Fully Spectral Partial Shape Matching

Or Litany[1,2] Emanuele Rodolà[3]
Alexander Bronstein[1,2,4] Michael Bronstein[1,2,3]

[1]Tel Aviv University [2]Intel [3]USI Lugano [4]Technion

Eurographics, 27 April 2017

3D sensing applications

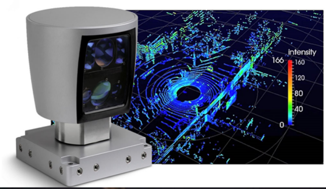

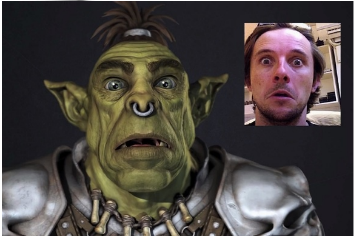

LIDAR Velodyne HDL-64E (as in the Google Car); Intel RealSense R200 3D camera; FaceShift Inc. ; Me ; A cute baby

3D sensing applications

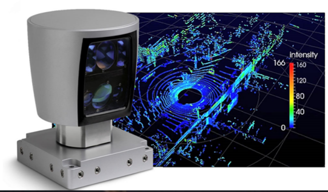

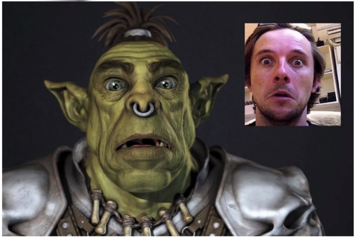

- Non-rigid deformations
- Limited view points

LIDAR Velodyne HDL-64E (as in the Google Car); Intel RealSense R200 3D camera; FaceShift Inc. ; Me ; A cute baby

