


CPSC 340 Assignment 2 (due October 2nd)

Frequency-Based Supervised Learning, K-Means Clustering


- You can work in groups on the assignments. However, please hand in your own assignments and state the group members that you worked (as well as other sources of help like online material).
- Place your name and student number on the first page, and submit all answers as a single PDF file to handin.
- Please organize your submission sequentially according to the sections used in this document.
- All Sections (1-4) are equally weighted.
- There may be updates/clarifications to the assignment after the first version is put online. Any modifications will be marked in **red**.

1 K-Nearest Neighbours

Download and expand *a2.zip*, then run *classifier2Ddemo* in Matlab. This loads a dataset with 2 features, fits a decision stump, and then produces a plot similar to the one below:



This is a scatterplot of the feature values across all the training examples. The x-axis is given by the first feature, the y-axis is given by the second feature, and the symbol ('+' or 'o') gives the class label. The green region shows the area of the space that the decision stump labels '+', and the blue region shows the area of the space that the decision stump labels 'o'. In this case, you can see the decision stump is using a rule that tests whether the first feature is larger than a number that is close to 0.5. You can also see that it incorrectly classifies some of the 'o' points where the first feature is less than 0.5. If you change the depth of the decision tree in the demo and re-run it, you can see how the depth of the decision tree affects the predictions that the model will make. For example, below are the plots with depths of 3, 5, and 7.



1.1 Prediction Function

In this dataset, nearby points seem to receive the same class label. This indicates that a k -nearest neighbours classifier might be a better choice than a decision tree. The file `knn.m` has implemented the training function for a k -nearest neighbour classifier (which is to just memorize the data). Fill in the `predict` function in `knn.m` so that the model file implements the k -nearest neighbour prediction rule (based on Euclidean distance). Hand in your `predict` function as well plots similar to the above for $k = 1, 5,$ and 10 .

Hint: Matlab can be slow at executing operations in ‘for’ loops, but allows extremely-fast hardware-dependent vector and matrix operations. By taking advantage of SIMD registers and multiple cores, vector and matrix operations in Matlab will often be several times faster than if you implemented them yourself in a fast language like C. If you find that calculating the Euclidean distances between all pairs of points takes too long, the following code will form a matrix containing the squared Euclidean distances between all training and test points:

```
[N,D] = size(X);
[T,D] = size(Xtest);
D = X.^2*ones(D,T) + ones(N,D)*(Xtest').^2 - 2*X*Xtest';
```

Element $D(i, j)$ gives the squared Euclidean distance between training point i and testing point j .

1.2 Choosing the Number of Neighbours

Recall the two parts of the fundamental trade-off in machine learning:

1. How small you can make the training error?
2. How well does the training error approximates the test error?

In a k -nearest neighbour classifier, how does the parameter k affect each of the two parts of the trade-off? What is one way that we could choose k in practice?

1.3 Runtime

If you have pre-computed the distances to all N training examples, the cost of classifying a new example by sorting these distances is bounded by $O(N \log N)$. Show how we can classify a new example more quickly, even for large values of k , if we use an $O(N)$ -time *selection* algorithm (a *selection* algorithm is an algorithm that returns the k th largest element in an unsorted list).

2 Naive Bayes

If you run the function `example_trainValidate` it will perform the following:

1. Load the `newsgroups` data from Assignment 1, except we have now renamed `Xtest` to `Xvalidate` and `ytest` to `yvalidate` since we have already looked at them in the previous assignment.
2. Train a decision tree of depth 20, and report its validation error.
3. Call the `knn` function and report its validation error. Once you have finished Question 1.1, it should report a lower validation error than the decision tree.
4. Call the `naiveBayes` function and report its validation error. Note that the current `naiveBayes` function is extra-naive, because it assumes that $p(x_{ij}|y_i) = 1/2$ for all features j .

2.1 Estimating the Conditional Probabilities

The `predict` function of the naive Bayes classifier is already implemented, but the calculation of the variable `p_xy` is incorrect (right now, it just sets all values to 1/2). [Modify this function so that `p_xy` computes the conditional probability of these values based on the frequencies in the data set.](#) Hand in your code and report the validation error that you obtain.

