Learning kernels to match: Efficient Deformable Shape Correspondence via Kernel Matching

Amit Boyarski

July 2018

Technion
Israel Institute of Technology

erc

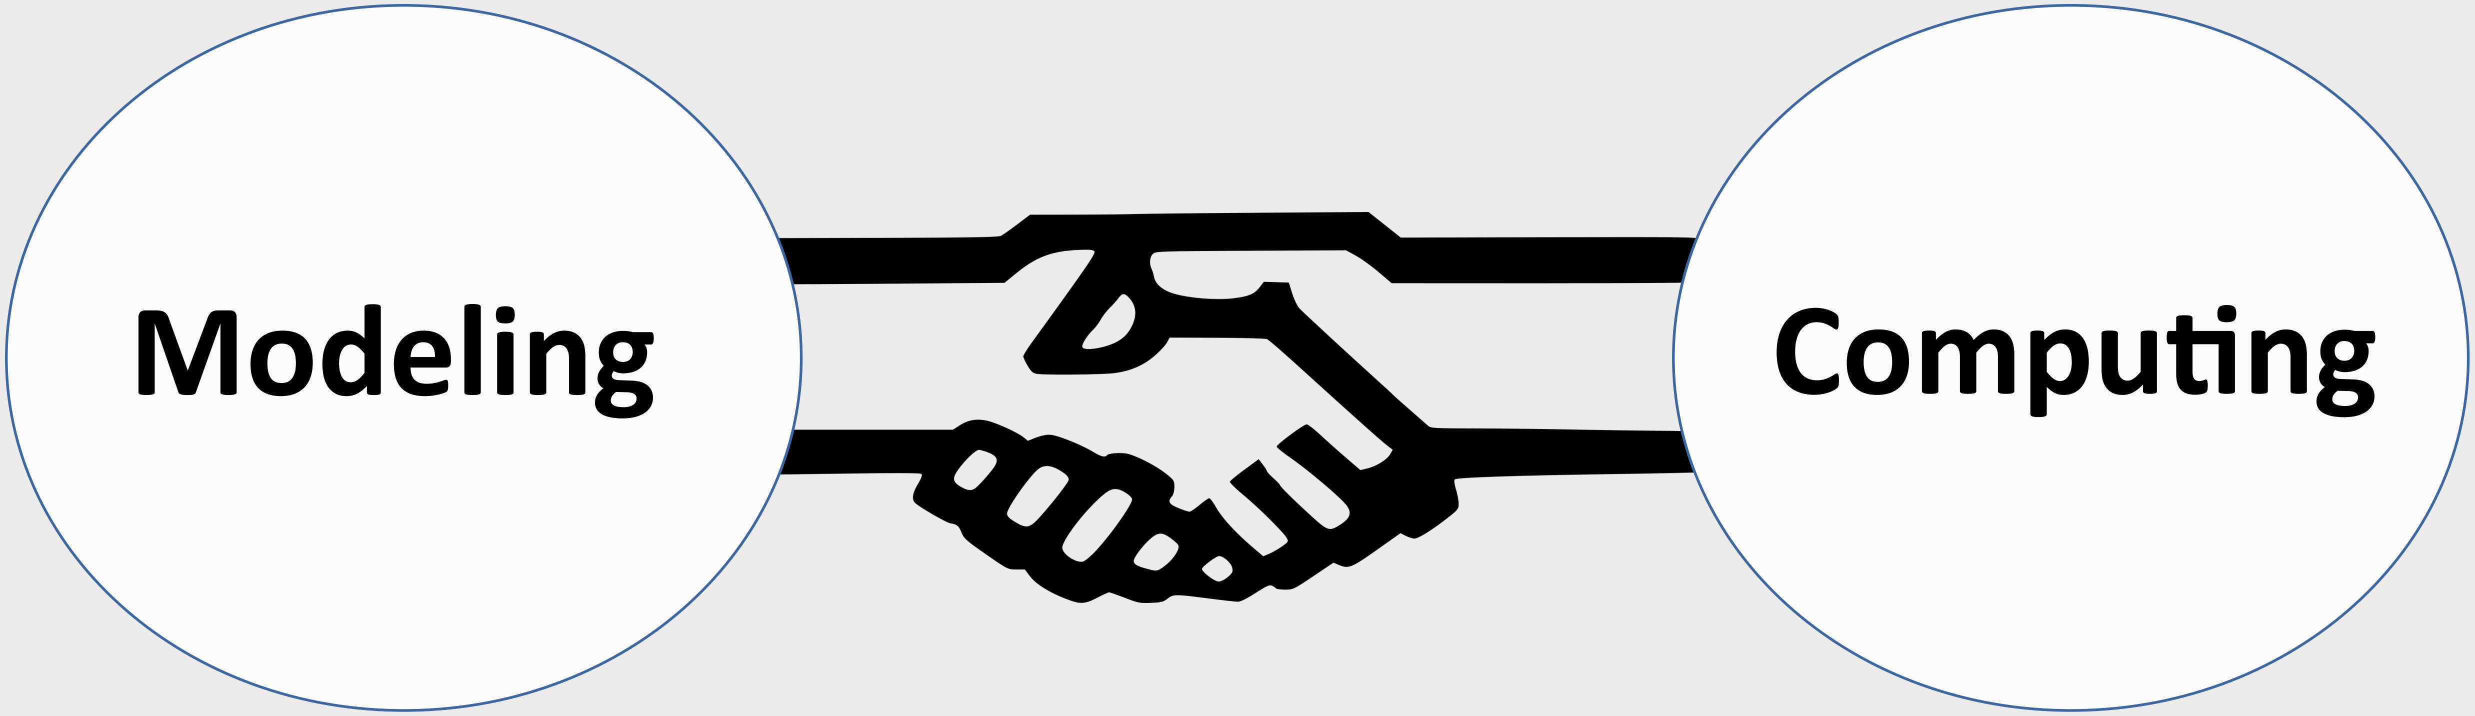
Modeling
Computing

Shape Correspondence

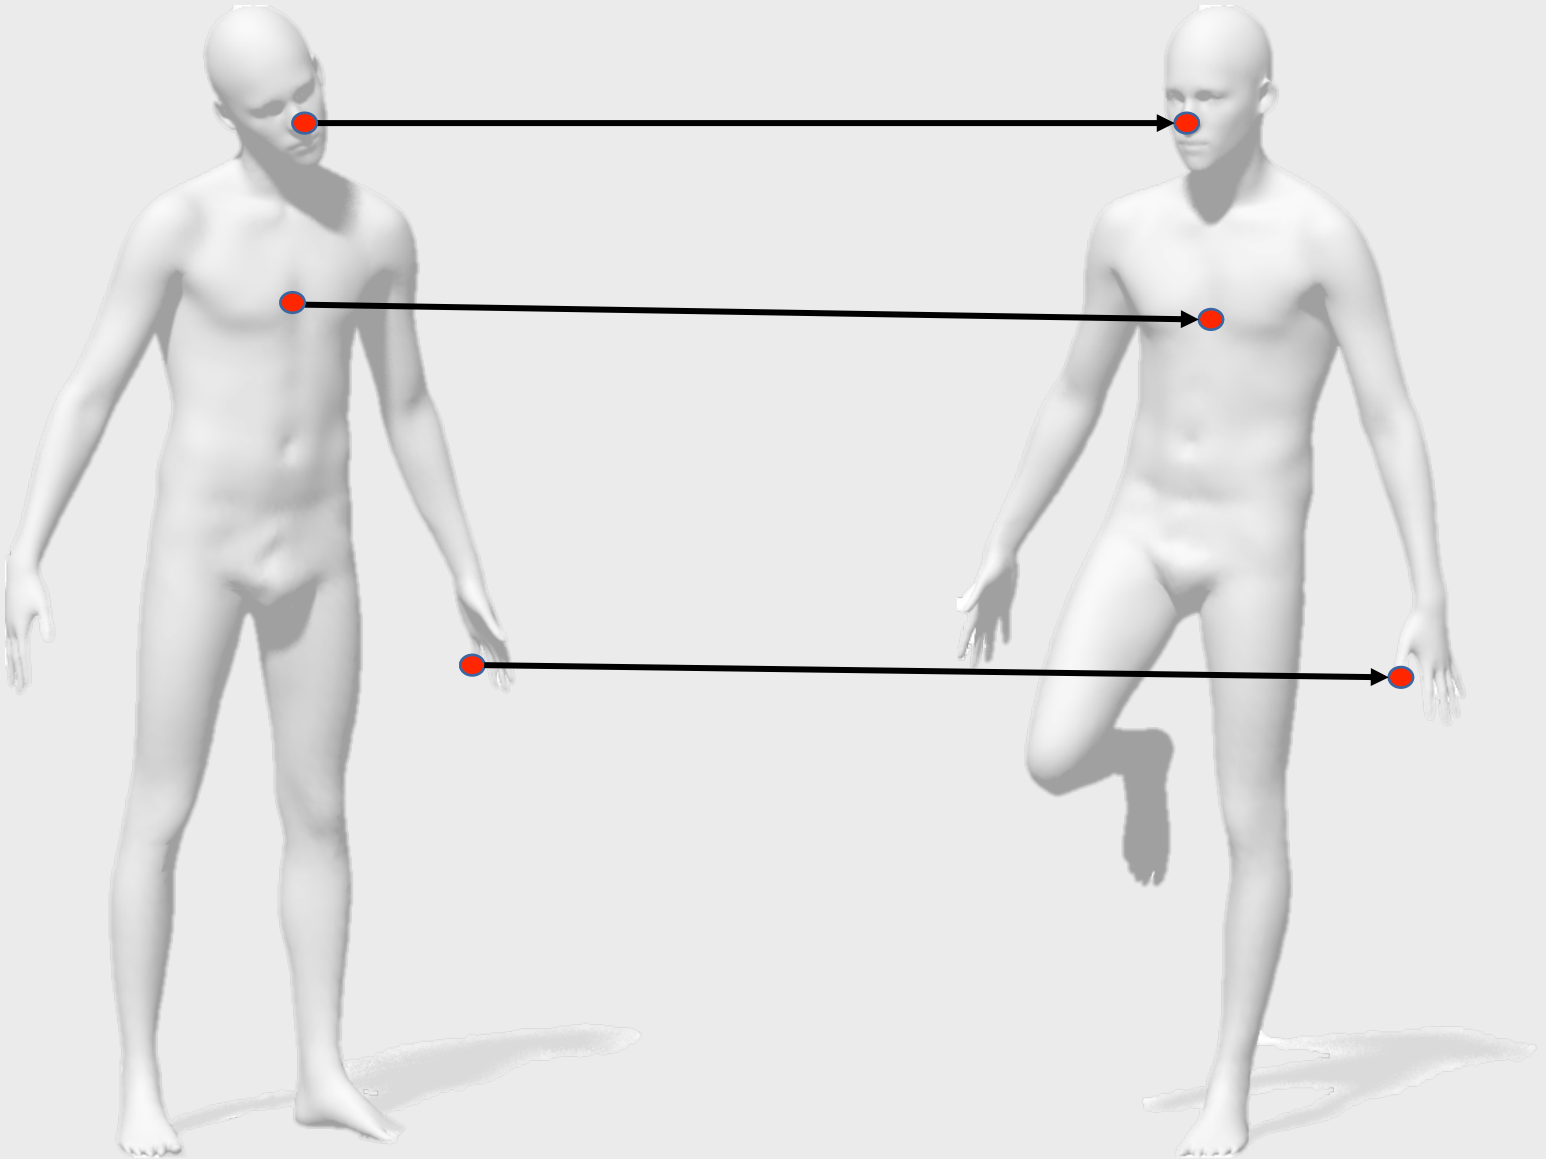

Shape Correspondence

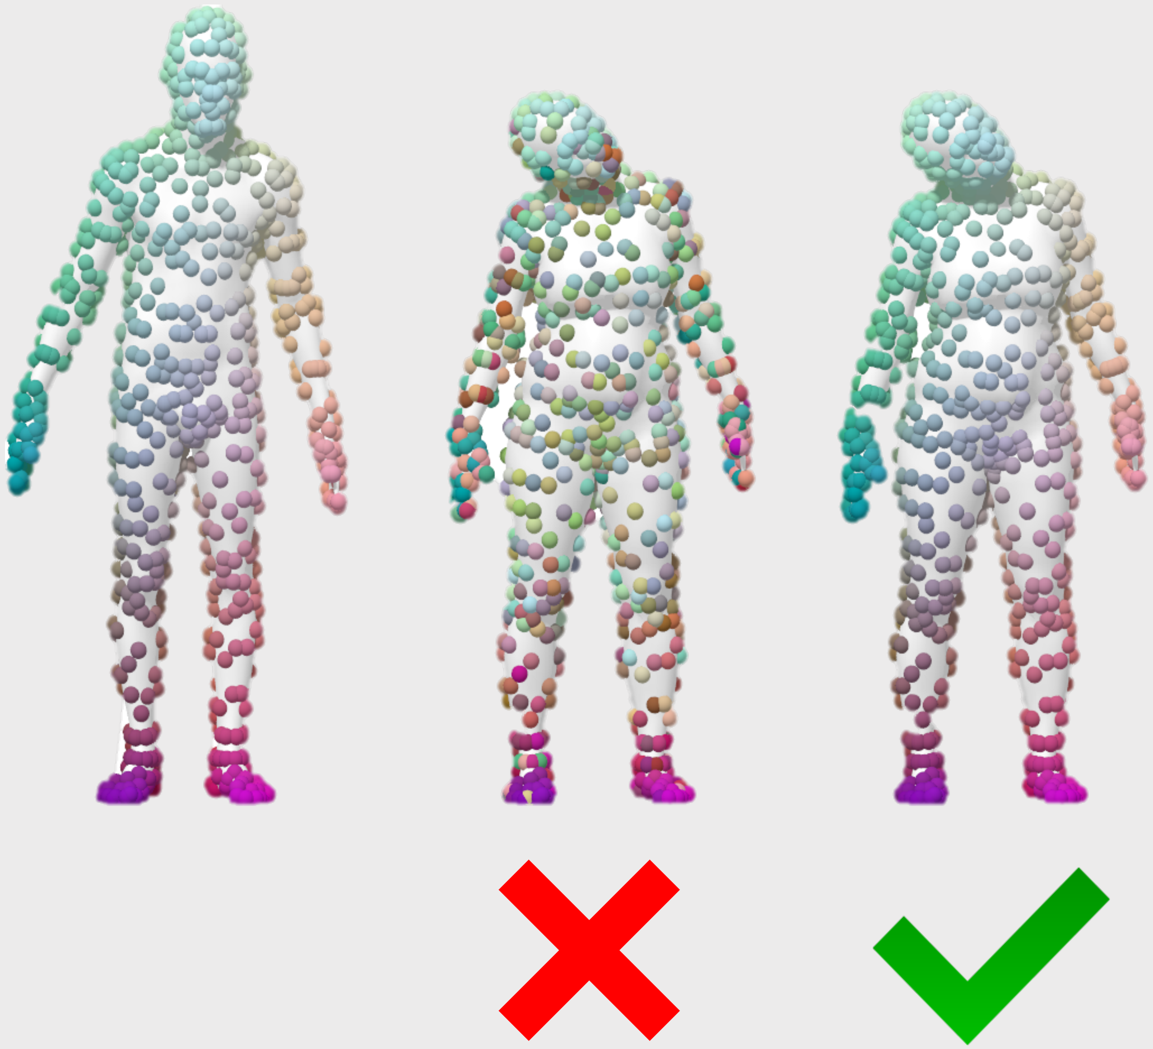

General overview

In the discrete setting, a (bijective) correspondence can be represented as a permutation matrix

$$\mathbf{\Pi}\mathbf{1} = \mathbf{1}$$

$$\mathbf{\Pi}^\top\mathbf{1} = \mathbf{1}$$

$$\Pi_{ij} \in \{0, 1\}$$

$$\underbrace{\hphantom{\Pi_{ij} \in \{0, 1\}}}_{\mathcal{P}_n}$$

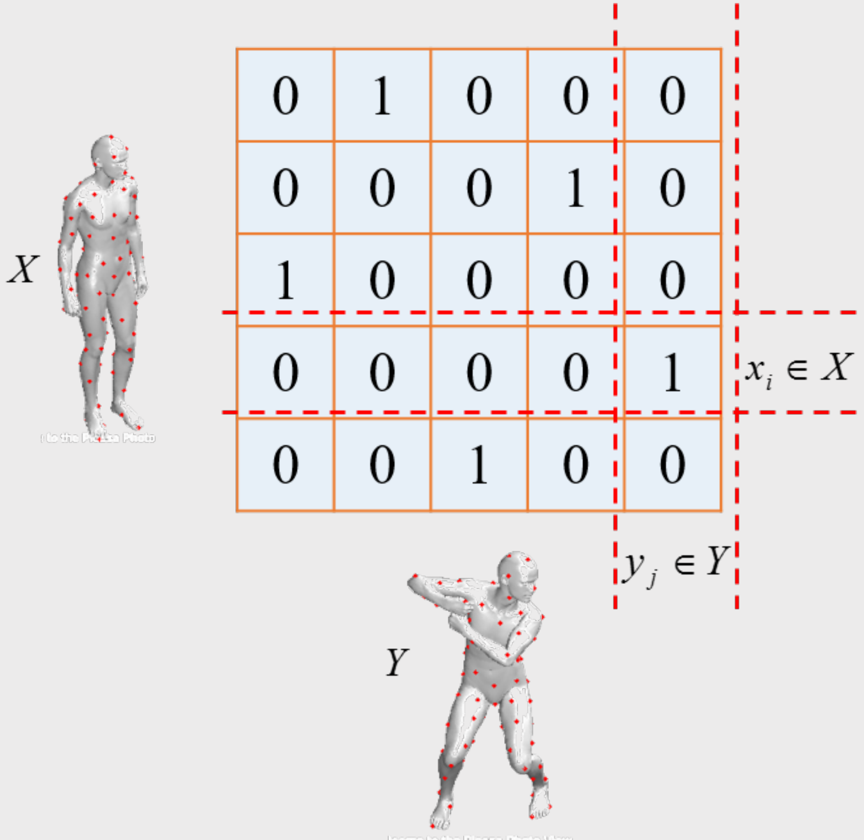

General overview

$$\min_{\mathbf{\Pi} \in \mathcal{P}_n} \langle \mathbf{\Pi}, \mathbf{F} \rangle = \min_{\mathbf{P} \in \mathcal{B}_n} \langle \mathbf{P}, \mathbf{F} \rangle$$

$$\mathbf{\Pi 1} = \mathbf{1}$$

$$\mathbf{\Pi}^\top \mathbf{1} = \mathbf{1}$$

$$\Pi_{ij} \in \{0, 1\}$$

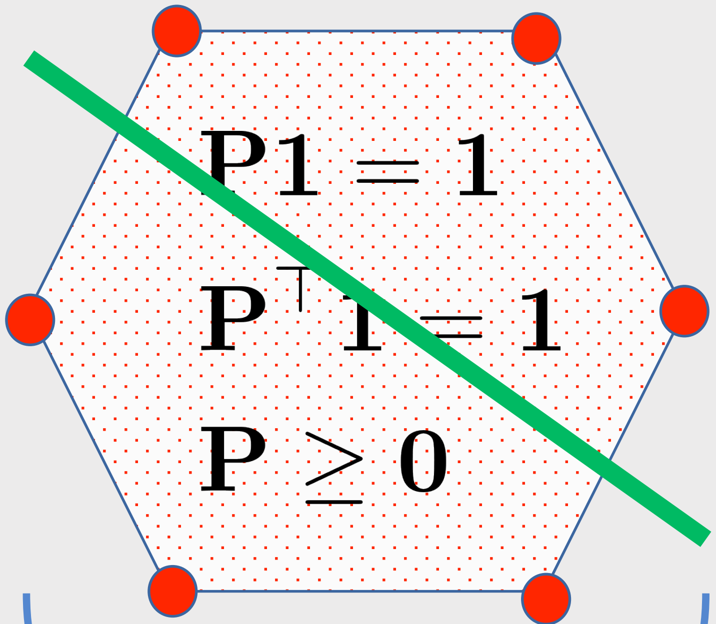

