Chapter 1: Introduction to AI and Everyday Examples

What is Intelligence?

Intelligence is a term that opens a room for discussions since it involves the ability to learn, think, understand, solve problems, and apply knowledge effectively. It is not just a cognitive skill, but a combination of many mental abilities. The level of intelligence differs in humans and animals in different ways based on how they perceive and behave. Therefore, it can be defined as the ability to use knowledge, interpretation, and insight to do any task successfully. It also includes the ability to adapt to new situations and make appropriate decisions based on experience.

Intelligence can be classified into several different types based on how individuals think, learn, and interact with the world. Some of them are explained here:

The Class Leader (Interpersonal Intelligence)

This is a student who naturally brings people together during a group project. When two classmates cannot agree on what to do, this person steps in and says something like, *"Tom, you're great at drawing, so you can make the poster. Jerry, you found the best information online, so you can write the report."* The student **understands** how others feel, communicate well, and helps the team work smoothly.

The Nature Lover (Naturalistic Intelligence)

This student is the one who, on a field trip to the garden, while everyone else is running around, huddled down over a long line of ants carrying a large leaf. He is aware of different types of plants and can even tell you which birds sing. The student is **smart** in the way they make sense of the natural world around them.

The Daydreamer (Intrapersonal Intelligence)

This student is self-aware. Before a big test, rather than panicking like everyone else, they might say to themselves, *I know I get distracted by noise so I'm going to put on my headphones. I also remember things better if I draw little pictures next to my notes.* He understands his own **feelings** and knows how is learn in the best positive manner.

What is Artificial Intelligence?

Artificial Intelligence (AI) is the field of computer science that makes machines intelligent, enabling machines to perform tasks that usually require human intelligence. These machines can conduct complex thinking processes such as data analysis, pattern recognition, prediction

of trends, solving problems, and decision-making by using data and learning from experience. AI systems are designed to learn, adapt, and improve their performance over a period of time.

Computing and **Intelligence** are the two important terms that help us understand the concepts of Artificial Intelligence.
Computing is all about using computers to conduct mathematical, logical, or relational calculations. Intelligence involves learning, reasoning, decision-making, and adapting to new situations. When these two concepts are combined together, then it forms an AI, which enables machines to simulate human intelligence and assist in decision-making.

History of Artificial Intelligence

Artificial intelligence has developed from simple rule-based programs into powerful systems capable of learning, reasoning, and creating. Here's a brief journey through its evolution:

