

STATISTICAL SHAPE MODELS (SSM)

Medical Image Analysis

Serena Bonaretti

serena.bonaretti@istb.unibe.ch

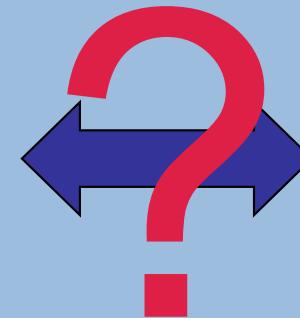
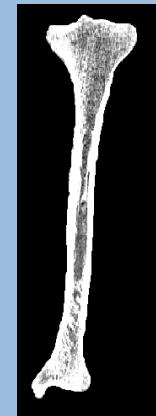
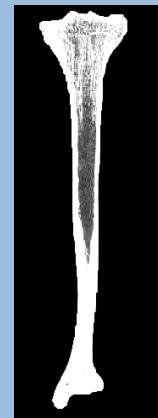
ISTB - Institute for Surgical Technology and Biomechanics
University of Bern

Overview

- > Introduction
 - Why are we interested in human shapes?
- > How do we describe shapes and their variability?
 - Shape description
 - From shape representation to shape space
 - Adding shapes to the shape space
 - Variability description
 - From variability description to space reduction
 - Principal Component Analysis
 - New instances creation

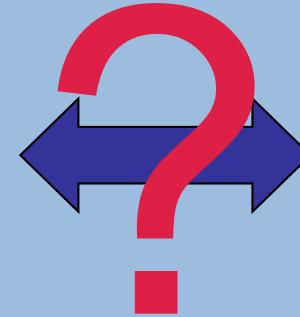
Why are we interested in human shapes?

- > Example 1: bone implant design



Why are we interested in human shapes?

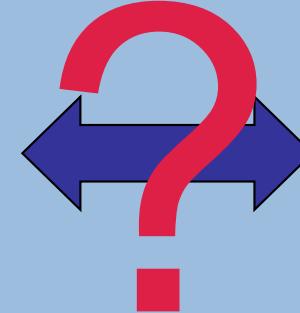
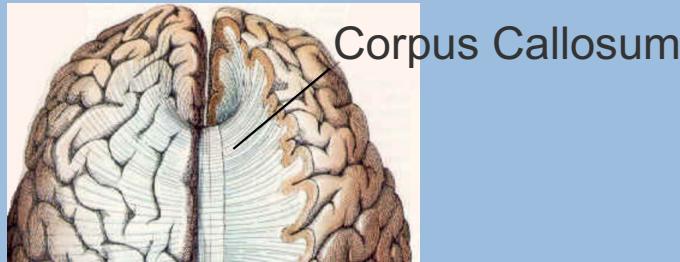
- > Example 2: hearing aid design



Additional problem:
Ear canals change shape
when people chew

Why are we interested in human shapes?

- > Example 3: anatomy-physiology correlation



Cognitive
abilities

Why are we interested in human shapes?

- > Large variability among human beings
 - Ethnicity
 - Age
 - Height
 - Weight
 - ...

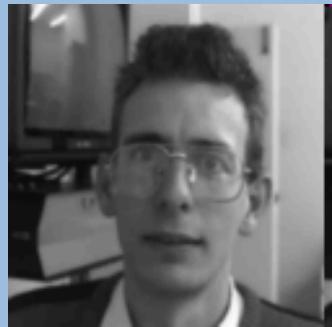


- > We want to find the “secrets” contained in shapes
- > We want to describe shapes and their variability



How do we describe shapes and their variability?

> University of Manchester, 1995



Timothy Cootes

COMPUTER VISION AND IMAGE UNDERSTANDING
Vol. 61, No. 1, January, pp. 38–59, 1995

Active Shape Models—Their Training and Application

T. F. COOTES, C. J. TAYLOR, D. H. COOPER, AND J. GRAHAM*

Department of Medical Biophysics, University of Manchester, Oxford Road, Manchester M13 9PT, England

Received July 29, 1992; accepted April 12, 1994



Chris Taylor

1. INTRODUCTION

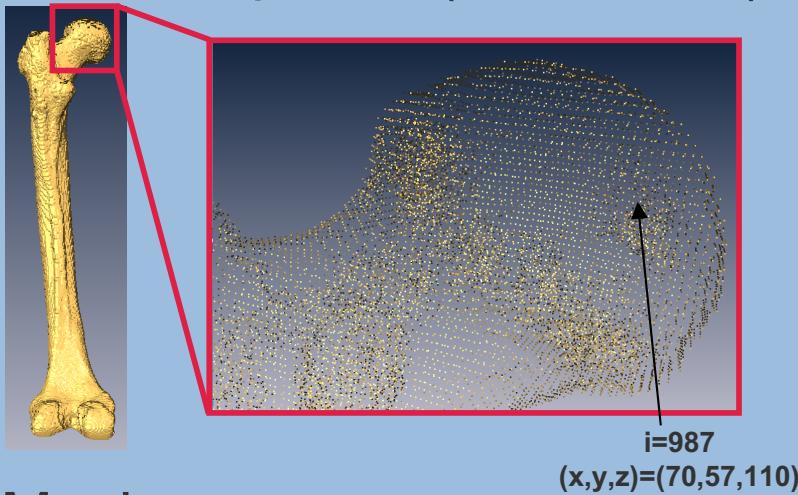
We address the problem of locating examples of known objects in images. Image interpretation using rigid models is well established [1, 2]. However, in many practical situations objects of the same class are not identical and rigid models are inappropriate. In medical applications, for instance, the shape of organs can vary considerably through time and between individuals. In addition, many industrial applications involve assemblies with moving parts, or components whose appearance can vary. In such cases flexible models, or deformable templates, can be used to allow for some degree of variability in the shape of the imaged objects [3–23].

In this paper we present new methods of building and using flexible models of image structures whose shape can vary. The models are able to capture the natural variability within a class of shapes and can be used in image search to find examples of the structures that they represent. Previous approaches have allowed models to deform, but have not tailored the variability to the class of shapes concerned—the models are not specific. Our main contribution is to describe how to create models which allow for considerable variability but are still specific to the class of structures they represent.

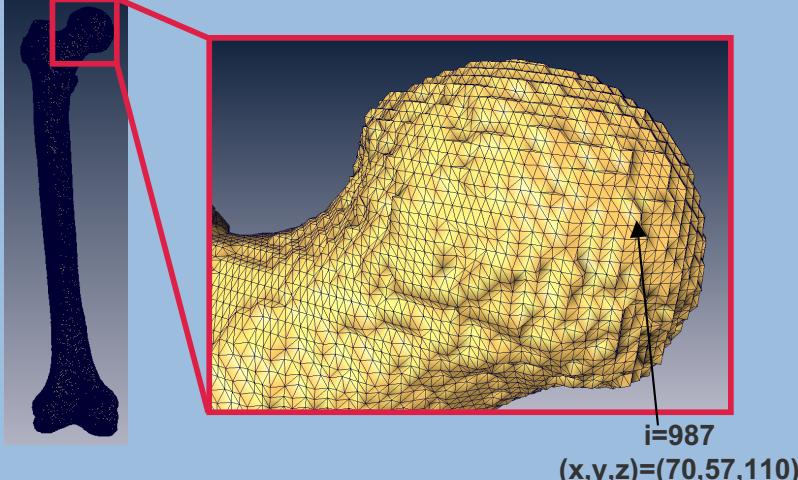
Shape representation



> Cloud of points (landmarks)



