

02502 Image Analysis Exam Spring 2025

02502 Image Analysis Exam Spring 2025

Written exam, May 20, 2025

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/ba9da061-7277-4895-9026-8f6181d9e4d4>)

Pedestrian registration

In order to find and classify pedestrians in road images several components are needed. One component is a landmark based registration.

A first attempt involves applying a geometrical transformation using landmark annotation between humans. An algorithm is used to automatically identify a region of landmarks at the end of each limb of the human.

The figure below shows landmarks of found on two humans. These landmarks are encoded in an image called **CarTraining_Landmarks.png**, where **landmark #1** is marked with a region of pixels with value 1, **landmark #2** is marked with a region of pixel with value 2 and so on. The position of the landmark is defined as the center of mass of the region of pixels.

Human 1 has landmarks 1–4 and **Human 2** has landmarks 5–8.

There is correspondence between:


Landmark #1 and Landmark #5

Landmark #2 and Landmark #6

Landmark #3 and Landmark #7

Landmark #4 and Landmark #8

An optimal translation is computed that brings the landmarks of **Human 1** to match those of **Human 2**.



All data for the exam can be downloaded here

(<https://designer.mcq.eksamen.dtu.dk/api/images/ba9da061-7277-4895-9026-8f6181d9e4d4>)

What is the absolute value of the first component of the found translation vector?

Vælg en svarmulighed

- ☐ Between 170 and 175
- ☐ Between 165 and 170
- ☐ Between 160 and 165
- ☒ Between 155 and 160
- ☐ Do not know
- ☐ Between 175 and 180

What is the sum of squared distance between the four corresponding landmarks of the two humans after applying the optimal geometric translation?

Vælg en svarmulighed

- ☐ Between 3000 and 4000
- ☒ Between 2000 and 3000
- ☐ Between 1000 and 2000
- ☐ Between 5000 and 6000
- ☐ Between 4000 and 5000
- ☐ Do not know

Animal similarity

As a small part of a large system to analyze animal photos, there is a need for computing how similar two images are.

We have two images, **ImageA.png** and **ImageB.png**, that we wish to explore the similarity of.

The procedure is the following:

- Load the two images and convert them to grayscale
- Intensity-normalize the two images by setting the mean to zero and the standard deviation to one. Then, add 50 to each pixel. The normalization to a mean of 50 is done separately for each image
- Select every second pixel in the images for calculating the similarity index, and organize them as two vectors. The same procedure for selecting the pixels is to be used for both images

As the similarity index, we calculate the normalized correlation coefficient (NCC).

All data for the exam can be downloaded here

(<https://designer.mcq.eksamen.dtu.dk/api/images/ba9da061-7277-4895-9026-8f6181d9e4d4>)

The NCC can be expressed as the angle between the two vectors. What is this angle in degrees?

Vælg en svarmulighed

- ☐ Between 70 and 75
- ☐ Between 60 and 65
- ☒ Between 55 and 60
- ☐ Between 75 and 80
- ☐ Between 65 and 70
- ☐ Do not know



What is the NCC between the two images?

Vælg en svarmulighed


- ☐ between 0.40 and 0.45
- ☐ between 0.45 and 0.50
- ☒ Between 0.50 and 0.55
- ☐ Between 0.55 and 0.60
- ☐ Between 0.60 and 0.65
- ☐ Do not know

You perform a quality check of the histogram of **ImageA.png** using 50 bins after normalization and after sampling every second pixel. Which figure best correspond to the histogram?


Vælg en svarmulighed



☐




☐



☐ Do not know

○



Letter recognition

You are working on a letter recommendation system and test your system on a supplied photo of a Teabag: **TeaBag.png**

The system start by converting the image from color to grayscale and turning it into a binary image using a threshold of 0.6, where pixel above are set to foreground and the rest to background.

Secondly, a BLOB analysis is performed and the area and perimeters of all the BLOBS are computed.

