

02502 Image Analysis Exam Spring 2024

02502 Image Analysis Exam Spring 2024

Written exam, May 17, 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 download here. (<https://designer.mcq.eksamen.dtu.dk/api/images/8ae658c5-dbc1-4515-a46d-7111564b02a7>)

Statistical analysis of wine measurements

To improve wine production in Italy, three producers of wine measured 13 variables for a selection of their wines. They measured:

- Alcohol
- Malic acid
- Ash
- Alcalinity of ash
- Magnesium
- Total phenols
- Flavanoids
- Nonflavanoid phenols
- Proanthocyanins
- Color intensity
- Hue
- OD280/OD315 of diluted wines
- Proline

for 178 wines. They are stored in a text file, where each row contains the measurements . The last column of the data is the producer of the wine (1, 2 or 3).

In the following, only the columns containing the measurements should be used for the PCA. The last column should be used to keep track of the producer of the wine.

You can for example do:

```
x_org = np.loadtxt(data_name, comments="%")  
x = x_org[:, :13]  
producer = x_org[:, 13]
```

The goal is to learn more about the data using a principal component analysis (PCA). Before the PCA analysis, the data is scaled by first computing the mean value of each measurement (for example the mean of the alcohol) and then subtracting the mean from the measurements. Secondly, the difference between the minimum and maximum of each measurement is computed and the measurement is divided by this value.

To do the PCA, the covariance is first computed followed by a call to **np.linalg.eig**.

The normalized measurements are also projected to the PCA space.

Data: All data for the exam can be download here.

(<https://designer.mcq.eksamen.dtu.dk/api/images/8ae658c5-dbc1-4515-a46d-7111564b02a7>)

After projection to PCA space the wines from producer 1 and producer 2 are compared. The average projected value on the first principal component is computed for the wines from producer 1 and from producer 2. What is the difference between the two average values?

Choose one answer

- Between 0.70 and 1.0
- Do not know
- Between 0.50 and 0.70
- Between 0 and 0.20
- Between 0.20 and 0.30
- Between 0.30 and 0.50

After projection to PCA space, the coordinates of the wines in PCA space is examined. The minimum and maximum projected coordinates on the first principal component is found and the difference is computed. What is it?

Choose one answer

- Between 3 and 4
- Do not know
- Between 1 and 2
- Between 2 and 3
- Between 4 and 5
- Between 0 and 1

What is the alcohol level of the first wine after normalization (subtraction of mean and division by max-min)

Choose one answer

- Do not know
- between 0.40 and 0.50
- between 0.20 and 0.30
- between 0.0 and 0.10
- between 0.30 and 0.40
- between 0.10 and 0.20

What is the average value of the elements in the covariance matrix?

Choose one answer

- between 0.004 and 0.006
- between 0.002 and 0.003
- between 0.001 and 0.002
- Do not know
- between 0.003 and 0.004
- between 0.006 and 0.008

How many percent of the total variation in the data set is explained by the first five principal components?

Choose one answer

- Do now know
- Between 75% and 80%
- between 70% and 75%
- Between 85% and 90%
- Between 80% and 85%
- Between 90% and 95%

Flower design using principal component analysis (PCA)

A newly formed company, Bulbz, has approached you and asked you to help with their new line of flowers. They have provided you with 15 photos of their flower named **flower01.jpg** to **flower15.jpg**.

They would like you to use a data driven approach, principal component analysis (PCA) to analyze their flowers.

You start by computing an average flower image and then the PCA from **sklearn.decomposition** to compute 5 principal components from the flower images. You also project all the flower images onto the principal components to find their positions in PCA space.

Data: All data for the exam can be download here.

