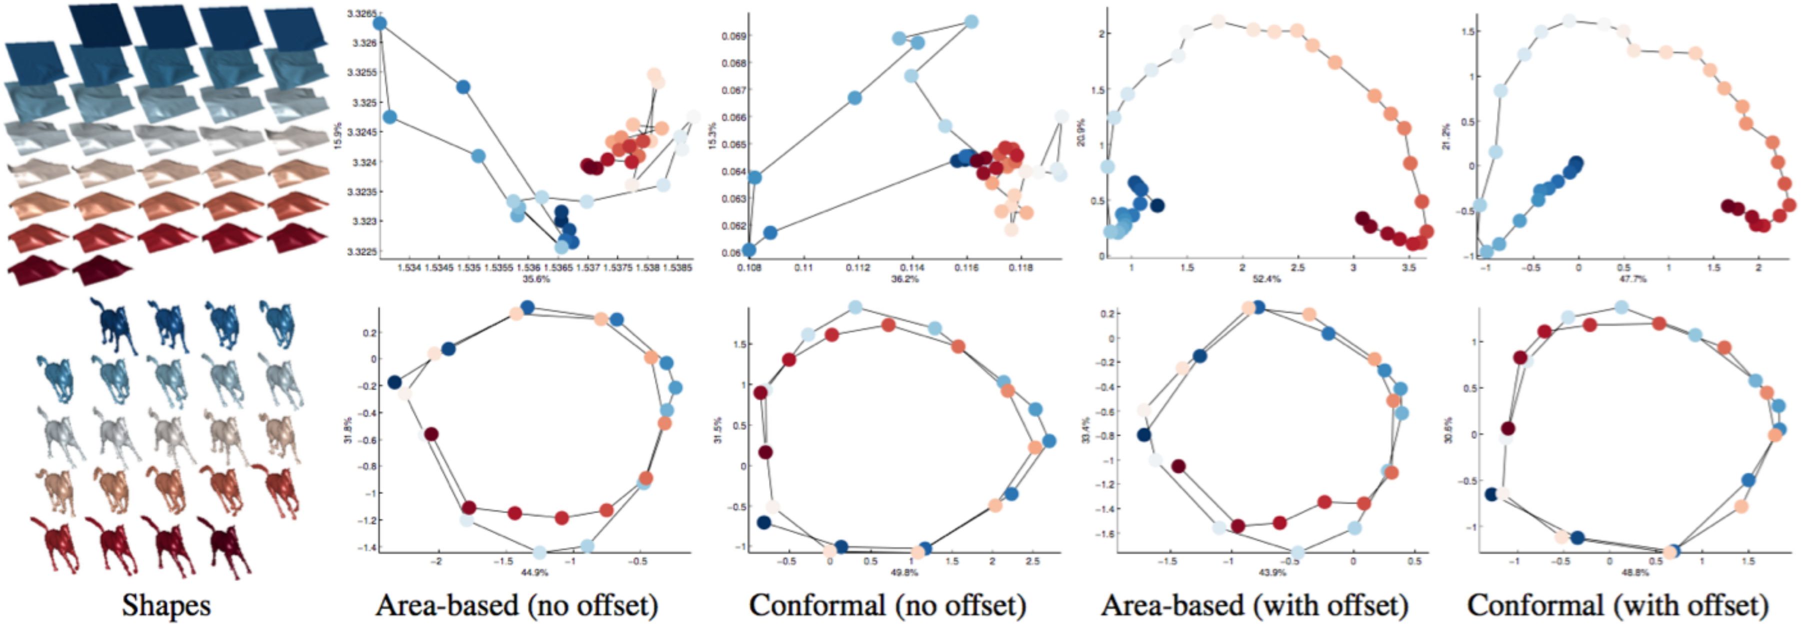

PCA of various shape difference operators

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.

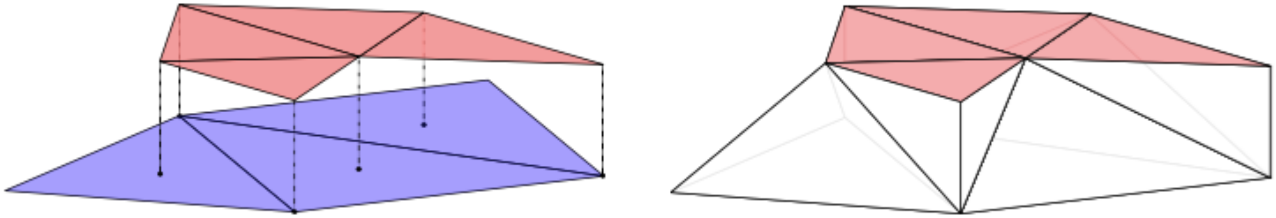

Mesh (blue) and offset (red)

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.

