Correspondence Problems

Justin Solomon

6.838: Shape Analysis

Spring 2021

Surface Correspondence Problems

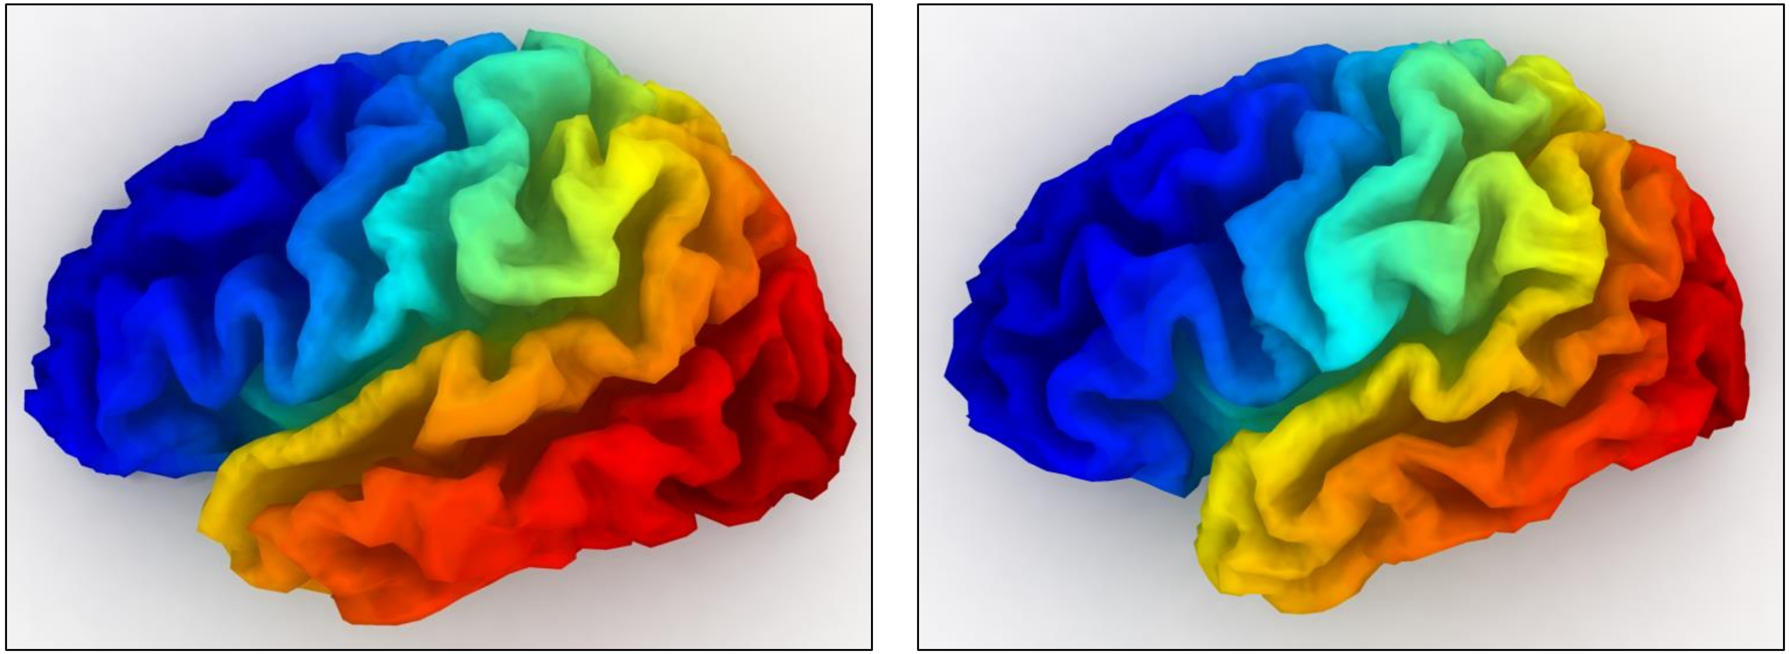

Which points on one object correspond to points on another?

Typical Distinction from Registration

Seek shared structure instead of alignment

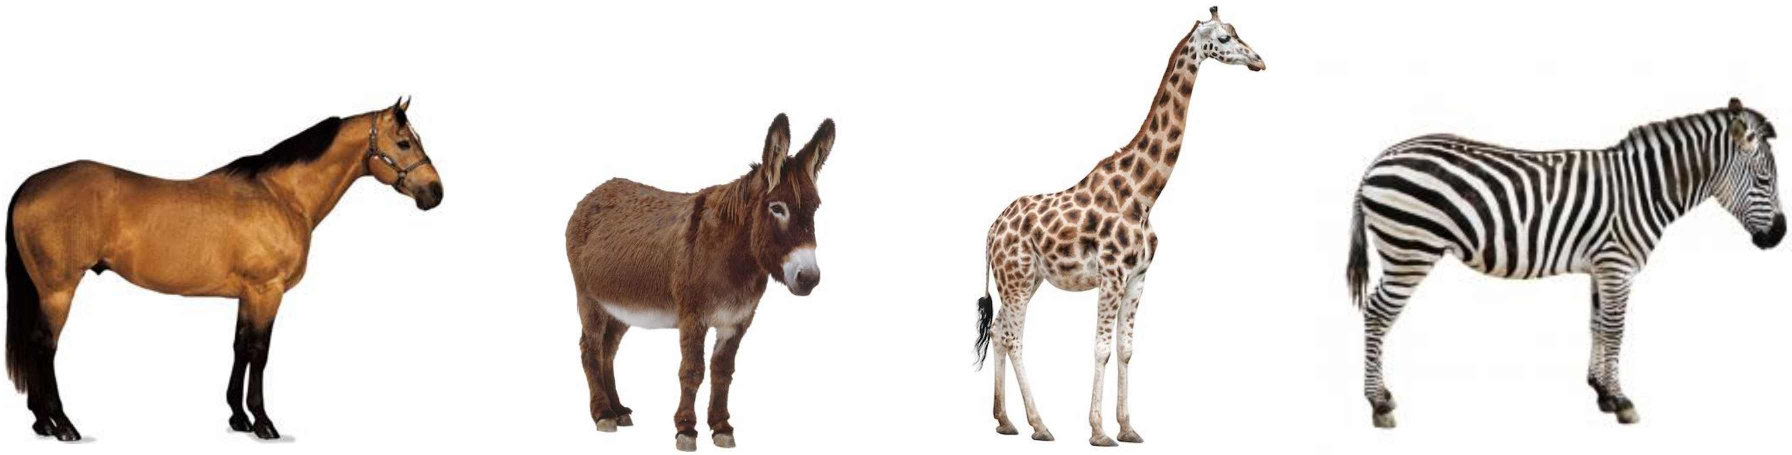

Applications

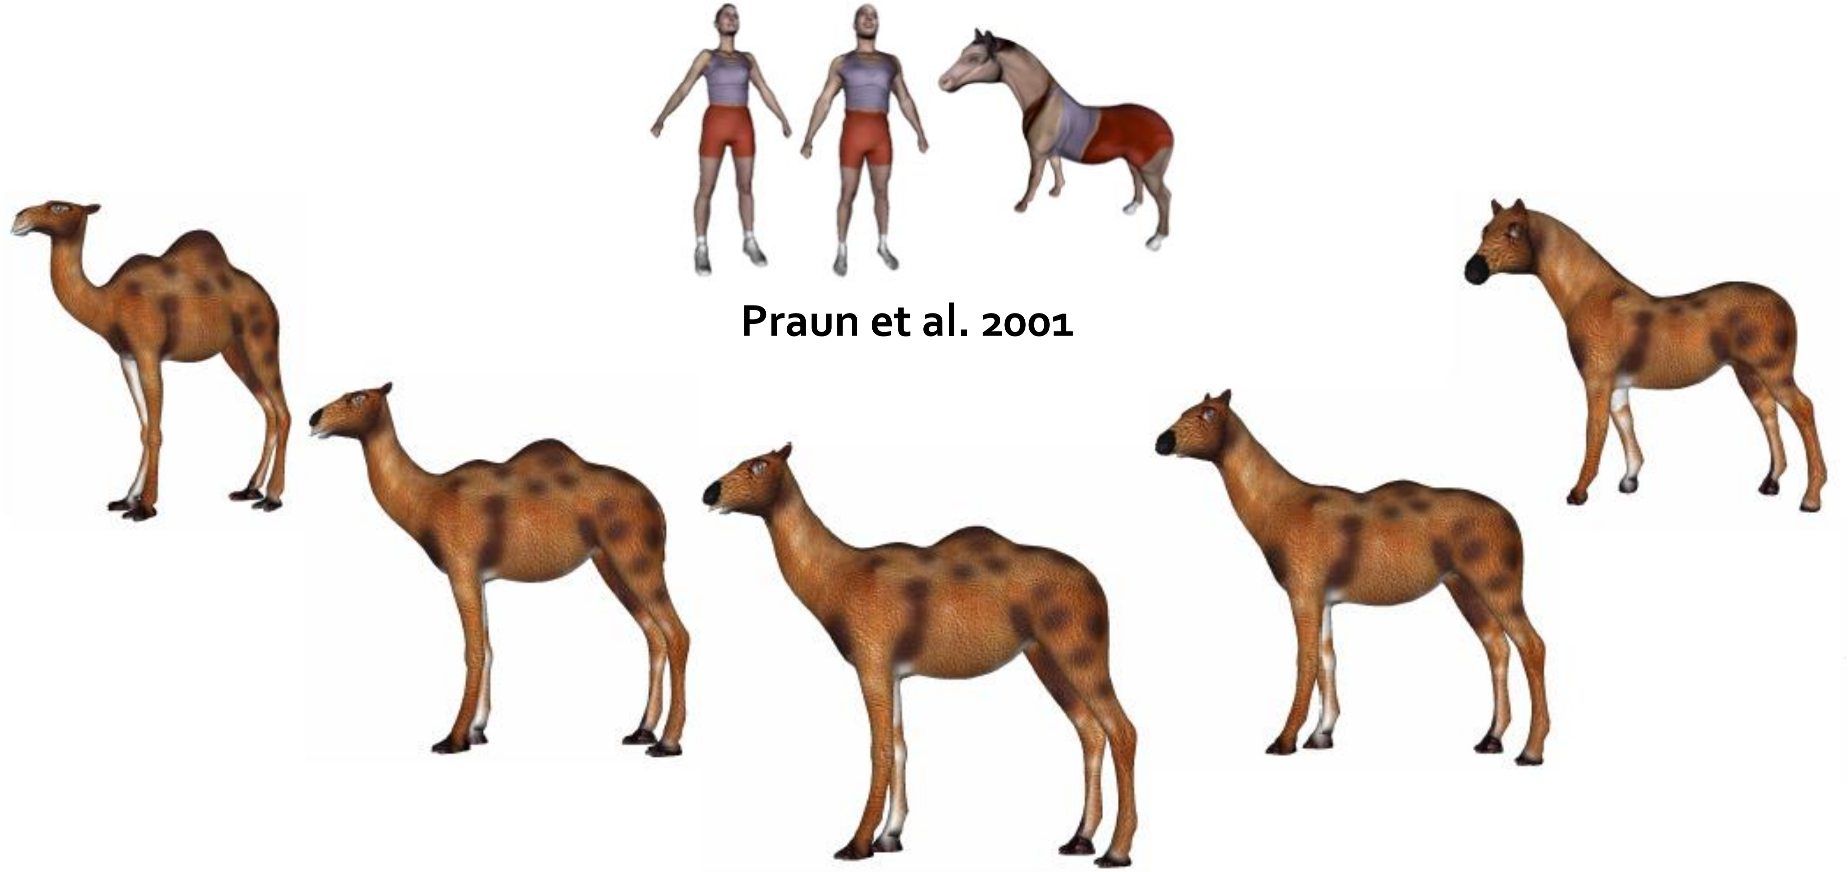

Praun et al. 2001

Kraevoy and Sheffer 2004

Texture transfer

Applications

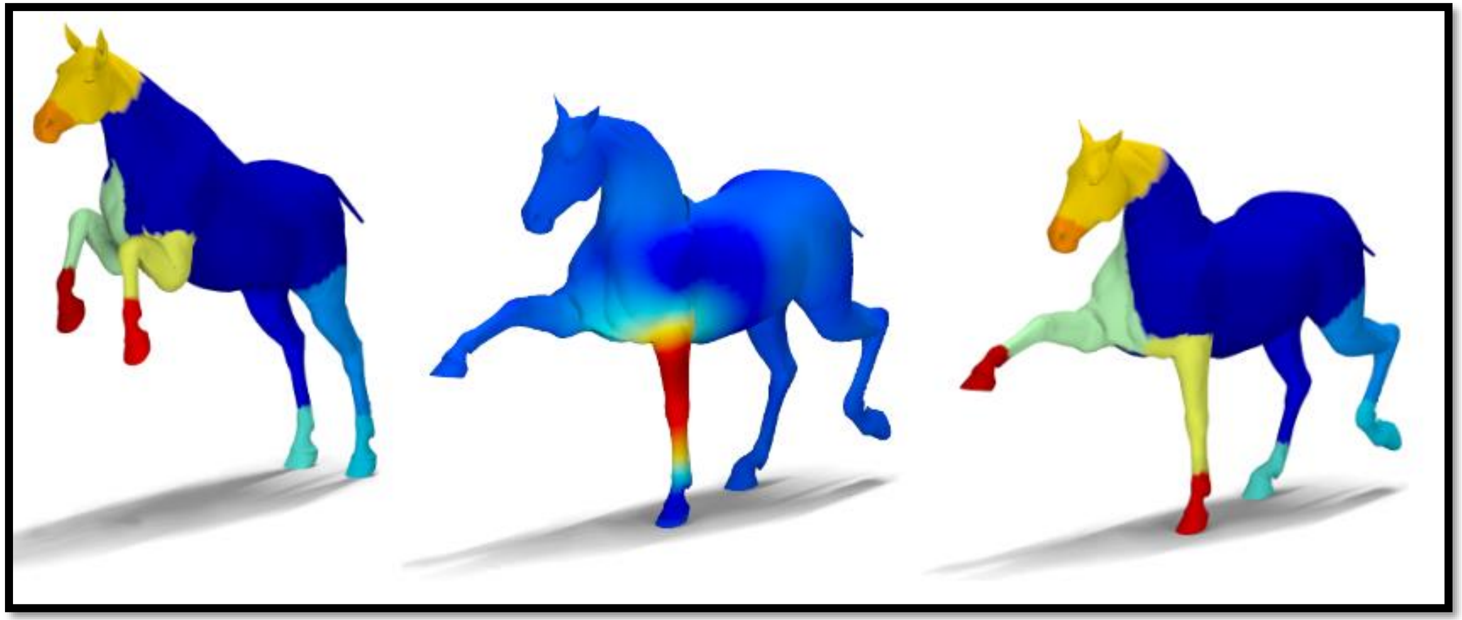

Ovsjanikov et al. 2012

Segmentation transfer

Applications

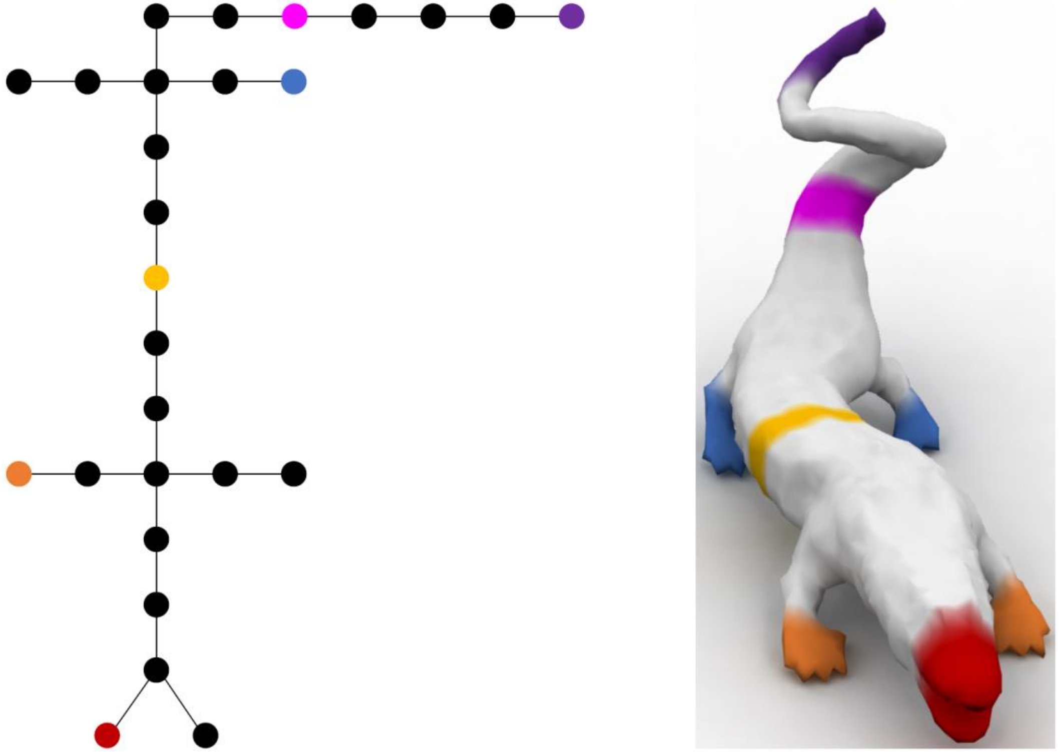

Solomon et al. 2016

Abstraction

Applications

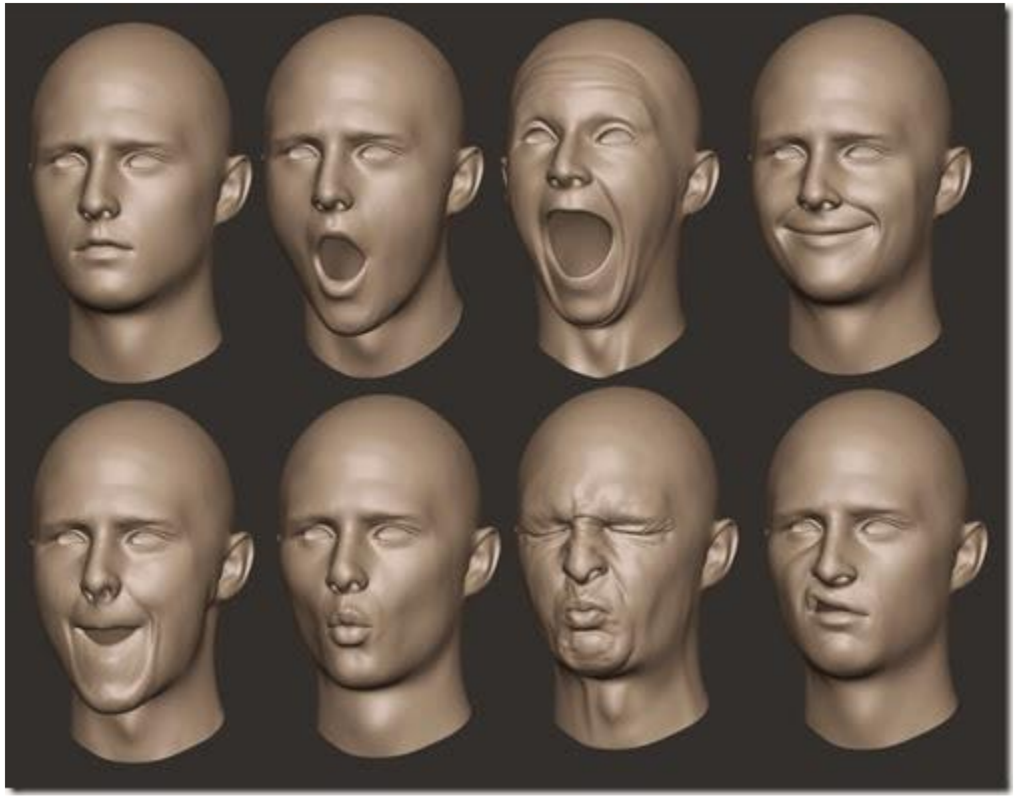

Image from "Shape Interpolations: Blendshape Math for Meshes" (https://graphicalanomaly.wordpress.com/)

Blendshape modeling

Applications

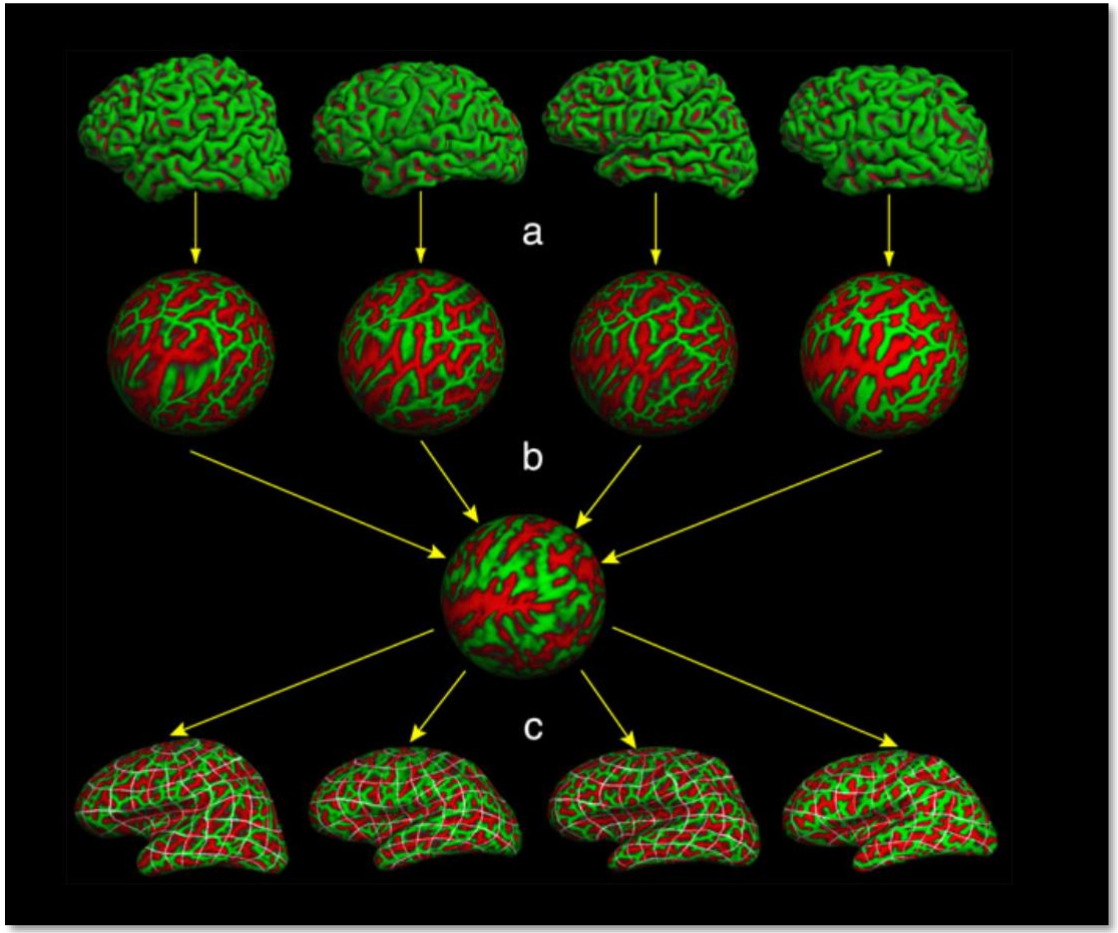

Image from “Freesurfer” (Wikipedia)

Statistical shape analysis

Applications

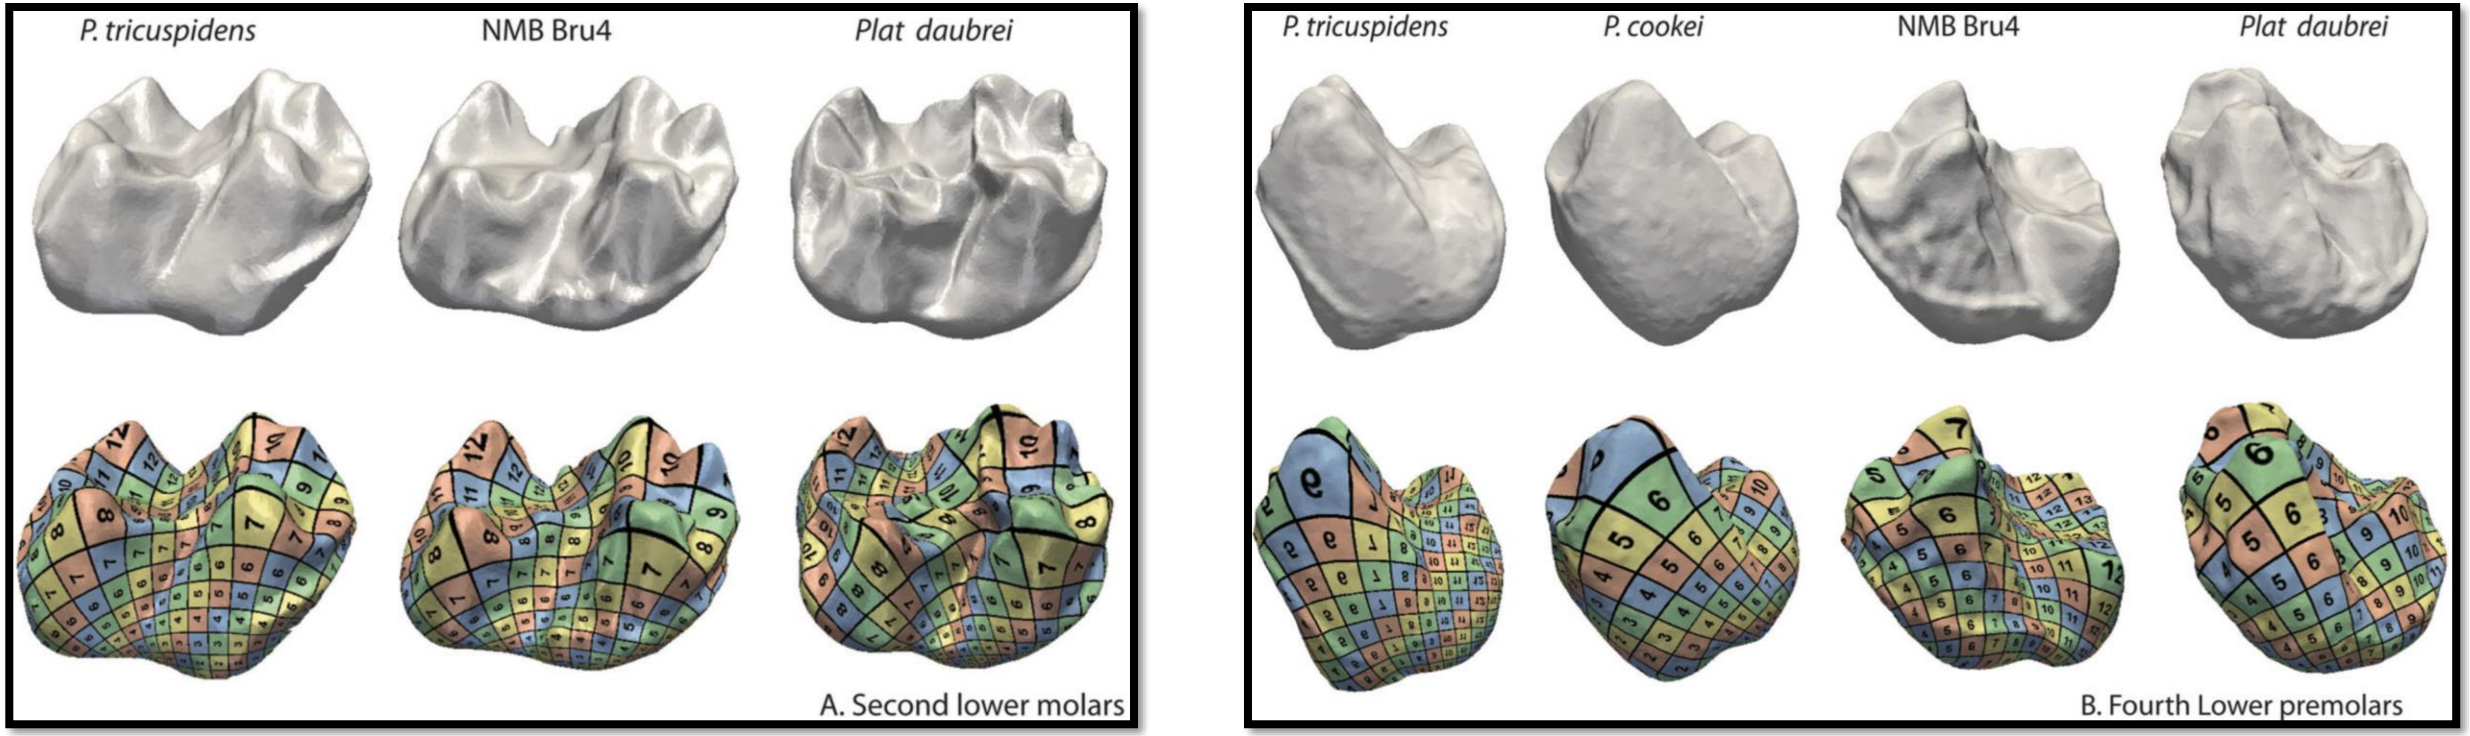

A. Second lower molars

B. Fourth Lower premolars

"Earliest Record of Platychoerops, A New Species From Mouras Quarry, Mont de Berru, France"
Boyer, Costeur, and Lipman 2012

Paleontology

Mapping problem

Given two (or more) shapes
Find a map f, satisfying the following properties:

- Fast to compute
- Bijective
  (if we expect global correspondence)
- Low-distortion
- Preserves important features

Adapted from slides by Q. Huang, V. Kim

Geometric Quality of Mappings

What do we need the map for?

Shape interpolation and texture transfer require highly accurate maps

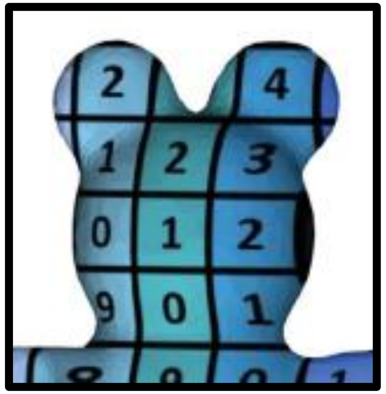

Target Texture
(projection)

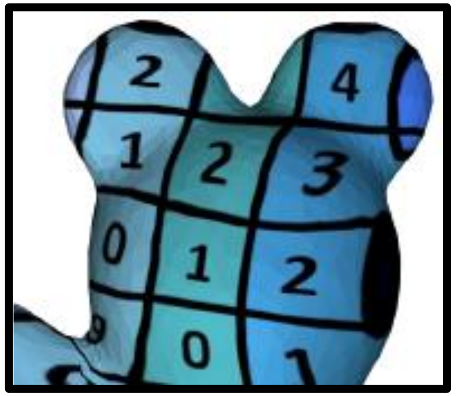

Locally and globally
accurate map

Globally accurate,
locally distorted map

Slide courtesy Danielle Ezuz

Geometric Quality of Mappings

How can we evaluate map quality?

Given a ground truth map, compute the cumulative error graph

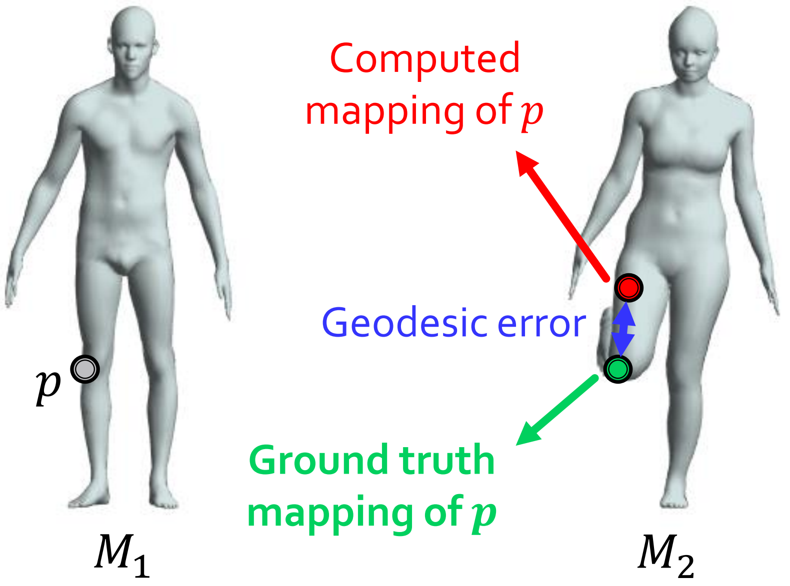

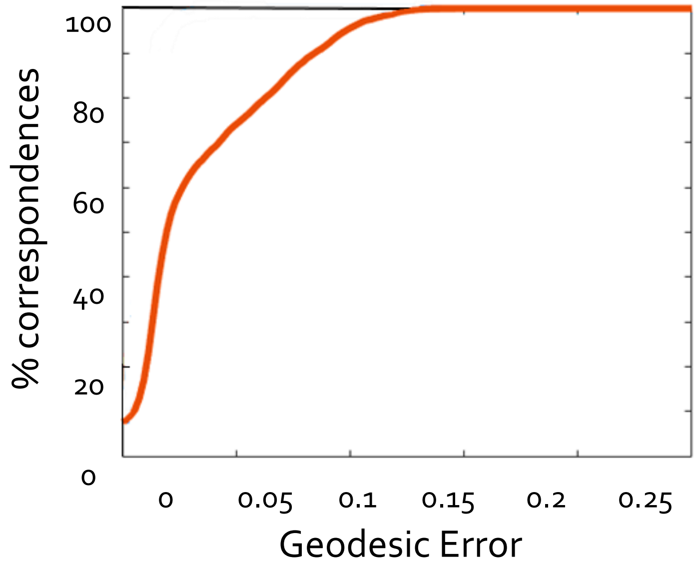

Slide courtesy Danielle Ezuz

Geometric Quality of Mappings

How can we evaluate map quality?

Given a ground truth map, compute the cumulative error graph

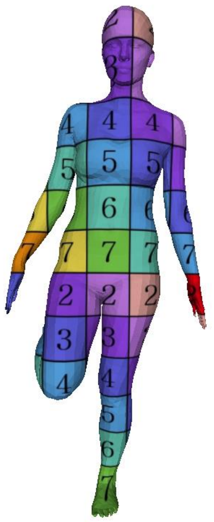

Target Texture

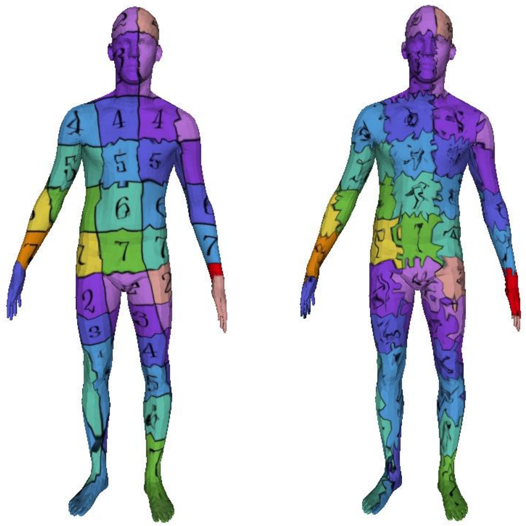

Which one is map #1?

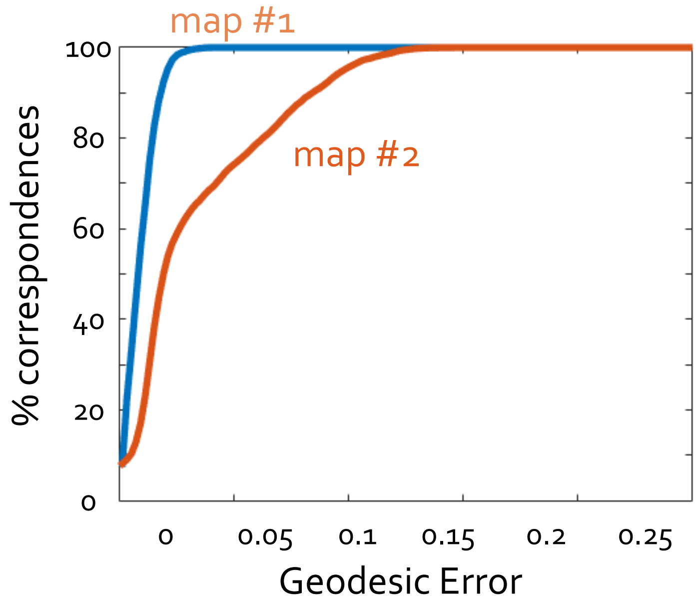

Slide courtesy Danielle Ezuz

Geometric Quality of Mappings

How can we evaluate map quality?

Given a ground truth map, compute the cumulative error graph

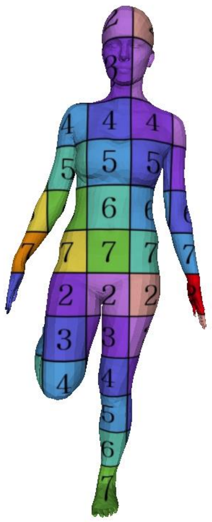

Target Texture

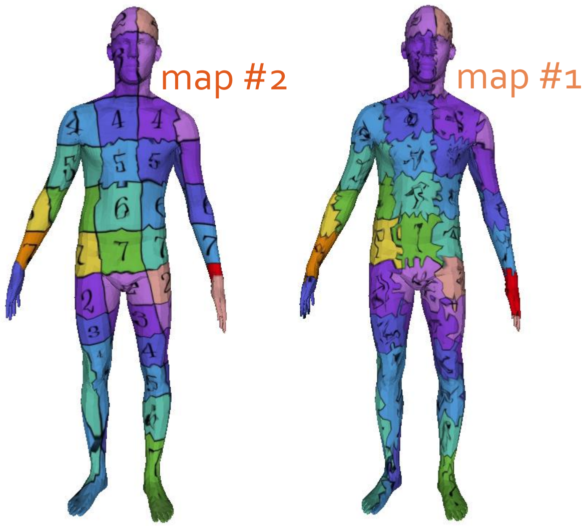

Which one is map #1?

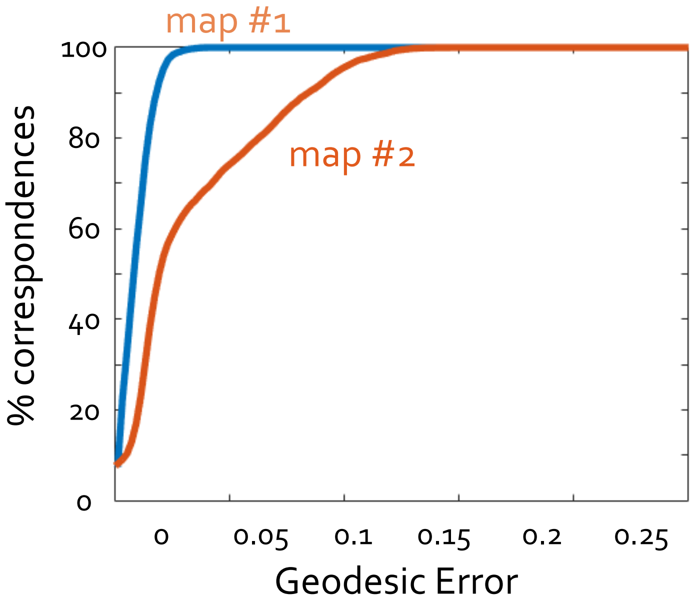

Slide courtesy Danielle Ezuz

Geometric Quality of Mappings

How can we evaluate map quality?

