



Fundamentals of AI

Core Concepts

Artificial Intelligence (AI)	The simulation of human intelligence processes by computer systems.
Machine Learning (ML)	A subset of AI that allows systems to learn from data without being explicitly programmed.
Deep Learning (DL)	A subset of ML using artificial neural networks with multiple layers to analyze data with complex structures.
Supervised Learning	Learning from labeled data to predict outcomes on new, unseen data.
Unsupervised Learning	Learning from unlabeled data to discover patterns and relationships.
Reinforcement Learning (RL)	Training an agent to make sequences of decisions in an environment to maximize a reward.

AI Agents

Definition	An entity that perceives its environment through sensors and acts upon that environment through actuators.
Rationality	An agent is rational if it chooses actions that maximize its expected performance measure, given its percept sequence, built-in knowledge, and possible actions.
Types of Agents	Simple reflex agents, model-based reflex agents, goal-based agents, and utility-based agents.
PEAS Description	Performance measure, Environment, Actuators, Sensors.
Example PEAS - Self-Driving Car	Performance: Safety, travel time, comfort, legal compliance. Environment: Roads, other traffic, pedestrians, weather conditions. Actuators: Steering, accelerator, brakes, signals. Sensors: Cameras, radar, GPS, speedometers.

Search Algorithms

Breadth-First Search (BFS)	Explores all the neighbor nodes at the present depth prior to moving on to the nodes at the next depth level. Complete and optimal if step cost is 1.
Depth-First Search (DFS)	Explores as far as possible along each branch before backtracking. Not complete and not optimal.
<i>A Search*</i>	An informed search algorithm, or a best-first search, meaning that it is formulated in terms of weighted graphs: starting from a specific starting node of a graph, it aims to find a path to the given goal node having the smallest cost. Complete and optimal if heuristic is admissible.
Heuristic Function	Estimates the cost from the current state to the goal state. Used in informed search algorithms like A*.
Admissible Heuristic	A heuristic is admissible if it never overestimates the cost to reach the goal. $h(n) \leq h^*(n)$

Machine Learning Algorithms

Supervised Learning

Linear Regression	Models the relationship between a dependent variable and one or more independent variables by fitting a linear equation to observed data. Formula: $y = mx + b$
Logistic Regression	A statistical model that uses a logistic function to model a binary dependent variable. Used for classification problems. Formula: $p = \frac{1}{1 + e^{-z}}$ where $z = mx + b$
Support Vector Machines (SVM)	Finds the optimal hyperplane that maximizes the margin between different classes in the data. Kernel functions can be used for non-linear data.
Decision Trees	A tree-like model that makes decisions based on features of the data. Easy to interpret, but prone to overfitting.
Random Forest	An ensemble learning method that operates by constructing a multitude of decision trees at training time and outputting the class that is the mode of the classes (classification) or mean prediction (regression) of the individual trees. Reduces overfitting.
K-Nearest Neighbors (KNN)	Classifies a data point based on the majority class of its k nearest neighbors. Simple but computationally expensive for large datasets.

Unsupervised Learning

K-Means Clustering	Partitions n observations into k clusters, in which each observation belongs to the cluster with the nearest mean (cluster centers or cluster centroid), serving as a prototype of the cluster.
Hierarchical Clustering	Builds a hierarchy of clusters by iteratively merging or splitting them. Can be agglomerative (bottom-up) or divisive (top-down).
Principal Component Analysis (PCA)	A statistical procedure that uses an orthogonal transformation to convert a set of observations of possibly correlated variables into a set of values of linearly uncorrelated variables called principal components. Used for dimensionality reduction.
Association Rule Learning	Discovers interesting relations between variables in large databases. Example: Market Basket Analysis (Apriori algorithm).

