Analysis of 3D Shapes (IN2238)

Frank R. Schmidt

Matthias Vestner

Summer Semester 2017

Heat kernel signature

We define the **heat kernel signature** at a point $x \in S$ as the vector

$$HKS(x) = (k_{t_1}(x,x), \ldots, k_{t_T}(x,x)) \in \mathbb{R}^T$$
$$k_t(x,x) = \sum_{k=0}^{\infty} e^{\lambda_k t} \phi_k^2(x)$$

In this view, each evaluation of the heat kernel in the vector above describes **the amount of heat staying at point** x after time t, when starting with a unit heat source (dirac) at x itself.

The HKS also has an informative property. If the eigenvalues of the Laplacians on S_1 and S_2 are not repeated, then:

$$\Phi : S_1 \to S_2 \text{ is an isometry iff } k_t^{S_1}(x,x) = k_t^{S_2}(\Phi(x), \Phi(x))$$

Metric spaces

Let M be a set. The tupel set (M, d_M), $d_M : M \times M \to \mathbb{R}_{\geq 0}$ is a metric space if

- identity of indiscernibles: $d_M(x, y) = 0 \Leftrightarrow x = y$
- symmetry: $d_M(x, y) = d_M(y, x)$
- triangle inequality: $d_M(x, y) \leq d_M(x, z) + d_M(z, y)$ for all $x, y, z \in M$

Satisfying a subset of these properties leads to the definition of "semi"-metric spaces, "pseudo"-metric spaces, etc.

Distance to set, diameter

The distance from a point x to a set S in a metric space X is defined by

$$\mathrm{dist}_X(x, S) = \inf_{y \in S} d_X(x, y)$$

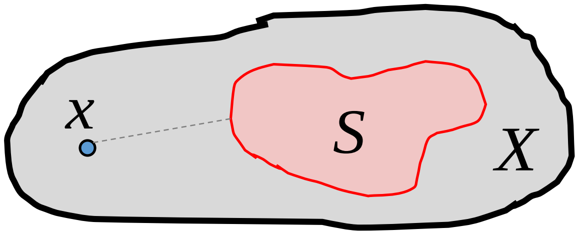

The diameter of a set S in a metric space X is defined by

$$\mathrm{diam}(S) = \sup_{x,y \in S} d_X(x, y)$$

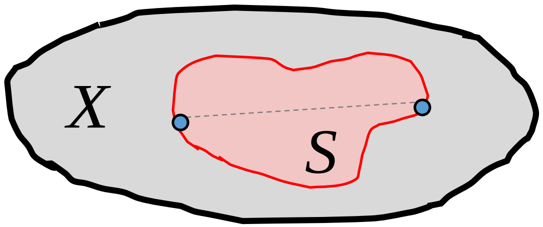

Hausdorff distance

The **Hausdorff distance** between two compact subsets $X, Y \subset (Z, d_Z)$ is defined by

$$d_H^Z(X,Y) = \max\{\sup_{x \in X} \mathrm{dist}_Z(x,Y), \sup_{y \in Y} \mathrm{dist}_Z(y,X)\}$$

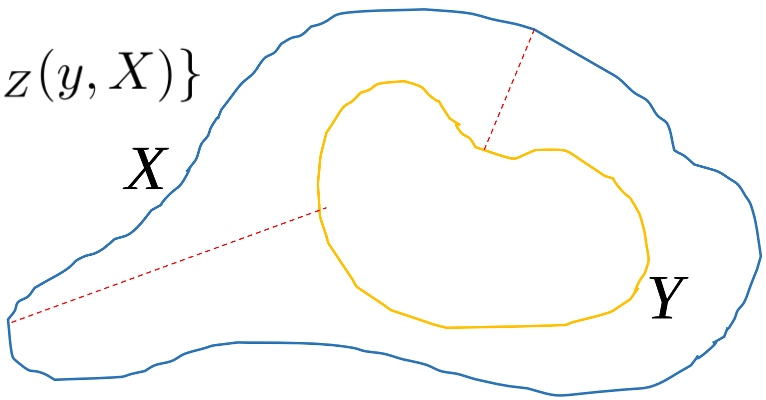

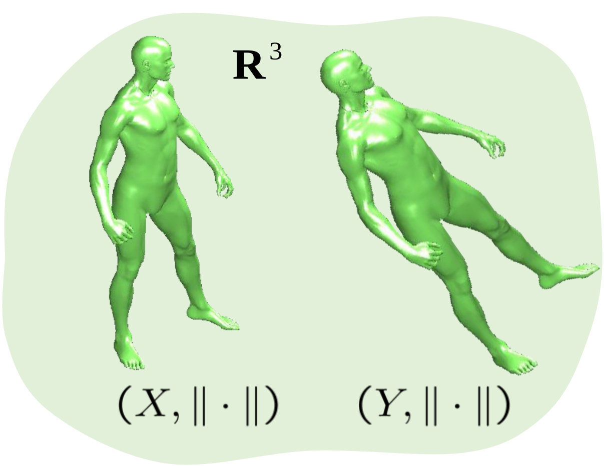

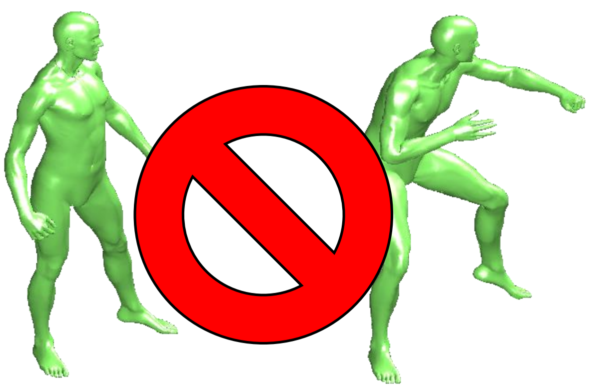

(X, d_X) (Y, d_Y)

Gromov Hausdorff distance

Can we define a Hausdorff distance between metric spaces?

