TALOS ERA CHAIR IN ARTIFICIAL INTELLIGENCE FOR HUMANITIES AND SOCIAL SCIENCES

ΠΑΝΕΠΙΣΤΗΜΙΟ ΚΡΗΤΗΣ
UNIVERSITY OF CRETE

What is Artificial Intelligence?

Prof Christophe Roche
ERA Chair Holder - University of Crete
https://christophe-roche.fr/

May 2025

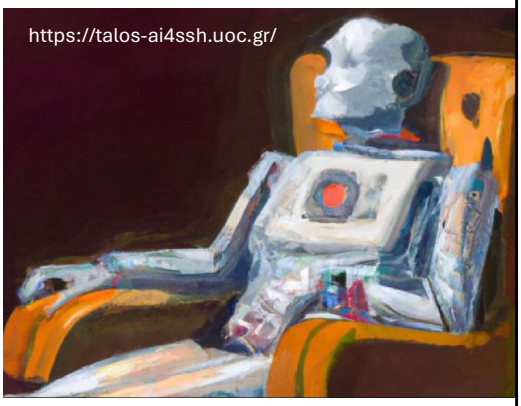

Horizon ERA Chair TALOS AI4SSH Project funded by the European Commission
Grant Agreement n° 101087269, https://cordis.europa.eu/project/id/101087269

This MOOC is an introduction to Artificial Intelligence that doesn't require any prerequisites in computer science.

What is Artificial Intelligence?

Contents

"The question of whether computers can think is just like the question of whether submarines can swim."

— EDSGER W. DIJKSTRA

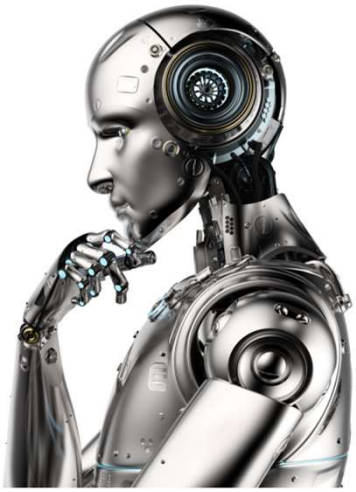

It is divided into 9 points. We will begin by introducing the notion of intelligence to arrive at a definition of Artificial Intelligence. We will then say a few words about the history of this relatively old discipline and its many applications, before introducing the two traditional approaches to AI: connectionist and symbolic, and the potential of combining them. We will conclude by stressing the importance of promoting AI ethics.

1. Introduction

What is intelligence?

Is a chess computer intelligent?

What is intelligence? And more specifically, what might artificial intelligence be? To help us, let's begin with a first example: "Is a chess-playing program intelligent?"

1. Introduction

What is intelligence?

Is a chess computer intelligent?

<u>It can</u>:

- **Represent (the world)**
- **Reason (on the representation of the world)**
- **Act (on the representation of the world)**

This program is capable of 1)Representing the world, and more precisely, <u>its</u> world, that is, the chessboard, its pieces, their value, their position, and the rules of the game. 2) Reasoning about this representation 3) Acting on the world by changing the game configuration.
It has all the features of intelligence...

1. Introduction

What is intelligence?

Is a chess computer intelligent?

It can:

- **Represent (the world)**
- **Reason (on the representation of the world)**
- **Act (on the representation of the world)**

But… dedicated to one task

Except that it is dedicated to a single task and cannot play a different game.
Adaptability is an essential quality.

1. Introduction

What is Human Intelligence?

"The ability to acquire and apply knowledge and skills" Oxford Dictionary

"the ability to learn or understand or to deal with new or trying situations" Merriam-Webster

"the ability to learn, understand, and make judgments or have opinions that are based on reason" Cambridge

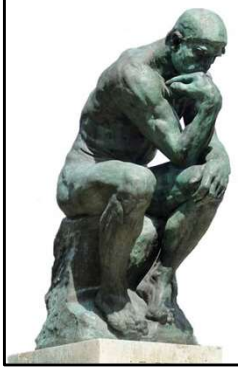

Rodin Thinker
Paris (France)
04/08/2023

ChatGPT

Human intelligence is a complex and multifaceted concept that refers to the cognitive abilities and mental capacity of human beings.

It involves the ability to reason, think abstractly, solve problems, learn from experience, understand complex ideas, communicate effectively, and adapt to new situations.