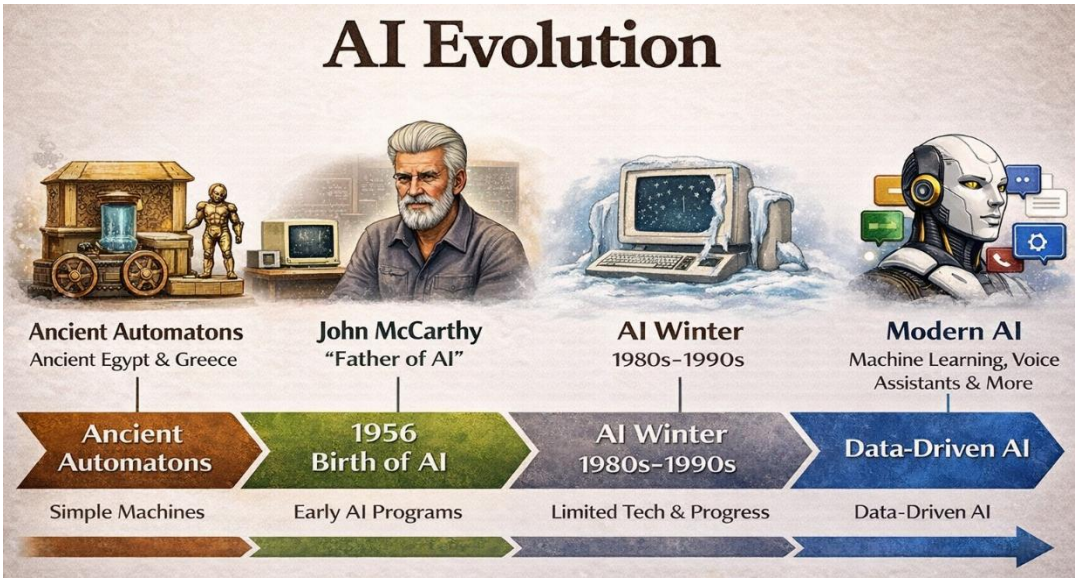

Ancient ideas of Intelligent Machines

Long before computers existed, people built simple machines called **automatons** which were created using water, gears, and levers. These simple machines could perform fixed tasks repeatedly but could not learn from previous experience or data. Although these machines were not truly intelligent, those days these simple machines helped shape the idea of making a machine that can **think** and make **decisions** in future.

One famous example is of the mechanical water clock used in ancient Egypt and Greece, which used the flow of water to operate levers and gears.

Alan Turing and Machine Intelligence

The theory of intelligent machines began in the 1940s, and one of the significant contributions was from the great mathematician, **Alan Turing**. He asked a question, "***Can machines think***?" which inspired many researchers to explore intelligent computing.

Turing further introduced the **Turing Test**, a method to calculate machine intelligence. This test would state whether a machine can perform in such a way that a person cannot tell whether they are talking to a machine or a human being. Alan Turing's ideas laid a strong foundation for Artificial Intelligence, that shaped the development of AI systems that are widely used today in applications such as search engines, voice assistants, and recommendation systems.

Birth of AI

In 1956, Dartmouth held the first Artificial Intelligence Summer Research Project. This was the first step toward AI. John McCarthy, a computer science professor at Stanford University who is known as the 'Father of Artificial Intelligence,' came up with the word **Artificial Intelligence** to describe the field of science and engineering that makes tools, especially computer programs, that can think for themselves. At this time, scientists started to make computers that could solve problems and play simple games. However, AI progress halted in the 1980s and 1990s due to slow processors, limited memory, and high cost. This period was referred to as the **AI Winter**.

Founding Fathers of Artificial Intelligence

1956 is the year associated with the birth of AI. Some of the most notable scientists who made important contributions include John McCarthy, Marvin Minsky, Claude Shannon, Herbert A. Simon, and Alan Newell. Their contribution includes their work in problem-solving, machine learning, and computational theory that laid the foundation for modern AI systems. These researchers are known to develop early AI programs and explore how machines could ignite human thinking and decision-making.

