PVP 23

Code: 23ME3602

III B.Tech - II Semester - Regular Examinations – APRIL 2026

ARTIFICIAL INTELLIGENCE AND MACHINE LEARNING (MECHANICAL ENGINEERING)

Duration: 3 hours Max. Marks: 70

Note: 1. This question paper contains two Parts A and B.
2. Part-A contains 10 short answer questions. Each Question carries 2 Marks.
3. Part-B contains 5 essay questions with an internal choice from each unit. Each Question carries 10 marks.
4. All parts of Question paper must be answered in one place.

BL – Blooms Level CO – Course Outcome

PART – A

		BL	CO
1.a)	What is an Intelligent Agent?	L1	CO1
b)	What is First Order Logic (FOL)?	L1	CO1
c)	List the difference between supervised and unsupervised learning.	L1	CO2
d)	What is underfitting?	L1	CO2
e)	What is dimensionality reduction?	L1	CO3
f)	Define clustering.	L1	CO3
g)	What are genetic operators?	L1	CO4
h)	What is an Artificial Neural Network?	L1	CO4
i)	What is boosting?	L1	CO5
j)	What is deep learning?	L1	CO5

PART – B

			BL	CO	Max. Marks
UNIT-I					
2	a)	Explain the characteristics of intelligent agents with examples.	L2	CO1	5 M
	b)	Distinguish between propositional logic and first-order logic.	L2	CO1	5 M
OR					
3	a)	Explain the concept of rational agents and rationality.	L2	CO1	5 M
	b)	Explain the concept of Wumpus world logic.	L2	CO1	5 M
UNIT-II					
4	a)	Explain Support Vector Machine (SVM) and its advantages.	L2	CO2	5 M
	b)	Compare regression and classification problems with examples.	L2	CO2	5 M
OR					
5	a)	Illustrate Decision Trees and their working principle.	L3	CO2	7 M
	b)	Explain supervised learning with suitable examples.	L2	CO2	3 M
UNIT-III					
6	a)	Describe the steps of the K-means clustering algorithm.	L2	CO3	5 M

	b)	Explain dimensionality reduction and its importance.	L2	CO3	5 M
OR					
7	Illustrate the working principle of Principal Component Analysis.		L3	CO3	10 M
UNIT-IV					
8	a)	Explain the working principle of a perceptron model.	L2	CO4	5 M
	b)	Discuss genetic operators such as selection, crossover and mutation.	L2	CO4	5 M
OR					
9	Illustrate the Convolutional Neural Networks and their applications.		L3	CO4	10 M
UNIT-V					
10	a)	Explain the working of Random Forest algorithm.	L2	CO5	5 M
	b)	Summarize the working of deep auto-encoders.	L2	CO5	5 M
OR					
11	a)	Describe the applications of deep neural networks in real-world problems.	L2	CO5	5 M
	b)	Explain ensemble learning methods in machine learning.	L2	CO5	5 M

PVP 23

Code: 23ME3602

III B.Tech - II Semester - Regular Examinations – APRIL 2026

ARTIFICIAL INTELLIGENCE AND MACHINE LEARNING (MECHANICAL ENGINEERING)

Duration: 3 hours Max. Marks: 70

Note: 1. This question paper contains two Parts A and B.
2. Part-A contains 10 short answer questions. Each Question carries 2 Marks.
3. Part-B contains 5 essay questions with an internal choice from each unit. Each Question carries 10 marks.
4. All parts of Question paper must be answered in one place.

BL – Blooms Level CO – Course Outcome

Scheme of Evaluation

PART–A

1.a. What is an Intelligent Agent?**BL:** L1 **CO:1** **Marks:** 2
- Definition of Intelligent Agent – 1 Mark
- Example (robot, chatbot, self-driving car, etc.) – 1 Mark

1.b. What is First Order Logic (FOL)?**BL:** L1 **CO:**1 **Marks:** 2
- Definition of FOL – 1 Mark
- Mention predicates, variables, quantifiers – 1 Mark

1.c. List the difference between supervised and unsupervised learning?**BL:** L1 CO2 **Marks:** 2
- Supervised learning definition – 1 Mark
- Unsupervised learning definition – 1 Mark

1.d. What is underfitting?**BL:** L1 CO2 **Marks:** 2
- Definition – 1 Mark
- Cause or example – 1 Mark

1.e. What is dimensionality reduction?**BL:** L1 CO3 **Marks:** 2
- Definition – 1 Mark
- Purpose/benefit – 1 Mark

1.f. Define clustering?**BL:** L1 CO3 **Marks:** 2
- Definition – 1 Mark
- Example or purpose – 1 Mark

1.g. What are genetic operators?**BL:** L1 CO4 **Marks:** 2
- Definition – 1 Mark
- Types (selection, crossover, mutation) – 1 Mark

1.h. What is an Artificial Neural Network?**BL:** L1 CO4 **Marks:** 2
- Definition – 1 Mark
- Structure mention – 1 Mark

1.i. What is boosting?**BL:** L1 CO5 **Marks:** 2

- Definition – 1 Mark
- Example (AdaBoost, XGBoost) – 1 Mark

1.j. What is deep learning?**BL:** L1 CO5 **Marks:** 2

- Definition – 1 Mark
- Use of multiple hidden layers – 1 Mark

PART–B

UNIT–I

2.a Explain the characteristics of intelligent agents with examples? L2 CO-1 **Marks:** 5

- Definition – 1 Mark
- Characteristics (autonomy, reactivity, proactiveness, learning) – 3 Marks
- Example – 1 Mark

2.b Distinguish between propositional logic and first-order logic L2 CO-1 **Marks:** 5

- Propositional Logic explanation – 2 Marks
- First Order Logic explanation – 2 Marks
- Comparison table/examples – 1 Mark

