02502 Image Analysis Exam Fall 2024

02502 Image Analysis Exam Fall 2024

Written exam, December 12, 2024

Course name: Image Analysis

Course number: 02502

Number of Questions: 25

Aids allowed: All aids allowed.

Duration: 4 hours

Weighting: All questions are equally weighted

Notes: There are five possible answers to each question and a "do not know" option. A correct answer will be equivalent to 5 points. An incorrect answer will be equivalent to -1 points. Questions unanswered (equivalent to "do not know") will not produce points. The final grade is determined by the examiners.

Appendix: Remember to submit your code (Python files, Notebooks, PDF, AI prompts or similar) to the "mellemregninger "/appendix part of the exam.

Data: All data for the exam can be downloaded here (https://designer.mcq.eksamen.dtu.dk/api/images/77f4b631-158c-411a-b472-cc9a567522fa) .

Finding the kidneys in computed tomography scans

As part of a diagnostic system, we need to locate the kidneys in computed tomography scans. We have been provided with an example DICOM slice (1-189.dcm) that we can use to develop our algorithm.

We start by reading the scan slice and create a first binary estimate by applying a lower threshold of 100 and a high threshold of 250. Meaning that pixels with a value that is within these limits are set to foreground (1) and the rest to background (0). This results in a binary image with several structures.

We do a BLOB analysis of the binary image and computes the area and perimeter of all BLOBS. The kidneys are identified by filtering the BLOBS using the area and the perimeter.

After the filtering, where only the two kidneys are left a morphological closing with a disk-shaped structuring element (radius=3) is performed on the binary image.

Start by answering the question about the BLOB area estimation.

All data for the exam can be downloaded here (https://designer.mcq.eksamen.dtu.dk/api/images/77f4b631-158c-411a-b472-cc9a567522fa) .

When all BLOBs are found and the area and the perimeter are computed for all of them. We set the minimum perimeter to 400 and the maximum perimeter to 600. The maximum BLOB area is set to 5000. What should the minimum BLOB area be set to, so only the kidneys are found?

Vælg en svarmulighed

- 3000
- 500
- Do not know
- 4000
- 1000
- 2000

An expert has drawn an approximate outline of the kidneys.

You compare your segmentation (after the morphological closing) with the expert segmentation using the DICE score. What is the DICE score?

Vælg en svarmulighed

- ◯ Between 0.75 and 0.80
- ◯ Between 0.90 and 0.95
- ◯ Between 0.95 and 1.00
- ◯ Between 0.80 and 0.85
- ◯ Between 0.85 and 0.90
- ◯ Do not know

One pixel in the image corresponds to a physical size of 0.78 mm x 0.78 mm.

When the kidneys have been isolated and the morphological closing is performed, the total area of the kidneys is computed. It is:

Vælg en svarmulighed

- ◯ Between 25 and 35 cm^2
- ◯ Between 35 and 45 cm^2
- ◯ Between 15 and 25 cm^2
- ◯ Between 55 and 65 cm^2
- ◯ Between 45 and 55 cm^2
- ◯ Do not know

When the kidneys have been isolated and the morphological closing is performed, the Hounsfield units in the original image under the segmentation mask is sampled. What is the median HU value?

Vælg en svarmulighed

- ◯ Between 120 and 130
- ◯ Between 130 and 140
- ◯ Between 140 and 150
- ◯ Between 150 and 160
- ◯ Between 160 and 170
- ◯ Do not know

Traffic jam analysis

Your company has been awarded a tender by the Danish Ministry of Transportation to analyze traffic jams on the M3 motorway. Traffic cameras near the exit for the Technical University of Denmark have recorded unusual afternoon traffic patterns, with heavy congestion in the mornings but not in the afternoons.

The Ministry has provided a dataset for analysis. The files contains three features: **density** (number of cars), **speed** (Km/h), and **weather** (labeled as "0" for a blue sky and "1" for rain).

In the file, **traffic_train.txt** the first 140 rows are from morning traffic and the remaining 140 rows are from afternoon traffic.
In the file, **traffic_test.txt** the first 60 rows are from morning traffic and the remaining 60 rows are from afternoon traffic.

All data for the exam can be downloaded here (https://designer.mcq.eksamen.dtu.dk/api/images/77f4b631-158c-411a-b472-cc9a567522fa) .

The dataset is divided into two classes and you visualize the first 100 training samples for each class:

Class 1: Morning traffic samples (represented by green dots).

Class 2: Afternoon traffic samples (represented by blue dots).

Which of the figures below matches your scatter plot of the **density** and **speed** features?

Vælg en svarmulighed

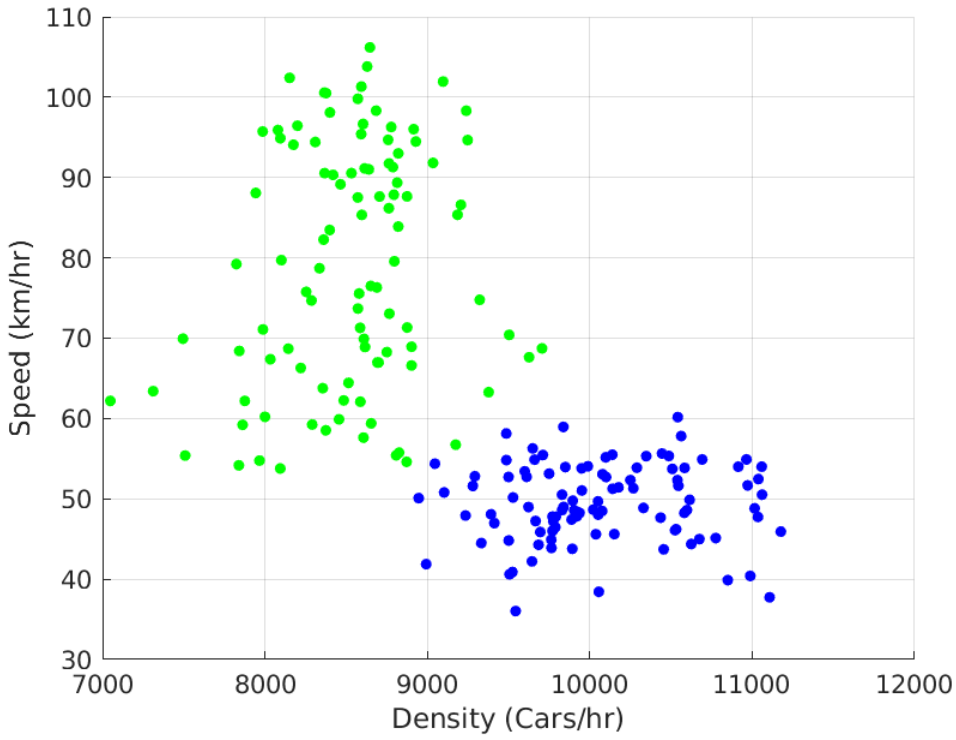

○

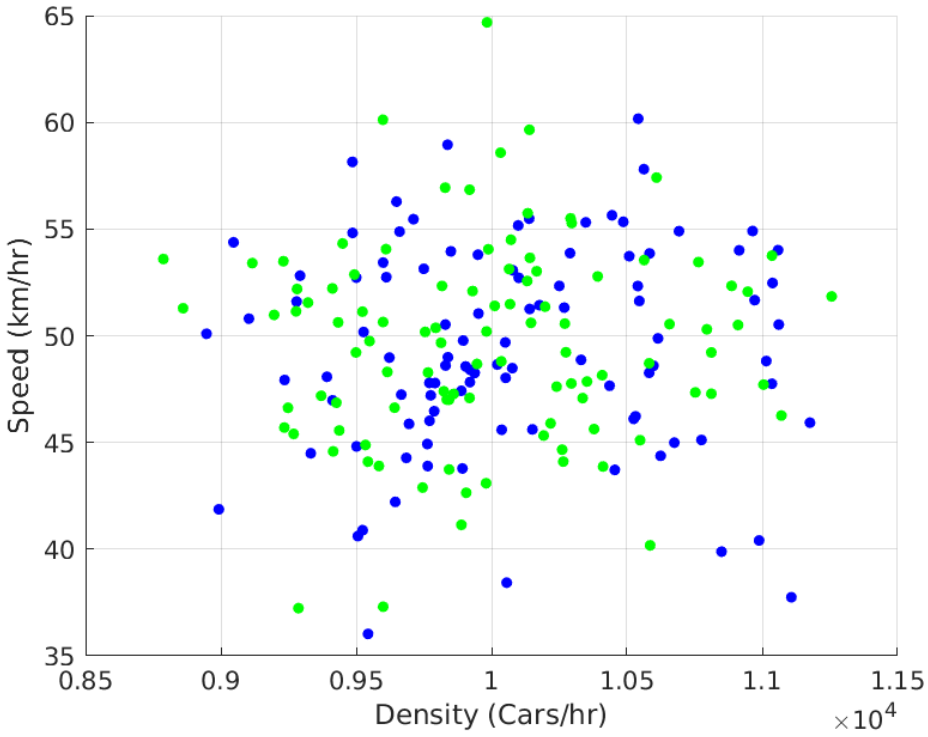
65
60
55
50
45
40
35
Speed (km/hr)
0.85
0.9
0.95
1
1.05
1.1
1.15
Density (Cars/hr)
×10^4

