Isospectralization

or how to hear shape, style and correspondence

Emanuele Rodolà

Sapienza Università di Roma

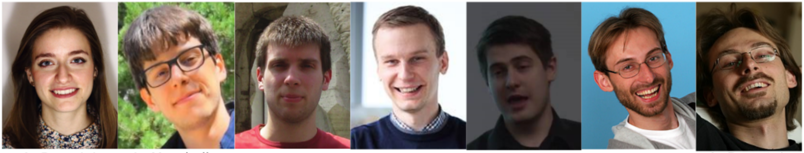

A. Rampini · L. Moschella · L. Cosmo · M. Ovsjanikov · M. Panine · M. Bronstein · A. Bronstein

Sapienza Università di Roma · Università Ca'Foscari Venezia · École Polytechnique Université Paris-Saclay · Imperial College London

IPAM Workshop on Shape Analysis · UCLA, 16 April 2019

Marvel's Daredevil, 2003

CAN ONE HEAR THE SHAPE OF A DRUM?

MARK KAC, The Rockefeller University, New York

To George Eugene Uhlenbeck on the occasion of his sixty-fifth birthday

> "La Physique ne nous donne pas seulement l'occasion de résoudre des problèmes . . . , elle nous fait presentir la solution." H. POINCARÉ.

Wave equation

Wave equation

$$\frac{\partial^2 u(x,y;t)}{\partial t^2} = \Delta u(x,y;t)$$

Wave equation

$$\frac{\partial^2 u(x,y;t)}{\partial t^2} = \Delta u(x,y;t)$$

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

Wave equation

$$\frac{\partial^2 u(x,y;t)}{\partial t^2} = \Delta u(x,y;t)$$

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

$$u(x,y;t) = \sum_j d_j \phi_j(x,y) \left(\cos\left(\sqrt{\lambda_j} t\right) + i \sin\left(\sqrt{\lambda_j} t\right)\right)$$

$\lambda =$ vibration frequencies

$$\lambda = \text{vibration frequencies}$$

$$\phi = \text{vibration modes}$$

isometric $\implies$ isospectral

$$\textbf{isometric} \overset{?}{\Longleftrightarrow} \textbf{isospectral}$$

Shape DNA

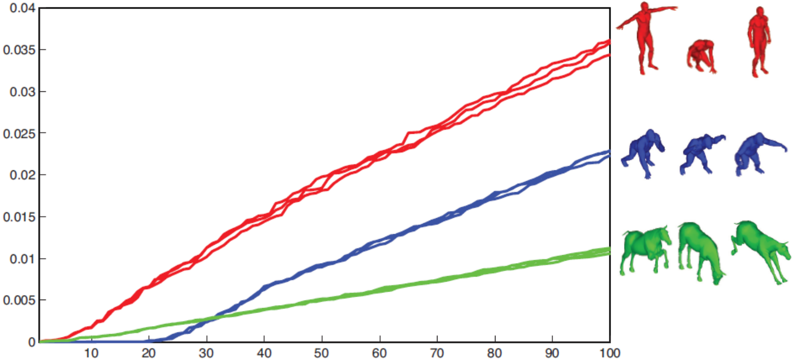

Laplacian spectrum used as deformation-invariant shape descriptor

Reuter et al. 2005

Can we reconstruct the shape from the sequence (λ_i)?

Can we reconstruct the shape from the sequence (λ_i)?

$$\int_{\mathcal{M}} \|\nabla \phi_i(x)\|^2 dx = \lambda_i$$

Not technically...

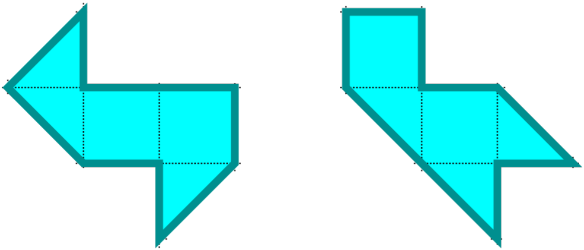

"One cannot hear the shape of the drum", Gordon et al., 1991

Not technically...

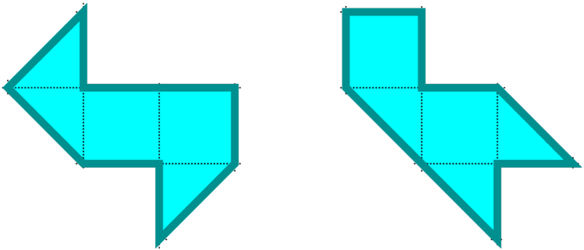

Claim: Isospectral non-isometric shapes are rare

"One cannot hear the shape of the drum", Gordon et al., 1991

Mickey-from-spectrum

In practice, it seems we can!

