Unit 1 - Intro to AI

Unit 1
Intro to AI

1.1 Unit Introduction

Welcome to the first unit of the AI Fundamentals course!

In this unit you will start immersing yourself in the world of **artificial intelligence**.

You will learn:

- the different categories that compose AI and their applications
- how AI works and the learning paradigms

You can't wait to get started, right?

Note that real humans worked on this course, it has not been generated by AI. Maybe it helped a bit with the phrasing of some sentences, but not much.

Continue to 1.2: What is AI?

1.2 What is AI?

Artificial intelligence or AI is everywhere right now: from content creation to virtual assistants. Let's not forget personalized recommendations and automated customer service. The list could go on for hours!

But what exactly is AI?

AI is a **technology** that **enables computers to imitate human intelligence**, allowing them to perform tasks like **reasoning**, **learning**, and **problem-solving**, similar to how humans do.

Do not worry! Computers don't have an actual brain or think like humans, they simply **use math, statistics, and logic to process information and**

make decisions.

AI systems work by taking in a lot of information, learning from it, and then using what they learned to **guess** what might happen in the future.

Believe it or not, AI isn't new! It's been around for a long time and has evolved significantly, especially in recent years.

Let's take a look at the history of AI.

Continue to 1.2.1: History of AI

1950s

Turing test

Alan Turing laid the **foundation for AI** by proposing that machines could mimic human intelligence and introduced the **Turing test** to see if a **machine's behavior could be mistaken for a human's**.

> Fun fact: This test was finally passed by a machine in 2013!

1956

Field of study

The term **artificial intelligence** was coined by John McCarthy, marking the official birth of AI as a field of study.

1980s

Machine Learning

The introduction of **machine learning** allows computers to learn from data and improve over time.

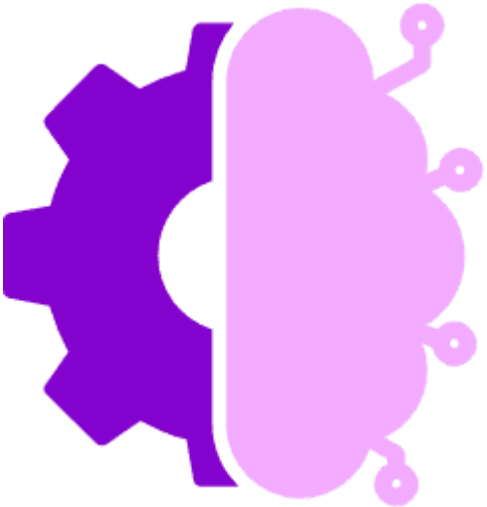

1997

Deep Blue

IBM's Deep Blue, a **chess playing computer**, beats the world chess champion Garry Kasparov.

2010s

Deep Learning

This decade saw major advances in **deep learning**, where AI uses multiple layers to recognize complex patterns, for example enabling self-driving cars to understand their surroundings.

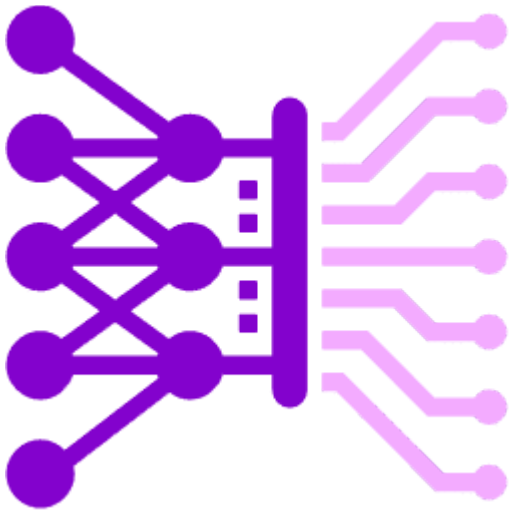

2017

Transformer model

The **Transformer** model is introduced, changing how computers understand and work with human language by capturing the relationships between distant words.

