Introduzione sul significato di AI, terminologia e funzionalità, limiti e sviluppi delle applicazioni

Stefano Gualandi Università di Pavia, Dipartimento di Matematica

email: stefano.gualandi@unipv.it
bluesky: @famo2spaghi.bsky.social
blog: http://stegua.github.com
web: http://matematica.unipv.it/gualandi/

22 maggio 2024, Aula Magna – Collegio A. Volta

Outline

1 Introduction

2 ML tasks

3 Neural Networks

4 Generative AI

5 AI4Pharma

For an introduction? Ask to ChatGPT 1/2

ChatGPT

Come definiresti l'Intelligenza Artificiale ad un pubblico di studenti di master in chimica, e rappresentanti di aziende farmaceutiche?

Per un pubblico di studenti di master in chimica e rappresentanti di aziende farmaceutiche, definirei l'intelligenza artificiale (IA) in modo pratico e orientato all'applicazione. Ecco una possibile definizione:

L'intelligenza artificiale (IA) è un insieme di tecnologie che permettono a un computer o a un sistema di eseguire compiti che richiederebbero intelligenza umana, come l'apprendimento dai dati, il riconoscimento di schemi complessi, il ragionamento e la presa di decisioni. Nell'ambito chimico e farmaceutico, l'IA viene impiegata per accelerare la scoperta di farmaci, ottimizzare processi di sintesi chimica, analizzare grandi volumi di dati sperimentali e simulare le interazioni molecolari. In pratica, grazie all'IA, è possibile analizzare dati più velocemente, prevedere la reattività di nuove molecole, individuare nuovi bersagli terapeutici e personalizzare trattamenti in base ai profili genetici dei pazienti.

L'IA può ridurre il tempo e i costi di ricerca e sviluppo farmaceutico, aiutando a trovare soluzioni più rapidamente rispetto ai metodi tradizionali.

For an introduction? Ask to ChatGPT 2/2

ChatGPT

Si ma io sono un matematico, e mi piacciono le definizioni formali. Come posso definire l'intelligenza artificiale?

Memory updated

In un contesto più formale, possiamo definire l'intelligenza artificiale in termini matematici come segue:

L'intelligenza artificiale (IA) è lo studio e la costruzione di agenti intelligenti, dove un agente è un'entità che percepisce l'ambiente attraverso sensori e agisce su di esso attraverso attuatori. Formalmente, un **agente** può essere descritto come una funzione f che mappa una sequenza di percezioni $p_1, p_2, \ldots, p_n$ in un'azione a, cioè:

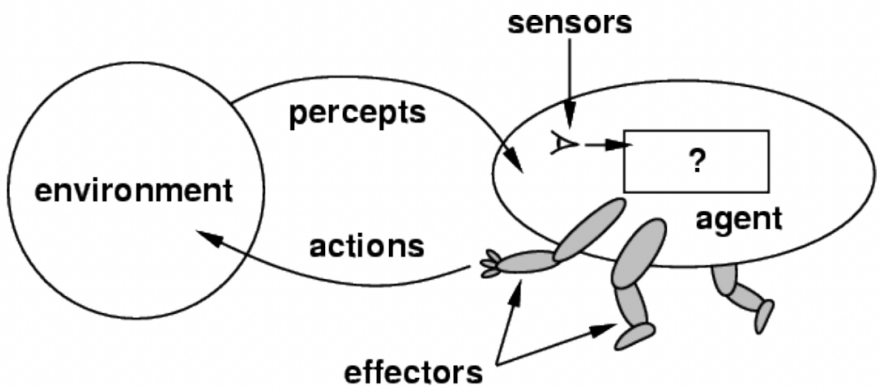
$$f : (p_1, p_2, \ldots, p_n) \to a$$

L'obiettivo dell'IA è progettare agenti che massimizzino una misura di utilità U, in base a una funzione obiettivo definita, in modo che l'azione selezionata a ottimizzi i risultati in relazione all'ambiente. In altre parole, un sistema di IA cerca di eseguire il calcolo più efficiente per risolvere problemi in base alle informazioni disponibili, ottimizzando una funzione obiettivo in un contesto dinamico.

Intro to AI during my master (years 2002-04)

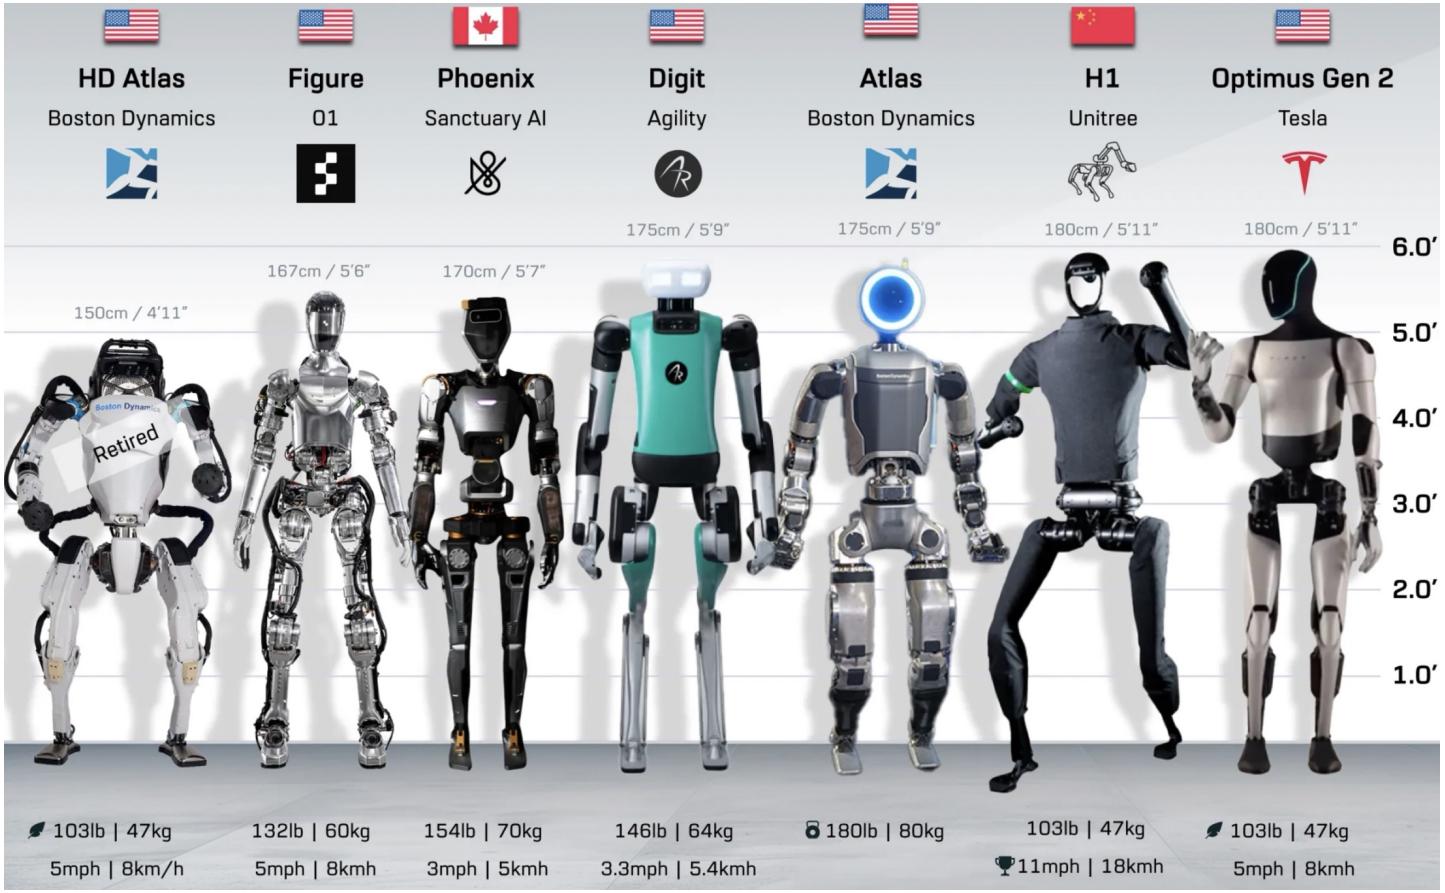
How do you design an intelligent agent?

- Definition: An **intelligent agent** perceives its environment via **sensors** and acts rationally upon that environment with its **effectors**.
- A discrete agent receives **percepts** one at a time, and maps this percept sequence to a sequence of discrete **actions**.
- Properties
  - –Autonomous
  - –Reactive to the environment
  - –Pro-active (goal-directed)
  - –Interacts with other agents via the environment



MAIN TEXTBOOK ON AI: `https://aima.cs.berkeley.edu/`
Artificial Intelligence: A Modern Approach. Stuart Russell and Peter Norvig

Humanoid Robots (Isaac Asimov's style)



https://www.youtube.com/watch?v=F_7IPm7f1vI

https://www.youtube.com/watch?v=DrNcXgoFv20

https://www.youtube.com/watch?v=-9EM5_VFlt8

Industrial Robotic Arm (mature technology!)



QUESTION: First robot playing chess ever?

QUESTION: First robot playing chess ever?



MAELZEL'S EXHIBITION,

No. 29, St. James's Street.

The Automaton Chess Player

Being returned from *Edinburgh* and *Liverpool*, where (giving the Pawn and Move) it baffled all Competition, in upwards of 200 Games, although opposed by ALL THE BEST PLAYERS.

Has opened its Second Campaign,

WITH THE ADDITION OF THE

AUTOMATON TRUMPETER,

AND THE

Conflagration of Moscow,

In which Mr. M. has endeavoured to combine the ARTS of DESIGN, MECHANISM, and MUSIC, so as produce, by a novel Imitation of Nature, a perfect Fac Simile of the real Scene. The View is from an elevated Station on the Fortress of the *Kremlin*, at the Moment when the Inhabitants are evacuating the Capital of the Czars, and the Head of the French Columns commences it Entry. The gradual Progress of the Fire, the hurrying Bustle of the Fugitives, the Eagerness of the Invaders, and the Din of warlike Sounds, will tend to impress the Spectator with a true Idea of a Scene which baffles all Powers of Description.

The MORNING EXHIBITIONS begin at 1 and 3 o'Clock, and the EVENING EXHIBITION at 8 precisely, when GAMES will be played AGAINST ANY OPPONENT, to whom the double Advantage of A PAWN AND THE MOVE WILL BE GIVEN.

Admission 2*s*.6*d*. *Children* 1*s*.6*d*. *each.*

☞ Each Exhibition lasts One Hour. Should a Game not be finished in that Time, the Party will be at Liberty to take it down with a View to its being resumed at another Opportunity.

Mr. M. begs leave to announce that the ORCHESTRION, *the* AUTOMATON TRUMPETER, *the* CONFLAGRATION OF MOSCOW, *and the Patent for the* METRONOMES, *are to be disposed of.*



https://en.wikipedia.org/wiki/Mechanical_Turk

QUESTION: First robot playing chess ever?



MAELZEL'S EXHIBITION,

No. 29, St. James's Street.

The Automaton Chess Player

Being returned from *Edinburgh* and *Liverpool*, where (giving the Pawn and Move) it baffled all Competition, in upwards of 200 Games, although opposed by ALL THE BEST PLAYERS.

Has opened its Second Campaign,

WITH THE ADDITION OF THE

AUTOMATON TRUMPETER,

AND THE

Conflagration of Moscow,

In which Mr. M. has endeavoured to combine the ARTS of DESIGN, MECHANISM, and MUSIC, so as produce, by a novel Imitation of Nature, a perfect Fac Simile of the real Scene. The View is from an elevated Station on the Fortress of the *Kremlin*, at the Moment when the Inhabitants are evacuating the Capital of the Czars, and the Head of the French Columns commences it Entry. The gradual Progress of the Fire, the hurrying Bustle of the Fugitives, the Eagerness of the Invaders, and the Din of warlike Sounds, will tend to impress the Spectator with a true Idea of a Scene which baffles all Powers of Description.

The MORNING EXHIBITIONS begin at 1 and 3 o'Clock, and the EVENING EXHIBITION at 8 precisely, when GAMES will be played AGAINST ANY OPPONENT, to whom the double Advantage of A PAWN AND THE MOVE WILL BE GIVEN.

Admission 2*s*.6*d*. *Children* 1*s*.6*d*. *each.*

☞ Each Exhibition lasts One Hour. Should a Game not be finished in that Time, the Party will be at Liberty to take it down with a View to its being resumed at another Opportunity.

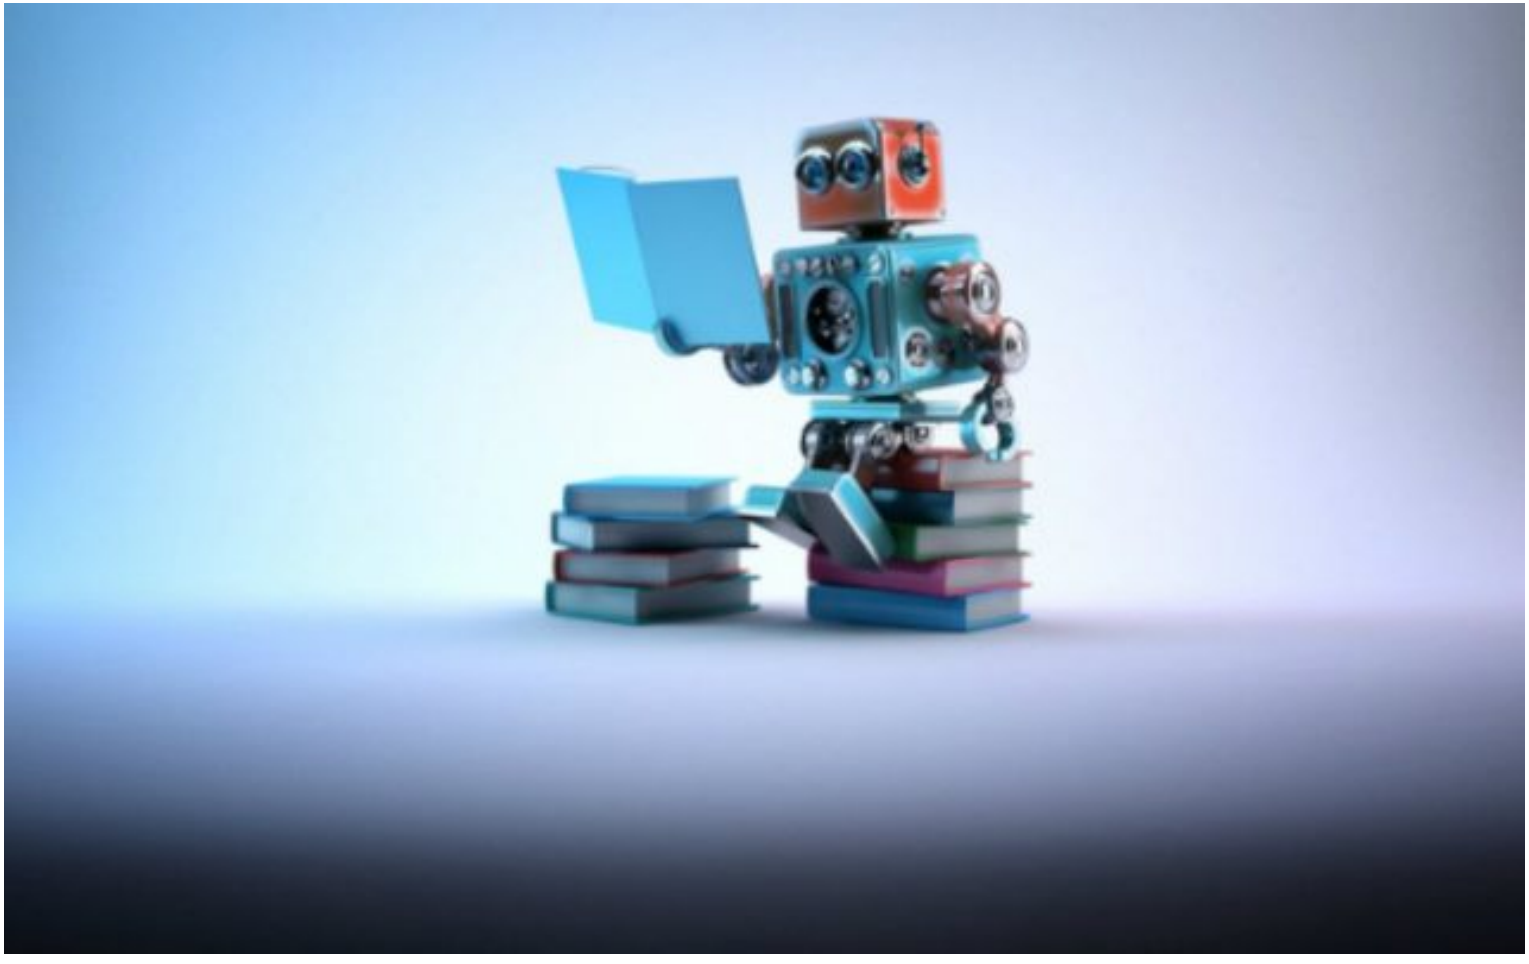
Mr. M. begs leave to announce that the ORCHESTRION, *the* AUTOMATON TRUMPETER, *the* CONFLAGRATION OF MOSCOW, *and the Patent for the* METRONOMES, *are to be disposed of.*



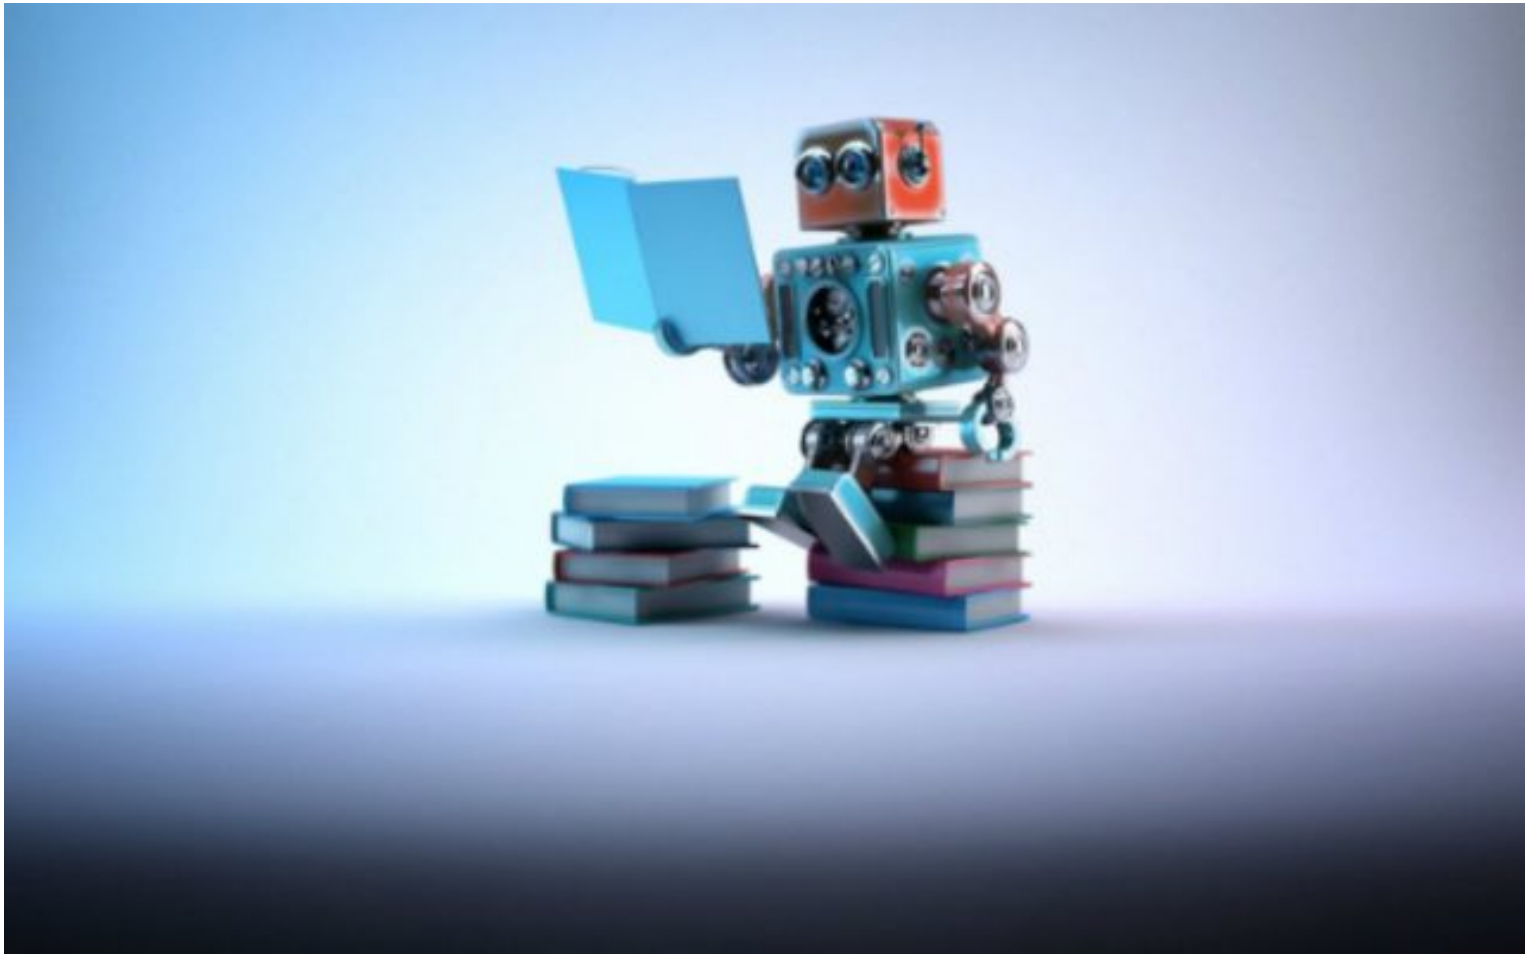
https://en.wikipedia.org/wiki/Mechanical_Turk

1995-1997: Deep Blue was a **chess-playing expert system** run on a unique purpose-built IBM supercomputer

Artificial Intelligence and Machine Learning



Artificial Intelligence and Machine Learning



"Machine take me by surprise with great frequency"
– A.M. Turing

What does it mean "Learning"?

