

The background features a dark silhouette of a human figure on the left and a robot figure on the right. The human figure is shown in profile, facing right, with its right hand extended towards the robot. The robot figure is also in profile, facing left, with its right hand reaching out to meet the human's. The robot has a boxy head with a single antenna on top and a cylindrical body. The overall scene is set against a dark, gradient background.

# Artificial Intelligence

Eric Postma

[eric.postma@gmail.com](mailto:eric.postma@gmail.com)

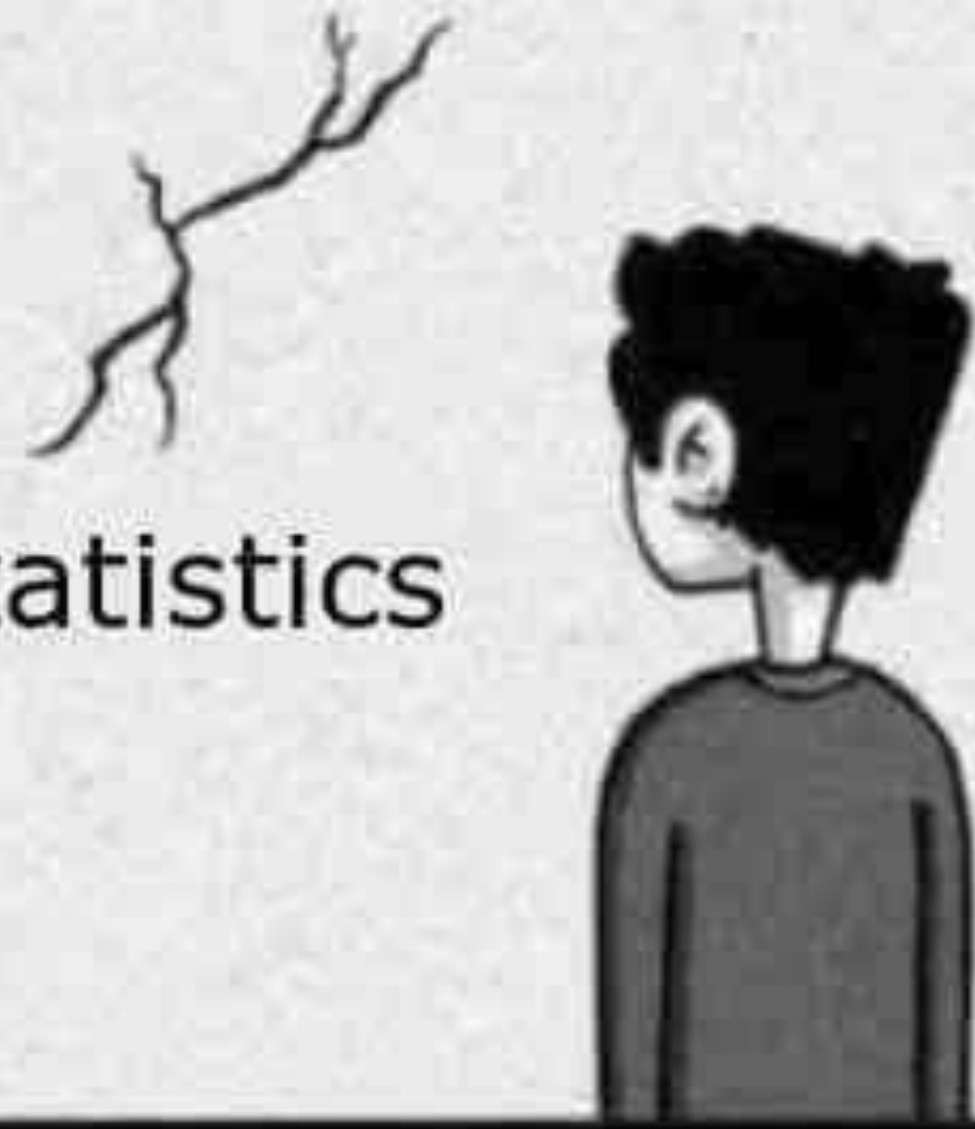
[www.ericpostma.nl](http://www.ericpostma.nl)

CSAI/MindLabs, Tilburg University

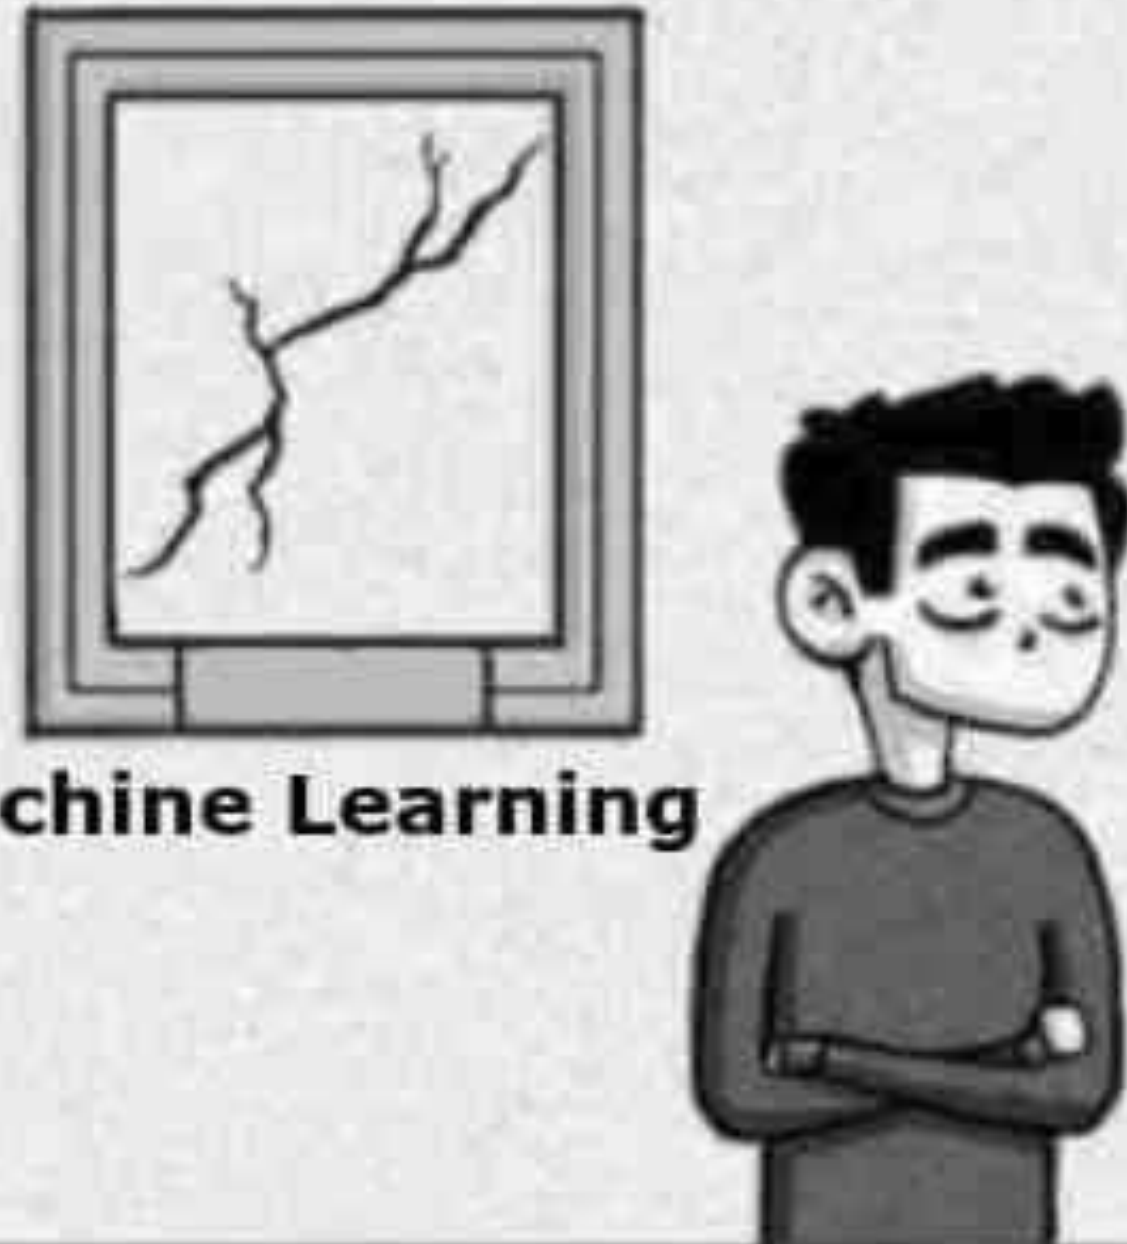
JADS, 's-Hertogenbosch

Natuurkundig Genootschap Wessel Knoops, Arnhem 10 Maart 2025

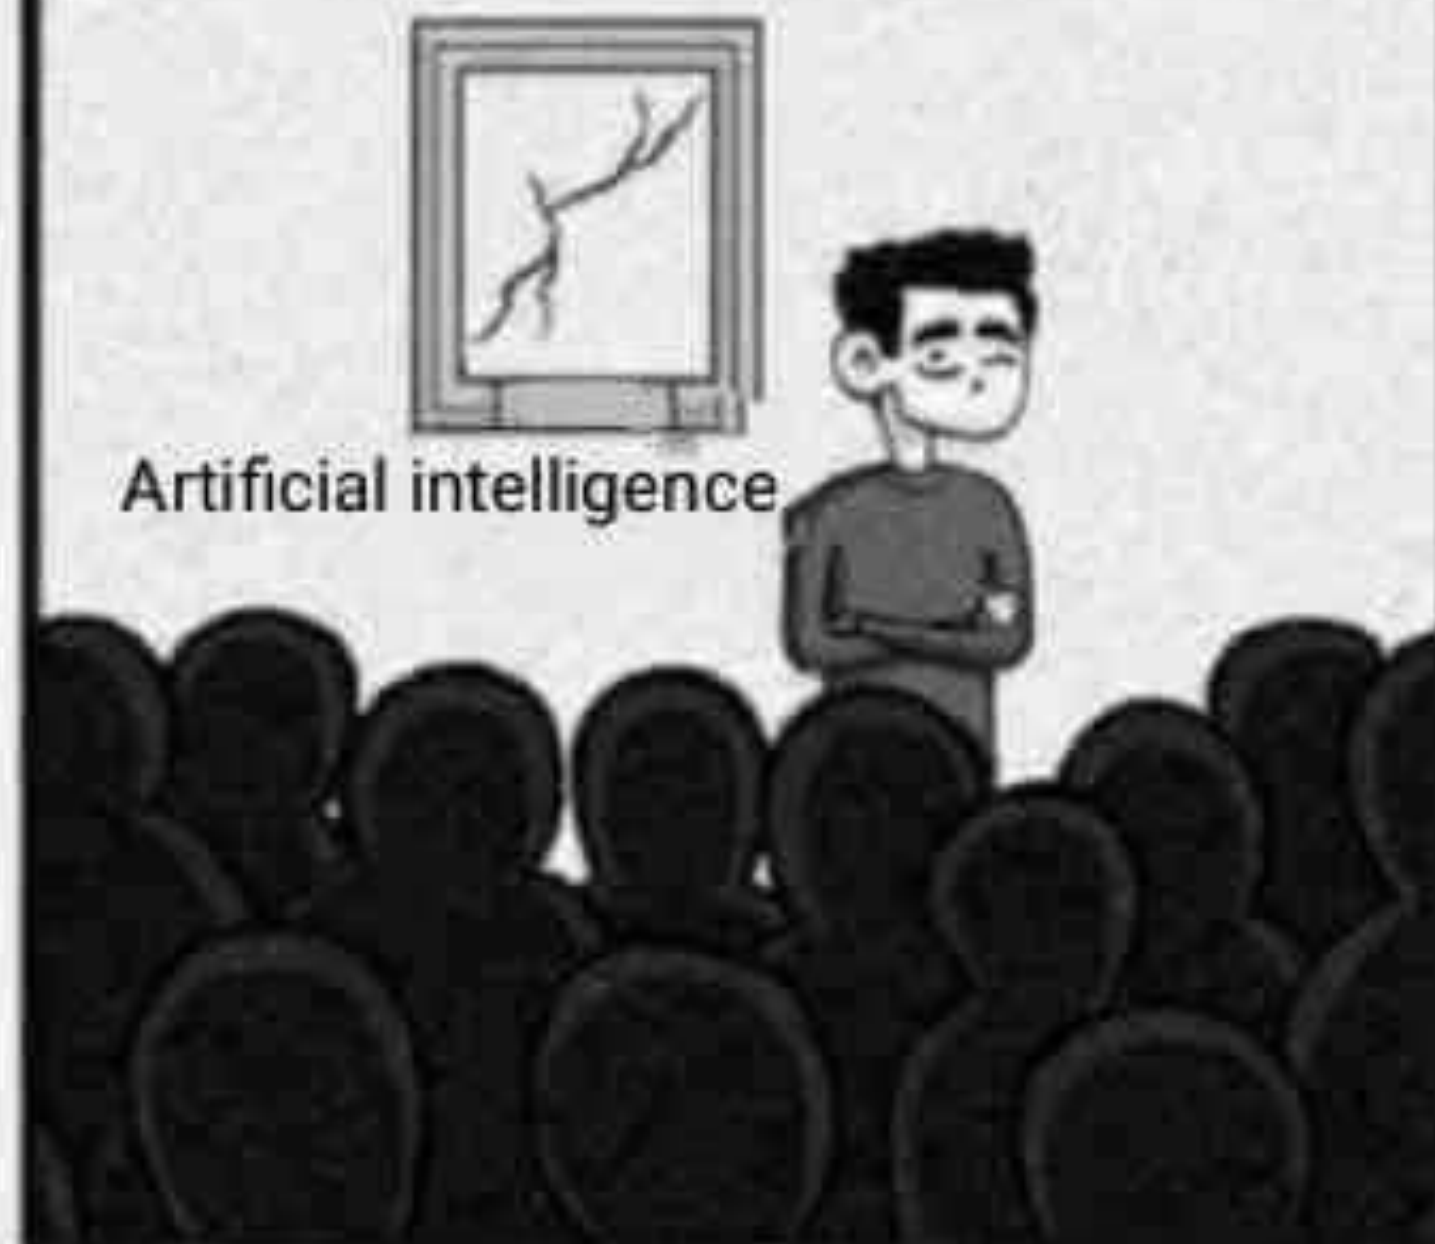
statistics



Machine Learning

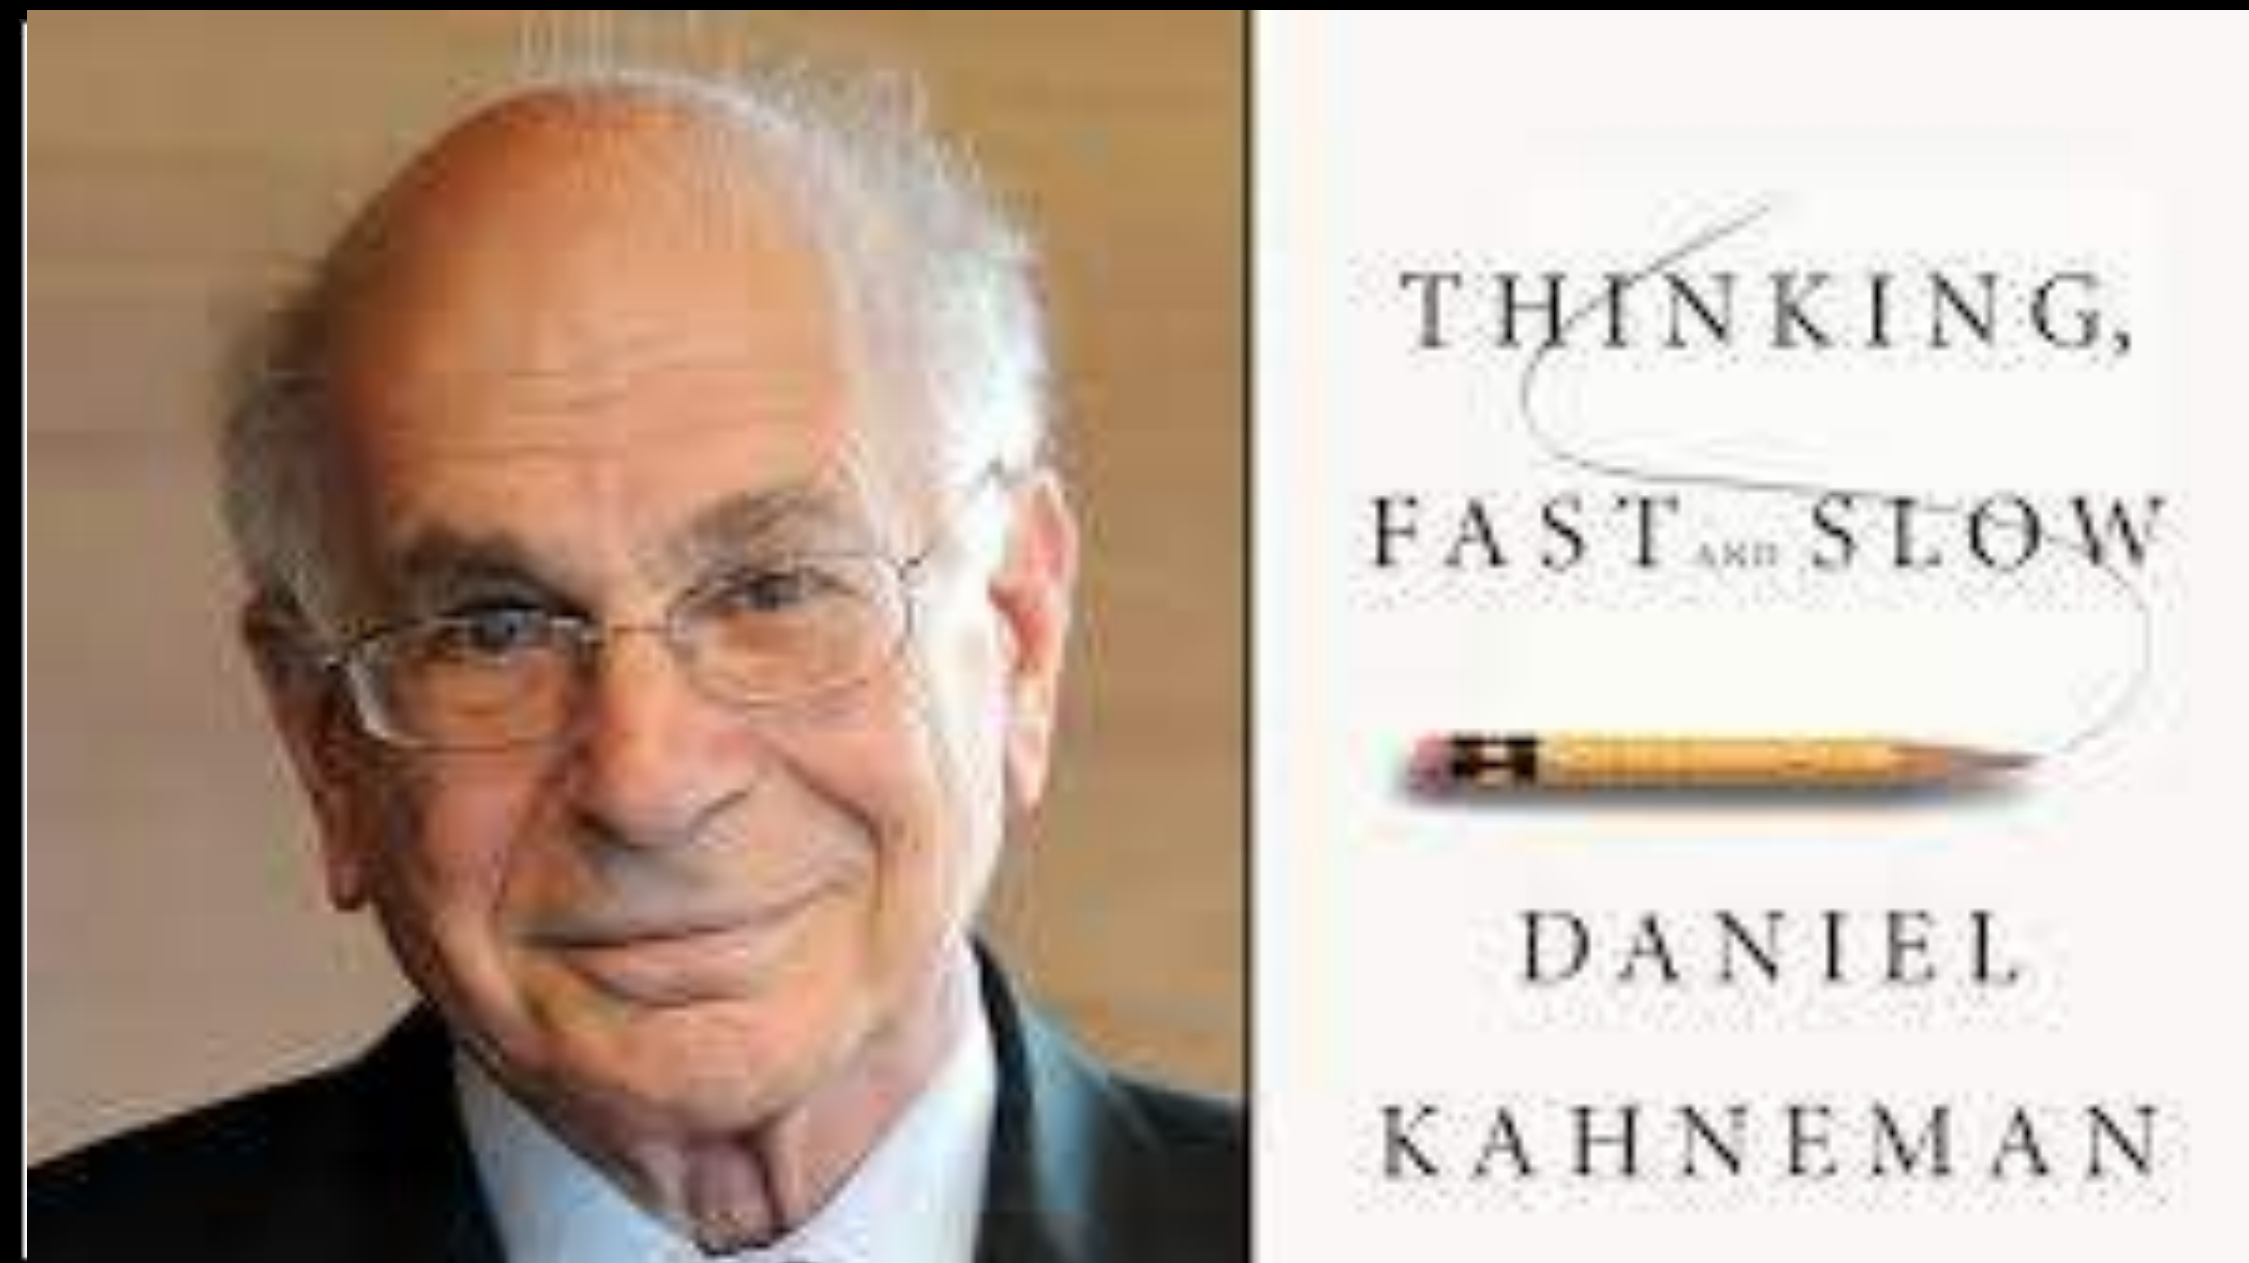


Artificial intelligence



# ***PATTERNS vs RULES***

## **FAST vs SLOW**



# **A brief history of Artificial Intelligence**

$+$  means' union of set with afterwords  
 $+$  also means  $+$   
 ii. Associativity  
 $a + = a + 1$   
 $N_0$  is the class of natural numbers.  
Transitivity  
 $N_0 + \mathbb{N}$ ,  $0 \in N_0$ ,  $a \in N_0 \Rightarrow a +$   
 $\mathbb{N} \subseteq \mathbb{N}$ .  $0 \in \mathbb{N} = \mathbb{N}$ .  $x + 2 \in \mathbb{N}$ .  $x + 2 \in \mathbb{N}$ .  $N_0 \subseteq \mathbb{N}$   
 Also (also)  $a + = b + \Rightarrow a = b$ .  
 $a \in N_0 \Rightarrow a + = 0$   
  
 Transitive something with union of elements by  
 type of defining  $a \cup 0 = a$   
 $a \cup 1 = a \cup$   
 $a \cup 2 = a \cup$   
 -----  
 Hence  $a \cup (b \cup c) = (a \cup b) \cup c$

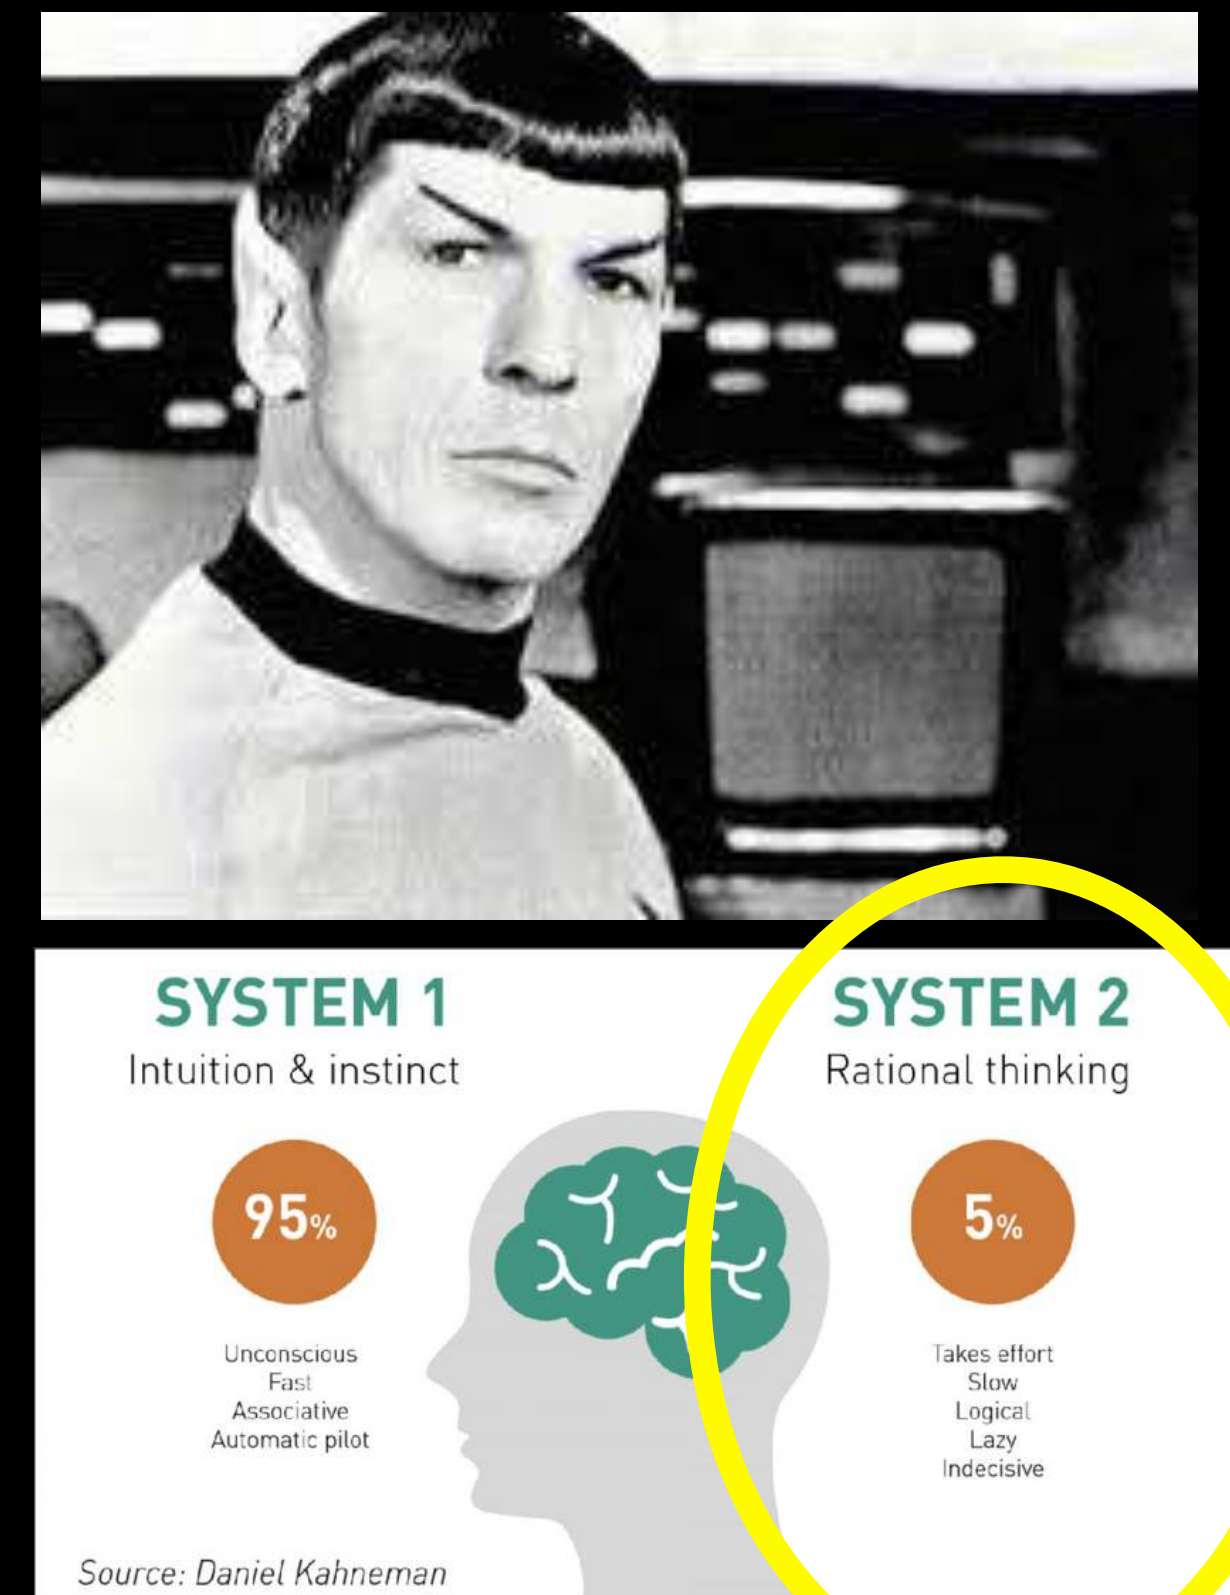


Alan Turing

# GOFAI: AI with logic (1957-2000)

Good Old-Fashioned AI

- Acquire knowledge with logic (IF-THEN rules)
- Body vs mind —> hardware vs software
- Knowledge about the brain is irrelevant
- Emphasis on *slow thinking*

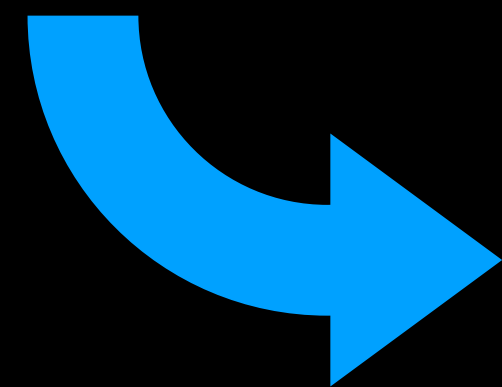


# GOFAI: Knowledge in IF-THEN rules

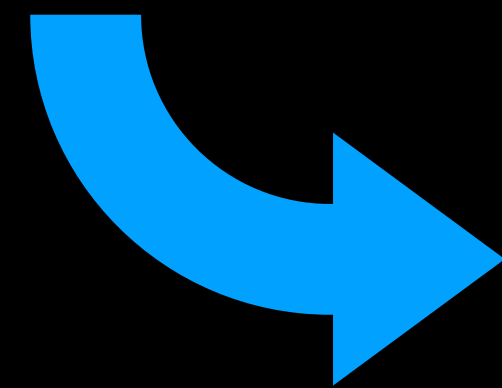


“Expert systems”

**IF go outside THEN check forecast**

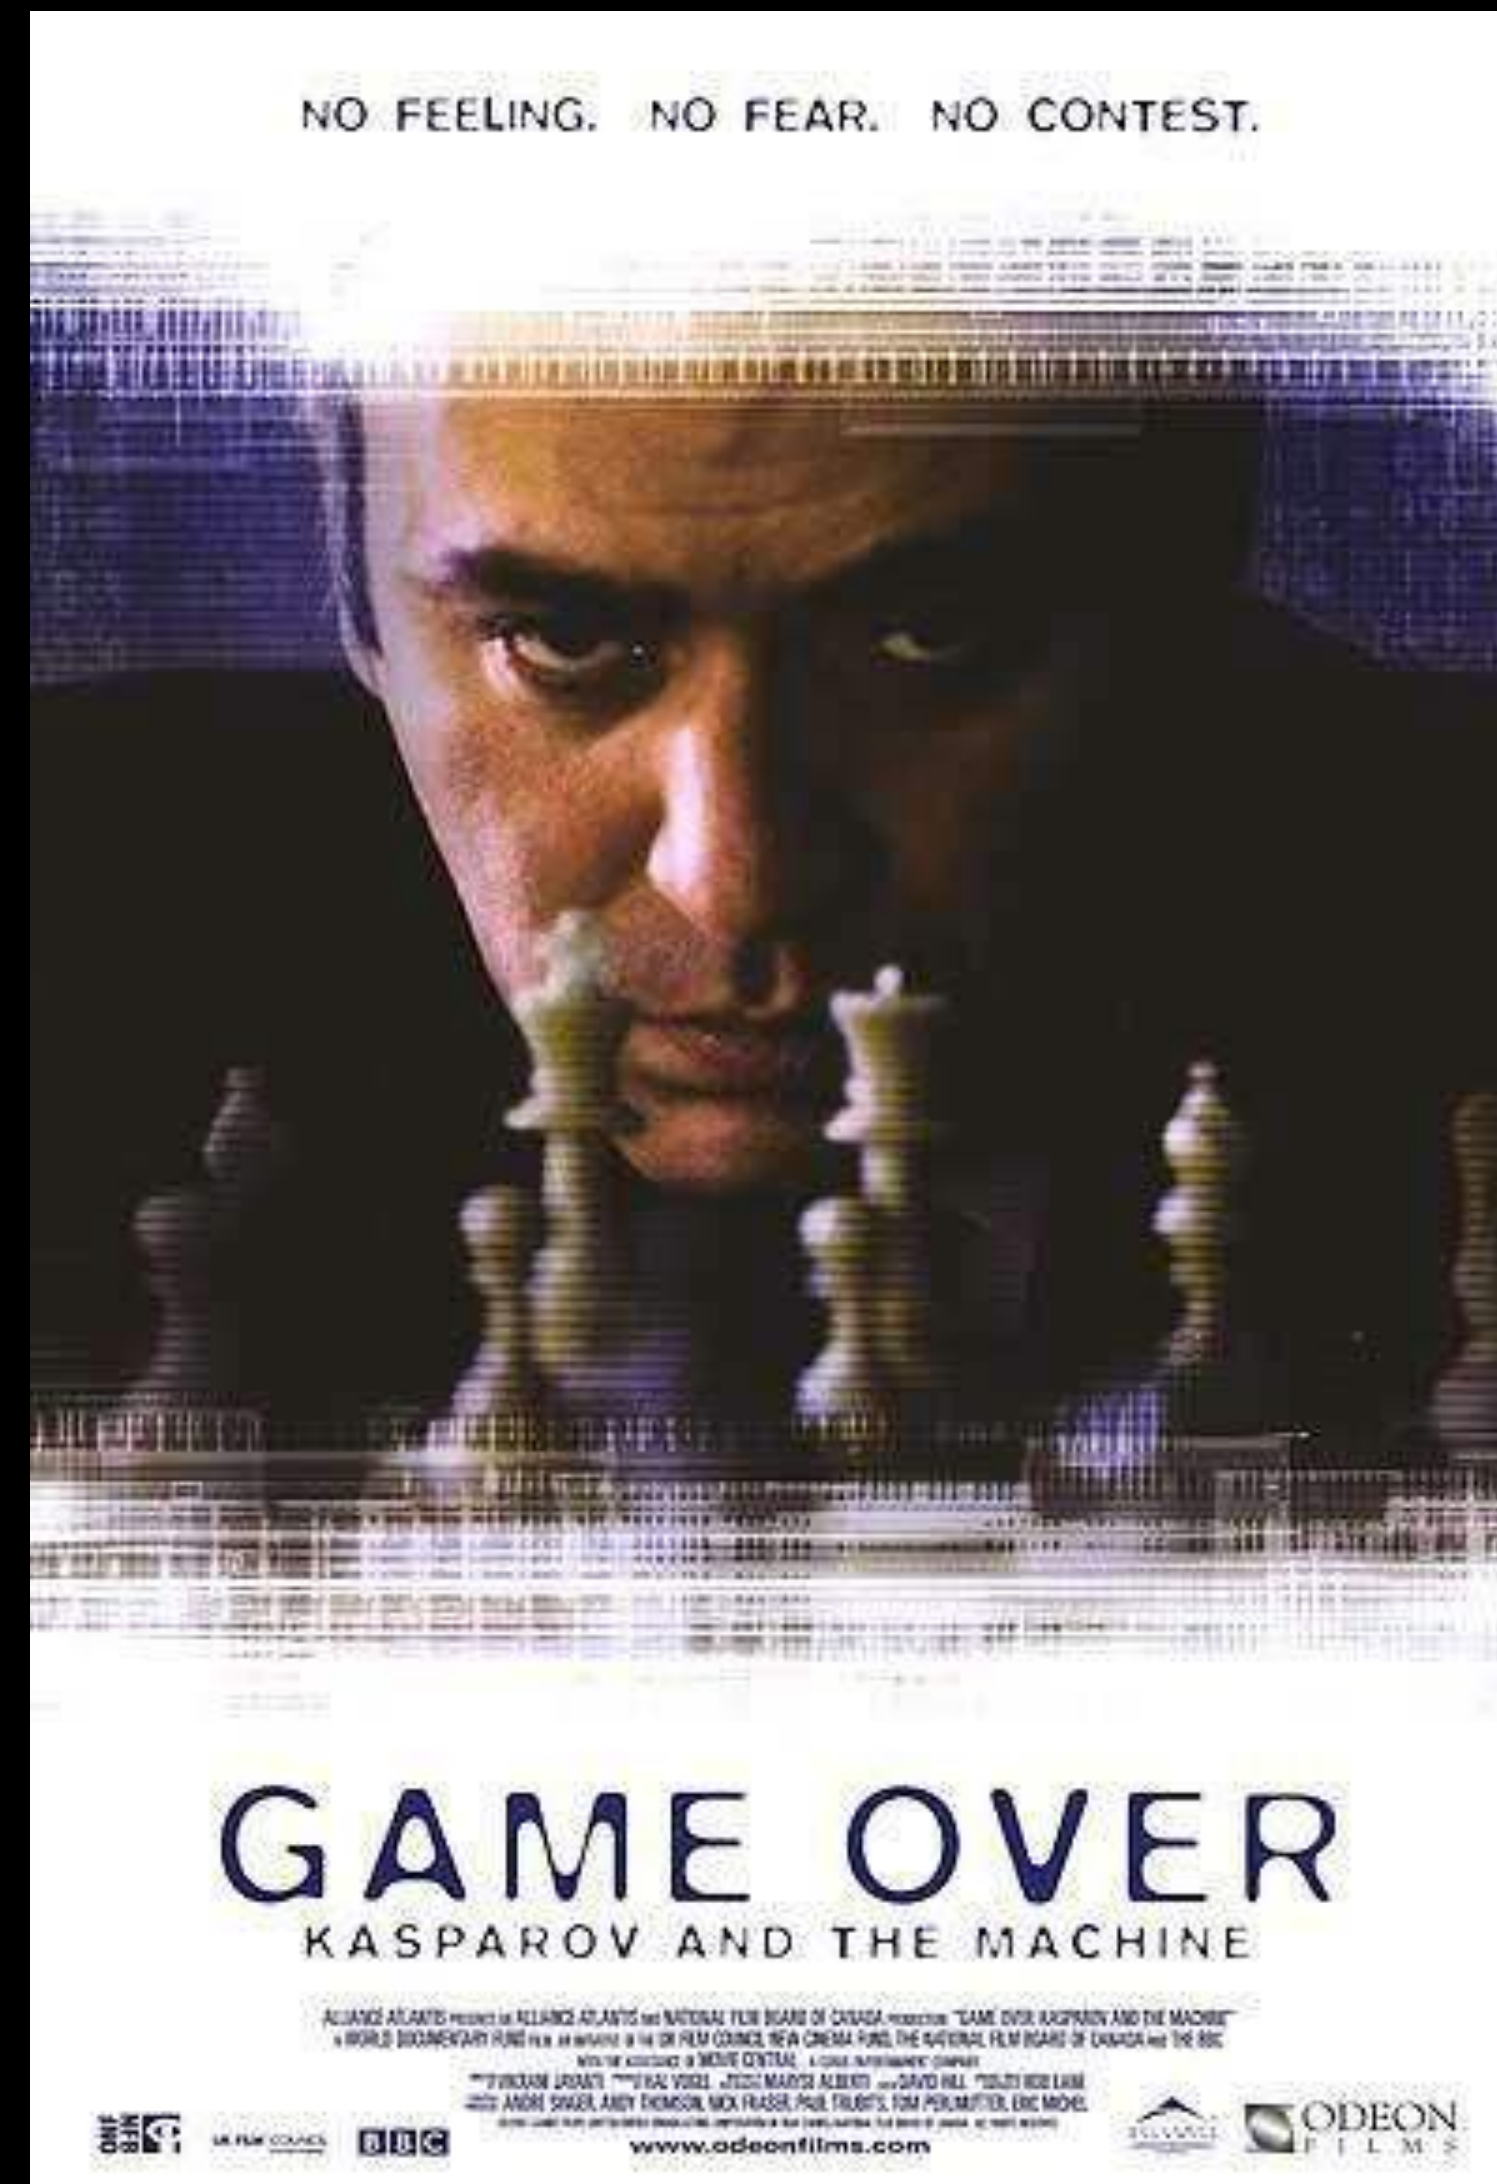


**IF forecast is rain THEN take umbrella**



**IF rain THEN open umbrella**

# Success of GOFAI: Deep Blue (1997)





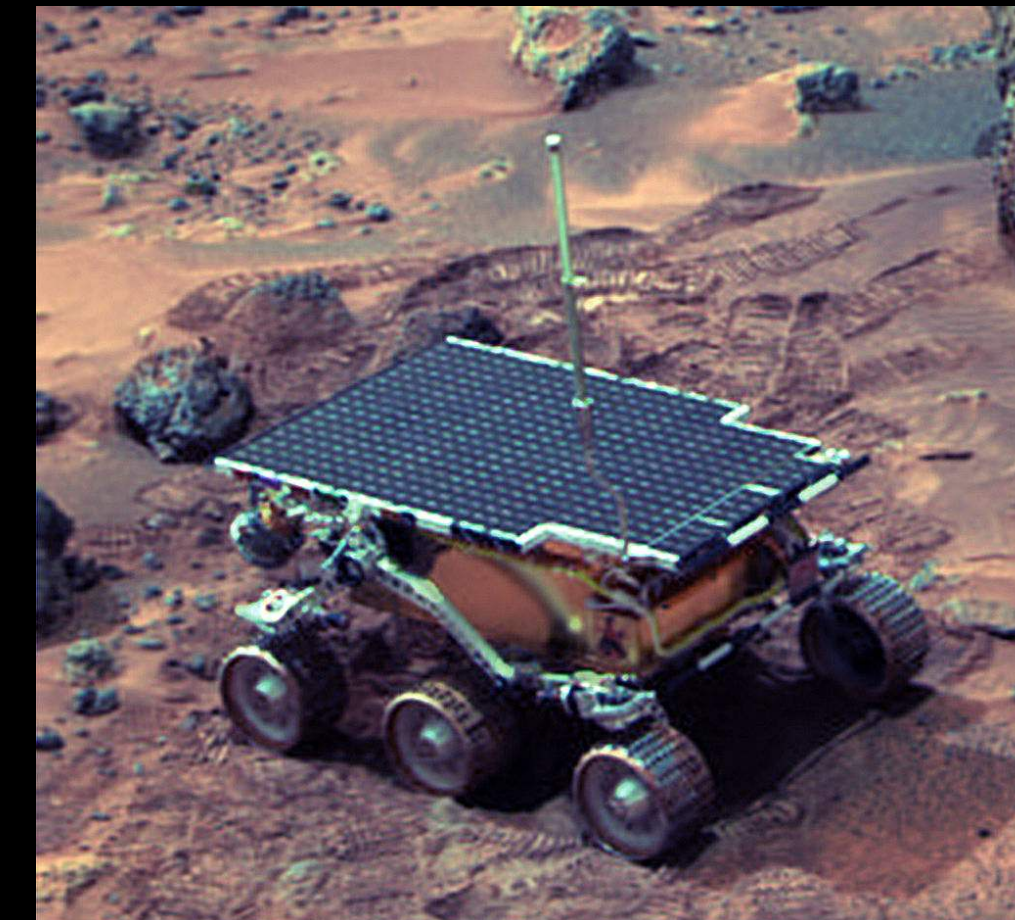
# Autonomous Robots

## EARTH, MOON AND MARS

— All distances to scale; bodies x20 larger —



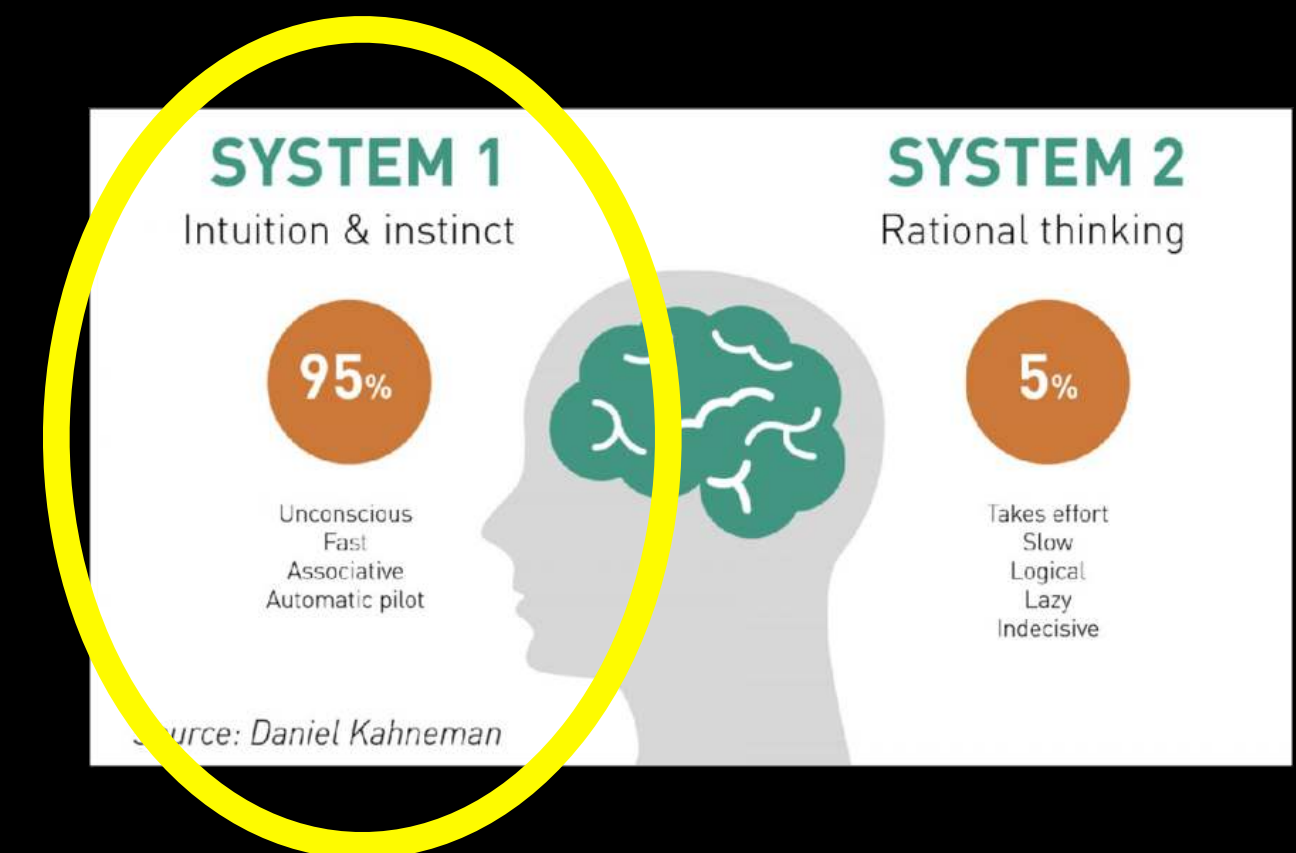
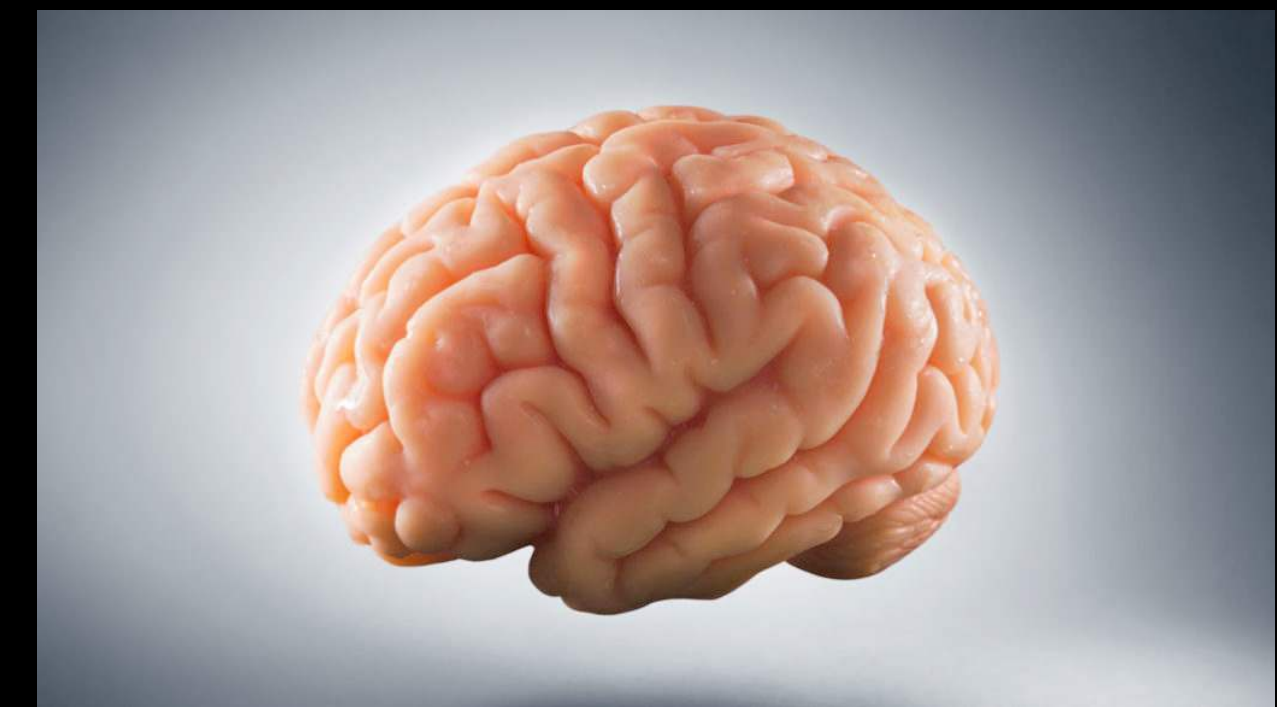
James O'Donoghue / NASA imagery — T: @physicsJ IG: jameslikesspace



