Geometry and learning in 3D shape correspondence

Alex Bronstein

Department of Computer Science, Technion

IPAM · UCLA, 2019


T. Remez


O. Litany


E. Rodolà

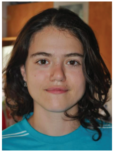
O. Halimi

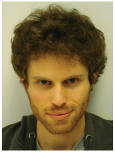
A. Boyarski


R. Slossberg


R. Kimmel

TEL AVIV UNIVERSITY אוניברסיטת תל אביב

SAPIENZA UNIVERSITÀ DI ROMA


M. Vestner


Z. Lähner


D. Cremers


M. Bronstein

USI

Supported by

erc intel

Deformable 3D correspondence

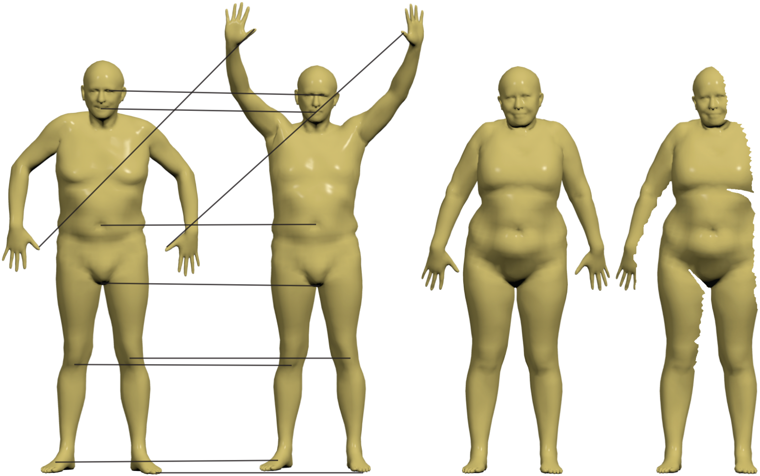

Pointwise correspondence

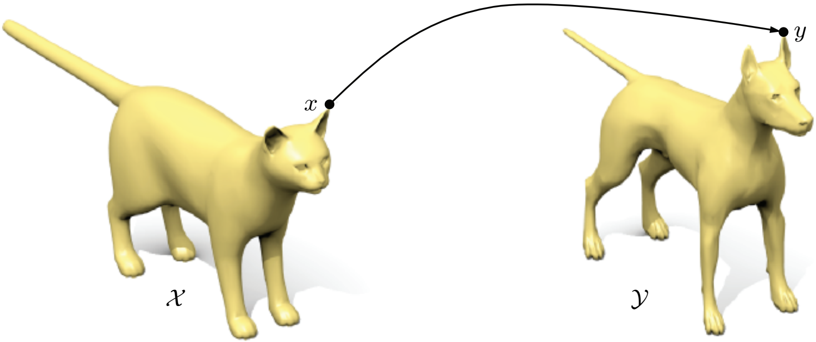

Point-wise map $t\colon \mathcal{X} \to \mathcal{Y}$

Functional correspondence

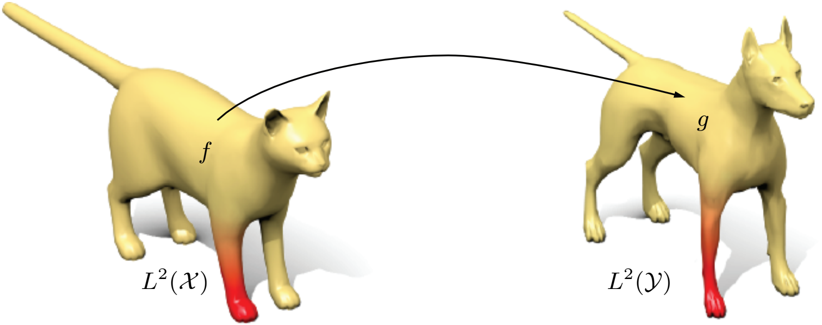

Functional map $T\colon L^2(\mathcal{X}) \to L^2(\mathcal{Y})$

Ovsjanikov et al. 2012

Functional correspondence in spectral domain

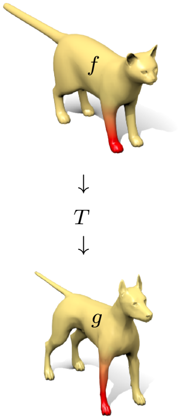

Ovsjanikov et al. 2012

Functional correspondence in spectral domain

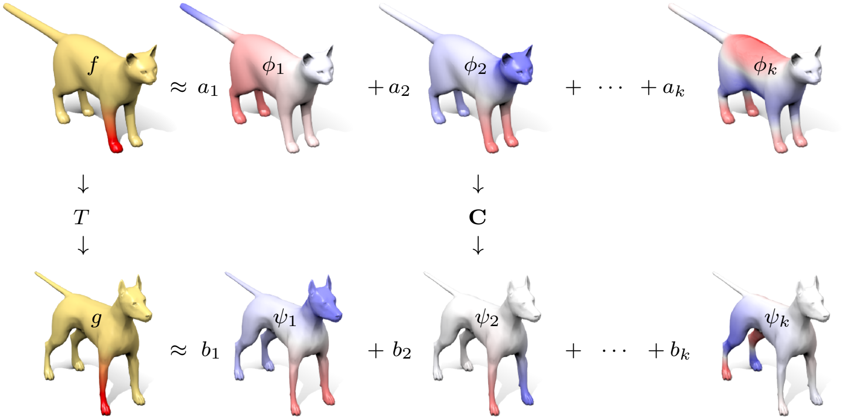

Ovsjanikov et al. 2012

Functional correspondence in spectral domain

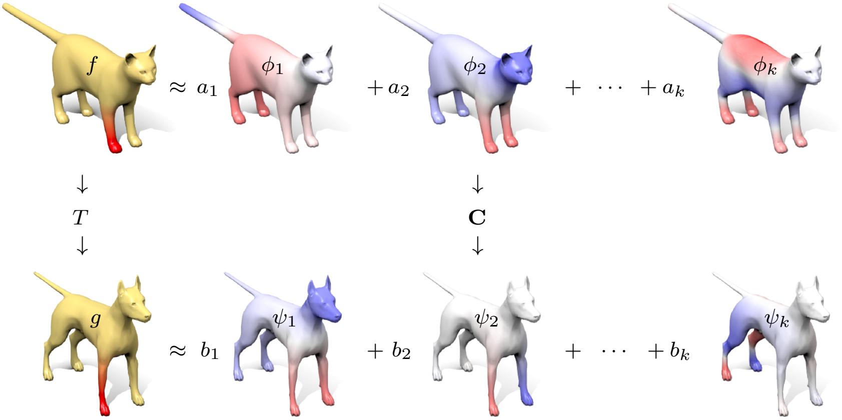

Functional correspondence boils down to a linear equation w.r.t. $\mathbf{C}$

$$g = Tf \qquad \Longleftrightarrow \qquad \mathbf{b} = \mathbf{C}\mathbf{a}$$

Ovsjanikov et al. 2012

Functional correspondence in spectral domain

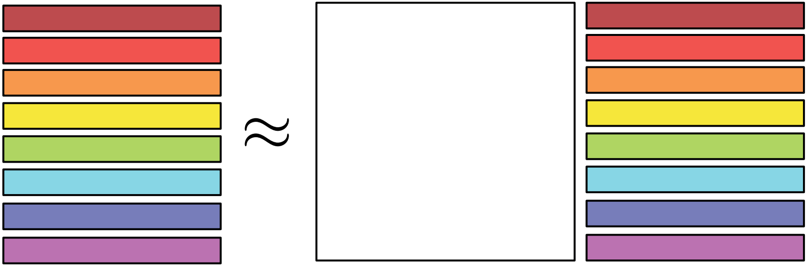

$\mathbf{B}_{k\times q} = \left(\langle \psi_i, g_j \rangle_{L^2(\mathcal{Y})}\right)$ $\quad\approx\quad$ $\mathbf{C}_{k\times k}$ $\quad$ $\mathbf{A}_{k\times q} = \left(\langle \phi_i, f_j \rangle_{L^2(\mathcal{X})}\right)$

where $\mathbf{A}$, $\mathbf{B}$ are Fourier coefficients of corresponding 'probe' functions

$$g_i \approx T f_i \qquad i = 1, \dots, q \geq k$$

Ovsjanikov et al. 2012

Point-wise descriptors

- Assign a vector-valued descriptor $\mathbf{h} : \mathcal{X} \to \mathbb{R}^k$ to every point

Point-wise descriptors

- Assign a vector-valued descriptor $\mathbf{h} : \mathcal{X} \to \mathbb{R}^k$ to every point
- Each descriptor dimension $h_i : \mathcal{X} \to \mathbb{R}$ is a probe function

Siamese descriptor learning

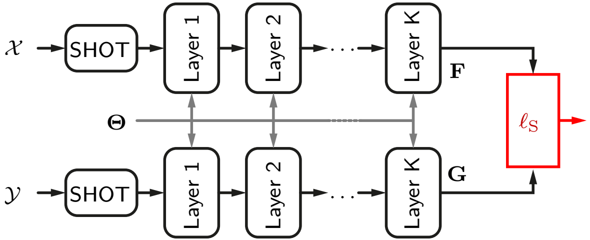

- Two instances of the same network with shared parameters

Litany, Remez, Rodolà, Bronstein^2 2017

Siamese descriptor learning

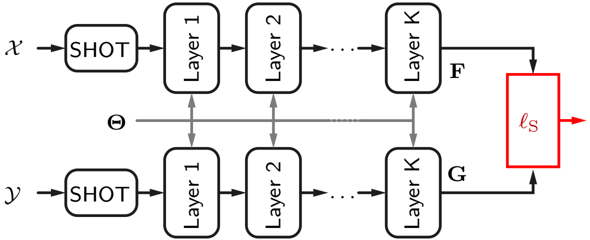

- Two instances of the same network with shared parameters
- Siamese loss on pairs of similar and dissimilar points $S, D \subset \mathcal{X} \times \mathcal{Y}$

$$\ell_{\text{siam}} = \sum_{x, x^+ \in S} \gamma \|\mathbf{F}(x) - \mathbf{F}(x^+)\|_2^2 + \sum_{x, x^- \in D} (1 - \gamma) \left[\mu - \|\mathbf{F}(x) - \mathbf{F}(x^-)\|_2^2\right]_+$$

