Tutorial 2
Biological shape descriptors

Patrice Koehl and Joel Hass

University of California, Davis, USA
http://www.cs.ucdavis.edu/~koehl/IMS2017/

Deciphering Biological Shapes

-How do we understand shapes?

The Mumford experiments

-Shape Descriptors

Deciphering Biological Shapes

-How do we understand shapes?

The Mumford experiments

-Shape Descriptors

Start with an easy case:

Before moving to the problem of comparing surfaces in R^3, we ask a simpler question:

Problem: How similar are two *regions in the plane*?

This is already an important problem.

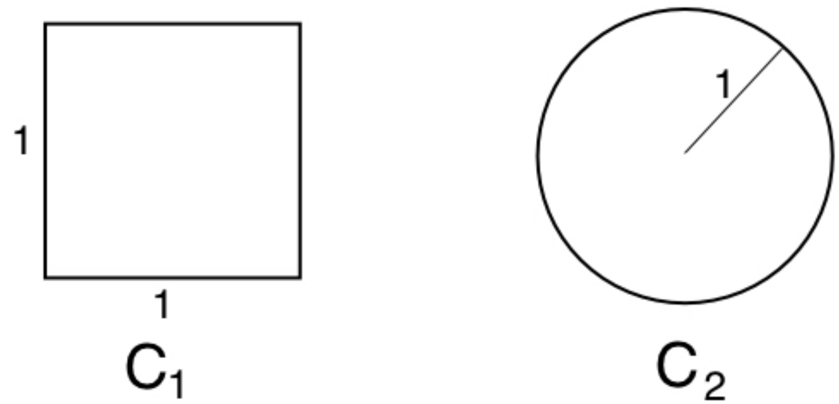

Question: How close is a square to a circle?

Distance between shapes

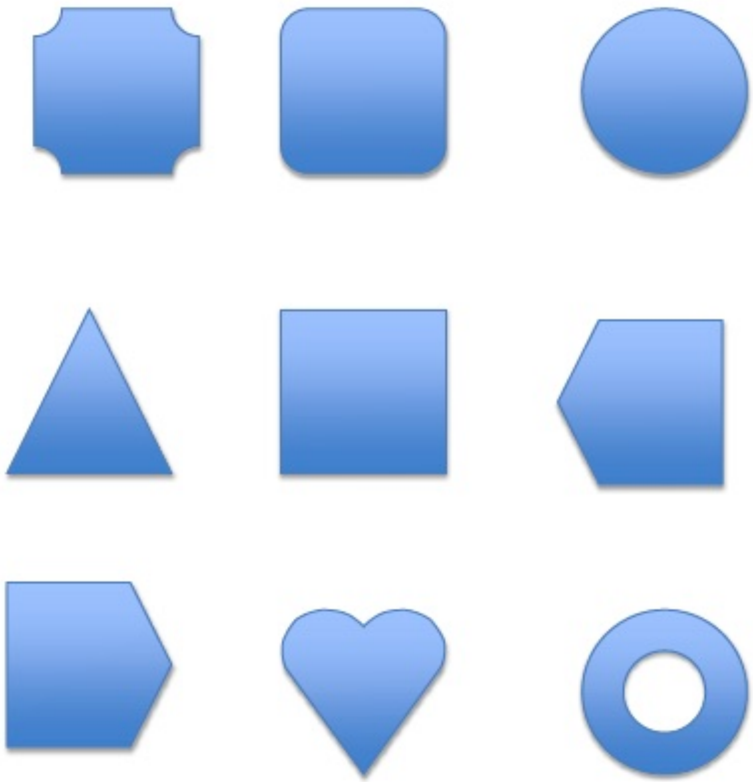

Which of these nine shapes is closest to

?

Which is second closest?

Application - Facial Recognition

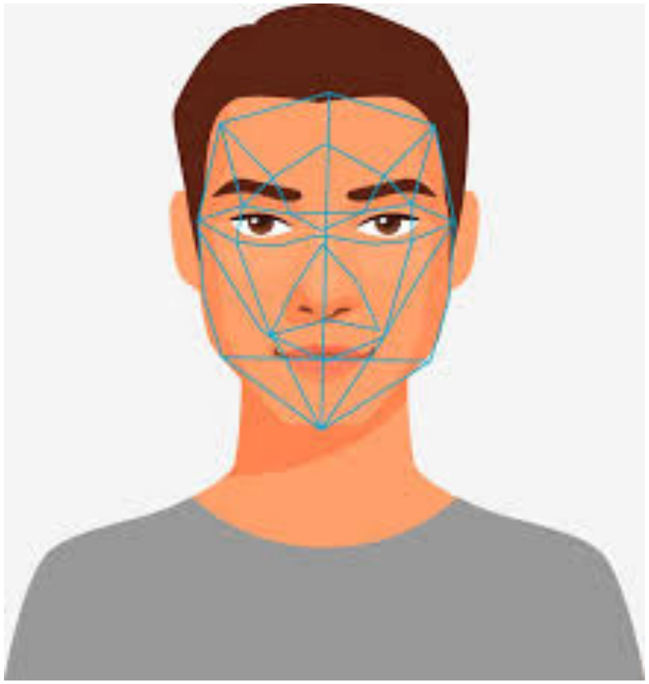

Start with a 2D photograph.
Create some planar regions from a face.
Compare their shapes.

Application - Computer Vision

"Purring Test" Cat or Dog?

Flip a coin - correct 50% of the time
Software fifteen years ago - not much better
Today - 99%

Application - Computer Vision

Dog or Muffin?   Still a challenge

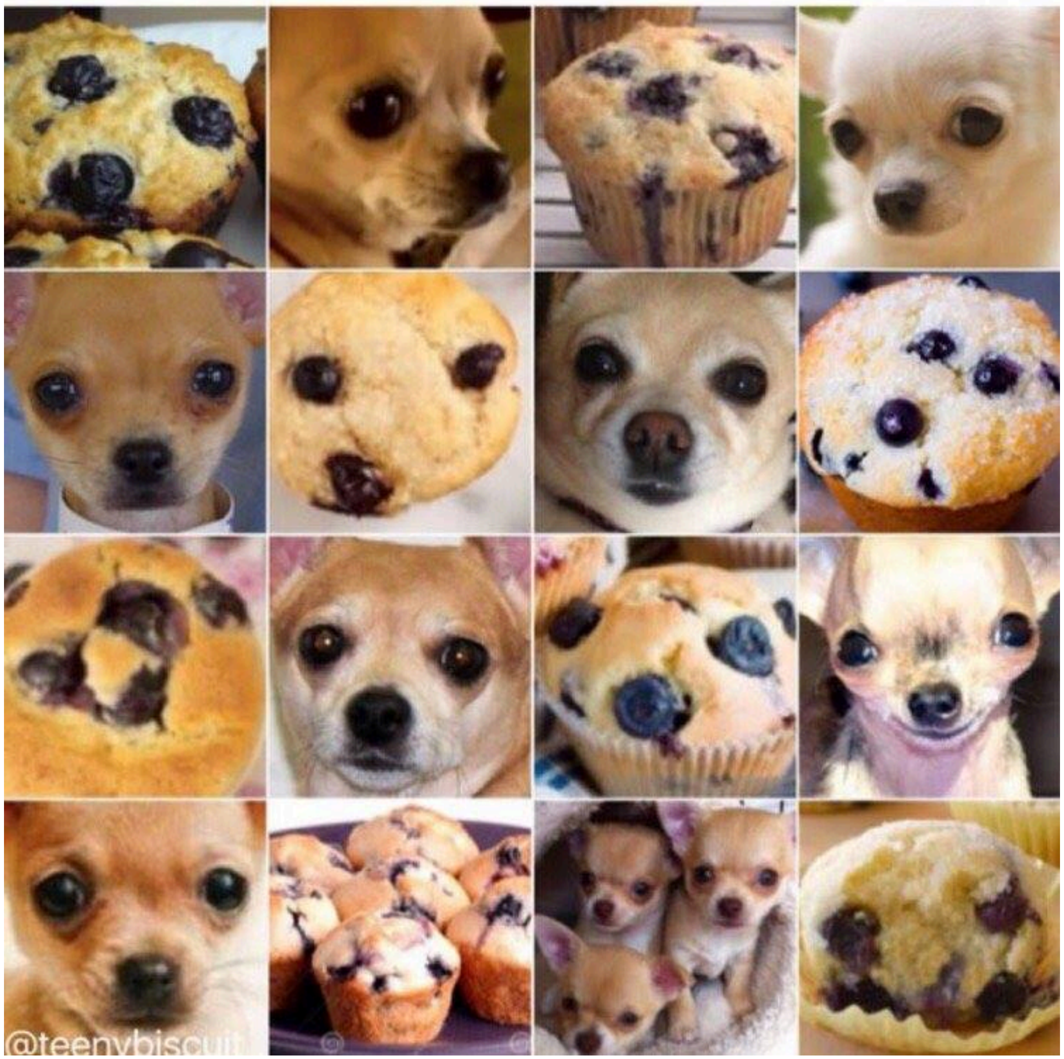

Application - Computer Vision

Puppy or Bagel?

Application - Character Recognition

What letter is this?

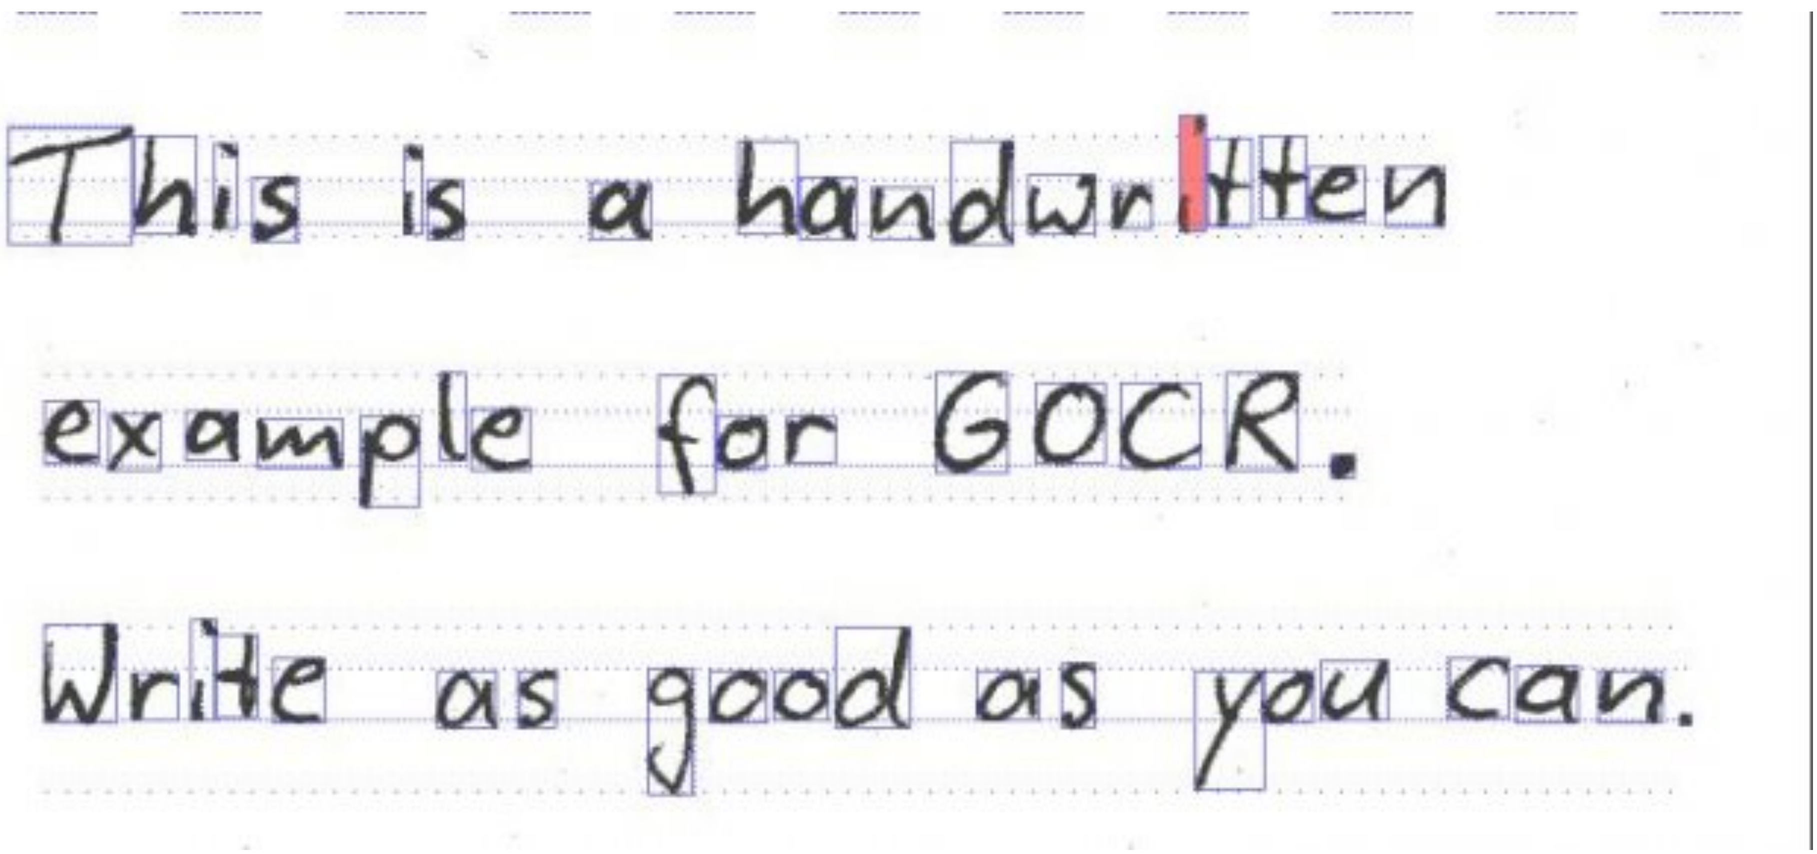

Test Case

How close are these two shapes?

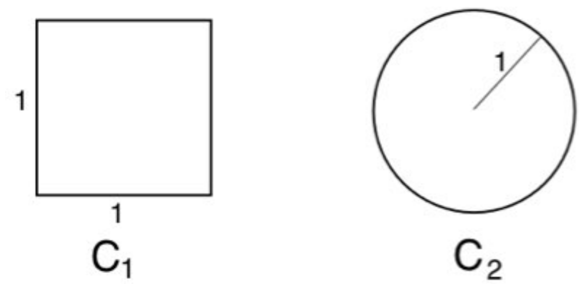

Can either compare curves or enclosed regions:

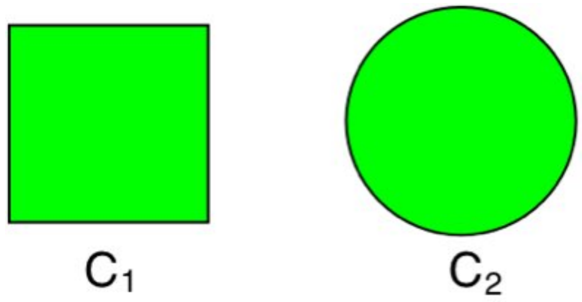

Our Goal: Find a mathematical framework to measure the similarity of two shapes.

Goal for 2D shapes:
A metric on curves in the plane

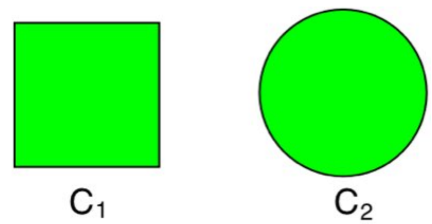

1. $d(C_1, C_2) = 0 \Longleftrightarrow C_1$ *is isometric to* C_2 (isometry)

2. $d(C_1, C_2) = d(C_2, C_2)$ (symmetry)

3. $d(C_1, C_3) \leq d(C_1, C_2) + d(C_2, C_3)$ (triangle inequality)

Why these three metric properties?

1. $d(C_1, C_2) = 0 \Longleftrightarrow C_1$ *is isometric to* C_2 (isometry)

2. $d(C_1, C_2) = d(C_2, C_2)$ (symmetry)

3. $d(C_1, C_3) \leq d(C_1, C_3) + d(C_1, C_3)$ (triangle inequality)

Each property plays an important role in applications.

