Image Analysis

Rasmus R. Paulsen
Tim B. Dyrby
DTU Compute

rapa@dtu.dk

http://courses.compute.dtu.dk/02502

Lecture 12 – Statistical models of shape and appearance

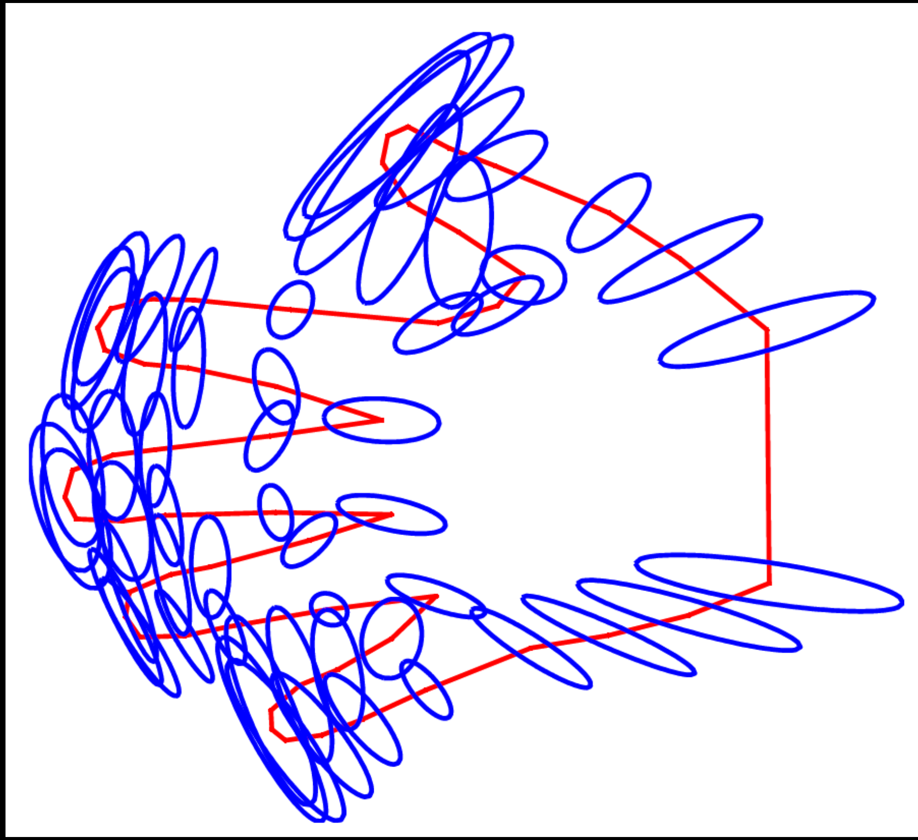

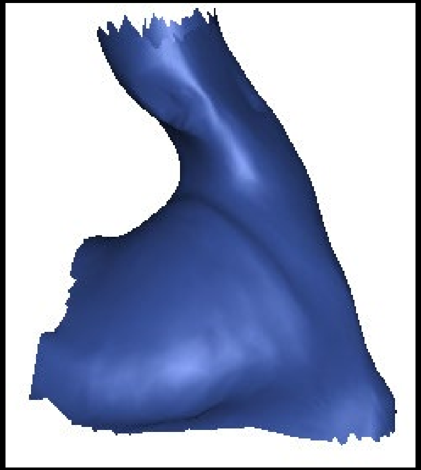

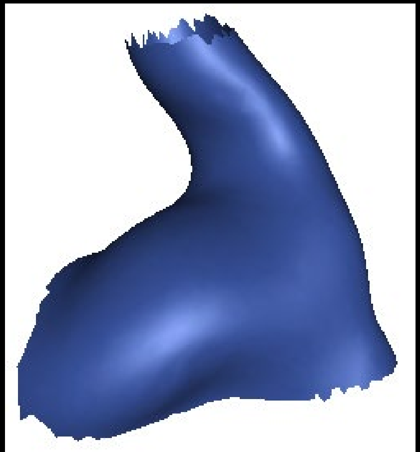

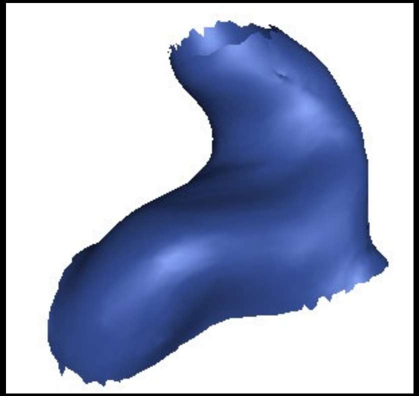

Today's Learning Objectives

- Describe the concept of shape models
- Define the shape of an object using landmarks
- Describe point correspondence
- Describe and use the vector representation of a shape
- Describe how a shape can be seen as a point in high-dimensional space
- Explain how shapes can be aligned
- Describe how principal component analysis can be used to model shape variation
- Explain the similarity between Eigenfaces and shape and appearance models

A typical scenario

- Doctor X believes that he can “see” on a hand X-ray if the patient is in risk of arthritis!
- Specifically Doctor X is sure that the *shape of the joints* is an estimator for arthritis!

Can we verify that?

Scenario II

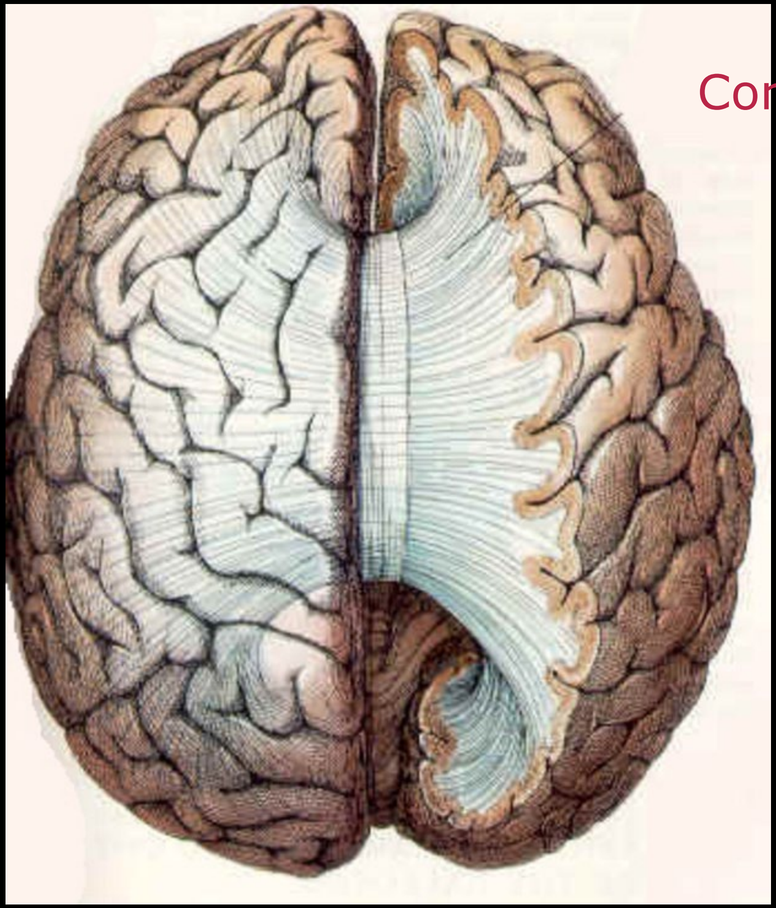

Corpus Callosum

- MR images have been captured of a large group of people
- Cognitive abilities measured as well
- Is there a correlation between *how the brain looks* and how we behave?
- Does the shape of corpus callosum tell us something?

Scenario II

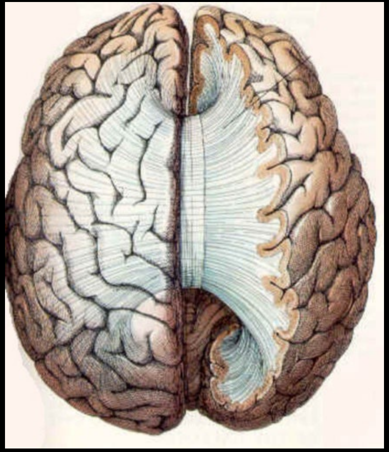

Corpus Callosum

- We can get the MR slice with the corpus callosum from all the patients

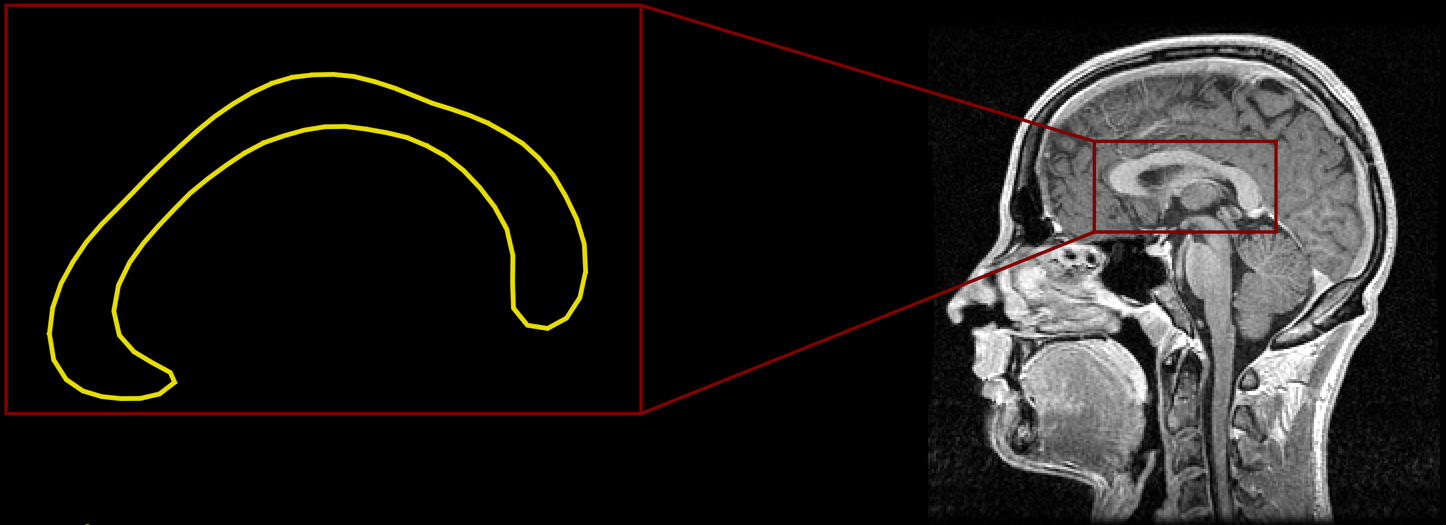

Scenario III


- An experienced hearing aid fitter has seen a lot of ears!
- Some hearing aid users are very difficult to fit. Why?
- Large variation in the shape of ears
- Ear canals change shape when people chews
- Is it possible to learn about the shape and use it?






Shape Analysis

- What can we learn from shape?
- What can we use it for?
- How do we do it?

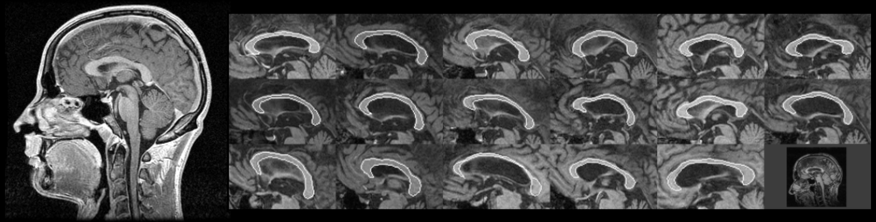

600 MR scans and behavioural data

1000 historical X-rays

A boxful of something that look like ear canals

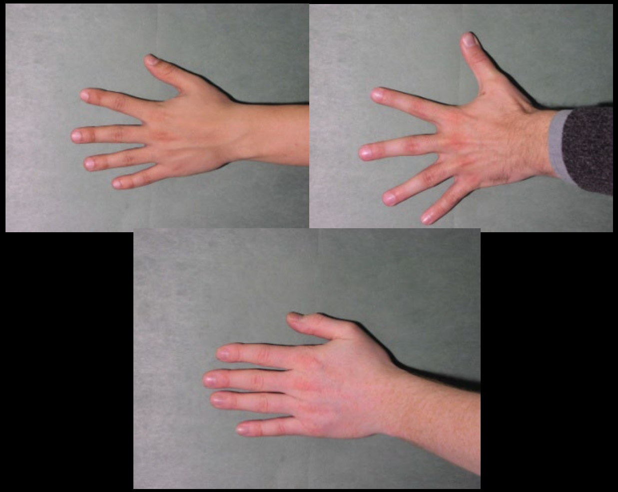

A set of hand photographs

What is shape?