TRECCANI ISTITUTO MAGAZINE CATALOGO SCUOLA LIBRI

imparare

Vocabolario on line

Crea un ebook con questa voce | Scaricalo ora (0) Condividi

imparare v. tr. [lat. **imparare*, comp. di *in-*[1] e *parare* «procurare»; propr. «procacciarsi una nozione», o sim.]. – **1. a.** Acquistare cognizione di qualche cosa, o fare propria una serie di cognizioni (relative a un'arte, a una scienza, a un'attività, ecc.), per mezzo dello studio, dell'esercizio, dell'osservazione, della pratica, attraverso l'esempio altrui, ecc. Con sign. generico: *è un fannullone che non ha voglia d'imparar nulla*; *da lui, abbiamo tutti qualcosa da i.*; *sbagliando s'impara* (prov.). Determinando: *i. la grammatica, la letteratura, la fisica, il disegno*; *i. una lingua straniera, i. un gioco, i. le buone maniere. I. a memoria* o *a mente*, fissare nella memoria (un brano di prosa, dei versi, ma anche brevi frasi, serie di nomi, di numeri, ecc.) in modo da saper ripetere testualmente le parole; in questo senso, anche il semplice *imparare*: *ho già imparato la poesia* (talora con la particella pron.: *mi sono già imparato le prime due strofe*); *i. la lezione*, mandarla a memoria sia testualmente sia assimilando genericamente i concetti. **b.** Con *a* e l'infinito, apprendere a saper fare qualche cosa: *i. a leggere, a scrivere, a contare, a suonare, a cantare, a ballare, a dipingere, a cucinare*; con sign. più ampio (acquistare

Informal definition

"Machine Learning is the field of study that gives computers the ability to learn without being explicitly programmed"

Reinforcement Learning plays ATARI (Nature, 2015)



https://www.youtube.com/watch?v=V1eYniJ0Rnk

Formal definition

Definition 1 (Mitchell, 1997)

A computer program is said to **learn** from experience E with respect to some class of tasks T and performance measure P, if its performance at tasks in T, as measured by P, improves with experience E.

Formal definition

Definition 1 (Mitchell, 1997)

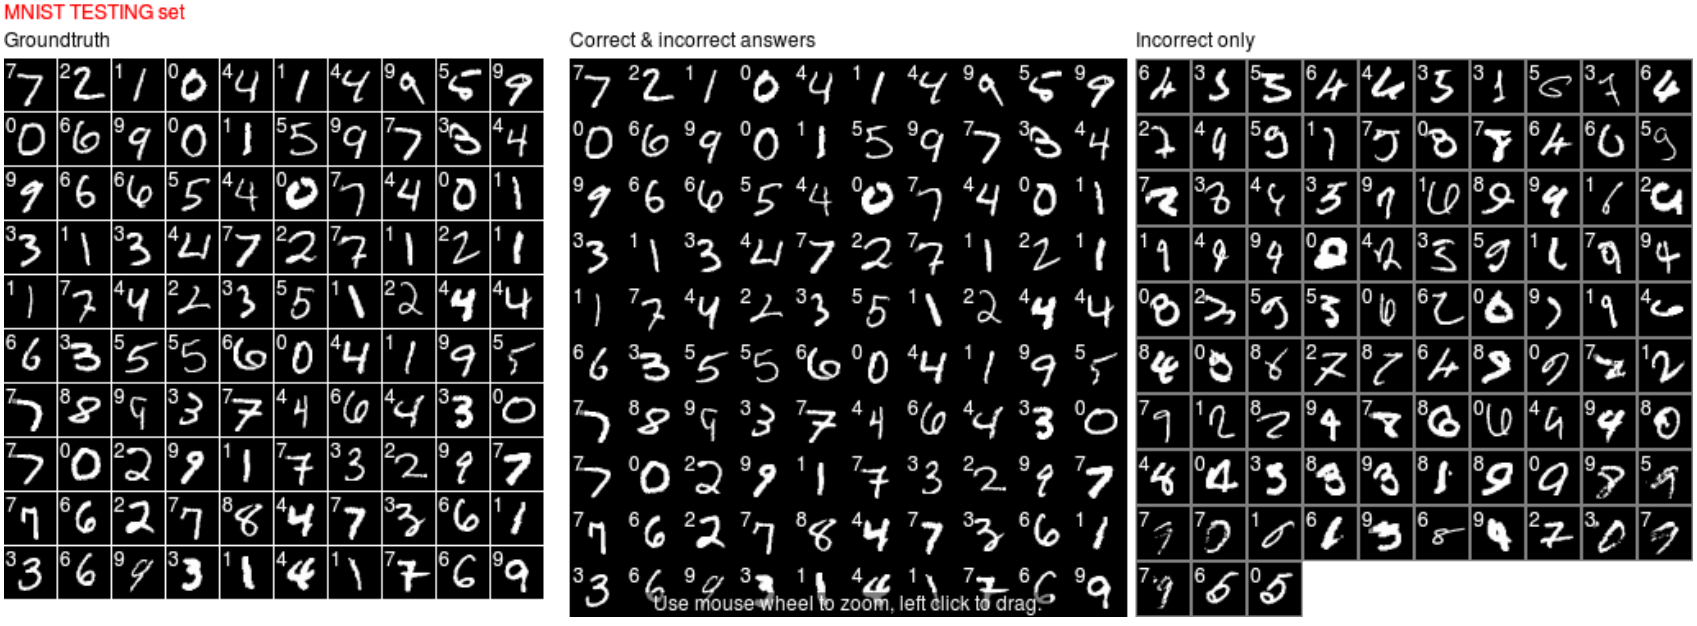
A computer program is said to **learn** from experience E with respect to some class of tasks T and performance measure P, if its performance at tasks in T, as measured by P, improves with experience E.

Example 2

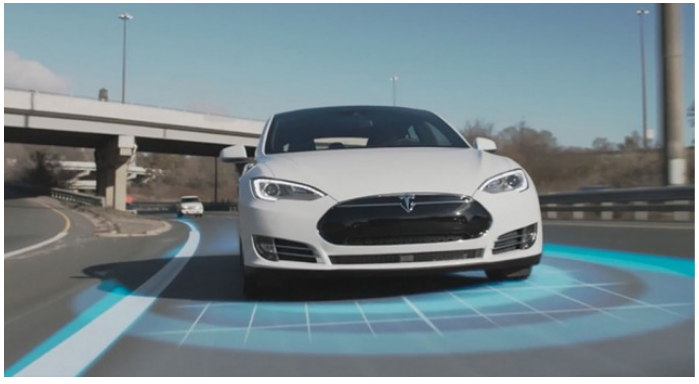
A handwriting recognition learning problem:

- **Task** T: recognizing and classifying handwritten words within images
- **Performance measure** P: percent of words correctly classified
- **Training experience** E: a dataset of handwritten words with given classifications

Handwriting recognition learning problem



Self-driving cars

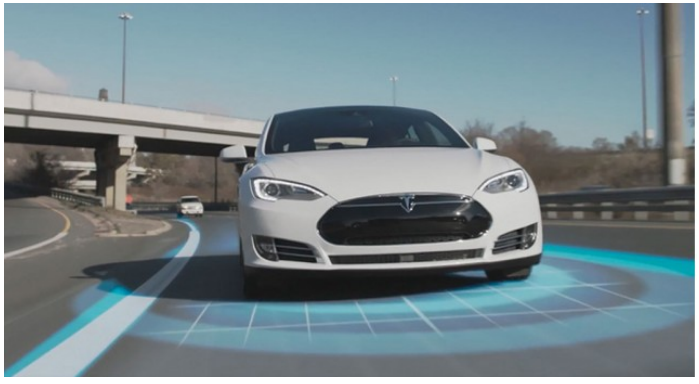


Example 3

A robot driving learning problem:

- **Task** T: driving on public streets using vision sensors
- **Performance measure** P: average distance traveled before an error (as judged by human operator)
- **Training experience** E: a sequence of images and steering commands recorded while observing a human driver

Self-driving cars



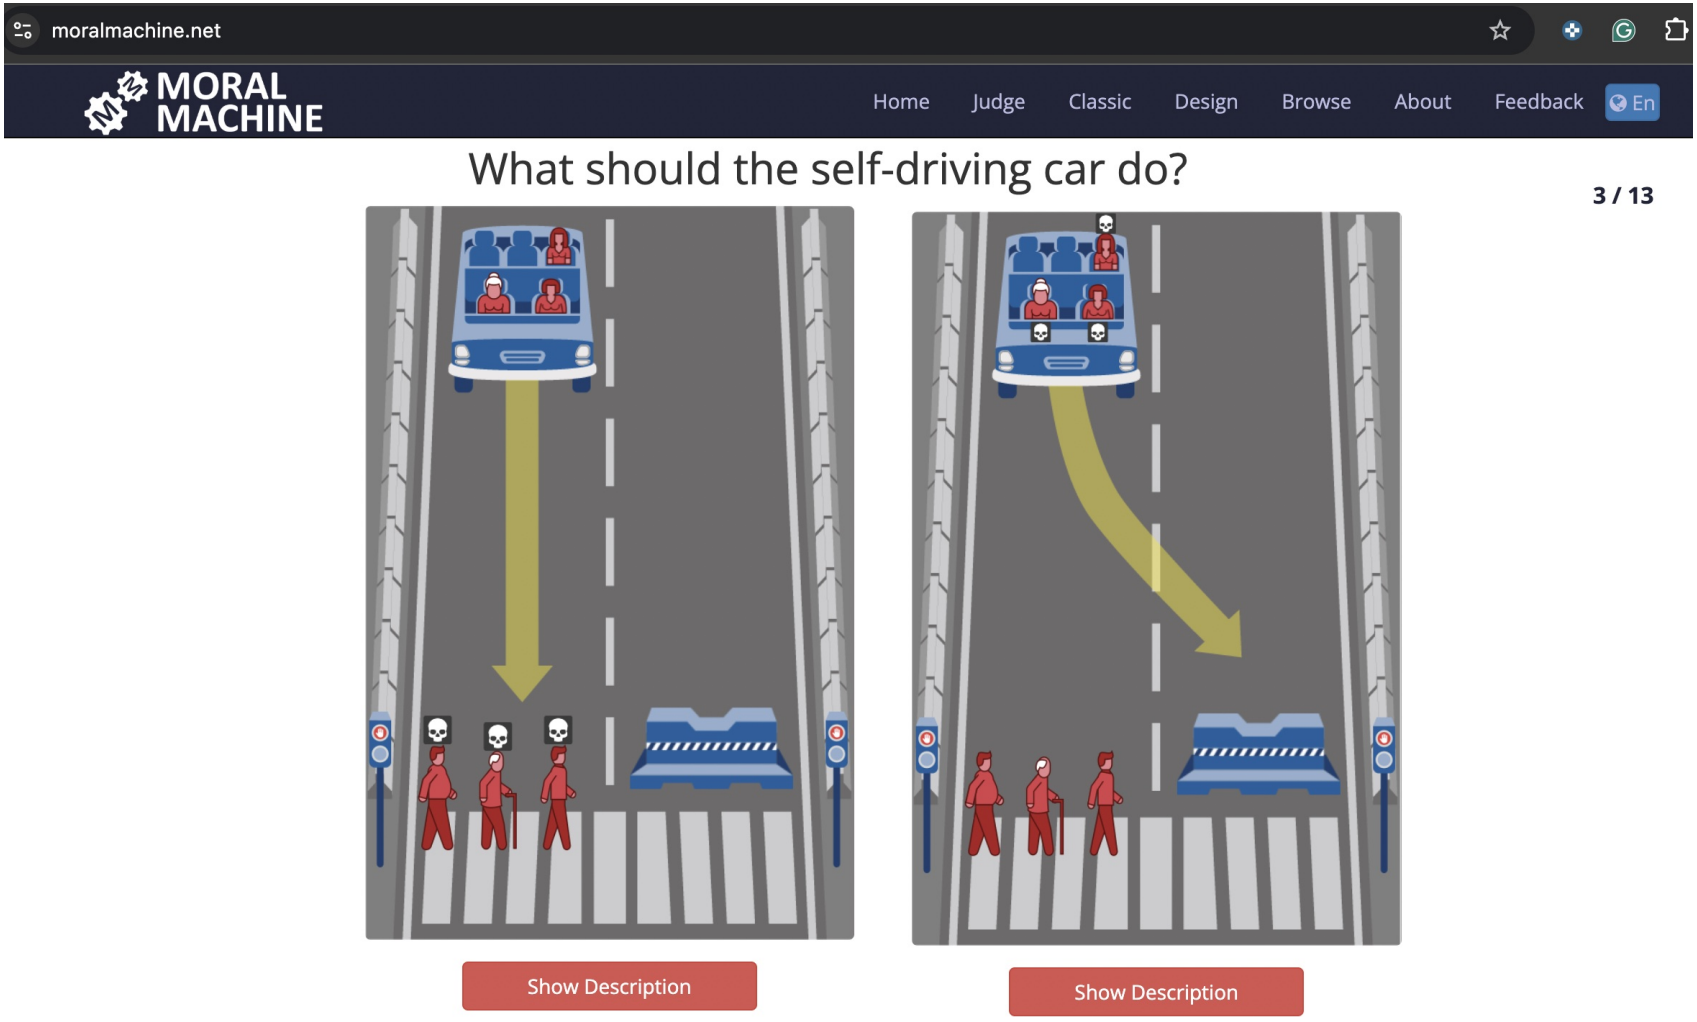


Example 3

A robot driving learning problem:

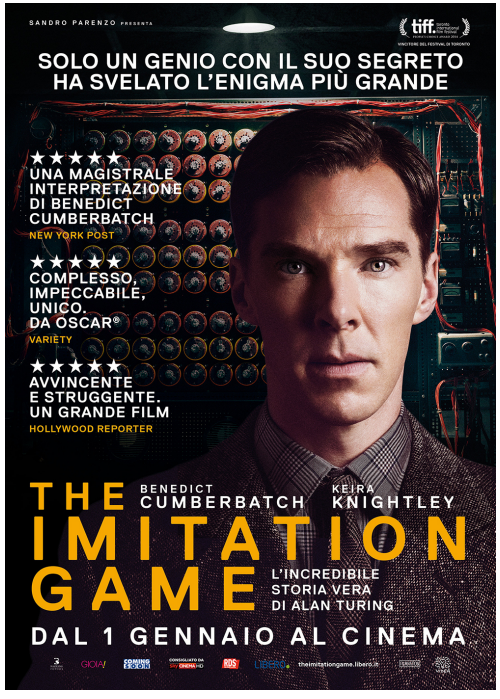
- **Task** T: driving on public streets using vision sensors
- **Performance measure** P: average distance traveled before an error (as judged by human operator)
- **Training experience** E: a sequence of images and steering commands recorded while observing a human driver

Handwriting recognition learning problem



https://www.moralmachine.net/

"Can Machine Think?"
– *A.M. Turing*



VOL. LIX. NO. 236.] [October, 1950

MIND

A QUARTERLY REVIEW

OF

PSYCHOLOGY AND PHILOSOPHY

I.—COMPUTING MACHINERY AND INTELLIGENCE

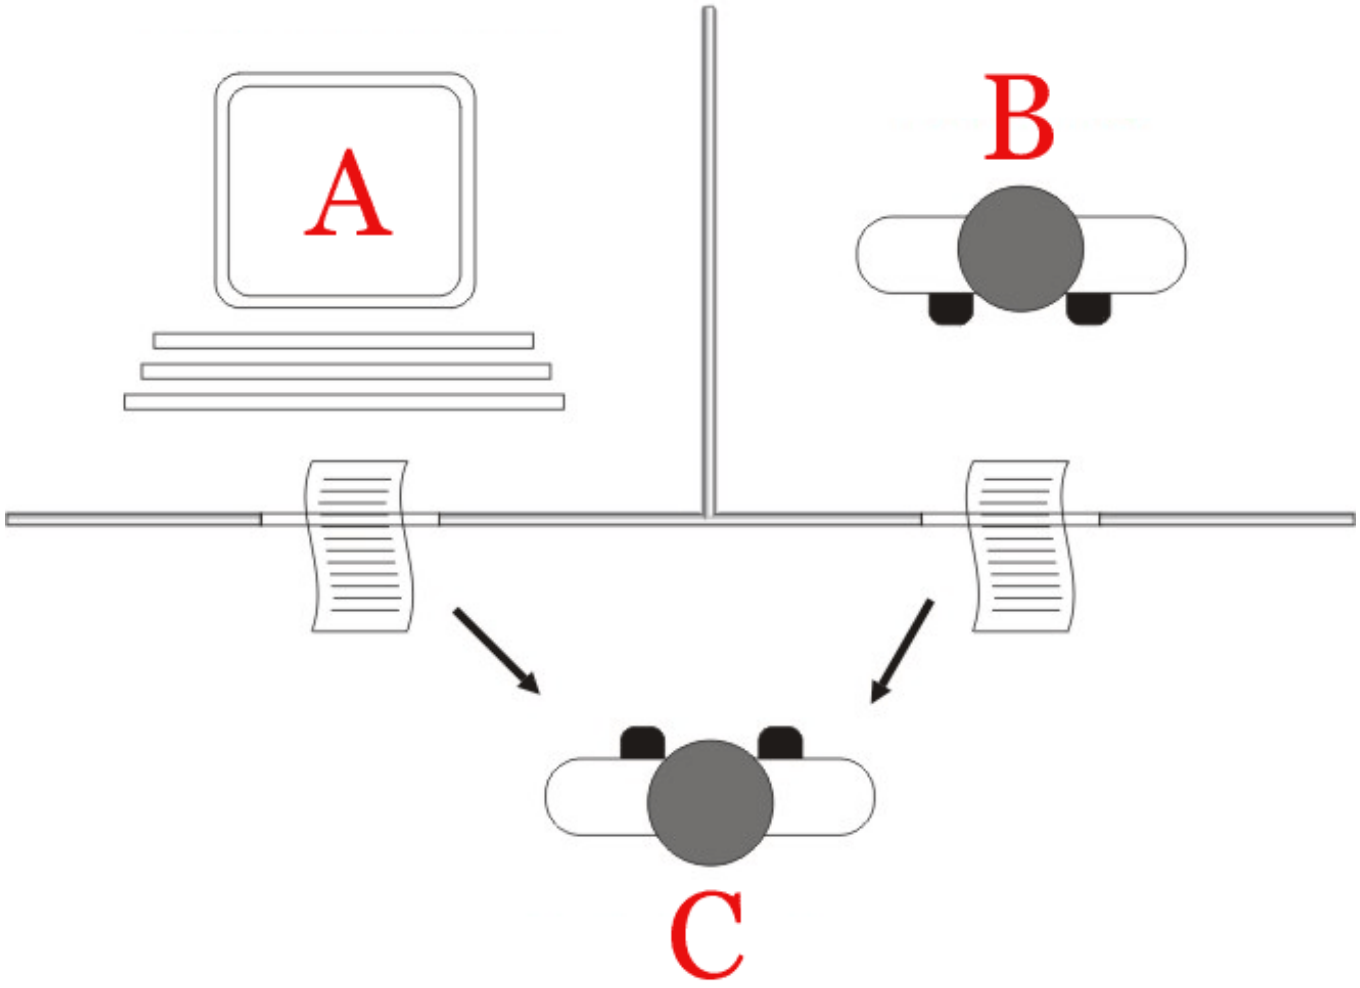
BY A. M. TURING

1. *The Imitation Game.*

I PROPOSE to consider the question, 'Can machines think?' This should

Imitation Game: Turing's Test

https://en.wikipedia.org/wiki/Turing_test



Essential Filmography

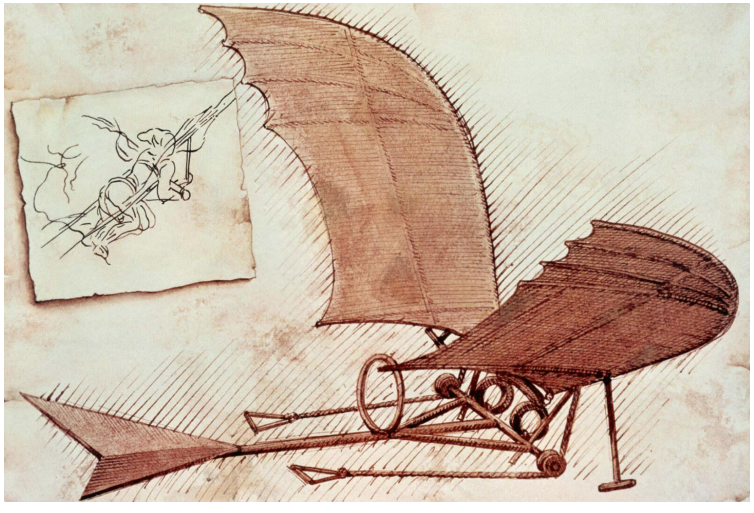
- **War games, 1983** (computer che gioca contro se stesso)
- **Terminator, 1984** (prima rete neurale addestrate in robot umanoide)
- **L'uomo bicentenario, 1999** (dovrebbe essere una miniserie Netflix!)
- **The imitation game, 2014** (Alan Turing, padre dell'AI moderna)