○

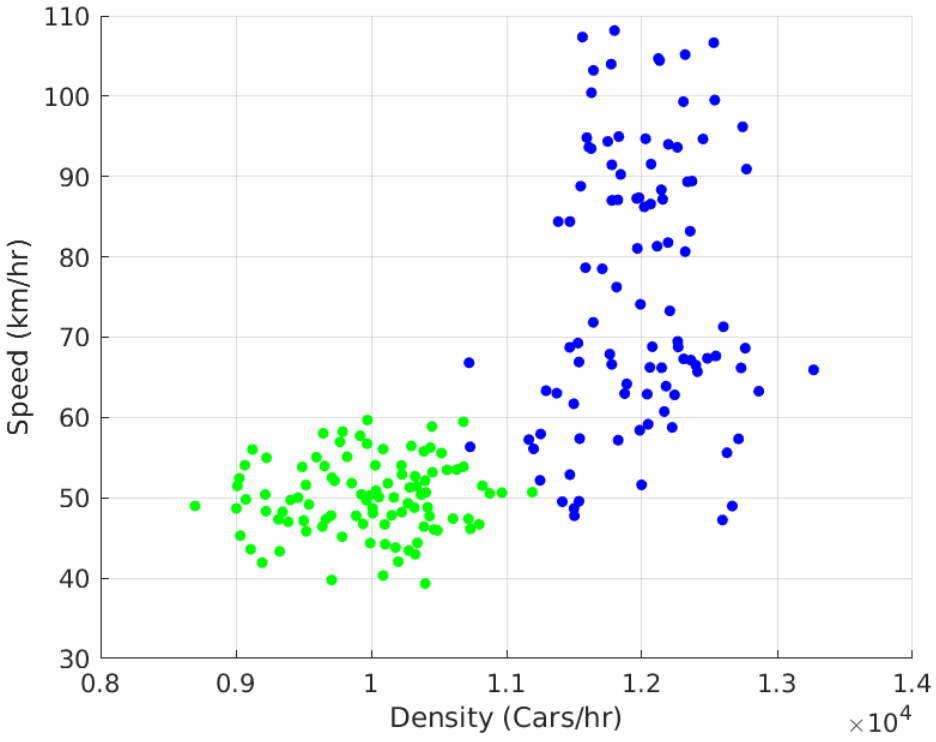
110
100
90
80
70
60
50
40
30
Speed (km/hr)
0.8
0.9
1
1.1
1.2
1.3
1.4
Density (Cars/hr)
×10^4

○

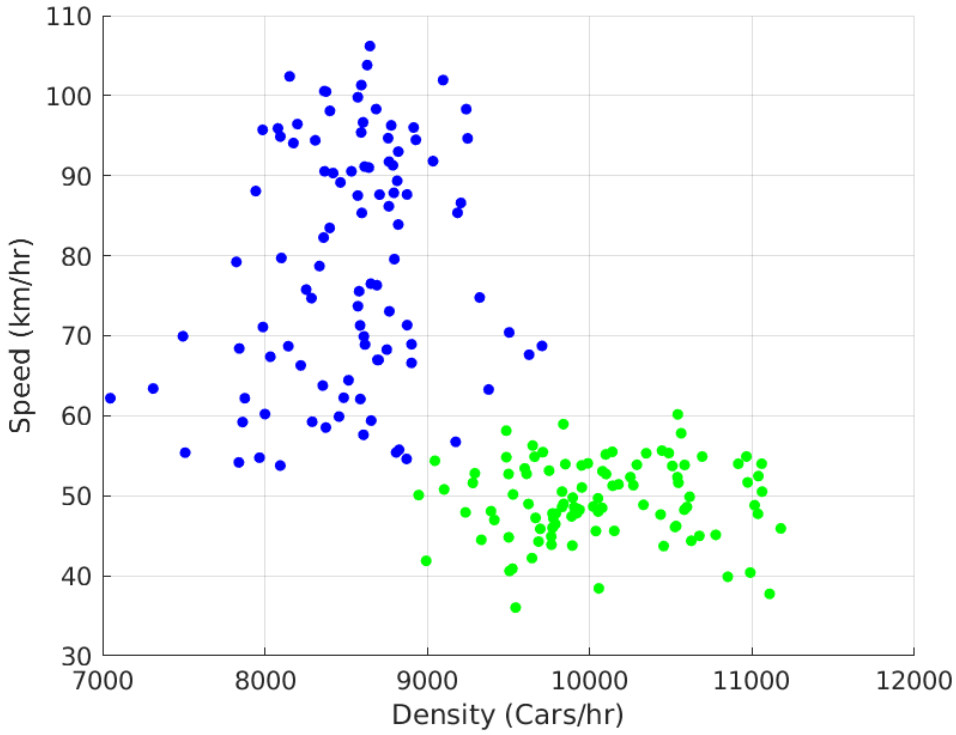
110
100
90
80
70
60
50
40
30
Speed (km/hr)
7000
8000
9000
10000
11000
12000
Density (Cars/hr)

○

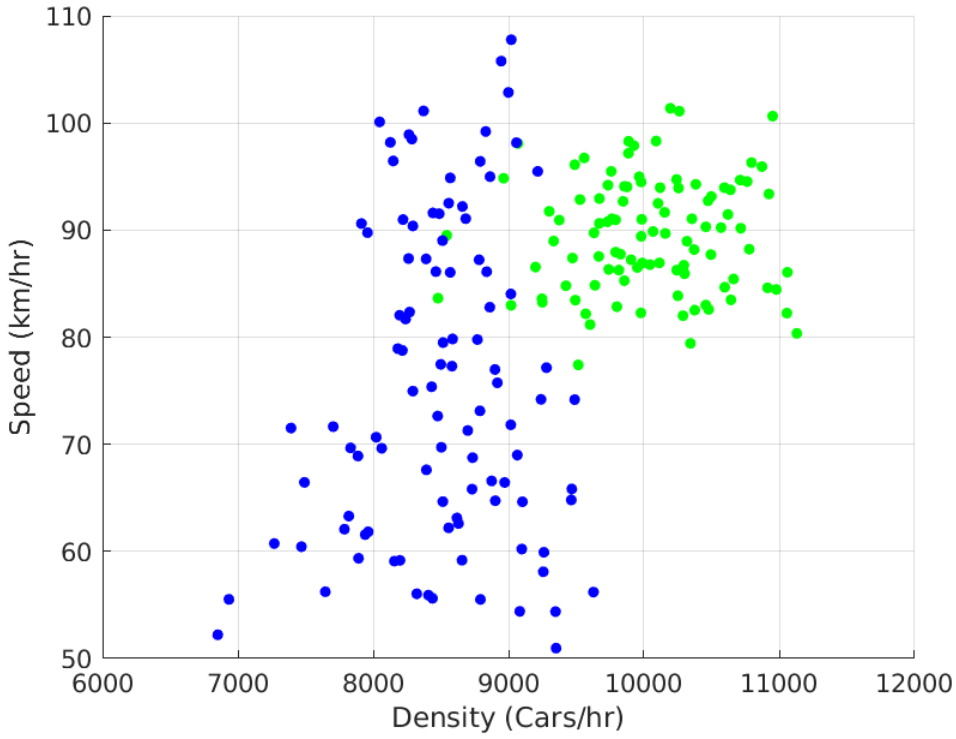
110
100
90
80
70
60
50
Speed (km/hr)
6000
7000
8000
9000
10000
11000
12000
Density (Cars/hr)

○ Do not know

To detect traffic jams, you use a Fisher's Linear Discriminant Analysis (LDA) model. The training set consists of the the samples in traffic_train.txt. The model is trained using two features: density and speed.

The trained LDA model is then applied to the test dataset traffic_test.txt.

How many afternoon samples were classified as morning samples, indicating slowed, high-density traffic (i.e., traffic jams)?

