UNIT 1: Introduction: Artificial Intelligence for Everyone

Title: Introduction: AI for Everyone	**Approach**: Example-based learning**,** Hands-on activities, Discussion

Summary:

This unit covers various aspects of Artificial Intelligence (AI), including its definition, evolution, types, domains, terminologies, and applications. It explains the fundamental concepts of AI, such as supervised learning, natural language processing (NLP), computer vision etc. Additionally, it delves into machine learning (ML) and deep learning (DL) and discusses their differences, types, and applications. The content also outlines the benefits and limitations of AI, addressing concerns such as job displacement, ethical considerations, explainability, and data privacy.

Learning Objectives:

1. Understand the basic concepts and principles of Artificial Intelligence.
2. Explore the evolution of AI and identify the different types of AI.
3. Learn about the domains of AI, such as statistical data, natural language processing, and computer vision.
4. Understand the terminologies associated with AI, including machine learning, deep learning, and reinforcement learning.

Key Concepts:

1. What is Artificial Intelligence?
2. Evolution of AI
3. Types of AI
4. Domains of AI
5. AI Terminologies
6. Benefits and limitations of AI

Learning Outcomes:

Students will be able to -

1. Communicate effectively about AI concepts and applications in written and oral formats.
2. Describe the historical development of AI.
3. Differentiate between various types and domains of AI, including their applications.
4. Recognize the key terminologies and concepts related to machine learning and deep learning.
5. Formulate informed opinions on the potential benefits and limitations of AI in various contexts.

Pre-requisites: Reasonable fluency in English language and basic computer skills

What is Artificial Intelligence (AI)?

Artificial Intelligence (AI), has evolved drastically over the years, touching various aspects of our lives. It is a technology that has not only fascinated us but also significantly impacted how we live, work, and interact with the world around us. Within the vast landscape of AI, there exist several distinct Domains of Artificial Intelligence, each with its unique characteristics and applications. According to Statista, the global AI market, with a value of billion 113.60 GBP in 2023, is on a continuous growth trajectory, primarily fueled by substantial investments.

Artificial intelligence (AI) refers to the ability of a machine to learn patterns and make predictions.



In its simplest form, Artificial Intelligence is a field that combines computer science and robust datasets to enable problem-solving. AI does not replace human decisions; instead, AI adds value to human judgment. Think of AI as a smart helper that can understand things, learn from examples, and do tasks on its own without needing to be told exactly what to do each time. For example, AI can:

- **Understand Language:** AI can understand and respond to what you say, like virtual assistants such as Siri or Alexa.
- **Recognize Images:** AI can look at pictures and recognize what is in them, like identifying animals in photos.
- **Make Predictions:** AI can analyze data to make predictions, like predicting the weather or suggesting what movie you might like to watch next.
- **Play Games:** AI can play games and learn to get better at them, like playing chess or video games.
- **Drive Cars:** AI can help cars drive themselves by sensing the road and making decisions to stay safe.

What is not AI?

When we talk about machines, not all of them are considered Artificial Intelligence (AI). Here are some examples:

- **Traditional Rule-Based Systems**: These machines follow set rules without learning from data.
- **Simple Automation Tools**: Basic tools like timers or calculators do specific tasks but do not think or learn.
- **Mechanical Devices**: Machines like pulleys or gears work based on physics but do not learn or think.

- **Fixed-Function Hardware**: Devices like microwave ovens perform tasks without learning or thinking.
- **Non-Interactive Systems**: Machines that do not change based on new information, like a basic electric fan.
- **Basic Sensors**: Sensors collect data but do not analyze or understand it.

Artificial Intelligence machines are different. They learn from data and can make decisions on their own. For example, a smart washing machine can adjust its settings based on what it is washing. AI goes beyond just following rules; it can learn, adapt, and make decisions based on data and context.

Evolution of AI

The history of AI can be traced back to ancient times, with philosophical discussions about the nature of intelligence and the possibility of creating artificial beings. However, the modern era of AI began in the mid-20th century with significant developments and milestones:

