Analysis of Three-Dimensional Shapes
Functional Maps

(IN2238)

Frank R. Schmidt
Matthias Vester
Zorah Lähner

21st June 2016

Shape Matching

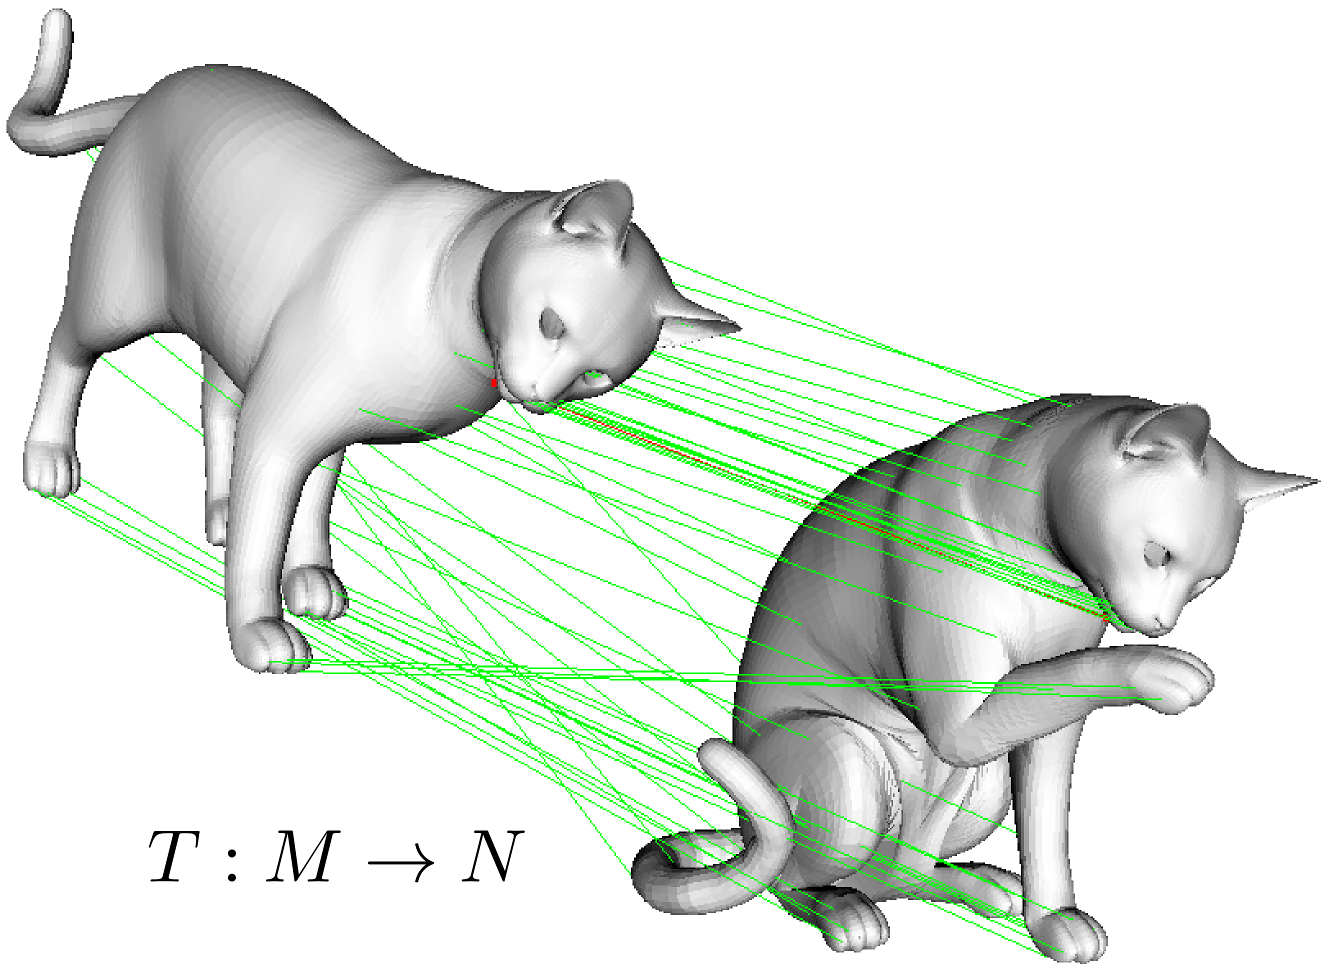

$$T : M \to N$$

Representation

We already saw for 2D Matchings that a correspondences can be represented as permutation matrices if both shapes have the same number of vertices.