My view on AI: Natural Flight vs. Artificial Flight



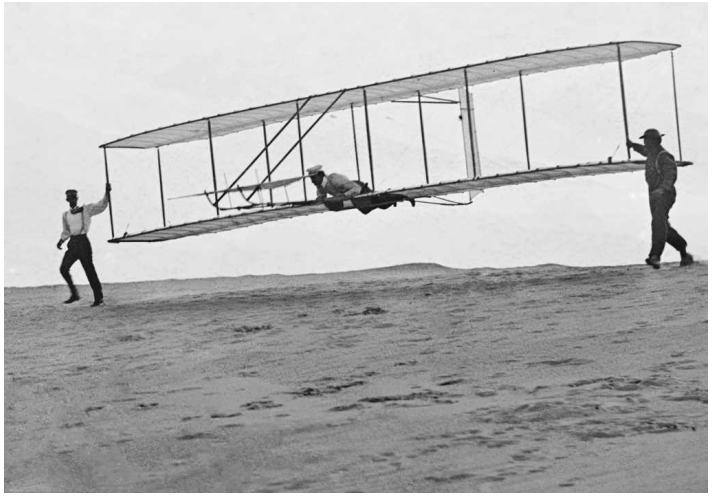
My view on AI: Natural Flight vs. Artificial Flight

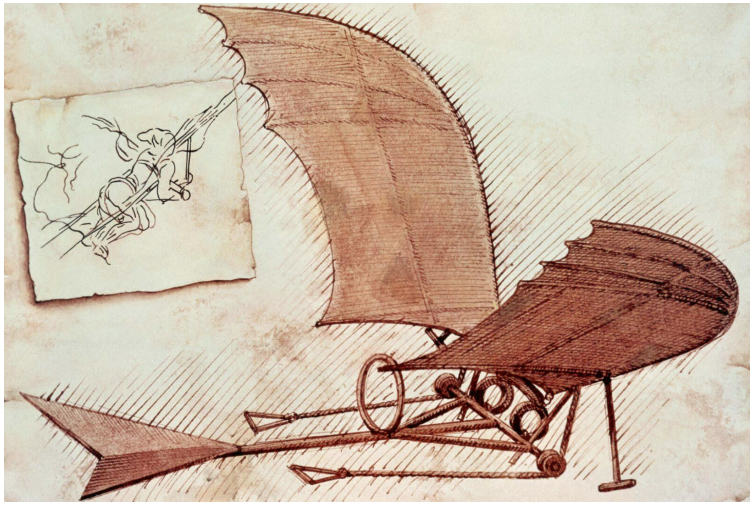




My view on AI: Natural Flight vs. Artificial Flight

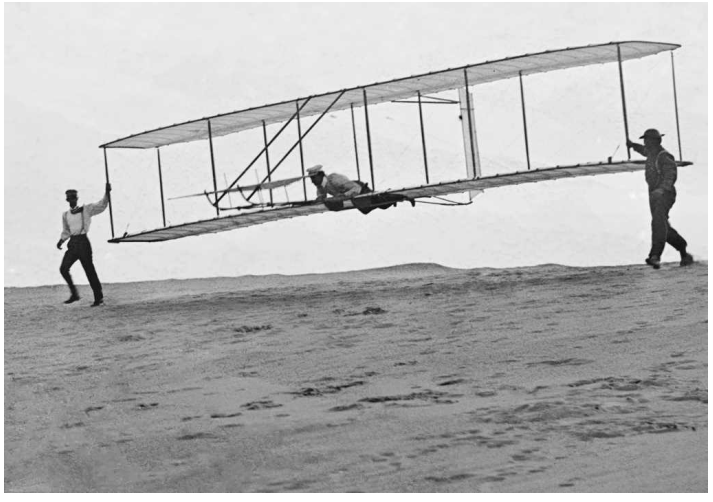


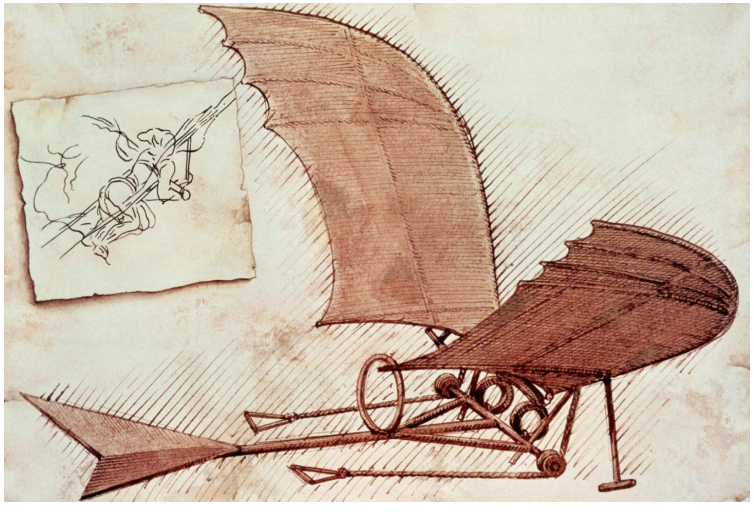




My view on AI: Natural Flight vs. Artificial Flight









Learning Machines, by A.M. Turing

In the process of trying to imitate an adult human mind we are bound to think a good deal about the process which has brought it to the state that it is in. We may notice three components,

(*a*) The initial state of the mind, say at birth,

(*b*) The education to which it has been subjected,

(*c*) Other experience, not to be described as education, to which it has been subjected.

Learning Machines, by A.M. Turing

In the process of trying to imitate an adult human mind we are bound to think a good deal about the process which has brought it to the state that it is in. We may notice three components,

(*a*) The initial state of the mind, say at birth,

(*b*) The education to which it has been subjected,

(*c*) Other experience, not to be described as education, to which it has been subjected.

456 A. M. TURING:

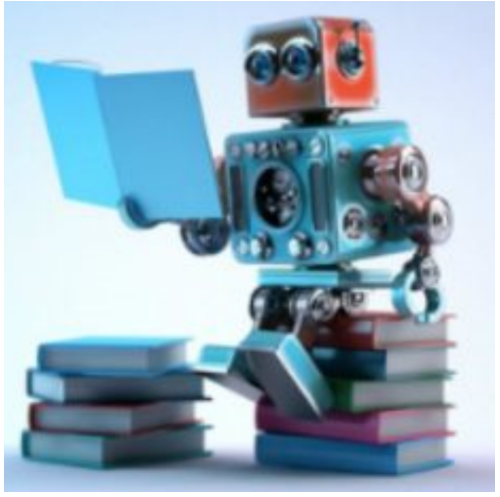
Instead of trying to produce a programme to simulate the adult mind, why not rather try to produce one which simulates the child's? If this were then subjected to an appropriate course of education one would obtain the adult brain. Presumably the child-brain is something like a note-book as one buys it from the stationers. Rather little mechanism, and lots of blank sheets. (Mechanism and writing are from our point of view almost synonymous.) Our hope is that there is so little mechanism in the child-brain that something like it can be easily programmed. The amount of work in the education we can assume, as a first approximation, to be much the same as for the human child.

Learning Machines, by A.M. Turing

the different genetical combinations that had been tried, so as to avoid trying them again?

We may hope that machines will eventually compete with men in all purely intellectual fields. But which are the best ones to start with? Even this is a difficult decision. Many people think that a very abstract activity, like the playing of chess, would be best. It can also be maintained that it is best to provide the machine with the best sense organs that money can buy, and then teach it to understand and speak English. This process could follow the normal teaching of a child. Things would be pointed out and named, etc. Again I do not know what the right answer is, but I think both approaches should be tried.

We can only see a short distance ahead, but we can see plenty there that needs to be done.

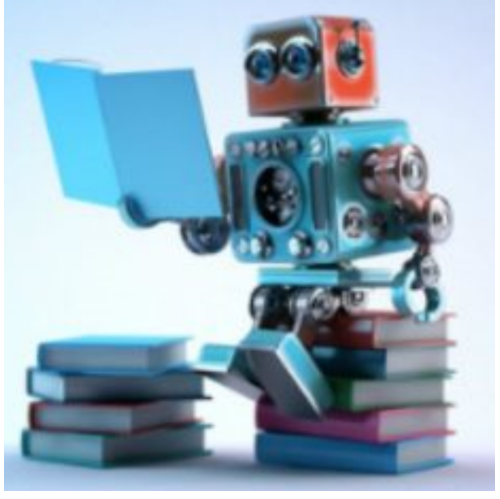


Learning Machines, by A.M. Turing

the different genetical combinations that had been tried, so as to avoid trying them again?

We may hope that machines will eventually compete with men in all purely intellectual fields. But which are the best ones to start with? Even this is a difficult decision. Many people think that a very abstract activity, like the playing of chess, would be best. It can also be maintained that it is best to provide the machine with the best sense organs that money can buy, and then teach it to understand and speak English. This process could follow the normal teaching of a child. Things would be pointed out and named, etc. Again I do not know what the right answer is, but I think both approaches should be tried.

We can only see a short distance ahead, but we can see plenty there that needs to be done.



My doubt: Will a computer *love* reading?

In filosofia: AI debole vs. AI forte

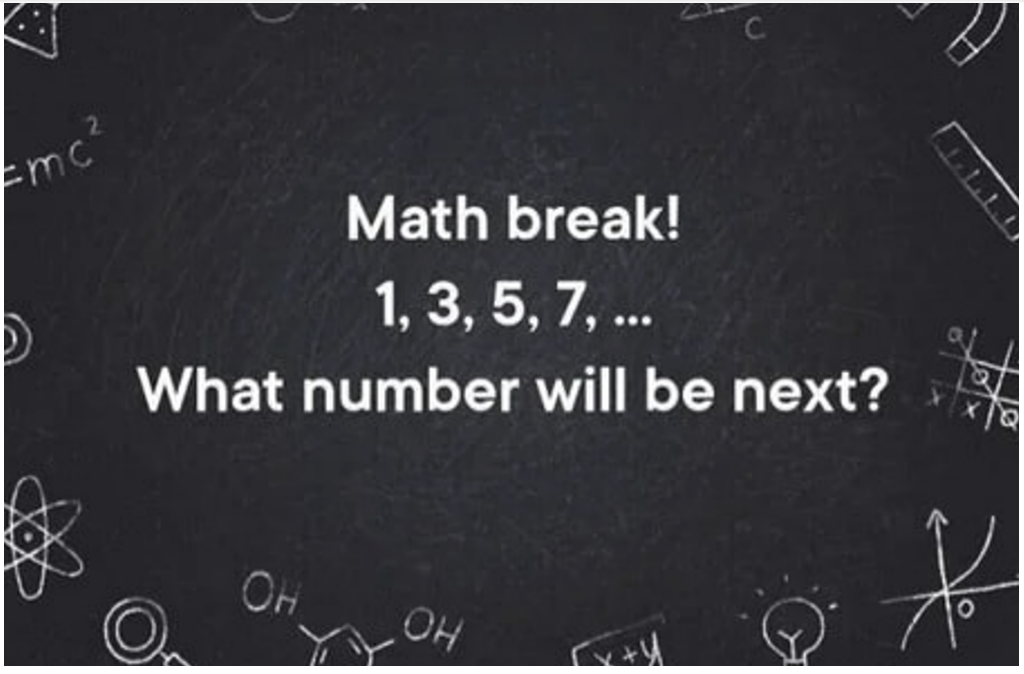
- **AI DEBOLE**: Si riferisce a sistemi di intelligenza artificiale che sono progettati per risolvere problemi specifici o svolgere compiti predefiniti (ad esempio, riconoscimento vocale, giocare a scacchi, analizzare dati). L'AI debole non possiede una vera comprensione o coscienza, ma simula solo l'intelligenza in compiti limitati.

In filosofia: AI debole vs. AI forte

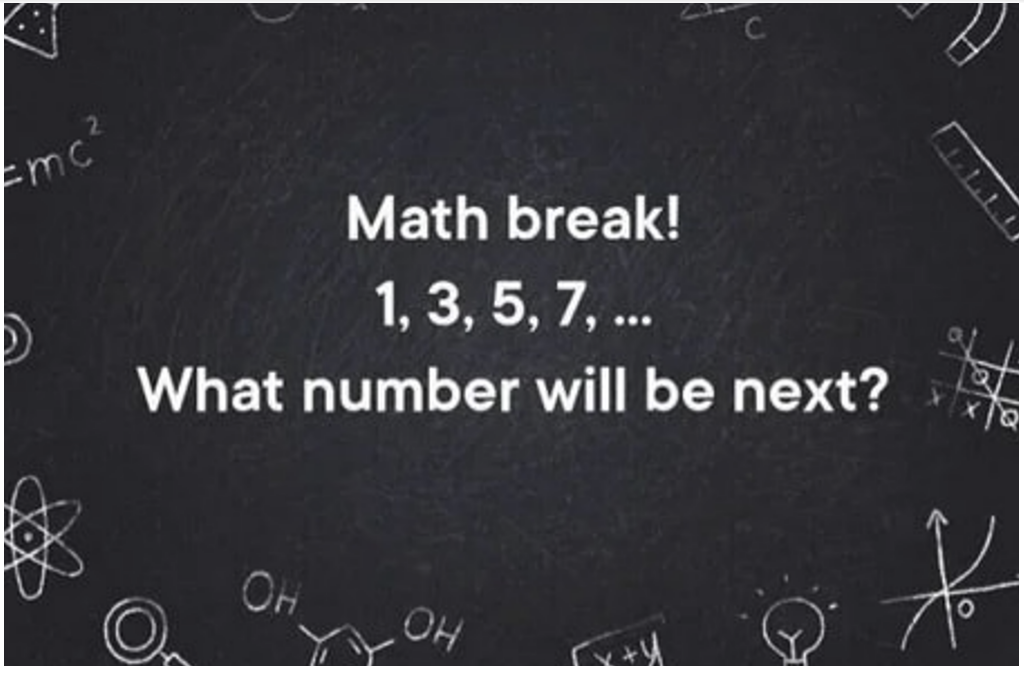
- **AI DEBOLE**: Si riferisce a sistemi di intelligenza artificiale che sono progettati per risolvere problemi specifici o svolgere compiti predefiniti (ad esempio, riconoscimento vocale, giocare a scacchi, analizzare dati). L'AI debole non possiede una vera comprensione o coscienza, ma simula solo l'intelligenza in compiti limitati.

- **AI FORTE**: Si riferisce a una forma di intelligenza artificiale teorica che non solo simula l'intelligenza, ma possiede una vera coscienza, autoconsapevolezza e comprensione. Secondo i sostenitori dell'AI forte, un sistema di AI forte sarebbe in grado di pensare, ragionare e avere esperienze soggettive, come un essere umano.

Quiz time!

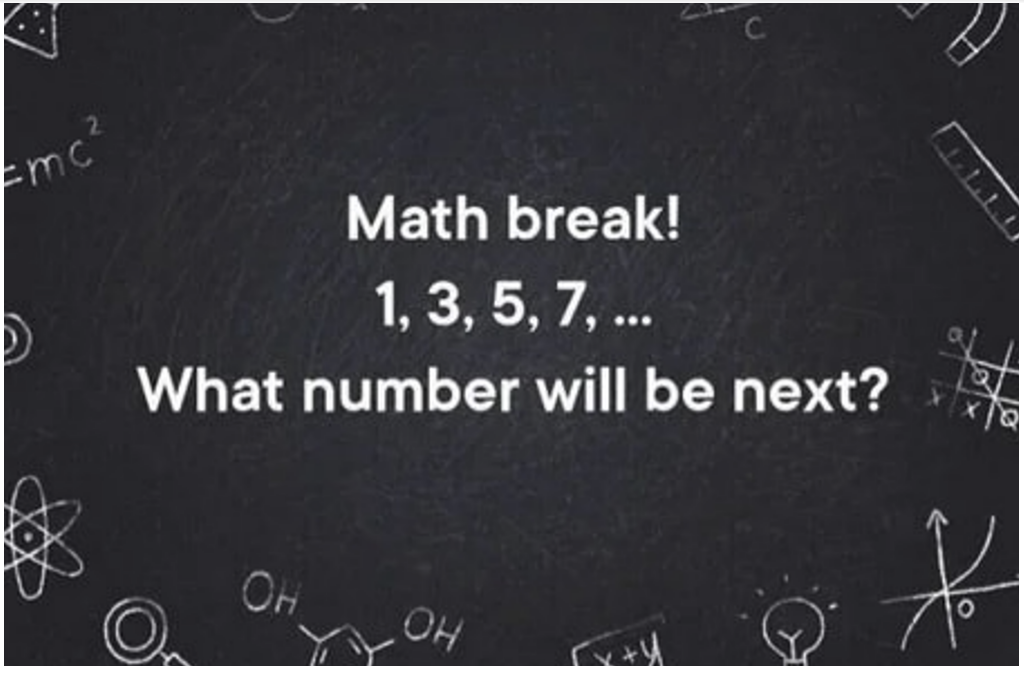


Quiz time!



The next number is **217341**

Quiz time!

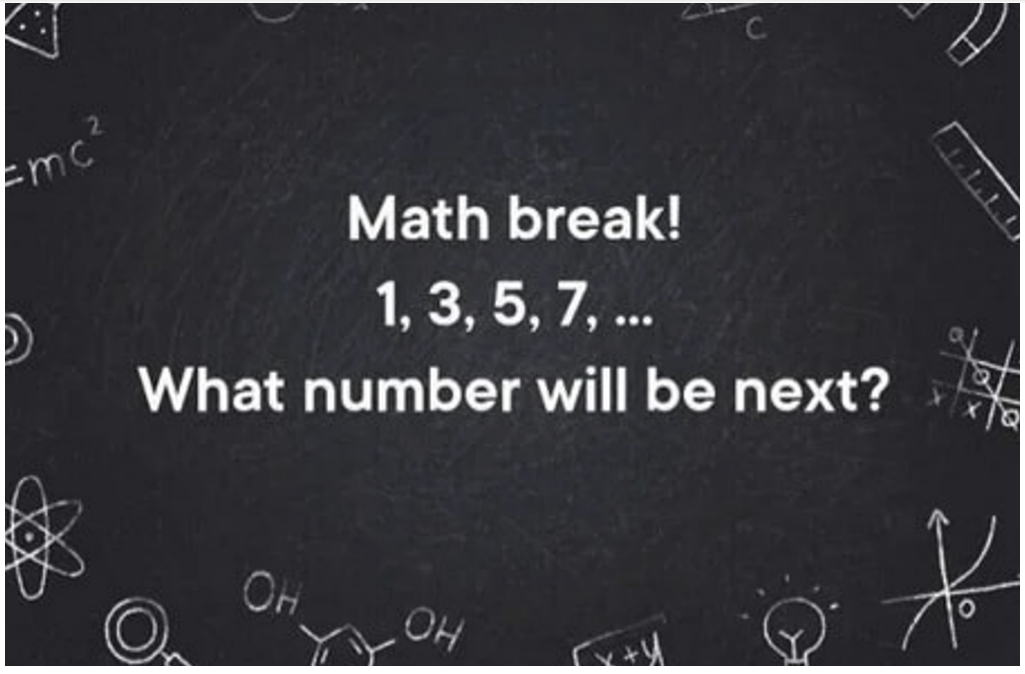


The next number is **217341** IF

$$f(x) = \frac{18111}{2}x^4 - 90555x^3 + \frac{633885}{2}x^2 - 452773x + 217331$$

since then

Quiz time!



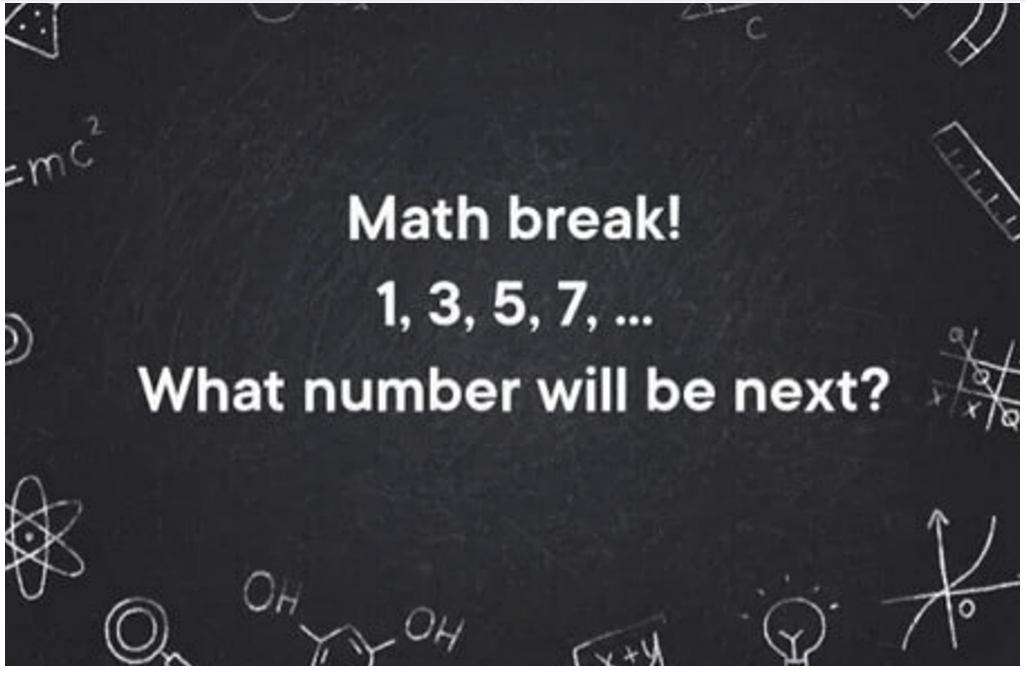
The next number is **217341** IF

$$f(x) = \frac{18111}{2}x^4 - 90555x^3 + \frac{633885}{2}x^2 - 452773x + 217331$$

since then

$$f(1) = 1,$$

Quiz time!



