Machine Learning

Introduction to the Conceptual Framework of Artificial Intelligence (AI)

Author: Zsolt Lengyel, iTStudy Hungary Training and Research Centre

Reviewer: Mária Hartyányi, iTStudy Hungary Training and Research Centre

Learning Goal

The rapid rise of artificial intelligence has triggered extreme reactions in people, ranging from boundless enthusiasm to the fear that it will soon take our jobs. The purpose of this learning material is to review the theoretical background needed for a realistic assessment of what AI actually is, and to clarify the basic concepts essential for informed AI use (such as machine learning, training dataset, the model, etc.).

By the end of the learning material, we will find the answers to the following questions:

- How does a computer learn and communicate?
- What do we mean by the concept of machine learning?
- What does it mean that artificial intelligence systems do not operate deterministically?
- Is “artificial” intelligence a machine version of human intelligence?
- Can artificial intelligence be smarter than humans?

Introduction

The famous mathematician Alan Turing[1] put it this way in a 1947 lecture he gave to London Mathematical Society:

“What we want is a machine that can learn from experience.”

However, the computers available at the time did not make it possible to build such a program, and it took more than seventy years to get closer to Turing's vision: the creation of machines that can learn.

Following the 1956 Dartmouth conference, where it became an independent field of research, the mathematical models and algorithms needed for machine learning were gradually developed. At the same time, the technological conditions also fell into place: huge volumes of digital data became available, often referred to as “big data,” and high-performance computers, servers, storage systems, and processors needed to process them emerged as well. By the early 21st century, artificial intelligence already played a significant role in industry and IT, but it was only ChatGPT's public debut that brought it widespread social recognition.

November 30, 2022 is a notable day in the history of artificial intelligence. On that day, ChatGPT was launched, and something happened that had never happened before: the number of registered users reached one million within

[1] During the Second World War, he was a key member of the Bletchley Park codebreaking team, which cracked the Nazis' notorious encryption code, the Enigma.

five days — compared, for example, to Netflix, which needed three and a half years to reach the same milestone.

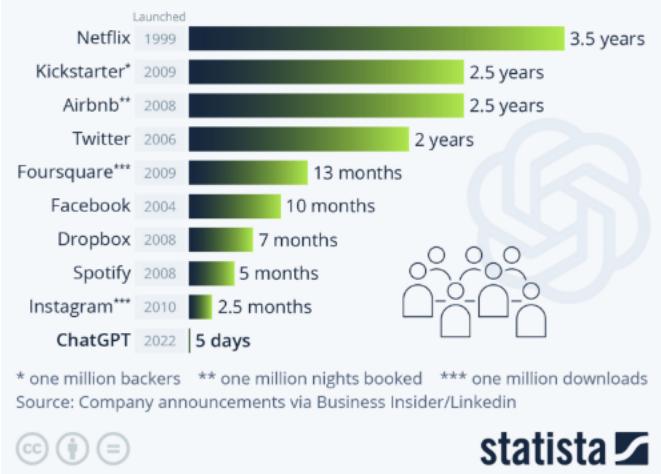

How long did it take for various applications to reach 1 million users? (Source: statista.com)

What sparked this enormous interest?

While AI had previously been seen mainly as a field of research in industry, or for IT specialists, ChatGPT's interface let anyone experience that it is possible to communicate with a computer through natural human language. ChatGPT answered questions posed in natural language quickly, coherently, and convincingly, and this simultaneously sparked enthusiasm, surprise, and concern in people.

We encounter many similar images online. What do these images suggest? (Source: ChatGPT)

Over the past decades we have got used to computers as part of our lives. They are indispensable tools capable of executing commands. We are not even surprised when robots in sci-fi movies speak in a strange mechanical voice. But when “someone” on the phone speaks in a perfectly human voice, we are still surprised, even today, discovering that it is a machine, not a person, on the other end of the line.

