4.5 Overview of Supervised and Unsupervised Learning, Reinforcement Learning Basics.

Machine learning (ML) is a subset of artificial intelligence (AI). It enables systems to learn from data, identify patterns and make decisions with minimal human intervention. The three primary types of ML are:

Supervised Learning: Learning from labelled data.

Unsupervised Learning: Discovering patterns in unlabeled data.

Reinforcement Learning: Learning through interactions with environment. Each approach has unique characteristics, advantages and real-world applications.

Supervised Learning

Supervised learning is like learning with a teacher. The model is trained on a labeled dataset, meaning each input has a corresponding output. The key characteristics of supervised learning are:

Labeled Data: Training data has predefined labels.

Types of Problems: Used for classification task like spam detection and regression task like predicting house prices.

Algorithms: Linear Regression, Logistic Regression, SVM, Decision Trees, Neural Networks.

Unsupervised Learning

Unsupervised learning works with data that has no predefined labels. The model identifies patterns, clusters or associations independently. The key characteristics of unsupervised learning are:

Unlabeled Data: No predefined outputs.

Types of Problems: Used for Clusteringtask like customer segmentation and association task like market basket analysis.

Algorithms: K-Means, Hierarchical Clustering, PCA, Autoencoders.

Reinforcement Learning (RL)

Reinforcement learning involves an agent that interacts with an environment, learning through rewards and penalties to maximize long-term success. The key characteristics of reinforcement learning are:

Interaction-Based Learning: The agent learns by taking actions and receiving feedback.

No Labeled Data: Learns from trial and error.

Algorithms: Q-learning, SARSA, Deep Q-Networks (DQN)

Applications

1. Supervised Learning (learns with examples)

Email spam filter → marks mails as spam or not spam

Medical diagnosis → predicts if a patient has a disease

House price prediction → guesses price based on size & location

Voice assistants → recognize speech commands

2. Unsupervised Learning (finds patterns without answers)

Customer grouping → shops like Amazon group people by buying habits

Market basket analysis → "People who bought X also buy Y"

Anomaly detection → finds unusual credit card transactions (fraud)

Data compression → reduces large images or files (PCA)

3. Reinforcement Learning (learns by trial & error)

Self-driving cars → learn to drive safely by rewards (safe moves) and penalties (accidents)

Robotics → robots learning to walk, pick, or place items

Games → AI beating humans in Chess, Go, or video games Dynamic pricing → airlines adjusting ticket prices based on demand Quick recap in one line each:

Supervised → Prediction (future from past examples)

Unsupervised → Grouping & patterns (hidden structure)

Reinforcement → Learning by doing (rewards & penalties)