Isometry: $d(C_1, C_2) = 0 \Longleftrightarrow C_1$ is isometric to C_2

Allows for identifying different views of the same object.

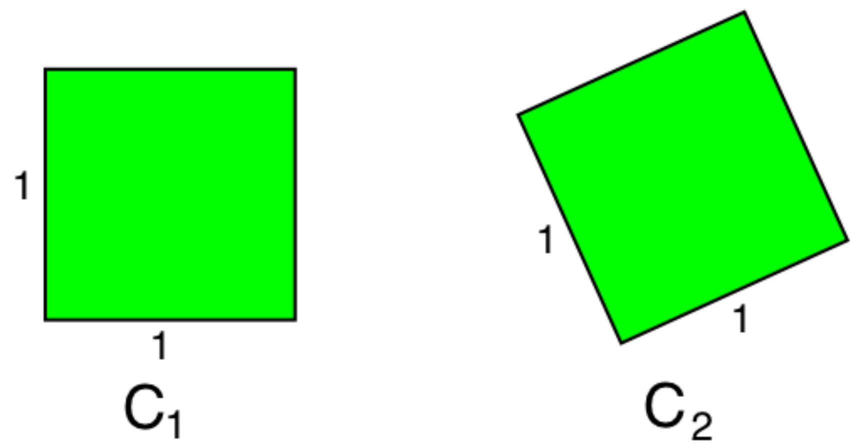

We want to consider these to be the same object.
Our distance measure should not change if one shape is moved by a Euclidean Isometry.

Symmetry: $d(C_1, C_2) = d(C_2, C_2)$

The distance between two objects does not depend on the order in which we find them.

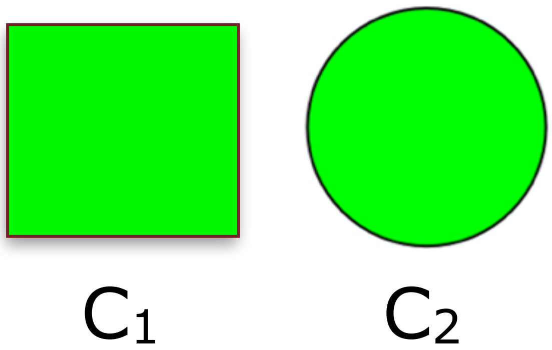

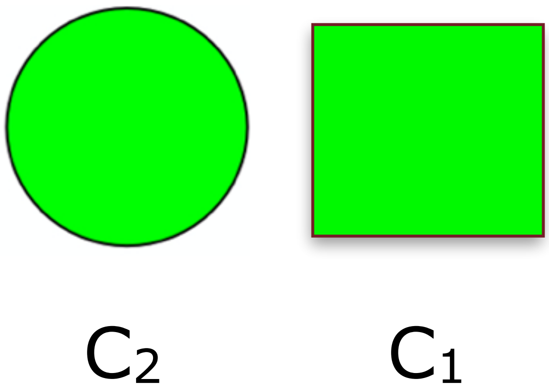

If I own the square, and you own the circle, we can agree on the distance between them.

Triangle inequality: $d(C_1, C_3) \leq d(C_1, C_2) + d(C_2, C_3)$

Measurements should be stable under small errors.

$$d(C_1, C_3) - d(C_2, C_3) \leq d(C_1, C_2)$$

If C_1 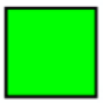 and C_2  are close, so $d(C_1, C_2)$ is small, then

the distance of C_1 and C_2 to a third shape C_3 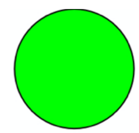 is about the same.

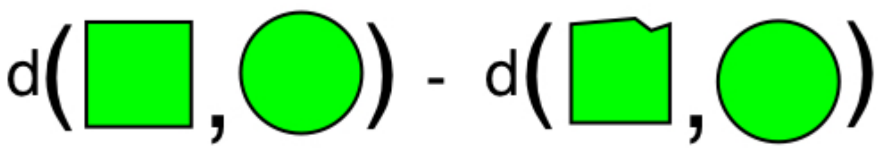

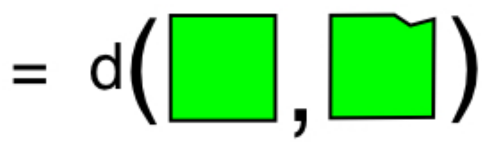

This means that noise, or a small error, does not affect distance measurements very much.

What is a good metric on the shapes in R^2?

David Mumford examined this question.

D. Mumford, 1991
Mathematical Theories of Shape: do they model perception?

There are many natural candidates for metrics giving distances between shapes.

We look at some of these metrics.

Hausdorff metric

d_H = Maximal distance of a point in one set from the other set, after a rigid motion.

$$d_H(A, B) = \min_{\text{rigid motions}} \{\sup_{x \in A} d(x, B) + \sup_{y \in B} d(y, A)\}$$

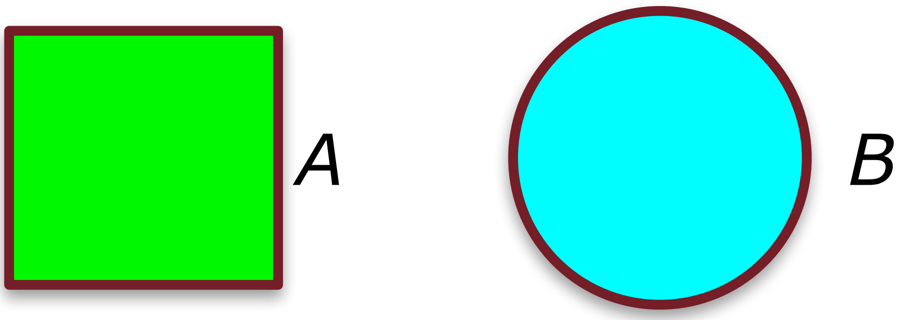

What is the Hausdorff distance?

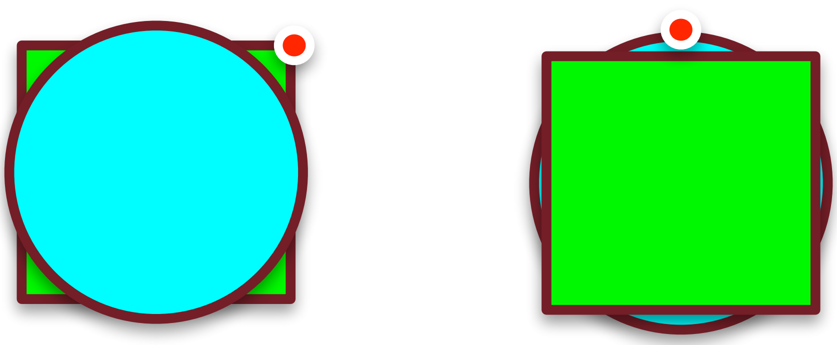

Add the distances of each red dot from the other set.

Gives a metric on {compact subsets of the plane}.

Drawbacks: Hausdorff metric

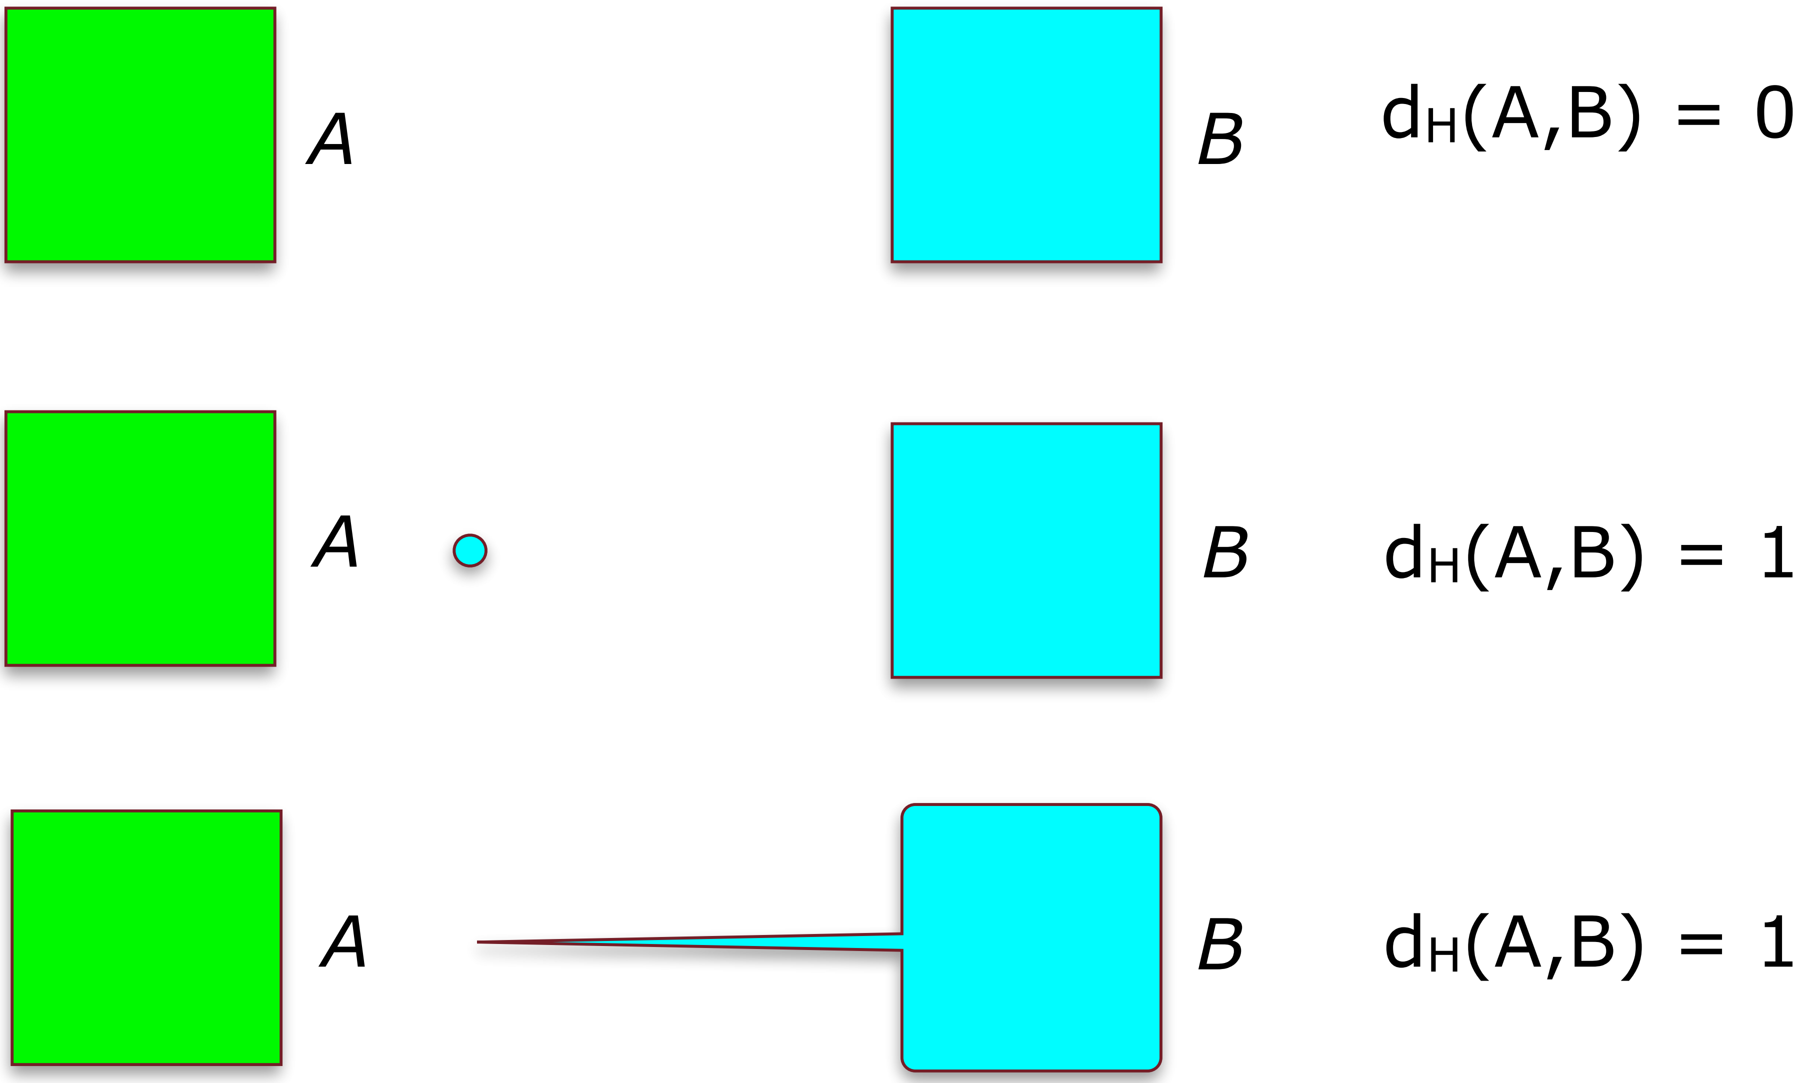

Drawbacks: Hausdorff metric

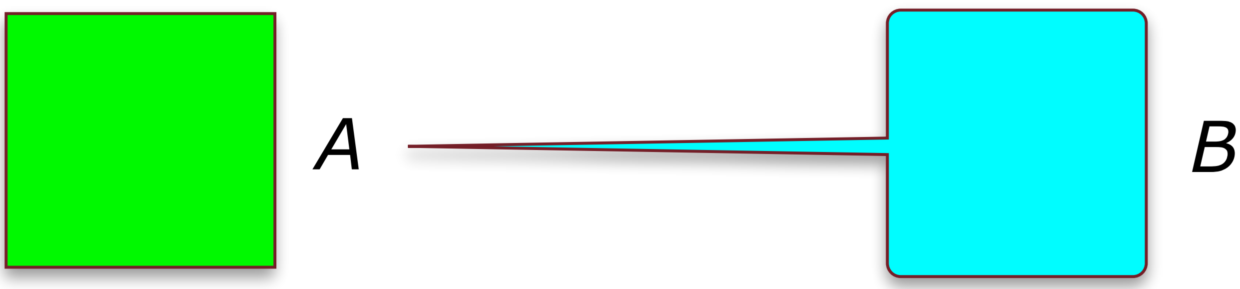

$d_H(A,B) = 1$

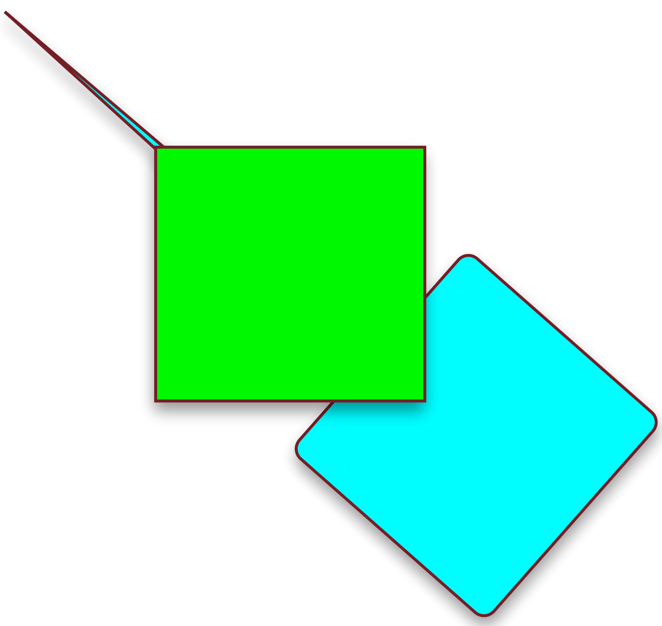

The alignment that minimizes Hausdorff distance may not give the correspondence we want.

Can we fix this with a different metric?