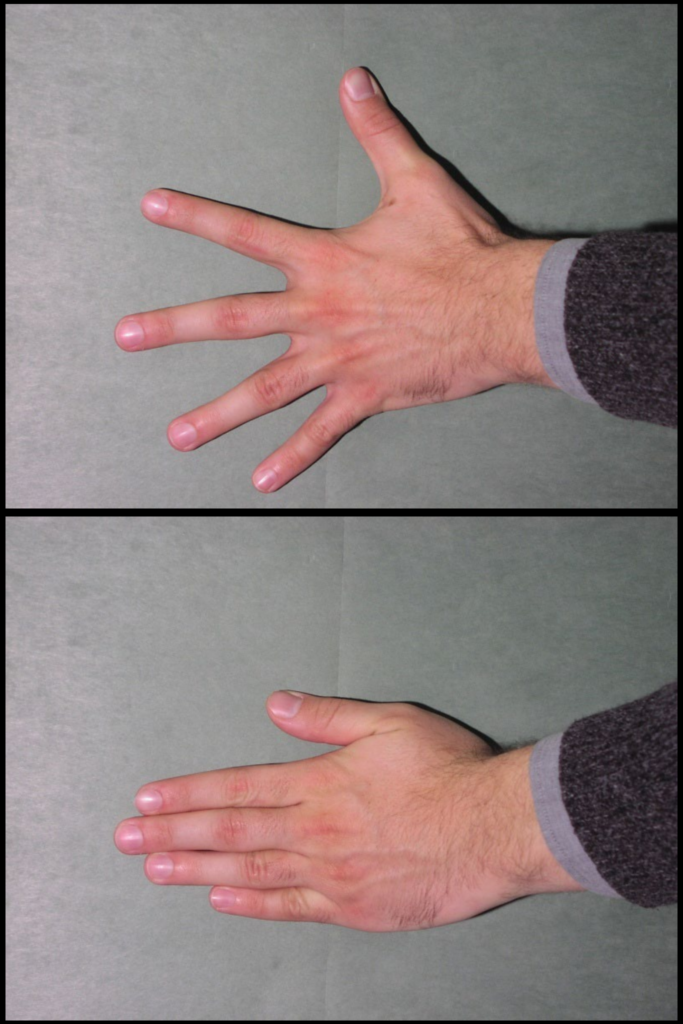

- How do we define the *shape* of this hand?
- What is the shape difference between the two hands?

Shape definition

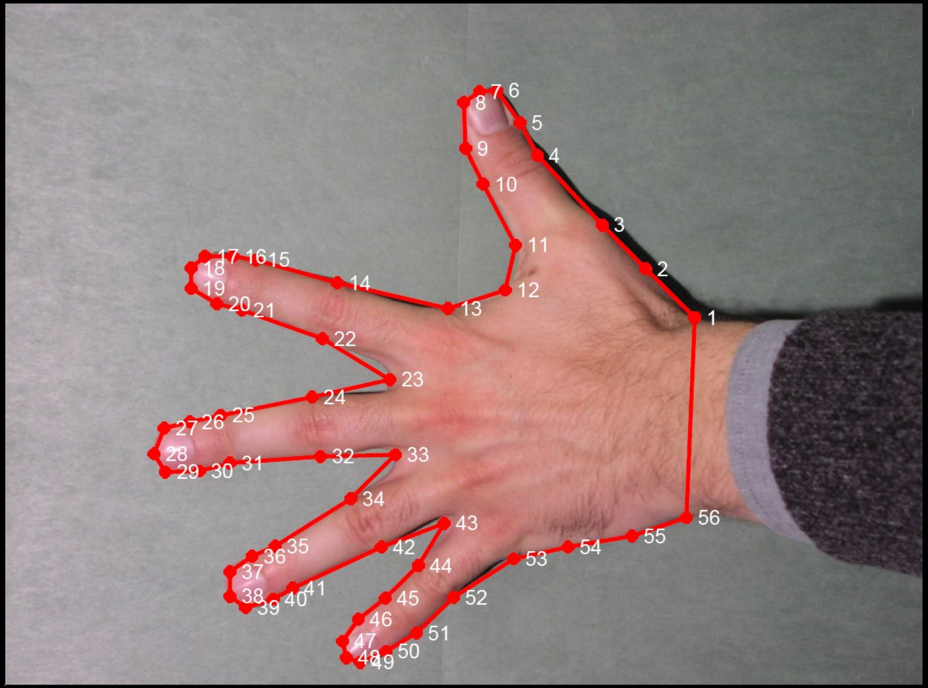

- Shape is defined using landmarks
  - Placed by an expert
- In this case the outer contour of the hand
- Just one of many ways!

Shape definition

Shape is all geometrical information that remains when location, scale, and rotational effects are removed

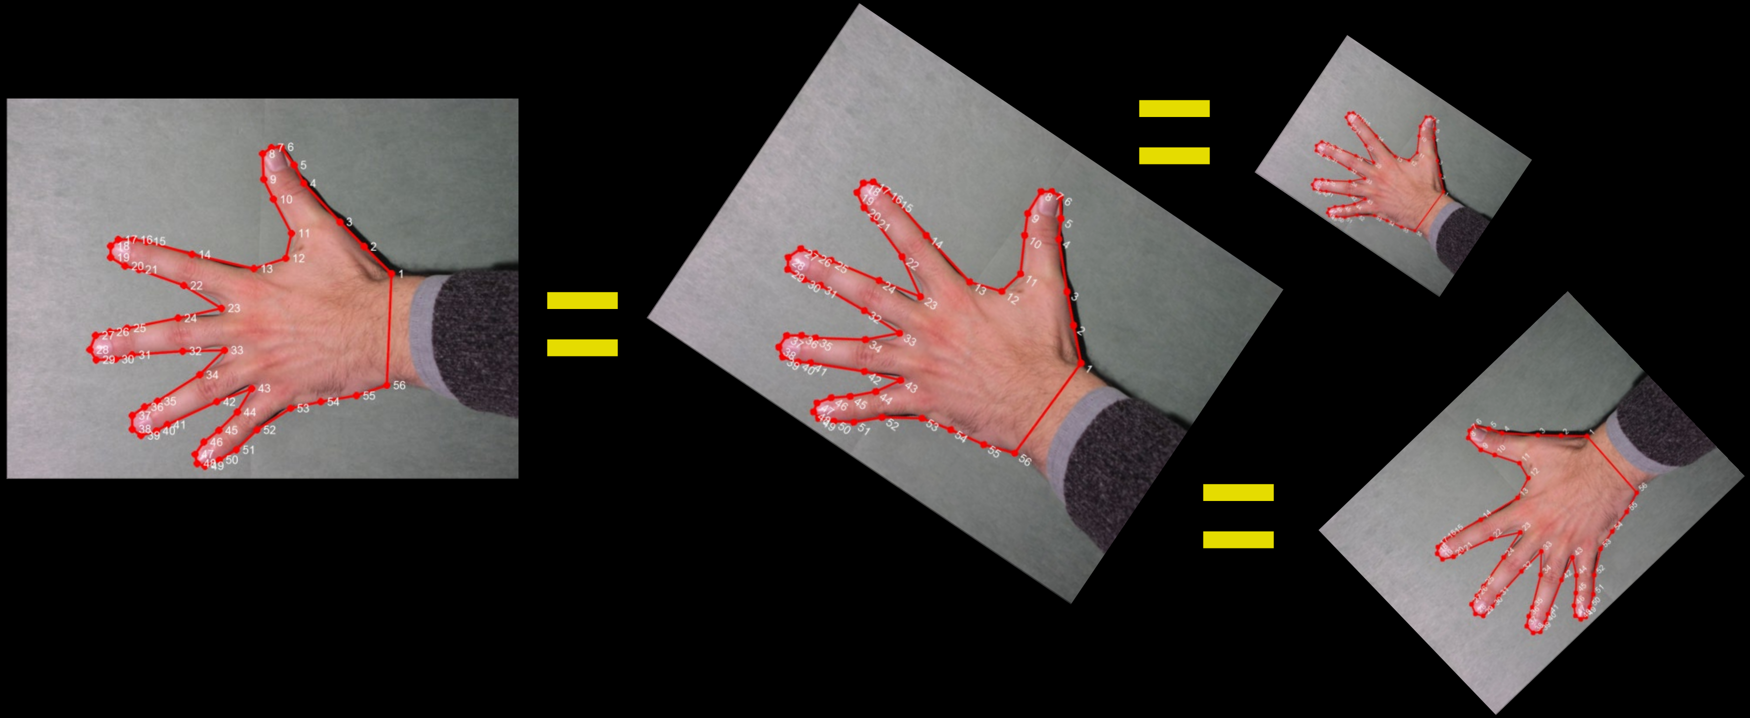

Shape definition

Shape is all geometrical information that remains when location, scale, and rotational effects are removed

