CS233: Geometric and Topological Data Analysis

Shape Differences, Cycle Consistency.

4 June 2018

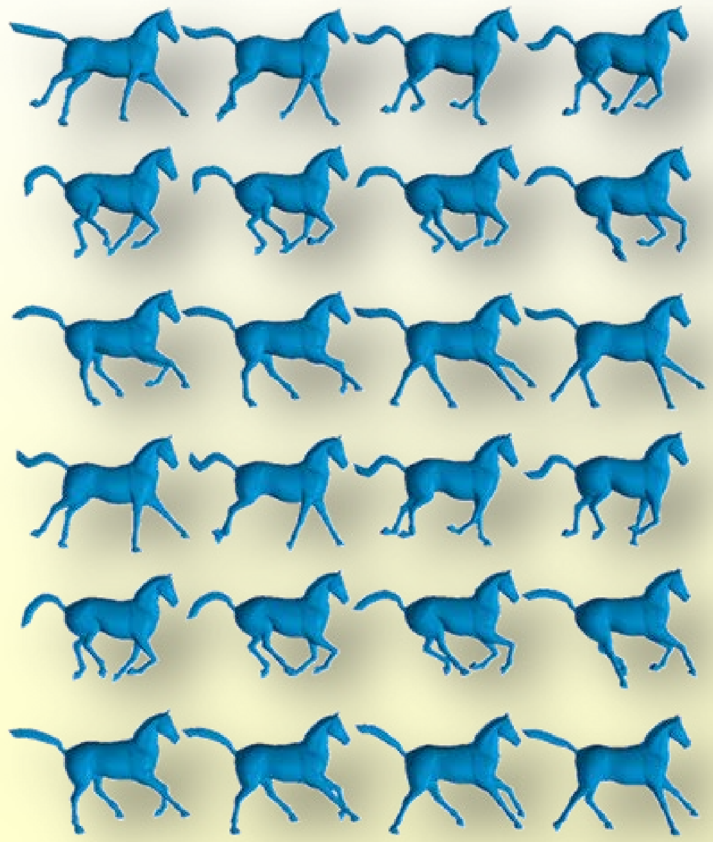

Slides ack: Peter Huang, Raif Rustamov

Last Time:
Function Spaces
and Functional Maps

Functional Maps

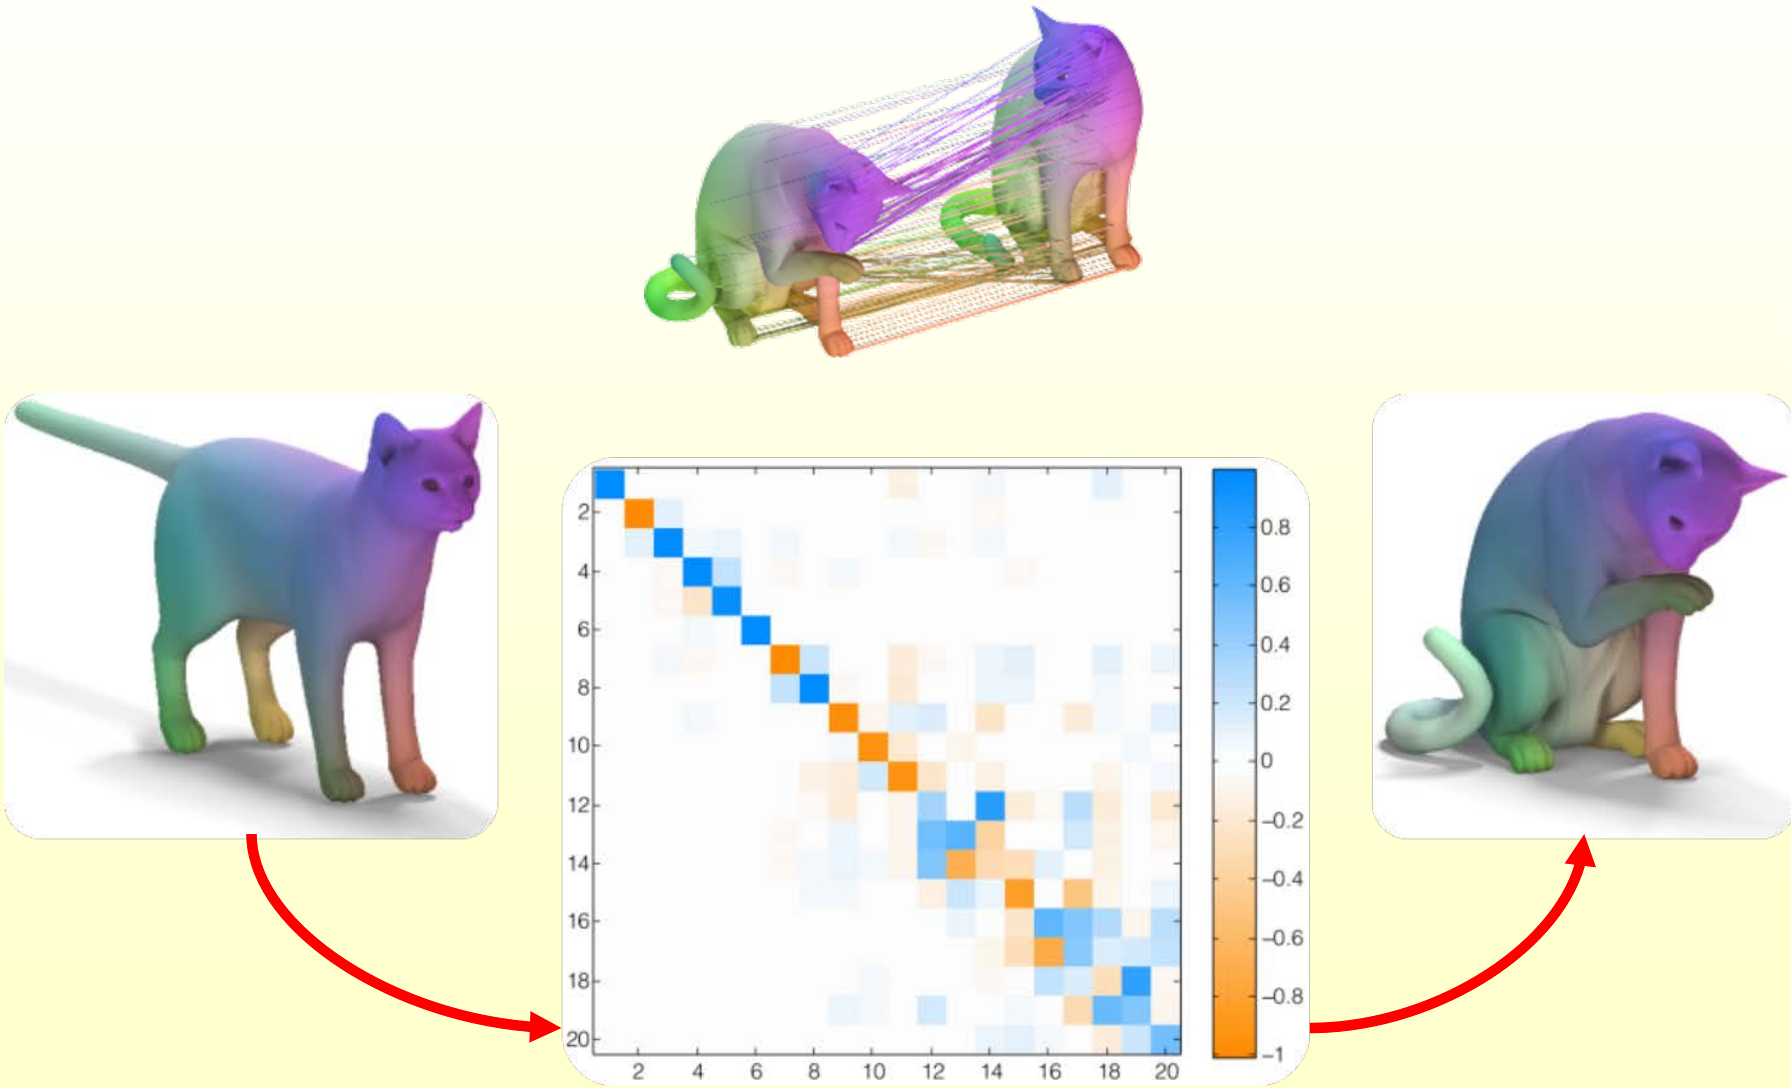

A Contravariant Functor

from cat to lion

Functions on cat are transferred to lion using T_φ

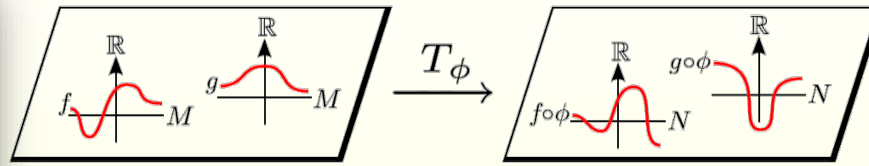

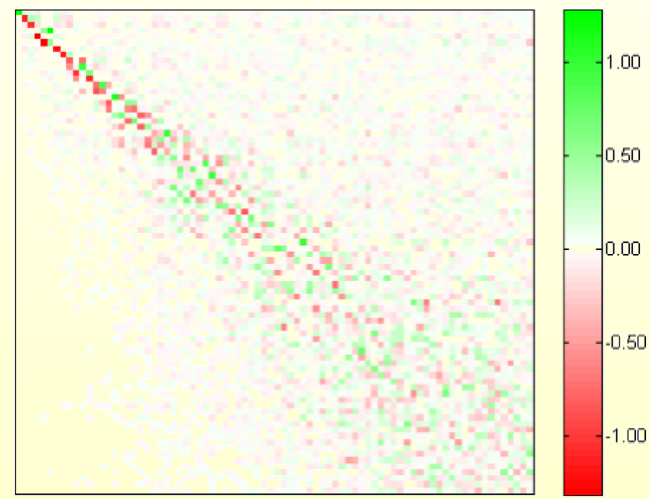

T_φ is a linear operator (matrix)

$$T_\phi : L^2(\text{cat}) \to L^2(\text{lion})$$

Maps as Linear Operators

- An ordinary shape map lifts to a linear operator mapping the function spaces
- With a truncated hierarchical basis, compact representations of functional maps are possible as ordinary matrices
- Map composition becomes ordinary matrix multiplication
- Functional maps can express many-to-many associations, generalizing classical 1-1 maps

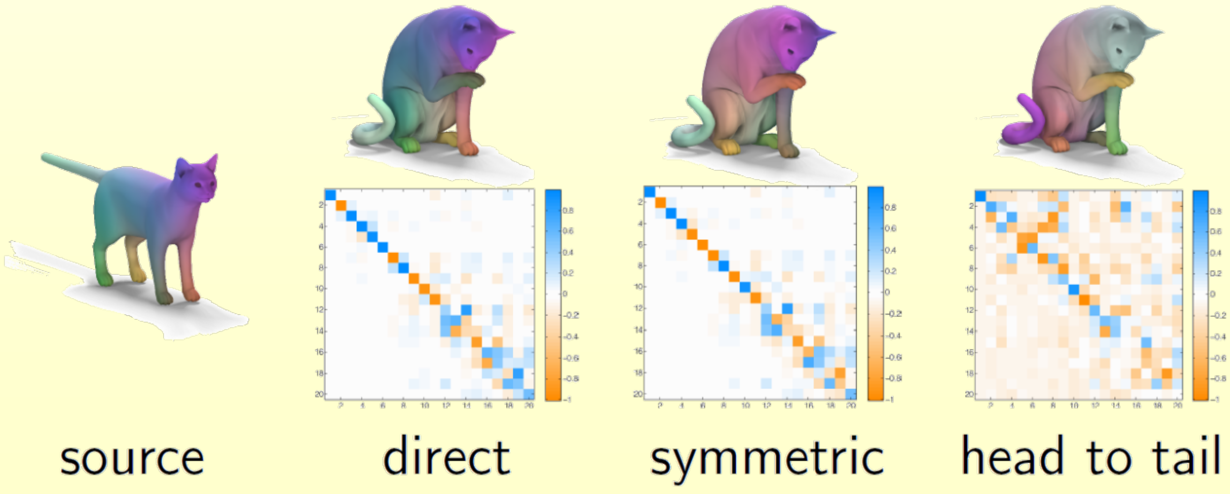

Using truncated Laplace-Beltrami basis

Estimating the Mapping Matrix

Suppose we don't know *C*. However, we expect a pair of functions $f : M \to \mathbb{R}$ and $g : N \to \mathbb{R}$ to correspond. Then, C must satisfy:

$$C\mathbf{a} \approx \mathbf{b}$$

where $f = \sum_i \mathbf{a_i} \phi_i^M, \quad g = \sum_i \mathbf{b}_i \phi_i^N$

Given enough $\{\mathbf{a_i}, \mathbf{b}_i\}$ pairs in correspondence, we can recover C through a linear least squares system.

Commutativity Regularization

In addition, we can phrase an operator commutativity constraint: given two operators $S_1 : \mathcal{F}(M, \mathbb{R}) \to \mathcal{F}(M, \mathbb{R})$ and $S_2 : \mathcal{F}(N, \mathbb{R}) \to \mathcal{F}(N, \mathbb{R})$.

$$\begin{array}{ccc} \mathcal{F}(M,\mathbb{R}) & \xrightarrow{C} & \mathcal{F}(N,\mathbb{R}) \\ {\scriptstyle S_1}\downarrow & & \downarrow{\scriptstyle S_2} \\ \mathcal{F}(M,\mathbb{R}) & \xrightarrow{C} & \mathcal{F}(N,\mathbb{R}) \end{array}$$

Thus: $CS_1 = S_2 C$ or $\|CS_1 - S_2 C\|$ should be minimized

Note: this is a linear constraint on C. S_1 and S_2 could be symmetry operators or e.g. Laplace-Beltrami or Heat operators.

Isometry Regularizer

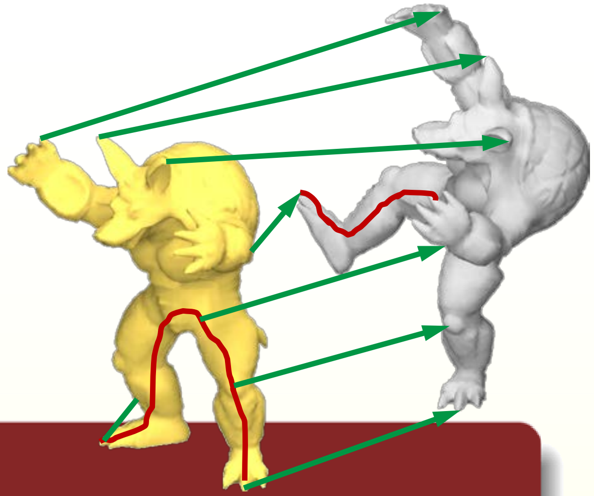

Lemma 1:

The mapping is *isometric*, if and only if the functional map matrix commutes with the Laplacian:

$$C\Delta_1 = \Delta_2 C$$

Differentiate and then transport

Transport and then differentiate

Basic FMaps Pipeline

Given a pair of shapes $\mathcal{M}, \mathcal{N}$:

1. Compute the first *k* (~80-200) eigenfunctions of the Laplace-Beltrami operator. Store them in matrices: $\Phi_{\mathcal{M}}, \Phi_{\mathcal{N}}$

2. Compute descriptor functions (e.g., Heat Kernel Signatures) on $\mathcal{M}, \mathcal{N}$. Express them in $\mathbf{A}, \mathbf{B}$, as columns of : $\Phi_{\mathcal{M}}, \Phi_{\mathcal{N}}$

3. Solve $C_{\text{opt}} = \arg\min_{C} \|C\mathbf{A} - \mathbf{B}\|^2 + \|C\Delta_{\mathcal{M}} - \Delta_{\mathcal{N}} C\|^2$

   $\Delta_{\mathcal{M}}, \Delta_{\mathcal{N}}$: diagonal matrices of eigenvalues of LB operator

4. Convert the functional map C_{opt} to a point to point map *T*.

Map Visualization

Can show that singular vectors of the functional representation C of T identify most distorted regions in a multi-scale way.

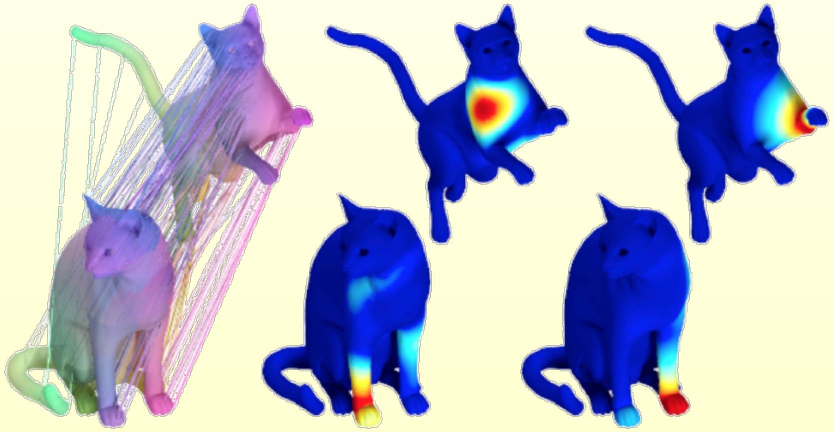

Partial Functional Maps

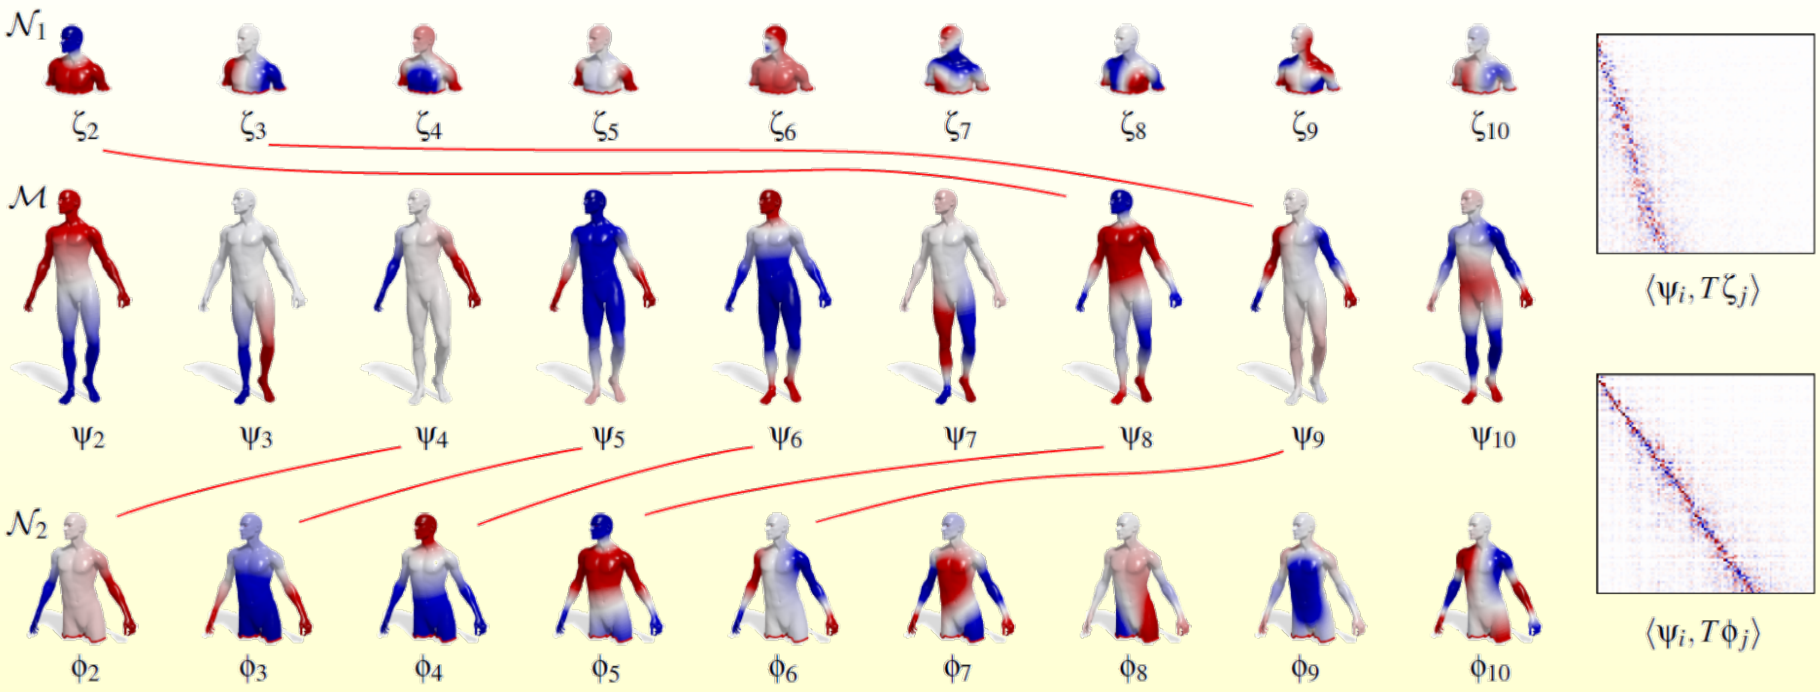

Shape Differences

App: Shape Differences

vs.

vs.

Measures of Distortion

- Shape correspondences are often computed by minimizing some measure of distortion

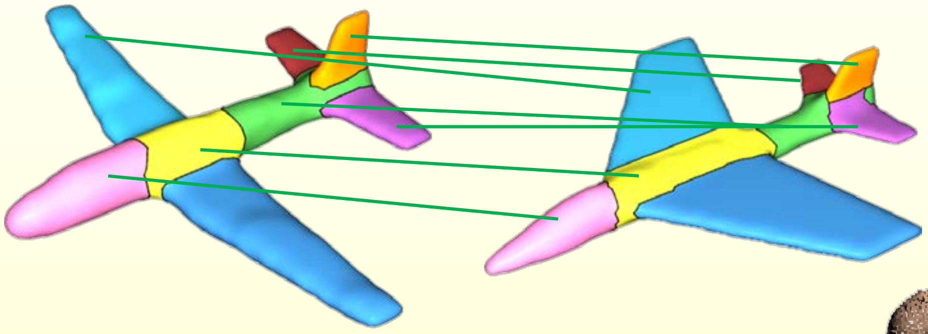

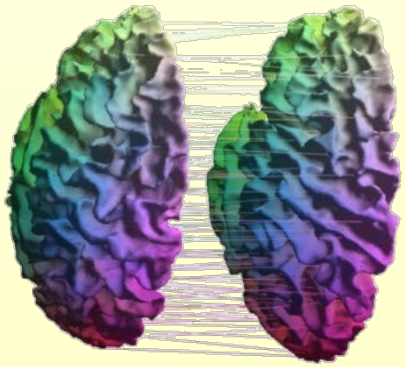

Yet the distortion estimates are forgotten as soon as the correspondences are computed

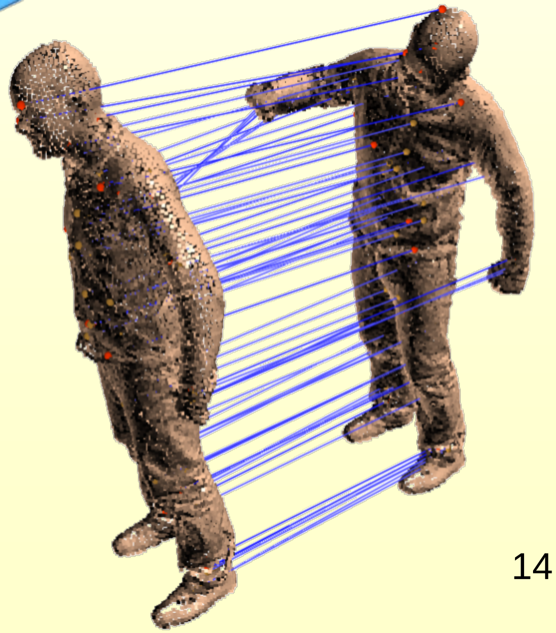

Understanding Intrinsic Distortions

- Where and how are shapes different, locally and globally, irrespective of their embedding

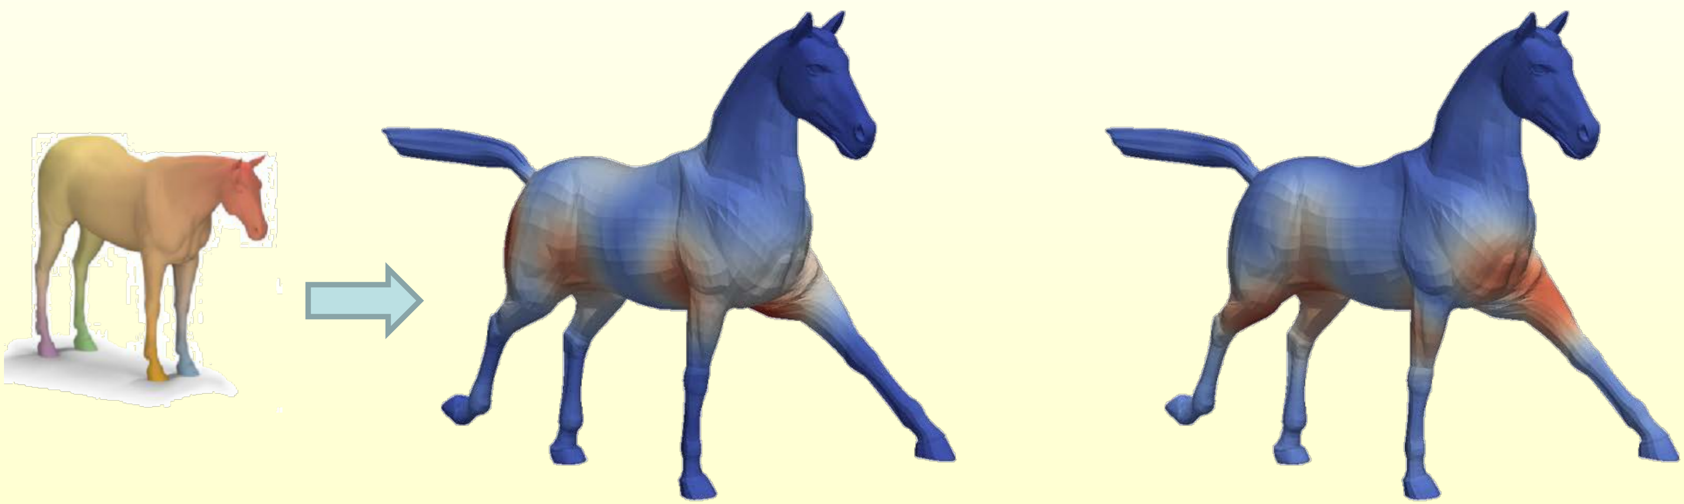

Area distortion

Conformal distortion

We saw how we can localize distortions using functional maps

Classical Approach to Relating Shapes

To measure distortions induced by a map, we track how inner products of vectors change after transporting

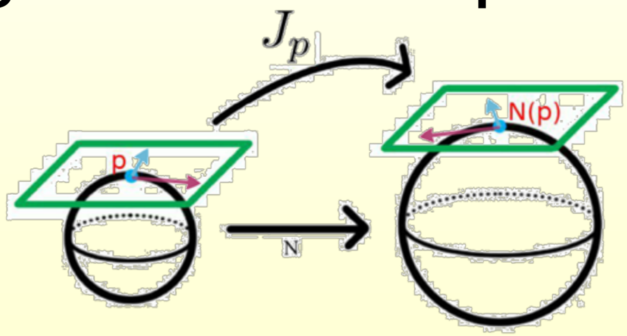

Riemann

Challenges:

- point-wise information only, hard to aggregate
- noisy

A Functional View of Distortions

To measure distortions induced by a map, track how inner products of vectors change after transporting.

To measure distortions induced by a map, track how inner products of functions change after transporting.