**The Sojourner Rover (1997)**

# AI through learning from examples (2000-now)

- Knowledge acquisition through statistics (machine learning)
- Body = mind  $\rightarrow$  materialism
- Knowledge about the brain is relevant
- Emphasis on *fast thinking*

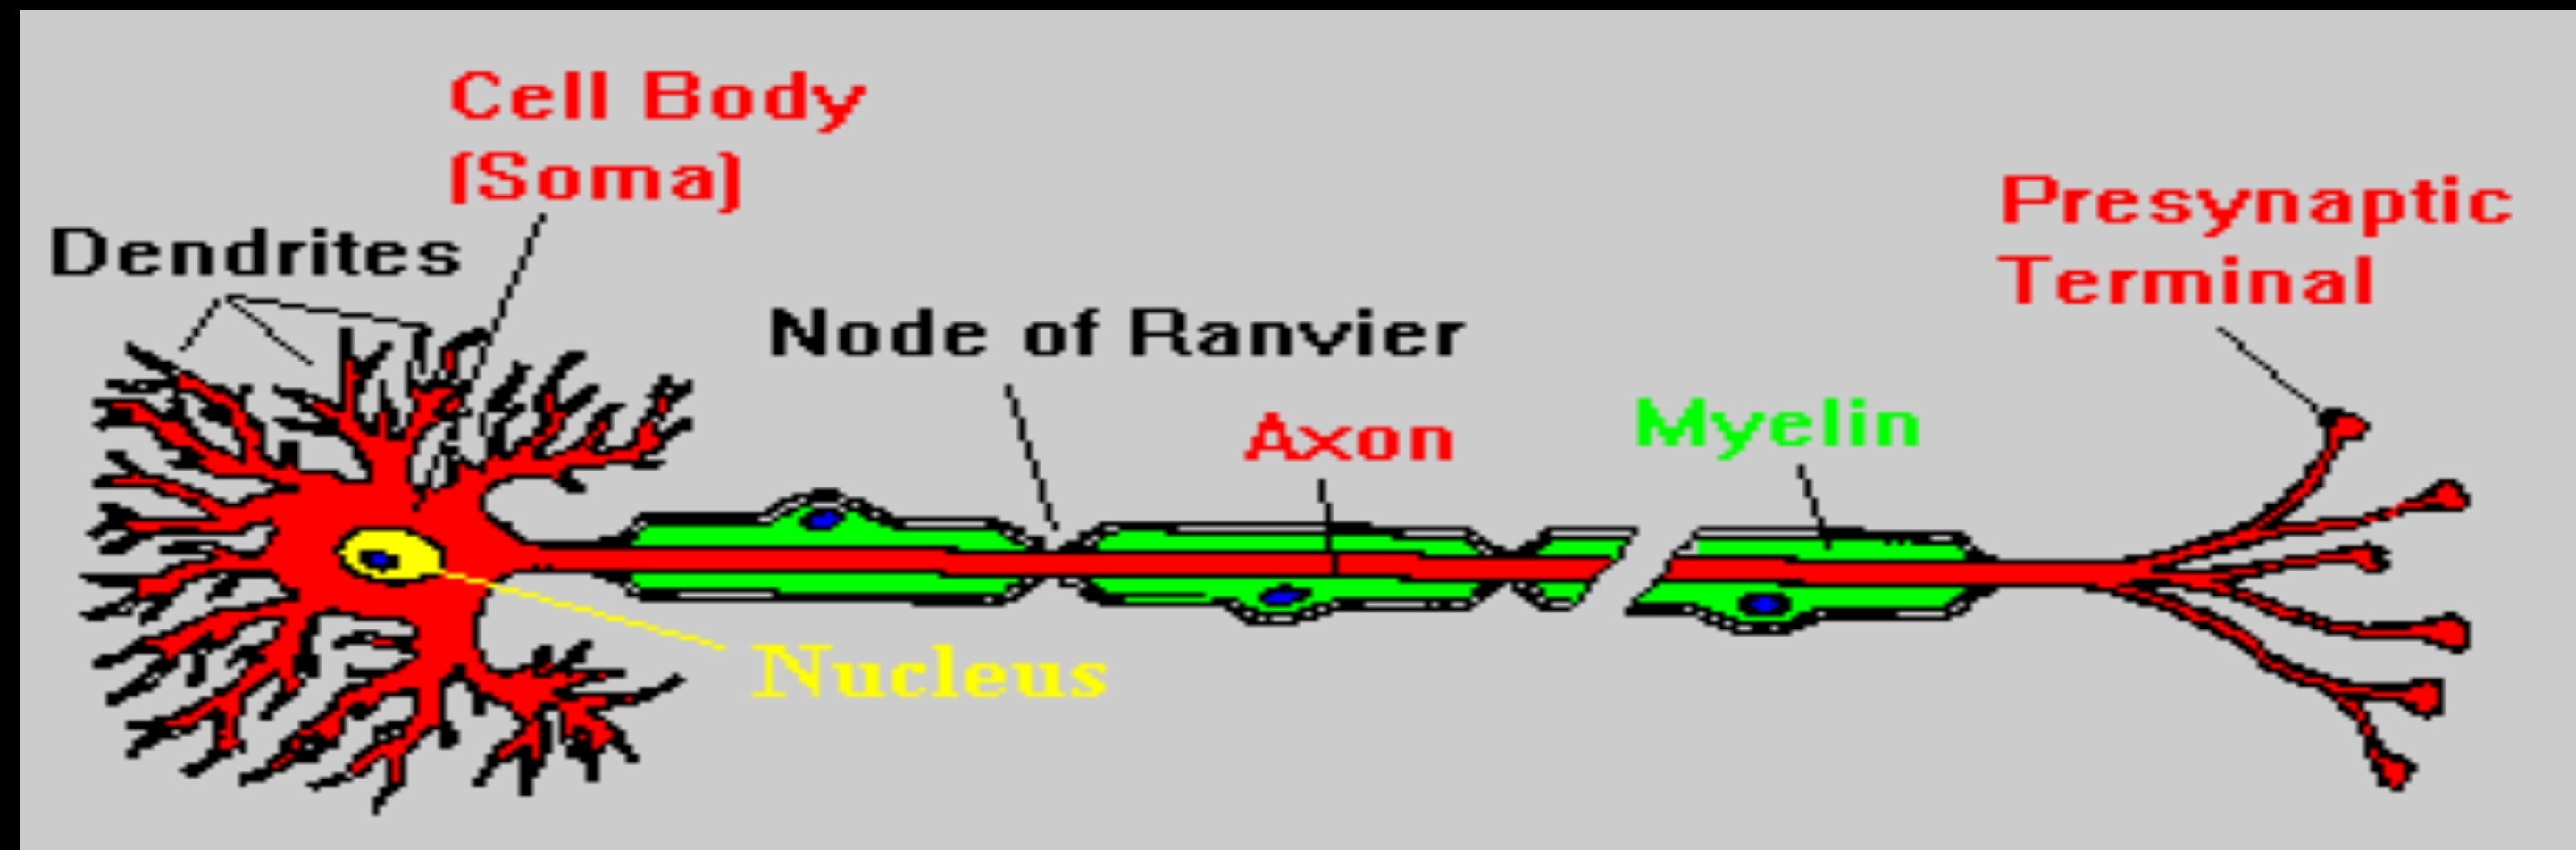
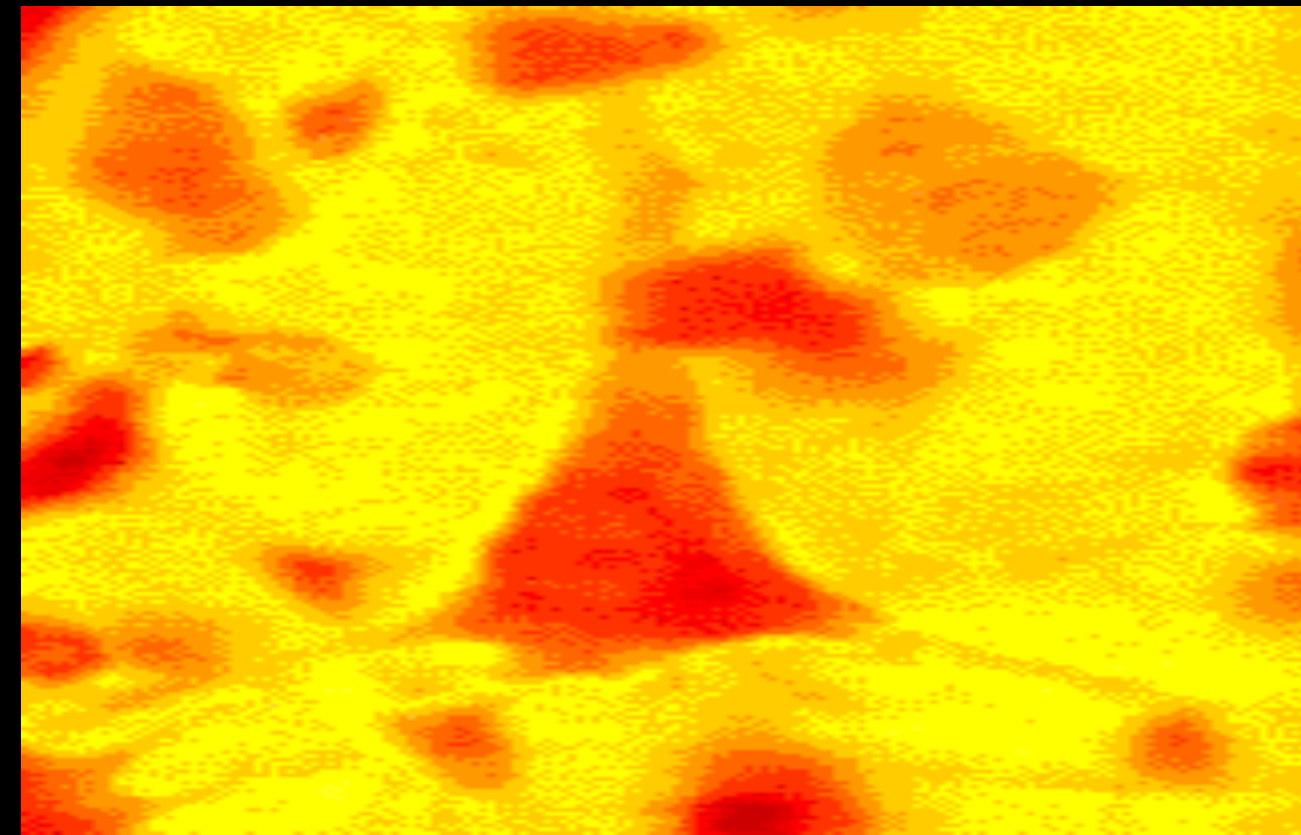


# Pattern Recognition

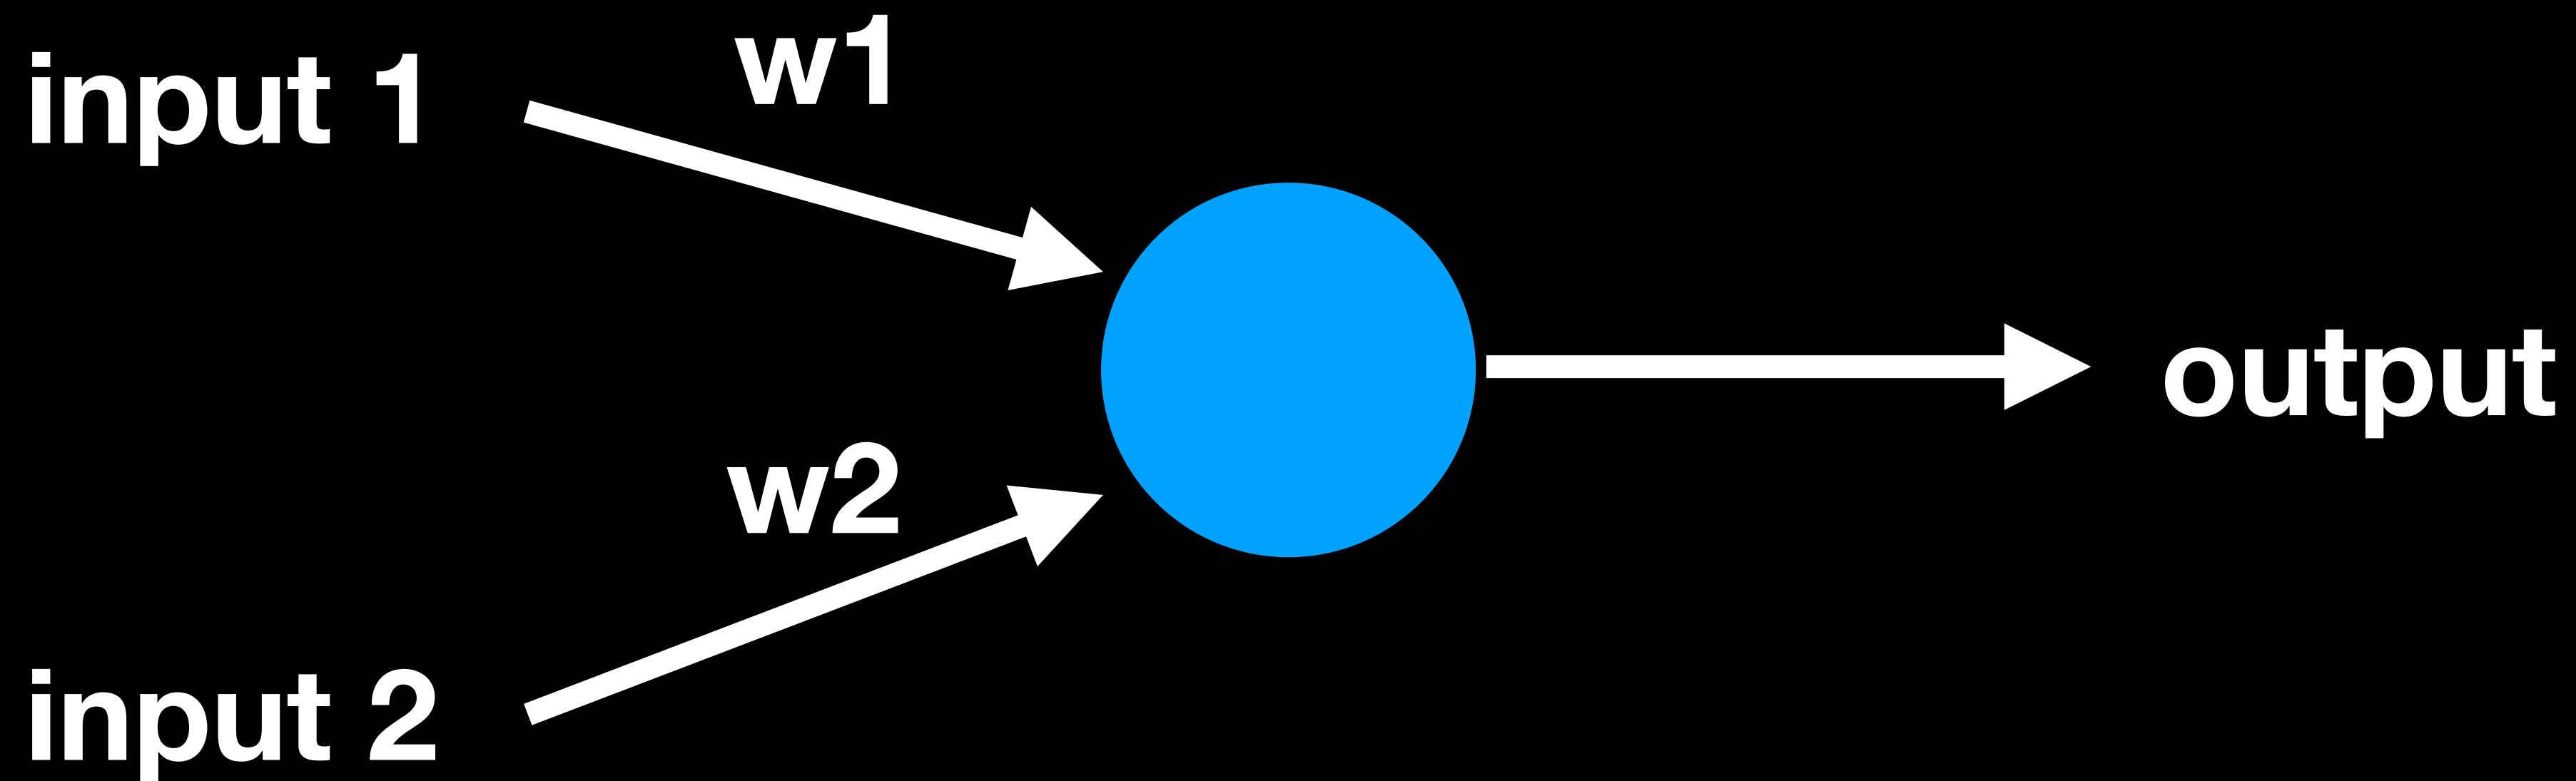
# How to recognise a dog?



# Artificial Neural Networks



# Perceptron 1958

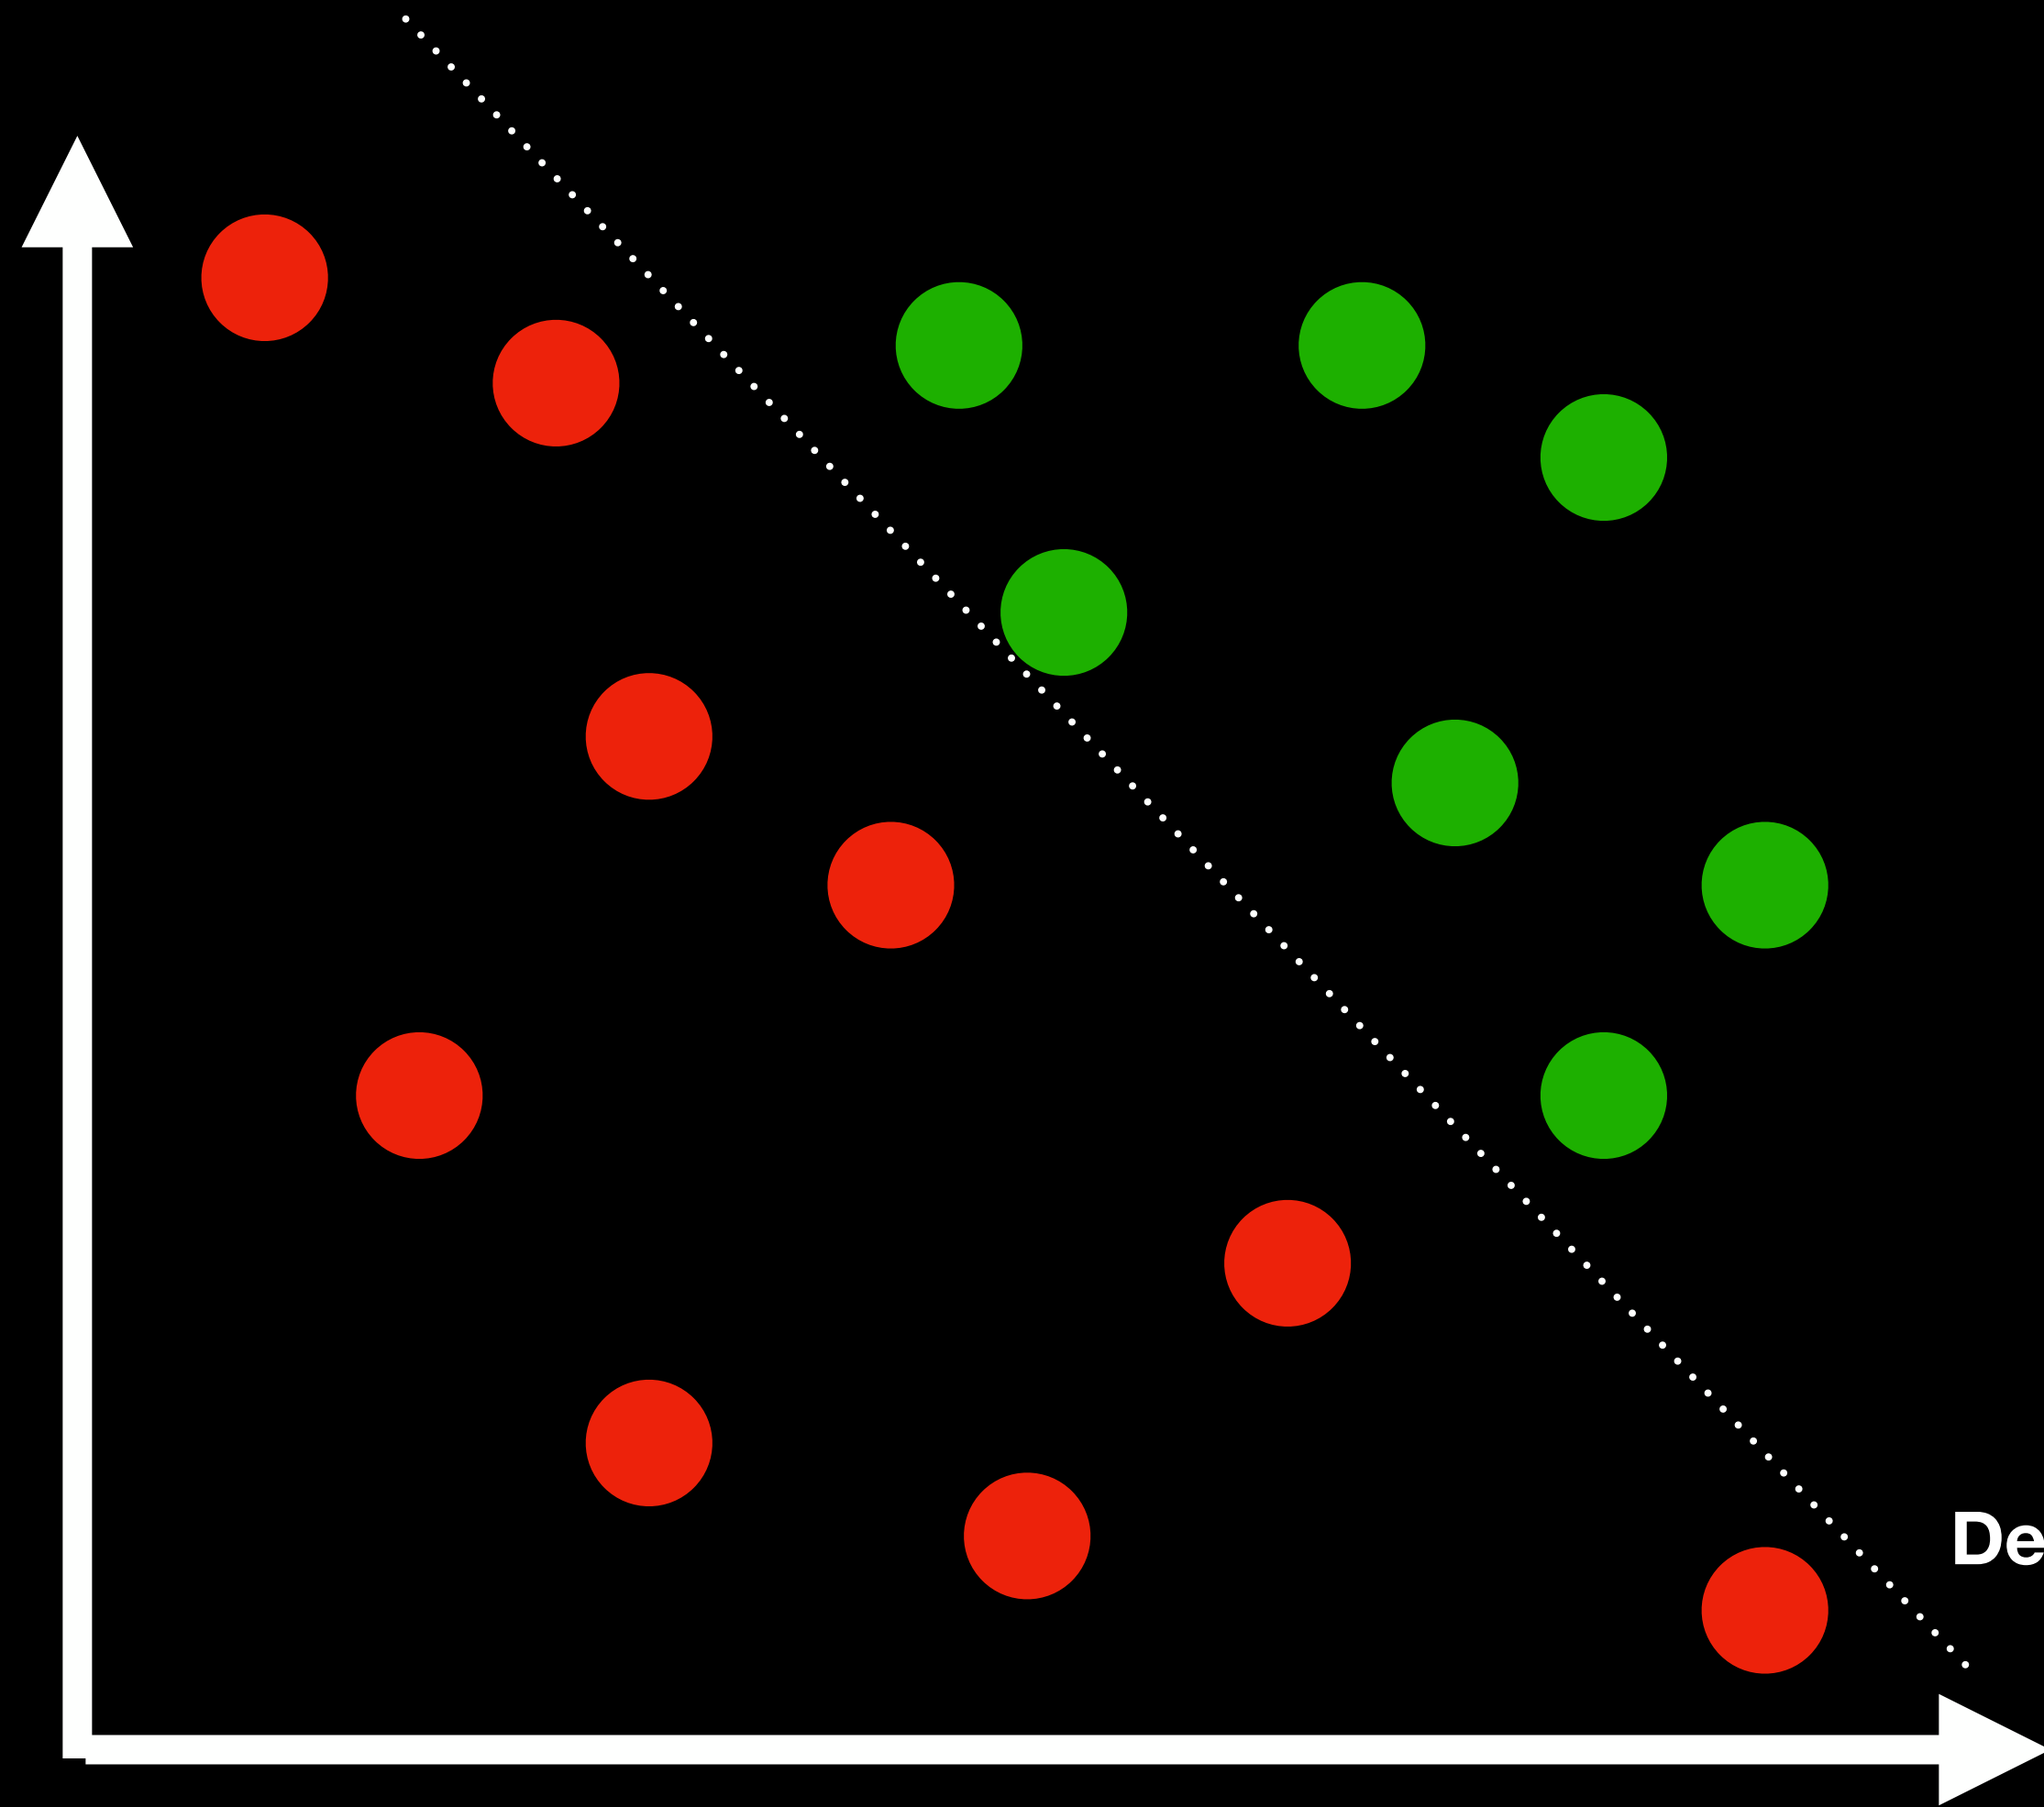


**ALS**  $(w1 \times \text{input 1} + w2 \times \text{input 2}) > \text{drempelwaarde DAN}$  **output = 1**  
**ANDERS** **output = 0**