The general idea is to embed the two metric spaces (X, d_X) and (Y, d_Y) into a new metric space (Z, d_Z) and compute the Hausdorff distance in the resulting embeddings.

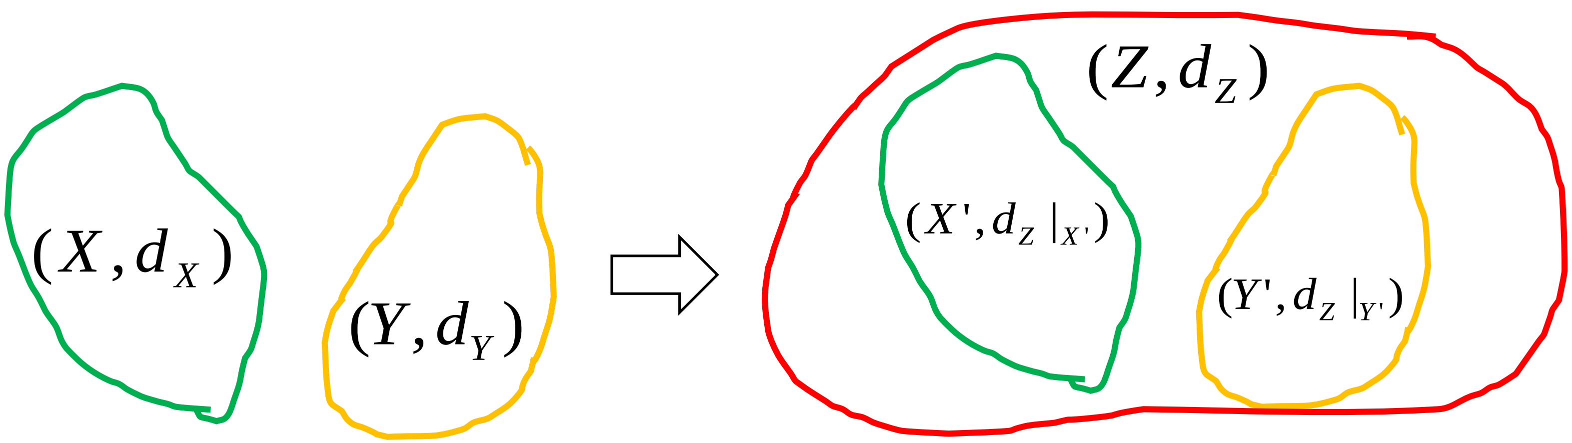

Further we define $d_{GH}(X,Y) < r$ if and only if there exists a metric space (Z, d_Z) and subspaces $X', Y' \subset Z$ which are isometric to X and Y such that $d_H^Z(X', Y') < r$.

Gromov Hausdorff distance

The **Gromov Hausdorff distance** between two metric spaces $(X, d_X), (Y, d_Y)$ is defined by

$$d_{\mathcal{GH}}(X,Y) = \inf_{Z,f,g} d_{\mathcal{H}}^{Z}(f(X), g(Y))$$

The infimum is taken over **all** ambient spaces Z and isometric embeddings $f : X \to Z, g : Y \to Z$.

The Gromov Hausdorff distance is a metric on the space of equivalence classes of metric spaces.

$X \equiv Y$ iff X and Y are isometric.

Fixed embedding space

The idea is closely related to multidimensional scaling (MDS). There however the metric space $Z = \mathbb{R}^k$ is fixed (and euclidean).