M

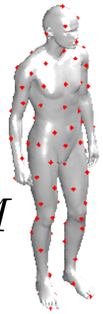

0	1	0	0	0
0	0	0	1	0
1	0	0	0	0
0	0	0	0	1
0	0	1	0	0

N

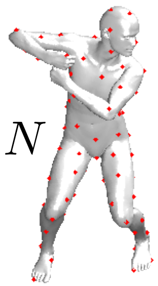

$$\min_{P} \; E(P)$$
$$\text{s.t. } P \text{ is permutation}$$

This does not scale well with the size of the shapes

Problems

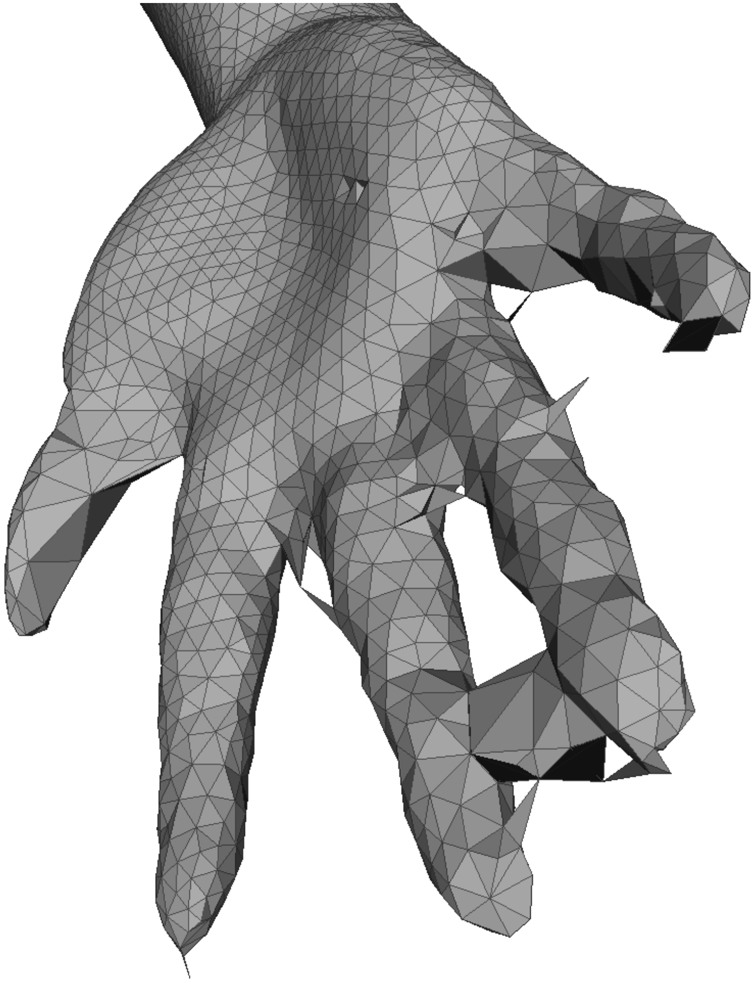

Functional Map

Assume we were given a bijection $T : M \to N$. Given any scalar function $f : M \to \mathbb{R}$ on M we can induce $g : N \to \mathbb{R}$ by composition.

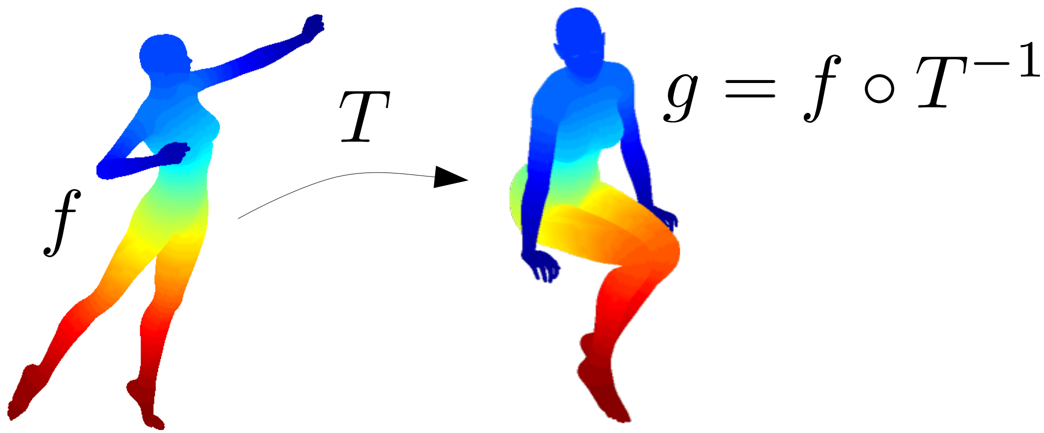

We can denote this transformation by a functional T_F such that

$$T_F(f) = f \circ T^{-1}$$

No Information Loss

If we know T, we can obviously construct T_F by its definition
$T_F(f) = f \circ T^{-1}$.

Can we also reconstruct T if we only know T_F? Yes, we can!

Let $\delta_a : M \to \mathbb{R}$ be an indicator function on M such that

$$e_a(x) = \begin{cases} 1 & \text{if } x = a \\ 0 & \text{otherwise} \end{cases}$$

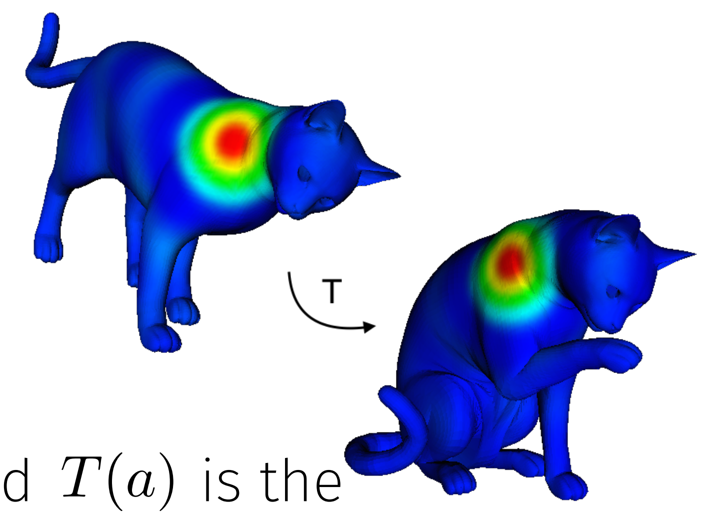

Then if we call $g = T_F(e_a)$, it must be
$g(y) = (e_a \circ T^{-1})(y) = 0$ whenever $T^{-1}(y) \neq a$
and $g(y) = 1$ otherwise.
Since T is a bijection, this happens only once and $T(a)$ is the unique point $y \in N$ such that $g(y) = 1$