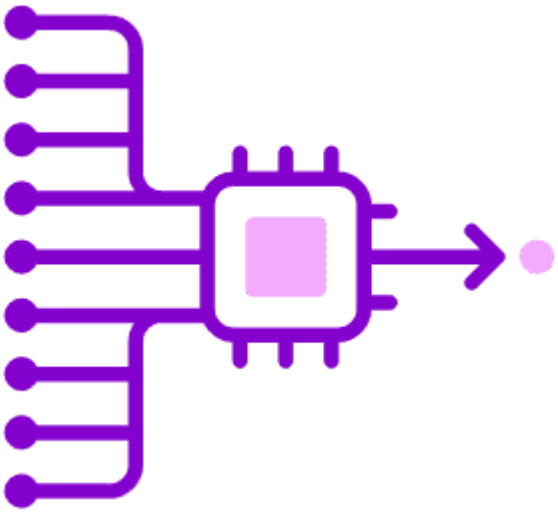

2020s

GPTs

OpenAI introduced the Generative Pre-trained Transformer (GPT) series, enhancing how computers understand and generate human language.

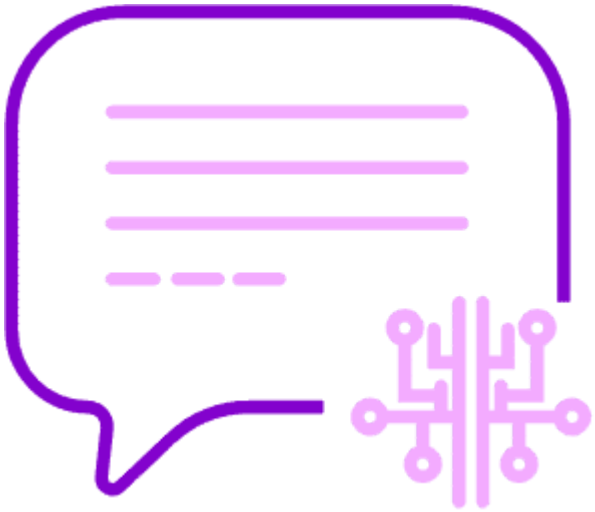

Continue to 1.2.2: Categories of AI

1.2.2 Categories of AI

AI has six main categories, each focusing on specific tasks and techniques. They can work both separately and together to create more advanced systems.

These categories are:

1. Machine Learning
2. Natural Language Processing
3. Expert Systems
4. Computer Vision
5. Speech Recognition
6. Robotics

Imagine AI as a school called 'AI will always love you' with different classes, each dedicated to a unique skill.

Let’s dive deeper into each category.

1: Machine Learning

1

Machine Learning (ML)

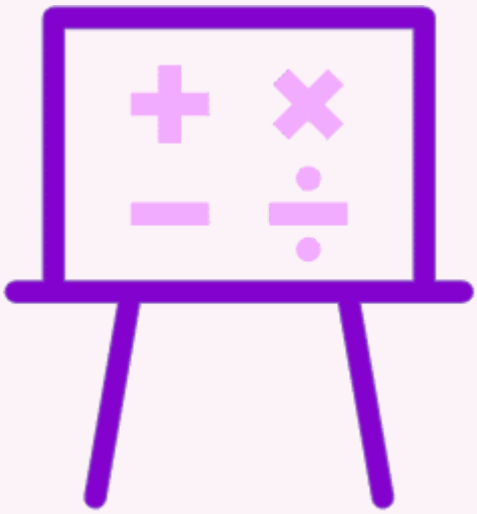

Machine Learning (ML) is a part of AI **that helps computers learn from data, find patterns**, and improve over time without needing humans to tell them what to do. By looking at lots of existing information, ML algorithms can spot trends, make choices, and improve their skills based on new examples.

It is like the math class where students learn to solve problems by practicing with examples.

NLP is focused on understanding and working with human language, while generative AI (used by ChatGPT for example) specifically creates new content, like text or images.

Click + for more information.

Real world application: _

Recommendations on what to watch on Netflix.

