

02502 Image Analysis Spring 2021

Technical University of Denmark

Written online exam, May 21, 2021

Course name: Image Analysis

Course number: 02502

Number of Questions: 20

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.

Data: All data for the exam can be downloaded from <http://people.compute.dtu.dk/rapa/02502F2021>
(<http://people.compute.dtu.dk/rapa/02502F2021>)

The **irisdata.txt** file contains measurements from 150 iris flowers. The measurements are the sepal length, sepal width, petal length and petal width. So you have $M=4$ features, $N=150$ observations. A principal component analysis (PCA) should be done on these data. How many percent of the total variation do the two first principal components explain?

Choose one answer

- 97.8%
- 89.3%
- 94.7%
- 92.4%
- 87.2%
- Do not know

The **irisdata.txt** file contains measurements from 150 iris flowers. The measurements are the sepal length, sepal width, petal length, and petal width. So you have $M=4$ features, $N=150$ observations. A principal component analysis (PCA) should be done on these data.

After the PCA, the flower data (sepal lengths and widths, petal lengths and widths) are projected into PCA space. What are the projected values of the first flower?

Choose one answer

- 2.68, -0.32, -0.028, 0.0023
- 0.023, -0.12, -1.22, 0.11
- 1.23, -0.44, 0.17, 2.33
- 1.32, 0.37, -2.31, 1.33
- 2.33, -1.23, 0.56, -1.23
- Do not know

The photo called sky_gray.png is loaded and a linear histogram stretching is performed so the new image has a maximum pixel value of 200 and a minimum pixel value of 10. What is the average pixel value of the new image?

Choose one answer

- 87
- 98
- 113
- 167
- 103
- Do not know

The photo called sky.png is loaded and an RGB threshold is performed with the limits $R < 100$, $G > 85$, $G < 200$, and $B > 150$. Pixels with values within these limits are set to foreground and the rest of the pixels are set to background.

The resulting 2D binary image is morphologically eroded using a disk-shaped structuring element with radius=5.

When doing an erosion the pixels beyond the image border are assigned a value of 1 (the default Matlab behavior).

How many foreground pixels are there in the final image?

Choose one answer

- 19977
- 50132
- 12311
- 23211
- 17654
- Do not know

The photo called flower.png is loaded and it is converted from the RGB color space to the HSV color space. Secondly, a threshold is performed on the HSV values with the limits $H < 0.25$, $S > 0.8$ and $V > 0.8$. Pixels with values within these limits are set to foreground and the rest of the pixels are set to background.

Finally, a morphological opening is performed on the binary image using a disk-shaped structuring element with radius=5. When doing a dilation, pixels beyond the image border are assigned a value of 0 and when doing an erosion the pixels beyond the image border are assigned a value of 1 (the default Matlab behavior).

What is the number of foreground pixels in the resulting image?

Choose one answer

- 5665
- 3421
- 8753
- 6712
- 1340
- Do not know

Five photos have been taken. They are named car1.jpg - car5.jpg and they have the dimensions (W=800, H=600). A principal component analysis (PCA) is performed on the grey values of the five images. You can use the two helper functions `pc_evectors.m` and `sortem.m` to compute the PCA. How much of the total variation in the images is explained by the first principal component?

Choose one answer

- 40.1%
- 27.3%
- 55.8%
- 78.3%
- 67.4%
- Do not know

The photo called sky_gray.png is transformed using a gamma mapping with $\gamma=1.21$. The output image is filtered using a 5x5 median filter. What is the resulting pixel value in the pixel at row=40, column=50 (when using a 1-based matrix-based coordinate system)?

Choose one answer

- 68
- 12
- 30
- 123
- 233
- Do not know

The photo called flowerwall.png is filtered using an average filter with a filter size of 15. The filtering is performed with border replication. What is the resulting pixel value in the pixel at row=5 and column=50 (when using a 1-based matrix-based coordinate system)?

