Artificial Intelligence

A Course About Foundations

UNIVERSITÀ DI PAVIA

Introduction

Marco Piastra

Prologue:
Apropos Artificial Intelligence

Artificial Intelligence: *then and now*

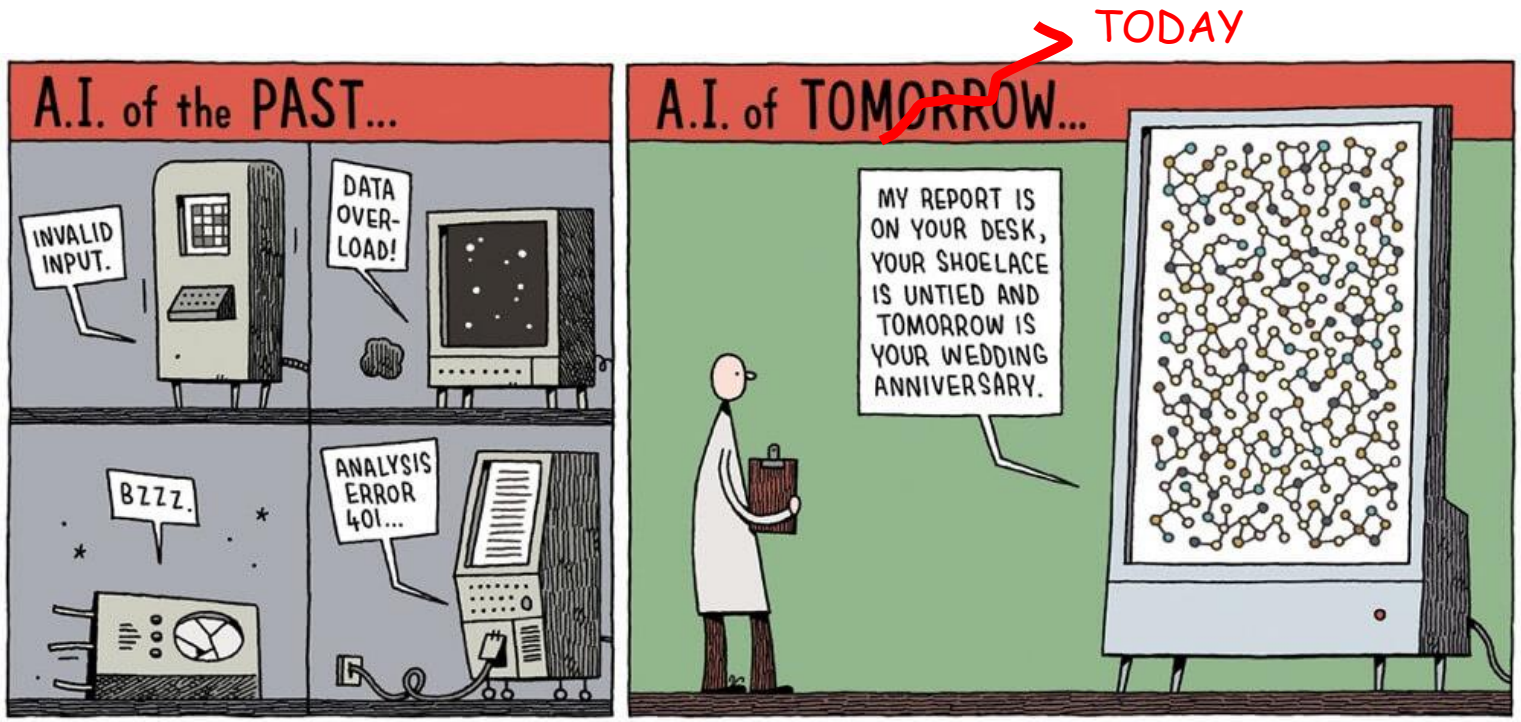

[Image from https://www.tomgauld.com/portfolio]

AI revolution?

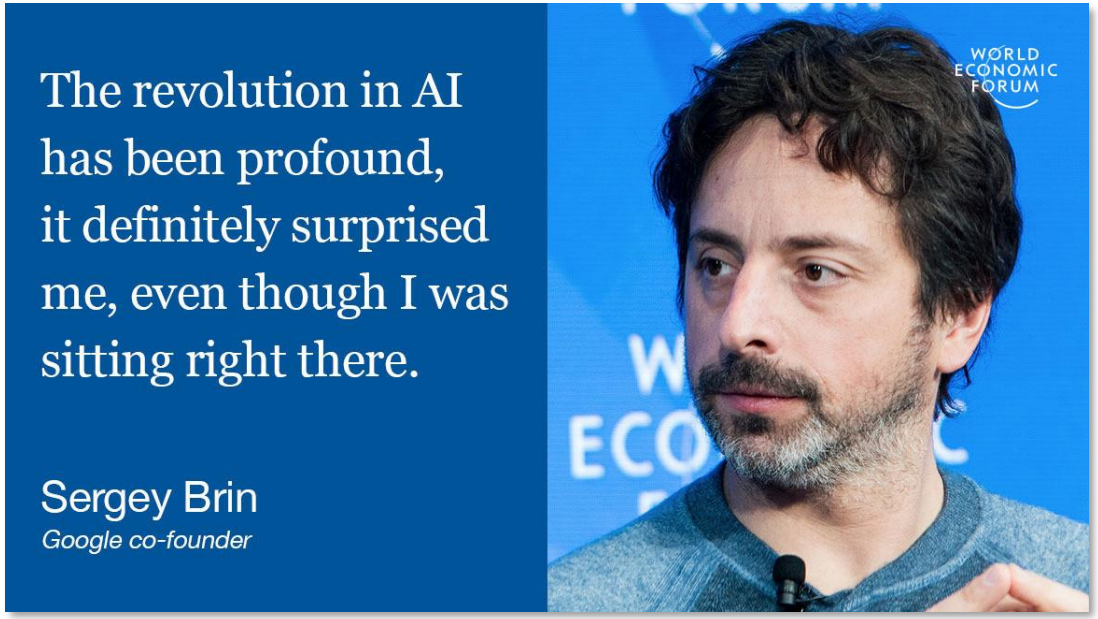

- Sergey Brin [Google Co-Founder, January 2017]

  "I didn't pay attention to it [i.e. Artificial Intelligence] at all, to be perfectly honest."

  "Having been trained as a computer scientist in the 90s, everybody knew that AI didn't work. People tried it, they tried neural nets and none of it worked."

[Quote and image from https://www.weforum.org/agenda/2017/01/google-sergey-brin-i-didn-t-see-ai-coming/]

The Deep Learning Revolution (in its early days)

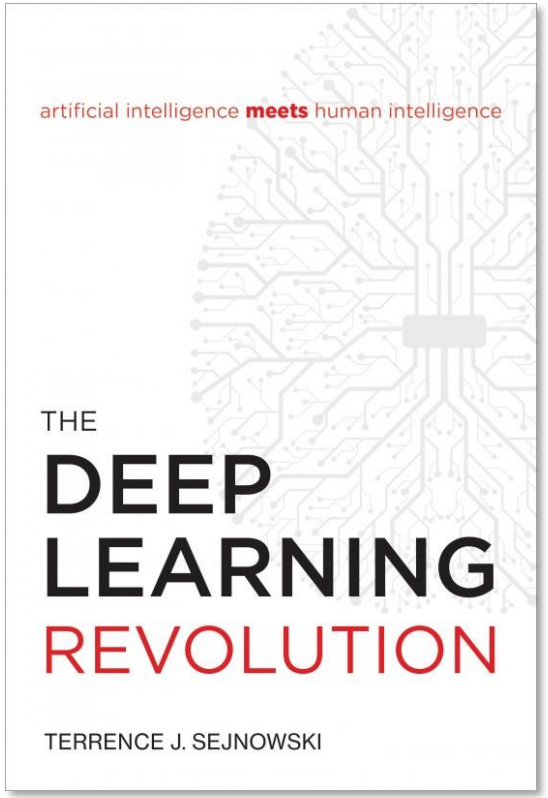

- Terrence J. Sejnowski [President of the Neural Information Processing (NIPS) Foundation, October 2018]

  "How deep learning—from Google Translate to driverless cars to personal cognitive assistants—is changing our lives and transforming every sector of the economy."

  "AI is now awakening and transforming our world.
  Driving these breakthroughs is the deep learning revolution [...]"

[Quote and image from https://mitpress.mit.edu/books/deep-learning-revolution]

Is Artificial Intelligence being oversold?

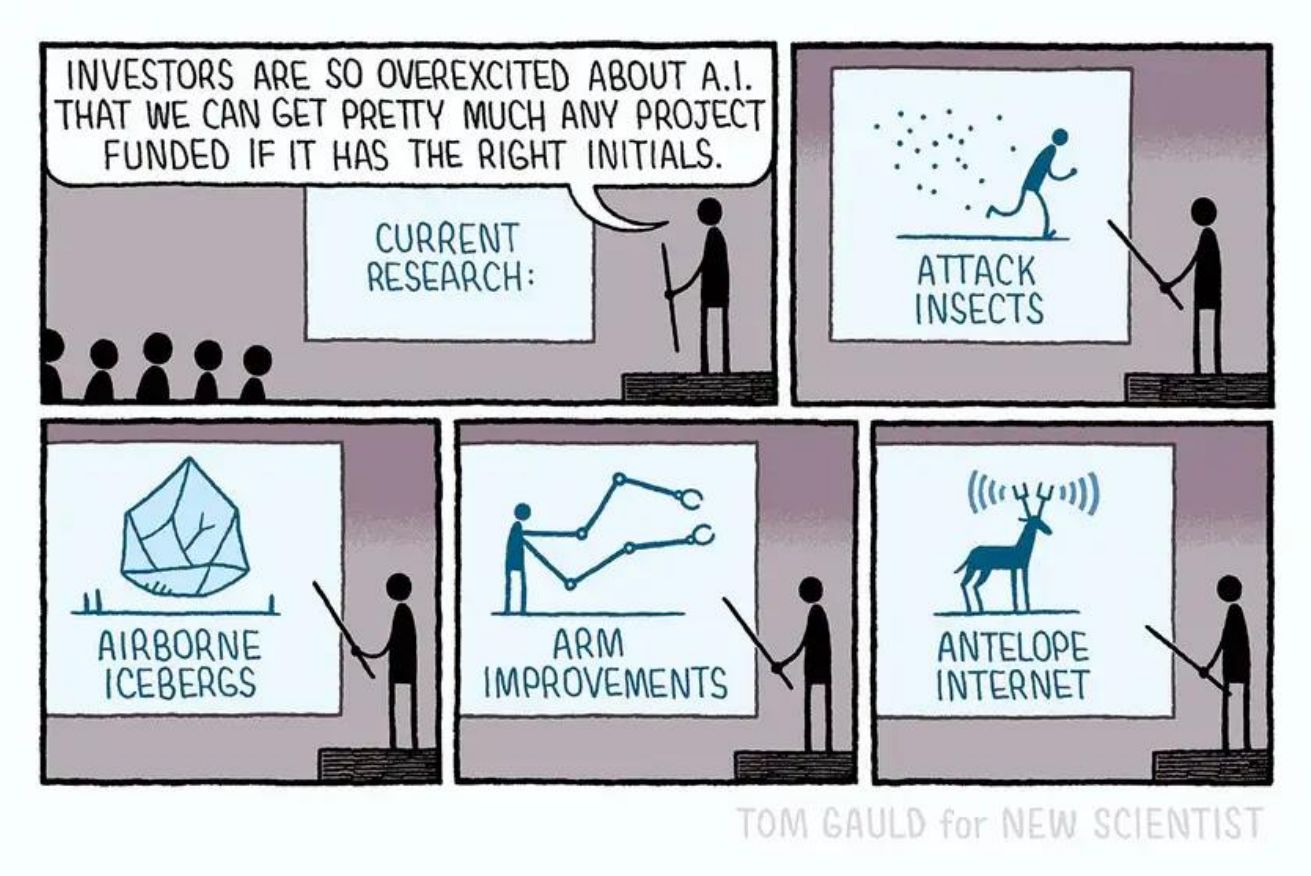

[Image from https://www.newscientist.com/article/2429313-tom-gauld-on-investing-in-ai/]

Artificial Intelligence hysteria?