2: Natural Language Processing

2

Natural Language Processing (NLP)

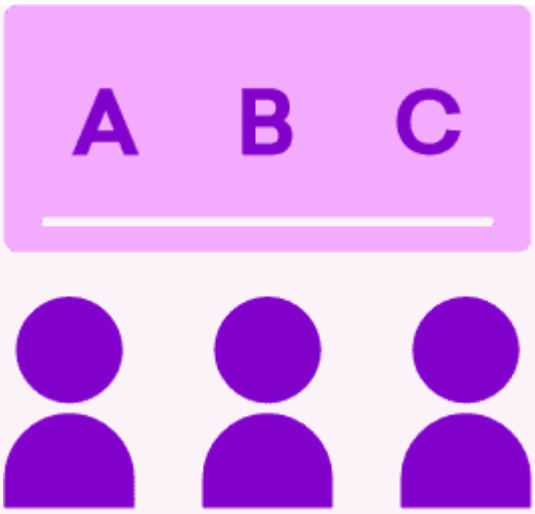

Natural Language Processing (NLP) is a part of AI that helps computers **understand, interpret, and create human language**. It lets machines **process and analyze text or speech like people do**, making it easier for them to work with language.

NLP is the language class, where students learn how to read, write, and talk.

Click + for more information.

Real world application: _

Autocorrect and autocomplete text.

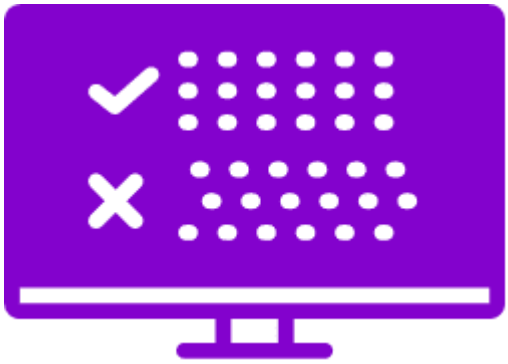

3: Expert Systems

3

Expert Systems

An **Expert System** is a computer program that **mimics how a human expert makes decisions in a specific area**. It uses **set rules to analyze information and draw conclusions**, applying knowledge from that field to help solve problems and provide advice.

It is similar to a science class where students apply specific rules and knowledge from their field to solve problems.

Click + for more information.

Real world application: _

Diagnosis of bacterial infections.

4

Computer Vision

Computer Vision is a field of artificial intelligence that **enables computers to interpret and understand visual information** from the world, such as images and videos, allowing them to perform tasks like object recognition, image classification, and scene understanding.

It is like the art class that helps students learn to interpret images and recognize patterns in visuals.

Click + for more information.

Real world application:

Facial recognition technology.

5: Speech Recognition

5

Speech Recognition

Speech Recognition is a technology that focuses on **converting spoken language into written text**, allowing machines to understand and transcribe audio input, such as voice commands and conversations. This technology **facilitates human-computer interaction** by enabling users to communicate with machines through their voices.

It is like the music class where students learn to interpret sounds.

Click + for more information.

Real world application: _

Virtual assistants.

6: Robotics

6

Robotics

Robotics is a part of AI that **builds machines to interact with the real world**, using other AI tools like machine learning and computer vision to help them work and make decisions on their own.

It is like the PE class, where students learn through action and physical tasks.

Click + for more information.

Real world application: _

Autonomous vehicles such as the Mars Rover.