Choose one answer

- 208
- 198
- 167
- 234
- 123
- Do not know

A photo has been taken of a set of floorboards (floorboards.png) and the goal is to measure the amounts of knots in the wood. First, a threshold of 100 is used, so pixels below the threshold are set to foreground and the rest is set to background. To remove noise a morphological closing is performed with a disk-shaped structuring element with radius=10 followed by a morphological opening with a disk-shaped structuring element with radius=3. When doing a dilation, pixels beyond the image border are assigned a value of 0 and when doing an erosion the pixels beyond the image border are assigned a value of 1 (the default Matlab behavior). Finally, all BLOBs that are connected to the image border are removed. How many foreground pixels are remaining in the image?

Choose one answer

- 6735
- 4321
- 7632
- 1300
- 2311
- Do not know

A photo has been taken of a set of floorboards (floorboards.png) and the goal is to measure the amounts of knots in the wood. First, a threshold of 100 is used, so pixels below the threshold are set to foreground and the rest is set to background. To remove noise a morphological closing is performed with a disk-shaped structuring element with radius=10 followed by a morphological opening with a disk-shaped structuring element with radius=3. When doing a dilation, pixels beyond the image border are assigned a value of 0 and when doing an erosion the pixels beyond the image border are assigned a value of 1 (the default Matlab behavior). A BLOB analysis is performed where all BLOBS are found using 8-connectivity. All BLOBs that are connected to the image border are removed.

The area of the found BLOBs are computed and only the BLOBs with an area larger than 100 pixels are kept. How many BLOBs are found in the final image?

Choose one answer

- 16
- 13
- 12
- 25
- 9
- Do not know

The binary image books_bw.png contains letters. A BLOB analysis is performed using 8-connectivity. For each BLOB, the area and the perimeter is computed. The BLOBs with area > 100 and perimeter > 500 are kept. Which letters are visible in the final image?

Choose one answer

- G and m
- C, a and V
- m and b
- J, G, and V
- b, a, and m
- Do not know

Seven corresponding landmarks have been placed on two images (cat1.png and cat2.png). The landmarks are stored in the files catfixedPoints.mat and catmovingPoints.mat. What is the sum of squared differences between the fixed and the moving landmarks?

Choose one answer

- 46553
- 21223
- 76321
- 33765
- 56412
- Do not know

Seven corresponding landmarks have been placed on two images (cat1.png and cat2.png). The landmarks are stored in the files catfixedPoints.mat and catmovingPoints.mat. A similarity transform (translation, rotation, and scaling) has been performed that aligns the moving points to the fixed points. The computed transform is applied to the cat2.png photo. How does the resulting image look like?

Choose one answer





Do not know

An abdominal scan has been acquired on a CT scanner. One of the slices of the scan is stored as a DICOM file called 1-179.dcm. An expert has marked a part of the liver as a binary mask (region of interest). The binary mask is stored as the file LiverROI.png. By using the DICOM image and the mask image, the image values in the DICOM image inside the mask (the liver) can be extracted.

The average value and the standard deviation of the extracted pixel values are computed. A low threshold, T_1 , is defined as the average value minus the standard deviation and a high threshold, T_2 , is defined as the average value plus the standard deviation.

Finally, a segmentation of the DICOM image (1-179.dcm) is made where all pixels with values $> T_1$ and $< T_2$ are set to foreground and the rest are set to background. How many foreground pixels are there?

Choose one answer

- 8984
- 7645
- 10221
- 6745
- 5482
- Do not know

An abdominal scan has been acquired on a CT scanner. One of the slices of the scan is stored as a DICOM file called 1-179.dcm. A low threshold, $T1 = 90$, and a high threshold, $T2 = 140$, are defined. The pixel values of the DICOM image are segmented by setting all pixel values that are $>T1$ and $<T2$ to foreground and the rest are set to background.