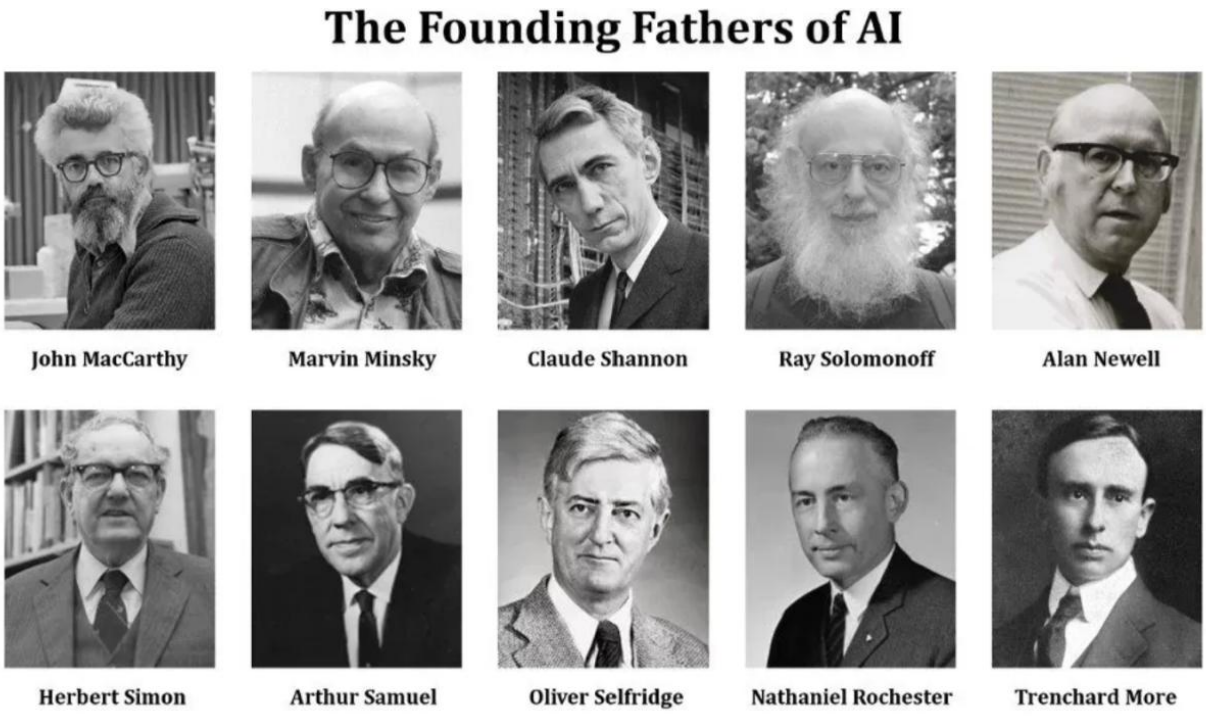

Modern Artificial Intelligence

The 2000s witnessed significant growth in technology, which made computers faster, cheaper, and more powerful. By now, AI systems could store and process a large amount of data. Instead of following only set rules, AI systems began to learn patterns from data and improve their performance over time. Modern AI systems can now recognize images, understand speech, and make simple decisions.

AI in our Daily Liver

AI has become an important part of our daily lives. AI technology is used in many ways, most of the time people don't even know it. AI helps make our tasks easier, faster, and more efficient by enabling machines and applications to learn, analyse data, and make smart decisions.

The following images illustrate a few real-life examples of AI in daily life.

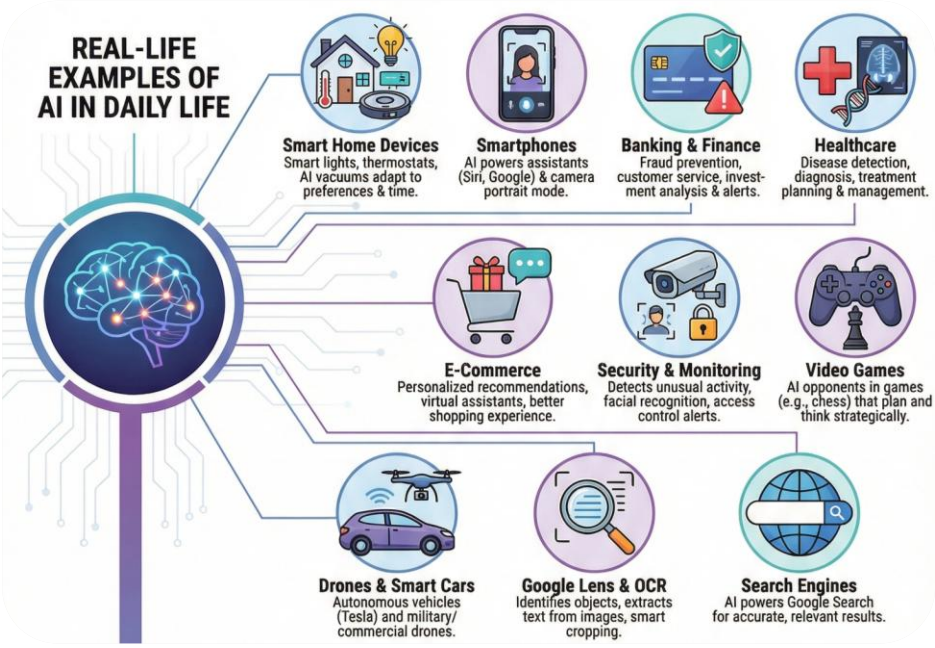

- **Smart Home Devices:** There are smart lights that change colour and intensity based on time, smart thermostats that change the temperature based on your preferences, AI-based vacuum cleaners, and much more. In the near future, AI will impact most of our interactions with smart home appliances.