What about human intelligence? Let's open a few dictionaries:
The *Oxford Dictionary* defines it as the ability to acquire and apply knowledge and skills.
Merriam-Webster adds the ability to "deal with new situations."
The *Cambridge Dictionary* includes reasoning.

ChatGPT, a conversational agent based on generative AI, summarises the core characteristics of human intelligence as: reasoning, thinking, problem-solving, understanding, communicating, and adapting.

However, nothing is said about knowledge and its representations, without which nothing is possible.

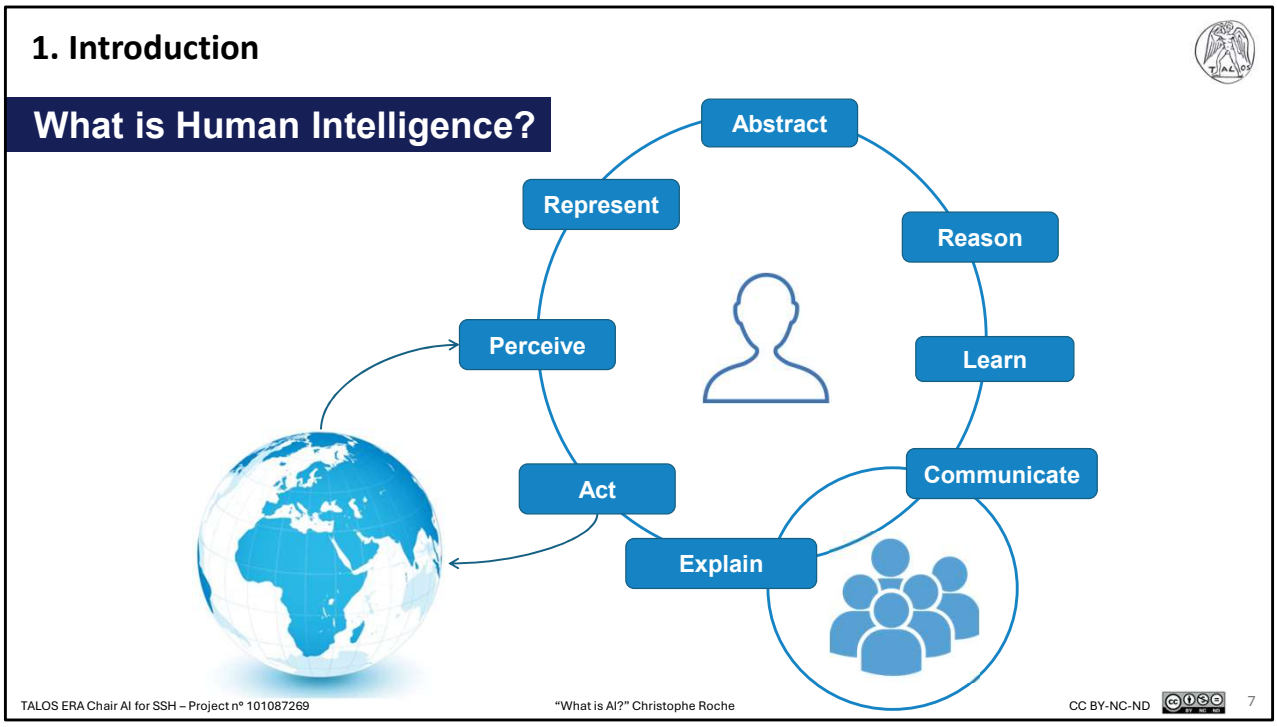

This slide identifies the main functions of human intelligence and the involved agents: the world, an agent assumed to be intelligent, and other agents, human or not. Perceiving the world, representing it, reasoning, learning to adapt to new situations, communicating, explaining, and acting are all essential functions.

1. Introduction

What is Artificial Intelligence?

"**Artificial Intelligence** (AI) is the part of **computer science** concerned with designing intelligence computer systems, that is, systems that exhibit the **characteristics we associate with intelligence in human behavior** – understanding language, learning, reasoning, solving problems, and so on" (*The Handbook of Artificial Intelligence)*

Artificial Intelligence is defined based on characteristics associated with human intelligence. According to the *Handbook of AI*, AI is the part of computer science concerned with designing programs whose behaviour could be called intelligent if exhibited by a human being.

1. Introduction

What is Artificial Intelligence?