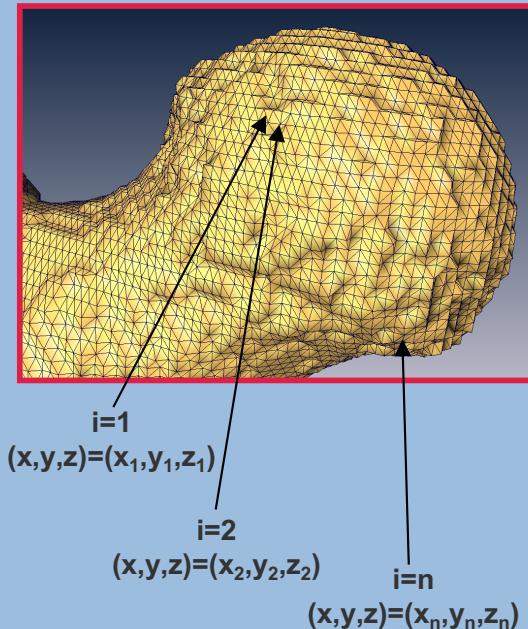
> Mesh



In any representation each point is characterized by:

- Id number
 $1 \leq i \leq n$, $n = \# \text{ points}$
- 3 coordinates
(x , y , z)

From shape representation to shape space



$$\underbrace{(x_1, y_1, z_1)}_{i=1}, \underbrace{(x_2, y_2, z_2)}_{i=2}, \dots, \underbrace{(x_n, y_n, z_n)}_{i=n}$$



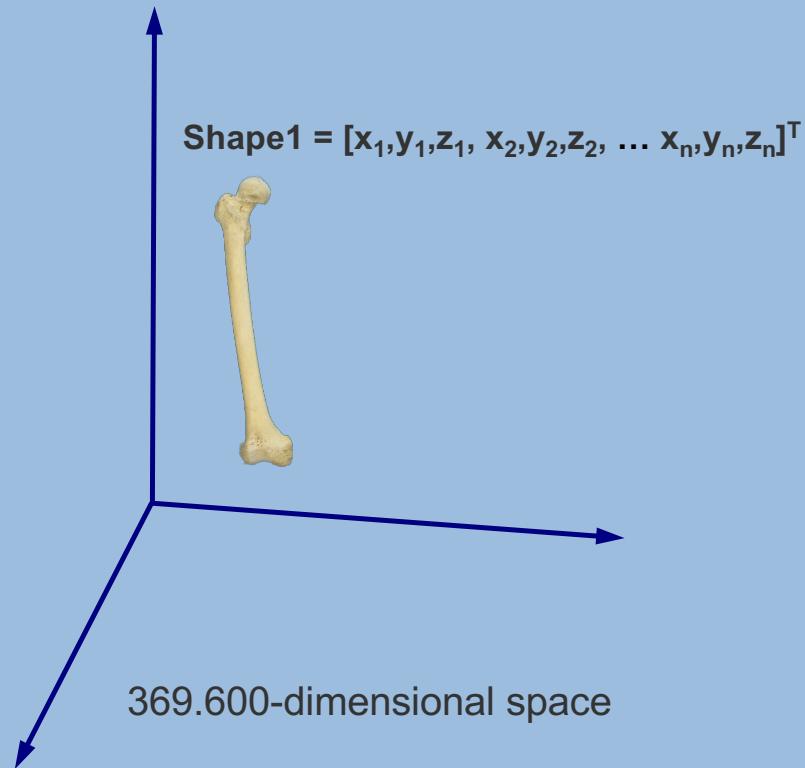
$$\text{shape1} = [x_1, y_1, z_1, x_2, y_2, z_2, \dots, x_n, y_n, z_n]^T$$



Key concepts:

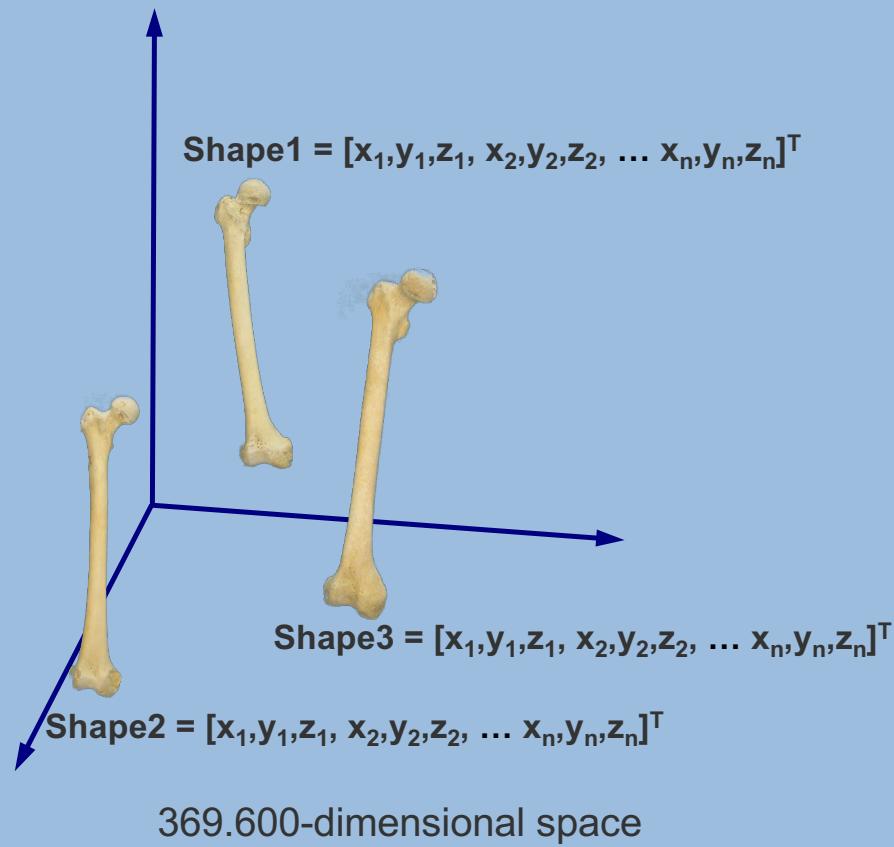
- The shape is represented as an array of coordinates (x, y, z)
- i.e. all the coordinates are put in one vector
- A vector can be represented as a point in its space

Shape space



- > The bone is described as a coordinate vector
- > This vector can be seen as a point in the space
- > But the space is not 2D, not 3D, not 4D, but... nD!!!
- > So if the shape is described by n points, the dimension of the space is 3xn
(in this case: n=123.200, so the space dimension is 369.600!)

Adding shapes to the space



- > We can add other shapes to our space in order to describe their variability (our initial aim)

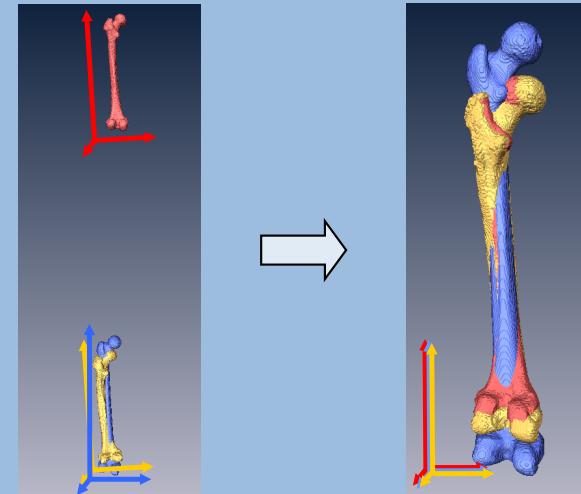


Shape characteristics:

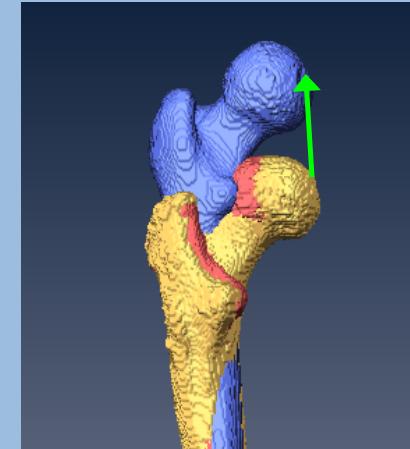
- same number of points
- points must be *aligned* in the same coordinate system
- points must be *correspondent* (same id for points in the same anatomical position)

