

02502 Image Analysis Exam Fall 2023

Written exam, December 11, 2023

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 or similar) to the "mellemregning" /appendix part of the exam.

Video image analysis frame rate

You have designed a new video image analysis system for analyzing the behavior of turtles in a zoo.

A camera is connected to a computer using a connection that can transfer 35 megabytes of data per second. Each image is an RGB image (one byte per channel) with a size of 2400 x 1200 pixels. Your image analysis algorithm takes 130 milliseconds to analyze one image.

What is the system frame rate?

Vælg en svarmulighed

- 9 frames per second
- Do not know
- 4 frames per second
- 16 frames per second
- 21 frames per second
- 12 frames per second

Pixelwise operations

You have taken a photo from your last holidays and you would like to see if you enhance or modify the appearance of this photo. For this you are experimenting with some pixelwise operations. You do:

1. Load the image as an RGB image.
2. Convert the image to gray scale. Now the image is a floating point image where the values are in the range of $[0, 1]$.
3. Do a linear gray scale histogram stretch. The new image should have a minimum value of 0.2 and a maximum value of 0.8.
4. Computing the average value of the histogram stretched image.
5. Use the **prewitt_h** filter to extract edges in the image.
6. Computing the maximum absolute value of the Prewitt filtered image.
7. Creating a binary image from the histogram stretched image by using a threshold with a value that is equal to the average value of the image.
8. Computing the number of foreground pixels in the binary image.

The photo for the questions: ardeche_river.jpg

(<https://designer.mcq.eksamen.dtu.dk/api/images/7ede9aad-0399-4428-85c3-6144d3ea9e17>)

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

Vælg en svarmulighed

- Do not know
- Between 400000 and 500000
- Between 300000 and 400000
- Between 200000 and 300000
- Between 100000 and 200000
- between 0 and 100000

What is the maximum absolute value in the Prewitt filtered image?

Vælg en svarmulighed

- 0.66
- 0.42
- 0.34
- 0.53
- 0.12
- Do not know

What is the average value of the histogram stretched image?

Vælg en svarmulighed

- 0.23
- 0.37
- 0.48
- 0.53
- Do not know
- 0.65

3D medical image registration

A medical doctor has sought your expertise to conduct image analysis on 3D multi-modal MRI scans. The multi-modality MRI scans involve the acquisition of two distinct MRI brain volumes, namely `ImgT1_v1.nii.gz` and `ImgT1_v2.nii.gz`. The image analysis will employ affine image registration of the two MRI scans, combining 3D geometrical transformations including translation (T_x, T_y, T_z), rotation (R_x, R_y, R_z), scaling (S_x, S_y, S_z), and a shearing (S_{xy}, S_{xz}, S_{yz}) matrix. We define that the 3D rotation matrix follows the **right-hand rule**:

$$R_x(\theta) = \begin{pmatrix} 1 & 0 & 0 \\ 0 & \cos\theta & -\sin\theta \\ 0 & \sin\theta & \cos\theta \end{pmatrix} \quad R_y(\theta) = \begin{pmatrix} \cos\theta & 0 & \sin\theta \\ 0 & 1 & 0 \\ -\sin\theta & 0 & \cos\theta \end{pmatrix} \quad R_z(\theta) = \begin{pmatrix} \cos\theta & -\sin\theta & 0 \\ \sin\theta & \cos\theta & 0 \\ 0 & 0 & 1 \end{pmatrix}$$

You can use the function “`array = sitk.GetArrayFromImage(sitkImage)`” to get a numpy array from a SimpleITK image object. After creating a `sitk.AffineTransform(3)` object, you can manually set up a translation using the function `transform.SetTranslation(...)`.

Unfortunately, the subject did not lay still during the scanning session, and the two MRI scans are to be registered.

The data for the questions: `ImgT1_v1.nii.gz`

(<https://designer.mcq.eksamen.dtu.dk/api/images/3284e8f1-2d74-46d2-a065-42983c29afc3>)

and `ImgT1_v2.nii.gz` (<https://designer.mcq.eksamen.dtu.dk/api/images/6e82e33d-21ef-44d5-a7fb-cd9c39716c87>)

Python tip: When you manually should set the **translation**, you should use a **list** and not a numpy array.