"**Artificial Intelligence** (AI) is the part of **computer science** concerned with designing intelligence computer systems, that is, systems that exhibit the **characteristics we associate with intelligence in human behavior** – understanding language, learning, reasoning, solving problems, and so on" (*The Handbook of Artificial Intelligence*)

Strong A.I.
Computer is a model of brain
AI system can ***think*** as human does

Weak A.I.
AI System whose ***behaviour*** can be said to be intelligent

"The question of whether computers can think is just like the question of whether submarines can swim."

— EDSGER W. DIJKSTRA

Two visions of AI exist, each with important implications. So-called *strong AI* assumes that computers are suitable models of the brain and that AI systems can think like humans. In contrast, *weak AI* focuses more on the behaviour of these systems, which may be called intelligent, treating intelligence as a purely human property. The quote from Dijkstra is particularly relevant.

2. History

AI begins with Computers

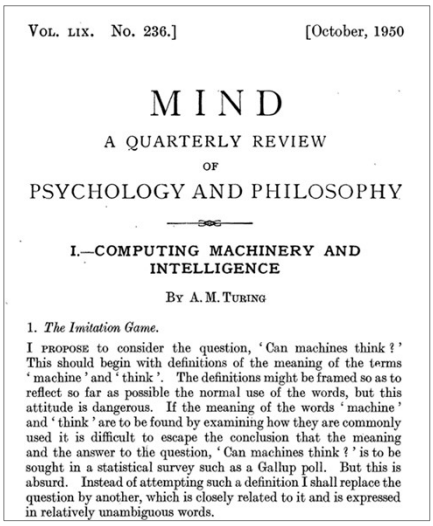

Vol. lix. No. 236.] [October, 1950

MIND

A QUARTERLY REVIEW

OF

PSYCHOLOGY AND PHILOSOPHY

I.—COMPUTING MACHINERY AND INTELLIGENCE

By A. M. Turing

1. *The Imitation Game.*

I propose to consider the question, 'Can machines think?' This should begin with definitions of the meaning of the terms 'machine' and 'think'. The definitions might be framed so as to reflect so far as possible the normal use of the words, but this attitude is dangerous. If the meaning of the words 'machine' and 'think' are to be found by examining how they are commonly used it is difficult to escape the conclusion that the meaning and the answer to the question, 'Can machines think?' is to be sought in a statistical survey such as a Gallup poll. But this is absurd. Instead of attempting such a definition I shall replace the question by another, which is closely related to it and is expressed in relatively unambiguous words.

'Can machines think?'

"Computing Machinery and Intelligence"
A.M. Turing
Mind, Volume LIX, Issue 236,
1st October 1950, Pages 433–460.

Turing is widely considered to be the father of theoretical computer science and artificial intelligence

Alan Turing

1912-1954

Alan Mathison Turing was an English mathematician, computer scientist, logician, cryptanalyst, philosopher, and theoretical biologist. Turing was highly influential in the development of theoretical computer science, providing a formalisation of the concepts of algorithm and computation with the Turing machine, which can be considered a model of a general purpose computer.

AI began with computers. It is a field of computer science. A foundational article by Alan Turing in 1950 posed the question: "Can machines think?"

2. History

AI begins with Computers

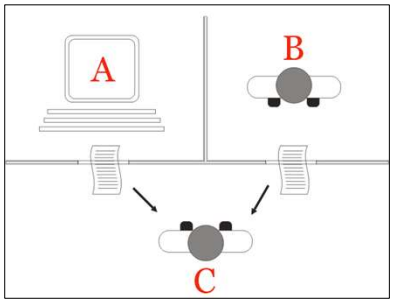

Alan Turing

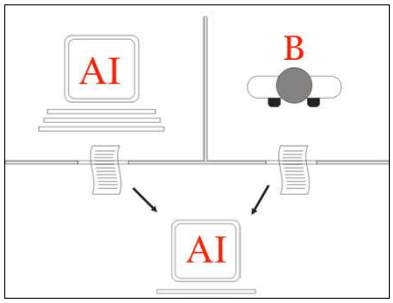

The Turing test, developed by Alan Turing in 1950, is a test of a machine's ability to exhibit intelligent behavior equivalent to, or indistinguishable from, that of a human.

Turing introduced the test that bears his name, which assesses whether a program can be deemed intelligent based on its responses. Ironically, today it is computers that test us to confirm we are human.

2. History

Games & Logic

Game software are used as a measure of progress in AI.

In 1951, Christopher Strachey wrote a checkers program and Dietrich Prinz wrote a program for chess.

The computer is the Mark I machine of the University of Manchester.

