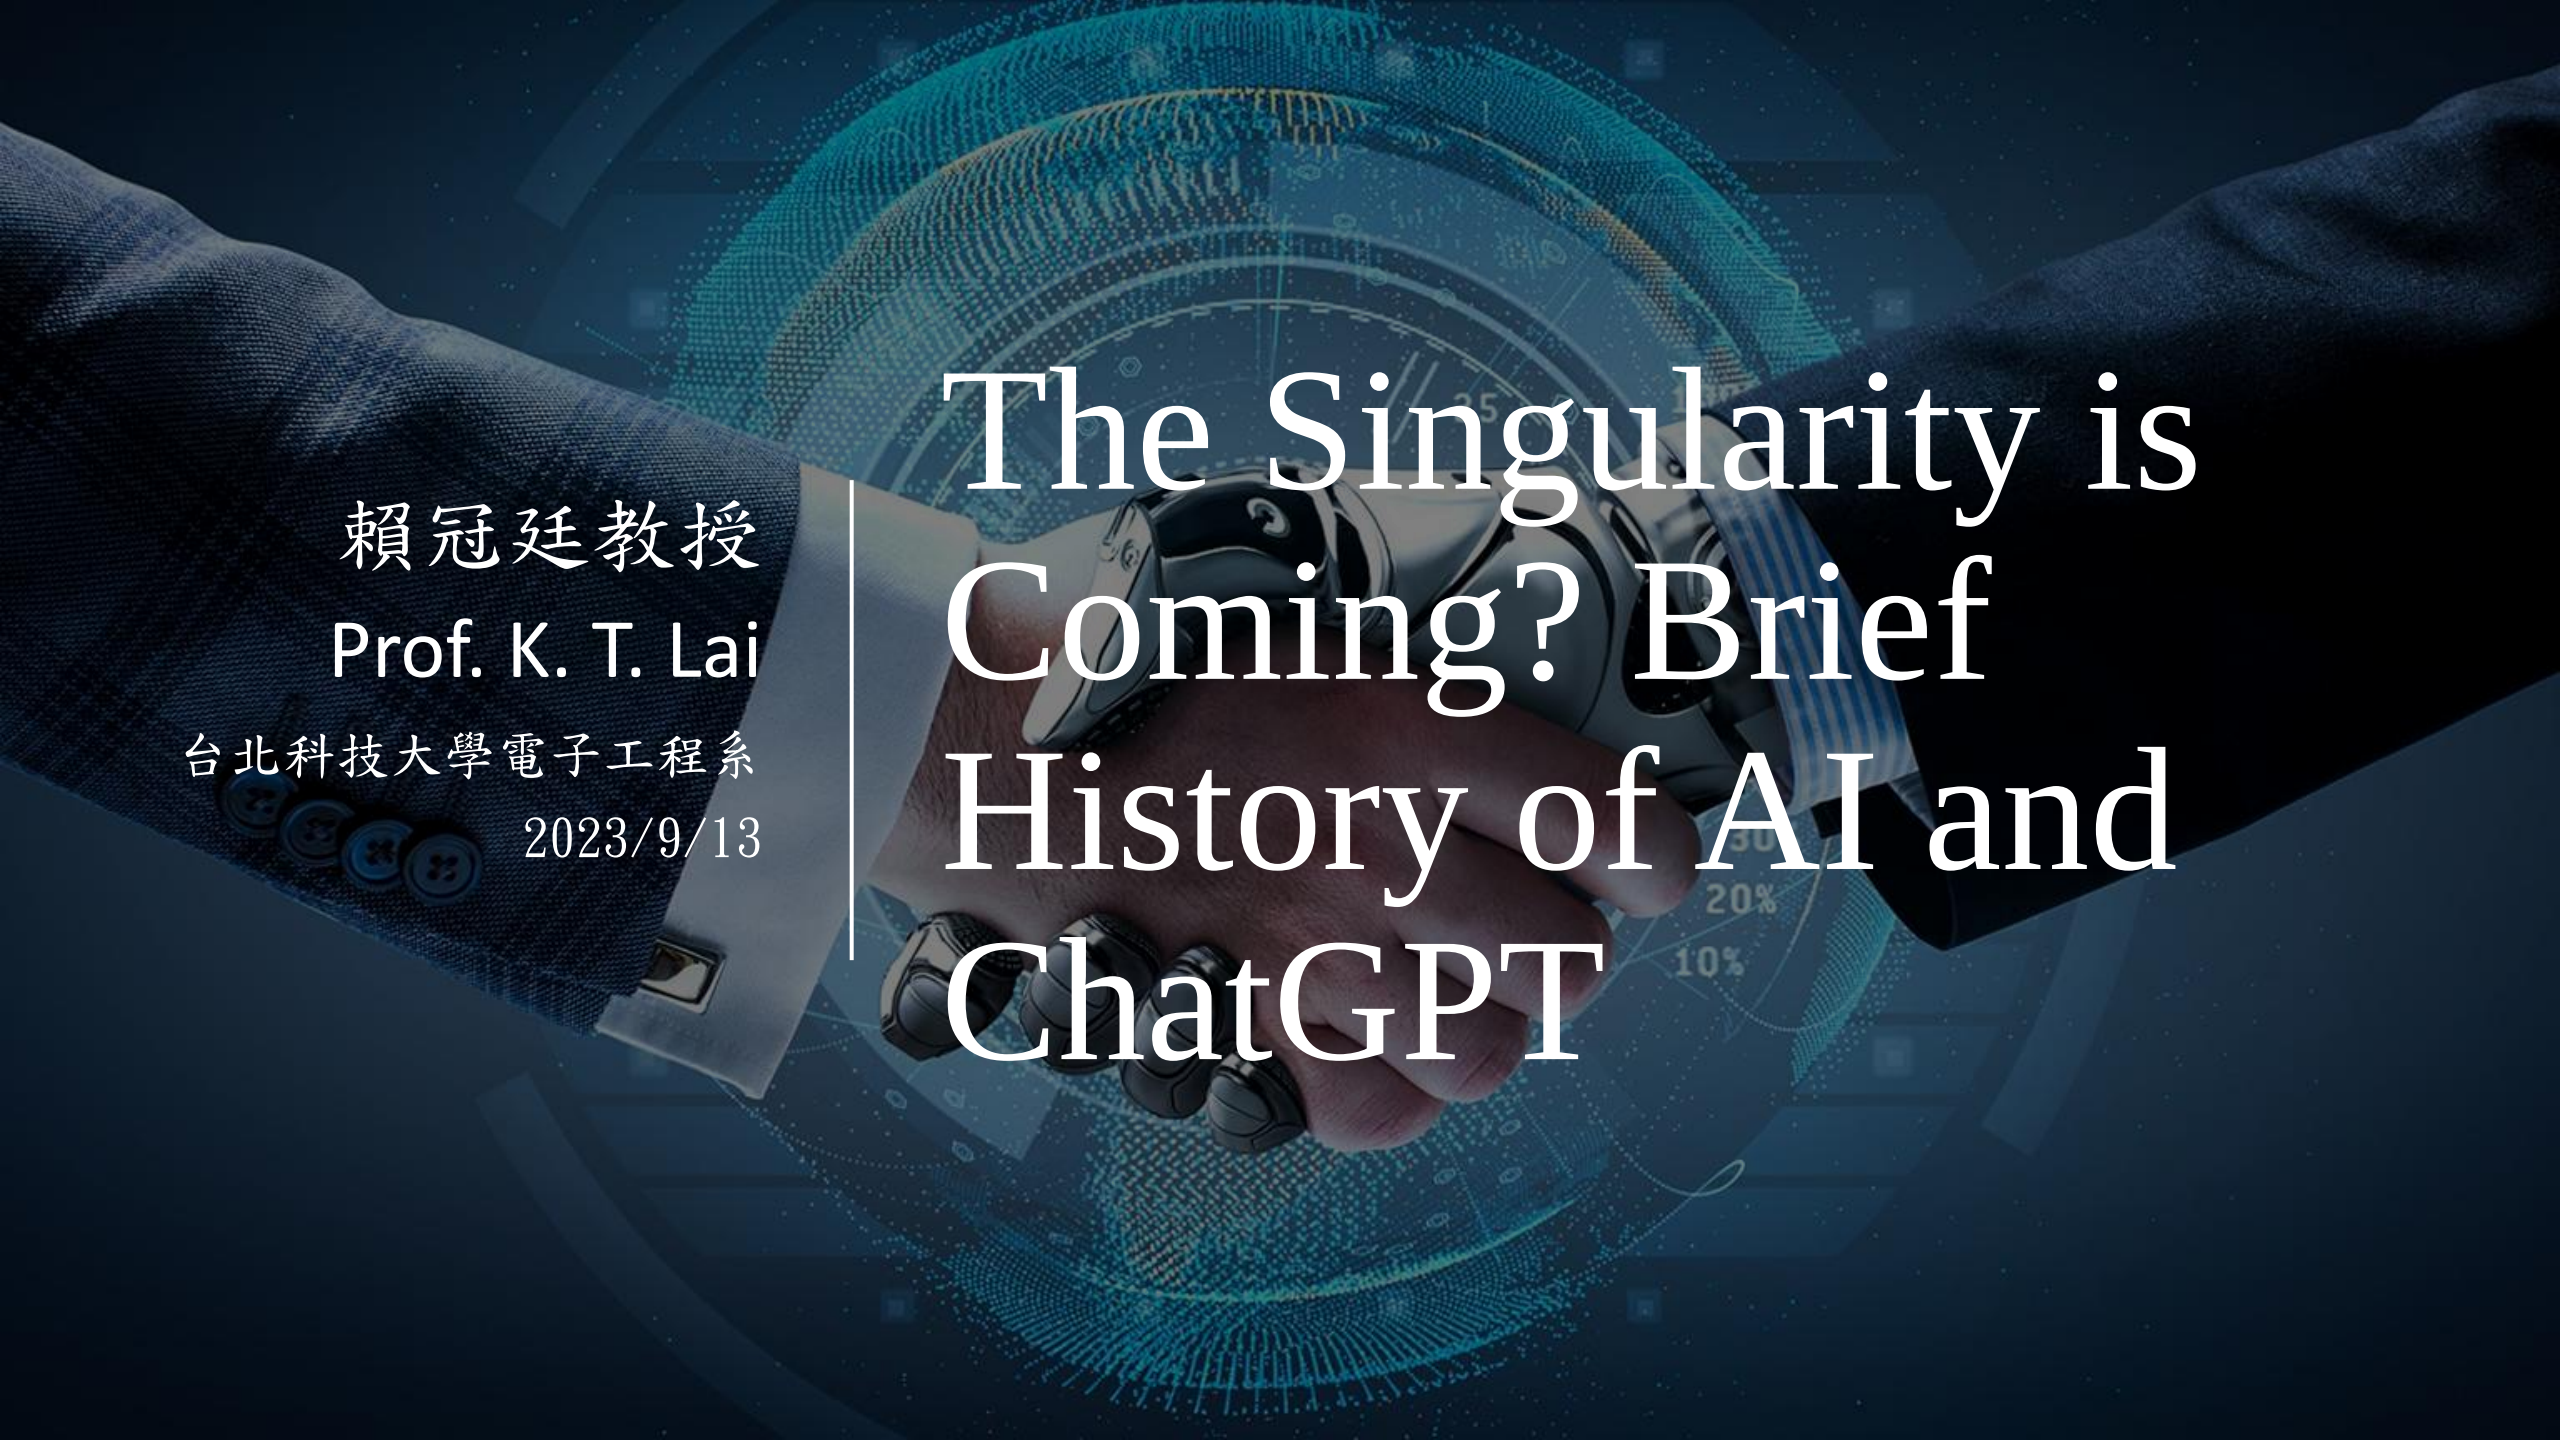

The Singularity is Coming? Brief History of AI and ChatGPT

賴冠廷教授
Prof. K. T. Lai
台北科技大學電子工程系
2023/9/13

Back to the Future

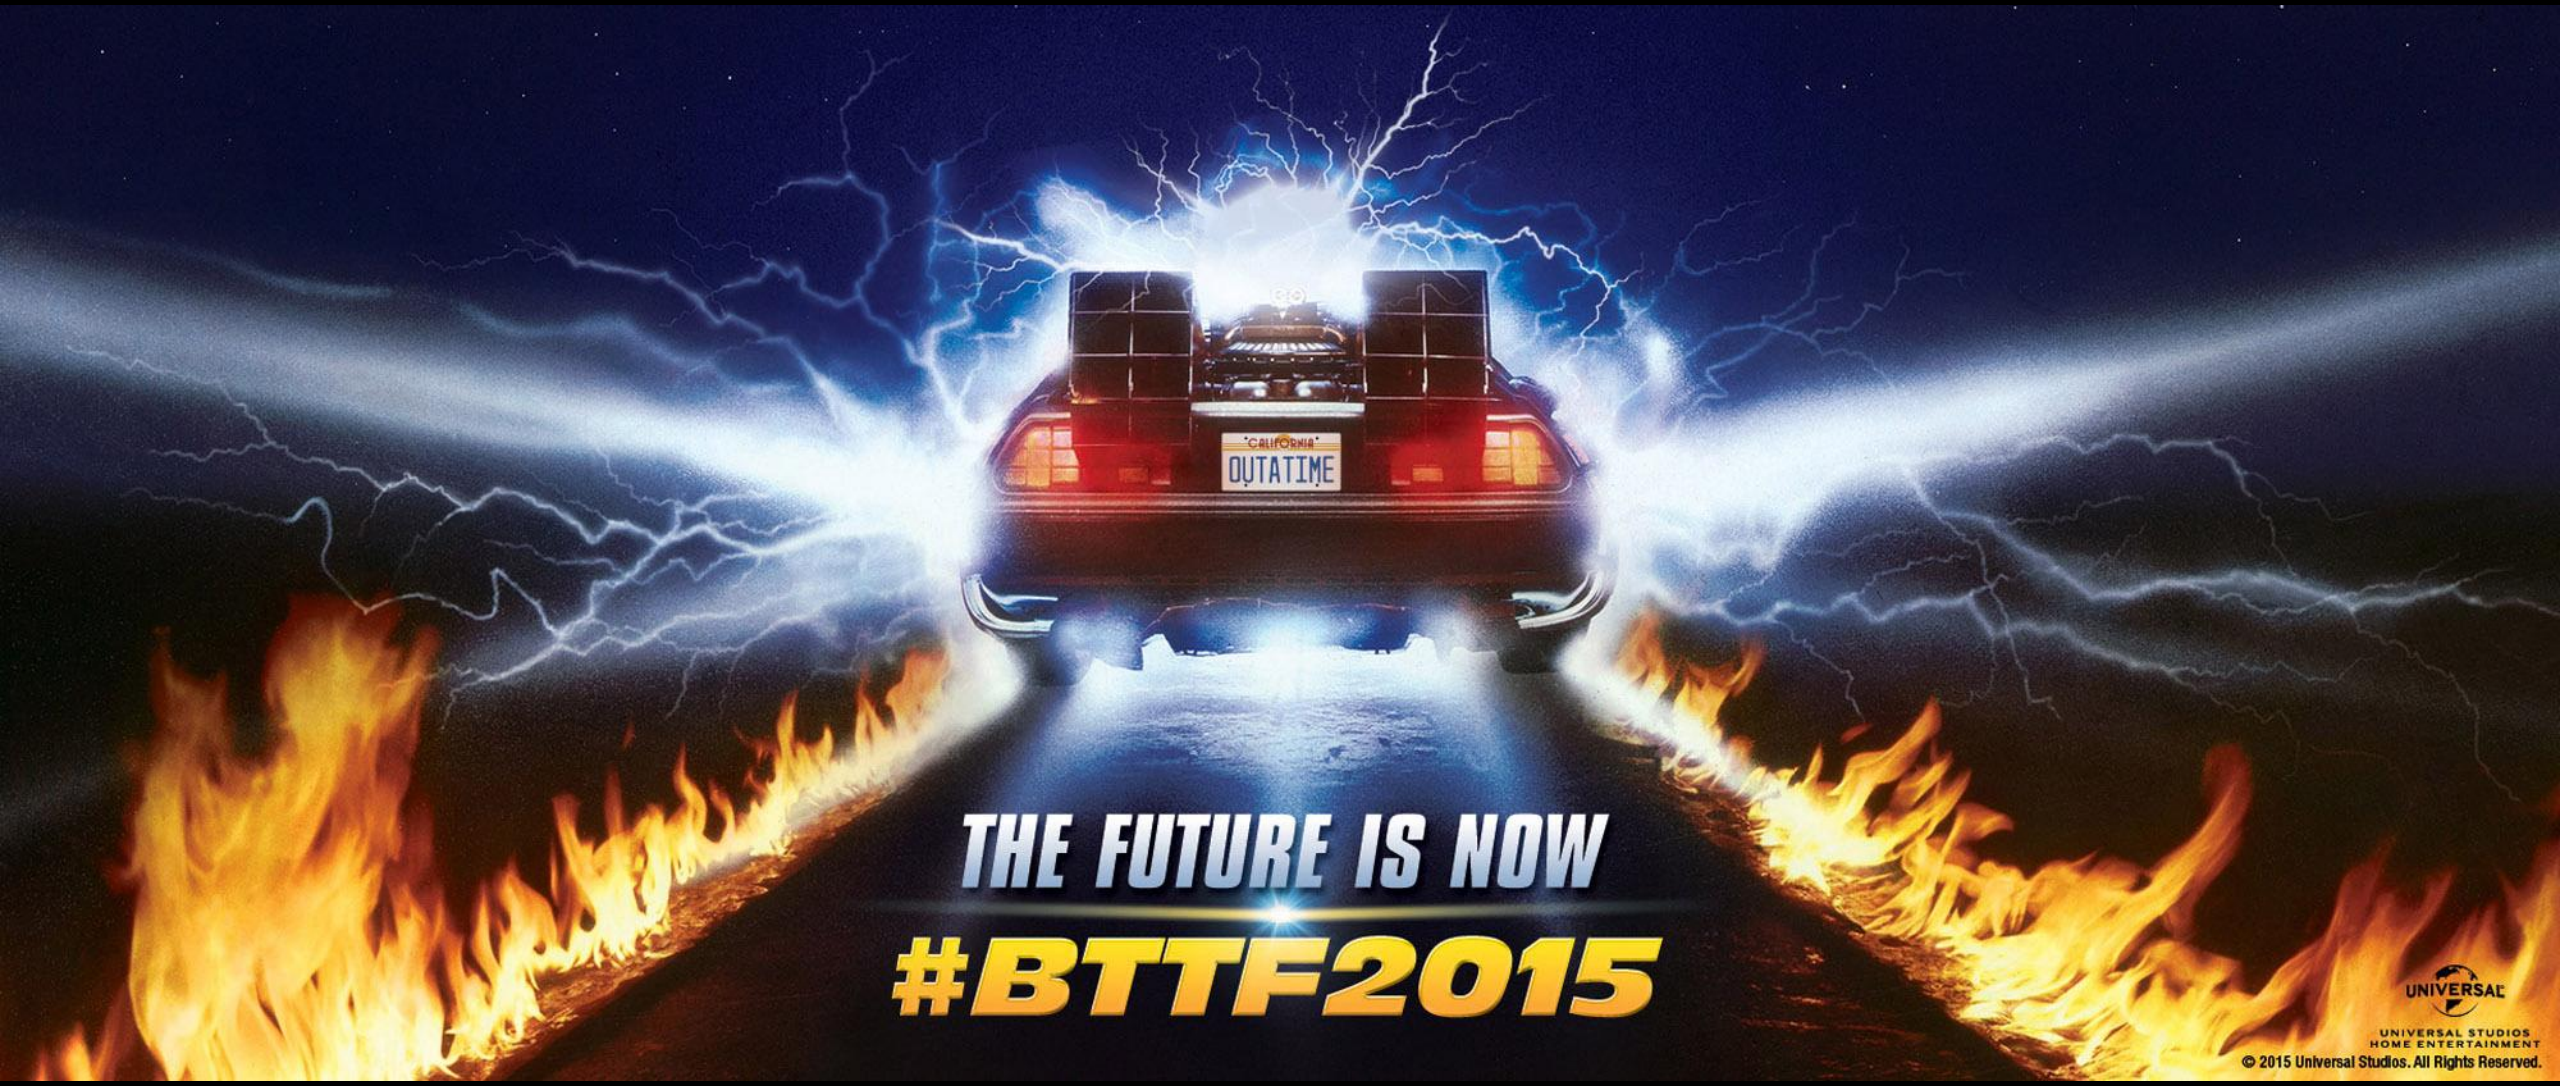

Sci-fi Movie Technologies that are Real Now

https://www.buzzfeed.com/kasiagalazka/science-fiction-things-that-actually-exist-now

1. The Jetsons (1962) – Flatscreen TV

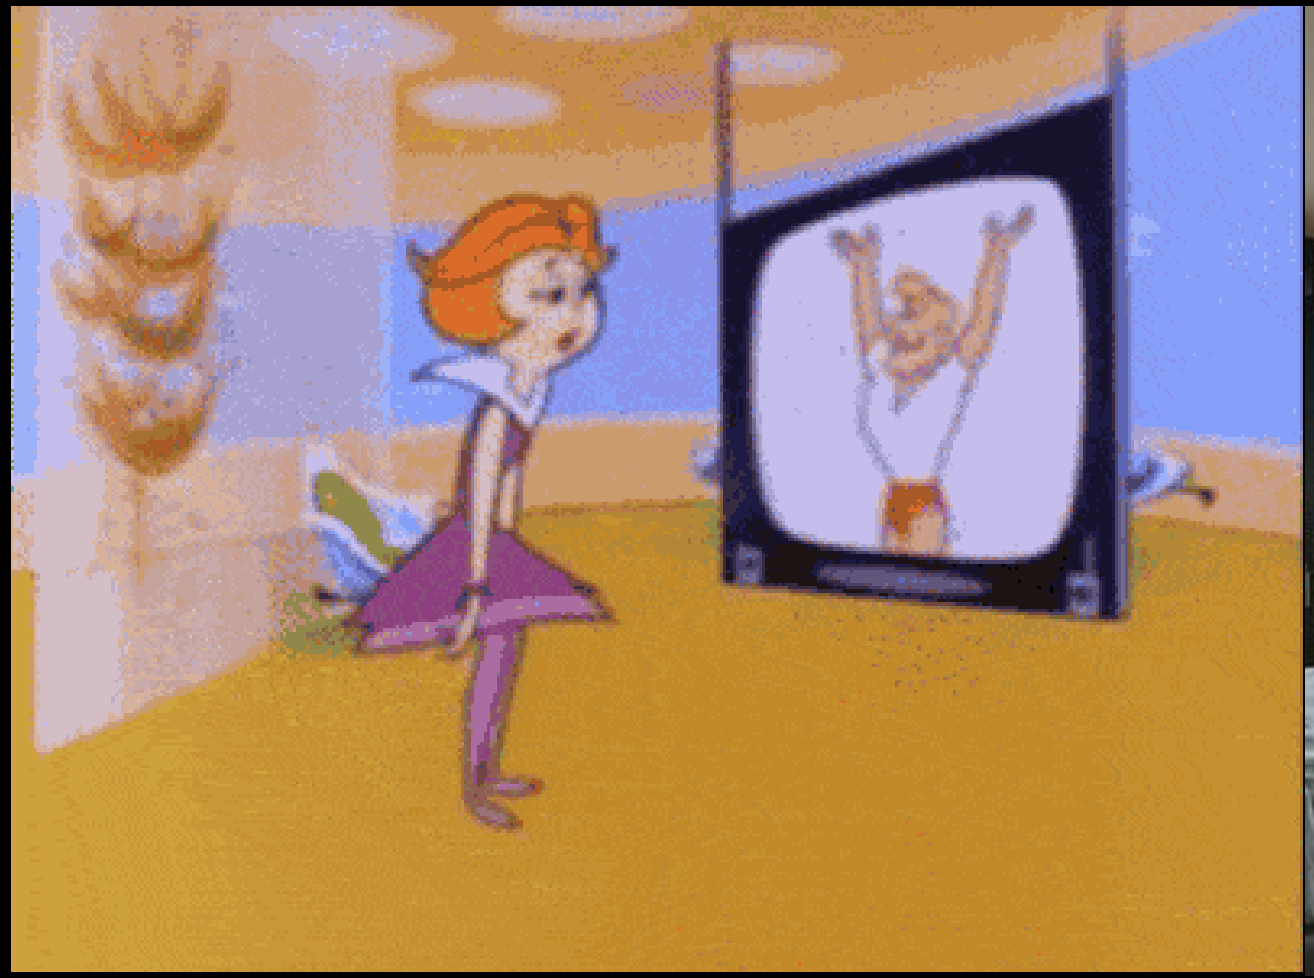

2. The Jetsons (1962) – Smart Watch

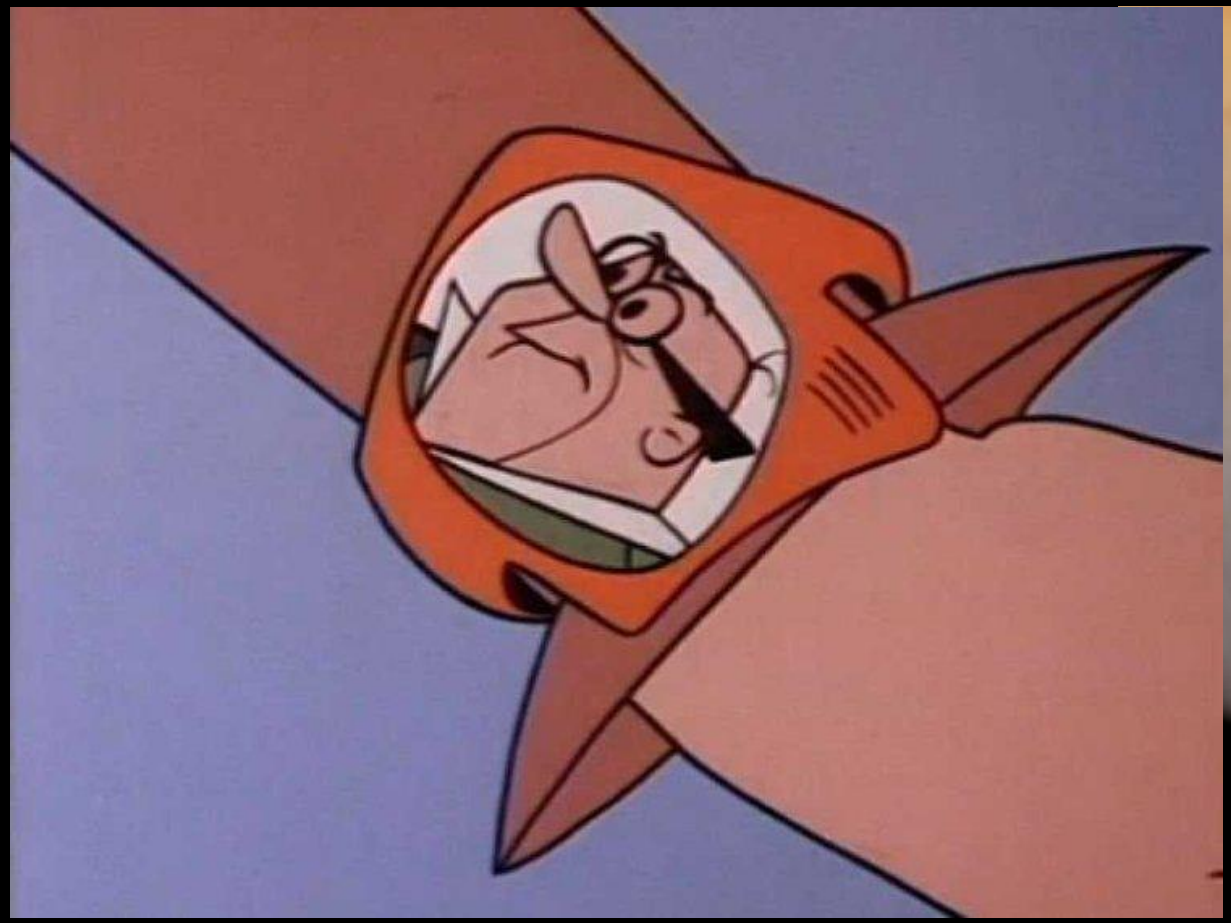

3. The Jetsons (1962) – Roomba

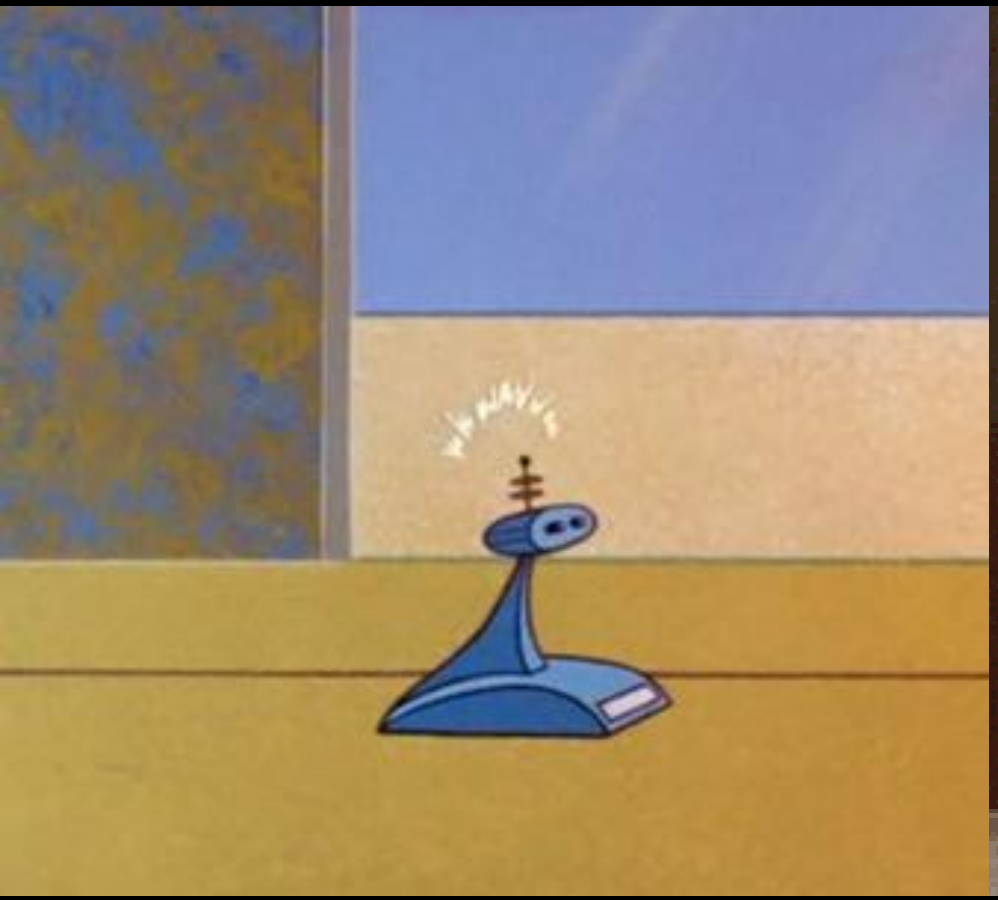

4. Visit to the World's Fair 2014 (1964) – Coffee Maker

5. Star Trek Communicator (1966) – Flip Phone

6. Star Trek Pad (1968) – iPad

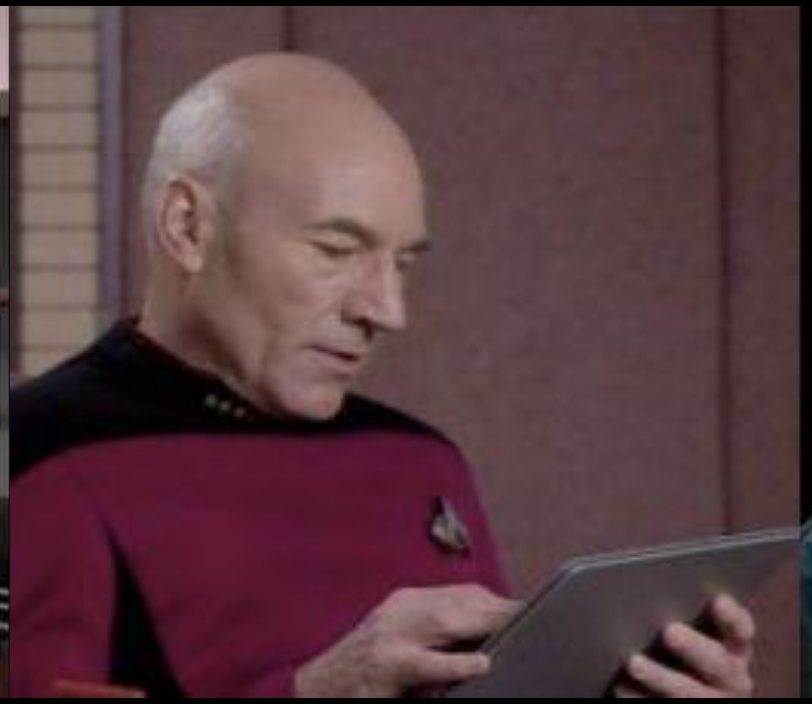

7. Star Trek (1993) – Google Glasses

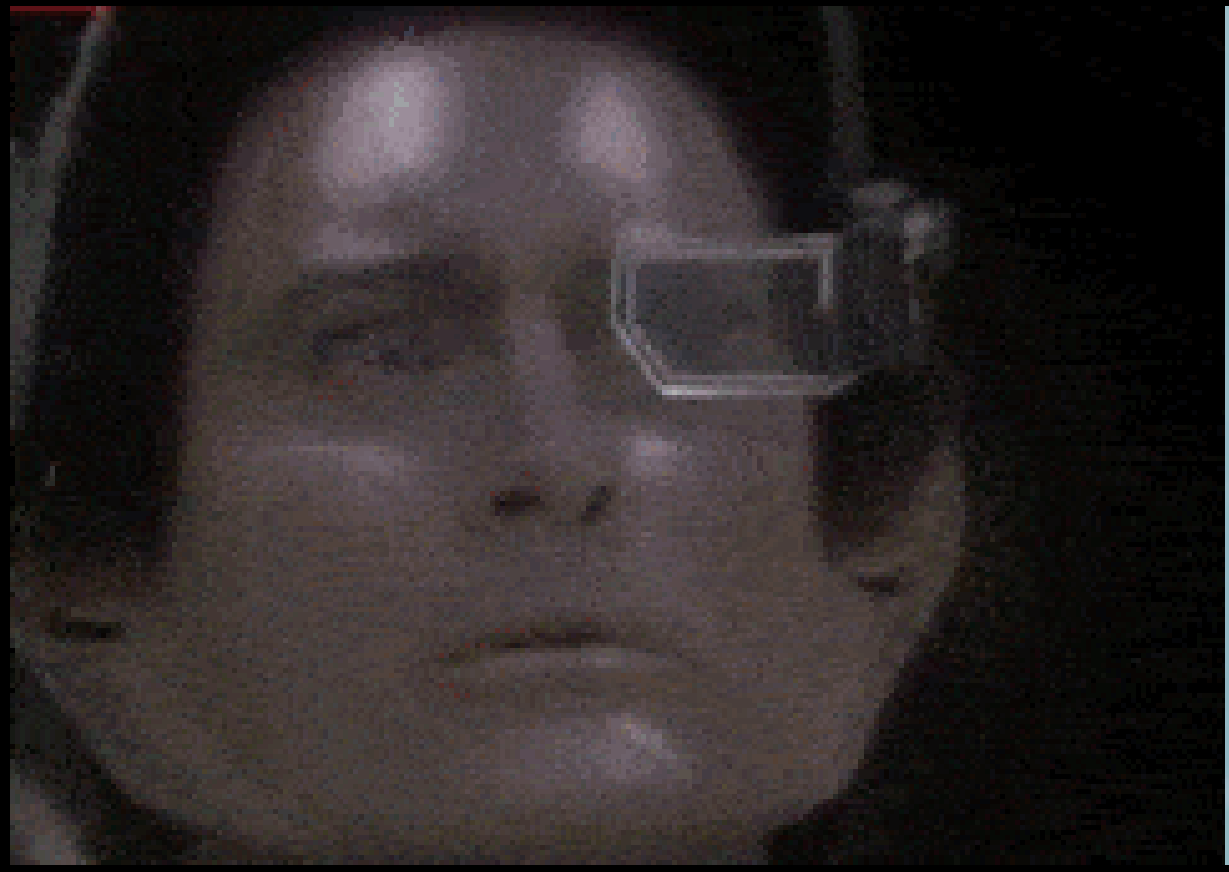

8. 2001 Space Odyssey (1968) – Video Conf.