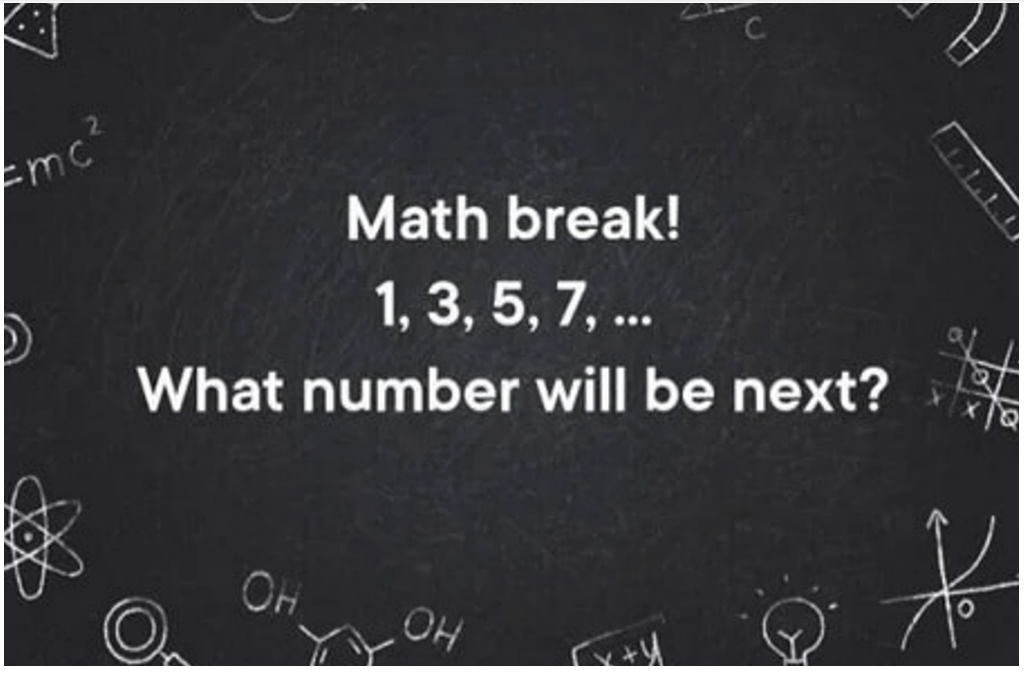
The next number is **217341** IF

$$f(x) = \frac{18111}{2}x^4 - 90555x^3 + \frac{633885}{2}x^2 - 452773x + 217331$$

since then

$$f(1) = 1, \quad f(2) = 3,$$

Quiz time!



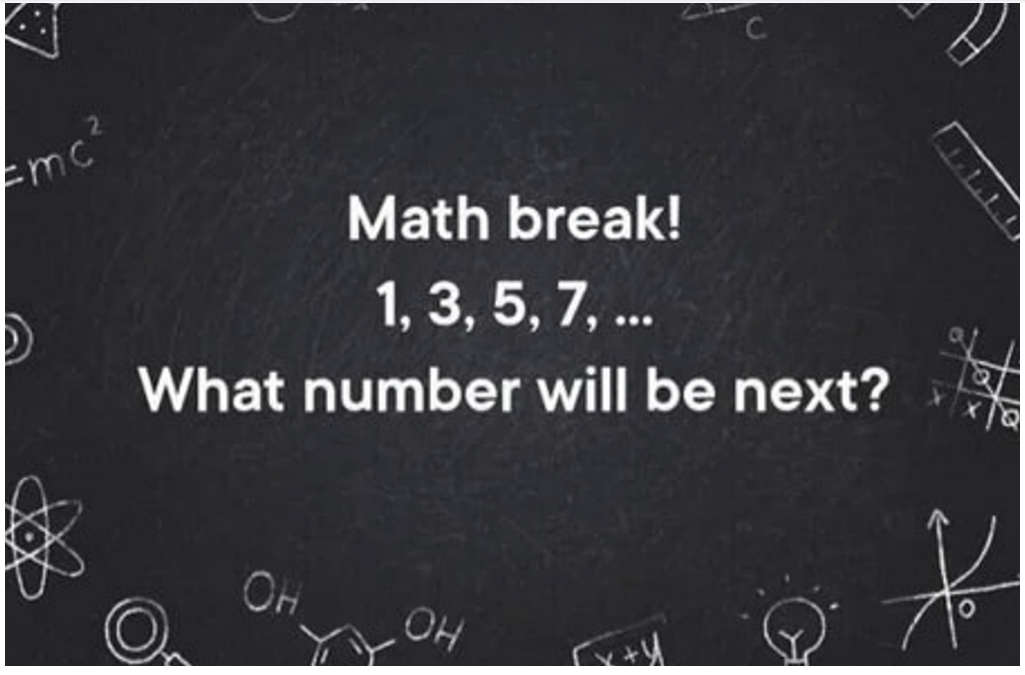
The next number is **217341** IF

$$f(x) = \frac{18111}{2}x^4 - 90555x^3 + \frac{633885}{2}x^2 - 452773x + 217331$$

since then

$$f(1) = 1, \quad f(2) = 3, \quad f(3) = 5,$$

Quiz time!



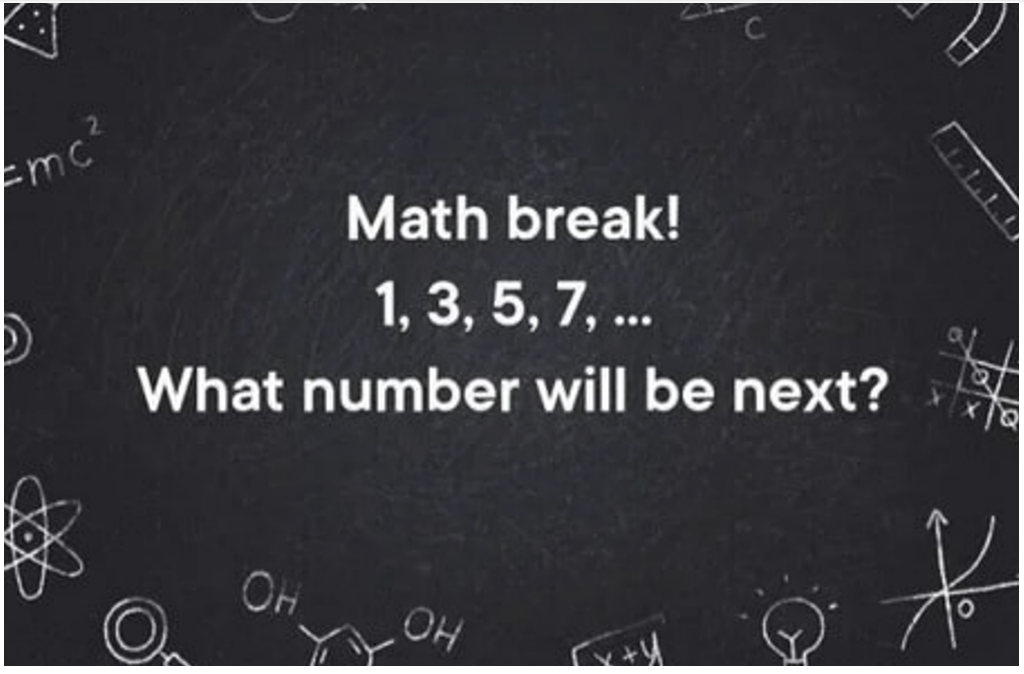
The next number is **217341** IF

$$f(x) = \frac{18111}{2}x^4 - 90555x^3 + \frac{633885}{2}x^2 - 452773x + 217331$$

since then

$$f(1) = 1, \quad f(2) = 3, \quad f(3) = 5, \quad f(4) = 7,$$

Quiz time!



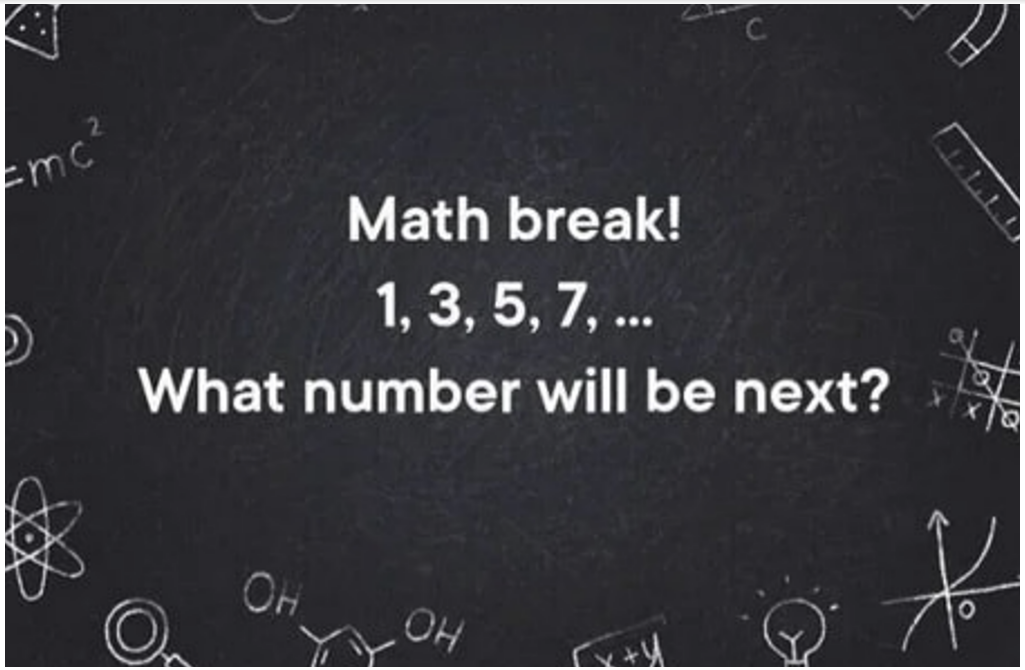
The next number is **217341** IF

$$f(x) = \frac{18111}{2}x^4 - 90555x^3 + \frac{633885}{2}x^2 - 452773x + 217331$$

since then

$$f(1) = 1, \quad f(2) = 3, \quad f(3) = 5, \quad f(4) = 7, \quad f(5) = 217341$$

Quiz time!



The next number is **217341** IF

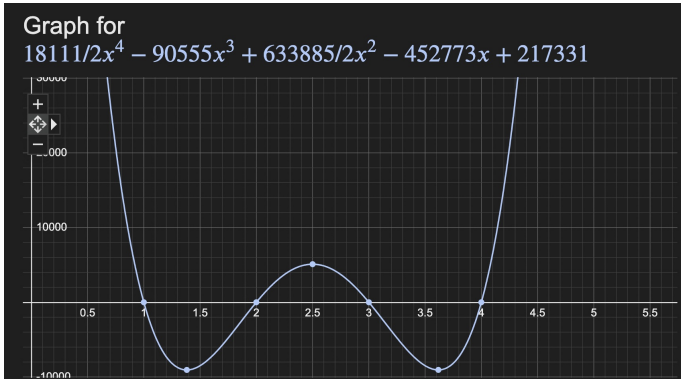
$$f(x) = \frac{18111}{2}x^4 - 90555x^3 + \frac{633885}{2}x^2 - 452773x + 217331$$

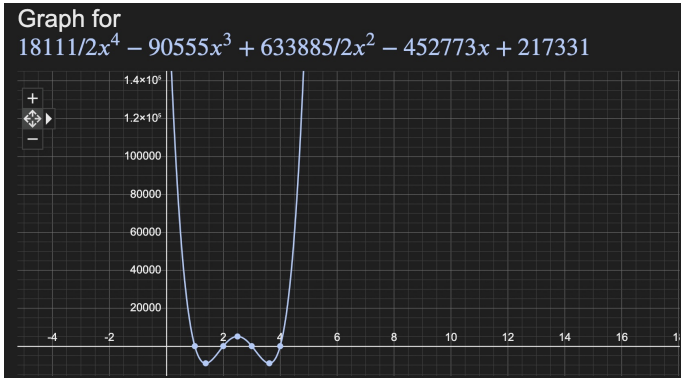
since then

$$f(1) = 1, \quad f(2) = 3, \quad f(3) = 5, \quad f(4) = 7, \quad f(5) = 217341$$

REMARK: Learning is finding a function that "*explains*" the data.

Fitting polynomial





$$f(x) = w_4 x^4 + w_3 x^3 + w_2 x^2 + w_1 x + w_0$$

Learning: trovare i valore di w_4, w_3, w_2, w_1, w_0 che minimizzano l'errore sui dati.

Outline

1 Introduction

2 ML tasks

3 Neural Networks

4 Generative AI

5 AI4Pharma

Math and Images

Disclaimer: "Images are just numbers!"

Math and Images

Disclaimer: "Images are just numbers!"



64x64 pixels
MedNIST Challenge

Math and Images

Disclaimer: "Images are just numbers!"



64x64 pixels
MedNIST Challenge



16x16 pixels
Downsampling

Math and Images

Disclaimer: "Images are just numbers!"



64x64 pixels
MedNIST Challenge



16x16 pixels
Downsampling

17	14	13	12	15	14	14	87	82	16	14	14	13	14	16	21
15	12	11	10	104	77	10	123	116	23	101	13	10	12	14	19
13	10	10	11	115	108	13	146	146	63	117	29	8	10	11	17
12	9	90	64	106	140	11	159	145	114	148	21	6	7	10	16
11	7	99	123	85	152	98	184	184	165	161	17	8	8	13	16
11	6	44	135	110	166	152	174	177	168	171	5	8	93	114	15
11	6	5	139	162	153	185	195	189	196	163	5	102	171	139	16
10	5	6	46	161	168	208	218	214	227	177	6	165	173	54	15
9	4	25	22	170	188	212	210	210	230	170	158	207	156	6	11
8	3	241	147	119	217	224	225	225	244	182	215	214	86	6	11
5	1	91	249	228	211	247	253	253	254	234	216	158	5	5	11
5	1	16	186	81	195	255	255	255	255	252	152	5	4	4	11
6	2	1	1	3	234	255	255	255	253	138	4	0	1	3	10
6	2	2	2	5	255	255	255	255	218	27	5	1	2	3	10
6	3	2	2	95	255	251	255	255	115	2	2	1	1	3	10
8	4	4	3	171	255	255	255	255	5	3	3	3	3	4	10

Image as matrix of numbers

Math, Color Images, and Tensors



RGB Image

Math, Color Images, and Tensors



RGB Image



Red channel

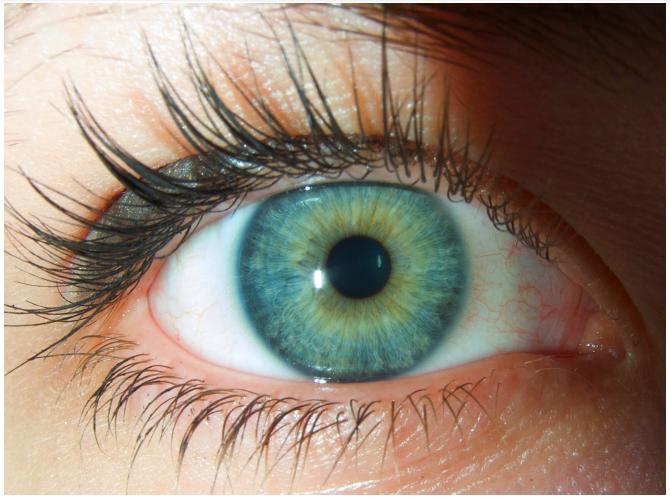


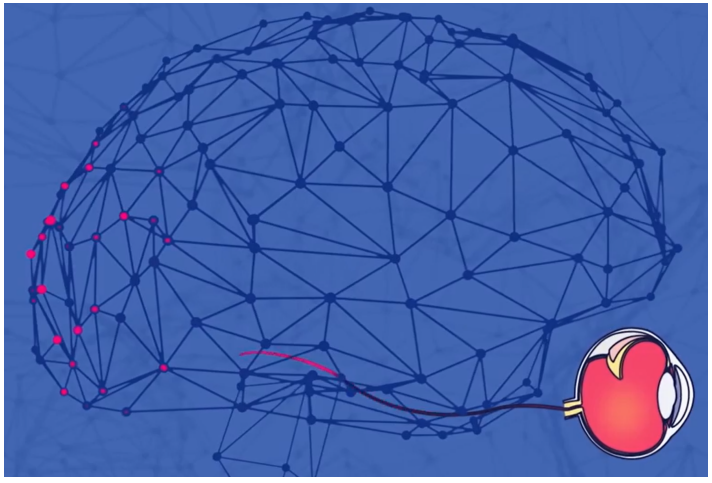
Green channel

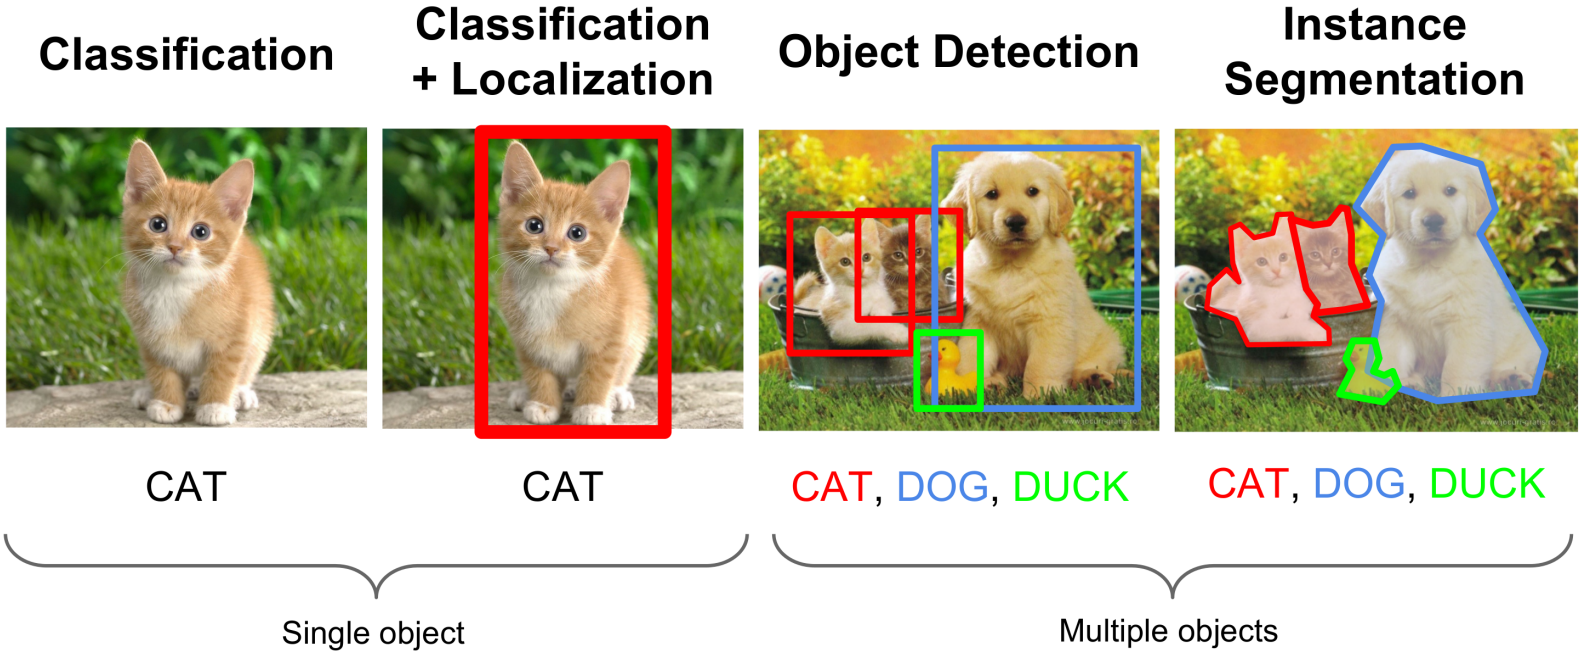


Blu channel

Human Vision vs. Artificial Vision: Where are we?







