Analysis of Three-Dimensional Shapes

(IN2238, TU München, Summer 2014)

Dr. Emanuele Rodolà

rodola@in.tum.de

Room 02.09.058, Informatik IX

07.04.2014

Computer Vision Group

4 Post-docs
14 PhD students
Master and bachelor students welcome!

Prof. Dr. Daniel Cremers

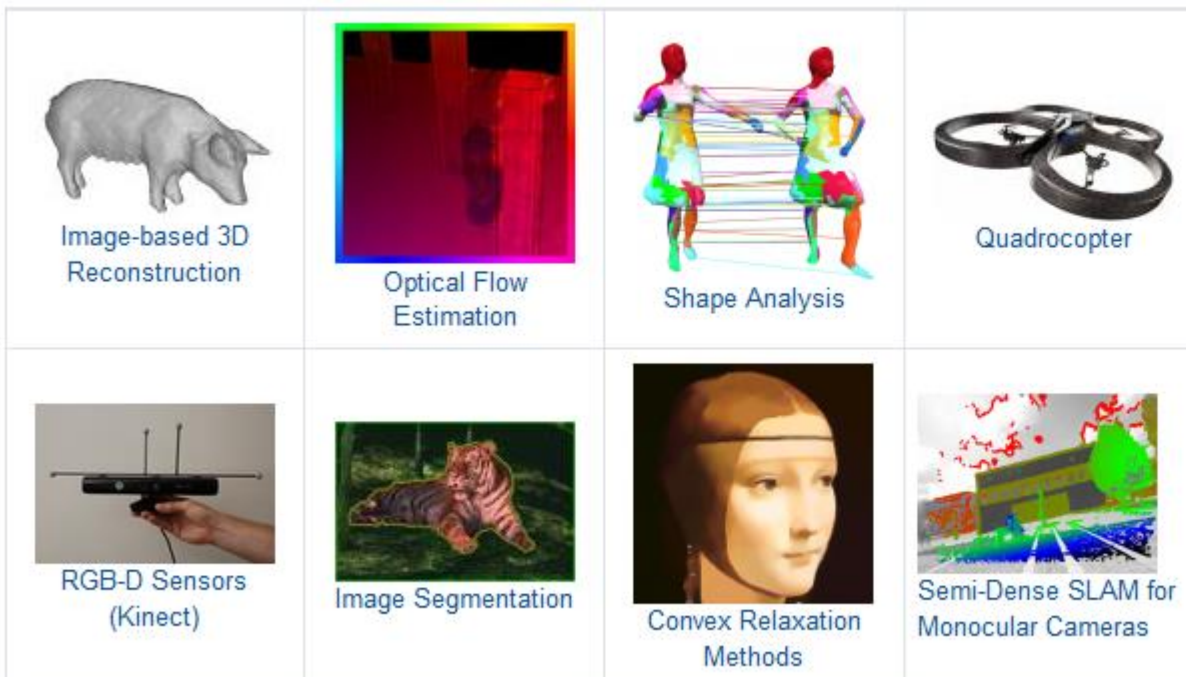

Formalities

- **Who?**

TA

Dr. Emanuele Rodolà

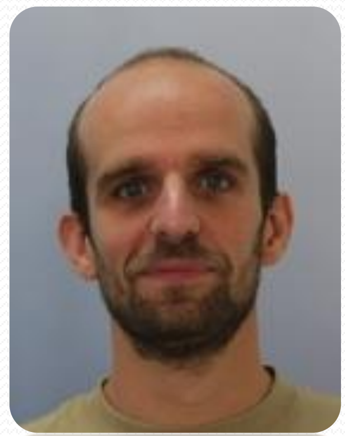

Thomas Windheuser

Matthias Vestner

- **Where?** Room 02.09.023, Informatik IX
- **When?** Mondays 10:00-12:00 *lecture*

  Tuesdays 14:00-15:00 *exercises*

Other formalities

- **Mathematical problems**
- **Programming exercises** (Matlab, C++)
- **Final exam** (written or oral or both)
- **Office hours:** send me an e-mail to set up a meeting

- **Textbook** (just a suggestion)

Bronstein, Bronstein, Kimmel. *Numerical geometry of non-rigid shapes*, Springer 2008

- **Scientific papers** will be suggested throughout the lecture

Topics

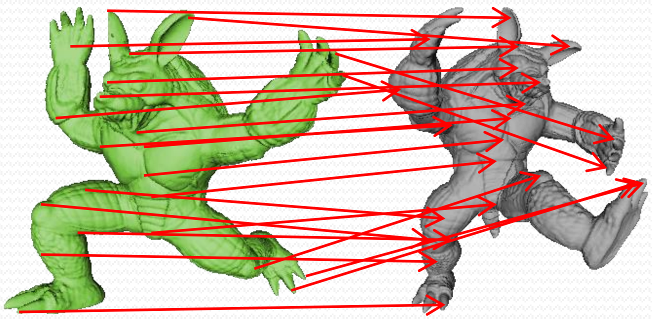

Correspondence

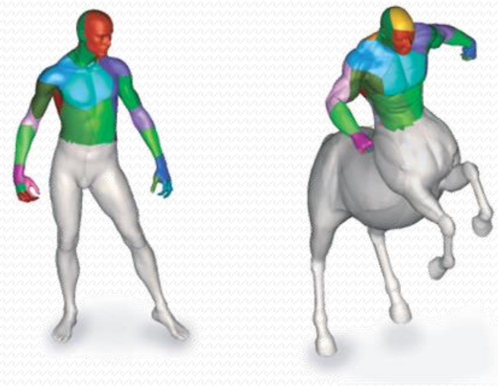

Partial similarity

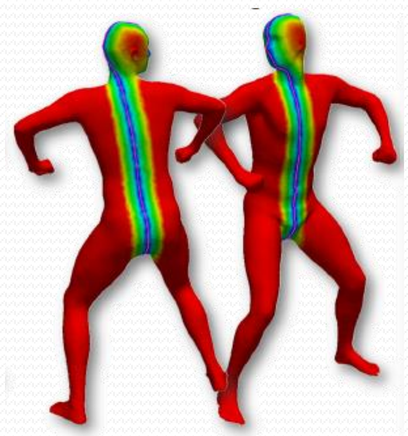

Symmetry

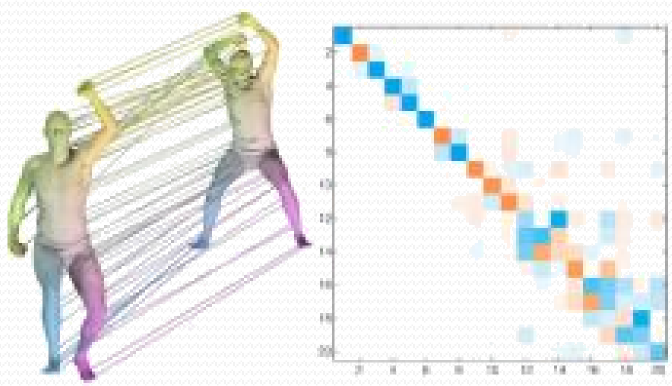

Representation

Topics

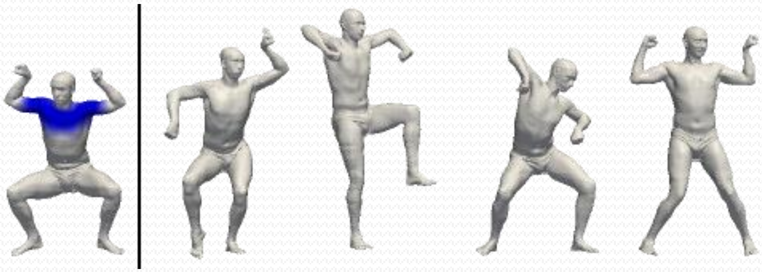

Analysis of shape collections

Feature detection

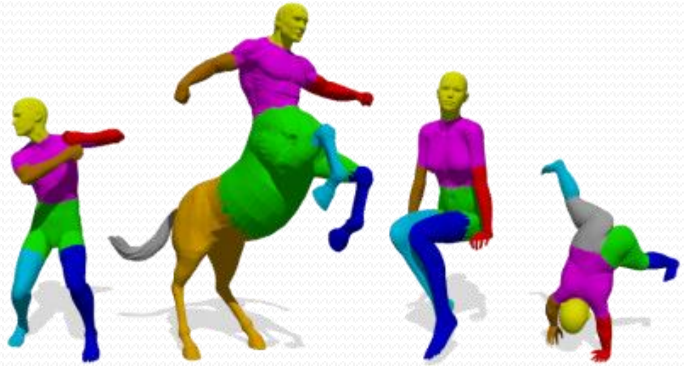

Segmentation

Description

Tools

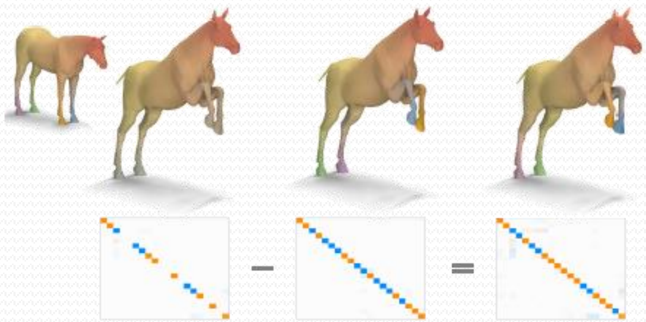

Linear algebra

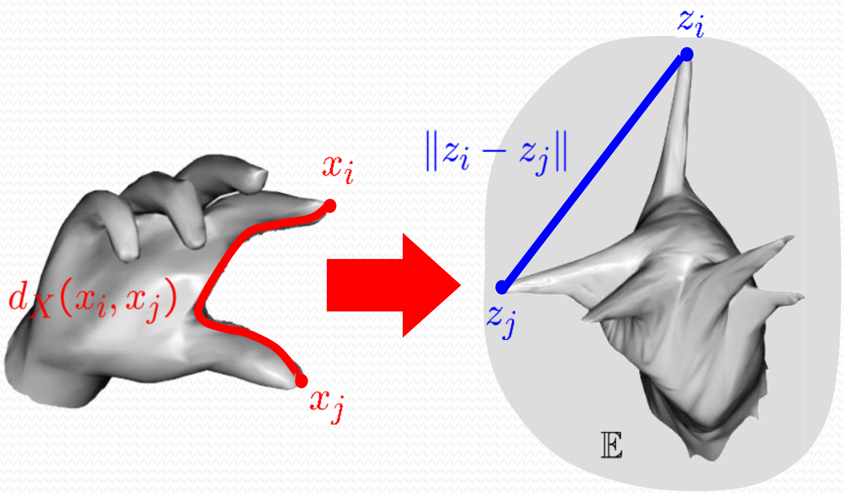

Metric spaces

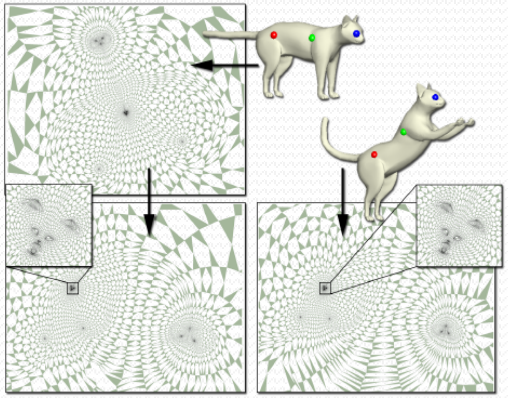

Conformal geometry

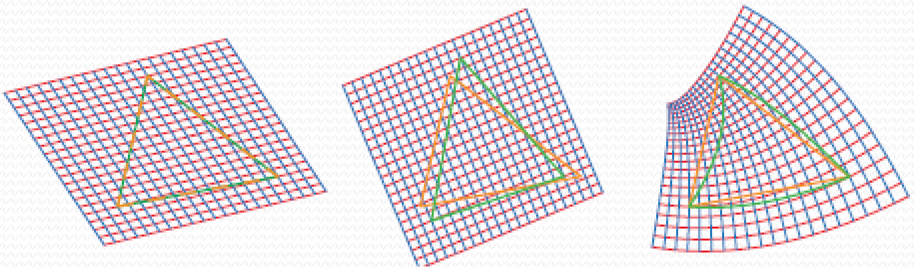

Differential geometry

Tools

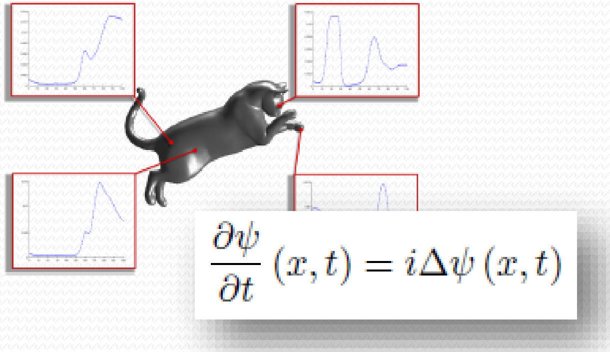
$$\frac{\partial\psi}{\partial t}(x,t) = i\Delta\psi(x,t)$$

PDEs

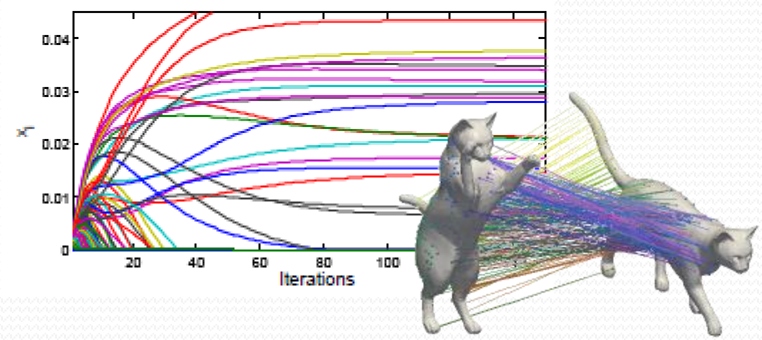

Optimization

Tools

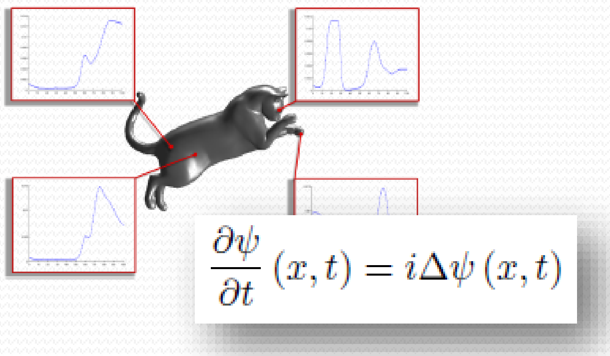

PDEs

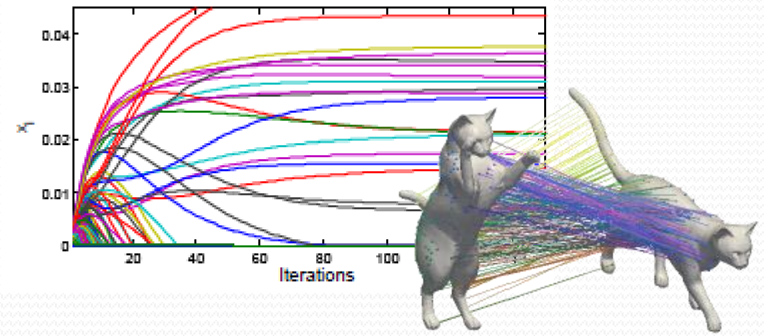

Optimization

Good news:

90% of the time we will be able to have a visualization of what we are doing!

Seminar

Recent Advances in the Analysis of 3D Shapes (IN2107)

When? Wednesdays, 14:00

Where? 02.09.023

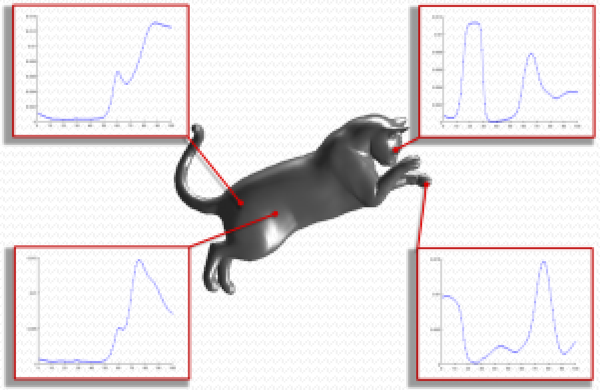

First meeting: Apr 16, 14:00

Topic: Laplace-Beltrami Operator on manifolds

What is a shape?

"There can be no such thing as a mathematical theory of shape. The very notion of shape belongs to the natural sciences."

J. Koenderink. *Solid Shape*. MIT Press 1990.

What is a shape?

- Proteins
- Molecules
- 2D Images
- **3D models** (coming from a 3D scanner)
- **3D models** (coming from CAD software)
- Volumetric models (medical imaging)
- More complicated structures (things you can't even visualize)

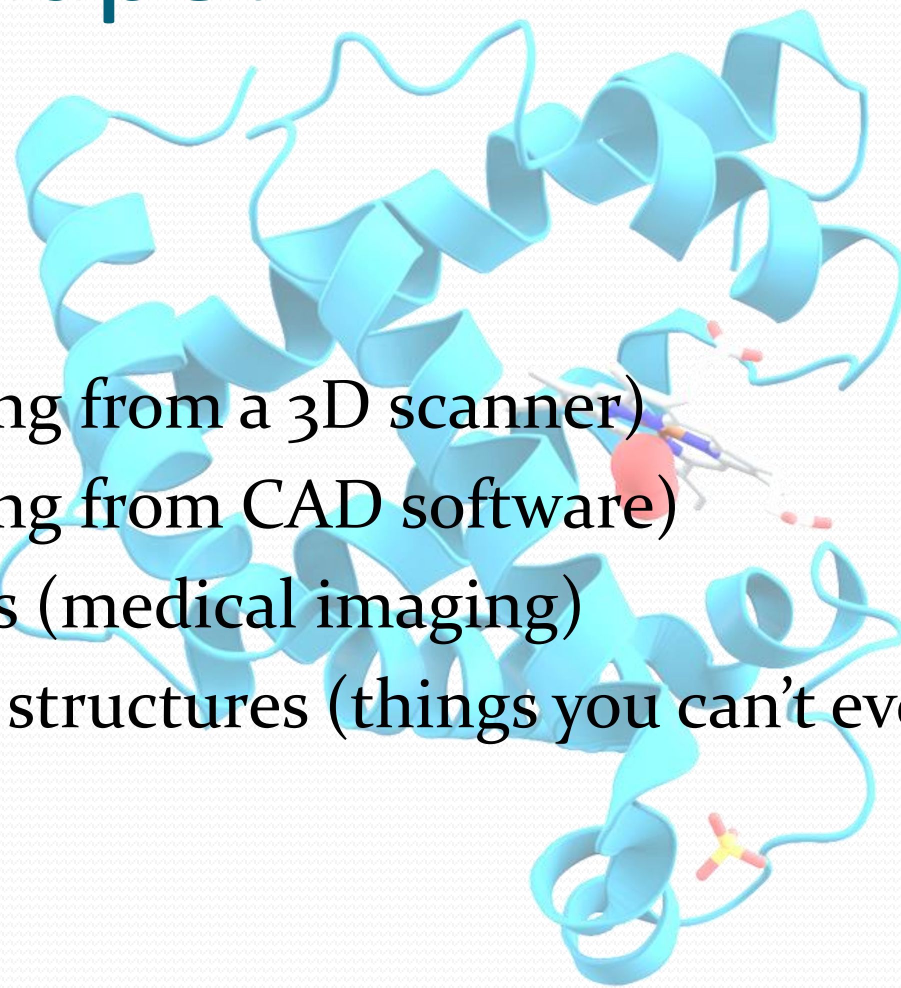

Shapes vs images: domain

Euclidean
(flat)

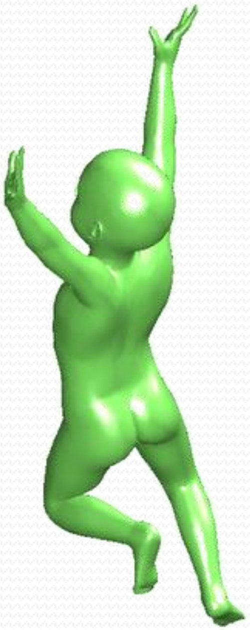

Non-Euclidean
(curved)

Shapes vs images: representation

Array of pixels
(uniform grid)

