Pattern Recognition Lab

FRIEDRICH-ALEXANDER UNIVERSITÄT ERLANGEN-NÜRNBERG

FACULTY OF ENGINEERING

A Practical Guide to Statistical Shape Models Featuring Hands-on Examples in CONRAD

Anna Gebhard, Tobias Geimer, Mathias Unberath, Frank Schebesch, Mena Abdelmalek, Andreas Maier

https://www5.cs.fau.de/conrad/

Outline

I. Theory
   1. Why Statistical Shape Models (SSM)?
   2. Data Representation
   3. Model Creation Pipeline
      a) Alignment and Scaling
      b) Computation of a Mean Shape
      c) Extraction of Variation
   4. Principal Component Analysis (PCA)
   5. PCA Using Singular Value Decomposition
   6. Trained SSM & Functionality

II. Practice: Tutorial Code

III. Limitations of SSMs

Why SSMs?

Statistical Shape Models (SSMs) can express a range of expected, evidence-based **variation** on top of a **mean shape** derived from a cohort or population.

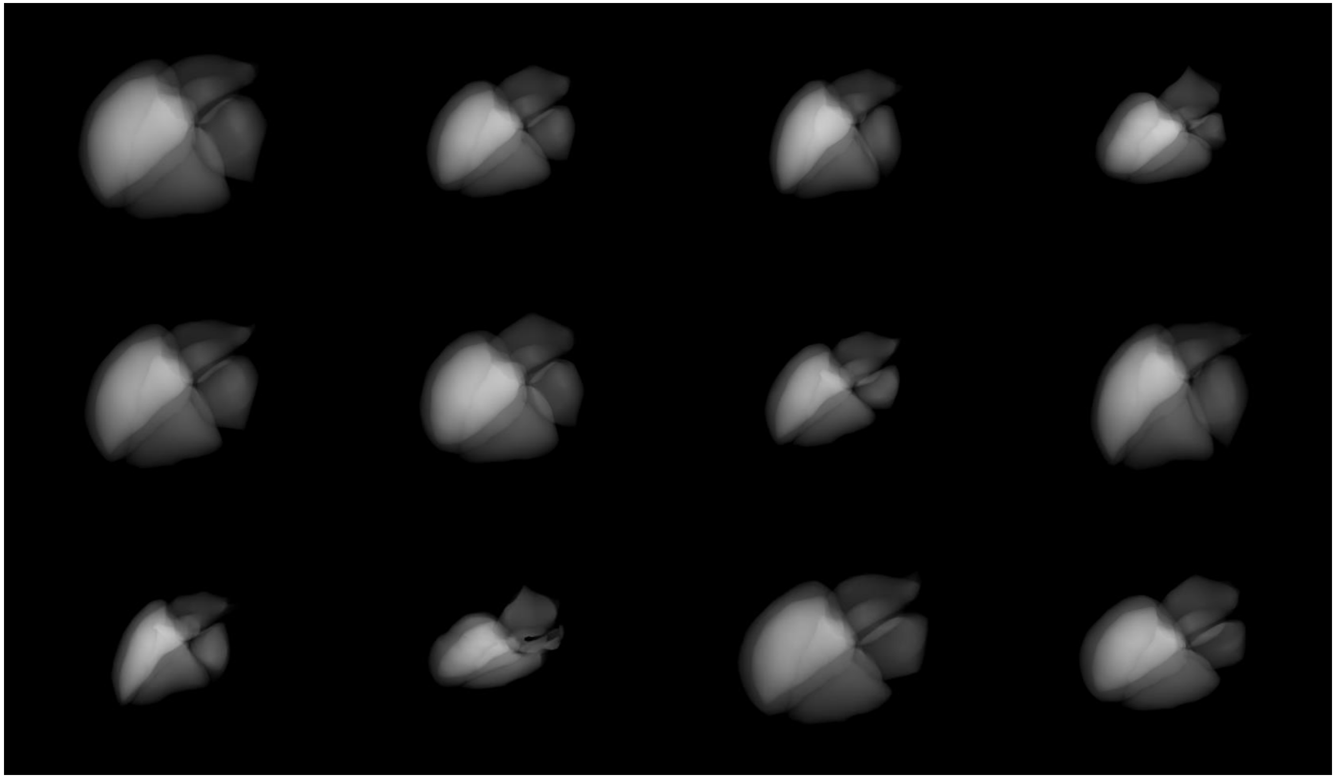

Figure 1: Training set of cardiac model variations

Why SSMs?

Mean Shape:
Knowledge about the general shape

Variation:
Knowledge about how much the shape can differ between subjects

This can be used in different applications of SSMs:
- Classification
- Segmentation
- Phantom generation

Applications of SSMs

Classification:
Use the variation to differentiate between classes

Segmentation:
Use prior knowledge about the shape to segment images

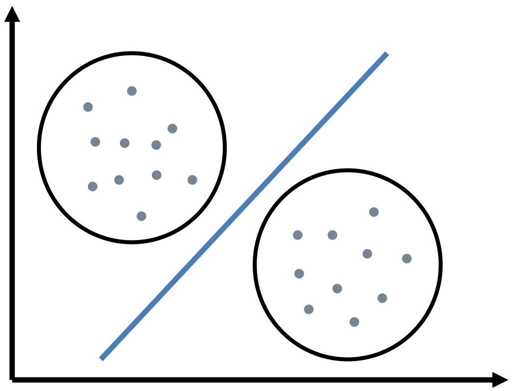

Figure 2: Two variation clusters

Figure 3: Segmented lung contour in a CT slice

Applications of SSMs

Phantom Creation:
Use shape extrapolation to create realistic representations within given variation

