Deep Learning

A course about theory & practice

UNIVERSITÀ DI PAVIA

Introduction

Marco Piastra

Prologue: "Deep Learning"?

The Artificial Intelligence Cosmos

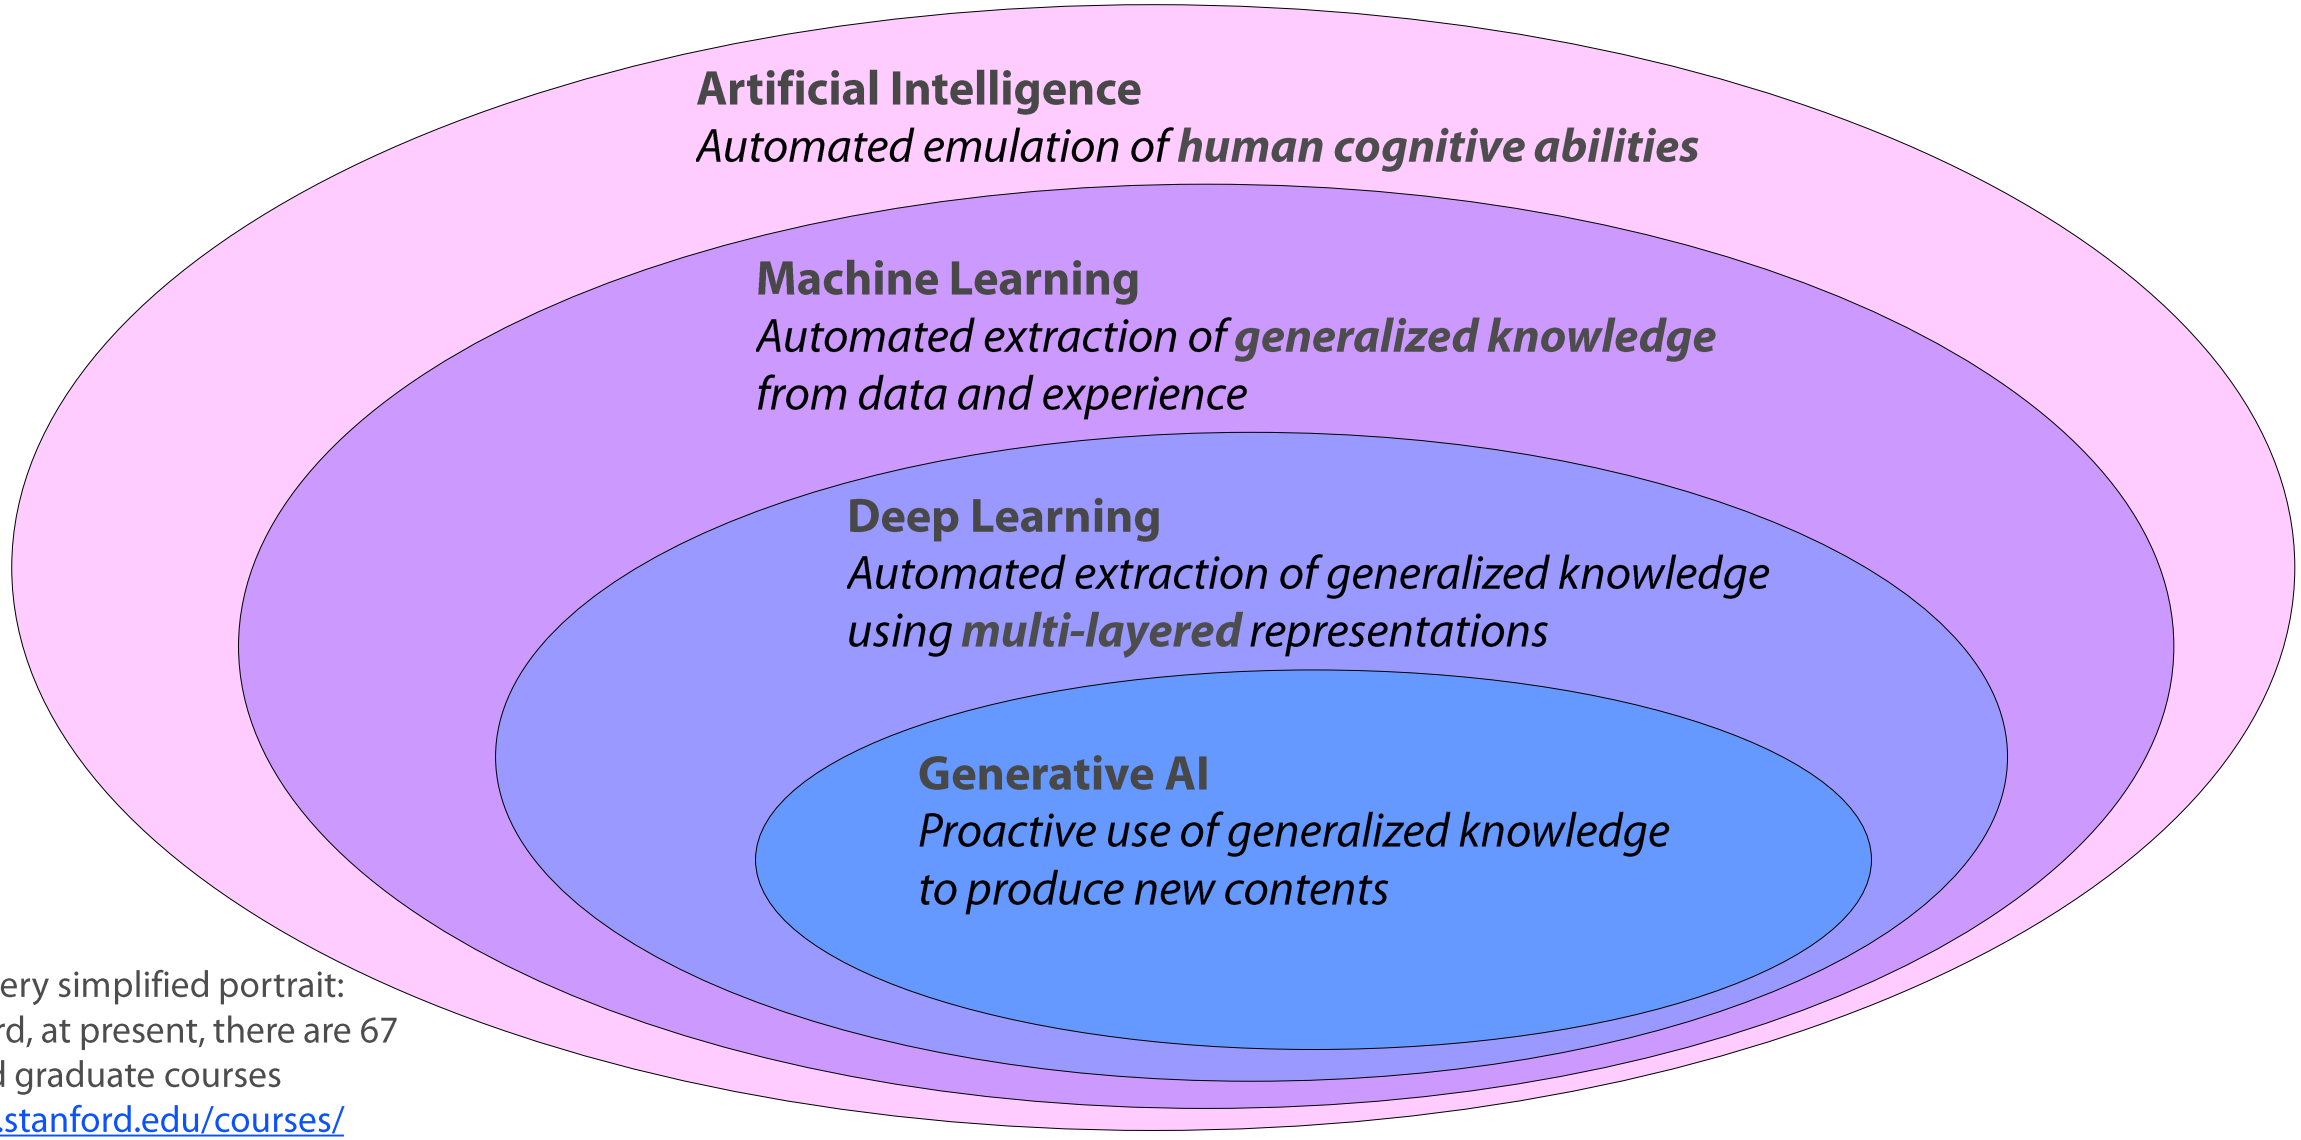

- This is a very simplified portrait: at Stanford, at present, there are 67 AI-related graduate courses https://ai.stanford.edu/courses/