Shape correspondence problem

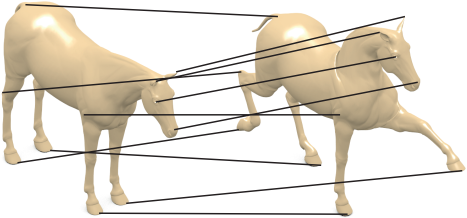

Isometric

Shape correspondence problem

Isometric Partial

Shape correspondence problem

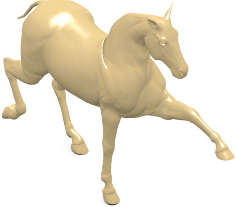

Isometric

Partial

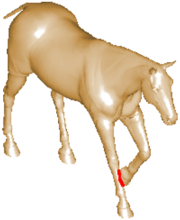

Topological noise

Shape correspondence problem

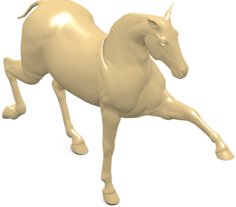

Isometric

Partial

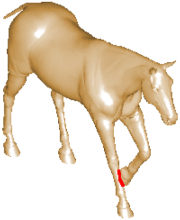

Topological noise

Different representation

Shape correspondence problem

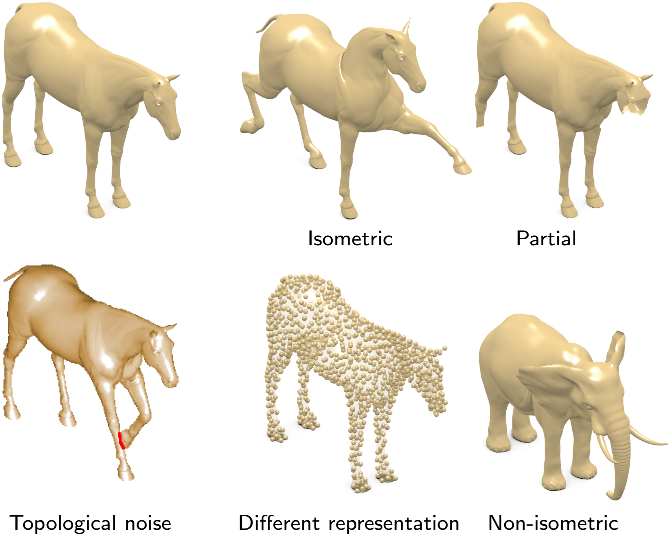

Point-wise maps

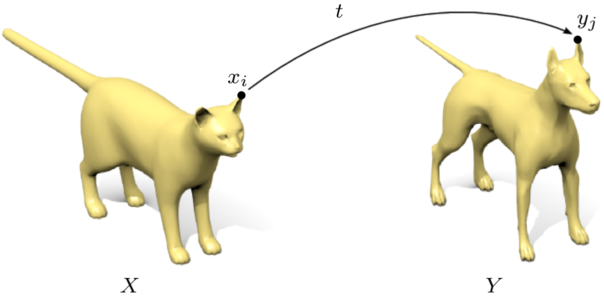

Point-wise maps $t\colon X \to Y$

Functional maps

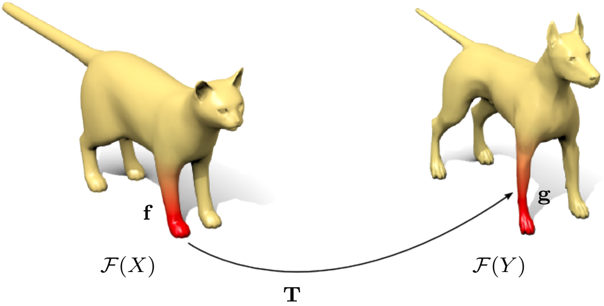

Functional maps $\mathbf{T}\colon \mathcal{F}(X) \to \mathcal{F}(Y)$

Ovsjanikov et al. 2012

Functional correspondence

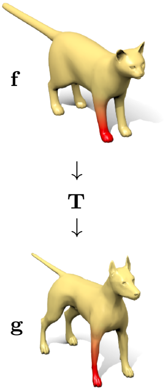

Ovsjanikov et al. 2012

Functional correspondence

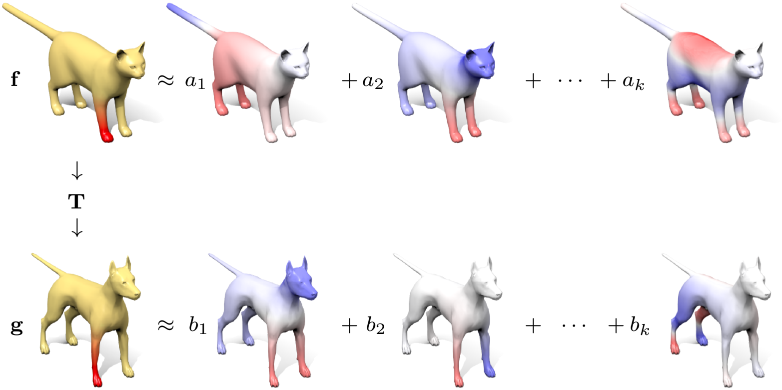

Ovsjanikov et al. 2012

Functional correspondence

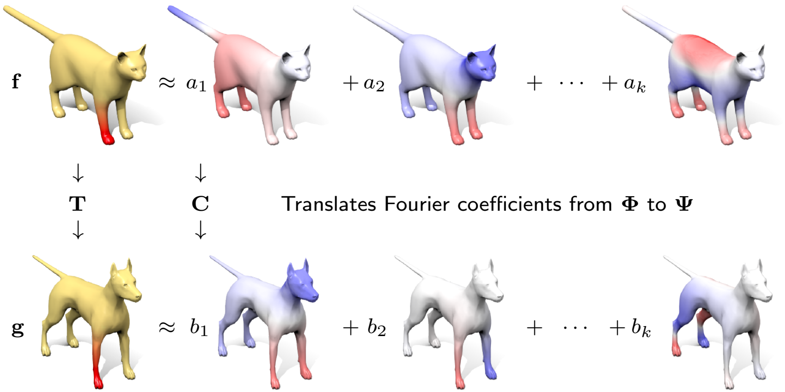

Ovsjanikov et al. 2012

Functional correspondence

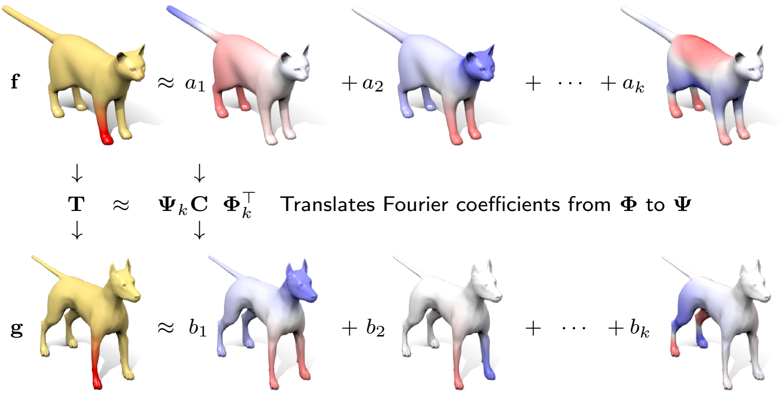

$$\boldsymbol{\Psi}_k^\top \mathbf{g} = \mathbf{C}\boldsymbol{\Phi}_k^\top \mathbf{f}$$

where $\boldsymbol{\Phi}_k = (\phi_1, \ldots, \phi_k)$, $\boldsymbol{\Psi}_k = (\psi_1, \ldots, \psi_k)$ are Laplace-Beltrami eigenbases

Ovsjanikov et al. 2012

Fourier analysis (non-Euclidean spaces)

The Laplacian is invariant to isometries

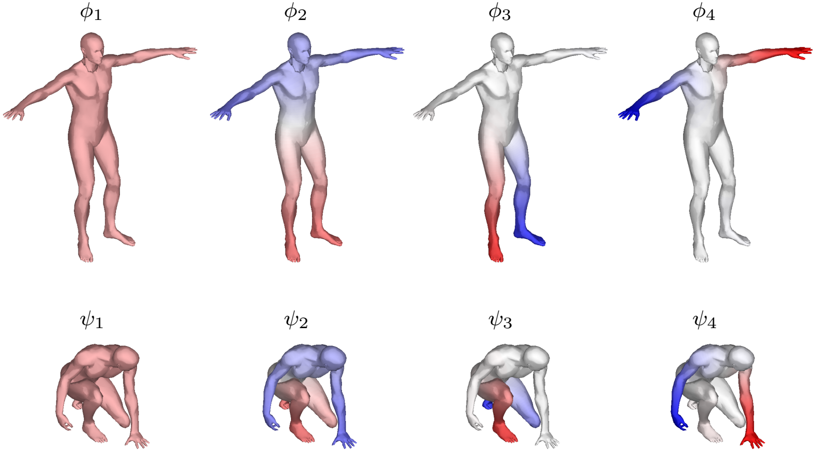

Functional correspondence in Laplacian eigenbases

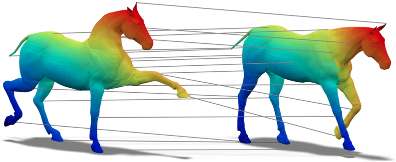

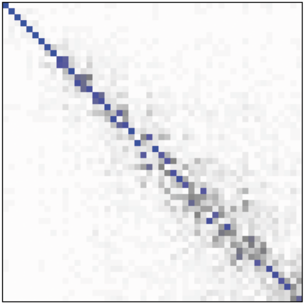

$$\mathbf{C} = \mathbf{\Psi}_k^\top \mathbf{T} \mathbf{\Phi}_k \Rightarrow c_{ij} = \langle \psi_i, T\varphi_j \rangle$$

For isometric simple spectrum shapes, $\mathbf{C}$ is diagonal since $\boldsymbol{\psi}_i = \pm \mathbf{T} \boldsymbol{\phi}_i$