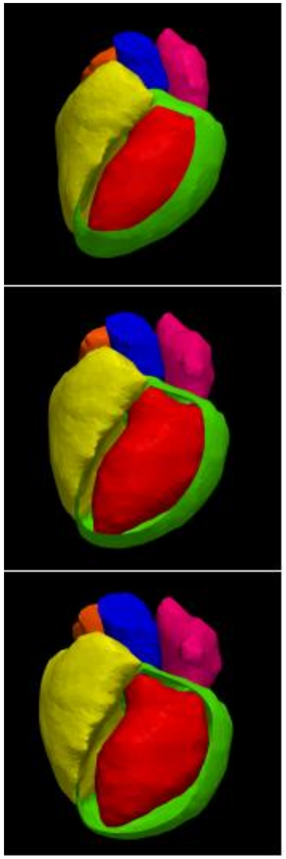

Figure 4: Different heart shapes[1]

[1] M. Unberath et al. (2015-July), "Open-source 4D statistical shape model of the heart for x-ray projection imaging", Proc ISBI 2017, pp. 739–742

Data Representation

Point Cloud

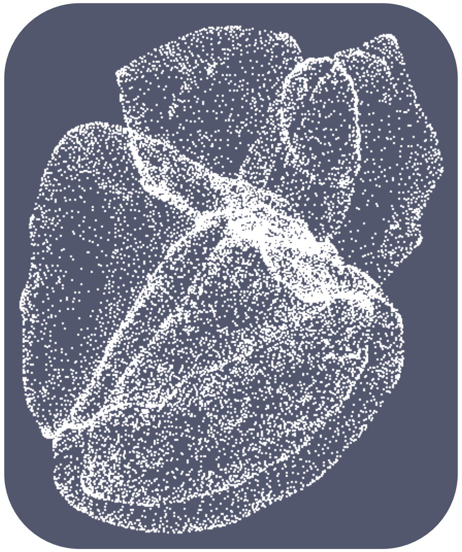

Set of K 3-D points
$\boldsymbol{p}_i \in \mathbb{R}^3, \mathrm{i} \in \{1, \dots, K\}$

Mesh

…with connectivity information (edges)

…

- B-splines
- NURBS
- Segmented volumetric images

Point correspondence or ways to establish it are crucial.[2]

[2] A. Sinha et al. (2017-February), „Simultaneous segmentation and correspondence improvement using statistical modes“, Med Imag: Imag Proc, vol. 10133, p. 101331B

Model Creation Pipeline

To create a statistical shape model, representative **training data** of shapes within a given **population** is needed. This data is used in the following pipeline:

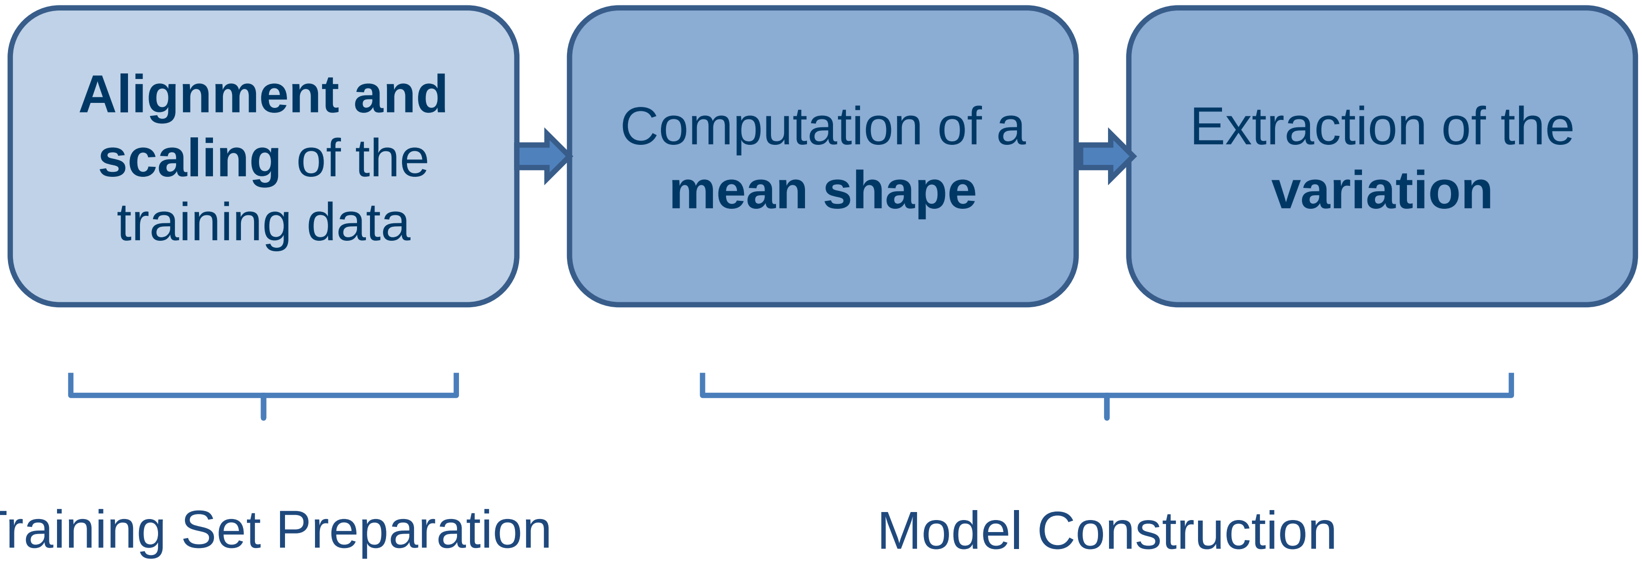

Alignment and Scaling

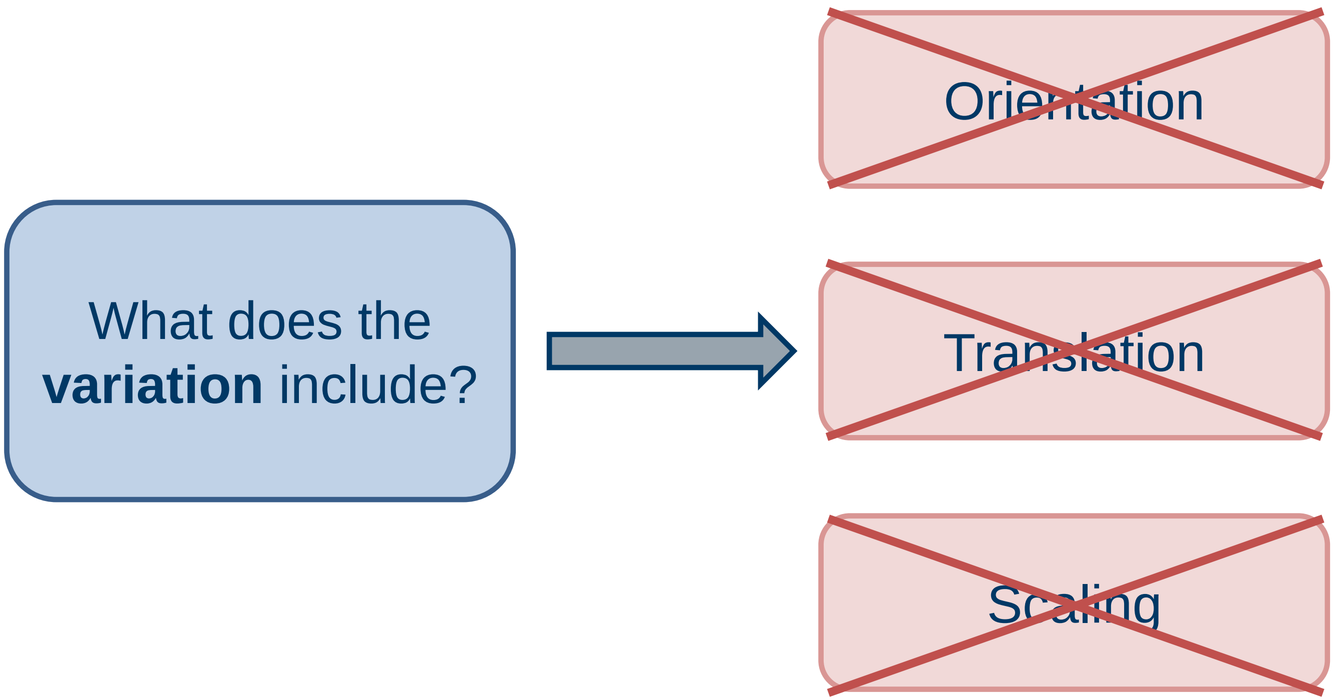

Any variation that is not supposed to be modeled needs to be removed prior to model construction.

Alignment and Scaling

What does the **variation** include? → Only actual differences between the shapes.

Remove rotation, translation and scaling from the training data!

