CSci 4968 and 6270
Computational Vision,
Fall Semester, 2009-2010
Homework 2
Due: Wednesday 9/23/09 and Friday 9/25/09

Overview

Your answer for question number 3 is due on or before Wednesday, 9/23/09 11:59:59 pm. For this question, no late submissions will be accepted.

Your solutions to problems 1 and 2 are due by 5 pm on Friday, 9/25/09. The same guidelines will be used for these problems as for HW 1, including submission. Be sure to include the phrase "HW 2 submission" on the subject line, so that I can automatically sort your email.

Problems

1. (**20 points**) Implement the motion detection algorithm outlined at the end of the Lecture 3 notes and discussed in class. The images you use will be pre-aligned. The function should take the two images as input arguments. Its output should be a summary description of each region (center, area, 2nd moments) and the final connected components image.

   Two issues must be addressed in your implementation and your write-up:

   (a) What size(s) did you use for the structuring elements in your morphological operations? Why did you pick these?

   (b) Develop an algorithm for automatically choosing a threshold to apply to the image of absolute intensity differences. Describe and justify your algorithm.

2. (**30 points**) Write a function that first implements the gradient computation and non-maximum suppression steps of edge detection, and then thresholds the surviving edges and extracts connected components. The input parameters to the function are the image, the Gaussian smoothing σ value, and the threshold on the gradient. The results of the function must be two images, one binary and one color. Here are more details:

   (a) The binary image must show the pixel locations that survived the non-maximum suppression step. At this point, no thresholding should have been be applied: instead there should be 1 in the binary image if and only if the pixel survives non-maximum suppression

   (b) The color image should be generated as follows. First threshold the gradient of the non-maximum-suppressed image (note: you will still need the gradient directions and magnitudes after non-maximum suppression).

Next, initialize a matrix of components labels so that each pixel is considered unlabeled. (Non-edge pixels will remain unlabeled, of course.) Then, scan through the image in row-major order and for each surviving edge pixel that is currently unlabeled, label a connected component of edge pixels in a breadth-first or depth-first manner. Finally, using `label2rgb` convert the label matrix to an rgb image. The most important issue will be deciding when two 8-connected pixels that have survived non-maximum suppression should be connected to each other in the same component.

3. (**25 points**) On Thursday, September 24, we will begin our discussion of David Lowe's paper, Distinctive Image Features from Scale-Invariant Keypoints, *International Journal of Computer Vision*, **60:2**, 2004, a landmark in the field of computer vision. (Use Google Scholar to find the pdf of this paper.) Your assignment is to read this paper prior to class and send me a text email message, **by 11:59:59 pm on Wednesday, Sept 23rd**,

   (a) A short (250 words or less) summary of the most important points and contributions of the paper.

   (b) A list of three issues about the techniques used in the paper that you think should be discussed in detail in class.

   Some of the material you should find accesssible based on discussions we have already had in class, while other material (and notation!) may be new to you. Just do your best!

   I will build our class discussion of this paper around the issues you raise in this email.