ECE661: Homework 11

Fall 2020
Due Date: Dec 02, 2020 (11:59 PM)

Turn in typed solutions via BrightSpace. Additional instructions can be found at BrightSpace.

1 Introduction

This is the last homework for this semester and it consists of two parts:

1. Face Recognition with PCA and LDA for dimensionality reduction and the nearest-neighbor search for classification
2. Object detection with a cascaded AdaBoost classifier.

2 Task 1: Face Recognition

The goal of this part of the homework is to classify an unknown face image given a database of labeled face images. Your overall approach will involve the following steps:

1. For this homework, you will use a subset of the well-known face database known as FacePix(30)[1] which is made available by John Black of Arizona State University. Although you will not be using the full database, nonetheless you should visit the website that is mentioned at [2] in order to become familiar with it. Note that the number 30 in the name of the database refers to the number of people whose face images are in the database. The database contains three sets of images for each individual, with 181 images in each set.
2. The subset of the database you will need for your homework can be downloaded via BrightSpace. This is a carefully selected subset of [2]. Download and extract the archive to a folder. When you unzip this archive, you will see two directories named 'train' and 'test', each with 630 images. The name of each image in each directory is a string like "XX–YY.png" where 'XX' is the identity of the human subject and

[1]Note that you only have permission to use this dataset for educational, non-commercial use. The archive includes a copy of license agreement that governs the use of the dataset.

'YY' an integer index for the image for that human subject. Since the same names are used in both the 'train' and the 'test' directories, it may seem like that the two directories contain exactly the same image set. However, that is not the case, as you can verify by displaying the images with the same names (using, say, the 'display' command in Linux).

3. Vectorize the images in the training set and compute the covariance matrix for the image vectors. Do not forget to subtract the mean and to normalize each image vector to unity before you calculate the covariance matrix.

4. Use PCA and LDA to create a low-dimensional representation for the faces in the training set. With each approach, you will retain p eigenvectors. For PCA, the p eigenvectors are directly the eigenvectors of the covariance matrix. Make sure you use the computational trick discussed in class for calculating the eigenvectors of the covariance matrix. For LDA, the eigenvectors are for the matrix $S_W^{-1}S_B$ where S_W is the within-class scatter matrix and S_B is the between-class scatter matrix. Should it happen that your S_W is singular, then you must use the Yu and Yang's algorithm for finding the LDA eigenvectors (ref. pages 66 through 71 in Prof. Kak's tutorial on Optimal Subspaces [1]).

5. Project all of training set images into the p–dimensional subspace.

6. Now classify each image in the test set by vectorizing it, projecting the vector into the p dimensional subspace, and using NN (nearest neighbor) classification. The NN classification means that your test image gets the same person label as its nearest neighbor in the p dimensional space. **As mentioned earlier, the person label of a training image is the integer 'XX' in the name 'XX-YY.png' of the image.** By comparing the 'XX' of the nearest training image with the 'XX' of the test image, you can figure out whether or not the classification is correct.

7. Calculate your classification error rate by repeating this experiment for all the images in the test set.

For this assignment, show the classification accuracy for both PCA and LDA as a function of the subspace dimensionality p. Subsequently, compare the results for PCA and LDA and how they change with p. For your homework solution, please turn in a report in pdf format using the BrightSpace. The report should include

1. A brief outline of your classification algorithm for both PCA and LDA.
2. A plot showing the classification accuracy as a function of the subspace dimensionality p for both PCA and LDA. Compute the accuracy using
$$accuracy = \frac{\text{\# of test images correctly classified}}{\text{total \# of test images}}$$
3. Comparison of the results for PCA and LDA.
4. Your source code. Make sure that your source code files are adequately commented and cleaned up.

3 Task 2 : Object Detection with Cascaded AdaBoost Classification

The goal here is to use the Viola and Jones approach to the design of an object detector with an arbitrarily low false positive rate. (In object detection, you have a false positive when you declare non-object blob of pixels as the object.) The Viola and Jones algorithm carries out a cascaded implementation of the AdaBoost classifier, with each stage of the cascade carrying out classifications at a targeted false-positive and true-detection rate. At each stage, from a very large number of elementary features, the AdaBoost selects the best minimal set for the targeted false-positive and true-detection rates. For this homework, your job is to design a car detector. You will be sent by email an archive with a training and a test directory. Each directory has two sub-directories, one with positive examples of cars and the other with negative examples. Using the Viola and Jones algorithm, you will create a cascaded detector from the positive and the negative examples in the training directory. Subsequently, you will test the detector on the positive and the negative examples in the test directory.

Please do not share the data directories you will be sent for this homework and please do not make them public. The data directories are proprietary information and belong to a corporate entity. The best thing to do would be to delete them after you are done with this homework.

For your homework solution, please turn in a report in pdf format using the BrightSpace. The report should include

1. A brief outline of the classifier including the full training procedure.
2. A plot showing the false positive and false negative rate after the first k stages of the cascade as a function of k:
   (a) Compute the false positive rate using
   $$FP = \frac{\text{\# of misclassified negative test images}}{\text{\# of negative test images}}$$
   (b) Compute the false negative rate using
   $$FN = \frac{\text{\# of misclassified positive test images}}{\text{\# of positive test images}}$$
3. Your source code. Make sure that your source code files are adequately commented and cleaned up.

You are permitted to look at sample solutions from previous semesters. However, the work you turn in must be your own. Clearly identify the steps you have taken to solve the problem with your own words. Your grade depends on the completeness and clarity of your work as well as the results.

References

[1] Optimal Subspaces Tutorial. URL `https://engineering.purdue.edu/kak/Tutorials/OptimalSubspaces.pdf`.

[2] FacePix Database. URL `https://cubic.asu.edu/content/facepix-database`.