



BGS INSTITUTE OF TECHNOLOGY

(Affiliated to VTU, Belgaum; Approved by AICTE, New Delhi and Recognized by Govt. of Karnataka)
BG Nagara - 571 448 (Bellur Cross), Nagamangala Taluk, Mandya District

Certificate

This is to certify that Mr/Ms..... JEERAN R......
USN: 4BW.17.C.S.029..... has satisfactorily completed the course of
experiments in MACHINE LEARNING..... Laboratory (Course
Code...17.C.S.L.16.....) prescribed by the Visvesvaraya Technological
University, Belagavi forV.II..... Semester, BE Computer Science
Engineering, of this College in the year 2020-2021

Record Marks : 29 Test Marks : 09

IA Marks : 38

Date : 16/01/2021

H. J. Prasad
Staff Incharge

B. K. Ragh 16/1/21
Head of the Department

Name of the Student : Jeevan . RClass : A Sem. VI

Expt. No.	Date	Title of Experiment	Page No.	Marks Obtained					Sign. of the staff
				a	b	c	d	Total	
1	5/11/20	Find-S Algorithm	01	04	10	05	10	29	
2	7/11/20	Candidate - Elimination algorithm	02	04	10	05	10	29	
3	12/11/20	ED3 algorithm	04	04	10	05	10	29	
4	14/11/20	Back propagation algorithm	10	04	10	05	10	29	
5	19/11/20	Naive Bayesian classifier	13	04	10	05	10	29	
6	21/11/20	Naive Bayesian classifier using Built in Java API	18	04	10	05	10	29	
7	26/11/20	Bayesian network consistency medical DSN.	21	04	10	05	10	29	
8	28/11/20	K-mean algorithm	25	04	10	05	10	29	
9	31/12/20	K-nearest Neighbor algorithm	27	04	10	05	10	29	

INDEX

Name of the Student: Jeevan R

Class: D Sem. _____

Expt. No.	Date	Title of Experiment	Page No.	Marks Obtained				
				a	b	c	d	Total
10	5/12/20	Locally weighted Regression Algorithm	30	04	10	05	10	29

Input: find, CSV

Example	sky	airTemp	humidity	wind	water	forecast	temp
1	sunny	warm	normal	strong	warm	cast	70
2	sunny	warm	high	strong	warm	cast	75
3	warm	cold	light	strong	warm	cast	70
4	sunny	warm	high	strong	cast	cast	70

Output:

['sunny', 'warm', 'strong', '?', '70']



1. Implement and demonstrate the FIND-S algorithm for finding the most specific hypothesis based on a given set of training data sample.
Read the training data from a CSV file

```
import numpy as np
import pandas as pd
data = pd.DataFrame(data = pd.read_csv('finds.csv'))
concept = np.array(data.iloc[:, 0:-1])
target = np.array(data.iloc[:, -1])
def learn(concept, target):
    specific_h = concept[0].copy()
    for i, h in enumerate(concept):
        if target[i] == 'yes':
            for x in range(len(specific_h)):
                if h[x] != specific_h[x]:
                    specific_h[x] = '?'
    return specific_h
```

Input: kind.csv

Example	Sky	AirTemp	Humidity	Wind	Watts	Forecast	Engage
1	Sunny	warm	Normal	Steady	warm	Same	Yes
2	Sunny	warm	High	Steady	warm	Same	Yes
3	Rainy	cold	High	Steady	warm	Change	Yes
4	Sunny	warm	High	Steady	cool	Change	Yes

Output:

Final G:

[['Sunny', '?', '?', '?', '?', '?', '?', '?'], ['4', 'warm', '?', '?', '?', '?', '?']]

7/15/15

2) For a given set of training data examples stored in a .csv file, implement and demonstrate the candidate-Elimination algorithm to output a description of the set of all hypotheses consistent

```
import numpy as np
import pandas as pd
data = pd.DataFrame(data = pd.read_csv
                    ('find.csv'))
concept = np.array(data.iloc[:, 0:-1])
target = np.array(data.iloc[:, -1])
def learn(concept, target):
    specific_h = concept[0].copy()
    general_h = ["?" for i in range(len(specific_h))
                for i in range(len(specific_h))]
    for i, h in enumerate(concept):
        # checking if the hypothesis has a positive target
        if target[i] == "Yes":
            for x in range(len(specific_h)):
                # change values in S & G only if values change
```

```

if h[x] != specific_h[x]:
    specific_h[x] = '?'
    general_h[x] = '?'