Litany, Remez, Rodolà, Bronstein[2] 2017

Siamese descriptor learning

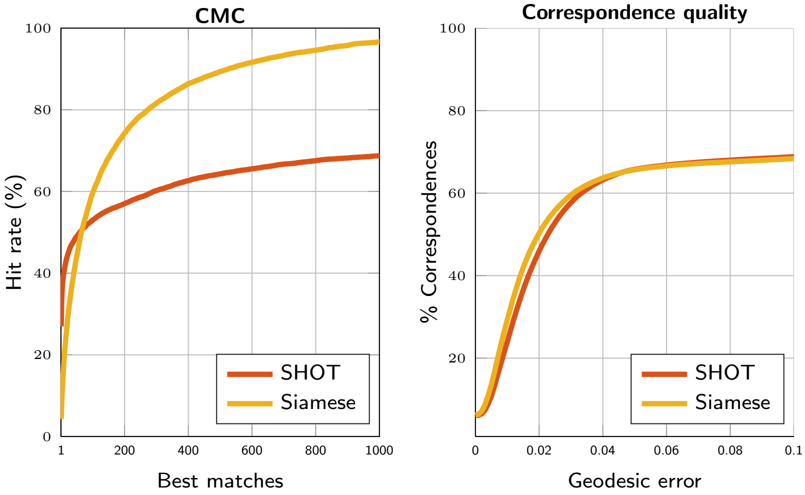

data: SCAPE (Anguelov et al. 2005); FMNet trained on FAUST

FMNet

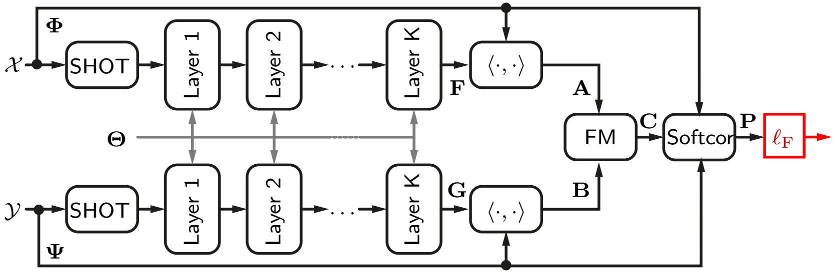

- **Functional map layer:** $\mathbf{C} = \arg\min \|\mathbf{CA} - \mathbf{B}\|_F^2$

Litany, Remez, Rodolà, B^2 2017

FMNet

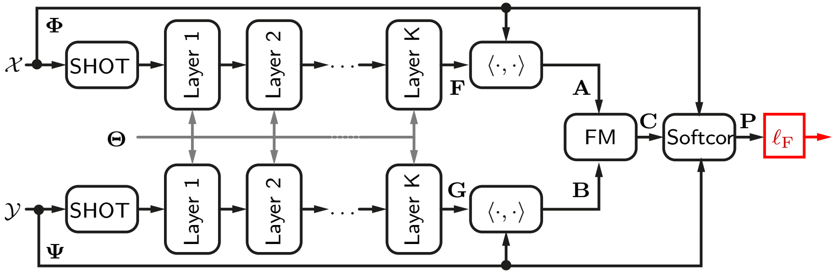

- **Functional map layer:** $\mathbf{C} = \arg\min \|\mathbf{CA} - \mathbf{B}\|_F^2$
- **Soft correspondence layer:** $\mathbf{P} = |\mathbf{\Psi C \Phi}^\top|_{\|\cdot\|}$

Litany, Remez, Rodolà, B[2] 2017

Soft correspondence

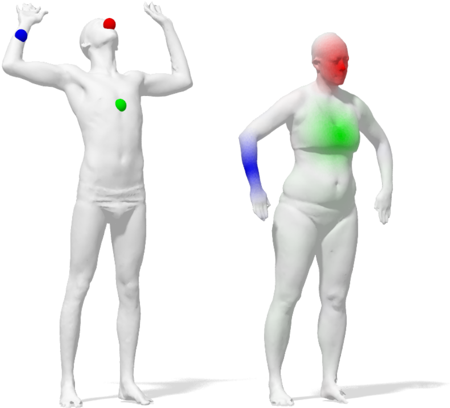

- FMNet outputs soft correspondence $\mathbf{P} = |\mathbf{\Psi} \mathbf{C} \mathbf{\Phi}^\top|_{\|\cdot\|}$

Litany, Remez, Rodolà, B^2 2017

Soft correspondence

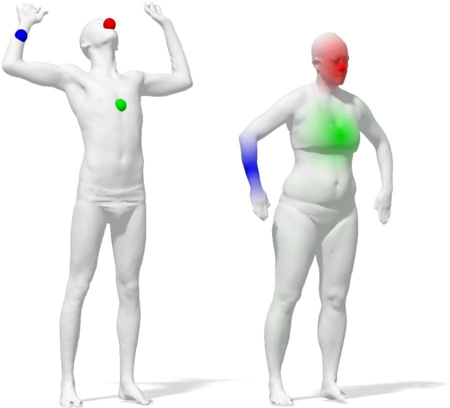

- FMNet outputs soft correspondence $\mathbf{P} = |\mathbf{\Psi C \Phi}^\top|_{\|\cdot\|}$
- $P(x, y)$ can be interpreted as the probability of point $x \in \mathcal{X}$ mapping to point $y \in \mathcal{Y}$

Litany, Remez, Rodolà, B^2 2017

FMNet

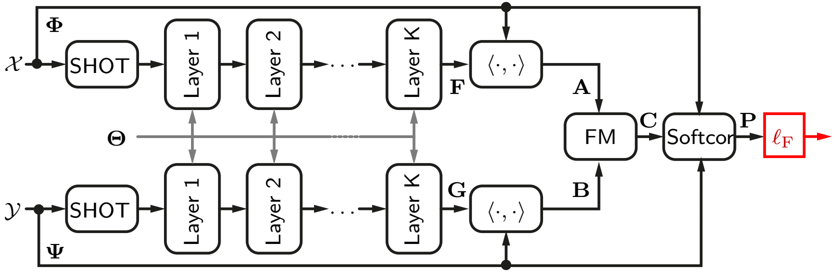

- **Functional map layer:** $\mathbf{C} = \arg\min \|\mathbf{CA} - \mathbf{B}\|_F^2$
- **Soft correspondence layer:** $\mathbf{P} = |\boldsymbol{\Psi}\mathbf{C}\boldsymbol{\Phi}^\top|_{\|\cdot\|}$
- Functional map loss:

$$\ell_{\mathsf{sup}} = \sum_{(x,y)\in(\mathcal{X},\mathcal{Y})} P(x,y)\, d^2_{\mathcal{Y}}(y, \pi^*(x))$$

Litany, Remez, Rodolà, B^2 2017

Learning to find correspondence

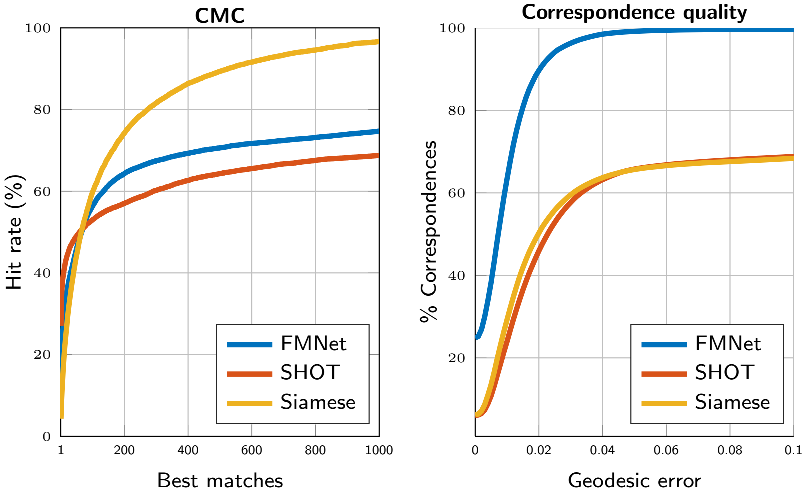

data: SCAPE (Anguelov et al. 2005); FMNet trained on FAUST

Training regime

- Supervised loss

$$\ell_{\mathsf{sup}}(\mathcal{X}, \mathcal{Y}) \;=\; \sum_{(x,y)\in(\mathcal{X},\mathcal{Y})} P(x, y)\, d^2_{\mathcal{Y}}(y, \pi^*(x))$$

Vestner, Litman, Rodola, B, Cremers 2017; Vestner *et al.* 2017

Training regime

- Supervised loss

$$\ell_{\mathsf{sup}}(\mathcal{X},\mathcal{Y}) = \sum_{(x,y)\in(\mathcal{X},\mathcal{Y})} P(x,y)\, d^2_{\mathcal{Y}}(y,\pi^*(x))$$

- Requires dense groundtruth correspondences $\pi^* : \mathcal{X} \to \mathcal{Y}$ on many shapes

Vestner, Litman, Rodola, B, Cremers 2017; Vestner *et al.* 2017

Training regime

- Supervised loss

$$\ell_{\mathsf{sup}}(\mathcal{X},\mathcal{Y}) \;=\; \sum_{(x,y)\in(\mathcal{X},\mathcal{Y})} P(x,y)\, d^2_{\mathcal{Y}}(y,\pi^*(x))$$

- Requires dense groundtruth correspondences $\pi^* : \mathcal{X} \to \mathcal{Y}$ on many shapes
- Can we use no groundtruth at all?

Vestner, Litman, Rodola, B, Cremers 2017; Vestner *et al.* 2017

Training regime

- Supervised loss

$$\ell_{\mathsf{sup}}(\mathcal{X},\mathcal{Y}) \;=\; \sum_{(x,y)\in(\mathcal{X},\mathcal{Y})} P(x,y)\, d^2_{\mathcal{Y}}(y,\pi^*(x))$$

- Requires dense groundtruth correspondences $\pi^* : \mathcal{X} \to \mathcal{Y}$ on many shapes
- Can we use no groundtruth at all?
- Can we train on the same pair of shapes being matched with no additional data?