Riemann

Lots of Inner Products

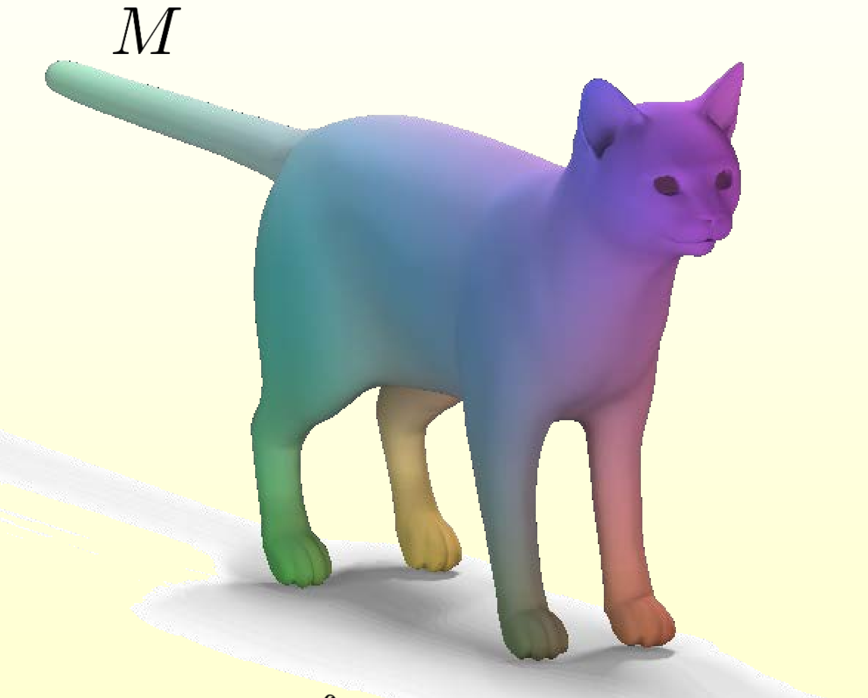

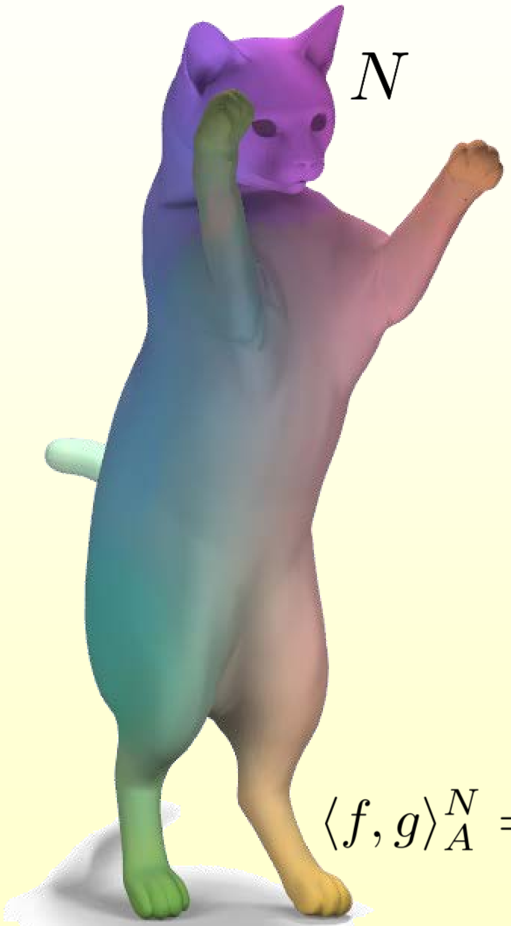

$$\langle f, g \rangle_A^M = \int_M f g \, dA$$

$$\langle f, g \rangle_C^M = \int_M \nabla f \cdot \nabla g \, dA$$

area inner product
conformal inner product

$$\langle f, g \rangle_A^N = \int_N f g \, dA$$

$$\langle f, g \rangle_C^N = \int_N \nabla f \cdot \nabla g \, dA$$

The Art of Measurement

- A metric is defined by a functional inner product

$$h^M(f,g) = \int_M f(x)g(x)d\mu(x)$$

- So we can compare *M* and *N* by comparing

$$h^M(f,g)$$

$$h^N(F(f),F(g))$$

The functional map *F* transports these functions to *N*, where we repeat this measurement with the inner product $h^N(F(f),F(g))$

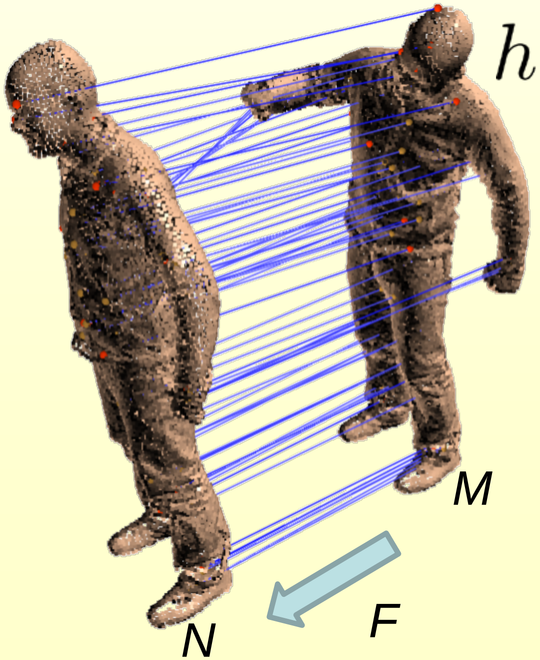

Riemann

Inner Products of Functions

$$\langle f, g \rangle = \left\langle \sum_i f_i h_i(x), \sum_j g_j h_j(x) \right\rangle$$

$$= \sum_{ij} f_i g_j \langle h_i(x), h_j(x) \rangle$$

$$= f^\top A g$$

Area weights matrix

Theorem

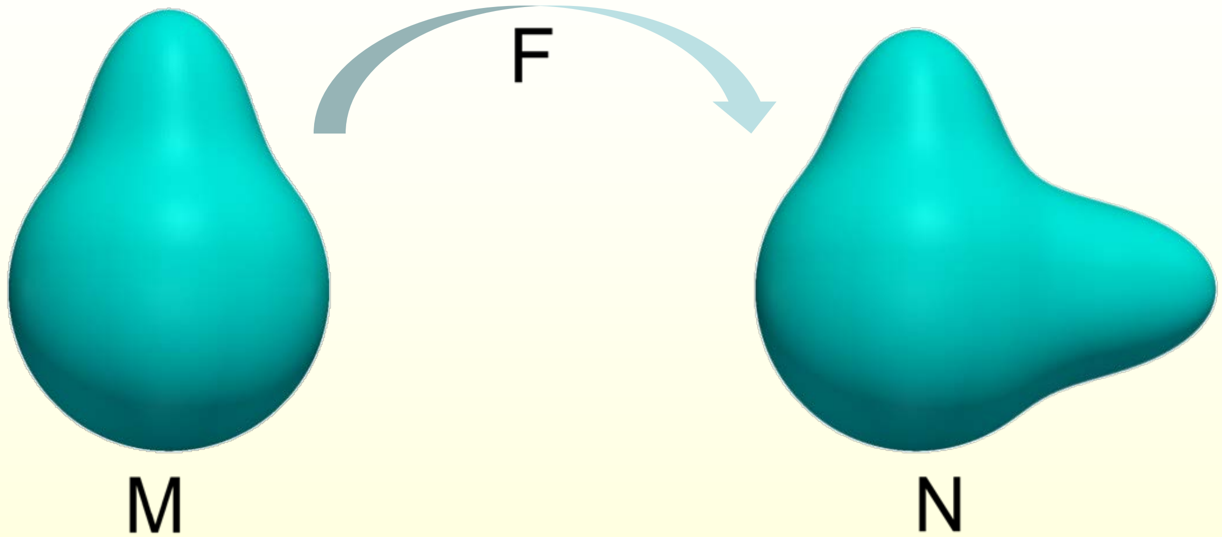

If *F* comes from bijective p2p map *T: N → M*, then

$$\int_N F(f)F(g)\, d\nu = \int_M fg\, d\mu \quad \forall f, g$$ ⟺ *T* is point-wise area preserving

Measurement Discrepancies

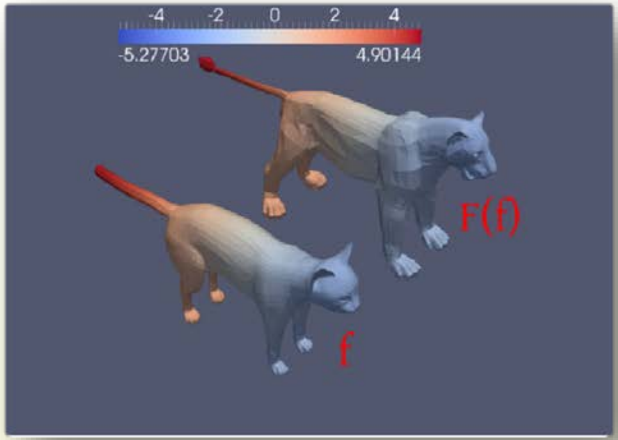

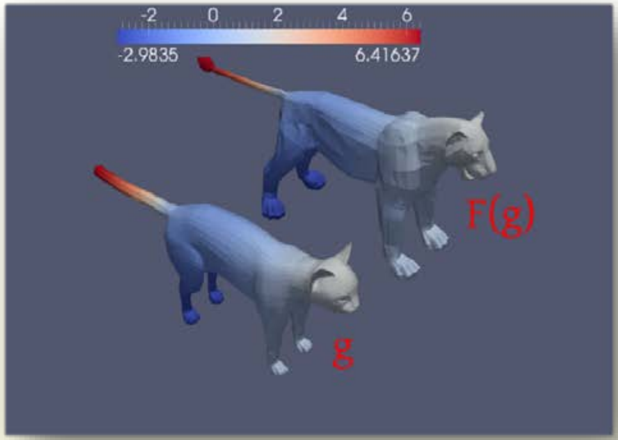

$$\underbrace{\int_{lion} F(f)F(g)\, d\mu_l}_{\text{after}} \neq \underbrace{\int_{cat} f g\, d\mu_c}_{\text{before}}$$

However, both can be considered as inner products on the cat

The Universal Compensator

Comptes Rendus Hebdomadaires des
Séances de l'Académie des Sciences de Paris

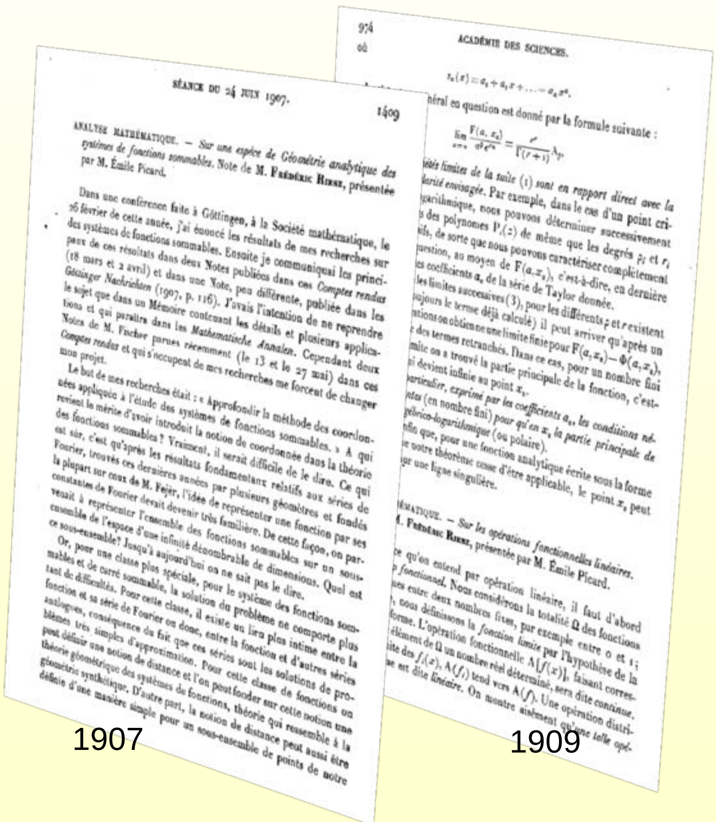

1907 1909

Riesz Representation Theorem

There exists a linear operator

$$V : L^2(cat) \to L^2(cat)$$

such that

$$\langle f, g \rangle_{\text{after}} = \langle f, V(g) \rangle_{\text{before}}$$

Frigyes Riesz

Riesz Representation Theorem

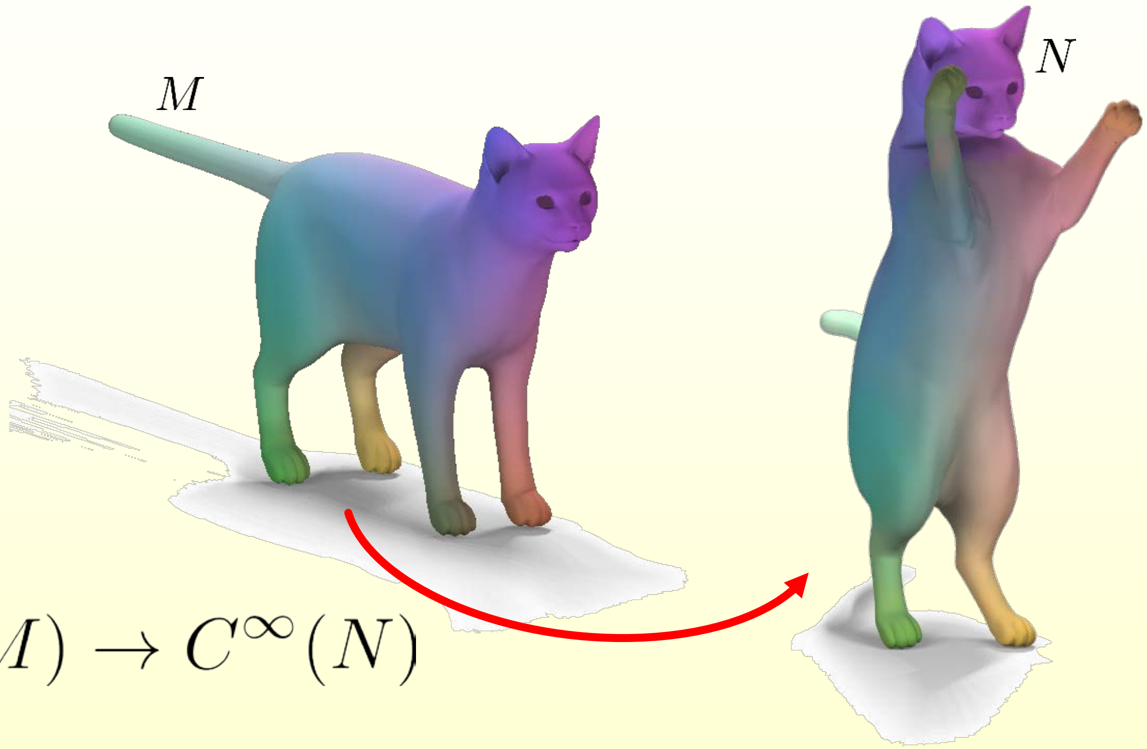

$$F : C^\infty(M) \to C^\infty(N)$$

$$\exists V \text{ s.t.}$$
$$\langle F(f), F(g) \rangle^N = \langle f, V(g) \rangle^M \;\forall f, g$$

Sanity Check

$$
\begin{aligned}
\langle f, g \rangle &\approx f^\top A g \\
\langle F(f), F(g) \rangle^N &\approx [F(f)]^\top A_N [F(g)] \\
&= f^\top \cdot F^\top A_N F \cdot g \\
&= f^\top \cdot (A_M A_M^{-1}) F^\top A_N F \cdot g \\
&= f^\top \cdot A_M (A_M^{-1} F^\top A_N F \cdot g) \\
&\approx \langle f, (A_M^{-1} F^\top A_N F) g \rangle
\end{aligned}
$$

Area-Based Shape Difference:

$$V = A_M^{-1} F^\top A_N F$$

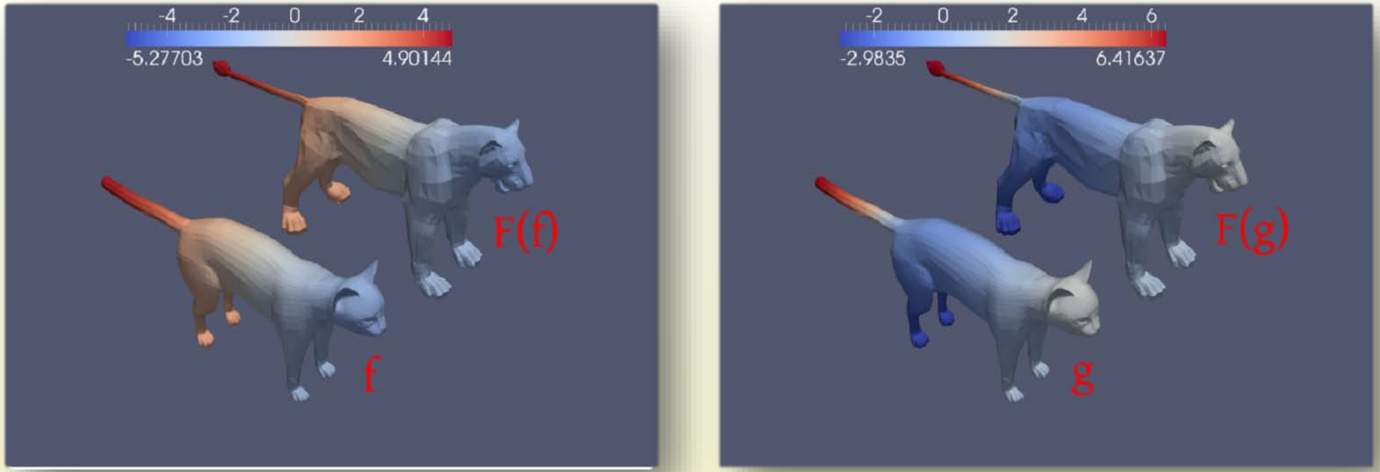

$$\int_{lion} F(f)F(g) \neq \int_{cat} fg$$

$$\downarrow$$

$$\int_{lion} F(f)F(g) = \int_{cat} f V(g)$$

A Small Example of V

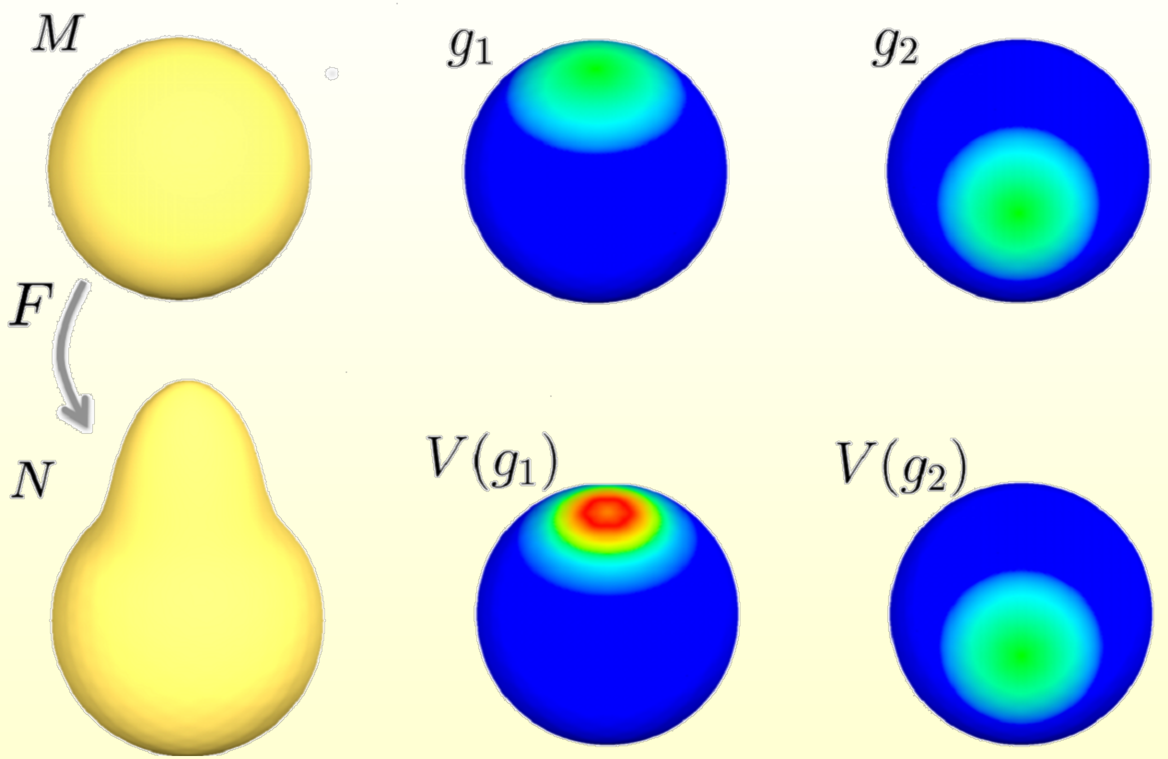

$$\int_N F(f)F(g) = \int_M fV(g)$$

Conformal Shape Difference: *R*

Consider a different inner-product of functions ...

get information about conformal distortion

$$\int_N \nabla F(f) \nabla F(g) = \int_M \nabla f \nabla R(g)$$

The choice of inner product should be driven by the application at hand.

Conformal Shape Difference: R

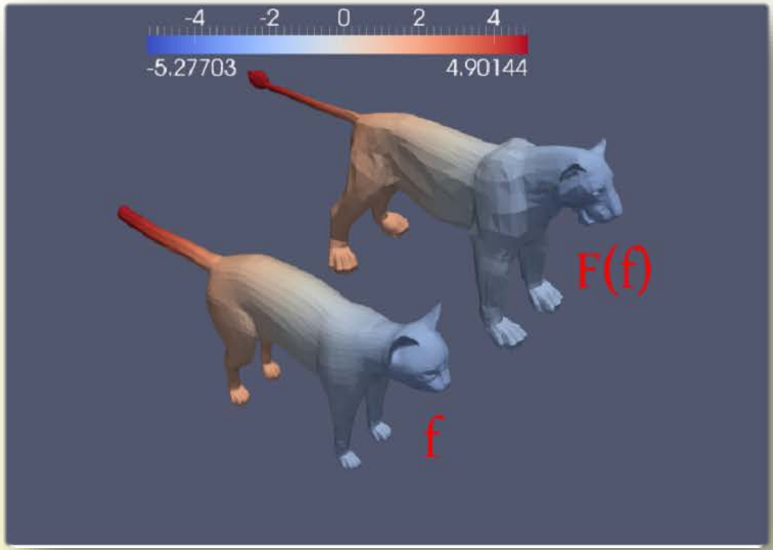

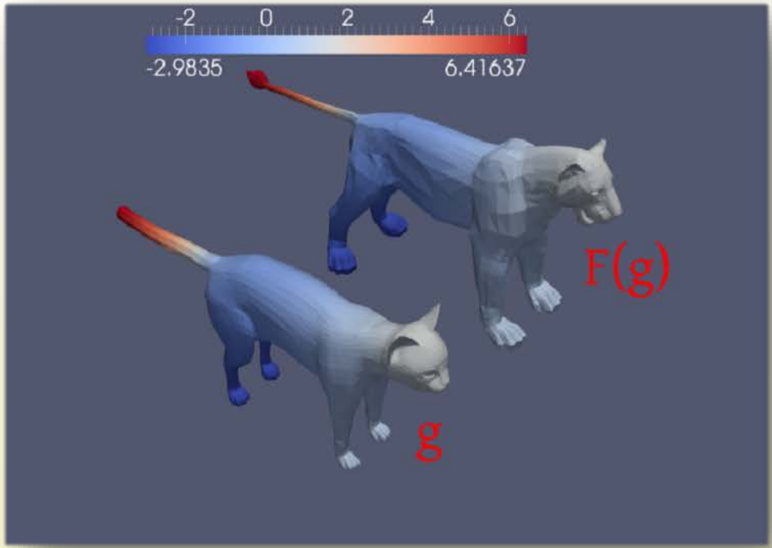

$$\int_{lion} \nabla F(f) \cdot \nabla F(g) \neq \int_{cat} \nabla f \cdot \nabla g \quad \forall f, g$$

$$\int_{lion} \nabla F(f) \cdot \nabla F(g) = \int_{cat} \nabla f \cdot \nabla R(g)$$

Input: Functional Map *F*

from cat to lion

Functions on cat are transferred to lion using *F*

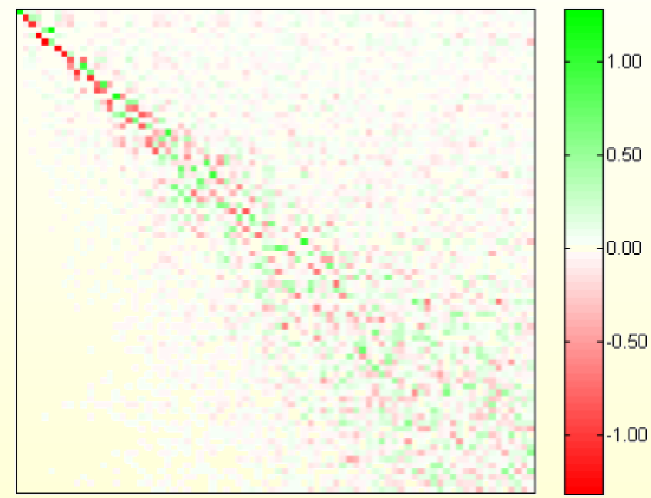

***F* is a linear operator (matrix)**

$$F : L^2(cat) \to L^2(lion)$$

Output: A Shape Difference

V – area-based shape difference

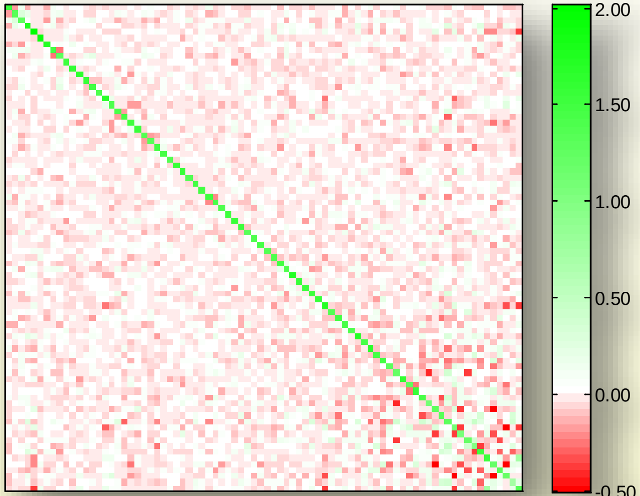

linear operator (matrix)

$V : L^2(cat) \to L^2(cat)$

R – conformal shape difference

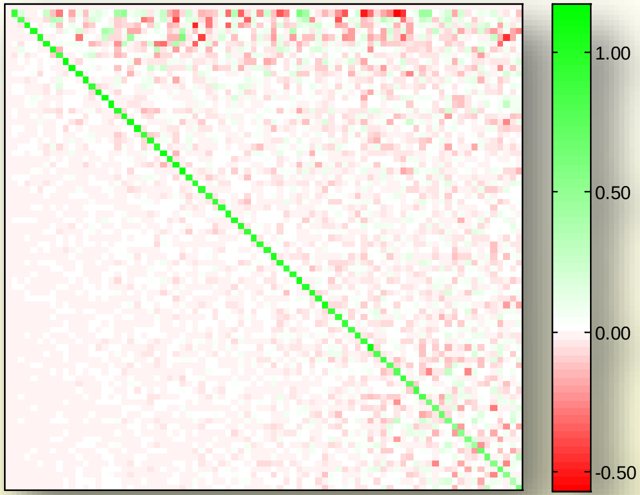

linear operator (matrix)

$R : L^2(cat) \to L^2(cat)$

Note: the are both operators on the map domain (the cat)

Summary

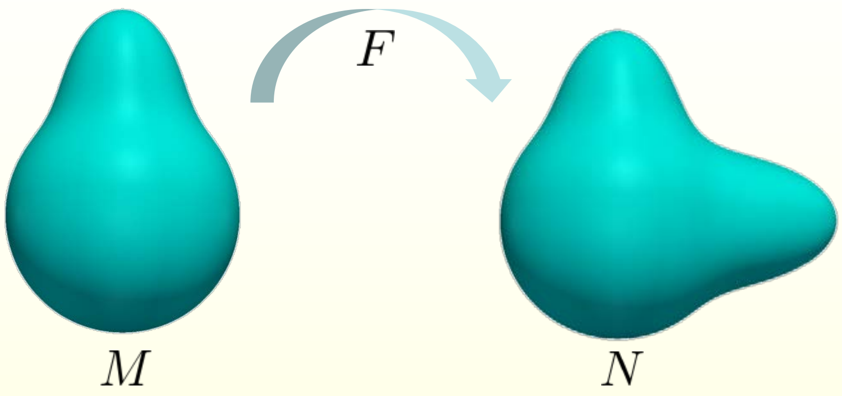

Assume *F* comes from a bijective p2p map *T: N→M*

V is identity ↔ map is point-wise area-preserving

R is identity ↔ map is conformal

V and R are both identity ↔ map is an isometry

Informativeness

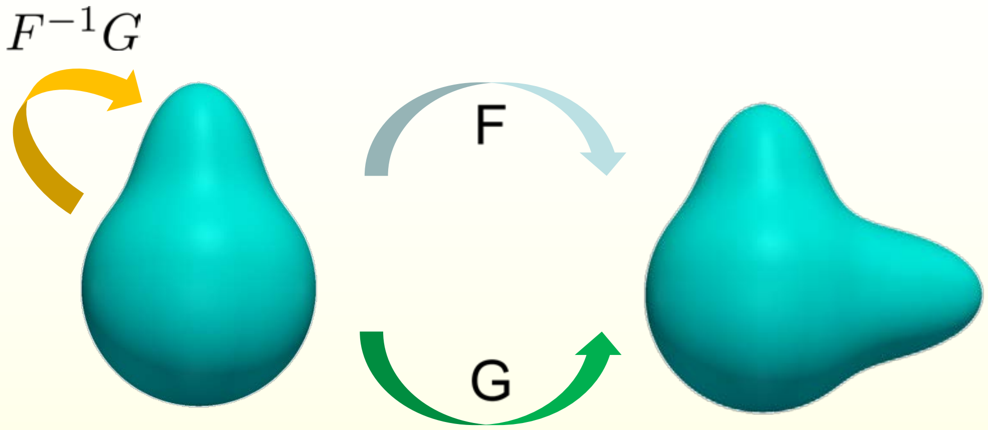

Assume *F & G* come from bijective p2p maps

$V_F = V_G$ ⟺ $F^{-1}G$ is area preserving

$R_F = R_G$ ⟺ $F^{-1}G$ is conformal

- Shape differences encode the map up to an area preserving or conformal self-map.
- Shape differences are fully informative up to the given notion of distortion.