Splines

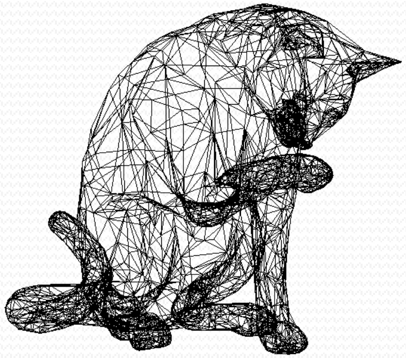

Graph

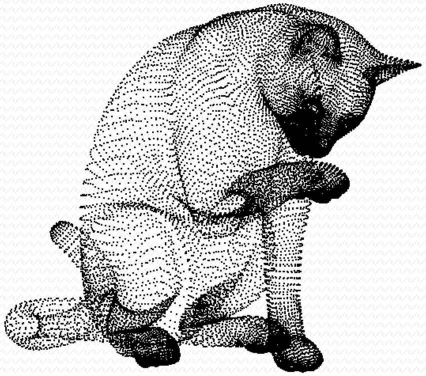

Point cloud

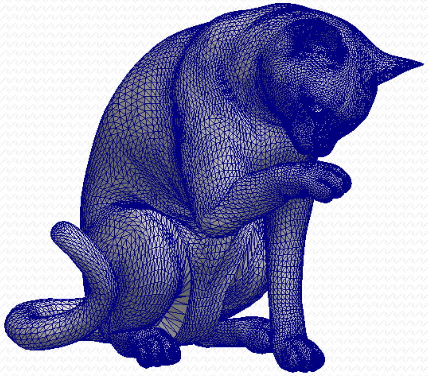

Triangular mesh

Shapes vs images: parametrization

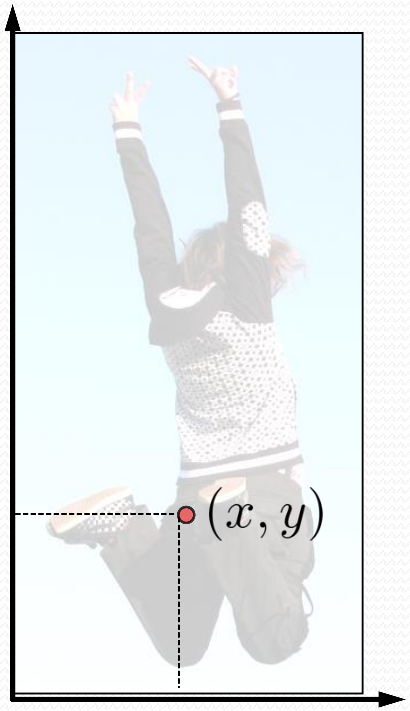

Global

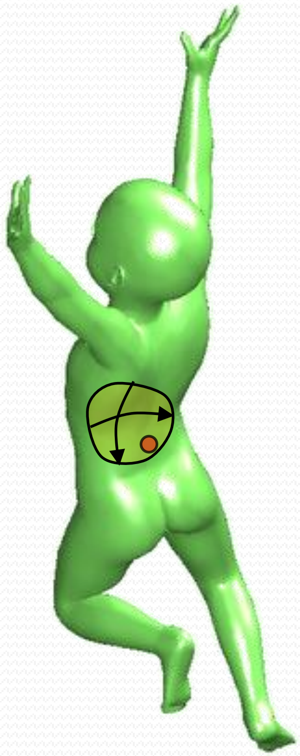

Local

Shapes vs images: sampling

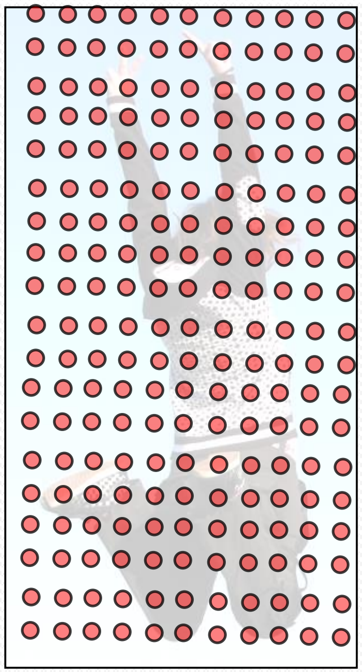

Uniform

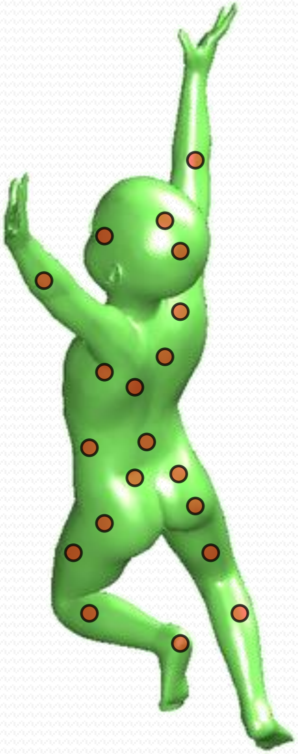

“Uniform” is not well-defined

Shapes vs images: transformations

Perspective

Affine

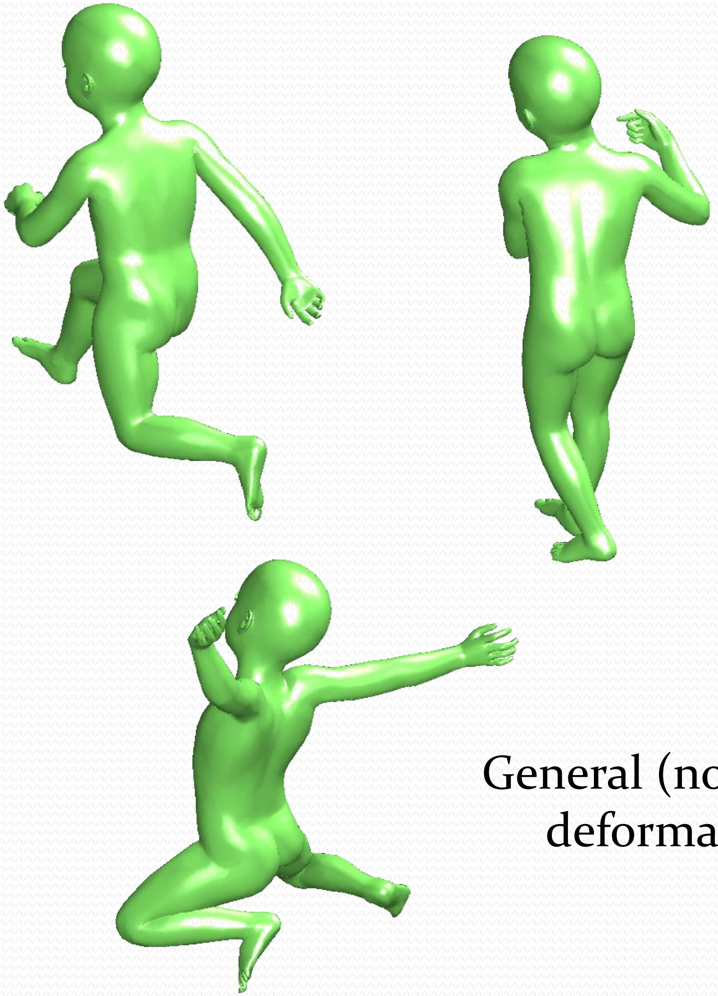

General (non-rigid) deformations

Shapes vs images: calculus

$$\frac{1}{2}$$ 

Shapes vs images: calculus

$$\frac{1}{2}$$  $$+\frac{1}{2}$$  $$=$$

Shapes vs images: calculus

$\frac{1}{2}$  $+\frac{1}{2}$  $=$ 

Shapes vs images: calculus

$\frac{1}{2}$  $+\frac{1}{2}$  $=$ 

$\frac{1}{2}$  $+\frac{1}{2}$  $= ?$

Shape similarity

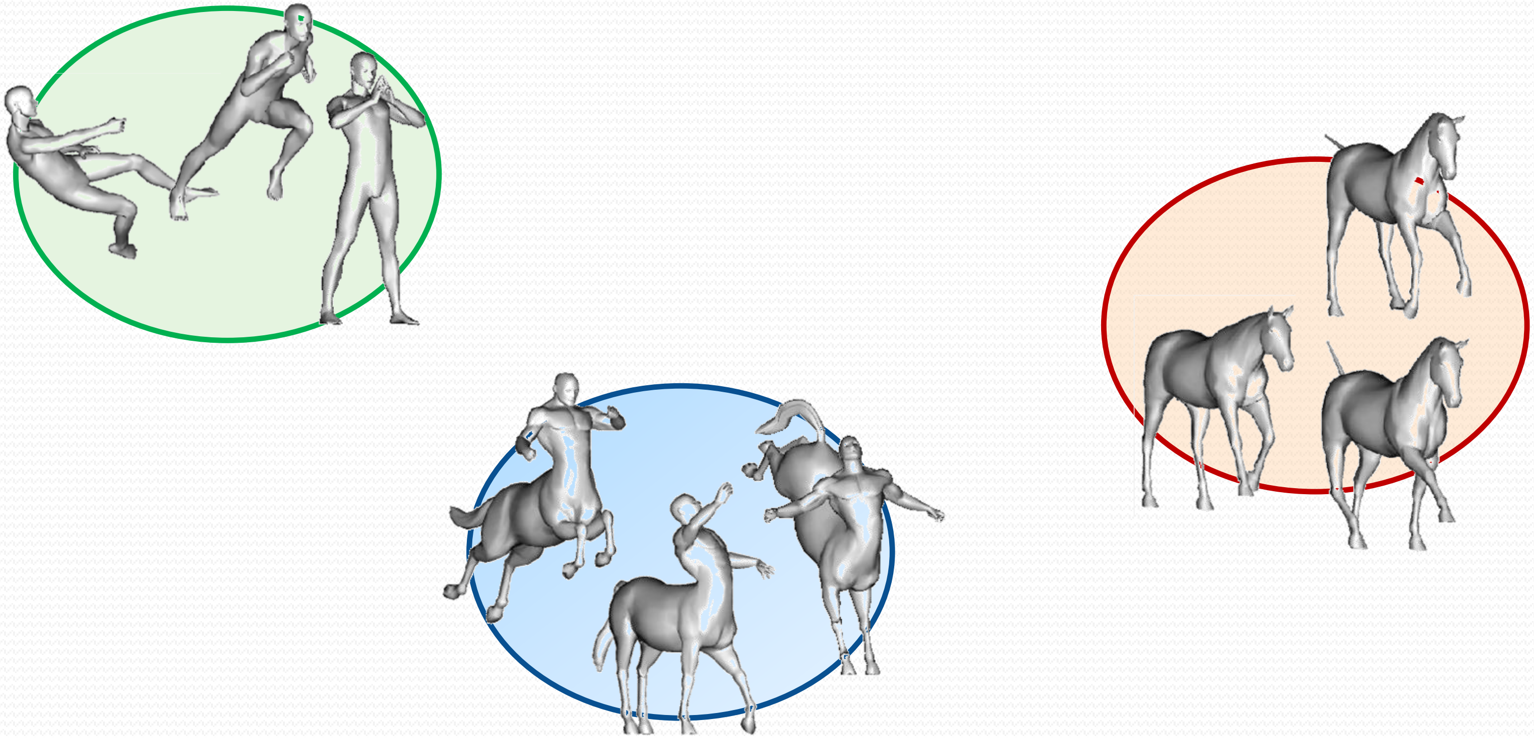

Is there something like a “space of shapes”?

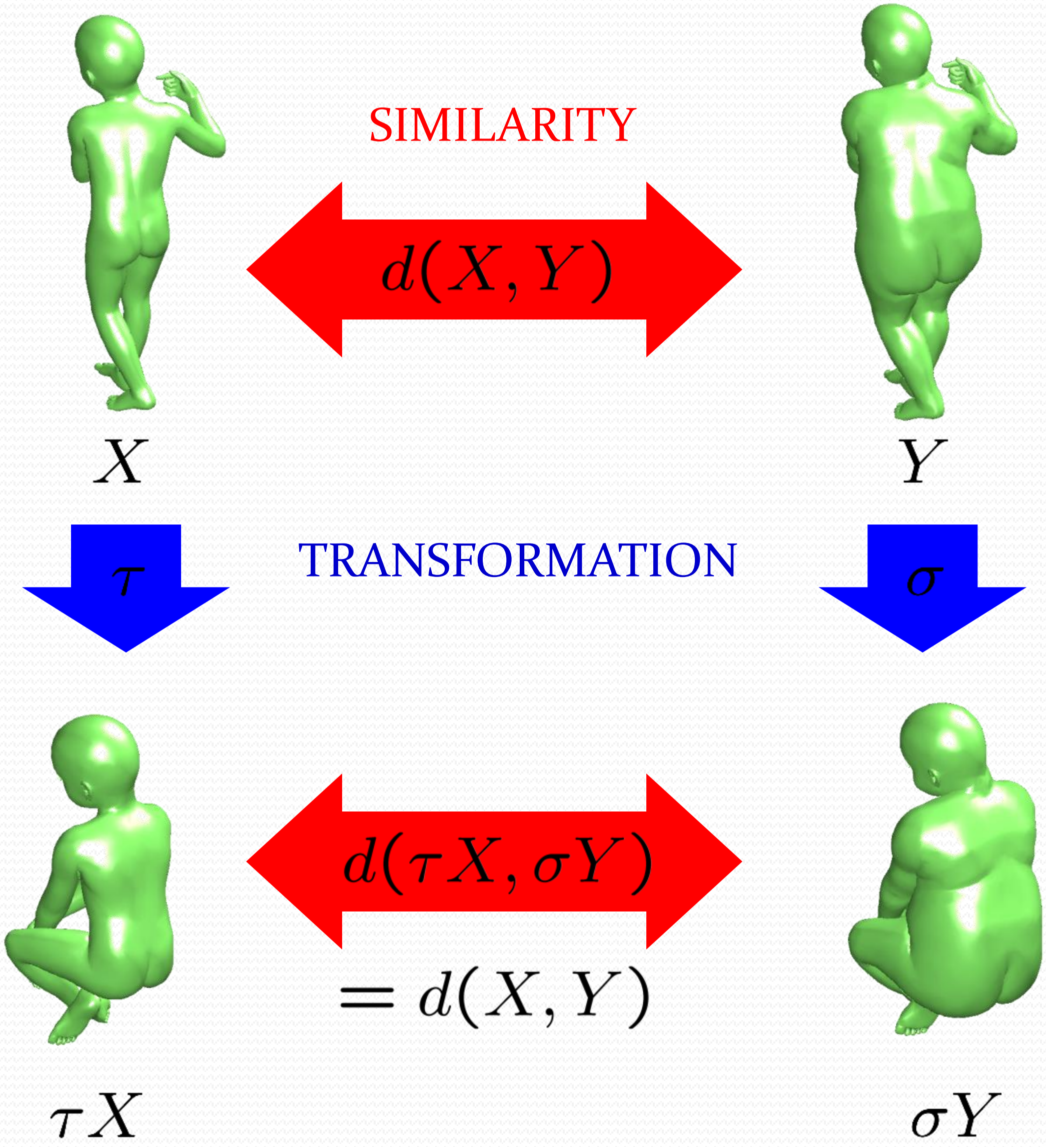
SIMILARITY
$d(X,Y)$
X
Y
TRANSFORMATION
τ
σ
$d(\tau X,\sigma Y)$
$= d(X,Y)$
τX
σY

Shape matching

- Given a pair of shapes, let's try to find a **correspondence** between them.

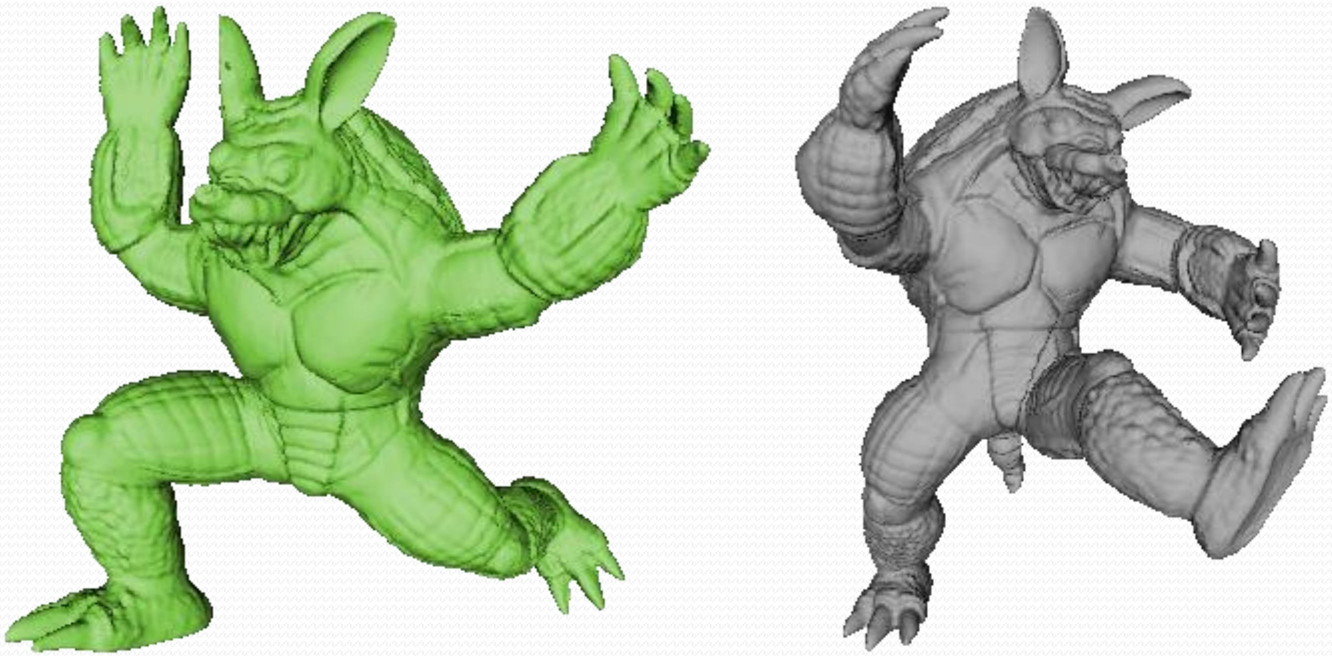

Shape matching

- Find the **best** alignment/map/correspondence.

