

HOMEWORK 1

NAIVE BAYES, DECISION TREES, MLE AND MAP

CMU 10-701: MACHINE LEARNING (SPRING 2021)

piazza.com/cmu/spring2021/10701

OUT: Monday, February 08 2021

DUE: February 17th 2021 11:59pm EST

START HERE: Instructions

- **Collaboration policy:** Collaboration on solving the homework is allowed, after you have thought about the problems on your own. It is also OK to get clarification (but not solutions) from books or online resources, again after you have thought about the problems on your own. There are two requirements: first, cite your collaborators fully and completely (e.g., “Jane explained to me what is asked in Question 2.1”). Second, write your solution *independently*: close the book and all of your notes, and send collaborators out of the room, so that the solution comes from you only. See the Academic Integrity Section on the course site for more information:
https://www.cs.cmu.edu/~aarti/Class/10701_Spring21/index.html
- **Extension Policy:** See the homework extension policy here:
https://www.cs.cmu.edu/~aarti/Class/10701_Spring21/index.html
- **Submitting your work:**
 - All portions of the assignments should be submitted to Gradescope (<https://gradescope.com/>).
 - **Programming:** We will autograde your Python code in Gradescope. After uploading your code, our grading scripts will autograde your assignment by running your program on a virtual machine (VM). We recommend debugging your implementation on your local machine (or the linux servers) and making sure your code is running correctly before any submission. **Our autograder requires that you write your code using Python 3.6.9 and Numpy 1.17.0.**
 - **Written questions:** **You must type the answers in the provided .tex file. Hand-written solutions will not be accepted. Make sure to answer each question in the provided box. DO NOT change the size of the boxes, because it may mess up the autograder.** Upon submission, make sure to label each question using the template provided by Gradescope. Please make sure to assign ALL pages corresponding to each question.

1 Course Policies [1pt]

1. [1pt] Have you read and understood the course policies?

☐ Yes

☐ No

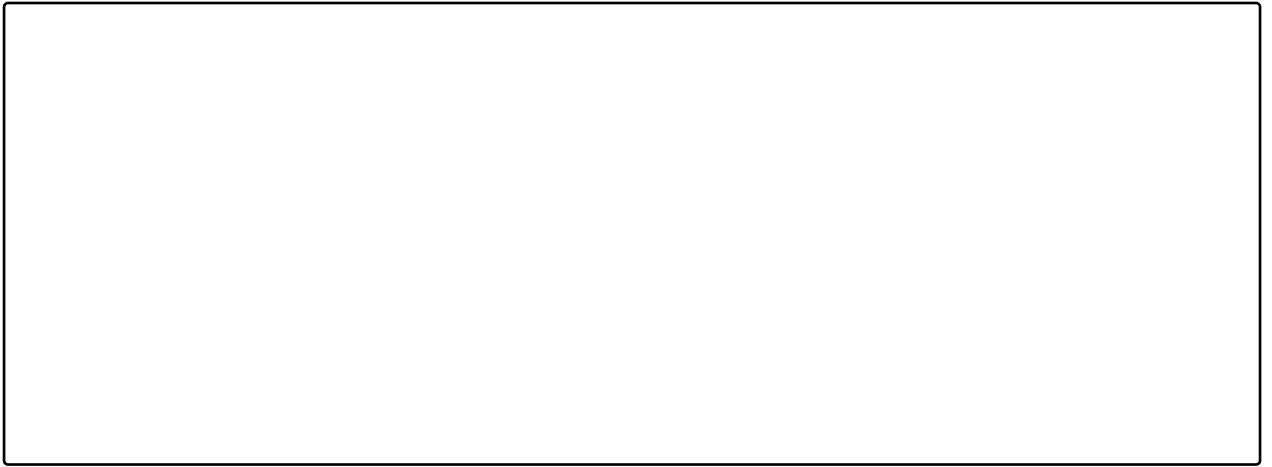
2 Probability Review [15pts]

A group of travellers find themselves lost in a cave. They come upon 3 tunnels A , B , C . Both tunnels A and B are closed loops that do not lead to an exit and in fact lead right back to the entrance of the 3 tunnels. Tunnel C is the tunnel which leads to the exit. If they go through tunnel A , then it takes 2 days to go through the tunnel. If they go through tunnel B , then it takes 1 day to go through the tunnel. If they go through tunnel C , then they immediately leave the cave. Suppose the travellers choose tunnels A , B and C with constant probability 0.3, 0.5, 0.2 every time. (For the following questions please round your answer to 4 digits.)