Part-to-full correspondence

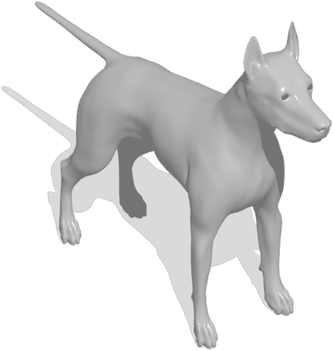

Full model

Partial query

Partial Laplacian eigenvectors

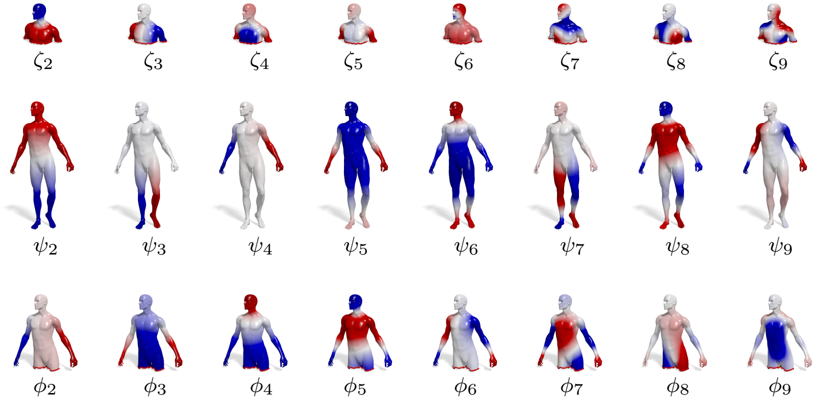

Laplacian eigenvectors of a shape with missing parts
(Neumann boundary conditions)

Rodolà, Cosmo, Bronstein, Torsello, Cremers 2016

Partial Laplacian eigenvectors

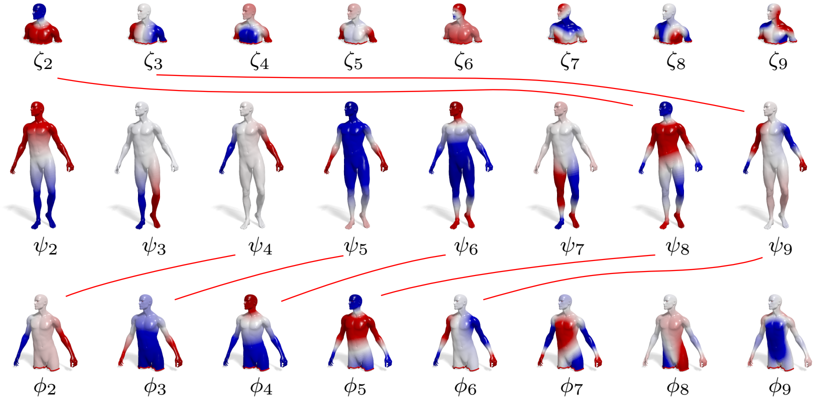

Laplacian eigenvectors of a shape with missing parts
(Neumann boundary conditions)

Rodolà, Cosmo, Bronstein, Torsello, Cremers 2016

Partial Laplacian eigenvectors

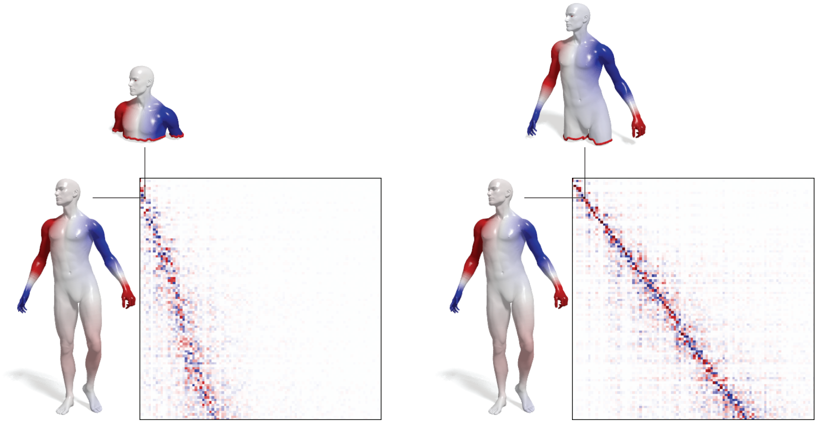

Functional correspondence matrix $\mathbf{C}$
Slope $\approx$ ratio of the two surface areas

Rodolà, Cosmo, Bronstein, Torsello, Cremers 2016

Going fully spectral

PFM has two major drawbacks:

- Explicit spatial indicator $\rightarrow$ runtime is $O(n)$

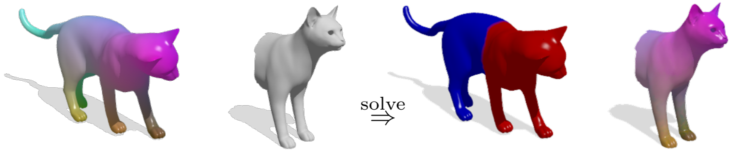

Going fully spectral

PFM has two major drawbacks:

- Explicit spatial indicator $\to$ runtime is $O(n)$
- The partiality prior requires heavy engineering

$$\rho_{\mathrm{part}}(v) = \mu_1 \left(\mathrm{area}(X) - \int_Y \eta(v) dx \right)^2 + \mu_2 \int_Y \xi(v) \| \nabla_Y \eta(v) \| dx$$

$$\xi(v) = \delta \left(\eta(v) - \frac{1}{2} \right)$$

$$\eta(v) = \frac{1}{2} (\tanh(2v - 1) + 1)$$

$$\rho_{\mathrm{corr}}(\mathbf{C}) = \mu_3 \| \mathbf{C} \circ \mathbf{W} \|_{\mathrm{F}}^2 + \mu_4 \sum_{i \neq j} (\mathbf{C}^\top \mathbf{C})_{ij}^2 + \mu_5 \sum_i ((\mathbf{C}^\top \mathbf{C})_{ii} - d_i)^2$$