AI isn't as scary as we imagine. AndreyZH/Shutterstock

The reality of AI is currently very different, particularly when you look at the threat of automation. Back in 2013, researchers estimated that, in the following ten to 20 years, 47% of jobs in the US could be automated. Six years later, instead of a trend towards mass joblessness, we're in fact seeing US unemployment at a historic low.

Current AI is good at finding **patterns** in large datasets, and not much else.

[Quote from https://theconversation.com/ais-current-hype-and-hysteria-could-set-the-technology-back-by-decades-120514]

If the Present is Confusing, the Future is Uncertain...

FINANCIAL TIMES

Artificial intelligence + Add to myFT

Elon Musk predicts AI will overtake human intelligence next year

Tesla chief says infrastructure will need to keep up with technology's demands as he seeks investment for own start-up

Elon Musk had previously suggested AI would surpass human intelligence by 2029 © Kirsty Wigglesworth/AP

George Hammond in San Francisco APRIL 8 2024 453

[https://www.ft.com/content/027b133f-f7e3-459d-95bf-8afd815ae23d]

OBSERVER

BUSINESS

Meta's A.I. Chief Yann LeCun Explains Why a House Cat Is Smarter Than The Best A.I.

"A cat can remember, can understand the physical world, can plan complex actions, can do some level of reasoning—actually much better than the biggest LLMs."

By Sissi Cao · 02/15/24 3:24pm

Yann LeCun testifies before the U.S. Senate Intelligence Committee on September 19, 2023 in Washington, D.C. Kevin Dietsch/Getty

[https://observer.com/2024/02/metas-a-i-chief-yann-lecun-explains-why-a-house-cat-is-smarter-than-the-best-a-i/]

The New York Times

The Ezra Klein Show

April 12, 2024

What if Dario Amodei Is Right About A.I.?

Anthropic's co-founder and C.E.O. explains why he thinks artificial intelligence is on an "exponential curve."

[https://www.nytimes.com/2024/04/12/opinion/ezra-klein-podcast-dario-amodei.html]

Computers vs. Humans

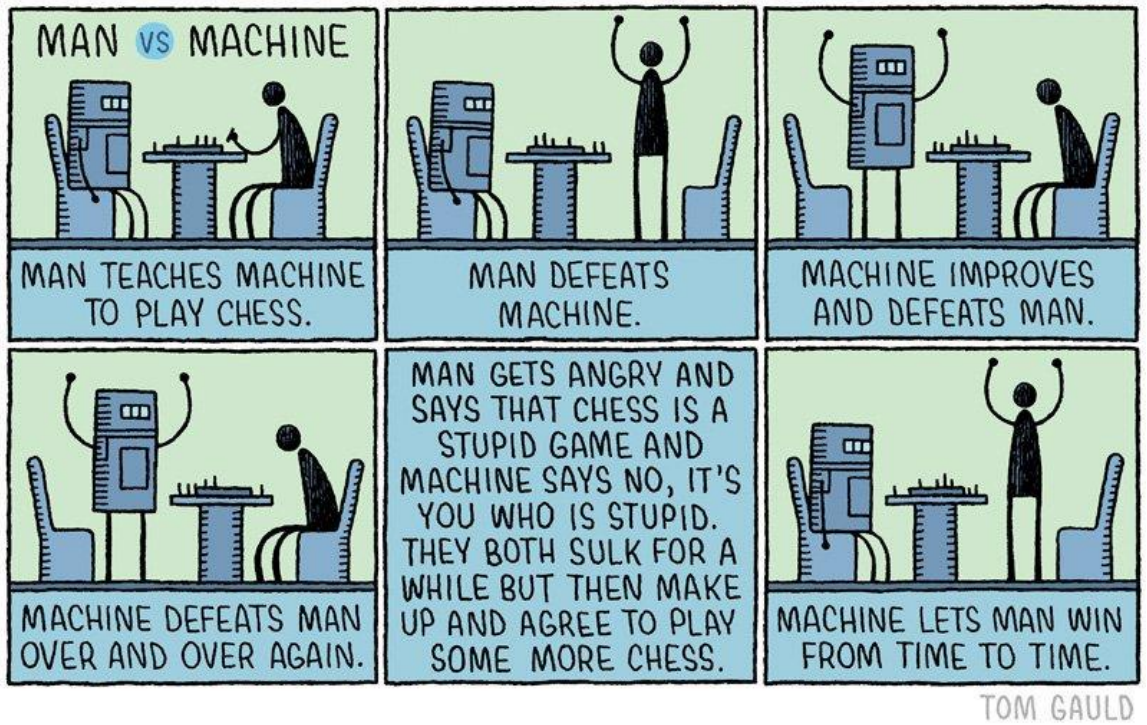

[Image from https://www.tomgauld.com/portfolio]

Artificial Intelligence

The Artificial Intelligence Cosmos

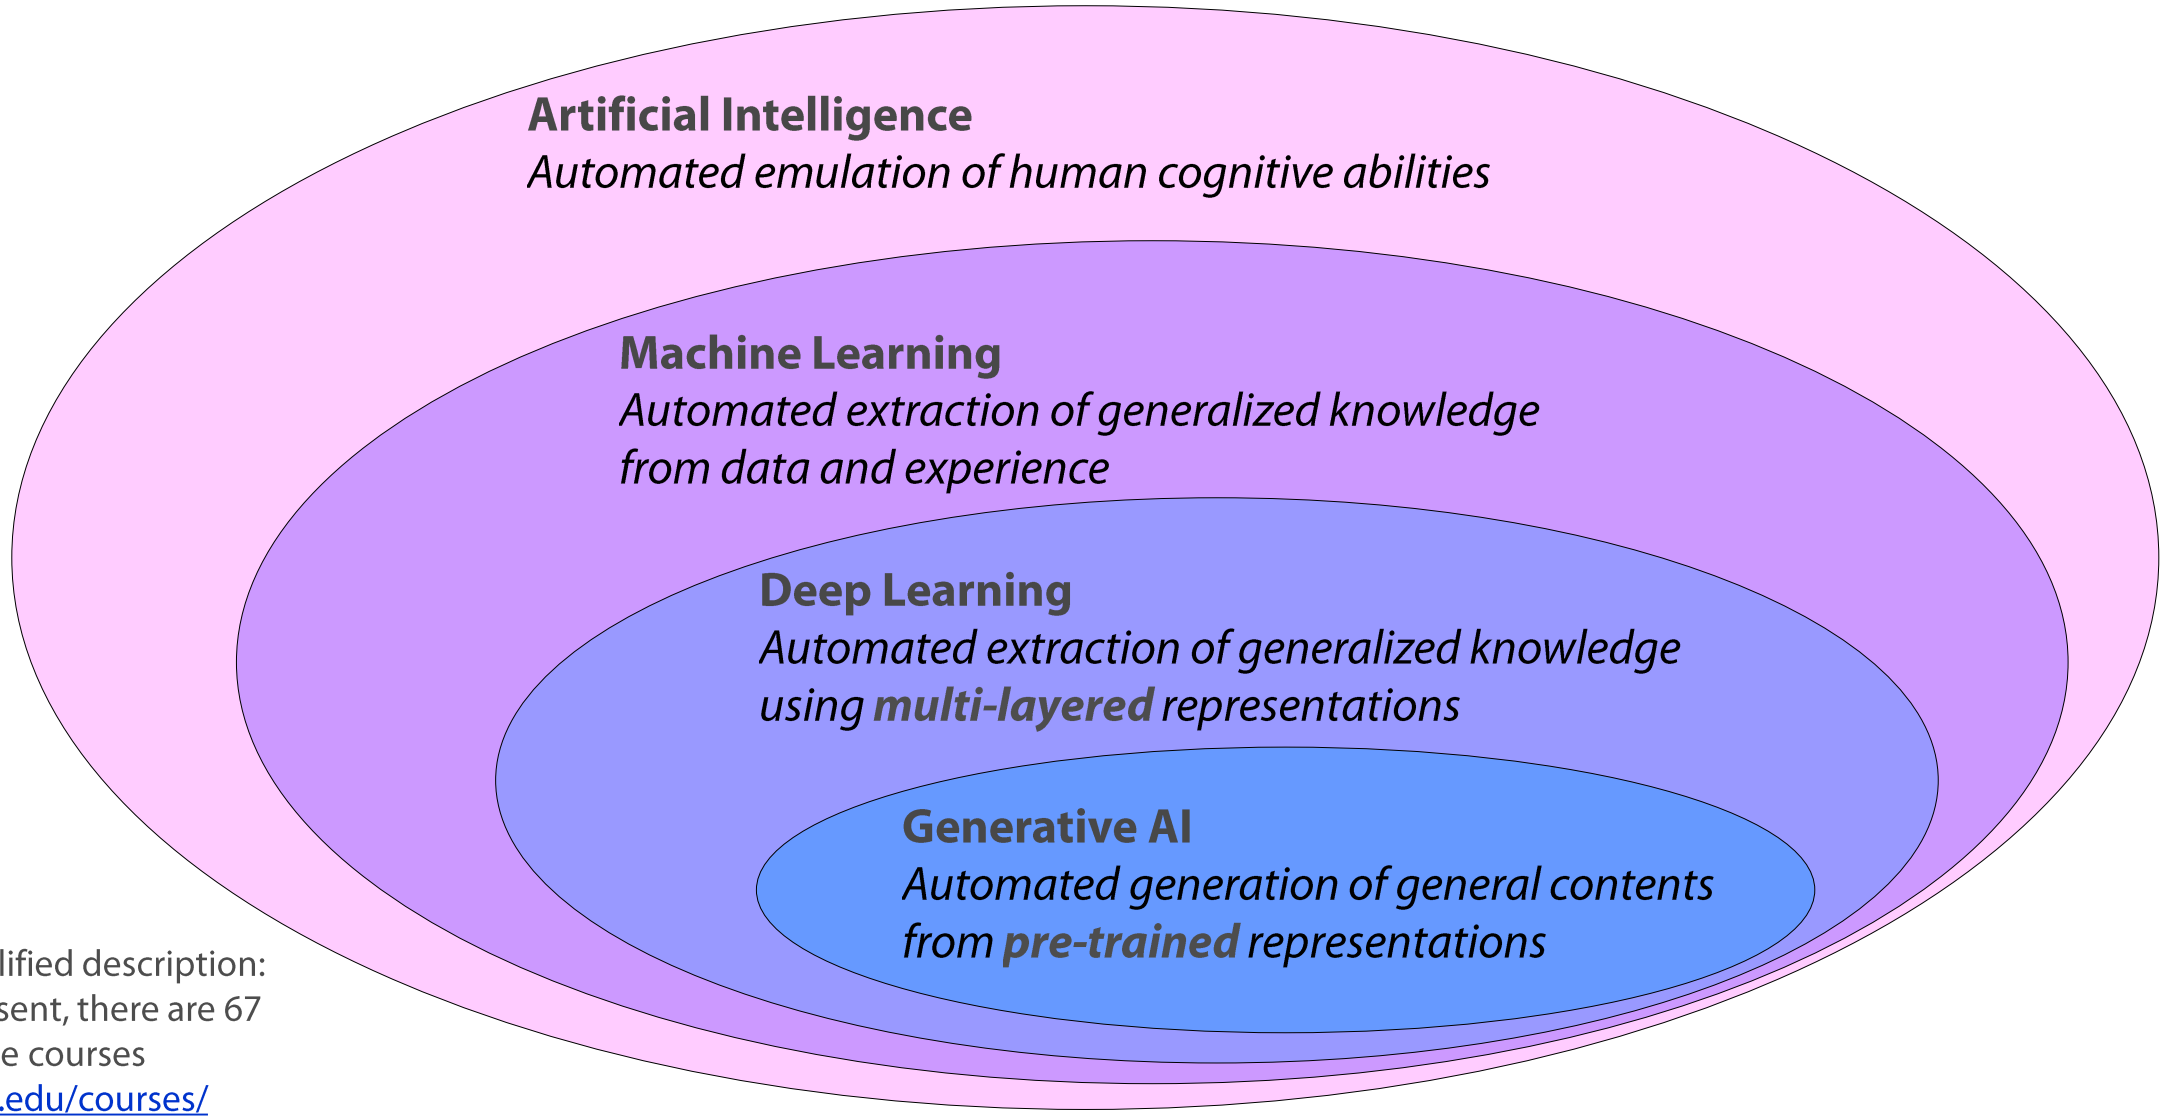

- This is a very simplified description: at Stanford, at present, there are 67 AI-related graduate courses https://ai.stanford.edu/courses/

"Artificial Intelligence" (first appearance of the term)

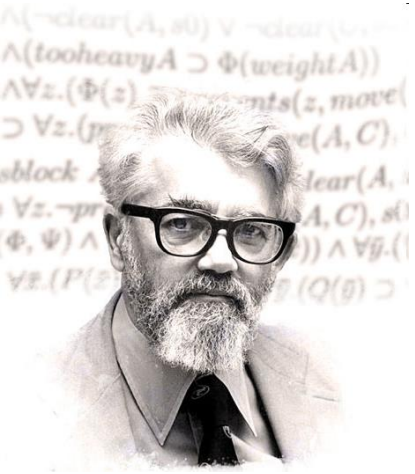

[Image from Wikipedia]

"We propose that a two-month, ten-man study
of **artificial intelligence** carried out during the summer of 1956 [...]

