Contact:
Christian Schellewald
Mobile Phone: 94123453

Examination in Course TDT4265

Computer Vision (Datasyn)

19th May 2012
Time: 09:00-13:00

Computer Vision 2012

Time allowed: 4 hours = 240 minutes.
No written or printed materials allowed.
No calculator allowed.

The point scores for each question part are shown in the text. Use these point scores to help you budget your time and effort in answering.

Total number of pages: 15
Total point score: 110

READ EACH QUESTION THROUGH CAREFULLY BEFORE BEGINNING YOUR ANSWER. ANSWER ALL QUESTIONS. SHOW ALL YOUR WORKING.
(If you find it convenient, you can answer the questions directly on the exam pages.)

1 Theme: Image processing fundamentals

1.1 Filters

1.1.1 [2 points] Provide a binomial approximation of a 1D Gaussian filter with 5 elements and give the normalization factor.

Answer:

Normalization factor:

1.1.2 [2 points] Padding: Name four options to extend the boundary of an array for convolution.

Answer:

1.1.3 [3 points] Use your favorite padding method for the 1d array and convolve it with the given filter mask.

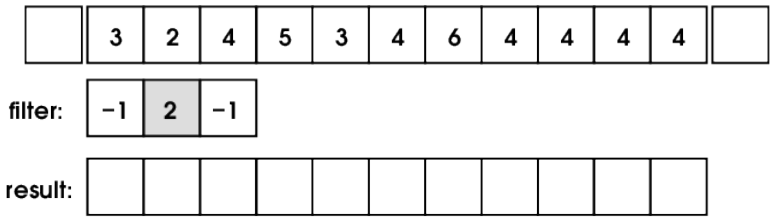

1.1.4 [2 points] When is correlation and convolution the same operation? (In other words: When is the output of both operations the same for arbitrary arrays?)

Answer:

1.1.5 [2 points] Which kind of filters are typically normalised to one (1.0) and which kind of filters have typically elements that sum up to zero (0.0)?

Answer:

1.1.6 [2 points] What happens with a grey value image if you filter it with a filter-mask that has a sum of elements larger than one ($s > 1.0$) or smaller than one ($0 \leq s < 1.0$)?

Answer:

1.1.7 [3 points] Determine the filter mask

Which filter mask defines the following operation? What size has your obtained filter mask? I_F is the filtered version of the image I. (Give the mask as an array.)

$$I_F(i,j) = I(i-1,j-1) + 2I(i,j-1) + I(i+1,j-1) - I(i-1,j+1) - 2I(i,j+1) - I(i+1,j+1)$$

Answer:

1.2 Match the shown filters with their corresponding filter results

An original image which is shown as part of figure 2 is filtered by six different filter masks which are depicted in figure 1. The randomly permuted filtering results are shown in figure 2. For the images and filter-masks the smallest value (which can be negative) is shown black while the largest value appears white. The outside boundary of the masks and filtered images represents the zero value.

