CS233: Geometric and Topological Data Analysis

Functional Maps and Applications, Map Visualization

30 May 2018

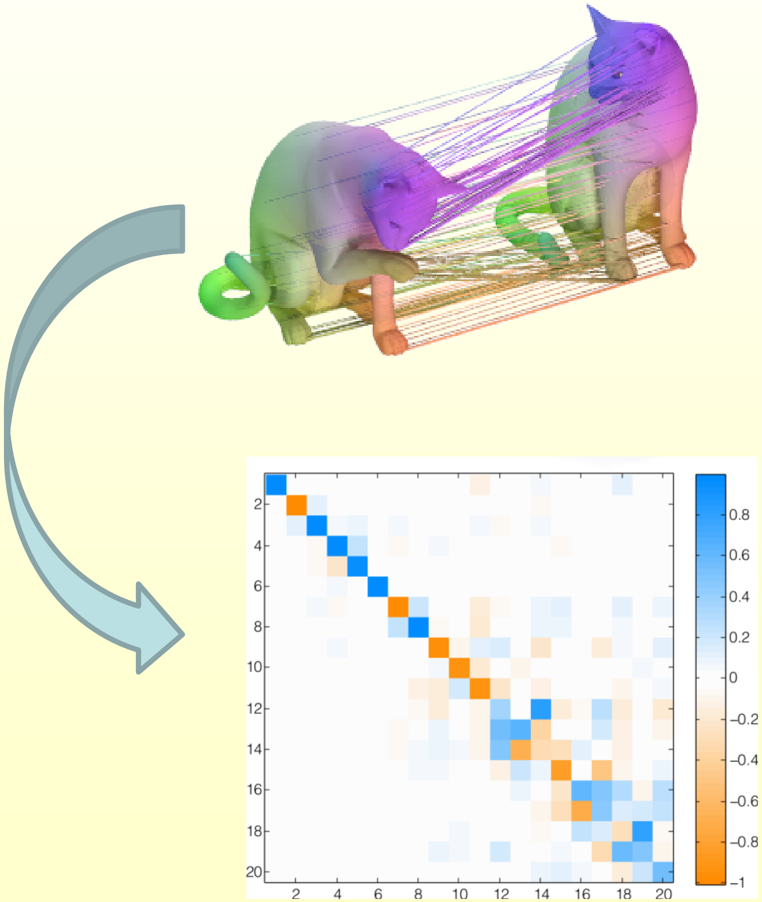

Slides ack: Maks Ovsjanikov, Mirela Ben Chen, Justin Solomon

Last Lecture: Deep Learning on Graph Dara

Graph CNNs

Input: Feature matrix $\mathbf{X} \in \mathbb{R}^{N \times E}$, preprocessed adjacency matrix $\hat{\mathbf{A}}$

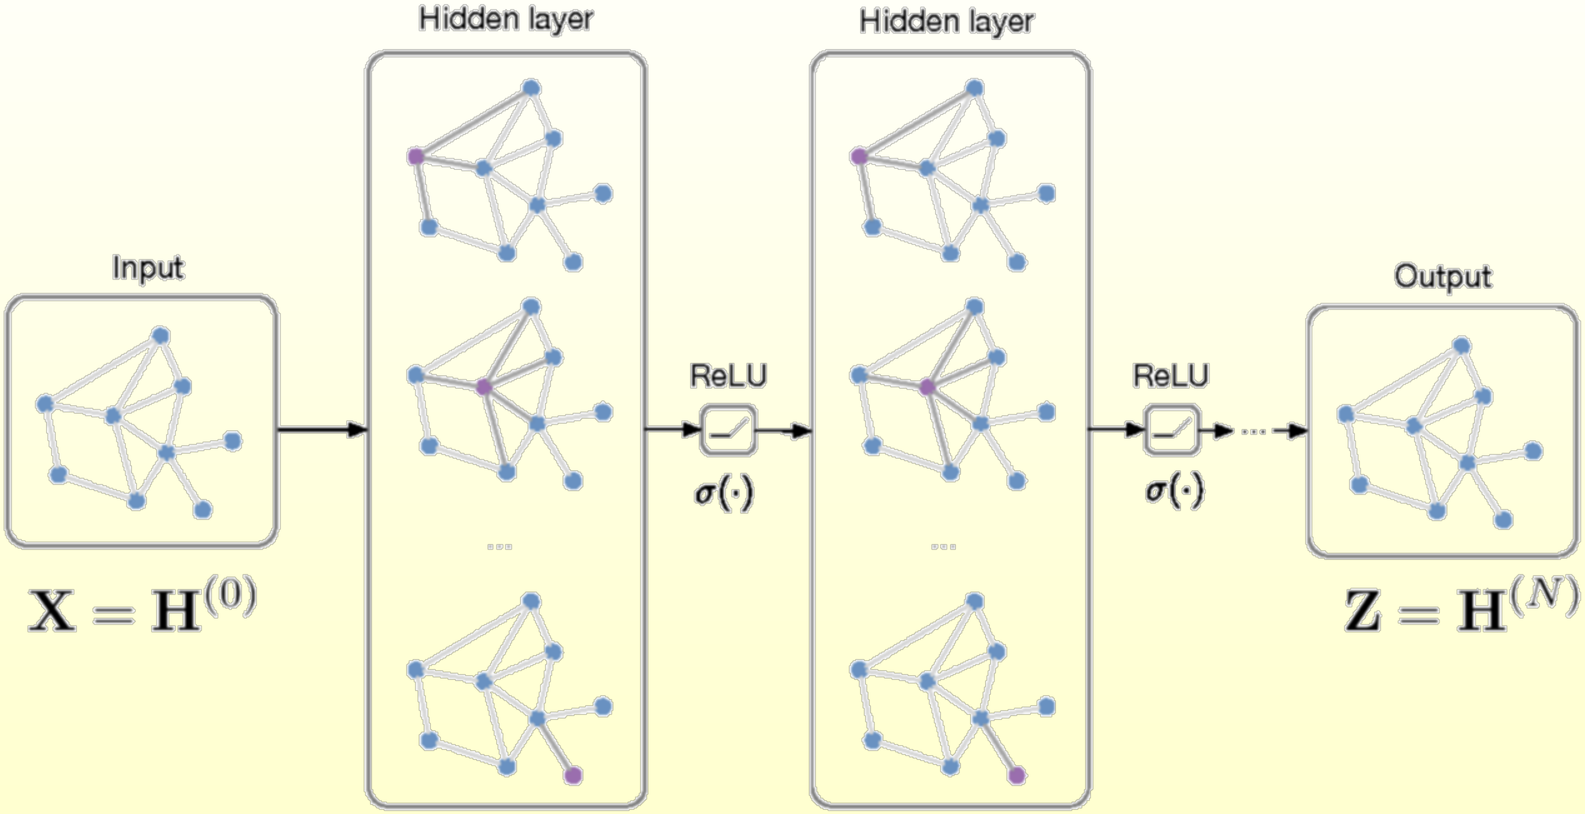

$$\mathbf{X} = \mathbf{H}^{(0)}$$

$$\mathbf{Z} = \mathbf{H}^{(N)}$$

$$\mathbf{H}^{(l+1)} = \sigma\left(\hat{\mathbf{A}}\mathbf{H}^{(l)}\mathbf{W}^{(l)}\right)$$

(or more sophisticated filters / basis functions)

[Kipf & Welling, ICLR 2017]

Geodesic CNNs

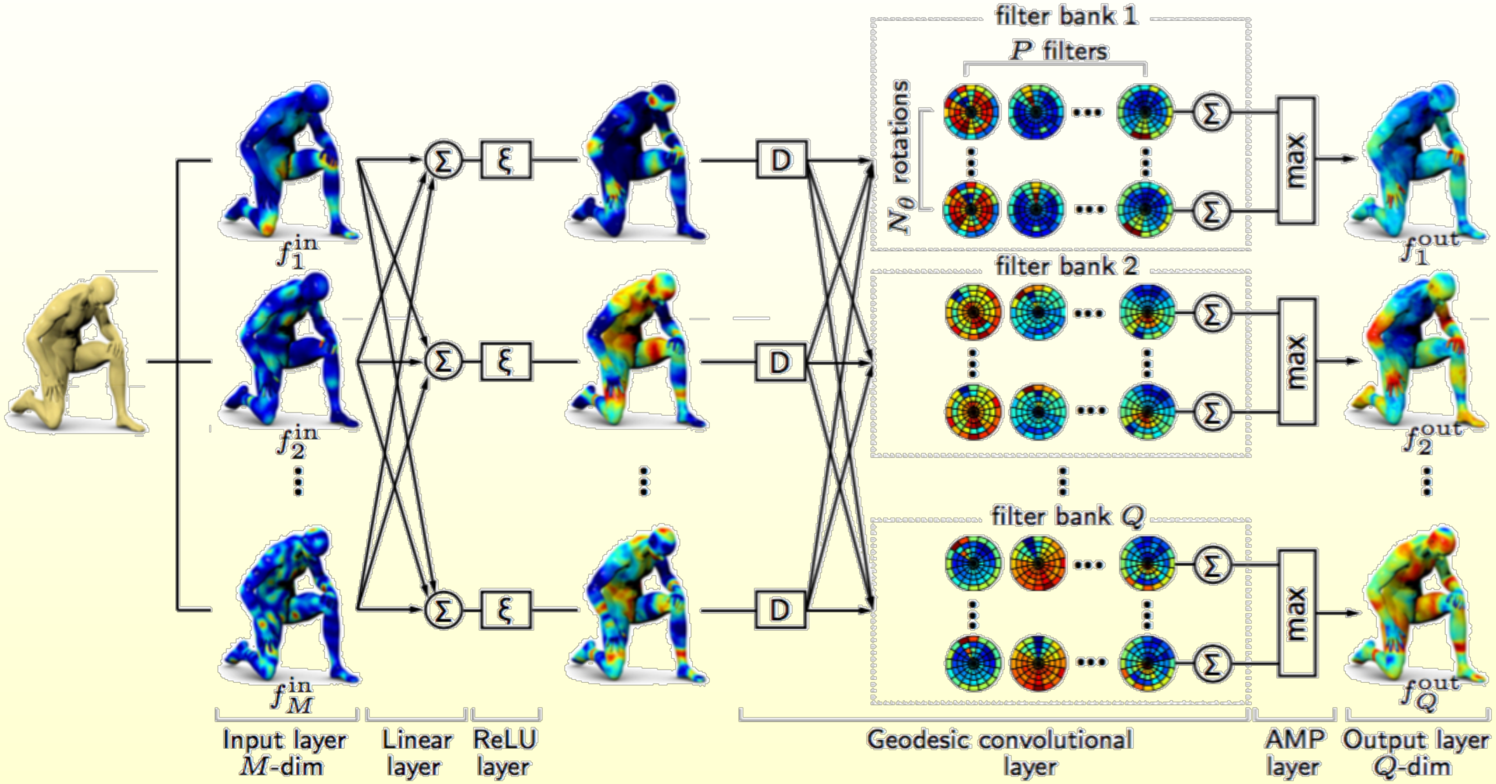

Jonathan Masci et al 2015

Spectral CNN

Convolutional layer expressed in the spectral domain

$$\mathbf{g}_l = \xi\left(\sum_{l'=1}^{p} \mathbf{\Phi}\hat{\mathbf{W}}_{l,l'}\mathbf{\Phi}^\top \mathbf{f}_{l'}\right) \qquad \begin{array}{l} l = 1,\ldots,q \\ l' = 1,\ldots,p \end{array}$$

where $\hat{\mathbf{W}}_{l,l} = n \times n$ diagonal matrix of filter coefficients

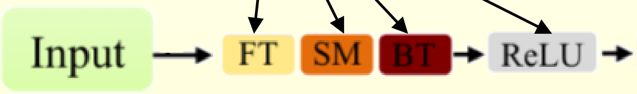

Bruna et al. 2014

Synchronized Spectral CNN

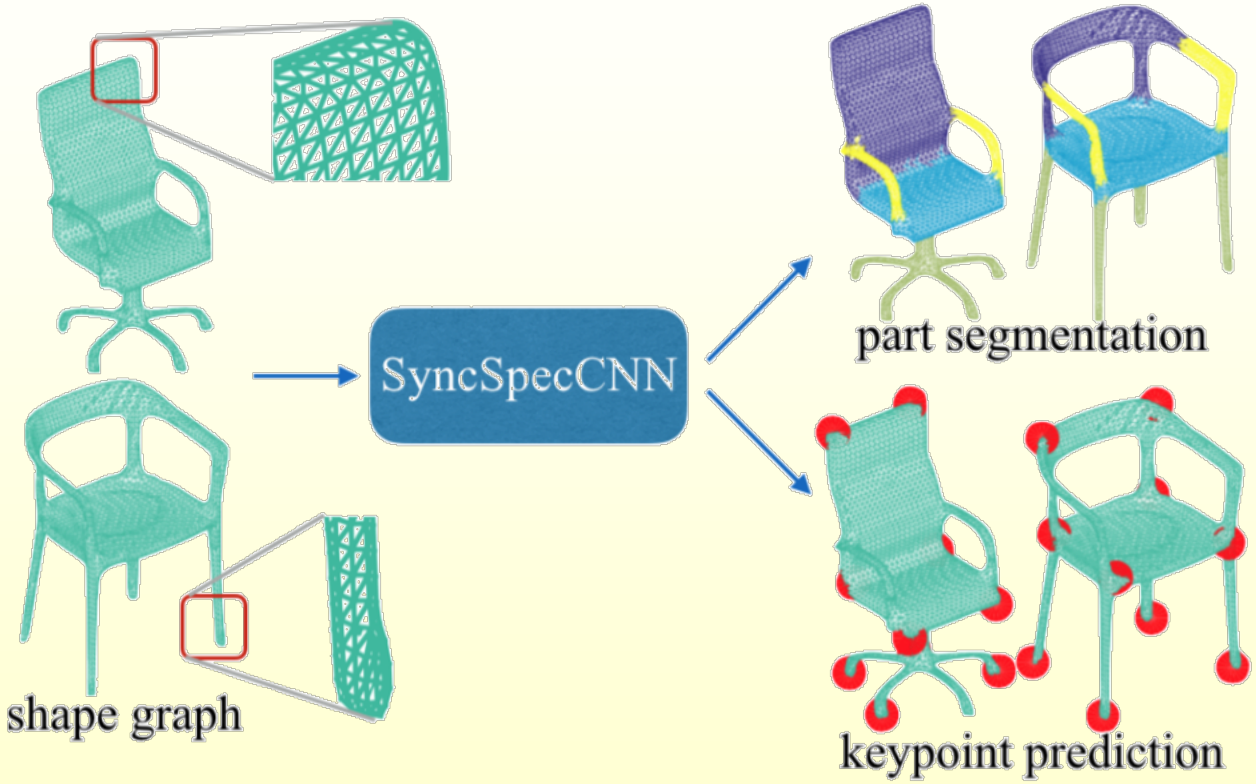

Input: shape graph
equipped with vertex functions

Output: vertex labels

Vertical and Horizontal Networks

Vertical Learning Networks

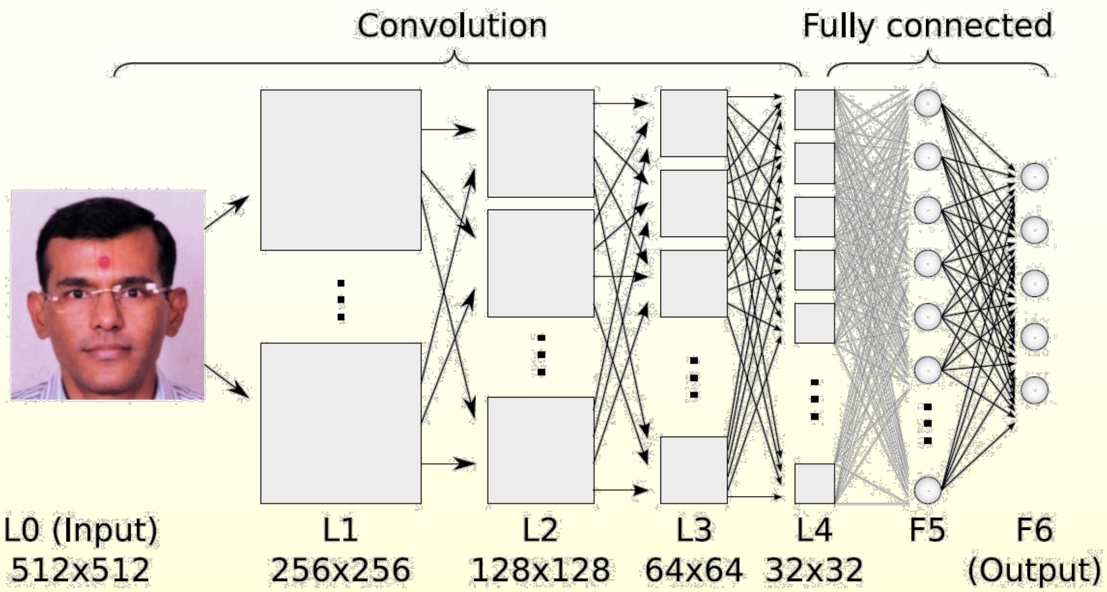

[Makwana, 2016]

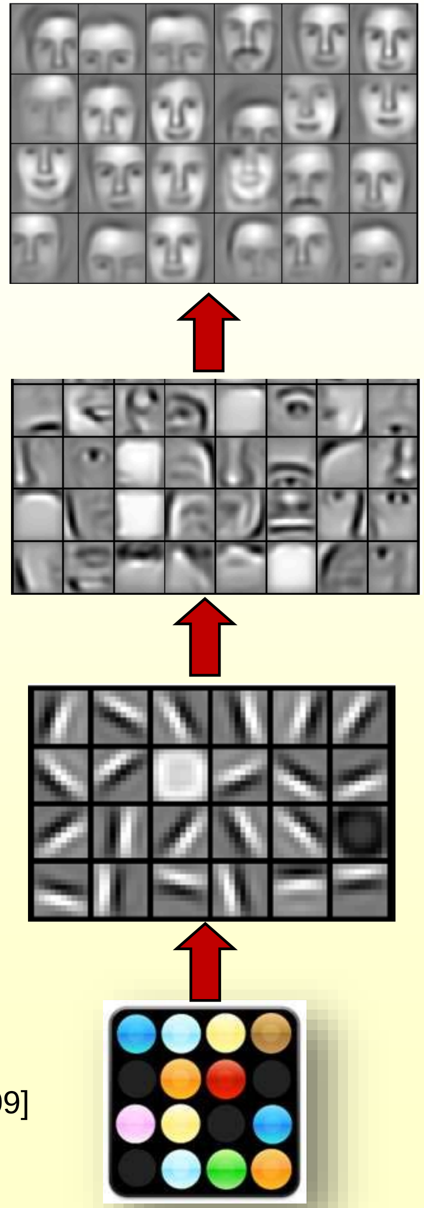

[Lee et al., 2009]

Data-driven feature learning at ascending abstraction layers

"Deep" nets

Horizontal Networks

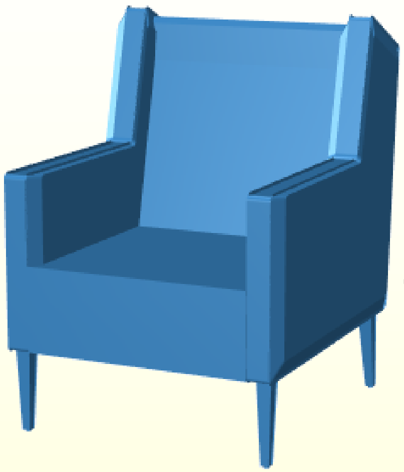

Similarity as a communications channel

34 *The Mathematical Theory of Communication*

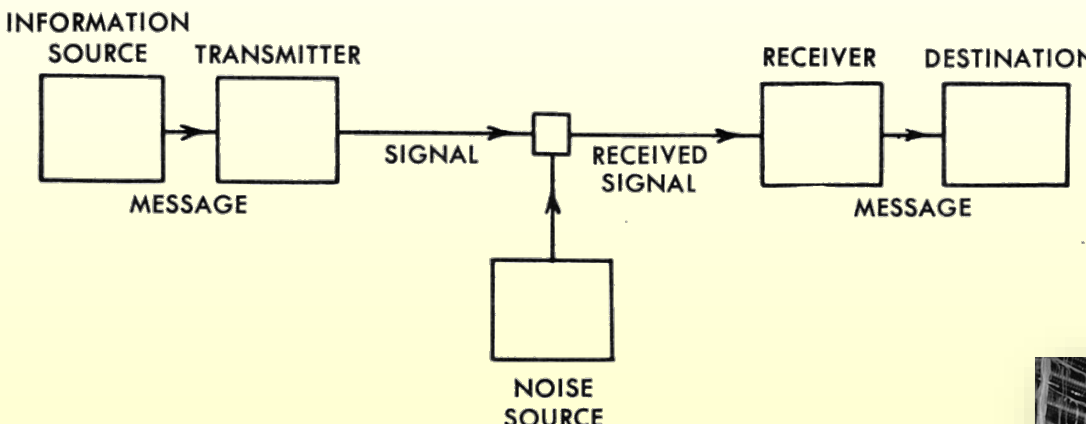

Fig. 1. — Schematic diagram of a general communication system.

