CSE190 – Image Processing – Homework #7
Instructor: Prof. Serge Belongie. T.A.: Josh Wills.
`http://www-cse.ucsd.edu/~sjb/classes/wi02/cse190`
Due (in class) 1:25pm Wed. Mar. 6, 2002.

Reading

- GW 10.2.2 and 11.4.
- GW Review Material Ch. 1.

Written exercises

1. GW, Problem 10.13.
2. GW, Problem 10.14.
3. GW, Problem 11.17.
4. GW, Problem 11.18.

Matlab exercises

1. Hough Transform.

   (a) Implement the Hough Transform (HT) using the (ρ, θ) parameterization as described in GW Section 10.2.2. Use accumulator cells with a resolution of 1° in θ and 1 pixel in ρ.

   (b) Produce a simple 11×11 test image made up of zeros with 5 ones in it, arranged like the 5 points in GW Figure 10.20(a). Compute and display its HT; the result should look like GW Figure 10.20(b). Now threshold the HT to find the (ρ, θ)-coordinates of cells with counts greater than 2 and plot the corresponding lines in (x, y)-space on top of the original image.

   (c) Load in the matchstick image in GW Figure 8.02(a) and shrink it to half its size using `I=imresize(I,0.5,'bil','crop');`. Compute and display its edges using the Sobel operator with default threshold settings, i.e. `BW=edge(I,'sobel');`. Now compute and display the HT of `BW`. As before, threshold the HT and plot the corresponding lines atop the original image; this time, use a threshold of 50% of the maximum accumulator count over the entire HT.

   (d) Repeat the previous step for another image of your choice. The image can be from the textbook or elsewhere, but its size must be at least 128×128 and it should contain several extended straight lines.

   Things to turn in:

   - Code listing for part 1a.
   - Code listing for generating results in parts 1b, 1c, and 1d.
   - Printouts of program output for parts 1b, 1c, and 1d.

2. Principal Components Analysis.

   This exercise makes use of the face dataset on `http://isomap.stanford.edu/datasets.html`, consisting of a large set of images of a single face under varying pose and lighting conditions. The total number of faces in the dataset is 698. Each face is stored as a column vector of length 4096 and can be reshaped into a 64×64 grayscale image. For purposes of this exercise, keep only the first 100 faces.

(a) Load this dataset into Matlab and display the first 12 images in a 3×4 subplot.

(b) Compute and display the mean face.

(c) Do PCA on the set of 100 faces, using the trick based on the eigenvectors of the small covariance matrix discussed in class. Make a plot of the eigenvalues sorted in descending order. Display the first 20 eigenfaces in a subplot, and title each image with the corresponding eigenvalue.

(d) Compute the reconstruction of face no. 1 based on the first 50, 75, and 90 principal components. How much of the variance is captured in each of these cases? Make a 2×2 subplot showing the original image followed by the three reconstructions.

(e) Repeat the previous step for face no. 50.

Things to turn in:

- Printouts of program output for steps 2a, 2b, 2c, 2d, 2e.
- Written answer for part 2d.
- Code listing for steps 2b, 2c, 2d.