# Classification (2 inputs)

- AI expert
- AI non expert

EP uren

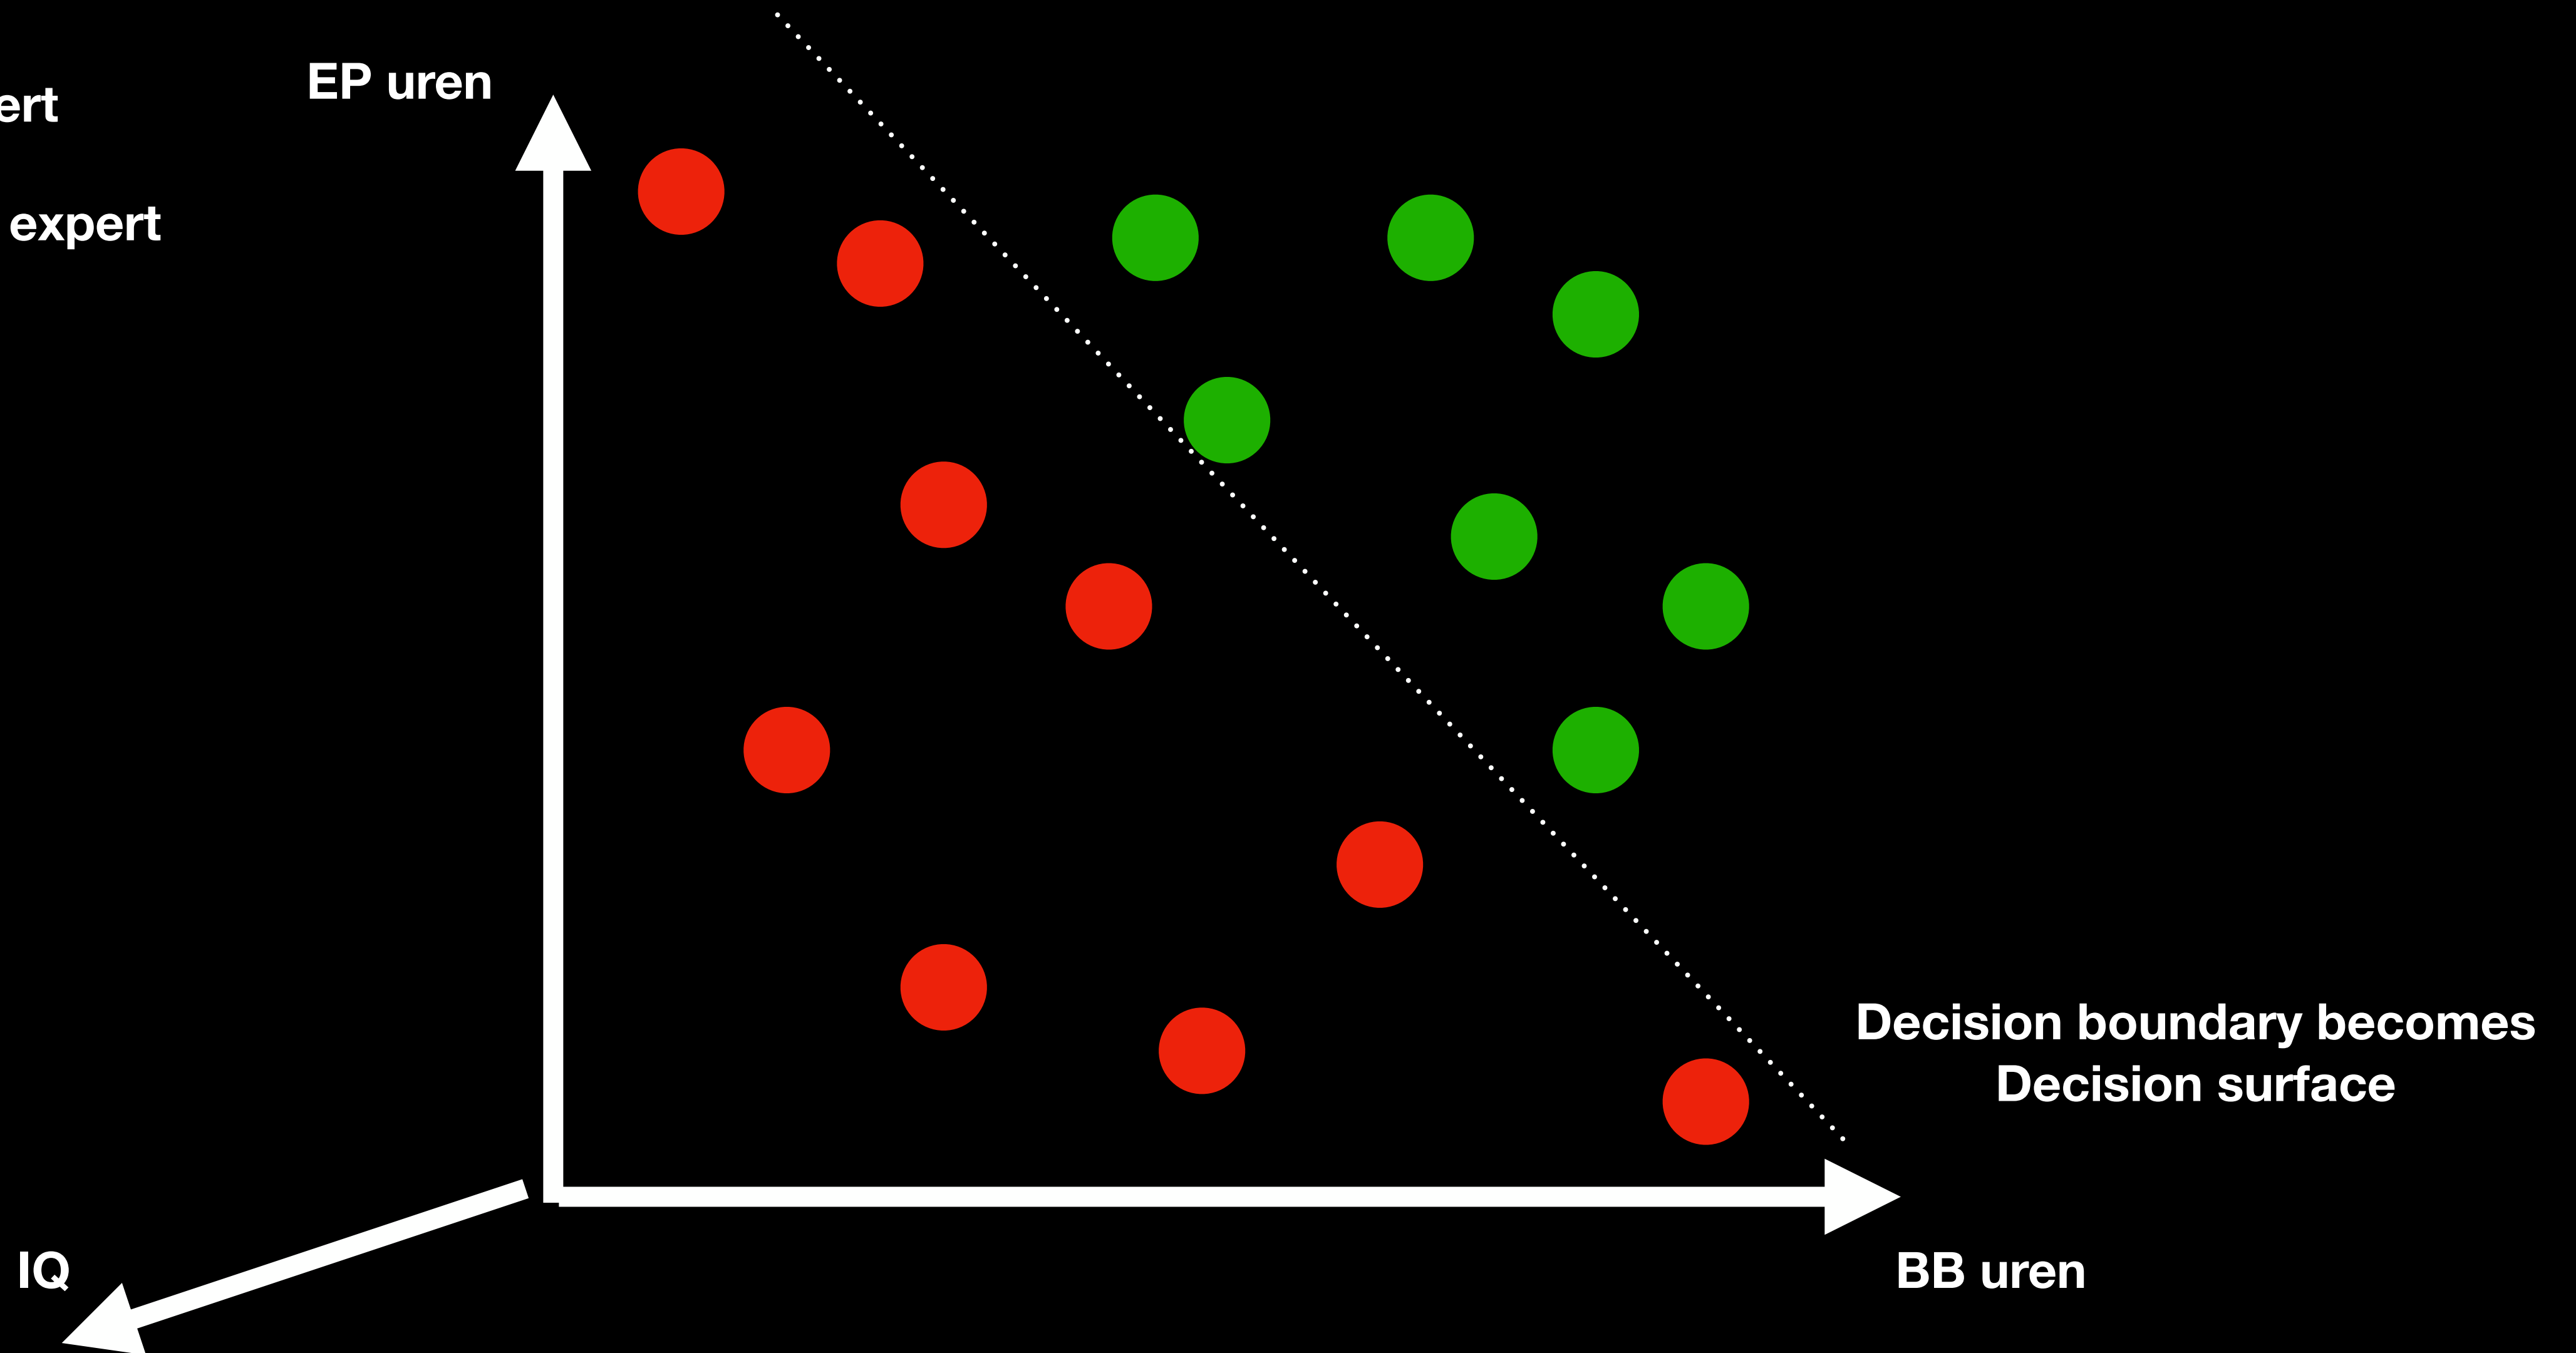


Decision boundary

BB uren

# Classification (3 inputs)

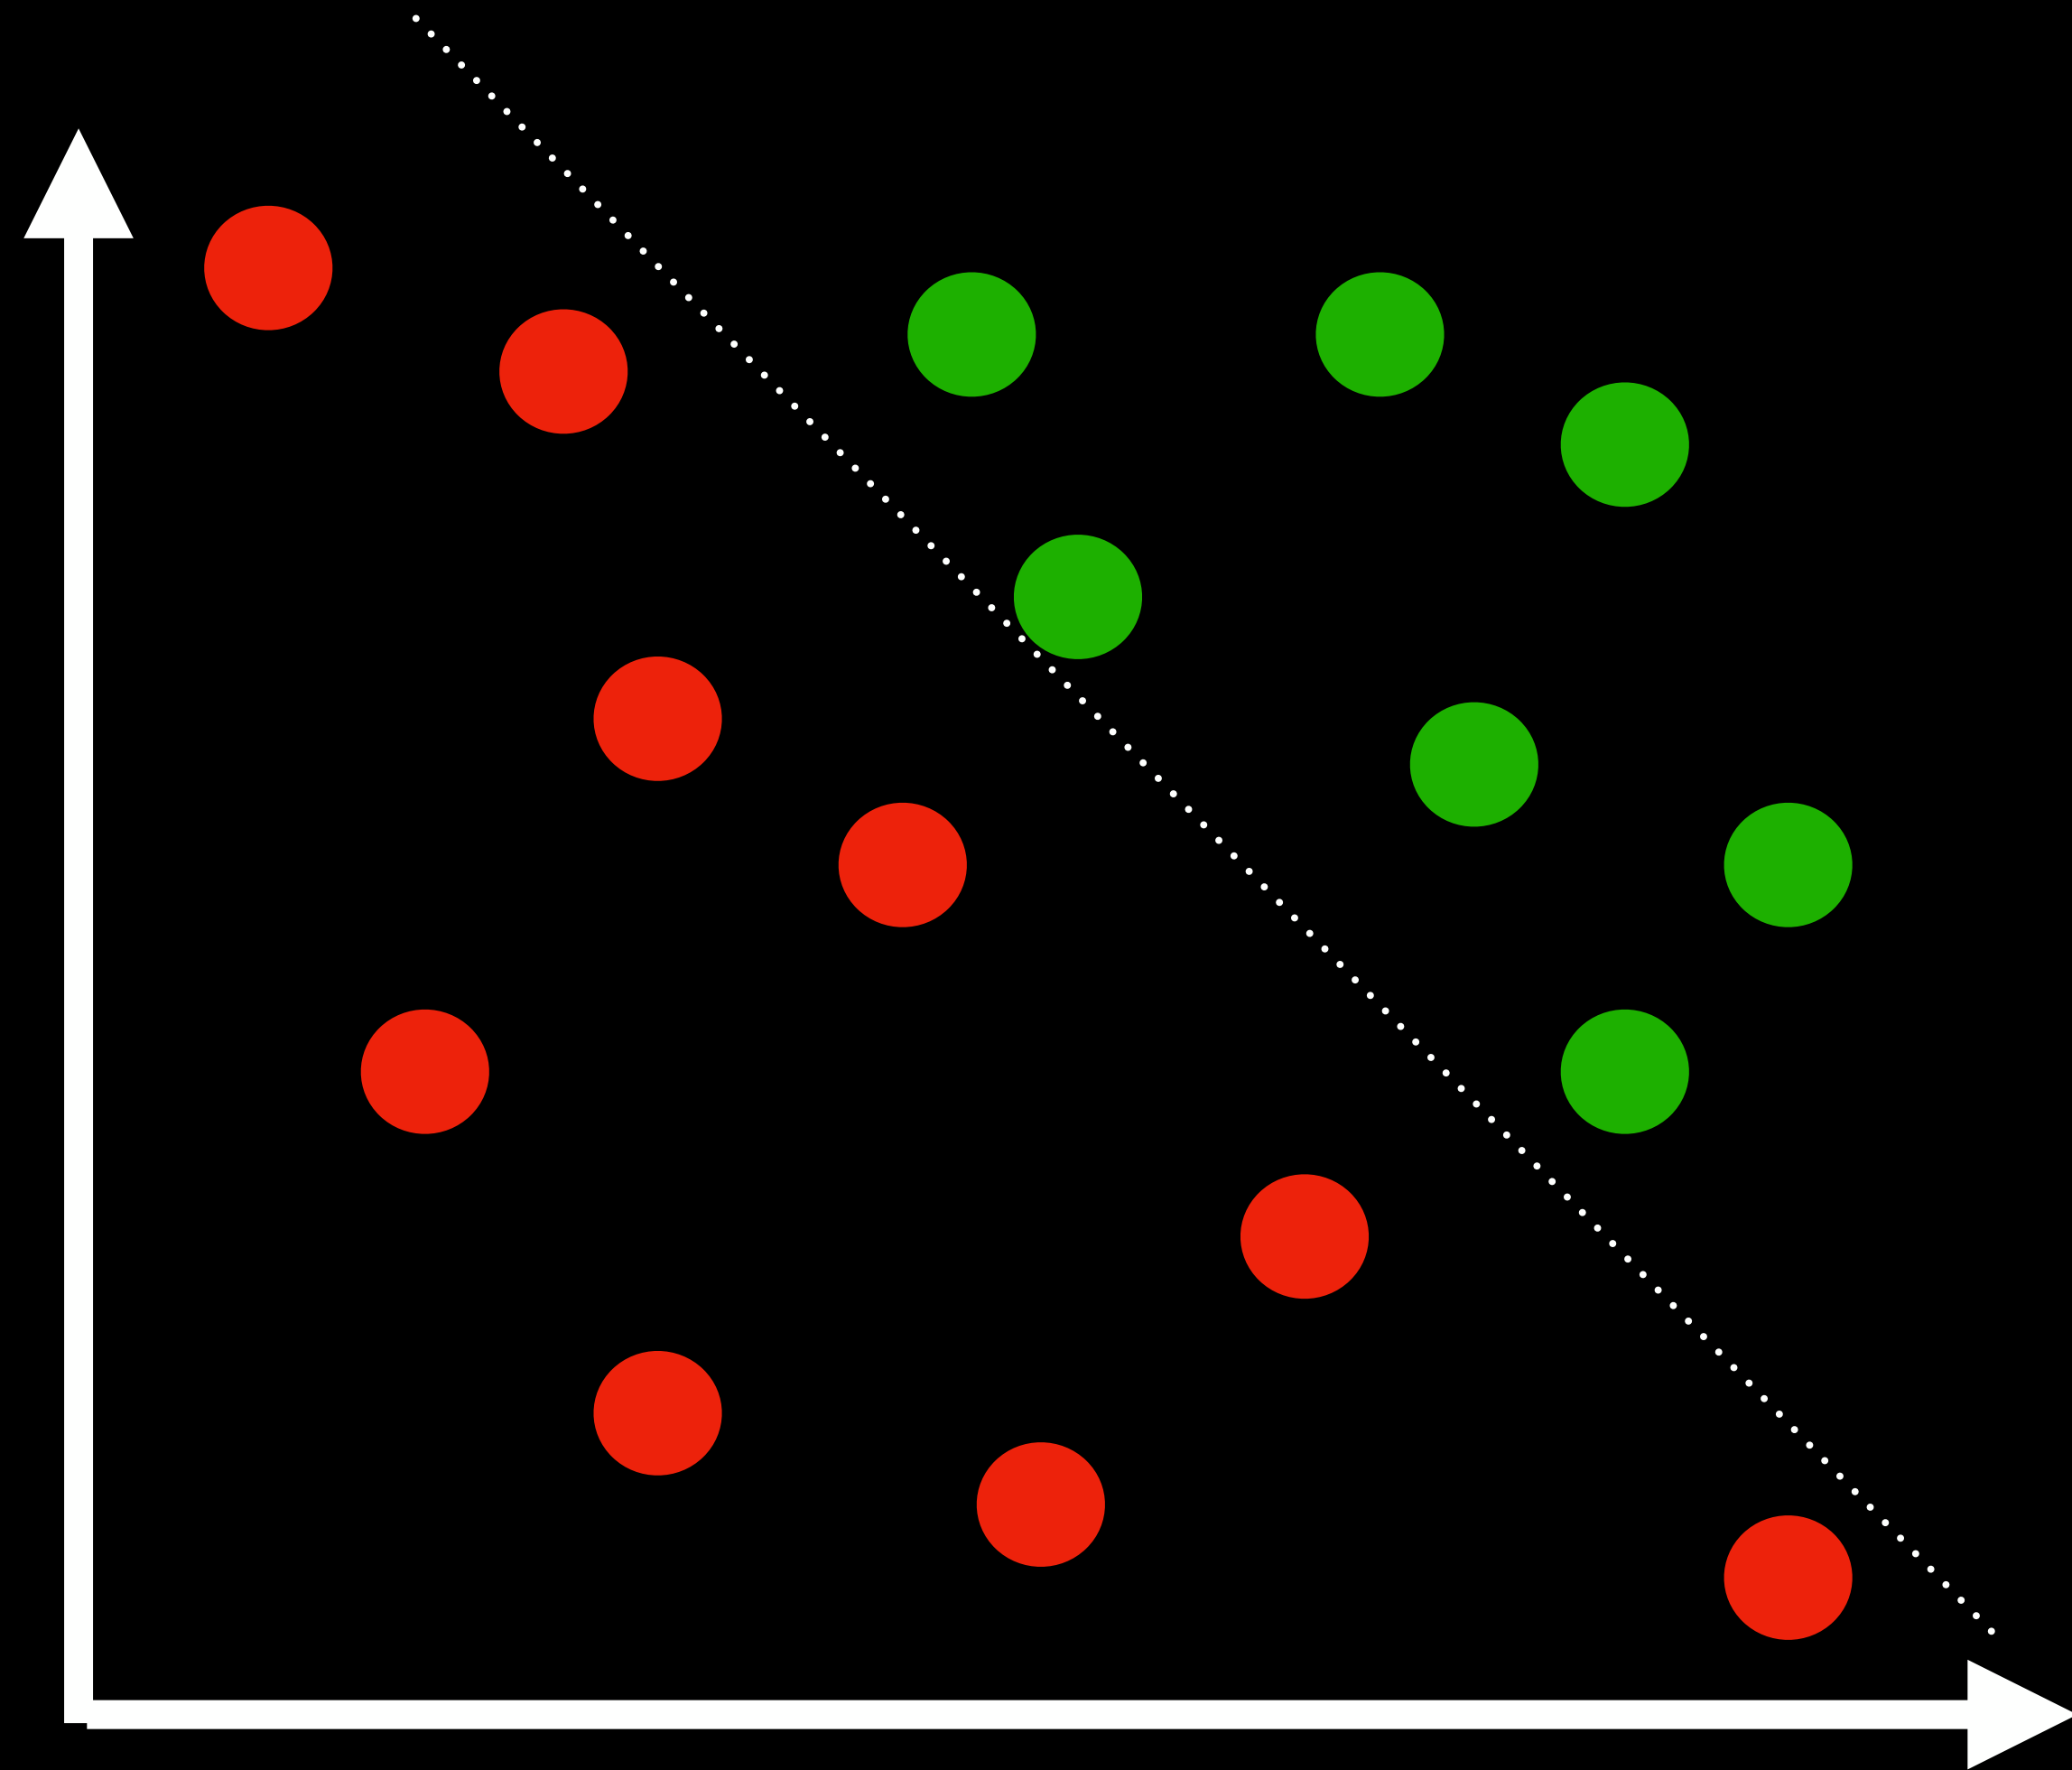
- AI expert
- AI non expert



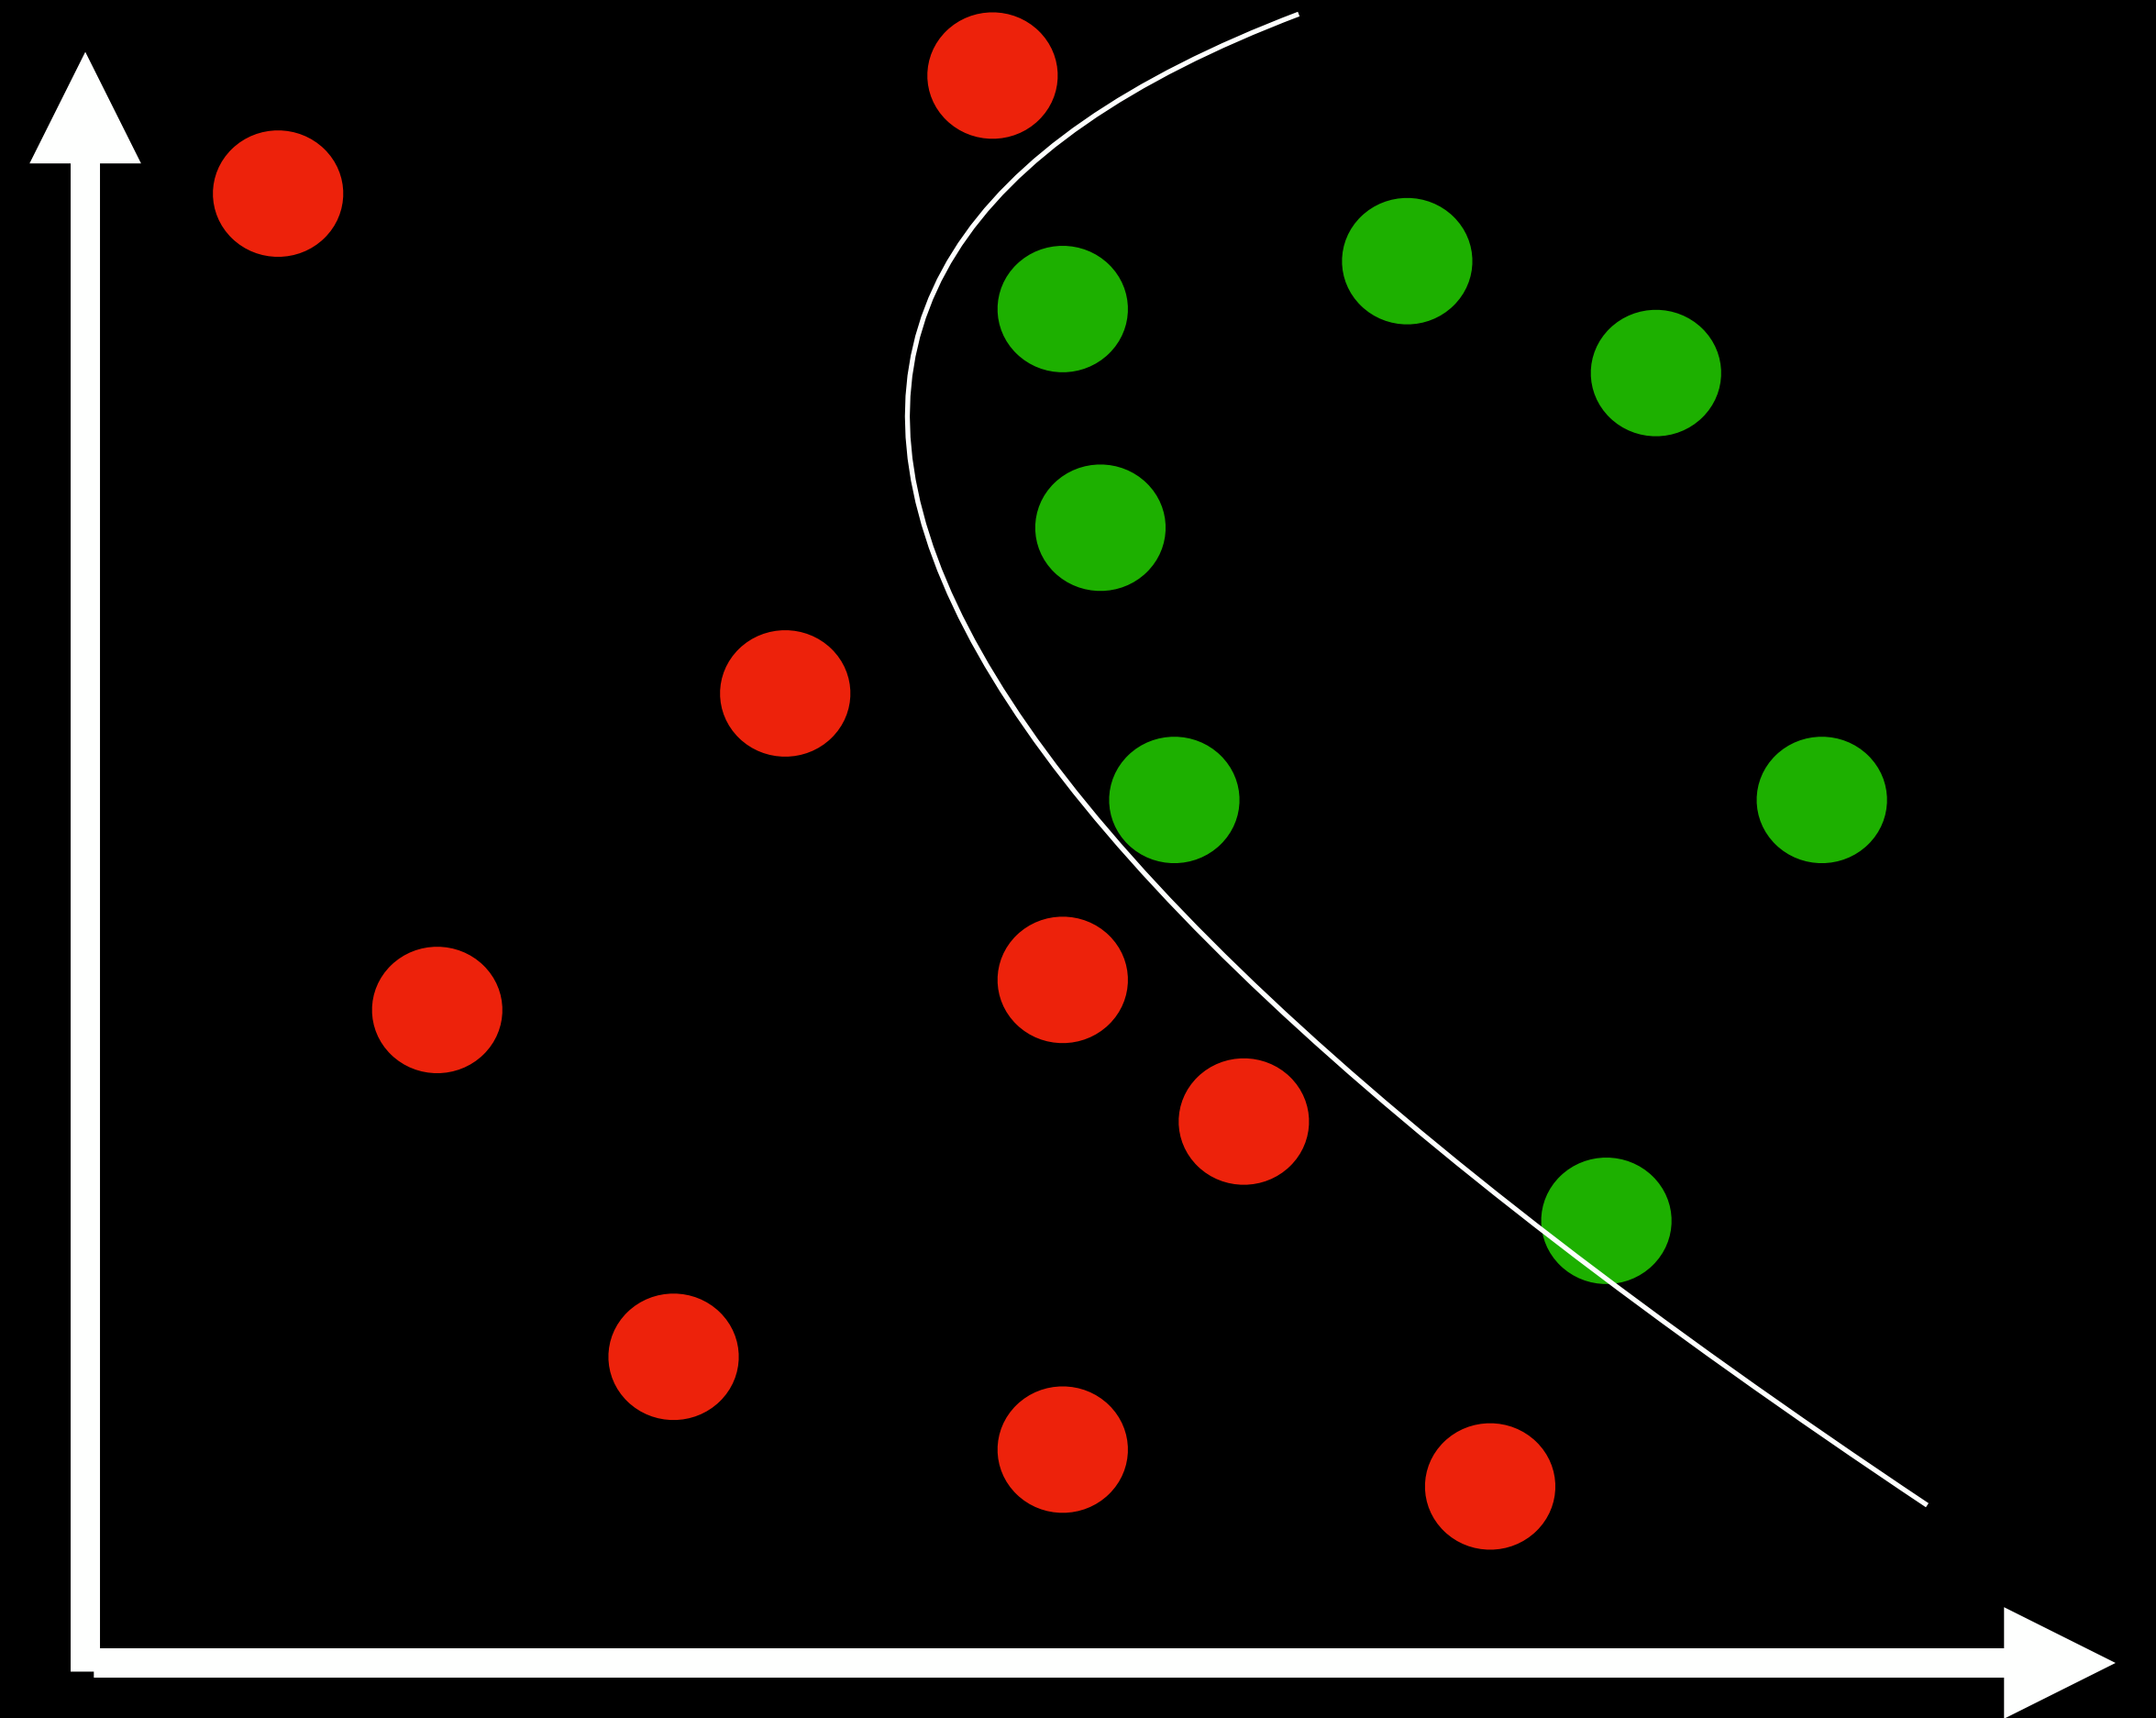




# Classification (2 inputs)

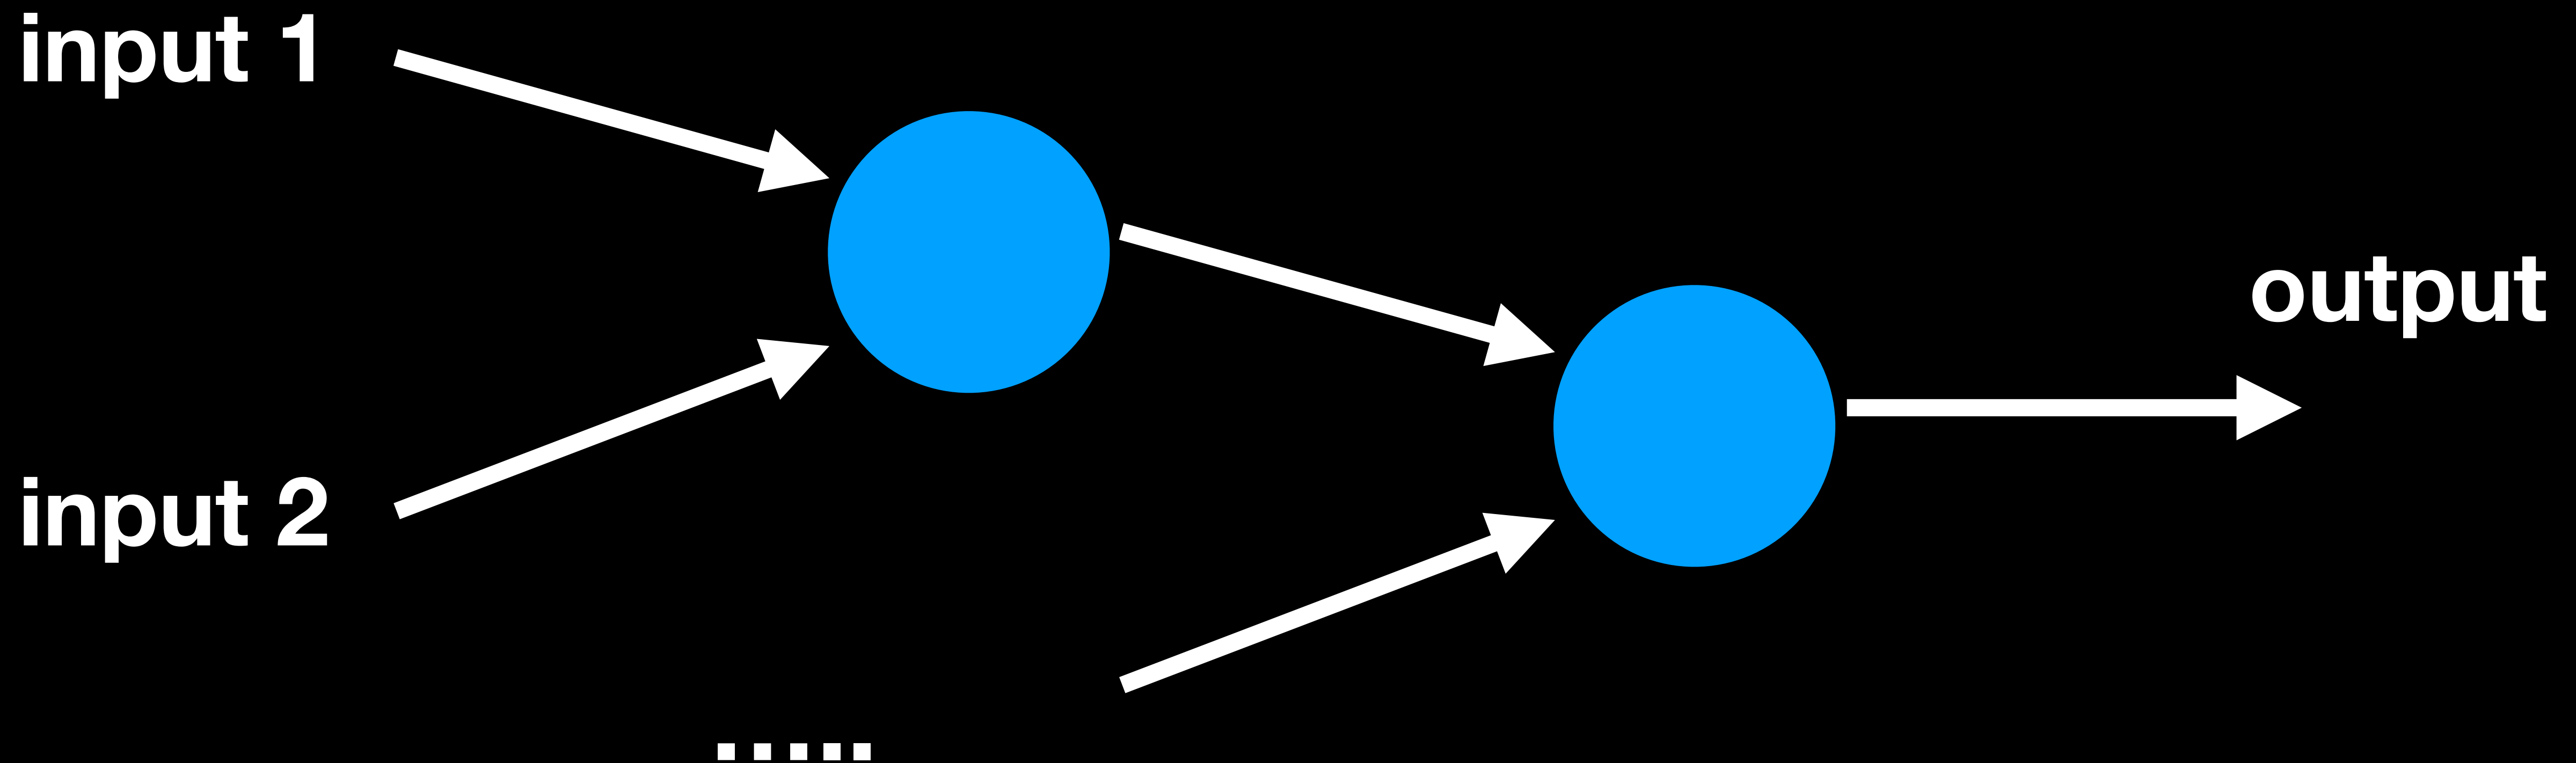


Linear decision boundary  
Perceptron can learn this

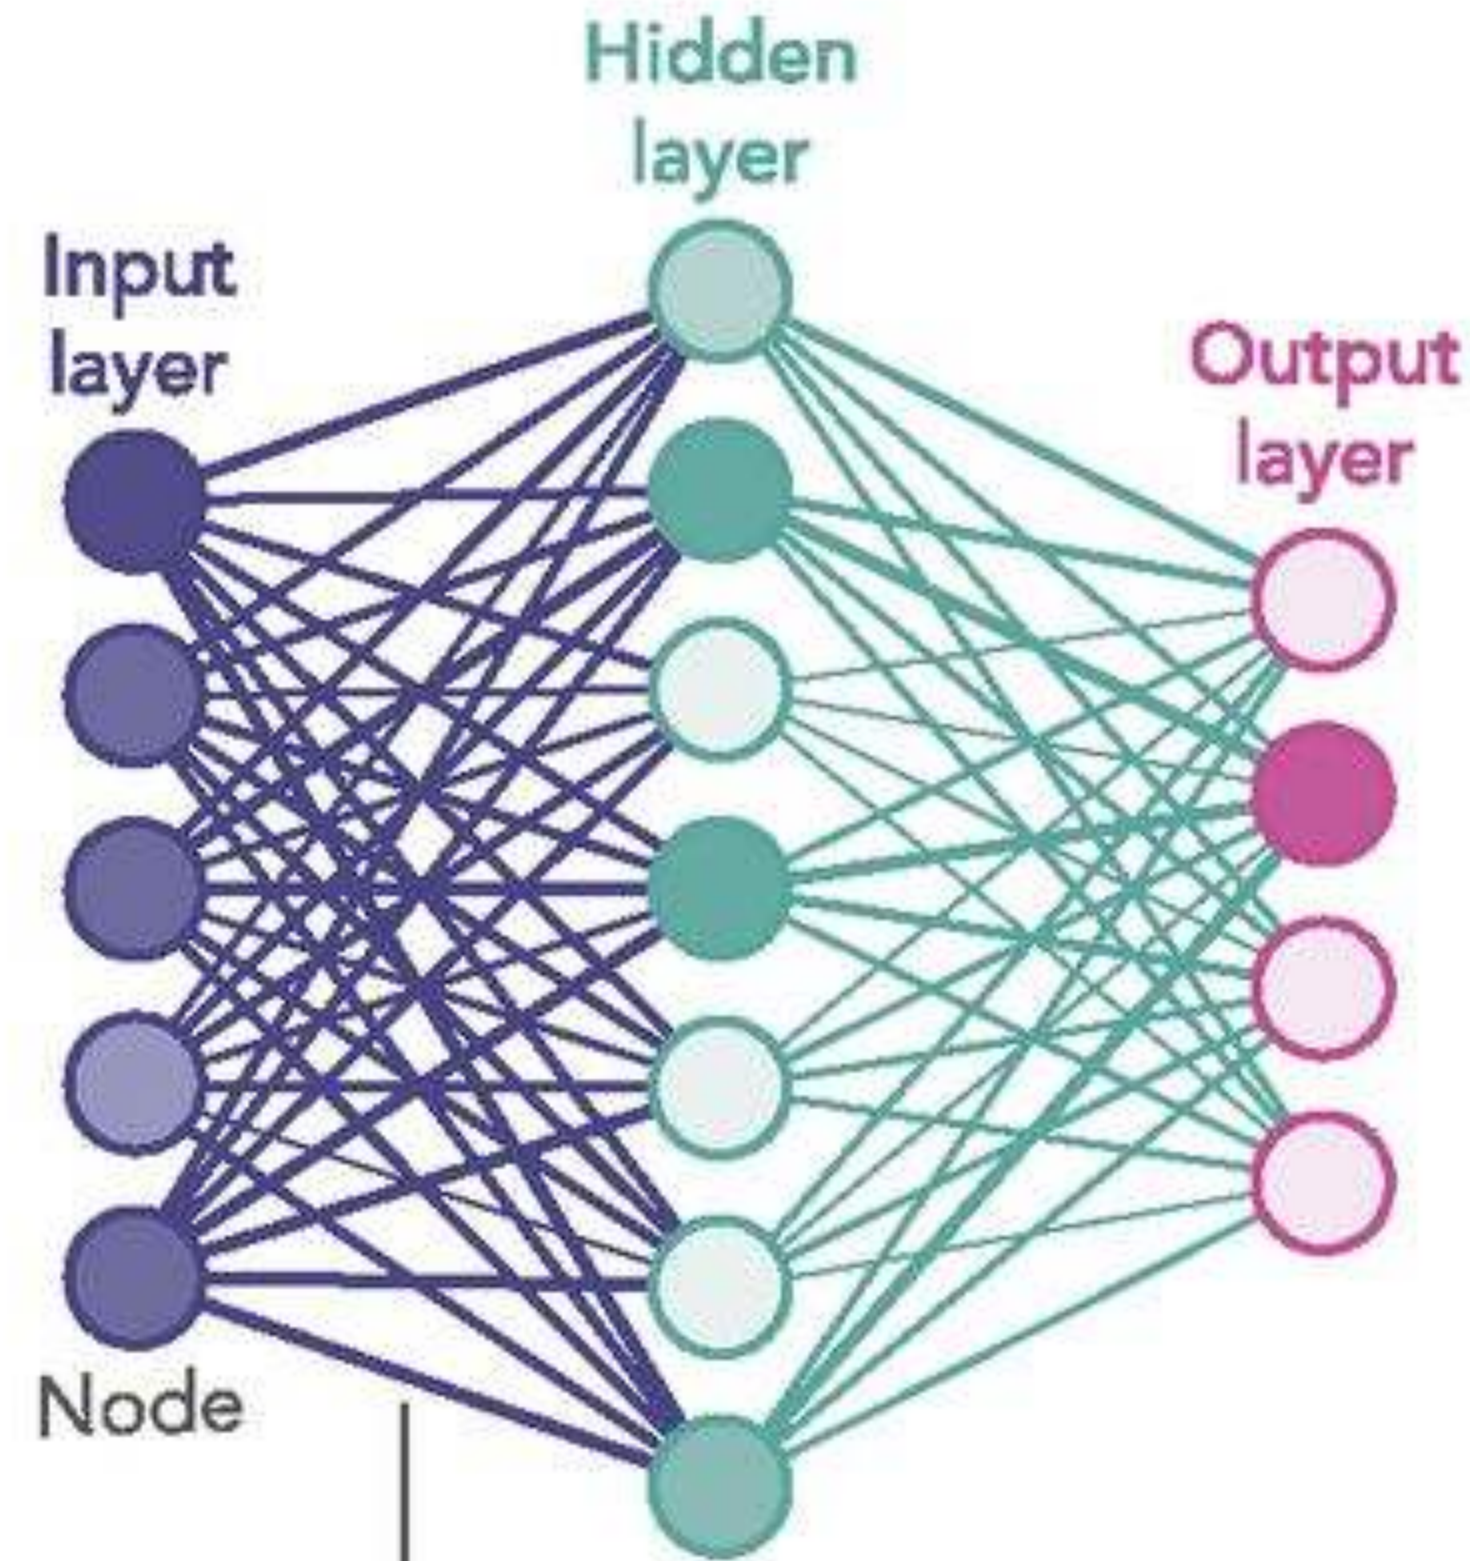


nonlinear decision boundary  
Perceptron can not learn this

# Perceptrons stapelen



# 1980's

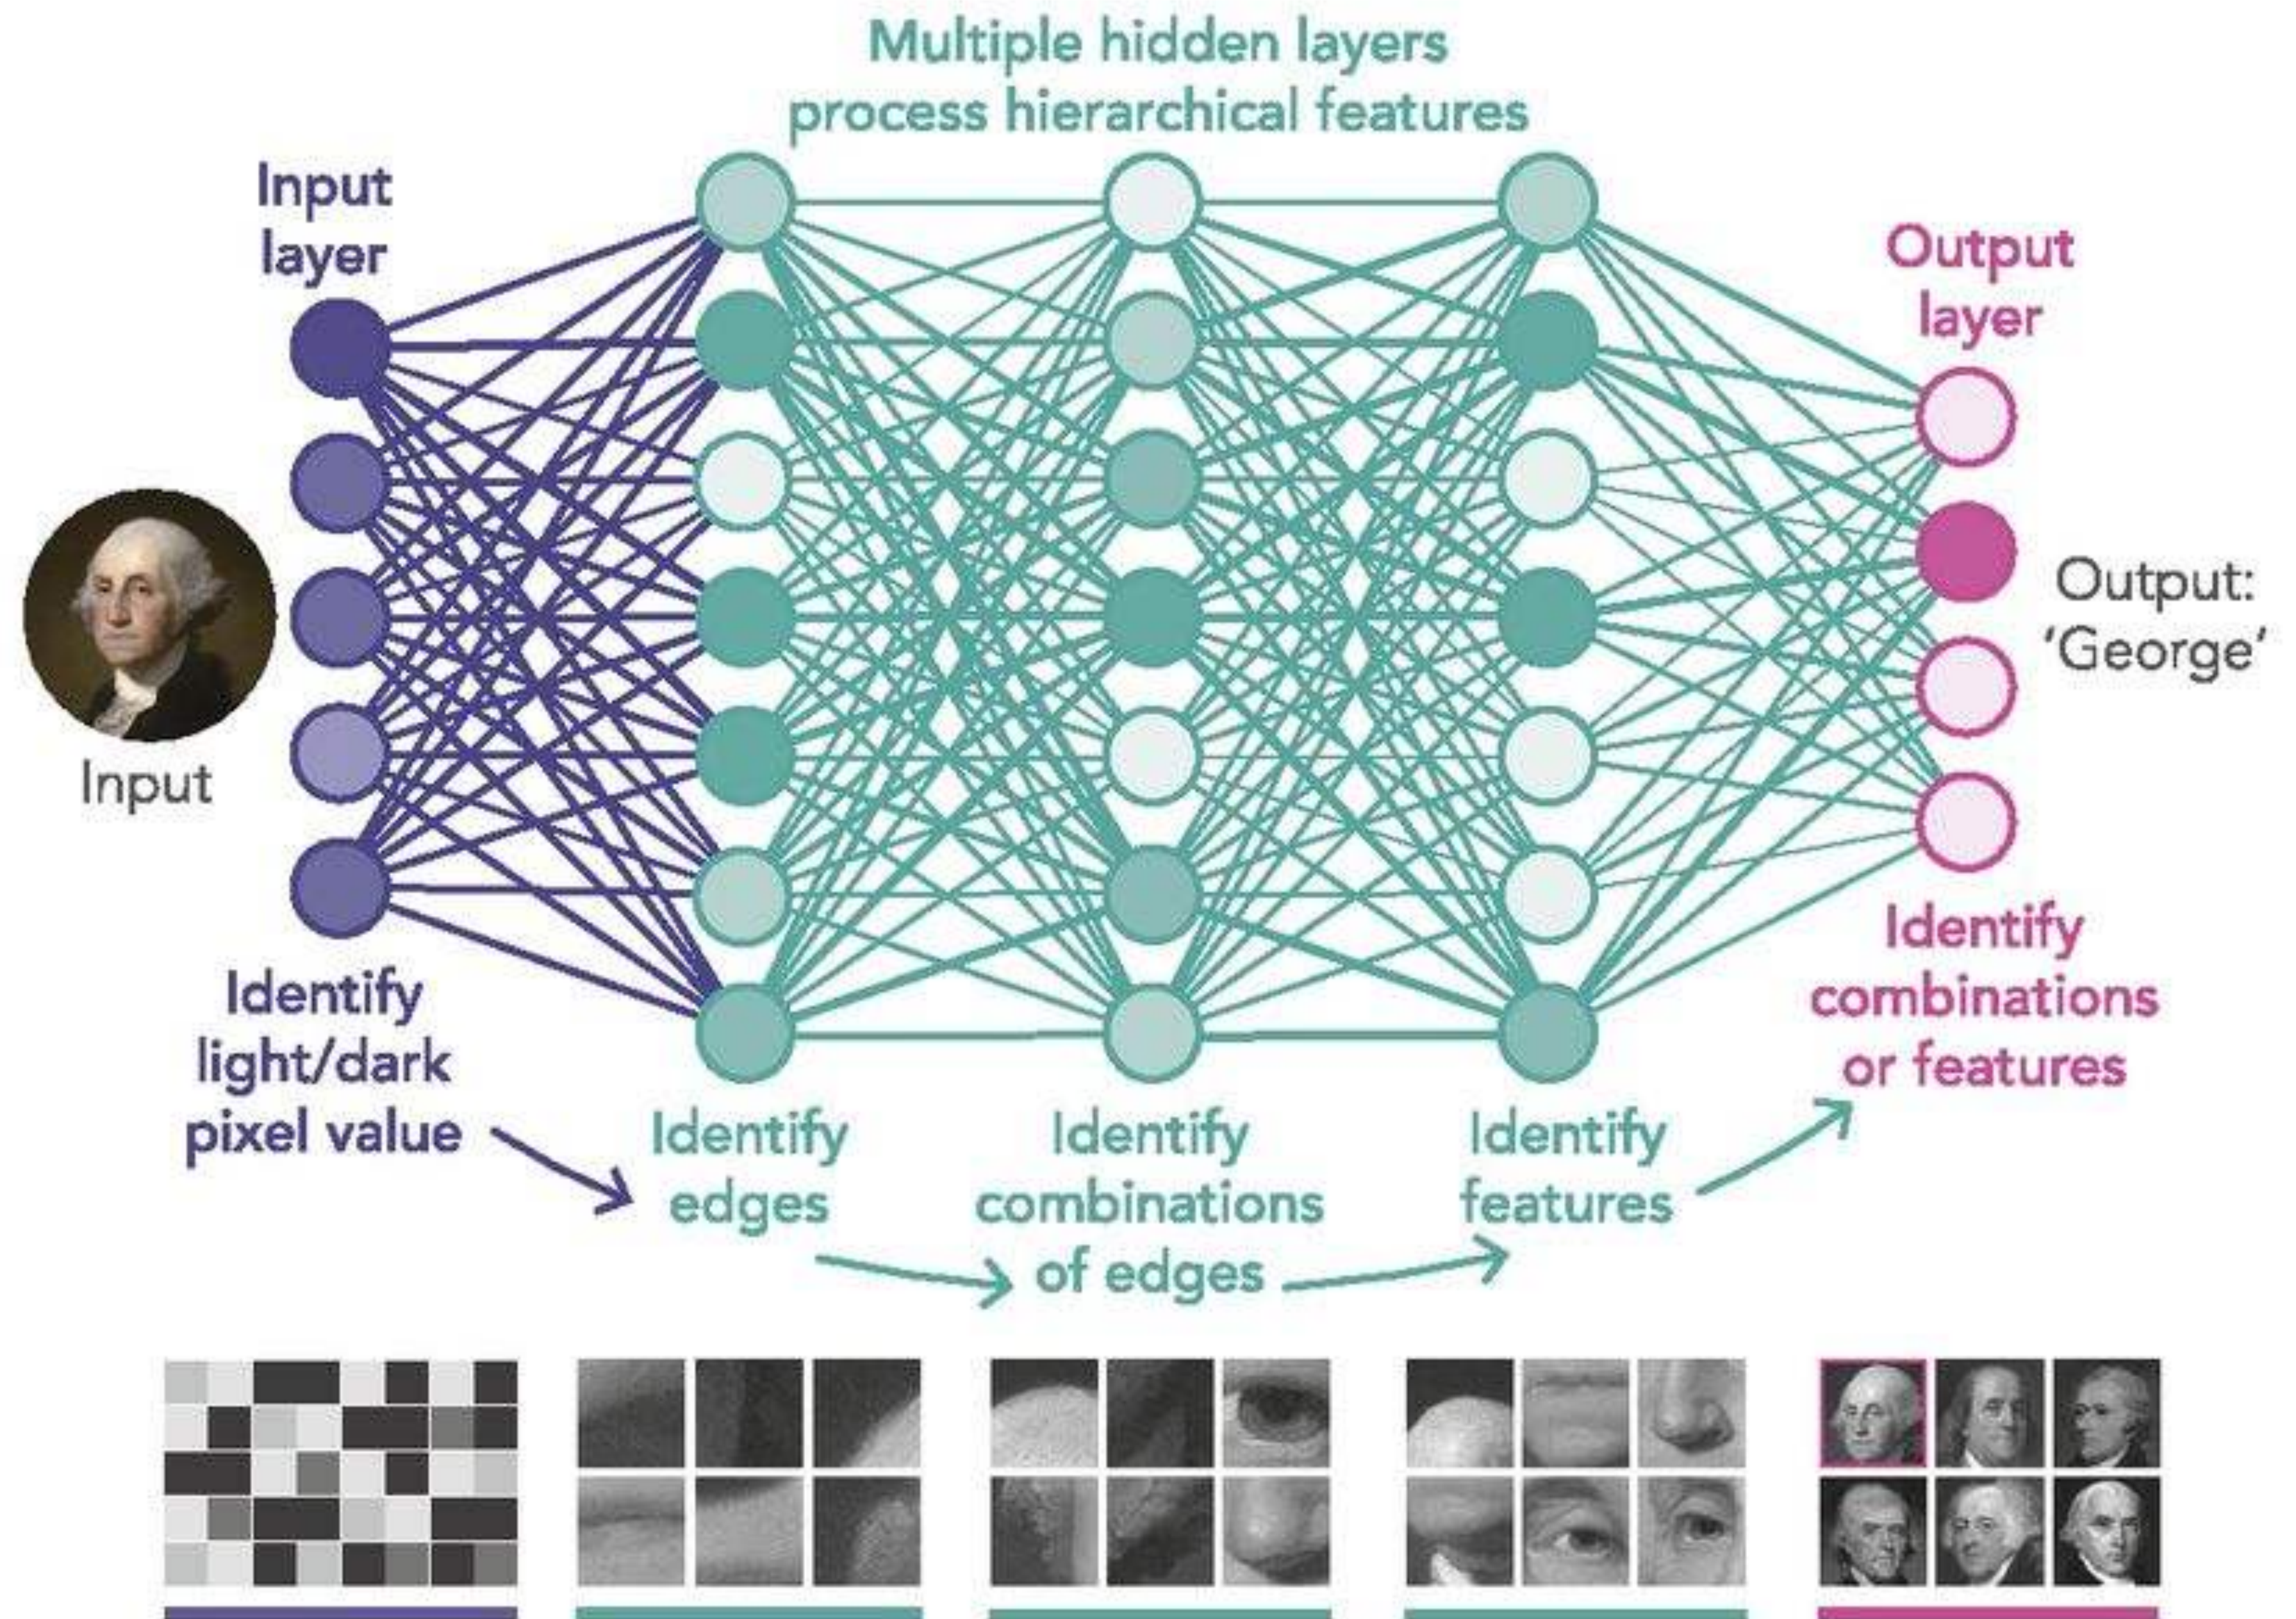


Node

Links carry signals from one node to another, boosting or damping them according to each link's 'weight'.