Vestner, Litman, Rodola, B, Cremers 2017; Vestner *et al.* 2017

Training regime

- Supervised loss: mismatch with groundtruth correspondence

$$\ell_{\mathsf{sup}}(\mathcal{X},\mathcal{Y}) \;=\; \sum_{(x,y)\in(\mathcal{X},\mathcal{Y})} P(x,y)\, d^2_{\mathcal{Y}}(y,\pi^*(x))$$

Halimi *et al.* 2019

Training regime

- Supervised loss: mismatch with groundtruth correspondence

$$\ell_{\mathsf{sup}}(\mathcal{X},\mathcal{Y}) = \sum_{(x,y)\in(\mathcal{X},\mathcal{Y})} P(x,y)\, d^2_{\mathcal{Y}}(y,\pi^*(x))$$

- Unsupervised loss: geodesic distance distortion

$$\ell_{\mathsf{uns}}(\mathcal{X},\mathcal{Y}) = \sum_{x,x'\in\mathcal{X}} \left(d_{\mathcal{X}}(x,x') - \sum_{y,y'\in\mathcal{Y}} P(x,y)P(x',y')\, d_{\mathcal{Y}}(y,y') \right)^2$$

Halimi *et al.* 2019

Training regime

- Supervised loss: mismatch with groundtruth correspondence

$$\ell_{\mathsf{sup}}(\mathcal{X},\mathcal{Y}) = \sum_{(x,y)\in(\mathcal{X},\mathcal{Y})} P(x,y)\, d^2_{\mathcal{Y}}(y,\pi^*(x))$$

- Unsupervised loss: geodesic distance distortion

$$\begin{aligned}\ell_{\mathsf{uns}}(\mathcal{X},\mathcal{Y}) &= \sum_{x,x'\in\mathcal{X}} \left(d_{\mathcal{X}}(x,x') - \sum_{y,y'\in\mathcal{Y}} P(x,y)P(x',y')\, d_{\mathcal{Y}}(y,y') \right)^2 \\ &\approx \left\| \mathbf{D}_{\mathcal{X}} - \mathbb{E}_P \left\{ \mathbf{D}_{\mathcal{Y}} \circ (P \times P) \right\} \right\|_{\mathsf{F}}^2\end{aligned}$$

Halimi *et al.* 2019

Training regime

- Supervised loss: mismatch with groundtruth correspondence

$$\ell_{\mathsf{sup}}(\mathcal{X},\mathcal{Y}) = \sum_{(x,y)\in(\mathcal{X},\mathcal{Y})} P(x,y)\, d^2_{\mathcal{Y}}(y,\pi^*(x))$$

- Unsupervised loss: geodesic distance distortion

$$\begin{aligned}\ell_{\mathsf{uns}}(\mathcal{X},\mathcal{Y}) &= \sum_{x,x'\in\mathcal{X}} \left(d_{\mathcal{X}}(x,x') - \sum_{y,y'\in\mathcal{Y}} P(x,y)P(x',y')\, d_{\mathcal{Y}}(y,y') \right)^2 \\ &\approx \left\| \mathbf{D}_{\mathcal{X}} - \mathbb{E}_P \left\{ \mathbf{D}_{\mathcal{Y}} \circ (P \times P) \right\} \right\|_{\mathsf{F}}^2\end{aligned}$$

- Both trained on some representative class of shapes

Halimi *et al.* 2019

Training regime

- Supervised loss: mismatch with groundtruth correspondence

$$\ell_{\mathsf{sup}}(\mathcal{X},\mathcal{Y}) = \sum_{(x,y)\in(\mathcal{X},\mathcal{Y})} P(x,y)\, d^2_{\mathcal{Y}}(y,\pi^*(x))$$

- Unsupervised loss: geodesic distance distortion

$$\begin{aligned}\ell_{\mathsf{uns}}(\mathcal{X},\mathcal{Y}) &= \sum_{x,x'\in\mathcal{X}} \left(d_{\mathcal{X}}(x,x') - \sum_{y,y'\in\mathcal{Y}} P(x,y)P(x',y')\, d_{\mathcal{Y}}(y,y') \right)^2 \\ &\approx \left\| \mathbf{D}_{\mathcal{X}} - \mathbb{E}_P \left\{ \mathbf{D}_{\mathcal{Y}} \circ (P \times P) \right\} \right\|^2_{\mathsf{F}}\end{aligned}$$

Halimi *et al.* 2019

Training regime

- Supervised loss: mismatch with groundtruth correspondence

$$\ell_{\mathsf{sup}}(\mathcal{X},\mathcal{Y}) = \sum_{(x,y)\in(\mathcal{X},\mathcal{Y})} P(x,y)\, d^2_{\mathcal{Y}}(y,\pi^*(x))$$

- Unsupervised loss: geodesic distance distortion

$$\ell_{\mathsf{uns}}(\mathcal{X},\mathcal{Y}) = \sum_{x,x'\in\mathcal{X}} \left(d_{\mathcal{X}}(x,x') - \sum_{y,y'\in\mathcal{Y}} P(x,y)P(x',y')\, d_{\mathcal{Y}}(y,y') \right)^2$$

$$\approx \|\mathbf{D}_{\mathcal{X}} - \mathbb{E}_P\{\mathbf{D}_{\mathcal{Y}} \circ (P\times P)\}\|^2_{\mathsf{F}}$$

- Self-supervised: unsupervised training on the same pair of shapes for which correspondence is being inferred

Halimi *et al.* 2019

Training regime

- Supervised loss: mismatch with groundtruth correspondence

$$\ell_{\mathsf{sup}}(\mathcal{X},\mathcal{Y}) = \sum_{(x,y)\in(\mathcal{X},\mathcal{Y})} P(x,y)\, d^2_{\mathcal{Y}}(y,\pi^*(x))$$

- Unsupervised loss: geodesic distance distortion

$$\begin{aligned}\ell_{\mathsf{uns}}(\mathcal{X},\mathcal{Y}) &= \sum_{x,x'\in\mathcal{X}} \left(d_{\mathcal{X}}(x,x') - \sum_{y,y'\in\mathcal{Y}} P(x,y)P(x',y')\, d_{\mathcal{Y}}(y,y') \right)^2 \\ &\approx \left\| \mathbf{D}_{\mathcal{X}} - \mathbb{E}_P \left\{ \mathbf{D}_{\mathcal{Y}} \circ (P \times P) \right\} \right\|^2_{\mathsf{F}}\end{aligned}$$

- Self-supervised: unsupervised training on the same pair of shapes for which correspondence is being inferred
- Semi-supervised: combination of $\ell_{\mathsf{uns}} + \ell_{\mathsf{sup}}$ on a small set of shapes or on sparse groundtruth correspondences

Halimi *et al.* 2019

Unsupervised vs. supervised loss

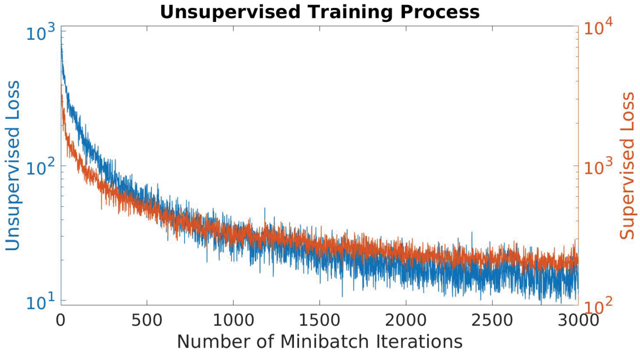

Unsupervised vs. supervised

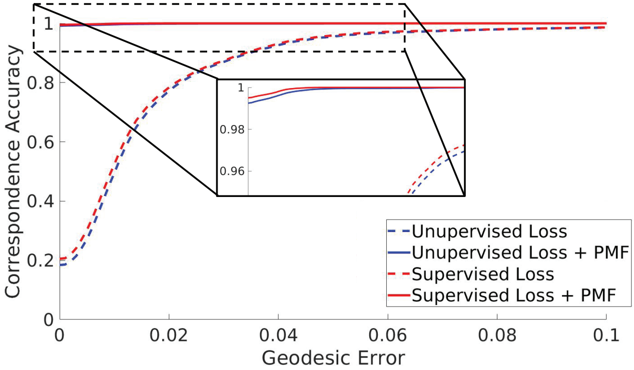

Unsupervised training

Train & Test: synthetic FAUST

Generalization

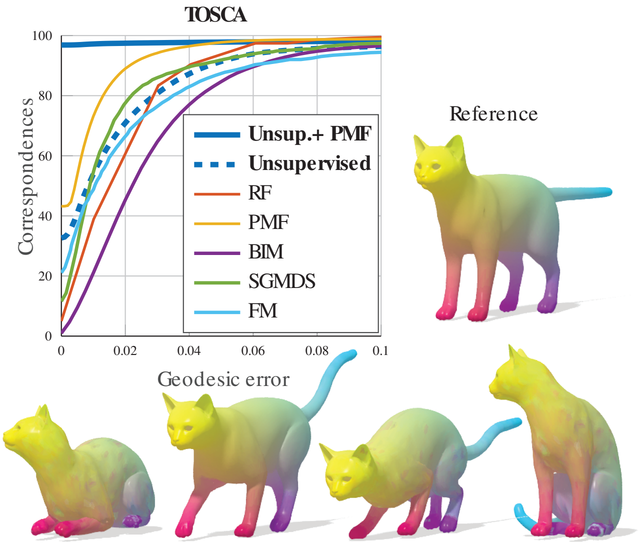

Train: 80 synthetic FAUST; Test: TOSCA animals

Generalization

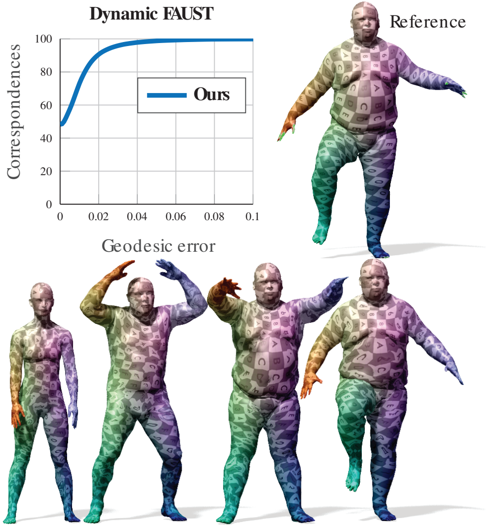

Train: 80 shapes from synthetic FAUST; Test: DFAUST

Partial correspondence

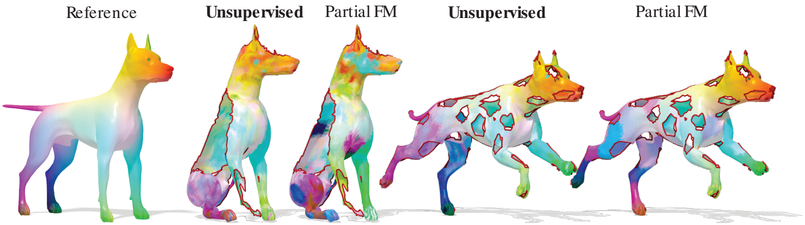

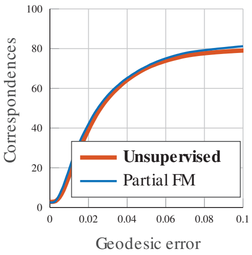

Self-supervised correspondence

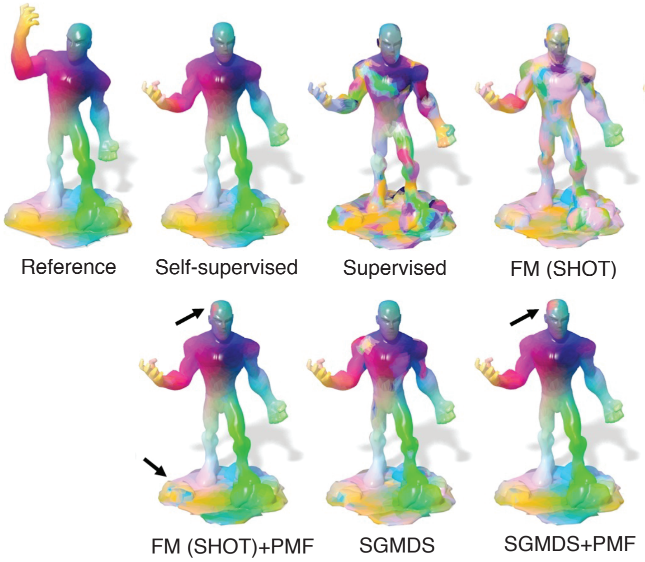

Variations on the unsupervised loss

- Unsupervised loss

$$\ell_{\mathsf{uns}} = \left\| \mathbf{D}_{\mathcal{X}} - \mathbf{P}^{\top} \mathbf{D}_{\mathcal{Y}} \mathbf{P} \right\|_{\mathsf{F}}^{2}$$