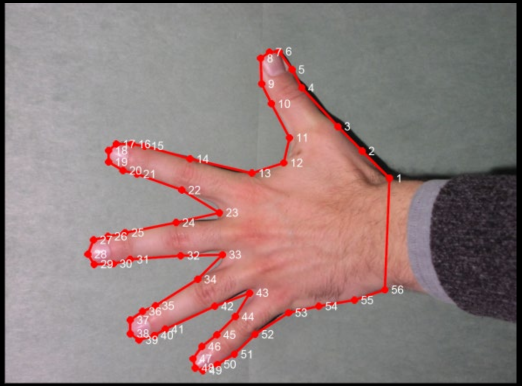

≠

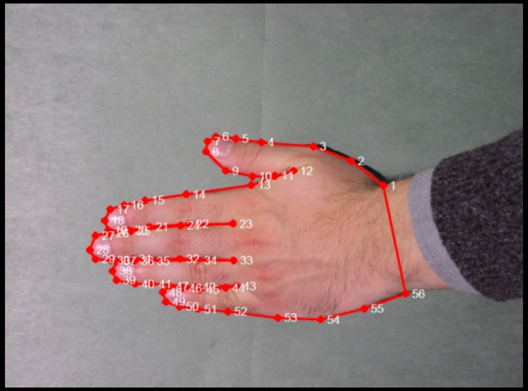

Landmarks and point correspondence

Landmarks are placed on the same place on all shapes in the training set

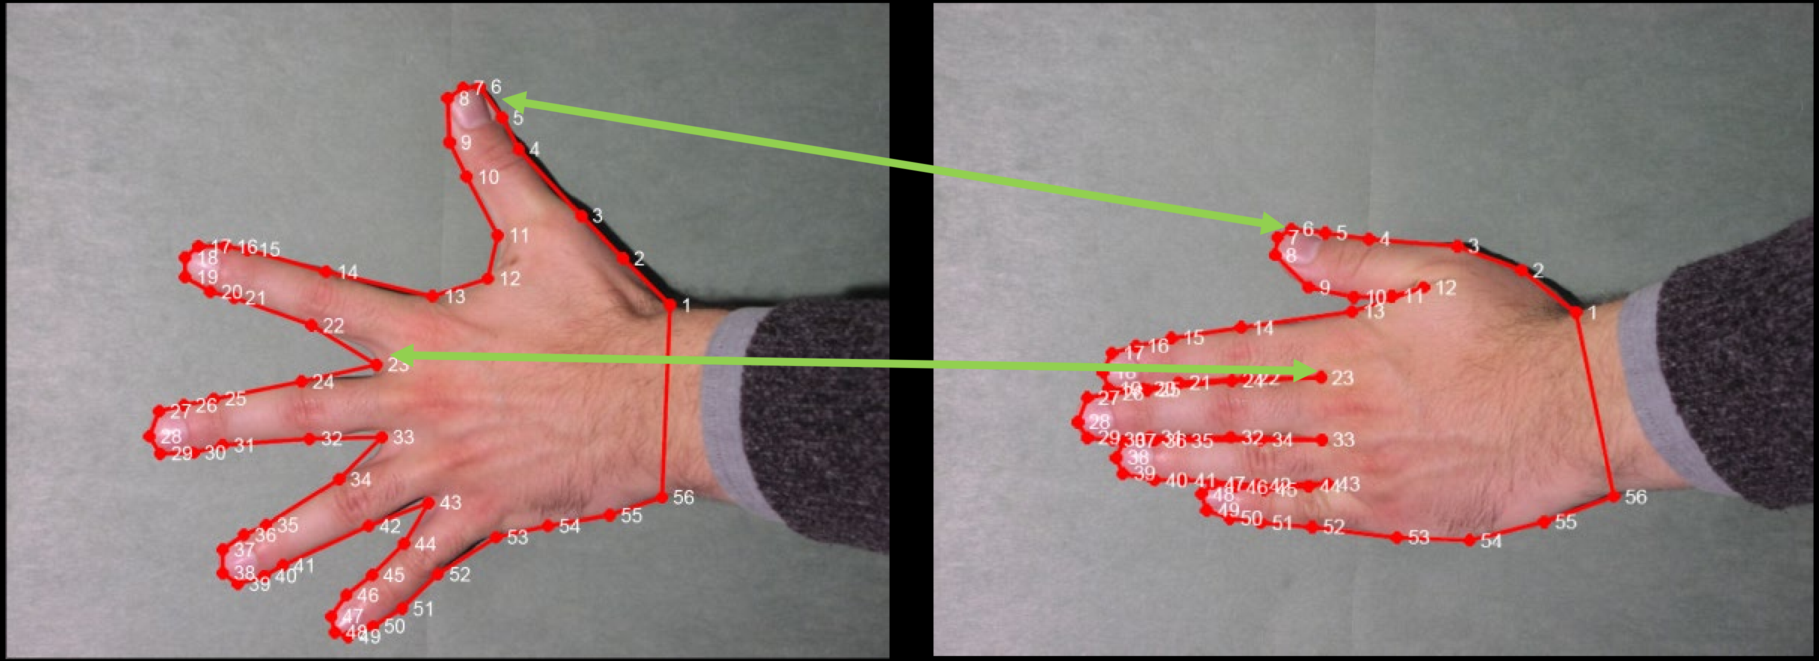

Shape as a vector

$$1 : (x_1, y_1)$$
$$2 : (x_2, y_2)$$
$$\vdots$$
$$N : (x_n, y_n)$$

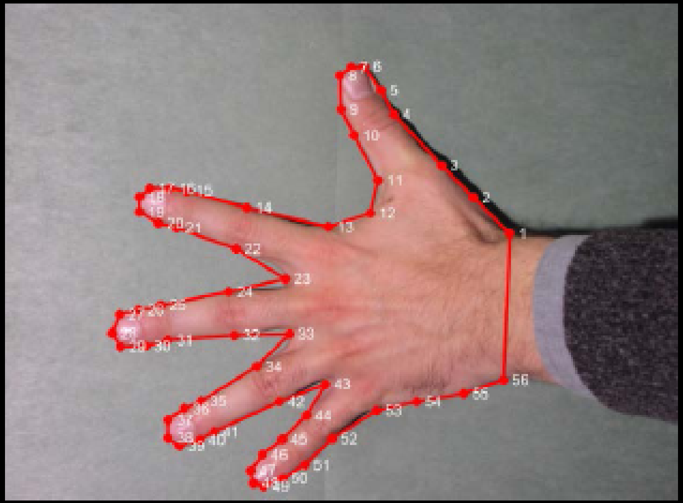

- The shape is represented as an array of (x,y) coordinates
- Trick number one!
  All coordinates are put into one vector!
- n=56 points
  - Vector with 112 elements!

$$\mathbf{x} = [x_1, x_2, \ldots, x_n, y_1, y_2, \ldots, y_n]^T$$

Shapes in high-dimensional space

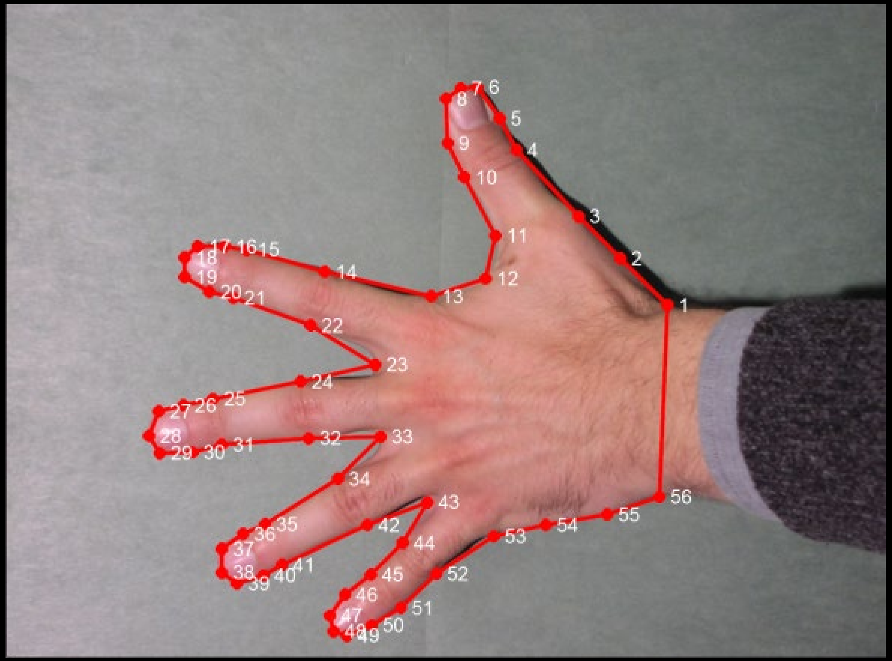

$$\mathbf{x} = [x_1, x_2, \ldots, x_n, y_1, y_2, \ldots, y_n]^T$$

- One hand is now described using one vector
- A vector can also be seen as a point in space!

Trick number two!

Coordinates in space

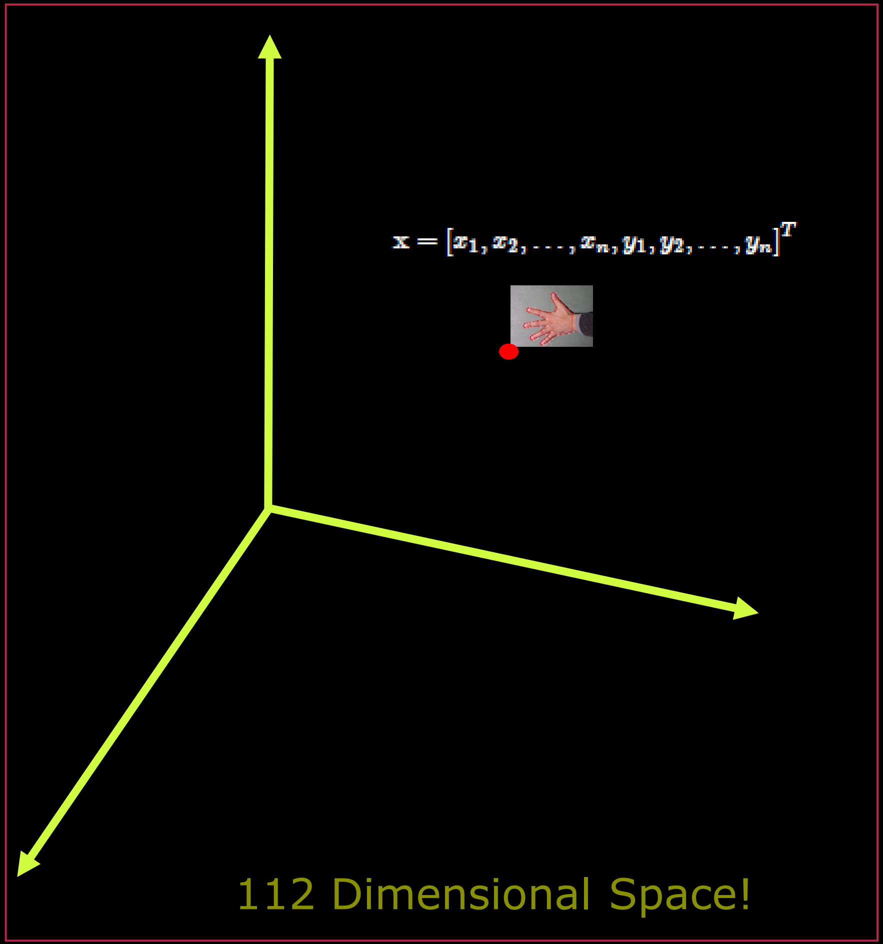

- On hand is now described using one vector
- A vector can also be seen as a coordinate in space!
- Not 2D space, not 3D space, not 4D space…
- 112 Dimensional Space!
- A hand has a position in this space!

Hands in Space

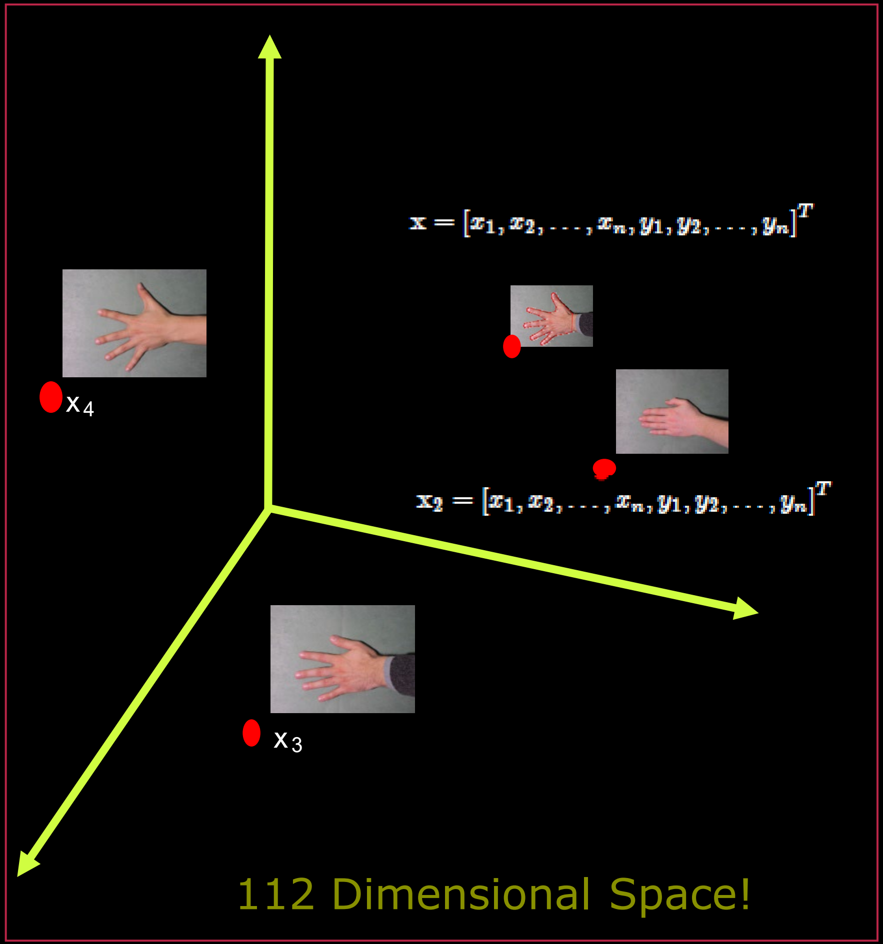

- A hand has a position in space!
- Another hand appears
  - in the same space
  - different position = different shape
- All hands have a place in this space!

Shape Analysis

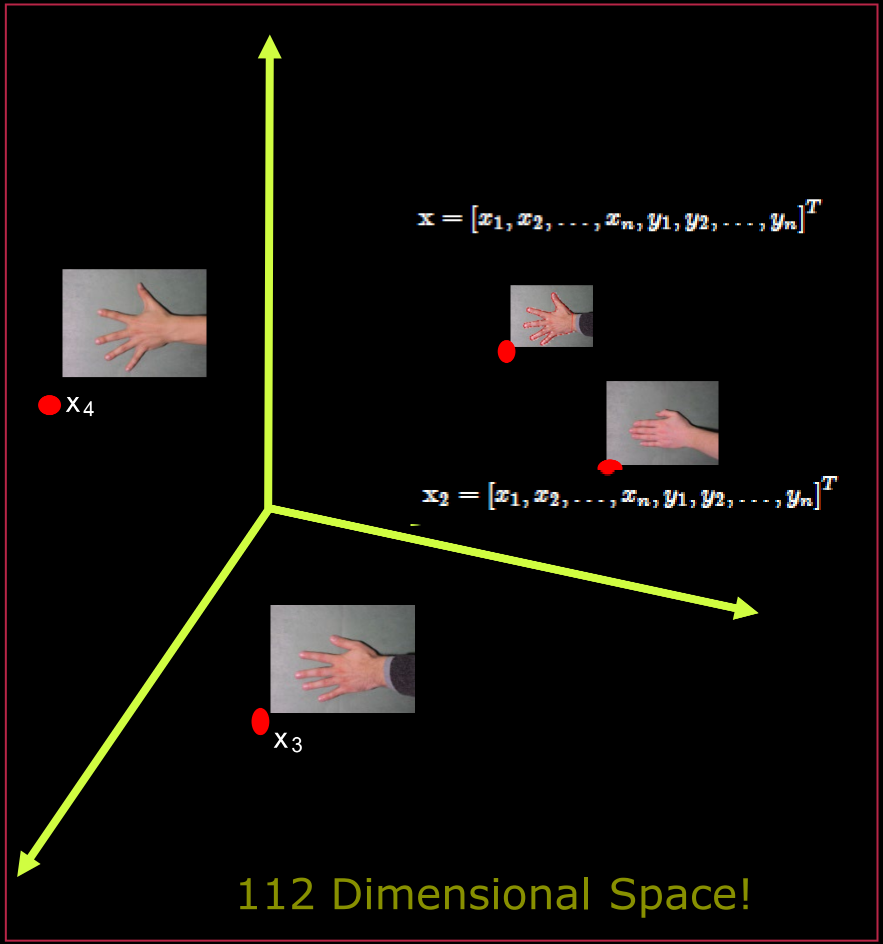

- Shape analysis
  - Similar shapes are placed on “planes” in the shape-space
  - Also called a manifold