```

checking if the hypothesis has a negative target
if target[i] == 'No':

```

for x in range(len(specific_h)):

```

for negative hypothesis change value in G

```

if h[x] != specific_h[x]:

```

```

    general_h[x][x] = specific_h[x]

```

else:

```

    general_h[x][x] = '?'

```

find indices where we have empty rows, meaning that that are unchanged.

```

indices = [i for i, val in enumerate(general_h)

```

```

    if val == ['?', '?', '?', '?', '?', '?']]

```

```

for i in indices:

```

```

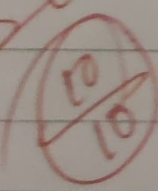
    general_h.remove(['?', '?', '?', '?', '?', '?'])

```

```

return specific_h, general_h.

```

Sum


Input: playTennis.csv

Outlook	Temp	Humidity	Wind	Description
Sunny	Hot	High	weak	NO
Sunny	Hot	High	strong	NO
overcast	Hot	High	weak	Yes
Rain	Mild	High	weak	Yes
Rain	cool	Normal	weak	Yes
Rain	cool	Normal	strong	NO
overcast	cool	Normal	strong	Yes
Sunny	mild	High	weak	no
Sunny	cool	Normal	weak	Yes
Rain	mild	Normal	weak	Yes
Sunny	mild	Normal	strong	Yes
overcast	mild	High	strong	Yes
overcast	Hot	Normal	weak	Yes
Rain	mild	High	strong	no

output:

{ 'ROOT': { 'outlook': { 'overcast': 'Yes',
 'Rain': { 'wind': { 'strong': 'NO', 'weak': 'Yes',
 'Sunny': { 'humidity': { 'High': 'NO', 'Normal': 'Yes' }

- 2) Write a program to demonstrate the working of the decision tree based ID3 algorithm. Use an appropriate data set for building the decision tree and apply this knowledge to classify a new sample.

```
import numpy as np
import pandas as pd
from pprint import pprint
data = pd.read_csv("playTennis.csv")
data_size = len(data)
tree_node = {}
tree = {"ROOT": data}

def total_entropy(data, col):
    mydict = {}
    for elem in data[col]:
        if elem in mydict.keys():
            mydict[elem] += 1
        else:
            mydict[elem] = 1
    total = sum(mydict.values())
    E = 0
```

```

for key in mydict.keys():
    E += entropy(mydict[key], total)
return E

```

```

def entropy(num, denom):
    return -(num/denom) * np.log2(num/denom)

```

```

def get-sorted-data(data, column):
    sort = {}
    for column_name in get-attribute(data, column):
        sort[column_name] = data.loc[data
            [column] == column_name]
    return sort

```

```

def get-attribute(data, column):
    return data[column].unique().tolist()

```

```

def InfoGain(total-entropy, sorted-data,
    entropy-by-attribute):
    length = data-size
    total = 0
    for col, df in sorted-data.items():
        total += (len(df) / length) * entropy-by-
            attribute[col]

```

```
return total_entropy - total
```

```
def get_entropy_by_attribute(sorted_data):  
    entropies = {}  
    for key, df in sorted_data.items():  
        entropies[key] = total_entropy(df, 'Decision')  
    return entropies
```

```
def drop_node(data, column):  
    return data.drop(column, axis=1)
```

```
def id3(tree):  
    for branch, data in tree.items():  
        # make sure it's a DataFrame  
        if not isinstance(data, pd.DataFrame):  
            continue  
        # Fetch column names so you can use them to  
        # iterate later  
        columns = data.columns  
        # Calculate the Entropy for the entire dataset  
        total_entropy_for_data = total_entropy(data,  
                                                values, -1)  
        # If only one column is left, it means we're  
        done.
```

```

if len(columns) == 1:
    break
# keep track of information gain to choose the attribute
with maximum info gain.
info_gain_list = []
# Now iterate over each column to choose the
attribute with maximum info gain
# calculate information gain w.r.t of column
for i in range(0, len(data.columns)-1):
    # sort the rows w.r.t of
    sorted_rows = get_sorted_data(data, column[i])
# Calculate the entropy w.r.t to each attribute based
on sorted column.
entropy_by_attribute = get_entropy_by_attribute(sorted_rows)
# get the info gain
info_gain_list.append(info_gain)
# save it
info_gain = InfoGain(total_entropy_for_data, sorted_rows, entropy_by_attribute)
# Find index of max info gain
node = info_gain_list.index(max(info_gain_list))
# sort the data into branches based on the new
node