Template metric

distance = Area of non-overlap after rigid motion.

$$d_T(A, B) = \min_{\text{rigid motions}} \{Area(A-B) + Area(B-A)\}$$

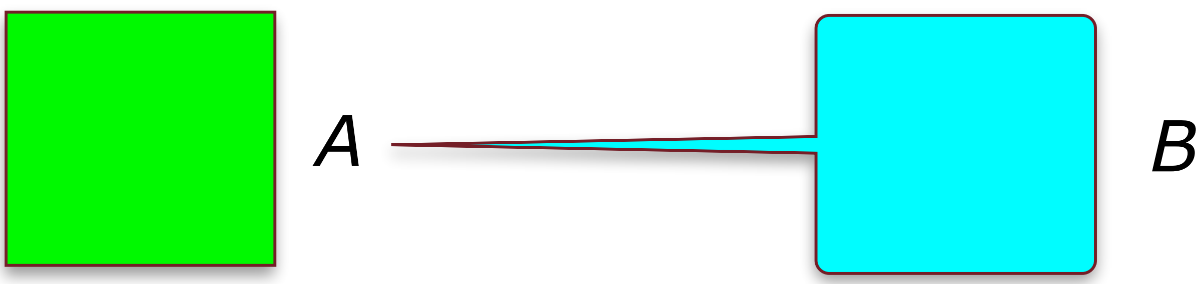

$d_T(A,B) \approx 0$

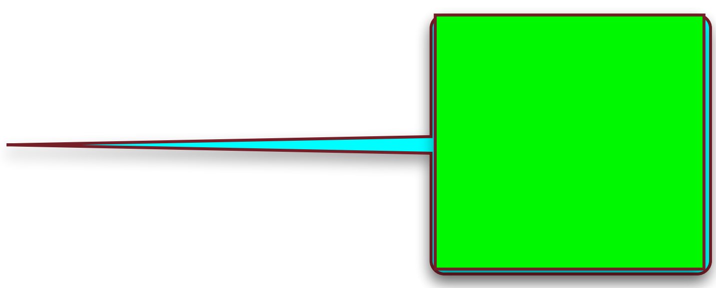

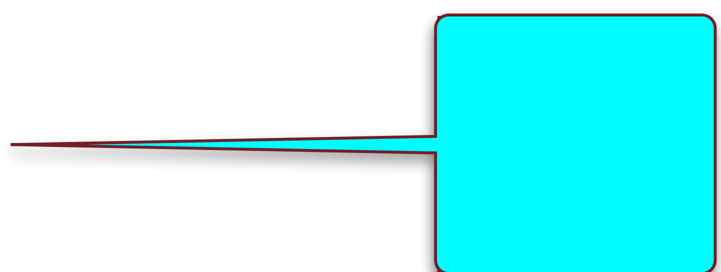

Blue area at left + green area at right

Drawbacks: Template metric

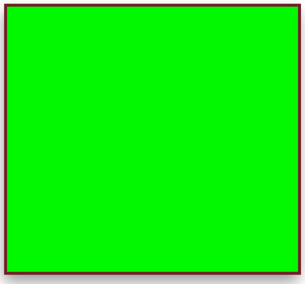

A

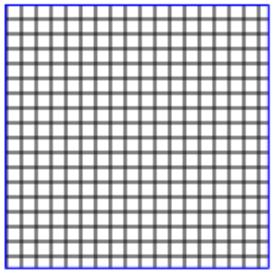

B

The area overlap is small. $d_T(A,B) \approx 1$

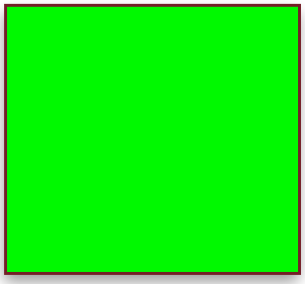

A

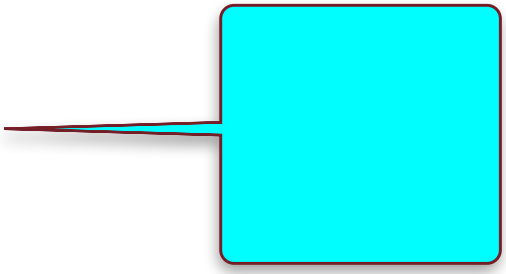

B

The area overlap is large. $d_T(A,B) \approx 0$

Challenge- Intrinsic geometry.

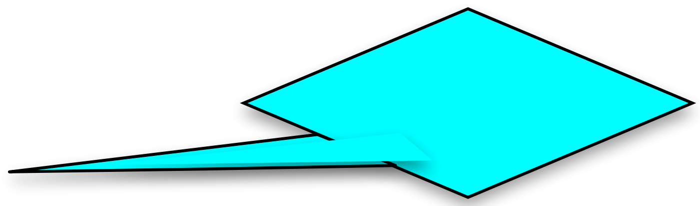

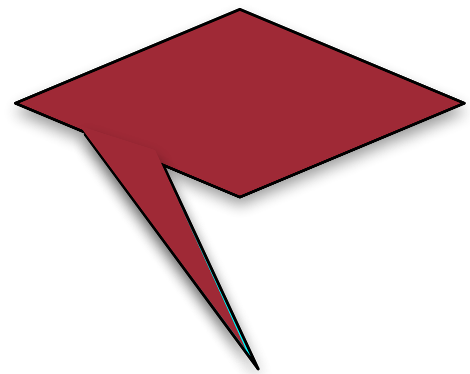

These shapes are *intrinsically* close. Not picked up by Hausdorff *or* template metrics.

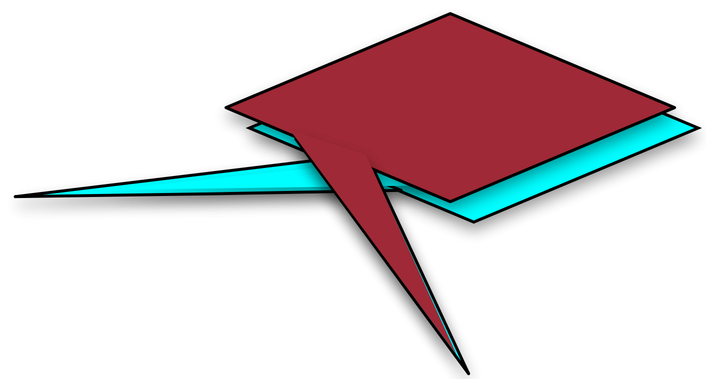

How can we see this?

Gromov-Hausdorff metric

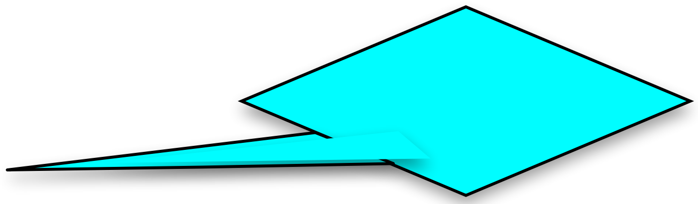

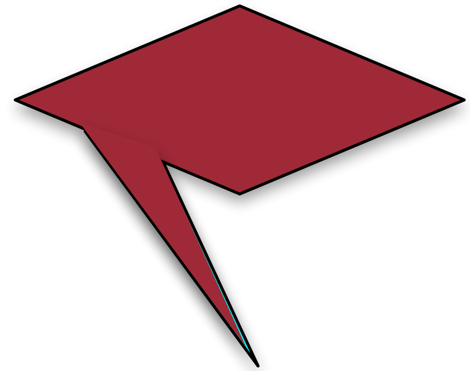

One way to see that these are close:
Bend them in R^3, and then use R^3-Hausdoff metric.
This gives the *Gromov-Hausdorff* metric.

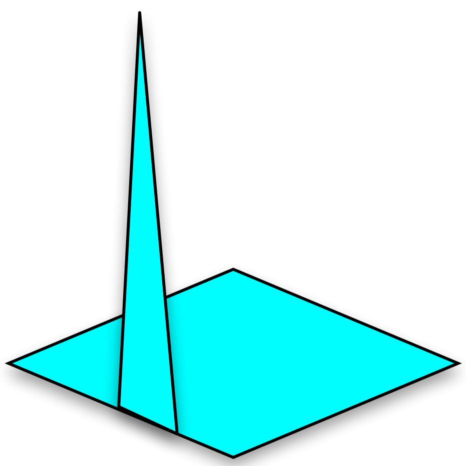

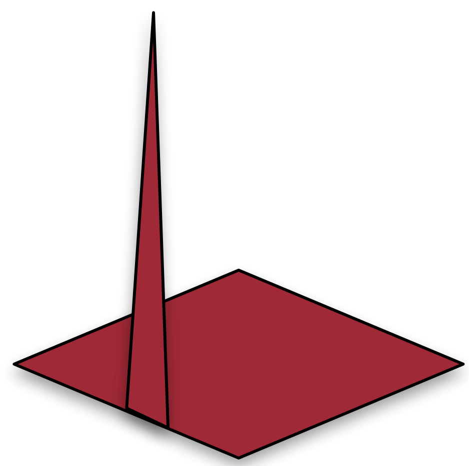

Optimal transport metric

Also called the *Wasserstein* or *Monge-Kantorovich* metric.
Distance between two shapes is the cost of moving one shape to the other:

$$\text{Distance} = \int \text{(area of subregion)} \times \text{(distance moved)}$$

Drawback - Optimal Transport

Can be discontinuous

Can be hard to compute

Optimal diffeomorphism metric

Define an *energy* that measures the stretching between two shapes.

This energy defines a distance between two spaces that are diffeomorphic.

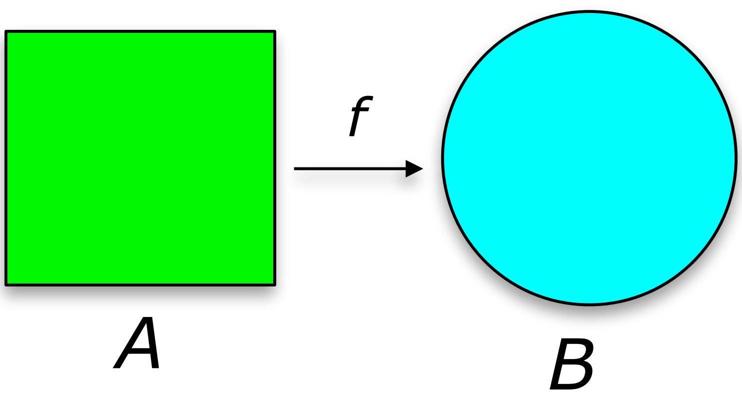

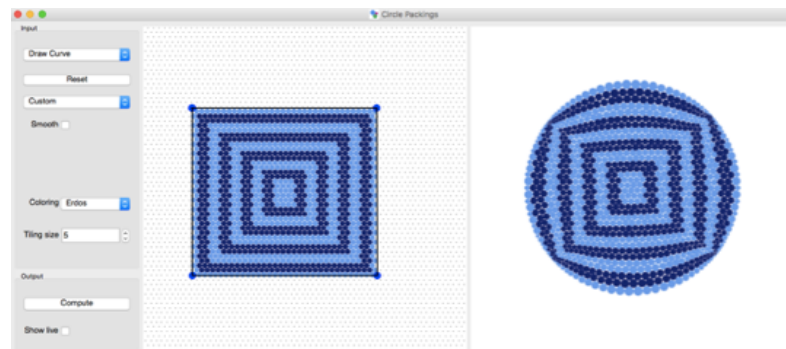

$$E(f) = \int\int (\partial f/\partial x)^2 + (\partial f/\partial x)^2 \, dxdy$$

$$d_D(A, B) = \min_{\text{diffeomorphisms}} \{E(f)\}$$

Drawback: Optimal diffeomorphism

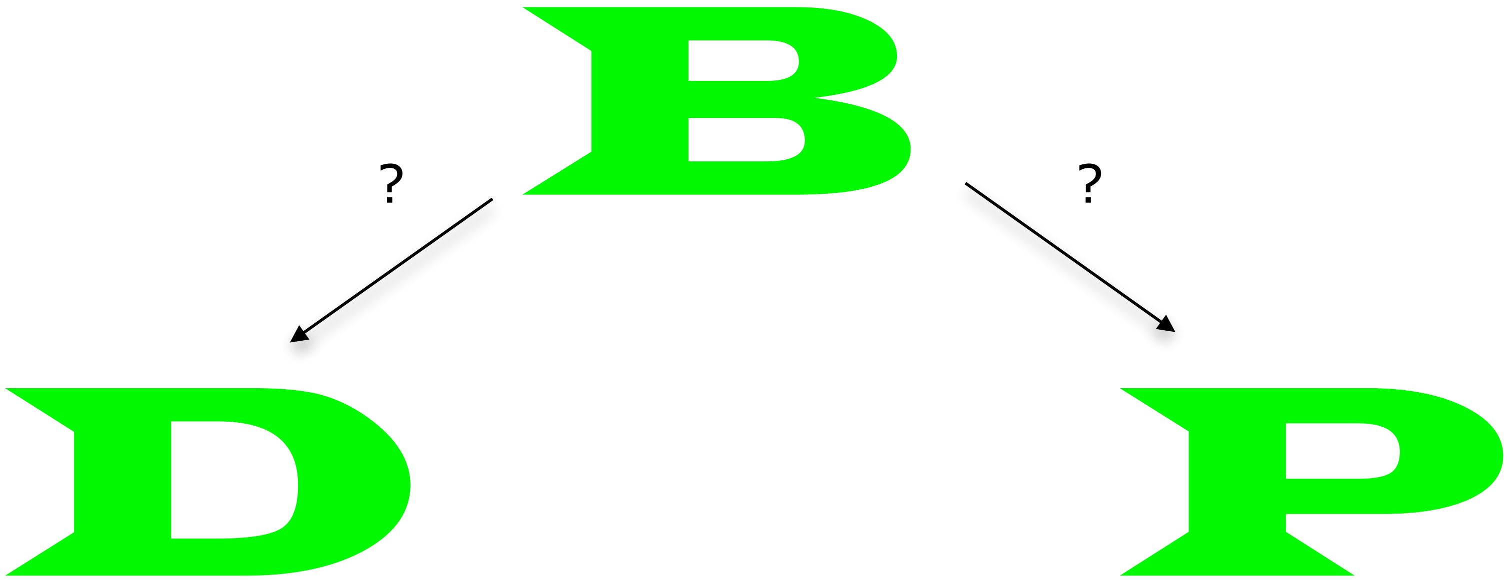

Requires diffeomorphic shapes

Maps with tears

Optimal diffeomorphism but allowing some tears.

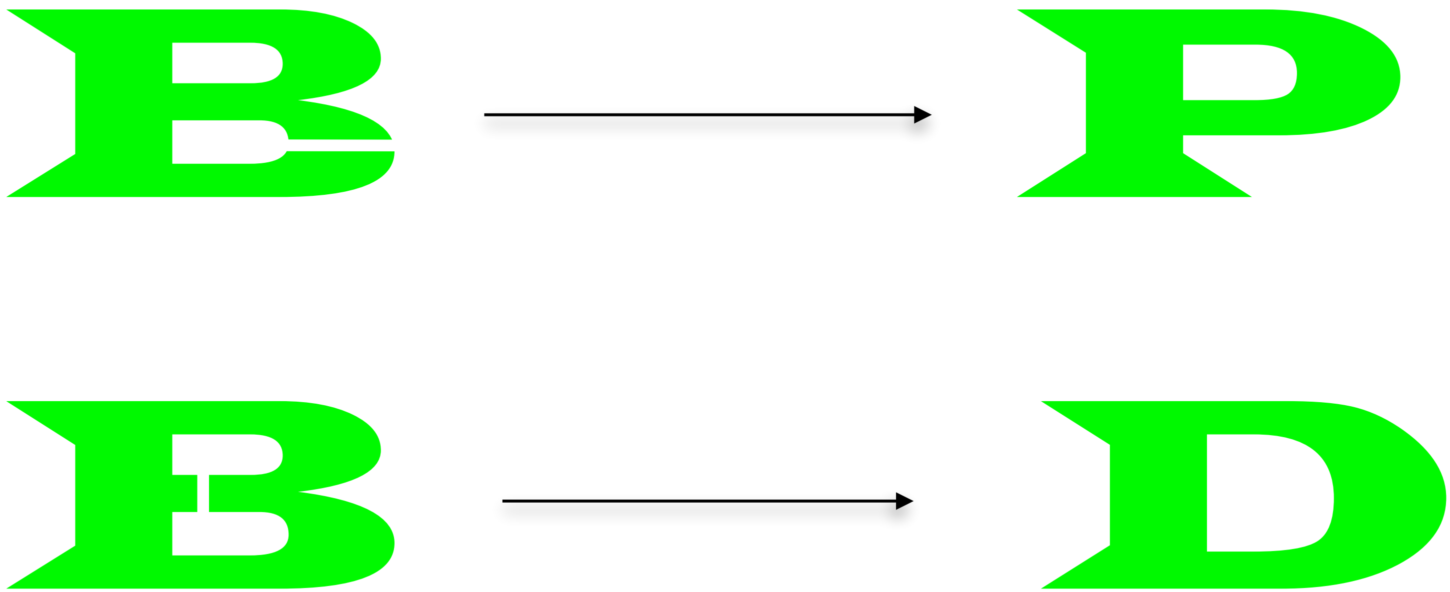

Hard to compute.