Shape alignment

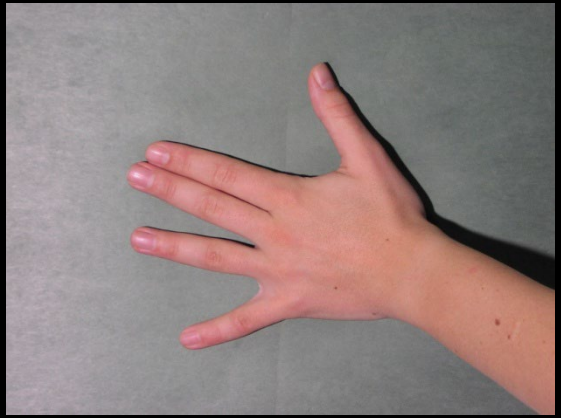
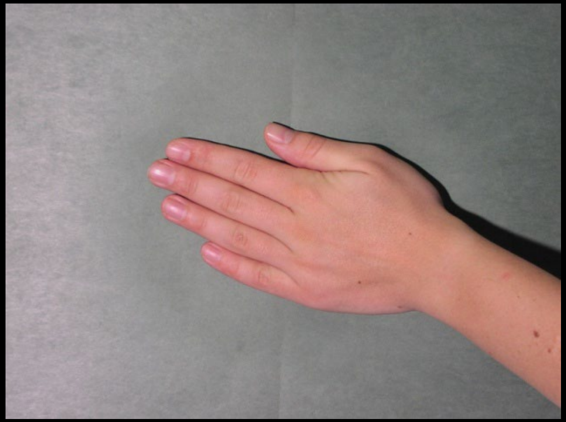
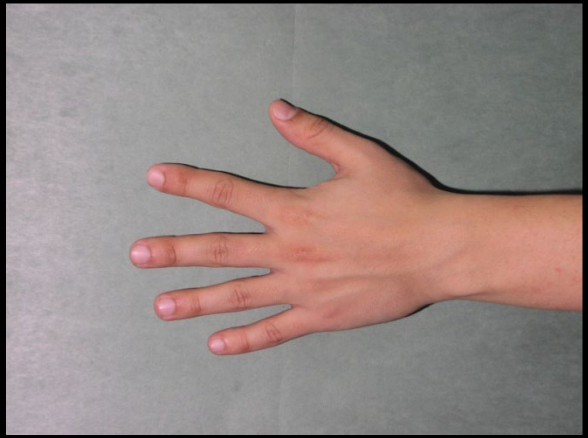

- 40 training images of hands
- 56 landmarks on each
- Placed in random location (translation+rotation)

Shape alignment

Before alignment

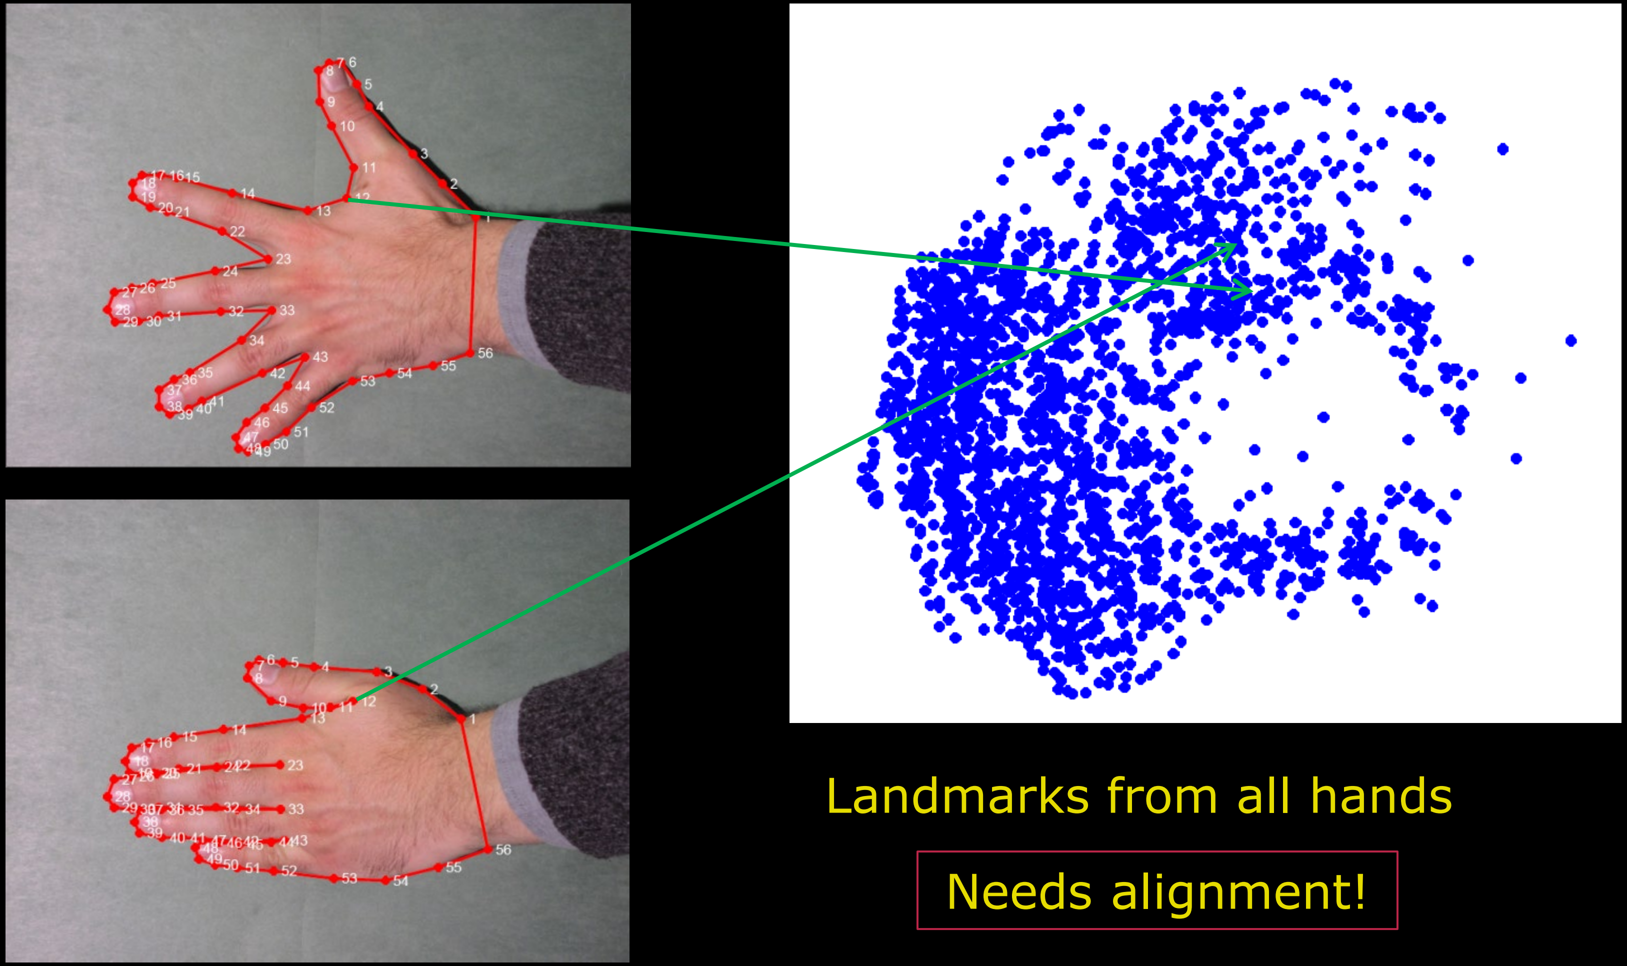

Landmarks from all hands

Needs alignment!

What is alignment?

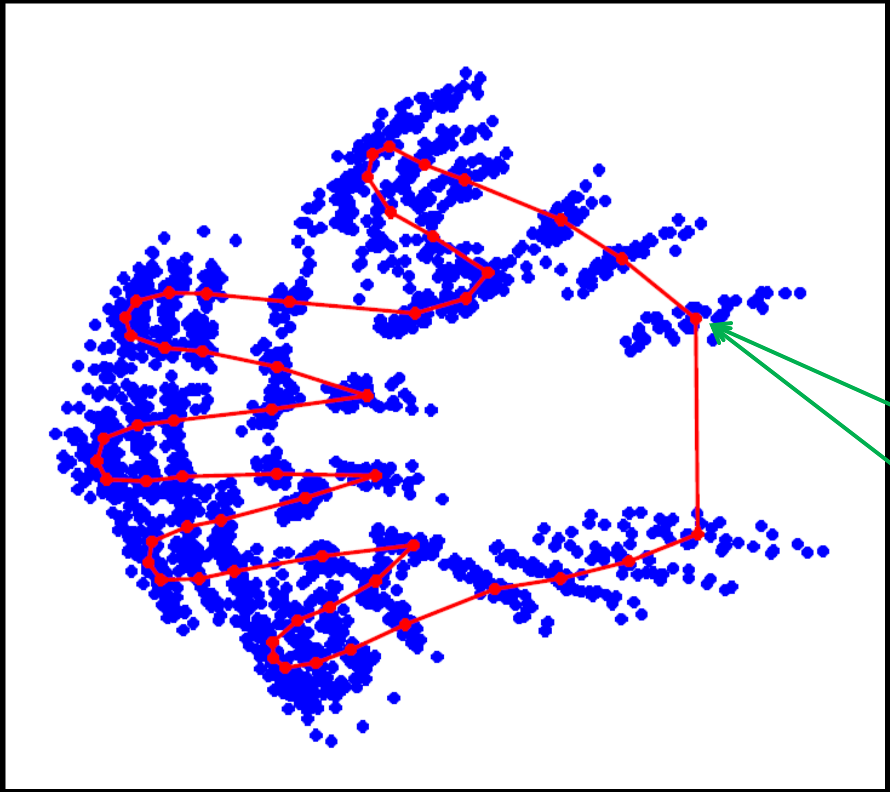

Average shape

- Group wise registration
  - Not one-to-one
  - All to the average shape

$$\bar{\mathbf{x}} = \frac{1}{N}\sum_{i=1}^{N}\mathbf{x}_i$$

$$\bar{\mathbf{x}} = [\bar{x_1}, \bar{x_2}, \ldots, \bar{x_n}, \bar{y_1}, \bar{y_2}, \ldots, \bar{y_n}]^T$$

But hey! We do not have an average shape?

Procrustes Analysis (alignment)

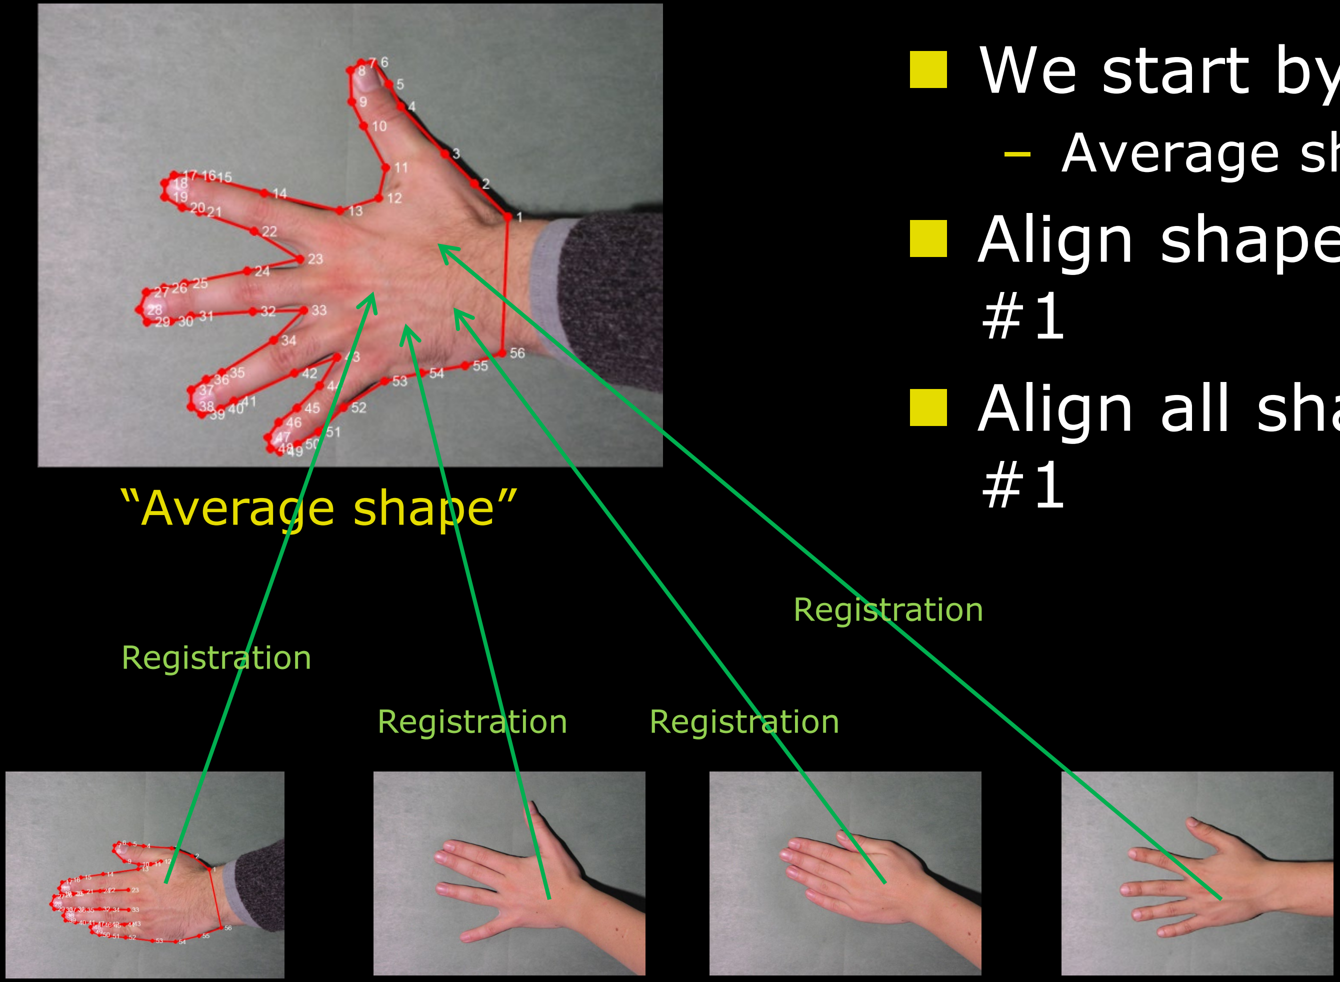

Shape #2

- We start by defining
  - Average shape = Shape #1
- Align shape #2 to shape #1
- Align all shapes to shape #1