1. [6 pts] Suppose we record the travellers' choices as a sequence (e.g., $ABBA \dots C$). What is the probability that the pattern AAB appears in the sequence before any BAA appears?

Note: You should also count cases where AAB appears in the sequence and BAA does not.

2. [4 pts] What is the expected number of days that the travellers will be lost in the cave?



3. [5 pts] What is the variance of days that the travellers will be lost in the cave?
(Hint: To compute $Var(T)$ for a random variable T , compute $E[T^2]$ first and apply the definition $Var(T) = E[T^2] - E[T]^2$.)

3 MLE and MAP [16 pts]

3.1 MLE [8 pts]

An exponential distribution with parameter λ has the probability density:

$$p(x) = \begin{cases} \lambda e^{-\lambda x} & \text{if } x \geq 0 \\ 0 & \text{if } x < 0 \end{cases}$$

- Given some i.i.d. data $\{x_i\}_{i=1}^n \sim \text{Exp}(\lambda)$, derive the maximum likelihood estimate (MLE) $\hat{\lambda}_{MLE}$.
- An estimator is unbiased if its expected value is equal to the true parameter it's estimating. Is this estimator biased?

3.2 MAP [8 pts]

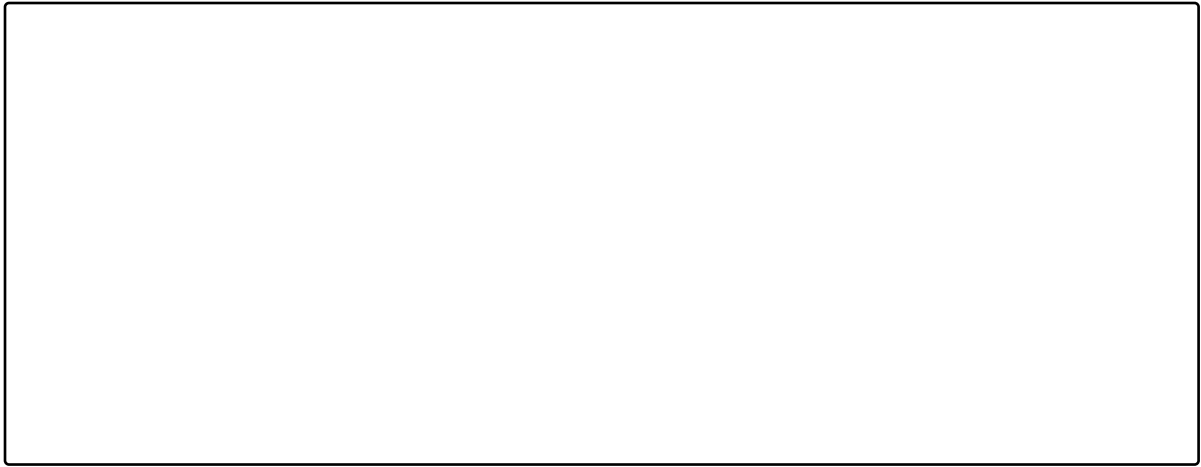
A gamma distribution with parameters α, β has a density function:

$$p(x) = \frac{\beta^\alpha}{\Gamma(\alpha)} x^{\alpha-1} e^{-\beta x},$$

where $\Gamma(t)$ is the gamma function (see https://en.wikipedia.org/wiki/Gamma_distribution).

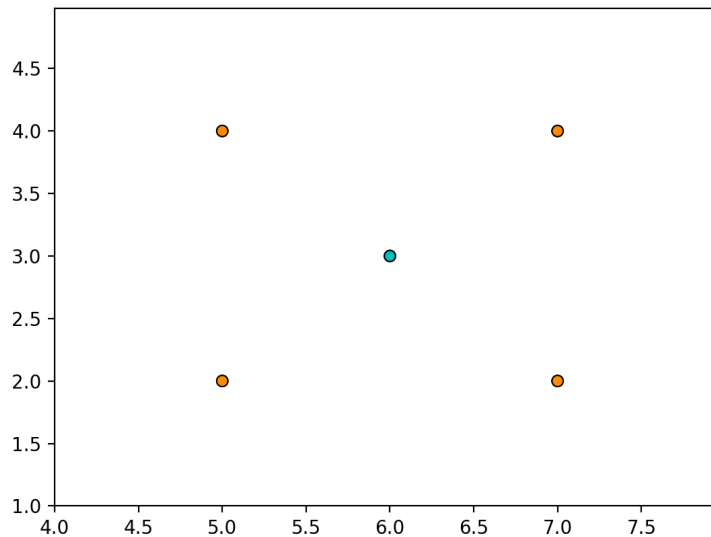
Suppose we start with a prior distribution for some parameters θ , and observe some data $\{x_i\}_{i=1}^n$ with likelihood $P(\text{data} \mid \theta)$. If the posterior for θ has the same form as the prior, then we say that the given prior is conjugate for the given likelihood.

