

Roll No.....

Total No. of Printed Pages: 1

Total No. of Questions: [11]

M. Tech. CSE (Semester – 1st)

MACHINE LEARNING

Subject Code: MCSCE1156

Paper ID: [190204]

Time: 03 Hours

Maximum Marks: 60

Instruction for candidates:

1. Section A is compulsory. It carries 16 marks. It consists of 4 questions of 4 marks each.
2. Section B consist of 4 questions of 8 marks each. The student has to attempt any 3 questions out of it.
3. Section C consist of 3 questions of 10 marks each. The student has to attempt any 2 questions.

Section – A

(4 marks each)

- Q1. What is the Gini index, and how is it used to construct decision trees?
- Q2. Explain the difference between K-means and K-means++ initialization. Why is K-means++ preferred?
- Q3. What are autoencoders, and how can they be used for unsupervised feature learning?
- Q4. What is reinforcement learning, and how is it applied to optimize IoT networks?

Section – B

(8 marks each)

- Q5. What are Generalized Linear Models (GLMs), and how do they extend linear regression?
- Q6. How does Kernel PCA differ from PCA, and in what scenarios would Kernel PCA be preferred over PCA?
- Q7. What is semi-supervised learning, and how does it utilize both labeled and unlabeled data to improve learning performance?
- Q8. What are the different types of machine learning models that can be used in IoT applications, and how do they address the challenges of IoT environments?

Section – C

(10 marks each)

- Q9. Implement a multi-class SVM using the one-vs-all (OvA) strategy, where a separate binary classifier is trained for each class. Test it on a multi-class dataset like the Iris dataset.
- Q10. Describe how the AdaBoost algorithm works and how it reduces model error. Create the AdaBoost algorithm where each weak learner is a decision stump (a decision tree with a depth of 1) as the base classifier.
- Q11. Discuss the trade-offs involved in choosing between online gradient descent and stochastic gradient descent for model optimization. Implement online gradient descent in Python. Given a dataset, apply the algorithm to train a simple linear regression model and evaluate its performance using a test set.