(<https://designer.mcq.eksamen.dtu.dk/api/images/8ae658c5-dbc1-4515-a46d-7111564b02a7>)

You want to explore the major visual variation in the data set. To do that you synthesize two flower images by following the direction of the first principal component. You can for example do that by:

```
synth_image_plus = average_image + 3 * np.sqrt(image_pca.explained_variance_[0])  
* image_pca.components_[0, :]  
synth_image_minus = average_image - 3 *  
np.sqrt(image_pca.explained_variance_[0]) * image_pca.components_[0, :]
```


After showing the average image and the two synthetic images, you asses that the major visual variation is:

Choose one answer

- The position of the flower in the image
- The number of leaves on the flower
- Do not know
- The color of the flower
- The size of the flower
- The length of the leaves


After you have computed the average flower image, you show it. It looks like:

Choose one answer

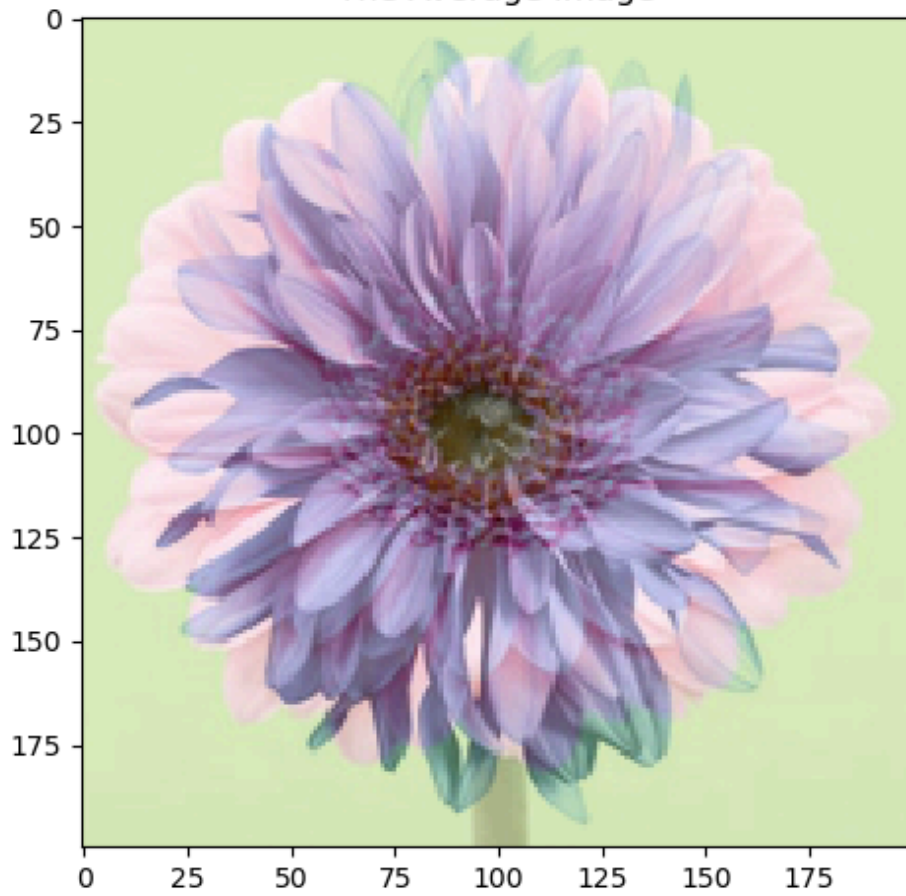




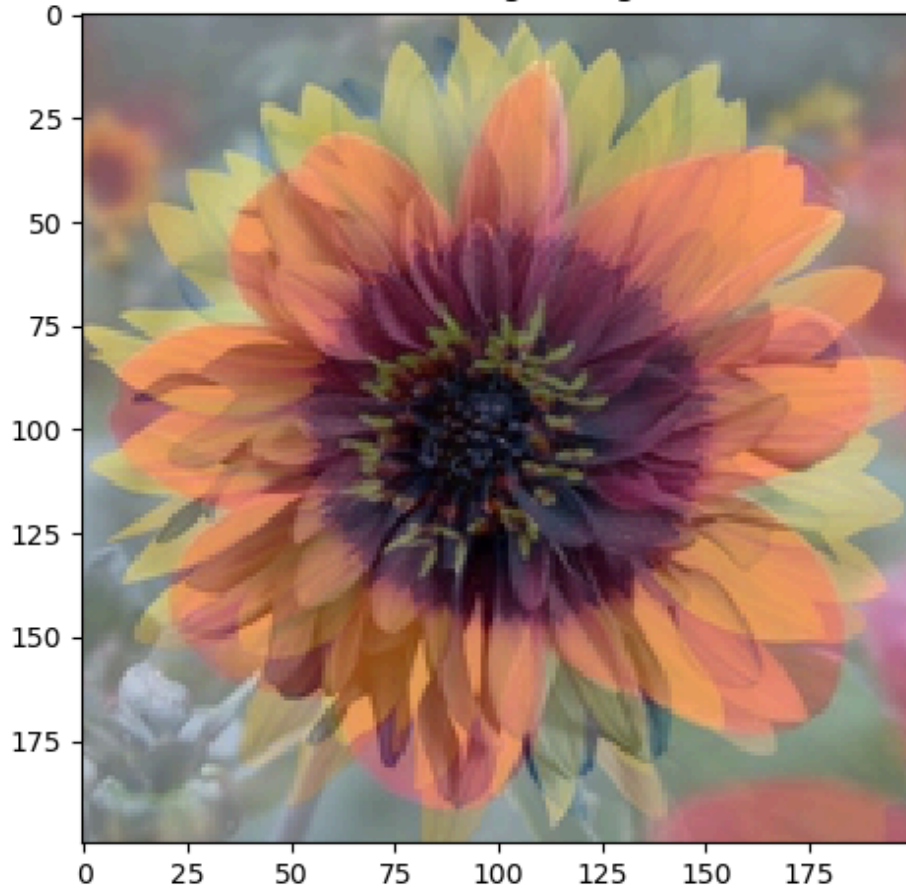
The Average Image




The Average Image



The Average Image



The Average Image

 Do not know

You would like to find out which two flowers are furthest away from each other when only projected to the first principal component. They are:

Choose one answer

- Do not know
- flower01 and flower12
- flower10 and flower12
- flower 02 and flower05
- flower08 and flower09
- flower03 and flower 07

How much of the total variation in the dataset is explained by the first principal component?

Choose one answer

- Between 20% and 30%
- Between 30% and 40%
- Do not know
- Between 40% and 50%
- Between 50% and 60%
- Between 60% and 70%

The company has given you an image of a new type of flower, **idealflower.jpg**. You would like to find a good match to the ideal flower in the set of 15 flowers. To do that, you project **idealflower.jpg** onto the principal components. Secondly, you only use the second principal component and find the flower that is closest to the ideal flower when projected onto the 2nd principal component. Select the image with that flower.

Choose one answer





Do not know

Spine analysis on computed tomography scans

The first step in diagnosing spine related diseases is to locate and segment the individual elements of the spine - the vertebrae. To develop a segmentation algorithm, we first try with a single DICOM slice (1-353.dcm) to get an idea of the appearance of the vertebra.

An expert has also created a mask (vertebra_gt.png) of the vertebra seen in the DICOM slice.

We start by reading the DICOM file and extracting the pixel values (in Hounsfield units).

Bone has a high x-ray attenuation and we manually select a threshold, so all pixel with HU value above 200 are set to foreground and the rest of the pixels are set to background. To remove small noise structures, a morphological closing with a disk-shaped structuring element (radius=3) is performed on the binary image.