OR

3.a Explain rational agents and rationality L2 CO-1 **Marks:** 5

- Definition of Rational Agent – 2 Marks
- Rationality explanation – 2 Marks
- Example – 1 Mark

3.b Explain Wumpus world logic L2 CO-1 **Marks:** 5

- Definition of Wumpus World – 2 Marks
- Rules and environment – 2 Marks
- Agent objective – 1 Mark

UNIT–II

4.a Explain Support Vector Machine (SVM) and advantages L2 CO-2 **Marks:** 5

- Definition – 2 Marks
- Working principle – 2 Marks
- Advantages – 1 Mark

4.b Compare regression and classification with examples L2 CO-2 **Marks:** 5

- Regression explanation – 2 Marks
- Classification explanation – 2 Marks
- Examples – 1 Mark

OR

5.a Illustrate Decision Trees and working principle L2 CO-2 **Marks:** 7

- Definition – 2 Marks
- Tree structure – 2 Marks
- Splitting criteria – 2 Marks
- Example – 1 Mark

5.b Explain supervised learning with examplesL2 CO-2 **Marks:** 3

- Definition – 1 Mark
- Process – 1 Mark
- Example – 1 Mark

UNIT–III

6.a Describe steps of K-means clustering L2 CO-3 **Marks:** 5

- Definition – 1 Mark
- Step-by-step algorithm – 3 Marks
- Example/result – 1 Mark

6.b Explain dimensionality reduction and importance L2 CO-3 **Marks:** 5

- Definition – 2 Marks
- Importance/benefits – 2 Marks
- Example – 1 Mark

7. Explain Principal Component Analysis (PCA)L3 CO-3 **Marks:** 10

- Definition – 2 Marks
- Working principle – 4 Marks
- Steps involved – 3 Marks
- Applications – 1 Mark

UNIT–IV

8.a Explain perceptron model L2 CO-4 **Marks:** 5

- Definition – 1 Mark
- Structure – 2 Marks
- Working principle – 2 Marks

8.b Genetic operators: selection, crossover, mutation L2 CO-4 **Marks:** 5

- Selection – 2 Marks
- Crossover – 2 Marks
- Mutation – 1 Mark

OR

9. Explain CNN and applications L3 CO-4 **Marks:** 10

- Definition – 2 Marks
- Architecture – 4 Marks
- Working – 2 Marks
- Applications – 2 Marks

UNIT–V

10.a Explain Random Forest algorithm L2 CO-5 **Marks:** 5

- Definition – 1 Mark
- Working principle – 3 Marks
- Advantages – 1 Mark

10.b Deep auto-encoders L2 CO-5 **Marks:** 5

- Definition – 2 Marks
- Encoder-decoder concept – 2 Marks
- Applications – 1 Mark

OR

11.a Applications of deep neural networks L2 CO-5 **Marks:** 5

- Definition – 1 Mark
- Real-world applications – 3 Marks
- Example – 1 Mark

11.b Ensemble method in machine learning L2 CO-5 **Marks:** 5

- Definition – 2 Marks
- Types (bagging, boosting, stacking) – 2 Marks
- Example – 1 Mark

KEY Solutions

PVP 23

Code: 23ME3602

III B.Tech - II Semester - Regular Examinations – APRIL 2026

ARTIFICIAL INTELLIGENCE AND MACHINE LEARNING (MECHANICAL ENGINEERING)

Duration: 3 hours Max. Marks: 70

Note: 1. This question paper contains two Parts A and B.
2. Part-A contains 10 short answer questions. Each Question carries 2 Marks.
3. Part-B contains 5 essay questions with an internal choice from each unit. Each Question carries 10 marks.
4. All parts of Question paper must be answered in one place.

BL – Blooms Level CO – Course Outcome

1. a. What is an Intelligent Agent? BL L1 CO1 2Marks

Ans. In the realm of AI, Intelligent Agents drive automation and decision-making with cognitive abilities.Intelligent agents represent a subset of AI systems demonstrating intelligent behaviour, including adaptive learning, planning, and problem-solving.

1. b. What is First Order Logic (FOL)? BL L1 CO1 2Marks

Ans. First-Order Logic (FOL) also known as predicate logic helps us representknowledge, reason through problems and understand language. By building on propositional logic and adding quantifiers and predicates, FOL allows us to express more complex relationships and make decisions based on logical reasoning.

1. c. List the difference between supervised and unsupervised learning? BL L1 CO2 2Marks

Ans.

Aspect	Supervised Learning	Unsupervised Learning
Data	Labeled (input + output)	Unlabeled (only input
Goal	Predict outputs accurately	Discover patterns/relationships
Common Tasks	Classification, Regression	Clustering, Dimensionality Reduction

1. d.What is underfitting? BL L1 CO2 2Marks

Ans: Underfitting happens when the model fails to learn important patterns. It performs poorly on both training and testing data. Underfitting happens due to:

Underfitting = High Bias + Low Variance

1. e.What is dimensionality reduction? **BL L1 CO3 2Marks**

Ans. Dimensionality Reduction is the process of reducing the number of input variables (features) in a dataset while preserving as much important information as possible.

1. f. Define clustering? **BL L1 CO3 2Marks**

Ans: Clustering is an unsupervised machine learning technique used to group similar data points together without using labelled data. It helps discover hidden patterns or natural groupings in datasets by placing similar data points into the same cluster.