Distortion Localization

- p2p map *f*: *M* → *N*, f2f map *T*: $L^2(N) \to L^2(N)$
- Region Ω 2 *M*, χ_Ω characteristic function of Ω
- Then

$$\frac{\alpha(\Omega)}{\alpha(T^{-1}(\Omega))} = \frac{\int \chi_\Omega \cdot \chi_\Omega}{\int \chi_\Omega \cdot V(\chi_\Omega)}$$

Where α denotes area.

Singular vectors of *f*, eigenvectors of *V*, localize regions of extreme distortion

Shape Differences in Collections

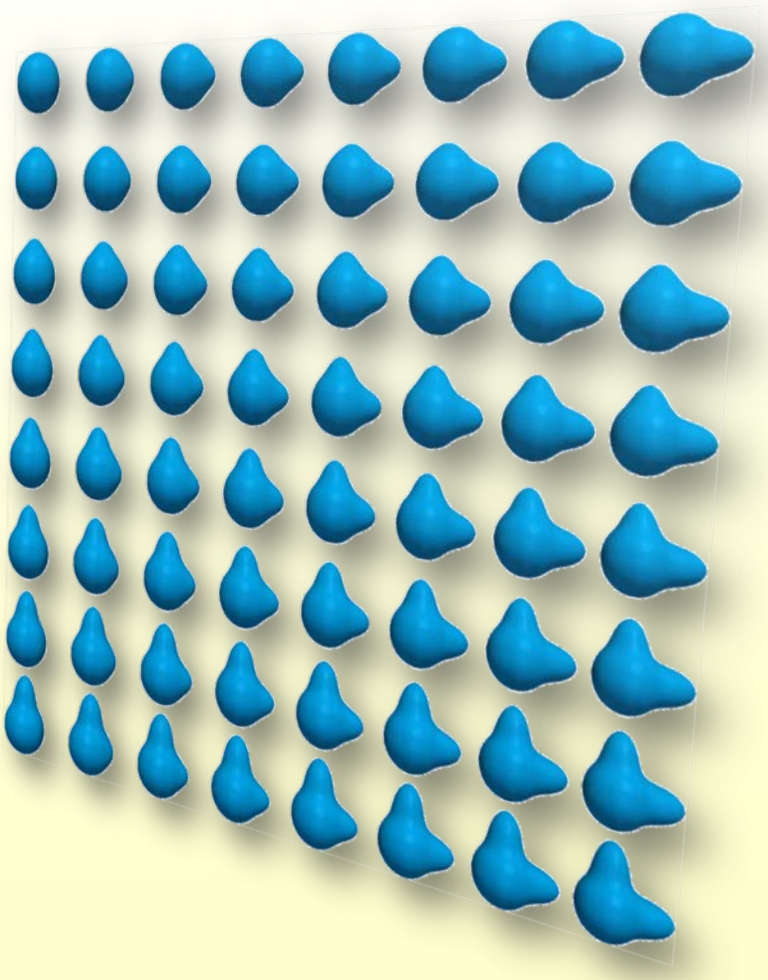

Comparing Differences I

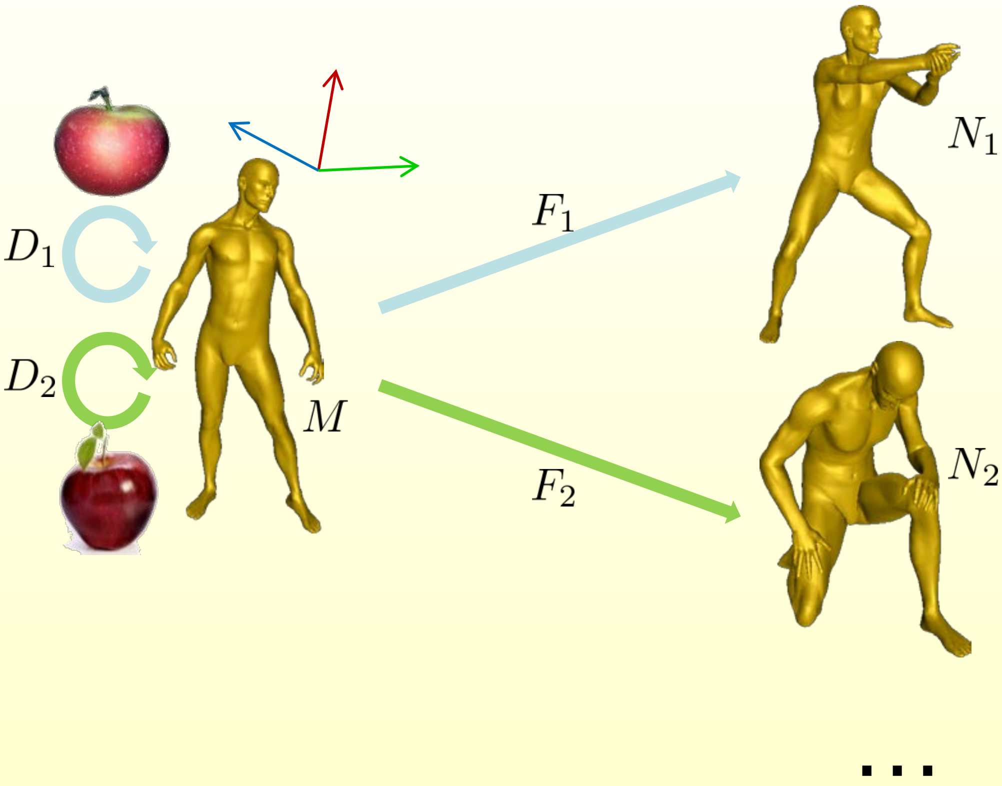

Intrinsic Shape Space

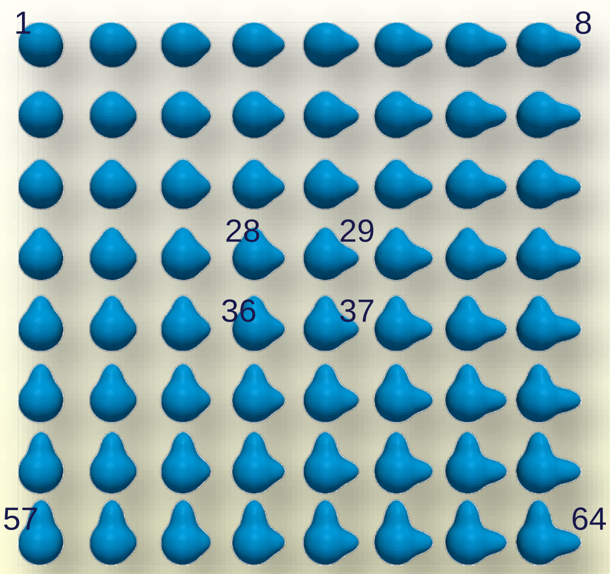

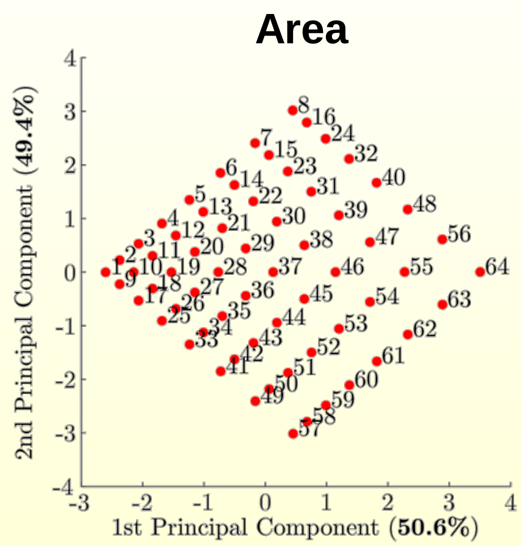

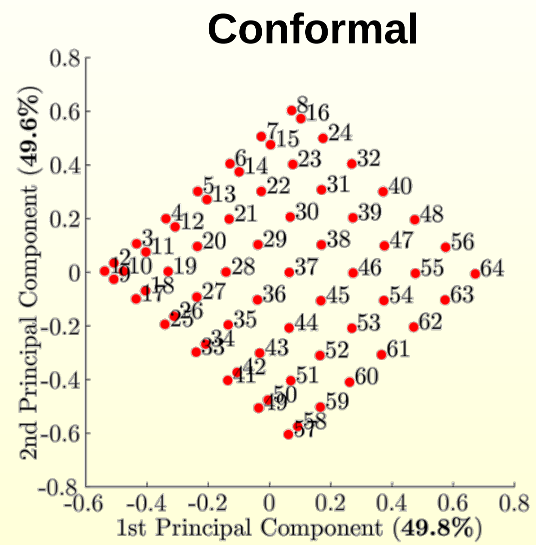

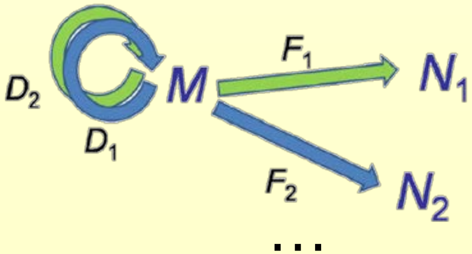

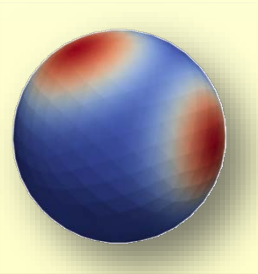

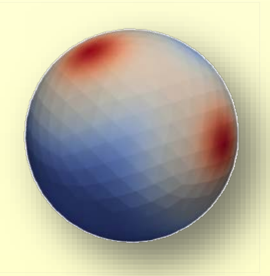

Intrinsic Shape Space

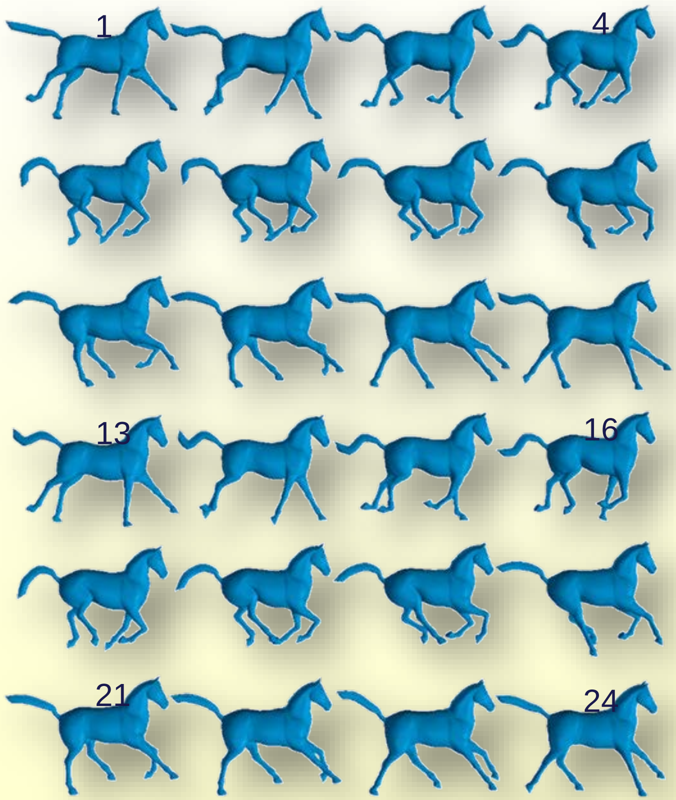

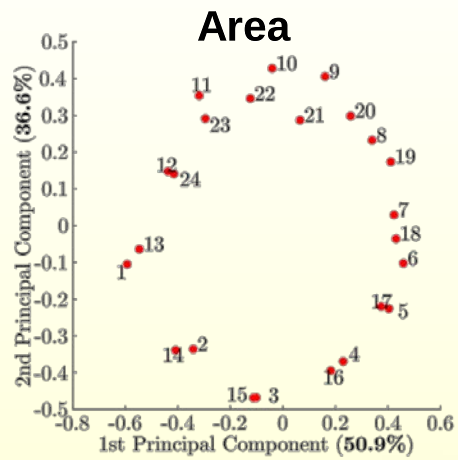

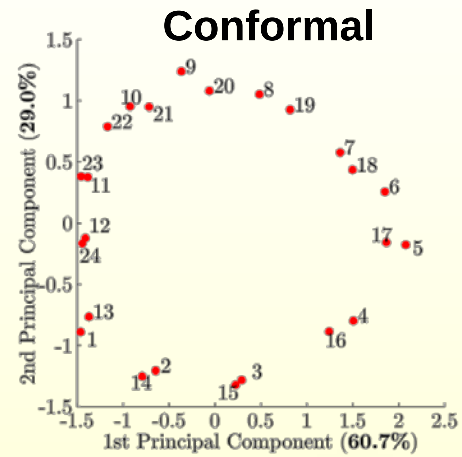

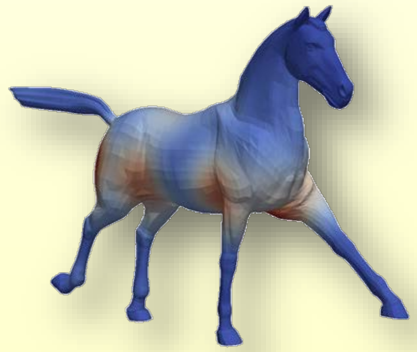

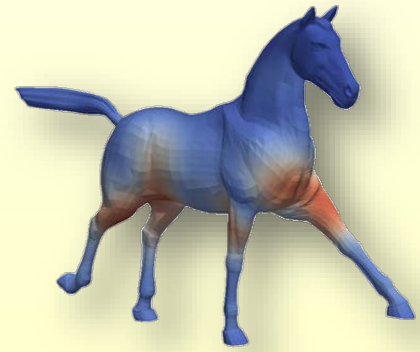

Medical Dataset

16 autistic, 11 control subjects

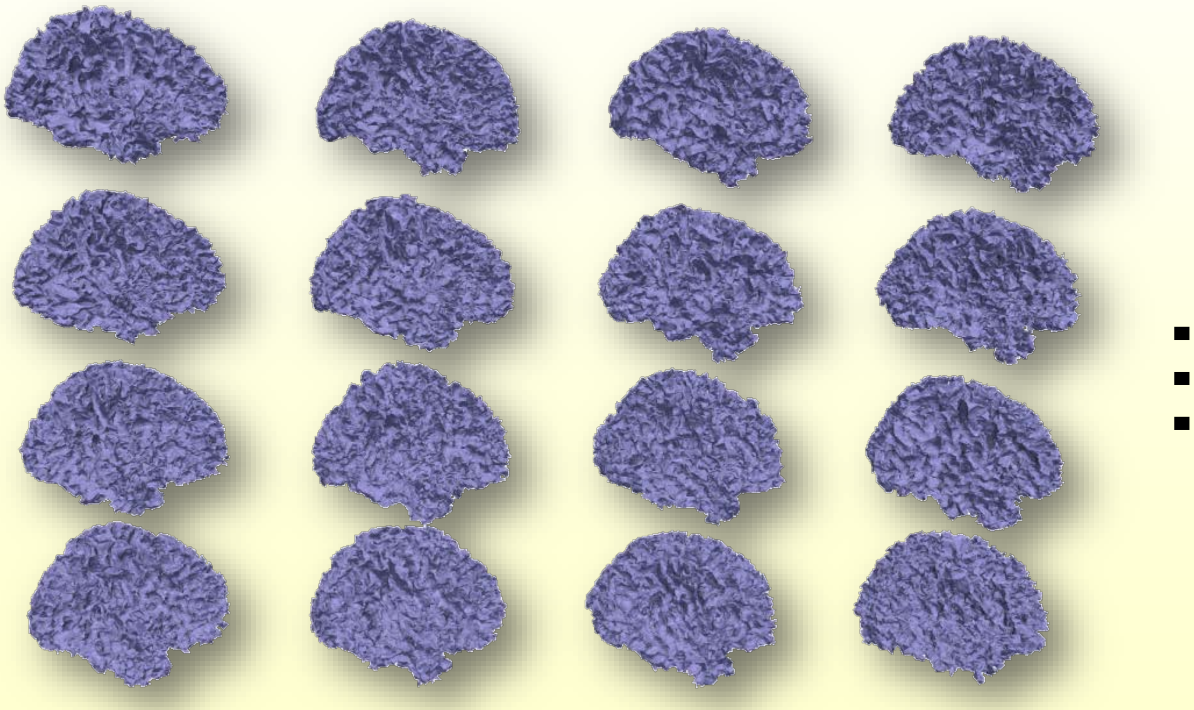

inner cortical surfaces

Medical Dataset

Area

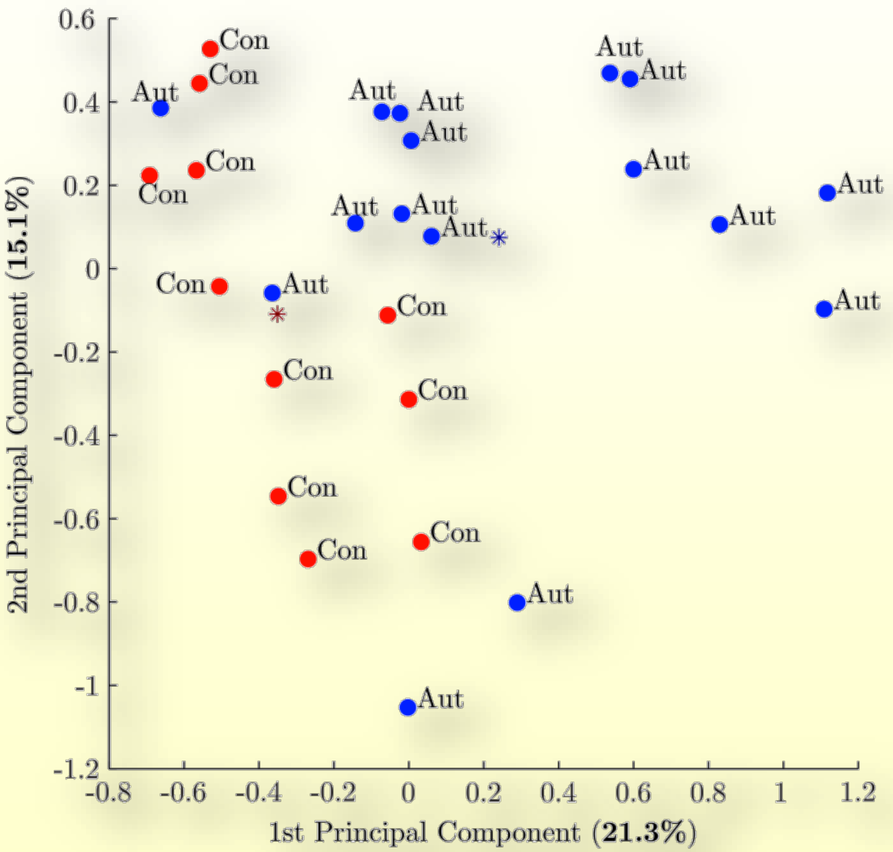

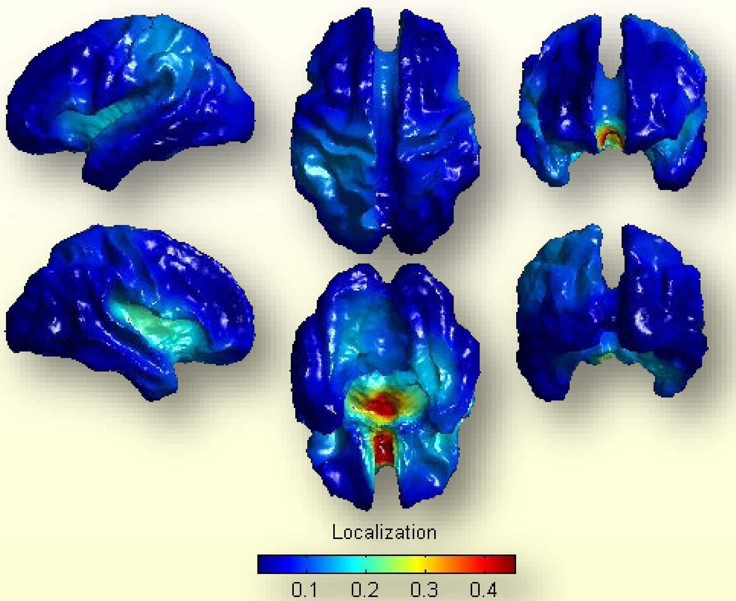

Means are stat. sign. different (p = 0.026)