Landmark based registration

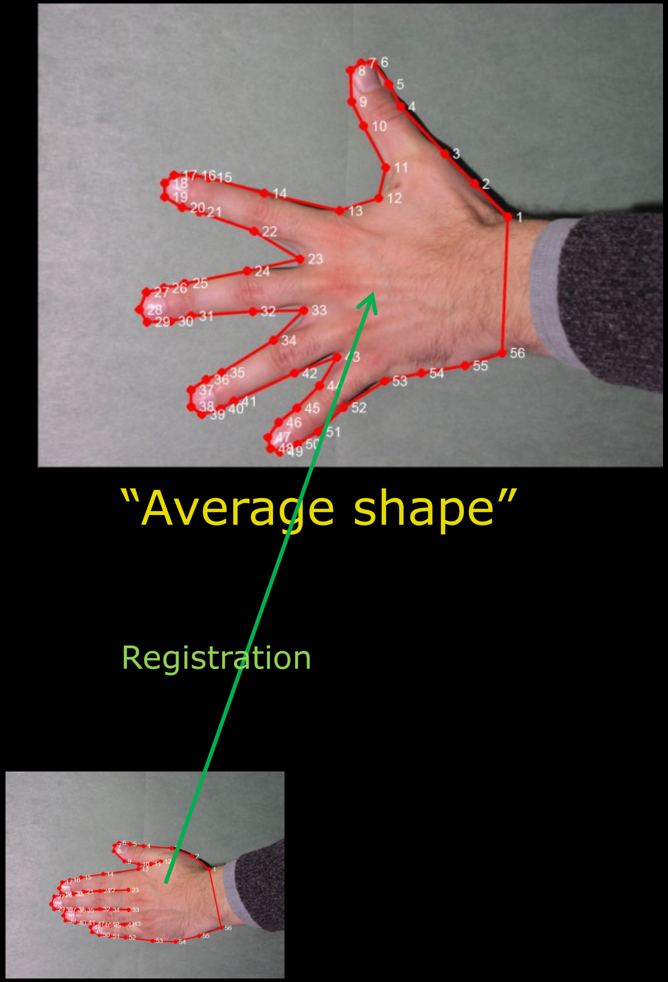

"Average shape"

Shape #2

- Shape #2 is transformed to fit the average shape
  - Translation
  - Rotation
  - Scaling
  - = Similarity Transform
- Result
  - Shape #2 is placed *on top of* the average shape

Procrustes Analysis

1. Average shape is set to shape #1
2. Register all shapes to the average shape
   – Landmark based registration
3. Recompute the average shape
4. If average shape changed return to step 2.

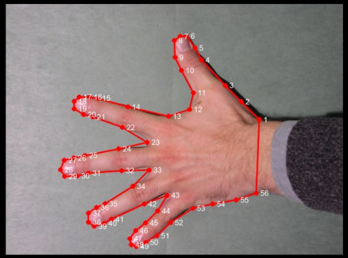

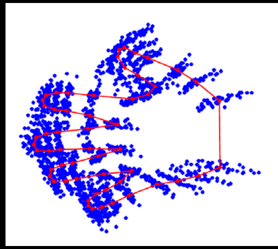

$$\bar{\mathbf{x}} = \frac{1}{N}\sum_{i=1}^{N}\mathbf{x}_i$$

Aligned shapes – what now

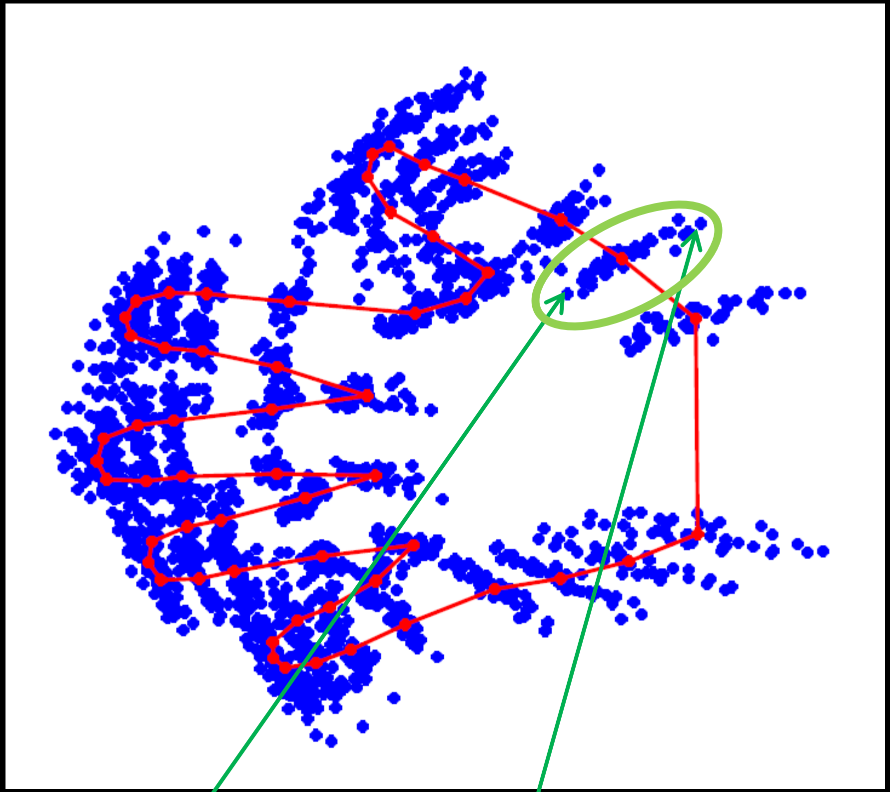

Shape #16 Shape #27

- Individual landmark variation
  - Over the training set
- What shape is the variation?

Principal Component Analysis (PCA)

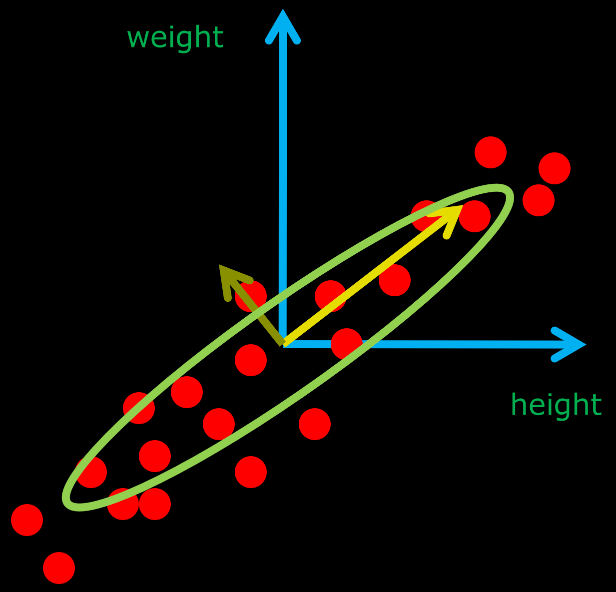

- PCA
  - Main axis in data
  - Eigenvectors
  - Eigenvalues
- Size of Eigenvalues describe explained variance

Principal Component Analysis (PCA)

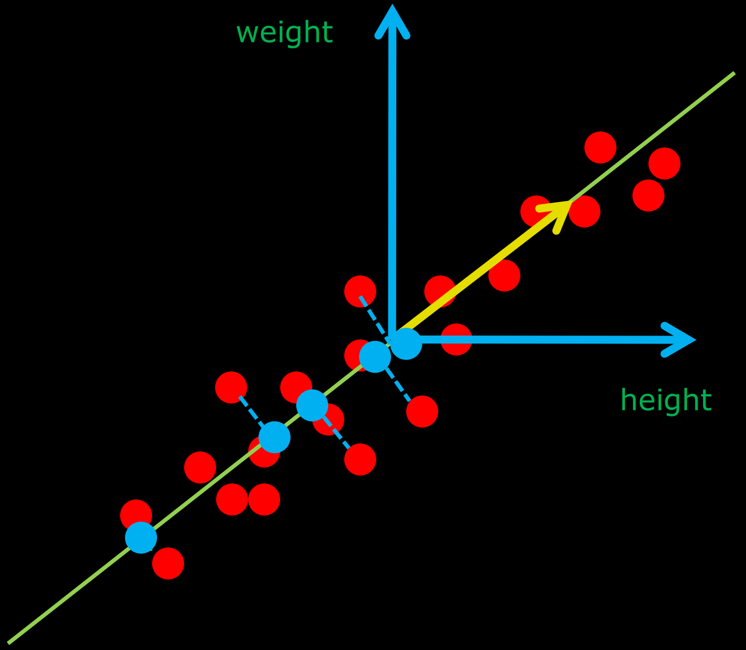

- We throw away the *noise dimensions*
- Points projected to the line

Principal Component Analysis (PCA)

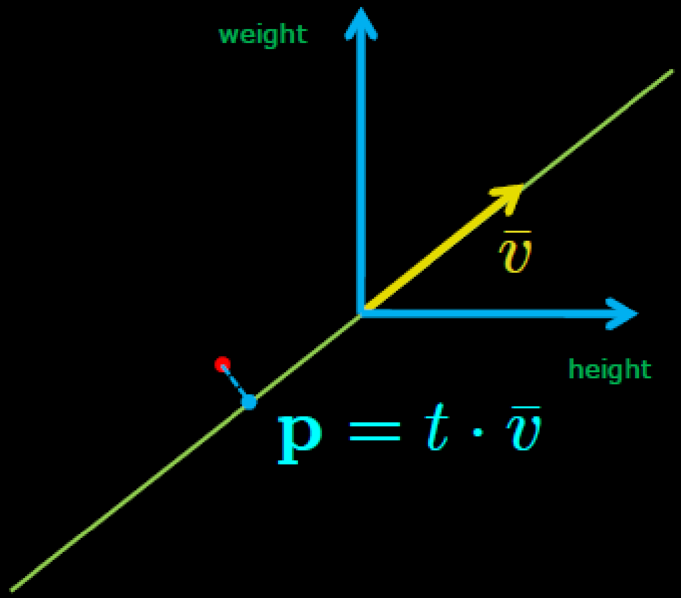

- We throw away the *noise dimensions*
- Points projected to the line
- A point can now be described by one parameter *t*
- We have reduced the number of dimensions

How many dimensions should we keep?