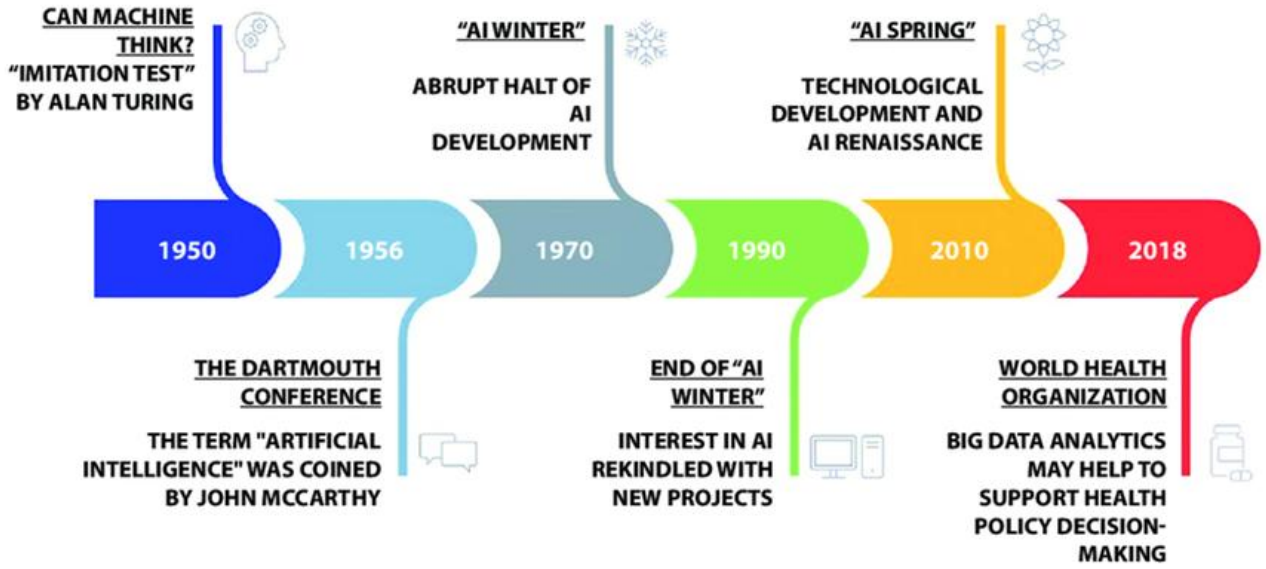

Source:https://www.researchgate.net/figure/Timeline-diagram-showing-the-history-of-artificial-intelligence_fig1_364826401

Time Period	Key Events and Developments
1950	1950 was a landmark year for the question of machine intelligence because of Alan Turing's famous paper "Computing Machinery and Intelligence." In this paper, Turing proposed a thought experiment called the "imitation game" (later known as the Turing test).
1956	The Dartmouth Conference was organized by McCarthy that marked the birthplace of AI as a field. The term "Artificial Intelligence" was coined by John McCarthy. McCarthy, along with Turing, Minsky, and Simon, laid the foundation for AI.
1960-1970	Significant progress in AI research that led to the development of expert systems, early neural networks, exploration of symbolic reasoning, and problem-solving techniques.

Time Period	Key Events and Developments
1980-1990	Mixed optimism and skepticism about AI with breakthroughs in machine learning, and neural networks led to "AI winter".
21st Century	Resurgence of interest and progress in AI with advancements in computing power, data availability, and algorithmic innovation. Also, there were breakthroughs in machine learning, deep learning, and reinforcement learning. That led to transformative applications of AI in healthcare, finance, transportation, and entertainment.

Types of AI

Computer scientists have identified three levels of AI based on predicted growth in its ability to analyze data and make predictions.

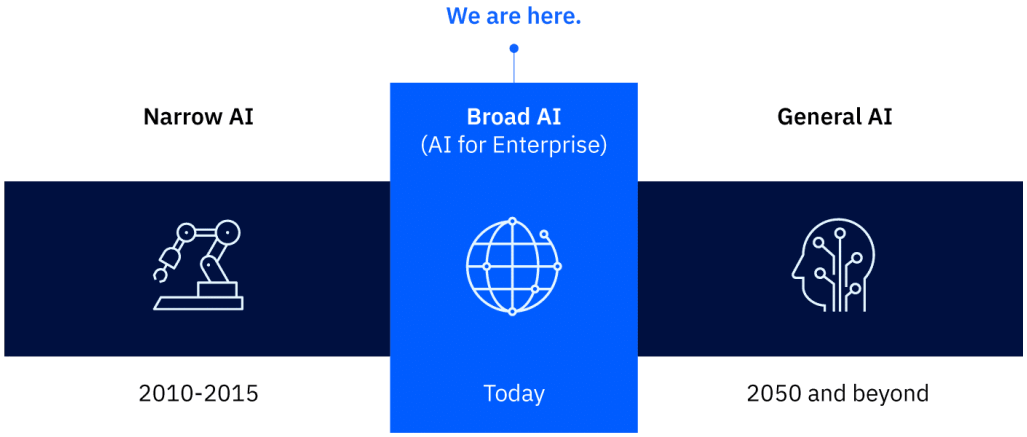

1. Narrow AI:
   - Focuses on single tasks like predicting purchases or planning schedules.
   - Rapidly growing in consumer applications, such as voice-based shopping and virtual assistants like Siri.
   - Capable of handling specific tasks effectively, but lacks broader understanding.
2. Broad AI:
   - Acts as a midpoint between Narrow and General AI.
   - More versatile than Narrow AI, capable of handling a wider range of related tasks.
   - Often used in businesses to integrate AI into specific processes, requiring domain-specific knowledge and data.
3. General AI:
   - Refers to machines that can perform any intellectual task a human can.
   - Currently, AI lacks abstract thinking, strategizing, and creativity like humans.
   - Artificial Superintelligence (ASI) may emerge, potentially leading to self-aware machines, but this is far from current capabilities.

Types of Data

Data might be facts, statistics, opinions, or any kind of content that is recorded in some format. This could include voices, photos, names, and even dance moves! It surrounds us and shapes our experiences, decisions, and interactions.

It is often referred to as the "**new oil**" of the 21st century. Did you know? 90% of the world's data has been created in just the last 2 years, compared to the previous 6 million years of human existence.

Data can be classified into three different types which are as follows:

- Structured Data
- Unstructured Data
- Semi-structured Data

Structured data is like a neatly arranged table, with rows and columns that make it easy to understand and work with. It includes information such as names, dates, addresses, and stock prices. Because of its organized nature, it is straightforward to analyze and manipulate, making it a preferred format for many data-related tasks.

On the other hand, **unstructured data** lacks any specific organization, making it more challenging to analyze compared to structured data. Examples of unstructured data include images, text documents, customer comments, and song lyrics. Since unstructured data does not follow a predefined format, extracting meaningful insights from it requires specialized tools and techniques.

Semi-structured data falls somewhere between structured and unstructured data. While not as organized as structured data, it is easier to handle than unstructured data. Semi-structured data uses metadata to identify certain characteristics and organize data into fields, allowing for some level of organization and analysis. An example of semi-structured data is a social media video with hashtags used for categorization, blending structured elements like hashtags with unstructured content like the video itself.

