

- This is an individual assignment. However, you are allowed to discuss the problems with other students in the class. But you should write your own code and report.
 - If you have any discussion with others, you should acknowledge the discussion in your report by mentioning their name.
 - Be precise with your explanations in the report. Unnecessary verbosity will be penalized.
 - You have to hand in the report as a hard-copy in the assignments hand-in box opposite to MC317. You have to submit the code in myCourses. Note that you should do both these submissions before the deadline.
 - After the deadline, you have one week to submit your assignment with 30% penalty.
 - You are free to use libraries with general utilities for linear algebra, data handling, and plotting, such as numpy, scipy and matplotlib for python. However, you should implement the algorithms yourself, which means **you should not use pre-existing implementations of the nearest neighbor and linear discriminant analysis algorithms as found in SciKit learn, Tensorflow, etc.!**
 - If you have questions regarding the assignment, you can ask for clarifications in the class discussion forum or go to the following office hours: Prasanna, Philip (section 1), Ali, Lucas (section 2).
-

Linear Classification and Nearest Neighbor Classification

1. You will use a synthetic data set for the classification task that you'll generate yourself. Generate two classes with 20 features each. Each class is given by a multivariate Gaussian distribution, with both classes sharing the same covariance matrix. You are provided with the mean vectors (DS1-m0 for mean vector of negative class and DS1-m1 for mean vector of positive class) and the covariance matrix (DS1-cov). Generate 2000 examples for each class, and label the data to be positive if they came from the Gaussian with mean m1 and negative if they came from the Gaussian with mean m0. Randomly pick 30% of each class (i.e., 600 data points per class) as a test set, and train the classifiers on the remaining 70% data. When you report performance results, it should be on the left out 30%. Call this dataset as DS1, and submit it with your code.
2. We first consider the probabilistic LDA model as seen in class: given the class variable, the data are assumed to be Gaussians with different means for different classes but with

the same covariance matrix. This model can formally be specified as follows:

$$Y \sim \text{Bernoulli}(\pi), \quad X | Y = j \sim \mathcal{N}(\mu_j, \Sigma).$$

Estimate the parameters of the probabilistic LDA model using the maximum likelihood approach. For DS1, report the best fit accuracy, precision, recall and F-measure achieved by the classifier, along with the coefficients learnt.

3. For DS1, use k -NN to learn a classifier. Repeat the experiment for different values of k and report the performance for each value. We will compare this non-linear classifier to the linear approach, and find out how powerful linear classifiers can be. Do you do better than LDA or worse? Are there particular values of k which perform better? Report the best fit accuracy, precision, recall and f-measure achieved by this classifier.
4. Now instead of having a single multivariate Gaussian distribution per class, each class is going to be generated by a mixture of 3 Gaussians. For each class, we'll define 3 Gaussians, with the first Gaussian of the first class sharing the covariance matrix with the first Gaussian of the second class and so on. For both the classes, fix the mixture probability as (0.1,0.42,0.48) i.e. the sample has arisen from first Gaussian with probability 0.1, second with probability 0.42 and so on. Mean for three Gaussians in the positive class are given as DS2-c1-m1, DS2-c1-m2, DS2-c1-m3. Mean for three Gaussians in the negative class are gives as DS2-c2-m1, DS2-c2-m2, DS2-c2-m3. Corresponding 3 covariance matrices are given as DS2-cov-1, DS2-cov-2 and DS2-cov-3. Now sample from this distribution and generate the dataset similar to question 1. Call this dataset as DS2, and submit it with your code.
5. Now perform the experiments in questions 2 and 3 again, but now using DS2. Report the same performance measures as before. What do you observe?
6. Comment on any similarities and differences between the performance of both classifiers on datasets DS1 and DS2?

Instruction for code submission

1. Submit a single zipped folder with your McGill id as the name of the folder. For example if your McGill ID is 12345678, then the submission should be 12345678.zip.
2. If you are using python, you must submit your solution as a jupyter notebook.
3. Make sure all the data files needed to run your code is within the folder and loaded with relative path. We should be able to run your code without making any modifications.

Instruction for report submission

1. Your report should be brief and to the point. When asked for comments, your comment should not be more than 3-4 lines.

2. Do not include your code in the report!
3. If your report consists of more than one page, make sure the pages are stapled.