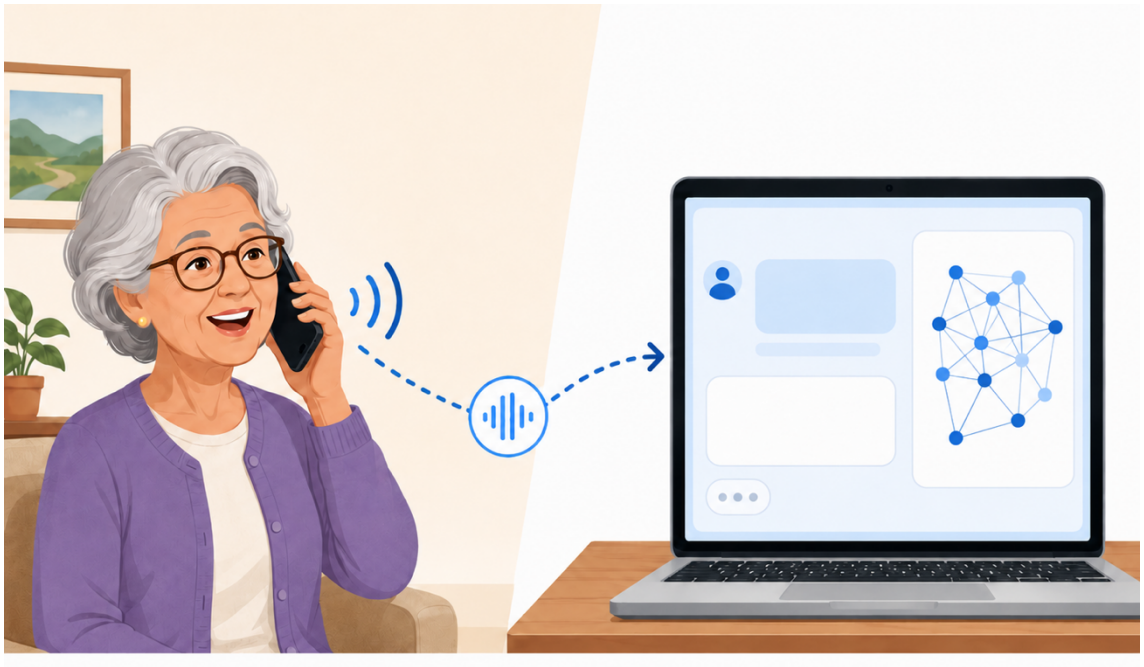

Gramma has a chat to AI (Source: ChatGPT 5.6 Sol)

A few years have passed since ChatGPT first appeared, but the confusion and extreme views surrounding it are still with us. Even today, many people accept its answers unconditionally and without doubt, while others reject it completely and do not trust its reliability at all.

Is the machine really learning?

When we say that a computer “learns”, we certainly do not mean the same process by which a human being, especially a young child, learns. Science today still cannot fully and precisely describe the complex processes of human learning.

Behind machine learning there lie decades of research and technological development.

> The collective term for algorithms and methods that learn from data is **machine learning**, which has now become one of the most important subfields of artificial intelligence. Using machine learning algorithms, the computer recognizes patterns based on data and examples and uses them to generate predictions and decision recommendations.

The essence of machine learning is determined by three things: **data, the model, and the learning paradigm**. The “recipe” is simple: you need digitally stored **data**, from which the AI learns; **a learning algorithm** that specifies how the learning takes place; and **a model** in which the learned relationships are stored.