Mumford Experiments

Two groups of subjects, and 15 polygons

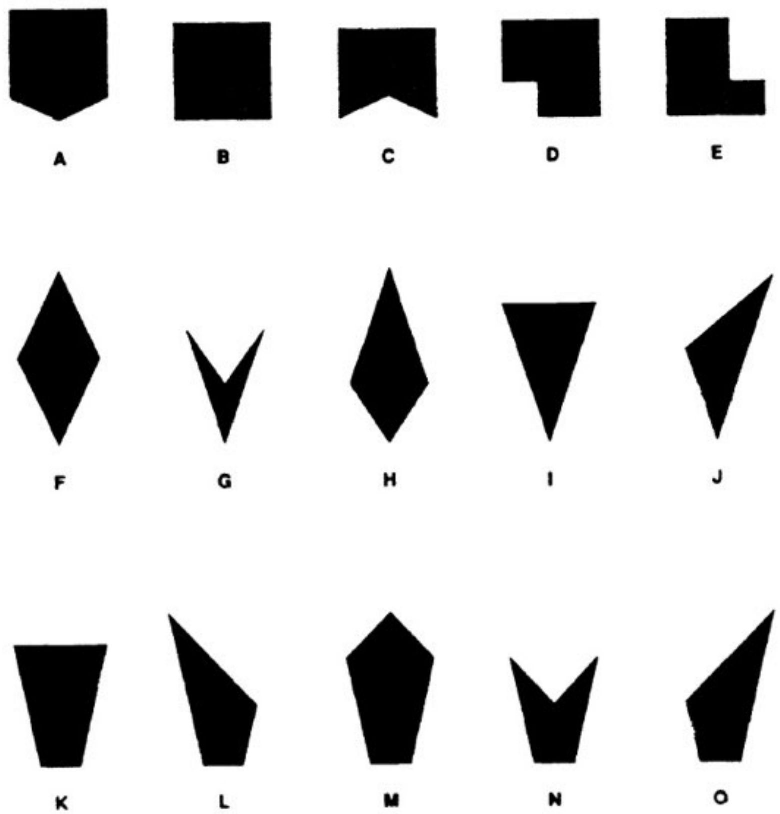

a. Pigeons

b. Harvard undergraduates

Experiment Conclusion: Human and pigeon perception of shape similarity do not indicate an underlying mathematical metric.

Deciphering Biological Shapes

-How do we understand shapes?

The Mumford experiments

-Shape Descriptors

Now look at surfaces and shapes in R^3

How similar are these two shapes?

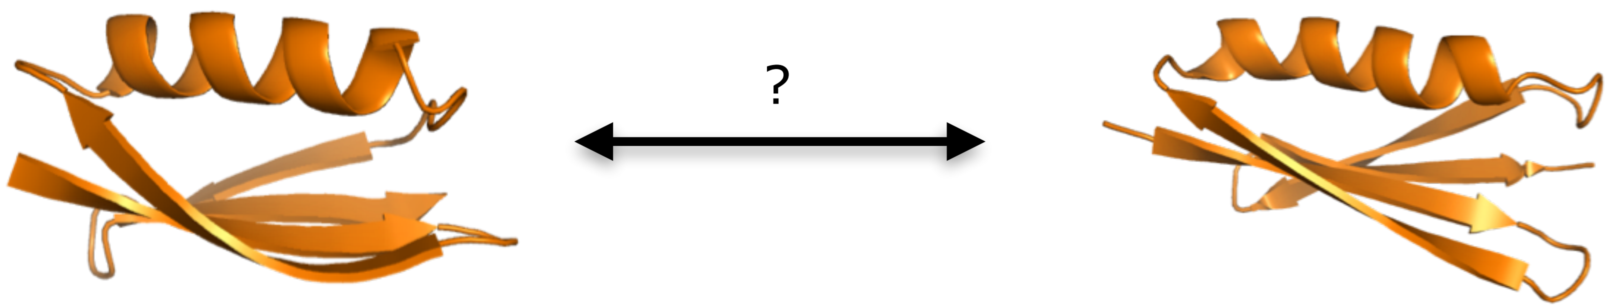

P1

P2

How do we compare two proteins?

"Feature space"

P1

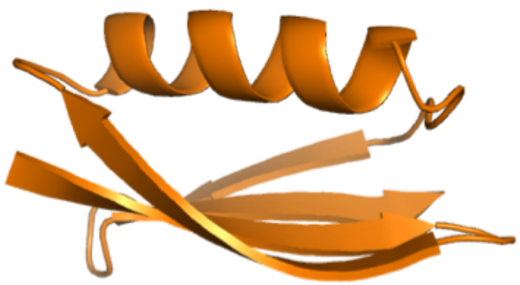

V1=(a1,b1,…..)

?

$$d = \|V1 - V2\|$$

P2

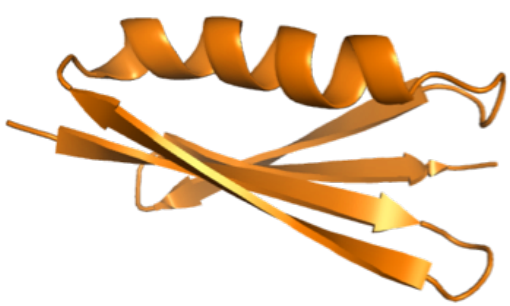

V2=(a2,b2,…..)

Fourier Analysis of Time Signal

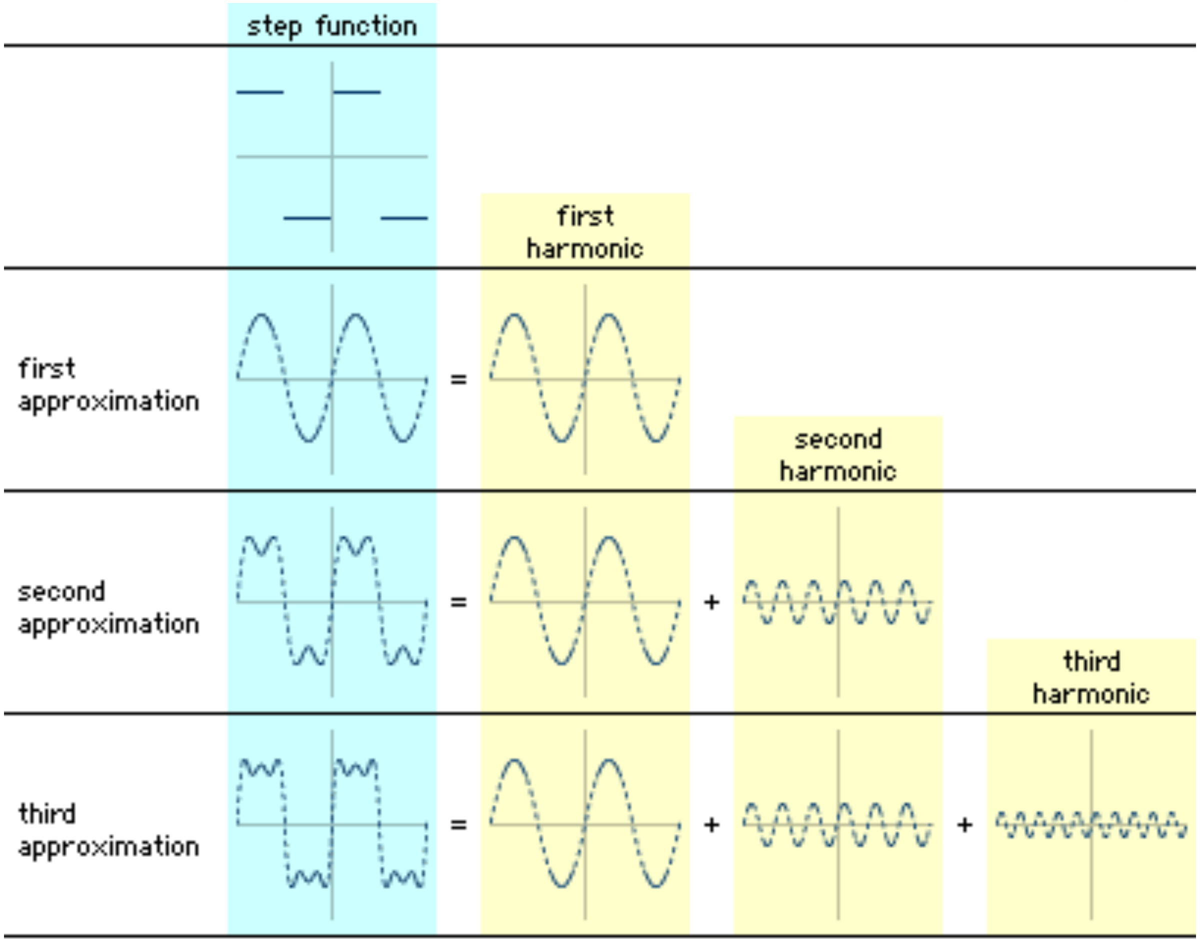

$$f_r = \frac{1}{N}\sum_{s=0}^{N-1} F_s e^{2\pi rs/N} \quad \longleftrightarrow \quad F_s = \frac{1}{N}\sum_{r=0}^{N-1} f_r e^{-2\pi rs/N}$$

Harmonic Representation of Shapes

1. Surface-based shape analysis
   Spherical harmonics

2. Volume-based shape analysis
   3D-Zernike moments

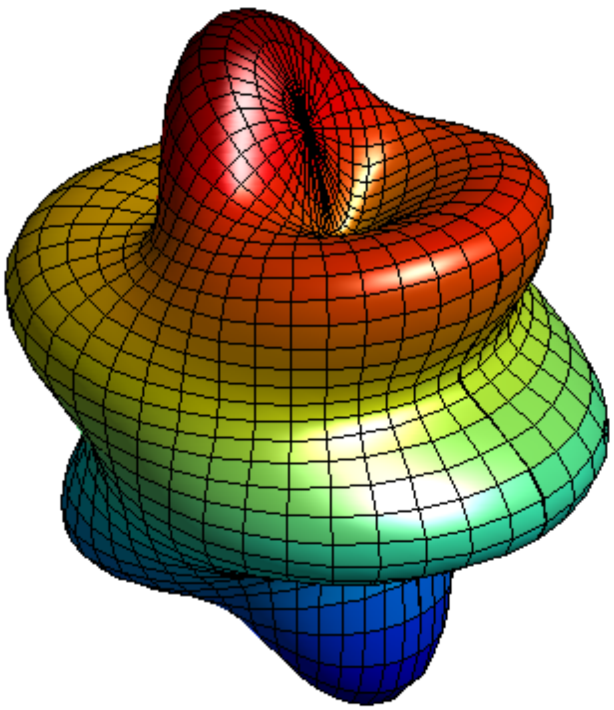

The challenge of the elephant…

Enrico Fermi once said to Freeman Dyson:

"I remember my friend Johnny von Neumann used to say, with four parameters I can fit an elephant, and with five I can make him wiggle his trunk."

(F. Dyson, Nature (London) 427, 297,2004)

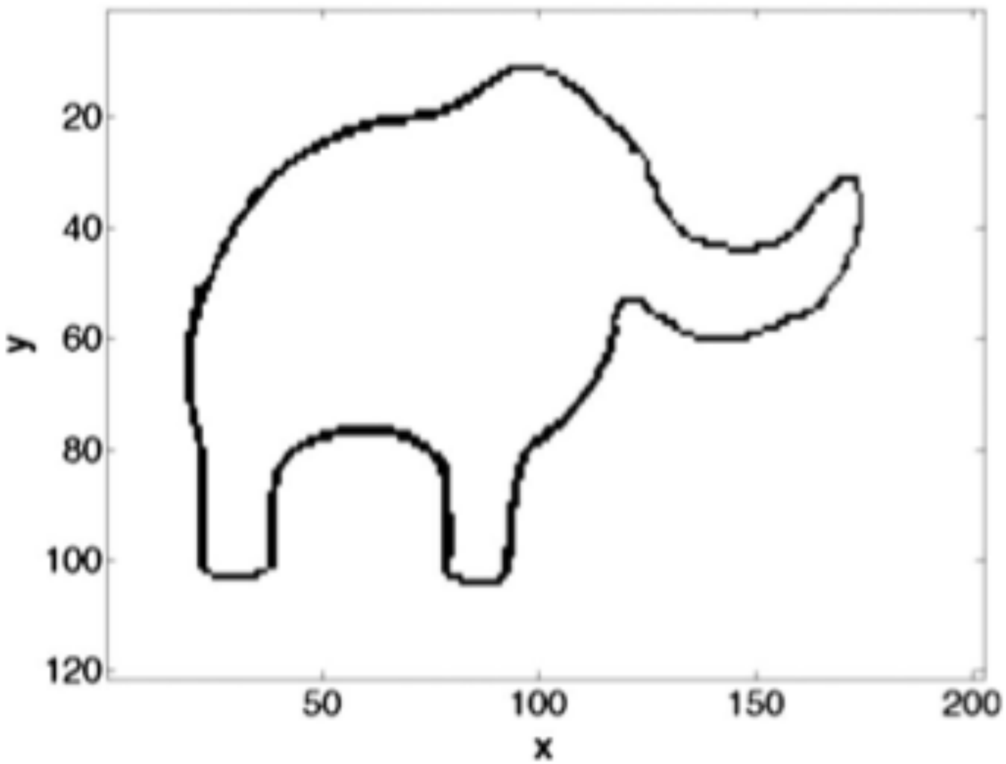

The challenge of the elephant…

The "best" solution, so far… (Mayer et al, Am. J. Phys. 78, 648-649,2010)

$$x(t) = \sum_{k=0}^{K} \left(A_k^x \cos(kt) + B_k^x \sin(kt) \right)$$

$$y(t) = \sum_{k=0}^{K} \left(A_k^y \cos(kt) + B_k^y \sin(kt) \right)$$

k	A_k^x	B_k^x	A_k^y	B_k^y
0	0	0	0	0
1	0	50	-60	-30
2	0	18	0	8
3	12	0	0	-10
4	0	0	0	0
5	0	50	0	0

