UC San Diego

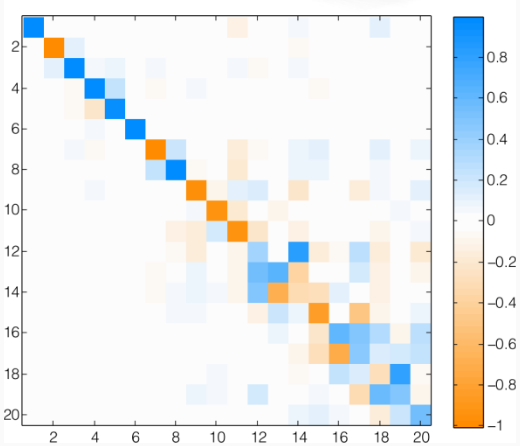

Lecture 16: Functional Maps

Instructor: Hao Su

Mar 6, 2018

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

Maps and Correspondences

Maps

$$\phi : X \longrightarrow Y$$

Map from *X* to *Y*

Maps and Correspondences

- Multiscale mappings
  - Point/pixel level
  - part level

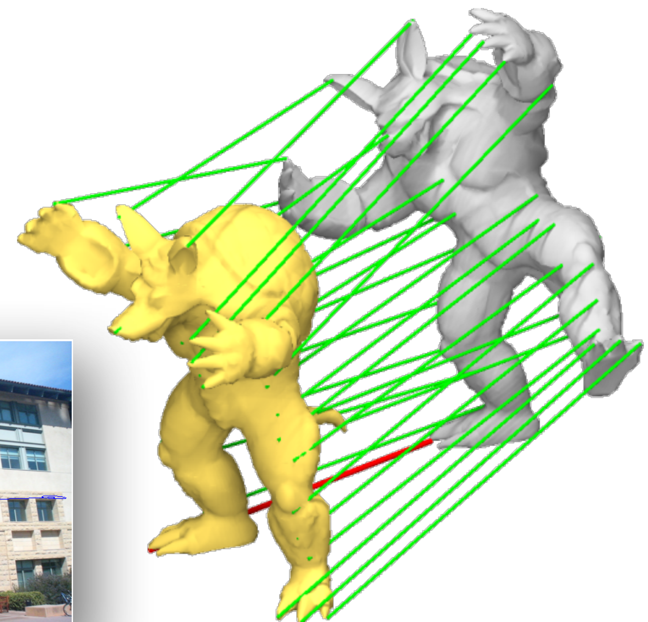

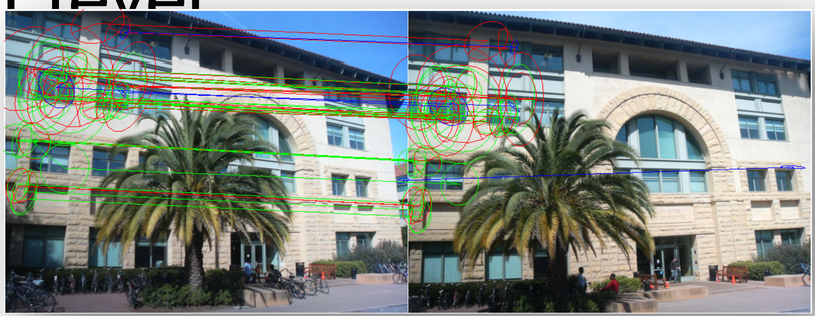

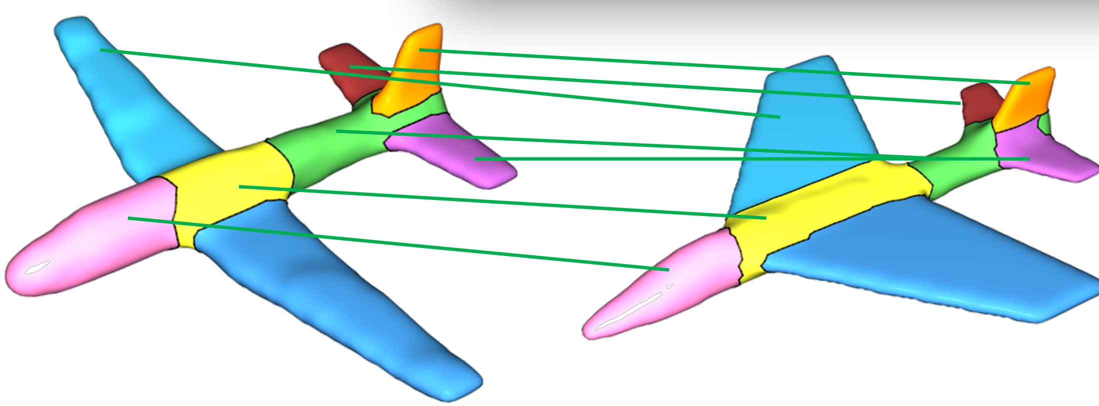

Maps capture what is the same or similar across two data sets

Map Composition

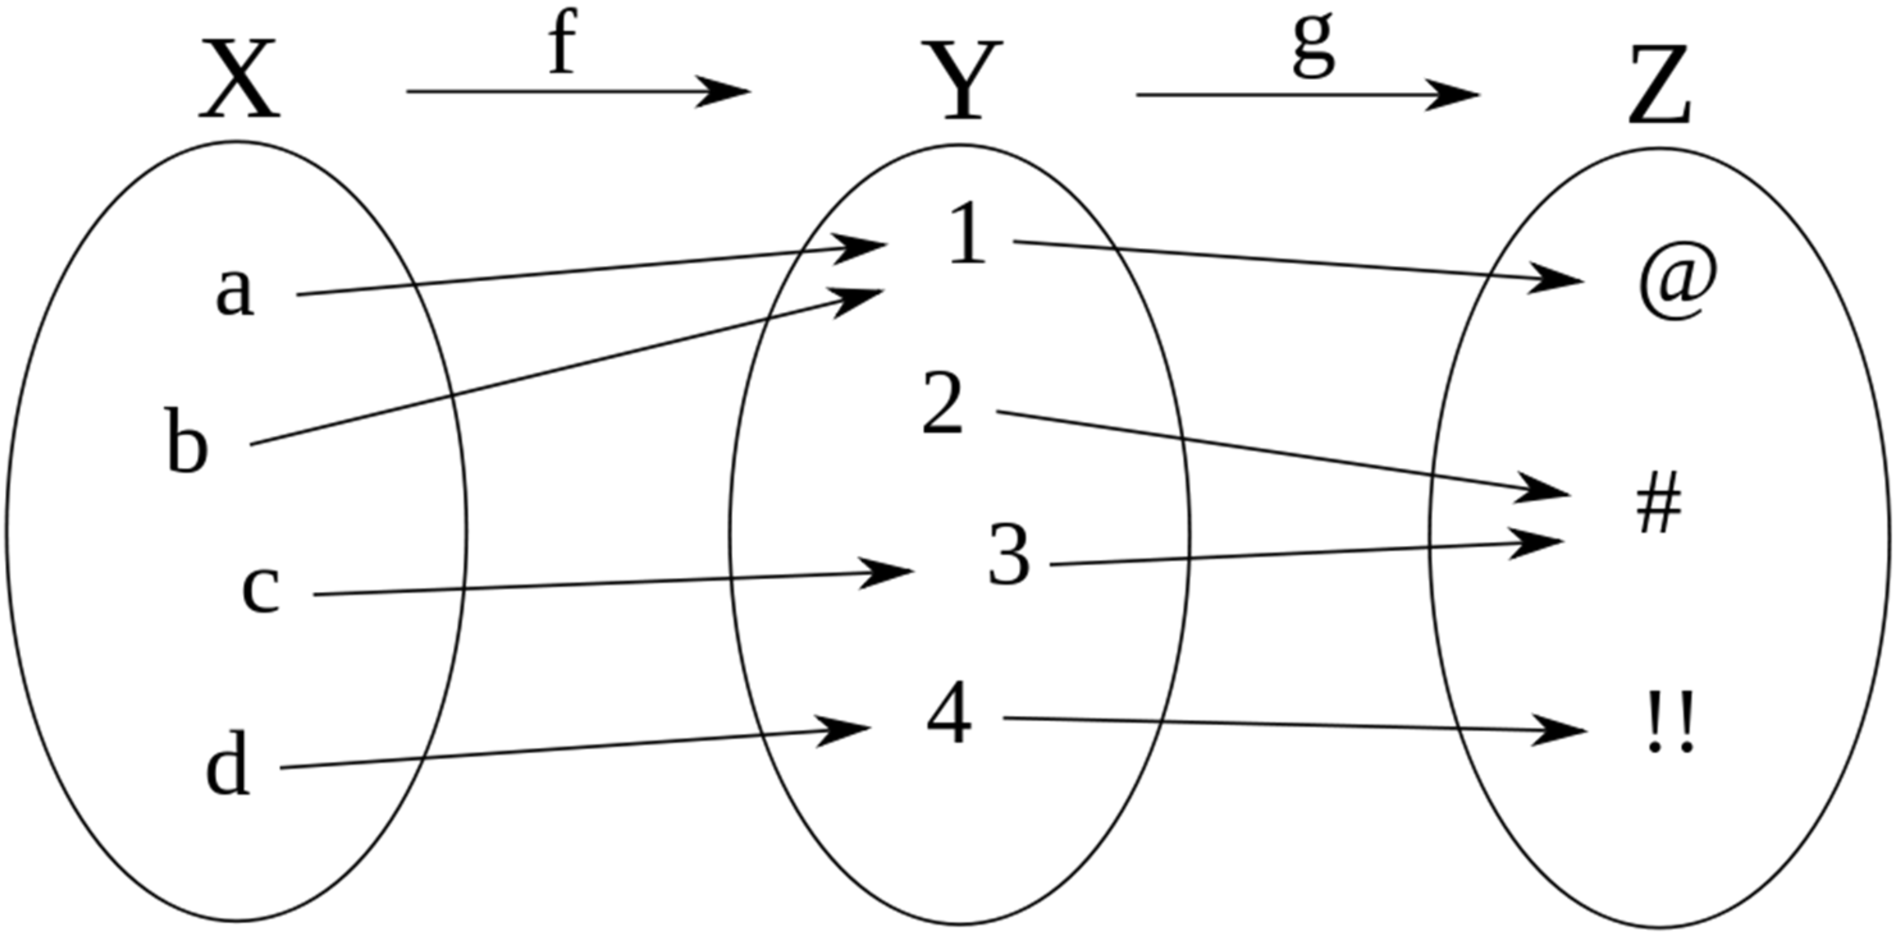

Problems and Issues

Symmetry, ambiguity, scale, bad data

Non-Convex, Combinatorial Optimization

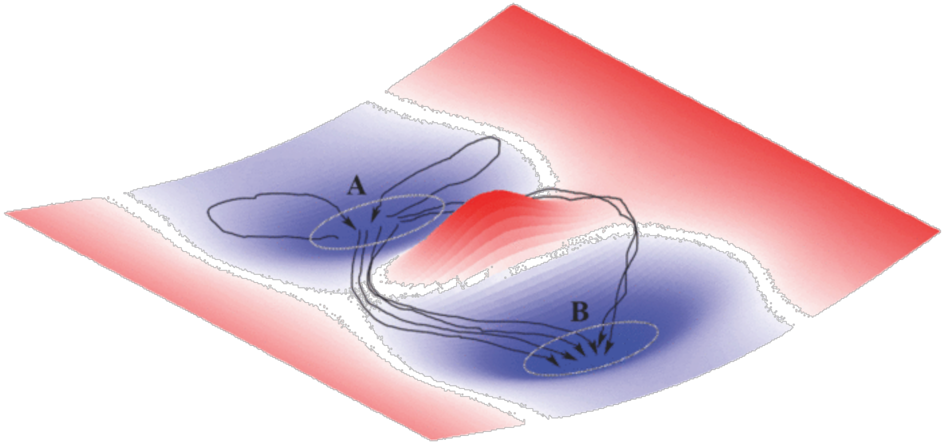

multiple minima

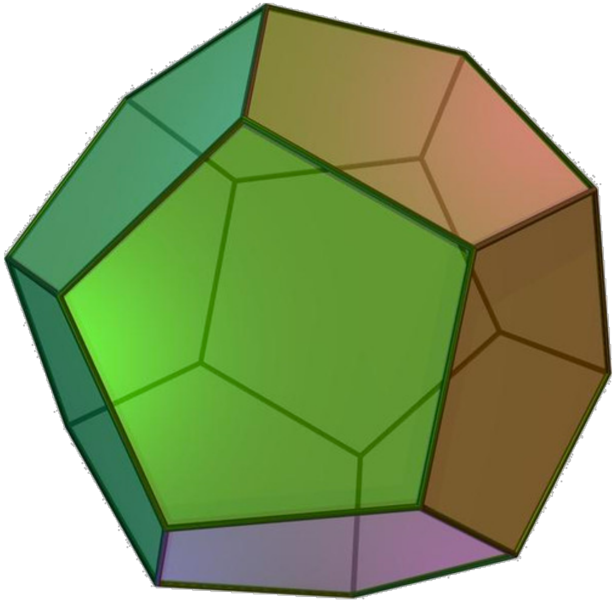

n! permutations

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