Going fully spectral

PFM has two major drawbacks:

- Explicit spatial indicator $\rightarrow$ runtime is $O(n)$
- The partiality prior requires heavy engineering

Our idea: "reorder" and spatially localize the eigenfunctions

Going fully spectral

PFM has two major drawbacks:

- Explicit spatial indicator $\rightarrow$ runtime is $O(n)$
- The partiality prior requires heavy engineering

Our idea: "reorder" and spatially localize the eigenfunctions

- No indicator $\rightarrow$ Runtime is $O(k^2)$

Going fully spectral

PFM has two major drawbacks:

- Explicit spatial indicator $\rightarrow$ runtime is $O(n)$
- The partiality prior requires heavy engineering

Our idea: "reorder" and spatially localize the eigenfunctions

- No indicator $\rightarrow$ Runtime is $O(k^2)$
- One-to-one correspondence yields a simple prior

Localized basis functions

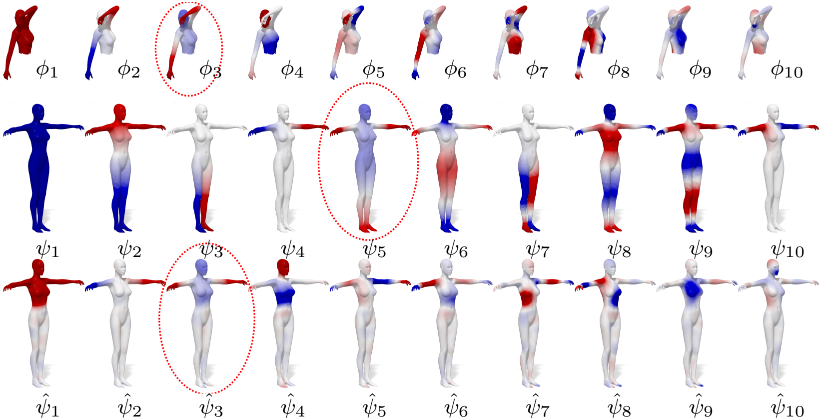

Localization

- Energy minimized in PFM

$$\min_{\mathbf{C},v} \|\mathbf{CA} - \mathbf{B}(v)\| + \rho_{\mathrm{corr}}(\mathbf{C}) + \rho_{\mathrm{part}}(v)$$

$$\begin{aligned} &v : \mathcal{N} \to [0,1] \\ &\mathbf{A} = (\langle \phi_i, f_j \rangle_{\mathcal{M}}) \\ &\mathbf{B}(v) = (\langle \psi_i, v \cdot g_j \rangle_{\mathcal{N}}) \end{aligned}$$

Localization

- Energy minimized in PFM

$$\min_{\mathbf{C},v} \|\mathbf{CA} - \mathbf{B}(v)\| + \rho_{\text{corr}}(\mathbf{C}) + \rho_{\text{part}}(v)$$

- Satisfying the data-term induces a localizing map C

Localization

Localization

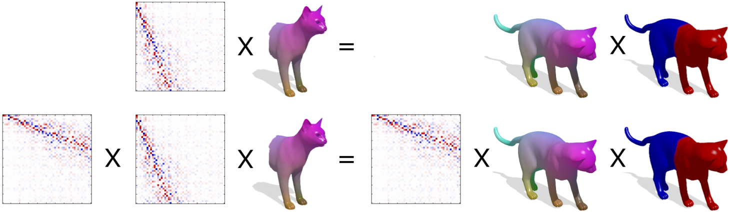

Localization

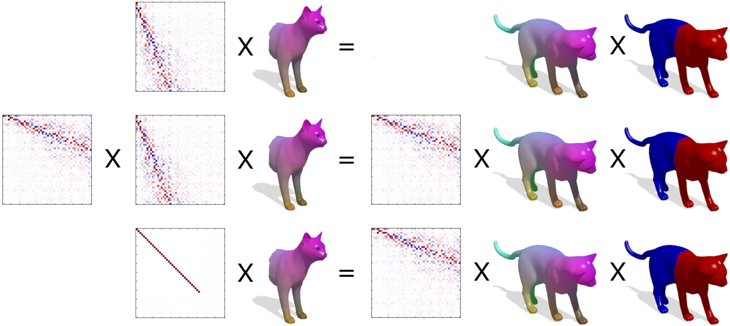

Localization

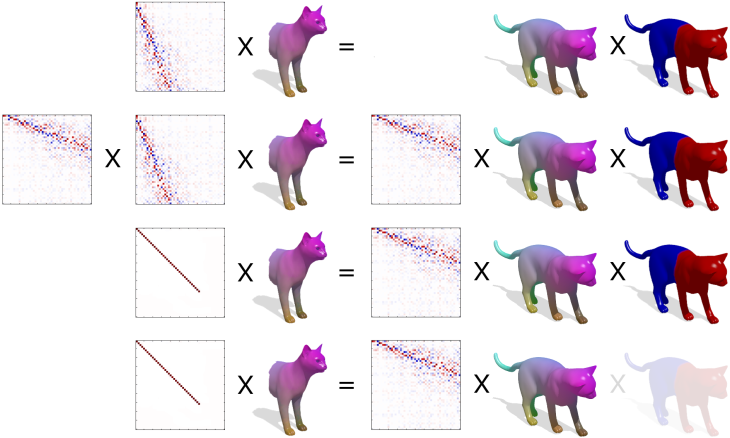

Localization

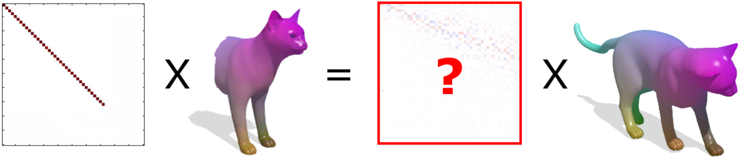

Localization

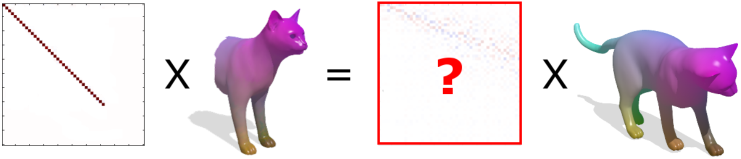

$$\mathbf{W}_r \times \langle \mathbf{\Psi}, \mathbf{F} \rangle = \mathbf{Q}^\top \times \langle \mathbf{\Phi}, \mathbf{G} \rangle$$

