Matching Deformable Objects in Clutter

Emanuele Rodolà
USI Lugano

Joint work with

   

L. Cosmo A. Torsello J. Masci M.M. Bronstein

Università Ca'Foscari Venezia

USI

3DV 2016, Stanford, 25 October 2016

Shape correspondence problem

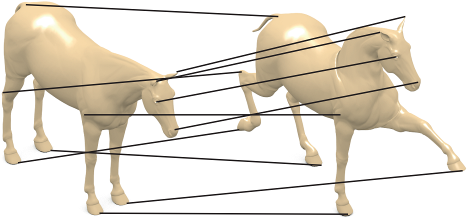

Isometric

Shape correspondence problem

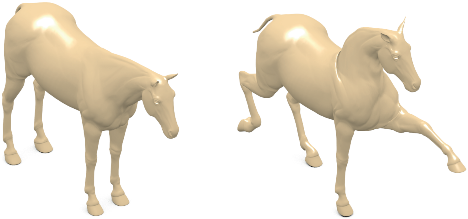

Isometric Partial

Shape correspondence problem

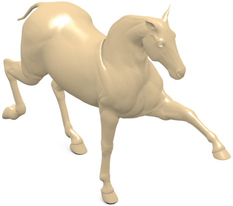

Isometric

Partial

Different representation

Point-wise maps

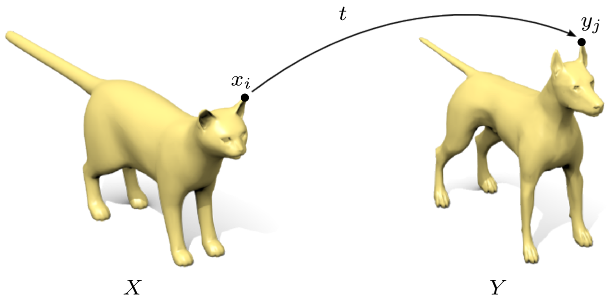

Point-wise maps $t\colon X \to Y$

Functional maps

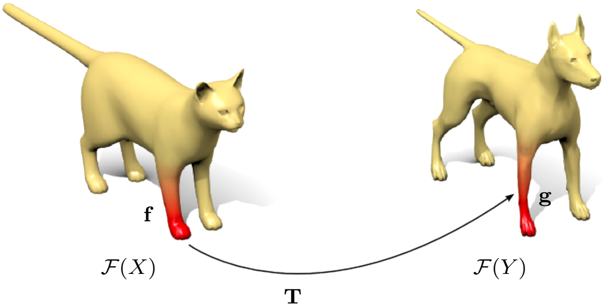

Functional maps $\mathbf{T}\colon \mathcal{F}(X) \to \mathcal{F}(Y)$

Ovsjanikov et al., 2012

Functional correspondence

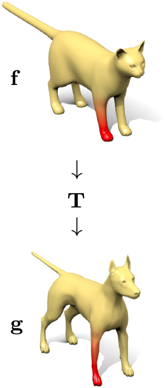

Ovsjanikov et al., 2012

Functional correspondence

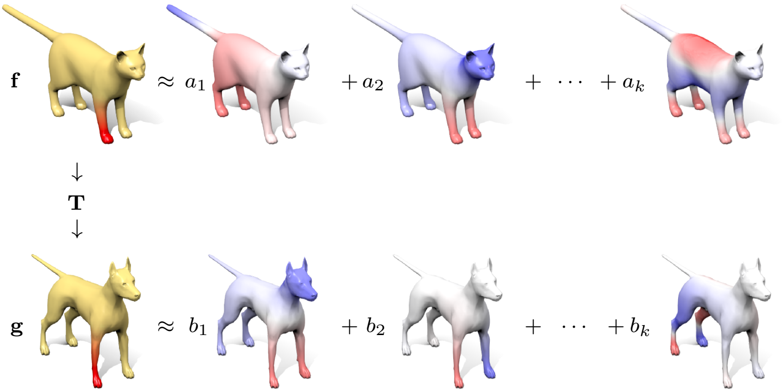

Ovsjanikov et al., 2012

Functional correspondence

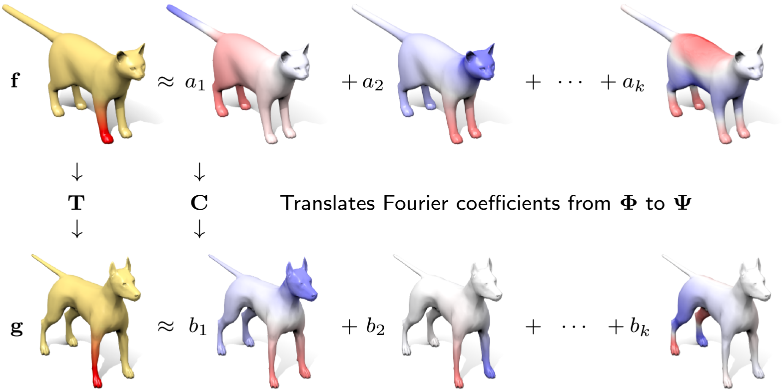

where $\mathbf{\Phi}_k = (\phi_1, \ldots, \phi_k)$, $\mathbf{\Psi}_k = (\psi_1, \ldots, \psi_k)$ are Laplace-Beltrami eigenbases

Ovsjanikov et al., 2012

Functional correspondence

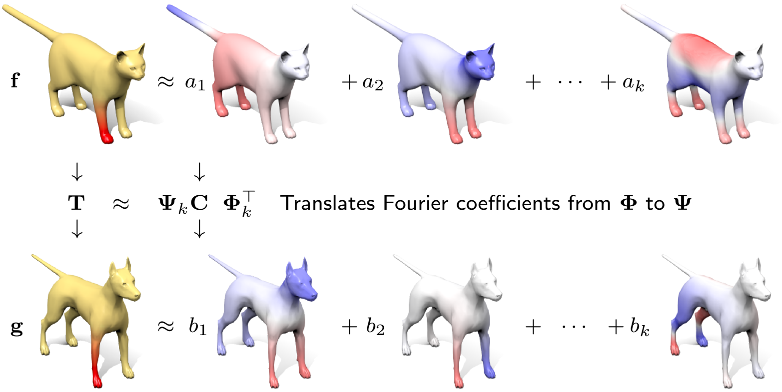

$$\boldsymbol{\Psi}_k^\top \mathbf{g} = \mathbf{C} \boldsymbol{\Phi}_k^\top \mathbf{f}$$