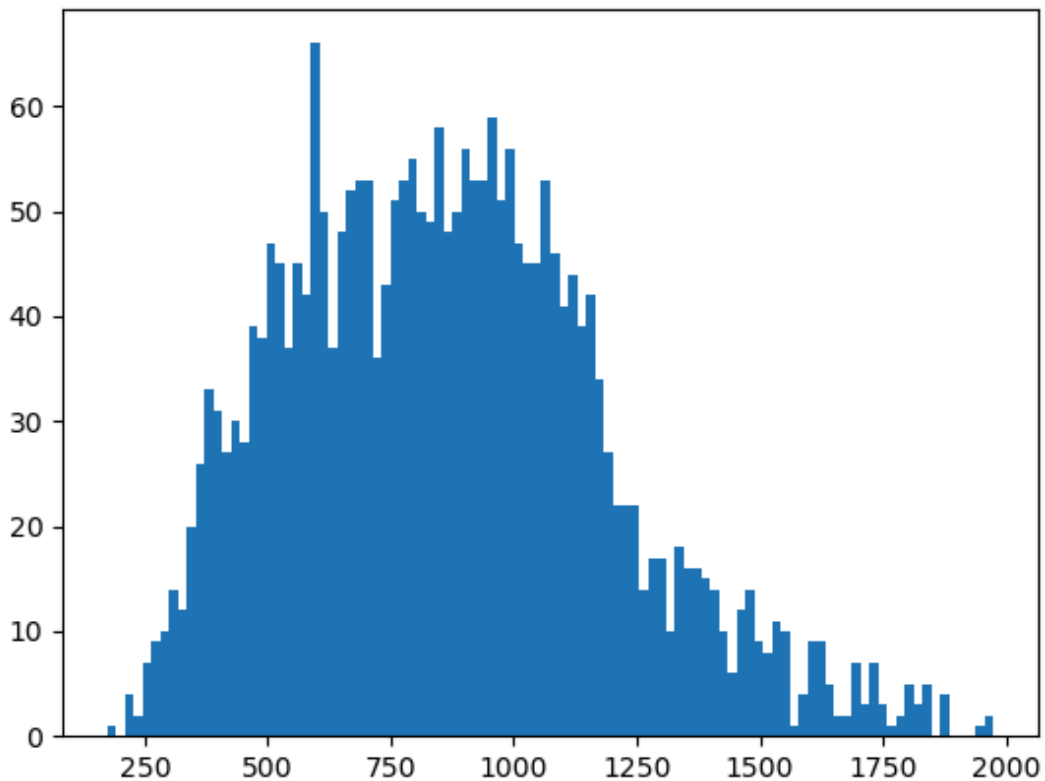
We do a BLOB analysis of the cleaned binary image to be able to recognize the vertebra. The area (number of pixels) of all the found BLOBs are computed. Finally, only the BLOBs with an area larger than 500 pixels are kept. We consider this the result of our segmentation algorithm.



Data: All data for the exam can be download here.


(<https://designer.mcq.eksamen.dtu.dk/api/images/8ae658c5-dbc1-4515-a46d-7111564b02a7>)

You notice that the expert mask contains both the hard bone of the vertebra and the softer bone inside. To get an idea of the distribution of the Hounsfield units in the masked part of the vertebra, you use the mask to extract the underlying pixel values from the DICOM slice. You then plot a histogram with 100 bins of these values. How does this histogram look like?


Choose one answer







Do not know



To get an idea of the true distribution of Hounsfield units in the vertebra, you sample the original pixel values in the mask that you found by the algorithm. What is the mean and standard deviation of the HU values?

Choose one answer

- Average is in [350, 400] and standard deviation is in [220, 230]
- Do not know
- Average is in [450, 500] and standard deviation is in [290, 300]
- Average is in [350, 400] and standard deviation is in [150, 160]
- Average is in [400, 450] and standard deviation is in [120, 130]
- Average is in [300, 350] and standard deviation is in [210, 220]

After the BLOB analysis, the areas of all the BLOBs are computed. What is the minimum and maximum areas?