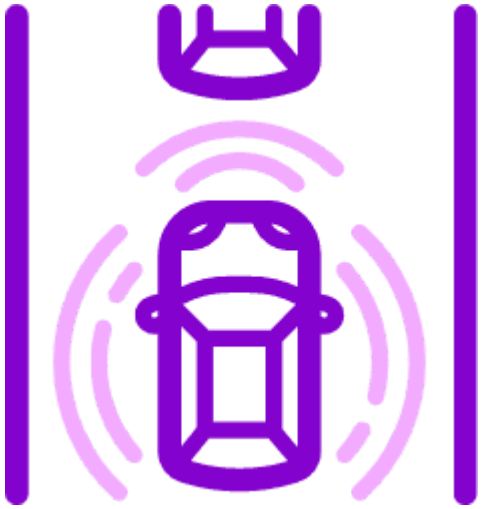

Continue to 1.2.3: Machine learning vs Deep learning

1.2.3 Machine learning vs Deep learning

Machine learning

Deep learning

Machine learning is the most widely used category of AI because it forms the **foundation for many AI applications**. You will focus on this

Deep learning is a part of machine learning that uses **layered networks** (similar to neurons in the brain) to learn from large amounts of

over the next few lessons.

raw data, which helps AI function like the human brain.

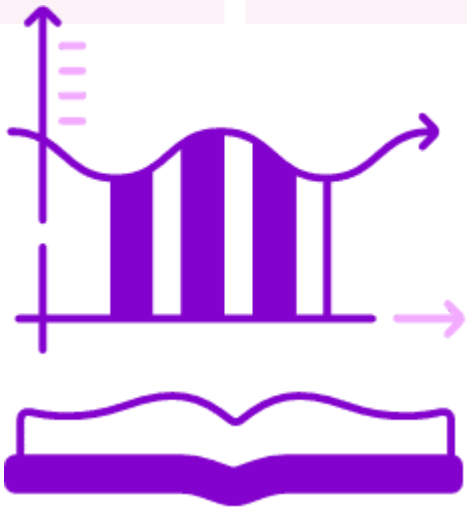

Deep learning is like an advanced math class where students dive deep into complex subjects, using layered understanding to master intricate topics and spot patterns.

Continue to 1.3: Learning paradigms

Machine learning involves training an algorithm to learn patterns from data by providing examples and making predictions on new data.

🔑 Key Concept

An **algorithm** is a set of instructions or rules.

There are two different phases:

Training phase

During training, the algorithm analyzes the data to **identify patterns**. From these patterns it develops a **mathematical model** that represents the **relationship between this data**.

Prediction phase

Once trained, the mathematical model is used to **make predictions** on new data. It applies the **patterns** learned during training to interpret or predict outcomes from new information.

🔑 Key Concept

A **mathematical model** is a way to represent real-world things or problems using numbers and formulas to make predictions or understand how things work.

Continue to 1.3.1: Types of learning paradigms

1.3.1 Types of learning paradigms

Machine learning includes three main types of learning, known as learning paradigms. Each paradigm offers a different method for how machines learn from data in terms of input provided and possible outputs.

