

# 02502 Image Analysis Exam Fall 2022

Technical University of Denmark

Written exam, December 12, 2022

**Course name:** Image Analysis

**Course number:** 02502

**Number of Questions:** 26

**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.

## Abdominal analysis

There is a growing interest in the functions of the kidneys and how the shape of the kidneys relates to diseases. You have been asked to help analyzing kidneys as seen on computed tomography (CT) scans of the human abdominal area.

You have been given a slice of a CT scan as a DICOM file (1-162.dcm) and expert annotations of background (BackROI.png, LiverROI.png, KidneyROI.png and the AortaROI.png).

Your first algorithm works like this:

1. Reads the DICOM image
2. Uses the annotated regions to extract pixel values in the liver, kidney and aorta regions
3. Determines two threshold using minimum distance classification. Threshold number ( $t_1$ ) separates the liver and the kidney. The second threshold ( $t_2$ ) separates the kidney and the aorta.
4. Segments the DICOM image by setting all pixel that have a value  $>t_1$  and  $<t_2$  to foreground and the rest to background.
5. Computes the DICE score between the resulting segmented image and the KidneyROI.png image.

You have tested the algorithm on the 1-162.dcm image.

Images and annotations for the kidney analysis.

(<https://resources.mcq.eksamen.dtu.dk/v1/4697d7d8-4319-4777-9753-7a014bad6ef6>)

What are the two thresholds  $t_1$  and  $t_2$ ?

Choose one answer

- 135.5 and 196.7
- Do not know
- 123.1 and 178.3
- 145.7 and 203.1
- 167.1 and 203.3
- 155.1 and 192.7

What is the computed DICE score?

Choose one answer

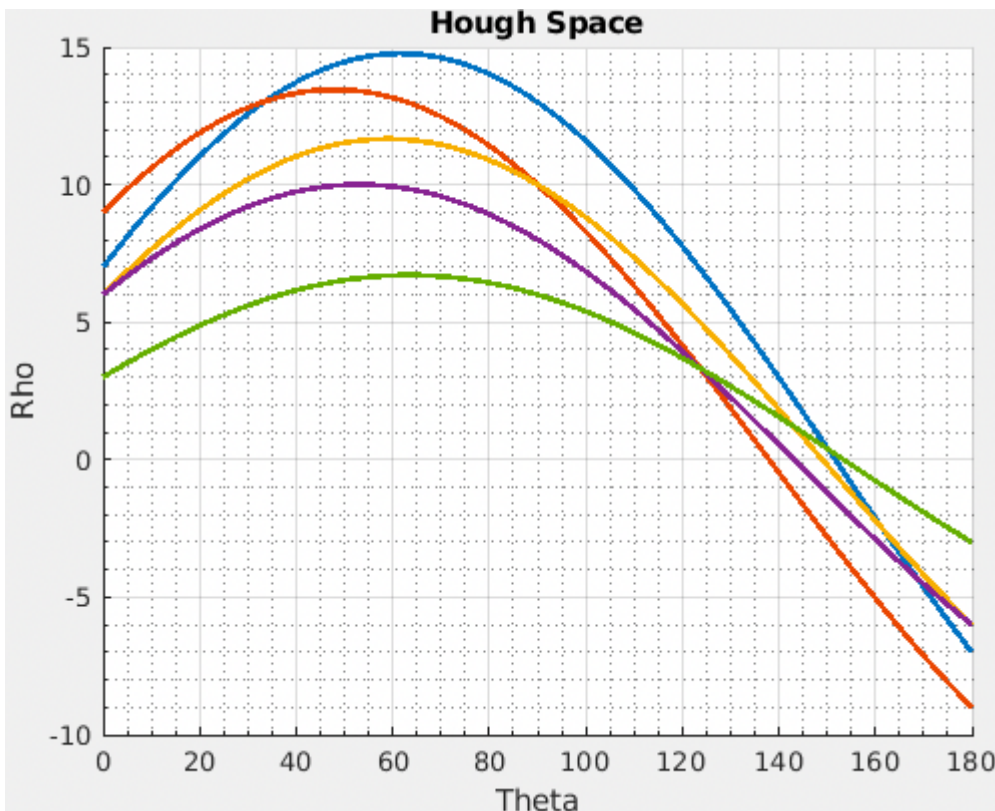
- 0.563
- 0.356
- 0.241
- 0.743
- Do not know
- 0.419

The five data points  $(x, y)$ :

$(7,13), (9,10), (6,10), (6,8), (3,6)$

result in the corresponding five Hough space sinograms (curves in Hough space) shown below.

Two Hough space sinograms are crossing at  $(\rho = 0.29, \theta = 151^\circ)$ . What are the corresponding data points for the two crossing curves in Hough space? (approximately)



Choose one answer

- $(9, 10)$  and  $(6, 10)$
- $(7, 13)$  and  $(3, 6)$
- $(9, 10)$  and  $(7, 13)$
- $(6, 10)$  and  $(3, 6)$
- Do not know
- $(7, 13)$  and  $(6, 10)$

We are searching for the shortest path going from the top to the bottom of the image below. You are using dynamic programming for the task and as part of that you are computing an *accumulator image*. What is the minimal value of the last row of the computed accumulator image?