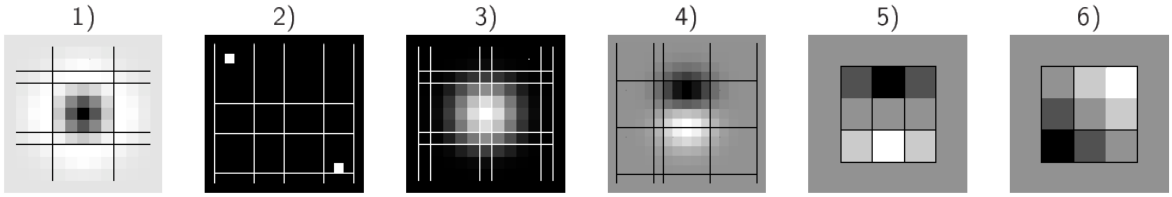

Figure 1: Six filters

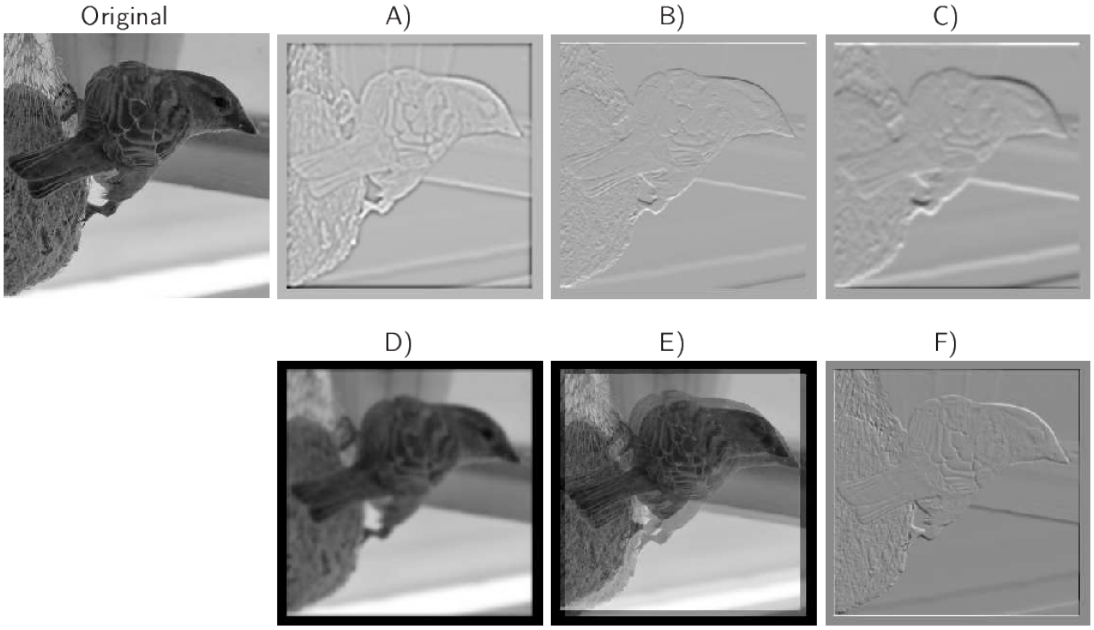

Figure 2: Six filtered images

1.2.1 [3 points] Determine which masks were used to compute the six output images. Use the second column in the answer-table below to assign a single filtered image to the mask number given in the first column. Briefly indicate a reason for your choice.

Answer:

Mask:	Filter result (A-H):	reason:
1)		
2)		
3)		
4)		
5)		
6)		

1.2.2 [2 points] Which of the filters is a Gaussian , LoG (Laplacian of Gaussian), derivative of Gaussian or a Sobel filter?

Answer:

Filter name	Filter number
Gaussian	
LoG	
derivative of Gaussian	
Sobel	

2 Theme: Transformations

2.1 Projections

The perspective projection (with focal length f) of a 3D point $(x, y, z)^\top$ to the image plane $(u, v)^\top$ is:

$$(u, v)^\top = \left(\frac{xf}{z}, \frac{yf}{z}\right)^\top$$

2.1.1 [2 points] How do you obtain an orthographic projection (=parallel projection) from the perspective projection?

Answer:

2.1.2 [1 points] Give the orthographic projection of the 3D point $(12, 24, 36)^T$.

Answer:

$$(u, v)^\top =$$

2.2 Interpolation

Assume you want to interpolate a grey value between the discrete pixel positions at the location shown in figure 3.

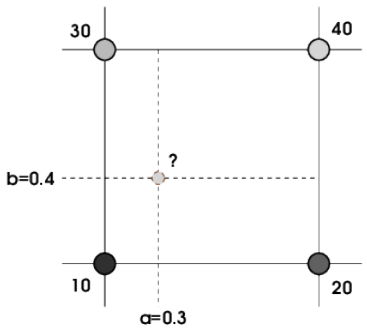

Figure 3: Perform a bi-linear interpolation.