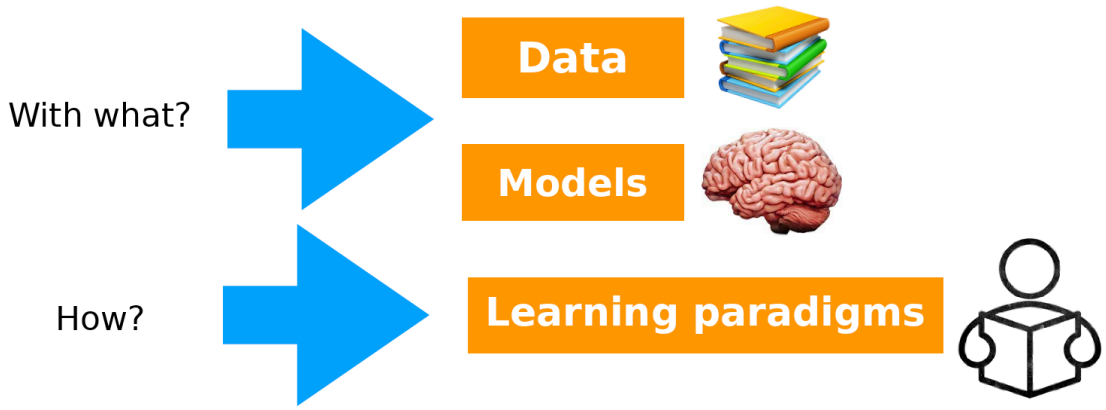

The building blocks of machine learning

Large amounts of data are used for training; this dataset is called the **training database** or **training dataset**. As a result of the learning process, a **model** is created, which is able to identify unknown data based on the relationships it recognised during training. However, the model can only draw conclusions from what it has learned from the training data: it has no independent knowledge or experience beyond that.

Machine learning algorithms (also known as learning paradigms) are divided into three main groups:

- supervised learning,
- unsupervised learning, and
- reinforcement learning.

Once the learning process is complete, **the model will serve as the core of the AI application,** drawing on the knowledge it has acquired.

Types of Machine Learning (ML)

Supervised Learning

In **supervised learning**, the machine learns from examples in which the correct answers are provided along with the input data. An example of this could be millions of photographs labelled with the names of the objects shown in them.

This is a dog! (Forrás: ChatGPT)

The essence of the learning algorithm is the easiest to understand through a concrete example. Imagine that we want to train a model to decide whether a picture shows an apple or a banana.

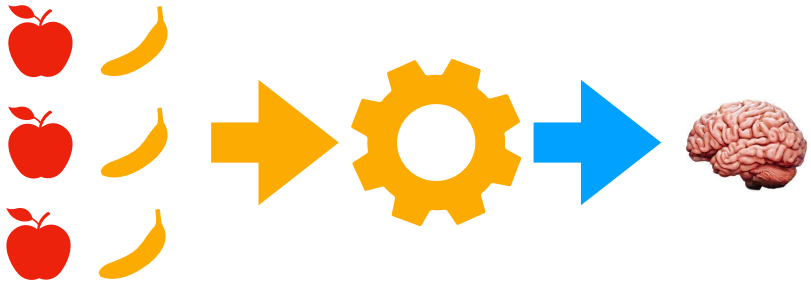

Input data Learning algorithm Model

Train and build the model

During training, we show the algorithm thousands of images – and for each image, the correct answer is provided: “this is an apple”, “this is a banana”. These predefined correct answers are called labels – hence the name: labelled data. From these examples, the model learns which visual features – shape, colour, etc. – indicate which fruit!

If the training is successful, the model will also be able to decide, **with high probability**, for an unknown object it has never seen before, whether it is an apple or a banana – without anyone telling it the answer.

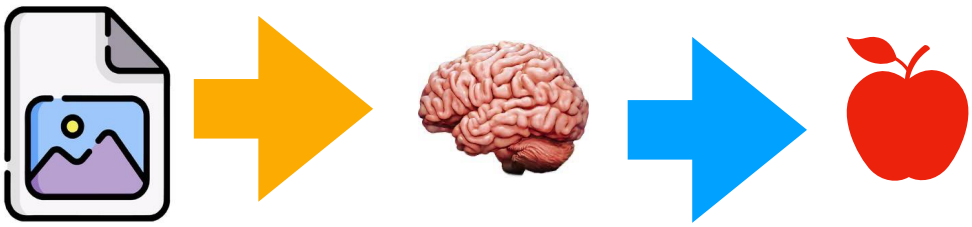

The working model (the model is illustrated by a diagram of the human brain)

The same principle applies to speech recognition, handwriting recognition, and facial recognition, except that the model is trained not with visual information but with other types of data (for example, text or audio).

Speech recognition Text generation, translation

Unsupervised Learning

During **unsupervised learning**, there is no predefined correct answer: the machine is not given the right answers in advance, and the observed objects have no predefined names (labels).

Cinderella did it in a flash! (Source: ChatGPT)

The algorithm's task is to independently discover the hidden patterns in the input dataset (clusters, similarities, differences), look for relationships in the dataset and organize (e.g. group) the data.