The challenge of the elephant…

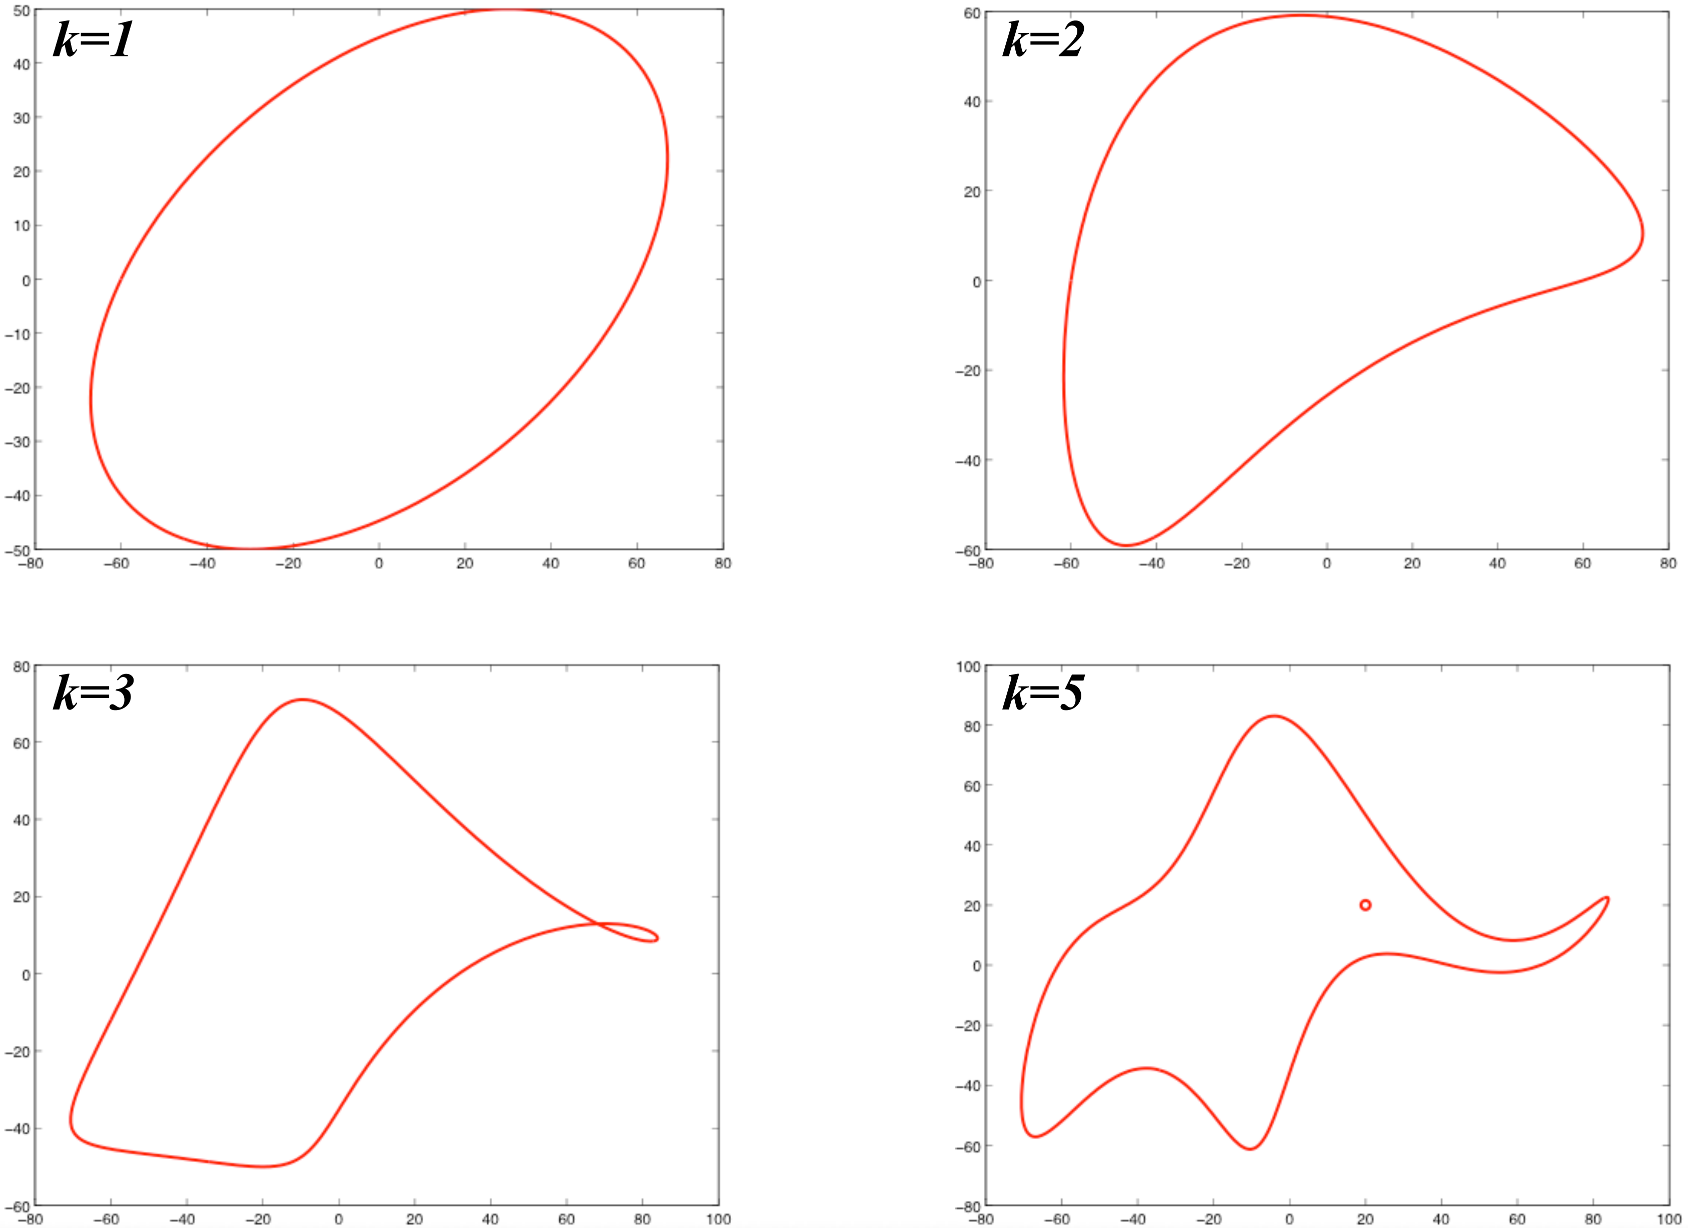

3D: Spherical harmonics

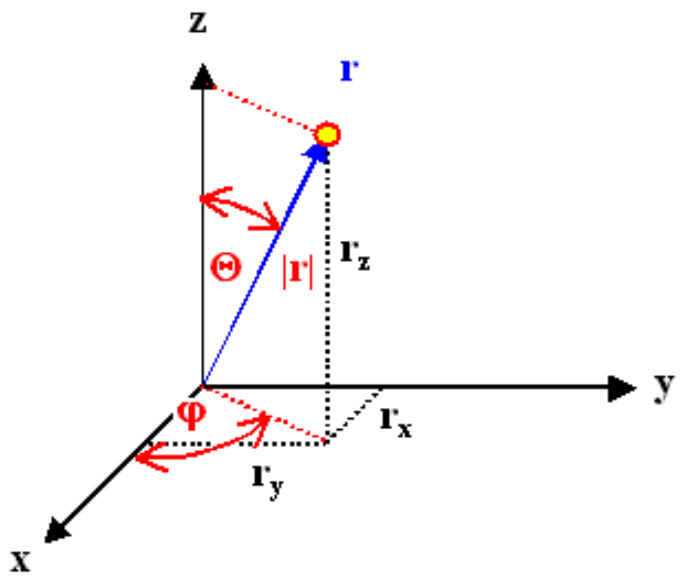

Any function f on the unit-sphere can be expanded into spherical harmonics:

$$f(\theta,\varphi) = \sum_{l=0}^{+\infty} \sum_{m=-l}^{l} c_{l,m} Y_l^m(\theta,\varphi)$$

where the basis functions are defined as:

$$Y_l^m(\theta,\varphi) = \sqrt{\frac{2l+1}{4\pi}\frac{(l-m)!}{(l+m)!}} P_l^m(\cos\theta) e^{im\varphi}$$

The coefficients $c_{l,m}$ are computed as:

$$c_{l.m} = \int_0^{2\pi}\int_0^{\pi} f(\theta,\varphi)\left(Y_l^m(\theta,\varphi)\right)^* \sin(\theta)d\theta d\varphi$$

3D: Spherical harmonics

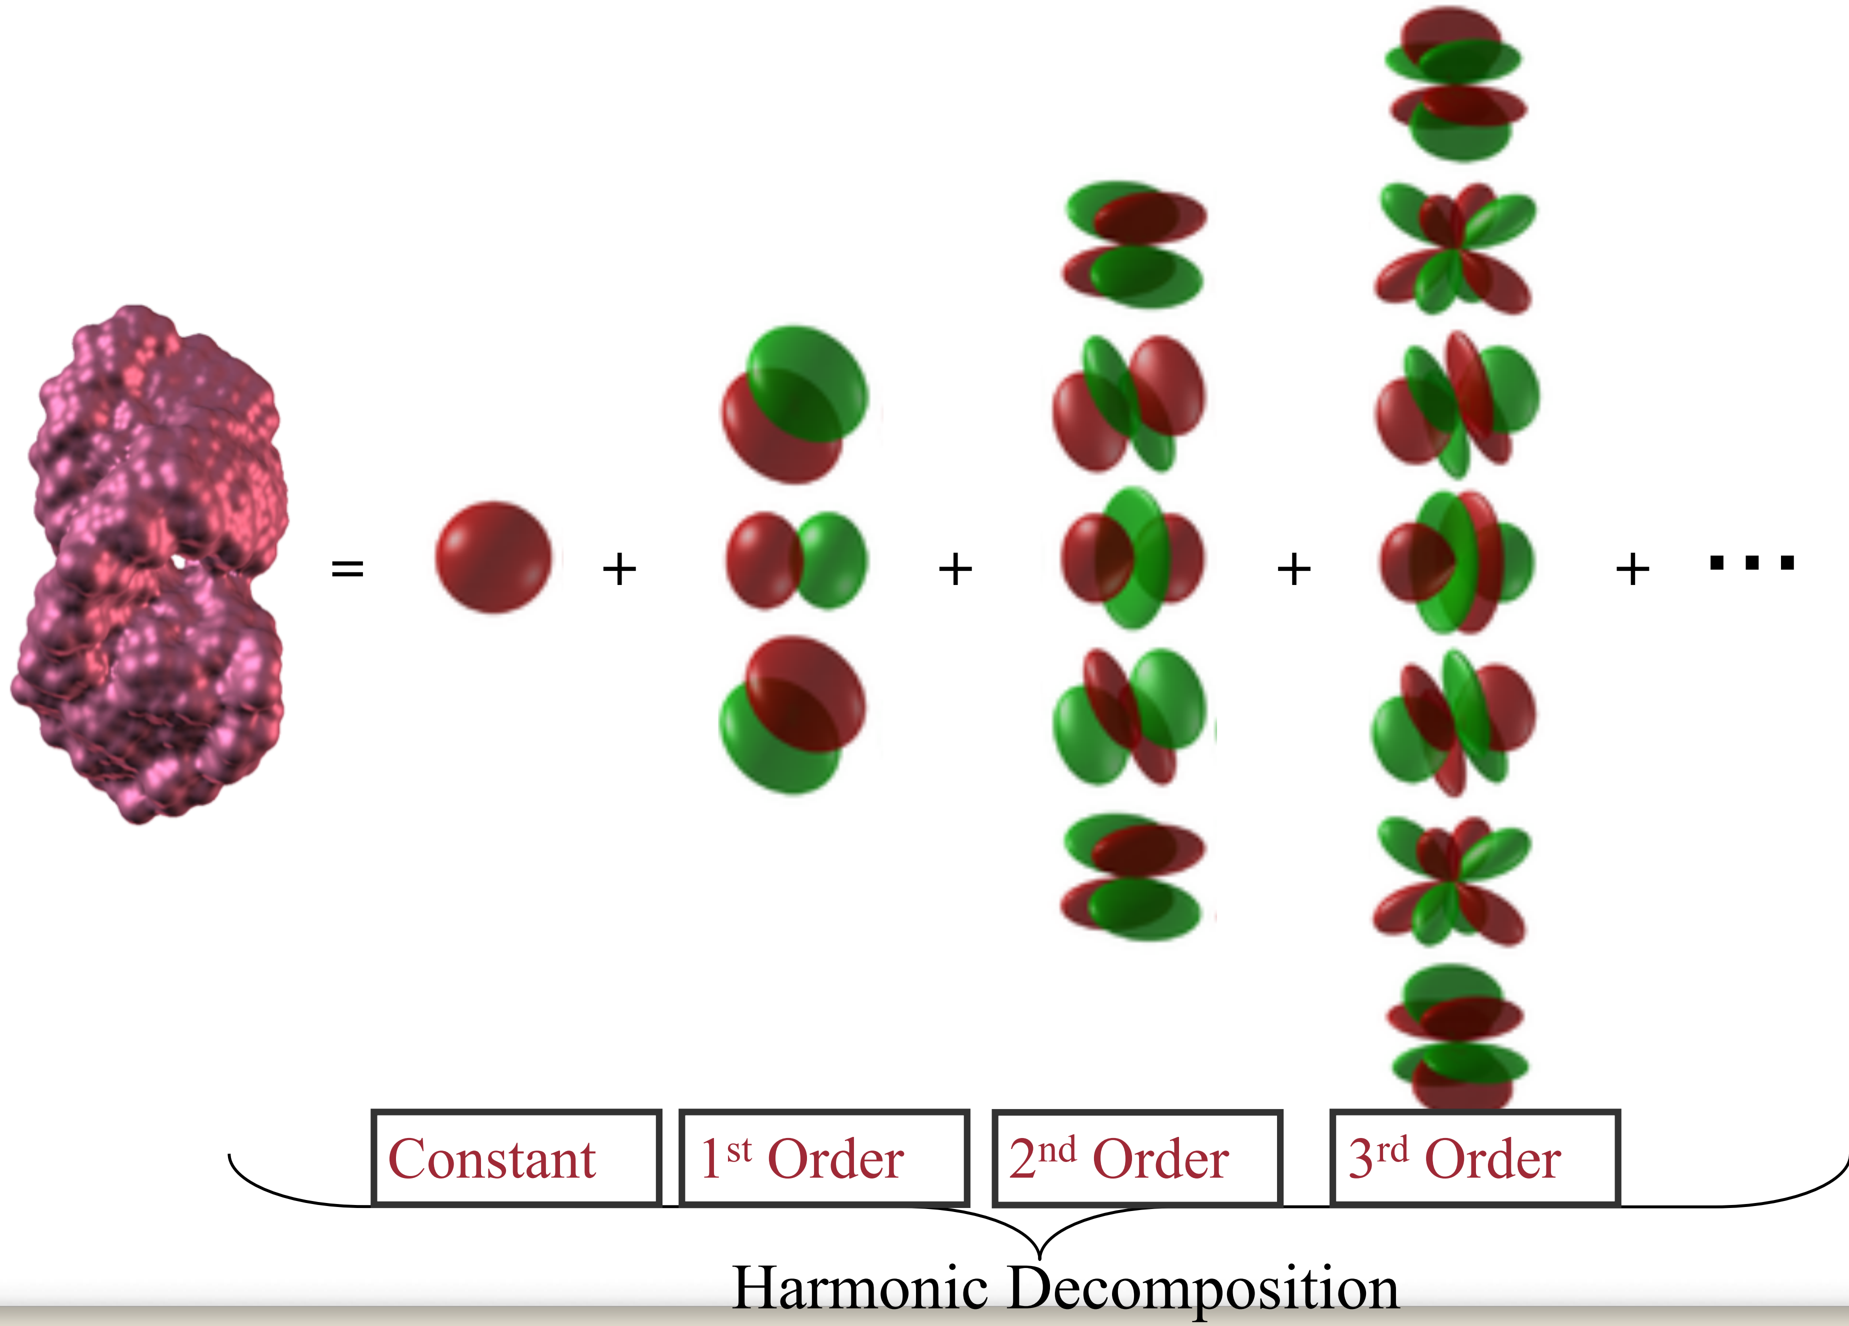

What are the spherical harmonics Y_l^m ?