- Show that the Gamma distribution (that is $\lambda \sim \text{Gamma}(\alpha, \beta)$) is a conjugate prior of the $\text{Exp}(\lambda)$ distribution. In other words, show that if the data points $x_i \sim \text{Exp}(\lambda)$ and $\lambda \sim \text{Gamma}(\alpha, \beta)$ then $P(\lambda \mid \text{data}) \sim \text{Gamma}(\alpha^*, \beta^*)$ for some values α^*, β^* .
- Derive the maximum a posteriori estimator (MAP) $\hat{\lambda}_{MAP}$ as a function of α, β . What happens as the number of data points n gets large?



4 K-Nearest Neighbors [15 Points]

1. [3pt] Consider K-NN using Euclidean distance on the following data set (each point belongs to one of two classes: orange or green).



- (a) [2pt] Draw the decision boundary for the above data set using the 1-NN classification rule.
- (b) [1pt] Would the decision boundary be different if another green point was added to the data set at (6.1, 3.1)? Why or why not?

2. [6pt] k -NN Black Box

- (a) [3pt] In a k -NN classification problem, assume that the distance measure is not explicitly specified to you. Instead, you are given a “black box” where you input a set of instances P_1, P_2, \dots, P_n and a new example Q , and the black box outputs the nearest neighbor of Q , say P_i , and its corresponding class label C_i . Is it possible to construct a k -NN classification algorithm (w.r.t the unknown distance metric) based on this black box alone? If so, how? If not, why not? You may use the black-box more than once on any subset of P_1, P_2, \dots, P_n .

- (b) [3pt] If the black box returns the j nearest neighbors (and their corresponding class labels) instead of the single nearest neighbor (assume $j \neq k$), is it possible to construct a k -NN classification algorithm based on the black box? If so how, and if not why not? Consider the cases when $j < k$ and $j > k$. Justify your argument. Again, you may use the black-box more than once on any subset of P_1, P_2, \dots, P_n .

3. [6 pts] Figure 1 shows K-NN classification decision boundaries with various values of K and distance metrics. Match each plot from Figure 1 to one of the corresponding K-NN settings.

To obtain full credit, you do not need to provide any justification, but only provide the correct pairing.