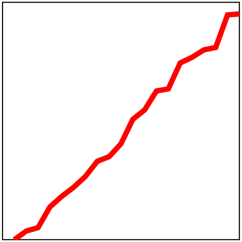

Input

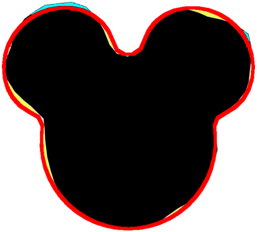

Output

Cosmo, Panine, Rampini, Ovsjanikov, Bronstein, Rodolà (CVPR 2019)

Discretization: shapes

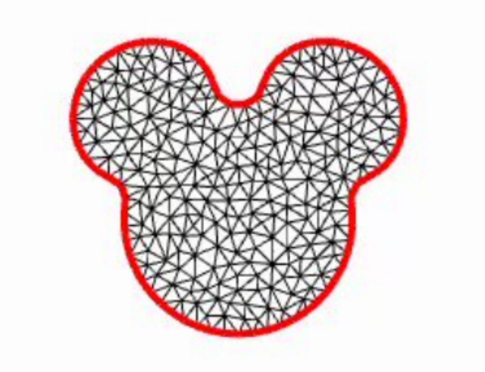

Discretization: Laplace-Beltrami operator

1^{st} order FEM discretization yields:

$$\mathbf{L} = \mathbf{A}^{-1}\mathbf{S}$$

where $\mathbf{S} = (s_{ij})$ is the stiffness matrix and $\mathbf{A} = (a_{ij})$ is the mass matrix.

Discretization: Laplace-Beltrami operator

1^{st} order FEM discretization yields:

$$\mathbf{L} = \mathbf{A}^{-1}\mathbf{S}$$

where $\mathbf{S} = (s_{ij})$ is the stiffness matrix and $\mathbf{A} = (a_{ij})$ is the mass matrix.

$$s_{ij} = \begin{cases} \frac{-\ell_{ij}^2 + \ell_{jk}^2 + \ell_{ki}^2}{8A_{ijk}} + \frac{-\ell_{ij}^2 + \ell_{jh}^2 + \ell_{hi}^2}{8A_{ijh}} & \text{if } e_{ij} \in E_i \\ \frac{-\ell_{ij}^2 + \ell_{jh}^2 + \ell_{hi}^2}{8A_{ijh}} & \text{if } e_{ij} \in E_b \\ -\sum_{k \neq i} s_{ik} & \text{if } i = j \end{cases}$$

$$a_{ij} = \begin{cases} \frac{1}{3} \sum_{lk:ilk \in F} A_{ilk} & \text{if } i = j \\ 0 & \text{if } i \neq j \end{cases}$$

Discretization: Laplace-Beltrami operator

1^{st} order FEM discretization yields:

$$\mathbf{L} = \mathbf{A}^{-1}\mathbf{S}$$

where $\mathbf{S} = (s_{ij})$ is the stiffness matrix and $\mathbf{A} = (a_{ij})$ is the mass matrix.

$$s_{ij} = \begin{cases} \frac{-\ell_{ij}^2 + \ell_{jk}^2 + \ell_{ki}^2}{8A_{ijk}} + \frac{-\ell_{ij}^2 + \ell_{jh}^2 + \ell_{hi}^2}{8A_{ijh}} & \text{if } e_{ij} \in E_i \\ \frac{-\ell_{ij}^2 + \ell_{jh}^2 + \ell_{hi}^2}{8A_{ijh}} & \text{if } e_{ij} \in E_b \\ -\sum_{k \neq i} s_{ik} & \text{if } i = j \end{cases}$$

$$a_{ij} = \begin{cases} \frac{1}{3} \sum_{lk:ilk \in F} A_{ilk} & \text{if } i = j \\ 0 & \text{if } i \neq j \end{cases}$$

$$\ell_{ij}(\mathbf{X}) = \|\mathbf{X}^i - \mathbf{X}^j\|_2$$

Optimization

$$\min_{\mathbf{X}\in\mathbb{R}^{n\times d}} \|\boldsymbol{\lambda}(\boldsymbol{\Delta}_X(\mathbf{X})) - \boldsymbol{\mu}\|_\omega + \rho_X(\mathbf{X})$$

Optimization

$$\min_{\mathbf{X}\in\mathbb{R}^{n\times d}} \|\boldsymbol{\lambda}(\boldsymbol{\Delta}_X(\mathbf{X})) - \boldsymbol{\mu}\|_\omega + \rho_X(\mathbf{X})$$

We call this procedure isospectralization

Optimization

$$\min_{\mathbf{X}\in\mathbb{R}^{n\times d}} \|\boldsymbol{\lambda}(\boldsymbol{\Delta}_X(\mathbf{X})) - \boldsymbol{\mu}\|_\omega + \rho_X(\mathbf{X})$$

We call this procedure isospectralization

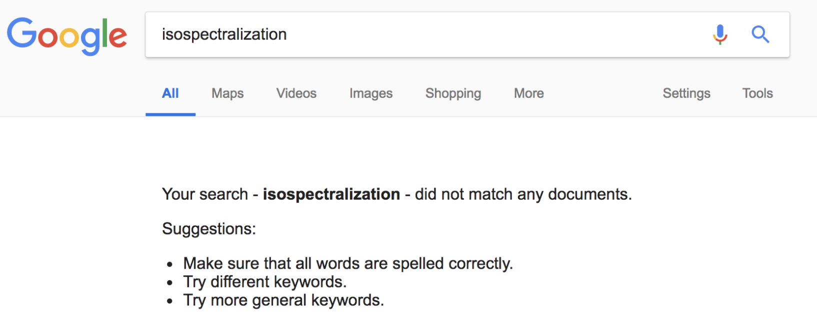

Optimization

$$\min_{\mathbf{X}\in\mathbb{R}^{n\times d}} \|\boldsymbol{\lambda}(\boldsymbol{\Delta}_X(\mathbf{X})) - \boldsymbol{\mu}\|_\omega + \rho_X(\mathbf{X})$$