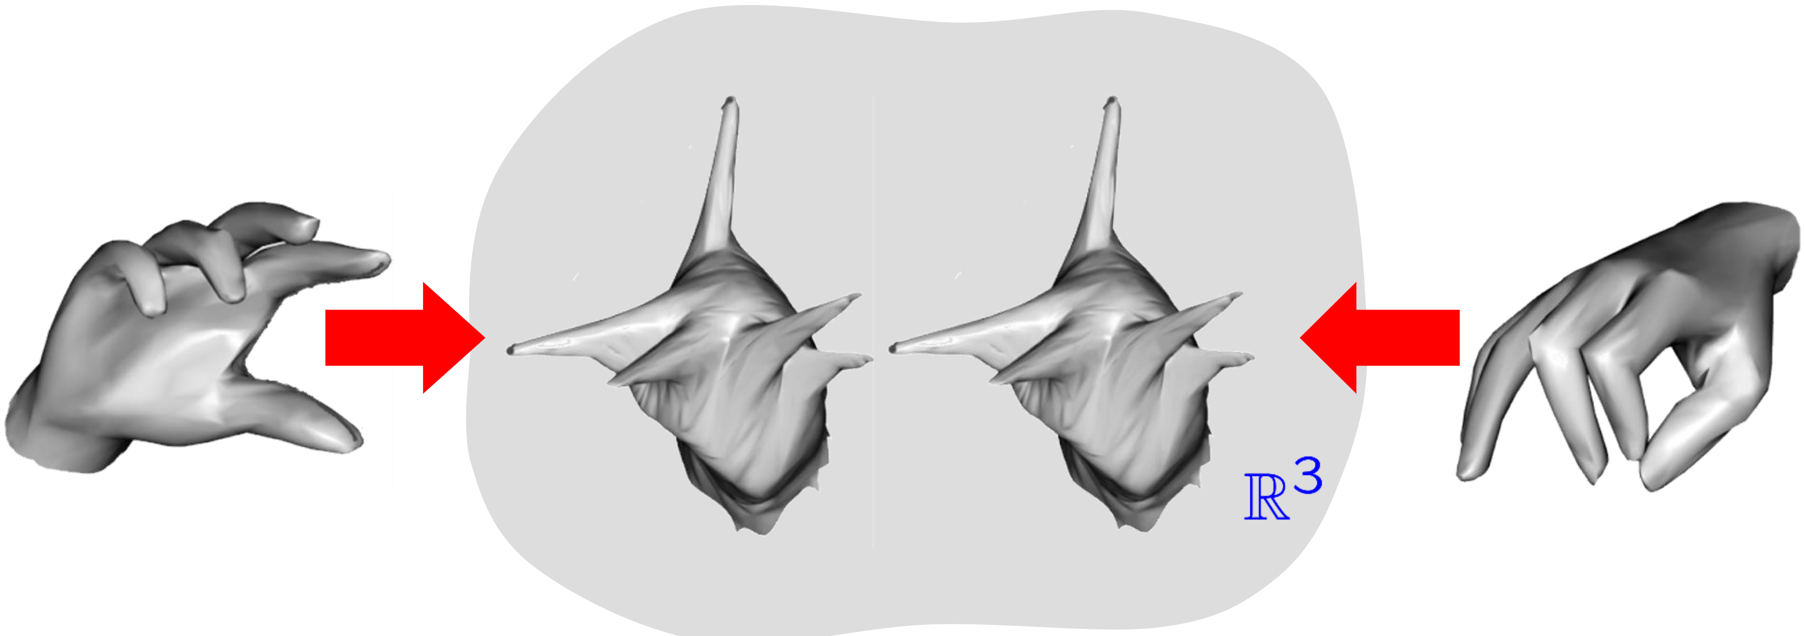

Coverings

Let $x \in X$. An open ball of radius $r > 0$ centered at x is defined by

$$B_r(x) = \{z \in X : d_X(x, z) < r\}$$

For a subset $A \subset X$, we define

$$B_r(A) = \cup_{a \in A} B_r(a)$$

A set $C \subset X$ is an **r-covering** of X if $B_r(C) = X$.

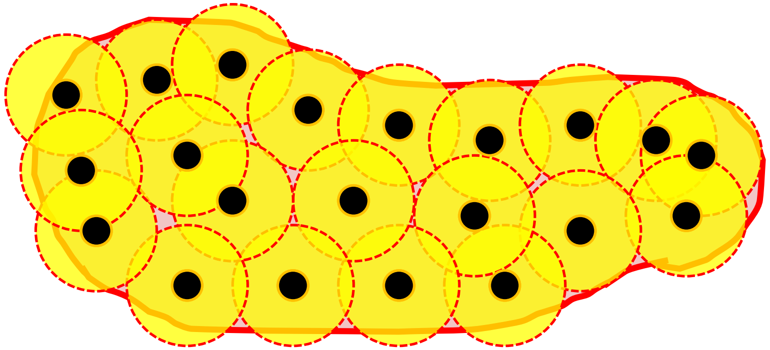

Let $\{x_1, \ldots, x_n\} \subset X$ be a r-covering of the compact metric space (X, d_X). Then

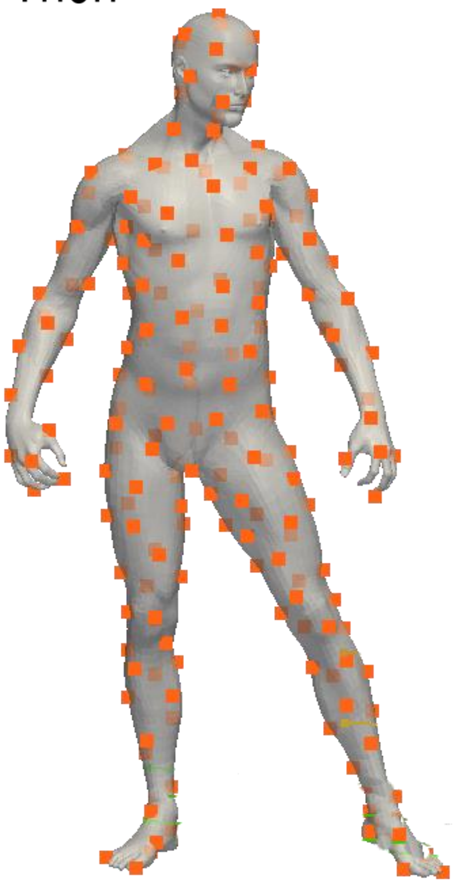

$$d_{GH}(X, \{x_1, \ldots, x_n\}) \leq r$$

This tells us that "shape samplings" are close to the underlying shapes in the Gromov-Hausdorff sense.