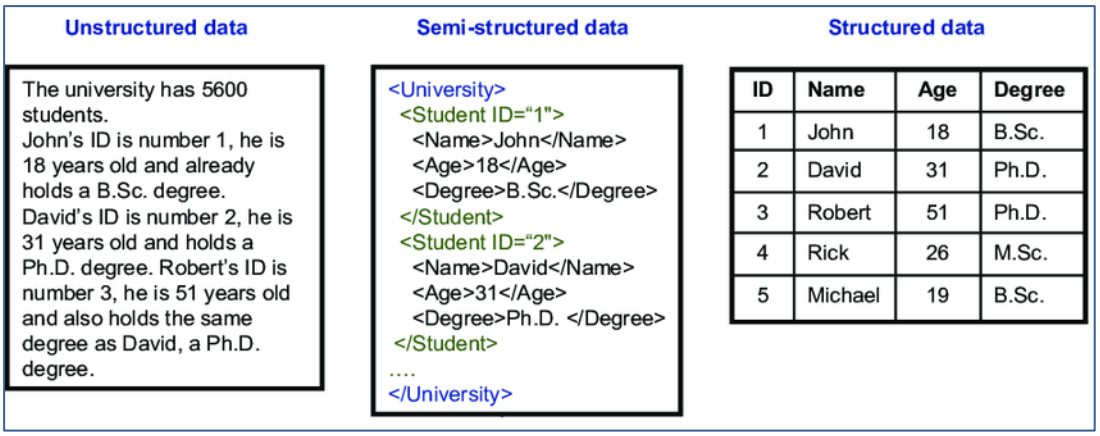

Source: https://www.researchgate.net/figure/Unstructured-semi-structured-and-structured-data_fig4_236860222

Domains of AI

Artificial Intelligence (AI) encompasses various fields, each focusing on different aspects of replicating human intelligence and performing tasks traditionally requiring human intellect. These fields are classified based on the type of data input they handle:

a) **Statistical Data**: Statistical data includes numerical, categorical, and alphanumeric inputs. This domain involves the collection, analysis, and interpretation of large datasets to extract insights and patterns using statistical methods, machine learning algorithms, and data visualization techniques. Examples include:

- **Search recommendations and Google Maps history**, which are based on your previous interactions and location data.
- **Amazon's personalized recommendations**, which are influenced by your shopping habits and browsing history.
- **Social media activity, cloud storage, and digital textbooks**, all of which generate and store vast amounts of data that can be analysed for insights.

b) **Natural Language Processing (NLP):** NLP focuses on processing text and speech inputs to enable computers to understand, interpret, and generate human language. It involves tasks such as language translation, sentiment analysis, text summarization, and speech recognition, facilitating communication between humans and machines through natural language interfaces.

The goal of an NLP-Trained model is to be capable of "understanding" the contents of documents, including the slangs, sarcasm, inner meaning, and contextual definitions of the language in which the text was written.

Differences Between NLP, NLU, and NLG:

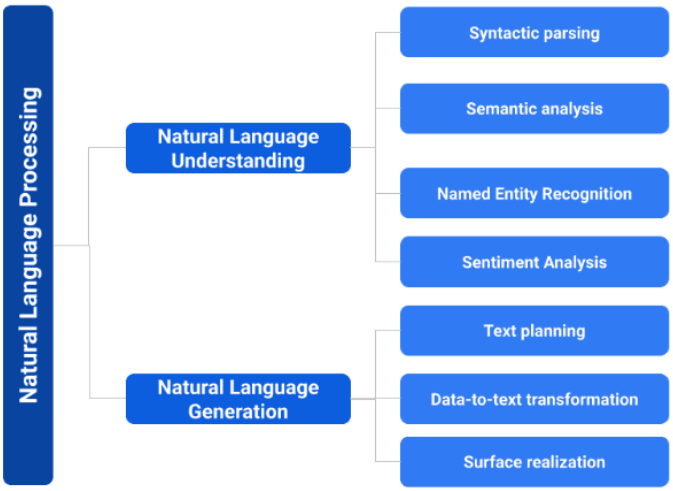

Source: https://www.baeldung.com/cs/natural-language-processing-understanding-generation

Natural Language Processing (NLP): This is the broad umbrella term encompassing everything related to how computers interact with human language. Think of it as the "what" - what computers can do with human language. *It is like the whole library - filled with different tools and techniques for working with language data.*

Natural Language Understanding (NLU): This is a subfield of NLP that focuses on understanding the meaning of human language. It analyzes text and speech, extracting information, intent, and sentiment. NLU helps computers understand the language and what it means. *Imagine finding a specific book in the library.*

Natural Language Generation (NLG): This is another subfield of NLP, but instead of understanding, it focuses on generating human language. It takes structured data as input and turns it into coherent and readable text or speech. *Think of this as writing a new book based on the information gathered in the library.*

c) **Computer Vision:** Computer Vision deals with visual data inputs, primarily images and videos. It enables computers to interpret and understand visual information, and perform tasks such as object detection, image classification, facial recognition, and scene understanding, enabling applications such as autonomous vehicles, medical imaging, and augmented reality.

Computer Vision is like giving computers the ability to see and understand the world through digital images and videos, much like how humans use their eyes to perceive their surroundings. In this domain, computers analyze visual information from images and videos to recognize objects, understand scenes, and make decisions based on what they "see."

When we take a digital image, it is essentially a grid of tiny coloured dots called **pixels**. Each pixel represents a tiny portion of the image and contains information about its colour and intensity.