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

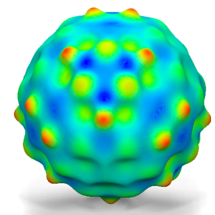

Curvature

Parts

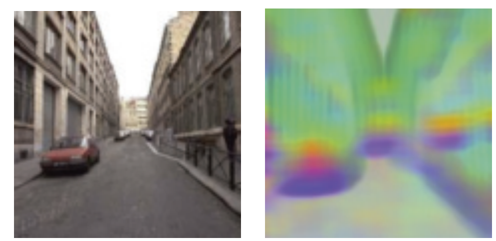

SIFT flow, C. Liu 2011

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

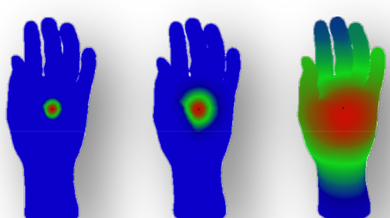

$$\frac{\partial u}{\partial t} = -\Delta u$$

heat diffusion

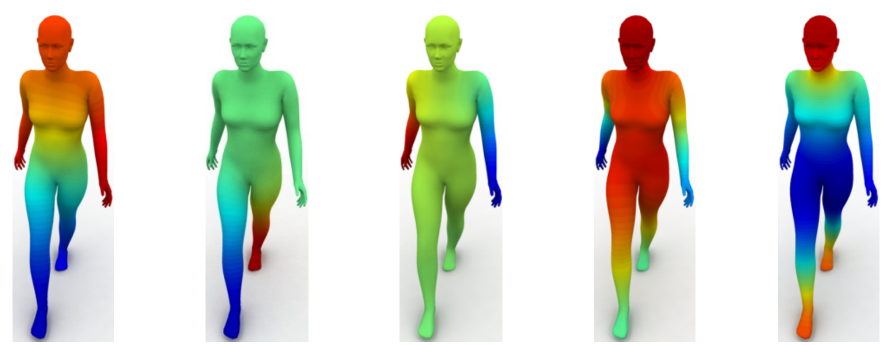

Laplace Beltrami eigenfunctions

Functional Maps

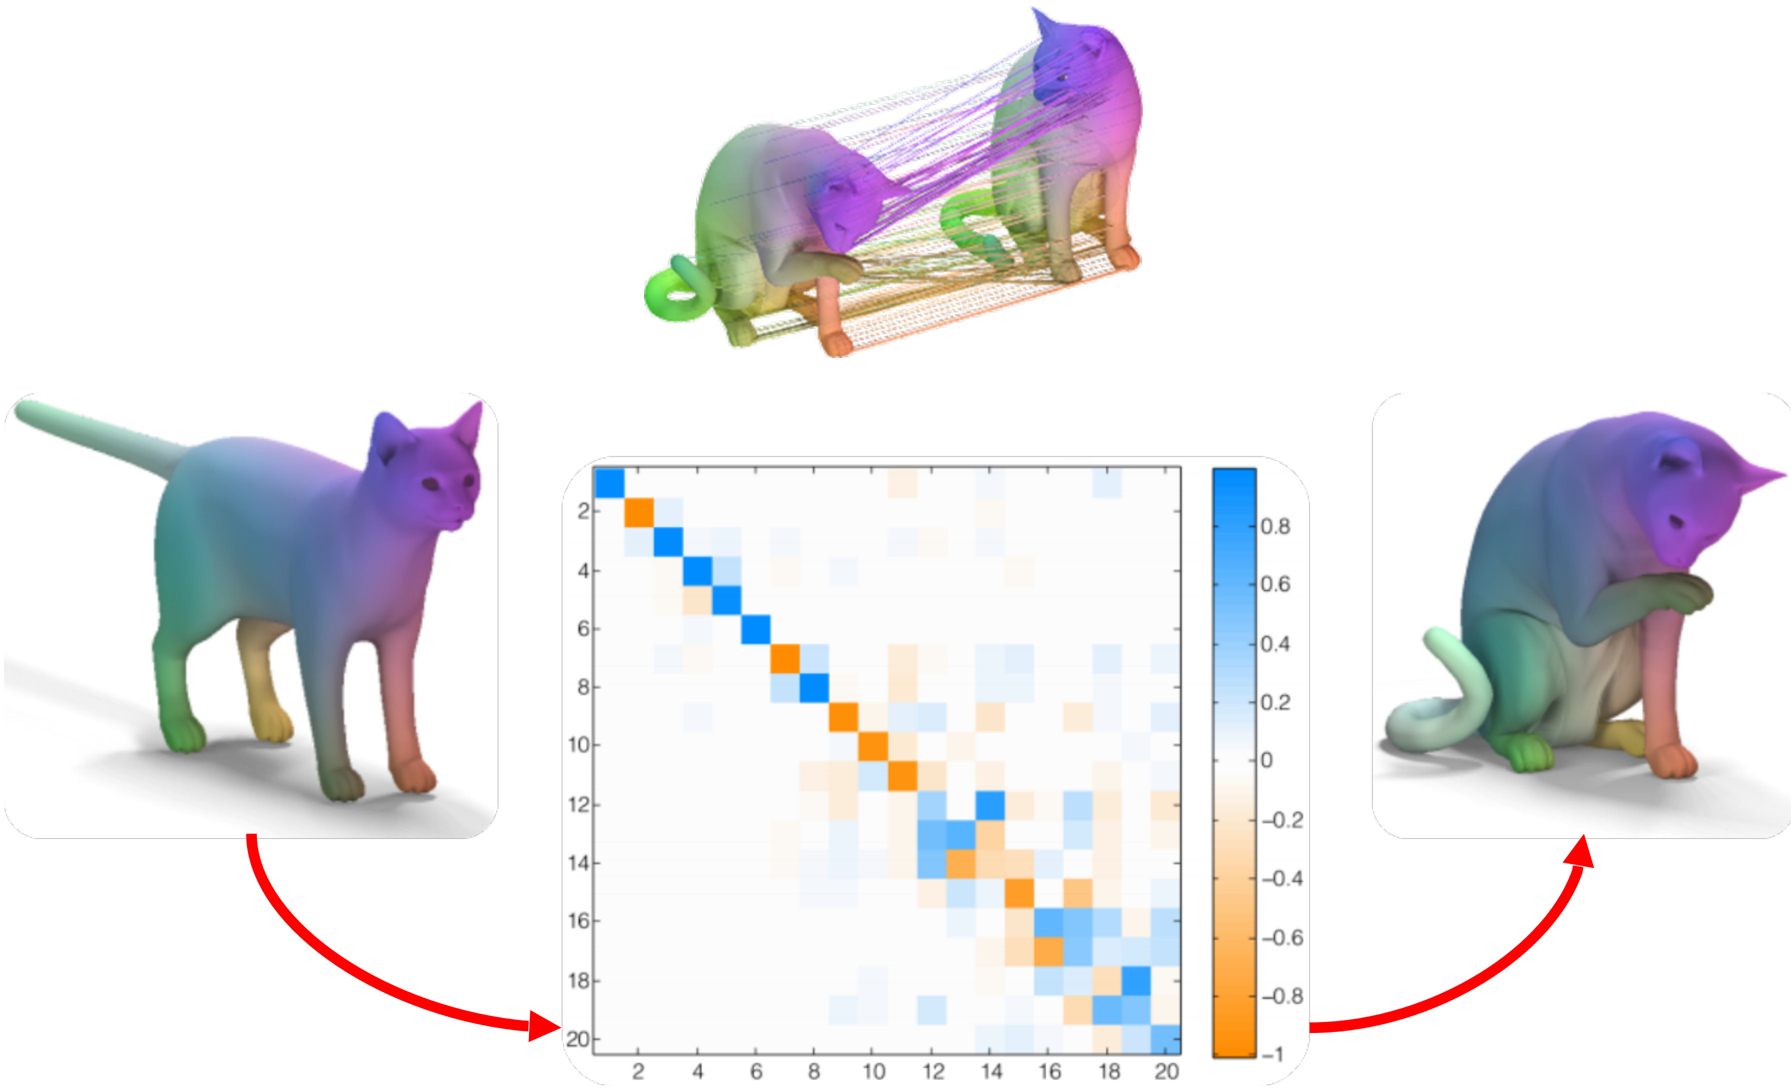

Starting from a Regular Map φ

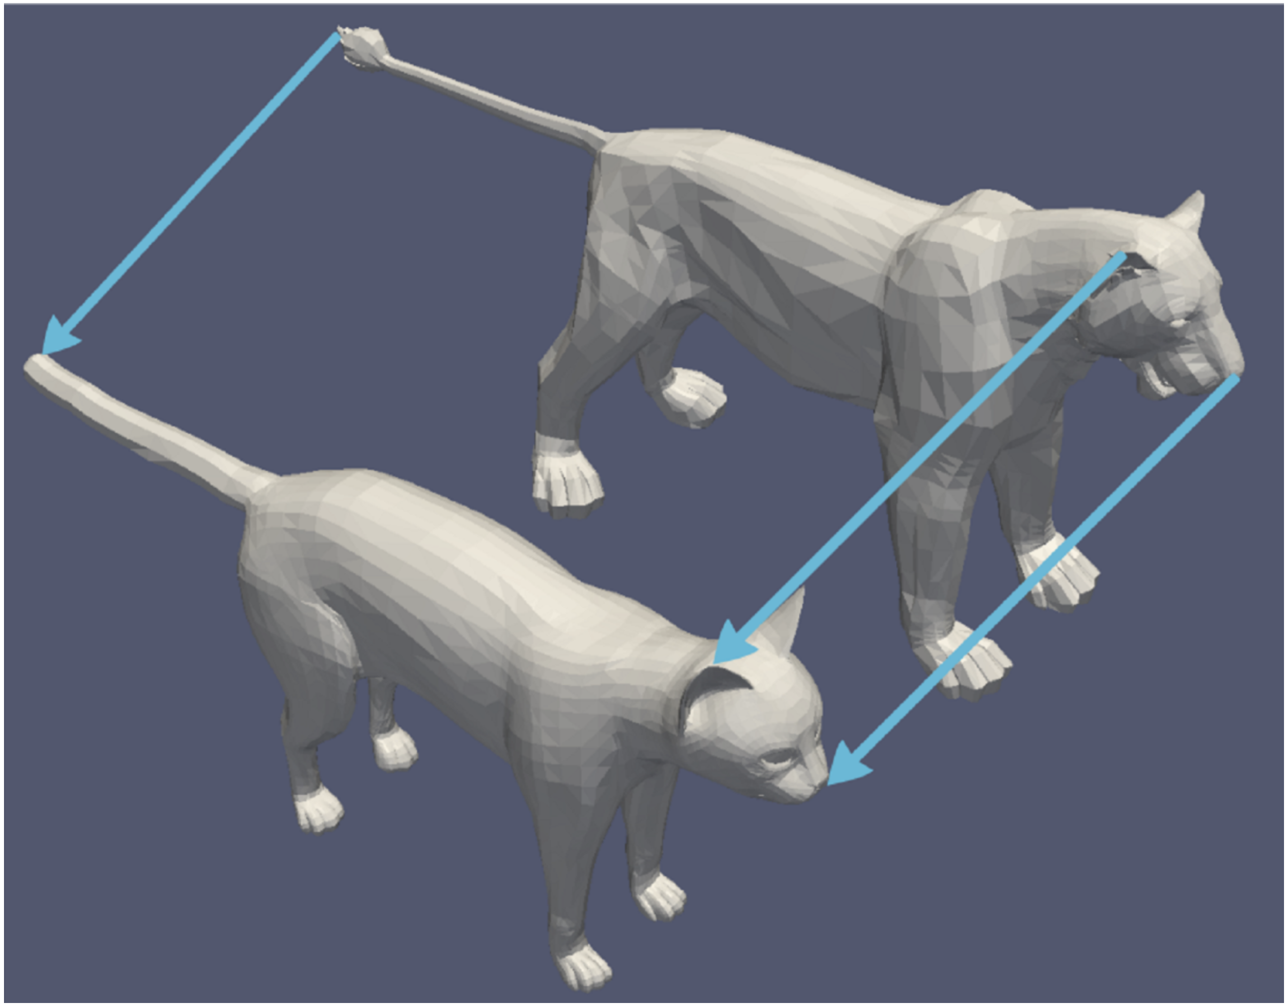

φ: lion → cat

Attribute Transfer via Pull-Back

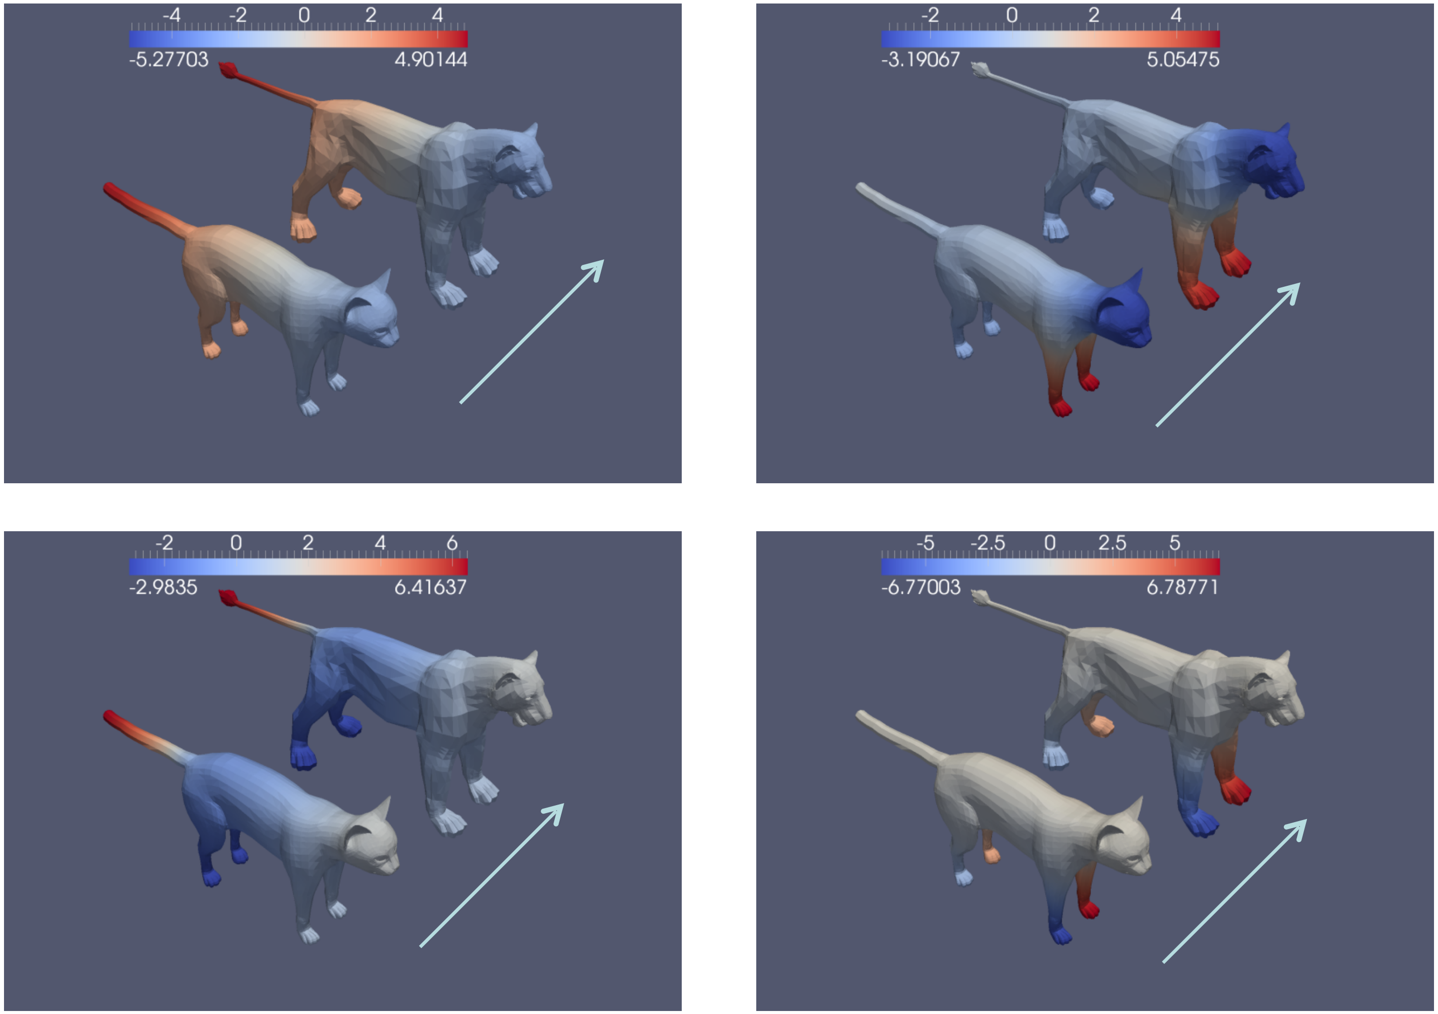

T_φ: cat → lion

A Contravariant Functor

from cat to lion

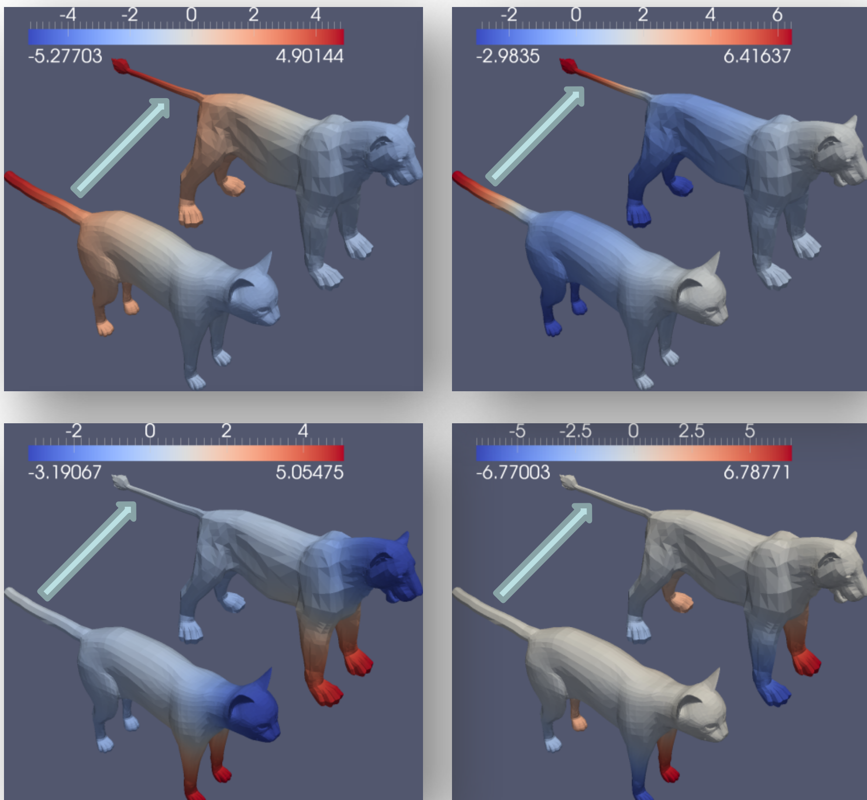

Functions on cat are transferred to lion using T_φ

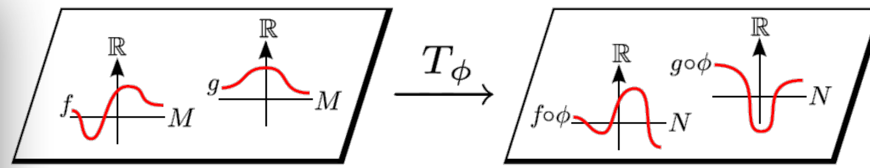

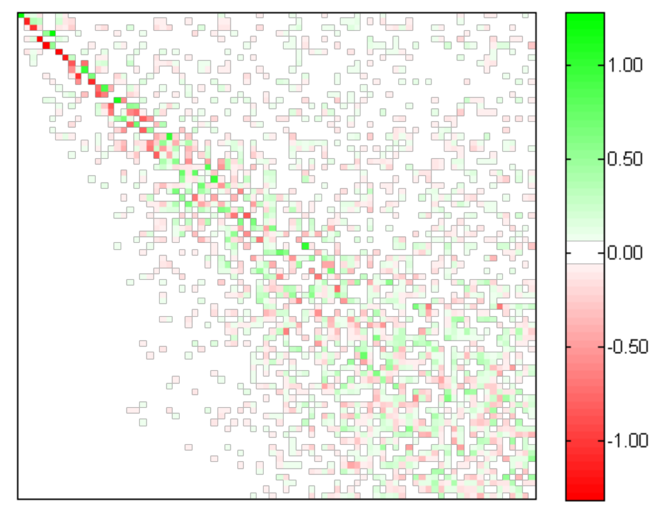

T_φ is a linear operator (matrix)