64	94	21	19	31
38	88	30	23	92
81	55	47	17	43
53	62	23	23	18
35	59	84	44	90

Choose one answer

- 186
- 121
- 35
- Do not know

90

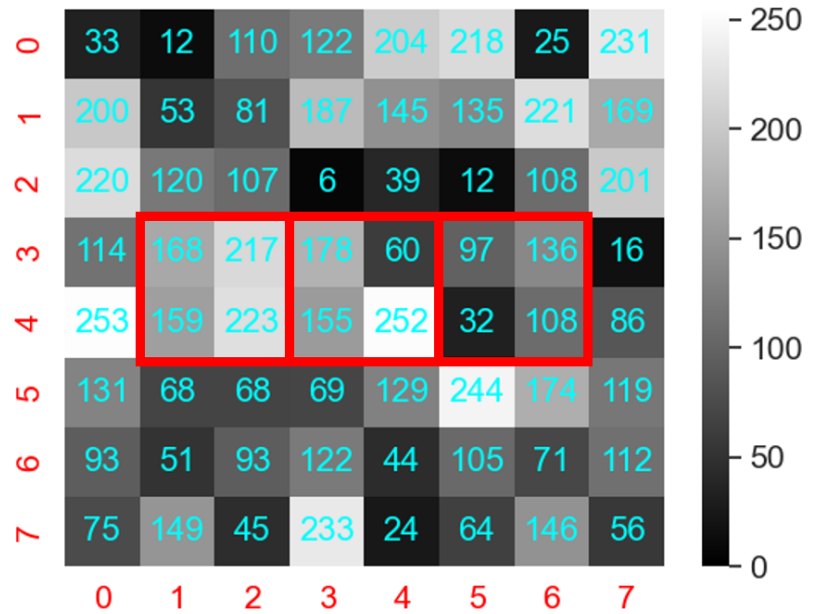
19

## Fast face detection

You want to create a new app that can detect faces and put funny hats on them. You are basing your method on the well known Viola Jones face detector.



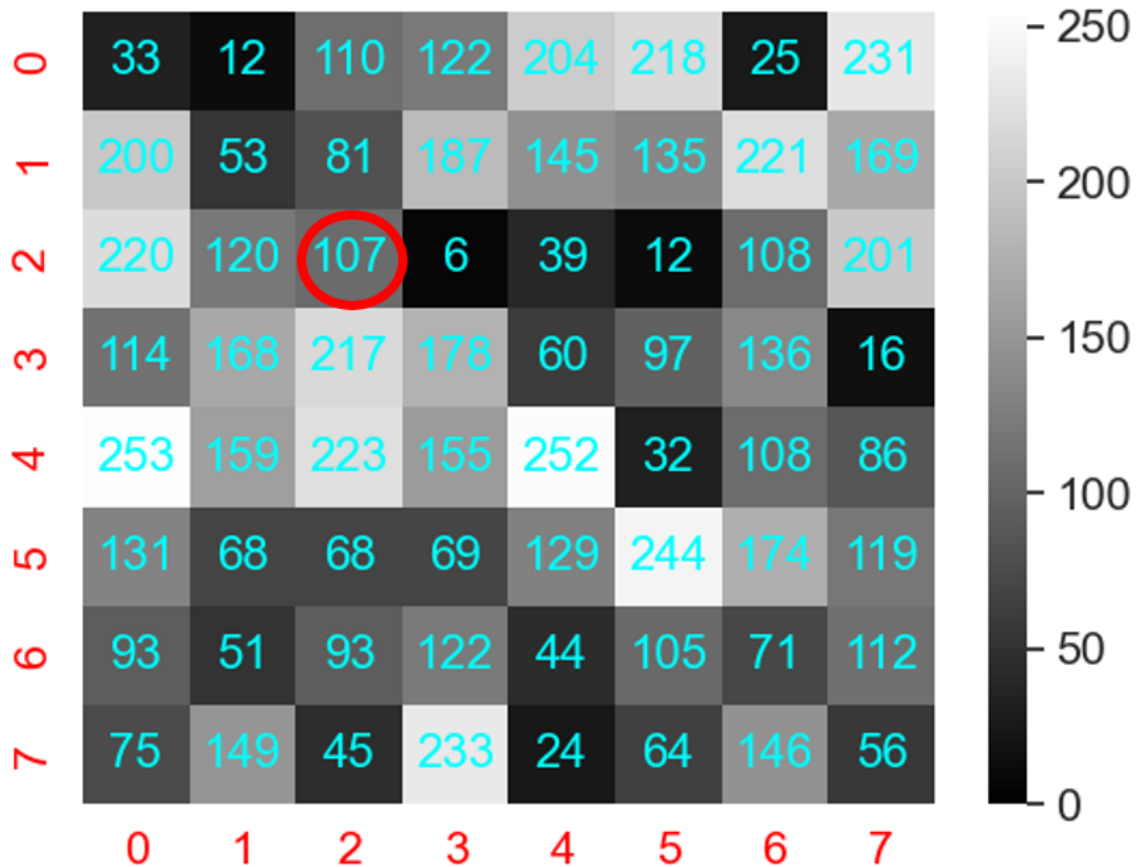
You need to compute many image features fast and have chosen to use the Haar features. In the image, you can see the *three rectangle Haar feature* and how it is placed in an image. What is the resulting feature value?