A collaborator sent you an affine matrix A , generated by SimpleITK. What is the result of applying A to `ImgT1_v1`?


Set the centre of rotation in the middle of the image.

The following plots are generated with the function `imshow_orthogonal_view(sitk_image, origin = None)` from exercise 9, where `sitk_image` is your transformed image.


$$A = \begin{pmatrix} 0.98 & -0.16 & 0.17 & 0 \\ 0.26 & 0.97 & 0 & -15 \\ -0.17 & 0.04 & 0.98 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}$$

Vælg en svarmulighed


Axial



Coronal




Sagittal




Do not know


Axial



Coronal




Sagittal




○


Axial



Coronal




Sagittal




○


Axial



Coronal




Sagittal




○


Axial



Coronal



Sagittal



You try to automatically estimate the registration parameters using a rigid (Euler) 3D transformation using SimpleITK.

The fixed image is `ImgT1_v1` and the moving image is `ImageT1_v2`.

You set the center of rotation to the middle of the image.

What are the estimated rigid geometrical transformation parameters, as found by SimpleITK (Rotations are given in degrees)?

Vælg en svarmulighed

- $(T_x, T_y, T_z, R_x, R_y, R_z, S_x, S_y, S_z, S_{xy}, S_{xz}, S_{yz}) = (50, 0, 30, 0, 12, 20, 1.2, 1, 1, 0, 0, 0)$
- Do not know
- $(T_x, T_y, T_z, R_x, R_y, R_z, S_x, S_y, S_z, S_{xy}, S_{xz}, S_{yz}) = (0, 0, 30, 0, 12, 0, 1, 1, 1, 0, 0, 0)$
- $(T_x, T_y, T_z, R_x, R_y, R_z) = (50, 0, 30, 0, 12, 20)$
- $(T_x, T_y, T_z, R_x, R_y, R_z) = (6, 0, 29, 0, 12, 0)$
- $(T_x, T_y, R_x, R_y) = (50, 0, 0, 12)$

You try to register the volumes manually. Considering the `ImgT1_v1` as the fixed image and the `ImgT1_v2` as the moving image. You do the following steps:

1. Apply -20 degrees roll to `ImgT1_v2` using the default center of rotation.
2. Make a brain mask including only voxel intensities values > 50 in image `ImgT1_v1`.
3. Compute the mean squared error as a similarity metric within the brain mask.

What is the mean squared error between the transformed `ImgT1_v2` and `ImgT1_v1`?

Vælg en svarmulighed

- Do not know
- Between 10000 and 20000
- Between 20000 and 30000
- Between 30000 and 40000
- Between 50000 and 60000
- Between 40000 and 50000

Change detection and color space conversions

We are trying to implement a first prototype of a video based change detection system. The input images are RGB images and we would like to test if converting the images to the HSV color space could aid in the change detection.

Your first version of the system is working like this:

1. You have two images: `frame_1.jpg` and `frame_2.jpg`
2. After reading the images you convert them to the HSV color space using `rgb2hsv`
3. You extract the **S** channel of both the HSV images and scale the channel with 255
4. You compute the absolute difference image between the two S images
5. You compute the average value and the standard deviation of the values in the difference image
6. You compute a threshold as the average value plus two times the standard deviation
7. You compute a binary *change image* by setting all pixel in the difference image that are higher than the threshold to foreground (1) and the rest of the pixels to background (0)
8. You compute the number of changed pixels
9. You perform a BLOB analysis on the binary change image
10. You find the BLOB with the largest area

Date for the question: ChangeDetection.zip

(<https://designer.mcq.eksamen.dtu.dk/api/images/ec9114d1-91d8-4783-a65f-c34c30ea57fc>)

What is the area (in pixels) of the largest region that you find using BLOB analysis?

Vælg en svarmulighed

- between 3000 and 6000
- Between 6000 and 9000
- Between 9000 and 12000
- Between 12000 and 15000
- Do not know
- Between 15000 and 18000

What is the threshold you use to convert the difference image into a binary image?

Vælg en svarmulighed

- 78
- 83
- Do not know
- 34
- 99
- 56

What is the number of changed pixels that you find in the binary image?