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

Functional Map

$$\boxed{\phi : M \to N}$$

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

Dual of a point-to-point map

Bases for a Function Space

Point basis
Finite-element basis

Local bases

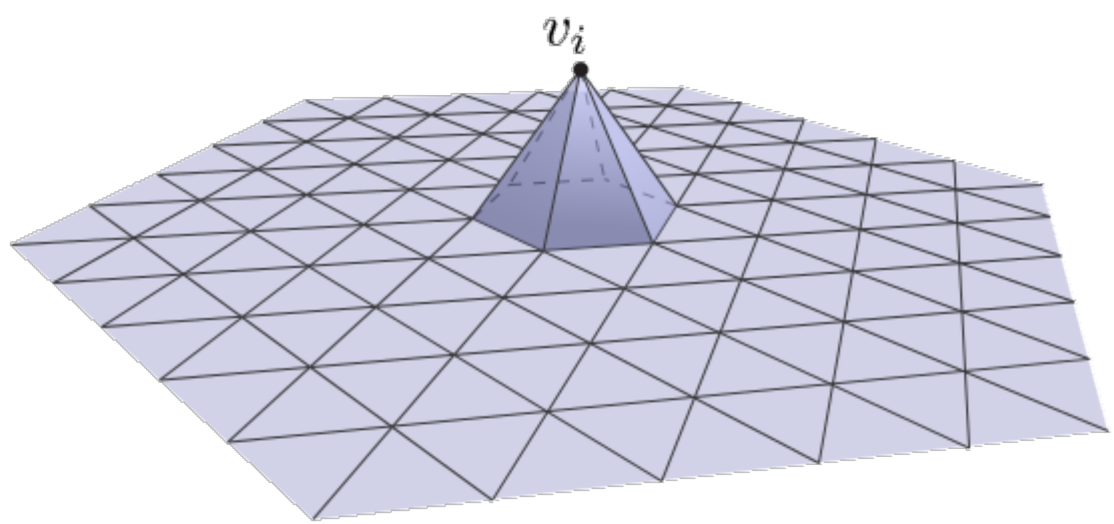

Bases for a Function Space

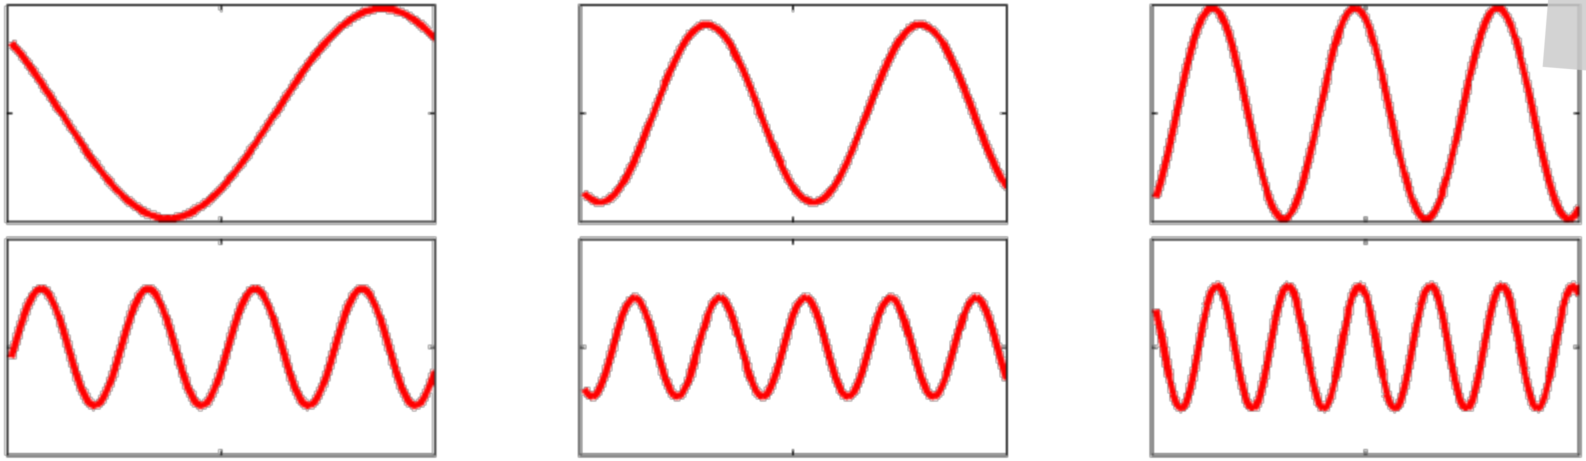

Fourier

Laplace-Beltrami

global support

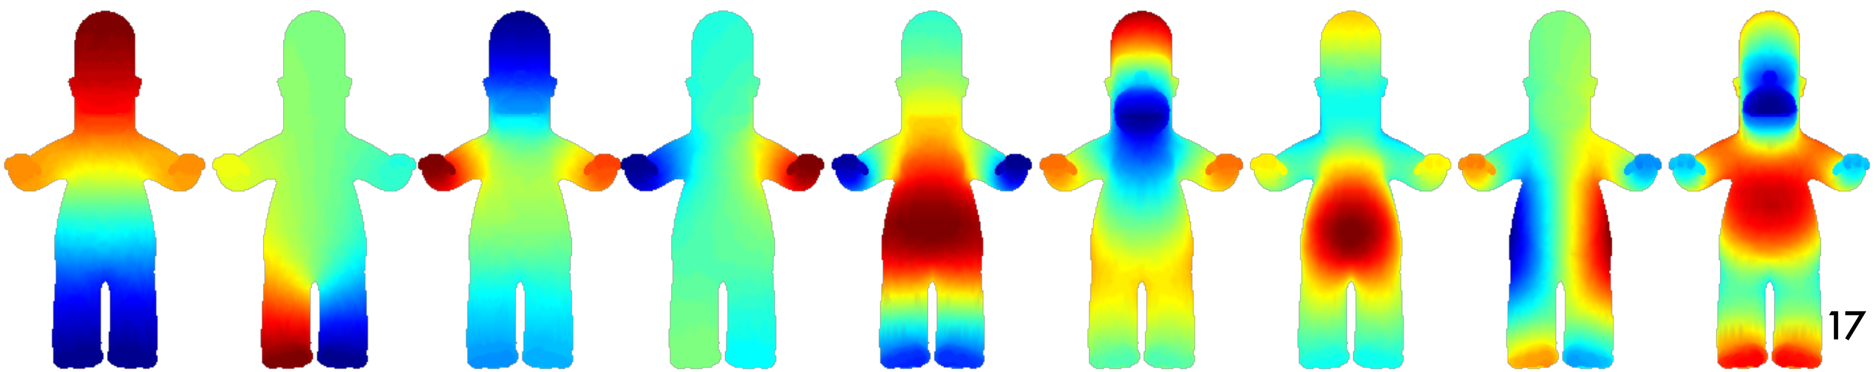

More Exotic Bases Possible

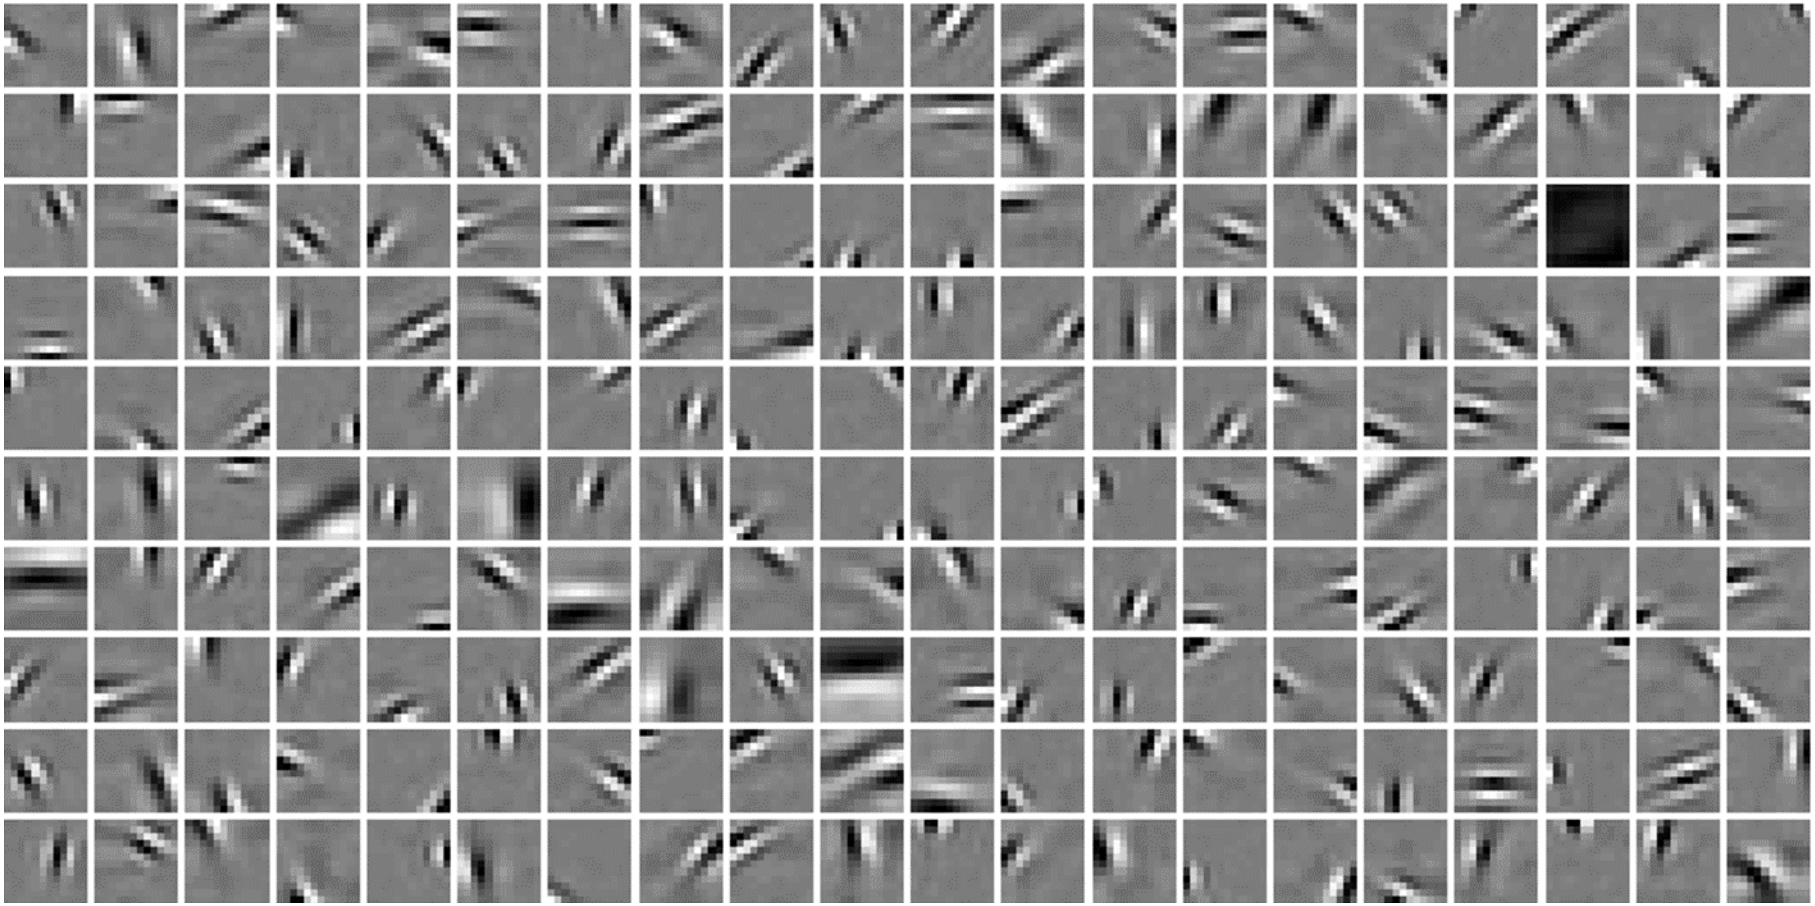

Textons, wavelets, …

Exploit Linearity

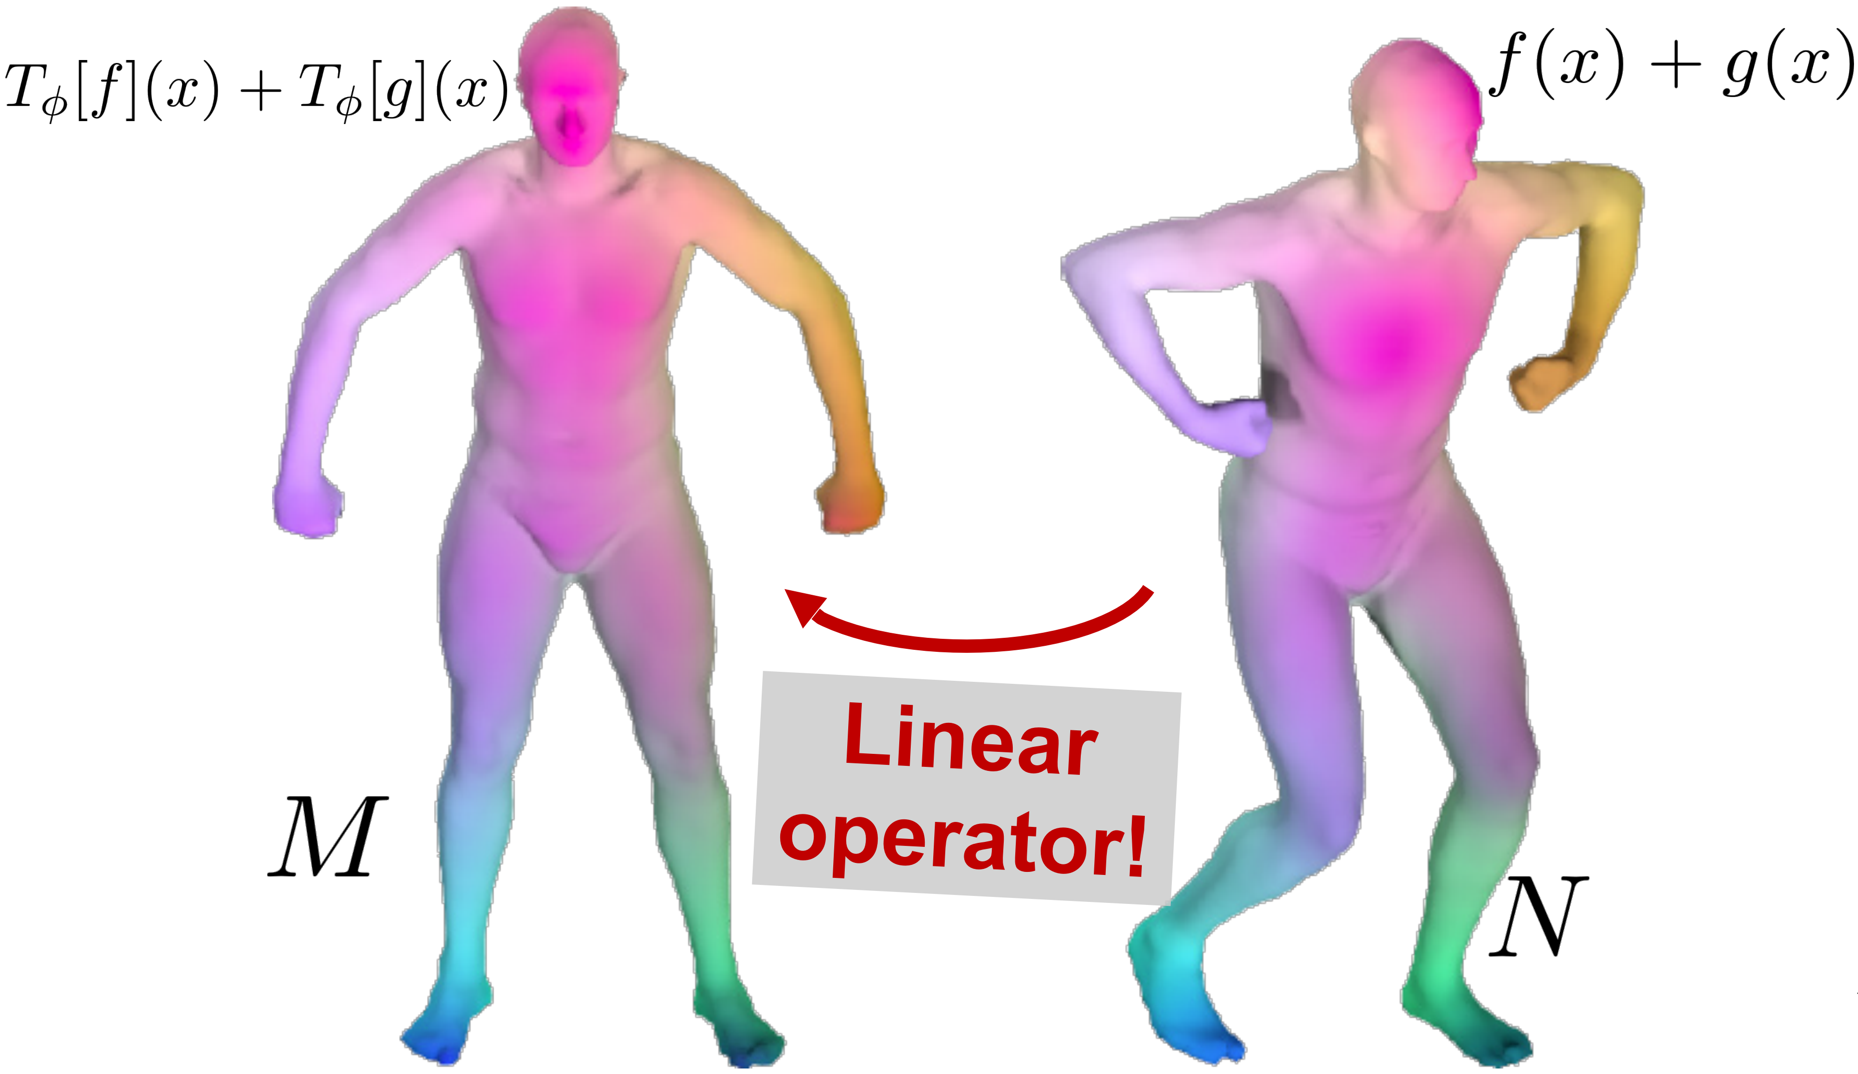

Application of Basis

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

$\phi :$ 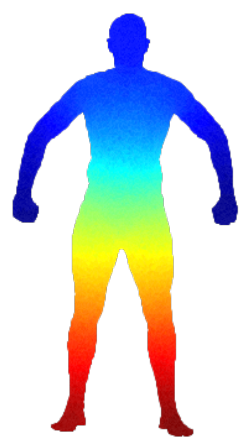 $\longrightarrow$ 

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

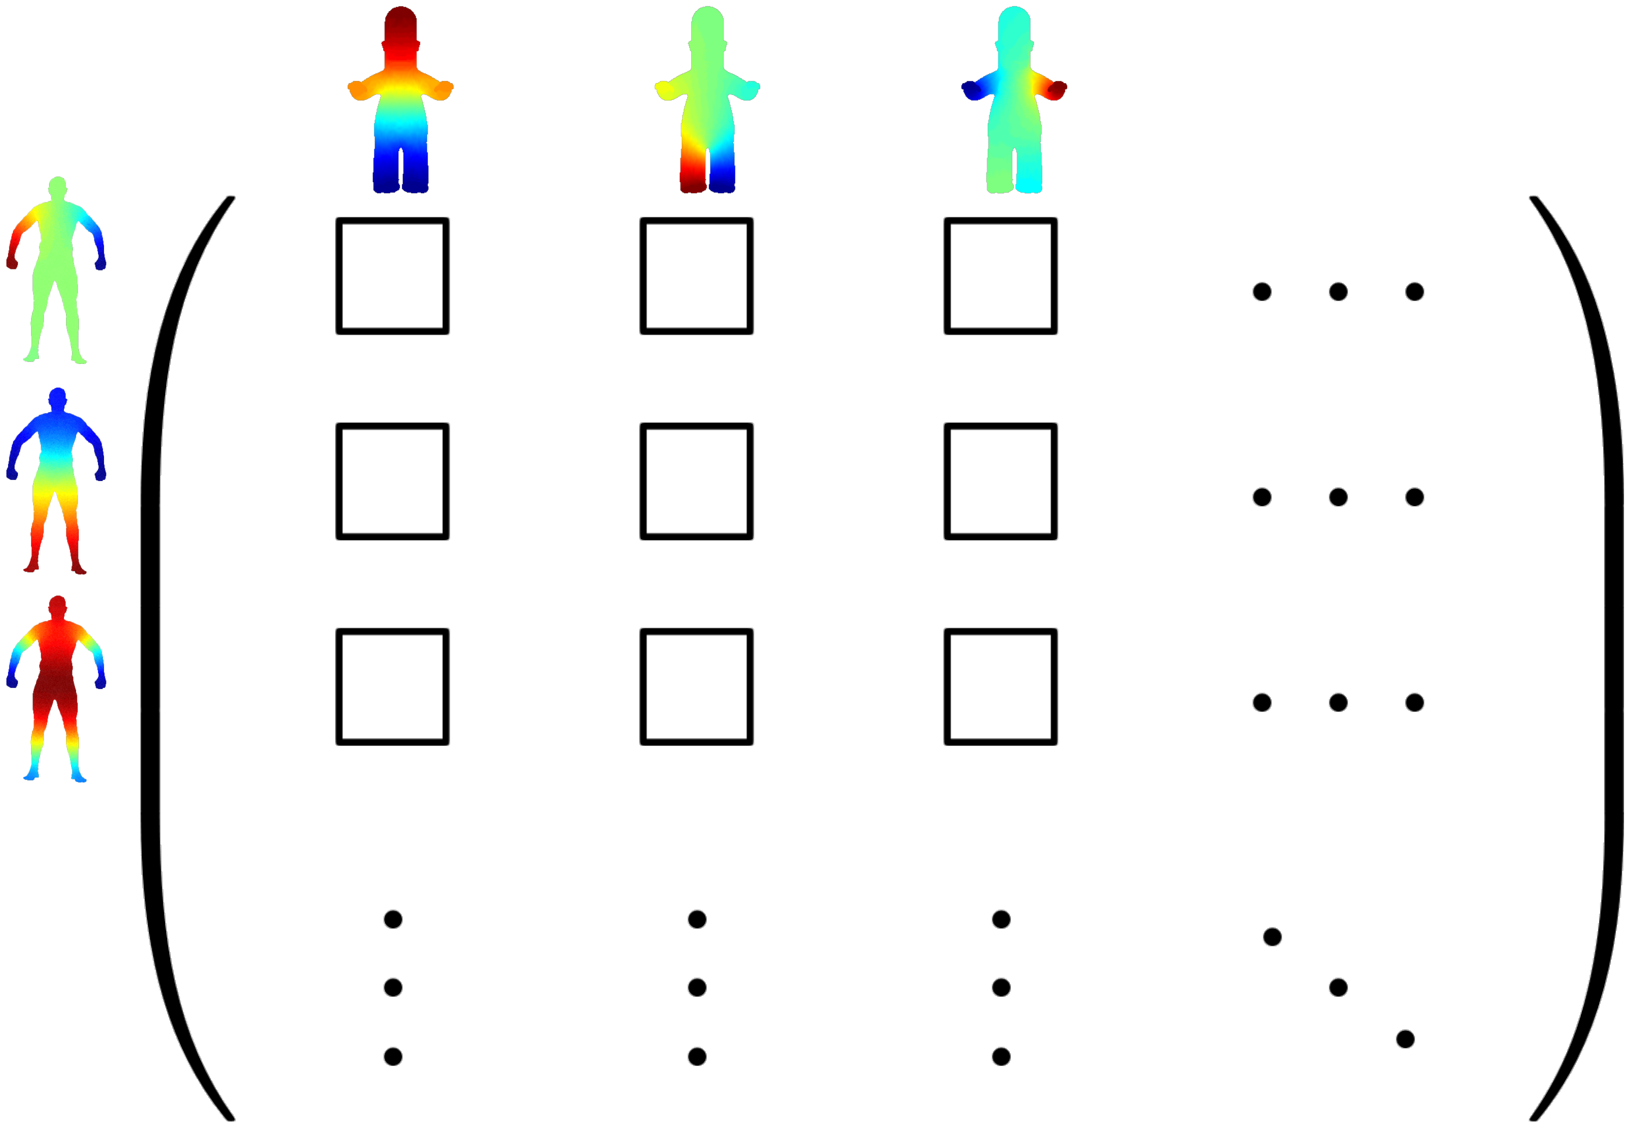

Functional Map Representation

Definition

For a fixed choice of basis functions $\{\phi^M\}$ and $\{\phi^N\}$, and a bijection $T : M \to N$, 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

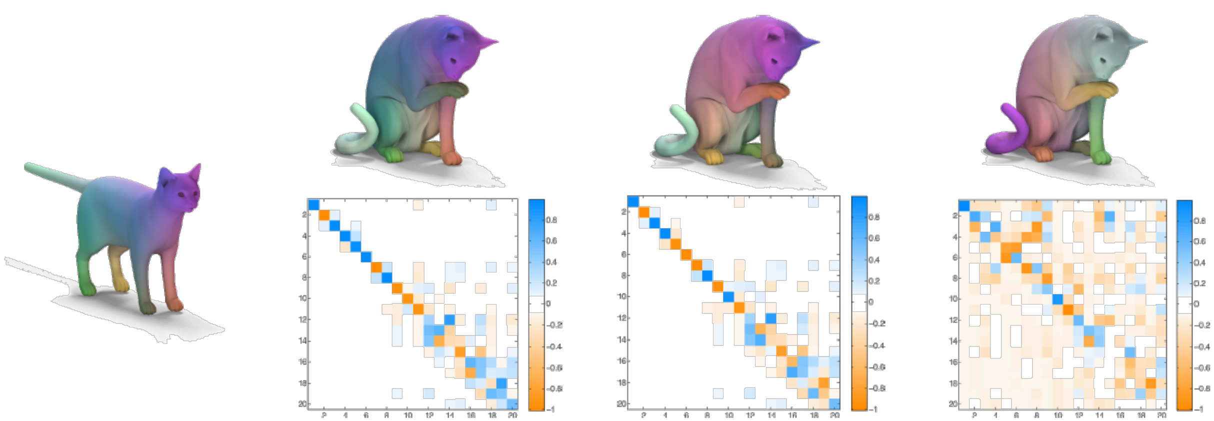

source — direct — symmetric — head to tail

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

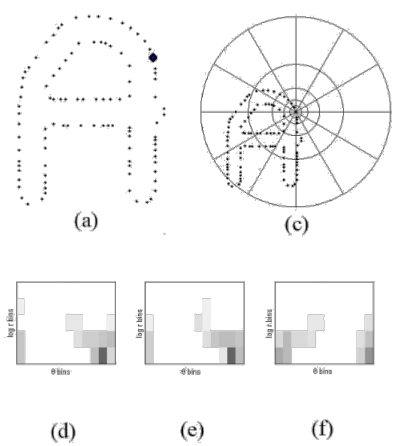

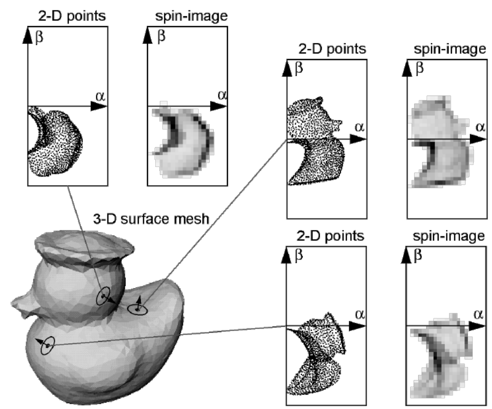

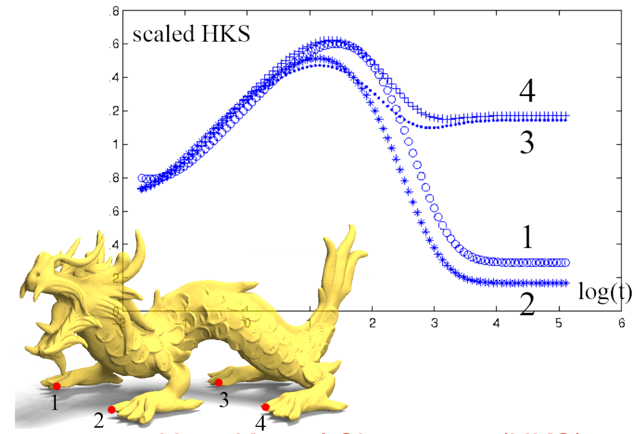

Heat Kernel Signatures (HKS):
[Sun, Ovsjanikov, G. ’08]]

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

