Computer Assisted Image Analysis II

VT 2016

Exam 2016-03-18

Time:	14:00–19:00	
Place:	Bergsbrunnagatan 15, room 1	
Tools:	Pen or pencil and paper, 1 hand-written A4 sheet with notes	
Grades:	0–17	fail
	18–24	3
	25–32	4
	33–40	5

OBS: Please use drawings and figures to illustrate your answers when suitable. Answer each question on a different sheet. Make sure that you write your exam code (or name) on each page, and remember that the top left corner will be hidden by the staple. It is OK to use both front and back sides of the paper.

Results Will be forwarded to the student office no later than April 8.

GOOD LUCK!

1. Filtering and edge detection (Nataša)

a) List 3 edge-preserving filters and describe one of them. (1p)

b) The image gradient is used to detect edges. What features of the edge are encoded in the gradient? How? (1p)

c) How does the Canny edge detector use the gradient? (1p)

d) How does Canny's algorithm distinguish which edges are relevant and which are not? (2p)

2. Mathematical morphology (Maxime)

a) What sequence of transforms/operations constitutes an opening? How is it used in grey-scale image processing? (1p)

b) What sequence of transforms/operations constitutes the Top-hat transform? What is it used for in image processing? (1p)

c) A binary image contains disks with diameter of 5 pixels (D_5), disks with diameter of 10 pixels (D_{10}), segments of length 10 pixels (S_{10}) and segments of length 15 pixels (S_{15}). The segments can have different orientations. Which (and why) are the operations to:

- Remove only D_5?
- Remove D_5 and D_{10} (only the disks)?
- Select S_{15}?
- What kind of image can this be? (3p)

3. Graph cuts (Filip)

a) Use the Ford-Fulkerson algorithm to calculate the maximum s-t flow on the graph shown in below. The numbers in the figure indicate the capacity of each edge. Show your calculations. (2p)

b) Based on the results from the max-flow calculations, what is the capacity of the minimum s-t cut on the graph? Which edges are included in the minimum s-t cut? (1p)

c) In the course, we have seen that image segmentation can be formulated as an optimization problem. With this approach, we first define an *energy functional* that measures the "goodness" of any segmentation. Then, we use an optimization method to find a segmentation that is as good as possible, according to the energy functional. Some optimization methods find *locally optimal* solutions, while others are guaranteed to find *globally optimal* solutions. Discuss the differences between these two types of optimization methods. What are their advantages and disadvantages? (2p)

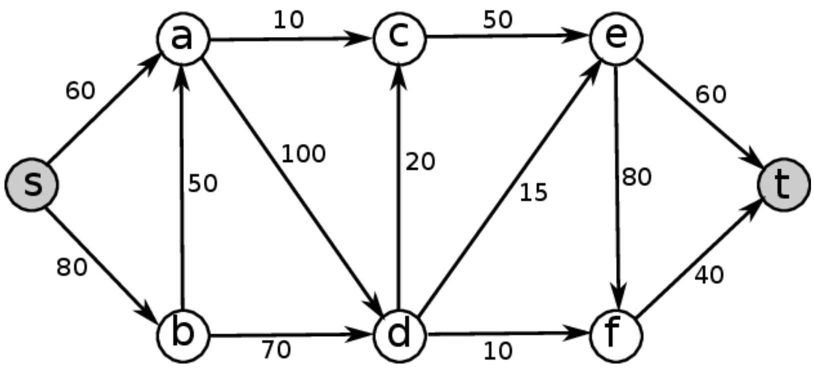

4. Evaluation of segmentation results (Carolina)

The image below shows a small part of a microscopy image of worms, where the goal of the image analysis is to measure the fat content per worm. The fatty regions are outlined in black. The image also shows three different outlines of the worms:

Green = ground truth = GT

Blue = automated approach 1 = AA1

Red = automated approach 2 = AA2

As can be seen from the image, AA1 misses the tip of one worm, and AA2 over-estimates the area of the worm body by constantly including slightly more than the ground truth.