Vælg en svarmulighed

- ◯ Between 15 and 20
- ◯ Between 25 and 30
- ◯ Between 5 and 10
- ◯ Between 0 and 5
- ◯ Do not know
- ◯ Between 10 and 15

To get an impression of the weathers influence on the traffic you count the number of days where there was rain in the morning in the training set. How many days?

Vælg en svarmulighed

- ◯ Between 55 and 65
- ◯ Between 35 and 45
- ◯ Do not know
- ◯ Between 25 and 35
- ◯ Between 45 and 55
- ◯ Between 65 and 75

You are asked to design a system that can track skiers on a ski slope. Due to the weather, you can expect that the lightning conditions slowly changes during the slope opening times. What elements do you need in your image analysis system?

Vælg en svarmulighed

- ◯ Estimation of a slowly changing background image, computing difference between current video frame and background, detection of significant changes in the image, analysis of moving objects.
- ◯ Conversion of image from RGB to grey scale, parametric classification of pixel values, image registration to a template figure
- ◯ Dynamic programming to find slope, image filtering to sort birds from the image, fast facial recognition to find skiers
- ◯ Do not know
- ◯ Image registration of slopes to a template slope, LDA classification of registration results, morphological closing of outlines
- ◯ Morphological opening of binary image from camera, median filtering of binary images, image rotation to get slope in the right camera view

Sorting screws and bolts using machine vision

Your friend has renovated his apartment and all the screws and bolts is now all messed up. She calls for your help to use machine vision to sort them. She also provides 20 example images of the screws and bolts that is in her box.

Perhaps principal component analysis (PCA) can be a part of the solution. Therefore you start by computing an average photo of the collection and then use PCA from **sklearn.decomposition** to compute the 7 principal components from the photos. You also project all photos to the PCA spaces, so you have their positions in PCA space.

All data for the exam can be downloaded here (https://designer.mcq.eksamen.dtu.dk/api/images/77f4b631-158c-411a-b472-cc9a567522fa) .

How many PCA components are needed to explain at least 44% of the total variation in the set of photos?

Vælg en svarmulighed

- ◯ Do not know
- ◯ 4
- ◯ 3
- ◯ 5
- ◯ 2
- ◯ 1

It would be nice to group the screws and bolts together in boxes. Two do that it is necessary to identify the ones that are most similar. First you try to find the two photos that are most similar in PCA space, meaning the two photos that have the smallest distance to each other in PCA space. They are:

Vælg en svarmulighed

- ◯ screws_000. jpg and screws_005.jpg
- ◯ screws_011. jpg and screws_013.jpg
- ◯ screws_012. jpg and screws_016.jpg
- ◯ screws_006. jpg and screws_009.jpg
- ◯ screws_010. jpg and screws_017.jpg
- ◯ Do not know

To get an idea of how mixed up the screws are you examine the positions on the first principal component and show the two photos that have the smallest and the largest value. Select the corresponding image.

Vælg en svarmulighed

- ◯ Do not know
- ◯ 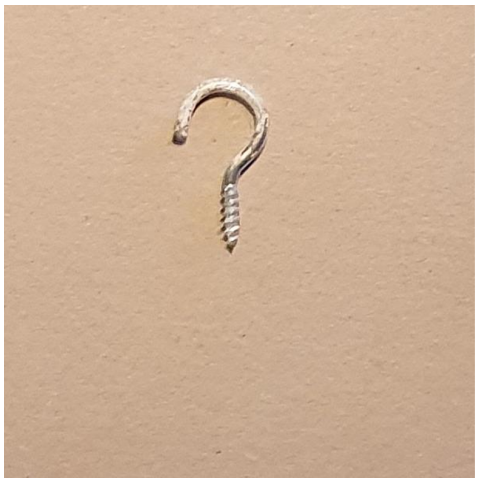 
- ◯  
- ◯  

○ 

○ 

To get an idea of the distances in PCA space, you select two photos screws_007.jpg and screws_008.jpg and compute the Euclidean distance between them in PCA space by using all the PCA coordinates. What is this distance?

Vælg en svarmulighed

- ◯ Do not know
- ◯ Between 0 and 5000
- ◯ Between 5000 and 10000
- ◯ Between 10000 and 15000
- ◯ Between 15000 and 20000
- ◯ Between 20000 and 25000

You recognize the screw in the photo **screws_007.jpg** as a plaster wall screw and you are curious about the position in PCA space. You plot all screws positions on the first two principal components and mark the position of **screws_007.jpg**. Mark the best matching plot (it might be flipped left and right and/or up and down)