Birkhoff Von-Neumann Theorem:

$$\mathcal{P}_n = \text{extremePts}(\mathcal{B}_n)$$

$$\mathcal{P}_n$$

$$\mathcal{B}_n = \text{conv}(\mathcal{P}_n)$$

General overview

$$\arg\max_{\mathbf{\Pi}\in\mathcal{P}_n} E(\mathbf{\Pi}) =$$

$$\arg\max_{\mathbf{\Pi}\in\mathcal{P}_n} \alpha \langle \mathbf{\Pi}, \mathbf{F}_{\mathcal{Y}}\mathbf{F}_{\mathcal{X}}^{\top} \rangle + \langle \mathbf{\Pi}, \mathbf{K}_{\mathcal{Y}}\mathbf{\Pi}\mathbf{K}_{\mathcal{X}} \rangle$$

LAP

QAP

Bijectivity.

Data term.
Pointwise descriptors.
"Embeddings".

Regularization term.
Pairwise descriptors.

A general framework

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

- Relaxation: $$\mathbf{P}^k = \arg\max_{\mathbf{P}\in\mathcal{RP}_n} f(\mathbf{P}, \mathbf{\Pi}^k)$$
- Projection on $\mathcal{P}_n$: $$\mathbf{\Pi}^{k+1} = \mathcal{P}\left(\mathbf{P}^k\right)$$

Relaxation and **projection** can be really simple if you choose $\mathbf{K}_{\mathcal{Y}}, \mathbf{K}_{\mathcal{X}}$ wisely!

