Analysis of Three-Dimensional Shapes
(IN2238, TU München, Summer 2014)

The Assignment Problem
(24.04.2014)

Dr. Emanuele Rodolà
rodola@in.tum.de
Room 02.09.058, Informatik IX

Announcement

On popular demand, we moved the lecture to:

Thursday 16:00-18:00

However, during <u>holidays</u> we will move back to:

Monday 10:00-12:00

<u>Check the calendar on the course web page!</u>

Seminar

“Heat Kernel Signature”
Thomas Hörmann

Wednesday, April 30
14:00 Room 02.09.023

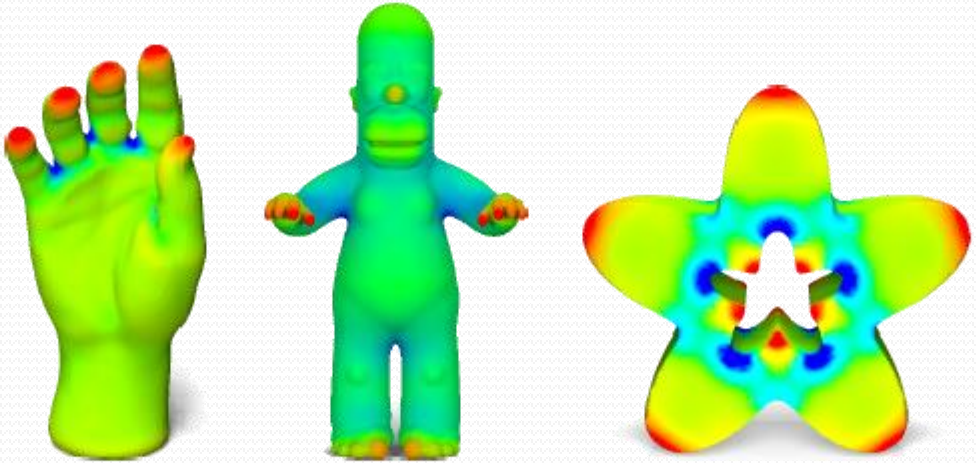

The matching game

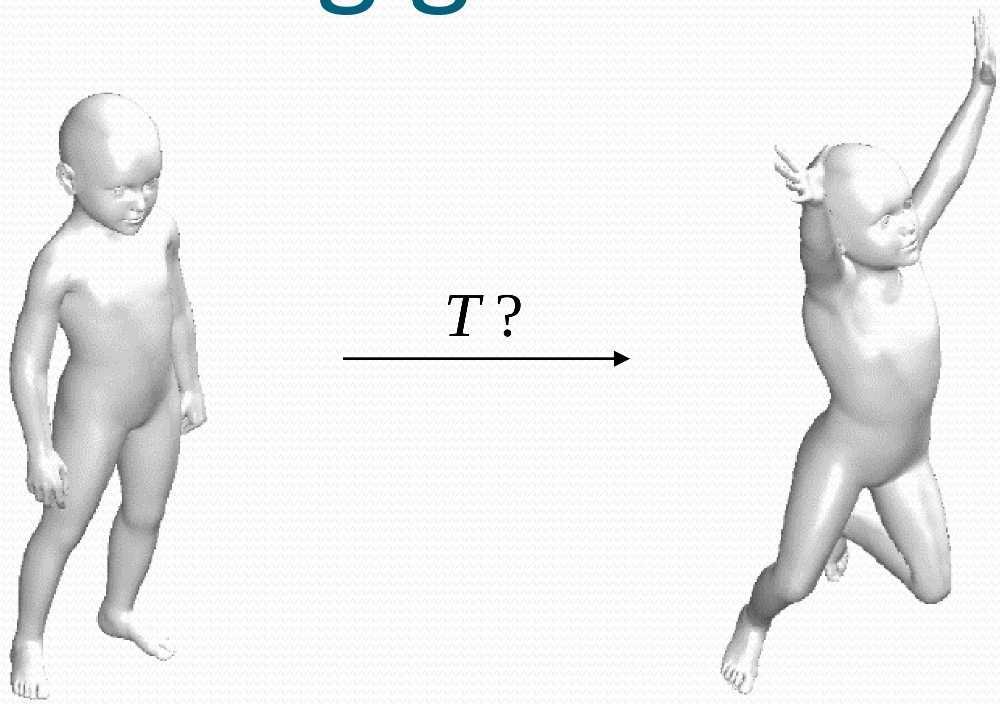

You will be given two shapes. Find the best correspondence you can!

- Do not bother looking for them on the web
- You can use *whatever* technique you want, or mixtures thereof
- The best solution will get a **prize**!

The space of shapes

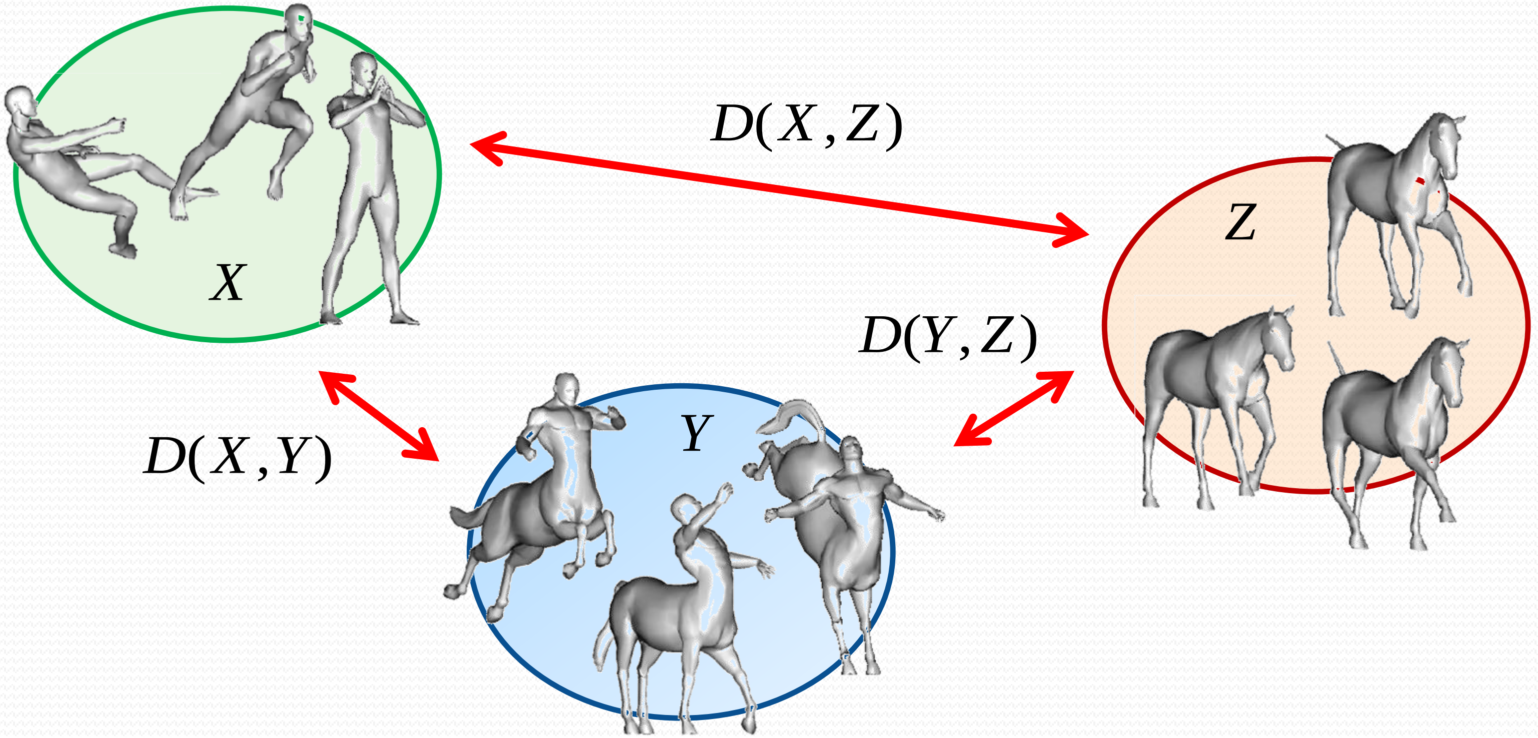

Is there something like a "space of shapes"? There are <u>many</u>!

