

ECE 661: Homework 7

Due November 18th

Problem:

This homework is quite a bit different than the preceding ones. The previous homeworks have been focused on understanding the geometrical issues involved in computer vision. This homework and the next one will focus more on machine learning.

The goal of this homework is to classify faces. Given images of the faces of m different people, you wish to determine which of the m people is shown in each image. Here is one common way to this:

1. Using training images of each person, determine an appropriate p -dimensional subspace (where p is much smaller than the dimensionality of the input space) and project the training images into this space
2. Project the test image into the same p -dimensional subspace
3. Compute the distance between the projected test image and each of the projected training images. The closest training sample gives the class of the test image (i.e. the identity of the person). This is called nearest neighbor classification).

The only difficult step is the first one, how do we determine an appropriate p -dimensional subspace? In class you have studied two dimensionality reduction techniques, PCA and LDA. For this homework you will try both methods and compare their performance. Some questions you will need to answer in order to do this:

- What is the input space? The easiest option is to convert each pixel to a grayscale value and then stack the columns of the image to form one long vector; however, there are other possibilities and you are free to use anything you would like.
- What should p (the subspace dimensionality) be? Something to think about: what are the theoretical limits on p for PCA and LDA?
- Implementation issues related to the high dimensionality of the input space. Refer to Prof. Kak's tutorial at <http://cobweb.ecn.purdue.edu/~kak/OptimalSubspaces.pdf>. In particular, you will find pages 8-13 useful for PCA and pages 28-32 and 41-47 useful for LDA.

Finally, you will need some data. The CUbiC FacePix(30) database contains images of 30 different people taken under various poses and lighting. First, download a subset of

the database from

http://engineering.purdue.edu/ece661/homework/ECE661_hw7_images.zip

[Note that you only have permission to use this dataset for educational, non-commercial use]. This archive has a readme file with a copy of the license agreement for the dataset and also two folders, train and test. Within each folder there are 21 images of each person with the filename formatted by:

[# of person (01-30)]_[# of image (01-21)].png

There should be a total of $30 \times 21 = 630$ images in each folder. Use these images for training and testing your recognition algorithm.

Solution:

You should turn in a report in pdf format of your homework solution using electronic turn-in. The report should include:

1. An overview of your classification algorithm for both PCA and LDA. Make sure you explain how you answered the questions raised in the problem statement.
2. A plot showing the classification accuracy for both PCA and LDA as a function of the subspace dimensionality p .
 - (a) Compute the classification accuracy using

$$accuracy = \frac{\# \text{ of test images correctly classified}}{\text{total } \# \text{ of images}}$$

3. Compare the results for PCA and LDA and how they change with p . Do they match what you would expect?
4. Your source code (you are strongly encouraged to use MATLAB).

Finally, this homework is new for this year so it may not work very well. Don't get stressed out if the results are poor. Just do the best you can with it.