Localized Comparisons

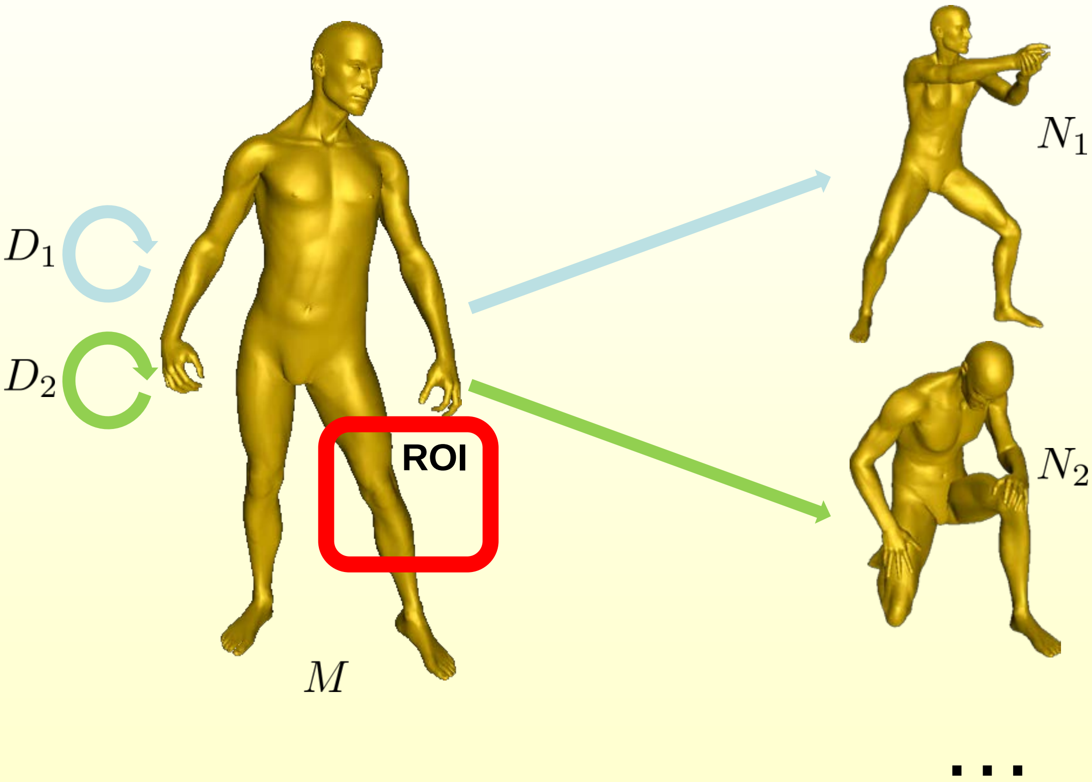

$$\rho : M \to \mathbb{R}$$

supported in RoI

$$D_1\rho \text{ to } D_2\rho$$

Exaggeration of Difference in RoI

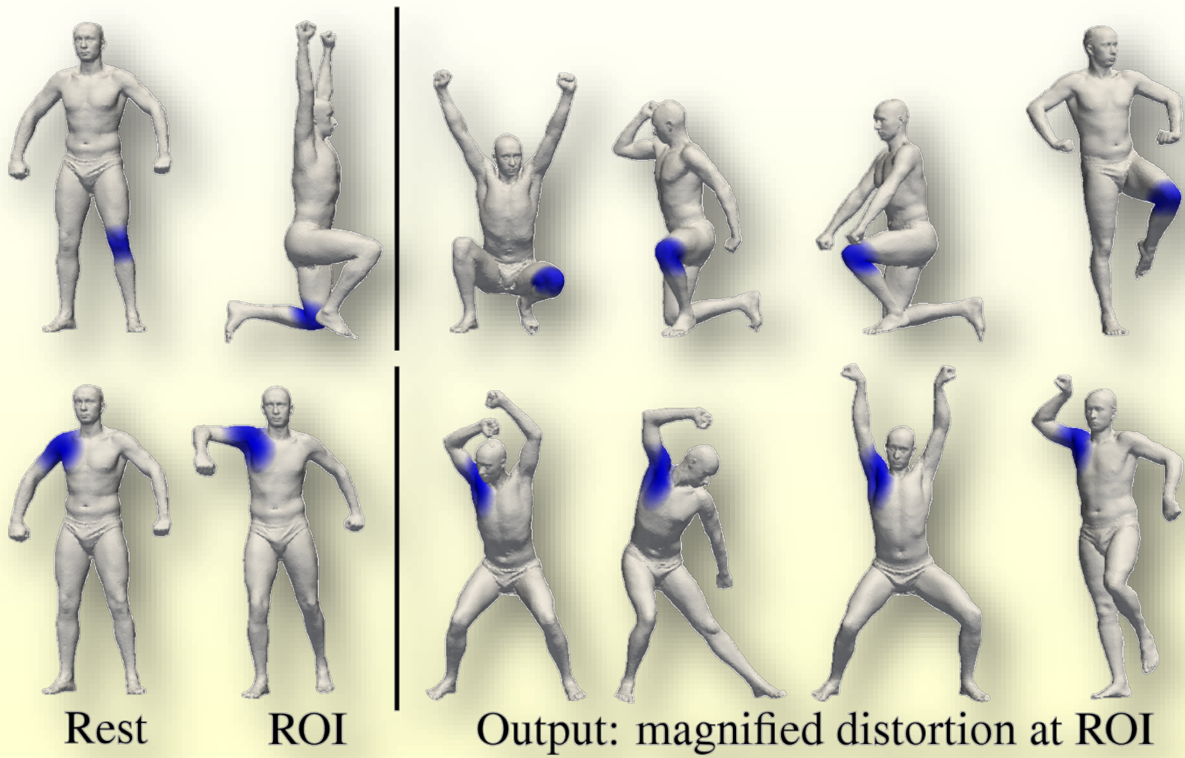

Comparing Differences II

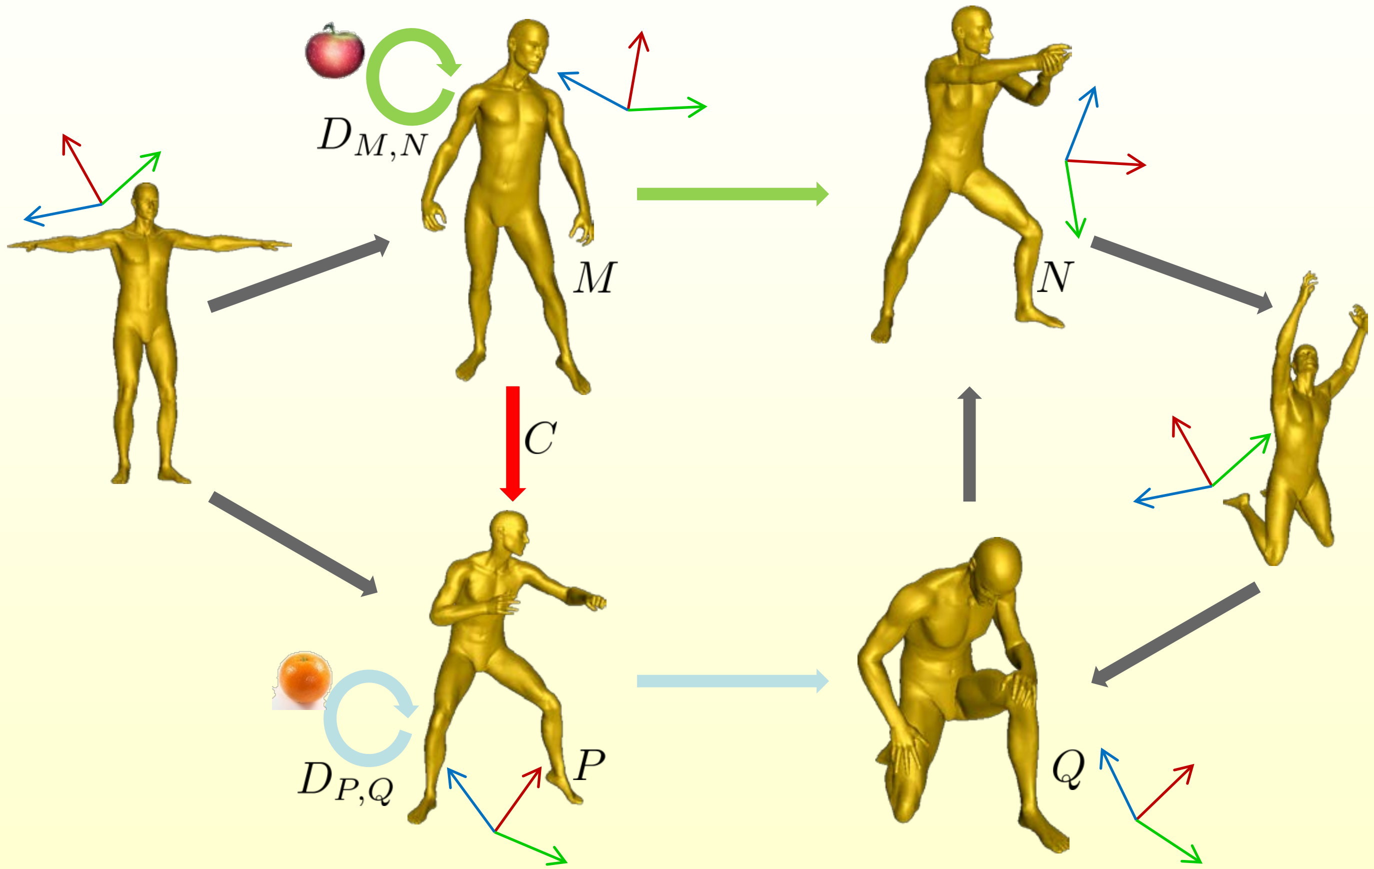

$$D_{M,N} \sim C^{-1} D_{P,Q} C$$

Analogies: D relates to C as B relates to A

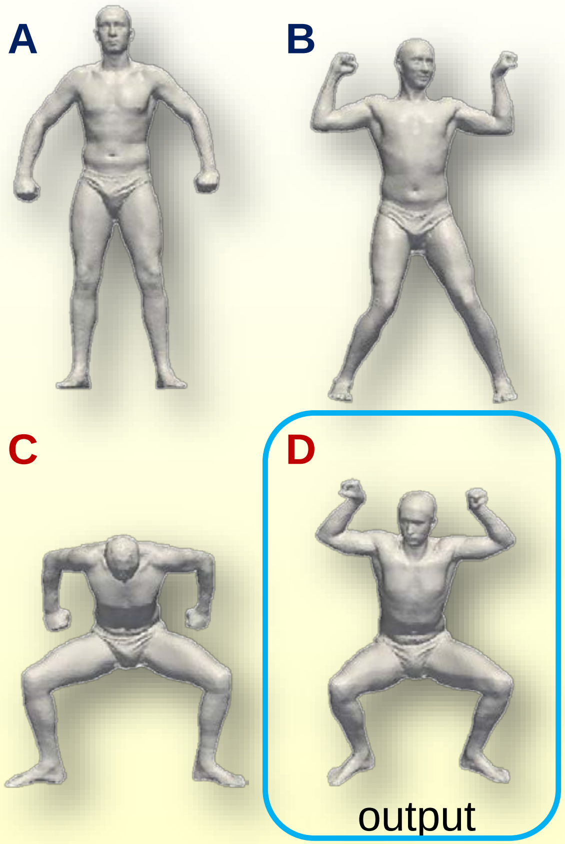

D = C + (B – A)

hands raised up (B – A)

Analogies: D relates to C as B relates to A

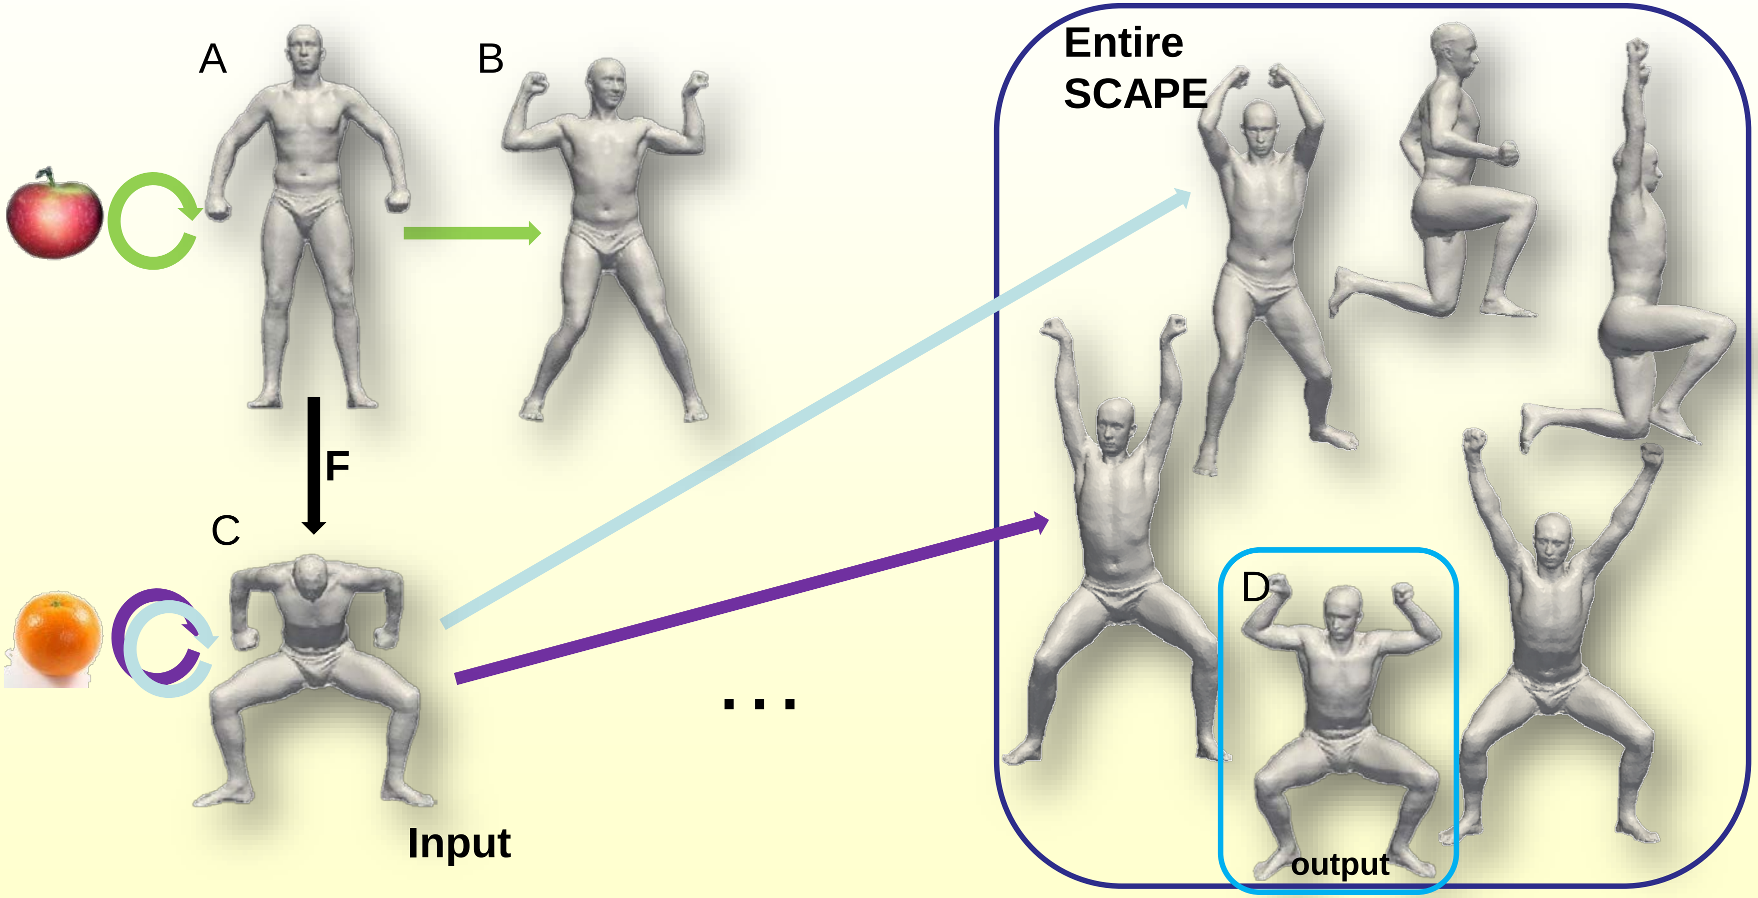

Shape Analogies

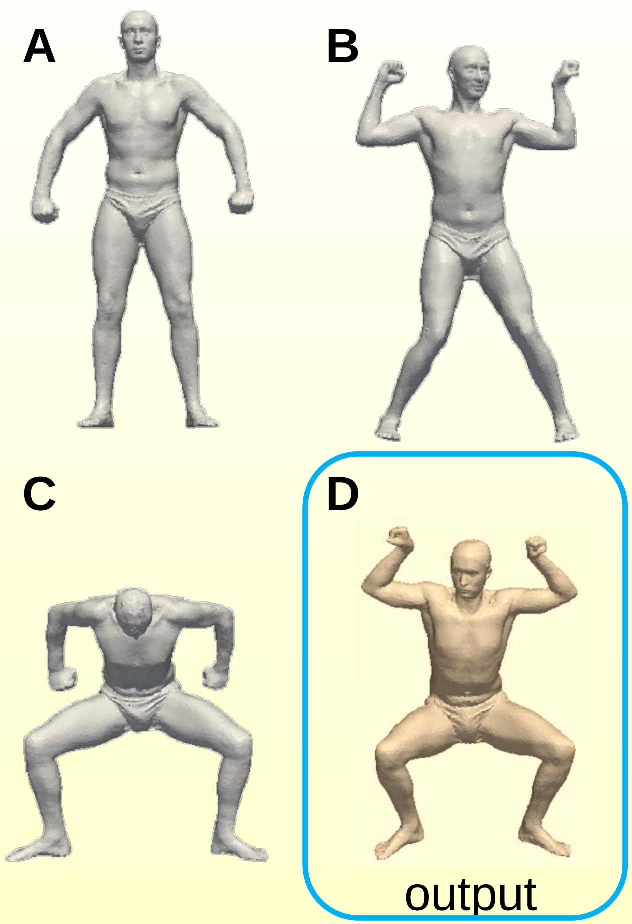

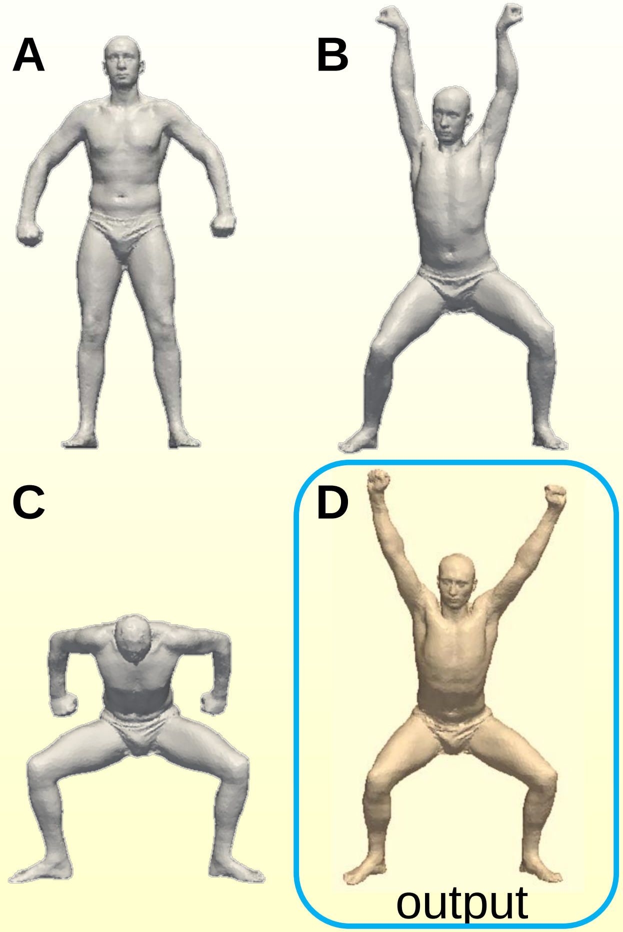

Extrinsic Shape Differences
Shape Synthesis

[E. Corman, J. Solomon, M. Ben-Chen, L. J. Guibas, and M. Ovsjanikov, 2016]

Intrinsic differences of an offset surface
capture extrinsic distortions of the original surface

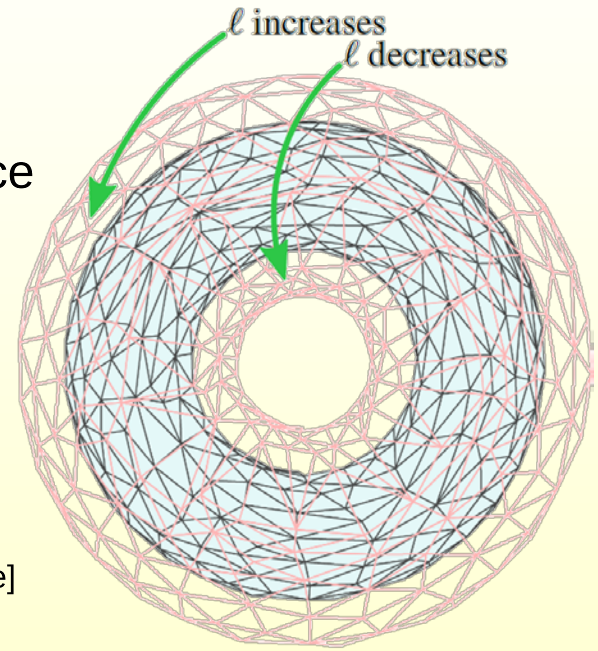

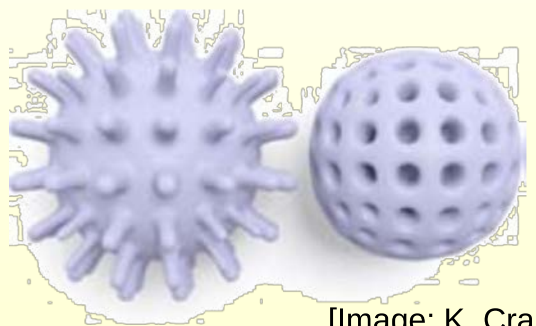

[Image: K. Crane]

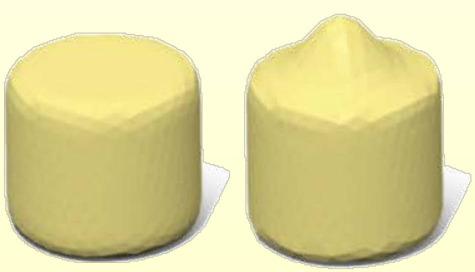

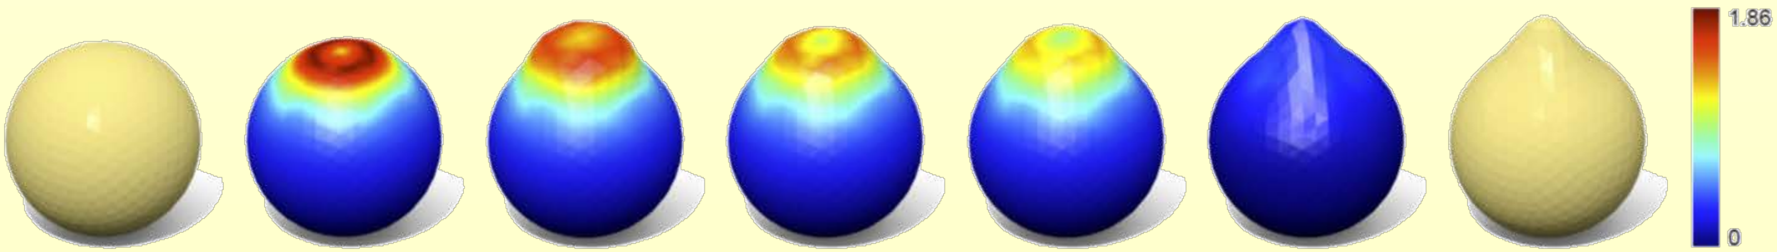

Comparing Differences III

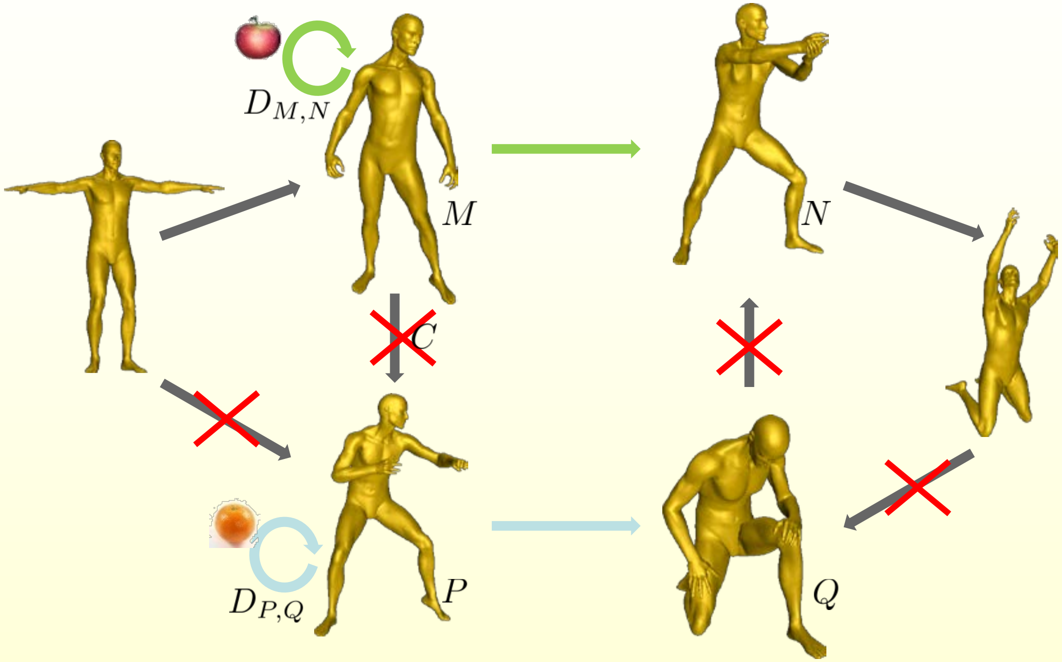

$$D_{M,N} \sim C^{-1} D_{P,Q} C$$

$$Spec(D_{M,N}) \sim Spec(D_{P,Q})$$

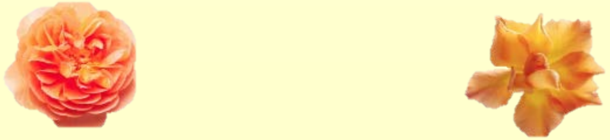

Aligning Disconnected Collections

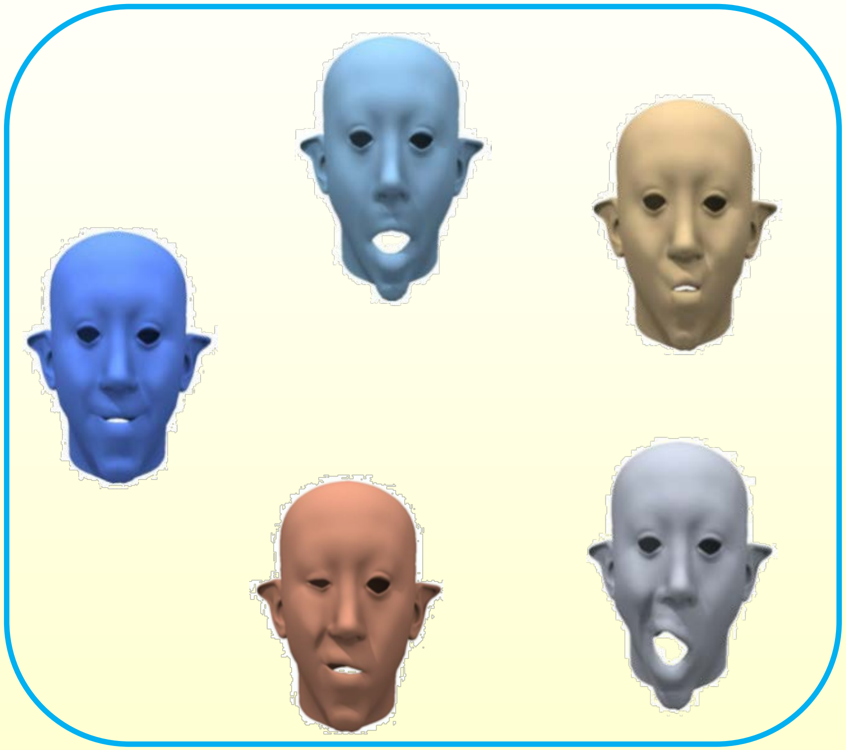

First Collection

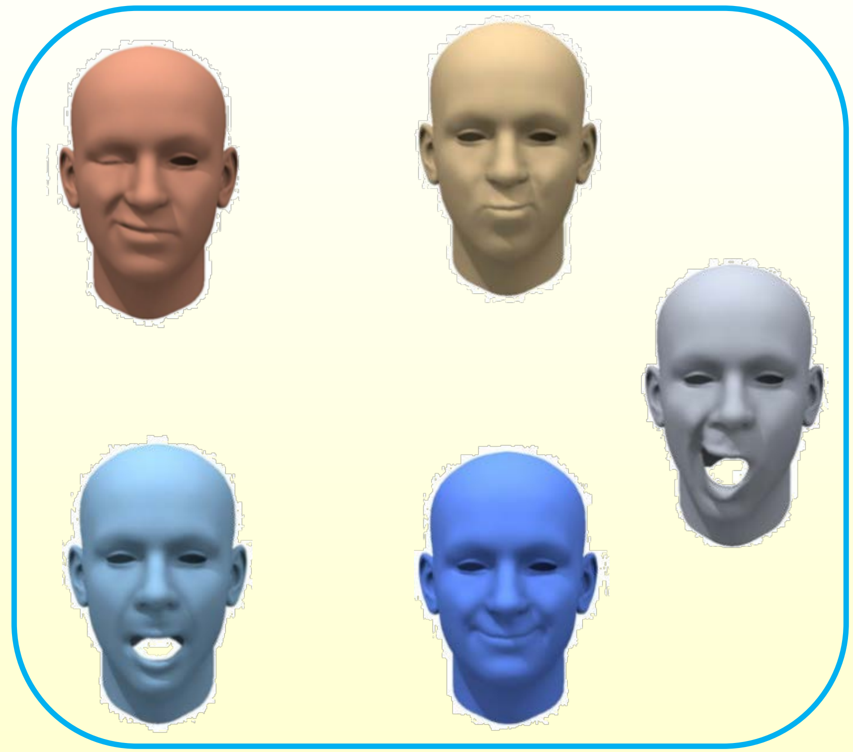

Second Collection

Aligning Disconnected Collections

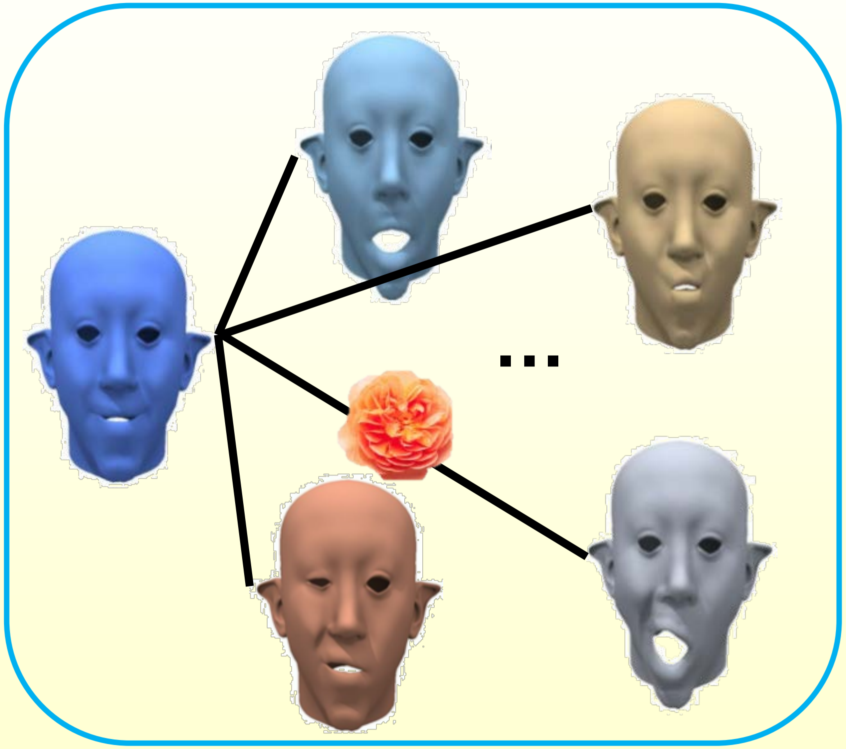

Complete graph

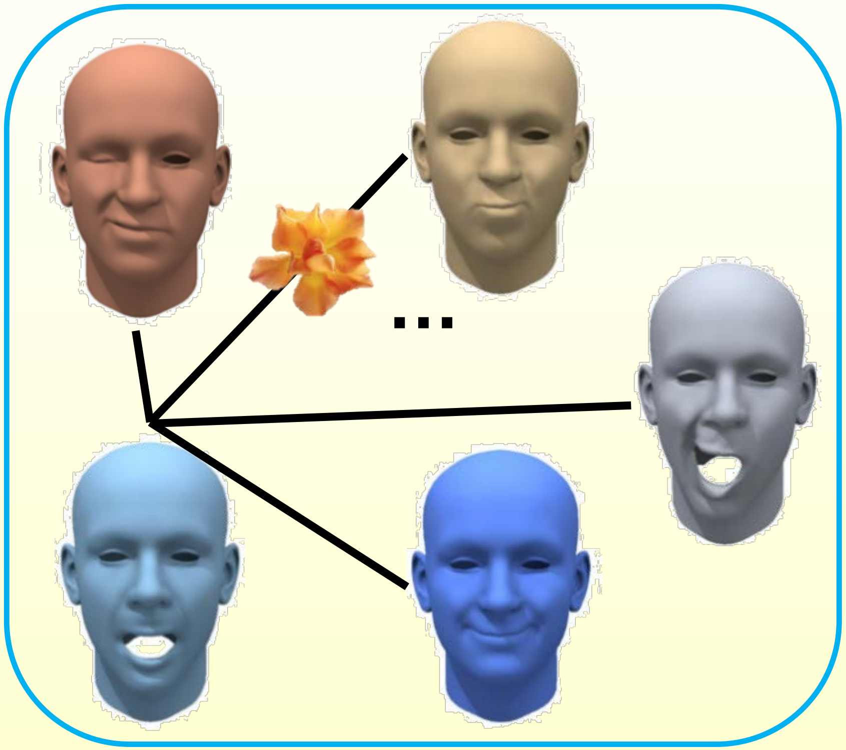

Complete graph

Aligning, Without "Crossing the River"