```

```

branches = get_sorted_data(data, column[node])
# if we've reached the end of iteration, just
  assign the value,
# else drop the sorted column
  for attr, df in branches.items():
    if (total_entropy(df, column[-1]) == 0):
      branches[attr] = df.iloc[0, -1]
    else:
      branches[attr] = df.drop(column[node],
                               axis=1)


# keep track of nodes already done
  treenode.append(column[node])
# add the new branches to the tree
  child = {column[node]: {}}
  tree[branch] = child
  tree[branch][column[node]] = branches

# ID3
  id3(tree[branch][column[node]])

x = id3(tree)
pprint(tree, depth=5)

```

Seen



Input:

Output:


(0.4444444444444444))
 (0.3333333333333333) 0.0000000000000000)
 (0.2222222222222222) 0.0000000000000000)

Actual Output:

(0.4444444444444444)
 (0.3333333333333333)
 (0.2222222222222222)

Expected Output:

(0.2410424444444444)
 (0.8766666666666666)
 (0.2410424444444444)



u) Build an Artificial Neural Network by implementing the Back propagation algorithm and test the same using appropriate dataset

```
import numpy as np.
```

```
X = np.array([[2,9],[1,5],[3,6]], dtype=float)
```

```
Y = np.array([[92],[86],[89]], dtype=float)
```

```
X = X / np.amax(X, axis=0)
```

```
# maximum of X array longitudinally
```

```
Y = Y / 100
```

```
# Sigmoid Function
```

```
def sigmoid(x):
```

```
    return 1 / (1 + np.exp(-x))
```

```
# Derivative of sigmoid function
```

```
def derivative_sigmoid(x):
```

```
    return x * (1 - x)
```

```
# variable initialization
```

```
epoch = 1000 # setting training iteration
```

```
lr = 0.1 # setting learning rate
```

```
input_layer_neurons = 2 # no. of feature in dataset
```

```
hidden_layer_neurons = 3 # no. of hidden layer neuron
```

```
output_neurons = 1 # no. of neuron of output layer
```


#weight and bias initialization

$w_h = \text{np.random.uniform}(\text{size} = (\text{input_layer_neurons}, \text{hidden_layer_neurons}))$

$b_h = \text{np.random.uniform}(\text{size} = (1, \text{hidden_layer_neurons}))$

$w_{out} = \text{np.random.uniform}(\text{size} = (\text{hidden_layer_neurons}, \text{output_neurons}))$

$b_{out} = \text{np.random.uniform}(\text{size} = (1, \text{output_neurons}))$

#for i in range(epoch):

#forward Propagation

$h_{inp1} = \text{np.dot}(x, w_h)$

$h_{inp} = h_{inp1} + b_h$

$h_{layer_act} = \text{sigmoid}(h_{inp})$

$out_{inp1} = \text{np.dot}(h_{layer_act}, w_{out})$

$out_{inp} = out_{inp1} + b_{out}$

$output = \text{sigmoid}(out_{inp})$

#Back propagation

$E_D = y - output$

$outgrad = \text{derivative_sigmoid}(output)$

$d_output = E_D * outgrad$

$$EH = d_output \cdot \text{dot}(wout, T)$$

$$hiddengrad = \text{derivative_sigmoid}(hlayer - act)$$

#how much hidden layer wk contributed to error

$$d_hiddenlayer = EH * hiddengrad$$

$$wout += hlayer_act \cdot T \cdot \text{dot}(d_output) * lr$$

#dotproduct of nextlayererror and currentlayerop


$$wh += X \cdot T \cdot \text{dot}(d_hiddenlayer) * lr$$

print ("Input : \n" + str(x))

print ("Actual Output : \n" + str(y))

print ("Predicted Output : \n", output)

Seen




Input: 5 pima-indian-diabetes . data.csv

output:

split (768 rows into train = 514 and test = 254 rows)

Accuracy: 0.393700787401574778



print('Actual Output: %s' % y)

print('Predicted Output: %s' % y_hat)