2.2.1 [3 points] Perform a bi-linear interpolation for the gray value at the (fractional) location $(a = 0.3, b = 0.4)$ between the four corner pixels with grey values 30, 40, 20 and 10 (enumerated in clockwise order) . Which grey value do you compute?

(Hint: You do not need to recall the bi-linear interpolation formula and can obtain the result stepwise.)

Answer:

interpolated greyvalue=

3 Theme: Filtering in the Frequency Domain

3.1 Fourier transformation

3.1.1 [4 points] Can you match the shown image-spectrum pairs ?

The eight images in figure 4 are to be matched to the eight (absolute) images of their Fourier transformations ($|\mathcal{F}\{f\}|$) that are shown in figure 5 . Provide a short reason for each selected pair (Why did you match the particular pair? e.g. recall 1D Fourier transforms, and consider symmetries).

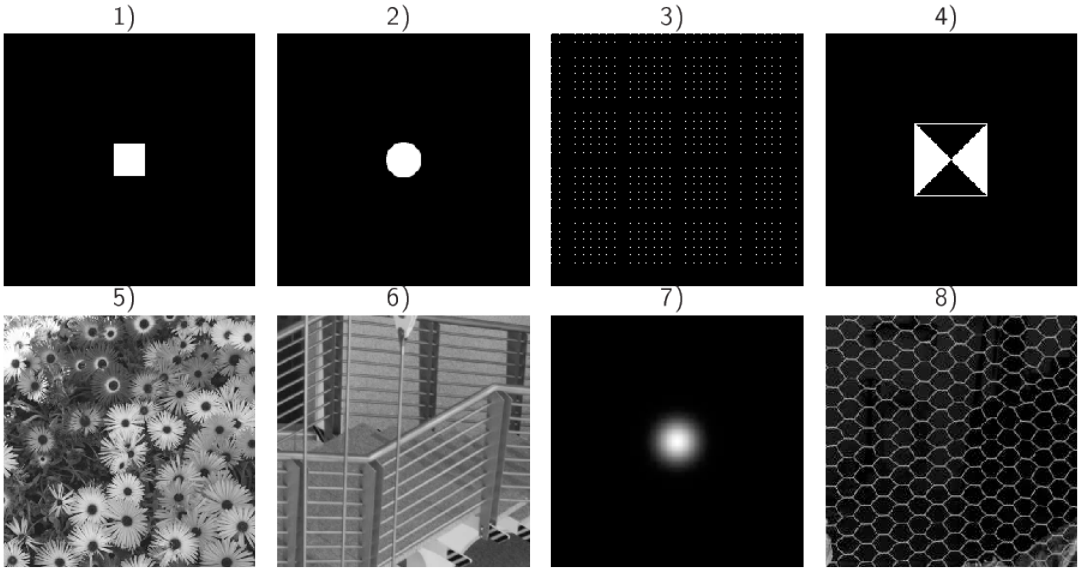

Figure 4: Eight images

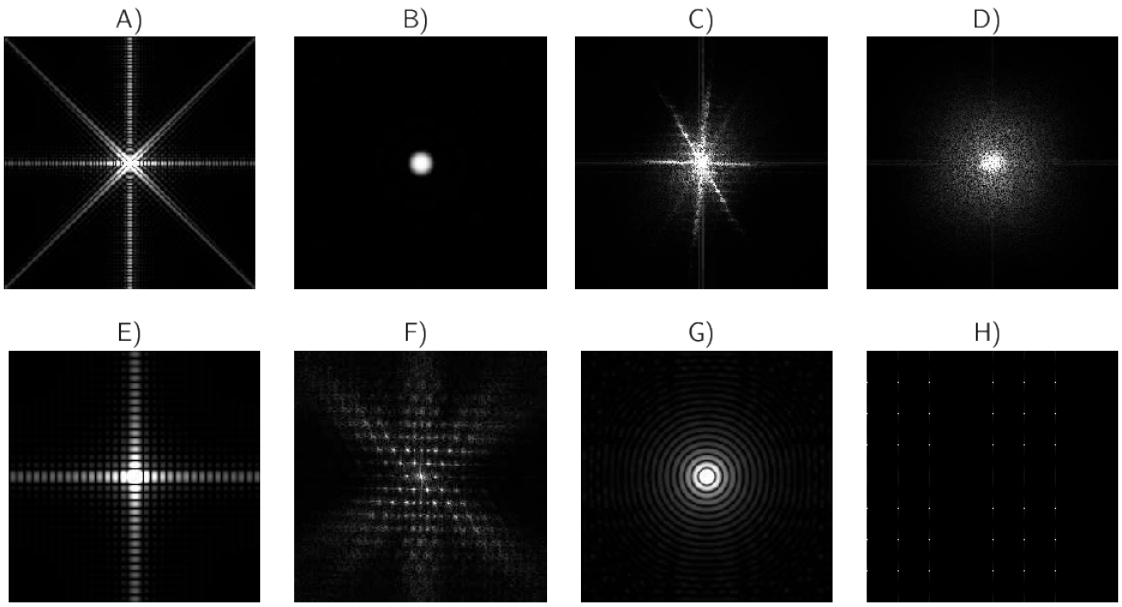

Figure 5: $|\mathcal{F}\{f\}|$ Amplitude of the Fourier transform of the eight images