9. 2001 Space Odyssey (1968) – HAL 9000

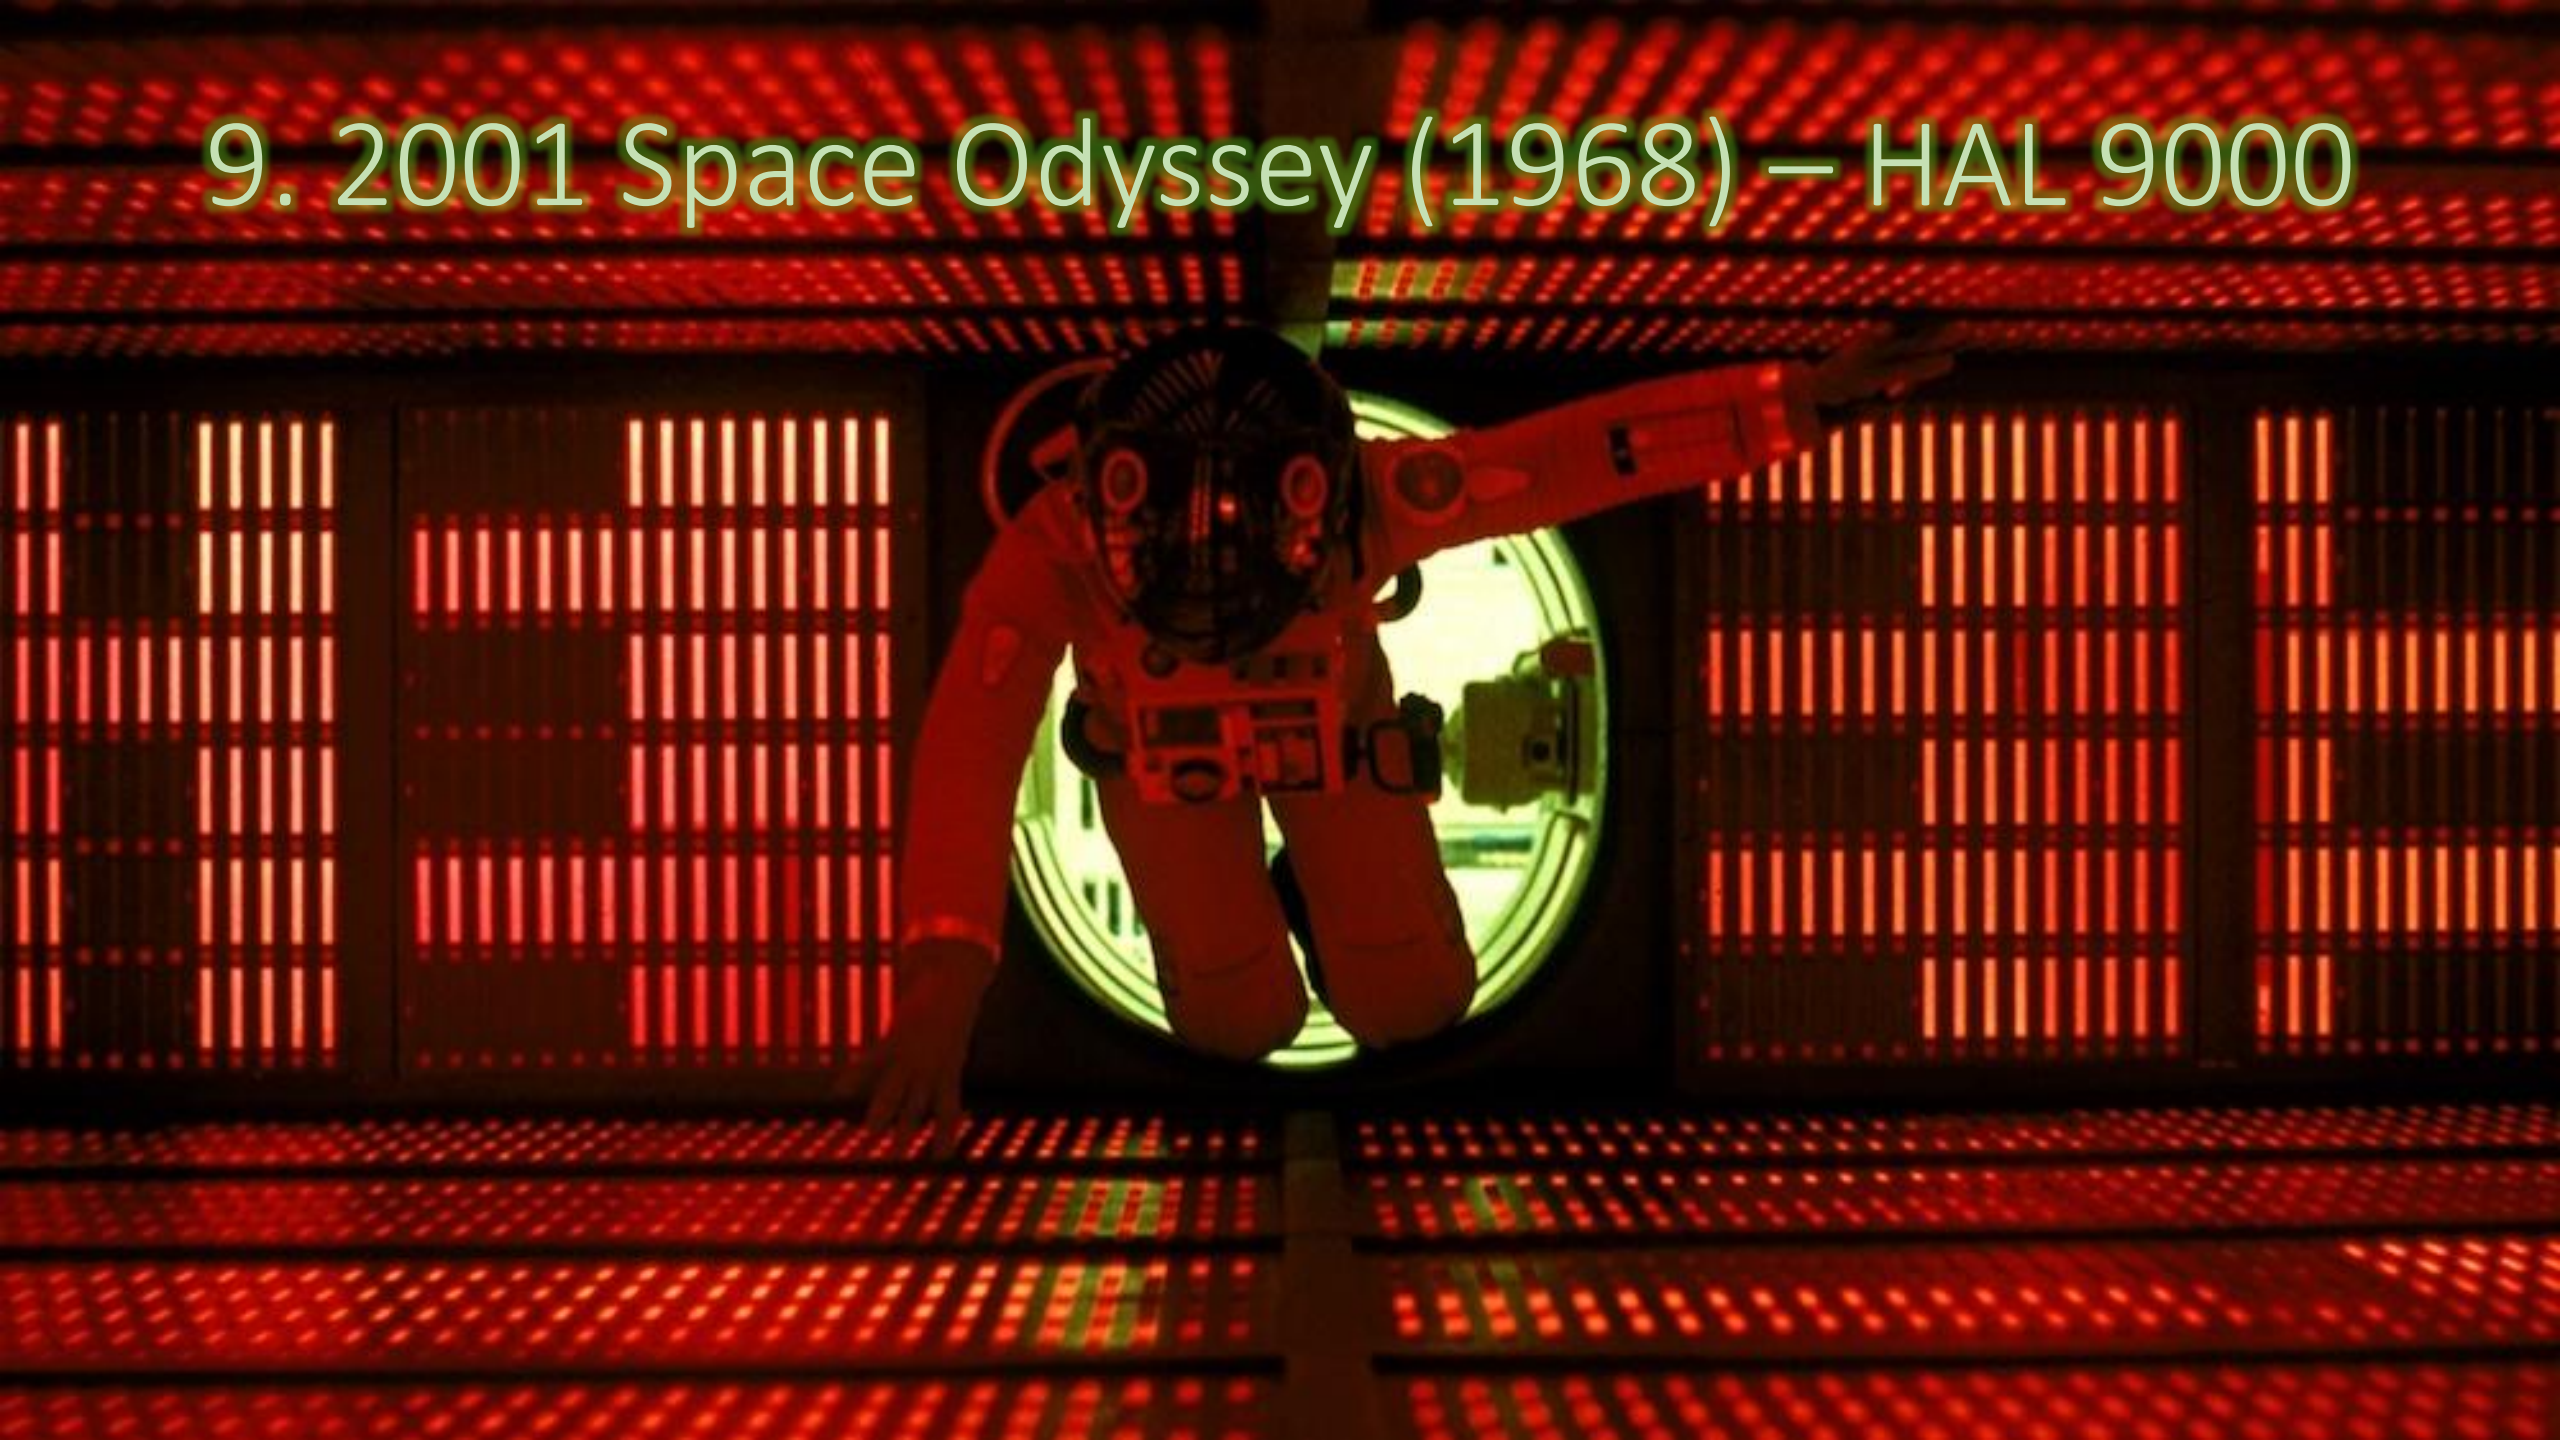

BEZOS - BLUE ORIGIN - BRANSON VIRGIN GALACTIC - MUSK SPACEX

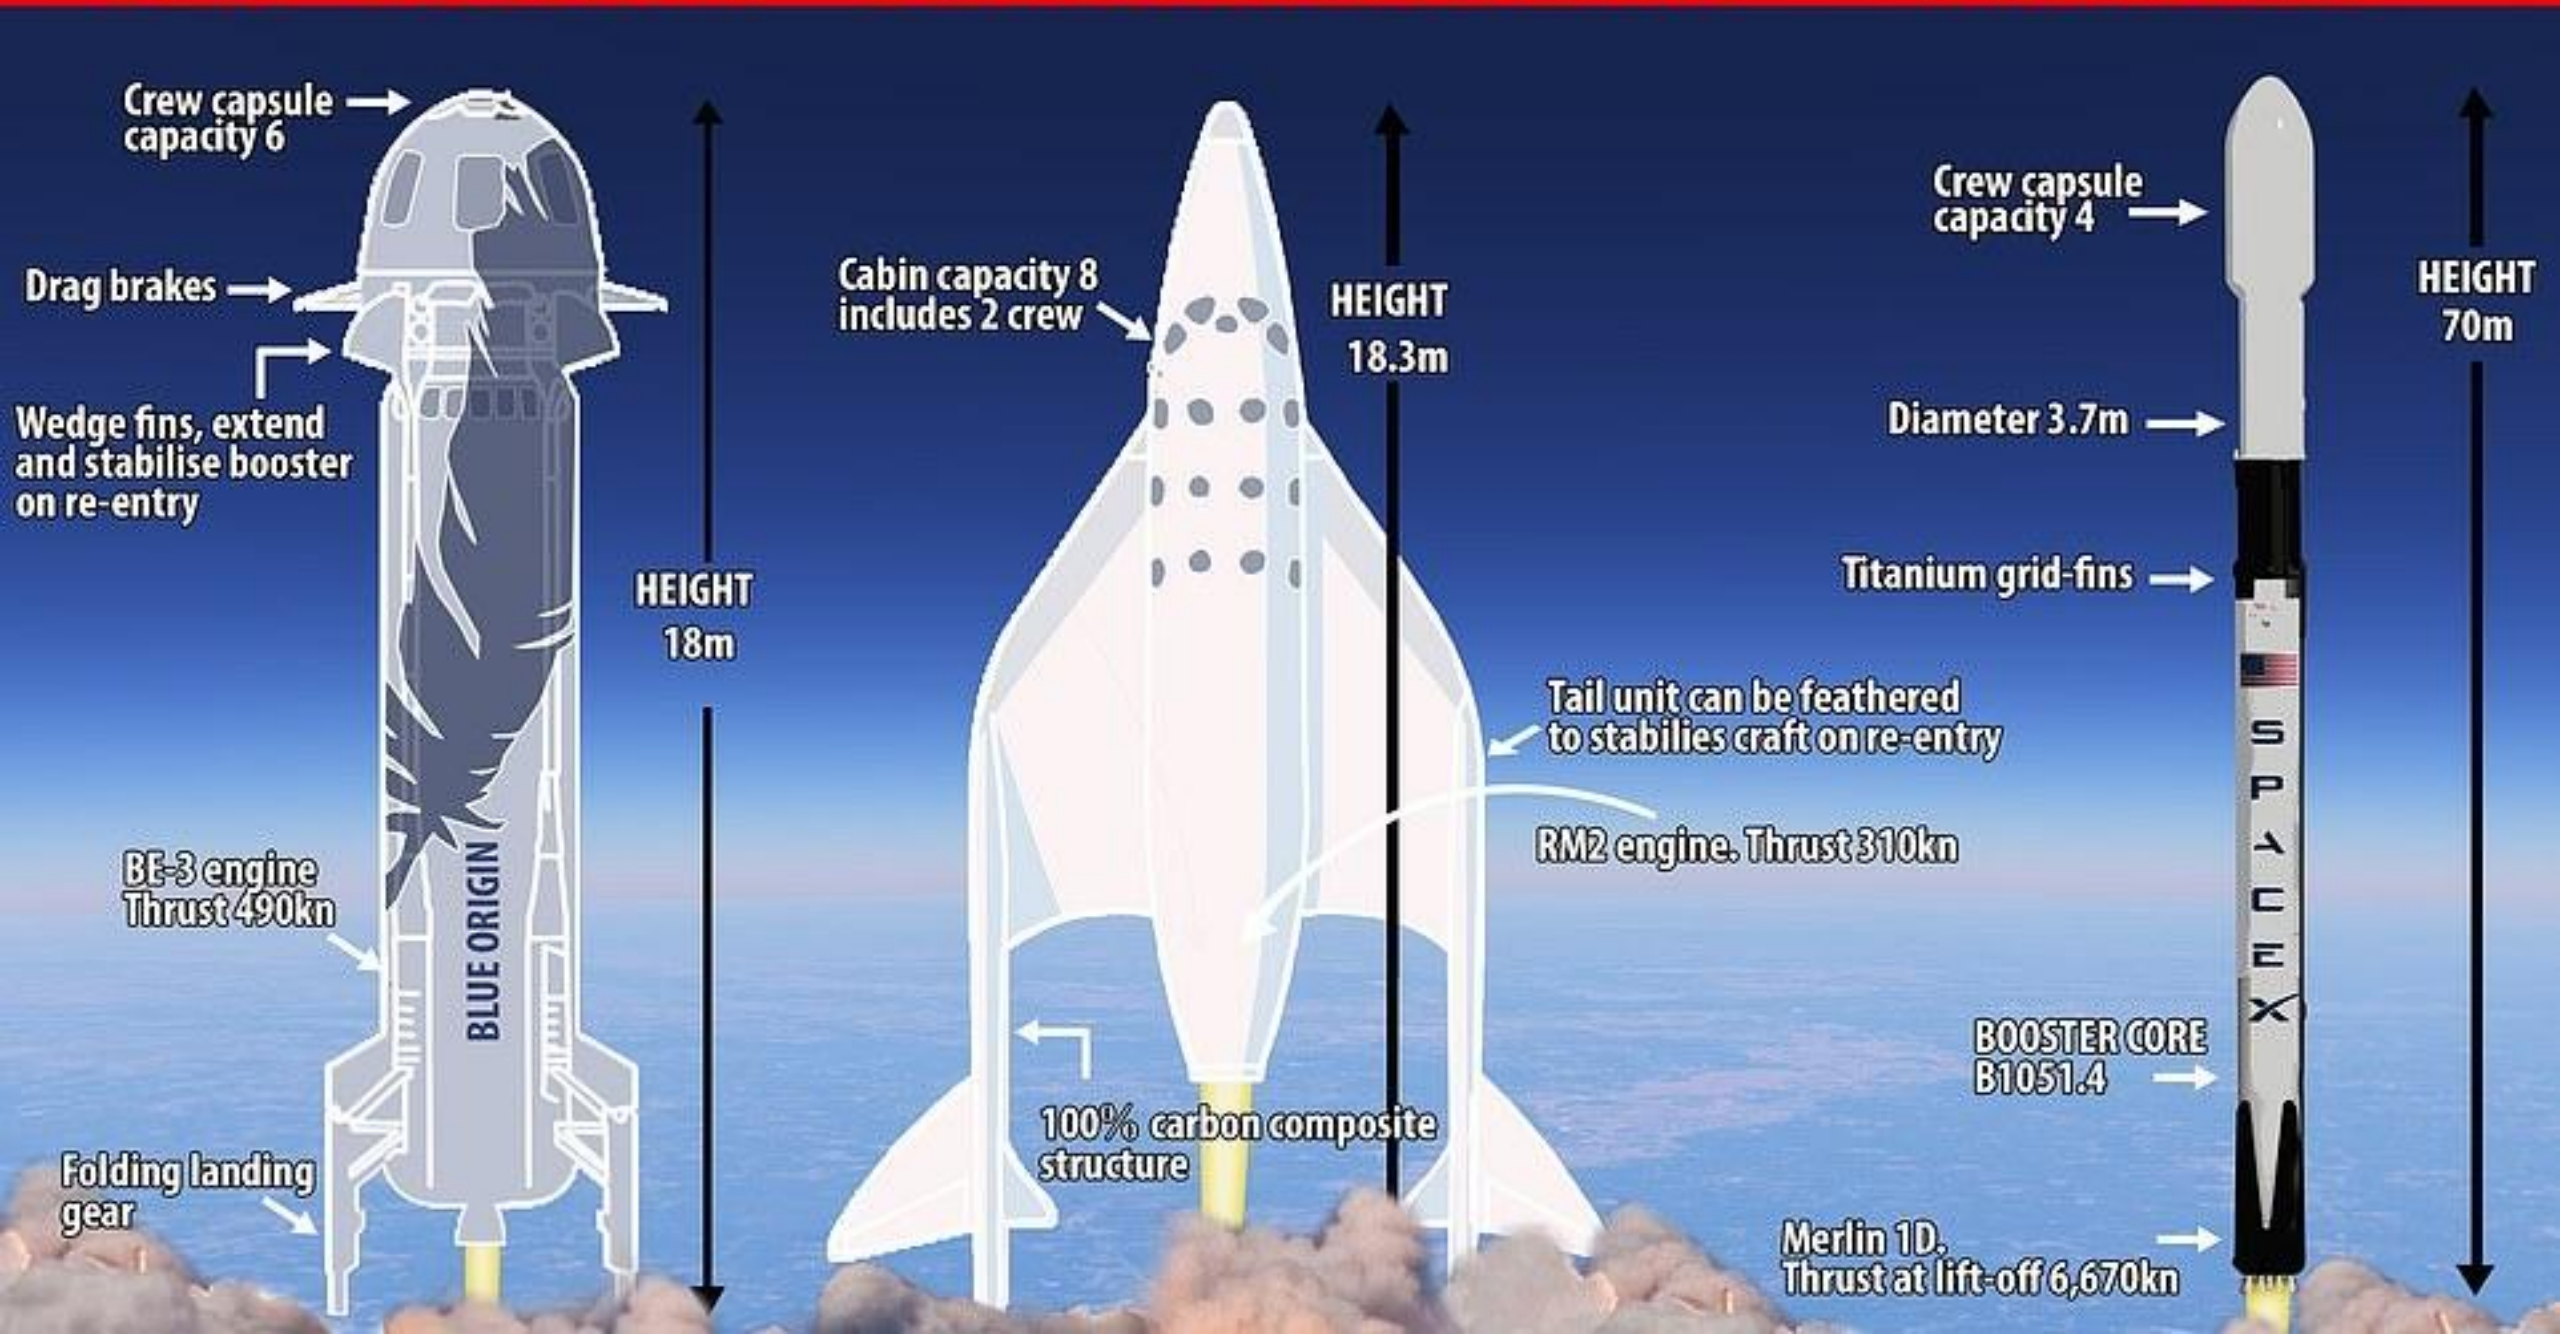


looking up to the stars

10. Knight Rider (1985) – K.I.T.T.

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

Tesla Full Self-Driving (FSD)

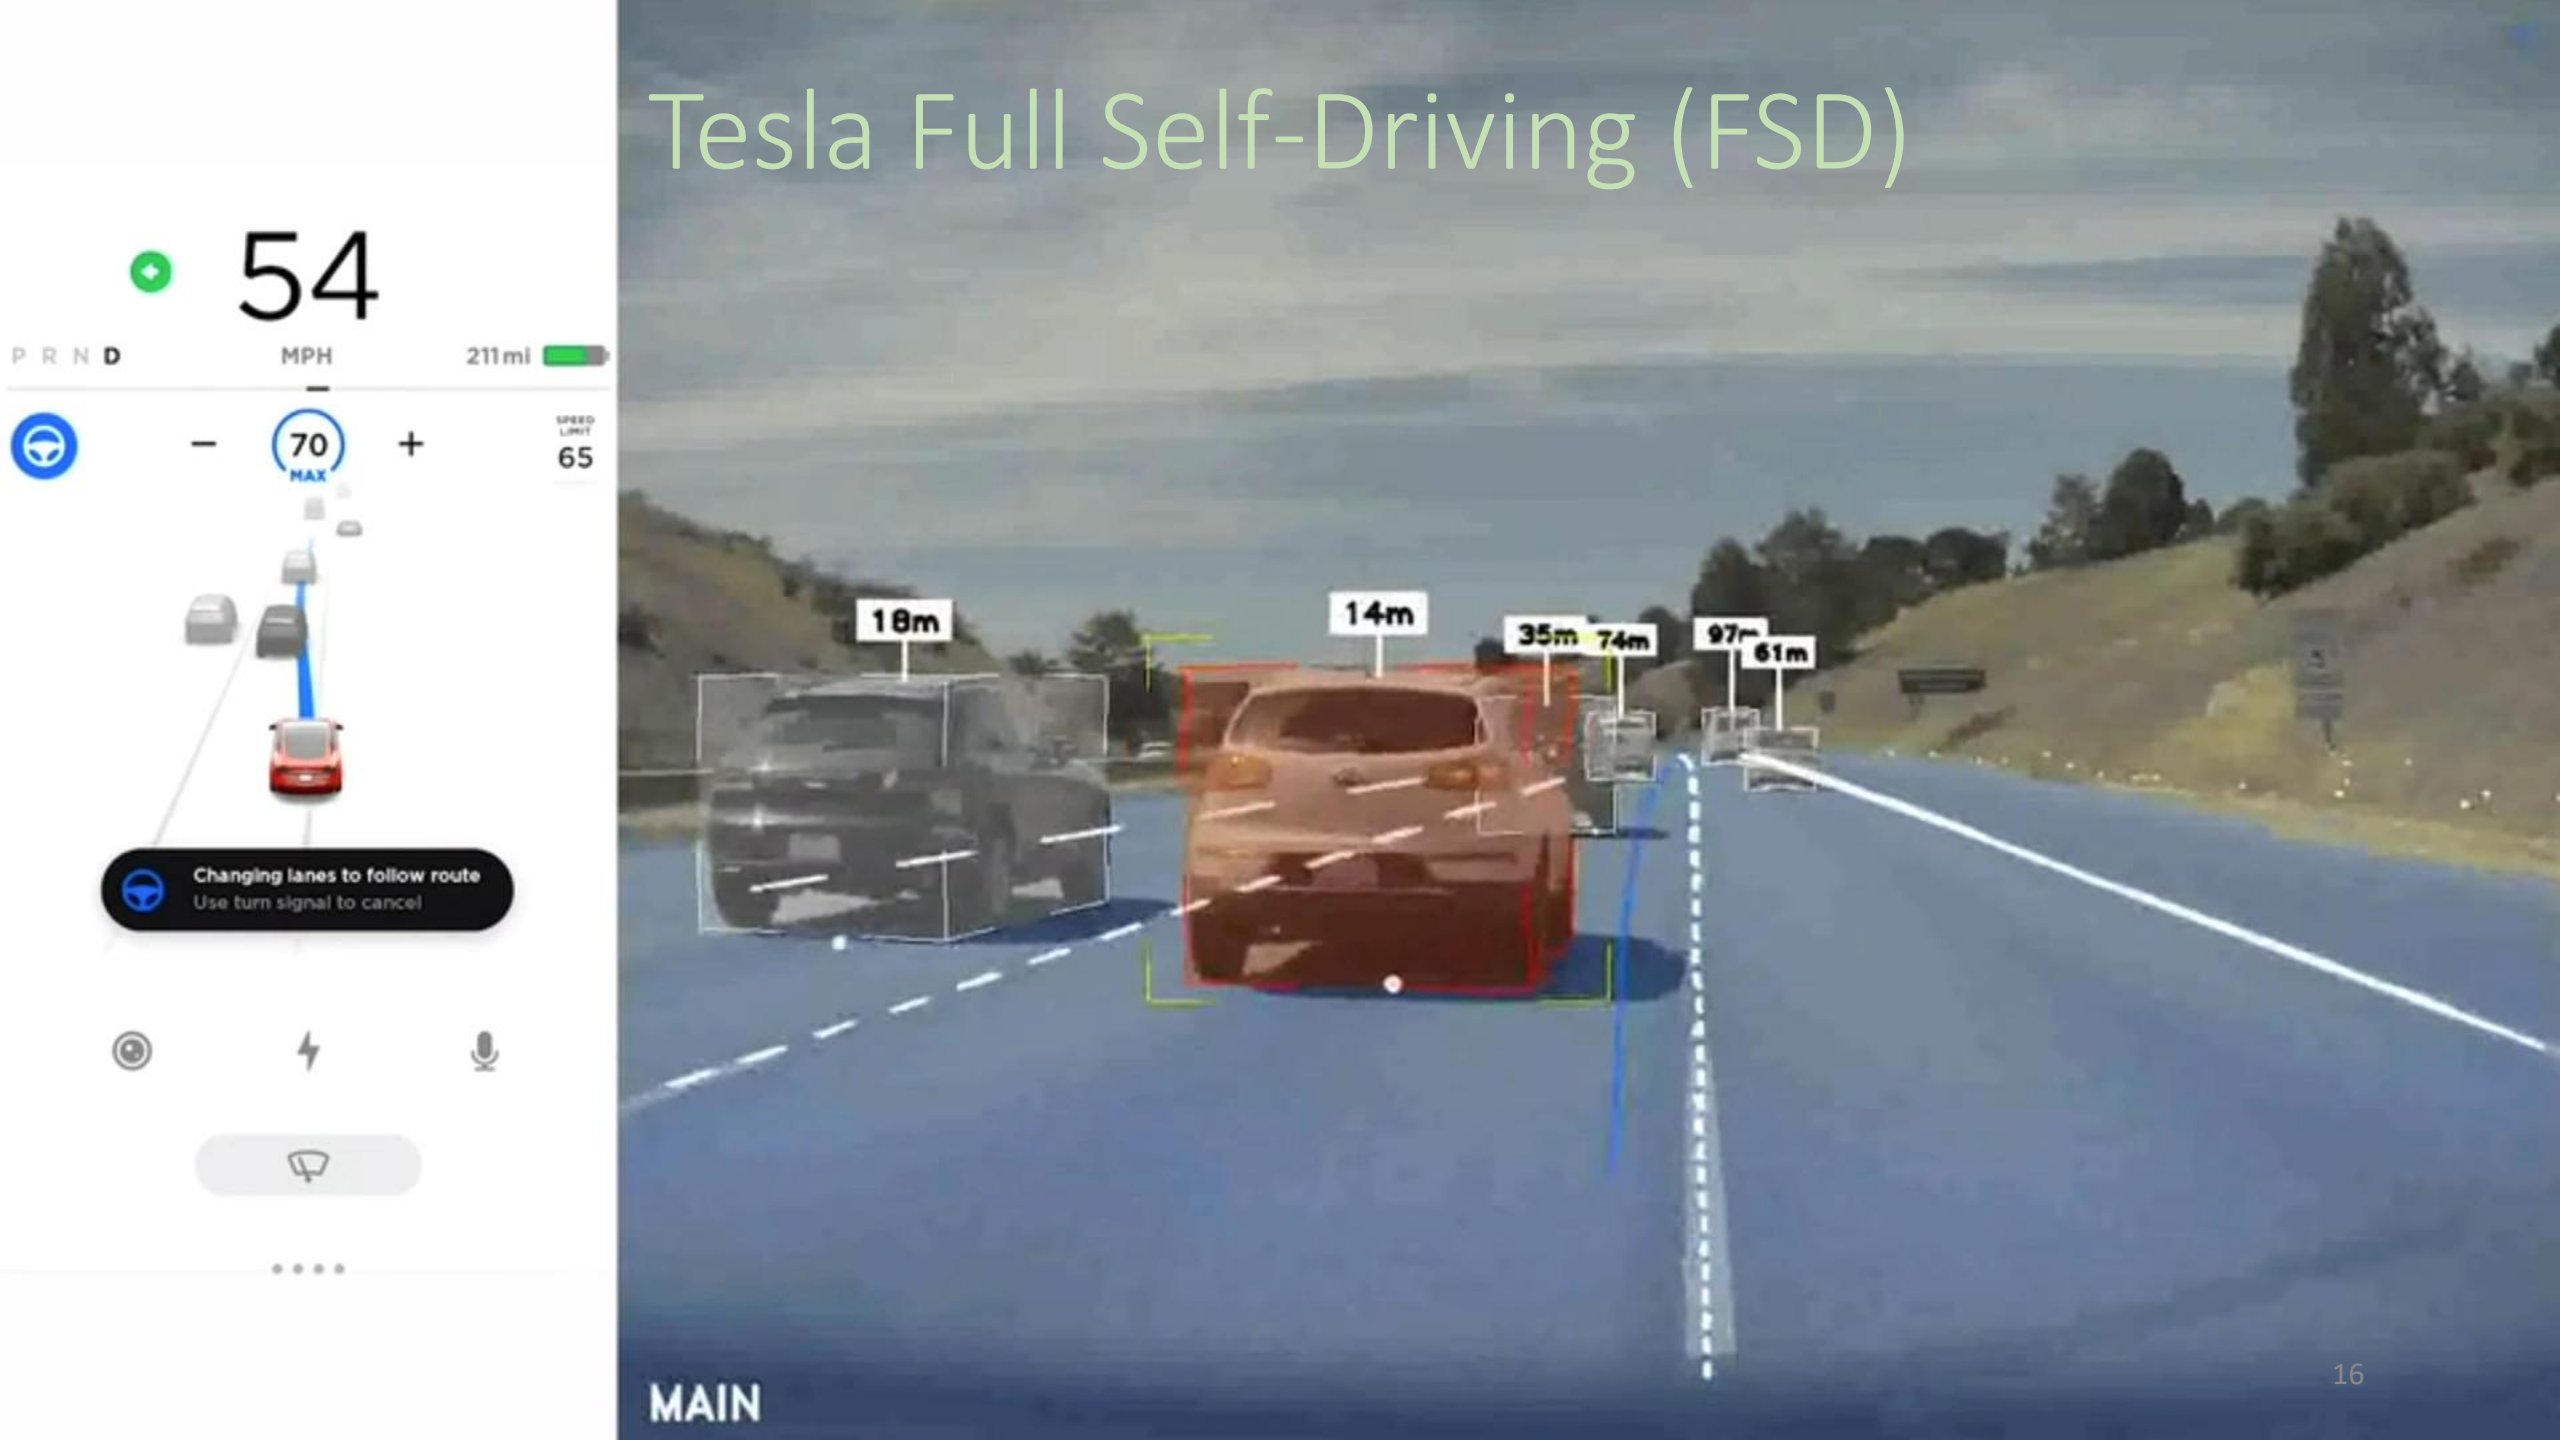

https://www.hk01.com/數碼生活/867807/chatgpt-ai作畫早40年前已預言-多啦a夢6件-神預言-科技產物

What do we have now?

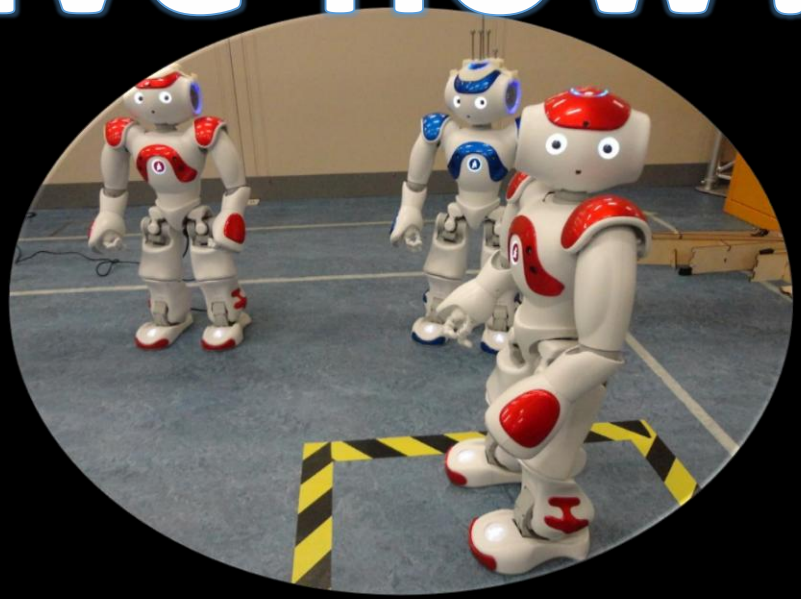


BostonDynamics

Boston Dynamics Spot


Boston Dynamics
CNBC
The deal valued Boston Dynamics at $1.1 billion

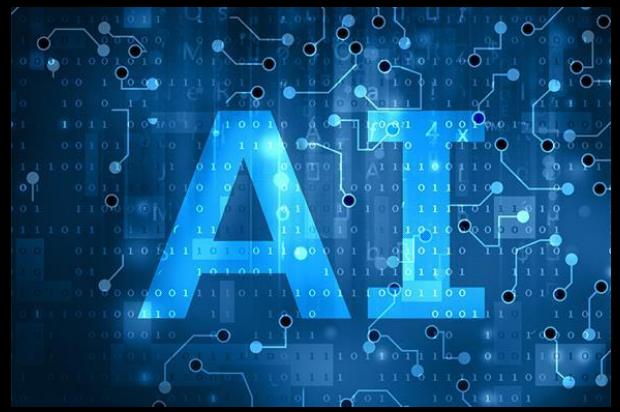
AI


IoT


Boston Dynamics

The Singularity

I. J. Good

1 The accelerating pace of change ...

Agricultural Revolution ← 8,000 years → Industrial Revolution ← 120 years → Light-bulb ← 90 years → Moon landing ← 22 years → World Wide Web ← 9 years → Human genome sequenced

2 ... and exponential growth in computing power ...

Computer technology, shown here climbing dramatically by powers of 10, is now progressing more each hour than it did in its entire first 90 years

3 ... will lead to the Singularity

2045
Surpasses brainpower equivalent to that of all human brains combined

Surpasses brainpower of human in 2023

Surpasses brainpower of mouse in 2015

COMPUTER RANKINGS
By calculations per second per $1,000

Analytical engine
Never fully built, Charles Babbage's invention was designed to solve computational and logical problems

UNIVAC I
The first commercially marketed computer, used to tabulate the U.S. Census, occupied 943 cu. ft.

Colossus
The electronic computer, with 1,500 vacuum tubes, helped the British crack German codes during WW II

Apple II
At a price of $1,298, the compact machine was one of the first massively popular personal computers

Power Mac G4
The first personal computer to deliver more than 1 billion floating-point operations per second

10^{26}
10^{20}
10^{15}
10,000,000,000
100,000
1
0.00001

Hollerith Tabulator
IBM Tabulator
National Ellis 3000
Zuse 2
Zuse 3
ENIAC
BINAC
IBM SSEC
Whirlwind
EDVAC
DEC PDP-4
IBM 1130
Datamatic 1000
IBM 1620
DEC PDP-10
Data General Nova
Intellec-8
IBM PC
Compaq Deskpro 386
Pentium PC
Pentium II PC
Dell Dimension 8400
Mac Pro
Nvidia Tesla GPU & PC

ELECTROMECHANICAL → RELAYS → VACUUM TUBES → TRANSISTORS → INTEGRATED CIRCUITS

1900 1920 1940 1960 1980 2000 2011 2020 2045

120 Years of Moore's Law

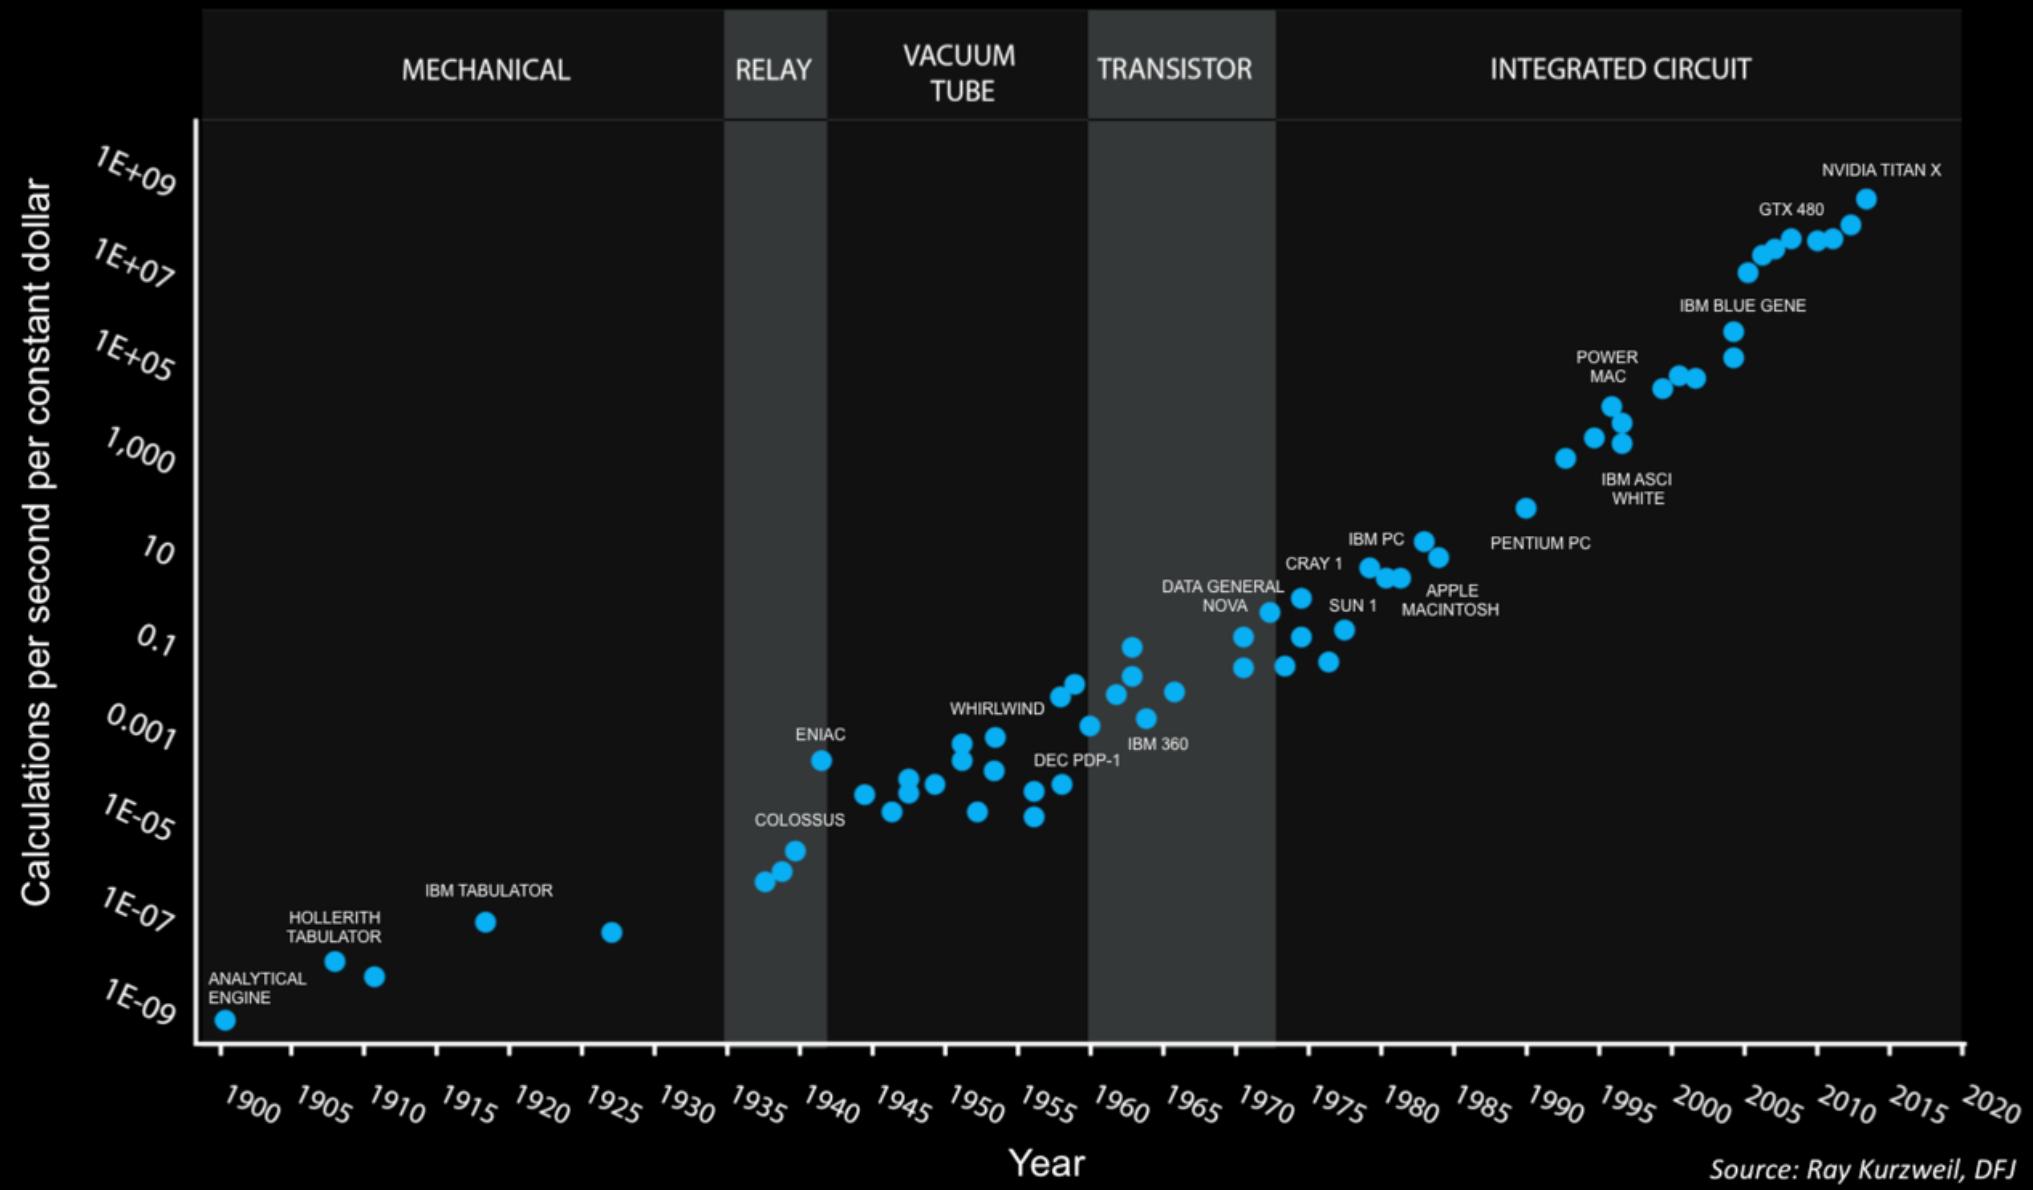

Source: Ray Kurzweil, DFJ

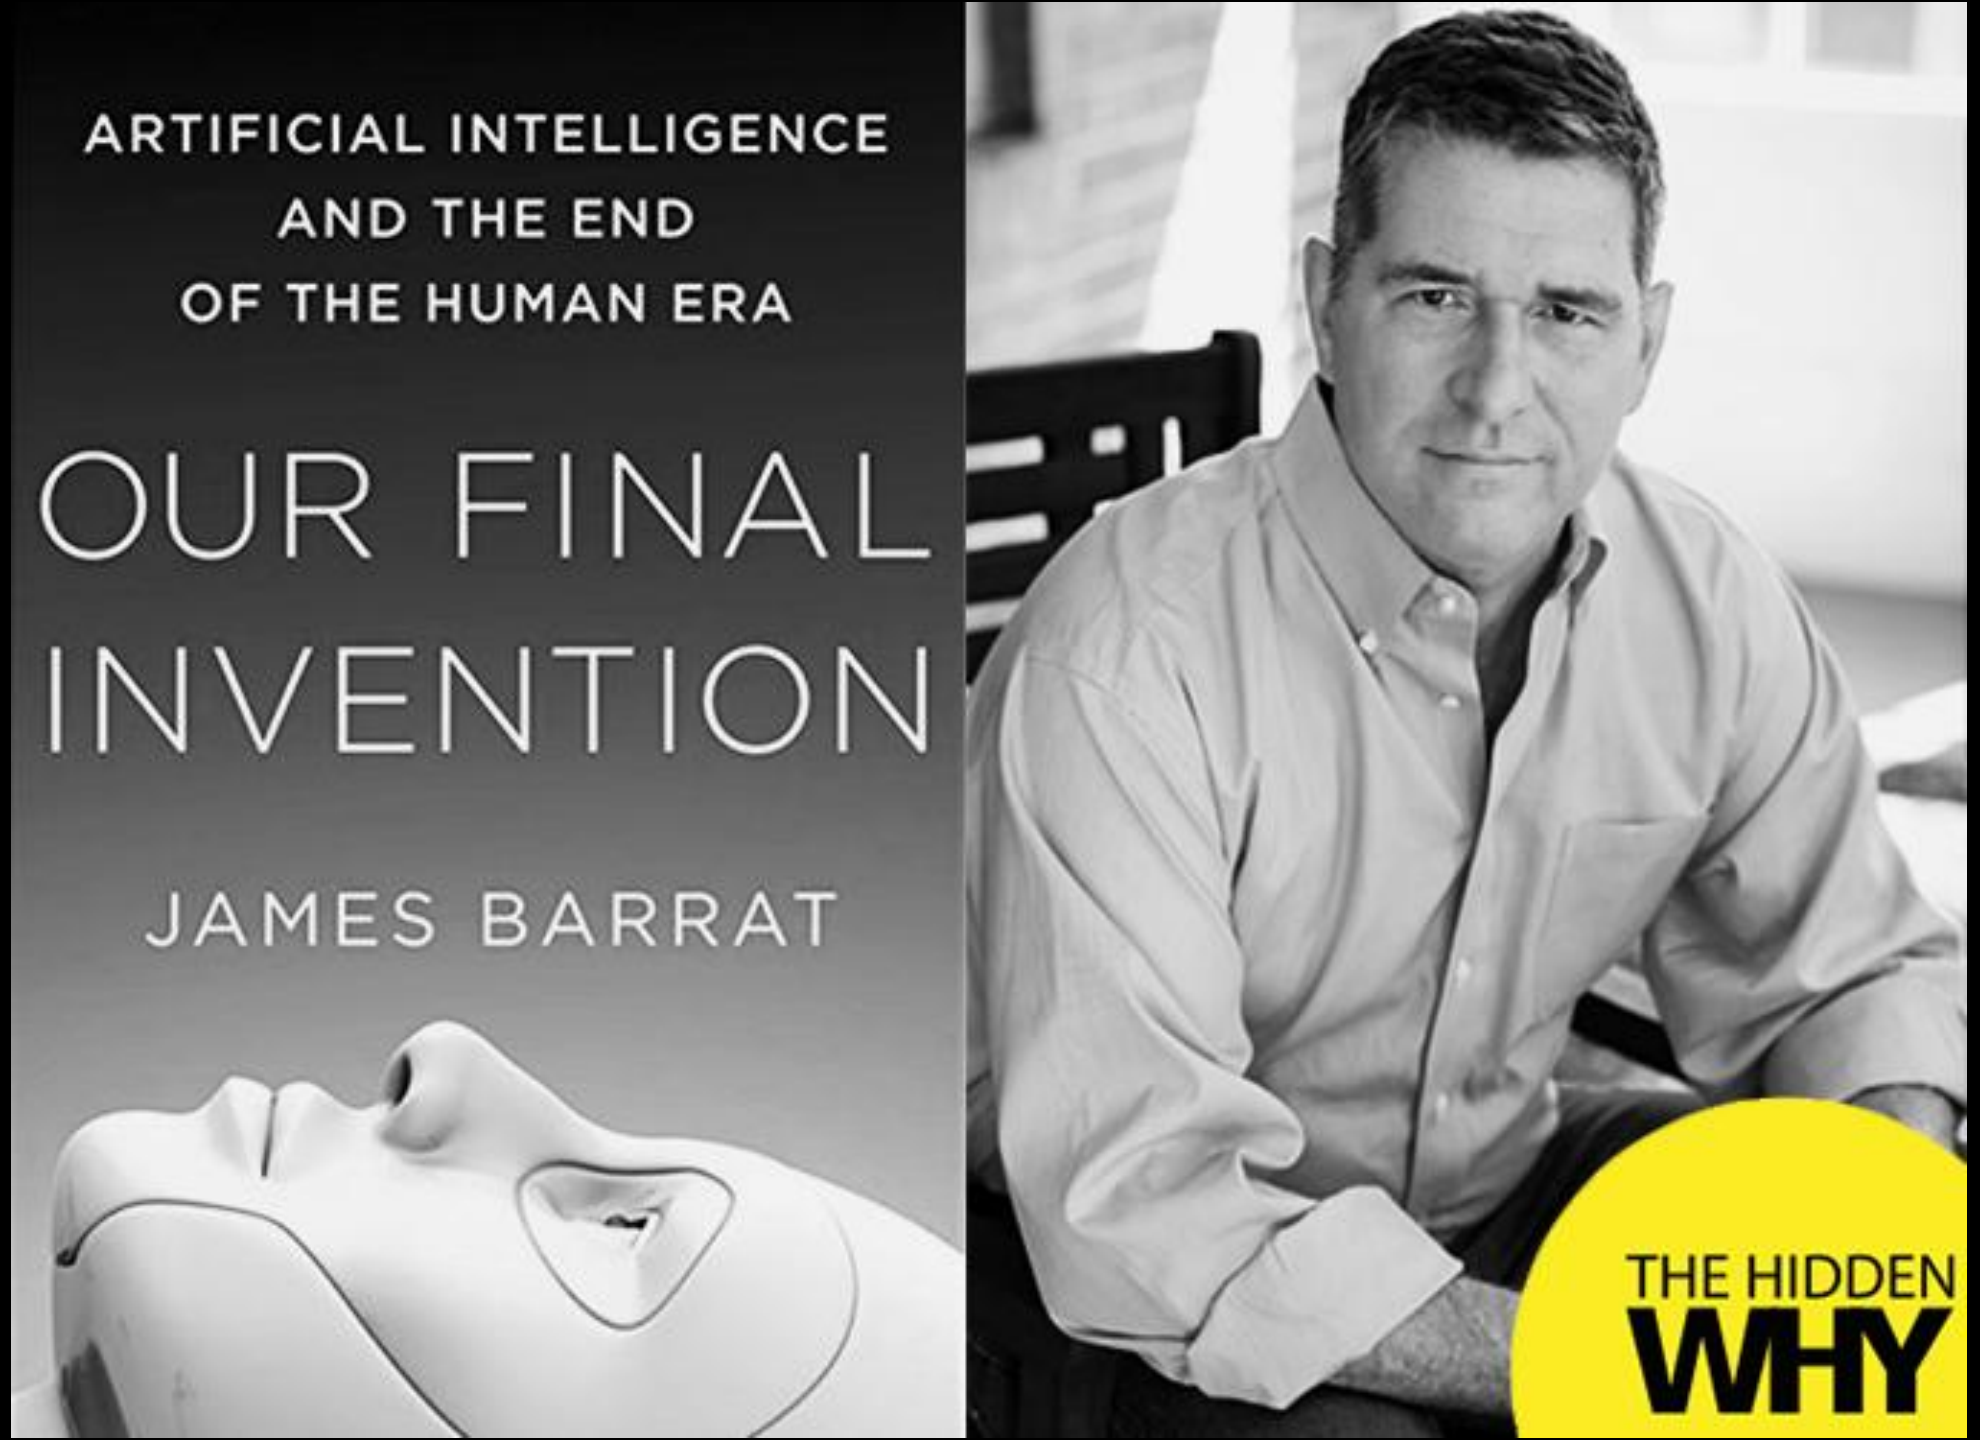
ARTIFICIAL INTELLIGENCE
AND THE END
OF THE HUMAN ERA
OUR FINAL
INVENTION
JAMES BARRAT
THE HIDDEN
WHY