Accuracy

0.39

- 5) Write a program to implement the naive Bayesian classifier for a sample training data set stored as a .CSV file. Compute the accuracy of the classifier, considering few test data sets.

```
import csv
```

```
import random
```

```
import math
```

```
def loadCsv(filename):
```

```
    lines = csv.reader(open(filename, "r"))
```

```
    dataset = list(lines)
```

```
    for i in range(len(dataset)):
```

```
        dataset[i] = [float(x) for x in dataset[i]]
```

```
    return dataset
```

```
def splitDataset(dataset, splitRatio):
```

```
    trainSize = int(len(dataset) * splitRatio)
```

```
    trainSet = []
```

```
    copy = list(dataset)
```

```
    while len(trainSet) < trainSize:
```

```
def separateByClass(dataset):  
    separated = {}  
    for i in range(len(dataset)):  
        vector = dataset[i]  
        if (vector[-1] not in separated):  
            separated[vector[-1]] = []  
            separated[vector[-1]].append(vector)  
    return separated
```

```
def mean(numbers):  
    return sum(numbers) / float(len(numbers))
```

```
def stdev(numbers):  
    avg = mean(numbers)  
    variance = sum([pow(x - avg, 2)  
                    for x in numbers]) / float(len(numbers) - 1)  
    return math.sqrt(variance)
```

```
def summarize(dataset):  
    summaries = [(mean(attribute), stdev(attribute))  
                 for attribute in zip(*dataset)]  
    del summaries[-1]  
    return summaries
```

```
def summarizeByClass(dataset):  
    separated, = separateByClass(dataset)  
    summaries = {}  
    for classValue, instances in separated.items():  
        summaries[classValue] = summarize(instances)  
    return summaries.
```

```
def calculateClassProbabilities(summaries, inputVector):  
    probabilities = {}  
    for classValue, classSummaries in summaries.items():  
        probabilities[classValue] = 1  
        for i in range(len(classSummaries)):  
            mean, stdev = classSummaries[i]  
            x = inputVector[i]
```

```
        probabilities[classValue] *= calculateProbability  
            (x, mean, stdev)
```

```
    return probabilities.
```

```
def predict(summaries, inputVector):  
    probabilities = calculateClassProbabilities  
        (summaries, inputVector)  
    bestLabel, bestProb = None, -1
```

```
for classValue, probability in probabilities.items():  
    if bestLabel is None or probability > bestProb:  
        bestProb = probability  
        bestLabel = classValue  
return bestLabel.
```

```
def getPredictions (summary, testSet):  
    predictions = []  
    for i in range (len (testSet)):  
        result = predict (summary, testSet[i])  
        predictions.append (result)  
    return predictions.
```


```
def getAccuracy (testSet, predictions):  
    correct = 0  
    for i in range (len (testSet)):  
        if testSet [i] [-1] == predictions [i]:  
            correct += 1  
    return correct / float (len (testSet)) * 100.0
```

```
def main():  
    filename = '5_pima-indian-diabetes.data.csv'  
    splitRatio = 0.67
```

```
dataset = loadCsv(filename)
trainingSet, testSet = splitDataset
                           (dataset, splitRatio)
print ('Split {0} rows into train = {1}
and test = {2} rows'.format(len(dataset),
                             len(trainingSet), len(testSet)))
summary1 = summarizeByClass(trainingSet)
prediction = getPrediction(summary1, testSet)
accuracy = getAccuracy(testSet, prediction)
print ('Accuracy: {0} %'.format(accuracy))
```

~~main()~~

Seen



Input : Gpg.csv

Output :

Accuracy metrics

Accuracy of the classifier is 0.4.

Confusion matrix



$\begin{bmatrix} 1 & 2 \end{bmatrix}$

$\begin{bmatrix} 1 & 1 \end{bmatrix}$

Recall and precision.

0.5

0.3333333333



6. Assuming a set of documents that need to be classified, use the naive Bayesian classifier model to perform this task. Built-in Java class/API can be used to write the program. Calculate the accuracy, precision and recall for your data set.

```
import pandas as pd.
```

```
msg = pd.read_csv('6pg.csv', names=['message',  
                                     'label'])
```

```
print('The dimension of the dataset', msg.shape)  
msg['labelnum'] = msg.label.map({'spam': 1,  
                                 'neg': 0})
```

```
x = msg.message
```

```
y = msg.labelnum
```

```
print(x)
```

```
print(y)
```

```
# splitting the dataset into train and test data  
from sklearn.model_selection import train_test_split  
x_train, x_test, y_train, y_test = train_test_split(x, y)
```

```
print(x_test.shape)
```

```
print(x_train.shape)
```

```
print(y_test.shape)
```

```
print(y_train.shape)
```

#output of count vectorizer is a sparse matrix from sklearn.feature_extraction.text

```

import CountVectorizer
count_vect = CountVectorizer()
x_train_dtm = count_vect.fit_transform(x_train)
x_test_dtm = count_vect.transform(x_test)
print(count_vect.get_feature_names())
df = pd.DataFrame(x_train_dtm.toarray(),
                  column=count_vect.get_feature_names())
print(df) # tabular representation
print(x_train_dtm) # sparse matrix representation
# Training Naive Bayes (NB) classifier on
# training data.