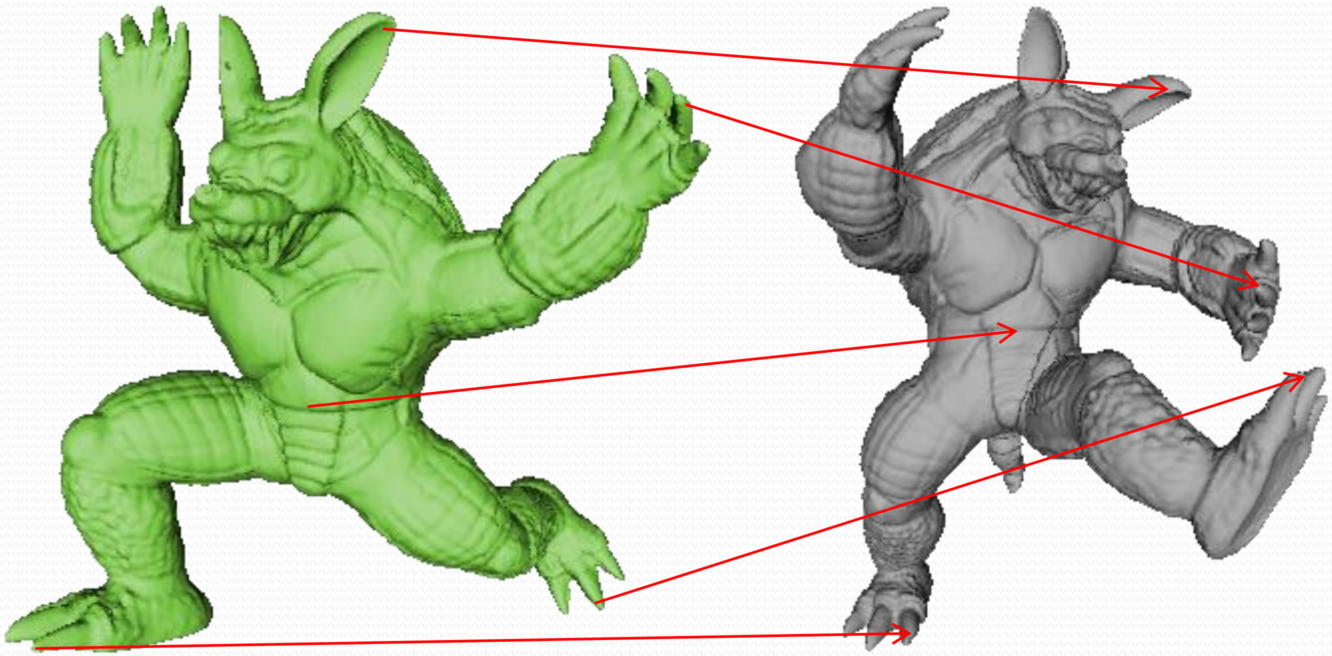

Shape matching

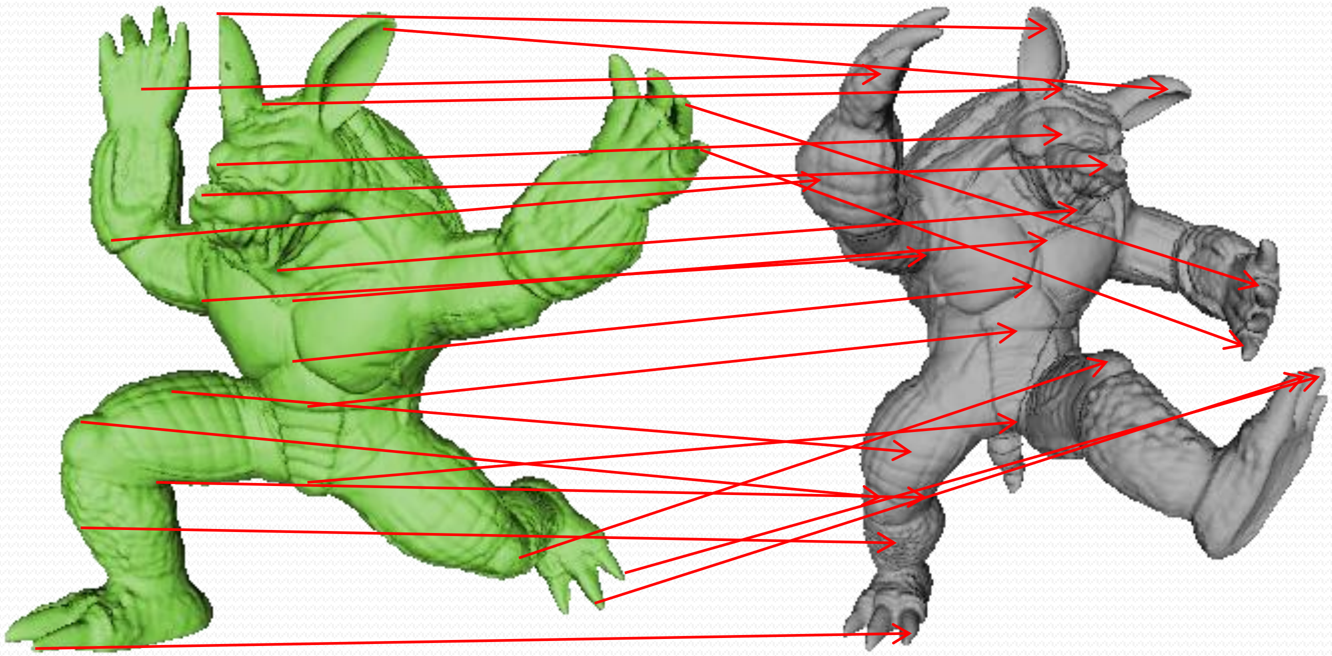

In the real world

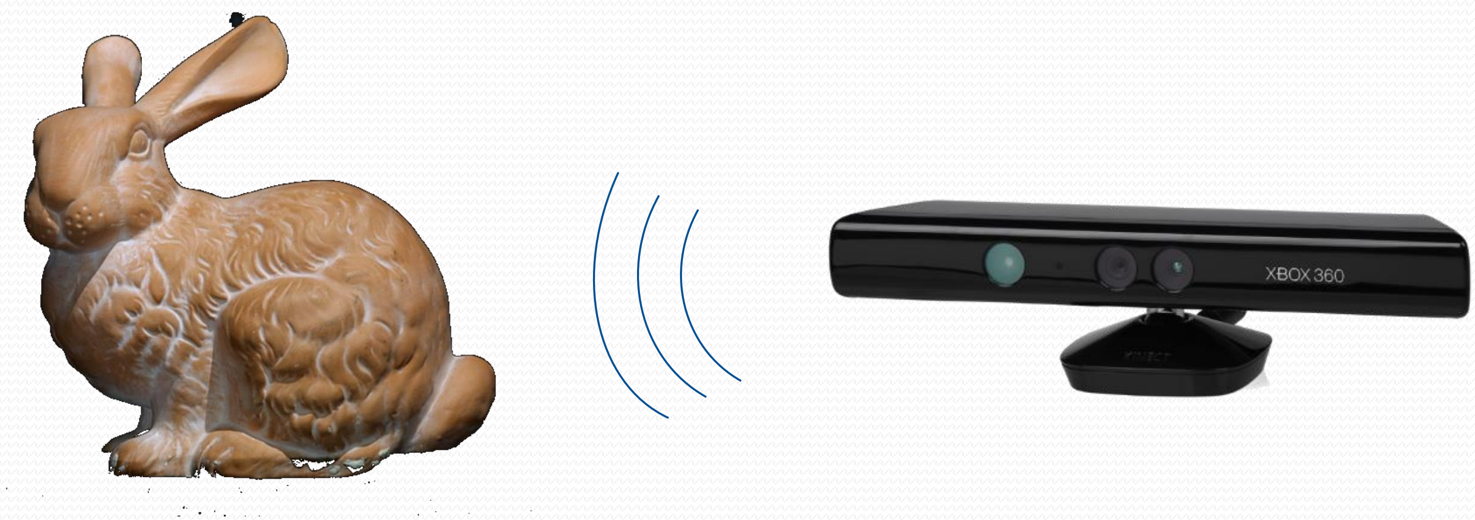

In the real world

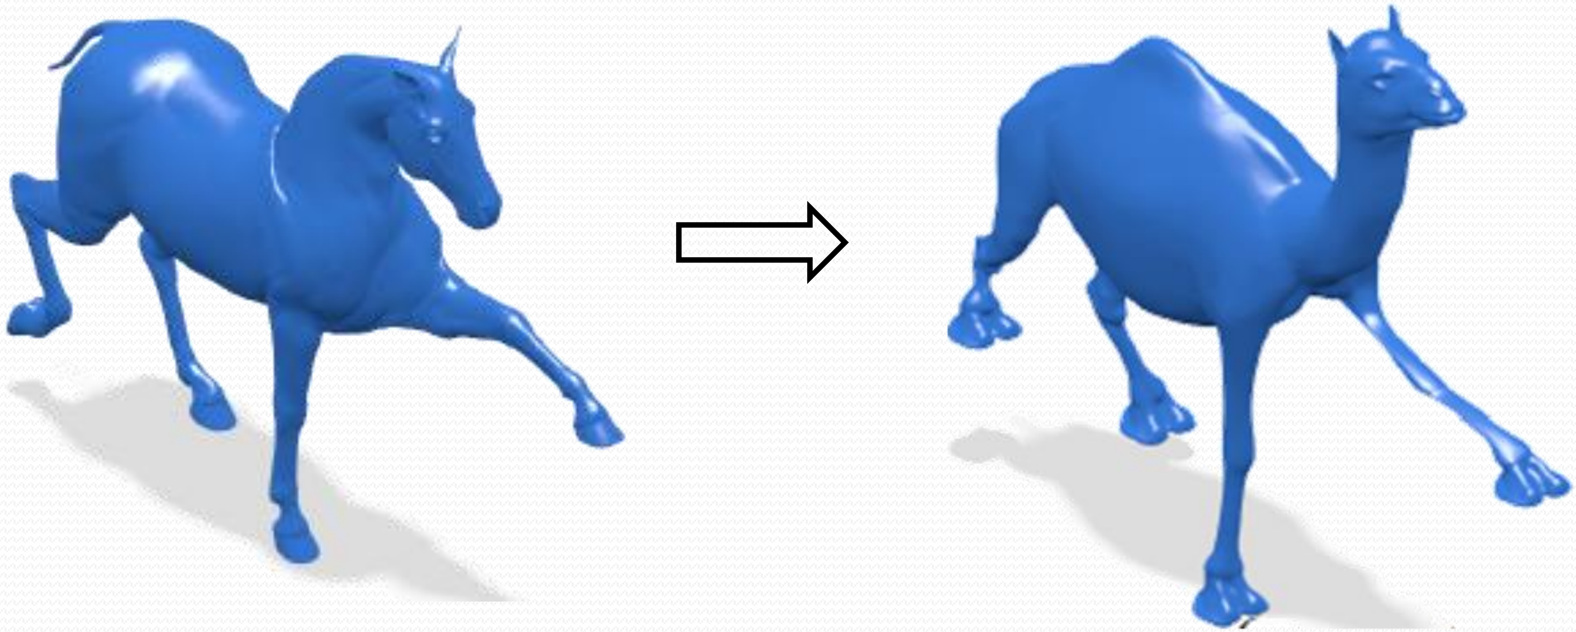

Taxonomy

Local refinement (e.g. ICP)	vs.	**Global** alignment (search)
Rigid rotation, translation	vs.	**Deformable** general deformation
Pair two shapes	vs.	**Collection** multiple shapes

Pairwise rigid correspondence

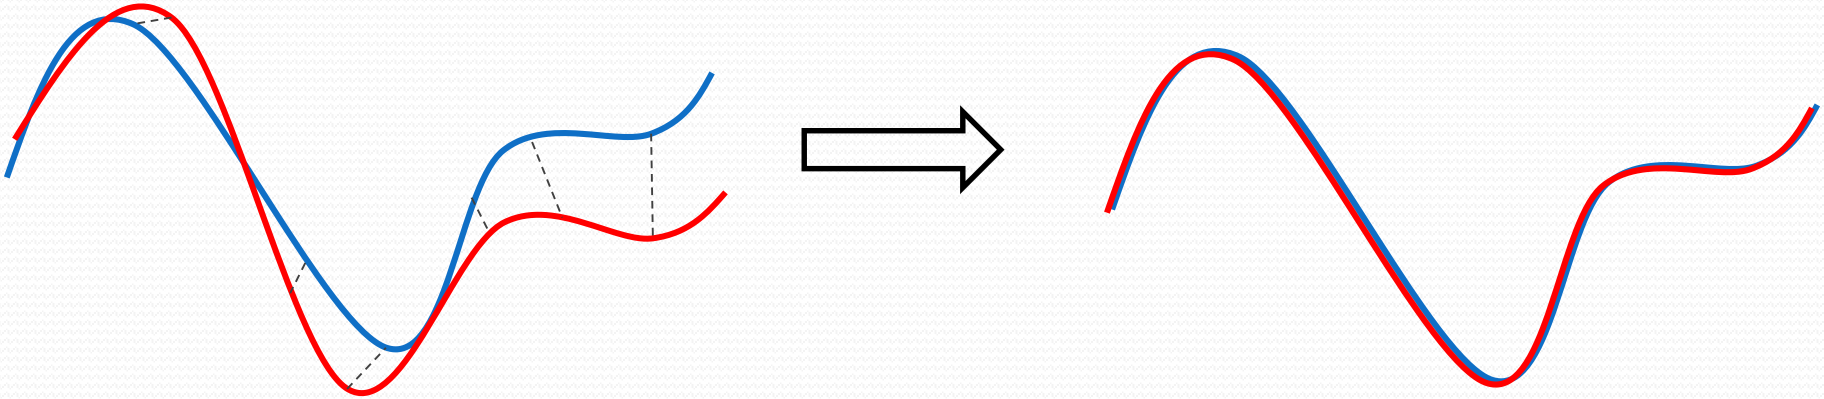

Iterative Closest Point

For a given pair of shapes *M* and *N*, **iterate**:

1. For each $x_i \in M$ find its nearest neighbor $y_i \in N$
2. Find the deformation *R*, *t* minimizing:

$$\sum_{x_i \in M} \left\| R x_i + t - y_i \right\|$$

Pairwise rigid correspondence

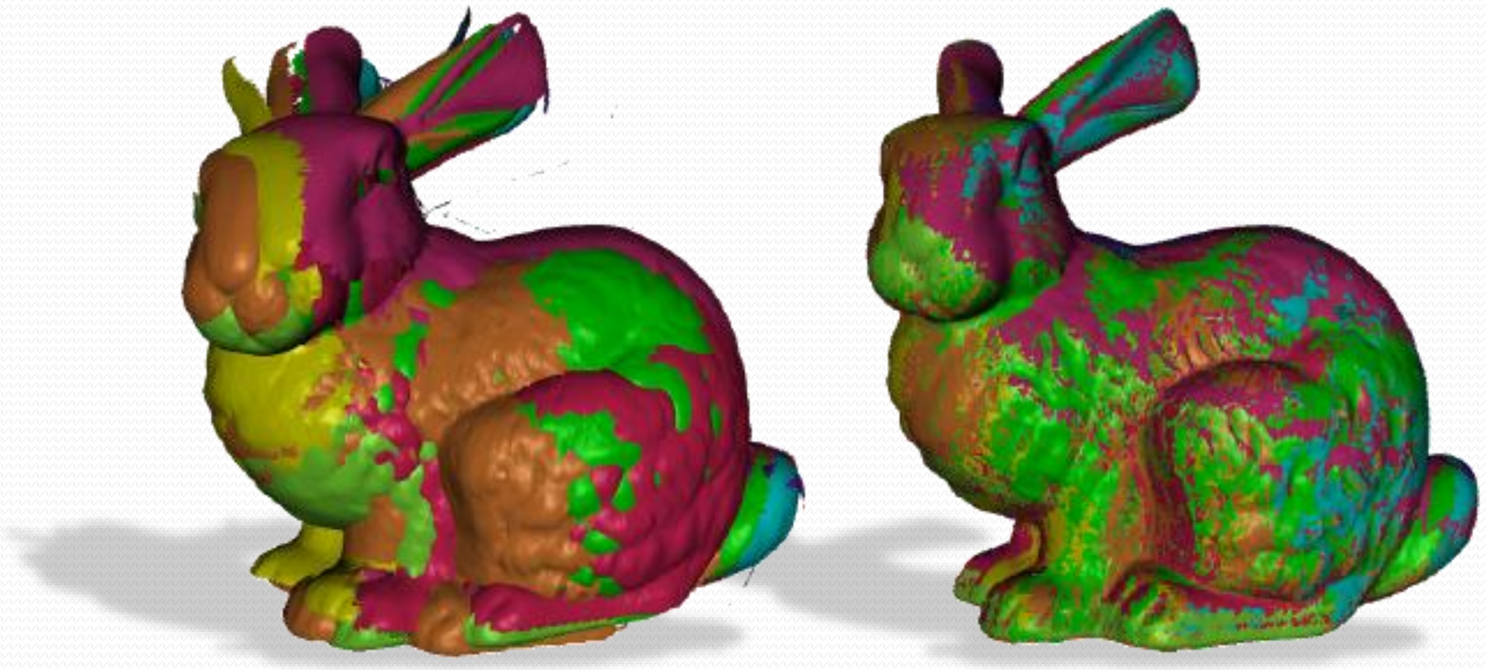

Taxonomy

Local refinement (e.g. ICP)	vs.	**Global** alignment (search)
Rigid rotation, translation	vs.	**Deformable** general deformation
Pair two shapes	vs.	**Collection** multiple shapes

Taxonomy

Local refinement (e.g. ICP)	vs.	**Global** alignment (search)
Rigid rotation, translation	vs.	**Deformable** general deformation
Pair two shapes	vs.	**Collection** multiple shapes

Pairwise rigid correspondence

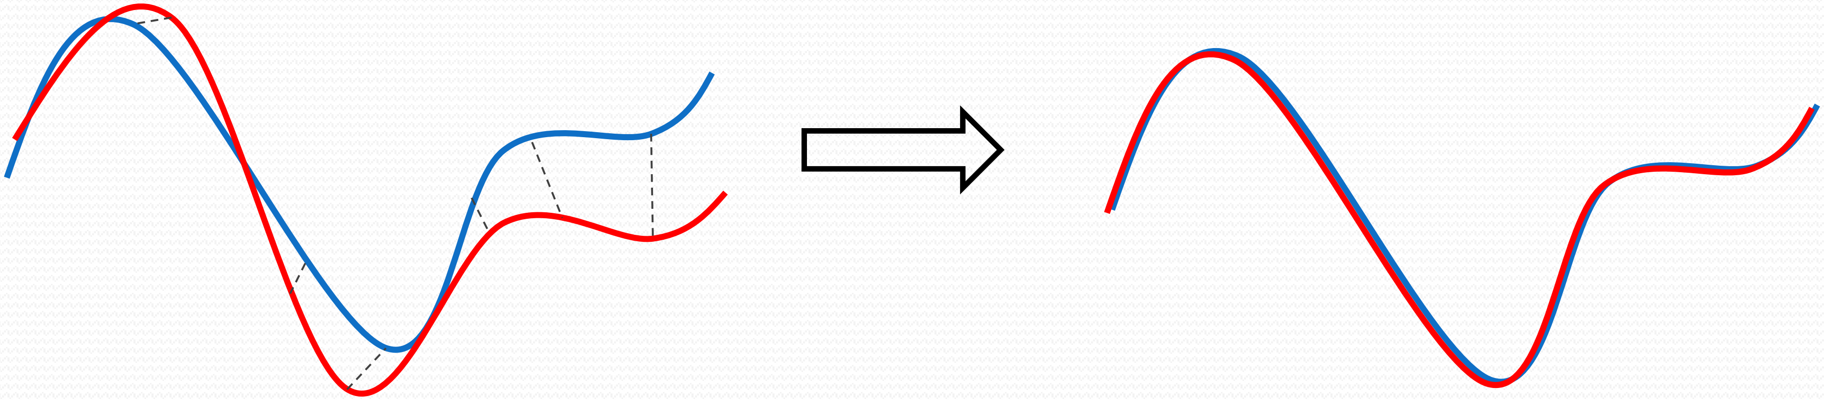

Iterative Closest Point

…

…

1. **Find the deformation *R*, *t* minimizing:**

$$\sum_{x_i \in M} \left\| R x_i + t - y_i \right\|$$

Deformable shape matching

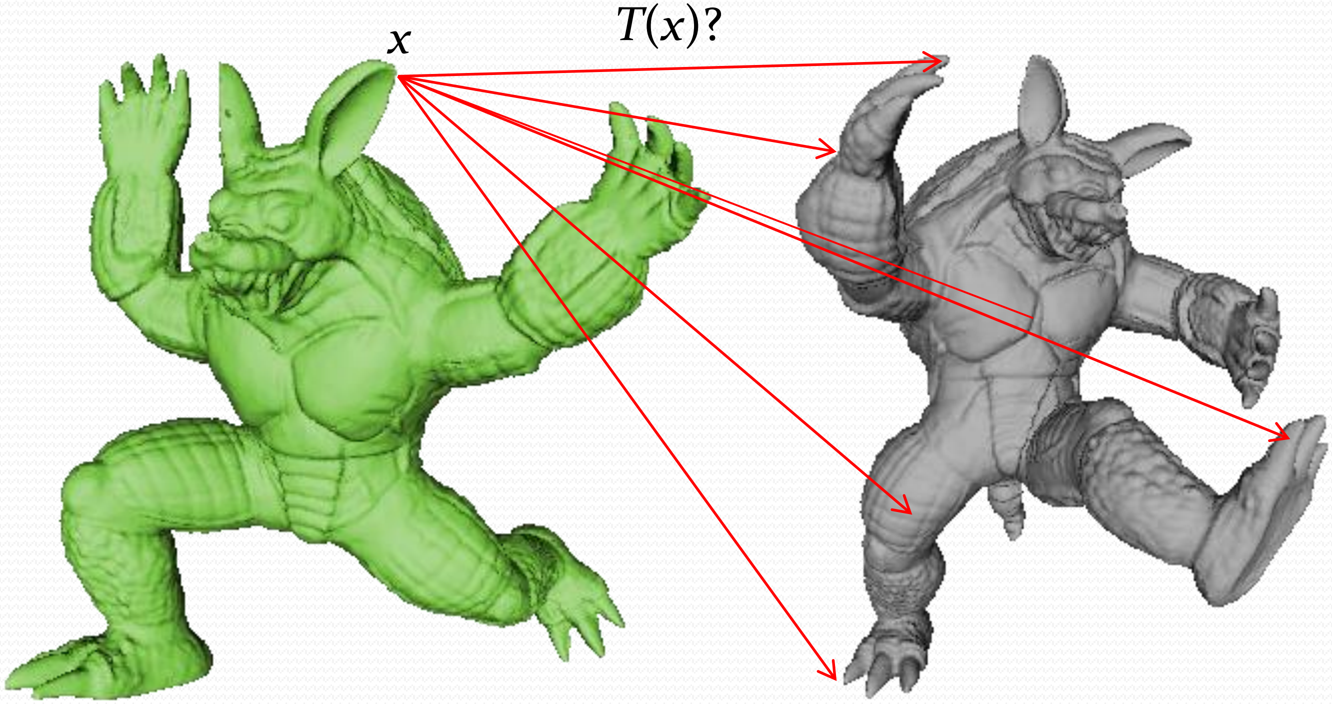

- Unlike rigid matching (rotation/translation), there is no *compact representation* to optimize for.

Deformable shape matching

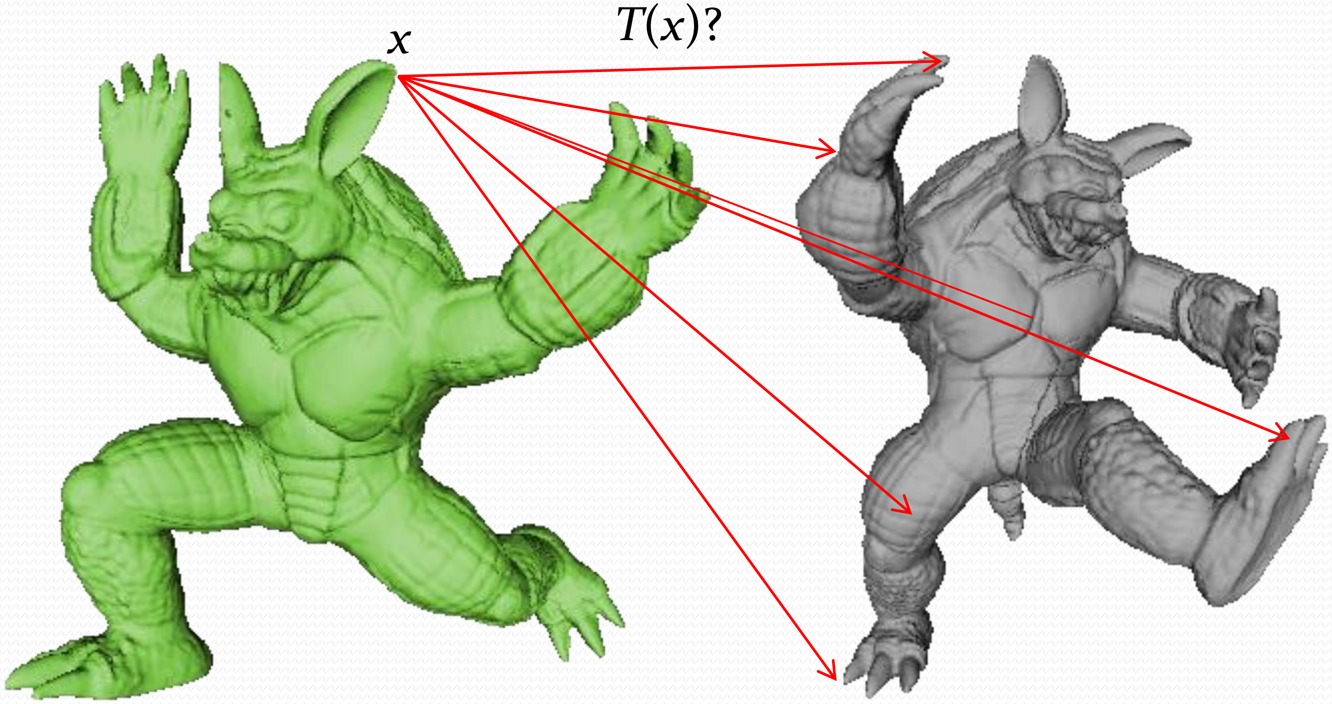

- Instead, directly optimize over all possible point-to-point correspondences.

Signature preservation

$$T_{opt} = \underset{T:M \to N}{\arg\min} \sum_{x_i \in M} \left\| S(x_i) - S(T(x_i)) \right\|$$

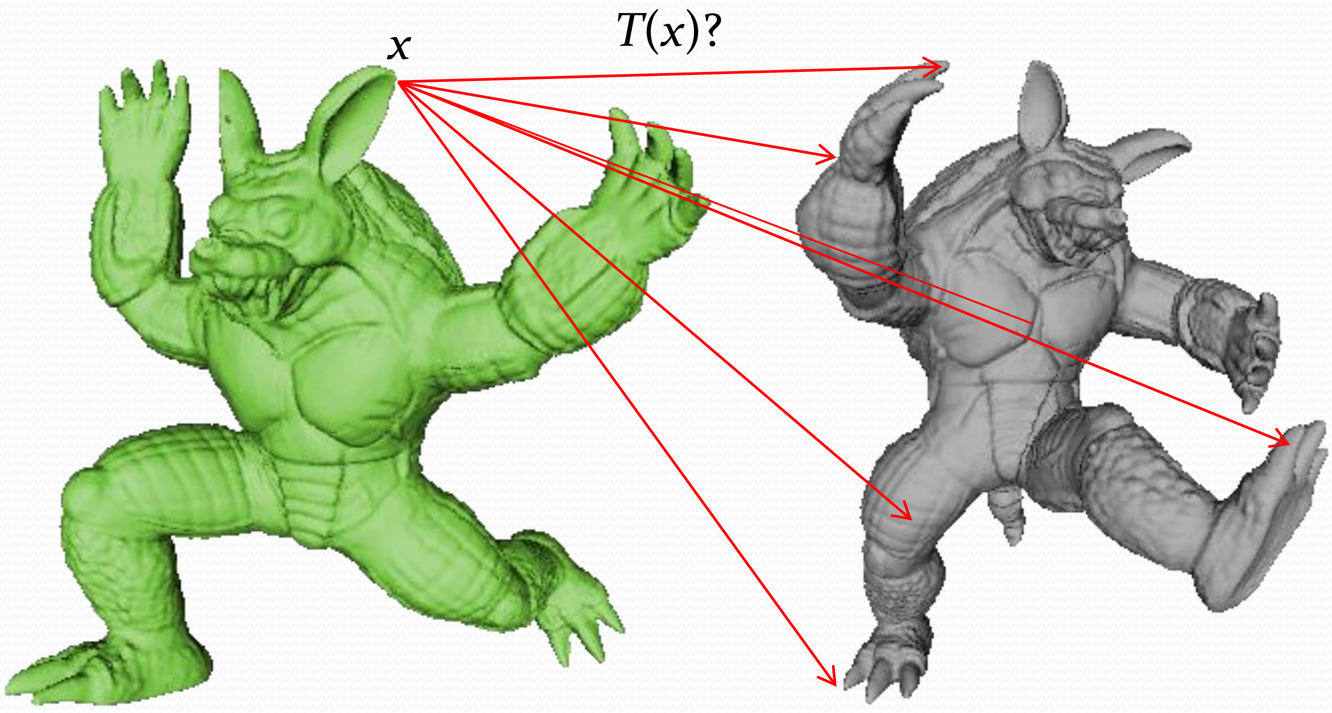

What signature?

One possibility: Look for similar textures

What signature?