Variations on the unsupervised loss

- Unsupervised loss

$$\ell_{\mathsf{uns}} \;=\; \left\|\mathbf{D}_{\mathcal{X}} - \mathbf{P}^\top \mathbf{D}_{\mathcal{Y}} \mathbf{P}\right\|_{\mathsf{F}}^2 \;\approx\; \left\|\mathbf{\Phi}^\top \mathbf{D}_{\mathcal{X}} \mathbf{\Phi} - \mathbf{C}^\top \mathbf{\Psi}^\top \mathbf{D}_{\mathcal{Y}} \mathbf{\Psi} \mathbf{C}\right\|_{\mathsf{F}}^2$$

can be fully expressed in spectral domain

Variations on the unsupervised loss

- Unsupervised loss

$$\ell_{\mathsf{uns}} \;=\; \left\|\mathbf{D}_{\mathcal{X}} - \mathbf{P}^{\top}\mathbf{D}_{\mathcal{Y}}\mathbf{P}\right\|_{\mathsf{F}}^{2} \;\approx\; \left\|\mathbf{\Phi}^{\top}\mathbf{D}_{\mathcal{X}}\mathbf{\Phi} - \mathbf{C}^{\top}\mathbf{\Psi}^{\top}\mathbf{D}_{\mathcal{Y}}\mathbf{\Psi}\mathbf{C}\right\|_{\mathsf{F}}^{2}$$

  can be fully expressed in spectral domain

- No need to compute point-wise soft correspondence $\mathbf{P}$ from functional map

Variations on the unsupervised loss

- Unsupervised loss

$$\ell_{\mathsf{uns}} = \left\| \mathbf{D}_{\mathcal{X}} - \mathbf{P}^\top \mathbf{D}_{\mathcal{Y}} \mathbf{P} \right\|_{\mathsf{F}}^2 \approx \left\| \mathbf{\Phi}^\top \mathbf{D}_{\mathcal{X}} \mathbf{\Phi} - \mathbf{C}^\top \mathbf{\Psi}^\top \mathbf{D}_{\mathcal{Y}} \mathbf{\Psi} \mathbf{C} \right\|_{\mathsf{F}}^2$$

  can be fully expressed in spectral domain

- No need to compute point-wise soft correspondence $\mathbf{P}$ from functional map

- Geodesic distances can be approximated via heat equation

Variations on the unsupervised loss

- Unsupervised loss

$$\ell_{\text{uns}} = \left\| \mathbf{D}_{\mathcal{X}} - \mathbf{P}^\top \mathbf{D}_{\mathcal{Y}} \mathbf{P} \right\|_{\mathsf{F}}^2 \approx \left\| \mathbf{\Phi}^\top \mathbf{D}_{\mathcal{X}} \mathbf{\Phi} - \mathbf{C}^\top \mathbf{\Psi}^\top \mathbf{D}_{\mathcal{Y}} \mathbf{\Psi} \mathbf{C} \right\|_{\mathsf{F}}^2$$

  can be fully expressed in spectral domain

- No need to compute point-wise soft correspondence $\mathbf{P}$ from functional map
- Geodesic distances can be approximated via heat equation
- Other (better) distances between distributions should be used