Describe how

a) the Dice similarity coefficient (2p)
b) the Hausdorff distance (2p)
are used to evaluate image segmentation results.

c) Discuss why it is of interest to include both measurements in an evaluation. (1p)

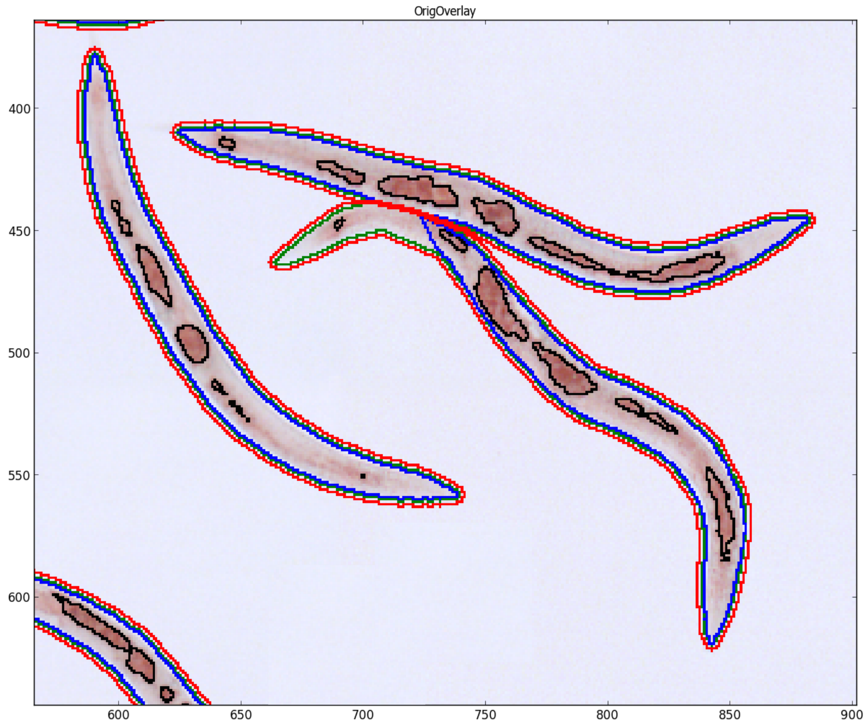

5. Digital geometry (Robin)

a) Give a skeletonization algorithm that preserves topology. Explain how it works.

Give a skeletonization algorithm that is reversible. Explain how it works. (3p)

b) How would you define a skeletonization algorithm for gray-scale images such that

- it is reversible?
- it preserves topology? (2p)

6. Computer vision, Motion, Classification (Anders)

a) Derive the equations for a pinhole camera and explain all variables involved when a 3-D point is projected into 2D image plane. (2p)

b) The apparent 2D velocity of an object is measured at two positions on the border of the object. The velocities are (1, 2), (1, -1). Give both a graphical and an exact solution to the true velocity. (2p)

c) Two sets of points, A and B, are used to train two linear classifiers. Plot the points and draw two decision boundaries, one corresponding to a SVM classifier and one that is a non-SVM classifier. A = {(1,1)}, B = {(2,2), (2,1))}. (1p)

7. Fuzzy sets and the Coverage model (Nataša)

A fuzzy set, which can be seen as a generalization of a classic (crisp) set, is defined by a fuzzy membership function. Coverage based object representation is a special case of a fuzzy object representation.

a) What are main characteristics of fuzzy membership functions? What is a membership function of a crisp set? How is the membership function defined for a coverage representation? (1p)

b) The fuzzy c-means clustering method uses a fuzzy representation of objects. Explain the difference between K-means clustering and fuzzy c-means clustering. (1p)

c) We often want to estimate measurements of real objects, using digital representations of these objects and some estimators for the measures we are interested in. What does it mean when we say that an estimator is accurate? What does it mean when we say that an estimator is precise? (1p)

d) How can we estimate area of an object from its pixel coverage representation? What can we gain, in comparison with utilizing crisp digital representation of the object? (2p)

8. Deep Learning (Sajith)

a) What is a convolutional neural network? (1p)

b) How does it help in improving the classification accuracy for deep neural networks? (2p)

c) What are the pitfalls to be taken care of while training a deep neural network? (2p)