Definition: Manhattan distance: $d_M(x, y) = \sum^D |\mathbf{x} - \mathbf{y}|$

where \mathbf{x}, \mathbf{y} are both vectors of size D

(a) $K = 1$, Euclidean Distance

(b) $K = 25$, Euclidean Distance

(c) $K = 25$, Manhattan Distance

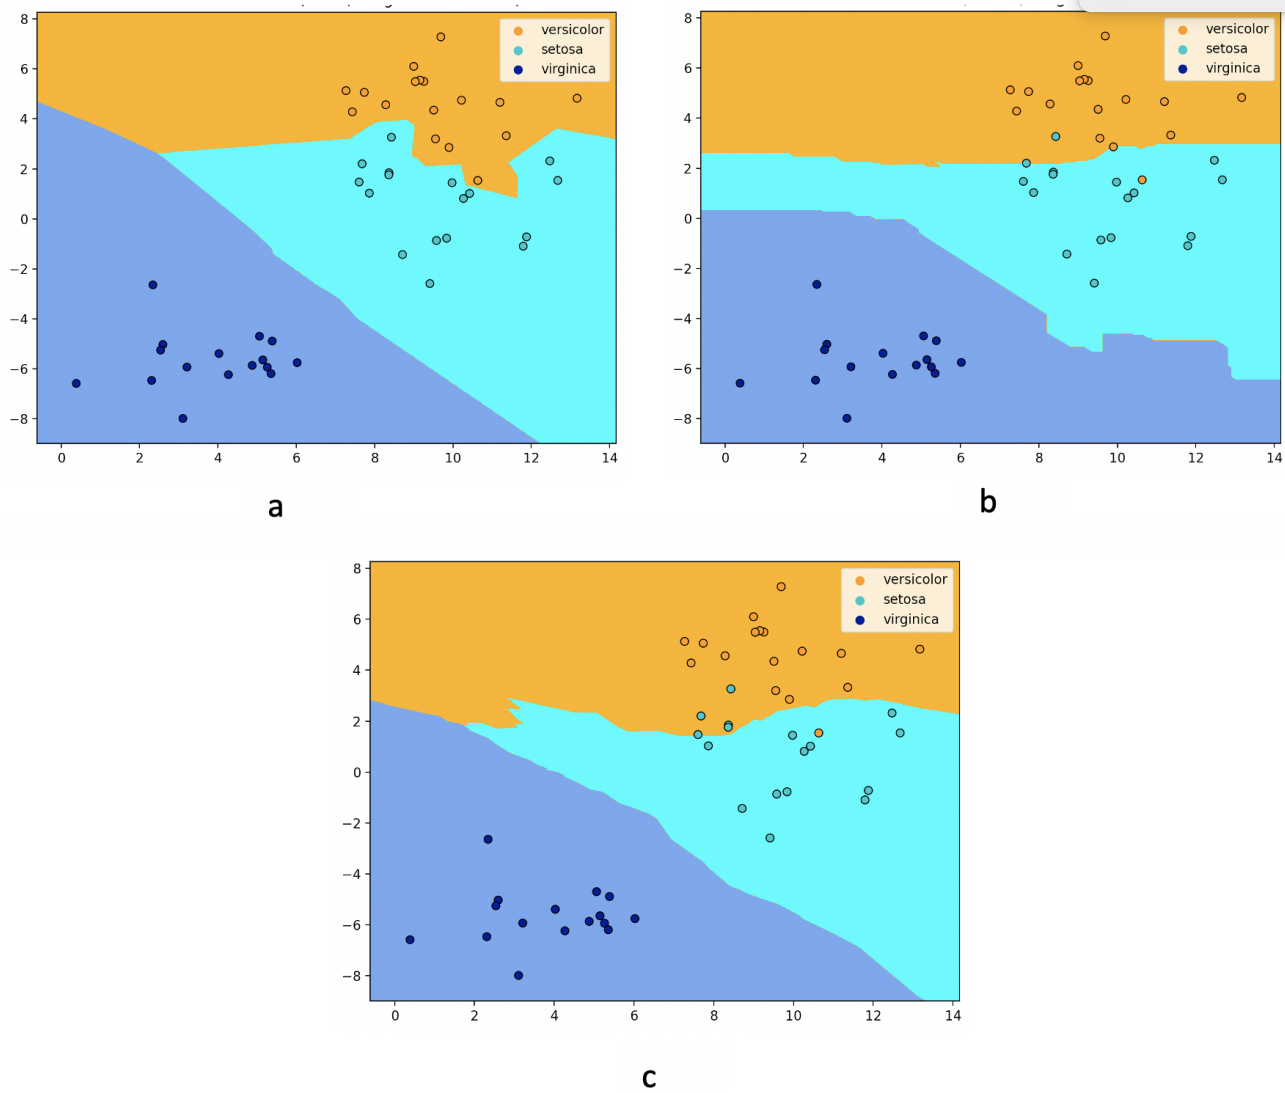


Figure 1: K-NN decision boundaries for various distance metrics and K values

5 Naive Bayes Programming [53 pts]

Programming Instructions

This part of the assignment will have you implement a Naive Bayes classifier. You will submit your completed `naive_bayes.py` file to Gradescope, where we will run your code against a suite of tests. Your grade will be automatically determined from the testing results. Since you get immediate feedback after submitting your code and you are allowed to submit as many different versions as you like (without any penalty), it is easy for you to check your code as you go.