Shape Recovery from operators

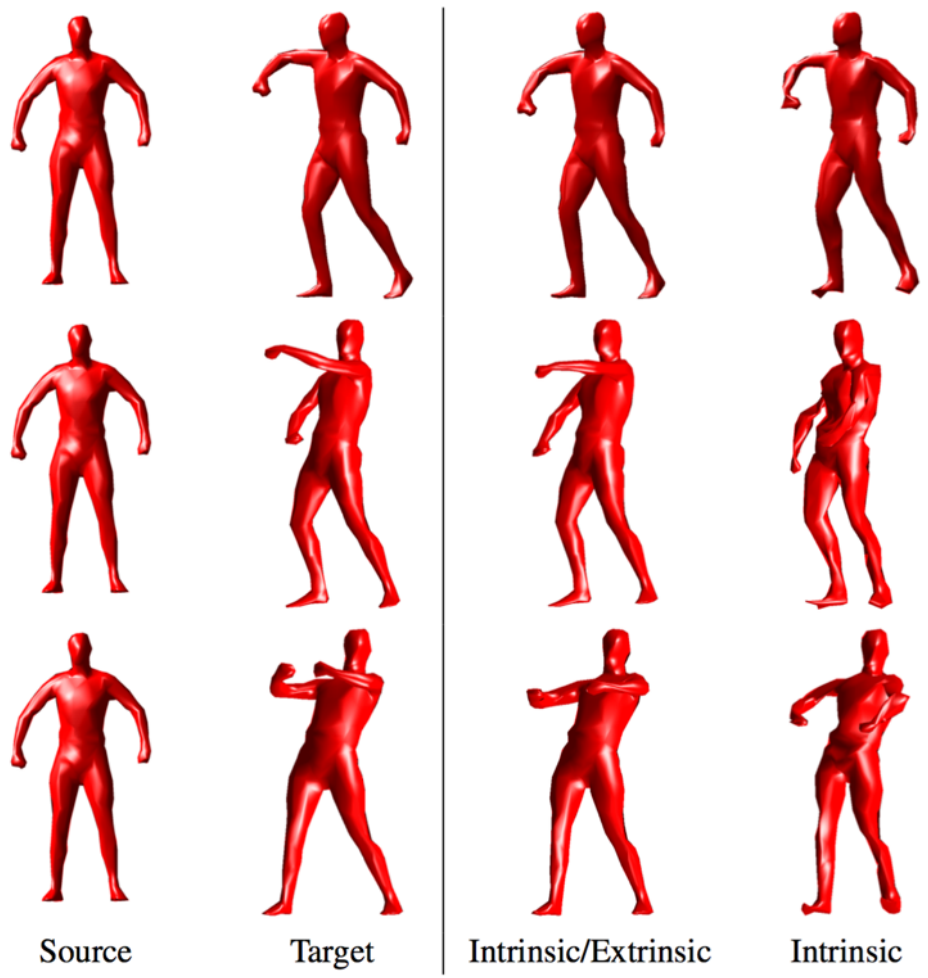

Shape Recovery from operators

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

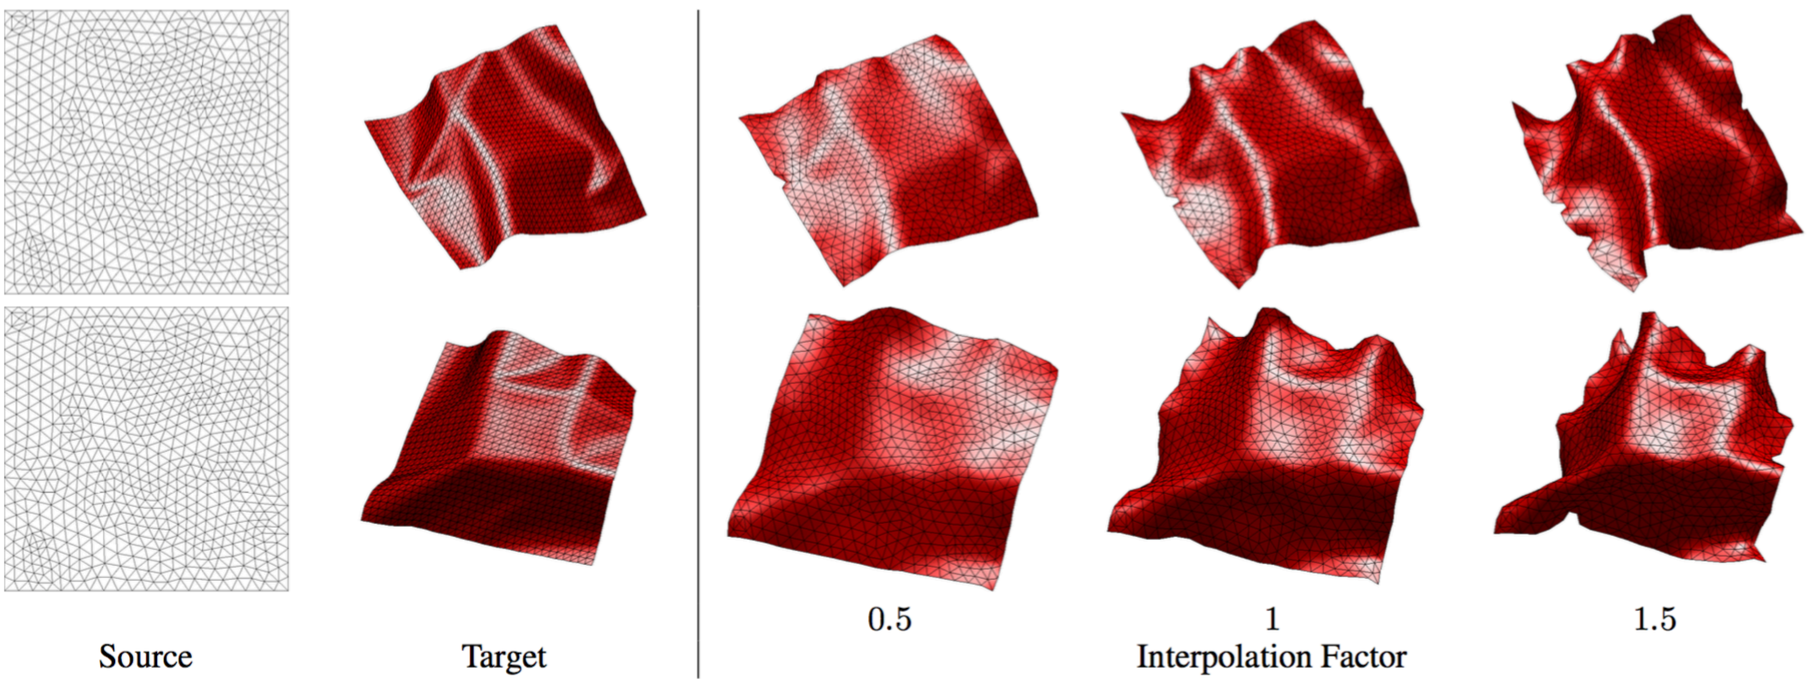

Shape Recovery from operators

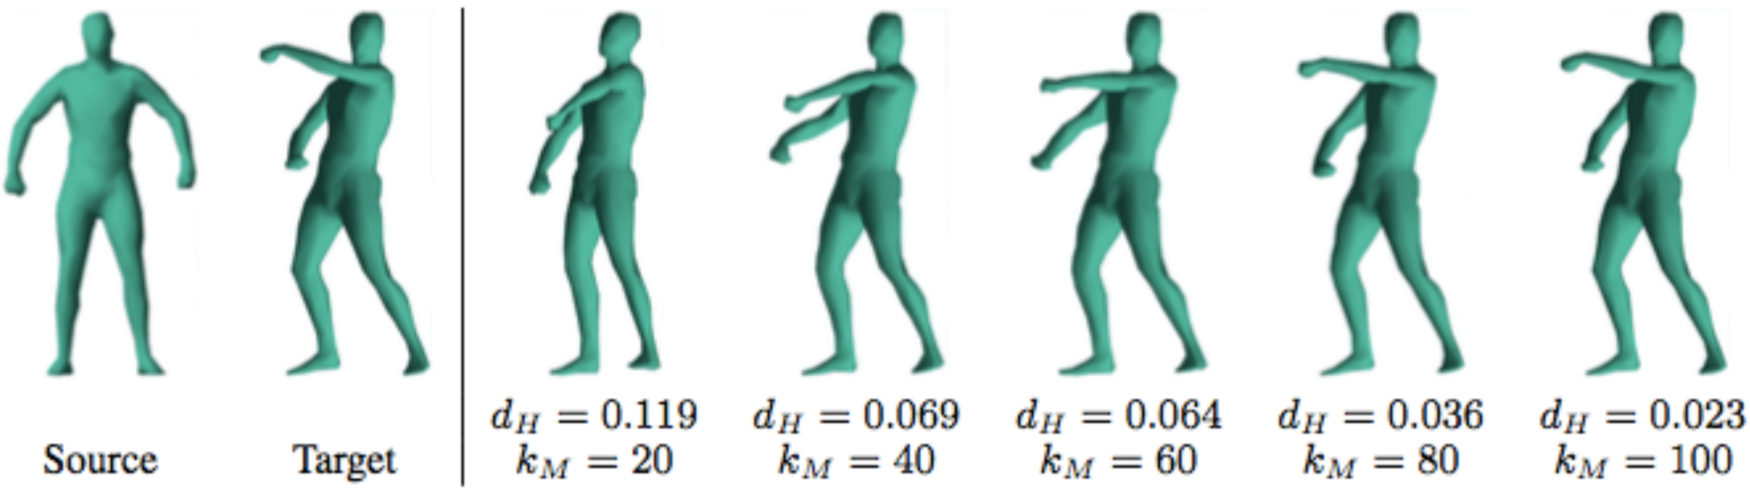

Conclusion

- Laplacian-based methods can be used for both similarity and difference (distortion).
- Can recover the metric from a Laplacian even in a noisy/approximate case.
- Shapes can be represented as sets of linear operators and recovered via “simple” optimization problems.
- Second fundamental form encoded via offsets.

Thank you!

Questions?