Laplace-Beltrami Eigenbasis

$$\Delta\Phi = \lambda\Phi$$

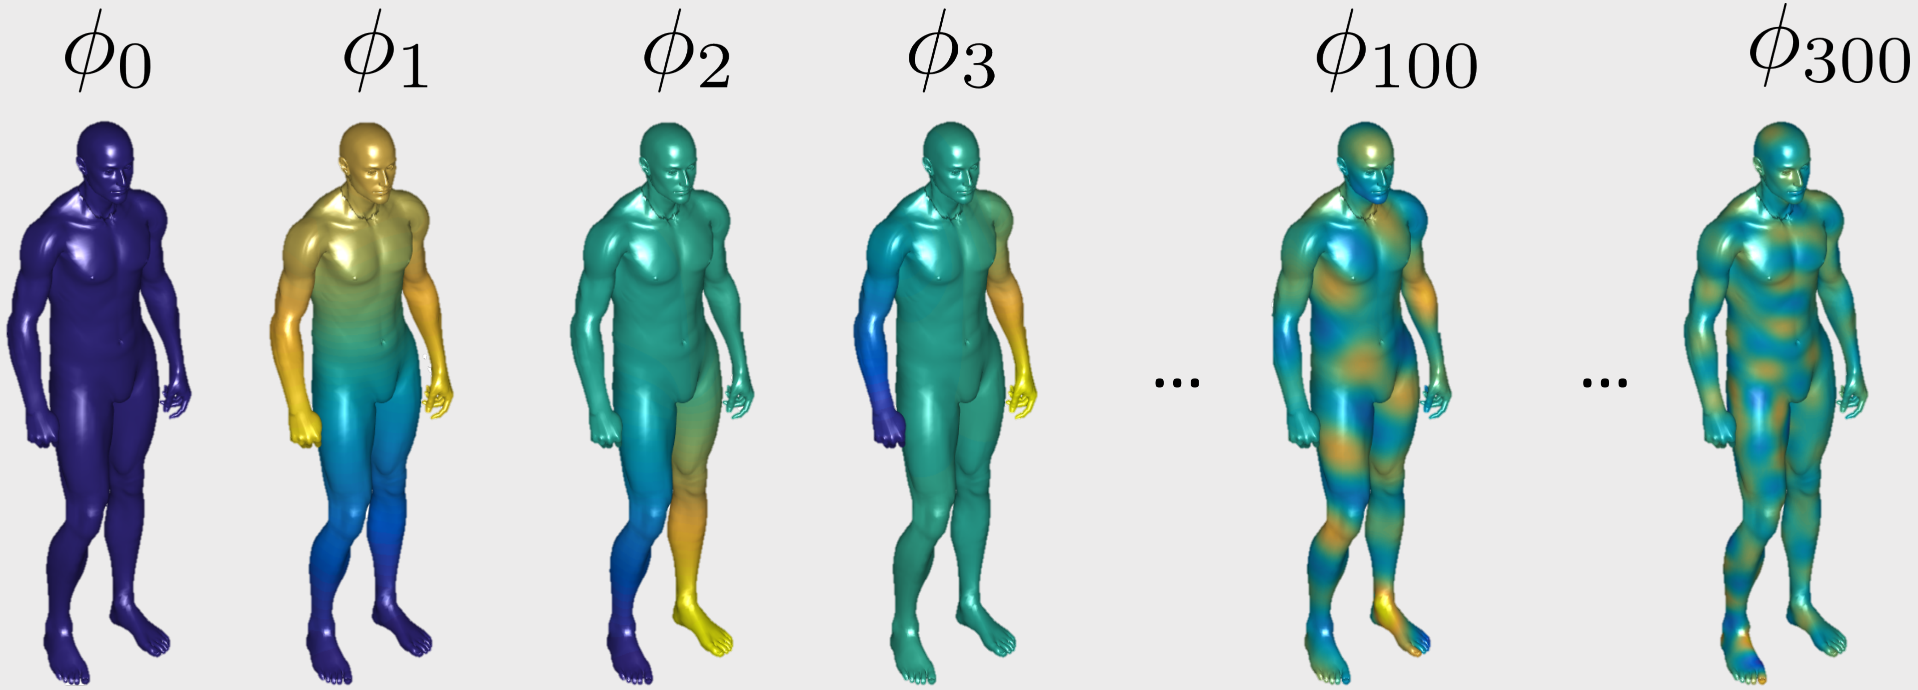

Distance Matching

- A common choice for pairwise descriptors are *geodesic distances*

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

Isometric deformations ⇔ invariant geodesic distances

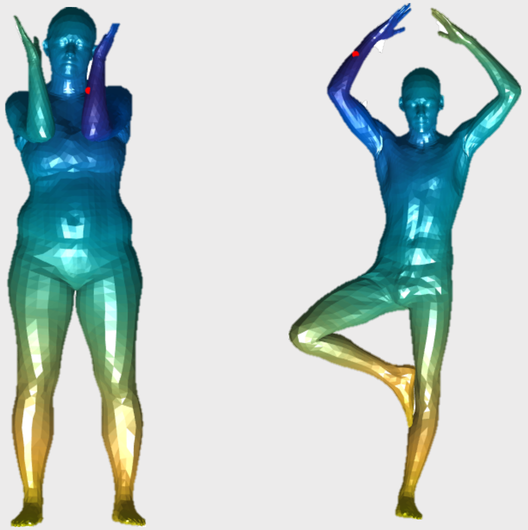

Cons:

- **Emphasize global similarity.**
- **Result in hard optimization problems (convex, smooth)**

BBK (2006), Chen, Q., Koltun (2015), Solomon, J., Peyré, G., Kim, V. G., & Sra, S. (2016).

Kernel Matching

- A good choice for pairwise descriptors are **positive definite kernels**

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

e.g. Heat Kernels

$$\mathbf{K}^t_{\mathcal{X}} = e^{-t\Delta} = \mathbf{\Phi} e^{-t\mathbf{\Lambda}} \mathbf{\Phi}^\top$$

$$\frac{du(t,x)}{dt} = \Delta_{\mathcal{X}} u(t,x), \;\; u(0,x) = u_0(x)$$

$$u(t,x) = \int_{\mathcal{X}} k(t,x,x') u_0(x') dx'$$

Shtern, Kimmel (2014) . Lähner, Z., Vestner, M., Boyarski, et. al (2017). Vestner, M., Litman, R., Rodolà, E., Bronstein, A., & Cremers, D. (2017).

Kernel Matching

$$\max_{\mathbf{\Pi} \in \mathcal{P}_n} \frac{1}{2} \langle \mathbf{\Pi}, \mathbf{K}_{\mathcal{Y}} \mathbf{\Pi} \mathbf{K}_{\mathcal{X}} \rangle = \max_{\pi \in \mathcal{P}_n} \pi^\top (\mathbf{K}_{\mathcal{X}} \otimes \mathbf{K}_{\mathcal{Y}}) \pi$$

$$\mathbf{K}^t_{\mathcal{X}} \otimes \mathbf{K}^t_{\mathcal{Y}} = e^{-t\Delta_{\mathcal{X}}} \otimes e^{-t\Delta_{\mathcal{Y}}}$$

$$= e^{-t(\Delta_{\mathcal{X}} + \Delta_{\mathcal{Y}})}$$

$$= e^{-t\Delta_{\mathcal{X} \times \mathcal{Y}}}$$

$$= \mathbf{K}^t_{\mathcal{X} \times \mathcal{Y}}$$

Heat kernel on product manifold

A kind of "Dirichlet energy" $\langle \mathbf{x}, \Delta \mathbf{x} \rangle$ in product space.

i.e., $(x, \pi^*(x))$ is a "smooth" sub-manifold in the product space.

Correspondence in product space

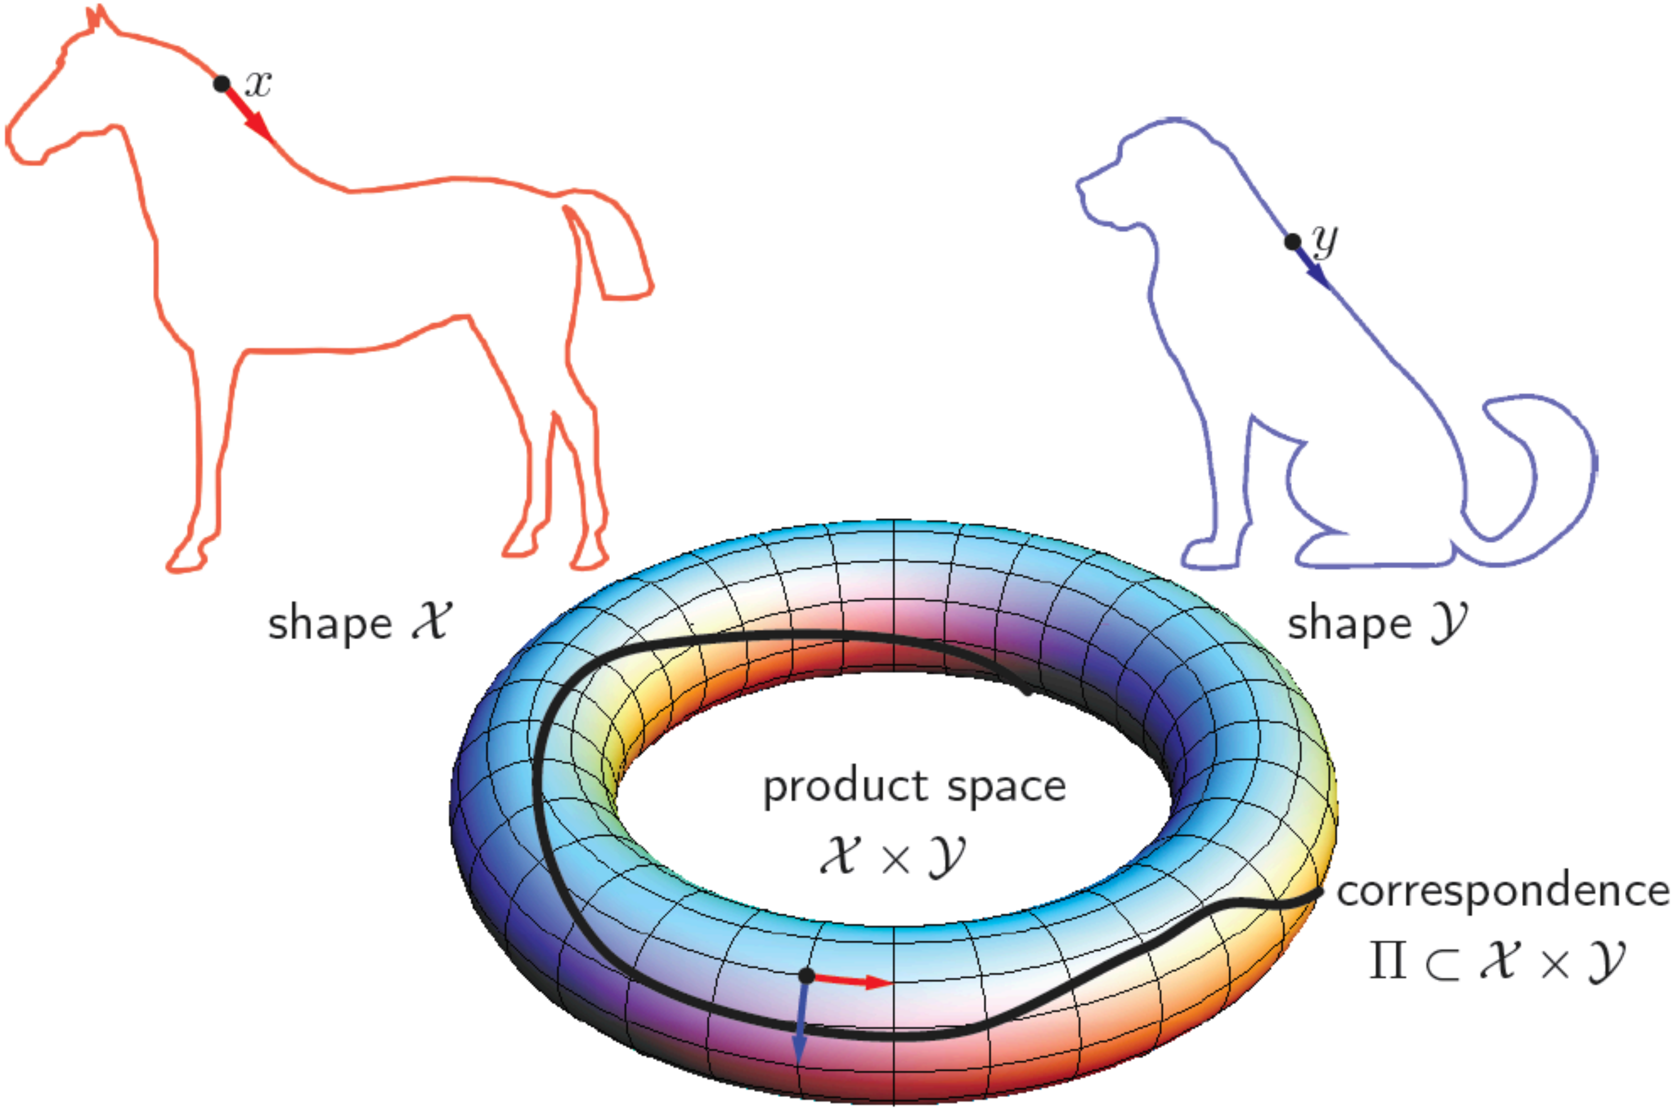

Vestner, Litman, Rodola, B, Cremers 2017

Kernel Matching - Algorithm

$$\max_{\mathbf{\Pi}\in\mathcal{P}_n} \langle \mathbf{\Pi}, \mathbf{K}_{\mathcal{Y}}\mathbf{\Pi}\mathbf{K}_{\mathcal{X}}\rangle \quad \Longleftrightarrow \quad \max_{\mathbf{P}\in\mathcal{B}_n} \langle \mathbf{P}, \mathbf{K}_{\mathcal{Y}}\mathbf{P}\mathbf{K}_{\mathcal{X}}\rangle$$

Exact relaxation!