Vælg en svarmulighed

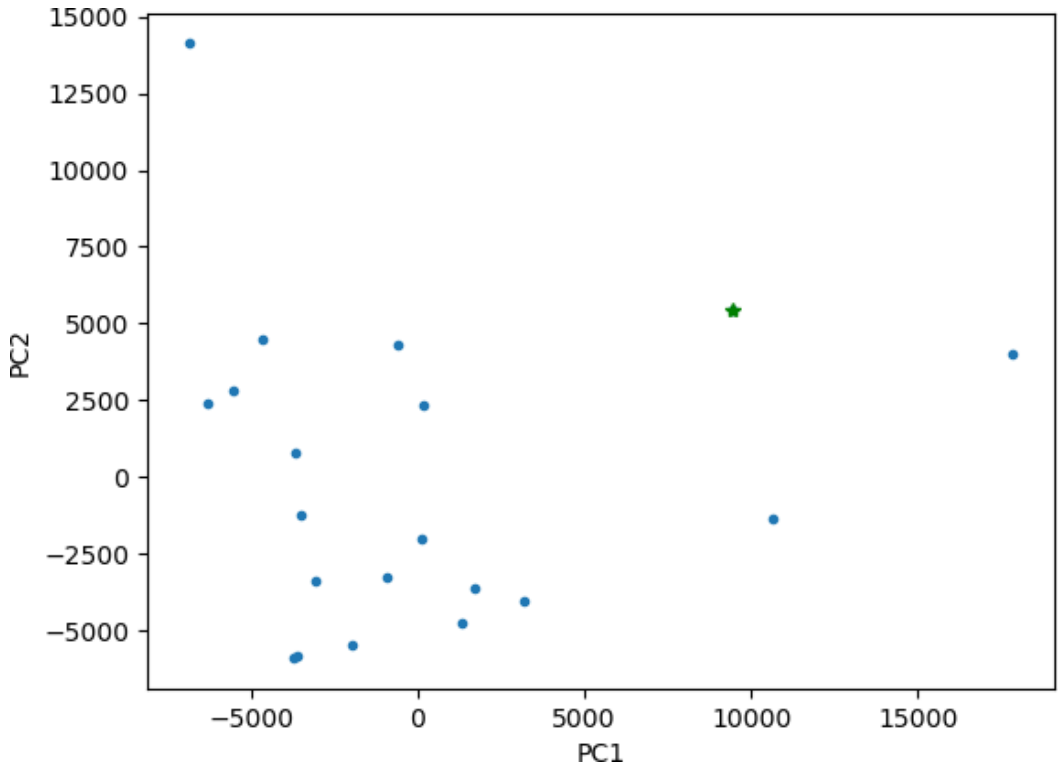

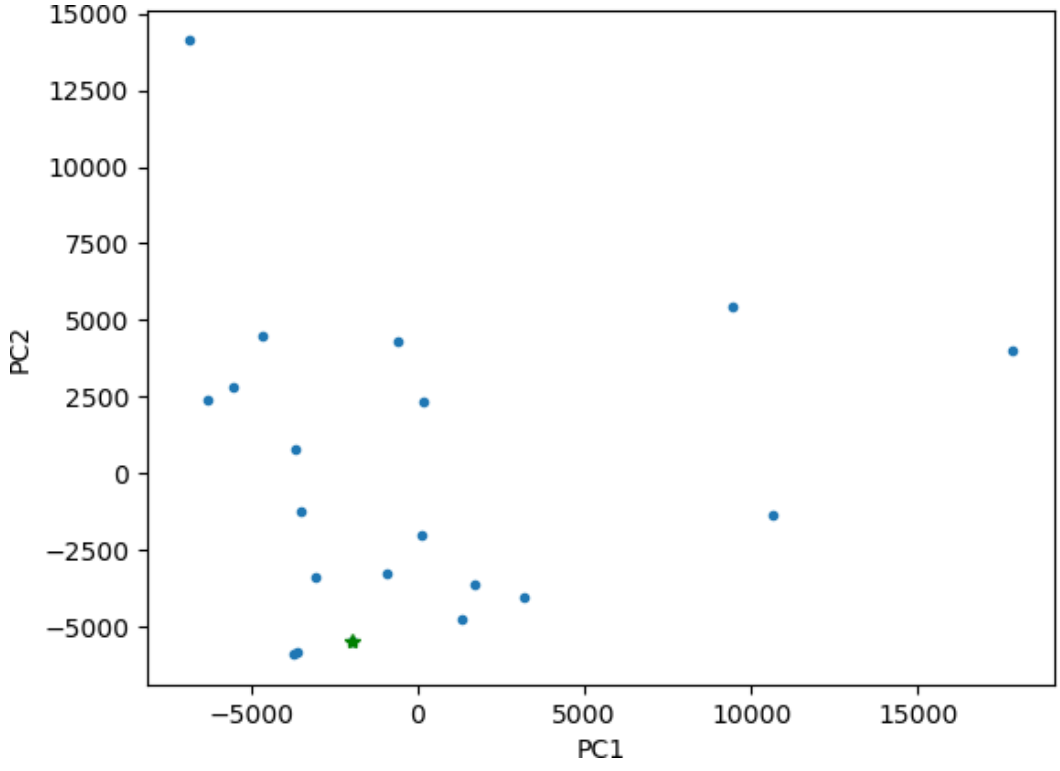

Do not know

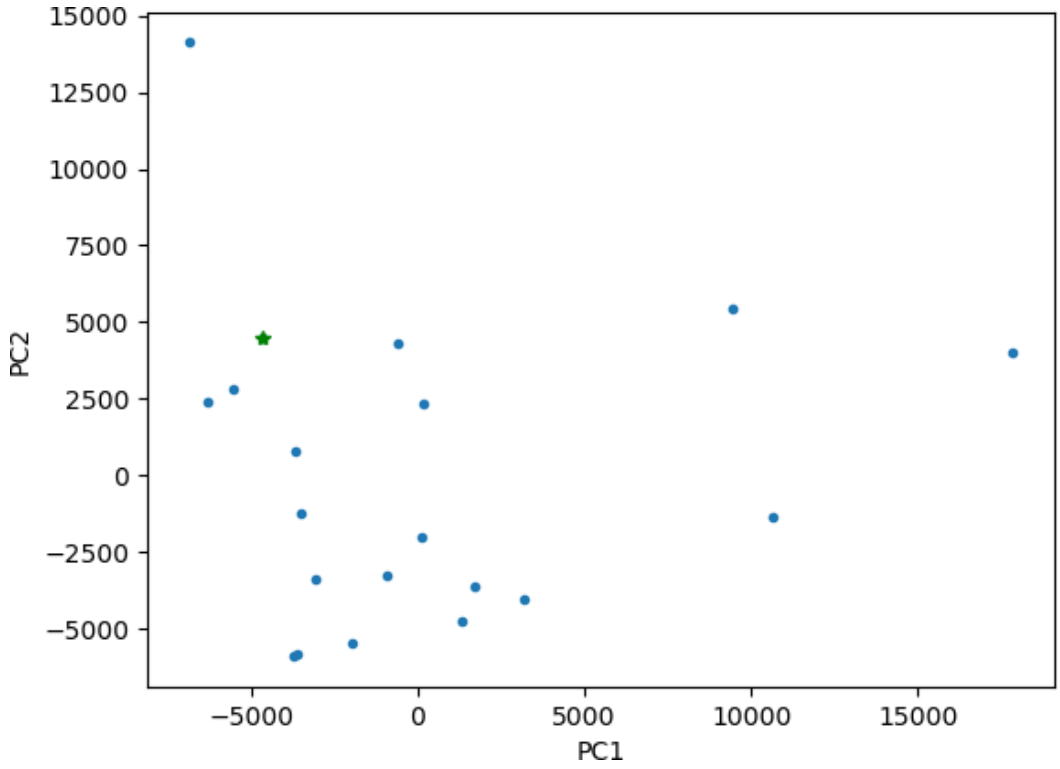
15000
12500
10000
7500
5000
2500
0
−2500
−5000
PC2
−5000
0
5000
10000
15000
PC1

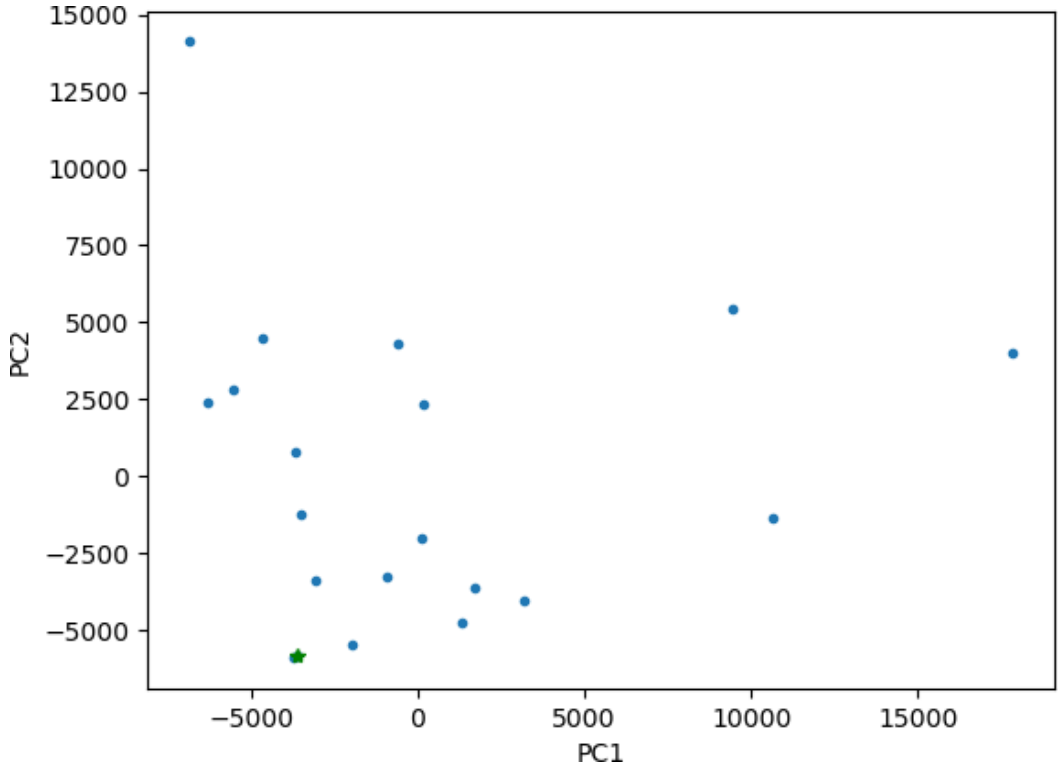
15000
12500
10000
7500
5000
2500
0
−2500
−5000
PC2
−5000
0
5000
10000
15000
PC1

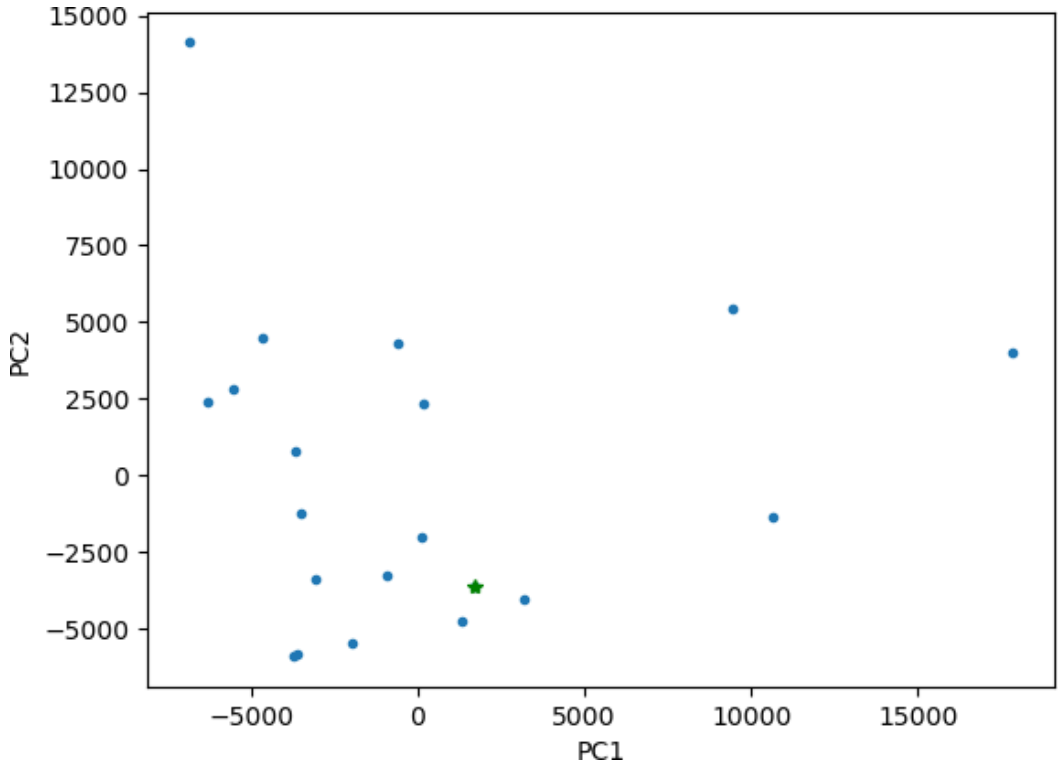
15000
12500
10000
7500
5000
2500
0
−2500
−5000
PC2
−5000
0
5000
10000
15000
PC1