The study is to proceed on the basis of the conjecture
that every aspect of learning or any other feature of **intelligence**
can in principle be ***so precisely described***
that a machine can be made to ***simulate*** it. […]

It may be speculated that a large part of human thought
consists of manipulating **words** according to **rules** of **reasoning**
and **rules** of **conjecture**."

[John McCarthy et al., 1955 – *all emphases added*]

Automated Symbolic Logic

"The only way we know of expressing abstractions [...] is in language.
That is why we have decided to program a system which reasons verbally. "

[John McCarthy, PROGRAMS WITH COMMON SENSE, 1959]

- **Formal, Symbolic Logic *as a viable candidate***

  In those days, Formal Logic seemed an obvious choice

  - its formalism is derived from natural language
    (since G. Frege, 1879)
  - it has a clear semantics
    (since 'The Semantic Theory of Truth', A. Tarski, 1930)
  - it is *compositional*
    (the meaning of a complex expression is determined by its constituent and the rules used to combine them)
  - it can be turned into a *computing system*
    (see Herbrand's Theorem, 1930)

Automated Symbolic Logic

"The only way we know of expressing abstractions [...] is in language.
That is why we have decided to program a system which reasons verbally. "

[John McCarthy, PROGRAMS WITH COMMON SENSE, 1959]

- ***Logicism* dominated the early period of AI**

  The logical approach in three theses:
  (from 'Logic and artificial intelligence', N. J. Nilsson, 1991 – *emphasis added*)

  1. Intelligent machines will have ***knowledge*** of their ***environments***.
  2. The most versatile intelligent machines will represent much of their knowledge about their environments ***declaratively***.
  3. For the most versatile machines, the language in which declarative knowledge is represented must be at least as expressive as ***first-order predicate calculus*** [i.e. ***symbolic logic***].

AI Symbolic Framework

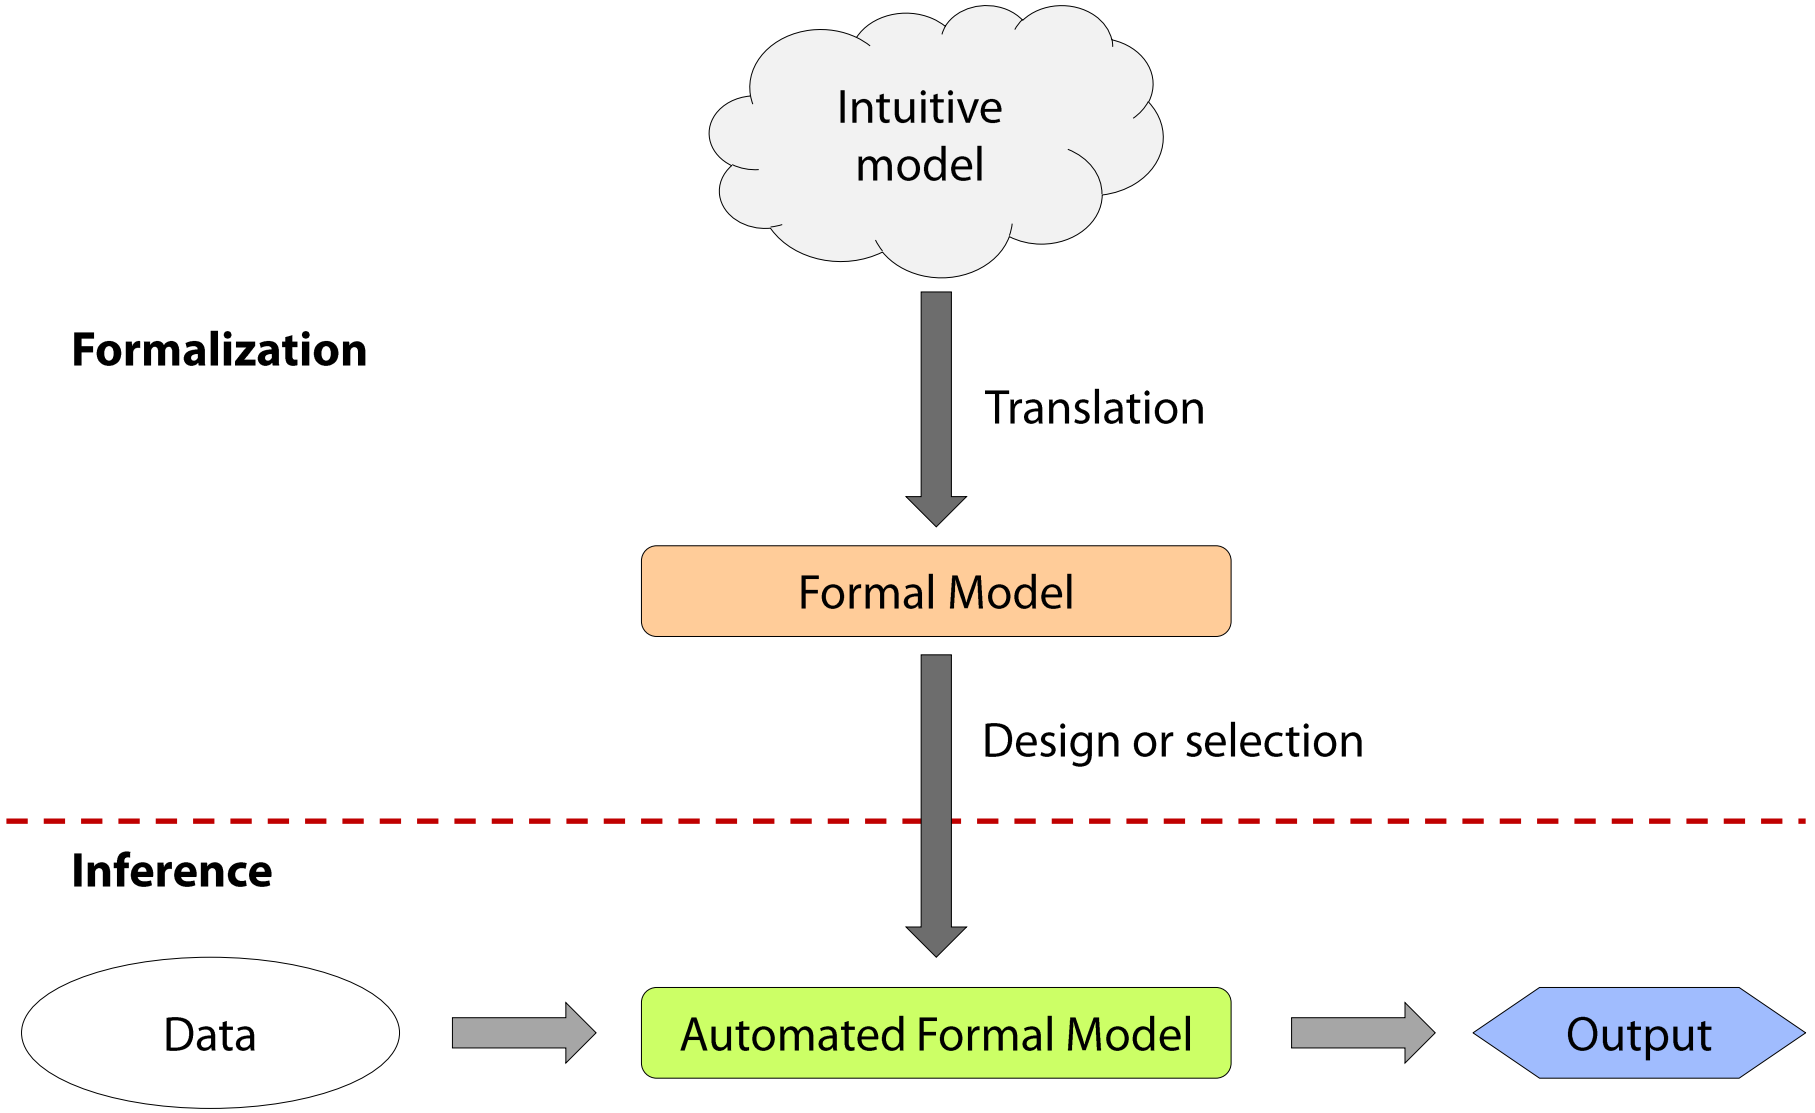

Criticism of the symbolic approach

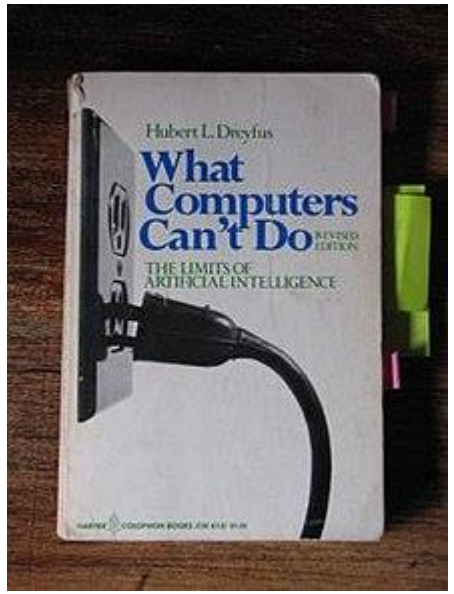

[Image from Wikipedia]

[H. Dreyfus, 1972]

Human intelligence and expertise depend primarily
on ***unconscious*** processes
rather than *conscious* symbolic manipulation,
and that these *unconscious* skills can never be fully captured in formal rules.

*(These ideas are now embedded in the so-called **sub-symbolic** approach)*

'Out of Control', BBC2 Horizon documentary, 2012

- *How big is the unconscious mind (in humans)?*

  In this BBC documentary, several senior neuro-scientists were asked to represent on a sheet of paper the extent of conscious activity (vs. unconscious) in the human brain

[Image from https://www.bbc.co.uk/programmes/p00pyhx2]

Connectionism: Artificial Neural Networks

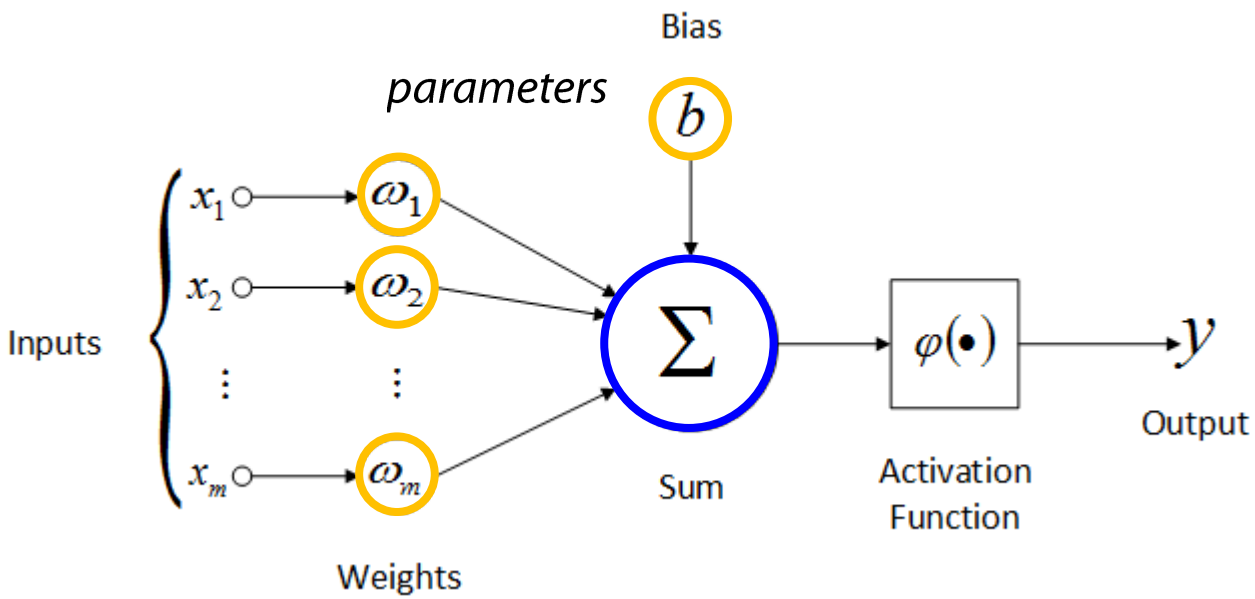

[Images from Wikipedia]