2.2 Interpreting the Naive Bayes Model

The predictions made by naive Bayes are bit harder to explain than the explicit rules used by decision trees, since they simultaneously combine information from all the features. However, we can get an idea of which features matter the most by looking at the conditional probabilities. [For each of the four classes 'c', report the three words that have the highest value of \$p\(x_{ij} = 1|y_i = c\)\$.](#)

2.3 Runtime of Naive Bayes for Binary Data

If you have D features for N objects, you can implement the training phase of a naive Bayes classifier in $O(ND)$ since the dominant cost is looking at each element $X(i, j)$ one time. (You do not have to actually implement it in this way for Question 3.1, but you should think about how this could be done). [What is the cost of classifying \$T\$ test examples with the model?](#)

2.4 Explaining the Performance of Different Methods

The naive Bayes model is much less powerful than a k-nearest neighbour classifier (or even a decision tree) because it can only model very simple distributions, yet on this dataset it actually gives lower error. [What theorem from class could explain why a less-flexible model worked better on a particular dataset?](#)

3 Random Forests

If you run the function `example_randomForest` it will perform the following:

1. Load Fisher's famous iris flower dataset, split into a training and validation set.
2. Fit a depth-4 decision tree, and report the validation error.

3. Fit a ‘random forest’ classifier using depth-4 decision trees, and report the validation error.

You will notice that the validation errors of the two models are the same, because the ‘random’ forest classifier doesn’t yet incorporate any randomization.

3.1 Bootstrap Aggregation

Currently, the *randomForest.m* function trains each tree on the entire N training examples. Modify this function so that each tree is trained on a bootstrap sample of the training data (N samples chosen with replacement). [Hand in your code, and report the validation error with this modification.](#)

3.2 Random Feature Selection

Using bootstrap samples introduces a small amount of variation among the trees, but they are still very similar to each other. In order to introduce a larger variation among the trees, modify the *randomStump.m* function so that it splits using *a random feature* instead of searching over all features. [Hand in your code, and report the validation error with this modification.](#)


3.3 Choosing the Number of Random Features/Trees

You can do either problem 1 or problem 2 below:

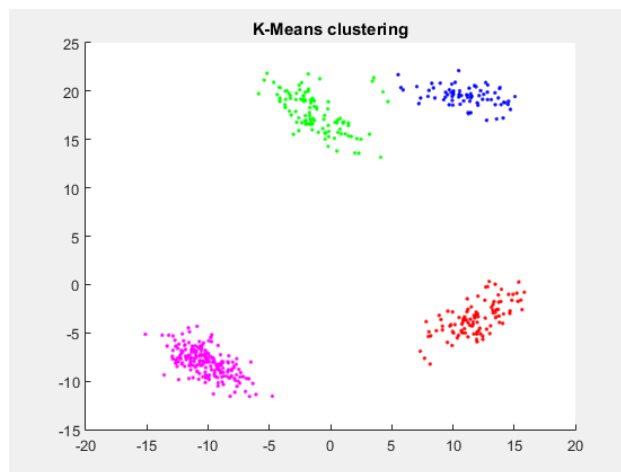
1. Consider the problem of learning a decision tree of fixed depth, where at each depth you consider a random subset of the features as possible candidates for the split variable. Let’s use k to denote the number of random features that each decision stump considers (so in the previous question, we had $k = 1$). [How does the value \$k\$ affect the two parts of the fundamental trade-off in machine learning?](#)
2. Consider the problem of learning a random forest, and let’s use k to denote the number of random trees that we learn. [How does the value \$k\$ affect the two parts of the fundamental trade-off in machine learning?](#)

4 K-Means Clustering

If you run the function *example_Kmeans*, it will load a dataset with two features and a very obvious clustering structure. It will then apply the k -means algorithm with a random initialization. The result of applying the algorithm will thus depend on the randomization, but a typical run might look like this:



But the 'correct' clustering (that was used to make the data) is this:



(Note that the colours are arbitrary.)

4.1 K-Means++

If you run the demo several times, you will see that k -means occasionally discovers the four clusters in the data. However, on many runs it does not. To try to improve its performance on a typical run, modify the `clusterKmeans.m` file so that it uses the k -means++ initialization. **Hand in your modified code. Report whether k -means++ increase the chances of finding the four true clusters, and comment on whether k -means runs faster with the k -means++ initialization.**

Hint: you can use the function `sampleDiscrete` to sample from a discrete probability distribution. For example, if you call `sampleDiscrete([0.2 0.3 0.5])`, then 50% of the time it will return 3, 30% of the time it will return 2, and 20% of the time it will return 1.