

KATHMANDU UNIVERSITY
End Semester Examination
March/April, 2025

Marks Scored:

Level : B.Tech.

Year : II

Exam Roll No. :

Time: 30 mins.

Course : AICC 201

Semester : I

F. M. : 10

Registration No.:

Date

1 APR 2025

SECTION "A"

[20Q. × 0.5 = 10 marks]

Choose and encircle the most appropriate answer.

1. What is the primary goal of Artificial Intelligence (AI)?
 - a. To replace human intelligence
 - b. To build machines that mimic human behavior and decision-making
 - c. To create systems that require no human intervention
 - d. To eliminate the need for software development
2. What does it mean for an AI system to be rational?
 - a. It always makes the correct decision
 - b. It acts in a way that maximizes expected performance based on available information
 - c. It mimics human emotions and feelings
 - d. It learns without any form of supervision
3. Which of the following algorithms is more memory-intensive?
 - a. Breadth-First Search (BFS)
 - b. Depth-First Search (DFS)
 - c. Hill Climbing
 - d. Simulated Annealing
4. Which of the following is an example of an uninformed search algorithm?
 - a. A* Search
 - b. Best-First Search
 - c. Breadth-First Search
 - d. Greedy Search
5. What is a key drawback of the hill climbing algorithm?
 - a. It requires a large amount of memory
 - b. It gets stuck in local optima
 - c. It ignores the cost function
 - d. It is slower than BFS
6. Which of the following is a valid proposition in propositional logic?
 - a. $\exists \mathbf{x} (P(\mathbf{x}) \rightarrow Q(\mathbf{x}))$
 - b. $\forall \mathbf{x} (P(\mathbf{x}) \wedge Q(\mathbf{y}))$
 - c. $\exists \mathbf{y} (R(\mathbf{y}) \rightarrow P(\mathbf{y}))$
 - d. $P \wedge (Q \vee \neg R)$

7. In predicate logic, which of the following quantifiers is used to represent "there exists at least one instance"?
- \forall (Universal Quantifier)
 - \exists (Existential Quantifier)
 - \Rightarrow (Implication)
 - \Leftrightarrow (Biconditional)
8. Which inference rule allows removing an existential quantifier by introducing a new constant?
- Existential Elimination
 - Universal Instantiation
 - Modus Ponens
 - Resolution
9. Given two expressions $P(x, y)$ and $P(a, z)$, what is the correct unifier?
- $\{x \rightarrow a, y \rightarrow z\}$
 - $\{x \rightarrow z, y \rightarrow a\}$
 - $\{x \rightarrow y, z \rightarrow a\}$
 - No unification possible
10. In probability theory, what does $P(A | B)$ represent?
- The probability of event A occurring independently
 - The probability of event A given that B has occurred
 - The probability of both A and B occurring together
 - The probability of event A given that B has occurred
11. Which of the following statements accurately defines the bias-variance trade-off in Machine Learning?
- A model with high bias and low variance has high generalization ability.
 - Increasing the complexity of a model will reduce both bias and variance simultaneously.
 - A model with low bias and high variance is likely to overfit the data.
 - Bias and variance are independent of each other in a learning problem.
12. Which normalization technique is most appropriate for a dataset where the features have different units, such as age in years and income in dollars?
- Min-Max scaling
 - Z-score normalization (standardization)
 - Log transformation
 - Data augmentation
13. Which of the following methods is suitable for encoding ordinal categorical data?
- One-Hot Encoding
 - Label Encoding
 - Binary Encoding
 - Feature Hashing

14. Which of the following is an example of interval data?
 - a. Rating scale from 1 to 5
 - b. Height of an individual
 - c. Type of fruit
 - d. Temperature in Celsius

15. Which of the following correctly represents a basic feedforward neural network structure?
 - a. The input layer is connected to the output layer only through a single hidden layer.
 - b. The input layer is not connected to the hidden layer but only to the output layer.
 - c. The number of hidden layers must always be greater than 2 for a network to function properly.
 - d. Each node in a layer is connected to all nodes in the subsequent layer.

16. What does the cost function of a model typically represent?
 - a. The sum of individual errors for each data point.
 - b. The total accumulated error across all data points in the training set.
 - c. The error for a single prediction.
 - d. The number of incorrect predictions made by the model.

17. In binary classification, what does the cross-entropy function measure?
 - a. The average of the squared differences between predicted and actual values.
 - b. The difference between predicted probabilities and the actual labels (0 or 1).
 - c. The number of misclassified samples.
 - d. The variance of the predictions.

18. In regression, what would be the effect of using Mean Absolute Error (MAE) instead of Mean Squared Error (MSE)?
 - a. MAE will be more sensitive to outliers compared to MSE.
 - b. MAE works better for classification problems than MSE.
 - c. MSE will penalize large errors more than MAE.
 - d. MSE does not account for the magnitude of errors, whereas MAE does.

19. In the context of hyperparameter tuning, grid search is best suited when:
 - a. The search space is small and you need to test every possible combination.
 - b. The search space is large and you want to find approximate optimal values quickly.
 - c. You do not know which hyperparameters influence model performance.
 - d. The model parameters do not significantly affect the model accuracy.

20. In gradient descent, if the learning rate is too small, which of the following is most likely to occur?
 - a. The algorithm will converge quickly but to a suboptimal solution.
 - b. The algorithm will converge too slowly and may get stuck in a local minimum.
 - c. The algorithm will fail to update the parameters properly.
 - d. The cost function will diverge, leading to poor model performance.

KATHMANDU UNIVERSITY
End Semester Examination
March/April, 2025

Level : B.Tech.
Year : II
Time : 2 hrs. 30 mins.

01 APR 2025

Course : AICC 201
Semester : I
F. M. : 40

SECTION "B"

[6Q. × 4 = 24 marks]

Attempt *ANY SIX* questions.

1. What is AI? Explain the four different approaches of AI.
2. Explain what you understand by the following terms: States, Initial States, Actions, Transition model, Goal test, and Path cost.
3. Differentiate between uninformed and informed search. Explain Greedy best-first search.
4. Differentiate between propositional and predicate logic. List and explain the Rules of Inference.
5. How do data quality and quantity impact machine learning models? What are common techniques used for normalizing numerical data? Explain with the required formula.
6. What is the Bayesian Network? Explain how the dependencies among variables are represented in Bayesian networks with the help of a proper diagram.
7. How is the cost function used to evaluate model performance? Explain the difference between Mean Squared Error (MSE) and Cross-Entropy Loss.

SECTION "C"

[2Q. × 8 = 16 marks]

Attempt *ANY TWO* questions.

8. What is the gradient descent algorithm, and how does it work? Compare Batch Gradient Descent, Stochastic Gradient Descent (SGD), and Mini-batch Gradient Descent.
9. What do you understand about the local search algorithm? Explain about Hill-climbing search.
10. Write short note on: [4Q × 2=8]
 - a. Minimax algorithm
 - b. Neural Network
 - c. Hyper-parameter Tuning
 - d. Quantifiers