[Rumelhart, D.E., J.L. McClelland 1986]

- ***Basic assumption***

  Mental phenomena can be described
  by interconnected networks of simple and often uniform units

Artificial Neural Networks

- ***Learning* is an *optimization* process of numerical parameters**

  Using a very large dataset, made of data plus *annotations* (*data items*)

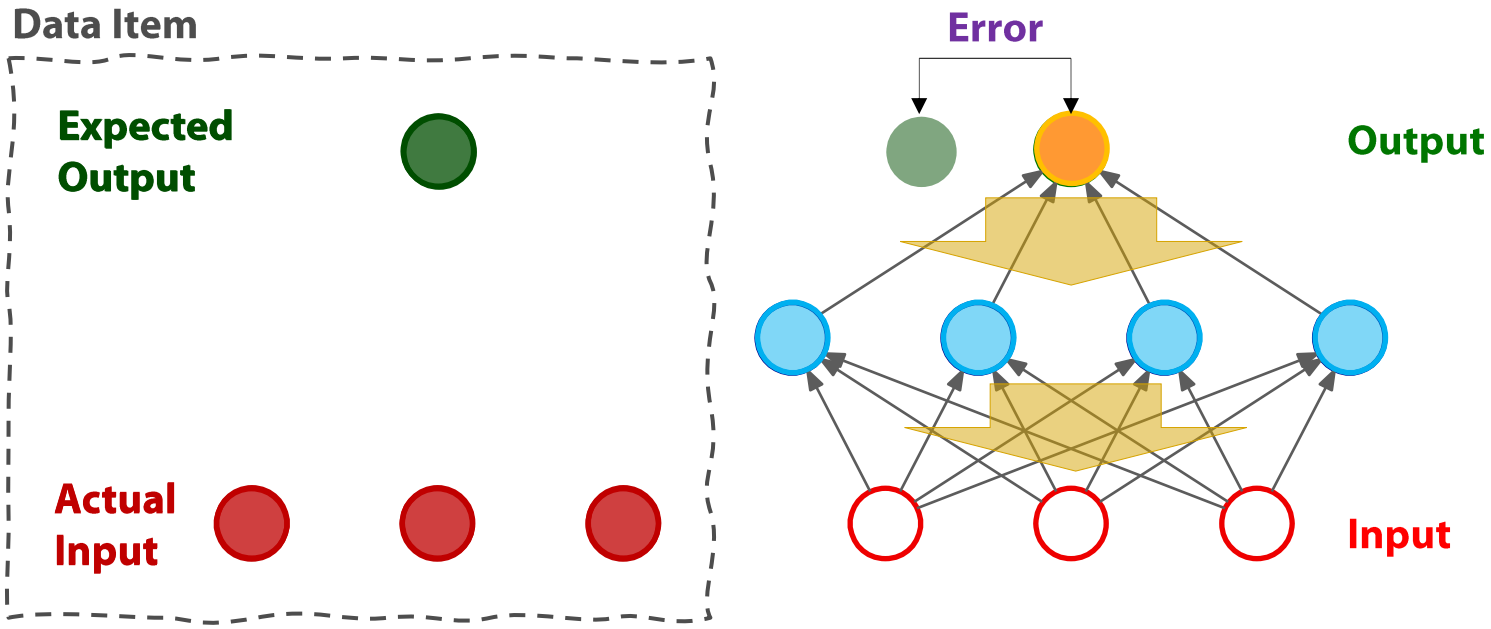

Input data
are fed as input

Input is propagated forward
to compute output

Error is propagated backward
to improve *parameters*

Overall method:

- *show a data item*
- *improve*
- *repeat*
- *a huge amount of times ...*

Machine Learning vs Deep Learning

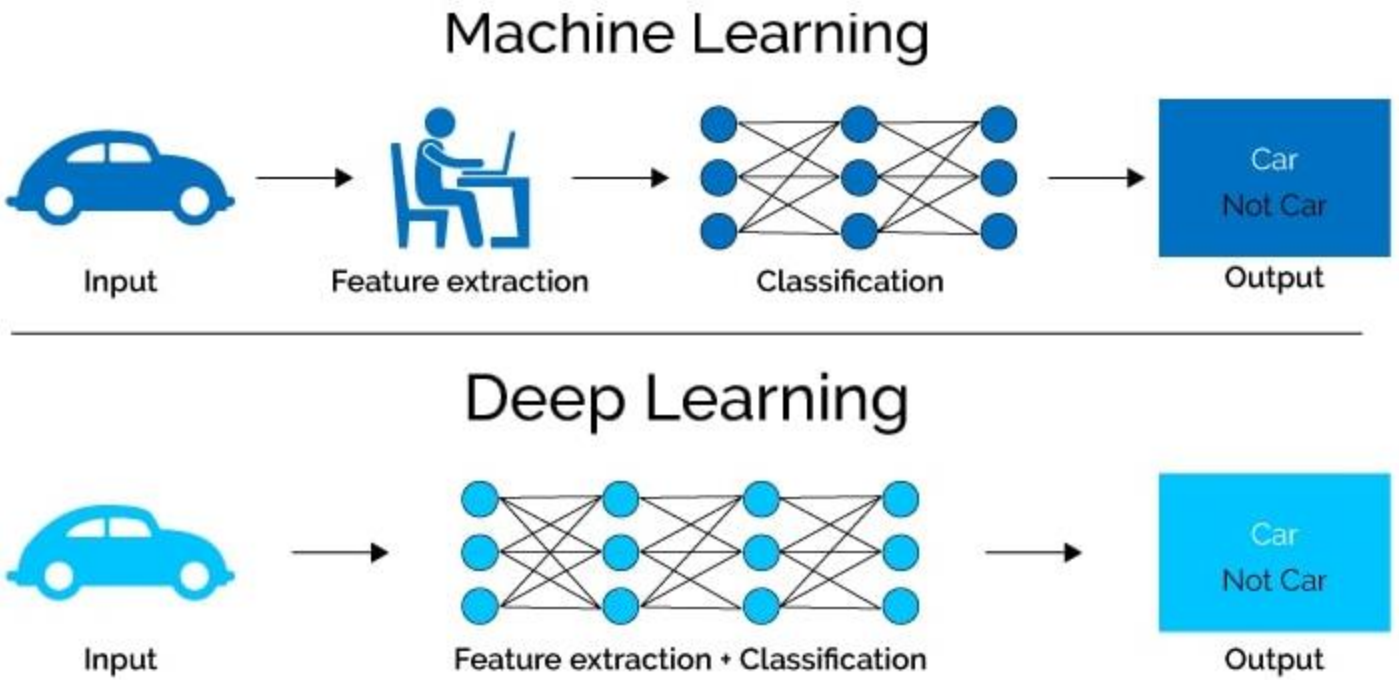

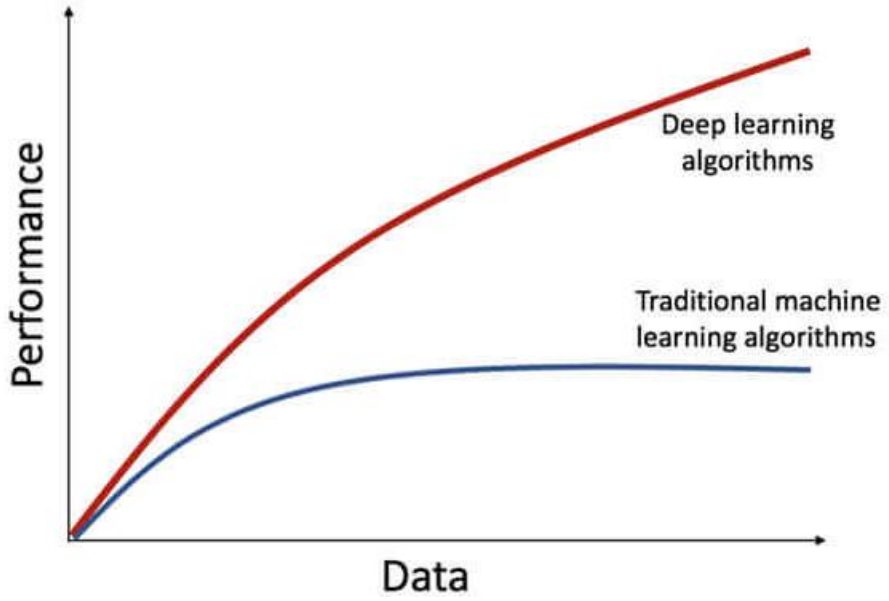

AI revolution?

An artificial neural network is capable to abstract an *associative function* (input-output) from a very large quantity of annotated data

- ***Why they work now (and did not before)***

  Recent progresses are due:
  - better mathematical assessement, for both representation and optimization
  - substantial improvement in parallel computational power
  - introduction of newer and more sophisticated, layered network architectures
  - vast number of numerical parameters
    (ChatGPT has 1.5 billions)

Deep Learning and Artificial Perception

Is there a cat
in this picture?

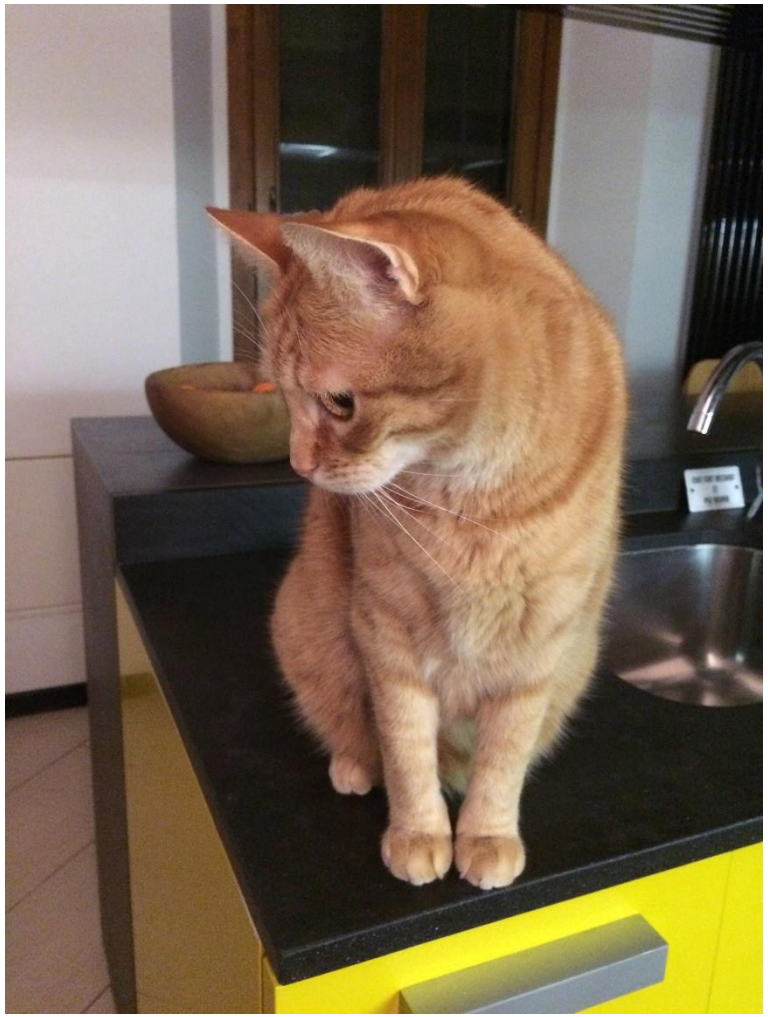

[this is *my* cat, Rabarbaro]