```

```

from sklearn.naive_bayes import MultinomialNB
clf = MultinomialNB().fit(x_train_dtm, y_train)
predicted = clf.predict(x_test_dtm)
# printing accuracy metrics
from sklearn import metrics
print('Accuracy metrics')
print('Accuracy of the classifier', metrics.
      accuracy_score(y_test, predicted)).
print('confusion matrix')
print(metrics.confusion_matrix(y_test, predicted))

```

```
print('Recall and Precision')  
print(metrics.recall_score(ytest, predicted))  
print(metrics.precision_score(ytest, predicted))
```

Goal

$$\frac{10}{10}$$

1) Write a program to construct a Bayesian network considering medical data. Use this patient to demonstrate the diagnosis of heart disease using standard Heart Disease Data Set. You can use Java/Python ML library class/API.

```
import bayespy as bp
import numpy as np
import csv
from colorama import init
from colorama import Fore, Back, Style
init()
```

```
ageEnum = {'SuperSeniorCitizen': 0, 'SeniorCitizen': 1,
           'MiddleAged': 2, 'Youth': 3, 'Teen': 4}
```

```
genderEnum = {'Male': 0, 'Female': 1}
```

```
familyHistoryEnum = {'Yes': 0, 'No': 1}
```

```
dietEnum = {'High': 0, 'Medium': 1, 'Low': 2}
```

```
lifestyleEnum = {'Athlete': 0, 'Active': 1, 'Moderate': 2,
                 'Sedentary': 3}
```

```
cholesterolEnum = {'High': 0, 'Borderline': 1, 'Normal': 2}
```

```
heartDiseaseEnum = {'Yes': 0, 'No': 1}
```

Output

```
Enter age = 4 { 'SuperSeniorCitizen': 0, 'SeniorCitizen': 0, 'MiddleAged': 2, 'Youth': 3, 'Teen': 4 }  
Enter Gender = { 'male': 0, 'Female': 13 }  
Enter Family history = { 'Yes': 0, 'No': 13 }  
Enter diet Num = { 'High': 0, 'Medium': 1, 'Low': 13 }  
Enter Lifesty = { 'Sedentary': 0, 'Active': 1, 'Moderate': 13 }
```

```
# reading csv file
with open('heart_disease_data.csv') as csvfile:
    # creating a csv reader object.
    lines = csv.reader(csvfile)
    dataset = list(lines)
    data = []
    for x in dataset:
        data.append([ageEnum[x[0]], genderEnum[x[1]],
                    familyHistoryEnum[x[2]], dietEnum[x[3]],
                    lifestyleEnum[x[4]], cholesterolEnum[x[5]],
                    heartDiseaseEnum[x[6]])
```

```
data = np.array(data)
```

```
N = len(data)
```

```
# Group assignment probabilities.
```

```
p_age = bp.nodes.Dirichlet(1.0 * np.ones(5))
```

```
# Group assignments for nodes.
```

```
age = bp.nodes.Categorical(p_age, plates=(N, 1))
```

```
age.observe(data[:, 0])
```

```
p_gender = bp.nodes.Dirichlet(1.0 * np.ones(2))
```

```
gender = bp.nodes.Categorical(p_gender, plates=(N, 1))
```

```
gender.observe(data[:, 1])
```

$p_family_history = bp_node . Dirichlet(1.0 * np . ones(2))$
 $family_history = bp_node . Categorical(p_family_history,$
 $plate1 = (N, 1))$
 $family_history . observe(data[:, 2])$

$p_diet = bp_node . Dirichlet(1.0 * np . ones(3))$
 $diet = bp_node . Categorical(p_diet, plate1 = (N, 1))$
 $diet . observe(data[:, 3])$

$p_lifestyle = bp_node . Dirichlet(1.0 * np . ones(4))$
 $lifestyle = bp_node . Categorical(p_lifestyle,$
 $plate1 = (N, 1))$
 $lifestyle . observe(data[:, 4])$

