



NORTHEASTERN UNIVERSITY, KHOURY COLLEGE OF COMPUTER SCIENCE

CS 6220 Data Mining — Assignment 7

Due: November 25, 2024(100 points)

YOUR NAME
YOUR E-MAIL

Evaluating Your Classifier Performance

In lecture, you will recall that you can calculate the ROC curve by computing the true positive rate (TPR = Probability of Detection, i.e. P_D) and false positive rate (FPR = Probability of False Alarm, i.e. P_{FA}) at every threshold. In this homework, our function will take in predicted classifier scores (`scores` from Livermore Laboratory sensors) and the true labels (`labels` of whether or not laser tubes are calibrated).

ROC/AUC from min-FPR to max-FPR

We will implement a special function that efficiently plots the ROC curve and calculates the AUC for a specified range of FPR. Where your function differs from traditional ROC curve calculations is that it will compute only the true positive rate (i.e., TPR, or P_D) from the range starting from some minimum FPR (i.e., min-FPR, or min- P_{FA} , defaulted to 0) to some specified maximum FPR (i.e., max-FPR, or max- P_{FA} , required parameter).

In particular, it will then return all the FPR values and TPR values that can create a ROC curve in the specified FPR range, as well as the area under it (AUC). For example, I can specify max-FPR = 0.4, and our function will give us all the ROC scores in the range FPR = [0, 0.4], inclusive of the end values. While there may be a way to leverage existing libraries (e.g., Scikit-Learn), please refrain from doing so (though you can feel free to check your answers with them.)

The function signature looks like the following:

```
def roc_range(scores, labels, maxfpr, minfpr = 0)
    '''
    Inputs:
```

```
scores - predictions from any given classifier
labels - the true label (either 0, 1) of the data
maxfpr - the maximum false positive rate (FPR)
minfpr (optional, default = 0) - the minimum false positive rate (FPR)
```

Outputs:

```
fpr_in_range - list of FPR values from minpfa to maxpfa
tpr_in_range - list of true positive rate (TPR) values from minfpr to maxfpr
auc_in_range - single value of area under curve from minpfa to maxpfa
'''
```

```
return fpr_in_range, tpr_in_range, auc_in_range
```

Try it Out On Some Data

We will check your answers against our data, which you can find [here at the course website](#). To read this data, feel free to use this code:

```
import numpy as np
data = np.load("assignment7.npz")

# The data that you will read in
scores_small = data['scores_small']
scores_large = data['scores_large']
labels_small = data['labels_small']
labels_large = data['labels_large']
```

This unpacks the data, which has some small test data that you can use to try out your algorithm before running the analysis on the larger set of data. If interested, this data has been drafted from [National Ignition Facility](#) readouts. Please make your code readable for any data with the above generic signature.

Homework Questions

1. Plot the ROC curve and calculate the AUC for the following ranges:
 - a) $P_{FA} \in [0, 1.0]$, the full range of thresholds
 - b) $P_{FA} \in [0, 0.4]$
 - c) $P_{FA} \in [0, 0.75]$
 - d) $P_{FA} \in [0.25, 0.75]$
2. Your implementation notes:
 - a) Describe your implementation. How would you sweep your thresholds? For each threshold, how would you calculate the PFA and PD? What is the runtime in big- \mathcal{O} notation?
 - b) Determine the runtime of your implementation in big- \mathcal{O} .

- c) Can you make your implementation run in $\mathcal{O}(N \log N)$?
3. What thresholds provide a *precision* of 0.9?
 4. At this threshold, what is the *accuracy* of the classifier?