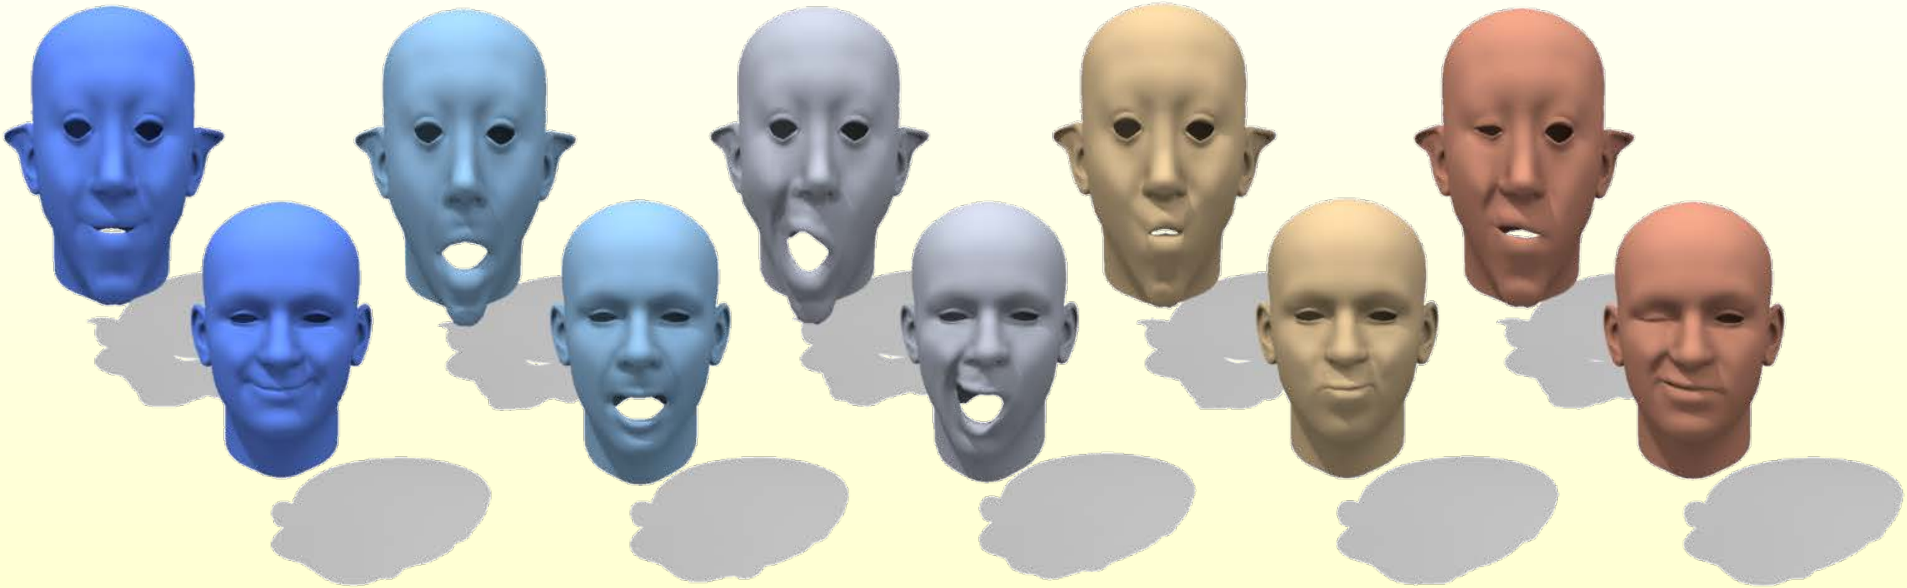

Comparing the differences is sometimes easier than comparing the originals

Shape Differences are a Change Recipe

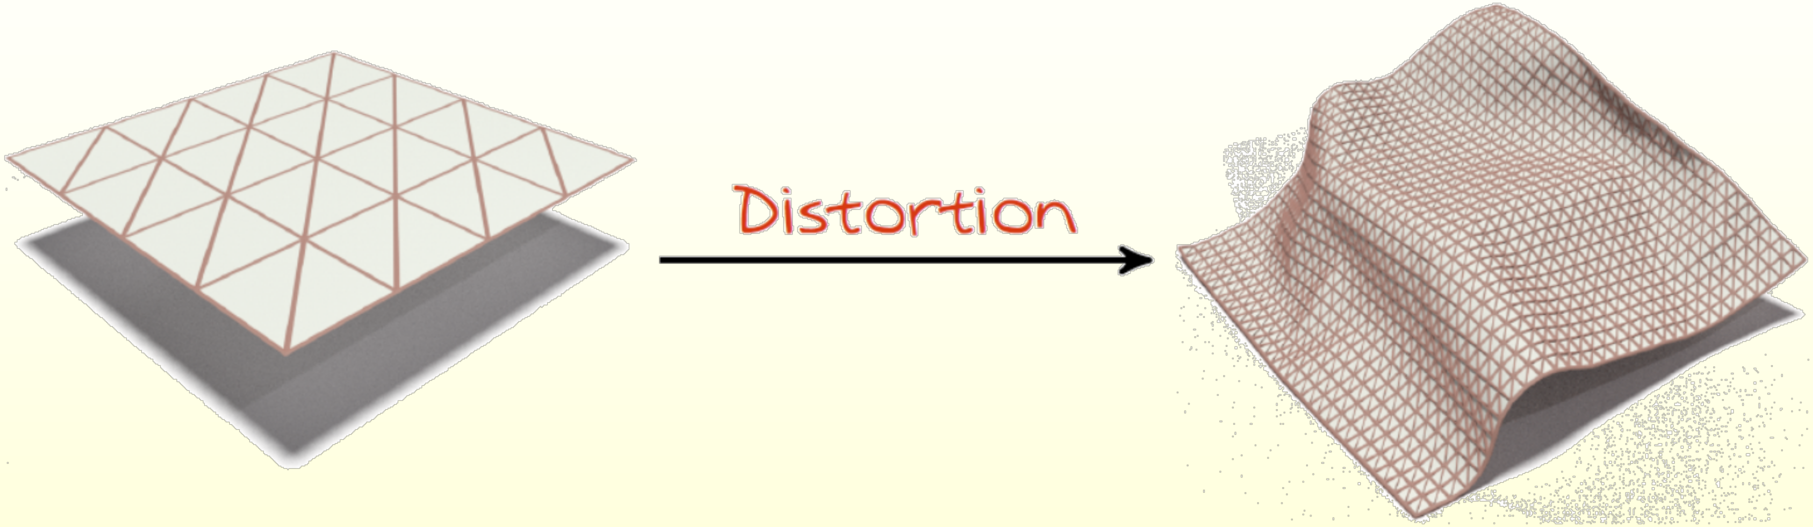

A recipe encoded as a matrix:

Area distortion
Conformal distortion

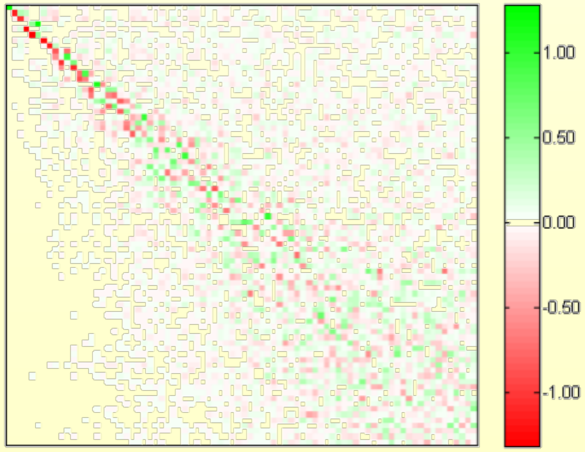

A novel type of latent space representation for 3D data

Large Networks: Consistency of Network Transport

Map Networks for Related Data

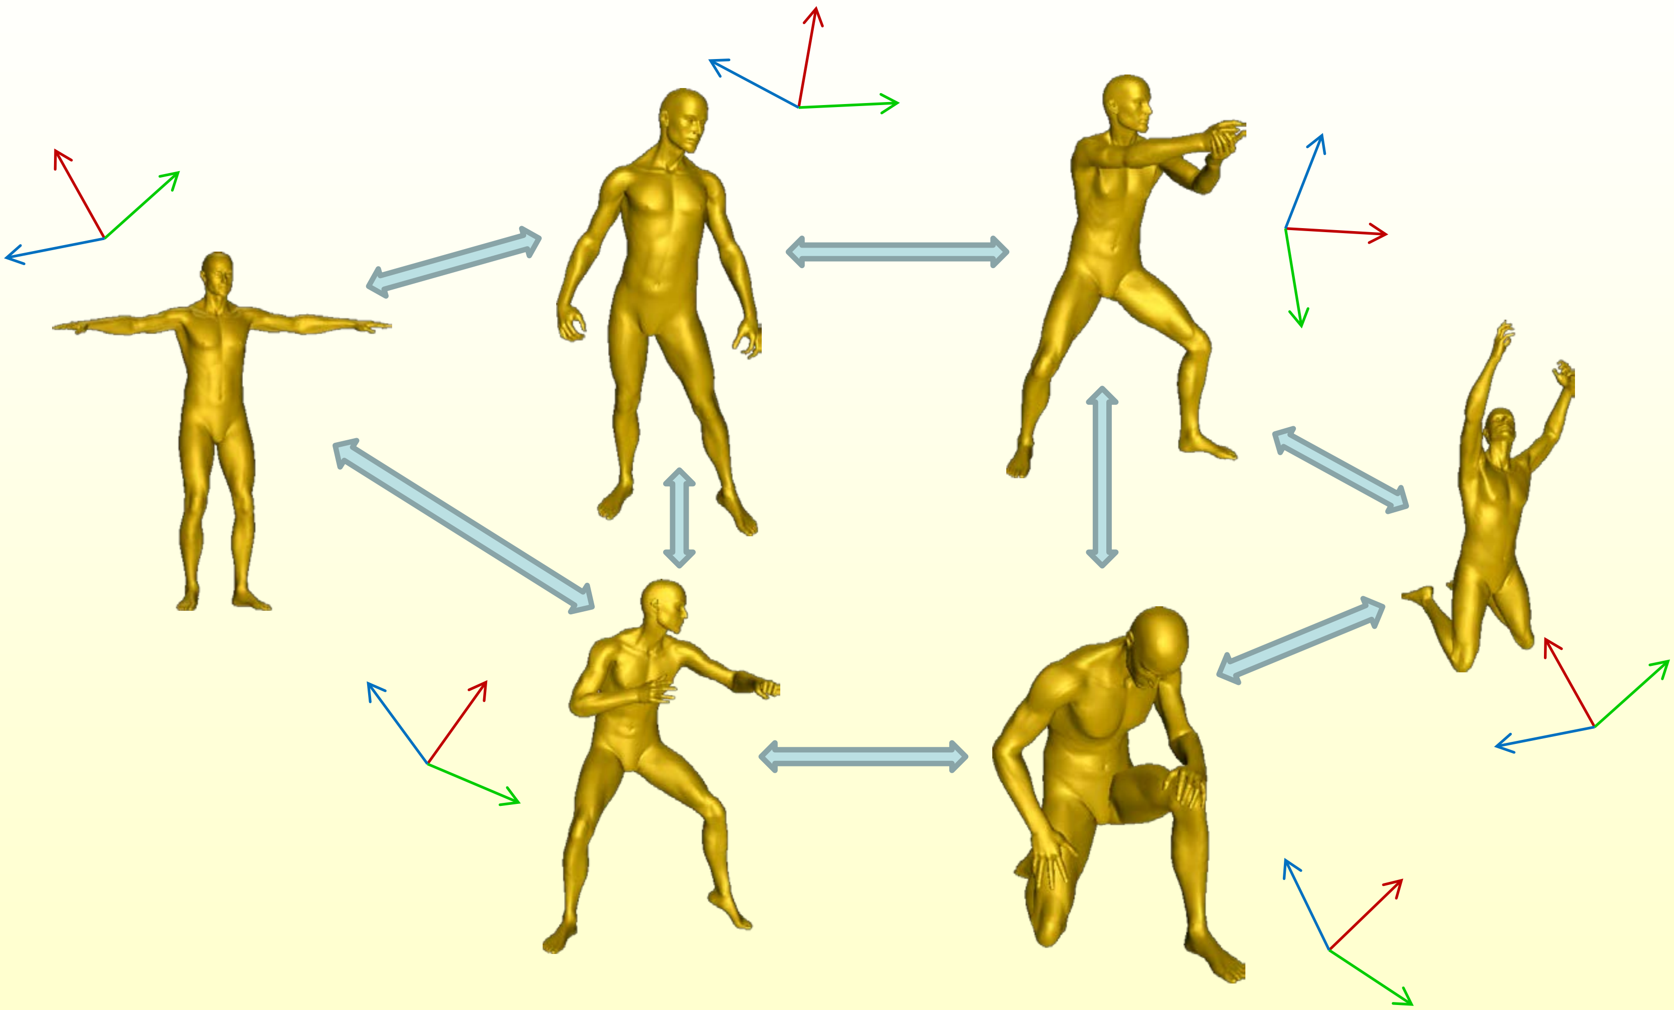

Networks of “samenesses”

A Functorial View of Data

Herni Cartan

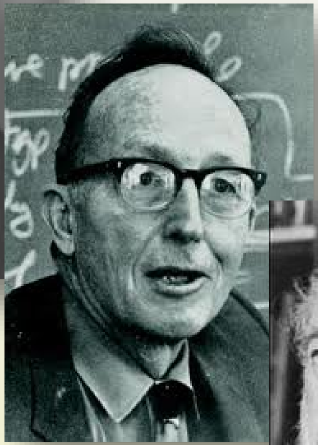

Saunders MacLane

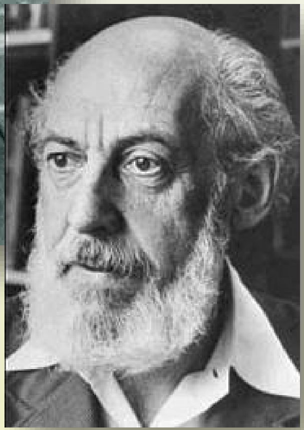

Samuel Eilenberg

The Information is in the Maps

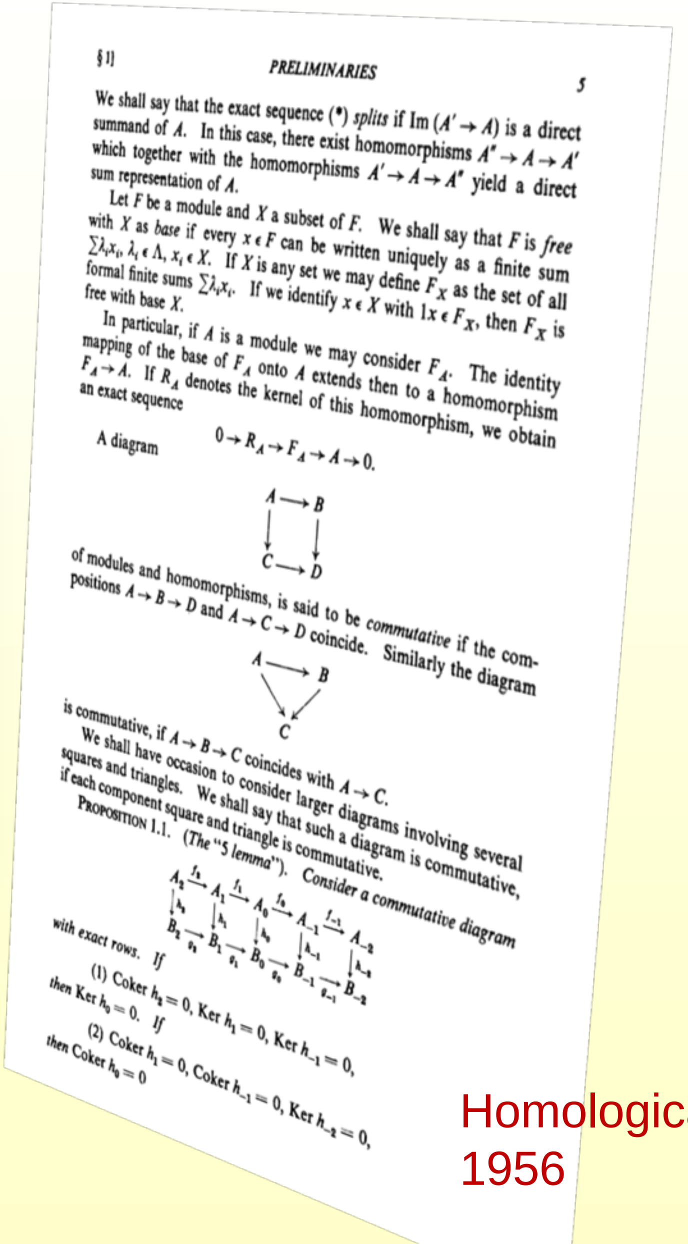

§ 1] PRELIMINARIES 5

We shall say that the exact sequence (*) *splits* if Im $(A' \to A)$ is a direct summand of A. In this case, there exist homomorphisms $A'' \to A \to A'$ which together with the homomorphisms $A' \to A \to A''$ yield a direct sum representation of A.

Let F be a module and X a subset of F. We shall say that F is *free* with X as *base* if every $x \in F$ can be written uniquely as a finite sum $\sum \lambda_i x_i$, $\lambda_i \in \Lambda$, $x_i \in X$. If X is any set we may define F_X as the set of all formal finite sums $\sum \lambda_i x_i$. If we identify $x \in X$ with $1x \in F_X$, then F_X is free with base X.

In particular, if A is a module we may consider F_A. The identity mapping of the base of F_A onto A extends then to a homomorphism $F_A \to A$. If R_A denotes the kernel of this homomorphism, we obtain an exact sequence

$$0 \to R_A \to F_A \to A \to 0.$$

A diagram

$$\begin{array}{ccc} A & \longrightarrow & B \\ \downarrow & & \downarrow \\ C & \longrightarrow & D \end{array}$$

of modules and homomorphisms, is said to be *commutative* if the compositions $A \to B \to D$ and $A \to C \to D$ coincide. Similarly the diagram

$$\begin{array}{ccc} A & \longrightarrow & B \\ & \searrow \quad \swarrow & \\ & C & \end{array}$$

is commutative, if $A \to B \to C$ coincides with $A \to C$.

We shall have occasion to consider larger diagrams involving several squares and triangles. We shall say that such a diagram is commutative, if each component square and triangle is commutative.

PROPOSITION 1.1. (*The "5 lemma"*). *Consider a commutative diagram*

$$\begin{array}{ccccccccc} A_2 & \xrightarrow{f_2} & A_1 & \xrightarrow{f_1} & A_0 & \xrightarrow{f_0} & A_{-1} & \xrightarrow{f_{-1}} & A_{-2} \\ \downarrow h_2 & & \downarrow h_1 & & \downarrow h_0 & & \downarrow h_{-1} & & \downarrow h_{-2} \\ B_2 & \xrightarrow[g_2]{} & B_1 & \xrightarrow[g_1]{} & B_0 & \xrightarrow[g_0]{} & B_{-1} & \xrightarrow[g_{-1}]{} & B_{-2} \end{array}$$

with exact rows. If

(1) Coker $h_2 = 0$, Ker $h_1 = 0$, Ker $h_{-1} = 0$, *then* Ker $h_0 = 0$. *If*

(2) Coker $h_1 = 0$, Coker $h_{-1} = 0$, Ker $h_{-2} = 0$, *then* Coker $h_0 = 0$

Homological Algebra 1956

Yes, But With a Statistical Flavor

- Yes, straight out of the playbook of homological algebra / algebraic topology
- But, the maps
  - are not given by canonical constructions
  - they have to be estimated and can be noisy
  - the network acts as a regularizer -- commutativity and cycle closure constraints
    - fix the maps
    - map recovery is possible even with 50% error!

Maps are composable, algebraic objects

Maps processing

Fixing Maps

[Q. Huang, G. Zhang, L. Gao, S. Hu, A. Bustcher, and L. Guibas, 2012]

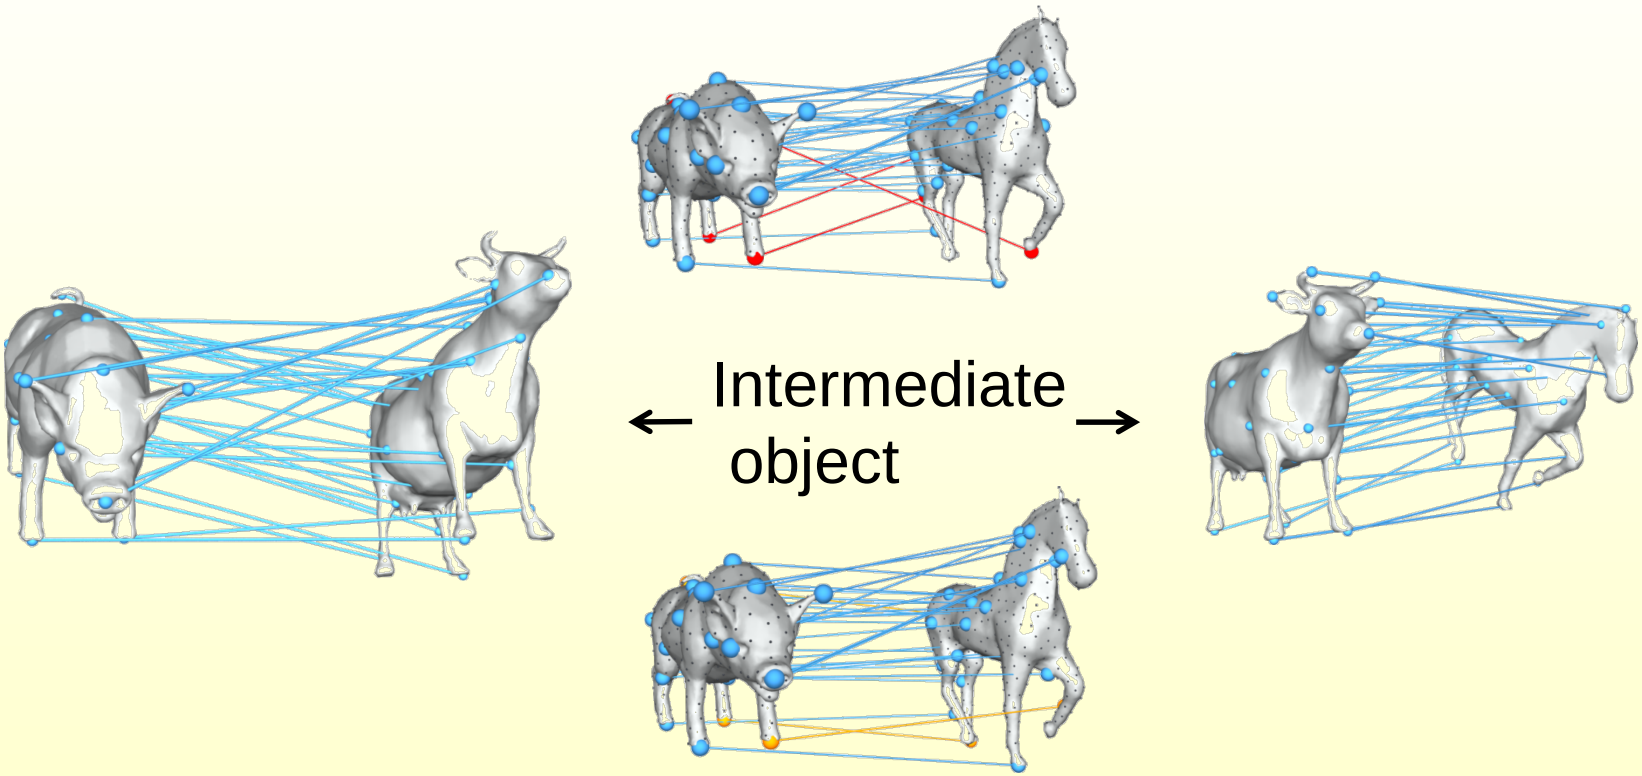

Cycle-Consistency ≡ Low-Rank

- In a map network, commutativity, path-invariance, or cycle-consistency are equivalent to a low rank or semidefiniteness condition on a big mapping matrix

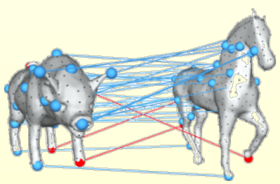

$$X = \begin{pmatrix} I_m & X_{1,2} & \cdots & X_{1,n} \\ X_{1,2} & I_m & \cdots & \cdots \\ \vdots & \vdots & I_m & X_{(n-1),n} \\ X_{n,1} & \vdots & X_{n,(n-1)} & I_m \end{pmatrix}.$$

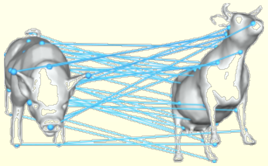

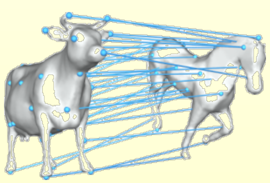

- Conversely, such a low-rank condition can be used to
  - regularize and clean up functional maps
  - extract shared structure

Map processing!