Alignment and correspondence

- > How to make *alignment*?
 - Rigid registration

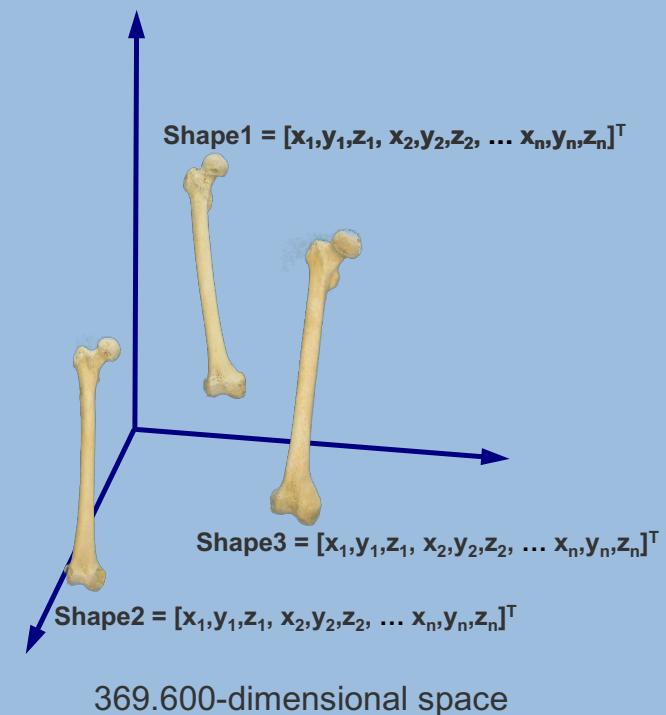


- > How to make *correspondence*?
 - Manually (landmarks selection)
 - Automatically (non-rigid registration)



Let's summarize a bit...

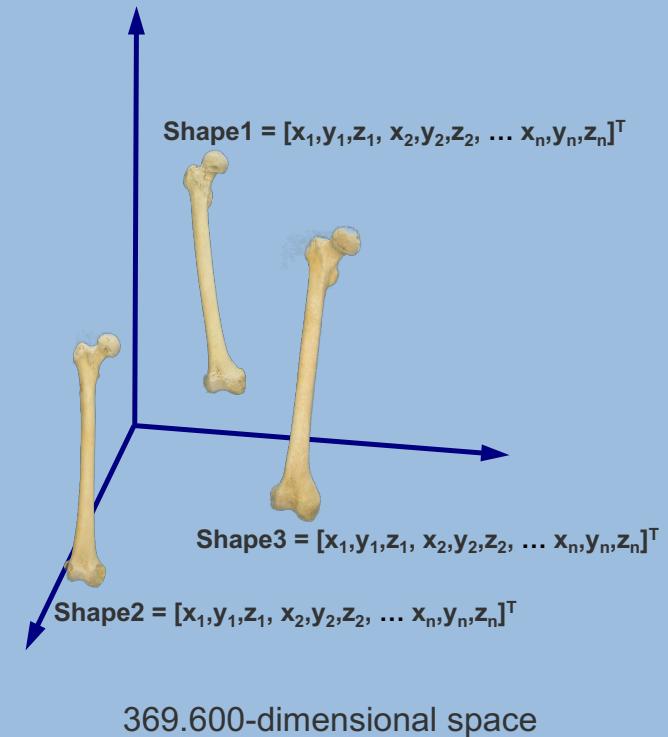
- > We are interested in describing shapes and their variability
- > Shape description
 - The shape is conceived as points
 - Every point is described by an id and 3 coordinates
 - All the coordinates can be put in one vector
 - This allows us to describe a shape as a point in a space (n-D space)
 - We can add other shapes (with defined characteristics) to the space
- > Variability description...



From shape description to variability description

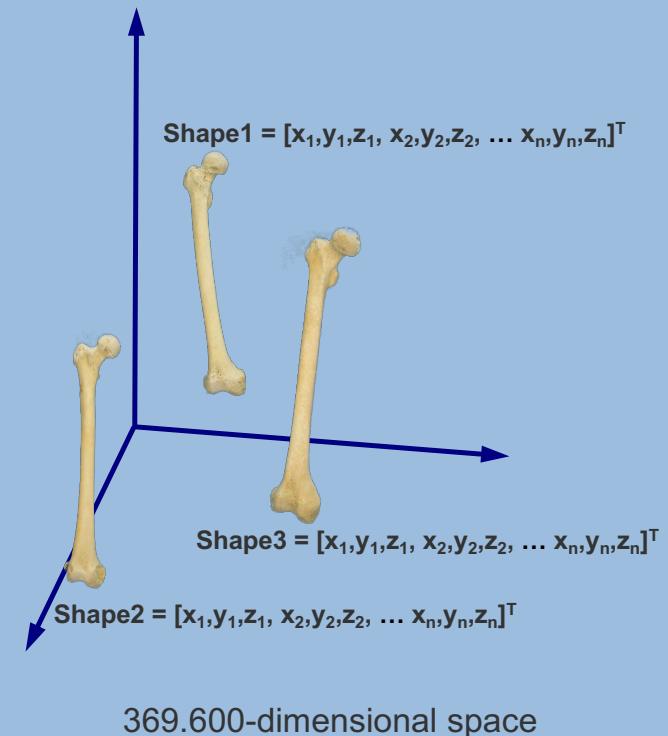
- > Shape description:
 - Shapes as n-D vectors
 - They are in a n-D space

- > We want to describe the variability of the shapes that are in this space
 - i.e. we want to find the patterns/common characteristics/“secrets” that are contained in the shapes



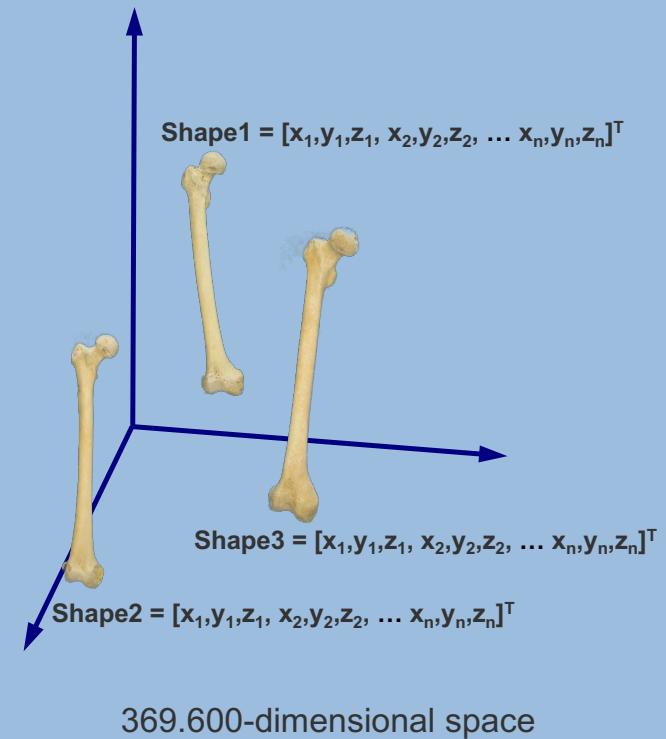
From variability description to space reduction

- > We want to describe the variability of the shapes that are in the n-D space
 - > But the n-D space is huge and difficult to manage
 - > It would be easier with a smaller space ($369.600D \rightarrow 7D / 5D / 3D, \dots$)
 - > So we would like to reduce our space to handle it in an easier way
- 
- How can we reduce a space?
Reducing its ***dimension***