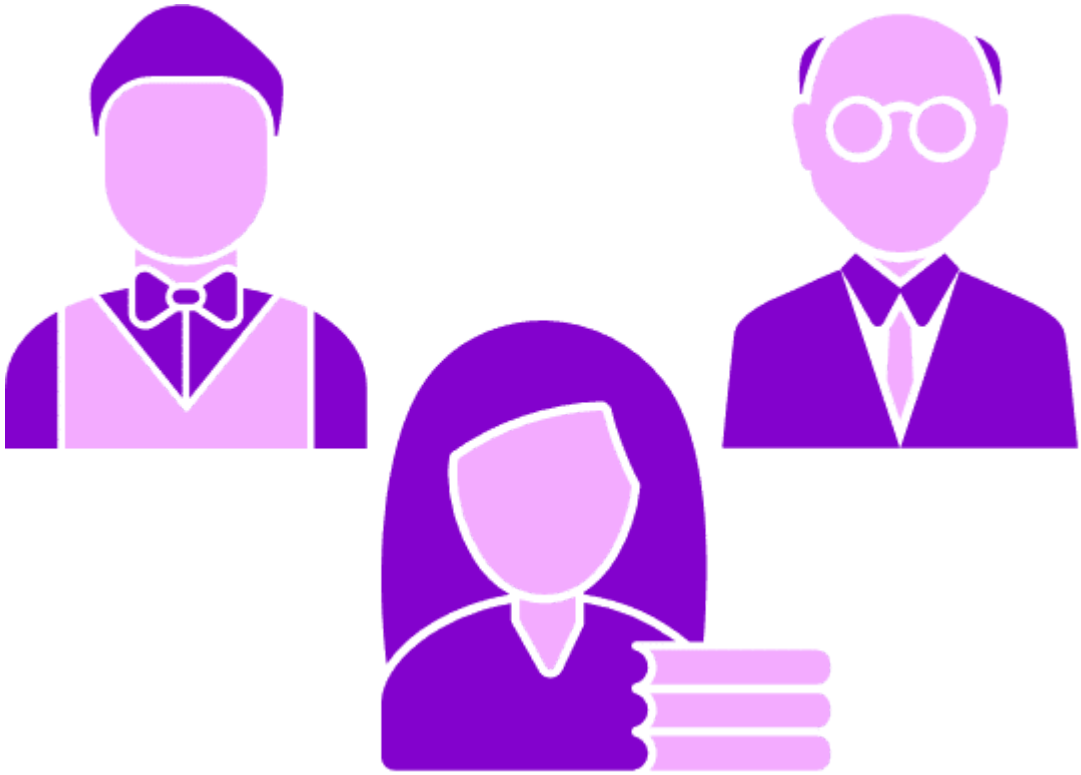

It is similar to three teachers who each have their own unique teaching style.

They are:

1. **supervised** learning
2. **unsupervised** learning
3. **reinforcement** learning

Let's have a look at each one in more detail.

Supervised learning

1

In supervised learning paradigms the inputs are labeled data.

Think of them as examples with the correct answers. The model learns by analyzing these labeled examples to **understand the relationships between inputs and outputs**, allowing it to make accurate predictions on new, unseen data.

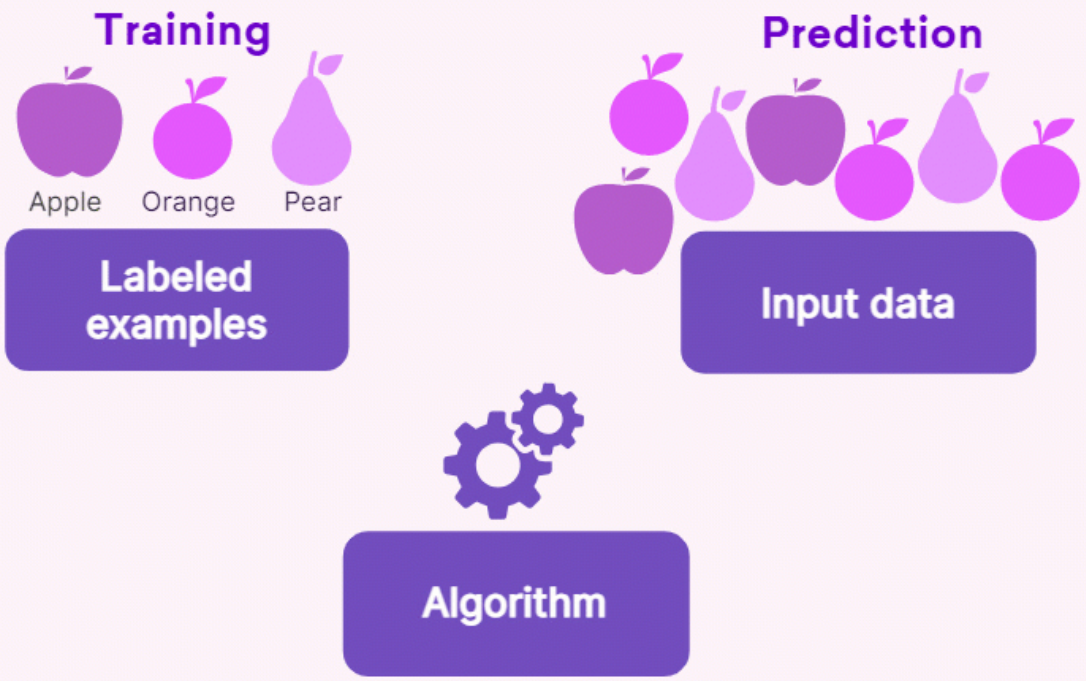

For example, if you want your machine learning model to recognize different kinds of **fruit**, you would provide labeled examples of **apples**, **pears**, and **oranges**. The model uses these examples to learn the differences and identify patterns associated with each fruit, e.g. shape, color, texture, size. When a new image is given to the algorithm, it can predict which of the three fruits it is.