Localization

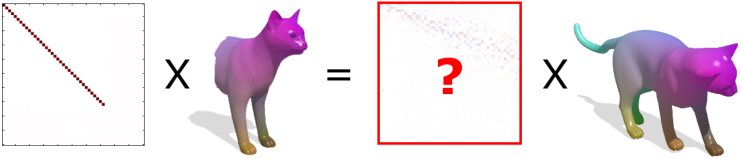

$$\mathbf{W}_r \times \langle \mathbf{\Psi}, \mathbf{F} \rangle = \mathbf{Q}^\top \times \langle \mathbf{\Phi}, \mathbf{G} \rangle$$
$$= \langle \mathbf{Q}^\top \mathbf{\Phi}, \mathbf{G} \rangle$$
$$= \langle \hat{\mathbf{\Phi}}, \mathbf{G} \rangle$$

Localized basis functions

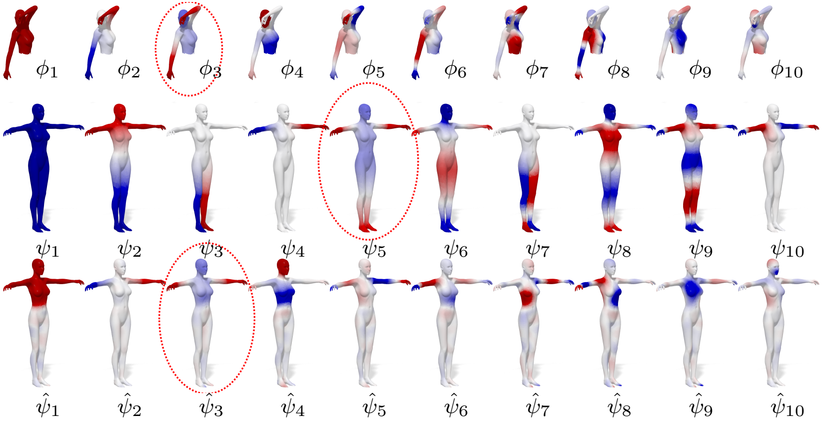

Fully spectral partial correspondence

- Our problem

$$\min_{\mathbf{Q}\in S(k,r)} \operatorname{off}(\mathbf{Q}^\top \mathbf{\Lambda}_{\mathcal{N}} \mathbf{Q}) + \mu \|\mathbf{W}_r \mathbf{A} - \mathbf{Q}^\top \mathbf{B}\|_{2,1}$$

Fully spectral partial correspondence

- Our problem

$$\min_{\mathbf{Q} \in S(k,r)} \operatorname{off}(\mathbf{Q}^\top \mathbf{\Lambda}_{\mathcal{N}} \mathbf{Q}) + \mu \|\mathbf{W}_r \mathbf{A} - \mathbf{Q}^\top \mathbf{B}\|_{2,1}$$

- Compute new basis functions as linear combinations of Laplace-Beltrami eigenfunctions

Fully spectral partial correspondence

- Our problem

$$\min_{\mathbf{Q}\in S(k,r)} \mathrm{off}(\mathbf{Q}^\top \mathbf{\Lambda}_{\mathcal{N}} \mathbf{Q}) + \mu \|\mathbf{W}_r \mathbf{A} - \mathbf{Q}^\top \mathbf{B}\|_{2,1}$$

- Compute new basis functions as linear combinations of Laplace-Beltrami eigenfunctions
- Non-smooth optimization on the Stiefel manifold with $k \times r$ variables

Fully spectral partial correspondence

- Our problem

$$\min_{\mathbf{Q}\in S(k,r)} \text{off}(\mathbf{Q}^\top \mathbf{\Lambda}_{\mathcal{N}} \mathbf{Q}) + \mu \|\mathbf{W}_r \mathbf{A} - \mathbf{Q}^\top \mathbf{B}\|_{2,1}$$

- Compute new basis functions as linear combinations of Laplace-Beltrami eigenfunctions
- Non-smooth optimization on the Stiefel manifold with $k \times r$ variables
- Rank r controls *amount* of partiality

Fully spectral partial correspondence

- Our problem

$$\min_{\mathbf{Q}\in S(k,r)} \operatorname{off}(\mathbf{Q}^\top \boldsymbol{\Lambda}_{\mathcal{N}} \mathbf{Q}) + \mu \|\mathbf{W}_r \mathbf{A} - \mathbf{Q}^\top \mathbf{B}\|_{2,1}$$

- Compute new basis functions as linear combinations of Laplace-Beltrami eigenfunctions
- Non-smooth optimization on the Stiefel manifold with $k \times r$ variables
- Rank r controls *amount* of partiality
- Descriptors control *location* of partiality

Fully spectral partial correspondence

- Our problem

$$\min_{\mathbf{Q}\in S(k,r)} \operatorname{off}(\mathbf{Q}^\top \boldsymbol{\Lambda}_{\mathcal{N}} \mathbf{Q}) + \mu \|\mathbf{W}_r \mathbf{A} - \mathbf{Q}^\top \mathbf{B}\|_{2,1}$$

- Compute new basis functions as linear combinations of Laplace-Beltrami eigenfunctions
- Non-smooth optimization on the Stiefel manifold with $k \times r$ variables
- Rank r controls *amount* of partiality
- Descriptors control *location* of partiality
- Two-sided partiality

$$\min_{(\mathbf{P},\mathbf{Q})\in S^2(k,r)} \operatorname{off}(\mathbf{P}^\top \boldsymbol{\Lambda}_{\mathcal{M}} \mathbf{P}) + \operatorname{off}(\mathbf{Q}^\top \boldsymbol{\Lambda}_{\mathcal{N}} \mathbf{Q}) + \mu \|\mathbf{P}^\top \mathbf{A} - \mathbf{Q}^\top \mathbf{B}\|_{2,1}$$