Choose one answer

- 116
- Do not know
- 317
- 495
- 324
- 9

To be able to compute image features very fast, you pre-compute an integral image. The input image can be seen below. What is the value of the integral image at the marked pixel?



Choose one answer

- 109
- 516
- 936
- 789
- Do not know
- 884

## Pixelwise operations

You are testing a prototype of a system that can monitor car parking areas. In order to locate the cars you want to create a binary image. Your algorithm does:



1. Converts the input image from RGB to gray scale and transforms it into a floating point image.
2. Does a linear gray scale transformation so the transformed image has a minimum pixel value of 0.1 and a maximum pixel value of 0.6
3. Computes a pixel value threshold using Otsu's method
4. Applies the found threshold, so all pixels with values above the threshold becomes foreground (white) and the rest background (black).


You test your algorithm on the supplied photo (pixelwise.png)

Example photo (pixelwise.png) (<https://resources.mcq.eksamen.dtu.dk/v1/a72b2e16-e776-44c8-9441-55b6a6766450>)


What is the resulting binary image, when applying your algorithm to the test photo?

Choose one answer






○



○



Do not know

What is the threshold value found by Otsu method.

Choose one answer

- 0.17
- 0.37
- 0.33
- Do not know
- 0.21
- 0.49

You would like to make a cool black and white illustration of a rocket. To do that you filter the image (rocket.png) with a Prewitt filter. Secondly, you threshold the filtered image, so all pixels with a value above 0.06 become white and the rest black. How many white pixels are there in the resulting image?

[Rocket image \(rocket.png\)](#).

<https://resources.mcq.eksamen.dtu.dk/v1/0534cdd8-40d1-4f7a-bab2-f59876ce2c57>

Choose one answer

- Between 1000 and 3000
- Between 9000 and 11000
- Between 3000 and 5000
- Between 5000 and 7000
- Between 7000 and 9000
- Do not know



## Linear Discriminant Analysis (LDA)

We have used data points from two classes in a two-dimensional space and trained a classifier model using two Gaussian distributions. The mean values are [24, 3] and [30, 7] for classes 1 and 2 respectively. As a classifier, we use the Linear Discriminant Classifier that assumes equal prior-probabilities and isotropic covariances.

The covariance matrix is:

$$\Sigma = \begin{bmatrix} 2 & 0 \\ 0 & 2 \end{bmatrix}$$

A general formulation for a hyperplane and points belonging to class 2 is:

$$y_{C \in 2}(\mathbf{x}) = \mathbf{x}^T \mathbf{w} + \mathbf{c}w$$

where  $\mathbf{w}$  is the weight vector and  $\mathbf{c}w$  is the threshold defining the hyperplane. So, if  $\mathbf{x}^T \mathbf{w} > \mathbf{c}w$  a sample point belongs to class 2.

What is  $y_{C \in 2}(\mathbf{x})$  for  $\mathbf{x}=(23, 5)$  and what class does the sample belong to?

[LDA helper function \(LDA.py\)](#)

<https://resources.mcq.eksamen.dtu.dk/v1/81cff419-4332-474d-be3d-243883b87ab9>

Choose one answer

- 12 and class 1
- 170 and class 2
- 30 and class 1
- 12 and class 1
- Do not know
- 1.7 and class 2

## Animal sorting machine

You have been asked by a farmer to develop a machine vision system that can detect if a cow or a sheep is passing by a camera. The camera just measure the average intensity of the image of the animal. To train your system, the farmer has given you these measured training data measured on 10 cows and 12 sheep:

Cows = [26, 46, 33, 23, 35, 28, 21, 30, 38, 43]

Sheep = [67, 27, 40, 60, 39, 45, 27, 67, 43, 50, 37, 100]

First, you try to compute the intensity threshold dividing cows and sheep using a minimum distance classifier. Secondly, you fit Gaussians to the two sets of the data to do a parametric classification.

What threshold did you find in the minimum distance classifier?

Choose one answer

- 42.6
- 39.4
- Do not know
- 42.1
- 38.1
- 41.2

You want to find the point where the two Gaussians cross. To do that you sample values of the two Gaussians for different values. What are the values of the Gaussians for the cows and the sheep when trying with value=38?