Measure *conformal distortion* (angle preservation)

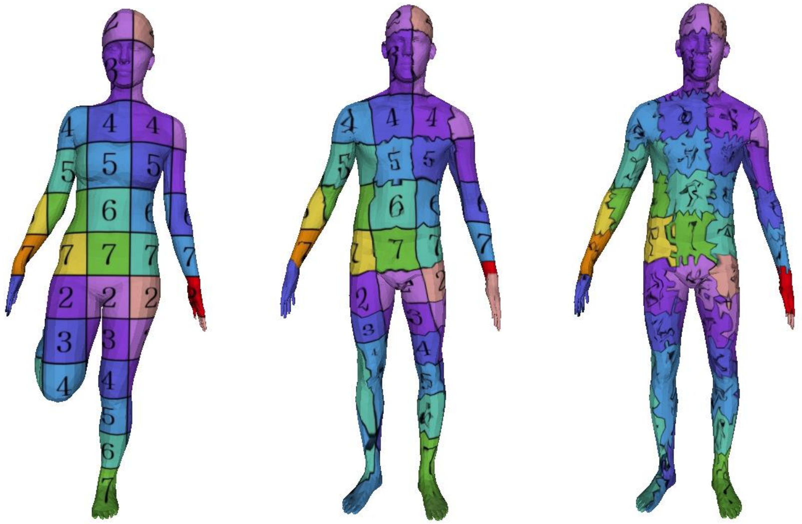

Target Texture

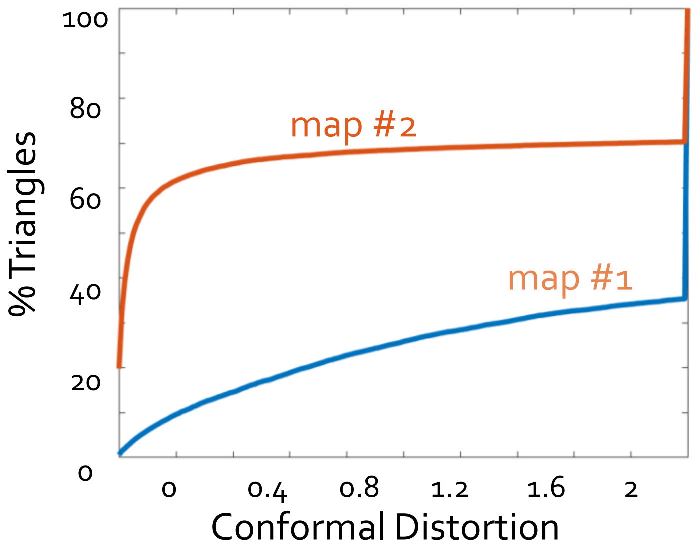

Slide courtesy Danielle Ezuz

Geometric Quality of Mappings

How can we evaluate map quality?

Measure *conformal distortion* (angle preservation)

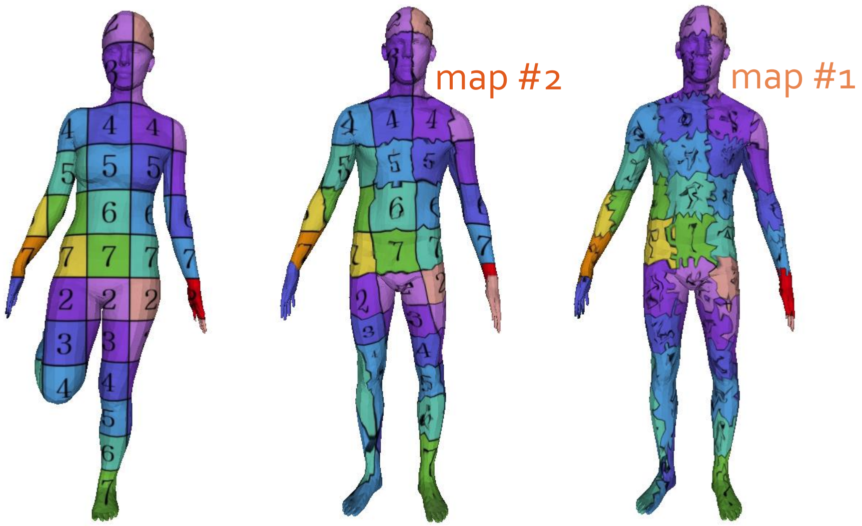

Target Texture

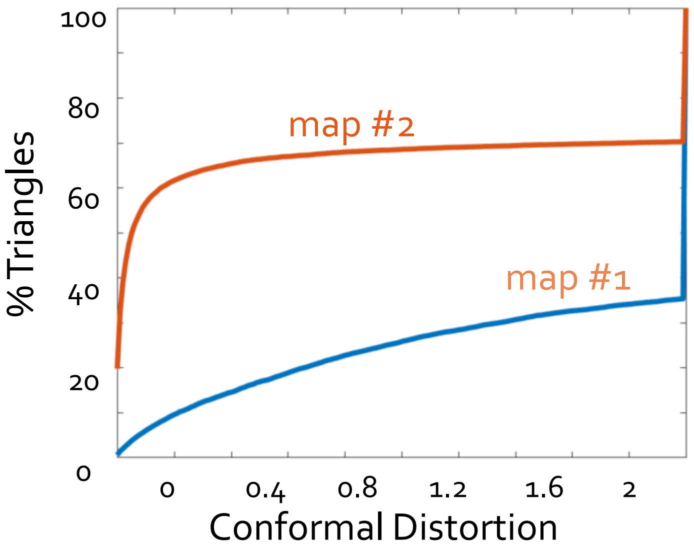

Slide courtesy Danielle Ezuz

Today's Plan

Sampling of surface mapping algorithms and models.

Graphics/vision bias!

Example: Consistent Remeshing (Co-Parameterization)

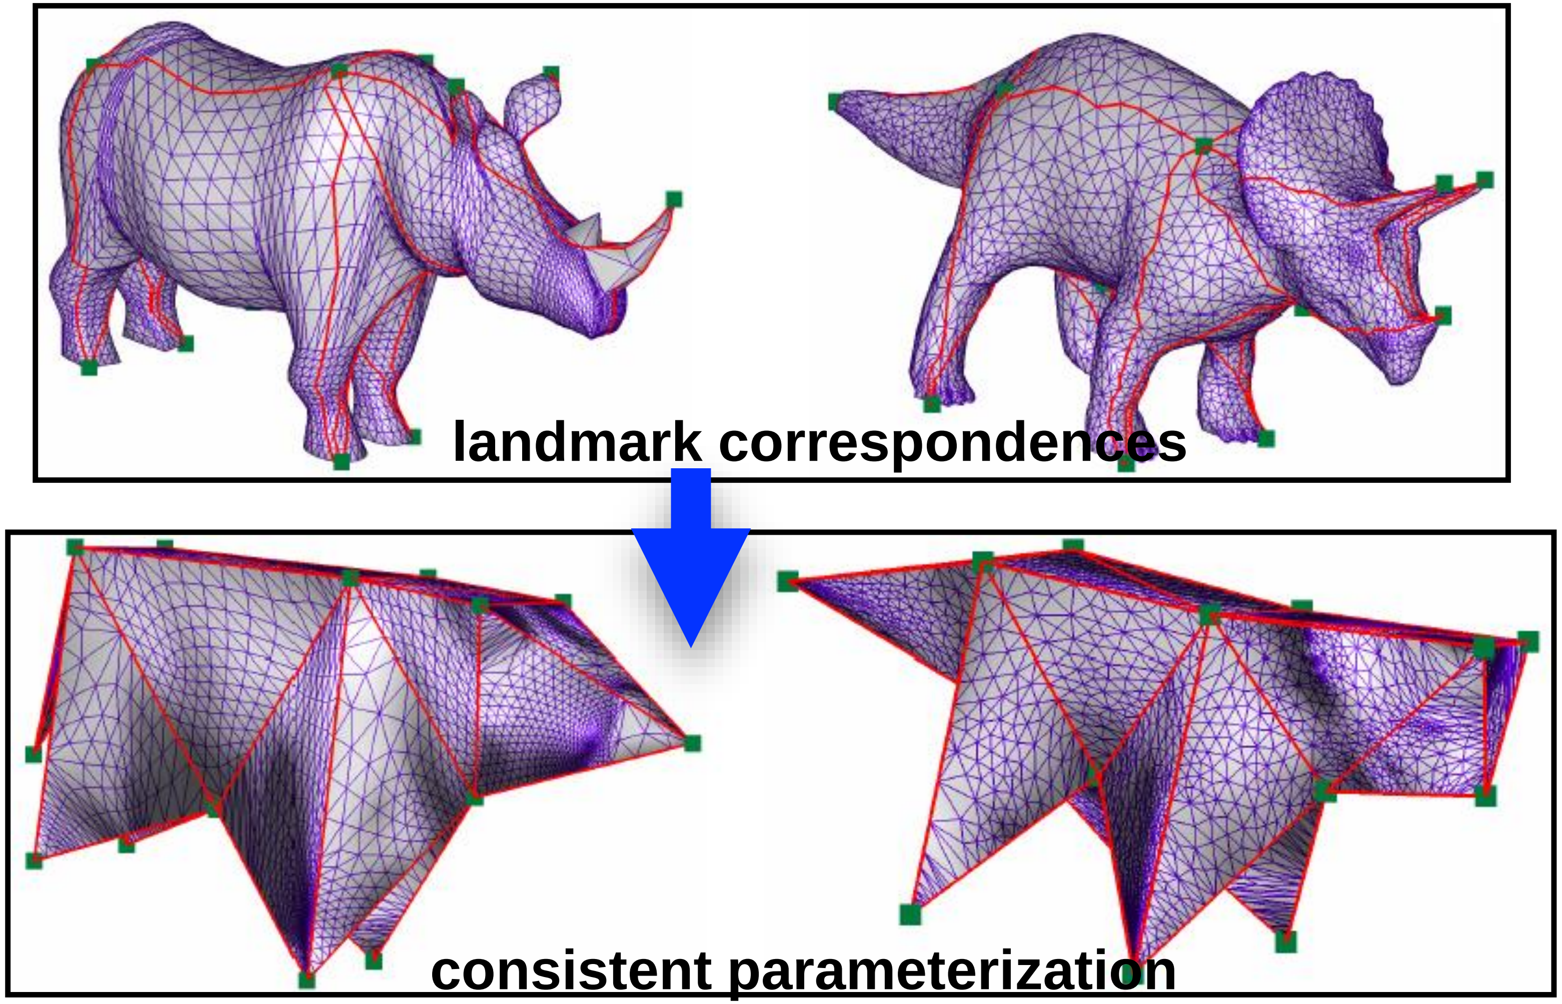
landmark correspondences

consistent parameterization

Kraevoy 2004

Adapted from slides by Q. Huang, V. Kim

Recall:

Example: Mesh Embedding

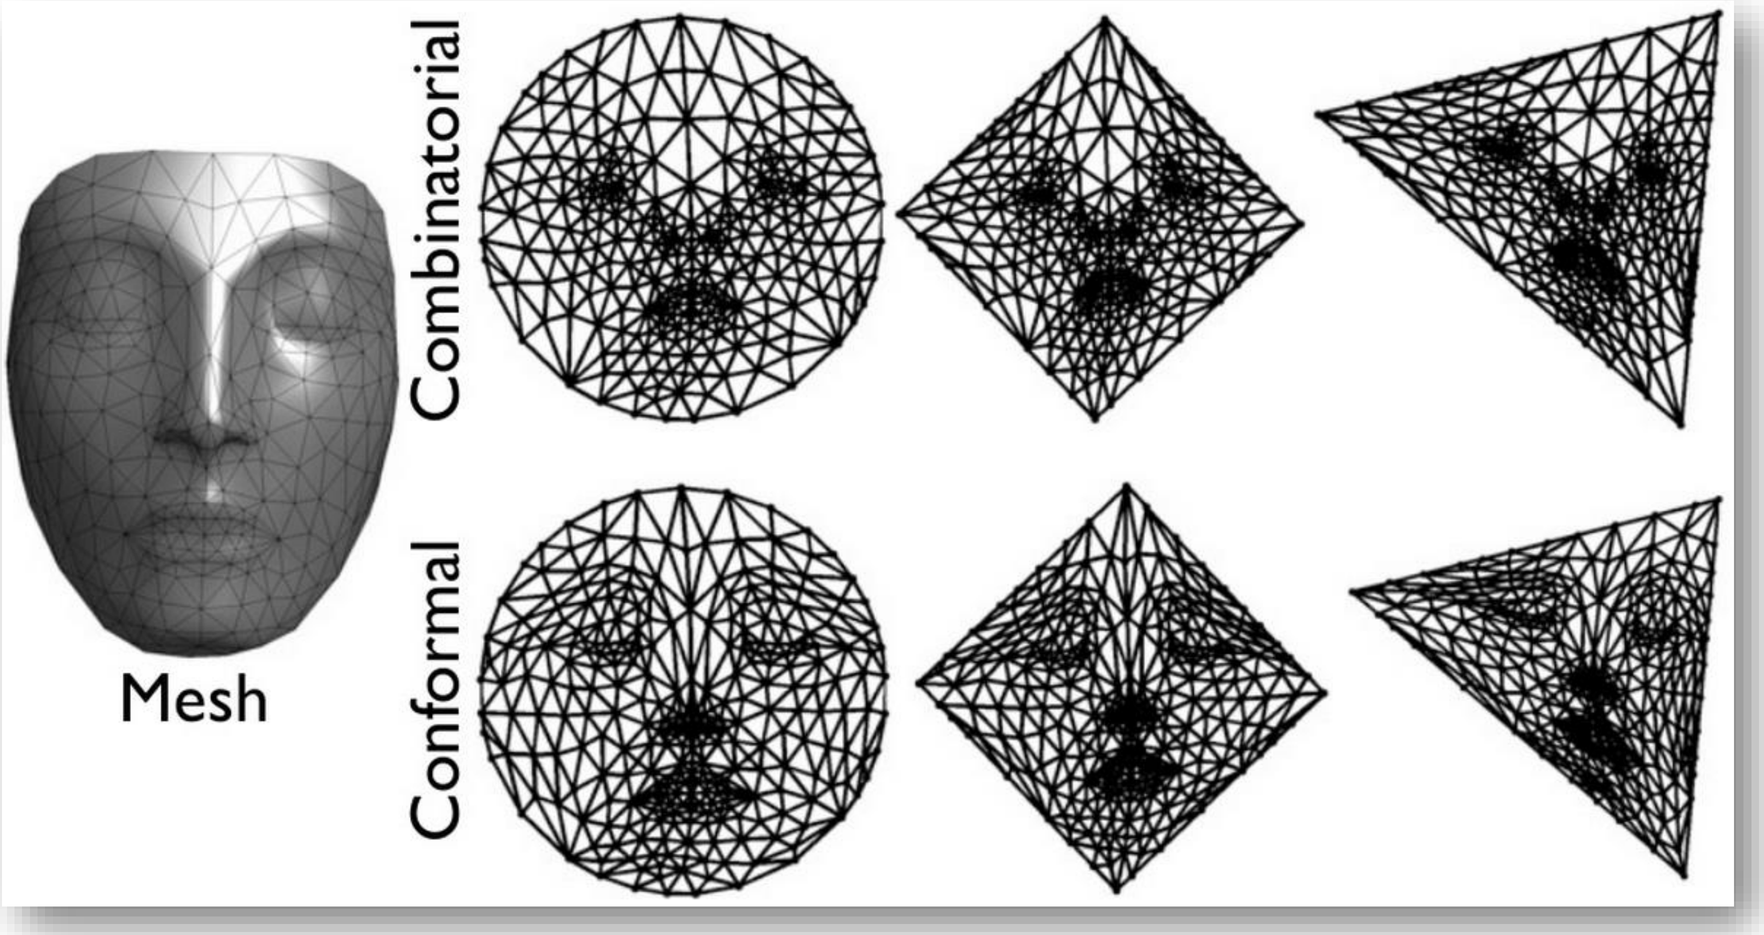

G. Peyré, mesh processing course slides

Recall:

Linear Solve for Embedding

$$\begin{aligned} \min_{\mathbf{x}_1,\ldots,\mathbf{x}_{|V|}} \quad & \sum_{(i,j)\in E} w_{ij}\|\mathbf{x}_i - \mathbf{x}_j\|_2^2 \\ \text{s.t.} \quad & \mathbf{x}_v \text{ fixed } \forall v \in V_0 \end{aligned}$$

- **$w_{ij} \equiv 1$:** Tutte embedding
- **w_{ij} from mesh:** Harmonic embedding

Assumption: w symmetric.

Tutte Embedding Theorem

$$\begin{array}{rl} \min_{\mathbf{x}_1,\ldots,\mathbf{x}_{|V|}} & \sum_{(i,j)\in E} w_{ij} \|\mathbf{x}_i - \mathbf{x}_j\|_2^2 \\ \text{s.t.} & \mathbf{x}_v \text{ fixed } \forall v \in V_0 \end{array}$$

Tutte embedding bijective if *w* nonnegative and boundary mapped to a convex polygon.

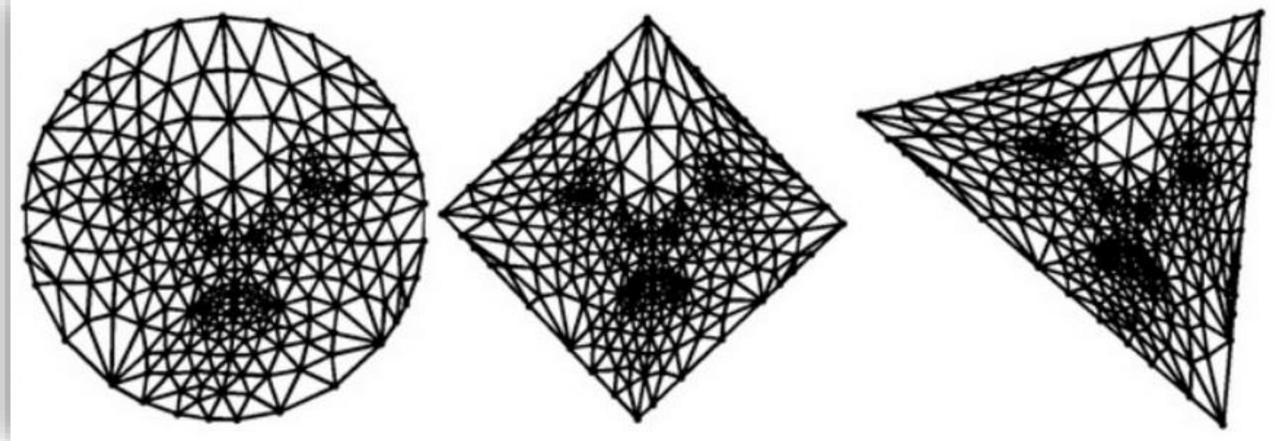

"How to draw a graph" (Proc. London Mathematical Society; Tutte, 1963)

Tradeoff: Consistent Remeshing

- **Pros:**
  - **Easy**
  - **Bijective**
- **Cons:**
  - **Need manual landmarks**
  - **Hard to minimize distortion**

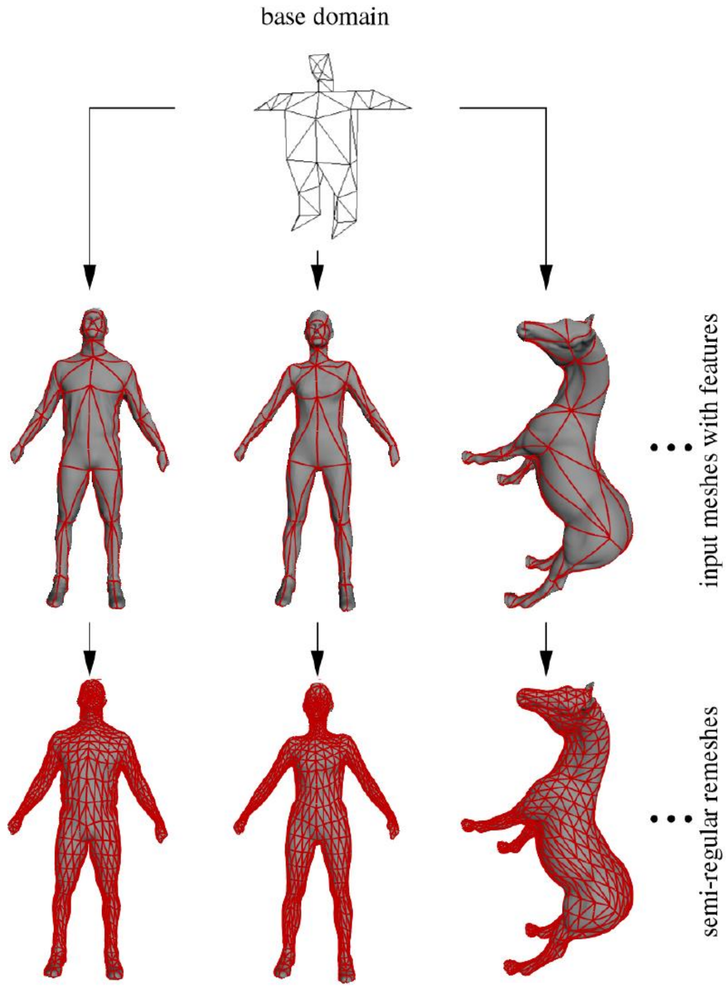

Praun et al. 2001

Adapted from slides by Q. Huang, V. Kim

Automatic Landmarks

- Simple algorithm:
  - Set landmarks
  - Measure energy
  - Repeat

- Possible metrics
  - Conformality
  - Area preservation
  - Stretch

E.g. small conformal distortion, large area distortion:

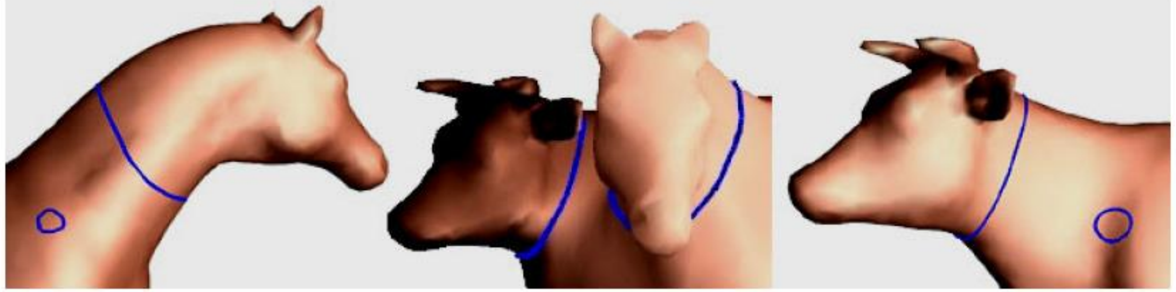

Schreiner et al. 2004

Adapted from slides by Q. Huang, V. Kim

Recent Coparameterization in Graphics

“Orbifold Tutte Embeddings” (Aigerman and Lipman, SIGGRAPH Asia 2015)

FreeSurfer: Spherical Coparameterization

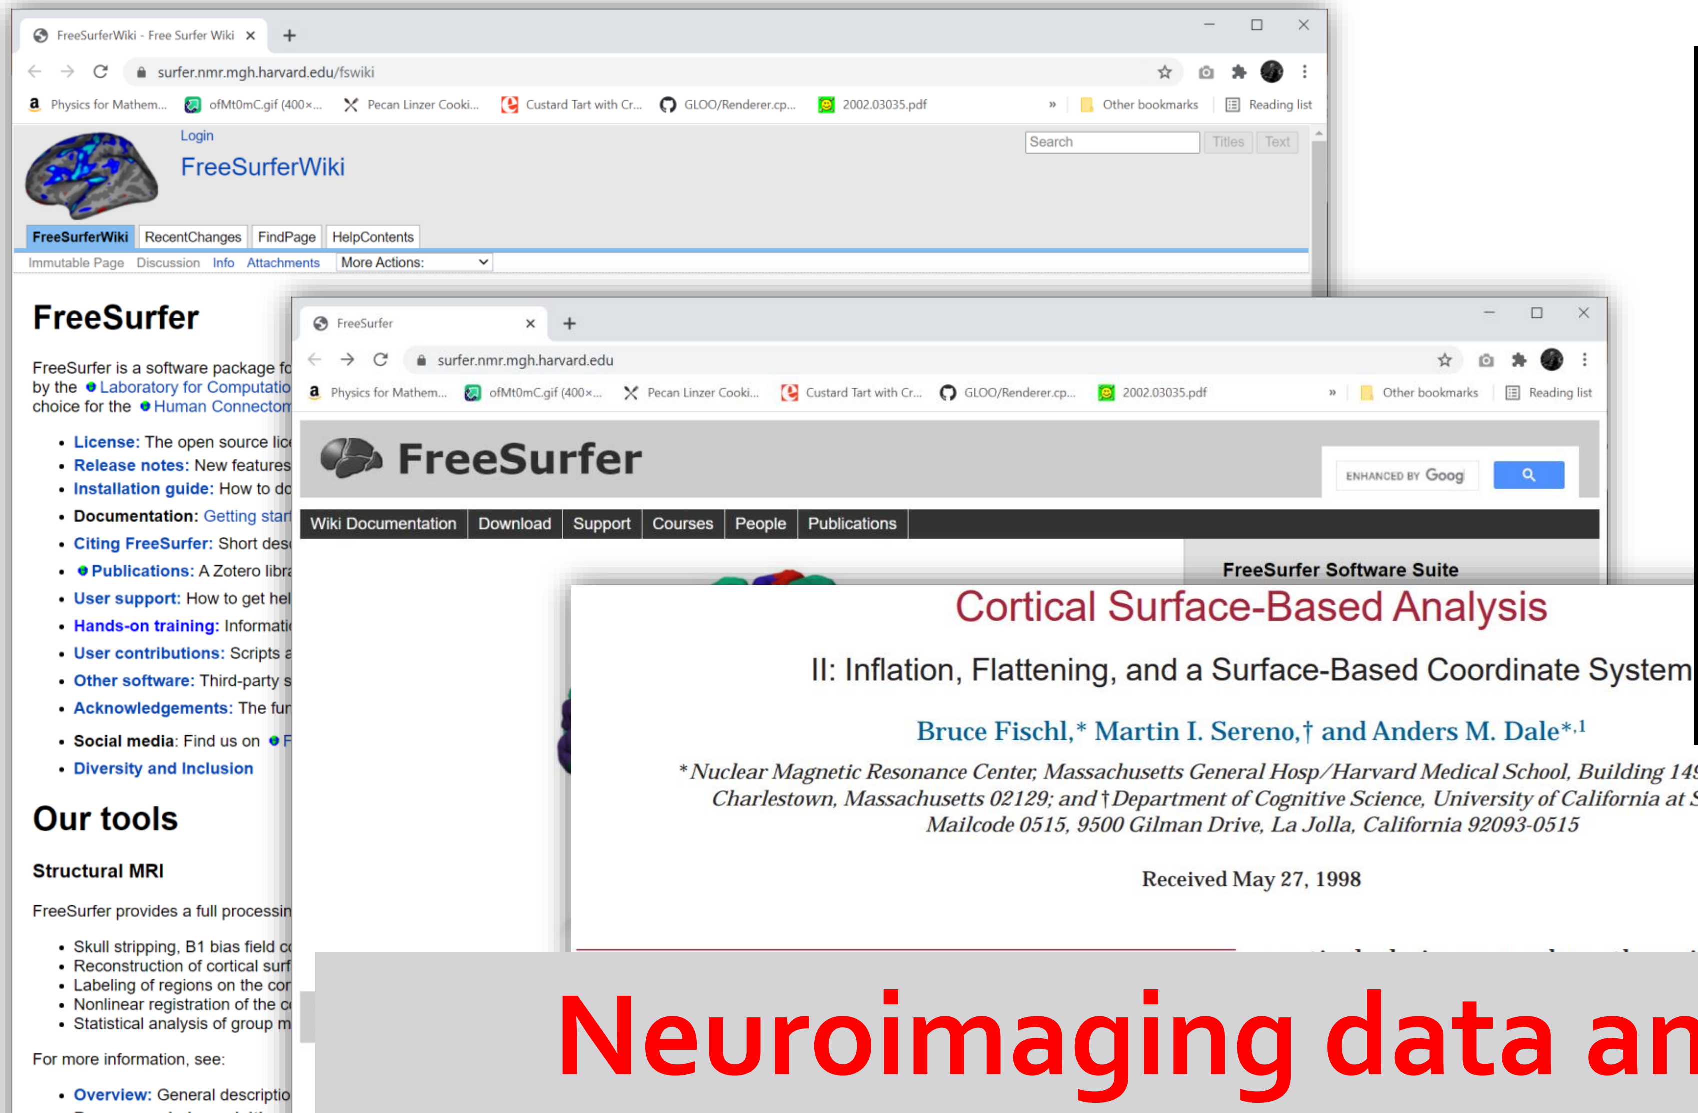

FreeSurfer

FreeSurfer is a software package fo
by the Laboratory for Computatio
choice for the Human Connectom

- **License:** The open source lic
- **Release notes:** New features
- **Installation guide:** How to do
- **Documentation:** Getting star
- **Citing FreeSurfer:** Short des
- **Publications:** A Zotero libra
- **User support:** How to get hel
- **Hands-on training:** Informati
- **User contributions:** Scripts a
- **Other software:** Third-party s
- **Acknowledgements:** The fu
- **Social media:** Find us on F
- **Diversity and Inclusion**

Our tools

Structural MRI

FreeSurfer provides a full processin

- Skull stripping, B1 bias field c
- Reconstruction of cortical surf
- Labeling of regions on the cor
- Nonlinear registration of the c
- Statistical analysis of group m

For more information, see:

- **Overview:** General descriptio

FreeSurfer Software Suite

Cortical Surface-Based Analysis

II: Inflation, Flattening, and a Surface-Based Coordinate System

Bruce Fischl,* Martin I. Sereno,† and Anders M. Dale*,[1]

Nuclear Magnetic Resonance Center, Massachusetts General Hosp/Harvard Medical School, Building 149, 13th Street, Charlestown, Massachusetts 02129; and †Department of Cognitive Science, University of California at San Diego, Mailcode 0515, 9500 Gilman Drive, La Jolla, California 92093-0515

Received May 27, 1998

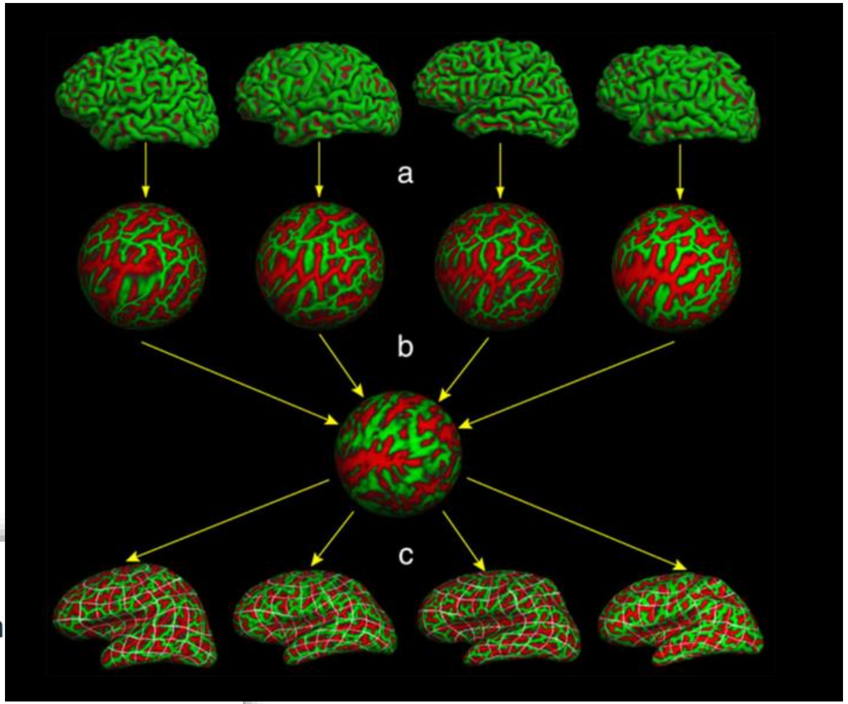

Neuroimaging data analysis

Digression:

Related Problem

Mapping specifically into the plane

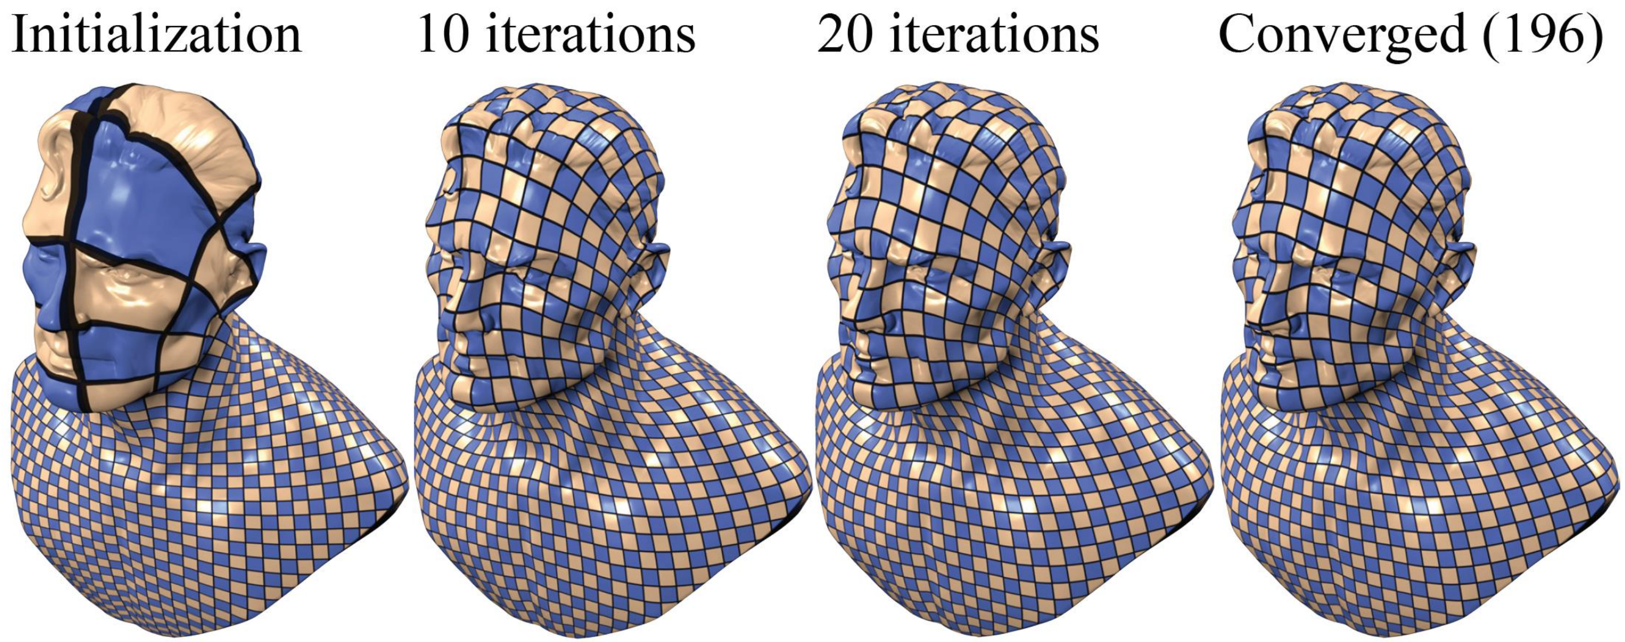

Image from "Scalable Locally Injective Mappings" (Rabinovich et al., 2017)

Parameterization

Local Distortion Measure

Tutte distortion:

$$\begin{array}{rl} \min_{\mathbf{x}_1,\ldots,\mathbf{x}_{|V|}} & \sum_{(i,j)\in E} w_{ij}\|\mathbf{x}_i - \mathbf{x}_j\|_2^2 \\ \text{s.t.} & \mathbf{x}_v \text{ fixed } \forall v \in V_0 \end{array}$$

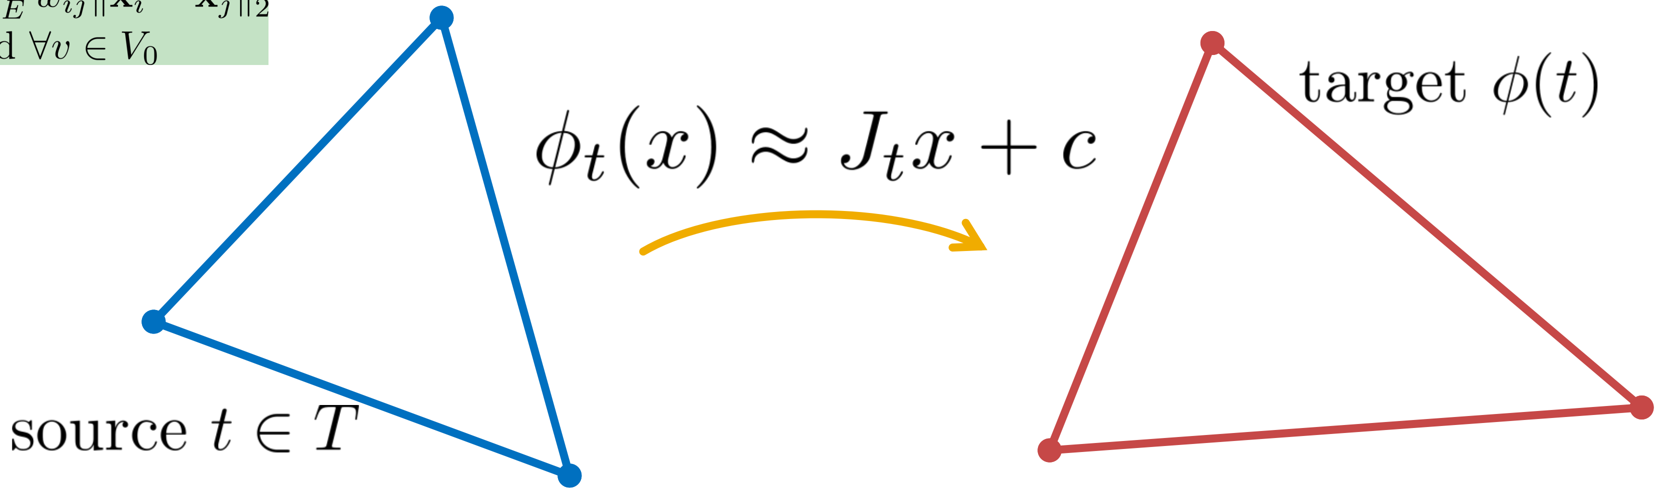

$$\text{Distortion} := \sum_{t\in T} A_t \mathcal{D}(J_t)$$

Triangle distortion measure

Notation from Rabinovich et al. 2017

Q:

How do you measure distortion of a triangle?

