

CSE152a – Computer Vision – Assignment 4 WI14

Instructor: Prof. David Kriegman.

Revision 0

Instructions:

- This assignment should be solved, *and written up* in groups of 2. Work alone *only* if you can not find a partner. No groups of 3 are allowed.
- Submit your assignment electronically by email to obeijbom@cs.ucsd.edu with the subject line *CSE152 Assignment 4*. The email should have two files attached.
 1. A pdf file with your writeup. This should have all code attached in the appendix. Name this file: CSE_152_hw4_writeup_lastname1_lastname2.pdf.
 2. A compressed archive with all your matlab code files. Name this file: CSE_152_hw4_code_lastname1_lastname2.zip.

The code is thus attached *both* as text in the writeup appendix and as m-files in the compressed archive.

- **Please make this a proper report, with methods, thoughts, comments and discussions. All code should be tucked away in an appendix.** This assignment allows for some more exploration than the previous ones. Please make sure you write a complete report and explain what you did, and why you did it.
- No physical hand-in for this assignment.
- You may do problems on pen an paper, just scan and include in the writeup pdf file.
- In general, MATLAB code does not have to be efficient. Focus on clarity, correctness and function here, and we can worry about speed in another course.

The Yale Face Database

In this assignment, we will have a look at some simple techniques for object recognition, in particular, we will try to recognize faces. The face data that we will use is derived from the Yale Face Database. For more information, follow link from <http://vision.ucsd.edu/~iskwak/ExtYaleDatabase/ExtYaleB.html>. The database consists of 5760 images of 10 individuals, each under nine poses and 64 different lighting conditions. The availability of such standardized databases is important for scientific research as they provide a common testing ground to test the efficacy of different algorithms.



Figure 1: The Yale face database B.

In this assignment, we will only use 640 images corresponding to a frontal orientation of the face. These faces are included in the file `yaleBfaces.zip`. You will find the faces divided into five different subsets. Subset 0 consists of images where the light source direction is almost frontal, so that almost all of the face is brightly illuminated. From subset 1 to 4, the light source is progressively moved toward the horizon, so that the effects of shadows increase and not all pixels are illuminated. The faces in subset 0 will be used as training images, and subsets 1 to 4 will be used as test images.

1 Naive recognition [10 points]

For this first assignment, you will do face recognition by comparing the raw pixel values of the images. Let subset 0 be the train set, and report classification accuracy on test sets 1 to 4. Use the nearest neighbor classifier. Once you have classified all images in a test set, report the average accuracy (nbr correct matches divided by nbr test images). Report the average accuracy for each test set, so *4 numbers*. Comment on the performance. Is it good? Is there any difference between the sets? Is this expected? Why or why not? Discuss and include motivating examples from the database. Notes:

1. You can load a dataset by using e.g. `[trainset trainlabels]=loadSubset(0)`.
2. You can use `drawFaces.m` to look at the images (e.g. `imshow(drawFaces(trainset, 10))`).

2 Recognition Using Eigenfaces [40 points]

- Write a function `[W,mu]=eigenTrain(trainset,k)` that takes as input a $N \times d$ matrix `trainset` of vectorized images from subset 0, where $N = 70$ is the number of training images and $d = 2500$ is the number of pixels in each training image. Perform PCA on the data and compute the top $k = 20$ eigenvectors. Return the $k \times d$ matrix of eigenvectors W , and a d dimensional vector `mu` encoding the mean of the training images. You should use the matlab command `svds` to find the first k eigenvectors. (10 points)
- Rearrange each of the top 20 eigenvectors you obtained in the previous step into a 2D image of size 50×50 . Display these images by appending them together into a 500×100 image (a 10×2 grid of images). (4 points)
- Select one image per person from subset 0 (e.g., the 10 images `person01_01.png`, `person02_01.png`, ... `person10_01.png`). Show what each of these images would look like when using only the top k eigenvectors to reconstruct them, for $k = 1, 2, 3, 4, 5, \dots, 10$. This reconstruction procedure should project each image into a k dimensional space, project that k dimensional space back into a 2500 dimensional space, and finally resize that 2500 vector into a 50×50 image. (6 points)
- Write a function called `testlabels=eigenTest(trainset,trainlabels,testset,W,mu,k)` that takes as input the same $N \times d$ matrix `trainset` of vectorized images from subset 0, an N dimensional vector `trainlabels` that encodes the class label of each training image (e.g., 1 for person01, 2 for person02, etc.), an $M \times d$ matrix `testset` of M vectorized images from one of the test subsets (1-4), the output of PCA W and `mu`, and the number of eigenvectors to use k . Project each image from `trainset` and `testset` onto the space spanned by the first k eigenvectors. For each test image, find the nearest neighbor in the training set using an L_2 distance in this lower dimensional space and predict the class label as the class of the nearest training image. Your function should return an M dimensional vector `testlabels` encoding the predicted class label for each test example. Evaluate `eigenTest` on each test subset 1-4 separately for values $k = 1 \dots 20$ (so it should be evaluated 4×20 times). Plot the error rate

(fraction of incorrect predicted class labels) of each subset as a function of k in the same plot, and use the matlab `legend` function to add a legend to your plot. (10 points)

- Repeat the experiment from the previous step, but throw out the first three eigenvectors. That is, use k eigenvectors starting with the fourth eigenvector. Produce a plot similar to the one in the previous step. How do you explain the difference in recognition performance from the previous part? (5 points)
- Explain any trends you observe in the variation of error rates as you move from subsets 1 to 4 and as you increase the number of eigenvectors. Use images from each subset to reinforce your claims. (5 points)

3 Recognition Using Fisherfaces [25 points]

- Write a function called `[W,mu]=fisherTrain(trainset,trainlabels,c)` that takes as input the same $N \times d$ matrix `trainset` of vectorized images from subset 0 and the corresponding class labels `trainlabels`, and the number of classes $c = 10$. Your function should do the following (10 points):
 - Compute the mean μ of the training data, and use PCA to compute the first $N - c$ principal components. Let this be W_{PCA} .
 - Use W_{PCA} to project the training data into a space of dimension $N - c$.
 - Compute the between-class scatter matrix S_B and the within class scatter matrix S_W on the $N - c$ dimensional space from the previous space.
 - Compute W_{fld} , by solving for the generalized eigenvectors of the $c - 1$ largest generalized eigenvalues for the problem $S_B w_i = \lambda_i S_W w_i$. You can use the matlab function `eig` to solve for the generalized eigenvalues of S_B and S_W .
 - The fisher bases will be a $W = W_{fld} W_{PCA}$, where W is $(c-1) \times d$, W_{fld} is $(c-1) \times (N-c)$, and W_{PCA} is $(N-c) \times d$.
- As in the Eigenfaces exercise, rearrange the top 9 Fisher bases you obtained in the previous part into images of size 50×50 and stack them into one big 450×50 image. (5 points)
- As in the eigenfaces exercise, perform recognition on the testset with Fisherfaces. As before, use a nearest neighbor classifier, and evaluate results separately for each test subset 1-4 for values $k = 1 \dots 9$. Plot the error rate of each subset as a function of k in the same plot, and use the matlab `legend` function to add a legend to your plot. Explain any trends you observe in the variation of error rates with different subsets and different values of k , and compare performance to the Eigenface method. (10 points)