Machine and Deep Learning

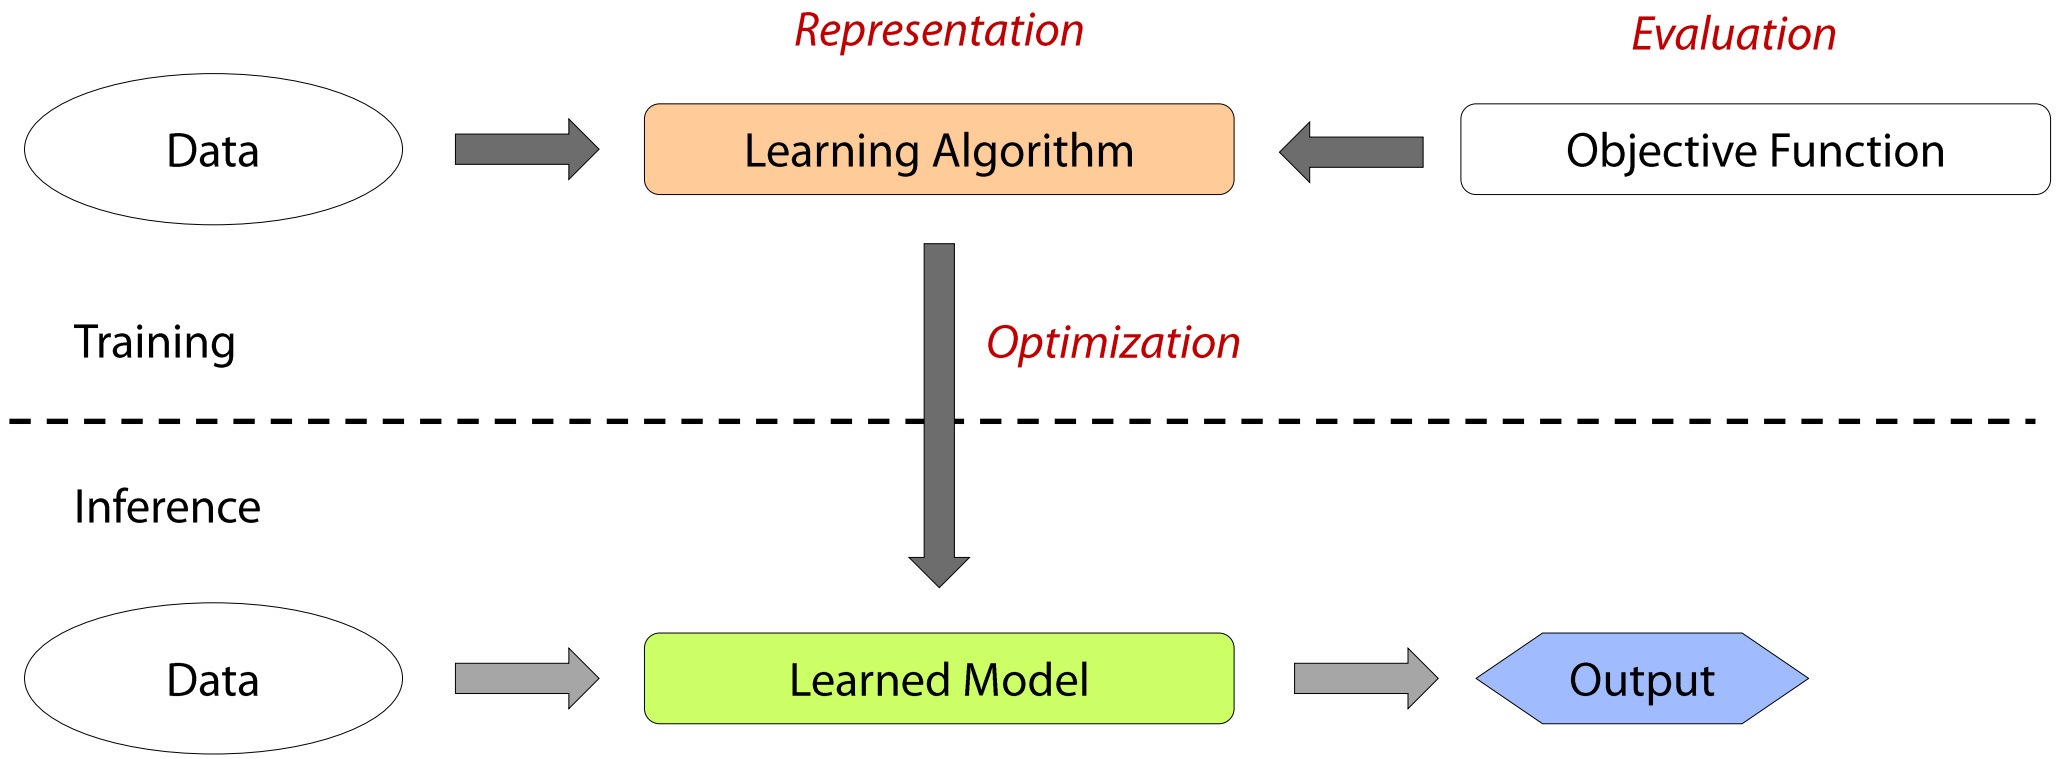

Types of learning algorithms and typical applications

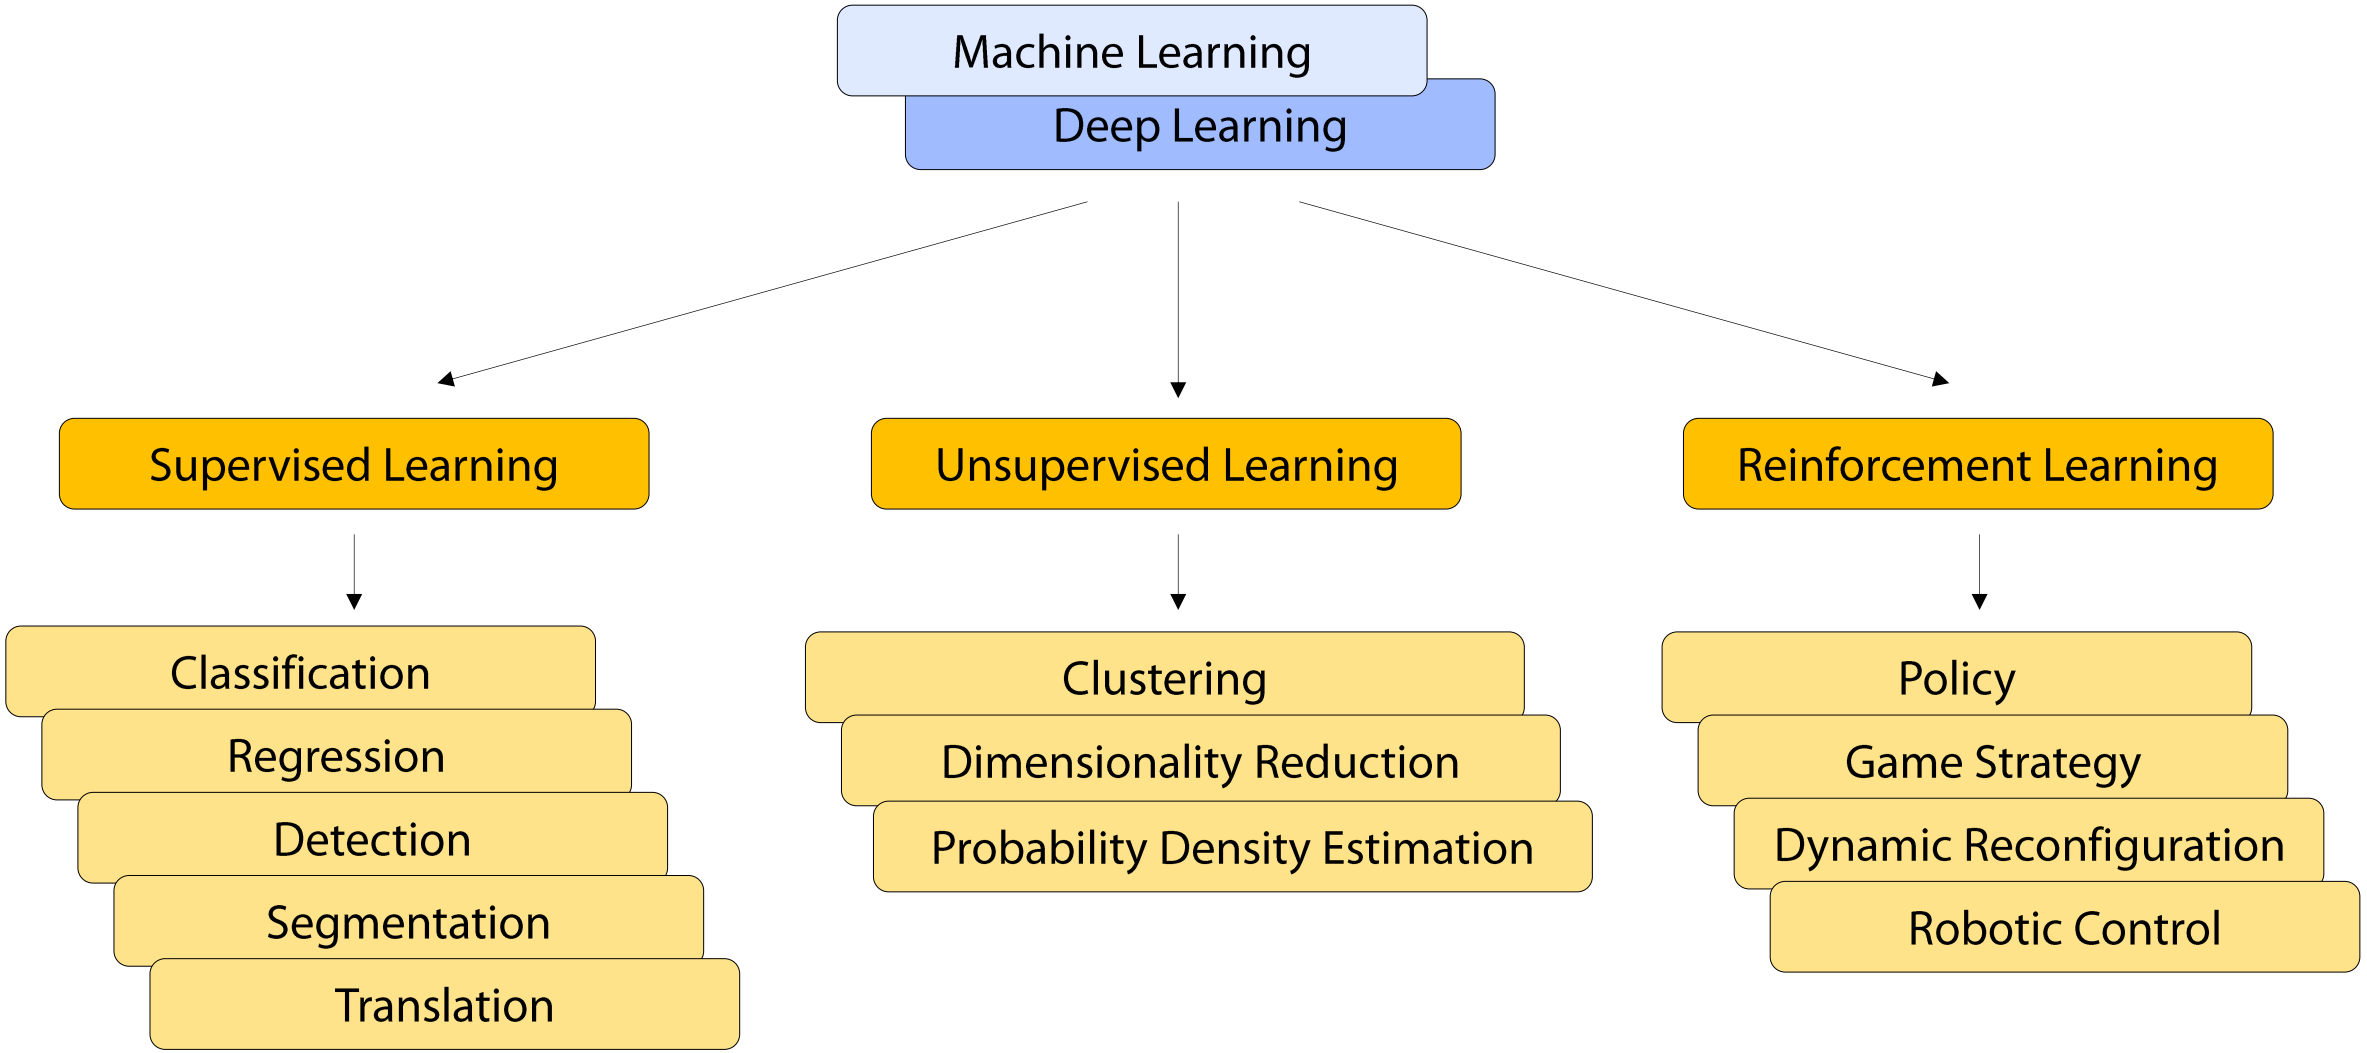

Machine Learning vs Deep Learning

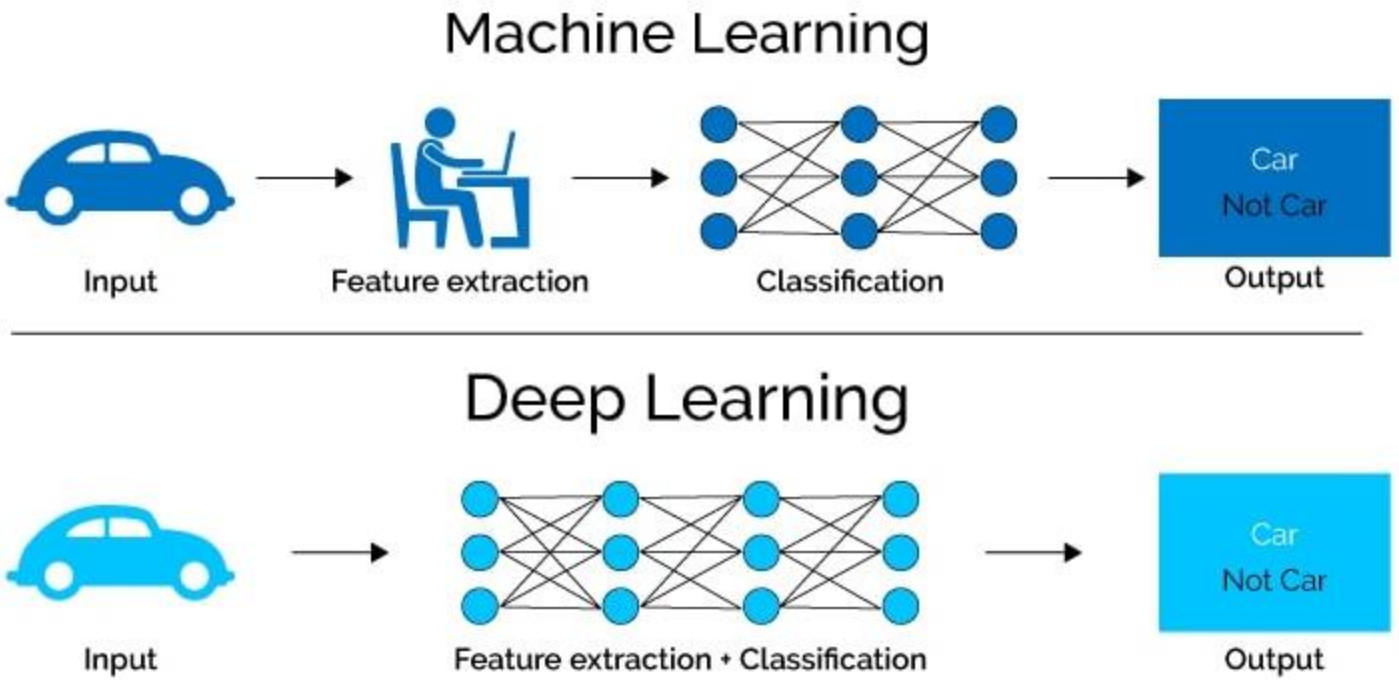

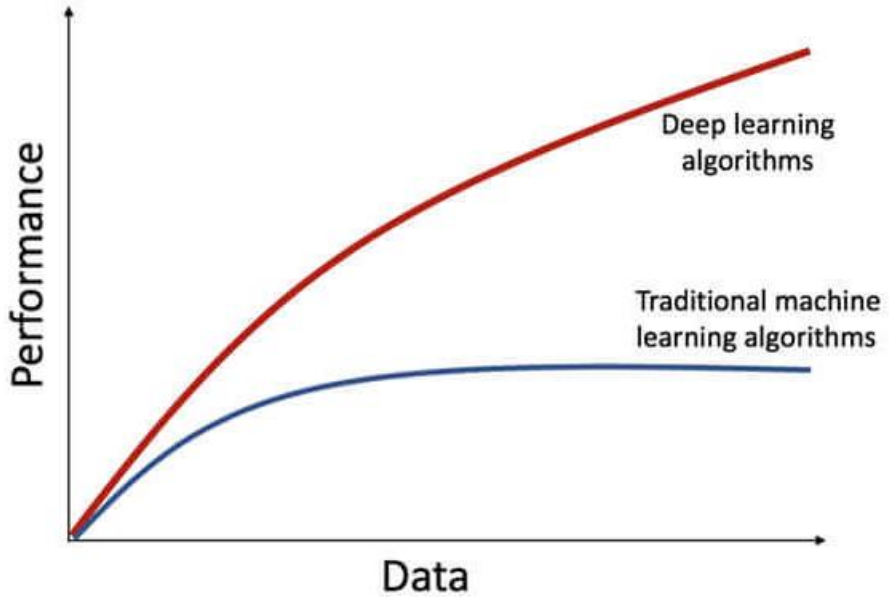

Artificial Neural Networks: synthesizing a (nonlinear) function from data

Artificial Neural Networks

- **An assembly of simple computational units**

  Each unit performs numerical multiplications (of weights) and summations followed by the application of a non-linear filter function

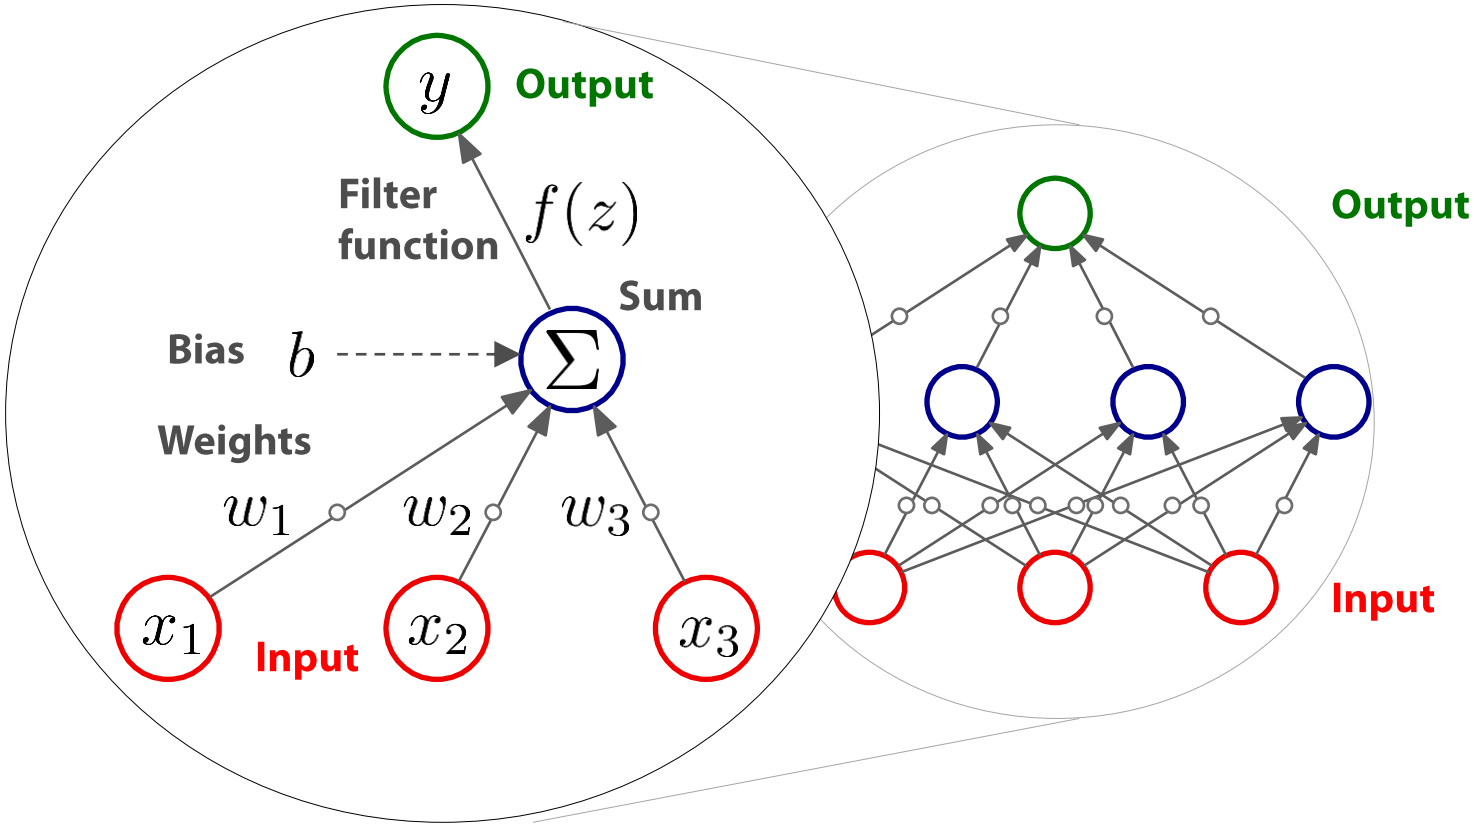

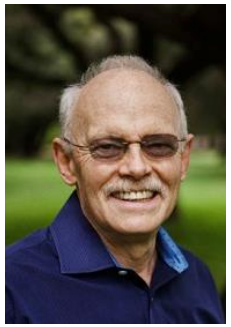

[Images from Wikipedia]

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

Artificial Neural Networks

- **The learning process is an incremental *optimization of numerical parameters***

  Using a vast dataset of input-output pairs (*data items*)

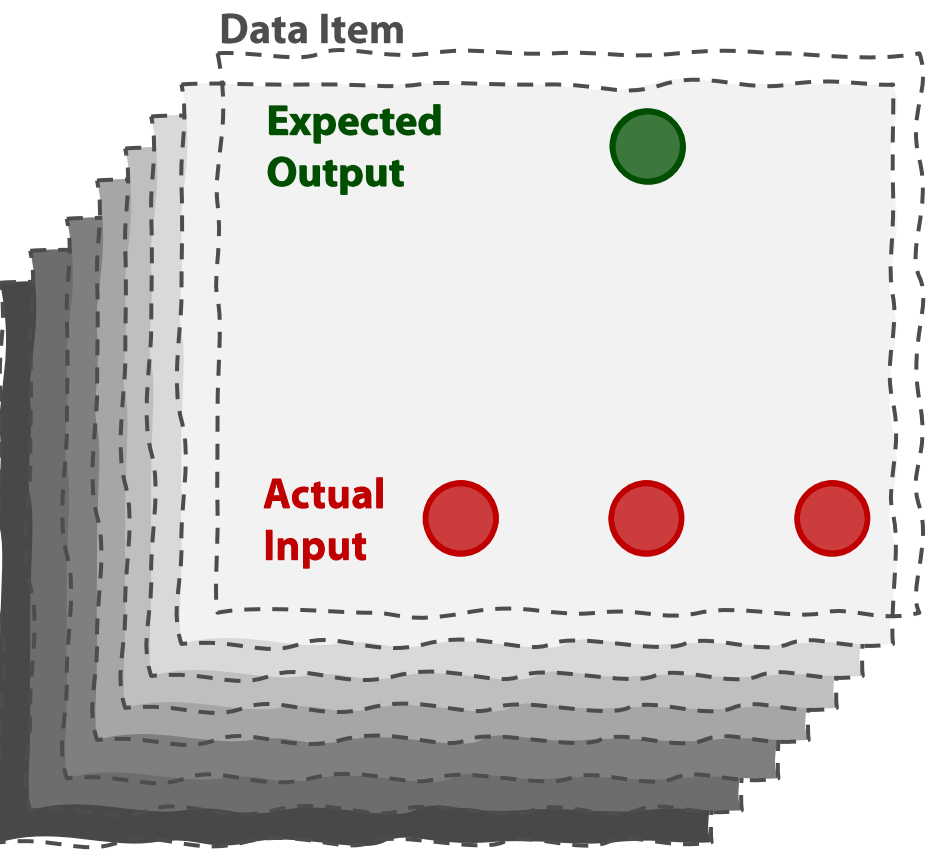

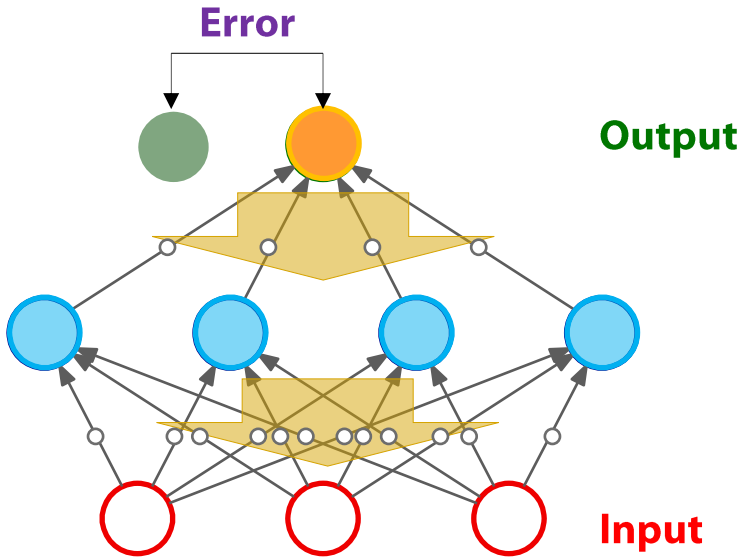

Actual data are presented as **input**

The **input** is propagated upwards to compute the **output**

The **output** is compared with **expected output**

The **error** is propagated downwards to improve ***parameters***

General method:

- *show one data item*
- *improve*
- *repeat*

a huge number of times ...