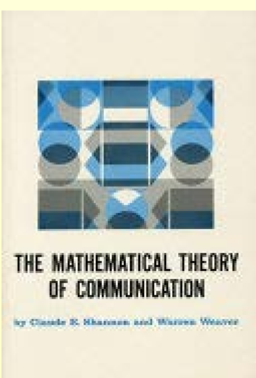

Claude Shannon

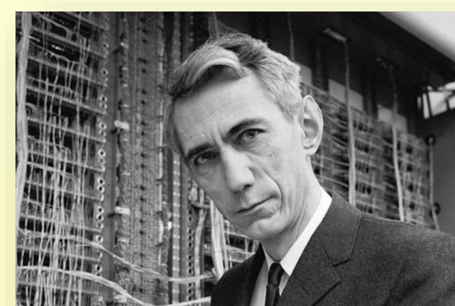

Networks of Images

Or of Shapes, Or of Both

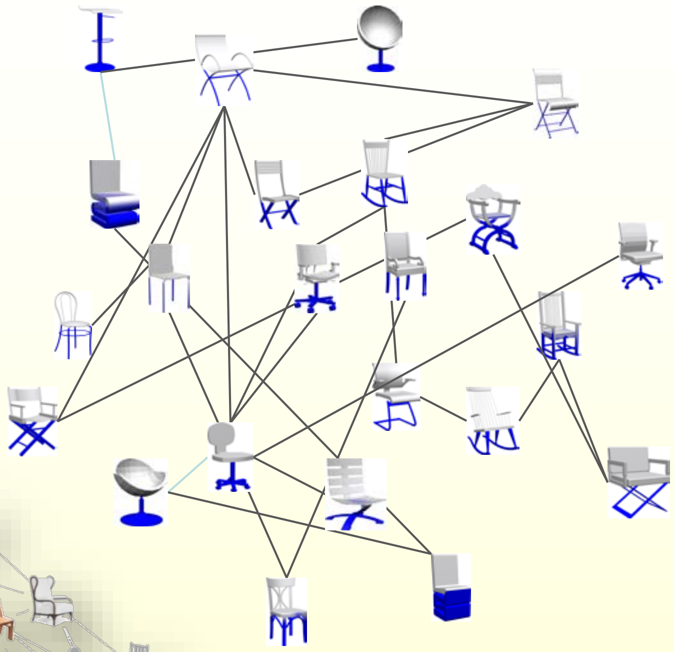

Matching Enables Information Transport

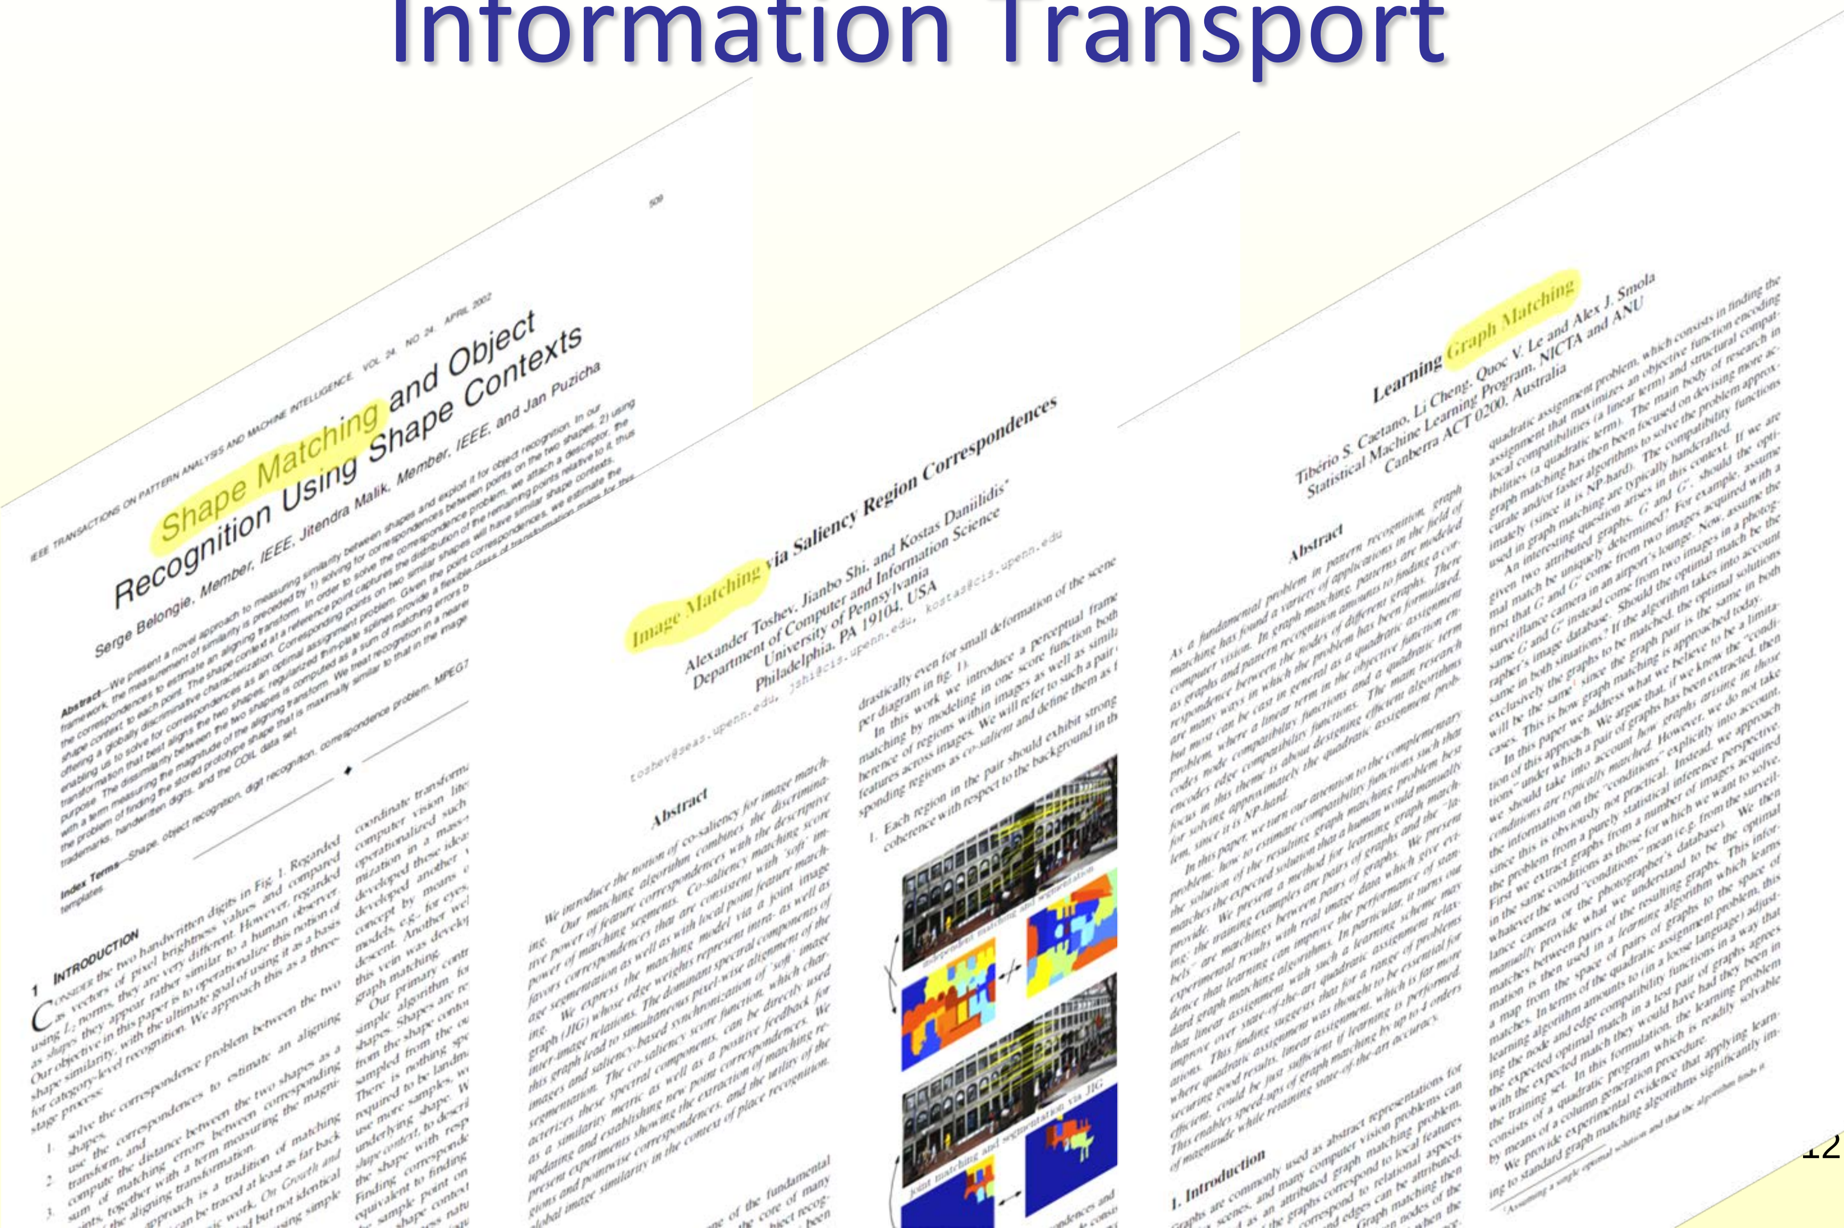

Good Correspondences or Maps are Information Transporters

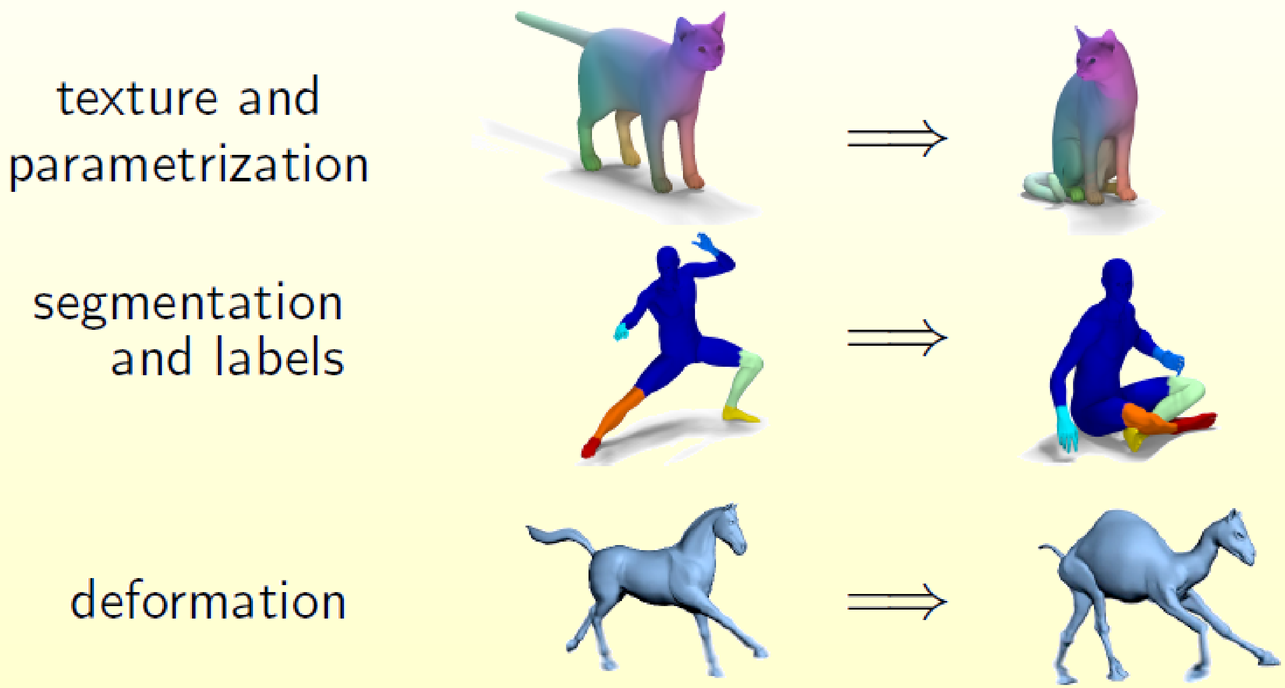

Maps vs. Distances/Similarities
Networks vs. Graphs

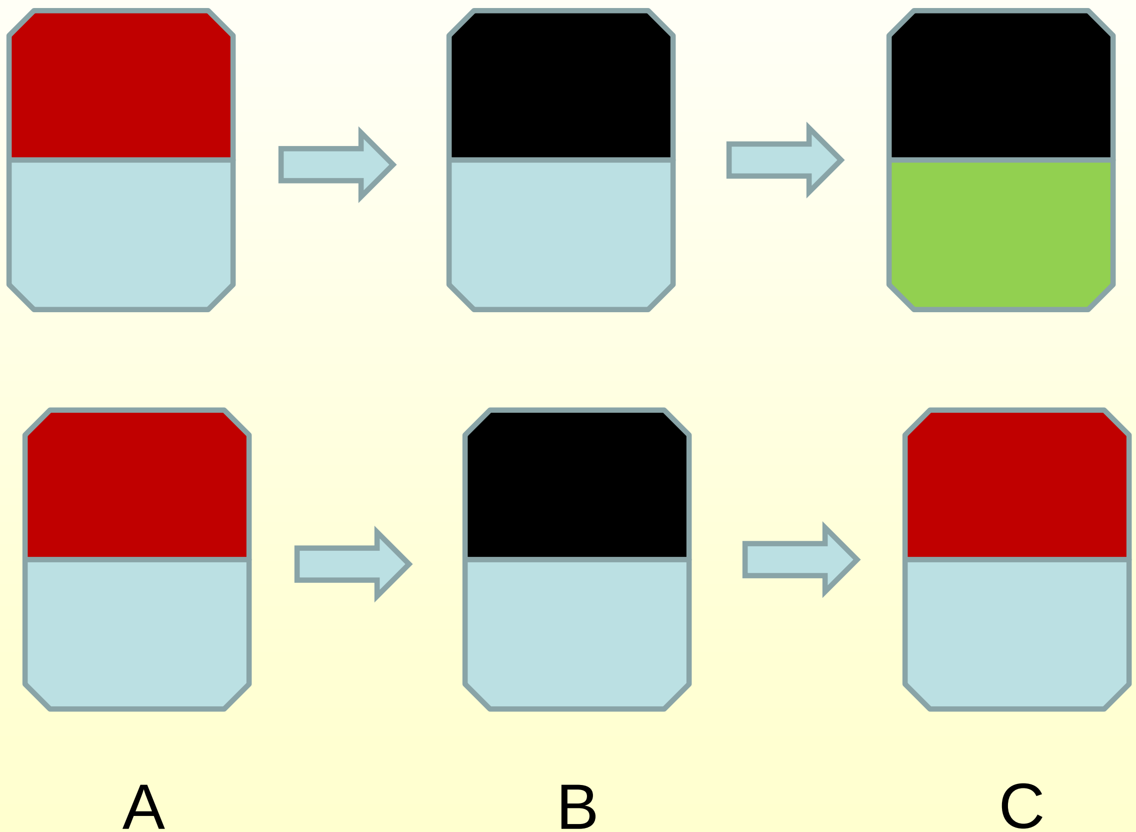

Persistence of correspondences

Each Data Set Is Not Alone

- The interpretation of a particular piece of geometric data is deeply influenced by our interpretation of other related data

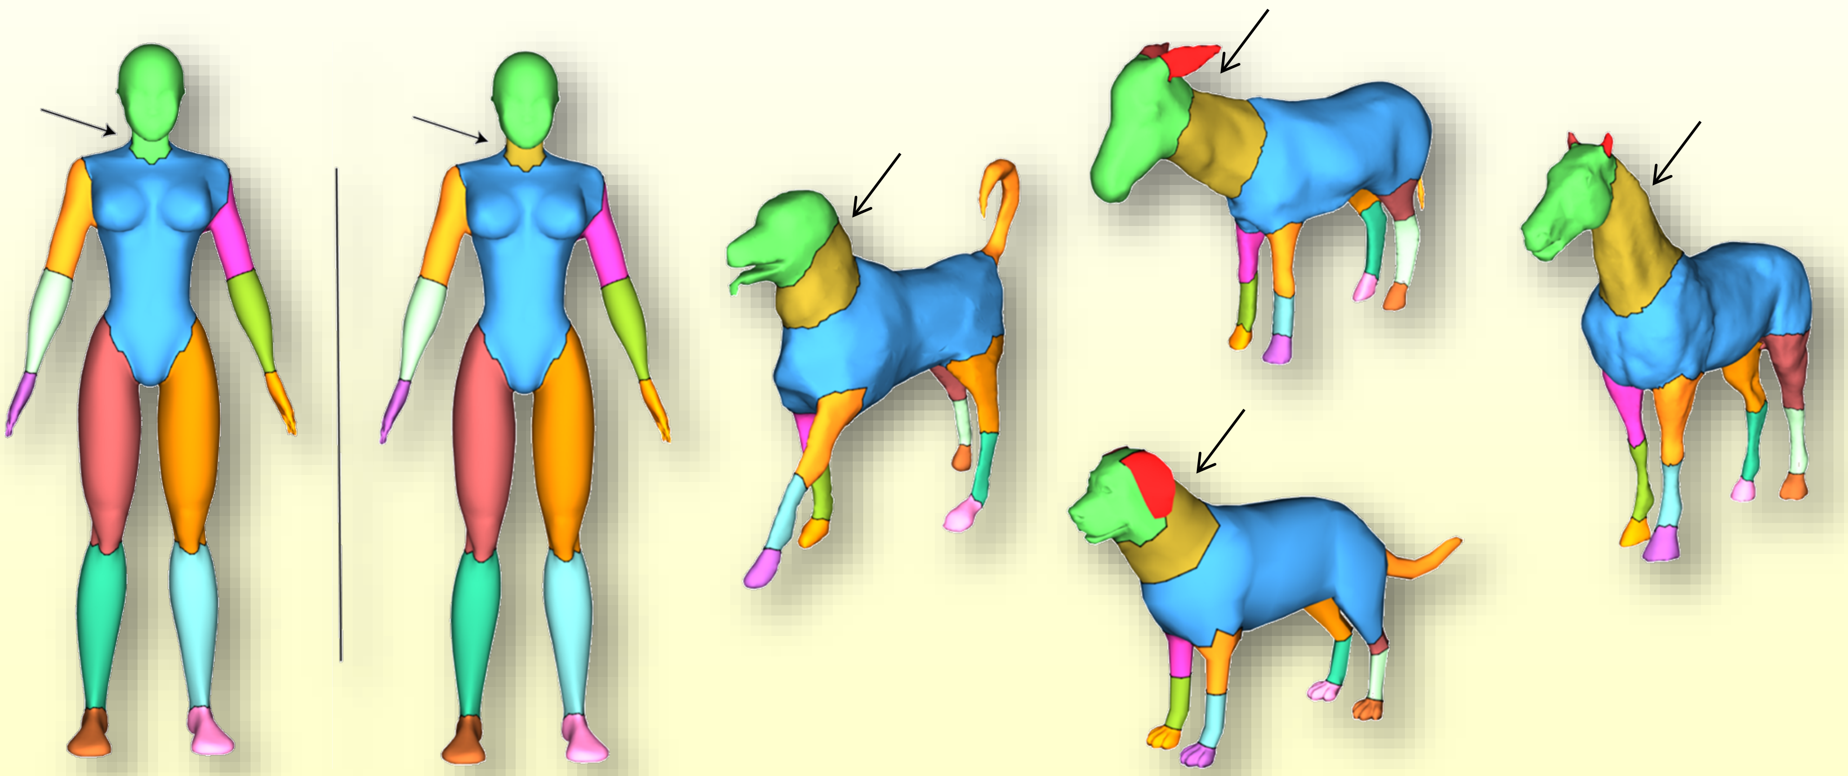

3D Segmentation

Societies, or
Social Networks of Data Sets

Our understanding of data can greatly benefit from extracting these relations and building relational networks.

We can exploit the relational network to

- transport information around the network
- assess the validity of operations or interpretations of data (by checking consistency against related data)
- assess the quality of the relations themselves (by checking consistency against other relations through cycle closure, etc.)
- extract shared structure among the data

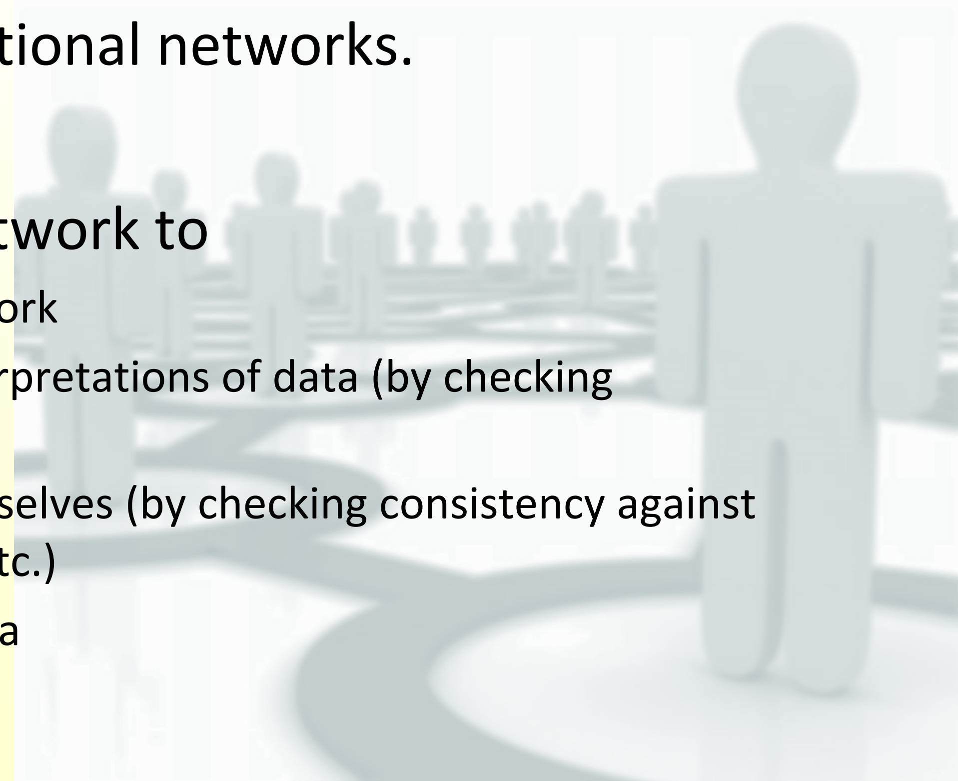

Thus the network becomes the great regularizer in joint data analysis.