“Logic Theorist is a computer program written in 1955 and 1956 by Allen Newell, Herbert A. Simon and Cliff Shaw. It was the first program deliberately engineered to mimic the problem solving skills of a human being and is called "**the first artificial intelligence program**”. It would eventually prove 38 of the first 52 theorems in Whitehead and Russell's Principia Mathematica, and find new and more elegant proofs for some” (Wikipedia)

Early AI programs focused on games and logic, stressing on knowledge representation and reasoning, laying the foundation for what is called *symbolic AI,* also known as classical AI.

2. History

The Dartmouth Conference

In 1956 John McCarthy regarded as the father of AI, organized a conference to draw the talent and expertise of others interested in machine intelligence for a month of brainstorming. He invited them to Vermont for "The Dartmouth summer research project on artificial intelligence."

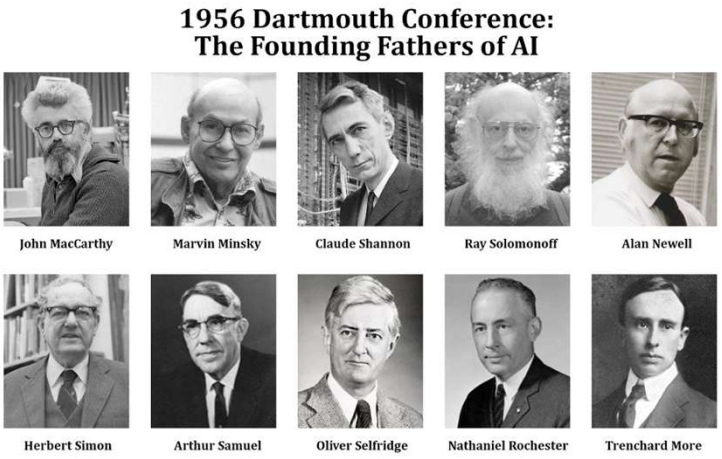

The term *Artificial Intelligence* was coined in 1956 at a conference that brought together the founding fathers of AI.

2. History

AI Periods

■ **Euphoric period (1952-1969)**

- General Problem Solver
- Lisp
- Robinson's Resolution Principle

■ **Return to reality (end of 60's – beginning of 70's)**

- Machine Translation (English-Russian)

■ **Expert Systems (1970's)**

- DENDRAL: molecule structure identification
- Mycin: medical diagnosis
- PROSPECTOR: drilling site choice

■ **Industry of AI (1980's)**

- Japan's 5th generation project
- Software tools for expert systems
- LISP-specific hardware

■ **Dark ages (end of 1980's – 1990's)**

Strong AI did not reach its promises

"*It is not my aim to surprise or shock you – but the simplest way I can summarize is to say that there are now world machine that think, that learn and that create*" Simon (1957)

■ **Triumph of AI**

- Return of AI mainly due to the hardware progress and the availability of massive data
- AI embedded in everyday life

AI has gone through different periods. After an initial euphoric period with promising results, such as the General Problem Solver and the symbolic programming language Lisp, reality struck with disappointing outcomes in machine translation. The success of expert systems in the 1970s and their industrial applications in the 1980s gave a second wind to AI, and Symbolic AI in particular. A new dark period followed, mainly due to unrealistic expectations around strong AI. Since the 2000s, AI has made a comeback and is enjoying a real triumph, thanks to the technological advances that have made connectionist AI operational.

3. Domains of Application

Medicine

Smart City

Facial recognition

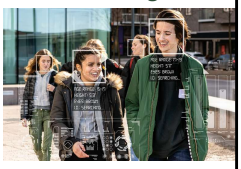

Robotics (Cobots)

Finance

AI Coding Assistant
AI Writing Assistant

Digital Humanities

Chatbot)

Today, AI is used in nearly every field, from medicine to robotics to finance. It is omnipresent in the assistants we use at all times and plays an increasingly important role in our societies. Smart cities have become a reality.

4. Scalability

Hardware & Data

Computers

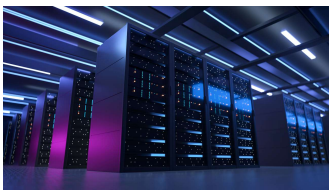

Big Data

Storage

Linked and Open Data

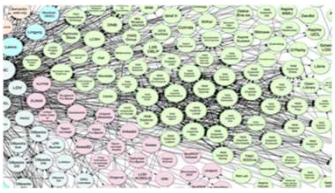

The current success of AI has been made possible by technological advances in computing, both in terms of computer processing and information storage capacity, and by the availability of large volumes of data, both structured and unstructured.