Artificial Neural Networks

- **From *shallow* to *deep* networks**

  A feed-forward neural network with one hidden layer

$$\tilde{y} = \boldsymbol{w} \cdot g(\boldsymbol{W}^{[1]}\boldsymbol{x} + \boldsymbol{b}^{[1]}) + b$$

  It can approximate any target function

$$y = f^*(\boldsymbol{x}), \quad \boldsymbol{x} \in \mathbb{R}^d$$

  (given enough units and proper parameters)

Deep Learning systems (e.g. PyTorch, TensorFlow) use this representation

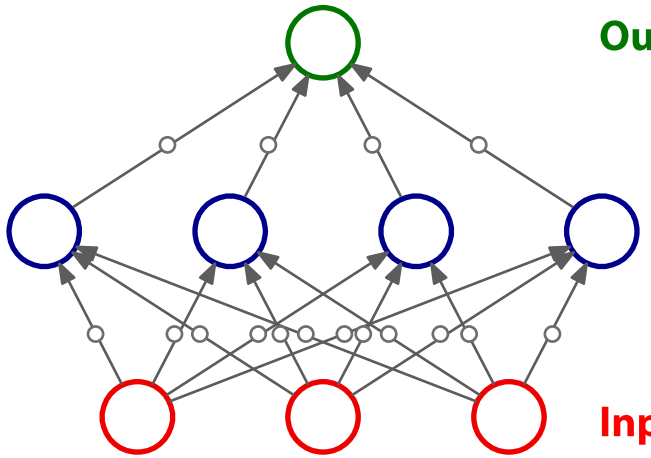

The two representations are equivalent

Artificial Neural Networks

Under mild mathematical constraints,
an artificial neural network can represent *any non-linear* function

$$\tilde{y} = \boldsymbol{w} \cdot g(\boldsymbol{W}^{[1]}\boldsymbol{x} + \boldsymbol{b}^{[1]}) + b$$

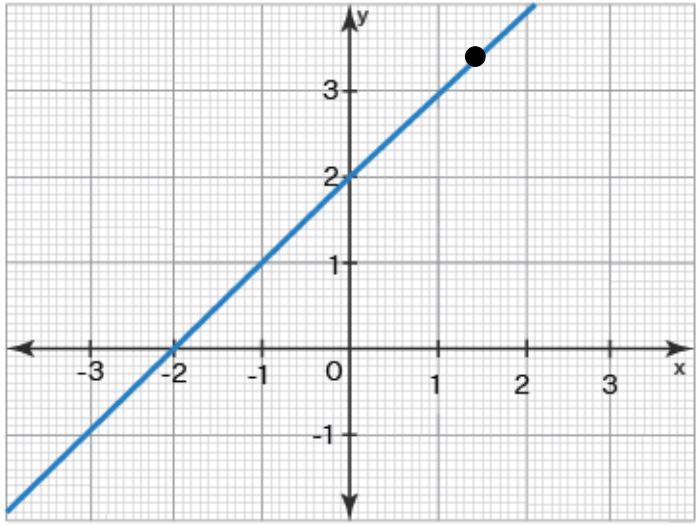

Linear function

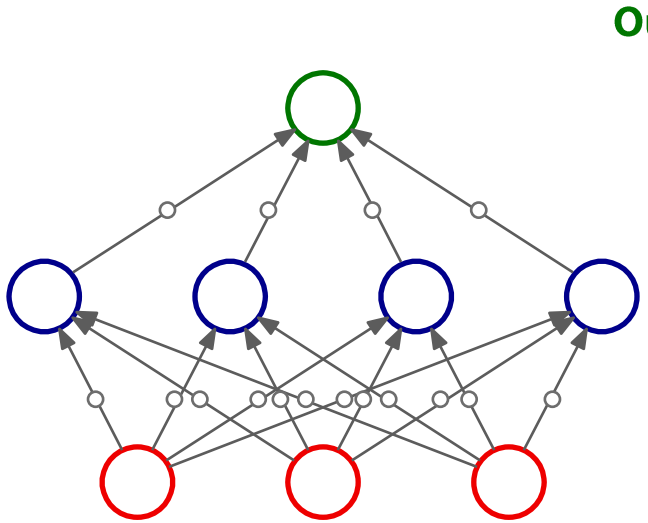

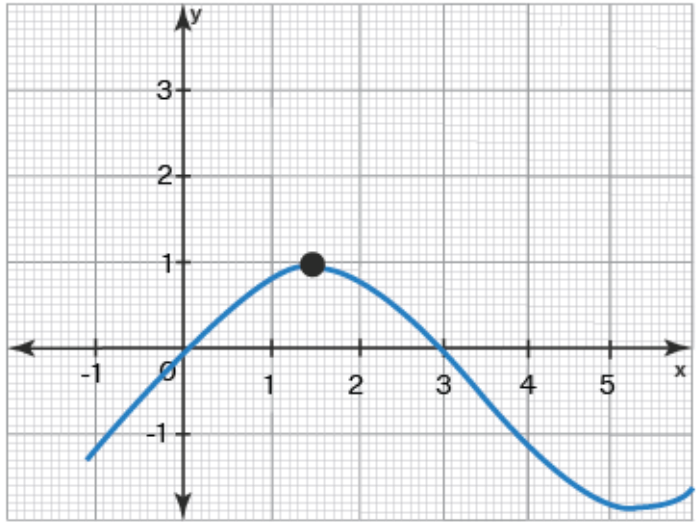

Non-linear function

Artificial Neural Networks

- **From *shallow* to *deep* networks**

  A feed-forward neural network with two hidden layers

$$\tilde{y} = \boldsymbol{w} \cdot g(\boldsymbol{W}^{[2]} g(\boldsymbol{W}^{[1]} \boldsymbol{x} + \boldsymbol{b}^{[1]}) + \boldsymbol{b}^{[2]}) + b$$

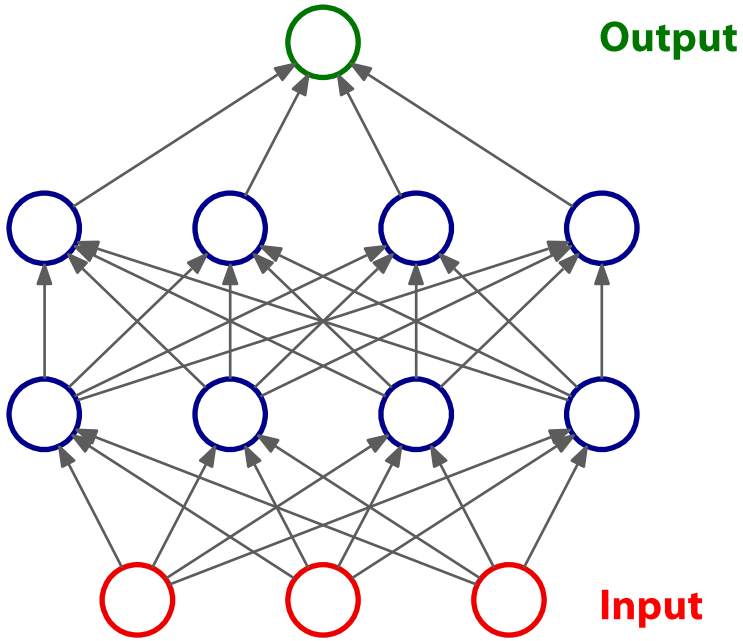

Artificial Neural Networks

- **From *shallow* to *deep* networks**

  A feed-forward neural network with three hidden layers

$$\tilde{y} = \boldsymbol{w} \cdot g(\boldsymbol{W}^{[3]} g(\boldsymbol{W}^{[2]} g(\boldsymbol{W}^{[1]} \boldsymbol{x} + \boldsymbol{b}^{[1]}) + \boldsymbol{b}^{[2]}) + \boldsymbol{b}^{[3]}) + b$$

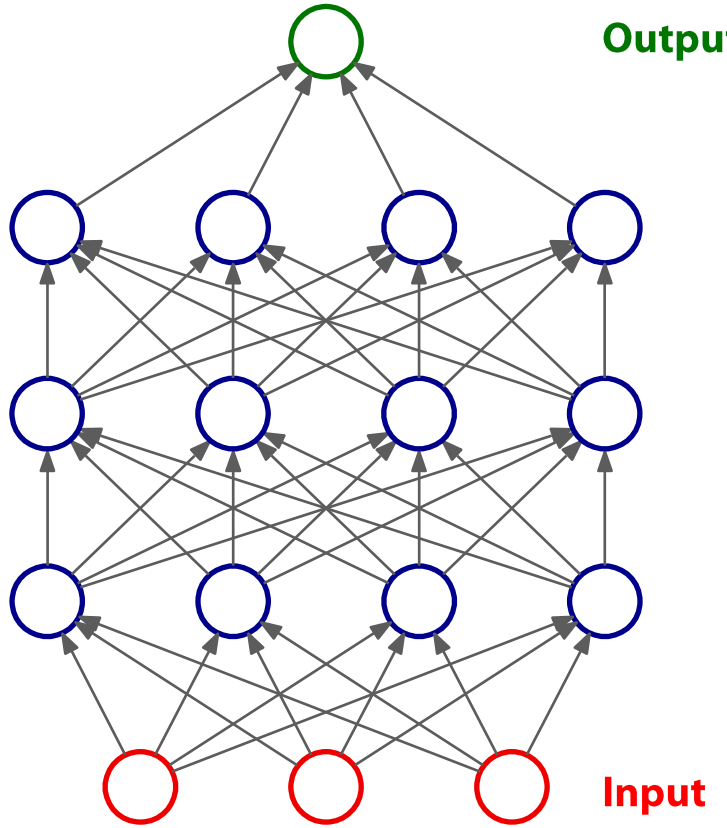

"AI Winter"
"AI Spring"
"AI Summer"?

AI strikes back?

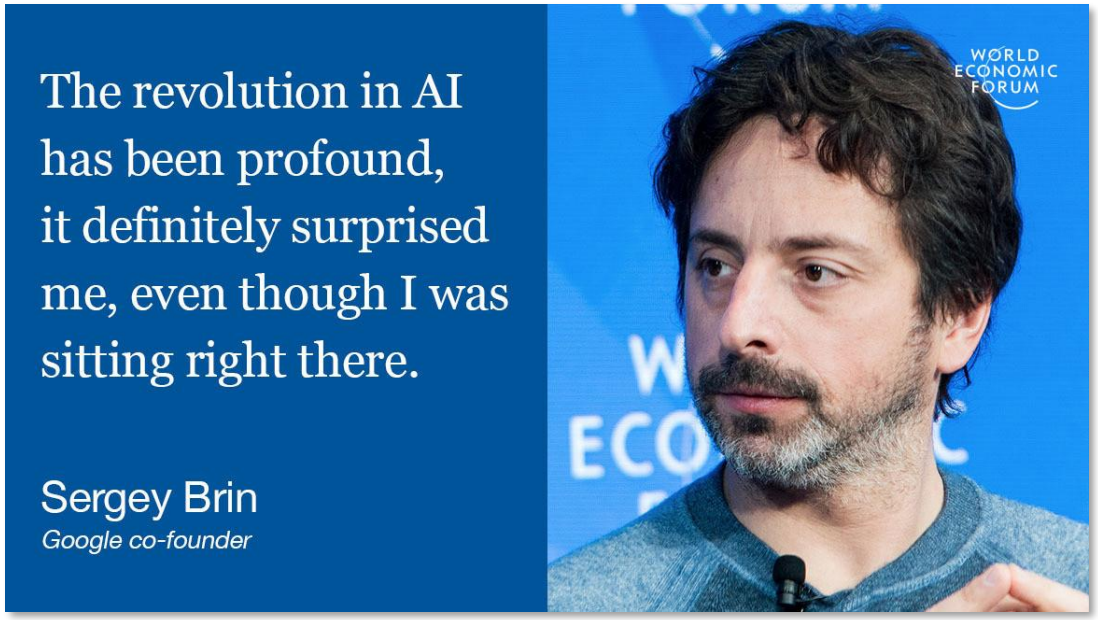

- Sergey Brin [Google Co-Founder, January 2017]

  "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/]

One Giant Leap for Mankind

Is there a cat
in this picture?

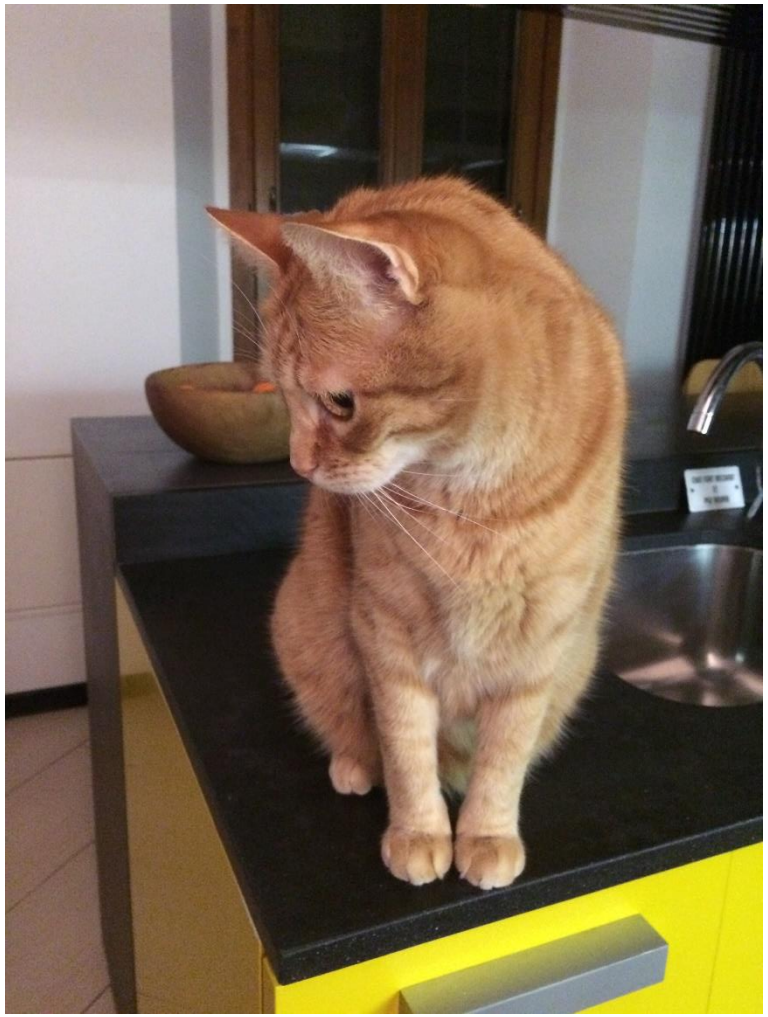

[this is *my* cat, Rabarbaro]

ImageNet Challenge

- The ImageNet Large Scale Visual Recognition Challenge

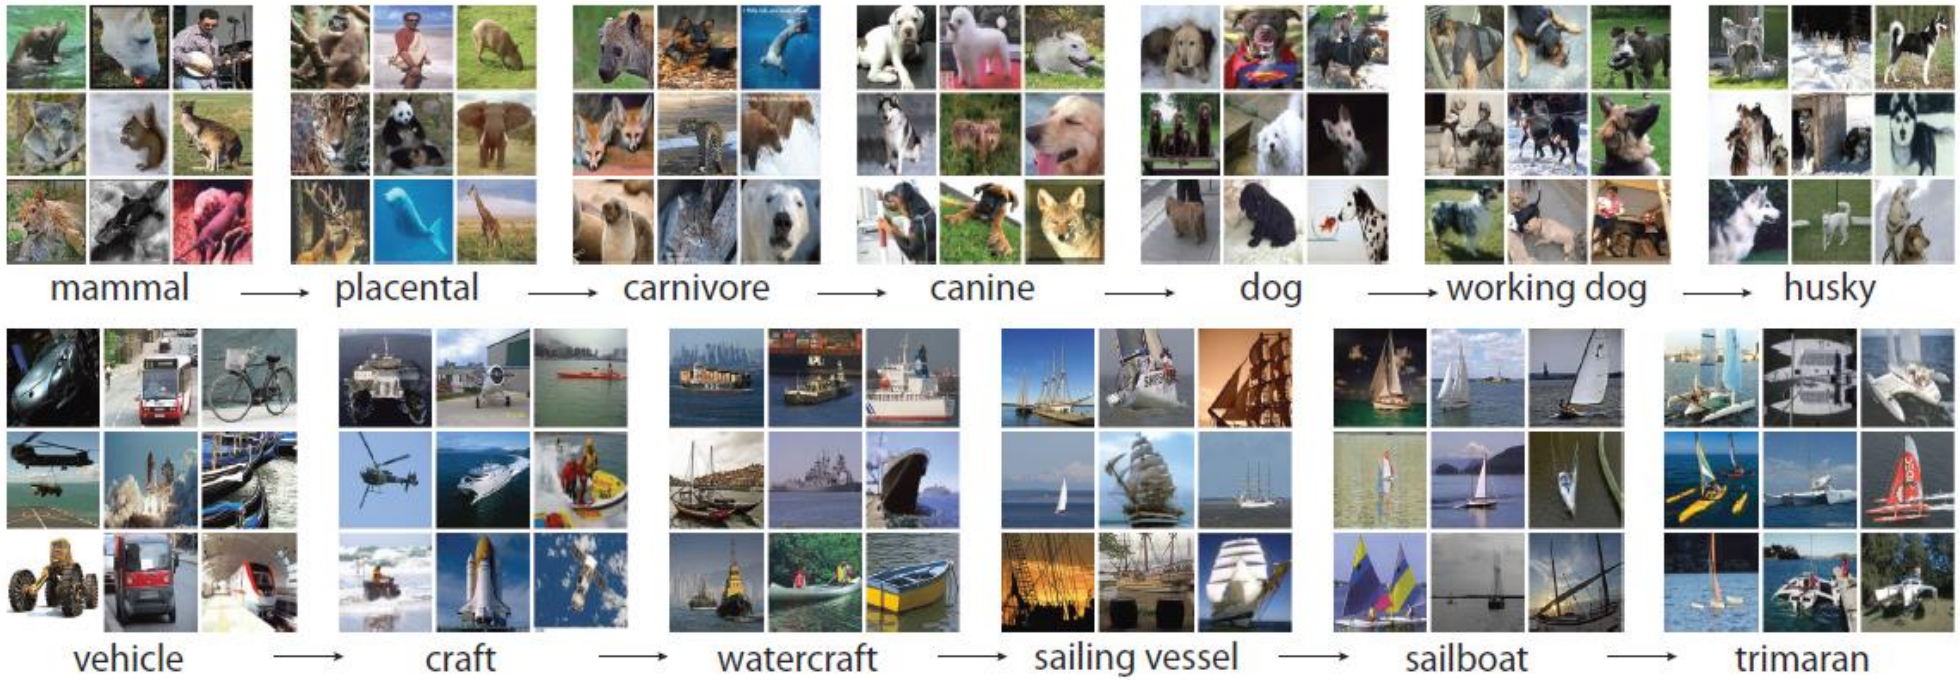

1,461,406 full resolution images

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

The image classification challenge was run annually from 2010 to 2017

[figures from www.nvidia.com]

ImageNet Challenge

- The ImageNet Large Scale Visual Recognition Challenge

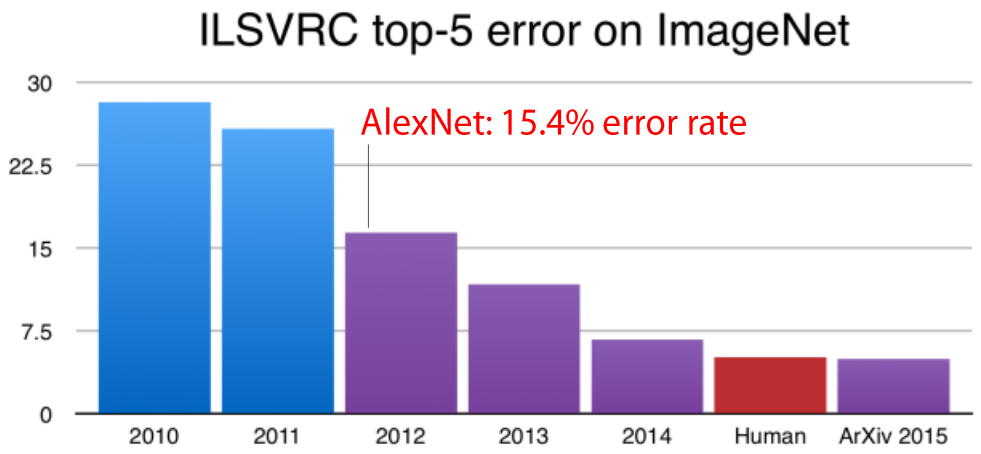

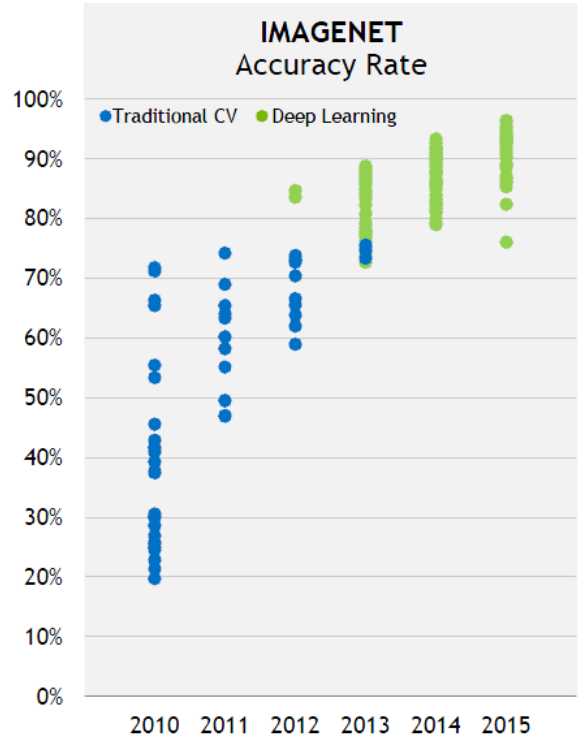

1,461,406 full resolution images

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 Neural Networks have evolved since then ...*

