Computing and Processing Correspondences with Functional Maps

SIGGRAPH Asia 2016 course

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



Università della Svizzera italiana







General Overview

Overall Objective:

Create tools for computing and analyzing *mappings* between geometric objects.



General Overview

Overall Objective:

Create tools for computing and analyzing *mappings* between geometric objects.



Rather than comparing *points* on objects it is often easier to compare *real-valued functions* defined on them. Such maps can be represented as matrices.

Course Overview

Course Website:

http://www.lix.polytechnique.fr/~maks/fmaps_course/

or http://bit.do/fmaps

Ourse Notes:



Linked from the website. Or use <u>http://bit.do/fmaps_notes</u> Attention: (significantly) more material than in the lectures

Sample Code:

See **Sample Code** link on the website.

Course Schedule

2:15pm – 3:10pm *Introduction* (Maks)

- Introduction to the functional maps framework.
- 2:15pm 3:10pm *Computing Functional Maps* (Michael)
 Optimization methods for functional map estimation.

04:00pm - 04:15pm Break



04:15pm - 05:05pm *Maps in Shape Collections* (Etienne)
Networks of Maps, Descriptor learning, Shape comparison.

3 05:10pm - 06:00pm *Conversion, Applications* (Emanuele)

• Pointwise map recovery, Applications

06:00pm - Wrapup, Q&A (all)

What is a Shape?

- Continuous: a surface embedded in 3D.
- Discrete: a triangle mesh.



5k – 200k triangles

Shapes from the SCAPE, TOSCA and FAUST datasets

What is a Shape?

A graph embedded in 3D: a triangle mesh.

- Connected.
- Manifold (each edge on at most 2 triangles).
- Without boundary.



Shapes from the SCAPE, TOSCA and FAUST datasets

Rigid Shape Matching



- Given a pair of shapes, find the optimal *Rigid Alignment* between them.
- The unknowns are the rotation/translation parameters of the source onto the target shape.

Iterative Closest Point (ICP)

• Classical approach: iterate between finding correspondences and finding the transformation:



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

1. For each $x_i \in \mathcal{M}$ find **nearest** neighbor $y_i \in \mathcal{N}$.

$$\underset{R,t}{\operatorname{arg\,min}} \ \sum_{i} \|Rx_{i} + t - y_{i}\|_{2}^{2}$$

• Classical approach: iterate between finding correspondences and finding the transformation:



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

1. For each $x_i \in \mathcal{M}$ find **nearest** neighbor $y_i \in \mathcal{N}$.

$$\underset{R,t}{\operatorname{arg\,min}} \sum_{i} \|Rx_i + t - y_i\|_2^2$$

• Classical approach: iterate between finding correspondences and finding the transformation:



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

1. For each $x_i \in \mathcal{M}$ find **nearest** neighbor $y_i \in \mathcal{N}$.

$$\underset{R,t}{\operatorname{arg\,min}} \ \sum_{i} \|Rx_{i} + t - y_{i}\|_{2}^{2}$$

• Classical approach: iterate between finding correspondences and finding the transformation:



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

1. For each $x_i \in \mathcal{M}$ find **nearest** neighbor $y_i \in \mathcal{N}$.

$$\underset{R,t}{\operatorname{arg\,min}} \ \sum_{i} \|Rx_{i} + t - y_{i}\|_{2}^{2}$$

• Classical approach: iterate between finding correspondences and finding the transformation:

$$\mathcal{M}$$
 \mathcal{N}

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

1. For each $x_i \in \mathcal{M}$ find **nearest** neighbor $y_i \in \mathcal{N}$.

$$\underset{R,t}{\operatorname{arg\,min}} \ \sum_{i} \|Rx_{i} + t - y_{i}\|_{2}^{2}$$

• Classical approach: iterate between finding correspondences and finding the transformation:



- 1. Finding nearest neighbors: can be done with spacepartitioning data structures (e.g., KD-tree).
- 2. Finding the optimal transformation **R**, *t* minimizing: $\underset{R,t}{\operatorname{arg\,min}} \sum_{i} ||Rx_i + t - y_i||_2^2$

Can be done efficiently via SVD decomposition.

Arun et al., Least-Squares Fitting of Two 3-D Point Sets

Non-Rigid Shape Matching



Unlike rigid matching with rotation/translation, there is no compact representation to optimize for in non-rigid matching.

Non-Rigid Shape Matching



Main Problems:

- What does it mean for a correspondence to be "good"?
- How to compute it efficiently in practice?