(Strictly) convex objective

$$\mathbf{\Pi}^0 = \operatorname*{argmax}_{\mathbf{\Pi}\in\mathcal{P}_n} \langle \mathbf{\Pi}, \mathbf{F}_{\mathcal{Y}}\mathbf{F}_{\mathcal{X}}^\top\rangle$$

$$\mathbf{\Pi}^{k+1} = \operatorname*{argmax}_{\mathbf{\Pi}\in\mathcal{P}_n} \langle \mathbf{\Pi}, \mathbf{K}_{\mathcal{Y}}^{t_k}\mathbf{\Pi}^k\mathbf{K}_{\mathcal{X}}^{t_k}\rangle$$

$$t^{k+1} = \mu t^k, \quad 0 < \mu < 1$$

(Optional : Include data term.)

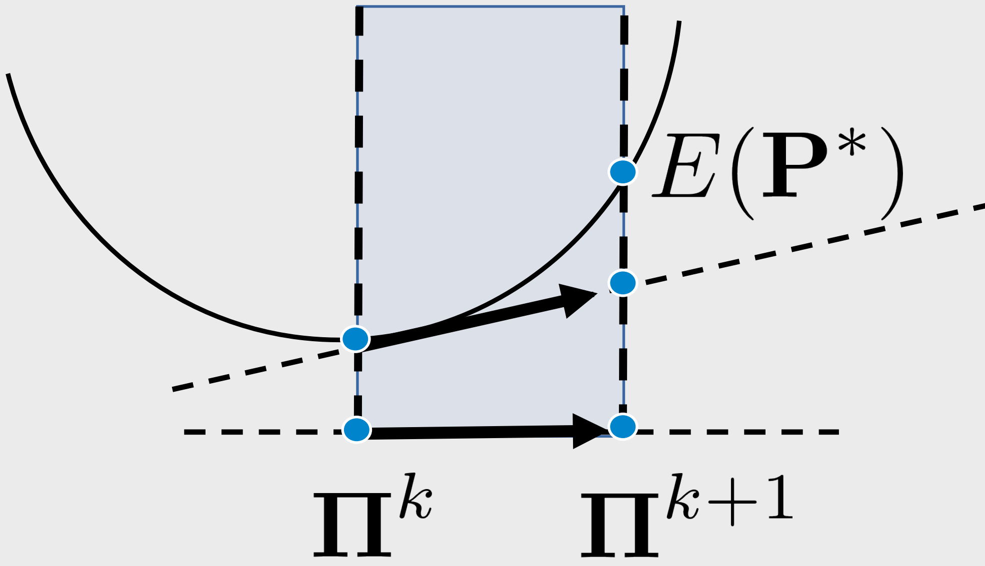

KM interpretations

$\mathcal{X}$ $\quad\quad \boldsymbol{\Pi}^* = Id \quad\quad$ $\mathcal{Y}$

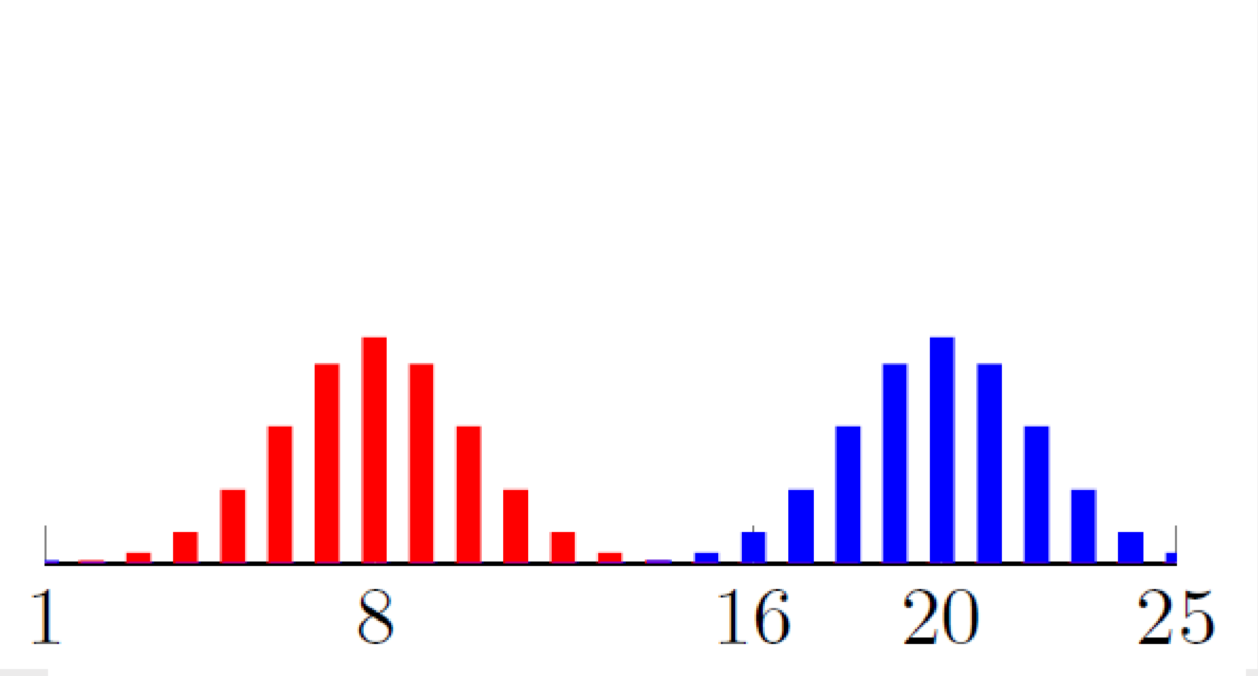

$$\underset{\boldsymbol{\Pi} \in \mathcal{P}_n}{\operatorname{argmax}} \langle \boldsymbol{\Pi}, \mathbf{K}_{\mathcal{Y}} \boldsymbol{\Pi}^k \mathbf{K}_{\mathcal{X}} \rangle$$

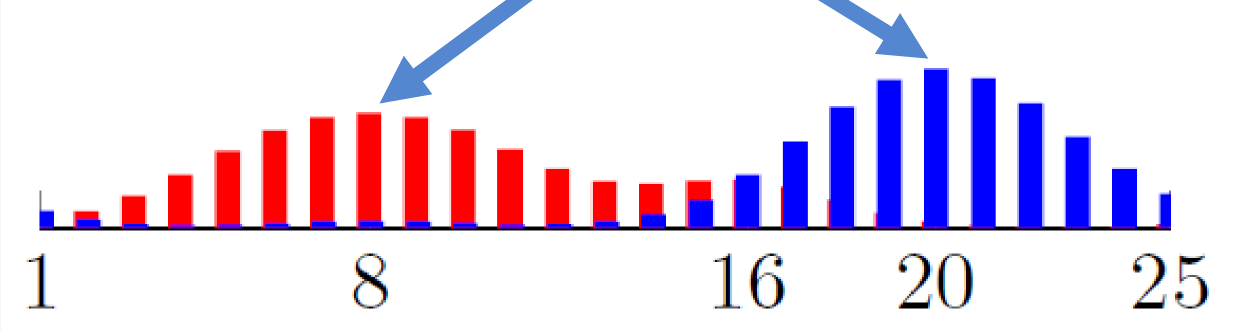

$\mathbf{K}_{\mathcal{X}} \delta$ $\quad \boldsymbol{\Pi}^k \quad$ $\mathbf{K}_{\mathcal{Y}} \boldsymbol{\Pi}^k \mathbf{K}_{\mathcal{X}} \delta$

Lederman, R. R., & Talmon, R. (2014).

KM interpretations

Spectral interpretation

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

Spectral rep. of correspondence

Low-pass filtering

KM interpretations

Spectral interpretation

$$\mathbf{\Pi}^{k+1} = \mathcal{P}_{\mathcal{P}_n}\left(\mathbf{K}_{\mathcal{Y}}\mathbf{\Pi}^k\mathbf{K}_{\mathcal{X}}\right)$$