$$p' = s\boldsymbol{R}\boldsymbol{p} + \boldsymbol{t}$$

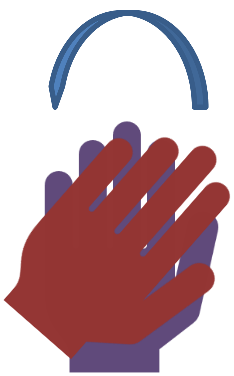

Rotation $\boldsymbol{R}$

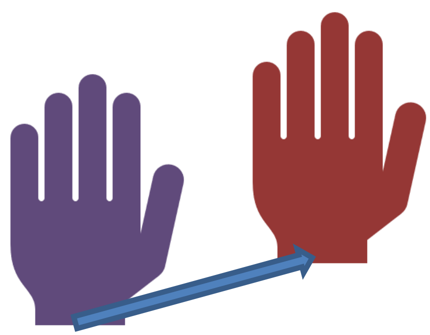

Translation $\boldsymbol{t}$

Scaling s

Computation of a Mean Shape

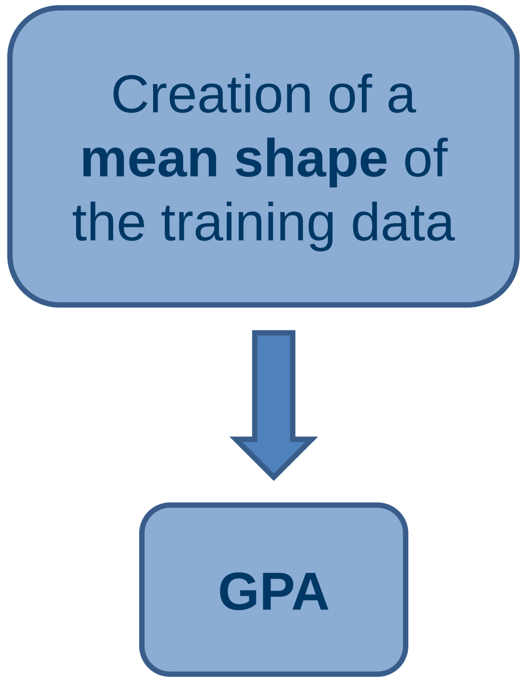

Generalized Procrustes Analysis (GPA)[3] is used to iteratively determine a mean shape to which all shapes in the training set have minimal distance (includes alignment and scaling).

[3] J. C. Gower et al. (1975), “Generalized Procrustes Analysis”, Psychometrika, vol. 40, pp. 33–51

Extraction of Variation

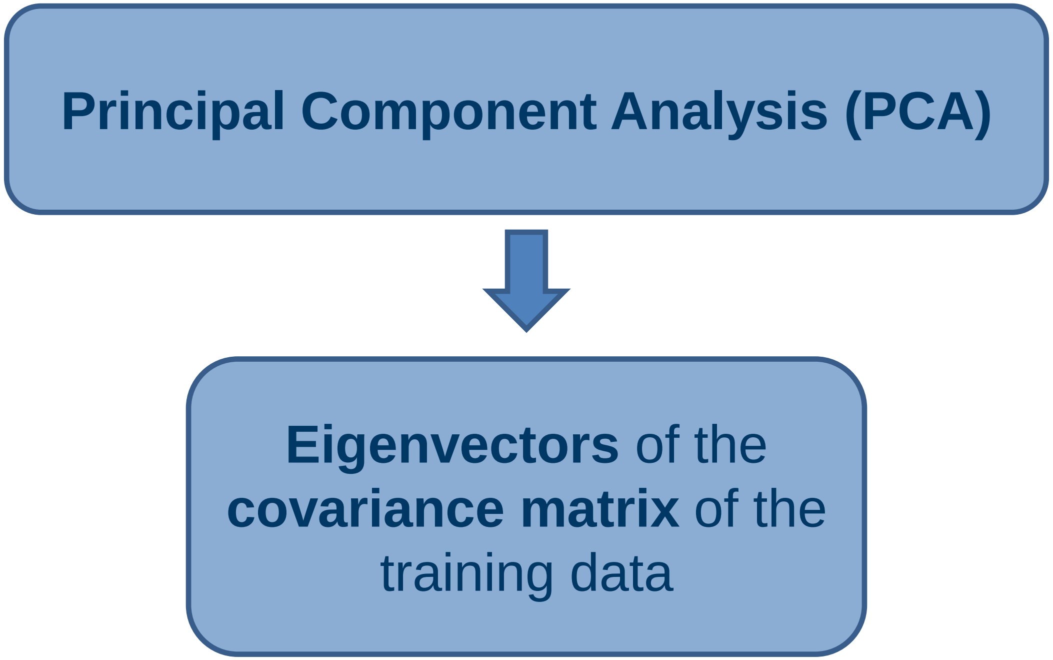

The idea of **Principal Component Analysis (PCA)** is to extract the principal modes of variation by computing the eigenvectors of the covariance matrix.

Principal Component Analysis (PCA)

Training data matrix

$$\boldsymbol{Y} = (\boldsymbol{x}_1, \boldsymbol{x}_2, \dots, \boldsymbol{x}_N) \in \mathbb{R}^{3K \times N}$$

with **zero-centered samples**

$$\mathbf{x}_\mathrm{i} \in \mathbb{R}^{3K}$$

Covariance matrix

$$\boldsymbol{C} = \frac{1}{N-1}\boldsymbol{Y}\boldsymbol{Y}^T \in \mathbb{R}^{3K \times 3K}$$

Principal Component Analysis (PCA)

Eigen-equation

$$\boldsymbol{C}\Phi_{\mathrm{j}} = \lambda_j \Phi_{\mathrm{j}}$$

The **eigenvectors** $\boldsymbol{\Phi}_j$ corresponding to the **eigenvalues** λ_j of the covariation matrix C represent the directions of variation present in the data. To get the **principal modes of variation** of the statistical shape model, the first M eigenvectors are chosen.

What is M?

Principal Component Analysis (PCA)

Cumulative variance

$$\mathrm{var}(M) = \frac{\sum_{j=1}^{M} \lambda_j}{\sum_{k=1}^{3K} \lambda_k}$$

The **number of components** M is determined by the **cumulative variance** exceeding a certain threshold, e.g. 90 %. This means that **90 % of the variation** in the training data can be explained by the first M eigenvectors.