Reinforcement Learning

Q-Learning	A model-free reinforcement learning algorithm to learn a policy telling an agent what action to take under what circumstances. It learns an optimal policy even when the agent is following a sub-optimal policy. Formula: $Q(s, a) = Q(s, a) + \alpha [R(s, a) + \gamma \max_{a'} Q(s', a') - Q(s, a)]$
SARSA	On-policy algorithm that updates the Q-value based on the action the agent actually takes. Stands for State-Action-Reward-State-Action.
Policy Gradient Methods	Directly optimize the policy without using a value function. Example: REINFORCE.
Markov Decision Process (MDP)	A mathematical framework for modeling decision-making in situations where outcomes are partly random and partly under the control of a decision maker. Defines: States, Actions, Transition probabilities, Rewards.

Neural Networks and Deep Learning

Basic Neural Network

Perceptron	The basic unit of a neural network that takes several inputs, weighs them, sums them up, and passes the result through an activation function to produce an output.
Activation Functions	Functions that introduce non-linearity to the output of a neuron. Examples: Sigmoid, ReLU, Tanh. Sigmoid: $\sigma(x) = \frac{1}{1 + e^{-x}}$ ReLU: $f(x) = \max(0, x)$
Feedforward Neural Network	A neural network where the connections between the nodes do not form a cycle. Information flows in one direction, from the input layer to the output layer.
Backpropagation	An algorithm used to train feedforward neural networks by calculating the gradient of the loss function with respect to the weights and biases, and updating them accordingly.
Loss Function	A function that quantifies the error between the predicted output and the actual output. Examples: Mean Squared Error (MSE), Cross-Entropy.

Convolutional Neural Networks (CNNs)

Convolutional Layer	Applies a filter (kernel) to the input to produce a feature map. Used for feature extraction.
Pooling Layer	Reduces the spatial size of the feature maps, reducing the number of parameters and computational complexity. Examples: Max Pooling, Average Pooling.
ReLU Layer	Applies the ReLU activation function to introduce non-linearity.
Fully Connected Layer	Connects every neuron in one layer to every neuron in the next layer. Used for classification.
Common Architectures	LeNet, AlexNet, VGGNet, ResNet, Inception.

Recurrent Neural Networks (RNNs)

Recurrent Layer	Processes sequential data by maintaining a hidden state that captures information about the past. The output of the hidden state is fed back into the network.
Long Short-Term Memory (LSTM)	A type of RNN architecture that addresses the vanishing gradient problem by using memory cells and gates to control the flow of information. Good for remembering long term dependencies.
Gated Recurrent Unit (GRU)	A simplified version of LSTM with fewer parameters, making it faster to train. Also effective at capturing long-term dependencies.
Applications	Natural Language Processing (NLP), speech recognition, time series analysis.

AI Ethics and Future Trends

Ethical Considerations

Bias - AI systems can perpetuate and amplify biases present in the data they are trained on, leading to unfair or discriminatory outcomes.
Transparency - Lack of transparency in AI models can make it difficult to understand how decisions are made, hindering accountability.
Privacy - AI systems can collect and process large amounts of personal data, raising concerns about privacy and data security.
Job Displacement - Automation driven by AI can lead to job losses in certain sectors.
Fairness - Ensuring that AI systems do not discriminate against individuals or groups based on protected characteristics.
Accountability - Establishing mechanisms to hold individuals and organizations accountable for the decisions made by AI systems.
Explainability - Developing AI models that are transparent and whose decisions can be easily understood and explained.
Safety - Ensuring that AI systems are safe and do not pose a risk to human health or well-being.

Future Trends

Explainable AI (XAI)	Developing AI models that are transparent and whose decisions can be easily understood and explained.
Federated Learning	Training machine learning models on decentralized data located on user devices or in data centers, without exchanging the data itself.
Generative AI	AI models that can generate new data instances that resemble the training data. Examples: Generative Adversarial Networks (GANs), Variational Autoencoders (VAEs).
Quantum Machine Learning	Combining quantum computing and machine learning to solve complex problems that are intractable for classical computers.
Edge AI	Running AI algorithms on edge devices, such as smartphones and IoT devices, to enable real-time processing and reduce latency.