Our autograder requires that you write your code using Python 3.6.9 and Numpy 1.17.0. Otherwise, when running your program on Gradescope, it may produce a result different from the result produced on your local computer. **Please do not include print statements outside the provided functions, as this may crash the autograder.**

The file `hw2data.pkl` contains data regarding words used in articles from The Economist and articles from The Onion. This programming assignment is focused on identifying which words are characteristic of which articles. You can load the pickle file into Python using `pickle`. After loading the data, you will see that there are 5 variables: `Vocabulary`, `XTrain`, `yTrain`, `XTest`, and `yTest`.

- `Vocabulary` is a $V \times 1$ dimensional array that contains every word appearing in the documents. When we refer to the j^{th} word, we mean `Vocabulary[j,0]`.
- `XTrain` is a $n \times V$ dimensional matrix describing the n documents used for training your Naive Bayes classifier. The entry `XTrain[i,j]` is 1 if word j appears in the i^{th} training document and 0 otherwise.
- `yTrain` is a $n \times 1$ dimensional matrix containing the class labels for the training documents. `yTrain[i,0]` is 1 if the i^{th} document belongs to The Economist and 2 if it belongs to The Onion.
- Finally, `XTest` and `yTest` are the same as `XTrain` and `yTrain`, except instead of having n rows, they have m rows. This is the data you will test your classifier on and it should not be used for training.

Logspace Arithmetic

When working with very large or very small numbers (such as probabilities), it is useful to work in *logspace* to avoid numerical precision issues. In logspace, we keep track of the logs of numbers, instead of the numbers themselves. For example, if $p(x)$ and $p(y)$ are probability values, instead of storing $p(x)$ and $p(y)$ and computing $p(x) * p(y)$, we work in log space by storing $\log(p(x))$, $\log(p(y))$, and we can compute the log of the product, $\log(p(x) * p(y))$, by taking the sum in logspace: $\log(p(x) * p(y)) = \log(p(x)) + \log(p(y))$.

If we want the sum of two probabilities, it's a little trickier: if $l(x) = \log p(x)$ and $l(y) = \log p(y)$, then $\log(p(x) + p(y)) = \log(\exp(l(x)) + \exp(l(y)))$. If we compute this expression

naively we risk overflow or underflow. A good workaround is to factor out $\exp(l(x))$ or $\exp(l(y))$, whichever is larger, before computing the sum.

Training Naive Bayes

1. **[8 Points]** Complete the function `D = NB_XGivenY(XTrain, yTrain, a=0.001, b=0.9)`. The output `D` is a $2 \times V$ matrix, where for any word index $w \in \{1, \dots, V\}$ and class index $y \in \{1, 2\}$, the entry `D[y-1, w-1]` is the MAP estimate of $\theta_{yw} = P(X_w = 1 | Y = y)$ with a $\text{Beta}(1.001, 1.9)$ prior distribution. Here we define $a = \alpha - 1$ and $b = \beta - 1$ where α, β are parameters of the Beta distribution. To help with numerical issues clip `D` to be in $[10^{-5}, 1 - 10^{-5}]$ before this function returns it.
2. **[8 Points]** Complete the function `p = NB_YPrior(yTrain)`. The output `p` is the MLE for $\rho = P(Y = 1)$.
3. **[8 Points]** Complete the function `yHat = NB_Classify(D, p, X)`. The input `X` is an $m \times V$ matrix containing m feature vectors (stored as its rows). The output `yHat` is a $m \times 1$ matrix of predicted class labels, where `yHat[i]` is the predicted label for the i^{th} row of `X`. So, the output vector should take the form `[[y0],[y1],...,[ym-1]]`. [Hint: In this function, you will want to use Logspace Arithmetic to avoid numerical problems.]
4. **[2 Points]** Complete the function `e = NB_ClassificationAccuracy(yHat, yTruth)` which measures the average number of times `yHat` agrees with `yTruth` as a performance metric for the Naive Bayes classifier.

Questions

5. **[4 Points]** Train your classifier on the data contained in `XTrain` and `yTrain` by running

```
D = NB_XGivenY(XTrain, yTrain)
p = NB_YPrior(yTrain)
```

Use the learned classifier to predict the labels for the article feature vectors in `XTrain` and `XTest` by running

```
yHatTrain = NB_Classify(D, p, XTrain)
yHatTest = NB_Classify(D, p, XTest)
```

Use the function `NB_ClassificationAccuracy` to measure and report the training and testing accuracy by running

```
trainAcc = NB_ClassificationAccuracy(yHatTrain, yTrain)
testAcc = NB_ClassificationAccuracy(yHatTest, yTest)
```

How do the train and test accuracies compare? Which is likely to be more representative of the performance of the trained classifier on a new collection of articles?