V+H: Functorial Data Analysis

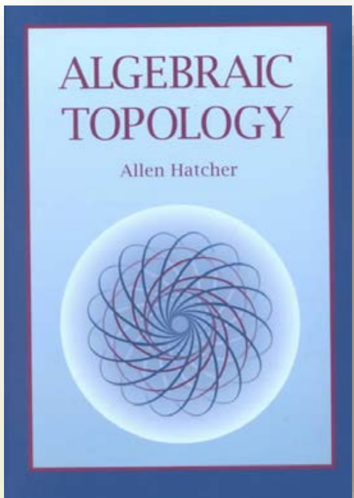

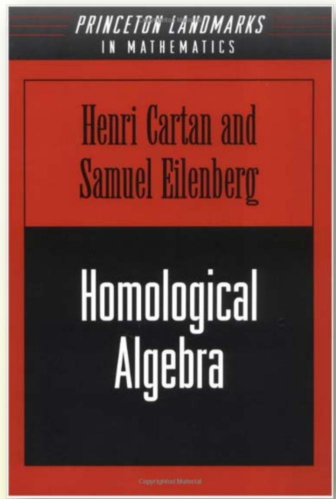

$$
\begin{array}{ccc}
H_*(X) & \xrightarrow{\phi} & H_*(Y) \\
{\scriptstyle H_*}\uparrow & & \uparrow{\scriptstyle H_*} \\
X & \xrightarrow[f]{} & Y
\end{array}
$$

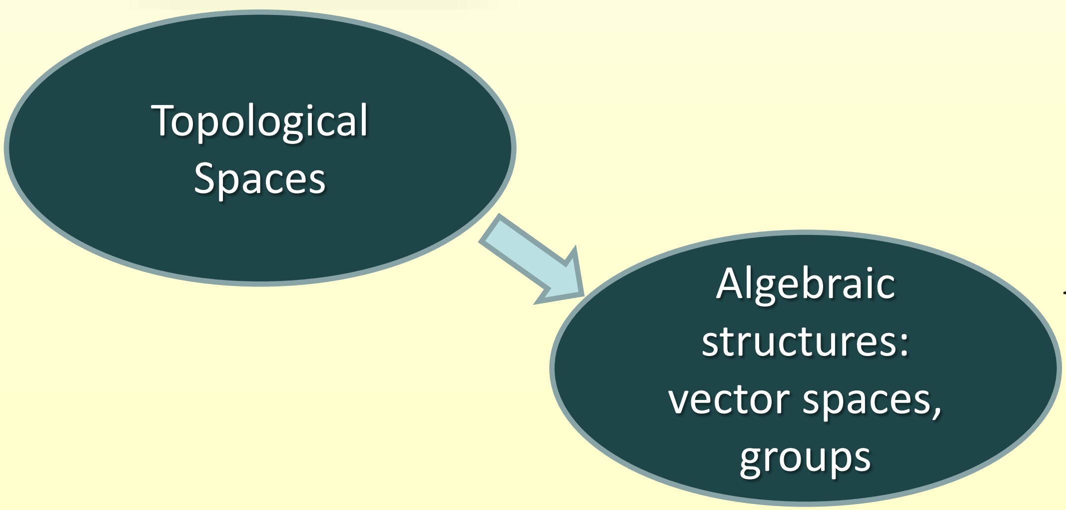

$$
\begin{array}{ccc}
L(X) & \xrightarrow{\phi} & L(Y) \\
{\scriptstyle DN}\uparrow & & \uparrow{\scriptstyle DN} \\
X & \xrightarrow[f]{} & Y
\end{array}
$$

Maps and Correspondences

Maps

$$\phi : X \to Y$$

Map from *X* to *Y*

Maps and Correspondences

- Multiscale mappings
  - Point/pixel level
  - part level

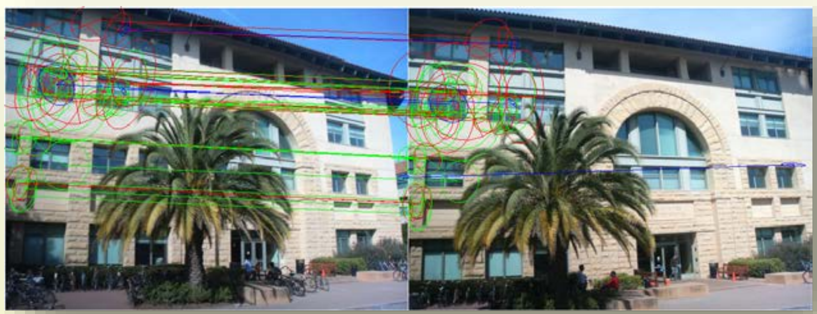

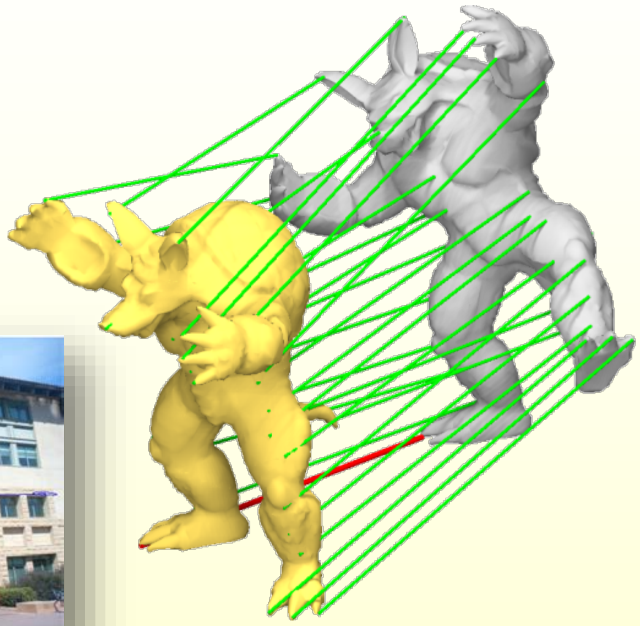

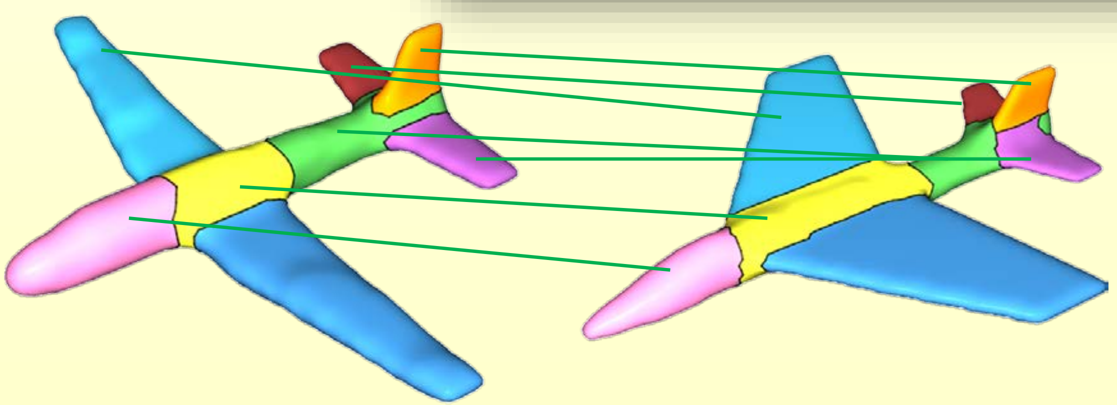

Maps capture what is the same or similar across two data sets

Algebraic Structure: Map Composition

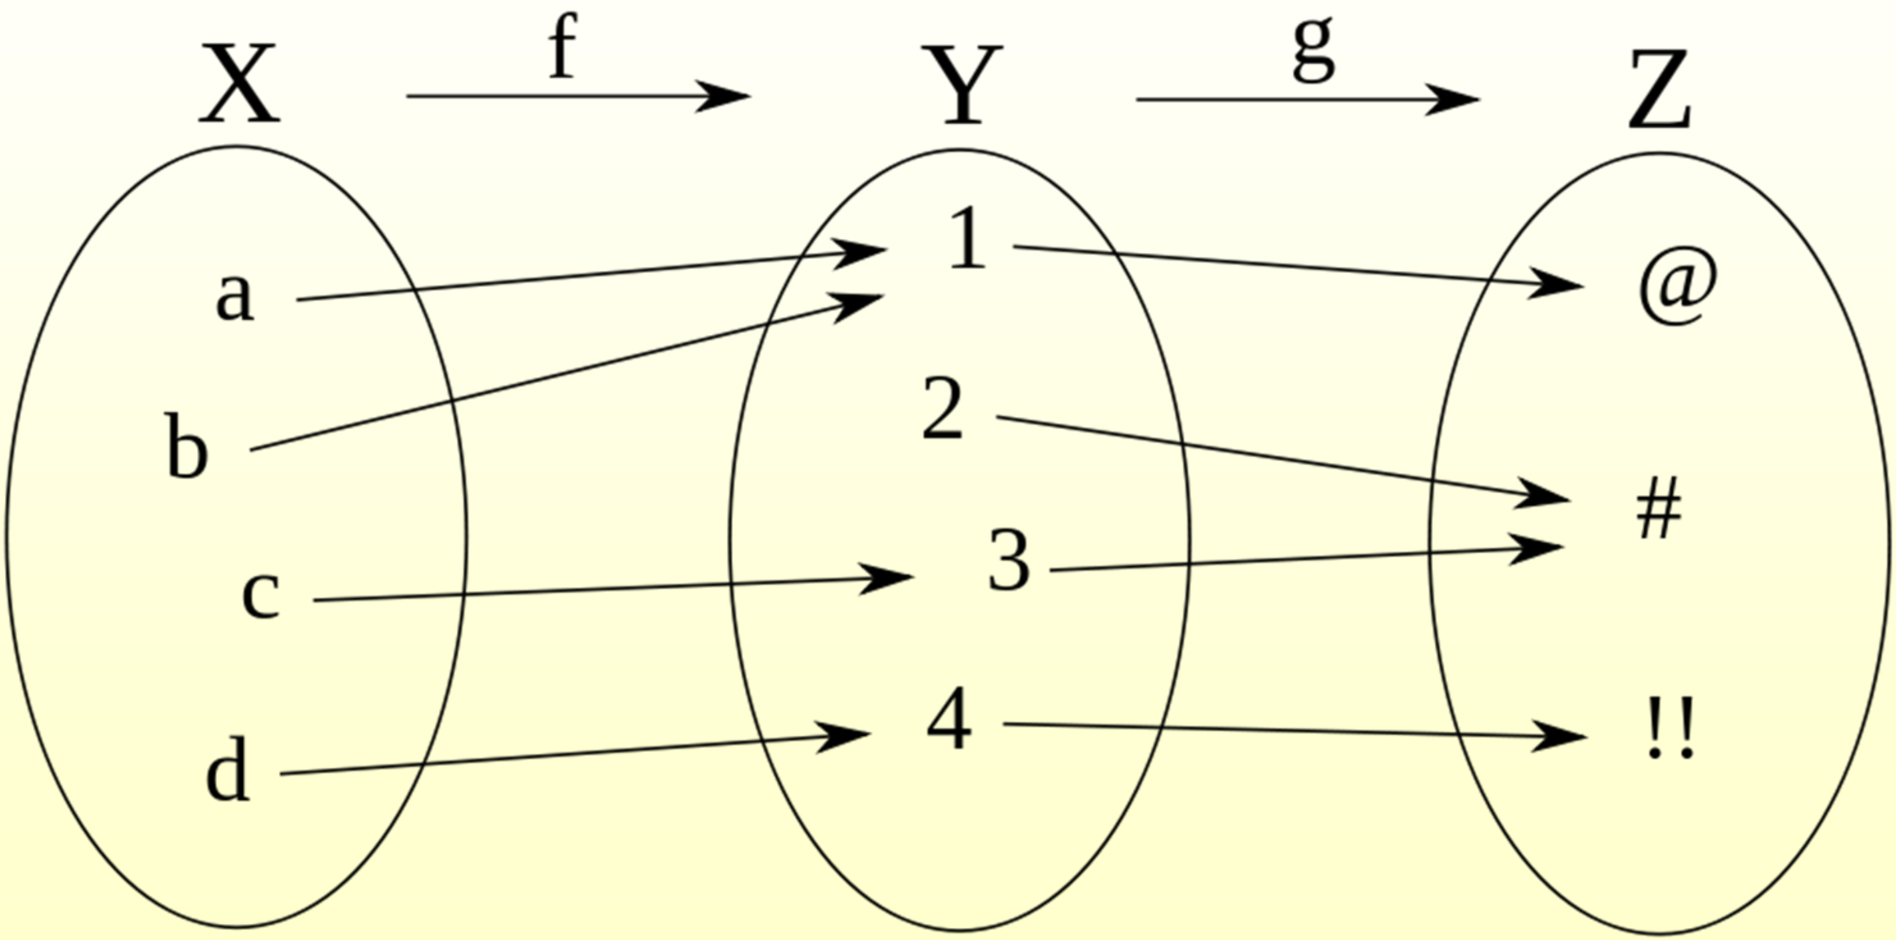

Problems and Issues

Symmetry, ambiguity, scale, bad data

Non-Convex, Combinatorial Optimization

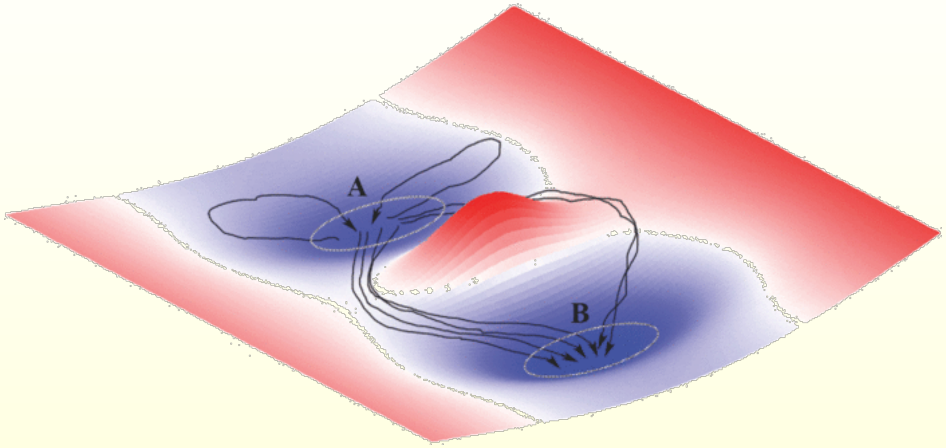

multiple minima

NP-hard quadratic assignments

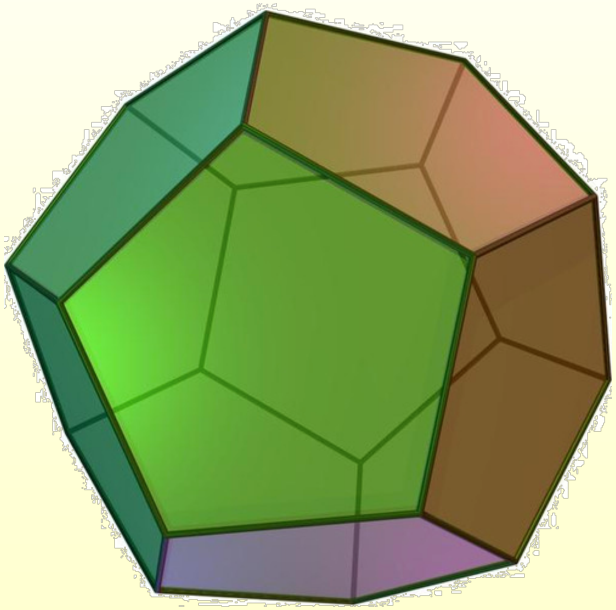

n! permutations

Symmetry, ambiguity, scale, bad data

A Potential Way Out

Find alternative representation more amenable to optimization

Redefine the notion of map

Function Spaces and Functional Maps

Knowledge as Functions

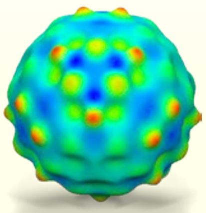

Curvature

Parts

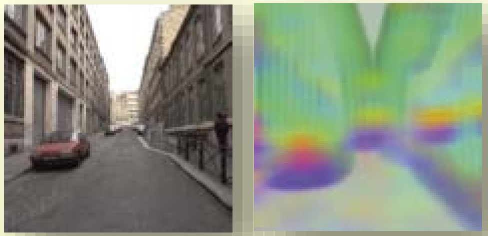

SIFT flow, C. Liu 2011

Knowledge towers over visual data: function spaces

A Dual View: Functions and Operators

- Functions on data
  - Properties, attributes, descriptors, part indicators, etc.
  - But also beliefs, opinions, etc
- Operators on functions
  - Maps of functions to functions
    - Laplace-Beltrami operator on a manifold M