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

Laplacian eigenbases

The Laplacian is invariant to isometries

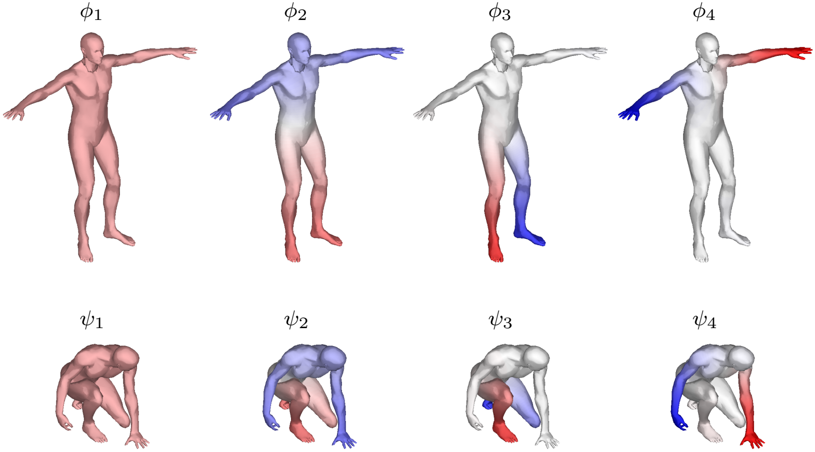

Functional correspondence in Laplacian eigenbases

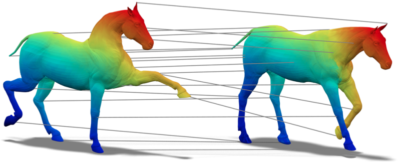

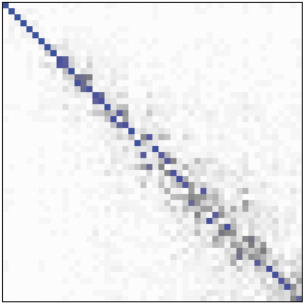

$$\mathbf{C} = \mathbf{\Psi}_k^\top \mathbf{T} \mathbf{\Phi}_k \Rightarrow c_{ij} = \langle \psi_i, T\varphi_j \rangle$$

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

Part-to-full correspondence

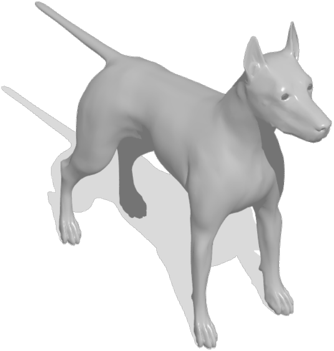

Full model

Partial query

Partial Laplacian eigenvectors

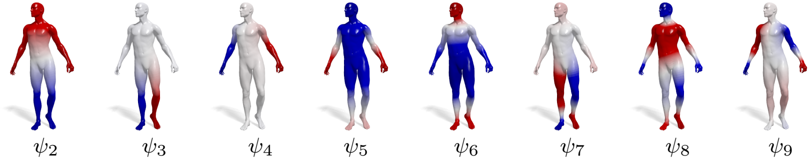

Rodolà, Cosmo, Bronstein, Torsello, Cremers 2016

Partial Laplacian eigenvectors

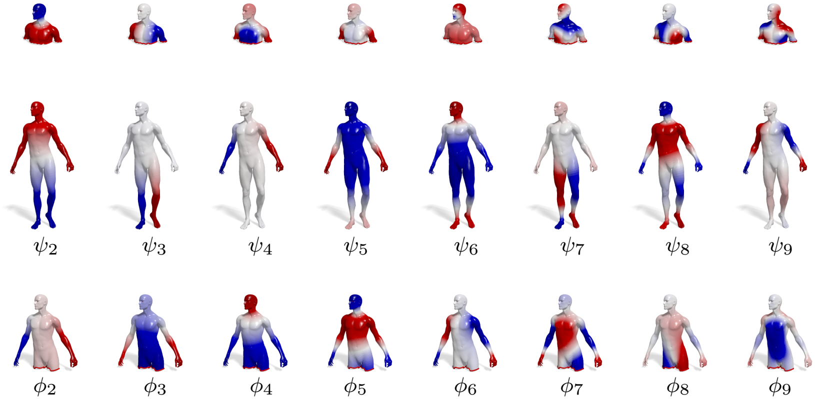

Rodolà, Cosmo, Bronstein, Torsello, Cremers 2016

Partial Laplacian eigenvectors

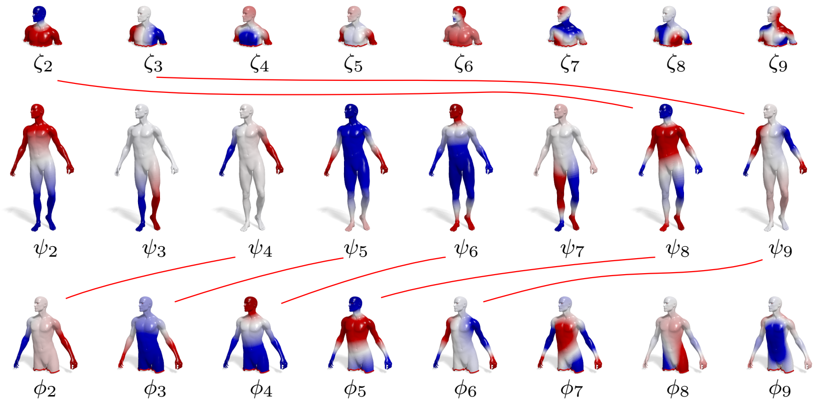

Laplacian eigenvectors of a shape with missing parts
(Neumann boundary conditions)