Lipschitz distance

$$d_{\mathcal{L}}(X,Y) = \inf_{f:X \to Y} \log\left(\max\left\{\mathrm{dil}(f), \mathrm{dil}(f^{-1})\right\}\right)$$

where $$\mathrm{dil}\, f = \sup_{x_1, x_2 \in X} \frac{d_Y(f(x_1), f(x_2))}{d_X(x_1, x_2)}$$

- The shapes are assumed to be bi-Lipschitz homeomorphic.
- $d_{\mathcal{L}} = \infty$ if not.

- $d_{\mathcal{L}}$ is a metric on the space of isometry classes of compact metric spaces.

Hausdorff distance

$$d_{\mathcal{H}}^{Z}(X,Y) = \max\left\{ \sup_{x\in X} \mathrm{dist}_{Z}(x,Y), \sup_{y\in Y} \mathrm{dist}_{Z}(y,X) \right\}$$

where $X, Y \subset (Z, d_Z)$

- $d_{\mathcal{H}}^{Z}$ is a *semi*-metric on the space of compact *subsets* of a metric space.
- A difference in a single point can make $d_{\mathcal{H}}^{Z}$ arbitrarily large.
- It only allows to compare subsets of a common metric space.

- It captures a more intuitive notion of distance among shapes

Gromov-Hausdorff distance

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

where $f : X \to Z,\ \ g : Y \to Z$ are isometric embeddings

- The infimum is taken over a huge feasible set.

- $d_{\mathcal{GH}}$ is a metric on the space of isometry classes of compact metric spaces.
- It encodes a natural notion of distance among shapes

Distances in the space of shapes

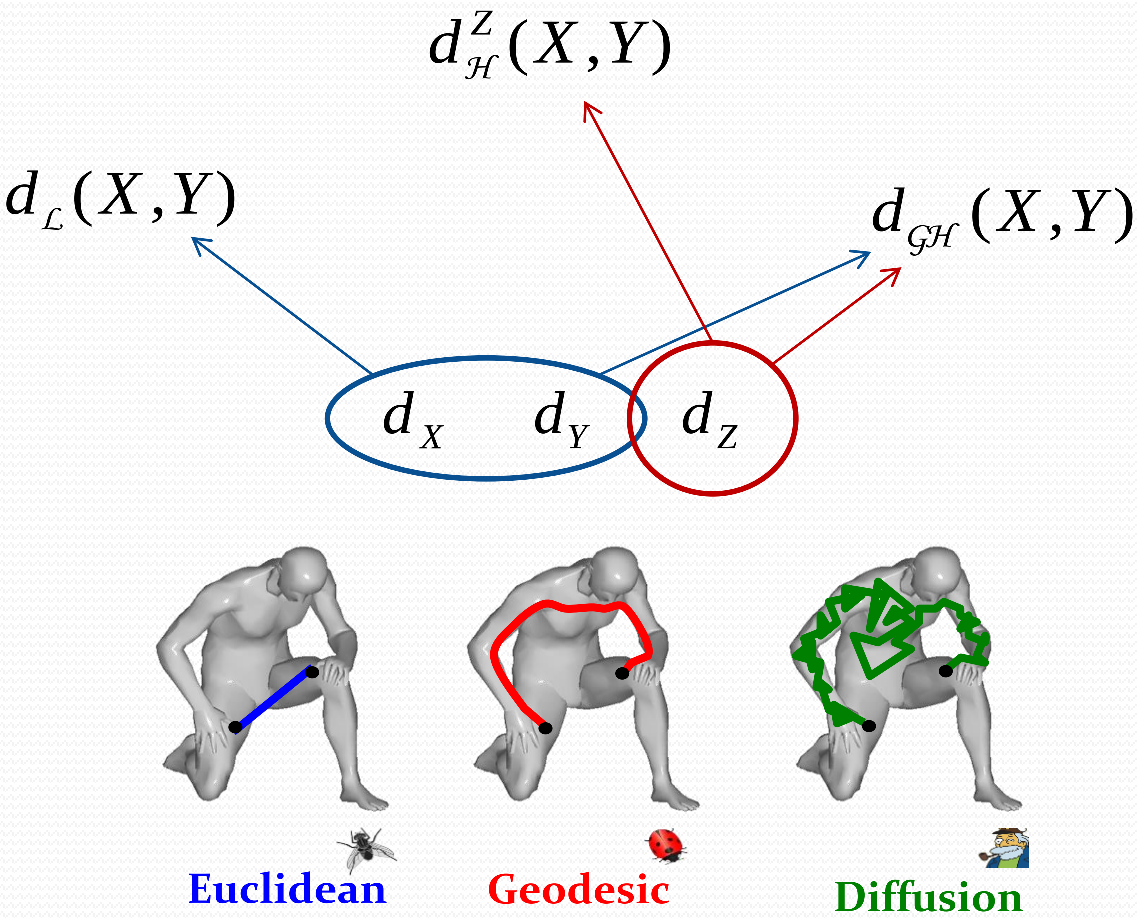

Gromov-Hausdorff distance

$$d_{\mathcal{GH}}(X,Y) = \frac{1}{2} \inf_{R \subset X \times Y} \operatorname{dis} R$$

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

The infimum is taken over all correspondences *R* between *X* and *Y*.

$d_{\mathcal{GH}}(X,Y) = 0$ if and only if *X* and *Y* are isometric.

Still impractical, but it gives an intuition on how to proceed in order to actually **compute** the distance.

Estimating the GH distance

$$d_{\mathcal{GH}}(X,Y) = \frac{1}{2} \inf_{R \subset X \times Y} \operatorname{dis} R$$

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

Let $\mathbf{X} = \{x_i\}_{i=1}^{n}$ be a r-covering of X, and $\mathbf{Y} = \{y_j\}_{j=1}^{n}$ be a r'-covering of Y. Then we can define an alternative distance

$$d_{\mathcal{P}}(\mathbf{X},\mathbf{Y}) = \frac{1}{2} \min_{\pi \in P_n} \max_{1 \le i,j \le n} \left| d_{\mathbf{X}}(x_i,x_j) - d_{\mathbf{Y}}(y_{\pi i}, y_{\pi j}) \right|$$

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

And one can prove $d_{\mathcal{GH}}(X,Y) \le r + r' + d_{\mathcal{P}}(\mathbf{X},\mathbf{Y})$

Estimating the GH distance

$$\mathbf{X} = \{x_i\}_{i=1}^{n} \quad \mathbf{Y} = \{y_j\}_{j=1}^{n} \qquad d_{\mathcal{P}}(\mathbf{X},\mathbf{Y}) = \frac{1}{2}\min_{\pi \in P_n}\max_{1 \le i,j \le n}\left| d_{\mathbf{X}}(x_i, x_j) - d_{\mathbf{Y}}(y_{\pi i}, y_{\pi j})\right|$$

(X, d_X)

(Y, d_Y)

Estimating the GH distance

$$\boxed{\mathbf{X} = \{x_i\}_{i=1}^{n} \quad \mathbf{Y} = \{y_j\}_{j=1}^{n}}$$

FPS

$$d_{\mathcal{P}}(\mathbf{X},\mathbf{Y}) = \frac{1}{2} \min_{\pi \in P_n} \max_{1 \le i, j \le n} \left| d_{\mathbf{X}}(x_i, x_j) - d_{\mathbf{Y}}(y_{\pi i}, y_{\pi j}) \right|$$