1. g. What are genetic operators? **BL L1 CO4 2Marks**

Ans:A Genetic Algorithm (GA) is a population-based evolutionary optimization technique inspired by the principles of natural selection and genetics. It works by iteratively evolving a population of candidate solutions using biologically motivated operators such as selection, crossover and mutation to find optimal or near-optimal solutions to complex problems where traditional optimization techniques are ineffective

1. h. What is an Artificial Neural Network? **BL L1 CO4 2Marks**

Ans: Artificial Neural Networks (ANNs) are computer systems designed to mimic how the human brain processes information. Just like the brain uses neurons to process data and make decisions, ANNs use artificial neurons to analyze data, identify patterns and make predictions. These networks consist of layers of interconnected neurons that work together to solve complex problemsKey Components of an ANN, Input Layer, Hidden LayersOutput Layer

1. i. What is boosting? **BL L1 CO5 2Marks**

Ans: Boosting is an ensemble learning technique that improves predictive accuracy by combining multiple weak learners into a single strong model. It works iteratively where each new model focuses on correcting the mistakes of its predecessors and gradually improves overall performance.Types of Boosting Algorithms:1.AdaBoost (Adaptive Boosting)2. Gradient Boosting3. XGBoost (Extreme Gradient Boosting)

1. j. What is deep learning? **BL L1 CO5 2Marks**

Ans. Deep Learning is transforming the way machines understand, learn and interact with complex data. Deep learning mimics neural networks of the human brain, it enables computers to autonomously uncover patterns and make informed decisions from vast amounts of unstructured data.

PART-B

Unit - 1

2. a Explain the characteristics of intelligent agents with examples? L2 CO-1 5 marks

Ans: In the realm of AI, Intelligent Agents stand as pivotal entities, driving automation and decision-making with cognitive abilities. This article explores the concept, architecture, functionalities, and real-world applications of these agents, shaping the modern AI landscape.

Understanding Intelligent Agents

Intelligent agents represent a subset of AI systems demonstrating intelligent behaviour, including adaptive learning, planning, and problem-solving. It operate in dynamic environments, where it makes decisions based on the information available to them. These agents dynamically adjust their behaviour, learning from past experiences to improve their approach and aiming for accurate solutions. The design of an intelligent agent typically involves four key components:

1. **Perception:** Agents have sensors or mechanisms to observe and perceive aspects of their environment. This may involve collecting data from the physical world, accessing databases, or receiving input from other software components.
2. **Reasoning:** Agents possess computational or cognitive capabilities to process the information they perceive. They use algorithms, logic, or machine learning techniques to analyze data, make inferences, and derive insights from the available information.
3. **Decision-Making:** Based on their perception and reasoning, agents make decisions about the actions they should take to achieve their goals. These decisions are guided by predefined objectives, which may include optimizing certain criteria or satisfying specific constraints.
4. **Action:** Agents execute actions in their environment to affect change and progress towards their goals. These actions can range from simple operations, such as sending a message or adjusting parameters, to more complex tasks, such as navigating a virtual world or controlling physical devices.

Examples of Intelligent Agents include self-driving cars, recommendation systems, virtual assistants, and game-playing AI.

Rational Agents and Rationality in Decision-Making

Intelligent agents are characterized by their rationality in decision-making, which aimto attain optimal outcomes or, in uncertain scenarios, the best-expected outcome.

A rational agent can be said to those, who do the right thing, It is an autonomous entity designed to perceive its environment, process information, and act in a way that maximizes the achievement of its predefined goals or objectives. Rational agents always aim to produce an optimal solution.

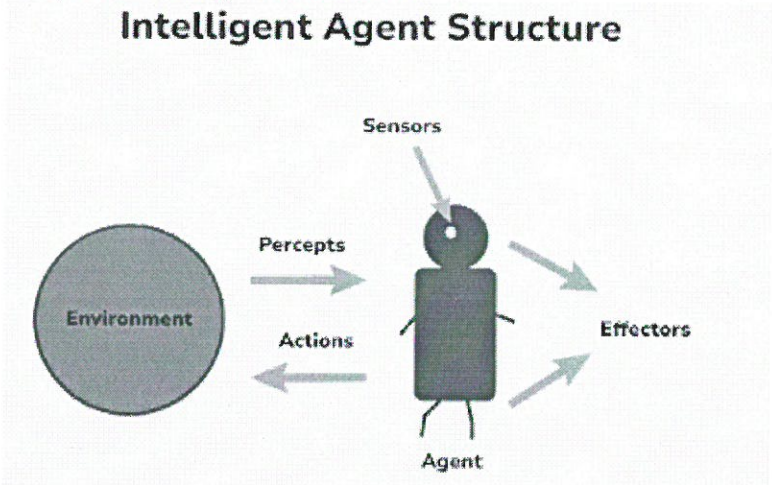

1. **Environment:** Environment is the area around the agent that it interacts with. An environment can be anything like a physical space, a room or a virtual space like a game world or the internet.
2. **Sensors:** Sensors are tools that AI agent uses to perceive their environment. They can be any physical like cameras, microphones, temperature sensors or a software sensor that read data from files.
3. **Actuators:** Actuators are tools that AI agent uses to interact with their environment through some actions. They can be any physical actuators like wheels, motors, robotic hands, or computer screens or they can be software actuators that send messages.
4. **Effectors:** Effectors take instructions from decision making mechanism and translates them into actions and these actions are performed through actuators.

2.b.Distinguish between propositional logic and first-order logic. CO-1, L2 5 Marks

Ans.

Aspect	Propositional Logic	First-Order Logic
Definition	Logic that deals with simple, complete statements (propositions) that are either true or false	Logic that represents objects, properties, and relationships using predicates
Basic Components	Propositions (P, Q), logical connectives (AND, OR, NOT)	Predicates, variables, constants, functions, quantifiers
Internal Structure	No internal structure (treated as atomic units)	Has structure (subject–predicate relationships)
Expressiveness	Limited (cannot represent relations or details)	Highly expressive (can model real-world scenarios)
Quantifiers	Not used	Uses universal (∀) and existential (□) quantifiers
Representation of Knowledge	Represents simple facts only	Represents complex facts, rules, and relationships

Aspect	Propositional Logic	First-Order Logic
Example Statement	**P: "It is raining"**	**Human(Ram), ∀x Human(x) → Mortal(x)**
Real-world Application	**Simple circuits, basic decision systems**	**Knowledge representation in Artificial Intelligence, expert systems**

OR

3.a.Explain the concept of rational agents and rationality? L2, CO-1, 5 Marks

A. In Artificial Intelligence, a rational agent is an entity that perceives its environment through sensors and acts upon that environment using actuators in a way that maximizes the achievement of its goals.