Importance of descriptors and rank

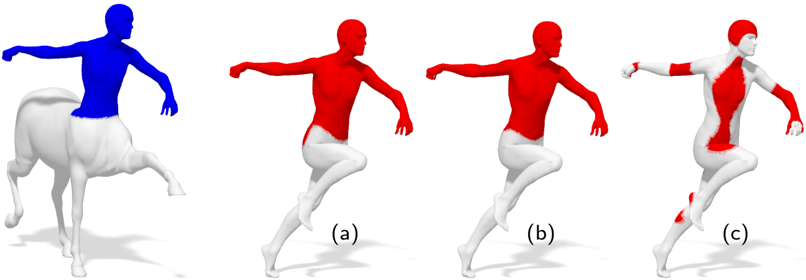

Importance of descriptors and rank

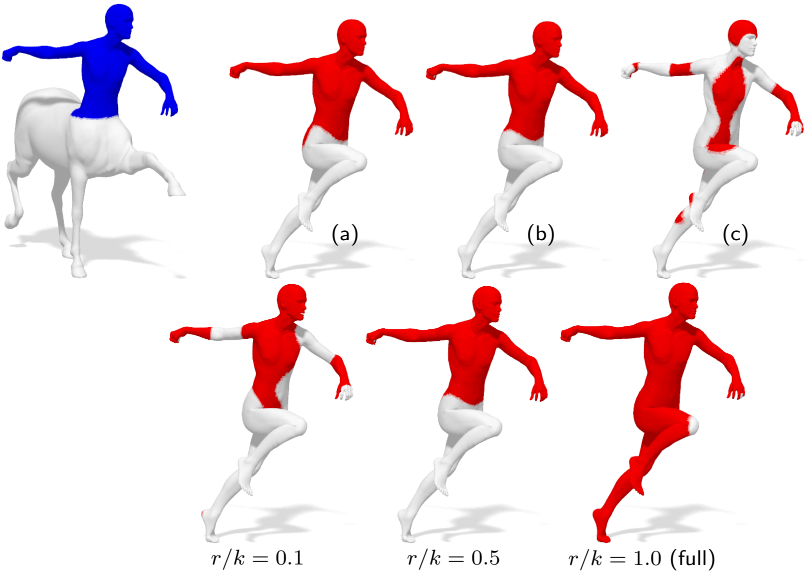

Geometric interpretation

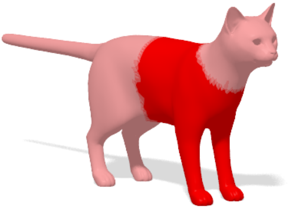

Full shape $\mathcal{N}$

Part $\mathcal{M}$

Geometric interpretation

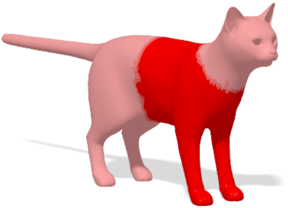

Full shape $\mathcal{N}$

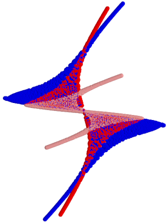

$\phi_2^{\mathcal{M}}, \phi_3^{\mathcal{M}}$ and $\phi_3^{\mathcal{N}}, \phi_5^{\mathcal{N}}$

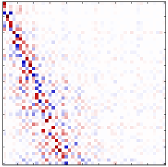

Laplacian eigenbasis

Part $\mathcal{M}$

Geometric interpretation

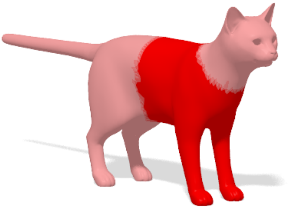

Full shape $\mathcal{N}$

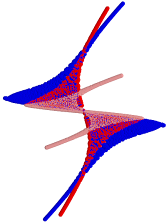

$\phi_2^{\mathcal{M}}, \phi_3^{\mathcal{M}}$ and $\phi_3^{\mathcal{N}}, \phi_5^{\mathcal{N}}$

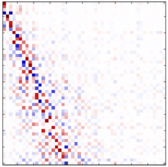

Laplacian eigenbasis

Part $\mathcal{M}$

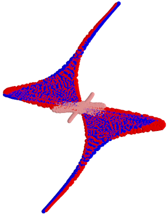

$\phi_2^{\mathcal{M}}, \phi_3^{\mathcal{M}}$ and $\hat{\phi}_2^{\mathcal{N}}, \hat{\phi}_3^{\mathcal{N}}$