- **Smartphones:** Whether you realise it or not, if you use a smartphone, you are interacting with AI. AI has immense impact on our lives, from the obvious features like the integrated smart assistants to the less obvious ones like the camera's portrait mode.
  For example, you may know that voice assistants, food apps, and reservation portals are based on AI. However, many people do not realize that AI also works behind features like portrait mode in smartphones. Have you ever wondered how your phone takes such impressive portrait photos? The answer is, surprisingly, AI.

- **Banking and Finance:** Artificial intelligence is heavily used in the banking and finance sector for tasks like customer service, fraud prevention, investment, and more. An easy illustration is the automated emails you get from banks whenever there is an unusual transaction. That is AI keeping an eye on your account and attempting to alert you to any fraud.

- **Healthcare:** AI-powered systems have made it much easier to find and diagnose diseases. AI also helps doctors plan treatments and take care of their patients. Hospitals and other healthcare facilities are therefore gradually using technologies that are powered by AI to ease everything from research to the detection of diseases.

- **E-Commerce:** AI in e-commerce makes the shopping experience more personalized and easier. It suggests products you may like, makes personal recommendations, and helps through virtual assistants and chatbots. Each of these features saves time, guides customers, and makes online buying simpler and more convenient.

- **Security and Monitoring:** AI helps make security systems faster and more efficient. It can detect unusual activities, recognise faces, and identify objects like unattended bags. It also sends alerts and helps control access through smart systems, allowing quick response to security situations.

- **Video Games:** AI is used in some video games to improve computer players' abilities to play. In games such as chess, the computer can play against humans by studying the moves and choosing the best ones. These games need planning and careful thinking. When a computer plays well, it shows that AI can ‘think’, ‘make decisions’, and ‘solve problems’ just like a human player.

- **Drones and Smart Cars:** The use of AI in smart cars and drones is one of the most noticeable examples of modern technology. Using a fully automatic self-driven car was once only a dream, but thanks to advancements made by the self-driven car companies, we now have a fleet of automatic vehicles on the road. Remember that successful drone programs are already being used by the defence forces across the world.

- **Smart Features in Scanning Apps:** AI-powered image recognition devices use advanced technology to identify objects captured by a camera, such as shoes, plants, or text, and provide useful information about them. With the help of Optical Character Recognition (OCR), users can easily extract text from scanned images of books or even hand-written documents. AI also improves features like smart cropping and edge detection in scanning apps.

- **Search Engines:** Finding information online is one of the most basic internet activities. Search engines help users quickly find the information they need. What makes search engines so effective is Artificial Intelligence? AI helps analyse search queries and provides the most relevant results. Whenever you search for information online, you are using AI technology.

Understanding what is not 'AI'

Artificial intelligence is used in a number of ways in our daily lives. But do all smart machines use AI? The answer is no! Some machines only follow fixed instructions. This brings us to the difference between **Automation** and **Artificial Intelligence**.

Automation machines do not think or learn. Instead, they just do what they are programmed to do and producing the same outcome every time. On the other hand, Artificial Intelligence machines can not only think and learn from data but also make decisions.

For example,

A microwave oven heats food for a set time, or a traffic light system changes at fixed intervals. They are **smart** in doing tasks, but they **do not think** or **learn**.