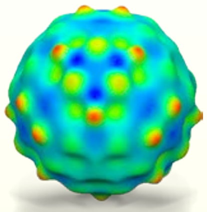

Curvature

Parts

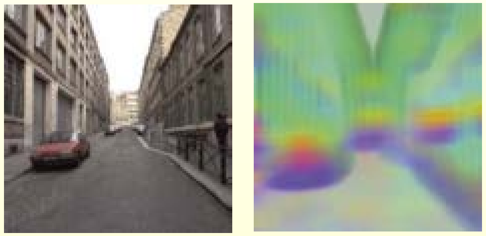

SIFT flow, C. Liu 2011

$$\Delta : C^\infty(M) \to C^\infty(M), \Delta f = \text{div } \nabla f$$

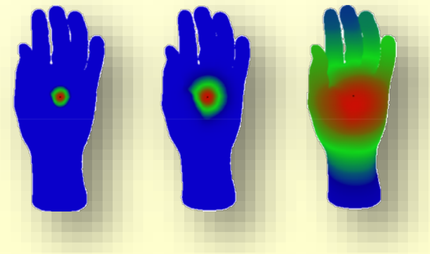

$$\frac{\partial u}{\partial t} = -\Delta u$$

heat diffusion

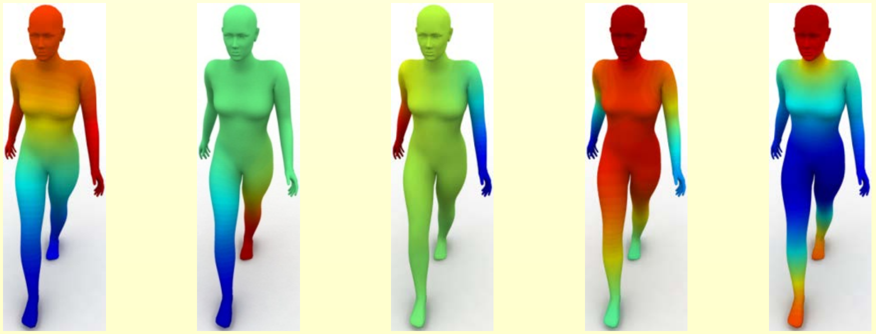

Laplace Beltrami eigenfunctions

Functional Maps

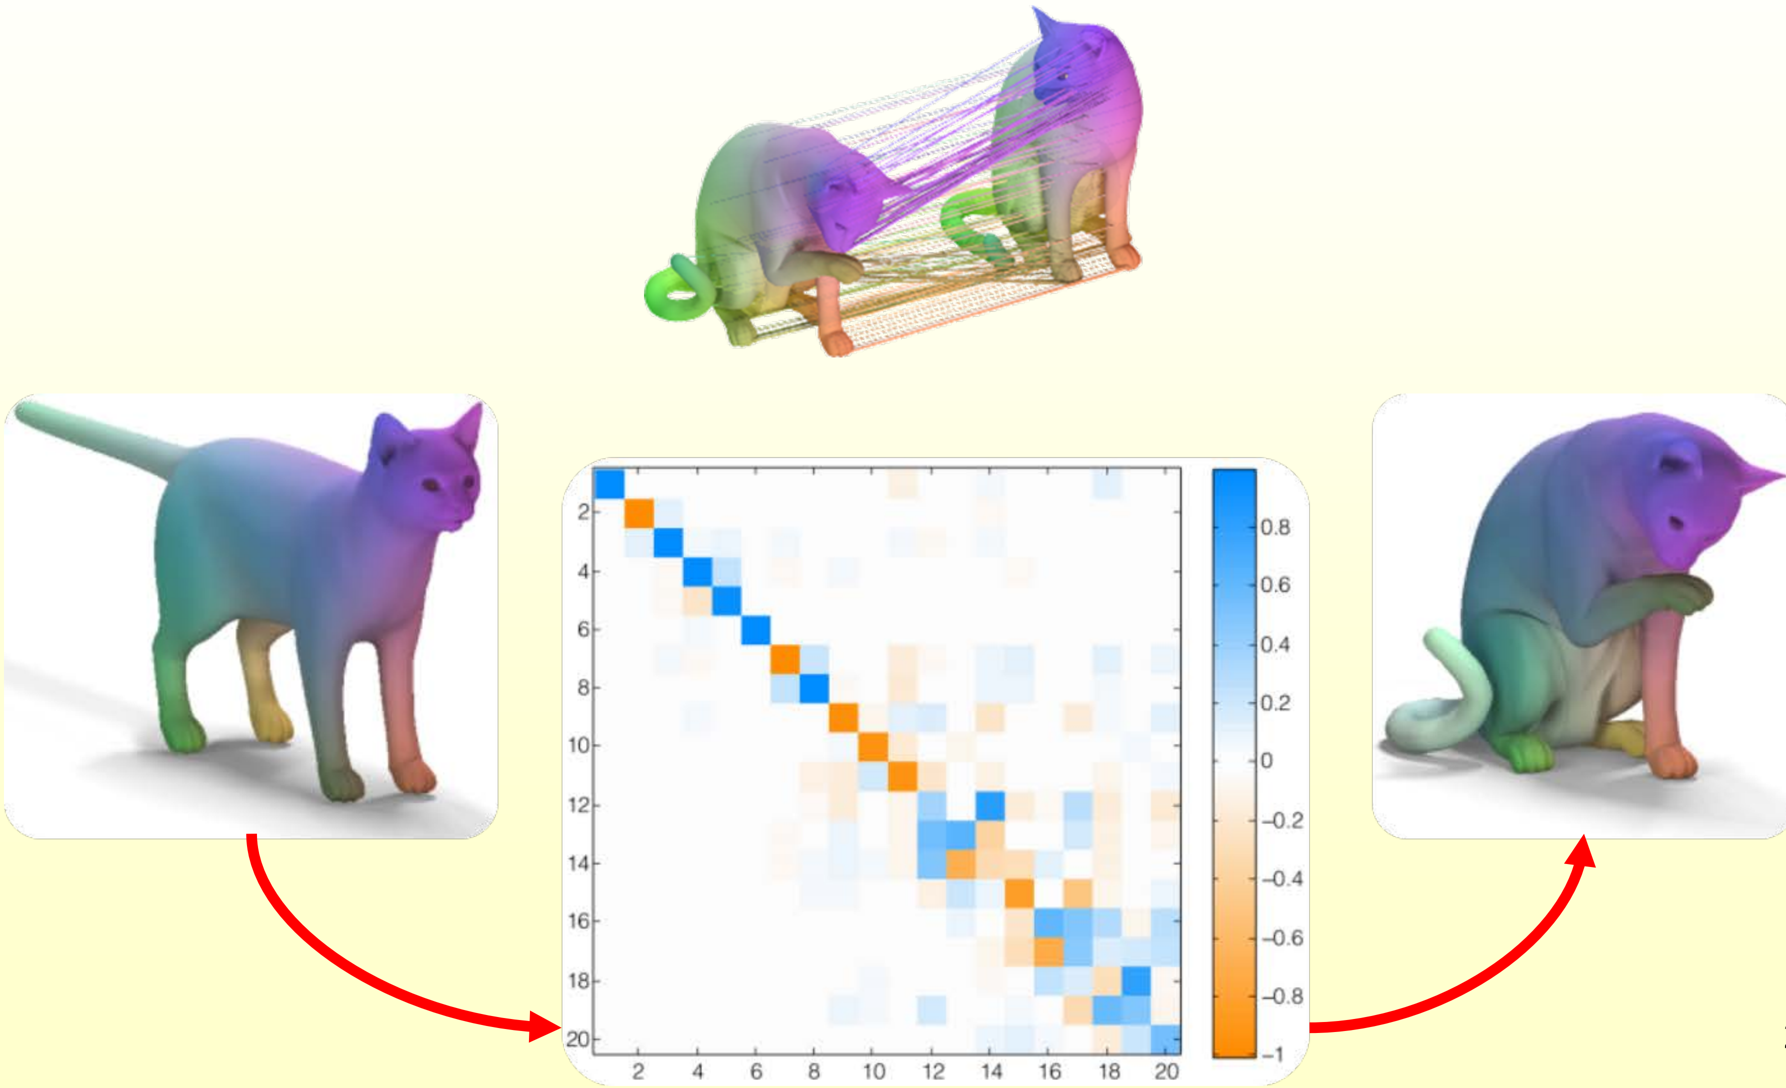

Starting from a Regular Map φ

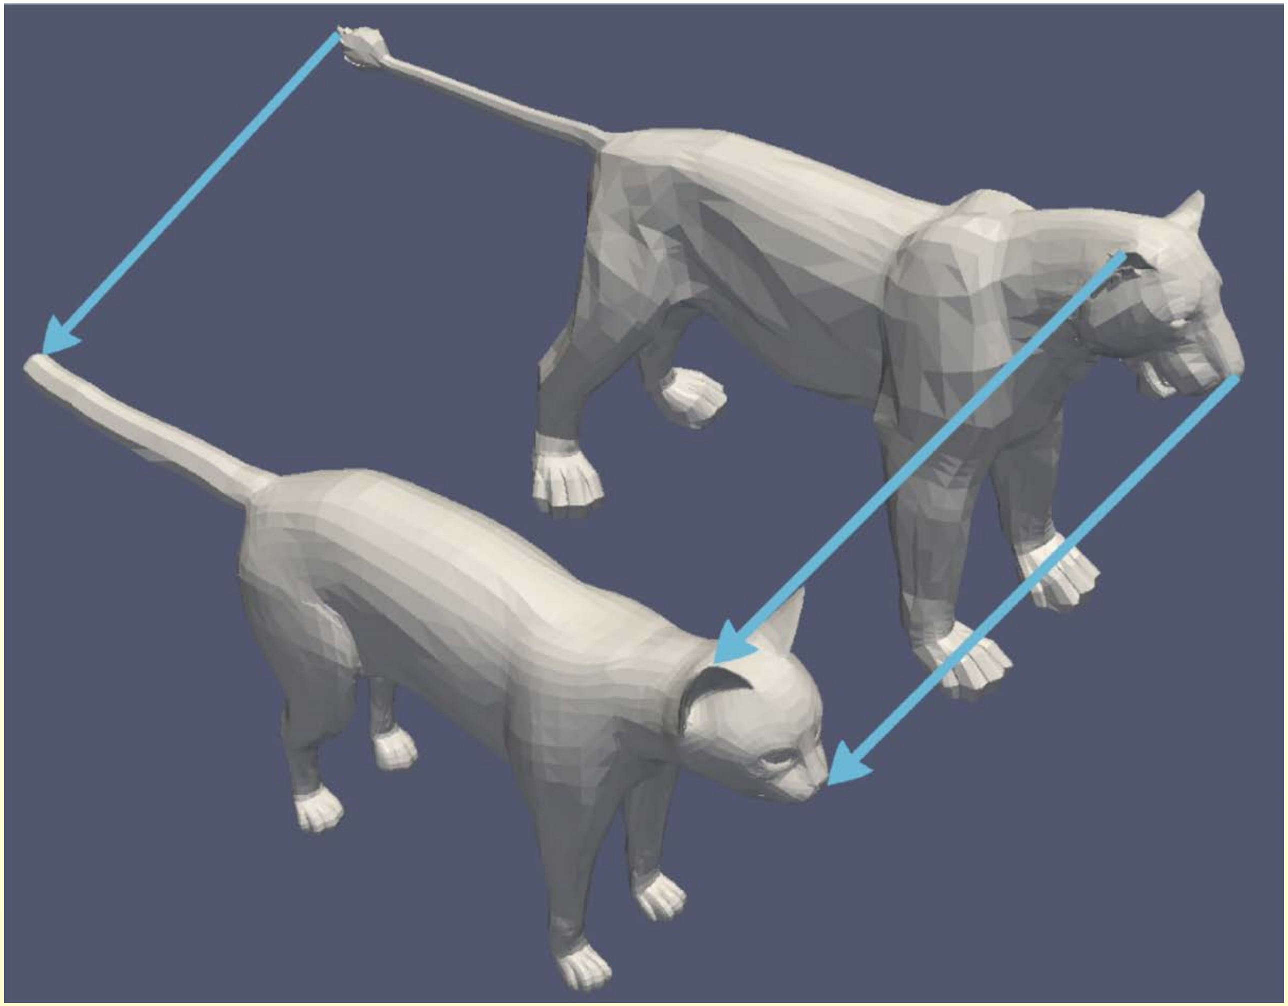

φ: lion → cat

Attribute Transfer via Pull-Back

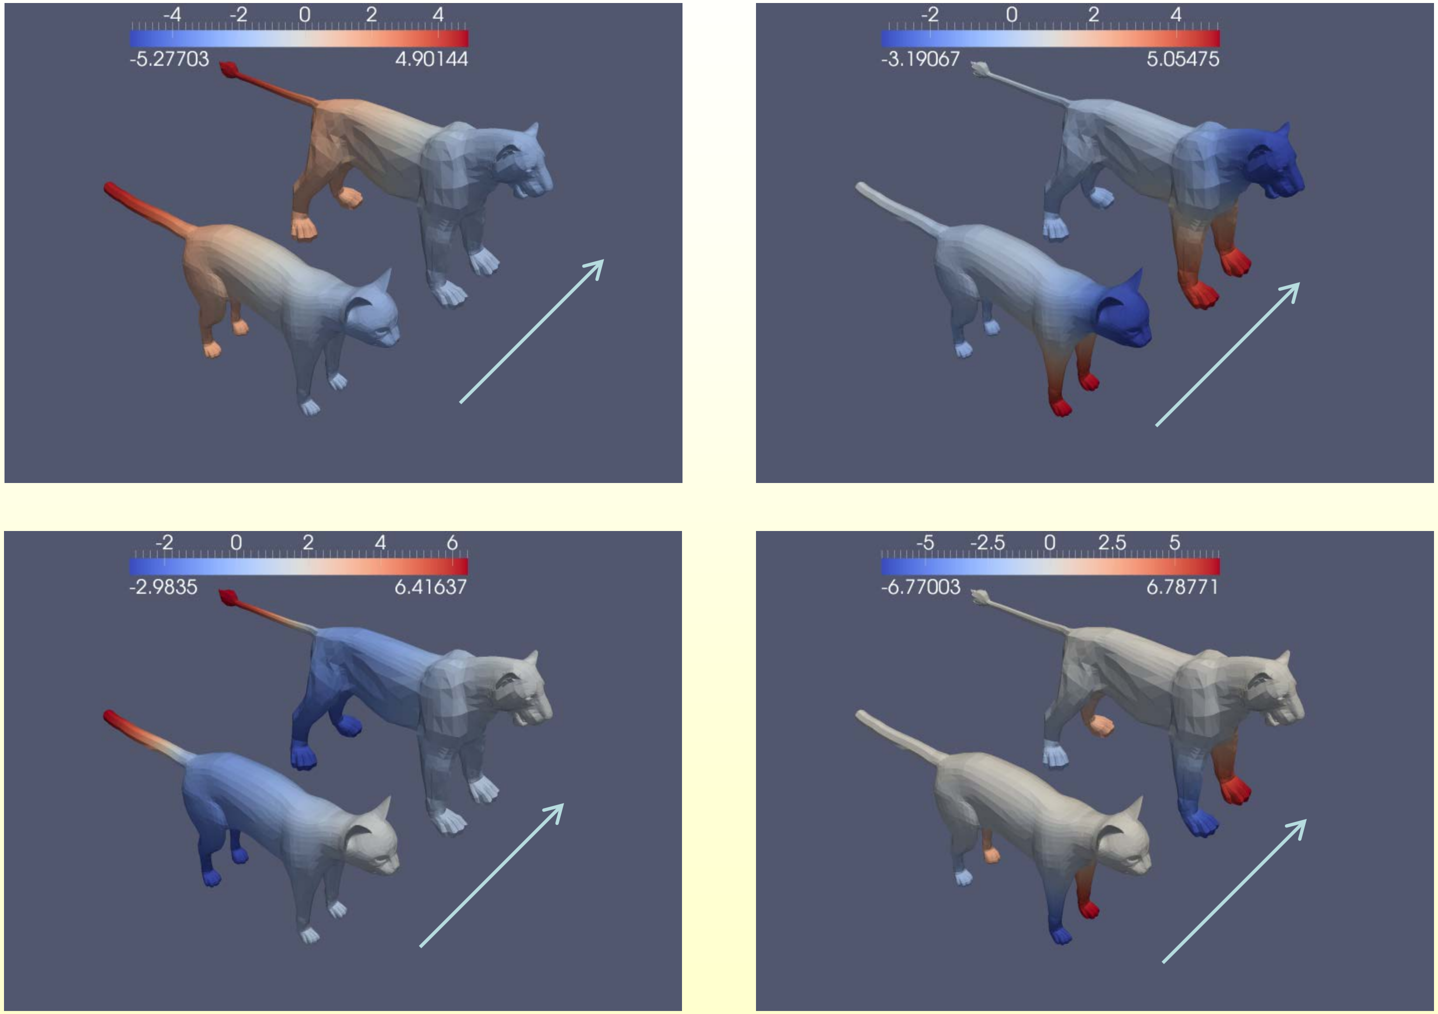

T_φ: cat → lion

A Contravariant Functor

from cat to lion

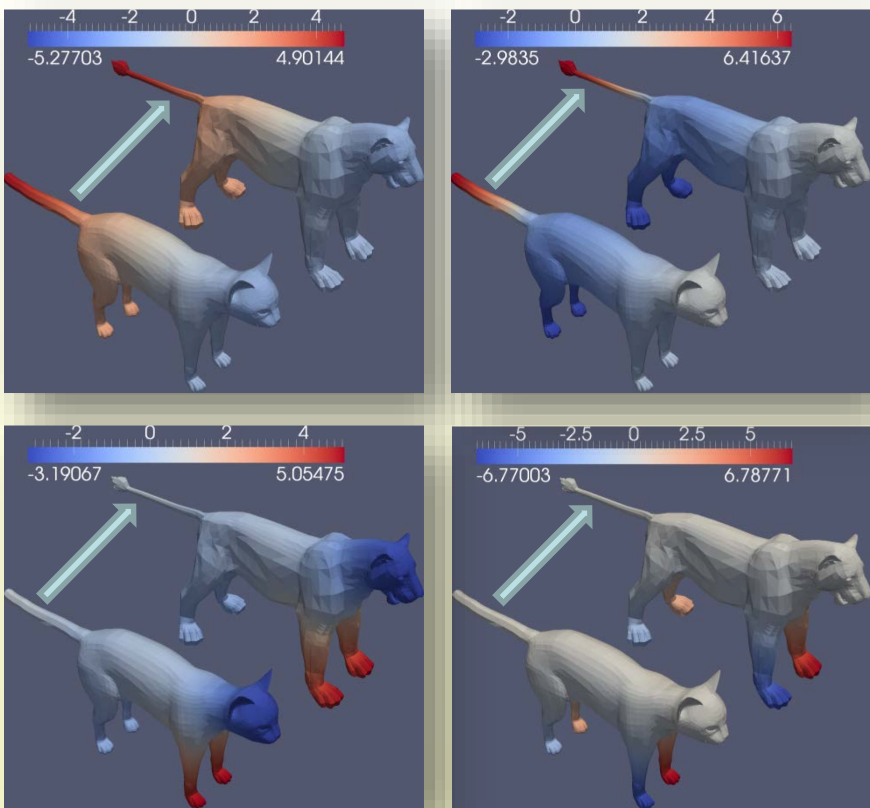

Functions on cat are transferred to lion using T_φ

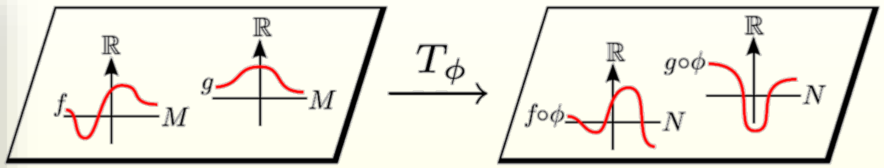

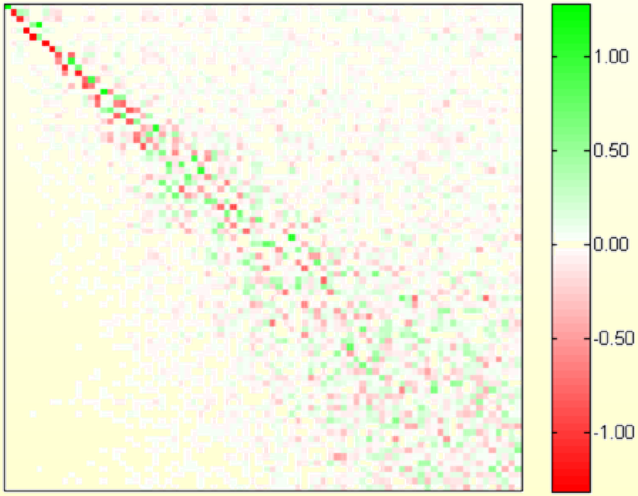

T_φ is a linear operator (matrix)