Typical Distortion Measures

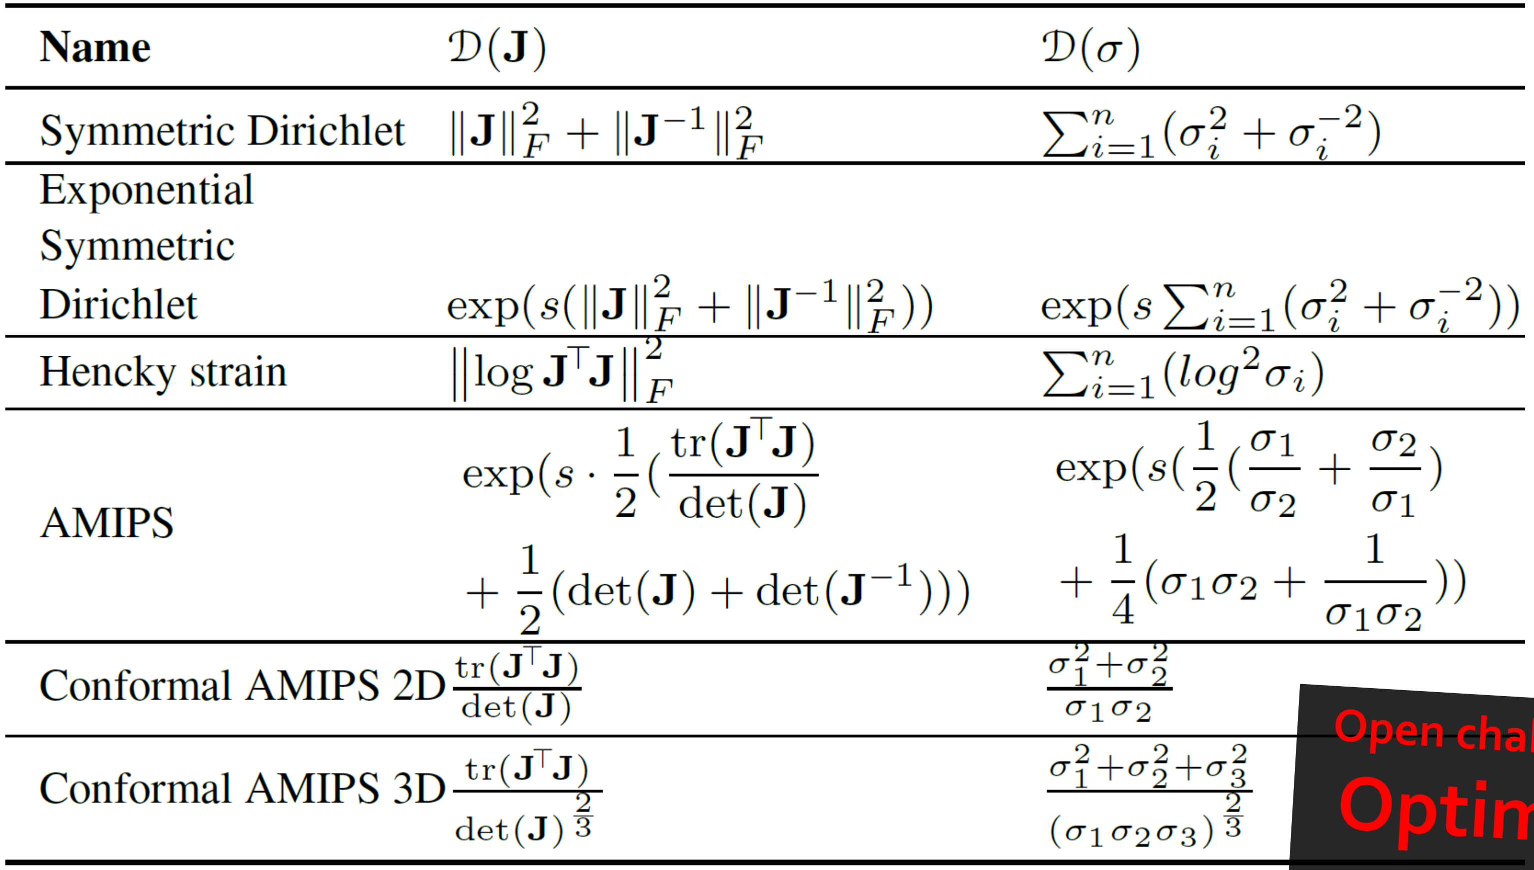

Name	$\mathcal{D}(\mathbf{J})$	$\mathcal{D}(\sigma)$
Symmetric Dirichlet	$\lVert\mathbf{J}\rVert_F^2 + \lVert\mathbf{J}^{-1}\rVert_F^2$	$\sum_{i=1}^n(\sigma_i^2 + \sigma_i^{-2})$
Exponential Symmetric Dirichlet	$\exp(s(\lVert\mathbf{J}\rVert_F^2 + \lVert\mathbf{J}^{-1}\rVert_F^2))$	$\exp(s\sum_{i=1}^n(\sigma_i^2 + \sigma_i^{-2}))$
Hencky strain	$\left\lVert\log \mathbf{J}^\top\mathbf{J}\right\rVert_F^2$	$\sum_{i=1}^n(log^2\sigma_i)$
AMIPS	$\exp(s \cdot \frac{1}{2}(\frac{\mathrm{tr}(\mathbf{J}^\top\mathbf{J})}{\det(\mathbf{J})} + \frac{1}{2}(\det(\mathbf{J}) + \det(\mathbf{J}^{-1})))$	$\exp(s(\frac{1}{2}(\frac{\sigma_1}{\sigma_2} + \frac{\sigma_2}{\sigma_1}) + \frac{1}{4}(\sigma_1\sigma_2 + \frac{1}{\sigma_1\sigma_2}))$
Conformal AMIPS 2D	$\frac{\mathrm{tr}(\mathbf{J}^\top\mathbf{J})}{\det(\mathbf{J})}$	$\frac{\sigma_1^2+\sigma_2^2}{\sigma_1\sigma_2}$
Conformal AMIPS 3D	$\frac{\mathrm{tr}(\mathbf{J}^\top\mathbf{J})}{\det(\mathbf{J})^{\frac{2}{3}}}$	$\frac{\sigma_1^2+\sigma_2^2+\sigma_3^2}{(\sigma_1\sigma_2\sigma_3)^{\frac{2}{3}}}$

Open challenge:
Optimize directly

Table from “Scalable Locally Injective Mappings” (Rabinovich et al., 2017)

End-to-End Coparameterization

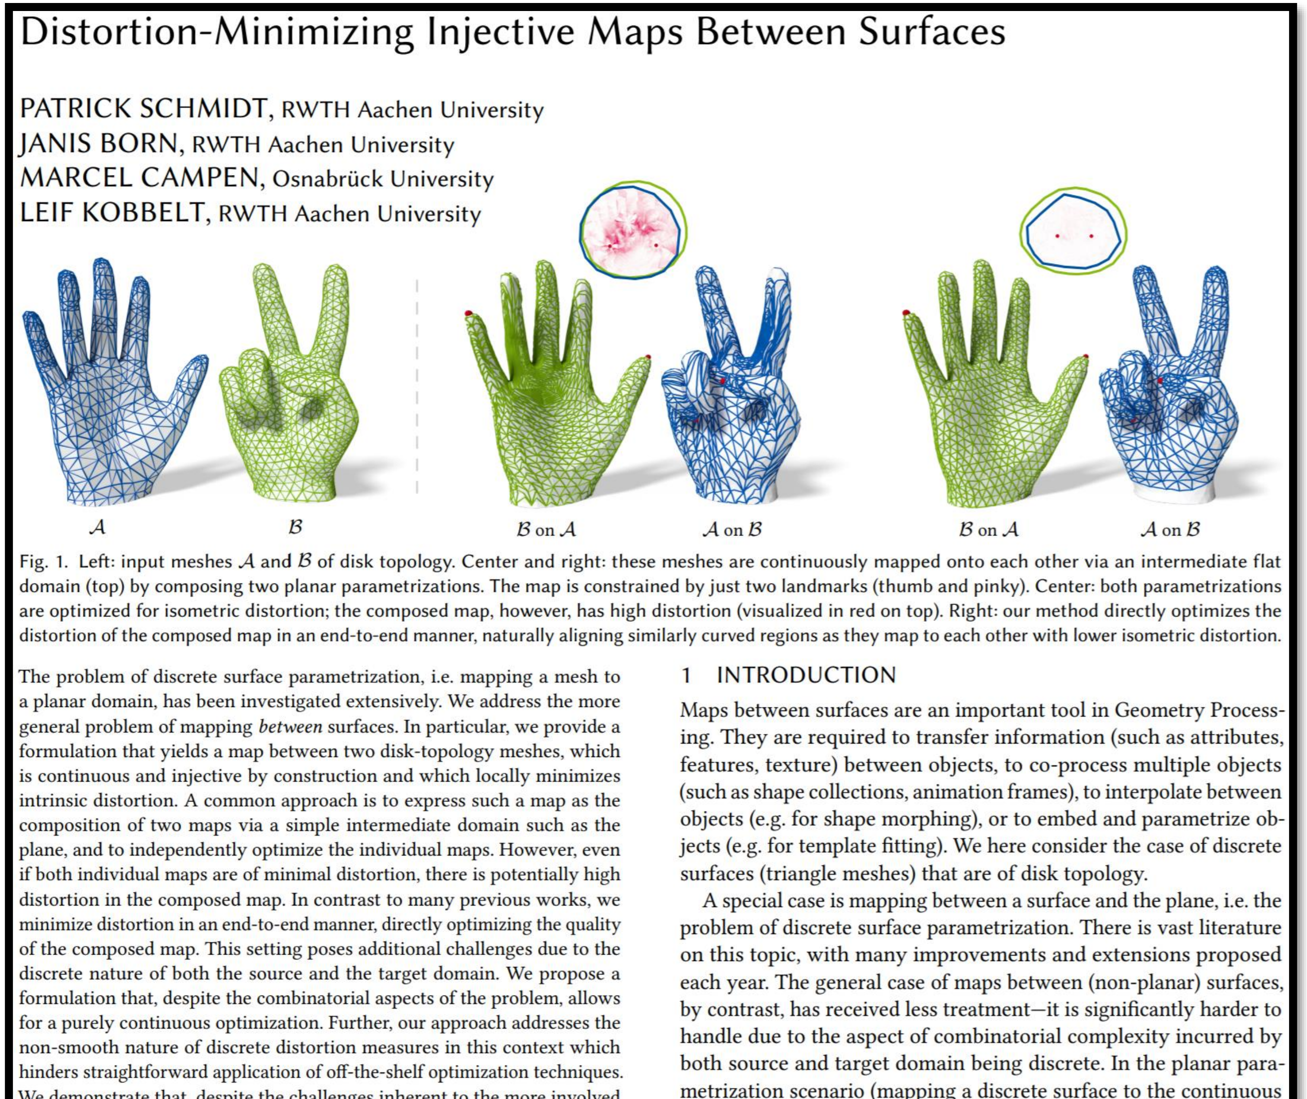

Distortion-Minimizing Injective Maps Between Surfaces

PATRICK SCHMIDT, RWTH Aachen University
JANIS BORN, RWTH Aachen University
MARCEL CAMPEN, Osnabrück University
LEIF KOBBELT, RWTH Aachen University

Fig. 1. Left: input meshes $\mathcal{A}$ and $\mathcal{B}$ of disk topology. Center and right: these meshes are continuously mapped onto each other via an intermediate flat domain (top) by composing two planar parametrizations. The map is constrained by just two landmarks (thumb and pinky). Center: both parametrizations are optimized for isometric distortion; the composed map, however, has high distortion (visualized in red on top). Right: our method directly optimizes the distortion of the composed map in an end-to-end manner, naturally aligning similarly curved regions as they map to each other with lower isometric distortion.

The problem of discrete surface parametrization, i.e. mapping a mesh to a planar domain, has been investigated extensively. We address the more general problem of mapping *between* surfaces. In particular, we provide a formulation that yields a map between two disk-topology meshes, which is continuous and injective by construction and which locally minimizes intrinsic distortion. A common approach is to express such a map as the composition of two maps via a simple intermediate domain such as the plane, and to independently optimize the individual maps. However, even if both individual maps are of minimal distortion, there is potentially high distortion in the composed map. In contrast to many previous works, we minimize distortion in an end-to-end manner, directly optimizing the quality of the composed map. This setting poses additional challenges due to the discrete nature of both the source and the target domain. We propose a formulation that, despite the combinatorial aspects of the problem, allows for a purely continuous optimization. Further, our approach addresses the non-smooth nature of discrete distortion measures in this context which hinders straightforward application of off-the-shelf optimization techniques. We demonstrate that, despite the challenges inherent to the more involved

1 INTRODUCTION

Maps between surfaces are an important tool in Geometry Processing. They are required to transfer information (such as attributes, features, texture) between objects, to co-process multiple objects (such as shape collections, animation frames), to interpolate between objects (e.g. for shape morphing), or to embed and parametrize objects (e.g. for template fitting). We here consider the case of discrete surfaces (triangle meshes) that are of disk topology.

A special case is mapping between a surface and the plane, i.e. the problem of discrete surface parametrization. There is vast literature on this topic, with many improvements and extensions proposed each year. The general case of maps between (non-planar) surfaces, by contrast, has received less treatment—it is significantly harder to handle due to the aspect of combinatorial complexity incurred by both source and target domain being discrete. In the planar parametrization scenario (mapping a discrete surface to the continuous

Back to Correspondence: New Idea

Not all calculations have to be at the triangle level!

Long-distance interactions

can stabilize geometric computations.

Gromov-Hausdorff Distance

Distance between metric spaces X, Y

$$d_{\mathrm{GH}}(X,Y) := \inf_{\phi : X \to Y} \sup_{x,x' \in X} |d_X(x,x') - d_Y(\phi(x),\phi(x'))|$$

Best map

Worst distortion

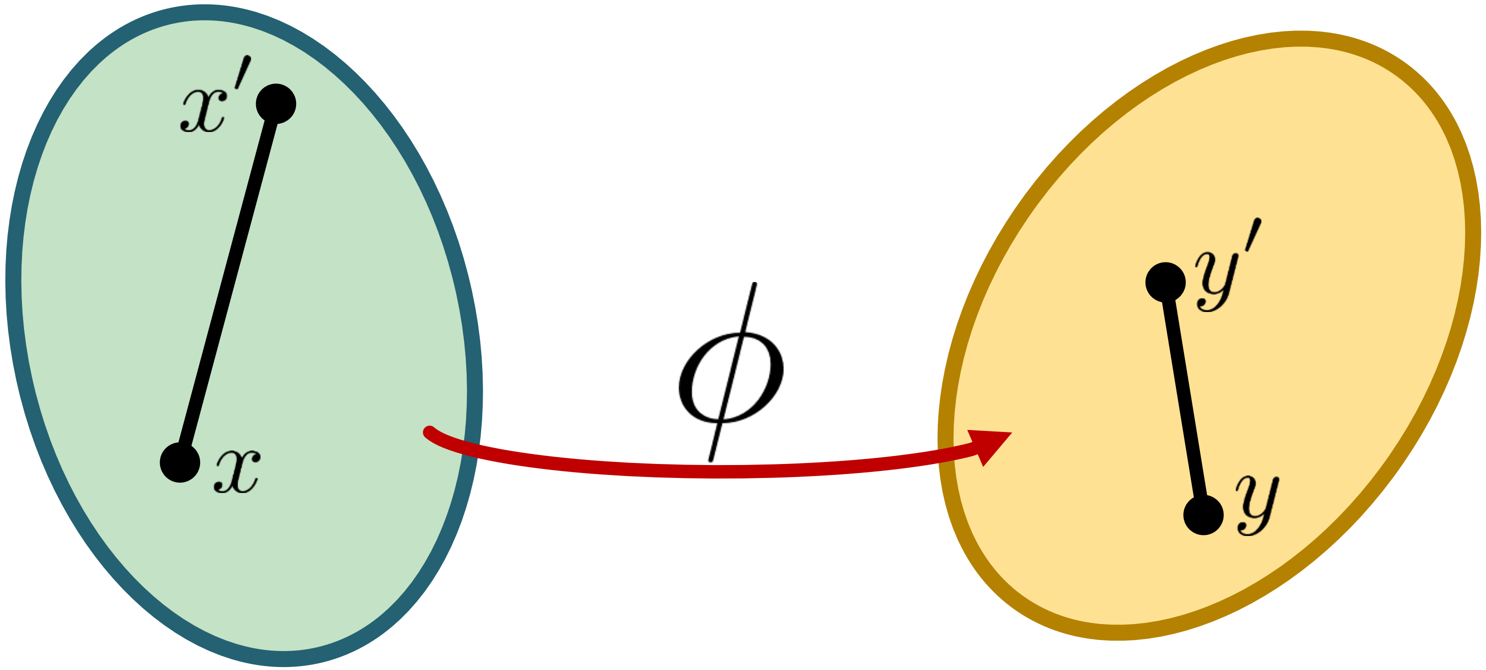

Recall:

Classical Multidimensional Scaling

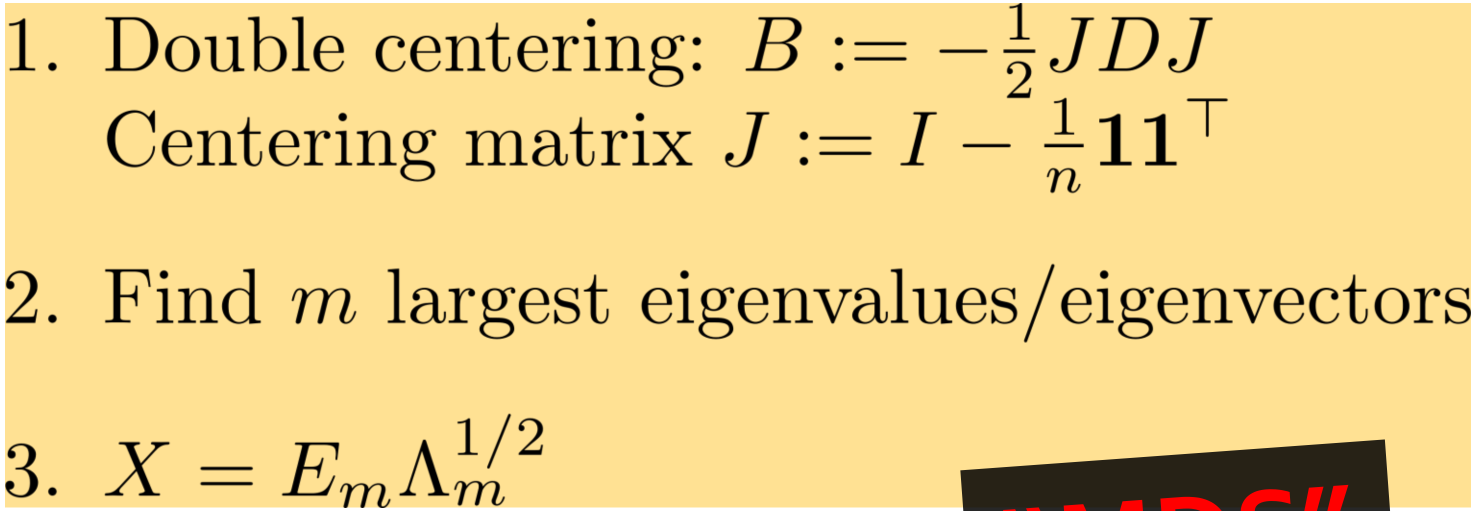

1. Double centering: $B := -\frac{1}{2} JDJ$
   Centering matrix $J := I - \frac{1}{n}\mathbf{1}\mathbf{1}^\top$

2. Find m largest eigenvalues/eigenvectors

3. $X = E_m \Lambda_m^{1/2}$

"MDS"

Torgerson, Warren S. (1958). *Theory & Methods of Scaling*.

Generalized MDS

Search for a permutation!

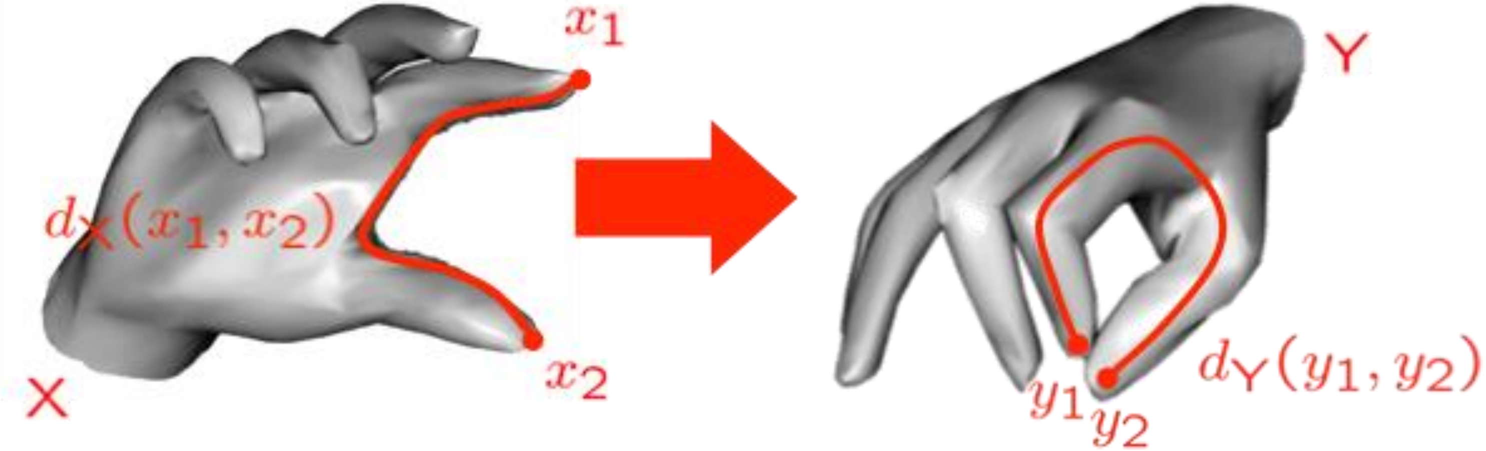

$$d_{\text{int}}(X, Y) := \min_{\{y_1,\ldots,y_n\} \subset Y} \|d_X(x_i, x_j) - d_Y(y_i, y_j)\|$$

Bronstein, Bronstein, and Kimmel; PNAS 2006

Problem: Quadratic Assignment

$$\begin{aligned}\min_T \quad & \langle M_0 T, T M_1 \rangle \\ \text{s.t.} \quad & T \in \{0,1\}^{n\times n} \\ & T\mathbf{1} = p_0 \\ & T^\top \mathbf{1} = p_1\end{aligned}$$

Nonconvex quadratic program!
NP-hard!

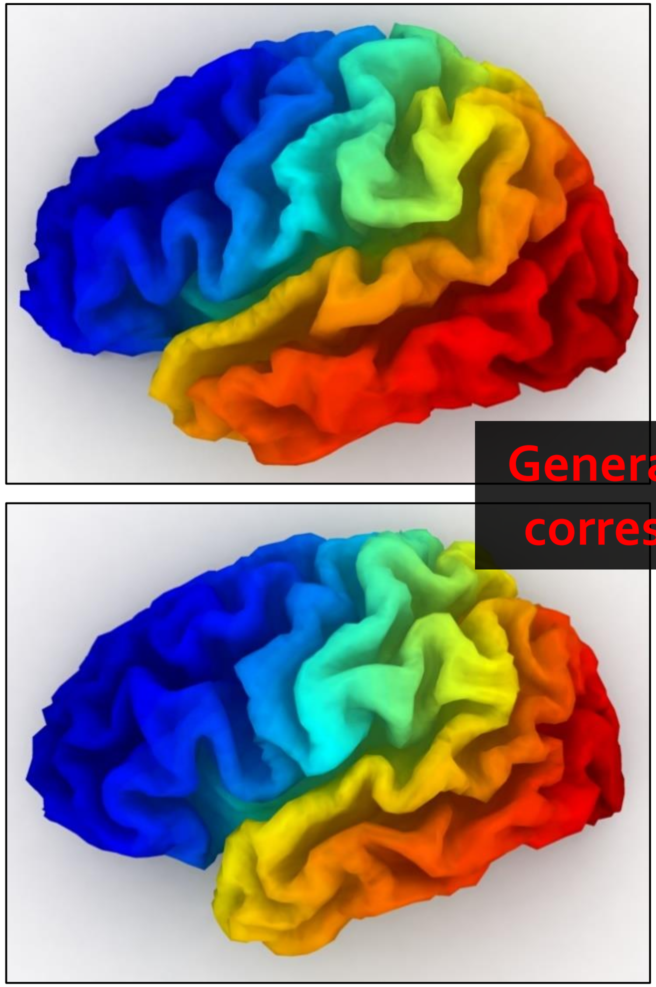

General notion of correspondence

What’s Wrong?

- Hard to optimize
- Multiple optima

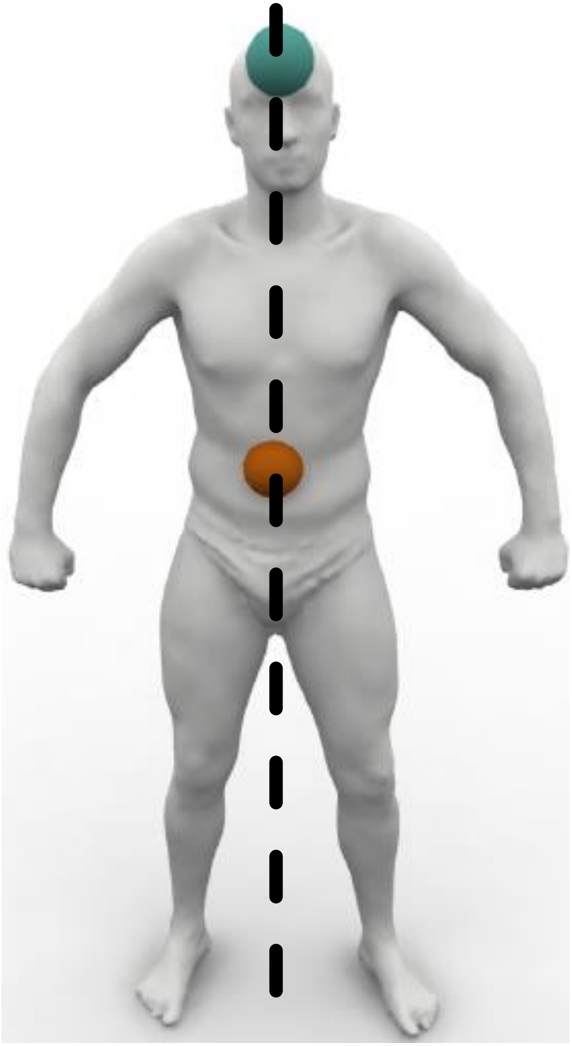

Tradeoff: GMDS

- Pros:
  - Good distance for non-isometric metric spaces
- Cons:
  - Non-convex
  - HUGE search space (i.e. permutations)

Adapted from slides by Q. Huang, V. Kim

GMDS in Practice

- Heuristics to explore the permutations
  - Solve at a very coarse scale and interpolate
  - Coarse-to-fine
  - Partial matching

Bronstein'08

Adapted from slides by Q. Huang, V. Kim

GMDS in Practice

- Heuristics to explore the permutations
  - Solve at a very coarse scale and interpolate
  - Coarse-to-fine
  - Partial matching

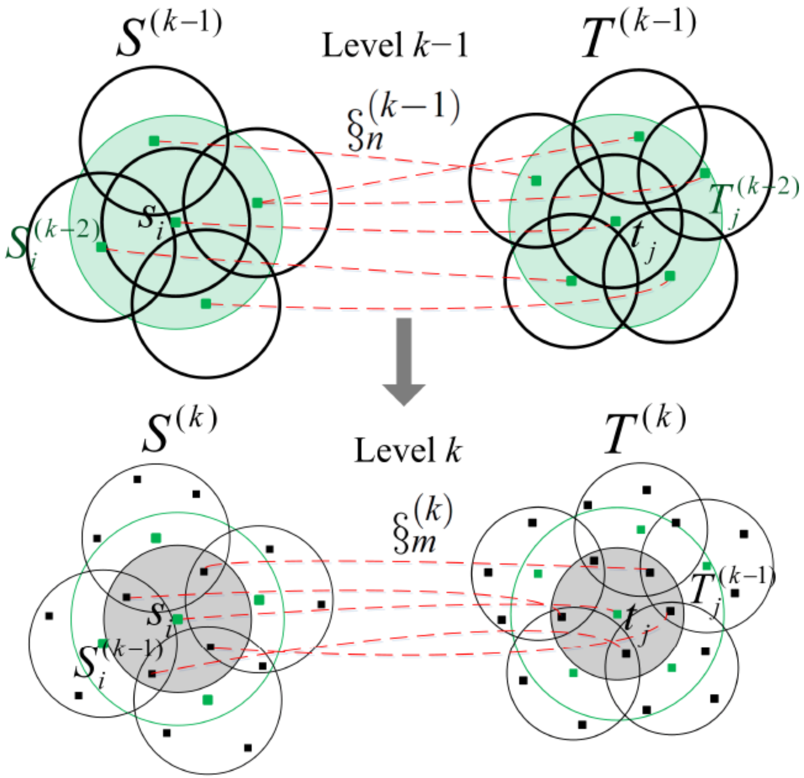

Sahillioglu'12

Adapted from slides by Q. Huang, V. Kim

GMDS in Practice

- Heuristics to explore the permutations
  - Solve at a very coarse scale and interpolate
  - Coarse-to-fine
  - Partial matching

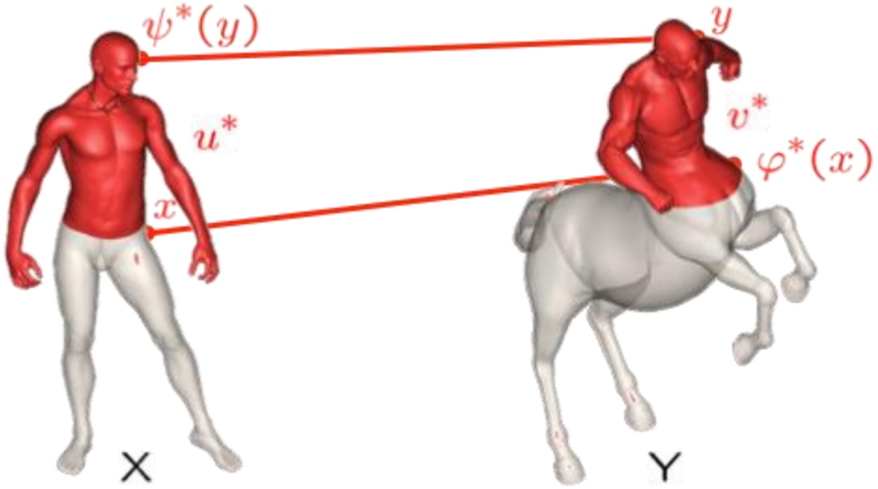

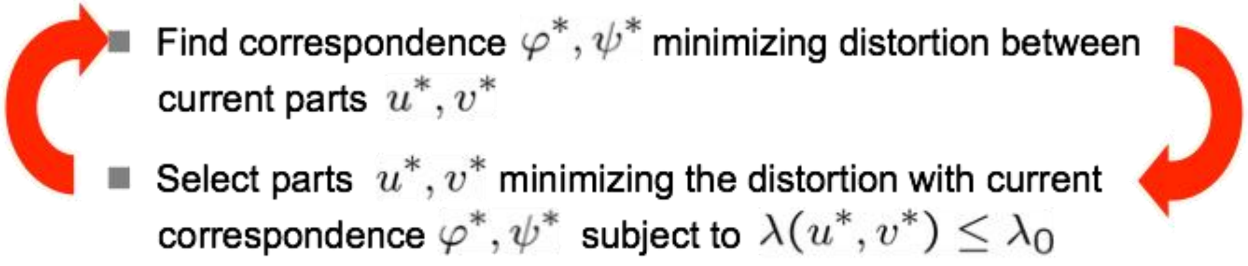

- Find correspondence φ^*, ψ^* minimizing distortion between current parts u^*, v^*
- Select parts u^*, v^* minimizing the distortion with current correspondence φ^*, ψ^* subject to $\lambda(u^*, v^*) \leq \lambda_0$

A. Bronstein, M. Bronstein, A. Bruckstein, R. Kimmel, *IJCV* 2008

Adapted from slides by Q. Huang, V. Kim

Returning to Desirable Properties

Given two (or more) shapes
Find a map f, satisfying the following properties:

- ~~Fast to compute~~
- ~~Bijective~~

(if we expect global correspondence)

- Low-distortion
- Preserves important features

(unless local optimum is bad)

Adapted from slides by Q. Huang, V. Kim

Recent idea:

Gromov-Wasserstein Distance

[Mémoli 2007]

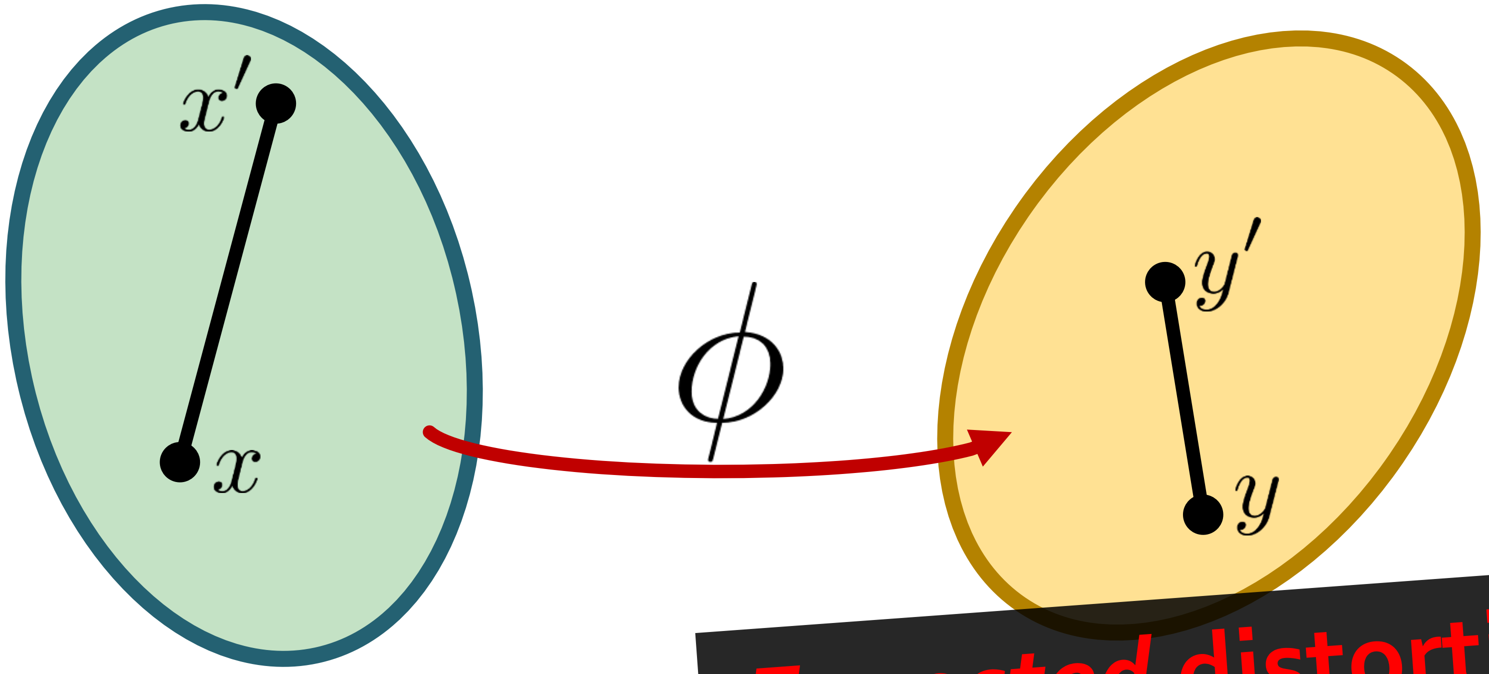

Expected distortion

$$\mathrm{GW}_2^2((\mu_0, d_0), (\mu, d)) :=$$

$$\min_{\gamma \in \mathcal{M}(\mu_0, \mu)} \iint_{\Sigma_0 \times \Sigma} [d_0(x, x') - d(y, y')]^2 \, d\gamma(x, y) \, d\gamma(x', y')$$

Recall:

Entropic Regularization

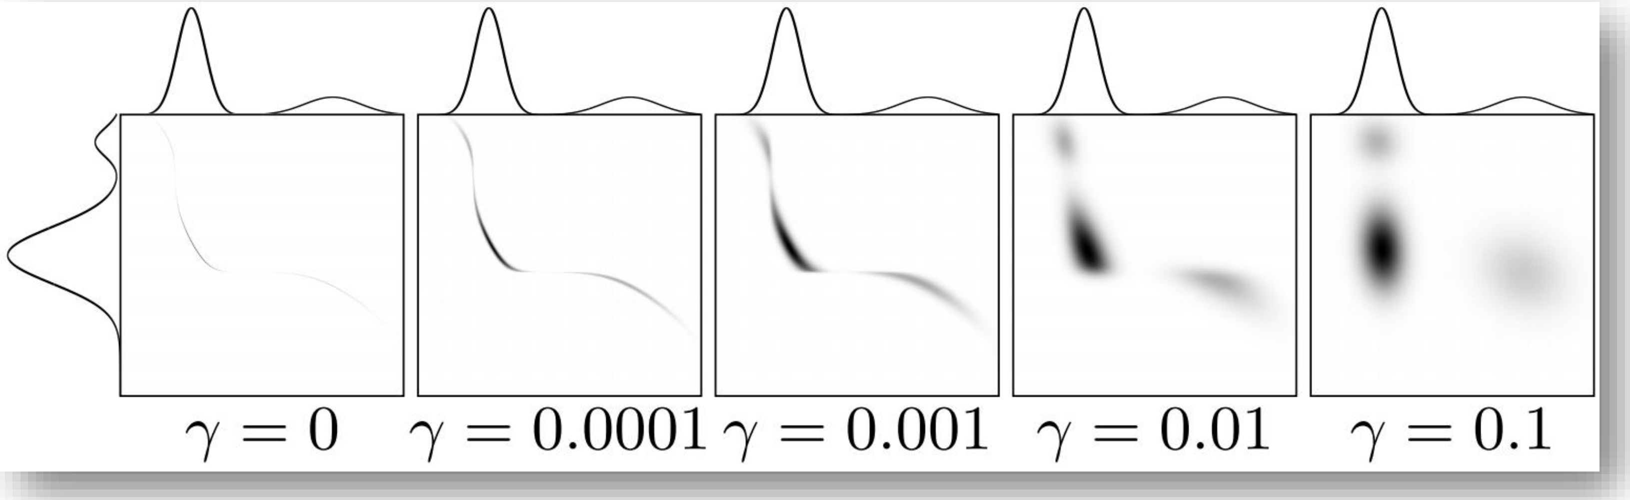

$$\begin{aligned} \min_T \quad & \textstyle\sum_{ij} T_{ij} d(x_i, x_j) - \gamma H(T) \\ \text{s.t.} \quad & \textstyle\sum_j T_{ij} = p_i \\ & \textstyle\sum_i T_{ij} = q_j \\ & T \geq 0 \end{aligned}$$

$$H(T) := -\sum_{ij} T_{ij} \log T_{ij}$$

Cuturi. "Sinkhorn distances: Lightspeed computation of optimal transport" (NIPS 2013)

Gromov-Wasserstein Plus Entropy

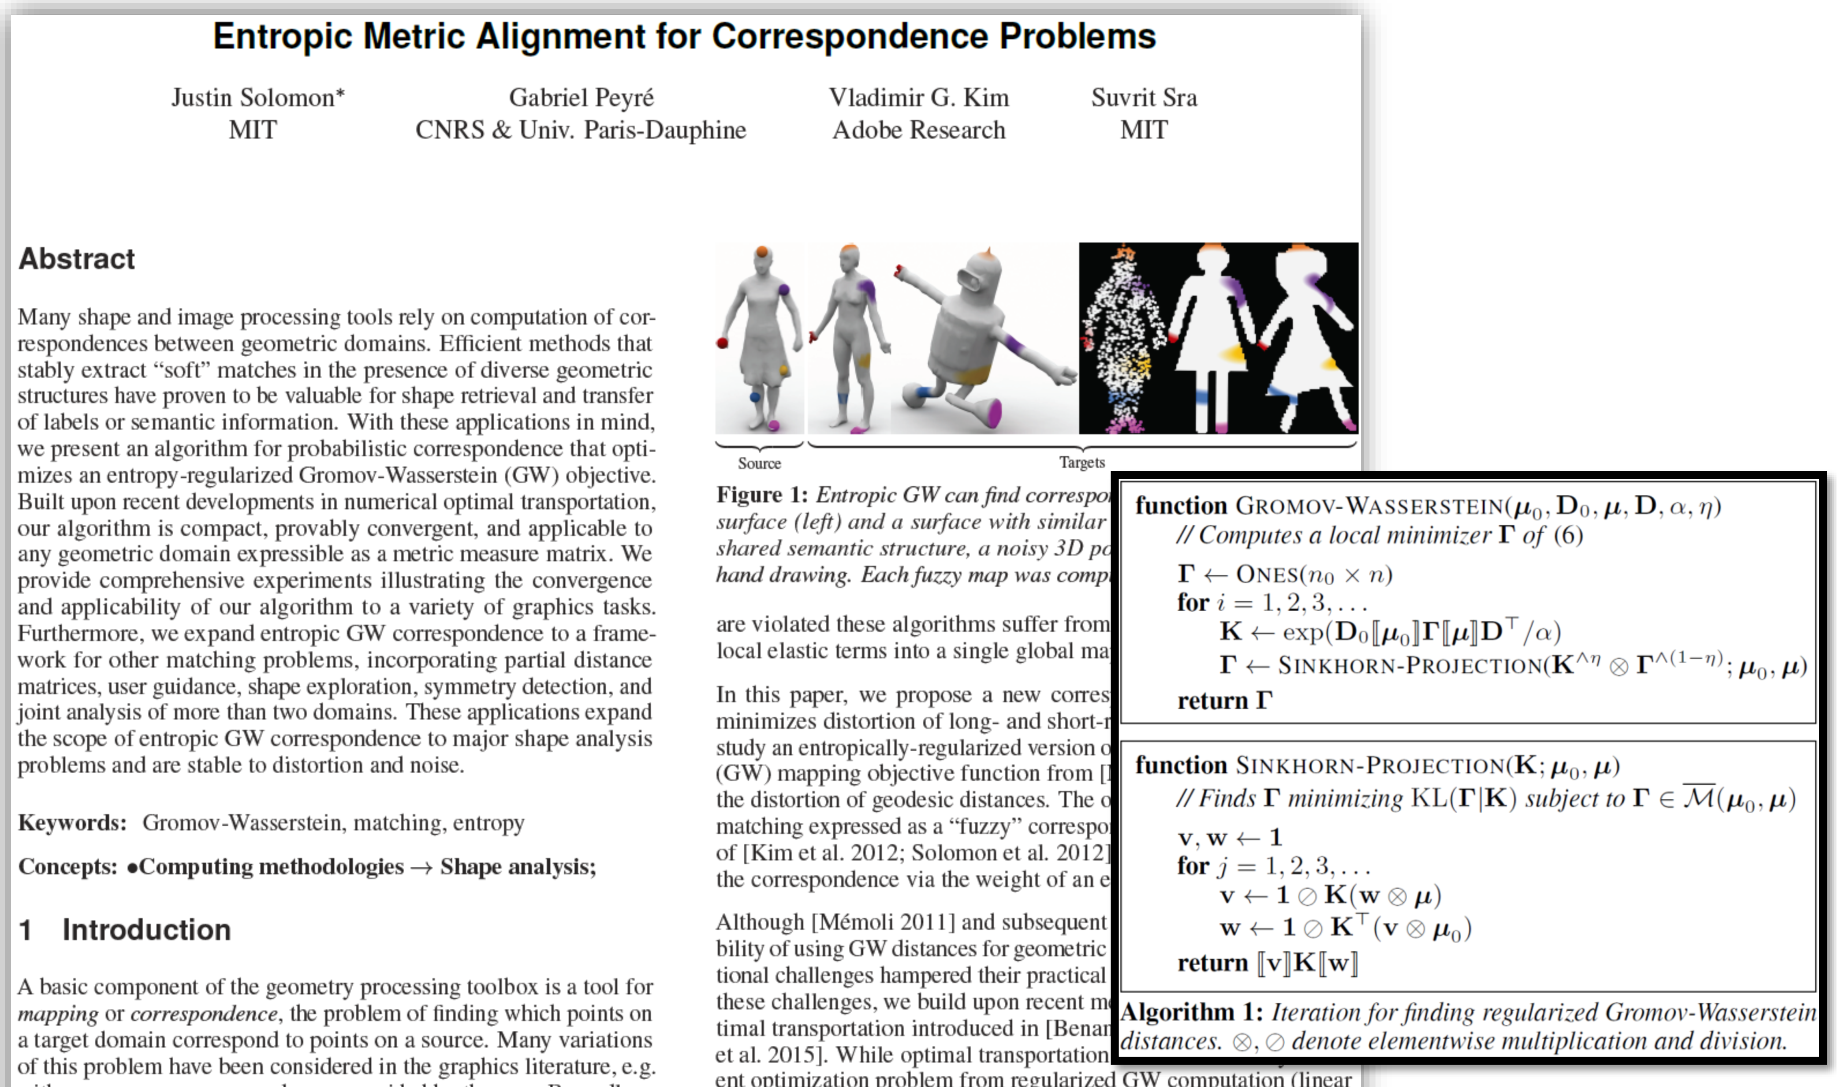

Entropic Metric Alignment for Correspondence Problems

Justin Solomon*
MIT

Gabriel Peyré
CNRS & Univ. Paris-Dauphine

Vladimir G. Kim
Adobe Research

Suvrit Sra
MIT

Abstract

Many shape and image processing tools rely on computation of correspondences between geometric domains. Efficient methods that stably extract "soft" matches in the presence of diverse geometric structures have proven to be valuable for shape retrieval and transfer of labels or semantic information. With these applications in mind, we present an algorithm for probabilistic correspondence that optimizes an entropy-regularized Gromov-Wasserstein (GW) objective. Built upon recent developments in numerical optimal transportation, our algorithm is compact, provably convergent, and applicable to any geometric domain expressible as a metric measure matrix. We provide comprehensive experiments illustrating the convergence and applicability of our algorithm to a variety of graphics tasks. Furthermore, we expand entropic GW correspondence to a framework for other matching problems, incorporating partial distance matrices, user guidance, shape exploration, symmetry detection, and joint analysis of more than two domains. These applications expand the scope of entropic GW correspondence to major shape analysis problems and are stable to distortion and noise.