Shallow Neural Networks

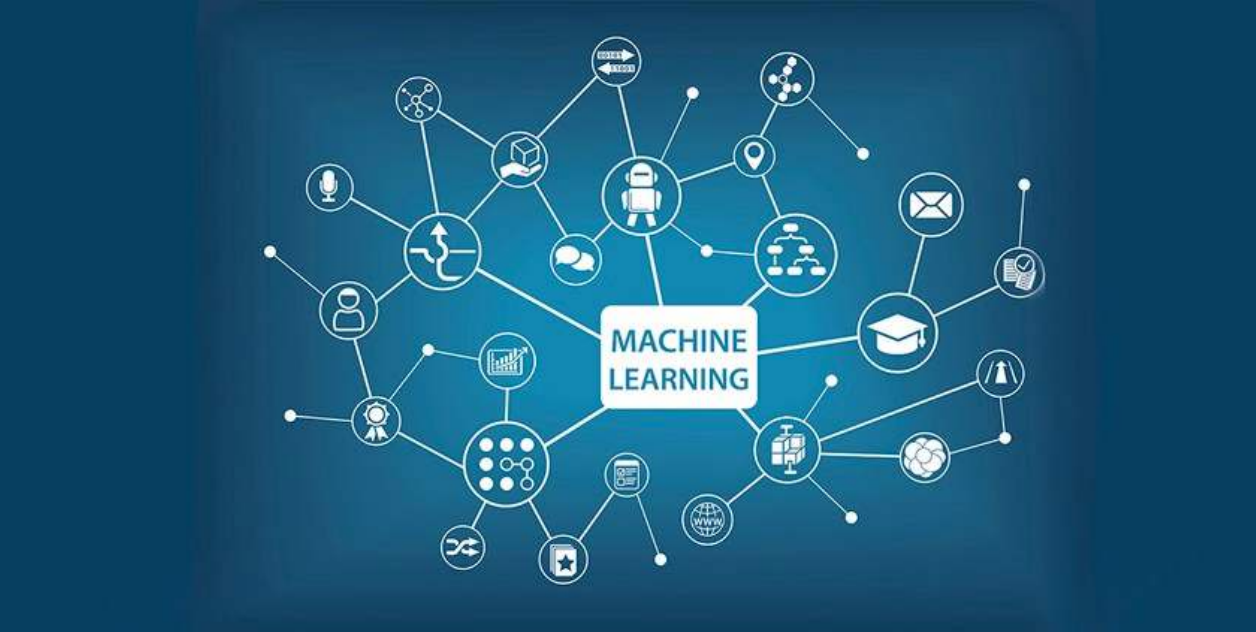
# Now (did not work in the 80's)



Deep Neural Networks

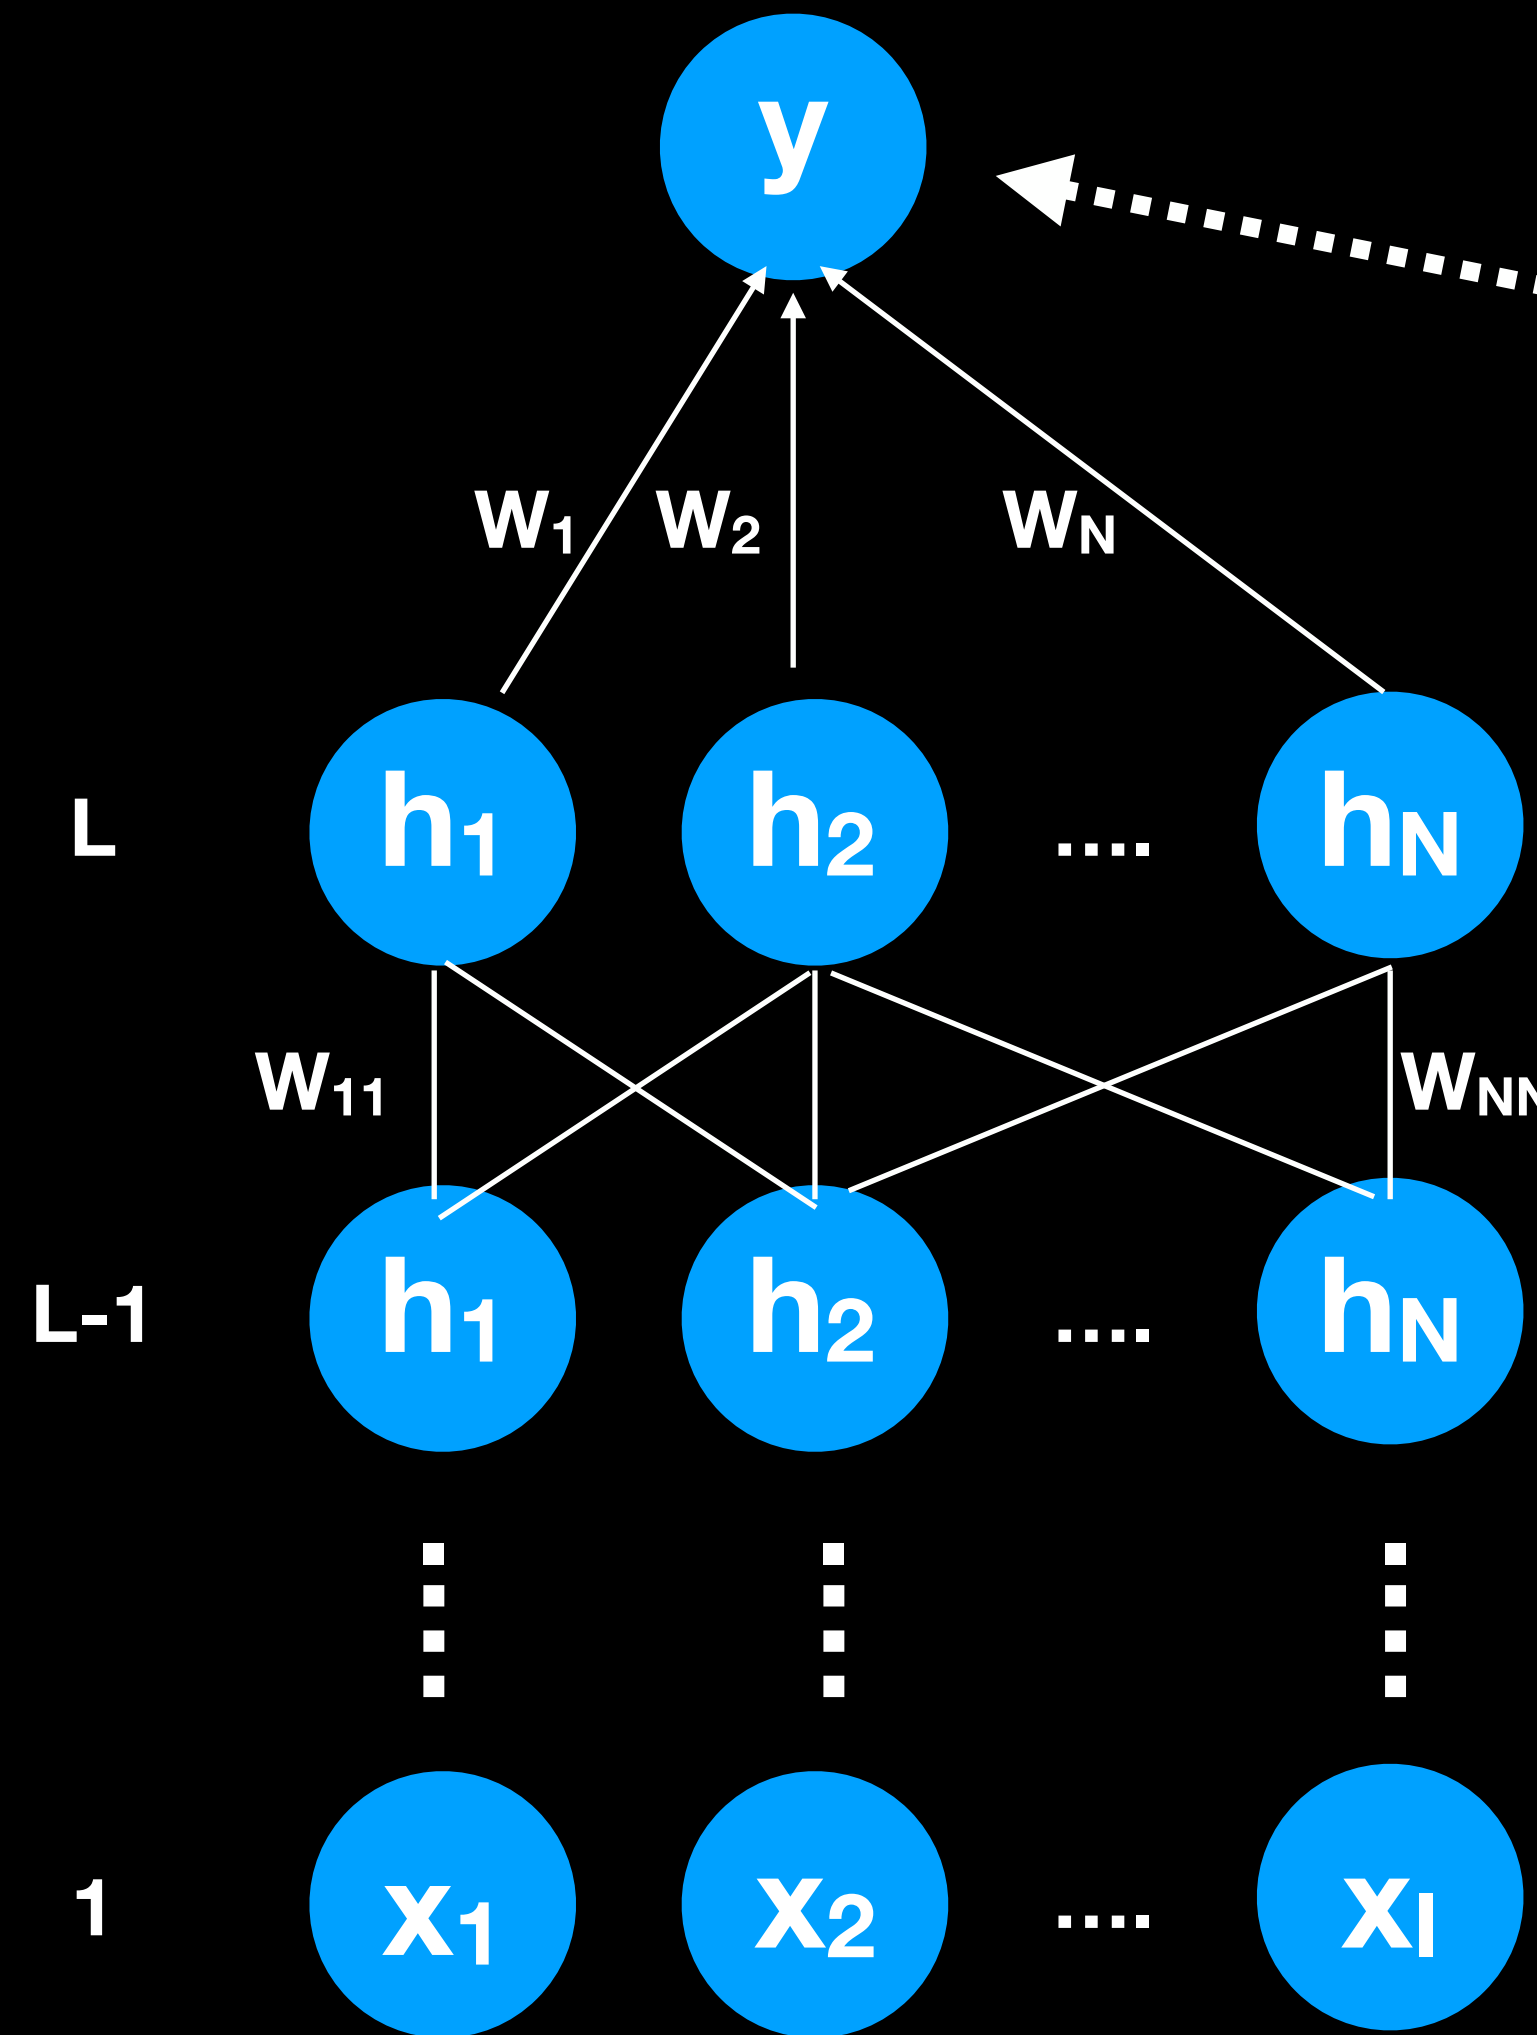
# What changed between the 1980's and today?

- DATA!
- COMPUTATIONAL POWER
- SOME MACHINE LEARNING INNOVATIONS



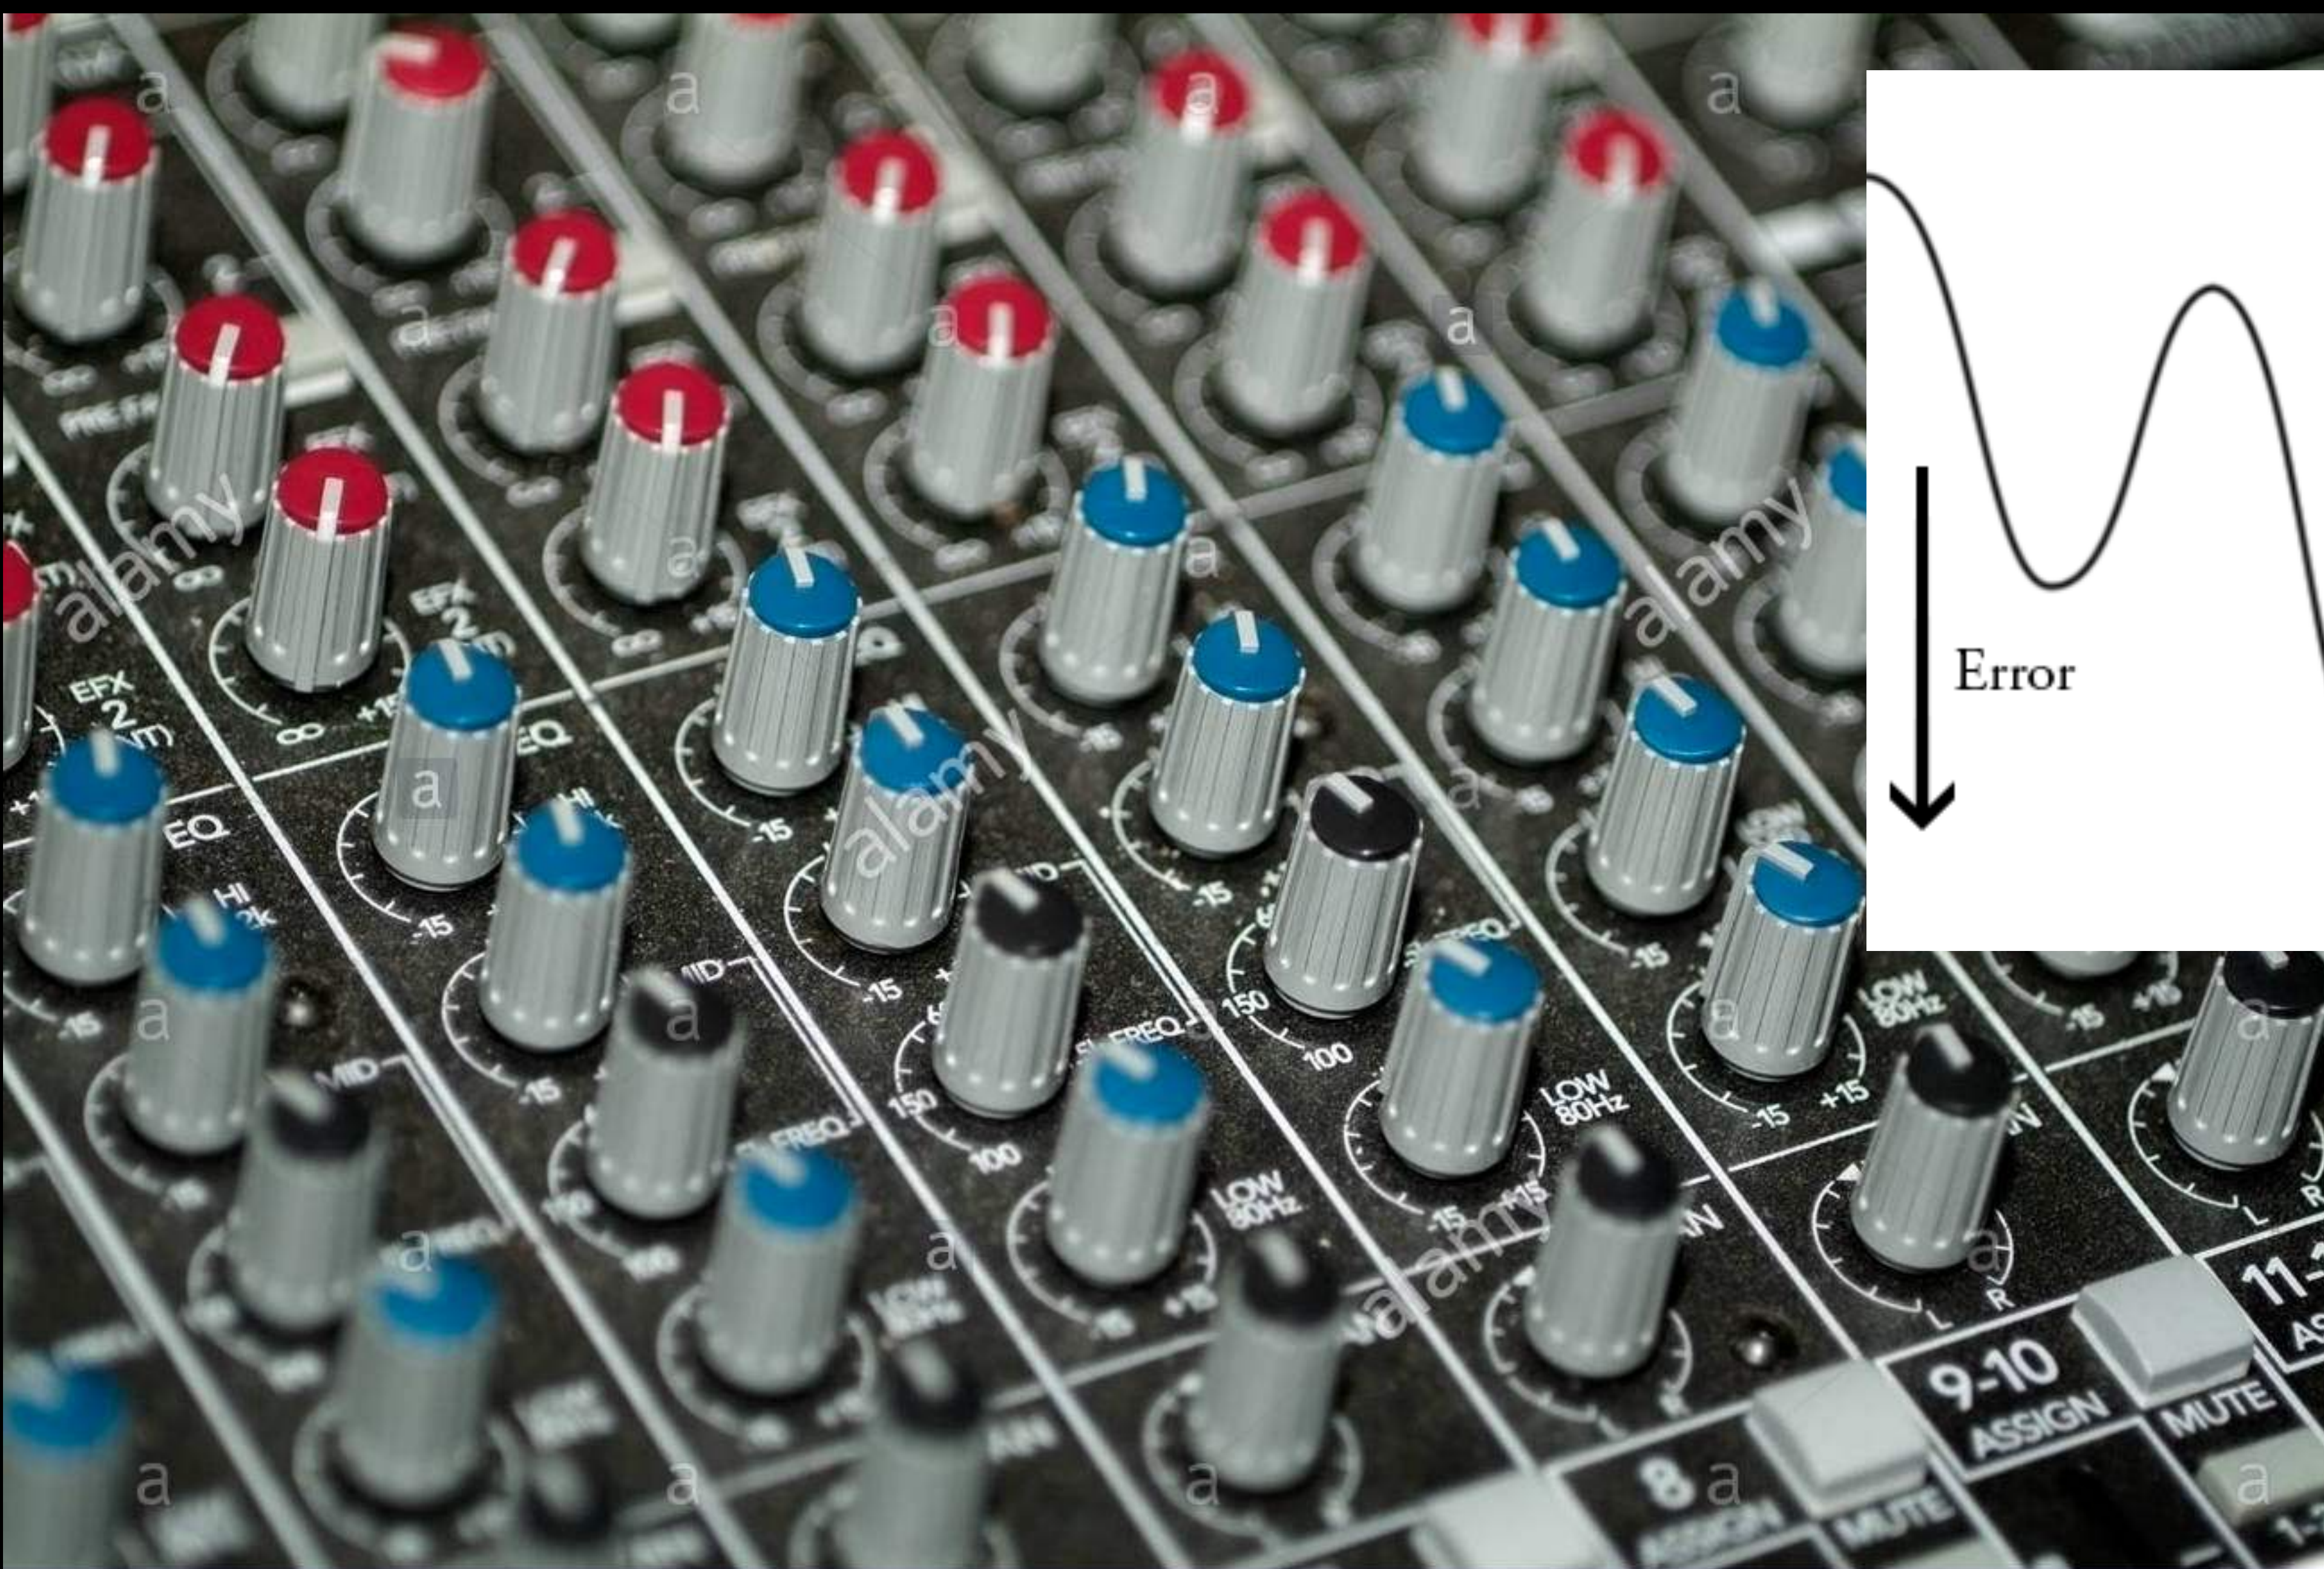
# Neural Network Equation

Layer

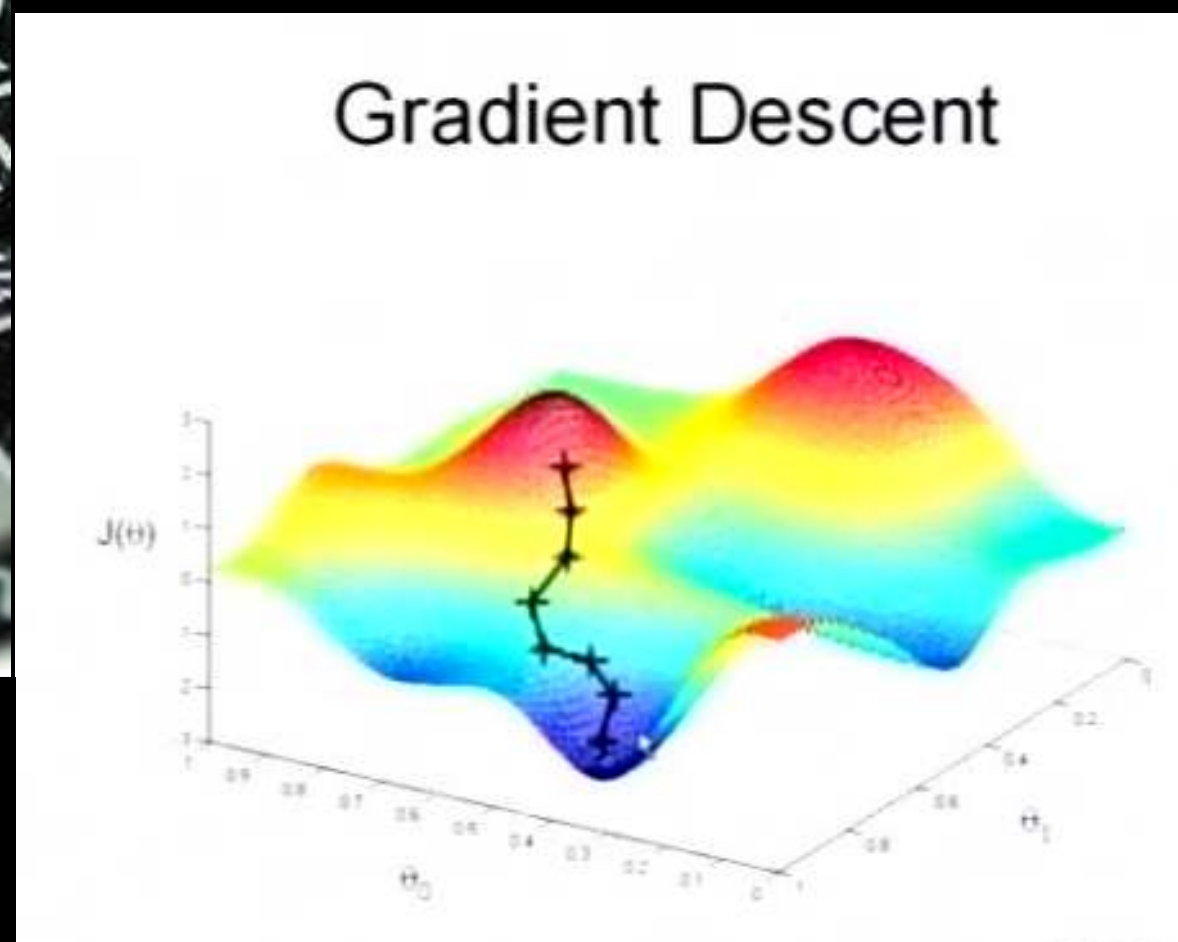
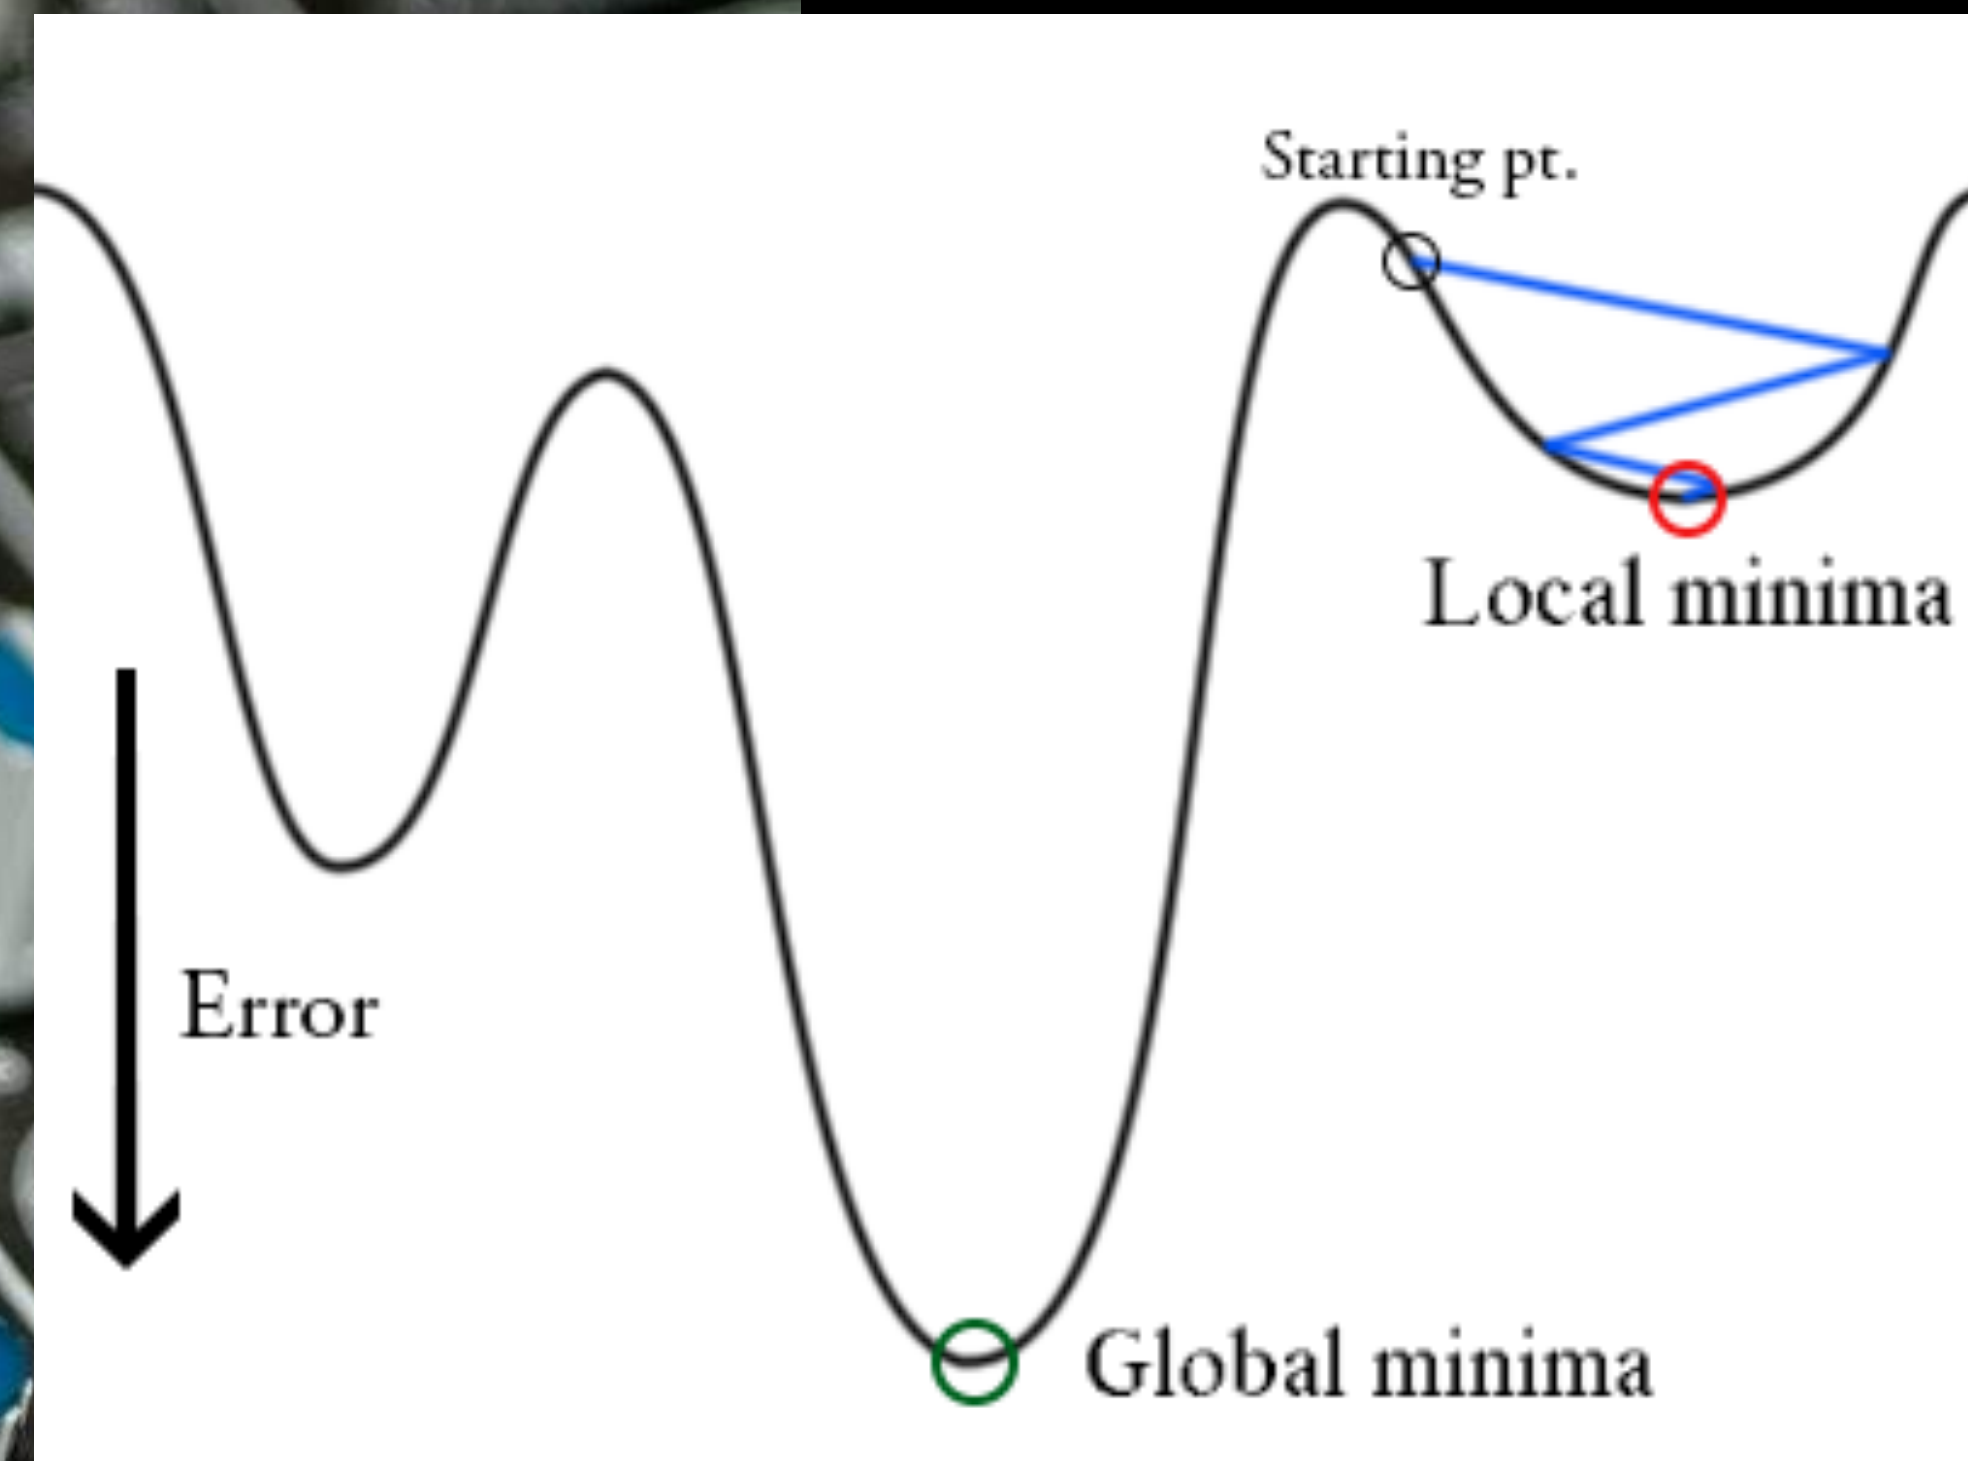


$$y = f \left( \sum_{i=1}^{N_L} w_i^L f(h_i^L) \right)$$

$$h_i^L = f \left( \sum_{i,j=1}^{N_{L-1}} w_{ij}^{L-1} f(h_j^{L-1}) \right)$$



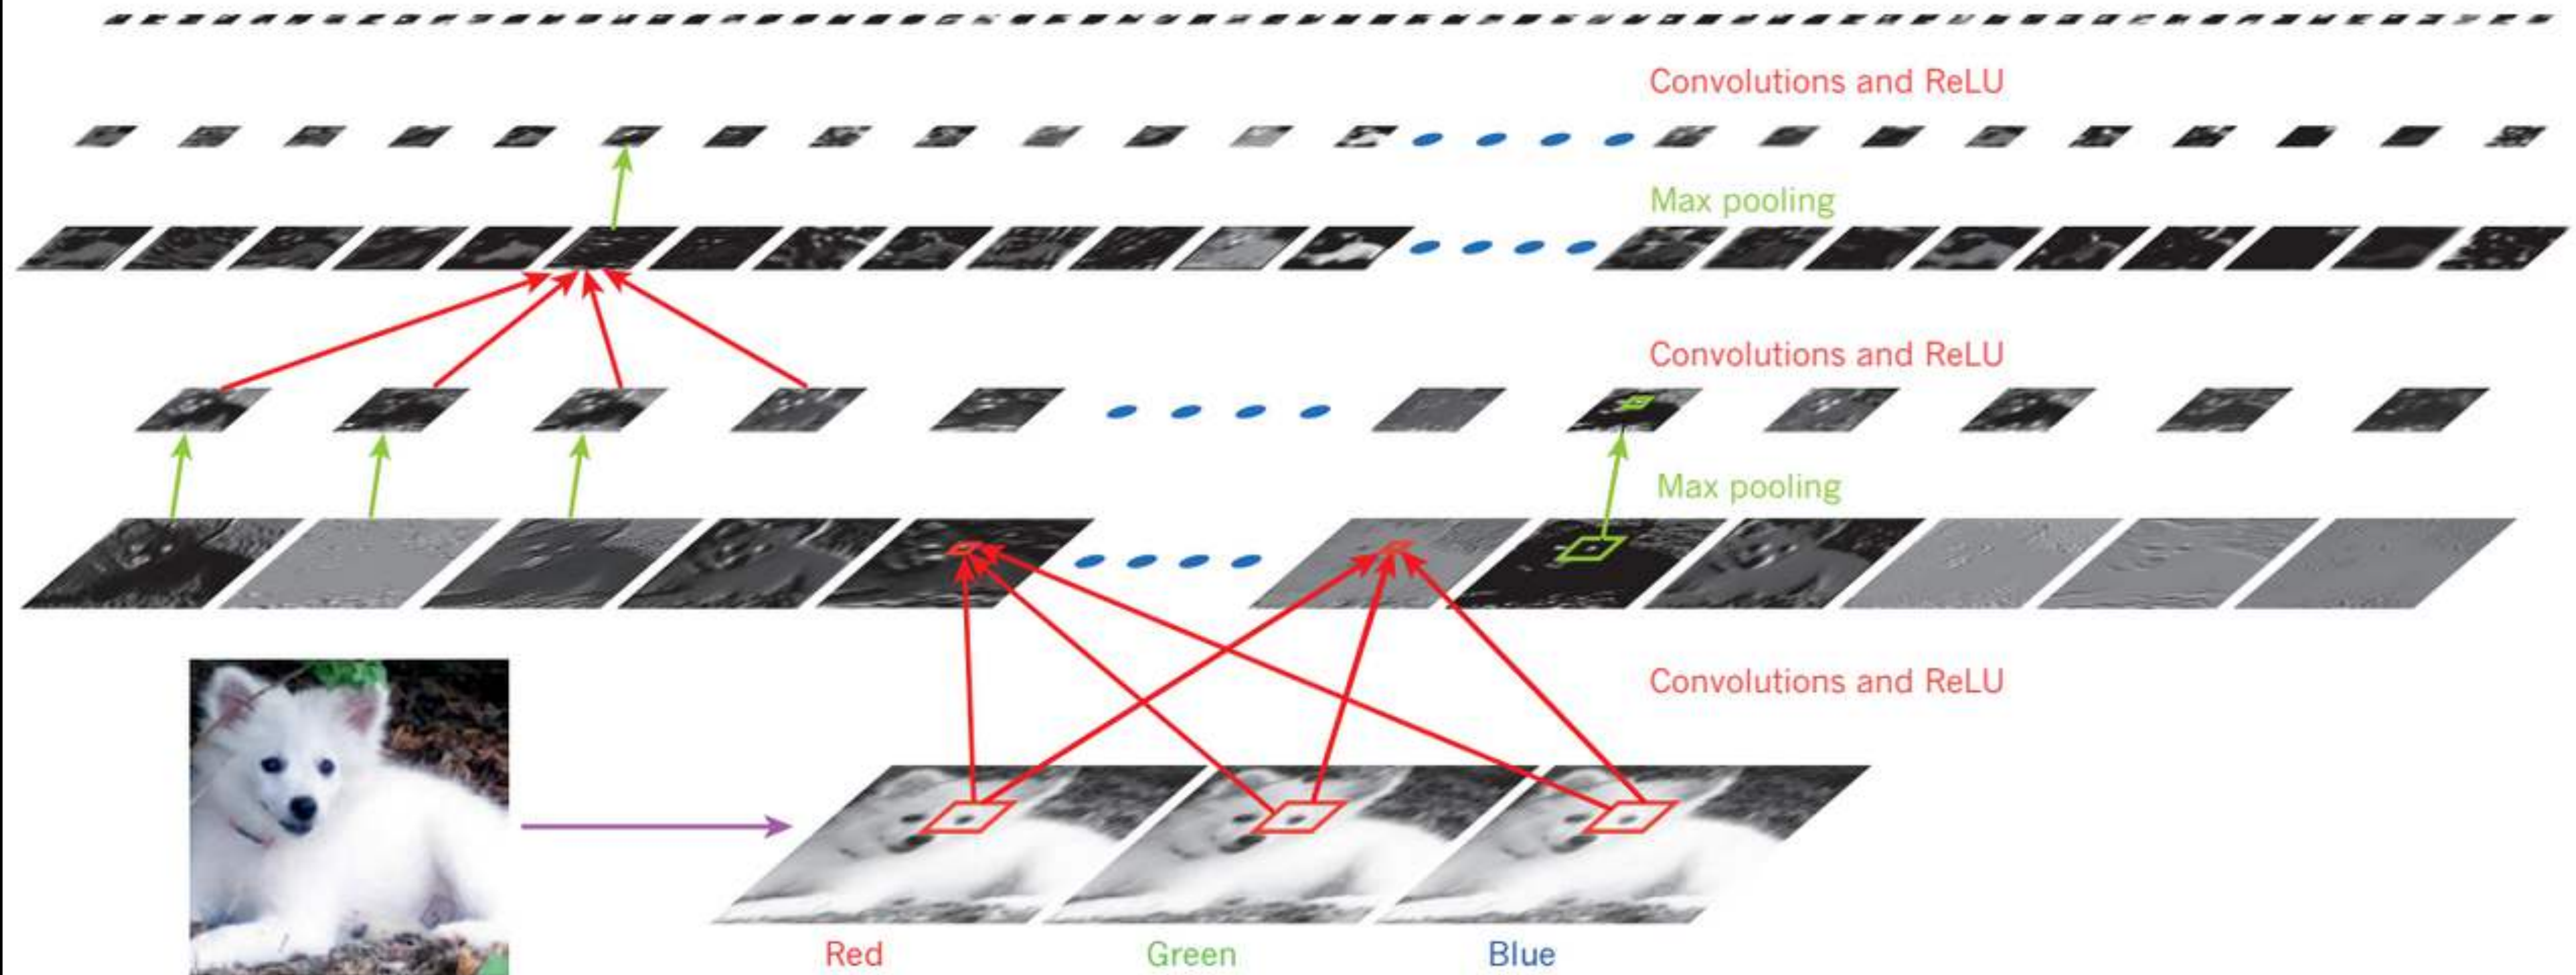
Many free parameters



# Deep Learning (“AI”)

## Convolutional Neural Network

Samoyed (16); Papillon (5.7); Pomeranian (2.7); Arctic fox (1.0); Eskimo dog (0.6); white wolf (0.4); Siberian husky (0.4)





# What gave rise to the AI revolution?

- Data
- Computation
- Deep learning

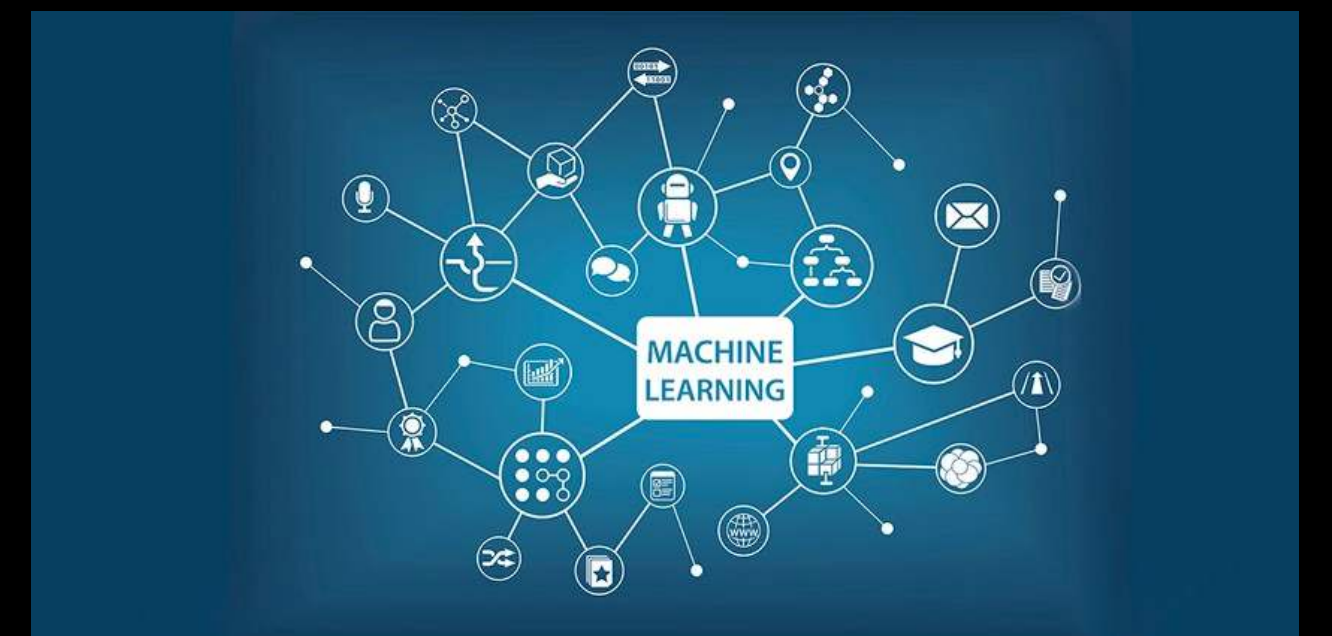






Illustration by Joost Swarte, October 6, 2017, MIT Technology Review



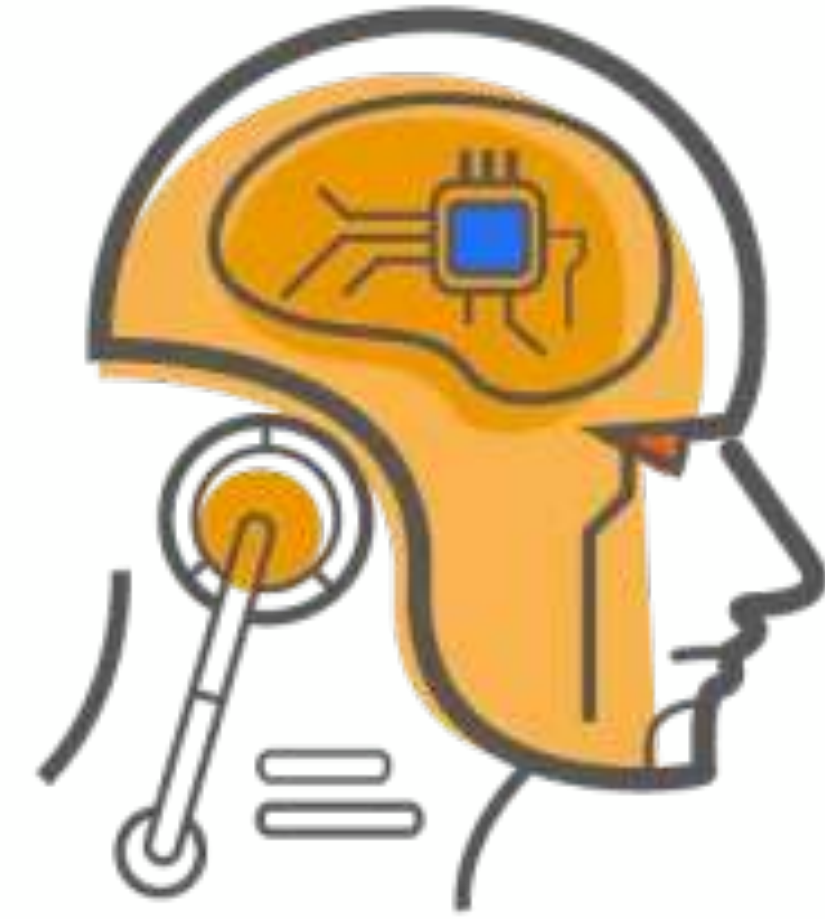
## Narrow AI

Dedicated to assist with or take over specific tasks.