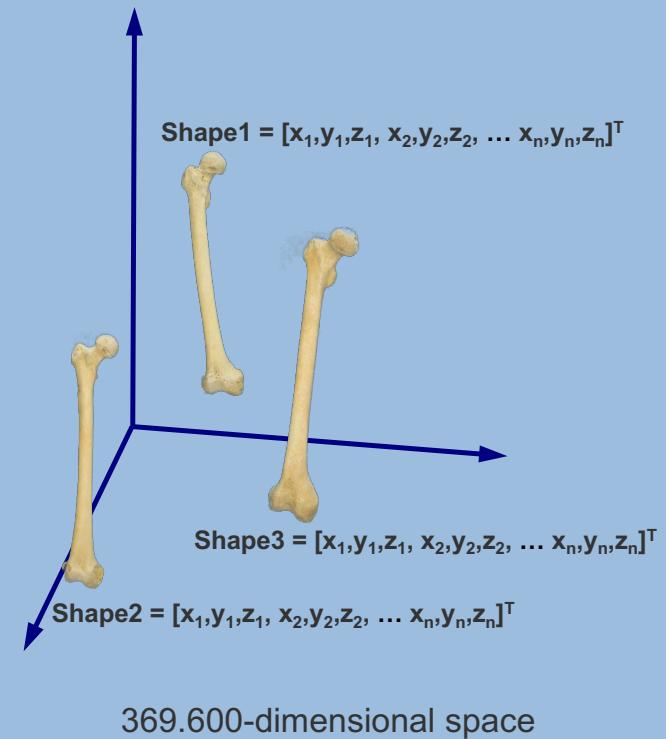
Space reduction

- > To describe the shape variability we reduce the shape-space dimension
 - > While reducing the space dimension we want to preserve data information
 - > The information that we are mostly interested in is the *variability* of the shapes
-  In a bit more formal way the variability can be described using the ***variance*** (The ***variance*** is a measure of the dispersion of the data around the mean value)



Space reduction technique

- > We want to reduce the shape space dimension preserving data information in terms of variability
- > The variability can be expressed as variance
- > We want to use the variance as a criterion/constraint for the shape space reduction
-  We want to reduce the space dimension *maximizing* the variance
- > The techniques we can use is called Principal Component Analysis (PCA)



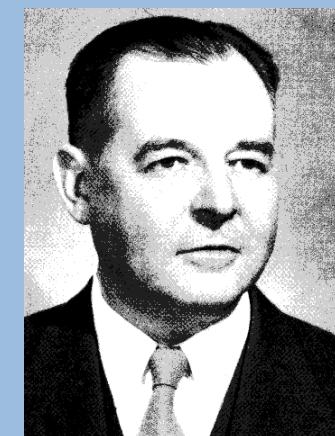
Principal Component Analysis (PCA)



Karl Pearson

On Lines and Planes of Closest
Fit to Systems of Points in Space.
1901, *Phil. Mag.*

Analysis of a Complex of Statistical
Variables with Principal Components
1933, *J of Educational Psychology*



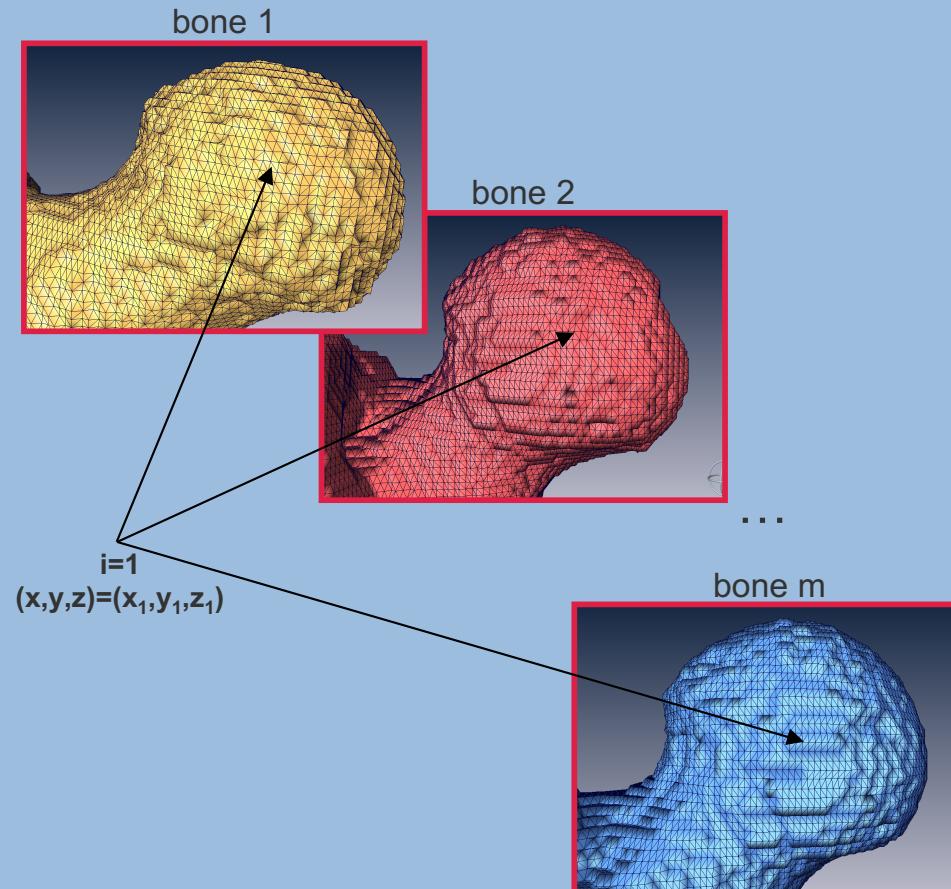
Harold Hotelling

Let's simplify for a while...

- > Every bone is described as a coordinate vector:

$$\text{Shape } i = [x_1, y_1, z_1, x_2, y_2, z_2, \dots x_n, y_n, z_n]^T$$

- > Let's consider the same point ($i=1$) for all bones:



	bone 1	bone 2	...	bone m		
x	12	20	16	10	12	15
y	11	15	8	10	7	9
z	6	14	10	12	8	15

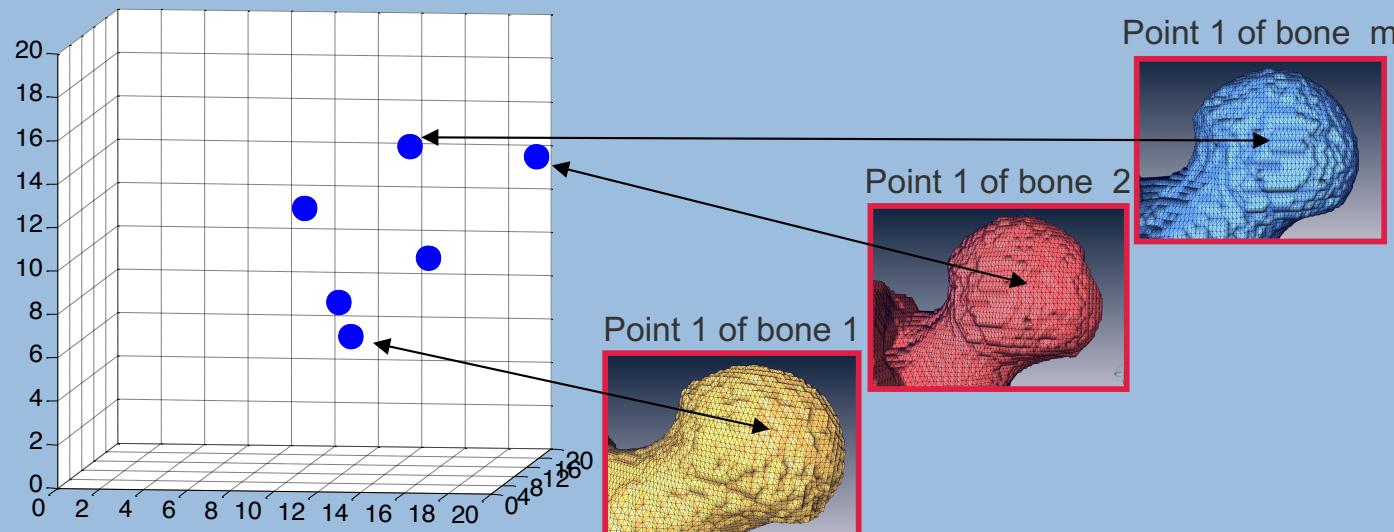
- > We have just created a matrix
- > It tells us that:
 - we are in a 3D space (x, y, z of 1 point)
 - we are using 6 observations ($m=6$)

Let's simplify for a while...