Since 2018, these system can identify objects and persons from videos, in real time

NOTE:
According to the recent EU AI Act,
***remote biometric identification** (RBI)*
in public places will require
a special authorization

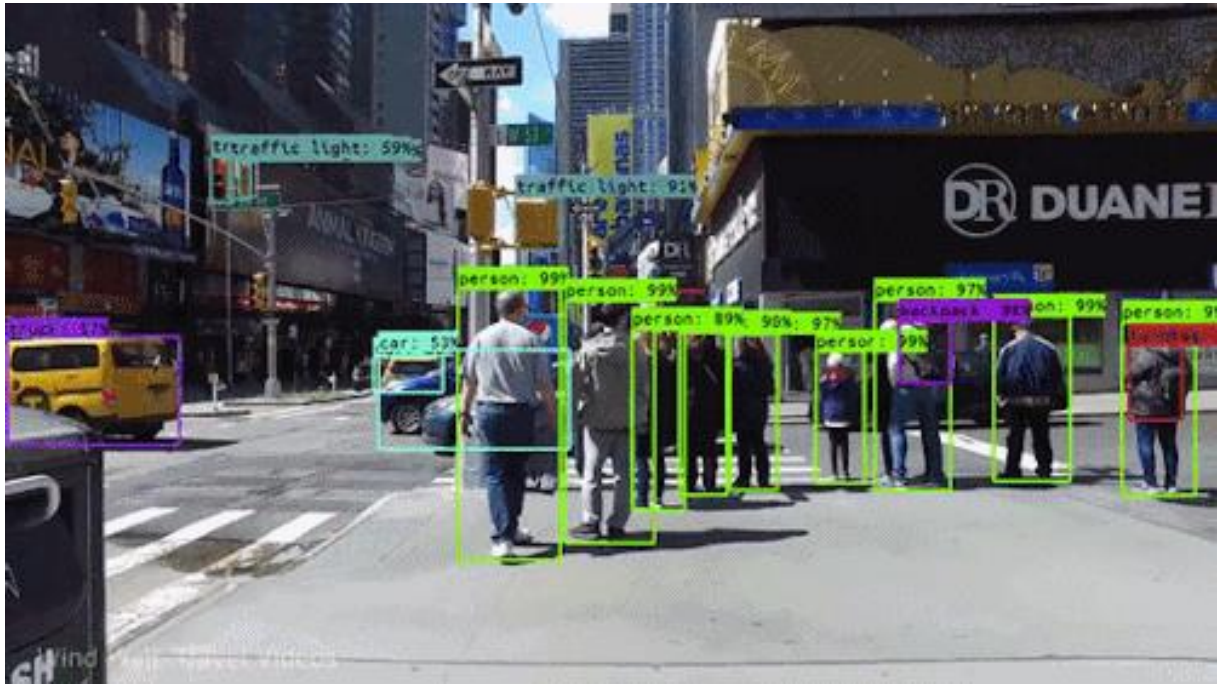

Image from: https://sgu.ac.id/id/computer-vision-artificial-intelligence-why-is-it-important/

Image Segmentation

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

They can also perform a complete scene analysis from videos, in real time

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

Generative AI

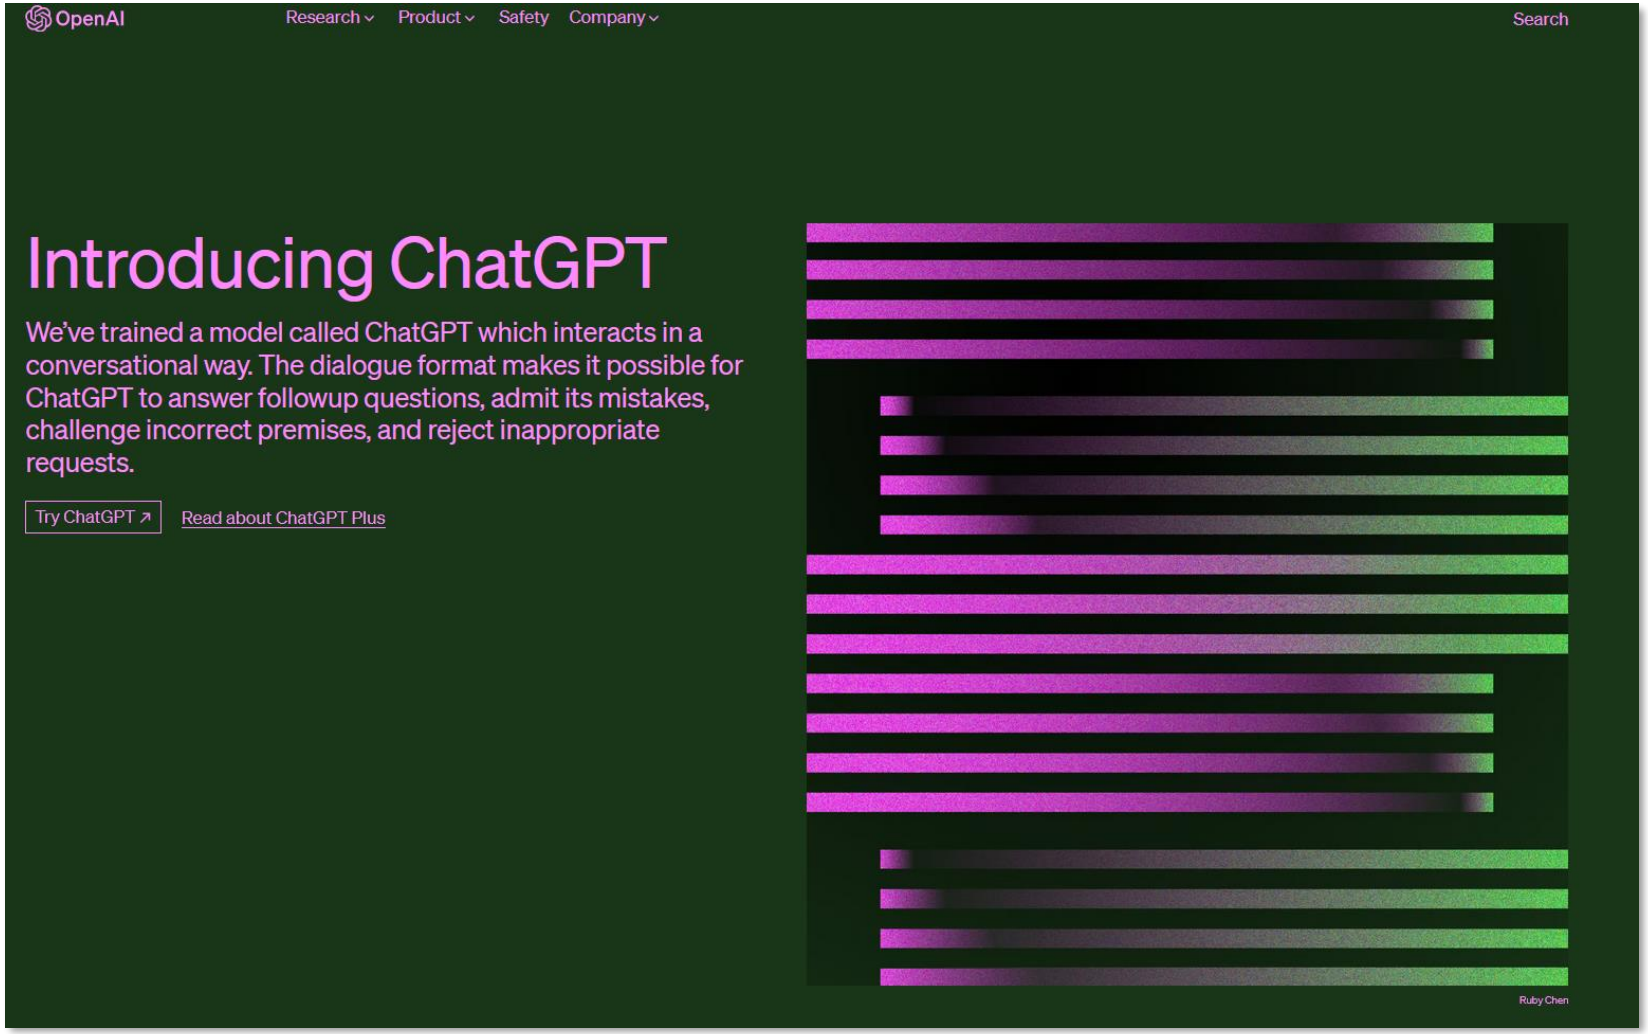

[Image from https://openai.com/blog/chatgpt, 09/03/2023]

Which may have been unexpected ...

The New York Times

How ChatGPT Kicked Off an A.I. Arms Race

Even inside the company, the chatbot's popularity has come as something of a shock.

By Kevin Roose

Feb. 3, 2023

[...] As ChatGPT has captured the world's imagination, Mr. Altman has been put in the rare position of trying to downplay a hit product.

He is worried that too much hype for ChatGPT could provoke a regulatory backlash or create inflated expectations for future releases, [...]

[Quote and images from https://www.nytimes.com/2023/02/03/technology/chatgpt-openai-artificial-intelligence.html/]

ChatGPT is not alone

- **DALL-E2**

  Diffusion Models: generating images from text

  «A teapot in the shape of an avocado»

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

Generating videos

ChatGPT is not alone

- **SORA**

  Generating videos from text prompts

«A stylish woman walks down a Tokyo street filled with warm glowing neon and animated city signage. She wears a black leather jacket, a long red dress, and black boots, and carries a black purse. She wears sunglasses and red lipstick. She walks confidently and casually. The street is damp and reflective, creating a mirror effect of the colorful lights. Many pedestrians walk about.»

[Video clip from https://openai.com/index/sora/]

Is Artificial Intelligence "X"?

Is Artificial Intelligence intelligent?

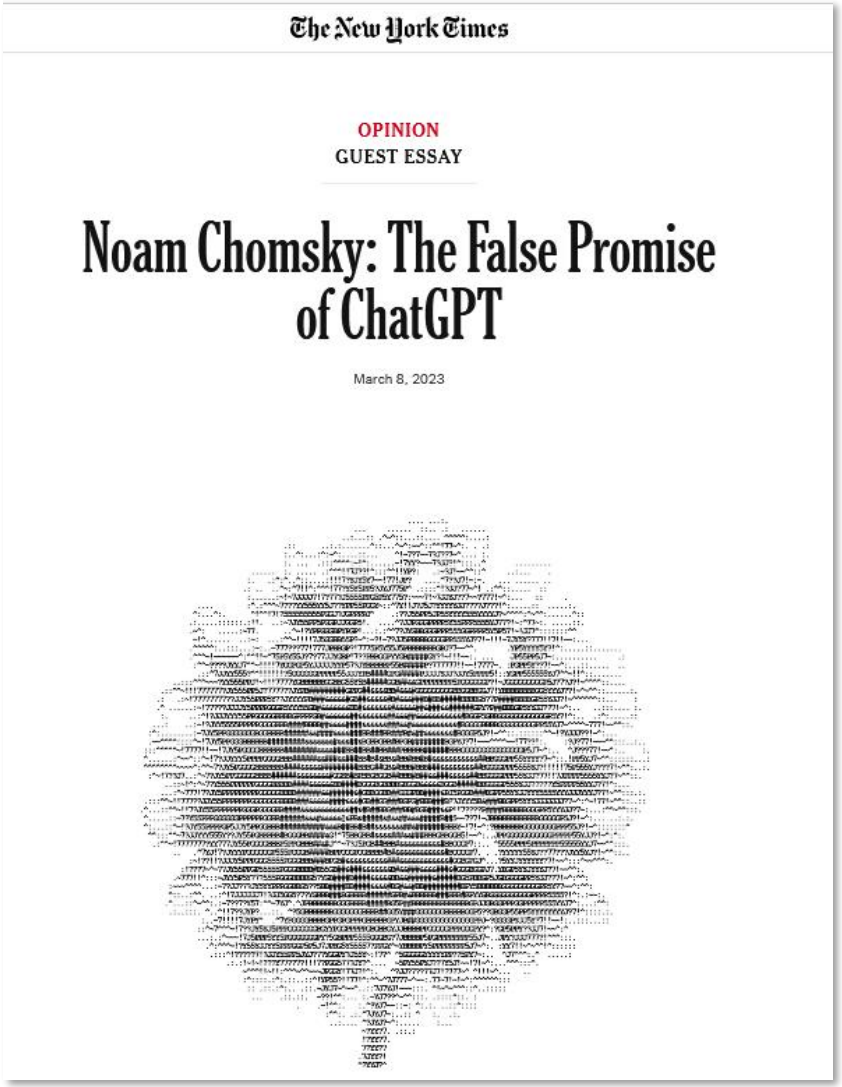

The New York Times

OPINION

GUEST ESSAY

Noam Chomsky: The False Promise of ChatGPT

March 8, 2023

[Image from https://www.nytimes.com/2023/03/08/opinion/noam-chomsky-chatgpt-ai.html]

Is Artificial Intelligence harmful?

The Guardian

Project Syndicate economists Global economy

Is AI really the biggest threat when our world is guided more by human stupidity?

Nouriel Roubini

There is both hope and hype for what artificial intelligence can do for growth – if politicians can tame its destructive potential

Mon 5 Feb 2024 14.06 CET

Share

[Image from https://www.theguardian.com/business/2024/feb/05/is-ai-really-the-biggest-threat-when-our-world-is-guided-more-by-human-stupidity]

Is Artificial Intelligence biased?

The New York Times

Artificial Intelligence › | An Unsettling Chat With Bing | Read the Conversation | How Chatbots Work | Spotting A.I.-Generated Text

Google Researcher Says She Was Fired Over Paper Highlighting Bias in A.I.

Timnit Gebru, one of the few Black women in her field, had voiced exasperation over the company's response to efforts to increase minority hiring.

Give this article | 276

Timnit Gebru, a respected researcher at Google, questioned biases built into artificial intelligence systems. Cody O'Loughlin for The New York Times

By Cade Metz and Daisuke Wakabayashi

Dec. 3, 2020

[Quote from https://www.nytimes.com/2020/12/03/technology/google-researcher-timnit-gebru.html]

Is Artificial Intelligence helpful?

The Guardian

It's cheap, quick and available 24/7, but is a chatbot therapist really the right tool to tackle complex emotional needs?

by Alice Robb

Illustration: Javier Jaén/The Guardian

Sat 2 Mar 2024 17.00 CET

[Image from https://www.theguardian.com/lifeandstyle/2024/mar/02/can-ai-chatbot-therapists-do-better-than-the-real-thing]

Is Artificial Intelligence legal?