Machine Learning Tasks

Machine Learning Tasks are typically classified into:

1. Supervised Learning
2. Unsupervised Learning
3. Reinforcement Learning

Definition 4 (1. Supervised Learning)

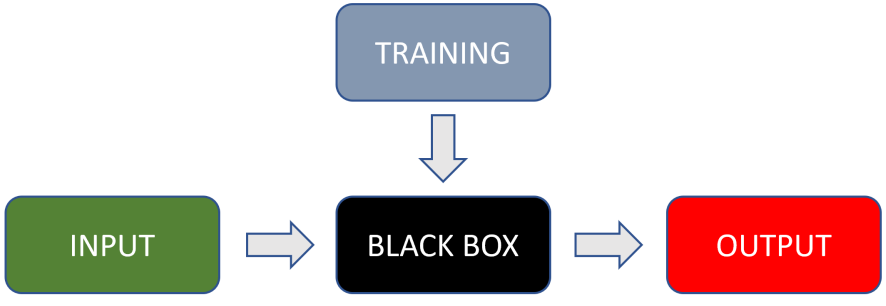
Supervised Learning is the task of learning (inferring) a function f that maps input vectors to their corresponding target vectors, by using a dataset containing a given set of pairs of $(input, output)$ samples.
Examples:

- Regression: the output vectors contain continuous values
- Classification: the output vectors contain discrete categories.
  Special case: **Binary Classification**.

Definition 4 (1. Supervised Learning)

Supervised Learning is the task of learning (inferring) a function f that maps input vectors to their corresponding target vectors, by using a dataset containing a given set of pairs of $(input, output)$ samples.
Examples:

- Regression: the output vectors contain continuous values
- Classification: the output vectors contain discrete categories.
  Special case: **Binary Classification**.



- Learns known patterns
- Labeled input data
- Predicts outcome

1. Supervised Learning: Binary Classification



$\Rightarrow$ { 0=NORMAL ; 1=PNEUMONIA }

1. Supervised Learning: Binary Classification



$\Rightarrow$ $\begin{cases} \text{0=NORMAL} \\ \text{1=PNEUMONIA} \end{cases}$

$\boldsymbol{x} \in X$: Input data

$y \in Y = \{0, 1\}$: Target

1. Supervised Learning: Binary Classification



$\Rightarrow$ { 0=NORMAL ; 1=PNEUMONIA }

$\boldsymbol{x} \in X$: Input data

$y \in Y = \{0, 1\}$: Target

CLASSIFICATION TASK: Find a function

$$f_w : X \to Y$$

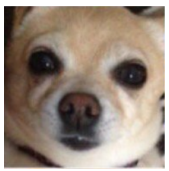
that maps each input to the *correct* target.

Dog or muffins?

Take $Y = \{-1, 1\}$, with $y = 1$ means **DOG**, and $y = -1$ means **MUFFINS**.

Dog or muffins?

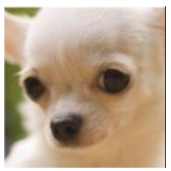
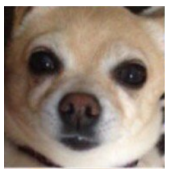
Take $Y = \{-1, 1\}$, with $y = 1$ means **DOG**, and $y = -1$ means **MUFFINS**.

$\boldsymbol{x}_1 =$  $\boldsymbol{x}_2 =$  $\boldsymbol{x}_3 =$  $\boldsymbol{x}_4 =$  $\boldsymbol{x}_5 =$ 

$y_1 = 1$ $y_2 = 1$ $y_3 = -1$ $y_4 = 1$ $y_5 = -1$

Dog or muffins?

Take $Y = \{-1, 1\}$, with $y = 1$ means **DOG**, and $y = -1$ means **MUFFINS**.

$\boldsymbol{x}_1 =$  $\boldsymbol{x}_2 =$  $\boldsymbol{x}_3 =$  $\boldsymbol{x}_4 =$  $\boldsymbol{x}_5 =$ 

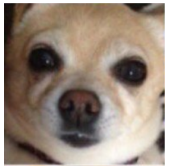
$y_1 = 1$ $y_2 = 1$ $y_3 = -1$ $y_4 = 1$ $y_5 = -1$

1. Use the Training Set $\Rightarrow$ run learning algorithm $\Rightarrow$ get

$$f_w : X \to Y$$

Dog or muffins?

Take $Y = \{-1, 1\}$, with $y = 1$ means **DOG**, and $y = -1$ means **MUFFINS**.

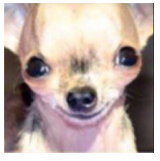
$\boldsymbol{x}_1 =$  $\boldsymbol{x}_2 =$  $\boldsymbol{x}_3 =$  $\boldsymbol{x}_4 =$  $\boldsymbol{x}_5 =$ 

$y_1 = 1$ $y_2 = 1$ $y_3 = -1$ $y_4 = 1$ $y_5 = -1$

1. Use the Training Set $\Rightarrow$ run learning algorithm $\Rightarrow$ get

$$f_w : X \to Y$$

2. Use a Test Set to validate f:

$f_w\left(\boldsymbol{x}_6 =\right.$  $\left.\right) = -1$ $f_w\left(\boldsymbol{x}_7 =\right.$  $\left.\right) = 1$

Is it a **4** or not?

We take $Y = \{-1, 1\}$, $y = 1$ means **YES, IS 4**, $y = -1$ **NO, IT'S NOT**.

Is it a **4** or not?

We take $Y = \{-1, 1\}$, $y = 1$ means **YES, IS 4**, $y = -1$ **NO, IT'S NOT**.

$\boldsymbol{x}_1 =$  $\boldsymbol{x}_2 =$  $\boldsymbol{x}_3 =$  $\boldsymbol{x}_4 =$  $\boldsymbol{x}_5 =$ 

$y_1 = 1$ $y_2 = -1$ $y_3 = 1$ $y_4 = 1$ $y_5 = -1$

Is it a **4** or not?

We take $Y = \{-1, 1\}$, $y = 1$ means **YES, IS 4**, $y = -1$ **NO, IT'S NOT**.

$\boldsymbol{x}_1 =$  $\boldsymbol{x}_2 =$  $\boldsymbol{x}_3 =$  $\boldsymbol{x}_4 =$  $\boldsymbol{x}_5 =$ 

$y_1 = 1$ $y_2 = -1$ $y_3 = 1$ $y_4 = 1$ $y_5 = -1$

1. Use the TRAINING SET $\Rightarrow$ run learning algorithm $\Rightarrow$ get

$$f_w : X \to Y$$

Is it a **4** or not?

We take $Y = \{-1, 1\}$, $y = 1$ means **YES, IS 4**, $y = -1$ **NO, IT'S NOT**.

$x_1 =$  $x_2 =$  $x_3 =$  $x_4 =$  $x_5 =$ 

$y_1 = 1$ $y_2 = -1$ $y_3 = 1$ $y_4 = 1$ $y_5 = -1$

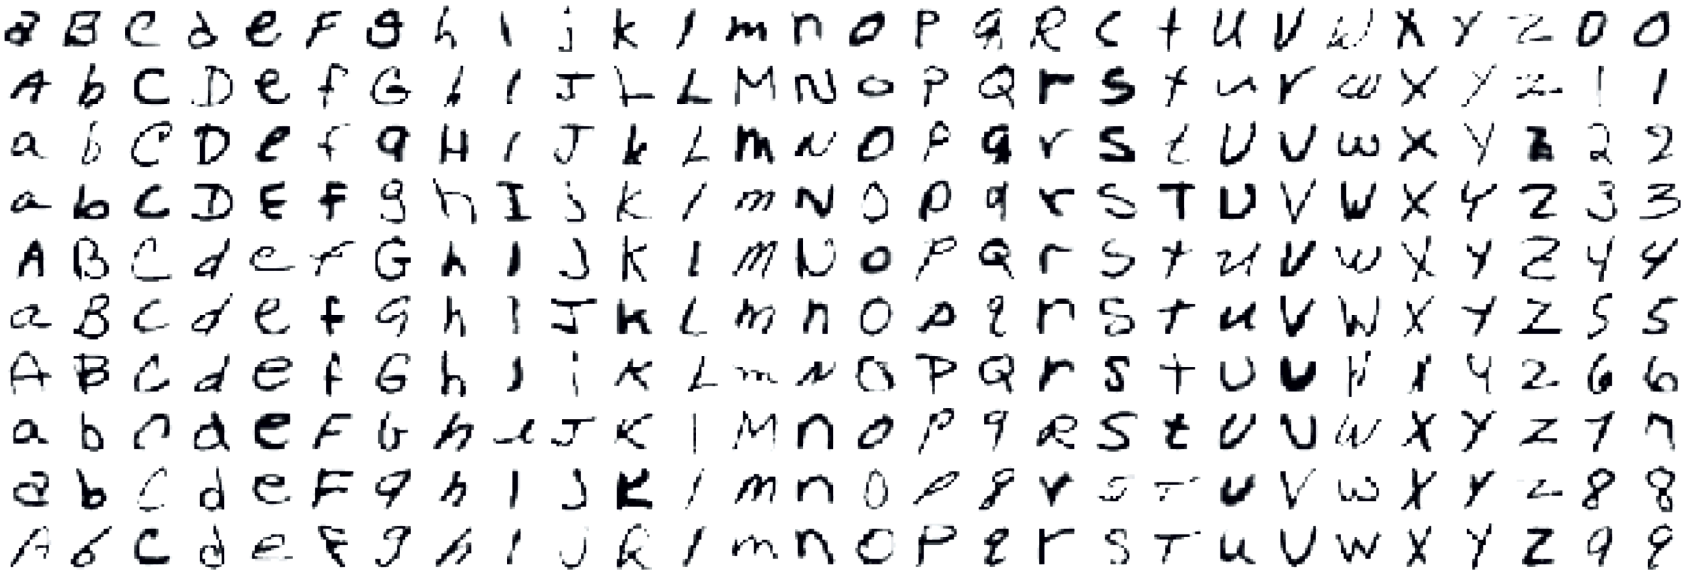
1. Use the TRAINING SET $\Rightarrow$ run learning algorithm $\Rightarrow$ get

$$f_w : X \to Y$$

2. Use a TEST SET to validate f:

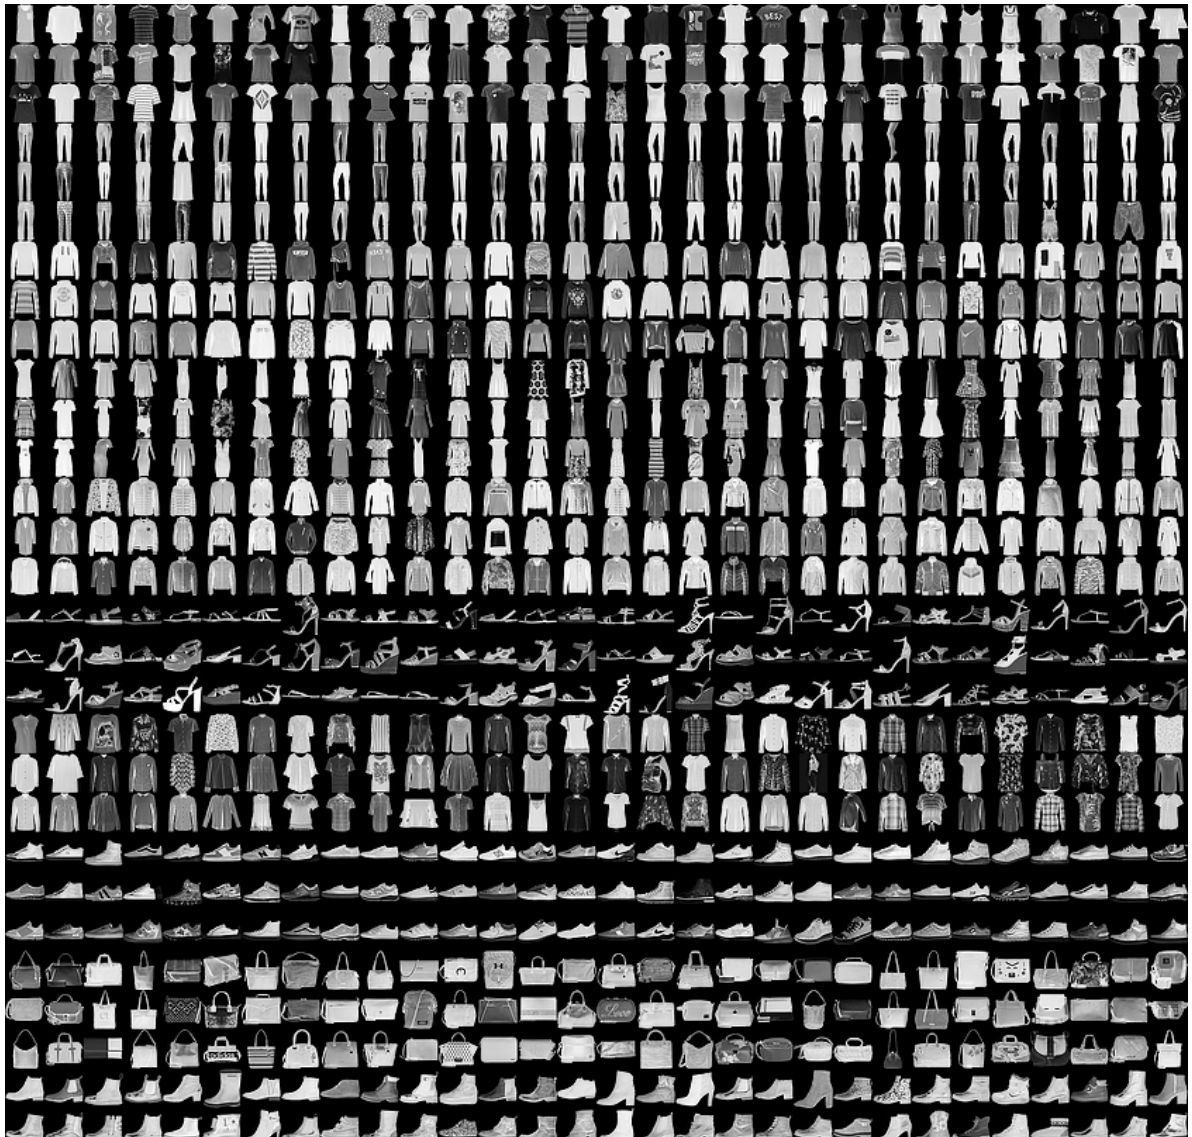
$f\left(x_6 =\right.$  $\left.\right) = -1$ $f\left(x_7 =\right.$  $\left.\right) = 1$

Extended MNIST database



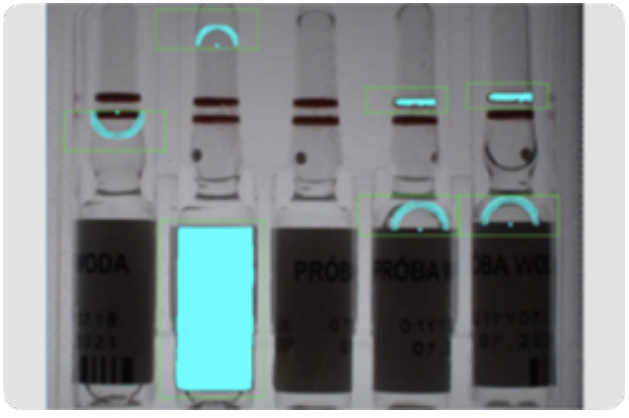
Fashion MNIST

https://github.com/zalandoresearch/fashion-mnist



See presentation by SeaVision this afternoon





https://seavision-latam.com/en/solutions/pharmaceutical-packaging-control/blister-control/

Definition 5 (2. Unsupervised Learning)

Unsupervised Learning is the task of learning (inferring) a function f that maps input vectors to their corresponding target vectors, but without any a priori knowledge about the correct mapping. Examples:

- Clustering: To group or partition the input vectors into k groups or clusters, with the vectors in each group close to each other.
- Dimensionality Reduction: The goal is to project the data from a high dimensional space down to two or three dimensions.

Definition 5 (2. Unsupervised Learning)

Unsupervised Learning is the task of learning (inferring) a function f that maps input vectors to their corresponding target vectors, but without any a priori knowledge about the correct mapping. Examples:

- Clustering: To group or partition the input vectors into k groups or clusters, with the vectors in each group close to each other.
- Dimensionality Reduction: The goal is to project the data from a high dimensional space down to two or three dimensions.



- Learns unknown patterns
- Unlabeled input data
- Find hidden patterns

2. Unsupervised Learning: Clustering

A selection from the 64-dimensional digits dataset

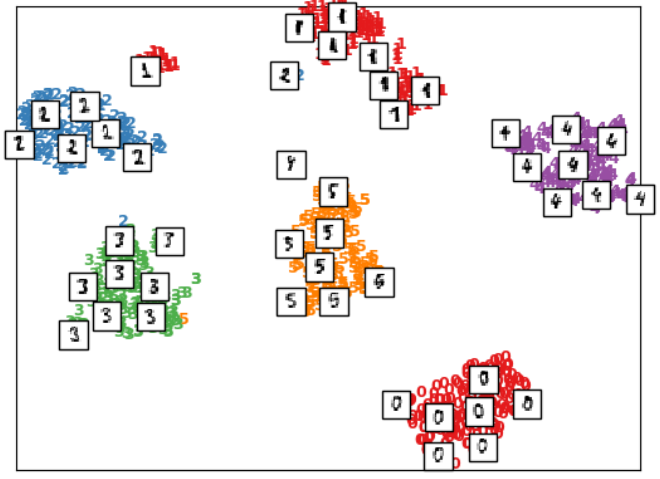


2. Unsupervised Learning: Clustering

A selection from the 64-dimensional digits dataset



t-SNE embedding of the digits (time 5.87s)

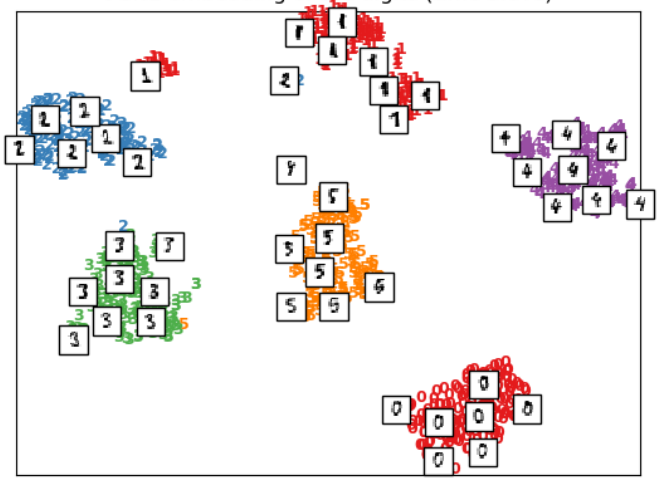


2. Unsupervised Learning: Clustering

A selection from the 64-dimensional digits dataset



t-SNE embedding of the digits (time 5.87s)



$\boldsymbol{x} \in X$: Input data

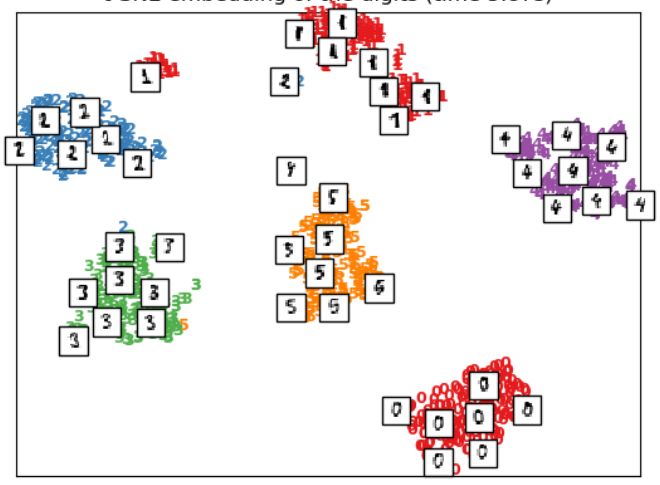
$c_k :\in C$: Output Cluster

2. Unsupervised Learning: Clustering

A selection from the 64-dimensional digits dataset



t-SNE embedding of the digits (time 5.87s)



