

Introduction to AI and ML

Artificial Intelligence (AI):

- **Definition:** AI is the simulation of human intelligence processes by machines.
- **Types:**
 - **Narrow AI:** Performs a specific task (e.g., voice assistants).
 - **General AI:** Can perform any intellectual task a human can do (still theoretical).
 - **Super AI:** Surpasses human intelligence (hypothetical).

Machine Learning (ML):

- **Definition:** A subset of AI where machines learn from data and improve over time without being explicitly programmed.
 - **Types:**
 - **Supervised Learning**
 - **Unsupervised Learning**
 - **Reinforcement Learning**
-

□ Supervised Learning

Overview:

- The algorithm is trained on a **labeled dataset** (input-output pairs).
- Goal: Learn a mapping from inputs to outputs.

1. Regression:

- **Goal:** Predict a continuous output.
- **Examples:**
 - Predicting house prices.
 - Estimating temperature.
- **Common Algorithms:**
 - Linear Regression
 - Polynomial Regression
 - Decision Trees (for regression tasks)

2. Classification:

- **Goal:** Predict discrete labels.
- **Examples:**
 - Spam detection (spam/not spam)
 - Image recognition (cat/dog/etc.)

- **Common Algorithms:**
 - Logistic Regression
 - k-Nearest Neighbors (k-NN)
 - Support Vector Machines (SVM)
 - Decision Trees, Random Forests
 - Neural Networks
-

□ Unsupervised Learning

Overview:

- No labeled data; the model finds structure in the data.

1. Clustering:

- **Goal:** Group similar instances.
- **Examples:**
 - Customer segmentation
 - Document/topic grouping
- **Common Algorithms:**
 - k-Means
 - DBSCAN
 - Hierarchical Clustering

2. Dimensionality Reduction:

- **Goal:** Reduce the number of features while preserving structure.
 - **Examples:**
 - Data visualization
 - Noise reduction
 - **Common Techniques:**
 - Principal Component Analysis (PCA)
 - t-SNE
 - Autoencoders
-

□ Deep Learning

Overview:

- Subset of ML focused on neural networks with many layers (hence "deep").
- Excels at processing unstructured data (images, audio, text).

1. Neural Networks:

- **Structure:** Layers of neurons (input → hidden → output).
- **Activation Functions:** ReLU, Sigmoid, Tanh.
- **Training:** Via backpropagation and gradient descent.

2. Convolutional Neural Networks (CNNs):

- **Used for:** Image and video recognition.
- **Key components:**
 - Convolutional layers (feature extraction)
 - Pooling layers (dimensionality reduction)
 - Fully connected layers (classification)

3. Recurrent Neural Networks (RNNs):

- **Used for:** Sequential data like time series and natural language.
 - **Challenges:** Vanishing gradient problem.
 - **Variants:**
 - LSTM (Long Short-Term Memory)
 - GRU (Gated Recurrent Units)
-

□ Reinforcement Learning (RL)

Overview:

- An agent learns to make decisions by interacting with an environment.
- Receives **rewards** or **penalties** to improve its strategy (policy).

Key Concepts:

- **Agent:** Learner/decision maker.
- **Environment:** Where the agent acts.
- **State:** Current situation.
- **Action:** Move made by the agent.
- **Reward:** Feedback from the environment.

1. Markov Decision Processes (MDPs):

- Formal framework for modeling RL problems.
- Defined by:
 - **States (S)**
 - **Actions (A)**

- **Transition probabilities (P)**
- **Rewards (R)**
- **Discount factor (γ)**

2. Q-Learning:

- A value-based method in RL.
- Learns the **Q-values** (expected future rewards) for each action in each state.
- **Update Rule:**

$$Q(s, a) \leftarrow Q(s, a) + \alpha [r + \gamma * \max_{a'} Q(s', a') - Q(s, a)]$$

Where:

- α is the learning rate,
- γ is the discount factor,
- r is the reward.