LUMS
A Not-for-Profit University

Department of Electrical Engineering
School of Science and Engineering

EE514/CS535 Machine Learning

HOMEWORK 1

Due Date: 23:55, Tuesday, March 02, 2021 (Submit online on LMS)
Format: 7 problems, for a total of 100 marks
Contribution to Final Percentage: 2.5%
Instructions:

- Each student must submit his/her own hand-written assignment, scanned in a single PDF document.
- You are allowed to collaborate with your peers but copying your colleague's solution is strictly prohibited. Anybody found guilty would be subjected to disciplinary action in accordance with the university rules and regulations.
- Note: Vectors are written in lowercase and bold in the homework, for your written submission kindly use an underline instead. In addition, use capital letters for matrices and lowercase for scalars.

Problem 1 (20 marks)

(**Note:** Compile and submit screenshots of your plots for this question.)

Polynomial Regression - Polynomial regression is a form of regression analysis in which the relationship between the independent variable and the dependent variable is modeled as an n-th degree polynomial. The model equation that relates the input to the output is of the form:

$$y_i(x_i) = \theta_0 + \theta_1 x_i + \theta_2 x_i^2 + \cdots + \theta_M x_i^M.$$

(a) [**4 marks**] Assuming we use an M^{th} degree polynomial for our regression analysis and we have N outputs $y_1, y_2, \ldots, y_N$, express the regression model in the matrix form $\boldsymbol{y} = A\boldsymbol{\theta}$.

(b) [**8 marks**] We will now implement the regression model on one feature variable using programming tools. Use the following code to generate some toy data:

```
import pandas as pd
xdic={'x': {11: 300, 12: 170, 13: 288, 14: 360, 15: 319, 16: 330,
17: 520, 18: 345, 19: 399, 20: 479}}
ydic={'y': {11: 305, 12: 270, 13: 360, 14: 370, 15: 379, 16: 400,
17: 407, 18: 450, 19: 450, 20: 485}}
x=pd.DataFrame.from_dict(xdic)
```

```
y=pd.DataFrame.from_dict(ydic)
import numpy as np
x_seq = np.linspace(x.min(),x.max(),300).reshape(-1,1)
```

Next, use the scikit-learn library to implement polynomial regression of degree 2 on the data. Finally, use the following code to display the data points and the regression curve:

```
import matplotlib.pyplot as plt
plt.figure()
plt.scatter(x,y)
plt.plot(x_seq,polyreg.predict(x_seq),color="black")
plt.title("Polynomial regression with degree "+str(degree))
plt.show()
```

(c) [**8 marks**] Repeat the regression process for degrees 1 (i.e., linear regression), 3, 4, and 5. What problem do you observe as the degree of the polynomial increases? (Hint: think about the trained model's accuracy when implemented on a test dataset.)

Problem 2 (10 marks)

Least Squares Formulation - In this question we will derive the least squares solution from the perspective of an optimization problem.

(a) [**2 marks**] First let us formulate the problem in terms of an optimization problem. Consider the following system: The correct output is given by $\boldsymbol{y} \in \mathbb{R}^n$. The measured output $\hat{\boldsymbol{y}} \in \mathbb{R}^n$ for a total of m instances is given by $A = [\hat{\boldsymbol{y}}_{\boldsymbol{1}}, \hat{\boldsymbol{y}}_{\boldsymbol{2}}, \ldots, \hat{\boldsymbol{y}}_{\boldsymbol{m}}] \in \mathbb{R}^{n \times m}$ (where $n > m$). Using these m instances, we want to determine the best approximation of $\boldsymbol{y}$ by taking a weighted sum of these m measured outputs, where the weights are given by the vector $\boldsymbol{\theta} \in \mathbb{R}^m$.