Choose one answer

- Do not know
- 0.03 and 0.02
- 0.04 and 0.04
- 0.03 and 0.04
- 0.03 and 0.03
- 0.04 and 0.02

## Finding mini figures

The company, PaintMyMiniz, would like to have a system that can automatically count the number of mini figures standing on a white plate. They have therefore called on your expertise. You decide to solve the task using BLOB analysis. Your first system works like this:

1. Converts the input photo from RGB to gray scale
2. Computes a threshold using Otsu's method
3. Computes a binary image by setting all pixels below the threshold to foreground and the rest to background
4. Removes BLOBs that are connected to the edges of the image
5. Computes all BLOBs in the image
6. Computes the area and the perimeter of all found BLOBs

You start by testing the system on an example photo (figures.png).

Test photo for BLOB analysis (figures.png)

(<https://resources.mcq.eksamen.dtu.dk/v1/59dd563f-0a09-4fb2-9890-e74c81f8ef07>)

You compute the area of all the BLOBs in the image. How many BLOBs have an area larger than 13000 pixels?

Choose one answer

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

You find the BLOB with the largest area. What is the perimeter of this BLOB?

Choose one answer

- 1589
- 1234
- 2034
- 1998
- 1679
- Do not know

## Car data analysis

A company, Carz2U, is selling used cars and needs a better way of estimating used car prices. They have gathered data on 203 used cars and they want you to do some analysis, to see how much information they can get out of the data they have gathered. For each car, they have recorded: wheel-base, length, width, height, curb-weight, engine-size, horsepower and highway-mpg.

To find out how much data they need, they want you to do a Principal Component Analysis (PCA) of the used car data. The measurements are of very different scales. Therefore, Carz2U asks you to normalize the data before doing the PCA. Normalizing is done by dividing each measurement by the standard deviation of the measurement before doing the PCA (the car width should be divided by the standard deviation computed over all the car widths for example).

The data for the question:

car\_data.txt (<https://resources.mcq.eksamen.dtu.dk/v1/bd8571b0-11e8-4232-8877-87f01a44c081>)



When you want to compute the PCA, you gather all data in a data matrix. What is the value of the data matrix at row=0, col=0 after subtraction of the mean and normalization?

Choose one answer

- Do not know
- 0.57
- 1.69
- 2.03
- 2.89
- 1.34

After computing the PCA of the car data, you compute the amount of the total variation explained by each of the principal components. How much of the total variation is explained by the two first principal components?

Choose one answer

- Do not know
- 57 %
- 46 %
- 69 %
- 87 %
- 91 %

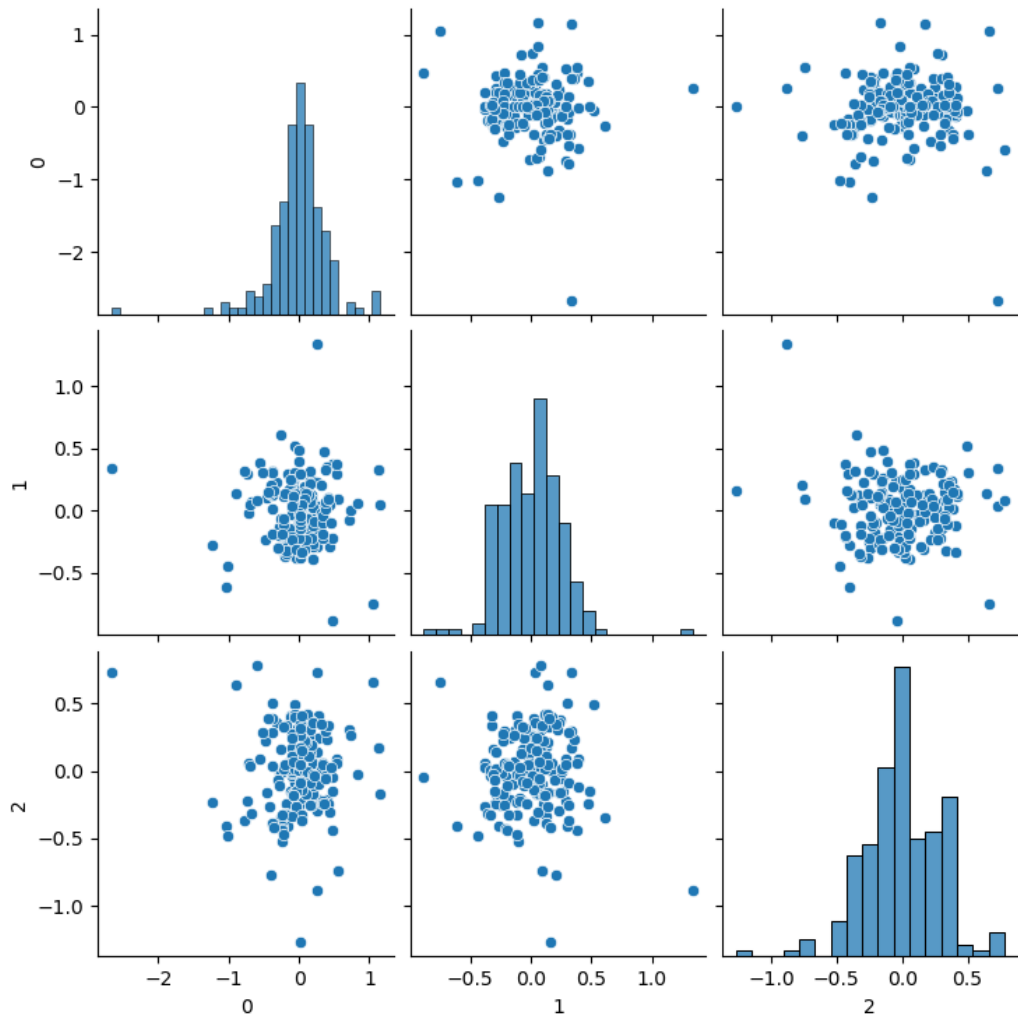
After computing the PCA, the data is projected onto the principal components to find their locations in PCA space. What is the absolute value of the first coordinate of the first car in the PCA space?