Keywords: Gromov-Wasserstein, matching, entropy

Concepts: •**Computing methodologies → Shape analysis;**

1 Introduction

A basic component of the geometry processing toolbox is a tool for *mapping* or *correspondence*, the problem of finding which points on a target domain correspond to points on a source. Many variations of this problem have been considered in the graphics literature, e.g.

Figure 1: *Entropic GW can find correspo
surface (left) and a surface with similar
shared semantic structure, a noisy 3D po
hand drawing. Each fuzzy map was comp*

are violated these algorithms suffer from
local elastic terms into a single global ma

In this paper, we propose a new corres
minimizes distortion of long- and short-r
study an entropically-regularized version o
(GW) mapping objective function from [
the distortion of geodesic distances. The o
matching expressed as a "fuzzy" correspo
of [Kim et al. 2012; Solomon et al. 2012]
the correspondence via the weight of an e

Although [Mémoli 2011] and subsequent
bility of using GW distances for geometric
tional challenges hampered their practical
these challenges, we build upon recent m
timal transportation introduced in [Benar
et al. 2015]. While optimal transportation
ent optimization problem from regularized GW computation (linear

```
function GROMOV-WASSERSTEIN($\boldsymbol{\mu}_0, \mathbf{D}_0, \boldsymbol{\mu}, \mathbf{D}, \alpha, \eta$)
    // Computes a local minimizer $\boldsymbol{\Gamma}$ of (6)
    $\boldsymbol{\Gamma} \leftarrow$ ONES($n_0 \times n$)
    for $i = 1, 2, 3, \ldots$
        $\mathbf{K} \leftarrow \exp(\mathbf{D}_0[\![\boldsymbol{\mu}_0]\!]\boldsymbol{\Gamma}[\![\boldsymbol{\mu}]\!]\mathbf{D}^\top/\alpha)$
        $\boldsymbol{\Gamma} \leftarrow$ SINKHORN-PROJECTION($\mathbf{K}^{\wedge\eta} \otimes \boldsymbol{\Gamma}^{\wedge(1-\eta)}; \boldsymbol{\mu}_0, \boldsymbol{\mu}$)
    return $\boldsymbol{\Gamma}$

function SINKHORN-PROJECTION($\mathbf{K}; \boldsymbol{\mu}_0, \boldsymbol{\mu}$)
    // Finds $\boldsymbol{\Gamma}$ minimizing $\mathrm{KL}(\boldsymbol{\Gamma}|\mathbf{K})$ subject to $\boldsymbol{\Gamma} \in \overline{\mathcal{M}}(\boldsymbol{\mu}_0, \boldsymbol{\mu})$
    $\mathbf{v}, \mathbf{w} \leftarrow \mathbf{1}$
    for $j = 1, 2, 3, \ldots$
        $\mathbf{v} \leftarrow \mathbf{1} \oslash \mathbf{K}(\mathbf{w} \otimes \boldsymbol{\mu})$
        $\mathbf{w} \leftarrow \mathbf{1} \oslash \mathbf{K}^\top(\mathbf{v} \otimes \boldsymbol{\mu}_0)$
    return $[\![\mathbf{v}]\!]\mathbf{K}[\![\mathbf{w}]\!]$
```

Algorithm 1: *Iteration for finding regularized Gromov-Wasserstein distances.* $\otimes, \oslash$ *denote elementwise multiplication and division.*

Convex Relaxation

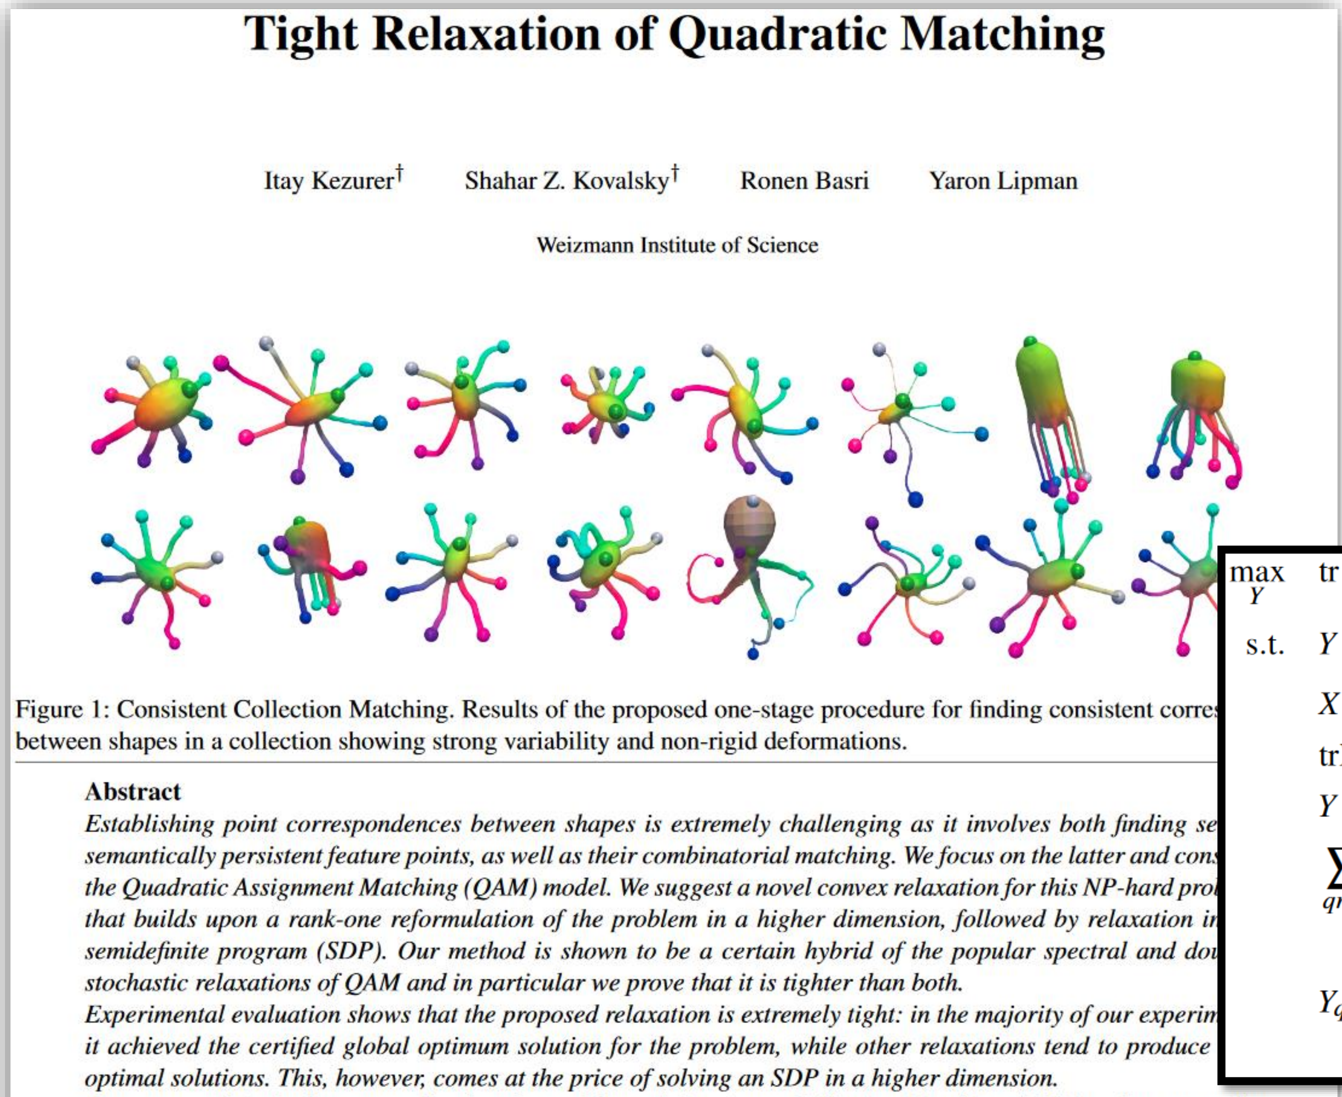

Tight Relaxation of Quadratic Matching

Itay Kezurer† Shahar Z. Kovalsky† Ronen Basri Yaron Lipman

Weizmann Institute of Science

Figure 1: Consistent Collection Matching. Results of the proposed one-stage procedure for finding consistent corres
between shapes in a collection showing strong variability and non-rigid deformations.

Abstract

Establishing point correspondences between shapes is extremely challenging as it involves both finding se semantically persistent feature points, as well as their combinatorial matching. We focus on the latter and cons the Quadratic Assignment Matching (QAM) model. We suggest a novel convex relaxation for this NP-hard pro that builds upon a rank-one reformulation of the problem in a higher dimension, followed by relaxation in semidefinite program (SDP). Our method is shown to be a certain hybrid of the popular spectral and dou stochastic relaxations of QAM and in particular we prove that it is tighter than both.

Experimental evaluation shows that the proposed relaxation is extremely tight: in the majority of our experim it achieved the certified global optimum solution for the problem, while other relaxations tend to produce optimal solutions. This, however, comes at the price of solving an SDP in a higher dimension.

$$
\begin{aligned}
\max_{Y} \quad & \mathrm{tr}\,(WY) \\
\text{s.t.} \quad & Y \succeq [X]\,[X]^T \\
& X \in \mathrm{conv}\,\Pi_n^k \\
& \mathrm{tr} Y = k \\
& Y \geq 0 \\
& \sum_{qrst} Y_{qrst} = k^2 \\
& Y_{qrst} \leq \begin{cases} 0, & \text{if } q = s,\ r \neq t \\ 0, & \text{if } r = t,\ q \neq s \\ \min\{X_{qr}, X_{st}\}, & \text{otherwise} \end{cases}
\end{aligned}
$$

Continuum

Weak assumptions

Strong assumptions

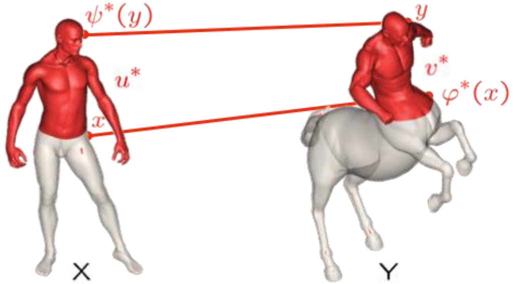

Low-distortion

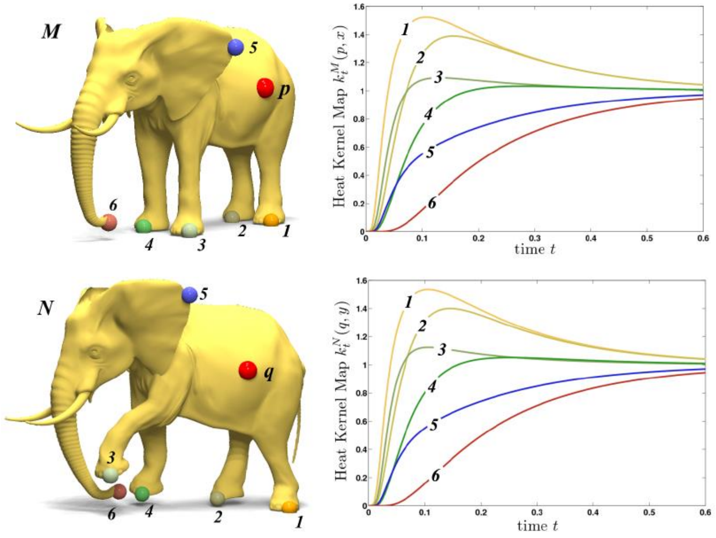

Isometry

Recall:

Heat Kernel Map

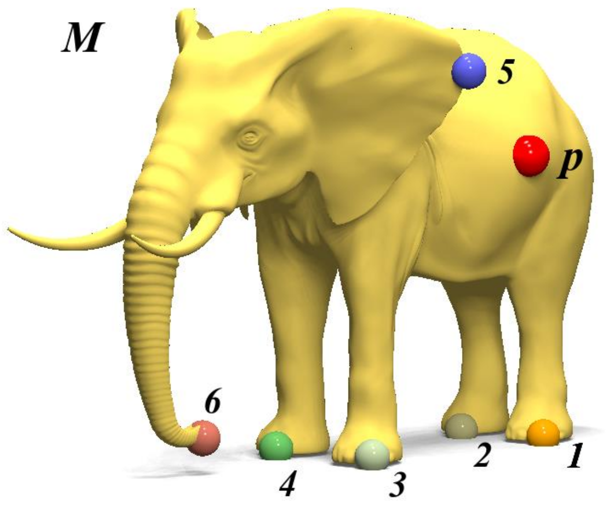

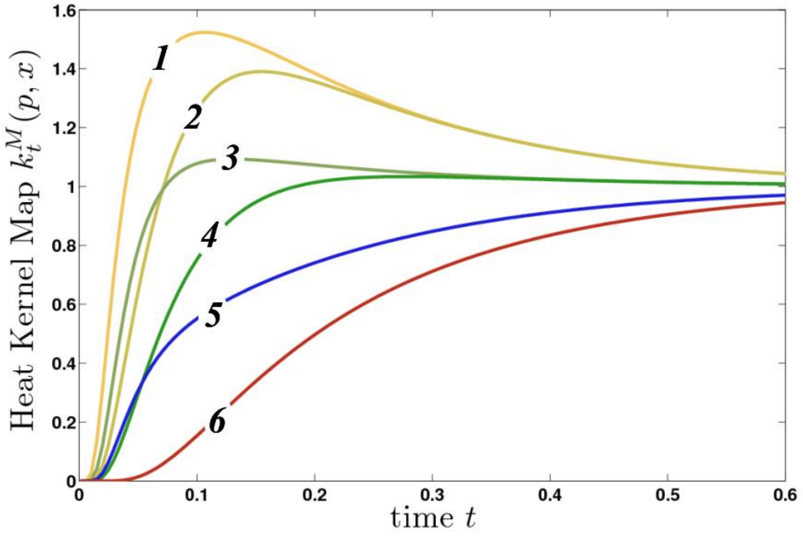

$$\mathrm{HKM}_p(x,t) := k_t(p,x)$$

Theorem: Only have to match one point!

KNN

One Point Isometric Matching with the Heat Kernel
Ovsjanikov et al. 2010

Tradeoff: Heat Kernel Map

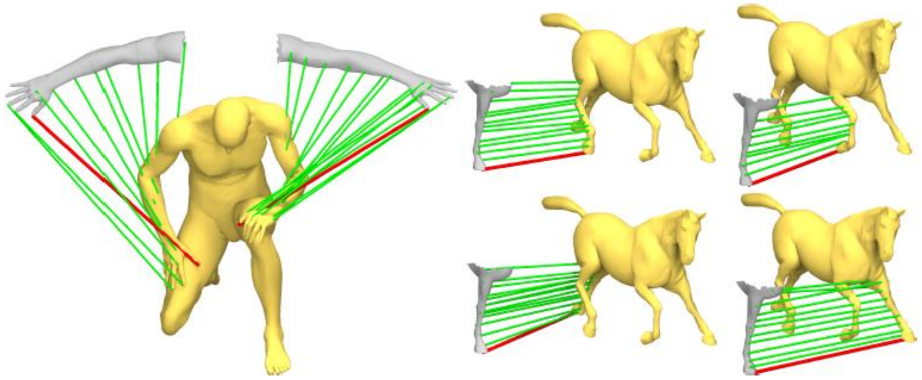

- Pros:
  - Tiny search space
  - Some extension to partial matching
- Cons:
  - (Extremely) sensitive to deviation from isometry

Adapted from slides by Q. Huang, V. Kim

Continuum

????

Weak assumptions

Strong assumptions

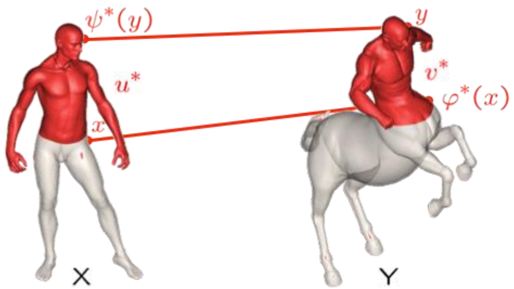

Low-distortion

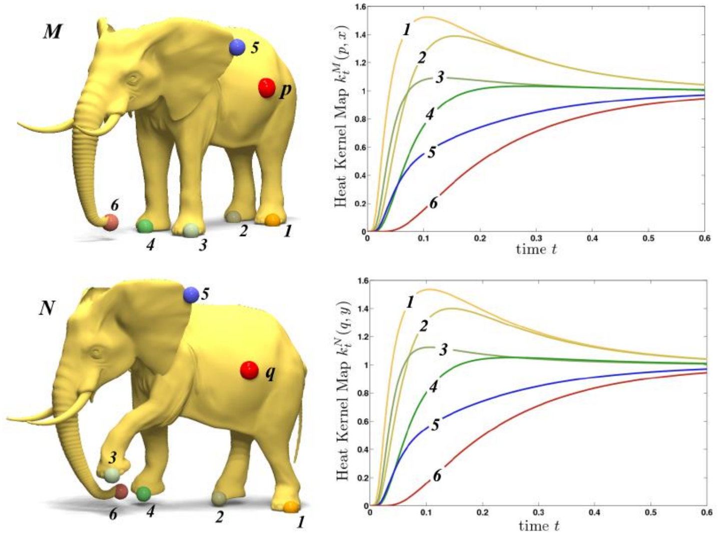

Isometry

Observation About Mapping

Angle and area preserving *Angle preserving*

$$\text{isometries} \subseteq \text{conformal maps}$$

Hard! **Easier**

Möbius Voting for Surface Correspondence
Lipman and Funkhouser 2009

$O(n^3)$ Algorithm for Perfect Isometry

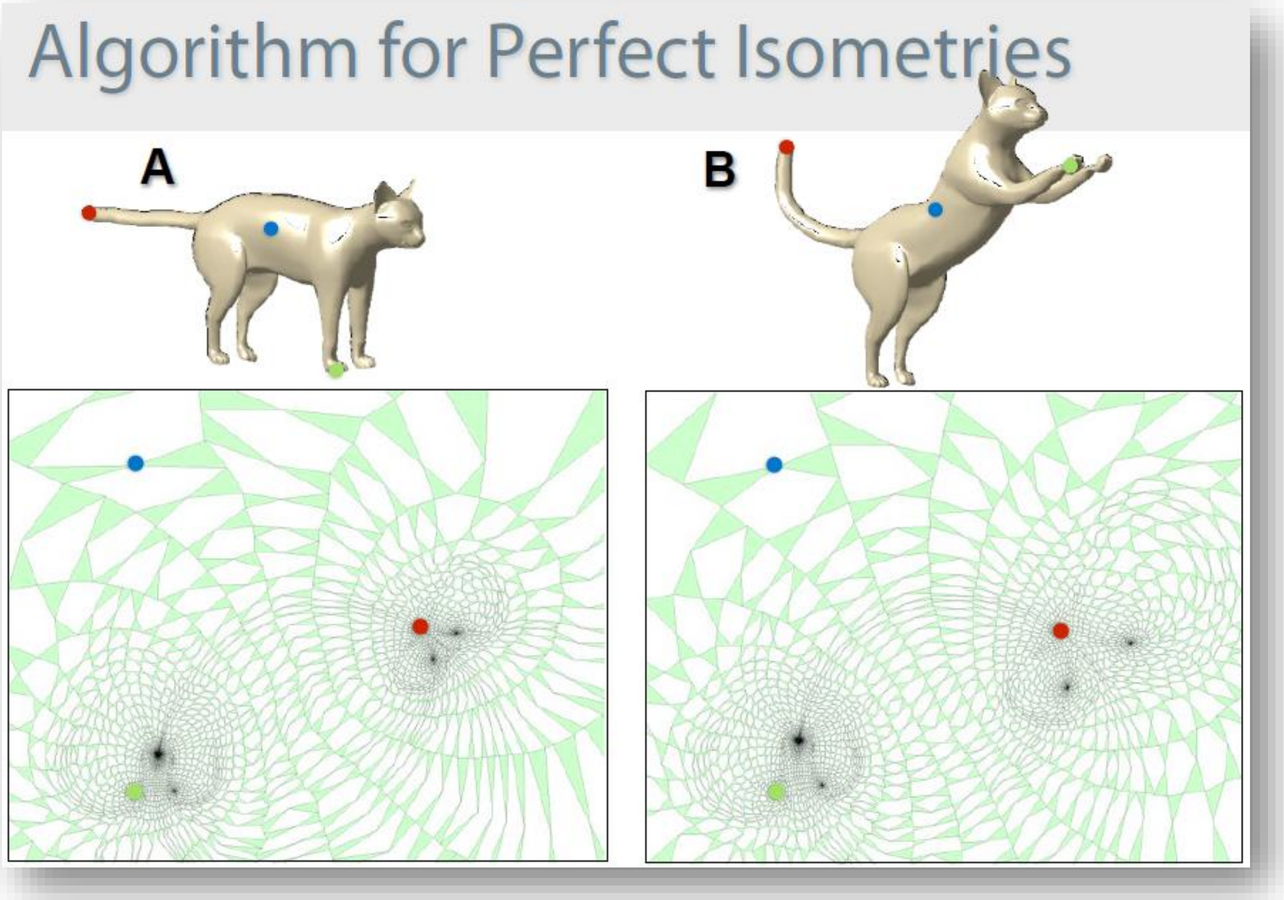

Algorithm for Perfect Isometries

http://www.mpi-inf.mpg.de/resources/deformableShapeMatching/EG2011_Tutorial/slides/4.3%20SymmetryApplications.pdf

Map triplets of points

Observation

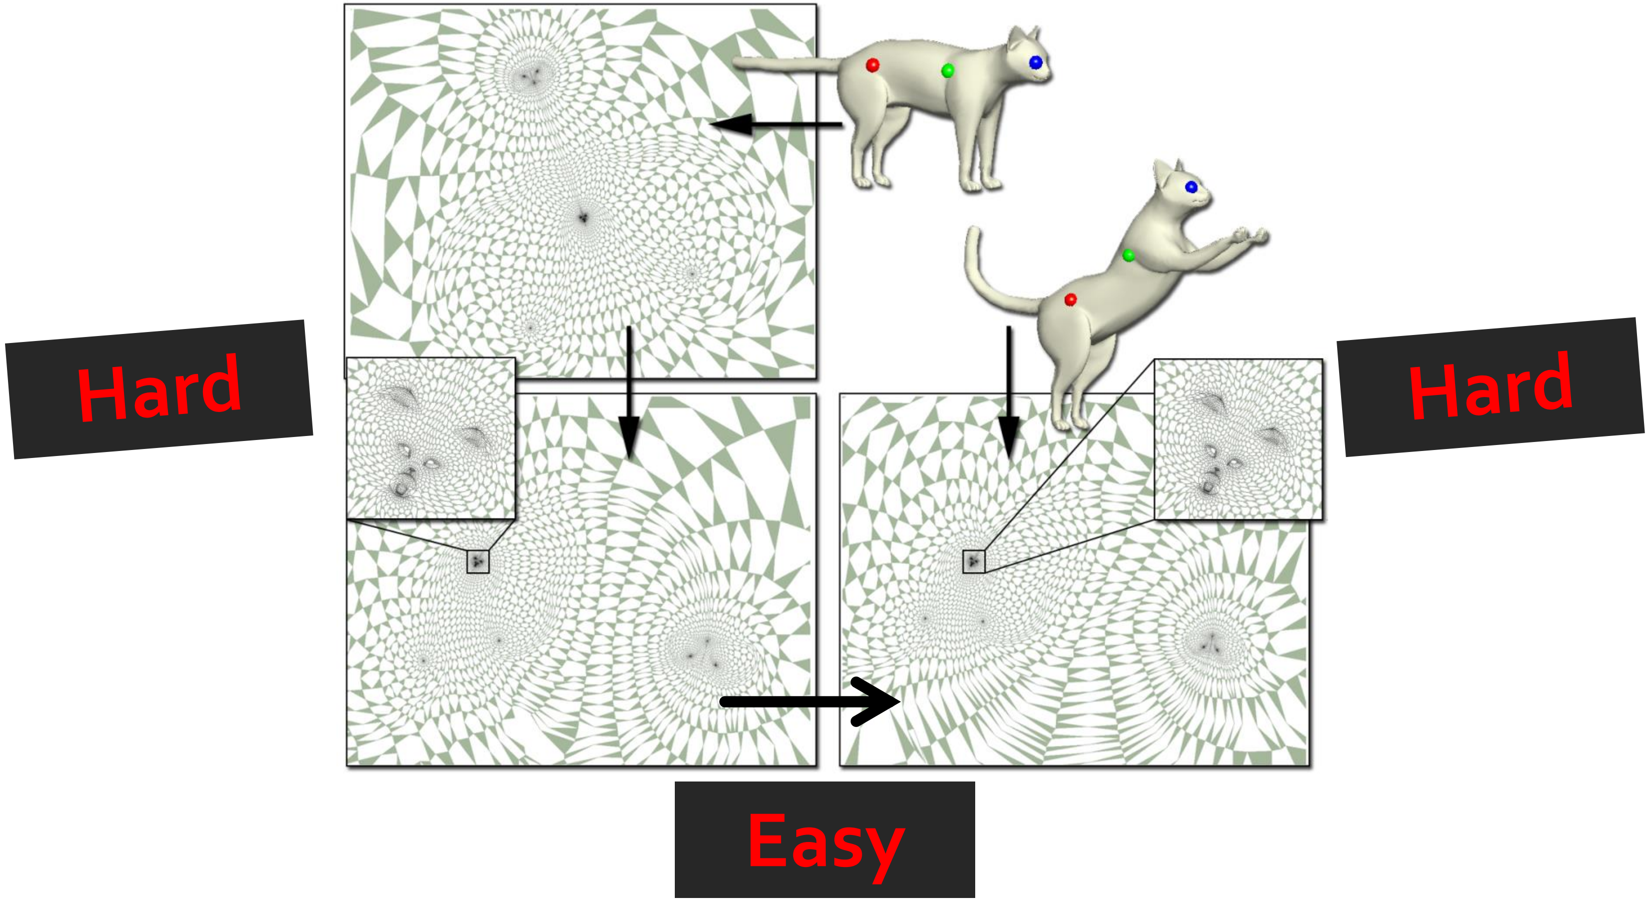

Hard work is per-surface, not per-map

Möbius Transformations

$$\frac{az+b}{cz+d}$$

http://www.ima.umn.edu/~arnold//moebius

Bijective conformal maps of the extended complex plane

Mid-Edge Uniformization

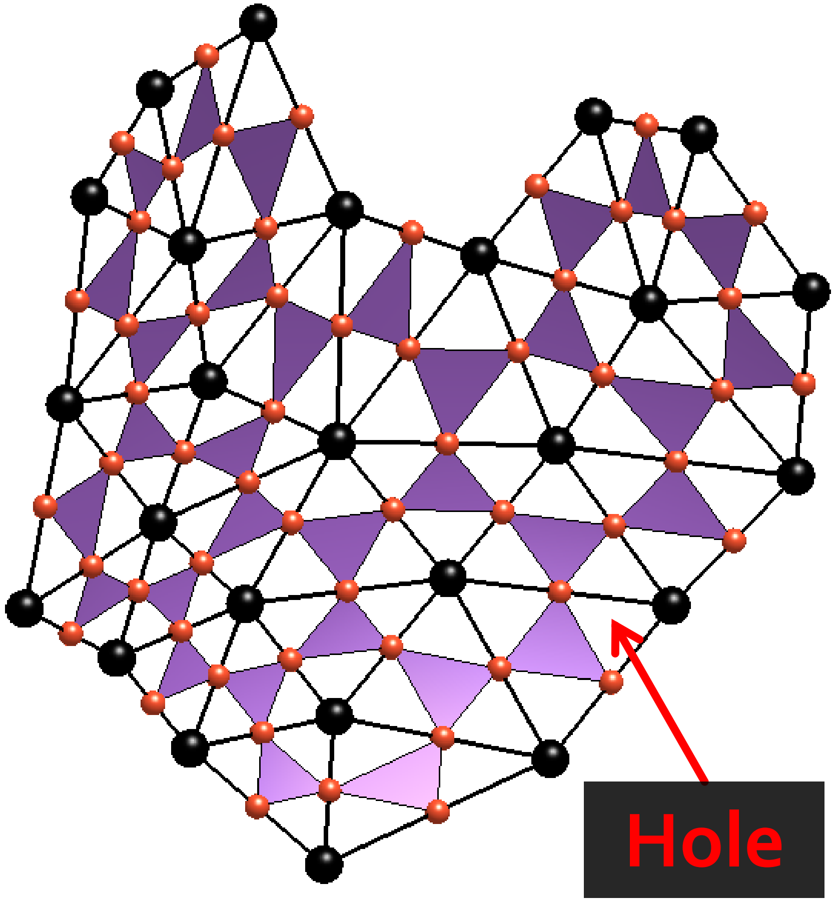

$$\Phi(v) = u(v) + iu^*(v)$$

PL, continuous

$$\Delta u = 0$$

PL, continuous at midpoints

Rotate gradient of *u* 90°

Cannot scale triangles to flatten

Möbius Voting

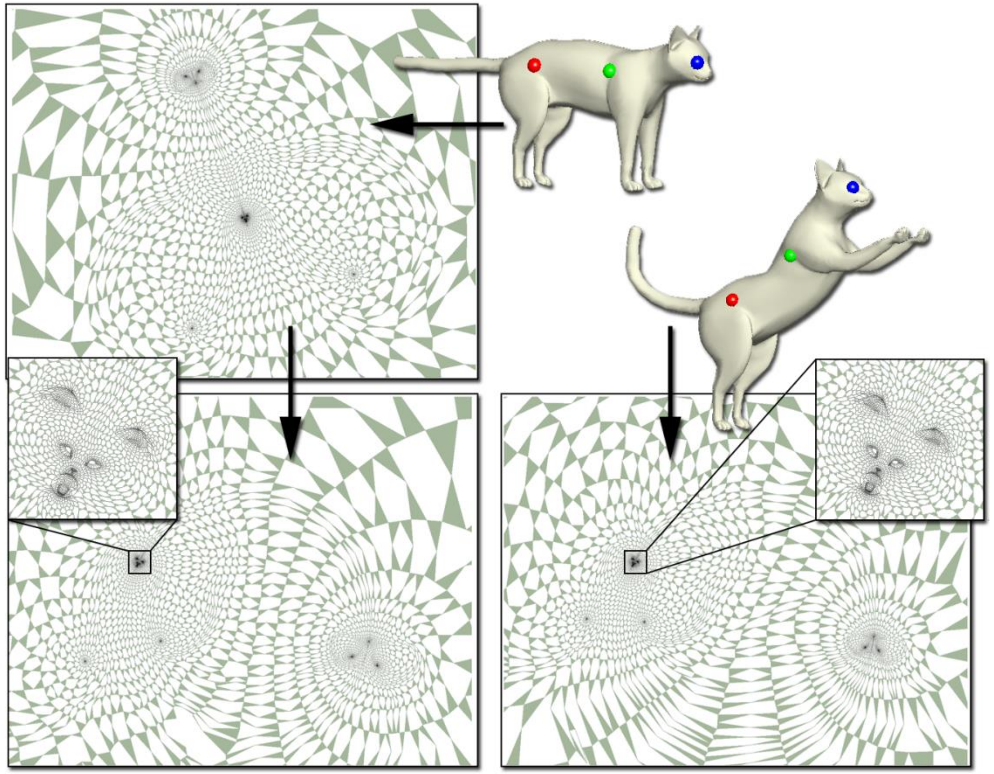

1. Map surfaces to complex plane
2. Select three points
3. Map plane to itself matching these points
4. Vote for pairings using distortion metric to weight
5. Return to 2

Möbius Voting for Surface Correspondence
Lipman and Funkhouser 2009

Voting Algorithm

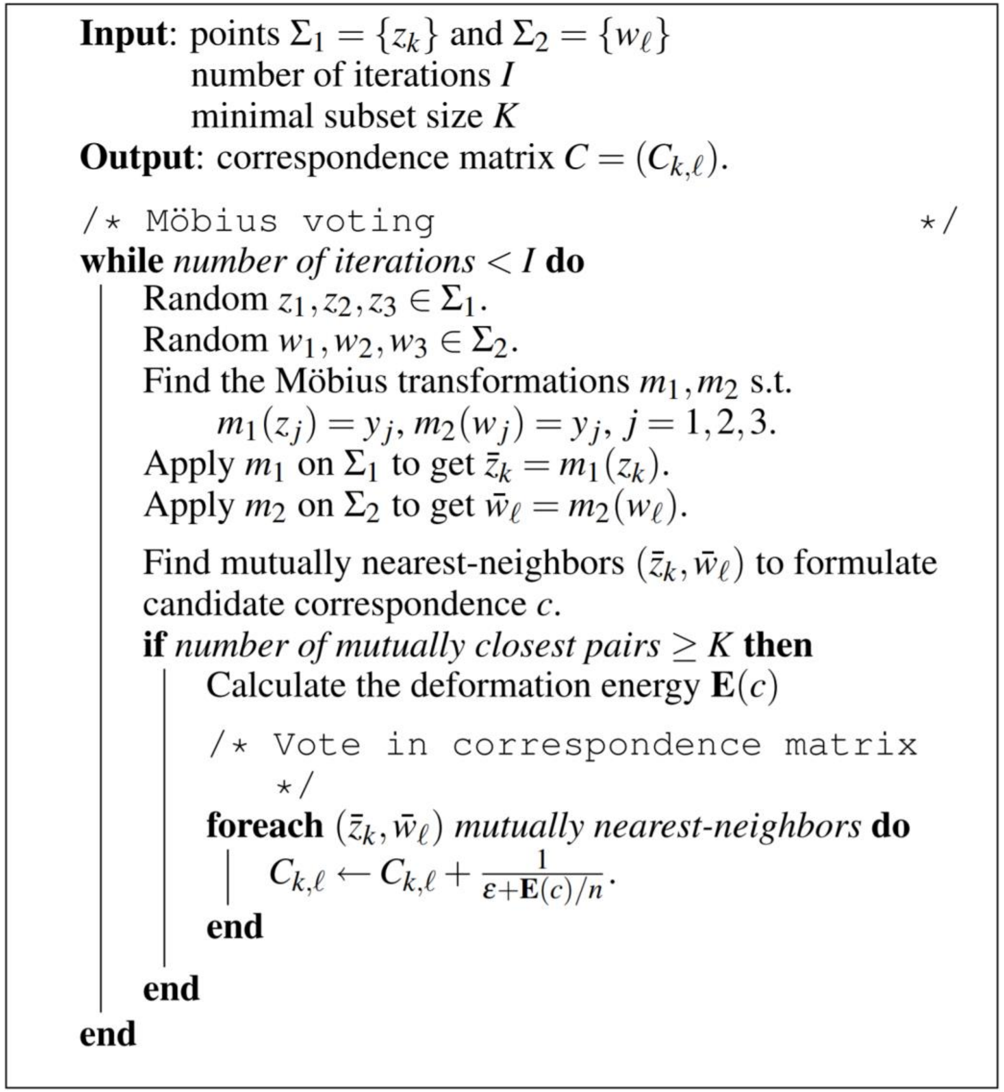

```
Input: points $\Sigma_1 = \{z_k\}$ and $\Sigma_2 = \{w_\ell\}$
       number of iterations $I$
       minimal subset size $K$
Output: correspondence matrix $C = (C_{k,\ell})$.

/* Möbius voting                                          */
while number of iterations $< I$ do
    Random $z_1, z_2, z_3 \in \Sigma_1$.
    Random $w_1, w_2, w_3 \in \Sigma_2$.
    Find the Möbius transformations $m_1, m_2$ s.t.
        $m_1(z_j) = y_j,\ m_2(w_j) = y_j,\ j = 1,2,3.$
    Apply $m_1$ on $\Sigma_1$ to get $\bar{z}_k = m_1(z_k)$.
    Apply $m_2$ on $\Sigma_2$ to get $\bar{w}_\ell = m_2(w_\ell)$.

    Find mutually nearest-neighbors $(\bar{z}_k, \bar{w}_\ell)$ to formulate
    candidate correspondence $c$.
    if number of mutually closest pairs $\geq K$ then
        Calculate the deformation energy $\mathbf{E}(c)$

        /* Vote in correspondence matrix
            */
        foreach $(\bar{z}_k, \bar{w}_\ell)$ mutually nearest-neighbors do
            $C_{k,\ell} \leftarrow C_{k,\ell} + \frac{1}{\varepsilon + \mathbf{E}(c)/n}$.
        end
    end
end
```

Tradeoff: Möbius Voting

- Pros:
  - Efficient
  - Voting procedure handles some non-isometry
- Cons:
  - Does not provide smooth/continuous map
  - Does not optimize global distortion
  - Only for genus 0

Adapted from slides by Q. Huang, V. Kim

Blended Intrinsic Maps

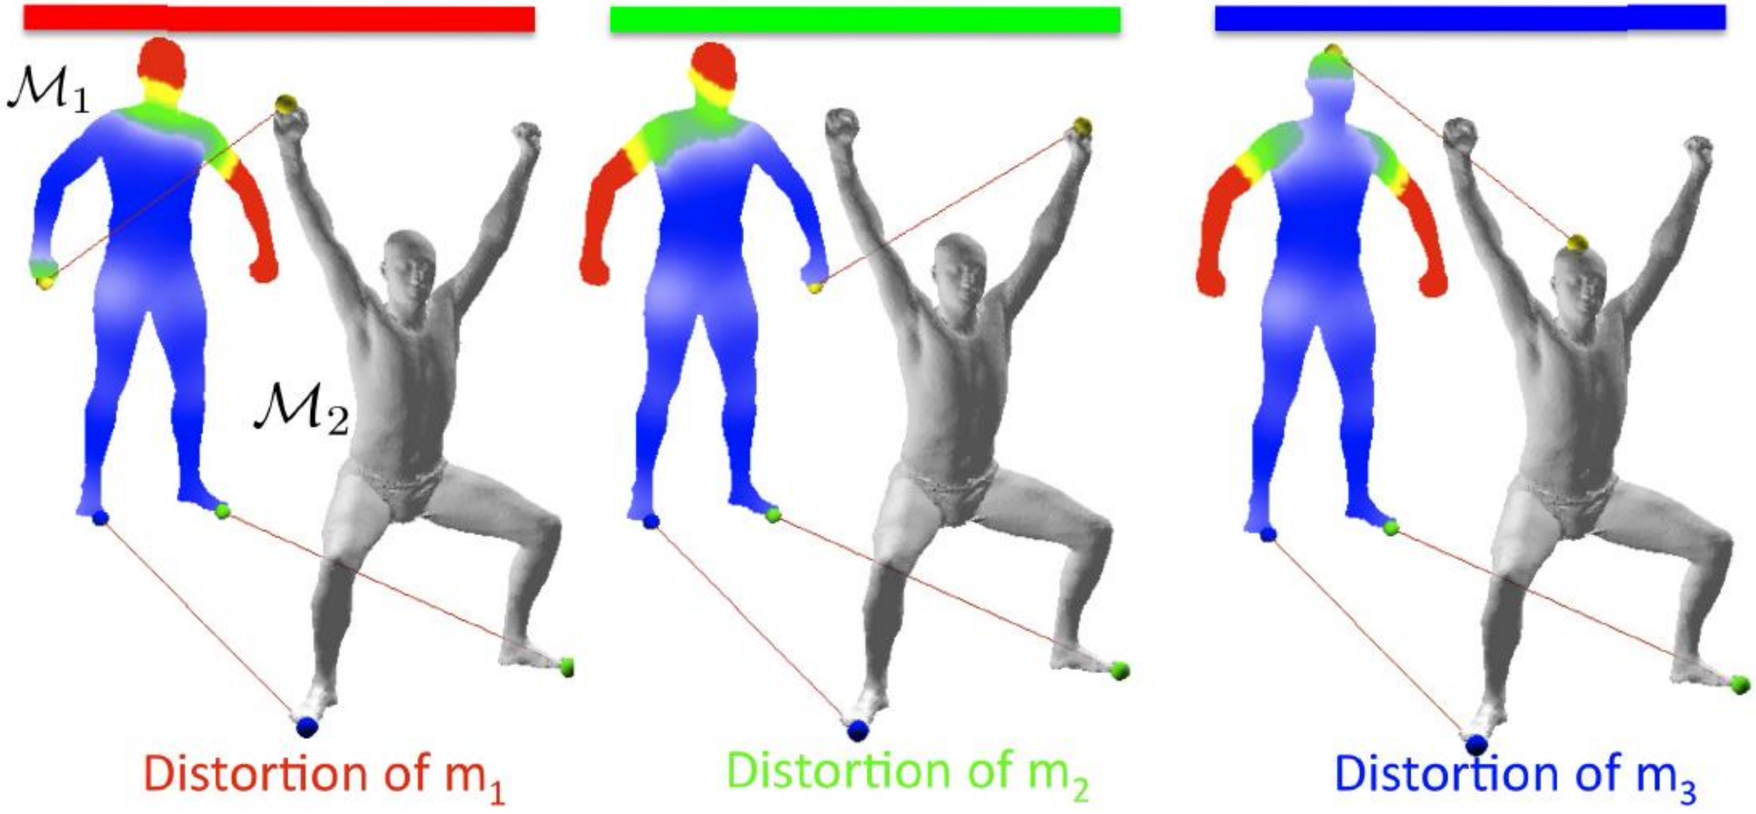

Different conformal maps distorted in different places.