Resolution is expressed as the total number of pixels along the width and height of the image. For example, an image with a resolution of 1920x1080 pixels has 1920 pixels horizontally and 1080 pixels vertically. Higher resolution images have more pixels, providing more detail.

Now, here's where AI comes in. To make sense of these images, computers convert them into numbers. They break down the image into a series of numbers that represent the colour and intensity of each pixel. This numerical representation allows AI algorithms to process the image mathematically and extract meaningful information from it. For instance, AI algorithms might learn to recognize patterns in these numbers that correspond to specific objects, like cars or faces. By analyzing large amounts of labeled image data, AI systems can "learn" to identify objects accurately.

Activity:

Divide the students into groups and provide them with a list of real-world applications without specifying which domain each application belongs to. Ask each group to categorize the applications into the three domains: Statistical Data, Natural Language Processing (NLP), and Computer Vision.

1. *Gesture recognition for human-computer interaction*
2. *Chatbots for customer service*
3. *Spam email detection*
4. *Autonomous drones for surveillance*
5. *Google Translate*

6. *Fraud detection in financial transactions*
7. *Augmented reality applications (e.g., Snapchat filters)*
8. *Sports analytics for performance optimization*
9. *Object detection in autonomous vehicles*
10. *Recommendation systems for e-commerce platforms*
11. *Customer segmentation for targeted marketing*
12. *Text summarization for news articles*
13. *Automated subtitles for videos*
14. *Medical image diagnosis*
15. *Stock prediction*

Statistical Data	Natural Language Processing	Computer Vision

AI Terminologies

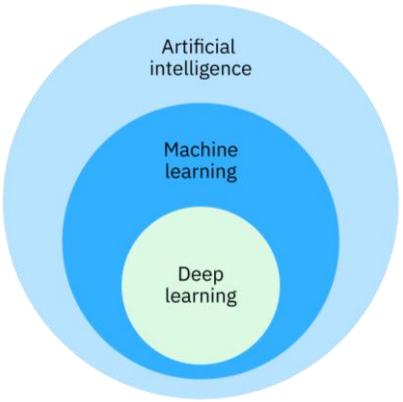

- **Artificial intelligence** machines don't think. They calculate. They represent some of the newest, most sophisticated calculating machines in human history. It is a computer system that can perform tasks that ordinarily require human intelligence or human interference.
- Some can perform what is called **machine learning** as they acquire new data. Machine learning is a subset of artificial intelligence (AI) that focuses on developing algorithms and models that enable computers to learn from data and make predictions or decisions without being explicitly programmed.
- Others, using calculations arranged in ways inspired by neurons in the human brain, can even perform **deep learning** with multiple levels of calculations. Deep learning is an AI function that imitates the working of the human brain in processing data and creating patterns for use in decision making.
    - The structure of Deep Learning is inspired by the structure of the neurons and neuron connections in the human brain.

- Neural networks, also known as Artificial Neural Networks (ANNs), are a subset of Machine Learning and the core heart and concept of Machine Learning.
- They comprise of node layers, containing an input layer, one or multiple hidden layers, and an output layer.
- If the output of any node is above a specified threshold, that node is activated, sending data to the next layer of the network.
- Otherwise, no data is passed along to the next layer of the network.
- If the number of Layers including the Input and Output Layer is more than three, then it is called a Deep Neural Network.

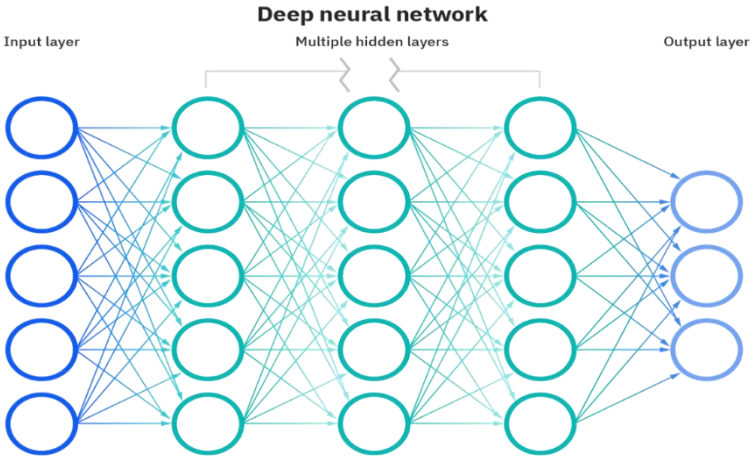

MACHINE LEARNING	DEEP LEARNING
1. Works on small dataset for accuracy	1. Works on Large dataset
2. Dependent on Low-end machine	2. Heavily dependent on high-end machine
3. Divides the tasks into sub-tasks, solves them individually and finally combine the results	3. Solves problem end to end
4. Takes less time to train	4. Takes longer time to train
5. Testing time may increase	5. Less time to test the data

Example: Imagine you are given the job to sort items in the electronic items department at an electronic store. You realize that there are dozens of products and very less time to sort them manually. How will you use artificial intelligence, machine learning, and deep learning to help with your work?