All data for the exam can be downloaded here

(<https://designer.mcq.eksamen.dtu.dk/api/images/ba9da061-7277-4895-9026-8f6181d9e4d4>)

You filter all BLOBs by removing all BLOBs with an area below 100 or above 1000 and also removes BLOBs with a perimeter below 10 or above 200. You then visualize the results. What letters are left in the segmented image?

Vælg en svarmulighed

- ☐ BLA C U
- ☐ Do not know
- ☒ LAC C A T
- ☐ LAC C U RR
- ☐ L CC R T
- ☐ BLA K C U RR T

How many BLOBs are found in the binary image?

Vælg en svarmulighed

- ☐ Between 1100 and 1300
- ☒ Between 900 and 1100
- ☐ Between 700 and 900
- ☐ Between 1300 and 1500
- ☐ Between 500 and 700
- ☐ Do not know

How many BLOBs in the binary image have a perimeter above 15?

Vælg en svarmulighed

- ☐ Between 30 and 50
- ☒ Between 110 and 130
- ☐ Between 70 and 90
- ☐ Between 50 and 70
- ☐ Between 90 and 110
- ☐ Do not know

Change detection in video

You are developing a small change detection algorithm and have acquired a set of six test images for simulating a video-stream input. The test images are named in chronological order:

movie000.jpg to movie005.jpg

All images are converted to gray scale using **rgb2gray** before further processing.

Your algorithm works by estimating a slowly changing background image. The algorithm starts by setting movie000.jpg as the first background estimate. For each new frame the background is updated using $\alpha = 0.80$.

You start by simulating a short video using **movie000.jpg** to **movie004.jpg** and inspecting the estimated background image at that point.

Secondly, you use **movie005.jpg** as a new frame. This frame is compared to the background image using the pixel-wise absolute difference.

The absolute difference image is converted to a binary image using a threshold of 0.2. Finally, a BLOB analysis is performed on the binary image to find the number of connected moving objects.

All data for the exam can be downloaded here

(<https://designer.mcq.eksamen.dtu.dk/api/images/ba9da061-7277-4895-9026-8f6181d9e4d4>)

How many BLOBs are found in the binary image?

Vælg en svarmulighed

- ☒ Between 80 and 90
- ☐ Between 70 and 80
- ☐ Between 90 and 100
- ☐ Between 60 and 70
- ☐ Between 50 and 60
- ☐ Do not know



After estimating the background image using the first images, you compute the maximum pixel value in the image. What is this value?

Vælg en svarmulighed

- ☐ Do not know
- ☐ Between 0.75 and 0.80
- ☒ Between 0.90 and 0.95
- ☐ Between 0.95 and 1.0
- ☐ Between 0.85 and 0.90
- ☐ Between 0.80 and 0.85


You visually inspect the binary change image. How does it look like?

Vælg en svarmulighed


☐☐

☐ Do not know


○



●



○



How many pixels are classified as changed in the binary image?

Vælg en svarmulighed

- ☐ Between 3000 and 3500
- ☐ Between 1500 and 2000
- ☐ Do now know
- ☒ Between 3500 and 4000
- ☐ Between 2000 and 2500
- ☐ Between 2500 and 3000

One of the following statements is not correct. Which one?

Vælg en svarmulighed

- ☐ Do not know
- ☐ Haar features can be computed fast using an integral image.
- ☐ Procrustes alignment can be used to do groupwise registration of several point clouds
- ☒ The Hough transform can find the shortest curved path through an image
- ☐ The Prewitt filter can be used to detect high gradients in an image
- ☐ A pixel with RGB values of [255, 255, 255] is white

Satellite image transformations

A company has recorded a satellite image named **Satellite_A.png**.

A geometric transformation must be applied to obtain the correct orientation of the landscape.

The geometric transformation is done by first applying a translation matrix A_T of 5 pixels along both the positive X and Y axes, then a clockwise rotation of 10 degrees (i.e., A_R), and finally, a scaling factor of 50% (i.e., A_S) is applied. The transformed image is named **Img_B**.