Another possibility: Let's look at the geometry!

$$\left(\Delta_X + \frac{\partial}{\partial t}\right) u = 0$$

Heat equation governs the diffusion of heat on manifold X over time

Heat diffusion on manifolds

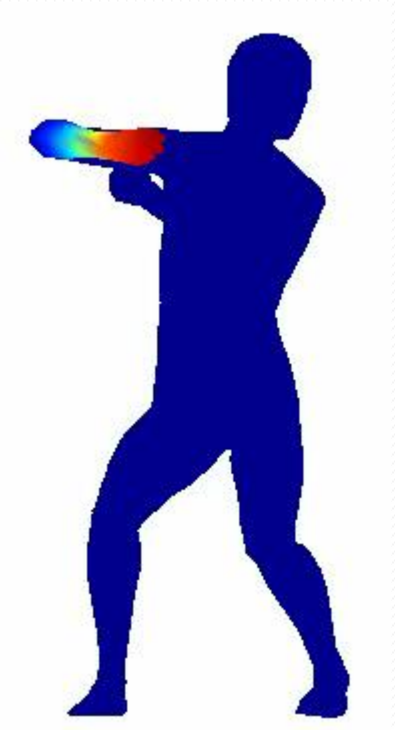

Heat Kernel Signature

Robust to pose variations

Signature preservation

$$T_{opt} = \arg\min_{T:M \to N} \sum_{x_i \in M} \left\| S(x_i) - S(T(x_i)) \right\|$$

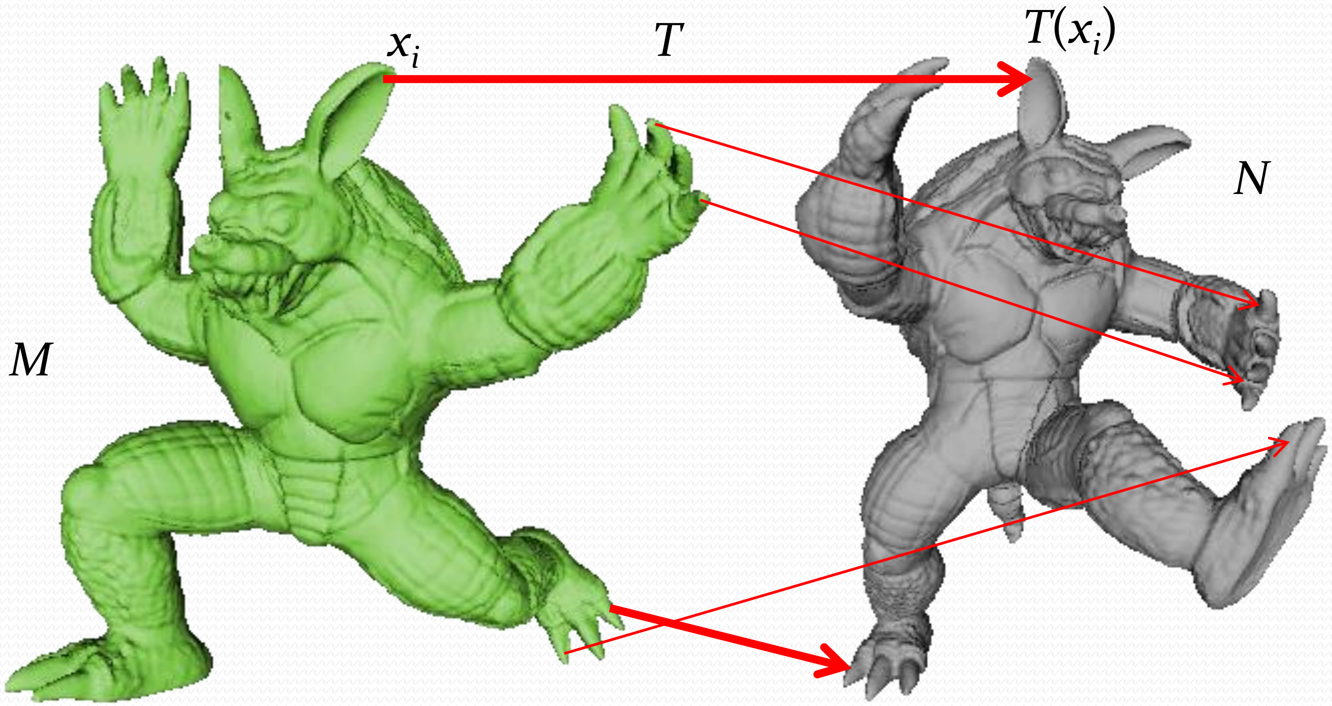

Metric preservation

$$T_{opt} = \arg\min_{T:M \to N} \sum_{(x_i, x_j) \in M \times M} \left\| d_M(x_i, x_j) - d_N(T(x_i), T(x_j)) \right\|$$

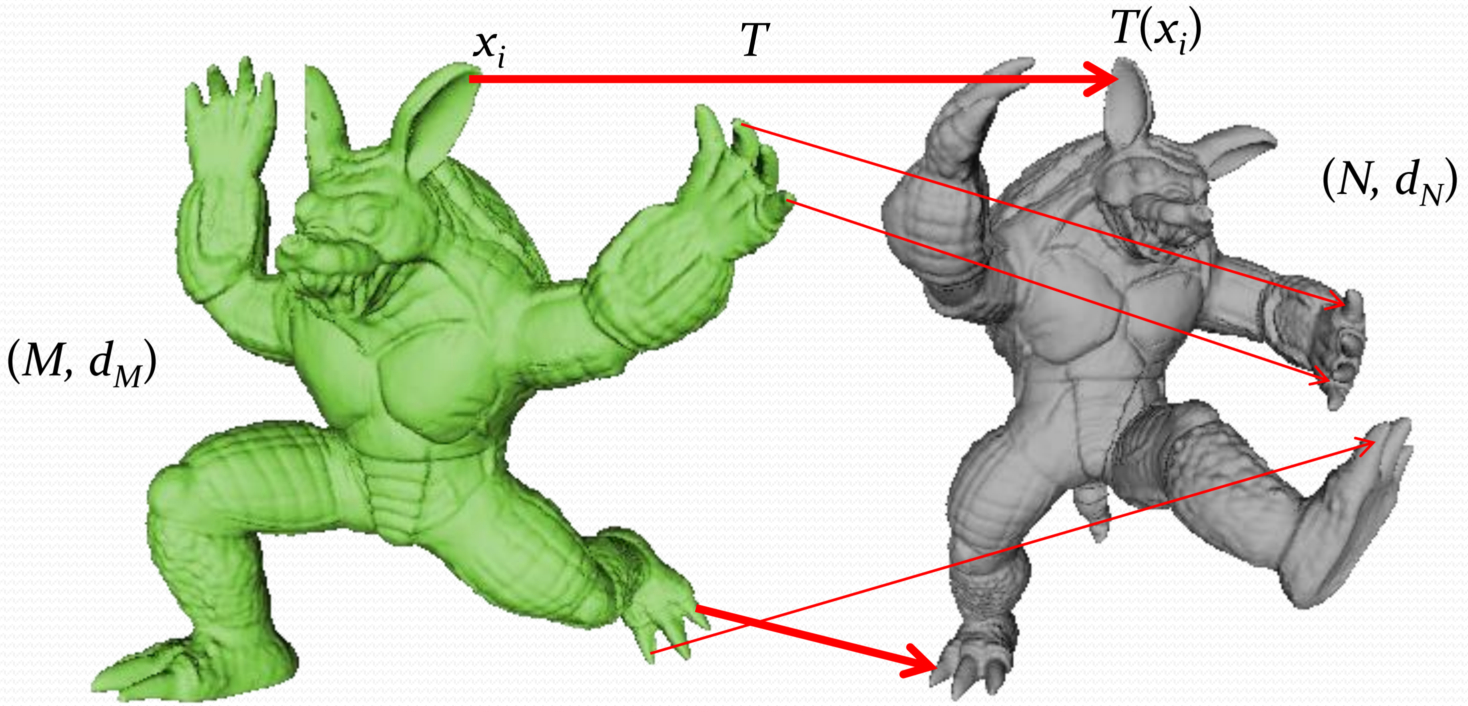

Metric preservation

$$T_{opt} = \arg\min_{T:M \rightarrow N} \sum_{(x_i, x_j) \in M \times M} \left\| d_M(x_i, x_j) - d_N(T(x_i), T(x_j)) \right\|$$

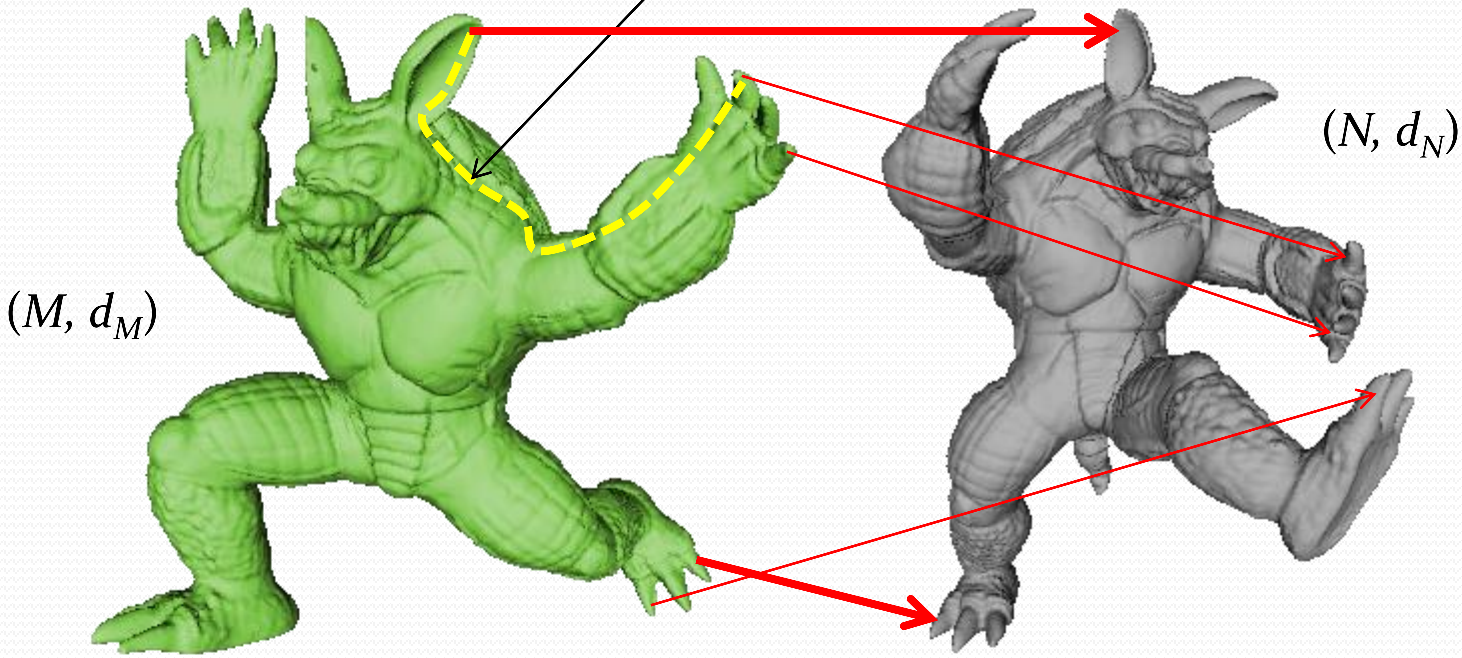

Metric preservation

$$T_{opt} = \arg\min_{T:M\to N} \sum_{(x_i, x_j)\in M\times M} \left\| d_M(x_i, x_j) - d_N(T(x_i), T(x_j)) \right\|$$

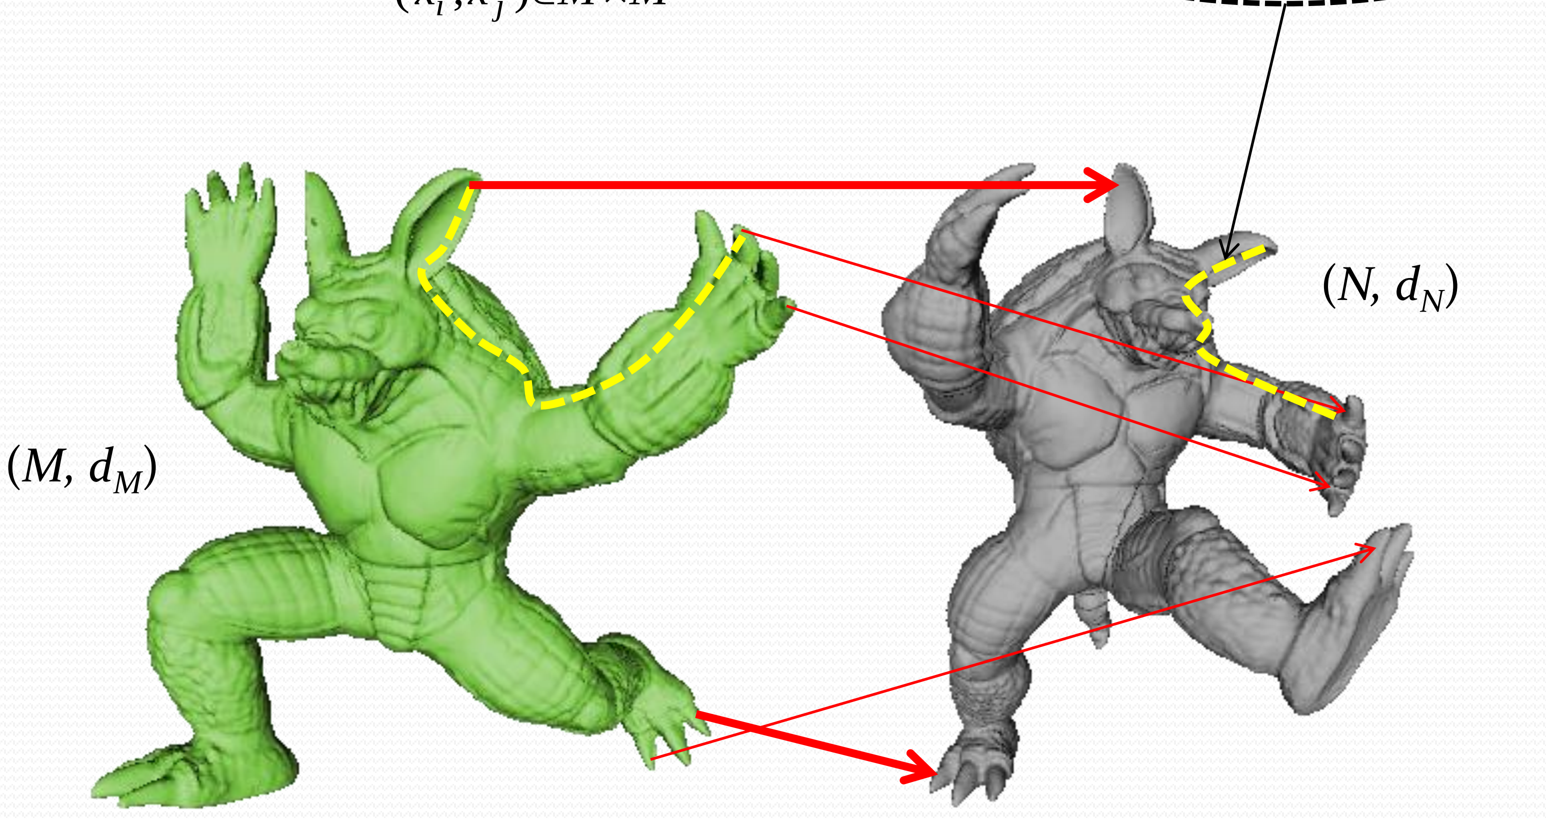

Metric preservation

$$T_{opt} = \underset{T:M \to N}{\arg\min} \sum_{(x_i, x_j) \in M \times M} \left\| d_M(x_i, x_j) - d_N(T(x_i), T(x_j)) \right\|$$

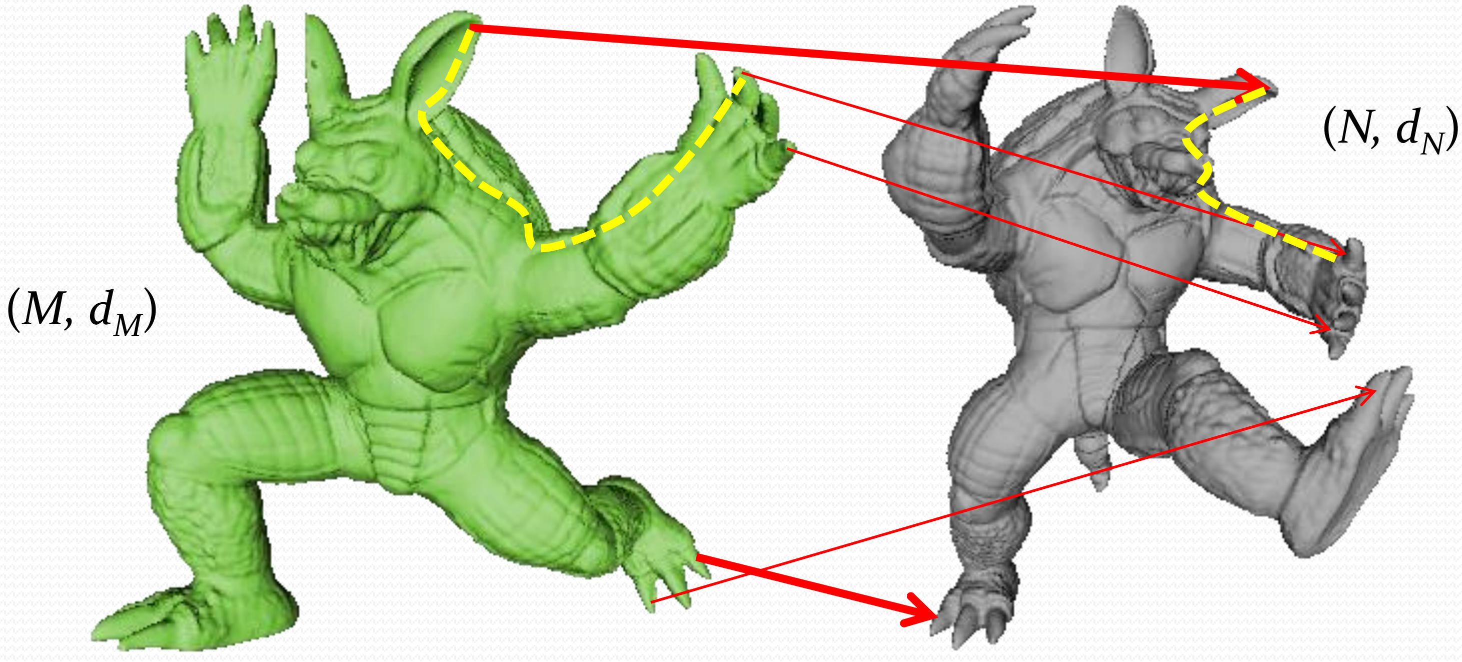

Metric preservation

$$T_{opt} = \arg\min_{T:M \to N} \sum_{(x_i, x_j) \in M \times M} \left\| d_M(x_i, x_j) - d_N(T(x_i), T(x_j)) \right\|$$

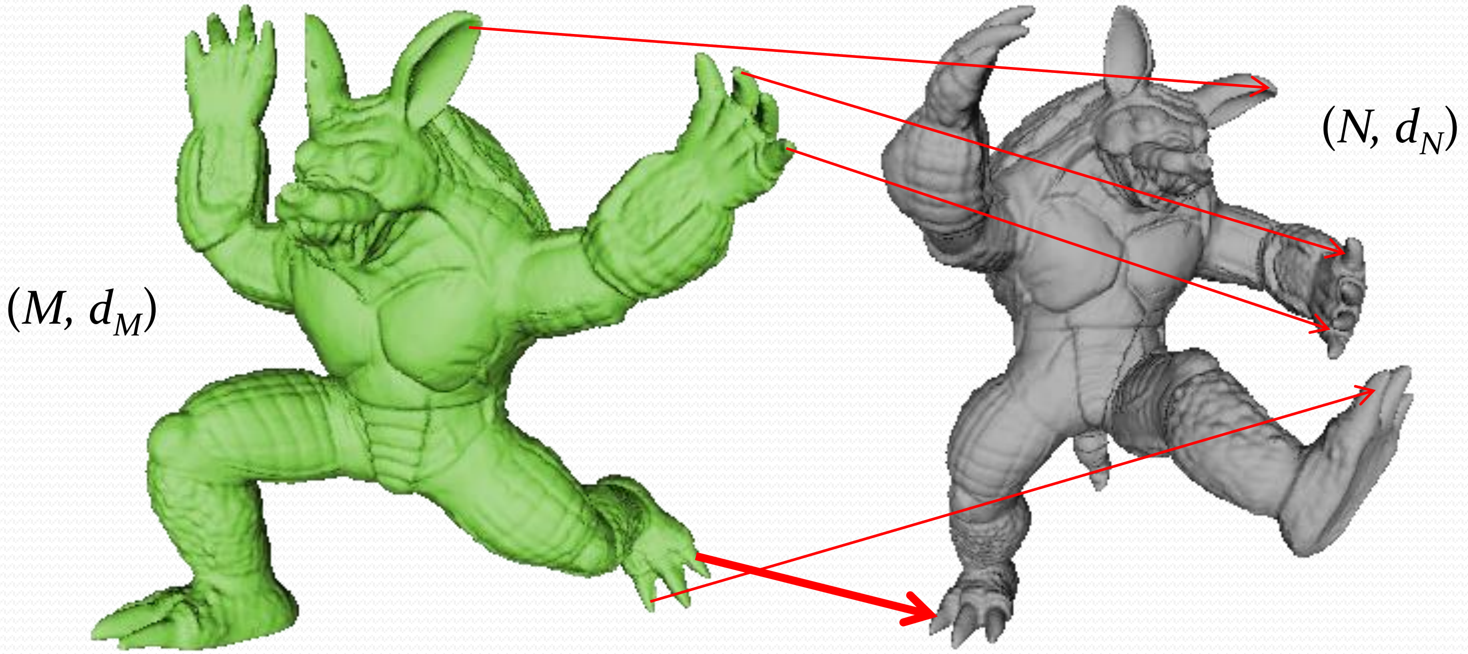

Metric preservation

$$T_{opt} = \arg\min_{T:M \to N} \sum_{(x_i, x_j) \in M \times M} \left\| d_M(x_i, x_j) - d_N(T(x_i), T(x_j)) \right\|$$