Isometric Shape Matching

Possible Model:

Good maps must preserve geodesic distances.



Geodesic: length of shortest path lying entirely on the surface.

Isometric Shape Matching



Find the point mapping by minimizing the distance distortion:

$$T_{\text{opt}} = \underset{T}{\operatorname{arg\,min}} \sum_{x,y} \| d^{\mathcal{M}}(x,y) - d^{\mathcal{N}}(T(x),T(y)) \|$$

The unknowns are point correspondences.

Isometric Shape Matching



Find the point mapping by minimizing the distance distortion:

$$T_{\text{opt}} = \underset{T}{\operatorname{arg\,min}} \sum_{x,y} \| d^{\mathcal{M}}(x,y) - d^{\mathcal{N}}(T(x),T(y)) \|$$

Problem:

The space of possible solutions is highly non-linear, non-convex.

Functional Map Representation

We would like to define a representation of shape maps that is more amenable to direct optimization.

- 1. A compact representation for "natural" maps.
- 2. Inherently global and multi-scale.
- 3. Handles uncertainty and ambiguity gracefully.
- 4. Allows efficient manipulations (averaging, composition).
- 5. Leads to simple (linear) optimization problems.



Background: Laplace-Beltrami Operator

Given a compact Riemannian manifold \mathcal{M} without boundary, the Laplace-Beltrami operator Δ :

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



Laplace-Beltrami Operator

Given a compact surface \mathcal{M} without boundary, the Laplace-Beltrami operator Δ :

- 1. Is invariant under isometric deformations.
- 2. Has a countable eigendecomposition:

$$\Delta \phi_i = \lambda_i \phi_i$$

that forms an orthonormal basis for $L^2(\mathcal{M})$.

3. Characterizes the geodesic distances fully.

Laplace-Beltrami Eigenfunctions

The Laplace-Beltrami operator Δ has an eigendecomposition: $\Delta \phi_i = \lambda_i \phi_i$



Ordered from low frequency (smoothest) to higher frequency (oscillating).

Laplace-Beltrami Eigenfunctions

.... that forms an orthonormal basis for $L^2(\mathcal{M})$: Any (square-integrable) $f : \mathcal{M} \to \mathbb{R}$ can be represented as a linear combination of the LB eigenfunctions.



In the Discrete World

- Functions are defined at vertices of the mesh.
- Integration is defined with respect to a discrete volume measure: ||f||₂² = f^TAf
 A diagonal matrix of area weights.
- Laplacian is discretized as a matrix $L = A^{-1}W$



$$L_{ij} = \frac{1}{2A(j)} \left(\cot(\alpha_{ij}) + \cot(\beta_{ij}) \right)$$

Can be derived from 1st order FEM.

In the Discrete World

• Computing the eigenfunctions of the Laplacian reduces to solving the generalized eigenvalue problem:

$$L\phi = \lambda\phi \quad \Leftrightarrow \quad W\phi = \lambda A\phi$$

- **eigs** function in Matlab
- Both A and W are sparse positive semidefinite.

Number of	Computation
triangles	time (in s)
5000	0.65
25000	2.32
50868	3.6
105032	10

Time to compute 100 basis functions.

Functional Approach to Mappings

Given a pair of shapes and a point-to-point map $T : \mathcal{N} \to \mathcal{M}$



The induced functional correspondence:

$$T_F(f) = g, \quad g = f \circ T$$

Attention: the functional map is in the opposite direction.

Approach

Given a pair of shapes and a point-to-point map $T : \mathcal{N} \to \mathcal{M}$

$$f: \mathcal{M} \to \mathbb{R} \xrightarrow{T_F} g: \mathcal{N} \to \mathbb{R}$$

$$f: \mathcal{M} \to \mathbb{R} \xrightarrow{T_F} \mathcal{N}$$

$$\mathcal{M} \xrightarrow{T} \mathcal{N}$$

$$T_F(f) = g, \quad g = f \circ T$$

Note that $T_F(f)$ is:

1) Linear $T_F(\alpha_1 f_1 + \alpha_1 f_1) = \alpha_1 T_F(f_1) + \alpha_2 T_F(f_2)$

2) Complete (recover *T* from indicator functions)