Spin Images:
[Johnson, Hebert ’99]

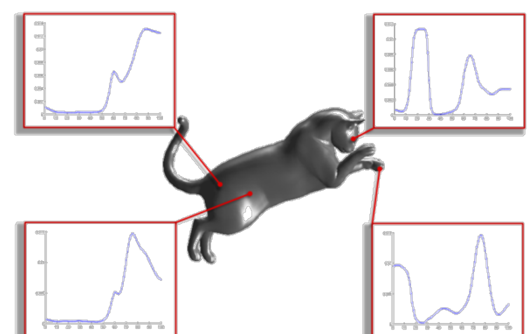

Rigid invariance
(extrinsic)

Isometric invariance
(intrinsic)

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

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 $S_2 : \mathcal{F}(N, \mathbb{R}) \to \mathcal{F}(N, \mathbb{R})$.

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

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

$$\boxed{C\Delta_1 \approx \Delta_2 C}$$

Differentiate and then transport

Transport and then differentiate

Isometry Regularizer

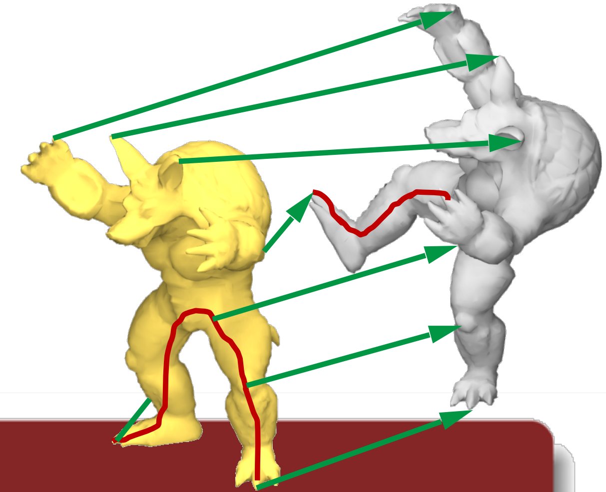

Lemma 1:

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.

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

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.

Sparcity in a Localized Basis

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

Sum of Euclidean norms of cols

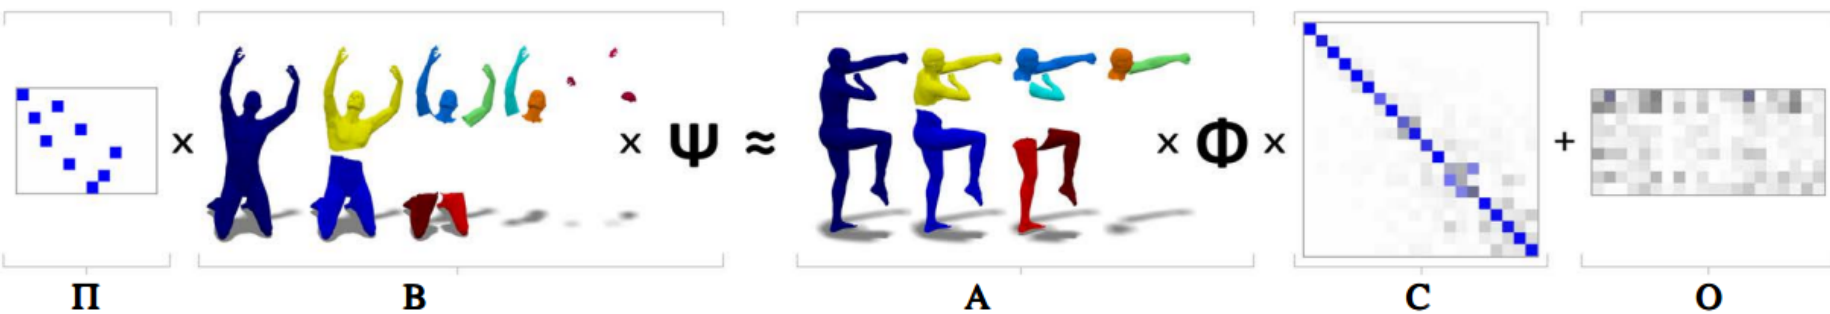

Sparse Modeling of Intrinsic Correspondences (Pokrass, Bronstein2, Sprechmann, Sapiro)

General Optimization for Maps

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

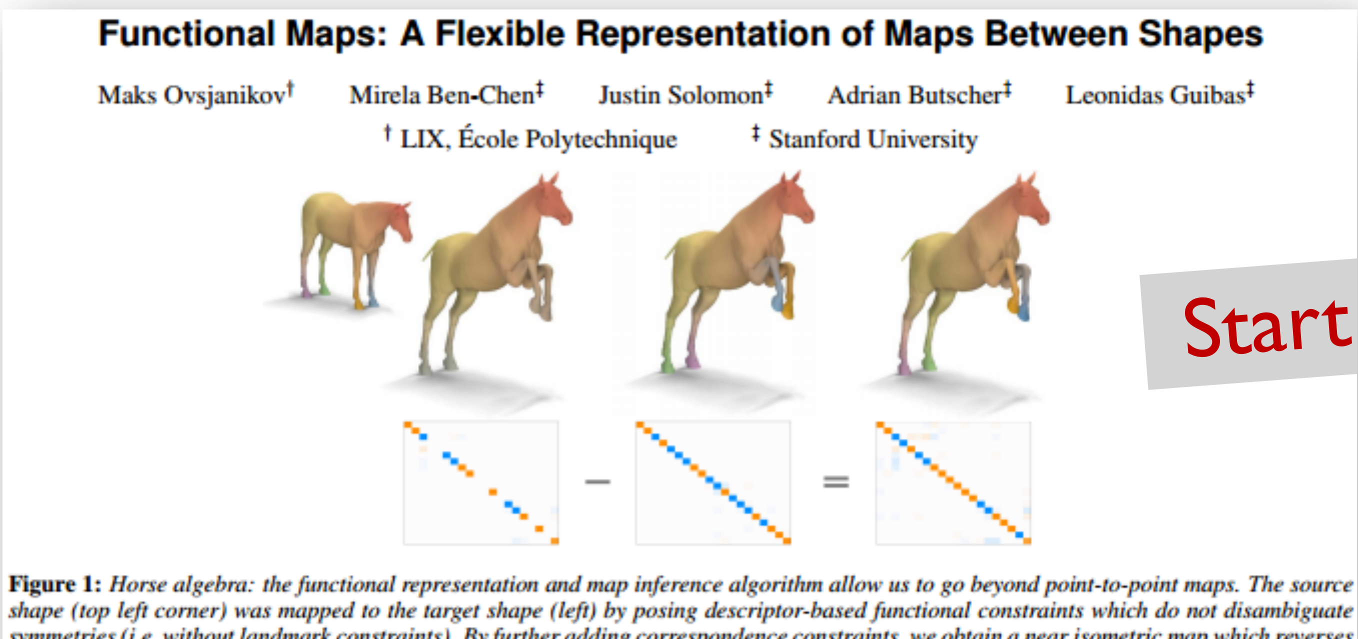

Functional Maps: A Flexible Representation of Maps Between Shapes

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 symmetries (i.e. without landmark constraints). By further adding correspondence constraints, we obtain a near isometric map which reverses*