Vælg en svarmulighed

- Do not know
- Between 25000 and 30000
- Between 15000 and 20000
- Between 10000 and 15000
- Between 5000 and 10000
- Between 20000 and 25000


Hough transform

A green matchstick person is illustrated below using five straight lines in a Cartesian coordinate system (x, y) with arbitrary units.

One long line is for the body, two shorter lines represent the arms, and two longer lines represent the legs. At the end of each stick, there is a dot representing the head, the two hands, and the two feet, respectively.

Each line can be represented as a point set $(\text{Theta}, \text{Rho})$ in the Hough space, where Theta is an angle in degrees in the range: $]-90, 90]$ and Rho is a vector length.

Which of the sets of Hough-space point sets $(\text{Theta}, \text{Rho})$ represent the green matchstick person?



Vælg en svarmulighed

- Do not know
- $(-45, 4.9), (-45, -0.7), (0, 3), (-45, 0.7), (45, 3.5)$
- $(90, 1.4), (45, 5.7), (0, 3), (135, -0.7), (110, 3.5)$
- $(135, -0.7), (45, 4.9), (0, 3), (135, -0.7), (45, 3.5)$
- $(-45, -1.4), (45, 5.7), (0, 3), (-45, 0.7), (45, 3.5)$
- $(0, 2), (-45, -0.7), (45, 4.9), (135, -0.7), (90, 3.5)$

Heart analysis on computed tomography scans

A radiologist has asked you for help to assess the function of the heart as seen on a computed tomography scan. The main interest is the heart muscle (myocardium) and the blood inside the myocardium (blood). We try first on a single DICOM slice (1-001.dcm), where we also have some manual annotations. We have a region-of-interest (ROI) in the myocardium (MyocardiumROI.png), a region-of-interest in the blood (bloodROI.png) and a full ground truth annotation of the blood (bloodGT.png).

Our initial analysis consist of:

- Read the DICOM file and get the pixel values (as Hounsfield units).
- Extract the pixel values of the ROI of the myocardium and the blood using the manual annotations.

To make an automatic segmentation of the blood, the average Hounsfield unit value (μ) and the standard deviation (σ) of the Hounsfield units in the blood are computed using the manual annotations.

A class range is computed, so all pixels with:

$$\mu - 3 \cdot \sigma < \text{value} < \mu + 3 \cdot \sigma$$

Are set to foreground (1) and the rest of the pixels are set to background (0).

To clean the binary image, a morphological closing with disk shaped structuring element of radius 3 is performed followed by a morphological opening with a structuring element of radius 5. A BLOB analysis is performed on the resulting image and the area of all the BLOBs are computed. The BLOBs are filtered so only BLOBs with an area larger than 2000 pixels and an area less than 5000 pixels are kept.

Finally, the found BLOB(s) are compared to the manual segmentation found in BloodGT.png and the DICE score is computed.

The data for the question: HeartCT.zip

(<https://designer.mcq.eksamen.dtu.dk/api/images/dd016337-9dda-421e-b6c3-8bc4d9e33ce5>)

A class range is computed, so all pixels with:

$$\mu - 3 \cdot \sigma < \text{value} < \mu + 3 \cdot \sigma$$

Are set to foreground (1) and the rest of the pixels are set to background (0).

What is the class range (in Hounsfield units)?

Vælg en svarmulighed

- Do not know
- [219, 256]
- [172, 213]
- [234, 342]
- [198, 312]
- [210, 296]

How many BLOBs are found before they are filtered using the area?

Vælg en svarmulighed

- Between 4 and 7
- Between 8 and 11
- Between 16 and 19
- Between 12 and 15
- Do not know
- Between 1 and 3

You compare your found blood estimate with the manual drawn blood estimate in BloodGT.png. What is the DICE score?

Vælg en svarmulighed

- 0.87
- 0.85
- 0.78
- 0.81
- 0.94
- Do not know

To be able to do a pixelwise classification, we would like to use a minimum-distance classifier to separate the myocardium and the blood. What is the class limit between the two, when using a minimum-distance classifier?