We call this procedure isospectralization

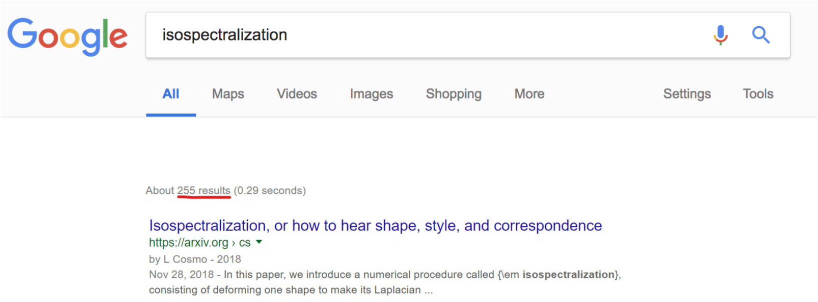

Cosmo, Panine, Rampini, Ovsjanikov, Bronstein, Rodolà (CVPR 2019)

Implementation details

- Data term: weighted norm (higher frequencies attenuated)

$$\|\boldsymbol{\lambda} - \boldsymbol{\mu}\|_\omega^2 = \sum_{i=1}^{k} \frac{1}{\mu_i^2} (\lambda_i - \mu_i)^2$$

Cosmo, Panine, Rampini, Ovsjanikov, Bronstein, Rodolà (CVPR 2019)

Implementation details

- Data term: weighted norm (higher frequencies attenuated)

$$\|\boldsymbol{\lambda} - \boldsymbol{\mu}\|_\omega^2 = \sum_{i=1}^{k} \frac{1}{\mu_i^2}(\lambda_i - \mu_i)^2$$

- Boundary length regularizer:

$$\rho_1(\mathbf{X}) = \sum_{ij \in \mathcal{E}_\mathrm{b}} \ell_{ij}^2(\mathbf{X})$$

Cosmo, Panine, Rampini, Ovsjanikov, Bronstein, Rodolà (CVPR 2019)

Implementation details

- Data term: weighted norm (higher frequencies attenuated)

$$\|\boldsymbol{\lambda} - \boldsymbol{\mu}\|_\omega^2 = \sum_{i=1}^{k} \frac{1}{\mu_i^2}(\lambda_i - \mu_i)^2$$

- Boundary length regularizer:

$$\rho_1(\mathbf{X}) = \sum_{ij \in \mathcal{E}_\mathrm{b}} \ell_{ij}^2(\mathbf{X})$$

- Avoid triangle flips:

$$\rho_2(\mathbf{X}) = \max\left\{ \sum_{ijk \in \mathcal{F}} (\mathbf{R}_{\frac{\pi}{2}}(\mathbf{x}_j - \mathbf{x}_i))^\top (\mathbf{x}_k - \mathbf{x}_i), 0 \right\}$$

Cosmo, Panine, Rampini, Ovsjanikov, Bronstein, Rodolà (CVPR 2019)

Implementation

- Few eigenvalues ($k \leq 30$)

Cosmo, Panine, Rampini, Ovsjanikov, Bronstein, Rodolà (CVPR 2019)

Implementation

- Few eigenvalues ($k \leq 30$)
- Alternate optimization over boundary and interior points

Cosmo, Panine, Rampini, Ovsjanikov, Bronstein, Rodolà (CVPR 2019)

Implementation

- Few eigenvalues ($k \leq 30$)
- Alternate optimization over boundary and interior points
- Re-sampling step once every 200 iterations

Cosmo, Panine, Rampini, Ovsjanikov, Bronstein, Rodolà (CVPR 2019)

Implementation

- Few eigenvalues ($k \leq 30$)
- Alternate optimization over boundary and interior points
- Re-sampling step once every 200 iterations
- Adam optimizer with auto-differentiation

Cosmo, Panine, Rampini, Ovsjanikov, Bronstein, Rodolà (CVPR 2019); Kingma and Ba, 2014

Topology

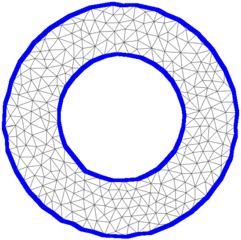

Initialization

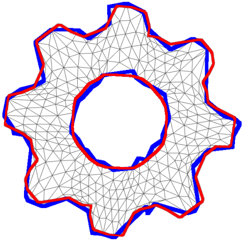

Reconstruction

Reconstruction of surfaces

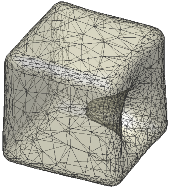

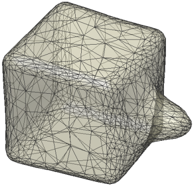

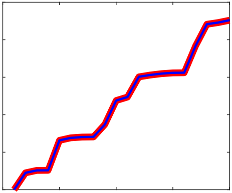

Reconstruction of surfaces

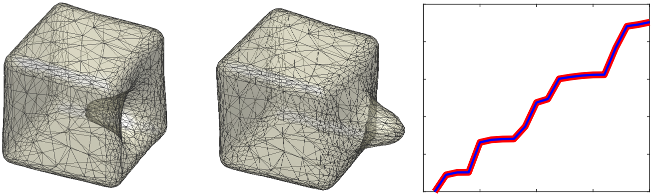

- Smoothness-promoting regularizer:

  $\rho_1(\mathbf{X}) = \|\mathbf{L}\mathbf{X}\|_F^2$

Reconstruction of surfaces

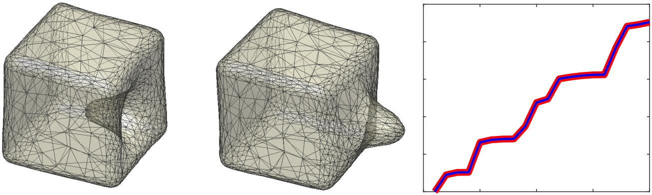

- Smoothness-promoting regularizer:

  $\rho_1(\mathbf{X}) = \|\mathbf{L}\mathbf{X}\|_F^2$

- Volume expansion regularizer to avoid isometric ambiguities:

  $\rho_2(\mathbf{X}) = \begin{pmatrix} 1 \\ 1 \\ 1 \end{pmatrix}^\top \sum_{ijk \in F} ((\mathbf{X}^j - \mathbf{X}^i) \times (\mathbf{X}^k - \mathbf{X}^j))(\mathbf{X}^i + \mathbf{X}^j + \mathbf{X}^k)$

Application: non-isometric correspondence

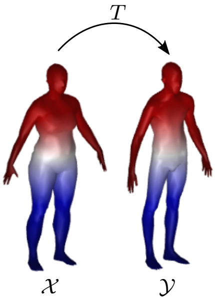

Application: non-isometric correspondence

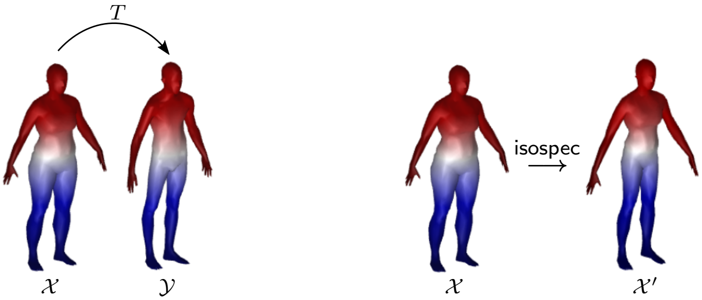

1. Isospectralization: deform $\mathcal{X}$ into $\mathcal{X}'$ whose spectrum $\boldsymbol{\lambda}_{\mathcal{X}'}$ is aligned with $\boldsymbol{\lambda}_{\mathcal{Y}}$

Application: non-isometric correspondence

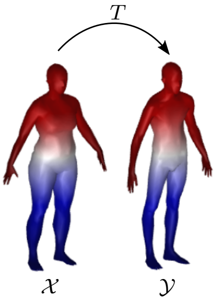

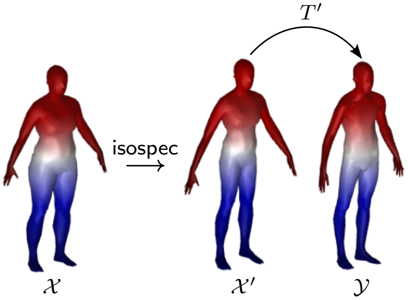

1. Isospectralization: deform $\mathcal{X}$ into $\mathcal{X}'$ whose spectrum $\boldsymbol{\lambda}_{\mathcal{X}'}$ is aligned with $\boldsymbol{\lambda}_{\mathcal{Y}}$
2. Compute the correspondence $T' : \mathcal{X}' \to \mathcal{Y}$ using an existing isometric matching algorithm

Application: non-isometric correspondence

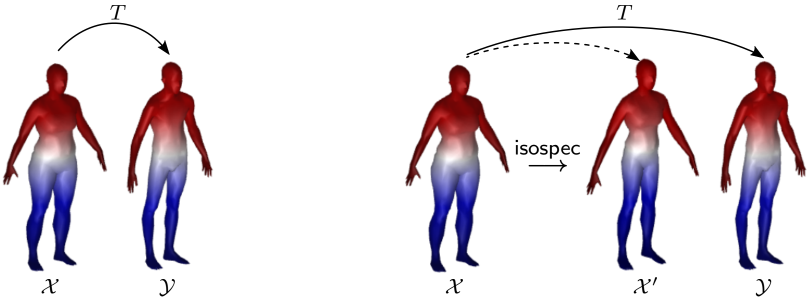

1. Isospectralization: deform $\mathcal{X}$ into $\mathcal{X}'$ whose spectrum $\boldsymbol{\lambda}_{\mathcal{X}'}$ is aligned with $\boldsymbol{\lambda}_{\mathcal{Y}}$
2. Compute the correspondence $T' : \mathcal{X}' \rightarrow \mathcal{Y}$ using an existing isometric matching algorithm
3. Convert T' to $T : \mathcal{X} \rightarrow \mathcal{Y}$ using the identity map between $\mathcal{X}$ and $\mathcal{X}'$.