Gradient descent

You must find the optimal solution to a registration problem that spans a 2D parameter space (x_1,x_2) . The cost function is defined as:

$$c = 7 * x_1 * x_1 + x_1 * x_2 + 3 * x_2 * x_2$$

Since there is no analytic solution, you have selected a numerical optimization approach based on the iterative gradient descent method. With the step size set to 0.1 and the starting position indicated by the green dot shown in the figure.

Which of the following results matches your iterative optimization process in the 2D parameter space?

Vælg en svarmulighed

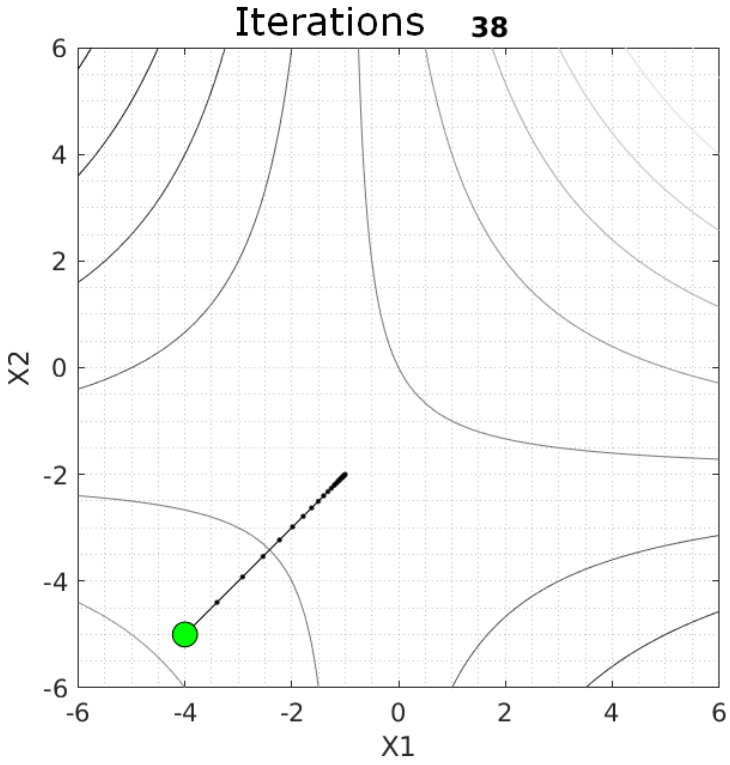

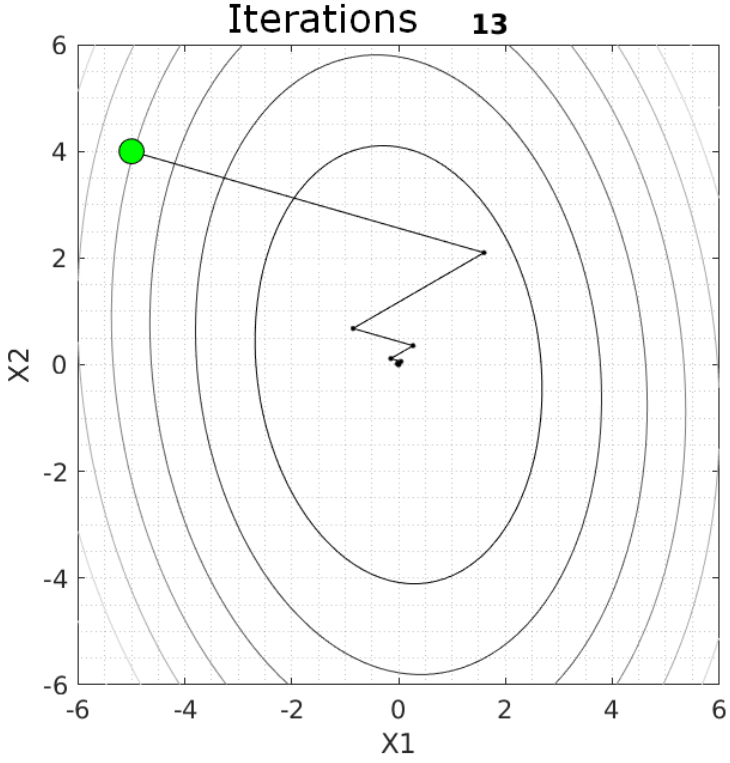

Do not know

○

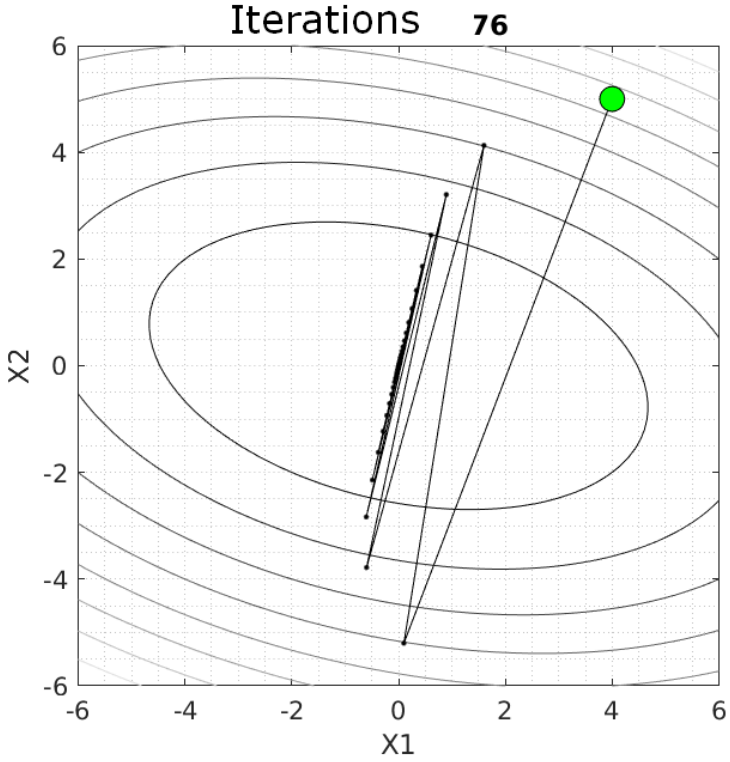
Iterations 76
X2
X1
6
4
2
0
-2
-4
-6
-6
-4
-2
0
2
4
6

○

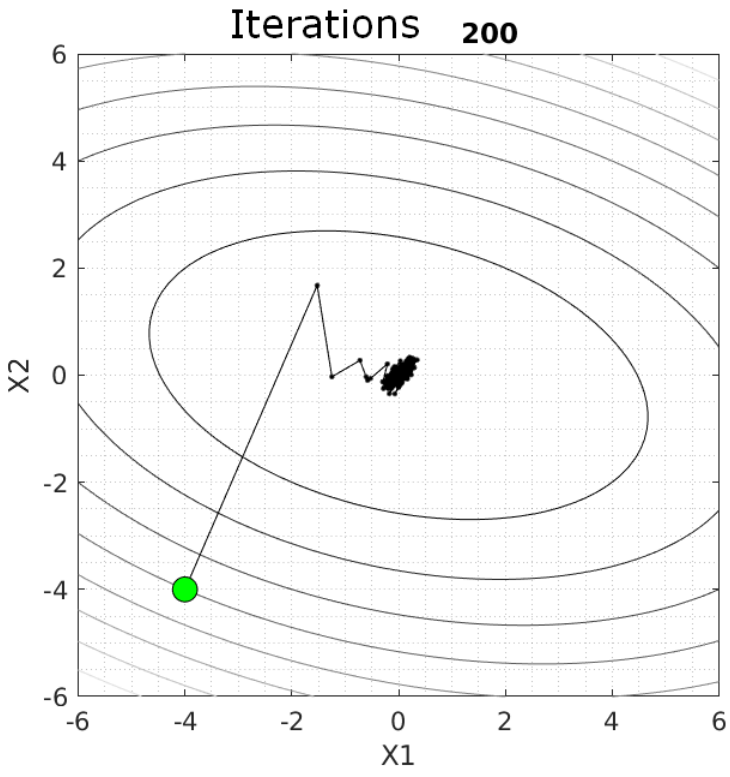
Iterations 200
X2
X1
6
4
2
0
-2
-4
-6
-6
-4
-2
0
2
4
6