It's like Sarah the Teacher giving her students examples to help them understand and learn to classify different concepts.

Here are some examples in which supervised learning is used.

Spam detection

Classifying emails as "spam" or "not spam" based on labeled examples of previous emails.

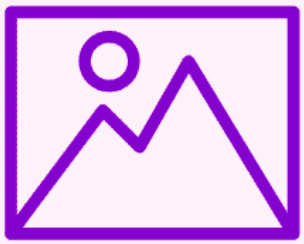

Image recognition

Identifying objects in images, such as recognizing cats and dogs, using labeled images for training.

Credit scoring

Predicting whether a loan applicant is expected to fail to repay the loan based on historical data of previous applicants with known outcomes.

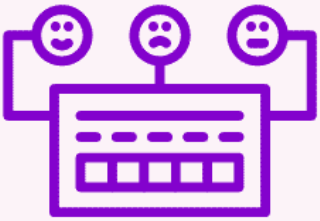

Medical diagnosis

Assisting doctors by predicting diseases based on patient data and previous diagnoses.

Sentiment analysis

Determining whether a piece of text (like a review) is positive, negative, or neutral using labeled examples of texts.

2: Unsupervised learning

Unsupervised learning

In unsupervised learning paradigms the inputs are not labeled, the model doesn't know any information about them. Instead, it finds patterns by grouping similar items on its own. This approach is useful when you have data but don't know what you're looking for, allowing the model to discover hidden connections or relationships within the information.

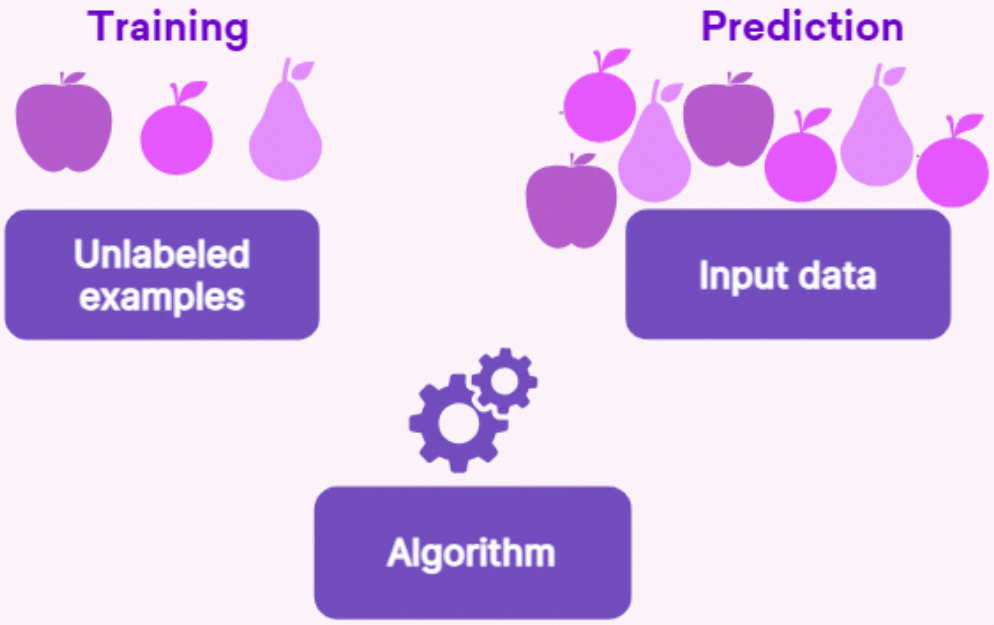

For example, if you give the model various images of fruits, it will find relationships and sort them into groups based on the fruit characteristics. When you show it a new image, it uses what it learned to place the fruit into one of the groups.