Consistency to sampling

Let $\{x_i\}_{i=1}^{n}$ be a r-covering of X and $\{y_i\}_{i=1}^{n'}$ be a r-covering of Y. Then

$$\left| d_{\mathcal{GH}}(X,Y) - d_{\mathcal{GH}}(\{x_i\}_{i=1}^{m}, \{y_j\}_{j=1}^{m'}) \right| \leq r + r'$$

This means d_{GH} is **consistent to sampling.**

If we have a way to compute d_{GH} for dense enough (small r) samplings of X and Y, then it would give us a good approximation to what happens in the continuous spaces.

Optimal coverings

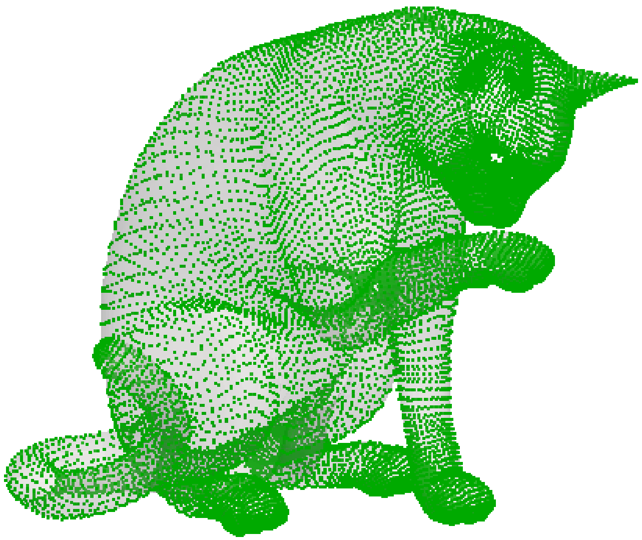

Can we devise an optimal sampling scheme in a metric sense?

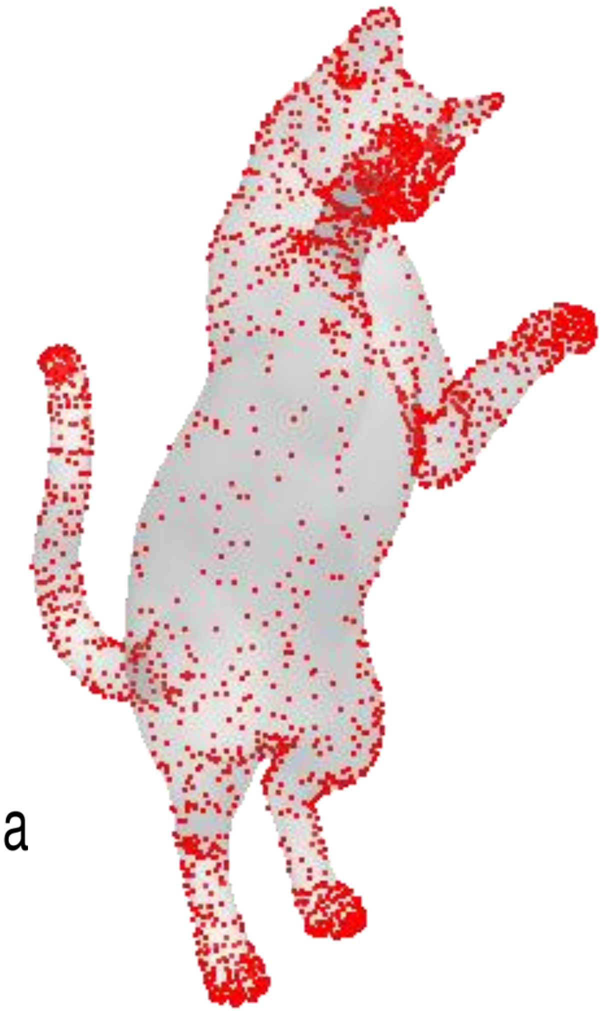

Farthest point sampling

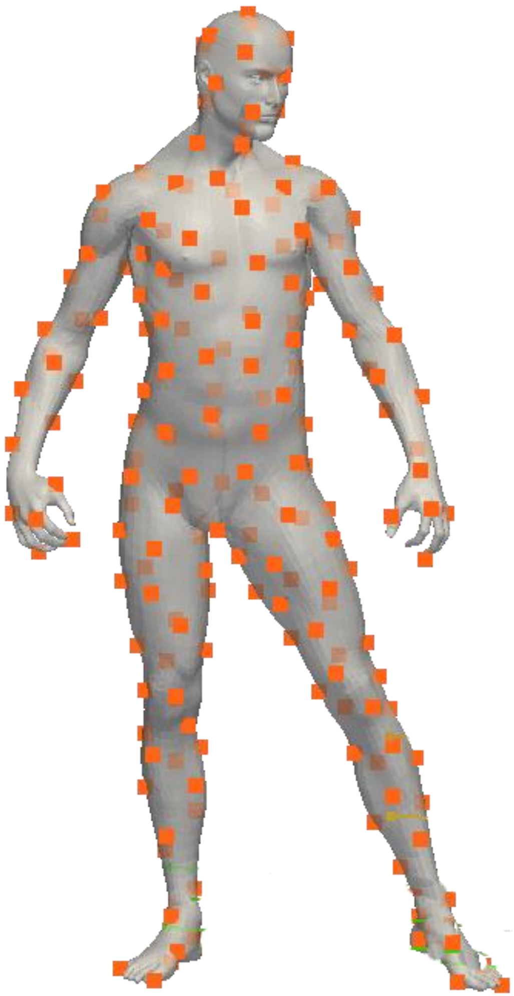

Fix n the number of points we want to have in our final covering X_n.