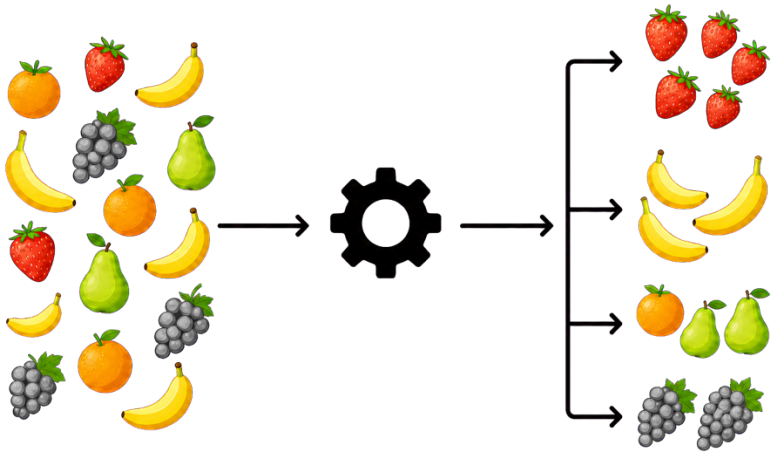

Unsupervised ML – for example, clustering

Such algorithms are used to build the models behind media recommendation systems: they group users' interests based on their behaviour, and recommend content accordingly – Netflix suggests personalised films, Spotify suggests music, and Amazon suggests products.

amazon NETFLIX Spotify

Reinforcement Learning

For supervised learning, the machine is given data; for reinforcement learning, it is given rules. **Reinforcement learning** is a learning method in which a virtual "agent" interacts with its environment and tries to make optimal decisions. It learns, through rewards or penalties, which actions lead to the best outcome.

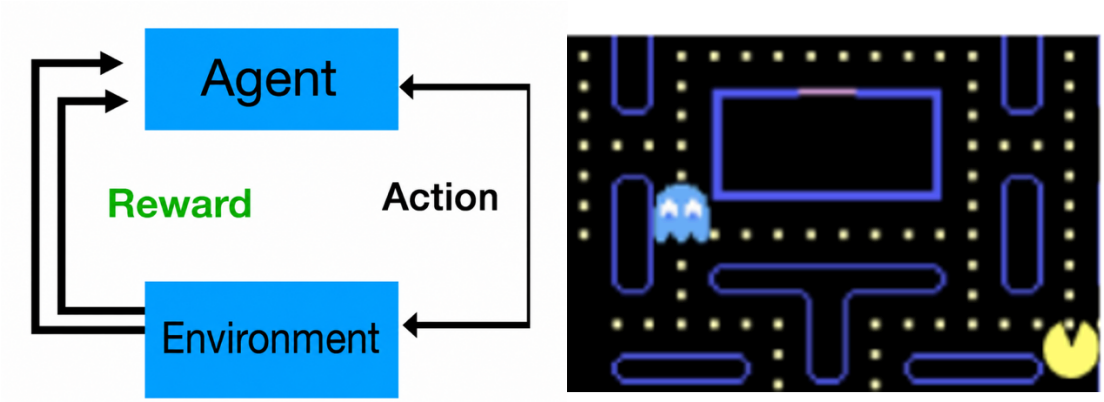

Reinforcement learning: the modell receives a reward or punishment

Chess, Go, and Reinforcement Learning

In 1997, a computer defeated Kasparov, the reigning world chess champion, for the first time. IBM's Deep Blue program was able to analyse millions of moves very quickly using brutal computational power and a strong chess-evaluation system. The program was built using traditional programming and did not run on artificial intelligence.

However, the game of Go could not be taught to a machine using classical methods, because in Go there are more possible positions and game sequences than the number of atoms in the universe. Even today, there is no powerful machine that would be able to store all possible states of the game. DeepMind's system (AlphaGo), *through reinforcement learning, by practising against itself,* learned to play at a level that allowed it to defeat Lee Sedol, one of the world's best players, in 2015.