Answer:

Image:	match (A-H):	reason:
1)		
2)		
3)		
4)		
5)		
6)		
7)		
8)		

3.2 Filtering

3.2.1 [3 points] Draw a sketch of a low-, high-, and band-pass filter into the shown frequency coordinate systems.

Indicate areas of high values with 1.0 and areas with low values with 0.0.

Answer:

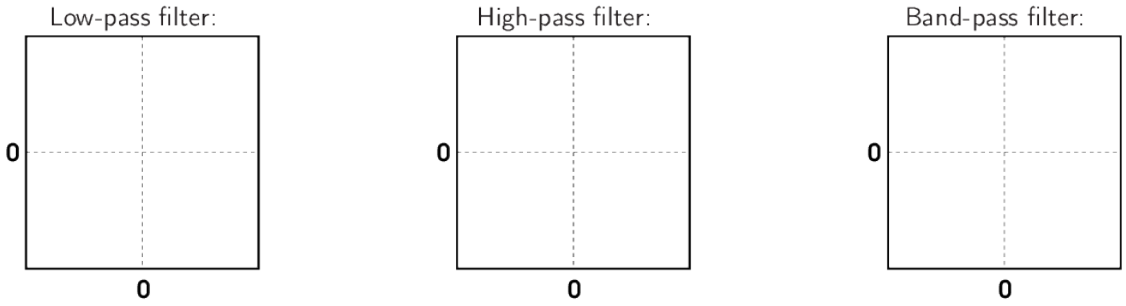

Figure 6: Sketch of three different filter-types

3.2.2 [2 points] Determine the filter type

In figure 7 the Fourier transformation of a filter is depicted (Black values correspond to low (zero) values and white corresponds to high values.) What kind of filter do you consider this ? (High-pass, Band-pass or Low-pass filter)? Give your reasoning.

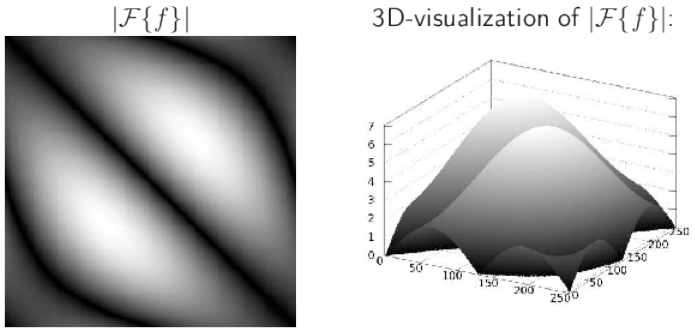

Figure 7: Fourier transformation of a filter

Answer:

3.2.3 [2 points] If you compute the Fourier transform of the two images shown in figure 8 can you distinguish the amplitude images of their Fourier transform $|\mathcal{F}\{f\}|$?