(X, d_X)

$\mathbf{X}$

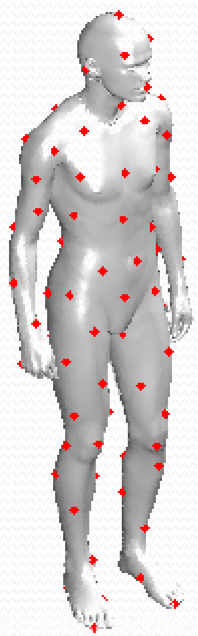

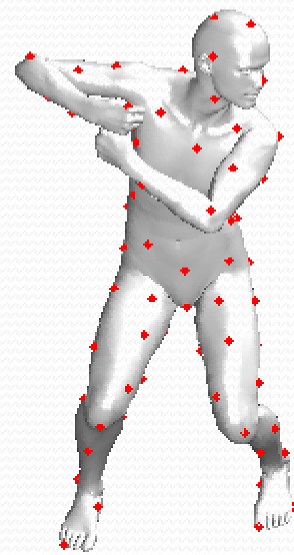

(Y, d_Y)

$\mathbf{Y}$

Estimating the GH distance

$$\mathbf{X} = \{x_i\}_{i=1}^{n} \quad \mathbf{Y} = \{y_j\}_{j=1}^{n} \qquad d_{\mathcal{P}}(\mathbf{X},\mathbf{Y}) = \frac{1}{2} \cancel{\min_{\pi \in P_n}} \max_{1 \le i,j \le n} \left| d_{\mathbf{X}}(x_i, x_j) - d_{\mathbf{Y}}(y_{\pi i}, y_{\pi j}) \right|$$

Assume we know the *true* matches between X and Y:

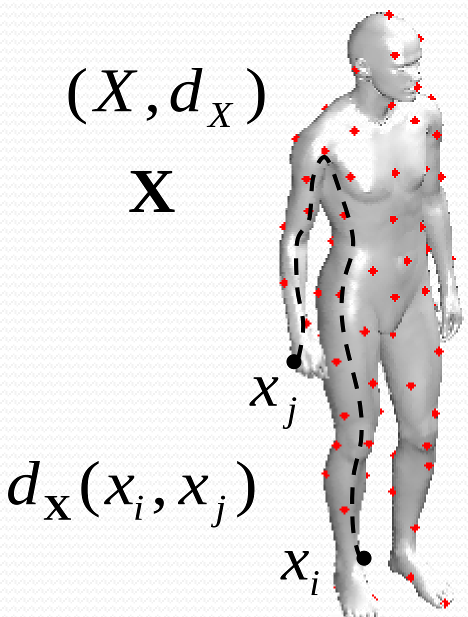

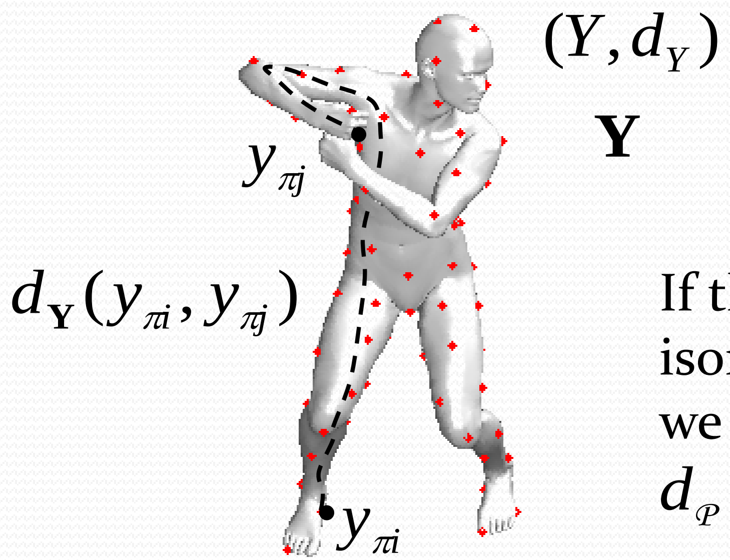

If the shapes are isometric, then we would expect $d_{\mathcal{P}}(\mathbf{X},\mathbf{Y}) = 0$

Estimating the GH distance

$$\mathbf{X} = \{x_i\}_{i=1}^{n} \quad \mathbf{Y} = \{y_j\}_{j=1}^{n} \qquad d_{\mathcal{P}}(\mathbf{X},\mathbf{Y}) = \frac{1}{2} \min_{\pi \in P_n} \max_{1 \le i,j \le n} \left| d_{\mathbf{X}}(x_i, x_j) - d_{\mathbf{Y}}(y_{\pi i}, y_{\pi j}) \right|$$

Keep in mind that:

- In practice, we can only expect near-isometries, i.e. $d_{\mathcal{P}}(\mathbf{X},\mathbf{Y}) < \varepsilon$
- $d_{\mathcal{GH}}$ is consistent to sampling, that is $\left| d_{\mathcal{GH}}(X,Y) - d_{\mathcal{P}}(\mathbf{X},\mathbf{Y}) \right|$ is bounded above
- In general, we don't know the true matches between X and Y!

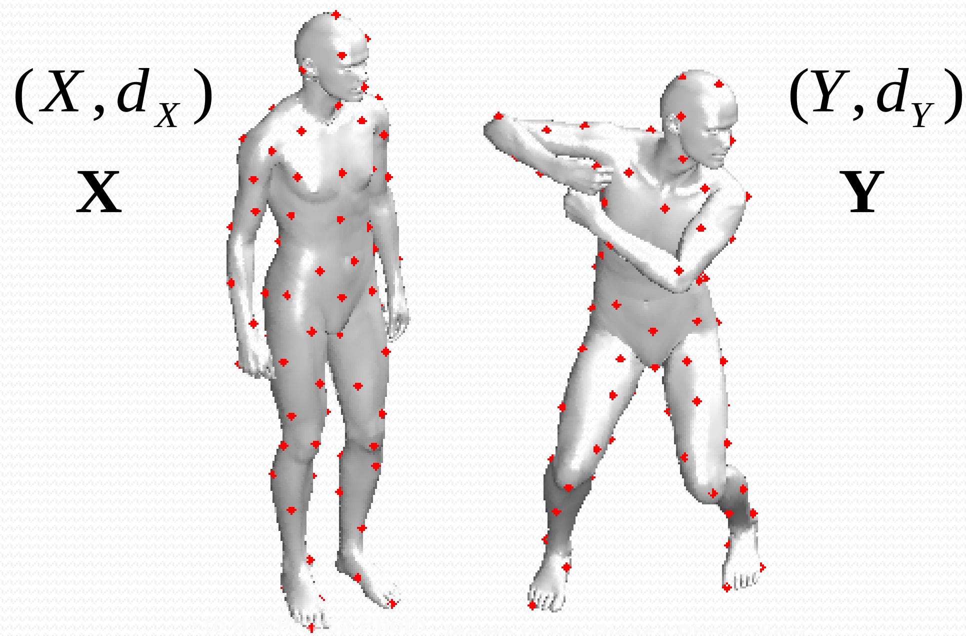

Example

Assume we are given a collection of shapes and the *true correspondence* among them…