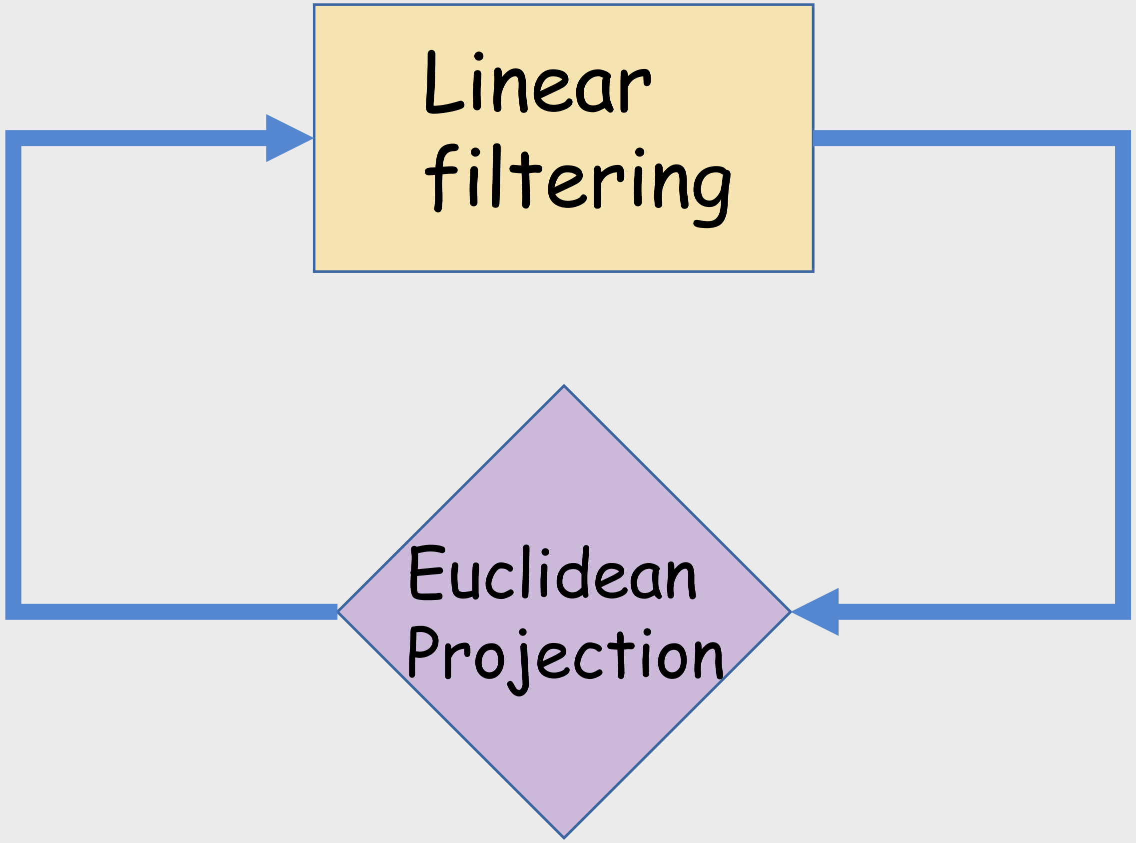

$$\mathbf{K}_{\mathcal{Y}}\mathbf{\Pi}^k\mathbf{K}_{\mathcal{X}}$$

Partiality

$$\mathbf{\Pi}\mathbf{1} \leq \mathbf{1}$$

$$\mathbf{\Pi}^\top\mathbf{1} = \mathbf{1}$$

$$\Pi_{ij} \in \{0, 1\}$$

Multiscale

Multiscale

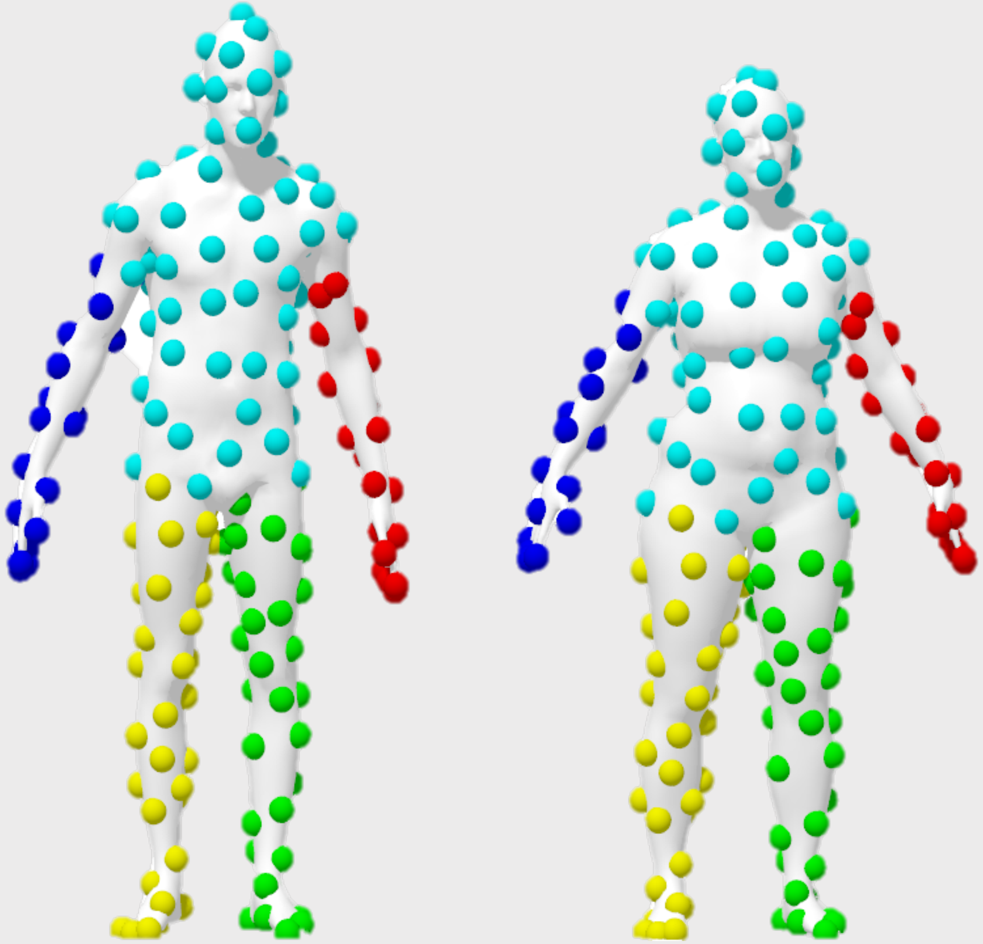

Multiscale

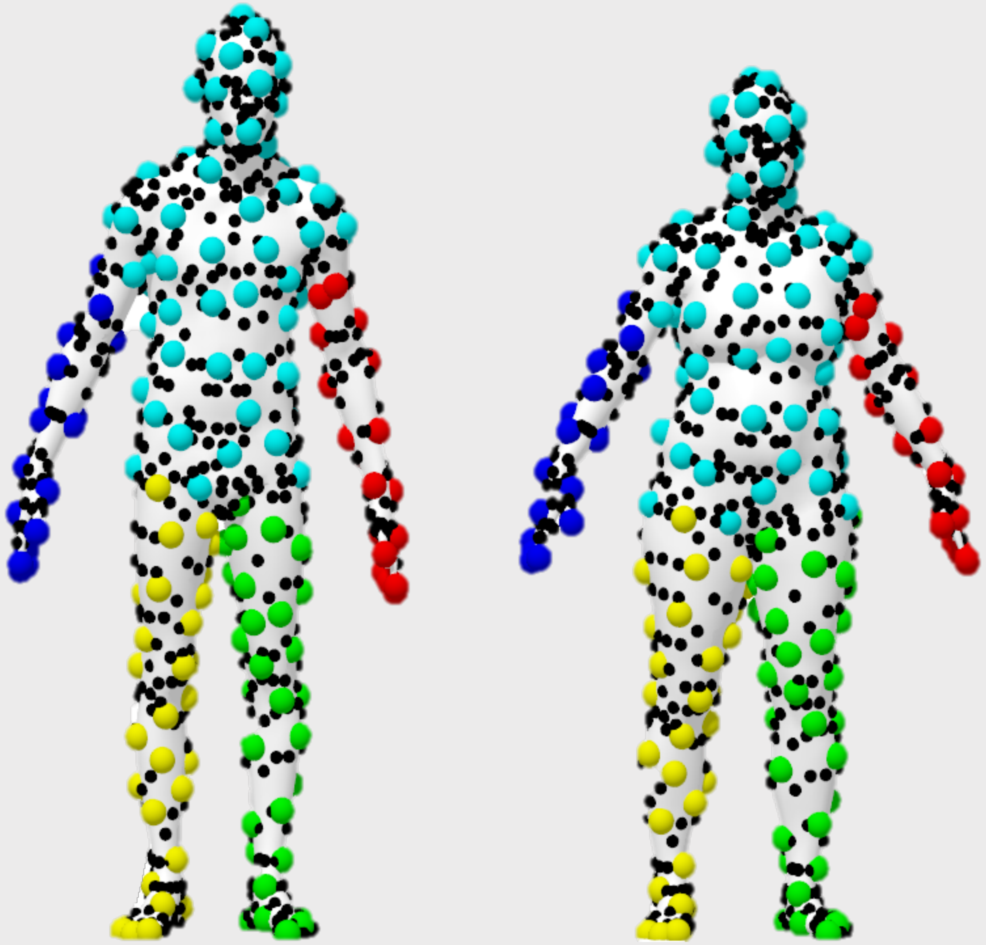

Multiscale

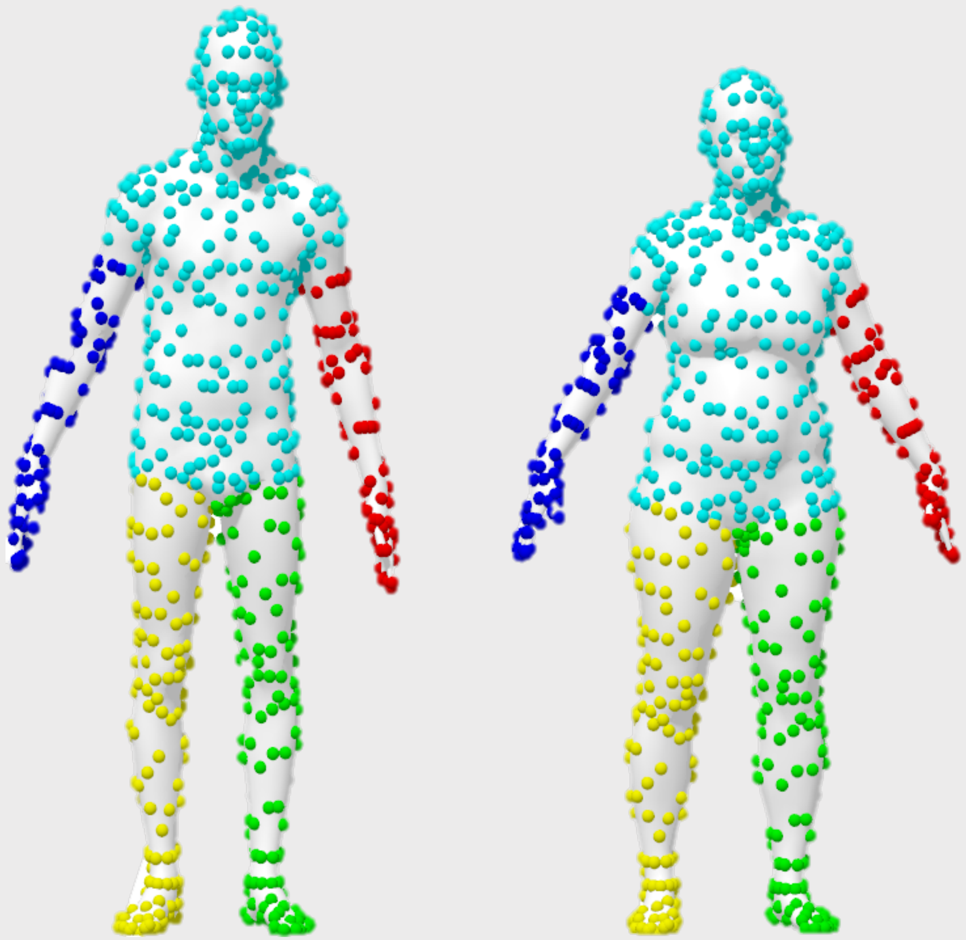

KM Results

SCAPE (*Anguelov* et al. 2005)

FAUST (Bogo et al. 2015)

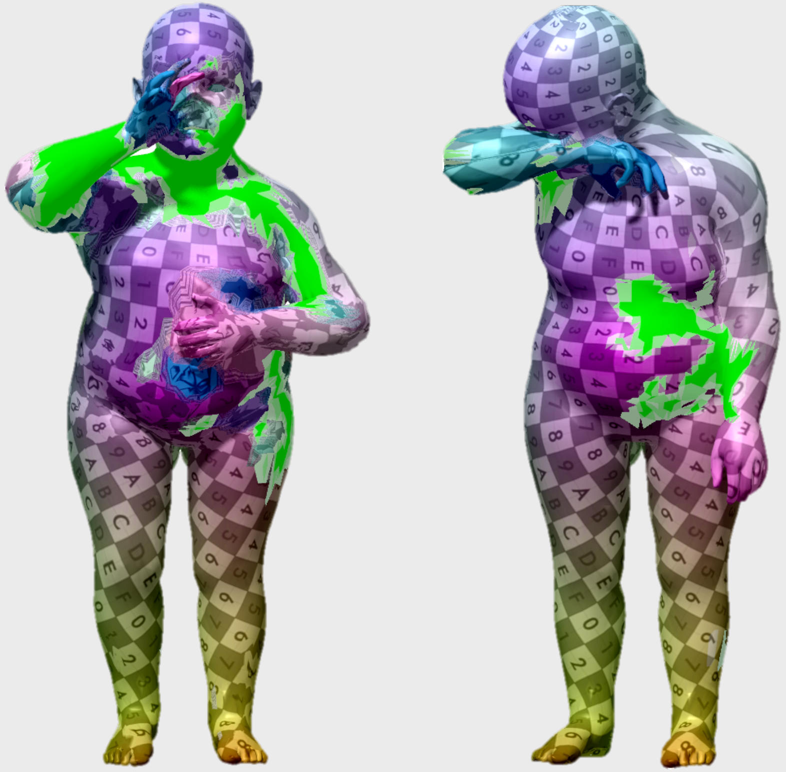

SHREC16 (*Cosmo* et al. 2016)