Linearity

We can show that T_F is a linear map:

$$\begin{aligned} T_F(\alpha f + g) &= (\alpha f + g) \circ T^{-1} \\ &= \alpha f \circ T^{-1} + g \circ T^{-1} \quad \text{(by linearity of the composition)} \\ &= \alpha T_F(f) + T_F(g) \end{aligned}$$

The key observation is that, while T can be a very complex transformation, T_F always acts *linearly*.

This means we can give T_F a matrix representation after choosing a basis for two function spaces on M and N .

Matrix Notation

Let $\{\phi_i^M, \phi_j^N\}$ be bases for function spaces $\mathcal{F}(M), \mathcal{F}(N)$ on M, N such that $f = \sum_i a_i \phi_i^M, f \in \mathcal{F}(M)$. Then we can write:

$$T_F(f) = T_F\left(\sum_i a_i \phi_i^M\right) = \sum_i a_i T_F\left(\phi_i^M\right)$$

and

$$T_F\left(\phi_i^M\right) = \sum_j c_{ji} \phi_j^N$$

Putting both together, we get:

$$T_F(f) = \sum_i a_i \sum_j c_{ji} \phi_j^N = \sum_{j,i} a_i c_{ji} \phi_j^N$$

Matrix Notation

$$T_F(f) = \sum_j \sum_i a_i c_{ji} \phi_j^N$$
$$= \sum_j b_j \phi_j^N$$

We can represent each function f on M by its coefficients a_i , and similarly $T_F(f)$ on N by the coefficients b_j .

Rewriting in matrix notation, we have:

$$T_F(a) = b = Ca$$

If the bases are orthogonal with respect to some inner product $\langle \cdot, \cdot \rangle$, then we can simply write

$$a_i = \langle f, \phi_i^M \rangle \qquad c_{ij} = \langle T_F(\phi_i^M), \phi_j^N \rangle$$

CONSTRUCTING C

Lets take a closer look at $c_{ij} = \langle T_F(\phi_i^M), \phi_j^N \rangle$

We know it holds: $Pe_x = e_{T(x)}$

Indicator function for vertex $T(x) \in N$

Indicator function for vertex $x \in M$

$a = \Phi_M^{-1} e_x$ Indicator function in the chosen basis of M

Ca Indicator function mapped to the basis of N

$\Phi_N C a$ Indicator function on N in the indicator basis

$$\Rightarrow \quad \Phi_N C \Phi_M^{-1} e_x = e_{T(x)}$$

$$\Phi_N C \Phi_M^{-1} = P \quad \Rightarrow \quad C = \Phi_N^{-1} P \Phi_M$$

Lets take a closer look at

$$C = \Phi_N^{-1} P \Phi_M$$

permutes rows

each column is an eigenfunction

Simply put, the Functional map C contains **all the inner products between the basis functions of the two shapes**, after the vertex ordering has been disambiguated by the bijection P.

This relation was actually already visible here:

$$c_{ij} = \langle T_F(\phi_i^M), \phi_j^N \rangle$$

LINEAR MAP

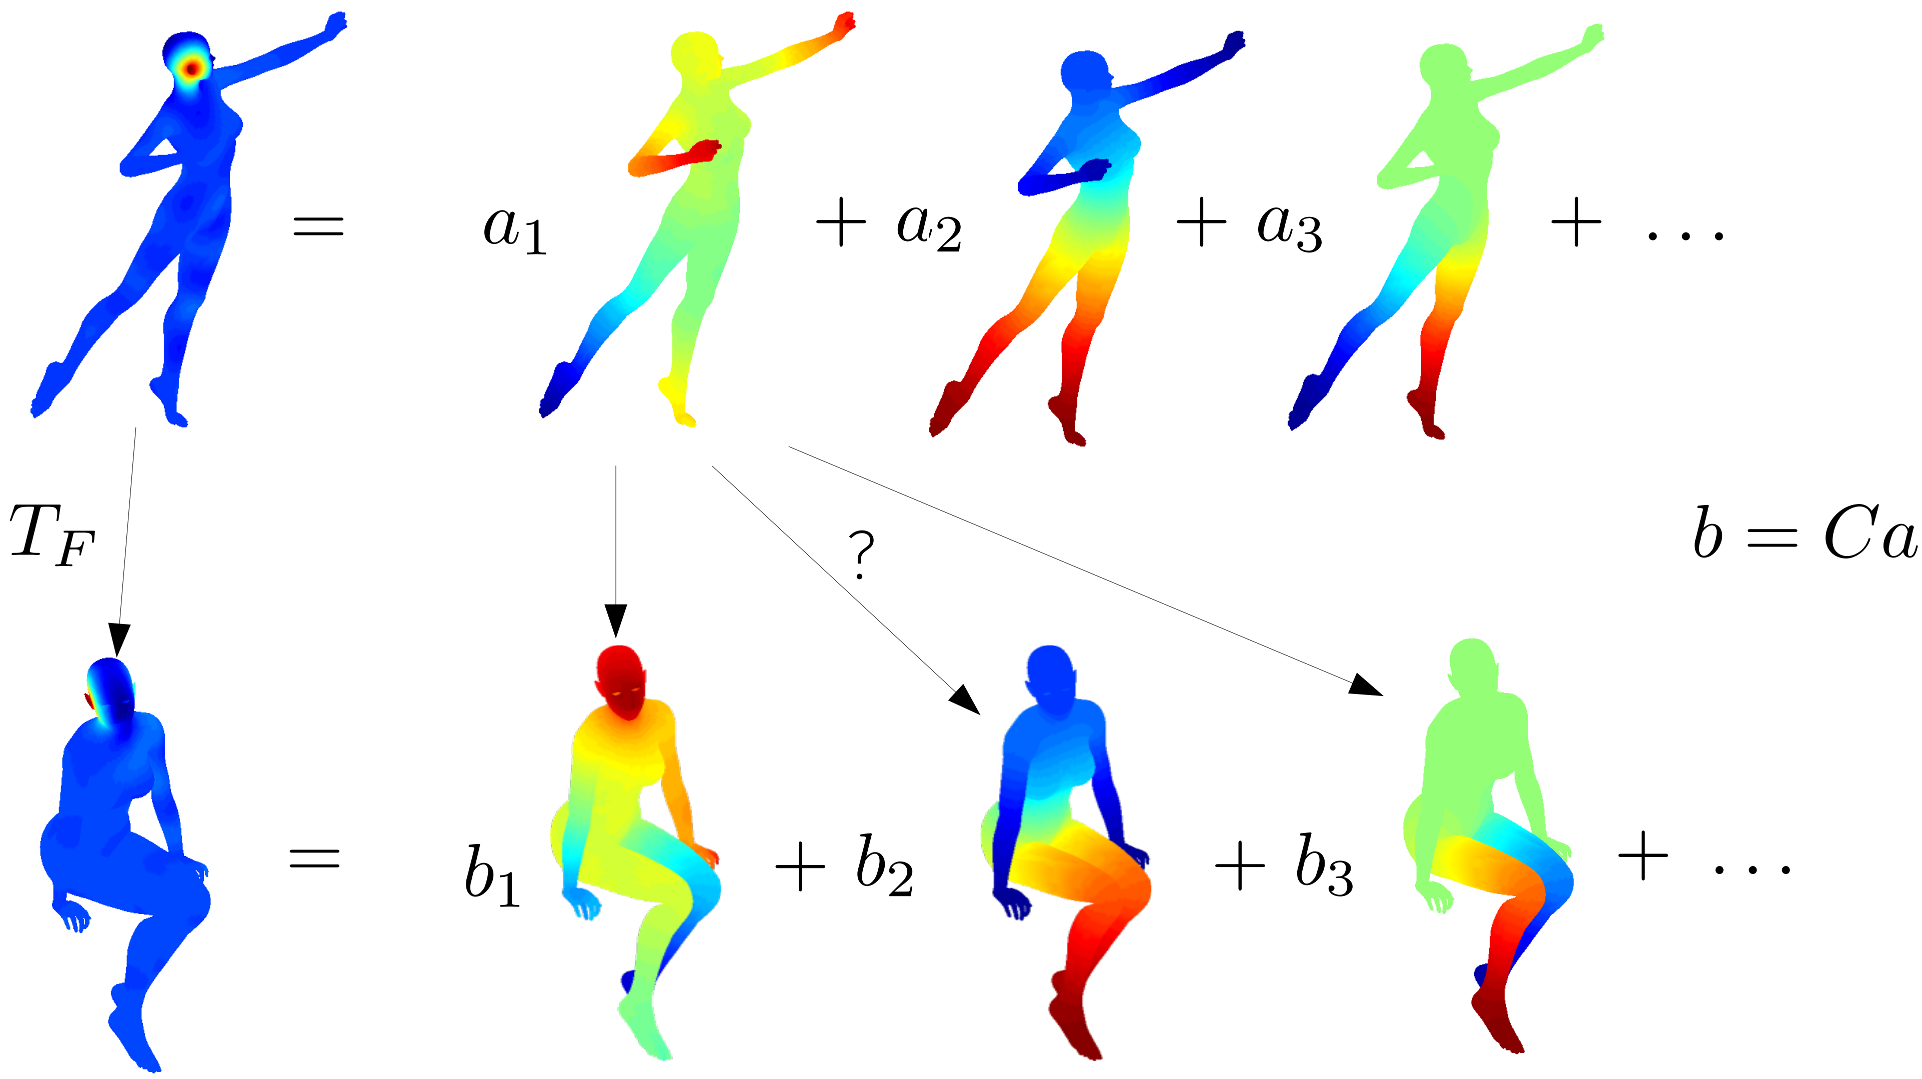

Choice of Basis

Up until now we have been assuming the presence of a basis for functions defined on the two shapes. The first possibility is to consider the indicator basis on each shape:

$$\phi_i^M(x) = \begin{cases} 1 & , \ x = i \\ 0 & , \ \text{otherwise} \end{cases}$$

$$C = \Phi_N^{-1} P \Phi_M$$

$$C = P$$

$$Ca = b \qquad Pa = b$$

P permutation matrix

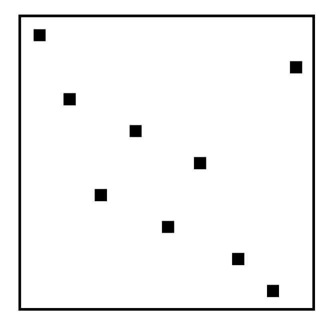

LBO Eigenfunction Basis

But we already learned about another possibility last week!