It's like Ulysses the Teacher providing his students with a variety of materials and asking them to explore and find patterns on their own, encouraging them to group similar items without any guidance.

Here are some examples in which unsupervised learning is used.

Customer segmentation

Market basket analysis

Anomaly detection

Grouping customers based on purchasing behavior and preferences without predefined labels.

Identifying items frequently purchased together to optimize product placements and promotions.

Detecting unusual patterns in data, such as fraud detection in financial transactions.

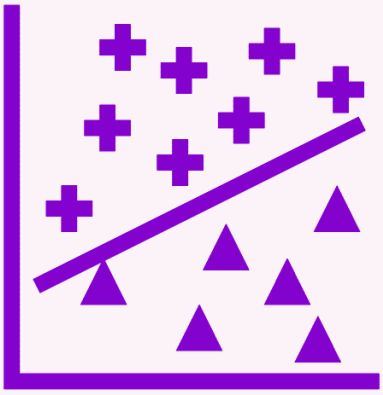

Document clustering

Organizing a collection of documents into topics or themes without prior categorization.

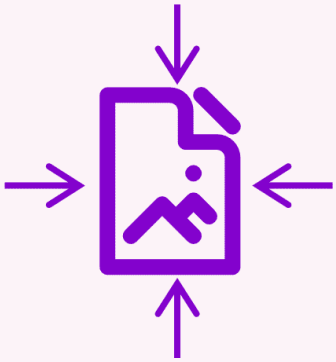

Image compression

Reducing the size of images by grouping similar pixel values to maintain quality while saving space.

3: Reinforcement learning

Reinforcement learning

In reinforcement learning paradigms, the algorithm (called agent in this case) learns to make decisions by interacting with an environment and getting feedback as rewards or penalties.

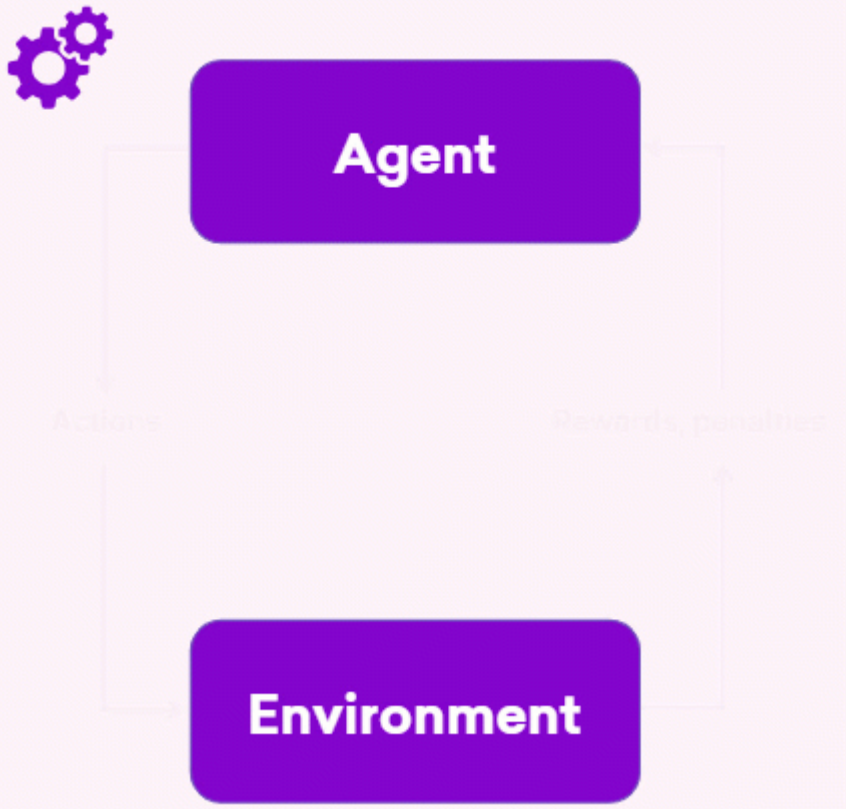

This method uses **trial and error** to get better over time as the agent tries different options to find the best outcomes. This is especially helpful in complicated situations where the best choices aren't clear.