The company has also provided some *points of interests* on **Satellite_A.png**. These are provided as a binary mask, **ROI_A.png**, that contains three pixels marked with value 1. The rest of the pixel have value 0.

The first task is to apply the geometric transform to the *point of interests* so their positions can be found in **Img_B**.

The pixel values in **Img_B** are computed using bilinear interpolation of the values in **Satellite_A.png**. A binary mask, **ROI_B.png**, corresponding to **Img_B** is provided. We would like to compute the interpolated values, where the value of **ROI_B.png** is 1. So the task is to find the non-zero pixel values in **ROI_B.png** and sample their values correctly in **Satellite_A.png** taking the transformation into account.

All data for the exam can be downloaded here

(<https://designer.mcq.eksamen.dtu.dk/api/images/ba9da061-7277-4895-9026-8f6181d9e4d4>)

What is the value of the element on position (0,0) in the final transformation matrix?

Vælg en svarmulighed

- ☐ Between 0.55 and 0.65
- ☐ Between 0.65 and 0.75
- ☐ Do not know
- ☐ Between 0.15 and 0.25
- ☒ Between 0.45 and 0.55
- ☐ Between 0.75 and 0.85

After sampling the interpolated values from **Satellite_A.png** given the positions marked in **ROI_B.png**, the average of the values is computed. This average value is:

Vælg en svarmulighed

- ☐ Between 170 and 175
- ☐ Do not know
- ☐ Between 180 and 185
- ☐ Between 185 and 190
- ☒ Between 175 and 180
- ☐ Between 165 and 170

After applying the transformation to the *points of interest*, you compare their transformed position with some given reference points. Which set of points are they most similar to (the order might be shuffle)?

Vælg en svarmulighed

- ☐ (174, 47), (285, 159), (45, 12)
- ☒ (70, 60), (97, 134), (185, 99)
- ☐ Do not know
- ☐ (84, 31), (135, 91), (206, 28)
- ☐ (38, 125), (123, 11), (232, 32)
- ☐ (4, 333), (45, 181), (190, 376)

Statistical analysis of chest CT scans

A new start-up has gained access to a database of CT scans of the chest and would like to develop a screening tool for vertebra bone fracture. They expect that a statistical analysis of single slices of CT scans can be helpful. You help exploring the possibilities.

You have been given the following training CT slices:

```
training_images = ["1-100.dcm", "1-110.dcm", "1-120.dcm", "1-130.dcm", "1-140.dcm",  
"1-150.dcm", "1-160.dcm", "1-170.dcm", "1-180.dcm", "1-190.dcm", "1-200.dcm"]
```

And two test slices: **1-105.dcm** and **1-115.dcm**

You start by computing the average image from the training images. Secondly, a binary segmentation of the average image is made by setting pixel with values in]100,500[to foreground and the rest to background. A morphological closing with disk shaped structuring element of radius=5 is applied to the binary image and a BLOB analysis is performed. Finally, only the BLOBs with an area in [2000,5000] and a perimeter in [400, 600] are kept.



To explore the statistical properties of the images, you do an image based principal component analysis (PCA) of the training images, where you only compute the first 5 principal components. You do not normalize the image before PCA.

All data for the exam can be downloaded here


(<https://designer.mcq.eksamen.dtu.dk/api/images/ba9da061-7277-4895-9026-8f6181d9e4d4>)

After doing the PCA, you plot the projected values of the training images on the first and the second principal component. How does this plot look like?


Vælg en svarmulighed

☐☐


☐



☐



☐ Do not know



How many BLOBs do you find before you do the area and perimeter filtering of the BLOBs?

Vælg en svarmulighed

- ☐ Between 10 and 15
- ☐ Between 20 and 25
- ☐ Between 25 and 30
- ☐ Between 5 and 10
- ☒ Between 15 and 20
- ☐ Do not know

After BLOB filtering with area and perimeter the remaining BLOB(s) is in a binary image. This image is compared to the test slice **1-115.dcm** using the DICE score. Before computing the DICE score, the test slice is made binary by by setting pixel with values in **]100,500[** to foreground and the rest to background.

