CS223B Homework Assignment 1 Solution Guide

Sebastian Thrun, Dan Maynes-Aminzade, Mitul Saha
cs223b@gmail.com

1: Programming a Feature Detector (65 Points)

Overview of Assignment Submissions

A summary of the grading statistics for the programming portion of the assignment can be found here:

`http://cs223b.stanford.edu/homework/hw1/hw1results.ppt`

A detailed analysis of various solutions to the problem has been posted here:

`http://cs223b.stanford.edu/homework/hw1/detailed-hw1-recap.ppt`

In addition, writeups for two of the highest scoring student submissions have been posted in the `hw1/SOLUTIONS` directory.

Rechecking Your Submission

If you'd like to continue to improve the performance of your car detector, you can now evaluate its performance on all 100 source images. Check the `LABELED` directory, which has been updated with 100 hand-labeled images, corresponding to the 100 source images. The scoring script is available here:

`http://cs223b.stanford.edu/homework/hw1/hw1scorer.zip`

2: Perspective Geometry (20 Points)

Consider a scene that contains 5 collinear features, A, B, C, D, and E, where collinearity is defined in 3D coordinates. We know that under perspective projection, the projected features are also collinear in the 2D image plane.

(a) Suppose A, B, C, D, and E are all equidistant. Will the *order* of the projected points (along the projected line) always be the same in the camera image? If yes, argue why this must be the case. If no, provide a counterexample.

(b) Will the observed features also be equidistant in the image plane? If yes, argue why. If not, argue why not.

(c) What are the implications of the previous answer for stereo vision? Specifically, suppose we find five collinear features a, b, c, d, e in the left image, will they also be collinear in the right image? If yes, argue why; if not, provide a counterexample.

(d) Suppose the projected points a, b, c, d, e are collinear in both stereo images. Can we expect that the order is preserved? Again, argue the correctness or provide a counterexample.

(a) Yes, they will. It's the direct result of the fact that the projection is linear.

(b) No, with the exception of uninteresting degenerate cases. Distant objects appear smaller under perspective projection.

(c) Collinearity in the 2D image space does not imply collinearity in 3D coordinates. Non-collinear points in 3D space can often be viewed from a direction such that their 2D projection is collinear.

(d) Unfortunately not. The answer was already given in class when we observed that a foreground object can shift relative to the background.

3: Hough Transform (15 Points)

Suppose you use the Hough Line Transform (described in Section 5.2 textbook) to map the edges of a planar hexagon in the (ρ, θ) space.

(a) What will be the arrangement of the "peaks" in the (ρ, θ) space?

(b) How does the arrangement change when you rotate and translate the polygon?

(c) How does your answer in (a) compare with the answer for 6 lines all intersecting at the same point?

(a) Since the hexagon is regular, all the sides are an equal distance from the origin, so the ρ components of the peaks will be the same. Each side of the hexagon is rotated 60 degrees relative to the previous side, so the θ components of the peaks will be at 60 degree intervals. Hence the peaks will be equidistant along the line $\rho = r$, where r is the distance from the origin to each side of the hexagon.

(b) When the hexagon is rotated about the origin, the θ component of each peak will shift by the same amount, while its ρ component stays fixed. So the peaks will remain on the line $\rho = r$, shifted by a fixed amount corresponding to the amount of rotation of the hexagon.

(c) Since the lines all intersect at the origin, that their ρ values will all be 0. In the Hough space we will see 6 peaks along the line $\rho = 0$, with the θ coordinate of each point matching the orientation of its corresponding line.