CSci 4968 and 6270
Computational Vision,
Fall Semester, 2009-2010
Homework 3
Due: Thursday, October 8, 2009, 5 pm

Overview

The same guidelines will be used for these problems as for HW 1, including submission. Be sure to include the phrase "HW 3 submission" on the subject line, so that I can automatically sort your email.

Problems

1. (**20 points**) Given a point (x_0, y_0) and a circle centered at the origin, with radius r, use Lagrange multipliers to find the closest point on the circle to (x_0, y_0), and then find the minimum distance from (x_0, y_0) to this closest point.

   Aside: You may submit a hand-written solution to this problem provided that you give it to me by the end of class on the due date.

2. (**50 points**) Implement and test a single-scale version of SIFT. Your function (call it `sss`) should accept as input arguments the image (not the file name), the sigma value, and the threshold on the Laplacian magnitude. Its output should be an $M \times 132$ matrix, where M is the number of keypoints detected, and there are 132 results values produced for each keypoint. The first two values should be the position of the keypoint. The third and fourth should be the strength and orientation (in degrees starting at 0 for positive x and rotating clockwise) of the keypoint, and the remaining 128 values should be the descriptor. In addition to this output, you should create several intermediate displays to show your results — try to be a bit creative here.

   Here are a few simplifications:

   - Apply Gaussian smoothing first. Be sure to save the results as a floating point image.
   - Use the Laplacian function instead of the difference of gaussians, but be sure you normalize by σ^2. Use the Matlab function `del2` to compute the Laplacian.
   - Use the Matlab function `gradient` to do the gradient computation.
   - Obviously, peak detection does not need to be applied across different scales, so you should only be concerned about non-maximum suppression in the spatial domain.
   - You do not need to compute the keypoint locations to subpixel accuracy.
   - You must compute the keypoint orientation and descriptor.

Review for the Midterm

The following problems and topics are suggestive of the kinds of issues that will be addressed on the mid-term exam. The actual questions will be much more specific than most of these are. I will continue to add review questions in future homework assignments.

1. What is the effect of different types of smoothing kernels on a variety of image structures, such as a single line of bright pixels against a dark background, a corner, and a planar intensity surface? Consider box filters and Gaussian filters. What effect do bilateral filters and median filters have on these?

2. Why is it important that the smoothing kernel values sum to 1? Why do we not need to worry about this for differentiation kernels?

3. Derive second derivative kernels. Include the cross-partial,

$$\frac{\partial^2 f}{\partial x \partial y}.$$

4. Under the assumption that a morphological structuring element K contains the $(0,0)$ location, show that

   (a) For dilation, $B \oplus K \subset (B \oplus K) \oplus K$

   (b) For erosion, $(B \ominus K) \ominus K \subset B \ominus K$

   (c) For opening, $(B \bullet K) \bullet K = B \bullet K$

   (d) For closing, $(B \circ K) \circ K = B \circ K$

5. Show that applying two consecutive smoothing Gaussian smoothing convolutions with standard deviations σ_1 and σ_2 is equivalent to applying a single Gaussian smoothing convolution using standard deviation

$$(\frac{\sigma_1^2 + \sigma_2^2}{\sigma_1^2 \sigma_2^2})^{1/2}.$$

   Which is more efficient computationally and why?

6. What do opening and closing with disk structuring elements do to corners? What about other image structures?

7. Give an example showing that the simple edgel contour splitting technique given at the start of the Lecture 6 notes can give poor results (too many lines) even when a good poly-line approximation exists.

8. Using the parametric form of the line, show how to compute the minimum distance of a point to a line. This does not require use of Lagrange multipliers.

9. Outline a least-squares technique line estimation technique that weights the points by their gradient magnitude. How might you incorporate the gradient direction into the least-squares estimation?

10. Outline a Hough transform technique to find circles in images.