ImageNet Challenge

- The ImageNet Large Scale Visual Recognition Challenge

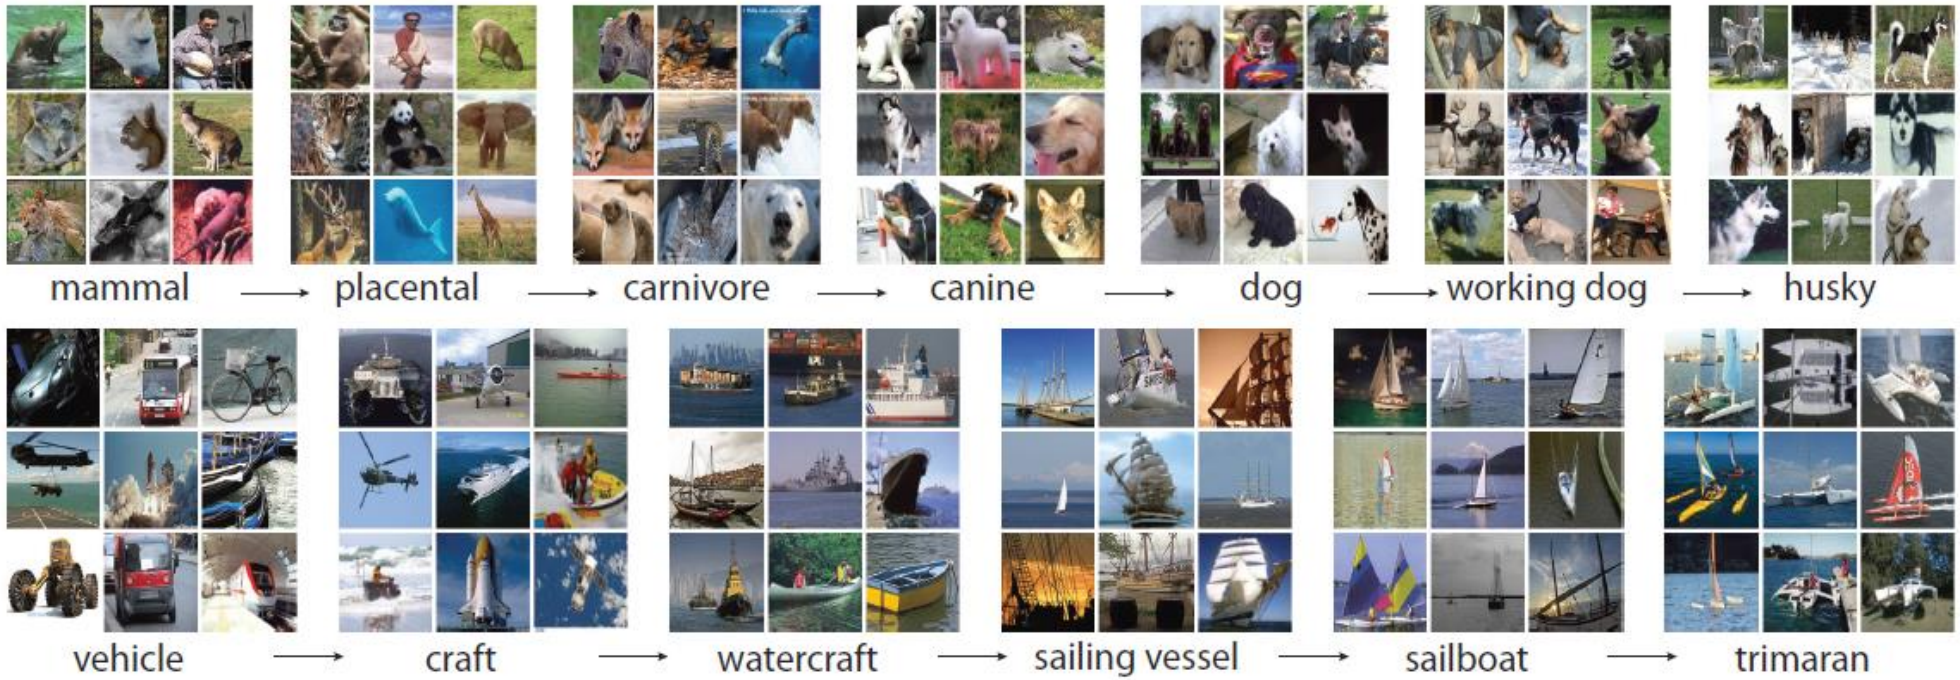

1,461,406 full resolution images

Complex and multiple textual annotation,
hierarchy of 1000 object classes along several dimensions

The image classification challenge is run annually since 2010

[figures from www.nvidia.com]

ImageNet Challenge

- The ImageNet Large Scale Visual Recognition Challenge

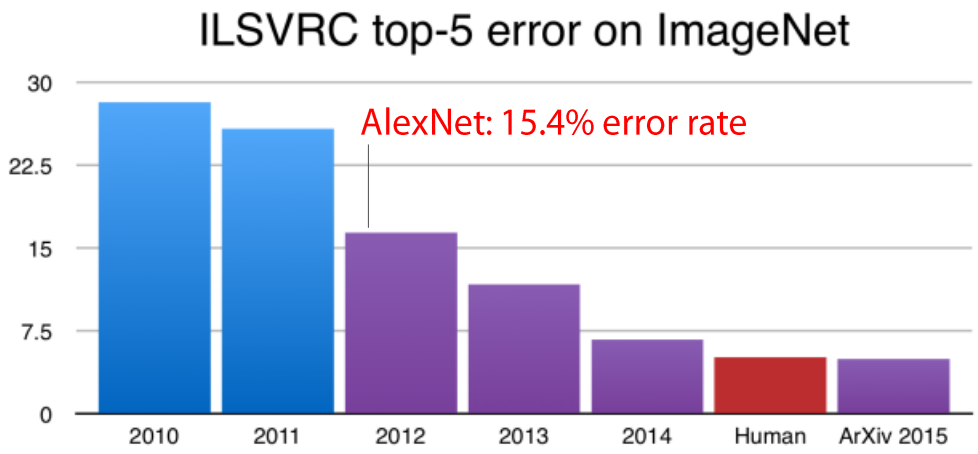

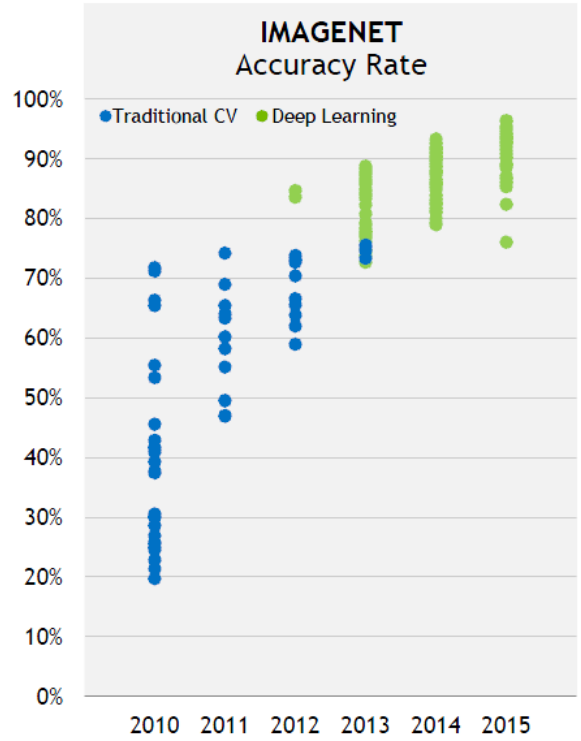

1,461,406 full resolution images

Complex and multiple textual annotation,
hierarchy of 1000 object classes along several dimensions

The image classification challenge was run annually from 2010 to 2017

[figures from www.nvidia.com]

Object (and People) Real-Time Detection

- *Deep Convolutional Networks have evolved since then ...*

Now these system can identify objects
and people in real-time from videos

NOTE:
According to the recent
EU Artificial Intelligence Act,
***remote biometric identification** (RBI)*
in public places will require
a special authorization
(more on this subject, later on)

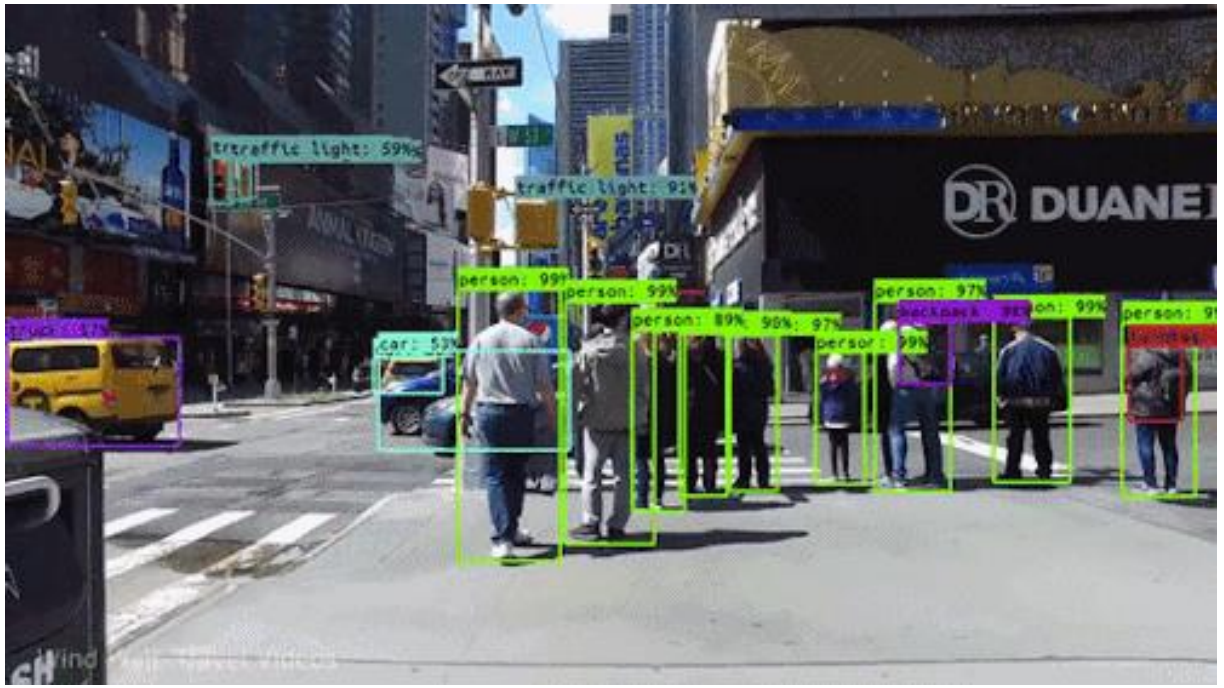

[Image from: https://sgu.ac.id/id/computer-vision-artificial-intelligence-why-is-it-important/]