- > We have put the coordinates of the point $i=1$ of some bones in a matrix:

	bone 1	bone 2	...	bone m		
x	12	20	16	10	12	15
y	11	15	8	10	7	9
z	6	14	10	12	8	15

- > We can plot them in a 3D coordinate system:

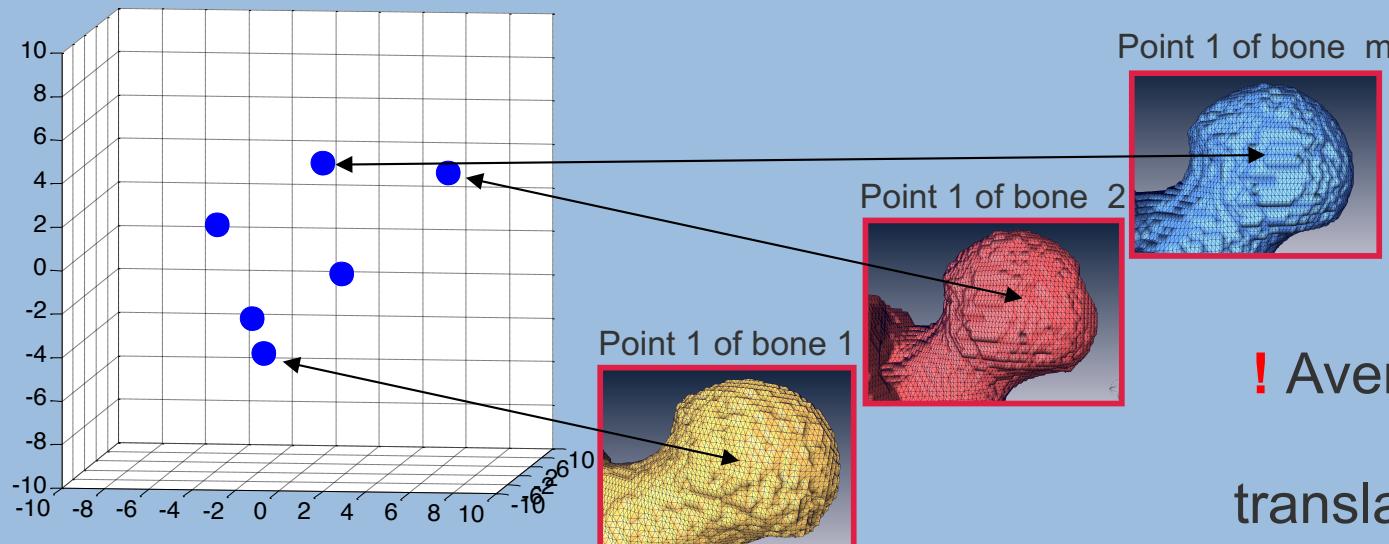


Let's simplify for a while...

- > We can also subtract the average to the coordinates...

$$\begin{array}{ccccccc} & \text{bone 1} & \text{bone 2} & \dots & & \text{bone m} & \\ \text{x} & 12 & 20 & 16 & 10 & 12 & 15 \\ \text{y} & 11 & 15 & 8 & 10 & 7 & 9 \\ \text{z} & 6 & 14 & 10 & 12 & 8 & 15 \end{array} - \begin{array}{c} \text{mean} \\ \hline 14.17 \\ 10.00 \\ 10.83 \end{array} = \begin{array}{ccccccc} & \text{bone 1} & \text{bone 2} & \dots & & \text{bone m} & \\ \text{-2.17} & 5.83 & 1.83 & -4.17 & -2.17 & 0.83 \\ 1.00 & 5.00 & -2.00 & 0.00 & -3.00 & -1.00 \\ -4.83 & 3.17 & -0.83 & 1.67 & -2.83 & 4.17 \end{array}$$

- > ...and represent them again

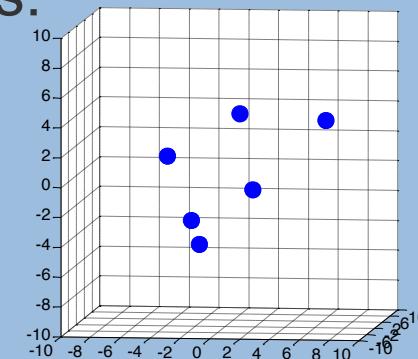


! Average subtraction
=
translation to the origin

Let's simplify for a while...

- > We have mean-subtracted coordinates:

	bone 1	bone 2	...	bone m		
x	-2.17	5.83	1.83	-4.17	-2.17	0.83
y	1.00	5.00	-2.00	0.00	-3.00	-1.00
z	-4.83	3.17	-0.83	1.67	-2.83	4.17



- > We want to describe their variability
- > So we can compute their **covariance matrix**
(The covariance matrix describes the degree of linear dependence of two variables)

$$\text{cov} = \frac{1}{m-1} X X^T$$

(X = mean-subtracted data matrix)

variance	12.97	5.80	6.43
	5.80	8.00	3.40
covariance (x,y)	6.43	3.40	12.17

Let's simplify for a while...

