ARTIFICIAL INTELLIGENCE

ENCT 305

Lecture : 3 **Year : III**
Tutorial : 1 **Part : I**
Practical : 3/2

Course Objectives:
The objective of this course is to build a strong foundation in Artificial Intelligence (AI), covering intelligent agents, search techniques, knowledge representation, machine and learning. It aims to equip students with both theoretical knowledge and practical skills to apply AI techniques in solving real-world problems, while also fostering awareness of the ethical considerations associated with AI applications.

1 Introduction (4 hours)

1.1 Definition, foundation, history of AI
1.2 AI tree: Branches and interdisciplinary nature
1.3 Definition and importance of knowledge and learning
1.4 Human intelligence and machine intelligence
1.5 Intelligent agents and types

2 Problem Solving and Search (9 hours)

2.1 Formal problem definition: States, actions, transitions, well-defined problems
2.2 Constraint satisfaction problems: Node consistency, path consistency, backtracking
2.3 Search algorithms, strategies, and evaluations
2.4 Uninformed: BFS, DFS, iterative deepening
2.5 Informed search: Best first search, greedy search, A* algorithm
2.6 Adversarial search: Minimax algorithm, alpha-beta pruning
2.7 Local search and optimization: Hill climbing, simulated annealing
2.8 Evolutionary optimization: Genetic algorithm

3 Knowledge Representation and Probabilistic Reasoning (7 hours)

3.1 Knowledge-based agent
3.2 Knowledge representation techniques and issues in representation
3.3 Propositional and predicate logic
3.4 Semantic networks and frames
3.5 Probabilistic reasoning, inference using Bayes' theorem
3.6 Fuzzy logic: Membership functions, fuzzy inference systems

4 Machine Learning Fundamentals (10 hours)

4.1 Forms of learning: Supervised, unsupervised, semi-supervised, reinforcement learning
4.2 Machine learning pipeline
  4.2.1 Preprocessing and cleaning
  4.2.2 Model development
  4.2.3 Training, testing, and hyperparameter tuning
4.3 Overfitting, underfitting, bias-variance tradeoff
4.4 Regression and classification with linear model
4.5 Classification using decision tree and Naive Bayes
4.6 Evaluation with confusion matrix

5 Neural Networks and Deep Learning Algorithms (6 hours)

5.1 Neural networks: Structures and activation functions
5.2 Perceptron, multilayer perceptron, and backpropagation
5.3 Introduction to deep learning
5.4 Concepts on recurrent and generative neural networks

6 AI Applications (7 hours)

6.1 Expert systems: Characteristics, architecture, development and various applications
6.2 NLP: Level of analysis and application
6.3 Robotics and computer vision: Fundamental, components and applications
6.4 VLSI testing, fault diagnosis, and sustainable systems

7 AI Ethics and Recent Trends (2 hours)

7.1 Responsible AI: Bias, fairness, explainability
7.2 Ethical and societal considerations
7.3 Privacy and security in AI
7.4 Emerging trends: Federated learning, embedded AI

Tutorial (15 hours)

1. Design a simple agent model (Vacuum cleaner, chess player)
2. Formulate a route-finding problem using states, actions, and transitions
3. Solve the problem with uninformed strategies
4. Apply the A* algorithm with the given heuristic values
5. Solve the minimax tree and demonstrate alpha–beta pruning
6. Solve a CSP like map coloring or mini-sudoku
7. Perform one iteration of the genetic algorithm (Selection, crossover, mutation)
8. Represent facts using semantic networks and frames
9. Solving the problem related to linear and logistic regression
10. Solving problems related to Bayes' theorem and fuzzy logic

11. Understanding machine learning basics through model training, cross-validation, and performance evaluation
12. Apply the Naive Bayes theorem to classify an instance
13. Model evaluation: Build a confusion matrix. Calculate sensitivity, specificity, precision, recall, and F1-score
14. Solve perceptron weight updates for a small dataset
15. Derive forward and backward propagation numerically for a 2-layer NN
16. Case study: Identify ethical risks in real-world AI use. Responsible AI in Nepal's context

Practical **(22.5 hours)**

1. Knowledge-based agents and search
2. Adversarial search and CSP
3. Symbolic and probabilistic reasoning
4. Machine learning – Supervised and unsupervised
5. Neural networks basics
6. Mini project and AI applications

Final Exam

The questions will cover all the chapters in the syllabus. The evaluation scheme will be as indicated in the table below:

Chapters	Hours	Marks distribution*
1,6 and 7	13	16
2	9	12
3	7	10
4	10	12
5	6	10
Total	**45**	**60**

* There may be minor deviation in marks distribution.

References

1. Russell, S., Norvig, P. (2020). Artificial intelligence: A modern approach. Pearson.
2. Rich, E., Knight, K., Nair, S. B. (2009). Artificial intelligence. McGraw-Hill.
3. Bishop, C. M. (2006). Pattern recognition and machine learning. Springer.
4. Deisenroth, M. P., Faisal, A. A., Ong, C. S. (2020). Mathematics for machine learning. Cambridge University Press.