Generative AI: ChatGPT

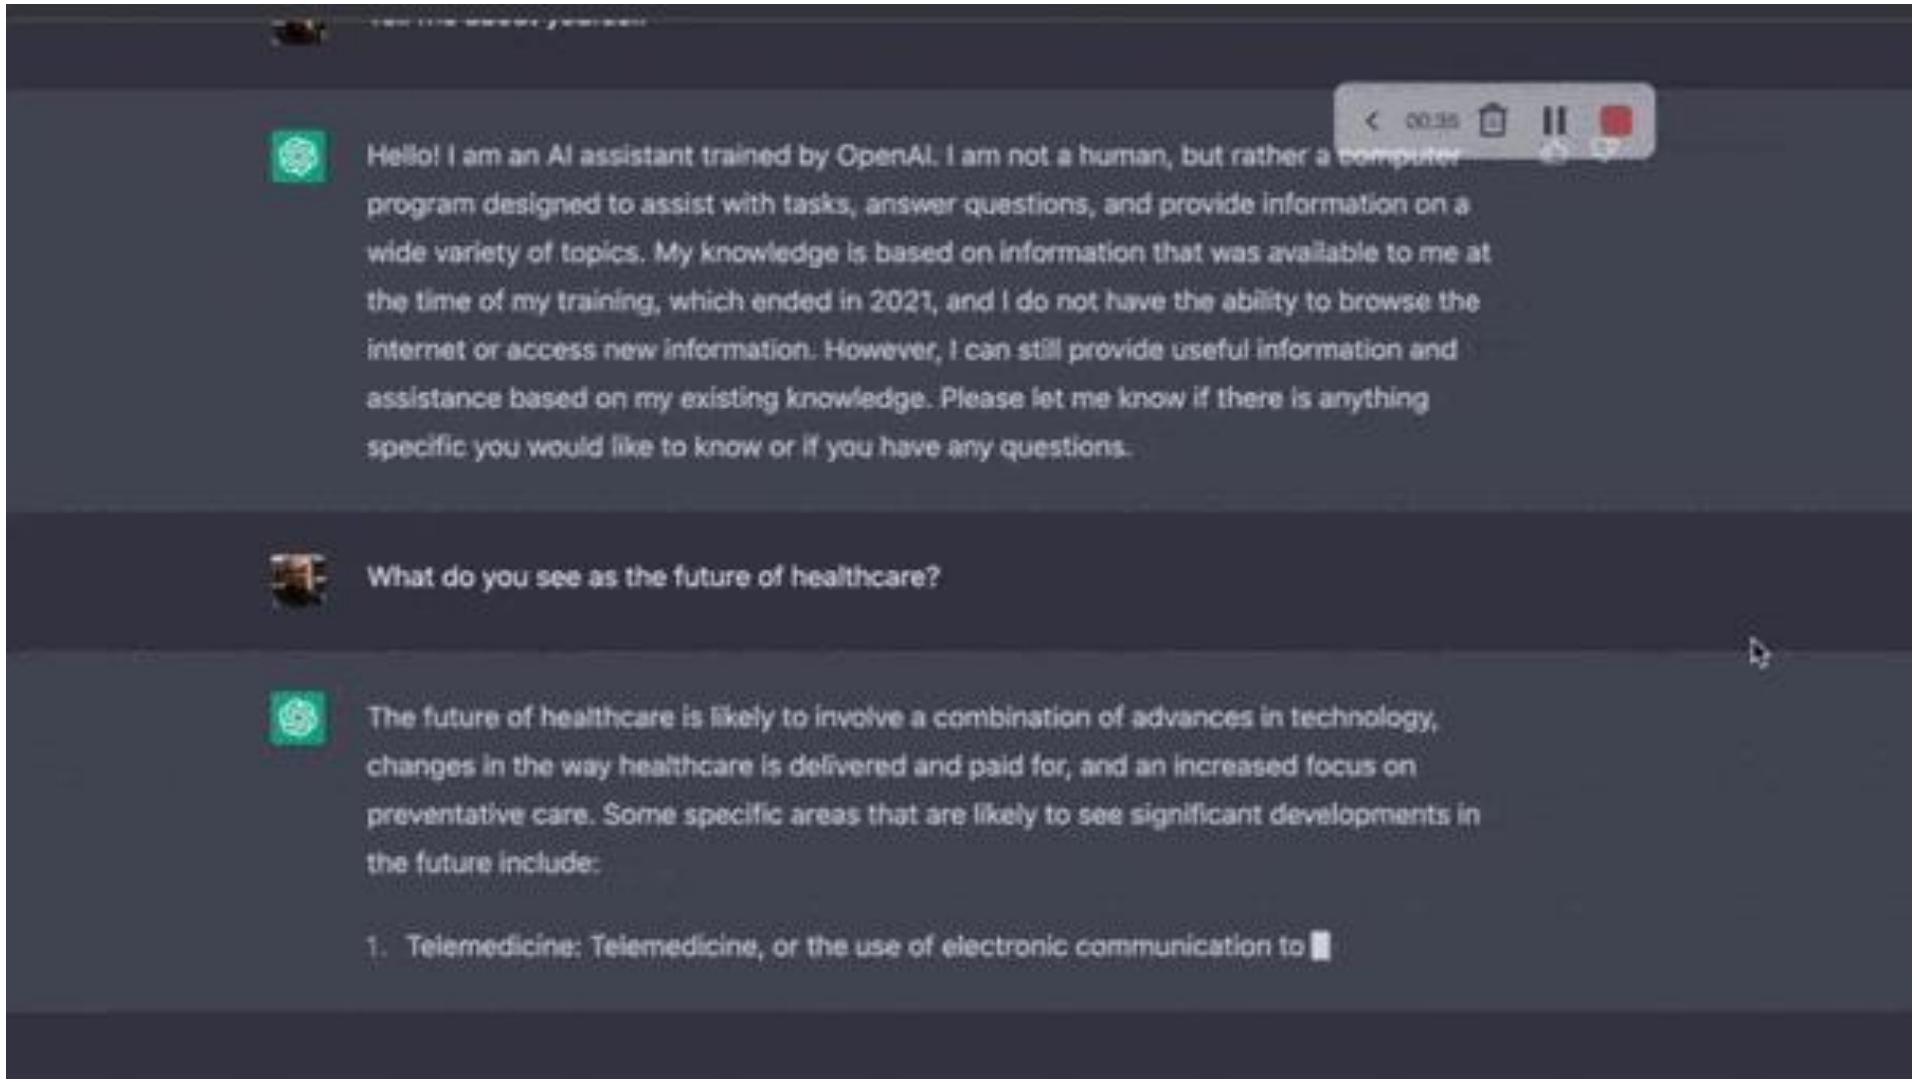

Image from https://info.homecareinnovationforum.com/the-potential-of-chatgpt-in-home-health-care

Language is not the only generative way

- **DALL-E**

  Image generation from text prompts

  «A teapot in the shape of an avocado»

[Image from https://www.nytimes.com/2022/04/06/technology/openai-images-dall-e.html]

Machine/Deep Learning

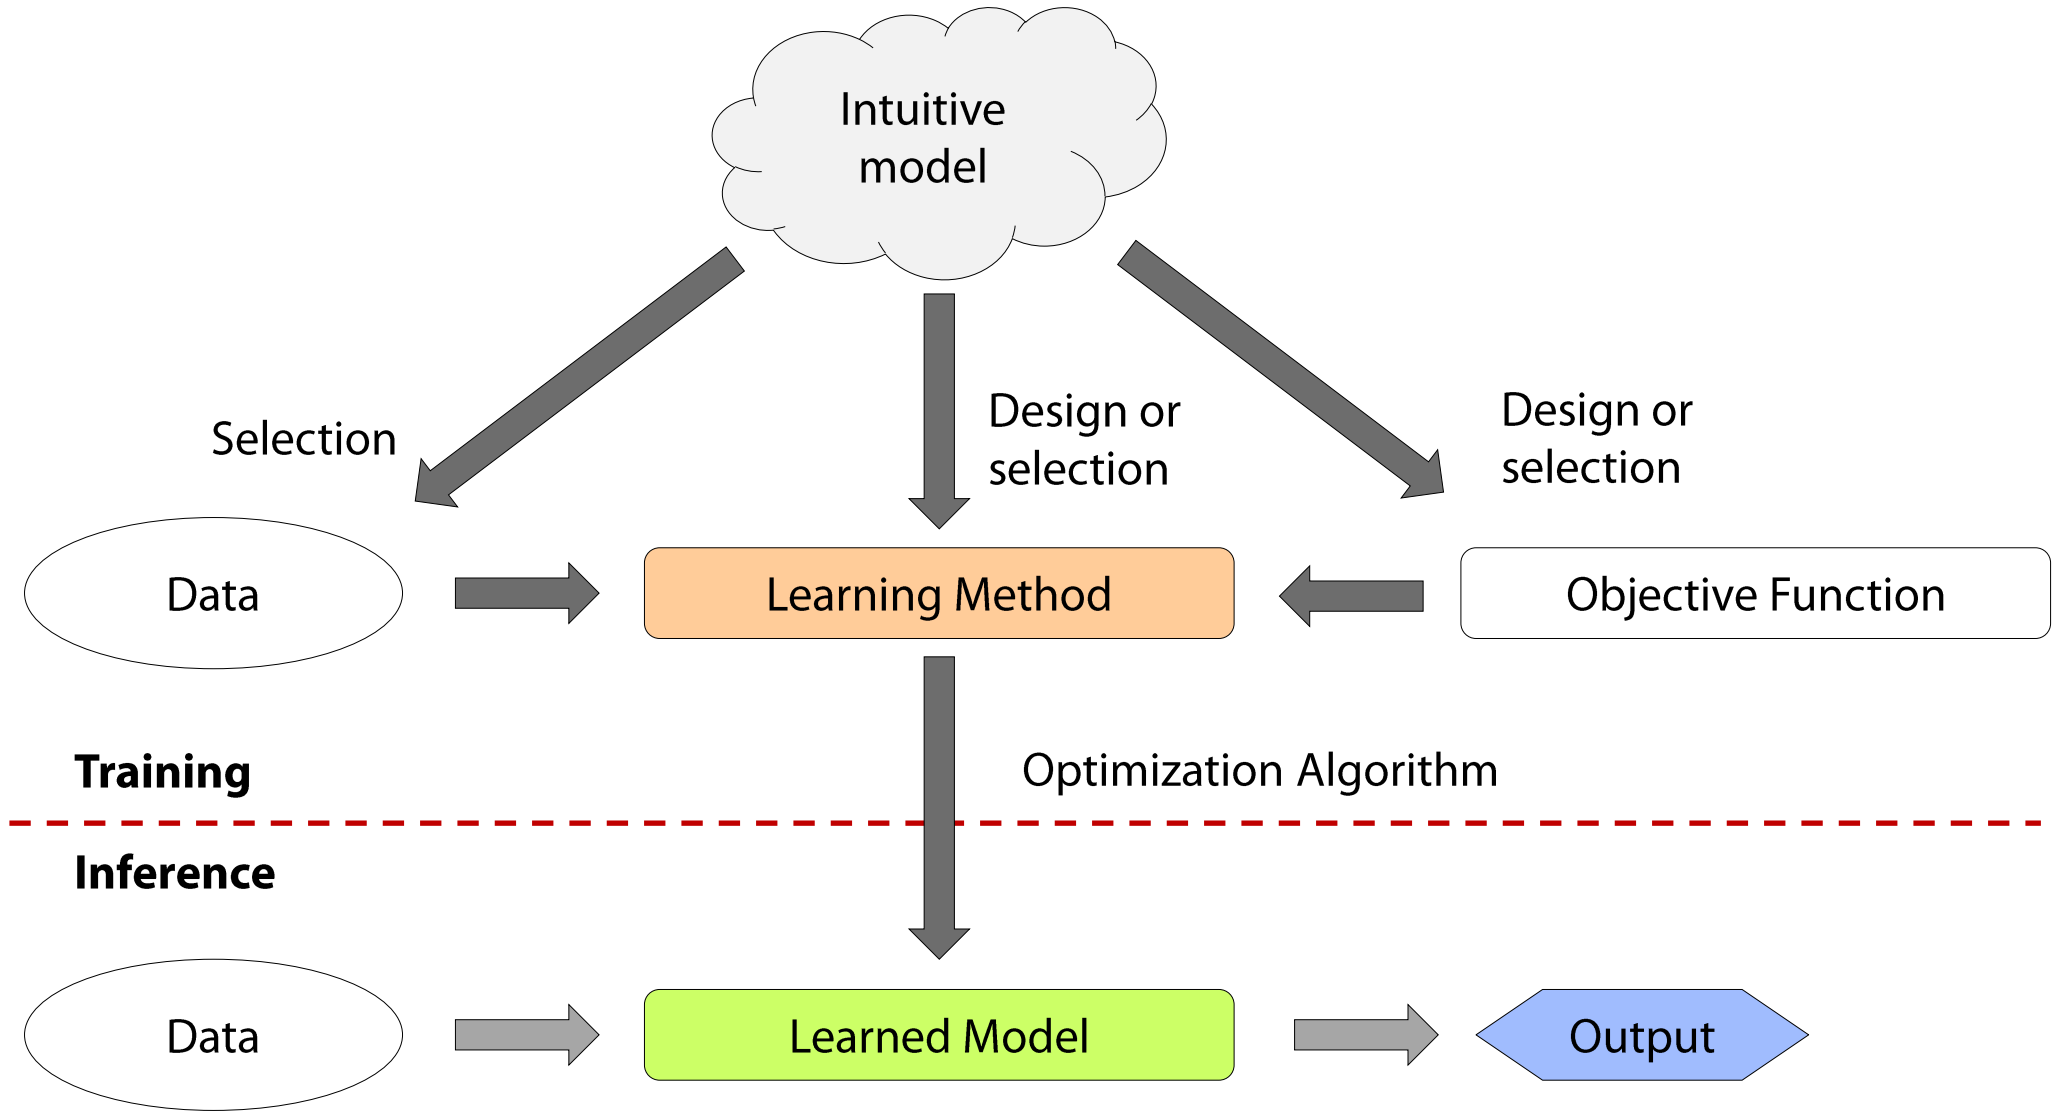

Artificial Mind?