Start here

Map Continuity

- Not explicitly enforced
- Implicit in the choice of basis

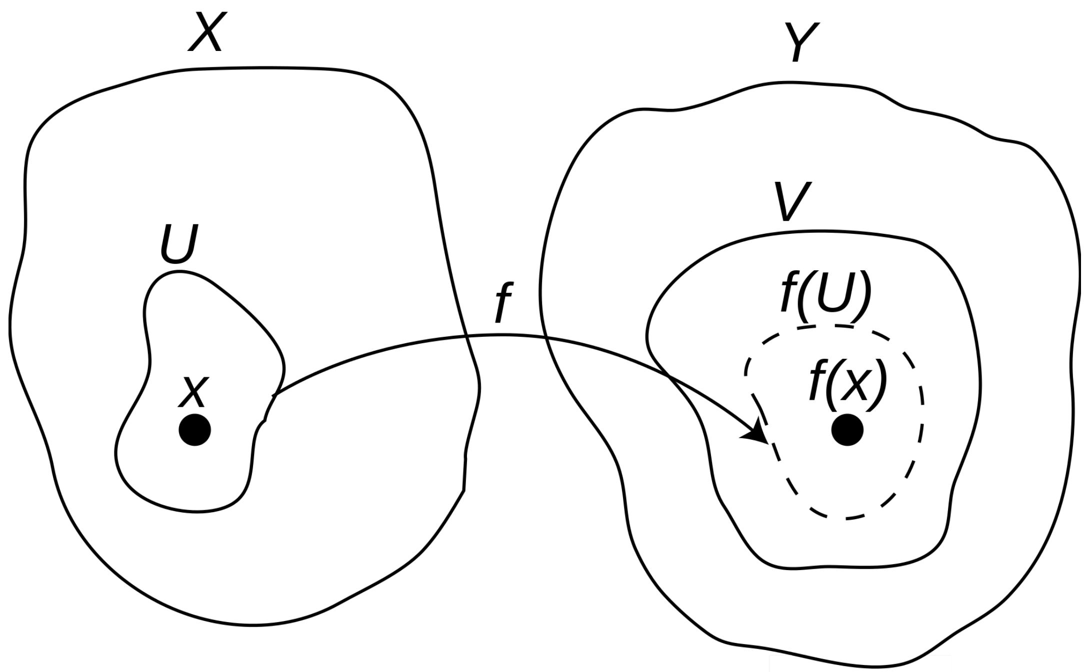

From Functional to Point-to-Point Maps

- Can try transporting delta functions individually -- expensive

Dual

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

Primal

$$p \mapsto q$$

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

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.

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

Functional Maps

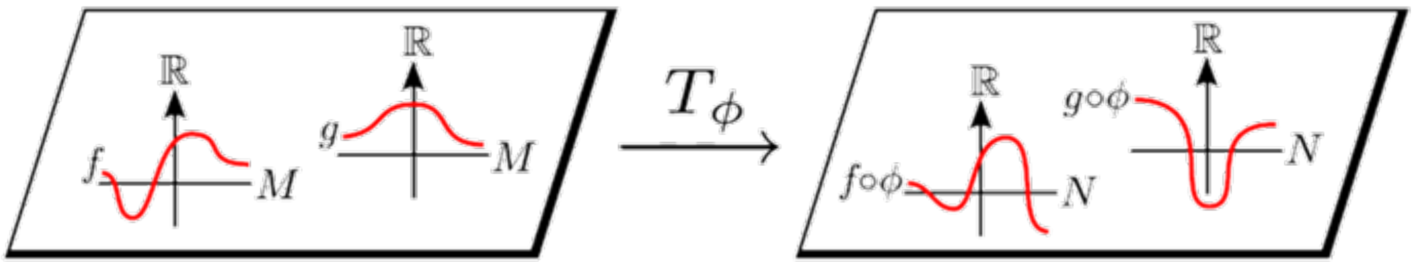

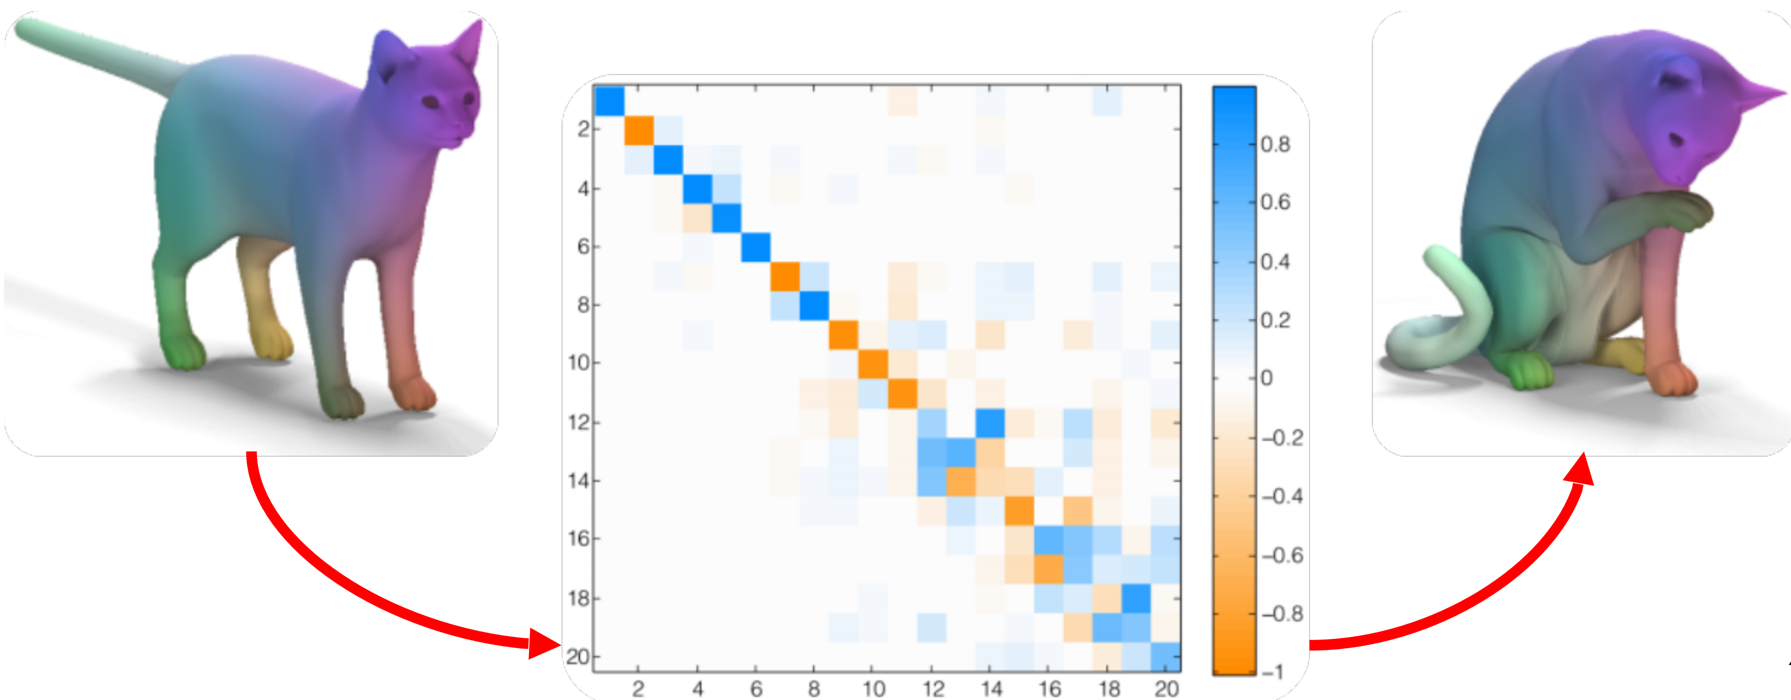

Joint Data Analysis

Joint Data Analysis

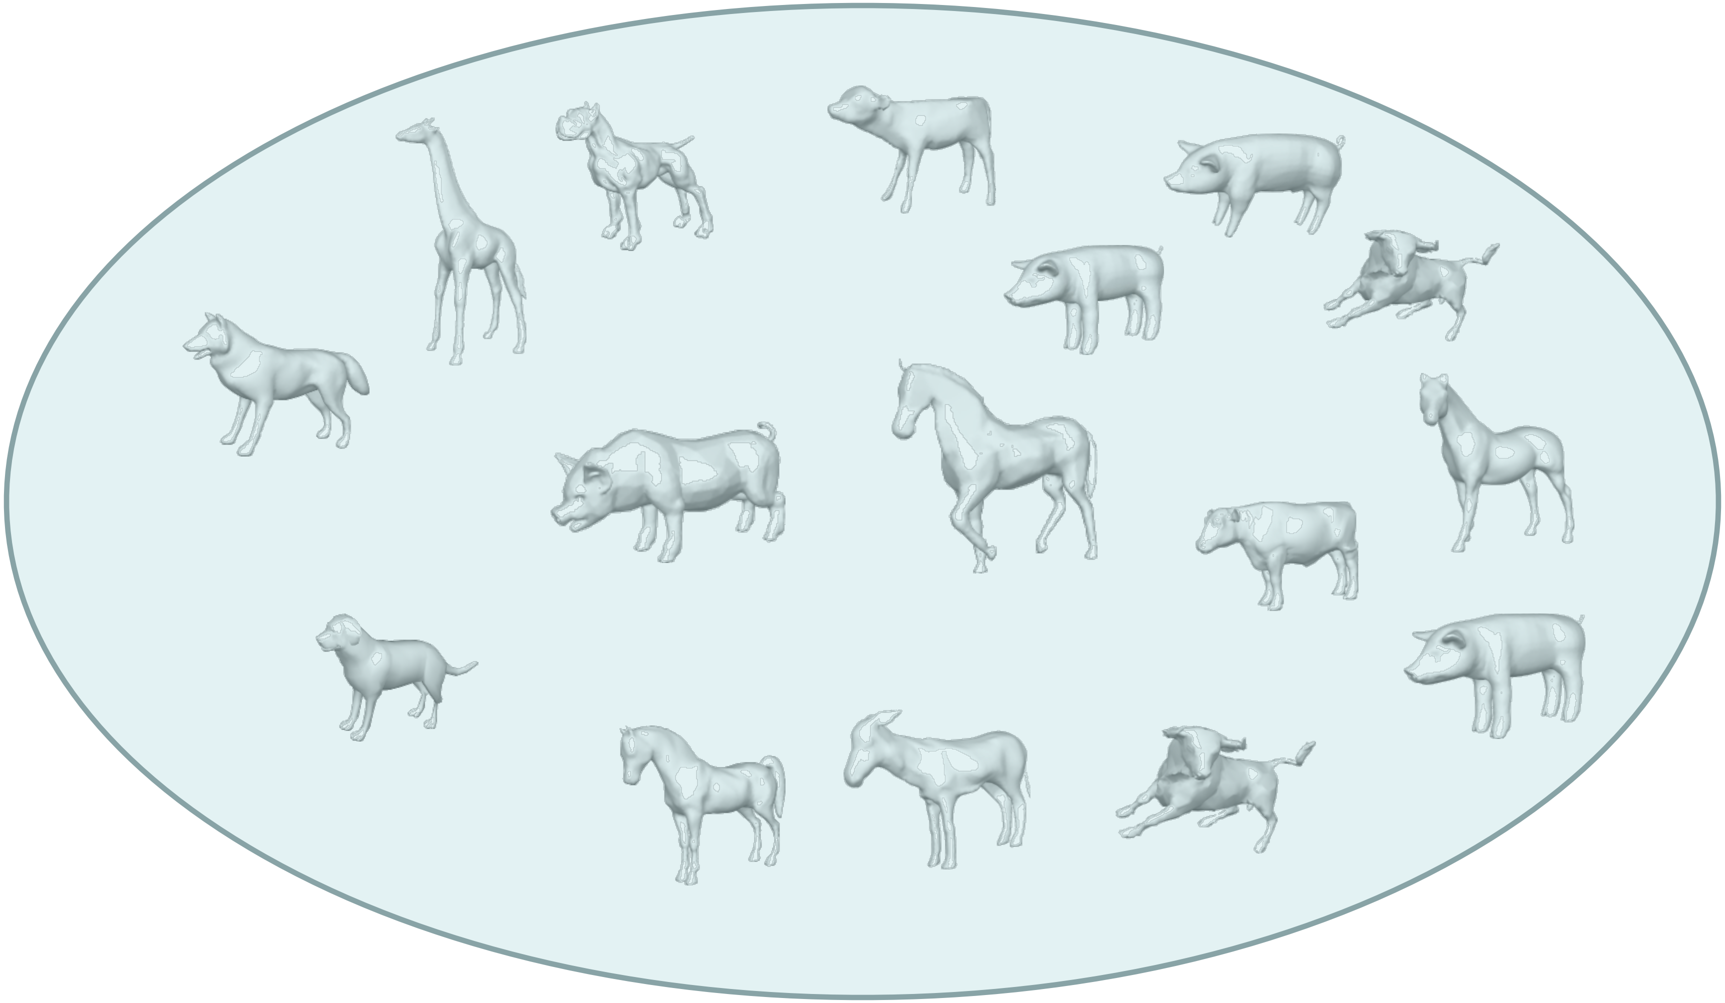

Maps Join Data Together

Individual Maps Can Have Errors

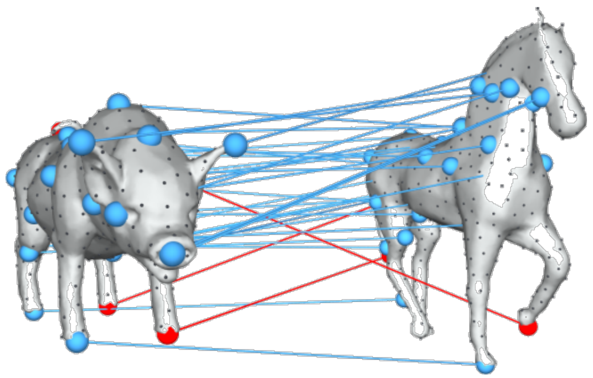

Blended intrinsic maps
[Kim et al. 11]

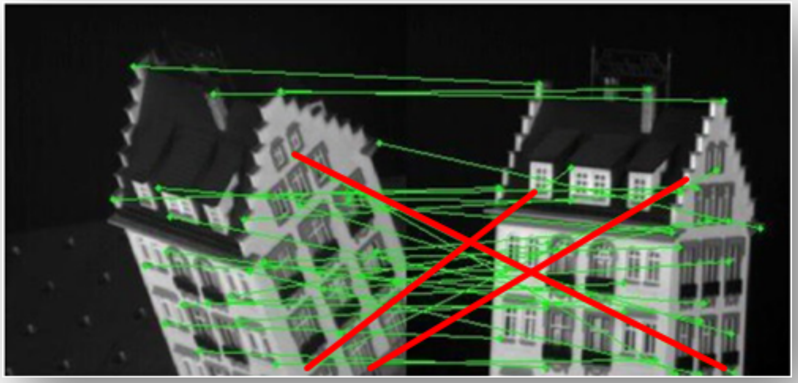

Learning-based graph matching
[Leordeanu et al. 12]