Functional correspondence

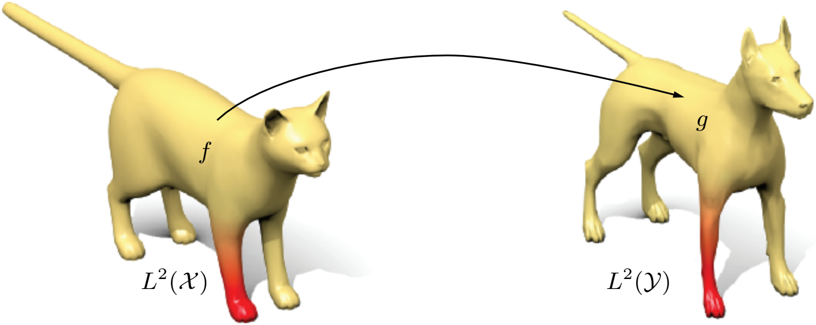

Functional map $T\colon L^2(\mathcal{X}) \to L^2(\mathcal{Y})$

Ovsjanikov et al. 2012

Pointwise correspondence

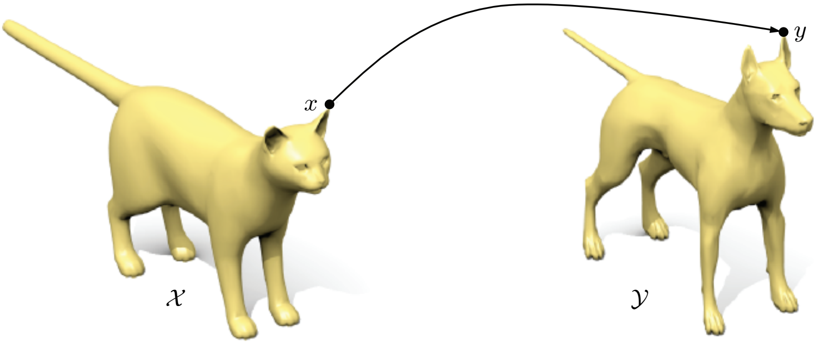

Point-wise map $t\colon \mathcal{X} \to \mathcal{Y}$

Pointwise correspondence

Recovered point-wise maps may be discontinuous and non-bijective

Correspondence in product space

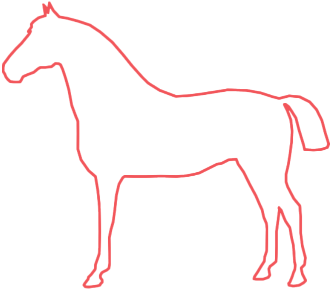

shape $\mathcal{X}$

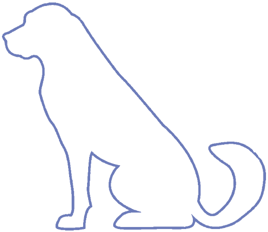

shape $\mathcal{Y}$

Vestner, Litman, Rodola, B, Cremers 2017

Correspondence in product space

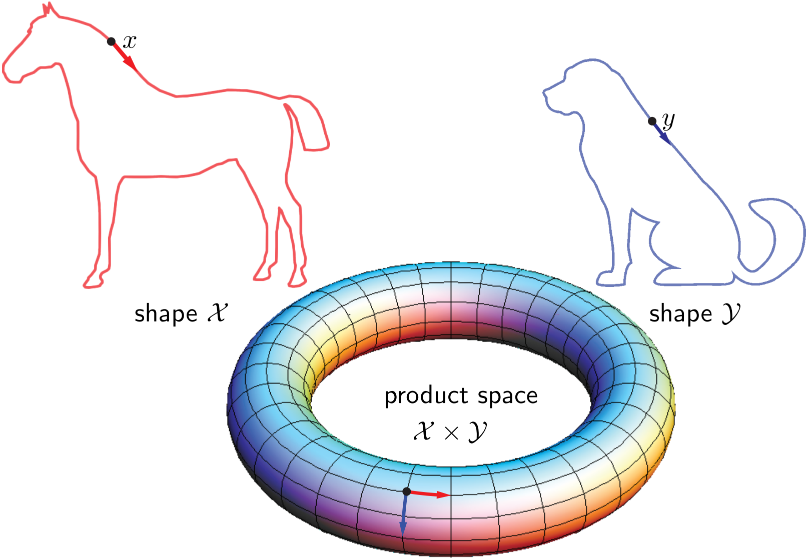

Vestner, Litman, Rodola, B, Cremers 2017

Correspondence in product space

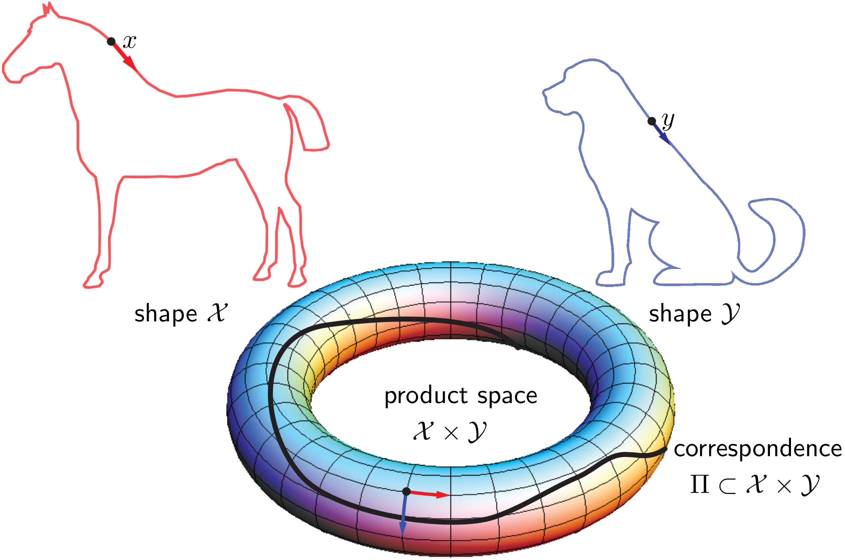

Vestner, Litman, Rodola, B, Cremers 2017

Alternating diffusion in product space

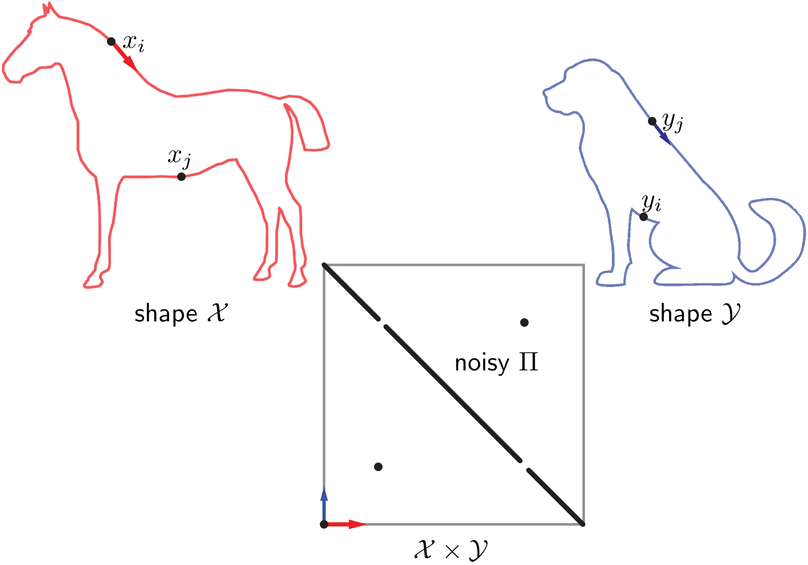

Lederman, Talmon 2015; Vestner, Litman, Rodola, B, Cremers 2017

Alternating diffusion in product space

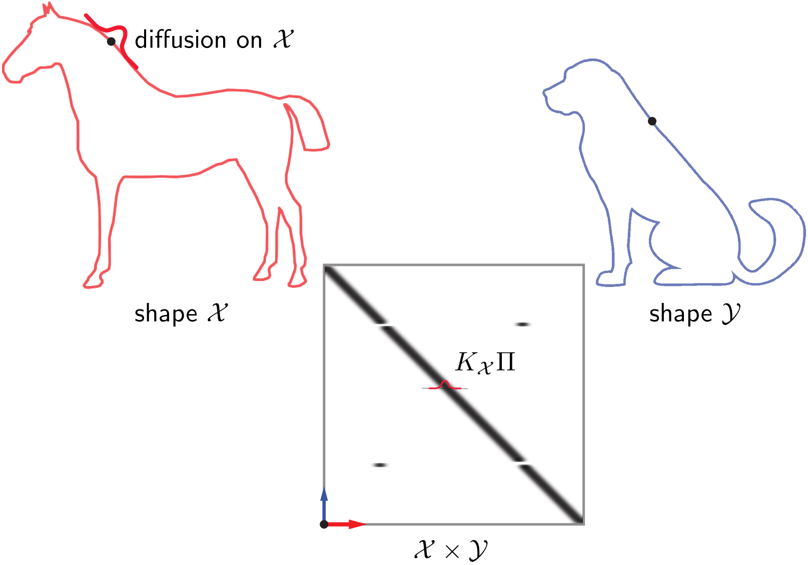

Lederman, Talmon 2015; Vestner, Litman, Rodola, B, Cremers 2017

Alternating diffusion in product space

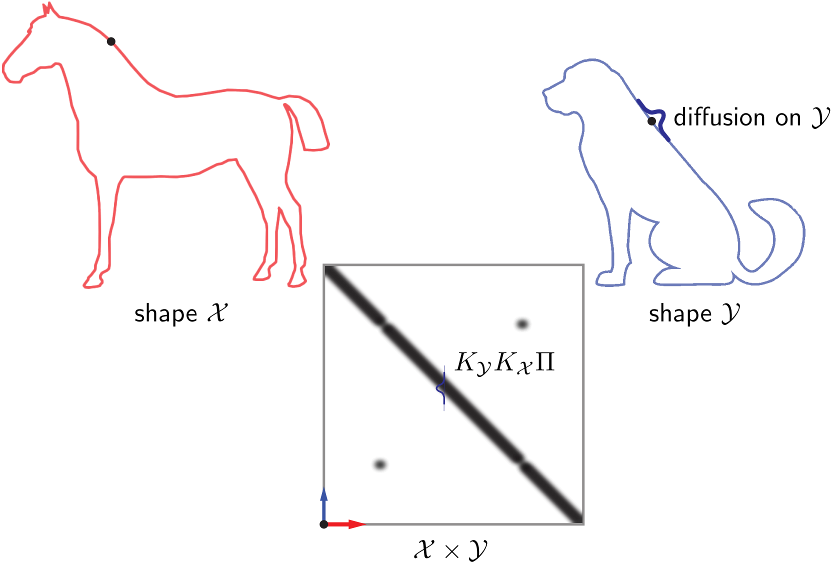

Lederman, Talmon 2015; Vestner, Litman, Rodola, B, Cremers 2017

Alternating diffusion in product space

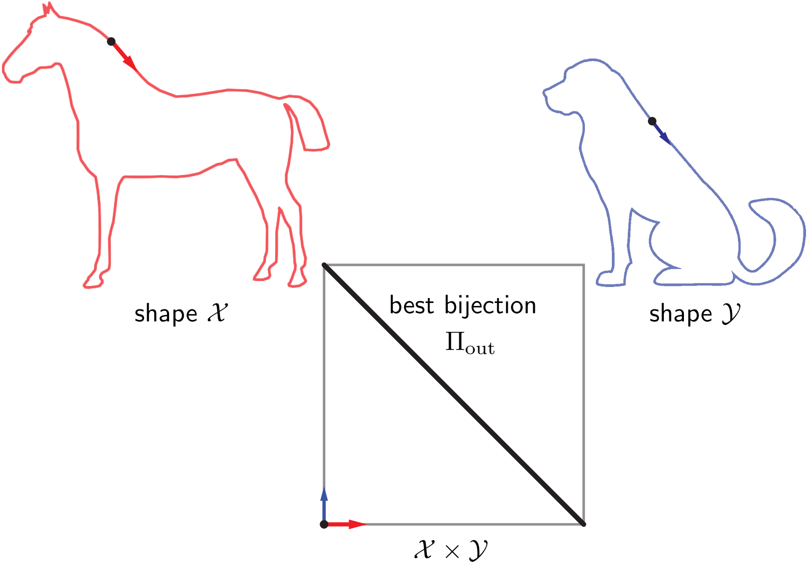

Lederman, Talmon 2015; Vestner, Litman, Rodola, B, Cremers 2017

1D illustration

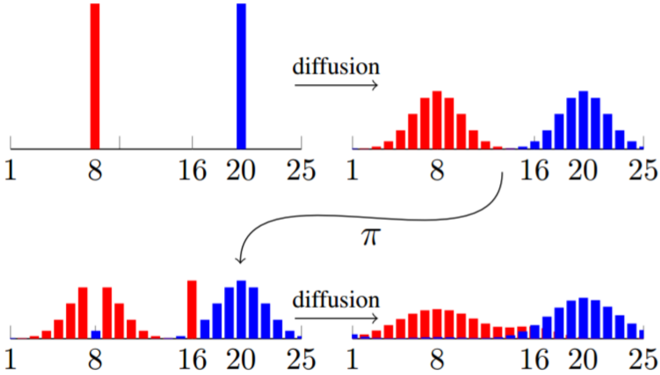

Vestner, Litman, R, Bronstein, Cremers 2017; Vestner, R *et al.* 2017

Product manifold filter

$$\boldsymbol{\Pi}_{\text{in}}$$

Input: initial correspondence

Vestner, Litman, Rodola, B, Cremers 2017; Vestner *et al.* 2017

Product manifold filter

$$\mathbf{\Pi}_{\text{in}}$$

Input: initial correspondence
(noisy, sparse, or a distribution)

Vestner, Litman, Rodola, B, Cremers 2017; Vestner *et al.* 2017

Product manifold filter

$$\mathbf{\Pi}_{\text{in}}\mathbf{K}_{\mathcal{X}}$$

Input: initial correspondence
(noisy, sparse, or a distribution)

Step 1: diffusion on $\mathcal{X}$

Vestner, Litman, Rodola, B, Cremers 2017; Vestner *et al.* 2017

Product manifold filter

$$\mathbf{K}_{\mathcal{Y}}\mathbf{\Pi}_{\mathrm{in}}\mathbf{K}_{\mathcal{X}}$$

Input: initial correspondence
(noisy, sparse, or a distribution)

Step 1: diffusion on $\mathcal{X}$

Step 2: diffusion on $\mathcal{Y}$

Vestner, Litman, Rodola, B, Cremers 2017; Vestner *et al.* 2017

Product manifold filter

$$\mathbf{\Pi}_{\text{out}} = \arg\max_{\mathbf{\Pi} \in \mathcal{P}_n} \langle \mathbf{\Pi}, \mathbf{K}_{\mathcal{Y}} \mathbf{\Pi}_{\text{in}} \mathbf{K}_{\mathcal{X}} \rangle$$

Input: initial correspondence
(noisy, sparse, or a distribution)

Step 1: diffusion on $\mathcal{X}$

Step 2: diffusion on $\mathcal{Y}$

Step 3: projection onto $\mathcal{P}_n$

Vestner, Litman, Rodola, B, Cremers 2017; Vestner *et al.* 2017

Product manifold filter

$$\mathbf{\Pi}_{k+1} = \arg\max_{\mathbf{\Pi} \in \mathcal{P}_n} \langle \mathbf{\Pi}, \mathbf{K}_{\mathcal{Y}} \mathbf{\Pi}_k \mathbf{K}_{\mathcal{X}} \rangle$$

Input: initial correspondence
(noisy, sparse, or a distribution)

Step 1: diffusion on $\mathcal{X}$

Step 2: diffusion on $\mathcal{Y}$

Step 3: projection onto $\mathcal{P}_n$

Iterate until convergence...

