Homework #3

Face Recognition with SIFT Features

Please present a report with all your answers, explanations, and sample images or plots. Submit also a soft copy of the source code and binaries used to generate these results. Please note that copying of any results or source code will result in ZERO credit for the whole homework.

In this homework you will implement a simple face recognition application using SIFT features. The idea is to extract SIFT features from a set of "training" images that represent the persons in the database, and then given a test image, we count how many SIFT features are matched in each database image and consider the image with the maximum number of matches the matched person. The general steps should be as follows:

1. Download the AT&T Face Dataset that will be used for the experiments. It has 10 images for each of 40 different persons with different viewpoint and lighting conditions, for a total of 400 images.
2. Make sure you have a library for computing SIFT features, for example VL Feat or use OpenCV's implementation.
3. Divide the images into a training set (i.e. the set of images stored for every person), and a test set (i.e. the set of images on which we "test" the system's performance). Use the first 4 images for each person as training images for a total of 160 images, and the 6 remaining images per person are used for testing for a total of 240 images.
4. Compute the SIFT features for all the training images.
5. For each test image:
    1. Compute its SIFT features
    2. For each feature, find its "matching" features from the training features. For this you will need to define a distance measure and a threshold. Try the following and report on which is better:
        a) Use the Euclidean distance and a fixed distance threshold, such that when the distance is

below the threshold $d(a,b) \leq t$ the two features are considered matching.

b) Use the Euclidean distance and the *Nearest Neighbor Distance Ratio*, such that when the ratio of the distance to the nearest feature to the distance to the second nearest feature is below a threshold $d_1/d_2 \leq t$ the two features are considered matching.

3. Return the image with the largest number of matching features as the matching image.
4. Decide whether the returned person is the correct person in the test image.

Requirements:

1. Implement the steps above.
2. For a few sample images, extract their SIFT features and plot them on top of the image with marker sizes (e.g. circles) proportional to the feature scale.
3. In order to complete part (5.2) above, you will need to choose a threshold. Try a few threshold values, measure the false positive and false negative ratios on the training set (with the ground truth image removed) or plot the distribution of distances for the nearest neighbors of some features. In either case, explain and justify your choice of the threshold value.
4. Compare between (5.2a) and (5.2b) based on precision, false positives, false negatives and explain which is better and why.