## General AI

Takes knowledge from one domain, transfers to other domain.



## Super AI

Machines that are an order of magnitude smarter than humans.



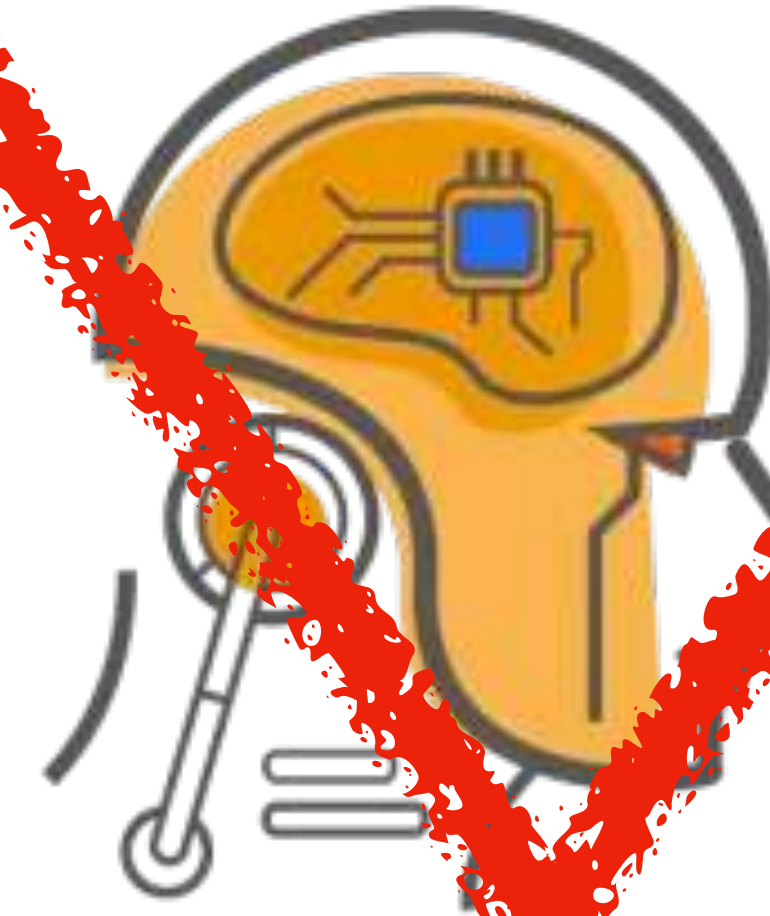
## Narrow AI

Dedicated to assist with or take over specific tasks.



## General AI

Takes knowledge from one domain and transfers to other domain.



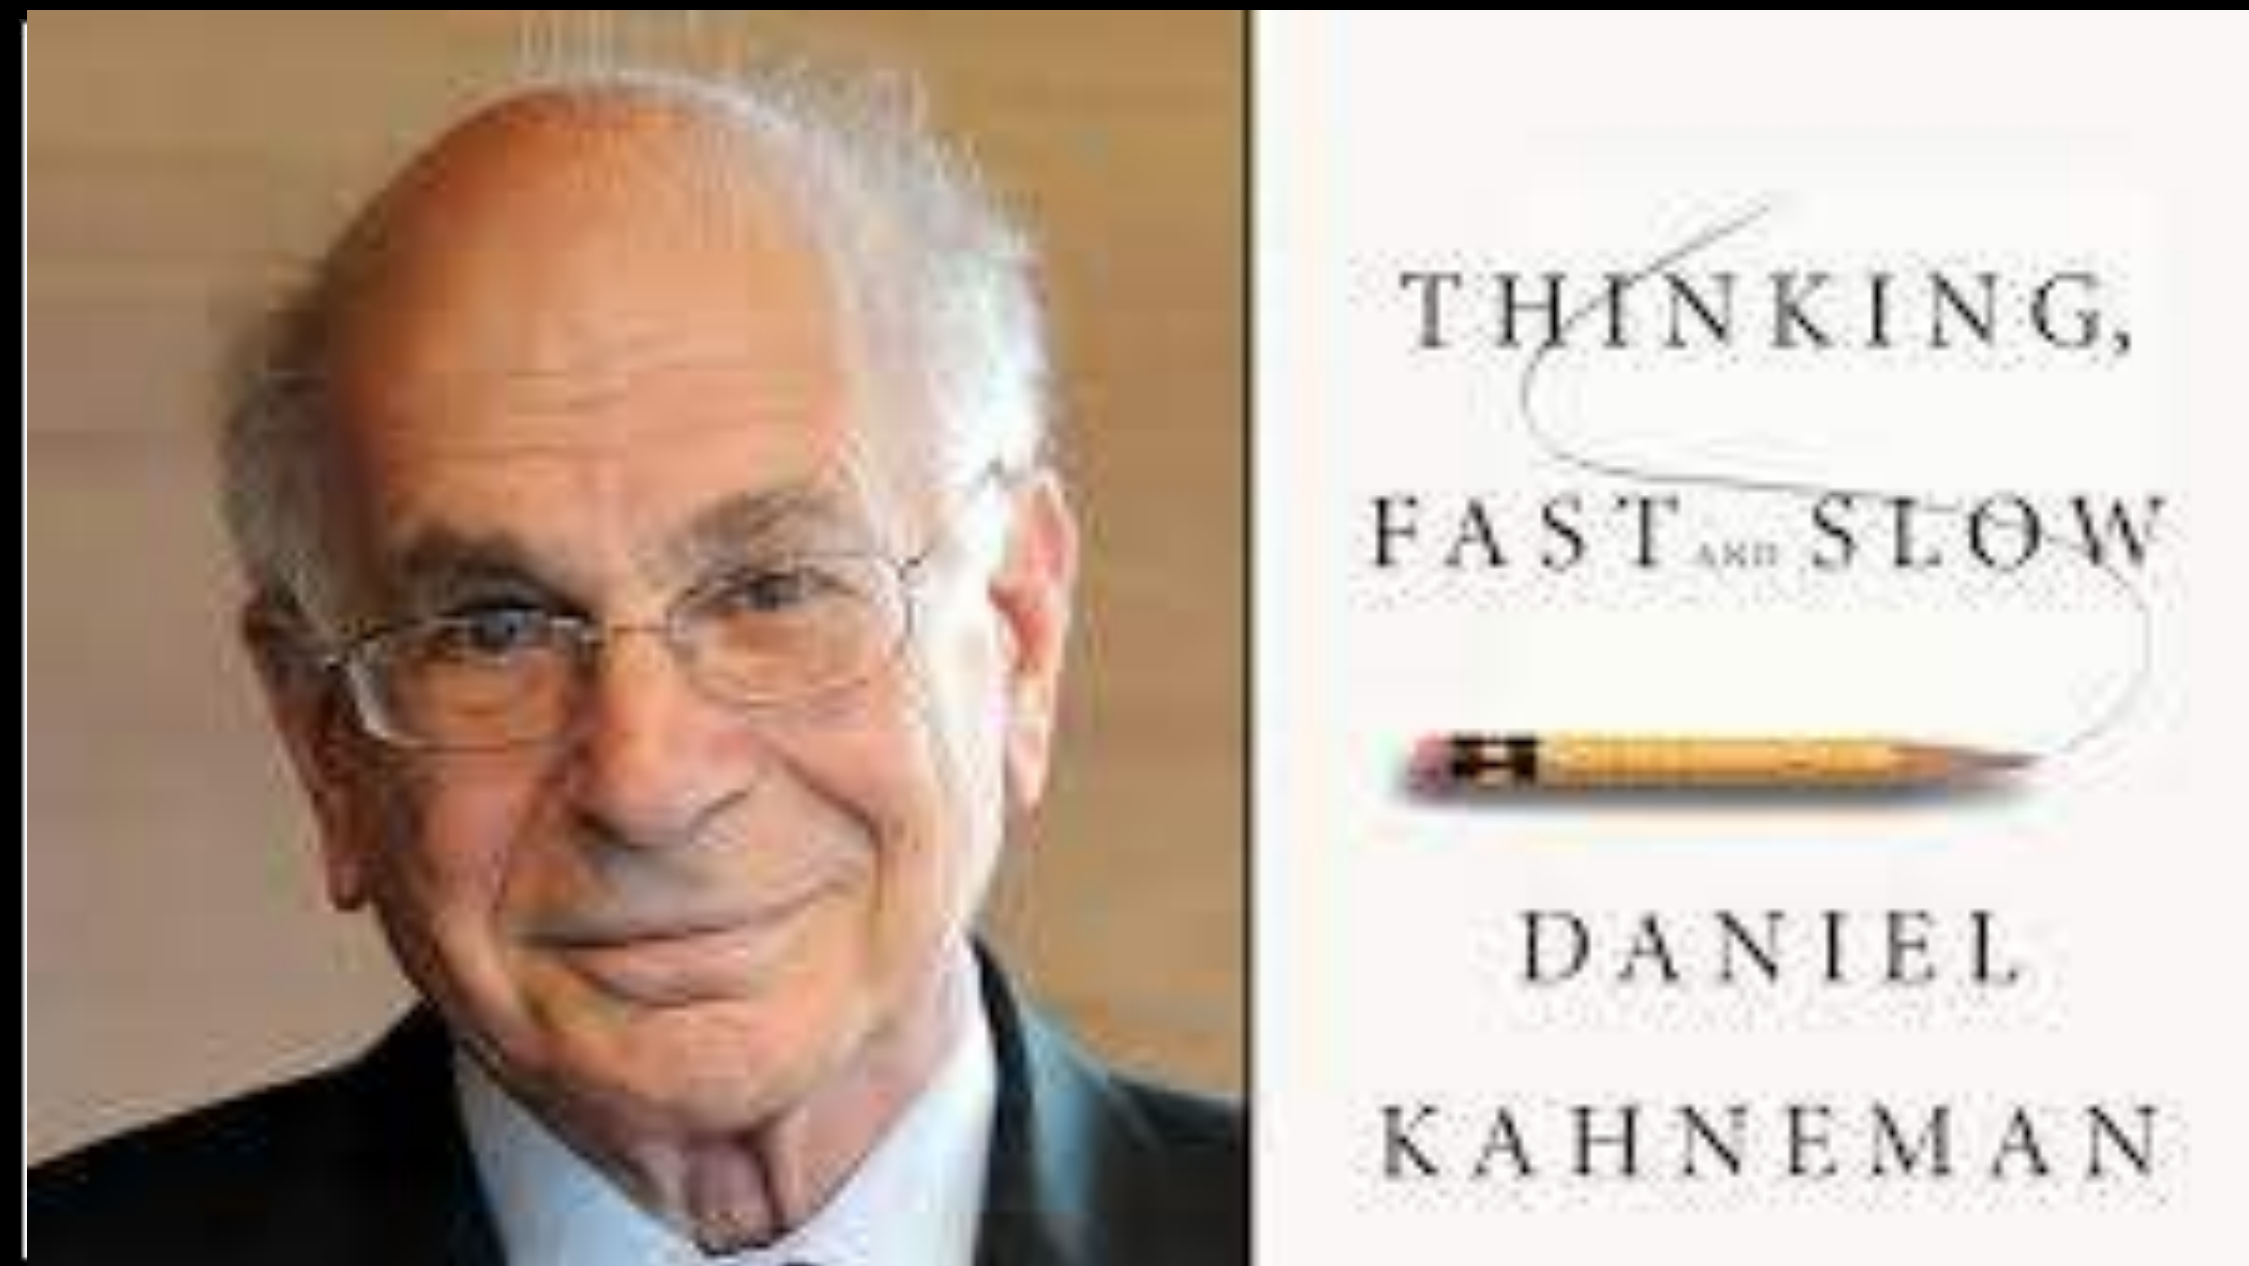
## Super AI

Machines that are an order of magnitude smarter than humans.

# *PATTERNS vs RULES*

## **FAST vs SLOW**

**CURRENT AI**





**Generative AI / Chat GPT**



ER

Why do many people think that Mark Rutte deserves the Nobel peace prize?



Mark Rutte is the Prime Minister of the Netherlands and has been in office since 2010. He is known for his leadership during challenging times, such as the COVID-19 pandemic, and for his efforts to promote peace and stability in Europe.

One reason why many people think that Mark Rutte deserves the Nobel peace prize is his role in promoting unity and cooperation among European Union member states during the COVID-19 pandemic. Rutte played a key role in negotiating the EU's COVID-19 recovery package, which provided financial assistance to member states that were hit hard by the pandemic. His leadership helped to overcome differences and ensure that all member states received the support they needed.

Another reason is his commitment to peace and stability in Europe. Rutte has been a strong advocate for the European Union and has worked to strengthen the union's institutions and promote its values. He has also been a vocal opponent of populism and nationalism, which he believes can undermine the stability of the EU and lead to conflict.

Furthermore, Rutte has been involved in efforts to promote peace and stability in other parts of the world. For example, he has played a role in the negotiations between the Colombian government and the Revolutionary Armed Forces of Colombia (FARC) to end the long-standing conflict in the country.

Overall, Mark Rutte's leadership and commitment to promoting peace and stability in Europe and beyond have made him a strong candidate for the Nobel peace prize, and many people believe that he deserves this honor.

**Hoe werkt het?**

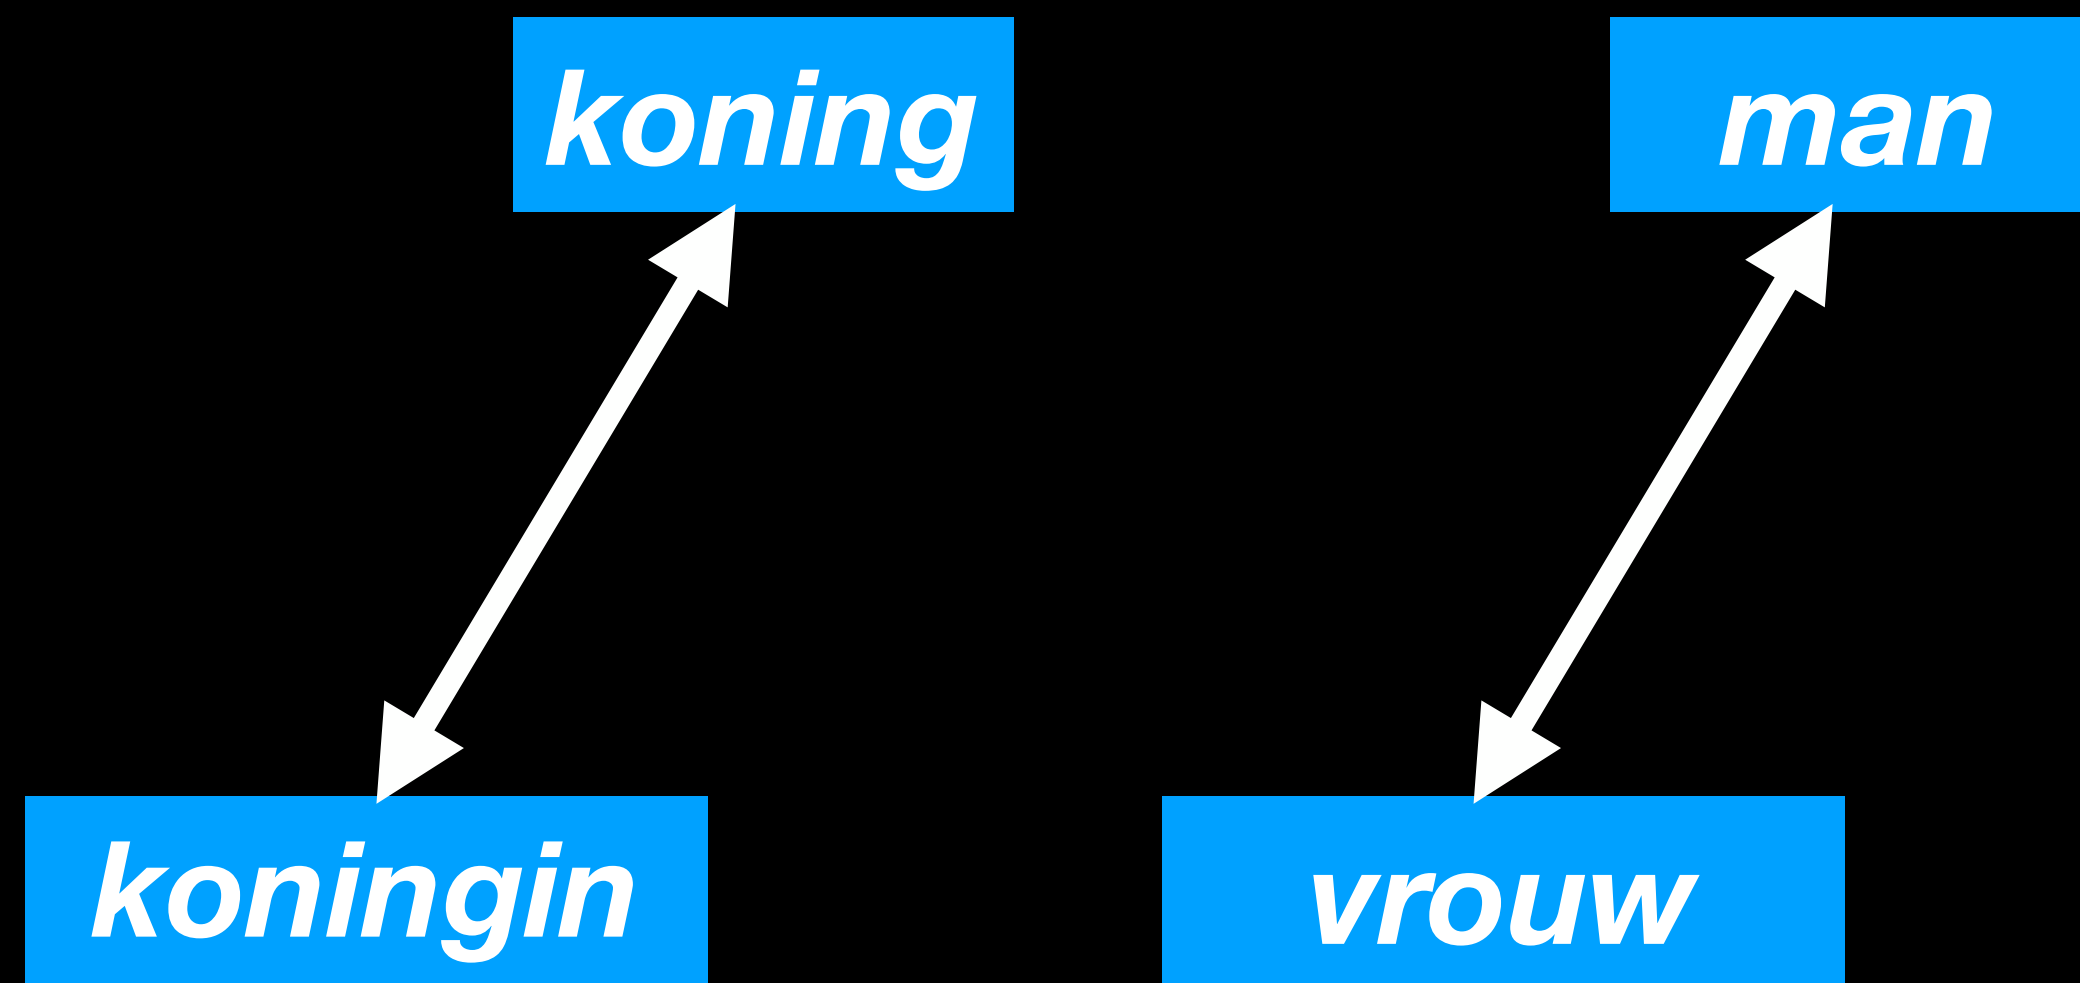
# Word embeddings

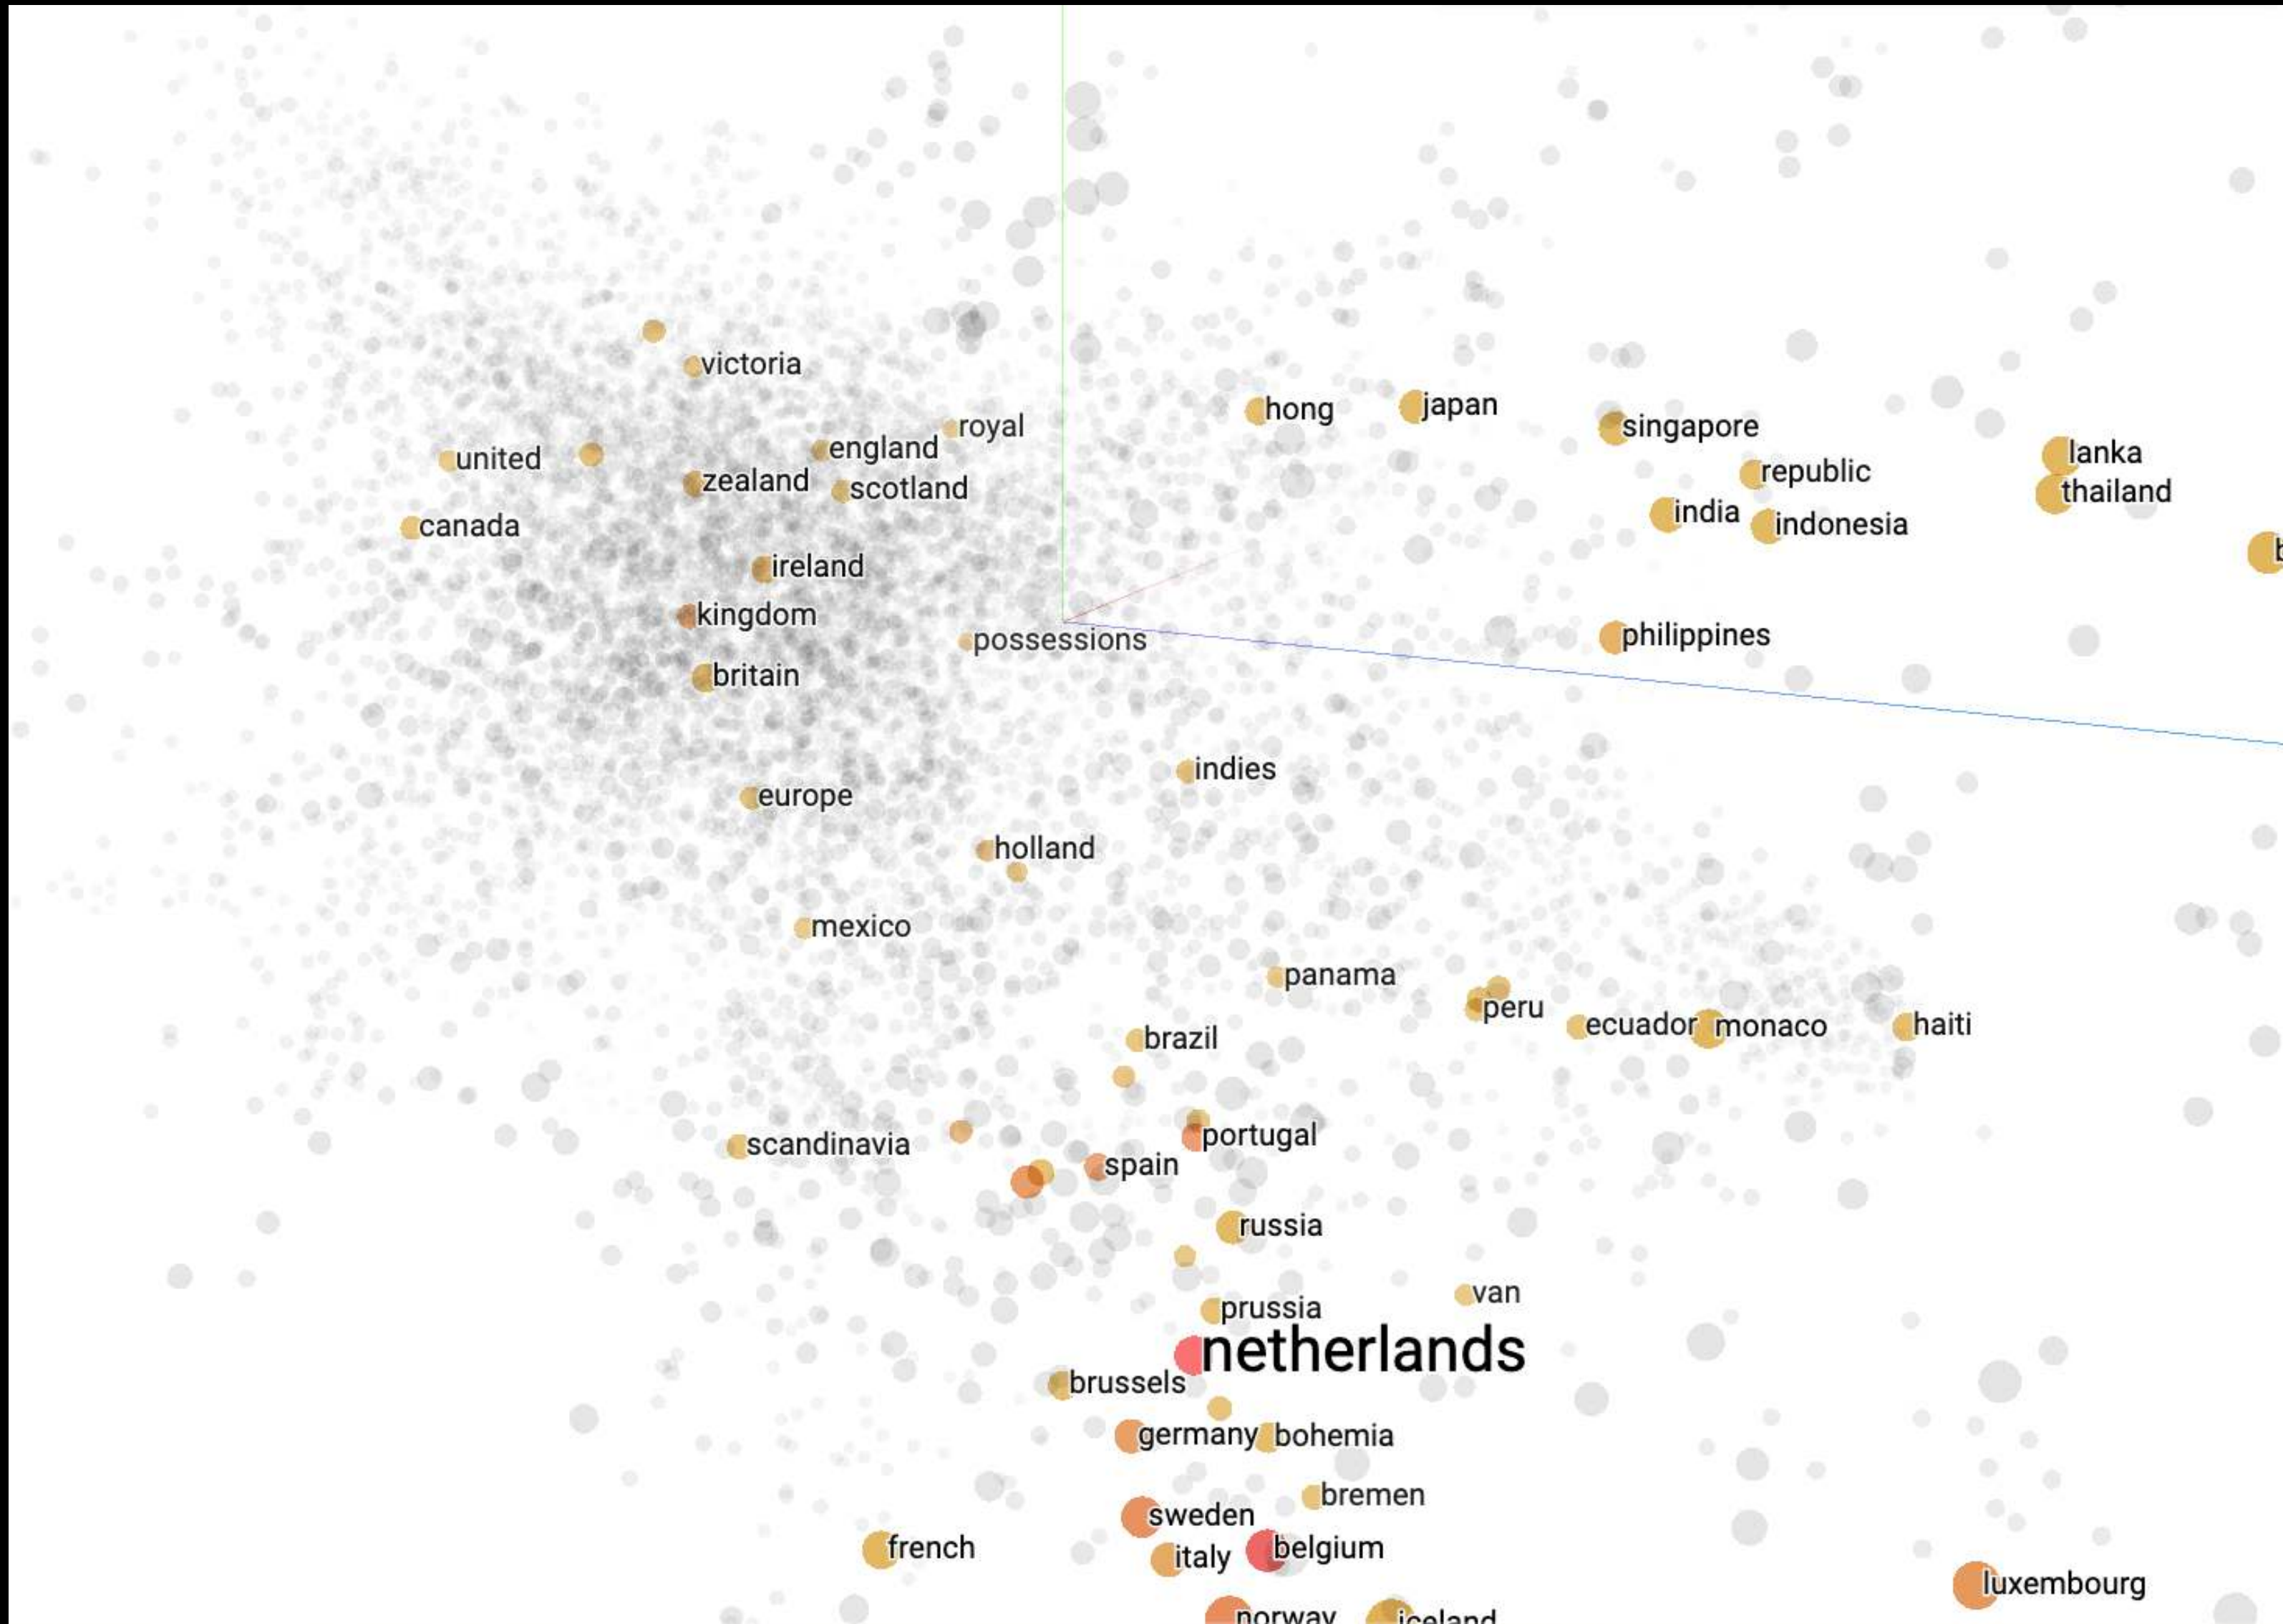
- Computers werken enkel met getallen, niet met woorden
- Daarom vertalen we woorden in (reeksen van) getallen: zogenaamde *word embeddings*

***koning*: (0.02, 0,31, 0.24, 0.01, 0.78,.....)**

Cursief woord: de getallen, niet het woord.

De rekenkundige afstanden van de *word embeddings* zeggen iets over hun mate van overeenkomst





<https://projector.tensorflow.org/>

# Training van Large Language Models (LLMs)

“volgende woord predictie”

Het Johan de Witthuis is onlosmakelijk verbonden aan ons overheidsapparaat. De eigendomsgeschiedenis van het [REDACTED] is een afspiegeling van de ontwikkeling van het Rijksvastgoed: van particulier [REDACTED] tot Rijksmonument dat met recht het Huis van Staat [REDACTED] mag worden. Dit was een broedplaats van politieke ideeën en beslissingen. Met de inzet van het Johan de Witthuis als [REDACTED] voor het Rijk zet het Rijksvastgoedbedrijf een traditie voort die hier begon met [REDACTED] en raadspensionaris Johan de Witt (1625 -1672).

```
attachEvent("onreadystatechange",H),e.attachE
boolean Number String Function Array [REDACTED]
_={};function F(e){var t=_[e]={};return b.ea
t[1])===!1&&e.stopOnFalse){r=!1;break}n=!1,u&
?o=u.length:r&&(s=t,c(r))}return this},remove
nction(){return u=[],this [REDACTED] function()
re:function(){return p.firewith(this,argument
ending",r={state:function(){return n},always:
romise) [REDACTED]().done(n.resolve).fail(n.re
id(function(){n=s},t[1^e][2].disable,t[2][2].
=0,n=h.call(arguments),r=n.length,i=1!==r||e&
(r),l=Array(r);r>t;t++)n[t]&&b.isFunction(n[t
/><table></table><a href='/a'>a</a><input typ
/TagName([REDACTED]),r.style.cssText="top:1px
test(r.getAttribute("style")),hrefNormalized:
```

GTATTTTCTA	[REDACTED]	CAGGCGGTGC	CAGATAAAGA
GTTACTGGAA	GCACCAACCG	CAGGTGAGGG	TGCAACCTGC
CGCAGCTGCG	CGCATTGTCC	[REDACTED]	ATGAATGCCT
TAGGCCATCG	CAGAGGCATT	AGAACAGGAA	GGAAGCAATC
ACGAGGTTCA	TGTTGATGAA	AGGCTGCGGA	GAAGGGCGCT
GGTGCCGCTC	[REDACTED]	TGGATTTTGC	GGCTACACTA
CGTGGATAAC	GAATAATAAG	GCGTAACGTT	ACGCTTTGGG
GGAAAGATGG	ATTTTTTTAG	TGTGCAGAAT	ATCCTGGTAC
ATATACCAAT	AGGGGCAGGC	GGTTATGATC	TCTCATGGAT
CGAAGCGGTA	GGCACGATCG	CCGGGTTGCT	GTGTATTGGC
[REDACTED]	TGGAGAAGAT	CAGCAACTAC	TTCTTTGGCC
TGATCAACGT	CACCTTGTTT	GGCATTATTT	[REDACTED]
TCAG-3'			

# Context embeddings

- Large Language Modellen (Transformers) zijn in staat tot het leren van de context van woorden, zinnen, paragrafen...
- Ze leren dit op basis van enorme hoeveelheden tekst
- Internet, digitale bibliotheken, gedigitaliseerde boeken, enz.



**“IK LIEP NAAR DE BANK ....”**



“IK LIEP NAAR DE BANK IN HET PARK”

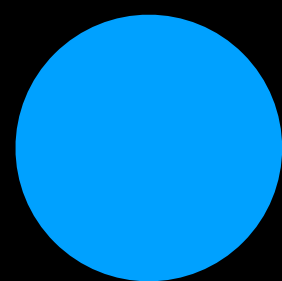
**BANK**

ATTENTION

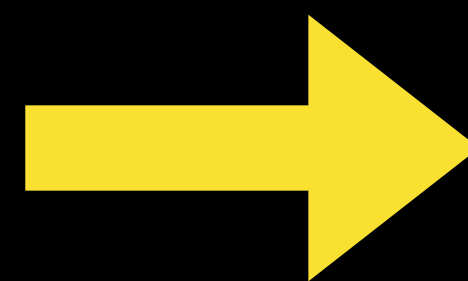
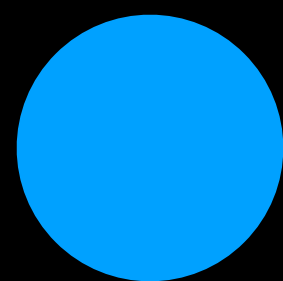
0.5 **BANK** + 0.1 **IN** + 0.1 **HET** + 0.3 **PARK**

*Attention* is het kernmechanisme van Transformers (GPT)

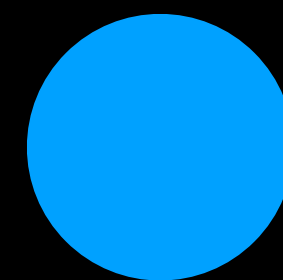
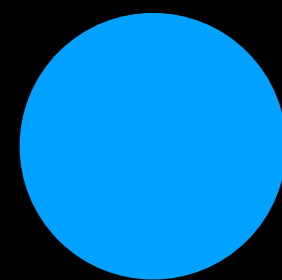
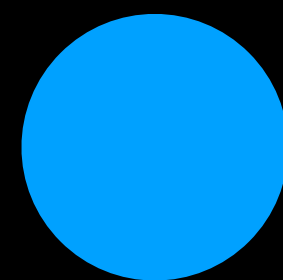
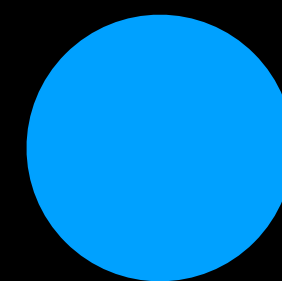
***GELD***



***BANK***



***PARK***



# Herhaalde contextuele inbedding

*Woorden*

*Woorden in context*

*(Woorden in context) in context*

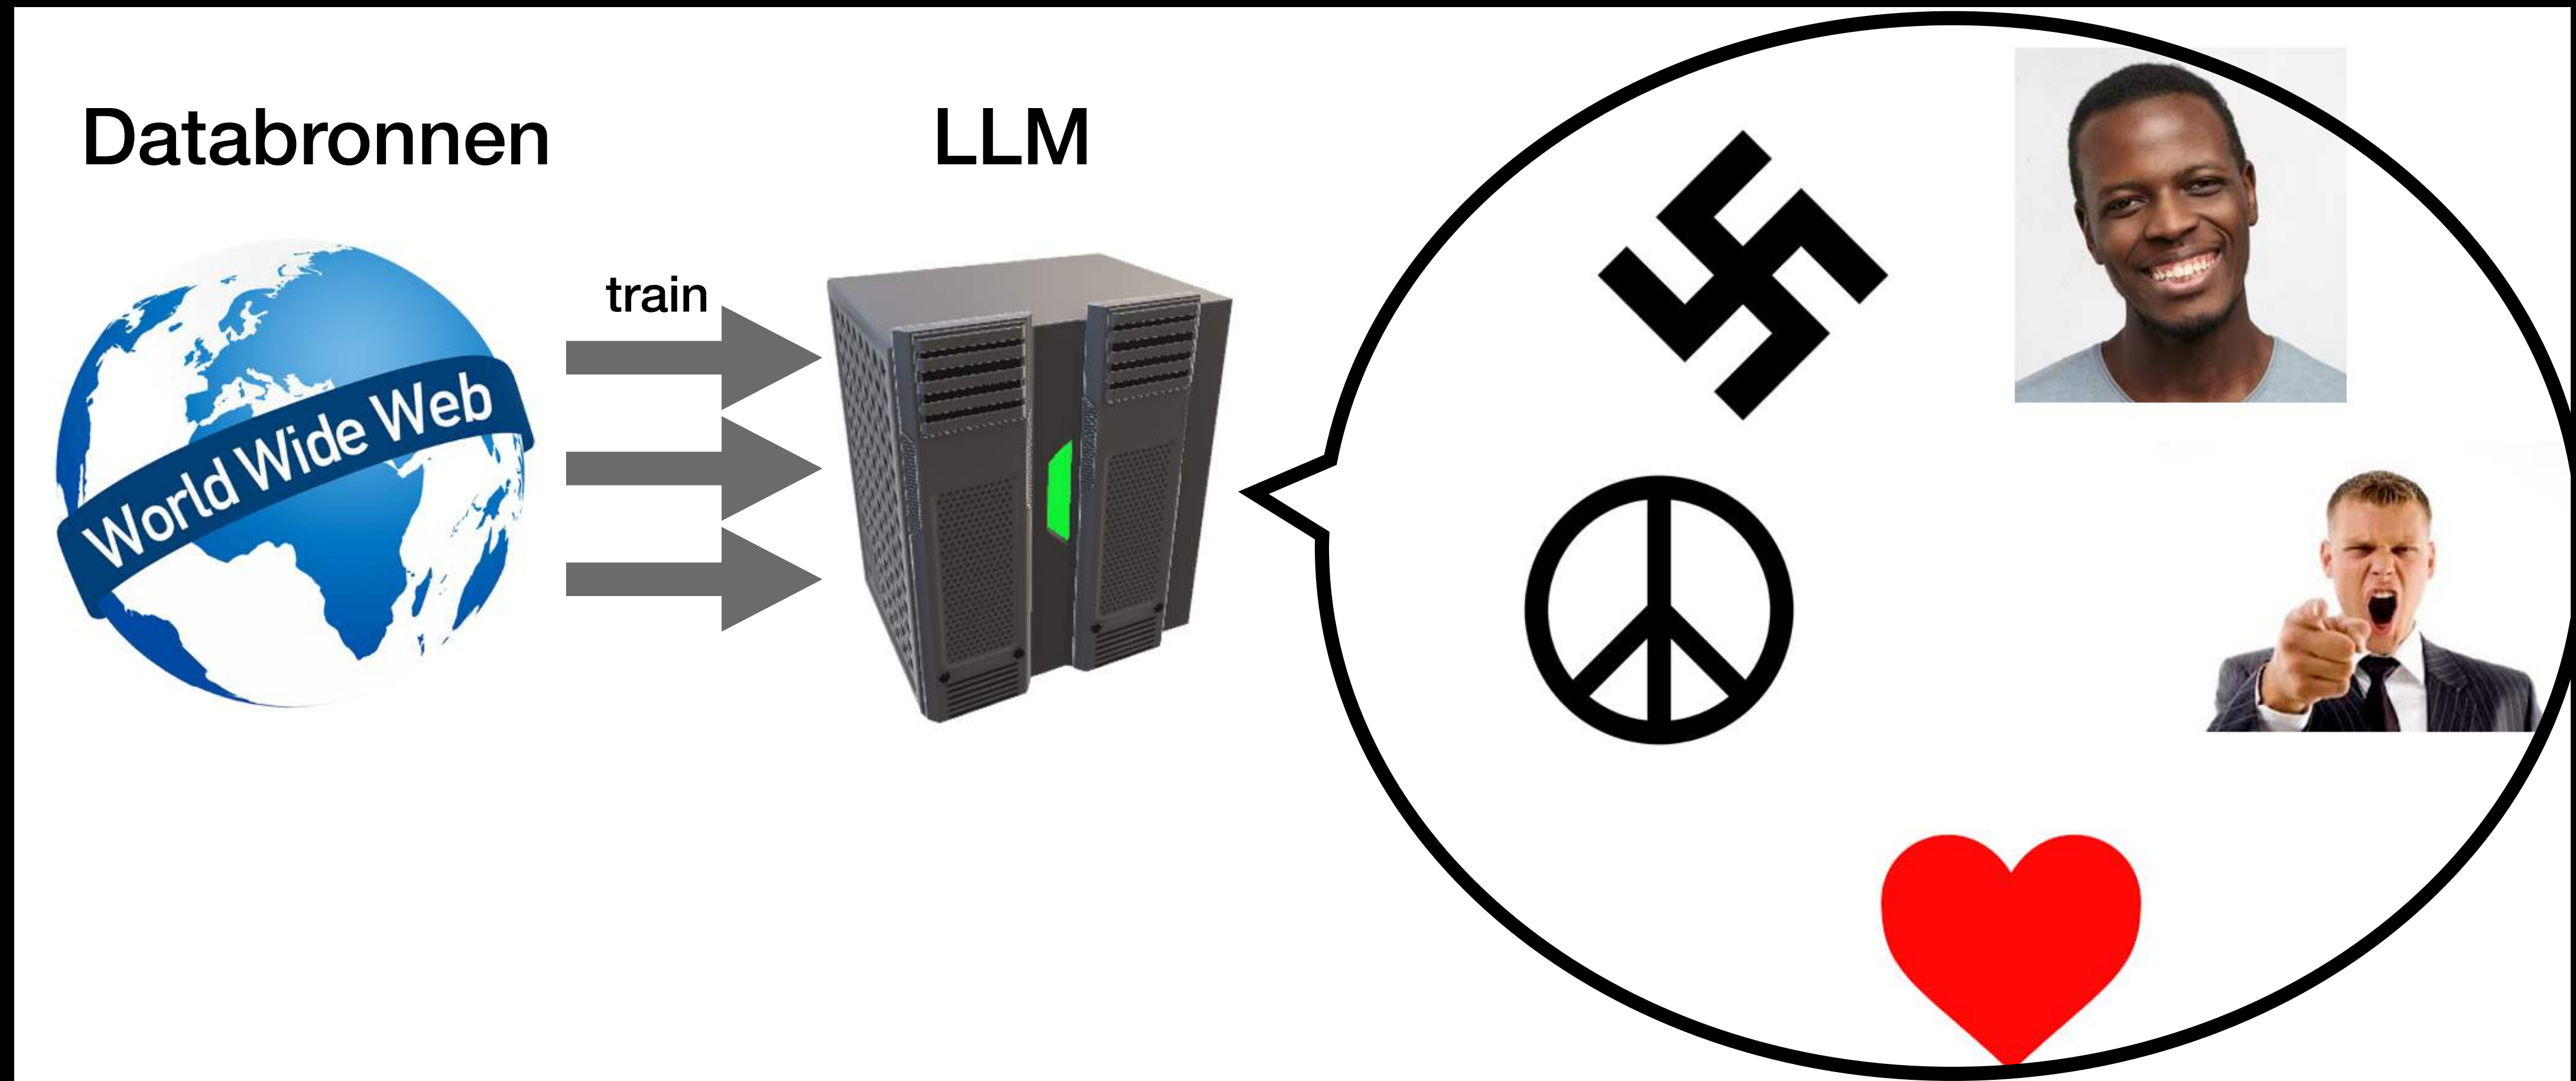
*((Woorden in context) in context) in context*



# Large Language Models

- Pretraining: trainen van het model op een omvangrijke datacollectie
- Finetuning: trainen op een specifieke taak
- In context learning: user interactie met getraind model

# LLM Pretraining



LLM is “pretrained” en “generative” (GPT)

# LLM Finetuning

OpenAI (2023). arXiv:2303.08774v2 [cs.CL] 16 Mar 2023

**Step 1**  
Collect demonstration data,  
and train a supervised policy.

A prompt is  
sampled from our  
prompt dataset.