$$Y_0^0(\theta,\varphi) = \frac{1}{2}\sqrt{\frac{1}{\pi}}$$

$$Y_1^{-1}(\theta,\varphi) = \frac{1}{2}\sqrt{\frac{3}{2\pi}}\,\sin\theta\,e^{-i\varphi}$$

$$Y_1^0(\theta,\varphi) = \frac{1}{2}\sqrt{\frac{3}{\pi}}\,\cos\theta$$

$$Y_1^1(\theta,\varphi) = \frac{-1}{2}\sqrt{\frac{3}{2\pi}}\,\sin\theta\,e^{i\varphi}$$

$$Y_2^{-2}(\theta,\varphi) = \frac{1}{4}\sqrt{\frac{15}{2\pi}}\,\sin^2\theta\,e^{-2i\varphi}$$

$$Y_2^{-1}(\theta,\varphi) = \frac{1}{2}\sqrt{\frac{15}{2\pi}}\,\sin\theta\,\cos\theta\,e^{-i\varphi}$$

$$Y_2^0(\theta,\varphi) = \frac{1}{4}\sqrt{\frac{5}{\pi}}\,(3\cos^2\theta - 1)$$

$$Y_2^1(\theta,\varphi) = \frac{-1}{2}\sqrt{\frac{15}{2\pi}}\,\sin\theta\,\cos\theta\,e^{i\varphi}$$

$$Y_2^2(\theta,\varphi) = \frac{1}{4}\sqrt{\frac{15}{2\pi}}\,\sin^2\theta\,e^{2i\varphi}$$

Importance of Rotational Invariance

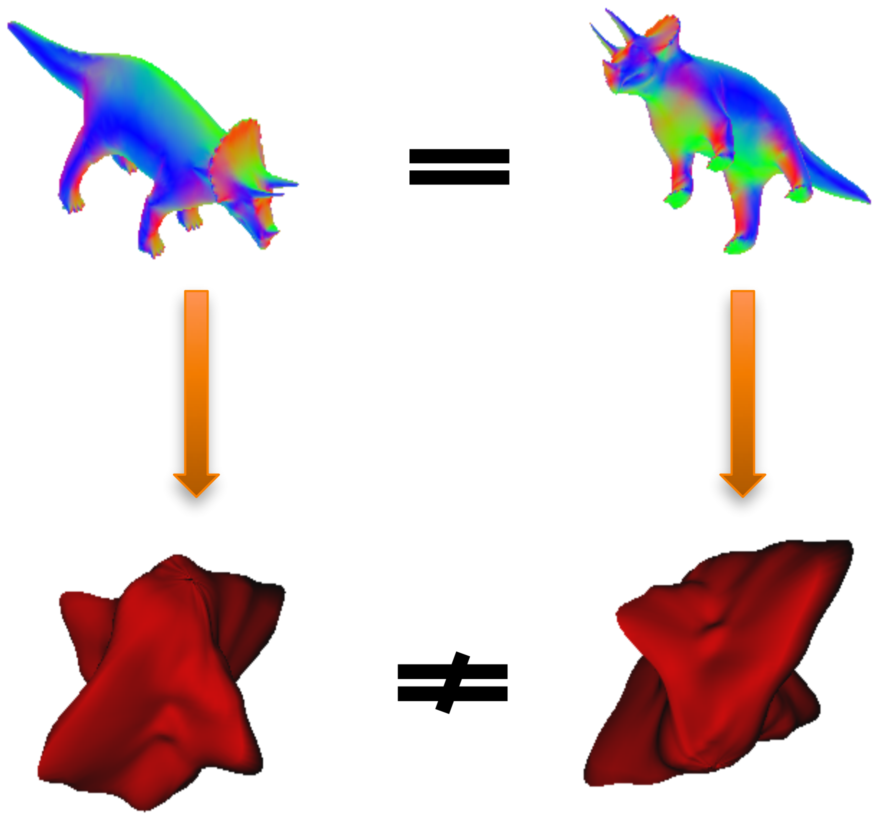

Shapes are unchanged by rotation

Shape descriptors may be sensitive to rotation: for example, the $c_{l,m}$ *are not rotation invariant*

Restoring Rotational Invariance

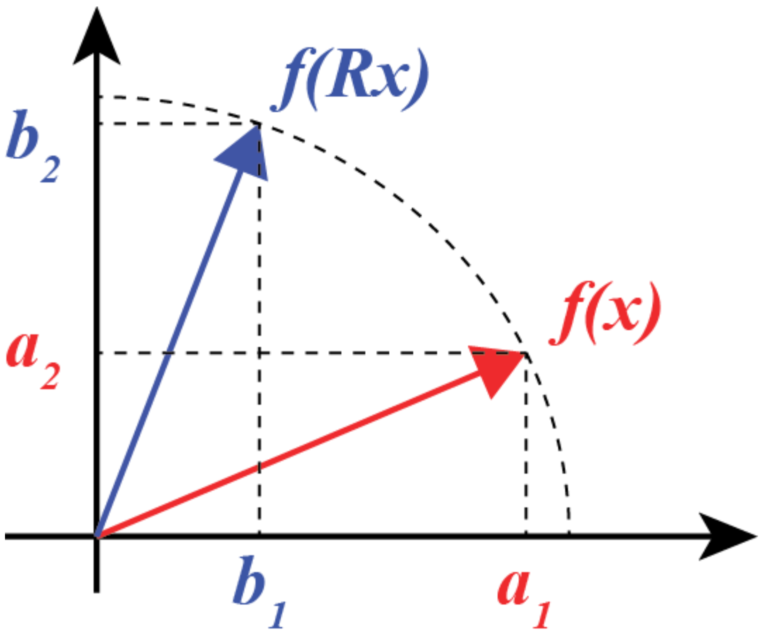

Note that:

$$f(x) = \begin{bmatrix} a_1 \\ a_2 \end{bmatrix} \neq \begin{bmatrix} b_1 \\ b_2 \end{bmatrix} = f(Rx)$$

However:

$$\|f(x)\| = \sqrt{a_1^2 + a_2^2} = \sqrt{b_1^2 + b_2^2} = \|f(Rx)\|$$

Invariant spherical harmonics descriptors:

$c_{l,m}$ for all l, m ➡ $g_l = \sqrt{\sum_{m=-l}^{l} c_{l,m}^2}$

Invariant spherical harmonics descriptors

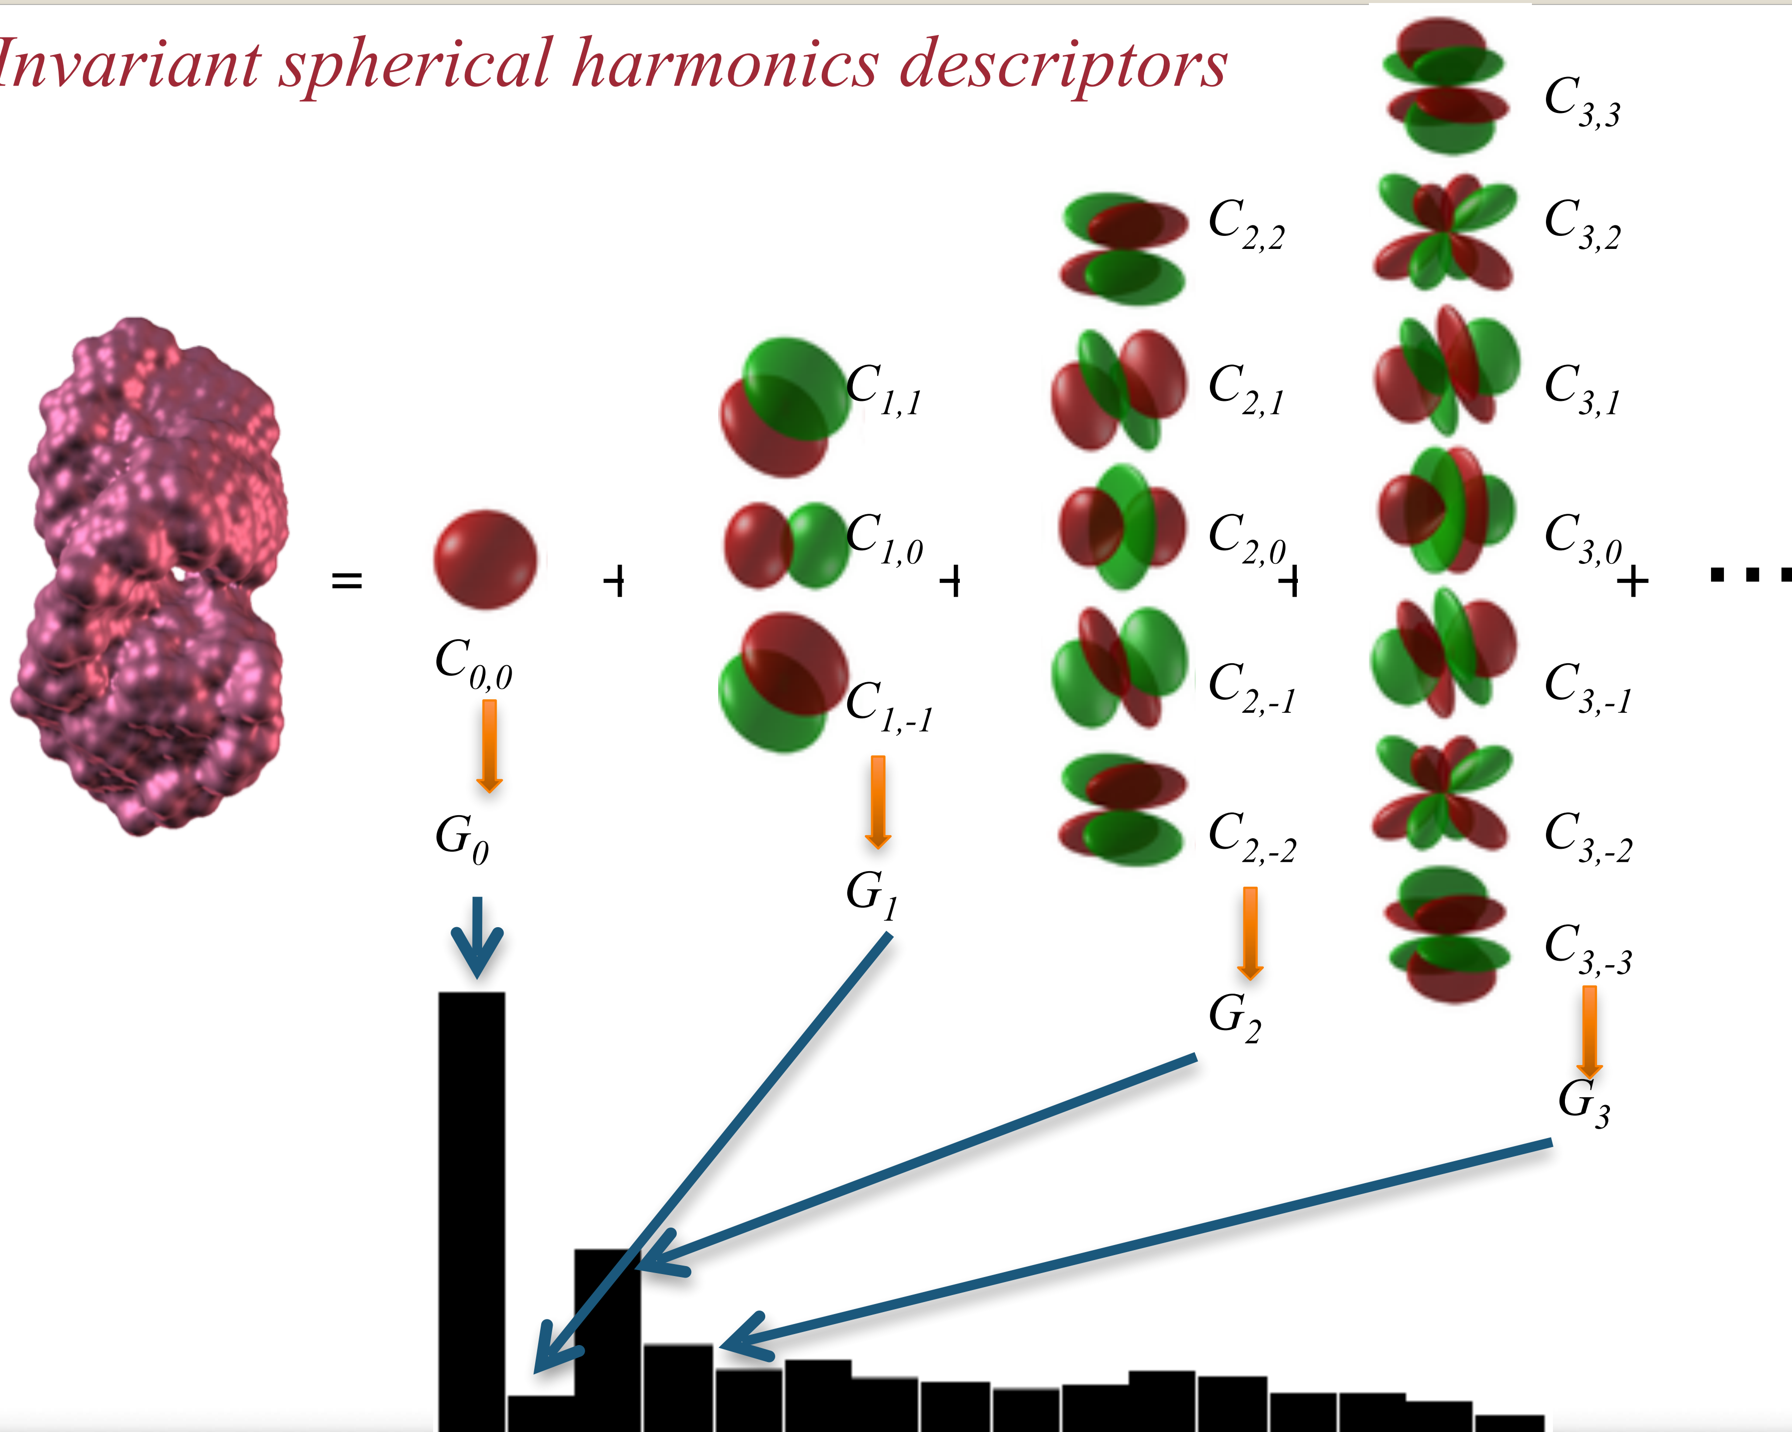

Some issues with Spherical Harmonics

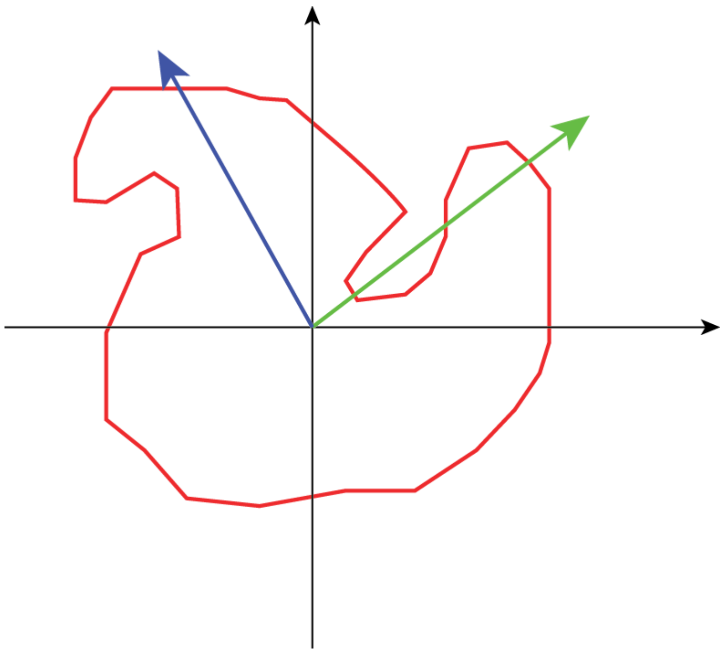

Spherical harmonics are surface-based:

-They require a parametrization of the surface (usually triangulation)

-They are appropriate for star-shaped objects

-They lose content information

From Surface to Volume

- Consider a set of concentric spheres over the object
- Compute harmonic representation of each sphere independently

