


USN



21AI71

Seventh Semester B.E./B.Tech. Degree Examination, June/July 2025 Advanced AI and ML

Time: 3 hrs.

Max. Marks: 100

Note: Answer any FIVE full questions, choosing ONE full question from each module.

Module-1

- 1 a. Explain the task environment and properties of task environments. (10 Marks)
- b. Describe and compare the four basic types of agent programs. (10 Marks)

OR

- 2 a. Explain the concept of adversarial search in game playing agents. How does minimax algorithm help an agent make decisions in a competitive environment? (10 Marks)
- b. Explain the alpha beta search algorithm also write the effectiveness of alpha beta pricing. (10 Marks)

Module-2

- 3 a. Describe the quantification problem in the context of logical agents acting under uncertainty. (10 Marks)
- b. How does an agent uses degrees of belief to handle uncertainty in decision making? (10 Marks)

OR

- 4 a. Describe the naïve Bayes model and its assumptions about conditional independence. (10 Marks)
- b. What is Bayes 'theorem' and how it is applied in combining multiple source of evidence? (10 Marks)

Module-3

- 5 a. Explain the concept of a Perceptron with a neat diagram. Discuss the Perceptron training rule. (10 Marks)
- b. Explain the derivation of the Back propagation rule. (10 Marks)

OR

- 6 a. Explain the prototypical of a genetic algorithm. (10 Marks)
- b. Explain the Hypothesis space search in the genetic algorithm. (10 Marks)

Module-4

- 7 a. Explain the association rule mining with an example. (10 Marks)
- b. Explain the application of collaborative filtering in building recommendation systems. Illustrate with python code Implementation. (10 Marks)

Important Note : 1. On completing your answers, compulsorily draw diagonal cross lines on the remaining blank pages.
2. Any revealing of identification, appeal to evaluator and /or equations written eg, 42+8 = 50, will be treated as malpractice.

OR

- 8 a. Explain the Naïve Baye's Model for sentiment classification in Python? Illustrate with code implementation. (10 Marks)
- b. What is the TF-IDF vectorizer, and how is it used in text analytics? Discuss its role in improving the performance of text classification tasks. What challenges are associated with text analytics? (10 Marks)

Module-5

- 9 a. Explain the following clustering technique with an algorithm.
i) K-means method ii) K-medoids method. (10 Marks)
- b. Explain how hierarchical clustering works, including the differences between building clustering from the bottom up (agglomerative) and top down (divisive). (10 Marks)

OR

- 10 a. What is instance – based learning? Explain K-nearest Neighbour learning algorithm. (10 Marks)
- b. Explain the following :
i) Locally weighted Linear Regression
ii) Radial basis functions. (10 Marks)

* * * * *