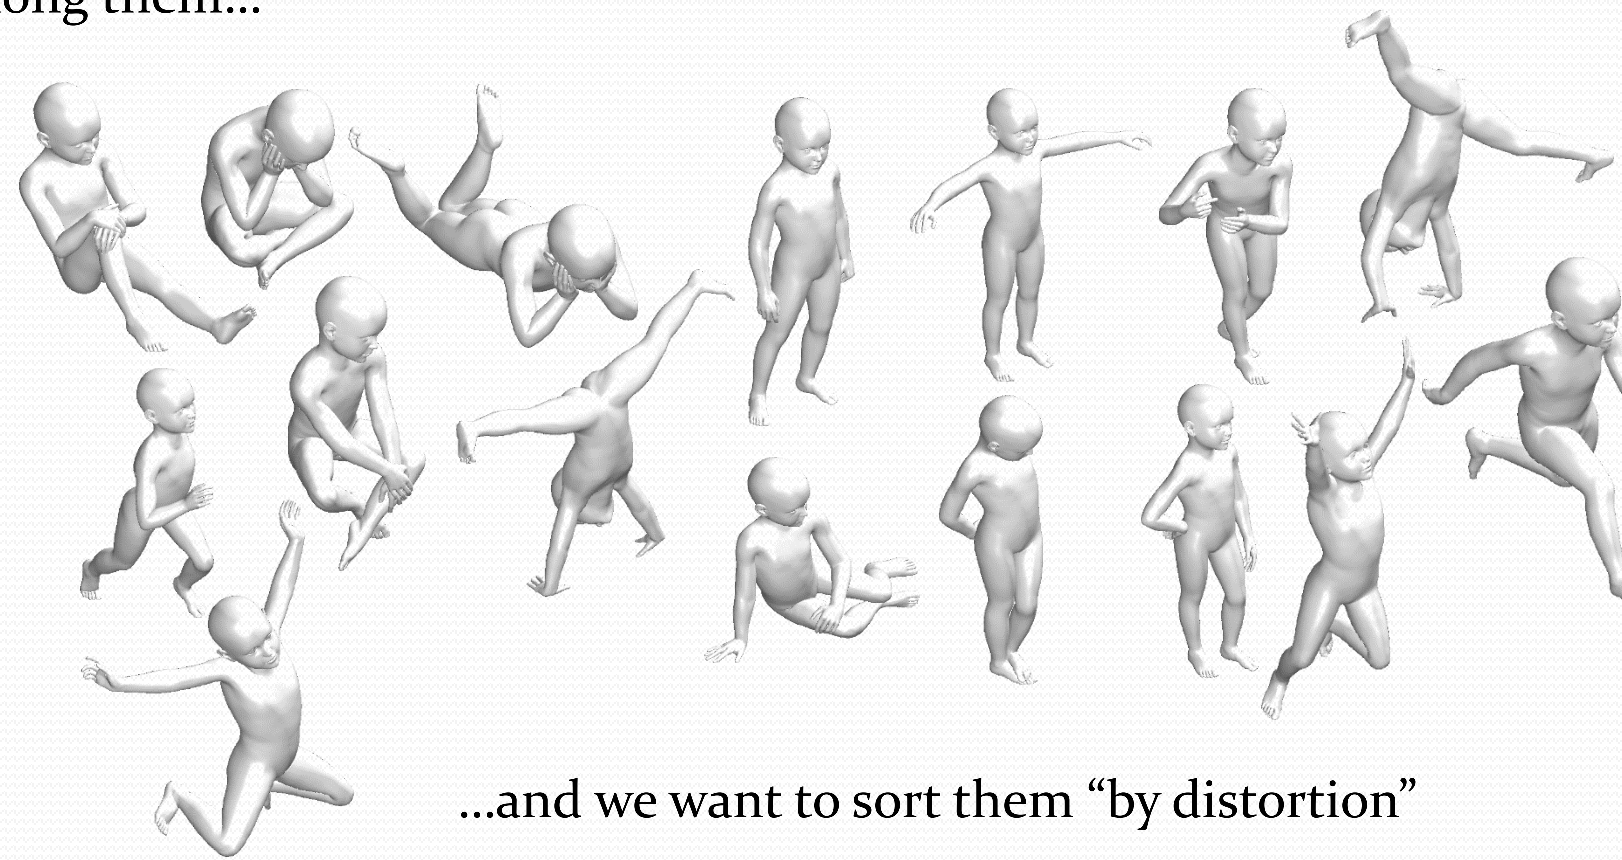

…and we want to sort them "by distortion"

Example

We can approximate their GH distance $d_{\mathcal{P}}(\mathbf{X},\mathbf{Y})$ to a preselected "null" pose:

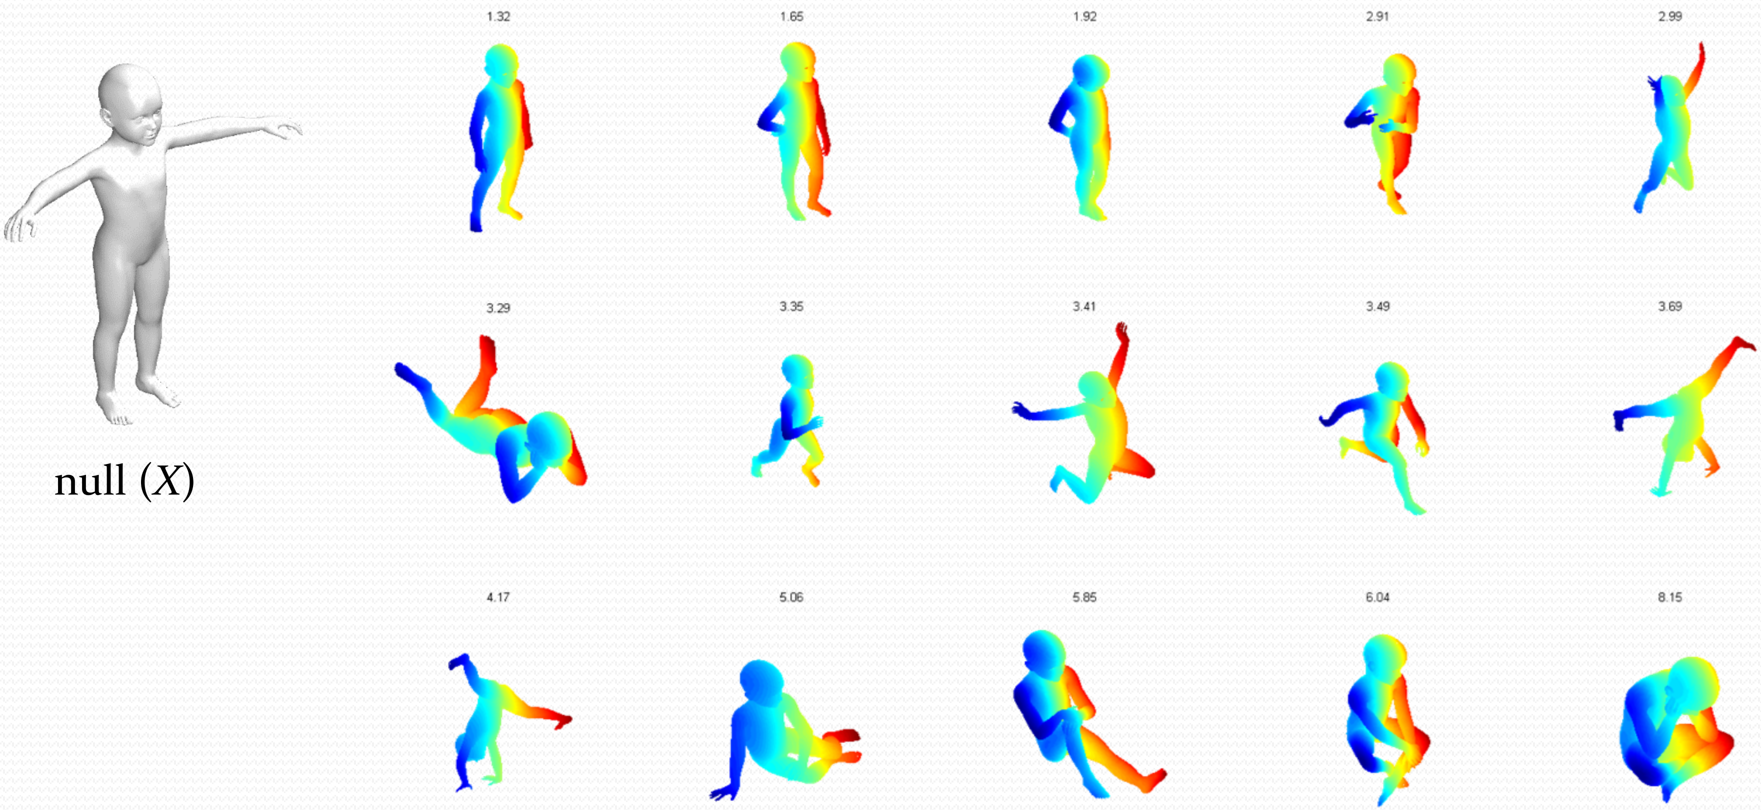

$$d_{\mathcal{P}}(\mathbf{X},\mathbf{Y}) = \frac{1}{2}\cancel{\min_{\pi\in P_n}}\max_{1\le i,j\le n}\left|d_{\mathbf{X}}(x_i,x_j) - d_{\mathbf{Y}}(y_{\pi i},y_{\pi j})\right|$$

Finding a correspondence

In general, we don't know the true correspondence between the two shapes.

$$d_{\mathcal{P}}(\mathbf{X},\mathbf{Y}) = \frac{1}{2} \min_{\pi \in P_n} \max_{1 \le i, j \le n} \left| d_{\mathbf{X}}(x_i, x_j) - d_{\mathbf{Y}}(y_{\pi i}, y_{\pi j}) \right|$$

In order to compute (an approximation to) the GH distance, we have to minimize over all possible correspondences P_n

This corresponds to minimizing over all possible surjective maps $f : X \to Y$ in our original formulation:

$$d_{\mathcal{GH}}(X,Y) = \frac{1}{2} \inf_{f : X \to Y} \operatorname{dis} f$$

Discretization

Let us represent shapes by their corresponding *ordered* collection of points:

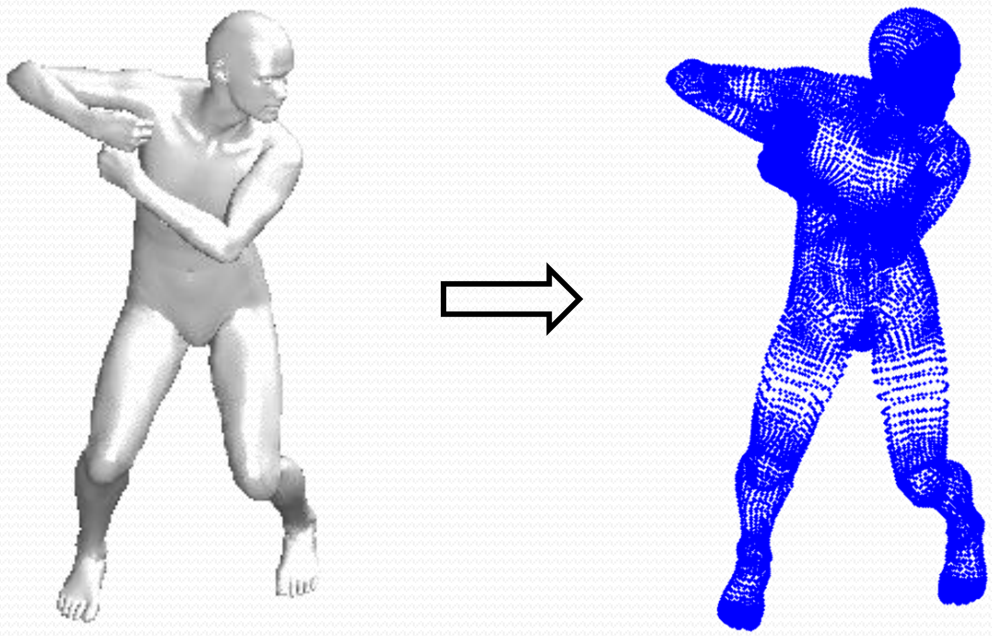

$$X = \{p_i\}_{i=1}^{n} \subset \mathbf{R}^3$$

$$\mathbf{X} \subset X \quad \text{(FPS)}$$

Discretization

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

X

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

Y

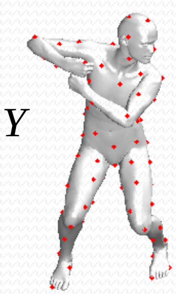

$R_{ij} = 1$
if x_i and y_j are
in correspondence

Discretization

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

X

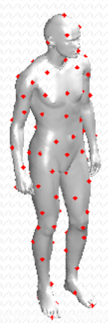

0	1	0	0	0
0	0	0	1	0
1	0	0	0	0
0	0	0	0	1
0	0	1	0	0

Y

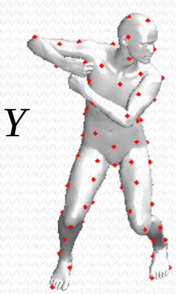

Asking for a bijection corresponds to require R to be a *permutation matrix*.

In other words, we are optimizing over all permutations of $\{1, \ldots, n\}$

$$X = \begin{bmatrix} 1 & 2 & 3 & 4 & 5 \end{bmatrix}$$

$$\big\downarrow \pi$$

$$Y = \begin{bmatrix} 3 & 5 & 2 & 1 & 4 \end{bmatrix}$$

Discretization

$$d_{\mathcal{P}}(\mathbf{X},\mathbf{Y}) = \frac{1}{2} \min_{\pi \in P_n} \max_{1 \le i,j \le n} \left| d_{\mathbf{X}}(x_i, x_j) - d_{\mathbf{Y}}(y_{\pi i}, y_{\pi j}) \right|$$

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

$$C_{(i\ell)(jm)} = \left| d_{\mathbf{X}}(x_i, x_j) - d_{\mathbf{Y}}(y_\ell, y_m) \right|$$

where the notation (ab) represents the match $(x_a, y_b) \in X \times Y$

$$\begin{array}{c|ccccc|c}
(x_1,y_1) & 0 & 13.5 & 23.4 & 104.6 & 7.64 & \\
(x_1,y_2) & 13.5 & 0 & 13.52 & 11.2 & 71.1 & \\
(x_1,y_3) & 23.4 & 13.52 & 0 & 0.22 & 23.44 & \\
\vdots & 104.6 & 11.2 & 0.22 & 0 & 16.5 & \\
 & 7.64 & 71.1 & 23.44 & 16.5 & 0 & \\
 & (x_1,y_1) & (x_1,y_2) & (x_1,y_3) & \cdots & & \ddots
\end{array}$$

Quite big!

Discretization

$$d_{\mathcal{P}}(\mathbf{X},\mathbf{Y})=\frac{1}{2}\min_{\pi\in P_n}\max_{1\le i,j\le n}\left|d_{\mathbf{X}}(x_i,x_j)-d_{\mathbf{Y}}(y_{\pi i},y_{\pi j})\right|$$

Rewriting with matrix notation, we get:

$$d_{\mathcal{P}}(\mathbf{X},\mathbf{Y})=\frac{1}{2}\min_{R}\max_{i,j,\ell,m}C_{(i\ell)(jm)}R_{ij}R_{\ell m}$$

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

Discretization

$$d_{\mathcal{P}}(\mathbf{X},\mathbf{Y}) = \frac{1}{2}\min_{R}\max_{i,j,\ell,m} C_{(i\ell)(jm)} R_{ij} R_{\ell m}$$

As we have already mentioned, this distance is sensitive to outliers.