Choose one answer

- Minimum is in [10, 20] and maximum is in [1000, 1500]
- Minimum is in [0, 10] and maximum is in [1000, 1500]
- Minimum is in [10, 20] and maximum is in [500, 1000]
- Do not know
- Minimum is in [0, 10] and maximum is in [500, 1000]
- Minimum is in [50, 100] and maximum is in [2000, 3000]

You compare your found segmentation with the mask provided by the expert. What is the DICE score?

Choose one answer

- Between 0.80 and 0.85
- Between 0.75 and 0.80
- Do not know
- Between 0.70 and 0.75
- Between 0.65 and 0.70
- Between 0.60 and 0.65

Analysis of zebra stripe patterns

A scientist wants to investigate the morphology of zebra stripes and their possible genetic relation and has collected many images of zebras. As an image analysis expert, you are asked to establish a two-class parametric pixel classification algorithm. Class 1 contains the white stripes and class 2 the black stripes. The zebra expert has manually drawn regions of interest (ROI) on a zebra image (Zebra.png) indicating which pixels to use for training.

Two binary masks are provided. One with examples of white stripes (Zebra_whiteStripes.png) and one with example of black stripes (Zebra_blackStripes.png).

Using the original image and the masks, train the parametric classifier and find the optimal threshold to separate white and black stripes.

The original image is stored as a one-channel gray scale image and should not be converted from RGB.

Data: All data for the exam can be download here.

(<https://designer.mcq.eksamen.dtu.dk/api/images/8ae658c5-dbc1-4515-a46d-7111564b02a7>)

The expert has also provided a round mask (zebra_MASK.png) indicating the part of the zebra that should be analyzed. Using your classifier, you classify all the pixel inside the mask as being either black or white stripe. How many pixels are classified as white stripe?

Choose one answer

- Between 2000 and 3000
- Between 6000 and 7000
- Between 5000 and 6000
- Do not know
- Between 3000 and 4000
- Between 4000 and 5000

What is the class range for the black stripes?

Choose one answer

- [0, 124]
- [0, 96]
- [0, 102]
- Do not know
- [0, 117]
- [0, 137]

What are the parameters (mean and standard deviation) of the Gaussian distribution for the trained classifier for the white stripes?

Choose one answer

- Mean is in range [190, 200] and standard deviation is in range [35, 45]
- Mean is in range [160, 170] and standard deviation is in range [50, 60]
- Mean is in range [210, 220] and standard deviation is in range [10, 20]
- Mean is in range [140, 150] and standard deviation is in range [25, 35]
- Mean is in range [180, 190] and standard deviation is in range [20, 30]
- Do not know

Image registration using gradient descent

You are performing an image registration with two images. You are only doing a simple translation of one of the images and this translation is parameterized by (x_1, x_2) . It turns out that the cost function can be expressed as:

$$C(x_1, x_2) = x_1^2 - x_1 \cdot x_2 + 3x_2^2 + x_1^3$$

To find the optimal translation, a gradient descent optimization algorithm is used. The starting guess is $(x_1, x_2) = (4, 3)$ and a constant step length of 0.07 is used.

What is x_1 after 5 iterations?

Choose one answer

- Between 1.0 and 5.0
- Between 0.00 and 0.15
- Between 0.15 and 0.25
- Between 0.25 and 0.50
- Between 0.50 and 1.0
- Do not know

How many iterations are needed before $C(x_1, x_2)$ is below 0.20 (the first time it is below)?

Choose one answer

- 12
- Do not know
- 6
- 2
- 48
- 20

Sorting plastic using linear discriminant analysis

A computer vision system is used to sort plastic into two types (class 1 and class 2). Each camera looks at a piece of plastic and computes an overall number based on the color and shape of the plastic piece. That means that there is only two numbers for each piece of plastics. We can plot these values into a 2-dimensional features space (x_1, x_2).