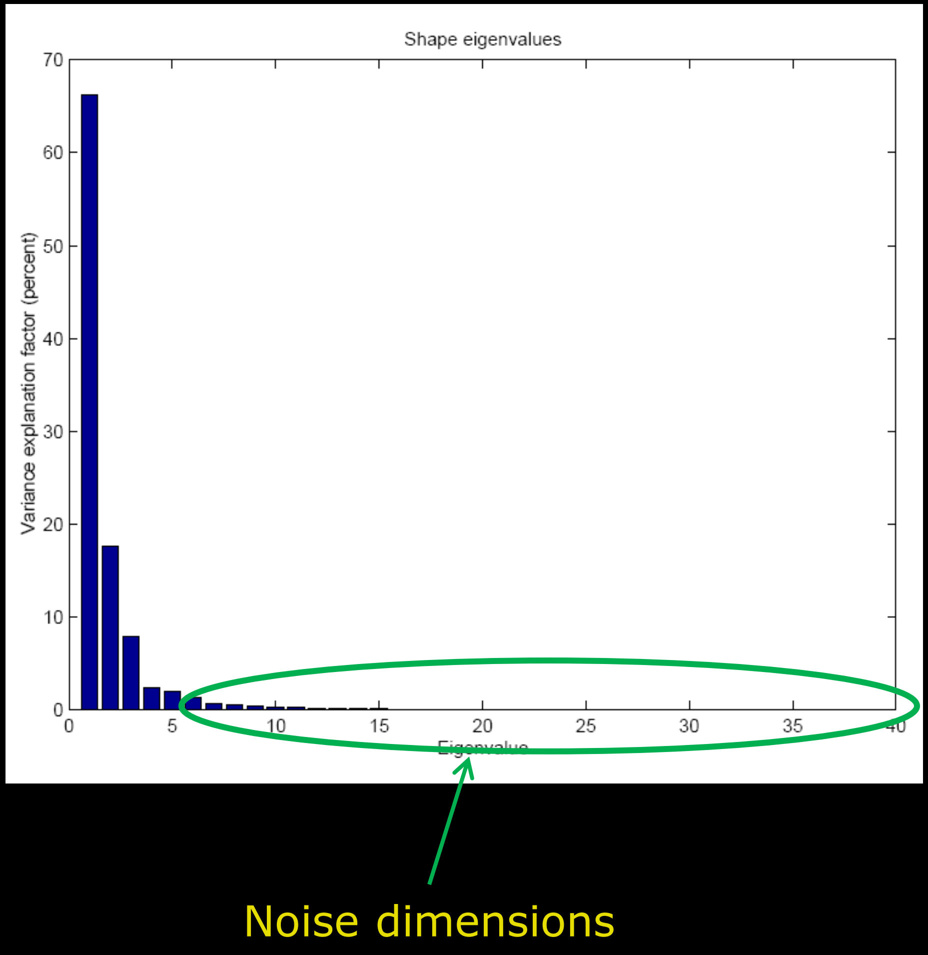

- Plot the Eigenvalues
- Explains how *important* each dimension is
- Cut away noise dimensions

Aligned shapes – what now

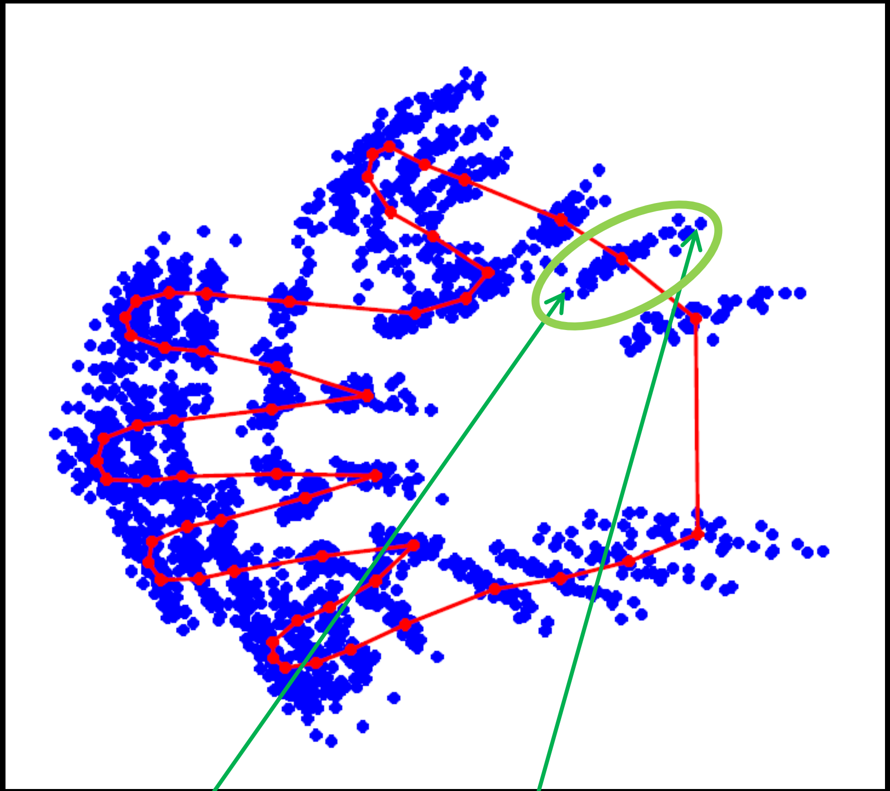

Shape #16

Shape #27

- Individual landmark variation
  - Over the training set
- What shape is the variation?

PCA Analysis

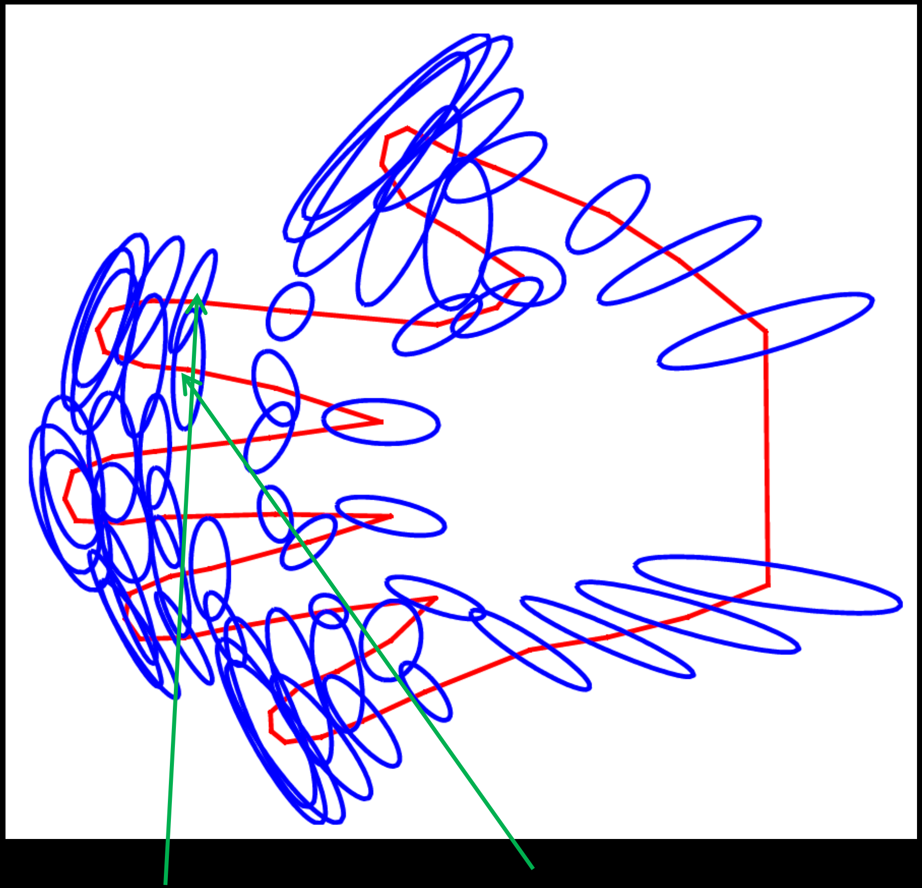

- PCA analysis on individual landmarks
- Describes the major direction of variation
- Landmarks are correlated!
- The movement over the shape is connected
- Return to shape space

PCA in shape space

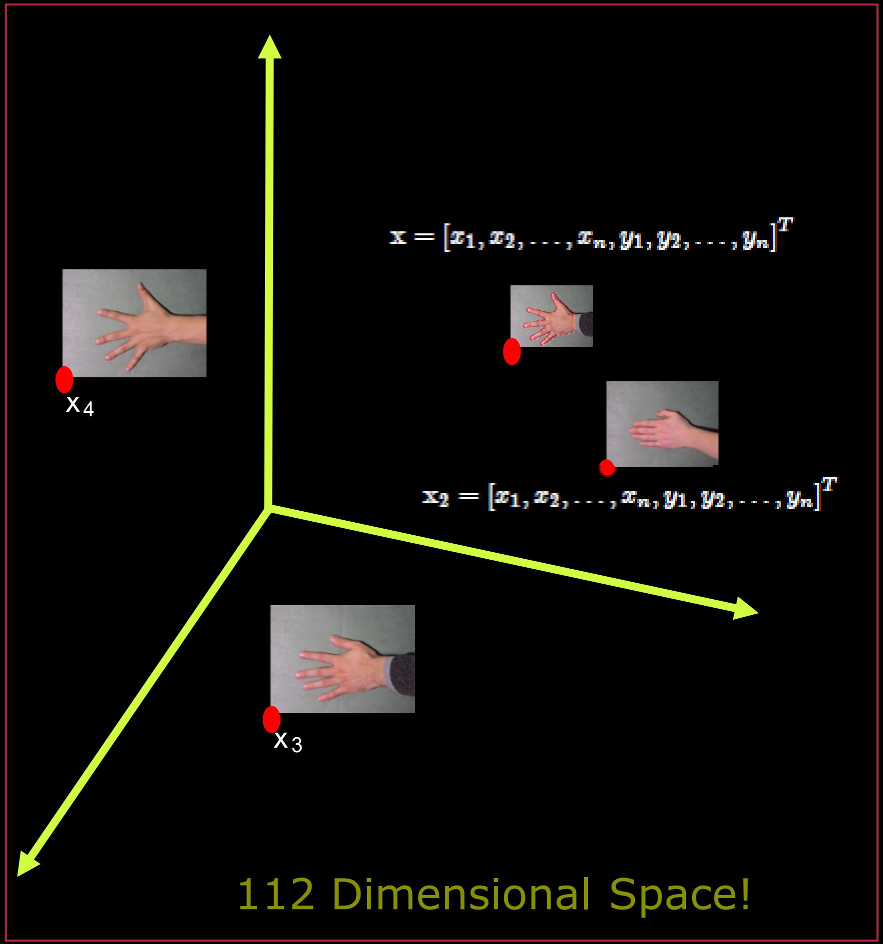

- Instead of doing PCA on 2D points we do it on 112D points
- Examine if our 40 *shapes is lying on a plane* in 112D space
- We find the directions that spans the maximum variation in shape space

Start by computing the shape average

$$\bar{\mathbf{x}} = \frac{1}{s}\sum_{i=1}^{s}\mathbf{x}_i$$

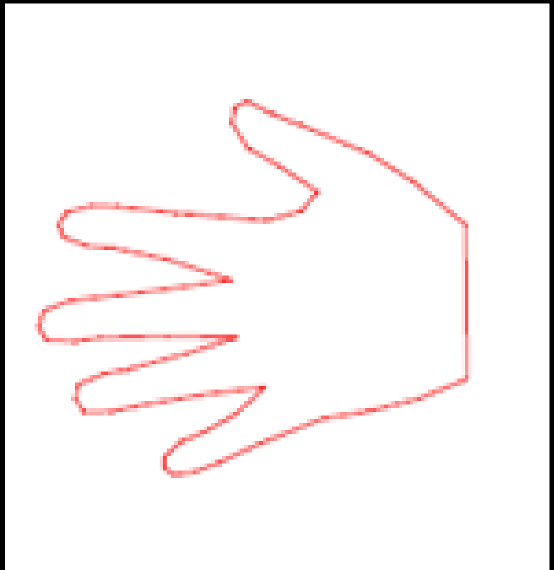

Since we do this on the aligned shapes – this is the Procrustes average

Do the eigenvector analysis

$$\mathbf{S} = \frac{1}{s-1}\sum_{i=1}^{s}(\mathbf{x}_i - \bar{\mathbf{x}})(\mathbf{x}_i - \bar{\mathbf{x}})^T$$