Intitialize $X_1 = \{p_1\}$

For $k = 2 : n$

$$p = \operatorname{argmax}\, d(x, X_{k-1})$$

$$X_k = X_{k-1} \cup \{p\}$$

end

Non-uniqueness due to

- choice of starting point p_1
- non-unique maximizer in iterations

Voroni cells

Each sampling $\{x_i\}$ of a shape X induces a set of regions $\{V_i\}$

$$V_i(X) = \{x \in X : d_X(x, x_i) < d_X(x, x_j) \forall i \neq j\}$$

These regions are known as **Voronoi regions or Voronoi cells.**

Each point x_i from the sampling can be seen as a representative for its Voronoi region.

Nearest neighbor search corresponds to identification of Voronoi cell $\Rightarrow$ connection to kd-trees.

Optimal sampling

The optimal sampling (with n samples) is the one minimizing the **maximum cluster radius:**

$$\varepsilon_\infty(\{x_i\}) = \max_i \max_{x \in V_i} d_x(x, x_i)$$

Optimal sampling is **NP hard** to compute.

However: FPS is "almost" optimal in the sense

$$\varepsilon_\infty(\{x_i^{fps}\}) \leq 2 \min_{\{x_i\}} \max_i \max_{x \in V_i} d_x(x, x_i)$$

Farthest point sampling

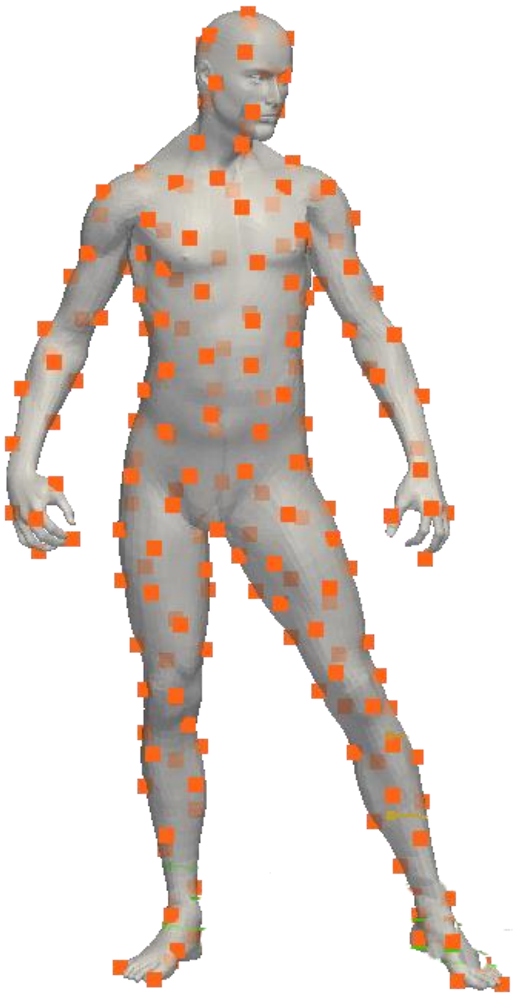

Final samling has progressively increasing density.

It is efficient to compute.

It is worse than optimal sampling by at most a factor of 2.

Correspondence

A **correspondence** between two sets X and Y is a subset of the product space $R \subset X \times Y$ satisfying

- for every $x \in X$ there exists at least one $y \in Y$ such that $(x, y) \in R$
- for every $y \in Y$ there exists at least one $x \in X$ such that $(x, y) \in R$

Any surjective map $f : X \to Y$ defines a correspondence:

$$R = \{(x, f(x), x \in X)\}$$

However not every correspondence is associated with a map.