Approach

Given a pair of shapes and a point-to-point map $T : \mathcal{N} \to \mathcal{M}$

$$f: \mathcal{M} \to \mathbb{R} \xrightarrow{T_F} g: \mathcal{N} \to \mathbb{R}$$

$$f: \mathcal{M} \to \mathbb{R} \xrightarrow{T_F} \mathcal{N}$$

$$\mathcal{M} \xrightarrow{T} \mathcal{N}$$

$$T_F(f) = g, \quad g = f \circ T$$

Note that $T_F(f)$ is:

1) Linear $T_F(\alpha_1 f_1 + \alpha_1 f_1) = \alpha_1 T_F(f_1) + \alpha_2 T_F(f_2)$

2) Complete (recover *T* from indicator functions)

Approach

Given a pair of shapes and a point-to-point map $T : \mathcal{N} \to \mathcal{M}$

$$f: \mathcal{M} \to \mathbb{R} \xrightarrow{T_F} g: \mathcal{N} \to \mathbb{R}$$

$$f: \mathcal{M} \to \mathbb{R} \xrightarrow{T_F} \mathcal{N}$$

$$\mathcal{M} \xrightarrow{T} \mathcal{N}$$

$$T_F(f) = g, \quad g = f \circ T$$

Note that $T_F(f)$ is:

1) Linear $T_F(\alpha_1 f_1 + \alpha_1 f_1) = \alpha_1 T_F(f_1) + \alpha_2 T_F(f_2)$

2) Complete (recover *T* from indicator functions)

Observation



Express both f and $T_F(f)$ in terms of basis functions:

$$f = \sum_{i} a_{i} \phi_{i}^{\mathcal{M}} \qquad g = T_{F}(f) = \sum_{i} b_{i} \phi_{i}^{\mathcal{N}}$$

Since T_F is linear, there is a linear transformation from $\{a_i\}$ to $\{b_j\}$.

Functional Map Definition



Functional map: matrix C that translates coefficients from $\Phi_{\mathcal{M}}$ to $\Phi_{\mathcal{N}}$.

Functional Maps

Definition:

For a fixed choice of basis functions $\{\phi^{\mathcal{M}}\}, \{\phi^{\mathcal{N}}\}, \text{ and a}$ linear transformation T_F between functions, a functional map is a matrix C, s.t. for any $f = \sum_i a_i \phi_i^{\mathcal{M}}$ if $T(f) = \sum_i b_i \phi_{i'}^{\mathcal{N}}$ then: $\mathbf{b} = C\mathbf{a}$

 C_{ij} : coefficient of $T_F(\phi_j^{\mathcal{M}})$ in the basis of $\phi_i^{\mathcal{N}}$. In an orthonormal basis: $C_{ij} = \int_{\mathcal{N}} T_F(\phi_j^{\mathcal{M}}) \phi_i^{\mathcal{N}} d\mu$

Example 1



Given two shapes with $n_{\mathcal{M}}, n_{\mathcal{N}}$ points and a map: $T : \mathcal{N} \to \mathcal{M}$

 $\mathbf{T}: n_{\mathcal{N}} \times n_{\mathcal{M}} \quad \text{matrix encoding the map } T, \\ \text{one 1 per column with zeros everywhere else.}$

If functions are represented as vectors (in the hat basis), the functional map is given by matrix-vector product:

$$g = \mathbf{T}^T f \qquad C = \mathbf{T}^T$$

Example 2

Given two shapes with $n_{\mathcal{M}}, n_{\mathcal{N}}$ points and a map: $T : \mathcal{N} \to \mathcal{M}$

 $\mathbf{T}: n_{\mathcal{N}} \times n_{\mathcal{M}} \quad \text{matrix encoding the map } T, \\ \text{one 1 per column with zeros everywhere else.}$

If functions are represented in the reduced basis:

 $\Phi_{\mathcal{M}}: n_{\mathcal{M}} \times k_{\mathcal{M}}$ matrix of the first $k_{\mathcal{M}}$ eigenfunctions of $\Delta_{\mathcal{M}}$ as columns. $\Phi_{\mathcal{N}}: n_{\mathcal{N}} \times k_{\mathcal{N}}$ matrix of the first $k_{\mathcal{N}}$ eigenfunctions of $\Delta_{\mathcal{N}}$ as columns.

