

Wks Mapped with the course

WK1	WK2	WK3	WK4	WK5	WK6	WK7	WK8	WK9
*	*	*	*	*	*	*	*	

Mark Distribution:

Total Marks	CIE Marks	ESE Marks	ESE Duration
100	40	60	3 hrs

Continuous Internal Evaluation Pattern:

Attendance	: 5 marks
Continuous Assessment Test	: 20 marks
Continuous Assessment Assignment	: 15 marks

SYLLABUS

Module 1 (12 hrs)

Introduction to Artificial Intelligence: Definition, history, and evolution of AI; Turing Test, rational agents, and intelligent systems; AI problem formulation: state space, goal state, operators; Uninformed search strategies: BFS, DFS, Depth-Limited Search, Iterative Deepening DFS; Informed search: Greedy Best-First Search, A* algorithm; Heuristics and admissibility; Introduction to local search: Hill Climbing, Simulated Annealing.

Module 2 (13 hrs)

Constraint Satisfaction, Logic and Reasoning: Constraint Satisfaction Problems (CSP): backtracking, arc consistency, forward checking; Adversarial search: Minimax algorithm, Alpha-Beta pruning; Game playing and game trees; Propositional logic: syntax, semantics, inference rules; First-Order Logic (FOL): predicates, quantifiers, unification; Forward and backward chaining; Introduction to probabilistic reasoning: Bayesian networks, conditional independence, Bayes theorem.

Module 3 (13 hrs)

Machine Learning Foundations: Overview of machine learning: supervised, unsupervised and reinforcement learning; Linear regression and logistic regression; Decision Trees and Random Forests; Support Vector Machines (SVM); Naive Bayes classifier; k-Nearest Neighbours (k-NN); Clustering: k-Means, hierarchical clustering; Dimensionality reduction: PCA; Bias-variance tradeoff, overfitting and underfitting; Model evaluation: cross-validation, confusion matrix, precision, recall, F1-score, ROC curve; Introduction to Reinforcement Learning: Q-learning, Markov Decision Processes.

Module 4 (15 hrs)

Neural Networks, Deep Learning and Applications: Biological and artificial neurons; Perceptron and Multi-Layer Perceptron (MLP); Activation functions: sigmoid, ReLU, tanh; Backpropagation and

gradient descent; Convolutional Neural Networks (CNN): architecture and applications in image recognition; Recurrent Neural Networks (RNN) and LSTM for sequential data; Introduction to Natural Language Processing (NLP): tokenisation, POS tagging, named entity recognition, word embeddings (Word2Vec, GloVe); Transformer architecture and BERT overview; AI Ethics: bias, fairness, explainability, privacy and societal impact; Current trends: Generative AI, Large Language Models (LLMs), and responsible AI.

Completion of an industrial visit or industrial training is desirable for this course.

Text Books:

1. Russell, S. J., & Norvig, P. (2020). Artificial intelligence: A modern approach (4th ed.). Pearson.
2. Rich, E., Knight, K., & Nair, S. B. (2009). Artificial intelligence (3rd ed.). Tata McGraw-Hill.

Reference Books:

1. Winston, P. H. (1992). Artificial intelligence (3rd ed.). Addison-Wesley.
2. Mitchell, T. M. (1997). Machine learning. McGraw-Hill.
3. Goodfellow, I., Bengio, Y., & Courville, A. (2016). Deep learning. MIT Press.
4. Géron, A. (2022). Hands-on machine learning with Scikit-Learn, Keras, and TensorFlow (3rd ed.). O'Reilly Media.

END SEMESTER EXAMINATION – MODEL QUESTION PAPER

Course Code	25UAMT403
Course Name	INTRODUCTION TO ARTIFICIAL INTELLIGENCE
Max. Marks: 40	Duration: 2.5 Hours

Part A (Answer all questions. Each question carries 2marks)				
Qn No.	Question	Marks	CO	BL
1.	Explain the concept of rational agents in AI. How does the PEAS framework help describe an agent's task environment?	2	CO1	L2
2.	Explain PEAS details of automated car.	2	CO2	L2
3.	Explain arc consistency in CSP. Apply AC-3 algorithm to the map-coloring problem of Australia.	2	CO2	L3
4.	Explain Minimax algorithm with Alpha-Beta pruning. How does it improve performance over plain Minimax.	2	CO2	L2
5.	Explain Bayesian networks with an example. How is conditional independence represented in a Bayesian network.	2	CO3	L2
6.	Distinguish supervised, unsupervised, and reinforcement learning with examples of real-world applications.	2	CO3	L2
7.	Explain backpropagation in neural networks. How does gradient	2	CO4	L2

	descent update the weights?			
8.	Describe ethical challenges in deploying AI systems. What is algorithmic bias and how can it be mitigated?	2	CO4	L2

Part B (Answer any one full question from each Module. Each carry 6 marks)				
Module I				
Qn No.	Question	Marks	CO	BL
9.1	Explain BFS, DFS, and Iterative Deepening DFS for search problems. Compare their time and space complexities.	3	CO1	L2
9.2	Apply A* algorithm to solve the 8-puzzle problem. Show the open and closed lists at each step.	3	CO2	L3
OR				
10.1	Explain the Hill Climbing and Simulated Annealing local search algorithms. What are the limitations of Hill Climbing and how does Simulated Annealing overcome them?	3	CO2	L2
10.2	What is the Turing Test? Check whether it is a sufficient measure of machine intelligence.	3	CO1	L3
Module II				
Qn No.	Question	Marks	CO	BL
11	Explain First-Order Logic with examples of predicates, quantifiers, and inference rules. Apply forward chaining and backward chaining to derive conclusions from the given knowledge base: (i) All humans are mortal. (ii) Socrates is human. Conclude: Is Socrates mortal?	3	CO3	L3
OR				
12.1	Explain the Minimax algorithm with Alpha-Beta pruning using a game tree of depth 3. Compute the optimal value at the root.	3	CO2	L3
12.2	Illustrate constraint satisfaction problems with backtracking search. Apply it to solve a 4-Queens problem.	3	CO2	L3
Module III				
Qn No.	Question	Marks	CO	BL
13	Explain the working of a Decision Tree classifier with the ID3 algorithm. Apply it to classify the following dataset using Information Gain. [Provide sample dataset with 4 attributes and 10 records in exam]. Discuss overfitting and pruning strategies.	6	CO3	L3

OR				
14.1	Describe Support Vector Machines (SVM) with the concept of maximum margin hyperplane and kernel trick.	3	CO3	L2
14.2	Explain the k-Means clustering algorithm with an example. What are its limitations?	3	CO3	L2
Module IV				
Qn No.	Question	Marks	CO	BL
15.1	Describe the implementation of Convolutional Neural Networks (CNN) architecture. How are CNNs applied to image recognition tasks?	3	CO4	L2
15.2	Illustrate Transformer architecture and its significance in NLP. How does BERT leverage transformers for language understanding?	3	CO4	L3
OR				
16	Discuss AI ethics in detail: bias, fairness, explainability, and privacy. Describe current regulatory frameworks for responsible AI deployment. How can Large Language Models (LLMs) be made more transparent and accountable?	6	CO4	L3