> Imagine a robot learning to navigate a maze. It receives a reward for reaching the exit and a penalty for hitting walls. Over time, the robot figures out the best path to take by trying different routes and learning from its mistakes.

It's like Ralph the Teacher guiding his students to try different activities and learn from their successes and mistakes, helping them discover the best strategies for achieving their goals.

Here are some examples in which reinforcement learning is used.

Game playing

Training AI agents to play video games, where the agent learns strategies by playing against itself and receiving rewards for winning.

Robotics

Teaching robots to perform tasks, such as walking or picking up objects, by rewarding them for successful actions and guiding them through trial and error.

Autonomous vehicles

Training self-driving cars to navigate by rewarding them for safe driving decisions and penalizing risky behaviors.

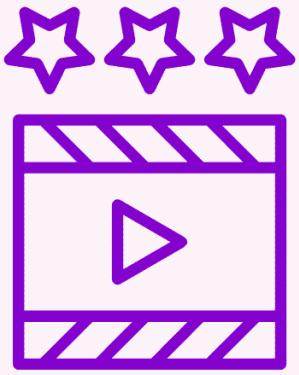

Personalized recommendations

Improving recommendation systems for movies or products by adjusting suggestions based on user interactions and feedback.

Healthcare treatment plans

Developing treatment strategies for patients by using past treatment outcomes to reward

effective actions and
improve future decisions.

Continue to the wrap up for this unit

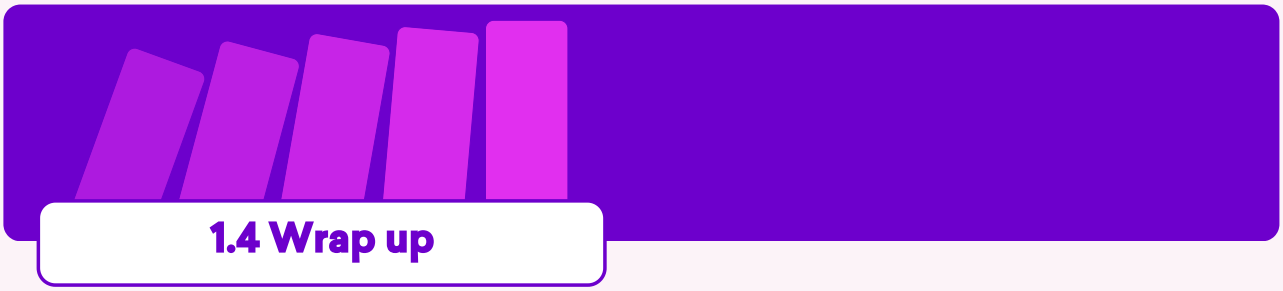

1.4 Wrap up

1

Artificial intelligence (AI) is a technology that gives computers the ability to **imitate human intelligence**, allowing them to perform tasks such as reasoning, learning, problem-solving solving and decision-making.

2

AI has several categories, each with specialized techniques for handling specific tasks. These include **machine learning** (teaching computers to learn from data), **natural language processing** (working with human language), **computer vision** (interpreting images and video), robotics (controlling machines), and **expert systems** (rule-based reasoning).

3

AI works in two steps: **training** and **prediction**. In the training step, the system learns from examples and creates rules that describe the relationships in the data. In the prediction step, it uses those rules to give answers or make predictions on new information.

Unit complete!

Well done!

You have completed the first unit!

In the next unit you will learn more about Generative AI.

make | academy

Mark this task complete to continue to the next unit.