Rational Agent

A rational agent chooses the best possible action based on:

- Its current percept (what it senses now)
- Its percept history (past experience)
- Its knowledge of the environment
- The actions available to it

Example:

A self-driving car selects actions like braking or turning based on traffic, road conditions, and safety goals.

Rationality

Rationality refers to doing the "right thing," meaning selecting actions that maximize the agent's expected performance measure given the available information.

It does not mean perfection or complete knowledge. Instead, it depends on:

- Performance measure (success criteria)
- Percept sequence (what the agent has seen so far)
- Prior knowledge of the environment
- Available actions

Example:

A chess-playing AI chooses the move that gives the highest chance of winning, even if it cannot predict all future moves.

3.b. Explain the concept of Wumpus world logic?

A.The Wumpus World is a simple logical environment used in Artificial Intelligence to demonstrate how an intelligent agent uses logic and reasoning to make decisions under uncertainty.It consists of a 4×4 grid cave where an agent starts at (1,1) and aims to find gold and exit safely. The environment contains hazards such as a Wumpus (monster) and pits, which the agent must avoid.

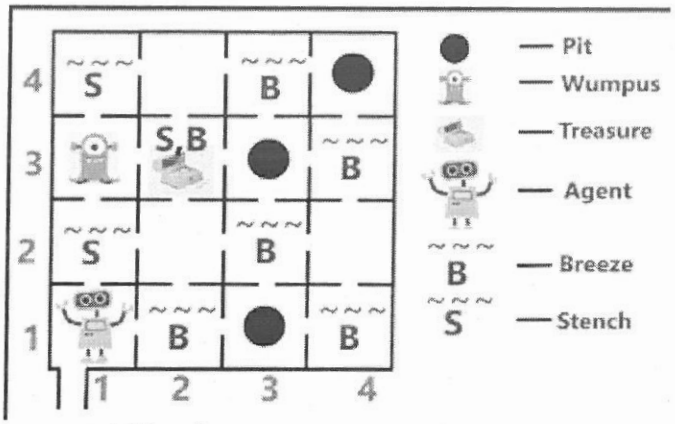

Key Concepts of Wumpus World Logic

1. Partial Observability

The agent cannot see the entire environment. It only perceives clues like:

- *Breeze* → indicates a nearby pit
- *Stench* → indicates a nearby Wumpus
- *Glitter* → indicates gold

2. Knowledge Base (KB)

The agent stores facts and rules in a knowledge base. For example:

- If there is no breeze in a cell → adjacent cells have no pits
- If there is stench → Wumpus is nearby

3. Logical Inference

Using propositional logic, the agent derives new information from known facts to identify safe and unsafe cells.

4. Decision Making

Based on inferred knowledge, the agent chooses actions such as moving, grabbing gold, or exiting safely.

5. Goal-Oriented Behavior

The agent's goal is to maximize success by collecting gold while avoiding danger.

UNIT-II

4.a Explain Support Vector Machine (SVM) and its advantages? L2 CO-2 5 Marks

A. Support Vector Machine (SVM) is a supervised learning algorithm in Machine Learning used for classification and regression. It works by finding the optimal hyperplane that separates data points of different classes with the maximum margin (distance between the nearest points of each class, called support vectors).

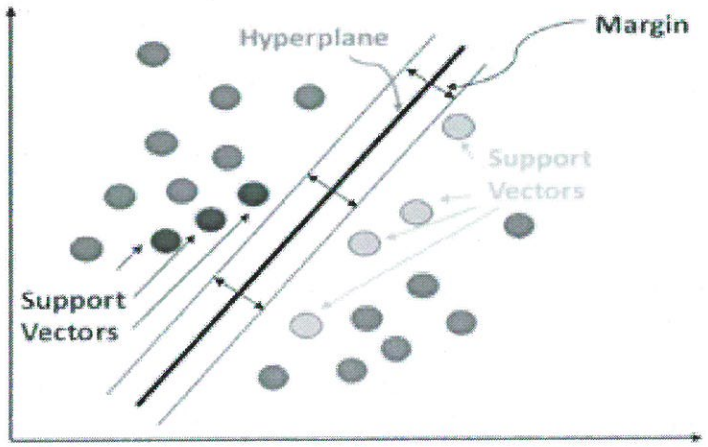

- The data points closest to the boundary are called **support vectors**
- SVM tries to **maximize the margin** for better generalization
- For non-linear data, SVM uses the **kernel trick** (e.g., polynomial, RBF) to map data into higher dimensions where it becomes separable

Advantages of SVM

1. **High Accuracy**
   Works well for both linear and non-linear classification problems
2. **Effective in High Dimensions**
   Performs well even when the number of features is large
3. **Robust to Overfitting**
   Maximizing margin reduces the risk of overfitting
4. **Works with Small Datasets**
   Gives good performance even with limited data
5. **Versatile (Kernel Trick)**
   Can handle complex data using different kernel functions

4.b. Compare regression and classification problems with examples? L2 CO-2 5 Marks

Ans.

