QPC: RN20BTECH635 AR 20 Reg. No | | | | | | | | | | |

GIET UNIVERSITY, GUNUPUR – 765022

B. Tech (Seventh Semester – Regular) Examinations, November – 2023

BPECS7011 – Statistical Machine Learning

(CSE)

Time: 3 hrs Maximum: 70 Marks

Answer ALL Questions

The figures in the right hand margin indicate marks.

PART – A: (Multiple Choice Questions) **(1 x 10 = 10 Marks)**

Q.1. Answer ***ALL*** questions		[CO#]	[PO#]
a.	Logistic regression is a _____ regression technique that is used to model data having a _____ outcome (i) Linear, Binary (ii) Linear , numeric (iii) Non-linear, binary (iv) Non-linear ,numeric	CO1	PO3
b.	What machine learning models are trained to make a series of decisions on based on the rewards and feedback they receive for their actions? (i) Supervised (ii) Unsupervised (iii) Reinforcement (iv) All the above	CO1	PO2
c.	What is the purpose of ridge regression in machine learning? (i) To handle multicollinearity in the dataset. (ii) To reduce the sum of squared residuals. (iii) To perform feature selection. (iv) To handle missing values in the dataset.	CO2	PO2
d.	Which assumption is commonly made in linear regression modeling? (i) The dependent variable is normally distributed. (ii) The dependent variable is categorical in nature. (iii) The residuals have a constant variance. (iv) The independent variables are strongly correlated with each other.	CO1	PO2
e.	What is the key assumption made by the Naive Bayes algorithm? (i) Independence of features (ii) High dimensionality (iii) Non-parametric nature (iv) Deterministic outcomes	CO3	PO3
f.	In K-nearest neighbors (KNN) algorithm, how is the majority voting performed? (i) Each neighbor has an equal vote. (ii) Closer neighbors have higher weights. (iii) Each neighbor votes based on its distance. (iv) The neighbor with the highest probability votes.	CO3	PO2
g.	In Artificial Neural Networks (ANN), what is the purpose of the activation function? (i)) Feature extraction (ii) Non-linearity introduction (iii) Dimensionality reduction (iv) Weight initialization	CO4	PO2
h.	In SVM, what is the role of support vectors in the decision boundary? (i) Determine the weights of the features (ii) Define the margin and decision boundary (iii) Act as outliers to be ignored (iv) Influence the regularization parameter	CO4	PO2
i.	What is the purpose of support vector machines (SVM)? (i) To maximize the margin between different classes. (ii) To minimize the training error in classification. (iii) To handle imbalanced datasets. (iv) To perform feature selection.	CO4	PO2
j.	In Bayes' theorem, what does P(A|B) represent? (i) Joint probability of A and B (ii) Conditional probability of A given B (iii) Marginal probability of A (iv) Prior probability of A	CO3	PO2

PART – B: (Short Answer Questions) **(2 x 10 = 20 Marks)**

Q.2. Answer ***ALL*** questions	[CO#]	[PO#]
a. Explain the concept of overfitting in machine learning.	CO1	PO2
b. What is the objective of dimensionality reduction techniques?	CO1	PO1
c. How does ridge regression differ from ordinary least squares regression?	CO2	PO1
d. How do logistic regression and random forest differ in terms of interpretability?	CO2	PO1
e. Discuss the importance of data cleaning and pre-processing in machine learning.	CO1	PO2
f. Explain the concept of hyper parameters in machine learning algorithms.	CO1	PO2
g. How does the KNN algorithm work in a classification task? Provide a brief overview	CO3	PO2
h. Explain the concept of joint probability and provide an example.	CO3	PO2
i. Describe the difference between the Support Vector Classifier and the Support Vector Machine.	CO4	PO1
j. How does deep learning software facilitate the development and deployment of complex neural network architectures?	CO4	PO1

PART – C: (Long Answer Questions) **(10 x 4 = 40 Marks)**

Answer ***ALL*** questions	Marks	[CO#]	[PO#]
3. a. Compare and contrast supervised learning and unsupervised learning. Provide examples of each.	5	CO1	PO3
b. Discuss the bias-variance trade-off in machine learning. How does it impact model performance?	5	CO1	PO2
(OR)			
c. Explain the concepts of underfitting and overfitting in machine learning. How can these issues be addressed?	5	CO1	PO2
d. Explain the steps involved in developing and deploying a machine learning model. Discuss the significance of each step in the model development lifecycle.	5	CO1	PO2
4. a. Compare and contrast logistic regression and random forest models for classification tasks.	5	CO2	PO3
b. Explain the variable importance plot in random forest models. Discuss how it helps identify the most influential features and how it can be used for feature selection. Provide an example of creating and interpreting a variable importance plot using a real-world dataset.	5	CO2	PO3
(OR)			
c. Compare and contrast the interpretability of logistic regression and random forest models. Discuss the trade-offs between the two models in terms of ease of interpretation and understanding the impact of features on the outcome. Support your answer with examples and real-world applications.	5	CO2	PO3
d. Consider the given dataset. Apply the Naïve Bayes algorithm and predict that if the fruit has the following properties then which type of fruit it is?	5	CO2	PO3

Fruit	Yellow	Sweet	Long	Total
Mango	350	450	0	650
Banana	400	300	350	400
Others	50	100	50	150
Total	800	850	400	1200

5. a.	Discuss the Curse of Dimensionality with examples in 1D, 2D, and 3D space. How does it impact the performance of machine learning algorithms?	5	CO3	PO1
b.	Discuss the significance of joint probability in understanding the relationships between variables. Provide real-world examples.	5	CO3	PO1
	(OR)			
c.	Discuss the challenges and limitations of using Naive Bayes for SMS spam classification, and propose potential solutions.	5	CO3	PO1
d.	Given a confusion matrix for a Naive Bayes SMS spam classifier with 150 true positives, 20 false positives, 30 false negatives, and 800 true negatives, calculate precision, recall, and F1 score.	5	CO3	PO2
6. a.	Perform a forward propagation for a simple artificial neural network with one input layer, one hidden layer with three neurons, and one output layer. Assume the activation function is sigmoid, and the input values are (0.5, 0.8, 0.2). Calculate the output of the neural network.	5	CO4	PO2
b.	Explain the concept of stochastic gradient descent (SGD) and its role in optimizing neural networks.	5	CO4	PO1
	(OR)			
c.	Discuss the impact of hyper parameter tuning on the performance of deep neural networks. Include specific examples of hyper parameters and their influence on model outcomes.	5	CO4	PO2
d.	Describe the working principles of artificial neural networks (ANN). Explain the concepts of forward propagation and backpropagation in the training process of neural networks. Discuss the optimization techniques used for training neural networks.	5	CO4	PO2

--- End of Paper ---