$p_cholesterol = bp_node . Dirichlet(1.0 * np . ones(3))$
 $cholesterol = bp_node . Categorical(p_cholesterol,$
 $plate1 = (N, 1))$
 $cholesterol . observe(data[:, 3])$

~~$p_heart_disease = bp_node . Dirichlet(np . ones(2),$~~
 ~~$plate1 = (5, 2, 2, 3, 4, 3))$~~

$heart_disease = bp_node . Multinomial([age, gender,$
 $family_history, diet, lifestyle, cholesterol],$
 $bp_nodes . Categorical, p_heart_disease)$


```
heart disease.observe (data [:,6])
```

```
p_heart disease.update()
```

```
m = 0
```

```
while m == 0:
```

```
    printf ("\n")
```

```
    res = bp.nodemultimixture([int(input
```

```
        'Enter Age: ' + str(ageEnum)),
```

```
        int(input('Enter Gender: ' + str(genderEnum))),
```

```
        int(input('Enter FamilyHistory: ' +
```

```
            str(familyHistoryEnum)), int(input('
```

```
            Enter dietEnum: ' + str(dietEnum))), int(
```

```
            input('Enter Lifestyle: ' + str(lifeStyleEnum))),
```

```
            int(input('Enter cholesterol: ' + str(cholesterolEnum)))]
```

```
    bp.nodem.categorical, p_heart disease).get-  
    moments()[0][heartDiseaseEnum['No']]
```

```
    print('Probability(HeartDisease) = ' + str(res))
```



```
    m = int(input('Enter for Continue: 0, Exit: 1'))
```

Seen
10/10


Input : Iris. data

output

Real cluster



Small clustering



8. Apply EM algorithm to cluster a set of data stored in a .csv file. Use the same data set for clustering using K-means algorithm. Compare the results of these two algorithms and comment on the quality of clustering. You can add Java/Python ML library class/API in the program.

```
import matplotlib.pyplot as plt
from sklearn import datasets
from sklearn.cluster
import KMeans
import pandas as pd
import numpy as np
```

```
#import some data to play with
```

```
iris = datasets.load_iris()
```

```
X = pd.DataFrame(iris.data)
```

```
X.columns = ['Sepal_Length', 'Sepal_Width', 'Petal_Length', 'Petal_Width']
```

```
Y = pd.DataFrame(iris.target)
```

```
Y.columns = ['Targets']
```

```
#Build the K means model
model = Kmeans(n_clusters = 3)
model.fit(x) #model.labels_ : Give cluster no
for which samples belongs to
```

```
#Visualise the clustering result
plt.figure(figsize = (14, 14))
colormap = np.array(['red', 'lime', 'black'])
```

```
#Plot the original classifications using Petal feature
plt.subplot(2, 2, 1)
plt.scatter(X.Petal_Length, X.Petal_Width,
            c = colormap[Y.Target], s = 40)
plt.title('K-Means Clustering')
plt.xlabel('Petal Length')
plt.ylabel('Petal Width')
```

```
#General EM for GMM
```

```
from sklearn import preprocessing
scaler = preprocessing.StandardScaler()
scaler.fit(X)
xsa = scaler.transform(X)
xs = pd.DataFrame(xsa, columns = X.columns)
```

```
from sklearn.mixture import GaussianMixture
gmm = GaussianMixture(n_components=2)
gmm.fit(xs)
gmm-y = gmm.predict(xs)
```

```
plt.subplot(2, 2, 3)
plt.scatter(x.Petal_Length, x.Petal_Width,
            c=colormap[gmm-y], s=40)
plt.title('GMM Clustering')
plt.xlabel('Petal Length')
plt.ylabel('Petal Width')
```

```
print('Observation: The GMM using EM
algorithm based clustering matched the
true labels more closely than the
Kmeans.')
```

~~Scanned~~

10/10

11/11/11

11/11/11

0.93333333

0.975


117.7, 26.69, 2.37

(11/11/11)

11/11/11

11/11/11

11/11/11



9. Write a program to implement K-nearest Neighbour algorithm to classify the iris dataset. Print both correct and wrong predictions. Java/Python ML library class can be used for this problem

```
import numpy as np
import matplotlib.pyplot as plt
import pandas as pd.
```

```
#Read dataset to pandas dataframe
dataset = pd.read_csv('iris.csv')
```

```
X = dataset.iloc[:, :-1].values
Y = dataset.iloc[:, 4].values
```

```
from sklearn.model_selection import train_test_split
X_train, X_test, Y_train, Y_test =
    train_test_split(X, Y, test_size = 0.20)
plt.plot(X_train, Y_train, 'b.', X_test, Y_test,
         'r.')
```

```

from sklearn.neighbors import KNeighborsClassifier
classifier = KNeighborsClassifier(n_neighbors=5)
classifier.fit(X_train, Y_train)