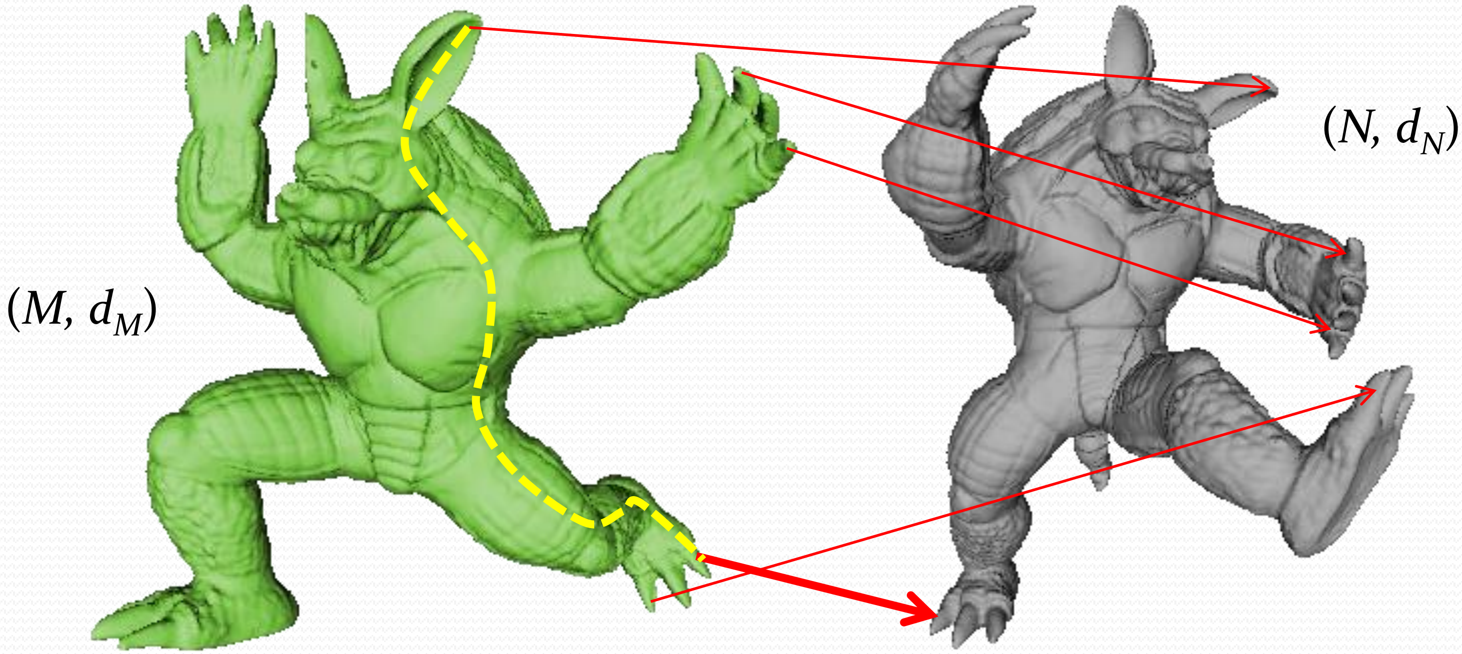

Metric preservation

$$T_{opt} = \underset{T:M \rightarrow N}{\arg\min} \sum_{(x_i, x_j) \in M \times M} \left\| d_M(x_i, x_j) - d_N(T(x_i), T(x_j)) \right\|$$

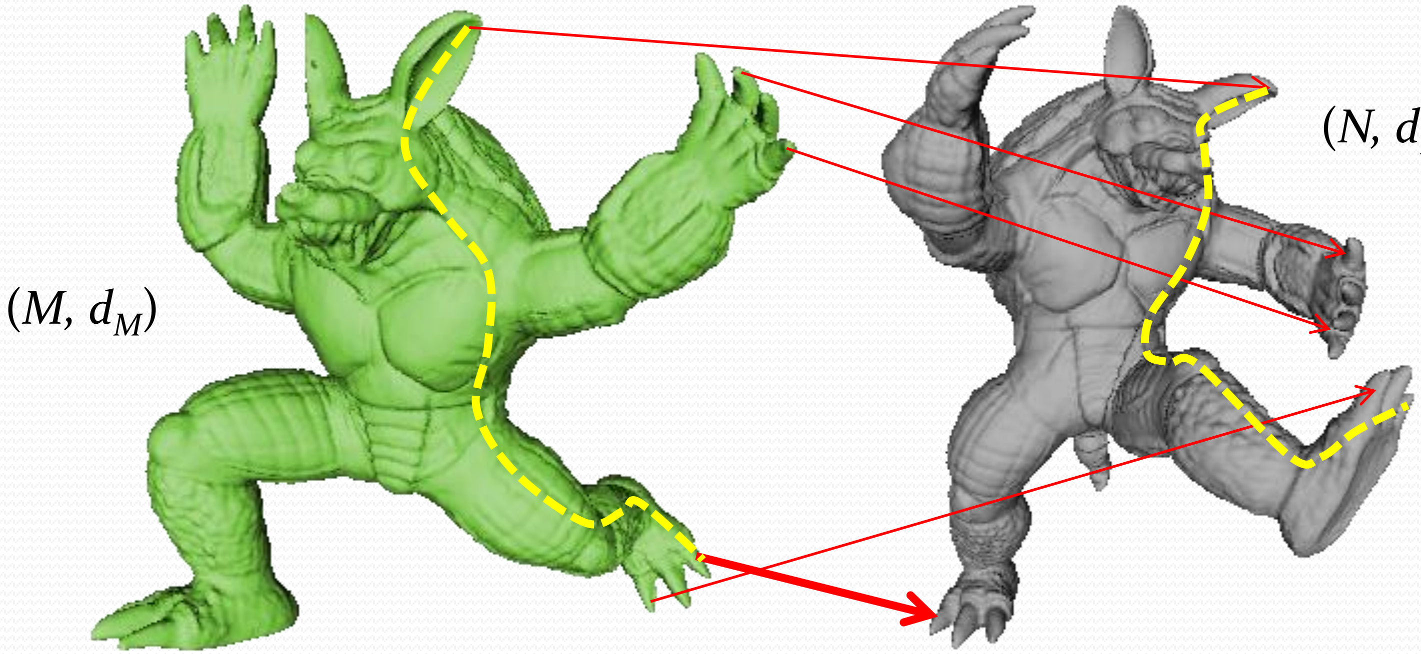

Metric preservation

$$T_{opt} = \arg\min_{T:M\to N} \sum_{(x_i,x_j)\in M\times M} \left\| d_M(x_i,x_j) - d_N(T(x_i),T(x_j)) \right\|$$

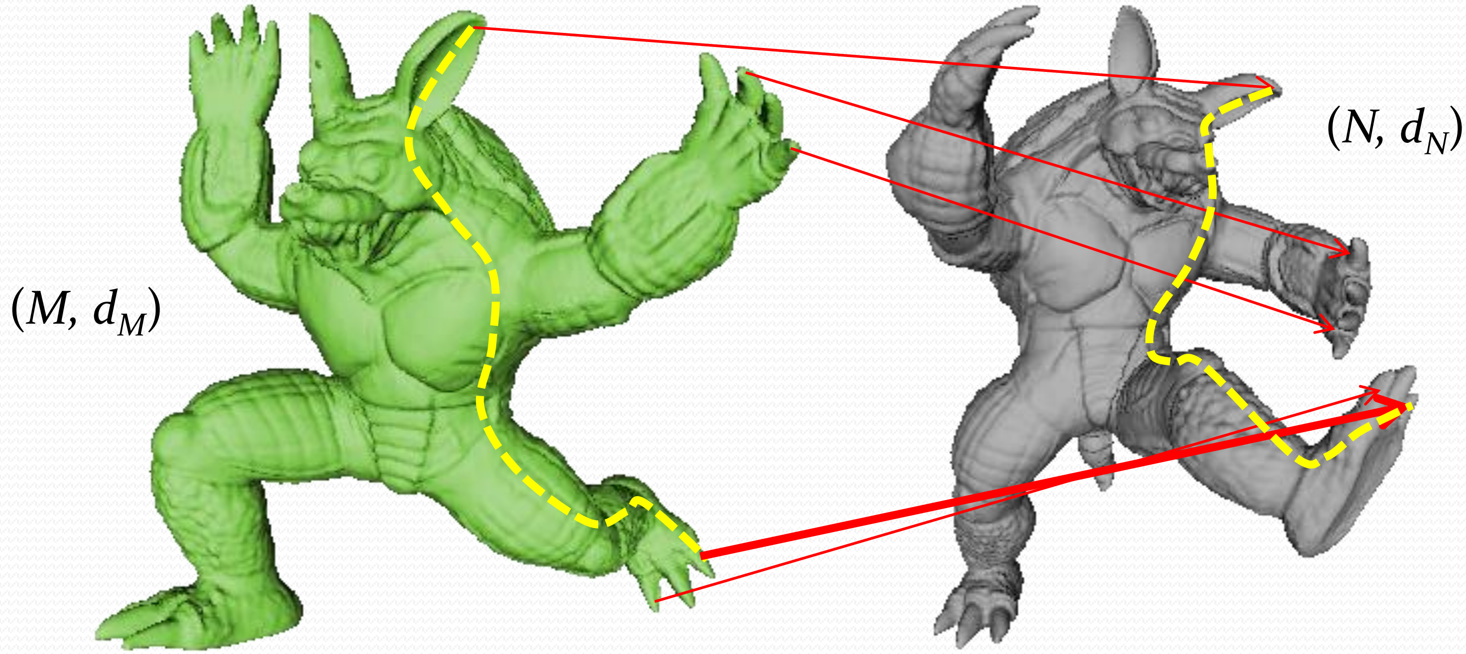

Metric preservation

$$T_{opt} = \arg\min_{T:M \to N} \sum_{(x_i, x_j) \in M \times M} \left\| d_M(x_i, x_j) - d_N(T(x_i), T(x_j)) \right\|$$

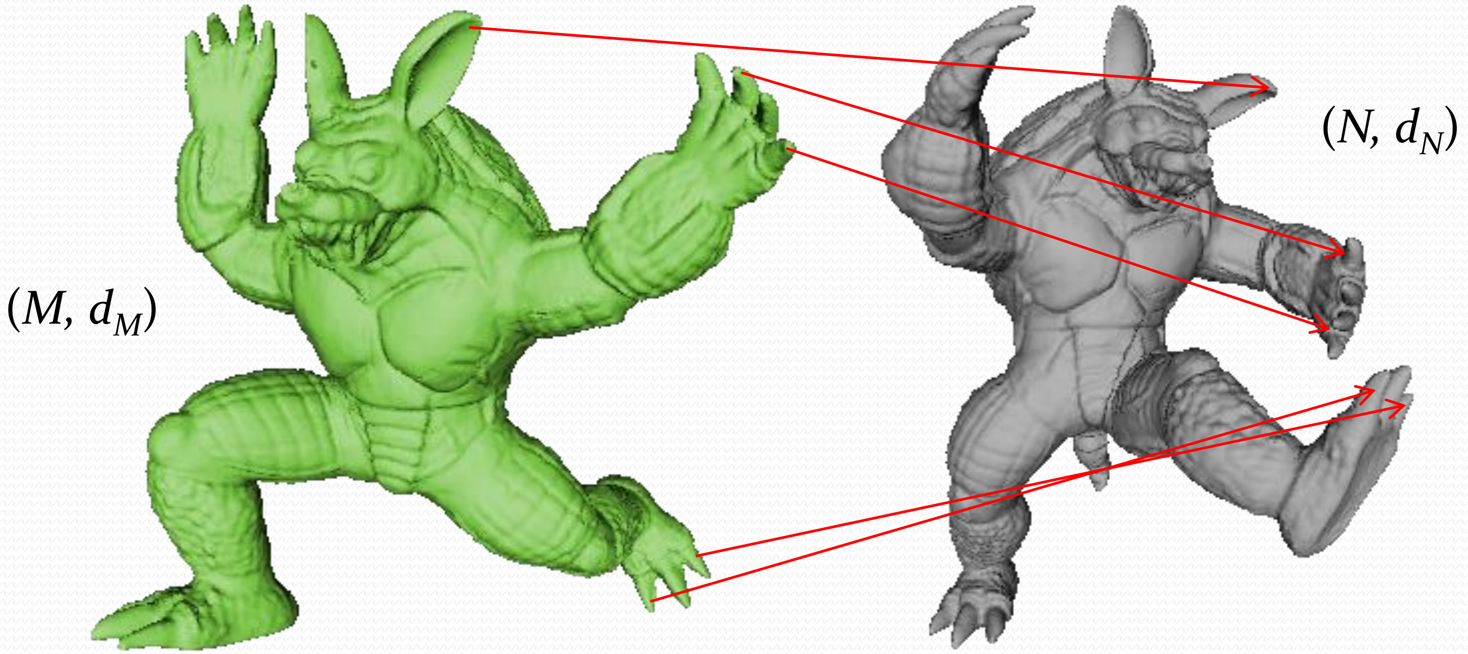

Examples of metrics

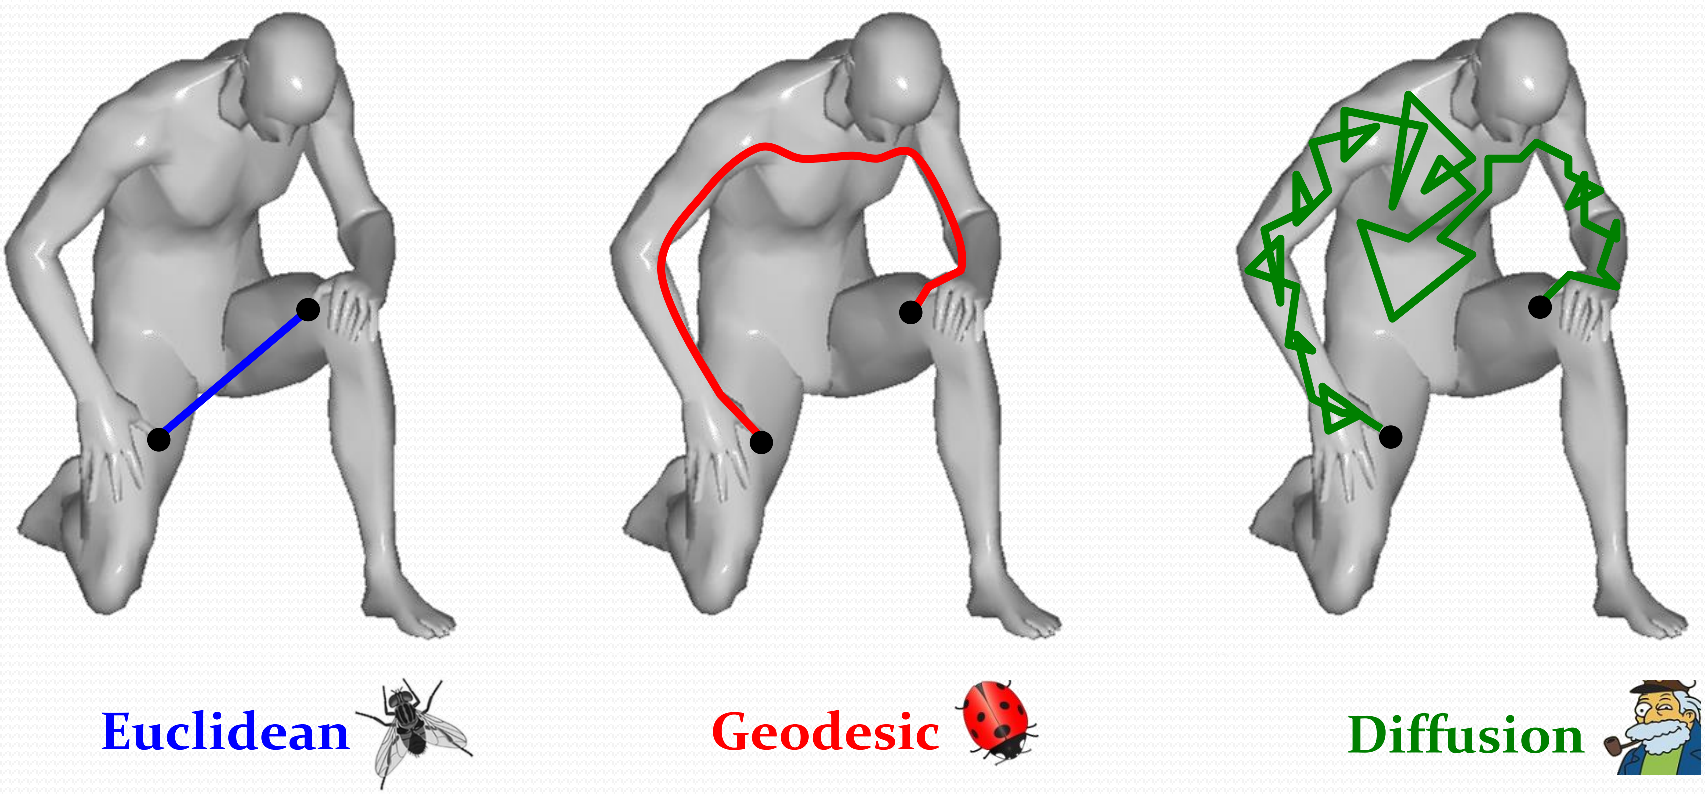

Invariance to what?

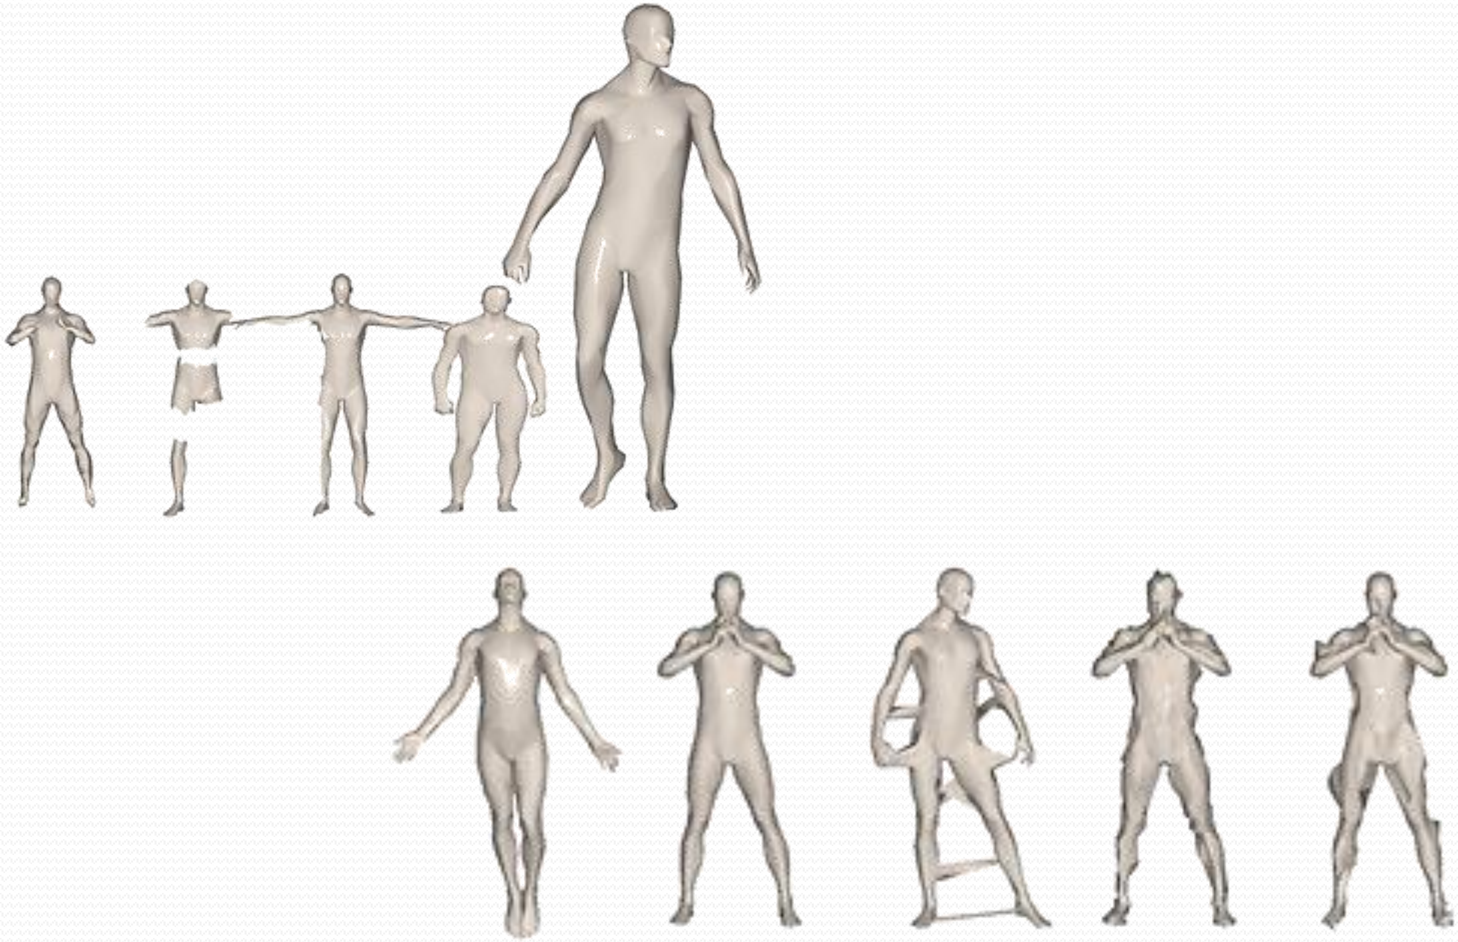

Invariance to what?

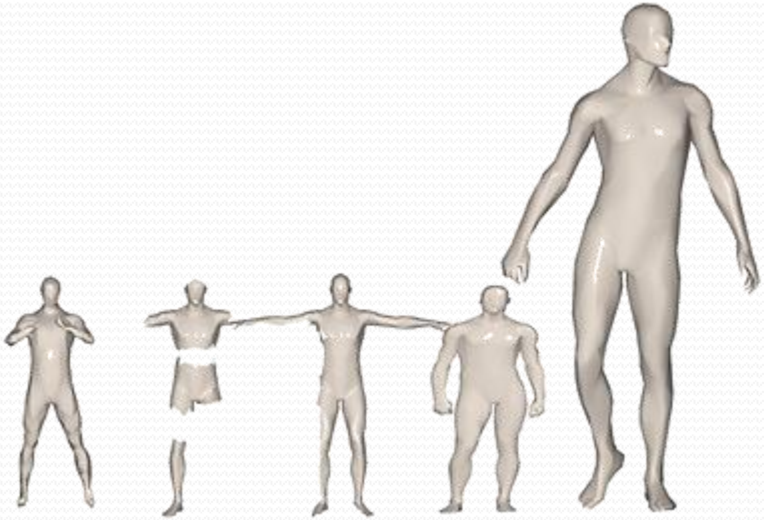

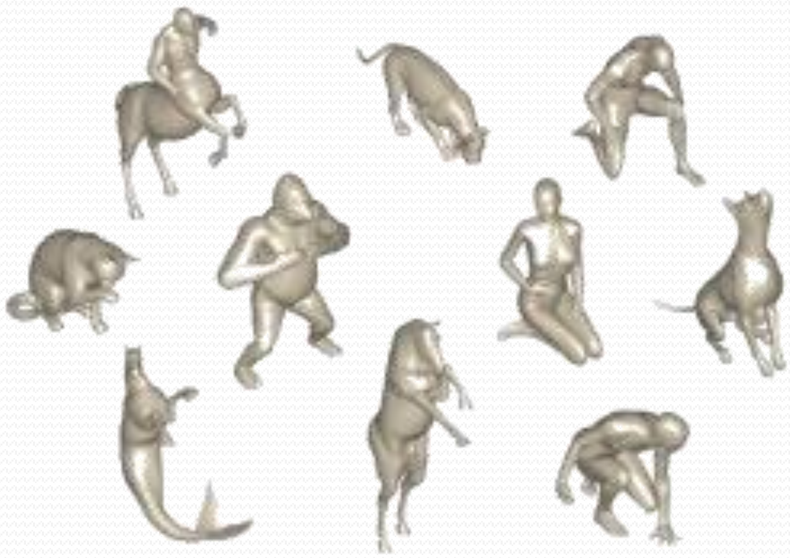

Invariance to what?

