

--	--	--	--	--	--	--	--	--	--	--	--	--	--	--

Question Paper Code : 91935

B.E./B.Tech. DEGREE EXAMINATIONS, APRIL/MAY 2025.

Sixth Semester

Computer Science and Engineering

CS 3491 — ARTIFICIAL INTELLIGENCE AND MACHINE LEARNING

(Common to : Biomedical Engineering/Computer Science and Design/Computer Science and Engineering (Cyber Security)/Computer and Communication Engineering/Electronics and Communication Engineering/Electronics and Telecommunication Engineering/Medical Electronics/Information Technology)

(Regulations 2021)

Time : Three hours

Maximum : 100 marks

Answer ALL questions.

PART A — (10 × 2 = 20 marks)

1. What is constraint satisfaction problem?
2. Define rational agent.
3. How Bayesian network works?
4. State the principle of acting under uncertainty.
5. Differentiate between soft margin and hard margin.
6. Mention the purpose of discriminant function.
7. How ensemble learning improves the performance of machine learning models?
8. What are the merit and demerits of instance based learning?
9. In a neural network, what is the purpose of the back propagation algorithm, and how does it use the gradient of the error to update the weights during training?
10. Why are activation functions important in neural networks, and how do they impact the model's ability to learn complex patterns?

PART B — (5 × 13 = 65 marks)

11. (a) Explain in detail the concept of heuristic search strategies. How do they differ from uninformed search strategies? Discuss various heuristic search techniques such as A* Search With their working principles, advantages, and limitations. Illustrate your explanation with suitable examples.

Or

- (b) Discuss the role of local search techniques and adversarial search strategies in solving optimization problems. Highlight the similarities and differences between local and adversarial search in terms of objective, strategy and application areas. Support your discussion with appropriate examples.
12. (a) Weather forecasting is a critical task in agriculture, disaster management, and daily planning. Meteorologists rely on historical weather data to make predictions about future conditions. Suppose a weather station wants to automate the prediction of whether it will rain on a particular day based on previous patterns of weather features such as humidity, temperature, wind speed and cloud cover. Using the data given below, apply the Naive Bayes classification model to predict whether it will rain or not using {Humidity = "Normal" Temperature = "mild", Wind = "weak", Sky = "Overcast"}, given its features. [Table: Humidity - Temperature - Wind - Sky - Rain (Class label): High - Hot - Strong - Clear - Yes; Normal - Mild - Weak - Cloudy - No; High - Cool - Strong - Overcast - Yes; Normal - Hot - Weak - Clear - No; High - Mild - Weak - Cloudy - Yes; Normal - Cool - Strong - Overcast - No; High - Hot - Weak - Clear - Yes; Normal - Mild - Strong - Overcast - No; High - Cool - Weak - Cloudy - Yes; Normal - Hot - Strong - Clear - No]
- (i) What evaluation metrics would you use to assess the performance of your model? Justify your choice. (9)
- (ii) What are the limitations of using Naive Bayes in weather forecasting, and how can the model be improved? (4)

Or

- (b) Explain the concept of probabilistic reasoning in artificial intelligence. How does it support decision-making under uncertainty? Provide relevant examples where probabilistic reasoning is essential, and explain how inference is performed in such models.
13. (a) Explain how a probabilistic discriminative model differs from a generative model in the context of this spam detection task. In your perspective which model is more advantageous and why?