State-of-the-art techniques

Wrong correspondences

Combining Maps

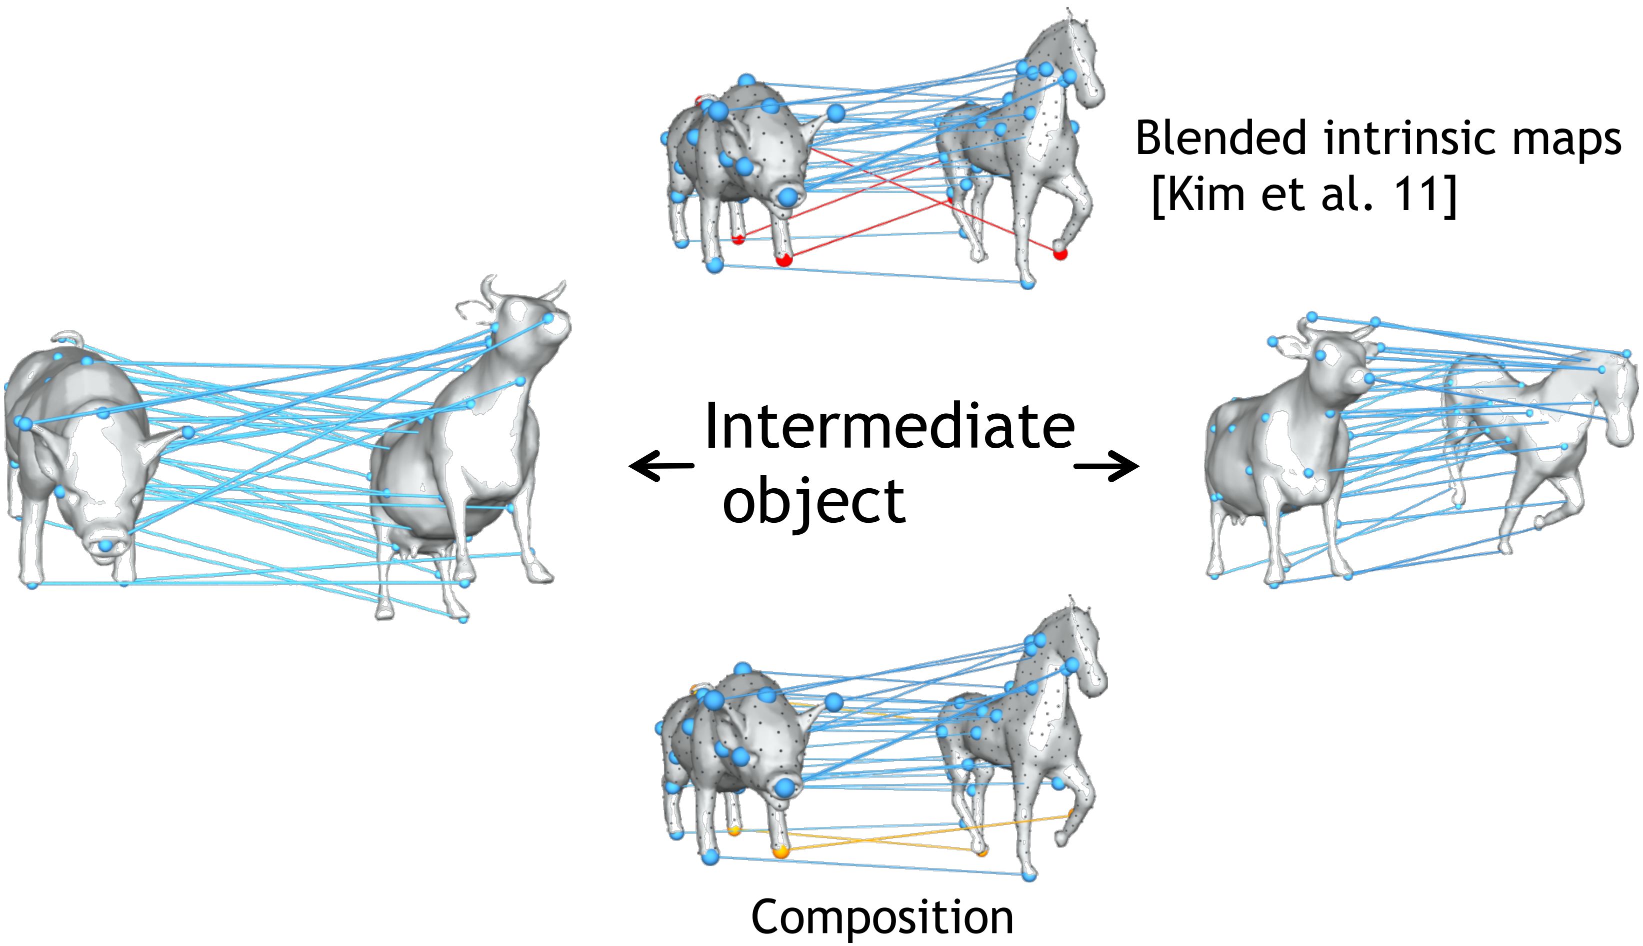

Composition can correct correspondences

Individual Data Set Operations Can Have Errors

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

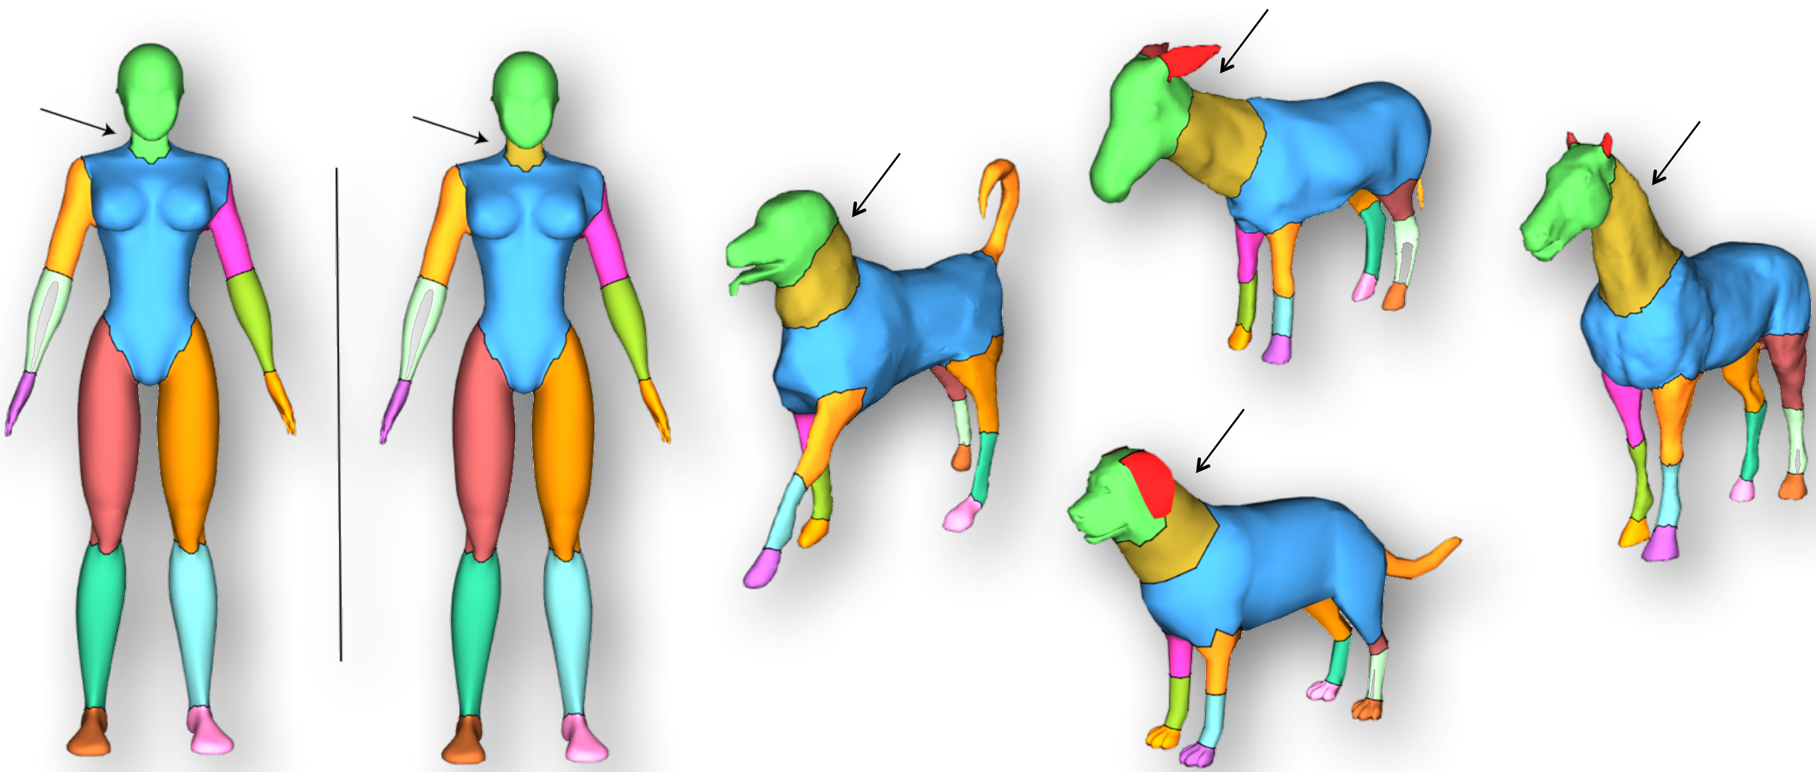

3D Segmentation

The Network View: Information Transport Between Visual Data

Networks of Images

Or of Shapes, Or of Both

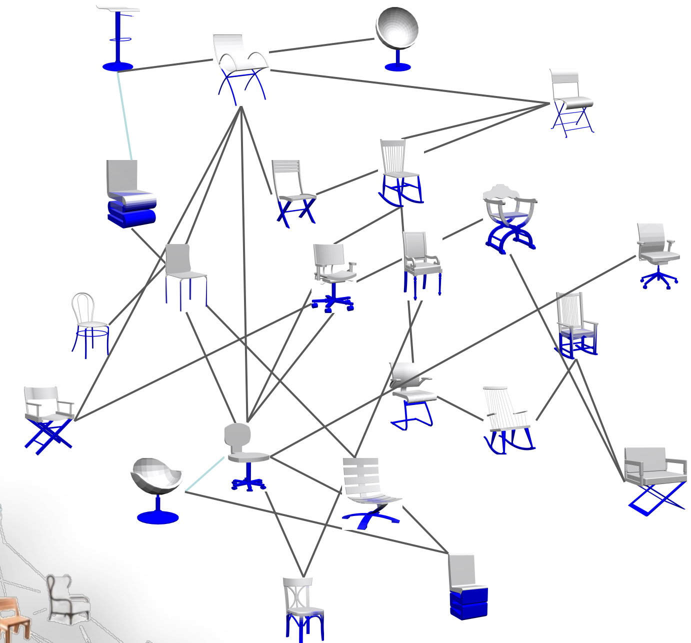

Good Correspondences or Maps are <u>Information Transporters</u>

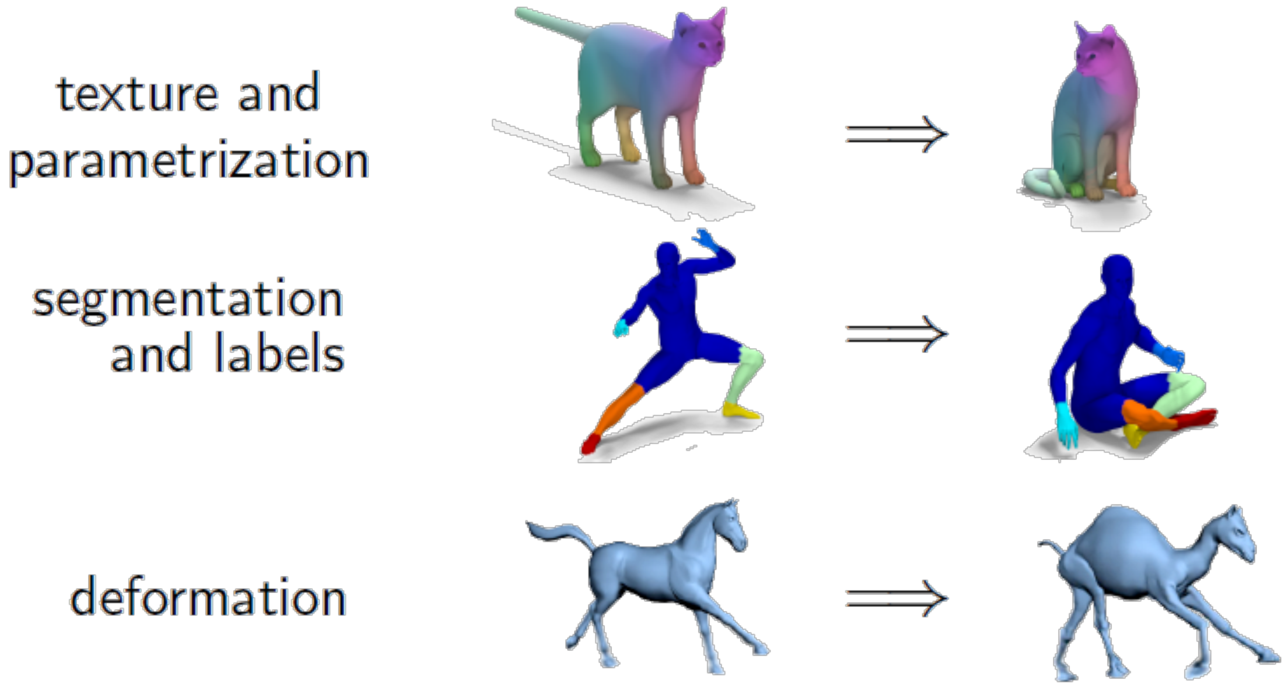

Maps are based on matching

Matching Has Been Extensively Studied

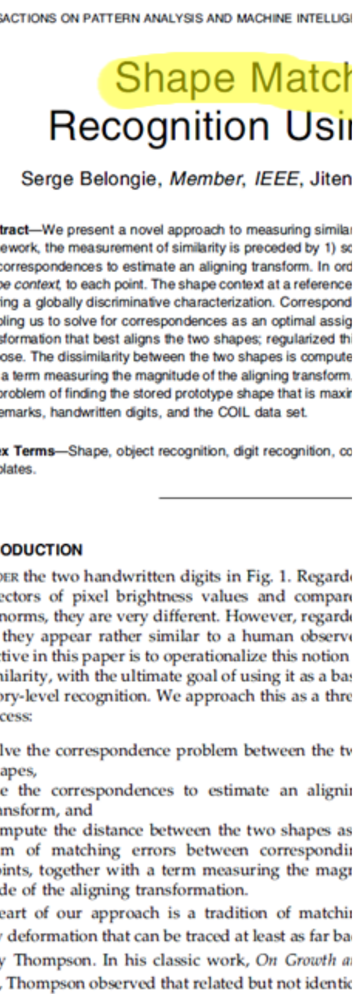

SACTIONS ON PATTERN ANALYSIS AND MACHINE INTELLIGI

Shape Match
Recognition Usi

Serge Belongie, *Member, IEEE*, Jiten

tract—We present a novel approach to measuring similar
ework, the measurement of similarity is preceded by 1) so
correspondences to estimate an aligning transform. In ord
pe context, to each point. The shape context at a reference
ing a globally discriminative characterization. Correspond
bling us to solve for correspondences as an optimal assig
sformation that best aligns the two shapes; regularized th
ose. The dissimilarity between the two shapes is compute
a term measuring the magnitude of the aligning transform.
problem of finding the stored prototype shape that is maxir
emarks, handwritten digits, and the COIL data set.

x Terms—Shape, object recognition, digit recognition, co
plates.

ODUCTION

DER the two handwritten digits in Fig. 1. Regard
ectors of pixel brightness values and compar
norms, they are very different. However, regard
they appear rather similar to a human observ
tive in this paper is to operationalize this notion
nilarity, with the ultimate goal of using it as a ba
ory-level recognition. We approach this as a thre
cess:

lve the correspondence problem between the tw
apes,
e the correspondences to estimate an aligni
nsform, and
mpute the distance between the two shapes as
m of matching errors between correspondi
oints, together with a term measuring the magn
de of the aligning transformation.

eart of our approach is a tradition of matchi
deformation that can be traced at least as far ba
Thompson. In his classic work, *On Growth a*
Thompson observed that related but not identic

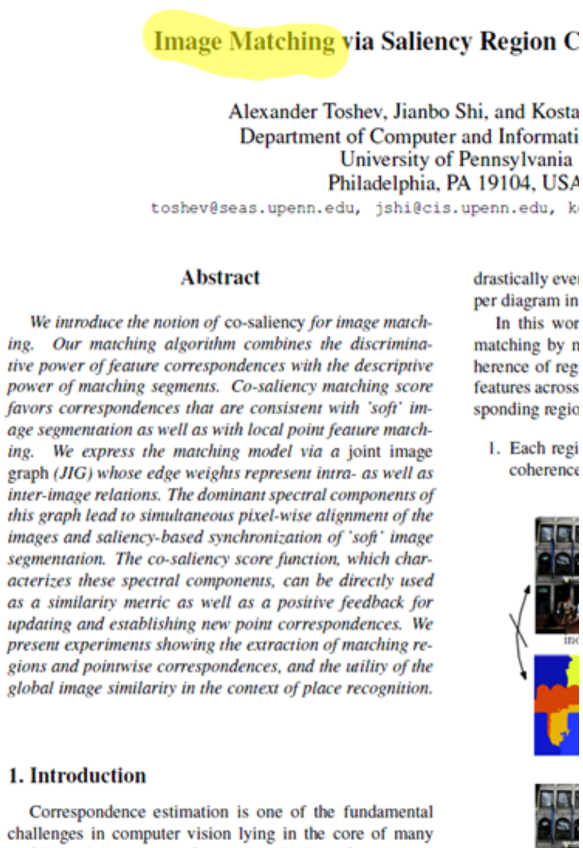

Image Matching via Saliency Region C

Alexander Toshev, Jianbo Shi, and Kosta
Department of Computer and Informati
University of Pennsylvania
Philadelphia, PA 19104, USA
toshev@seas.upenn.edu, jshi@cis.upenn.edu, k

Abstract

We introduce the notion of co-saliency *for image matching. Our matching algorithm combines the discriminative power of feature correspondences with the descriptive power of matching segments. Co-saliency matching score favors correspondences that are consistent with 'soft' image segmentation as well as with local point feature matching. We express the matching model via a joint image graph (JIG) whose edge weights represent intra- as well as inter-image relations. The dominant spectral components of this graph lead to simultaneous pixel-wise alignment of the images and saliency-based synchronization of 'soft' image segmentation. The co-saliency score function, which characterizes these spectral components, can be directly used as a similarity metric as well as a positive feedback for updating and establishing new point correspondences. We present experiments showing the extraction of matching regions and pointwise correspondences, and the utility of the global image similarity in the context of place recognition.*

1. Introduction

Correspondence estimation is one of the fundamental challenges in computer vision lying in the core of many

drastically ever
per diagram in
In this wor
matching by n
herence of reg
features across
sponding regio

1. Each regi
coherence

Learning Graph Matching

Tibério S. Caetano, Li Cheng, Quoc V. Le and Alex J. Smola
Statistical Machine Learning Program, NICTA and ANU
Canberra ACT 0200, Australia

Abstract

As a fundamental problem in pattern recognition, graph matching has found a variety of applications in the field of computer vision. In graph matching, patterns are modeled as graphs and pattern recognition amounts to finding a correspondence between the nodes of different graphs. There are many ways in which the problem has been formulated, but most can be cast in general as a quadratic assignment problem, where a linear term in the objective function encodes node compatibility functions and a quadratic term encodes edge compatibility functions. The main research focus in this theme is about designing efficient algorithms for solving approximately the quadratic assignment problem, since it is NP-hard.