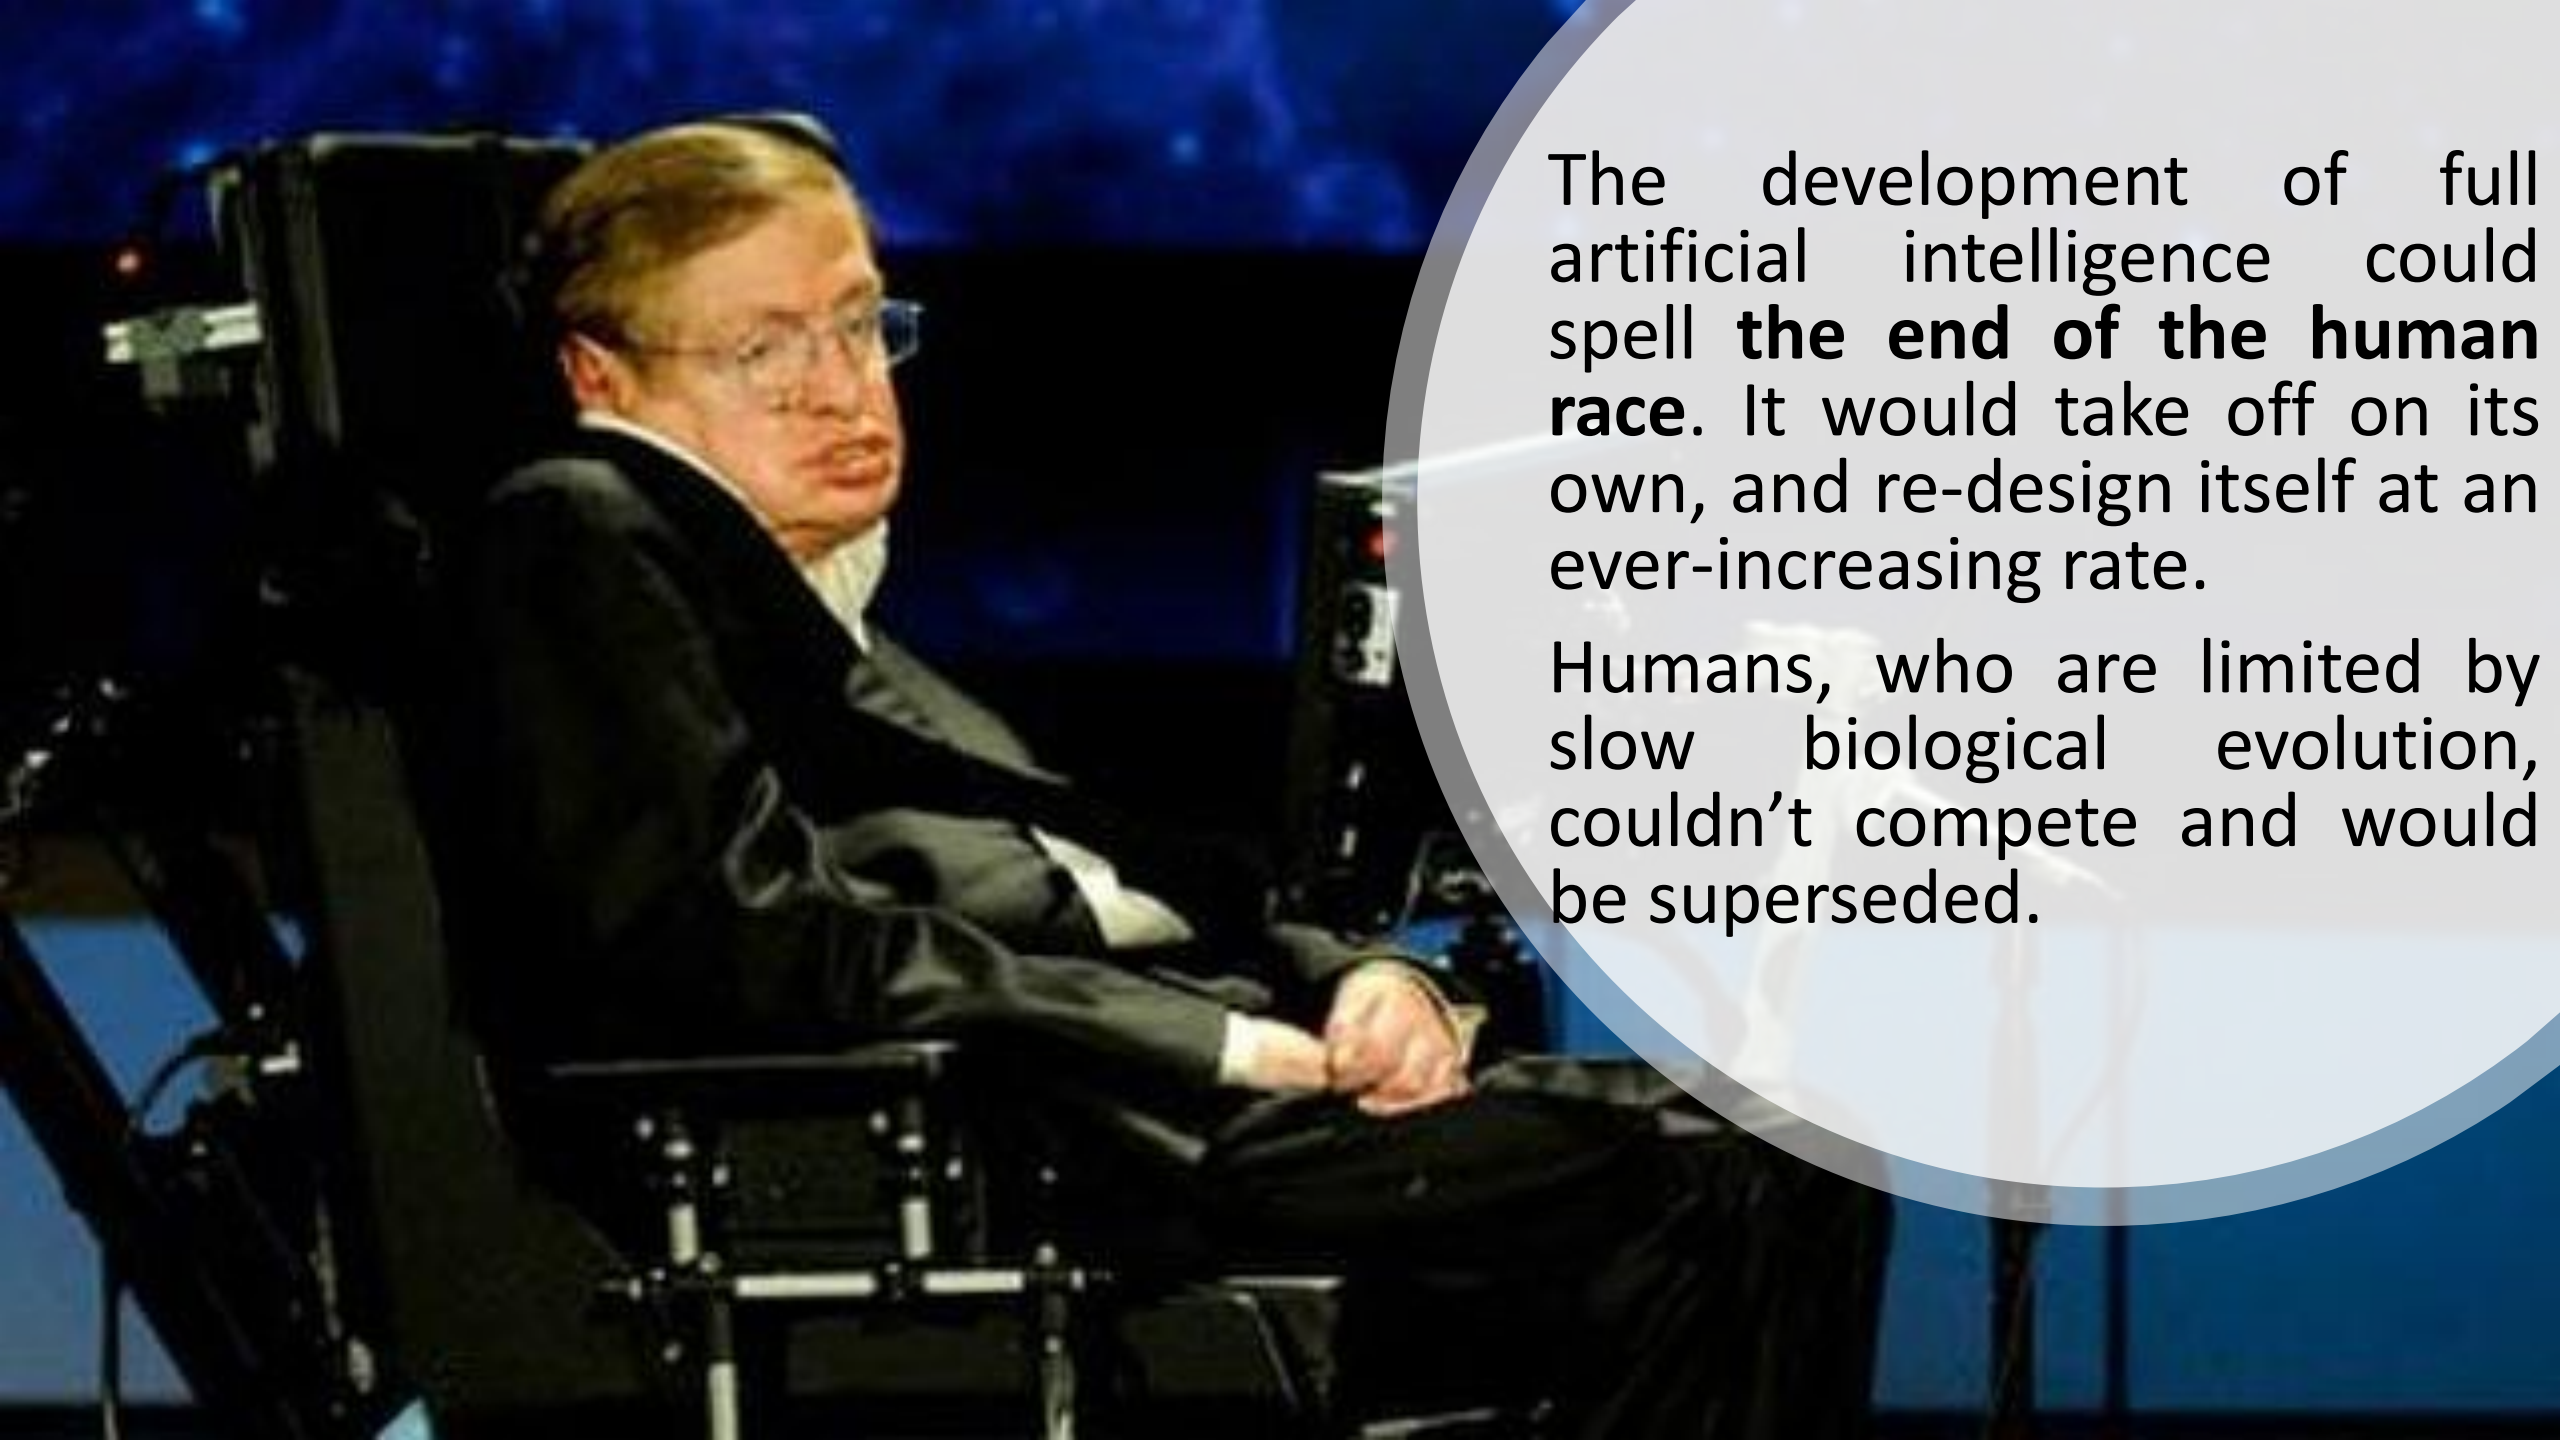

The development of full artificial intelligence could spell **the end of the human race**. It would take off on its own, and re-design itself at an ever-increasing rate.

Humans, who are limited by slow biological evolution, couldn’t compete and would be superseded.

- Robots will do everything better than us
- AI is a greater risk than North Korea
- AI is a fundamental risk to the existence of human civilization

AI is —
the last invention we'll ever make,
the last challenge we'll ever face!

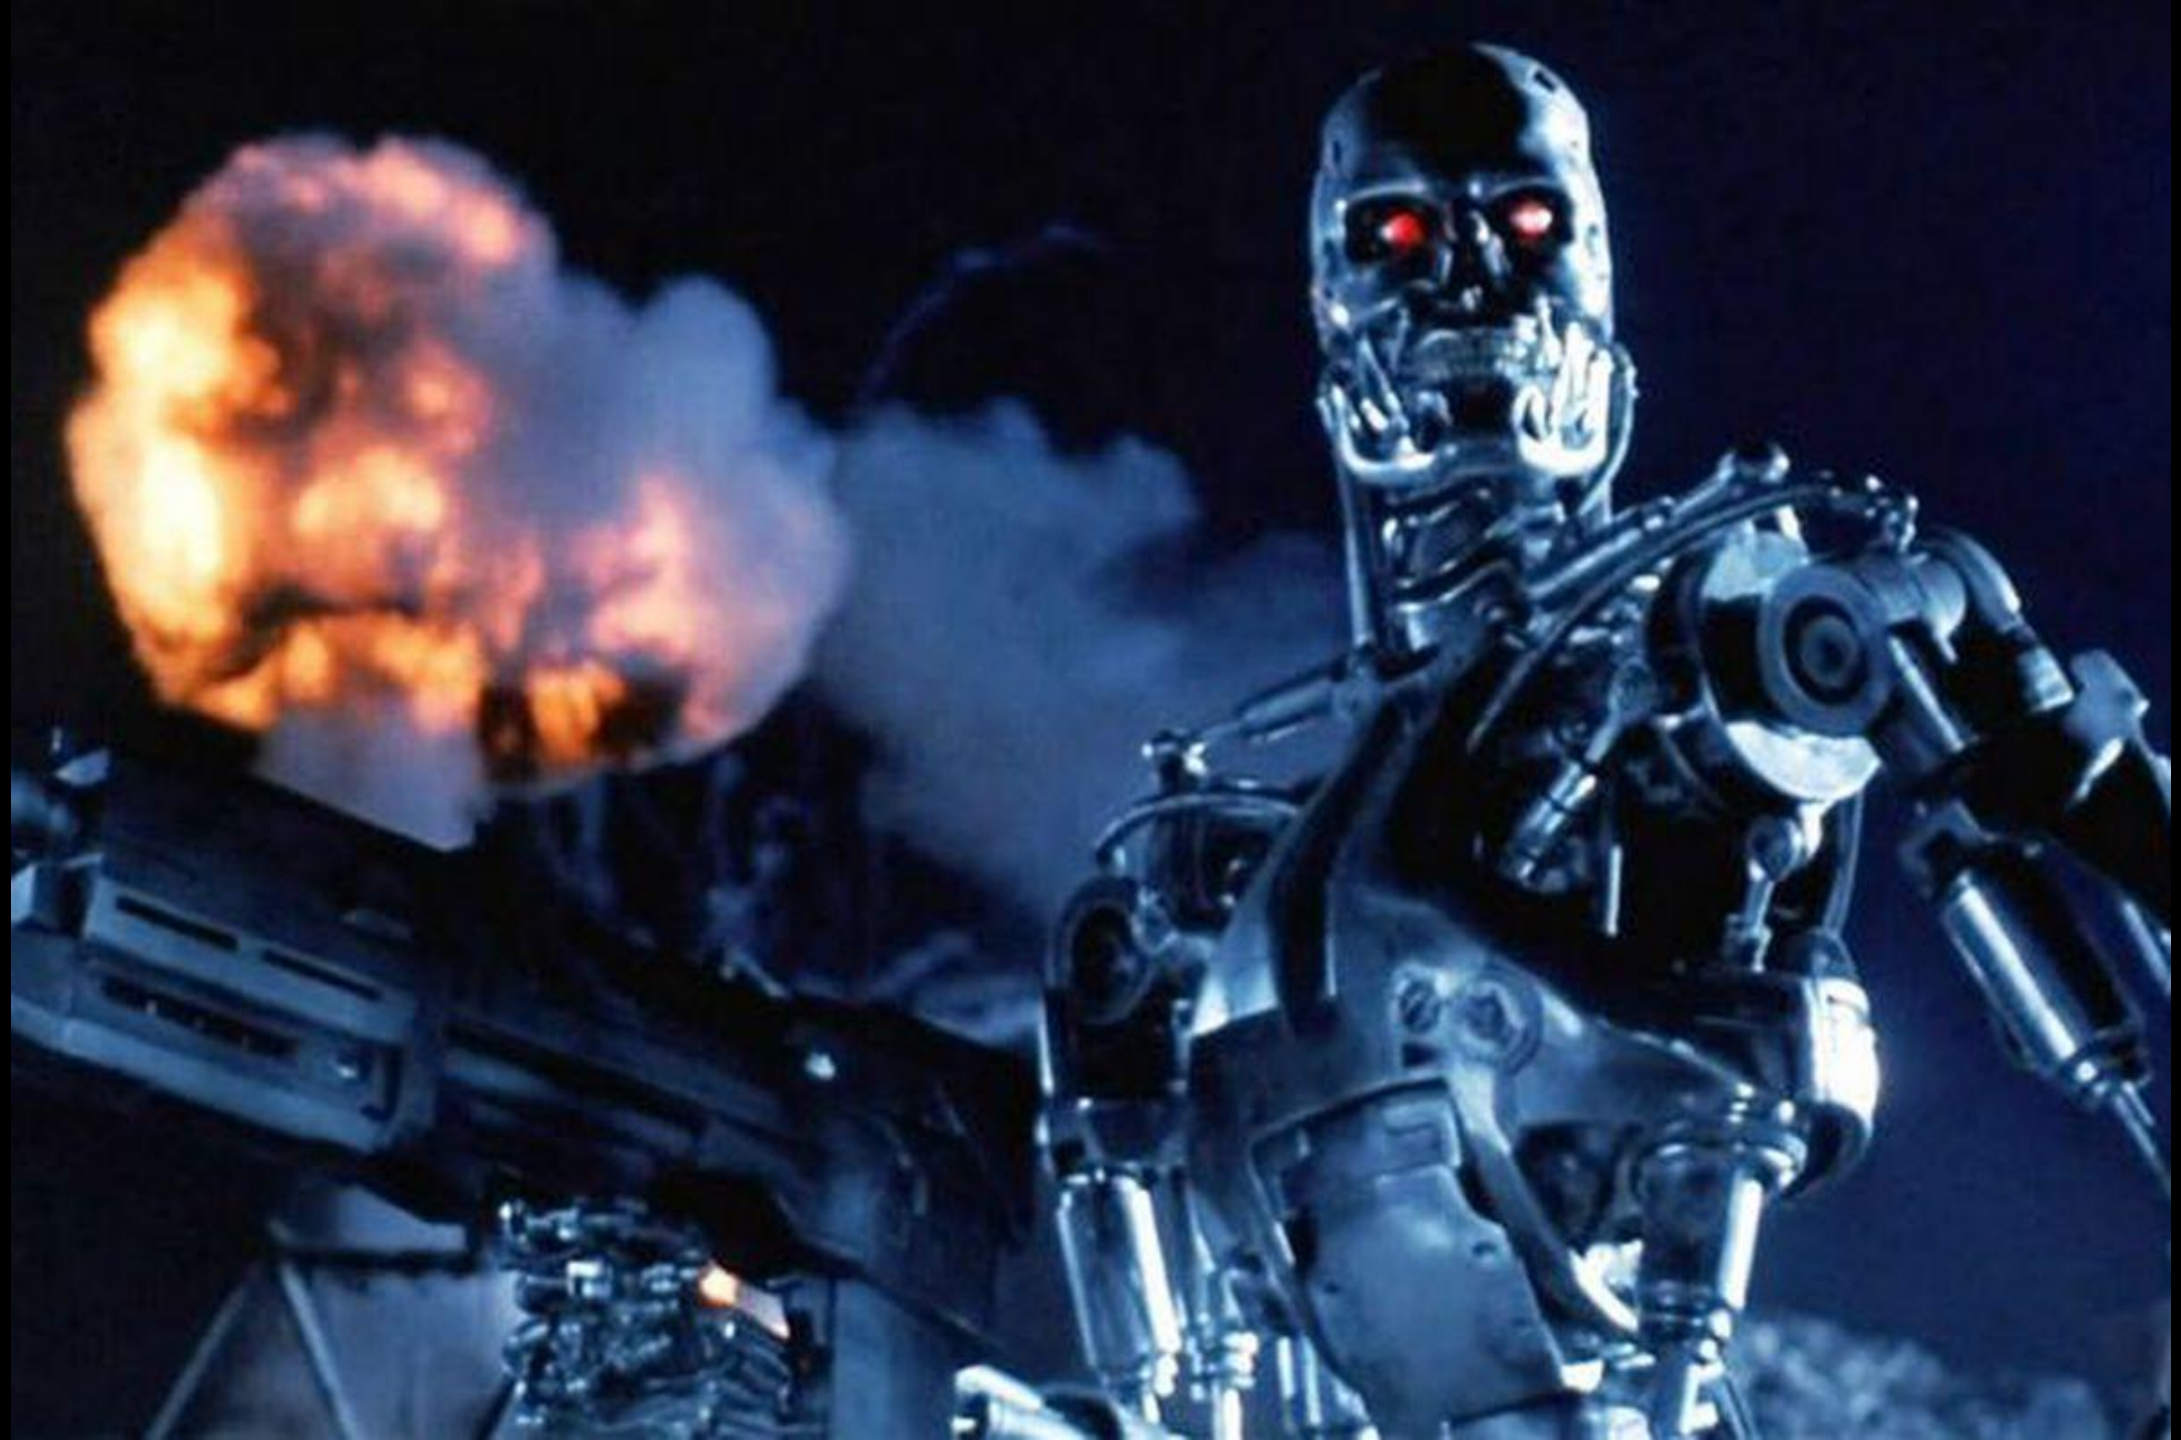

So, what is AI?

AI的起源:

1956 Dartmouth Conference: The Founding Fathers of AI

John MacCarthy

Marvin Minsky

Claude Shannon

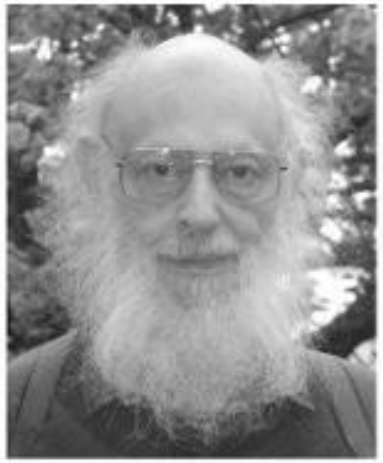

Ray Solomonoff

Alan Newell

Herbert Simon

Arthur Samuel

Oliver Selfridge

Nathaniel Rochester

Trenchard More

Courtesy of scienceabc.com

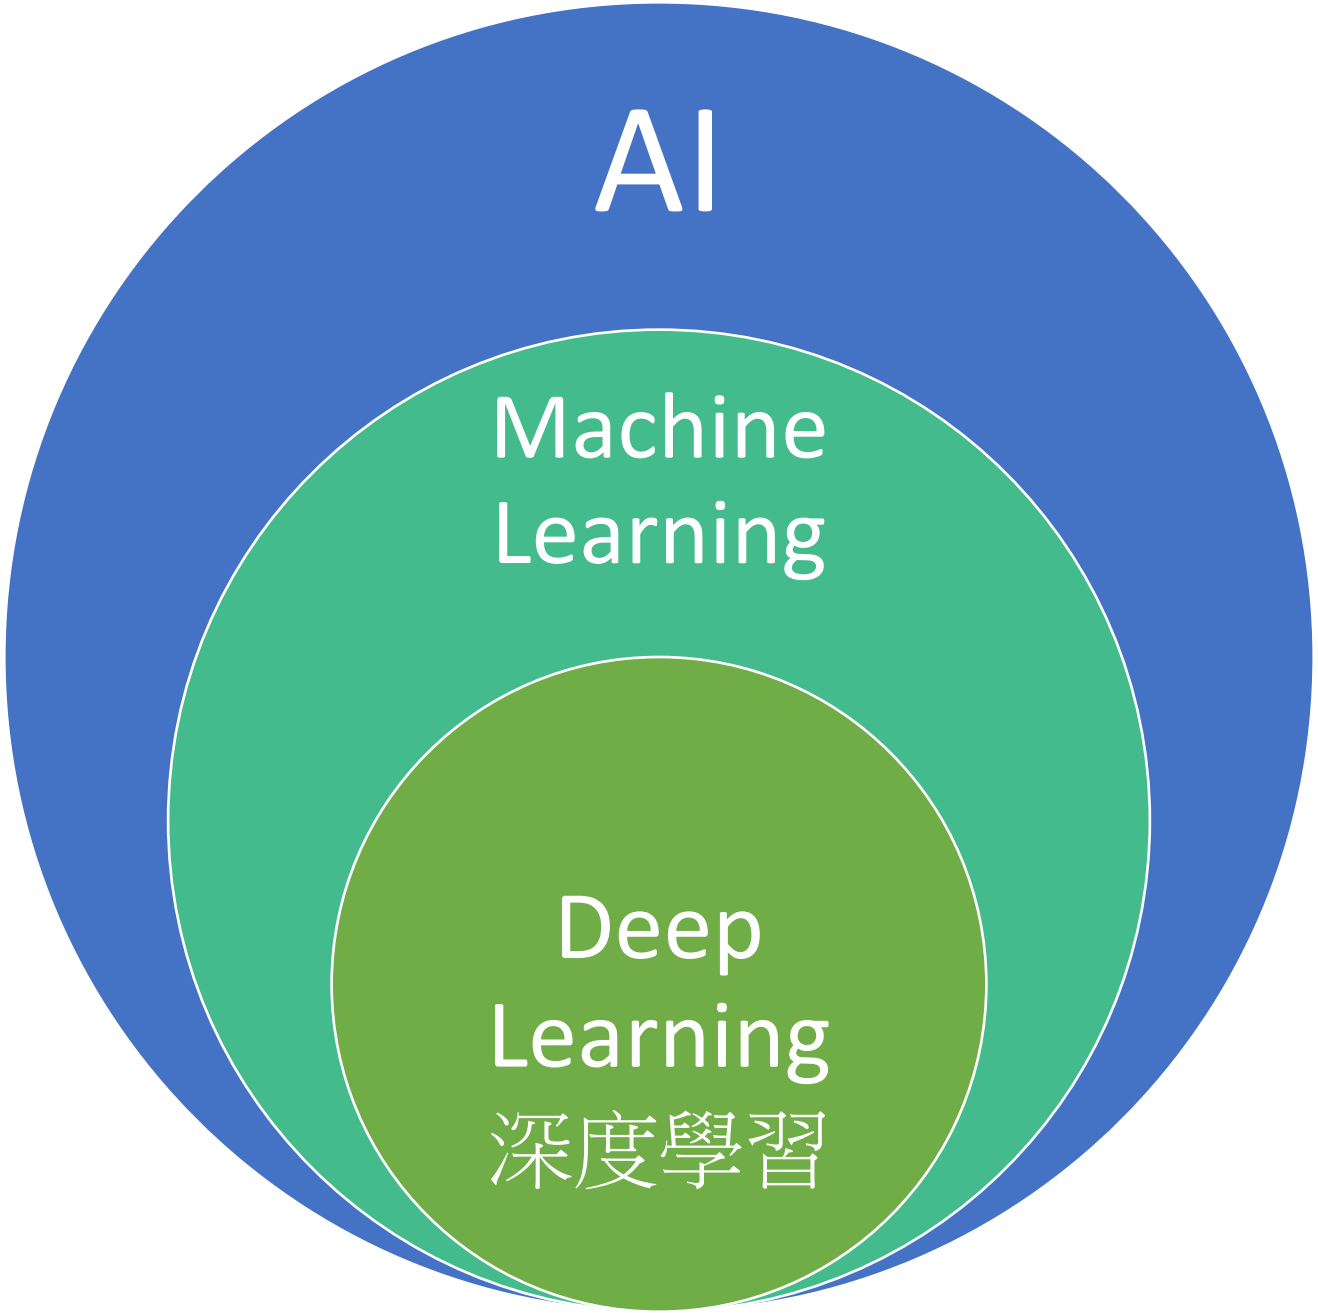
AI
Machine Learning
Deep Learning
深度學習

Machine Learning (Statistical Learning)
機器學習 vs. 程式設計

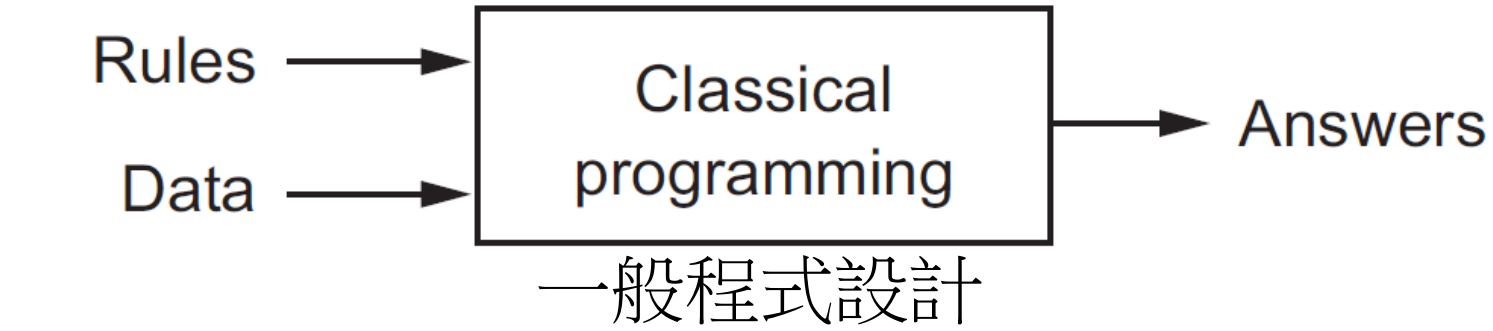

一般程式設計

機器學習

Francois Chollet, “Deep Learning with Python,” Manning, 2017

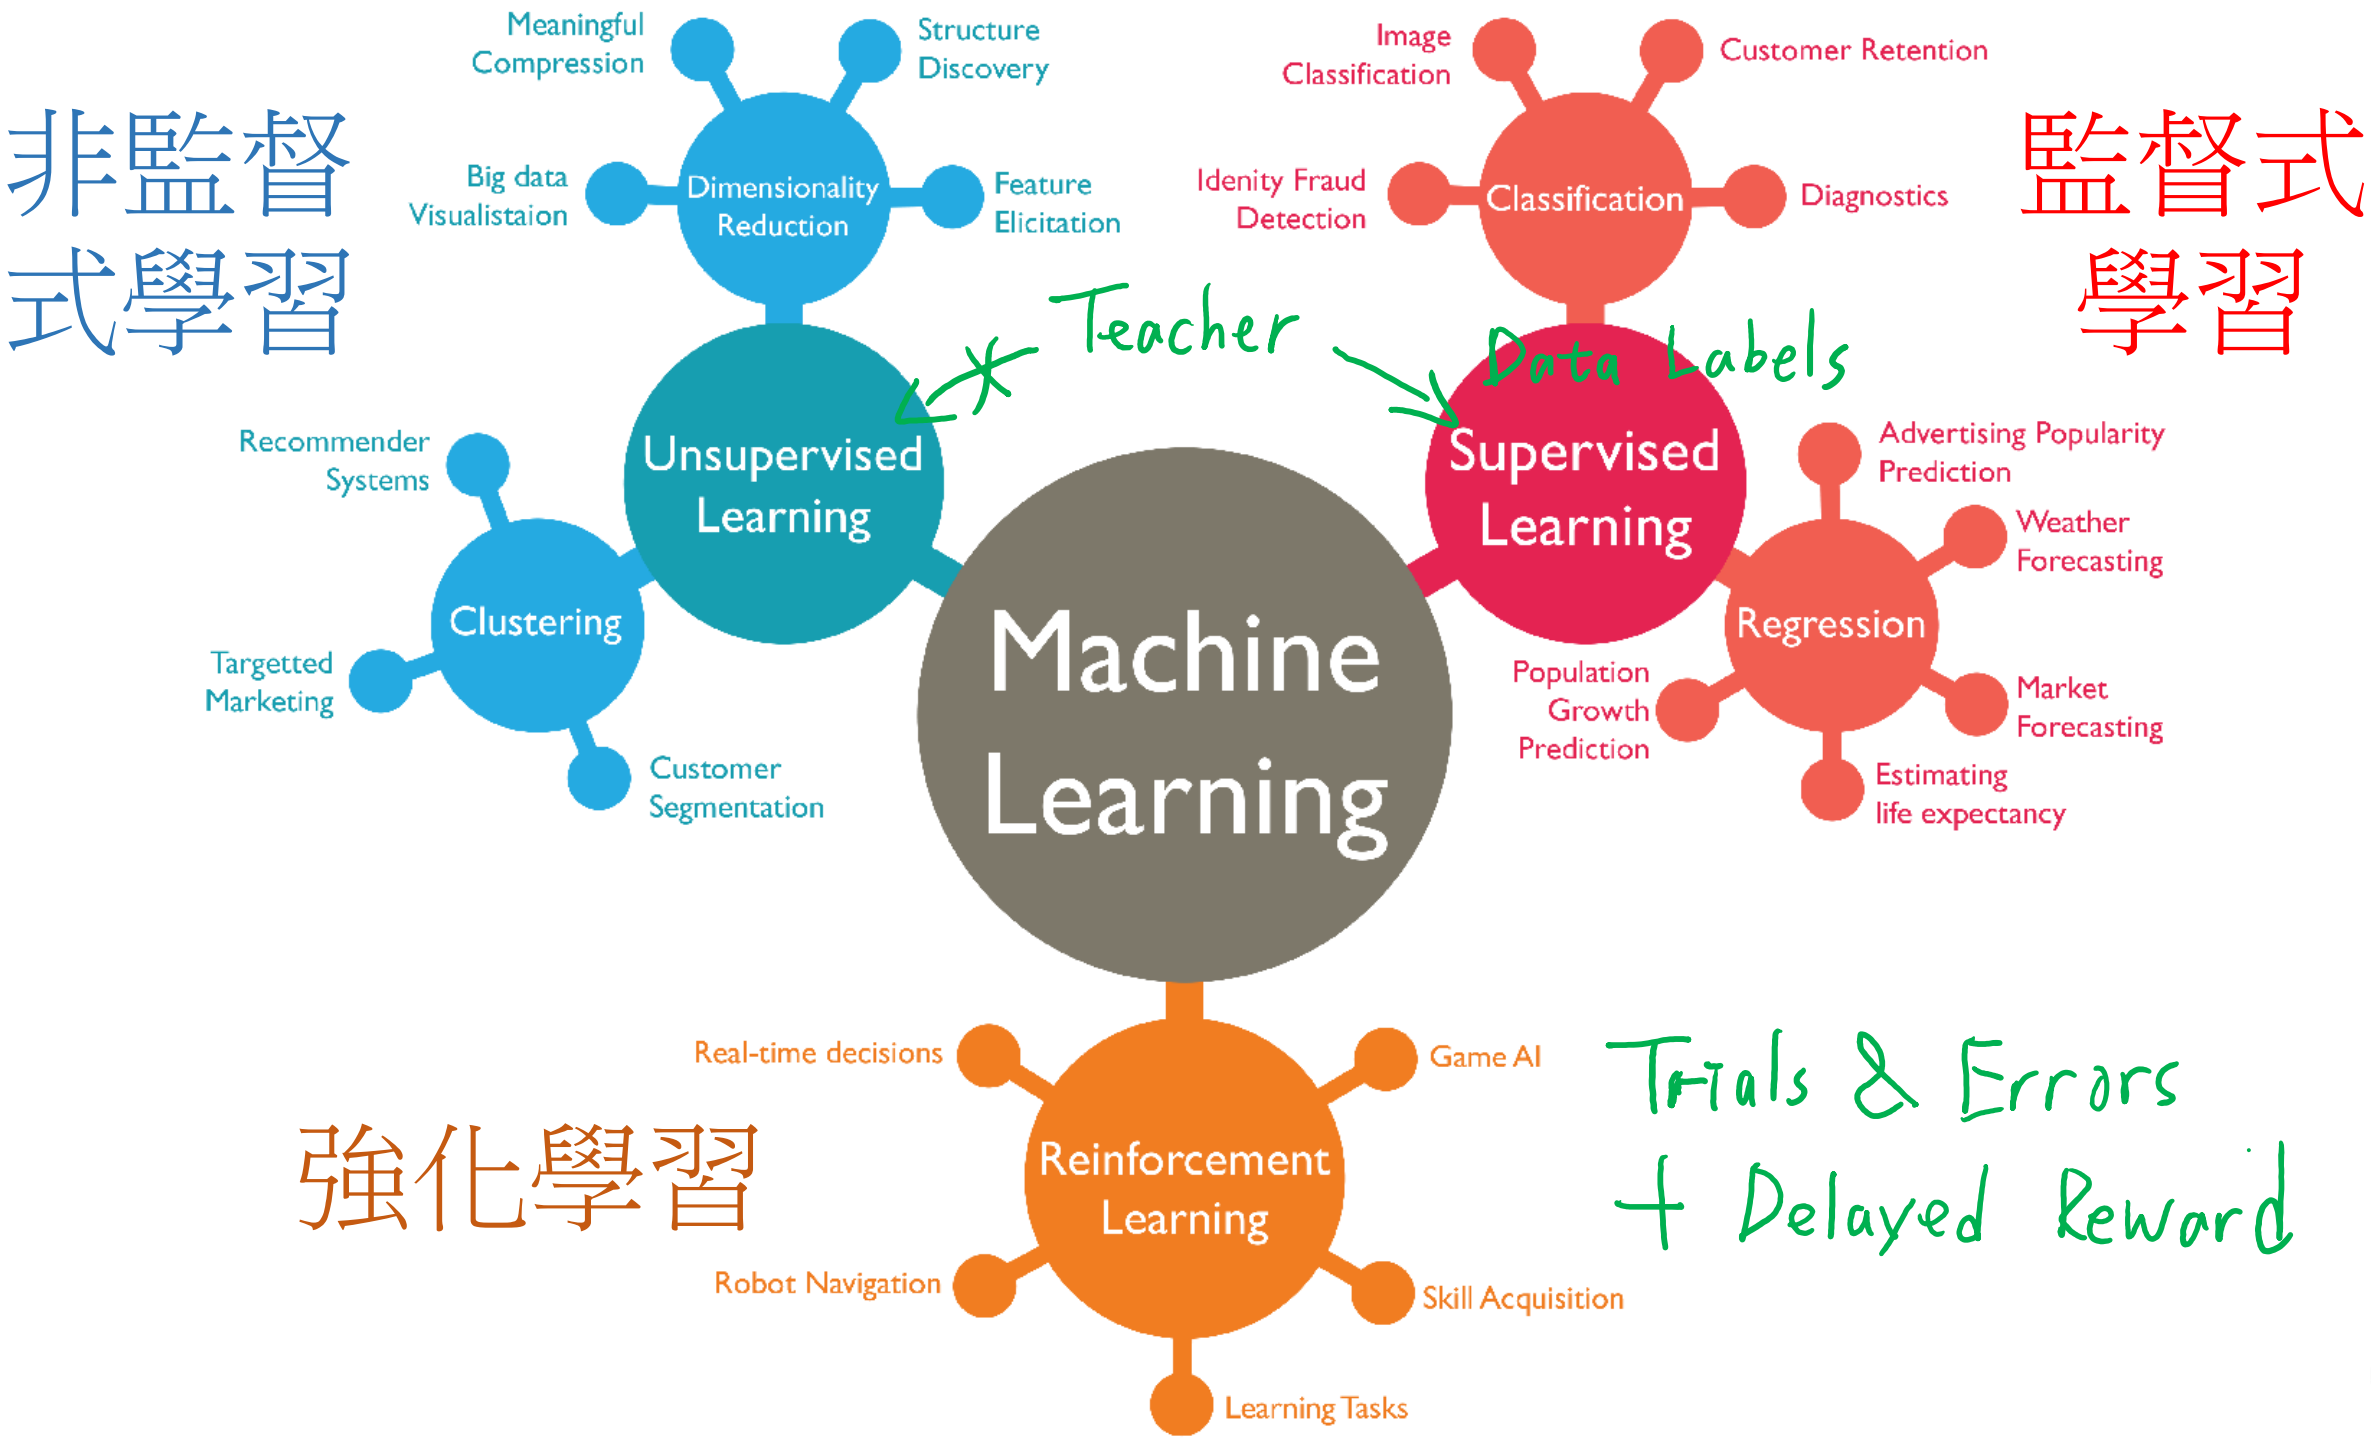
非監督式學習
Meaningful Compression
Structure Discovery
Big data Visualistaion
Dimensionality Reduction
Feature Elicitation
Unsupervised Learning
Recommender Systems
Clustering
Targetted Marketing
Customer Segmentation
Teacher
Data Labels
監督式學習
Image Classification
Customer Retention
Idenity Fraud Detection
Classification
Diagnostics
Supervised Learning
Advertising Popularity Prediction
Weather Forecasting
Regression
Population Growth Prediction
Market Forecasting
Estimating life expectancy
Machine Learning
強化學習
Real-time decisions
Game AI
Reinforcement Learning
Robot Navigation
Skill Acquisition
Learning Tasks
Trials & Errors + Delayed Reward


Data
Labelers

Self-Supervised Learning

- Predict any part of the input from any other part.
- Predict the future from the past.
- Predict the future from the recent past.
- Predict the past from the present.
- Predict the top from the bottom.
- Predict the occluded from the visible
- Pretend there is a part of the input you don't know and predict that.