Vestner, Litman, Rodola, B, Cremers 2017; Vestner *et al.* 2017

Input: sparse correspondence

Init $= 20$ sparse corresponding points

Input: sparse correspondence

Input: sparse correspondence

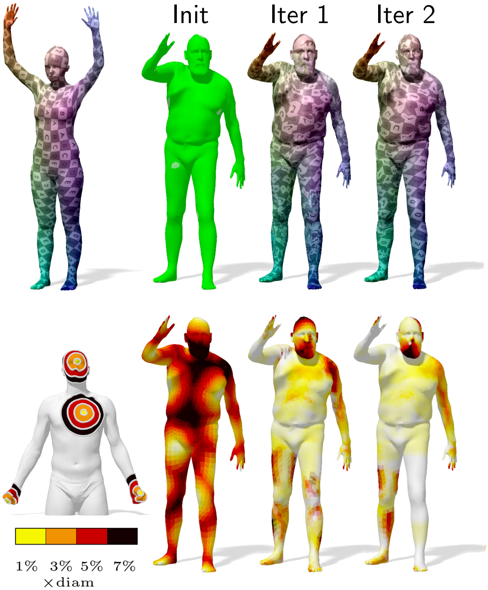

Input: sparse correspondence

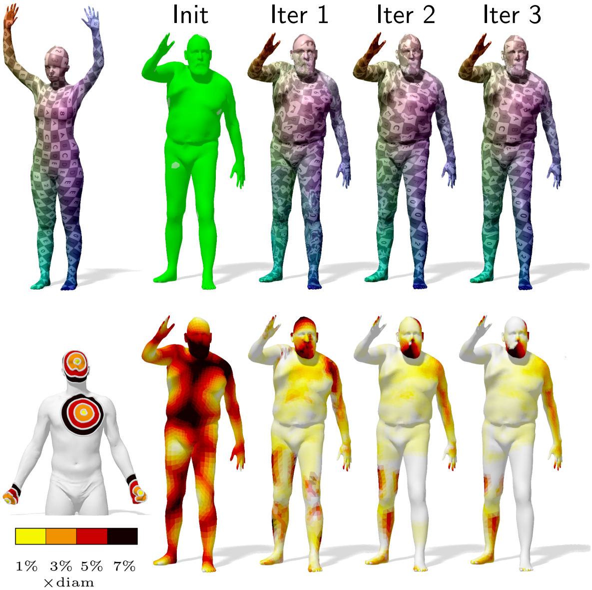

Input: sparse correspondence

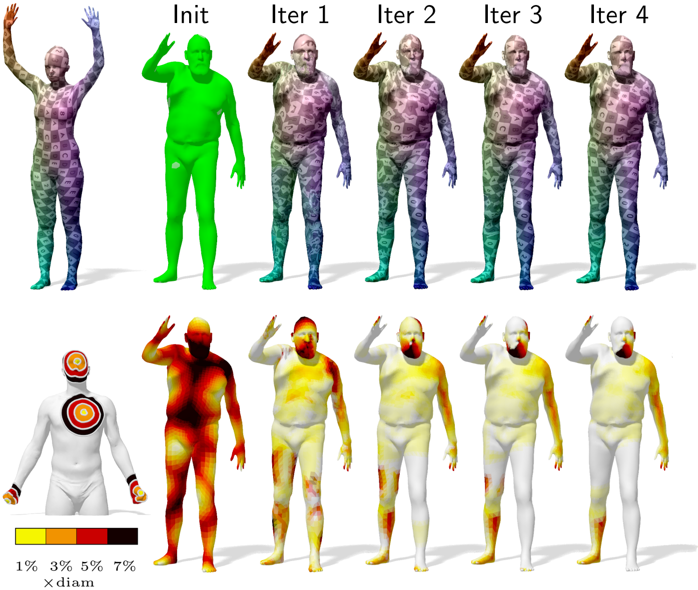

Input: sparse correspondence

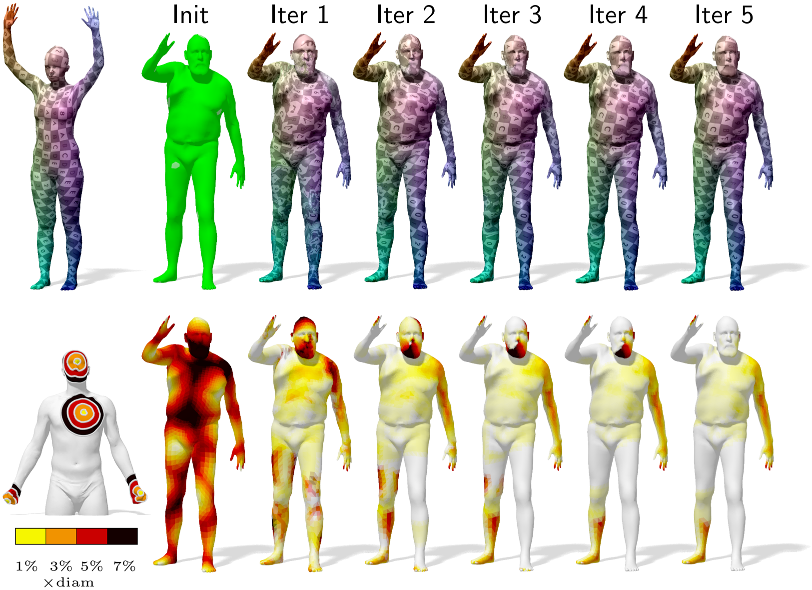

Input: sparse correspondence

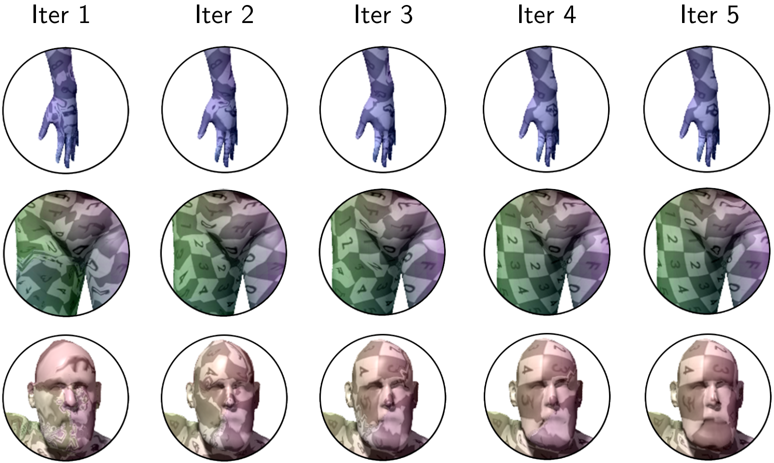

Input: functional correspondence

Input: functional correspondence

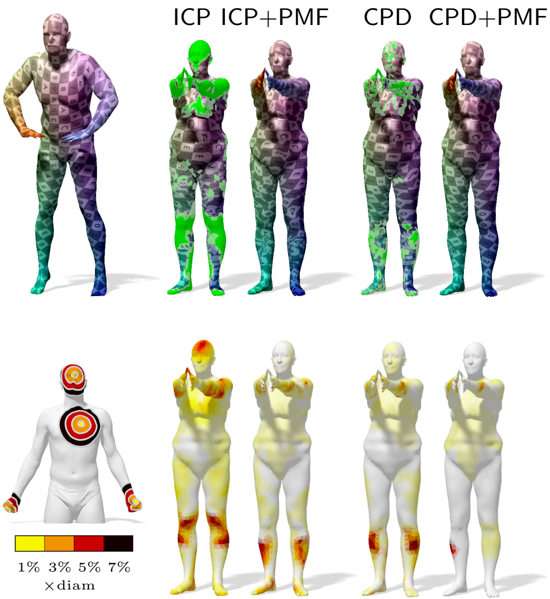

Input: functional correspondence

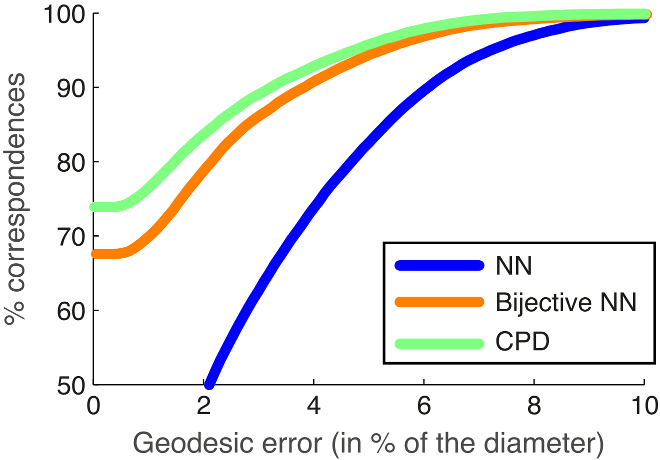

Data: SCAPE; Evaluation: Kim et al., 2011; CPD: Rodolà et al., 2015

Input: functional correspondence

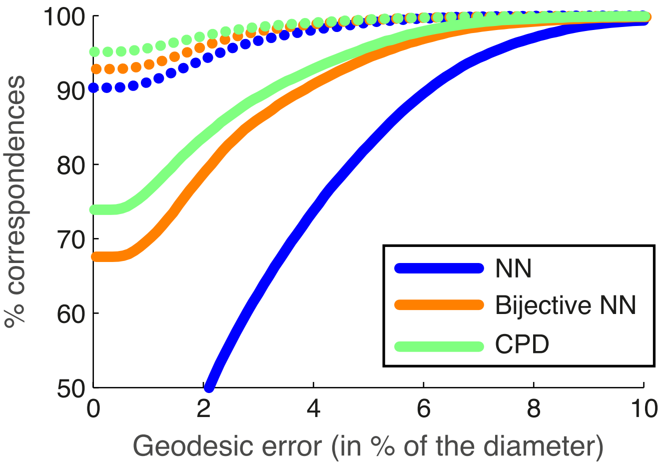

Data: SCAPE; Evaluation: Kim et al., 2011; CPD: Rodolà et al., 2015

FAUST benchmark

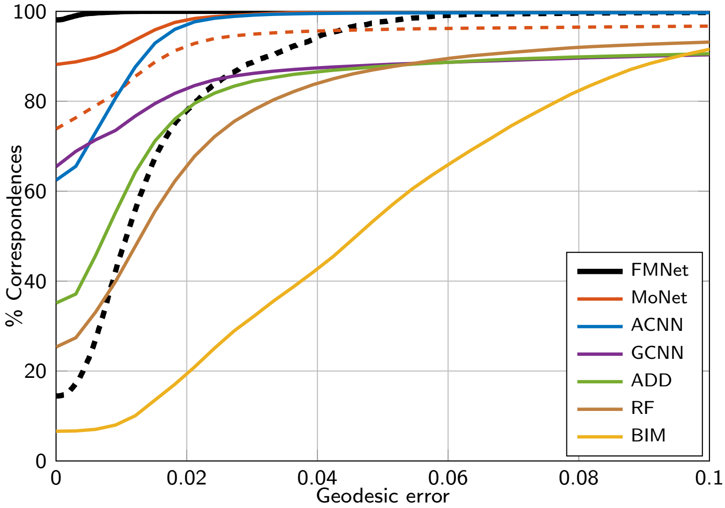

data: FAUST (Bogo et al. 2015)

Alternative interpretations

- Kernel matching **quadratic assignment problem**

Vestner, Litman, Rodola, B, Cremers 2017; Vestner *et al.* 2017

Quadratic assignment problem

$$\mathbf{\Pi}_{k+1} = \arg\max_{\mathbf{\Pi} \in \mathcal{P}_n} \langle \mathbf{\Pi}, \mathbf{K}_{\mathcal{Y}} \mathbf{\Pi}_k \mathbf{K}_{\mathcal{X}} \rangle$$

Vestner, Litman, Rodola, B, Cremers 2017; Vestner *et al.* 2017

Quadratic assignment problem

$$\mathbf{\Pi}_* = \arg\max_{\mathbf{\Pi}\in\mathcal{P}_n} \langle \mathbf{\Pi}, \mathbf{K}_{\mathcal{Y}}\mathbf{\Pi}\mathbf{K}_{\mathcal{X}} \rangle$$

A sequence of linearizations of a quadratic assignment problem

Vestner, Litman, Rodola, B, Cremers 2017; Vestner *et al.* 2017

Quadratic assignment problem

$$\mathbf{\Pi}_* = \arg\min_{\mathbf{\Pi} \in \mathcal{P}_n} \|\mathbf{\Pi}\mathbf{K}_{\mathcal{X}} - \mathbf{K}_{\mathcal{Y}}\mathbf{\Pi}\|_{\mathrm{F}}$$