The EU AI Act; A risk-based approach

Regulation - EU - 2024/1689, effective since 01/08/2024

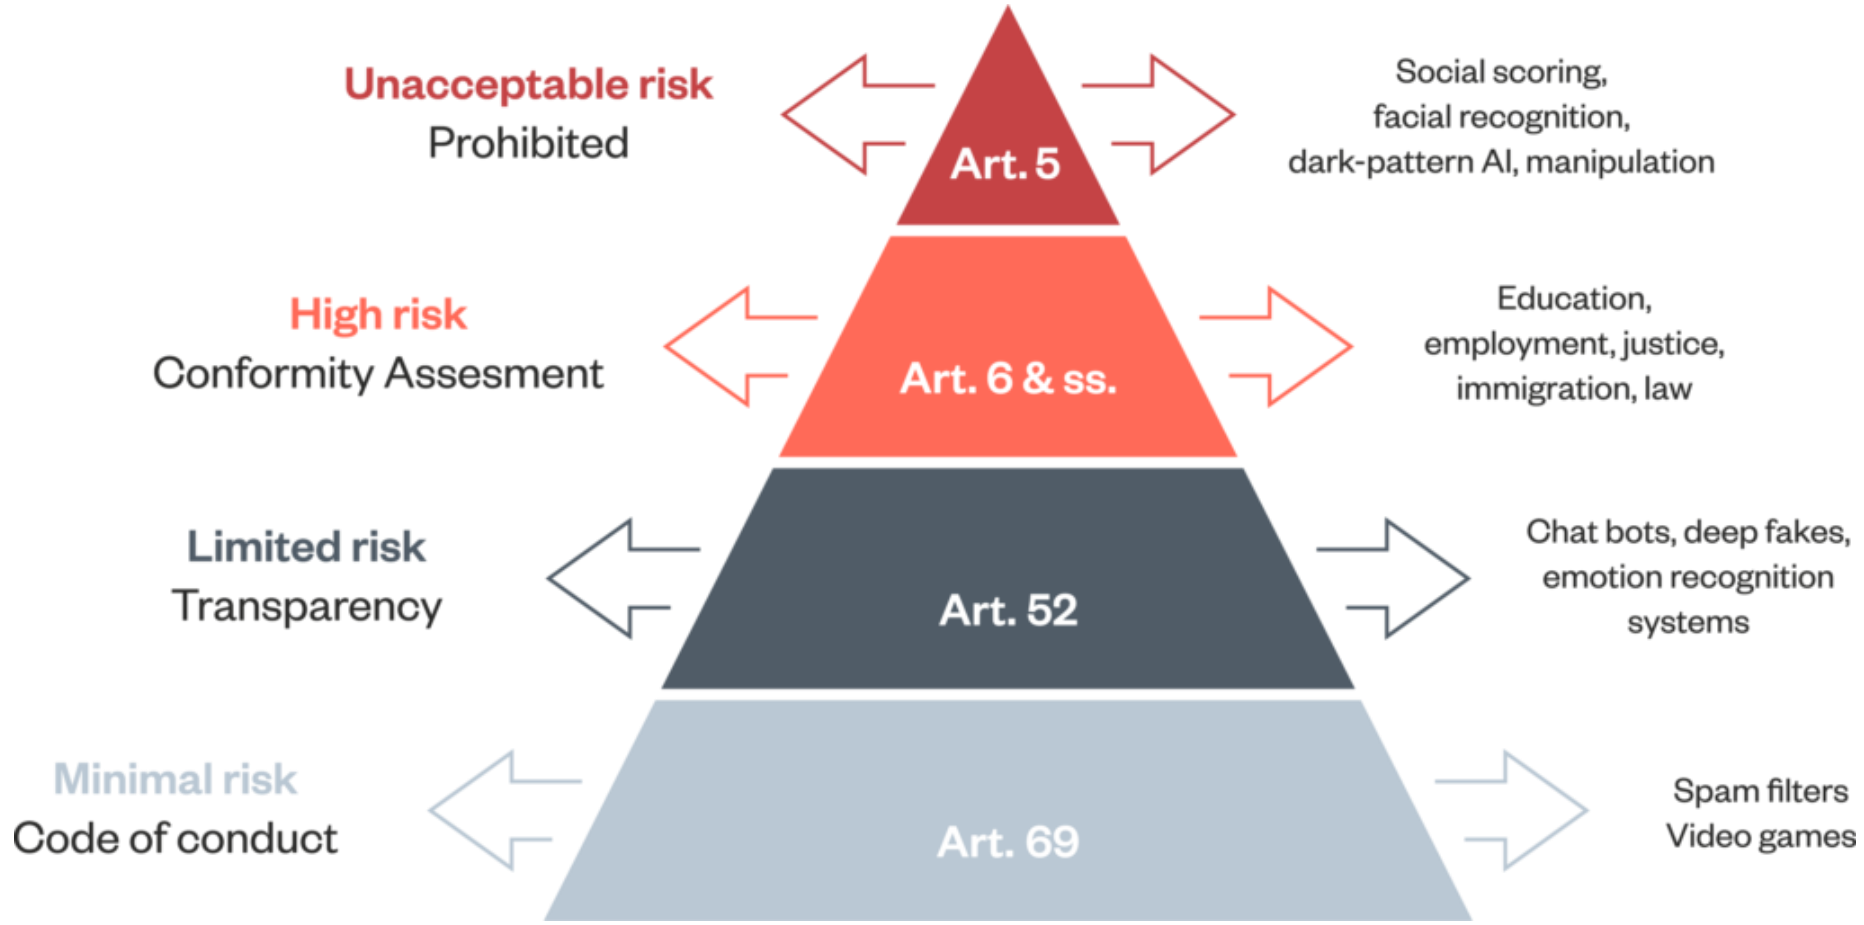

[Image from: https://www.adalovelaceinstitute.org/resource/eu-ai-act-explainer/]

Discovering Patterns (and Making Use of Them)

Artificial Intelligence is good at finding patterns ...

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]

Which patterns?

Do we ourselves represent a cat as a pattern?

Or a generalized pattern of a prototype?

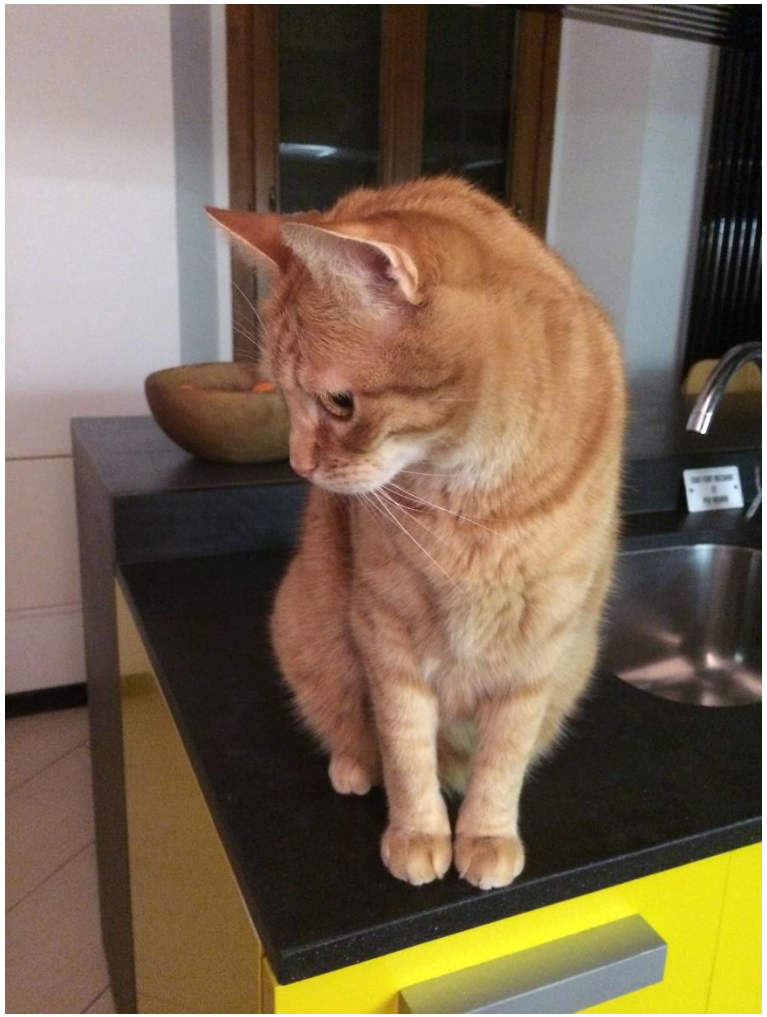

[this is *my* cat, Rabarbaro]

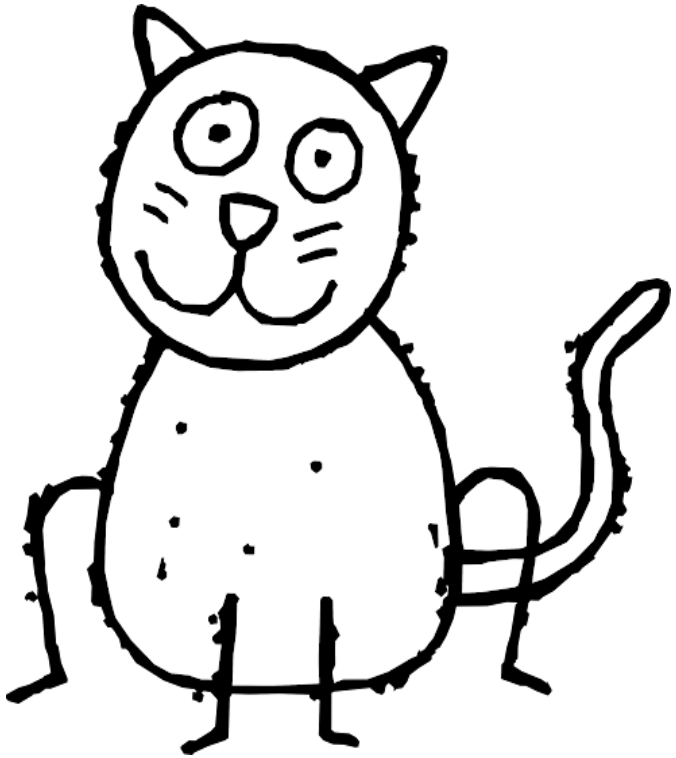

ChatGPT: how does it work?

- **Positional encoding (*Embedding*)**

Words (=*token*) from natural language are each translated into a high-dimensional *numerical vector*

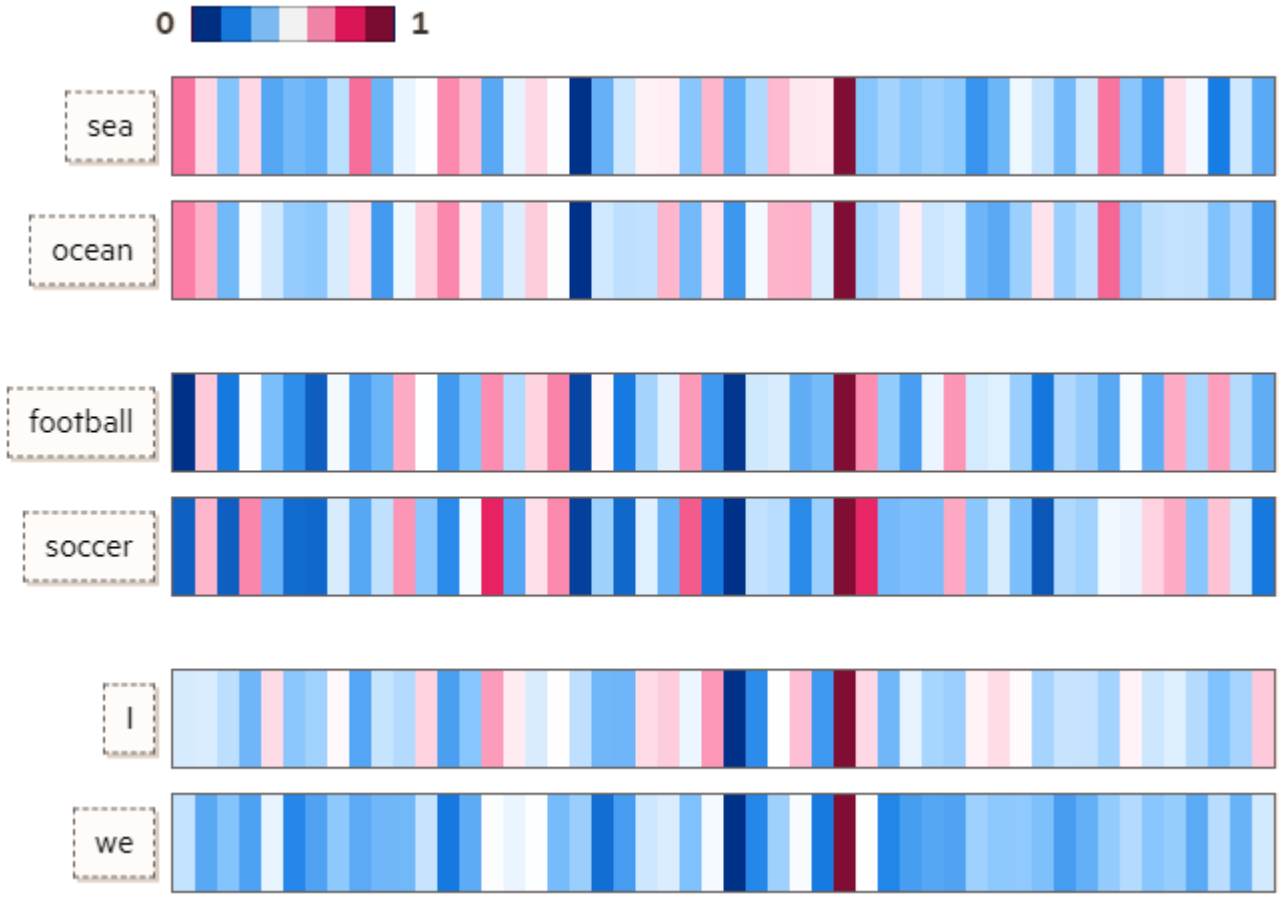

[images from https://ig.ft.com/generative-ai/]

ChatGPT: how does it work?

- **Positional encoding (*Embedding*)**

  Words (=*token*) from natural language are each translated into a high-dimensional *numerical vector*

  Such vector is computed by estimating the *probability of co-occurrence* in a context of other words in a (very) large text corpus

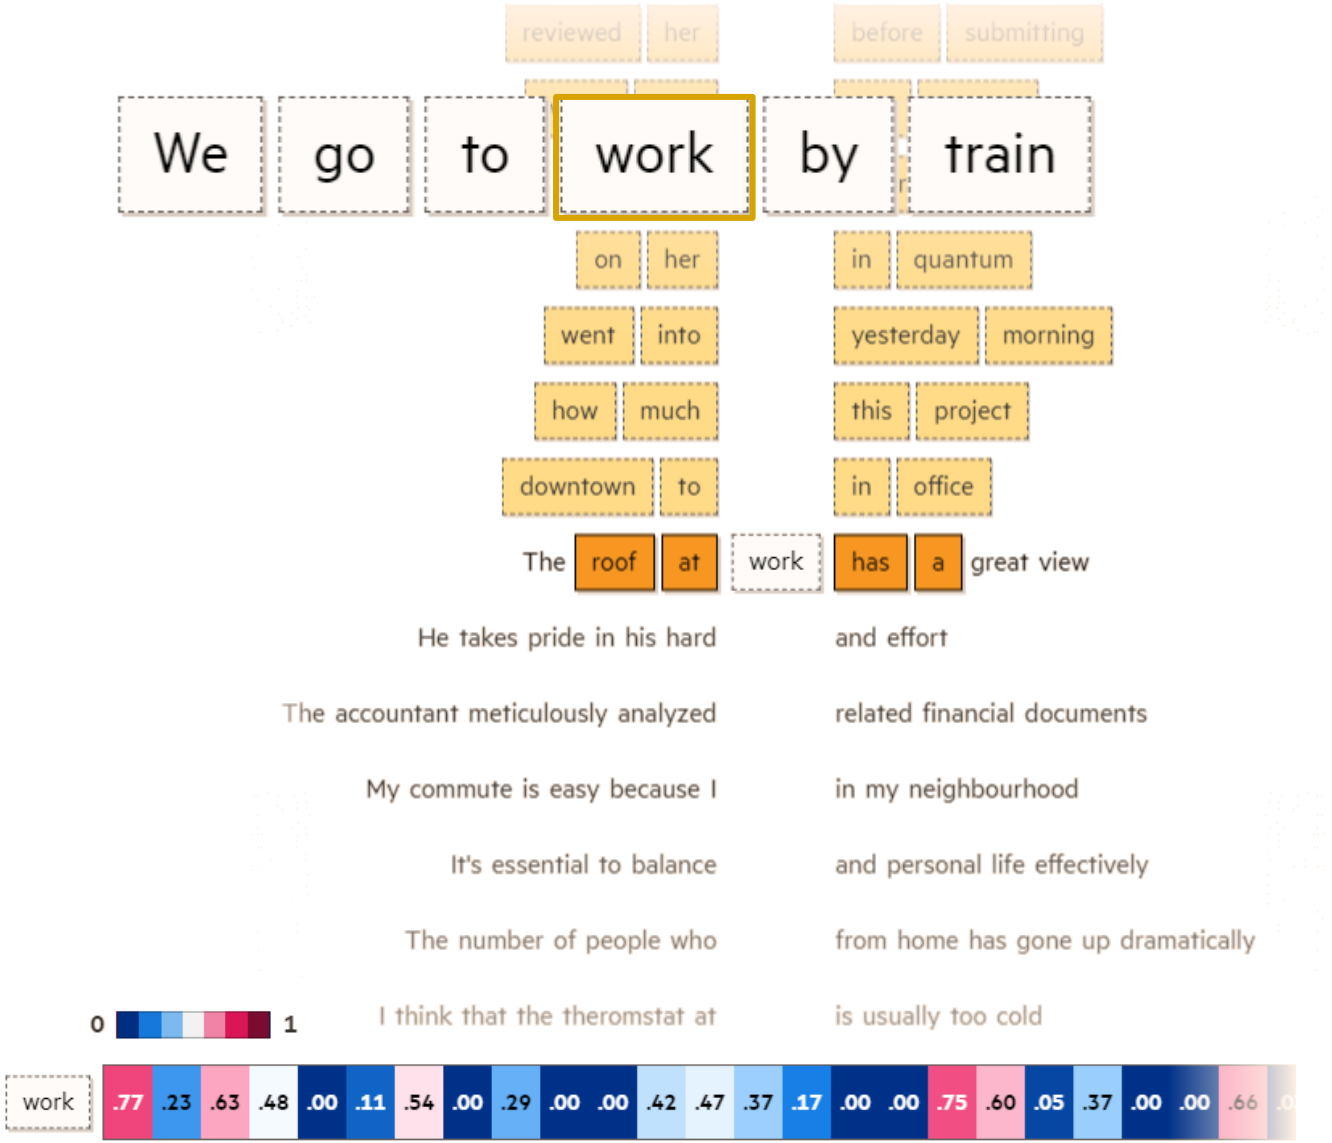

[Images from https://ig.ft.com/generative-ai/]

ChatGPT: how does it work?