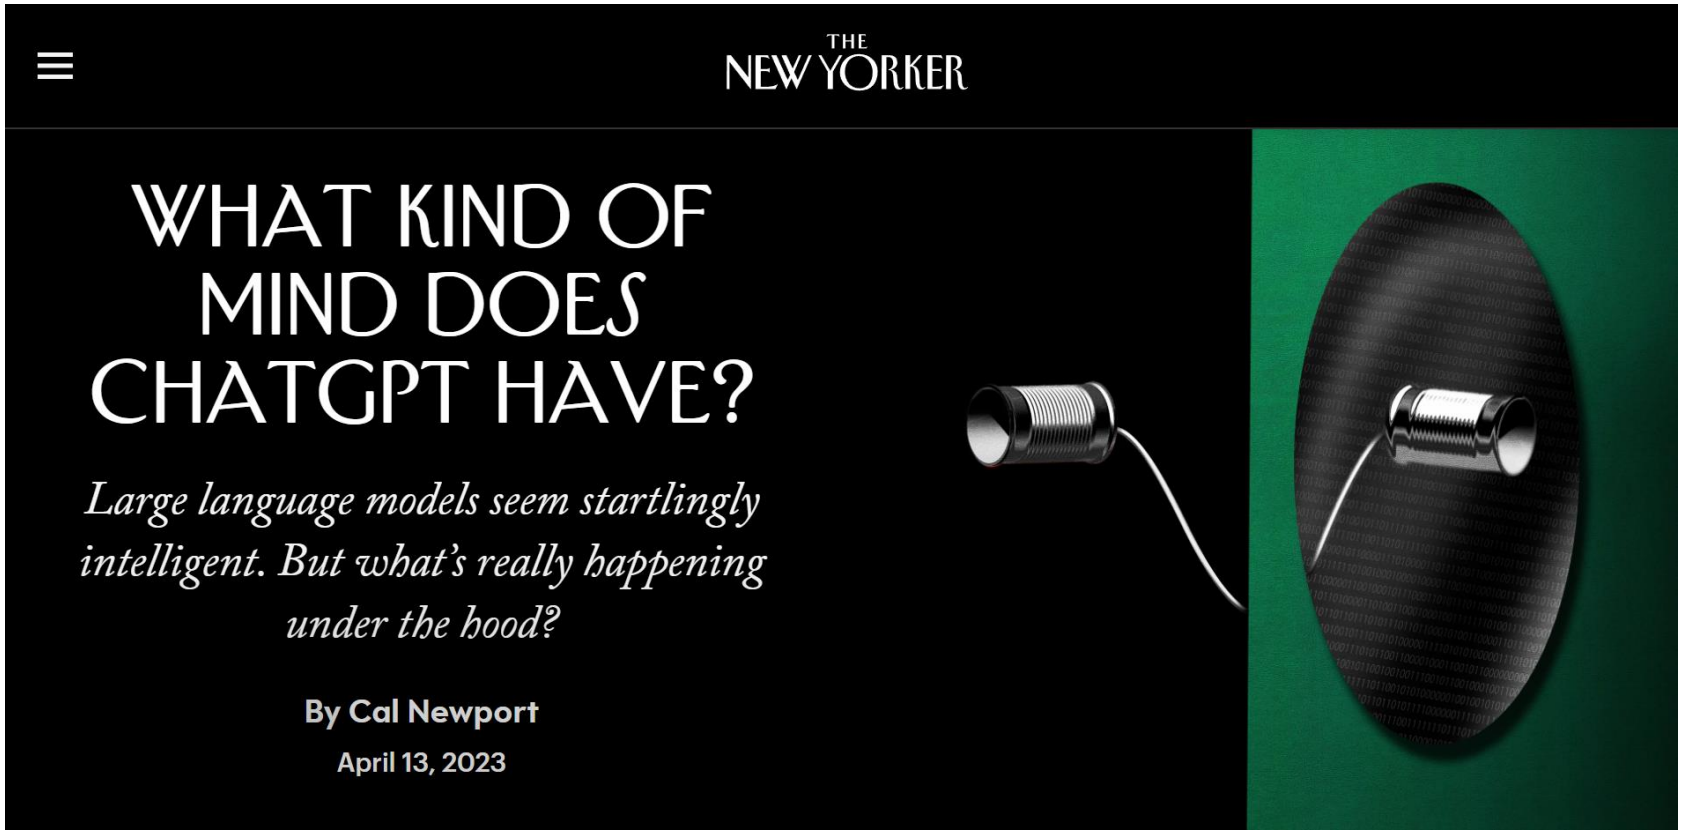

"A system like ChatGPT doesn't create, it imitates. [...] it instead copies, manipulates, and pastes together text that already exists."

"Even if ChatGPT isn't intelligent, couldn't it still take our jobs?"

"Although this ability can generate attention-catching examples, the technology is unlikely in its current form to significantly disrupt the job market. "

[Quote and images from https://www.newyorker.com/science/annals-of-artificial-intelligence/what-kind-of-mind-does-chatgpt-have

"[...] we discover that programs like ChatGPT don't represent an alien intelligence with which we must now learn to coexist; instead, they turn out to run on the well-worn digital logic of **pattern-matching**, pushed to a radically larger scale. "

Regulations: The EU AI Act

A risk-based approach to regulation

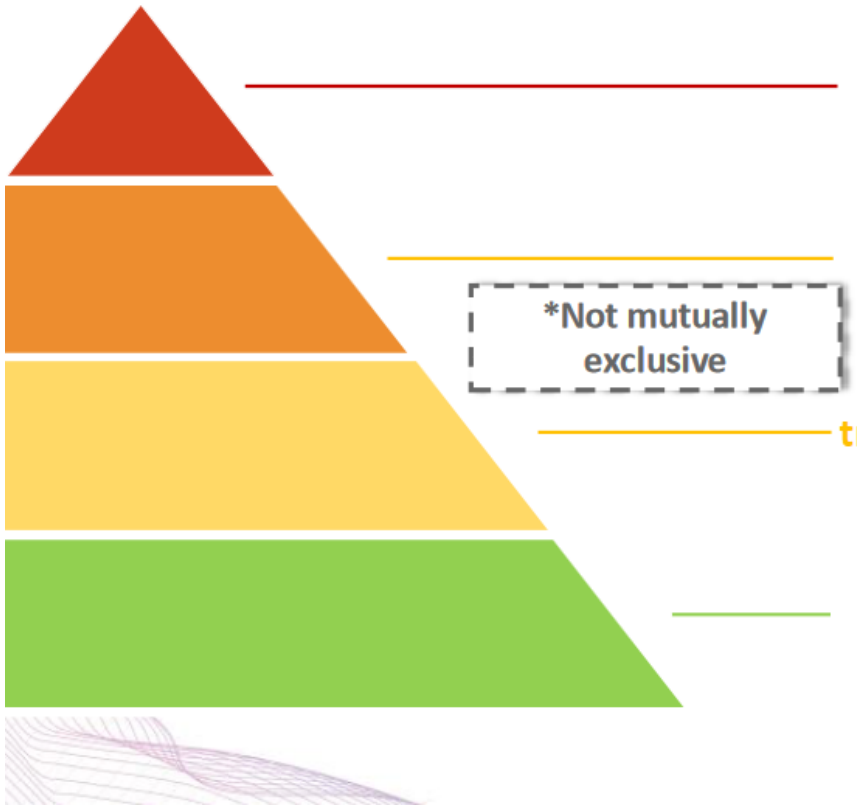

Unacceptable risk
e.g. social scoring
— **Prohibited**

High risk
e.g. recruitment, medical devices
— **Permitted** subject to compliance with AI requirements and ex-ante conformity assessment

AI with specific transparency obligations
'Impersonation' (bots)
— **Permitted** but subject to information/transparency Obligations

Minimal or no risk
— **Permitted** with no restrictions

European Commission

https://www.ceps.eu/wp-content/uploads/2021/04/AI-Presentation-CEPS-Webinar-L.-Sioli-23.4.21.pdf

A few words of caution ...

Observer Bias: Choosing the Right Metaphors

The Economist

Culture | Johnson

Talking about AI in human terms is natural—but wrong

When it comes to artificial intelligence, metaphors are often misleading

IMAGE: NICK LOWNDES

Jun 22nd 2023

Save Share Give

[https://www.economist.com/culture/2023/06/22/talking-about-ai-in-human-terms-is-natural-but-wrong]

AI-Specific Traits

- **Infinitely repeatable**

  Once trained, any AI model performs *deterministic* inferences

- **Completely observable mechanisms**

  Interpretation may be challenging, yet computations *can be observed* to the last bit

- **Duplicable**

  AI systems are *software systems*:
  they may require substantial hardware resources, but they can be replicated at will

- pause -

(This page intentionally left blank)

OK, what is this course about?

"Inference"
(the very idea)