Aspect	**Regression**	**Classification**
Definition	Predicts continuous numerical values	Predicts discrete class labels
Output Type	Continuous (real-valued)	Categorical (finite set of classes)
Nature of Problem	Quantitative prediction	Qualitative decision-making

Aspect	Regression	Classification
Goal	Estimate a function to map input → numeric output	Assign input to one of the predefined categories
Examples	House price prediction, rainfall estimation, salary prediction	Spam detection, disease diagnosis, image classification
Typical Algorithms	Linear Regression, Ridge, Lasso, SVR	Logistic Regression, Decision Tree, SVM, KNN, Naive Bayes
Mathematical Function	Fits a continuous function (e.g., line/curve)	Finds decision boundaries between classes
Output Range	Infinite/continuous range	Limited/discrete set
Error Measurement	Mean Squared Error (MSE), RMSE, MAE	Accuracy, Precision, Recall, F1-score, Confusion Matrix
Visualization	Best-fit line/curve	Decision boundary or regions
Data Requirement	Requires numeric dependent variable	Requires labeled categorical data
Complexity	Usually simpler	Can be more complex (multi-class problems)
Use Case Suitability	When exact value prediction is needed	When grouping or labeling is required

OR

5. a Illustrate Decision Trees and their working principle? L3 CO2-7Marks

A. A Decision Tree helps us to make decisions by mapping out different choices and their possible outcomes. It's used in machine learning for tasks like classification and prediction. In this article, we'll see more about Decision Trees, their types and other core concepts.

A Decision Tree helps us make decisions by showing different options and how they are related. It has a tree-like structure that starts with one main question called the root node which represents the entire dataset. From there, the tree branches out into different possibilities based on features in the data.

Root Node: Starting point representing the whole dataset.

Branches: Lines connecting nodes showing the flow from one decision to another.

Internal Nodes: Points where decisions are made based on data features.

Leaf Nodes: End points of the tree where the final decision or prediction is made

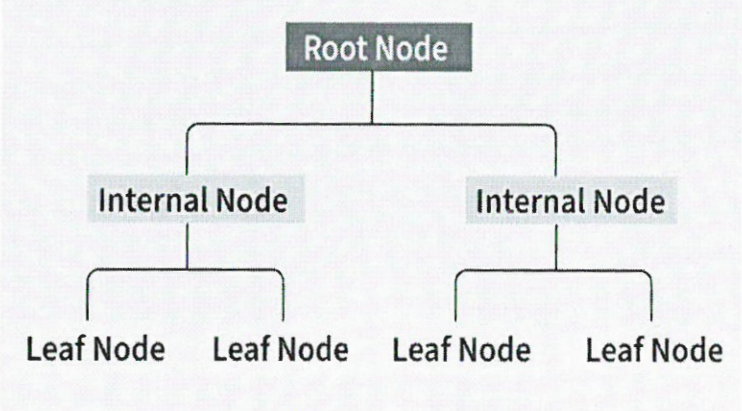

A Decision Tree also helps with decision-making by showing possible outcomes clearly. By looking at the "branches" we can quickly compare options and figure out the best choice.

There are mainly two types of Decision Trees based on the target variable:

Classification Trees: Used for predicting categorical outcomes like spam or not spam. These trees split the data based on features to classify data into predefined categories.

Regression Trees: Used for predicting continuous outcomes like predicting house prices. Instead of assigning categories, it provides numerical predictions based on the input features

5.b Explain supervised learning with suitable examples? L2, CO-2 3Marks

Ans: Supervised learning is a type of Machine Learning where a model is trained using labeled data—each input is paired with the correct output. The goal is to learn a mapping from inputs to outputs so the model can predict results for new, unseen data.

How It Works

- dataset contains input features (X) and target labels (Y)
- The algorithm learns the relationship between X and Y
- After training, it predicts outputs for new inputs

Types of Supervised Learning

1. Classification

Output is a category or class

Example: Email spam detection (Spam / Not Spam)

2. Regression

Output is a continuous value

Example: Predicting house prices based on size and location

Examples

- **House Price Prediction (Regression):**
  The model learns from past data (area, rooms, location → price) to predict future prices.
- **Email Spam Detection (Classification):**
  The model learns from labeled emails to classify new emails as spam or not.
- **Handwritten Digit Recognition:**
  The system learns from labeled images (0–9) to identify new digits.

UNIT-III

6.a Describe the steps of the K-means clustering algorithm?L2 CO-3, 5 Marls

Ans: K-means clustering is an unsupervised algorithm in Machine Learning that partitions data into **K clusters** such that points within a cluster are similar and points across clusters are dissimilar.

Step-by-Step Procedure

1. Select the Number of Clusters (K)
Choose the value of K (number of clusters) based on prior knowledge or methods like the elbow method.

2. Initialize Centroids
Randomly pick K data points or use methods like **K-means++** to select initial centroids more effectively.

3. Distance Calculation
Compute the distance between each data point and every centroid (commonly **Euclidean distance**).

4. Cluster Assignment (E-step)
Assign each data point to the cluster whose centroid is **closest** to it.

5. Update Centroids (M-step)
Recalculate each centroid as the **mean of all data points** assigned to that cluster.

6. Check for Convergence
Determine whether:
Centroids no longer change significantly, or
Data point assignments remain the same
If not, repeat Steps 3–5.

6. **Termination**
The algorithm stops when convergence is reached or after a fixed number of iterations.

Note: step by step process with example values also considered

6.b Explain dimensionality reduction and its importance? L2, CO-3 5 Marks

A. Dimensionality reduction is a technique in Machine Learning used to reduce the number of input features (variables) in a dataset while preserving as much important information as possible. It simplifies high-dimensional data into a lower-dimensional form.

Concept

In many real-world problems, datasets contain hundreds or thousands of features, which can make computation complex and inefficient. Dimensionality reduction transforms this data into fewer features by:

- Removing irrelevant or redundant data
- Combining features while retaining key patterns

Example: Reducing 100 features of an image dataset to 20 important features without losing significant information.