(Truncated) **basis of eigenvectors**

$$\boldsymbol{\Phi} = [\Phi_1, \ldots, \Phi_M] \in \mathbb{R}^{3K \times M}$$

PCA Using Singular Value Decomposition

Diagonalized covariance matrix

$$\boldsymbol{C} = \boldsymbol{U}\boldsymbol{L}\boldsymbol{U}^T$$

The covariance matrix is symmetric and can be diagonalized with $\boldsymbol{U} \in \mathbb{R}^{3K\times 3K}$, a **matrix of eigenvectors,** and $\boldsymbol{L} \in \mathbb{R}^{3K\times 3K}$, a **diagonal matrix of the eigenvalues** in decreasing order.

SVD of the data matrix Y

$$\boldsymbol{Y} = \boldsymbol{U}\boldsymbol{S}\boldsymbol{V}^T$$

$\boldsymbol{U} \in \mathbb{R}^{3K\times 3K}$ and $V \in \mathbb{R}^{N\times N}$ are **unitary matrices** and $\boldsymbol{S} \in \mathbb{R}^{3K\times N}$ contains the **singular values** on the diagonal.

PCA Using Singular Value Decomposition

$$\boldsymbol{C} = \frac{1}{N-1}(\boldsymbol{USV}^T)(\boldsymbol{USV}^T)^T$$

$$= \frac{1}{N-1}\boldsymbol{USV}^T\boldsymbol{VSU}^T$$

$$= \boldsymbol{U}\frac{\boldsymbol{S}^2}{N-1}\boldsymbol{U}^T$$

This shows that the **singular values of the data matrix** are related to the **eigenvalues of the covariance matrix** via $\lambda_i = \frac{s_i^2}{N-1}$.

Trained SSM

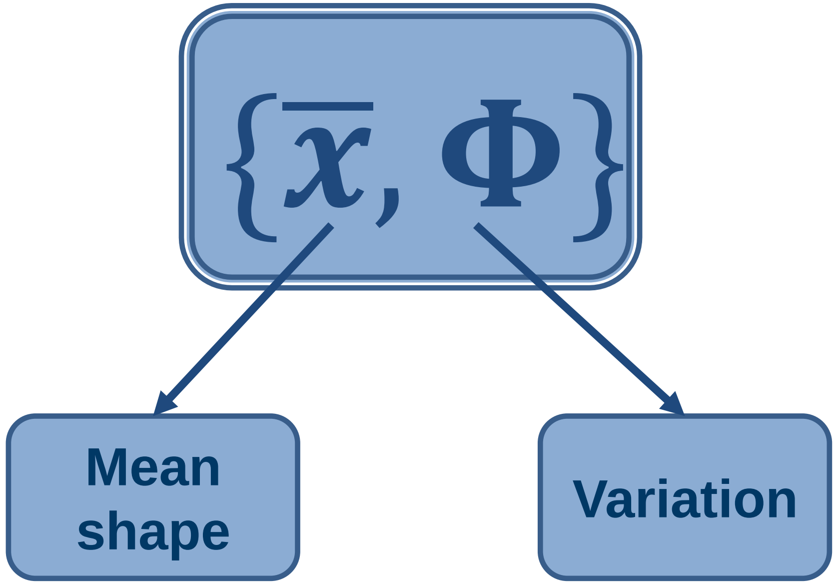

$$\boldsymbol{x}_i = \bar{\boldsymbol{x}} + \boldsymbol{\Phi}\boldsymbol{w}_i + \boldsymbol{\epsilon}$$

$$\boldsymbol{w}_i \in \mathbb{R}^M, \epsilon \in \mathbb{R}^{3K}$$

Shapes can be expressed as a linear combination of eigenvectors given the **feature weights** $\boldsymbol{w}_i$.
$\boldsymbol{\epsilon}$ is the **residual error** due to variance not explained by the model.

SSM: Functionality

reconstructShape (weights **w**)	**reduceDim** (shape **x**)
Reconstruct a shape $\boldsymbol{x}(\boldsymbol{w})$ from the model, given a set of feature weights $\boldsymbol{w} \in \mathbb{R}^{M}$:	A corresponding, scaled and aligned shape $\boldsymbol{x} \in \mathbb{R}^{3K}$ can be projected onto the model basis directly to obtain low-dimensional feature representation:
$\boldsymbol{x}(\boldsymbol{w}) = \overline{\boldsymbol{x}} + \boldsymbol{\Phi}\boldsymbol{w}$	$\boldsymbol{w}_{\boldsymbol{x}} = \boldsymbol{\Phi}^{\mathbf{T}}(\mathbf{x} - \overline{\boldsymbol{x}})$

SSM: Functionality

fitModel (shape **y**)

A scaled and aligned shape $\boldsymbol{y} \in \mathbb{R}^{3L\times 1}$ of different size without known point correspondence requires fitting by iteratively optimizing:

$$\min_{\boldsymbol{w}} ||\boldsymbol{f}\big(\boldsymbol{y}, \boldsymbol{x}(\boldsymbol{w})\big) - \boldsymbol{x}(\boldsymbol{w})||^2$$

with

- $\boldsymbol{x}(\boldsymbol{w}) = \overline{\boldsymbol{x}} + \boldsymbol{\Phi}\boldsymbol{w}$,
- and $\boldsymbol{f}(\boldsymbol{y}, \boldsymbol{x})$ being a matching operator (e.g. k-d tree) which yields a corresponding point in point cloud $\boldsymbol{y}$ to each point in $\boldsymbol{x}$.