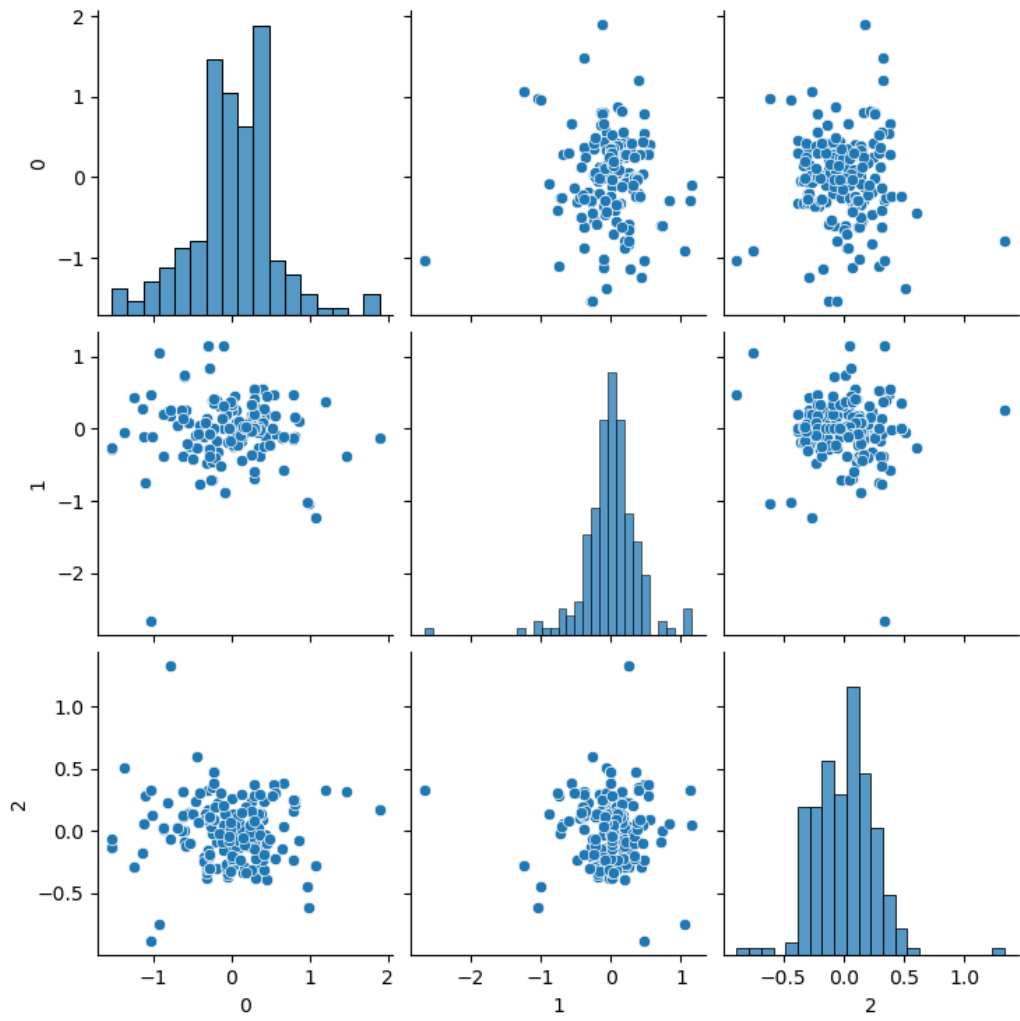
Choose one answer

- 4.12
- 0.34
- 0.97
- 1.13
- Do not know
- 2.31

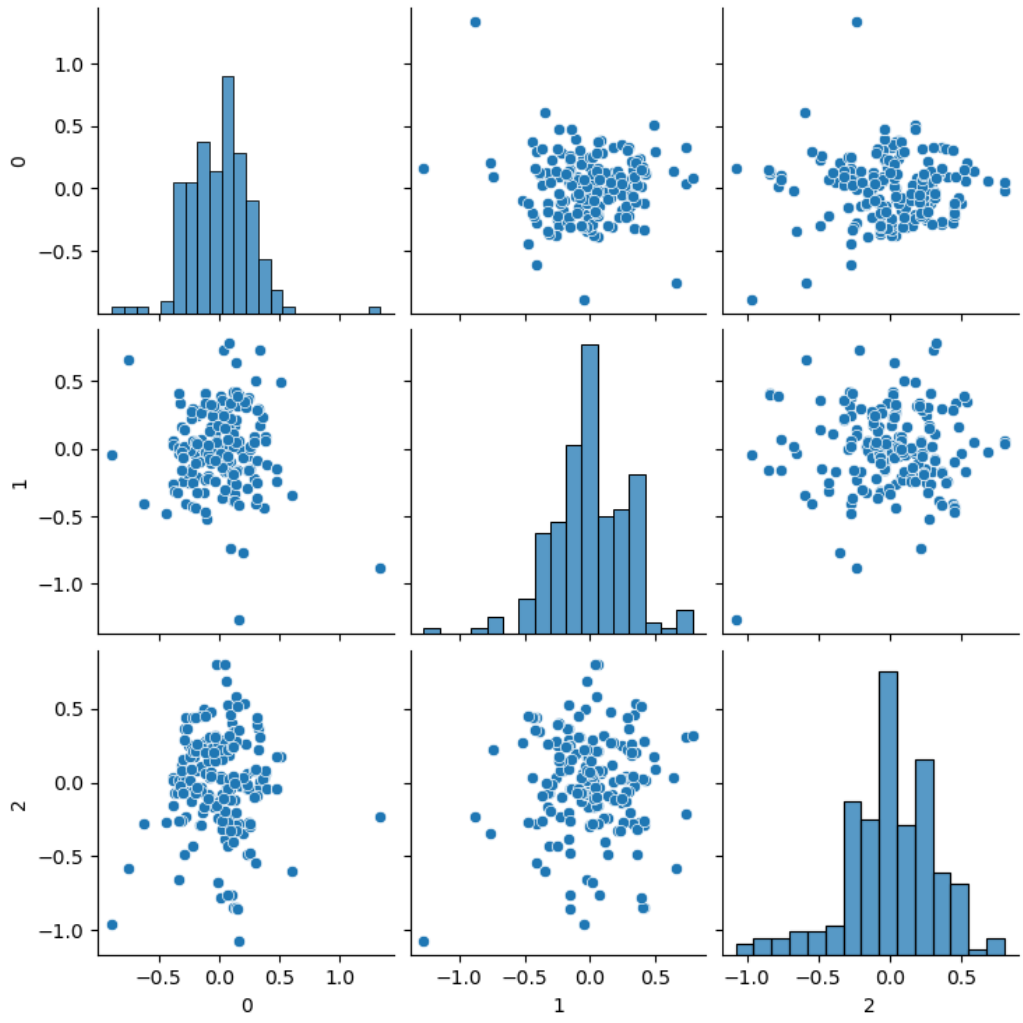
To get an overview of how the data is distributed after the PCA, you make a pairplot of the first three measurements (wheel-base, length and width) after they have been projected into PCA space. How does your pair plot look like?