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

Functional Map

$$\phi : M \to N$$

$$T_\phi : L^2(N) \to L^2(M)$$

Dual of a point-to-point map

Bases for a Function Space

Point basis
Finite-element basis

Local bases

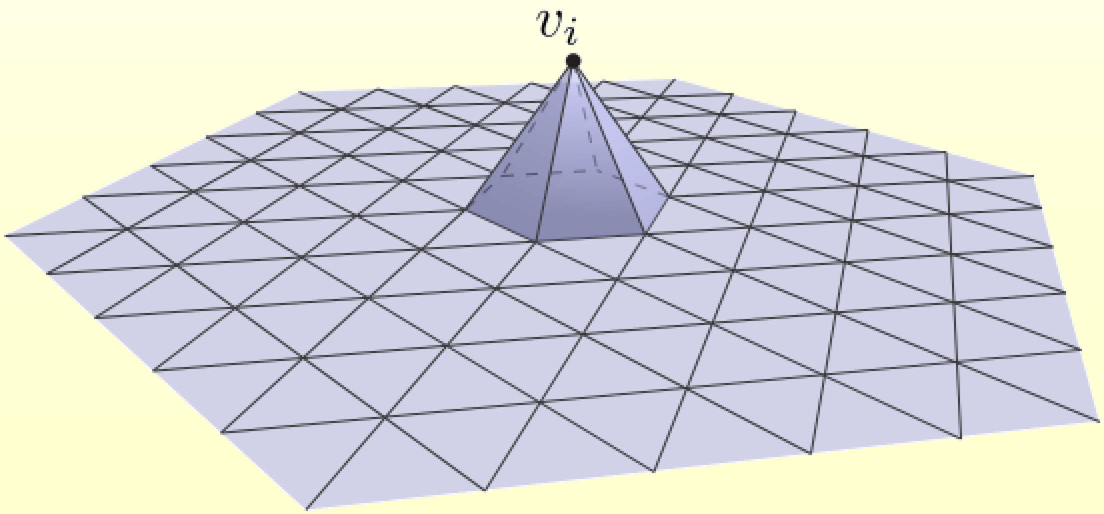

Bases for a Function Space

Fourier

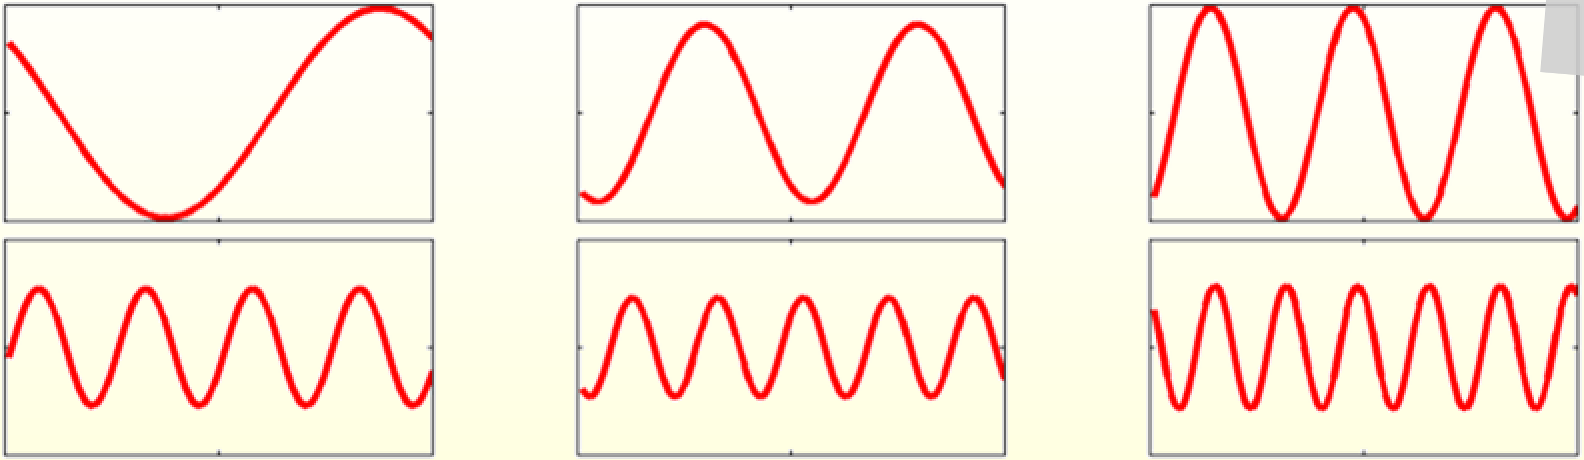

Laplace-Beltrami

global support

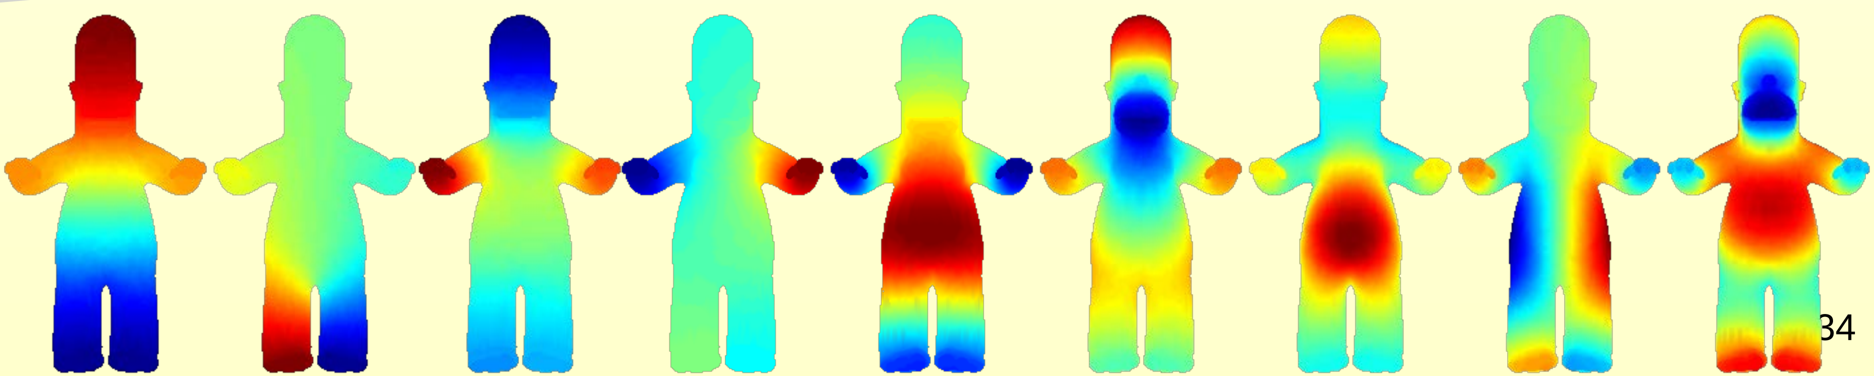

More Exotic Bases Possible

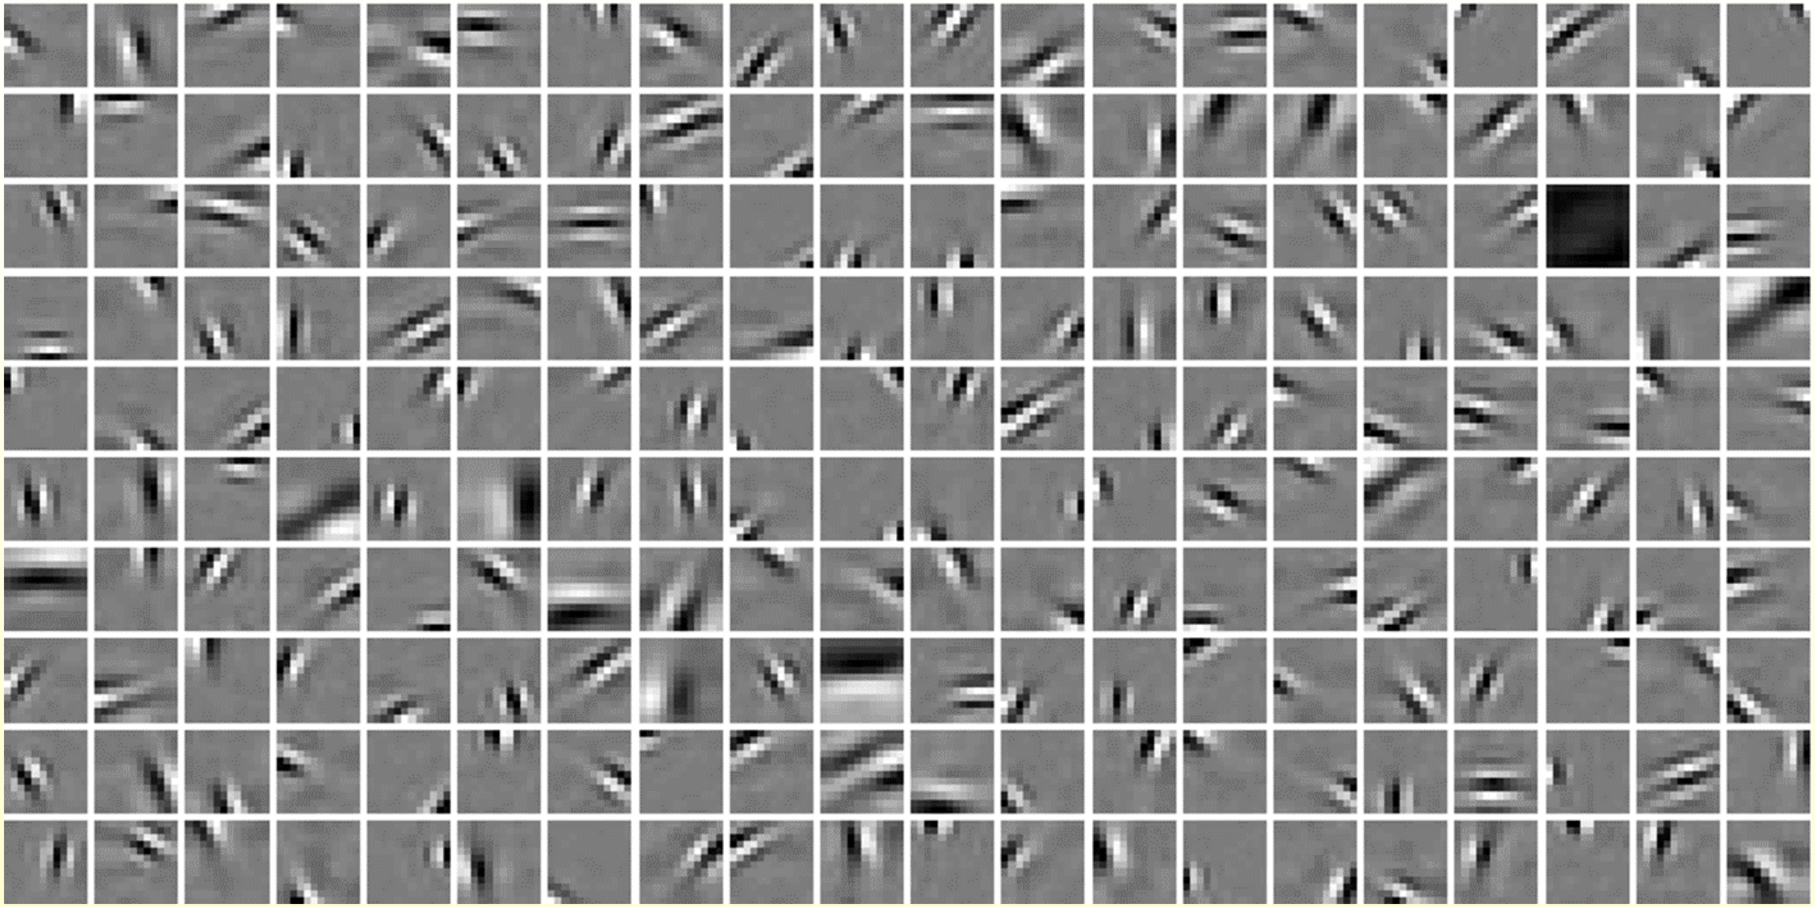

Textons, wavelets, …

Exploit Linearity

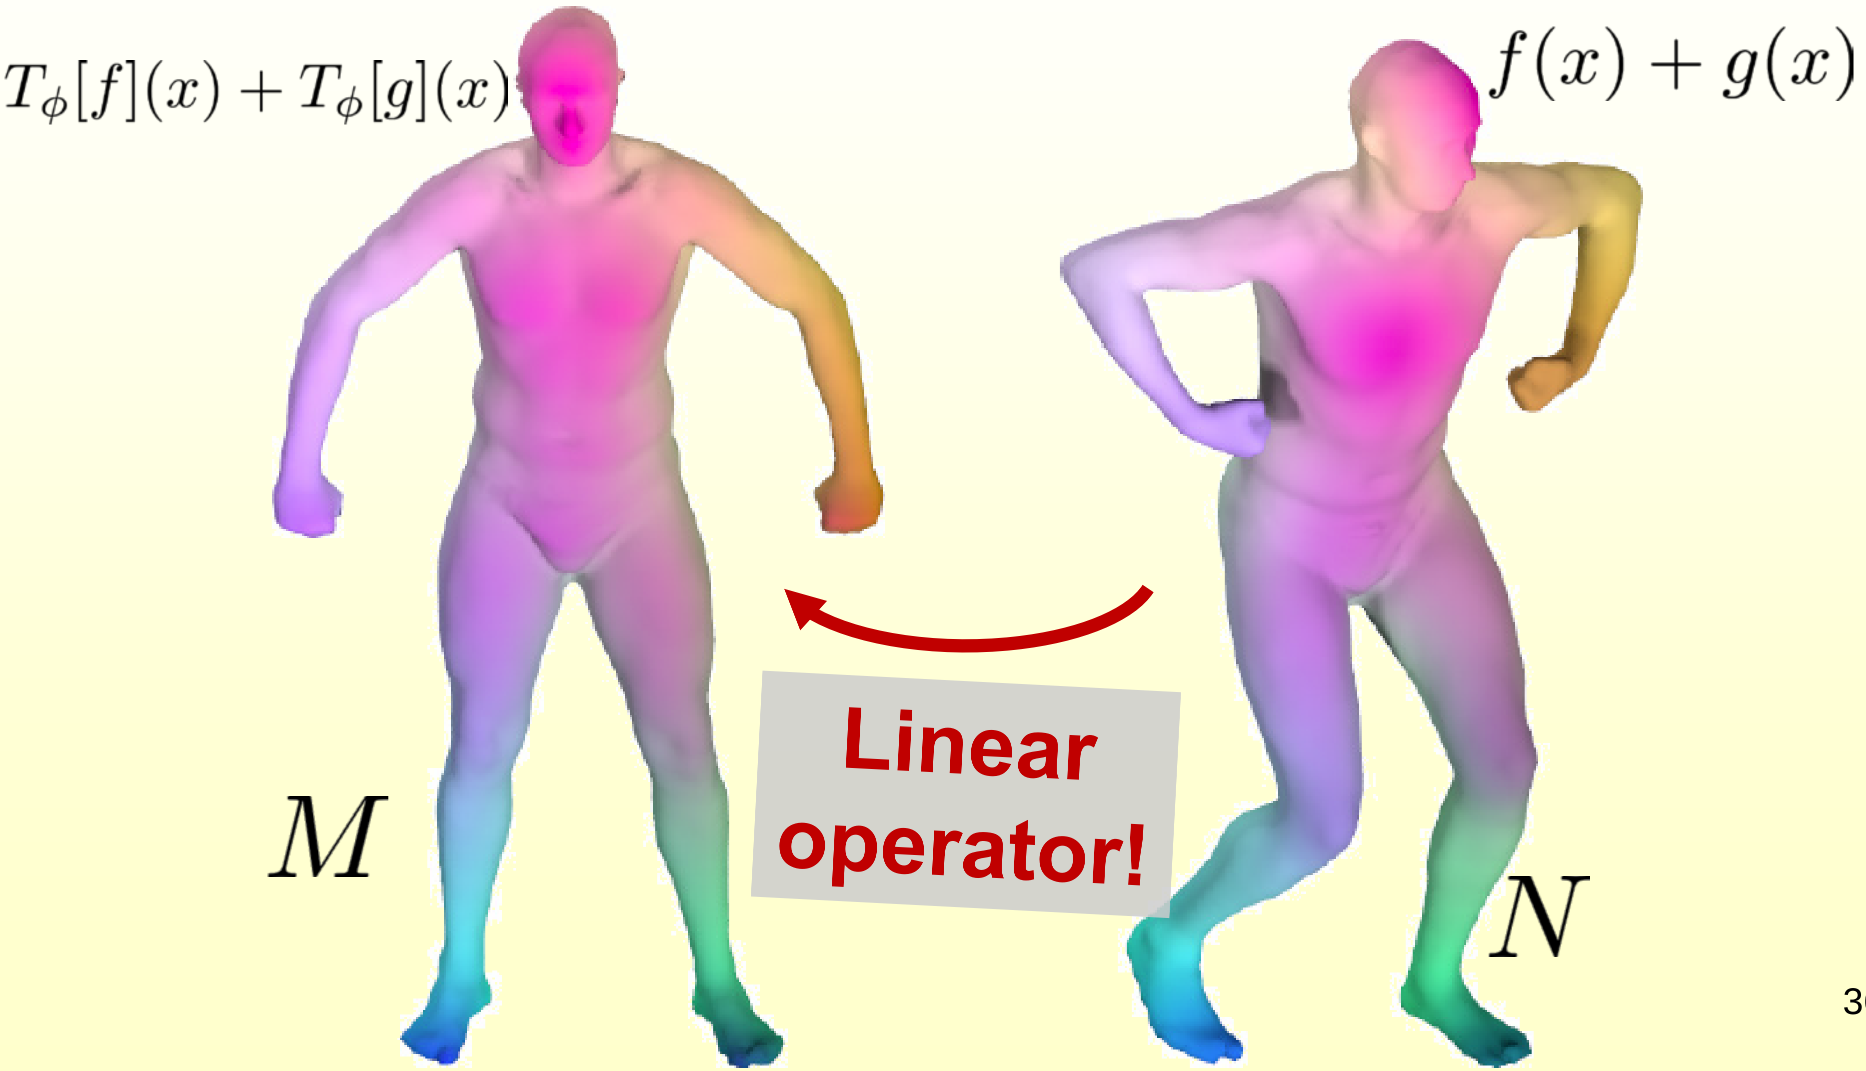

Application of Basis

$f(x) = a_1 \cdot$  $+ a_2 \cdot$  $+ a_3 \cdot$  $+ \cdots$

$\phi :$ 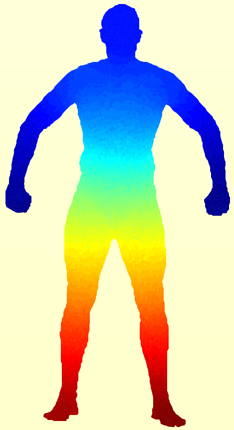 $\rightarrow$ 

Application of Basis

$T_\phi[f](x) = T_\phi[a_1 \cdot$  $+ a_2 \cdot$  $+ a_3 \cdot$  $+ \cdots]$

$= a_1 T_\phi[$  $] + a_2 T_\phi[$ 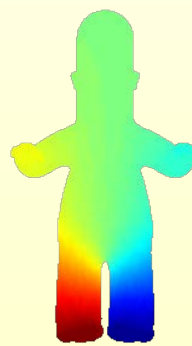 $] + a_3 T_\phi[$  $] + \cdots$

Enough to know these

Functional Map Matrix

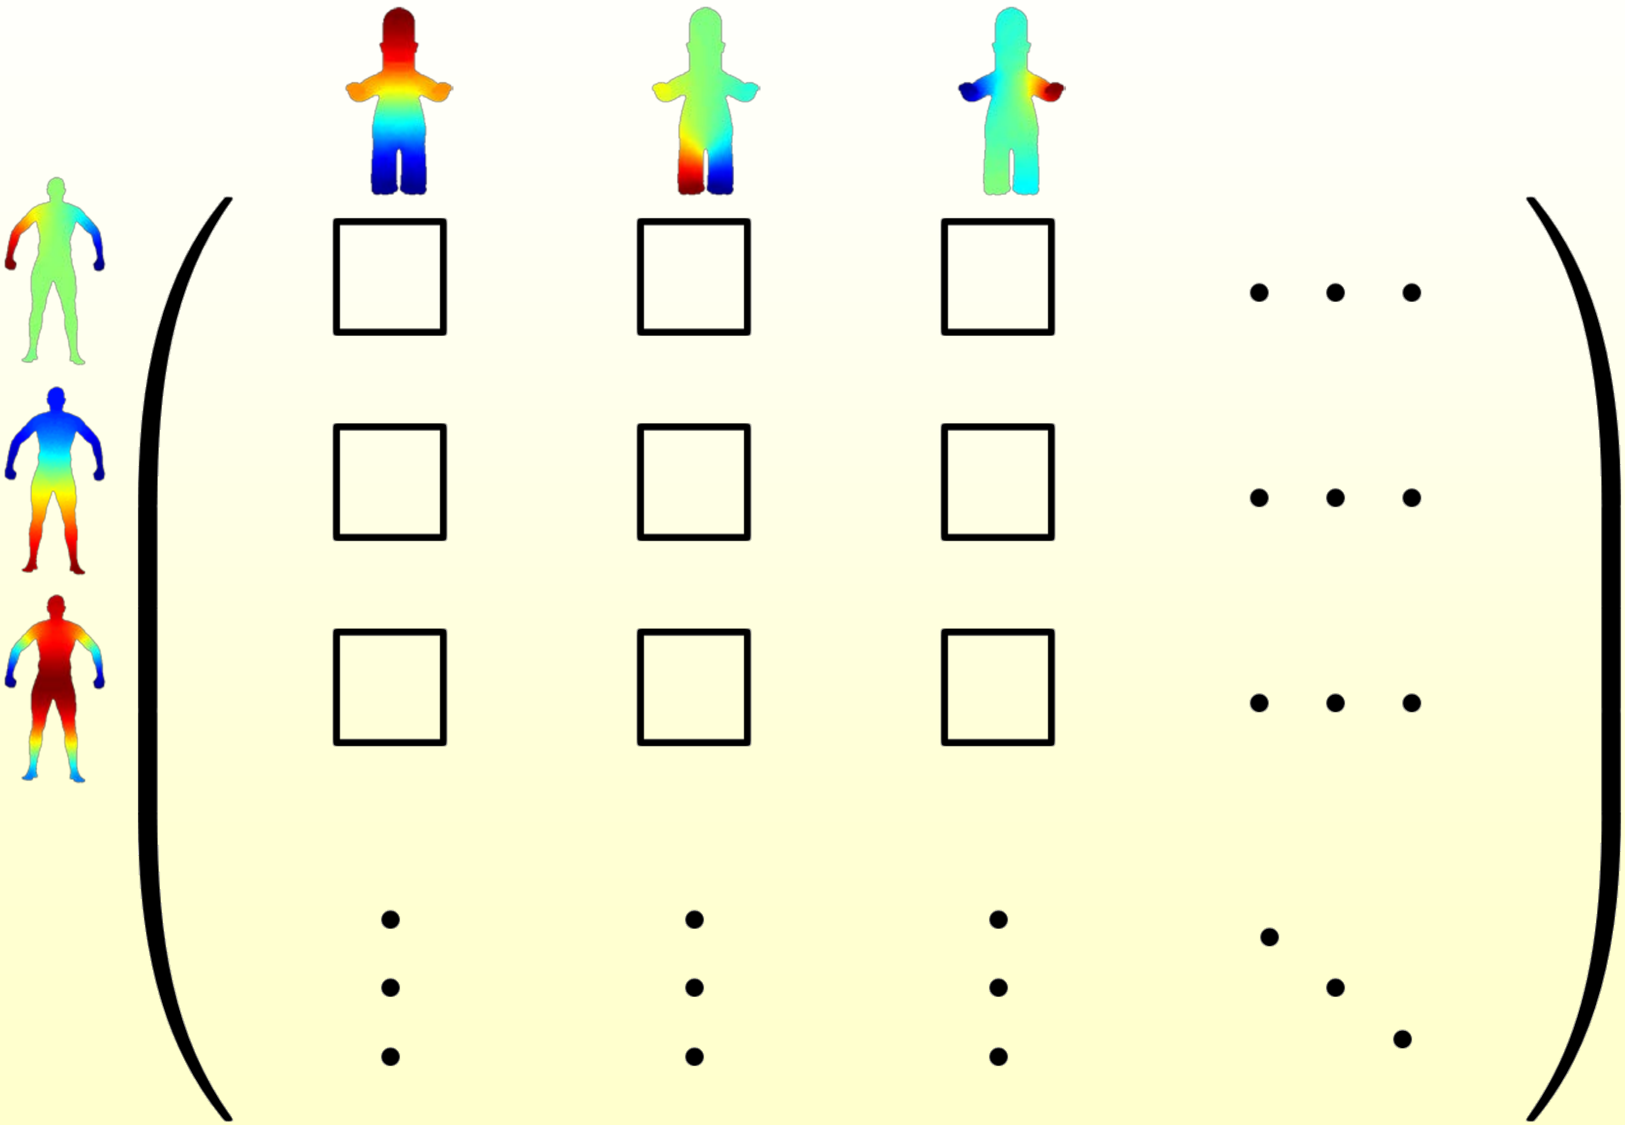

Functional Map Representation

Definition

For a fixed choice of basis functions $\{\phi^M\}$ and $\{\phi^N\}$, and a bijection $T : M \rightarrow N$, define its **functional representation** as a matrix C, s.t. for all $f = \sum_i a_i \phi_i^M$, if $T_F(f) = \sum_i b_i \phi_i^N$ then:

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

If $\{\phi^M\}$ and $\{\phi^N\}$ are both orthonormal w.r.t. some inner product, then

$$C_{ij} = \left\langle T_F(\phi_i^M), \phi_j^N \right\rangle .$$

Map Composition

$$\phi_1 : M \to N, \phi_2 : N \to P$$

$$T_{\phi_1} : L^2(N) \to L^2(M)$$

$$T_{\phi_2} : L^2(P) \to L^2(N)$$

$$T_{\phi_1}[T_{\phi_2}[f]]$$

Matrix multiplication

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

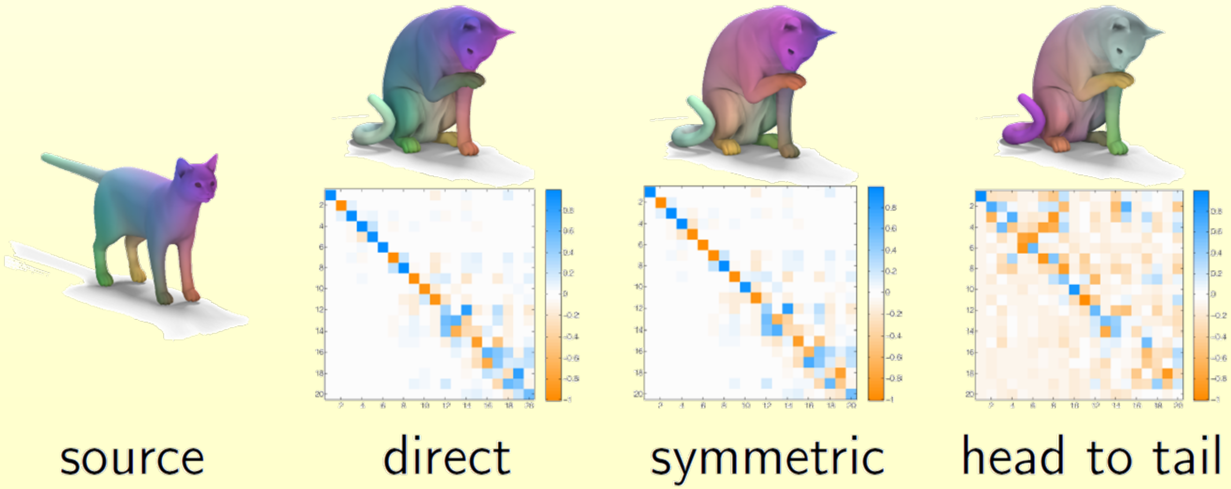

Using truncated Laplace-Beltrami basis

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 be s.t.

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

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.