Weight Restriction

Weights are usually restricted to only allow **“reasonable” shapes** to be constructed.
Assuming the data follows a normal distribution*, feature weights w_j are bounded within a certain range of the standard deviation $\sqrt{\lambda_j}$.

$$-3\sqrt{\lambda_j} \leq w_j \leq 3\sqrt{\lambda_j},$$
$$j \in \{1, \dots, M\}$$

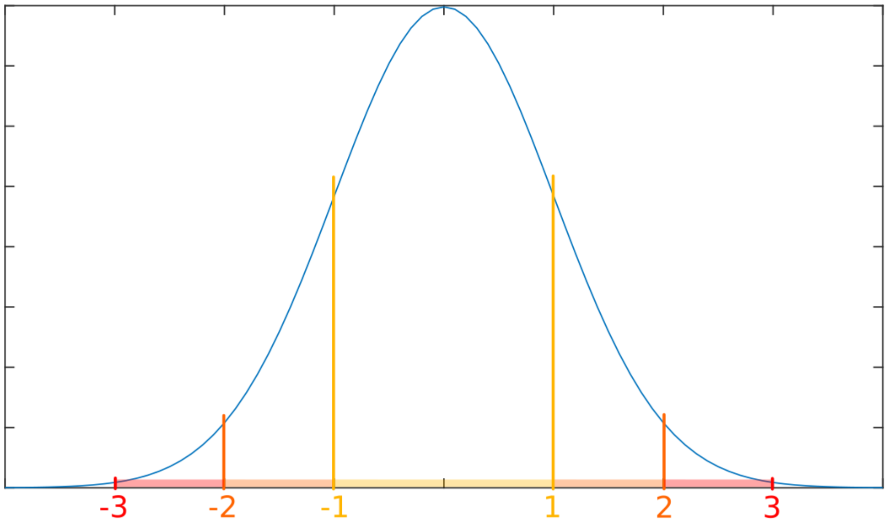

Figure 5: Gaussian bell curve with multiples of the standard deviation

* This, again, illustrates the need for many training data sets.