Choose one answer

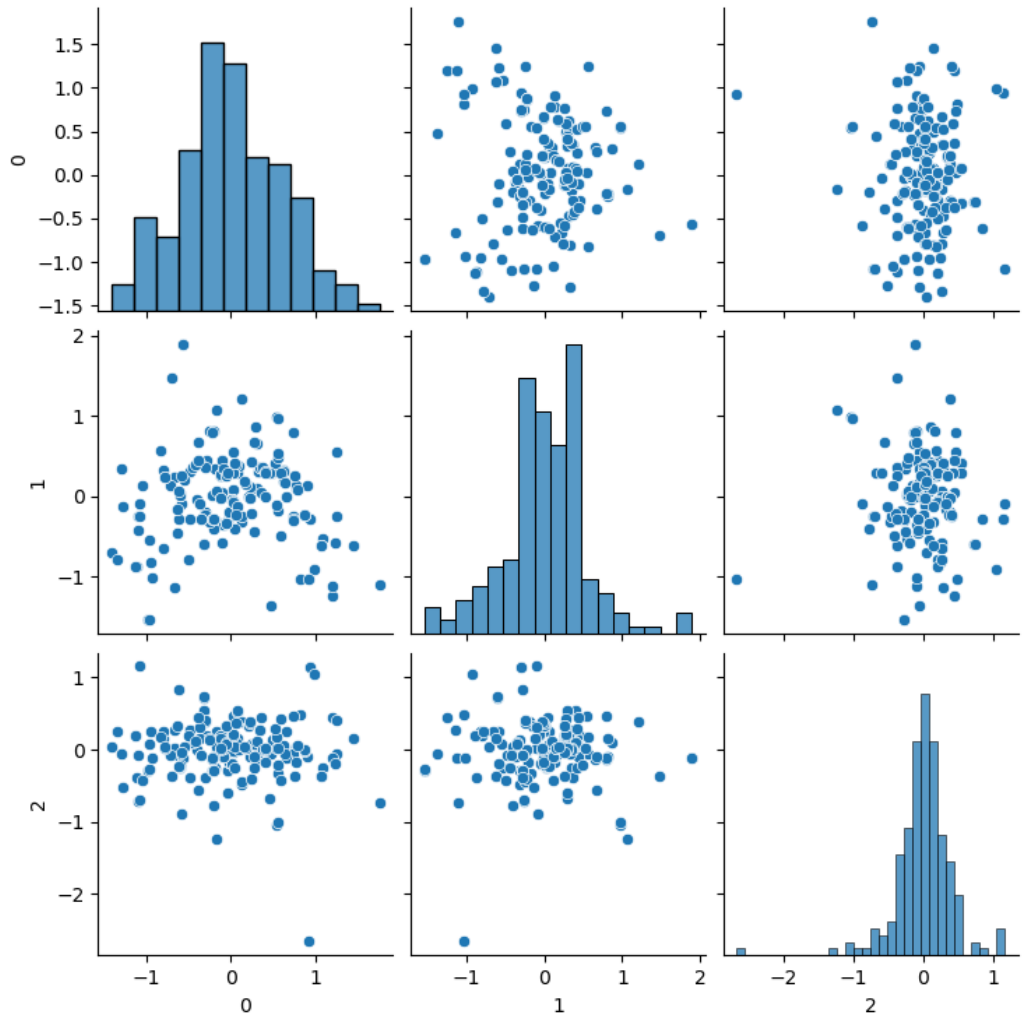




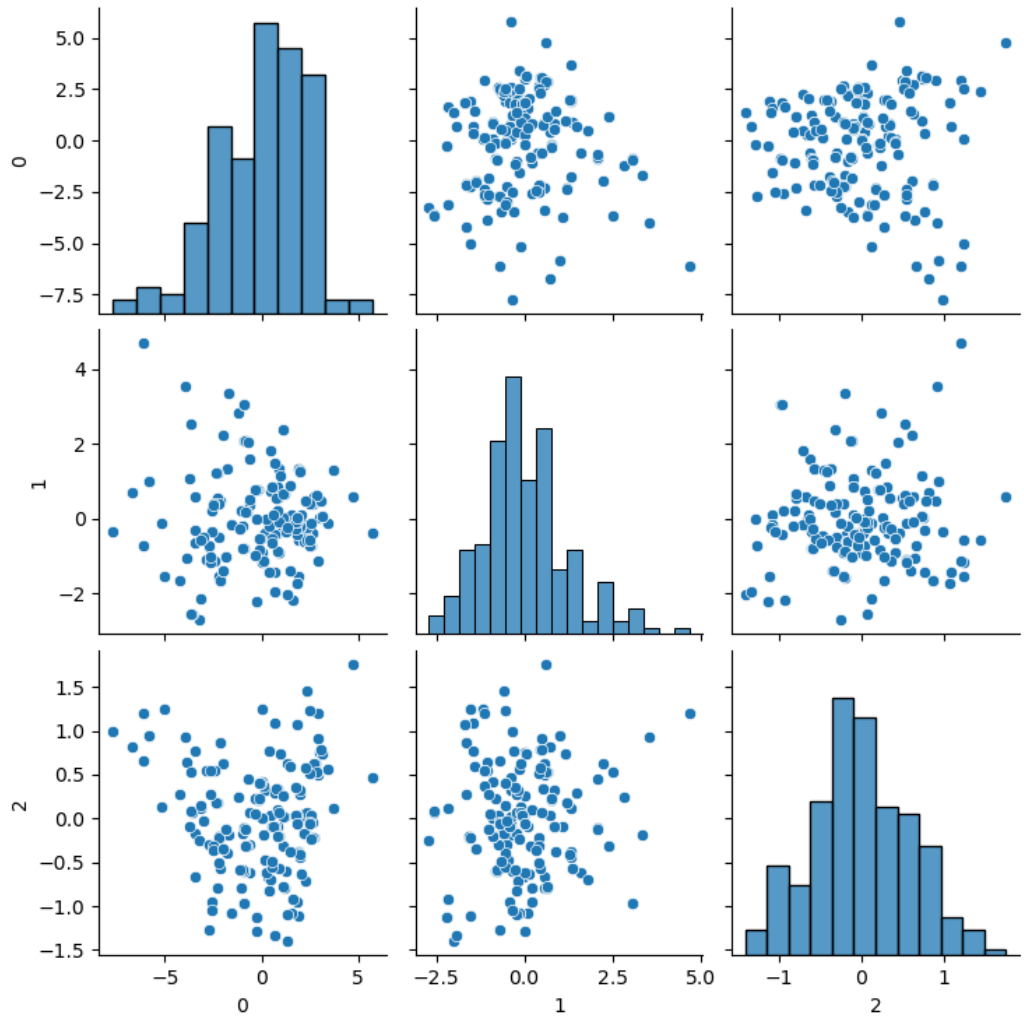
Do not know



○



○





## Landmark based registration

You are doing a landmark based registration of a photo of a toy rocket.