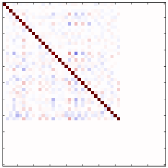

New basis

Animation

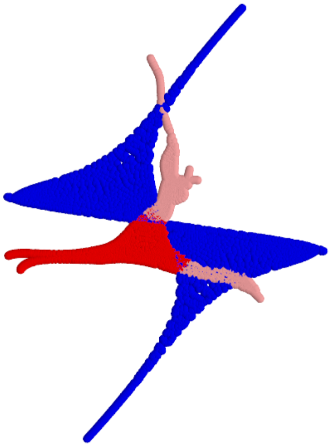

Convergence example

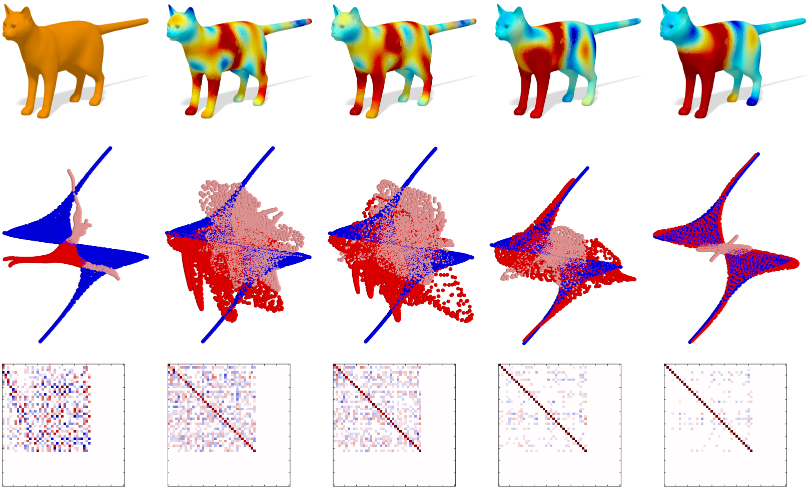

Initialization	75	150	700	4000

Increasing partiality

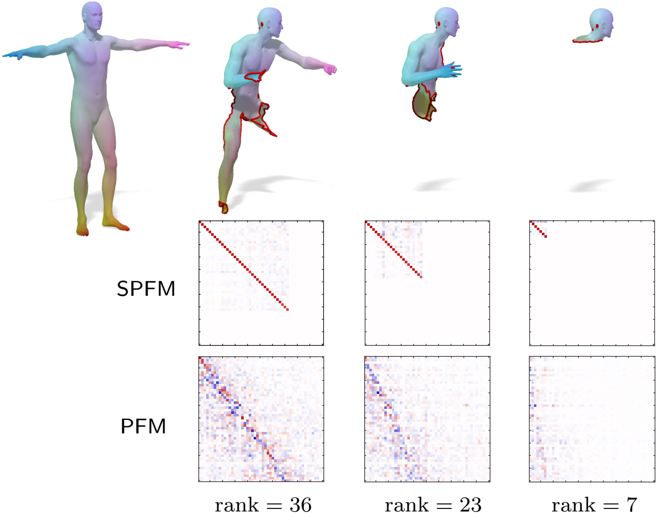

Robustness

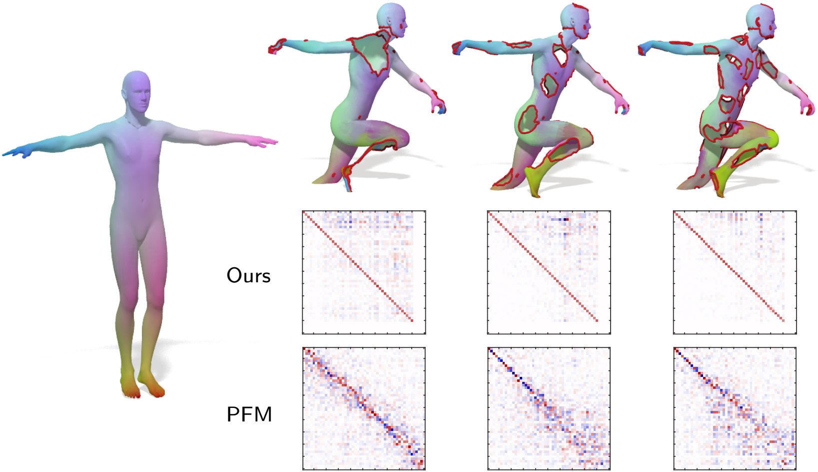

Runtime

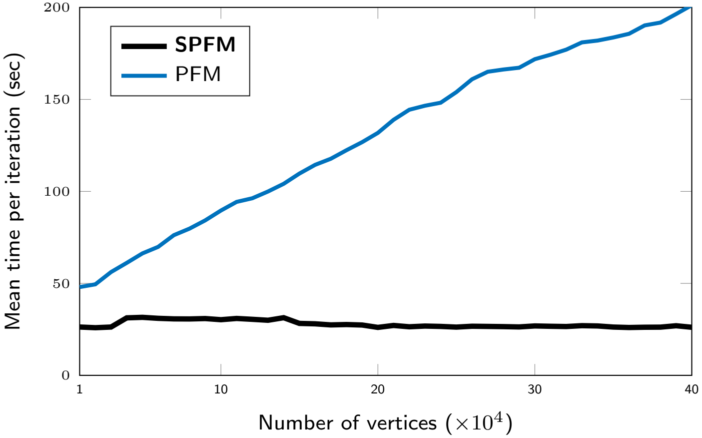
SPFM
PFM
Mean time per iteration (sec)
200
150
100
50
0
1
10
20
30
40
Number of vertices ($\times 10^4$)

SHREC'16 Partiality

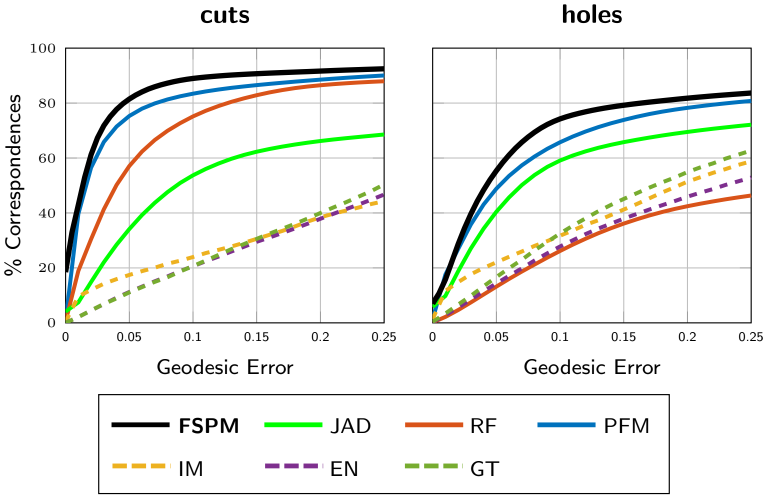

SHREC'16 Topology

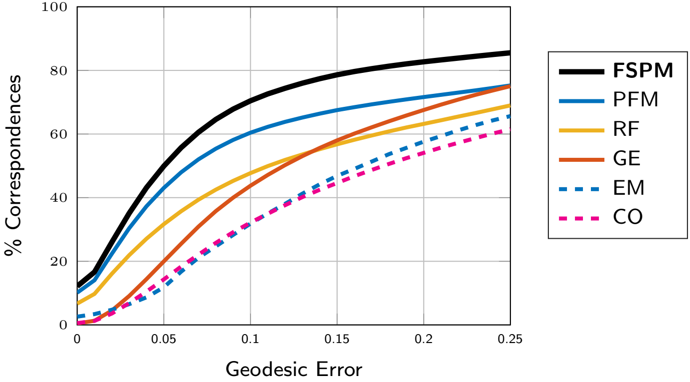

Correspondence examples: topological noise

data: Bogo et al. 2014 (FAUST)