```

```

accuracy = classifier.score(X_test, Y_test)
accuracy1 = classifier.score(X_train, Y_train)

```

```

print(accuracy)
print(accuracy1)

```

```

example = np.array([7, 7, 2, 6, 6, 9, 2, 3])
example = example.reshape(1, -1)
print(example)

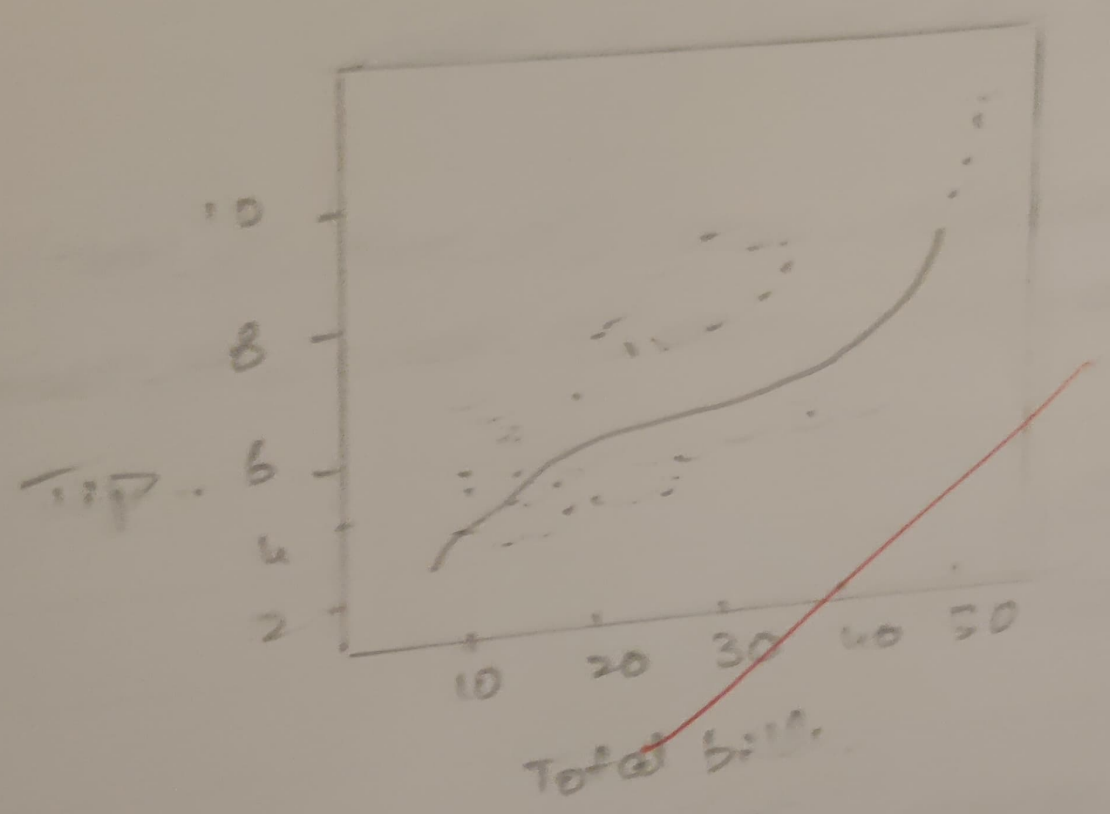
```

Get

10
10

Grand = data

output



10) Implement the non-parametric Locally Weighted Regression algorithm in order to fit data point. Selection appropriate data set for your experiment and draw graph.

```
import matplotlib.pyplot as plt
import pandas as pd
import numpy as np.
```

```
def kernel(point, xmat, k):
    m, n = np.shape(xmat)
    weights = np.mat(np.eye(m))
    #eye - identity matrix
    for j in range(m):
        diff = point - x[j]
        weights[j, j] = np.exp(-diff * diff / (2 * k ** 2))
    return weights
```

```
def local_weight(point, xmat, ymat, k):
    wei = kernel(point, xmat, k)
    w = (x.T * (wei * x)) * 2 * (x.T *
        (wei * ymat.T))
    return w
```

```
def localWeightRegression(xmat, ymat, k):  
    m, n = np.shape(xmat)  
    ypred = np.zeros(m)  
    for i in range(m):  
        ypred[i] = xmat[i] * localWeight  
            (xmat[i], xmat, ymat, k)  
    return ypred
```

```
def graphplot(x, ypred):  
    sortindex = x[:, 1].argsort(0)  
    # argsort - index of the smallest  
    xsort = x[sortindex][:, 0]  
    fig = plt.figure()  
    ax = fig.add_subplot(1, 1, 1)  
    ax.scatter(bill, tip, color = 'green')  
    ax.plot(xsort[:, -1], ypred[sortindex],  
            color = 'red', line width = 5)  
    plt.xlabel('Total bill')  
    plt.ylabel('Tip')  
    plt.show();
```

```
# load data points  
data = pd.read_csv('data10-tip.csv')
```

