

Stat 406 Spring 2007: Homework 6

Out Tue 13 March, back Mon 19 March

1 PCA for face recognition

In this question, you will compute a PCA basis (of size K) from a set of training images of faces, each of size 92×122 pixels. You will then project the training and test images into this low dimensional space, and compute Euclidean distance between the projected points. Each test image is then classified as the identity of the closest matching training image (a 1-nearest neighbor classifier).

Download the data `facesOlivetti_trainTest.mat`, which contains X , a 400×10304 matrix with all the faces, and y , a 400×1 vector containing the identity of each face (an integer between 1 and 40: each person has 10 images). See `pcaFaceDemo` (which calls `pcaHighDim`) for an example of how to compute a PCA basis using this data.

1. Find a PCA basis of size K using only the data in X_{train} (which contains the first 8 images of each person). Then project X_{train} and X_{test} into this basis, and compute the distance between the low dimensional representations of test image i and training image j . You can use the `sqdist` function for this (but beware that it assumes the data are stored columnwise instead of row-wise). Use a 1-NN classifier to predict the label of each test point and compare to y_{test} . Plot the misclassification rate as a function of K for $K = 2, 4, 6, 8, 10, 15, 20$. You should get a graph like Figure 1.
2. For $K = 20$, display the nearest neighbor for each test image. You should get a result like Figure 2. (I used the `montage` command to generate this: see `pcaFaceDemo` for an example of how to call this.) Indicate which matches are incorrect (I used the GUI to edit the picture, but you can just list the indices if you prefer).




Figure 1: Misclassification rate vs number of PCA dimensions.



Figure 2: Left: test set (2 images per person for 40 people). Right: the nearest neighbor in the training set (using $K = 20$ PCA dimensions). A red rectangle in the top right indicates an error. Notice how the person in the bottom left test set is wearing sunglasses, but get mapped to the same person without sunglasses in the training set.