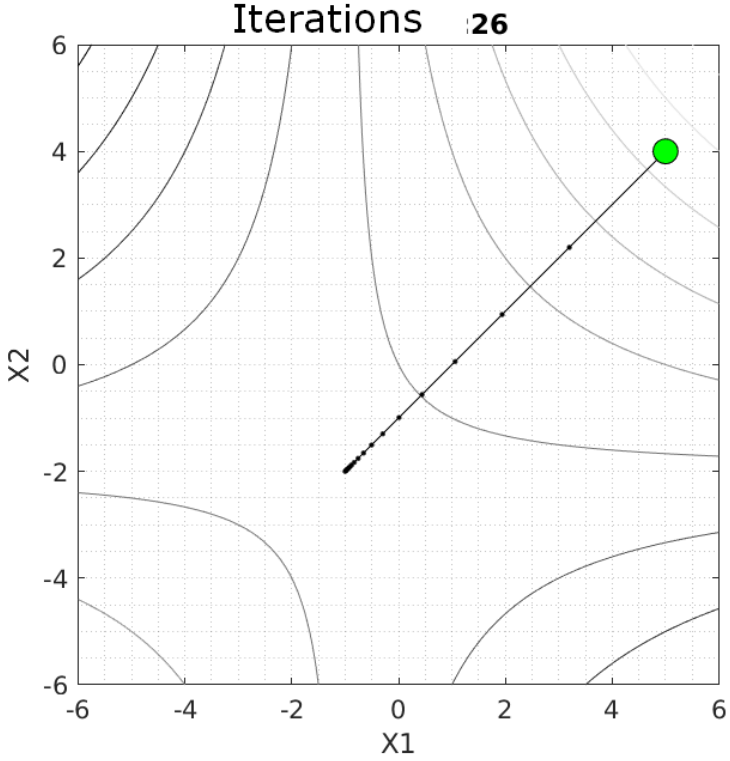
Iterations 26
X2
X1
6
4
2
0
-2
-4
-6
-6
-4
-2
0
2
4
6

You monitor the value of the cost function, c. How many iterations is needed before c is below 2.0 ?

Vælg en svarmulighed

- 1
- do now know
- 2
- 4
- 3
- 5

Cell analysis

A laboratory has collected two grey-scale images of tissue samples showing a Nissl stain of cell nuclei appearing as bright spots in the images. The laboratory has asked you for help to do different types of analysis of their cell images.

There are two images:
x_NisslStain_9-260.81.png
y_NisslStain_9-260.81.png

An expert has manually identified five corresponding landmarks in the two images. This provides us with two images where the landmarks are pixels with values 1 to 5. The rest of the pixels are 0.
You should find the (x, y) positions of the landmarks in the two images. The result should be two lists of corresponding landmarks. The two images are called:

LabelsFixedImg.png
LabelsMovingImg.png

All data for the exam can be downloaded here (https://designer.mcq.eksamen.dtu.dk/api/images/77f4b631-158c-411a-b472-cc9a567522fa) .

You have decided first applying a few pre-processing steps to identify cell nuclei as individual binary objects. You do this by a BLOB analysis on the two images:

x_NisslStain_9-260.81.png
y_NisslStain_9-260.81.png

The images can be made into binary images by setting of threshold of 30 and cleaned by applying a morphological opening with a disk-sized structuring element of size 3. How many individual cells do you find in the two images after the BLOB analysis (without removing BLOBS touching the image border).

Vælg en svarmulighed

- ◯ Between 0 and 5 cells per image
- ◯ Between 10 and 15 cells per image
- ◯ Do not know
- ◯ Between 5 and 10 cells per image
- ◯ Between 15 and 20 cells per image
- ◯ Between 20 and 25 cells per image

After finding the five landmarks, you compute the average landmark for both images and finally the Euclidean distance between the two average landmarks. What is this distance?

Vælg en svarmulighed

- ◯ Betwen 0 and 2
- ◯ Between 2 and 4
- ◯ Between 8 and 10
- ◯ Between 4 and 6
- ◯ Do not know
- ◯ Between 6 and 8

You have developed a fully automatic system for finding cells in image. The system is tested on an independent test image, where an expert has marked the true cells. After evaluation of the results, it is found that:

- 18 cells are correctly found by the system
- the system finds 7 objects that are not marked as cells by the expert
- the system fails to find X cells that the expert has marked as cells
- the system has a sensitivity of 0.82

What is X ?

Vælg en svarmulighed

- ◯ Do not know
- ◯ Between 8 and 10
- ◯ Between 0 and 2
- ◯ Between 2 and 5
- ◯ Between 10 and 13
- ◯ Between 5 and 8

Game dice

A gaming company wishes to segment the eyes from a grayscale image depicting three overlapping dice. The original image is ***CubesG.png***.

An expert has manually drawn five regions of interest (ROIs) to train segmentation models and has stored the pixel values for each ROI in separate text files.

The filenames are as follows:

ROI A: *A_Cubes.txt*
ROI B: *B_Cubes.txt*
ROI C: *C_Cubes.txt*
ROI D: *D_Cubes.txt*
ROI E: *E_Cubes.txt*

All data for the exam can be downloaded here (https://designer.mcq.eksamen.dtu.dk/api/images/77f4b631-158c-411a-b472-cc9a567522fa) .

A minimum distance classifier is trained with three classes. The class ranges are computed using three region-of-interests from the training data. After the class ranges has been computed, the segmentation is performed on **CubesG.png.** The pixels with value 0 (background) are kept as background. You visualize the segmentation and it looks like the picture below:

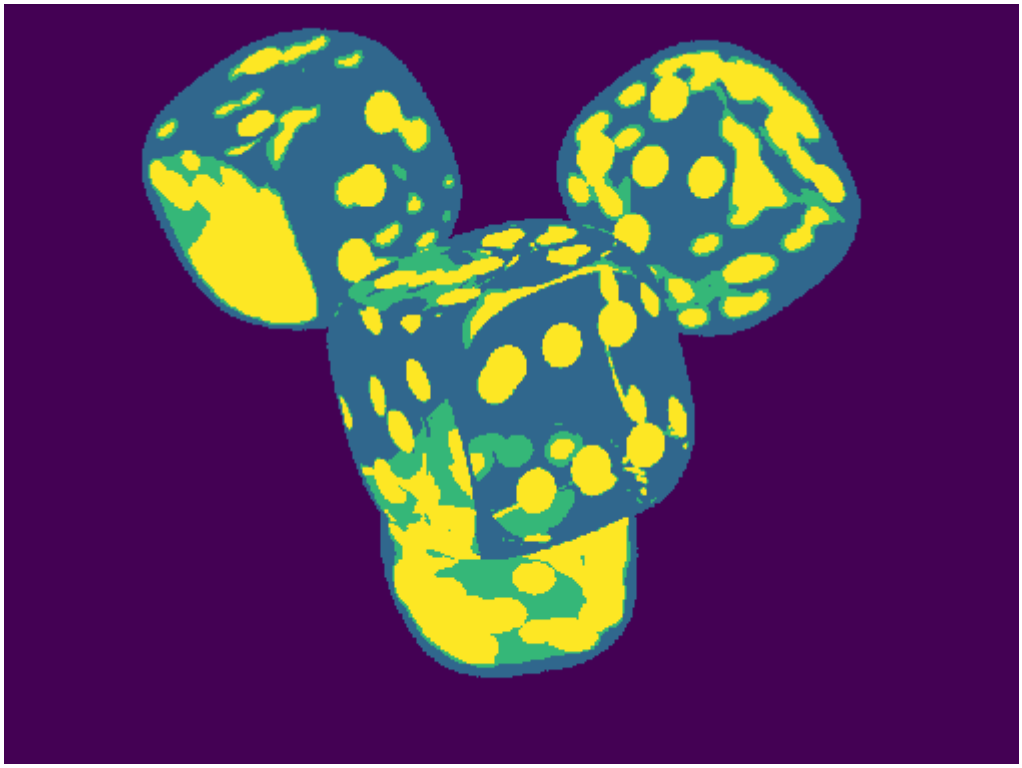

Which combinations of expert annotated ROI values were used?