Correspondence examples

Before isospectralization **After** isospectralization

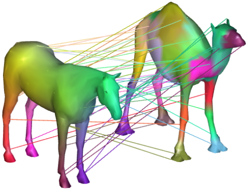

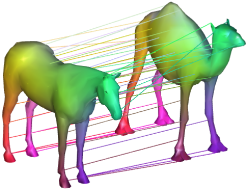

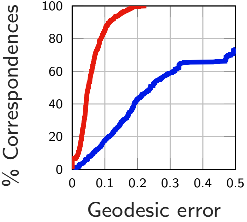

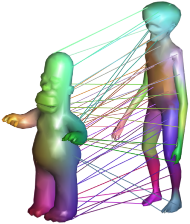

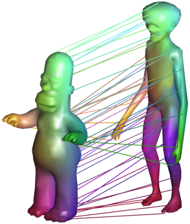

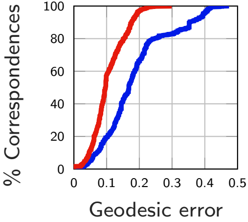

Results: average geodesic error

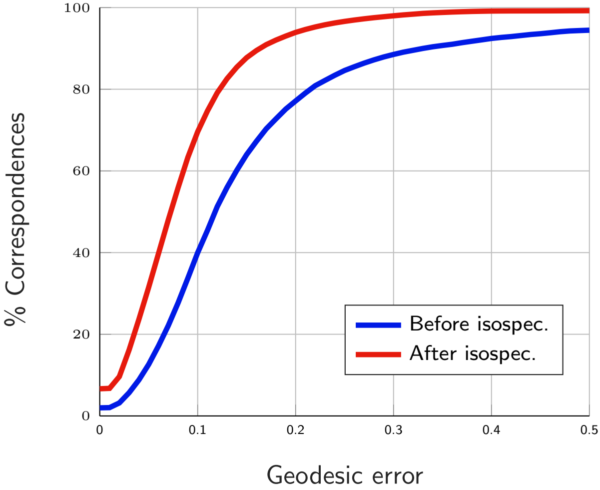

Correspondence examples

Correspondence examples

Why it works

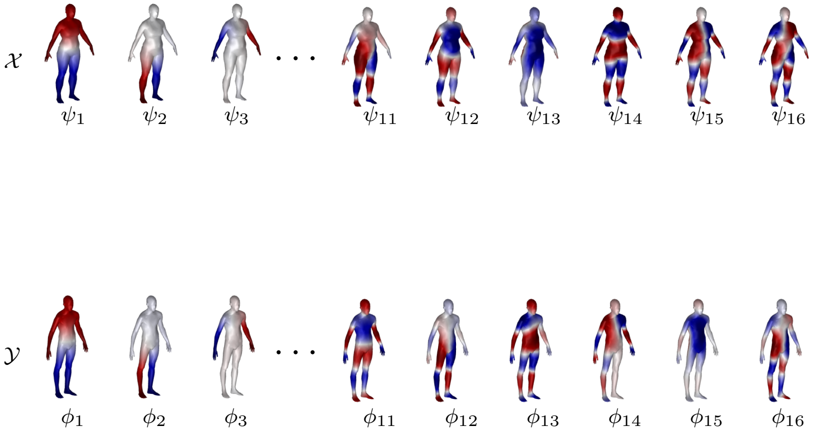
$\mathcal{X}$
ψ_1 ψ_2 ψ_3 $\cdots$ ψ_{11} ψ_{12} ψ_{13} ψ_{14} ψ_{15} ψ_{16}
$\mathcal{Y}$
ϕ_1 ϕ_2 ϕ_3 $\cdots$ ϕ_{11} ϕ_{12} ϕ_{13} ϕ_{14} ϕ_{15} ϕ_{16}

Why it works

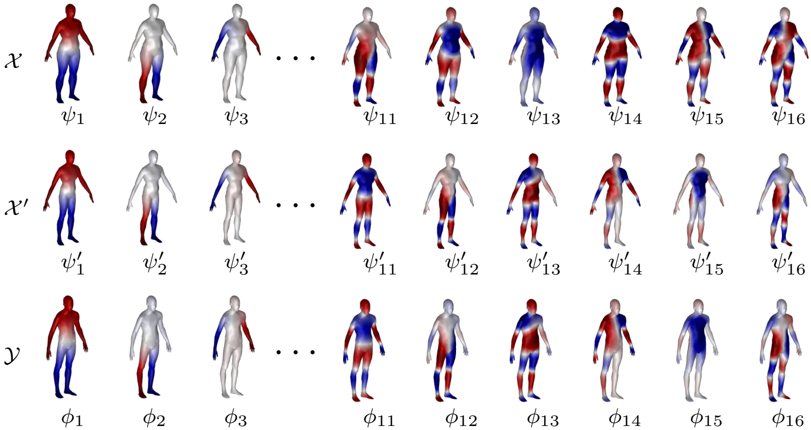

Isospectralization aligns the eigenspaces!