Vælg en svarmulighed

- 89
- 145
- 165
- 157
- 178
- Do not know

Principal component analysis on pistachio nuts

A manufacturer of pistachio nuts wants to be able to detect stones and other dirt when nuts are being sorted. The factory has installed a camera system and has recorded 200 images of nuts. An image analysis software has extracted 12 features per nuts, named: AREA, PERIMETER, MAJOR_AXIS, MINOR_AXIS, ECCENTRICITY, EQDIASQ, SOLIDITY, CONVEX_AREA, EXTENT, ASPECT_RATIO, ROUNDNESS and COMPACTNESS.

To remove redundancy in the data, a principal component analysis (PCA) is performed. Since the measurements are of very different scales, they are scales using the standard deviation before the PCA:

- 1) Load the data from the pistachio_data.txt file
- 2) Subtract the mean from the data
- 3) Compute the standard deviation of each measurement (for example the standard deviation of the AREA).
- 4) Divide each measurement by its own standard deviation.
- 5) Do the PCA

The data for the question pistachio_data.txt

(<https://designer.mcq.eksamen.dtu.dk/api/images/aa2d6e7f-c108-4987-aff6-94a03a033b4e>)

To find out how far away from *normal* a pistachio nut is, we can look at position of the nut in PCA space. We start by projecting the measurements of the first nut onto the principal components. The sum of squared projected values are now computed since they give a measure of how far away from the origin the nut is. What is the sum of squared projected values?

Vælg en svarmulighed

- 19.6
- Do not know
- 78.3
- 102.3
- 34.2
- 54.4

Which of the measurements has the smallest standard deviation?

Vælg en svarmulighed

- Area
- Compactness
- Do not know
- Solidity
- Eccentricity
- Roundness

After PCA has been computed, we would like to know how many components are needed to explain at least 97% of the total variation in the dataset. How many are needed?

Vælg en svarmulighed

- 2
- 6
- 3
- 4
- 5
- Do not know

The covariance matrix is computed before doing the PCA. What is the maximum absolute value found in the covariance matrix?

Vælg en svarmulighed

- 0.71
- 1.43
- 0.87
- Do not know
- 0.04
- 1.00

Landmark based registration

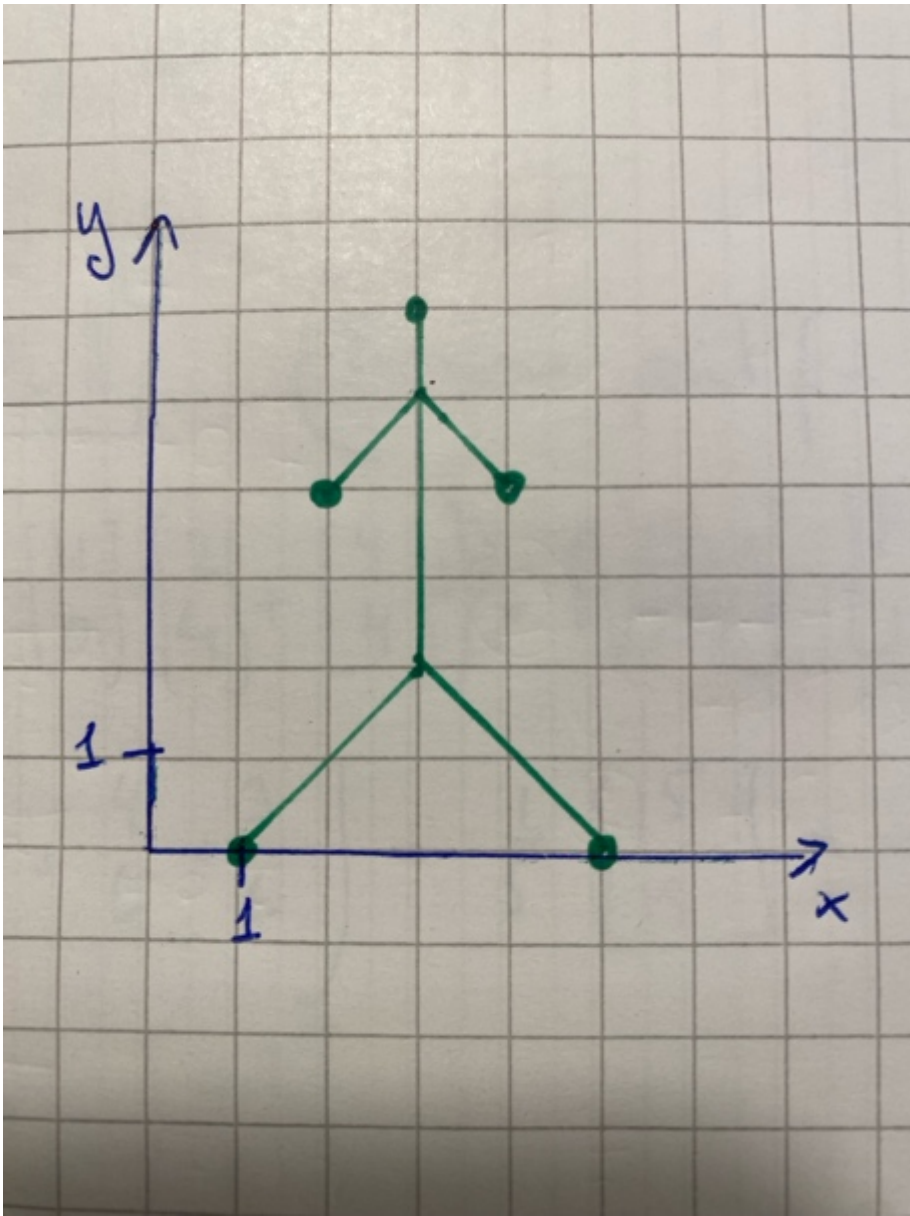
Below are two figures of a stick person. One where the person is standing still and one where the person is running.