- **Positional encoding (*Embedding*)**

  Words (=*token*) from natural language are each translated into a high-dimensional *numerical vector*

  Such vector is computed by estimating the *probability of co-occurrence* in a context of other words in a (very) large text corpus

  In this way, the *numerical similarity* among vectors is representative of words' affinity in terms of role or meaning (or both)

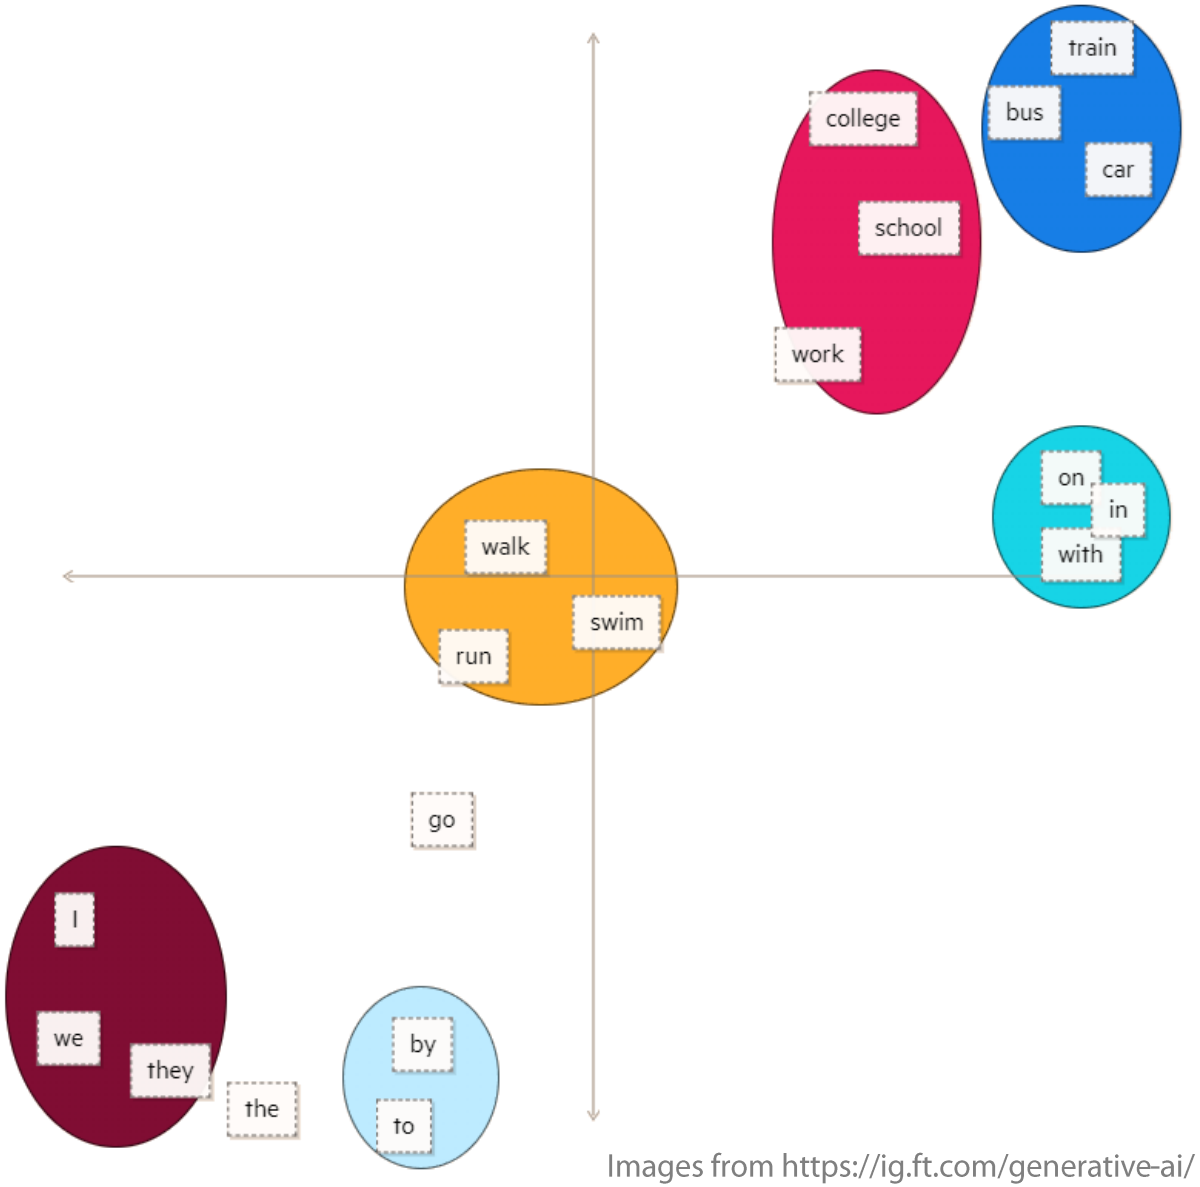

Images from https://ig.ft.com/generative-ai/

Chat LLM: completion-mode, one-token at time

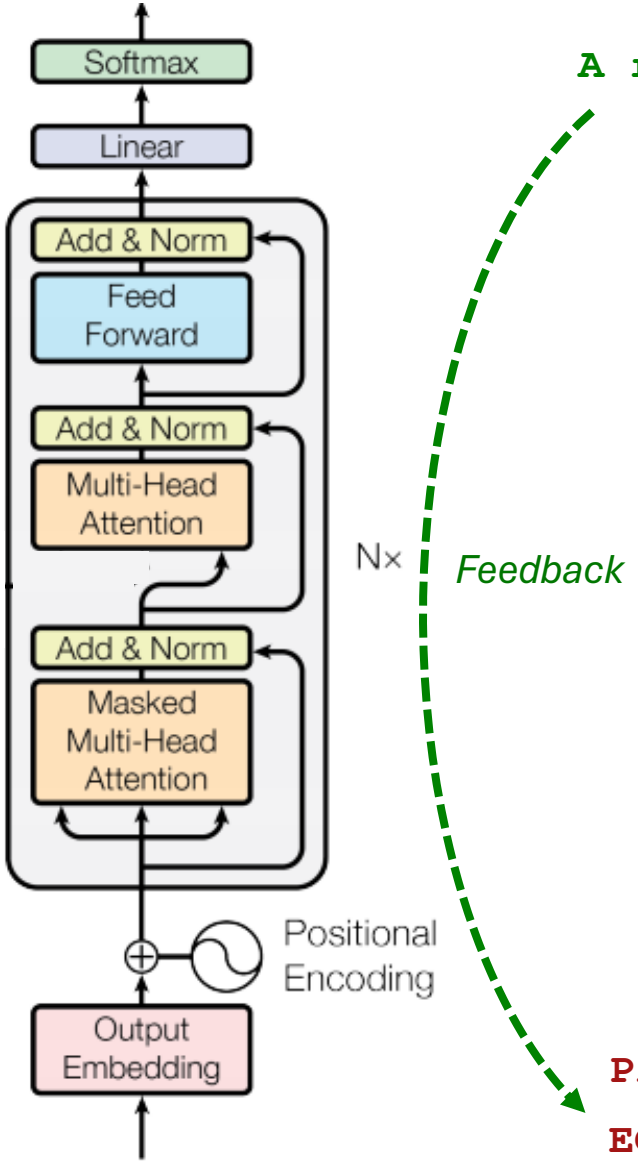

[image from https://arxiv.org/pdf/1706.03762.pdf]

A robot may not injure a human being

Feedback

Please recite the first law of robotics

EOS A robot may not injure a human

The system is designed to complete texts starting from a prompt

It has no memory: the context is represented by the feedback

Chat LLM: completion-mode, one-token at time

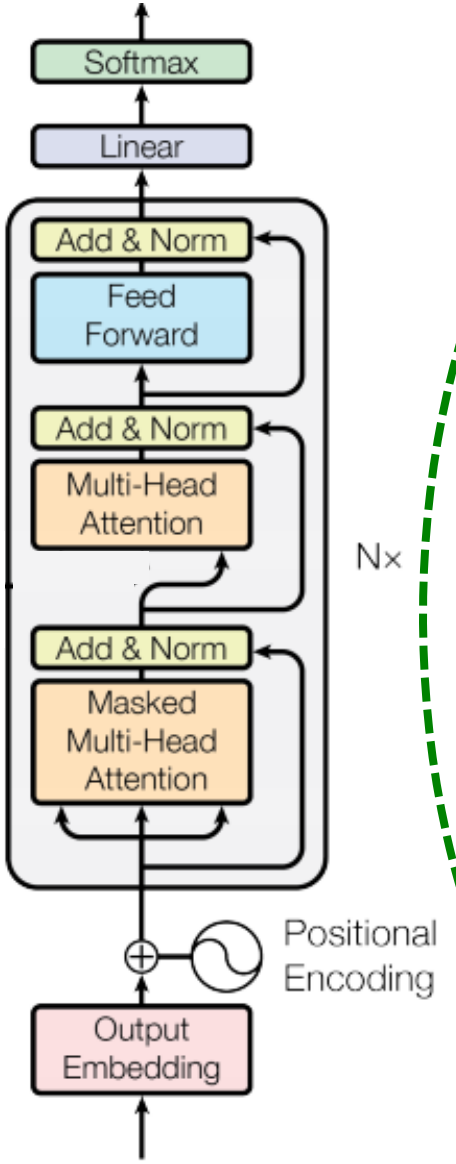

[image from https://arxiv.org/pdf/1706.03762.pdf]

Nx

Feedback

A robot may not ***injure*** **a human being**

Probability

\- +

Please recite the first law of robotics

EOS A robot may not injure a human

The system is designed to complete texts starting from a prompt

It has no memory: the context is represented by the feedback

At each step, the neural network generates a probability distribution of the next token

Chat GPT: learning in three steps

OpenAI

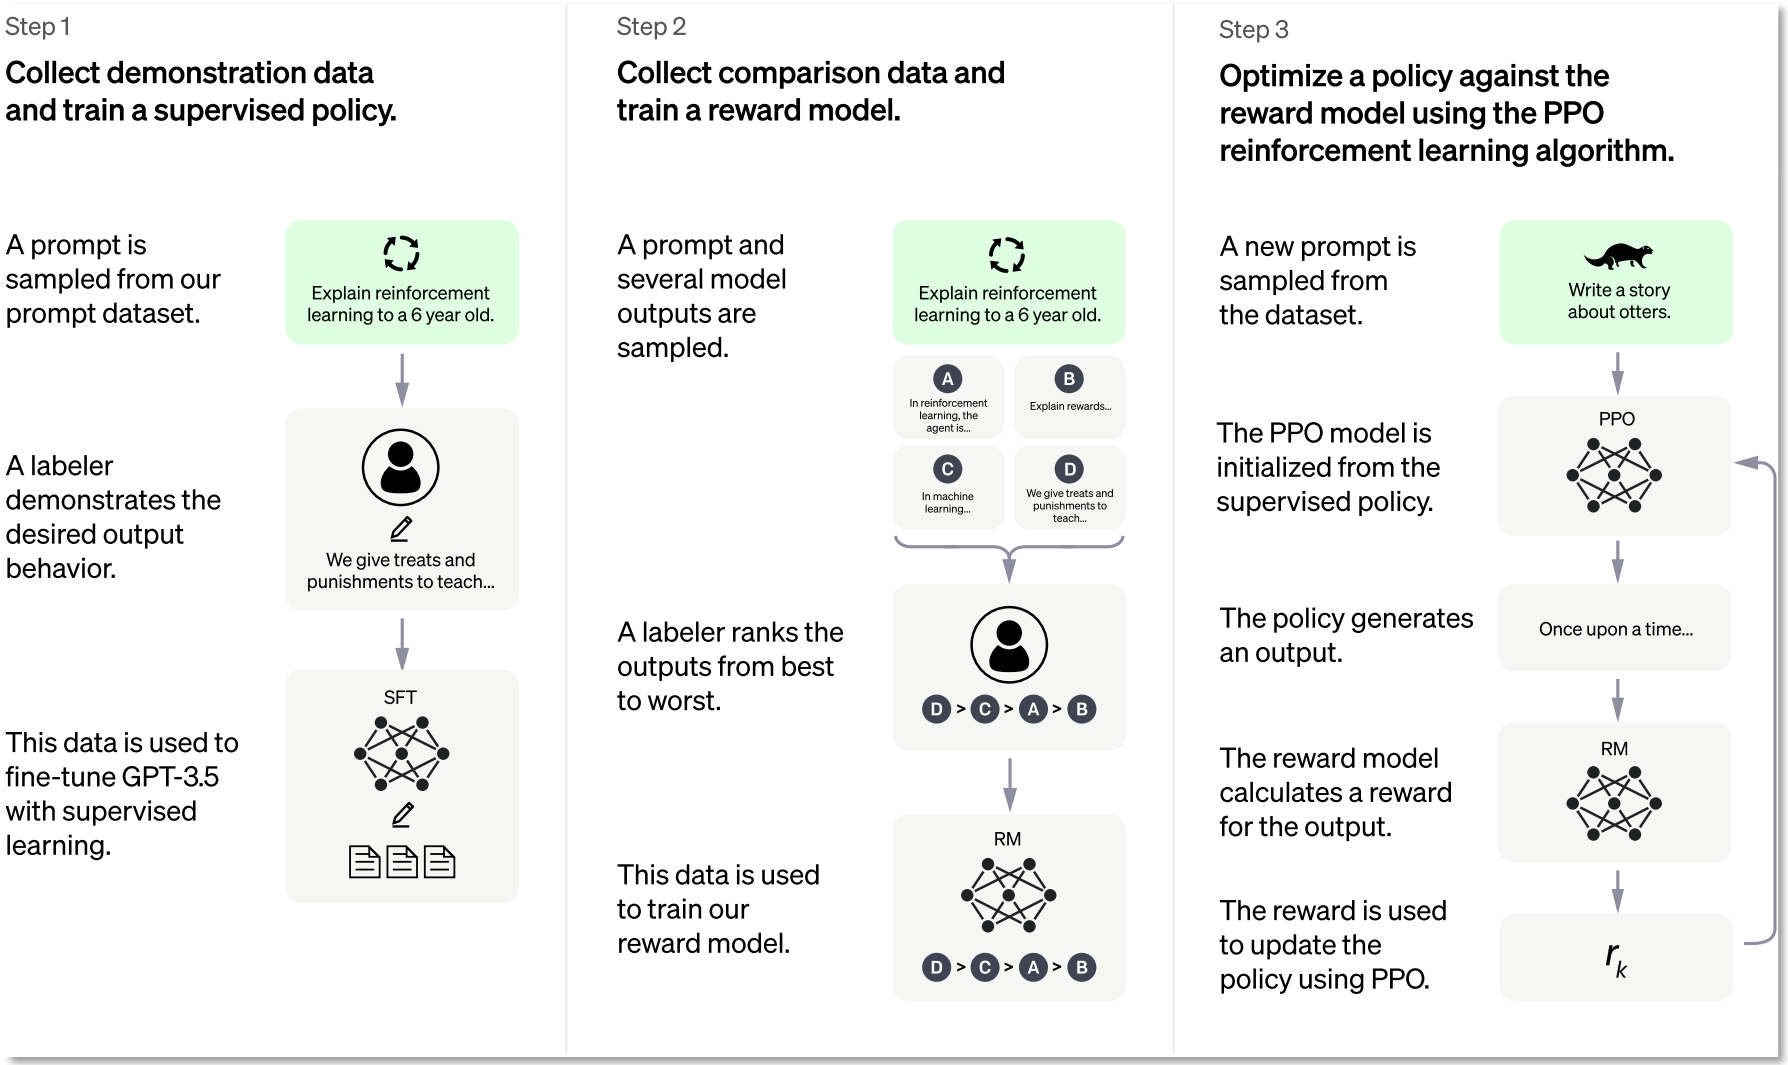

[image from https://openai.com/index/chatgpt/]

Putting together all the tricks in the book (at state-of-art)

LLMs in a Reasoning Engine

Google Gemini

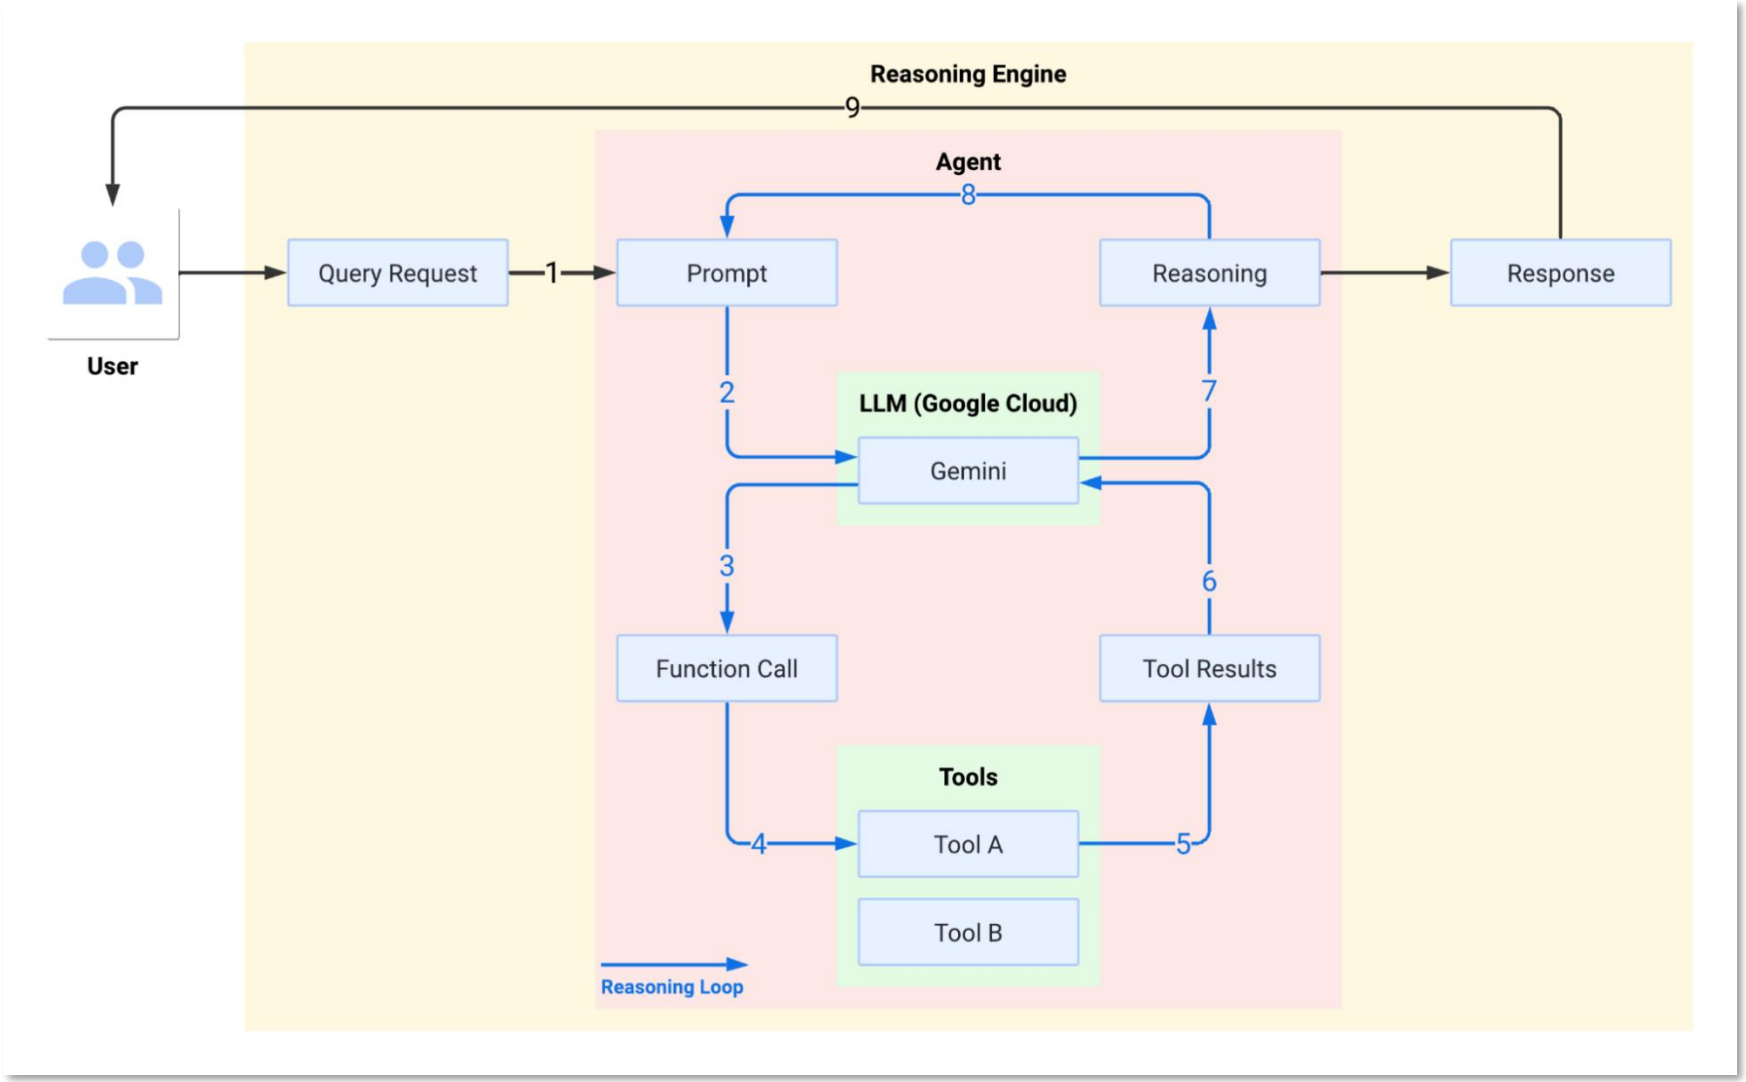

Integrating LLMs and tools in a reasoning loop

[image from https://www.googlecloudcommunity.com/gc/Community-Blogs/Building-and-Deploying-AI-Agents-with-LangChain-on-Vertex-AI/bc-p/796691]

Finding patterns in large datasets

Well, it's just a function anyway…

Games of strategy: AlphaGo (2016)

[image from: https://nikcheerla.github.io/deeplearningschool/2018/01/01/AlphaZero-Explained/]

- Mastering the game of Go with deep neural networks and tree search [2016, D. Silver, et al. (22 authors), http://www.nature.com/nature/journal/v529/n7587/full/nature16961.html]

There are more possible positions in Go than there are atoms in the universe

Sophisticated machine-learning techniques

Strategy selection via Monte Carlo Tree Search (MCTS)

Deep neural networks (trained on human matches) for both guidance and learning

Adversarial self-training:
playing again itself and improving via reinforcement learning

Super-human player?