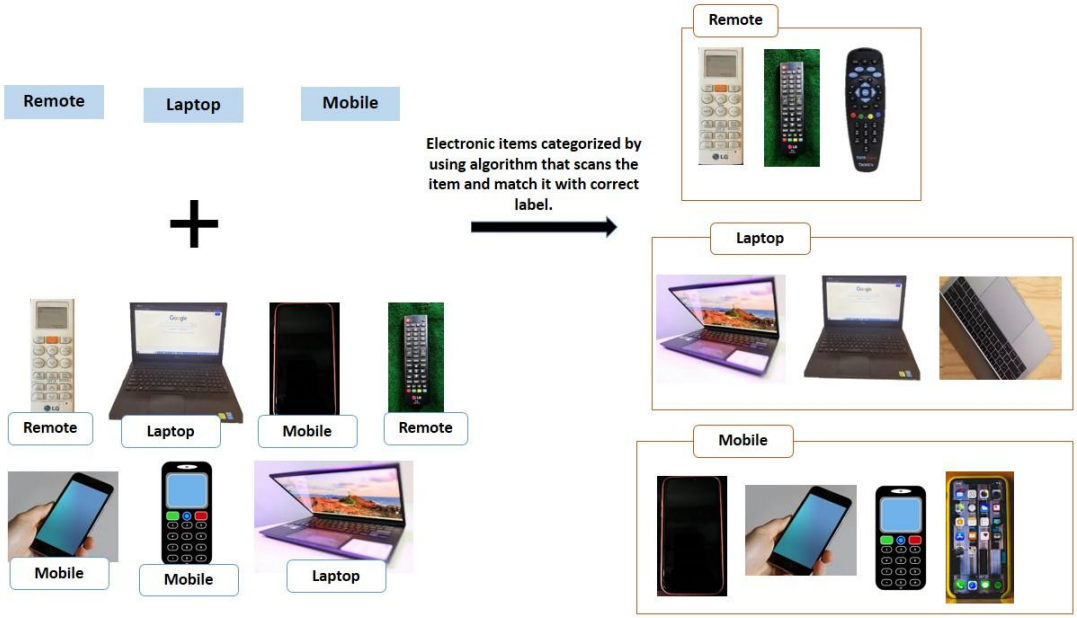

By training the model with labeled data, AI can match the visual characteristics of the electronic items with the correct category, significantly speeding up the sorting process and reducing human error.

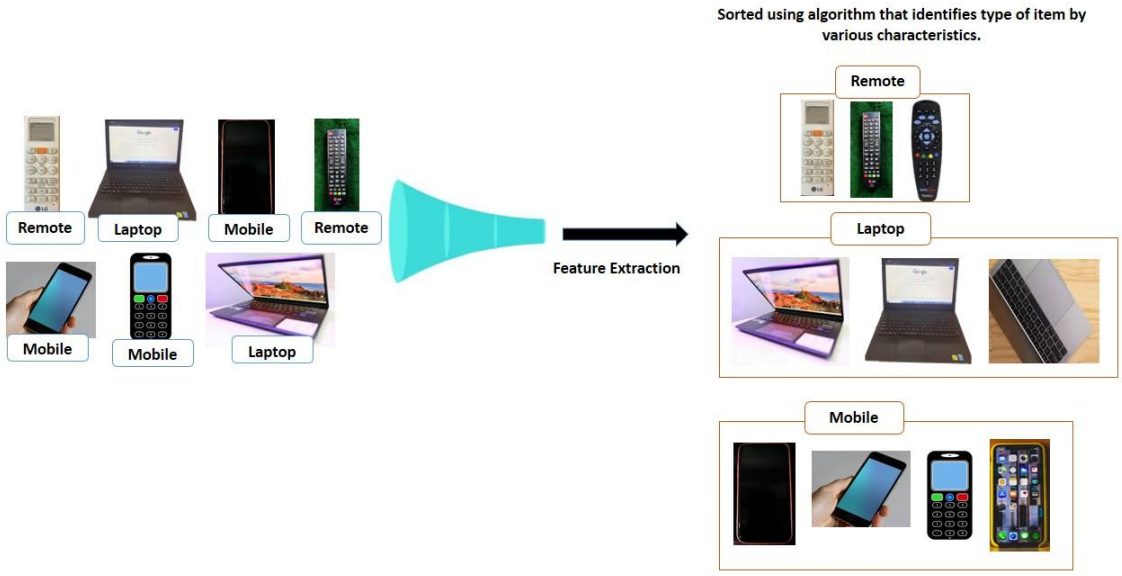

To improve the performance of the machine, you expose it to more data to ensure that the machine is trained on numerous characteristics of each type of electronic items, such as size, shape, and colour. The more data you provide for the algorithm, the better the model gets. By providing more data and adjusting parameters, the machine minimizes errors by repetitive guess work.