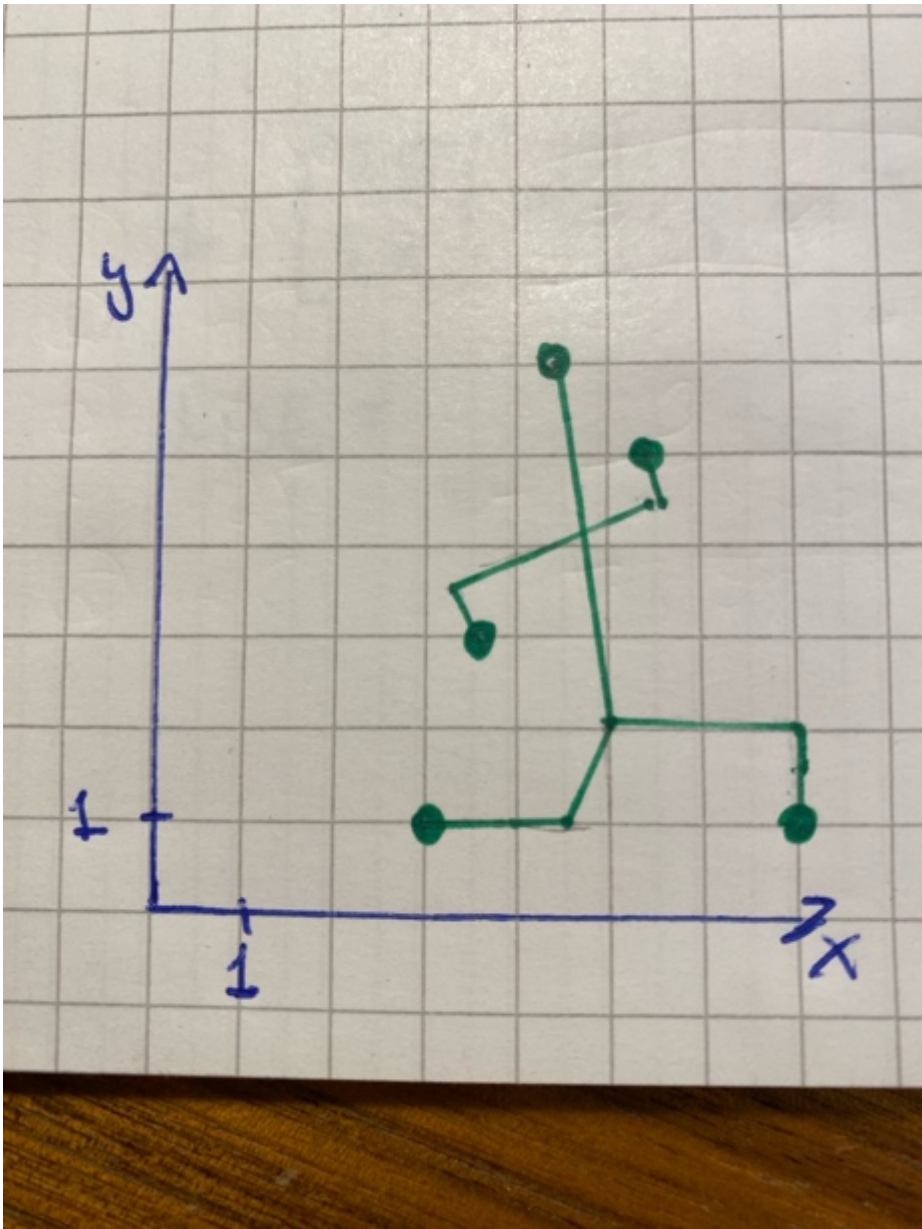
The dots representing the endpoints of the sticks i.e. the head, arms, and feet have changed position (in the running person) compared to the positions in standing person. We use the dots as landmarks to find the 2D registration that brings the landmarks from the template (the running person) over in the reference (the standing person).