Or

- (b) Explain the working principle of Random Forests. Describe the role of bootstrapping and feature randomness in constructing Random Forests. Also, discuss its advantages, limitations, and typical real-world applications where Random Forests perform well.
14. (a) All shopping malls have food courts which might include dishes like those given below. Use k-means for clustering into three clusters by considering the Manhattan distance metric. Suppose that the initial seeds are F4, F7, and F9. [Table: Food Id - Item - Calories - Carbohydrates - Proteins: F1 - Aloo Curry - 105 - 15 - 1; F2 - Cabbage - 131 - 7 - 2; F3 - Pumpkin - 67 - 7 - 2; F4 - Mutter Paneer - 147 - 11 - 9; F5 - Chana Dhal - 99 - 13 - 5; F6 - Moong Dhal - 211 - 31 - 13; F7 - Besan Khadi - 100 - 15 - 3; F8 - Kofta Curry - 147 - 13 - 3; F9 - Milk - 67 - 4 - 3; F10 - Cheese - 348 - 6 - 24] Apply k-means for 2 iterations. Trace the following intermediate results, at the end of each iteration. Find the new clusters formed and Calculate the Centroid of the new clusters.

Or

- (b) What is a Gaussian Mixture Model (GMM)? Explain how GMMs are used for clustering and density estimation. Describe in detail the role of the Expectation-Maximization (EM) algorithm in fitting a GMM.
15. (a) Discuss in detail about perceptron and multilayer perceptron with suitable illustrations.

Or

- (b) Explain the role of the ReLU (Rectified Linear Unit) activation function in deep neural networks. Further, elaborate on the importance of hyper parameter tuning in training deep learning models. What are the key hyper parameters that affect model performance, and what techniques can be used to optimize them? Provide examples to illustrate how improper tuning may lead to underfitting or overfitting.

PART C — (1 × 15 = 15 marks)

16. (a) Consider the database about customer's purchase of all electronics. The attributes are age, income, student category, rating and the class label Y or N (Y-Yes, buys computer, N-No, Don't buys computer). The training data set is given below: [Table: Age - Income - Student - Rating - Class Label: Youth - High - No - Fair - No; Youth - High - No - Excellent - No; Middle_aged - High - No - Fair - Yes; Senior - Medium - No - Fair - Yes; Senior - Low - Yes - Fair - Yes; Senior - Low - Yes - Excellent - No; Middle-aged - Low - Yes - Excellent - Yes; Youth - Medium - No - Fair - No; Youth - Low - Yes - Fair - Yes; Senior - Medium - Yes - Fair - Yes; Youth - Medium - Yes - Excellent - Yes; Middle_aged - Medium - No - Excellent - Yes; Middle_aged - High - Yes - Fair - Yes; Senior - Medium - No - Excellent - No] Determine the class label using decision tree induction method and describe about decision tree.

Or

- (b) A telecom company wants to predict whether a customer will churn (leave the service) in the next 6 months based on their usage data. The company has a rich dataset that includes customer demographics, service usage, contract type, and payment history. The objective is to build a predictive model to classify customers as either 'Churn' or 'No Churn' using the K-Nearest Neighbors (KNN) algorithm. Discuss about the need for K-Nearest Neighbors (KNN). [Table: Customer ID - Age - Tenure (Months) - Service Type - Monthly Spend (\$) - Payment Method - Contract Type - Customer Support Calls - Churn: 1 - 45 - 12 - Premium - 85 - Credit Card - Two-Year - 2 - 0; 2 - 29 - 3 - Standard - 50 - Bank Transfer - Month-to-Month - 8 - 1; 3 - 63 - 30 - Standard - 60 - Direct Debit - One-Year - 1 - 0; 4 - 40 - 8 - Premium - 90 - Credit Card - Month-to-Month - 5 - 1; 5 - 52 - 24 - Premium - 120 - Credit Card - Two-Year - 3 - 0; 6 - 35 - 18 - Standard - 70 - Credit Card - One-Year - 4 - 1; 7 - 50 - 60 - Premium - 100 - Bank Transfer - Two-Year - 0 - 0; 8 - 27 - 1 - Standard - 40 - Direct Debit - Month-to-Month - 10 - 1; 9 - 55 - 36 - Premium - 95 - Credit Card - One-Year - 2 - 0; 10 - 39 - 15 - Standard - 75 - Bank Transfer - One-Year - 6 - 0]