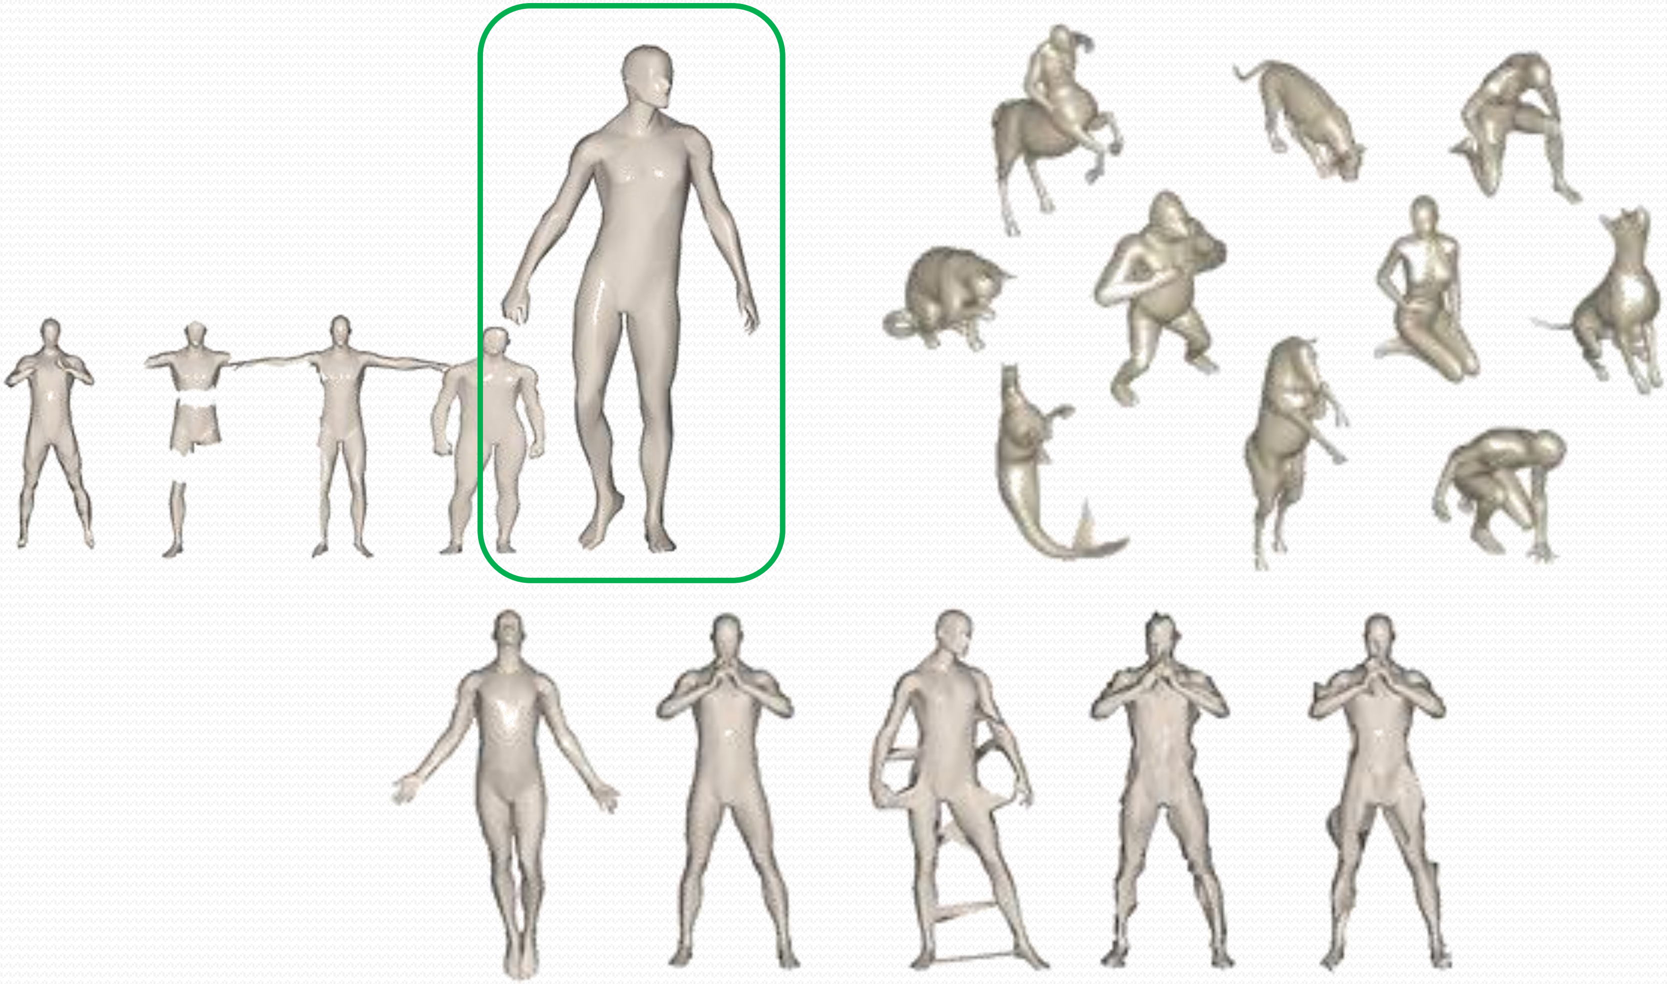

Invariance to what?

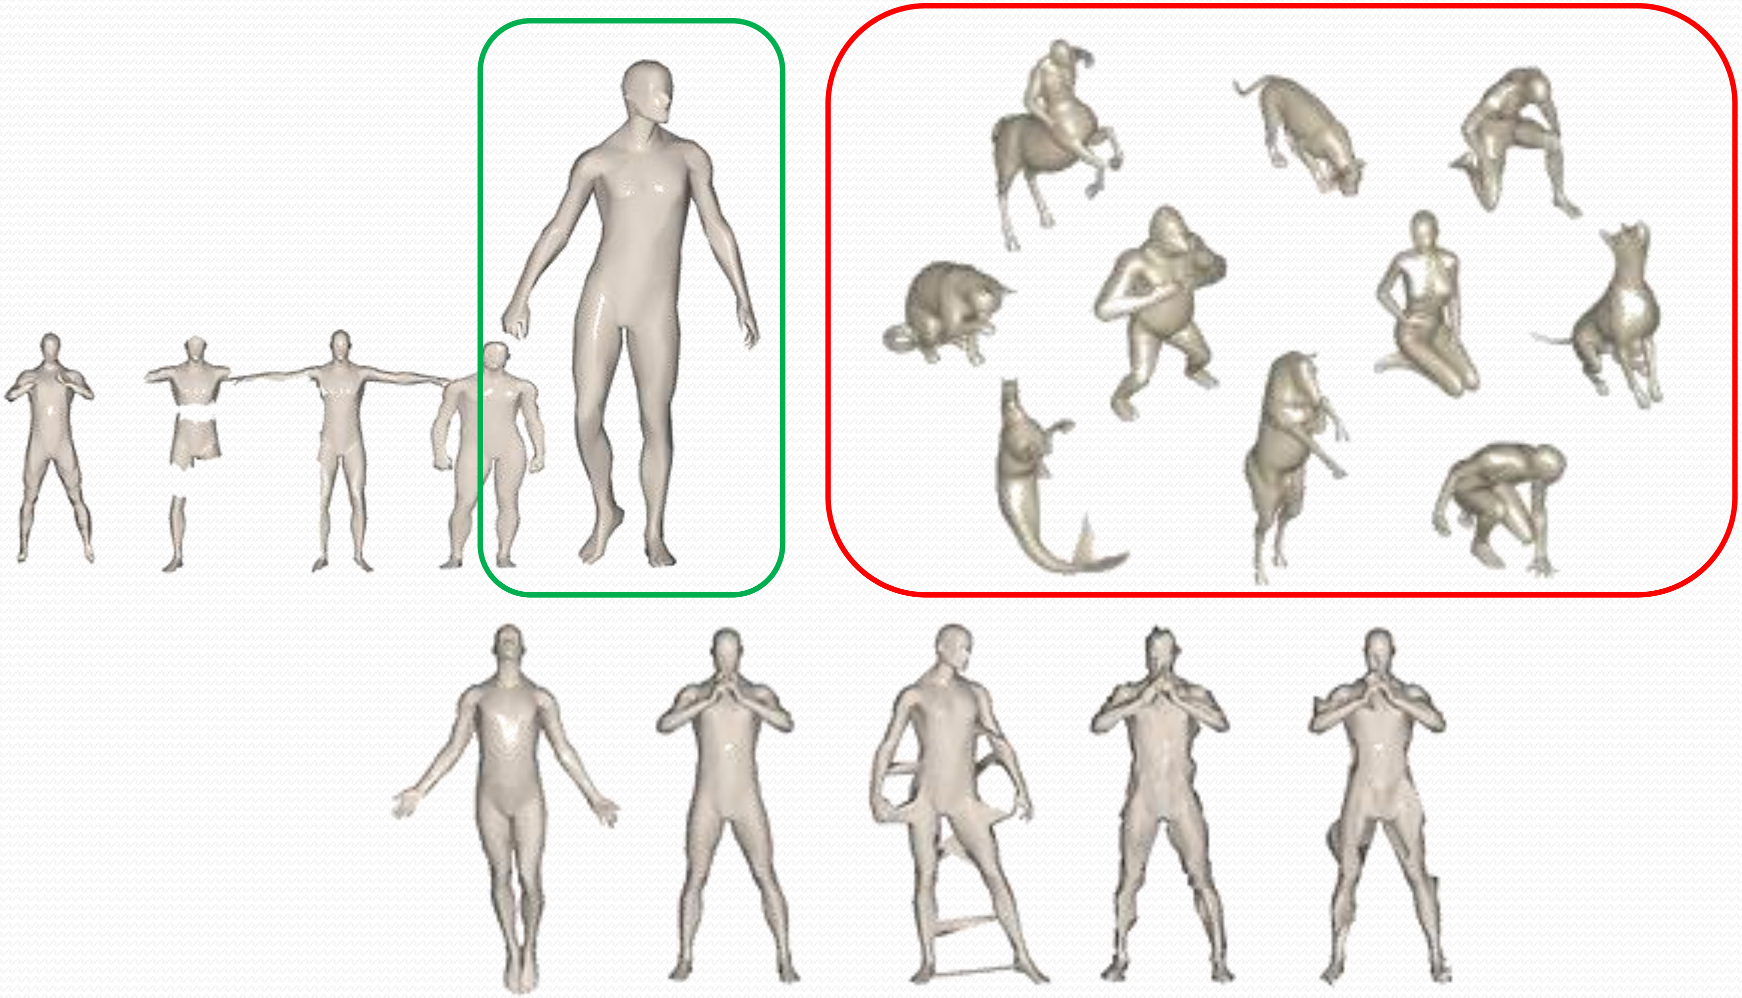

Invariance to what?

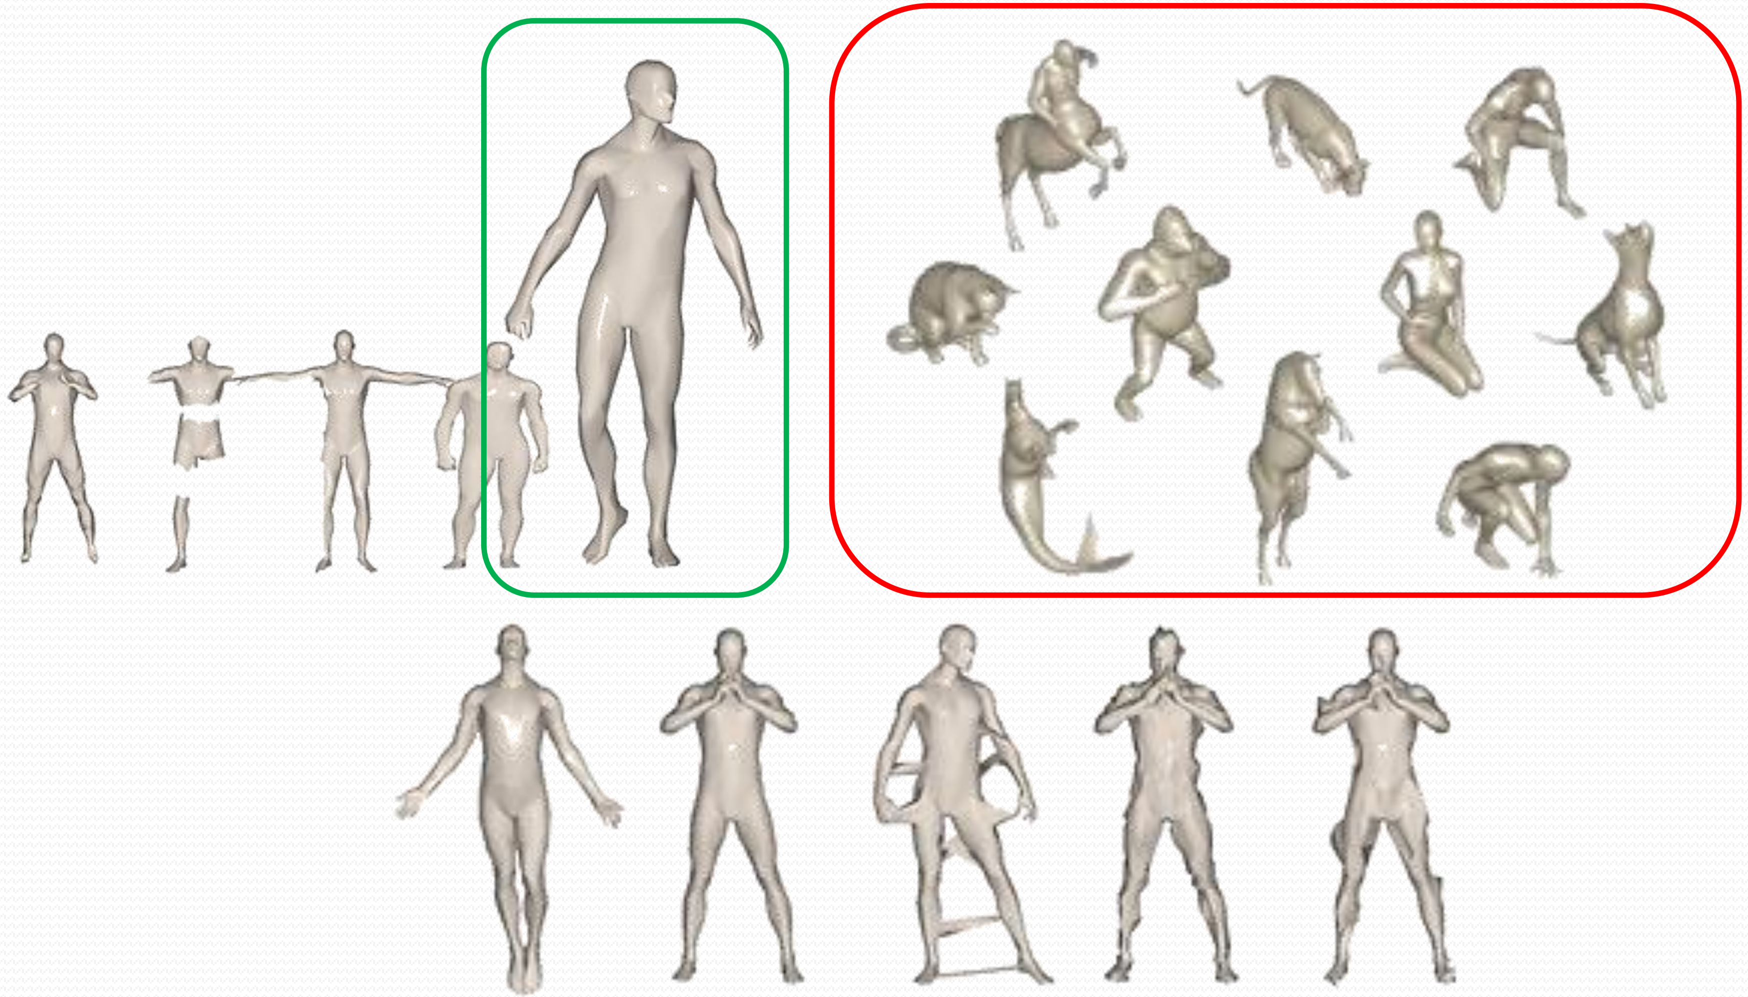

Inter-class matching, or…

Inter-class matching, or…

Matching a shark to a tornado

Inter-class matching, or…

Matching a shark to a tornado

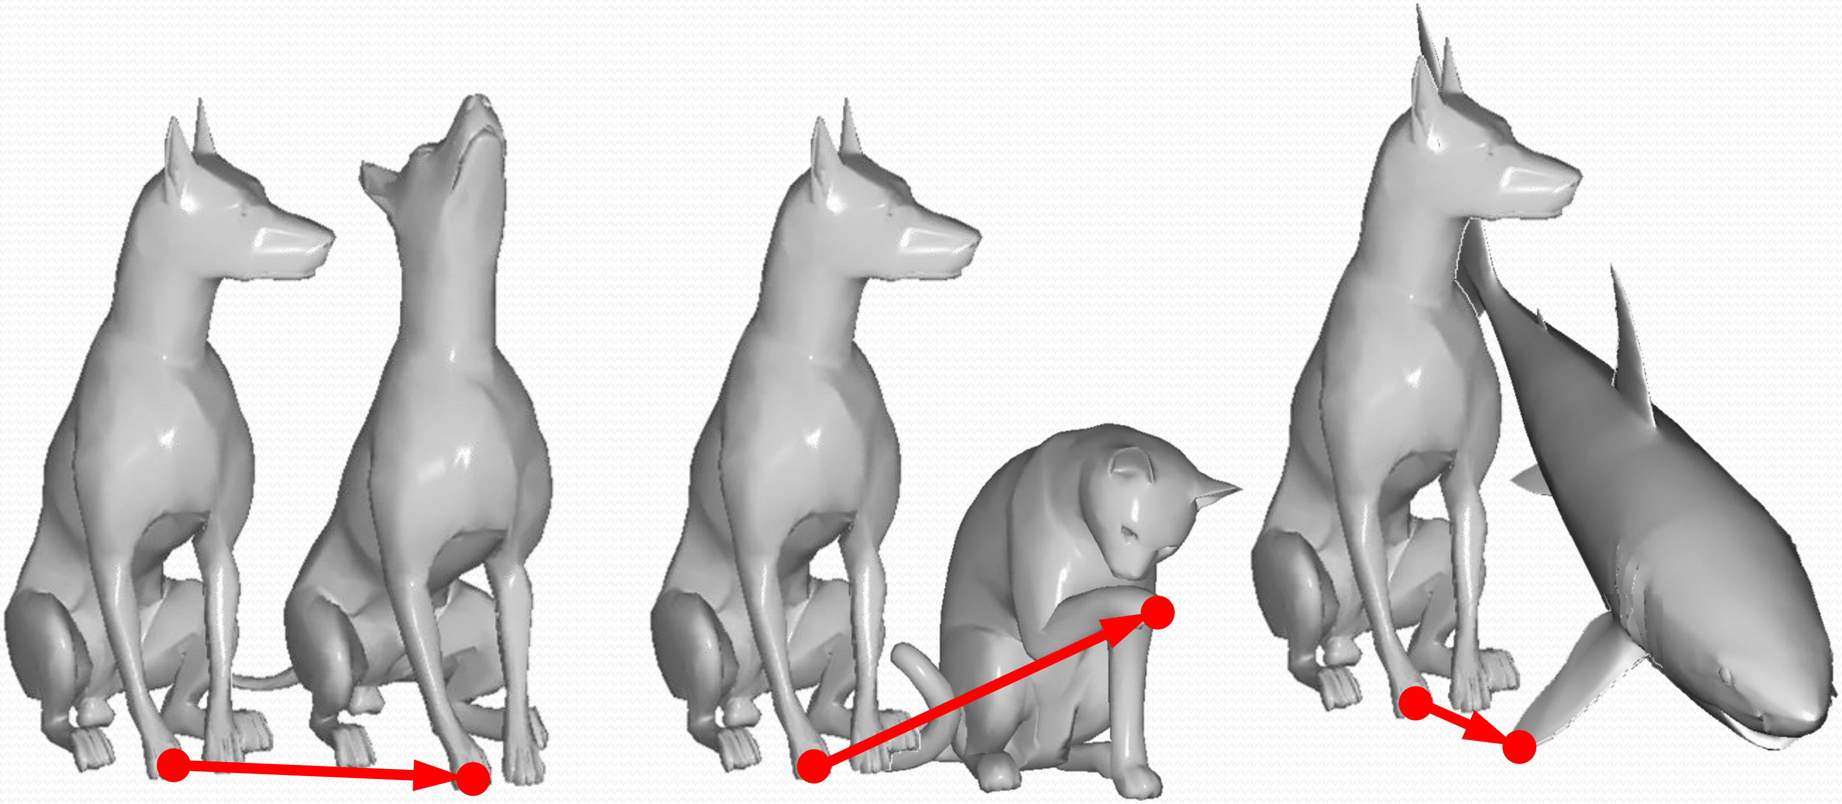

Geometric

“ ***accurate*** ”

Semantic

“ ***makes sense*** ”

Aesthetic

“ ***beautiful*** ”

Gromov-Hausdorff distance

$$T_{opt} = \arg\min_{T:M \to N} \sum_{(x_i, x_j) \in M \times M} \left\| d_M(x_i, x_j) - d_N(T(x_i), T(x_j)) \right\|$$