Rodolà, Cosmo, Bronstein, Torsello, Cremers 2016

Partial Laplacian eigenvectors

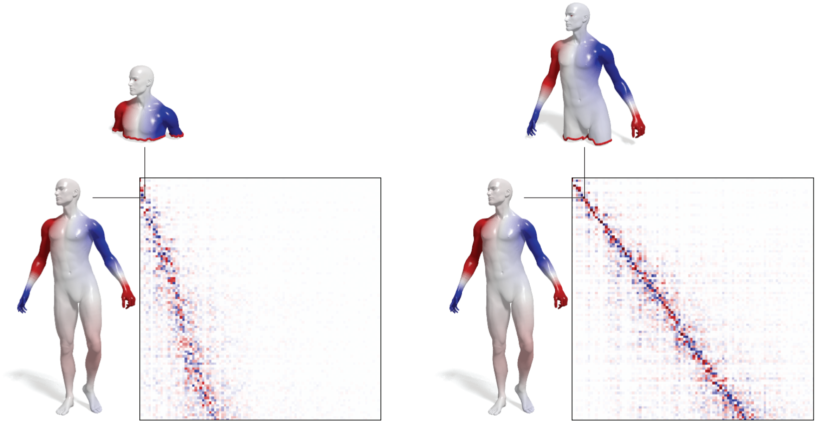

Functional correspondence matrix $\mathbf{C}$

Diagonal angle $\approx$ area ratio of surfaces

Rodolà, Cosmo, Bronstein, Torsello, Cremers 2016

Our setting: Objects in clutter

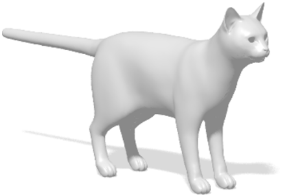

Full model

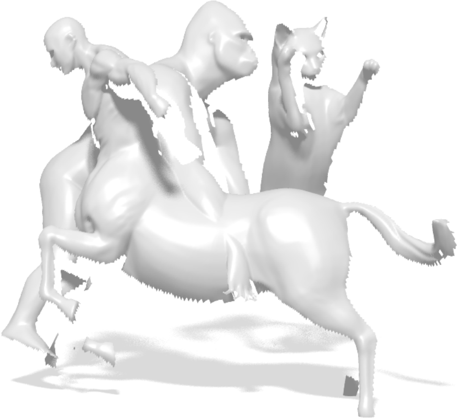

Cluttered partial view

Functional correspondence with clutter

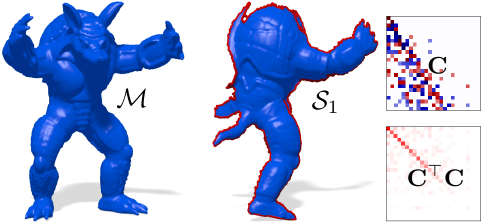

Functional correspondence with clutter

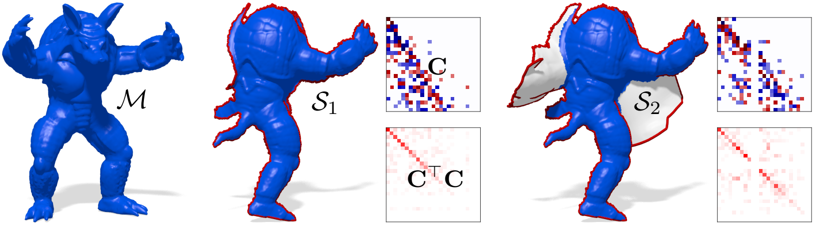

Functional correspondence with clutter

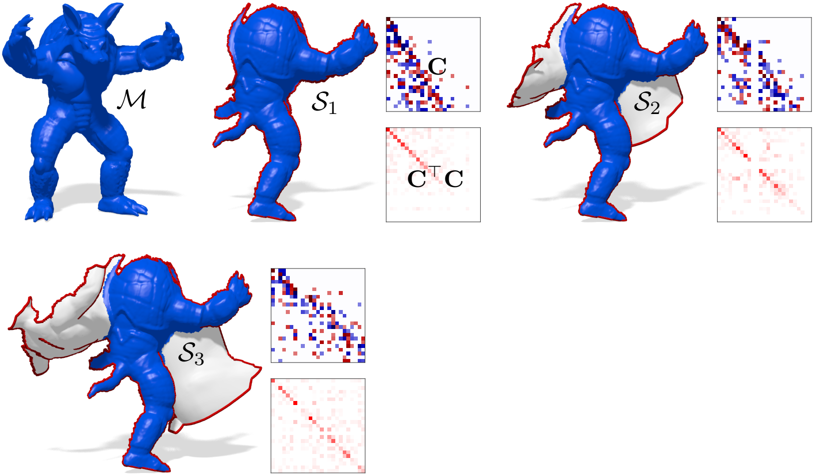

Functional correspondence with clutter

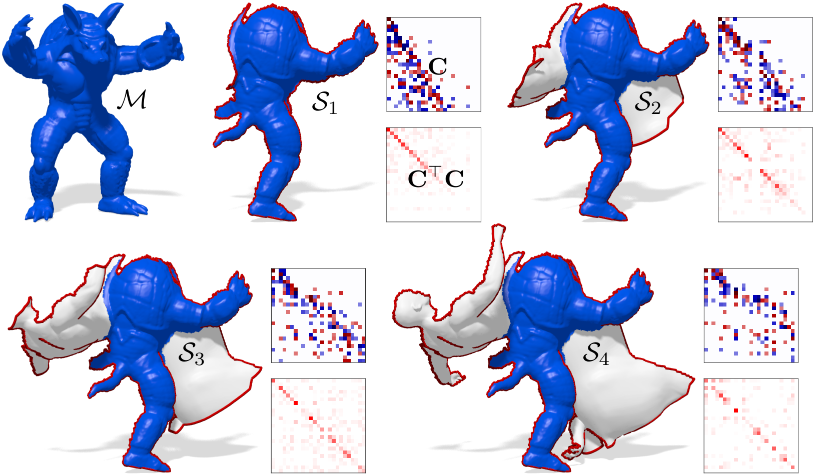

Functional object-in-clutter

$$\mathbf{T}\mathrm{diag}(u)\mathbf{f} = \mathrm{diag}(v)\mathbf{g}$$

$u : \mathcal{M} \to [0, 1]$

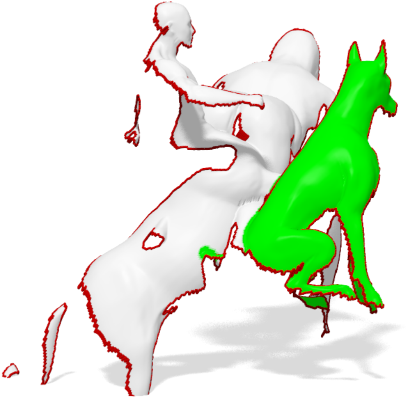

$v : \mathcal{S} \to [0, 1]$

Functional object-in-clutter

$$\mathbf{T}\mathrm{diag}(u)\mathbf{f} = \mathrm{diag}(v)\mathbf{g}$$

$\xrightarrow{\mathbf{T}}$

Functional object-in-clutter

$$\min_{\mathbf{C},\theta,u,v} \|\mathbf{C}\mathbf{\Phi}^\top \mathrm{diag}(u)\mathbf{F} - \mathbf{\Psi}^\top \mathrm{diag}(v)\mathbf{G}\|_{2,1} + \|\mathbf{C}\mathbf{\Phi}^\top u - \mathbf{\Psi}^\top v\|_2^2 + \rho_{\mathrm{corr}}(\mathbf{C},\theta) + \rho_{\mathrm{part}}(u,v)$$

Functional object-in-clutter

$$\min_{\mathbf{C},\theta,u,v} \|\mathbf{C}\mathbf{\Phi}^\top \mathrm{diag}(u)\mathbf{F} - \mathbf{\Psi}^\top \mathrm{diag}(v)\mathbf{G}\|_{2,1} + \|\mathbf{C}\mathbf{\Phi}^\top u - \mathbf{\Psi}^\top v\|_2^2$$
$$+ \rho_{\mathrm{corr}}(\mathbf{C},\theta) + \rho_{\mathrm{part}}(u,v)$$

- **Part regularization**

$$\rho_{\mathrm{part}}(u,v) = \mu_1 \left(\int_{\mathcal{M}} u dx - \int_{\mathcal{S}} v dx \right)^2 - \mu_2 \left(\int_{\mathcal{M}} u dx + \int_{\mathcal{S}} v dx \right)^2$$

$$+ \mu_3 \left(\int_{\mathcal{M}} \|\nabla_{\mathcal{M}} u\| dx + \int_{\mathcal{S}} \|\nabla_{\mathcal{S}} v\| dx \right)$$

Functional object-in-clutter

$$\min_{\mathbf{C},\theta,u,v} \|\mathbf{C}\mathbf{\Phi}^\top \mathrm{diag}(u)\mathbf{F} - \mathbf{\Psi}^\top \mathrm{diag}(v)\mathbf{G}\|_{2,1} + \|\mathbf{C}\mathbf{\Phi}^\top u - \mathbf{\Psi}^\top v\|_2^2$$
$$+ \rho_{\mathrm{corr}}(\mathbf{C},\theta) + \rho_{\mathrm{part}}(u,v)$$