On March 2016, AlphaGo won 4-1 against the legendary Lee Sedol,
the top Go player in the world over the past decade

Deep Reinforcement Learning (DRL)

- ***A Deep Neural Network learns a <u>policy</u>***

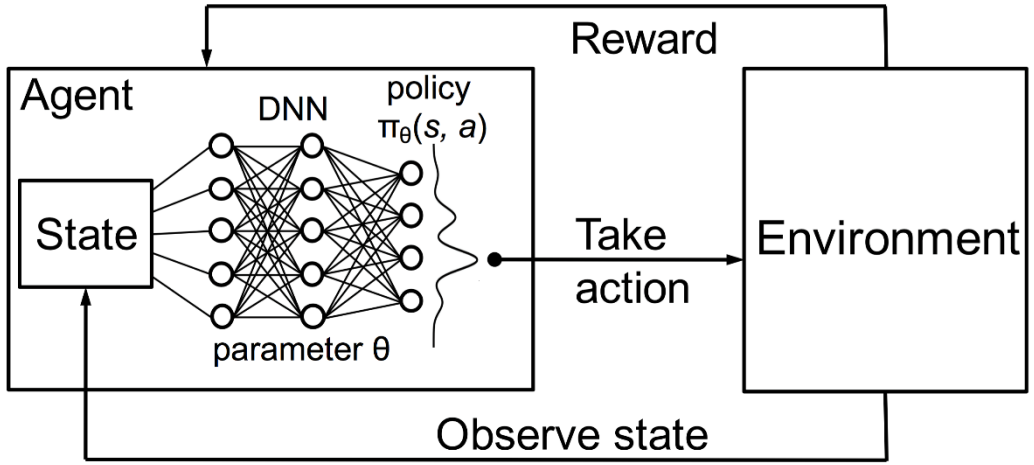

The agent interacts with an environment (it could be a copy of itself)

It selects an **action** in each **state** and receives a **reward** (possibly deferred) as a function of the results obtained

The DRL system optimizes its <u>policy</u>

Beyond emulating humans: AlphaZero (2018)

Image from: https://deepmind.com/blog/article/alphazero-shedding-new-light-grand-games-chess-shogi-and-go

AlphaGo was largely reliant
on the experience of human players

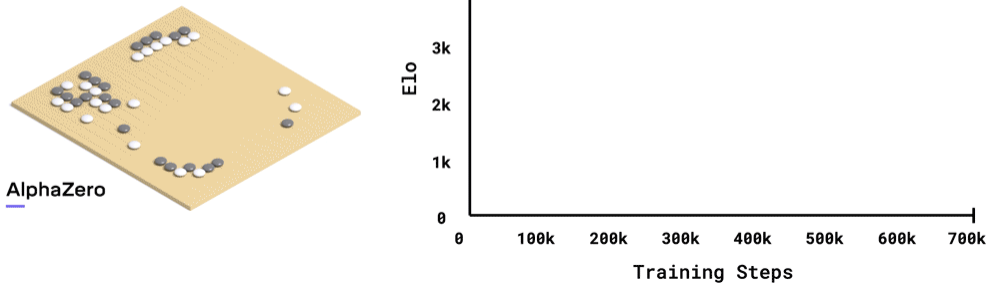

- **AlphaZero learns by itself**
  [2018, D. Silver, et al. (13 authors), https://science.sciencemag.org/content/362/6419/1140.full]

Basic Knowledge Only

It just knows the basic rules of the games

Learning via Self-Play

It plays against a (frozen) copy of itself

MCTS and DCNN in a closed loop

DeepMind

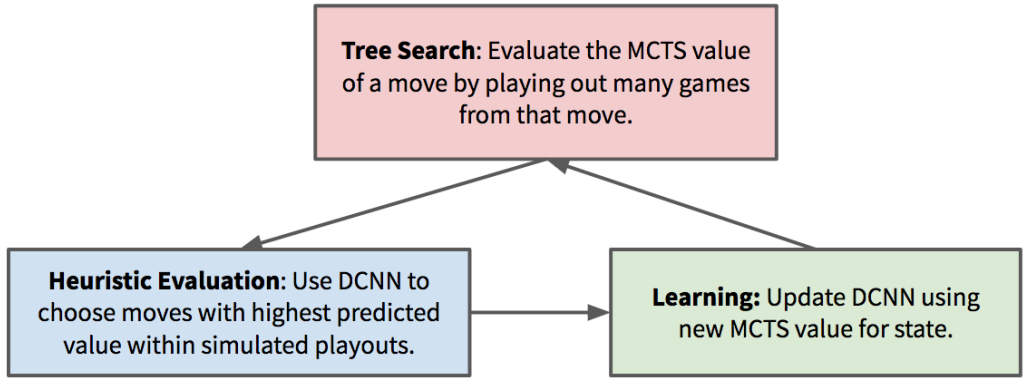

Image from: https://nikcheerla.github.io/deeplearningschool/2018/01/01/AlphaZero-Explained/

Beyond emulating humans: AlphaZero (2018)

Image from: https://deepmind.com/blog/article/alphazero-shedding-new-light-grand-games-chess-shogi-and-go

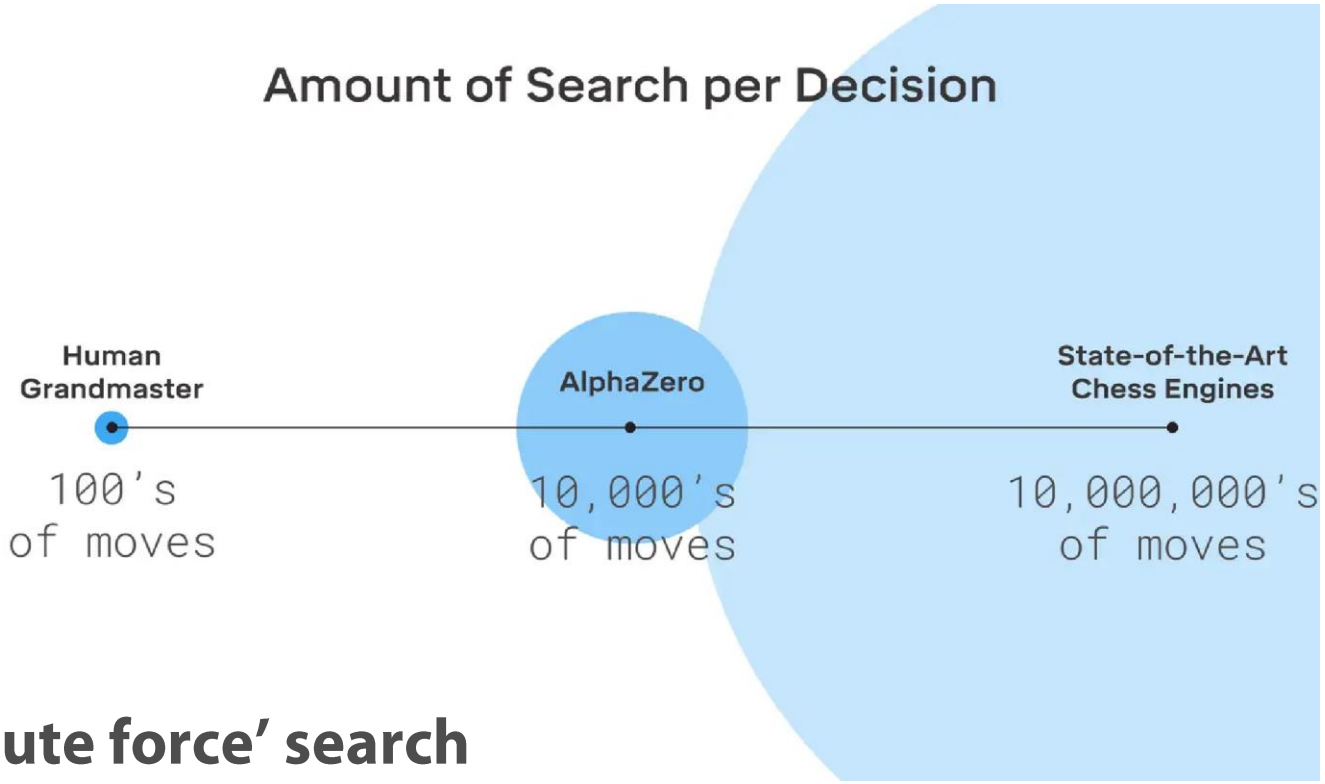

- **AlphaZero uses much less 'brute force' search**
  - When playing, the search process is driven by its neural network
    - *It acts like a memory of past experiences*
  - While training, it learns through a huge amount of self-playing
    - *But it is a faster learner than Alpha Go*

The true story of computers that play chess with humans

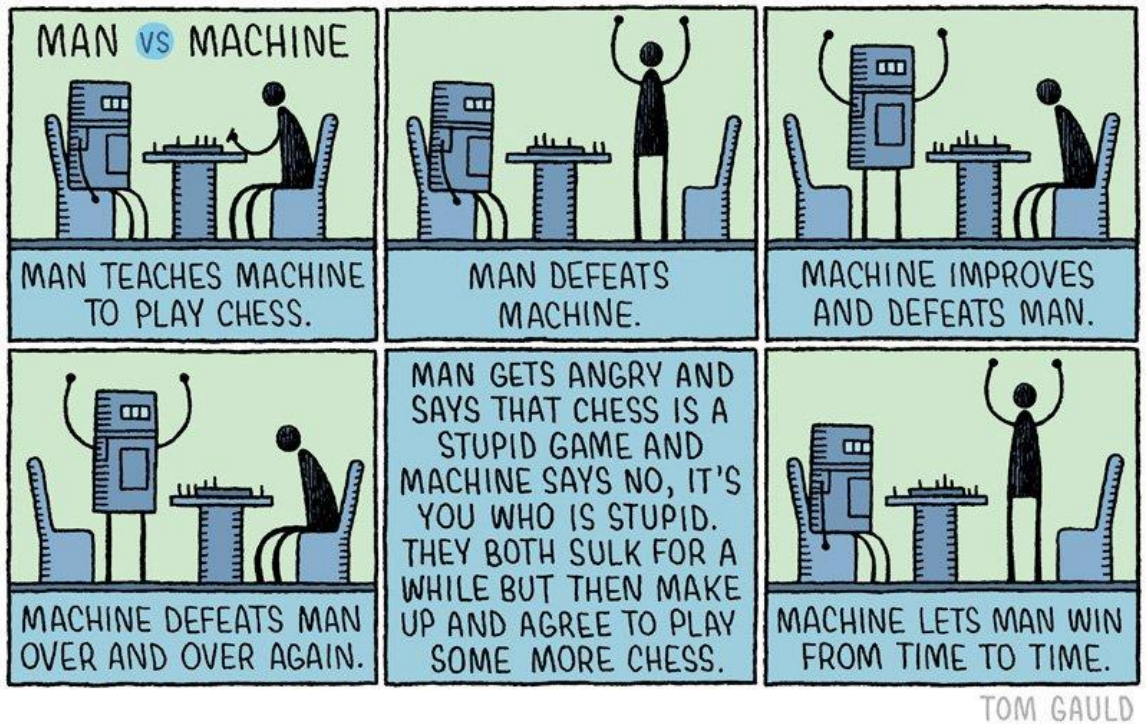

[Image from https://www.tomgauld.com/portfolio]

A Few Further Marvels

AlphaFold (2021)

Predicting 3D protein folding from primary structure

- **AlphaFold**

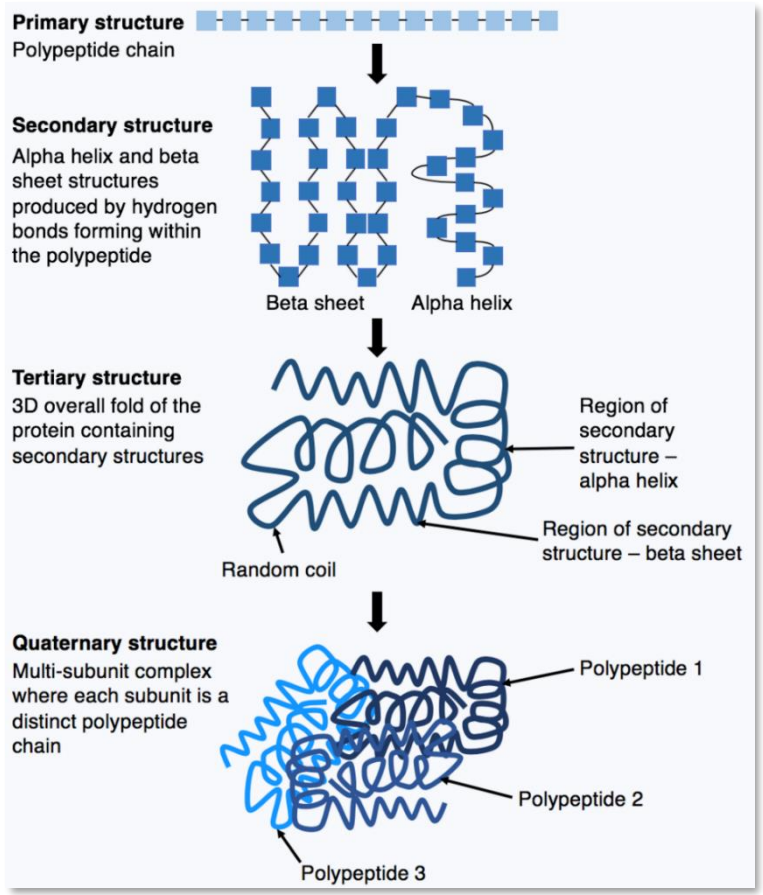

[Image from: https://en.wikipedia.org/wiki/AlphaFold]

Recycling iteration 0, block 03
Secondary structure assigned from the final prediction

[Video clip from: https://www.nature.com/articles/s41586-021-03819-2]

A Nobel Prize

The Nobel Prize in Chemistry 2024

They cracked the code for proteins' amazing structures

The Nobel Prize in Chemistry 2024 is about proteins, life's ingenious chemical tools. David Baker has succeeded with the almost impossible feat of building entirely new kinds of proteins. Demis Hassabis and John Jumper have developed an AI model to solve a 50-year-old problem: predicting proteins' complex structures. These discoveries hold enormous potential.

[Quote from: https://www.nobelprize.org/prizes/chemistry/

© Johan Jarnestad/The Royal Swedish Academy of Sciences

AlphaCode (2022)

Solving competitive coding problems using transformers

② AlphaCode reads the whole problem statement and produces code, analogous to how a human would approach the problem by reading it, coding a solution, and submitting.

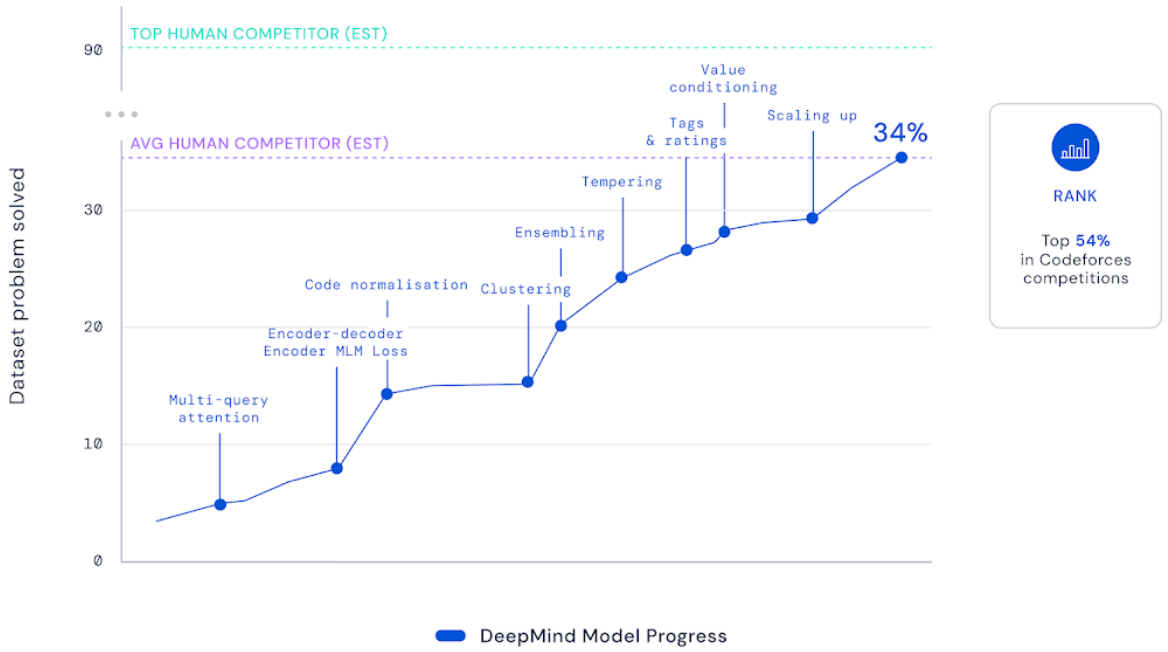

RANK

Top **54%** in Codeforces competitions

AlphaCode

② Solution (output)

```
t=int(input())
for i in range(t):
    s=input()
    t=input()
    a=[]
    b=[]
    for j in s:
        a.append(j)
    for j in t:
        b.append(j)
    a.reverse()
    b.reverse()
    c=[]
    while len(b)!=0 and len(a)!=0:
        if a[0]==b[0]:
            c.append(b.pop(0))
            a.pop(0)
        elif a[0]!=b[0] and len(a)!=1:
            a.pop(0)
            a.pop(0)
        elif a[0]!=b[0] and len(a)==1:
            a.pop(0)
    if len(b)==0:
        print("YES")
    else:
        print("NO")
```

First AlphaCode reads the two phrases.