Map Synchronization by Factorization

$$X = \begin{bmatrix} I_m & X_{12} & \cdots & X_{1n} \\ X_{21} & I_m & \cdots & \vdots \\ \vdots & \vdots & \ddots & X_{n-1,n} \\ X_{n1} & \cdots & X_{n,n-1} & I_m \end{bmatrix}$$

$$X_{ij} = X_{j1} X_{i1}^T$$

$$= \begin{bmatrix} I_m \\ \vdots \\ X_{n1} \end{bmatrix} \begin{bmatrix} I_m & \cdots & X_{n1}^T \end{bmatrix}$$

Primal Maps: Cycle Consistency via Convex Programming

Basic Primal Setting

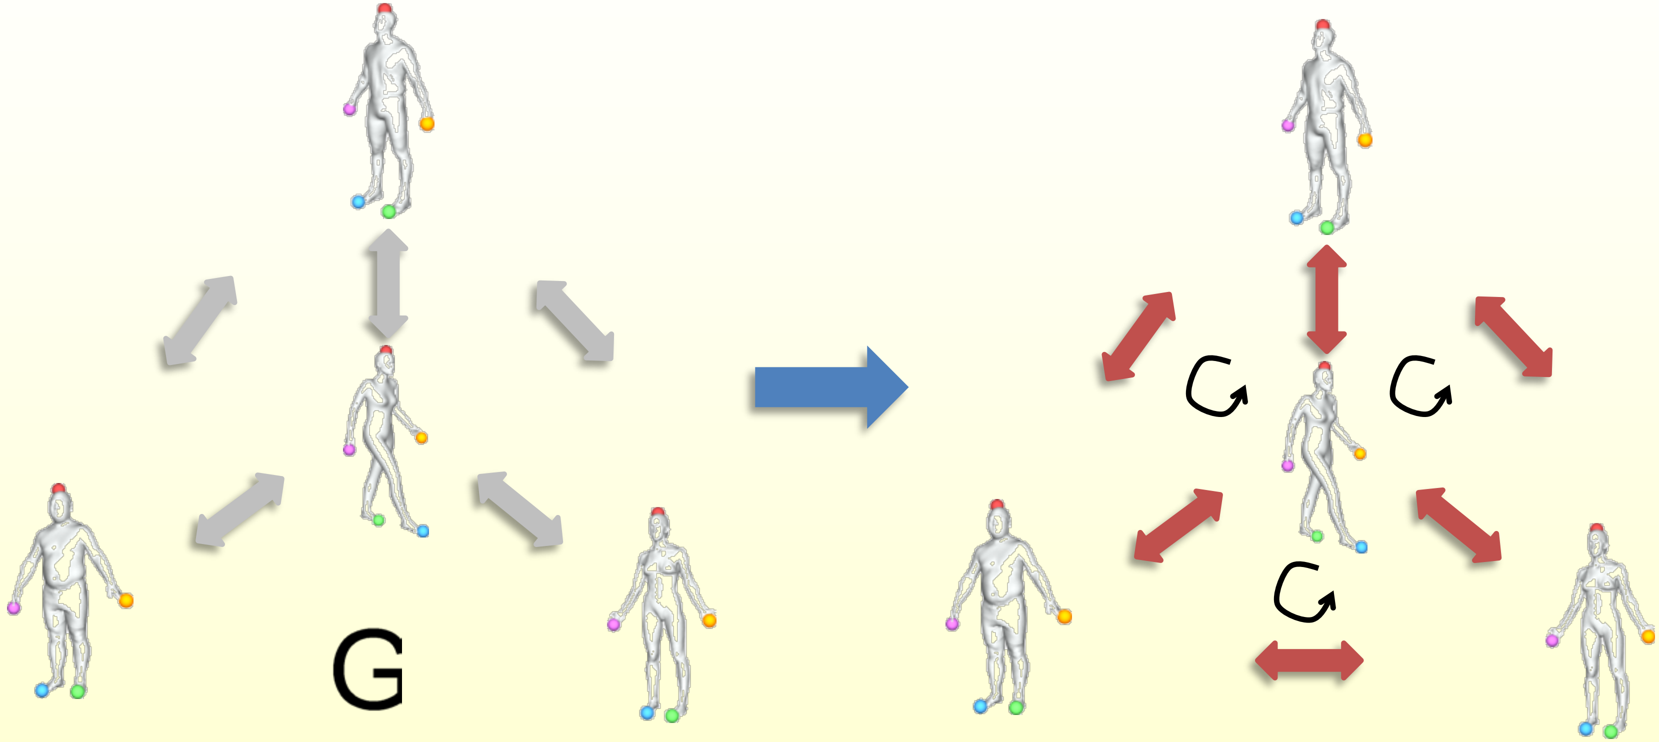

n objects, each object has *m* points (keypoints)

Matrix Representation of Primal Maps

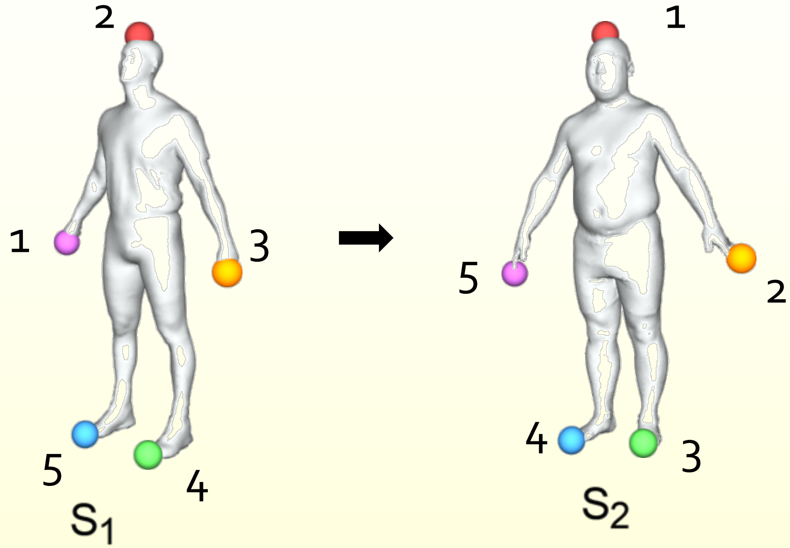

$$X_{12} = \begin{bmatrix} 0 & 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \end{bmatrix}$$

$$X = \begin{bmatrix} I_m & X_{12} & \cdots & X_{1n} \\ X_{12}^T & I_m & \cdots & \vdots \\ \vdots & \vdots & \ddots & X_{(n-1),n} \\ X_{1n}^T & \vdots & X_{(n-1),n}^T & I_m \end{bmatrix}$$

- Diagonal blocks are identity matrices
- Off diagonal blocks are permutation matrices
- Symmetric

Cycle Consistency Constraint

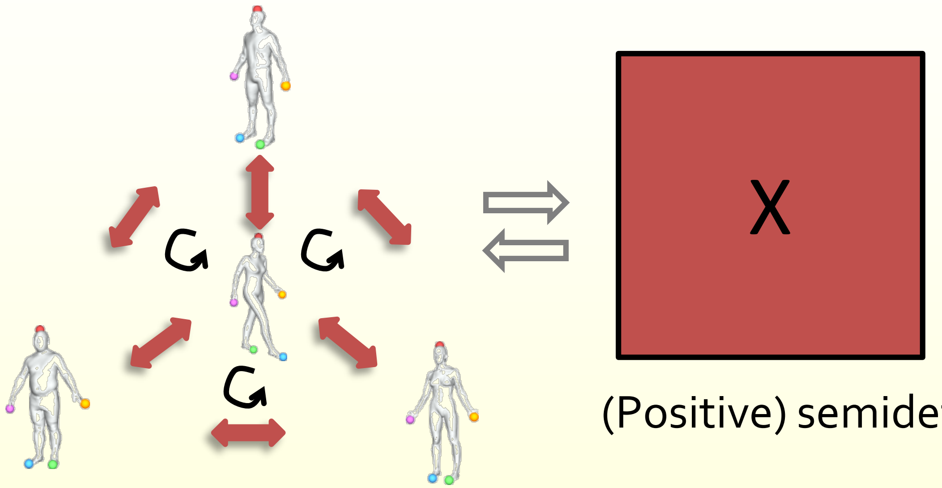

$\succeq 0$

(Positive) semidefiniteness

$$\mathbf{X}_{ij} = X_{j1}^T X_{i1} \quad \Longleftrightarrow \quad \mathbf{X} = \begin{bmatrix} \mathbf{I}_m \\ \vdots \\ \mathbf{X}_{n1}^T \end{bmatrix} \begin{bmatrix} \mathbf{I}_m & \cdots & \mathbf{X}_{n1} \end{bmatrix}$$

Low rank

Map Representation

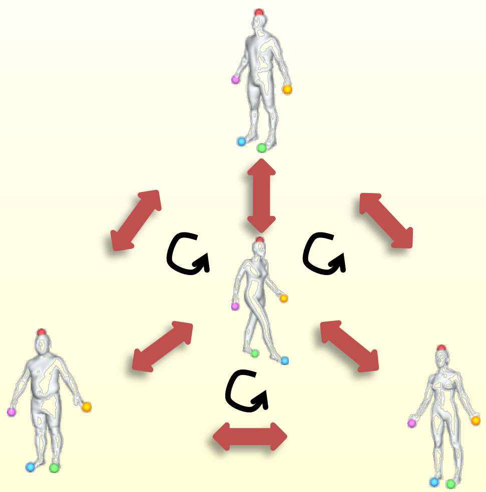

$$\mathbf{X} =$$

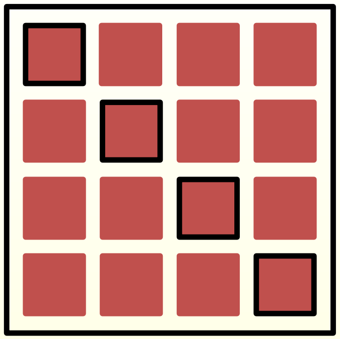

$$\mathbf{X}_{ii} = I_m, \quad 1 \le i \le n$$

$$\mathbf{X}_{ij}\mathbf{1} = \mathbf{1}, \mathbf{X}_{ij}^T\mathbf{1} = \mathbf{1}, \quad 1 \le i < j \le n$$

$$\mathbf{X} \succeq 0$$

$$\mathbf{X} \in \{0,1\}^{nm}$$

Relaxation / Convexification

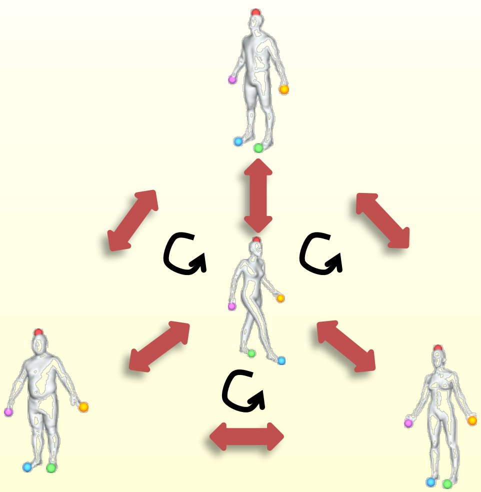

$\mathbf{X} =$ 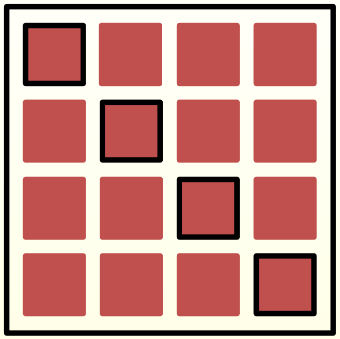

$$\mathbf{X}_{ii} = I_m, \quad 1 \le i \le n$$

$$\mathbf{X}_{ij}\mathbf{1} = \mathbf{1}, \mathbf{X}_{ij}^T\mathbf{1} = \mathbf{1}, \quad 1 \le i < j \le n$$

$$\mathbf{X} \succeq 0$$

$$1 \ge \mathbf{X} \ge 0$$

Convex Program: SDP Formulation

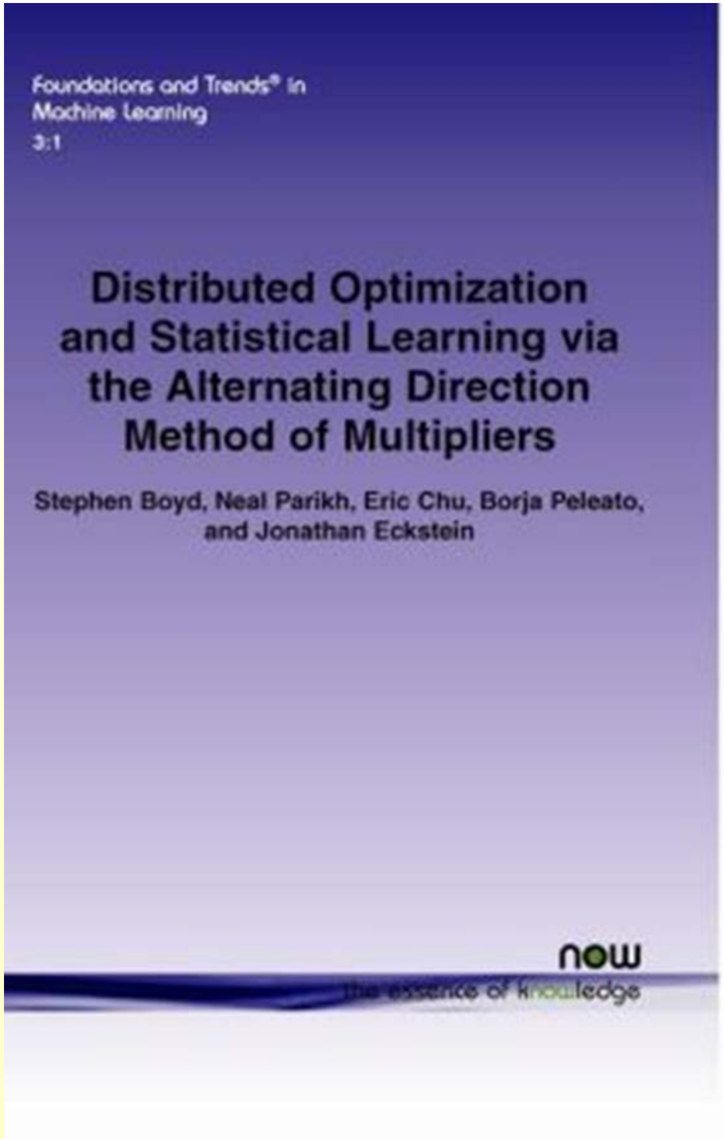

$$\text{minimize} \quad \sum_{(i,j)\in\mathcal{E}} \|\mathbf{X}_{ij}^{input} - \mathbf{X}_{ij}\|_1$$

$$\text{subject to} \quad \mathbf{X}_{ii} = I_m, \quad 1 \leq i \leq n$$

$$\mathbf{X}_{ij}\mathbf{1} = \mathbf{1}, \mathbf{X}_{ij}^T\mathbf{1} = \mathbf{1}, \quad 1 \leq i < j \leq n$$

$$\mathbf{X} \succeq 0$$

$$1 \geq \mathbf{X} \geq 0$$

ADMM [Boyd et al.11]

Map Correction

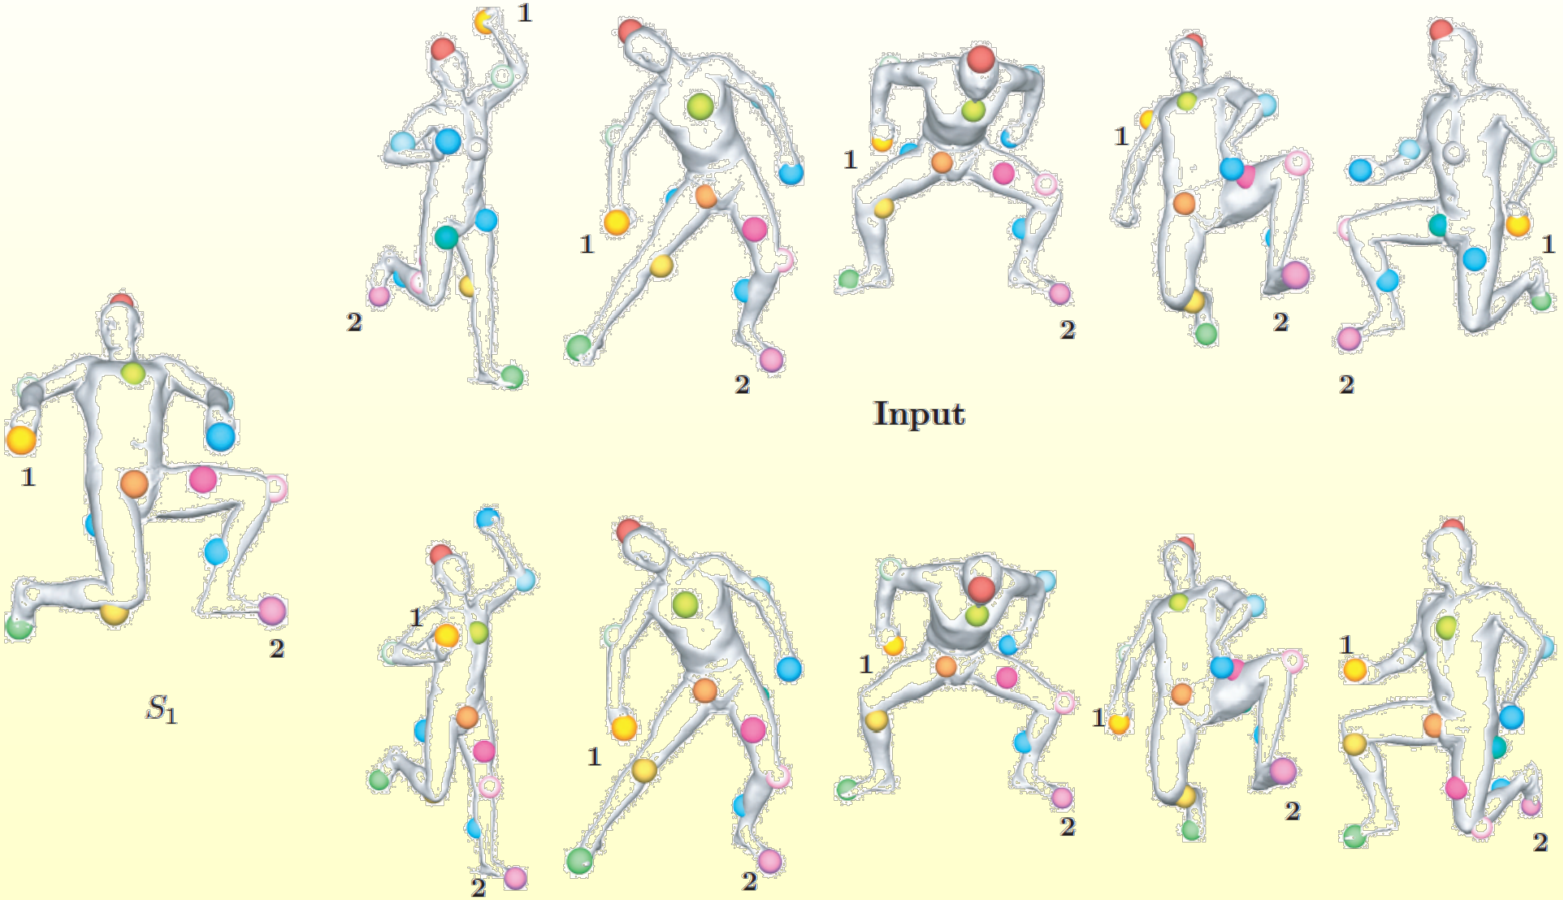

Exact Recovery Conditions

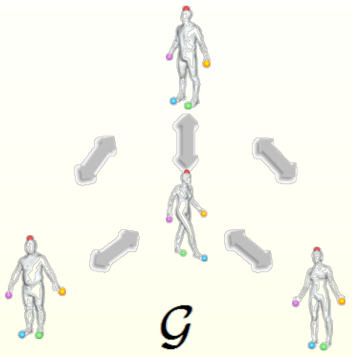

fraction incorrect corres. per point

< algebraic-connectivity(G)/4n

$\lambda_2(L_G)$

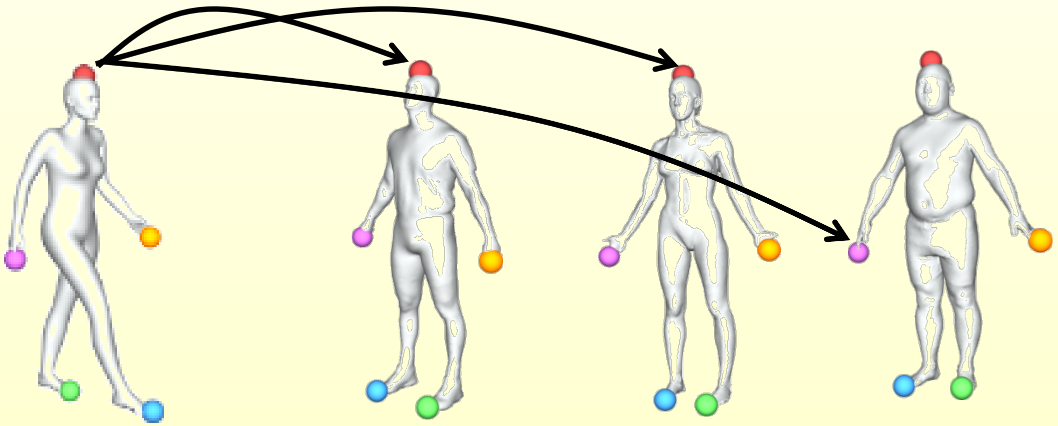

Algorithm error tolerance

For Complete Graph

- 25% incorrect correspondences
- Worst-case scenario
  - Two clusters of objects of equal size
  - Wrong correspondences between objects of different clusters only (50%)

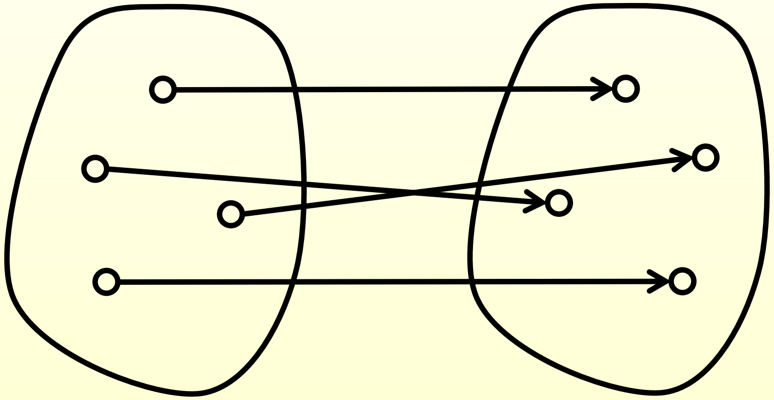

- Can recover even with 50% wrong correspondences in random case

Phase Transition

Numerical optimization: ADMM [Wen et al. 10, Boyd et al. 11]

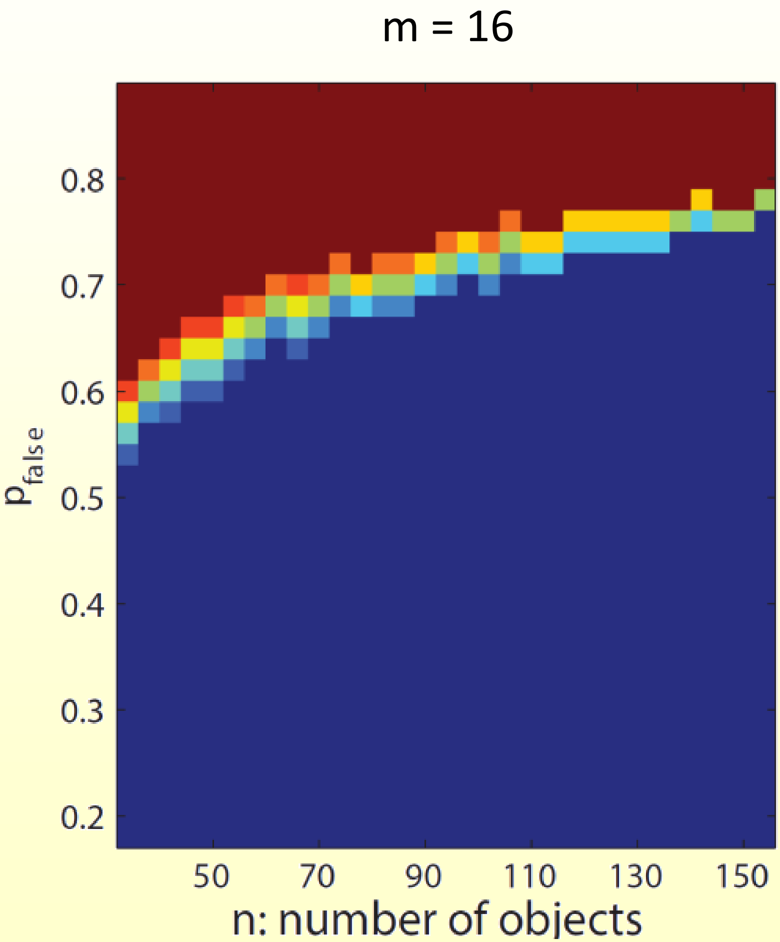

Red: never recovers

Blue: always recovers

Chair Data Set

16 images
Clustered SIFT features
+ RANSAC (60-120 points per image)

SDP Running time: 2m19s (3.2GHZ, single core)

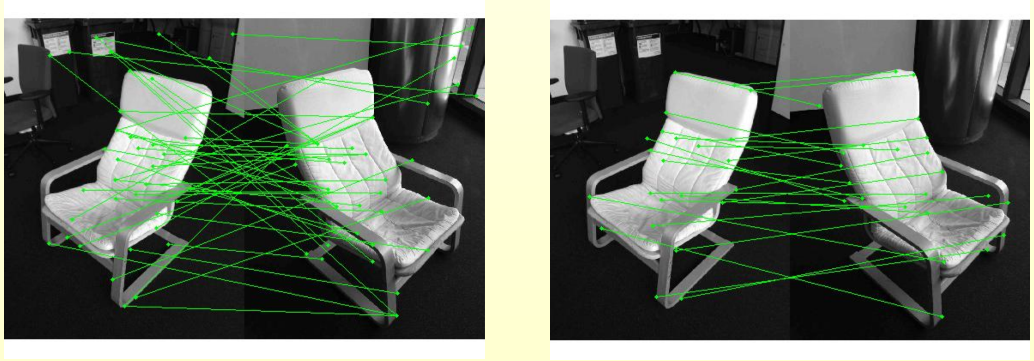

Input output

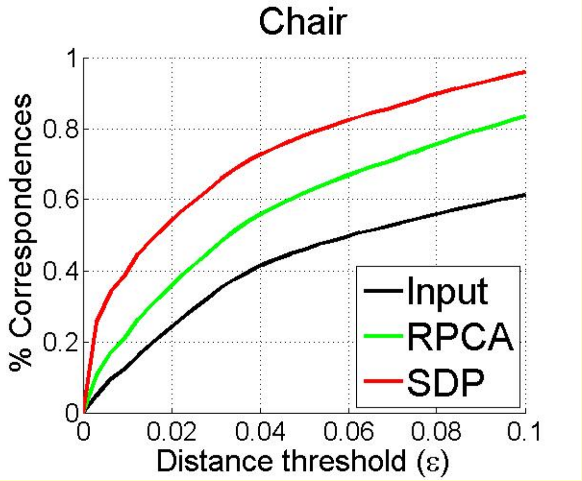

Functional Map Based Shared Structure Discovery

Entity Extraction in Images

[F. Wang, Q. Huang, L. G., ICCV ’13]

- Task: jointly segment a set of related images
  - same object, different viewpoints/scales:

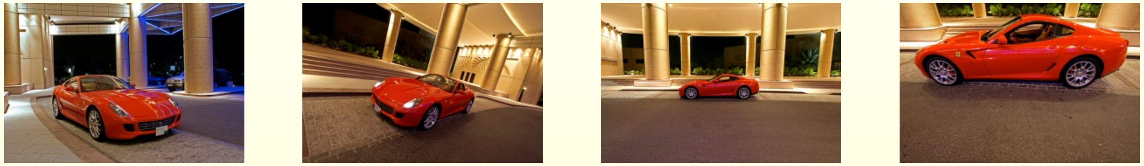

  - similar objects of the same class:

- Benefits and challenges:
  - Images can provide weak supervision for each other
  - But exactly how should they help each other? How to deal with clutter and irrelevant content?

Co-Segmentation via an Image Network

- Image similarity graph based on GIST
  - Each edge has global image similarity w_{ij} and functional maps in both directions;
  - Sparse if large.