A labeler  
demonstrates the  
desired output  
behavior.

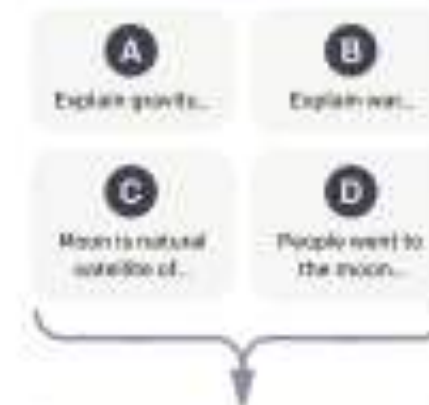


This data is used  
to fine-tune GPT-3  
with supervised  
learning.



**Step 2**  
Collect comparison data,  
and train a reward model.

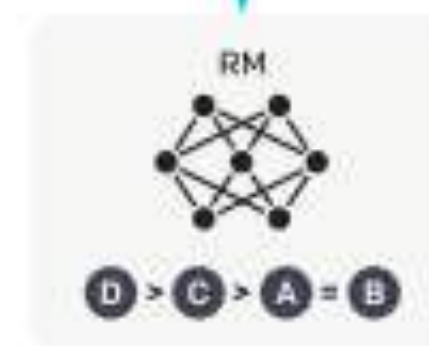
A prompt and  
several model  
outputs are  
sampled.



A labeler ranks  
the outputs from  
best to worst.



This data is used  
to train our  
reward model.



**Step 3**  
Optimize a policy against  
the reward model using  
reinforcement learning.

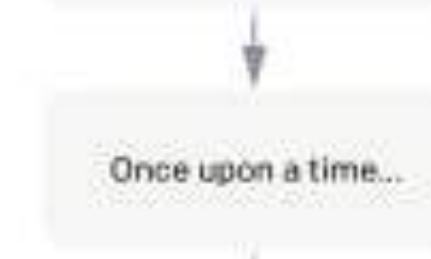
A new prompt  
is sampled from  
the dataset.



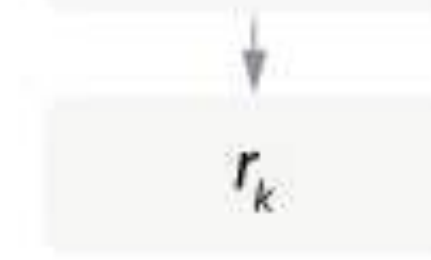
The policy  
generates  
an output.



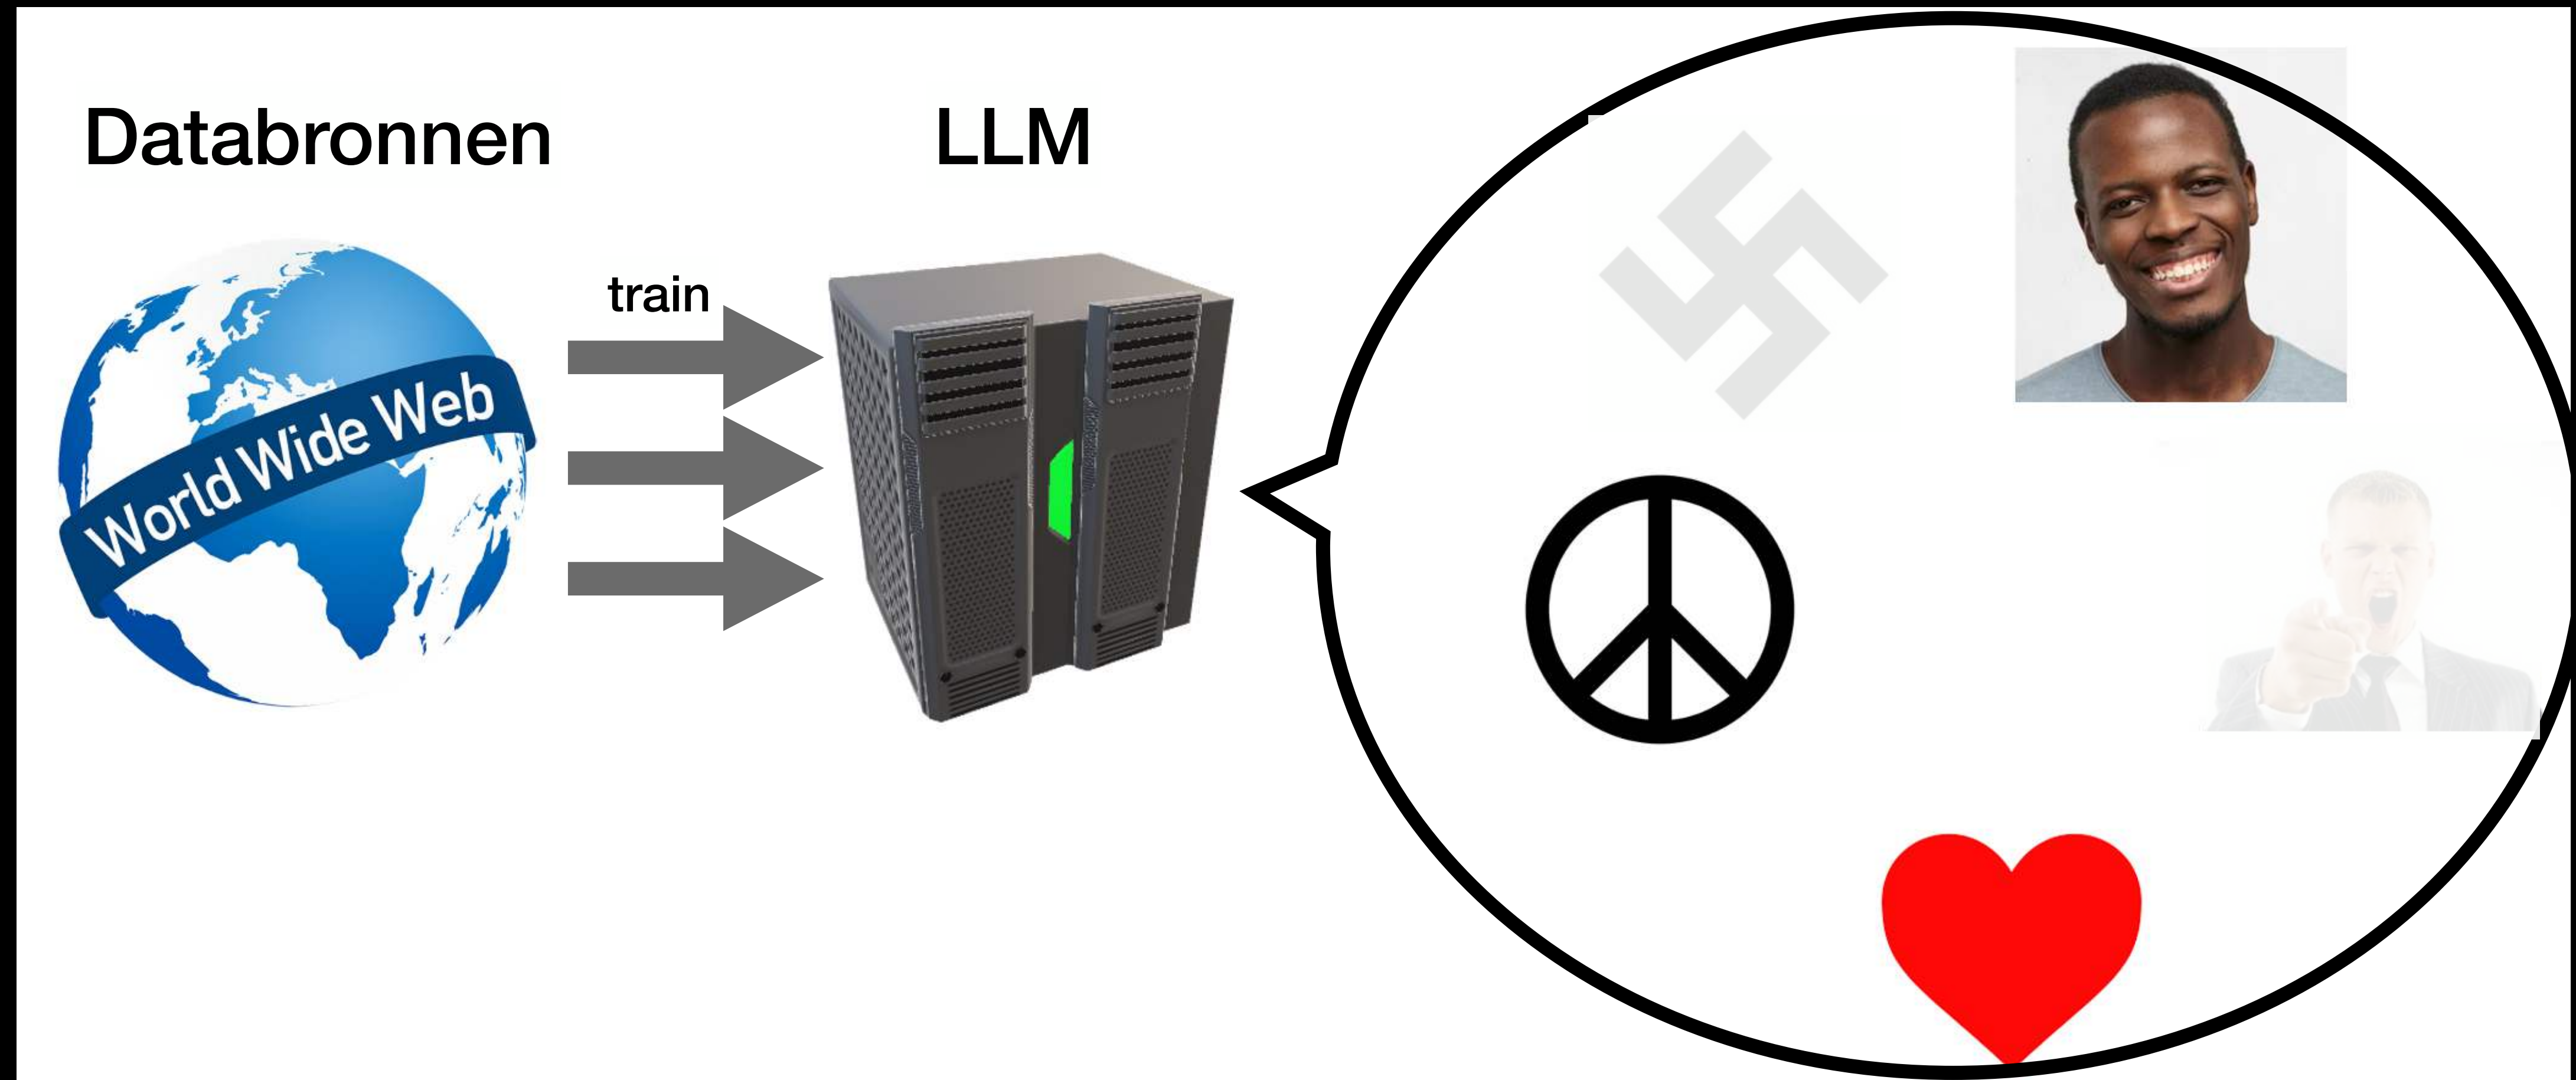
The reward model  
calculates a  
reward for  
the output.



The reward is  
used to update  
the policy  
using PPO.



# Resultaat van finetuning (GPT4)





# Drie stappen

Ontwikkeling

Toepassing

Leren actief

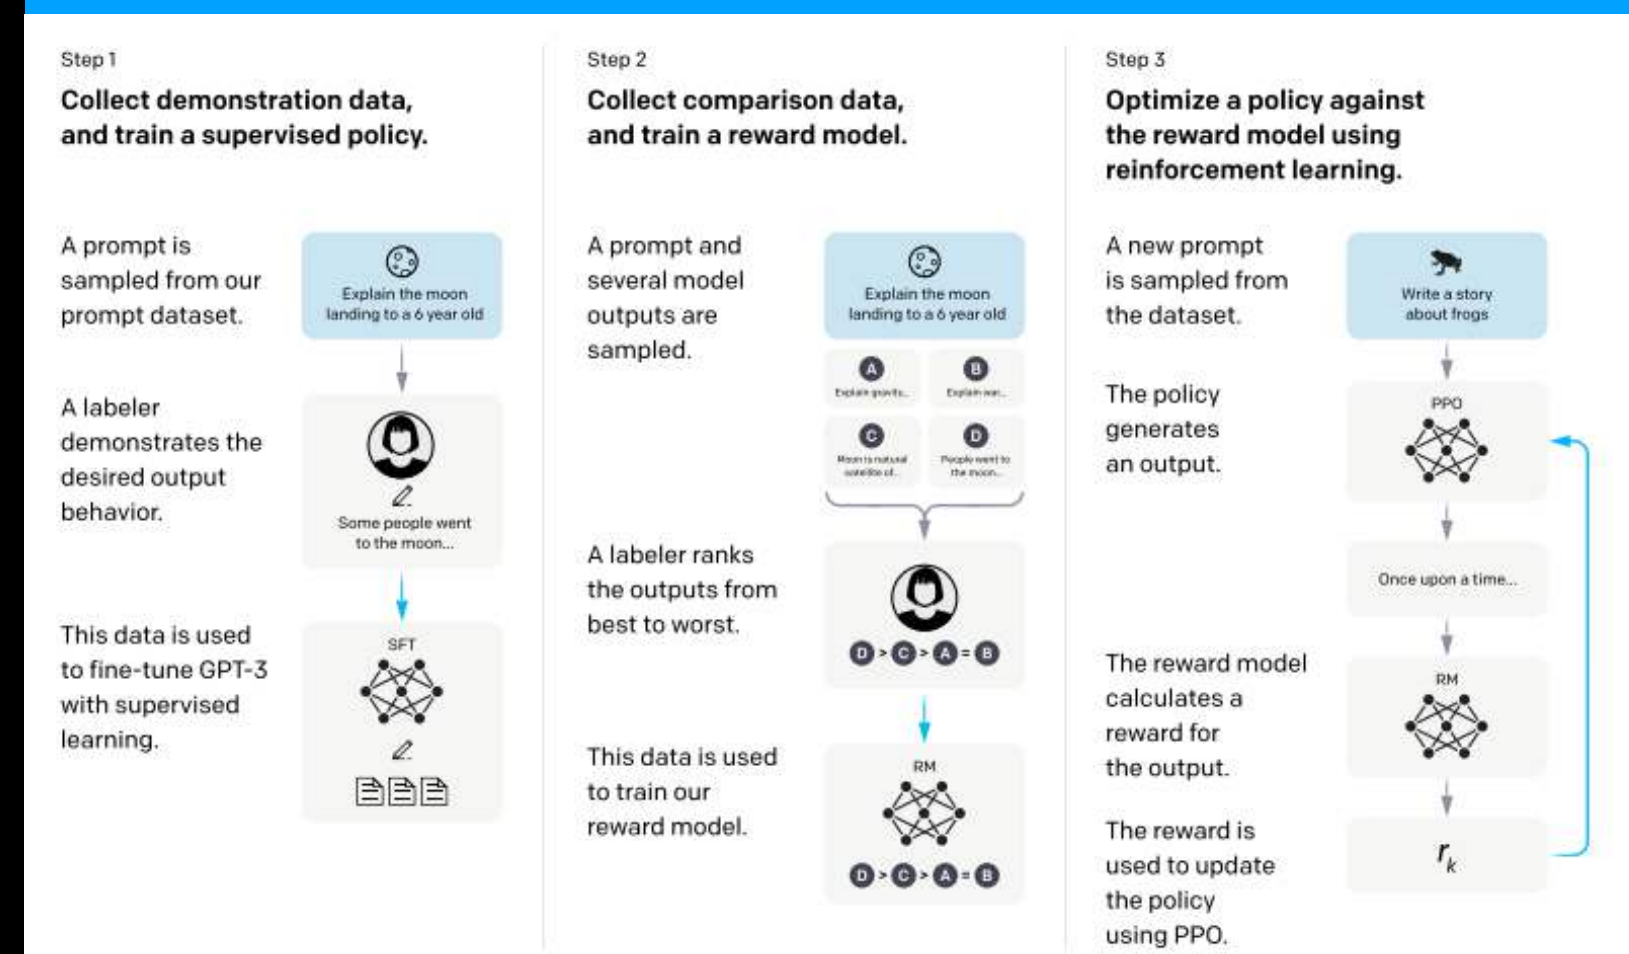
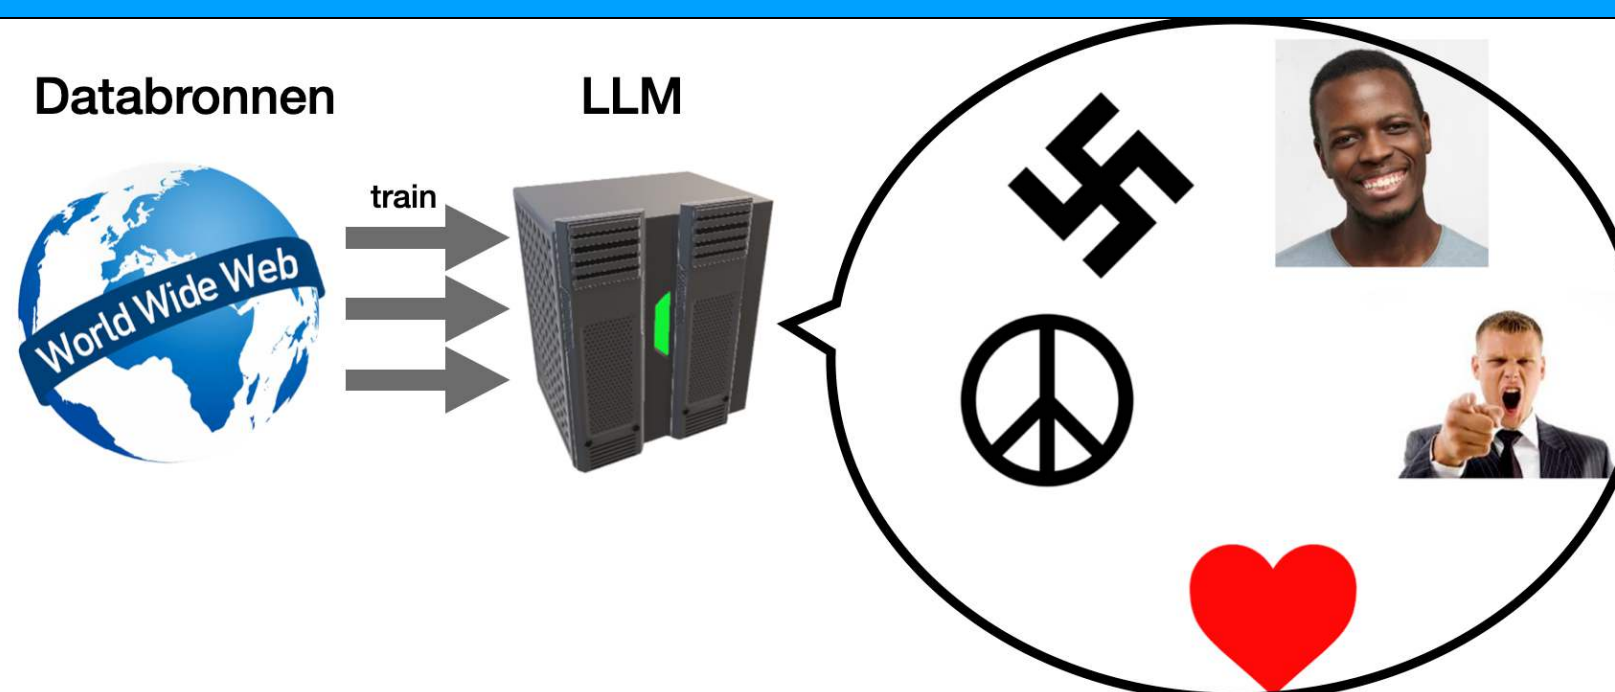
Leren actief

Leren niet actief

## 1. Pre-training

## 2. Finetuning

## 3. In-context learning



You  
what is in-context learning?

ChatGPT  
As of my last knowledge update in January 2022, "in-context learning" generally refers to a learning approach that emphasizes the acquisition of knowledge or skills within a specific context or real-world application. This approach aims to make learning more relevant, practical, and immediately applicable to specific situations or tasks.

User

What is unusual about this image?

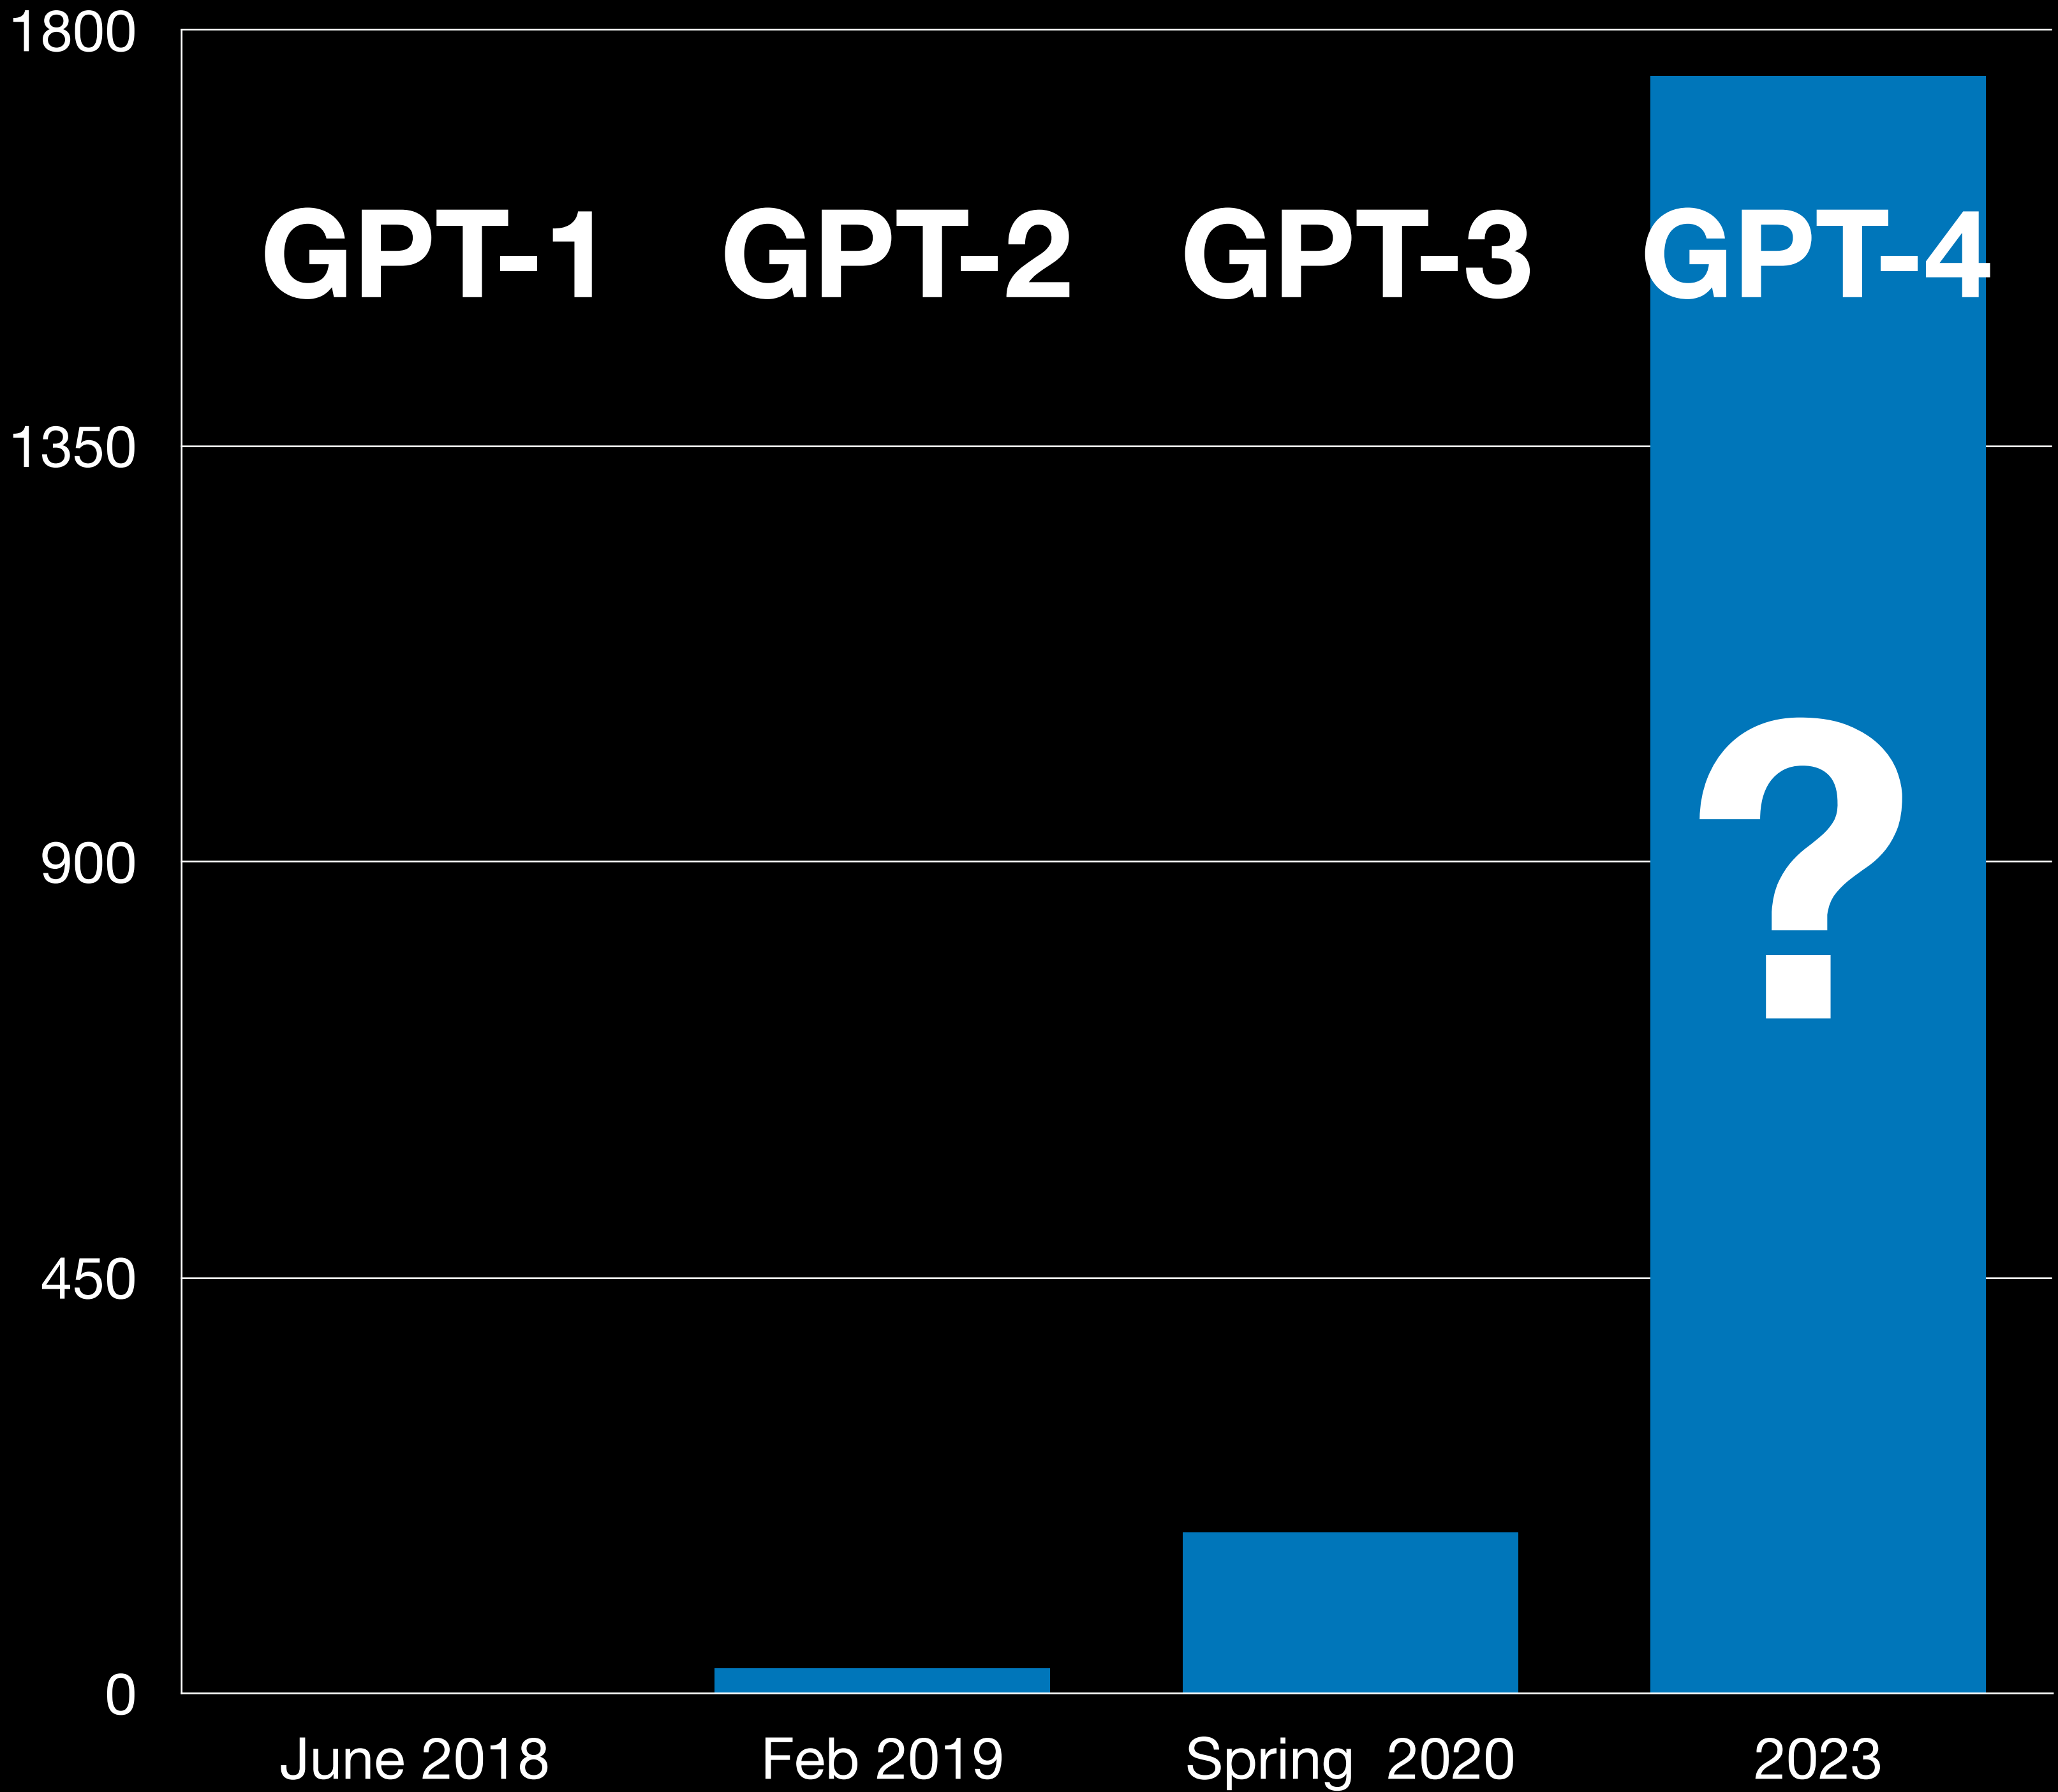


Source: <https://www.barnorama.com/wp-content/uploads/2016/12/03-Confusing-Pictures.jpg>

GPT-4

The unusual thing about this image is that a man is ironing clothes on an ironing board attached to the roof of a moving taxi.

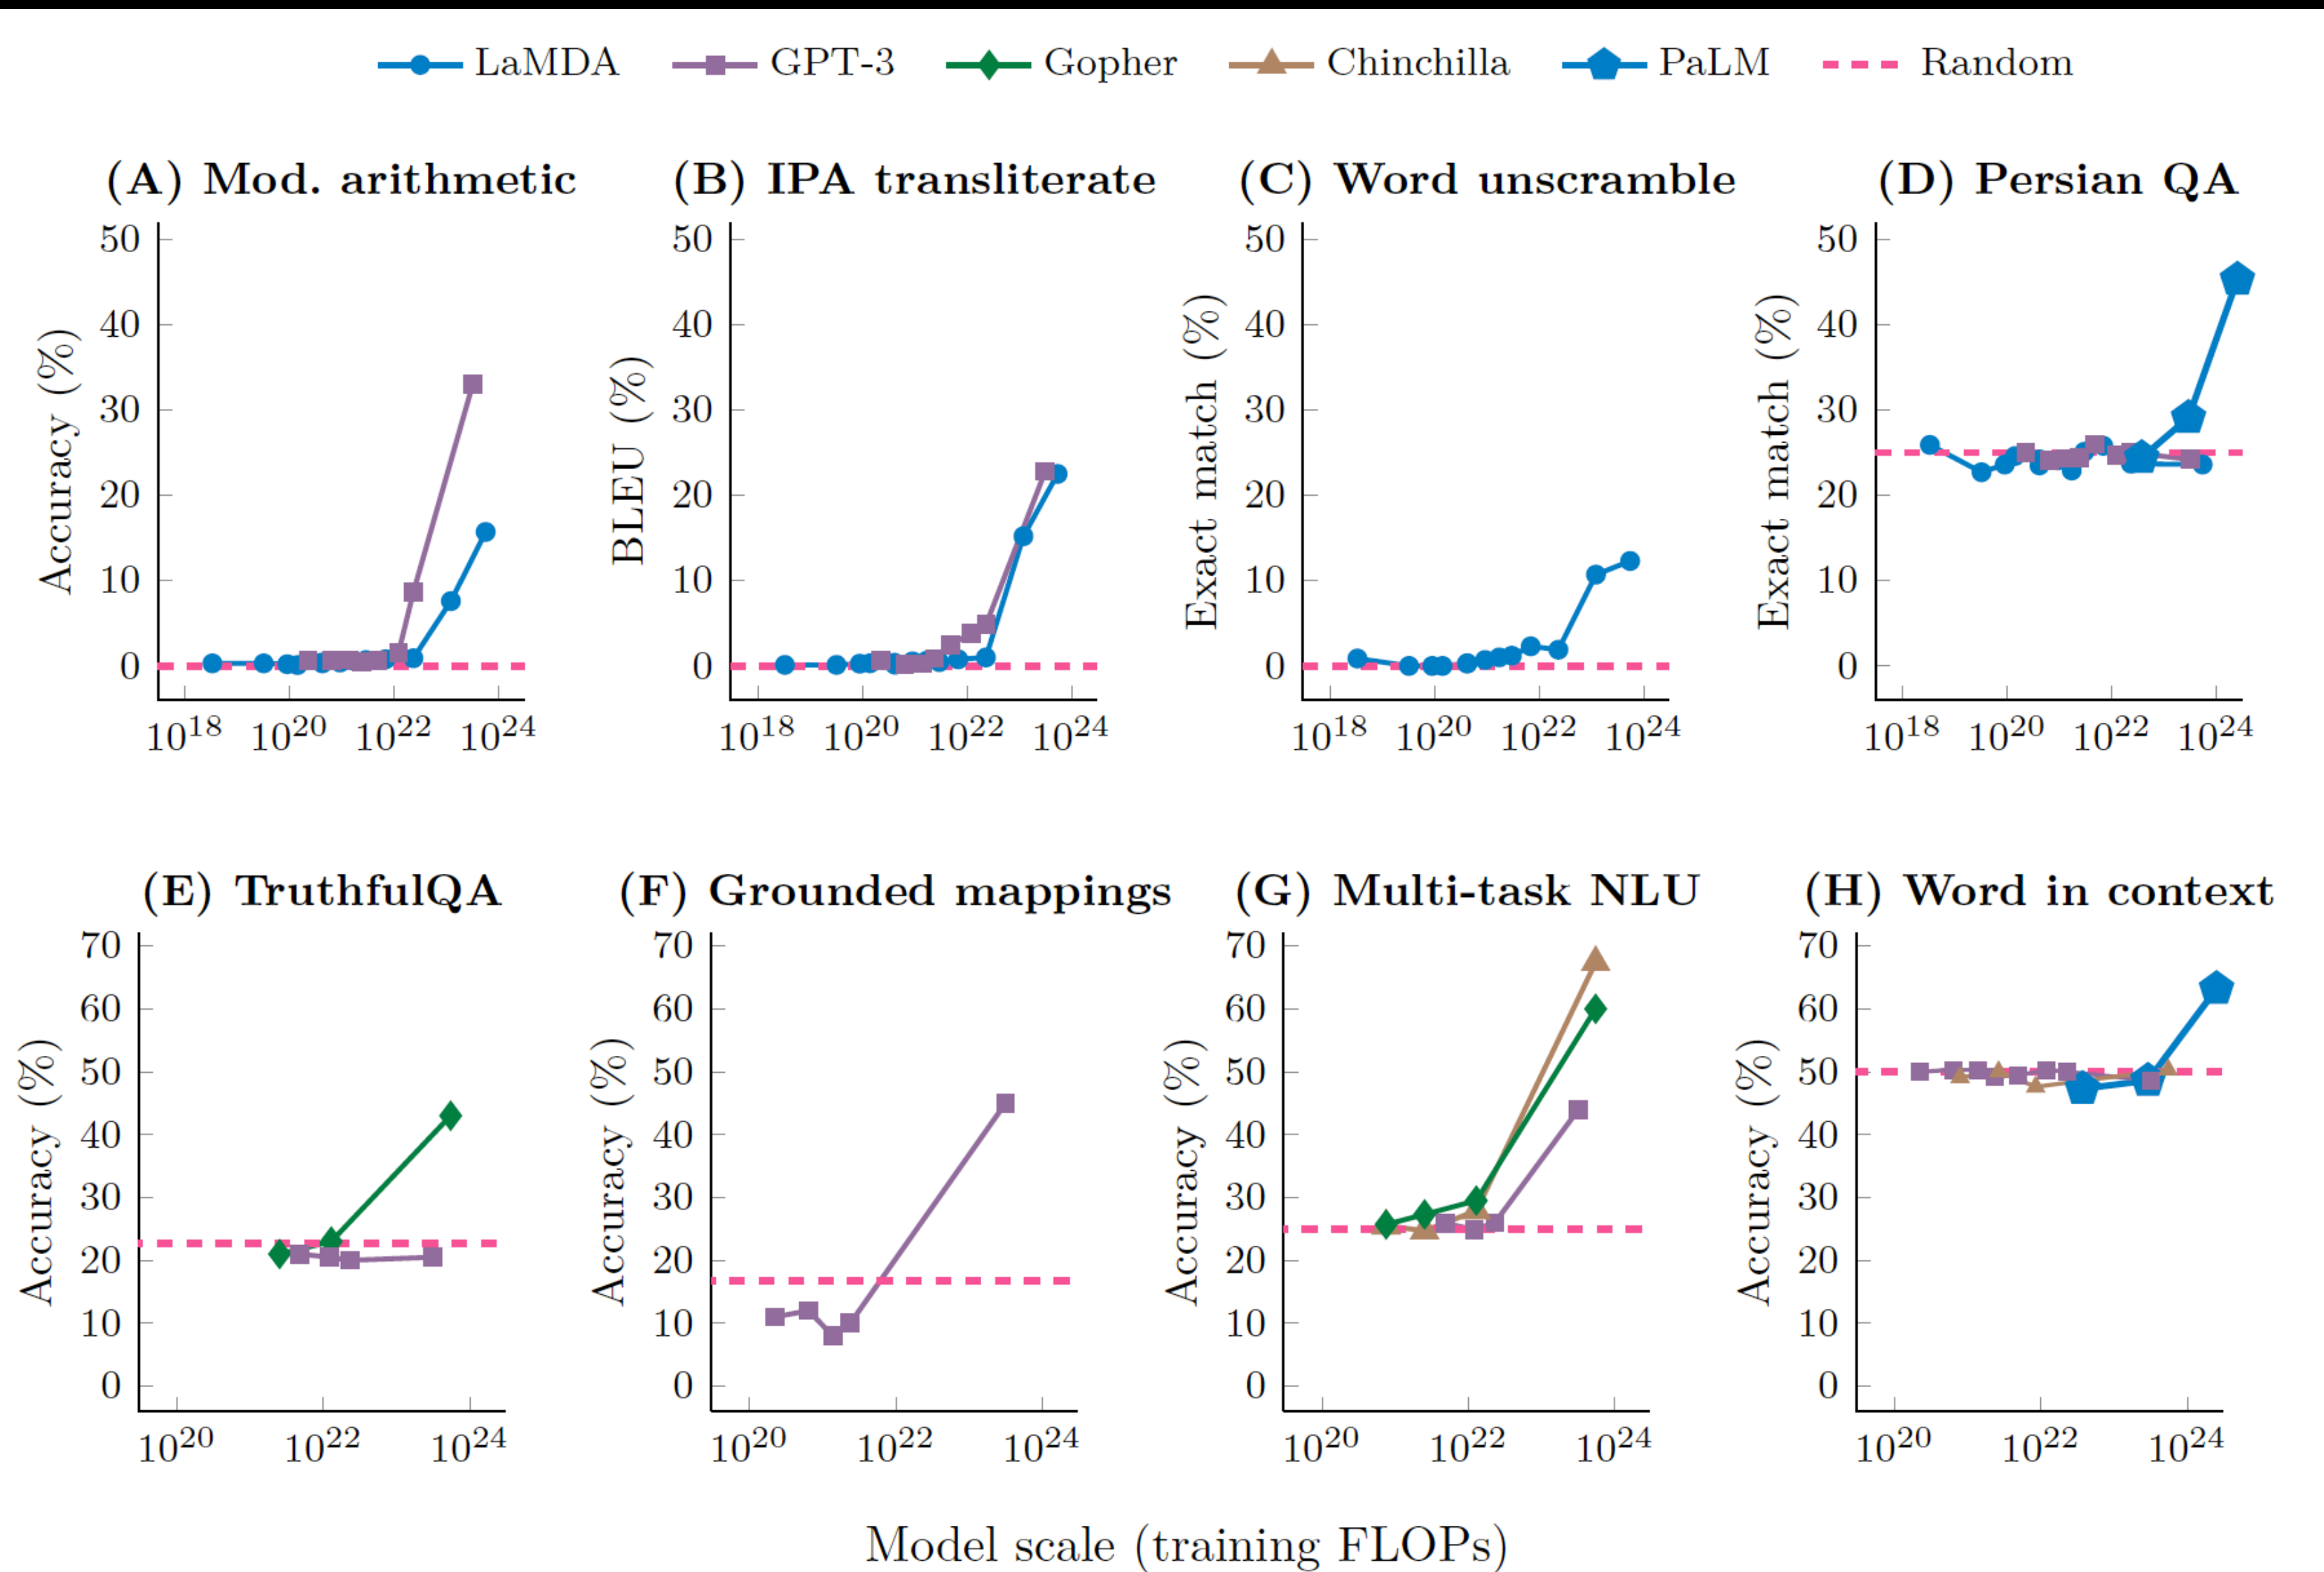
**Number of parameters  
(Billions)**



# Toekomst van Generative AI

- De kracht van LLMs is proportioneel aan het aantal parameters
- Het is niet duidelijk wat het effect is van verdere opschaling van het aantal parameters in GPT systemen
- Emergent properties?