If the letters at the end of both phrases don't match, the last letter must be deleted. If they do match we can move onto the second last letter and repeat.

Backspace deletes two letters. The letter you press backspace instead of, and the letter before it.

If we've matched every letter, it's possible and we output that.

[Images from https://deepmind.google/discover/blog/competitive-programming-with-alphacode/]

OK, about this course

The course in a nutshell

1) **Fundamentals**

2) **Deep Supervised Learning**

3) **Deep Convolutional Neural Networks**

4) **Deep Recurrent Networks**

5) **Generative Networks**

6) **Deep Reinforcement Learning**

Labs with Google Colab

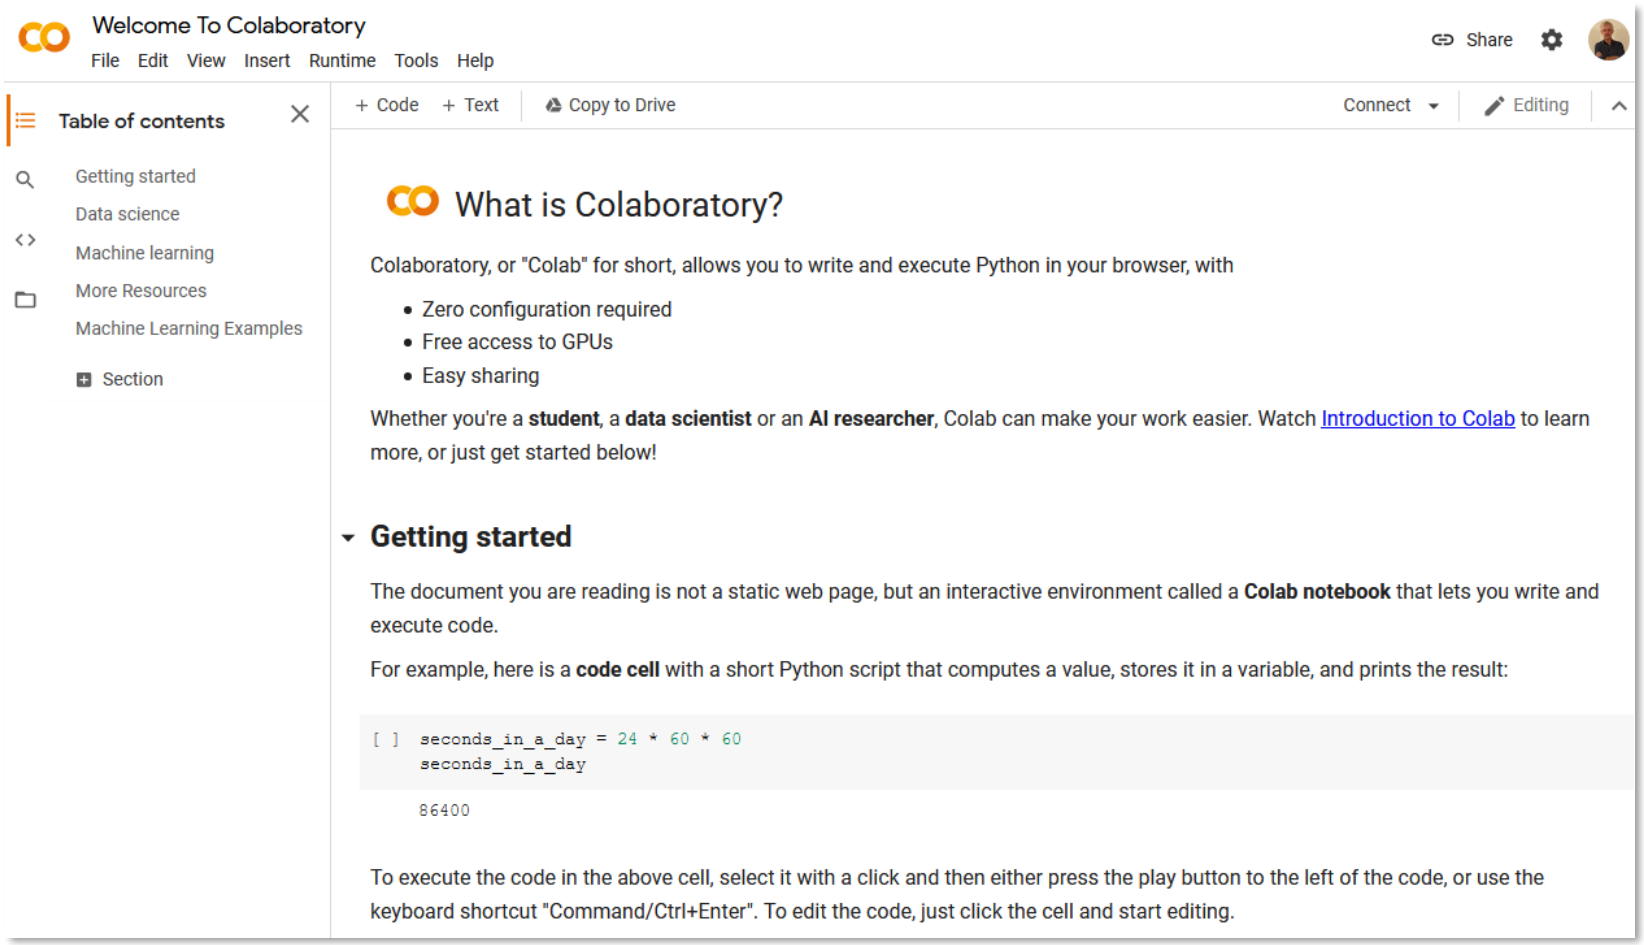

Make sure you have a look!

Better yet: follow the tutorial at https://colab.research.google.com/notebooks/intro.ipynb

PyCharm (optional)

However, if you really want to see how things work...

- ***Set it up on your computer***

  Python 3.10+
  Numpy 2.x+
  PyTorch 2.5.+
  (TensorFlow 2.+)

  ... and, above all ...

  Install PyCharm (Community Edition)
  https://www.jetbrains.com/pycharm/
  It's free
  Learn how to use the *debugger* :
  it will change the way you learn (and code)

CAUTION: no assistance whatsoever will be provided in the course, for this.
Legions of students succeeded on their own: you can make it too!

Please Register on KIRO

You will be able to receive any important messages relating to the course

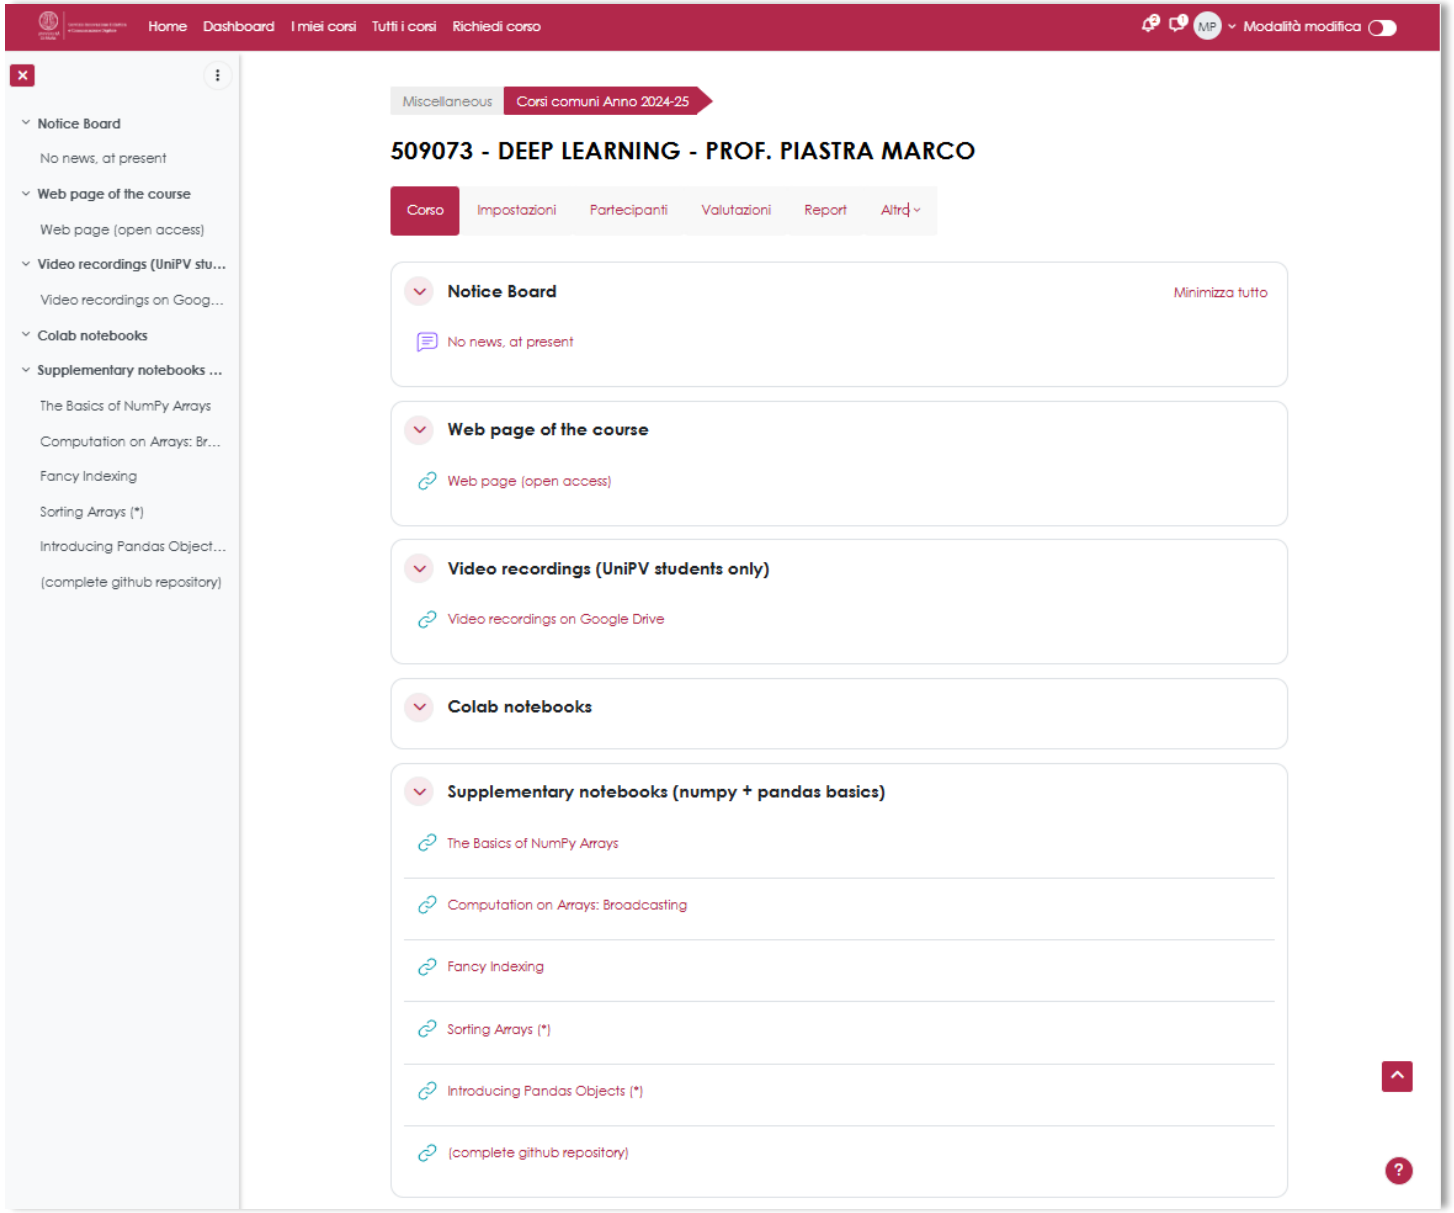

Course Page: https://elearning.unipv.it/course/view.php?id=**9221**

The Final Exam

1) Choose and propose a final project

The topic could be any of your choice
The techniques adopted must be (strongly) related with the course
Groupwork is allowed, with a maximum of two *(no exceptions)*

2) Develop and submit your project

Each final project must be submitted as a Google Colab notebook plus dataset *(if required)*

Submission must be made at least two days before the exams

3) Be strong on theory, nonetheless

Alone, not even a *phenomenal* final project will give you a final grade *(sorry)*

Appropriate knowledge about theory will be fundamental

4) Final Exam

Oral interview, about both final project and theory

Relative weights: final project 40%, theory 60%

Some reference textbooks for the course

- **Deep Learning**
  Foundations and Concepts

  Christopher Bishop, Hugh Bishop

  Springer, 2024

  See also
  https://www.bishopbook.com/

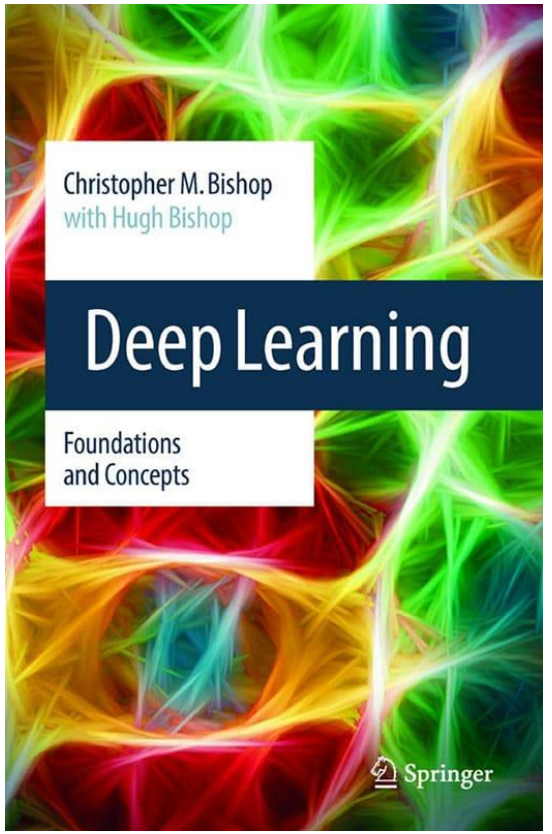

Some reference textbooks for the course

- **Deep Learning**

  Ian Goodfellow, Yoshua Bengio and Aaron Courville

  MIT Press, 2017

  See also
  http://www.deeplearningbook.org/

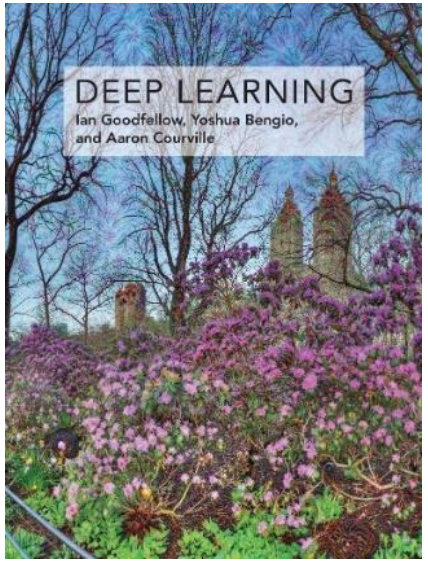

- **Hands-On Machine Learning with Scikit-Learn, Keras, and TensorFlow: Concepts, Tools, and Techniques to Build Intelligent Systems**, 3rd Edition

  Aurélien Géron

  O'Reilly, 2022

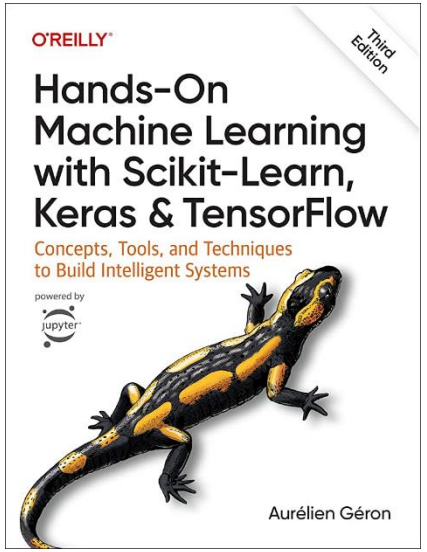

Some reference textbooks for the course

- **Dive into Deep Learning**

  Aston Zhang, Zachary Lipton, Mu Li and Alexander Smola

  Cambridge University Press, 2024

  See also
  https://d2l.ai/index.html
  A great accompanying site, with lots of multi-version examples (PyTorch, TensorFlow , JAX)

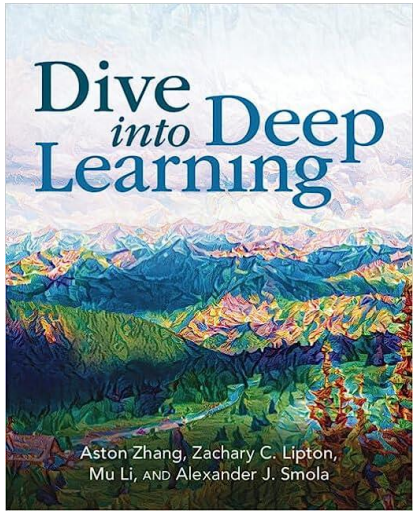

- **Reinforcement Learning, second edition**

  Richard S. Sutton, Andrew G. Barto

  The MIT Press, 2018