A sequence of linearizations of a quadratic assignment problem

Minimizes pairwise kernel distortion

Vestner, Litman, Rodola, B, Cremers 2017; Vestner *et al.* 2017

Alternative interpretations

- Kernel matching **quadratic assignment problem**

Vestner, Litman, Rodola, B, Cremers 2017; Vestner *et al.* 2017

Alternative interpretations

- Kernel matching **quadratic assignment problem**
- **Dirichlet energy minimization** in product space

Vestner, Litman, Rodola, B, Cremers 2017; Vestner *et al.* 2017

Dirichlet energy minimization in product space

Dirichlet energy minimization on $\mathcal{X}$:

$$\arg\min \mathbf{f}^\top \mathbf{L}_{\mathcal{X}} \mathbf{f}$$

Vestner et al. 2017

Dirichlet energy minimization in product space

Dirichlet energy minimization on $\mathcal{X}$:

$$\arg\min \mathbf{f}^\top \mathbf{L}_\mathcal{X} \mathbf{f} = \arg\max \mathbf{f}^\top e^{-t\,\mathbf{L}_\mathcal{X}} \mathbf{f}$$

Vestner et al. 2017

Dirichlet energy minimization in product space

Dirichlet energy minimization on $\mathcal{X}$:

$$\arg\min \mathbf{f}^\top \mathbf{L}_{\mathcal{X}} \mathbf{f} = \arg\max \mathbf{f}^\top \mathbf{K}_{\mathcal{X}} \mathbf{f}$$

Vestner et al. 2017

Dirichlet energy minimization in product space

Dirichlet energy minimization on $\mathcal{X}$:

$$\arg\min \mathbf{f}^\top \mathbf{L}_\mathcal{X} \mathbf{f} = \arg\max \mathbf{f}^\top \mathbf{K}_\mathcal{X} \mathbf{f}$$

Kernel matching energy

$$\arg\max \langle \mathbf{\Pi}, \mathbf{K}_\mathcal{Y} \mathbf{\Pi} \mathbf{K}_\mathcal{X} \rangle$$

Vestner et al. 2017

Dirichlet energy minimization in product space

Dirichlet energy minimization on $\mathcal{X}$:

$$\arg\min \mathbf{f}^\top \mathbf{L}_{\mathcal{X}} \mathbf{f} = \arg\max \mathbf{f}^\top \mathbf{K}_{\mathcal{X}} \mathbf{f}$$

Kernel matching energy

$$\arg\max \langle \mathbf{\Pi}, \mathbf{K}_{\mathcal{Y}} \mathbf{\Pi} \mathbf{K}_{\mathcal{X}} \rangle = \arg\max \boldsymbol{\pi}^\top (\mathbf{K}_{\mathcal{X}} \otimes \mathbf{K}_{\mathcal{Y}}) \boldsymbol{\pi}$$

Vestner et al. 2017

Dirichlet energy minimization in product space

Dirichlet energy minimization on $\mathcal{X}$:

$$\arg\min \mathbf{f}^\top \mathbf{L}_\mathcal{X} \mathbf{f} = \arg\max \mathbf{f}^\top \mathbf{K}_\mathcal{X} \mathbf{f}$$

Kernel matching energy

$$\arg\max \langle \mathbf{\Pi}, \mathbf{K}_\mathcal{Y} \mathbf{\Pi} \mathbf{K}_\mathcal{X} \rangle = \arg\max \boldsymbol{\pi}^\top \mathbf{K}_{\mathcal{X}\times\mathcal{Y}} \boldsymbol{\pi}$$

Vestner et al. 2017

Dirichlet energy minimization in product space

Dirichlet energy minimization on $\mathcal{X}$:

$$\arg\min \mathbf{f}^\top \mathbf{L}_{\mathcal{X}} \mathbf{f} = \arg\max \mathbf{f}^\top \mathbf{K}_{\mathcal{X}} \mathbf{f}$$

Kernel matching energy $=$ Dirichlet energy minimization on $\mathcal{X} \times \mathcal{Y}$

$$\arg\max \langle \mathbf{\Pi}, \mathbf{K}_{\mathcal{Y}} \mathbf{\Pi} \mathbf{K}_{\mathcal{X}} \rangle = \arg\max \boldsymbol{\pi}^\top \mathbf{K}_{\mathcal{X}\times\mathcal{Y}} \boldsymbol{\pi}$$

Vestner et al. 2017

Alternative interpretations

- Kernel matching **quadratic assignment problem**
- **Dirichlet energy minimization** in product space

Vestner, Litman, Rodola, B, Cremers 2017; Vestner *et al.* 2017

Alternative interpretations

- Kernel matching **quadratic assignment problem**
- **Dirichlet energy minimization** in product space
- **Iterated blurring/sharpening**

Vestner, Litman, Rodola, B, Cremers 2017; Vestner *et al.* 2017

Iterated blurring and sharpening

A diffusion process corresponds to low-pass filtering in the spectral domain:

$$\mathbf{K}_{\mathcal{Y}}\mathbf{\Pi}\mathbf{K}_{\mathcal{X}}$$

Vestner et al. 2017

Iterated blurring and sharpening

A diffusion process corresponds to low-pass filtering in the spectral domain:

$$\mathbf{K}_{\mathcal{Y}}\mathbf{\Pi}\mathbf{K}_{\mathcal{X}} = \mathbf{\Psi}e^{-t\mathbf{\Lambda}_{\mathcal{Y}}}\mathbf{\Psi}^{\top}\mathbf{\Pi}\mathbf{\Phi}e^{-t\mathbf{\Lambda}_{\mathcal{X}}}\mathbf{\Phi}^{\top}$$

Vestner et al. 2017

Iterated blurring and sharpening

A diffusion process corresponds to low-pass filtering in the spectral domain:

$$\mathbf{K}_{\mathcal{Y}}\mathbf{\Pi}\mathbf{K}_{\mathcal{X}} = \mathbf{\Psi} e^{-t\mathbf{\Lambda}_{\mathcal{Y}}} \underbrace{\mathbf{\Psi}^{\top}\mathbf{\Pi}\mathbf{\Phi}}_{\mathbf{C}} e^{-t\mathbf{\Lambda}_{\mathcal{X}}} \mathbf{\Phi}^{\top}$$

Vestner et al. 2017

Iterated blurring and sharpening

A diffusion process corresponds to low-pass filtering in the spectral domain:

$$\mathbf{K}_{\mathcal{Y}}\mathbf{\Pi}\mathbf{K}_{\mathcal{X}} = \mathbf{\Psi} e^{-t\mathbf{\Lambda}_{\mathcal{Y}}} \mathbf{C} e^{-t\mathbf{\Lambda}_{\mathcal{X}}} \mathbf{\Phi}^\top$$

Vestner et al. 2017

Iterated blurring and sharpening

A diffusion process corresponds to low-pass filtering in the spectral domain:

$$\mathbf{K}_{\mathcal{Y}}\mathbf{\Pi}\mathbf{K}_{\mathcal{X}} = \mathbf{\Psi} e^{-t\mathbf{\Lambda}_{\mathcal{Y}}} \mathbf{C} e^{-t\mathbf{\Lambda}_{\mathcal{X}}} \mathbf{\Phi}^{\top}$$

- Low-pass filtering of the functional map matrix $\mathbf{C}$

Vestner et al. 2017

Iterated blurring and sharpening

A diffusion process corresponds to low-pass filtering in the spectral domain:

$$\mathbf{K}_{\mathcal{Y}}\mathbf{\Pi}\mathbf{K}_{\mathcal{X}} = \mathbf{\Psi} e^{-t\mathbf{\Lambda}_{\mathcal{Y}}}\mathbf{C}e^{-t\mathbf{\Lambda}_{\mathcal{X}}}\mathbf{\Phi}^{\top}$$

- Low-pass filtering of the functional map matrix $\mathbf{C}$
- Conversion to point-wise map

Vestner et al. 2017

Iterated blurring and sharpening

A diffusion process corresponds to low-pass filtering in the spectral domain:

$$\mathbf{K}_{\mathcal{Y}}\mathbf{\Pi}\mathbf{K}_{\mathcal{X}} = \mathbf{\Psi} e^{-t\mathbf{\Lambda}_{\mathcal{Y}}}\mathbf{C} e^{-t\mathbf{\Lambda}_{\mathcal{X}}}\mathbf{\Phi}^{\top}$$

- Low-pass filtering of the functional map matrix $\mathbf{C}$
- Conversion to point-wise map

The smooth window reduces Gibbs oscillations due to truncation:

Vestner et al. 2017

Non-isometric deformations

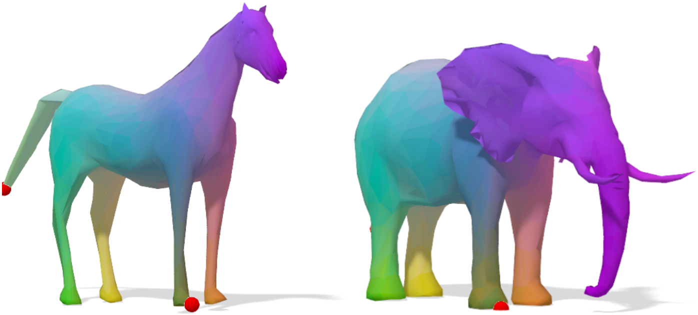

Input: 2 corresponding pairs

2D to 3D matching

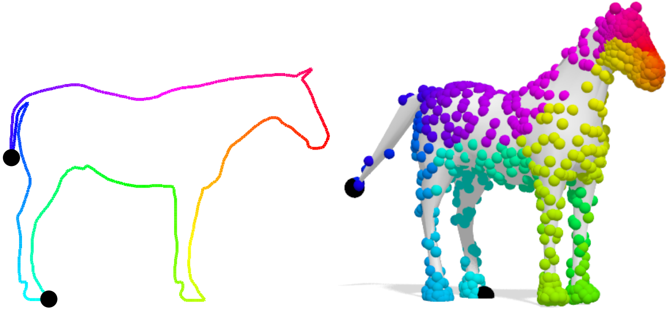

Input: 2 corresponding pairs

Discontinuous matching

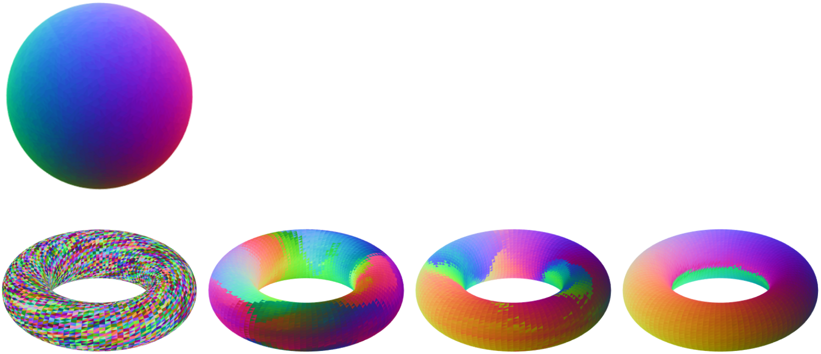

Random init 1 iteration 2 iterations 3 iterations

Minimizes discontinuity length?