The source landmarks are:

(220, 55), (105, 675), (315, 675)

and the destination landmarks are:

(100, 165), (200, 605), (379, 525)

Initially, you compute the landmark alignment error  $F$  (objective function):

$$F = \sum D(a_i, b_i)^2$$

Where  $D$  is the Euclidean distance between points and  $a_i$  and  $b_i$  are the source and target landmarks.

Now you compute the Euclidean transform that brings the source landmarks over in the target landmarks.

To check if the registration is a success, you compute  $F$  again but uses the transformed source landmarks against the target landmarks.

Finally, you apply the found transformation on the rocket photo and convert the warped image to bytes using `img_as_ubyte`. Please note that you should use the option `tform.inverse` when using the `warp` function.

Rocket photo (<https://resources.mcq.eksamen.dtu.dk/v1/78d877a8-ac93-4b64-8b53-e25d7f9d4cd7>)

What is the pixel value at (row=150, column=150) in the warped image?

Choose one answer

- 54
- Do not know
- 123
- 129
- 176
- 187

What is the landmark alignment error, F, before the registration?

Choose one answer

- Do not know
- 67021
- 61231
- 54323
- 45613
- 3421

What is the landmark alignment error, F, after the registration?

Choose one answer

- 21120
- 15122
- Do not know
- 23141
- 8765
- 12423

You want to remove noise from an image (rocket.png) and you try with a Gaussian image filter with  $\sigma=3$ . After the filtering, you convert the image back into a byte image using `img_as_ubyte`. What is the value of the pixel at (row=100, column=100) in the resulting image?

[Image to filter \(rocket.png\)](#)

<https://resources.mcq.eksamen.dtu.dk/v1/a58f55e5-9d4b-45ad-90ca-20f246ebc41d>

Choose one answer

- 212
- 123
- 104
- 221
- 32
- Do not know

The image (CPHSun.png) are rotated 16 degrees with a rotation center of (20, 20). Finally, the rotated image is converted to a byte image using `img_as_ubyte`. What is the pixel value at (200, 200) in the final image?

[Image to be rotated \(CPHSun.png\)](#)

<https://resources.mcq.eksamen.dtu.dk/v1/36c618c8-19a5-47e3-ae42-22aa3fdff4cd>

Choose one answer

- (13, 51, 210)
- Do not know
- (118, 156, 12)
- (99, 172, 34)
- (73, 23, 14)
- (88, 91, 95)

You are developing a system, that can track cars on a light background. Your system performs these operations:

1. Converts the input image from RGB to HSV
2. Creates a new image that only contains the S component of the HSV image
3. Uses Otsu's method to compute a threshold in the S image
4. Creates a binary image by applying the found threshold on the S image, so pixels above the threshold become foreground and the rest background.
5. Performs a morphological erosion with a disk shaped structuring element with radius=4 on the binary image.

You test your system on a test image (pixelwise.png) and counts the number of foreground pixels in the resulting image. How many foreground pixels are there?

[Test image \(pixelwise.png\)](#)

<https://resources.mcq.eksamen.dtu.dk/v1/9a06b624-37df-4ed0-a5cb-7b8c520ee2be>

Choose one answer

- Between 6000 and 9000
- Between 2000 and 4000
- Between 4000 and 6000
- Between 9000 and 11000
- Less than 2000
- Do not know

## Video change detection

You and your friends have created a startup company, weCU, where you offer a surveillance system for warehouses. The system consists of a camera connected to a computer that raises an alarm if the system detects visual changes in the warehouse.



The camera is connected to the computer using a USB-2 connection. The maximum transfer speed between the camera and the computer is 30 megabytes per second. Each image is RGB (3 bytes per pixel) with a size of 1024 x 768 pixels.

Your change detection algorithm uses 54 milliseconds to process one image.

What is the maximum frames-per-second your system can do?

Choose one answer

- 19
- Do not know
- 12
- 24
- 33
- 5

To test your system, you have taken two test photos (change1.png and change2.png) they are RGB images of size (360, 457).

The [images for the question \(change1.png and change2.png\)](https://resources.mcq.eksamen.dtu.dk/v1/58820466-0561-4c91-b00f-efddadd4b4b6) (<https://resources.mcq.eksamen.dtu.dk/v1/58820466-0561-4c91-b00f-efddadd4b4b6>).

Your algorithm works like this:

1. Converts both images to gray scale using **color.rgb2gray**. Both images are now floating point images where the pixel values are between 0 and 1.
2. Computes the absolute difference image between the two gray scale images.
3. Computes how many pixels in the difference image that have a value above 0.3. These are the *changed pixels*.
4. Computes the percentage of changed pixels compared to the total number of pixels in the image.

What is the percentage of changed pixels for the test images?

Choose one answer

- Do not know
- 18 %
- 12%
- 21 %
- 5 %
- 2 %