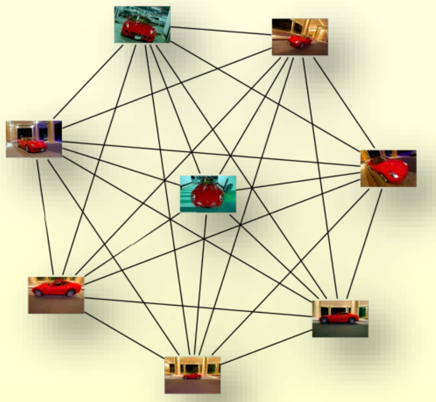

Graph for iCoseg-Ferrari

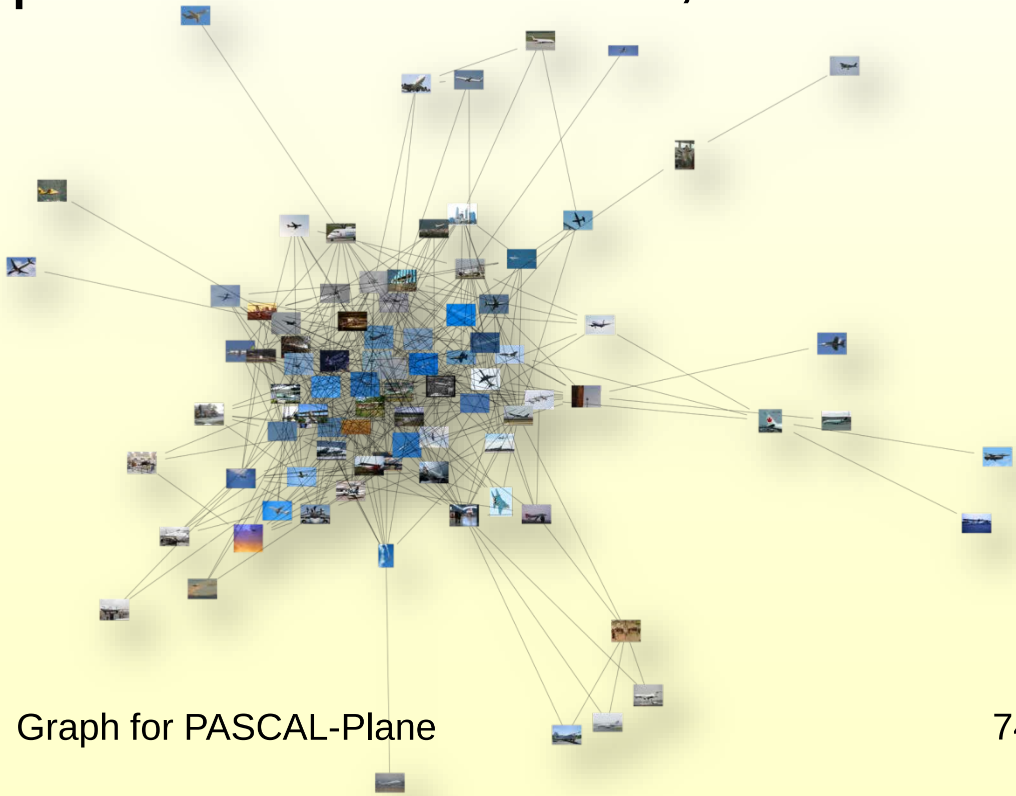

Graph for PASCAL-Plane

Superpixel Representation

- Over-segment images into super-pixels
- Build a graph on super-pixels
  - Nodes: super-pixels
  - Edges weighted by length of shared boundary

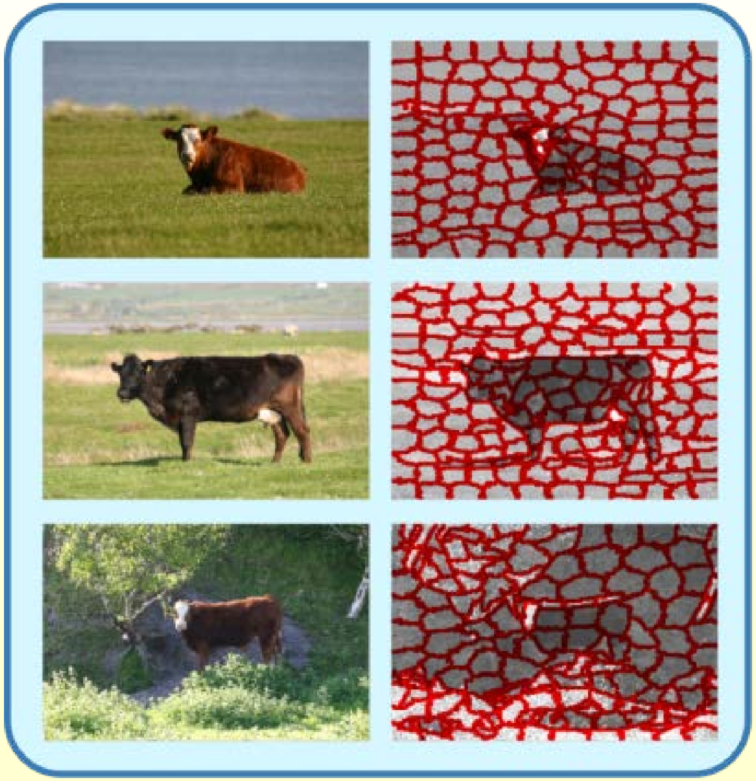

Encoding Functions over Graphs

- Basis of functional space
  - : First M Laplacian eigenfunctions of the graph

$$f = \sum_{j=1}^{M} f_j b_i^j = B_i \mathbf{f}$$

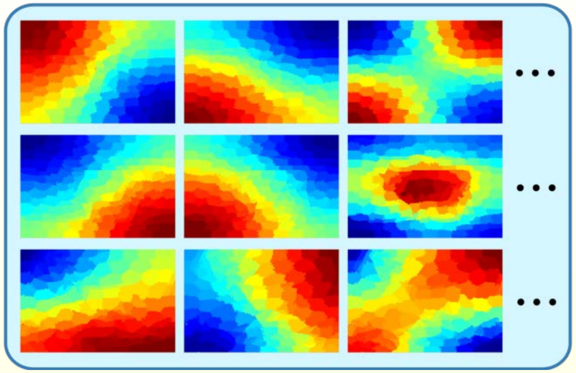

- Reconstruct any function with small error (M=30)

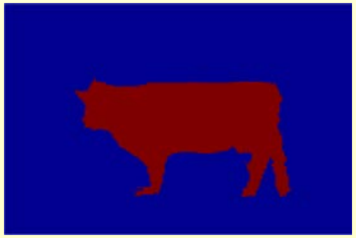

Binary indicator function

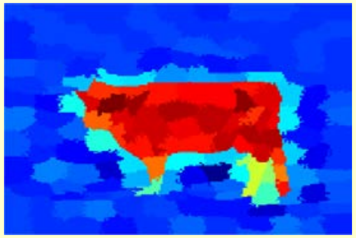

Reconstructed function

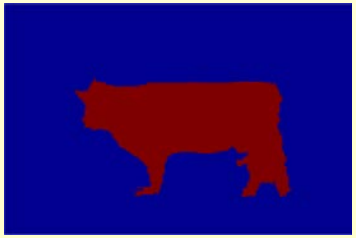

Thresholded reconstructed function

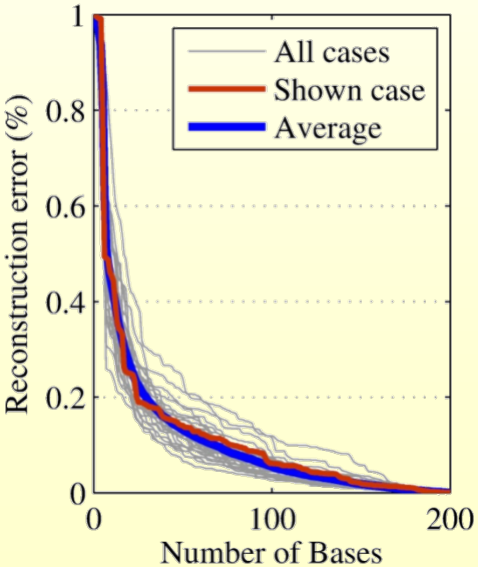

Reconstruction error

Joint Estimation of Functional Maps, I

- Functional map:
  - A linear map between functions in two functional spaces

$$\mathbf{f}' = X_{ij}\mathbf{f} \qquad X_{ij} \in \mathcal{R}^{M \times M}$$

  - Can be recovered by a set of probe functions

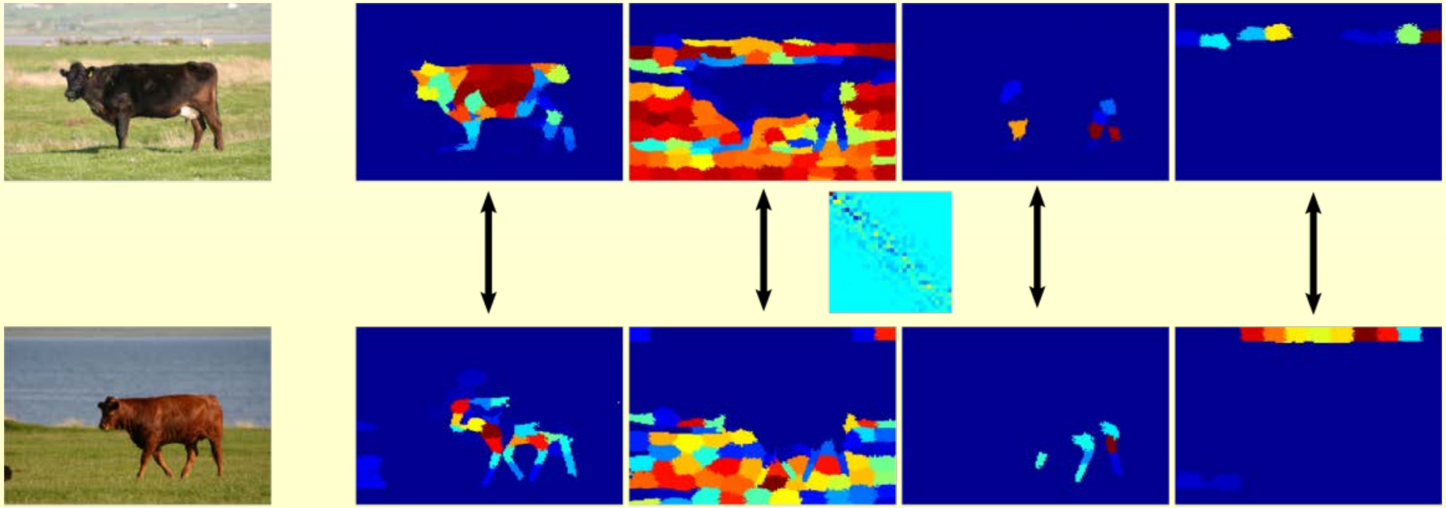

Joint Estimation of Functional Maps, I

- Recover functional maps by aligning image features:

$$f_{ij}^{\text{feature}} = \|X_{ij}D_i - D_j\|_1$$

- Features (probe functions) for each super-pixel:
  - average RGB color, 3-dimensional;
  - 64 dimensional RGB color histogram;
  - 300-dimensional bag-of-visual-words.

Joint Estimation of Functional Maps, II

- Regularization term:

Λ_i, Λ_j diagonal matrices of Laplacian eigenvalues

$$f_{ij}^{\text{reg}} = \|X_{ij}\Lambda_i - \Lambda_j X_{ij}\|^2$$

  - Correspond bases of similar spectra
  - Enforce sparsity of map

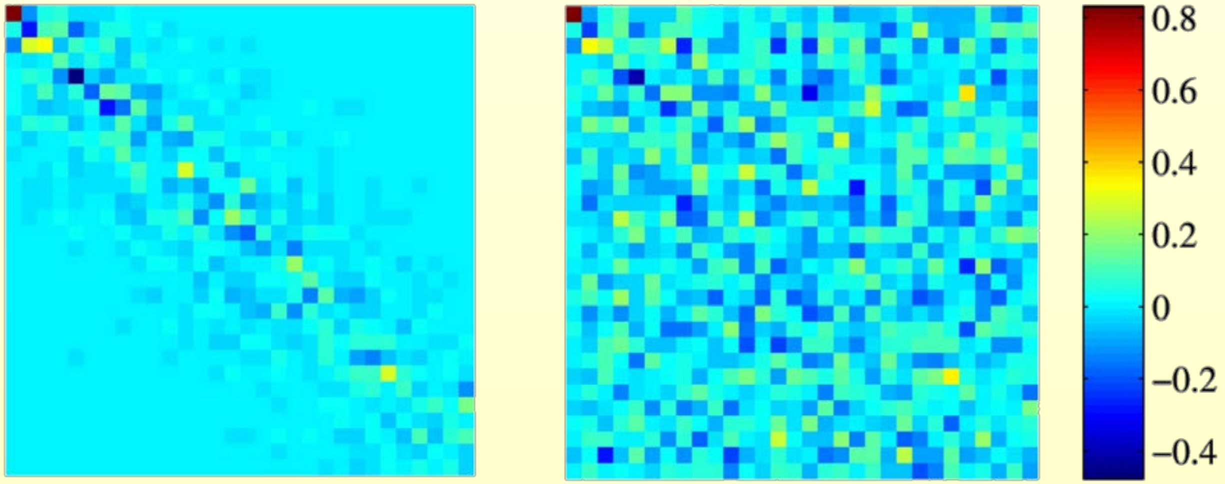

Map with regularization | Map without regularization

Joint Estimation of Functional Maps, III

- Incorporating map cycle consistency:
  - A transported function along any loop should be identical to the original function:

$$X_{i_k i_0} \cdots X_{i_1 i_2} X_{i_0 i_1} \mathbf{f} = \mathbf{f} \quad \Longleftrightarrow \quad X_{ij} Y_i = Y_j, \quad \forall (i,j) \in \mathcal{G}$$

  - Consistency term:

$$f^{\text{cons}} = \sum_{(i,j)\in\mathcal{G}} w_{ij} f_{ij}^{\text{cons}} = \sum_{(i,j)\in\mathcal{G}} w_{ij} \|X_{ij} Y_i - Y_j\|_{\mathcal{F}}^2$$

Image global similarity weight via GIST

Joint Estimation of Functional Maps, III

- Plato's allegory of the cave: a latent space

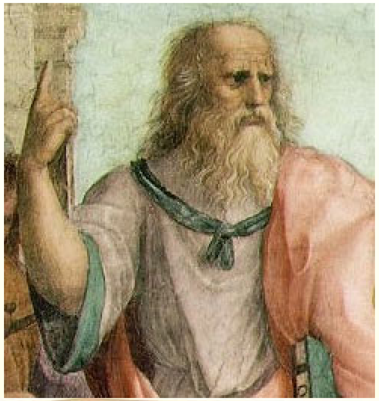

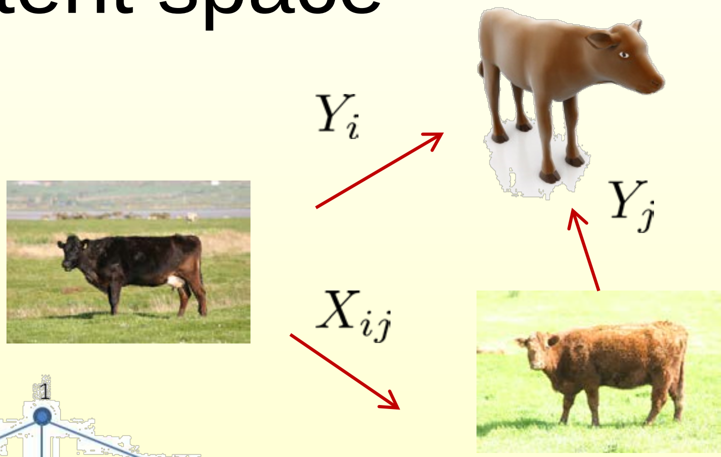

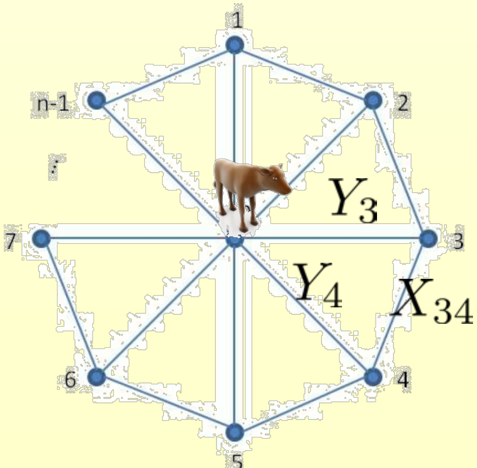

$$X_{ij} \approx Y_j^{-1} Y_i$$

X 30x30, *Y* 30x20

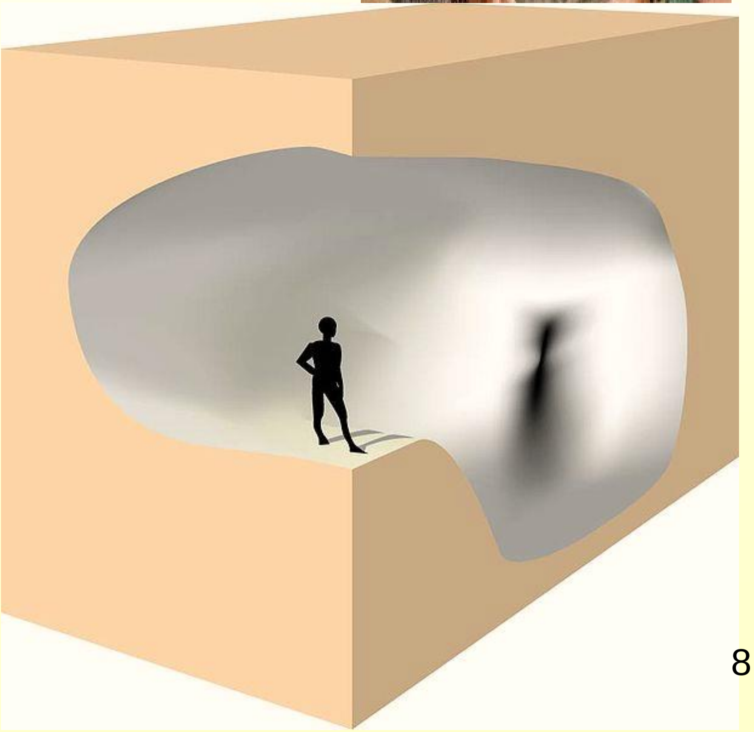

Joint Estimation of Functional Maps, IV

- Overall optimization

$$\min \sum_{(i,j)\in\mathcal{G}} w_{ij}\left(f_{ij}^{\text{feature}} + \mu f_{ij}^{\text{reg}} + \lambda f_{ij}^{\text{cons}}\right)$$
$$s.t. \quad Y^T Y = I_m$$

- Alternating optimization:
  - Fix Y, solve X ⟹ Independent QP problems

$$X_{ij}^{\star} = \arg\min_X \left(f_{ij}^{\text{feature}} + \mu f_{ij}^{\text{reg}} + \lambda f_{ij}^{\text{cons}}\right)$$

  - Fix X, solve Y ⟹ Eigenvalue problem

$$\min \quad \text{trace}(Y^T W Y) \qquad s.t. \quad Y^T Y = I_m$$

$$W_{ij} = \begin{cases} \sum_{(i,j')\in\mathcal{G}} w_{ij'}(I_m + X_{ij'}^T X_{ij'}) & i = j \\ -w_{ij}(X_{ji} + X_{ij}^T) & (i,j)\in\mathcal{G} \\ 0 & \text{otherwise.} \end{cases}$$

Consistency Matters

Source image

Target image

Without cycle consistency

With cycle consistency

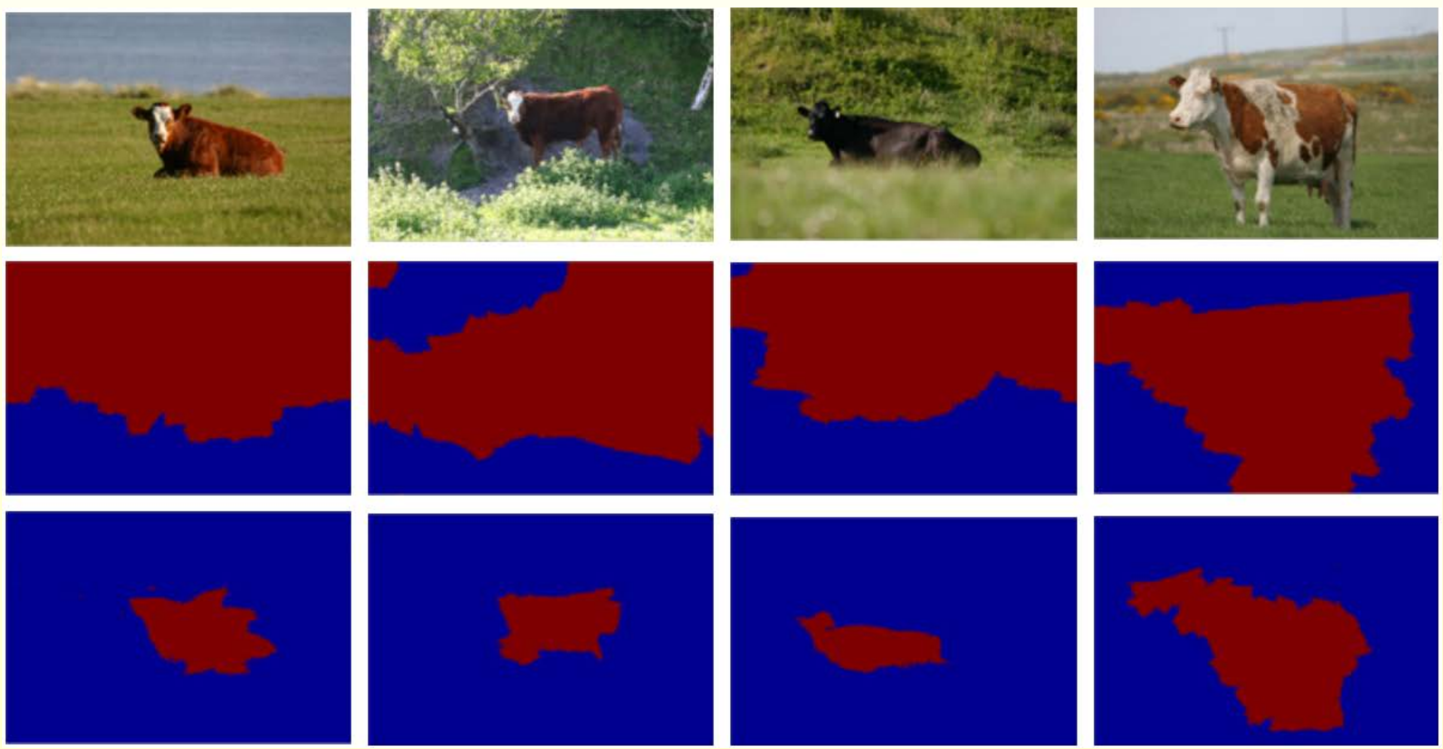

Generating Consistent Segmentations

- Two objectives for segmentation functions
  - consistent under functional map transportation

$$f^{\text{map}} = \sum_{(i,j)\in\mathcal{G}} w_{ij}\|X_{ij}\mathbf{f}_i - \mathbf{f}_j\|_{\mathcal{F}}^2$$

  - agreement with normalized cut scores:

We look for network fixed points!

$$f^{\text{seg}} = \sum_{i=1}^{N} \mathbf{f}_i^T B_i^T L_i B_i \mathbf{f}_i$$

Easy to incorporate labeled images with ground truth segmentation

- Joint optimization:

$$\min\ f^{\text{seg}} + \gamma f^{\text{map}} \quad s.t. \sum_{i=1}^{N}\|\mathbf{f}_i\|^2 = 1$$

Eigen-decomposition problem

Experiments

- iCoseg dataset
  - Very similar or the same object in each class;
  - 5~10 images per class.
- MSRC dataset
  - Similar objects in each class;
  - ~30 images per class.
- PASCAL data set
  - Retrieved from PASCAL VOC 2012 challenge;
  - All images with the same object label;
  - Larger scale;
  - Larger variability.

- iCoseg data set
- New unsupervised method
  - Mostly outperforms other unsupervised methods
  - Sometimes even outperforms supervised methods
  - Supervised input is easily added and further improves the results

Kuettel'12 (Supervised)		Unsupervised Fmaps
Image+transfer	Full model	
87.6	91.4	90.5

Supervised method

Class	Joulin '10	Rubio '12	Vicente '11	Fmaps -uns
Alaska Bear	74.8	86.4	90.0	**90.4**
Red Sox Players	73.0	90.5	90.9	**94.2**
Stonehenge1	56.6	87.3	63.3	**92.5**
Stonehenge2	86.0	88.4	88.8	87.2
Liverpool FC	76.4	82.6	87.5	**89.4**
Ferrari	85.0	84.3	89.9	**95.6**
Taj Mahal	73.7	88.7	91.1	**92.6**
Elephants	70.1	75.0	43.1	**86.7**
Pandas	84.0	60.0	92.7	88.6
Kite	87.0	89.8	90.3	**93.9**
Kite panda	73.2	78.3	90.2	**93.1**
Gymnastics	90.9	87.1	91.7	90.4
Skating	82.1	76.8	77.5	78.7
Hot Balloons	85.2	89.0	90.1	**90.4**
Liberty Statue	90.6	91.6	93.8	**96.8**
Brown Bear	74.0	80.4	95.3	88.1
Average	78.9	83.5	85.4	**90.5**

MSRC

Unsupervised performance comparison

Class	N	Joulin '10	Rubio '12	Fmaps -uns
Cow	30	81.6	80.1	**89.7**
Plane	30	73.8	77.0	**87.3**
Face	30	84.3	76.3	**89.3**
Cat	24	74.4	77.1	**88.3**
Car(front)	6	87.6	65.9	87.3
Car(back)	6	85.1	52.4	**92.7**
Bike	30	63.3	62.4	**74.8**

Supervised performance comparison

Class	Vicente '11	Kuettel '12	Fmaps -s
Cow	94.2	92.5	**94.3**
Plane	83.0	86.5	**91.0**
Car	79.6	88.8	83.1
Sheep	94.0	91.8	**95.6**
Bird	95.3	93.4	**95.8**
Cat	92.3	92.6	**94.5**
Dog	93.0	87.8	91.3

PASCAL

Class	N	L	Kuettel '12	Fmaps -s	Fmaps -uns
Plane	178	88	90.7	**92.1**	89.4
Bus	152	78	81.6	**87.1**	80.7
Car	255	128	76.1	**90.9**	82.3
Cat	250	131	77.7	**85.5**	82.5
Cow	135	64	82.5	**87.7**	85.5
Dog	249	121	81.9	**88.5**	84.2
Horse	147	68	83.1	**88.9**	87.0
Sheep	120	63	83.9	**89.6**	86.5

- New method mostly outperforms the state-of-the-art techniques in both supervised and unsupervised settings