The binary image is processed by first applying a morphological closing using a disk-shaped structuring element with $\text{radius}=3$ followed by a morphological opening with the same structuring element. When doing a dilation, pixels beyond the image border are assigned a value of 0 and when doing an erosion the pixels beyond the image border are assigned a value of 1 (the default Matlab behavior).

In the final step, a BLOB analysis is done using 8-connectivity. The largest BLOB is found. The area (in pixels) of the largest BLOB is:

Choose one answer

- 5944
- 4123
- 8293
- 3692
- 6032
- Do not know

NASA's Mars Perseverance rover has explored Mars since its landing at the beginning of 2021. To explore the surface of Mars, the rover uses a custom build camera. Now the rover has discovered three spectral peaks that might reflect different types of cosmic dust. Each dust spectra appears to follow a normal distribution. The parametric distributions of the three dust classes are $N(7,2^*2)$, $N(15,5^*5)$, and $N(3,5^*5)$.

NASA asks help to define the thresholds to perform robust classification.

They wish to perform a minimum distance classification of the three dust classes.

What signal thresholds should NASA use?

Choose one answer

- 5.0 and 11.0
- 7.2 and 11.0
- 12.3 and 15.6
- 4.2 and 9.0
- 7.0, 9.0 and 11.0
- Do not know

NASA's Mars Perseverance rover has explored Mars since its landing at the beginning of 2021. To explore the surface of Mars, the rover uses a custom build camera. Now the rover has discovered three spectral peaks that might reflect different types of cosmic dust. Each dust spectra appears to follow a normal distribution. The parametric distributions of the three dust classes are $N(7,2^*2)$, $N(15,5^*5)$, and $N(3,5^*5)$.

NASA asks help to define the thresholds to perform robust classification.

They wish to perform a parametric classification of the three dust classes.

What signal thresholds should NASA use?

Choose one answer

- 4.24 and 10.29
- 0.65 and 11.27
- 0.65 and 10.29
- 4.24 and 8.75
- 3.24 and 7.63
- Do not know

A template matching is done on the image:

246	204	245	173	44
40	36	167	193	180
248	108	9	189	8
244	234	217	100	71
124	202	238	167	12

using the template:

208	233	71
231	161	139
32	25	244

The normalised cross correlation (NCC) between the image and the template is computed. What is the NCC in the marked pixel in the image?

Choose one answer

- 0.67
- 0.89
- 0.45
- 0.92
- 0.75
- Do not know

A company is making an automated system for fish inspection. They are using a camera with a CCD chip that measures 5.4 x 4.2 mm and that has a focal length of 10 mm. The camera takes photos that have dimensions 6480 x 5040 pixels and the camera is placed 110 cm from the fish, where a sharp image can be acquired of the fish.

How many pixels wide is a fish that has a length of 40 cm?

Choose one answer

- 4364
- 4135
- 3213
- 5612
- 1872
- Do not know

Two types of mushrooms (A and B) have been grown in Petri dishes. It appears that the mushrooms only can grow in specific positions in the Petri dish. You are asked to train a linear discriminant analysis (LDA) classifier to estimate the probability of a mushroom type growing at a given position in the Petri dish. It is a very time-consuming experiment, so only five training examples for each type of mushroom were collected.

The training data are:

Class 0: Mushroom type A and their grow positions (x,y):

(1.00, 1.00)
(2.20, -3.00)
(3.50, -1.40)
(3.70, -2.70)
(5.00, 0)

Class 1: Mushroom type B and their grow positions(x,y):

(0.10, 0.70)
(0.22, -2.10)
(0.35, -0.98)
(0.37, -1.89)
(0.50, 0)

Note: To train the LDA classifier to obtain the weight-vector W for classification, use the provided Matlab function: LDA.m

What is the probability that the first training example of Mushroom Type A, with position (1.00, 1.00), actually belongs to class 1?

Choose one answer

- 0.81
- 0.18
- 0.73
- 0.99
- 0.65
- Do not know