Blended Intrinsic Maps
Kim, Lipman, and Funkhouser 2011

Use for Dense Mapping

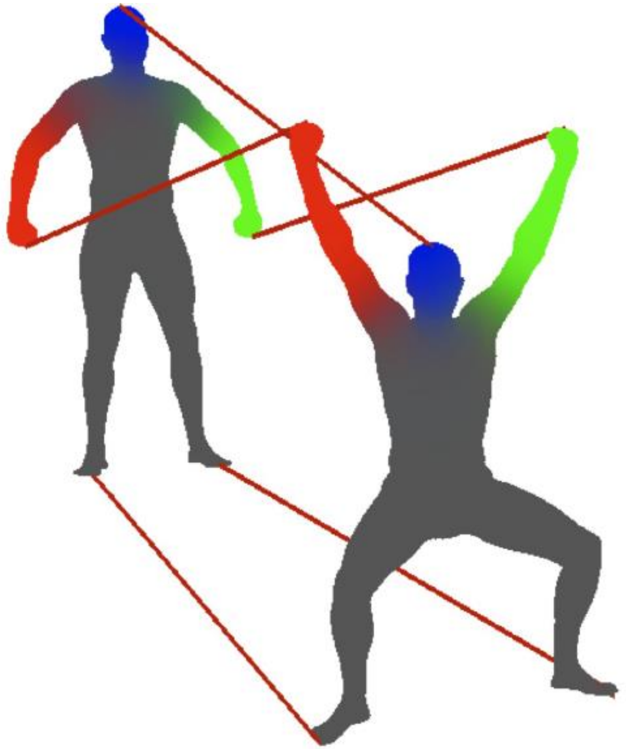

Blending Weights for m_1, m_2, and m_3

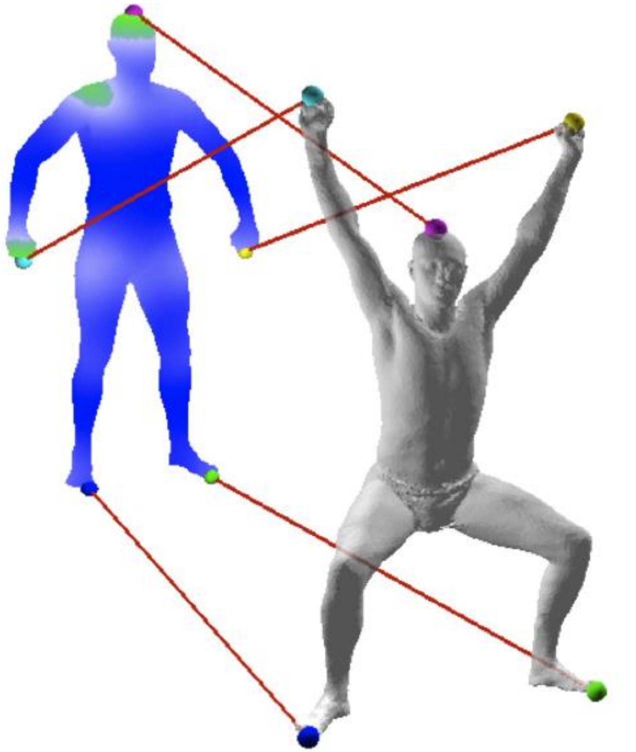

Distortion of the Blended Map

Combine good parts of different maps!

Blended Intrinsic Maps
Kim, Lipman, and Funkhouser 2011

Blended Intrinsic Maps

- Algorithm:
  - Generate consistent maps
  - Find blending weights per-point on each map
  - Blend maps

Adapted from slides by Q. Huang, V. Kim

Kim'11

Blended Intrinsic Maps

- Algorithm:
  - Generate consistent maps
  - Find blending weights per-point on each map
  - Blend maps

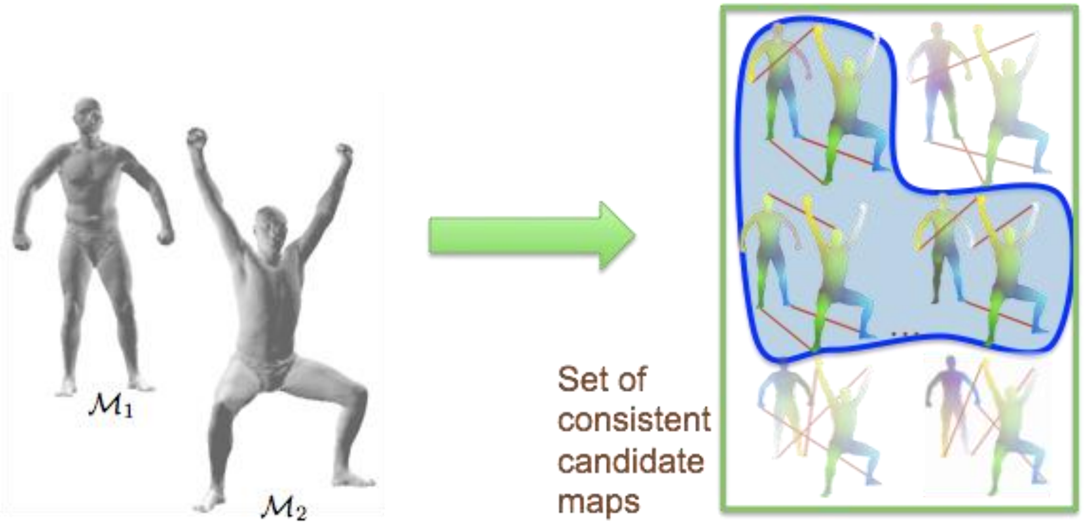

Adapted from slides by Q. Huang, V. Kim

Kim’11

Blended Intrinsic Maps

- Algorithm:
  - Generate consistent maps
  - Find blending weights per-point on each map
  - Blend maps

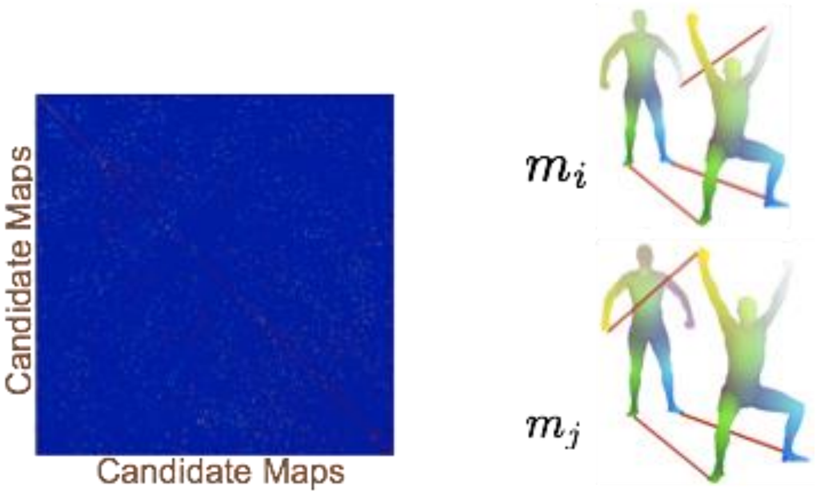

Map similarity matrix

Adapted from slides by Q. Huang, V. Kim

Kim'11

Blended Intrinsic Maps

- Algorithm:
  - Generate consistent maps
  - Find blending weights per-point on each map
  - Blend maps

Eigen-analysis:
“Block” of similar maps

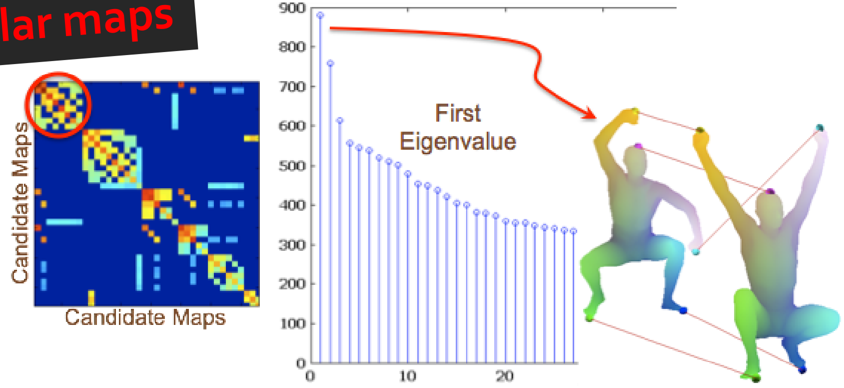

Adapted from slides by Q. Huang, V. Kim

Kim’11

Blended Intrinsic Maps

- Algorithm:
    - Generate consistent maps
    - Find blending weights per-point on each map
    - Blend maps

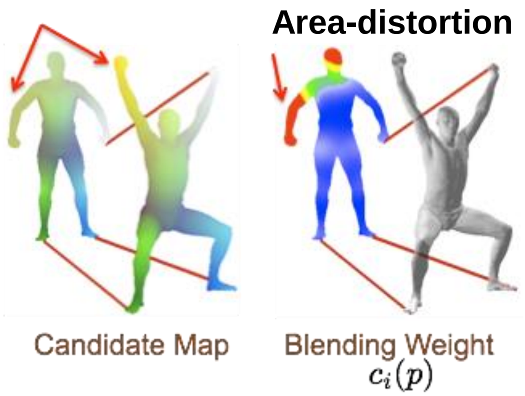

Adapted from slides by Q. Huang, V. Kim

Kim'11

Blended Intrinsic Maps

- Algorithm:
  - Generate consistent maps
  - Find blending weights per-point on each map
  - Blend maps

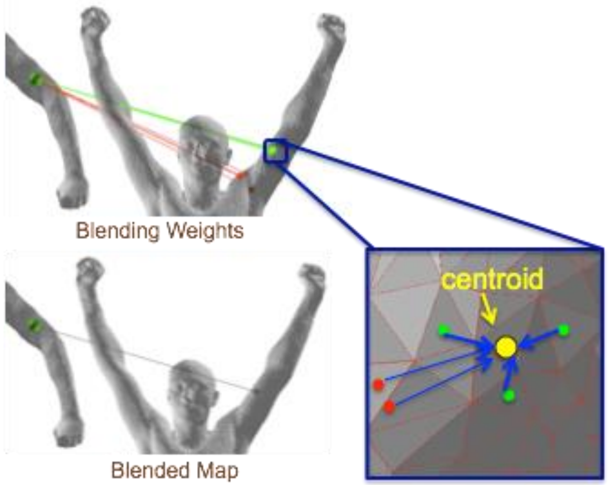

Adapted from slides by Q. Huang, V. Kim

Kim'11

Some Examples

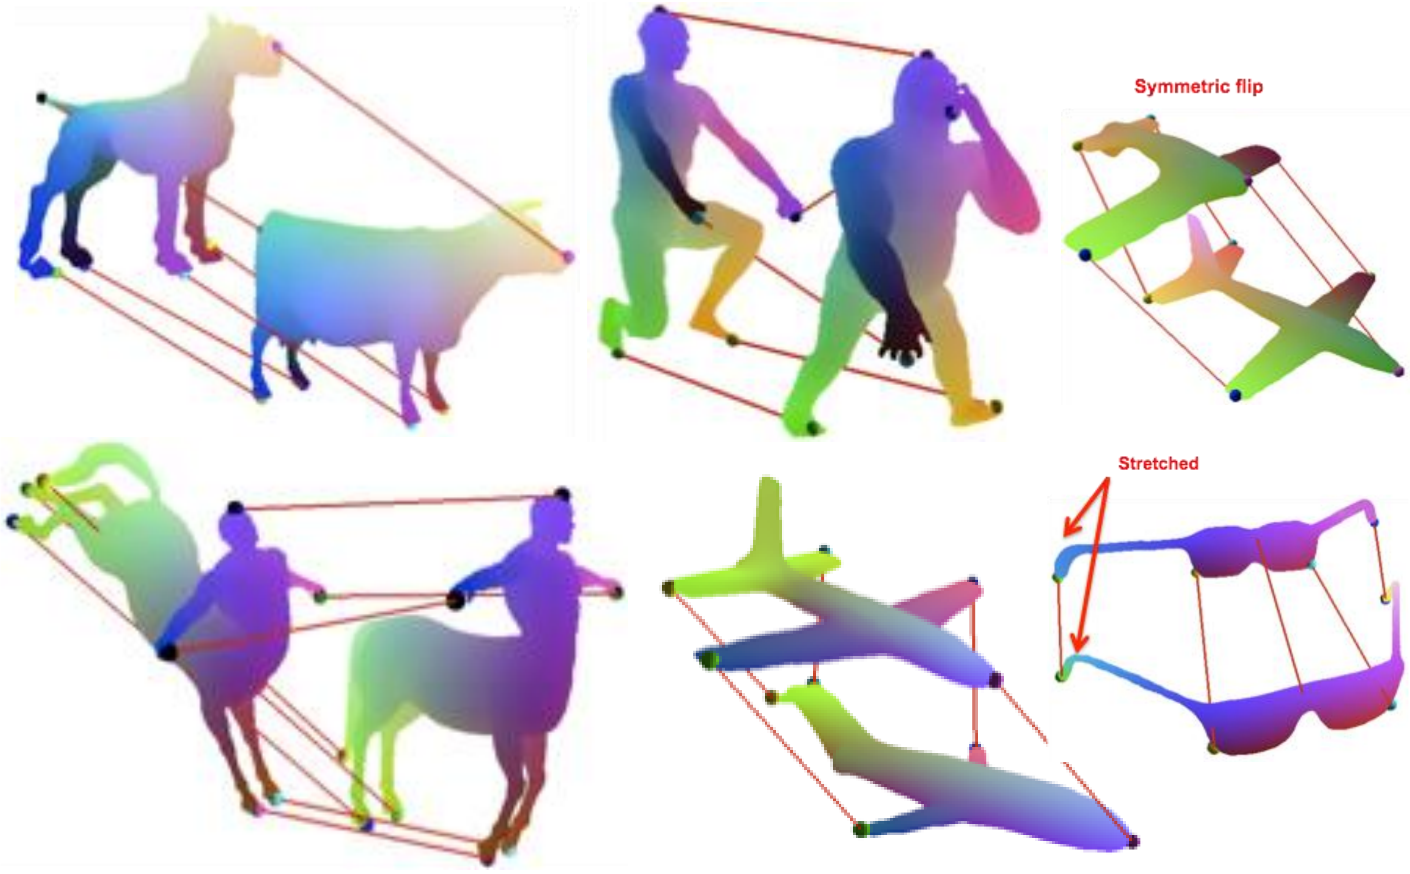

Adapted from slides by Q. Huang, V. Kim

Kim’11

Evaluation

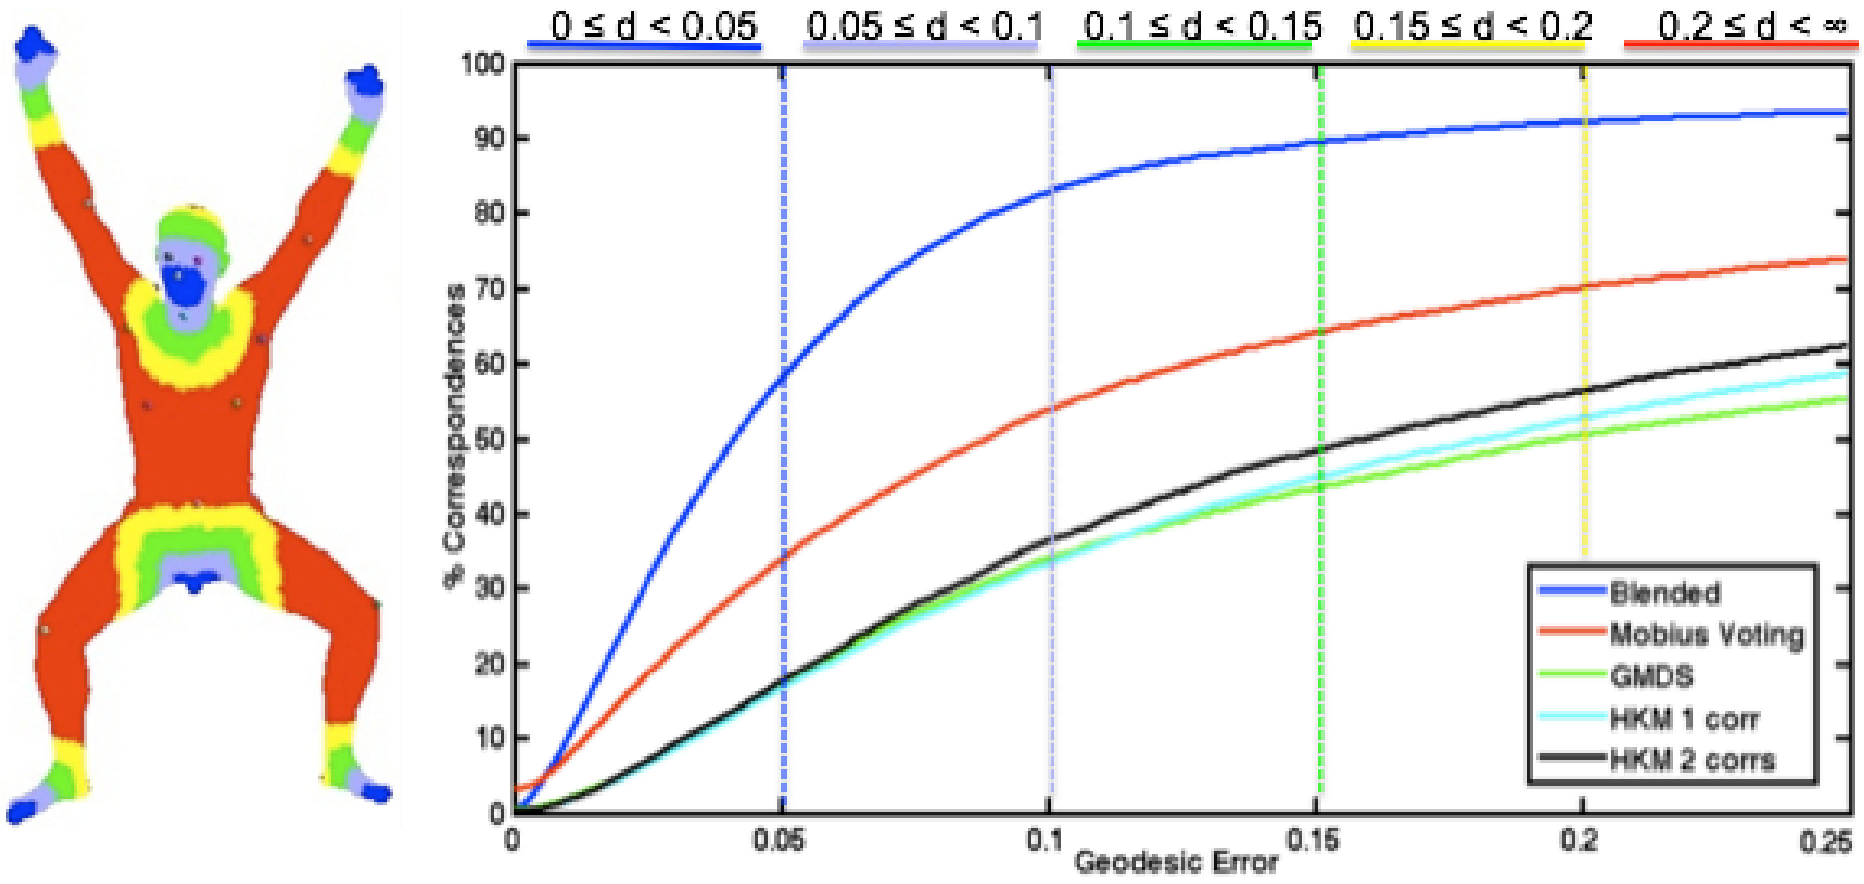

Kim'11

Tradeoff: Blended Intrinsic Maps

- Pros:
  - Can handle non-isometric shapes
  - Efficient
- Cons:
  - Lots of area distortion for some shapes
  - Genus o manifold surfaces

Adapted from slides by Q. Huang, V. Kim

Subtlety: Representation

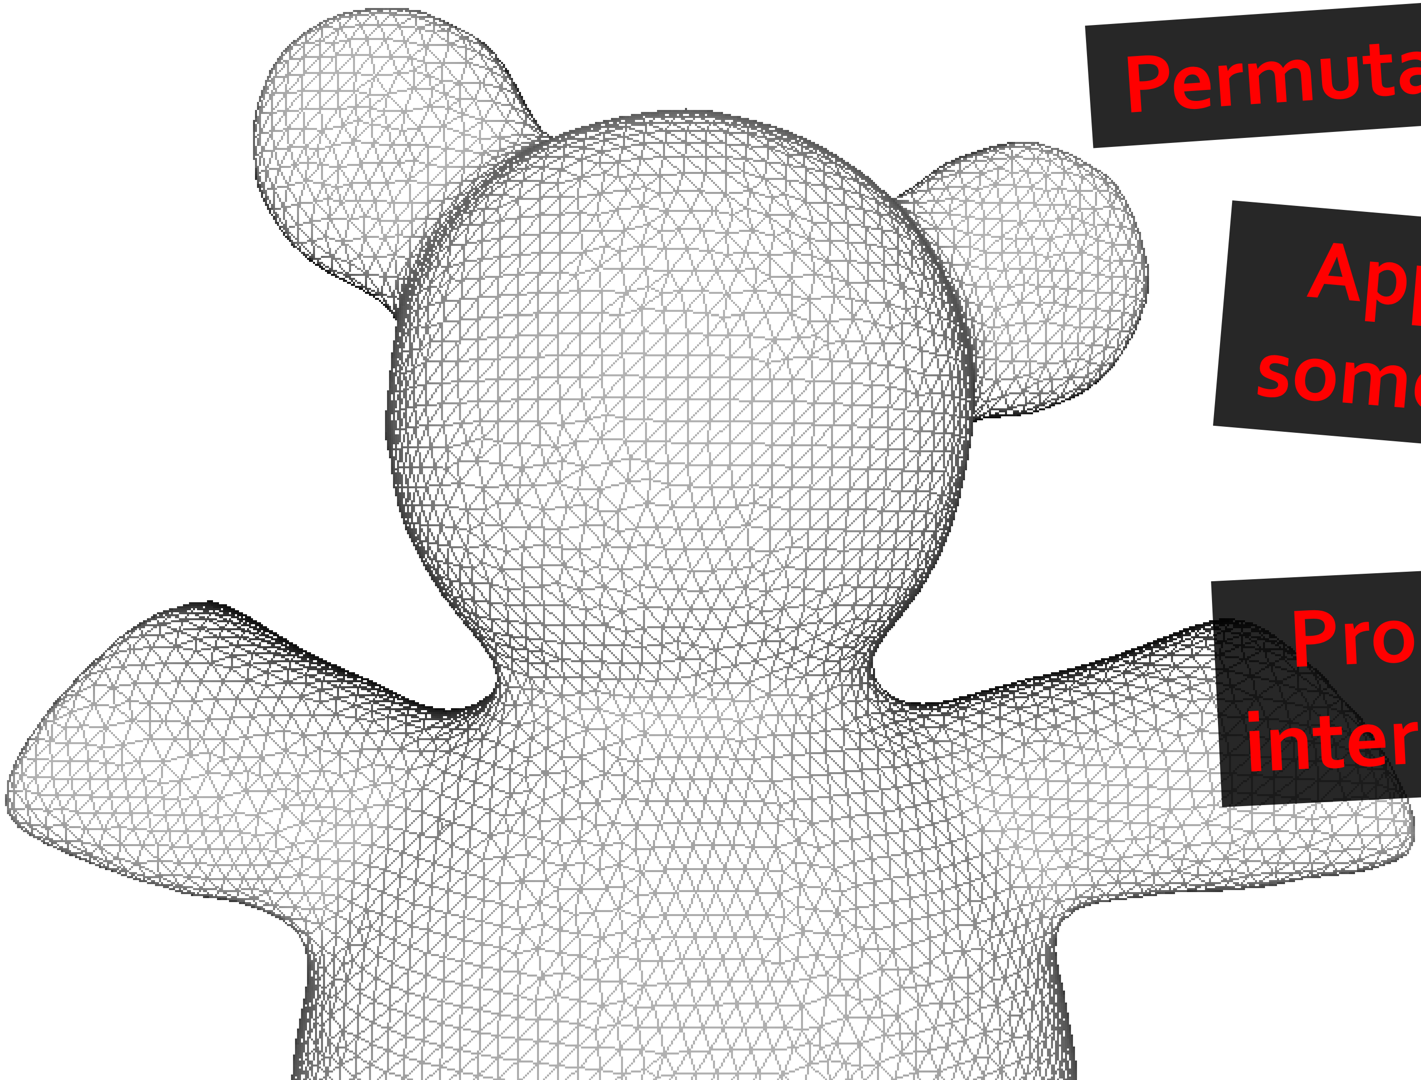

Permutation?

Approximation of something smooth?

Probabilistic interpretation?

Invertibility?

Functional Maps

[Ovsjanikov et al. 2012]

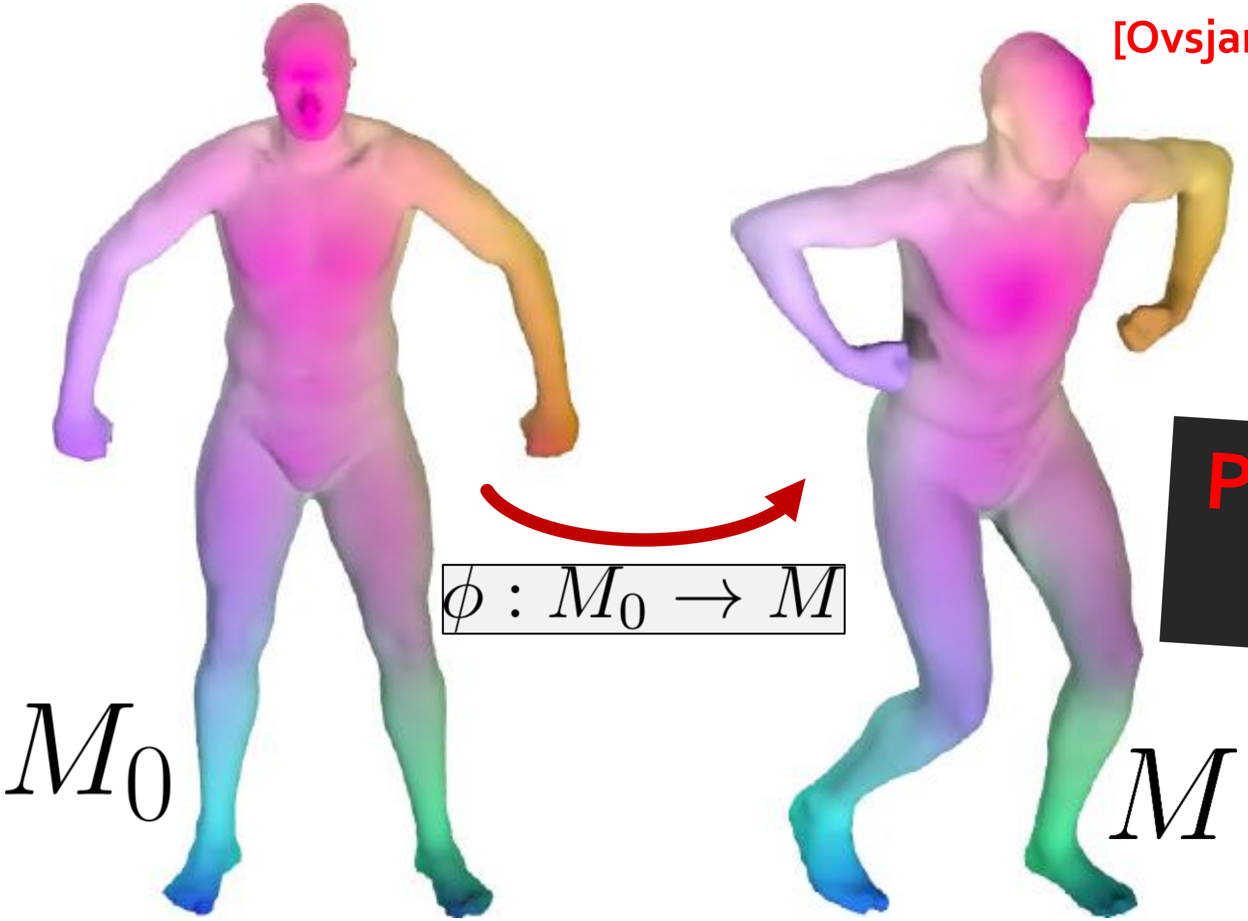

Points to points

Points on M_0 to points on M

Functional Maps

[Ovsjanikov et al. 2012]

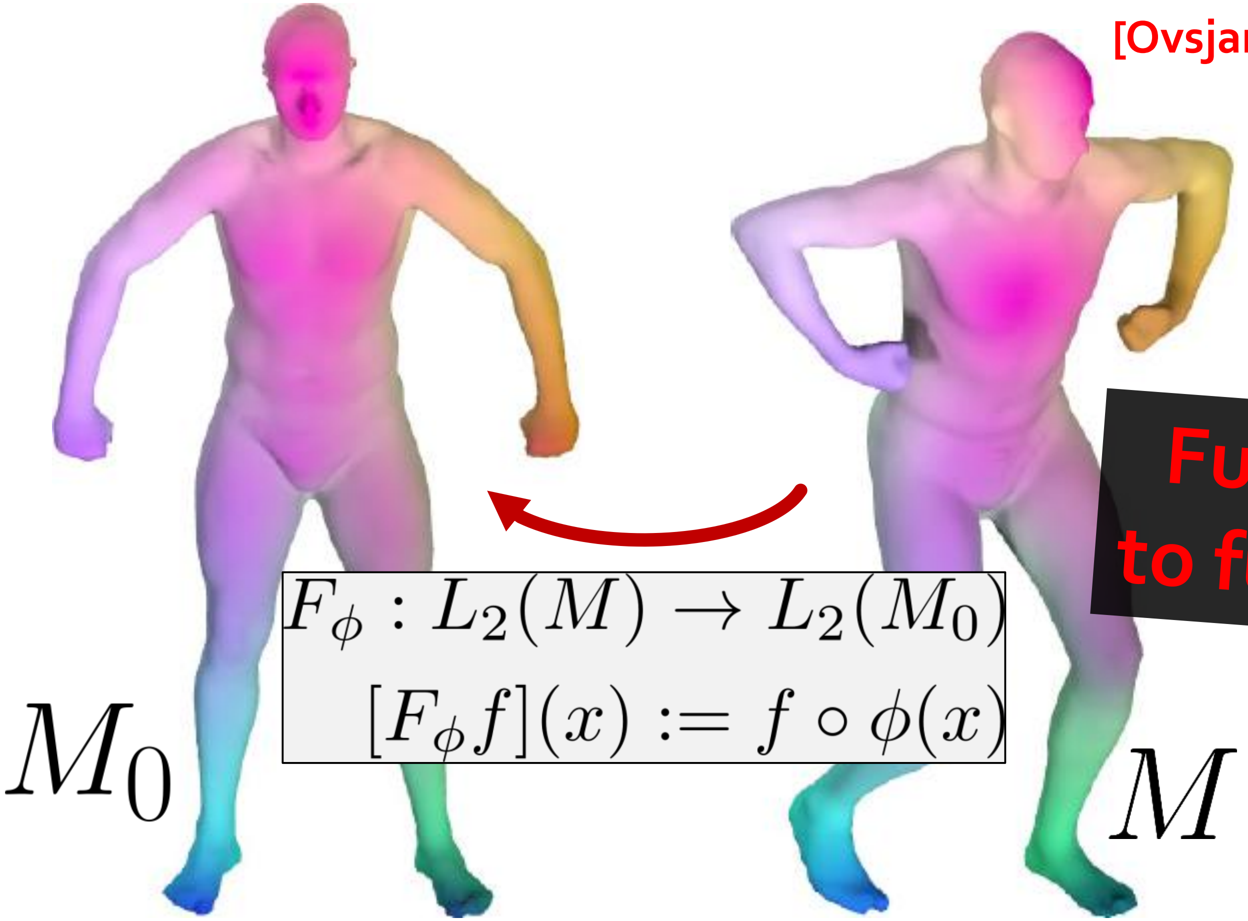

$$F_\phi : L_2(M) \to L_2(M_0)$$
$$[F_\phi f](x) := f \circ \phi(x)$$

Functions to functions

Functions on *M* to functions on *M*$_0$

Mathematical Sidebar

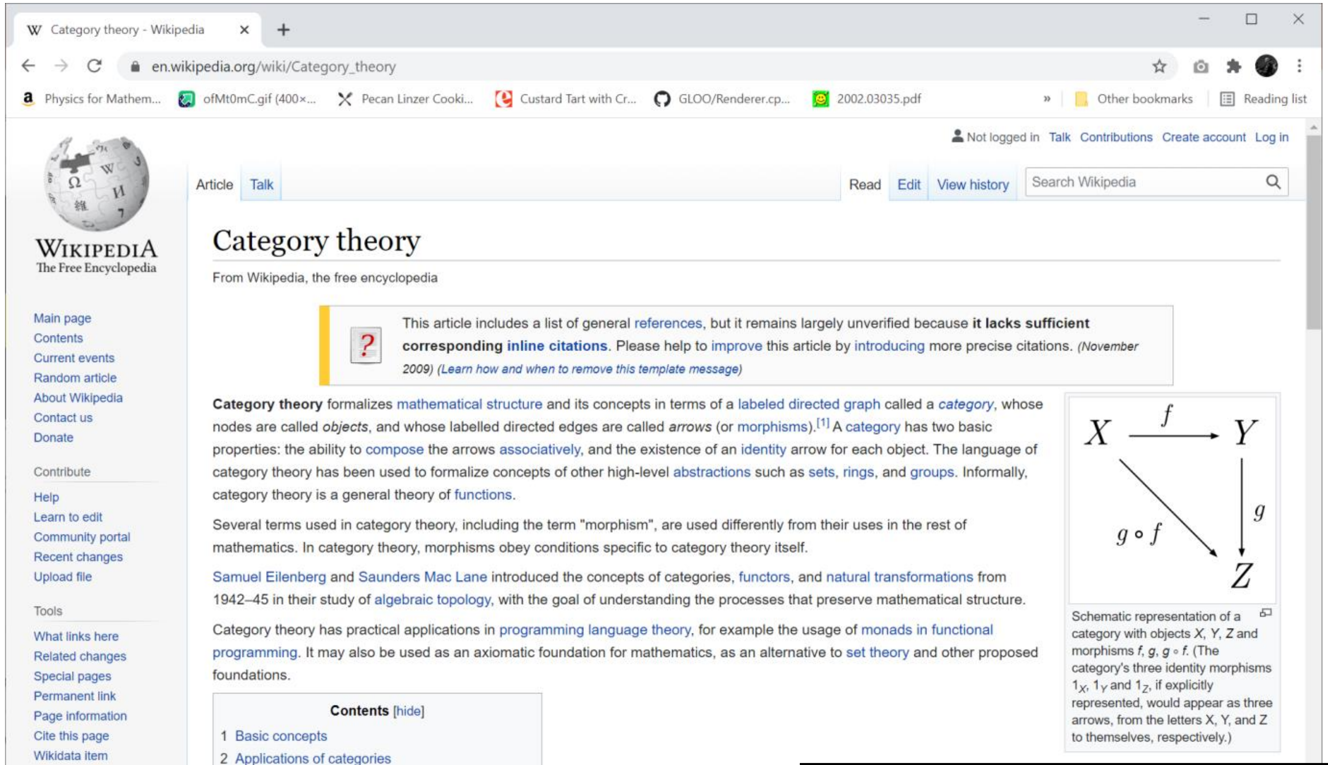

Main page
Contents
Current events
Random article
About Wikipedia
Contact us
Donate

Contribute

Help
Learn to edit
Community portal
Recent changes
Upload file

Tools

What links here
Related changes
Special pages
Permanent link
Page information
Cite this page
Wikidata item

Print/export

Download as PDF
Printable version

In other projects

Wikimedia Commons
Wikibooks
Wikiquote

Languages

Deutsch
Español
Français
한국어
Italiano

Category theory

From Wikipedia, the free encyclopedia

This article includes a list of general references, but it remains largely unverified because **it lacks sufficient corresponding inline citations**. Please help to improve this article by introducing more precise citations. *(November 2009)* *(Learn how and when to remove this template message)*

$$\begin{tikzcd}
X \arrow[r, "f"] \arrow[dr, "g \circ f"'] & Y \arrow[d, "g"] \\
 & Z
\end{tikzcd}$$

Schematic representation of a category with objects *X*, *Y*, *Z* and morphisms *f*, *g*, *g* ∘ *f*. (The category's three identity morphisms 1_X, 1_Y and 1_Z, if explicitly represented, would appear as three arrows, from the letters X, Y, and Z to themselves, respectively.)