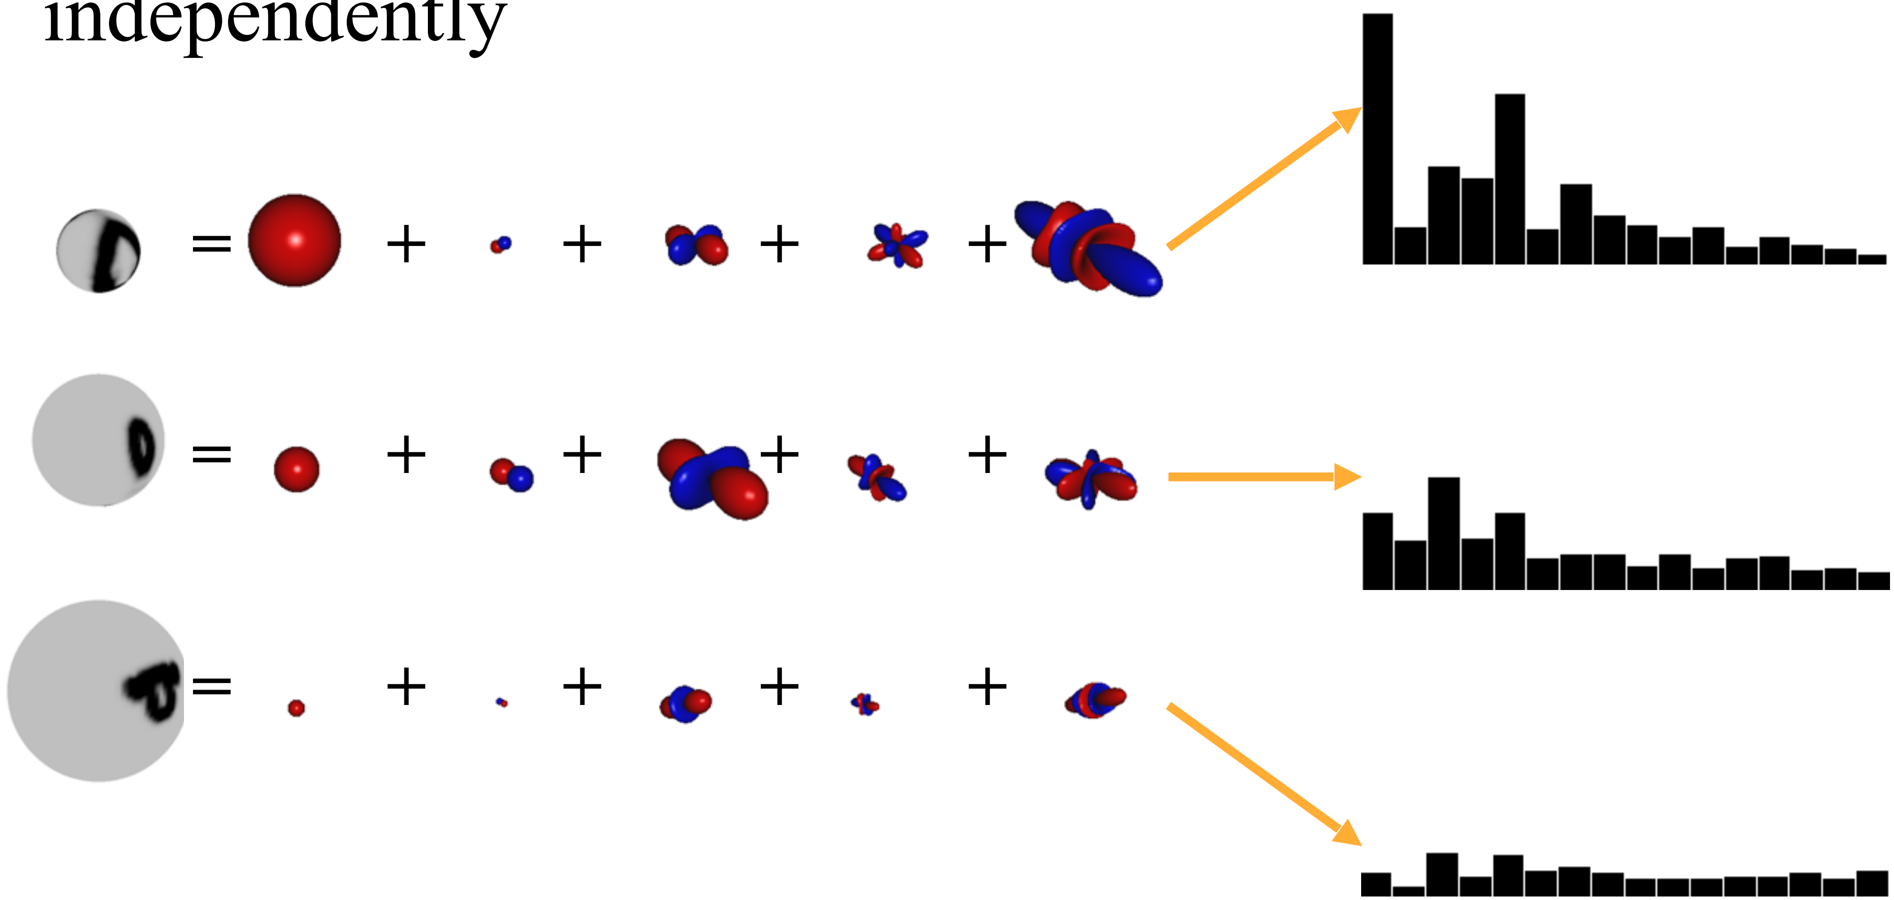

Problem: insensitive to internal rotations

A natural extension to Spherical Harmonics: The 3D Zernike moments

Surface-based

$$f(\theta,\varphi) = \sum_{l=0}^{+\infty} \sum_{m=-l}^{l} c_{l,m} Y_l^m(\theta,\varphi)$$

→

Volume-based

$$f(\theta,\varphi,r) = \sum_{n=0}^{+\infty} \sum_{l=0}^{+\infty} \sum_{m=-l}^{l} c_{l,m} R_{n,l} Y_l^m(\theta,\varphi)$$

with:

$$Y_l^m(\theta,\varphi) = \sqrt{\frac{2l+1}{4\pi}\frac{(l-m)!}{(l+m)!}} P_l^m(\cos\theta) e^{im\varphi}$$

and

$$R_{nl}(r) = \begin{cases} \sum_{k=0}^{(n-l)/2} N_{nlk} r^{n-2k} & n-l \quad even \\ 0 & n-l \quad odd \end{cases}$$

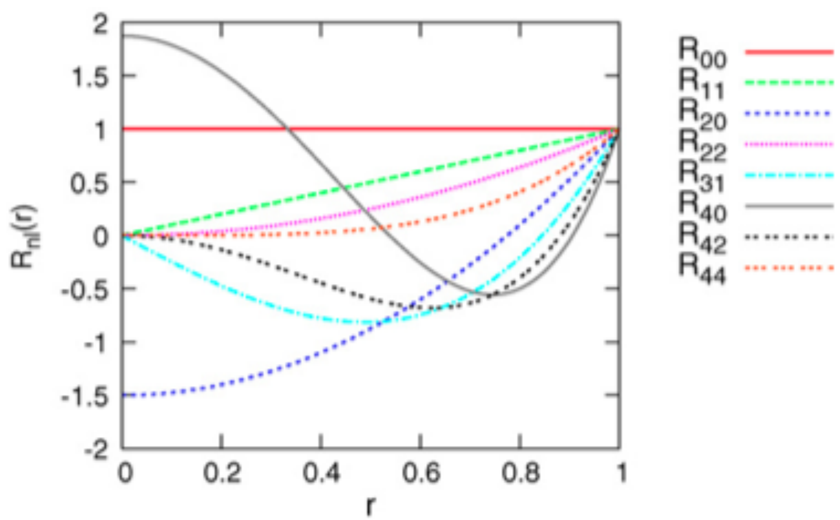

How does it work?

