

ARTIFICIAL INTELLIGENCE

ENCT 351

Lecture : 3
Tutorial : 1
Practical : 3

Year : III
Part : II

Course Objectives:

The objective of this course is to provide students with a foundation in Artificial Intelligence (AI), covering intelligent agents, search techniques, knowledge representation, machine learning, and AI ethics. It aims to equip students with both theoretical understanding and practical skills to apply AI techniques to real-world problems, while also developing awareness of the ethical and societal implications of AI systems.

1 Introduction (4 hours)

- 1.1 Definition, foundation, history of AI
- 1.2 Importance of knowledge and learning
- 1.3 Cognition and learning (Neuroscience)
- 1.4 Human intelligence and machine intelligence
- 1.5 AI tree: Branches and interdisciplinary nature
- 1.6 Intelligent agents and types
- 1.7 Agentic AI

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, frames and knowledge graph
- 3.5 Review of Bayes' Theorem and probabilistic reasoning

3.6 Fuzzy logic: Membership functions, fuzzy inference systems

4 Machine Learning Fundamentals (10 hours)

- 4.1 Foundations and four pillars of machine learning
- 4.2 Review of mathematics for machine learning
- 4.3 Continuation optimization; Unconstraint optimization, constraint optimization, convex optimization
- 4.4 Review of machine learning pipeline
 - 4.4.1 Model development: Generative versus discriminative
 - 4.4.2 Learning algorithm
 - 4.4.3 Capacity, overfitting and underfitting
 - 4.4.4 Hyperparameters and validation sets
 - 4.4.5 Estimators, bias and variance
 - 4.4.6 Review of MLE and MAP and cross entropy
- 4.5 Supervised learning algorithm: Decision tree and support vector machine
- 4.6 Unsupervised learning algorithm: t-SNE
- 4.7 Semi supervised and reinforcement learning
- 4.8 Model evaluation: Energy-based indicators

5 Neural Networks and Deep Learning Algorithms (8 hours)

- 5.1 Neural networks: Structures, activation functions and universal approximation theorem
- 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 (5 hours)

- 6.1 Expert systems: Characteristics, architecture, development and various applications
- 6.2 NLP: Level of analysis, challenges, modern approaches and applications
- 6.3 Robotics and computer vision: Fundamental, components and applications
- 6.4 Sustainable AI systems

7 Emerging Trends (2 hours)

- 7.1 Sequence to sequence models
- 7.2 Federated learning
- 7.3 Edge AI
- 7.4 Ethics and AI: Responsible 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 problems related to Bayes' theorem and Fuzzy Logic
10. Understanding machine learning basics through model training, cross-validation and performance evaluation
11. Calculation related to decision tree and support vector machine
12. Model Evaluation: Build a confusion matrix; Calculate sensitivity, specificity, precision, recall and f1-score
13. Solve perceptron weight updates for a small dataset
14. Derive forward and backward propagation numerically for a 2-layer NN
15. Case study: Identify ethical risks in real-world AI use-responsible AI in Nepal's context

Practical

(45 hours)

1. Intelligent agents and problem formulation
2. Uninformed search techniques
3. Informed (Heuristic) search techniques
4. Adversarial search and game playing
5. Constraint satisfaction problems (CSP)
6. Evolutionary computation (Genetic algorithms)
7. Knowledge representation (Logic, semantic networks, frames)
8. Probabilistic and fuzzy reasoning
9. Machine learning pipeline and data preprocessing
10. Supervised and unsupervised learning
11. Neural networks and deep learning basics
12. Mini project, AI applications and ethics

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	11	14
2	9	12
3	7	10
4	10	14
5	8	10
Total	45	60

* There may be minor deviation in marks distribution.

References

1. Russell, S., Norvig, P. (2021). Artificial intelligence: A modern approach. Pearson.
2. Rich, E., Knight, K., Nair, S. B. (2009). Artificial intelligence. McGraw-Hill Education.
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.
5. Goodfellow, I., Bengio, Y., Courville, A. (2016). Deep learning. MIT Press.