Inference as a Fundamental Trait

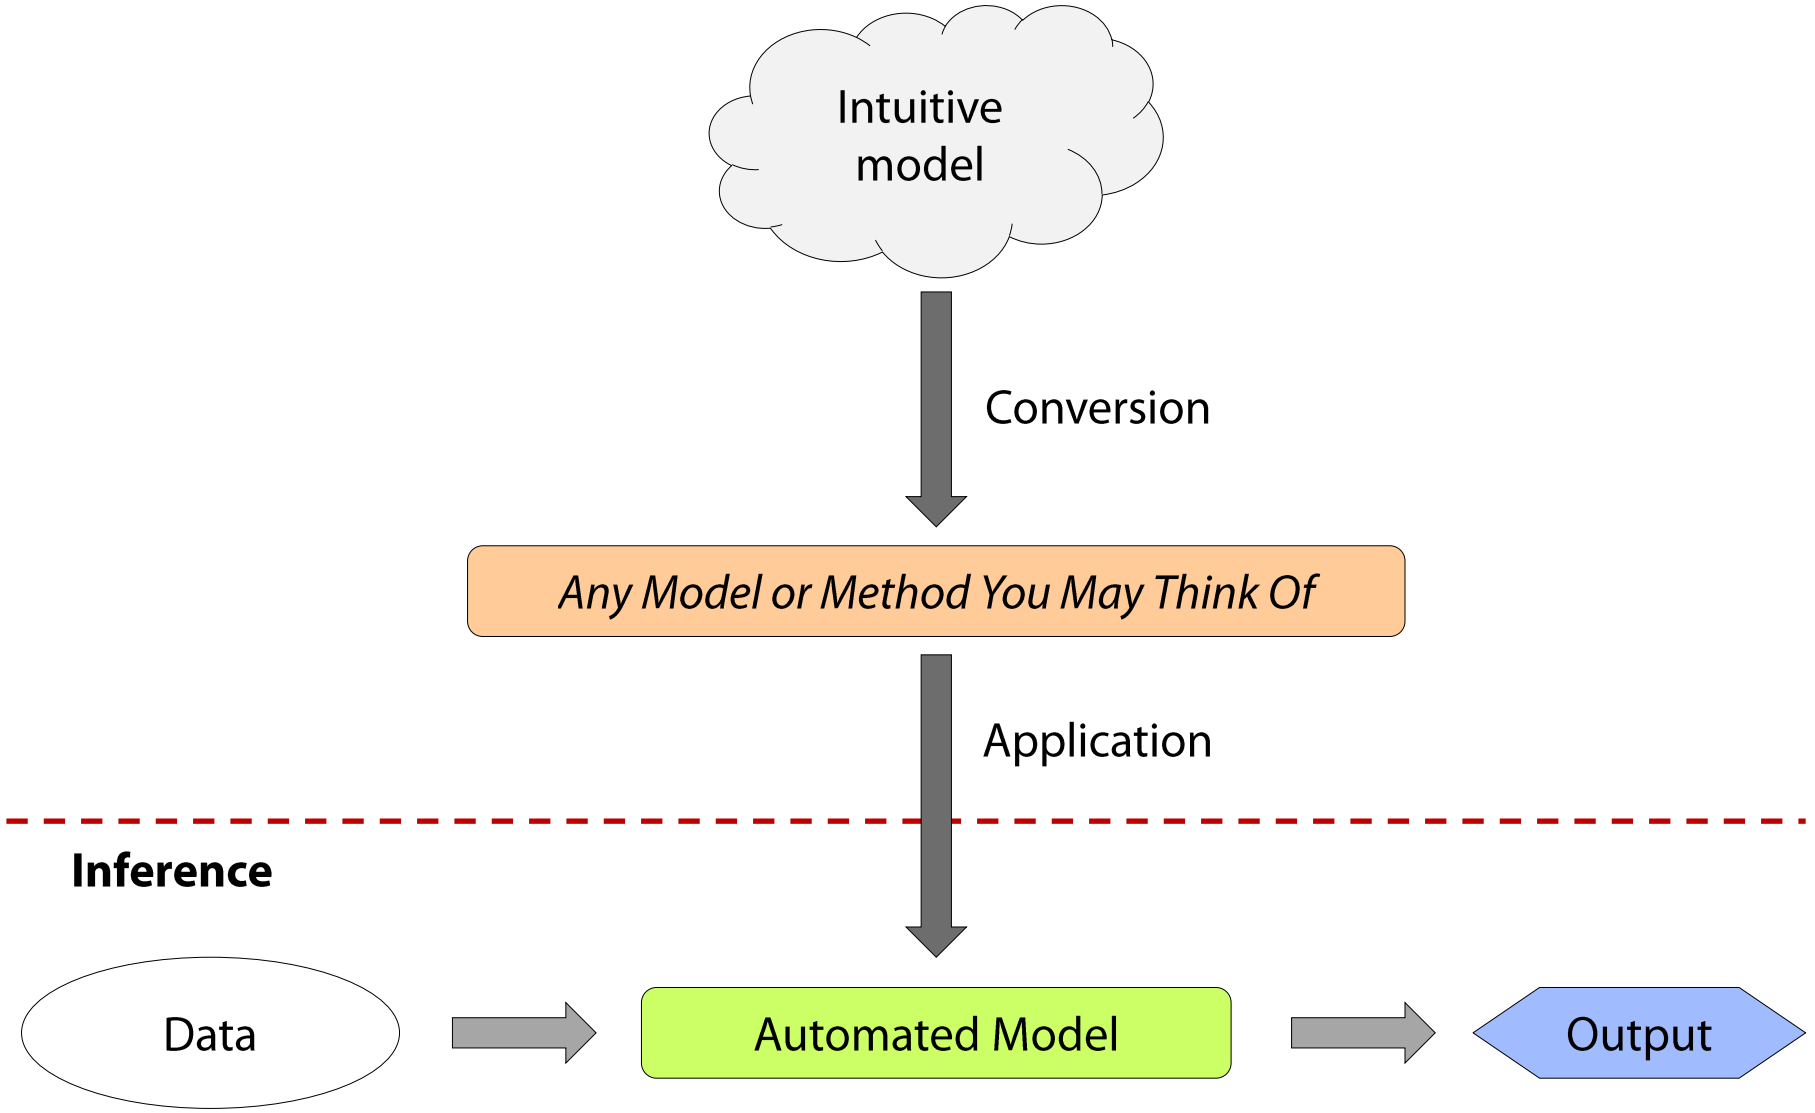

What Is Inference?

- **From a theoretical standpoint**

  What are the *connections* between input and output, in theory?
  (Can it be 'so precisely described' – see John McCarthy?)

  What are the justifications, strength and limitations of such connections?

  Can these connections be computed, at all?

- **From the point of view of computer engineering**

  To what extent, if any, can inference be translated into a computer *algorithm*?

  Are there practical limitations (computational complexity)?

  If so, what kind of compromises need we adopt to make it computationally feasible?

An aside question:
What is an algorithm, precisely?

Turing Machine (A. Turing, 1937)

- Informal description *(more to come, later on)*

  An infinite **tape**, made up of individual **cells**
  Each cell contains a **symbol**, from a finite **alphabet**

  A **read/write head**, which can move in each direction (one cell at time)

  A **finite state machine**:

  - The machine starts in an *initial state*
  - Each *state transition* is governed by the input symbol and the current state
  - The *next state* is stored into a register
  - The *output* is written to the cell
  - Then the head moves (i.e. *Left*, *None*, *Right*)

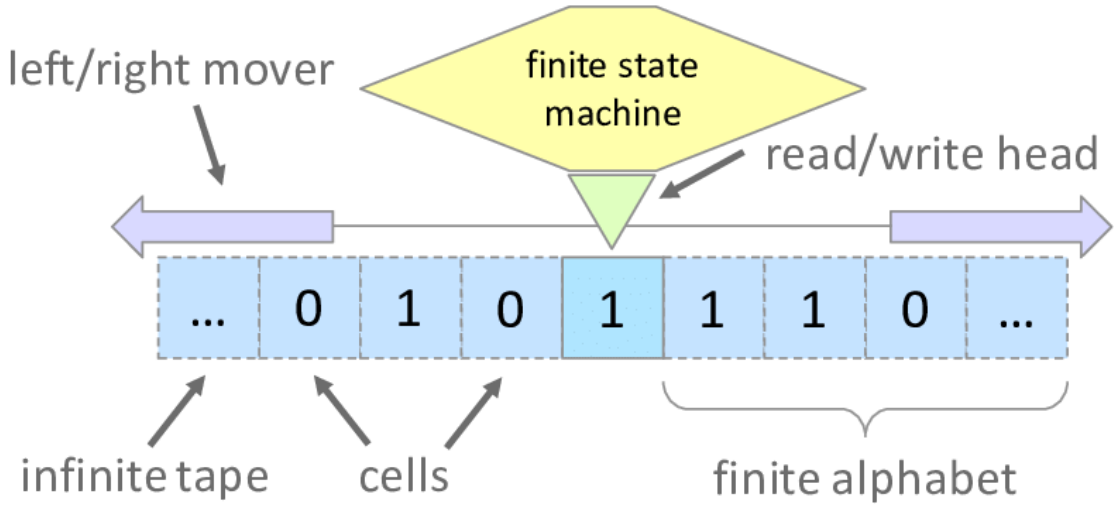

[https://www.researchgate.net/publication/341817215_Quantum_Accelerated_Estimation_of_Algorithmic_Information]

Turing Machine (A. Turing, 1937)

- What is the meaning of this?

  The Turing Machine is a mathematical model
  of a physical computing device

  Any given problem for which there is a Turing Machine
  that computes the solution is clearly computable by a physical machine

  Is the vice-versa also true?

  (*Whenever a problem is computable by a physical machine,
  does it exist a Turing Machine for it?*)

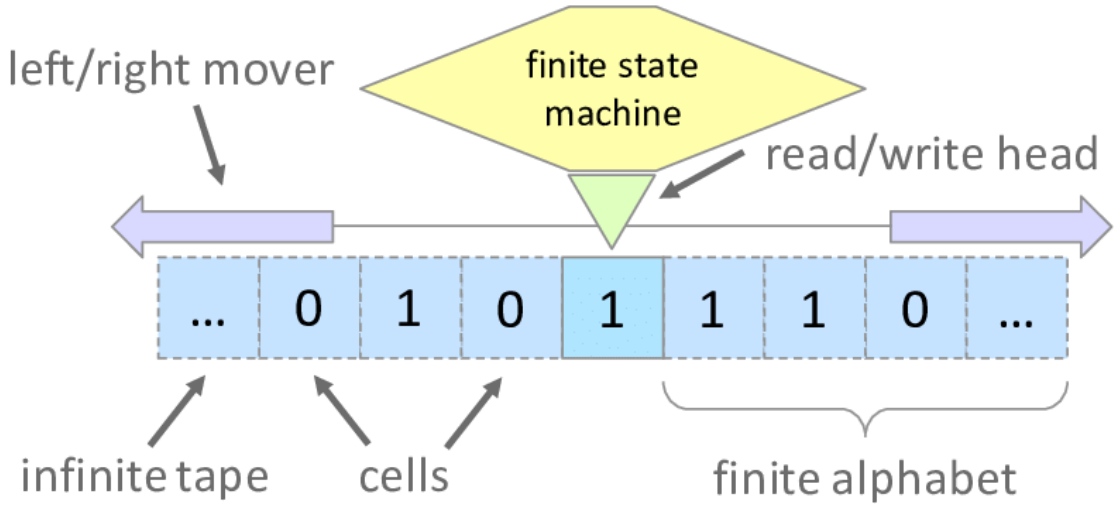

[https://www.researchgate.net/publication/341817215_Quantum_Accelerated_Estimation_of_Algorithmic_Information]

Church-Turing Thesis

- A possible formulation* (from Wikipedia):

  *"Every 'function which would naturally be regarded as computable'
  can be computed by a Turing machine."*

The vagueness in the above sentence gives raise to different interpretations.
One of these (though not entirely equivalent) is (from Wikipedia)*:*

*"Every 'function that could be physically computed'
can be computed by a Turing machine."* **

** Caution: there is no such a thesis in the original writings of either author.*
Its formulation could be extrapolated from both, hence the attribution (made by others)

*** Quantum computation shatters complexity theory, but is (almost) innocuous to computability theory*

Can machines think? (the Turing Test)

- **The Imitation Game** *(textual interactions only)*

"A man (A), a woman (B),
and an interrogator (C) who may be of either sex.

The interrogator stays in a room apart from the other two.

The object for the interrogator is to determine
which of the other two is the man and which is the woman.

The interrogator is allowed to put **questions** to A and B.
[…]

We now ask the question,
'What will happen when a machine takes the part of A in this game?'

[...]

These questions replace our original, 'Can machines think?' "

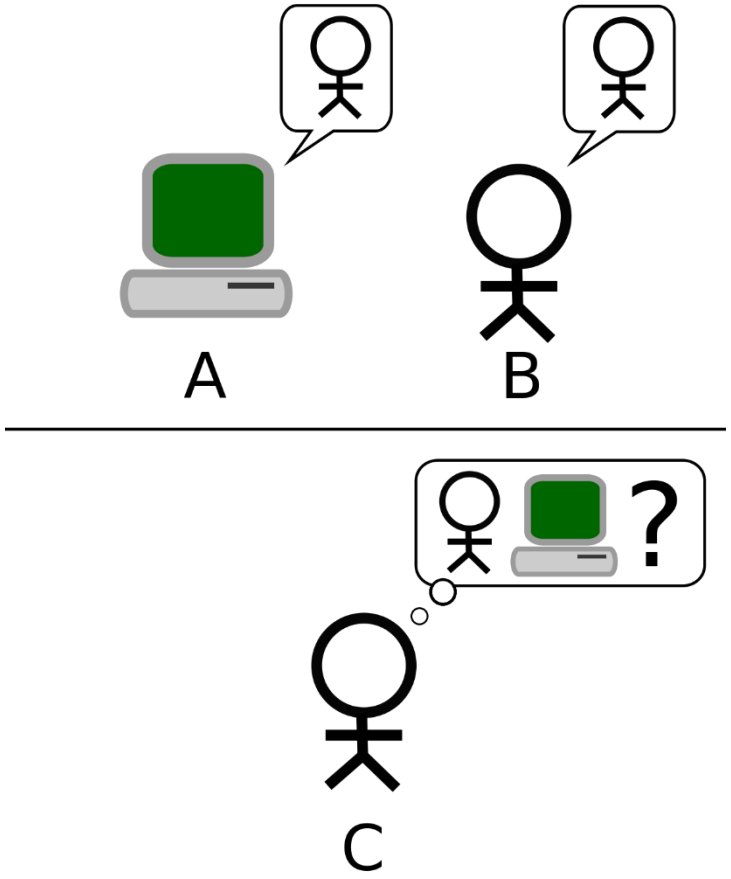

[Image from Wikipedia]

[Turing, A., Computing Machinery and Intelligence, 1950]

An aside question: Are we machines?

[Image from https://xkcd.com/329/]

A short summary of this course

Artificial Intelligence: *Part 1*

- *Reasoning with symbols*

  Propositional logic, first-order logic, logic programming

  Representation: language and semantics *(logic formulae and their meaning)*

  Inference: entailment, a semantic relation among formulae

  Automation: can machines compute entailment?

  Plausible reasoning: beyond the scope of logical entailment

Artificial Intelligence: *Part 2*

- *Reasoning with numbers*

  Machine learning *(in the probabilistic way)*

  Representation: probability, random variables, graphical models
  Inference: answers to queries from joint probability distributions
  Causality: can graphical models be interpreted as *causal* models, as well?

  Supervised learning: learning from *completely observed* and well-formed data items
  Unsupervised learning: when some parts of data items are either *missing* or *hidden*
  Reinforcement learning: learning from experience (even *online*)