Applications: style transfer

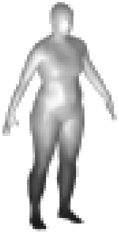

Source pose

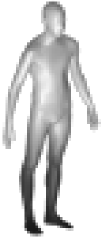

Target style

Applications: style transfer

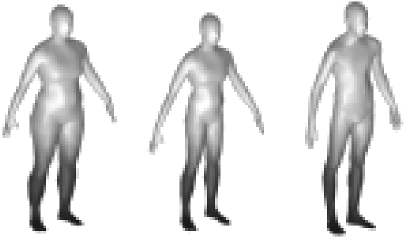

Source pose | Source pose + Target style | Target style

Embedding-free formulation

$$\underbrace{(\Delta + \rho)}_{\text{Hamiltonian } H}$$

Embedding-free formulation

$$\underbrace{(\Delta + \rho)}_{\text{Hamiltonian } H} f(x) = \Delta f(x) + \rho(x) f(x)$$

Embedding-free formulation

$$\underbrace{(\Delta + \rho)}_{\text{Hamiltonian } H} f(x) = \Delta f(x) + \rho(x) f(x)$$

$\rho : \mathcal{X} \to \mathbb{R}_+$ is the characteristic function of a region $\mathcal{R} \subset \mathcal{X}$

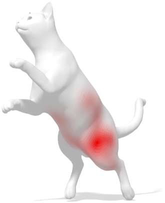

$\rho : \mathcal{X} \to [0, 1]$

Embedding-free formulation

$$\underbrace{(\Delta + \rho)}_{\text{Hamiltonian } H} f(x) = \Delta f(x) + \rho(x) f(x)$$

$\rho : \mathcal{X} \to \mathbb{R}_+$ is the characteristic function of a region $\mathcal{R} \subset \mathcal{X}$

$$H\psi = \lambda\psi$$

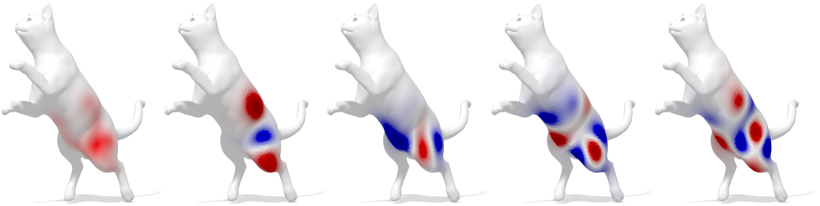

$\rho : \mathcal{X} \to [0, 1]$ ψ_4 ψ_5 ψ_9 ψ_{10}

Melzi, Rodolà, Castellani, Bronstein (CGF 2018); Choukroun et al 2017

Embedding-free formulation

$$\min_{\boldsymbol{\rho} \in [0,1]^{w \times h}} \|\boldsymbol{\lambda}(\boldsymbol{\Delta} + \boldsymbol{\rho}) - \boldsymbol{\mu}\|_w$$

Embedding-free formulation

$$\min_{\boldsymbol{\rho}\in[0,1]^{w\times h}} \|\boldsymbol{\lambda}(\boldsymbol{\Delta}+\boldsymbol{\rho})-\boldsymbol{\mu}\|_w$$

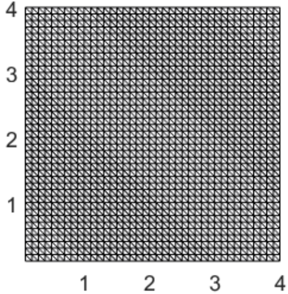

$\boldsymbol{\lambda}$

Eigenvalues of the Hamiltonian $\boldsymbol{\Delta}+\boldsymbol{\rho}$

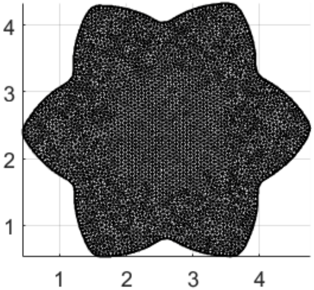

$\boldsymbol{\mu}$

Dirichlet eigenvalues of the Laplacian

Embedding-free formulation

$$\min_{\boldsymbol{\rho} \in [0,1]^{w \times h}} \|\boldsymbol{\lambda}(\boldsymbol{\Delta} + \boldsymbol{\rho}) - \boldsymbol{\mu}\|_w$$