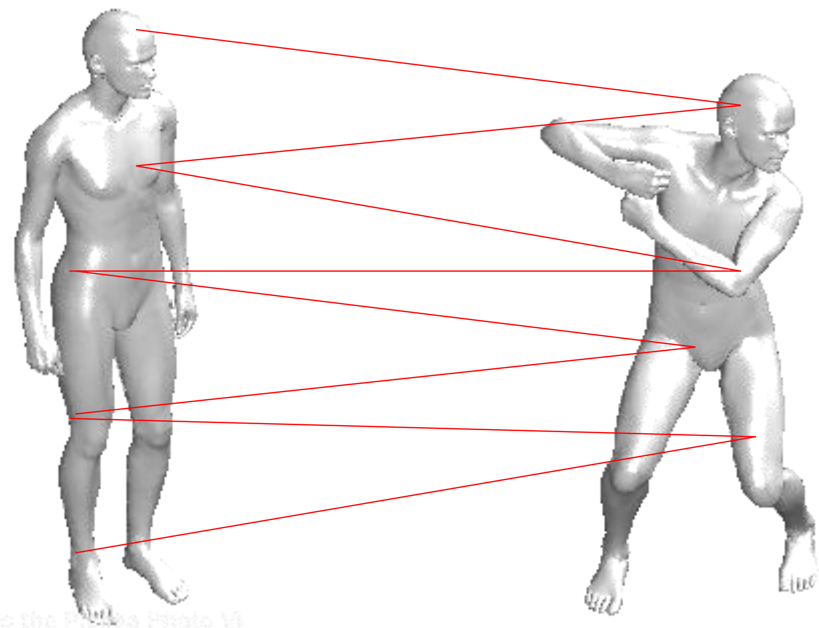

Metric distortion

The **distortion** of a correspondence $R \subset X \times Y$ is defined by

$$\mathrm{dis}(R) = \sup\{|d_X(x,x') - d_Y(y,y')| : (x,y),(x',y') \in R\}$$

Key observation:

$\mathrm{dis}(R) = 0$ if and only if R is associated with an isometry.

We say that R is an ε-isometry if $\mathrm{dis}\, R \leq \varepsilon$.

Correspondence and Gromov Hausdorff

$d_{GH}(X,Y) < r \quad \Leftrightarrow$ There exists a correspondence R such that $|d_X(x,x') - d_Y(y,y')| < 2r$ for all pairs $(x,y),(x',y') \in R$ of correspondence elements.

This allows us to speak about d_{GH} just by using correspondences R:

$$d_{GH}(X,Y) = \frac{1}{2} \inf_R \operatorname{dis} R$$

Intuition: Choose as embedding space (Z, d_Z) one of the metric spaces (X, d_X),(Y, d_Y).

A computational approach

We want to compute a correspondence $R \subset X \times Y$ minimizing

$$d_{GH}(X,Y) = \frac{1}{2}\inf_R \operatorname{dis} R$$

Let us rewrite

$$\begin{aligned} d_{GH}(X,Y) &= \frac{1}{2}\inf_R \operatorname{dis} R \\ &= \frac{1}{2}\inf_R \sup\{|d_X(x,x') - d_Y(y,y')| : (x,y),(x',y') \in R\} \\ &\left(= \frac{1}{2}\inf_{f:X\to Y} \sup_{x,x'} |d_X(x,x') - d_Y(f(x),f(x'))|\right) \end{aligned}$$

The last equality assumes that the optimal R is associated with a surjective map f.

A computational approach

For two coverings $\{x_i\}_{i=1}^n$ and $\{y_i\}_{i=1}^n$ (with sampling radii r and r') we can define a related distance

$$d_P(\{x_i\},\{y_i\}) = \frac{1}{2}\min_{\pi\in P_n}\max_{1\leq i,j\leq n}\left|d_X(x_i,x_j) - d_Y(y_{\pi(i)},y_{\pi(j)}\right|$$

where P_n denotes the set of all permutations of $\{1,\ldots n\}$.

From the bounds we have for r-coverings it can be shown that

$$d_{GH}(X,Y) \leq r + r' + d_P(\{x_i\},\{y_i\})$$

Discretization

A correspondence can be represented by a matrix $R \in \{0,1\}^{n \times n}$

X

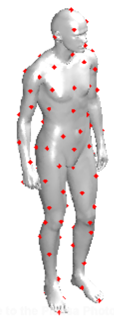

$$
\begin{array}{c|ccccc|c}
 & & & & & & \\
\hline
 & 0 & 1 & 0 & 0 & 0 & \\
 & 0 & 0 & 0 & 1 & 0 & \\
 & 1 & 0 & 0 & 0 & 0 & \\
 & 0 & 0 & 0 & 0 & 1 & x_i \in X \\
 & 0 & 0 & 1 & 0 & 0 & \\
\hline
 & & & & & y_j \in Y &
\end{array}
$$

Y

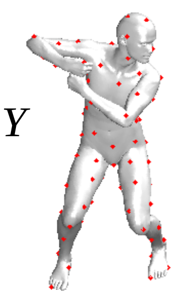

$R_{ij} = 1$ if x_i and y_j are in correspondence.

Asking for a bijection corrersponds to require R to be a permutation matrix.

Discretization

The metric distortion terms can be incorporated into a cost matrix $C \in \mathbb{R}^{n^2 \times n^2}$

$$C_{(il)(jm)} = \left| d_X(x_i, x_j) - d_Y(y_l, y_m) \right|$$

	(x_1, y_1)	(x_1, y_2)	(x_1, y_3)	$\cdots$		$\ddots$
(x_1, y_1)	0	13.5	23.4	104.6	7.64	
(x_1, y_2)	13.5	0	13.52	11.2	71.1	
(x_1, y_3)	23.4	13.52	0	0.22	23.44	
$\vdots$	104.6	11.2	0.22	0	17.5	
	7.64	71.1	23.44	17.5	0	

With this notation we can write the distance as

$$d_P(\{x_i\}, \{y_i\}) = \frac{1}{2} \min_R \max_{i,j,l,m} C_{(il)(jm)} R_{il} R_{jm}$$

where R is in the space of permutation matrices of size n.