In this paper, we turn our attention to the complementary problem: how to estimate compatibility functions such that the solution of the resulting graph matching problem best matches the expected solution that a human would manually provide. We present a method for learning graph matching: the training examples are pairs of graphs and the "labels" are matchings between pairs of graphs. We present experimental results with real image data which give evidence that learning can improve the performance of standard graph matching algorithms. In particular, it turns out that linear assignment with such a learning scheme may improve over state-of-the-art quadratic assignment relaxations. This finding suggests that for a range of problems where quadratic assignment was thought to be essential for securing good results, linear assignment, which is far more efficient, could be just sufficient if learning is performed. This enables speed-ups of graph matching by up to 4 orders of magnitude while retaining state-of-the-art accuracy.

1. Introduction

Graphs are commonly used as abstract representations for complex scenes, and many computer vision problems can be formulated as an attributed graph matching problem, where the nodes of the graphs correspond to local features of the image and edges correspond to relational aspects between features (both nodes and edges can be attributed,

quadratic assignment problem, which consists in finding
assignment that maximizes an objective function encod
local compatibilities (a linear term) and structural comp
ibilities (a quadratic term). The main body of researc
graph matching has then been focused on devising more
curate and/or faster algorithms to solve the problem appr
imately (since it is NP-hard). The compatibility functi
used in graph matching are typically handcrafted.
An interesting question arises in this context. If we
given two attributed graphs, G and G', should the o
mal match be uniquely determined? For example, assu
first that G and G' come from two images acquired wi
surveillance camera in an airport's lounge. Now, assume
same G and G' instead come from two images in a phot
rapher's image database. Should the optimal match be
same in both situations? If the algorithm takes into acc
exclusively the graphs to be matched, the optimal soluti
will be the same[1] since the graph pair is the same in b
cases. This is how graph matching is approached today.
In this paper we address what we believe to be a lim
tion of this approach. We argue that, if we know the "co
tions" under which a pair of graphs has been extracted, t
we should take into account *how graphs arising in th*
conditions are typically matched. However, we do not t
the information on the "conditions" explicitly into acco
since this is obviously not practical. Instead, we appro
the problem from a purely statistical inference perspect
First we extract graphs from a number of images acqu
in the same conditions as those for which we want to so
whatever the word "conditions" mean (e.g. from the surv
lance camera or the photographer's database). We t
manually provide what we understand to be the opti
matches between pairs of the resulting graphs. This in
mation is then used in a *learning* algorithm which lea
a map from the space of pairs of graphs to the space
matches. In terms of the quadratic assignment problem,
learning algorithm amounts to (in a loose language) adj
ing the node and edge compatibility functions in a way
the expected optimal match in a test pair of graphs agr
with the expected match they would have had they bee
the training set. In this formulation, the learning prob

Maps vs. Distances/Similarities
Networks vs. Graphs

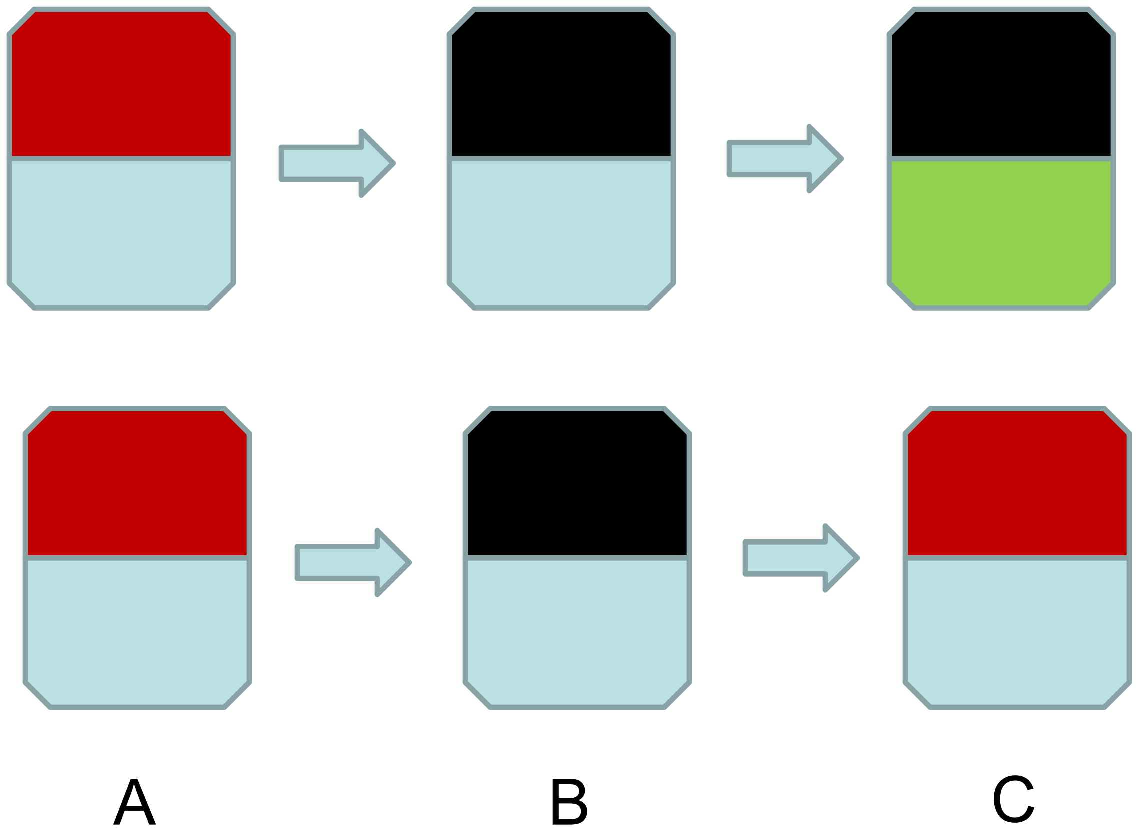

"Persistence" of Correspondences

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.

Consistency of Network Transport

Map Networks for Related Data

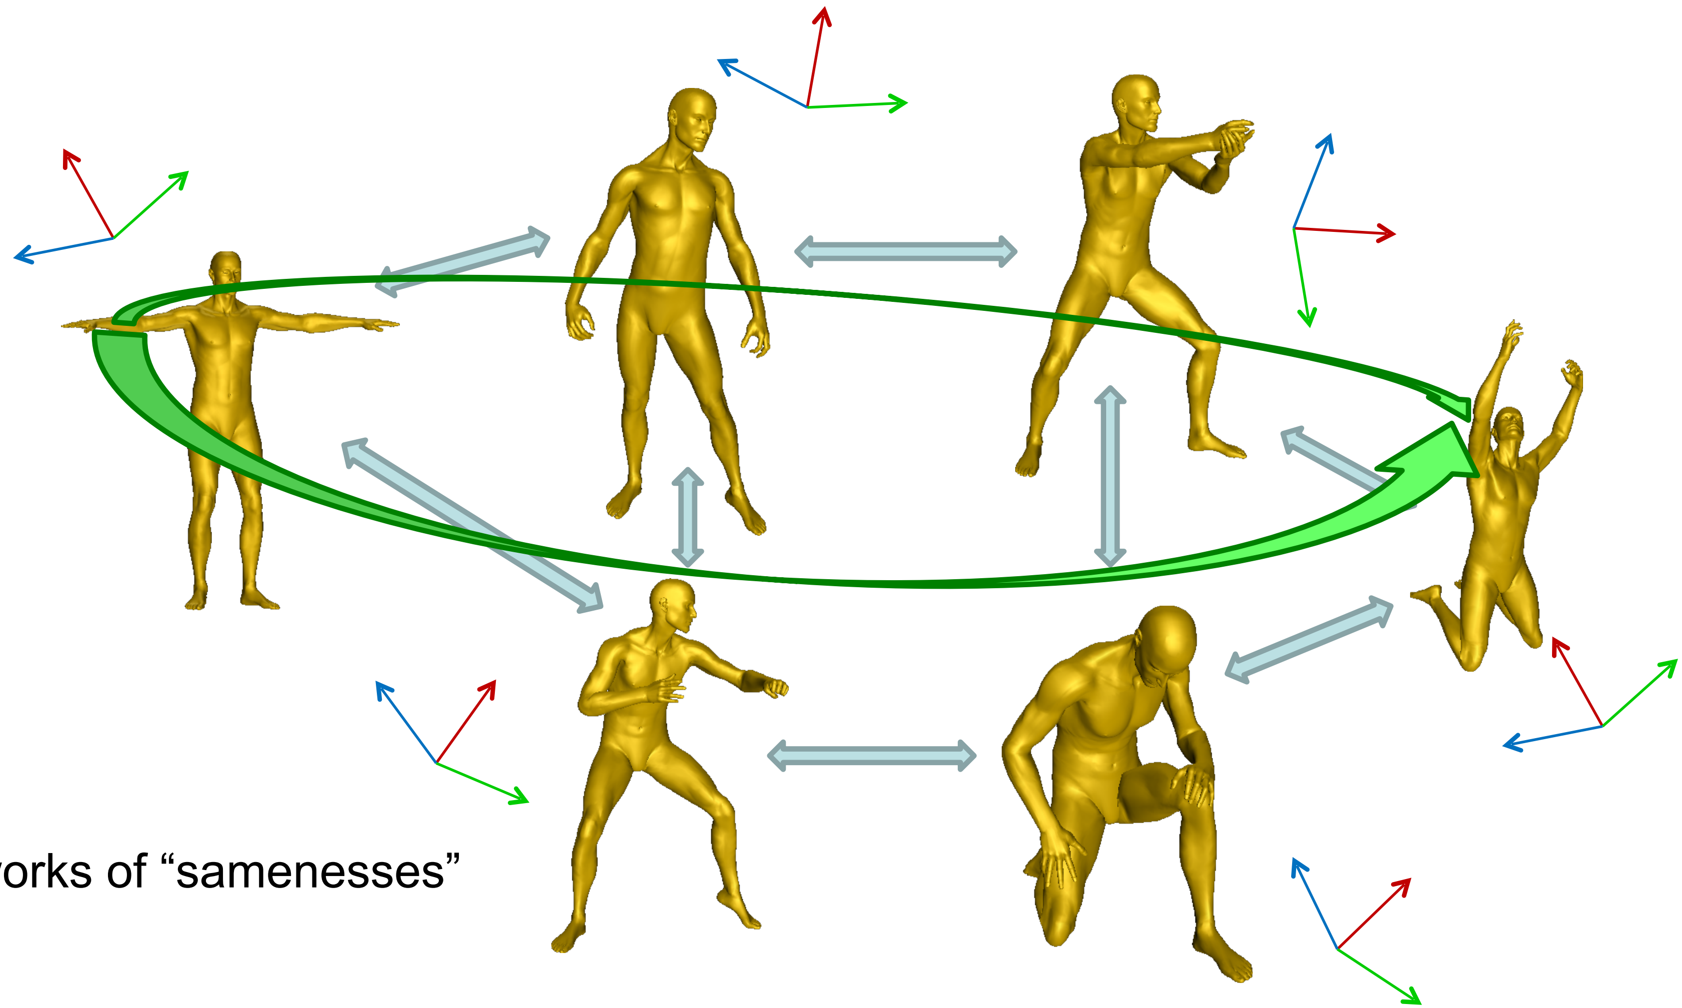

Networks of "samenesses"

Path Invariance

Transform Synchronization Problems

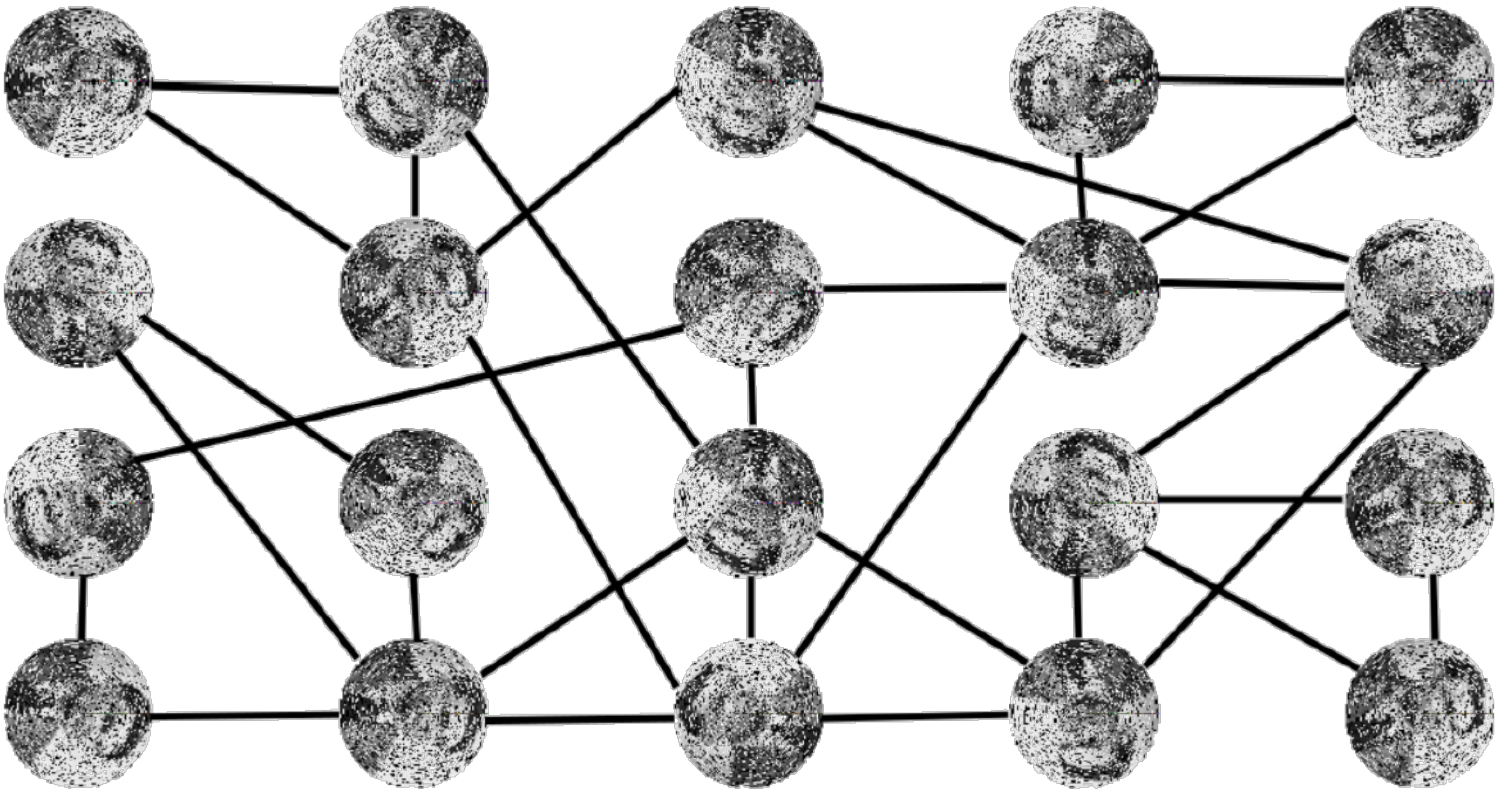

Bandeira, Afonso S. *Ten Lectures and Forty-Two Open Problems in the Mathematics of Data Science.* (2015).