$\boldsymbol{x} \in X$: Input data

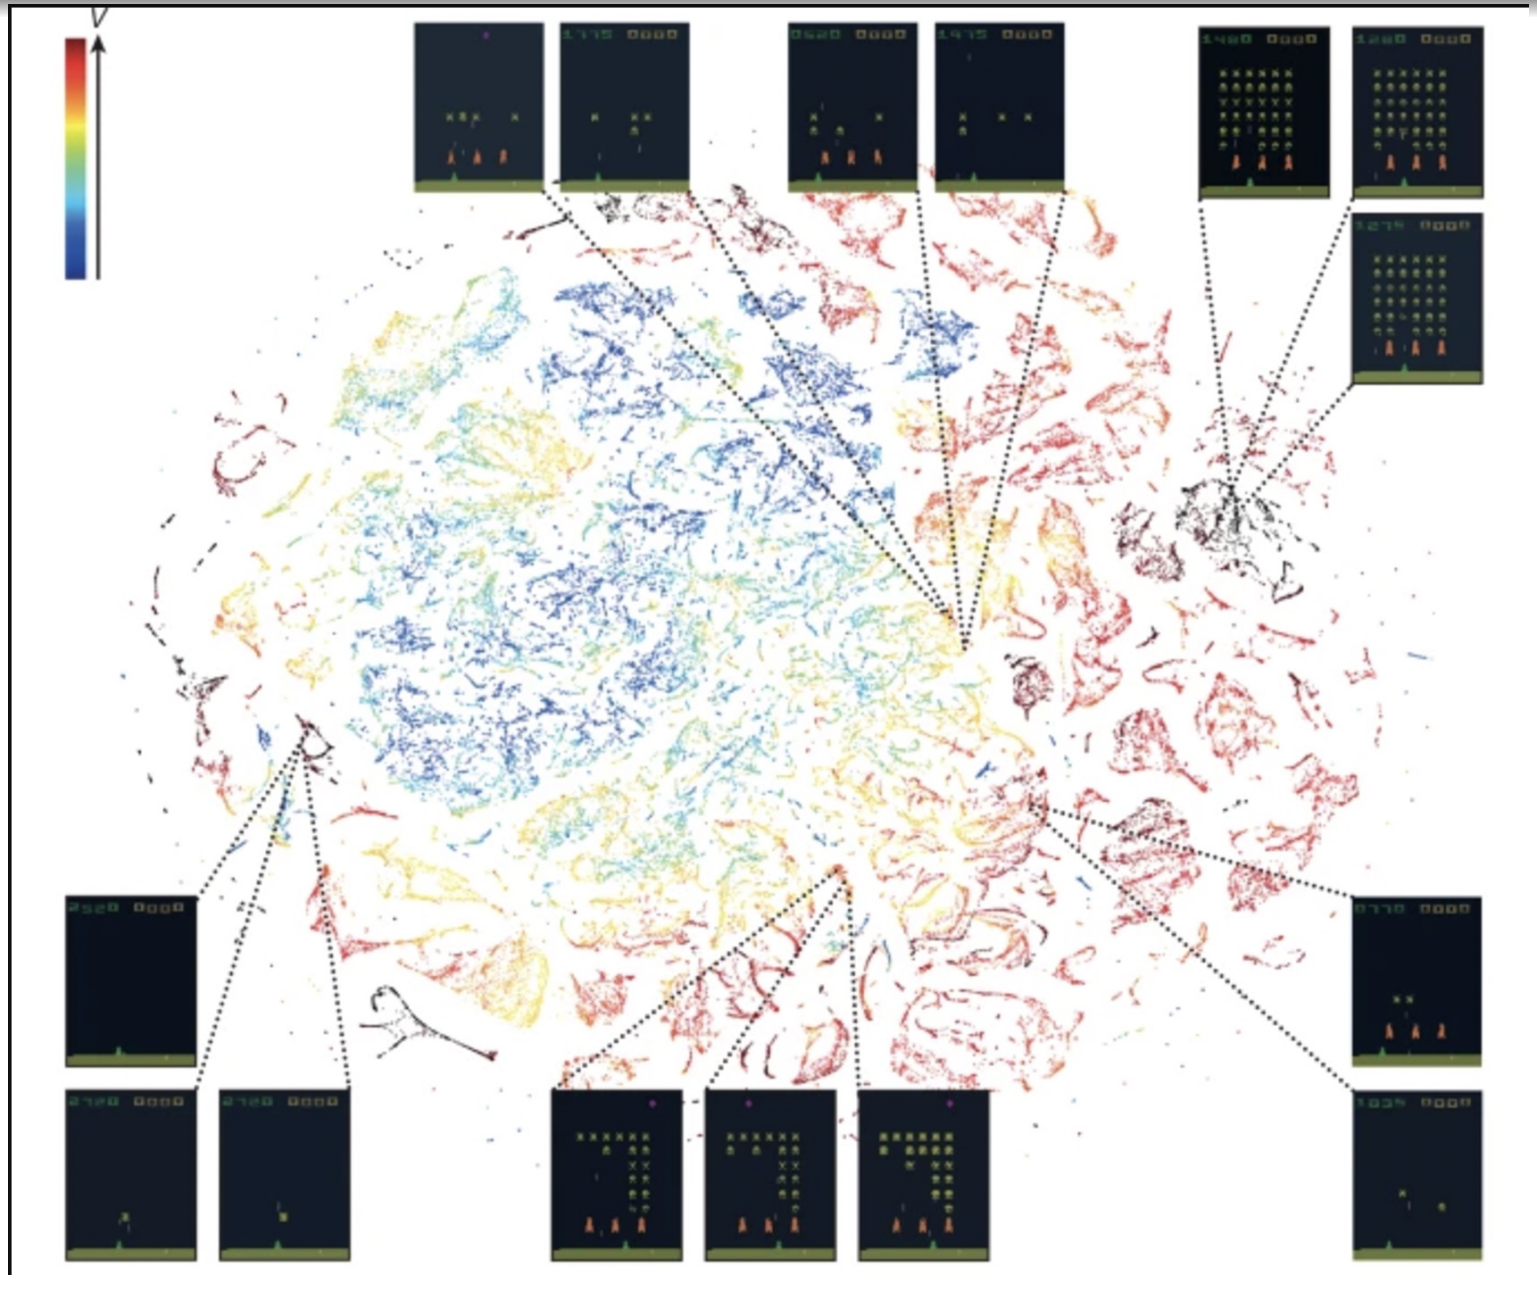
$c_k :\in C$: Output Cluster

CLUSTERING TASK: Find a function

$$f_w : X \to C$$

that maps similar inputs to the same cluster $c_k \in C$.

Atari 2-dimensional t-SNE embedding of game conditions



Definition 6 (3. Reinforcement Learning)

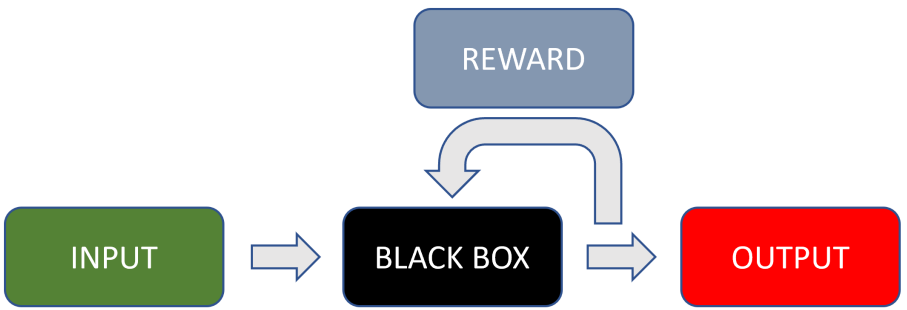
Reinforcement Learning is the task of learning from the consequences of interactions with an environment without being explicitly taught. Examples:

- Robotic Controllers: Several robots controller are trained with reinforcement learning.
- Gaming: The demo clip on the Wall game of ATARI

Definition 6 (3. Reinforcement Learning)

Reinforcement Learning is the task of learning from the consequences of interactions with an environment without being explicitly taught.
Examples:

- Robotic Controllers: Several robots controller are trained with reinforcement learning.
- Gaming: The demo clip on the Wall game of ATARI



- Learns unknown patterns
- Unlabeled input data
- Find hidden patterns

Finding the best weight vector w for $f_w(X)$

Finding the best weight vector w for $f_w(X)$

Given a data set, we divide the data into:

Training Set: it is used to learn a model, and contains usually 80% of the data.

Finding the best weight vector w for $f_w(X)$

Given a data set, we divide the data into:

Training Set: it is used to learn a model, and contains usually 80% of the data.

Test Set: it is used to evaluate the learned model, and it contains usually 20% of the data.

Finding the best weight vector w for $f_w(X)$

Given a data set, we divide the data into:

Training Set: it is used to learn a model, and contains usually 80% of the data.

Test Set: it is used to evaluate the learned model, and it contains usually 20% of the data.

Given a data set, to increase the performance, it is better to fit the model by iteratively changing the Training and Test set, and by recording the model that yields the BEST RESULTS (see n-fold cross validation).

Finding the best weight vector w for $f_w(X)$

Given a data set, we divide the data into:

Training Set: it is used to learn a model, and contains usually 80% of the data.

Test Set: it is used to evaluate the learned model, and it contains usually 20% of the data.

Given a data set, to increase the performance, it is better to fit the model by iteratively changing the Training and Test set, and by recording the model that yields the BEST RESULTS (see n-fold cross validation).

... but **BEST RESULTS** with respect to which METRIC?

Machine Learning Algorithms

SUPERVISED LEARNING ALGORITHMS

CLASSIFICATION	REGRESSION
K-Nearest Neighbors	Linear Regression
Naïve Bayes Classifier	Nonlinear Regression
Support Vector Machines	Ordinary Least Square
Neural Networks	Neural Networks
Decision Trees	
Random Forests	

Machine Learning Algorithms

SUPERVISED LEARNING ALGORITHMS

CLASSIFICATION	REGRESSION
K-Nearest Neighbors	Linear Regression
Naïve Bayes Classifier	Nonlinear Regression
Support Vector Machines	Ordinary Least Square
Neural Networks	Neural Networks
Decision Trees	
Random Forests	

Machine Learning Algorithms

UNSUPERVISED LEARNING ALGORITHMS

CLUSTER ANALYSIS

- K-means clustering
- Hierarchical clustering

DIMENSION REDUCTION

- Linear Discriminant Analysis
- Principal Component Analysis
- *Manifold Learning*
- *Autoencoders*

Outline

1 Introduction

2 ML tasks

3 Neural Networks

4 Generative AI

5 AI4Pharma

(Deep) Neural Networks



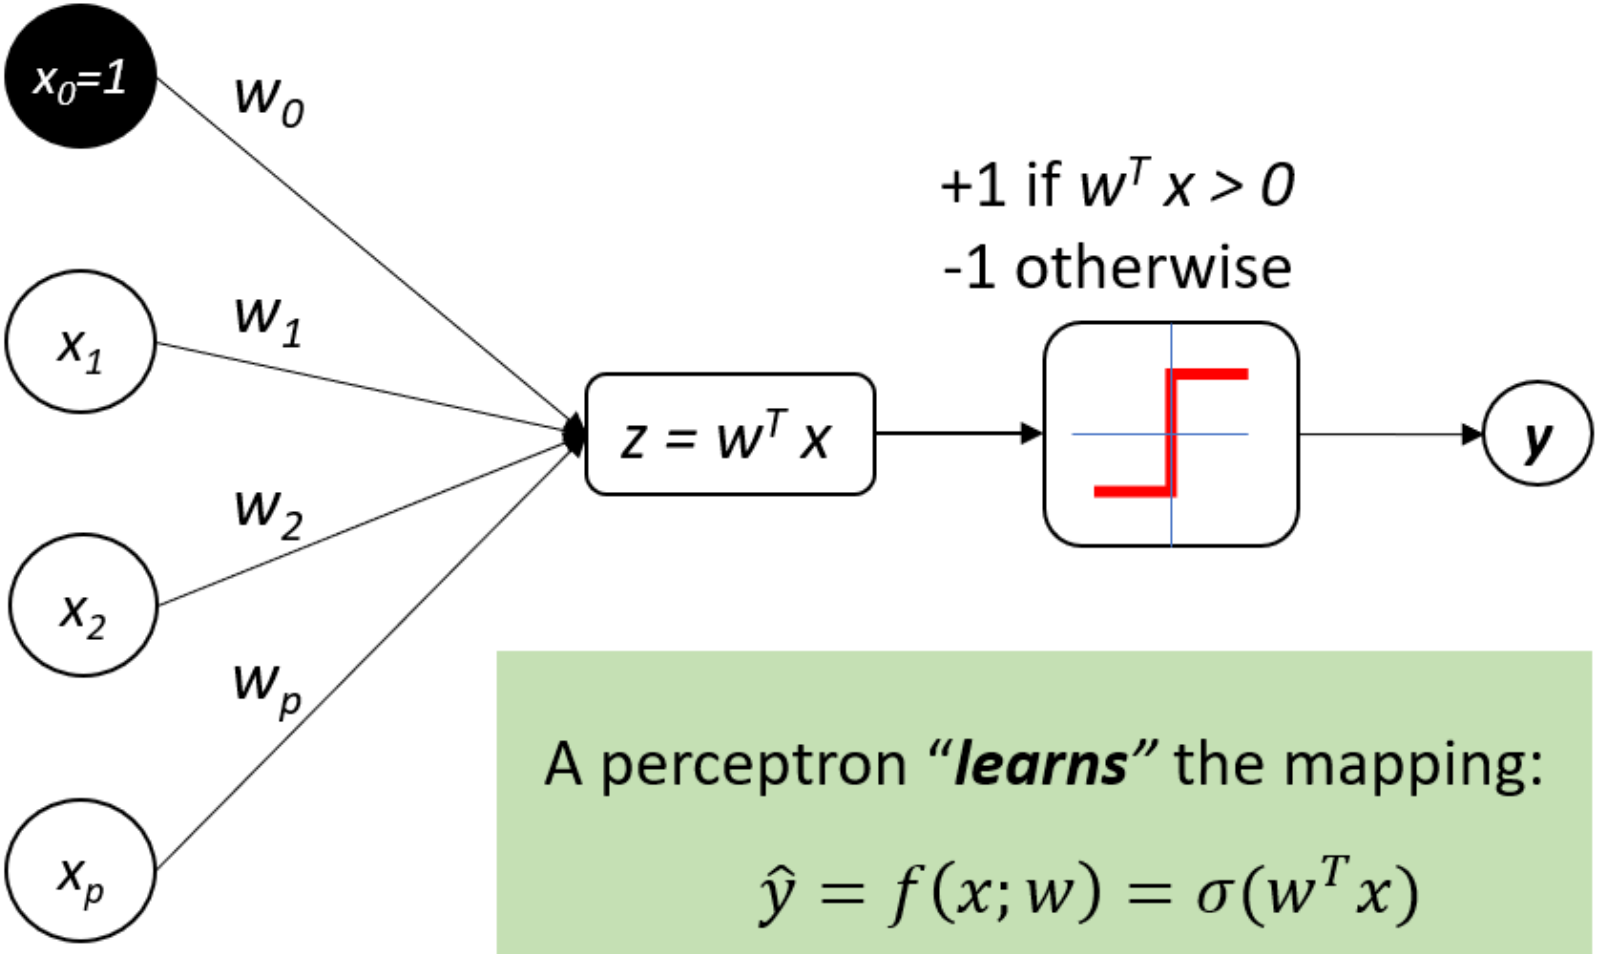
Biological Motivation

1. The human brain is estimated to contain a densely interconnected network of around 10^{11} neurons
2. Each neuron is connected to around others 10^4 neurons. Neurons activity is typically excited or inhibited through connections to other neurons.
3. The switching of neurons is of order of 10^{-3} seconds, much slower than modern computers, which have switching speeds of 10^{-10} seconds
4. Humans take around 10^{-1} seconds to recognize a friend (compare to switching speed...)
5. Human brain performs a highly parallel tasks