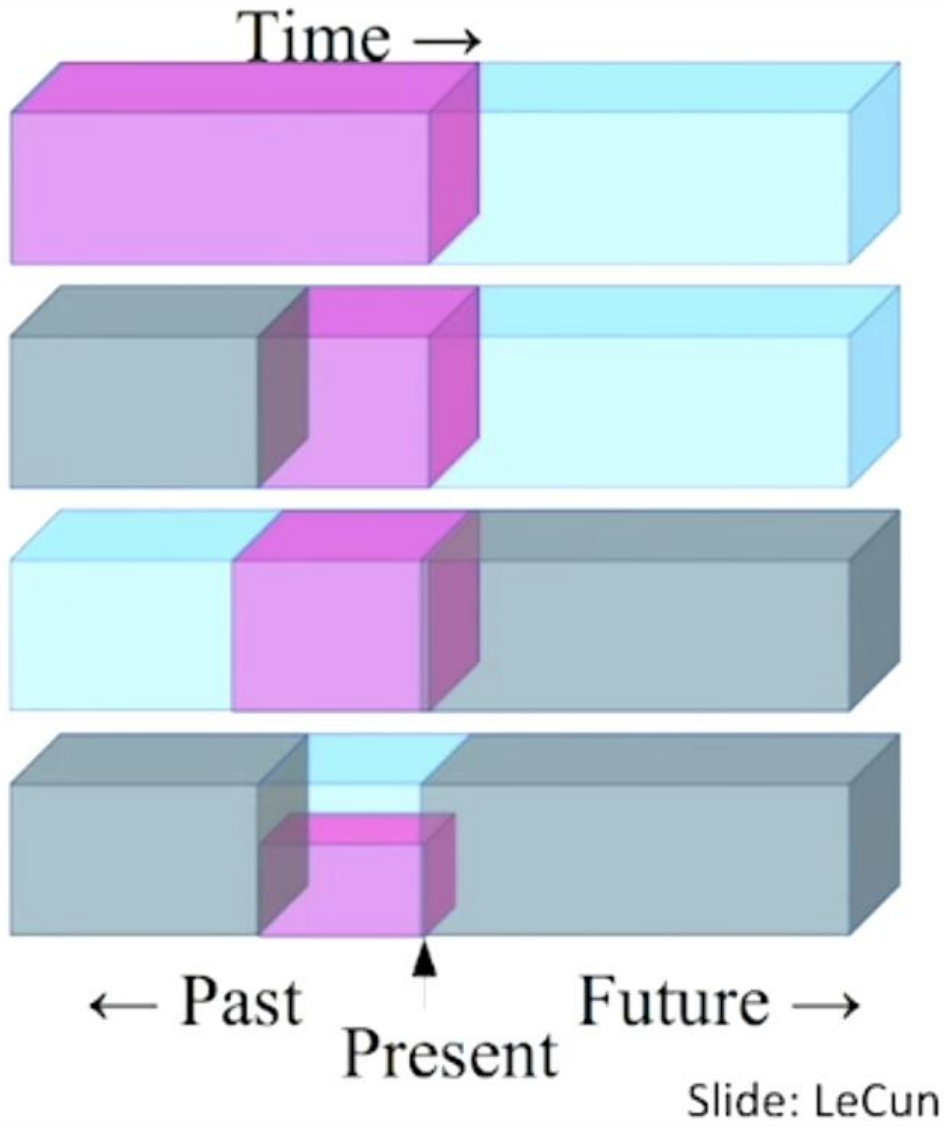

- https://www.youtube.com/watch?v=7I0Qt7GALVk

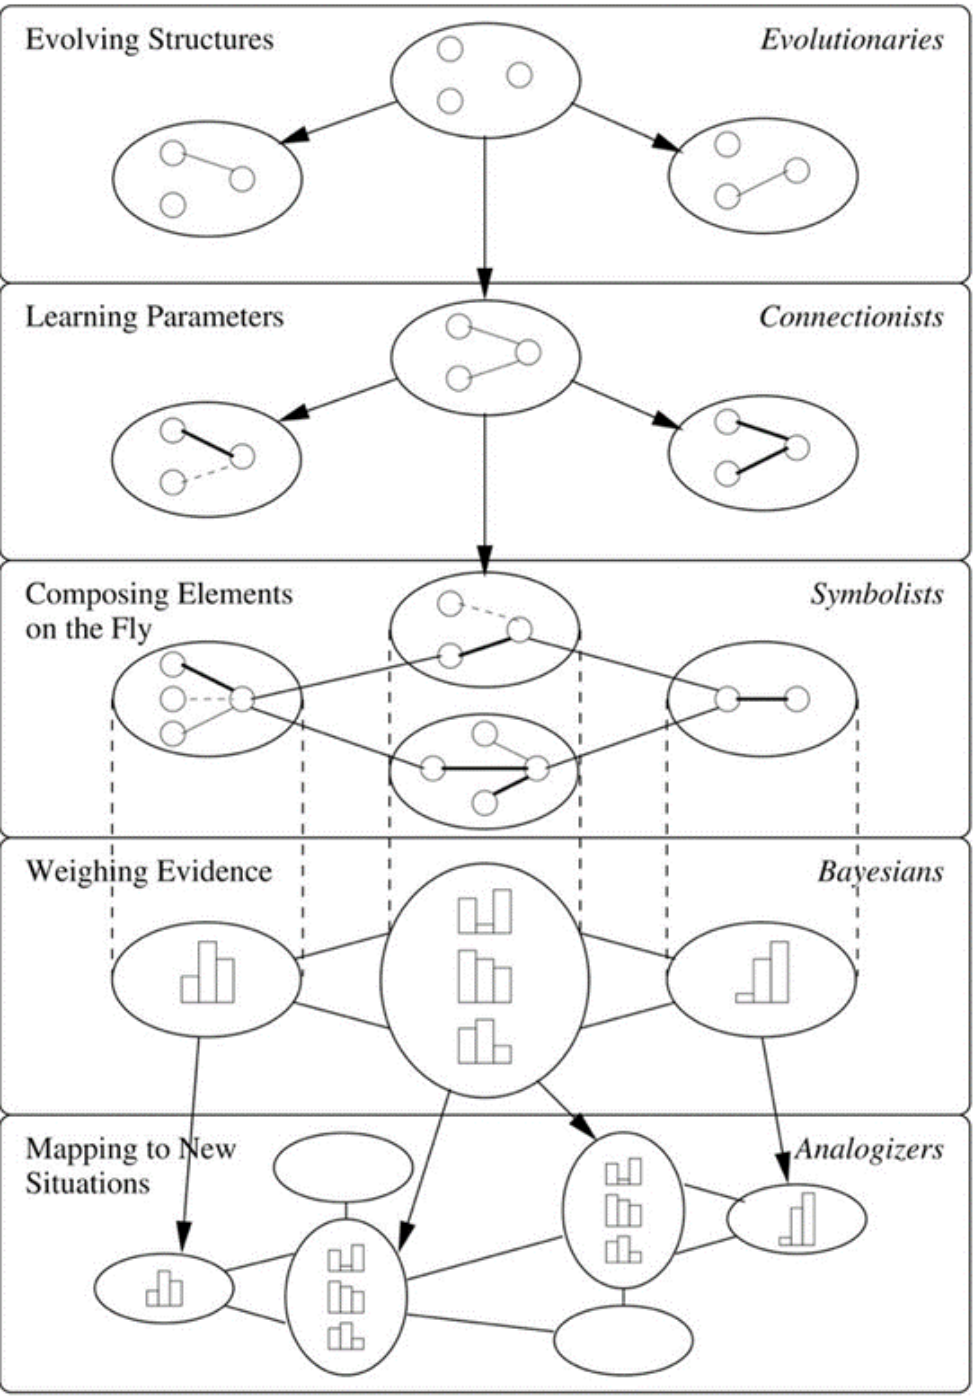

5 Tribes of Machine Learning

- Evolutionaries (演化法)
- Connectionists (類神經網路)
- Symbolists (歸納法)
- Bayesians (貝氏機率)
- Analogizers (類比近似)

The Master Algorithm – Pedro Domingos

5 Tribes of Machine Learning

- Symbolists: Decision Trees, Random Forest
- Bayesians: Naïve Bayesians
- Analogizers: SVM, k-NN
- Evolutionaries: Gene algorithms
- Connectionists: Deep Learning

All Algorithms can be Reduced to 3 Operations!

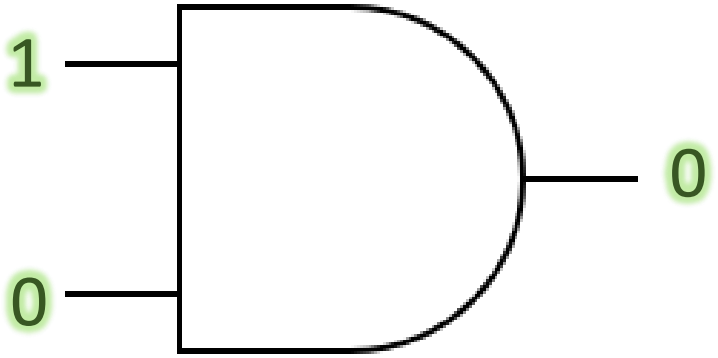

AND

A	B	Output
0	0	0
1	0	0
0	1	0
1	1	1

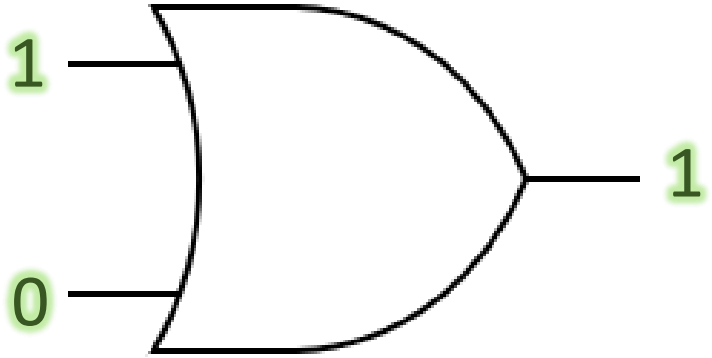

OR

A	B	Output
0	0	0
1	0	1
0	1	1
1	1	1

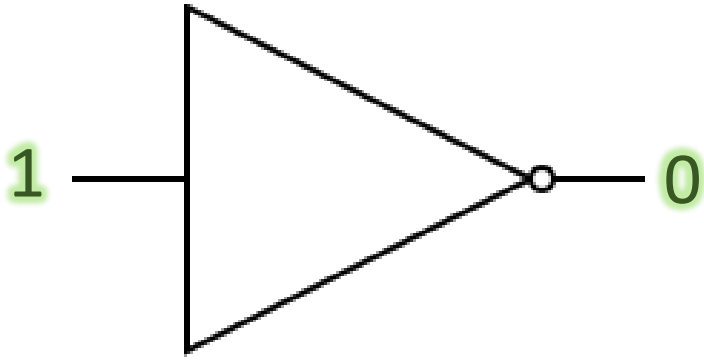

NOT

Input	Output
0	1
1	0

XOR

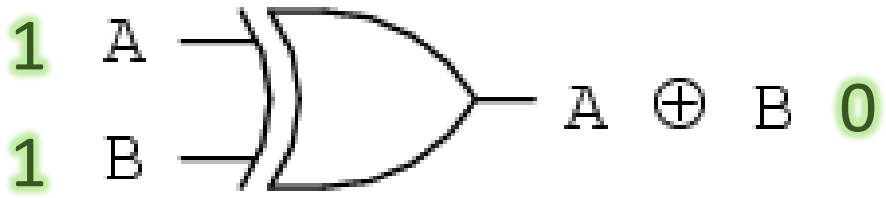

. . . is equivalent to . . .

A	B	Output
0	0	0
1	0	1
0	1	1
1	1	0

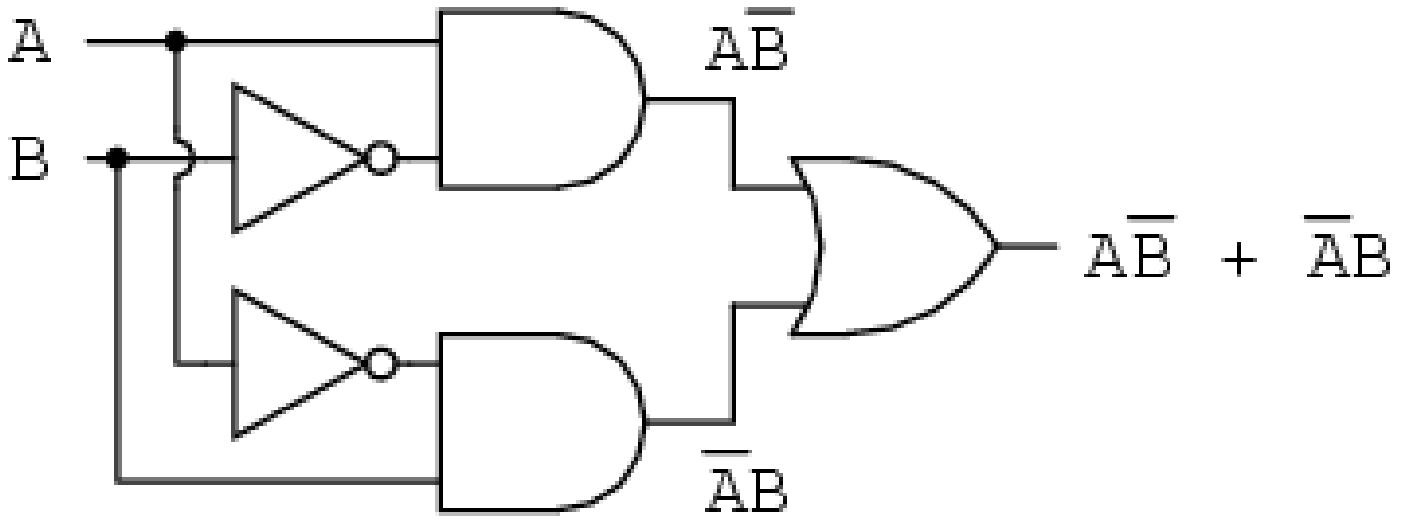

$$A \oplus B = A\overline{B} + \overline{A}B$$

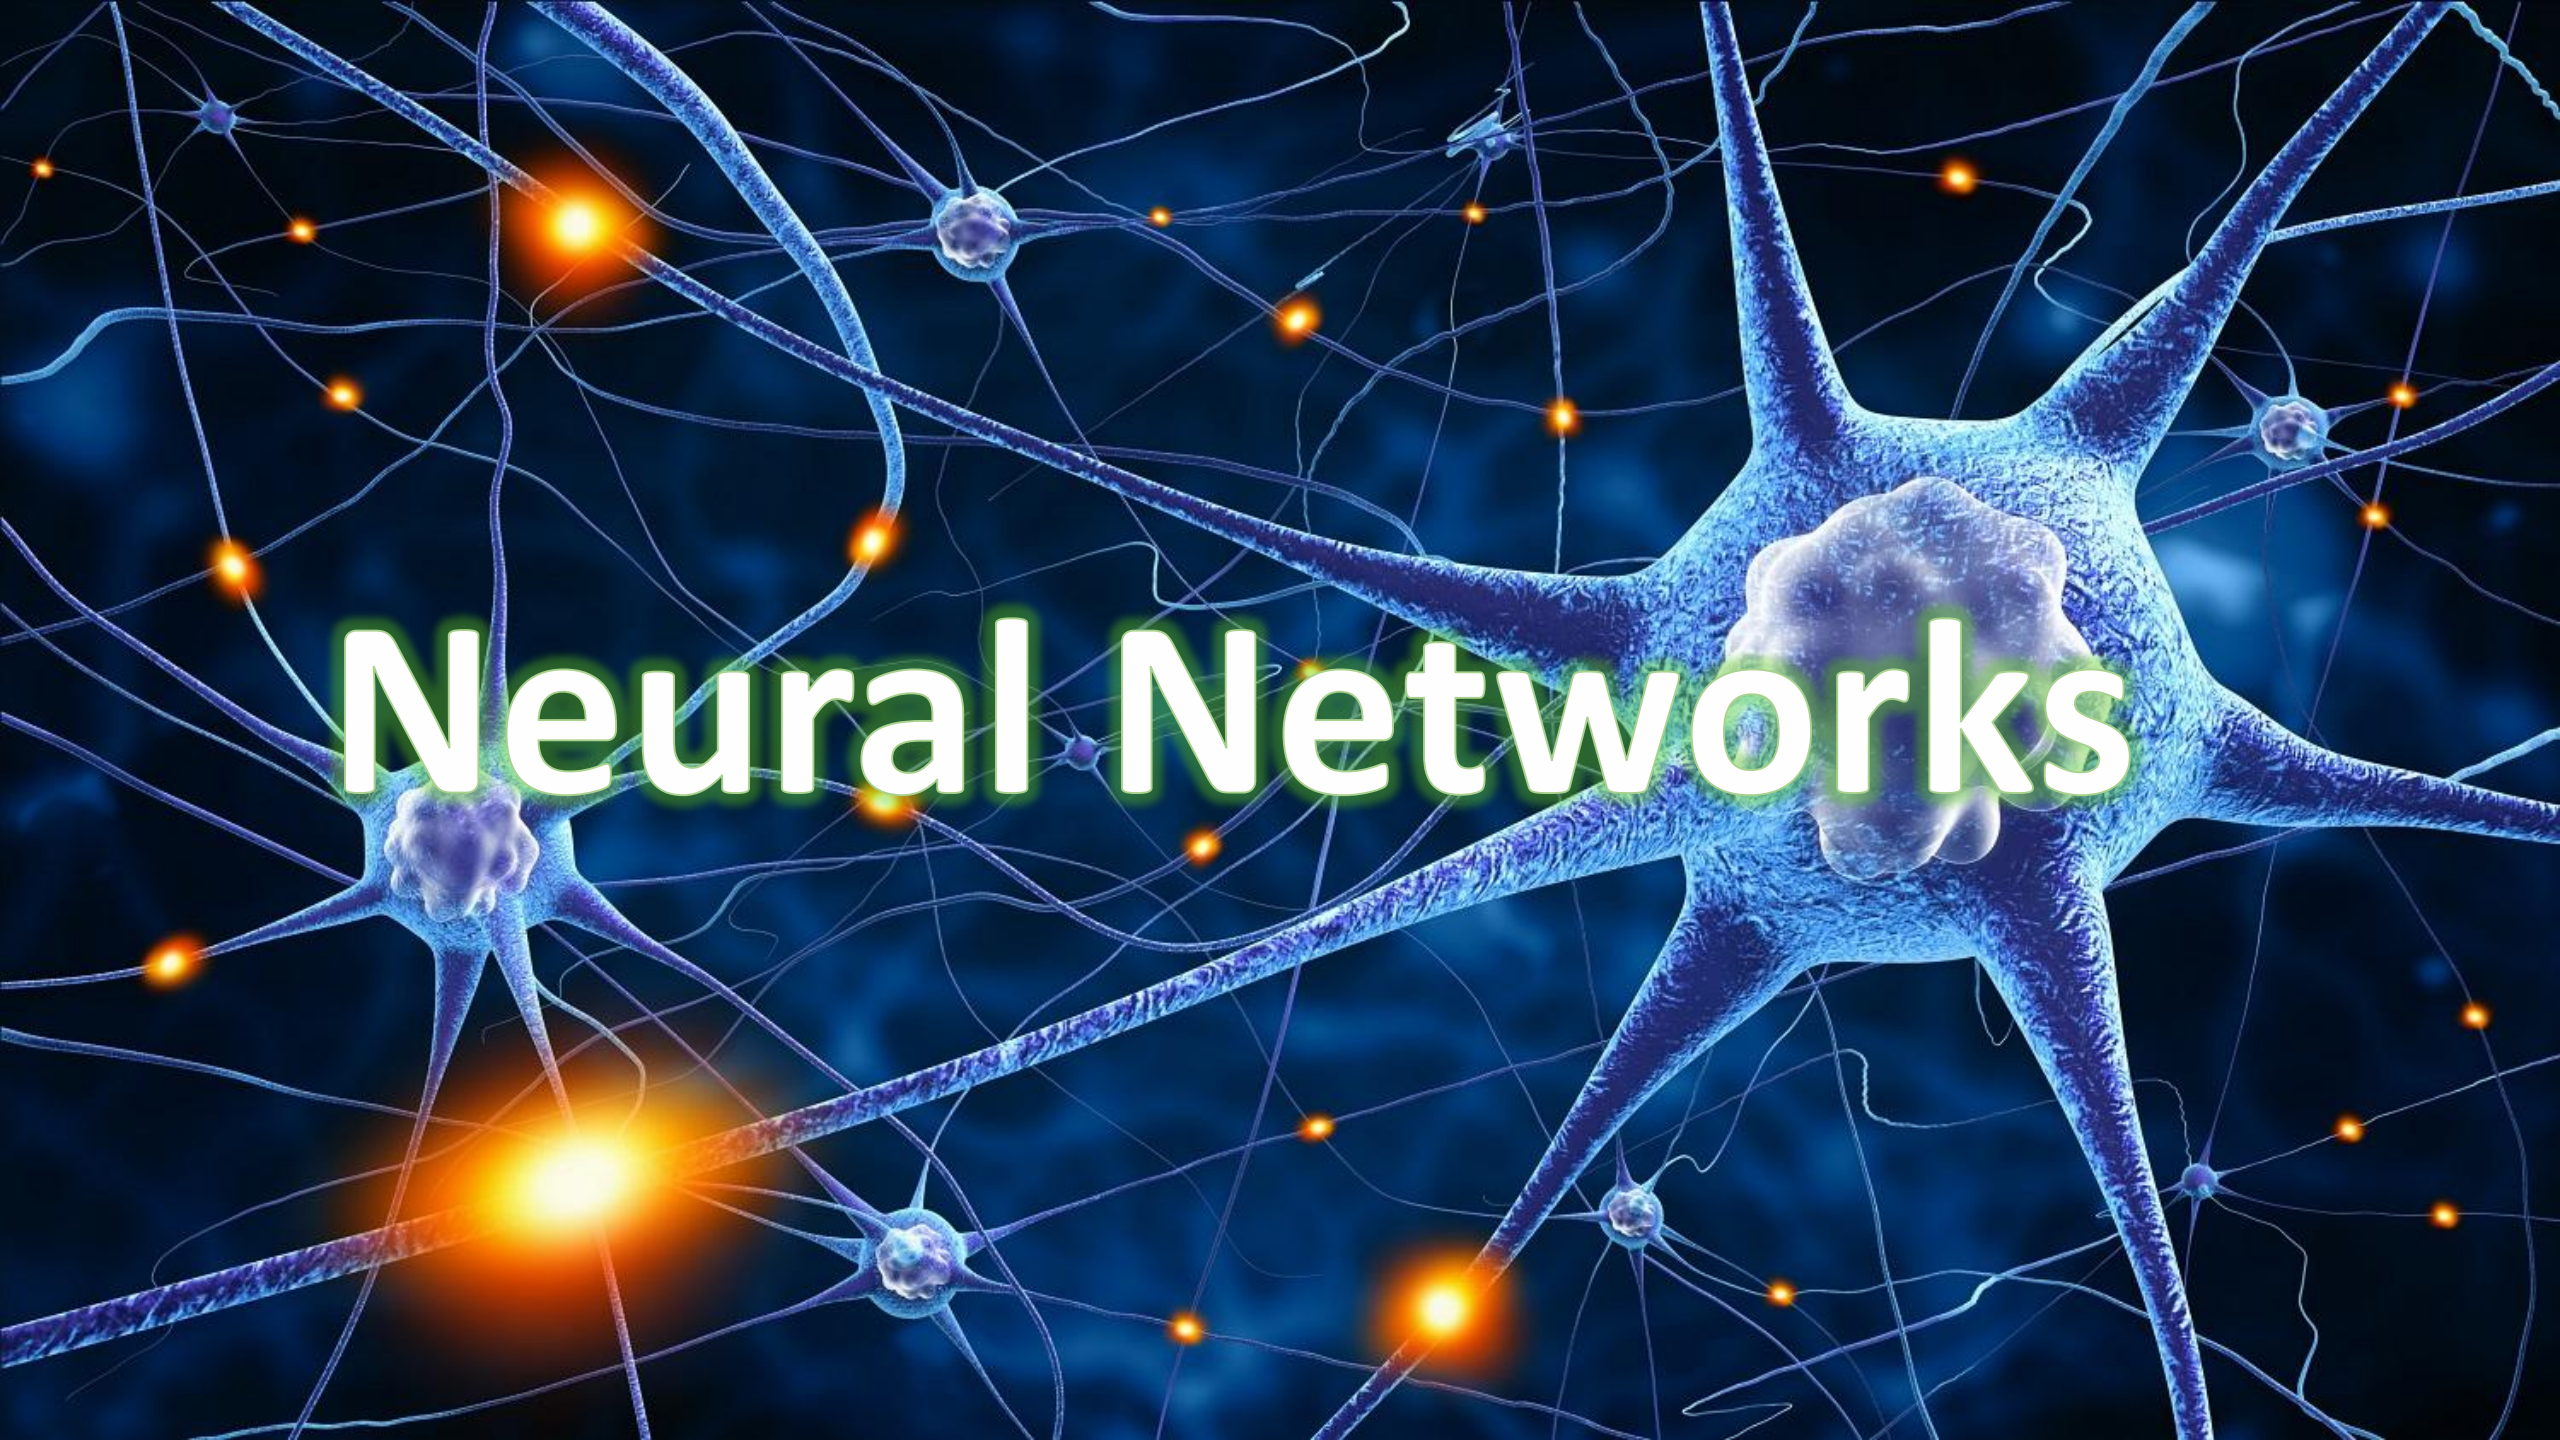
Neural Networks

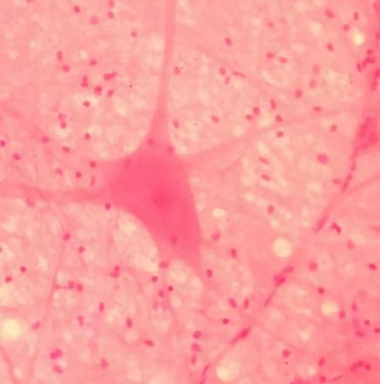

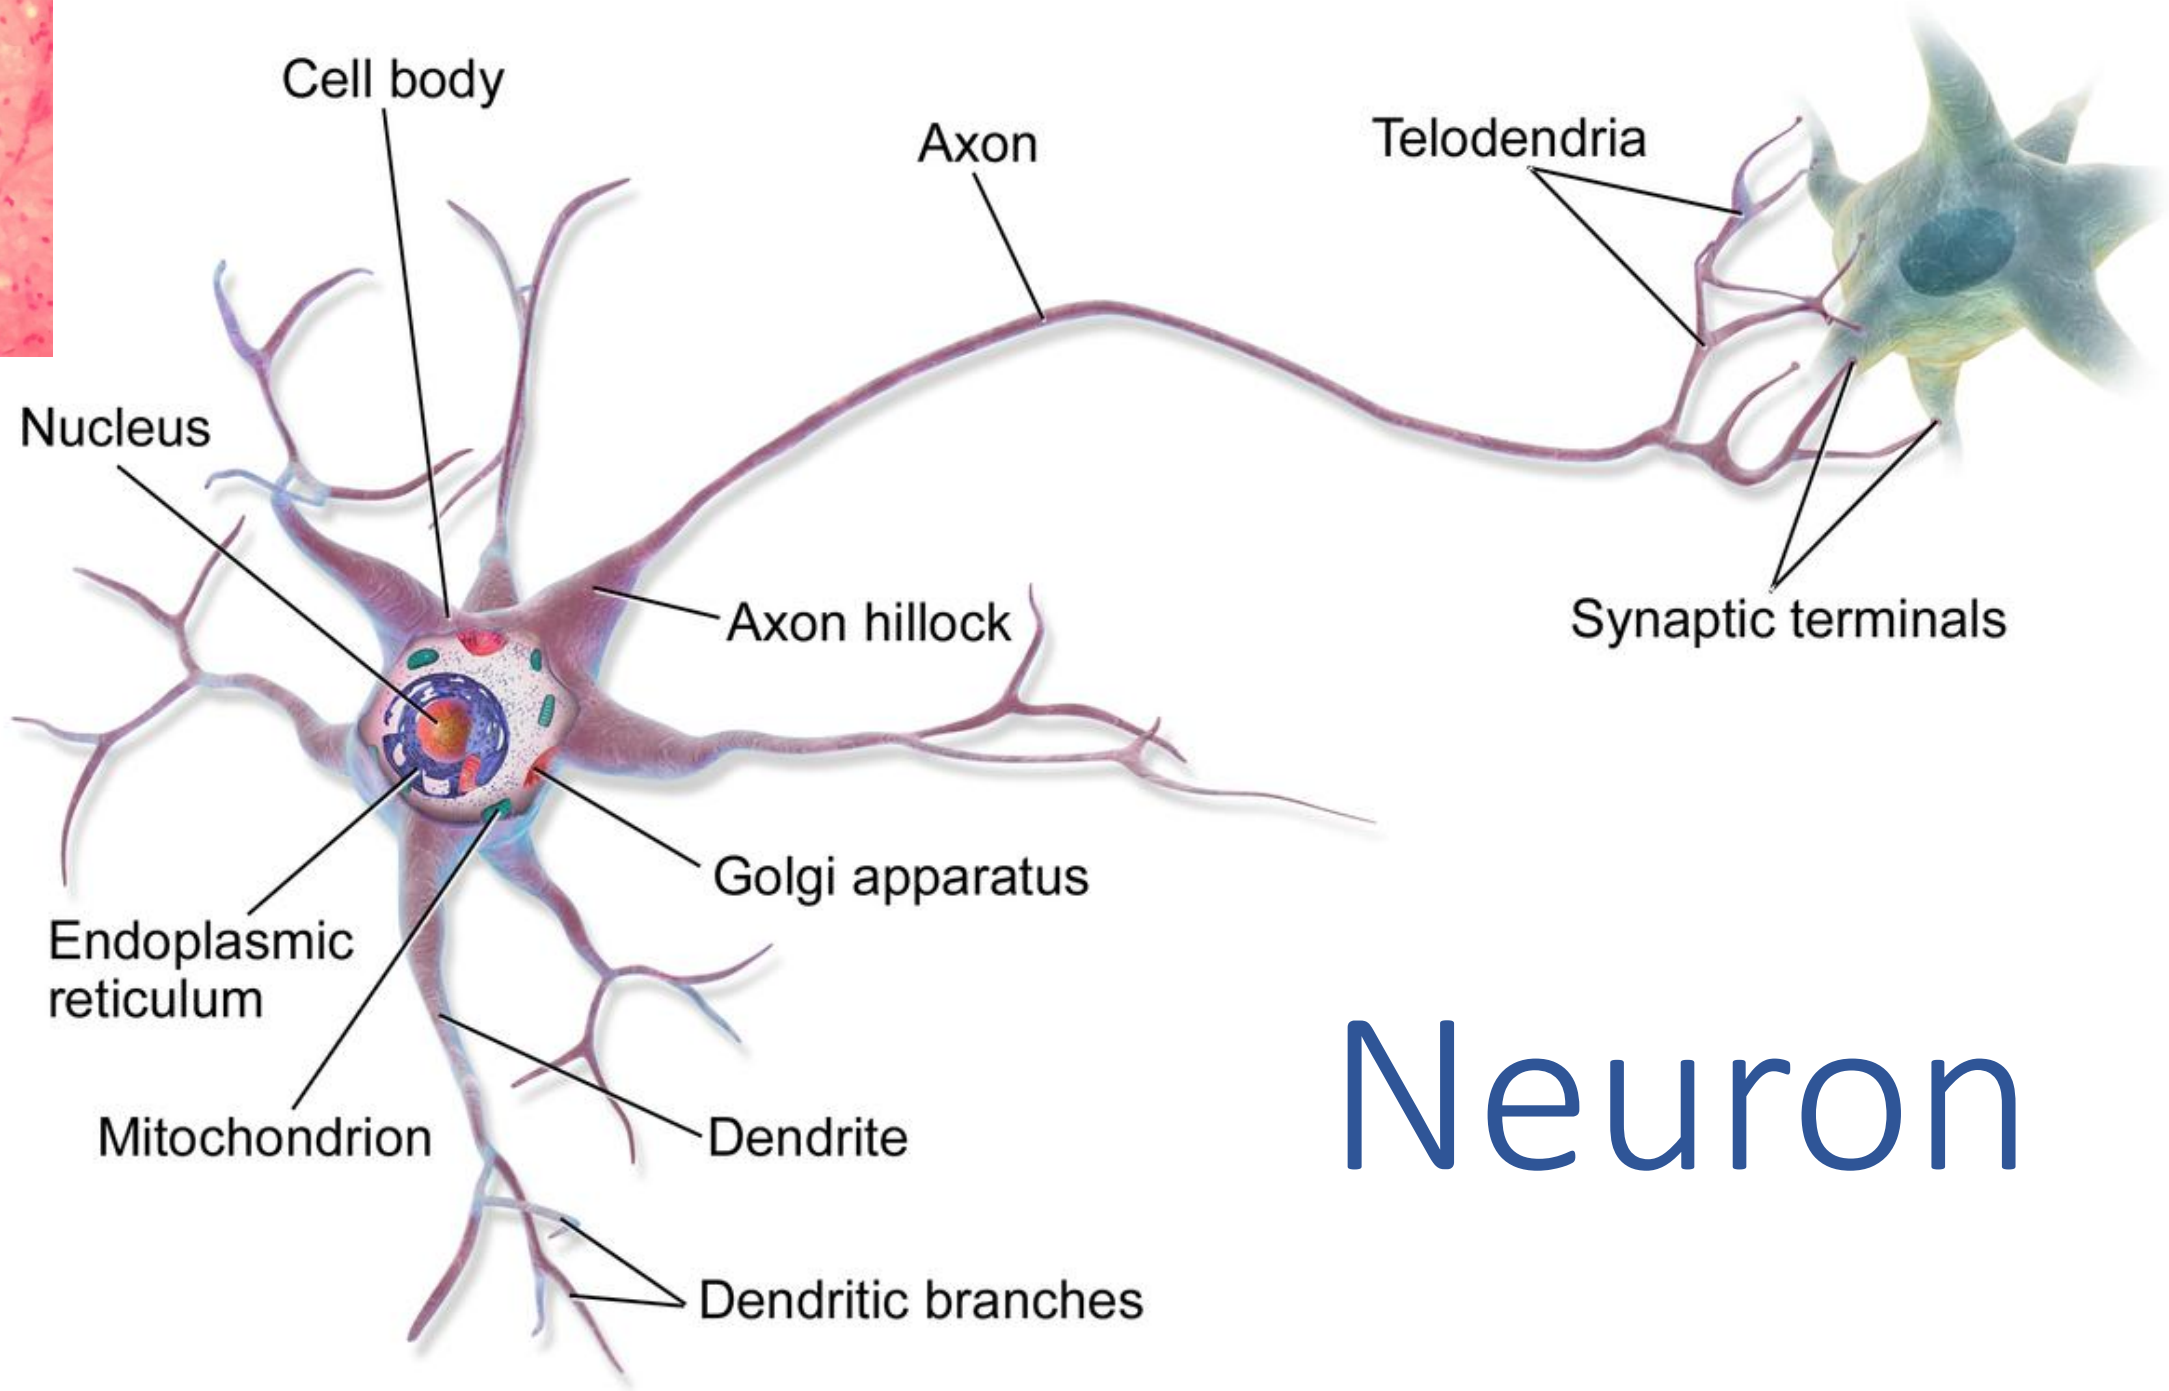
Cell body
Axon
Telodendria
Nucleus
Axon hillock
Synaptic terminals
Golgi apparatus
Endoplasmic reticulum
Mitochondrion
Dendrite
Dendritic branches