Main Modes of the Cardiac Model

$\sqrt{\lambda_1}$ $\sqrt{\lambda_2}$ $\sqrt{\lambda_3}$

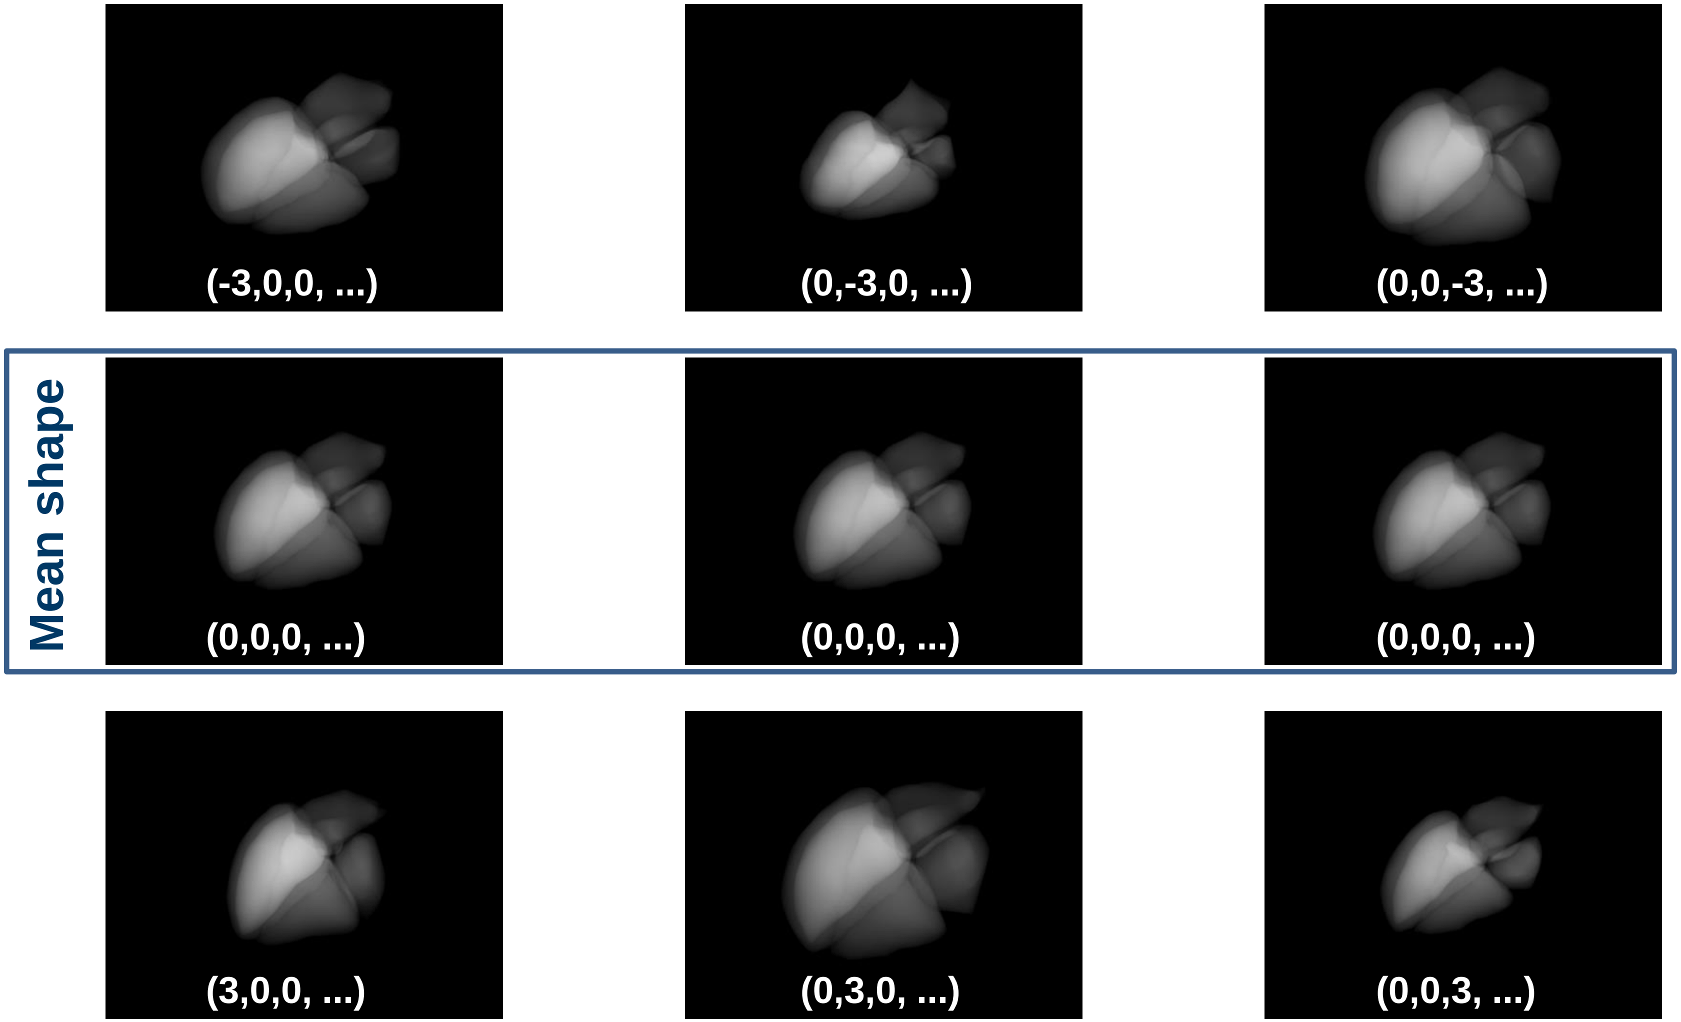

Tutorial Code

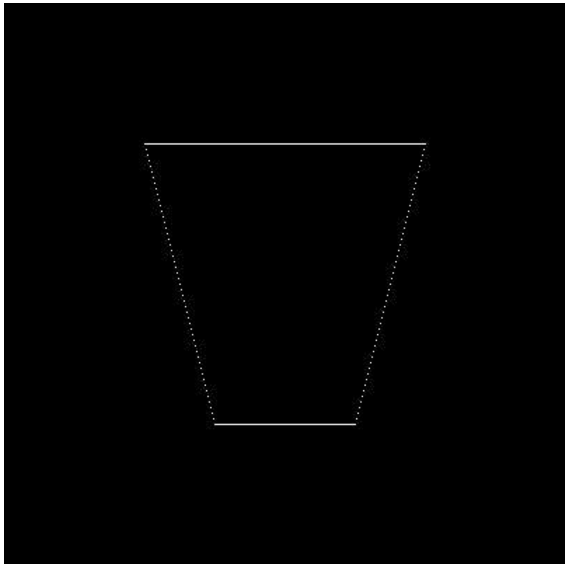

Figure 6: Shape 1

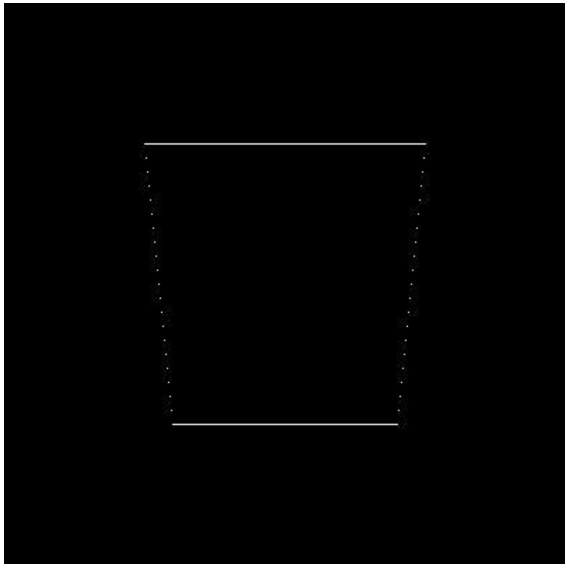

Figure 7: Shape 2

The tutorial code uses shape 1 and shape 2 to create a simple **Statistical Shape Model**. **

** In reality, much larger data sets are needed for training.

Tutorial Code

First, **GPA** is used to create the **mean shape** (**consenus**):

```
GPA gpa = new GPA(2);
gpa.addElement(0, shape1);
gpa.addElement(1, shape2);
gpa.runGPA();
SimpleMatrix consensus = gpa.getScaledAndShiftedConsensus();
```