Figure 8: Shifted images

Answer:

3.2.4 [4 points] You can perform a convolution ($f * g$) of an image $f(x, y)$ with a filter $g(x, y)$ by directly doing a spatial filtering or you can obtain the same result by using a frequency domain filtering. Draw a diagram and/or enumerate the steps that you need to follow for the frequency domain filtering.

Answer:

3.2.5 [2 points] Why does it make sense to do the computational detour described in 3.2.4 ?

(Do not compute it in detail, just indicate briefly what happens if the filter and/or image gets larger.)
Answer:

4 Theme: Morphological Image Processing

4.1 Morphological Operations

Recall that the **erosion** of an image A by a structuring element B is usually written as $A \ominus B$ and that the **dilation** of the image by the structuring element B is written as $A \oplus B$. Also recall that the morphological operation **opening** is a combination/concatenation of erosion and dilation (first erosion then dilation)

$$A \circ B = (A \ominus B) \oplus B$$

while for **closing** the operation order (first dilation then erosion) is just exchanged:

$$A \bullet B = (A \oplus B) \ominus B$$

4.1.1 [2 points] Which of the following statements describes an erosion and which describes a dilation?

1) Keep only the object pixels at which the structuring element fully overlaps with object pixels in image A.
2) All background pixels in the image that can be reached by the structuring element B (from object pixels) turn into object pixels

Answer:
Erosion is described by statement: Dilation is described by statement:

4.1.2 [3 points] Perform a dilation with the structuring element B on the shown image D

Use the following structuring element B:
Object/Foreground pixels have a value of 1 and background pixels are indicated by values of 0. The reference pixel is indicated by the circle.

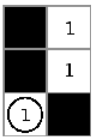

Image D:

0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0
0	0	0	0	1	0	0	0	0	0
0	0	0	1	1	1	1	0	0	0
0	0	0	1	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0

Answer: Result image (**dilation**):

Figure 9: Dilation

Mark all object pixels in the result image. (You can leave background pixels empty!)

4.1.3 [2 points] Perform an erosion with the structuring element B on the shown image E

Use the following structuring element B:
Object/Foreground pixels have a value of 1 and background pixels are indicated by values of 0. The reference pixel is indicated by the circle.

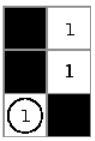

Image E:

0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0
0	0	0	0	1	0	0	0	0	0
0	0	0	1	1	1	1	0	0	0
0	0	0	1	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0

Answer: Result image (**erosion**):

Mark all object pixels in the result image. (You can leave background pixels empty!)

4.1.4 [3 points] Perform a hit-and-miss transform with the structuring element C on the shown image H

Use the following composite structuring element $C = (C_1, C_2)$:
Object/Foreground pixels have a value of 1 and background pixels are indicated by values of 0. The reference pixel is indicated by the circle. (Note: there might be pixel-positions where you do not care → these are shown black.)

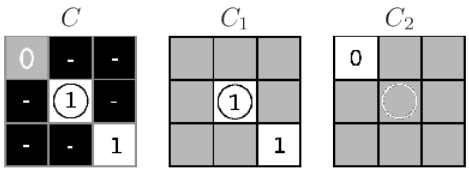

Image H:

0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0
0	0	0	0	1	0	0	0	0	0
0	0	0	1	1	1	1	0	0	0
0	0	0	1	1	0	0	0	0	0
0	0	0	0	1	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0

Answer: Result image (**hit-and-miss**):

Mark all object pixels in the result image. (You can leave background pixels empty!)

4.1.5 [1 point] Which image is the result of an opening and which is the result of a closing?

The shown results are obtained by using the shown structuring element S:
White pixels indicate object/foreground pixels.