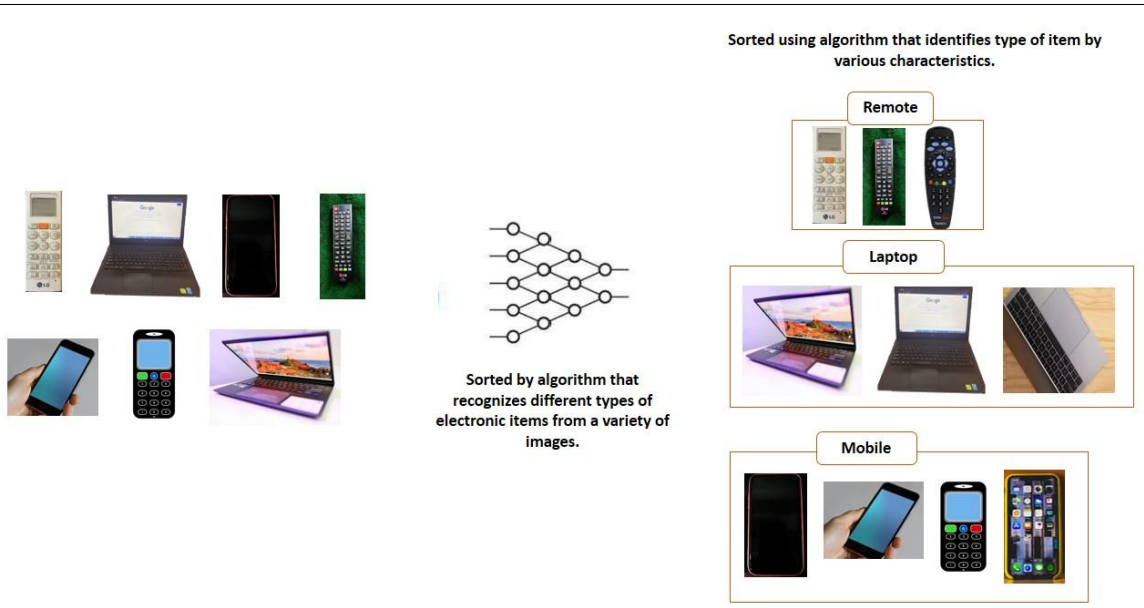

Deep learning models eliminate the need for feature extractions. Decide the algorithms based on deep learning to sort electronic items by removing the need to define what each product looks like. Feature extraction is built into the process without human input. Once you have provided the deep learning model with dozens of electronic items pictures, it processes the images through different layers of neural networks. The layers can then learn an implicit representation of the raw data on their own.

Types of Machine Learning

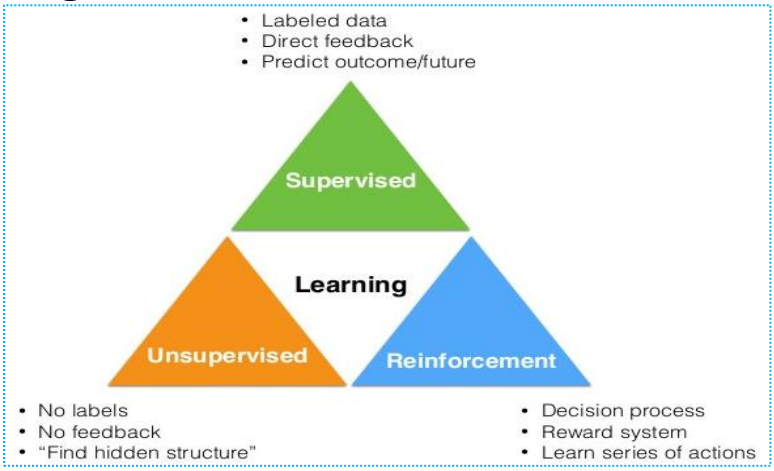

Supervised learning

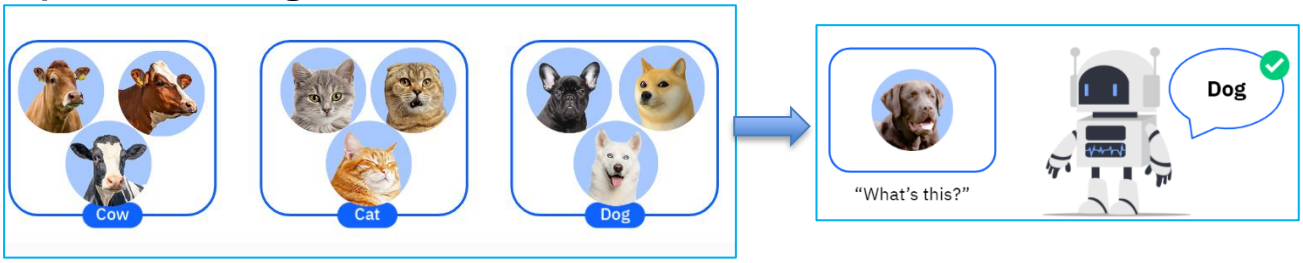

- Supervised learning is a type of machine learning where the model learns from labelled data, which means that the input data is accompanied by the correct output.

- In supervised learning, the algorithm learns to map input data to output labels based on example input-output pairs provided during the training phase.
- The goal of supervised learning is to learn a mapping function from input variables to output variables so that the model can make predictions on unseen data.
- Examples of supervised learning algorithms include linear regression, logistic regression, decision trees, support vector machines (SVM), and neural networks.

Unsupervised Learning:

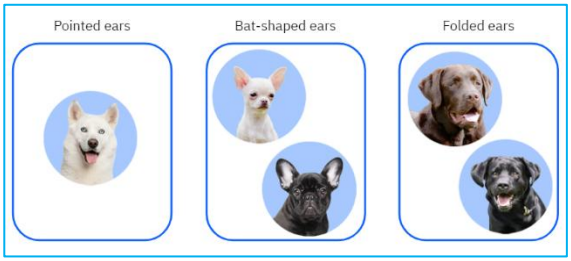

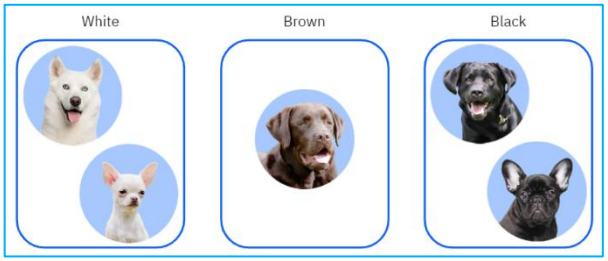

- Unsupervised learning is a type of machine learning where the model learns from unlabelled data, which means that the input data is not accompanied by the correct output.
- In unsupervised learning, the algorithm tries to find hidden patterns or structure in the input data without explicit guidance.
- The goal of unsupervised learning is to explore and discover inherent structures or relationships within the data, such as clusters, associations, or anomalies.
- Examples of unsupervised learning algorithms include k-means clustering, hierarchical clustering, principal component analysis (PCA), and autoencoders.