The eigenfunctions of the Laplace-Beltrami operator form an orthogonal basis (w.r.t. the weighted inner product $\langle\cdot,\cdot\rangle_M$) for l^2 functions on each shape.

In particular, we can approximate:

$$f = \sum_{i=0}^{\infty} a_i \phi_i^M \approx \sum_{i=0}^{m} a_i \phi_i^M$$

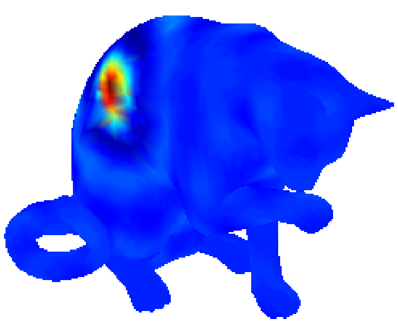

$m = 200$

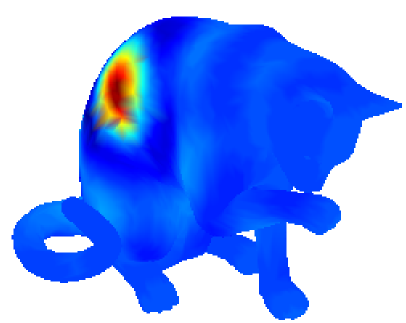

$m = 100$

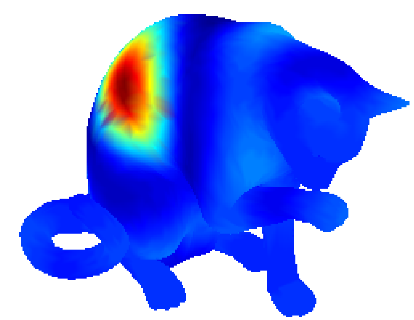

$m = 50$

LBO Eigenfunction Basis

This means we can also approximate:

$$f = \sum_{i,j=0}^{\infty} a_i c_{ij} \phi_j^N \approx \sum_{i,j=0}^{m} a_i c_{ij} \phi_j^N$$

Looking at matrix notation, we are reducing the size of C to $m \times m$

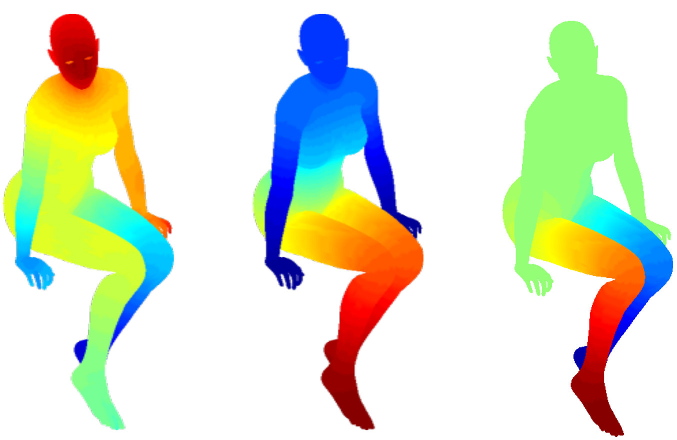

LBO eigenfunctions

$C =$

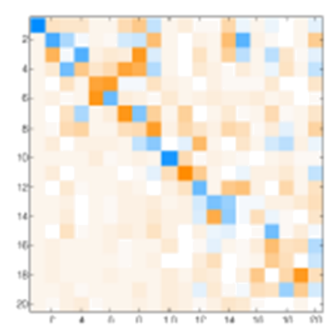

m

m

$\boxed{m \ll n}$

Structure in C

Isometries:

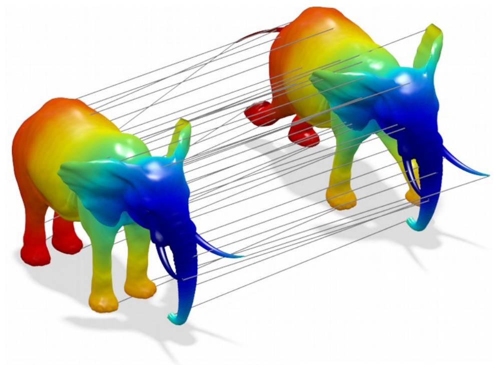

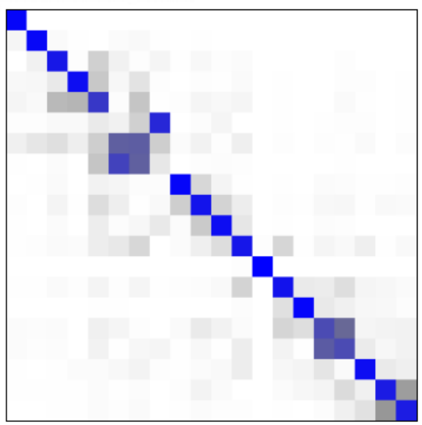

Non-Isometries:

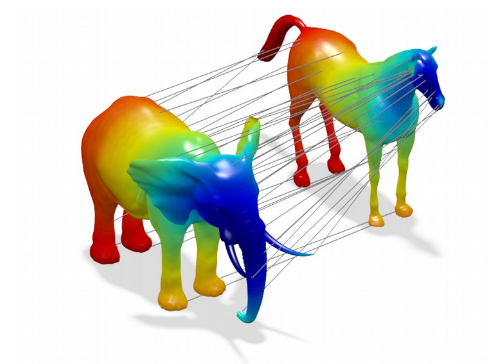

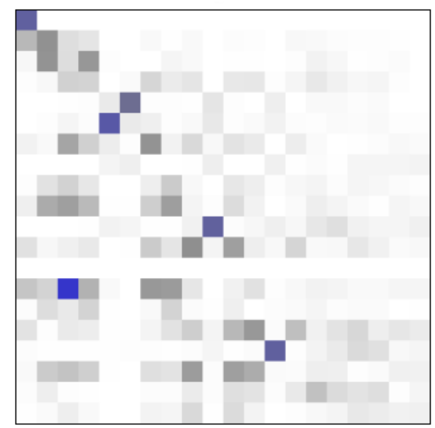

EXAMPLES

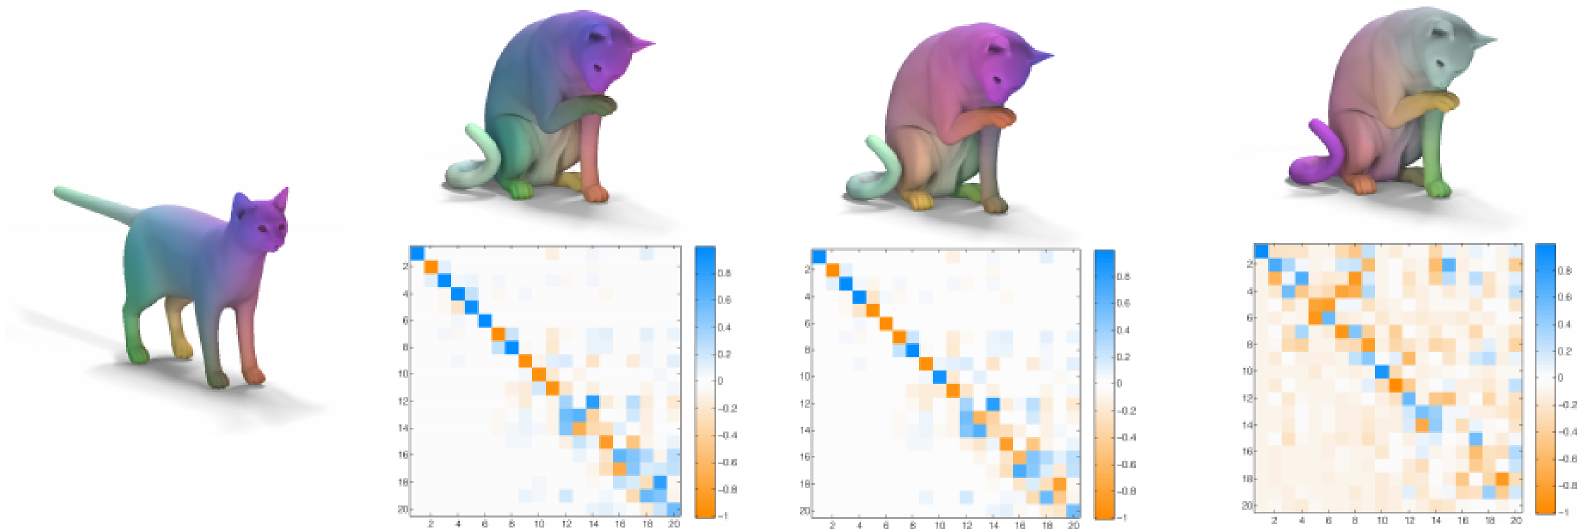

(a) *source* **(b)** *ground-truth map* **(c)** *left to right map* **(d)** *head to tail map*

Note that not every linear map corresponds to a (bijective) point-to-point correspondence.

Computing the Map

If we know enough compatible functions a and b we can deduce the linear relation by solving a bunch of linear equations:

$$Ca = b$$

Descriptor preservation

If we are given k descriptors, we can phrase k equations:

$$Ca_1 = b_1$$
$$Ca_2 = b_2$$
$$\dots$$
$$Ca_k = b_k$$

For instance, consider *curvature* or the *Heat Kernel Signature* from last week.