Path invariance ≡ Cycle consistency

Cycle Consistency

- Maps are consistent along cycles

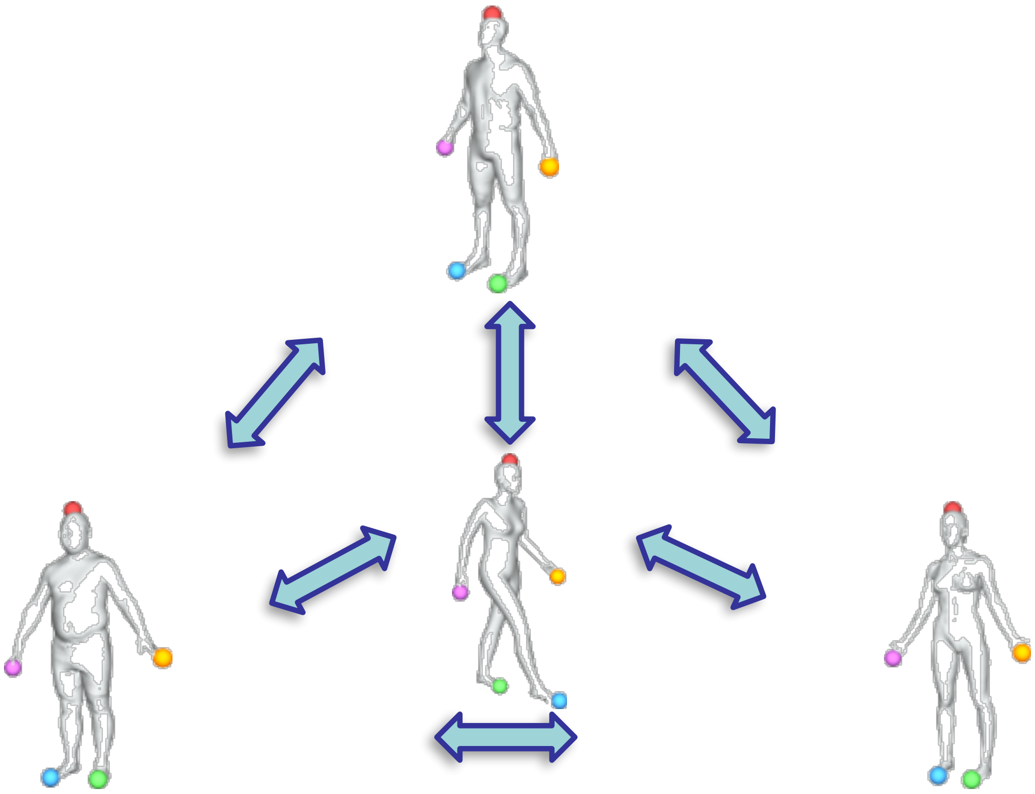

Cycle Consistency

- Maps are consistent along cycles

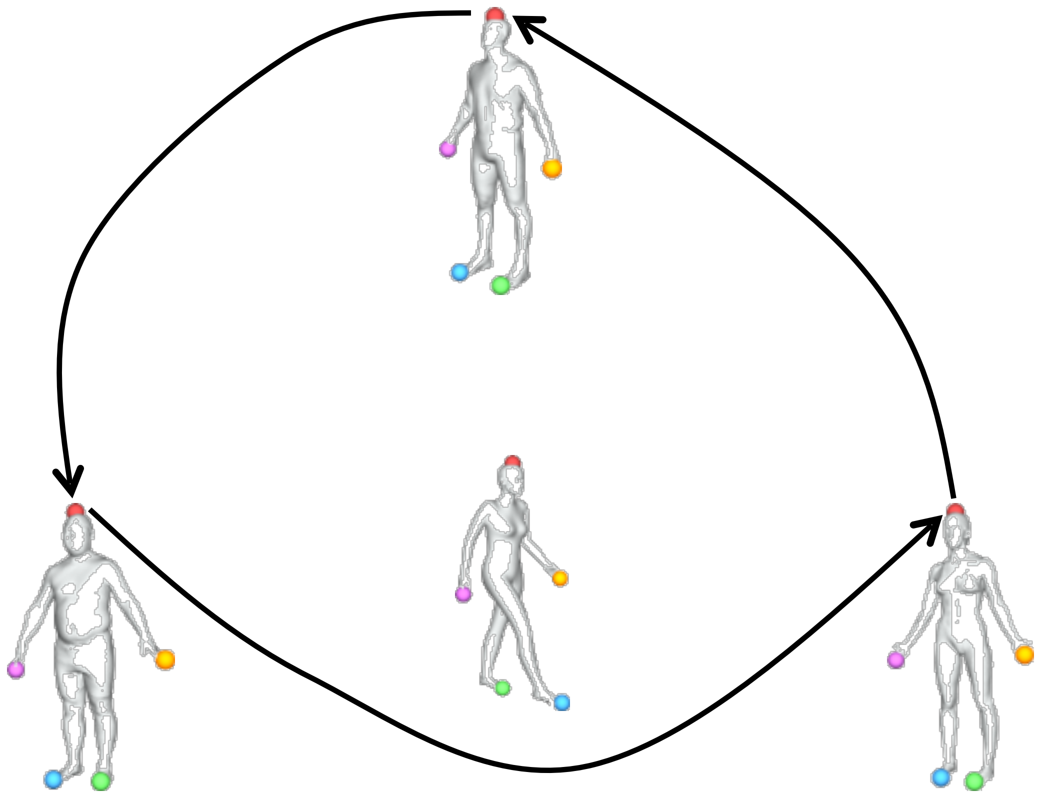

Consistent

Cycle Consistency

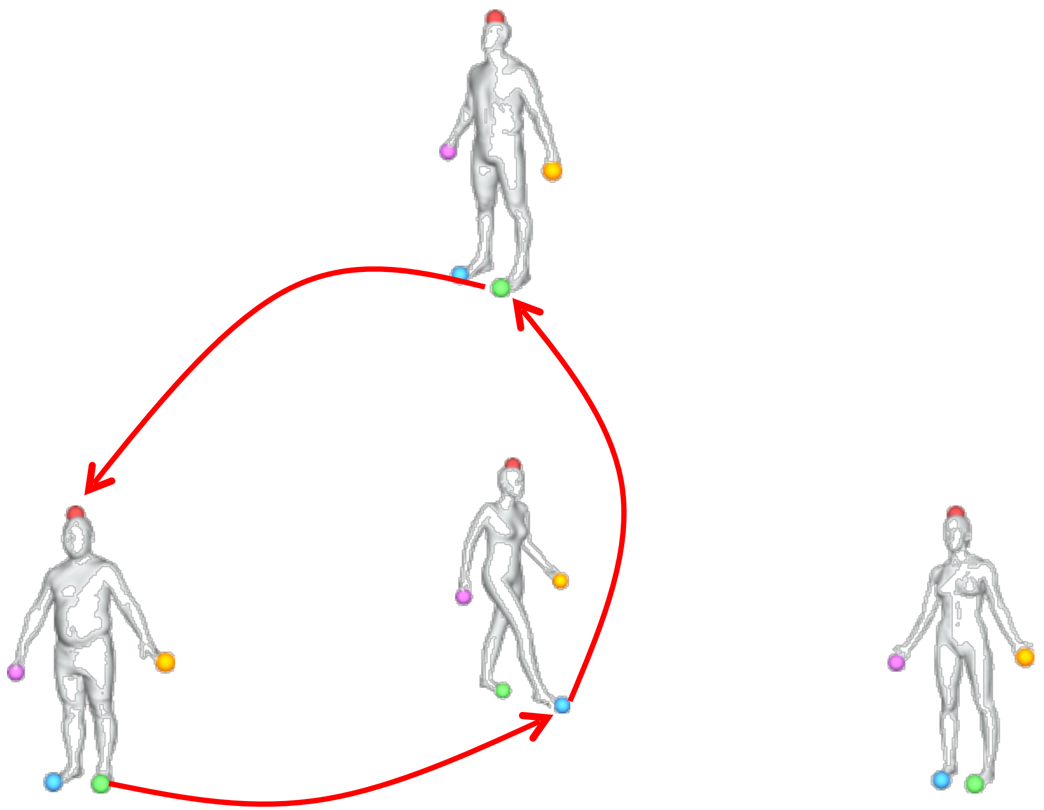

Inconsistent

Cycle Consistency

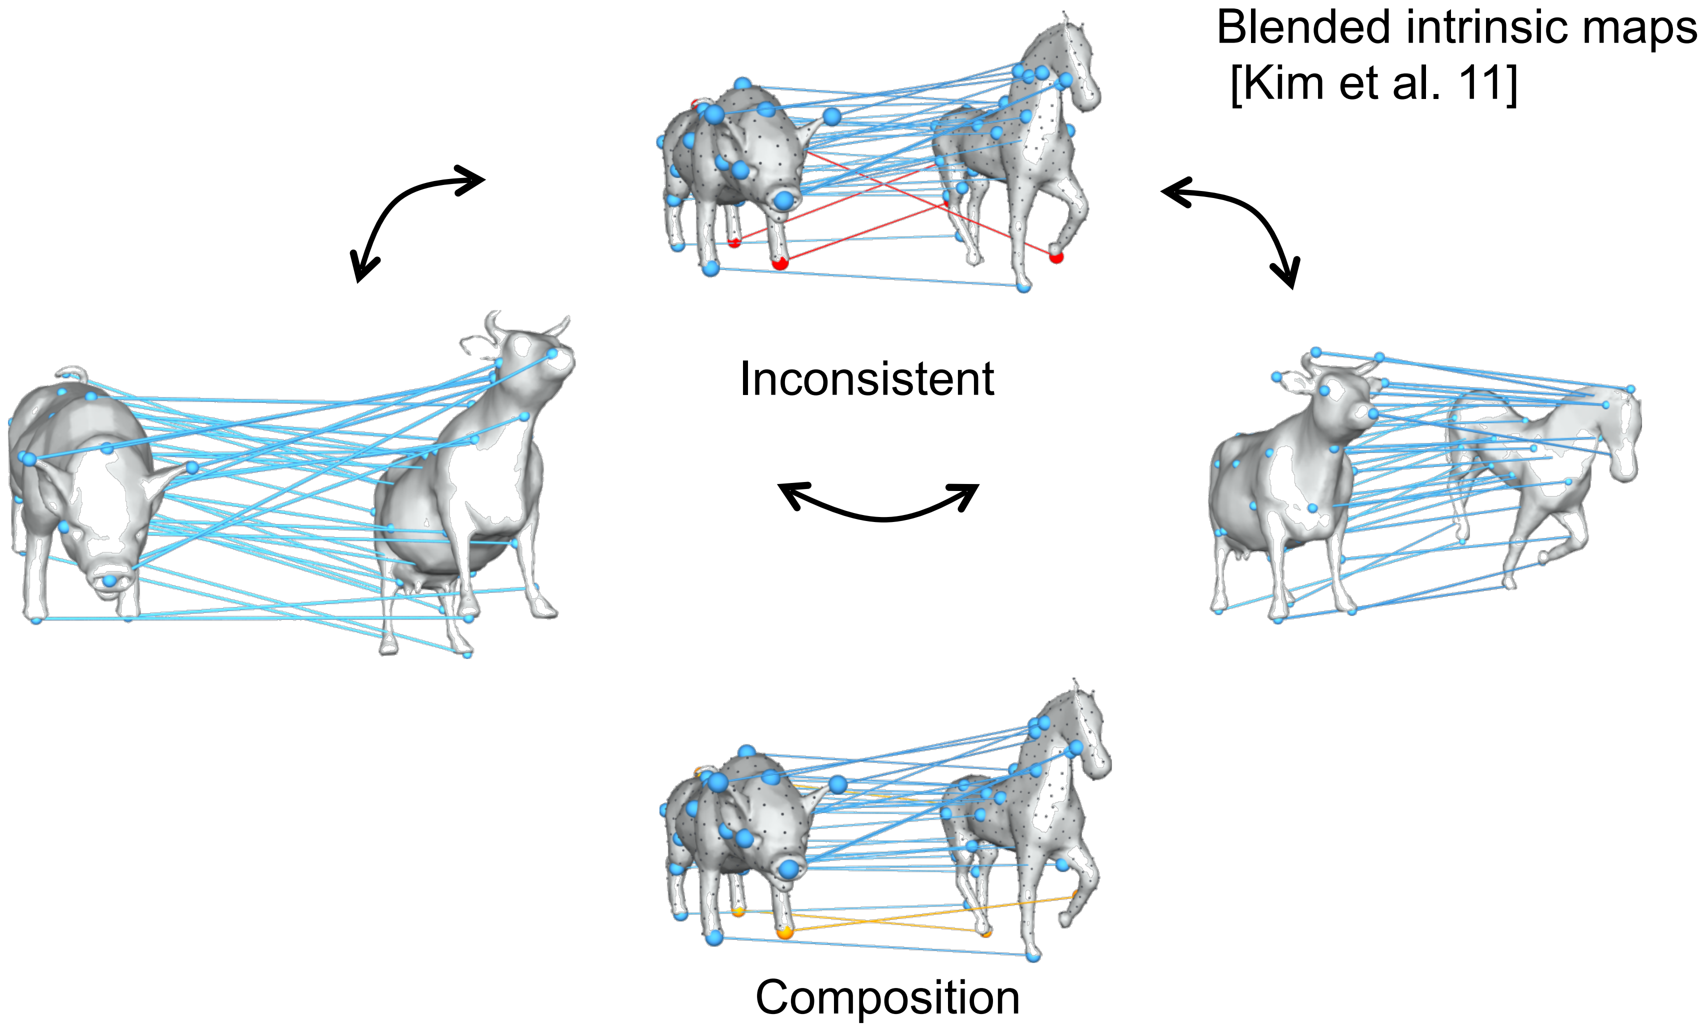

Cycle Consistency Can Help

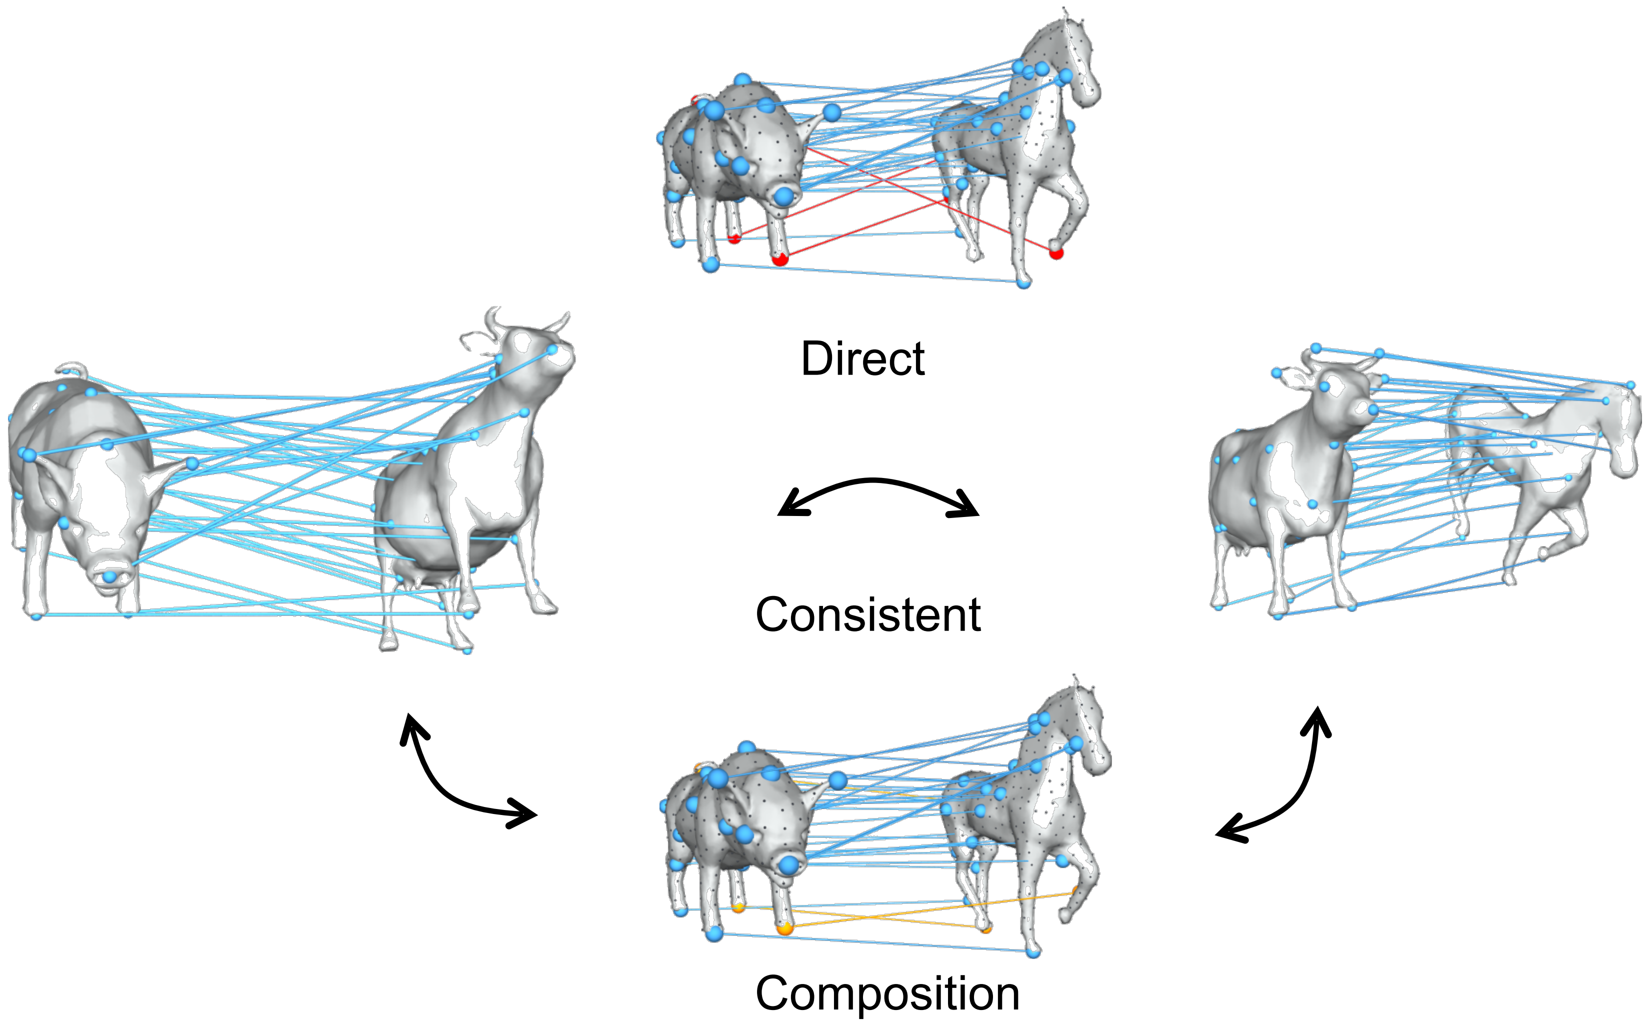

The End