Plenty of Functions: Descriptors for Points and Parts

- For shapes, there are many descriptors with various types of invariances

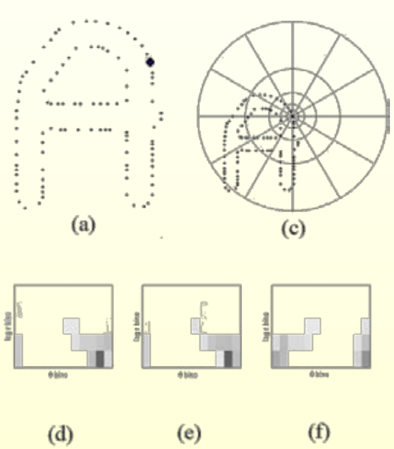

Shape Contexts:
[Belongie et al. ’00, Frome et al. ’04]

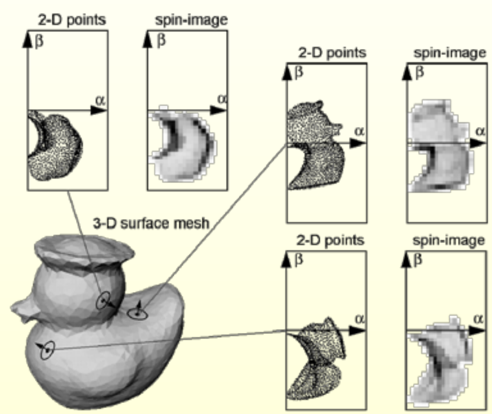

Spin Images:
[Johnson, Hebert ’99]

Rigid invariance (extrinsic)

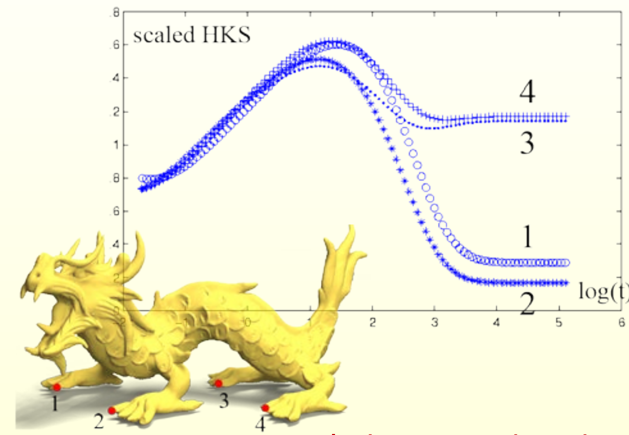

Heat Kernel Signatures (HKS):
[Sun, Ovsjanikov, G. ’08]]

Isometric invariance (intrinsic)

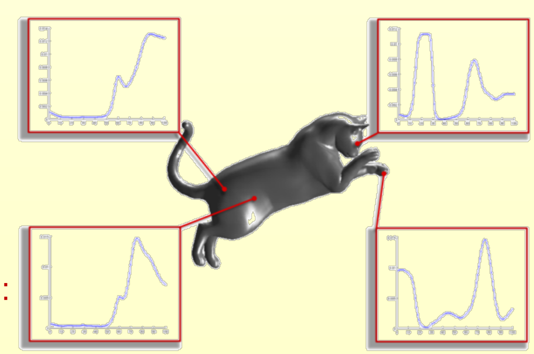

Wave Kernel Signatures (WKS):
[Aubry et. al. ‘11]

Map Estimation

$$CD_1 = D_2 \implies C = D_2 D_1^{-1}$$

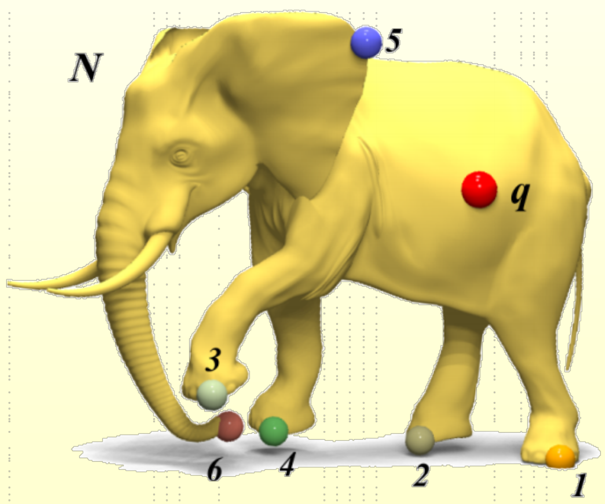

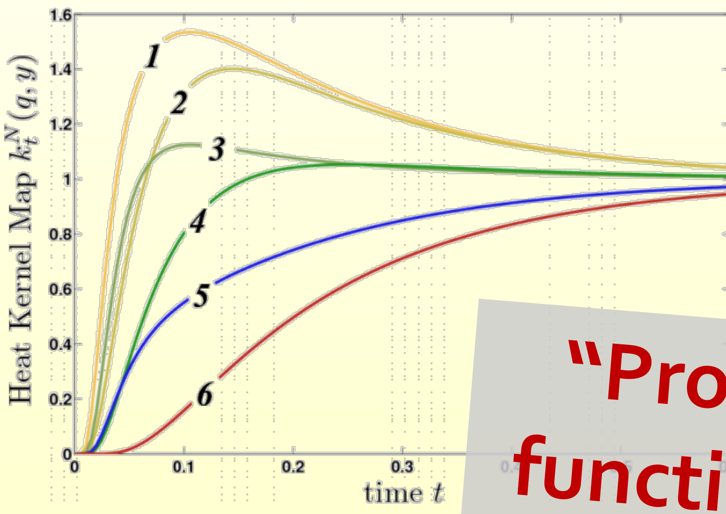

"Probe function"

Map from a linear solve

Function Preservation Constraints

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 be s.t.

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

Function preservation constraint is quite general and includes:

- Descriptor preservation (e.g. Gaussian curvature, spin images, HKS, WKS).
- Landmark correspondences (e.g. distance to the point).
- Part correspondences (e.g. indicator function).
- Texture preservation

injection of low-level knowledge or supervision

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.

Operator Commutativity

$$C\Delta_1 \approx \Delta_2 C$$

Differentiate and then transport

Transport and then differentiate

Isometry Regularizer

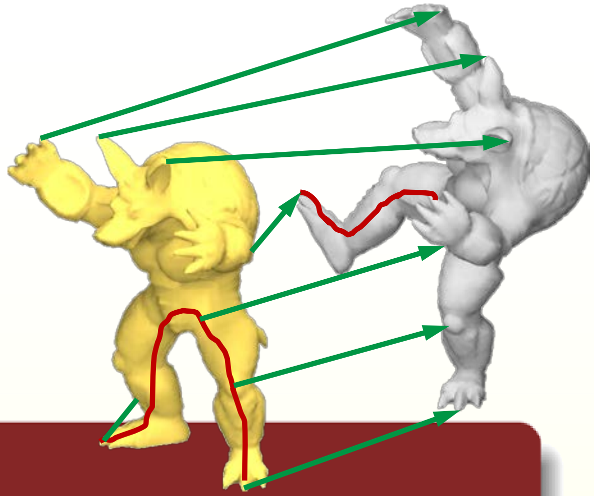

Lemma 1:

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

$$C\Delta_1 = \Delta_2 C$$

Also conformality, area or volume preservation, etc.

Conformal Regularization

Lemma 3:

If the mapping is *conformal* if and only if:

$$C^T \Delta_1 C = \Delta_2$$

Using these regularizations, we get a very efficient shape matching method.

Volume Preservation Regularizer

Lemma 2:

The mapping is *locally volume preserving*, if and only if the functional map matrix is *orthonormal*:

$$C^T C = \mathrm{I}$$

Rotations/reflections in functions space

Sparcity in a Localized Basis

$$\min \|C\|_{2,1}$$

Sum of Euclidean norms of cols

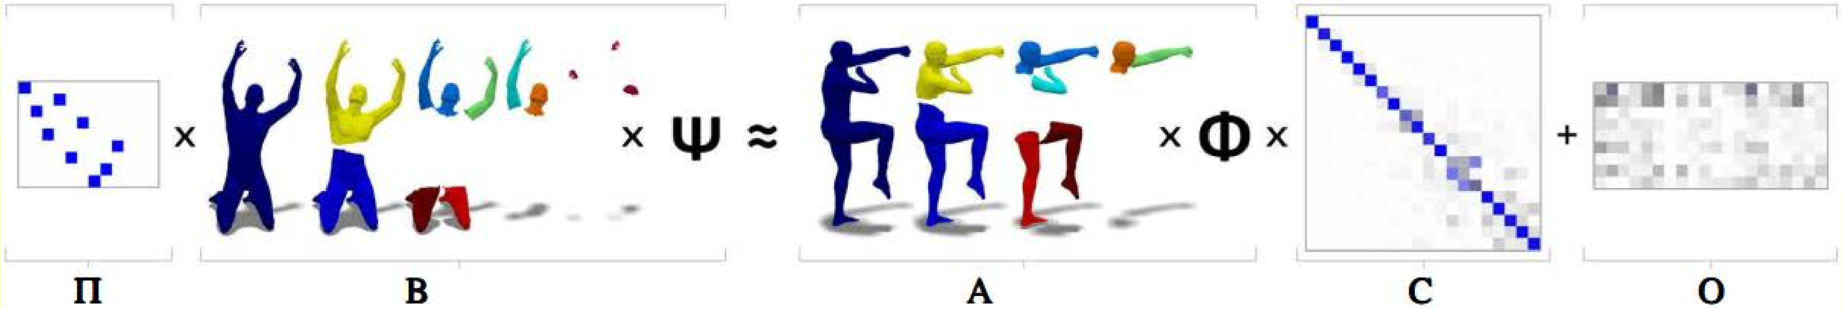

Sparse Modeling of Intrinsic Correspondences (Pokrass, Bronstein², Sprechmann, Sapiro)

General Optimization for Maps

$$
\begin{aligned}
\min_C \quad & \|CD_1 - D_2\|_2^2 \\
& [+\alpha\|C\Delta_1 - \Delta_2 C\|_{\text{Fro}}^2] \\
& [+\beta\|C\|_{2,1}] \\
\text{such that} \quad & [C^\top C = I]
\end{aligned}
$$

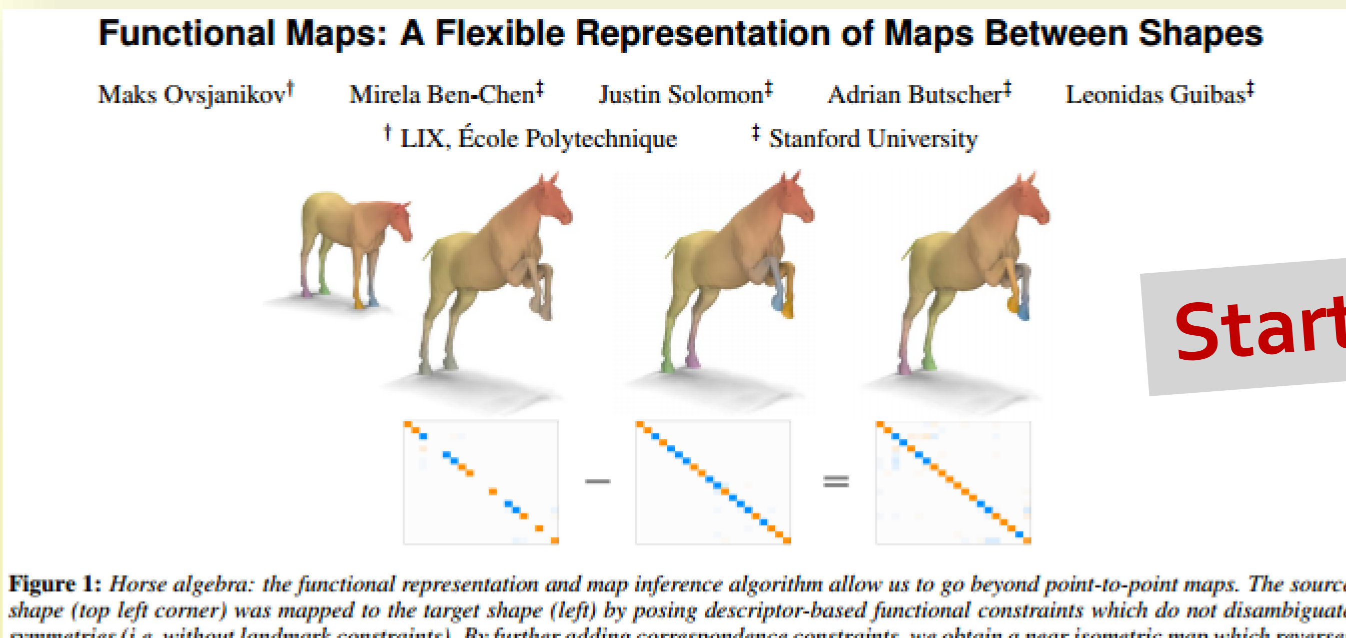

Functional Maps: A Flexible Representation of Maps Between Shapes

Maks Ovsjanikov[†] Mirela Ben-Chen[‡] Justin Solomon[‡] Adrian Butscher[‡] Leonidas Guibas[‡]

[†] LIX, École Polytechnique [‡] Stanford University

Figure 1: *Horse algebra: the functional representation and map inference algorithm allow us to go beyond point-to-point maps. The source shape (top left corner) was mapped to the target shape (left) by posing descriptor-based functional constraints which do not disambiguate*

Map Continuity

- Not explicitly enforced
- Implicit in the choice of basis

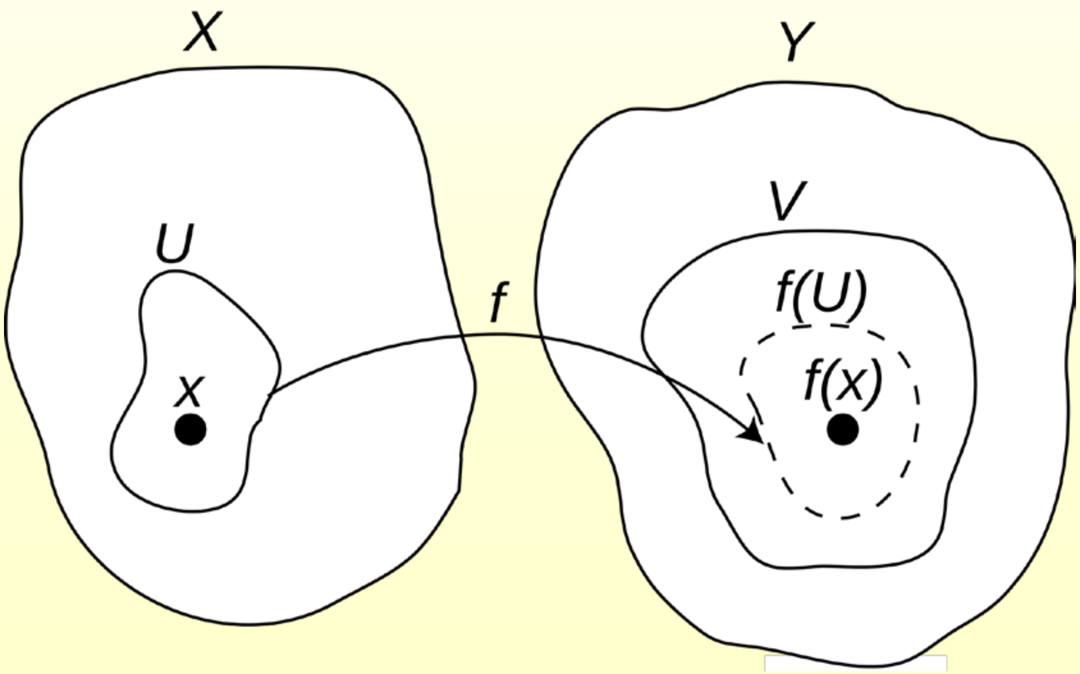

From Functional to Point-to-Point Maps

- Can try transporting delta functions individually -- expensive

Primal

$$p \mapsto q$$

Dual

$$p \begin{pmatrix} 0 \\ 0 \\ \vdots \\ 0 \\ 1 \\ 0 \\ \vdots \\ 0 \end{pmatrix} \mapsto \begin{pmatrix} 0 \\ 1 \\ \vdots \\ 0 \\ 0 \\ 0 \\ \vdots \\ 0 \end{pmatrix} q$$

$$\delta_x = (\phi_1^M(x), \phi_2^M(x), \phi_3^M(x), \ldots)$$

From Functional to Point-to-Point Maps

$$C\Phi_M^\top \leftrightarrow \Phi_N$$

Image of each point on surface M

Each point on surface N in LB basis

So transport, and then use nearest neighbor search

From Functional to Point-to-Point Maps

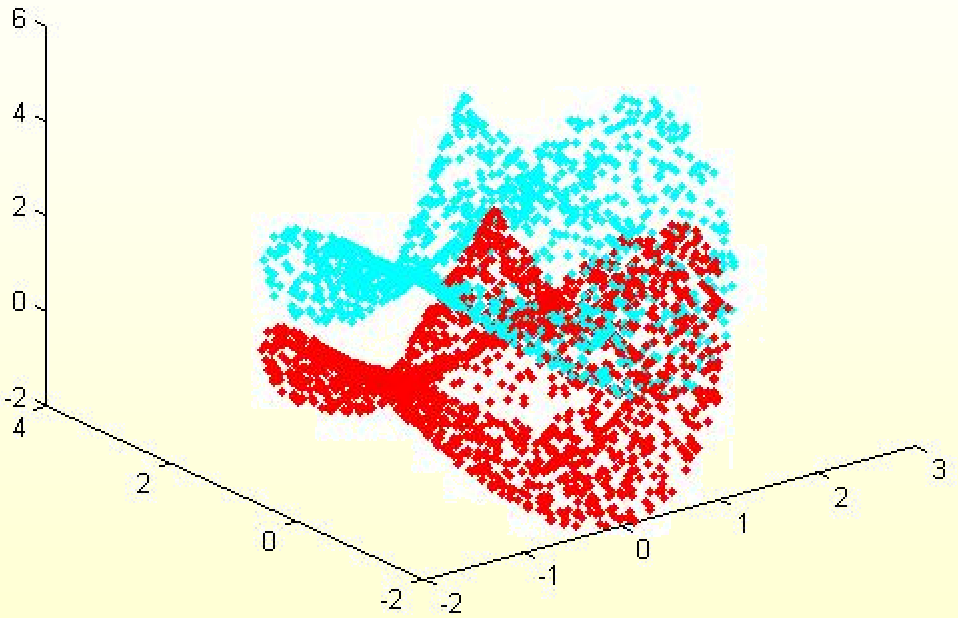

ICP in Function Space!

Ground Truth Comparison

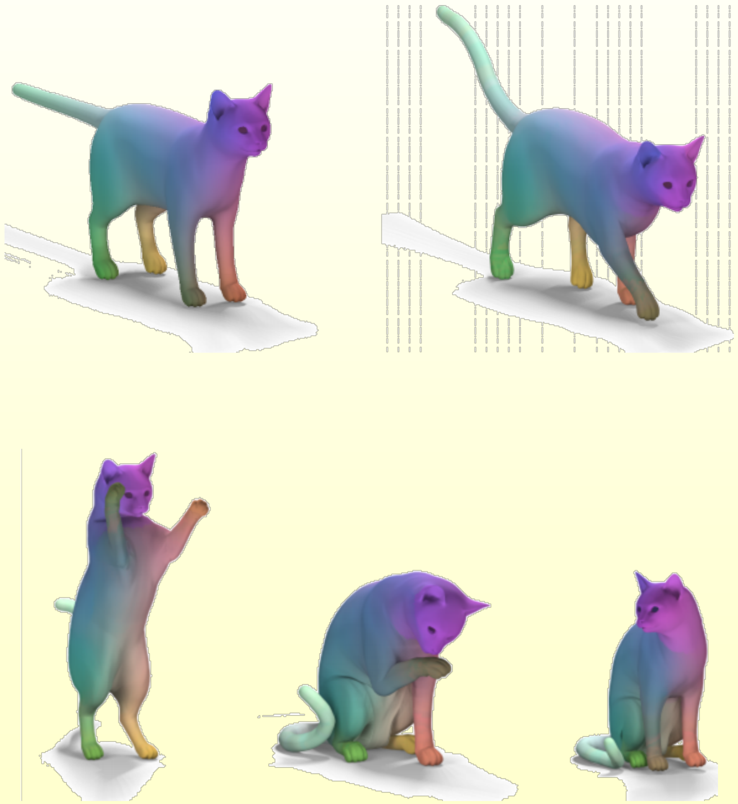

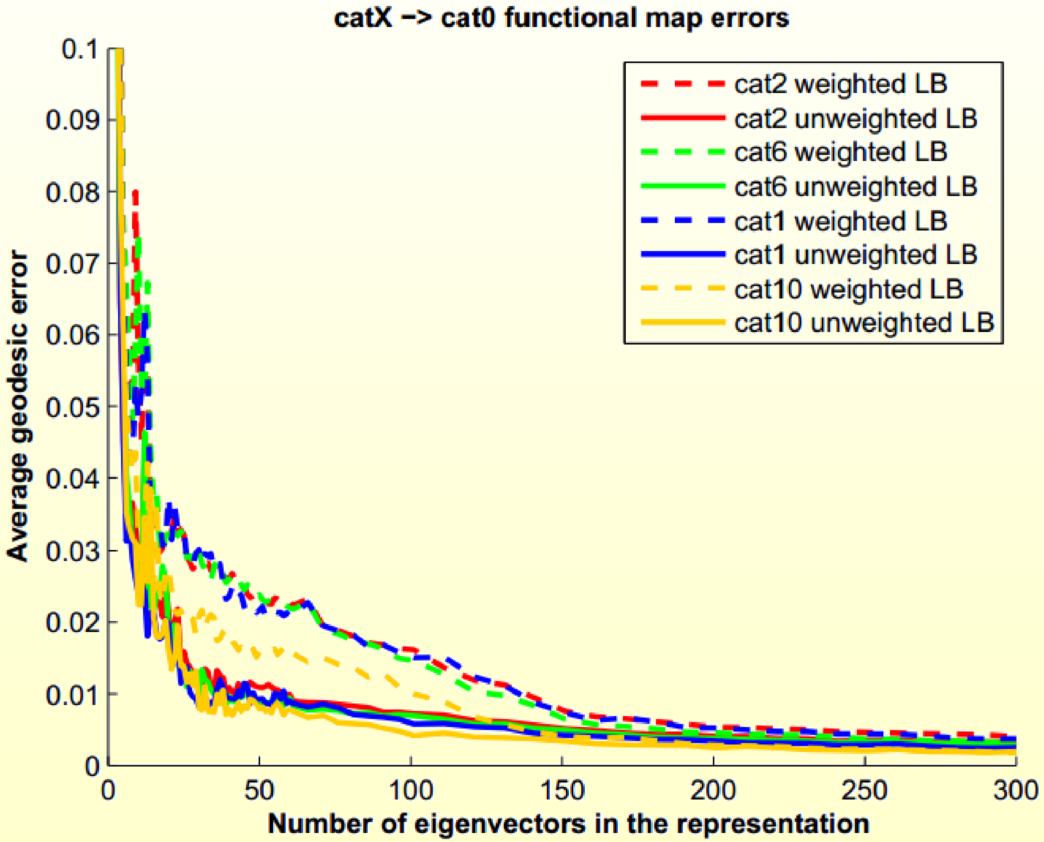

Large Scale Tests

A very simple method that puts together a modest set of constraints and uses 100 basis functions outperforms state-of-the-art:

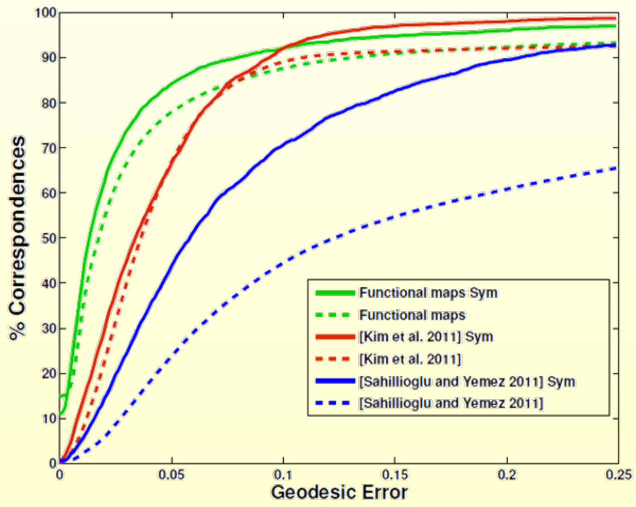

SCAPE

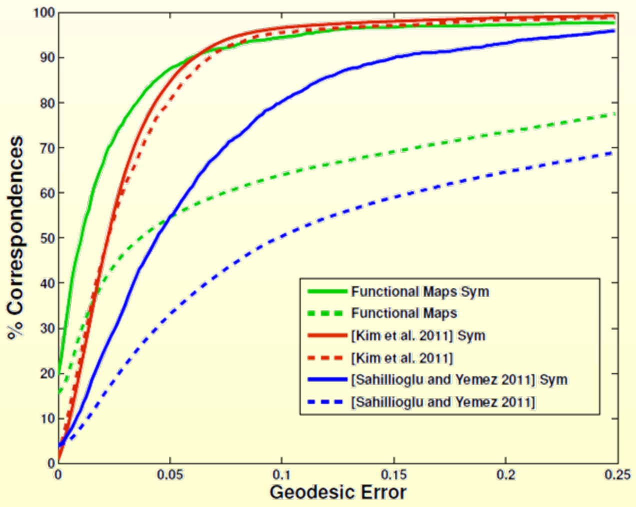

TOSCA

Roughly 10 probe functions + 1 part correspondence

Map Algebra

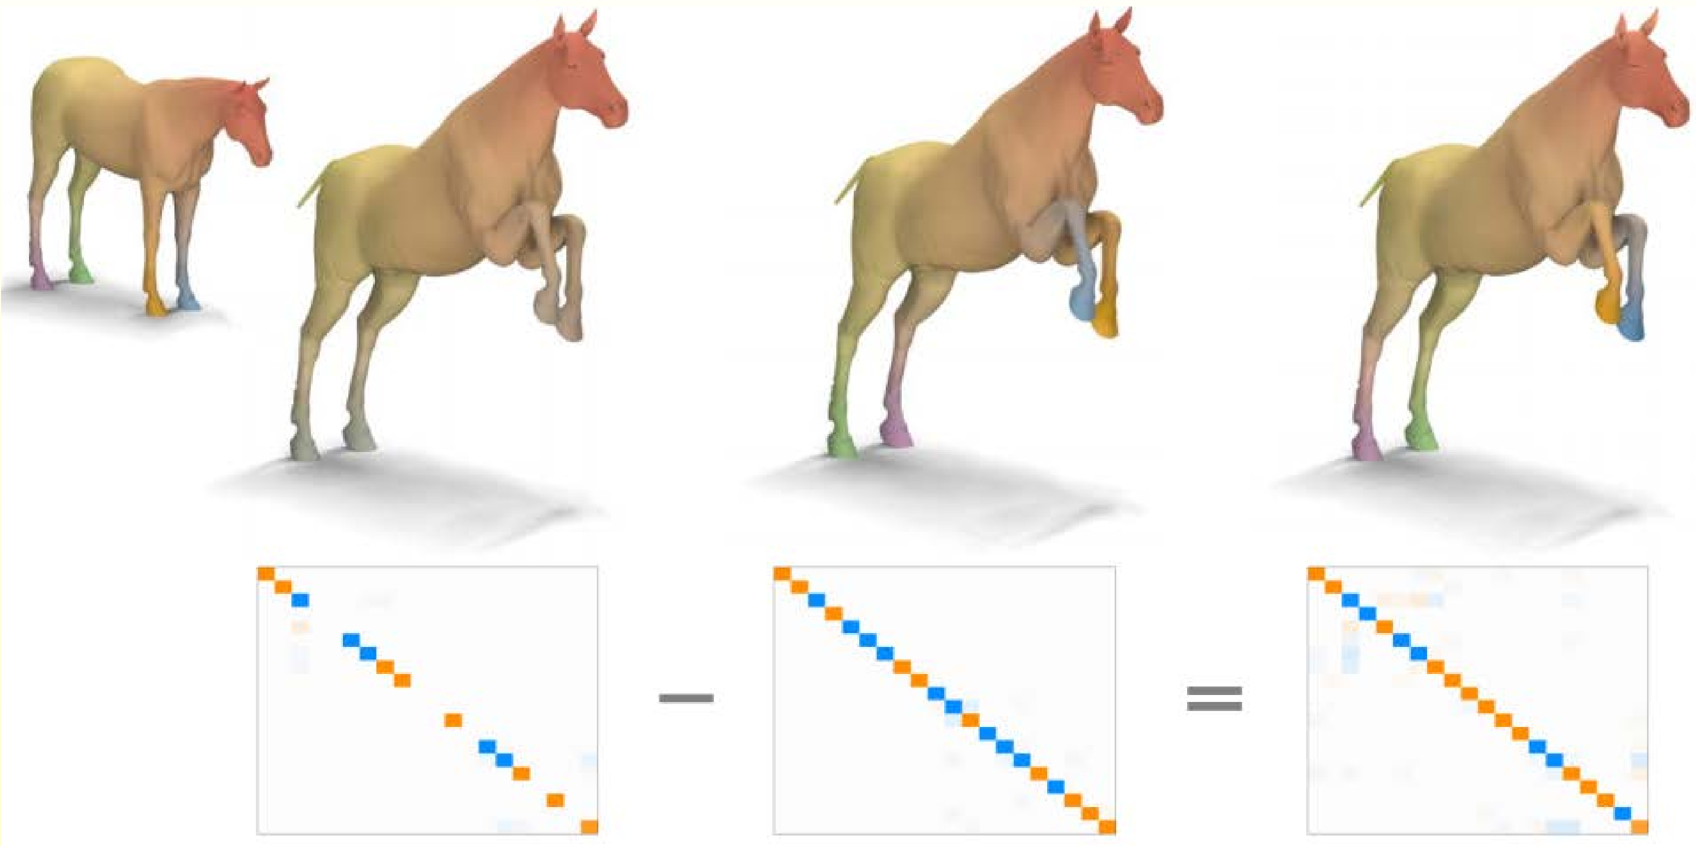

“Horse algebra”

Application: Segmentation Transfer

Map Visualization

Even given a map $T : M \to N$, it is often hard to visualize it.

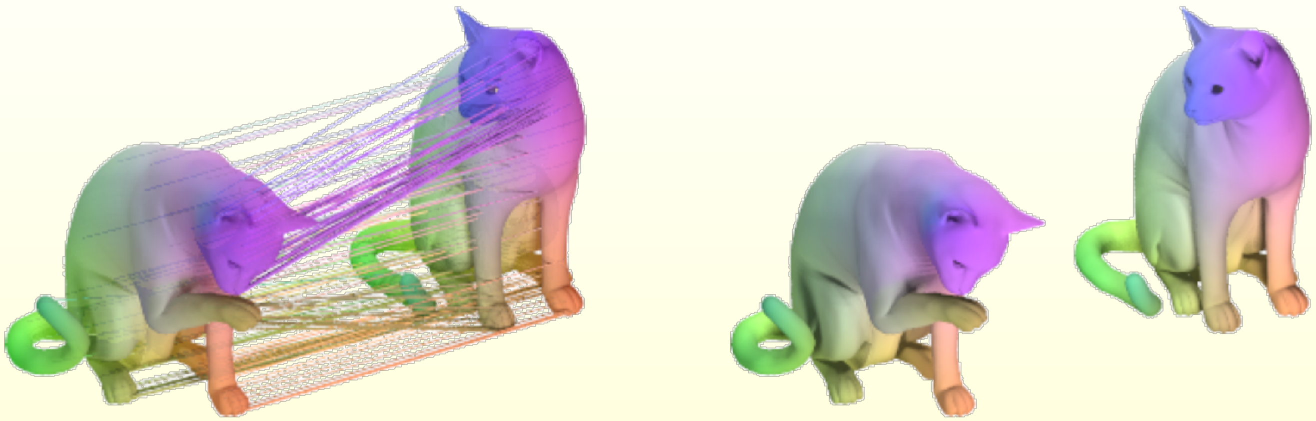

Common visualizations:

- Connecting (some) points by lines
- Plotting a function f on N and $f \circ T$ on M.

Question: how to pick a "good" function f.

Map Visualization

If f is an indicator of a region Ω, then:

$$\|f\|_2^2 = \sum_x f(x)^2 A(x) = \text{Area}(\Omega),$$
$$\|f \circ T\|_2^2 = \text{Area}\left(T^{-1}(\Omega)\right).$$

Finding the functions maximizing (minimizing) the ratio $\frac{\|f\|_2^2}{\|f \circ T\|_2^2}$ can identify regions most distorted by T.

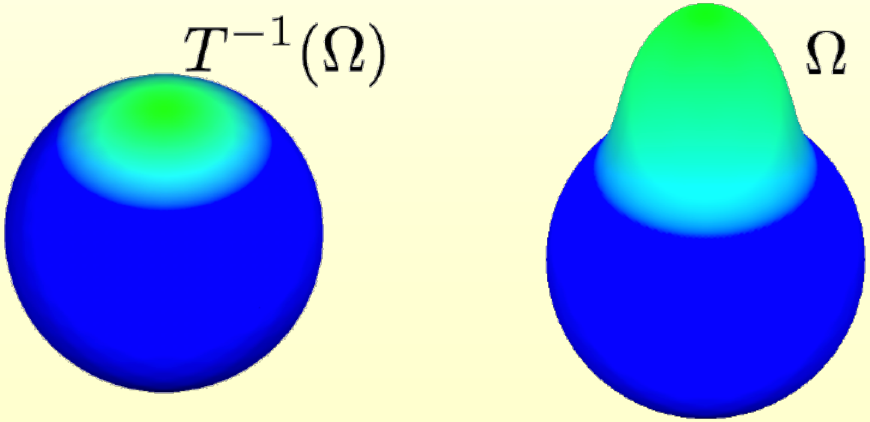

Map Visualization

In the functional representation and LB basis:

$$\frac{\|f\|_2^2}{\|f \circ T\|_2^2} = \frac{\|\mathbf{a}\|_2^2}{\|C\mathbf{a}\|_2^2}$$

and optimal functions are simply the *singular vectors* of C.

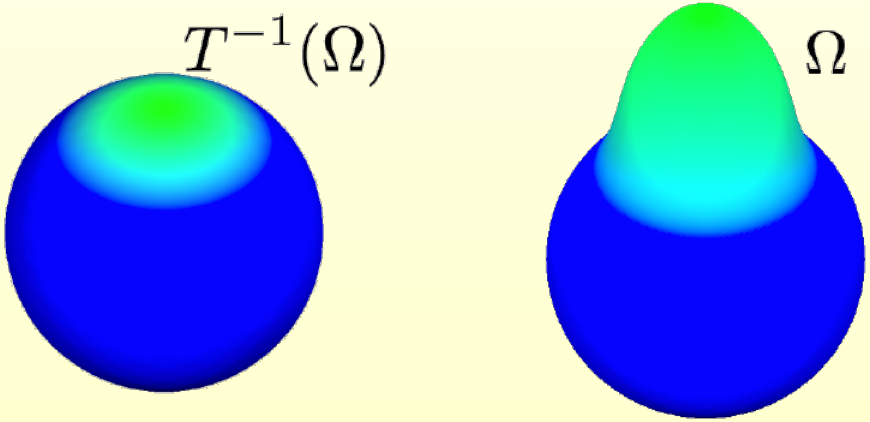

Map Visualization

Singular vectors of the functional representation C of T identify most distorted regions in a multi-scale way.

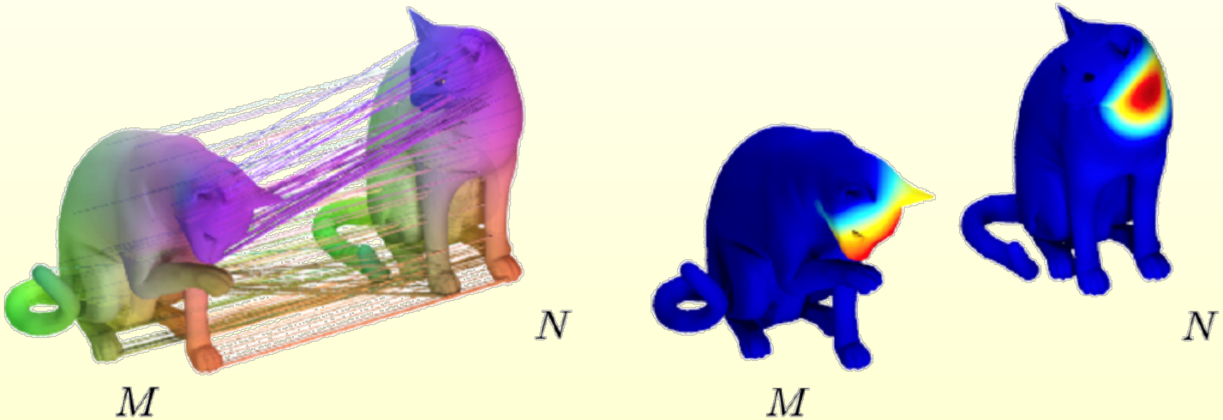

Map Visualization

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

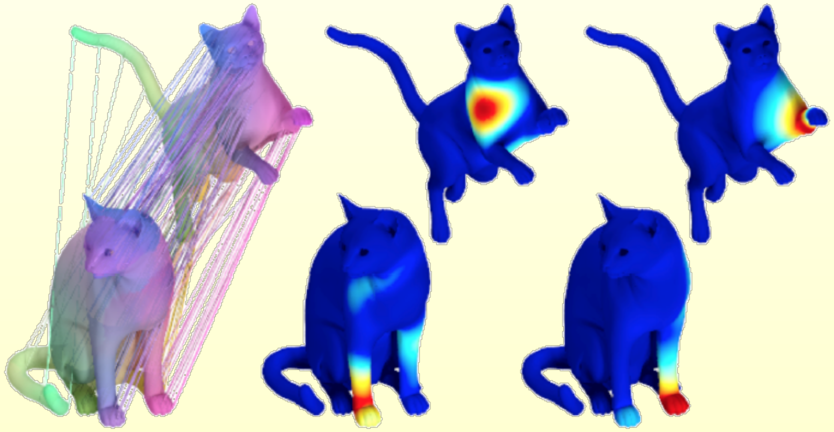

Multiple Shapes

With same method, can visualize maps to *multiple* shapes.

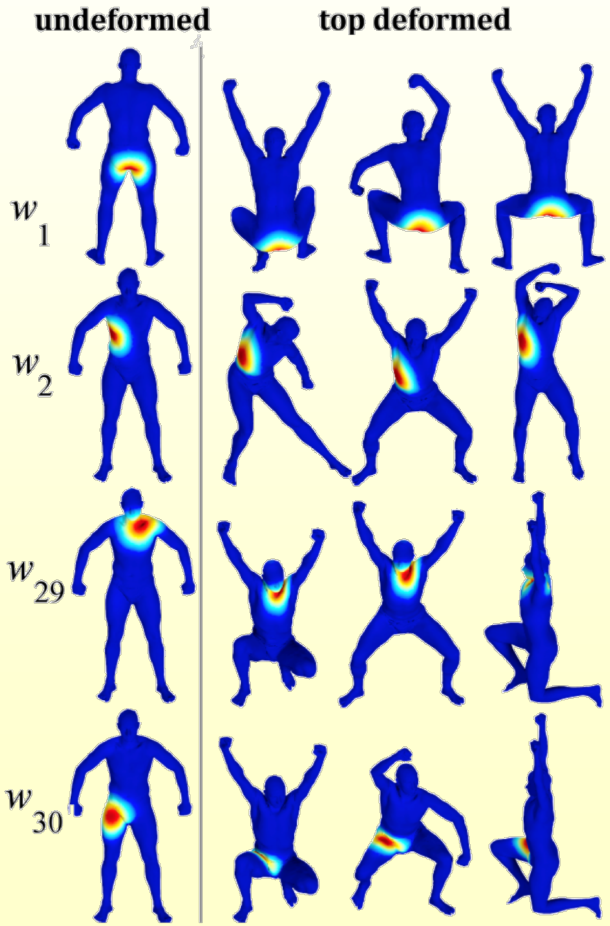

Coupled Bases

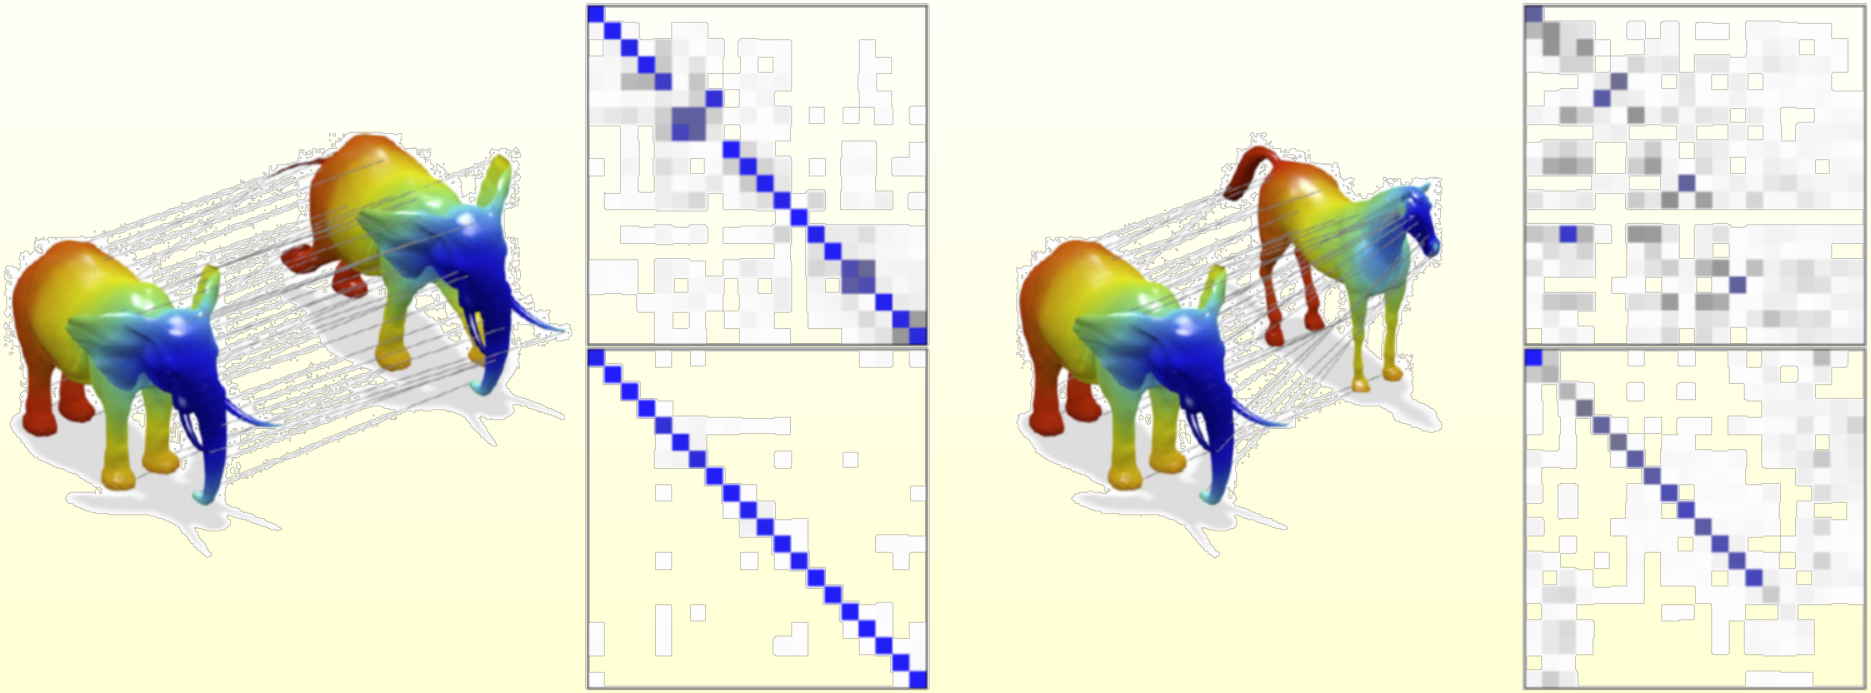

$$\min_{\mathbf{P},\mathbf{Q}} \quad \mathrm{trace}(\mathbf{P}^\top \mathbf{\Lambda}_X \mathbf{P}) + \mathrm{trace}(\mathbf{Q}^\top \mathbf{\Lambda}_Y \mathbf{Q}) + \mu \|\mathbf{P}^\top \mathbf{A} - \mathbf{Q}^\top \mathbf{B}\|$$

$$\text{s.t.} \quad \mathbf{P}^\top \mathbf{P} = \mathbf{I}, \mathbf{Q}^\top \mathbf{Q} = \mathbf{I}$$