S

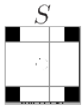

original image A

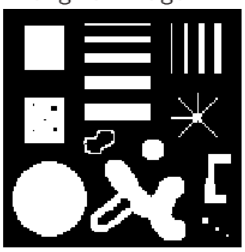

X

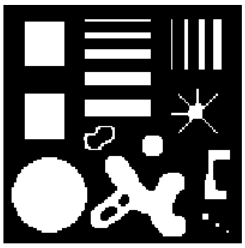

Y

Answer:
Image X is the result of
Image Y is the result of

4.1.6 [2 points] The morphological operation opening is idempotent. What does that mean?

Answer:

4.1.7 [2 points] Explain how you can extract the boundary of a binary image with the help of morphology.

Answer:

4.1.8 [2 points] How do erosion and dilation relate to each other when the structuring element is symmetric?

Answer:

4.1.9 [2 points] Shortly explain how the erosion of a grey value image is computed

Answer:

5 Theme: Segmentation

5.1 Histograms

5.1.1 [2 points] Sketch how the histogram of a low contrast image looks like (An example image is shown on the left side of figure 10).

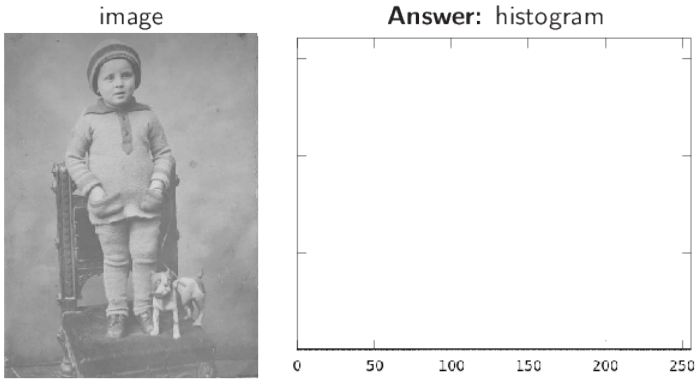

Figure 10: Low-contrast example, Draw a sketch of the histogram.

5.1.2 [1 points] Name a parameter-less method which you can use in order to enhance the image (i.e. increase the contrast)

Answer:

5.1.3 [2 points] Explain briefly the main idea behind that method.

Answer:

5.1.4 [2 points] Sketch a histogram of an image that is well suited for a binary (foreground/background) segmentation. Indicate where you would set the threshold.

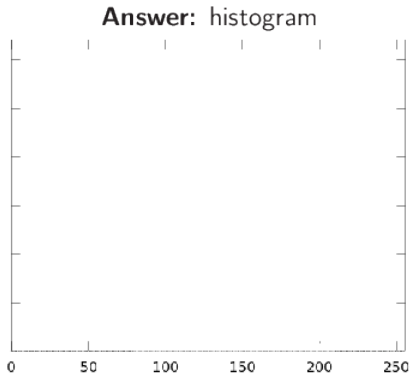

5.2 Region based segmentation

5.2.1 [3 points] An image is to be segmented into a number of regions R_i and we denote a Boolean homogeneity criteria of the region R_j with $H(R_j) = true/false$. Which conditions should the finally segmented regions fulfill?

(You might describe the desired properties of the resulting regions and/or provide a mathematical description of that.)
Answer:

5.3 Watershed segmentation

5.3.1 [2 points] What is meant by flooding in connection with Watershed segmentation?

Answer:

5.3.2 [2 points] Would you apply the Watershed algorithm on the image directly or would you apply it to a high-pass filtered version of the image (e.g. edge-image)? Provide a short explanation of your choice.

Answer:

5.3.3 [2 points] The Watershed segmentation leads typically to a highly over-segmented image. Name one pre- and one post-processing step that you can perform to obtain a more reasonable segmentation?

Answer:

5.4 Nearest Neighbor

5.4.1 [3 points] One of the simplest classification approaches is the K-Nearest-Neighbor method. Explain shortly how it works.