AlphaGO wins the game (Source: https://artsandculture.google.com/)

Transfer learning

The three basic algorithms are complemented by a fourth: knowledge-transfer-based learning, or **transfer learning**. This algorithm allows a completed, well-functioning model to be further developed by continuing its training in a narrower, specialised area — for example, if we want the model to use a source we provide (such as a textbook) when generating content.

Application Areas of Machine Learning Algorithms

Different machine learning algorithms are used for different tasks. Supervised learning is primarily suited to classification and prediction, such as object recognition, spam filtering, credit risk assessment, or supporting medical diagnostics.

Unsupervised learning is used to explore the internal structure of large datasets: it can help identify similar customer groups, detect anomalous events, or simplify complex data

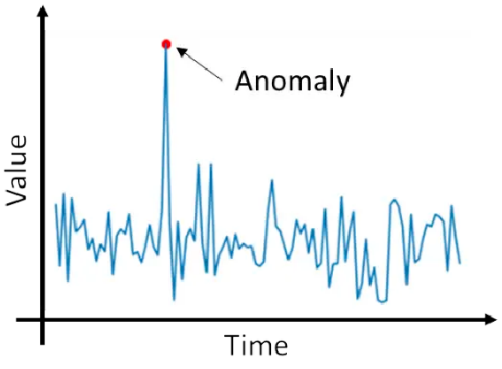

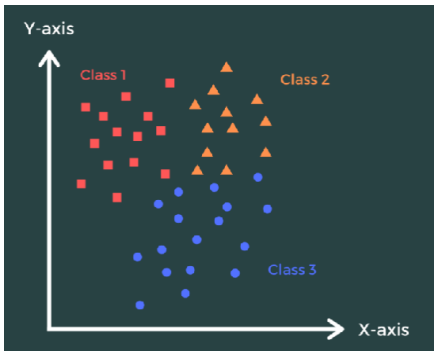

Reinforcement learning is applied, for example, in robot control, game programs, optimising transport systems, or certain automated control tasks – in situations where a series of successive decisions must be made.

The Resources and Costs of Machine Learning

The resource requirements of machine learning methods depend heavily on the amount of data, the complexity of the model, and the chosen learning method. With supervised learning, collecting and labelling training data can represent a significant cost, while with unsupervised learning it is primarily the storage and processing of large datasets that requires computing capacity.

Reinforcement learning is particularly costly, because the system has to gain experience through many trials, often in a simulated environment.

Why is "training" a model so expensive?
The *Hugging Face* AI company built a publicly available training dataset called *FineWeb*, which was created for pretraining language models. The text was collected from several billion web pages, after which faulty content, duplicates, and personal data were filtered out over several months of work. *The end result is about 54 terabytes of text, which anyone can use for free for training*. The finished model's knowledge is limited to what it managed to "learn" from this filtered, processed dataset.

Takeaways for Teachers

"Artificial intelligence does not think, it is talking." (Vilmos Csányi,[2] 2026)

The quote very clearly expresses what we've learned about machine learning in the previous chapters. *Trained models, of course, do not think in the sense that we, humans, think. They generate text using mathematical-statistical methods — they "talk" without understanding what they are "saying".*

What does it mean that artificial intelligence systems do not operate deterministically?

Traditional software gives the same result for the same input. Artificial intelligence systems, based on mathematically learned patterns, provide the result judged most probable by the algorithm. As a result, if several people ask the same question, they may get different answers from one another. The answer is not deterministic; only with a certain probability can it be considered true, and there is always some probability of error.

Does the "artificial" adjective mean that a machine version of natural human intelligence has been achieved?

No. Artificial intelligence is not a mechanical copy of human intelligence. AI systems can partially model certain abilities associated with intelligence — such as pattern recognition, learning, language-based responses, or prediction. However, they do not possess general understanding, consciousness, experience, and true comprehension characteristic of humans.

[2] Vilmos Csányi (1935–) Hungarian ethologist, university professor, academician, and writer.

Are AI systems smarter than humans?

While they may be faster and more accurate than humans in certain narrow tasks, they lack the general intelligence, experience, and judgment that humans possess.

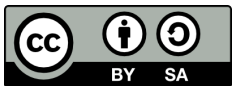

Funded by the European Union. Views and opinions expressed are those of the author(s) only and do not necessarily reflect those of the European Union or the National Agency.