Kernel choice

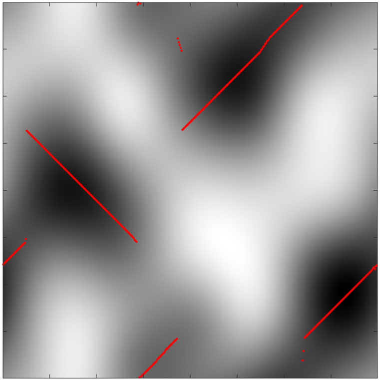

Random initialization, large kernel variance

Kernel choice

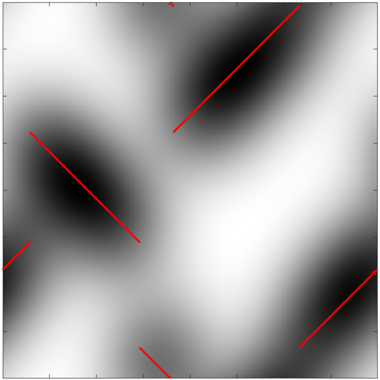

Random initialization, large kernel variance

Kernel choice

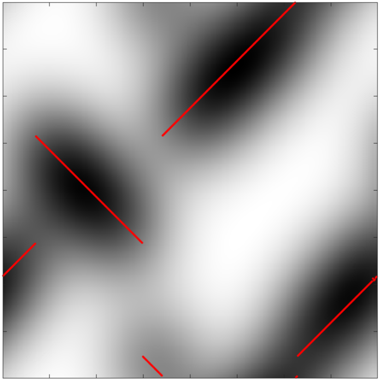

Random initialization, large kernel variance

Kernel choice

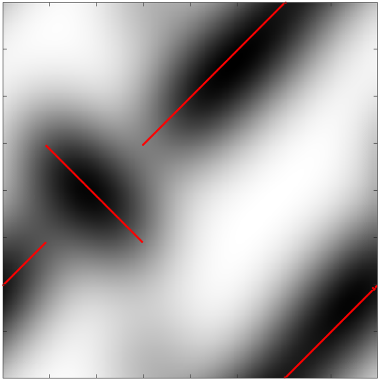

Random initialization, large kernel variance

Kernel choice

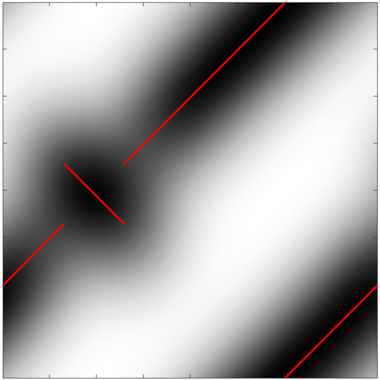

Random initialization, large kernel variance

Kernel choice

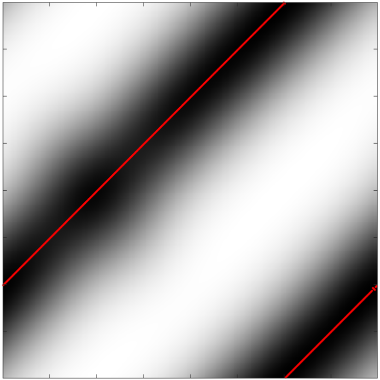

Random initialization, large kernel variance

Kernel choice

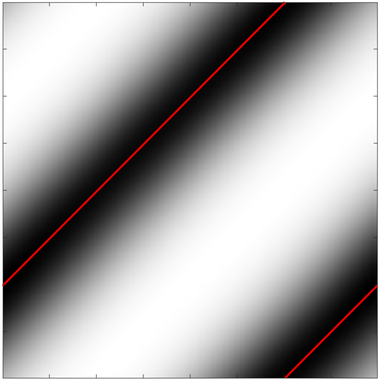

Random initialization, large kernel variance

Kernel choice

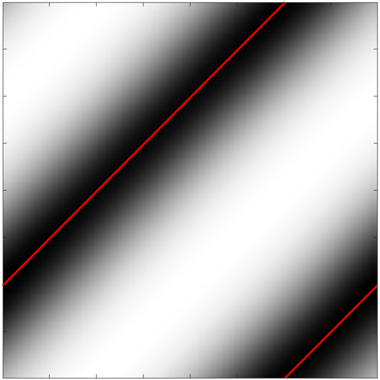

Random initialization, large kernel variance

Kernel choice

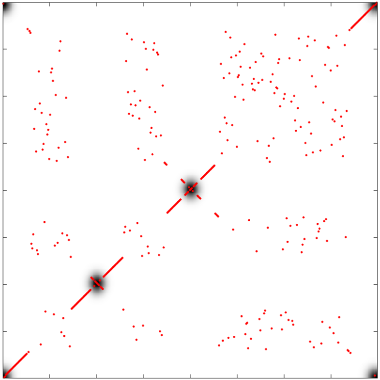

Random initialization, small kernel variance

Kernel choice

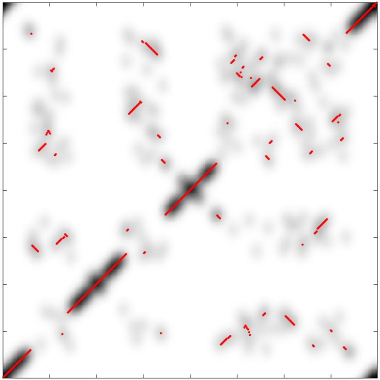

Random initialization, small kernel variance

Kernel choice

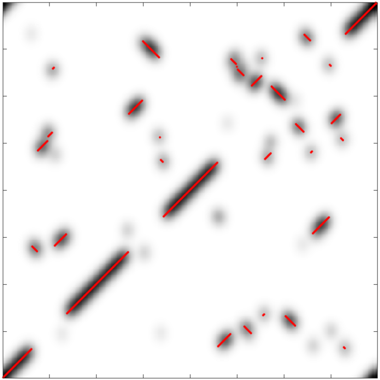

Random initialization, small kernel variance

Kernel choice

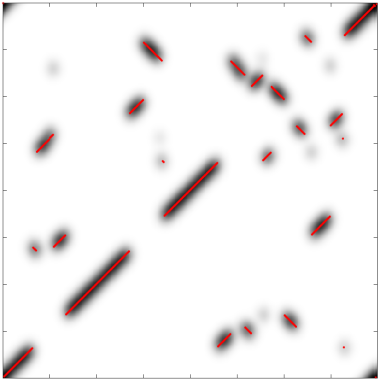

Random initialization, small kernel variance

Kernel choice

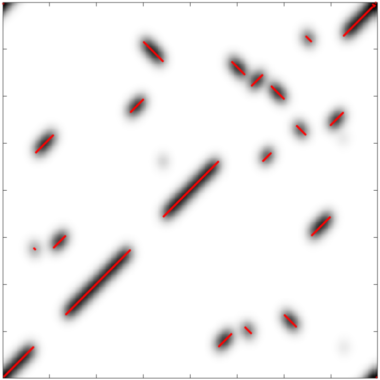

Random initialization, small kernel variance

Kernel choice

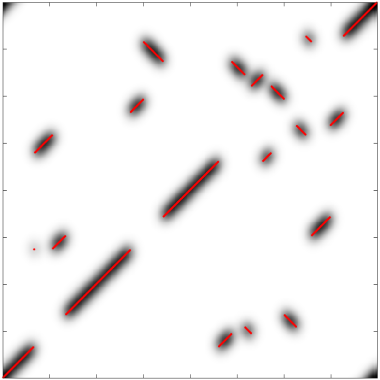

Random initialization, small kernel variance

Kernel choice

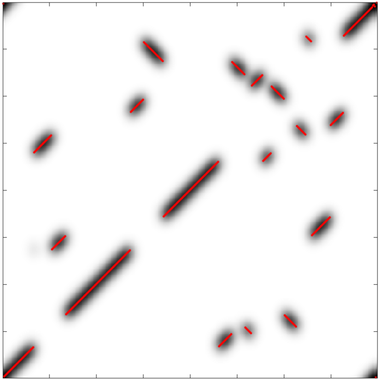

Random initialization, small kernel variance

Kernel choice

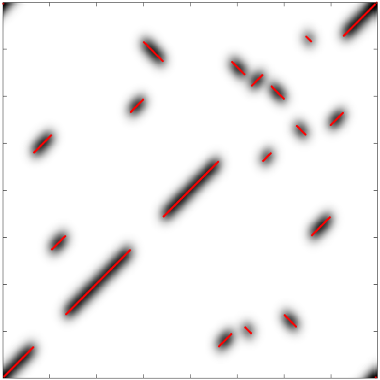

Random initialization, small kernel variance

Geometry and learning in deformable shape correspondence

- Incorporated **problem geometric structure** into descriptor learning

Vestner, Litman, Rodola, B, Cremers 2017; Vestner *et al.* 2017

Geometry and learning in deformable shape correspondence

- Incorporated **problem geometric structure** into descriptor learning: tiny net, little training data, state-of-the-art results

Vestner, Litman, Rodola, B, Cremers 2017; Vestner *et al.* 2017

Geometry and learning in deformable shape correspondence

- Incorporated **problem geometric structure** into descriptor learning: tiny net, little training data, state-of-the-art results
- Regarded bijective correspondence recovery as **denosing** in product space

Vestner, Litman, Rodola, B, Cremers 2017; Vestner *et al.* 2017

Geometry and learning in deformable shape correspondence

- Incorporated **problem geometric structure** into descriptor learning: tiny net, little training data, state-of-the-art results
- Regarded bijective correspondence recovery as **denosing** in product space
- Works surprisingly well and fast

Vestner, Litman, Rodola, B, Cremers 2017; Vestner *et al.* 2017

Geometry and learning in deformable shape correspondence

- Incorporated **problem geometric structure** into descriptor learning: tiny net, little training data, state-of-the-art results
- Regarded bijective correspondence recovery as **denosing** in product space
- Works surprisingly well and fast
- Many interesting interpretations

Vestner, Litman, Rodola, B, Cremers 2017; Vestner *et al.* 2017

Geometry and learning in deformable shape correspondence

- Incorporated **problem geometric structure** into descriptor learning: tiny net, little training data, state-of-the-art results
- Regarded bijective correspondence recovery as **denosing** in product space
- Works surprisingly well and fast
- Many interesting interpretations
- **Kernel learning** – work in progress

Vestner, Litman, Rodola, B, Cremers 2017; Vestner *et al.* 2017

Geometry and learning in deformable shape correspondence

- Incorporated **problem geometric structure** into descriptor learning: tiny net, little training data, state-of-the-art results
- Regarded bijective correspondence recovery as **denosing** in product space
- Works surprisingly well and fast
- Many interesting interpretations
- **Kernel learning** – work in progress
- Code available!

Vestner, Litman, Rodola, B, Cremers 2017; Vestner *et al.* 2017