Answer:

6 Theme: Edges and Boundaries

6.1 Edge detection

The output of high-pass filters (e.g. a Sobel filter) results usually in wide ridges around local maxima of the gradient and leads therefore to a thick not well localised edge response.

6.1.1 [3 points] Describe a technique that you can apply as an additional step to obtain a thin, localised response to an edge.

Answer:

6.2 Canny edge detection

6.2.1 [2 points] What are the three basic parameters of a Canny edge detector. Give a very brief description (Examples: r: radius, l: length, ...)

Answer:

Parameter:	Brief description

6.2.2 [3 points] Explain the double thresholding (=hysteresis thresholding) technique that occurs as one step of a Canny edge detector.

Answer:

6.3 Edges and second derivatives

6.3.1 [1 points] Draw a diagram to show that zero crossings in the second derivative of the image intensity are located at approximately the centre of an step-edge in the image.

Answer:

6.3.2 [2 points] You can create a simple zero crossing detector (to find zero crossing lines between negative and positive image areas) with a 2x2 moving window. Explain how this can be done.

Answer:

6.3.3 [3 points] Define a 4- and an 8-way chain code and provide the boundary contour of the shown object pixels as chain code strings for both. Start at the indicated pixel in clockwise direction.

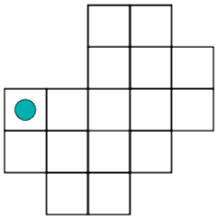

Answers:

4-way chain code:

8-way chain code:

6.3.4 [1 points] Which of the two codes (4-way or 8-way chain code) is shorter and why is that the case?

Answer:

6.4 Fourier descriptors

6.4.1 [1 points] Using only the first two Fourier descriptors (coefficients) for the reconstruction of an originally 2D curve with N points always leads to

a) a full reconstruction of the curve
b) a rectangular shaped curve
c) a circle shaped curve
d) an ellipse with different main axis
Answer:

7 Theme: Short Questions

Below are a number of statements about different topics of the Computer Vision Lecture. For most statements you just have indicate if the statement is true or false.

7.0.2 [1 point] Watershed: Watershed lines form continuous boundaries between regions.

Answer (true/false):

7.0.3 [1 point] Watershed: If n markers are provided in the Watershed algorithm the segmented image consists of $n+1$ catchment basins (segmentations).

Answer (true/false):

7.0.4 [1 point] Watershed: The flooding strategy lead to more efficient algorithms than approaches that search the downstream path to a local minimum for each pixel.

Answer (true/false):

7.0.5 [1 point] Hough Transform: A detected line corresponds to a maxima in the accumulator array.

Answer (true/false):

7.0.6 [1 point] Hough Transform: The accumulator array for the detection of circles has d dimensions. How large is d ?

Answer: d=

7.0.7 [1 point] The Fourier transform of a 1D function $f(x)$ with even symmetry $f(x) = f(-x)$ is only real valued.

Answer (true/false):

7.0.8 [1 point] The first principal component of the Principal Component Analysis (PCA) maximizes the variance of the data-points projected to it.

Answer (true/false):

7.0.9 [1 point] The computation of the first principal component of the Principal Component Analysis (PCA) corresponds to a computation of the largest eigenvalue and the corresponding eigenvector.

Answer (true/false):

7.0.10 [1 point] Dynamic Programming can drastically reduce the computation time for certain problems (like the shortest path through a number of layers) when compared to a brute-force approach.

Answer (true/false):

7.0.11 [1 point] Thresholding: Many parameters must be adjusted within Otsu's Thresholding method.

Answer (true/false):

7.0.12 [1 point] Thresholding: Optimal thresholding is achieved when the number of mis-segmented pixels reach the minimum.

Answer (true/false):

7.0.13 [1 point] The Harris Corner Detector is invariant to an image scale.

Answer (true/false):

——————— Congratulations this is the last page of the exam ! ———————