- **Part regularization**

$$\rho_{\mathrm{part}}(u,v) = \underbrace{\mu_1 \left(\int_{\mathcal{M}} u dx - \int_{\mathcal{S}} v dx \right)^2}_{\text{area preservation}} - \underbrace{\mu_2 \left(\int_{\mathcal{M}} u dx + \int_{\mathcal{S}} v dx \right)^2}_{\text{part size}}$$
$$+ \underbrace{\mu_3 \left(\int_{\mathcal{M}} \|\nabla_{\mathcal{M}} u\| dx + \int_{\mathcal{S}} \|\nabla_{\mathcal{S}} v\| dx \right)}_{\text{Mumford−Shah}}$$

Functional object-in-clutter

$$\min_{\mathbf{C},\theta,u,v} \|\mathbf{C}\mathbf{\Phi}^\top \mathrm{diag}(u)\mathbf{F} - \mathbf{\Psi}^\top \mathrm{diag}(v)\mathbf{G}\|_{2,1} + \|\mathbf{C}\mathbf{\Phi}^\top u - \mathbf{\Psi}^\top v\|_2^2$$
$$+ \rho_{\mathrm{corr}}(\mathbf{C},\theta) + \rho_{\mathrm{part}}(u,v)$$

- **Correspondence regularization**

$$\rho_{\mathrm{corr}}(\mathbf{C},\theta) = \mu_4\|\mathbf{C}\circ\mathbf{W}(\theta)\|_{\mathrm{F}}^2 + \mu_5\sum_{i\neq j}(\mathbf{C}^\top\mathbf{C})_{ij}^2 + \mu_6\sum_i |\mathbf{C}^\top\mathbf{C}|_{ii}$$

$\mathbf{W}(\theta) =$ 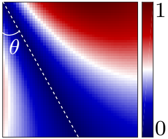

Functional object-in-clutter

$$\min_{\mathbf{C},\theta,u,v} \|\mathbf{C}\mathbf{\Phi}^\top \mathrm{diag}(u)\mathbf{F} - \mathbf{\Psi}^\top \mathrm{diag}(v)\mathbf{G}\|_{2,1} + \|\mathbf{C}\mathbf{\Phi}^\top u - \mathbf{\Psi}^\top v\|_2^2 + \rho_{\mathrm{corr}}(\mathbf{C},\theta) + \rho_{\mathrm{part}}(u,v)$$

- **Correspondence regularization**

$$\rho_{\mathrm{corr}}(\mathbf{C},\theta) = \underbrace{\mu_4 \|\mathbf{C} \circ \mathbf{W}(\theta)\|_{\mathrm{F}}^2}_{\text{slant}} + \underbrace{\mu_5 \sum_{i \neq j} (\mathbf{C}^\top \mathbf{C})_{ij}^2}_{\approx \text{orthogonality}} + \underbrace{\mu_6 \sum_i |\mathbf{C}^\top \mathbf{C}|_{ii}}_{\text{sparsity}}$$

$\mathbf{W}(\theta) =$ 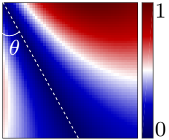

Learning descriptors

$$\min_{\mathbf{C},\theta,u,v} \|\mathbf{C}\mathbf{\Phi}^\top \mathrm{diag}(u)\mathbf{F} - \mathbf{\Psi}^\top \mathrm{diag}(v)\mathbf{G}\|_{2,1} + \cdots$$

For the data term we use dense descriptor fields.

Learning descriptors

$$\min_{\mathbf{C},\theta,u,v} \|\mathbf{C}\mathbf{\Phi}^\top \mathrm{diag}(u)\mathbf{F} - \mathbf{\Psi}^\top \mathrm{diag}(v)\mathbf{G}\|_{2,1} + \cdots$$

For the data term we use dense descriptor fields.

- Existing isometry-invariant descriptors (HKS, WKS) are affected by clutter and boundary effects

Sun et al. 2009; Aubry et al. 2011

Learning descriptors

$$\min_{\mathbf{C},\theta,u,v} \|\mathbf{C}\mathbf{\Phi}^\top \mathrm{diag}(u)\mathbf{F} - \mathbf{\Psi}^\top \mathrm{diag}(v)\mathbf{G}\|_{2,1} + \cdots$$

For the data term we use dense descriptor fields.

- Existing isometry-invariant descriptors (HKS, WKS) are affected by clutter and boundary effects
- Local descriptors (FPFH, SHOT) are not isometry invariant and sensitive to sampling

Sun et al. 2009; Aubry et al. 2011; Rusu et al. 2009; Tombari et al. 2010

Learning descriptors

$$\min_{\mathbf{C},\theta,u,v} \|\mathbf{C}\mathbf{\Phi}^\top \mathrm{diag}(u)\mathbf{F} - \mathbf{\Psi}^\top \mathrm{diag}(v)\mathbf{G}\|_{2,1} + \cdots$$

For the data term we use dense descriptor fields.

- Existing isometry-invariant descriptors (HKS, WKS) are affected by clutter and boundary effects
- Local descriptors (FPFH, SHOT) are not isometry invariant and sensitive to sampling

Our solution:

Perform metric learning upon 544-dim SHOT to derive 32-dim descriptors that are robust to clutter, missing parts, and near-isometries

Sun et al. 2009; Aubry et al. 2011; Rusu et al. 2009; Tombari et al. 2010; Hadsell et al. 2006; Masci, Boscaini, Bronstein, Vandergheynst 2015

Performance of learned descriptors

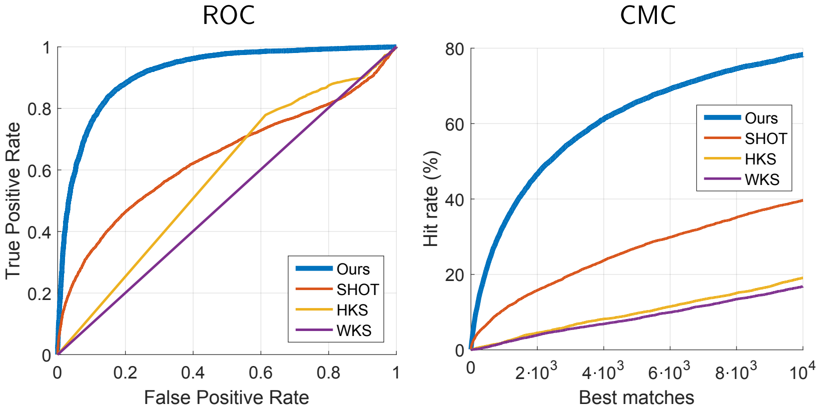

Tombari et al. 2010 (SHOT); Sun et al. 2009 (HKS); Aubry et al. 2011 (WKS)

Comparisons

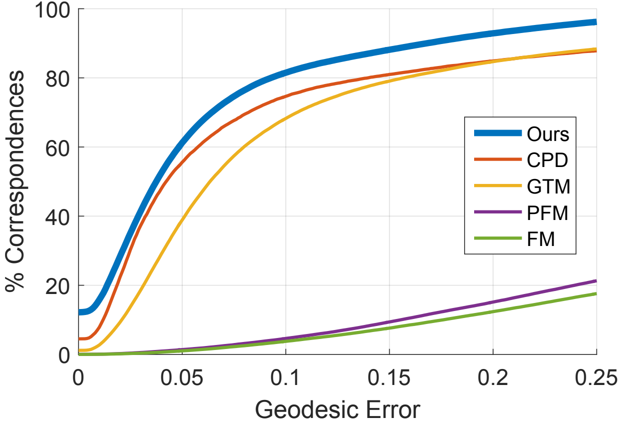

Methods: Myronenko et al. 2010 (CPD); Rodolà, Albarelli, Bergamasco, Torsello 2013 (GTM); Rodolà, Cosmo, Bronstein, Torsello, Cremers 2016 (PFM); Ovsjanikov et al. 2012 (FM)