- > And finally we can compute a matrix decomposition

12.97	5.80	6.43
5.80	8.00	3.40
6.43	3.40	12.17

- > We can decompose our covariance matrix into other matrices

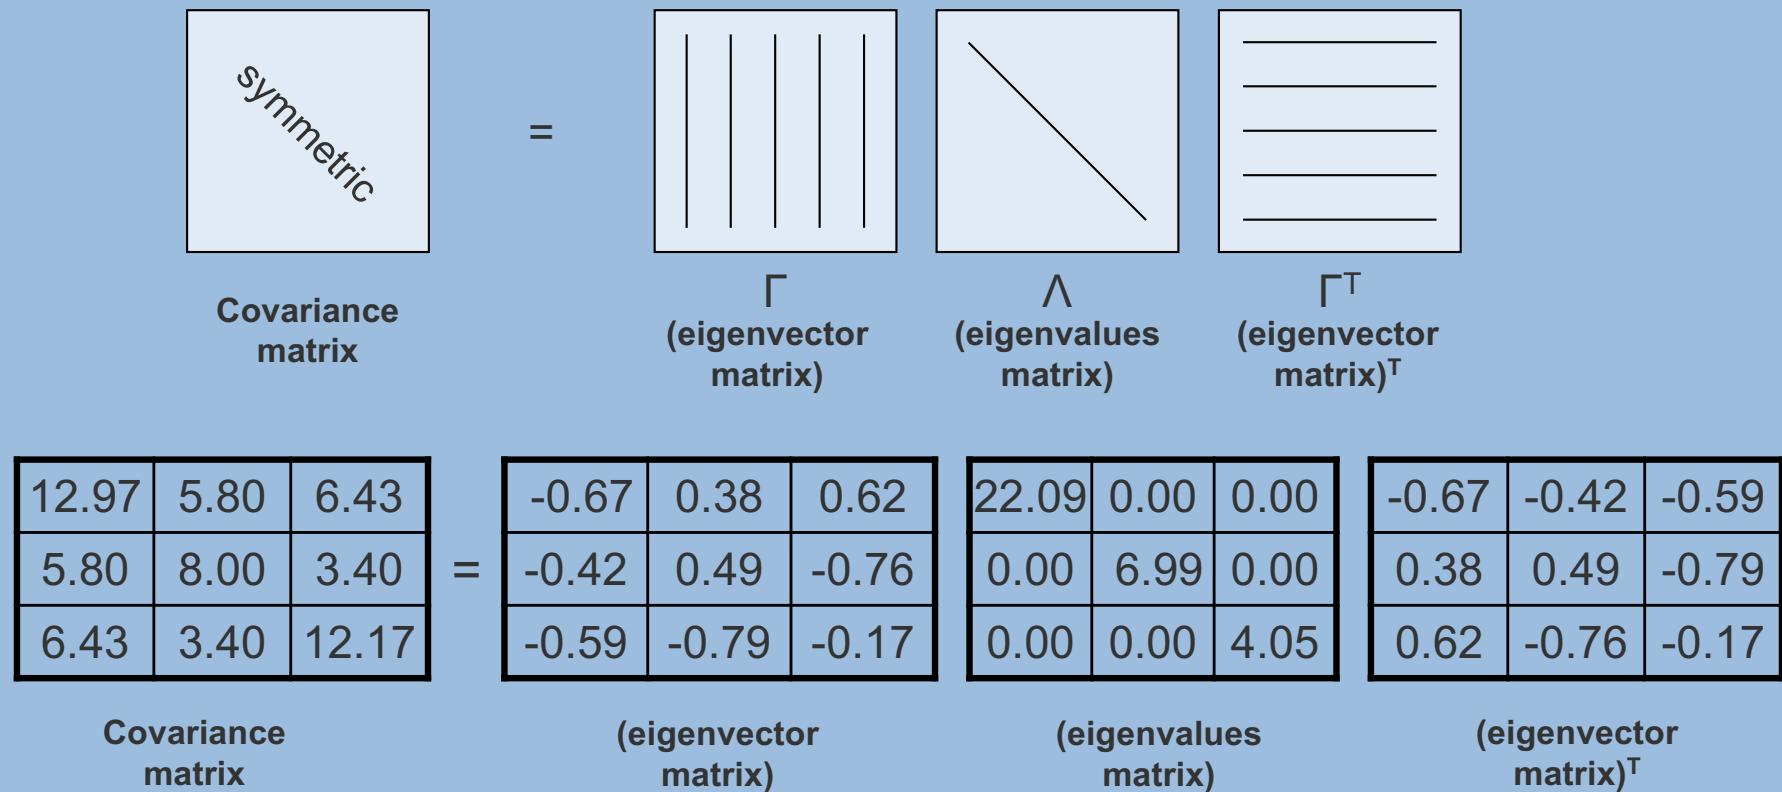
$$\begin{matrix} \text{symmetric} \\ \text{Covariance matrix} \end{matrix} = \begin{matrix} \Gamma \\ (\text{eigenvector matrix}) \end{matrix} \begin{matrix} \Lambda \\ (\text{eigenvalues matrix}) \end{matrix} \begin{matrix} \Gamma^T \\ (\text{eigenvector matrix})^T \end{matrix}$$

Spectral decomposition (or Jordan decomposition):

- Given a *symmetric* matrix A , we can decompose it as the product of three matrices: $A = \Gamma \Lambda \Gamma^T$

Let's simplify for a while...

- So we can decompose our covariance matrix:



- > And we can look at the meaning of the eigenvector and eigenvalue matrices

Let's simplify for a while...

> Meaning of the **eigenvectors** (1/2)

- They are three columns of numbers
- Each column is called *eigenvector* or *mode* or *principal component*
- Each mode has the same dimension of the shape space



Each mode can be considered as an axis of a new coordinate system

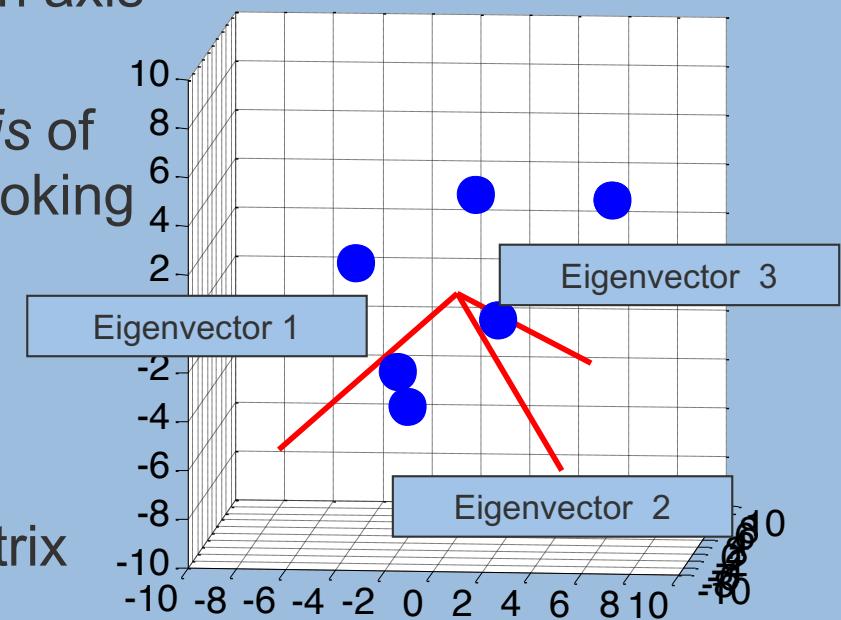


They can be considered as the *basis* of the new shape space that we are looking for (they are linearly independent)

- So we can move/rotate/project all our points into it
- To do that we can consider the eigenvector matrix as a rotating matrix

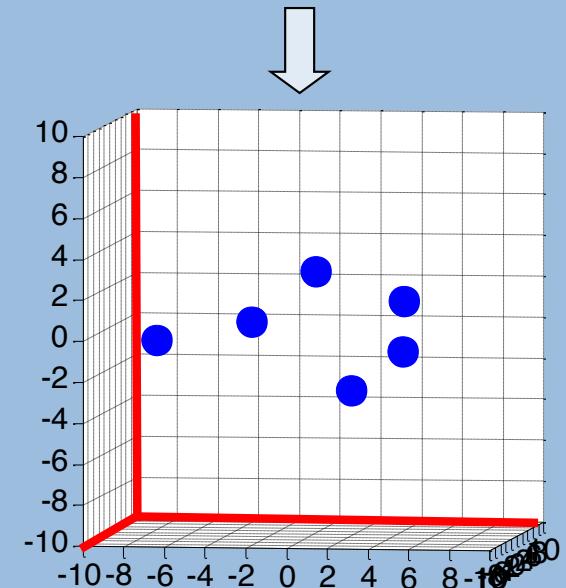
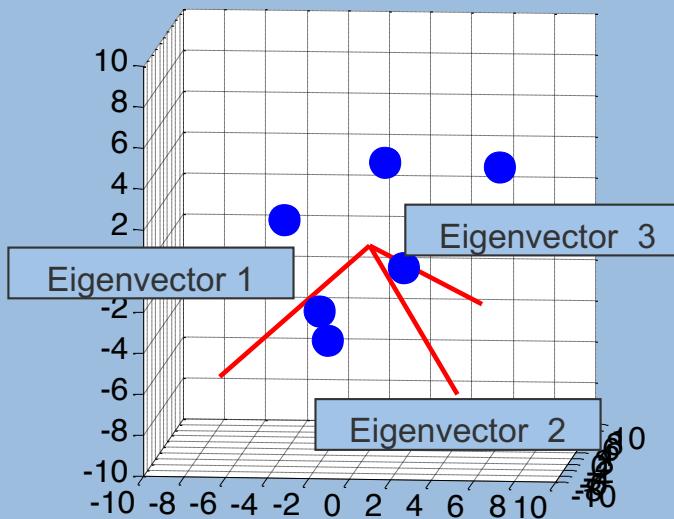
-0.67	0.38	0.62
-0.42	0.49	-0.76
-0.59	-0.79	-0.17

(eigenvector matrix)



Let's simplify for a while...