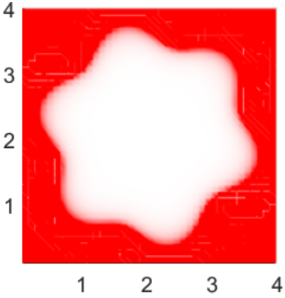

$\boldsymbol{\lambda}$

Eigenvalues of the Hamiltonian $\boldsymbol{\Delta} + \boldsymbol{\rho}$

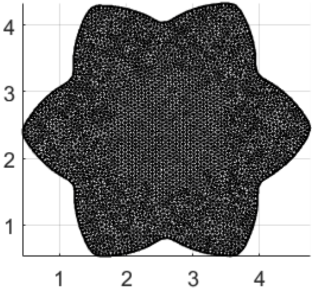

$\boldsymbol{\mu}$

Dirichlet eigenvalues of the Laplacian

Examples

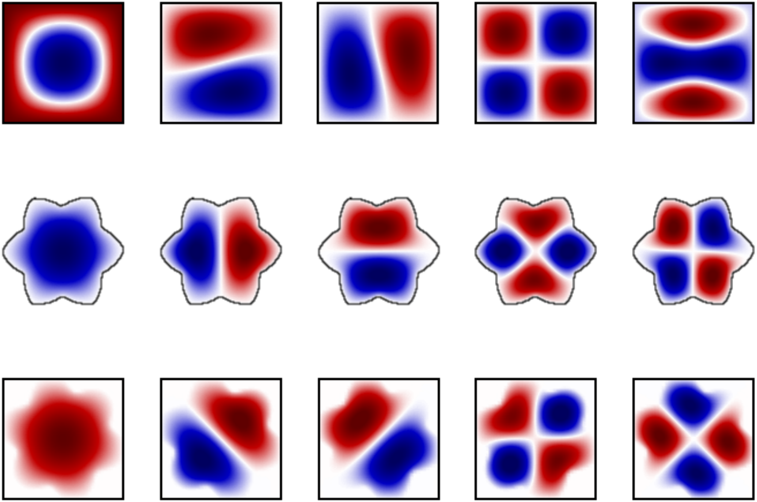

Examples

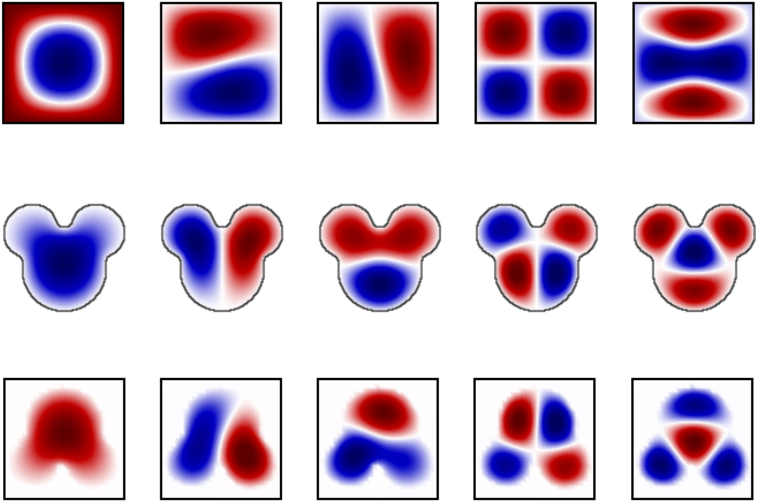

Recovering volumes

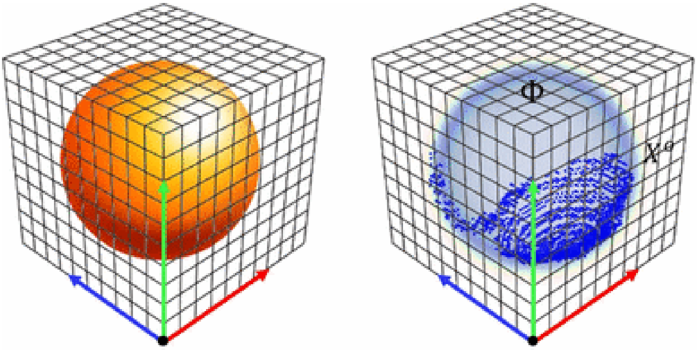

We may recover 3D shapes by modeling them as potentials in $\mathbb{R}^3$

Optimization

$$\min_{\boldsymbol{\rho} \in [0,1]^{w \times h}} \|\boldsymbol{\lambda}(\boldsymbol{\Delta} + \boldsymbol{\rho}) - \boldsymbol{\mu}\|_w$$

Algorithm:

- Initialize $\boldsymbol{\rho} = \mathbf{1}$
- Update $\boldsymbol{\rho}^{(t)} = \boldsymbol{\rho}^{(t-1)} - \alpha \nabla \|\boldsymbol{\lambda}(\boldsymbol{\Delta} + \boldsymbol{\rho}^{(t-1)}) - \boldsymbol{\mu}\|_w$

Optimization

$$\min_{\boldsymbol{\rho} \in [0,1]^{w \times h}} \|\boldsymbol{\lambda}(\boldsymbol{\Delta} + \boldsymbol{\rho}) - \boldsymbol{\mu}\|_w$$

Algorithm:

- Initialize $\boldsymbol{\rho} = \mathbf{1}$
- Update $\boldsymbol{\rho}^{(t)} = \boldsymbol{\rho}^{(t-1)} - 2\alpha(\boldsymbol{\Phi}^{(t-1)} \circ \boldsymbol{\Phi}^{(t-1)})\frac{\boldsymbol{\lambda}^{(t-1)} - \boldsymbol{\mu}}{\boldsymbol{\mu}^2}$