Examples with clutter

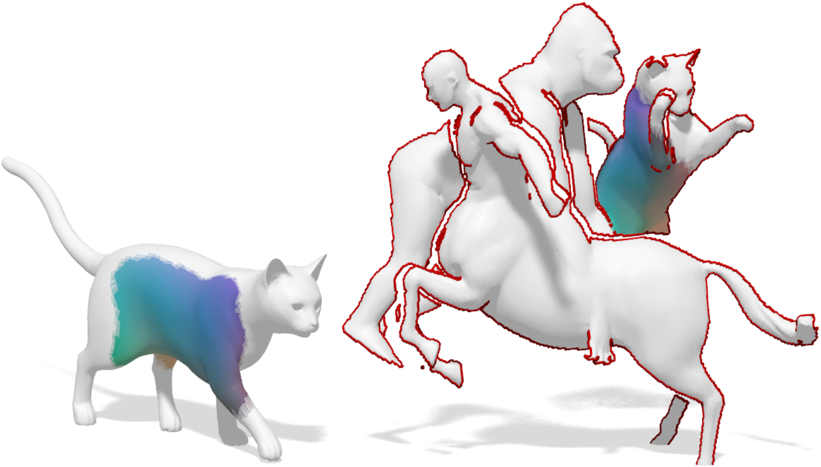

Examples with clutter

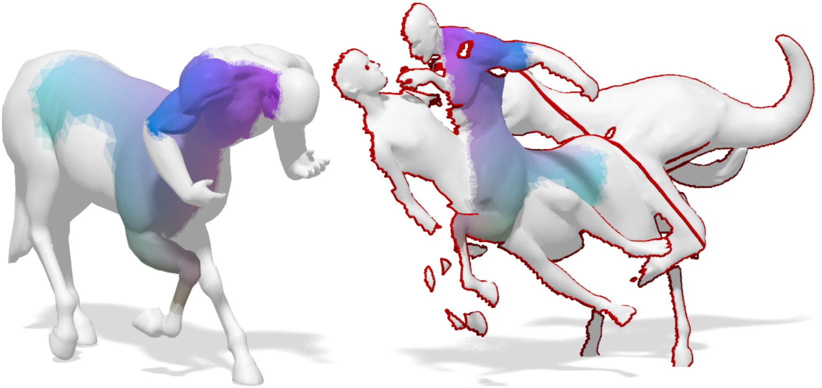

Examples with clutter

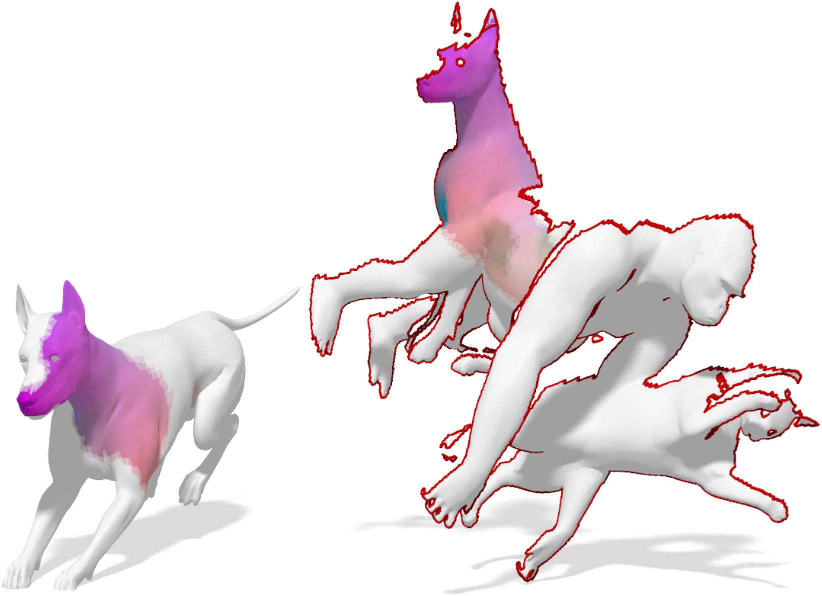

Failure case

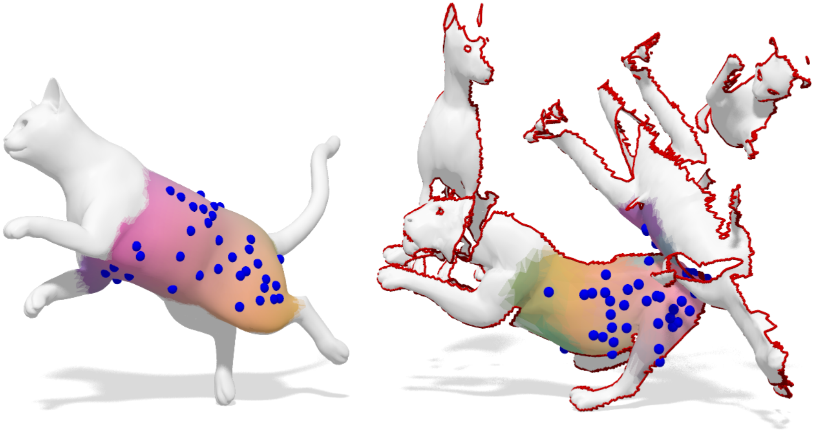

Conclusions

- Deformable object-in-clutter has been much less investigated than its rigid counterpart, and there is a lack of data and benchmarks.

Conclusions

- Deformable object-in-clutter has been much less investigated than its rigid counterpart, and there is a lack of data and benchmarks.
- We presented a spectral approach that works remarkably well despite the realistic setting.

Conclusions

- Deformable object-in-clutter has been much less investigated than its rigid counterpart, and there is a lack of data and benchmarks.
- We presented a spectral approach that works remarkably well despite the realistic setting.
- Existing descriptors do not behave well in this setting; we need new descriptors!

Conclusions

- Deformable object-in-clutter has been much less investigated than its rigid counterpart, and there is a lack of data and benchmarks.
- We presented a spectral approach that works remarkably well despite the realistic setting.
- Existing descriptors do not behave well in this setting; we need new descriptors!

Thank you!