Reinforcement Learning:

1.

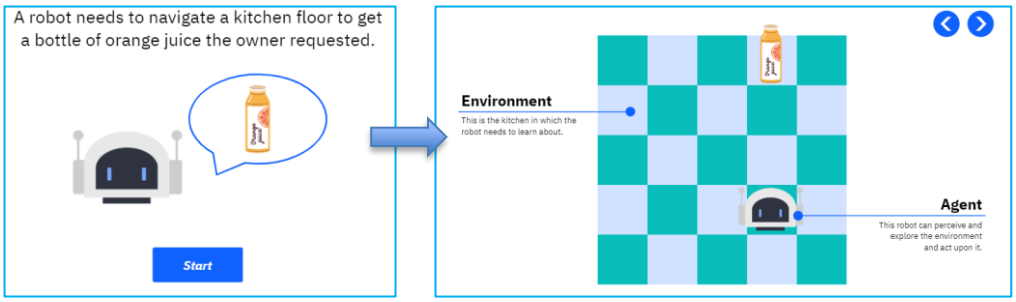

2.

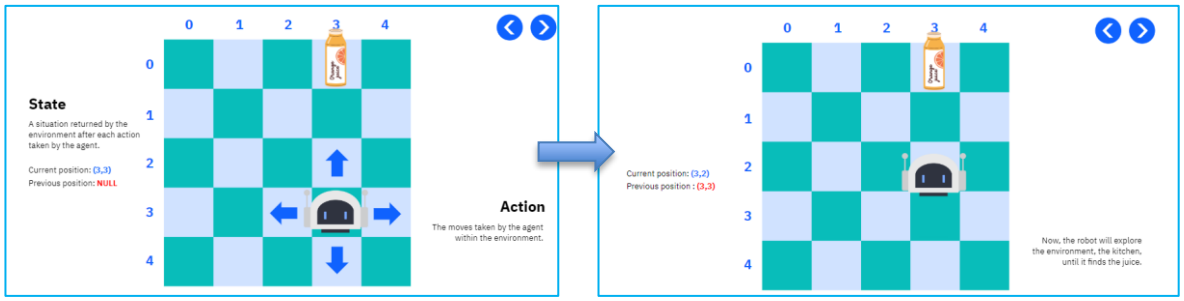

3. 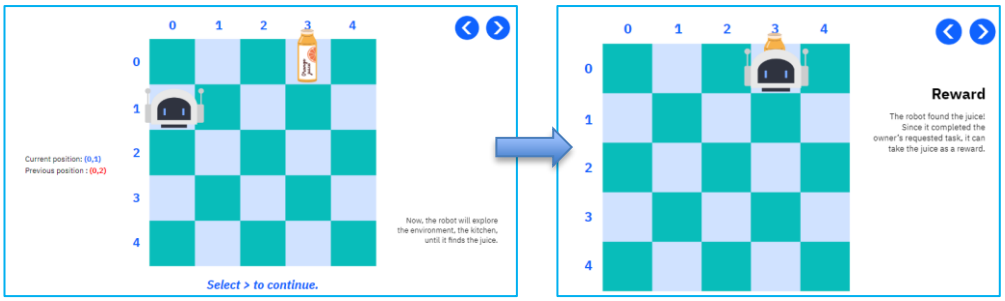

4. 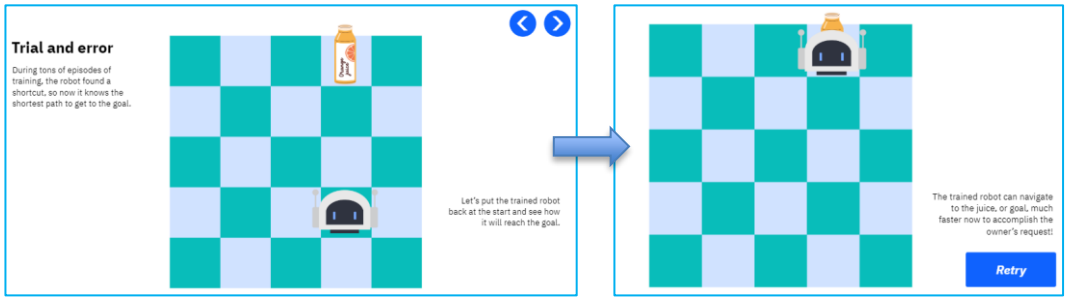

- Reinforcement learning is a type of machine learning where an agent learns to make decisions by *interacting with an environment* to maximize cumulative rewards.
- In reinforcement learning, the agent learns through *trial and error* by taking actions and receiving feedback from the environment in the form of rewards or penalties.
- The goal of reinforcement learning is to learn a policy or strategy that guides the agent to take actions that lead to the highest cumulative reward over time.
- Reinforcement learning is commonly used in scenarios where the agent must make a sequence of decisions over time, such as playing games, controlling robots, or managing financial portfolios.
- Examples of reinforcement learning algorithms include Q-learning, deep Q-networks (DQN), policy gradients, and actor-critic methods.

Benefits and limitations of AI

Benefits:

1. **Increased efficiency and productivity:** AI automates tasks, analyzes data faster, and optimizes processes, leading to increased efficiency and productivity across various sectors.
2. **Improved decision-making:** AI analyzes vast amounts of data and identifies patterns that humans might miss, assisting in data-driven decision-making and potentially leading to better outcomes.
3. **Enhanced innovation and creativity:** AI tools can generate new ideas, explore possibilities, and automate repetitive tasks, freeing up human resources for more creative pursuits and innovation.
4. **Progress in science and healthcare:** AI aids in drug discovery, medical diagnosis, and personalized medicine, contributing to advancements in healthcare and scientific research.