Automation	Artificial Intelligence
Works on fixed rules and present instructions	Works by learning from data
Does not think or learn	Can think, learn, and improve
Same input gives same output every time	Output may differ as it learns continuously
Cannot manage new situations	Can manage new and unfamiliar situations
Human intervention is needed for changes	Can adjust its behaviour on its own
Example: Traditional Washing Machine, Standard Traffic Signals	Example: Voice Assistance, Face recognition

Human Intelligence and Machine Intelligence

The ability of AI-powered machines to think and learn raises a question: how is the intelligence shown by these powerful machines different from human intelligence? Understanding this lets us compare human intelligence with machine intelligence.

Human Intelligence	Machine Intelligence
Human Intelligence develops naturally through biological evolution and experience.	Machine Intelligence is created using algorithms and learns patterns from data through training.
Learns from experience, senses, emotions, and reasoning.	Learns from data, algorithms, and pre-defined rules.
Highly adaptable to new and unpredictable situations.	Limited adaptability, depends on programming and data.
Capable of original thought, creativity, and imagination.	Relatively limited creativity, mostly follows patterns and data.
It can perceive and respond to emotions.	Lacks true emotional understanding, simulates responses.
Slower processing speed, prone to errors.	Faster processing speed, high accuracy in many tasks.
Has self-awareness and consciousness.	No consciousness or self-awareness.

How does AI learn?

AI systems become intelligent because they can learn from the data. Data means information such as pictures, text, numbers, or sounds. AI studies this data to find patterns and improve its performance. Data may be labelled or unlabelled.

- **Labelled data** means the data already has a **predefined label or tag** assigned to it.
- **Unlabelled data** means the data does not have any **labels or tags** assigned.

Let's say a teacher collects students' notebooks after a test in a class.

- If every book has the student's name written on it, the teacher knows whose book it is. This is like labelled data because each book already has the correct label (the student's name).
- If the books do not have names written on them, the teacher has to figure out whose book it is by looking at the handwriting or answers. This is like unlabelled data because there is no label to tell the correct owner.

What is Machine Learning?

The ability of AI to learn from data is called **Machine Learning** (ML). It is a branch of AI that enables machines to learn from data and improve their performance over time without being explicitly programmed for every task. In ML, algorithms are used to create models, which are trained using data. **Training** is the process by which the model learns patterns from data. By analysing large amounts of data, AI systems can make predictions, recognize patterns, and make decisions. Based on how data is used and how AI learns, Machine Learning is mainly divided into three types.

The three fundamental concepts of Machine Learning are:

- Supervised learning
- Unsupervised learning
- Reinforcement learning

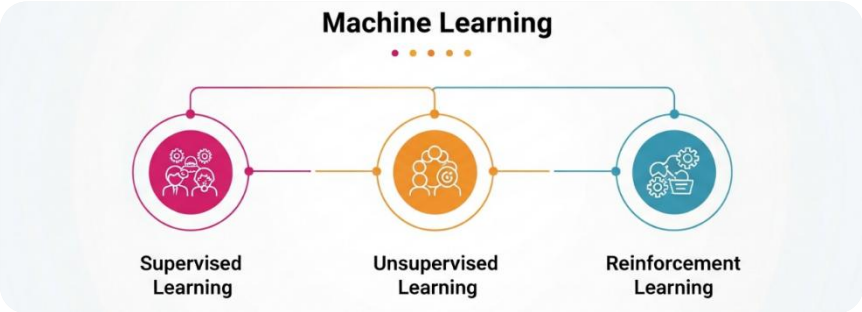

1. Supervised learning

Supervised learning is a type of machine learning where the machines learn from data that already has correct answers. The model learns the relationship between input data and their related output data, allowing it to make accurate predictions on new data. In this approach, the data provided to the model already includes the correct answers, enabling it to recognize patterns that lead to those answers. It is one of the simplest types of machine learning to implement and powers many modern applications.

Supervised Learning (Uses labelled data)

1. Email spam detection: Classifying emails as spam or not spam.
2. Image recognition: Identifying objects like cats, dogs, or cars in pictures.
3. House price prediction: Predicting house prices based on size, location, and features.
4. Handwritten digit recognition: Recognizing numbers written by hand (like in bank checks).
5. Medical diagnosis: Predicting diseases based on patient data and symptoms.