# Phase transitions?



# AI's Ostensible Emergent Abilities Are a Mirage

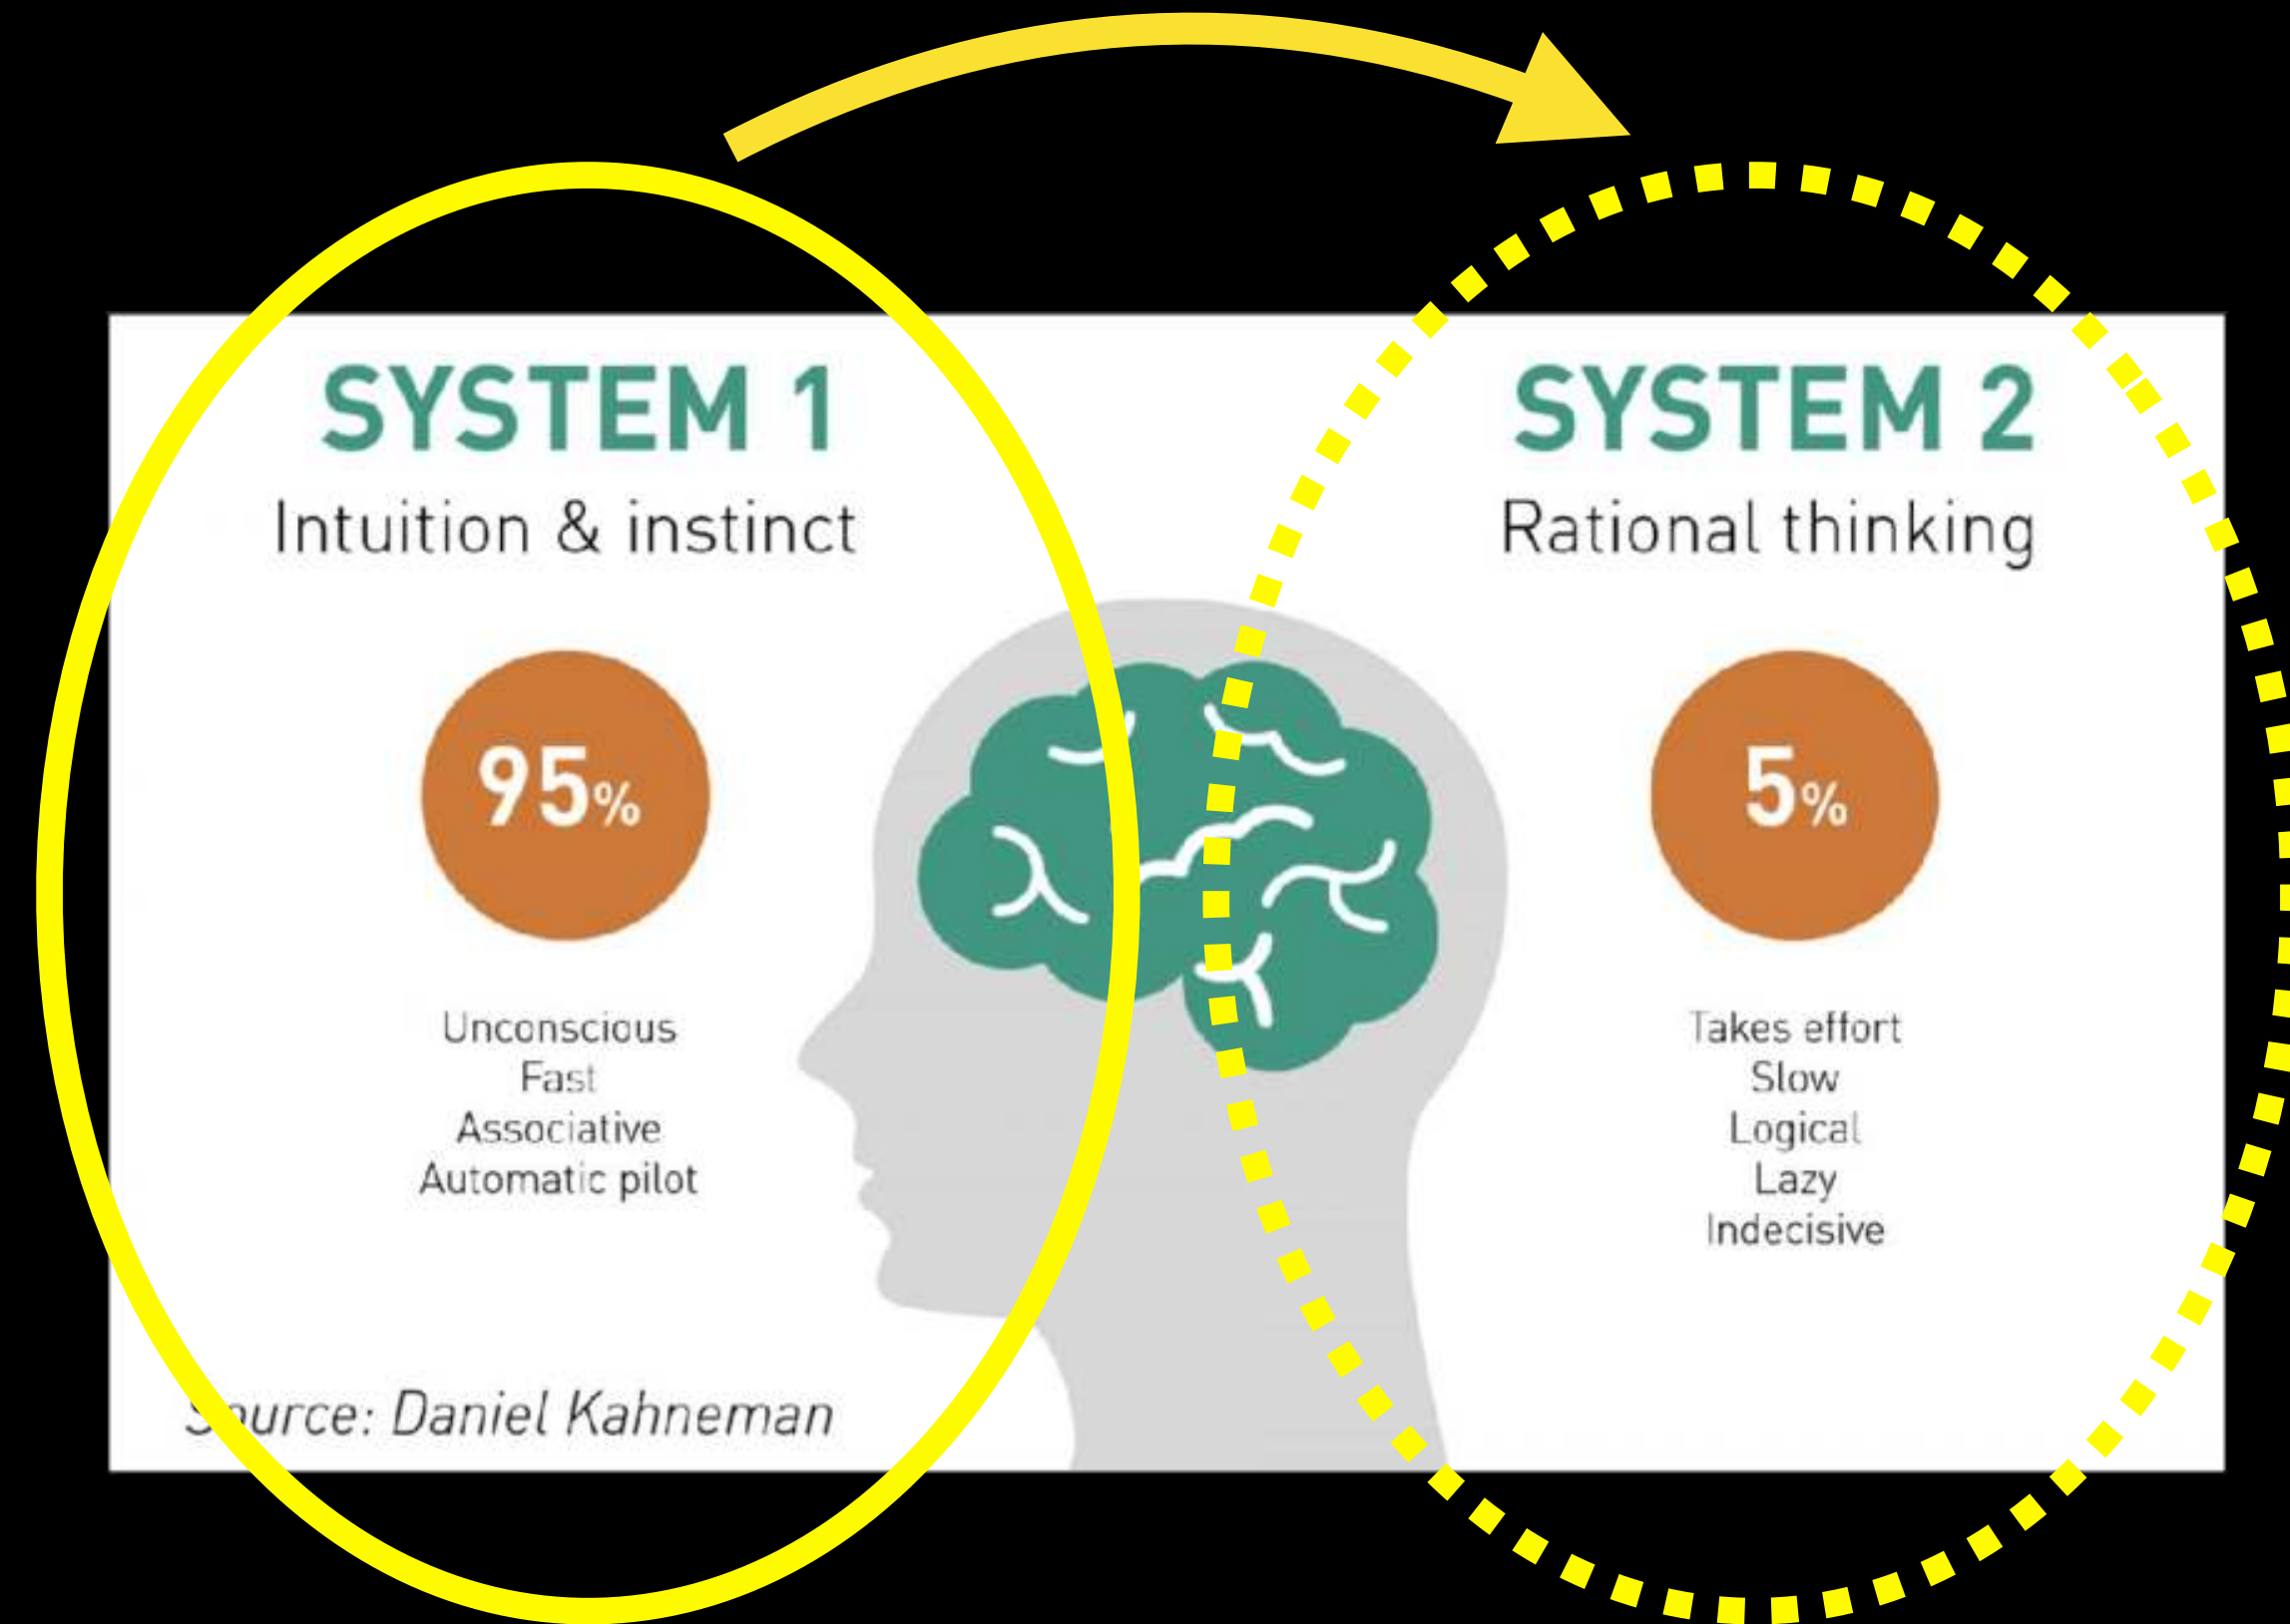
According to Stanford researchers, large language models are not greater than the sum of their parts.

May 8, 2023 | Katharine Miller     



# What LLMs are not able to do...

?



# Francois Chollet's ARC competition

(Abstraction and Reasoning Corpus)

## AGI

Progress toward artificial general intelligence (AGI) has stalled. LLMs are trained on unimaginably vast amounts of data, yet they remain unable to adapt to simple problems they haven't been trained on, or make novel inventions, no matter how basic.

Strong market incentives have pushed frontier AI research to go closed source. Research attention and resources are being pulled toward a dead end. You can change that.

## DEFINING AGI

Consensus but wrong:

AGI is a system that can automate the majority of economically valuable work.

Correct:

AGI is a system that can efficiently acquire new skills and solve open-ended problems.

Definitions are important. We turn them into benchmarks to measure progress toward AGI.

Without AGI, we will never have systems that can invent and discover alongside humans.



<https://arcprize.org/>

## ARC PRIZE

ARC Prize is a \$1,000,000+ public competition to beat and open source a solution to the ARC-AGI benchmark.

Hosted by Mike Knoop (Co-founder, Zapier) and François Chollet (Creator of ARC-AGI, Keras).

> [Start here](#)

## ARC-AGI

Most AI benchmarks measure skill. But skill is not intelligence. General intelligence is the ability to efficiently acquire new skills.

Chollet's unbeaten 2019 Abstraction and Reasoning Corpus for Artificial General Intelligence (ARC-AGI) is the only formal benchmark of AGI.

It's easy for humans, but hard for AI.

## PLAY

Try ARC-AGI. Given the examples, identify the pattern, solve the test puzzle.

Puzzle ID: 3aa6fb7a

Previous

1 of 6

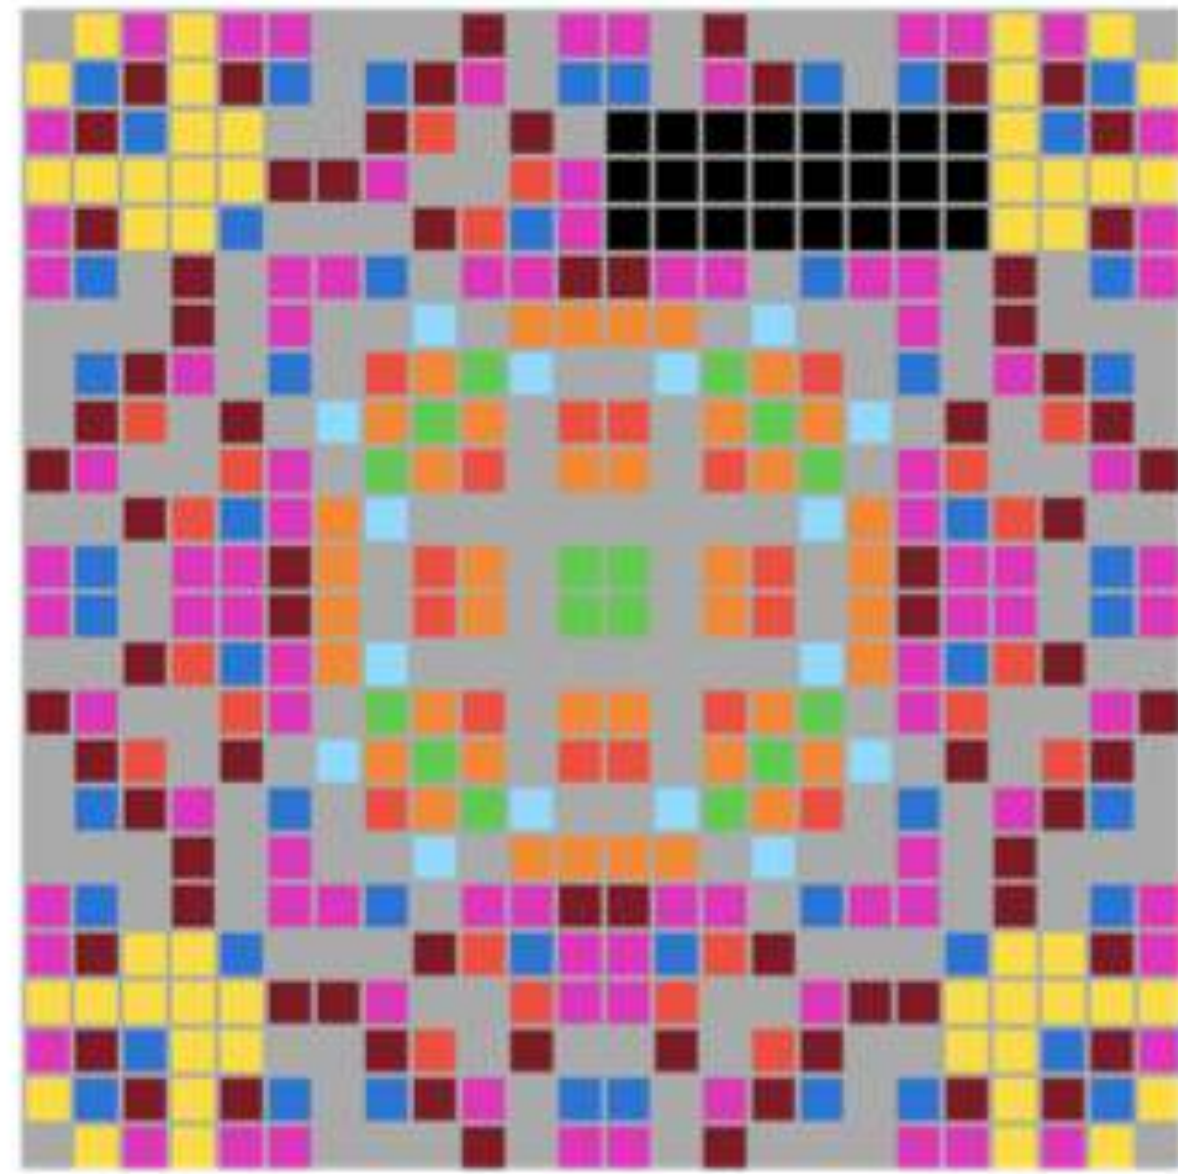
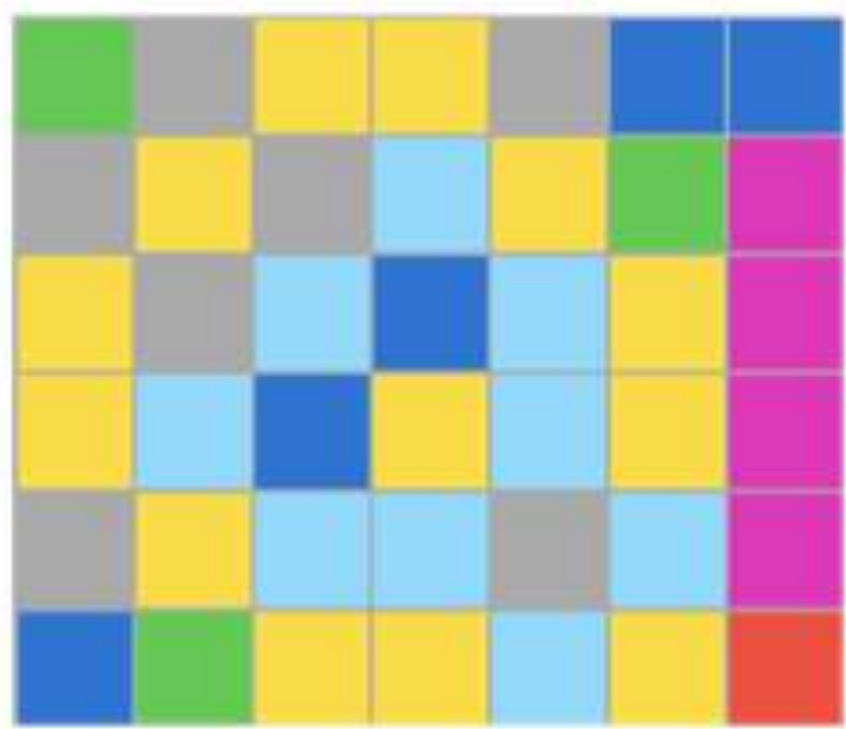
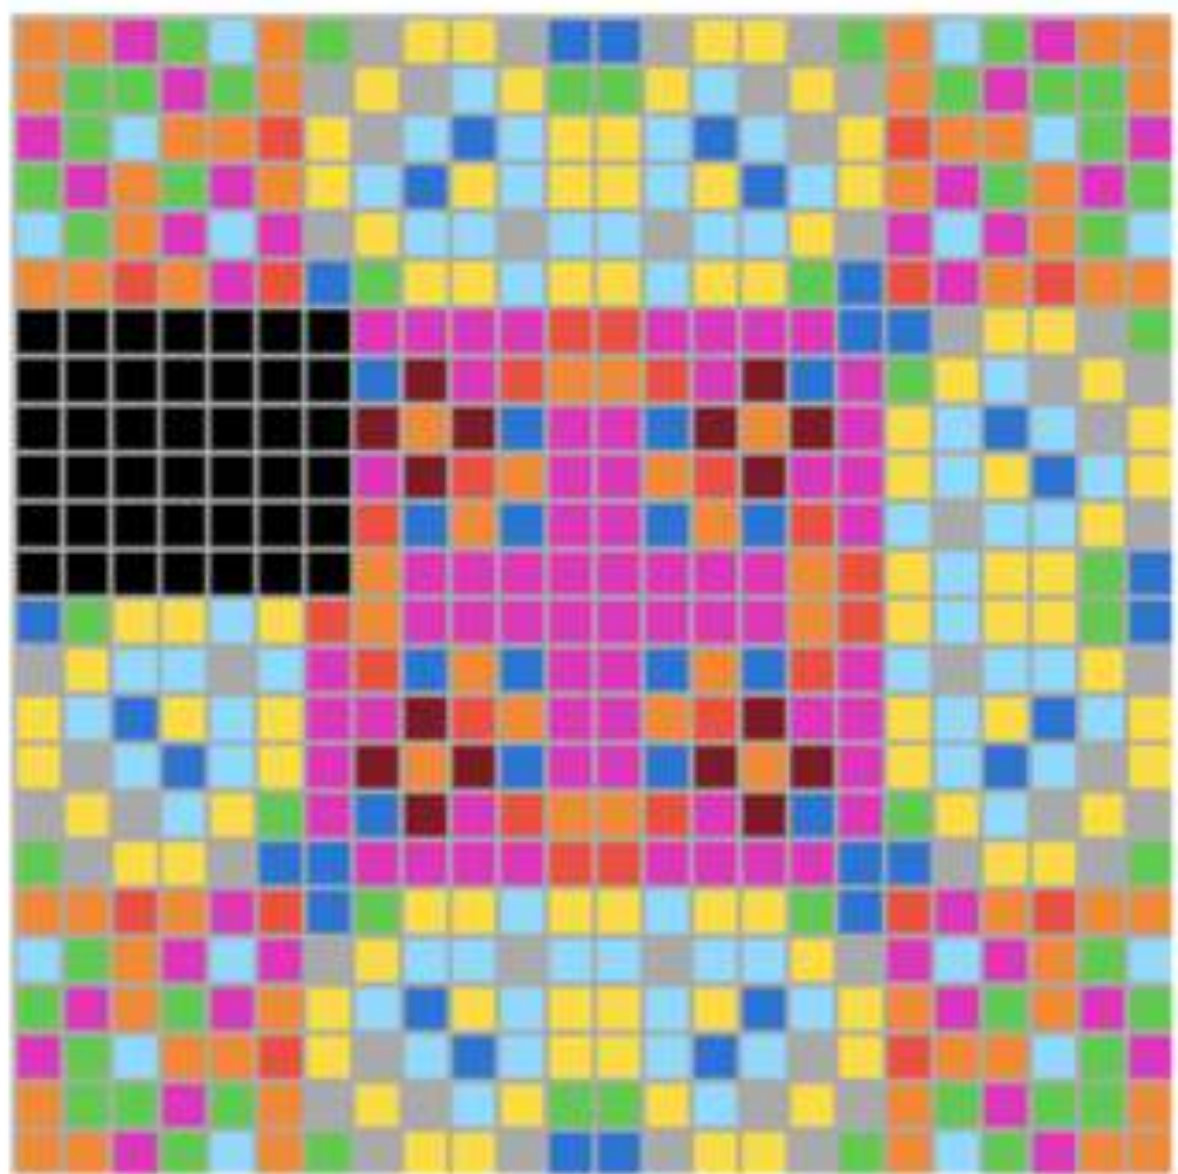
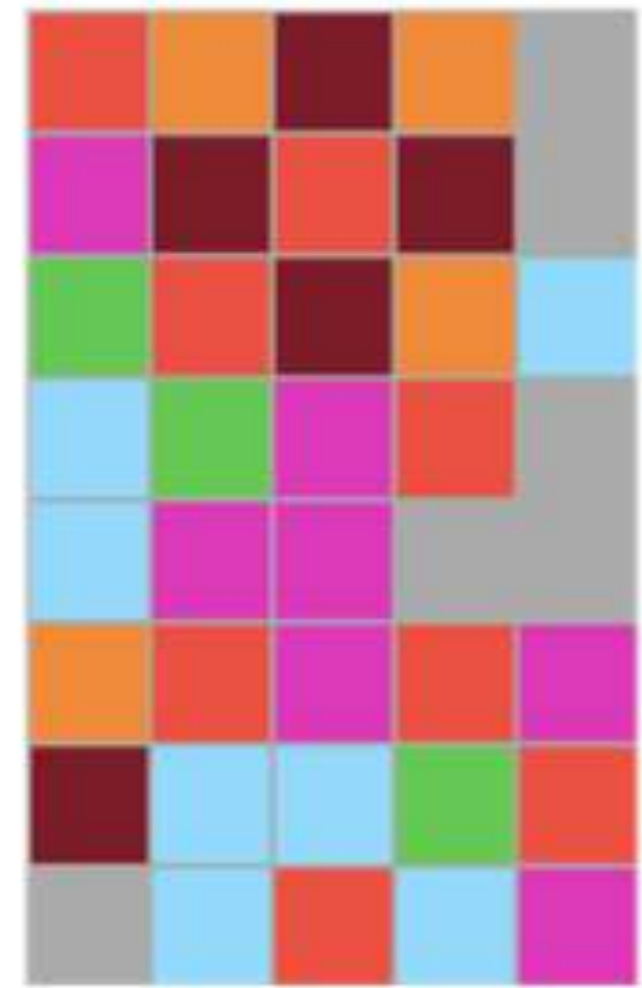
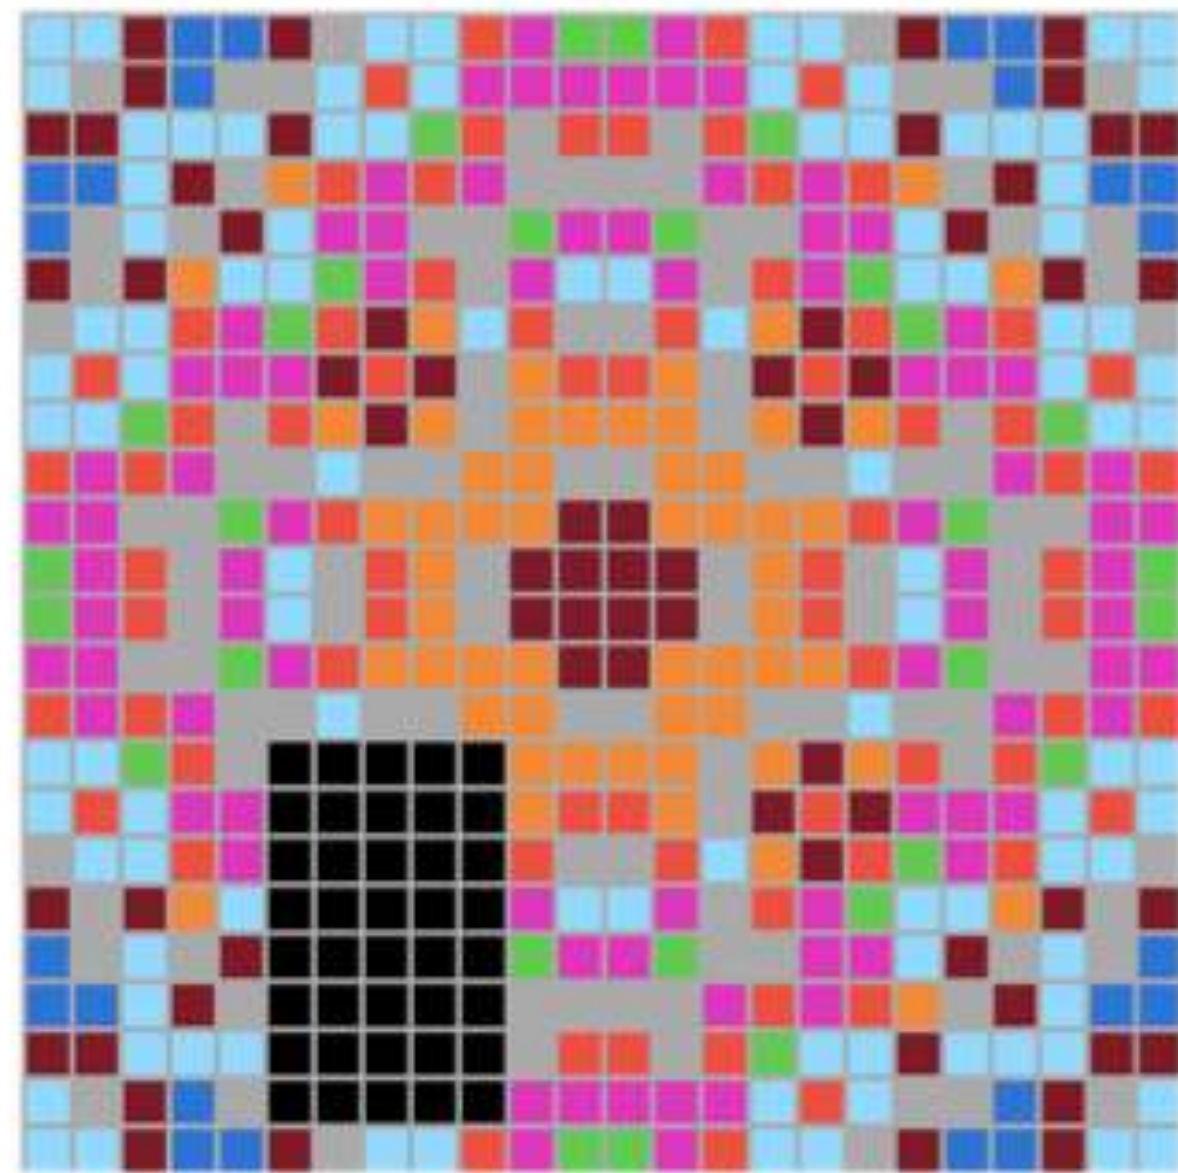
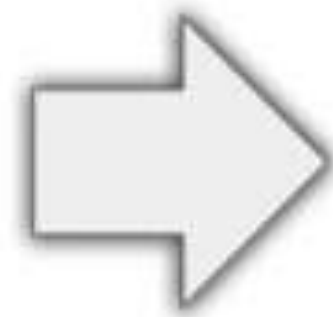
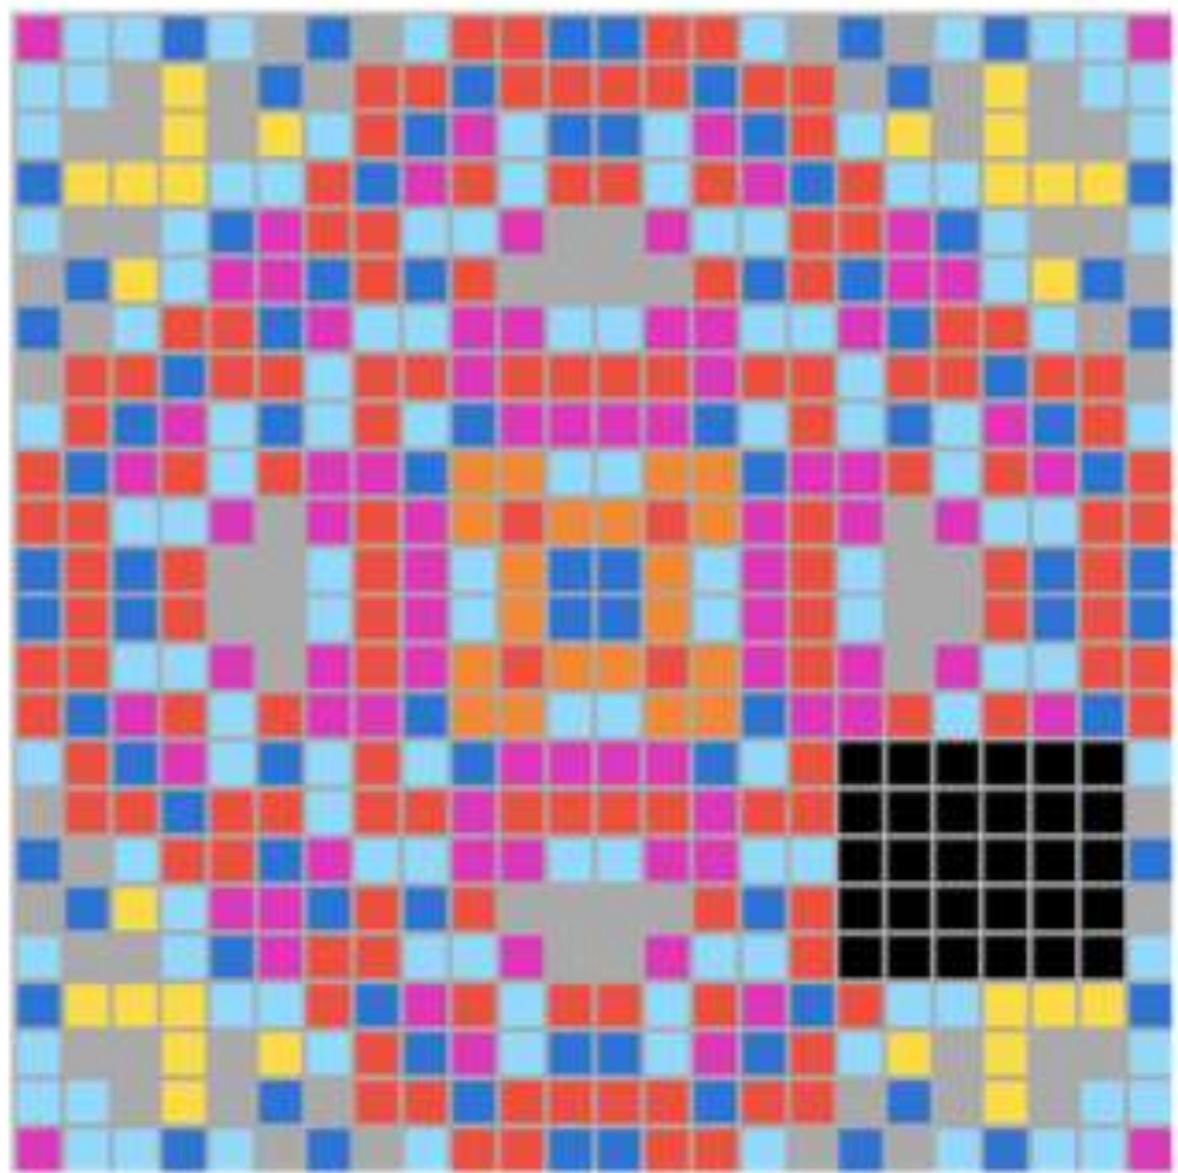
Next

## EXAMPLES

Ex.1 Input (7x7) → Ex.1 Output (7x7)

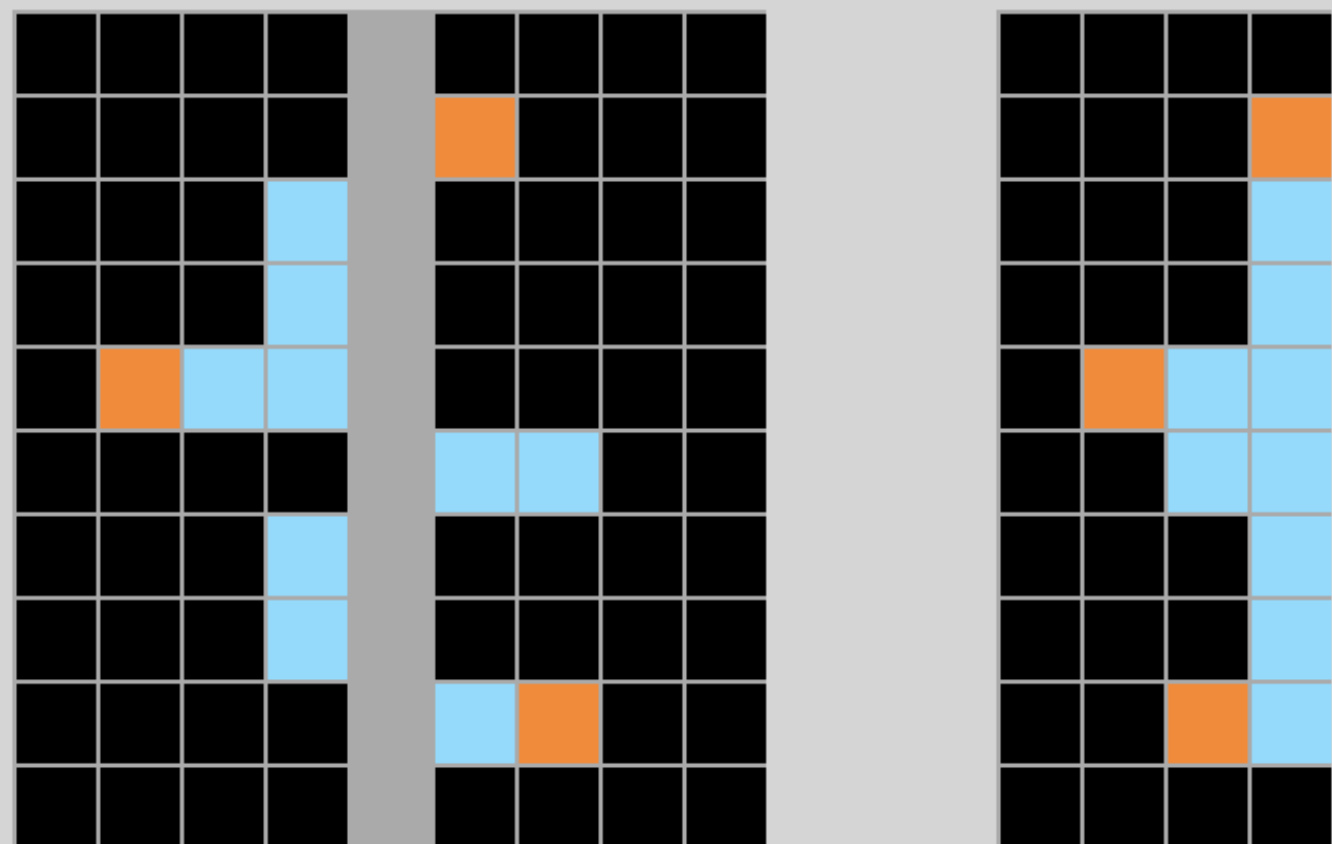
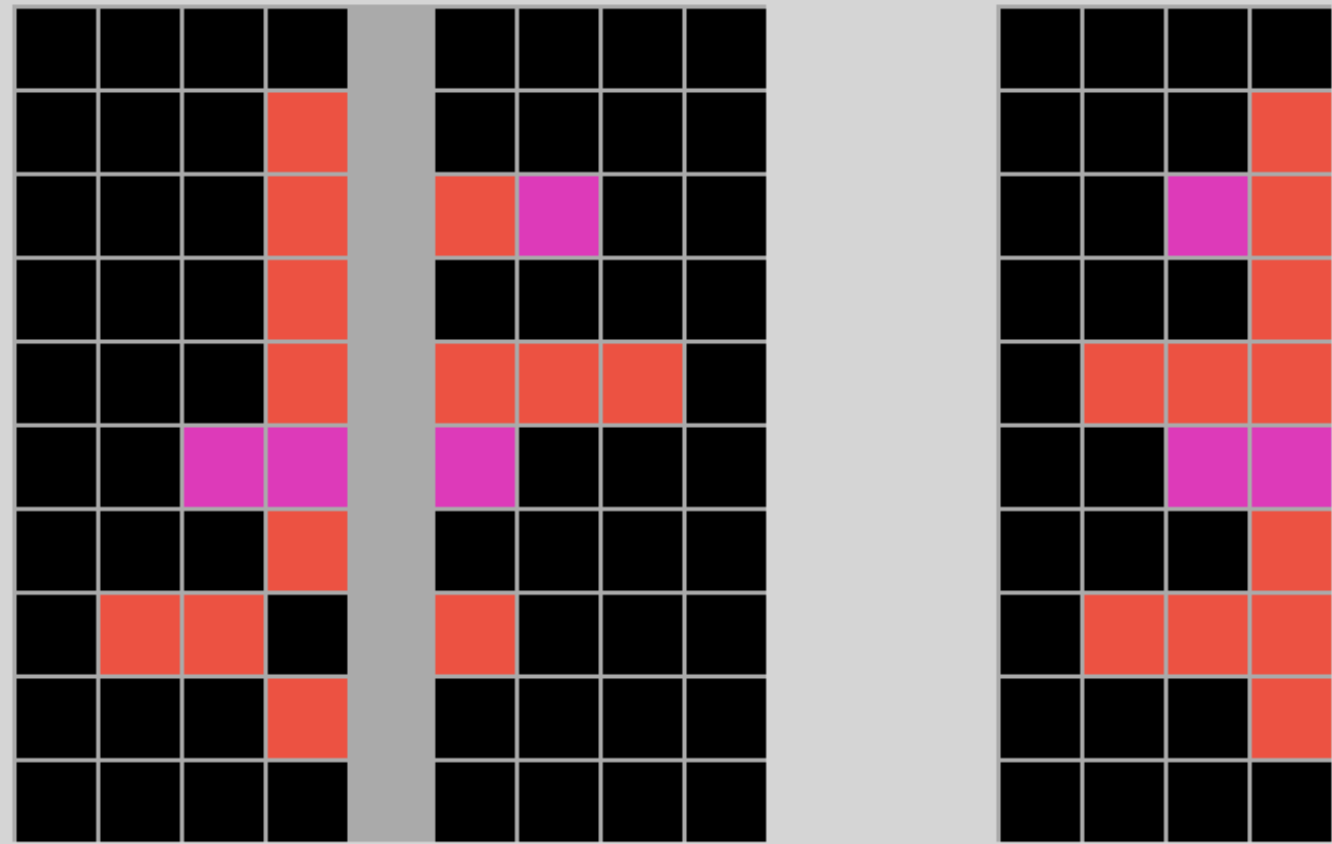
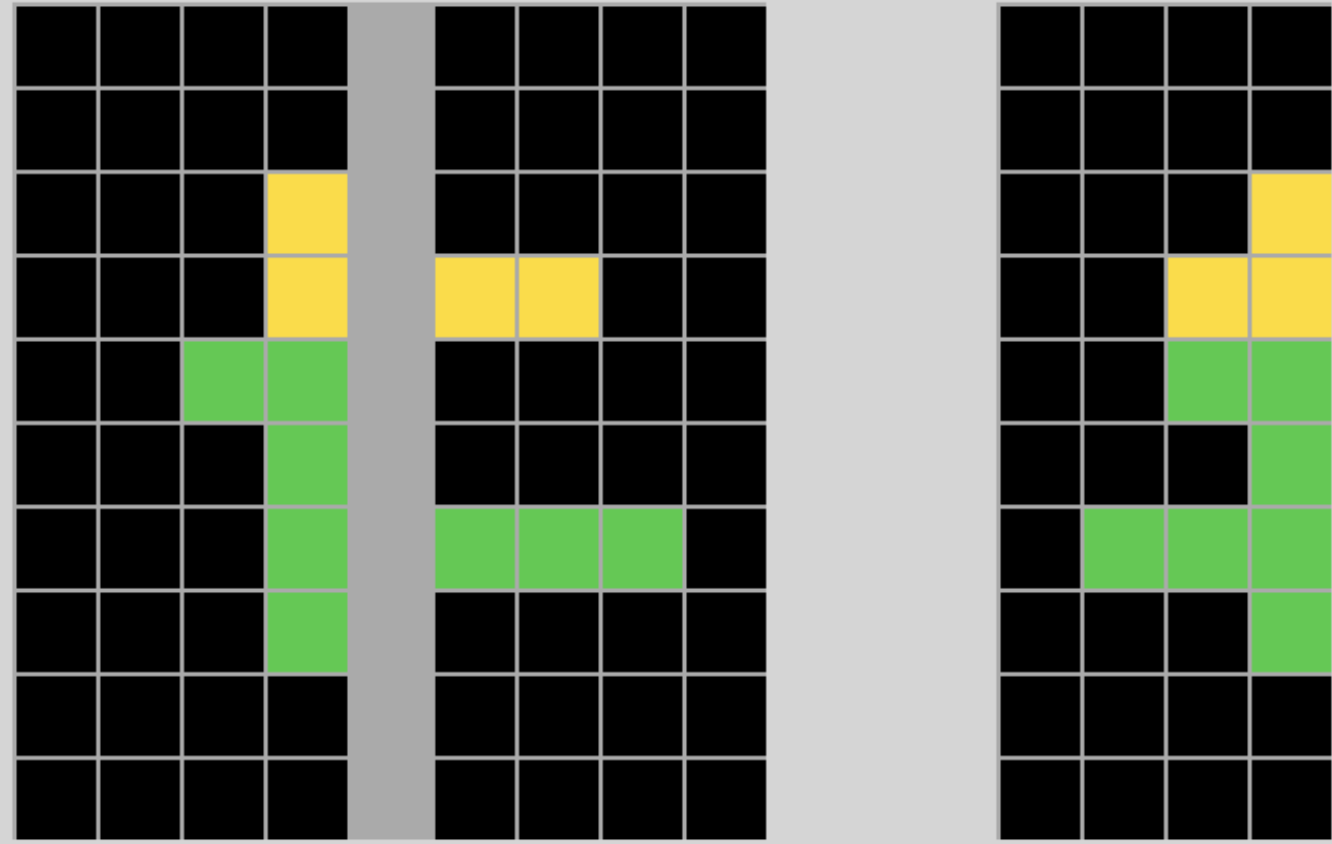
## TEST