5. Symbolic A.I.

Symbolic AI (or Classical AI) is the branch of artificial intelligence which aims to **explicitly represent** **human knowledge** in a **declarative form** (i.e. facts and rules)

Epistemological Principles

Human- & computer-readable form

- ✓ Symbolic AI was the dominant paradigm of AI from the mid-1950s until the late 1980s
- ✓ The most successful form of symbolic AI is **expert systems**, which use a network of production rules. Production rules connect symbols in a relationship similar to an If-Then statement.

Ontology –Terminology Knowledge Graph

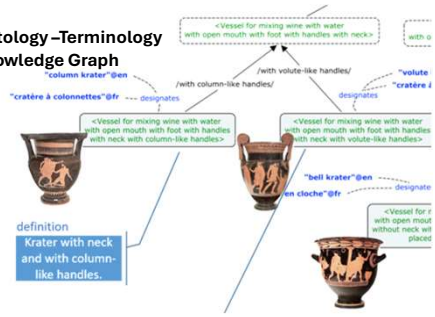

Expert Systems (reasoning)

If the soil pH *is less than* 6.0
and
the soil *has a high level of* aluminum
then it is not suitable for growing most crops.

Logic

Definition: Pocket-Watch(x) ::= Time-Piece(x) ∧ Portable(x) ∧ Pocket(x)
Properties: ⊨ ¬ (Pocket(x) ∧ Wrist(x))
Reasoning: Portable(x) ➞ (Wrist-Watch(x) ∨ Pocket-Watch(x))


Historically and conceptually, there are two main schools of thought in AI: *Symbolic AI* (also known as Classical AI) and *Connectionist AI* (also called sub-symbolic AI).

Symbolic AI is a branch of AI that aims to explicitly represent human knowledge in a declarative form. It means a form understandable by both humans and machines. Symbolic AI was the dominant AI paradigm until the late 1980s. Expert systems were among the most successful applications. Ontology and Knowledge Graph of Symbolic AI are now the semantic foundation of the Hybrid approach of AI.

6. Connectionist A.I.

Connectionist A.I. is the branch of artificial intelligence which aims to understand how the **human brain works at the neural level** and, in particular, how people learn and remember.

AI is nowadays most often associated with Deep Learning based on neural networks.

Deep Learning - Neural Network

Artificial Neural Networks are computing systems inspired by biological neural networks.

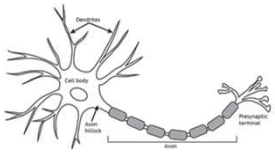

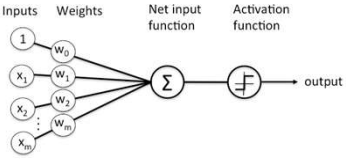

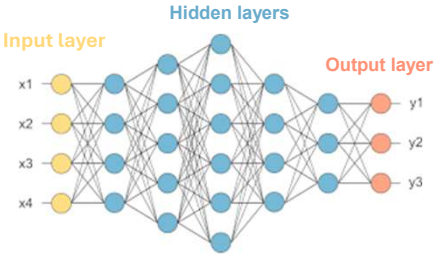

Such systems "**learn**" to perform tasks by considering **examples**. For example, in image recognition, they might learn to identify images that contain cats by analyzing example images that have been **manually labeled** as "cat".

Finding correlations between inputs and outputs

On the other hand, Connectionist AI is the branch of artificial intelligence inspired by how the brain works, particularly biological neural networks. Deep learning implementations use artificial neural networks arranged in layers and rely on statistical learning from large datasets. An artificial neuron computes an output by applying an activation function to a weighted sum of its inputs. These systems “learn” from examples by adjusting internal weights to minimize errors. Their goal is to model relationships between inputs and outputs, and the quality of results directly depends on the volume and quality of training data.

6. Connectionist A.I.

■ **Large Language Models (LLMs)**

Deep learning models trained on massive amounts of text data to:
- understand,
- generate (predict the next word in context),
- manipulate human language

Specific tasks: summarization, translation, Q&A, chat, etc.

■ **Generative A.I.**

A.I. systems to create new content:
- text,
- images, music,
- code, etc

Based on:
- deep neural networks
- LLMs

In addition to the first applications in image recognition, the most striking applications of neural networks include Large Language Models and Generative AI.