Non-isometric deformations

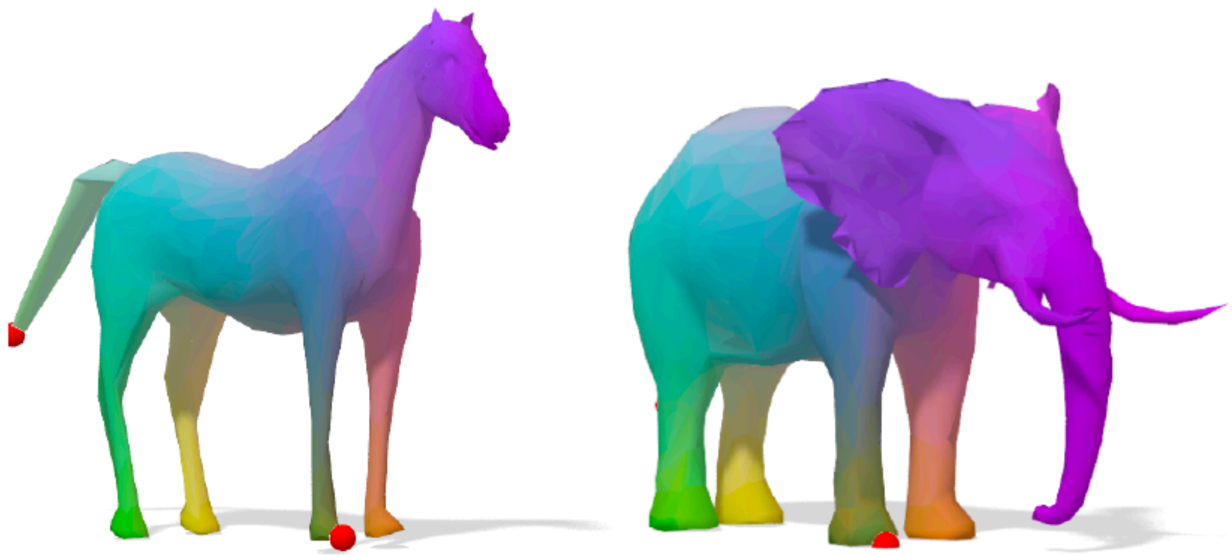

Input: 2 corresponding pairs

Smoothness

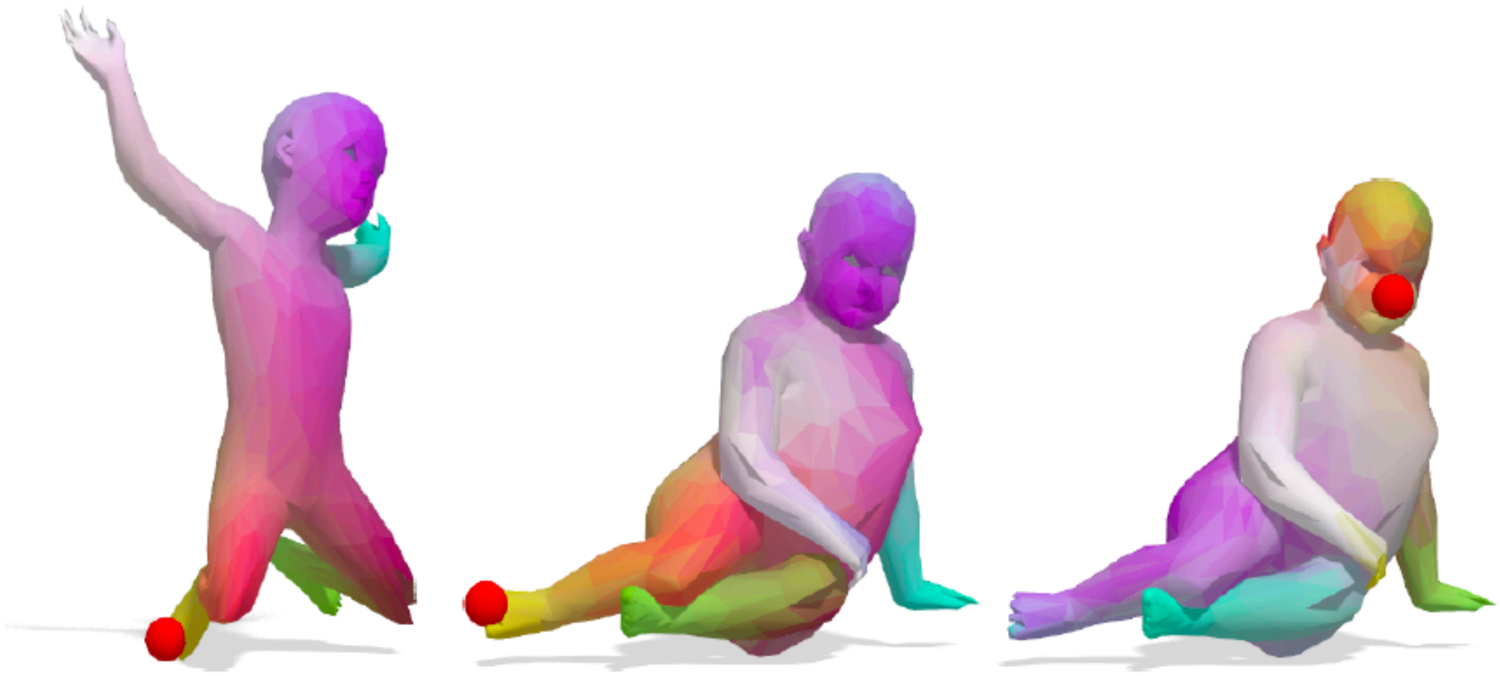

Input: 1 corresponding pair

2D to 3D matching

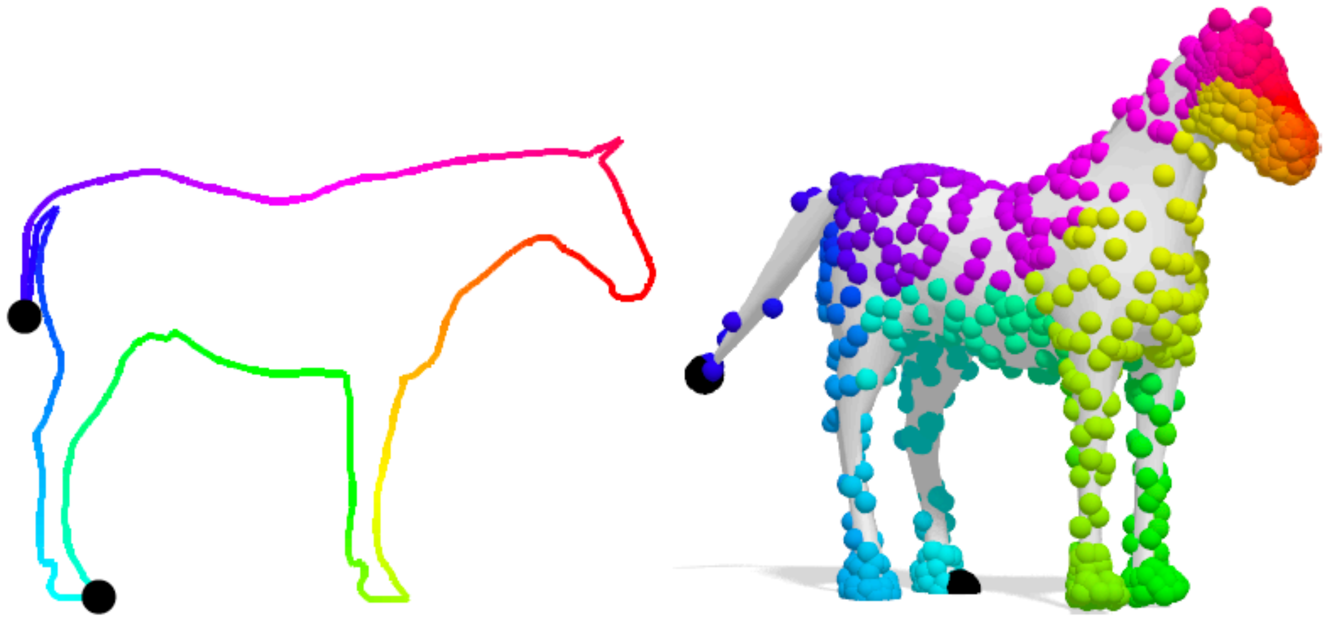

Input: 2 corresponding pairs

Discontinuous matching

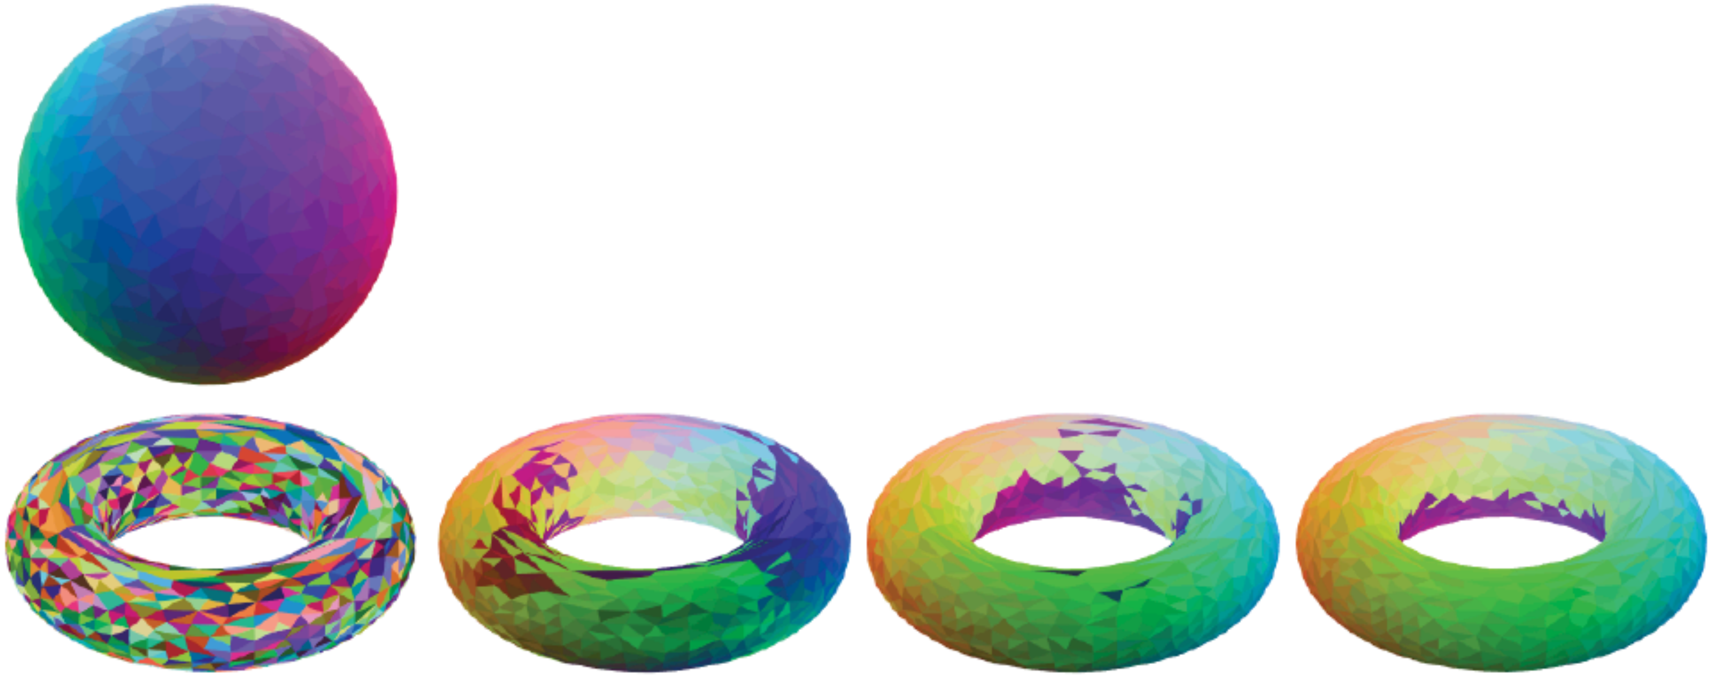

Random init | 1 iteration | 2 iterations | 3 iterations

Minimizes discontinuity length?