Vælg en svarmulighed

- A, B and C
- Do not know
- A, D and E
- C, D and E
- B, C and D
- B, D and E

A two class parametric classifier is trained using the values from ROI D and ROI E. What is the optimal threshold between the two classes?

Vælg en svarmulighed

- Between 145 and 155
- Do not know
- Between 175 and 185
- Between 195 and 205
- Between 165 and 175
- Between 155 and 165

Analysis of breast cancer data

To be able to use image analysis to diagnose breast cancer, a sample dataset has been created, where images of fine structures of breast tissues are provided together with information about the patients cancer status. There are samples from patients with without cancer (positive=1) and with cancer (negative=0).

For each patient a set of image features are computed and used in the further analysis.

Since the image features are probably correlated a principal component analysis is performed to reduce the dimensions of the feature space.

The data set can be loaded by first importing:

```
from sklearn.datasets import load_breast_cancer
```

and then by:

```
breast = load_breast_cancer()
x = breast.data
target = breast.target
```

here **x** contains the features per observation and **target** contains the indicator for with cancer (negative) (0) and without (positive) (1).

First the data should be scaled by subtracting the mean value from each measurements and then dividing by the standard deviation of each measurement.

The PCA should be computed using **np.linalg.eig.**

After the PCA, the scaled measurements are projected to PCA space for further analysis.

A PCA based classifier is tested by setting all samples that are projected to a negative value on the first principal component to **positive** (without cancer) and the rest to **negative** (with cancer).
How many samples are classified as positive (without cancer)?

Vælg en svarmulighed

- ◯ between 370 and 38
- ◯ between 100 and 120
- ◯ Do not know
- ◯ Between 305 and 320
- ◯ between 345 and 355
- ◯ between 250 and 260

A PCA based classifier is tested by setting all samples that are projected to a negative value on the first principal component to **positive** (without cancer) and the rest to **negative** (with cancer).
What is the accuracy of this classifier?

Vælg en svarmulighed

- between 0.90 and 1.0
- between 0.80 and 0.90
- do not know
- between 0.70 and 0.80
- between 0.60 and 0.70
- between 0.50 and 0.60

To better understand the data in PCA, the average values of the projection on the first principal component are computed for the negative (patients with cancer) and the positive (patients without cancer) samples. The two average values are:

Vælg en svarmulighed

- -2.21 and 3.71
- -3.23 and 4.23
- Do not know
- -4.98 and 4.87
- -5.43 and 1.87
- -1.23 and 2.67

The measurements projected to the first two principal components are plotted using **plt.scatter** where the color reflects if they are positive or negative.
The plot looks like:

Vælg en svarmulighed

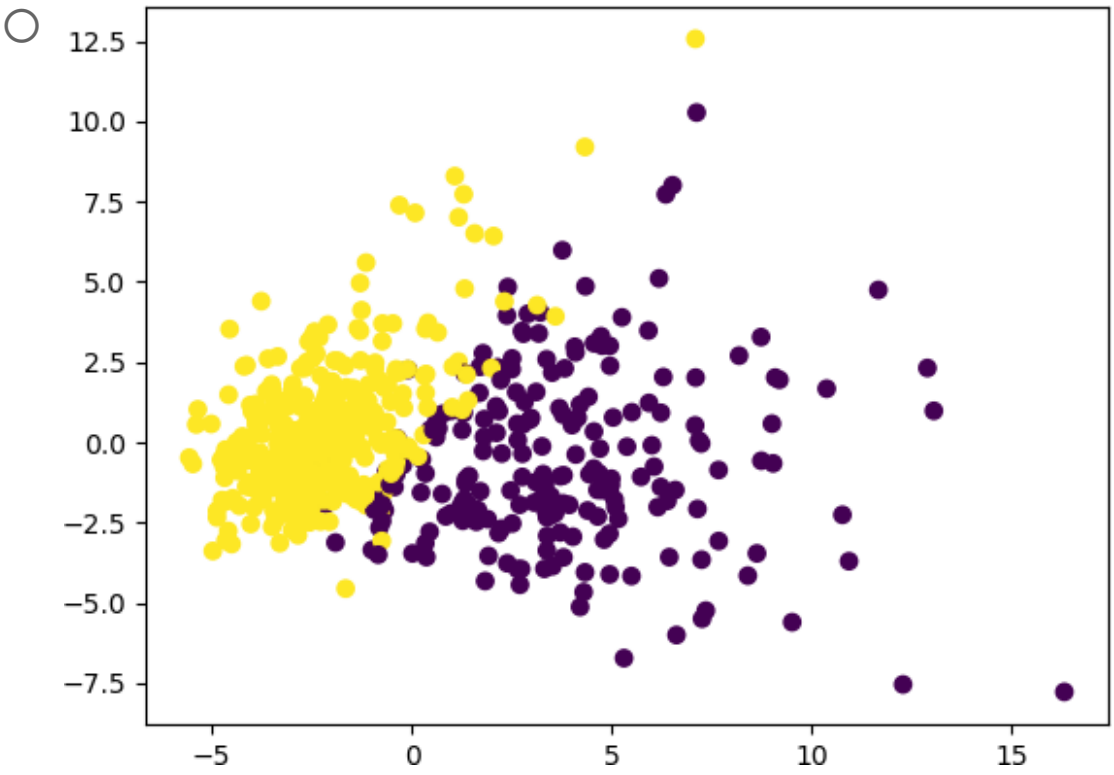

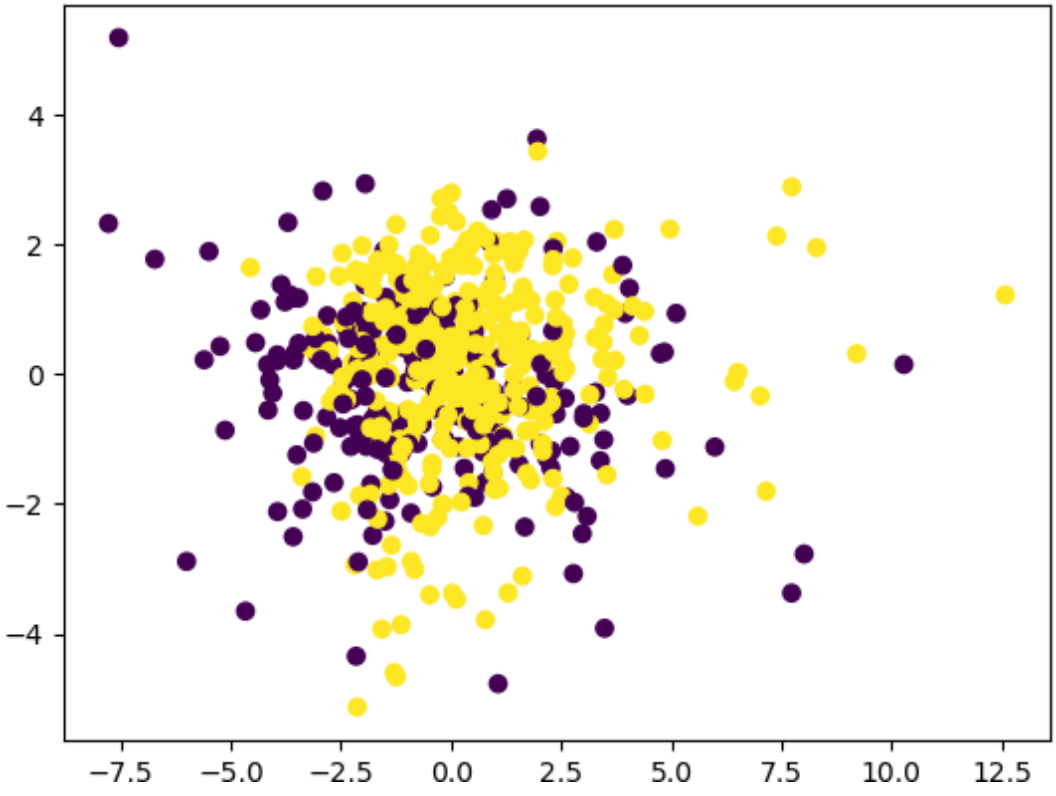

○

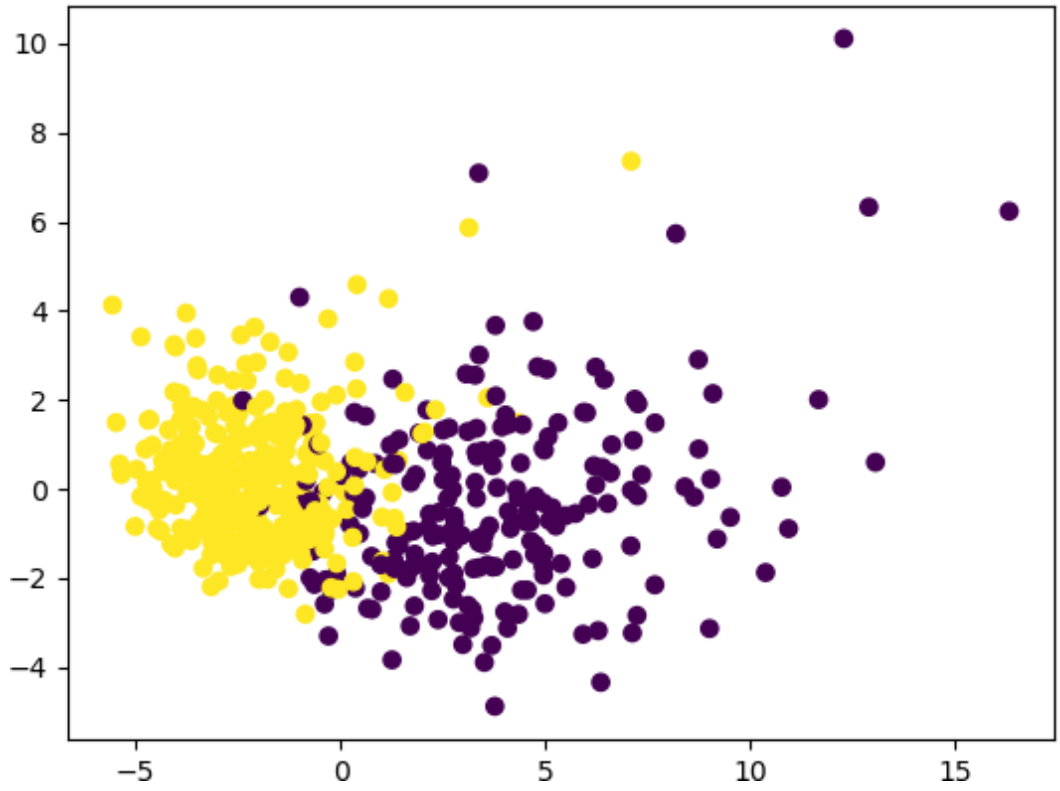

○

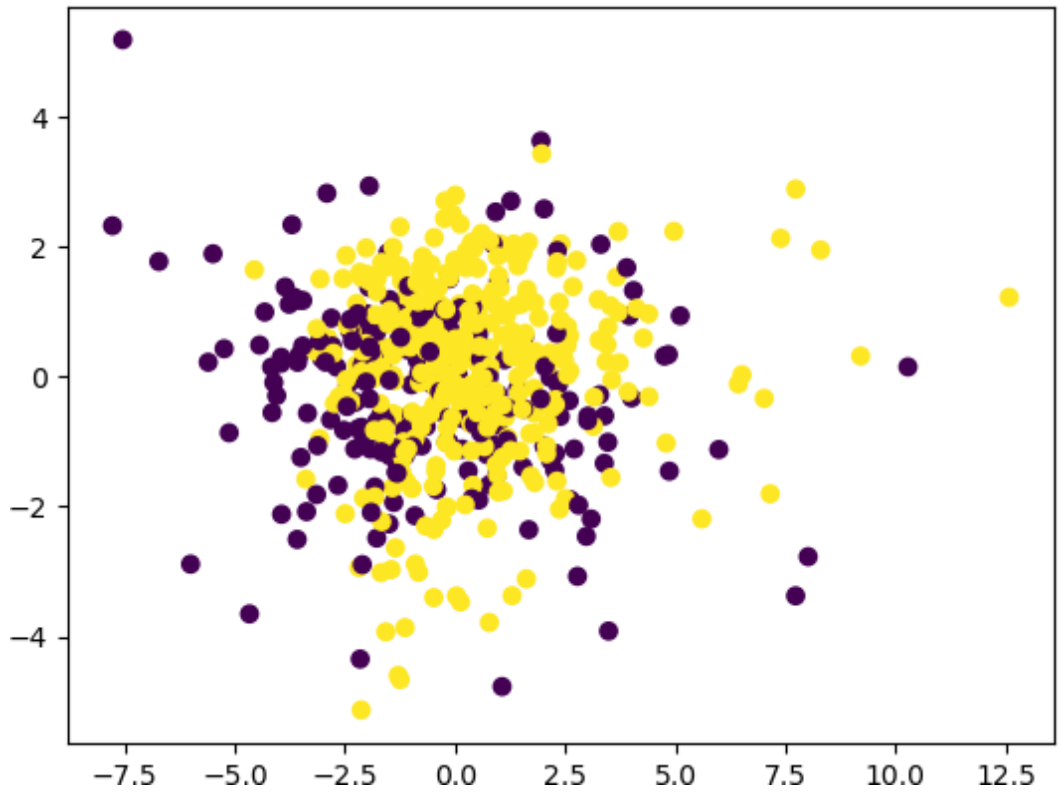

○ Do not know

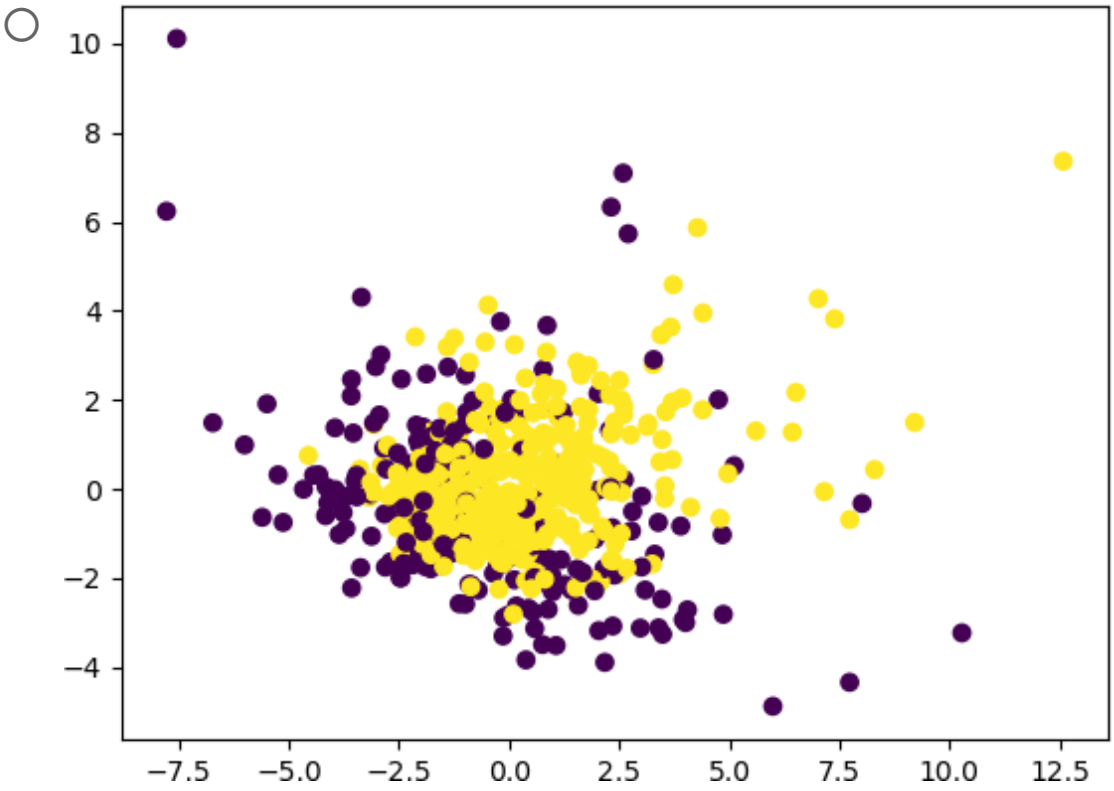
10
8
6
4
2
0
−2
−4
−7.5
−5.0
−2.5
0.0
2.5
5.0
7.5
10.0
12.5

To get a better understanding of the data, the number of features and the number of observations are computed. They are:

Vælg en svarmulighed

- Features=25, observations = 423
- Features=30, observations=569
- Features=41, observations=714
- Do not know
- Features=17, observations=265
- Features=47, observations=435