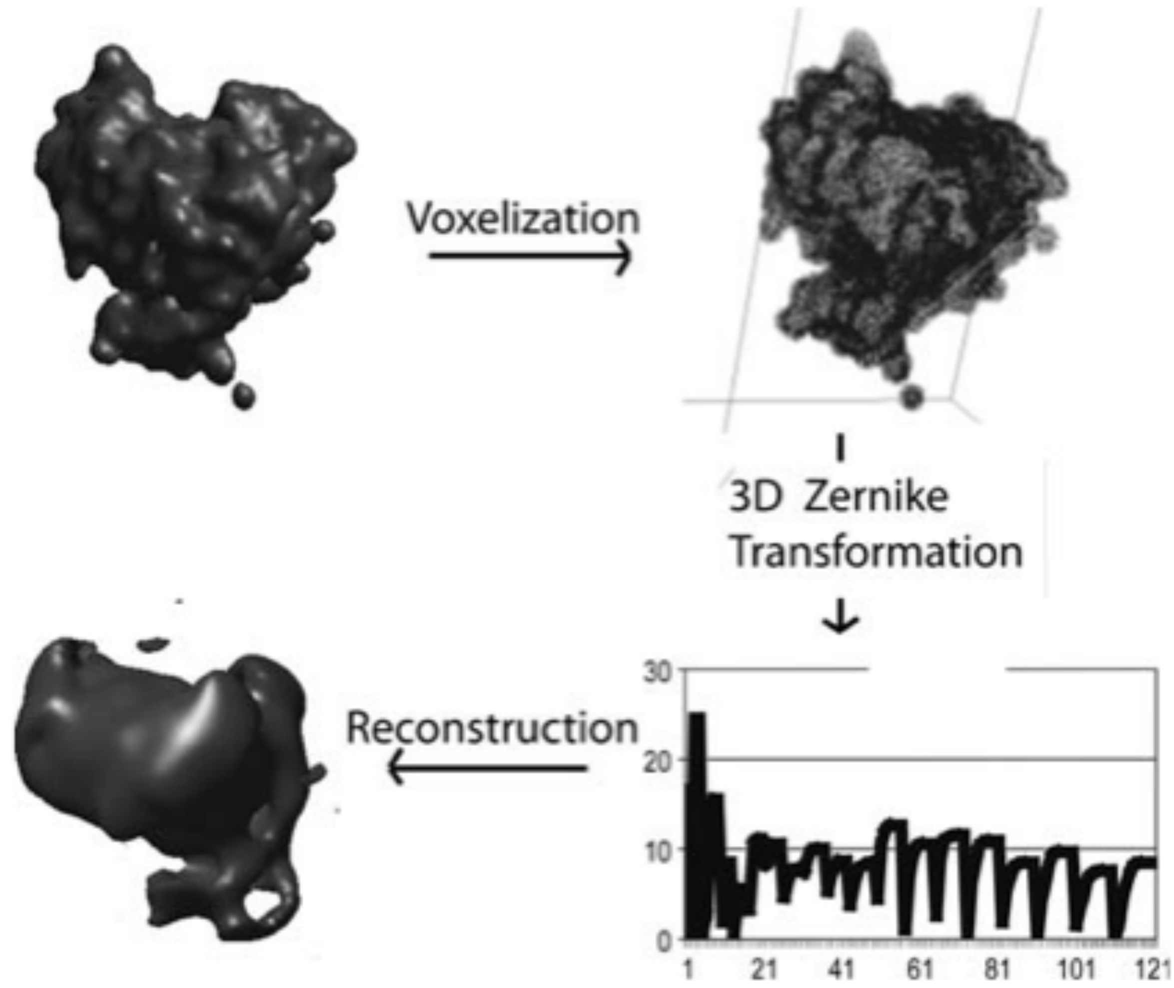

Applications

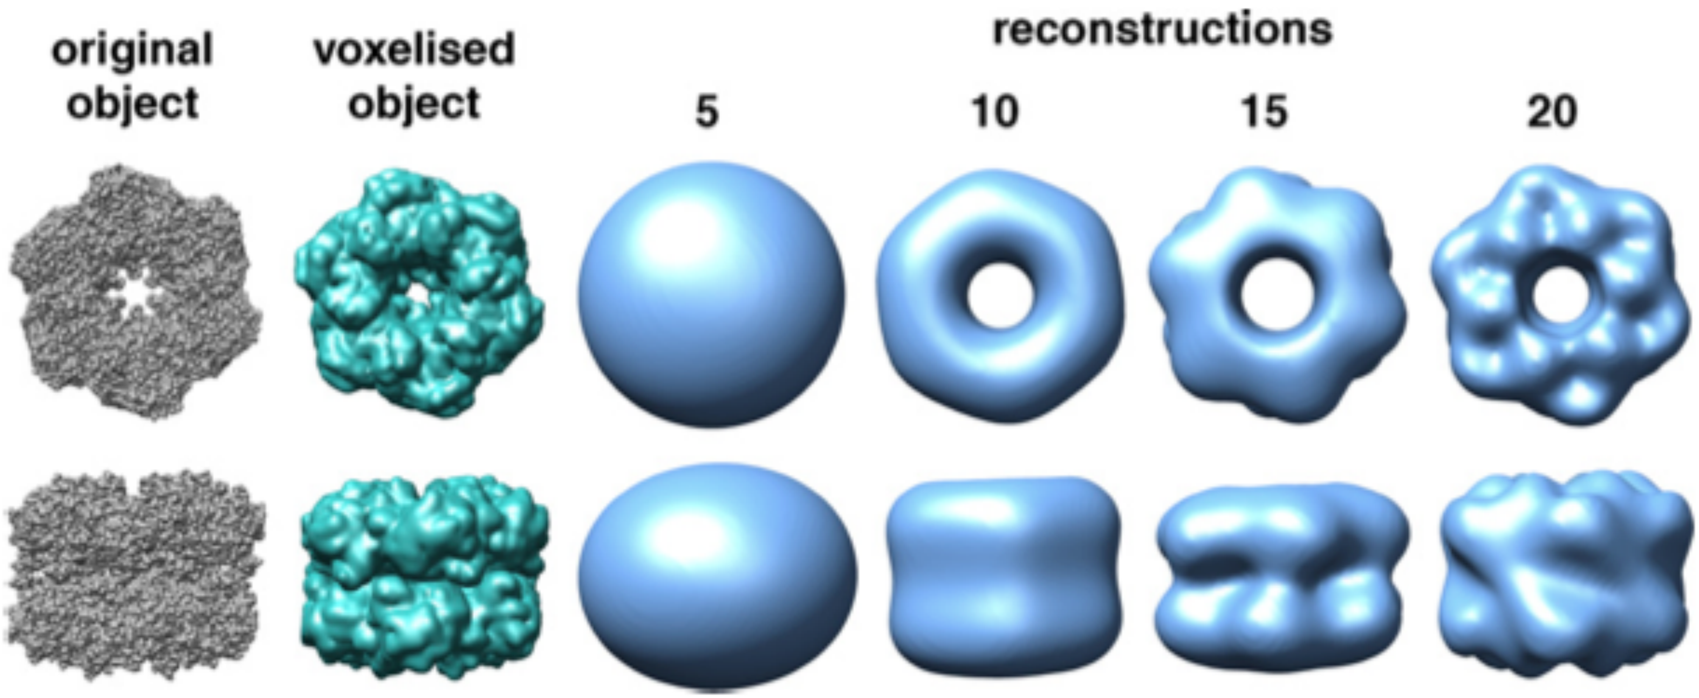

Comparing Old World Monkey Skulls

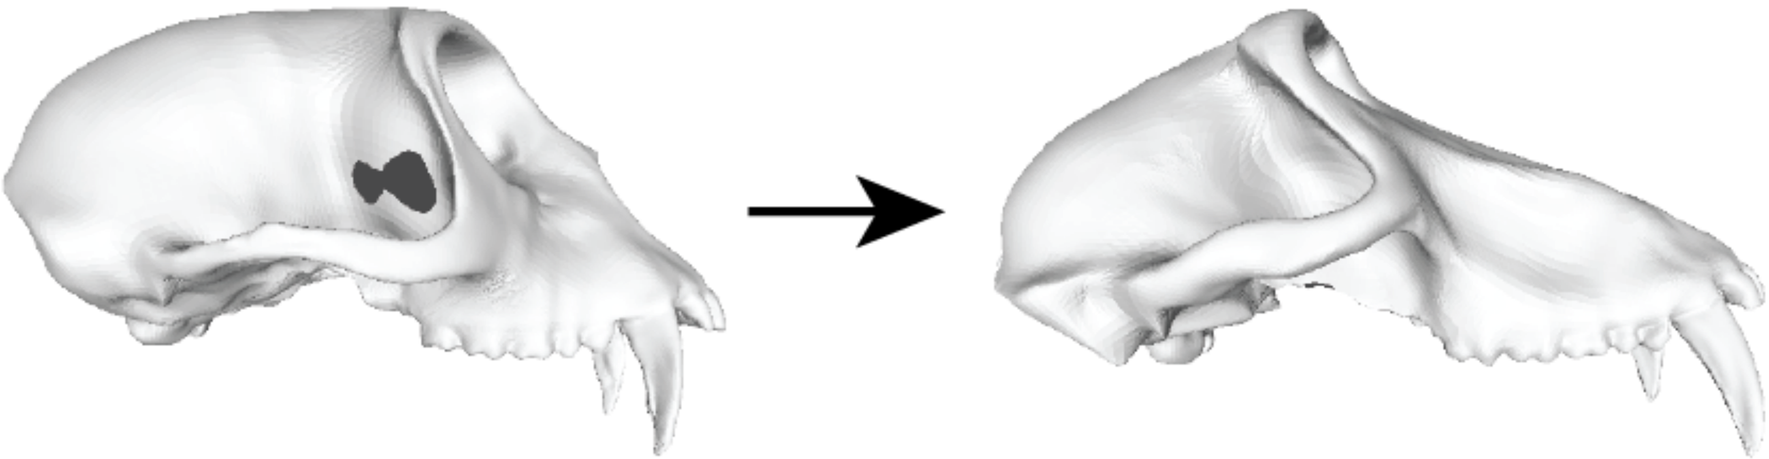

Old World Monkey Skulls: DNA Tree

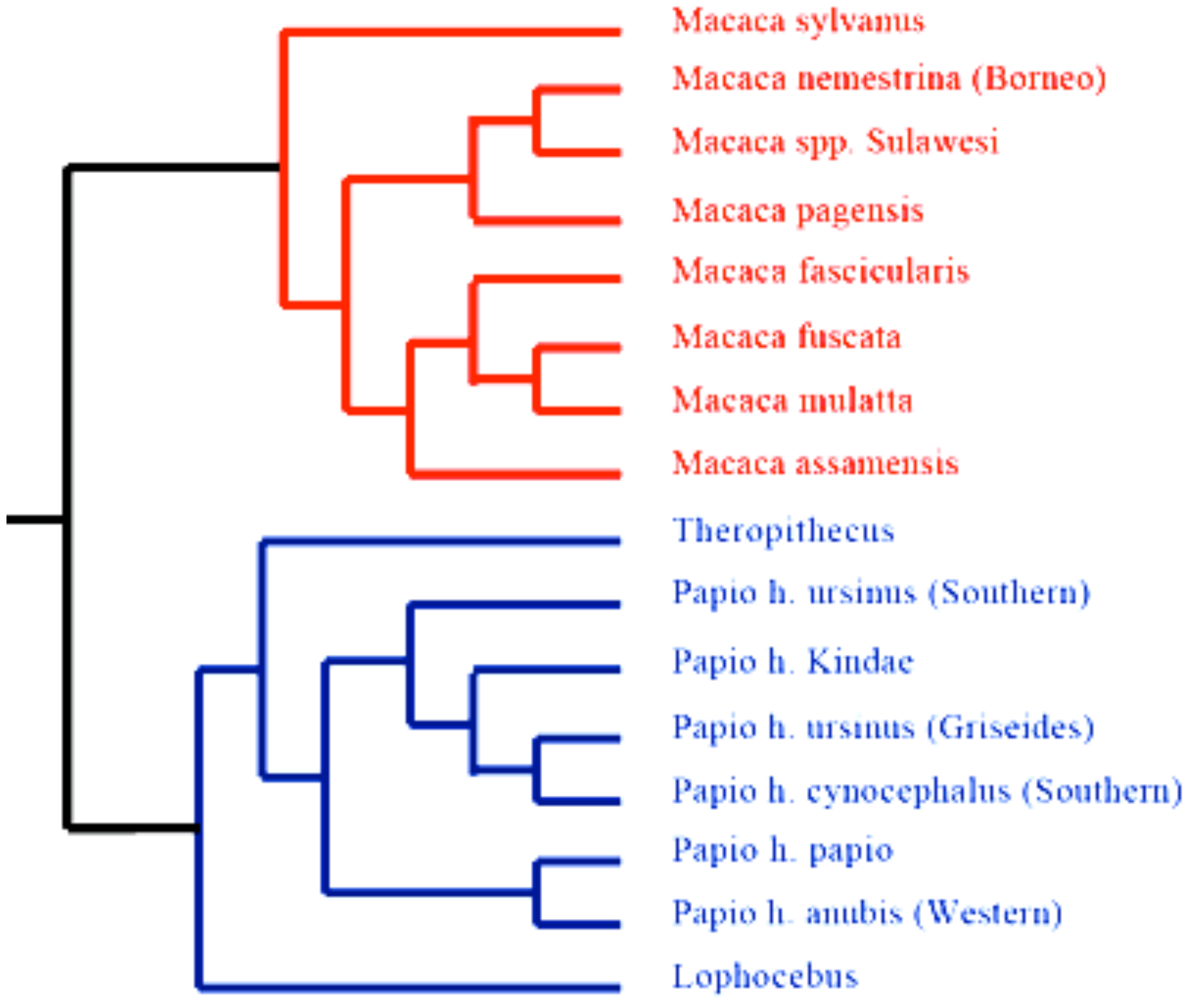

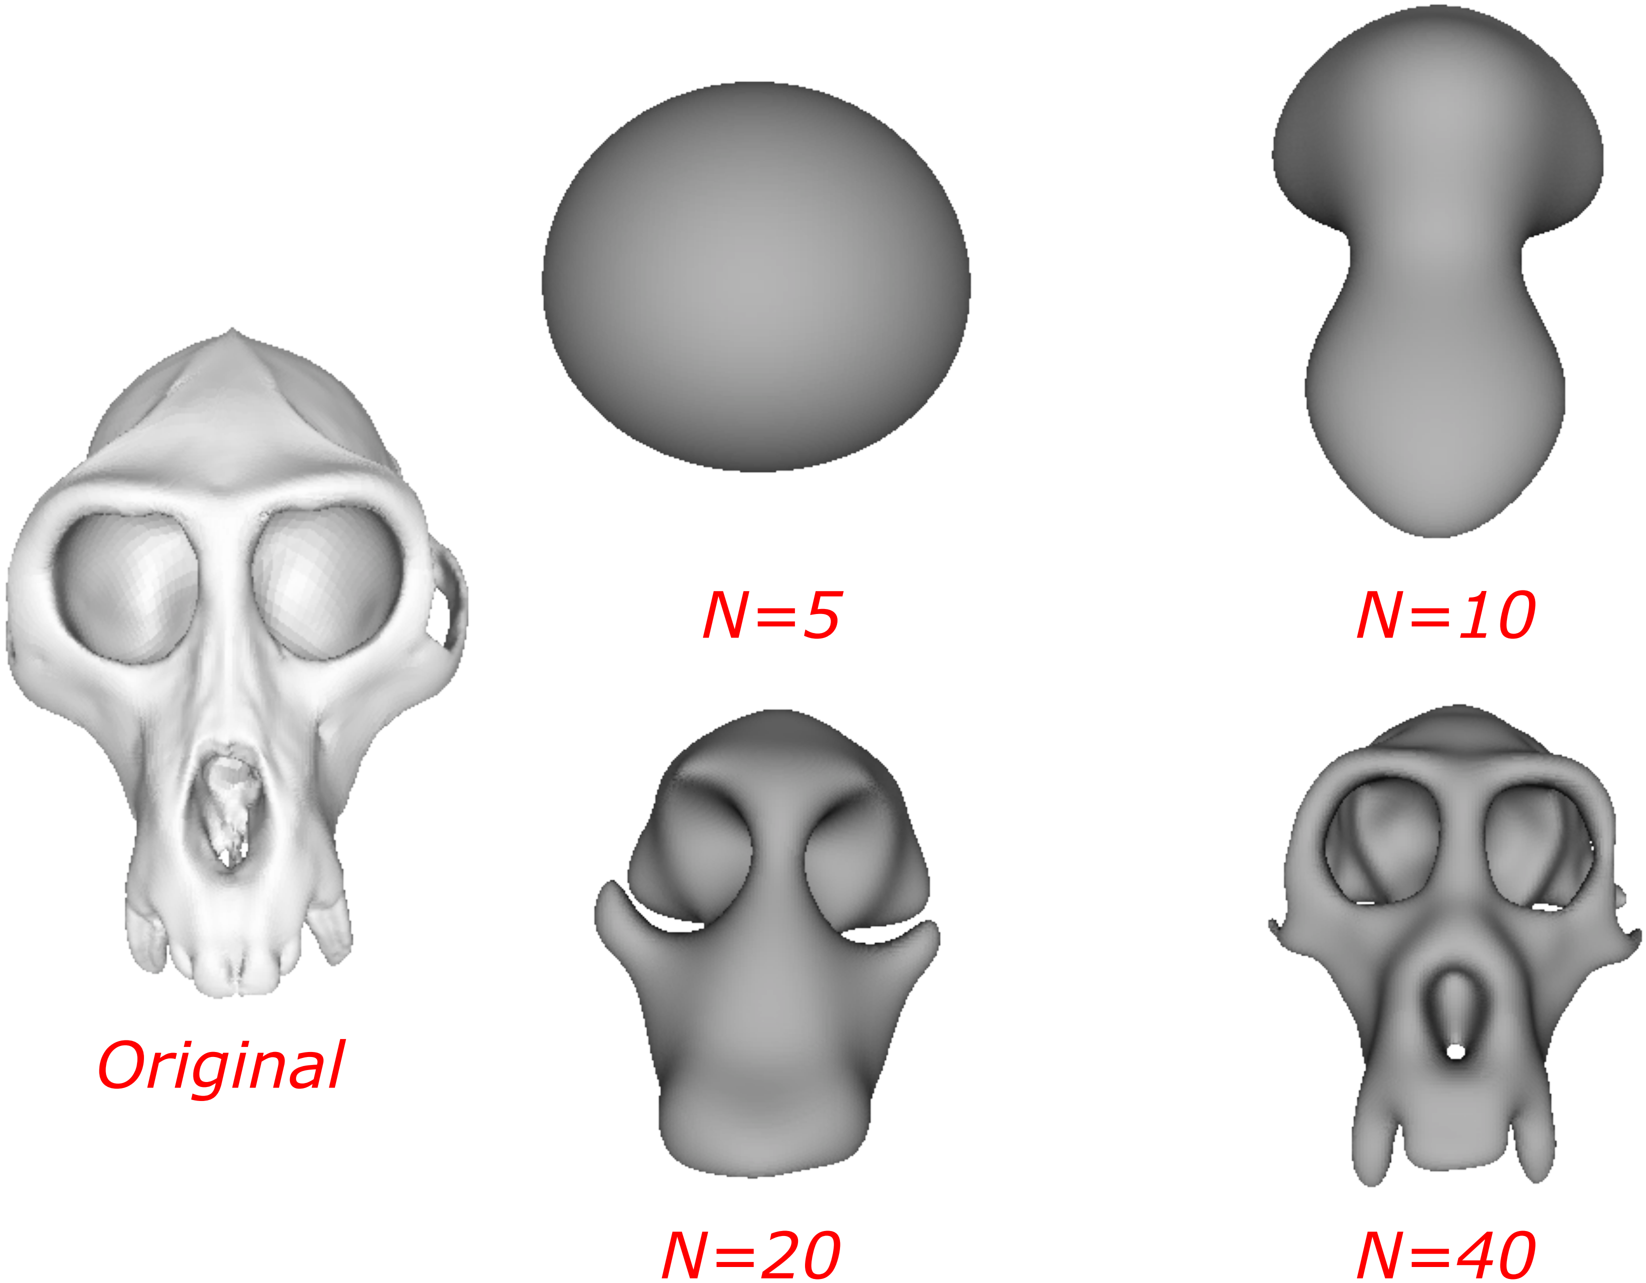

Original
N=5
N=10
N=20
N=40

Old World Monkey Skulls: Distance Tree

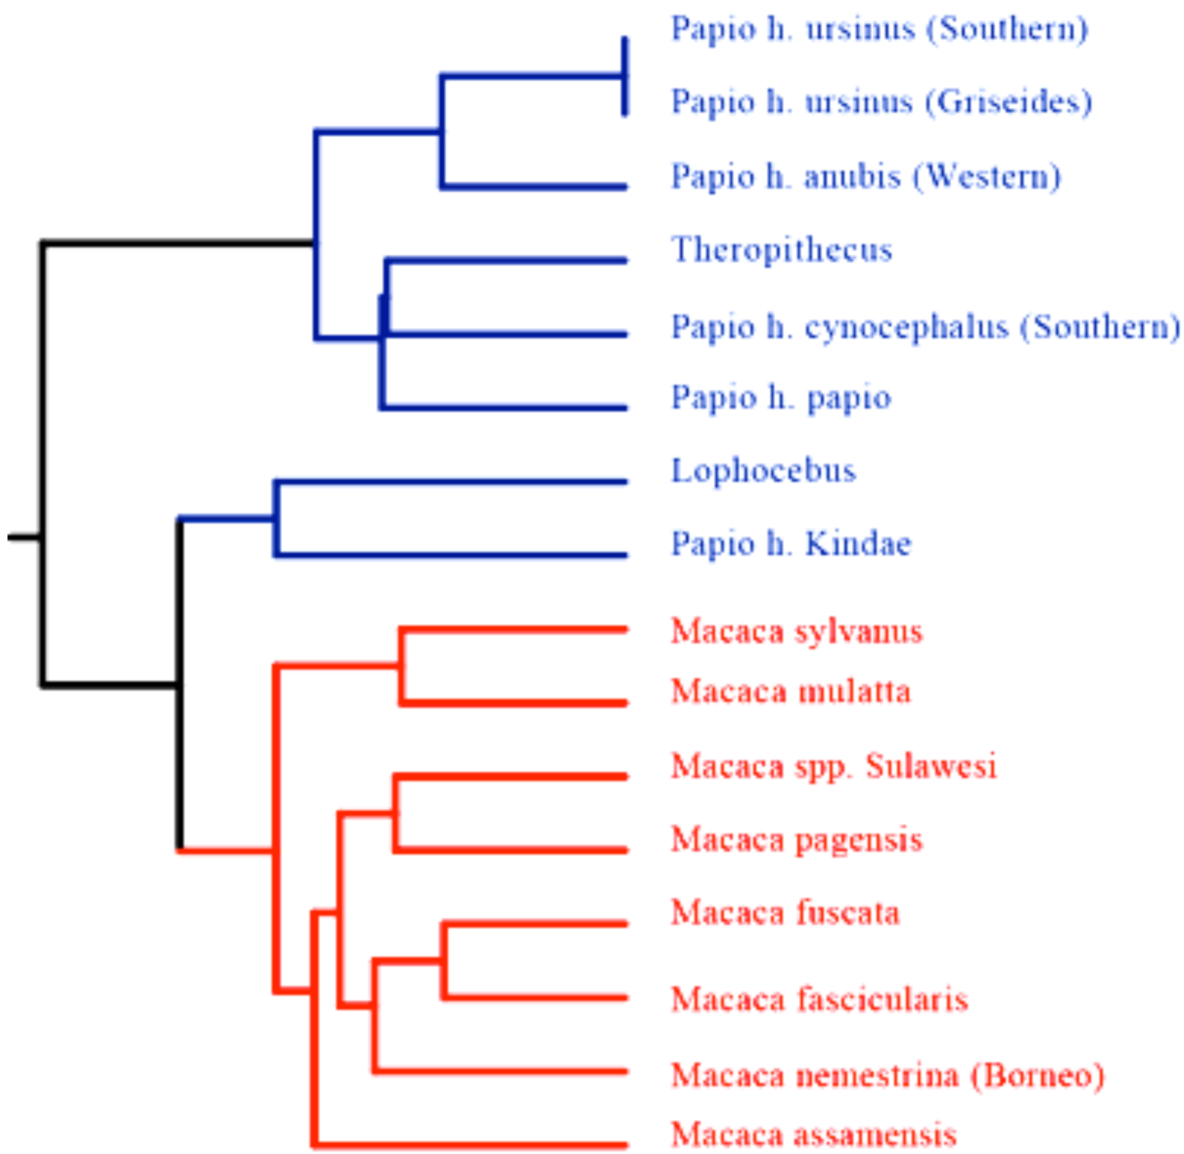

Analysis of the McGill Shape databases
458 objects, in 10 categories