We use a sum-of-squared-distances objective function to first compute how well the landmarks are aligned before any transformation.

Secondly, we compute the optimal translation $(\Delta X, \Delta y)$ that aligns the two sets of landmarks.

Finally, we try a similarity transform (translation, scaling and rotation) to align the two sets of landmarks.





What is the optimal translation that aligns the two sets of landmarks when using a transformations that only consist of translation?

Vælg en svarmulighed

- (3.1, 6.4)
- (-1.7, -0.4)
- (-5.2, 3.4)
- (1.8, 0.7)
- (2.7, -2.8)
- Do not know

What is the sum-of-squared-distances between the landmarks before any transformation?

Vælg en svarmulighed

- 13.25
- 15.75
- 18.75
- 21.45
- Do not know
- 31.25

After computing the similarity transform that optimally aligns the two sets of landmarks the rotation of the transform is shown using **tform.rotation**, where tform is the similarity transform. What is the absolute value of the found rotation (in degrees)?

Vælg en svarmulighed

- 2
- 12
- 4
- 5
- 8
- Do not know

Aquarium design using principal component analysis

Your friend has bought a new aquarium and is very excited to buy the fish that should live there. You suggest to use image based principal component analysis (PCA) to find the optimal selection of fish. In the local fish shop, you have taken photos of 10 different types of fish:

("discus.jpg", "guppy.jpg", "kribensis.jpg", "neon.jpg", "oscar.jpg", "platy.jpg", "rummy.jpg", "scalare.jpg", "tiger.jpg", "zebra.jpg")

First you compute an image of the average fish and then you use PCA from **sklearn.decomposition** to compute 6 principal components of the 10 photos.

The photos of the fishes: Fish.zip

(<https://designer.mcq.eksamen.dtu.dk/api/images/b136a28d-ed6d-4808-8430-0a2afb599798>)

After computing the PCA, you compute how much is component explains of the total variation. How much of the total variation is explained by the first two components?

Vælg en svarmulighed

- 61%
- 78%
- 34%
- 46%
- Do not know
- 53%

Your friend really likes the look of the neon fish and the guppy. To get an idea of how visually different they are, you compute the pixelwise sum of squared differences between the photo of the neon fish and the guppy. What is this value?

Vælg en svarmulighed

- Between 70.000.000 and 80.000.000
- Do not know
- Between 40.000.000 and 50.000.000
- Between 50.000.000 and 60.000.000
- Between 30.000.000 and 40.000.000
- Between 60.000.000 and 70.000.000

Your friend find the neon fish to be a personal favorite and decide to buy some of them. He also wants to have a flock of fish that are as visually different as possible from the neon fish.

You project all the fish photo on the principal components and find the fish that are furthest away from the neon fish in PCA space. Which fish should your friend choose?

Vælg en svarmulighed

- Kribensis
- Discus
- Oscar
- Tiger
- Do not know
- Platy