Landmark matches

Assume we know T(x) = y for some x. We can calculate the geodesic distance maps on both shapes and use them as constraints:

$$d_x^M(x') = d_M(x, x') = a \qquad d_y^N(y') = d_N(y, y') = b$$

Computing the Map

$$C \begin{pmatrix} | & & | \\ \mathbf{a}_1 & \cdots & \mathbf{a}_n \\ | & & | \end{pmatrix} = \begin{pmatrix} | & & | \\ \mathbf{b}_1 & \cdots & \mathbf{b}_n \\ | & & | \end{pmatrix}$$

$m \times m \qquad m \times n \qquad\qquad m \times n$

$n < m$ under-determined

$n = m$ full rank

$n > m$ over-determined

In the common case in which n > m, we can solve the resulting linear system in the least-squares sense:

$$CA = B \Rightarrow C^* = \arg\min_C \|CA - B\|_2^2$$

If we have prior knowledge about the structure of the map we can also add regularizers to the optimization term:

$$\arg\min_{C} \|CA - B\|_2^2 + \rho(C)$$

For example, diagonal structure for isometries:

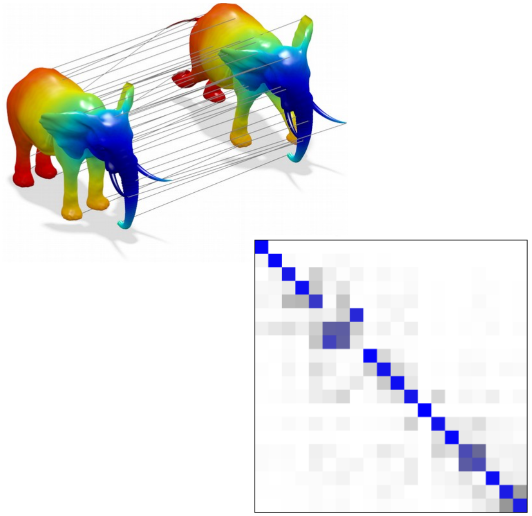

$$\arg\min_{C} \|CA - B\|_2^2 + \|C \circ W\|_2$$

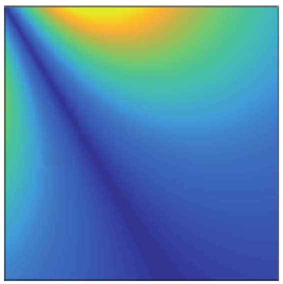

$\mathbf{W}$

Imagine the blue line is diagonal...

Functions to Correspondence

Once we have found an optimal Functional Map C^*, we may want to convert it back to a point-to-point correspondence.

Simplest idea: Map indicator functions *at each point*.

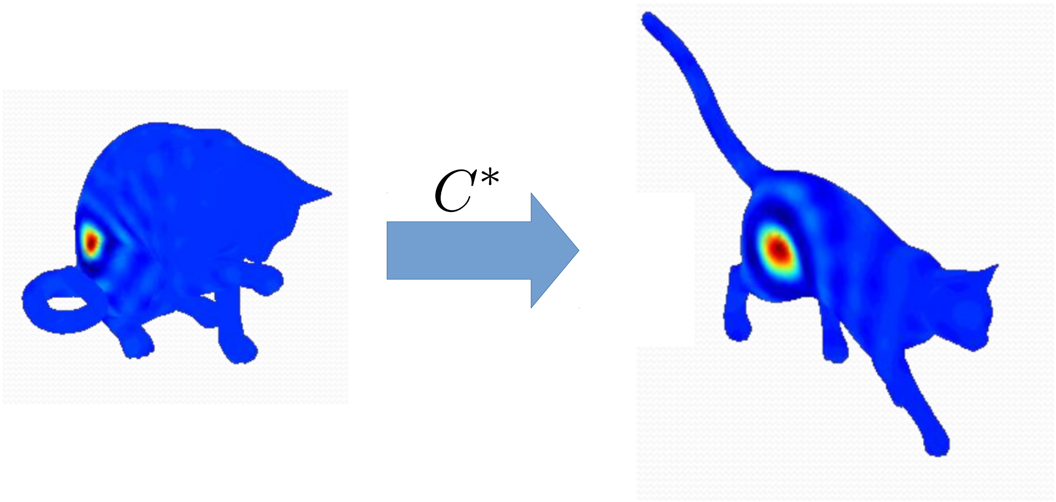

This is very inefficient and sensitive to numerical errors from truncation.

Functions to Correspondence

Observe that each indicator function around x, when represented in the eigenbasis, has as coefficients the k-th column of the matrix $\Phi_M^\top$ where k is the index of point x.