Types

1. Feature Selection

Selecting only the most important features (e.g., removing unnecessary columns)

2. Feature Extraction

Creating new features from existing ones

Example: Principal Component Analysis (PCA)

Importance of Dimensionality Reduction

1. Reduces Computational Cost
Less data → faster processing and training
2. Avoids Overfitting
Eliminates noise and irrelevant features
3. Improves Model Performance
Better accuracy due to reduced complexity
4. Handles Curse of Dimensionality
Improves performance when data has too many features
5. Better Visualization
Helps visualize high-dimensional data in 2D or 3D

OR

7. Illustrate the working principle of Principal Component Analysis? L3, CO-3 10 Mrls

Ans: Principal Component Analysis (PCA) is a widely used dimensionality reduction technique in Machine Learning that transforms a dataset with many correlated variables into a smaller set of uncorrelated variables called *principal components*, while preserving maximum variance (information).

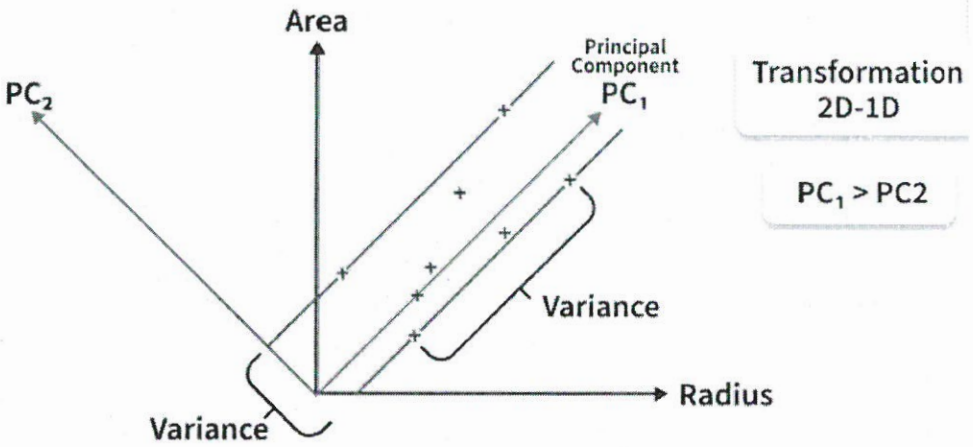

PCA works by **rotating the coordinate system** such that:

- The first axis (PC1) captures **maximum variance**
- The second axis (PC2) captures the **next highest variance**, orthogonal to PC1
- Remaining components follow similarly

Step-by-Step Working of PCA

1. Data Standardization

- Normalize data so each feature has **zero mean and unit variance**
- Ensures all features contribute equally

2. Compute Mean Vector

- Calculate mean of each feature
- Subtract mean from dataset (centering the data)

3. Covariance Matrix Calculation

- Compute covariance matrix to understand relationships between variables
- Shows how features vary with respect to each other

4. Eigenvalues and Eigenvectors

- Calculate eigenvalues and eigenvectors of the covariance matrix
- **Eigenvectors → directions (principal components)**
- **Eigenvalues → magnitude of variance along those directions**

5. Sort Eigenvalues

- Arrange eigenvalues in descending order
- Select top k eigenvectors corresponding to largest eigenvalues

6. Form Feature Vector (Projection Matrix)

- Construct a matrix using selected eigenvectors

7. Transform the Data

- Project original data onto new feature space
- Result is reduced-dimensional data with maximum retained variance

UNIT-IV

8.a. Explain the working principle of a perceptron model? L2, CO-4 5 Marks

A Perceptron is the simplest form of a neural network that makes decisions by combining inputs with weights and applying an activation function. It is mainly used for binary classification problems. It forms the basic building block of many deep learning models.

Takes multiple inputs and assigns weights
Computes a weighted sum and applies a threshold
Outputs either 0 or 1 (binary outcome)
Forms the foundation of larger neural networks

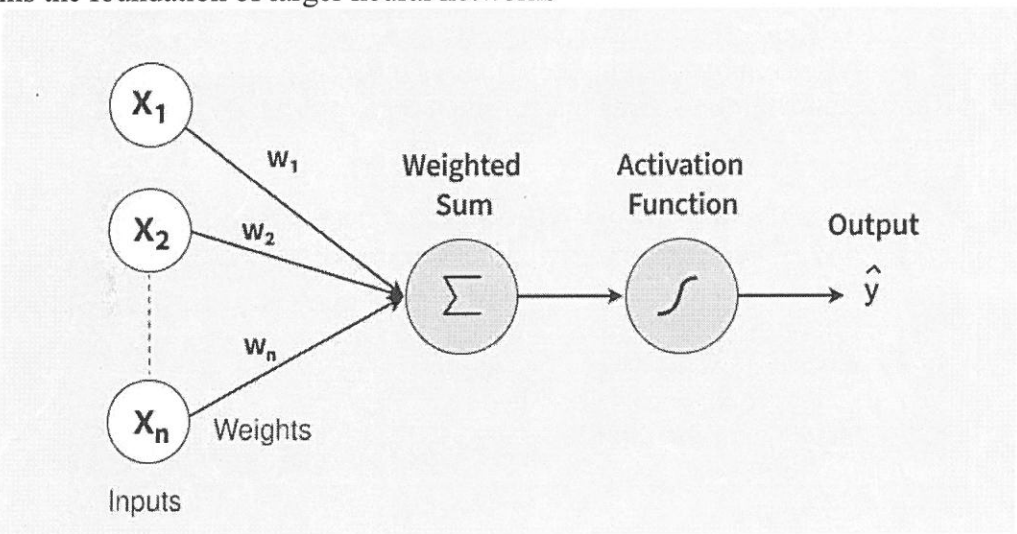

Step-by-Step Working

1. Input Layer

The perceptron receives multiple inputs:

$$x_1, x_2, x_3, \ldots, x_n$$

2. Assign Weights

Each input is multiplied by a weight:

$$w_1, w_2, w_3, \ldots, w_n$$

3. Weighted Sum + Bias

The perceptron computes a linear combination:

$$z = \sum(w_i x_i) + b$$

4. Activation Function

A **step function** (threshold function) is applied:

- If $z \geq 0 \rightarrow$ Output = 1
- If $z < 0 \rightarrow$ Output = 0

3. **Output**

The result is a **binary classification** (Yes/No, 0/1)

Advantages

- Simple and easy to implement
- Works well for linearly separable data

Limitations

- Cannot solve non-linear problems (e.g., XOR problem)
- Only supports binary classification

8.b Discuss genetic operators such as selection, crossover and mutation? L2 CO-2 5 Marks

A. In Genetic Algorithms (GA)—a search and optimization technique in Artificial Intelligence—genetic operators guide how solutions evolve from one generation to the next.

1. Selection (Survival of the Fittest)

- Purpose: Choose better individuals from the population to become parents
- Basis: Fitness value (higher fitness → higher selection probability)
- Common Methods:
  - *Roulette Wheel Selection* (probability-based)
  - *Tournament Selection* (best among a subset)
  - *Rank Selection* (based on ranking rather than raw fitness)
- Effect: Improves overall solution quality by favoring strong candidates while still allowing diversity

2. Crossover (Recombination)

- Purpose: Combine genetic information of two parents to create offspring
- How it Works: Portions of parent chromosomes are exchanged
- Types:
  - *Single-point crossover* (split at one point)
  - *Two-point crossover* (swap between two points)
  - *Uniform crossover* (mix genes randomly)
- Effect: Explores new regions of the solution space and produces potentially better offspring

3. Mutation (Random Variation)

- Purpose: Introduce random changes in offspring chromosomes
- How it Works: Small alterations such as flipping bits or modifying values
- Types:
  - *Bit-flip mutation* (binary representation)
  - *Swap mutation* (exchange positions)
  - *Gaussian mutation* (for real-valued data)

- Effect: Maintains diversity, avoids premature convergence, and helps escape local optima

OR

9. Illustrate the Convolutional Neural Networks and their applications? L3, CO-4 10 Marks

A. Convolutional Neural Networks (CNNs) are a class of deep learning models in Machine Learning designed to process grid-like data such as images. They automatically learn spatial features (edges, textures, shapes) using convolution operations.

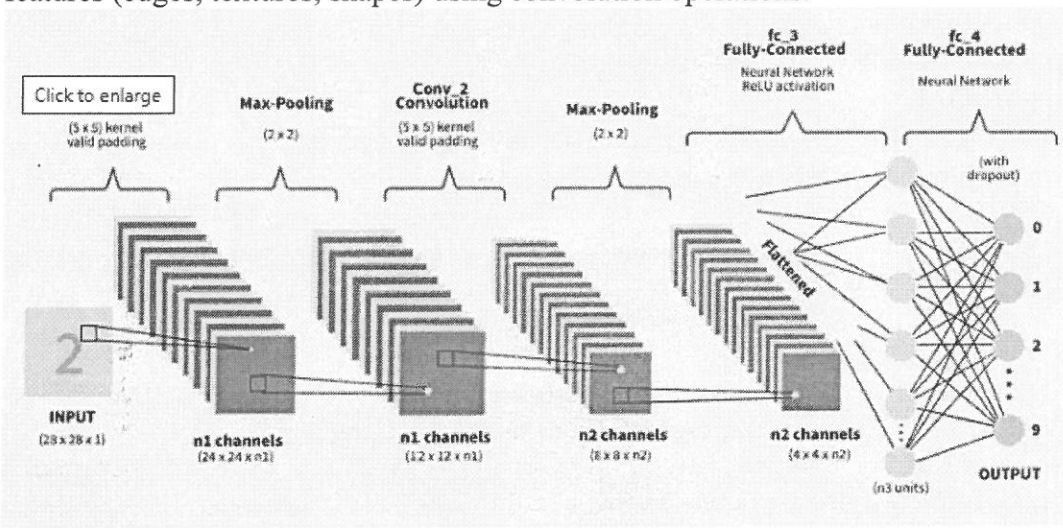

CNN method

Working Principle of CNN

1. Input Layer
   - Takes image data (height × width × channels, e.g., RGB)
2. Convolution Layer
   - Applies filters (kernels) that slide over the image
   - Extracts features like edges, corners, textures
   - Produces feature maps
3. Activation Function (ReLU)
   - Applies non-linearity: replaces negative values with zero
   - Helps the model learn complex patterns
4. Pooling Layer
   - Reduces spatial size (downsampling)
   - Common: Max Pooling (takes maximum value in region)
   - Reduces computation and prevents overfitting
5. Fully Connected Layer (FC Layer)
   - Flattens feature maps into a vector
   - Performs classification based on learned features
6. Output Layer
   - Produces final prediction (e.g., class probabilities using Softmax)

UNIT-V

10.a. Explain the working of Random Forest algorithm? L2 CO-5 5Marks

A. Random Forest is a supervised learning algorithm in Machine Learning that builds multiple **decision trees** and combines their outputs to improve accuracy and reduce overfitting. It is an **ensemble method** based on bagging (bootstrap aggregating).

Working of Random Forest Algorithm

1. Bootstrap Sampling (Bagging)

- From the original dataset, multiple **random samples (with replacement)** are created
- Each sample is used to train a separate decision tree

2. Build Decision Trees

- For each tree, a **random subset of features** is selected at each split
- This introduces diversity among trees

3. Tree Growth

- Each tree is grown independently (usually fully grown without pruning)

4. Prediction from Each Tree

- For a new input, every tree gives its own prediction

5. Aggregation (Voting/Averaging)

- **Classification:** Final output is based on **majority voting**
- **Regression:** Final output is the **average** of all tree predictions

Assumptions of Random Forest

Each tree makes its own decisions: Every tree in the forest makes its own predictions without relying on others.

