Image Analysis, Handin 3

These are distributed during the lecture September 20, 2013 and has to be finished by October 4, 2013. Written solutions are handed in either (i) at the lectures or (ii) to the box entitled 'inlämningsuppgift bildanalys' in the corridor on the third floor of the math building.

Note: Write your solutions neatly and explain your calculations. All exercises should be done *individually*.

1. Convergence of the K-means algorithm

From analysis of functions in one variable we learned that (see for instance page 155 in Persson-Böiers: Analys i en variabel or page 181 in Månsson-Nordbeck: Endimensionell analys.) 'every decreasing real valued function that is bounded from below converges'. This fact can be used to show that the value of the k-means function converges when using the k-means algorithm.

1) Prove that the k-means algorithm converges in a *finite* number of steps.

Hint: In every step of the k-means algorithms, the correspondences (the vector c in the lecture notes) are updated. Explain why there is only finitely many c. In an iteration two things may happen. Either the goal function is strictly less and then c is changed or the goal function does not change and then c does not change (or at least one can choose not to change c).

2. Segmentation with Graph Cuts

You are given the task of segmenting out two heart chambers in an image by Graph-Cuts. In `heart_data.mat` you are given a number of intensities which have been observed inside the two chambers (*chamber class*) and a number of intensities observed in the background (*background class*). You are also given the image of Figure 1, denoted by f below.

Hint: You can use the matlab routines given in laboratory session 3 to solve the min-cut problem.

1) Assume that the pixel intensities in the two classes are generated by two different Gaussian distributions. In order words, for pixel i, the likelihoods of observing intensity f_i, $P(f_i \,|\, \text{chamber class})$ and $P(f_i \,|\, \text{background class})$, respectively, are Gaussians. Recall, a Gaussian is given by $P(x) = \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{(x-\mu)^2}{2\sigma^2}}$. Estimate the mean and the standard deviation for these two distributions.

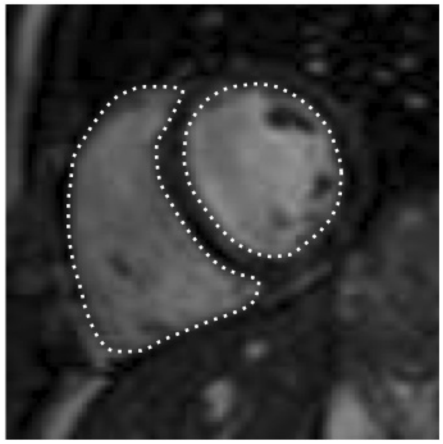

Figur 1: The data given in the task. A slice of a heart imaged by a MRI camera. This view is known as the "short-axis view" taken in a direction where the left and the right heart chambers are visible at the same time. The two chambers are shown inside the dotted lines, the left chamber is to the right in the image and vice versa.

2) Construct a graph of the heart image with a data-term consisting of the negative log likelihoods for the two classes. Solve it via max-flow/min-cut. Experiment with the prior/regularization weight (denoted by ν in the lectures) in order to obtain a reasonable segmentation. (Note: it is hard to get a perfect segmentation.)

3) Optional. Fine tune the segmentation by adding ground truth. This can be done by changing the data-term to ∞ (or a large number) for the pixels you know which region they belong to. This will force certain pixels to belong to a particular class. Such pixels can be obtained by manual inspection of the image.

Notes: Both 2) and 3) can be seen as adjustments a Doctor could perform. In 2) a slider could be added to give interactive response to different weights. Then a pen tool can be used to introduce the ground truth of 3).

In the written solution to these problems, supply both code (e.g. matlab code) and a printout of the results of using your algorithm, i.e. supply examples of input data (e.g. as an image) and result after applying your segmentation algorithm (e.g. also as an image).