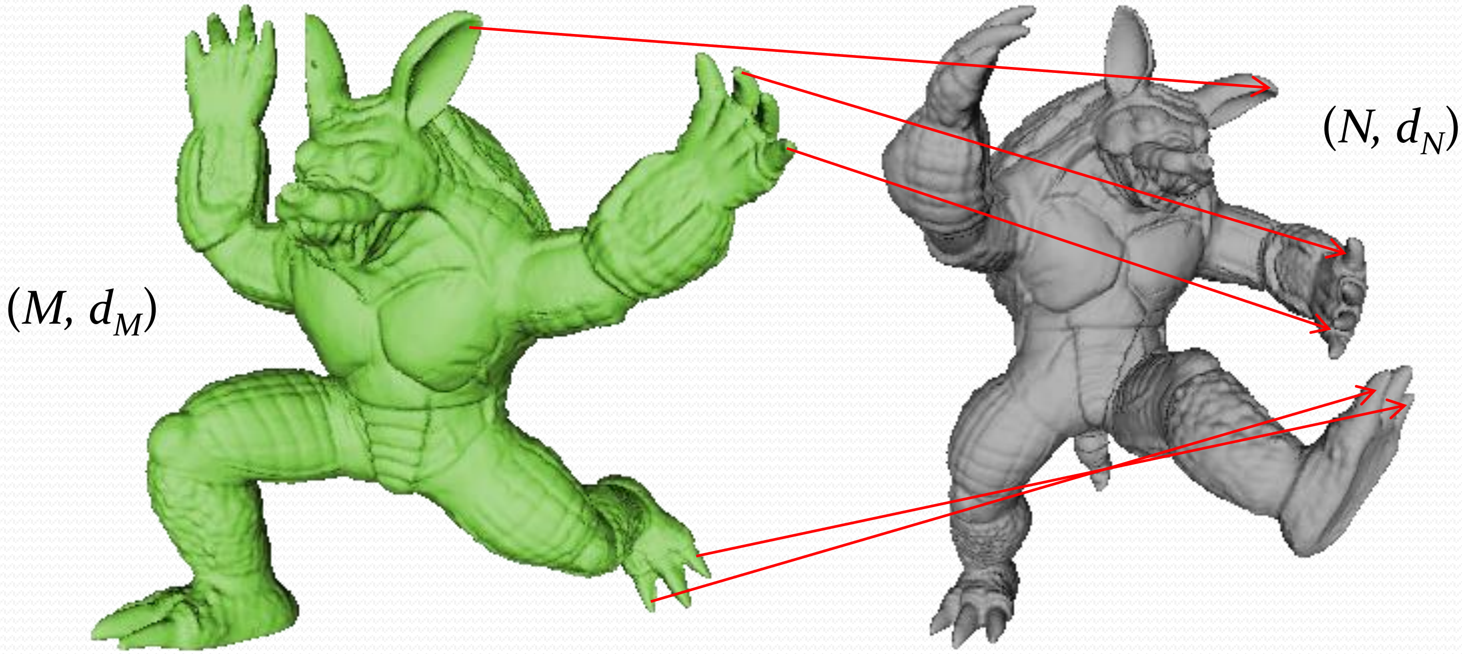

Gromov-Hausdorff distance

- Minimizing the worst-case distortion of the metric caused by the correspondence T is given by:

$$D_{\mathcal{GH}}(M,N) =$$

$$= \min_{T:M\to N} \max_{(x_i,x_j)\in M\times M} \left\| d_M(x_i,x_j) - d_N(T(x_i),T(x_j)) \right\|$$

This is a *true* distance among shapes ☺

Gromov-Hausdorff distance

$D_{\mathcal{GH}}\Bigg($  , 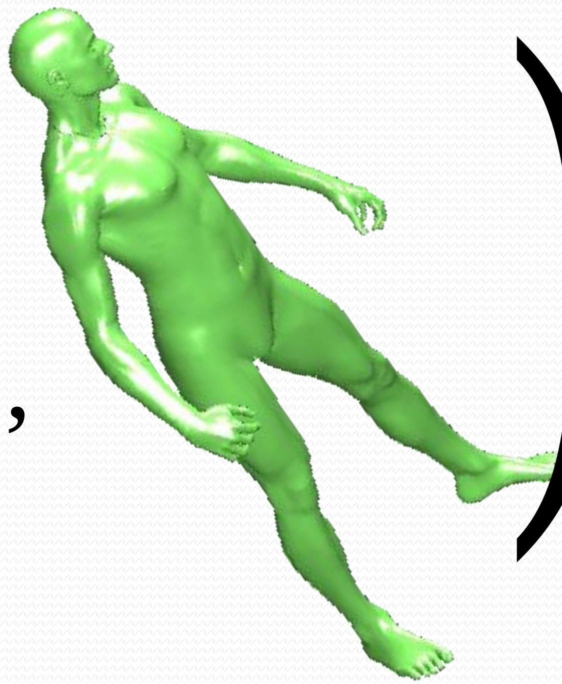 $\Bigg) = 0$

Gromov-Hausdorff distance

$D_{\mathcal{GH}}\Big($  ,  $\Big) = 0.82$

Gromov-Hausdorff distance

$D_{\mathcal{GH}}\Bigg($  ,  $\Bigg) = 10^3$

Space of shapes

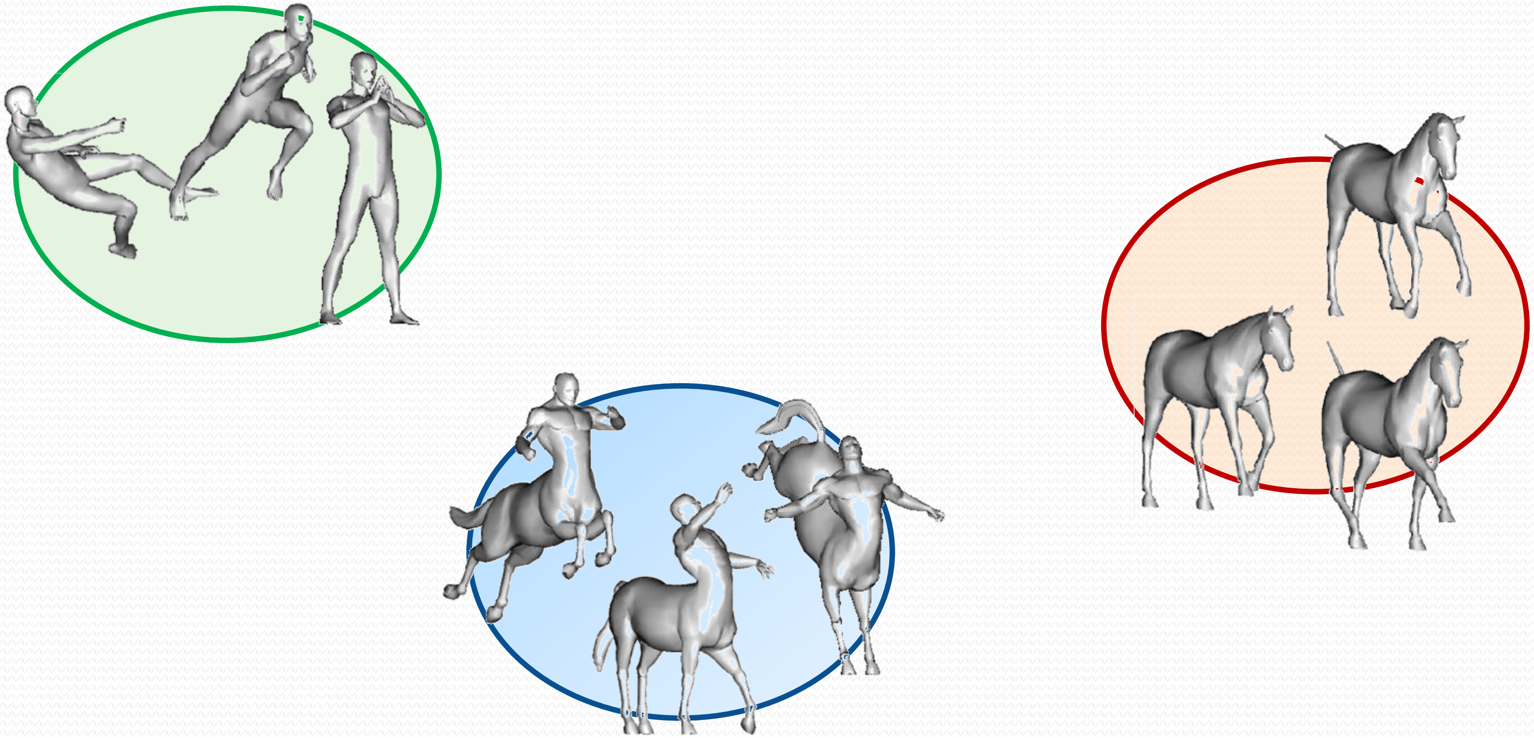

Is there something like a “space of shapes”?

Space of shapes

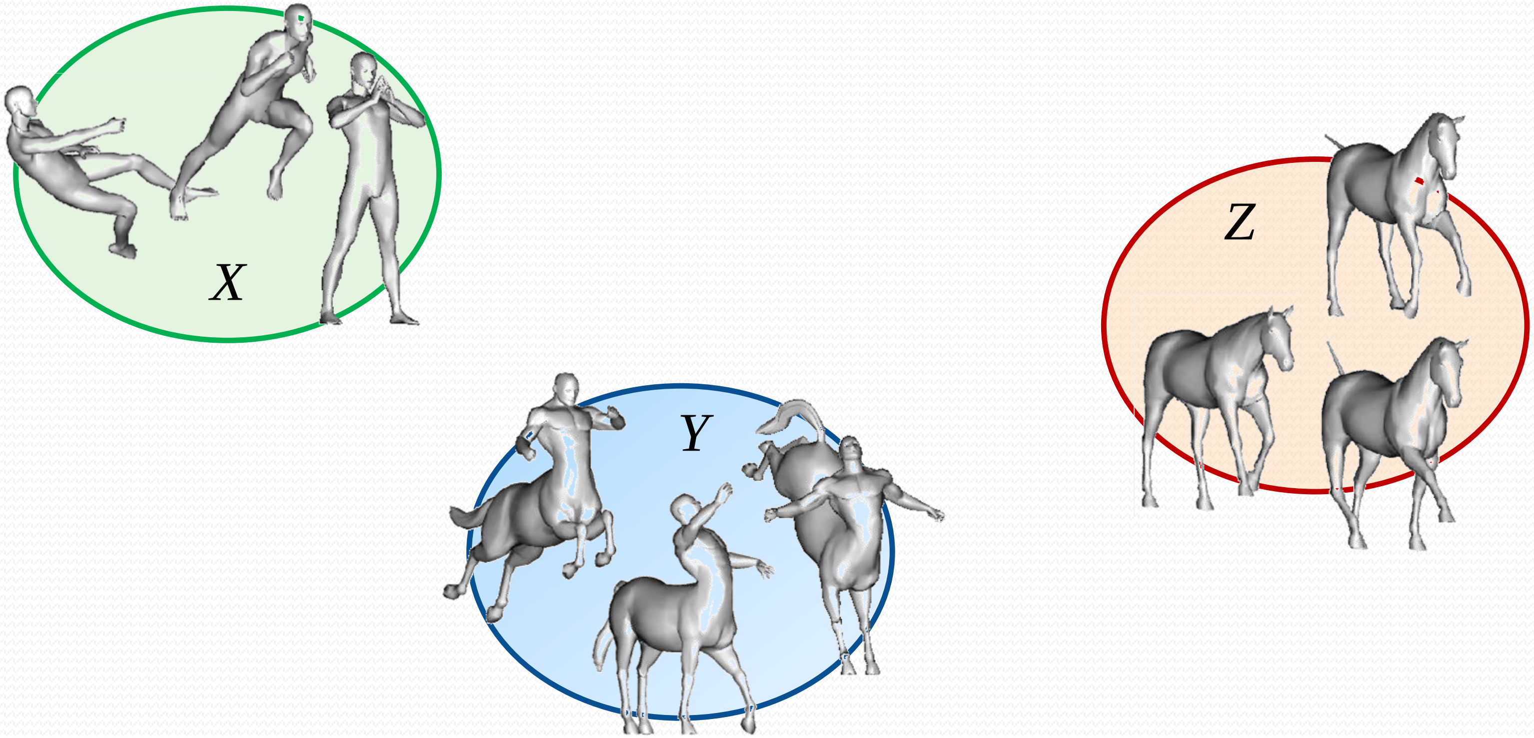

Is there something like a "space of shapes"? **Yes!**

Space of shapes

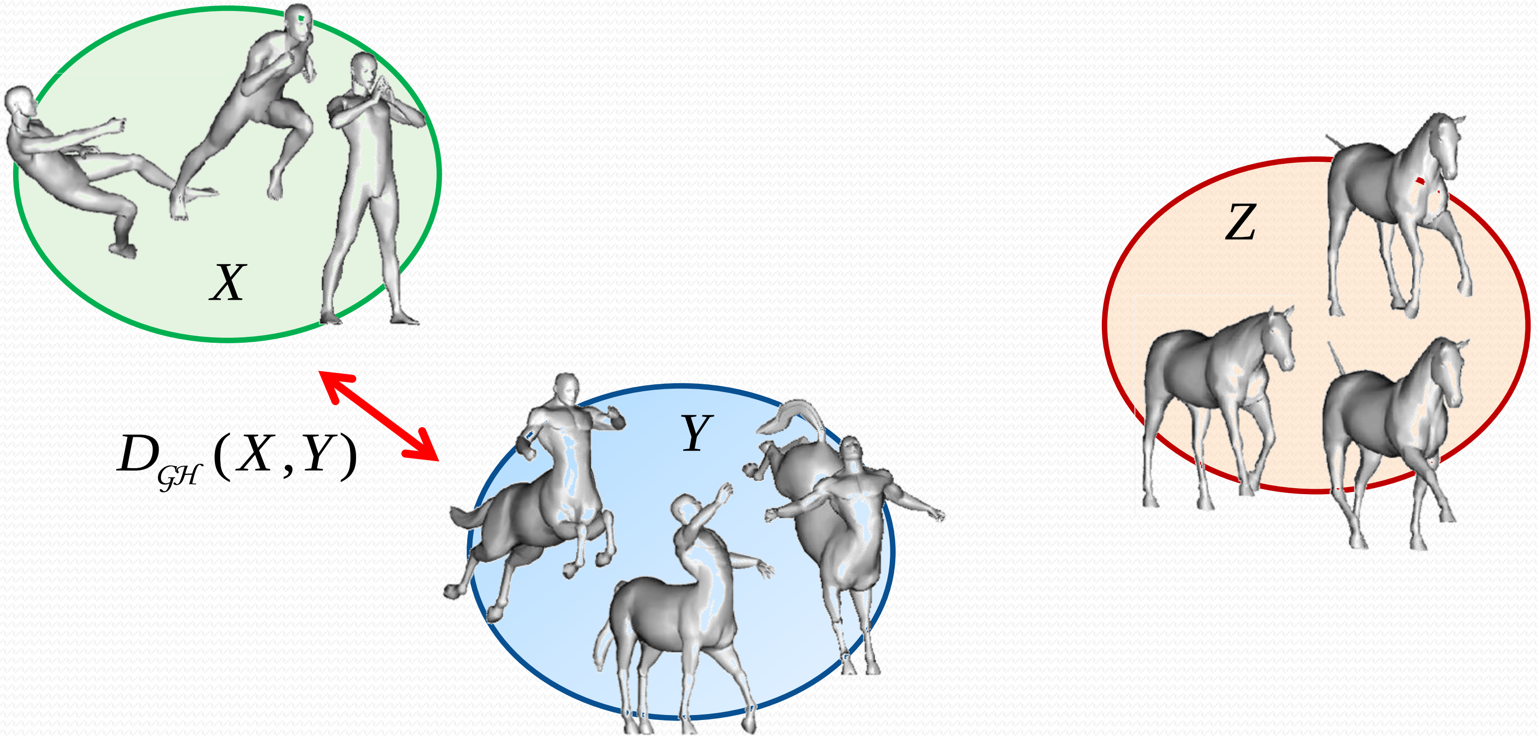

Space of shapes

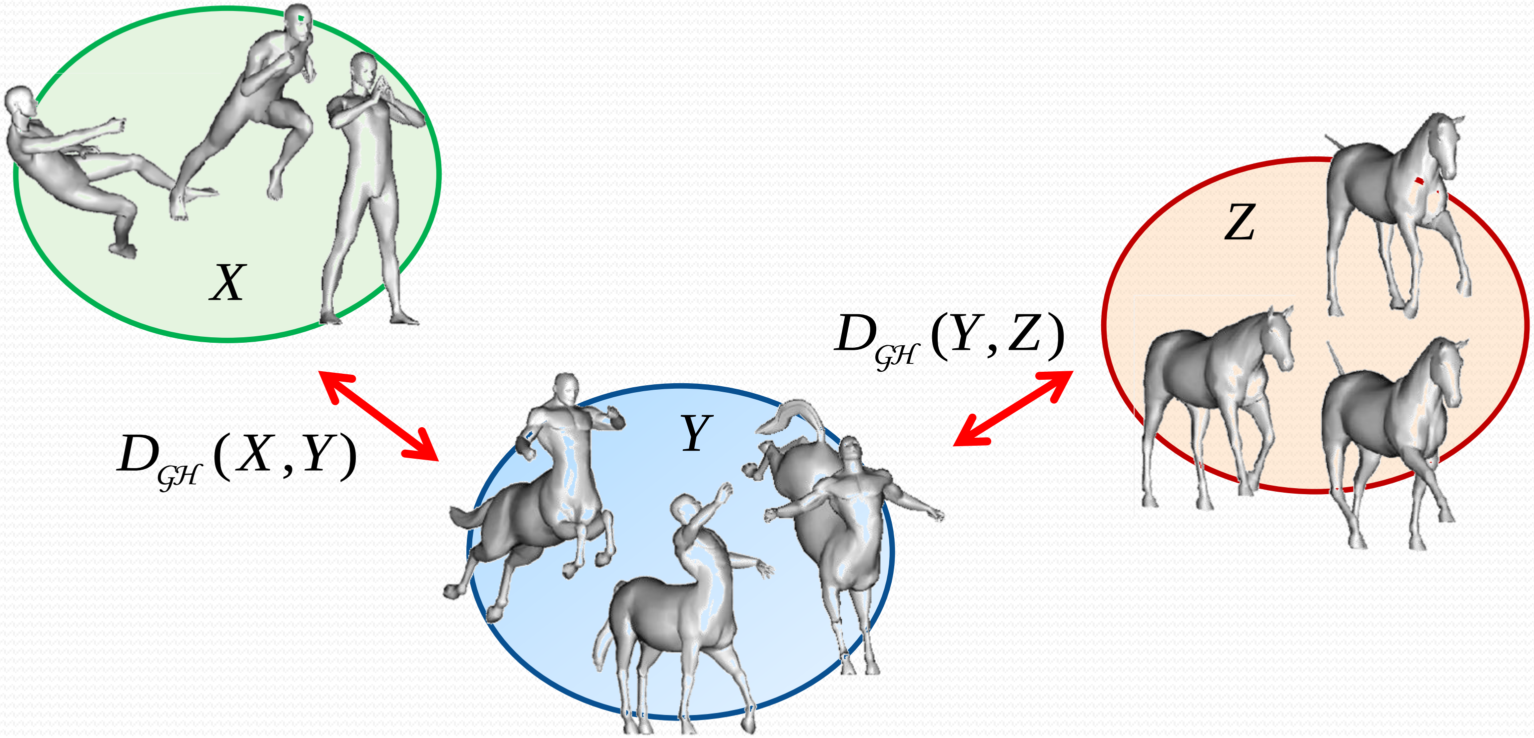

Space of shapes

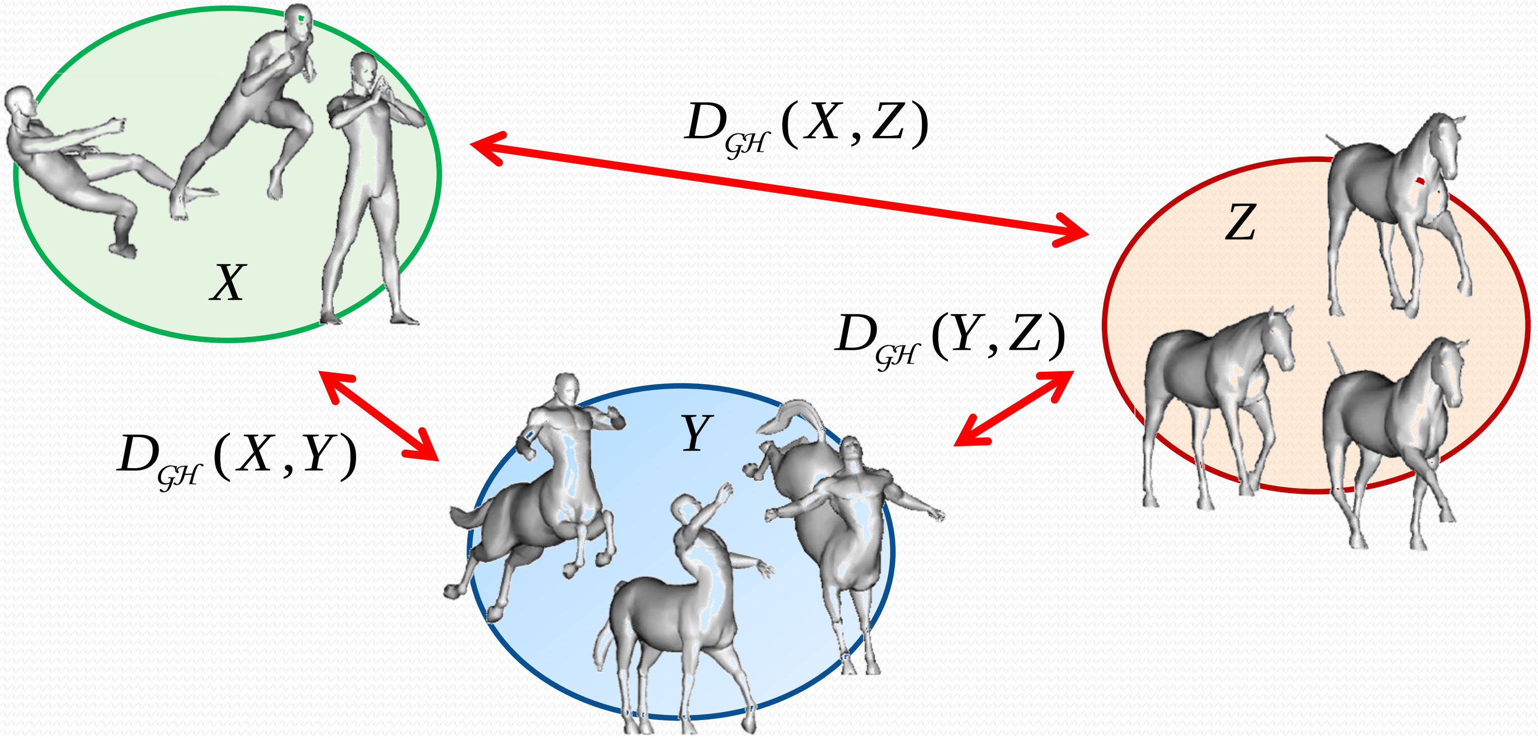

Space of shapes

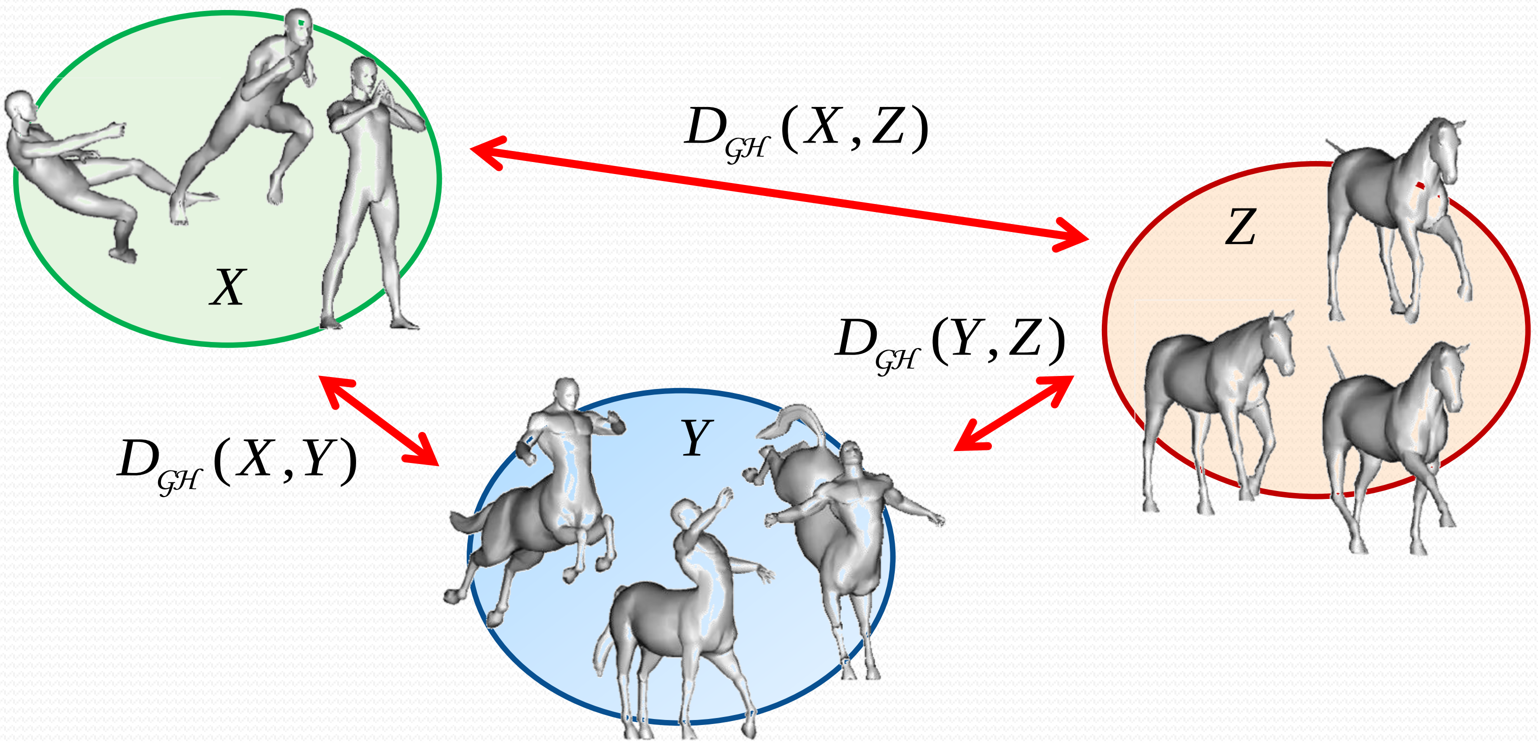

Triangle inequality: $D_{\mathcal{GH}}(X,Y) + D_{\mathcal{GH}}(Y,Z) \geq D_{\mathcal{GH}}(X,Z)$