Correspondence examples: topological noise

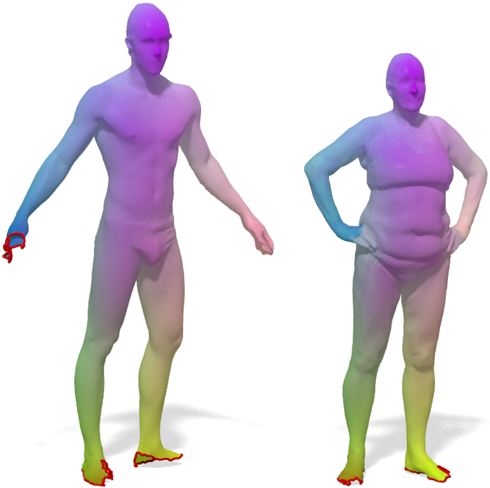

data: Bogo et al. 2014 (FAUST)

Correspondence examples: topological noise

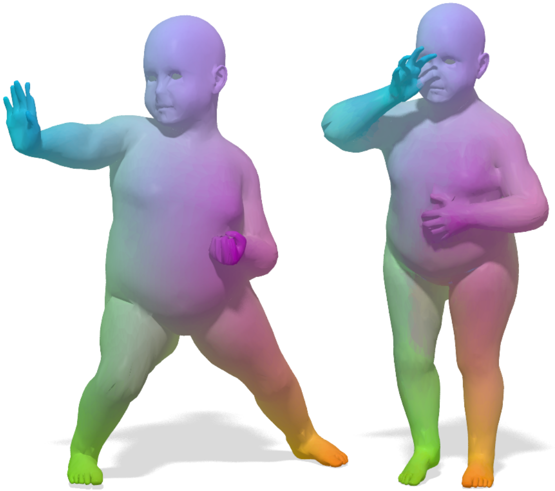

data: Lähner et al. 2016 (SHREC)

Correspondence examples: topological noise

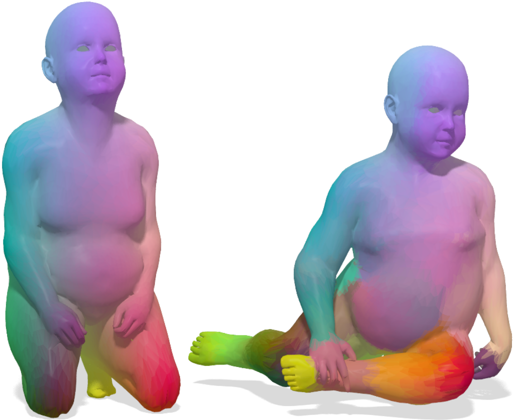

data: Lähner et al. 2016 (SHREC)

Partiality

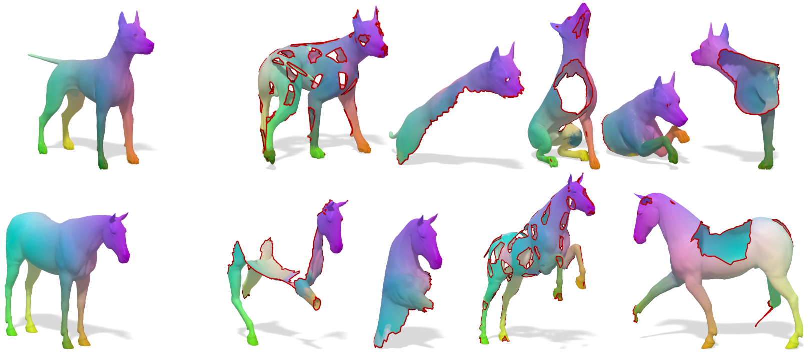

data: Cosmo et al. 2016 (SHREC)

Failure cases

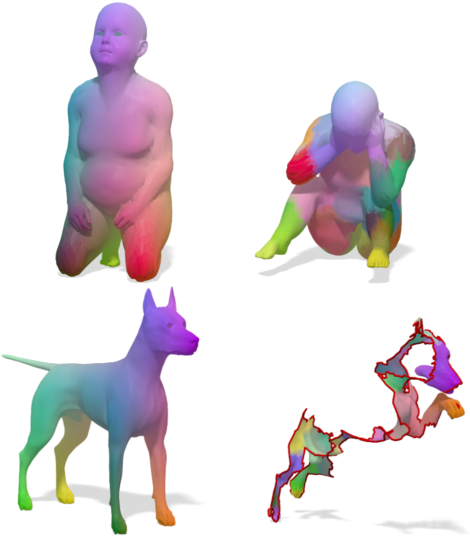

Summary

- Simpler: localization is attained in the spectral domain

Summary

- Simpler: localization is attained in the spectral domain
- Faster: constant complexity (does not depend on shape size)

Summary

- Simpler: localization is attained in the spectral domain
- Faster: constant complexity (does not depend on shape size)
- Better: state of the art results on challenging benchmarks

Summary

- Simpler: localization is attained in the spectral domain
- Faster: constant complexity (does not depend on shape size)
- Better: state of the art results on challenging benchmarks
- Potentially: a nifty end-to-end architecture for Deep Learning of descriptors

Thank you!

Code available at https://github.com/orlitany