Neuron

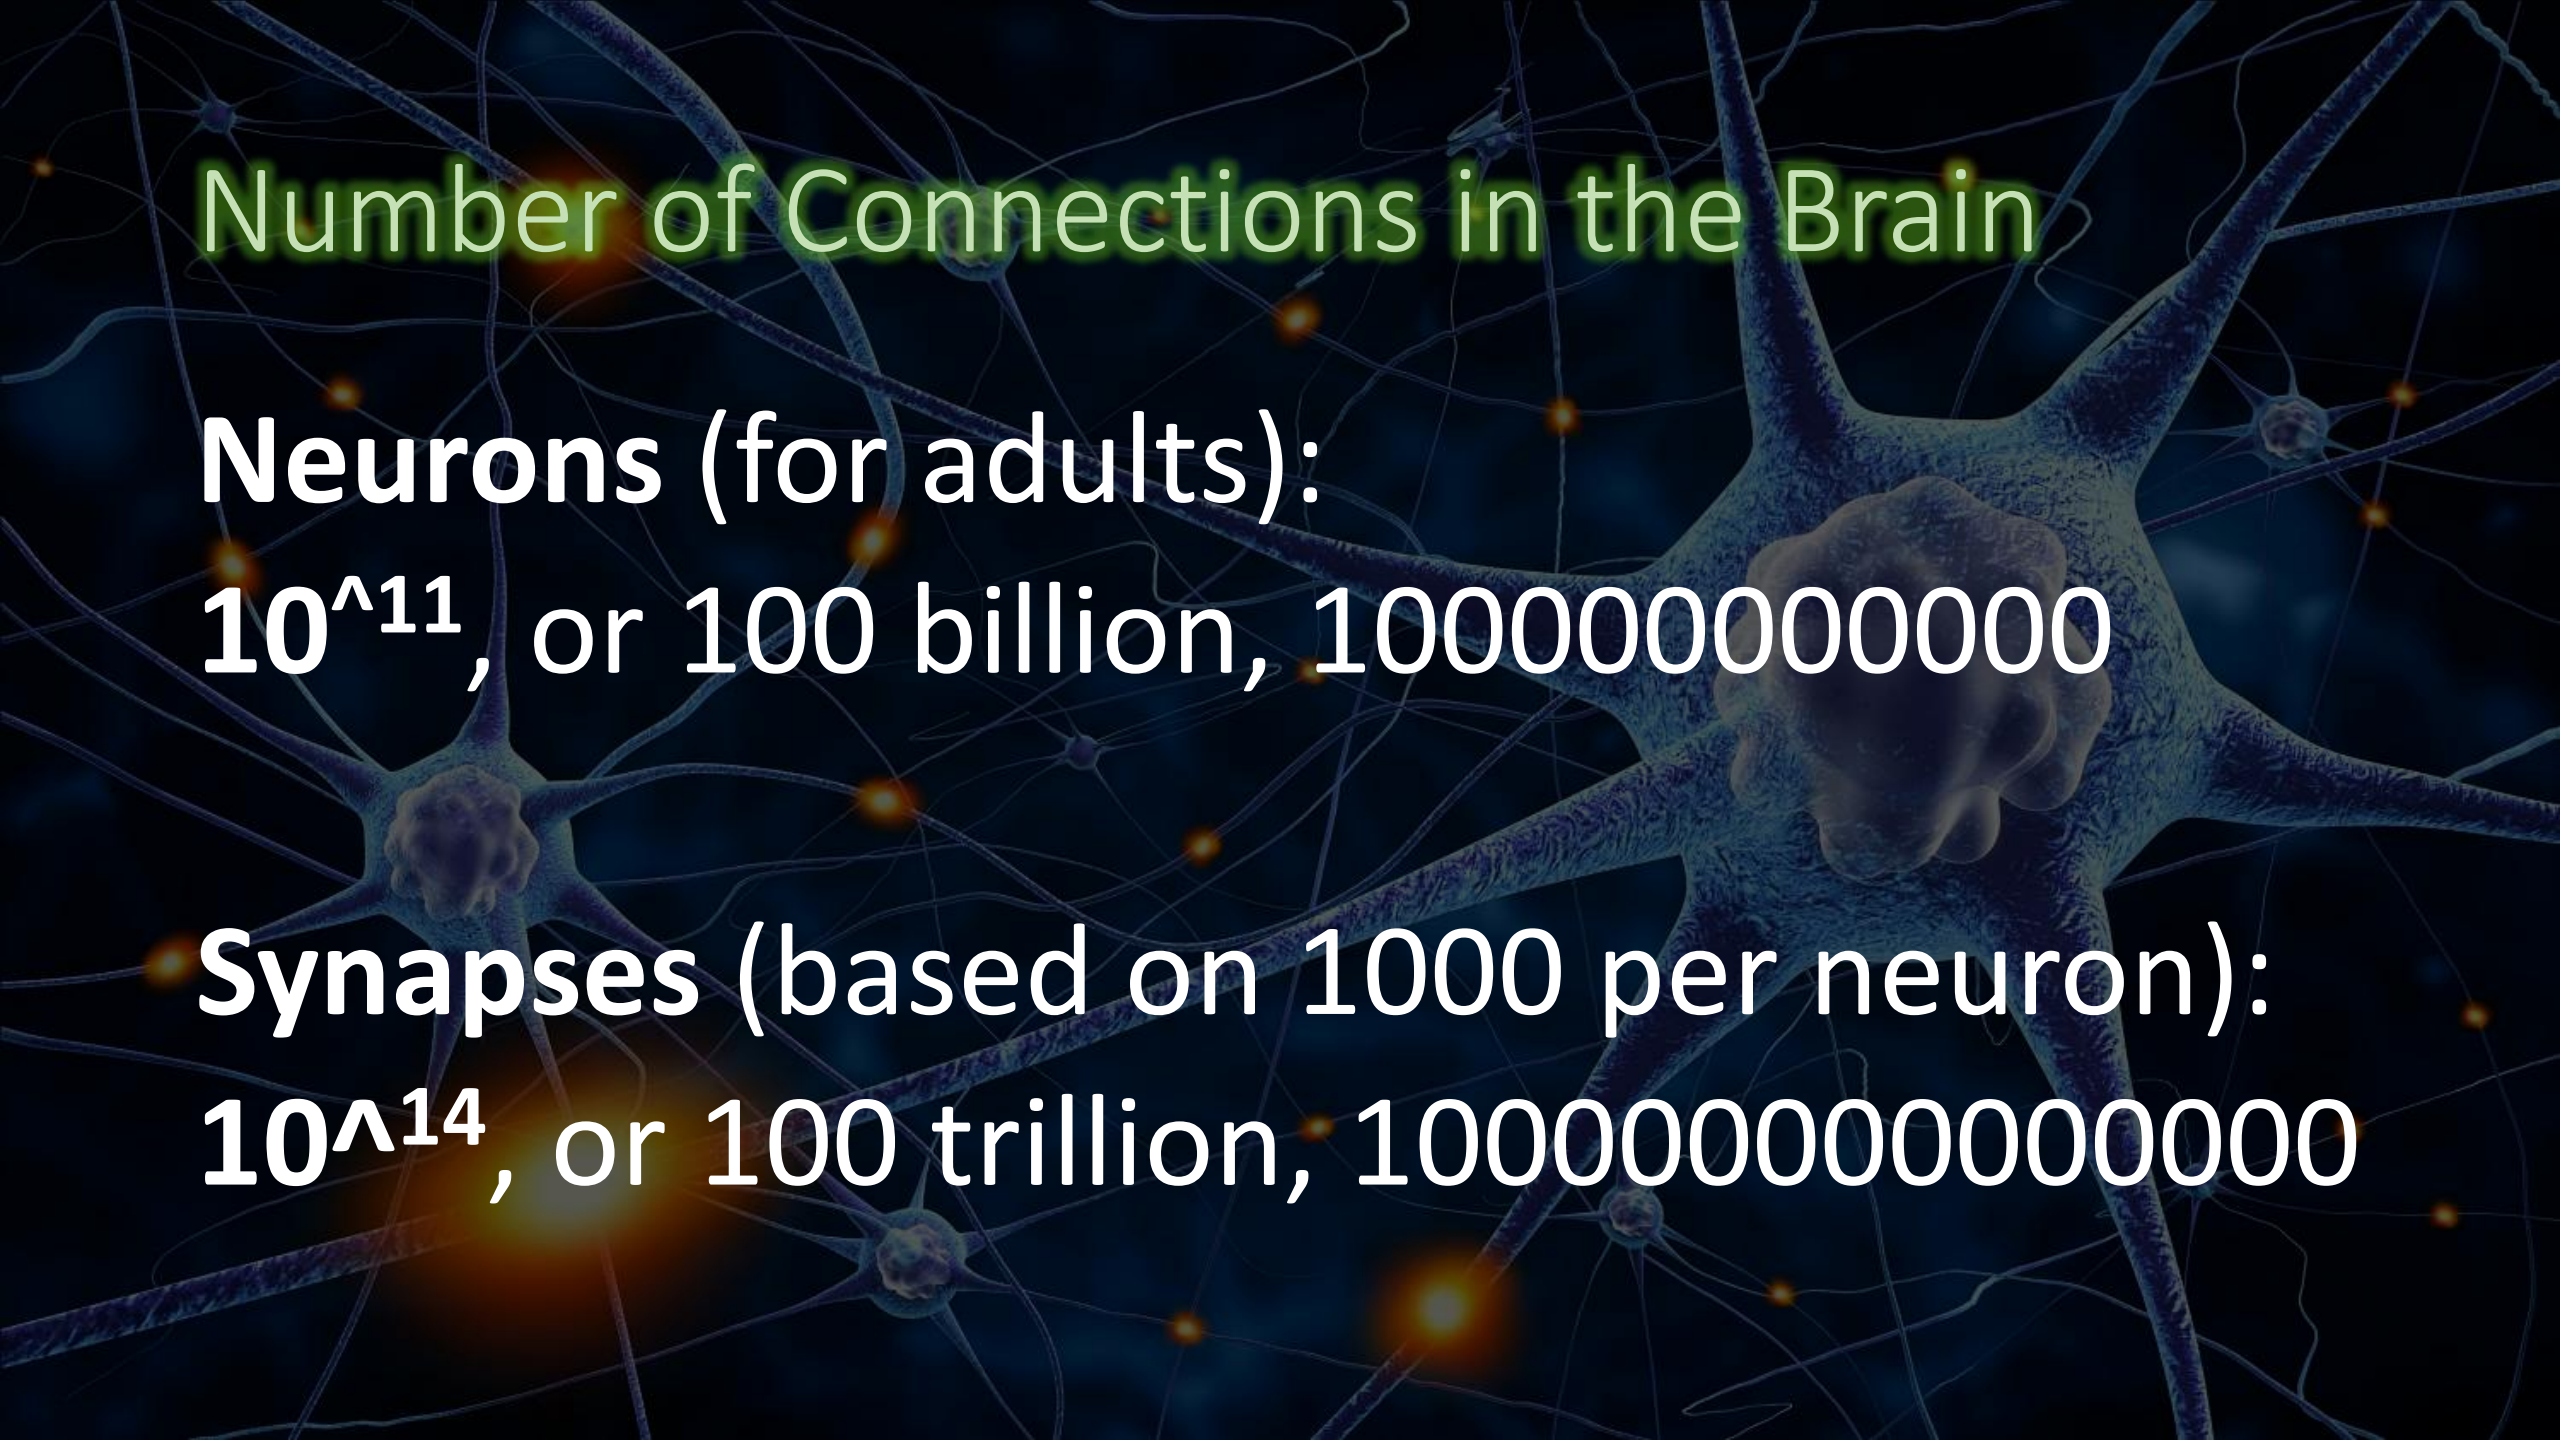

Number of Connections in the Brain

Neurons (for adults):
10^11, or 100 billion, 100000000000

Synapses (based on 1000 per neuron):
10^14, or 100 trillion, 100000000000000

Frank Rosenblatt's Perceptron (1957)

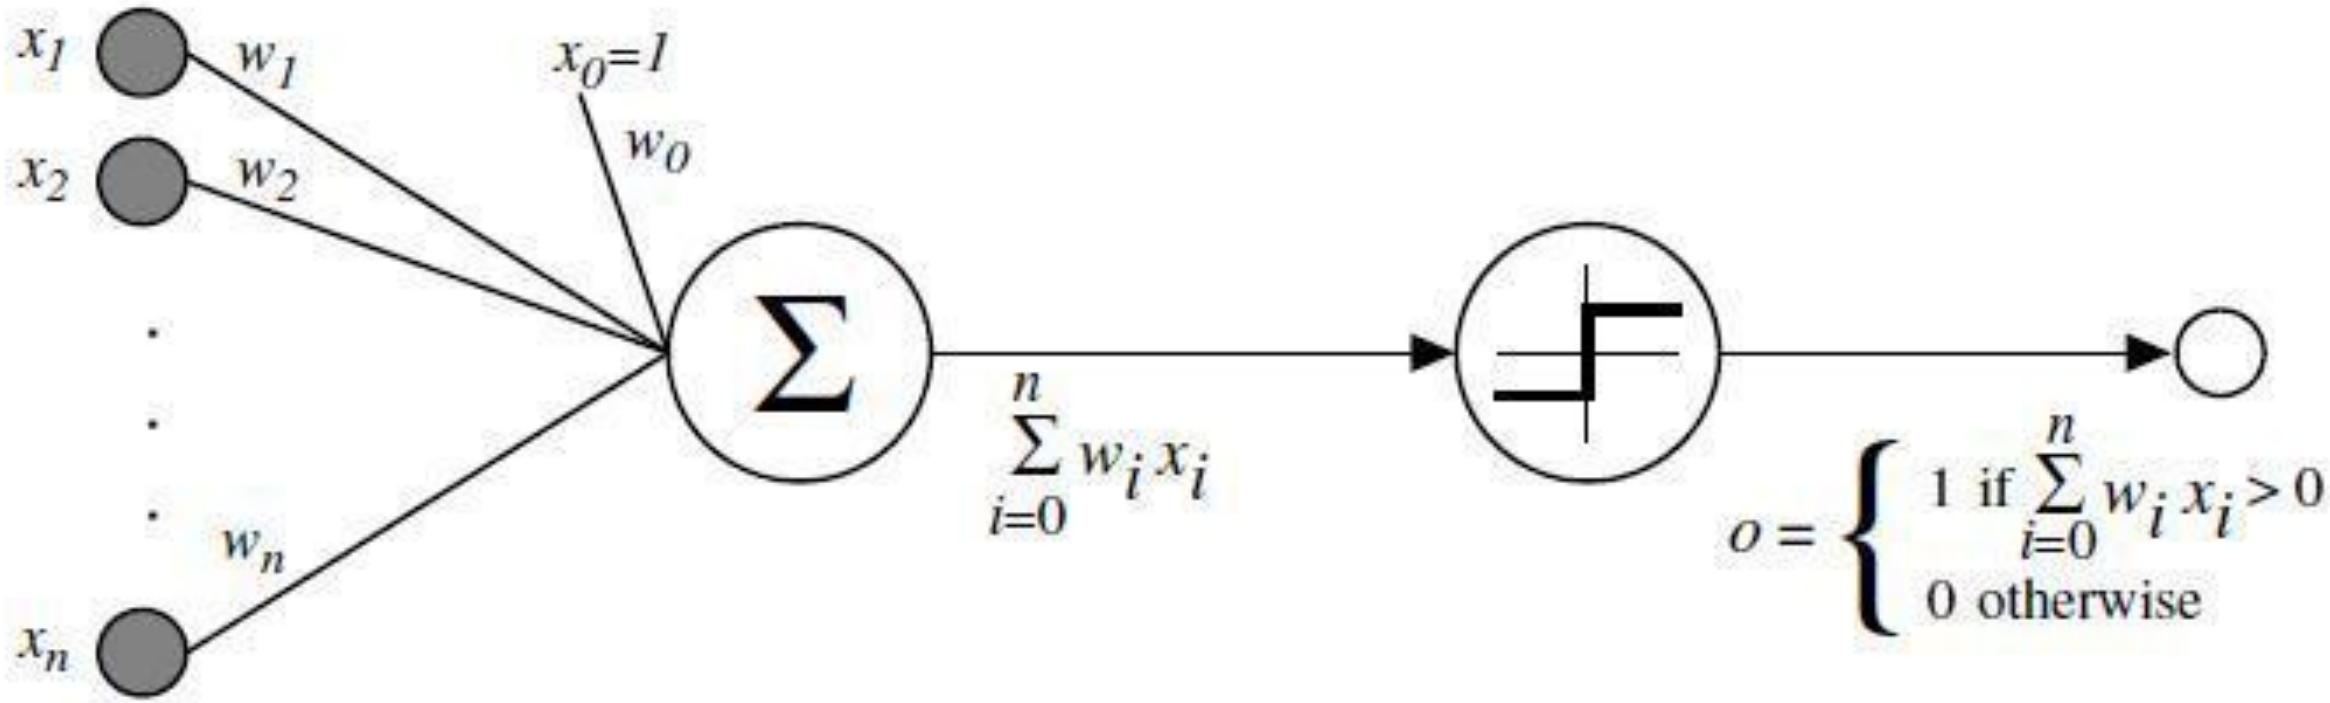

Modern Neuron

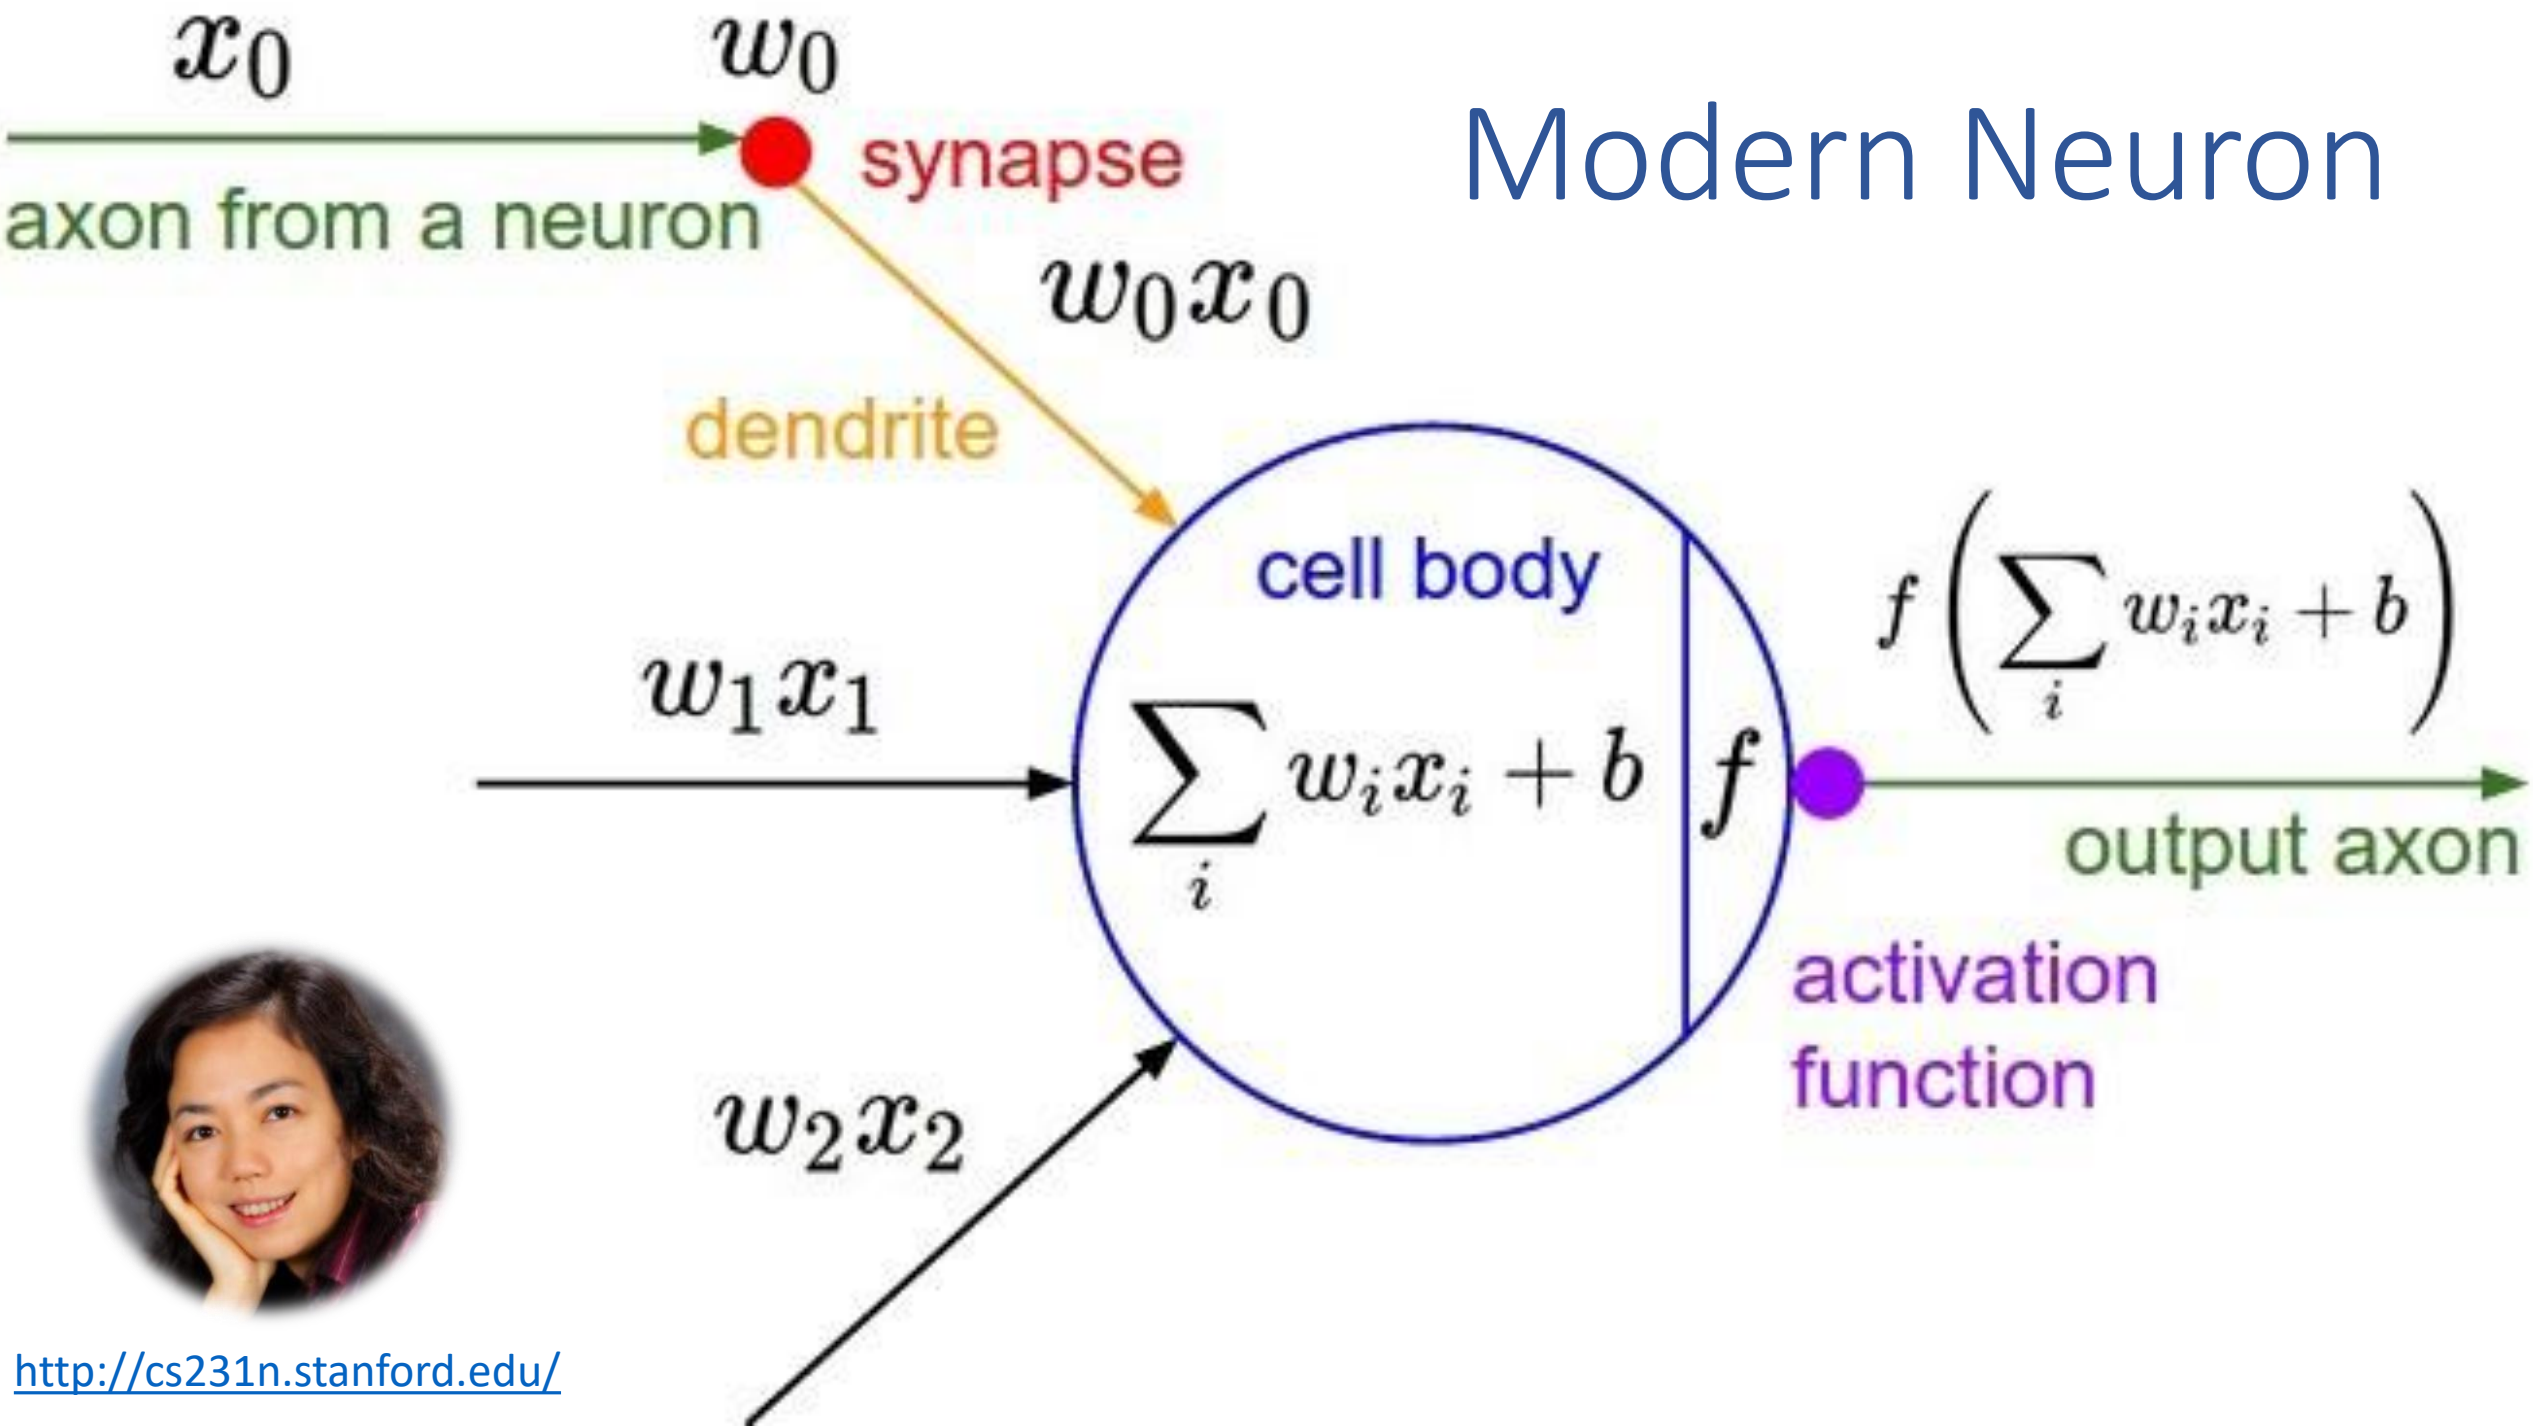

http://cs231n.stanford.edu/

1969: Perceptrons can't do XOR!

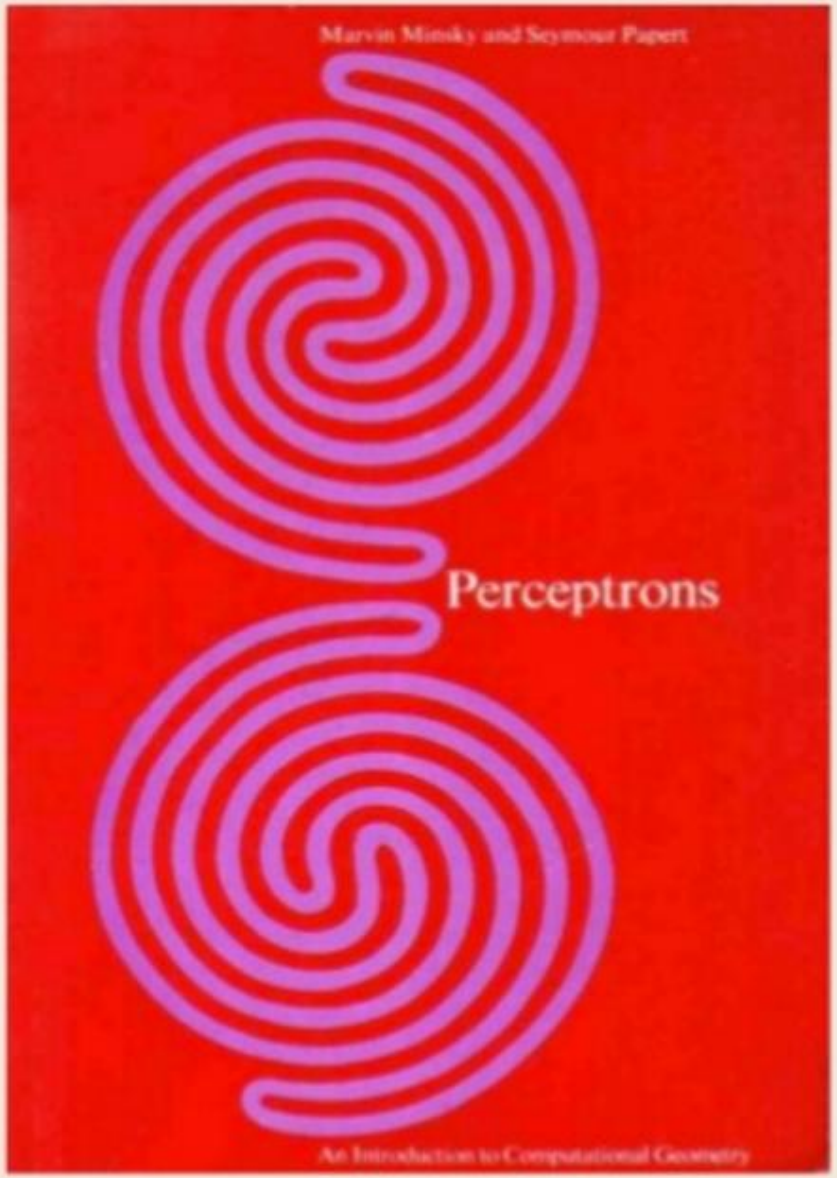

http://www.i-programmer.info/images/stories/BabBag/AI/book.jpg

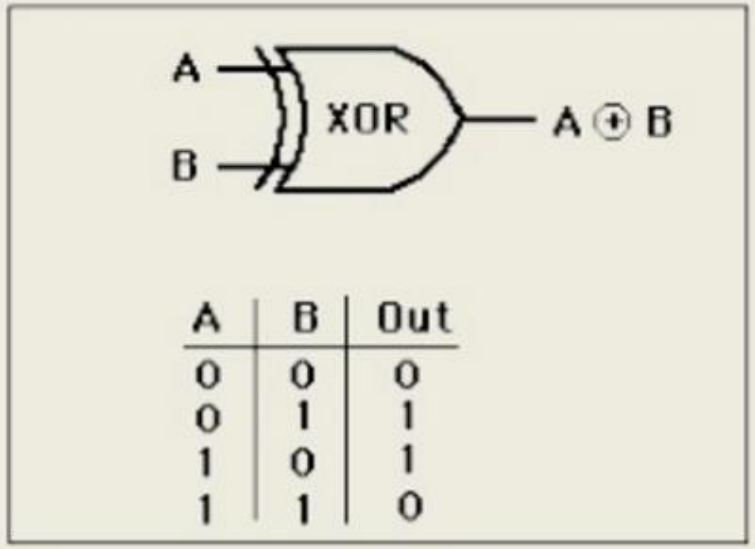

A	B	Out
0	0	0
0	1	1
1	0	1
1	1	0

http://hyperphysics.phy-astr.gsu.edu/hbase/electronic/ietron/xor.gif

Minsky & Papert

https://constructingkids.files.wordpress.com/2013/05/minsky-papert-71-csolomon-x640.jpg


AI Winter
1969 - 1990

Deep Learning

Geoffrey Hinton
(Toronto, Google)

Yann LeCun
(NewYork, Facebook)

Yoshua Bengio
(Montreal)

Learning XOR (1986)

Geoffrey Hinton

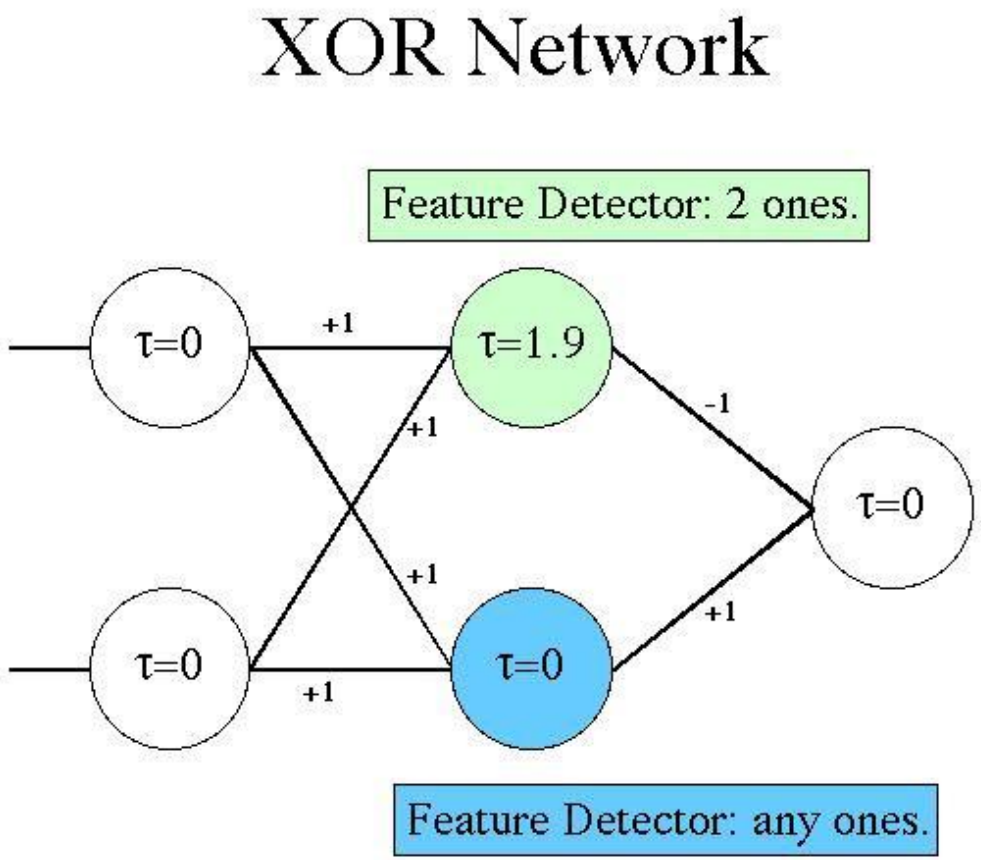

https://torontolife.com/life/ai-superstars-google-facebook-apple-studied-guy/

Backpropagation

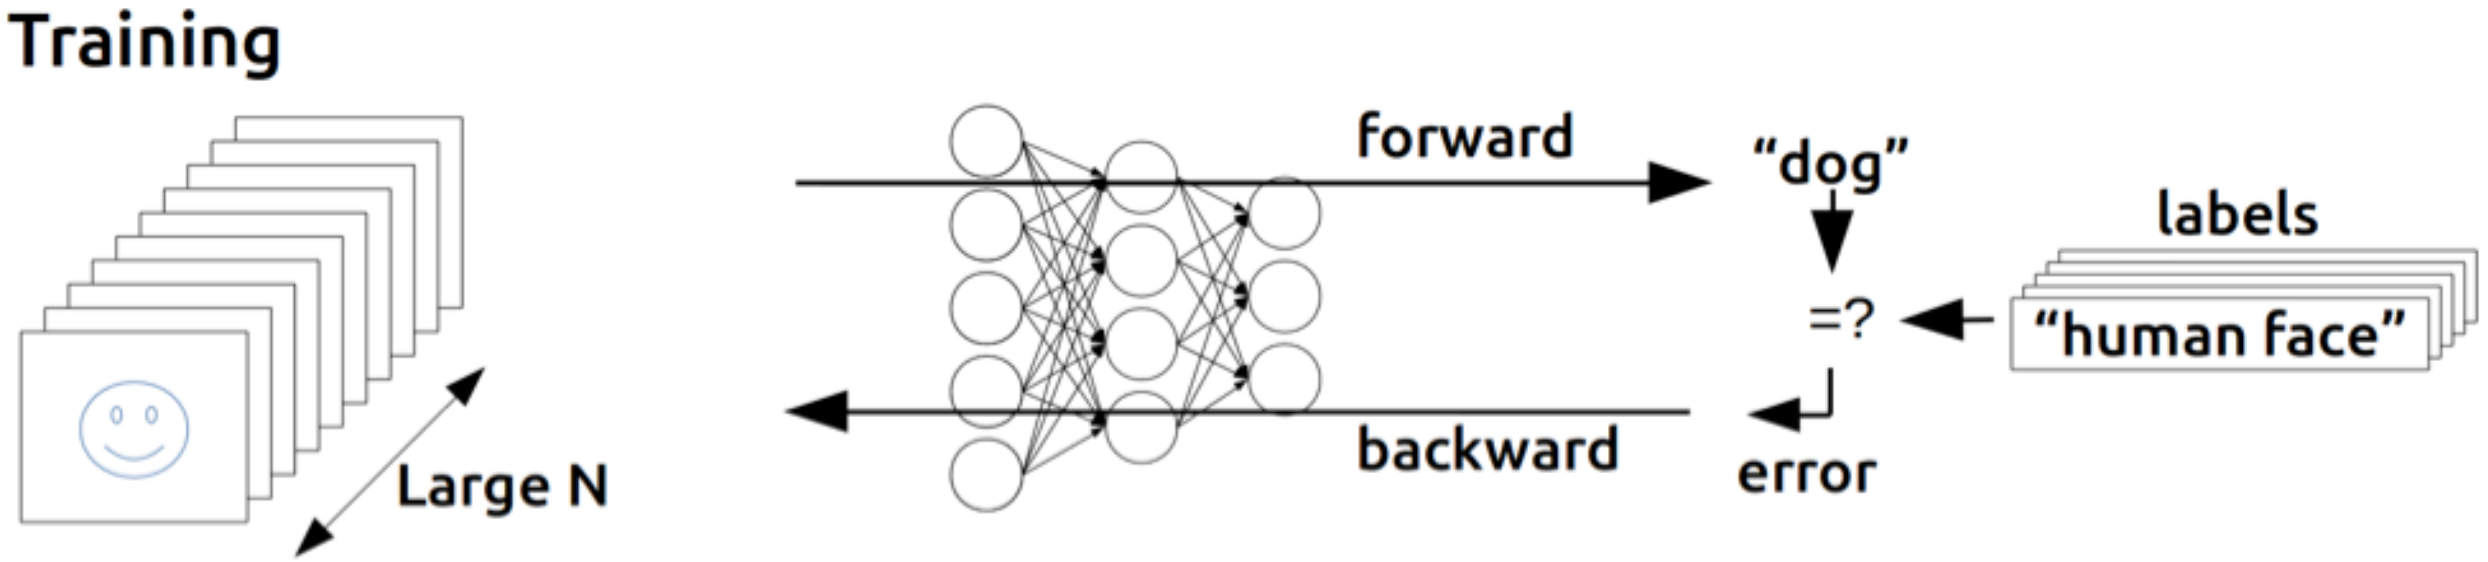

Inference

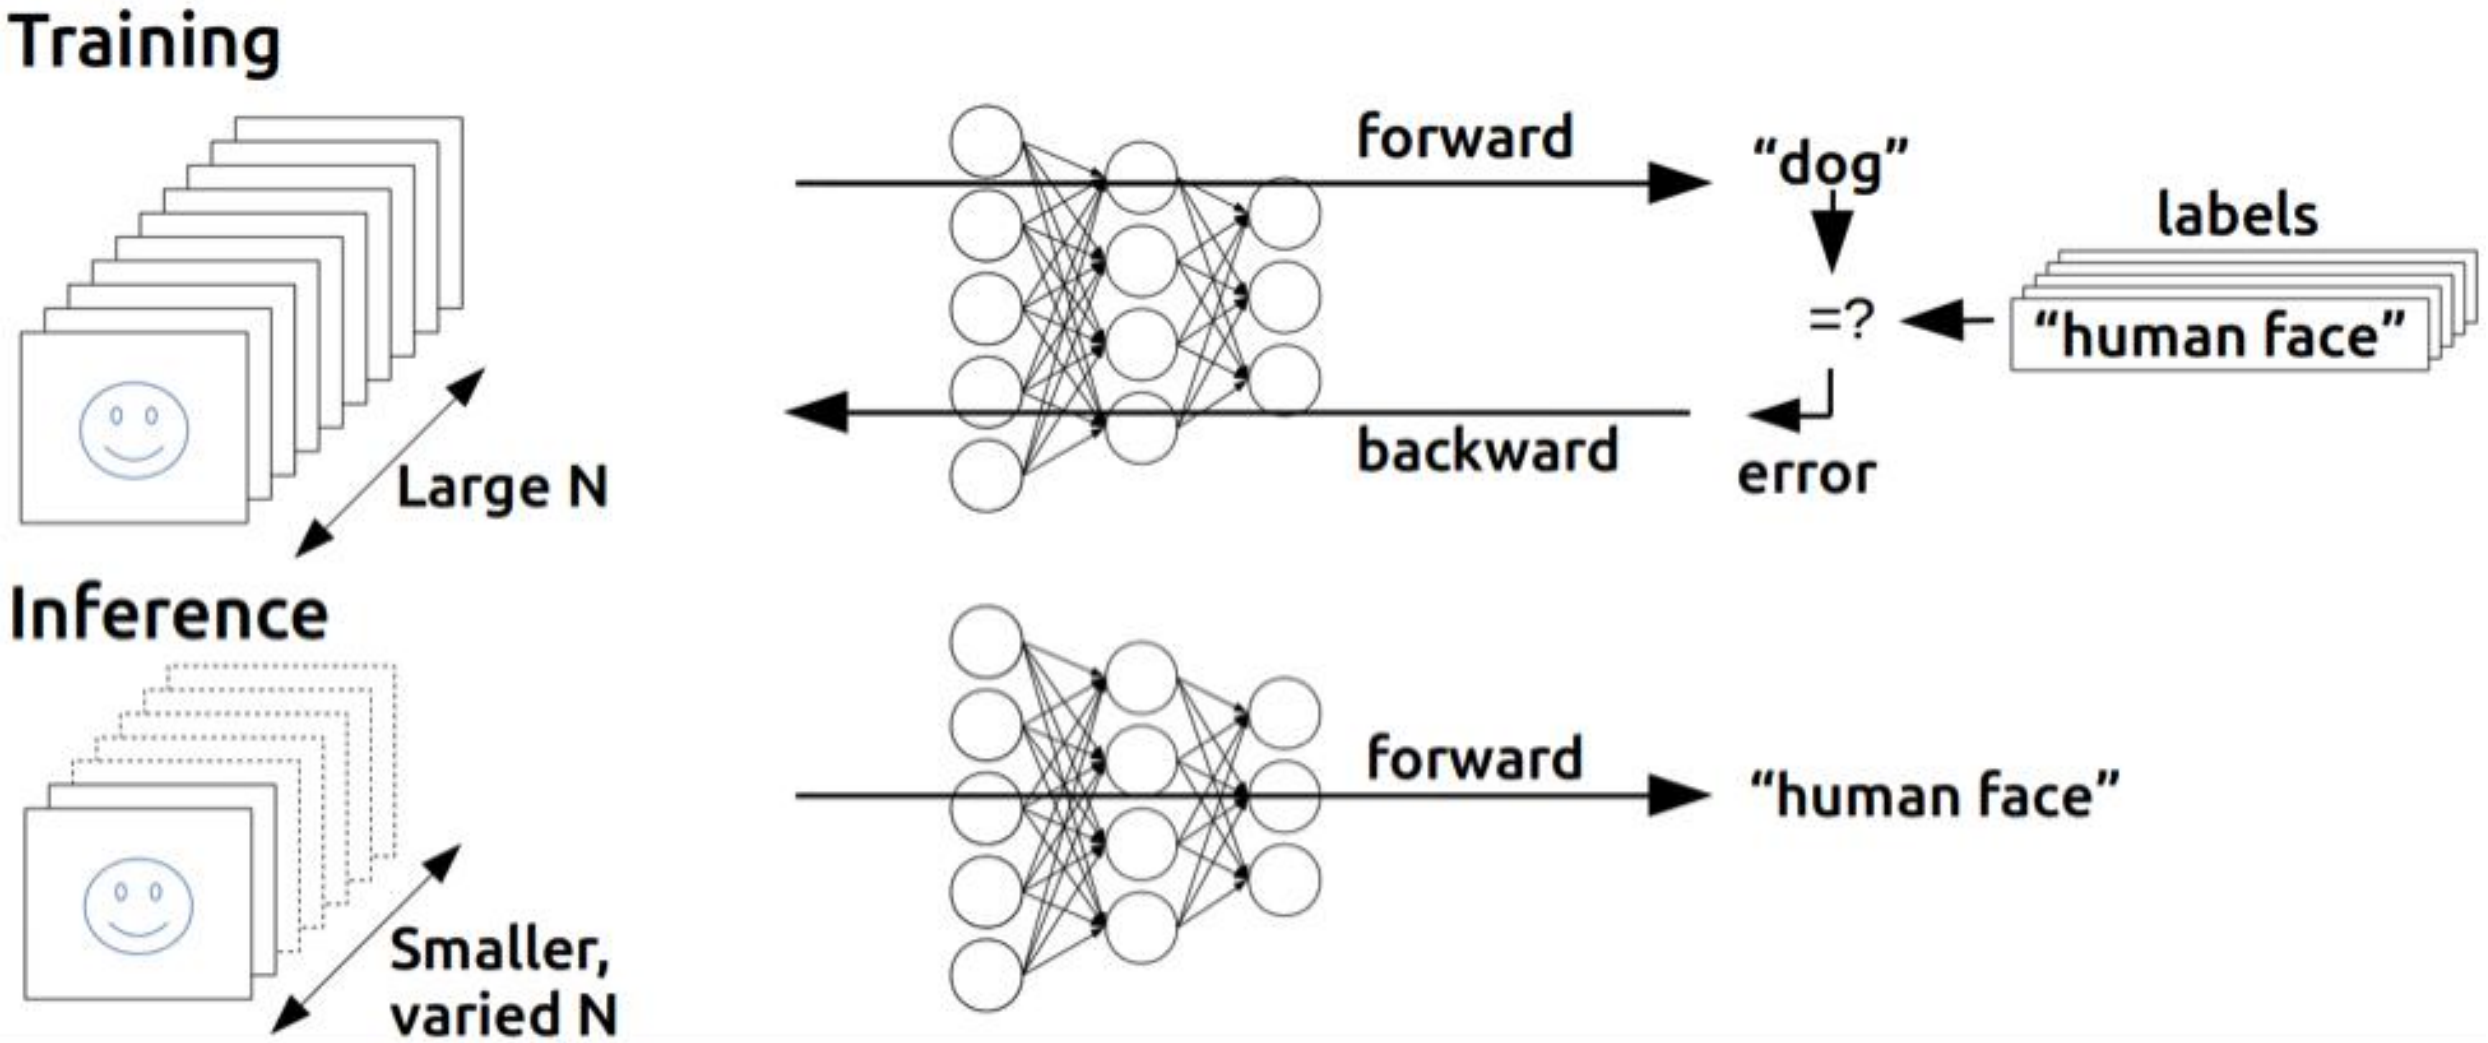
Training
Large N
forward
"dog"
labels
=?
"human face"
backward
error
Inference
Smaller,
varied N
forward
"human face"