Category theory formalizes mathematical structure and its concepts in terms of a labeled directed graph called a *category*, whose nodes are called *objects*, and whose labelled directed edges are called *arrows* (or morphisms).[1] A category has two basic properties: the ability to compose the arrows associatively, and the existence of an identity arrow for each object. The language of category theory has been used to formalize concepts of other high-level abstractions such as sets, rings, and groups. Informally, category theory is a general theory of functions.

Several terms used in category theory, including the term "morphism", are used differently from their uses in the rest of mathematics. In category theory, morphisms obey conditions specific to category theory itself.

Samuel Eilenberg and Saunders Mac Lane introduced the concepts of categories, functors, and natural transformations from 1942–45 in their study of algebraic topology, with the goal of understanding the processes that preserve mathematical structure.

Category theory has practical applications in programming language theory, for example the usage of monads in functional programming. It may also be used as an axiomatic foundation for mathematics, as an alternative to set theory and other proposed foundations.

Contents [hide]

1 Basic concepts
2 Applications of categories
3 Utility
  3.1 Categories, objects, and morphisms
  3.2 Functors
  3.3 Natural transformations
4 Categories, objects, and morphisms
  4.1 Categories
  4.2 Morphisms
5 Functors
6 Natural transformations
7 Other concepts
  7.1 Universal constructions, limits, and colimits
  7.2 Equivalent categories
  7.3 Further concepts and results
  7.4 Higher-dimensional categories

Functional Maps

[Ovsjanikov et al. 2012]

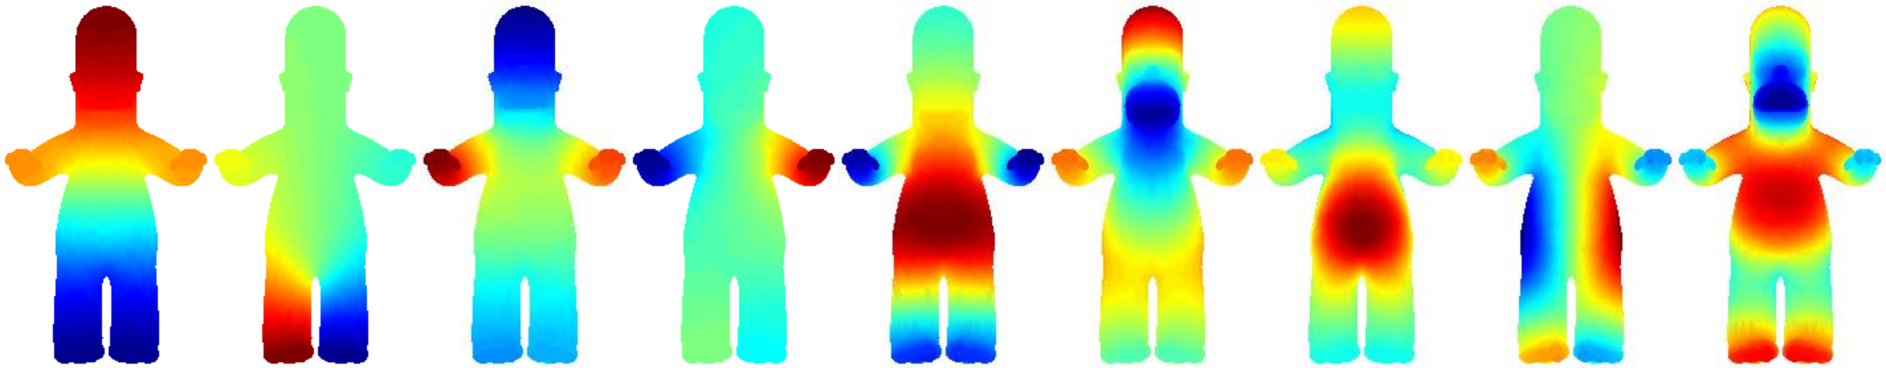

$$f(x) = \sum_i a_i \psi_i(x)$$

Functional map:

Matrix taking Laplace-Beltrami (Fourier) coefficients on *M* to coefficients on M_0

Example Maps

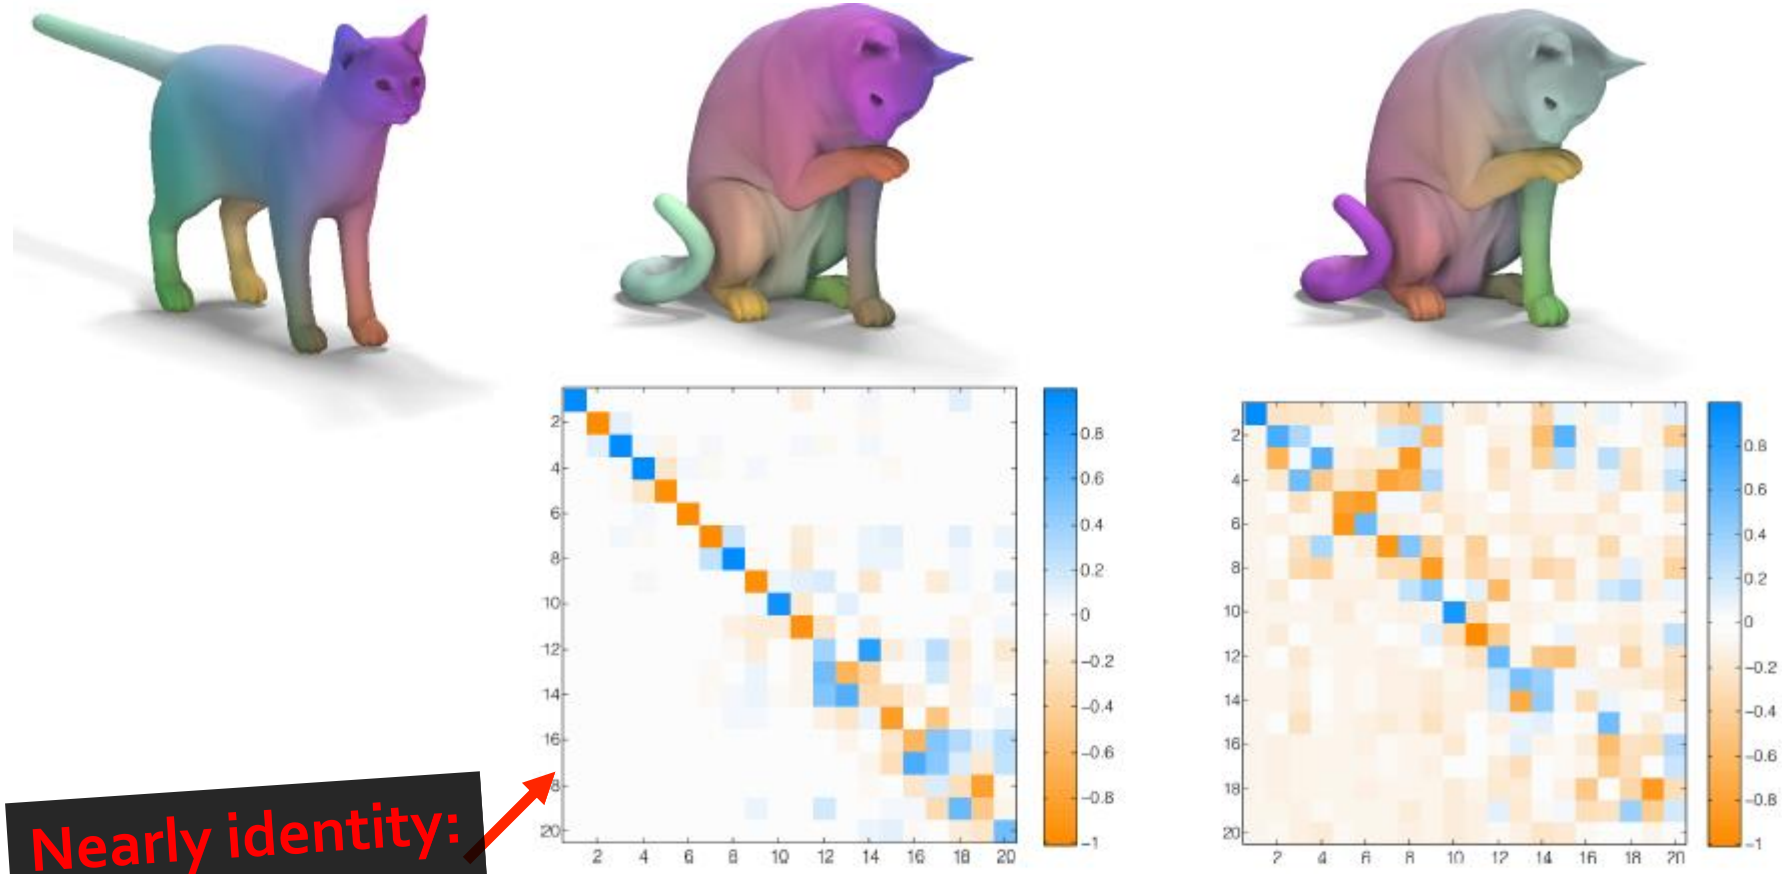

Nearly identity: Not a mistake!

(c) *left to right map*

(d) *head to tail map*

Adapted from slides by Q. Huang, V. Kim

Functional Maps

- Simple Algorithm
  - Compute some geometric functions to be preserved: A, B
  - Solve in least-squares sense for C: $B = C A$
- Additional Considerations
  - Favor commutativity
  - Favor orthonormality (if shapes are isometric)
  - Efficiently getting point-to-point correspondences

Adapted from slides by Q. Huang, V. Kim

Ovsjanikov'12

Tradeoff: Functional Maps

- Pros:
  - Condensed representation
  - Linear
  - Alternative perspective on mapping
  - Many recent papers with variations
- Cons:
  - Hard to handle non-isometry

    Some progress in last few years!

Adapted from slides by Q. Huang, V. Kim

Other Operators for Commutativity

- Compose with inverse map for identity [Eynard et al. 2016]
- Laplacian of displaced mesh [Corman et al. 2017]
- Diagonal operator from descriptor [Nogneng and Ovsjanikov 2017]
- Infinitesimal displacement rate of change of Laplacian [Corman and Ovsjanikov 2018]
- Kernel matrix [Wang et al. 2018]
- Operators built from matched curves [Gehre et al. 2018]
- Pointwise products of functions [Nogneng et al. 2018]
- Subdivision hierarchies [Shoham et al. 2019]
- Resolvent of Laplacian operator [Ren et al. 2019]

…and others

Example extension:

Coupled Quasi-Harmonic Basis

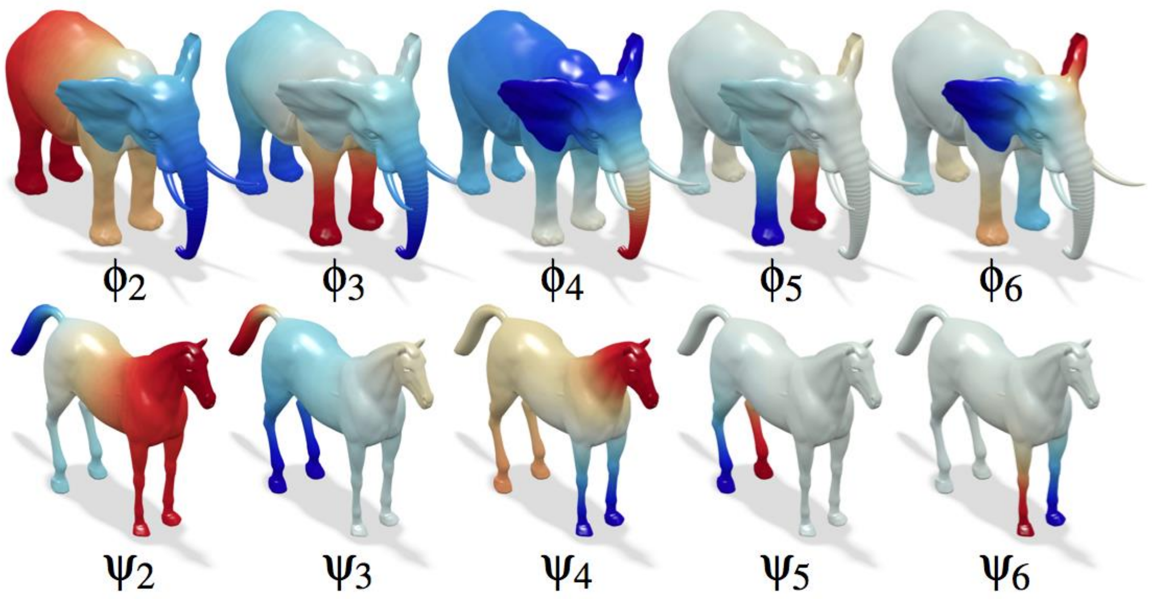

Laplacian eigenbases

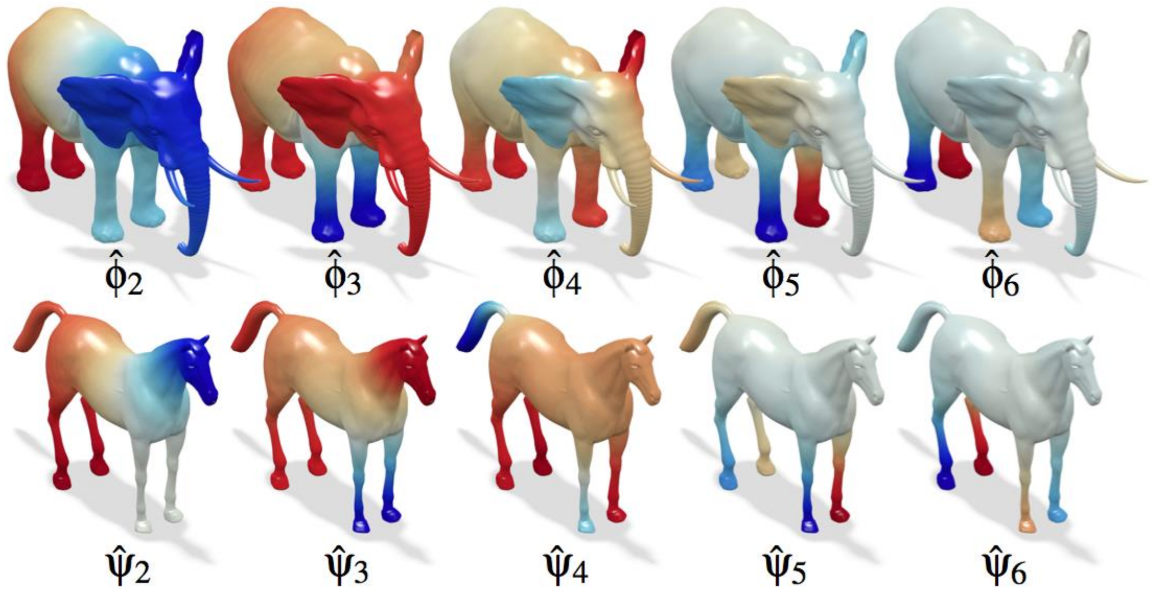

Coupled quasi-harmonic bases

$$\begin{aligned} \min_{\Phi,\Psi} \quad & \mathrm{off}(\Phi^\top W_X \Phi) + \mathrm{off}(\Psi^\top W_Y \Psi) + \mu \| F^\top \Phi - G^\top \Psi \|_{\mathrm{Fro}}^2 \\ \text{s.t.} \quad & \Phi^\top D_X \Phi = I \\ & \Psi^\top D_Y \Psi = I \end{aligned}$$

[Kovnatsky et al. 2012]

Example extension:

Leverage Symmetry

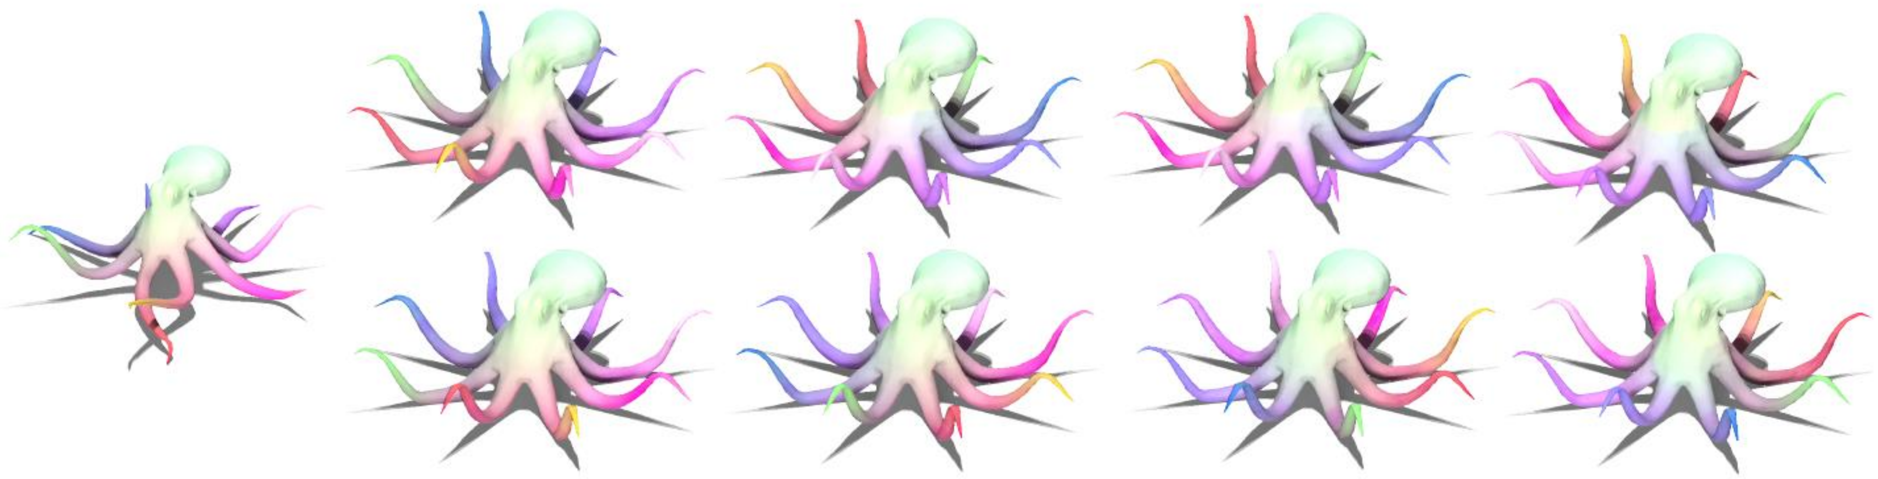

- Symmetry generators are self-maps
- Can quotient functional spaces by symmetries

"Shape Matching via Quotient Spaces," Ovsjankov et al. 2013

Example extension:

Map Upsampling

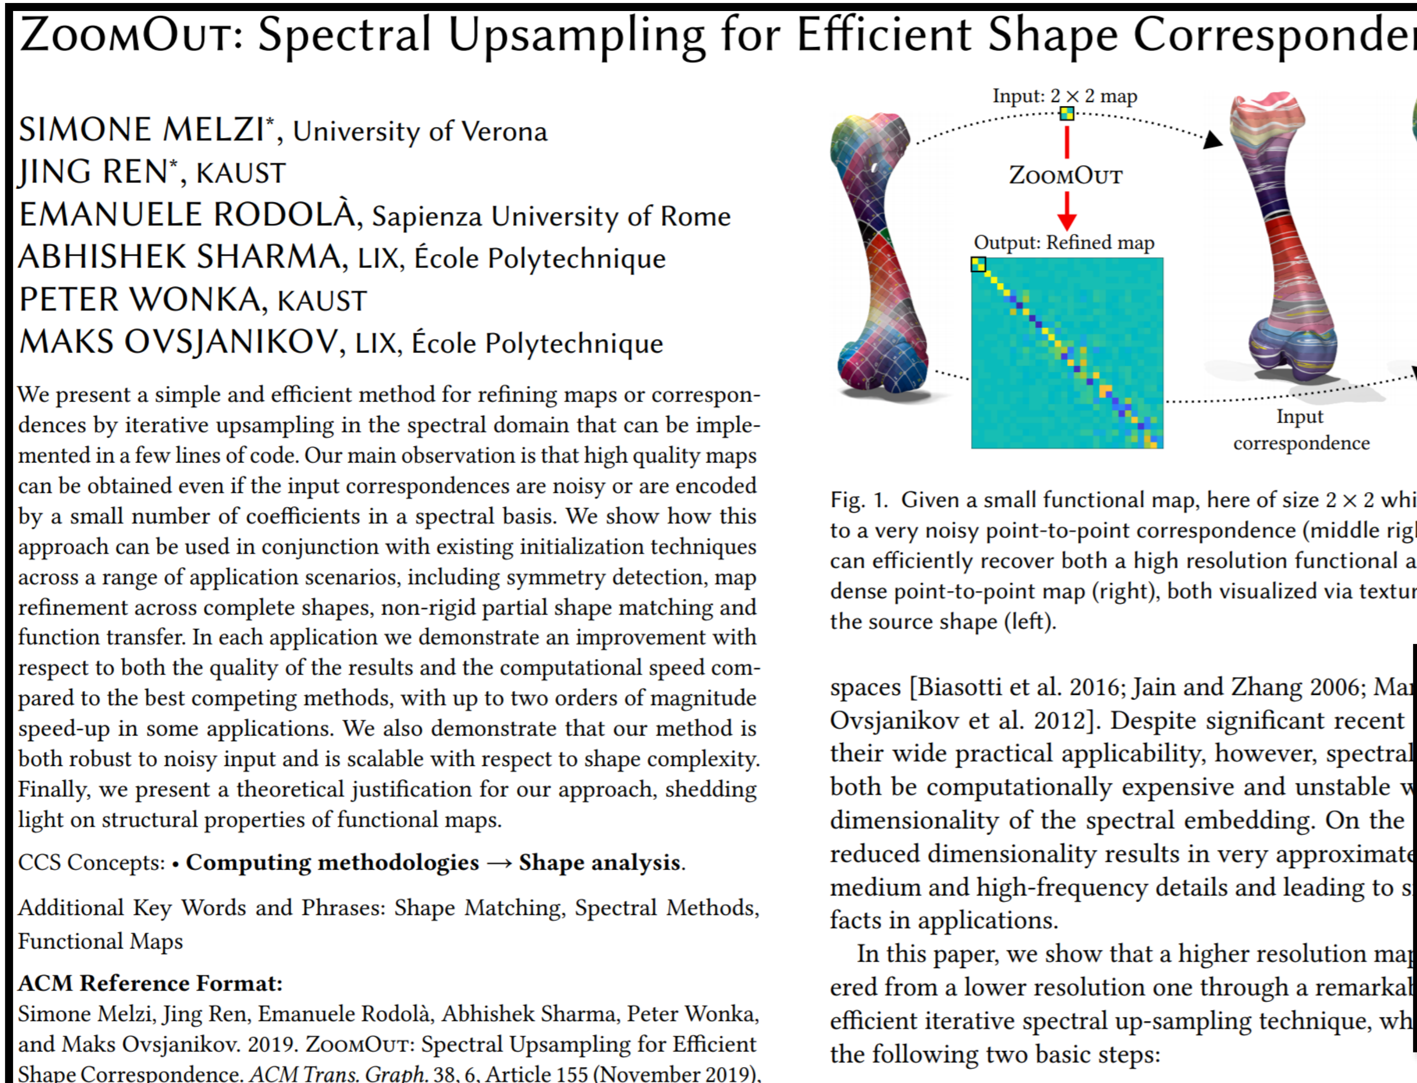

ZoomOut: Spectral Upsampling for Efficient Shape Correspondence

SIMONE MELZI*, University of Verona
JING REN*, KAUST
EMANUELE RODOLÀ, Sapienza University of Rome
ABHISHEK SHARMA, LIX, École Polytechnique
PETER WONKA, KAUST
MAKS OVSJANIKOV, LIX, École Polytechnique

We present a simple and efficient method for refining maps or correspondences by iterative upsampling in the spectral domain that can be implemented in a few lines of code. Our main observation is that high quality maps can be obtained even if the input correspondences are noisy or are encoded by a small number of coefficients in a spectral basis. We show how this approach can be used in conjunction with existing initialization techniques across a range of application scenarios, including symmetry detection, map refinement across complete shapes, non-rigid partial shape matching and function transfer. In each application we demonstrate an improvement with respect to both the quality of the results and the computational speed compared to the best competing methods, with up to two orders of magnitude speed-up in some applications. We also demonstrate that our method is both robust to noisy input and is scalable with respect to shape complexity. Finally, we present a theoretical justification for our approach, shedding light on structural properties of functional maps.

CCS Concepts: • **Computing methodologies** → **Shape analysis**.

Additional Key Words and Phrases: Shape Matching, Spectral Methods, Functional Maps

ACM Reference Format:
Simone Melzi, Jing Ren, Emanuele Rodolà, Abhishek Sharma, Peter Wonka, and Maks Ovsjanikov. 2019. ZoomOut: Spectral Upsampling for Efficient Shape Correspondence. *ACM Trans. Graph.* 38, 6, Article 155 (November 2019),

Fig. 1. Given a small functional map, here of size 2 × 2 which corresponds to a very noisy point-to-point correspondence (middle right) our method can efficiently recover both a high resolution functional and an accurate dense point-to-point map (right), both visualized via texture transfer from the source shape (left).

spaces [Biasotti et al. 2016; Jain and Zhang 2006; Ma
Ovsjanikov et al. 2012]. Despite significant recent
their wide practical applicability, however, spectral
both be computationally expensive and unstable w
dimensionality of the spectral embedding. On the
reduced dimensionality results in very approximate
medium and high-frequency details and leading to s
facts in applications.

In this paper, we show that a higher resolution map
ered from a lower resolution one through a remarkab
efficient iterative spectral up-sampling technique, wh
the following two basic steps:

(1) Given $k = k_0$ and an initial C_0 of size $k_0 \times k_0$.
(2) Compute $\arg\min_{\Pi} \|\Pi\Phi_{\mathcal{N}}^{k} C_{k}^{T} - \Phi_{\mathcal{M}}^{k}\|_F^2$.
(3) Set $k = k + 1$ and compute $C_k = (\Phi_{\mathcal{M}}^{k})^{+}\Pi\Phi_{\mathcal{N}}^{k}$.
(4) Repeat the previous two steps until $k = k_{\max}$.

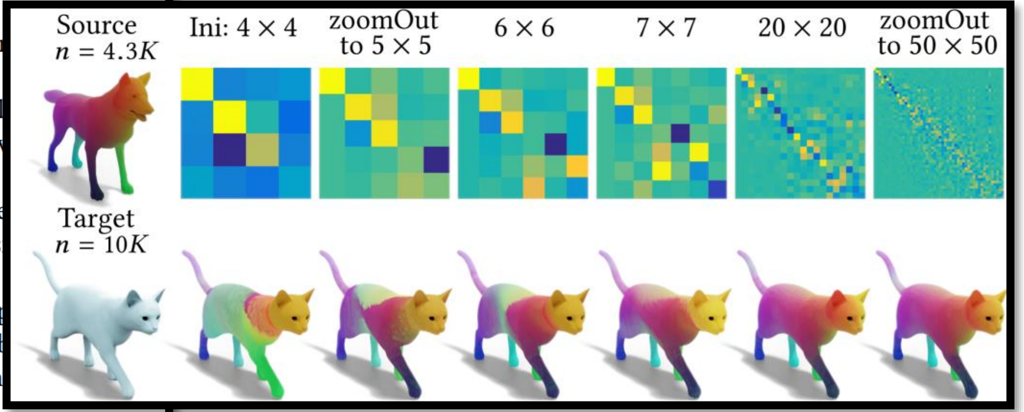

Example application:

Shape Differences

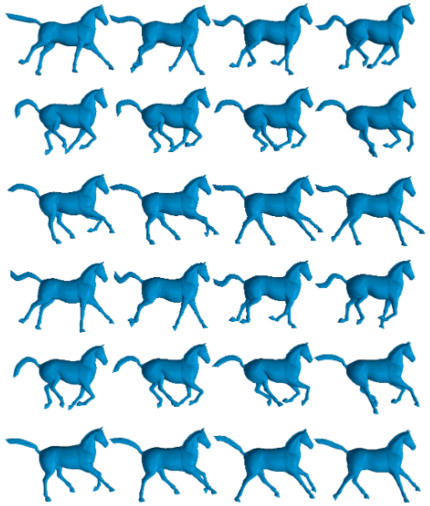

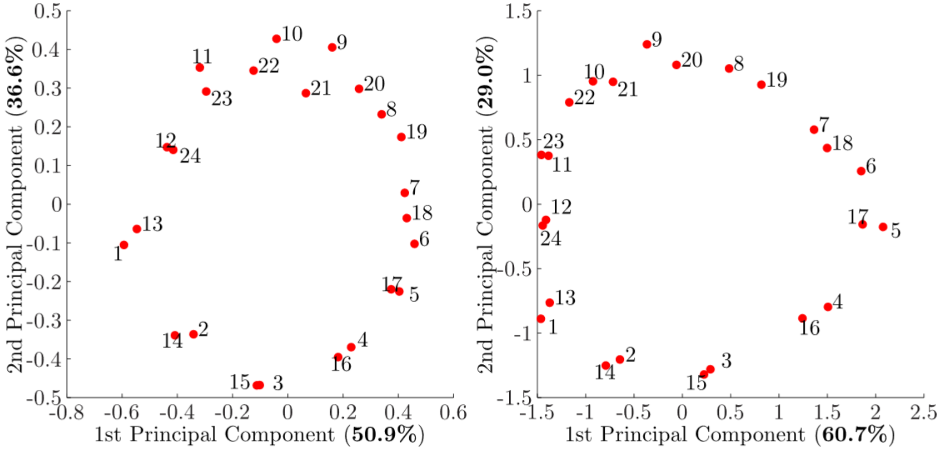

$$D = (H^M)^{-1} F^\top H^N F$$

“Map-based exploration of intrinsic shape differences and variability” (Rustamov et al., 2013)

Inner Products

[Rustamov et al. 2013]

$$\langle f, g \rangle_A := \int_M f(x) g(x) \, dA$$

$$\langle f, g \rangle_C := \int_M [\nabla f(x) \cdot \nabla g(x)] \, dA$$

Object of study:

Inner product matrix

$$M_{ij} := \langle \psi_i, \psi_j \rangle$$

Shape Differences

[Rustamov et al. 2013]

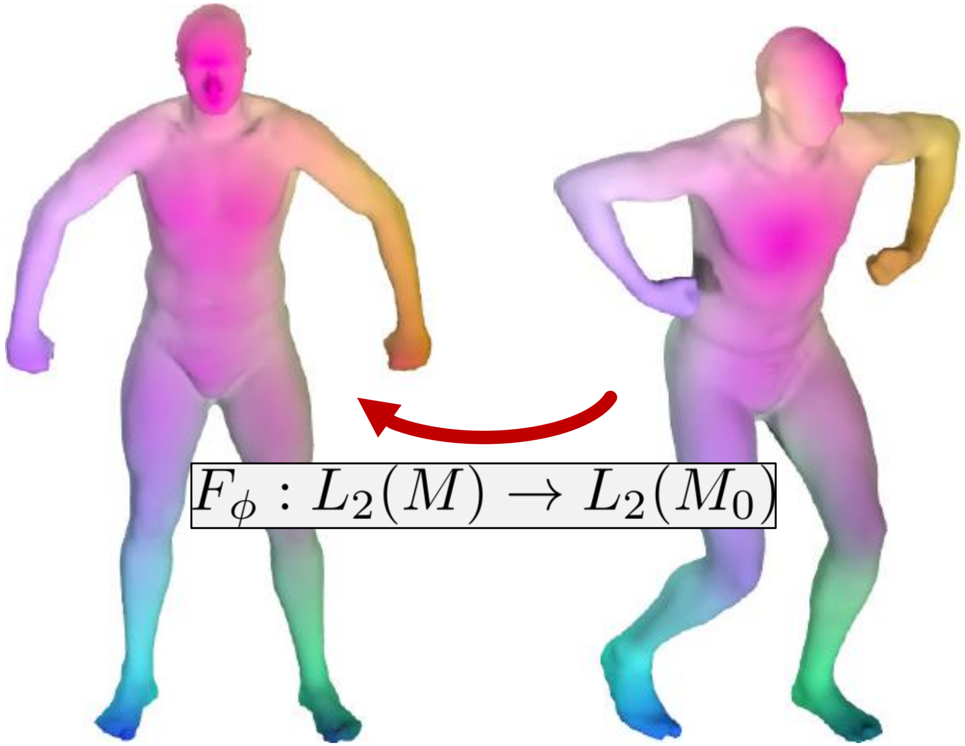

Trick:

Compare surfaces by comparing inner product matrices.

$$\langle f, g \rangle_F^M := \langle F_\phi[f], F_\phi[g] \rangle^{M_0} \quad D = (H^M)^{-1} F^\top H^N F$$

Functional map *pulls back* products

Continuous Question

[Rustamov et al. 2013]

Given

area-based and conformal inner product matrices,

can you compute

lengths and angles? ✓

Discrete Question

[Corman et al. 2017]

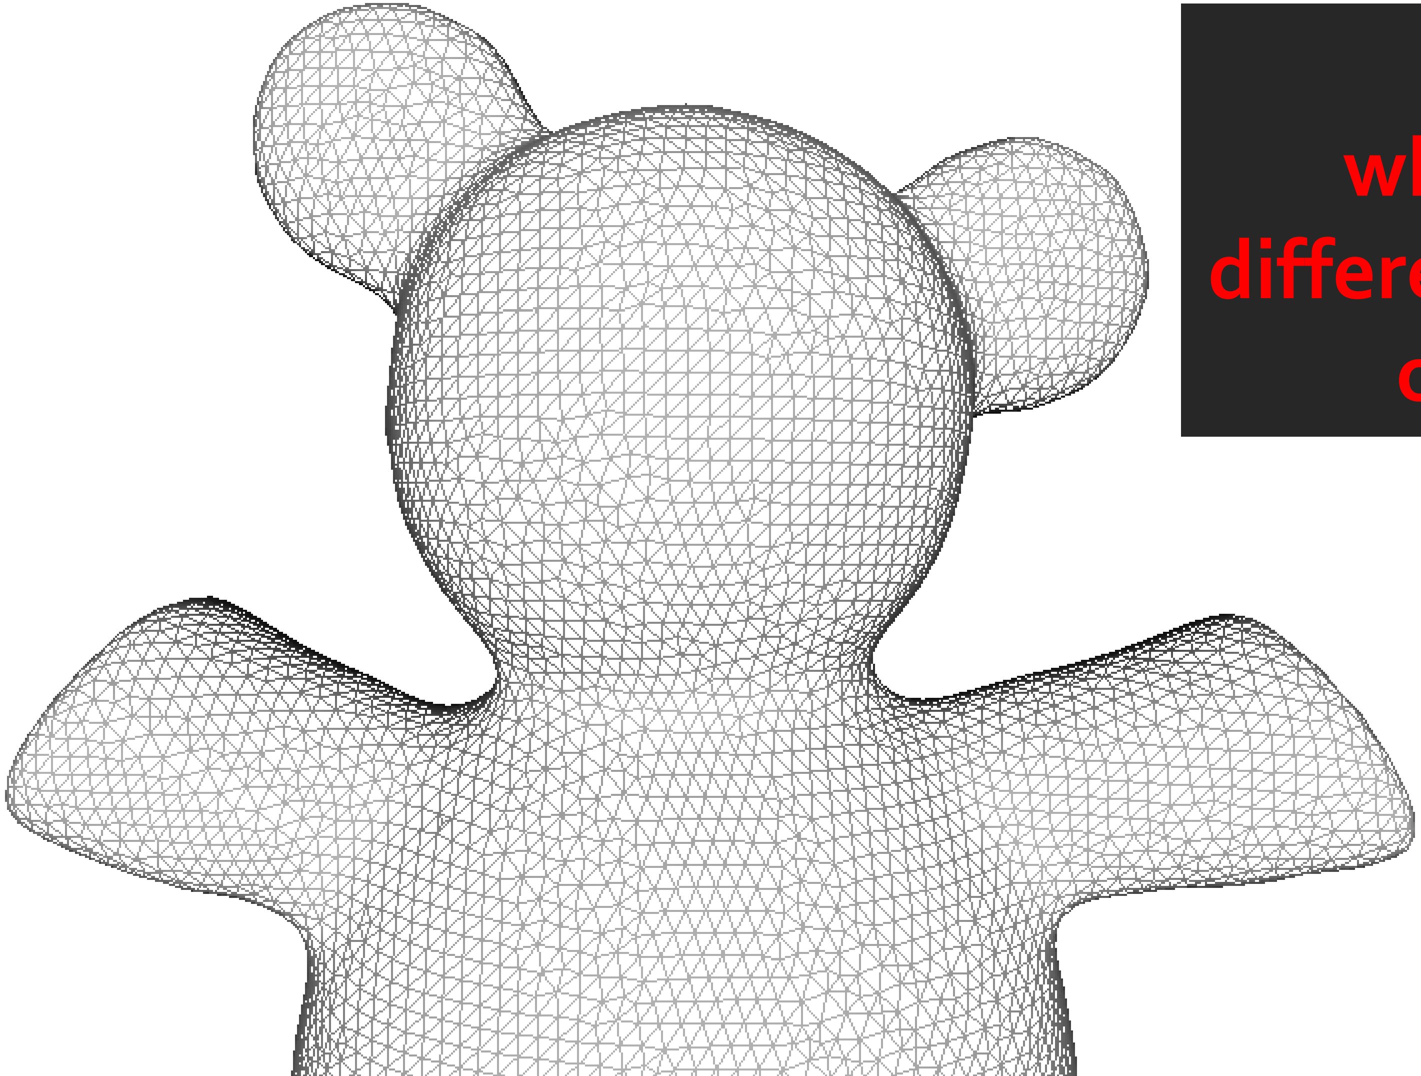

Precisely
what do shape differences determine on meshes?

✓ Edge lengths.