The weight vector $\boldsymbol{\theta} \in \mathbb{R}^m$ minimizes the mean-square error between the correct and measured output. We can define this as an optimization problem as shown:

$$\boldsymbol{\theta}^* = \underset{\boldsymbol{\theta}}{\text{minimize}} \quad f(\boldsymbol{\theta}),$$

where $f(\boldsymbol{\theta})$ is the cost function which we aim to minimize. For this problem, write down the expression for $f(\boldsymbol{\theta})$.

(b) [**8 marks**] Find the expression for $\boldsymbol{\theta}^*$ that minimises the cost function $f(\boldsymbol{\theta})$ by computing its gradient.
Note: You might find the following relationships useful for this proof:

$$\|\boldsymbol{a}\|_2^2 = \boldsymbol{a}^T \boldsymbol{a}$$
$$(AB)^T = B^T A^T$$

You will also need to review derivatives of matrices and vectors.

This final expression is also known as the **Left Pseudo-Inverse**, verify your answer by looking up the formula.

Problem 3 (15 marks)

Multivariate Gradient Descent - You are presented with the following feature vector x, the parameter vector θ, and the gold label y. The initial values of the parameter vector are also given, and assume x_0 to be the bias absorbed into the feature vector.

$$\boldsymbol{x} = \begin{bmatrix} x_0 \\ x_1 \\ x_2 \\ x_3 \\ x_4 \end{bmatrix} = \begin{bmatrix} 1 \\ 2 \\ 12 \\ 7 \\ 5 \end{bmatrix}, \quad \boldsymbol{\theta} = \begin{bmatrix} \theta_0 \\ \theta_1 \\ \theta_2 \\ \theta_3 \\ \theta_4 \end{bmatrix} = \begin{bmatrix} 0.7 \\ 2.0 \\ -0.2 \\ 0.3 \\ 0.7 \end{bmatrix}, \quad \boldsymbol{y} = 7$$

Perform 3 iterations of batch gradient descent on the dataset. Assume learning rate $\alpha = 0.001$ and round your answers to 4 decimal places. Show all your steps and the mathematical equations that you use.

Problem 4 (10 marks)

Hamming Distance - For categorical data in k-NN classification, we often use the Hamming distance as our distance metric. In the table provided below, you are provided the binary vectors for a set of colours, and their frequency in the training data:

Colour	Binary Vector				No. in Train Data
	1	2	3	4	
White	0	0	0	1	2
Pink	1	0	0	1	1
Red	1	0	0	0	2
Orange	1	1	0	0	3
Yellow	0	1	0	0	3
Green	0	1	1	0	1
Blue	0	0	1	0	1
Purple	1	0	1	0	2

Table 1: Binary coded vectors for different colours

(a) [**4 marks**] Suppose you are provided an unknown colour, defined by the vector: '1011'. Run the k-NN algorithm with $k = 3$ using Hamming distance to predict the colour. Show all of your working.

(b) [**6 marks**] We can also use Hamming distance to compute the error in the final prediction. For the test data tabulated below: Run the k-NN algorithm for $k = 5$ and

Test Data Point	Binary Vector				No. in Test Data
	1	2	3	4	
X	0	0	1	1	2
Y	0	1	0	1	1
Z	0	1	1	0	2

Table 2: Binary coded vectors for test data

calculate the prediction error using Hamming distance.
Hint: Make sure to average out the total error across all the test data points!

Problem 5 (15 marks)

(**Note:** Compile and submit screenshots of your plots for this question.)

Curse of Dimensionality - In this question we will look closely at the problem associated with higher dimensions in hypercubes when using k-NN. During your lectures, you learnt that as the dimensions increase, the neighbours in a constant vicinity decrease. Let us now visualise this phenomenon through simulations.
The expression for average distance in a unit cube $[0, 1]^3$ is given by:

$$\int_0^1 \int_0^1 \int_0^1 \int_0^1 \int_0^1 \int_0^1 \sqrt{(x_1 - x_2)^2 + (y_1 - y_2)^2 + (z_1 - z_2)^2} dx_1 dx_2 dy_1 dy_2 dz_1 dz_2$$

Solving this expression is really tedious, and as dimensions increase, it becomes harder to solve and visualise as well. Instead, we will numerically calculate average distance by plotting a frequency histogram of distances.