- > Meaning of the **eigenvectors** (2/2)
 - We can rotate all the points to the new coordinate system
 - It means that we can project all the points to the new coordinate system
-  But if we choose all of them we do not reduce the space dimension
- Therefore we will use just some of them
 - How do we choose just some of them? → Eigenvalues



Let's simplify for a while...

- > Meaning of the **eigenvalues** (1/2)
 - The number of eigenvalues is the same as the number of eigenvectors
 - They are ordered in a descending way
 - Each eigenvalue corresponds to the eigenvector that has the same position
-  Each eigenvalue represents the amount of variation that we can describe using the correspondent eigenvector
-  So we can describe a percentage of the variations using just some modes (the ones that have a correspondent meaningful eigenvalue)

22.09	0.00	0.00
0.00	6.99	0.00
0.00	0.00	4.05

(eigenvalues matrix)

Eigenvalues:
22.0920
6.9913
4.0501

-0.67	0.38	0.62
-0.42	0.49	-0.76
-0.59	-0.79	-0.17

(eigenvector matrix)

Let's simplify for a while...

> Meaning of the **eigenvalues** (2/2)

- Each eigenvalue represents the amount of variation
- Therefore we can calculate this variation in terms of percentage:
 $\text{variation\%} = \text{eigenvalues}/\text{sum(eigenvalues)} \times 100$
- And we can calculate the accumulated variation (how much variation the modes describe when they are taken into account together)
- The percentage of variation that we describe depends on the amount of eigenvectors that we take into account

22.09	0.00	0.00
0.00	6.99	0.00
0.00	0.00	4.05

(eigenvalues matrix)

Eigenvalues:
22.0920
6.9913
4.0501

Percentage of variation:	Accumulated variation:
66.68%	66.68%
21.10%	87.78%
12.22%	100.00%

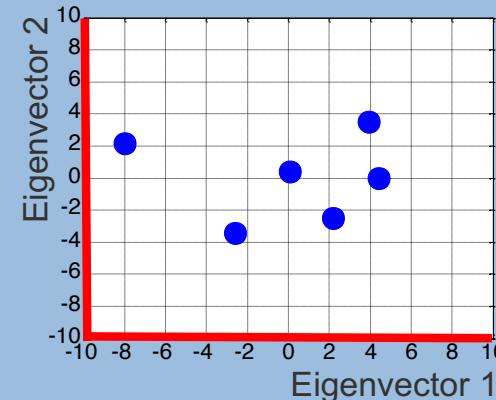
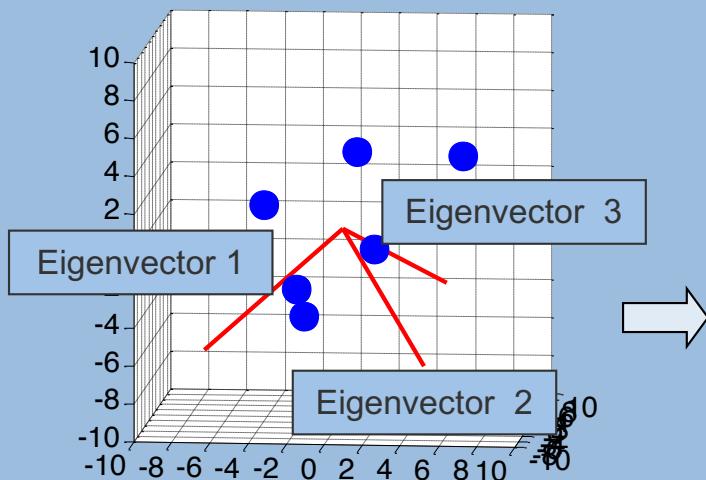
Let's simplify for a while...

- > Let's consider our example:

Eigenvalues:	Percentage of variation:	Accumulated variation:
22.0920	66.68%	66.68%
6.9913	21.10%	87.78%
4.0501	12.22%	100.00%

Eigenvectors:		
-0.67	0.38	0.62
-0.42	0.49	-0.76
-0.59	-0.79	-0.17

- > If we choose to project our data on 1st and 2nd mode of variation, we can describe the 87.78% of the whole variability of our dataset

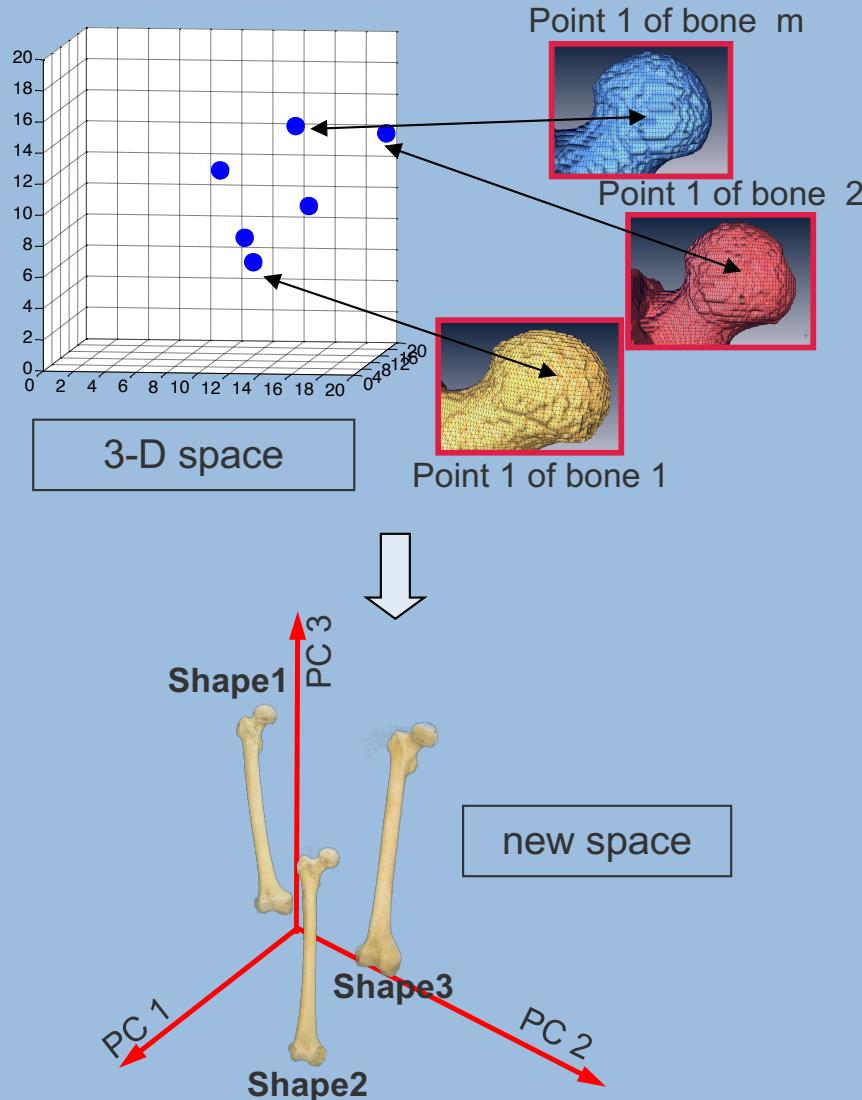


Projecting our data on meaningful modes is our way to reduce the space

Let's summarize a bit...