Coupled Bases

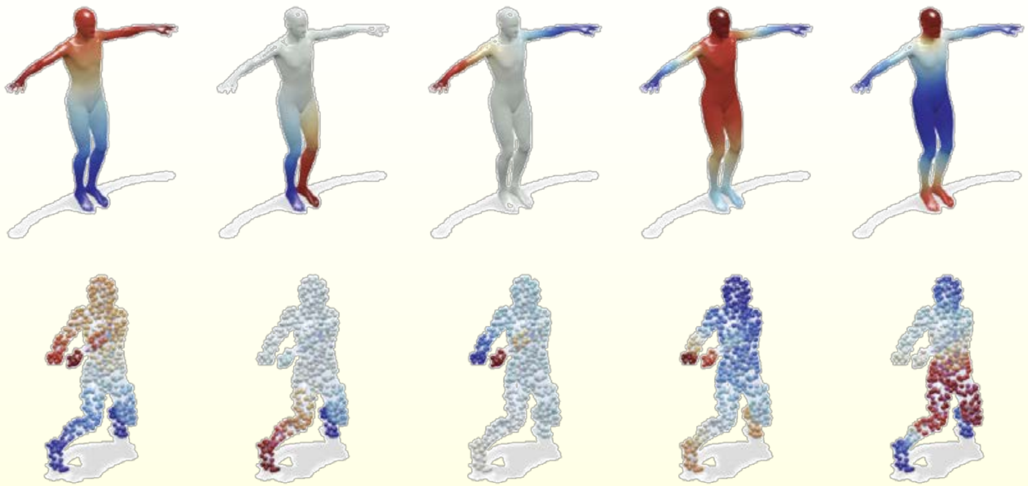

Laplacian eigenbases

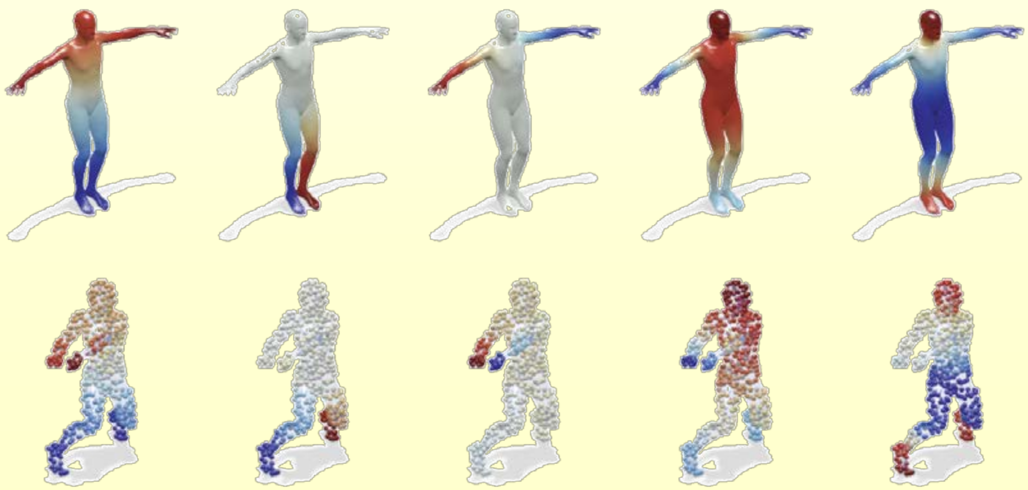

Coupled bases

Partial Functional Maps

Block diagonal case

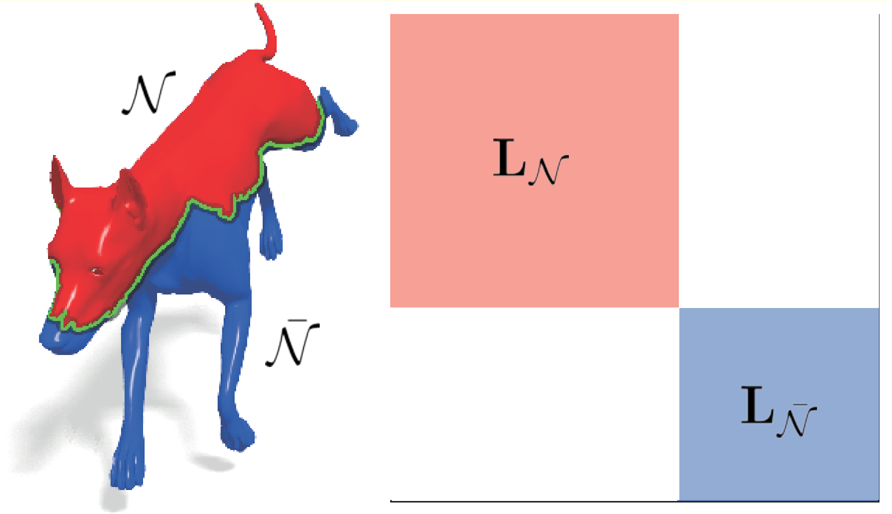

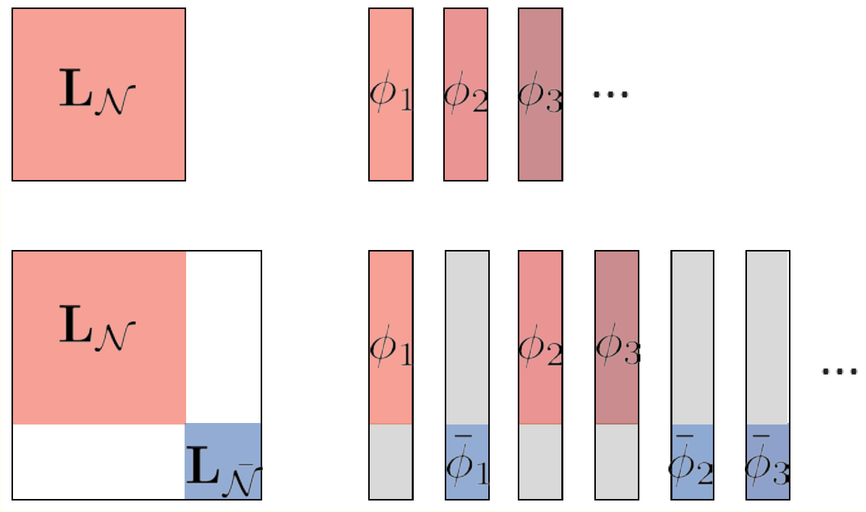

Partial Functional Maps

Weak coupling: eigenfunctions of the partial shape show up among those of the full shape, up to some bounded perturbation

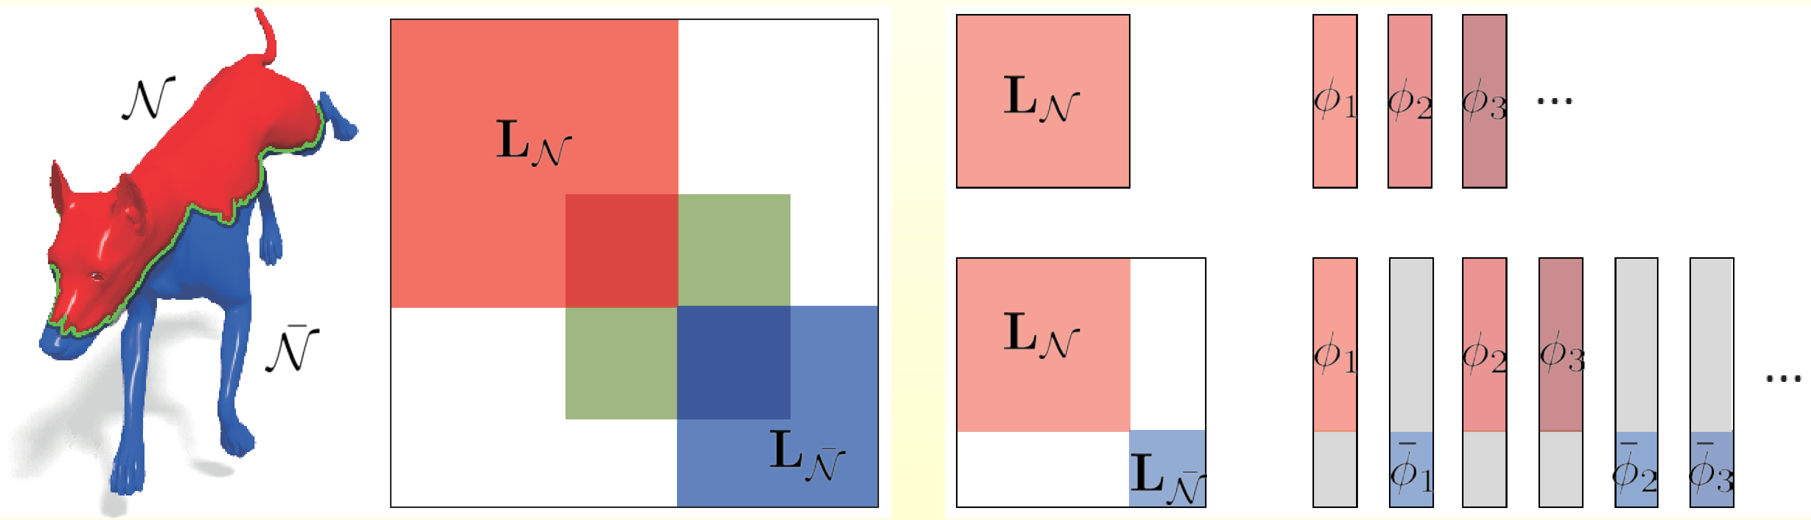

Eigenfunction Preservation

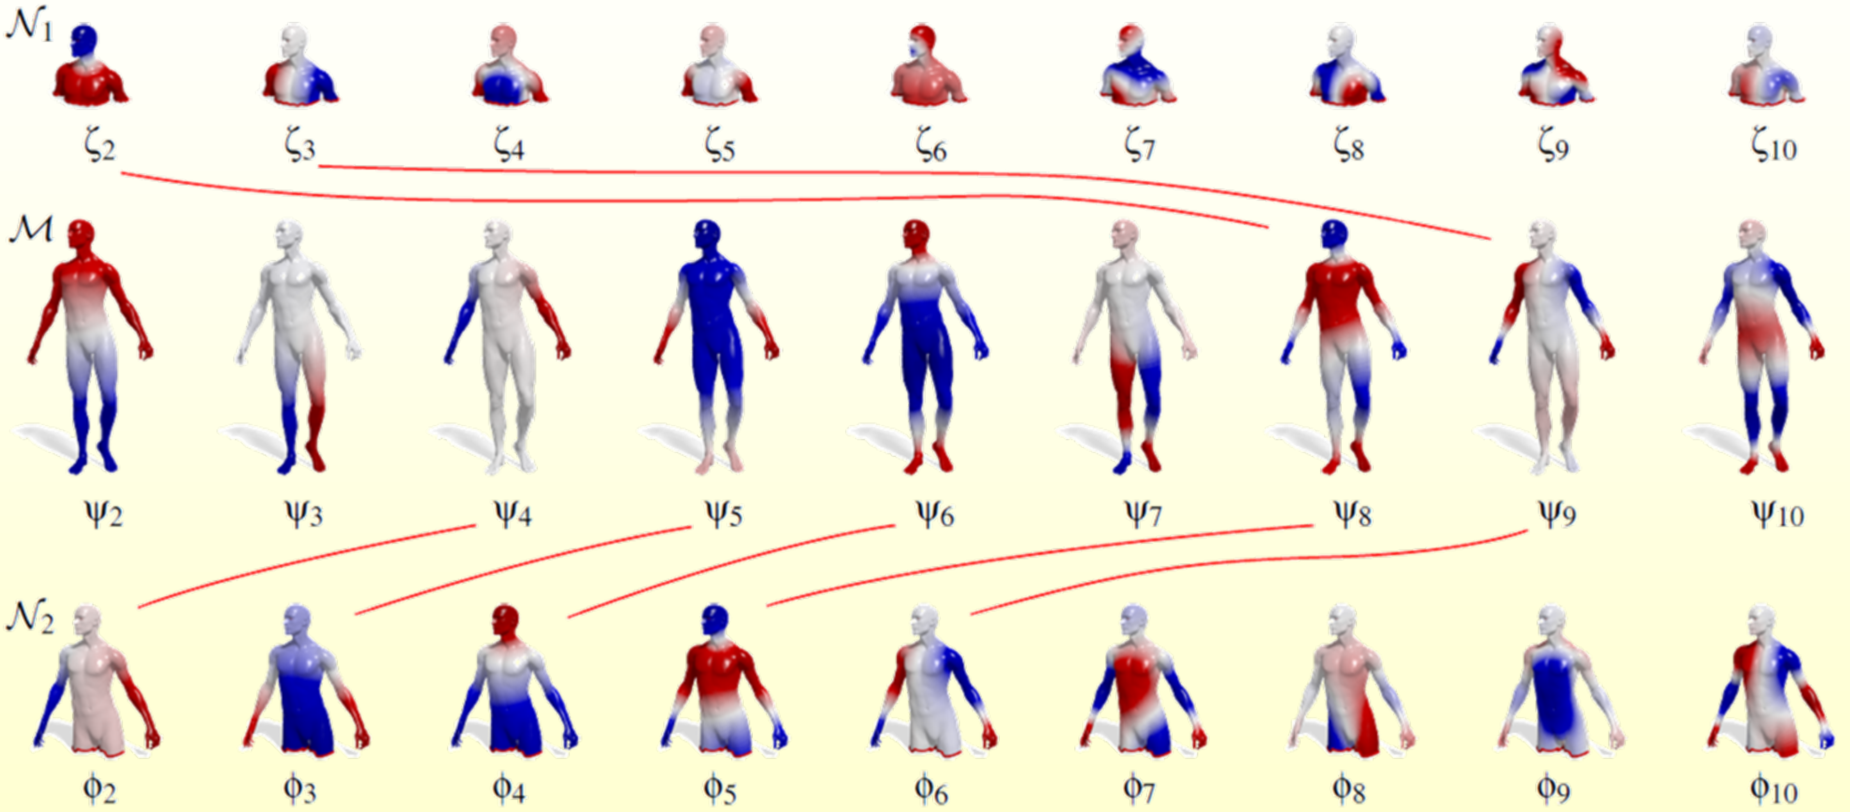

The Slope Rule

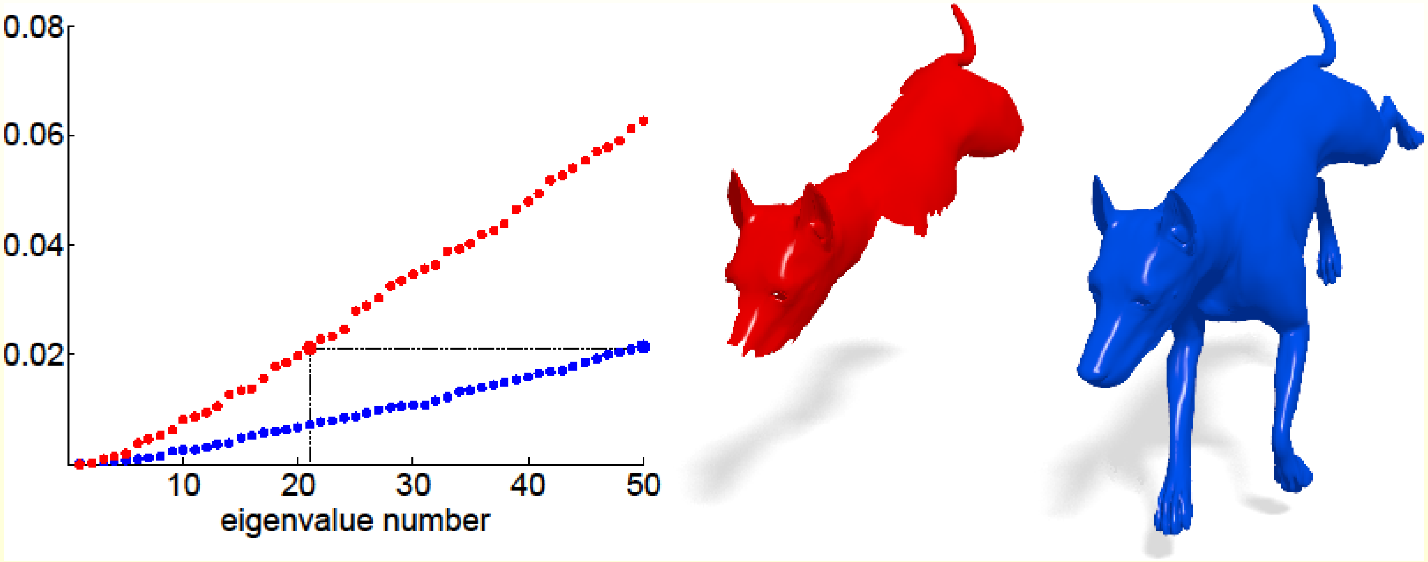

Weyl's law:

$$\lambda_j \approx \frac{1}{|S|} j$$

The Laplacian spectrum has slope inversely proportional to the surface area.

The Slope Rule

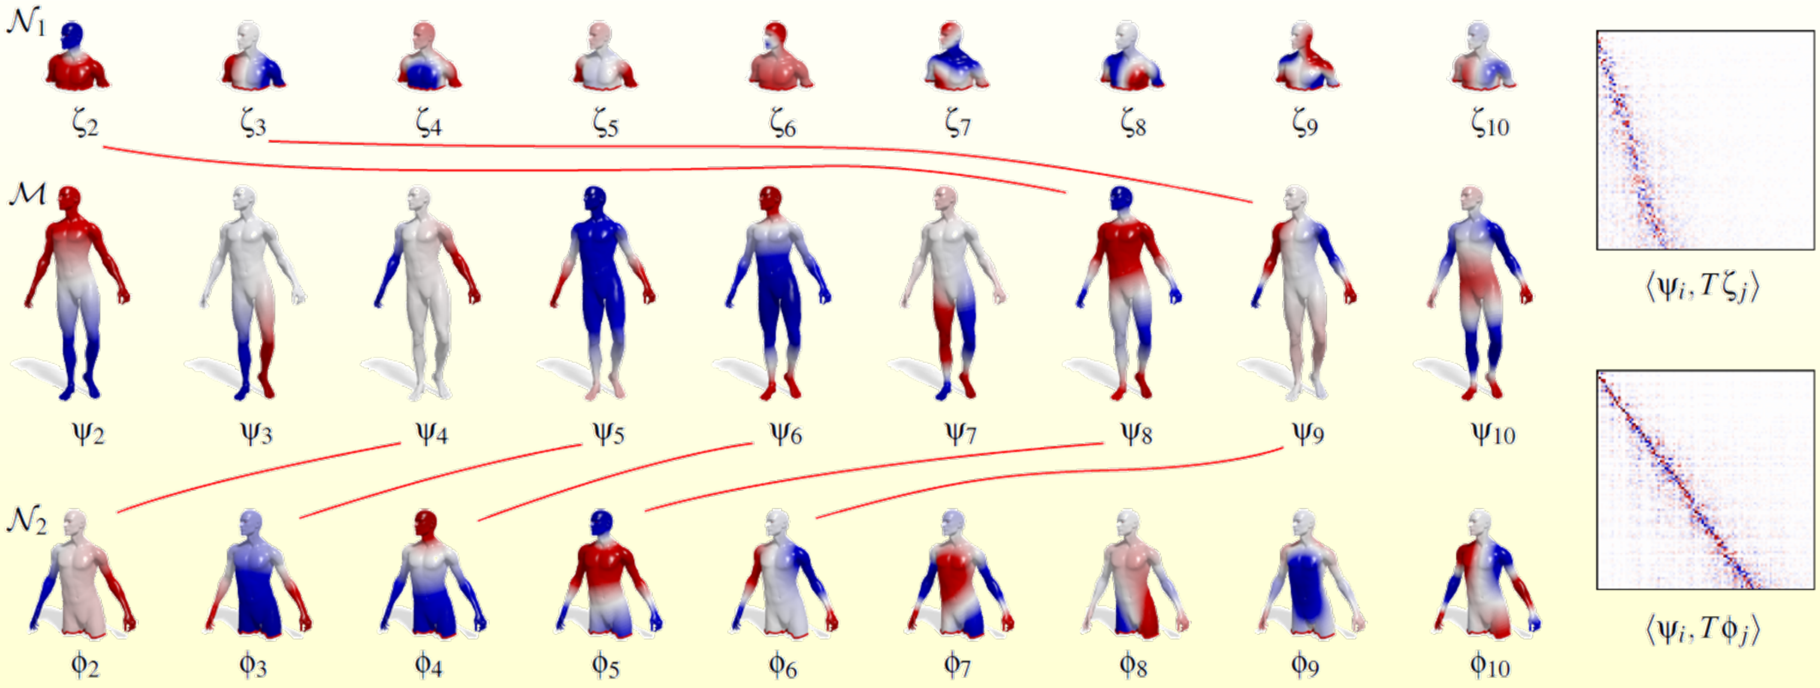

Partial Functional Maps

Descriptor and Subspace Learning

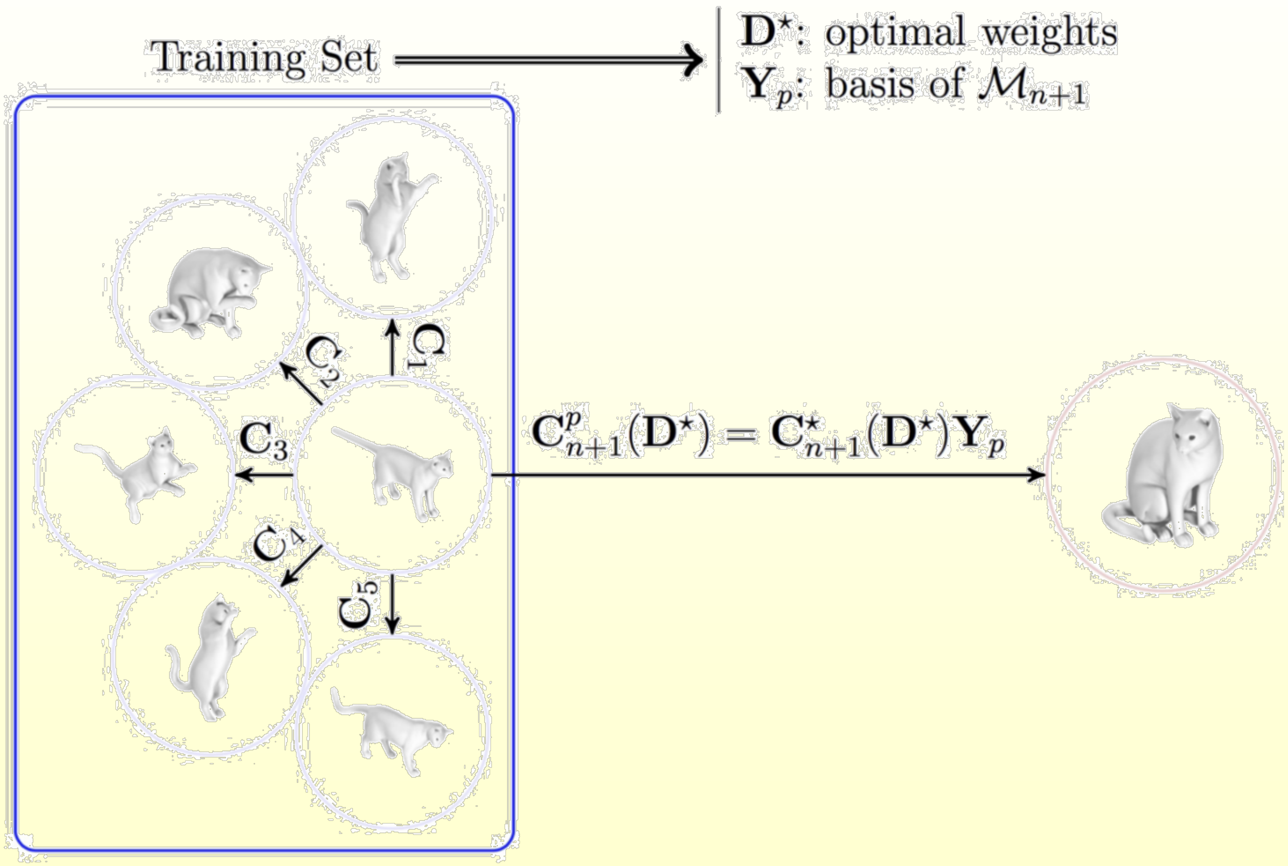

Supervised Learning. Given a collection of functional maps, a set of optimal weight $\mathbf{D}^\star$ and a basis of the p-best mapped functions $\mathbf{Y}_p$ are computed. When given a previously unseen shape (circled in red) we obtain an approximated functional map $\mathbf{C}^p_{n+1}(\mathbf{D}^\star)$ restricted to the most reliable function subspace.

Conclusion

- Many geometry processing tasks are best viewed as linear operators on functional spaces
- Operator composition, inversion and inference all lead to simple algebraic operations
- Using multiscale bases can improve compactness
- Performing spectral analysis on the operators can reveal the structure in a way that is easy to visualize

The End