Laplacian eigenvectors with clutter

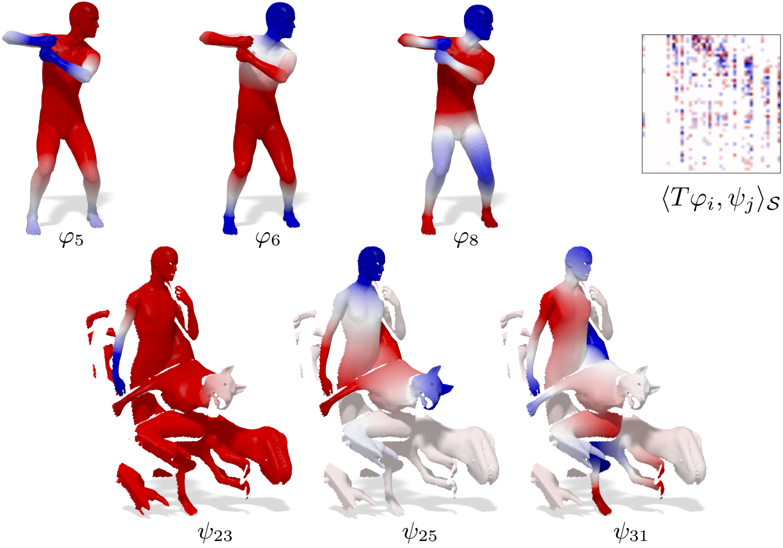

Examples (no clutter)

Perturbation analysis: intuition

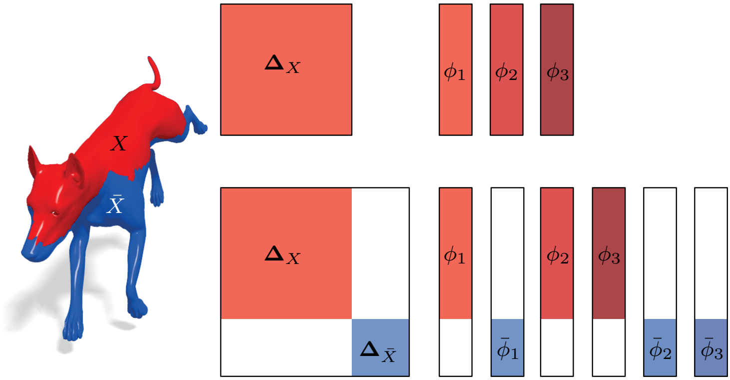

- Ignoring boundary interaction: disjoint parts (block-diagonal matrix)
- Eigenvectors = Mixture of eigenvectors of the parts

Perturbation analysis: eigenvalues

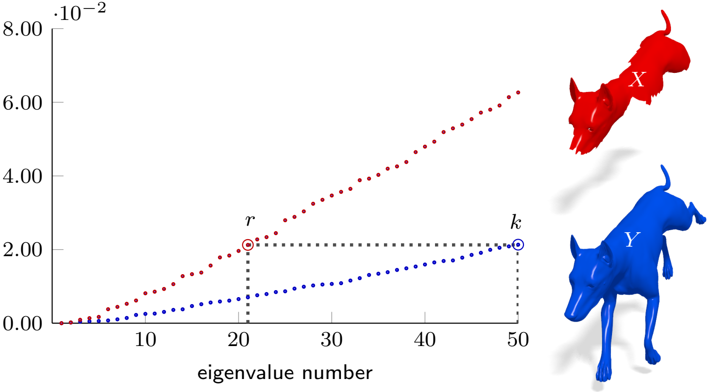

- Slope $\frac{r}{k} \approx \frac{\mathrm{area}(X)}{\mathrm{area}(Y)}$ (depends on the area of the cut)
- Consistent with Weyl's law for 2-manifolds

Rodolà, Cosmo, Bronstein, Torsello, Cremers 2016

Perturbation analysis: eigenvalues

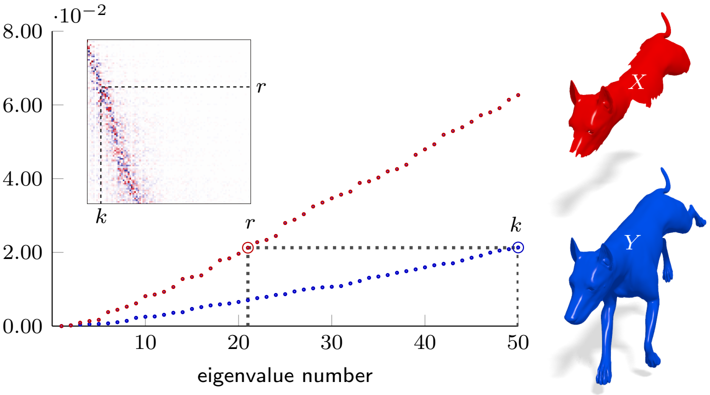

- Slope $\frac{r}{k} \approx \frac{\text{area}(X)}{\text{area}(Y)}$ (depends on the area of the cut)
- Consistent with Weyl's law for 2-manifolds

Rodolà, Cosmo, Bronstein, Torsello, Cremers 2016

Perturbation analysis: intuition

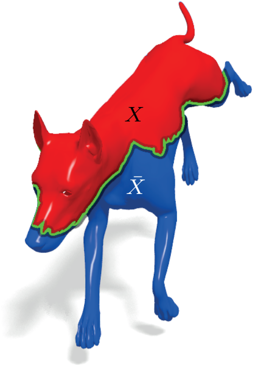

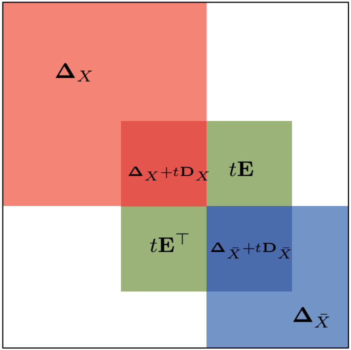

Rodolà, Cosmo, Bronstein, Torsello, Cremers 2016

Initialization

We solve a small minimum-distortion correspondence problem with sparsity constraints.

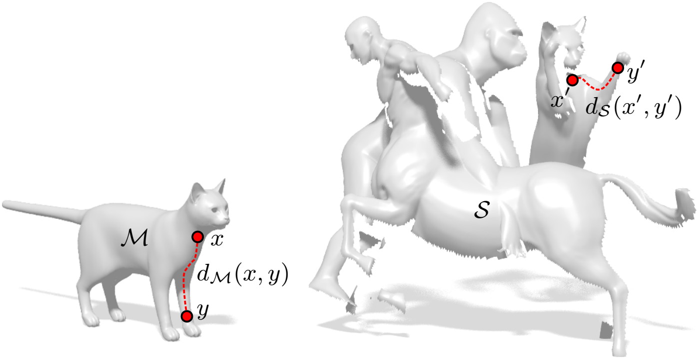

Initialization

We solve a small minimum-distortion correspondence problem with sparsity constraints.

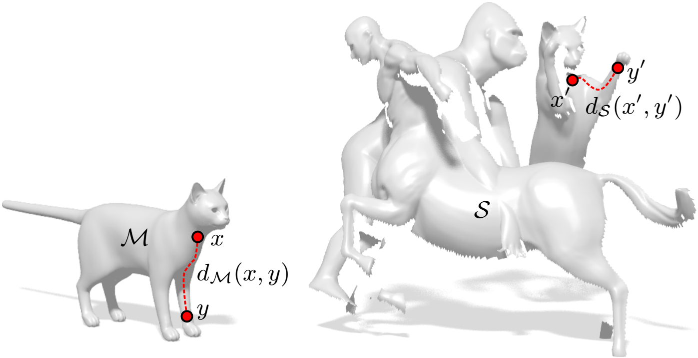

Formally, we find local solutions to a L_1-relaxed variant of the quadratic assignment problem (QAP).