$$\Phi_M^\top e_k \in \mathbb{R}^m$$

Representation of one indicator function in the eigenbasis

$$\Phi_M^\top \in \mathbb{R}^{m \times n}$$

Representation of **all** indicator functions in the eigenbasis

$\Phi_M^\top$ can be regarded as a set of n points in $\mathbb{R}^m$

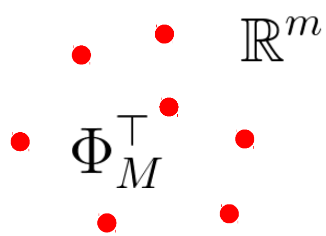

Clearly, the same can be done for the eigenfunctions on the second shape N

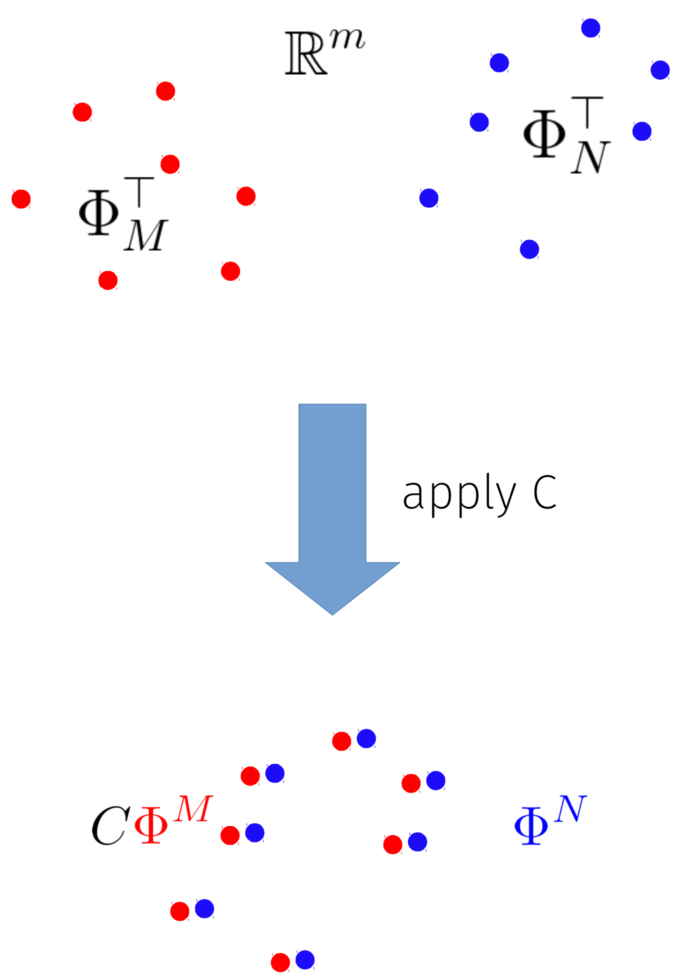

We can find correspondences by aligning both point clouds and searching for nearest neighbors

$$y^* = \arg\min_y \|T_F(e_x) - e_y\|_{L^2}^2$$

$$\approx \arg\min_y \|C \begin{pmatrix} \phi_1^M(x) \\ \vdots \\ \phi_m^M(x) \end{pmatrix} - \begin{pmatrix} \phi_1^N(y) \\ \vdots \\ \phi_m^N(y) \end{pmatrix} \|_{L^2}^2$$

$$\min_{P \in \{0,1\}^{n \times n}} \|C(\Phi^M)^\top - \Phi^N P\|_F^2$$

$$s.t. \quad P^\top 1 = 1$$
$$P1 = 1$$

EXAMPLE

Recovering correspondences from low-rank Functional Maps is a whole problem on its own.

ISSUES

Even when choosing a small k, a lot of compatible functions are necessary to reliably solve for C without imposing any regularization.

But the regularization terms heavily depend the basis.

Laplace-Beltrami eigenbasis: robust to nearly-isometric deformations only!

The recovered correspondences are often neither bijective nor continuous.

Matching Partial Shapes

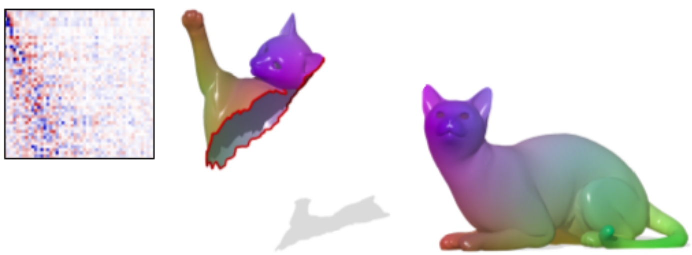

Moving away from isometries

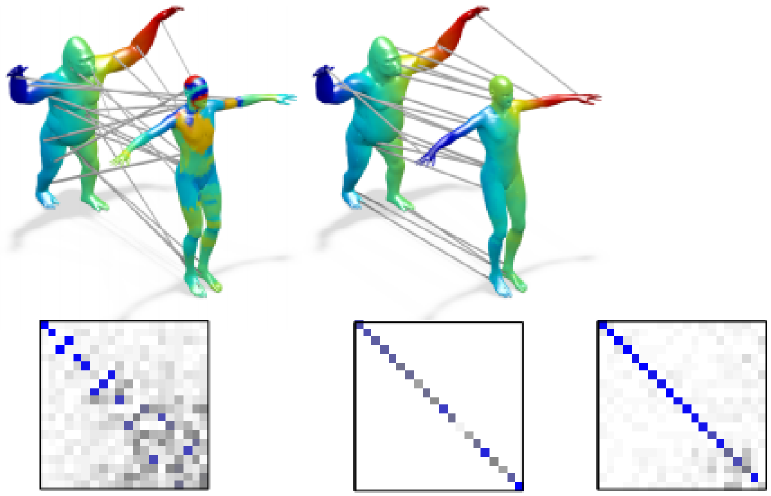

Extensions for vector fields...

References

- **Functional Maps: A Flexible Representation of Maps Between Shapes**. Ovsjanikov, Ben-Chen, Solomon, Butscher, Guibas. ACM SIGGRAPH, 2012.
- **Sparse Modeling of Intrinsic Correspondences.** Pokrass, Bronstein, Bronstein, Sprechmann, Sapiro. Computer Graphics Forum, 2014.
- **Partial Functional Correspondence.** Rodola, Cosmo, Bronstein, Torsello, Cremers. Computer Graphics Forum, 2016.
- **Coupled Quasi-Harmonic Bases.** Kovnatsky, Bronstein, Bronstein, Glashoff, Kimmel. Computer Graphics Forum, 2013.