(a) [**10 marks**] You are required to build a function that will randomly take two points from a $n-$dimensional hypercube and calculate the Euclidean distance for $i = 10000$ iterations. You are given the following starter code:

```
import matplotlib.pyplot as plt
%matplotlib inline
import math
import random

def freq_plot(n_dim, iterations = 10000):
dist = []
# code required here

plt.figure()
plt.hist(dist, range = [0, math.sqrt(n_dim)], bins=100,
density = True)
title = 'n = ' +  str(n_dim)
plt.gca().set(title=title, xlabel='Distances', ylabel='Frequency')
return
```

Use the **random.random()** function for generating a float between 0 and 1.
Attach a screenshot of your final working function code.

(b) [**5 marks**] Run the function in part (b) for n_dim $= 2, 3, 5, 10, 100, 1000$, and attach screenshots of your plots.

Problem 6 (15 marks)

KD Trees - One of the most commonly used k-NN algorithms, KD trees are space-partitioning data structures that divide a feature set based on some particular attributes, thereby reducing the time complexity of the k-NN algorithm.

(a) [**10 marks**] The following table contains data points $\mathbf{x} = (x_1, x_2)$ from a distribution $\in \mathbb{R}^2$. Use these data points to create a KD tree. Begin your solution by splitting using the x_1 feature, and then alternate between the 2 features at each level of the tree when determining the next split. Each child node should have 2-3 data points.

x_1	1	3	12	7	8	13	2	21	5	17	31	4	47	19	1	9	14
x_2	6	8	9	21	23	7	12	9	1	3	16	20	10	11	4	31	39

(b) [**5 marks**] What major issue can you observe when using KD trees to classify a test data point?

Problem 7 (15 marks)

Principal Component Analysis - In class we studied PCA as a means to reduce dimensionality, by minimizing the squared error of predictions in lower dimension, also known as projections onto a range space, as compared to the actual values. In this question, we will prove that this problem is equivalent to maximizing the squared length of these projections, i.e., maximizing variance of projected data.

Let us define the original problem from the class lecture. Considering n feature vectors of the form $\boldsymbol{x} \in \mathbb{R}^d$. Using only k basis vectors, we want to project $\boldsymbol{x}$ in a new space of lower dimensionality, from $\mathbb{R}^d$ to $\mathbb{R}^k$:

$$\boldsymbol{z} = W^T \boldsymbol{x},$$

where W is the orthogonal mapping matrix of size $d \times k$. To obtain a approximation of $\boldsymbol{x}$ we use the following reconstruction:

$$\hat{\boldsymbol{x}} = W\boldsymbol{z},$$

which is equivalent to:

$$\hat{\boldsymbol{x}} = WW^T\boldsymbol{x}$$
$$\hat{\boldsymbol{x}} = P\boldsymbol{x},$$

where $P = WW^T$ is the projection operator, that is idempotent: has the following properties: $P^2 = P = P^T$.
Our objective is to minimize the sum of squared error during reconstruction, we can write it as follows:

$$\text{minimize} \quad \|\boldsymbol{x} - P\boldsymbol{x}\|_2^2$$

(a) [**10 marks**] Show that this expression can be reduced to:

$$\text{minimize} \quad [\|\boldsymbol{x}\|_2^2 - \|P\boldsymbol{x}\|_2^2]$$

Note: You might find the following relationships useful for this proof:

$$\|\boldsymbol{a}\|_2^2 = \boldsymbol{a}^T\boldsymbol{a}$$
$$(AB)^T = B^T A^T$$

(b) [**5 marks**] Using the final expression in part (a) and the Pythagorean Theorem on $\|\boldsymbol{x}\|_2^2$ to show that this problem is equivalent to maximizing the variance of projected data. You may draw a diagram to augment your point.

— End of Homework —