Key Concepts:

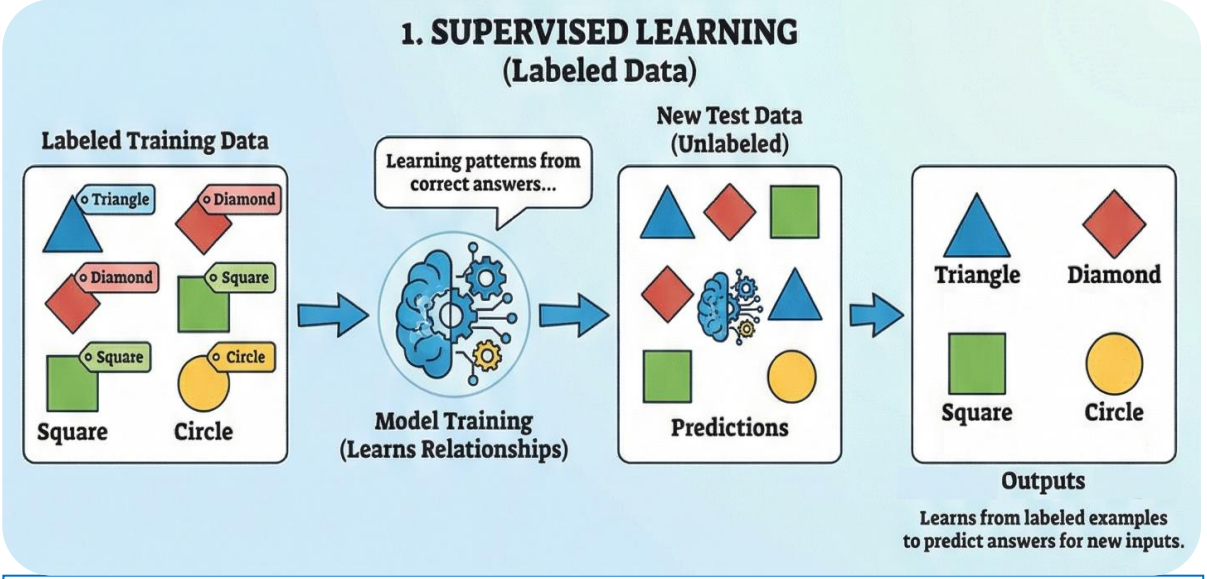

This diagram shows supervised learning. First, labelled data with different shapes (such as triangles, diamonds, squares, and circles) is used to train a model. Each shape has a correct label, helping the model learn how to tell them apart. After training, the model is given new data without labels. It uses what it has learned to identify the shapes and predict the correct labels. This shows how a trained model can classify new data based on patterns it has already learned.

2. Unsupervised learning

Unsupervised learning is a type of machine learning in which machines study data to identify hidden patterns, structures, or groupings. In this approach, the model learns from the patterns in data without any predefined labels or outputs.

Unlike supervised learning, unsupervised learning does not provide definite guidance or correct answers to the model. Instead, the model independently explores the data to discover meaningful patterns and relationships. It is commonly used for tasks such as clustering similar data points and detecting anomalies in datasets.

Unsupervised Learning (Uses unlabelled data)

1. Customer Segmentation: Grouping customers with similar buying habits.
2. Market Basket Analysis: Finding products that customers often buy together.
3. Document Clustering: Grouping similar articles or news topics.
4. Anomaly Detection: Detecting unusual transactions in banking or fraud detection.