bill = np.array (data, total-bill)

we use only Bill amount and Tips data.

tip = np.array (data, tip)

mbill = np.mat (bill)

.mat will convert nd array is converted
2D array.

mtip = np.mat (tip)

m = np.shape (mbill) [1]

one = np.mat (np.ones (m))

x = np.hstack ((one.T, mbill.T))

2x4 rows, 2 col.

ypred = localWeightRegression (x, mtip, 8)

increase k to get smooth curve

graphPlot (x, ypred).

~~10~~

10
10

Bayesian.

Jeevan R
4BW17CS029
CSE A

Assuming a set of documents that need to be classified, use the naive Bayesian classifier model to perform this task. Built-in Java class/API can be used to write the program. Calculate the accuracy, precision and recall for your data set.

```
import numpy as np
import pandas as pd
msg = pd.read_csv('6pg.csv', name = ['msg', 'label'])
x = np.array([[2.9], [1.5], [3.6]], dtype = float)
y = np.array([8.6, 8.9, 7.5], dtype = float)

x = x / (np.amax(x, axis=0) + np.exp(-x))
y = y / 100
epoch = 7000
lr = 0.1 # learning rate

input_neurons = 2
hidden_neurons = 3
output_neurons = 1

wh = np.uniform.random(size = (input_neurons, hidden_neurons))
bh = np.uniform.random(size = (1, hidden_neurons))
wout = np.uniform.random(size = (hidden_neurons, output_neurons))
bout = np.uniform.random(size = (1, output_neurons))

for i in range(epoch):
```

x = msg.msg
y = msg.label
print(x)
print(y)
from sklearn.model_selection import train_test_split

h_in_p1 = np.dot(Xt, w_h) from sklearn.linear_model

w_in_p = b_h + h_in_p1

h_out = accuracy(h_in_p)

b_in_p1 = np.dot(h_in_p, w_h)

b_in_p = h_out + b_in_p1

b_out = derivativu = accuracy(b_in_p)

E0 = h_out + b_out

h_out = E0 * b_out

d_output = E0 + h_out

print(x_test, shape)
print(x_train, shape)
print(y_test, shape)
print(y_train, shape)

def accuracy(x):

return np.mean(x) / (1 + np.exp(-x))

def derivativu - accuracy(x)

return sklearn.metrics.log_loss(1 - x) - multinomial

from sklearn.metrics import log_loss, multinomial

predict_d = dff.predicted_out_dry

print('input: ' + str(x))

print('Precision: ' + str(y))

print('Actual Output: ' + str(output))

print('Accuracy metrics')

print('Accuracy of the classifier: ' + str(accuracy))

print('Precision, confusion-matrix (y_test, predicted)')

print('Recall and precision')

print('media, recall, cobefit, predicted')

print('media, precision, recall, predicted')

The dimensions of the dataset (18, 21)

Name: `movie`, dtype: object

0	1
1	1
2	1
3	1
4	0
5	0
6	0
7	0
8	0
9	0
10	1
11	0
12	1
13	0
14	1
15	0
16	1
17	0

Name: `labelnum`, dtype: int64

- (5,)
- (13,)
- (5,)

['amagyu', 'an', 'awesome', 'bad', 'best', 'boy', 'can', 'dave', 'dead', 'no', 'energy', 'fun', 'good', 'great', 'have', 'he', 'holiday', 'horrible', 'how', 'is', 'like', 'locality', 'love', 'my']

Accuracy metrics

Accuracy of the classifier is 0.6

Confusion matrix

$\begin{bmatrix} 1 & 2 \\ 2 & 1 \end{bmatrix}$

$\begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$

Recall and Precision

1.0

0.5