Chain Rule

$$\frac{dy}{dx} = \frac{dy}{du}\frac{du}{dx}$$

$$\frac{d^2y}{dx^2} = \frac{d^2y}{du^2}\left(\frac{du}{dx}\right)^2 + \frac{dy}{du}\frac{d^2u}{dx^2}$$

$$\frac{d^3y}{dx^3} = \frac{d^3y}{du^3}\left(\frac{du}{dx}\right)^3 + 3\,\frac{d^2y}{du^2}\frac{du}{dx}\frac{d^2u}{dx^2} + \frac{dy}{du}\frac{d^3u}{dx^3}$$

$$\frac{d^4y}{dx^4} = \frac{d^4y}{du^4}\left(\frac{du}{dx}\right)^4 + 6\,\frac{d^3y}{du^3}\left(\frac{du}{dx}\right)^2\frac{d^2u}{dx^2} + \frac{d^2y}{du^2}\left(4\,\frac{du}{dx}\frac{d^3u}{dx^3} + 3\left(\frac{d^2u}{dx^2}\right)^2\right) + \frac{dy}{du}\frac{d^4u}{dx^4}.$$

Example: Recognizing Handwritten Digits

- MNIST dataset

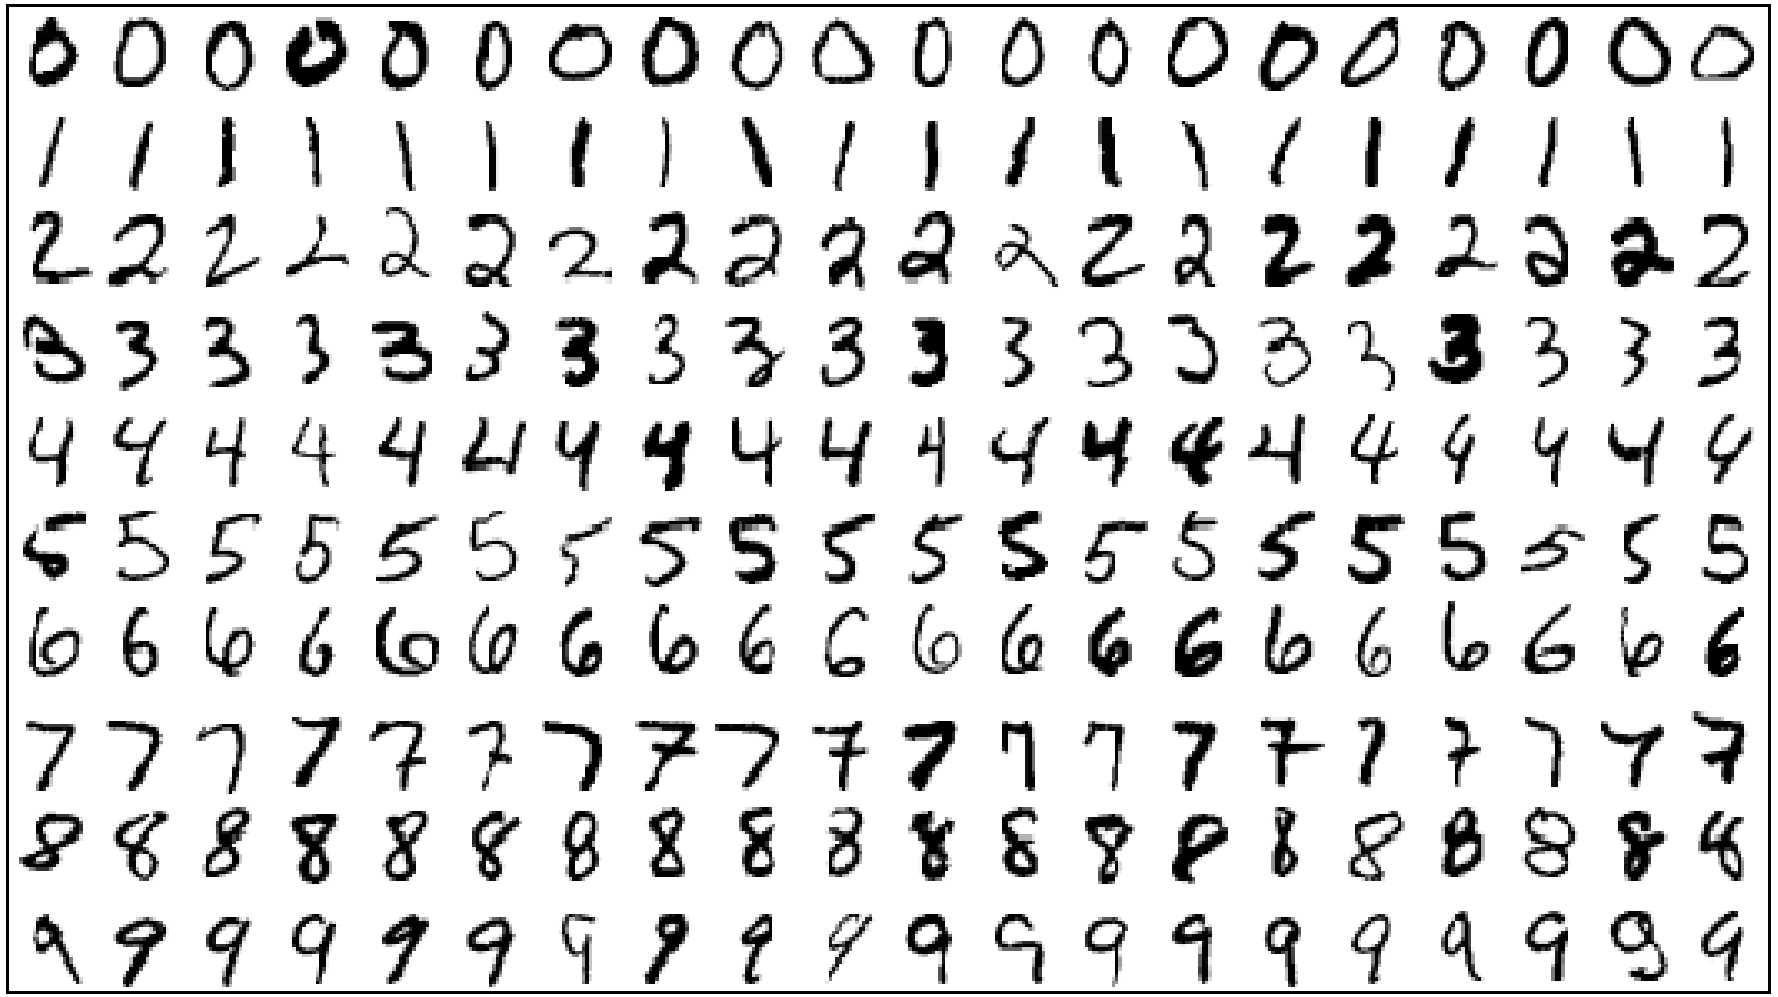

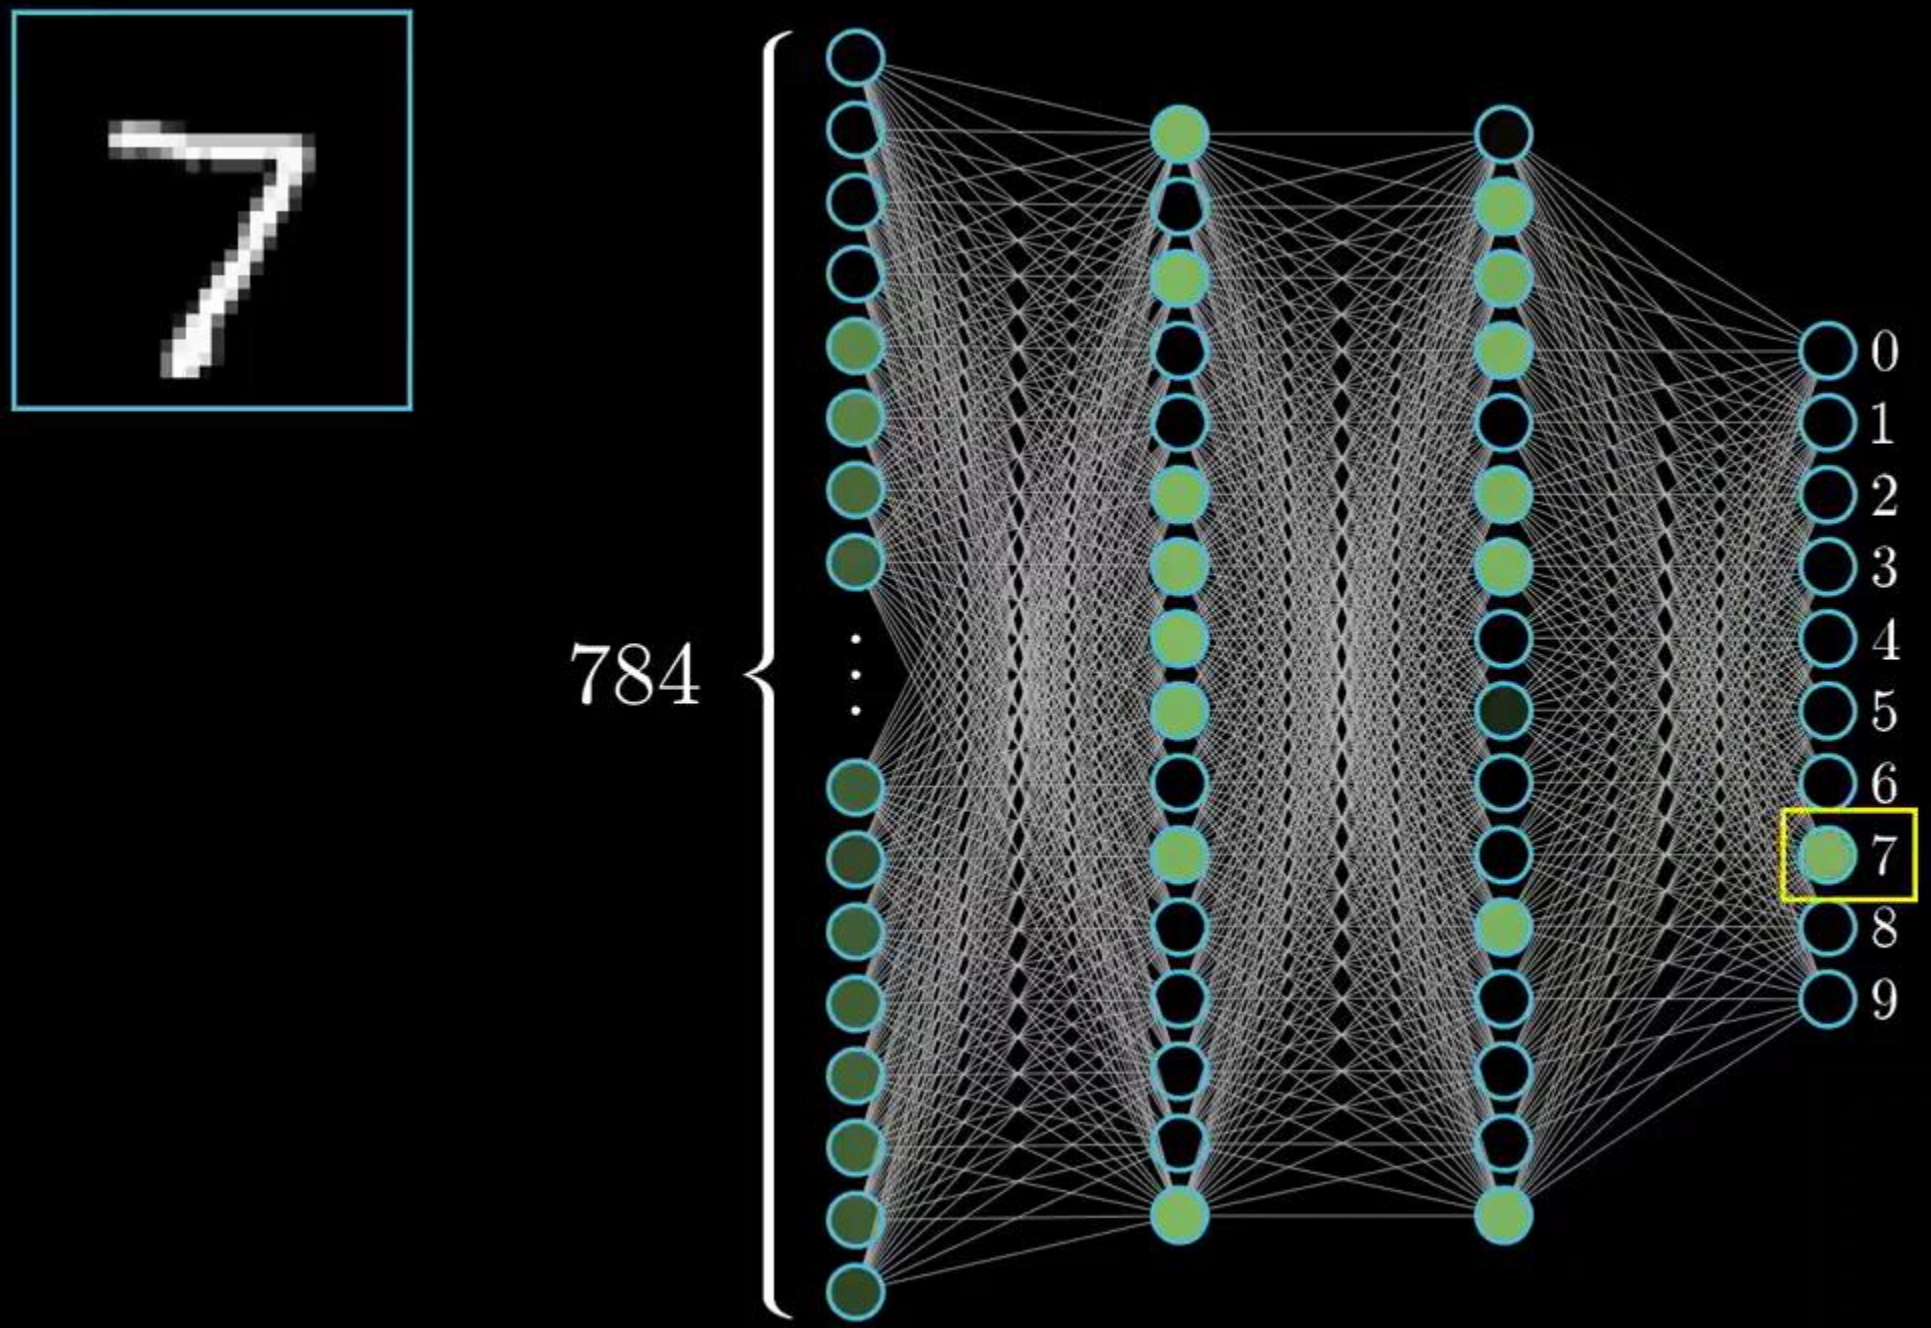
784
0
1
2
3
4
5
6
7
8
9

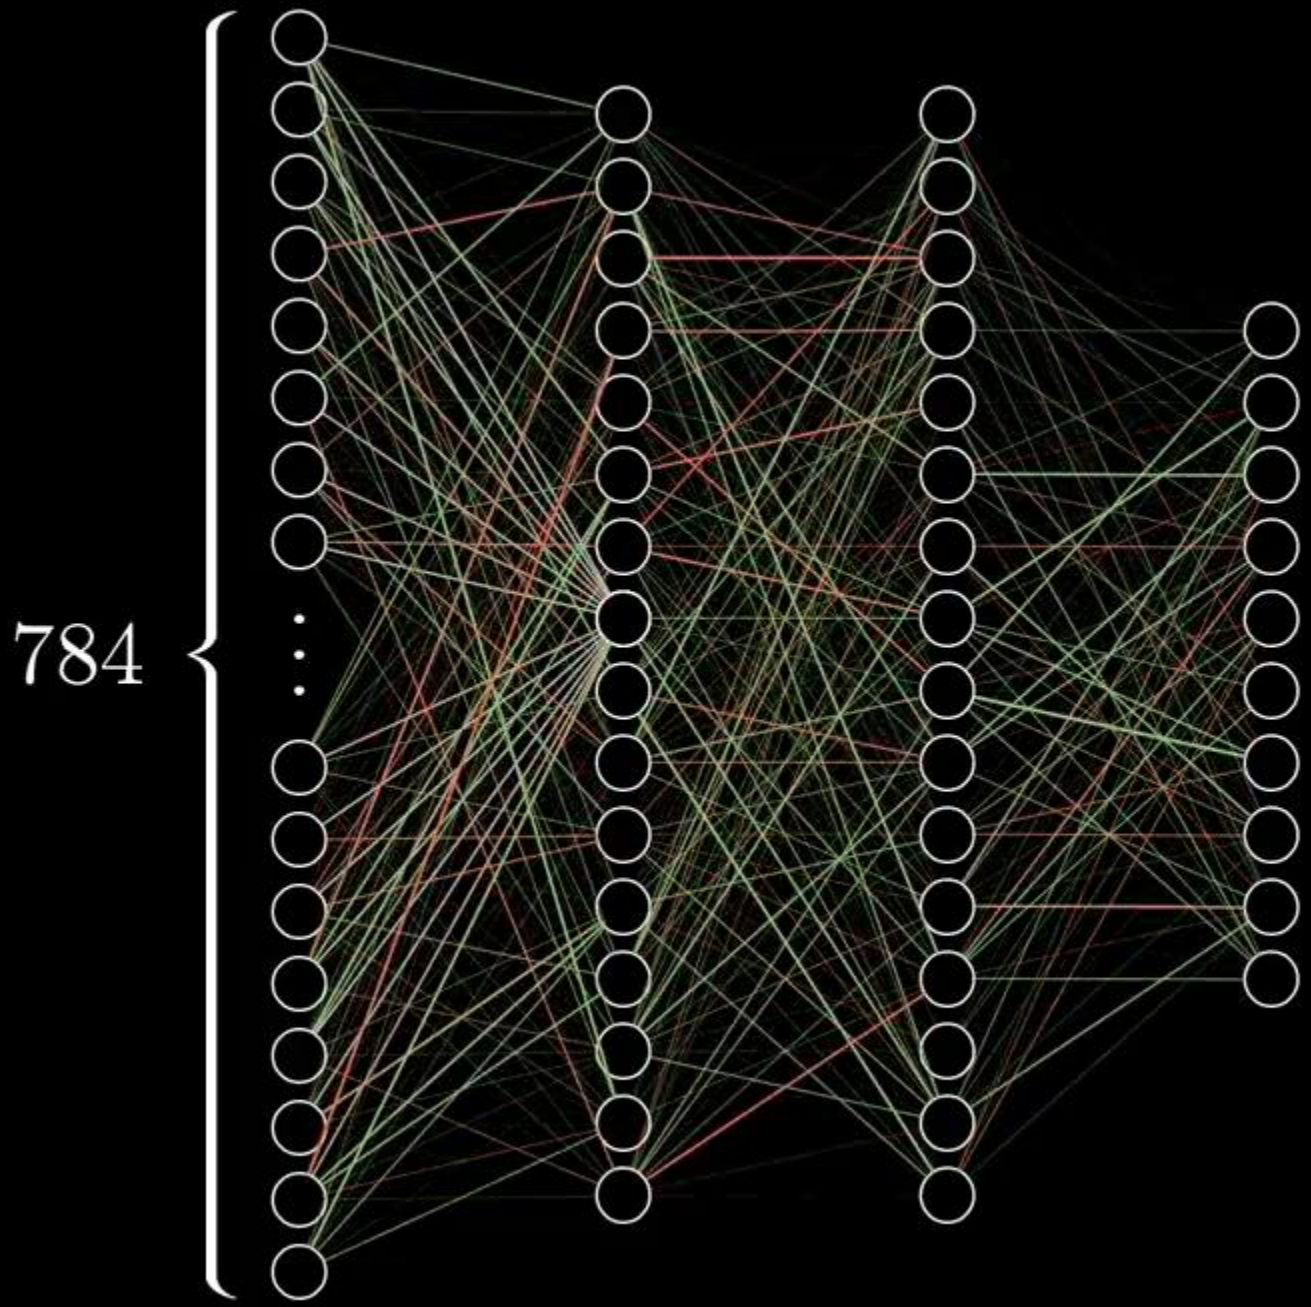

$784 \times 16 + 16 \times 16 + 16 \times 10$
weights

$16 + 16 + 10$
biases

$13{,}002$

Learning $\rightarrow$ Finding the right weights and biases

Gradient Descent

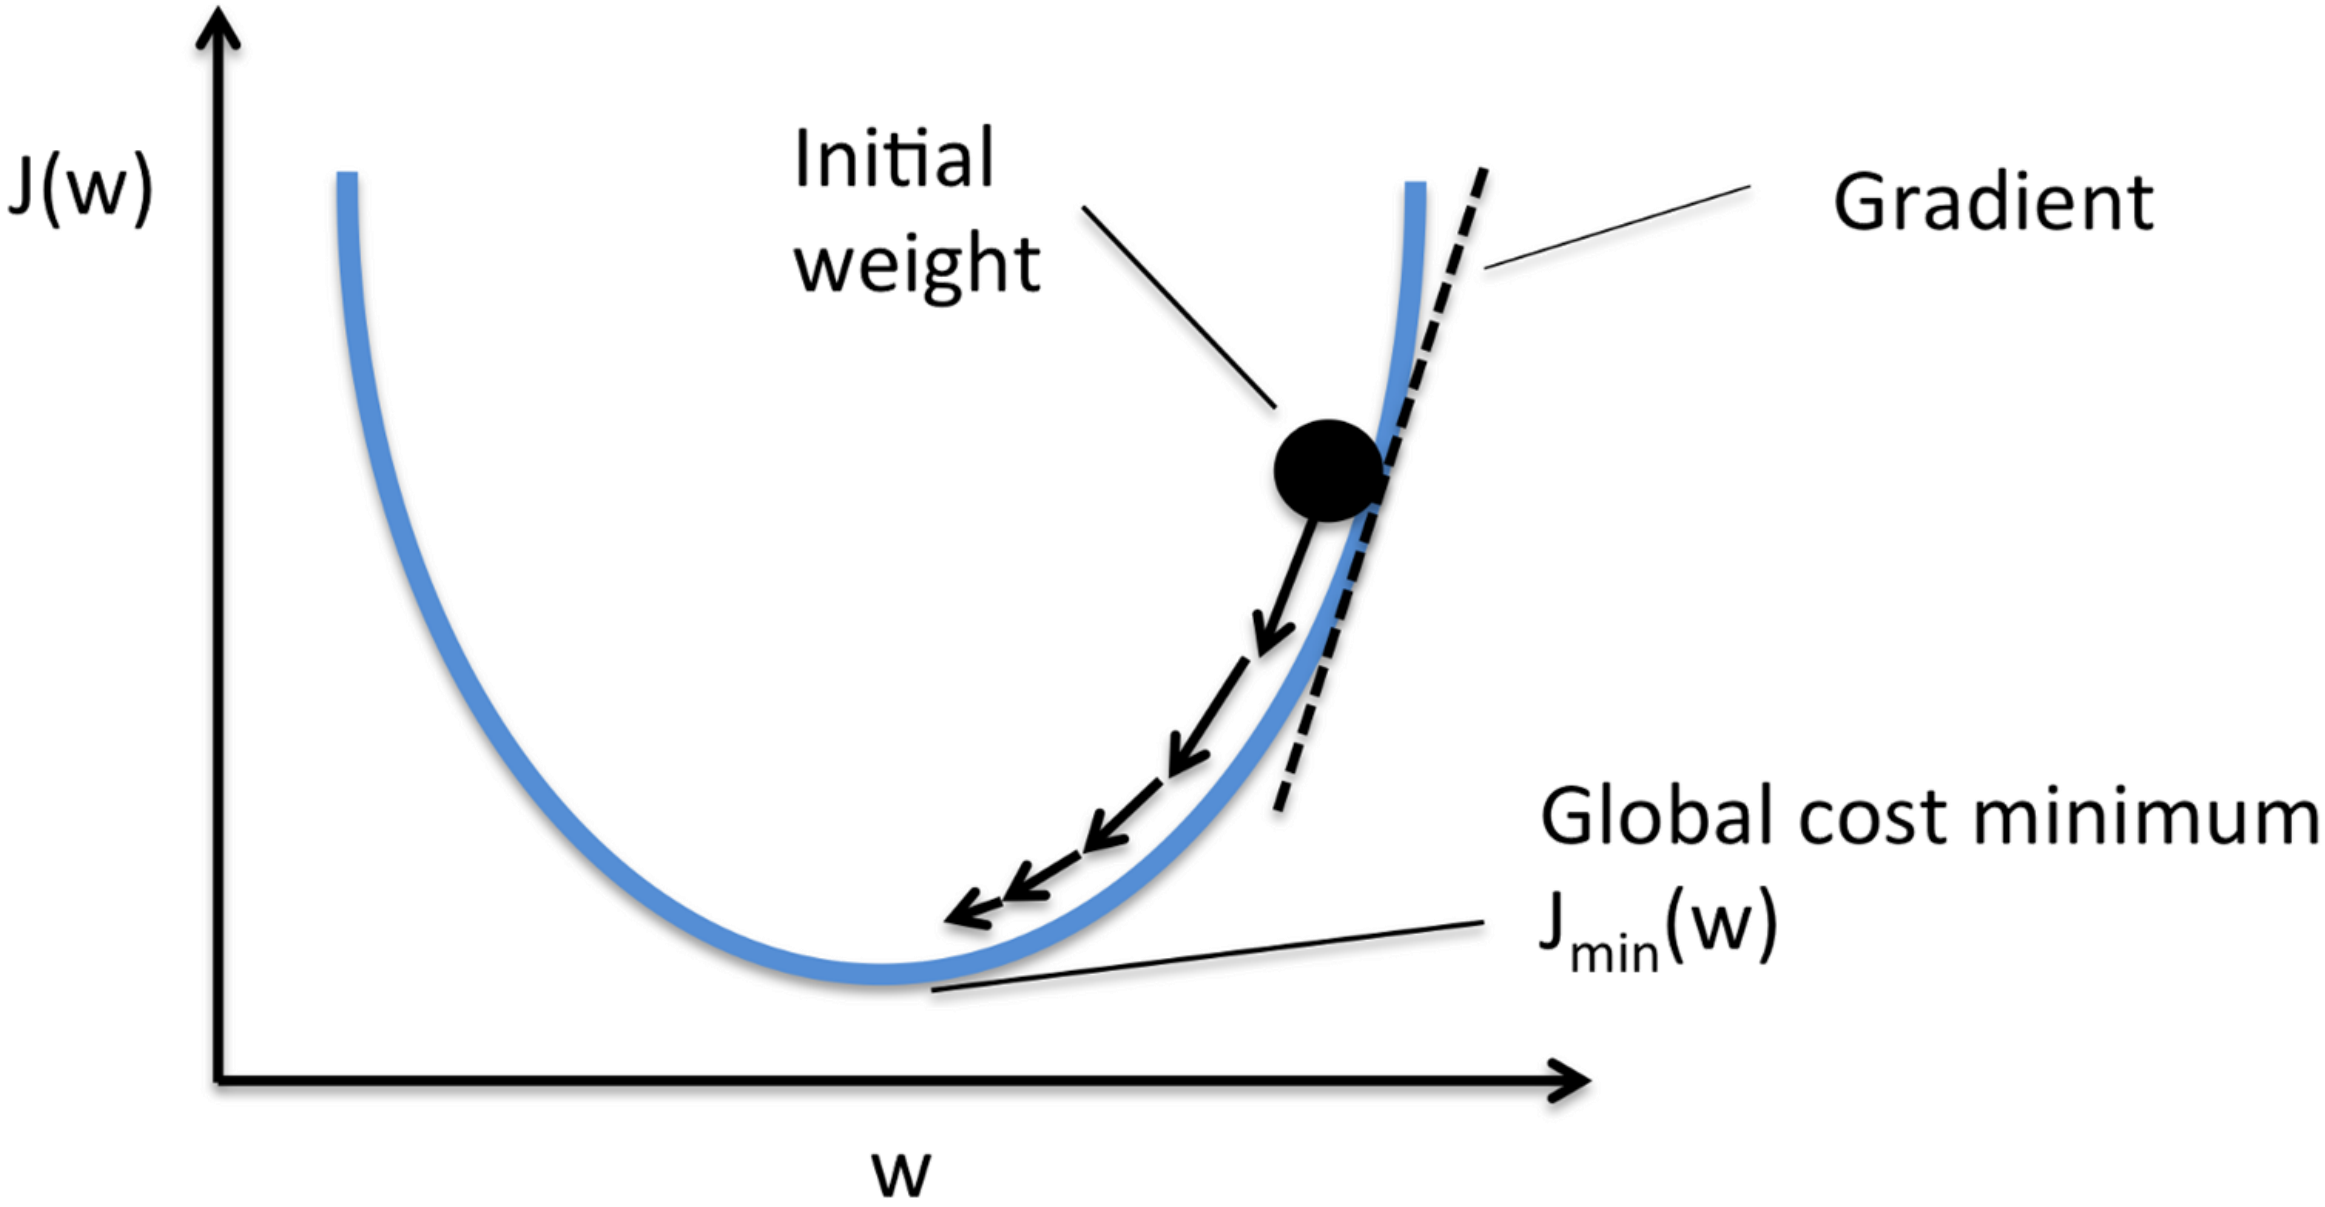

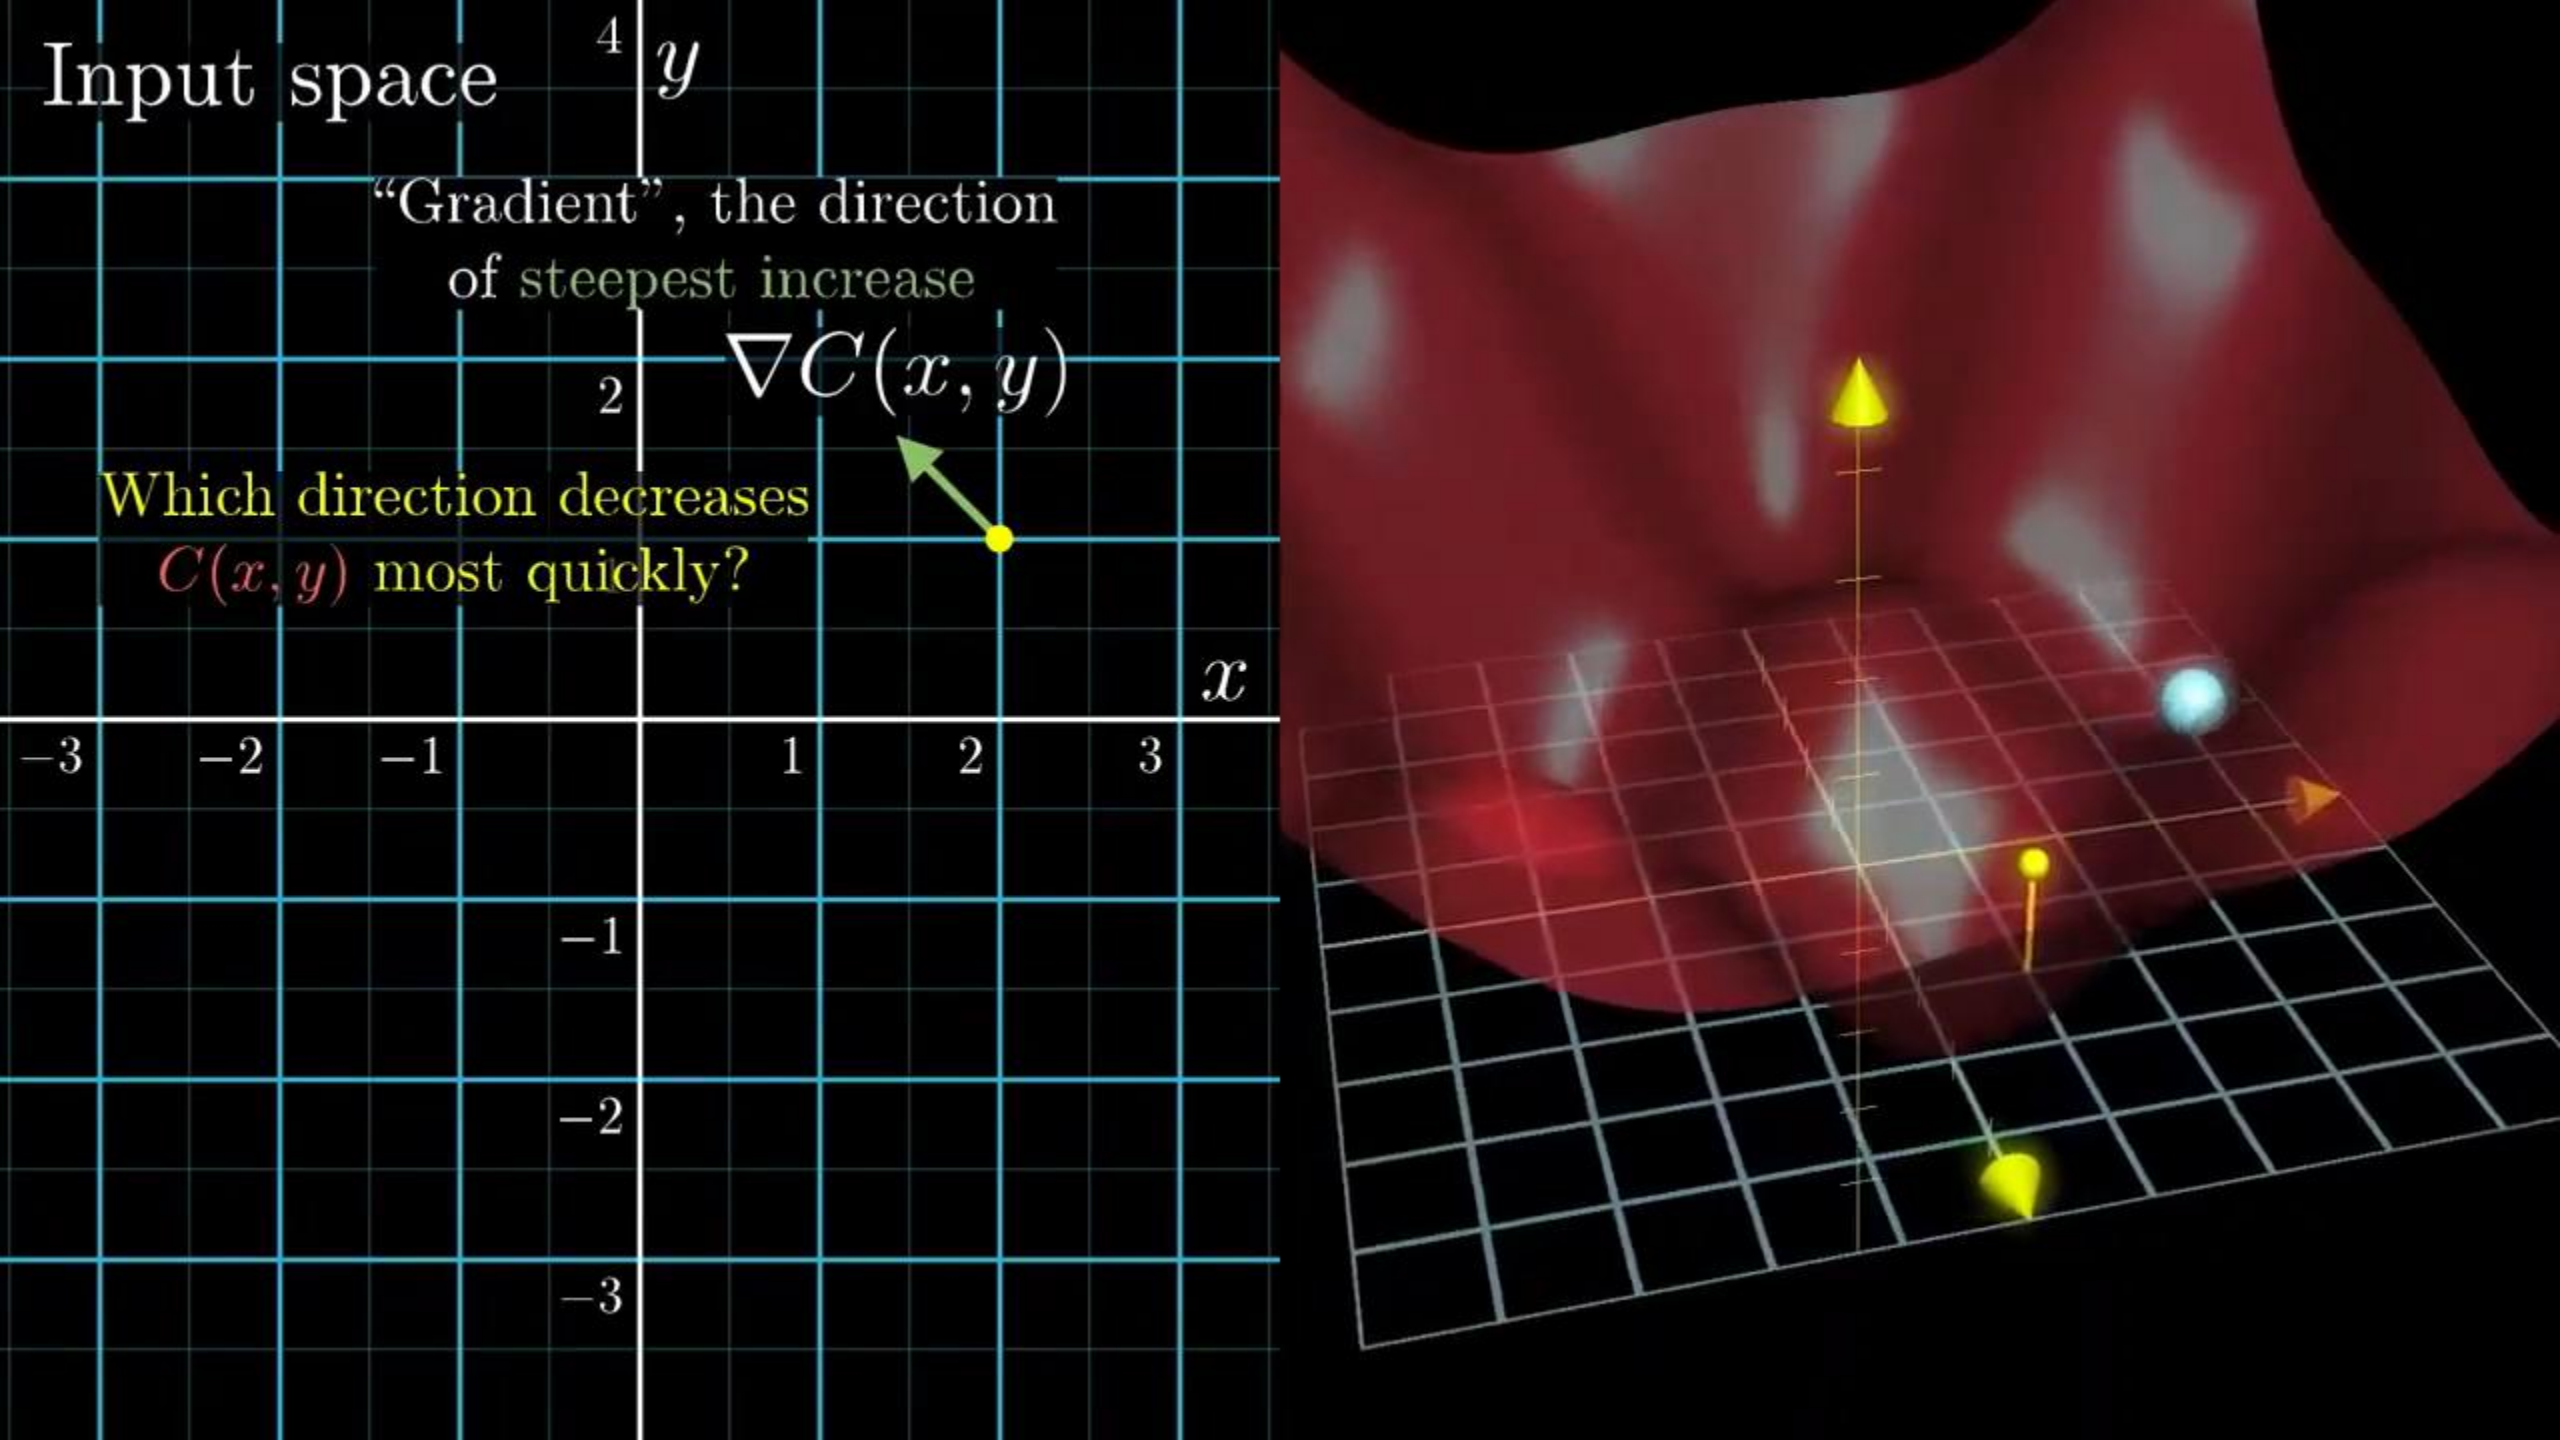
Input space
y
4
2
−1
−2
−3
x
−3
−2
−1
1
2
3
"Gradient", the direction
of steepest increase
$\nabla C(x, y)$
Which direction decreases
$C(x, y)$ most quickly?