KM Results

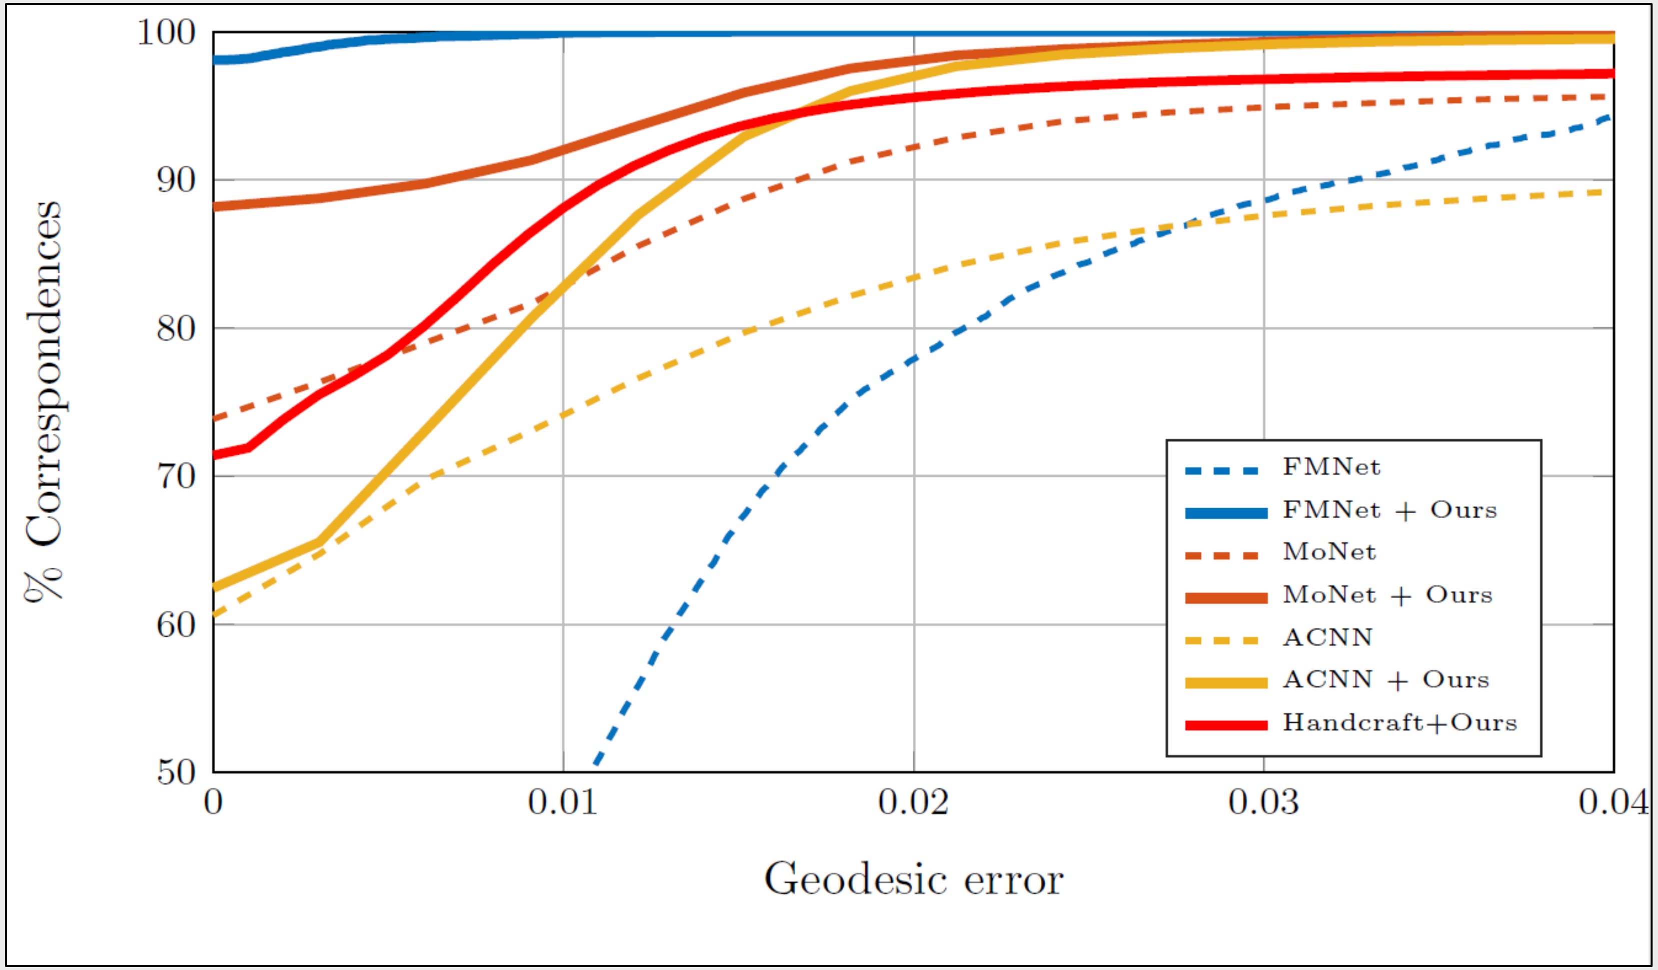

FAUST (Bogo et al. 2015). Evaluation: Kim et al. 2011.

KM Timings

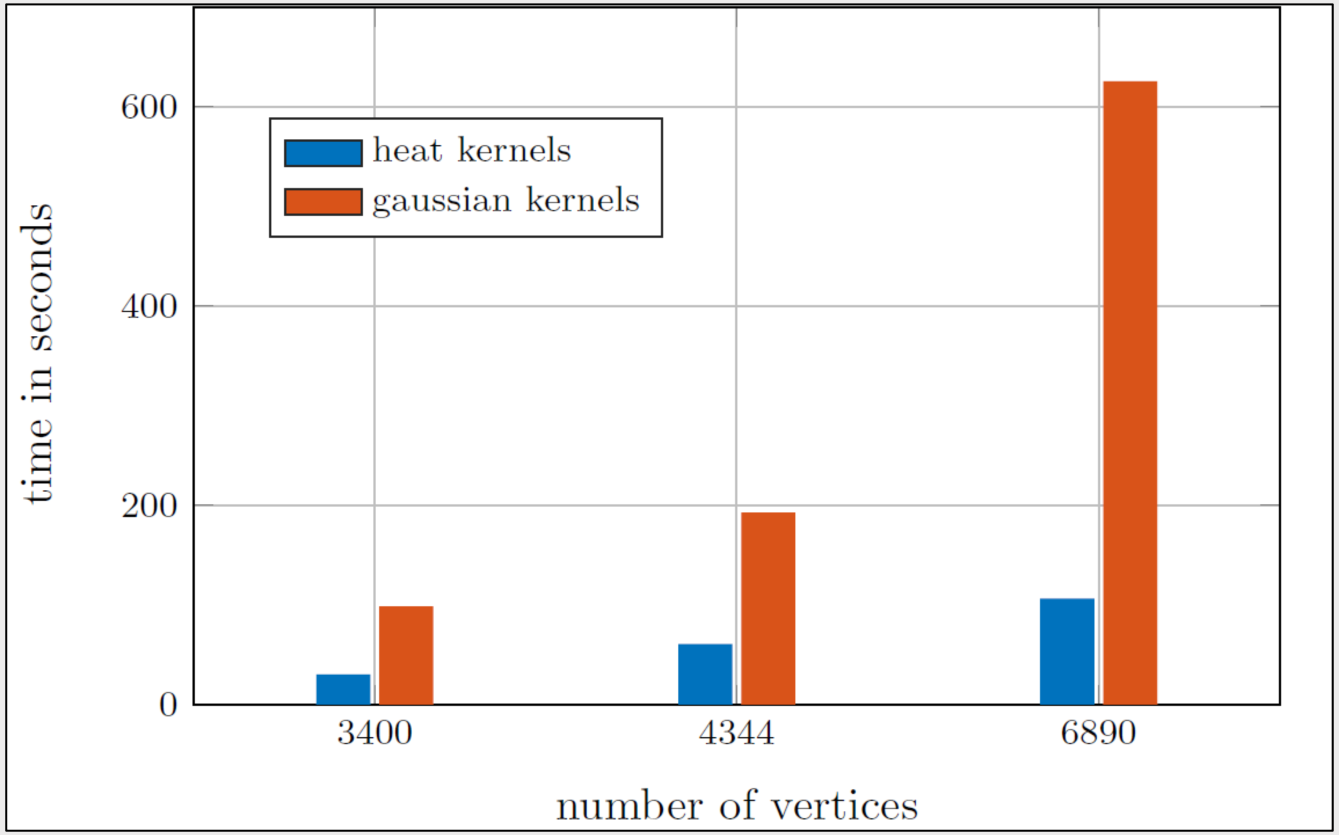

time in seconds
600
400
200
0
heat kernels
gaussian kernels
3400
4344
6890
number of vertices

Kernel Matching vs Distance Matching

$$\max_{\boldsymbol{\Pi} \in \mathcal{P}_n} \langle \boldsymbol{\Pi}, \mathbf{K}_{\mathcal{Y}} \boldsymbol{\Pi} \mathbf{K}_{\mathcal{X}} \rangle \quad \Longleftrightarrow \quad \max_{\mathbf{P} \in \mathcal{B}_n} \langle \mathbf{P}, \mathbf{K}_{\mathcal{Y}} \mathbf{P} \mathbf{K}_{\mathcal{X}} \rangle$$

Exact relaxation!

$$\max_{\boldsymbol{\Pi} \in \mathcal{P}_n} \langle \boldsymbol{\Pi}, \mathbf{D}_{\mathcal{Y}} \boldsymbol{\Pi} \mathbf{D}_{\mathcal{X}} \rangle \quad \Longleftrightarrow \quad \max_{\mathbf{P} \in \mathcal{B}_n} \langle \mathbf{P}, \mathbf{D}_{\mathcal{Y}} \mathbf{P} \mathbf{D}_{\mathcal{X}} \rangle$$

~~Exact relaxation!~~

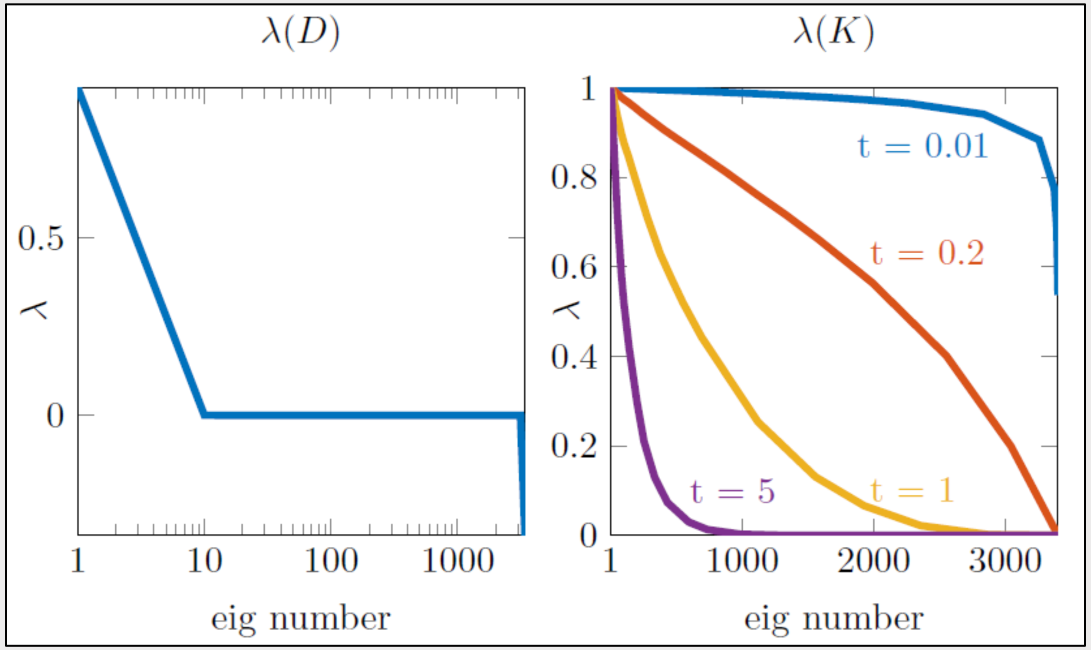

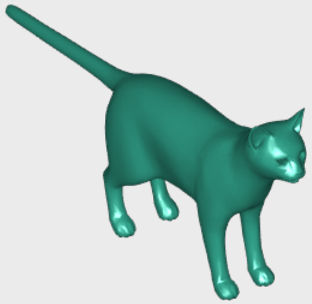

Distance matrices are Symmetric + zero trace.