6. **[5 Points]** In this question we explore how the size of the training data set affects the test and train accuracy. For each value of m in $\{100, 130, 160, \dots, 450\}$, train your Naive Bayes classifier on the first m training examples (that is, use the data given by `XTrain[0:m]` and `yTrain[0:m]`). Plot the training and testing accuracy for each such value of m . The x -axis of your plot should be m , the y -axis should be accuracy, and there should be one curve for training accuracy and one curve for testing accuracy.
- Explain the general trend of both the curves.
 - What would you expect to happen to the test accuracy of the classifier if the Naive Bayes assumption is satisfied and we have infinite training data?

7. [4 Points] We will try to interpret the learned parameters. Train your classifier on the data contained in `XTrain` and `yTrain`. For each class label $y \in \{1, 2\}$, create the lists according to the following criteria (Note that some of the words may look a little strange because we have run them through a stemming algorithm that tries to make words with common roots look the same. For example, “stemming” and “stemmed” would both become “stem”):

- Top five words that the model says are most likely to occur in a document from class y . That is, the top five words according to this metric:

$$P(X_w = 1|Y = y)$$

- Top five words w according to this metric:

$$\frac{P(X_w = 1|Y = y)}{P(X_w = 1|Y \neq y)}.$$

Which list of words is more informative about the class y ? Briefly explain your reasoning.



8. [4 Points] Having a metric to tell us the importance of a particular word in helping us classify text can be useful when wanting to shrink the dataset without affecting accuracy. Plot graphs of the training and testing accuracy of your model where you prune unimportant words from the feature set. You should compute the following importance values for every word:

$$I_{0/1} = \text{abs}(1 - \frac{P(X_w = 1|Y = 1)}{P(X_w = 1|Y = 2)}).$$

$$I_{1/0} = \text{abs}(1 - \frac{P(X_w = 1|Y = 2)}{P(X_w = 1|Y = 1)}).$$

For a given threshold value T , you should keep words where the importance I is greater than the threshold, and discard words below the threshold (You will need to do this separately for both values of I). Plot a graph of the test and train accuracy over $T \in [0.1, 1.0]$ with step size 0.1 [Hint: It may be useful to pre-compute a mask and apply it to $P(X|Y)$]. Report the two graphs (one for each I) in your writeup.

- What is the reason for the graph's shape?
- How many vocab words are able to be remove before the accuracy is affected?



9. [5 Points] The `augmentFeatures` function is designed to augment each feature vector of length V with a subset of size v' of new features. These features actually happen to be exact copies of existing features — although the classifier doesn't know that, and will treat them the same as all the other features. The feature augmentation process is identical in the training and test datasets.

We want to observe the effect of this augmentation on the accuracy of the Naive Bayes Classifier. The classifier is trained and tested separately both on the original dataset and the modified dataset with augmented feature vectors. For each value of m in $\{150, 180, \dots, 450\}$, the Naive Bayes classifier is trained on the first m training examples and only the test accuracies are plotted separately for the original and augmented

dataset. Generate the plot by running the `augmentFeatures` function and explain your observation. Does the augmentation of the feature vector violate any assumption for the Naive Bayes classifier? Why is the test accuracy of the trained Naive Bayes classifier affected by augmentation of the feature vector in this manner?

10. [5 Points] Use the `generateData` function to generate a new dataset of size 1000 from your learned model, and then report the classification accuracy of your model (which was trained on the original training set) on this new data set. Explain how the new dataset is constructed using the generative properties of the Naive Bayes model.

Collaboration Questions Please answer the following:

1. Did you receive any help whatsoever from anyone in solving this assignment?

Yes / No.

- If you answered 'yes', give full details: _____
- (e.g. "Jane Doe explained to me what is asked in Question 3.4")

2. Did you give any help whatsoever to anyone in solving this assignment?

Yes / No.

- If you answered 'yes', give full details: _____
- (e.g. "I pointed Joe Smith to section 2.3 since he didn't know how to proceed with Question 2")

3. Did you find or come across code that implements any part of this assignment ?

Yes / No. (See below policy on "found code")

- If you answered 'yes', give full details: _____
- (book & page, URL & location within the page, etc.).