iCoseg: 5 images per class are shown

iCoseg: 5 images per class are shown

iCoseg: 5 images per class are shown

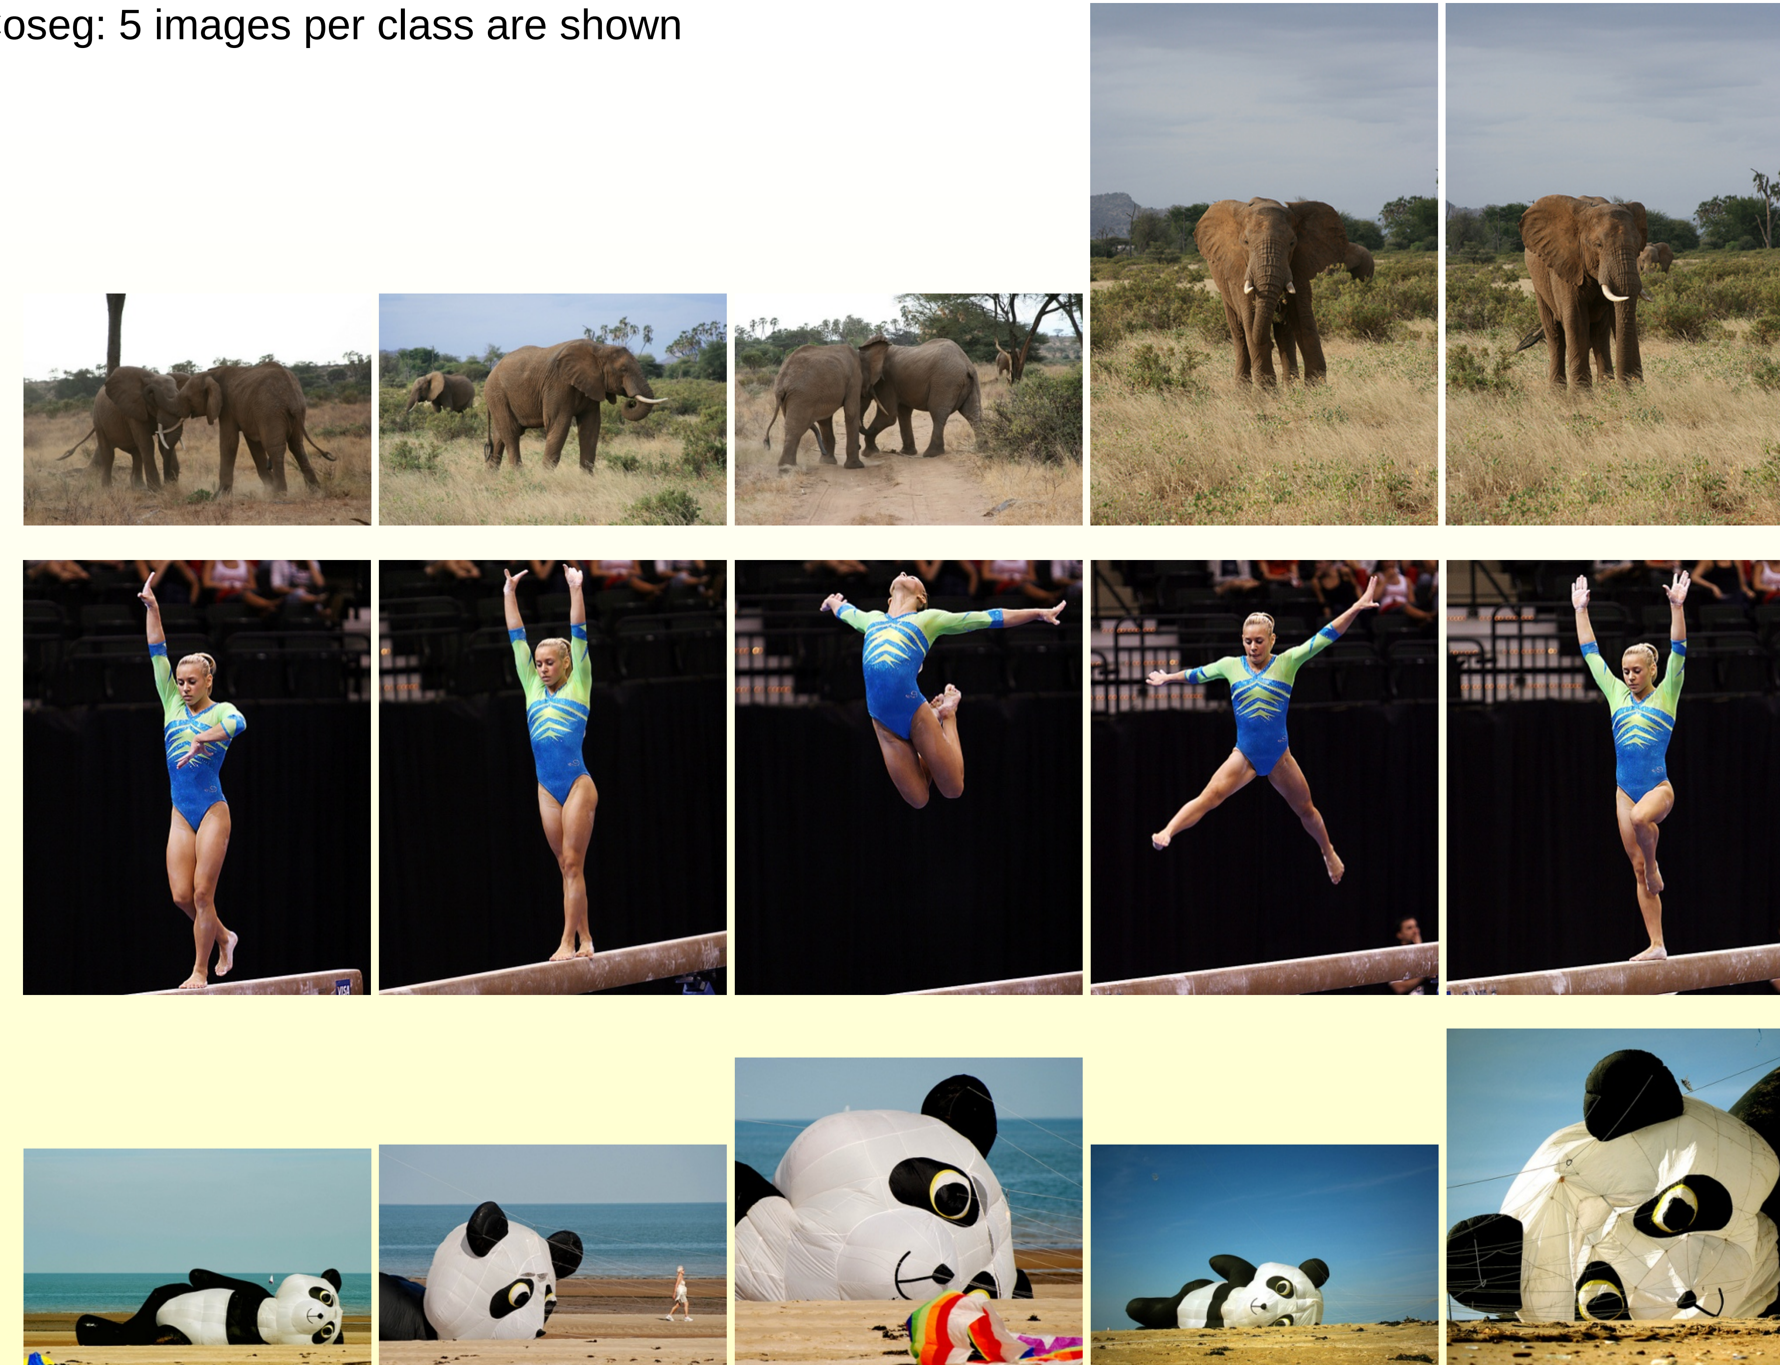

iCoseg: 5 images per class are shown

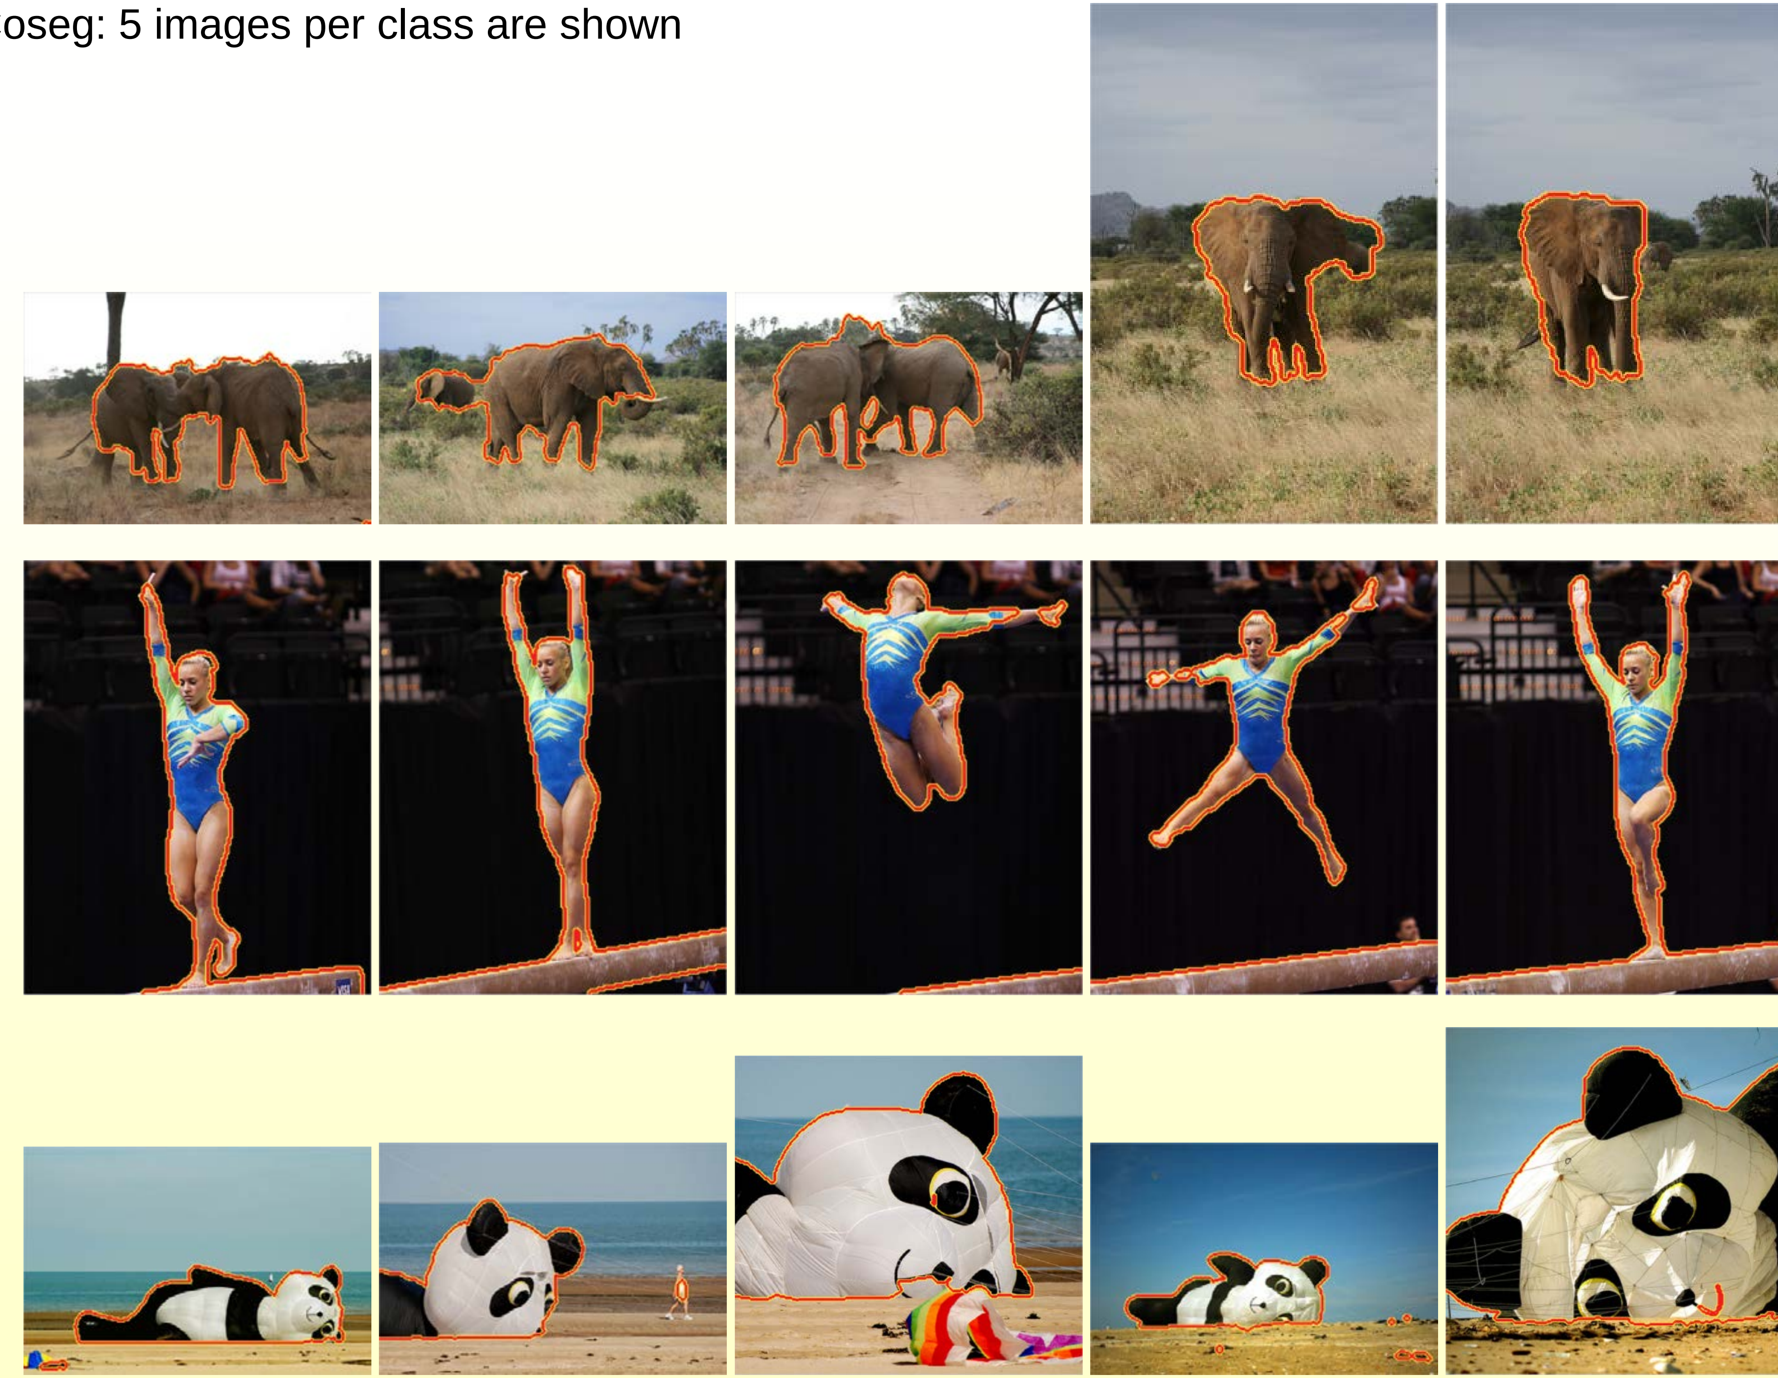

MSRC: 5 images per class are shown

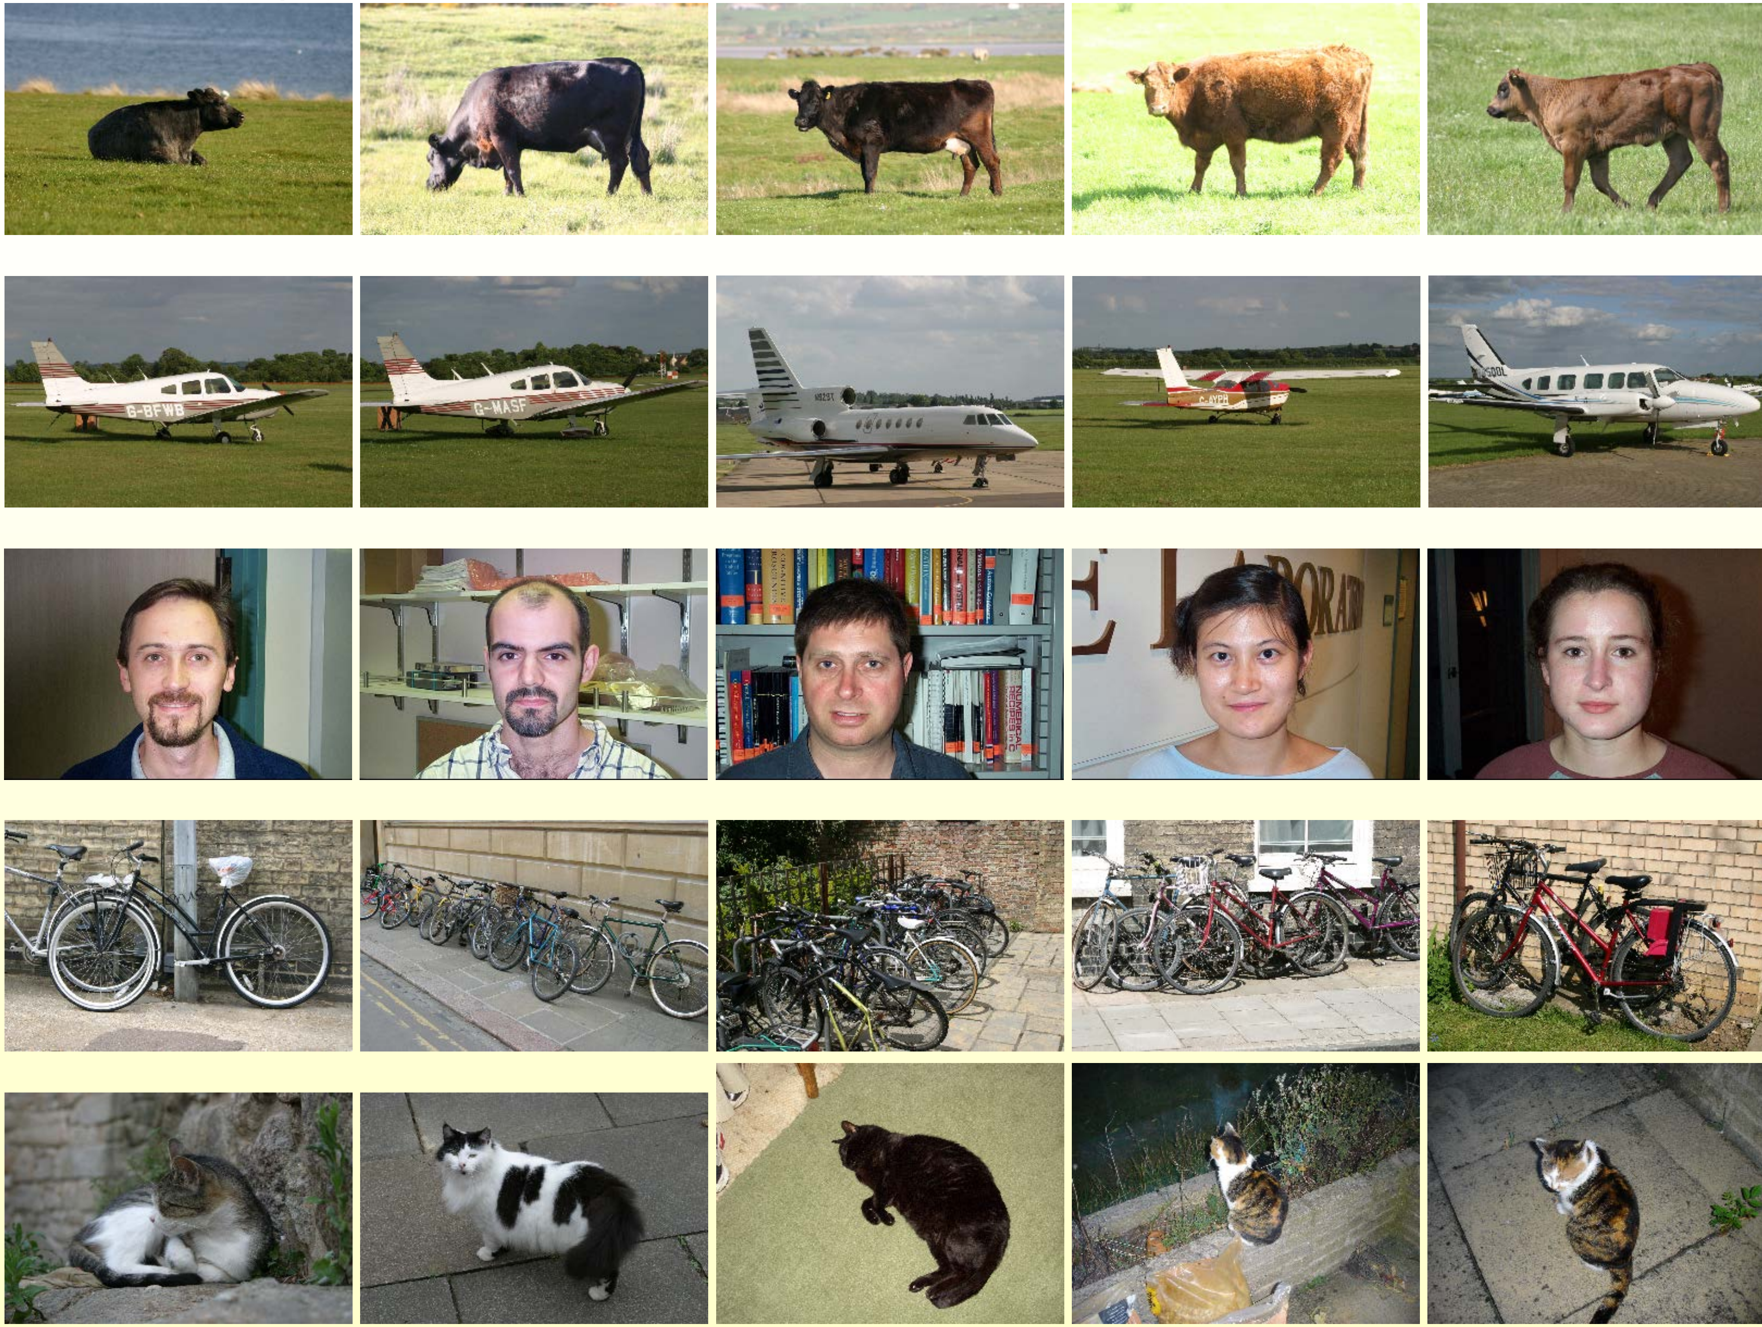

MSRC: 5 images per class are shown

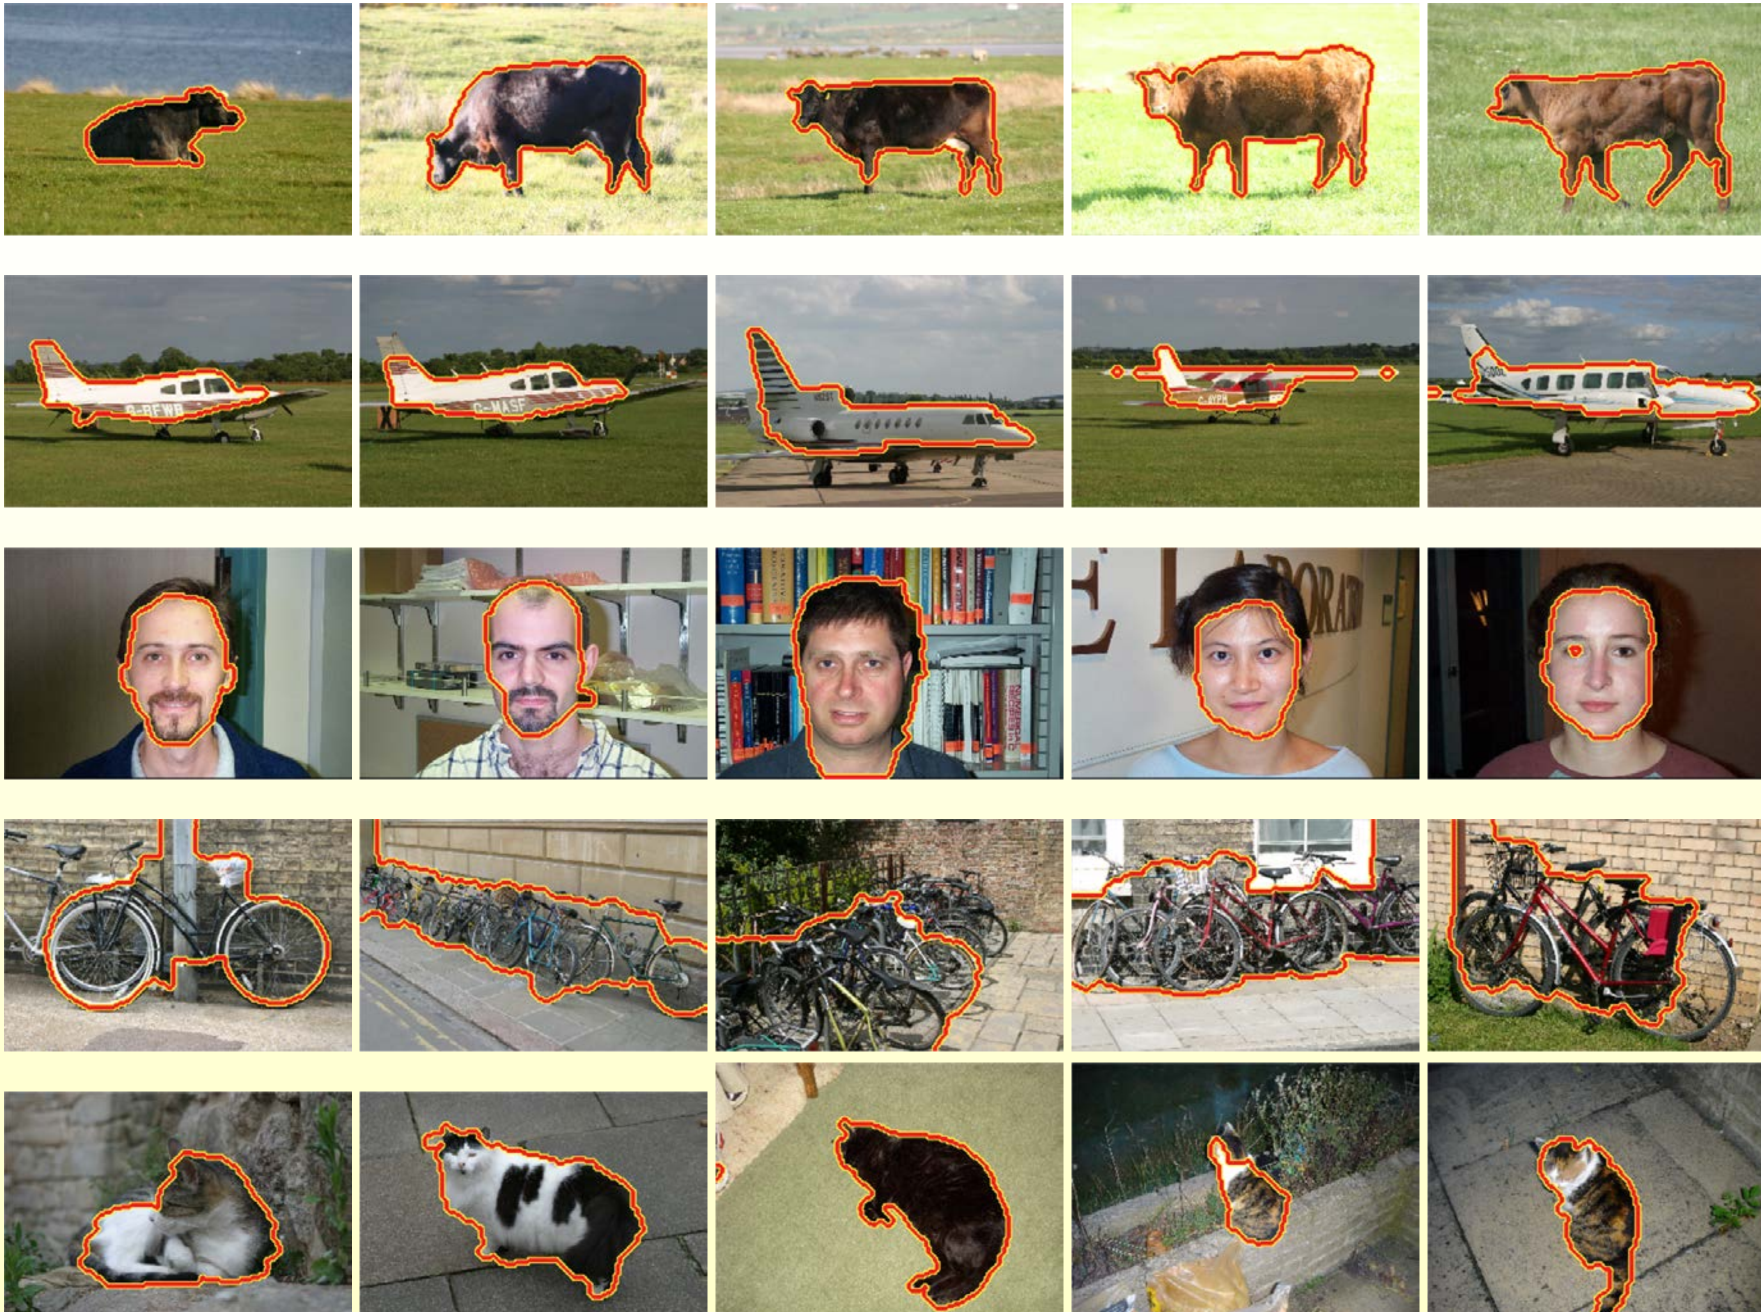

PASCAL: 10 images per class are shown

PASCAL: 10 images per class are shown

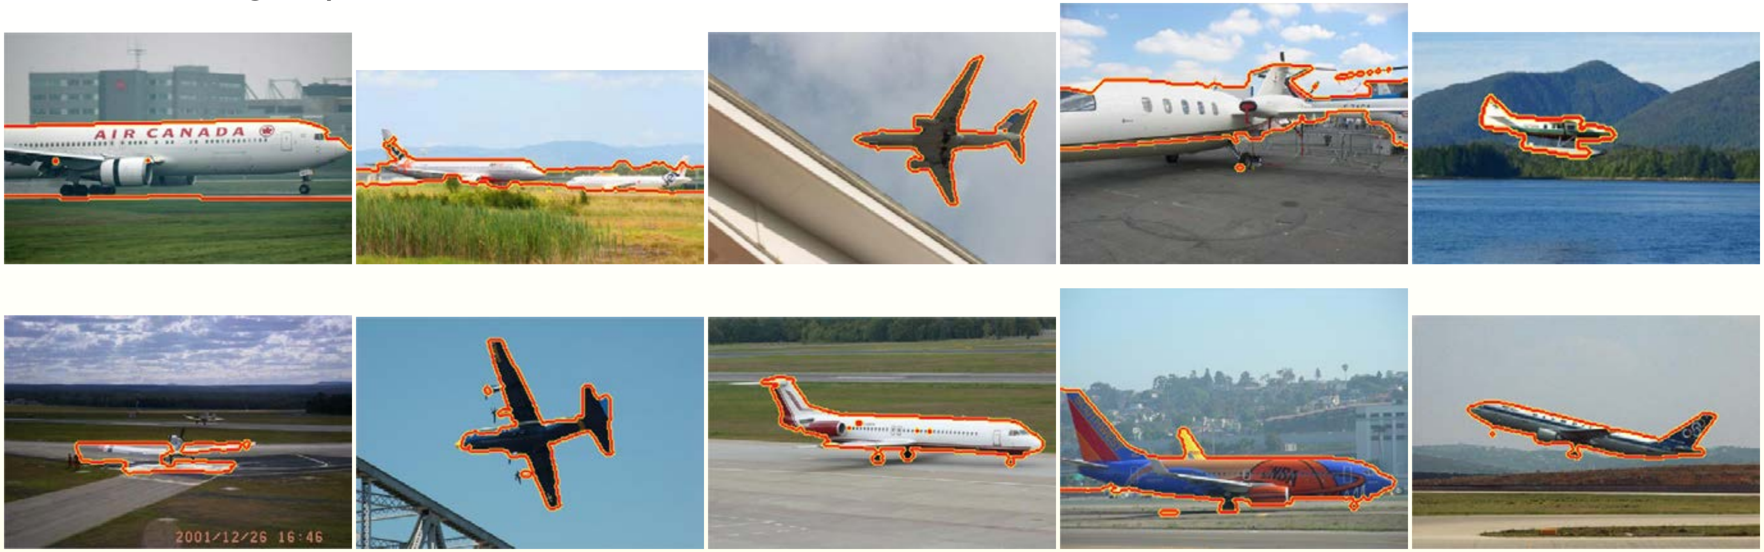

PASCAL: 10 images per class are shown

PASCAL: 10 images per class are shown

The End