Rodolà, Bronstein, Albarelli, Bergamasco, Torsello 2012; Rodolà, Torsello, Harada, Kuniyoshi, Cremers 2013

Metric learning

Learn an embedding function $F(x)$ parametrized by Θ, by minimizing the siamese loss:

$$L_{\mathsf{s}}(\Theta) = \sum_{x,x^+ \in S} \gamma \|F_\Theta(x) - F_\Theta(x^+)\|_2^2 + \sum_{x,x^- \in D} (1-\gamma)(m_s - \|F_\Theta(x) - F_\Theta(x^-)\|_2)_+^2$$

where (x, x^+), (x, x^-) are knowingly similar and dissimilar point pairs.

Cosmo, Rodolà, Masci, Torsello, Bronstein 2016; Bromley et al. 1994; Hadsell et al. 2006

Metric learning

Learn an embedding function $F(x)$ parametrized by Θ, by minimizing the siamese loss:

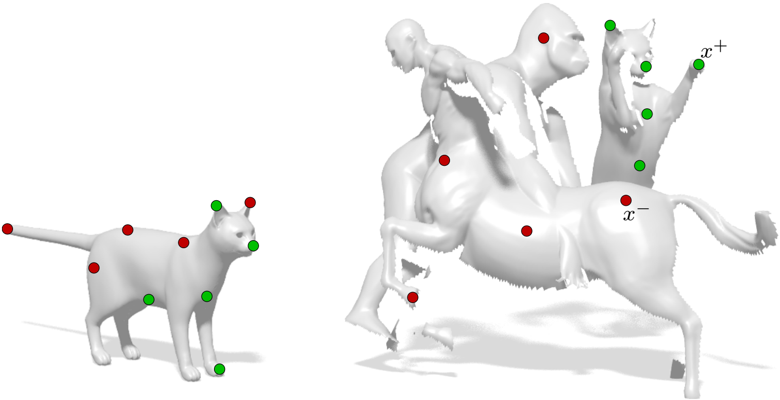

Cosmo, Rodolà, Masci, Torsello, Bronstein 2016; Bromley et al. 1994; Hadsell et al. 2006

Metric learning

Learn an embedding function $F(x)$ parametrized by Θ, by minimizing the siamese loss:

$$L_{\mathsf{s}}(\Theta) = \sum_{x, x^+ \in S} \gamma \|F_\Theta(x) - F_\Theta(x^+)\|_2^2 + \sum_{x, x^- \in D} (1-\gamma)(m_s - \|F_\Theta(x) - F_\Theta(x^-)\|_2)_+^2$$

where (x, x^+), (x, x^-) are knowingly similar and dissimilar point pairs.

Regularize with a global distribution penalty:

$$L_{\mathsf{g}}(\Theta) = \sigma_\Theta^+ + \sigma_\Theta^- + (m_g + \mu_\Theta^+ - \mu_\Theta^-)_+$$

Cosmo, Rodolà, Masci, Torsello, Bronstein 2016; Bromley et al. 1994; Hadsell et al. 2006; Kumar et al. 2015

Training data

- All points are used for training

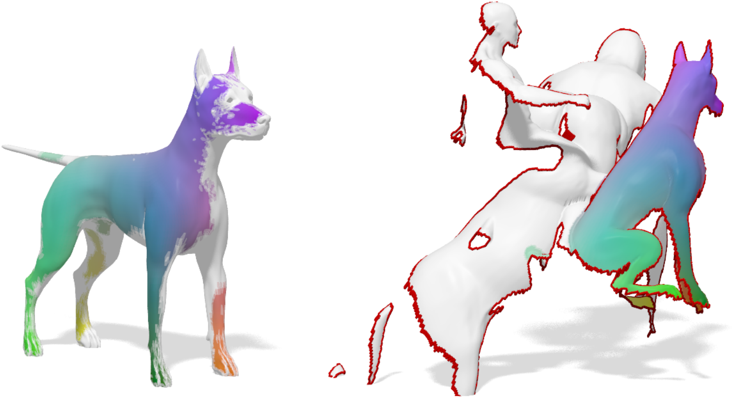

Cosmo, Rodolà, Masci, Torsello, Bronstein 2016

Training data

- All points are used for training
- A local extrinsic descriptor is attached to each point

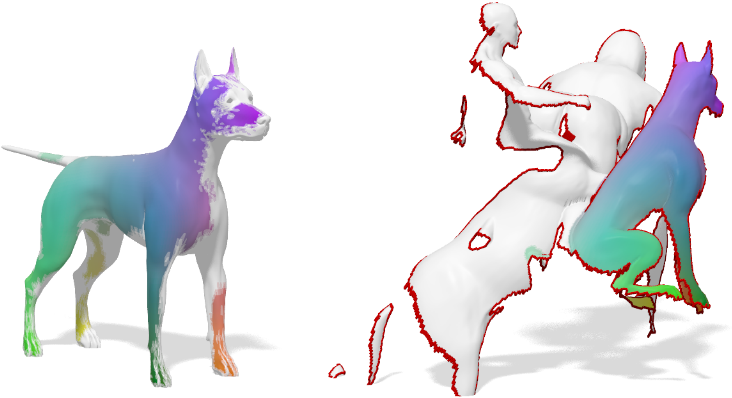

Cosmo, Rodolà, Masci, Torsello, Bronstein 2016; Tombari et al. 2010

ResNet

Function F_Θ is modeled as a deep residual network

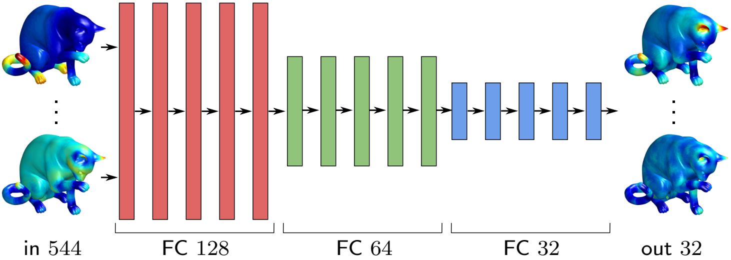

The geometric information lies in the descriptor fed as input

Cosmo, Rodolà, Masci, Torsello, Bronstein 2016; He et al. 2015

Perturbation analysis: details

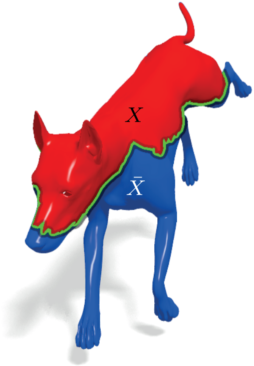

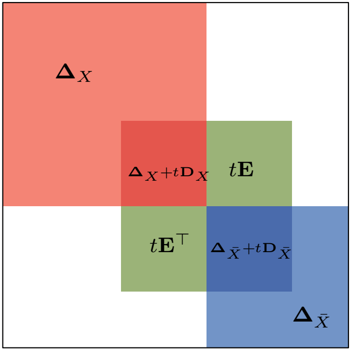

Perturbation analysis: details

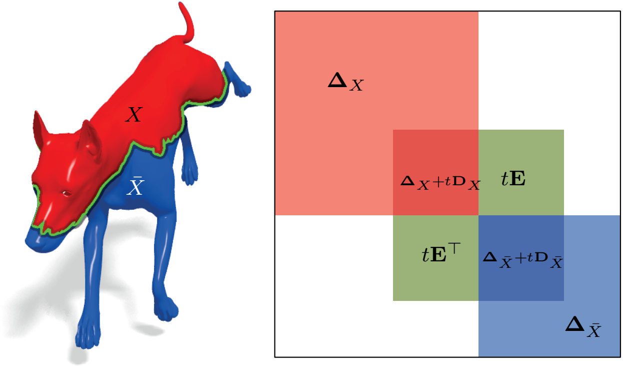

"How would the Laplacian eigenvalues and eigenvectors of the red part change if we attached a blue part to it?"