Input (7x7) → Output (7x7)



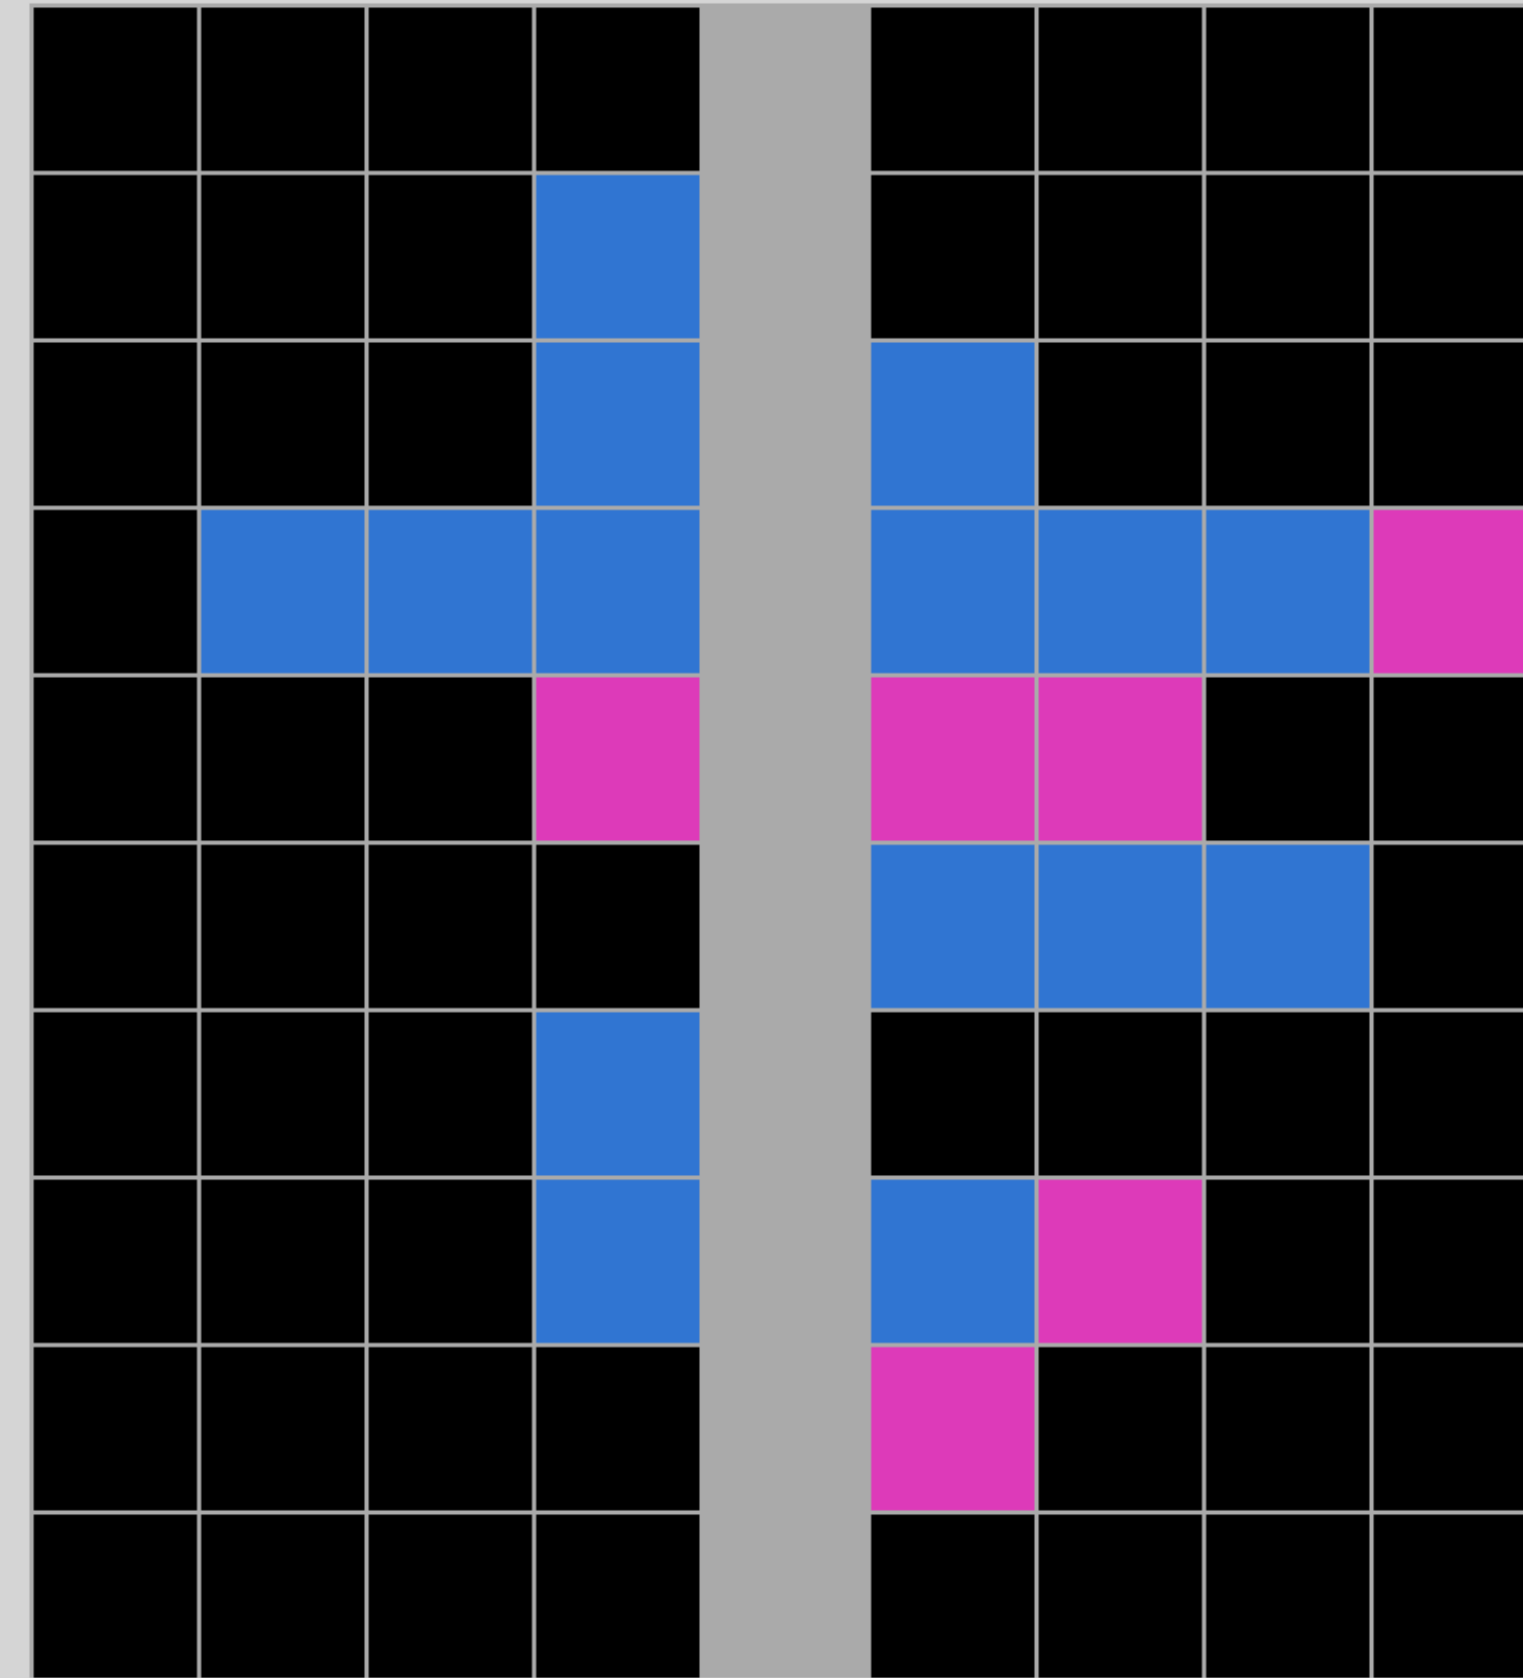
?

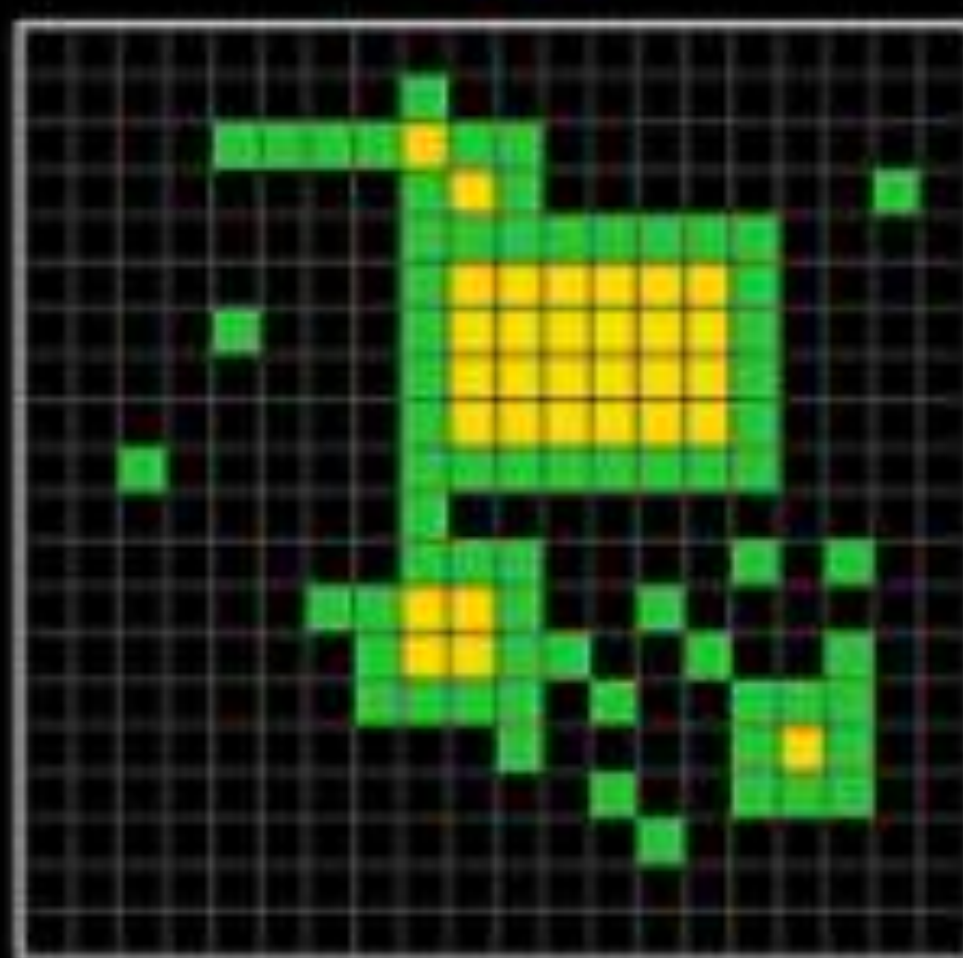
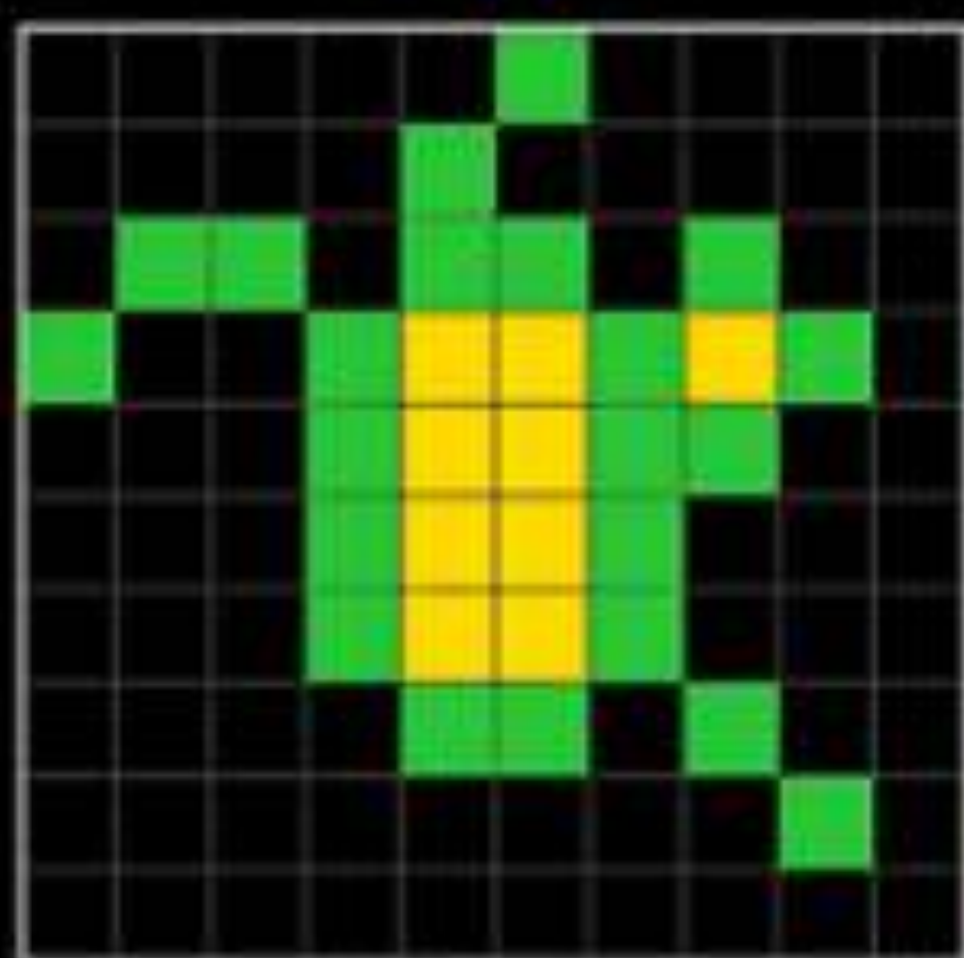
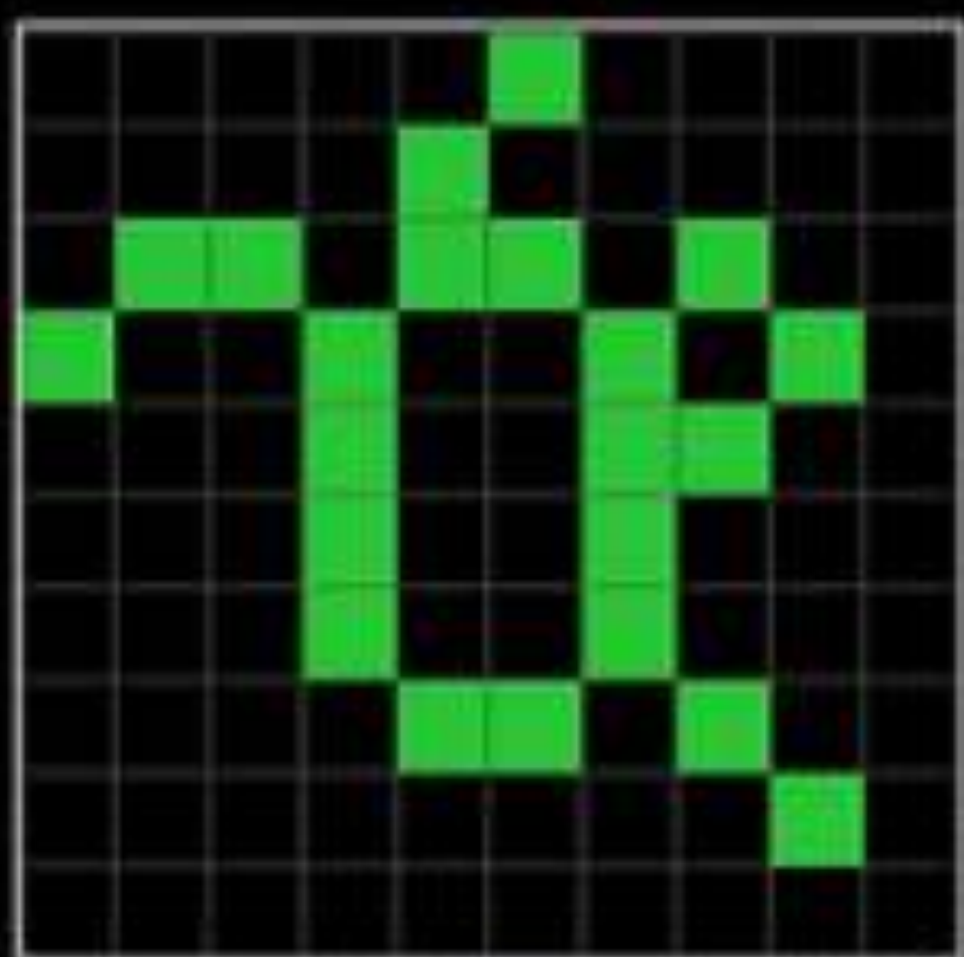
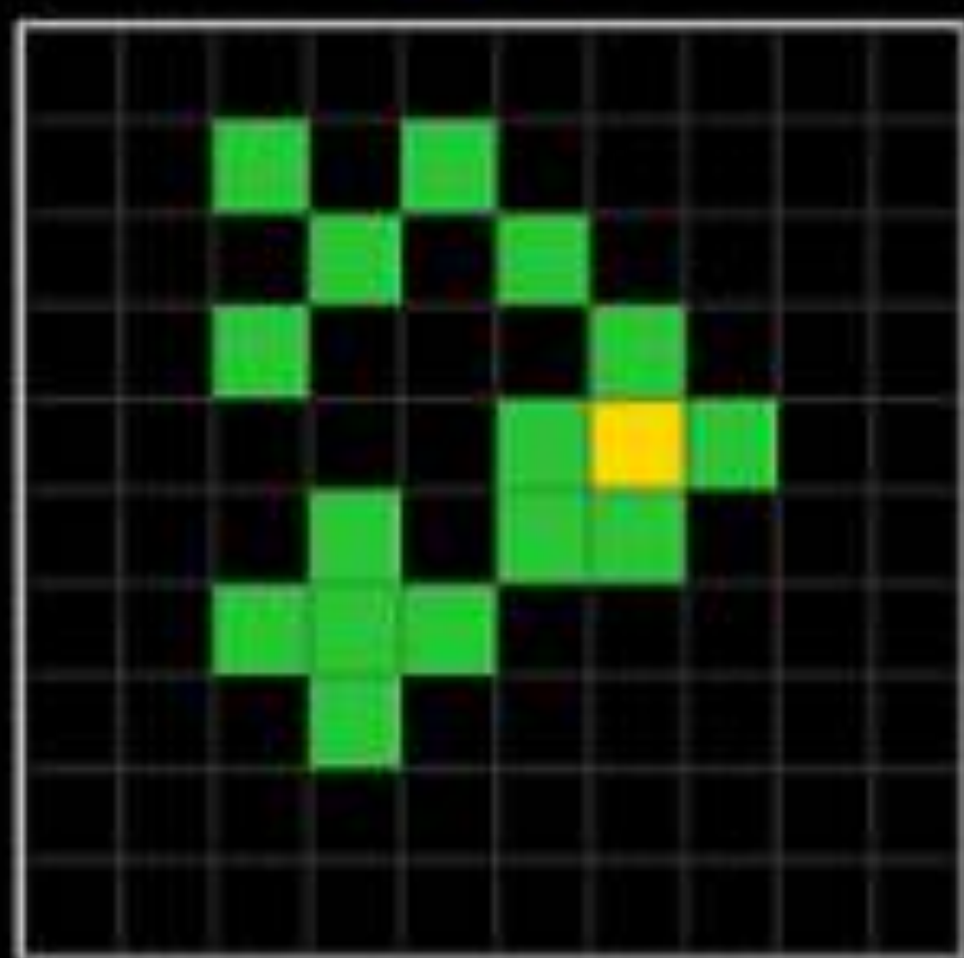
### Task demonstration



### Test input grid 1/1

Next test input





# Example of attempt on ARC challenge

## The Surprising Effectiveness of Test-Time Training for Abstract Reasoning

Ekin Akyürek Mehul Damani Linlu Qiu Han Guo Yoon Kim Jacob Andreas

Massachusetts Institute of Technology

Program Synthesizer	Fine-tuned LM	TTT Method	Score (pass@2)	
X	Ours	X	18.3%	
X	Ours	Ours	47.1%	
X	BARC	Ours	53.0%	
BARC	Ours	Ours	58.5%	
BARC	BARC	Ours	<b>61.9%</b>	
			Avg. Human	60.2%
			Best Human	<b>97.8%</b>
BARC (ensemble)			54.4%	
BARC (no synthesizer)			39.3%	
Claude - Few-shot prompting			21.0%	
GPT-4.0 - Few-shot prompting			9.0%	

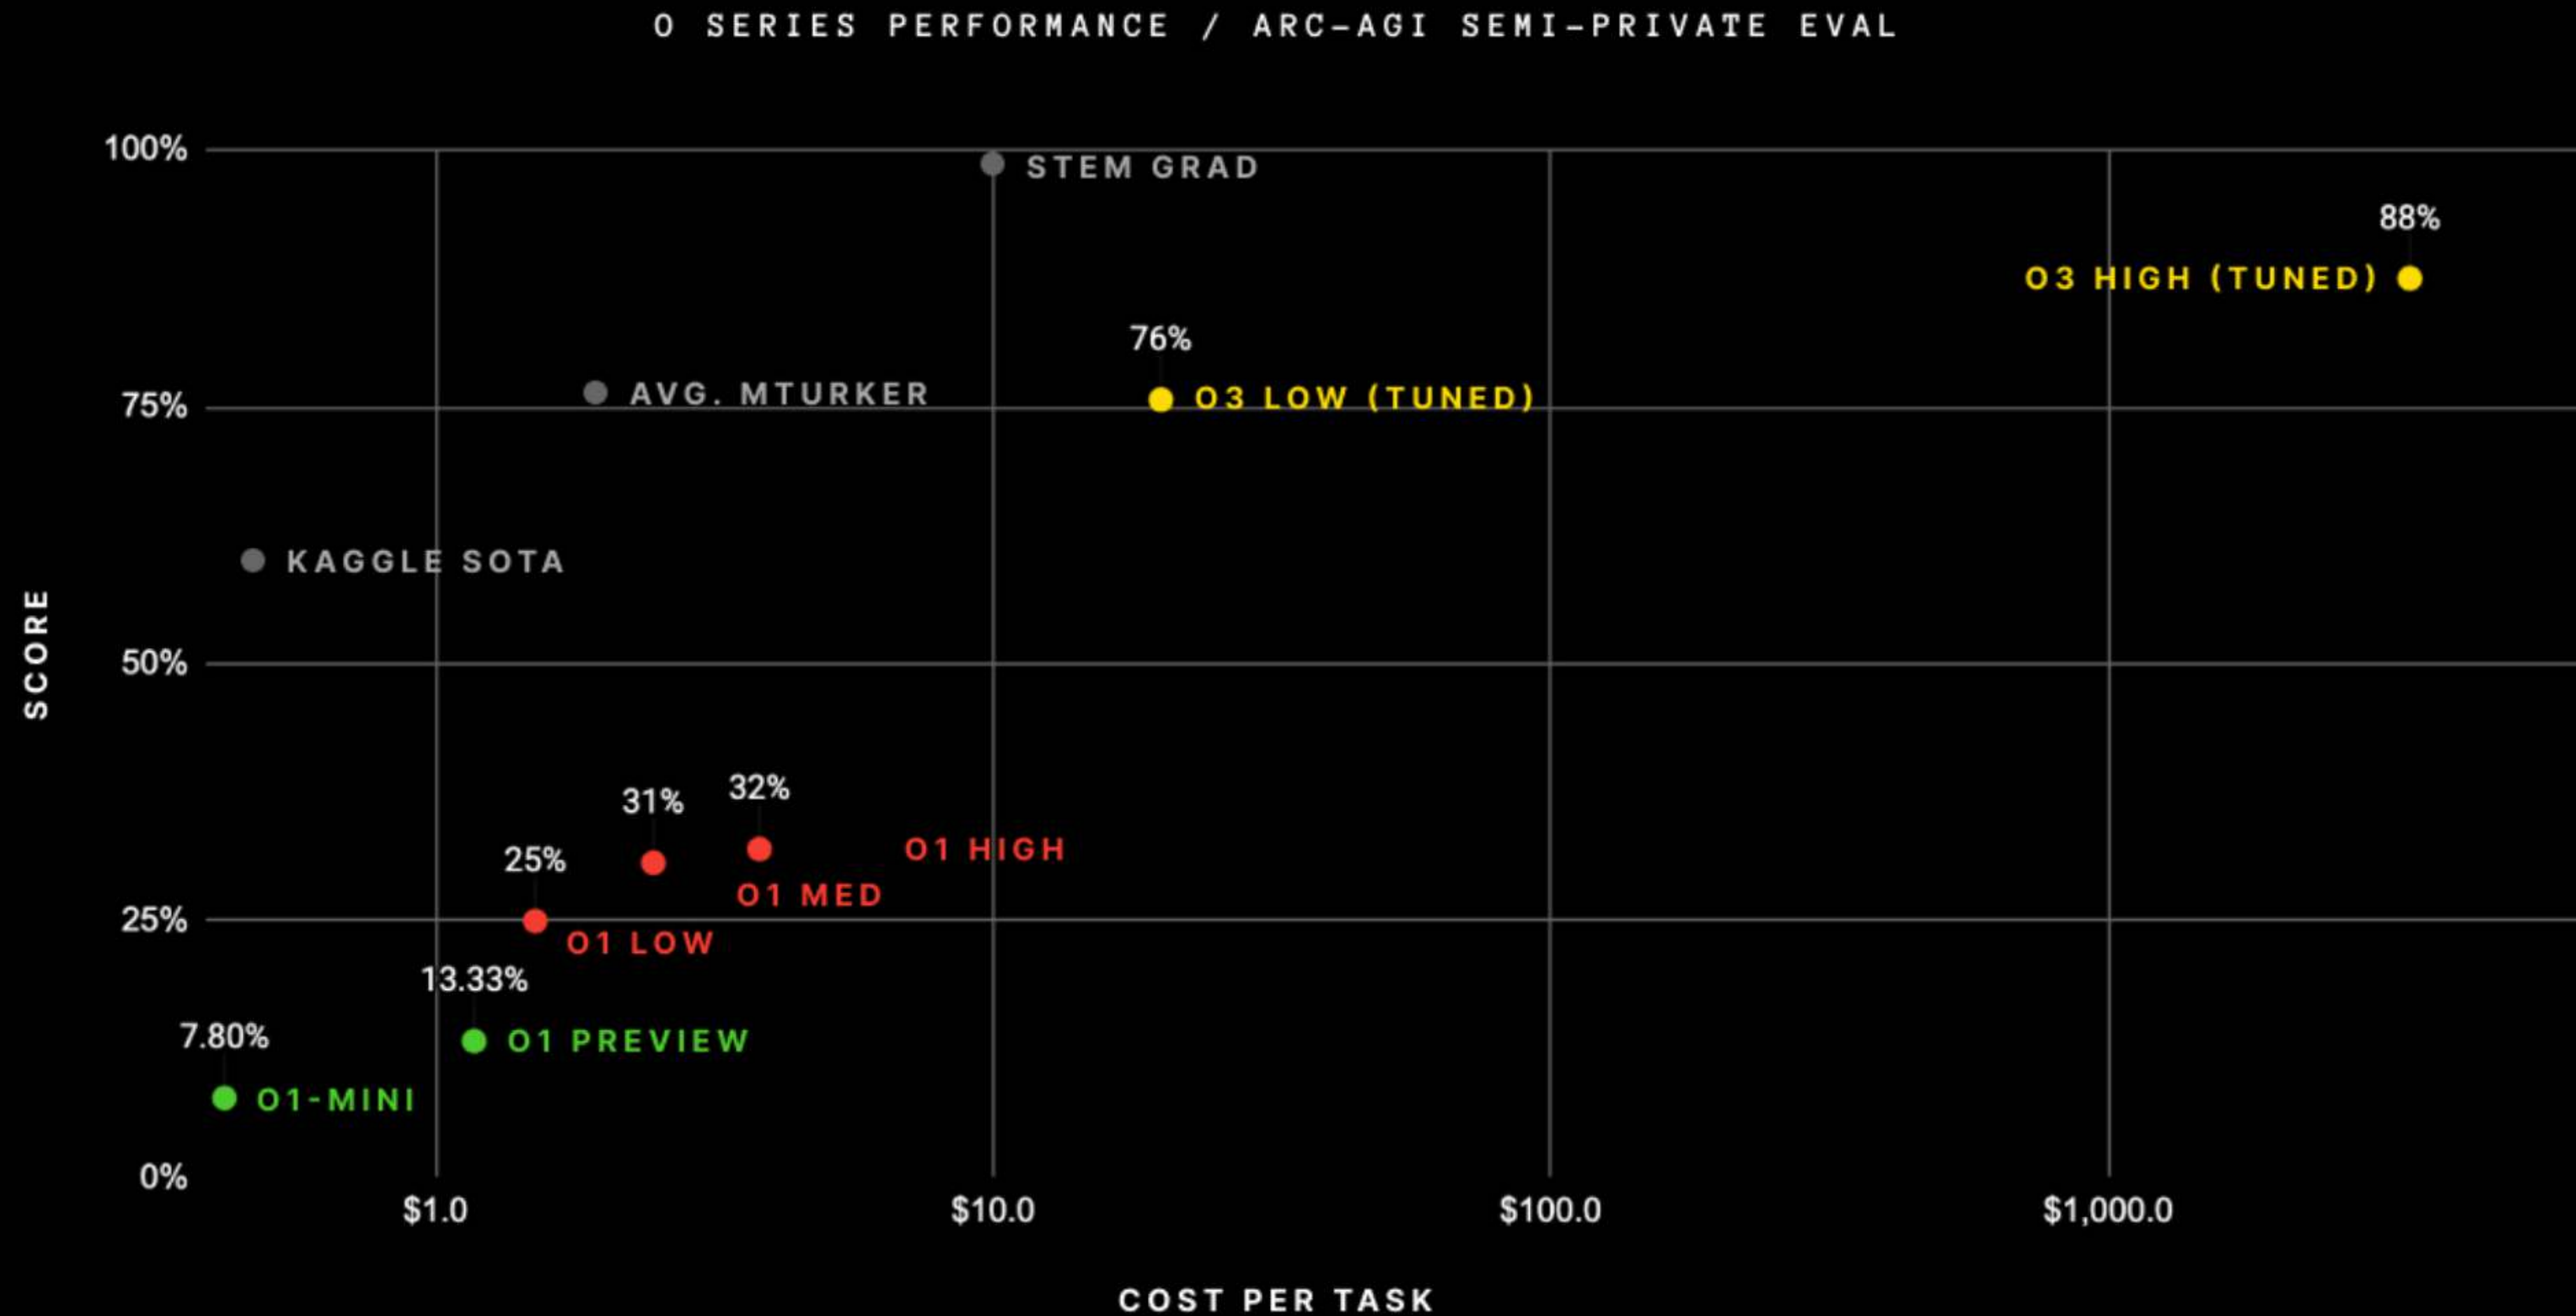
# Open AI's o3 breakthrough

20 December 2024

- Chain-of-Thought (CoT) sequences
- Training on ARC tasks
- Test-time search for appropriate combination of program sequences


# OPENAI O3 BREAKTHROUGH HIGH SCORE ON ARC-AGI-PUB

OpenAI's new o3 system – trained on the ARC-AGI-1 Public Training set – has scored a breakthrough 75.7% on the Semi-Private Evaluation set at our stated public leaderboard \$10k compute limit. A high-compute (172x) o3 configuration scored 87.5%.



TOGGLE ANIMATION

# Probeer het zelf: <https://arcprize.org/play>

	<p>AGI IS IN REACH. NEW IDEAS ARE NEEDED</p>	<ul style="list-style-type: none"><li>&gt; Home</li><li>&gt; ARC-AGI</li><li>&gt; 2024 Results</li><li>&gt; Technical Guide</li><li>&gt; Play</li><li>&gt; Blog</li><li>&gt; 2025 Sign up</li></ul> <p><b>DONATE</b></p>
--	--	--

## PLAY

Try ARC-AGI. Given the examples, identify the pattern, solve the test puzzle.

Puzzle ID: f8be4b64

Daily Puzzle

## DAILY PUZZLE!

Today is March 9, 2025.

Are you ready to try your luck at today's selected ARC-AGI task?



# Nieuwste ontwikkelingen

- Reasoning (CoT sequences)
- Agenten: combinatie van LLMs met programma's die actieve handelingen uitvoeren
- Vooralsnog zijn agenten meer een verkoophype dan een werkbare realiteit, maar dat zal gaandeweg veranderen

# De sterke en zwakke punten van ChatGPT

- Goede tekstuele kwaliteit. Overtuigende “faker”: de “kennis” komt uit teksten, niet uit ervaring
- Geweldig hulpmiddel voor onderwijs, onderzoek, overheid en industrie
- Vereist enig inzicht in de werking en ervaring met “prompts”
- Vertoont hallucinaties (die niet gemakkelijk te vermijden zijn)
- “Kennis” is gebaseerd op de ingevoerde data
- Krachtige desinformatie machine —> noodzaak van open source & development
- Krachtige hackmachine

# Ter afsluiting

- Basale en up-to-date kennis van AI is een noodzaak (regelmatig bijscholen)
- Moratorium op ontwikkeling van AI is een slecht idee, dit is geen “AGI”
- Wel schuilen er gevaren in misbruik en computer code generatie
- Onderken en anticipeer op de gevaren van AI toepassingen
- Onderken de zeer grote positieve bijdragen van AI onderzoek
- De integratie van AI in onze samenleving is een multidisciplinaire uitdaging

**Nobelprijs 2024**



## The Nobel Prize in Chemistry 2024

### Summary

### Laureates

David Baker

Demis Hassabis

John Jumper

Prize announcement

Press release

Popular information

Advanced information



Ill. Niklas Elmehed © Nobel Prize Outreach

**David Baker**

Prize share: 1/2



Ill. Niklas Elmehed © Nobel Prize Outreach

**Demis Hassabis**

Prize share: 1/4



Ill. Niklas Elmehed © Nobel Prize Outreach

**John Jumper**

Prize share: 1/4

Article

# Highly accurate protein structure prediction with AlphaFold

<https://doi.org/10.1038/s41586-021-03819-2>

Received: 11 May 2021

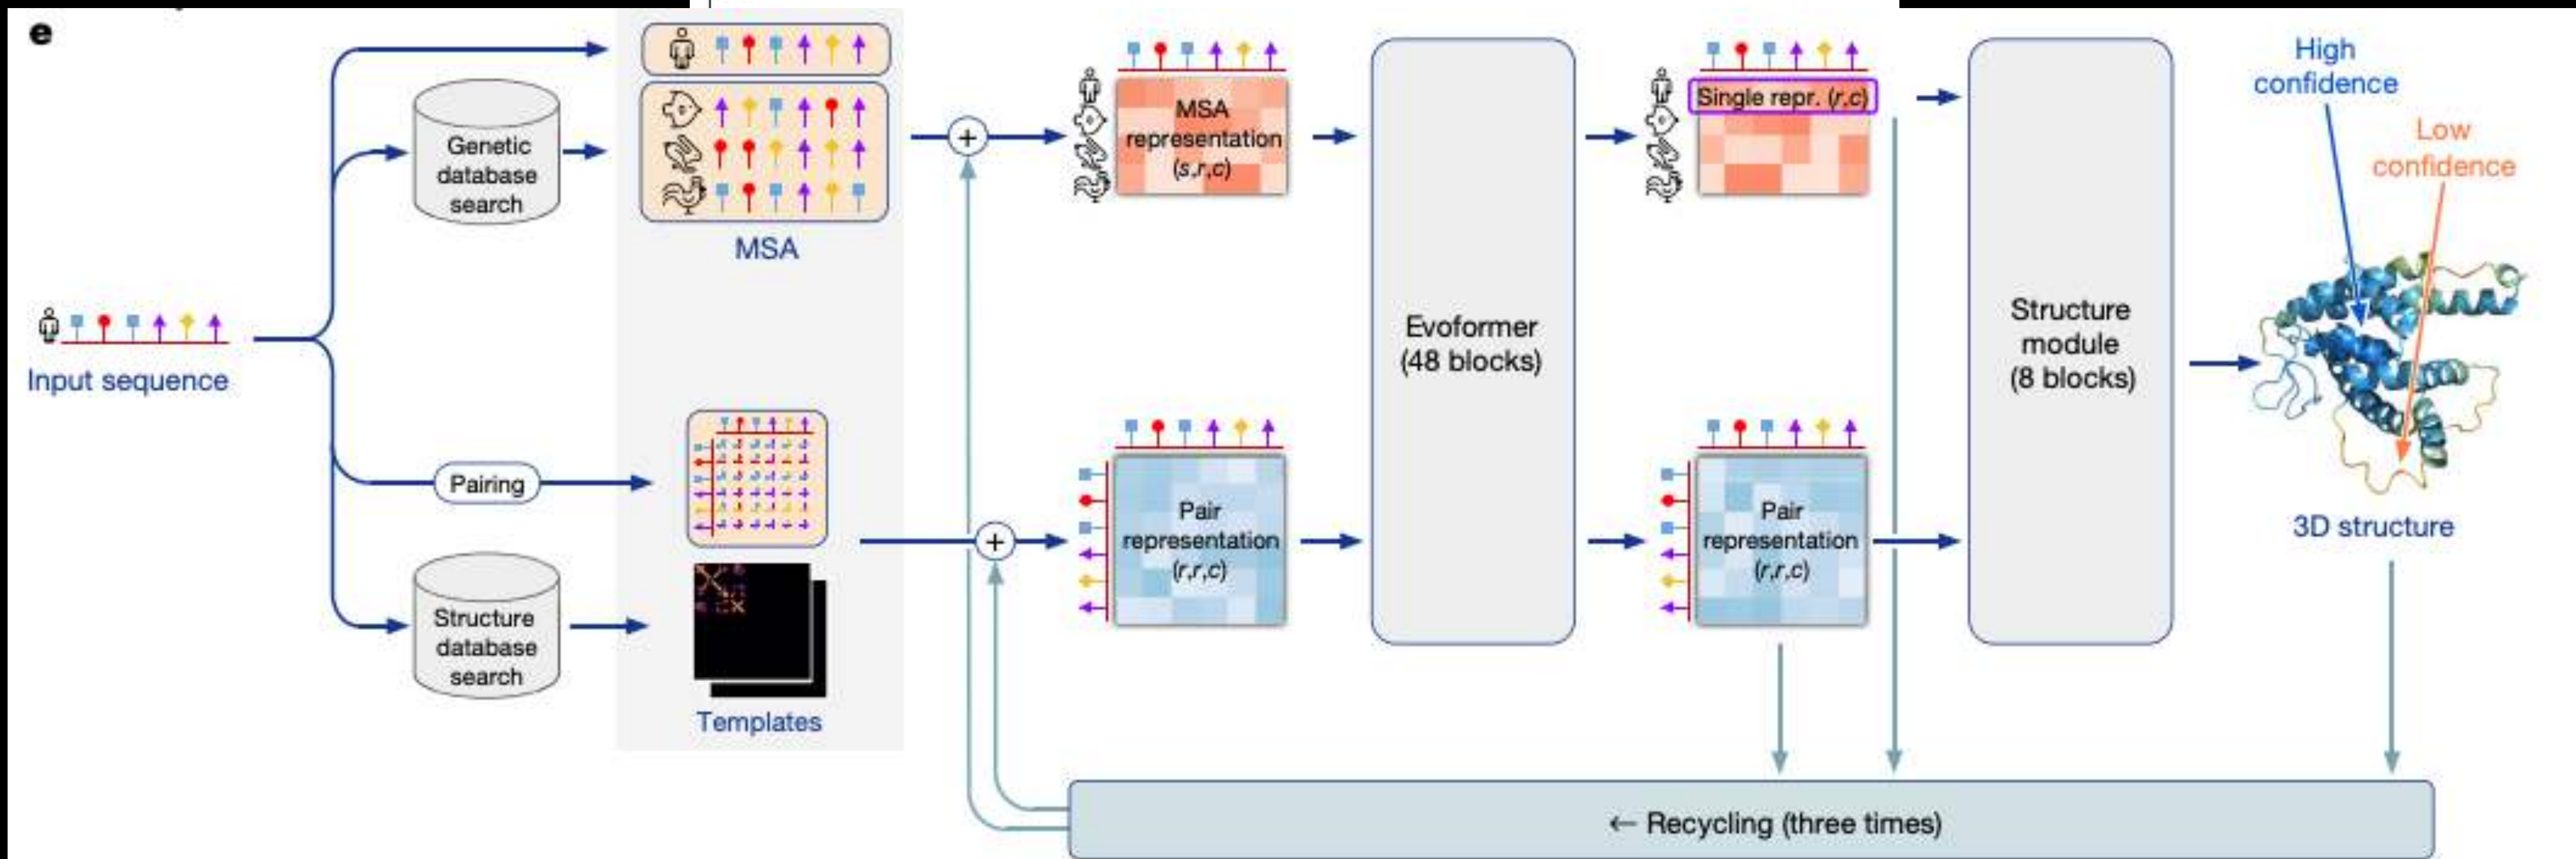
Accepted: 12 July 2021

Published online: 15 July 2021

Open access

 Check for updates

John Jumper<sup>1,4,✉</sup>, Richard Evans<sup>1,4</sup>, Alexander Pritzel<sup>1,4</sup>, Tim Green<sup>1,4</sup>, Michael Figurnov<sup>1,4</sup>,  
Olaf Ronneberger<sup>1,4</sup>, Kathryn Tunyasuvunakool<sup>1,4</sup>, Russ Bates<sup>1,4</sup>, Augustin Židek<sup>1,4</sup>,  
Anna Potapenko<sup>1,4</sup>, Alex Bridgland<sup>1,4</sup>, Clemens Meyer<sup>1,4</sup>, Simon A. A. Kohl<sup>1,4</sup>,  
Andrew J. Ballard<sup>1,4</sup>, Andrew Cowie<sup>1,4</sup>, Bernardino Romera-Paredes<sup>1,4</sup>, Stanislav Nikolov<sup>1,4</sup>,  
Rishub Jain<sup>1,4</sup>, Jonas Adler<sup>1</sup>, Trevor Back<sup>1</sup>, Stig Petersen<sup>1</sup>, David Reiman<sup>1</sup>, Ellen Clancy<sup>1</sup>,  
Michal Zielinski<sup>1</sup>, Martin Steinegger<sup>2,3</sup>, Michalina Pacholska<sup>1</sup>, Tamas Berghammer<sup>1</sup>,  
Sebastian Bodenstein<sup>1</sup>, David Silver<sup>1</sup>, Oriol Vinyals<sup>1</sup>, Andrew W. Senior<sup>1</sup>, Koray Kavukcuoglu<sup>1</sup>,  
Pushmeet Kohli<sup>1</sup> & Demis Hassabis<sup>1,4,✉</sup>



## The Nobel Prize in Physics 2024

---

### Summary

### Laureates

John J. Hopfield

Geoffrey E. Hinton

Prize announcement

Press release

Popular information

Advanced information

---

Share this



Ill. Niklas Elmehed © Nobel Prize  
Outreach

John J. Hopfield



Ill. Niklas Elmehed © Nobel Prize  
Outreach

Geoffrey E. Hinton

# Bijdragen

- Hopfield: “Hopfield netwerk” een neuraal netwerk dat beschreven kan worden als een Ising model (model van een magneet)
- Hinton: “Boltzmann machine” een variant die minder succesvol was dan convolutional neural networks
- Yan LeCunn: bedenker van convolutional neural networks (1989!)