Sensitivity to outliers

$$\frac{1}{2}\min_R \max_{i,j,l,m} C_{(il)(jm)} R_{ij} R_{lm}$$

$$\Downarrow$$

$$\frac{1}{2}\min_R \sum_{i,j,l,m} C_{(il)(jm)} R_{ij} R_{lm}$$

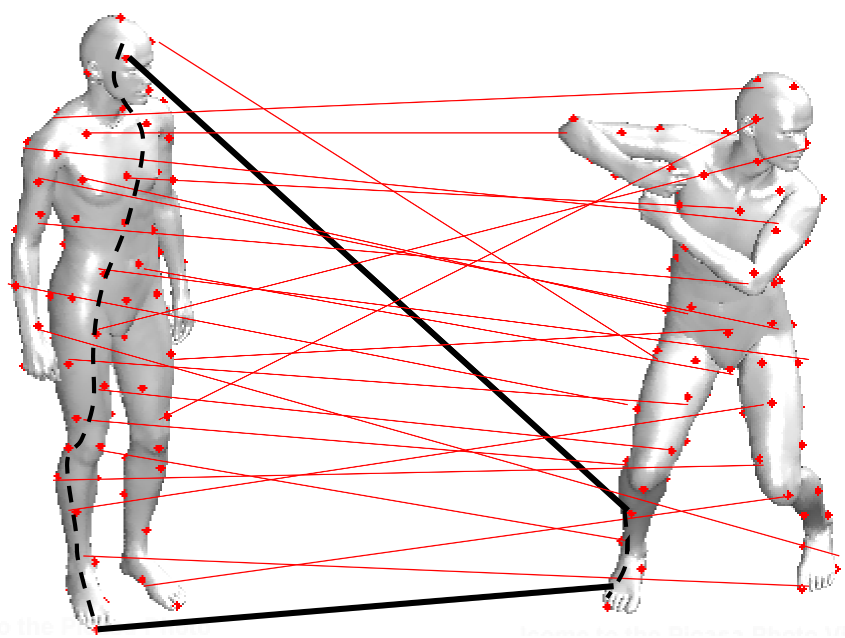

Gromov Hausdorff relaxed

We obtain a family of related problems by relaxing the max to a sum. Fix $p \geq 1$ and define the costs as

$$C^{(p)}_{(il)(jm)} = \left| d_X(x_i, x_j) - d_Y(y_l, y_m) \right|^p$$

Then we can consider the distance

$$d^{(p)}_P(\{x_i\}, \{y_i\}) = \frac{1}{2} \min_{\pi \in P_n} \sum_{1 \leq i,j \leq n} C^{(p)}_{(il)(jm)} R_{ij} R_{lm}$$

Quadratic Assignment Problem

$$d_P^{(p)}(\{x_i\},\{y_i\}) = \frac{1}{2}\min_{\pi\in P_n}\sum_{1\leq i,j\leq n} C^{(p)}_{(il)(jm)} R_{ij} R_{lm}$$

Rewriting in matrix notation , we get to the quadratic programm:

$$\min_{R\in\{0,1\}^{n\times n}} vec(R)^T C vec(R)$$

$$s.t. \quad R1 = 1,\ R^T 1 = 1$$

where $vec(R)$ is a column-stacked reshaping of R.

The quadratic optimization problem is also known as **Quadratic Assignment Problem (QAP).**

Quadratic Assignment Problem

$$\min_{R \in \{0,1\}^{n \times n}} vec(R)^T C vec(R)$$

$$s.t. \quad R1 = 1,\ R^T 1 = 1$$

This combinatorial optimization problem is unfortunately NP-hard.

In the literature there have been several attempts to relax the problem to make it more tractable. Int the following we will present some of these approaches.

Bistochastic relaxation

Leave the combinatorial setting by allowing the correspondence to take on continuous values.

$$\min_{R\in[0,1]^{n\times n}} vec(R)^T C vec(R)$$

$$s.t. \quad R1 = 1,\ R^T 1 = 1$$

Now each row and column can be regarded as discrete probability distributions.

X

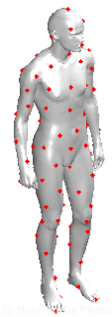

0.1	0.1	0.3	0.2	0.3
0.3	0.1	0.1	0.1	0.4
0.2	0.3	0.2	0.2	0.1
0.2	0.2	0.1	0.4	0.1
0.2	0.3	0.3	0.1	0.1