Limitations:

1. **Job displacement:** Automation through AI raises concerns about job displacement and the need for workforce retraining and upskilling.
2. **Ethical considerations:** Concerns exist around bias in AI algorithms, potential misuse for surveillance or manipulation, and the need for ethical guidelines and regulations.
3. **Lack of explainability:** Some AI models, particularly complex ones, lack transparency in their decision-making, making it difficult to understand how they arrive at their outputs.
4. **Data privacy and security:** Large-scale data collection and use for AI development raise concerns about data privacy and security vulnerabilities.

- ***Earn a credential on IBM Skills Build on the topic Artificial Intelligence Fundamentals using the link****:*
  https://students.yourlearning.ibm.com/activity/PLAN-CC702B39D429
- *Semantris*, is an NLP-Based game by Google based on Word association powered by semantic search.
  https://experiments.withgoogle.com/semantris
- This is a game built with machine learning. We **draw**, and a neural network tries to guess what you're **drawing**.
  https://quickdraw.withgoogle.com/
- *The experiment based on the computer vision domain of AI. It identifies that you draw and suggests the related images. To play the game, visit the following link on any computing device with speakers.*
  https://www.autodraw.com/

Extension Activities:

These activities provide opportunities for students to explore various aspects of artificial intelligence, develop critical thinking skills, and engage in hands-on learning experiences in the classroom.

1. **AI in the News:** Have students research recent news articles or stories related to artificial intelligence? They can explore topics such as AI advancements, ethical dilemmas, or AI applications in various industries. Students can then present their findings to the class and facilitate discussions on the implications of these developments.
2. **AI Applications Showcase:** Divide students into small groups and assign each group a specific AI application or technology (e.g., virtual assistants, self-driving cars, healthcare diagnostics). Ask students to research and create presentations or posters showcasing how their assigned AI technology works, its benefits, potential drawbacks, and real-world examples of its use.

3. **AI Coding Projects:** Introduce students to basic coding concepts and tools used in AI development, such as Python programming language and machine learning libraries like TensorFlow or scikit-learn. Guide students through hands-on coding projects where they can build simple AI models, such as image classifiers or chatbots. Encourage experimentation and creativity in designing and training their AI systems.
4. **AI Film Analysis:** Screen and analyze films or documentaries that explore themes related to artificial intelligence, such as "Ex Machina," "Her" "I, Robot," or "The Social Dilemma." After watching the films, facilitate discussions on how AI is portrayed, its potential impact on society, and ethical considerations raised in the narratives.

EXERCISE

A. Multiple-choice questions (MCQs):

1. Who is often referred to as the "Father of AI"?
    a. Alan Turing
    b. John McCarthy
    c. Marvin Minsky
    d. Herbert A. Simon
2. In which year was the term "Artificial Intelligence" first used by John McCarthy?
    a. 1930
    b. 1955
    c. 1970
    d. 2000
3. What does the term "Data is the new oil" imply?
    a. Data is as valuable as oil.
    b. Data is used as fuel for machines.
    c. Data is a non-renewable resource.
    d. Data and oil are unrelated.
4. Divya was learning neural networks. She understood that there were three layers in a neural network. Help her identify the layer that does processing in the neural network.
    a. Output layer
    b. Hidden layer
    c. Input layer
    d. Data layer
5. Which category of machine learning occurs in the presence of a supervisor or teacher?
    a. Unsupervised Learning
    b. Reinforcement Learning
    c. Supervised Learning
    d. Deep Learning

6. What does Deep Learning primarily rely on to mimic the human brain?
    a. Traditional Programming
    b. Artificial Neural Networks
    c. Machine Learning Algorithms
    d. Random Decision Making
7. What is the role of reinforcement learning in machine learning?
    a. Creating rules automatically
    b. Recognizing patterns in untagged data
    c. Rewarding desired behaviors and/or penalizing undesirable ones
    d. Mimicking human conversation through voice or text
8. Which AI application is responsible for automatically separating emails into "Spam" and "Not Spam" categories?
    a. Gmail
    b. YouTube
    c. Flipkart
    d. Watson

B. Fill in the Blanks:

1. To determine if a machine or application is AI-based, consider its ability to perform tasks that typically require _____ intelligence.
2. Artificial intelligence (AI) enables a machine to carry out cognitive tasks typically performed by _____.
3. Supervised, unsupervised, and reinforcement learning are three categories of _____.
4. _____ is a subset of artificial intelligence that is entirely based on artificial neural networks.
5. Machine learning can be used for online fraud detection to make cyberspace a _____ place.

C. True or False:

1. Chatbots like Alexa and Siri are examples of virtual assistants.
2. Supervised learning involves training a computer system without labeled input data.
3. Unstructured data can be easily analyzed using traditional relational database techniques.
4. Deep learning typically requires less time to train compared to machine learning.
5. Machine learning is not used in everyday applications like virtual personal assistants and fraud detection.

D. Short Answer Questions:

1. How is machine learning related to AI?
2. Define machine learning.
3. What do you mean by Reinforcement Learning? Write any two applications of Reinforcement Learning at School.
4. How do you understand whether a machine/application is AI based or not? Explain with the help of an example.

E. Long Answer Questions:

1. Define Data. Briefly explain the types of data.
2. What is deep learning, and how does it differ from traditional machine learning?

F. Case-study/Application Oriented Questions:

1. A hospital implemented an AI system to assist doctors in diagnosing diseases based on medical images such as X-rays and MRI scans. However, some patients expressed concerns about the accuracy and reliability of the AI diagnoses. How can the hospital address these concerns?