Major Types of Neural Networks

- Convolutional Neural Networks (CNN)
- Recurrent Neural Networks (RNN)
- Generative Adversarial Networks (GAN)
- Attention & Transformer

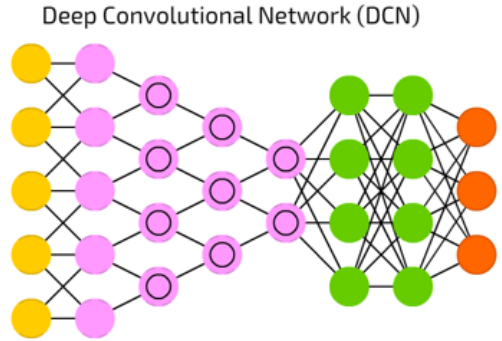

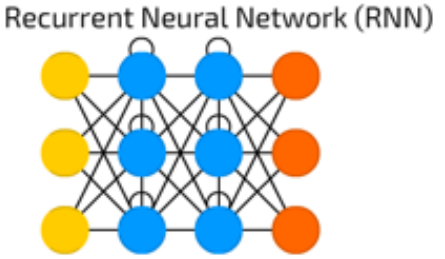

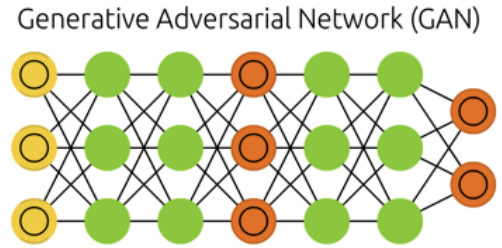

https://www.asimovinstitute.org/neural-network-zoo/

Convolutional Neural Network (LeNet-5)

- https://medium.com/@sh.tsang/paper-brief-review-of-lenet-1-lenet-4-lenet-5-boosted-lenet-4-image-classification-1f5f809dbf17

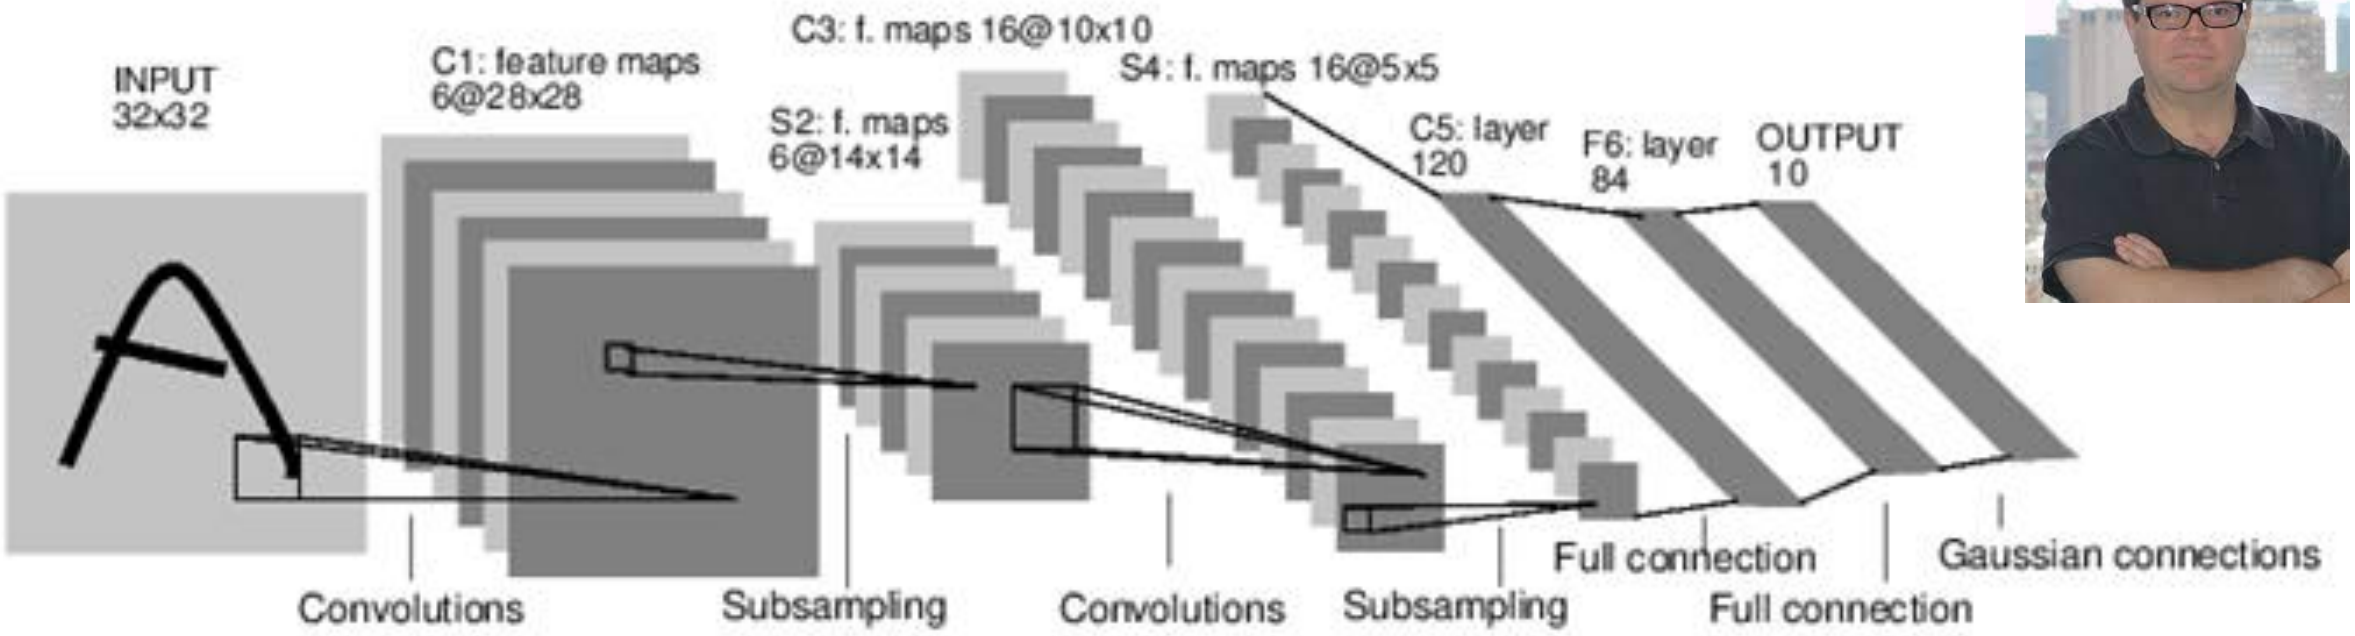

A Full Convolutional Neural Network (LeNet)


INTERNET

ImageNet Large Scale Visual Object Recognition Challenge (ILSVRC)

- 1000 categories
- For ILSVRC 2017
  - **Training images** for each category ranges from 732 to 1300
  - 50,000 validation **images** and 100,000 test **images**.
- Total number of images in ILSVRC 2017 is around 1,150,000

Error Rate on ImageNet Challenge (~2011)

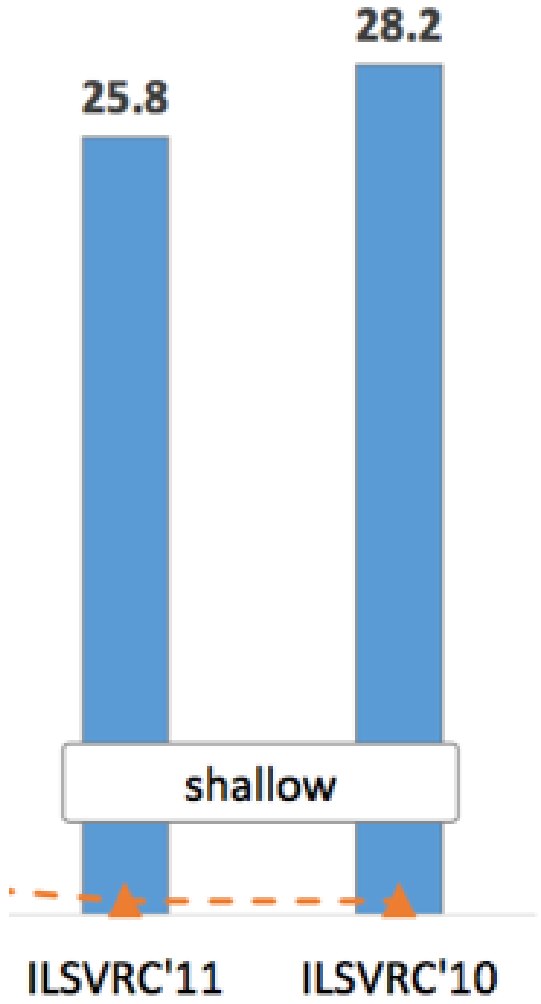

Deep Convolutional Neural Network (AlexNet)

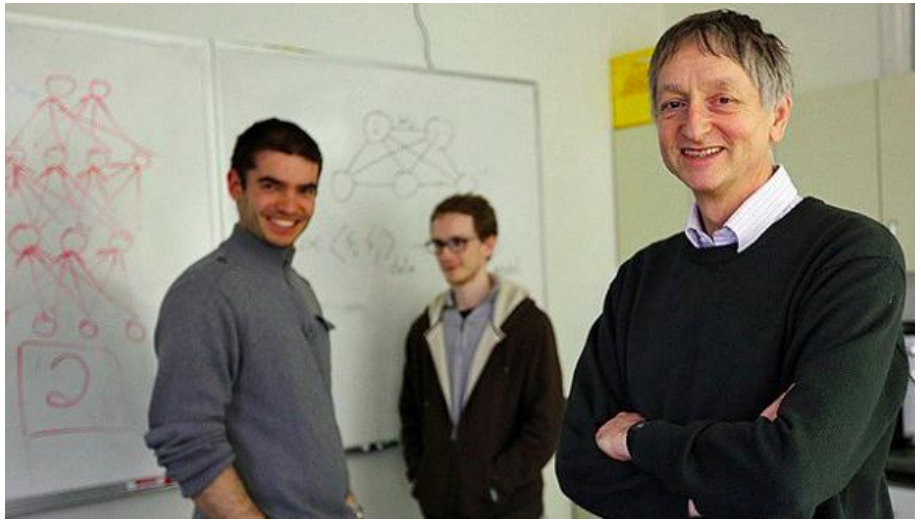

- Alex Krizhevsky, Ilya Suskever, Geoffrey Hinton, 2012

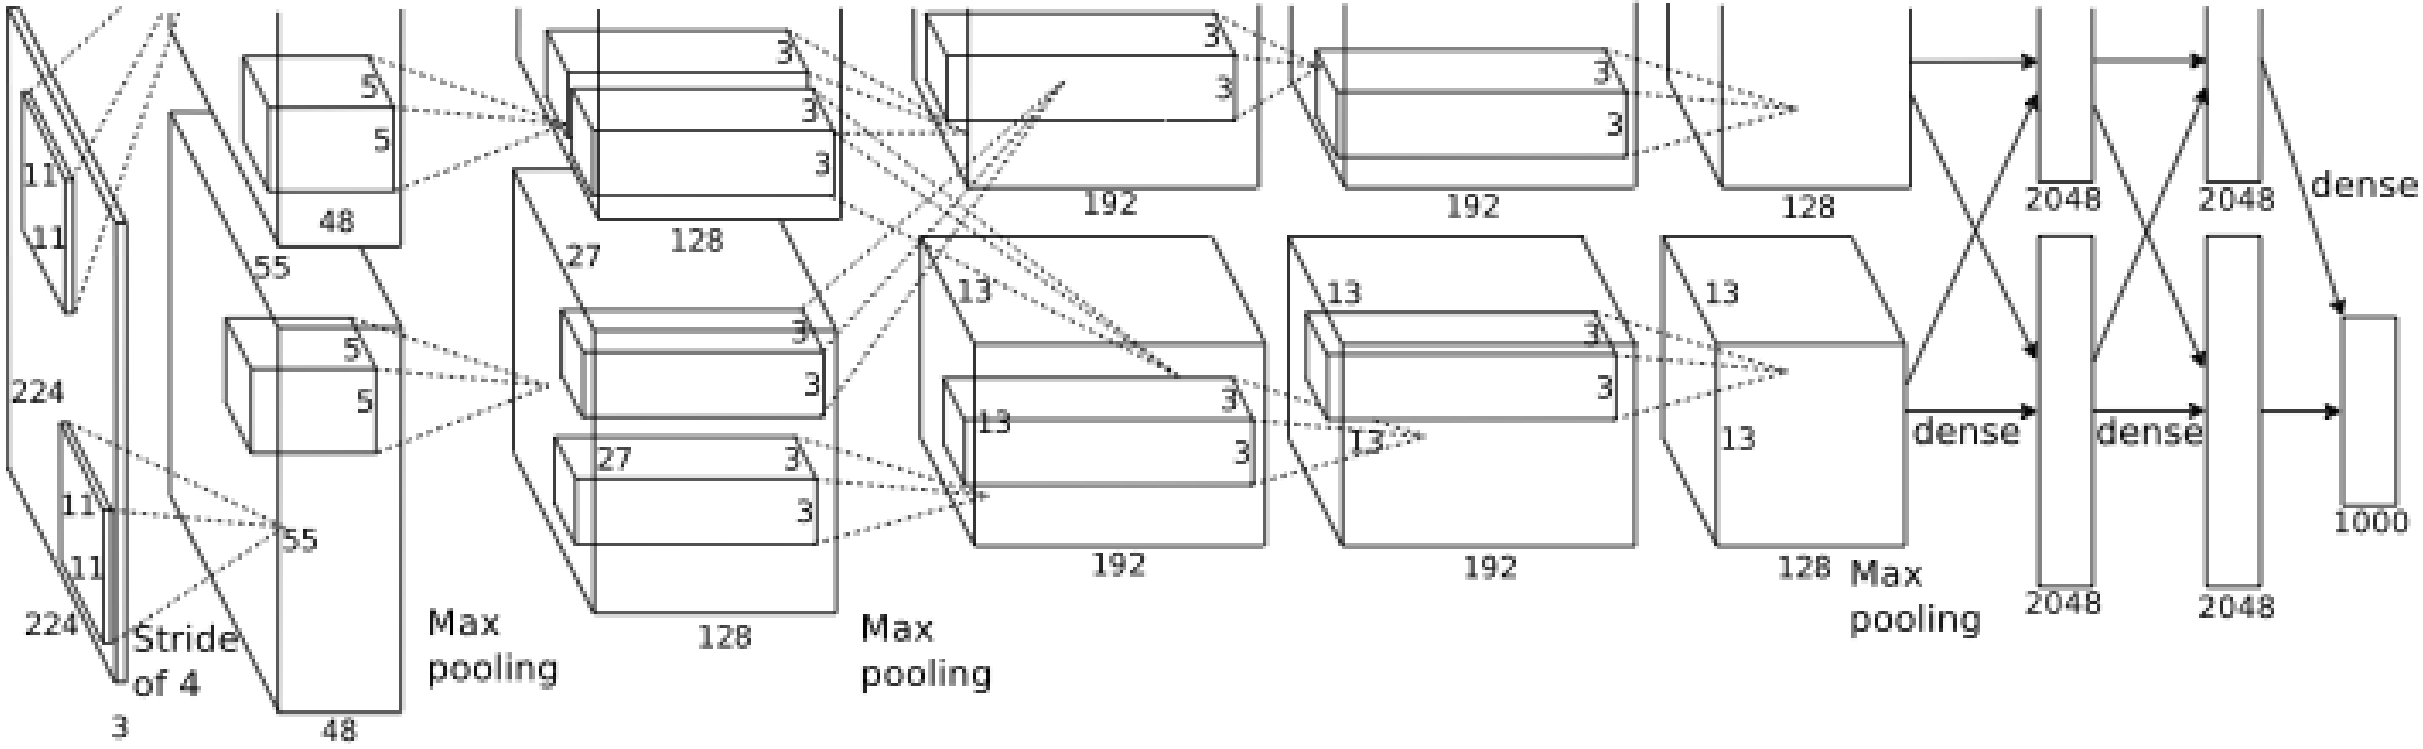

Error Rate on ImageNet Challenge (~2015)

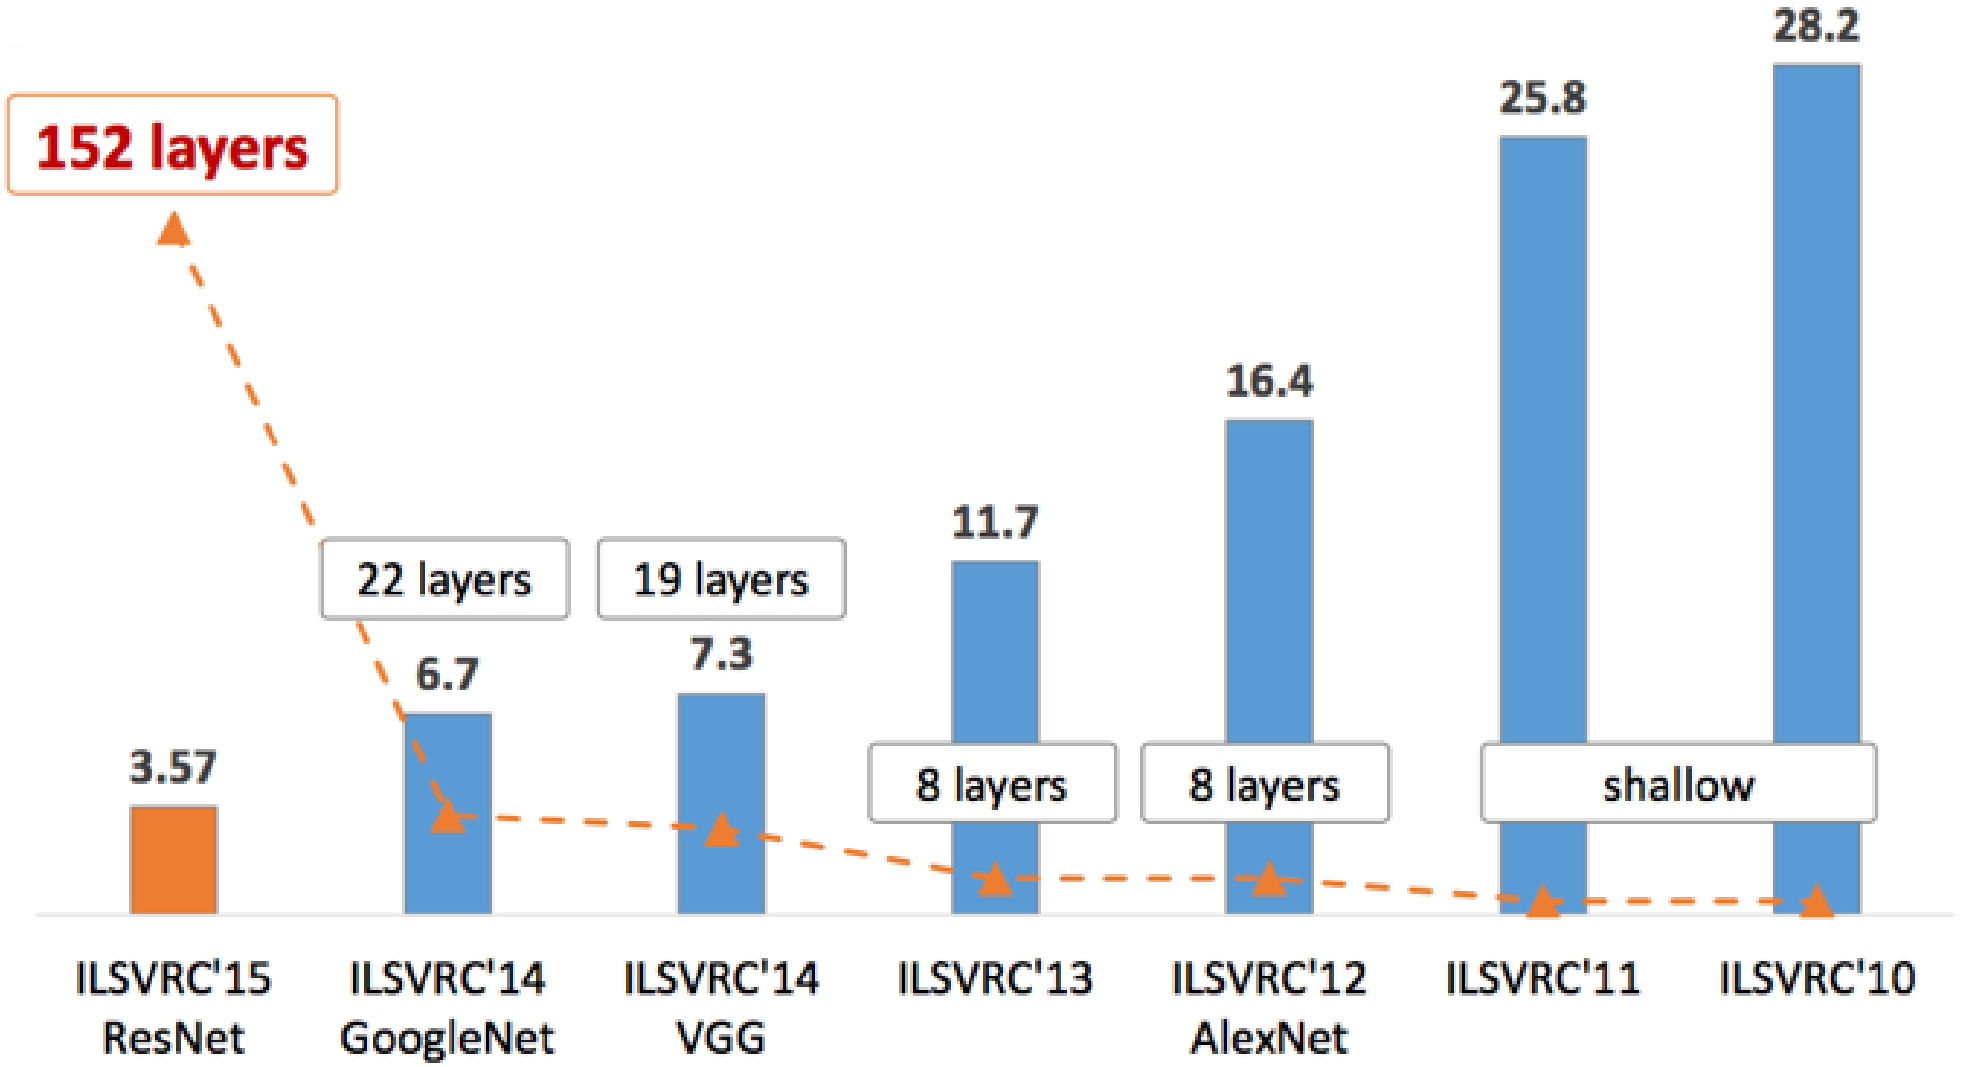

LeNet

AlexNet

VGG

GoogLeNet

Inception V3

Inception BN


WE NEED TO GO
DEEPER
memegenerator.net

Residual Network (ResNet)

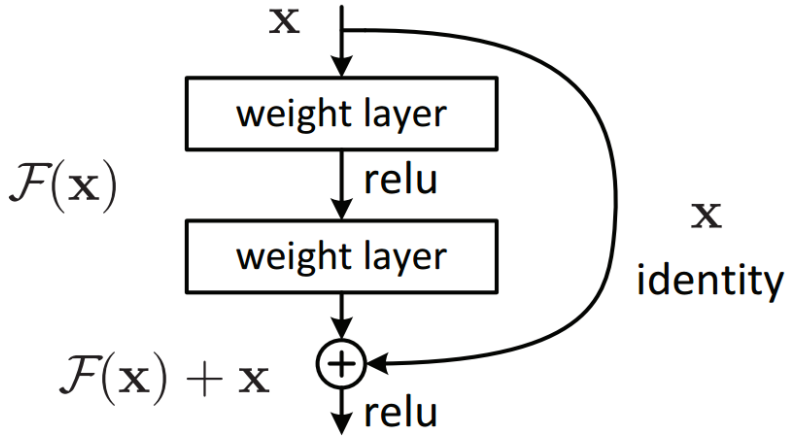

Figure 2. Residual learning: a building block.

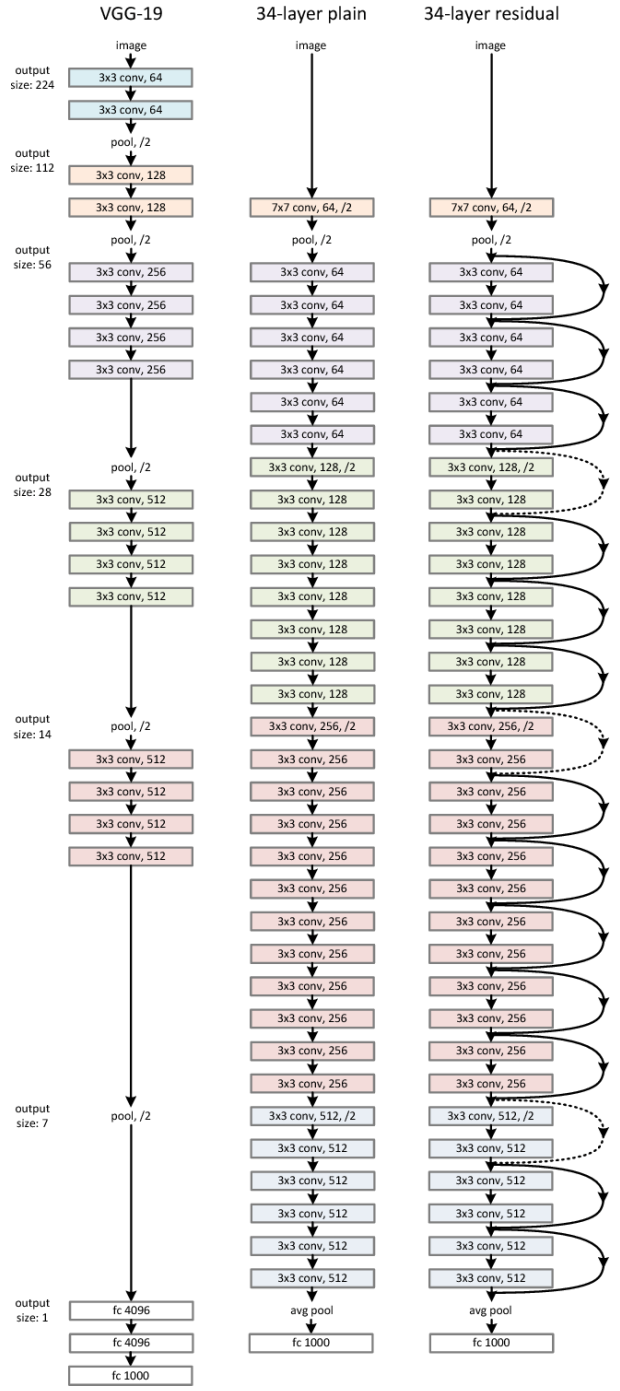

Comparison of Popular CNN Architectures

- ImageNet top 1 accuracy

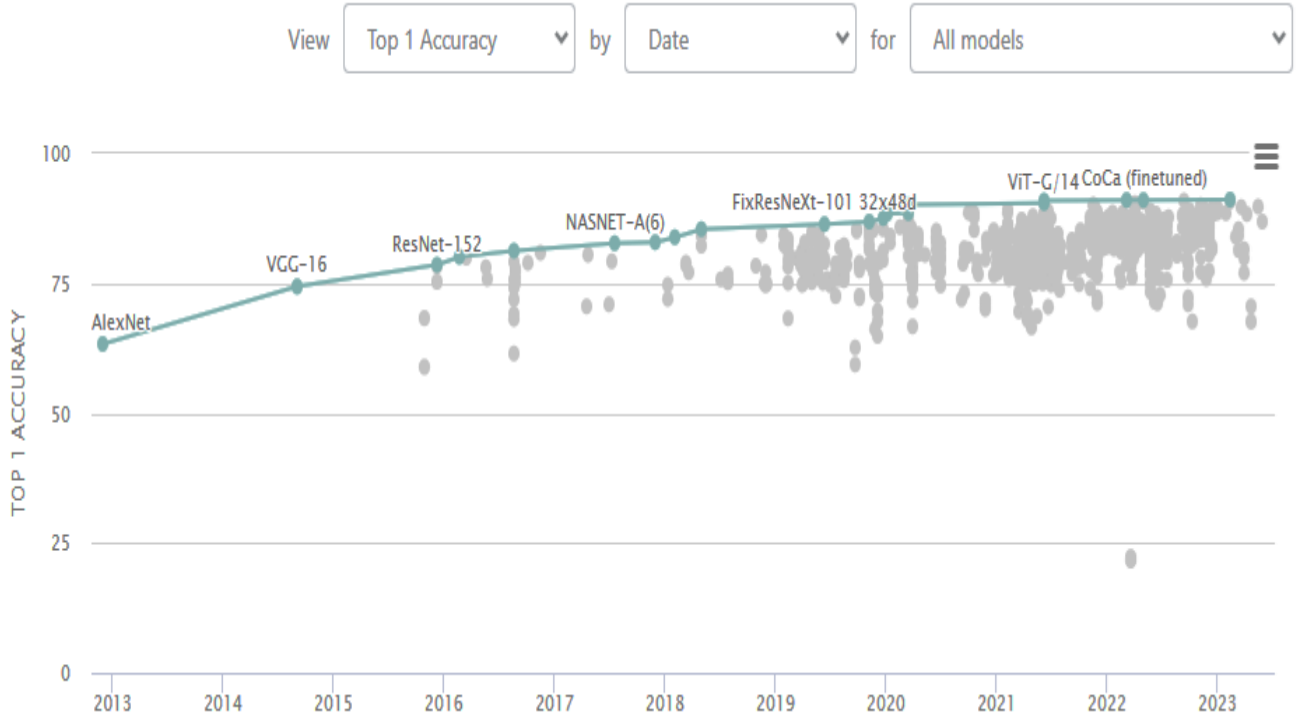

https://paperswithcode.com/sota/image-classification-on-imagenet

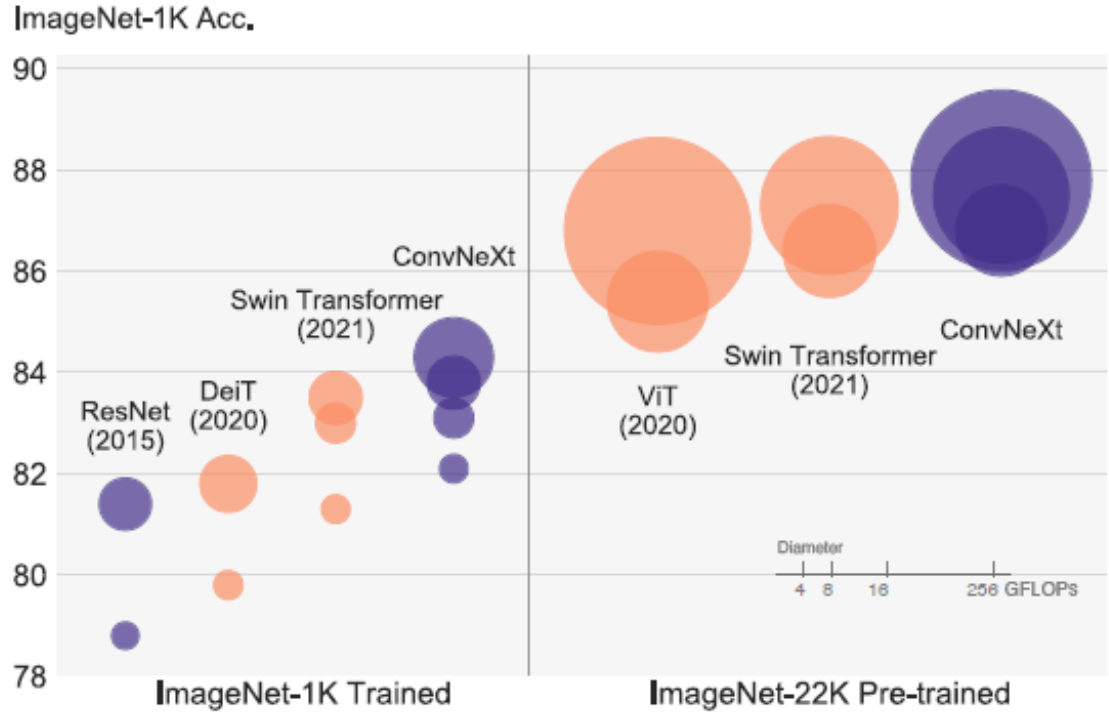

https://sh-tsang.medium.com/review-convnext-a-convnet-for-the-2020s-53b9ada30ab9

AUGMENTED STARTUPS
Computer Vision | AI | Drones

Real-time Object Detection

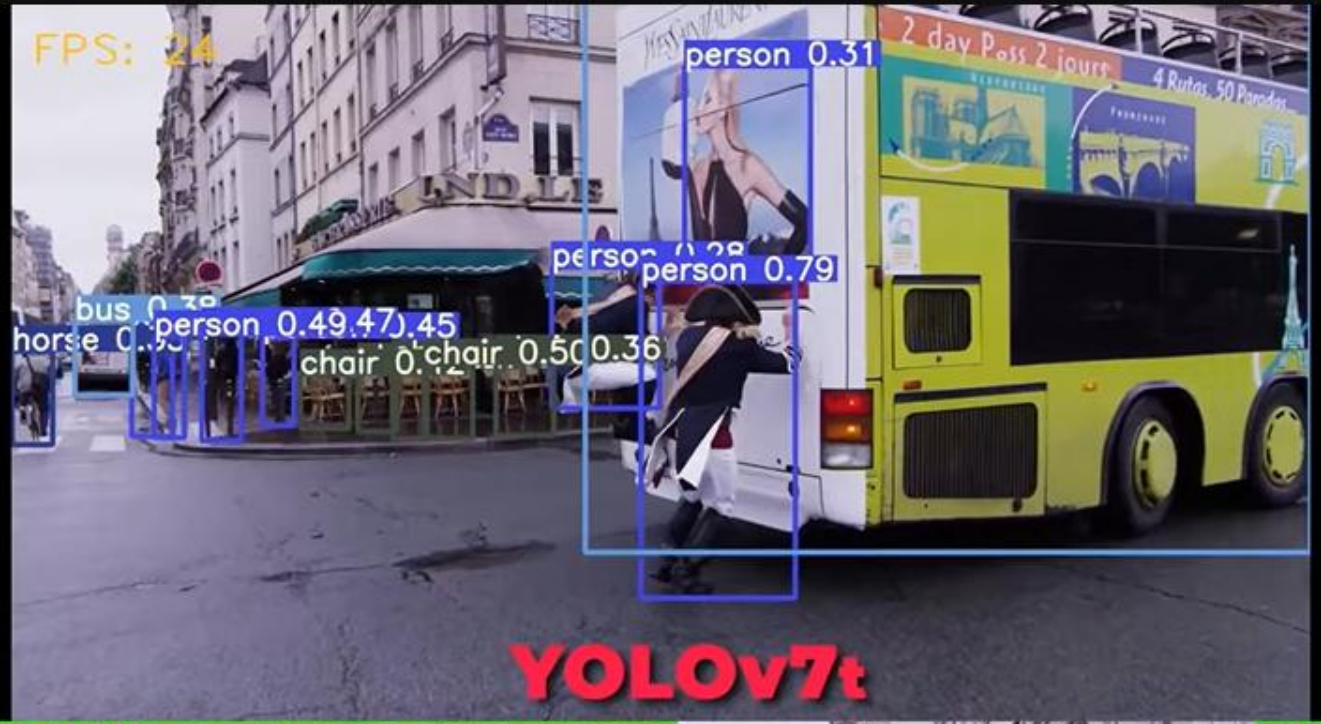

You Only Look Once (YOLO)

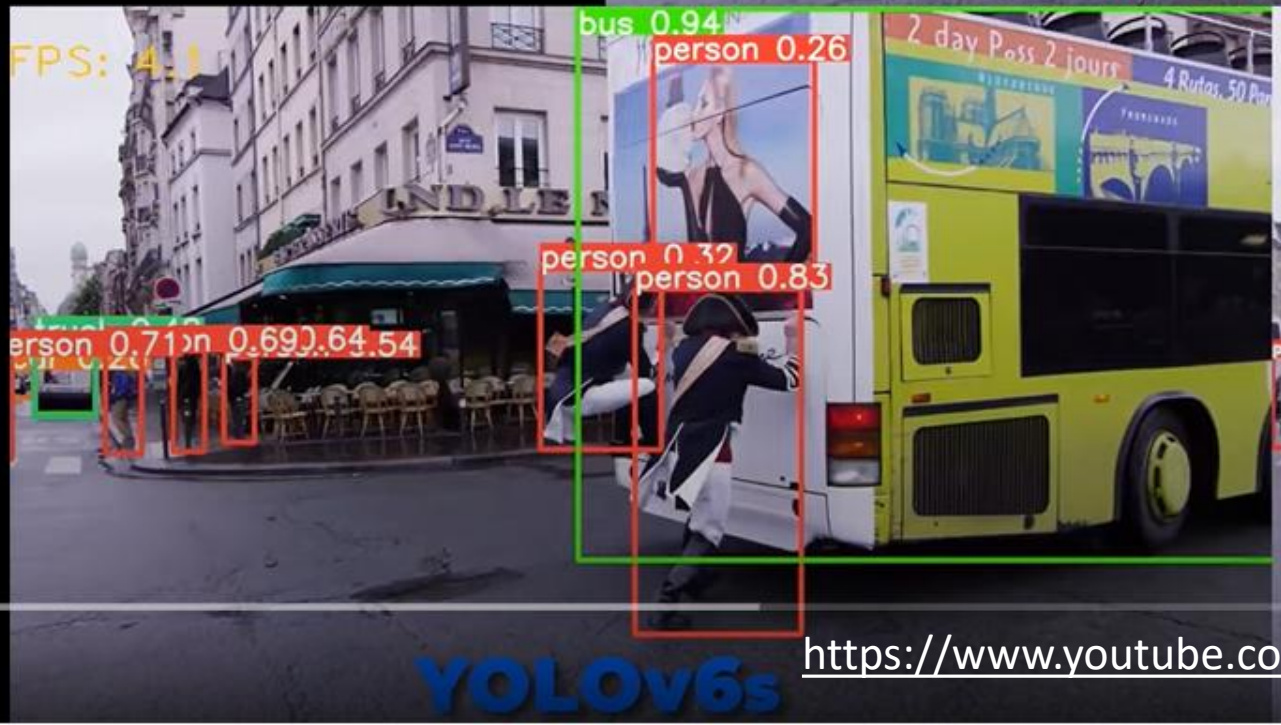

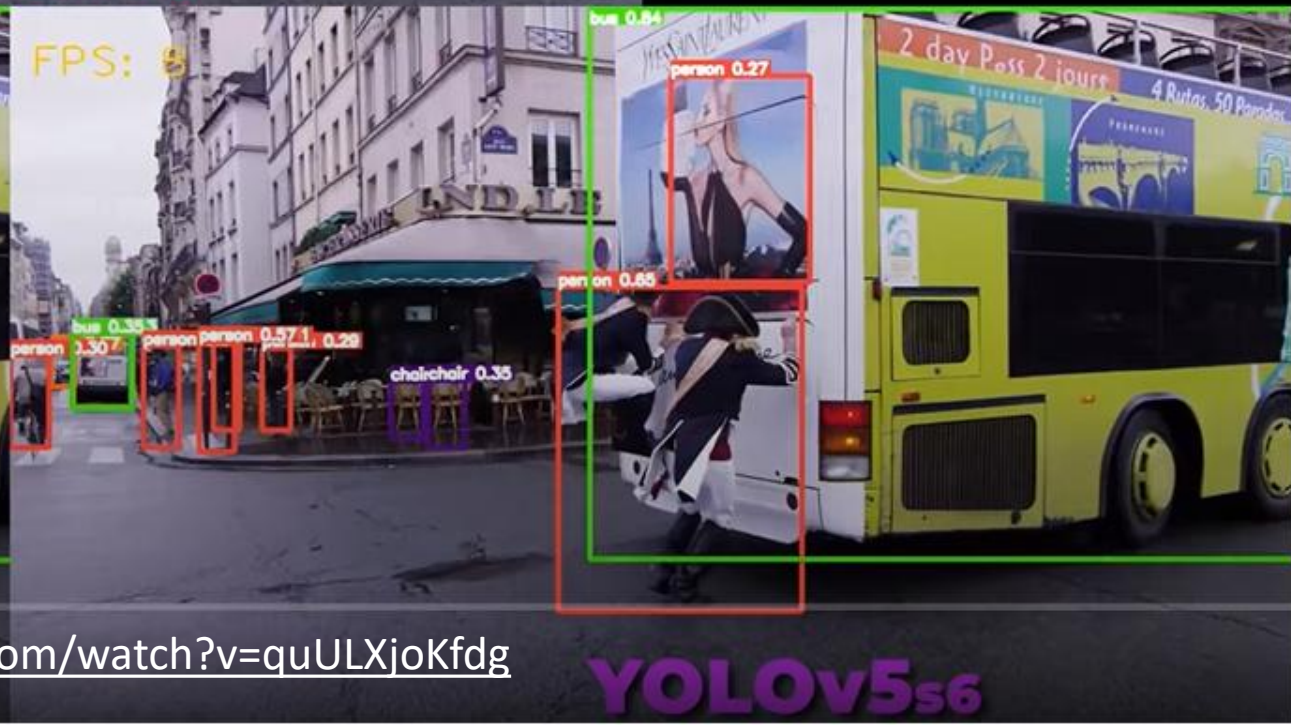

https://www.youtube.com/watch?v=quULXjoKfdg

YOLO v4, v7 Authors and Me

- From left to right
  - Me, Dr. Mark Liao, Dr. Wang

Recurrent Neural Networks (RNNs)

- An internal state (memory)
- Feedback loop
- Good for processing time-series data

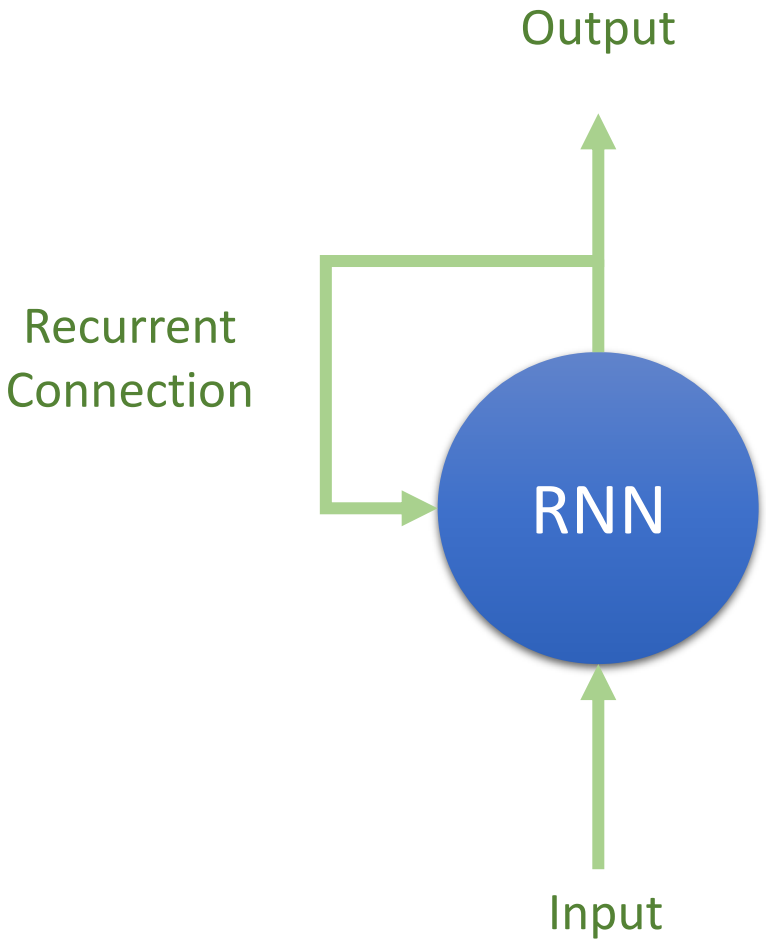

Unroll the RNN

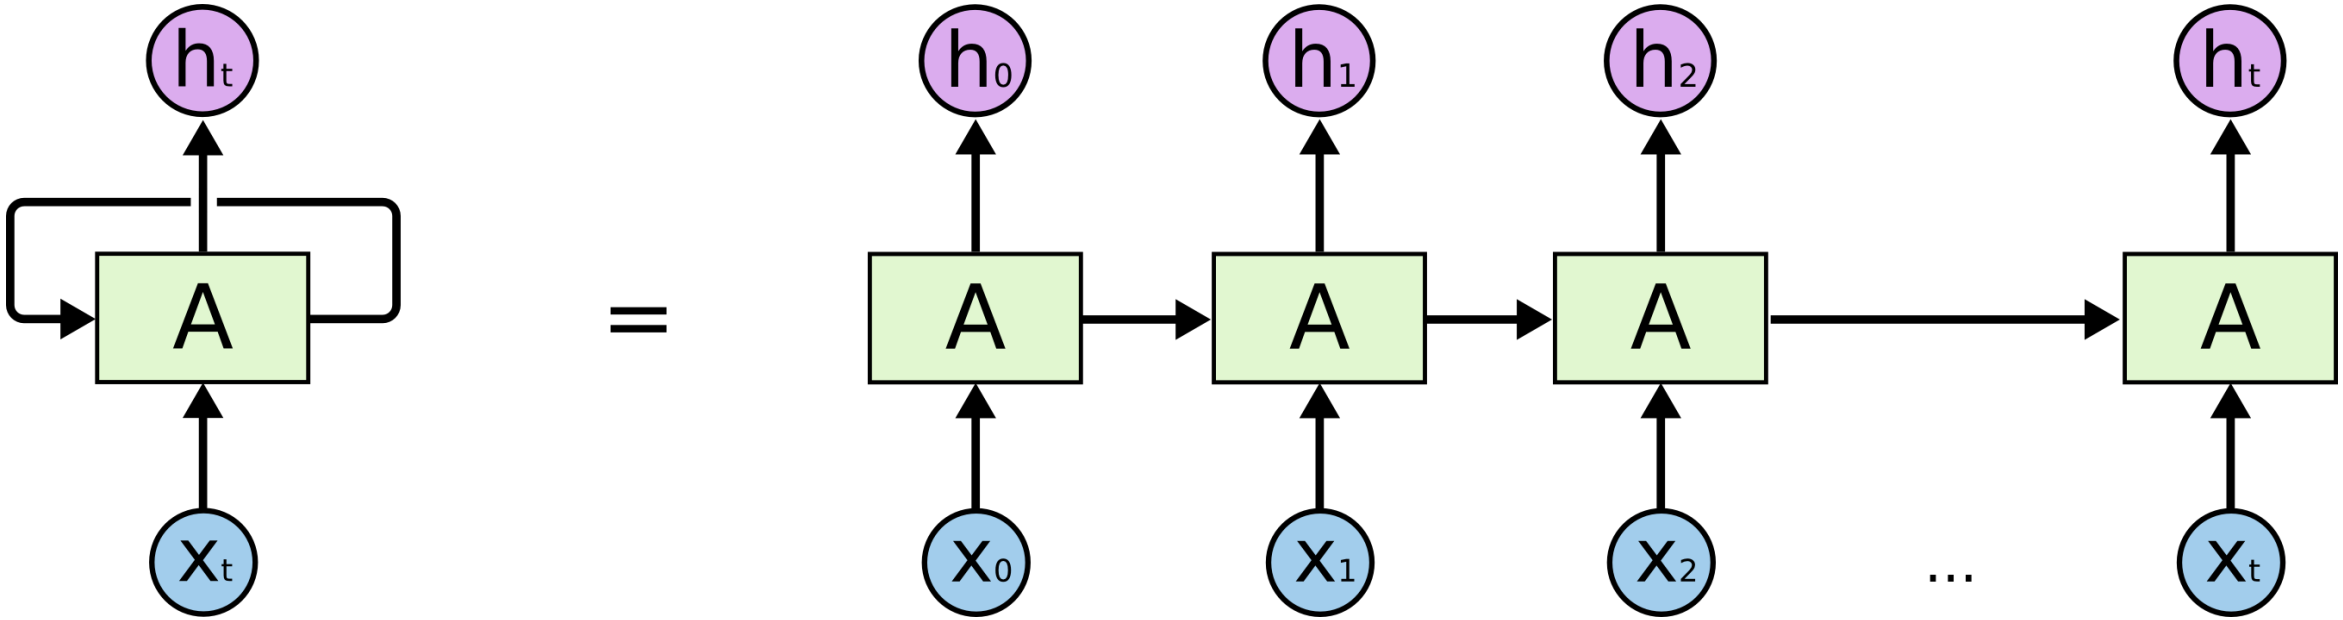

https://colah.github.io/posts/2015-08-Understanding-LSTMs/

Long Short-term Memory (LSTM)