Beyond two shapes

- Let us consider an entire **collection** of shapes

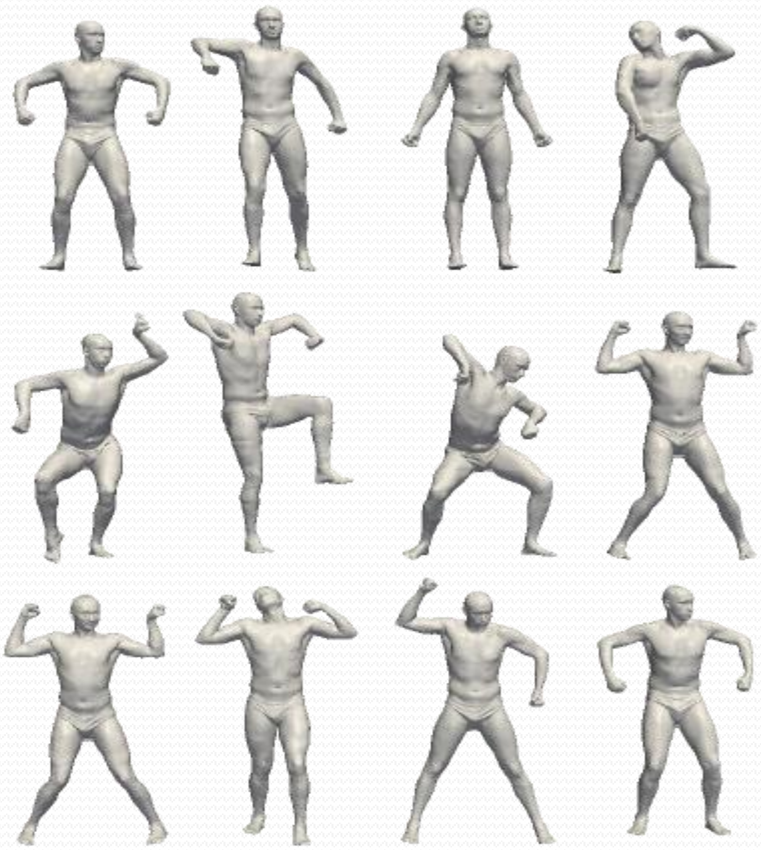

Beyond two shapes

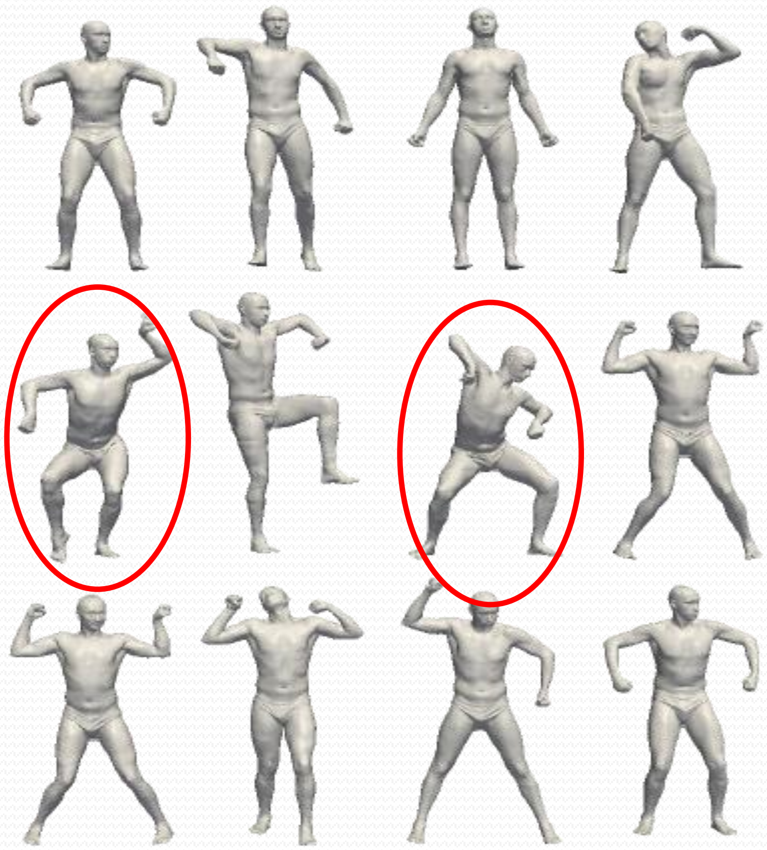

Difficult to match!

Beyond two shapes

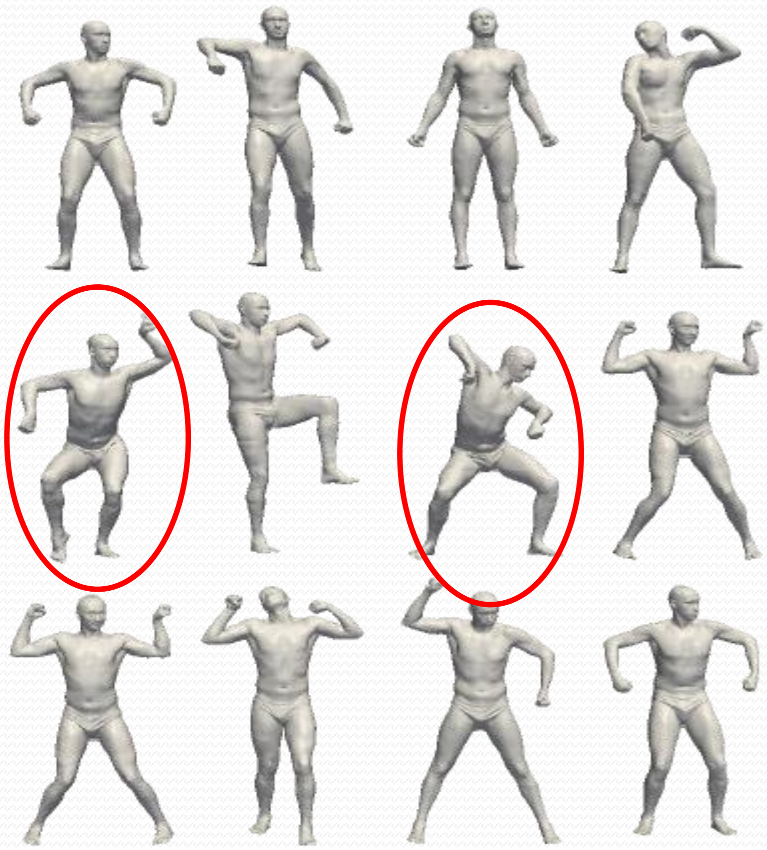

Difficult to match!

Can we use additional information to produce better correspondences?

Beyond two shapes

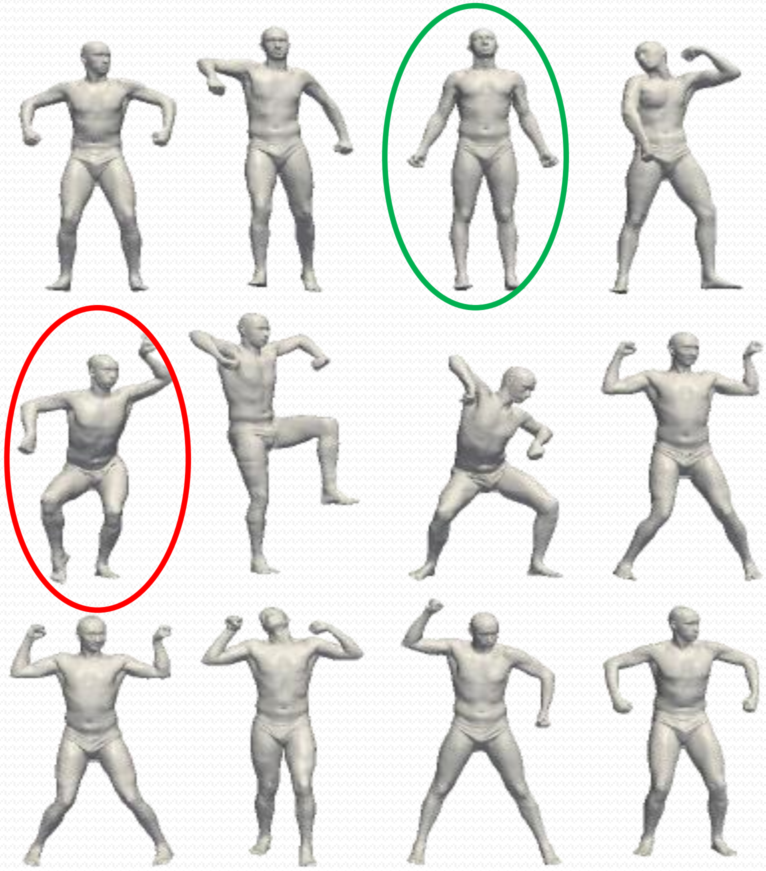

Easier to match!

Beyond two shapes

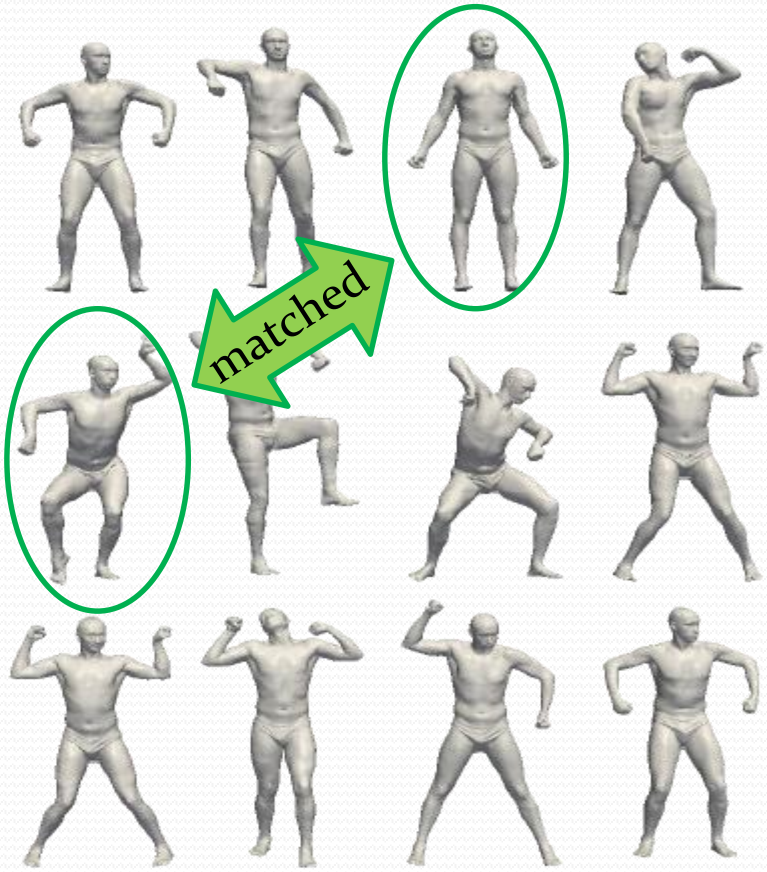

Beyond two shapes

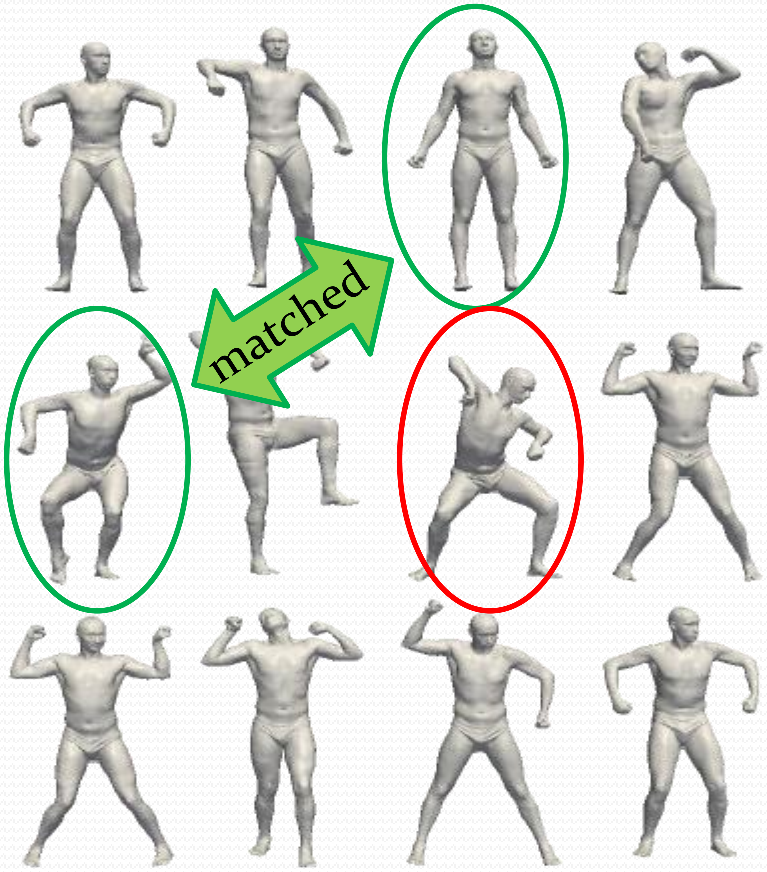

Beyond two shapes

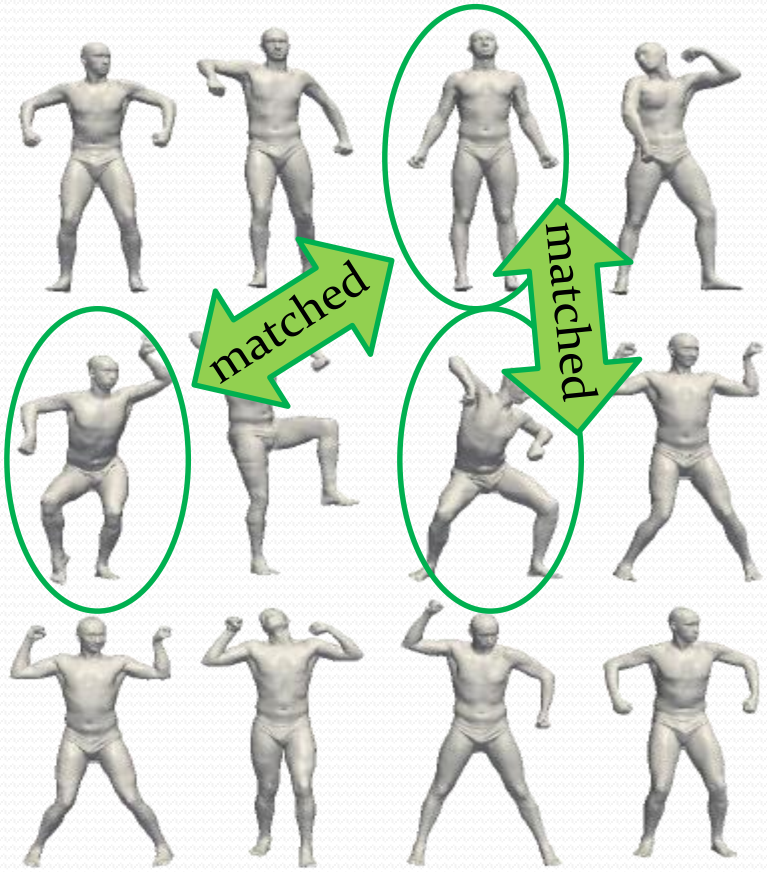

Beyond two shapes

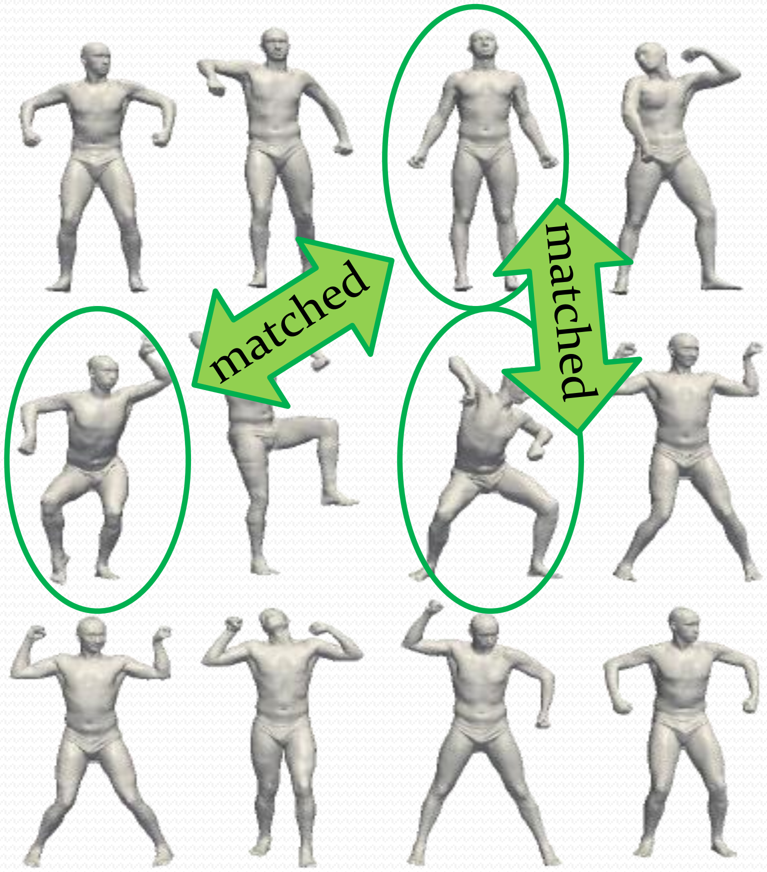

A correspondence can now be induced by transitivity or “triangle consistency”

Beyond two shapes

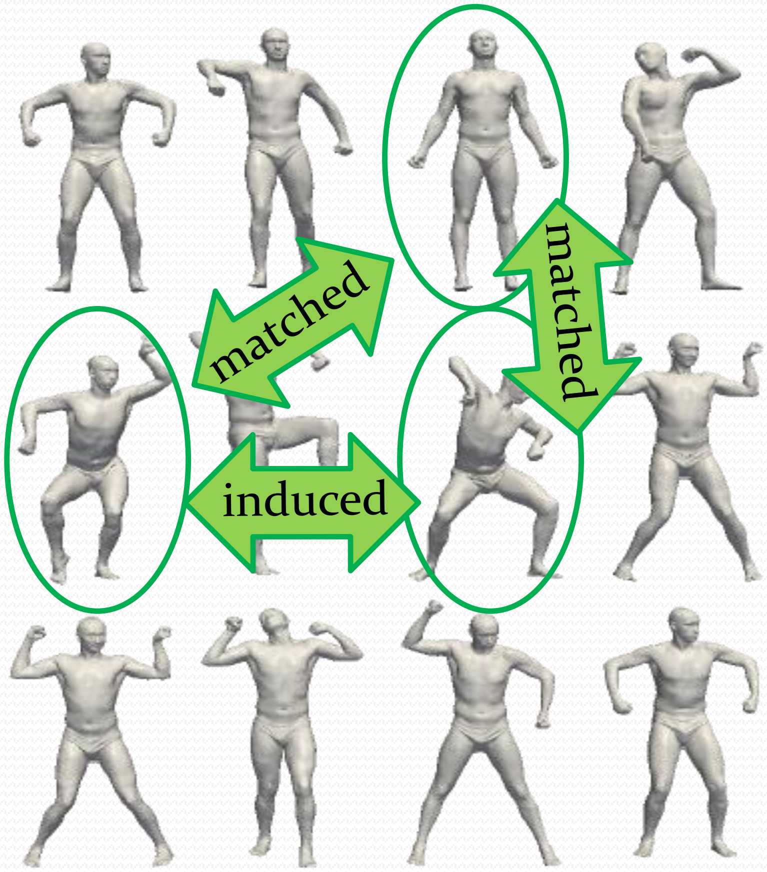

A correspondence can now be induced by transitivity or “triangle consistency”

Suggested readings

- *Numerical geometry of non-rigid shapes*. Chapter 1 – Introduction.