Perturbation analysis: details

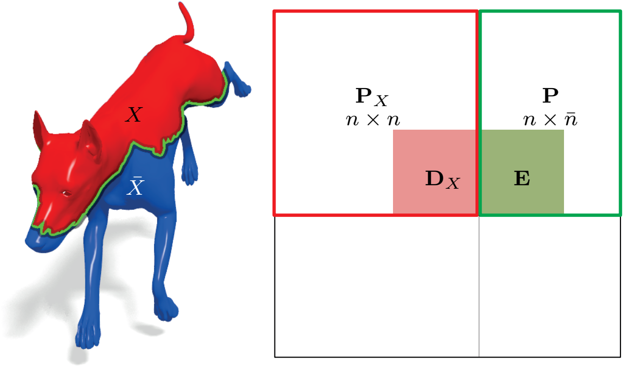

"How would the Laplacian eigenvalues and eigenvectors of the red part change if we attached a blue part to it?"

Perturbation analysis: details

Denote $\mathbf{\Delta}_X + t\mathbf{P}_X = \mathbf{\Phi}(t)\mathbf{\Lambda}(t)\mathbf{\Phi}(t)^\top$, $\mathbf{\Delta}_{\bar{X}} = \bar{\mathbf{\Phi}}\bar{\mathbf{\Lambda}}\bar{\mathbf{\Phi}}^\top$, $\mathbf{\Phi} = \mathbf{\Phi}(0)$, and $\mathbf{\Lambda} = \mathbf{\Lambda}(0)$.

Theorem 1 (eigenvalues) The derivative of the non-trivial eigenvalues is given by

$$\frac{d}{dt}\lambda_i = \boldsymbol{\phi}_i^\top \mathbf{P}_X \boldsymbol{\phi}_i \qquad \mathbf{P}_X = \begin{pmatrix} \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{D}_X \end{pmatrix}$$

Perturbation analysis: details

Denote $\mathbf{\Delta}_X + t\mathbf{P}_X = \mathbf{\Phi}(t)\mathbf{\Lambda}(t)\mathbf{\Phi}(t)^\top$, $\mathbf{\Delta}_{\bar{X}} = \bar{\mathbf{\Phi}}\bar{\mathbf{\Lambda}}\bar{\mathbf{\Phi}}^\top$, $\mathbf{\Phi} = \mathbf{\Phi}(0)$, and $\mathbf{\Lambda} = \mathbf{\Lambda}(0)$.

Theorem 1 (eigenvalues) The derivative of the non-trivial eigenvalues is given by

$$\frac{d}{dt}\lambda_i = \boldsymbol{\phi}_i^\top \mathbf{P}_X \boldsymbol{\phi}_i \qquad \mathbf{P}_X = \begin{pmatrix} \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{D}_X \end{pmatrix}$$

Theorem 2 (eigenvectors) Assuming $\lambda_i \neq \lambda_j$ for $i \neq j$ and $\lambda_i \neq \bar{\lambda}_j$ for all i, j, the derivative of the non-trivial eigenvectors is given by

$$\frac{d}{dt}\boldsymbol{\phi}_i = \sum_{\substack{j=1 \\ j \neq i}}^{n} \frac{\boldsymbol{\phi}_i^\top \mathbf{P}_X \boldsymbol{\phi}_j}{\lambda_i - \lambda_j}\boldsymbol{\phi}_j + \sum_{j=1}^{\bar{n}} \frac{\boldsymbol{\phi}_i^\top \mathbf{P}\,\bar{\boldsymbol{\phi}}_j}{\lambda_i - \bar{\lambda}_j}\bar{\boldsymbol{\phi}}_j \qquad \mathbf{P} = \begin{pmatrix} \mathbf{0} & \mathbf{0} \\ \mathbf{E} & \mathbf{0} \end{pmatrix}$$

Perturbation analysis: boundary interaction strength

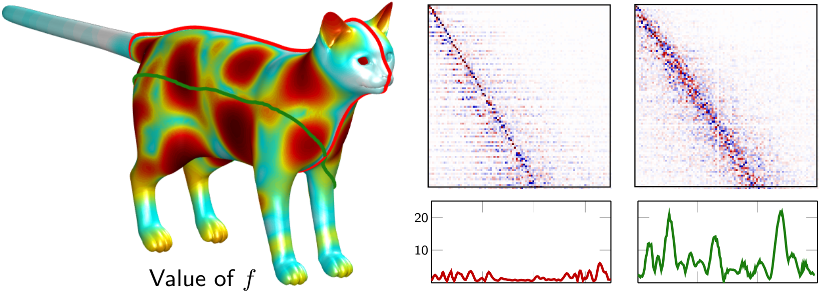

- Eigenvector perturbation depends on length and position of the boundary
- Perturbation strength $\leq c \int_{\partial X} f(x)dx$, where

$$f(x) = \sum_{\substack{i,j=1 \\ j \neq i}}^{n} \left(\frac{\phi_i(x)\phi_j(x)}{\lambda_i - \lambda_j} \right)^2$$

(bi-)Laplacian perturbation: typical picture

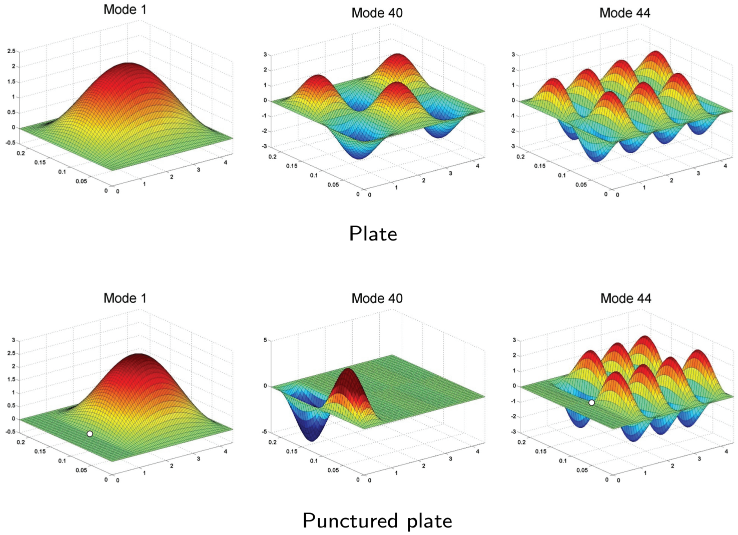

Figure: Filoche, Mayboroda 2009