The functional map matrix:

$$C = \Phi_{\mathcal{N}}^{+} \mathbf{T}^{T} \Phi_{\mathcal{M}} \qquad ^{+}: \text{left pseudo-inverse.}$$
$$C = \Phi_{\mathcal{N}}^{T} \mathbf{T}^{T} \Phi_{\mathcal{M}} \qquad \text{if} \qquad \Phi_{\mathcal{N}}^{T} \Phi_{\mathcal{N}} = Id$$
$$C = \Phi_{\mathcal{N}}^{T} A_{\mathcal{N}} \mathbf{T}^{T} \Phi_{\mathcal{M}} \qquad \text{if} \qquad \Phi_{\mathcal{N}}^{T} A_{\mathcal{N}} \Phi_{\mathcal{N}} = Id$$

Example Maps in a Reduced Basis

Triangle meshes with pre-computed ponitwise maps



"Good" maps are close to being diagonal

Functional Map algebra

- 1. Map composition becomes matrix multiplication.
- 2. Map inversion is matrix inversion (in fact, transpose).
- 3. Algebraic operations on functional maps are possible.
- E.g. interpolating between two maps with

$$C = \alpha C_1 + (1 - \alpha)C_2.$$
(a) $\alpha = 0$ (b) $\alpha = 0.25$ (c) $\alpha = 0.5$ (d) $\alpha = 0.75$ (e) $\alpha = 1$

Shape Matching

In practice we do not know *C*. Given two objects our goal is to find the correspondence.



How can the functional representation help to compute the map in practice?

Matching via Function Preservation

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

where $f = \sum_{i} a_i \phi_i^{\mathcal{M}}, \quad g = \sum_{i} b_i \phi_i^{\mathcal{N}}.$



Given enough $\{a, b\}$ pairs, we can recover *C* through *a linear least squares system*.

Map Constraints

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

Function preservation constraint is general and includes:

- Attribute (e.g., color) preservation.
- Descriptor preservation (e.g. Gauss curvature).
- Landmark correspondences (e.g. distance to the point).
- Part correspondences (e.g. indicator function).

Commutativity Constraints

Regularizations:

Commutativity with other operators:



Note that the energy: $||CS_{\mathcal{M}} - S_{\mathcal{N}}C||_F^2$ is *quadratic* in *C*.

Regularization

Lemma 1:

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

$$C\Delta_{\mathcal{M}} = \Delta_{\mathcal{N}}C$$

Implies that isometries result in diagonal functional maps.

Regularization

Lemma 2:

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

Basic Pipeline

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

- 1. Compute the first *k* (~80-100) eigenfunctions of the Laplace-Beltrami operator. Store them in matrices: Φ_M , Φ_N
- 2. Compute descriptor functions (e.g., Wave Kernel Signature) on \mathcal{M}, \mathcal{N} . Express them in $\Phi_{\mathcal{M}}, \Phi_{\mathcal{N}}$, as columns of : \mathbf{A}, \mathbf{B}

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

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

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



Conversion to point-to-point

Given a functional map *C*, we would like to convert to to a point-to-point map. $q: \mathcal{N} \to \mathbb{R}$



Option 1: declare $T(x) = \arg \max_{y} \Phi_{\mathcal{N}} C \delta_{x}$ Problems: high computational complexity $O(n_{\mathcal{M}} n_{\mathcal{N}})$, low accuracy.

Conversion to point-to-point

Given a functional map *C*, we would like to convert to to a point-to-point map. $q: \mathcal{N} \to \mathbb{R}$



Option 2: declare $T(x) = \underset{y}{\operatorname{arg\,min}} \|\delta_y - C\delta_x\|_2$

Advantages: computational complexity $O(n_M \log n_N)$, higher accuracy (e.g., works with the identity map).

Incorporating Orthonormality

In many practical situations we would expect a volumepreserving map, which implies:

$$C^T C = Id$$

Option: use post-processing to enforce this constraint. Iterate:

- 1. Compute the point-to-point map *T*.
- 2. Solve for the functional map: $\underset{C, \text{ s.t. } C^T C = Id}{\arg \min} \sum_{x \in \mathcal{M}} \|C\delta_x \delta_{T(x)}\|_2^2$

Exactly the same objective as ICP, but in higher dimension. Can use the same method!

Results







FAUST [Bogo et al. '14]

TOSCA [Bronstein et al. '08]

SCAPE [Anguelov et al. '05]

Results

A very simple method that puts together many constraints and uses 100 basis functions gives reasonable results:



O., Ben-Chen, Solomon, Butscher, Guibas, Functional maps: a flexible representation of maps between shapes, 2012

Segmentation Transfer without P2P

To transfer functions we do not need to convert functional to pointwise maps.

E.g. we can also transfer segmentations: for each segment, transfer its indicator function, and for each point pick the segment that gave the highest value.