Boils down to solving a sequence of eigenvalue problems

Application: Correspondence-free localization

Application: Correspondence-free localization

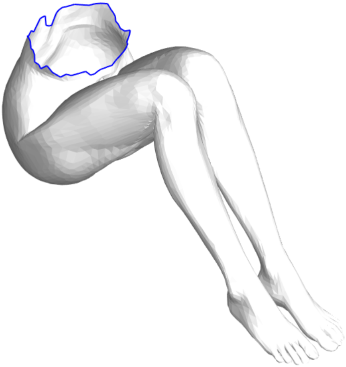

Application: Correspondence-free localization

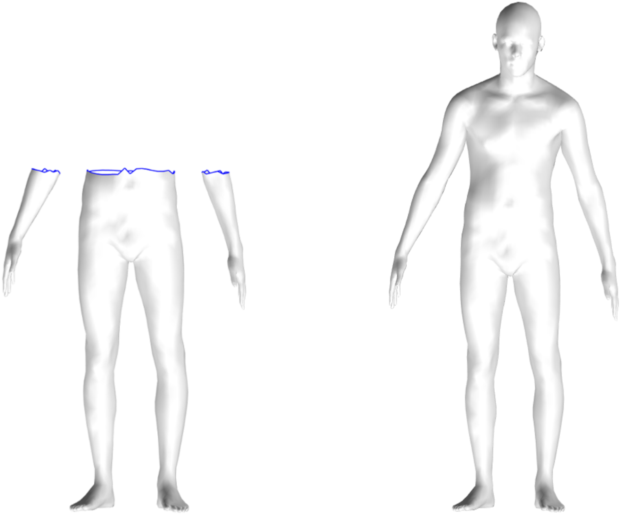

Application: Correspondence-free localization

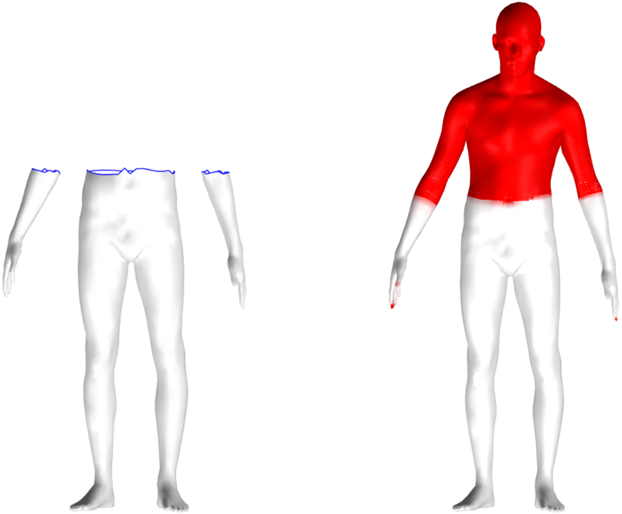

Main value:
Get rid of correspondence

Main value:
Get rid of correspondence

- One may hear the shape of the drum in practice by aligning Laplacian eigenvalues (isospectralization)

Main value:
Get rid of correspondence

- One may hear the shape of the drum in practice by aligning Laplacian eigenvalues (isospectralization)
- Applications in computer vision, graphics, learning?

Main value:
Get rid of correspondence

- One may hear the shape of the drum in practice by aligning Laplacian eigenvalues (isospectralization)
- Applications in computer vision, graphics, learning?
- May be useful for generative models and shape optimization

Thank you!

erc

StG 2018 SPECGEO

Bandwidth and mesh resolution

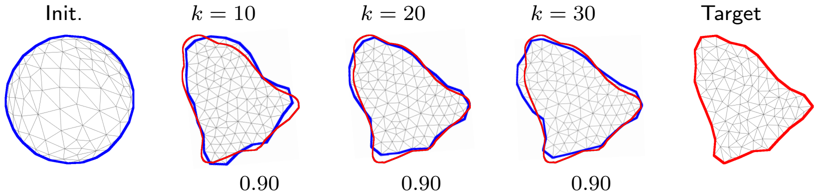

Bandwidth and mesh resolution

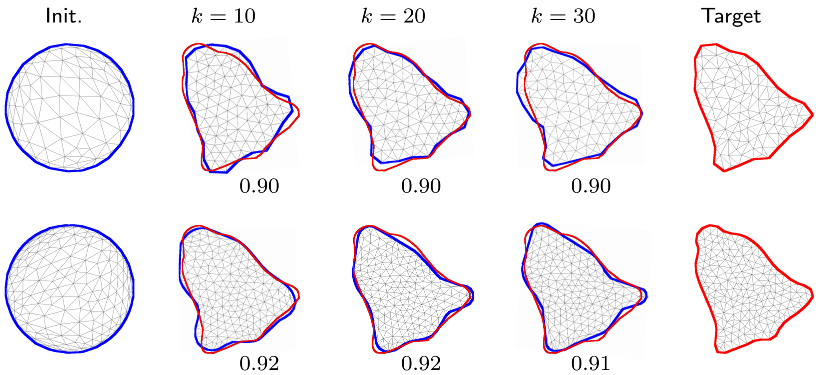

Bandwidth and mesh resolution

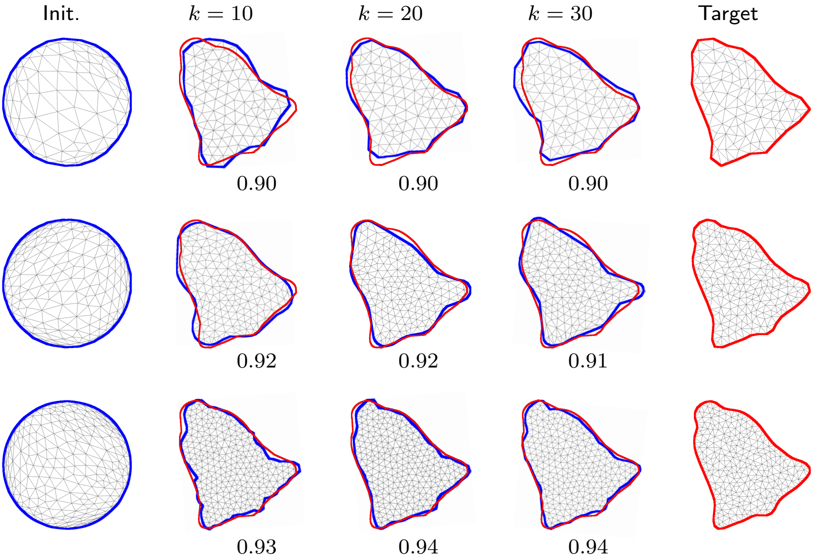

Initialization independence