Trained on vast text corpora, LLMs can understand and generate text in natural language . They work by predicting the next word in a given context by capturing linguistic patterns.

Generative AI refers to systems capable of creating original content from training data: text, images, music, code, videos, etc. Unlike traditional predictive AI, it generates new data instead of merely classifying or detecting. Today, LLMs are central to generative AI in language tasks, powering the creation of coherent and contextualized text.

7. Symbolic A.I. *versus* Connectionist A.I.

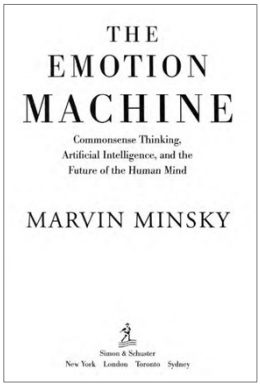

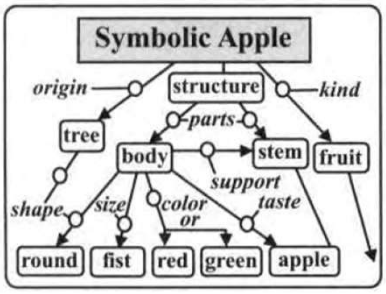

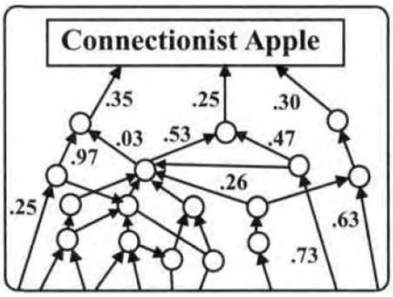

The diagram on the left shows a Semantic Network that describes various features and relationships between various aspects or parts of an apple. The diagram on the right shows an example of what is called a *"Connectionist Network,"* which also displays some aspects of an apple, but does not have any simple way to distinguish between different relationships; it only shows numbers that represent how closely those features are "associated." It would take too long here to explain how such networks

Knowledge representation is the main difference between Symbolic AI and Connectionist AI. Symbolic AI models knowledge in a human-readable form, while Connectionist AI uses numerical representations. That are weights or coefficients calculated during the learning phase and optimized for machine efficiency, especially in artificial neural networks. In the case of ChatGPT, these coefficients represent the connections between words. These connections are calculated from the source texts during the training phase.

The 'Minsky apple' is a clear illustration of these two approaches. The 'symbolic apple' is represented by a semantic network whose nodes and links carry meaning for humans, while the 'connectionist apple' is a network of unlabelled nodes and weighted links reflecting a notion of distance between nodes.

8. Hybrid A.I. (Neuro-Symbolic A.I.)

Combining the strengths of:

- ✓ **Symbolic AI:**
  - Knowledge Representation
  - Reasoning
  - Explainability
- ✓ **Connectionist AI :**
  - Learning from Data
  - scalability

- ➢ A Better Model for Leveraging Machine Learning and Human Expertise
- ➢ Solving the AI black box problem through transparency

Hybrid AI aims to combine the strengths of symbolic AI (explicit reasoning, formal representations, ontologies, knowledge graphs) and those of connectionist AI (deep learning, large-scale data processing, statistical robustness).

This hybrid approach makes it possible to overcome the limitations of each approach: the lack of explainability in neural networks and the scalability issues of symbolic systems. For example, current language models (LLMs) can be enriched by knowledge graphs to incorporate semantic reasoning. Hybrid AI thus paves the way for AI that is more reliable, explainable and capable of manipulating complex knowledge, while learning efficiently from data.

8. A.I. Ethics

There are many ethical challenges:

Lack of transparency of AI tools: AI decisions are not always intelligible to humans.

AI is not neutral: AI-based decisions are susceptible to inaccuracies, discriminatory outcomes, embedded or inserted bias.

Surveillance practices: data gathering and privacy.

Facial recognition

Protection?

Surveillance?

Autonomous weapons

select and engage targets without human intervention.

Autonomous weapons have been described as the third revolution in warfare, after gunpowder and nuclear arms.

Information processing raises ethical concerns, especially regarding personal data protection. AI systems exacerbate existing issues and introduce new ones, such as algorithmic bias producing discriminatory or inadequate outcomes, and the autonomy of AI systems whose opaque decisions raise problems of responsibility. AI must be regulated, and a legal framework defined to ensure that AI systems are safe, transparent and respectful of fundamental rights. The AI Act is the first comprehensive legislation in the world to regulate artificial intelligence within the European Union.