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

$$C^{(p)}_{(i\ell)(jm)} = \left| d_{\mathbf{X}}(x_i, x_j) - d_{\mathbf{Y}}(y_\ell, y_m) \right|^p$$

Then we can consider the distance

$$d^{(p)}_{\mathcal{P}}(\mathbf{X},\mathbf{Y}) = \frac{1}{2}\min_{R}\sum_{i,j}\sum_{\ell,m} C^{(p)}_{(i\ell)(jm)} R_{ij} R_{\ell m}$$

Quadratic Assignment Problem

$$d_{\mathcal{P}}^{(p)}(\mathbf{X},\mathbf{Y}) = \frac{1}{2}\min_{R}\sum_{i,j}\sum_{\ell,m} C_{(i\ell)(jm)}^{(p)} R_{ij} R_{\ell m}$$

In practice we will be interested in finding a *minimizer* rather than a minimum. Rewriting in matrix notation, we get to the quadratic program:

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

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

where $\operatorname{vec}\{R\}\in\mathbf{R}^{n^2}$ is a column-stacked reshaping of R.

This quadratic optimization problem is also known as the Lawler formulation of the **Quadratic Assignment Problem (QAP)**.

Quadratic Assignment Problem

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

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

This combinatorial optimization problem is unfortunately NP-hard.

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

Continuous relaxation

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

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

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

Continuous relaxation

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

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

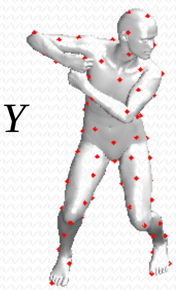

We can now regard each row and column as a discrete probability distribution associated to that point.

Continuous relaxation

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

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

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

The resulting feasible set is the set of *doubly-stochastic matrices*.

It forms a *convex* set known as the *Birkhoff polytope* (or *assignment* polytope). The n! vertices of this polytope are the permutation matrices.

Continuous relaxation

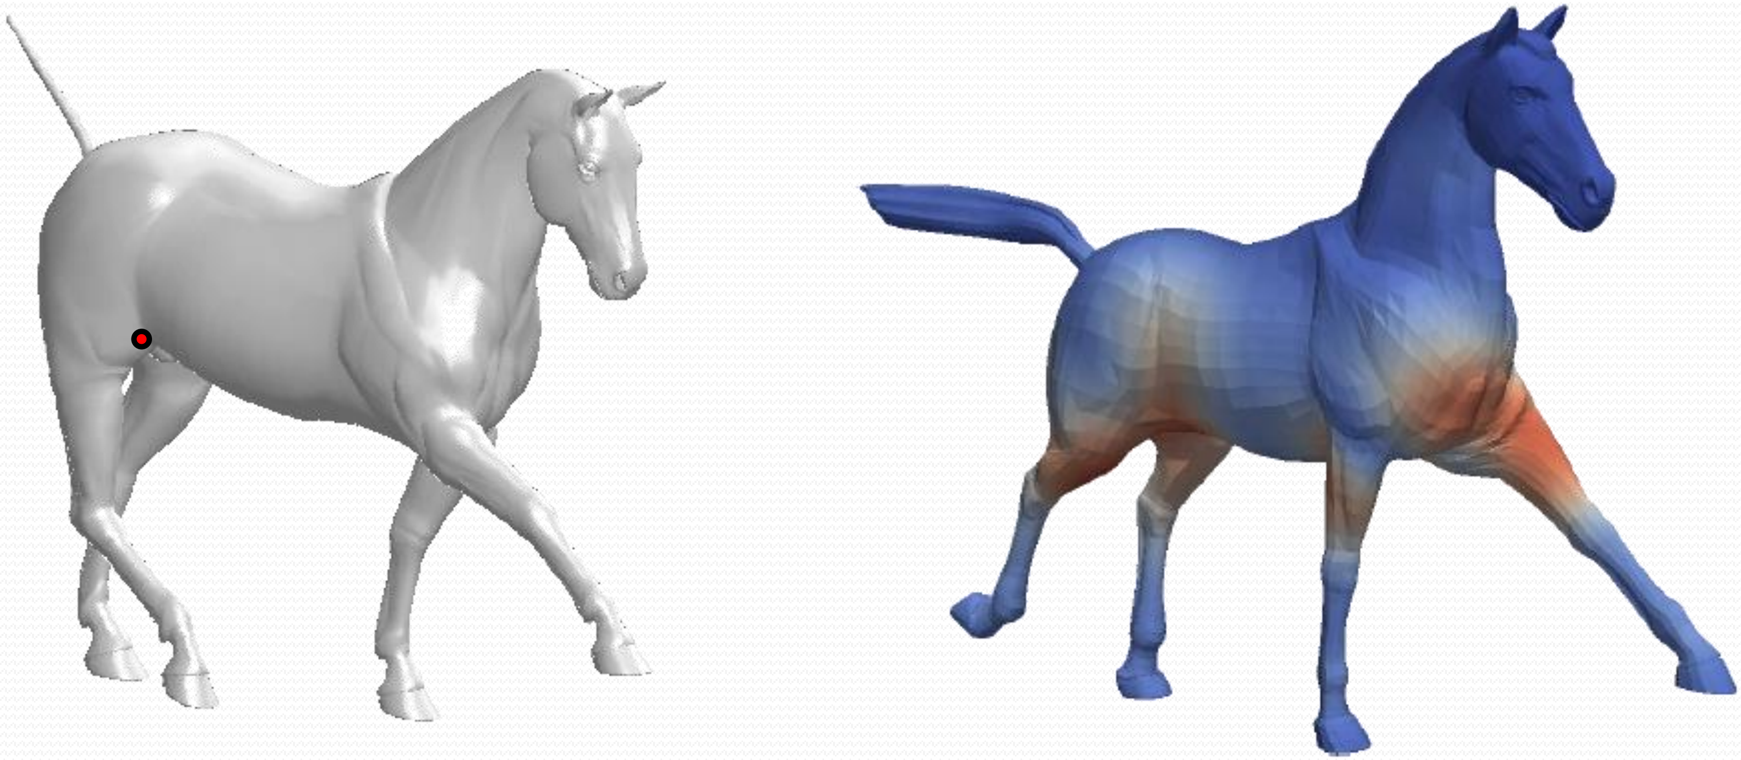

Continuous relaxation

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

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

Optimize by projected gradient descent. Let $\mathbf{x} \equiv \mathrm{vec}\{R\}$

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

$$\text{s.t.} \quad A\mathbf{x}=\mathbf{b}$$ also called *mapping constraints*

We can find a local optimum via the recursive equations:

$$\mathbf{x}^{(t+1)} = \Pi_{A\mathbf{x}=\mathbf{b}}\left(\mathbf{x}^{(t)} - \gamma C \, \mathbf{x}^{(t)}\right)$$

Where $\gamma > 0$ is the step length and $\Pi_{A\mathbf{x}=\mathbf{b}}$ is a projection operator.

Continuous relaxation

☹ 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

Other approaches further relax the QAP by taking the point of view of *regularization* theory.

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

$$\text{s.t.} \quad \|\mathbf{x}\|_2^2 = 1$$

Each value of **x** is interpreted as the *confidence* of the corresponding match.

Note that we are losing the connection with the Gromov-Hausdorff distance. <u>The optimal **x** is not even guaranteed to be a correspondence anymore</u>!

Spectral relaxation

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

$$\text{s.t.} \quad \|\mathbf{x}\|_2^2 = 1$$

If C is a Hermitian matrix, its *Rayleigh quotient* is defined as

$$R(C,\mathbf{x}) = \frac{\mathbf{x}^{*} C\,\mathbf{x}}{\mathbf{x}^{*}\mathbf{x}}$$

It can be shown that

$$R(C,\mathbf{v}_{\mathrm{min}}) = \lambda_{\mathrm{min}} \le R(C,\mathbf{x}) \le \lambda_{\mathrm{max}} = R(C,\mathbf{v}_{\mathrm{max}})$$

where λ and $\mathbf{v}$ are eigenvalues and eigenvectors of C.