↓

Never PD!

Kernel Matching ⇔ Distance Matching

$$\mathbf{P}\mathbf{1} = \mathbf{1}$$
$$\mathbf{P}^\top \mathbf{1} = \mathbf{1}$$
$$\mathcal{GB}_n$$

$$\mathbf{P} \in \mathcal{GB}_n \iff \mathbf{P} = \mathbf{P}_{\text{cent}} + \mathbf{C}_n \mathbf{W} \mathbf{C}_n$$

$$\mathbf{P}_{\text{cent}} = \frac{1}{n}\mathbf{1}\mathbf{1}^\top$$

Centroid of $\mathcal{B}_n$

$$\mathbf{C}_n = \mathbf{I} - \frac{1}{n}\mathbf{1}\mathbf{1}^\top$$

Centering matrix

Double centered kernel: $\tilde{\mathbf{K}} \equiv \mathbf{C}_n \mathbf{K} \mathbf{C}_n$

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

Kernel Matching ⇔ Distance Matching

$$\mathbf{I} - \frac{1}{n}\mathbf{1}\mathbf{1}^\top$$

$$\tilde{\mathbf{K}} = -\frac{1}{2}\mathbf{C}_n \mathbf{D} \mathbf{C}_n$$ Is a double centered PD kernel!

Matrix of **squared** Euclidean distances

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

Upcoming: learnable kernels

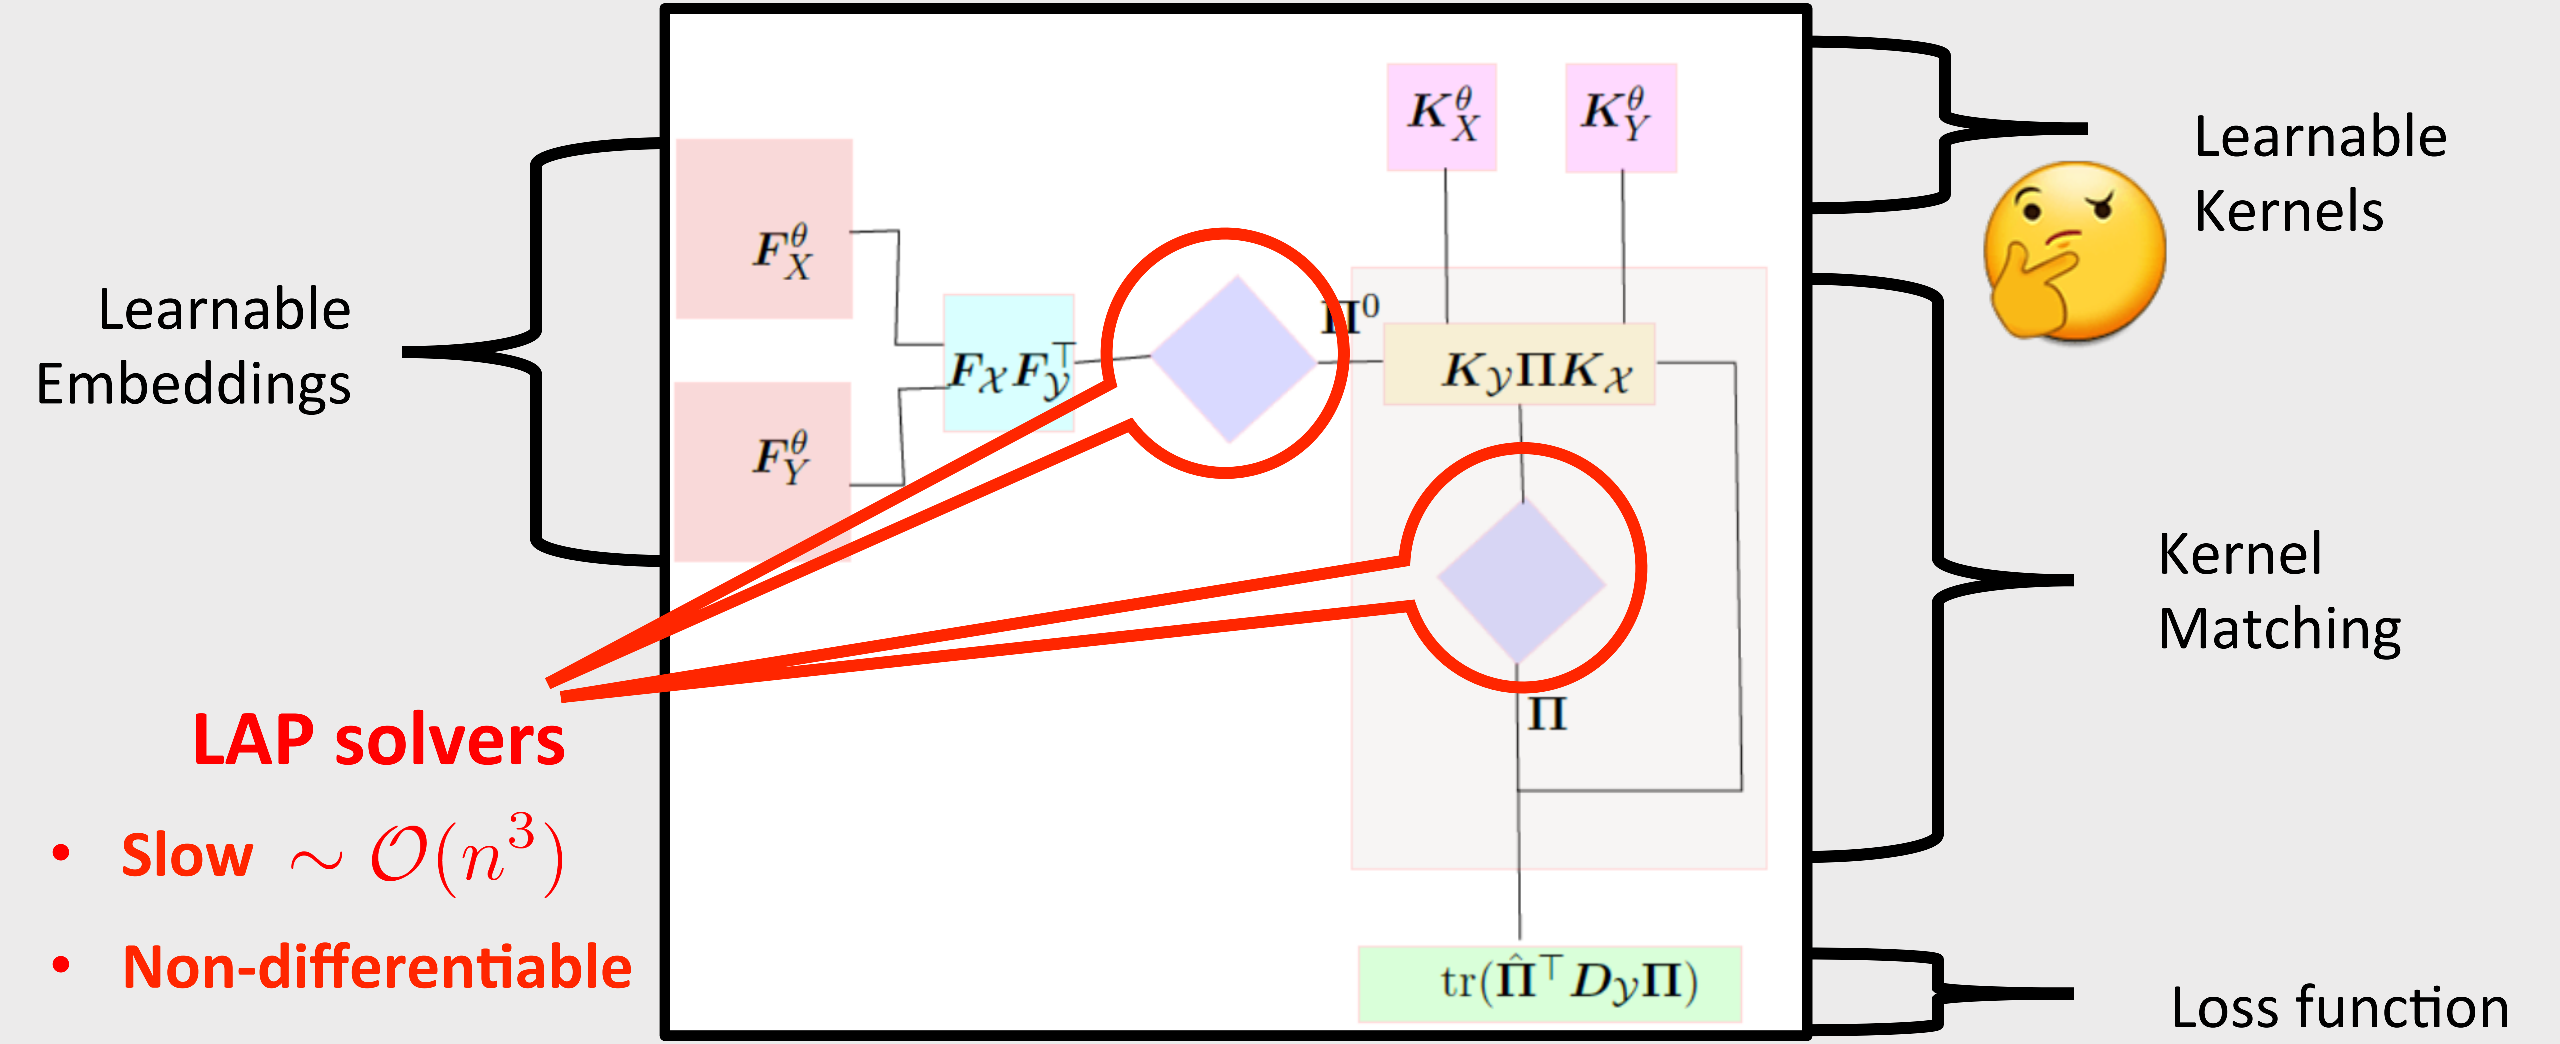

LAP solvers

- **Slow** $\sim \mathcal{O}(n^3)$
- **Non-differentiable**

Joint work with …

Matthias Vestner, Zorah Lähner, Daniel Cremers

Or Litany, Tal Remez

Ron Slossberg, Ron Kimmel, Alex Bronstein

Michael Bronstein

Emanuele Rodolà

Code : https://github.com/zorah/KernelMatching