- > We want to describe shape variability
- > We can describe the variability reducing the space dimension
- > We can reduce the space dimension using the PCA
- > The PCA is a multivariate statistical technique whose steps are:
 - Data alignment
 - Data mean subtraction
 - Covariance matrix calculation
 - Covariance matrix decomposition → eigenvectors and eigenvalues
 - The eigenvectors are the new basis of the space
 - The eigenvalues suggest us how many eigenvectors to take in order to reduce the space keeping a certain amount of variation among our shapes

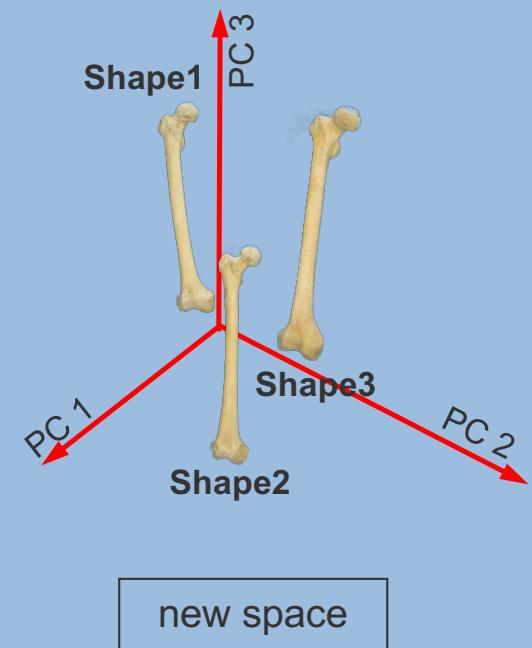
Extension to a bigger space



- > The whole procedure is exactly the same as in the 3D case
- > Number of data can be much higher → memory issue
- > The dimension reduction makes the model very compact and this allows us to manage our data in an easier way

What can we do with shapes in a reduced space?

- > We can create new instances to fill in the whole space
- > We can fill in along the direction of a chosen mode
- > We can fill in combining the modes together
- > In any case we know the percentage of the space that we are going to represent thanks to the eigenvalues



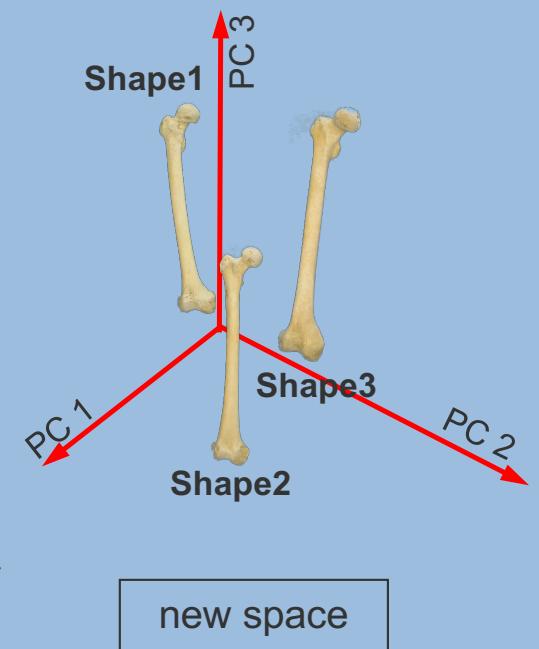
What can we do with shapes in a reduced space?

- > The formula that we can use to create new instances is:

$$\bar{x} = \bar{x} + \phi b$$

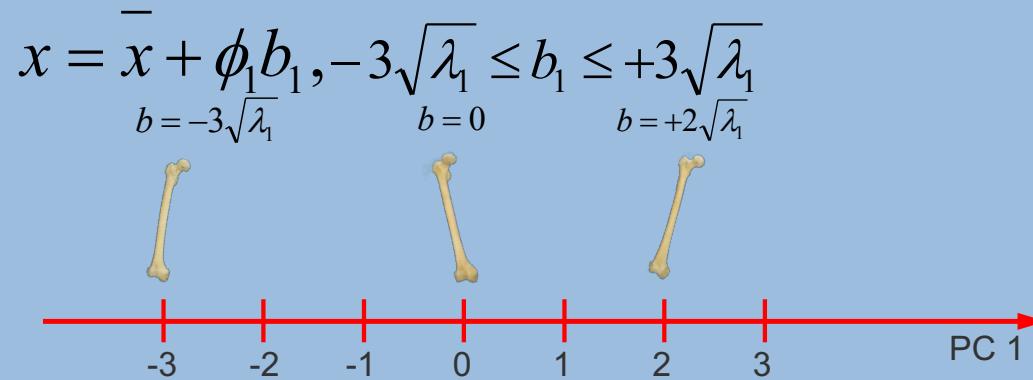
where:

- \bar{x} is the new shape vector
- \bar{x} is the average shape
- Φ is the chosen eigenvector
- b is a parameter whose value is: $-3\sqrt{\lambda_i} \leq b \leq +3\sqrt{\lambda_i}$
(λ is the corresponding eigenvalue)

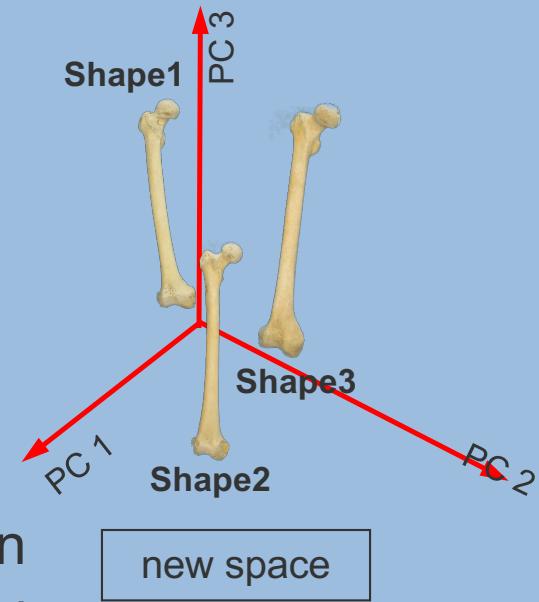


What can we do with shapes in a reduced space?

- > Creation of new instances using the 1st mode:



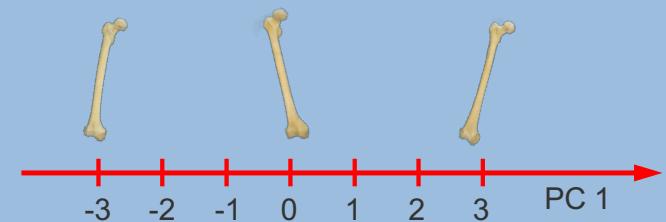
- > Comment on the value of b :
 - λ is the variance, so $\sqrt{\lambda}$ is the standard deviation
 - in a Gaussian distribution, $-3\text{s.d.} \leq \text{value} \leq +3\text{s.d.}$ represents the 99.7% of the distribution



What can we do with shapes in a reduced space?

- > We can create new instances using any mode (mode 1, 2, 3, ...)
- > We can create new instances using a linear combination of modes. For example:

$$x = \bar{x} + \phi_1 b_1 + \phi_2 b_2 + \dots, \quad -3\sqrt{\lambda_1} \leq b_1 \leq +3\sqrt{\lambda_1},$$
$$-3\sqrt{\lambda_2} \leq b_2 \leq +3\sqrt{\lambda_2}$$



- > Any mode usually represents a variation of some bone characteristics:
 - Mode 1: variation of length
 - Mode 2: variation of the antitorsion angle
 - ...
- > To understand this better, have a look at the exercise...