Random parts of the data are used: Each tree is built using random samples and features to reduce mistakes.

Enough data is needed: Sufficient data ensures the trees are different and learn unique patterns and variety.

Different predictions improve accuracy: Combining the predictions from different trees leads to a more accurate final result.

Advantages

- Reduces overfitting compared to a single decision tree
- Handles large datasets and high-dimensional data
- Provides good accuracy and robustness

10.b. Summarize the working of deep auto-encoders? L2 CO-5, 5 Marks

A deep auto-encoder is a multilayer neural network in Machine Learning that learns to compress data into a compact form and then reconstruct it, capturing the most important features.

Working Principle (Step-wise)

1. Input Layer
- Takes high-dimensional data (images, signals, etc.)

2. Encoder (Compression Phase)
- Multiple hidden layers progressively reduce dimensionality
- Learns a compact representation called the latent vector (bottleneck)

3. Bottleneck Layer
- Central layer with minimum neurons

- Stores the most important features of the input

4. Decoder (Reconstruction Phase)
- Expands the compressed data back to original dimensions
- Mirrors the encoder structure

5. Output Layer
- Produces reconstructed data similar to input

6. Training Objective
- Minimize reconstruction error (e.g., Mean Squared Error)
- Adjust weights using backpropagation

Types of Auto-Encoders
- Sparse Auto-Encoder → encourages fewer active neurons
- Denoising Auto-Encoder → reconstructs clean data from noisy input
- Variational Auto-Encoder (VAE) → learns probabilistic representations

Applications
- Data compression
- Feature extraction
- Image denoising
- Anomaly detection
- Pretraining deep networks

Advantages
- Reduces dimensionality efficiently
- Captures complex patterns
- Improves performance of downstream models

Limitations
- May learn trivial identity mapping if not constrained
- Requires large data for effective training

OR

11.a Describe the applications of deep neural networks in real-world problems? L2, co5, 5 Marks

Ans: Deep Neural Networks (DNNs) are powerful models in Machine Learning that learn complex patterns from large datasets. They are widely used across many real-world domains due to their high accuracy and ability to model non-linear relationships.

Visual & Image-Based Applications

1. Image Classification & Object Detection
- Identifying objects (cars, animals, people) in images
- Used in surveillance, retail analytics

2. Medical Imaging
- Detecting tumors, diseases from X-rays, MRI, CT scans
- Assists doctors in diagnosis

3. Face Recognition
- Used in smartphones, security systems, attendance systems

4. Autonomous Vehicles
- Detects lanes, pedestrians, traffic signs
- Enables safe navigation

5. Speech Recognition

- Converts speech to text (voice assistants)

6. Natural Language Processing (NLP)
- Chatbots, language translation, text summarization

7. Sentiment Analysis
- Analyzing opinions in reviews, social media

8. Recommendation Systems
- Suggesting products, movies (e.g., streaming platforms)

9. Fraud Detection
- Identifying suspicious financial transactions

10. Predictive Maintenance
- Predicting equipment failures in industries

11. Demand Forecasting
- Predicting sales, inventory needs

12. Drug Discovery & Healthcare Research
- Identifying new medicines and treatment methods

11.b Describe the ensemble method in machine learning L2, co5, 5 Marks

Ans: Ensemble learning is a method where multiple models are combined instead of using just one. Even if individual models are weak, combining their results gives more accurate and reliable predictions.

Multiple Models: Uses many small models together
Better Accuracy: Combined results improve performance
Reduced Errors: Mistakes of one model are balanced by those of others
Simple Idea: Like taking advice from a group instead of one person

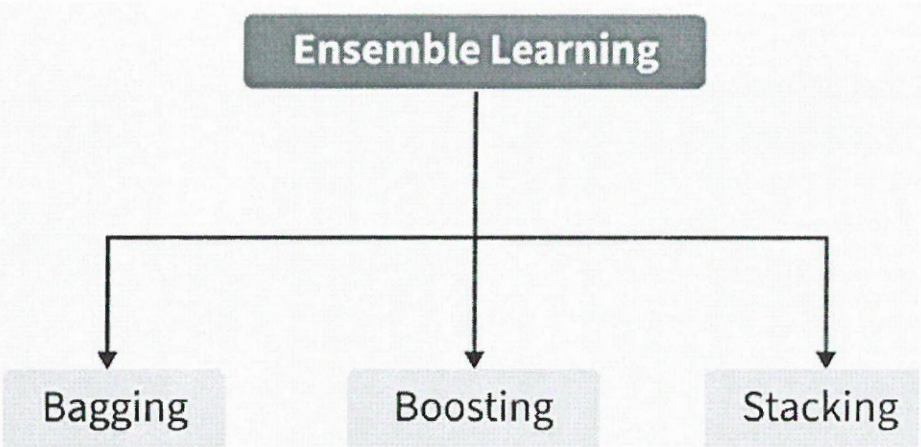

Types of Ensemble Learning

There are three main types of ensemble methods:

Bagging (Bootstrap Aggregating): Models are trained independently on different random subsets of the training data. Their results are then combined—usually by averaging (for regression) or voting (for classification). This helps reduce variance and prevents overfitting.

Boosting: Models are trained one after another. Each new model focuses on fixing the errors made by the previous ones. The final prediction is a weighted combination of all models, which helps reduce bias and improve accuracy.
Stacking (Stacked Generalization): Multiple different models (often of different types) are trained and their predictions are used as inputs to a final model, called a meta-model. The meta-model learns how to best combine the predictions of the base models, aiming for better performance than any individual model.