Spectral relaxation

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

$$\text{s.t.} \quad \|\mathbf{x}\|_2^2 = 1$$

The global optimum is given by the eigenvector associated to $\lambda_{\min}$.

It should be noted that, since C has non-negative entries, this eigenvector will have values in $[0,1]$ (Perron-Frobenius theorem).

Spectral relaxation

A common way to compute principal eigenvectors of a given matrix is via the *power iterations*:

$$\mathbf{x}^{(t+1)} = \frac{C\mathbf{x}^{(t)}}{\left\| C\mathbf{x}^{(t)} \right\|_2}$$

The sequence $\left(\mathbf{x}^{(t)}\right)$ converges to the dominant eigenvector under mild assumptions on C and $\mathbf{x}^{(0)}$.

Spectral relaxation

☹ 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

Spectral relaxation

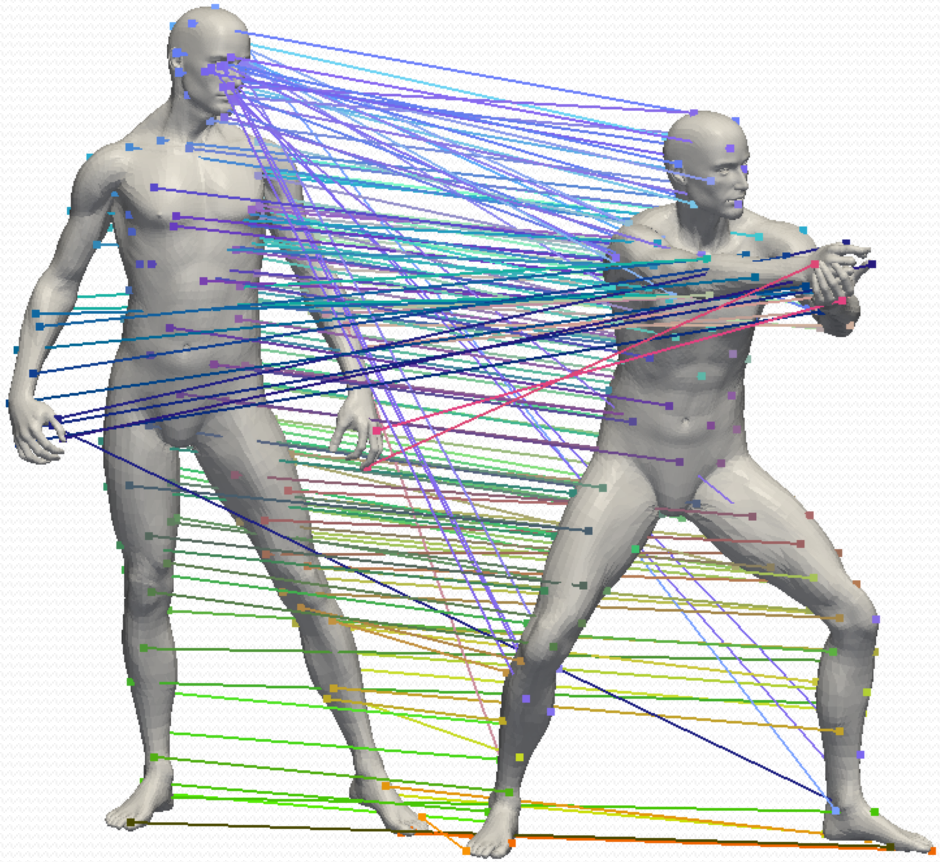

Partiality

Back to the continuous formulation:

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

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

The mapping constraints are imposing at least one match for each point in either shape.

This does not take into account *partiality*. See the following example.

Partiality

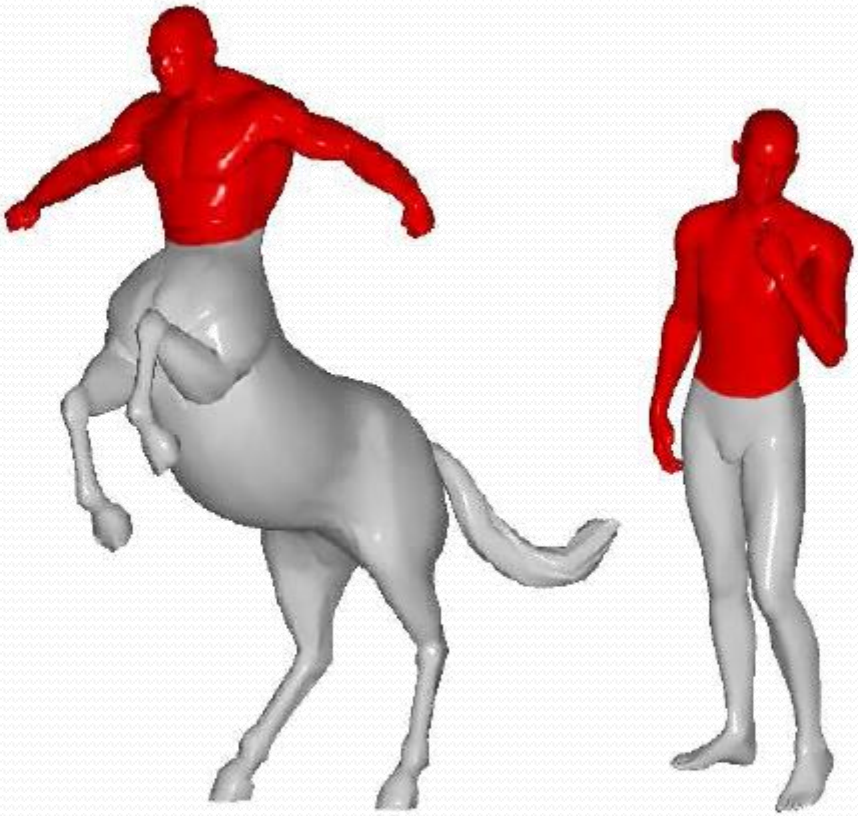

Partiality

If we want to account for partiality, we have to allow *unmatchable* points. One way to obtain this is by allowing the correspondence matrix R to have rows and columns summing up to zero:

$$\min_{R \in [0,1]^{n \times n}} \operatorname{vec}\{R\}^{\mathrm{T}} C \operatorname{vec}\{R\}$$
$$\text{s.t.} \quad \mathbf{1}^{\mathrm{T}} R \mathbf{1} = 1$$

Rewriting:

$$\min_{\mathbf{x} \in [0,1]^{n^2}} \mathbf{x}^{\mathrm{T}} C \, \mathbf{x}$$
$$\text{s.t.} \quad \|\mathbf{x}\|_1 = 1$$

In other words, we are now optimizing over *probability distributions over the space of possible matches*.

Game-theoretic relaxation

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

$$\text{s.t.} \quad \|\mathbf{x}\|_1 = 1$$

It turns out that we can interpret this relaxation using notions (and solvers!) from Game Theory.

Note that the L^1 regularizer will favor *sparse* solutions.