You are asked to train a linear discriminant analysis (LDA) classifier to see if you can create an optimal hyperplane separating the two types of plastic into class 1 and class 2. You have (previously, not used in these questions) been provided with 100 samples from each class and you have found out that each class have a prior probability of 0.5. Class 1 has a mean value of (24, 3) and class 2 has a mean value of (45, 7). The covariance matrix for both classes is:

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

Note: There are no further data needed for these questions.

In the sorting plant, a robot arm select plastic samples. One sample is analyzed and the computer vision system outputs $(x_1, x_2) = (30, 10)$.


What is the computed $y(x)$ value determining if the plastic piece belongs to class 2 and what class does it belong to?

Choose one answer



- (40.2,1)
- (-207.25,1)
- (707, 2)
- (-37.3, 1)
- Do not know
- (-1.7,2)



After the system has been installed, the sorting plant has sent you a figure with the camera measurements. You can see that the distribution of the samples fits with your training data. What figure did you receive?

Choose one answer



Do not know





3D image registration

You have developed a novel deep learning algorithm for performing 3D image registration. For testing purposes, you want to design some test cases to evaluate the performance of your algorithm. You will use a 3D MRI image of a brain (T1_brain_template.nii.gz).

To inspect the results, you use the ortho view that shows the brain in three planes: The axial view shows the x-y plane, the coronal views the x-z plan, and the sagittal view shows the y-z plan of the brain (for example the `imshow_orthogonal_view` from the advanced registration exercise).

You apply a rigid registration to “T1_brain_template” to generate your first test. The rigid transform is composed of a rotation with a yaw of 10 degrees followed by a pitch of -30 degrees. The output will be the moving image for your tests.

You also generate a mask by applying Otsu thresholding on the template volume. Then, apply a morphological closing with a `skimage.morphology.ball` as a structuring element with radius 5. Finally, apply erosion with a `ball` with a radius of 3.

The mask is applied to both the moving and template images. Compute the intensity based normalized correlation coefficient between the two (Equation 2.6 in Elastix notes).

Tip: For doing the exercise, you may need to work with numpy arrays. You can extract the numpy array from a SimpleITK object using the function `sitk.GetArrayFromImage(...)`. If needed, you can make a SimpleITK object from the numpy array with the function `sitk.GetImageFromArray(...)`.

Data: All data for the exam can be download here.

(<https://designer.mcq.eksamen.dtu.dk/api/images/8ae658c5-dbc1-4515-a46d-7111564b02a7>)

How does the binary mask look like?

Choose one answer

Binary image

Axial



Coronal



Sagittal



Binary image

Axial



Coronal



Sagittal



Binary image

Axial



Coronal



Sagittal



Binary image

Axial



Coronal



Sagittal



Do not know

Binary image

Axial



Coronal



Sagittal



What is the normalized correlation coefficient between the moving and the template image when masked with the binary mask?

Choose one answer


- Between 0.35 and 0.45
- Between 0.15 and 0.25
- Do not know
- Between 0.45 and 0.55
- Between 0.0 and 0.15
- Between 0.25 and 0.35

How does the moving image look like after the two-step rigid transformation?


Choose one answer

Moving image


Axial



Coronal




Sagittal




Moving image


Axial



Coronal




Sagittal




Moving image

Axial



Coronal




Sagittal




Moving image


Axial



Coronal




Sagittal




Moving image


Axial



Coronal



Sagittal



Do not know

You would like to estimate the amount of red flowers on your friends balcony. So you have taken a color image of the balcony (pots.jpg). From this image, you extract the red channel as an image and apply a median filter, with a square footprint of size 10. Finally, you threshold the median filtered image, so pixels with a value above 200 are set to foreground and the rest is set to background. How many foreground pixels are there in the final image?

Data: [All data for the exam can be download here.](#)

<https://designer.mcq.eksamen.dtu.dk/api/images/8ae658c5-dbc1-4515-a46d-7111564b02a7>

Choose one answer

- Between 8000 and 9000
- Between 6000 and 7000
- Do not know
- Between 5000 and 6000
- Between 10000 and 11000
- Between 7000 and 8000

