

Artificial Intelligence (AI)



What is AI?

Artificial Intelligence (AI) refers to machines designed to simulate human intelligence. These systems can perform tasks such as understanding language, making decisions, and recognising patterns.

Types of AI:

Narrow AI (Weak AI): AI that is designed for a specific task, such as voice assistants (Siri, Alexa), recommendation systems, or self-driving cars.

General AI (Strong AI): AI that can perform any intellectual task a human can. General AI remains theoretical and is not yet developed.

Real-World Applications:

Healthcare: AI helps in diagnosing diseases and assisting in surgeries.

Finance: AI detects fraud and assists with financial management.

Transportation: Self-driving cars and smart traffic systems.

Customer Service: Chatbots provide instant support.

Entertainment: AI personalises movie recommendations and creates realistic video effects.

Ethical Implications of AI



Key Issues:

Bias in AI: AI can inherit bias from the data it is trained on, leading to unfair decisions. Fairness is essential to avoid discrimination.

Privacy Concerns: AI often requires access to personal information. Ensuring that data is used responsibly and with consent is crucial.

Job Displacement: As AI and automation take over routine tasks, human workers might lose jobs. It's important to consider how society will handle this shift.

Scenarios:

Fairness: Ensuring AI systems select candidates fairly in hiring processes.

Privacy: Protecting users' personal data when AI analyses their online behaviour.

Jobs: Automation in factories replacing human workers with robots, leading to layoffs.

Machine Learning (ML)



What is Machine Learning (ML)?

Machine Learning is a subset of AI where computers learn from data patterns rather than being explicitly programmed. ML systems improve and adapt over time through experience.

Types of Machine Learning:

Supervised Learning: The machine is trained with labelled data (e.g., images labelled as "cat" or "dog").

Unsupervised Learning: The machine is given unlabelled data and must identify patterns and groups on its own.

Reinforcement Learning: The machine learns through trial and error by receiving rewards for correct actions and penalties for wrong ones (e.g., a robot navigating a maze).

Importance of Data:

Quality: Accurate, relevant, and complete data improves model predictions.

Quantity: More data allows the model to detect complex patterns.

Diversity: Diverse data ensures the model works well in various situations.

Future of AI and Its Implications



Potential Future Developments in AI:

AI in Education: AI-powered tutors could offer personalised learning experiences.

AI in Healthcare: AI could assist doctors in diagnosing diseases more accurately and faster.

AI and Jobs: AI could automate simple tasks (e.g., street cleaning), allowing humans to focus on more complex and creative jobs.

Ethical and Social Considerations:

Setting Rules: It's important to establish guidelines for the safe and ethical use of AI to avoid harmful outcomes.

Handling Mistakes: AI systems can make errors in important decisions (e.g., loan approvals), so there should be processes in place to manage these mistakes.

Reflection:

As AI continues to evolve, it's vital to consider both the potential benefits and challenges it might bring to society. AI's impact on jobs, privacy, and decision-making will require careful planning.

Key Vocabulary

Key Word	Definition
Artificial Intelligence (AI)	The ability of machines or computers to perform tasks that usually require human intelligence, such as decision-making, problem-solving, and understanding language.
Narrow AI	A type of AI that is designed to perform a specific task or a limited range of tasks, such as recognising faces or recommending products, but cannot learn or perform tasks outside its scope.
General AI	A type of AI that has the ability to understand, learn, and apply intelligence across a wide range of tasks, similar to human intelligence. It can adapt and solve problems without being specifically programmed for each one.
Machine Learning (ML)	A type of artificial intelligence that allows computers to learn from data and improve their performance over time without being explicitly programmed for every task.
Neural Network	A computational model inspired by the human brain, consisting of layers of interconnected nodes (neurons) that process data and make decisions based on learned patterns.
Bias (in AI)	When an AI system unfairly favours certain groups or individuals over others due to skewed data or flawed algorithms, leading to unequal or discriminatory outcomes.
Ethical Implications	The moral considerations and consequences of using AI, such as how decisions are made, the potential for mistakes, and the impact on people's jobs and privacy.

Neural Networks and Deep Learning



What are Neural Networks?


Neural Networks are computational models inspired by the human brain, consisting of layers of interconnected nodes (neurons). They process data, identify patterns, and make decisions.

Structure of a Neural Network:

Input Layer: Receives raw data (like pixels in a photo).

Hidden Layers: Process data and identify patterns.

Output Layer: Provides the final decision (e.g., identifying whether an image is a cat or dog).



How Neural Networks Learn:

Forward Propagation: Data passes through the network, generating a prediction.

Loss Calculation: The network checks how accurate the prediction was.

Backward Propagation: The network adjusts based on errors to improve future predictions.

Deep Learning:

A form of machine learning using very complex neural networks with many hidden layers, allowing for deep analysis of large datasets (e.g., identifying different dog breeds from thousands of images).