This question paper consists of TWO (2) sections. Answer ALL questions in the answer booklet provided. **[100 MARKS]**

SECTION A **(40 Marks)**

There are TWENTY (20) questions in this section. Answer ALL questions in the answer booklet.

1. In Python, how can you import a specific function or attribute from a module?

   A. import module
   B. from module import *
   C. from module import function
   D. import function from module

2. Which of the following is **NOT** a valid Python data type?

   A. Integer
   B. Float
   C. Character
   D. String

3. In Python, how do you create a list containing elements [1, 2, 3]?

   A. list = (1, 2, 3)
   B. list = [1, 2, 3]
   C. list = {1, 2, 3}
   D. list = "1, 2, 3"

4. What is the primary purpose of Pandas in data analysis?

   A. Data visualization
   B. Machine learning
   C. Data manipulation and analysis
   D. Web development

5. How can you import the Pandas library in Python?

   A. import pandas as pd
   B. from pandas import Pandas
   C. use Pandas
   D. import pandas

6. Which of the following is **NOT** a valid way to create a DataFrame in Pandas?

   A. pd.DataFrame(data)
   B. pd.DataFrame(data, columns=['A', 'B'])
   C. pd.DataFrame.from_csv('data.csv')
   D. pd.read_csv('data.csv')

7. Which of the following is **NOT** a valid Python comment?

   A. # This is a comment
   B. '''This is a comment'''
   C. // This is a comment
   D. """This is a comment"""

8. What is the purpose of the plt.title() function in Matplotlib?

   A. Set the title of the current Axes
   B. Plot a title on the x-axis
   C. Create a new figure
   D. Add a legend to the plot

9. How can you display a seaborn plot in Python?

   A. Using the show() function
   B. Using the plot() function
   C. Using the display() function
   D. Seaborn plots are automatically displayed

10. What does the head() function in Pandas do?

A. Returns the last 5 rows of a DataFrame
B. Returns the first 5 rows of a DataFrame
C. Sorts the DataFrame in ascending order
D. Computes the mean of the DataFrame

11. What does the len() function do in Python?

A. Returns the last element of a list
B. Returns the length of a list or string
C. Returns the maximum value in a list
D. Returns the square root of a number

12. Which of the following is an example of a Boolean value in Python?

A. "True"
B. 0
C. 1.5
D. True

13. What is the purpose of the range() function in Python?

A. Generate a sequence of numbers
B. Calculate the sum of a list
C. Create a new list
D. Calculate the average of a list

14. Which of the following is a supervised learning problem?

A. Image classification
B. Clustering customer behavior
C. Anomaly detection
D. Reinforcement learning

15. Which Python library is commonly used for implementing machine learning algorithms?

A. Pandas
B. NumPy
C. Scikit-Learn
D. Matplotlib

16. What is the difference between classification and regression in machine learning?

A. Classification predicts discrete labels; regression predicts continuous values.
B. Classification predicts continuous values; regression predicts discrete labels.
C. Classification and regression are the same thing.
D. Regression is only used for deep learning.

17. What is the purpose of the "train-test split" in machine learning?

A. To train the model on all available data
B. To test the model on the same data it was trained on
C. To split the data into a training set and a testing set
D. To increase model complexity

18. What is the term for the process of converting categorical data into numerical form in machine learning?

A. Binarization
B. Normalization
C. Encoding
D. Dimensionality reduction

19. What is the purpose of a confusion matrix in machine learning?

A. To evaluate the performance of a regression model
B. To visualize the decision boundary of a classification model
C. To assess the accuracy of a clustering algorithm
D. To summarize the performance of a classification model

20. Which machine learning algorithm is inspired by the way neurons work in the human brain?

A. Support Vector Machine
B. Decision Tree
C. K-Nearest Neighbors
D. Artificial Neural Network

SECTION B **(60 Marks)**

There are THREE (3) questions in this section. Answer ALL questions in the answer booklet.

Question 1 (20 marks)

Write a Python function called calculate_average that takes a list of numbers as input and returns the average (mean) of those numbers. Provide an example of how you would use this function with a list of numbers.

Question 2 (20 marks)

You have a CSV file containing a dataset with columns for "Name," "Age," and "Salary." Using Pandas, write Python code to read the CSV file into a DataFrame, filter the data to include only individuals aged 30 or older, and calculate the average salary for this group.

Question 3 (20 marks)

Discuss **THREE (3)** commonly used evaluation metrics for classification problems in machine learning (e.g., accuracy, precision, recall, F1-score). Explain the scenarios in which each metric is most appropriate and why.

***** END OF QUESTION PAPER *****