5. Social Network Analysis: Finding communities or groups with similar interest

Key Concepts:

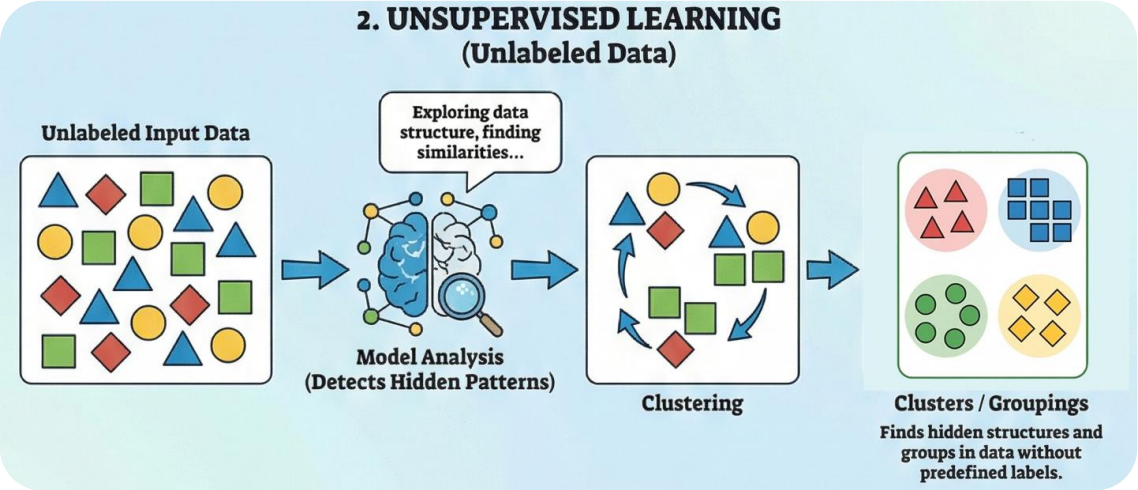

This diagram shows unsupervised learning. In this approach. In this method, the model is trained on unlabelled data, such as various shapes. Since there are no specified labels, the model learns by finding patterns and similarities in the data. It then organizes comparable objects into clusters. When new data is introduced, the model assigns it to the most appropriate group based on patterns it has learned. This method is widely used for clustering and learning hidden patterns in data.

3. Reinforcement learning

Reinforcement learning (RL) is based on trial and error. In this type of learning, machines learn by constantly searching for new and innovative approaches to solve a problem and improve the methods that maximize the defined reward. This type of learning is frequently used in robotics and gaming.

For example, a chess program based on RL learns by experimenting with several moves without knowing which one is the best. When it makes a successful move or wins a game, it receives a reward, and when it makes a wrong move or loses, it receives a penalty. By repeatedly playing and learning from rewards and penalties, the program gradually improves its strategy and learns how to play the game better.

Reinforcement Learning (Learning through rewards and penalties)

1. Game-playing AI: Learning to play chess, Go, or video games.
2. Self-driving cars: Learning safe driving decisions.
3. Robot navigation: Robots learning how to move through obstacles.
4. Recommendation Systems improvement: Improving suggestions based on user feedback.

Key Concepts:

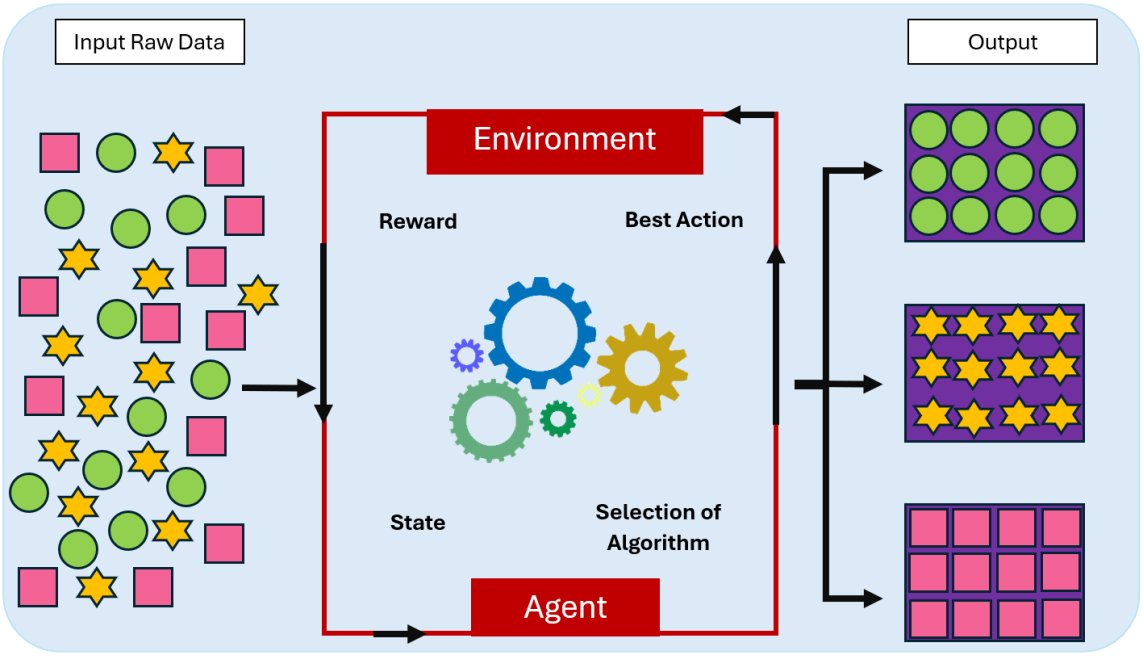

This diagram shows Reinforcement Learning, a form of machine learning, where a system learns by trial and error.

What happens is, first, the data system gets some data. Now, within this process we have two key components: the **agent** and the **environment**. The agent is the learner or decision-maker, and the environment holds everything in it.

The agent observes the current situation (referred to as the state) and takes an action. After this action, the environment will provide feedback. This feedback is either a reward (good action) or no reward or reduced reward (not so good action). The new situation is visible in the environment too.

The agent then modifies or improves its decision for the next time based on the reward it gets. As the agent continues to repeat this process over and over, it gradually learns which actions are most optimal. This process eventually allows it to make the right decision, such as correctly selecting or identifying a square, circle, triangle, or diamond.

Points to remember:

- ✓ Intelligence is the ability to learn, think, and solve problems.
- ✓ Not all automated machines use Artificial Intelligence.
- ✓ AI learns from data using Machine Learning.
- ✓ AI improves its performance over time through experience and feedback.
- ✓ There are three types of Machine Learning: Supervised, Unsupervised, and Reinforcement.

Exercise

A. Multiple Choice Questions.

1. What is labelled data?
   a) Data without any tags
   b) Data with predefined labels or tags
   c) Random data
   d) Incorrect data
2. Who introduced the Turing Test?
   a) John McCarthy
   b) Alan Turing
   c) Charles Babbage
   d) Isaac Newton
3. In supervised learning, data is:
   a) Unlabelled
   b) Random
   c) Labelled
   d) Deleted
4. Predicting exam marks based on the trained data uses:
   a) Supervised Learning
   b) Unsupervised Learning
   c) Reinforcement Learning
   d) All of the above
5. Which of the following is not AI?
   a) Voice assistant
   b) Face recognition
   c) Traditional traffic signal
   d) Smart chatbot

B. Fill in the blanks.

1. Intelligence includes the ability to learn and _____ problems.
2. AI is a technique to make intelligent_____.
3. Automation works on fixed _____ and preset instructions.
4. Unsupervised learning is mainly used for _____.
5. Machine Learning allows machines to learn from _____.

C. Short answer questions.

1. Define intelligence in your own words.
2. State two differences between automation and AI.
3. What is reinforcement learning?
4. Name the three types of Machine Learning.
5. Give two examples of AI used in daily life.

D. Think and apply.

Identify the Type of Machine Learning

Scenario	Type of Learning
Grouping customers based on shopping habits without label	
Predicting house prices based on past labelled data	
A game AI improving after winning or losing	
Sorting emails into spam and not spam	

E. Classify the following as AI or Automation.

Example	AI or Automation?
Traffic light changing at fixed intervals	
Face Recognition System	
Washing machine with preset timer	
Voice assistant answering questions