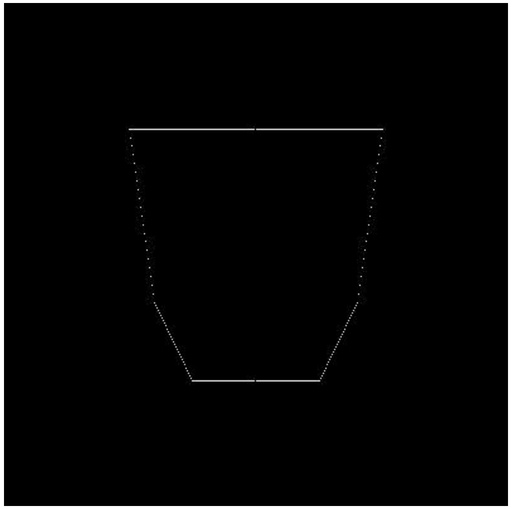

Figure 8: Mean shape

Tutorial Code

Then the **PCA** is run on the data matrix:

```
DataMatrix datam = new DataMatrix(gpa);
PCA pca = new PCA(datam);
pca.run();
```

And the **statistical shape model** is created:

```
ActiveShapeModel asm = new ActiveShapeModel(pca);
```

Tutorial Code

Now, **different weights** can be used to create new shapes:

```
double[] weights = {1, 0};
Mesh mesh1 = asm.getModel(weights);
SimpleMatrix shape3 = mesh1.getPoints();
```

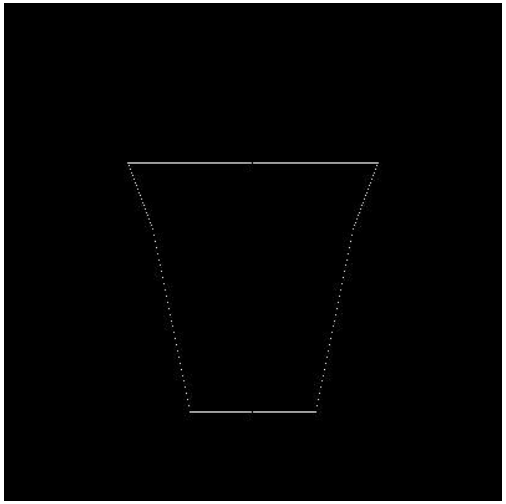

Figure 9: (1, 0)

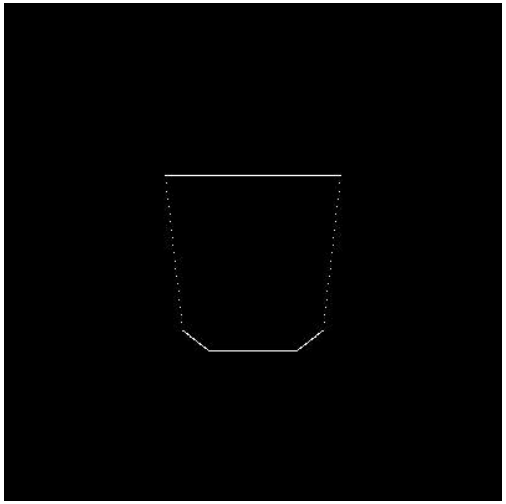

Figure 10: (0, 200)

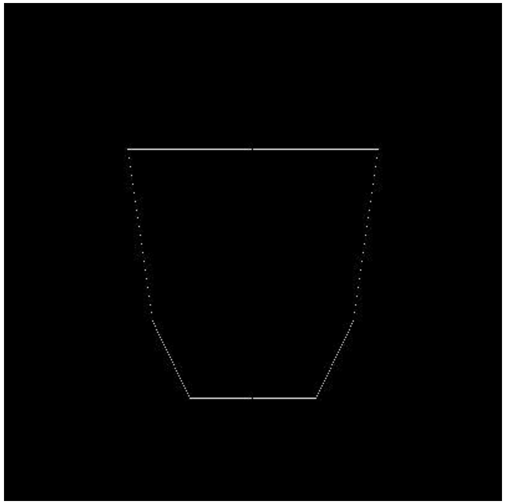

Figure 11: (0, 0)

Limitations of SSMs

- It requires a large population of representative training samples.

- Point correspondences need to be established…
  - …during preprocessing between samples in the training set.
  - …during model fit between mean shape and the new shape.
- ➢ *Many similarities to point set registration*

- Model fit requires initial guess and is susceptible to misalignments.

References and Further Readings

- M. Unberath et al. (2015-July), “Open-Source 4D Statistical Shape Model of the Heart for X-ray Projection Imaging”, Proc. ISBI 2015, pp. 739–742
- T. Cootes (2000), “An Introduction to Active Shape Models”, in Model-based Methods in Analysis of Biomedical Images (Oxford Univ Press) Chap. 7, pp. 223–248
- T. Heimann, H. Meinzer (2009), “Statistical Shape Models for 3D Medical Image Segmentation: A Review”, Medical Image Analysis, vol.13, pp. 543–563
- J. C. Gower et al. (1975), “Generalized Procrustes Analysis”, Psychometrika, vol. **40**, pp. 33–51
- A. Sinha et al. (2017-February), „Simultaneous segmentation and correspondence improvement using statistical modes“, in Medical Imaging: Image Processing, vol. 10133, p. 101331B