$\sum = 1$

$\sum = 1$

Y

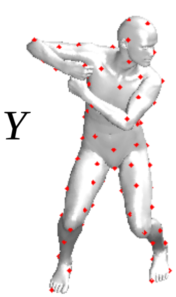

Probability distribution

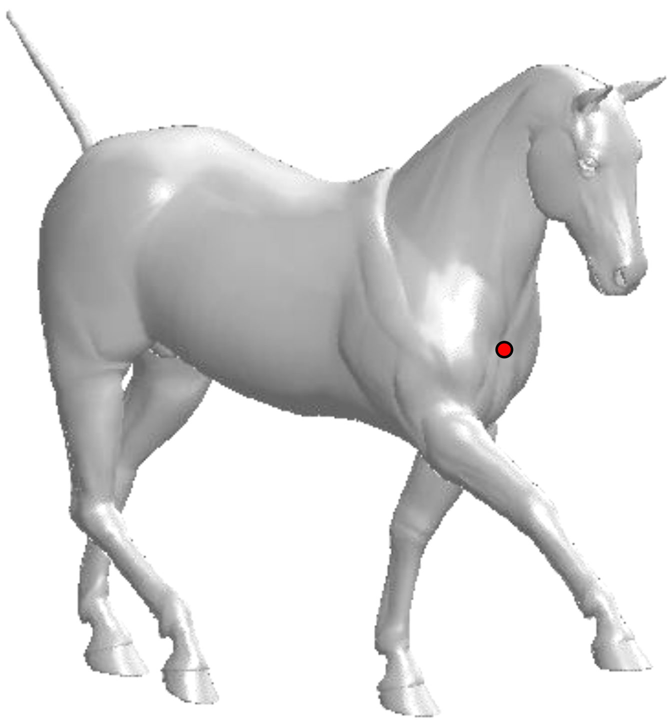

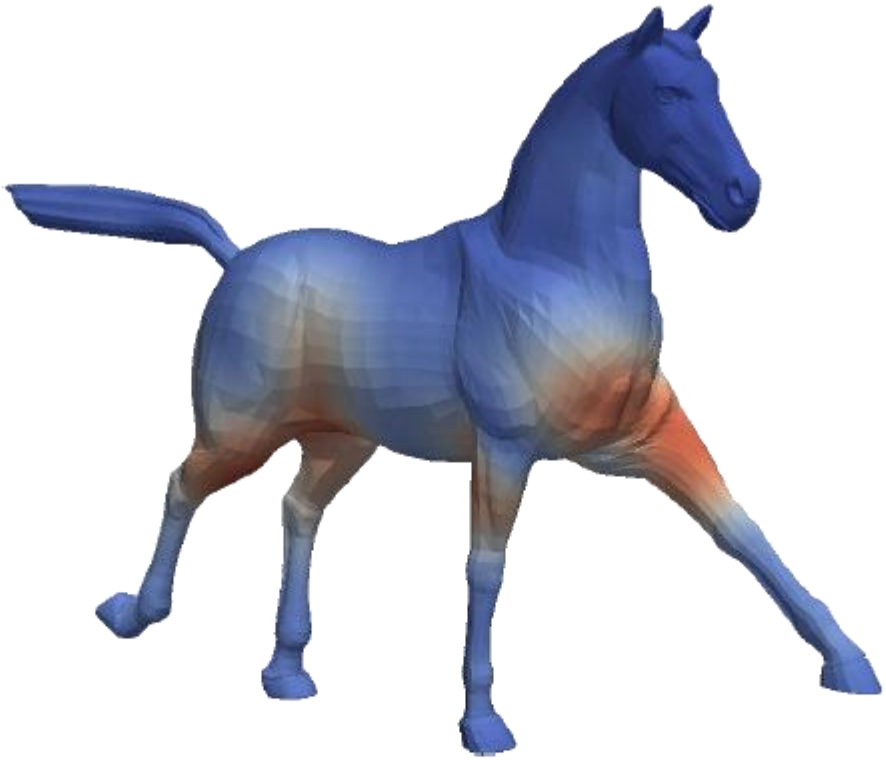

Optimization

$$\min_{R \in [0,1]^{n \times n}} vec(R)^T C vec(R)$$

$$s.t. \quad R\mathbf{1} = \mathbf{1},\ R^T\mathbf{1} = \mathbf{1}$$

Can be solved via projected gradient descent.

- ☹ Slow convergence
- ☹ Local optimum
- ☹ Implement efficient projection
- ☹ Choose good starting point
- ☹ Choose step size or do line search
- ☹ Binarize the final solution

- ☺ Easy to implement
- ☺ Local optima are usually good enough in practice

Spectral relaxation

An alternative characterization of permutation matrices

$$R \in \{0,1\}^{n \times n}, \quad R^T R = I$$

gives rise to the spectral relaxation

$$\min_{R \in [0,1]^{n \times n}} vec(R)^T C vec(R)$$

$$s.t. \quad R^T R = I$$

or even more relaxed:

$$\min_{x \in [0,1]^{n^2}} x^T C x$$

$$s.t. \quad x^T x = n$$

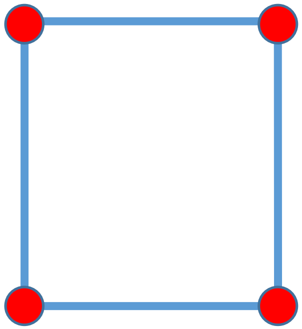

bistochastic

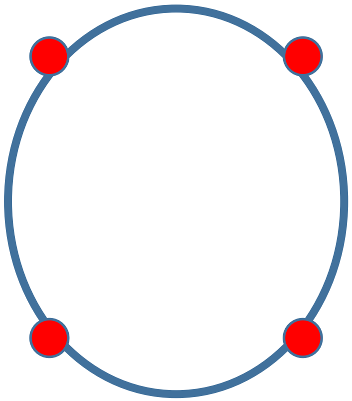

spectral

Spectral relaxation

$$\min_{x \in [0,1]^{n^2}} x^T C x$$

$$s.t. \quad x^T x = n$$

Global optimum given by eigenvector of C associated to smallest eigenvalue.

☹ The final solution is not a correspondence (needs post-processing)
☹ Needs binarization
☹ We are losing contact with the Gromov-Hausdorff…

☺ Easy to implement
☺ Global optimum
☺ Efficient