What is the DICE score?

Vælg en svarmulighed

- ☐ Between 0.20 and 0.30
- ☐ Between 0.30 and 0.40
- ☐ Do not know
- ☐ Between 0.50 and 0.60
- ☒ between 0.10 and 0.20
- ☐ Between 0.40 and 0.50

The test slice **1-105.dcm** should be compared to the training images. This is done by projecting the test slice into the 5 dimensional PCA space and finding the closest training image in this space. What image is that?

Vælg en svarmulighed

- ☐ 1-190.dcm
- ☐ 1-180.dcm
- ☒ 1-160.dcm
- ☐ 1-120.dcm
- ☐ 1-100.dcm
- ☐ Do not know


What is the perimeter of the BLOB that is found after the area and perimeter filtering of the BLOBs?

Vælg en svarmulighed

- ☐ Between 350 and 400
- ☐ Do not know
- ☒ Between 400 and 450
- ☐ Between 550 and 600
- ☐ Between 500 and 550
- ☐ Between 450 and 500

We are to establish an optimization procedure based on the gradient descent algorithm and have taken over the project from a colleague. The colleague managed to find an ideal parameter setting for the gradient descent, and you have an image that proves it works shown below. You know that the expression to optimize spans a two-parameter space :

$$f(x_1, x_2) = 12x_1^2 + x_1x_2 + 5x_2^2$$




Which of the following step-size constants best aligns with the one used in the gradient descent algorithm?

Vælg en svarmulighed

- ☐ 0.01
- ☐ Do not know
- ☒ 0.02
- ☐ 0.6
- ☐ 0.95

○ 1.3

You are working with quality assurance of printing of industrial products. As a test you want to see the color quality of printing on tea bags (TeaBag.png). You have developed a simple method that produces the image below. What are the steps in your pipeline?



All data for the exam can be downloaded [here](https://designer.mcq.eksamen.dtu.dk/api/images/ba9da061-7277-4895-9026-8f6181d9e4d4) (<https://designer.mcq.eksamen.dtu.dk/api/images/ba9da061-7277-4895-9026-8f6181d9e4d4>).

Vælg en svarmulighed

- ☐ Load image, rotate 30 degrees, Prewitt filter, threshold with 113
- ☒ Load image, RGB thresholding with $R < 100$, $G < 100$, and $B > 100$
- ☐ Load image, RGB to grayscale, Otsu's optimal threshold
- ☐ Do not know
- ☐ Load image, gradient filter, threshold with 23
- ☐ Load image, RGB to HSV, threshold with 230 on V component

Car navigation

Visionary is a new start-up company developing an object recognition system to be used in car systems. You are asked to develop the basic components for the system. The first component should segment a scene into relevant objects. You plan to first try applying a parametric classifier.

CarTraining.png is a standard image for the evaluation of such systems and represents a situation in traffic captured by a camera.

Training examples for each of the three classes have been provided to you:

- **Class 1** is non-humans (**Class1.png**)
- **Background** represents the background (**BG.png**)
- **Class 2** is humans (**Class2.png**)

Each class is represented as an image-mask with the same size as **CarTraining.png**, where non-zero values indicates where the relevant class pixel values can be sampled from **CarTraining.png**.

You use the training image and the mask images to train a three-class parametric classifier.

NOTE: The background class is between class 1 and class 2 and the class range should only capture the single range between class 1 and class 2. Start by computing the class range before applying it to the image.


All data for the exam can be downloaded here


(<https://designer.mcq.eksamen.dtu.dk/api/images/ba9da061-7277-4895-9026-8f6181d9e4d4>)


To visually test your classifier, you load the **CarTraining.png** image and set all pixels that are classified as background to 255. What is the resulting image?


Vælg en svarmulighed


☐ Do not know











You compute the class range for the background class. What is the closest class range?

Vælg en svarmulighed

- ☐ [40, 112]
- ☒ [123, 175]
- ☐ [10, 50]
- ☐ [97, 190]
- ☐ [156, 210]
- ☐ Do not know