Let us start simple: Single Perceptron



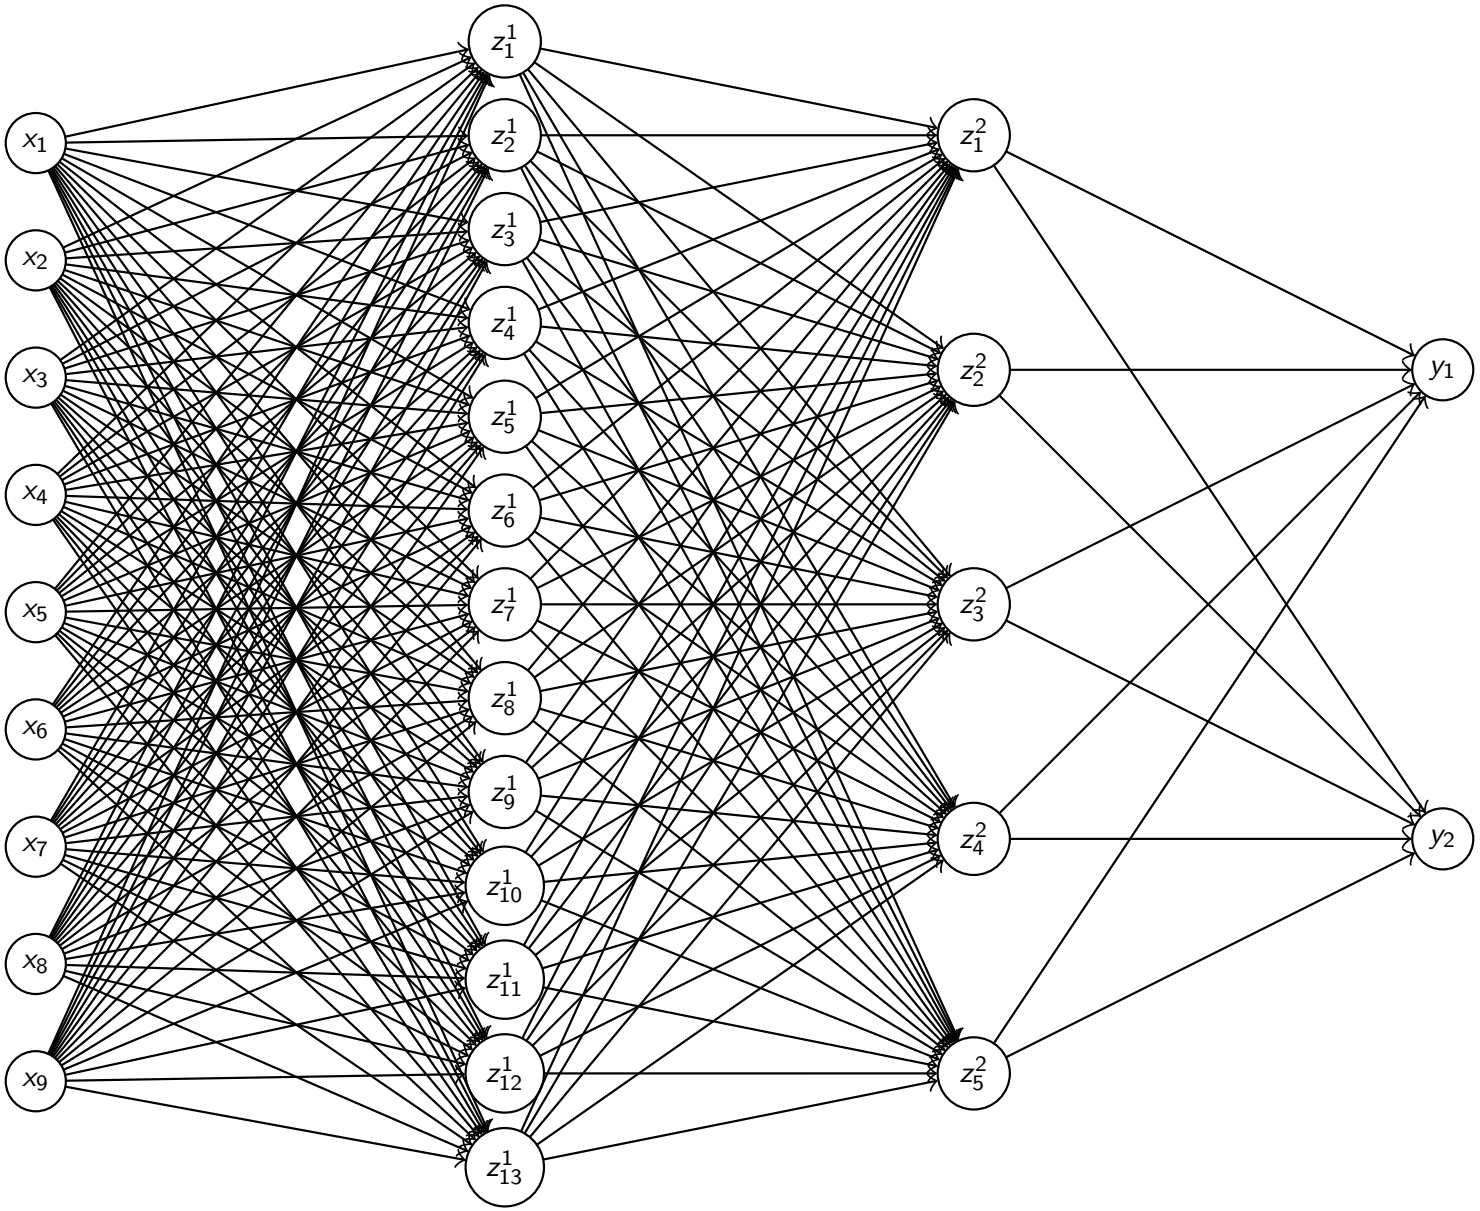
Perceptron



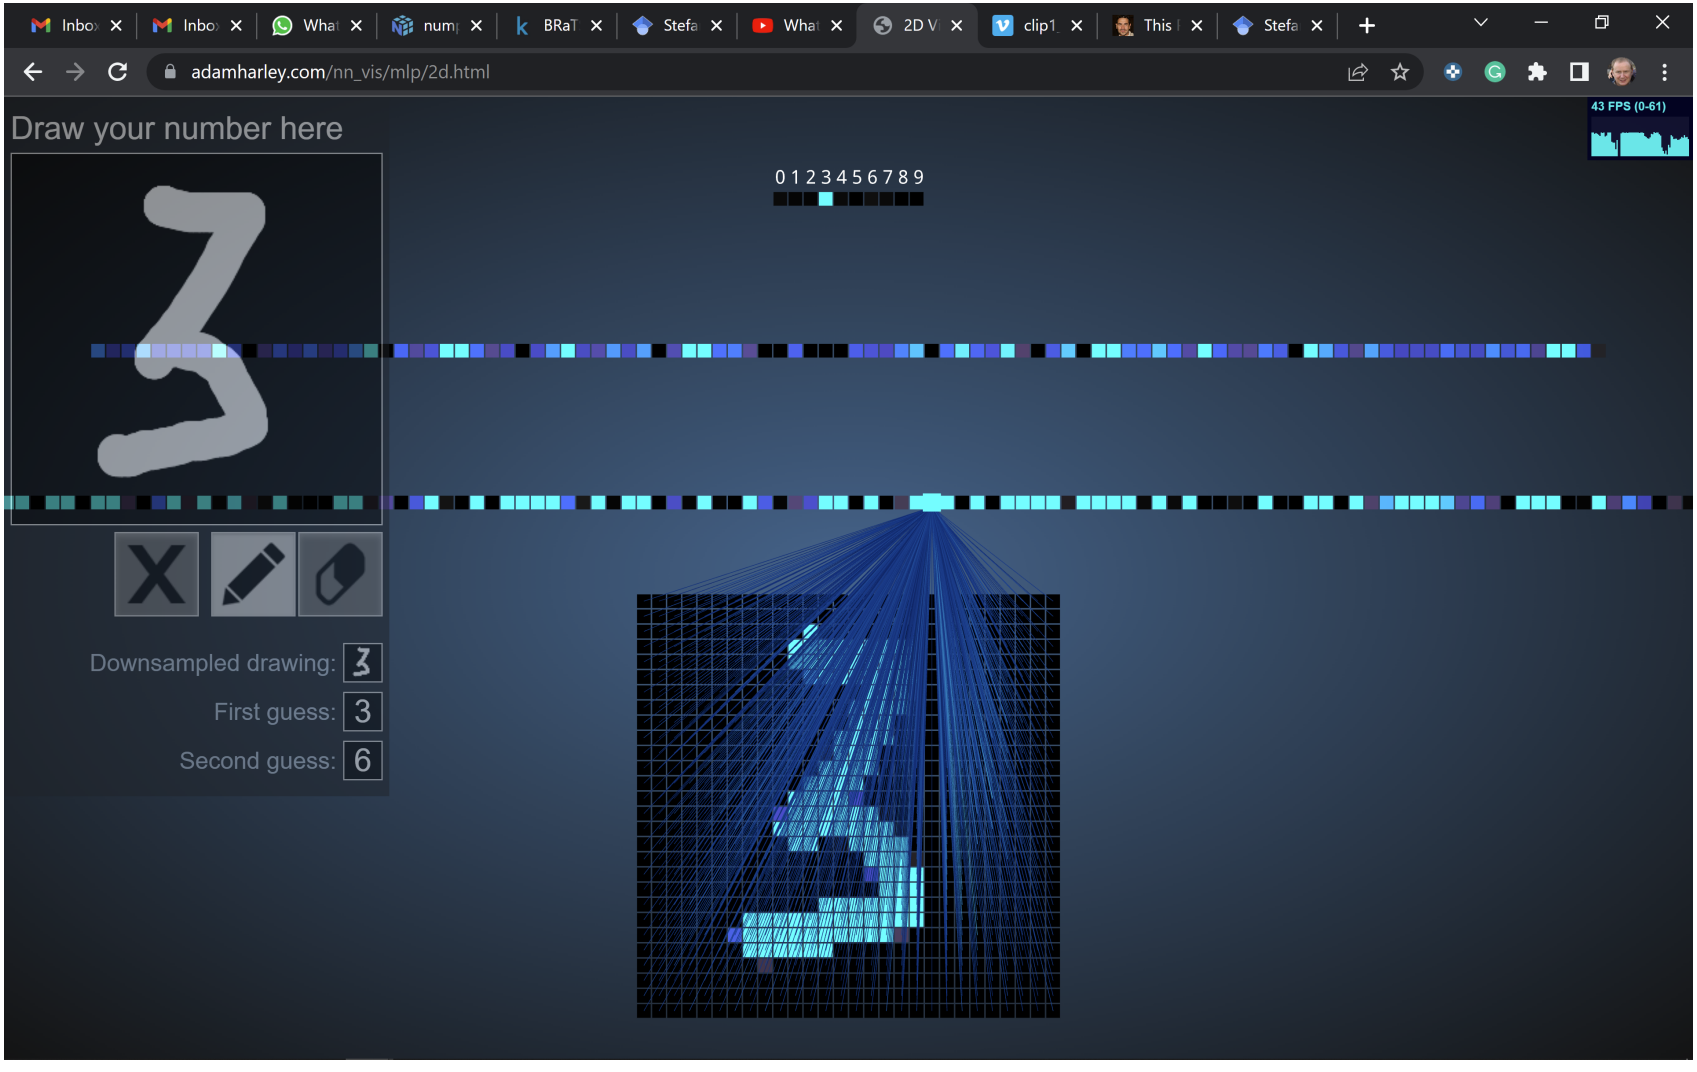
A perceptron "***learns***" the mapping:

$$\hat{y} = f(x; w) = \sigma(w^T x)$$

Multilayer Feedforward Neural Network



Multilayer Feedforward Neural Networks



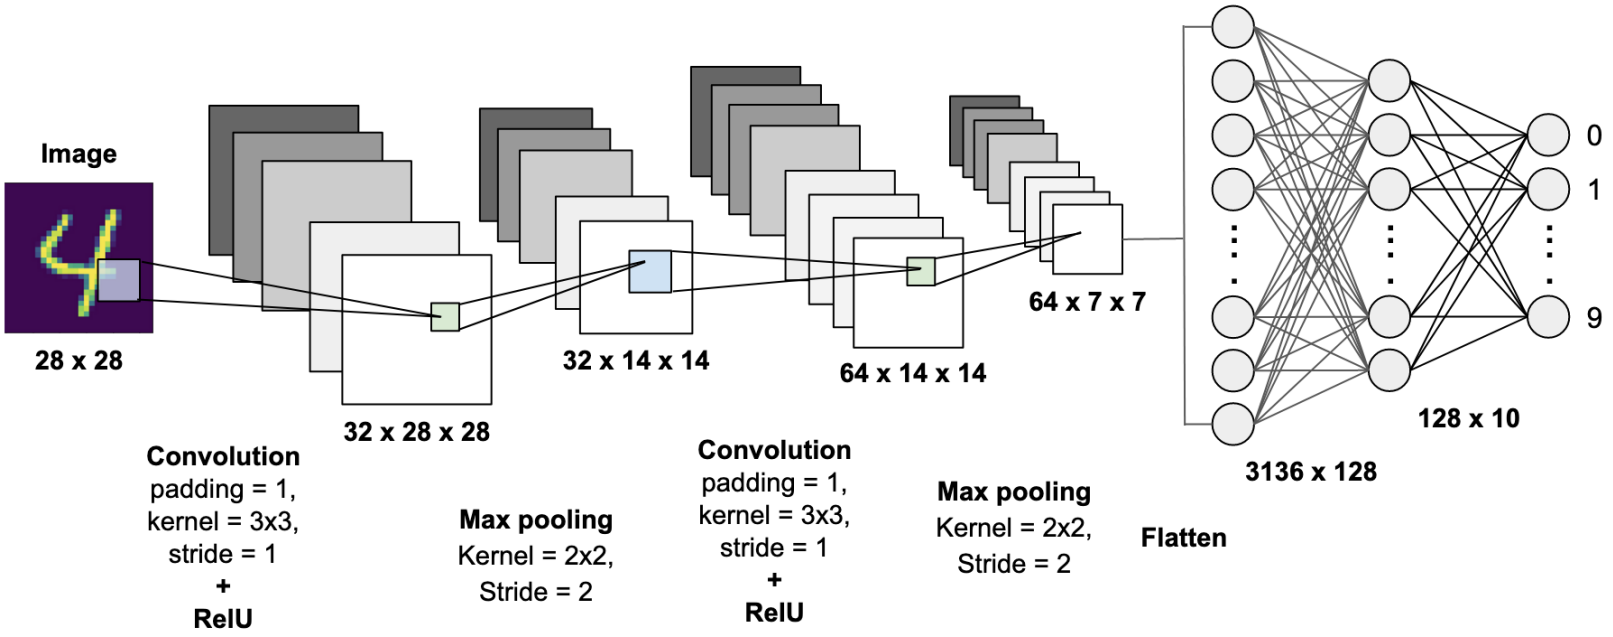
https://adamharley.com/nn_vis/

Backpropagation as Automatic Differentiation

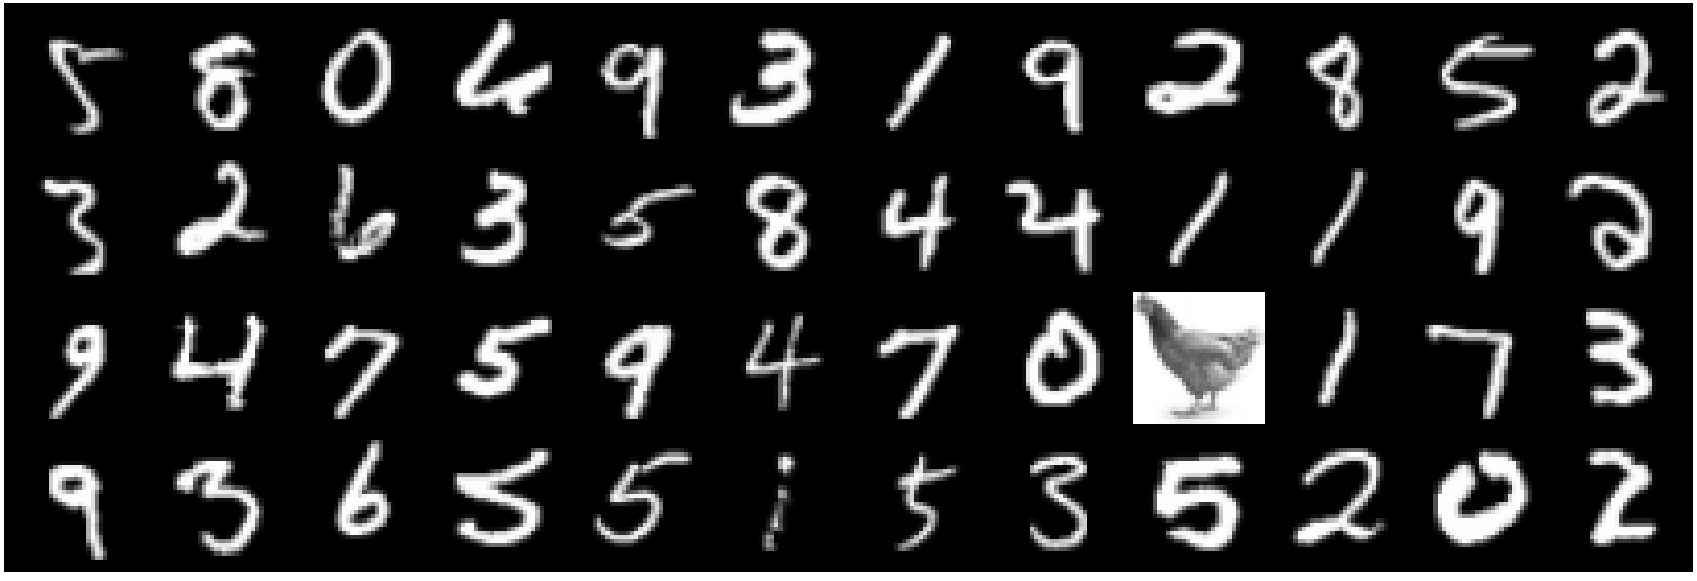
DEMO ON LEARNING WEIGHTS: Backpropagation Algorithms

https://www.youtube.com/watch?v=Ilg3gGewQ5U

Modern Convolutional Neural Network



WARNING: Limits and risks 1/3



WARNING: Limits and risks 2/3



76.4%
confident
that it's a 5



99.9%
confident
that it's a 5

WARNING: Limits and risks 3/3

Original image

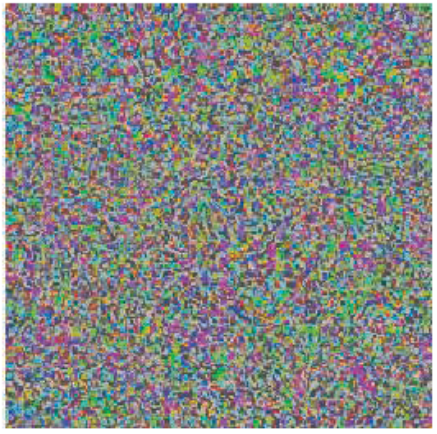


Dermatoscopic image of a benign melanocytic nevus, along with the diagnostic probability computed by a deep neural network.



+ 0.04 ×

Adversarial noise



Perturbation computed by a common adversarial attack technique. See (7) for details.

=

Adversarial example



Combined image of nevus and attack perturbation and the diagnostic probabilities from the same deep neural network.



GDPR 2018: The right to Explanations

The **data subject** *should have the right* **NOT TO BE** *subject to a decision, which may include a measure, evaluating personal aspects relating to him or her which is based solely on automated processing*

[...]

GDPR 2018: The right to Explanations

The **data subject** *should have the right* **NOT TO BE** *subject to a decision, which may include a measure, evaluating personal aspects relating to him or her which is based solely on automated processing*

[...]

(the data subject has) the right to obtain human intervention, to express his or her point of view, **to obtain an explanation of the decision reached after such assessment and to challenge the decision**.

Outline

1 Introduction

2 ML tasks

3 Neural Networks

4 Generative AI

5 AI4Pharma

Which is a fake photo?





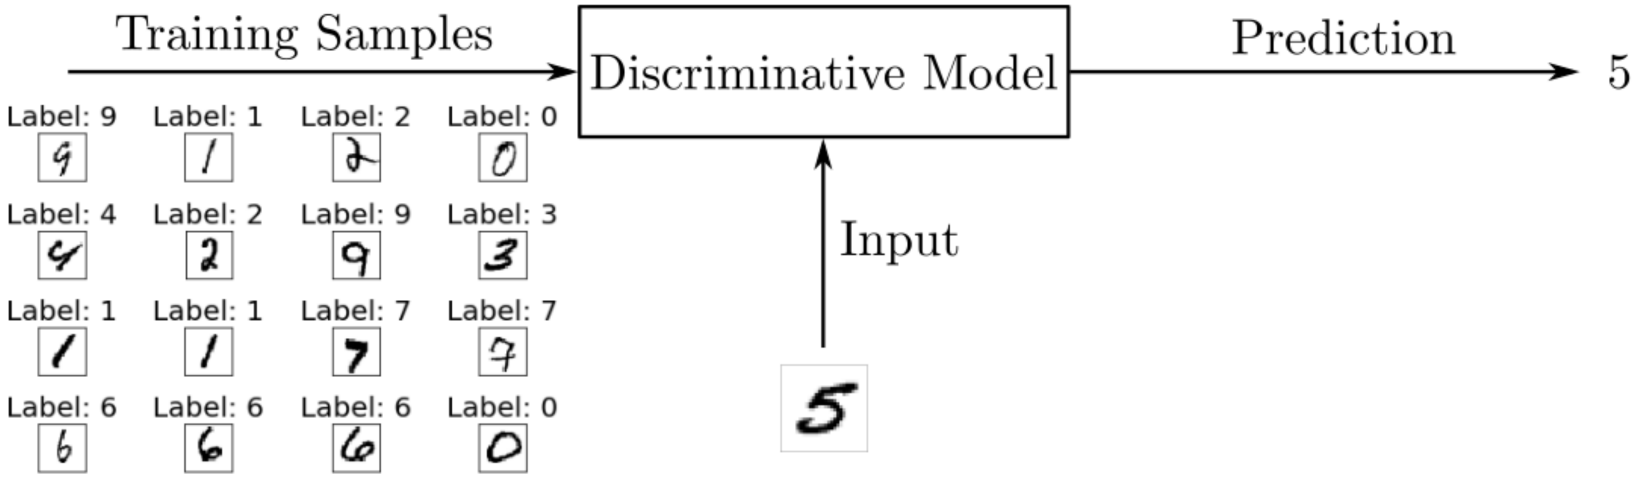
Which is a fake photo?





https://thispersondoesnotexist.com/

Using Deep NNs as Discriminative Models



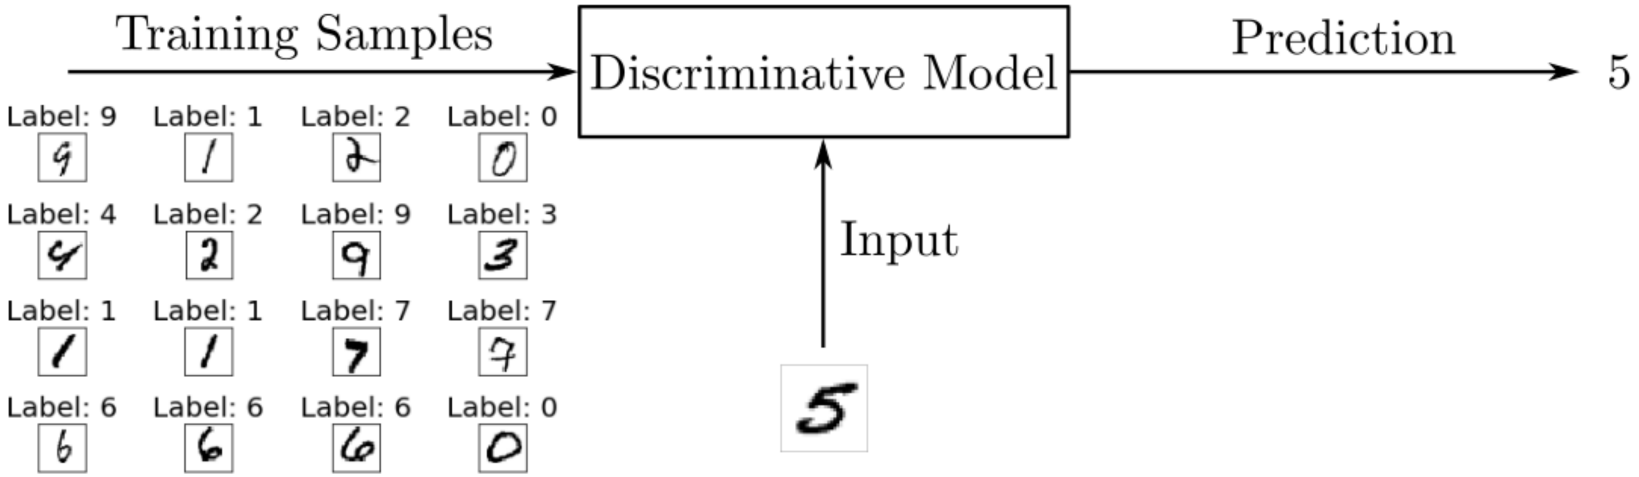
Step 1: **Train-test split of the dataset**.

Step 2: **Choice of architecture**.

Step 3: **Training**.

Step 4: **Testing**.

Using Deep NNs as Discriminative Models



Step 1: **Train-test split of the dataset**.

Step 2: **Choice of architecture**.

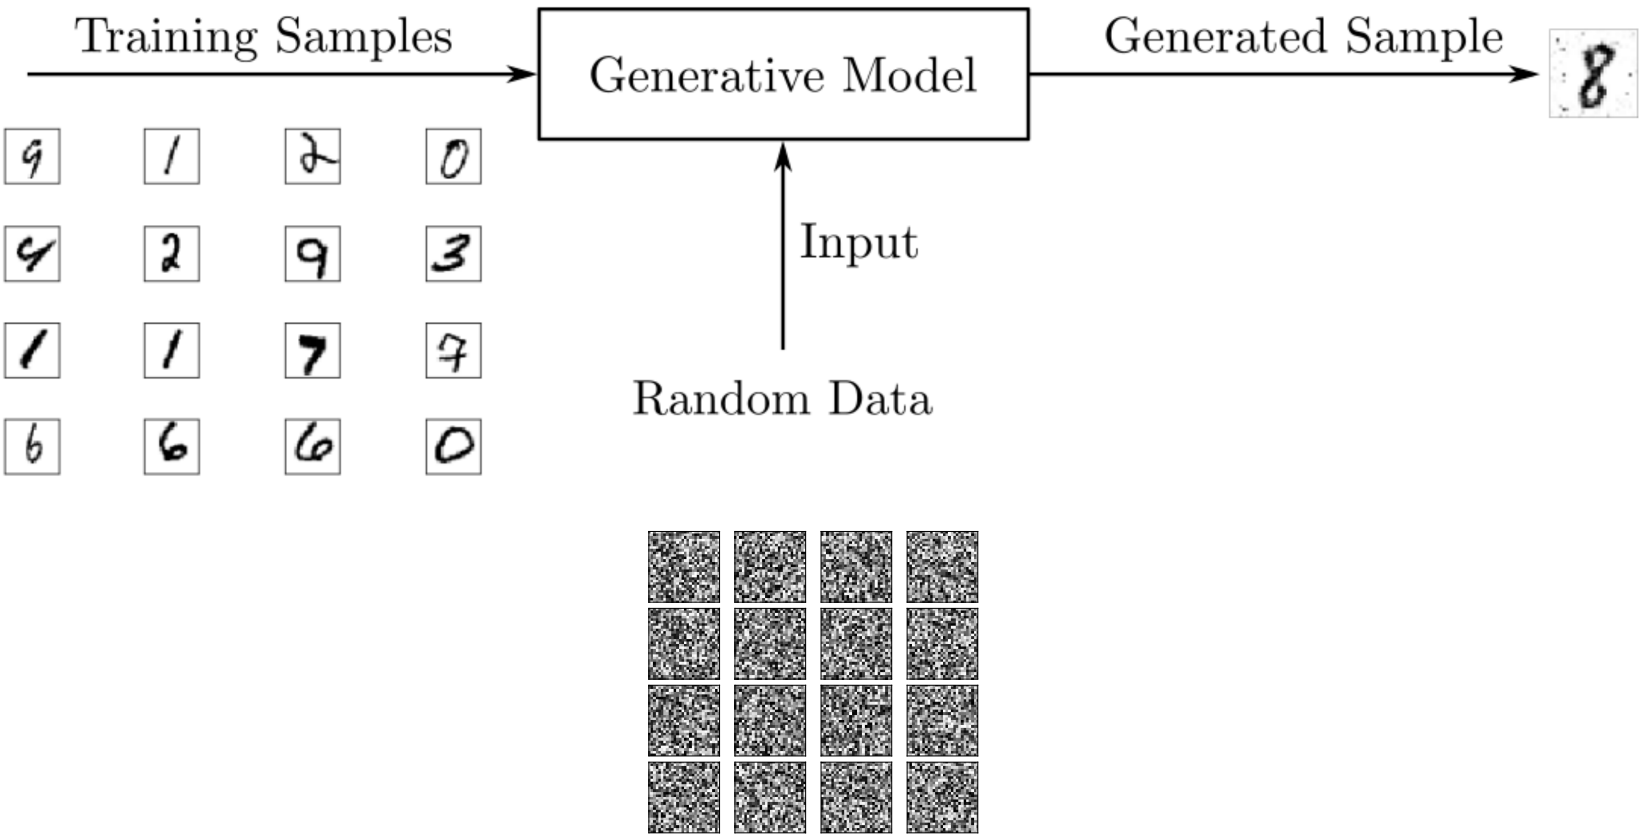
Step 3: **Training**.