Extension to Extrinsic Shape

[Corman et al. 2017]

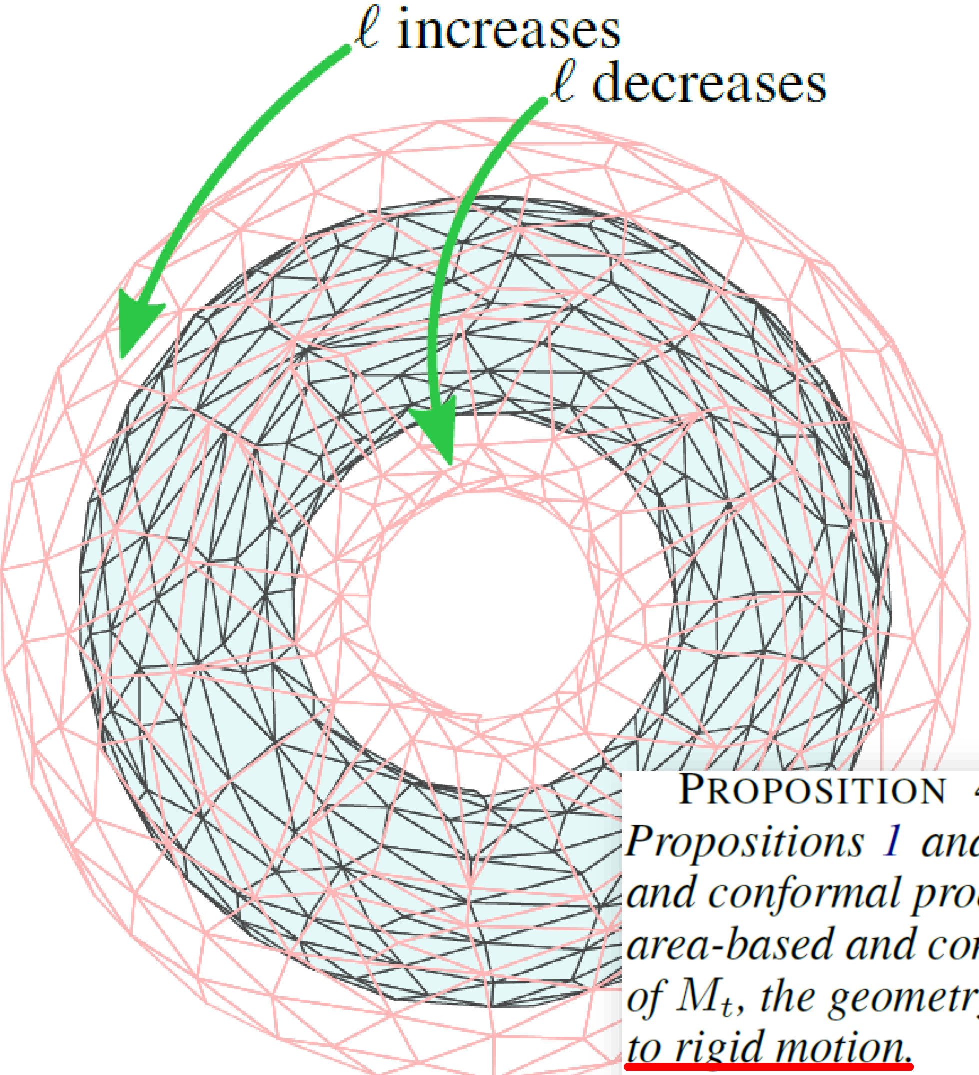

Throw in the offset surface.

Encodes mean curvature!

PROPOSITION 4. *Suppose a mesh M satisfies the criteria in Propositions 1 and 2. Given the topology of M, the area-based and conformal product matrices $A(\mu)$ and $C(\nu;\mu)$ of M, and the area-based and conformal product matrices $A_t(\mu_t)$ and $C(\nu_t;\mu_t)$ of M_t, the geometry of M can (almost always) be reconstructed up to rigid motion.*

Useful Survey

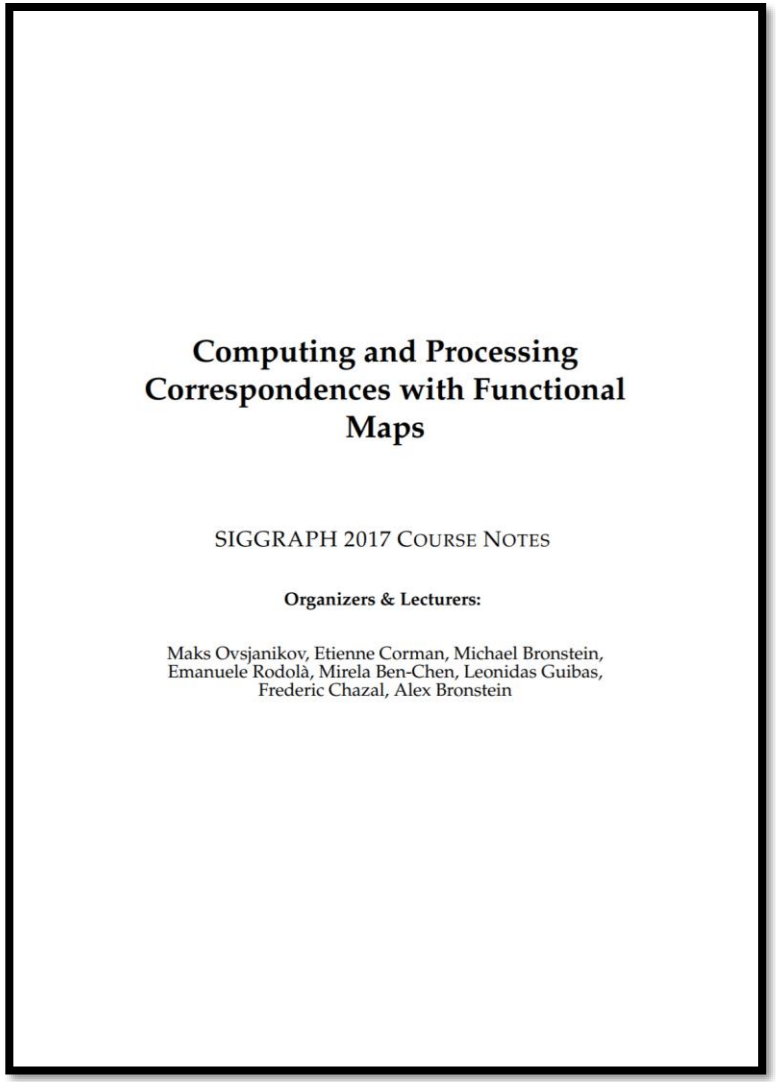

Computing and Processing Correspondences with Functional Maps

SIGGRAPH 2017 Course Notes

Organizers & Lecturers:

Maks Ovsjanikov, Etienne Corman, Michael Bronstein,
Emanuele Rodolà, Mirela Ben-Chen, Leonidas Guibas,
Frederic Chazal, Alex Bronstein

Deep Functional Maps

Many extensions:

- Incorporate commutativity
- Learn shape features
- Self-supervision

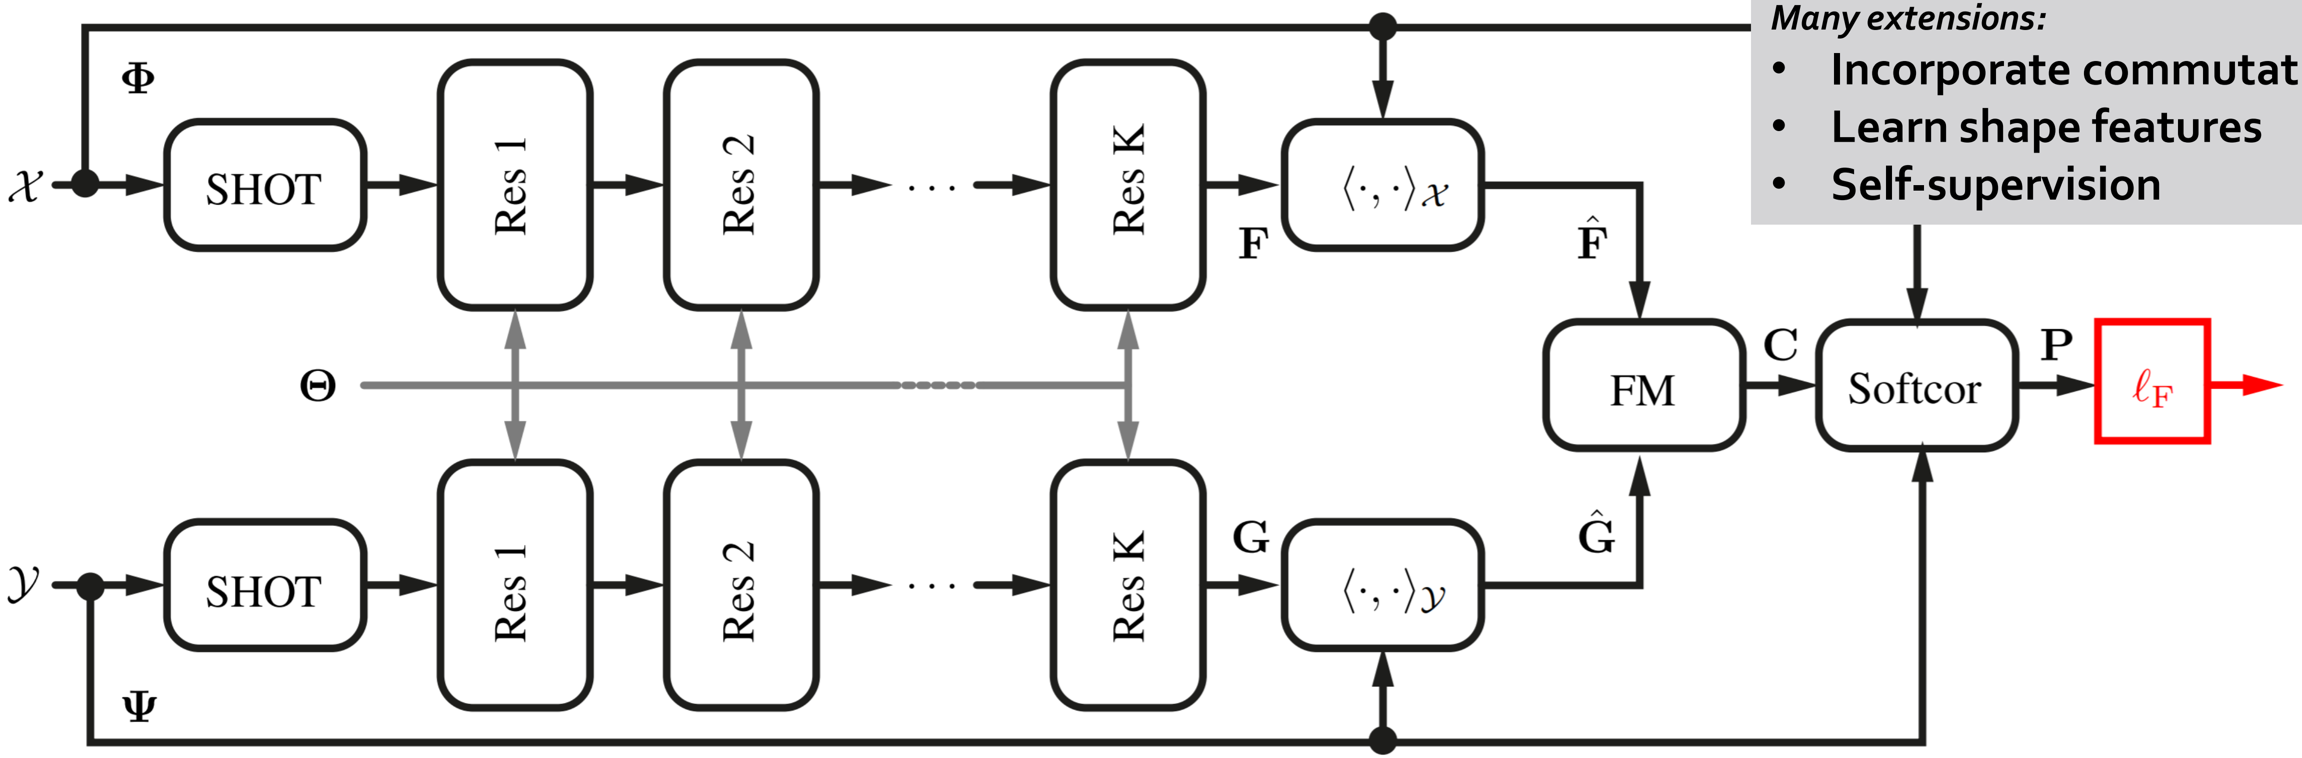

Figure 3. **FMNet architecture.** Input point-wise descriptors (SHOT [38] in this paper) from a pair of shapes are passed through an identical sequence of operations (with shared weights), resulting in refined descriptors $\mathbf{F}, \mathbf{G}$. These, in turn, are projected onto the Laplacian eigenbases $\boldsymbol{\Phi}, \boldsymbol{\Psi}$ to produce the spectral representations $\hat{\mathbf{F}}, \hat{\mathbf{G}}$. The functional map (FM) and soft correspondence (Softcor) layers, implementing Equations (3) and (6) respectively, are not parametric and are used to set up the geometrically structured loss ℓ_{F} (5).

“Deep functional maps: Structured prediction for dense shape correspondence” (Litany et al. 2017)

Correspondence Problems

Justin Solomon

6.838: Shape Analysis

Spring 2021

Extra: Reversible Harmonic Maps

Justin Solomon

6.838: Shape Analysis

Spring 2021

Reversible Harmonic Maps

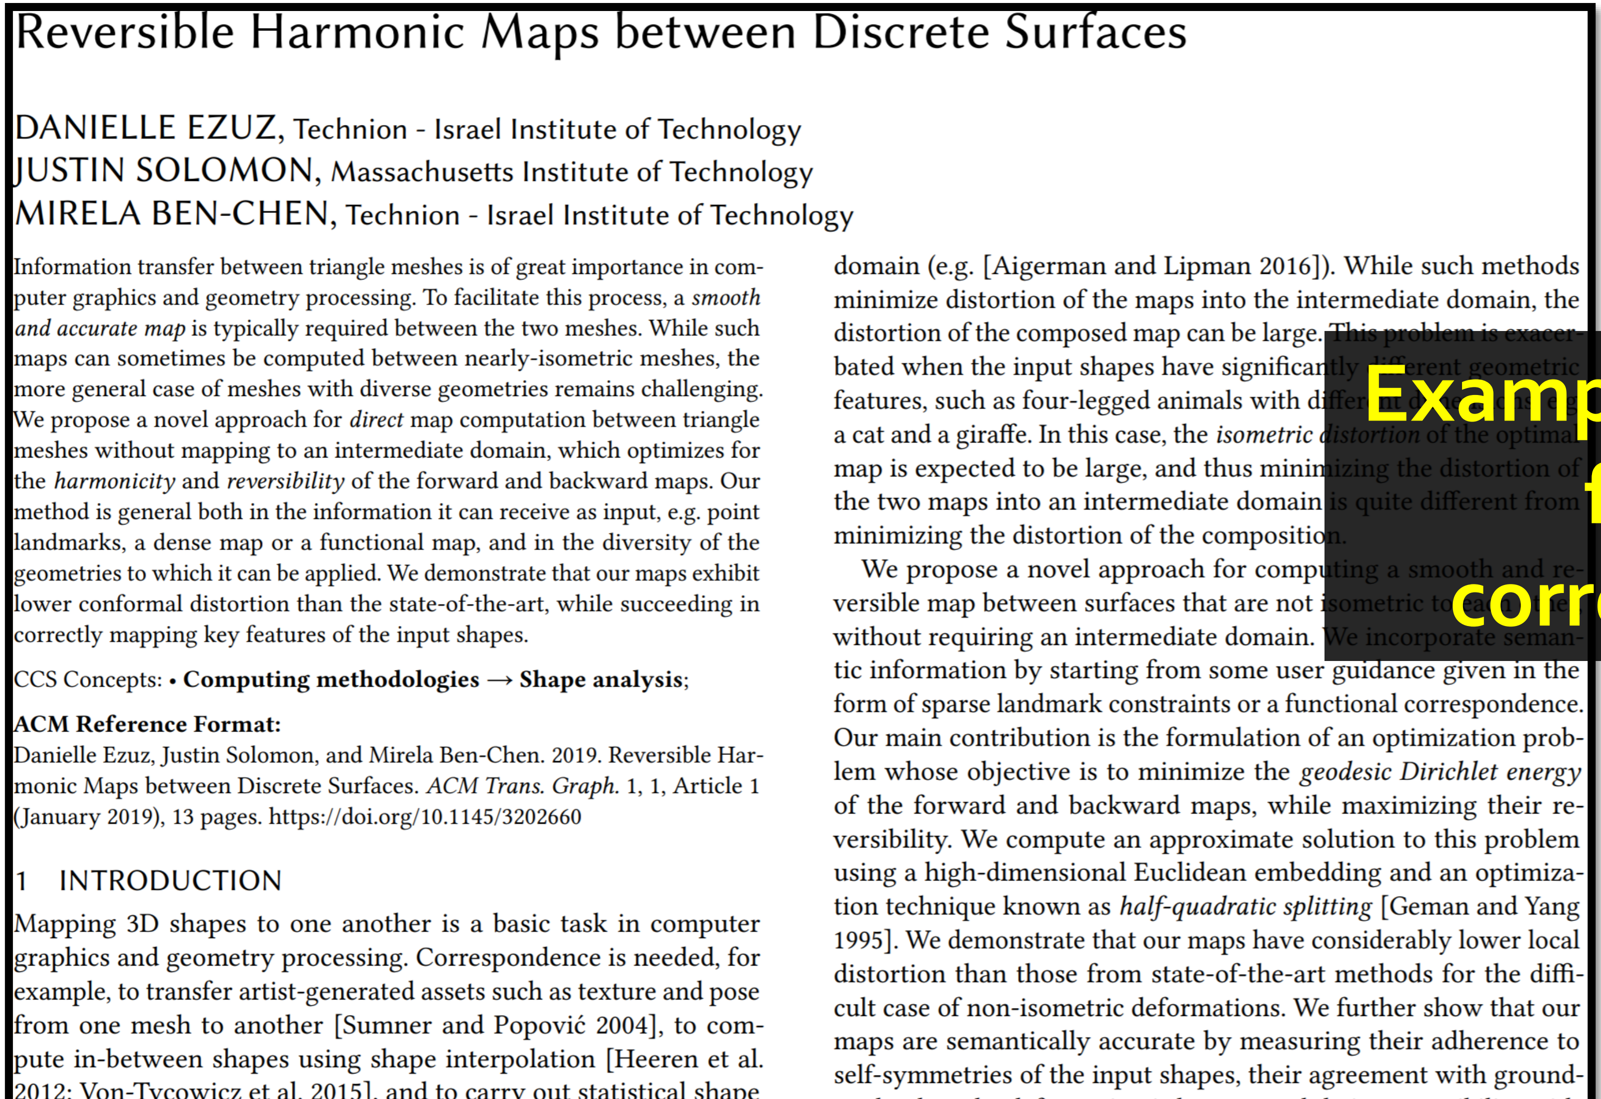

Reversible Harmonic Maps between Discrete Surfaces

DANIELLE EZUZ, Technion - Israel Institute of Technology
JUSTIN SOLOMON, Massachusetts Institute of Technology
MIRELA BEN-CHEN, Technion - Israel Institute of Technology

Information transfer between triangle meshes is of great importance in computer graphics and geometry processing. To facilitate this process, a *smooth and accurate map* is typically required between the two meshes. While such maps can sometimes be computed between nearly-isometric meshes, the more general case of meshes with diverse geometries remains challenging. We propose a novel approach for *direct* map computation between triangle meshes without mapping to an intermediate domain, which optimizes for the *harmonicity* and *reversibility* of the forward and backward maps. Our method is general both in the information it can receive as input, e.g. point landmarks, a dense map or a functional map, and in the diversity of the geometries to which it can be applied. We demonstrate that our maps exhibit lower conformal distortion than the state-of-the-art, while succeeding in correctly mapping key features of the input shapes.

CCS Concepts: • **Computing methodologies** → **Shape analysis**;

ACM Reference Format:
Danielle Ezuz, Justin Solomon, and Mirela Ben-Chen. 2019. Reversible Harmonic Maps between Discrete Surfaces. *ACM Trans. Graph.* 1, 1, Article 1 (January 2019), 13 pages. https://doi.org/10.1145/3202660

1 INTRODUCTION

Mapping 3D shapes to one another is a basic task in computer graphics and geometry processing. Correspondence is needed, for example, to transfer artist-generated assets such as texture and pose from one mesh to another [Sumner and Popović 2004], to compute in-between shapes using shape interpolation [Heeren et al. 2012; Von-Tycowicz et al. 2015], and to carry out statistical shape domain (e.g. [Aigerman and Lipman 2016]). While such methods minimize distortion of the maps into the intermediate domain, the distortion of the composed map can be large. This problem is exacerbated when the input shapes have significantly different geometric features, such as four-legged animals with different dimensions, e.g. a cat and a giraffe. In this case, the *isometric distortion* of the optimal map is expected to be large, and thus minimizing the distortion of the two maps into an intermediate domain is quite different from minimizing the distortion of the composition.

We propose a novel approach for computing a smooth and reversible map between surfaces that are not isometric to each other, without requiring an intermediate domain. We incorporate semantic information by starting from some user guidance given in the form of sparse landmark constraints or a functional correspondence. Our main contribution is the formulation of an optimization problem whose objective is to minimize the *geodesic Dirichlet energy* of the forward and backward maps, while maximizing their reversibility. We compute an approximate solution to this problem using a high-dimensional Euclidean embedding and an optimization technique known as *half-quadratic splitting* [Geman and Yang 1995]. We demonstrate that our maps have considerably lower local distortion than those from state-of-the-art methods for the difficult case of non-isometric deformations. We further show that our maps are semantically accurate by measuring their adherence to self-symmetries of the input shapes, their agreement with ground-

Example of a method for dense correspondence.

Slides courtesy D. Ezuz

Approach

Input: a sparse set of landmarks (p_i, q_i)

- Initialize the map by mapping **geodesic cells** of each landmark p_i to the corresponding landmark q_i

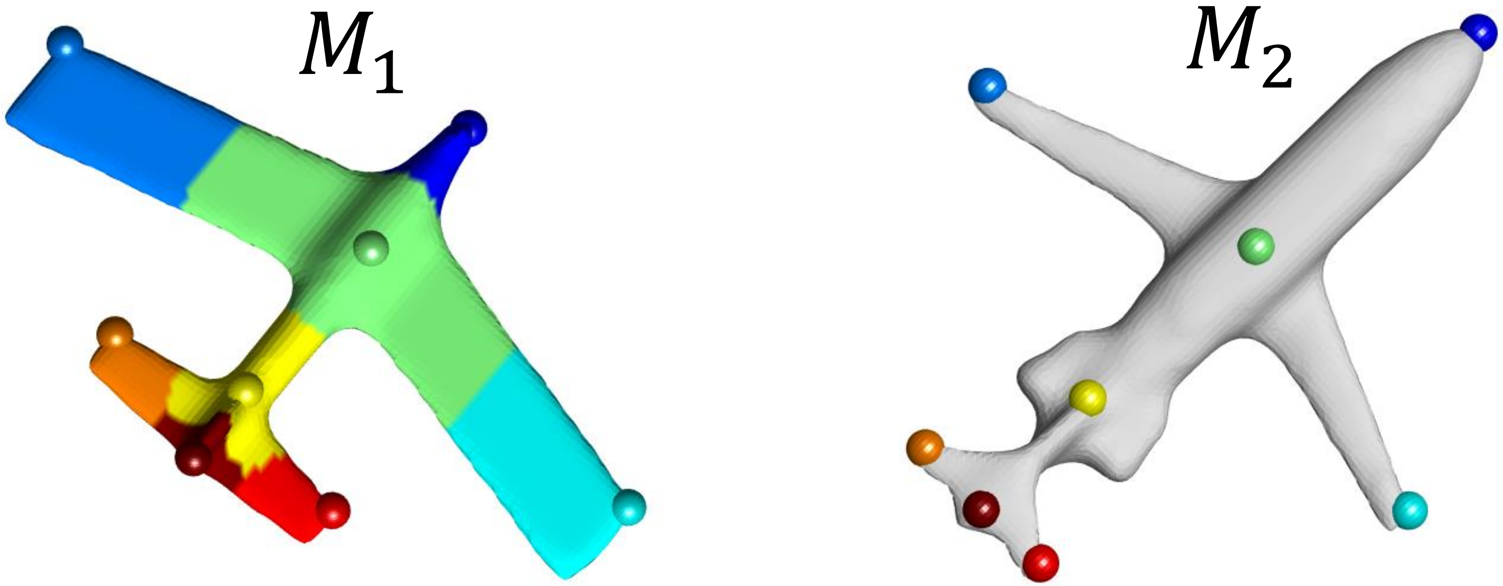

Approach

Input: a sparse set of landmarks (p_i, q_i)

- Initialize the map by mapping **geodesic cells** of each landmark p_i to the corresponding landmark q_i

- Optimize the map with respect to an **energy** that promotes **smoothness** and **bijectivity**

Approach

Measures **smoothness** of a map:

$$E(\phi_{12}) = \frac{1}{2}\int_{M_1} |d\phi_{12}|^2$$

A map is **harmonic** if it is a critical point of the Dirichlet energy

Approach

$$E_D(\phi_{12}) = \sum_{(u,v)\in E_1} w_{uv}\, d^2_{M_2}(\phi_{12}(u), \phi_{12}(v))$$

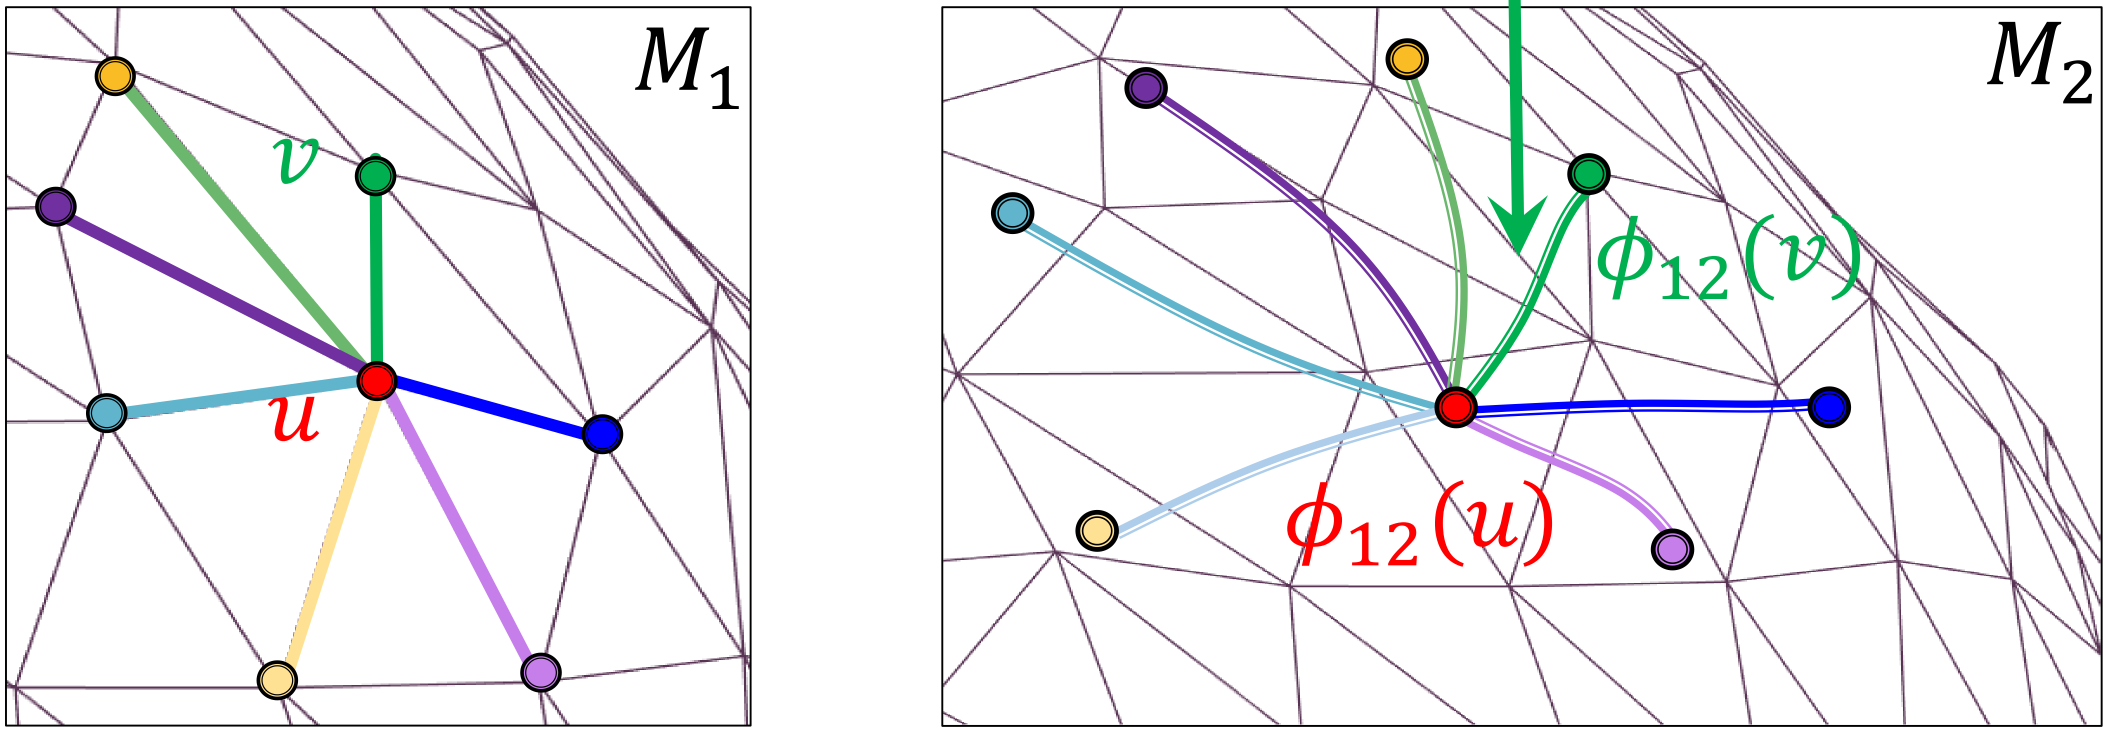

Discrete Precise Maps

Stochastic matrices with barycentric coordinates at each row:

$$P_{12} = \begin{array}{cc} \begin{array}{ccc} j & k & l \end{array} & \\ \left(\begin{array}{ccc} & \vdots & \\ \text{-}\,0.1\,\text{—} & 0.2\,\text{—} & 0.7\,\text{—} \\ & \vdots & \end{array}\right) & \text{row } i \end{array}$$

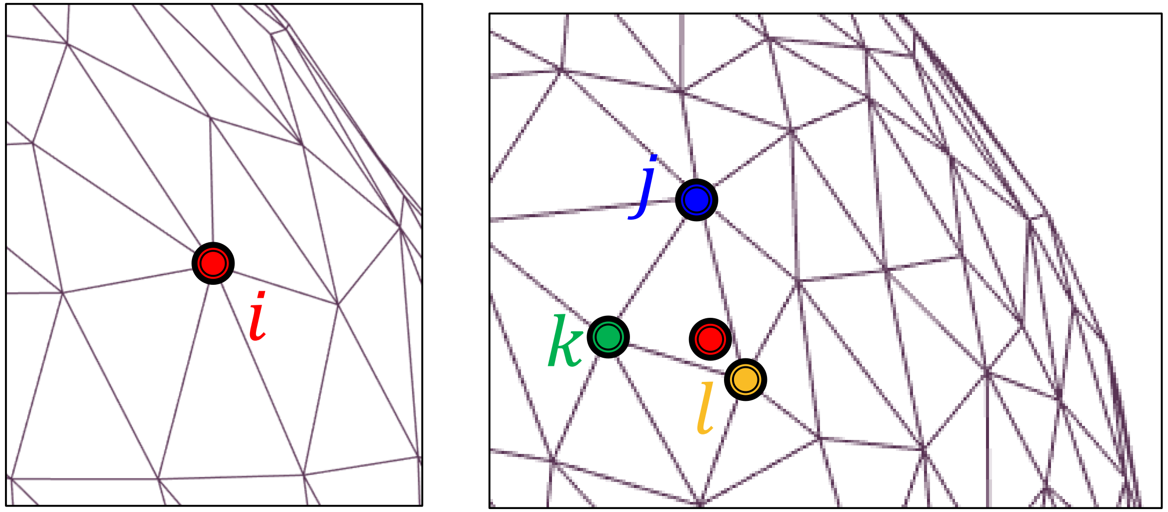

M_1 M_2

Discrete Precise Maps

Stochastic matrices with barycentric coordinates at each row:

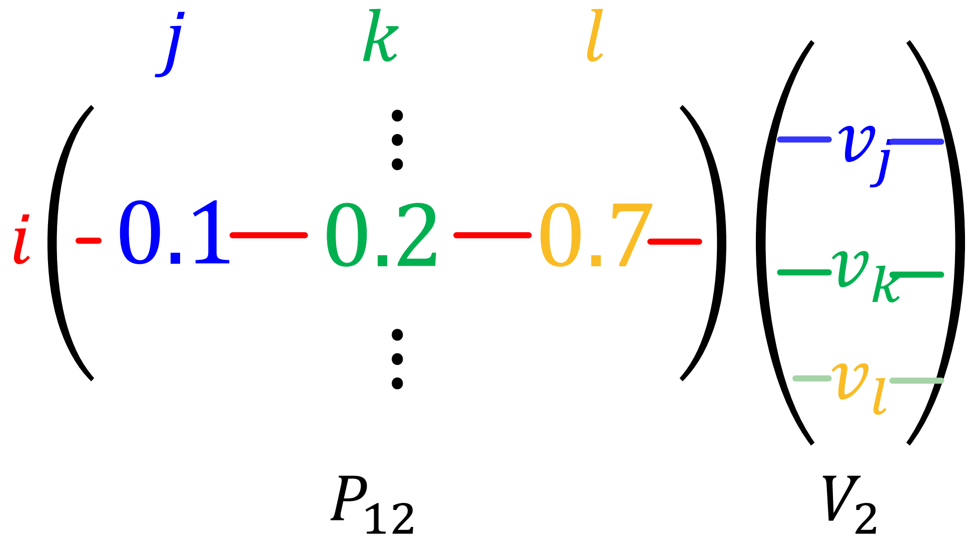

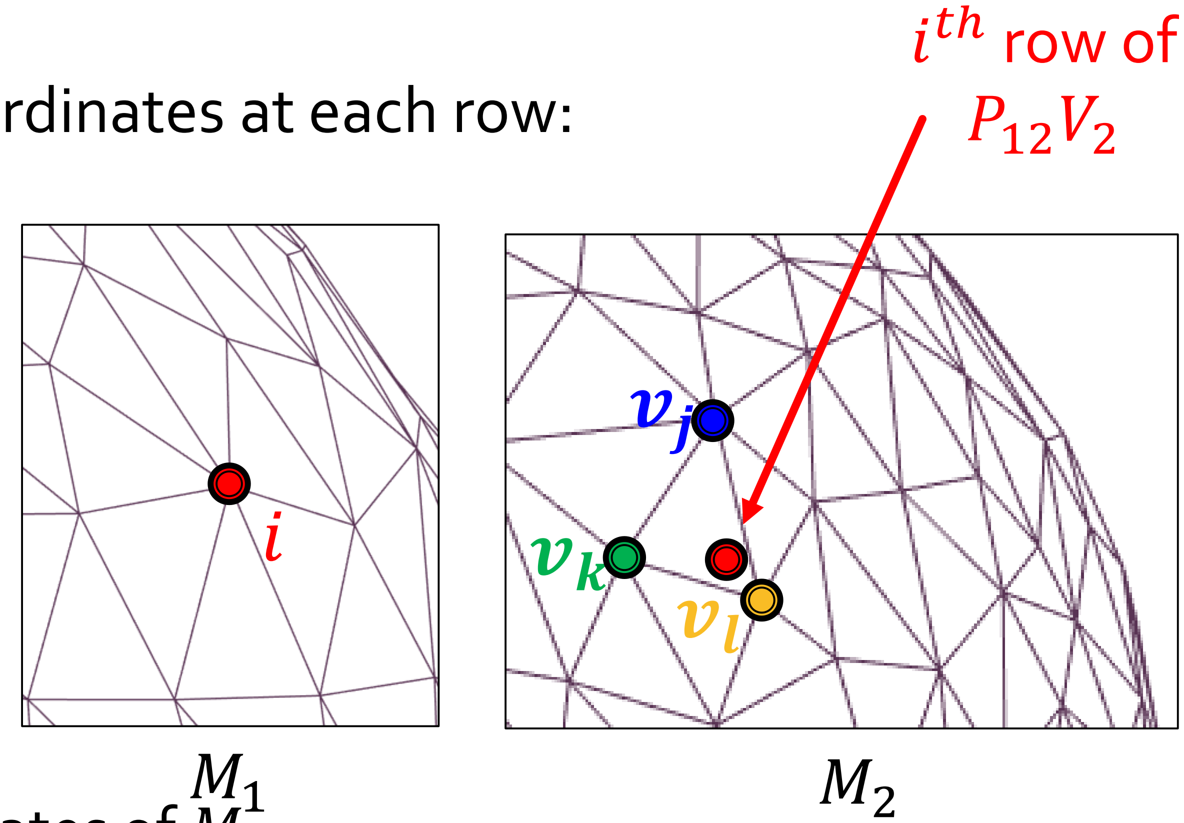

$V_2 \in \mathbb{R}^{n_2 \times 3}$ is a matrix with vertex coordinates of M_2

Discretization of Dirichlet Energy

If we replace the geodesic distances by Euclidean distances, the discrete Dirichlet energy is:

$$E_D^{Euc}(P_{12}) = \|P_{12}V_2\|_{W_1}^2 = Trace\left((P_{12}V_2)^\top W_1 P_{12}V_2\right)$$

W_1 is a matrix with $-w_{ij}$ at entry i, j, and the sum of the weights on the diagonal

$$\begin{array}{cc} & \begin{matrix} j & \quad i & \quad k \end{matrix} \\ i & \begin{pmatrix} -w_{ij} & \sum_v w_{iv} & -w_{ik} \\ & & \\ & & \end{pmatrix} \\ & W_1 \end{array}$$

Discrete Dirichlet Energy

We use a *high dimensional embedding* where Euclidean distances approximate geodesic distances (MDS)

$$X_2 \in \mathbb{R}^{n_2 \times 8}$$

Then, the discrete Dirichlet energy is approximated by:

$$E_D(P_{12}) = ||P_{12}X_2||_{W_1}^2$$

Minimizing the Dirichlet Energy

A map that maps all vertices to a single point is harmonic

Minimizing the harmonic energy “shrinks” the map:

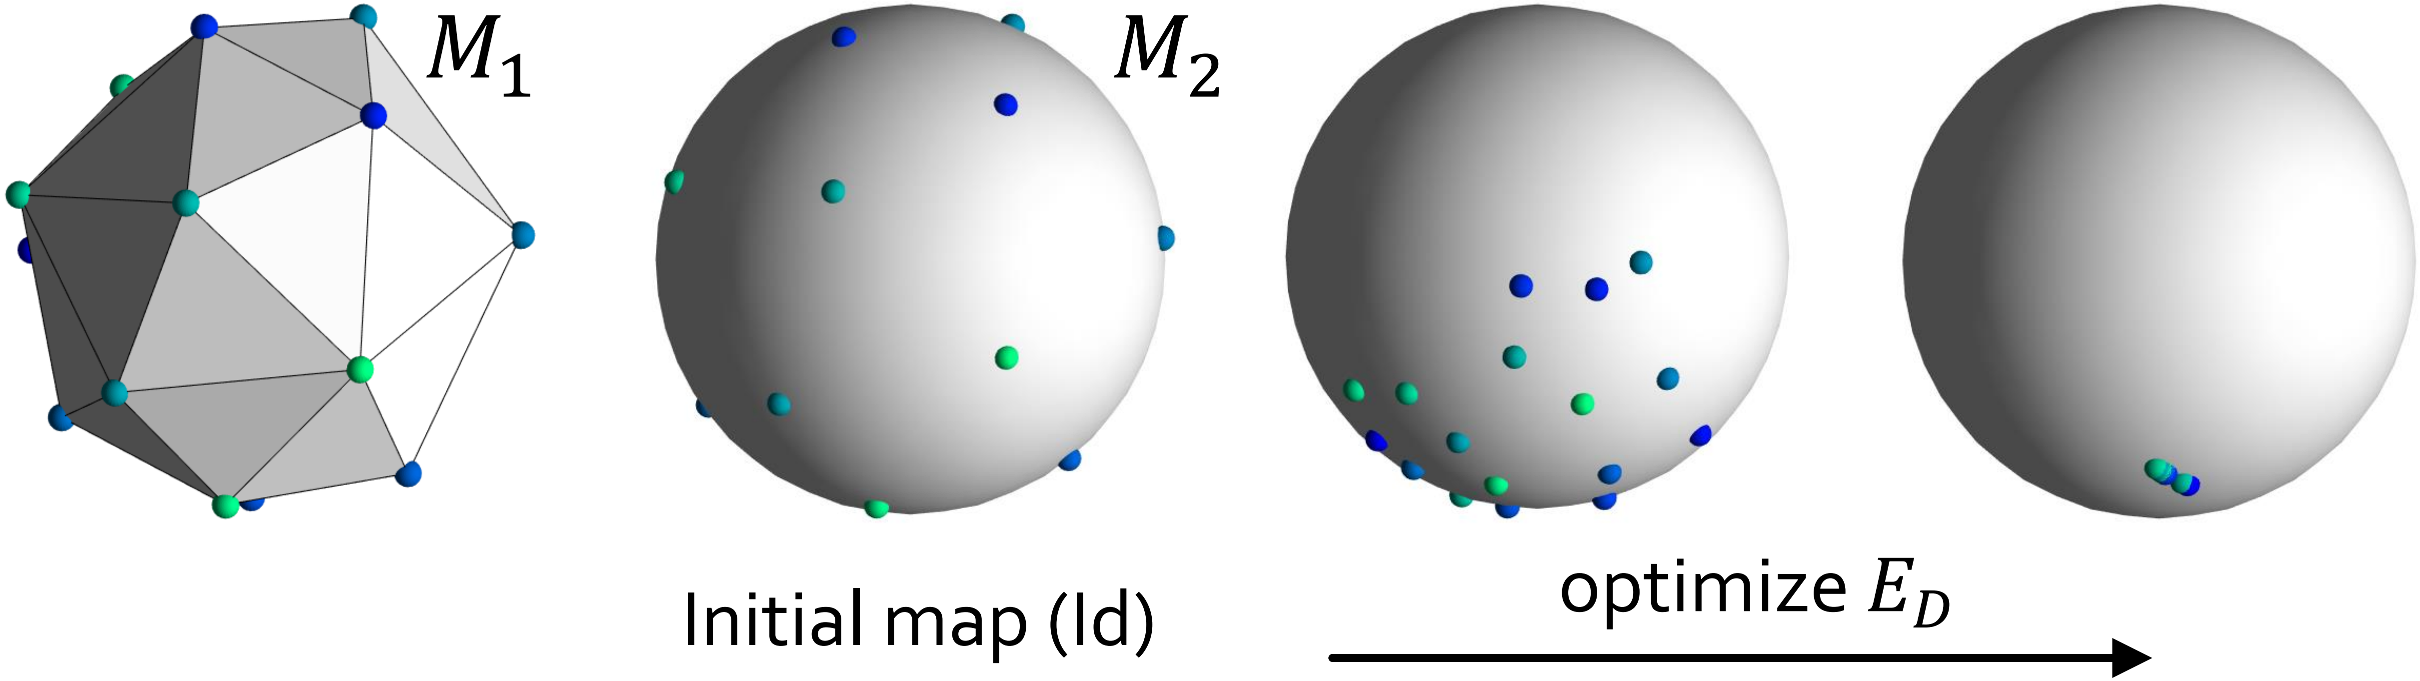

Initial map (Id)

optimize E_D

Reversibility

- We add a **reversibility term** to prevent the map from shrinking

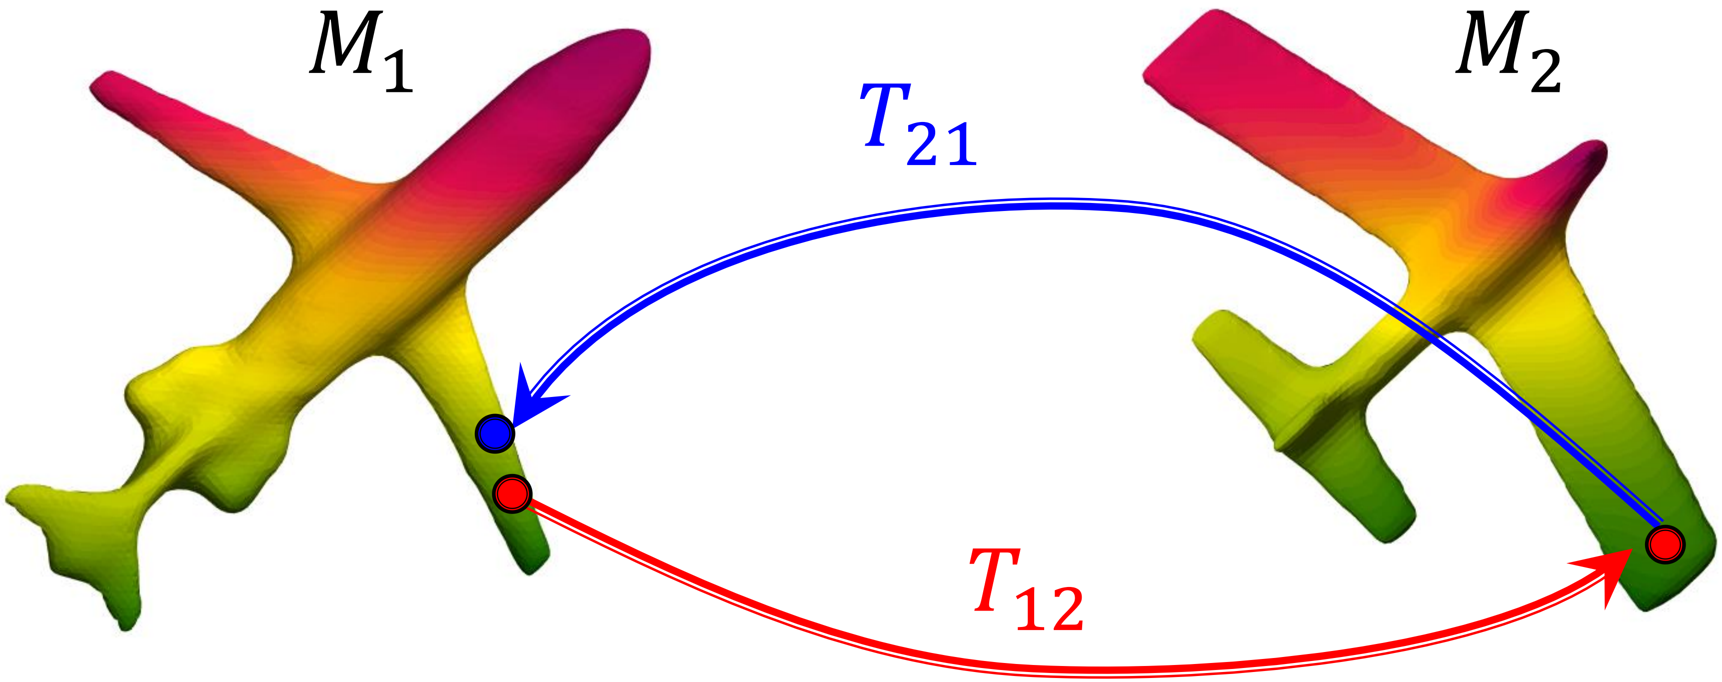

Reversibility

Continuous setting:

$$E_R(T_{12},T_{21}) = \sum_{v\in V_1} d_{M_2}\left(v, T_{21}\left(T_{12}(v)\right)\right) + \sum_{v\in V_2} d_{M_1}\left(v, T_{12}\left(T_{21}(v)\right)\right)$$

The term $E_R(T_{12},T_{21})$ promotes *injectivity* and *surjectivity*

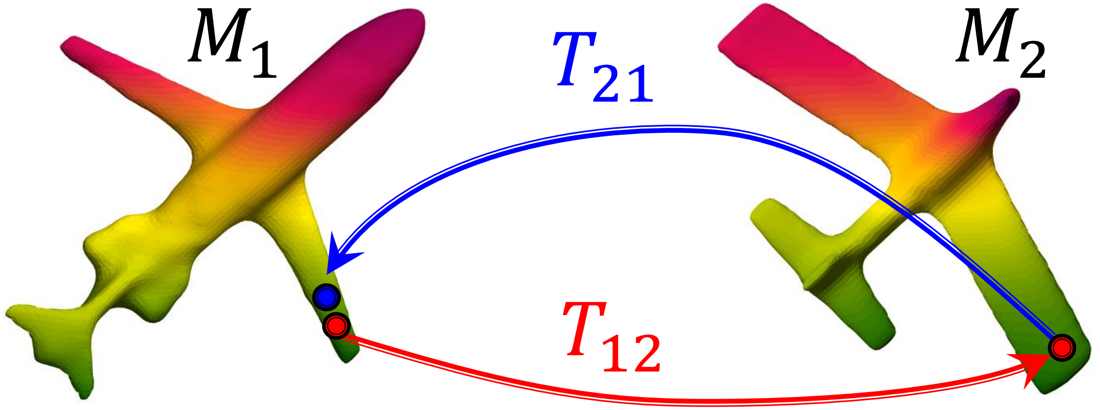

Reversibility

Discrete setting:

$$E_R(P_{12}, P_{21}) = \|P_{21}P_{12}X_2 - X_2\|_{M_2}^2 + \|P_{12}P_{21}X_1 - X_1\|_{M_1}^2$$

Again we use X_1, X_2 the high dimensional embedding of each shape to approximate geodesic distances

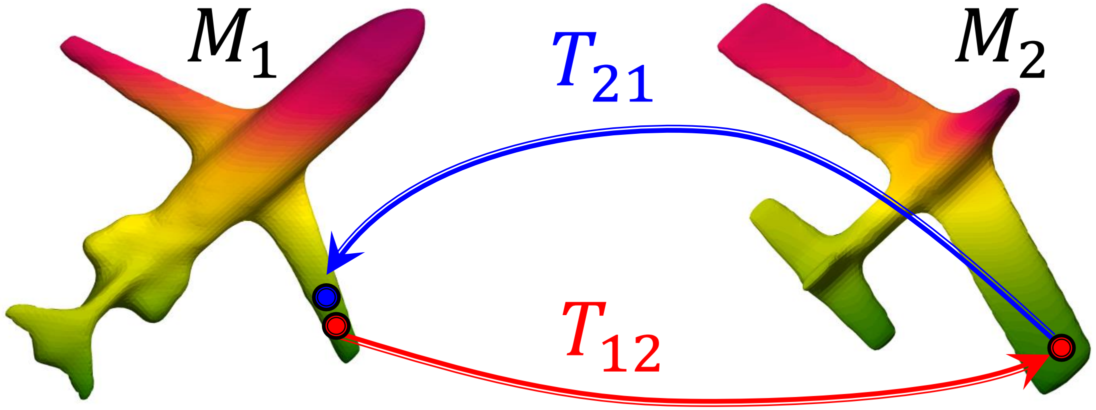

Total Energy

We combine the Dirichlet energy and the reversibility term:

$$E(P_{12}, P_{21}) = \alpha E_D(P_{12}) + \alpha E_D(P_{21}) + (1-\alpha)E_R(P_{12}, P_{21})$$

The parameter α controls the trade off between the terms

Optimization

All the terms are quadratic, but P_{12}, P_{21} are constrained to the feasible set of precise maps

$$P_{12} = \begin{array}{cc} \begin{array}{ccc} j & k & l \end{array} & \\ \left(\begin{array}{ccc} & \vdots & \\ 0.1 & 0.2 & 0.7 \\ & \vdots & \end{array}\right) & \text{row } i \end{array}$$

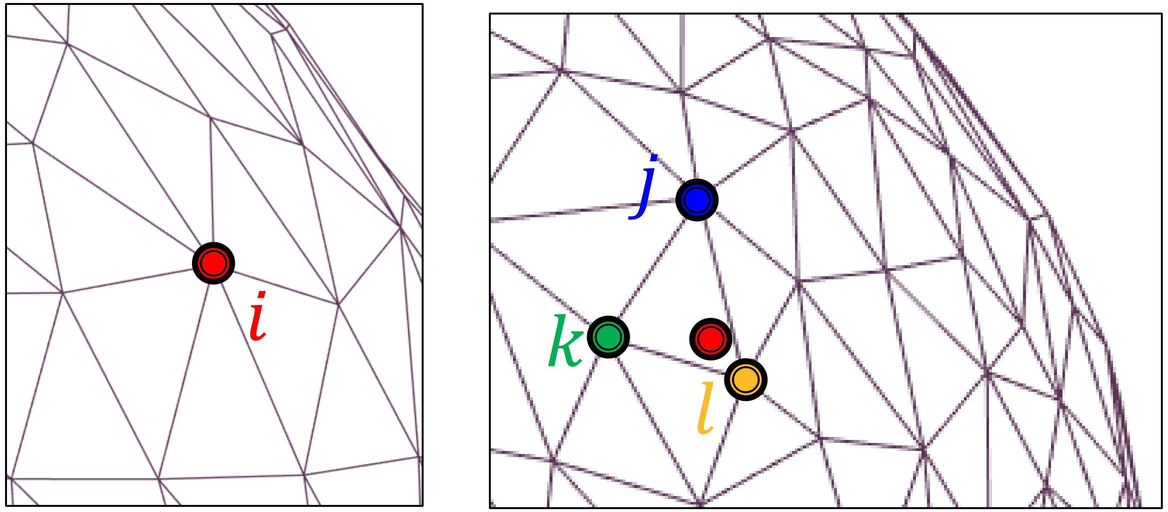

M_1 M_2

$$\boxed{E(P_{12}, P_{21}) = \alpha E_D(P_{12}) + \alpha E_D(P_{21}) + (1-\alpha) E_R(P_{12}, P_{21})}$$

Optimization

We know how to optimize functions of the form:

$$\arg \min_{\mathrm{P}_{12} \in S} \|P_{12}A - B\|^2$$

S is the feasible set of precise maps

Optimization

$$P_{12}^{*} = \arg\min_{P_{12}\in S} \|P_{12}A - B\|_{M_1}^{2}$$

If we constrain to **vertex-to-vertex** maps (subset of feasible set):

P_{12} is a binary stochastic matrix

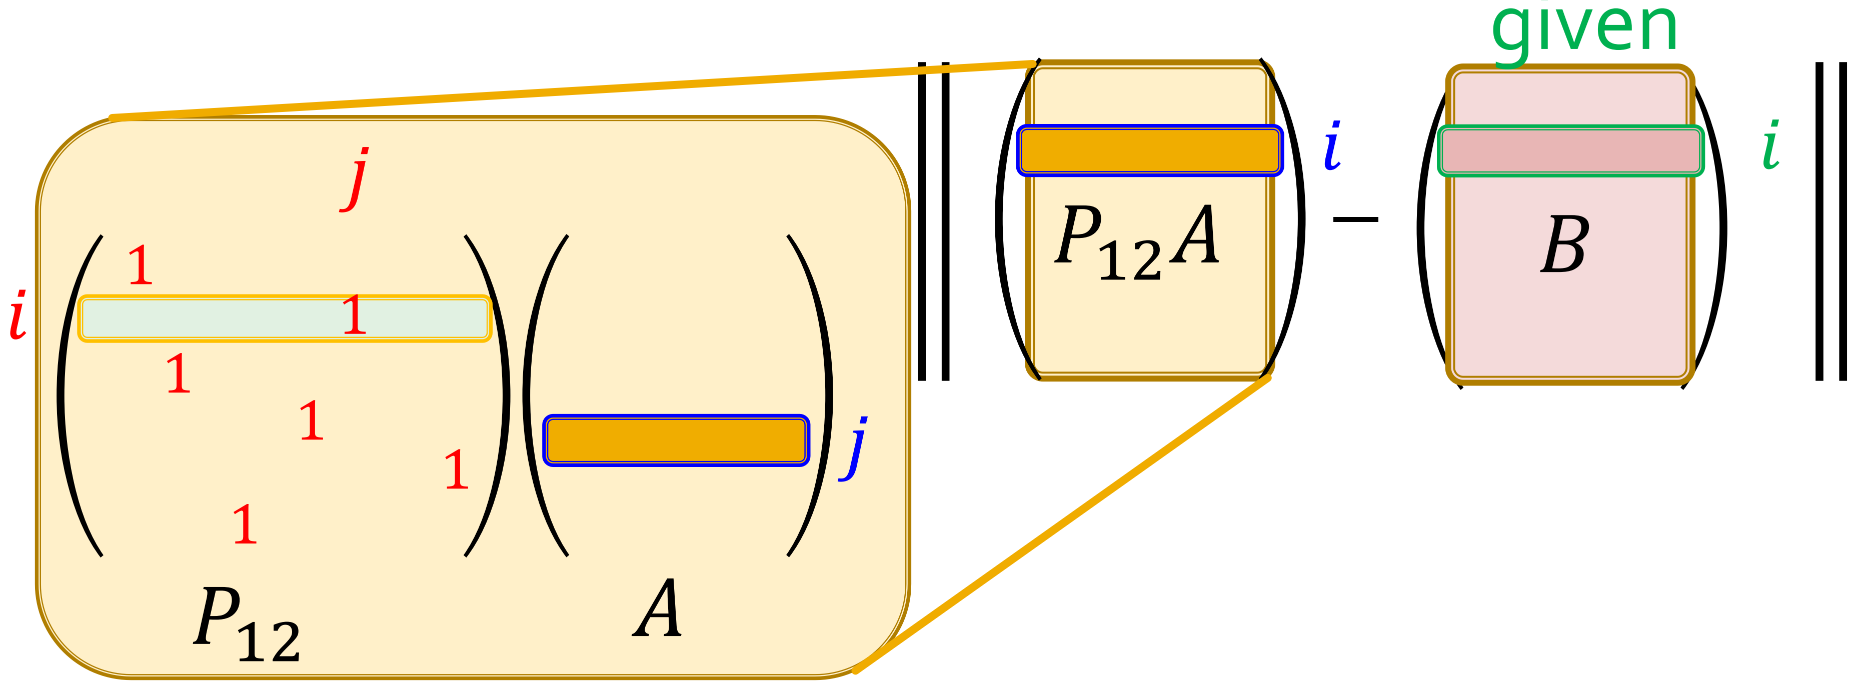

Optimization

$$P_{12}^{*}=\arg\min_{P_{12}\in S}\|P_{12}A-B\|_{M_1}^{2}$$

If P_{12} is any **precise** map:

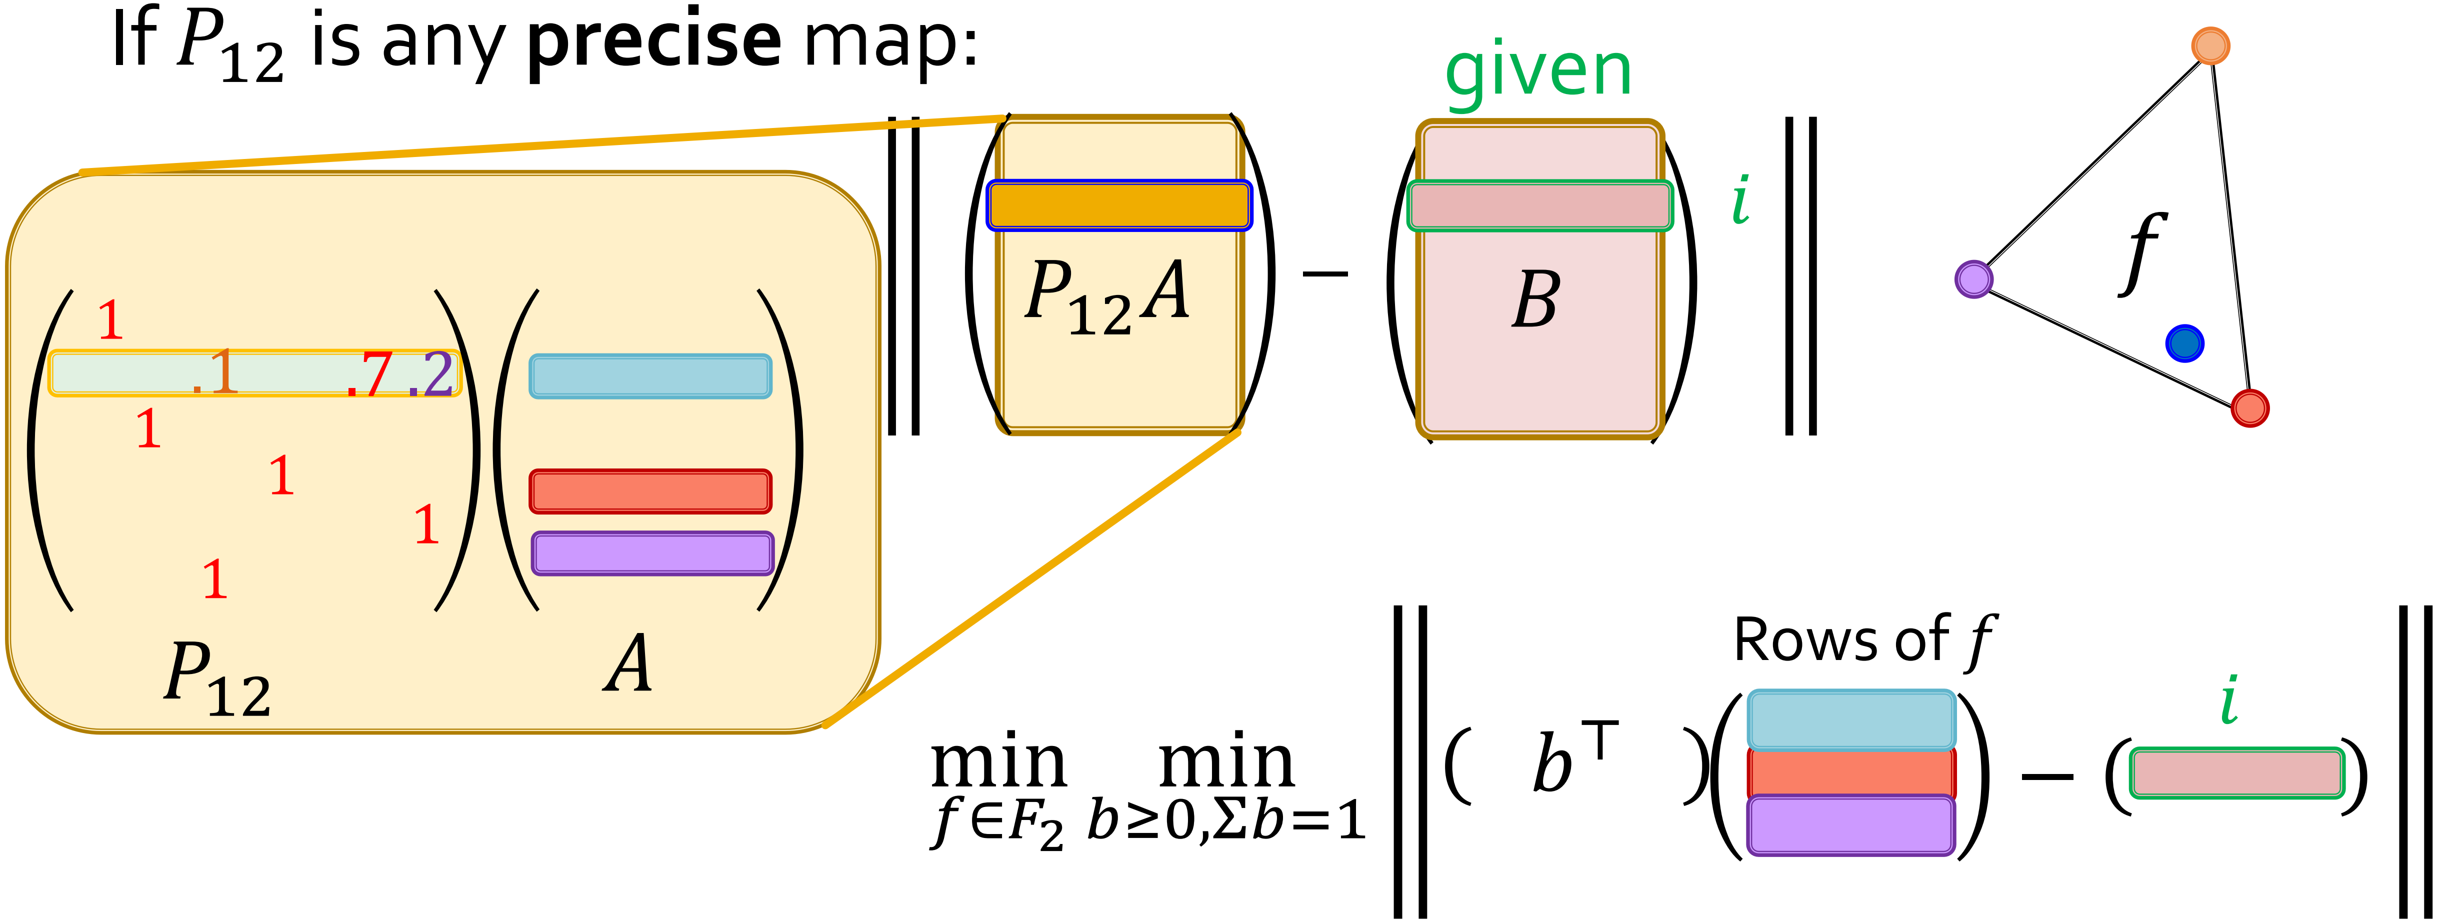

$\min_{f\in F_2}\ \min_{b\geq 0,\Sigma b=1}$

Optimization

$$P_{12}^{*}=\arg\min_{P_{12}\in S}\|P_{12}A-B\|_{M_1}^{2}$$

If P_{12} is any **precise** map:

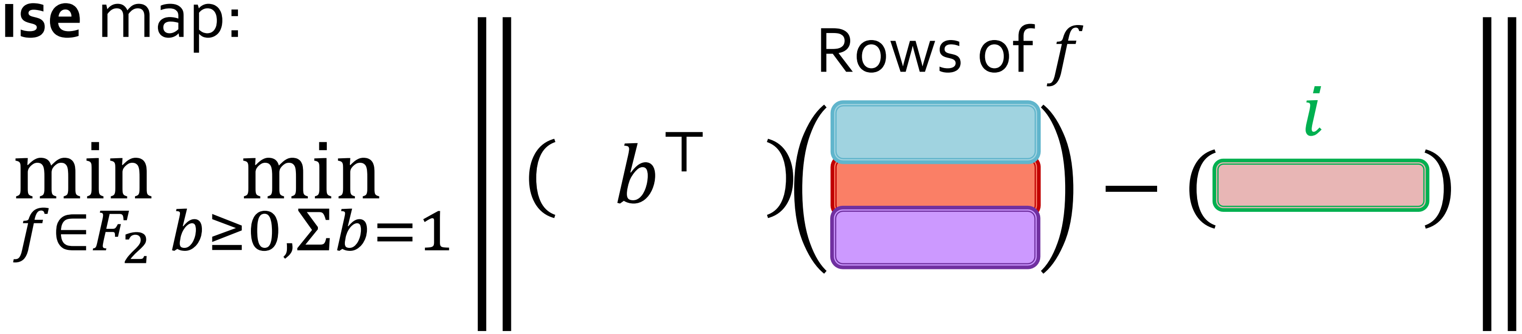

Seems expensive

- Optimize barycentric coordinates by projecting the i_{th} row to a triangle in $\mathbb{R}^{k_2}$ (geometric algorithm)
- Parallelizable!

Optimization

Our energies are not of this form exactly:

$$E_D(P_{12}) = Tr\left((P_{12}X_2)^\top W_1 \, P_{12} \, X_2\right)$$

$$E_R(\mathrm{P}_{12}, P_{21}) = \|P_{21}\underbrace{P_{12}X_2}_{} - X_2\|_{M_2}^2 + \|P_{12}P_{21}X_1 - X_1\|_{M_1}^2$$

We use “half quadratic splitting” such that our energy is of the desired form

Optimization

Introduce new variables

- X_{12} should approximate $P_{12}X_2$, so we add a term $||P_{12}X_2 - X_{12}||^2$
- X_{21} should approximate $P_{21}X_1$, so we add a term $||P_{21}X_1 - X_{21}||^2$

We replace $P_{12}X_2$ by X_{12} wherever it bothers our optimization

Optimization

We rewrite our energies with the new variables:

$$E_D(X_{12}) = Tr(X_{12}^{\top} W_1 \, X_{12})$$

$$E_R(X_{12}, X_{21}, \mathrm{P}_{12}, P_{21}) = \|P_{21} X_{12} - X_2\|_{M_2}^2 + \|P_{12} X_{21} - X_1\|_{M_1}^2$$

$$E_Q(X_{12}, P_{12}) = \|P_{12} X_2 - X_{12}\|_{M_1}^2$$

Optimization

We optimize the energy:

$$\begin{aligned} E(X_{12}, X_{21}, \mathrm{P}_{12}, P_{21}) = \alpha E_D(X_{12}) + \alpha E_D(X_{21}) + &\quad \text{Dirichlet} \\ + (1-\alpha) E_R(X_{12}, X_{21}, \mathrm{P}_{12}, P_{21}) + &\quad \text{Reversibility} \\ + \beta E_Q(X_{12}, P_{12}) + \beta E_Q(X_{21}, P_{21}) &\quad \text{Penalty} \end{aligned}$$

by alternatingly optimizing for each variable

- Optimize P_{12} or P_{21} using projection
- Optimize X_{12} or X_{21} by solving a linear system

Results

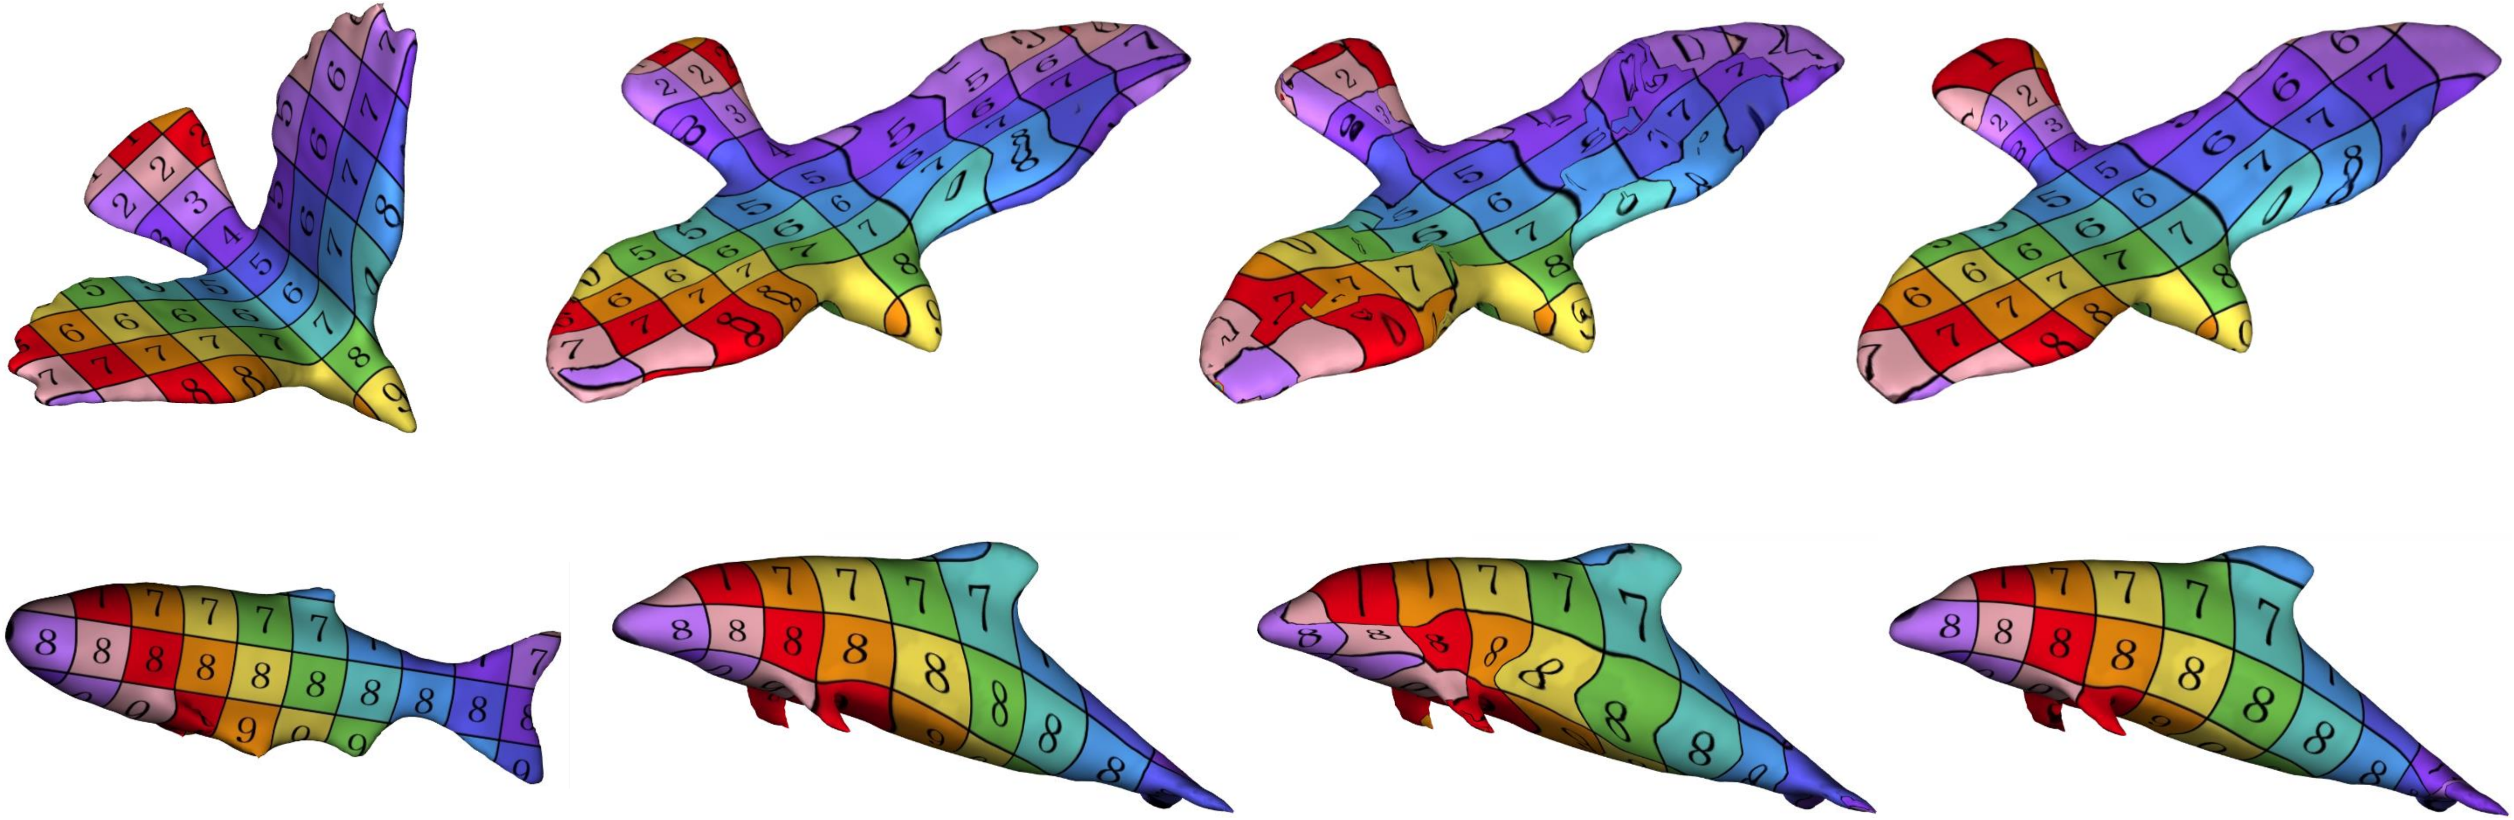

Target

Hyperbolic Orbifolds

Weighted Averages

Ours

Results

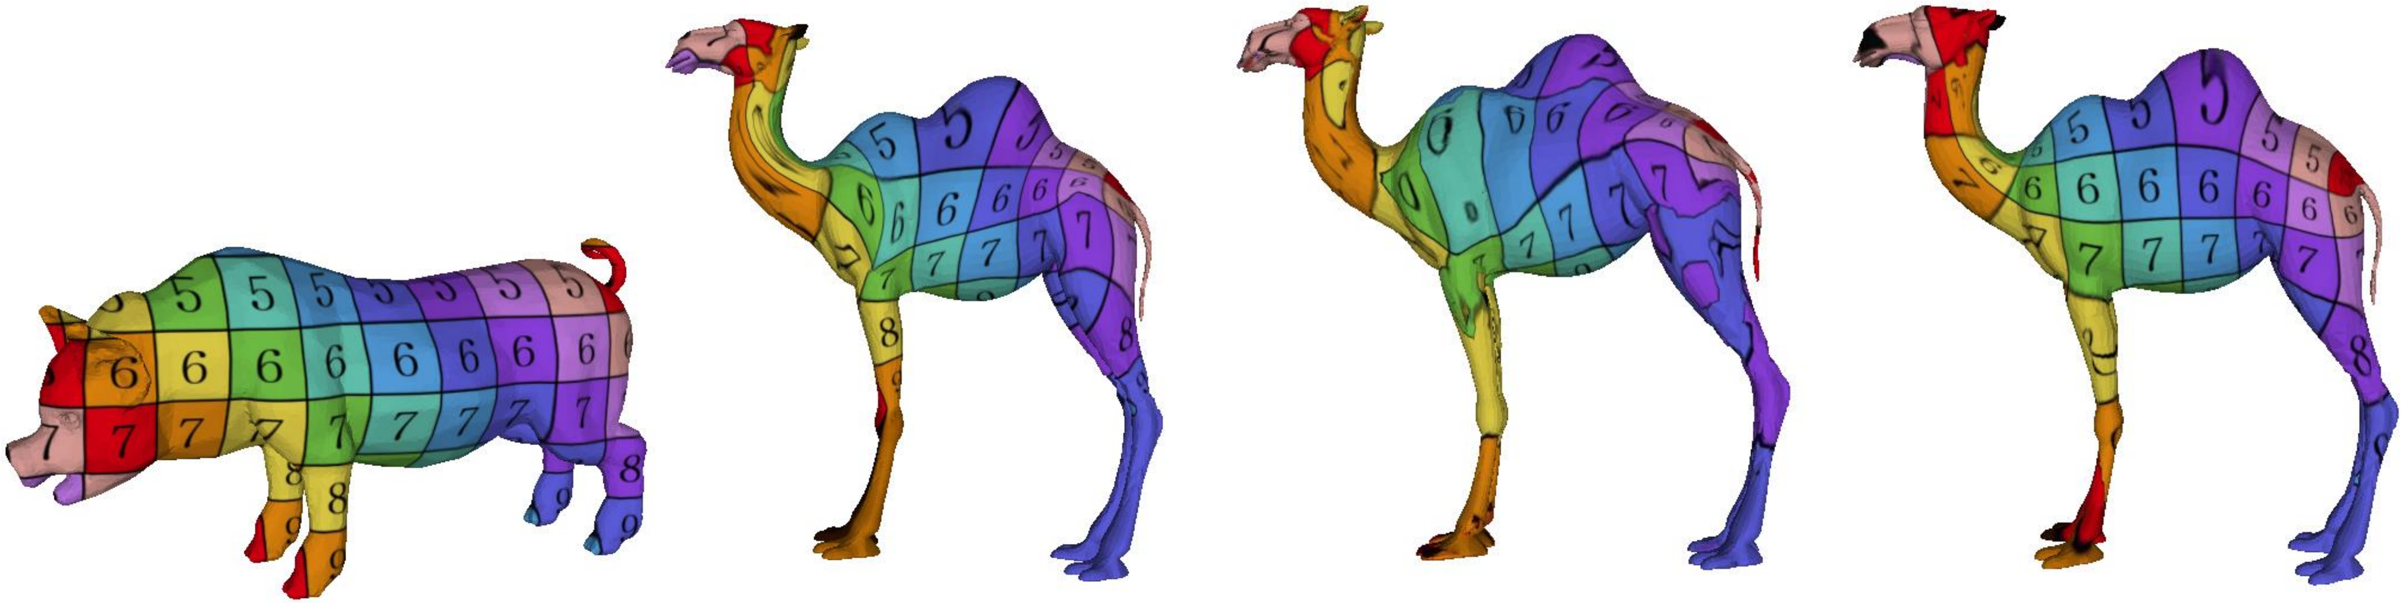

Target

Hyperbolic Orbifolds

Weighted Averages

Ours

Extra: Reversible Harmonic Maps

Justin Solomon

6.838: Shape Analysis

Spring 2021

Extra: Transfer Learning

Justin Solomon

6.838: Shape Analysis

Spring 2021