- Computing the covariance of the shape data

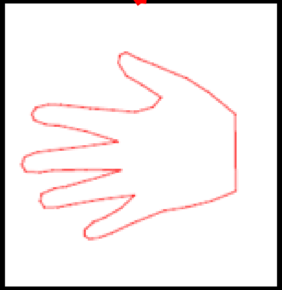

Average shape

Shape number i in the training set

Visualizing variation

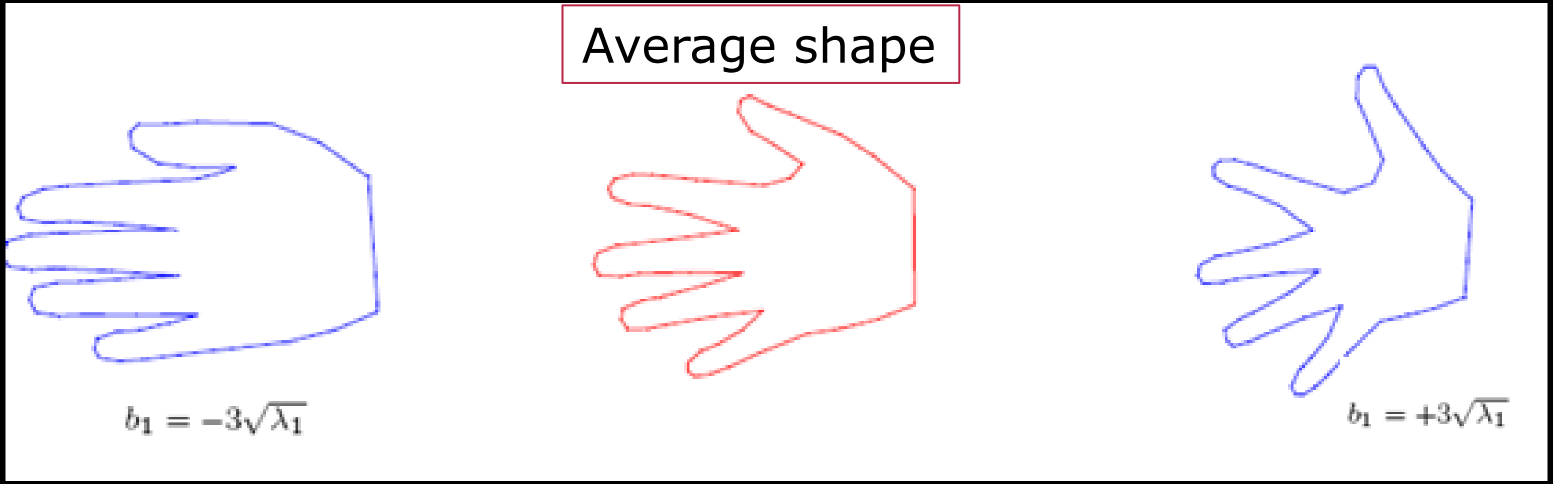

Visualizing the first principal component

$$\mathbf{x} \approx \bar{\mathbf{x}} + \Phi \mathbf{b}$$

Φ contains the t eigenvectors

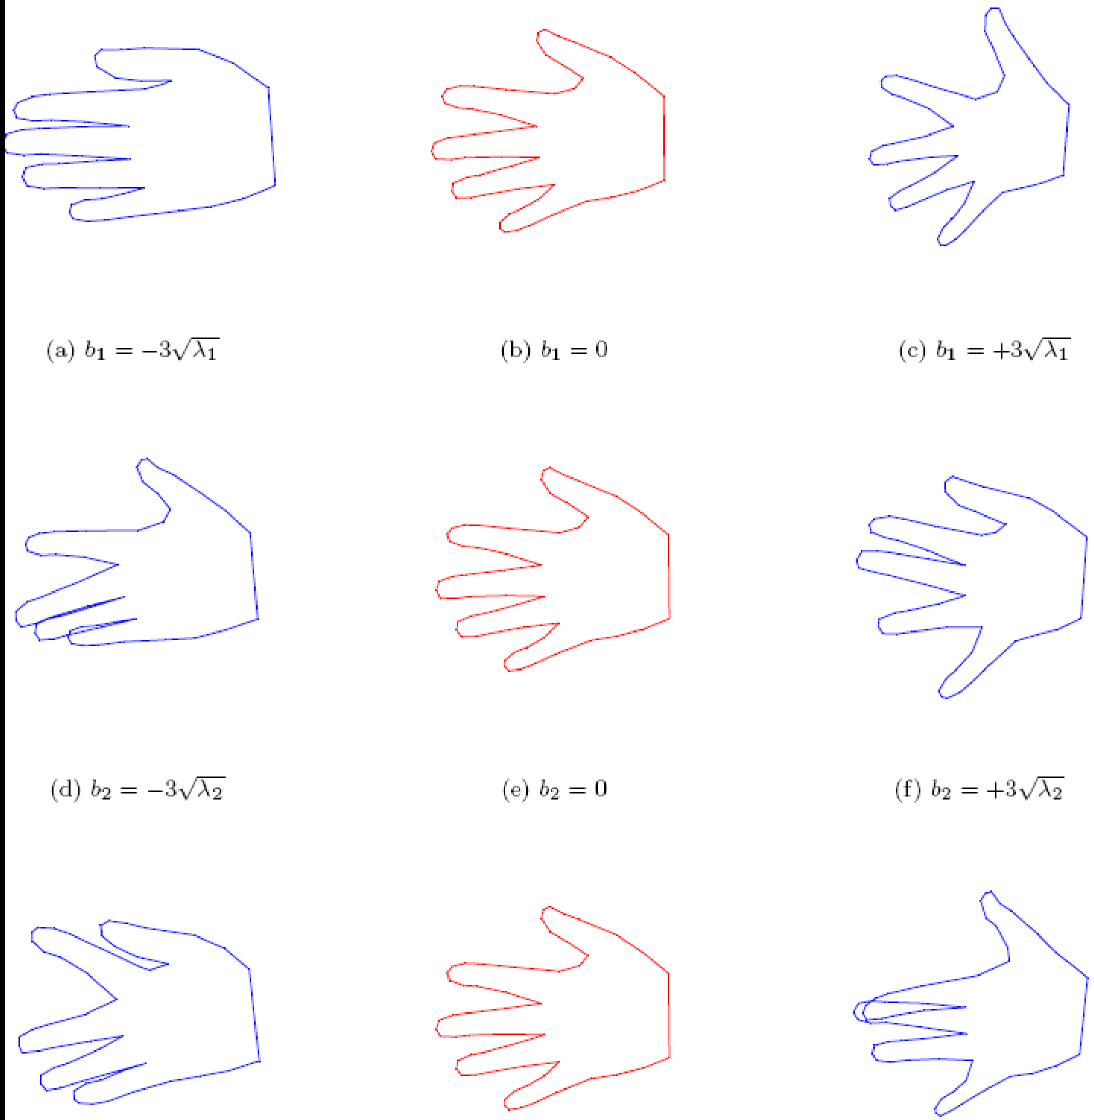

(a) $b_1 = -3\sqrt{\lambda_1}$ (b) $b_1 = 0$ (c) $b_1 = +3\sqrt{\lambda_1}$

(d) $b_2 = -3\sqrt{\lambda_2}$ (e) $b_2 = 0$ (f) $b_2 = +3\sqrt{\lambda_2}$

(g) $b_3 = -3\sqrt{\lambda_3}$ (h) $b_3 = 0$ (i) $b_3 = +3\sqrt{\lambda_3}$

Results of Shape Analysis

- Visualisation of the major variation of the shape over a population

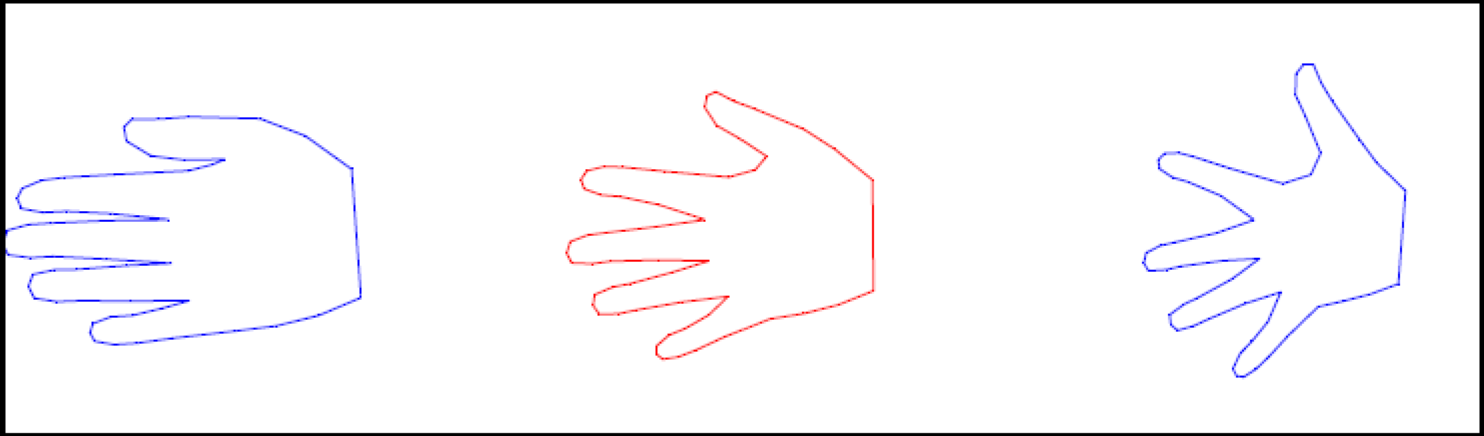

Hearing Aid Design

- Main variation of the shape of the ear canal
- Found using principal component analysis
- First mode of variation
- 7 modes explain 95% of the total variation

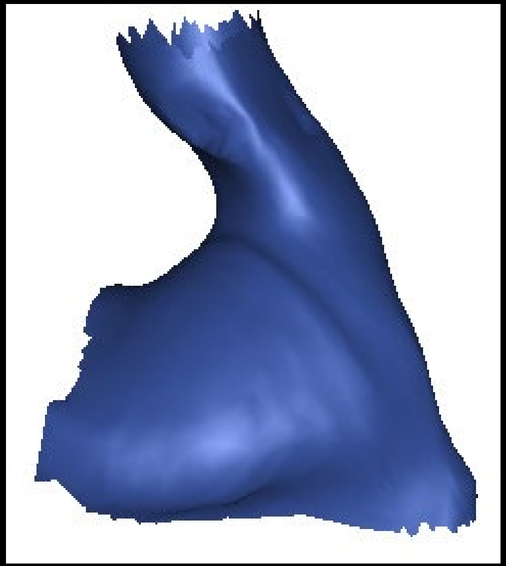

Average-1. mode

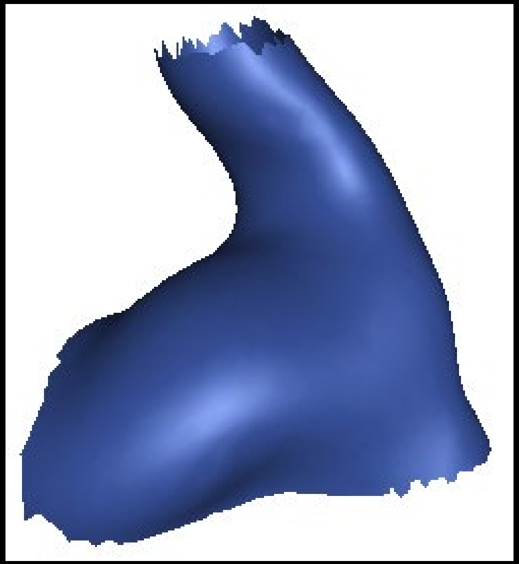

Average

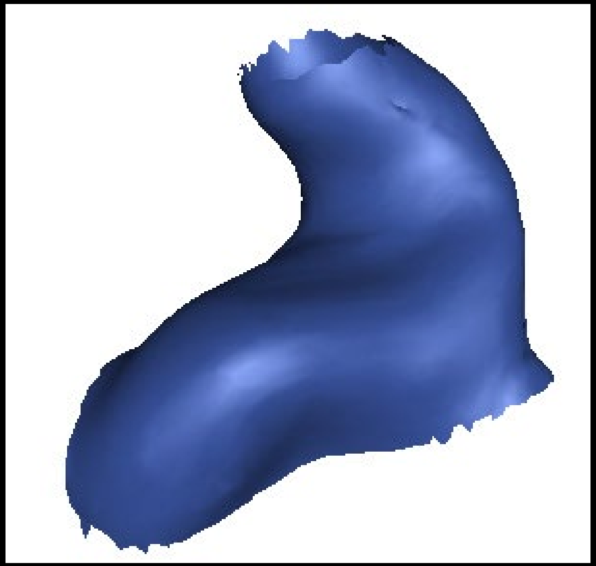

Average+1. mode

Modelling shape and appearance

- A model that can both model the shape of an object and the appearance (the texture)
- Texture: The pattern of intensities (or colors) across an image patch

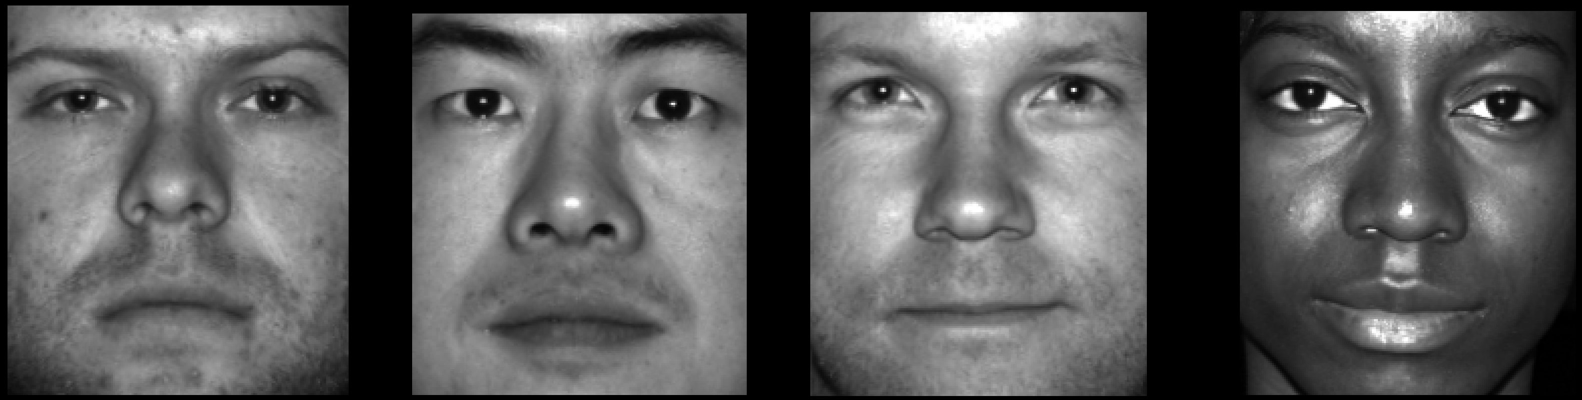

Back to lecture 3: Eigenfaces

Face data

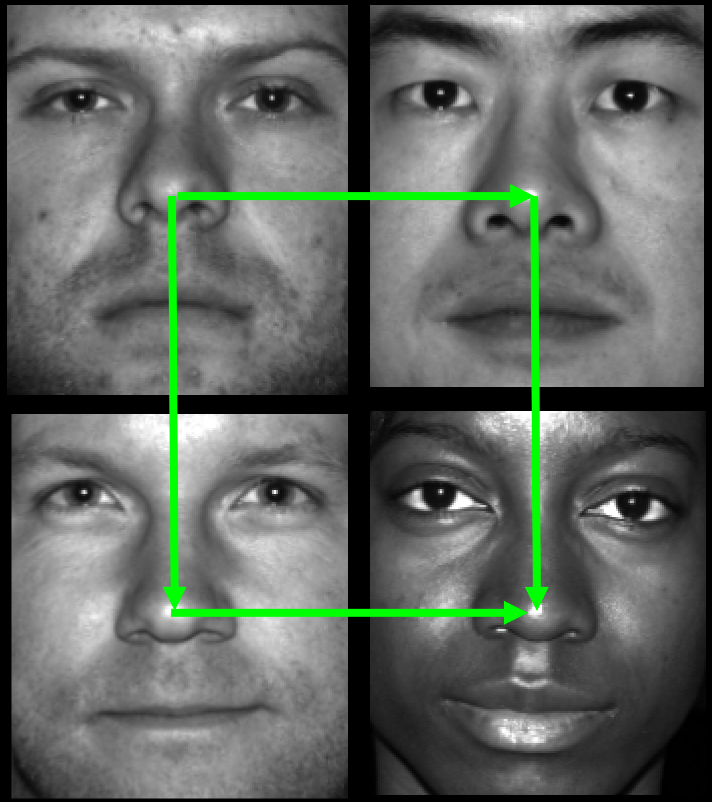

- 38 face images
  - 168 x 192 grayscale
- Aligned
  - The anatomy is placed ”in the same position in all image”
- Same illumination conditions on the images we use