Game-theoretic relaxation

One possible way to solve the relaxed problem is via the *replicator dynamics* equations:

$$\mathbf{x}_i^{(t+1)} = \mathbf{x}_i^{(t)} \frac{\left(C\mathbf{x}^{(t)}\right)_i}{\mathbf{x}^{(t)} C\, \mathbf{x}^{(t)}} \qquad \text{for} \quad i = 1, \ldots, n$$

Compare with the power iterations used in the L^2 case:

$$\mathbf{x}^{(t+1)} = \frac{C\mathbf{x}^{(t)}}{\left\| C\mathbf{x}^{(t)} \right\|_2}$$

Game-theoretic relaxation

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

$$\text{s.t.} \quad \|\mathbf{x}\|_1 = 1$$

It seems like we are losing the mapping constraints again…

In fact, it can be shown that the mapping constraints *can be incorporated into the cost matrix*, and still have the guarantee that they will be satisfied by the final solution.

$$C^{(p)}_{(i\ell)(jm)} = \left| d_{\mathbf{X}}(x_i, x_j) - d_{\mathbf{Y}}(y_\ell, y_m) \right|^p$$

Set $C^{(p)}_{(i\ell)(jm)} = 0$ whenever $i = j$ or $\ell = m$.

Note that you need to pass to a maximization problem!

Game-theoretic relaxation

☹ Very sparse
☹ Local optimum

☺ The mapping constraints can be incorporated into the cost matrix
☺ Easy to implement
☺ Does not need binarization
☺ Accurate
☺ Efficient
☺ Game-theoretic interpretation (why not consider different games?)

Game-theoretic relaxation

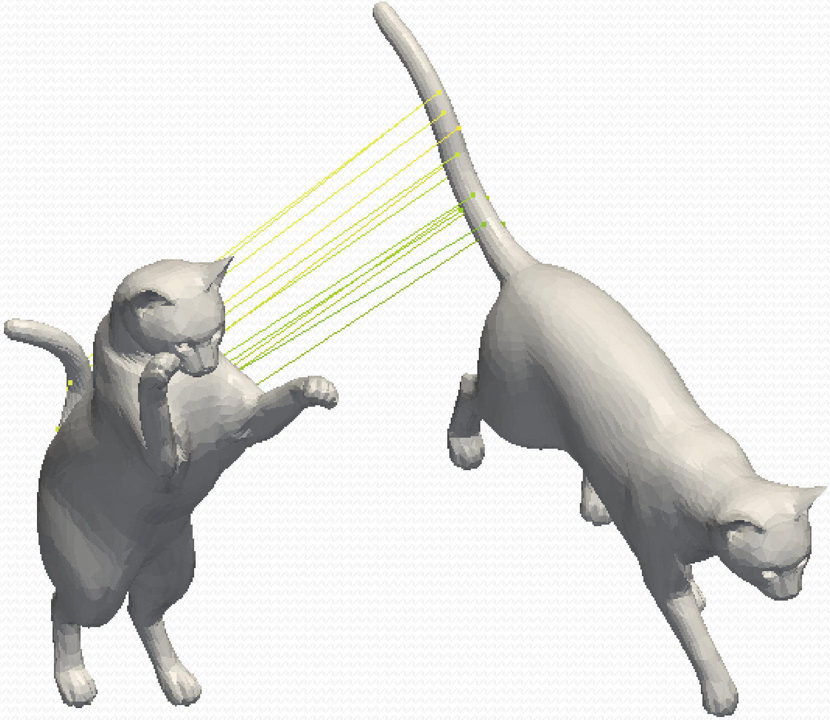

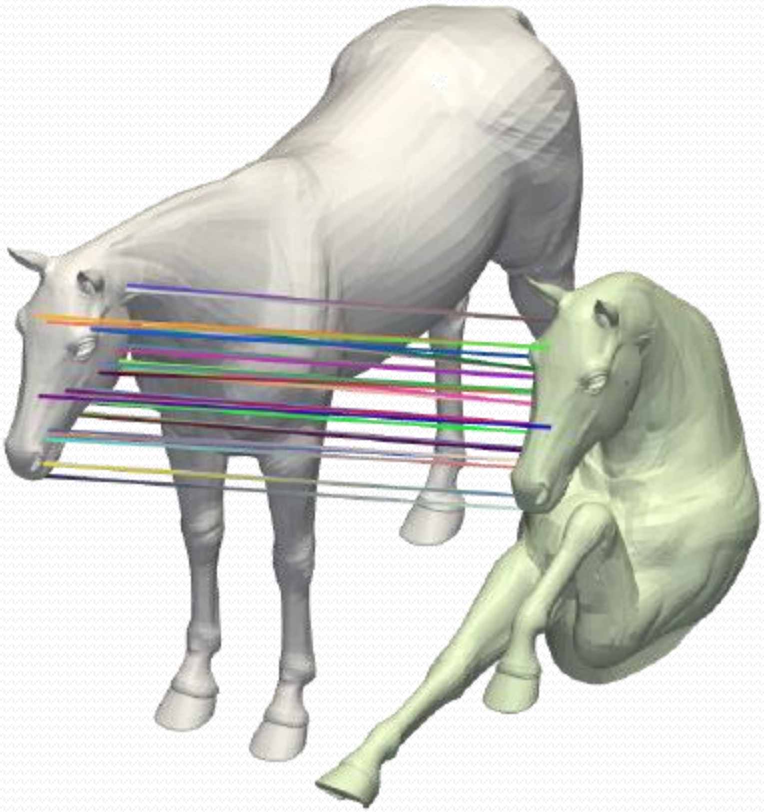

Similarity-based matching

We have considered problems of the form:

$$d_{\mathcal{P}}^{(p)}(\mathbf{X},\mathbf{Y}) = \frac{1}{2}\min_{R}\sum_{i,j}\sum_{\ell,m} C^{(p)}_{(i\ell)(jm)} R_{ij} R_{\ell m}$$

with

$$C^{(p)}_{(i\ell)(jm)} = \left| d_{\mathbf{X}}(x_i, x_j) - d_{\mathbf{Y}}(y_\ell, y_m) \right|^p$$

One could consider other cost functions, at the price of losing connections with the theory behind Gromov-Hausdorff distances.

One popular choice is the Gaussian similarity

$$C^{(p)}_{(i\ell)(jm)} = e^{-\beta \left| d_{\mathbf{X}}(x_i, x_j) - d_{\mathbf{Y}}(y_\ell, y_m) \right|^2}$$

Note that you need to pass to a maximization problem!

Linear Assignment Problem

A related problem is the Linear Assignment Problem (LAP).
Differently from the QAP, it does not impose preservation of the metric but rather of *pointwise* quantities. This results in a linear cost:

QAP

$$\min_{R} \sum_{i,j} \sum_{\ell,m} C_{(i\ell)(jm)} R_{ij} R_{\ell m}$$

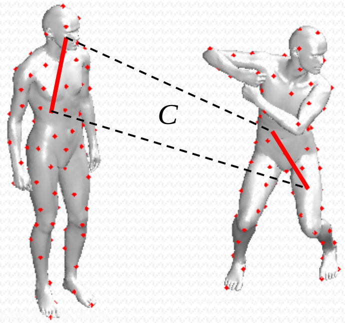

LAP

$$\min_{R} \sum_{i,j} \tilde{C}_{ij} R_{ij}$$

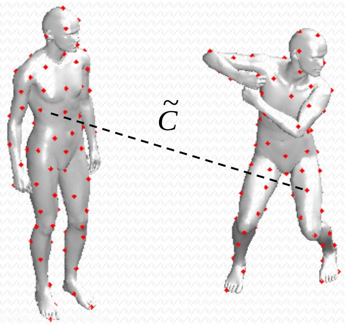

Mapping constraints for the LAP are the usual doubly-stochastic constraints on R

Suggested reading

The suggestions below follow closely the ideas we covered in this class.

- *On the use of Gromov-Hausdorff distances for shape comparison*. F.Mémoli. Proc. SGP 2007. Sections 4, 5, 8
- *A spectral technique for correspondence problems using pairwise constraints*. M.Leordeanu and M.Hebert. Proc. ICCV 2005.
- *A game-theoretic approach to deformable shape matching*. E.Rodolà et al. Proc. CVPR 2012.