Step 4: **Testing**.

This is an example of **Supervised Learning**

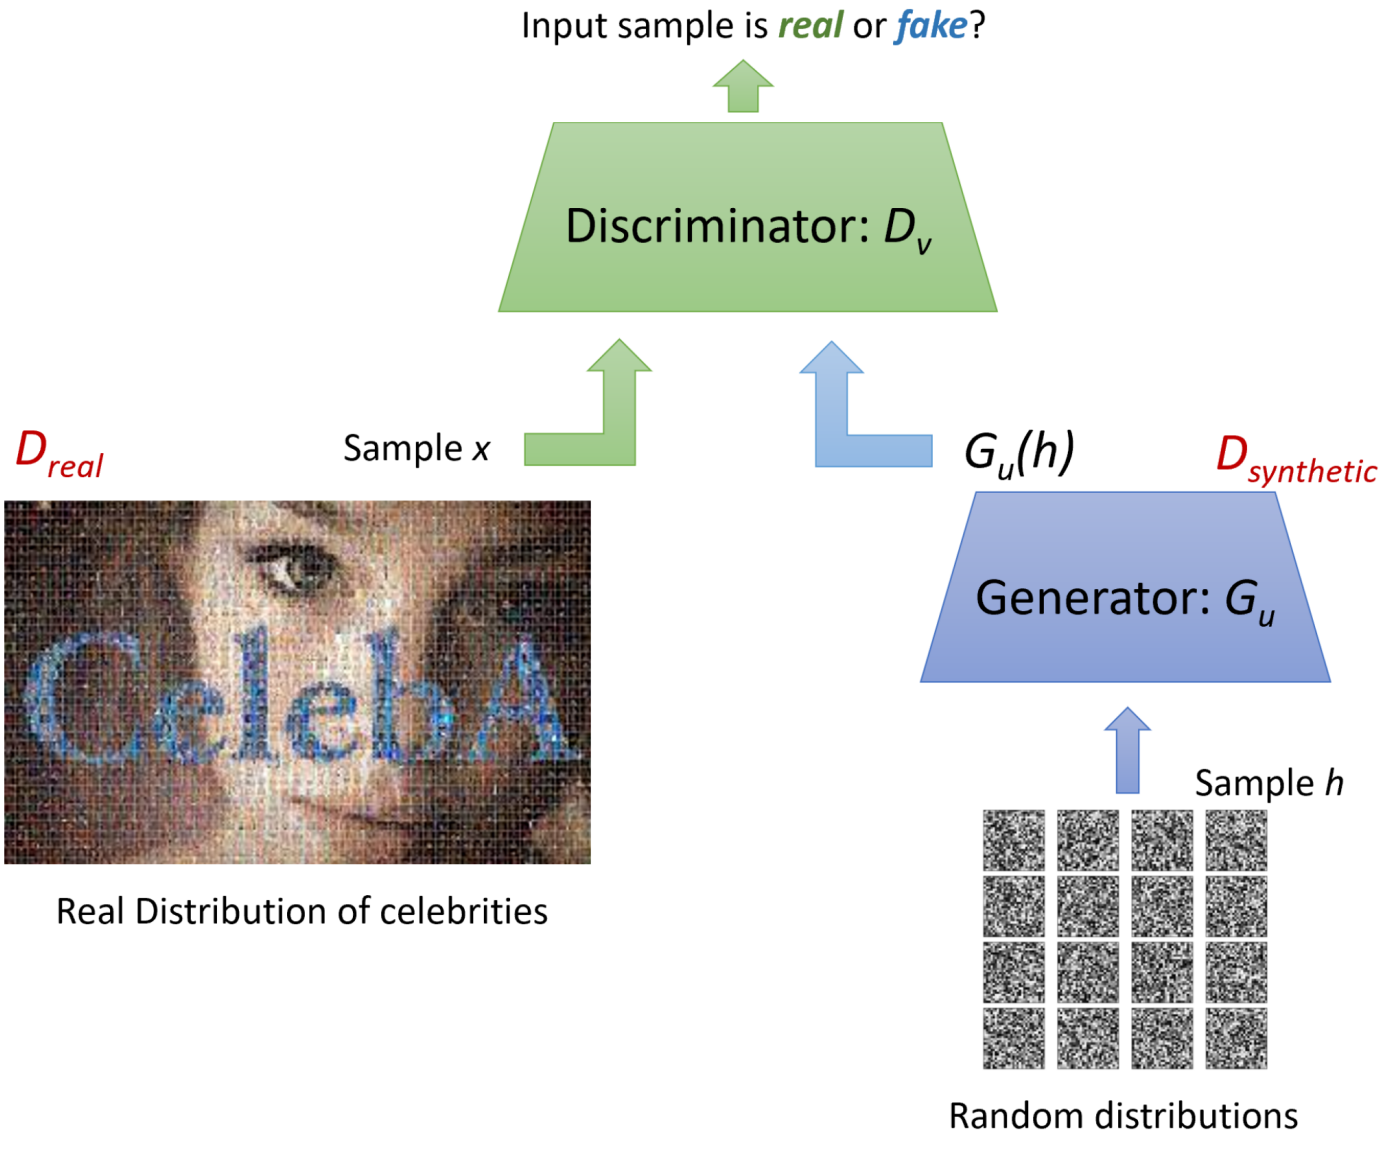
Using Deep NNs as Generative Models

This is an example of **Unsupervised Learning**

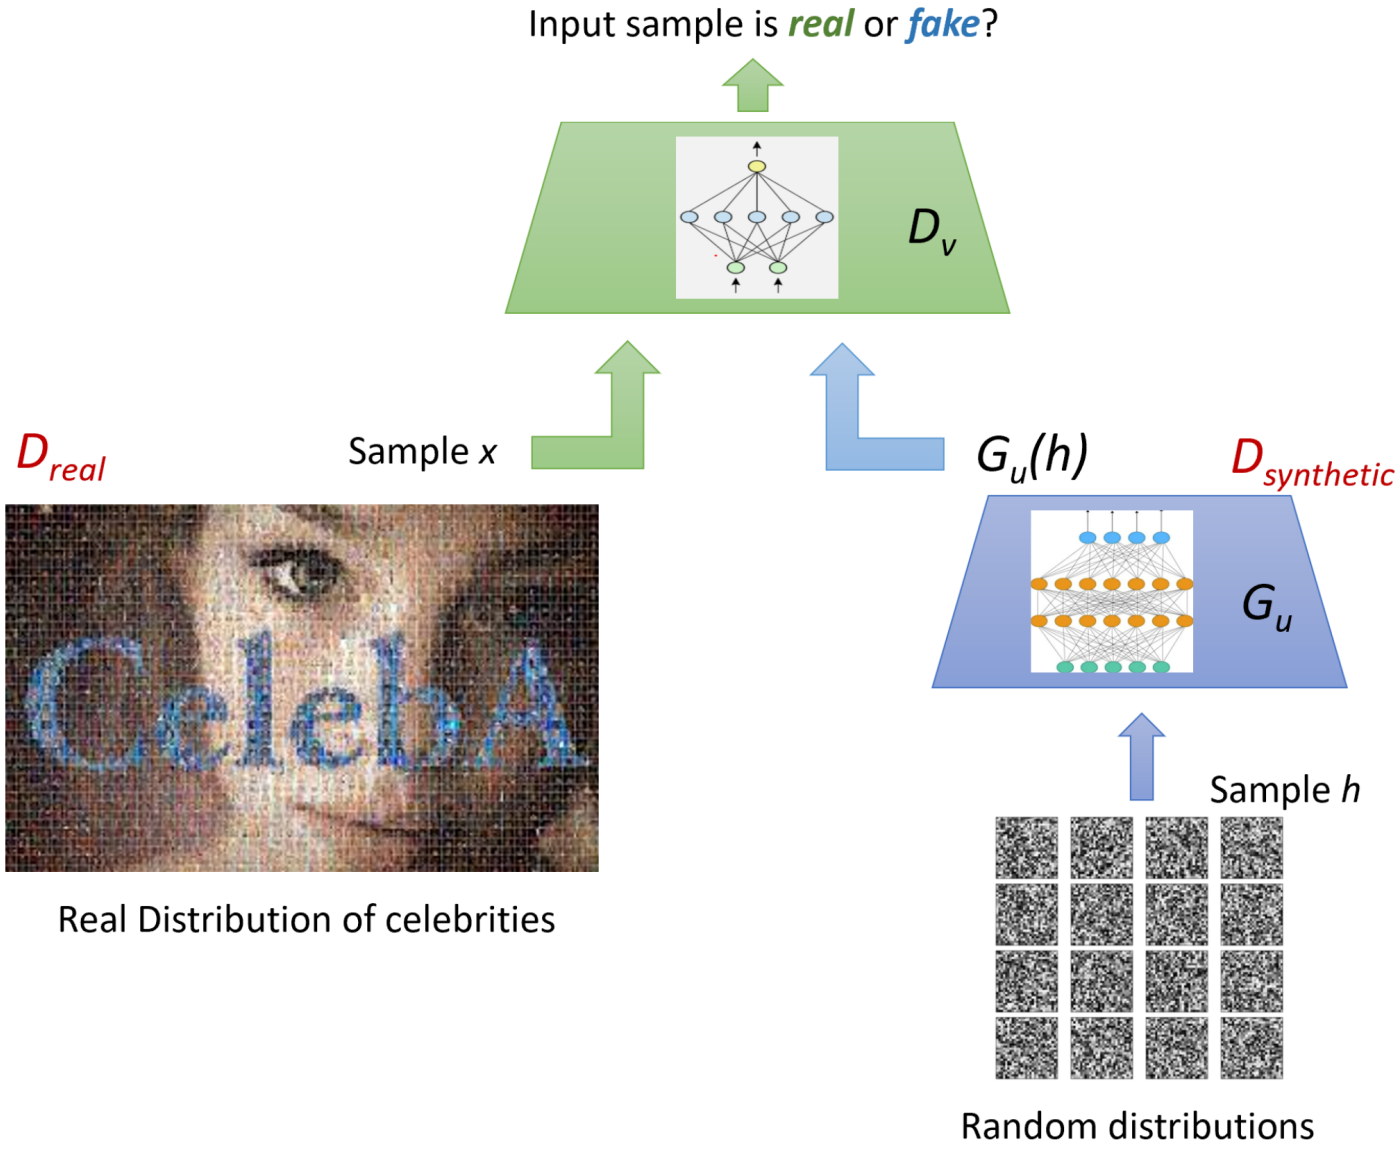


Picture from: `https://realpython.com/generative-adversarial-networks/`

The GAN game


Input sample is real or fake?
Discriminator: D_v
D_{real}
Sample x
$G_u(h)$
$D_{synthetic}$
Generator: G_u
Sample h
Real Distribution of celebrities
Random distributions

The GAN game


Input sample is real or fake?
D_v
D_{real}
Sample x
$G_u(h)$
$D_{synthetic}$
G_u
Sample h
Real Distribution of celebrities
Random distributions

Outline

1 Introduction

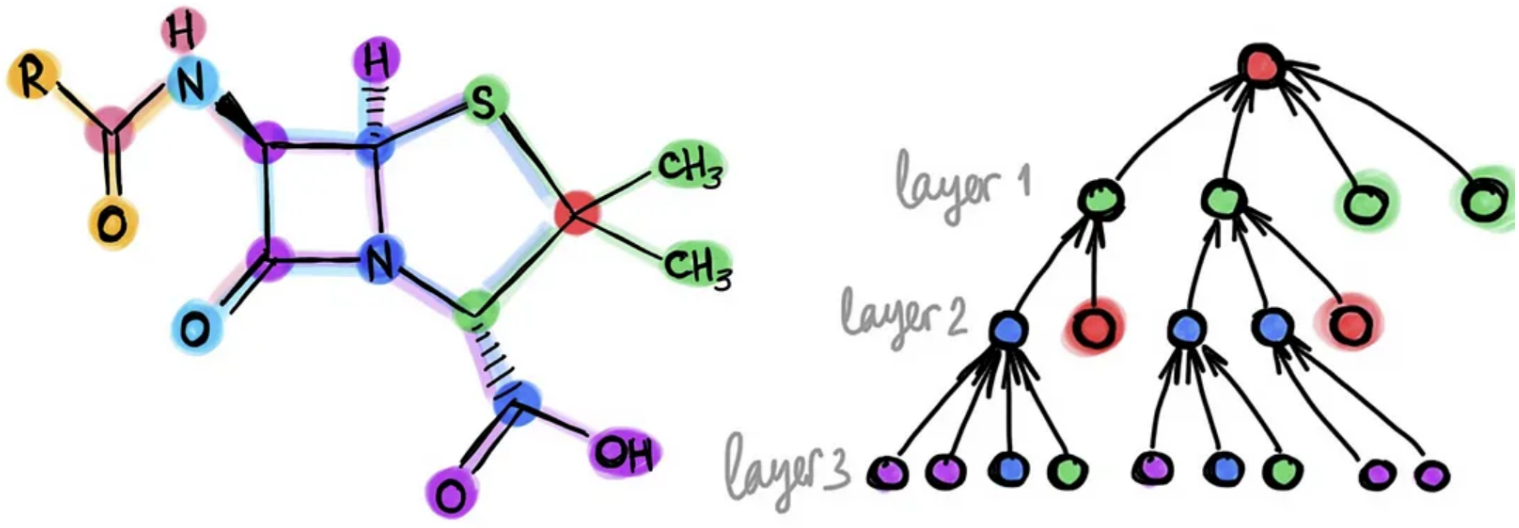
2 ML tasks

3 Neural Networks

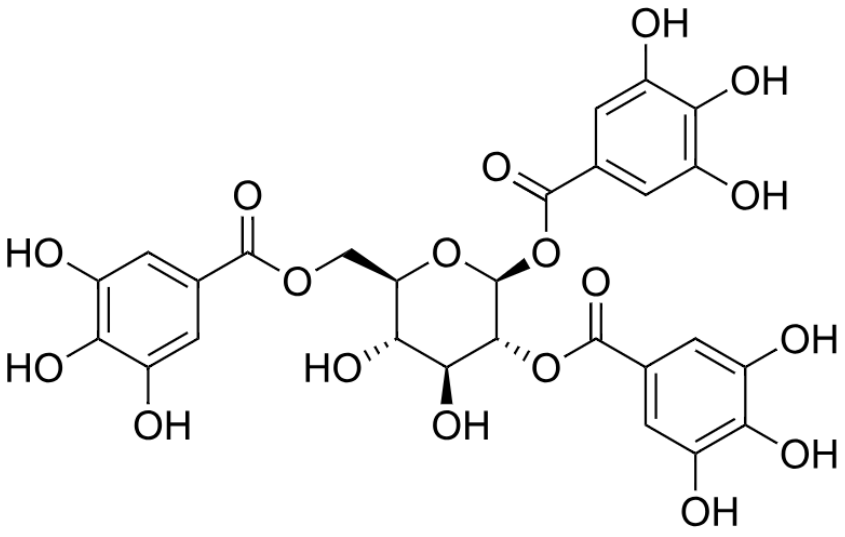
4 Generative AI

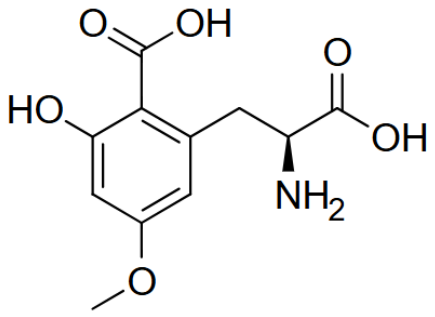
5 AI4Pharma

Molecules are graphs (and graphs are sparse matrices...)



Classification of molecules





Graph-level classification task: to predict whether a given chemical molecule is **non-toxic** (on the left) or **toxic** (on the right). This is an example of binary classification.

Dataset of interest for the pharmaceutical industry

Name	Short description
Mutag	4337 molecular structures with corresponding Ames tests
ogbg-molhiv	Graphs represent molecules (nodes atoms, edges bonds)
ogbg-molpcba	Graphs represent molecules (nodes atoms, edges bonds)
QM7	Chemical dataset of molecules with 7 heavy atoms
QM9	Chemical dataset of molecules with 9 heavy atoms
ZINC250k	List of drug-like compounds available for sale
ogbn-proteins	Graphs of protein interactions (e.g., co-expression)
Protein	Graphs of proteins to predict if the graph is an enzyme
Enzymes	Graphs of proteins to assign enzymes to classes
ogbl-ddi	Graphs of drug-drug interactions
Open Graph Bench.	Realistic, large-scale, and diverse datasets for ML on graphs
TUDatasets	Datasets for graph classification and regression
Graph Classification	Collection of graph datasets for classification tasks

History of AlphaFold (See Giorgio's talk this afternoon)

1. 2015: Atari's Deep Learning Network (first video)

History of AlphaFold (See Giorgio's talk this afternoon)

1. 2015: Atari's Deep Learning Network (first video)
2. 2016: AlphaGo beats the world champion in the game of GO.

History of AlphaFold (See Giorgio's talk this afternoon)

1. 2015: Atari's Deep Learning Network (first video)
2. 2016: AlphaGo beats the world champion in the game of GO.
3. 2018: AlphaZero learn the GO game from scratch playing against AlphaGo.

History of AlphaFold (See Giorgio's talk this afternoon)

1. 2015: Atari's Deep Learning Network (first video)
2. 2016: AlphaGo beats the world champion in the game of GO.
3. 2018: AlphaZero learn the GO game from scratch playing against AlphaGo.
4. 2018: AlphaFold wins the Critical Assessment of Structure Prediction (CASP) competion

History of AlphaFold (See Giorgio's talk this afternoon)

1. 2015: Atari's Deep Learning Network (first video)
2. 2016: AlphaGo beats the world champion in the game of GO.
3. 2018: AlphaZero learn the GO game from scratch playing against AlphaGo.
4. 2018: AlphaFold wins the Critical Assessment of Structure Prediction (CASP) competion
5. 2020: AlphaFold2 achieves a 92% accuracy in predicting the 3D structure of proteins.

History of AlphaFold (See Giorgio's talk this afternoon)

1. 2015: Atari's Deep Learning Network (first video)
2. 2016: AlphaGo beats the world champion in the game of GO.
3. 2018: AlphaZero learn the GO game from scratch playing against AlphaGo.
4. 2018: AlphaFold wins the Critical Assessment of Structure Prediction (CASP) competion
5. 2020: AlphaFold2 achieves a 92% accuracy in predicting the 3D structure of proteins.
6. 2024: AlphaFold3 can predict the structure of complexes created by proteins with DNA, RNA, various ligands, and ions

History of AlphaFold (See Giorgio's talk this afternoon)

1. 2015: Atari's Deep Learning Network (first video)
2. 2016: AlphaGo beats the world champion in the game of GO.
3. 2018: AlphaZero learn the GO game from scratch playing against AlphaGo.
4. 2018: AlphaFold wins the Critical Assessment of Structure Prediction (CASP) competion
5. 2020: AlphaFold2 achieves a 92% accuracy in predicting the 3D structure of proteins.
6. 2024: AlphaFold3 can predict the structure of complexes created by proteins with DNA, RNA, various ligands, and ions

`https://deepmind.google/technologies/alphafold/`

AlphaFold2 and the Nobel Prize

The Nobel Prize in Chemistry 2024

Summary

Laureates

David Baker

Demis Hassabis

John Jumper

Prize announcement

Press release

Popular information

Advanced information

Share this

David Baker
University of Washington, Seattle, WA, USA
Howard Hughes Medical Institute, USA

"for computational protein design"

and the other half jointly to

Demis Hassabis
Google DeepMind, London, UK

John Jumper
Google DeepMind, London, UK

"for protein structure prediction"

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.

Lettura consigliata divulgativa



Terence Tao's perspective on AI and teaching

mastodon

Back



Terence Tao
@tao@mathstodon.xyz

AI tools like #ChatGPT will soon be capable of answering a large fraction of traditional university homework type questions with reasonable accuracy. In the long term, it seems futile to fight against this; perhaps what we as lecturers need to do is to move to an "open books, open AI" mode of examination where we give the students full access to AI tools but ask them more challenging questions, both to teach the material and also to teach the students how best to use the AI tools of the future.

20 Dec 2022, 01:21 · · Web · 185 · 282