The Extended Yale Face Database B

Analyzing the deviation from the mean face

- We want to do the principal component analysis on the *deviations from the average face*

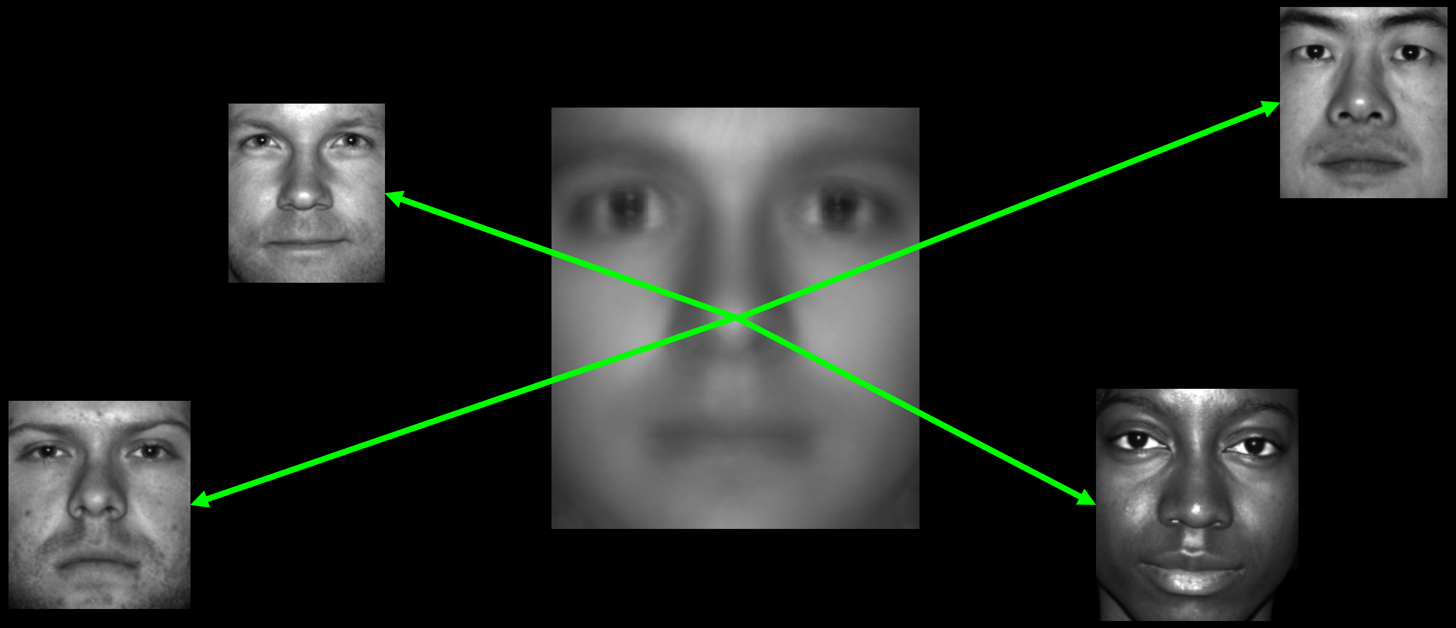

Visualizing the PCA faces

Main deviations from the average face

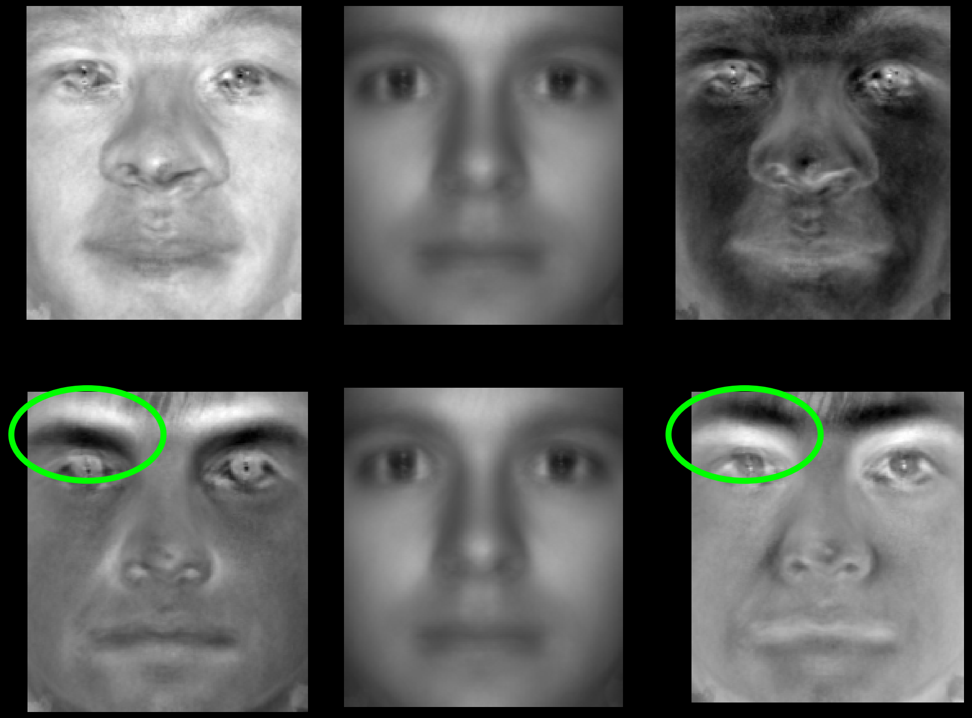

-PC | Average face | +PC

First PC – 40% of variation

Second PC – 8% of variation

A tool to see major variations – brow lifting

Eigenfaces: Modelling texture

- The modelling of the pure appearance
- Without removing variation in shape
- No *decoupling* of shape and appearance

Decoupling shape and texture

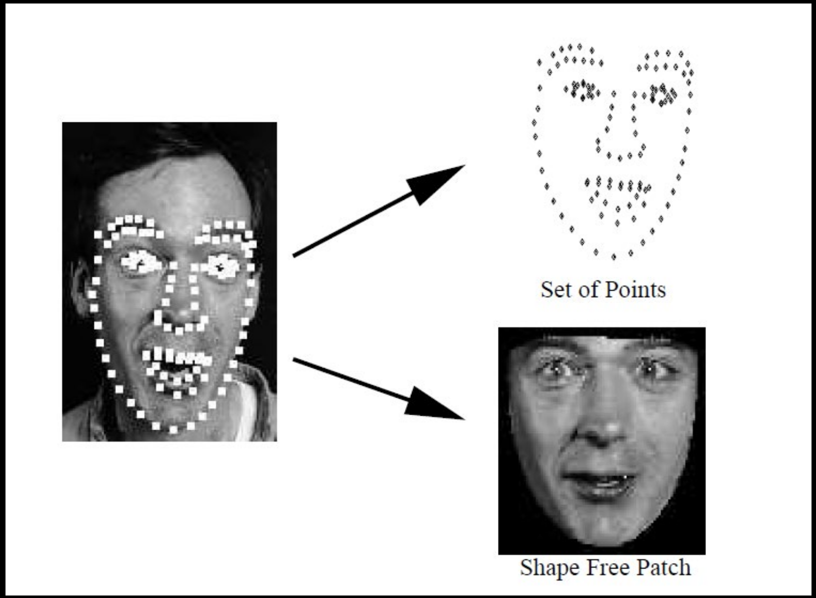

- Warp each face to average shape using the landmarks
- Non-linear geometrical transformation
- Sample the texture from the warped face

Eigenfaces on warped faces

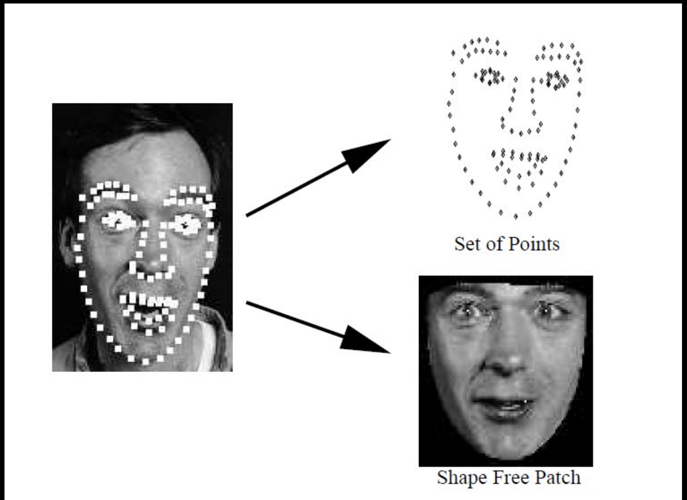

- Same PCA modelling as in the Eigenfaces approach
- Just slightly different notation

$$\mathbf{g} = \bar{\mathbf{g}} + \mathbf{P}_g \mathbf{b}_g$$

Combined shape and appearance model

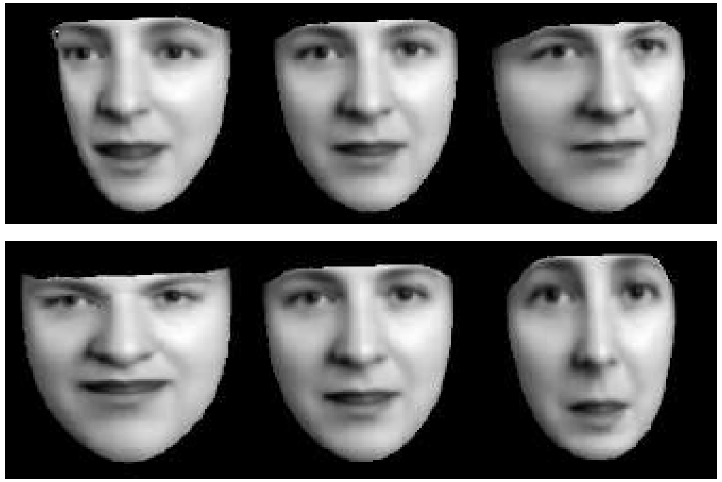

Figure 5.2: First two modes of shape variation (± 3 sd)

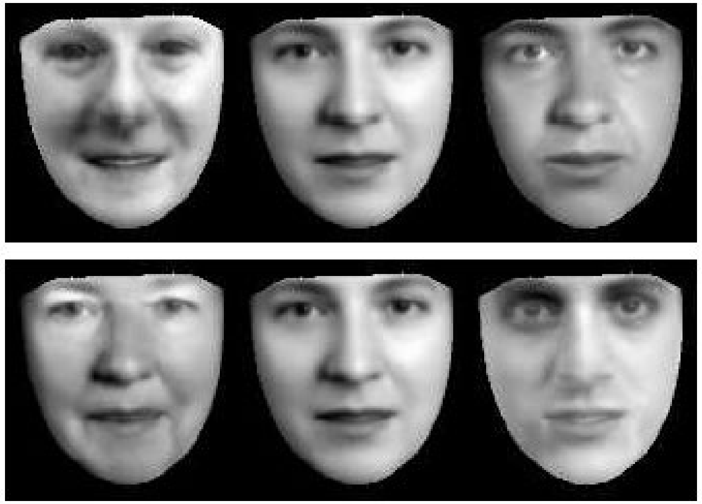

Figure 5.3: First two modes of grey-level variation (± 3 sd)

Facial Analysis

- Demo of AAM explorer