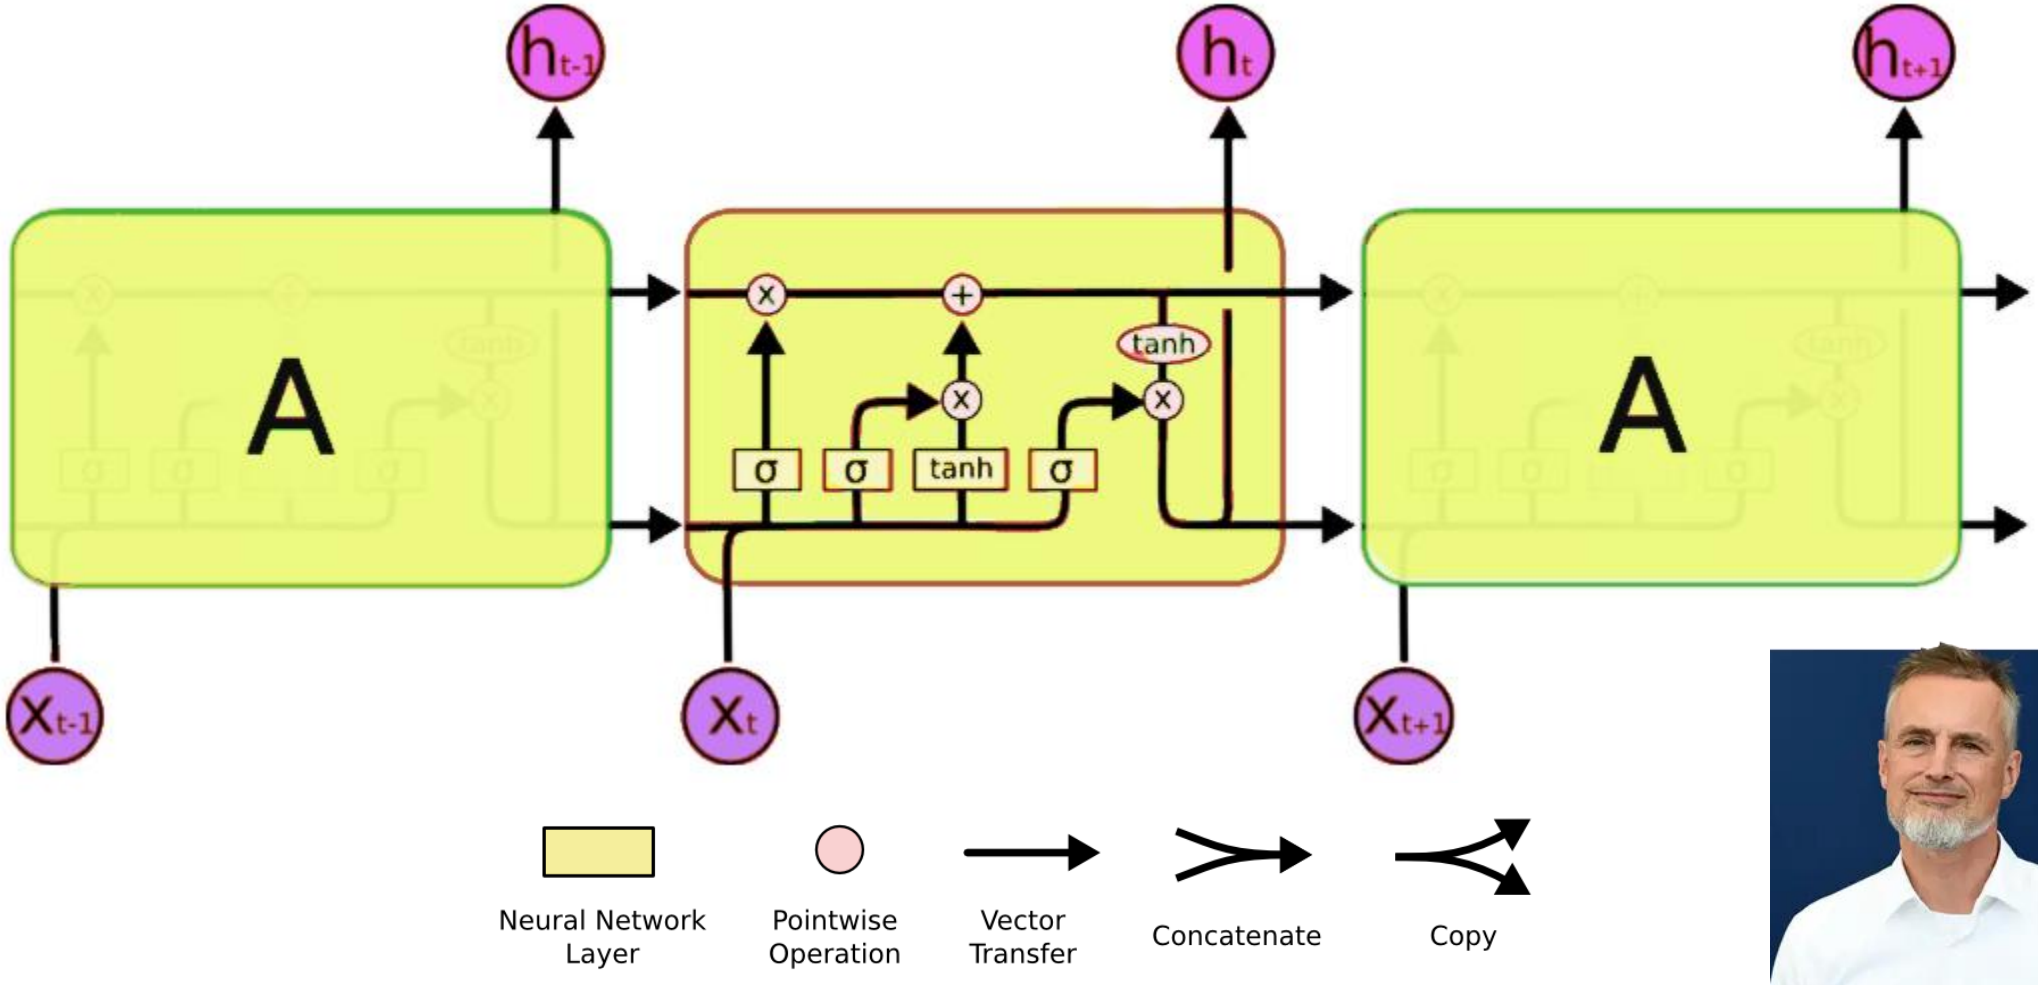

https://colah.github.io/posts/2015-08-Understanding-LSTMs/

Jürgen Schmidhuber

Sequence-2-Sequence model (Language Translation)

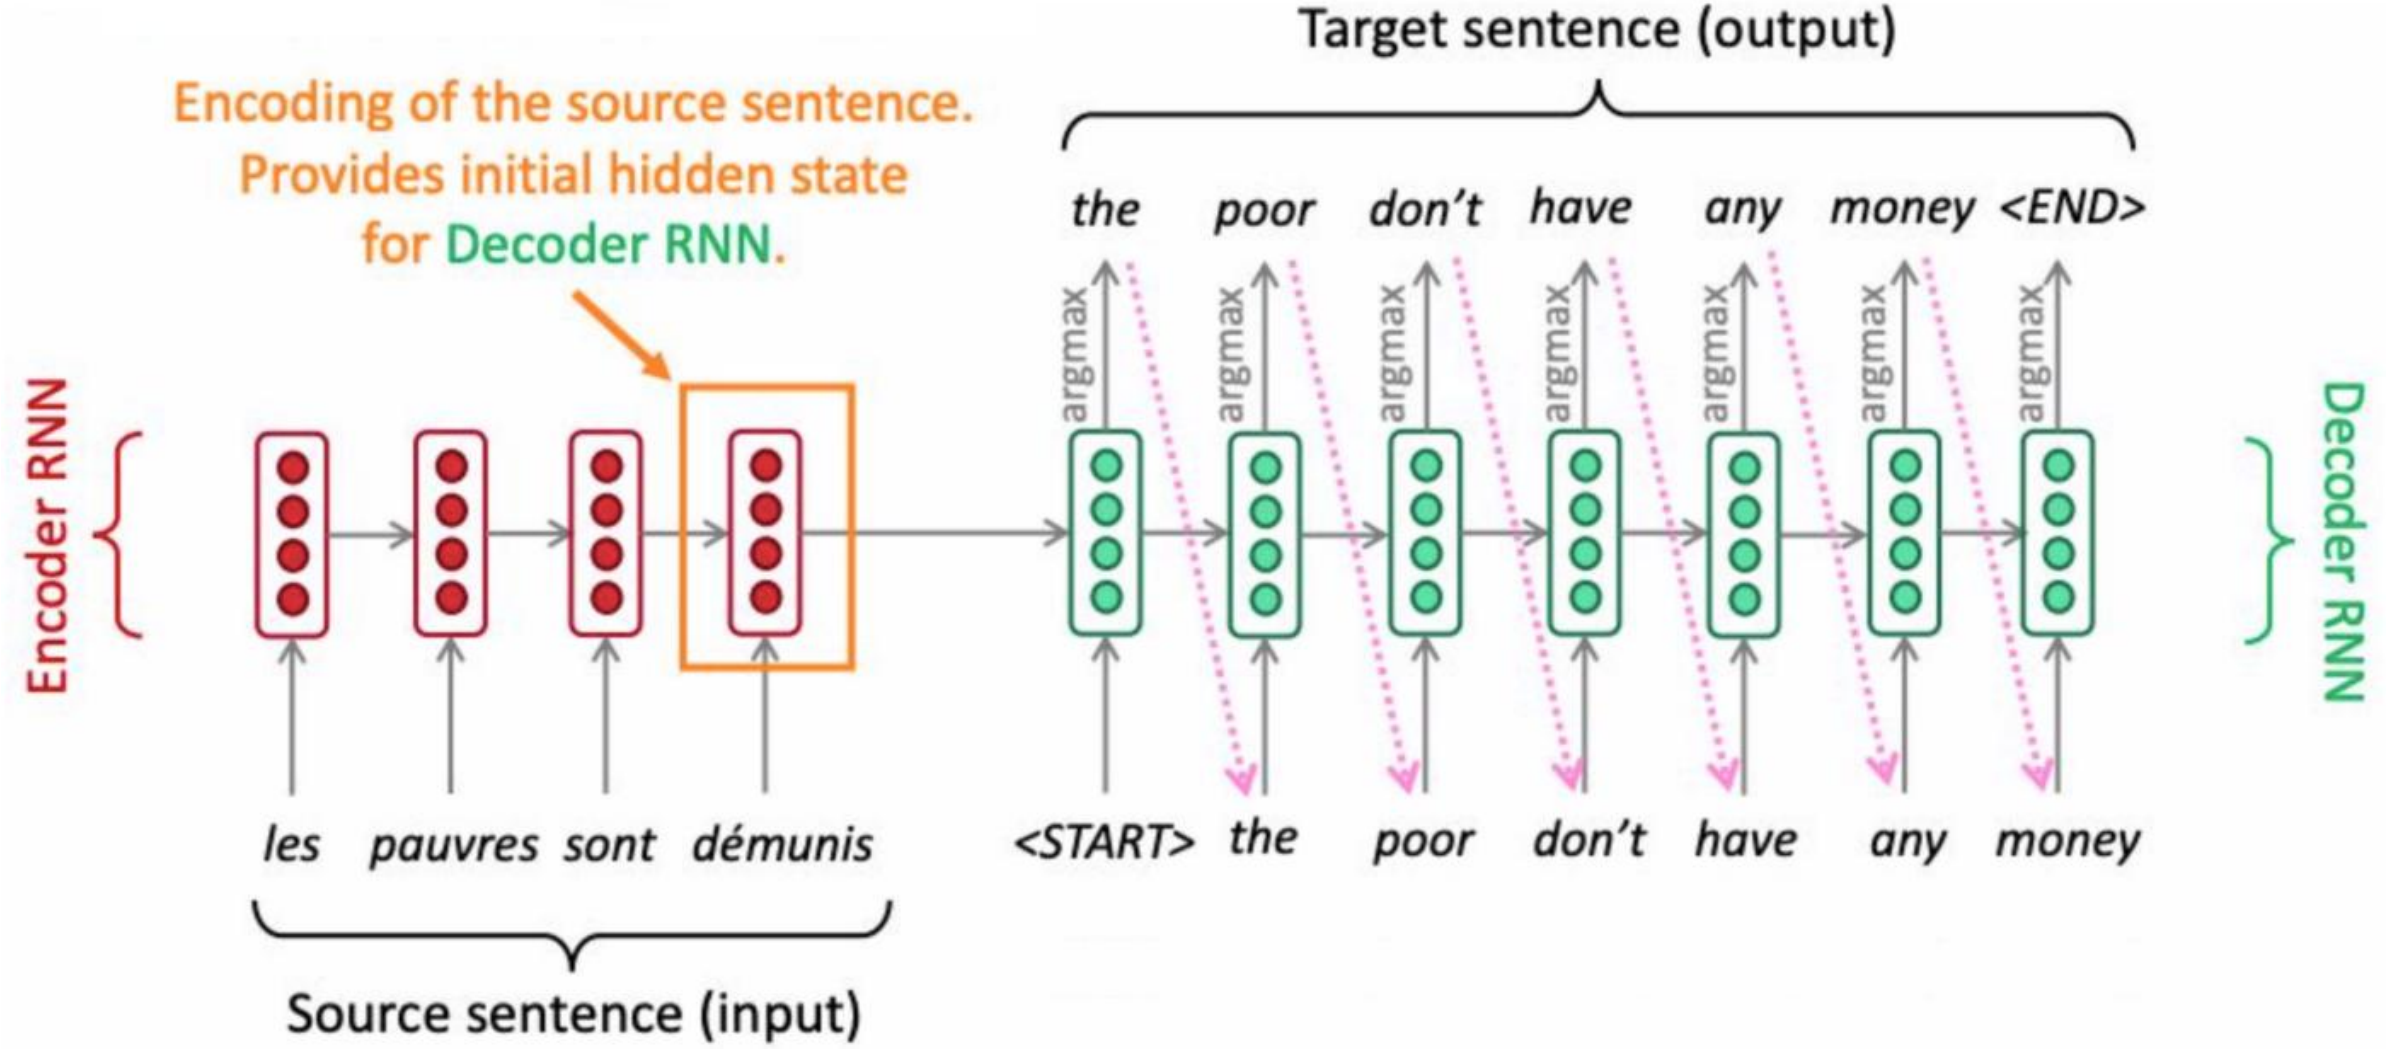

https://towardsdatascience.com/introduction-to-rnns-sequence-to-sequence-language-translation-and-attention-fc43ef2cc3fd

Attention is All You Need!

Ashish Vaswani* Google Brain avaswani@google.com

Noam Shazeer* Google Brain noam@google.com

Niki Parmar* Google Research nikip@google.com

Jakob Uszkoreit* Google Research usz@google.com

Llion Jones* Google Research llion@google.com

Aidan N. Gomez* † University of Toronto aidan@cs.toronto.edu

Łukasz Kaiser* Google Brain lukaszkaiser@google.com

Illia Polosukhin* ‡ illia.polosukhin@gmail.com

Google Brain & University of Toronto, *NIPS*, 2017

Attention Module in Transformer

- Query (Q), Key (K), Value (V) attention

$$\text{Attention}(Q, K, V) = \text{softmax}(\frac{QK^T}{\sqrt{d_k}})V$$

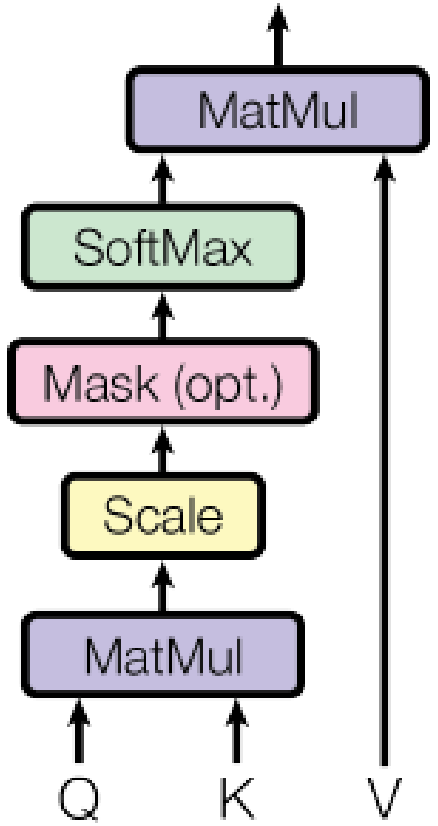

A. Waswani et al., "Attention is All You Need," NIPS, 2017

Query, Keys, Values

Query

"dogs on the beach"

Retrieving images from a database

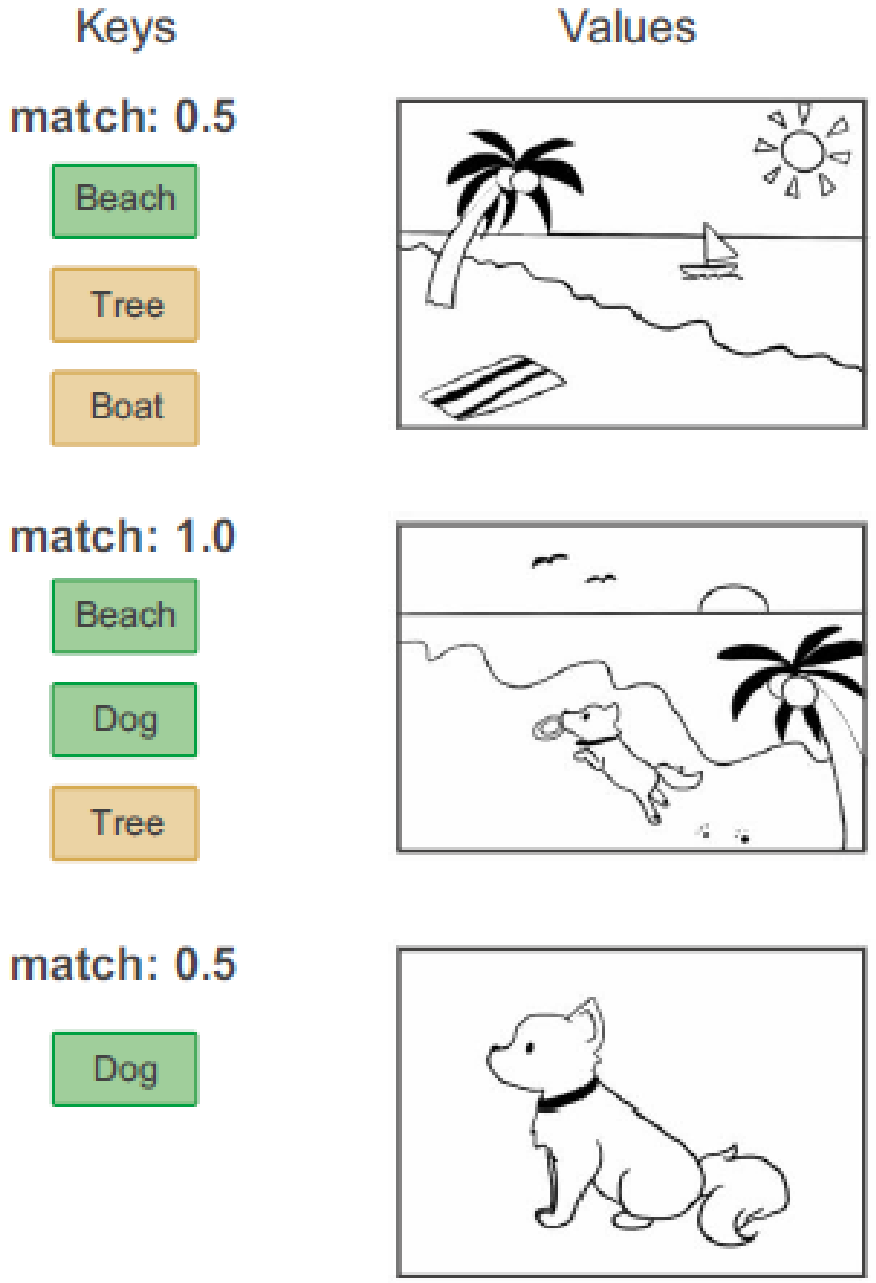

The Transformer Model

- Encoder-decoder architecture
- Multi-head attention
  - Self-attention in encoders
  - Masked Self-attention in decoders
  - Encoder-decoder attention
- Positional encoding

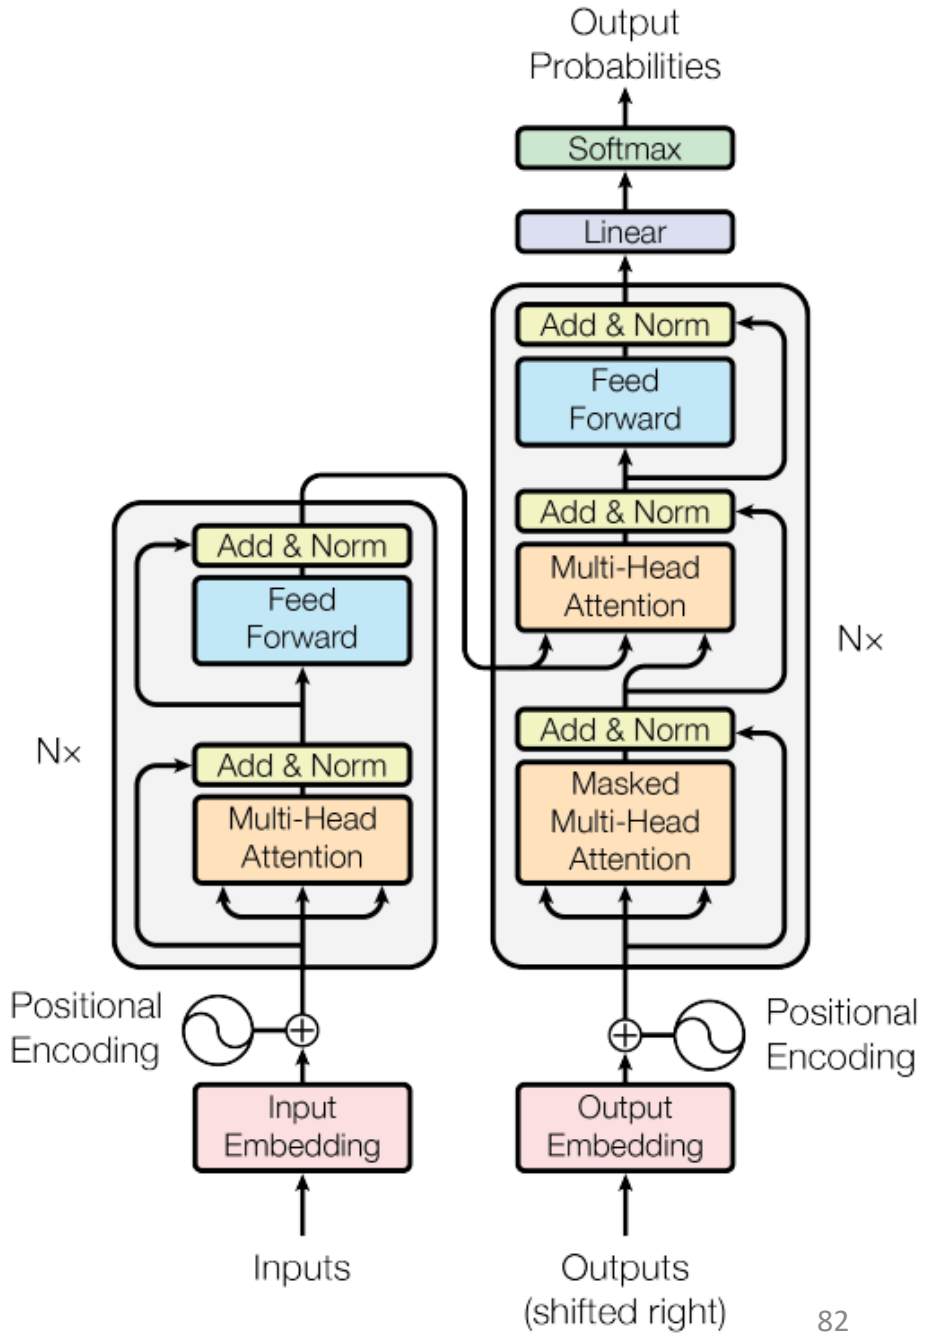

A. Waswani et al., "Attention is All You Need," NIPS, 2017

Visualizing Attention

- Tensor2Tensor Notebook

https://colab.research.google.com/github/tensorflow/tensor2tensor/blob/master/tensor2tensor/notebooks/hello_t2t.ipynb

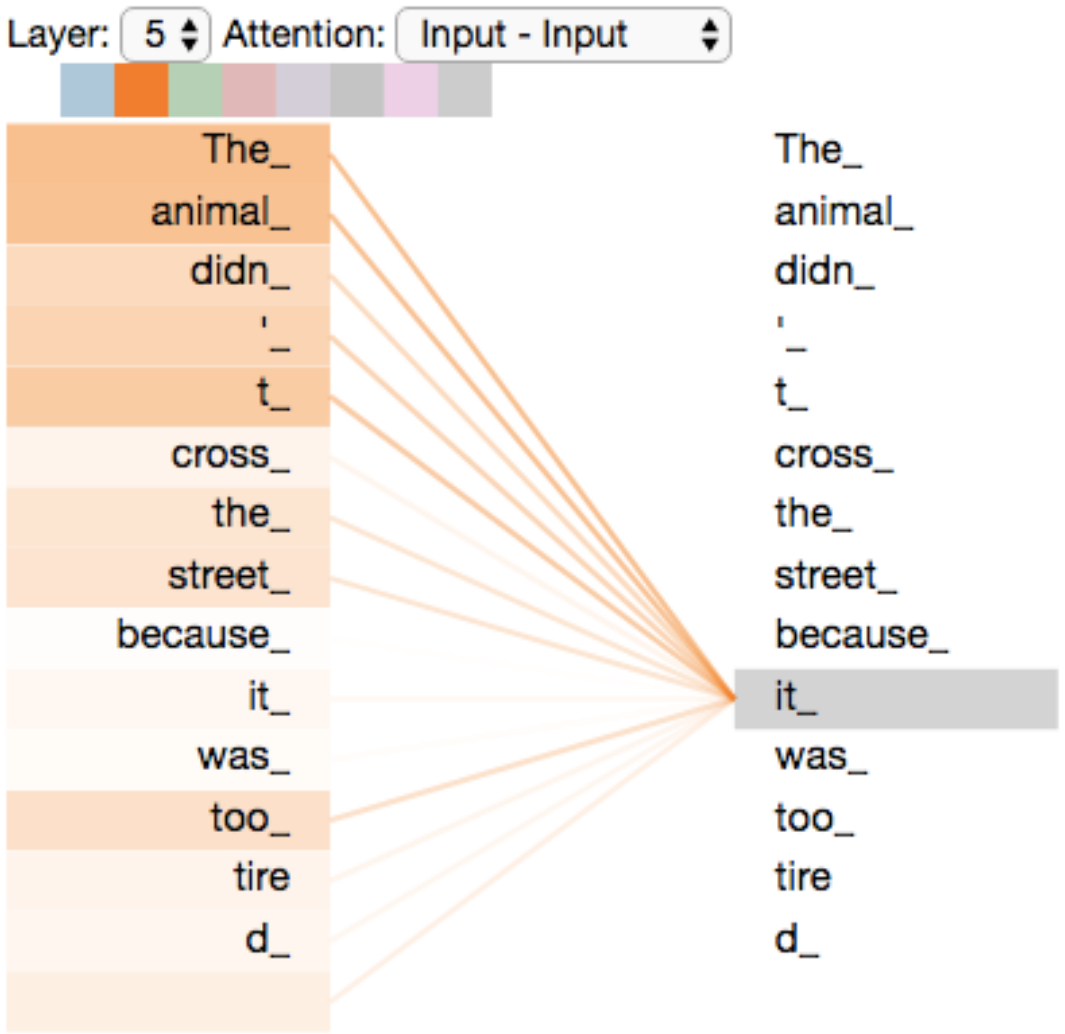


```
Inputs: The animal didn't cross the
street because it was too tired
```

```
Outputs: Das Tier überquerte die Straße
nicht, weil es zu müde war, weil es zu
müde war.
```

Deep Reinforcement Learning

Reinforcement Learning

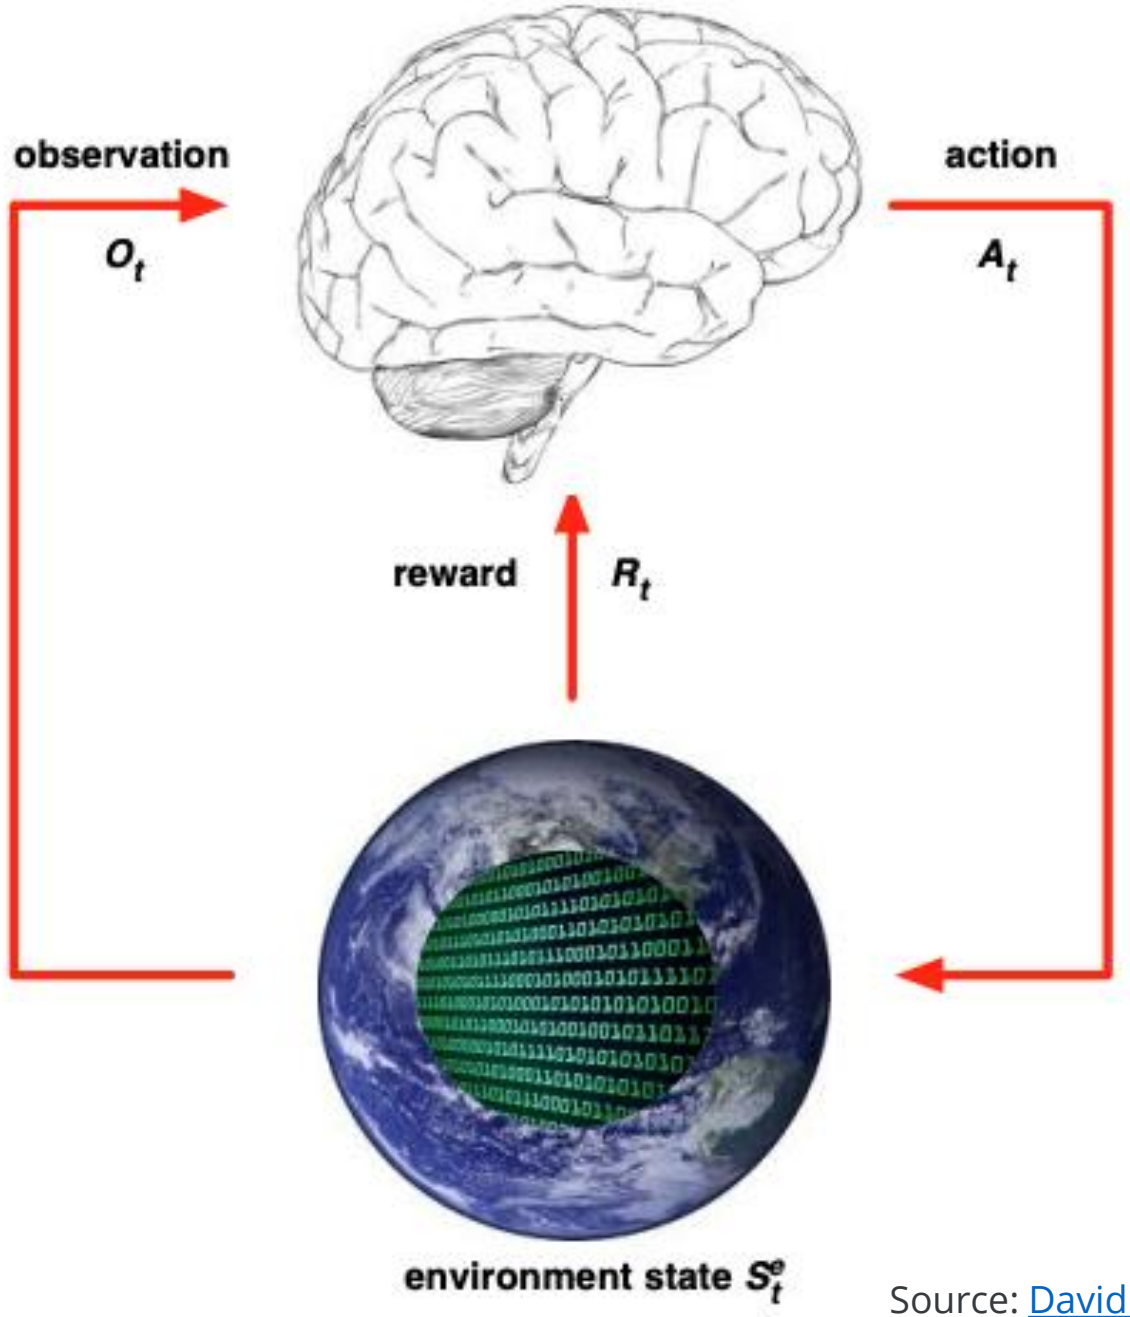

https://talkmarkets.com/content/deep-reinforcement-learning-for-trading-applications?post=252842

Source: David Silver, UCL

DeepMind: DRL in Atari

DeepMind

Demis Hassabis

Mustafa Suleyman

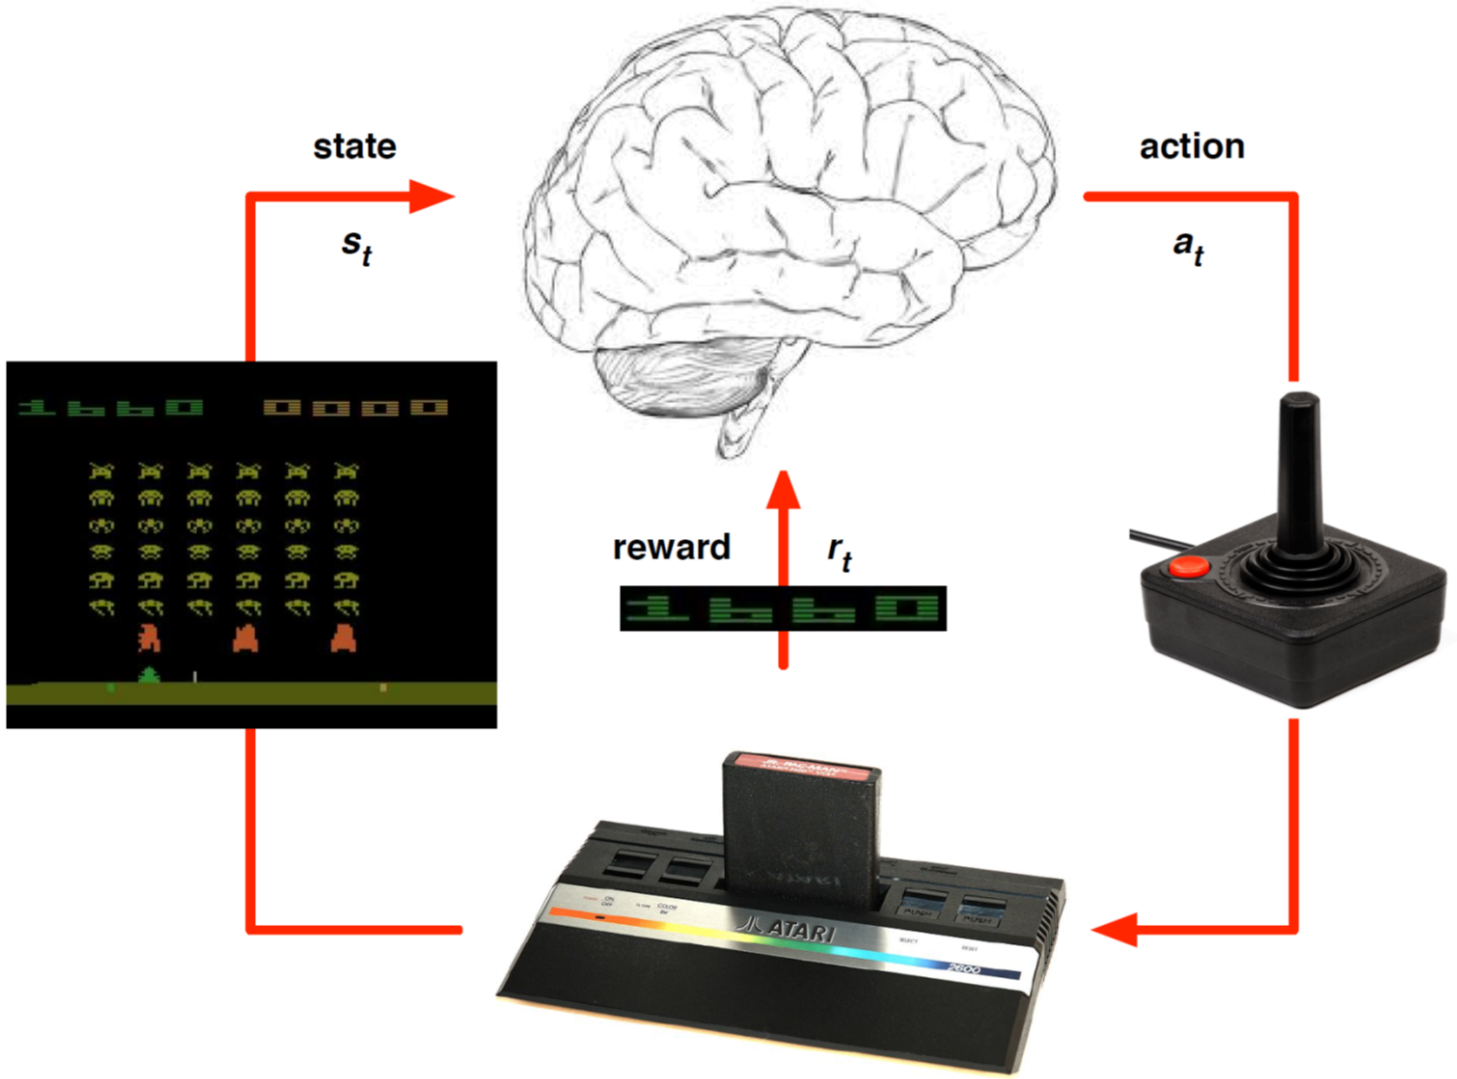

Mnih et al., "Human Level Control through Deep Reinforcement Learning," *Nature*, 2015

Learning to Play Atari Games

David Silver

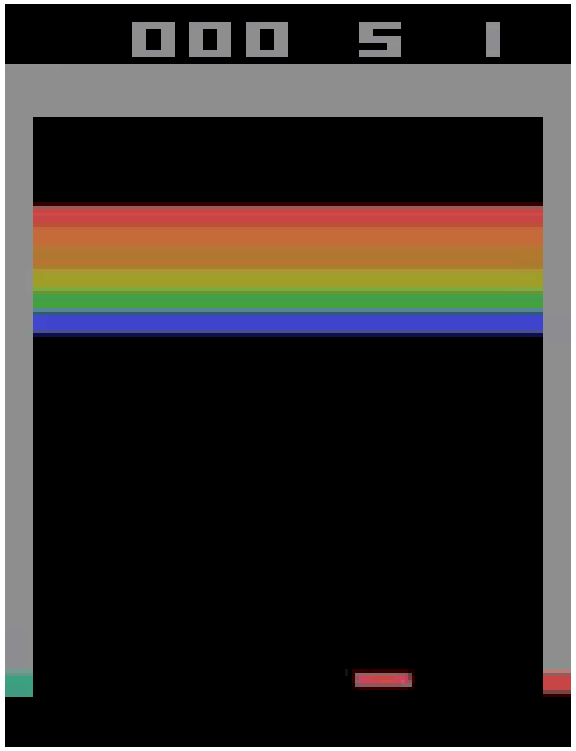

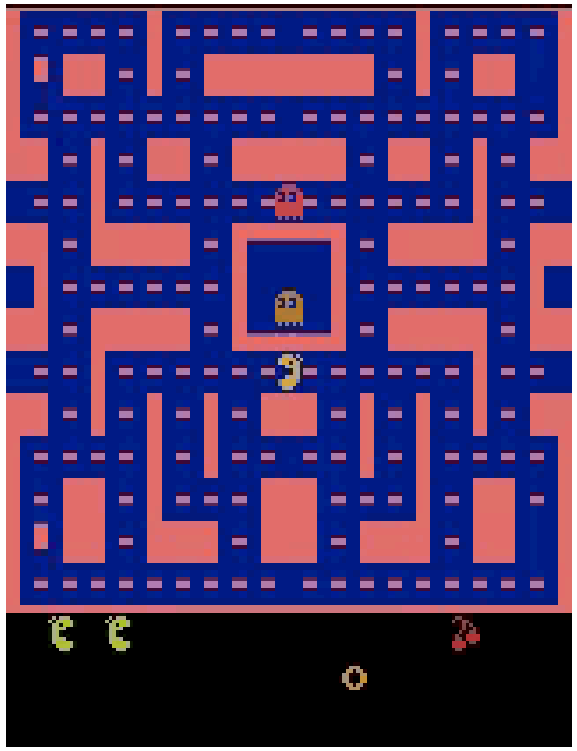

Complexity of Go vs. Chess

Game	Board size	State space	Game tree size
Go	19 x 19	10^{172}	10^{360}
Chess	8 x 8	10^{50}	10^{123}
Checkers	8 x 8	10^{18}	10^{54}

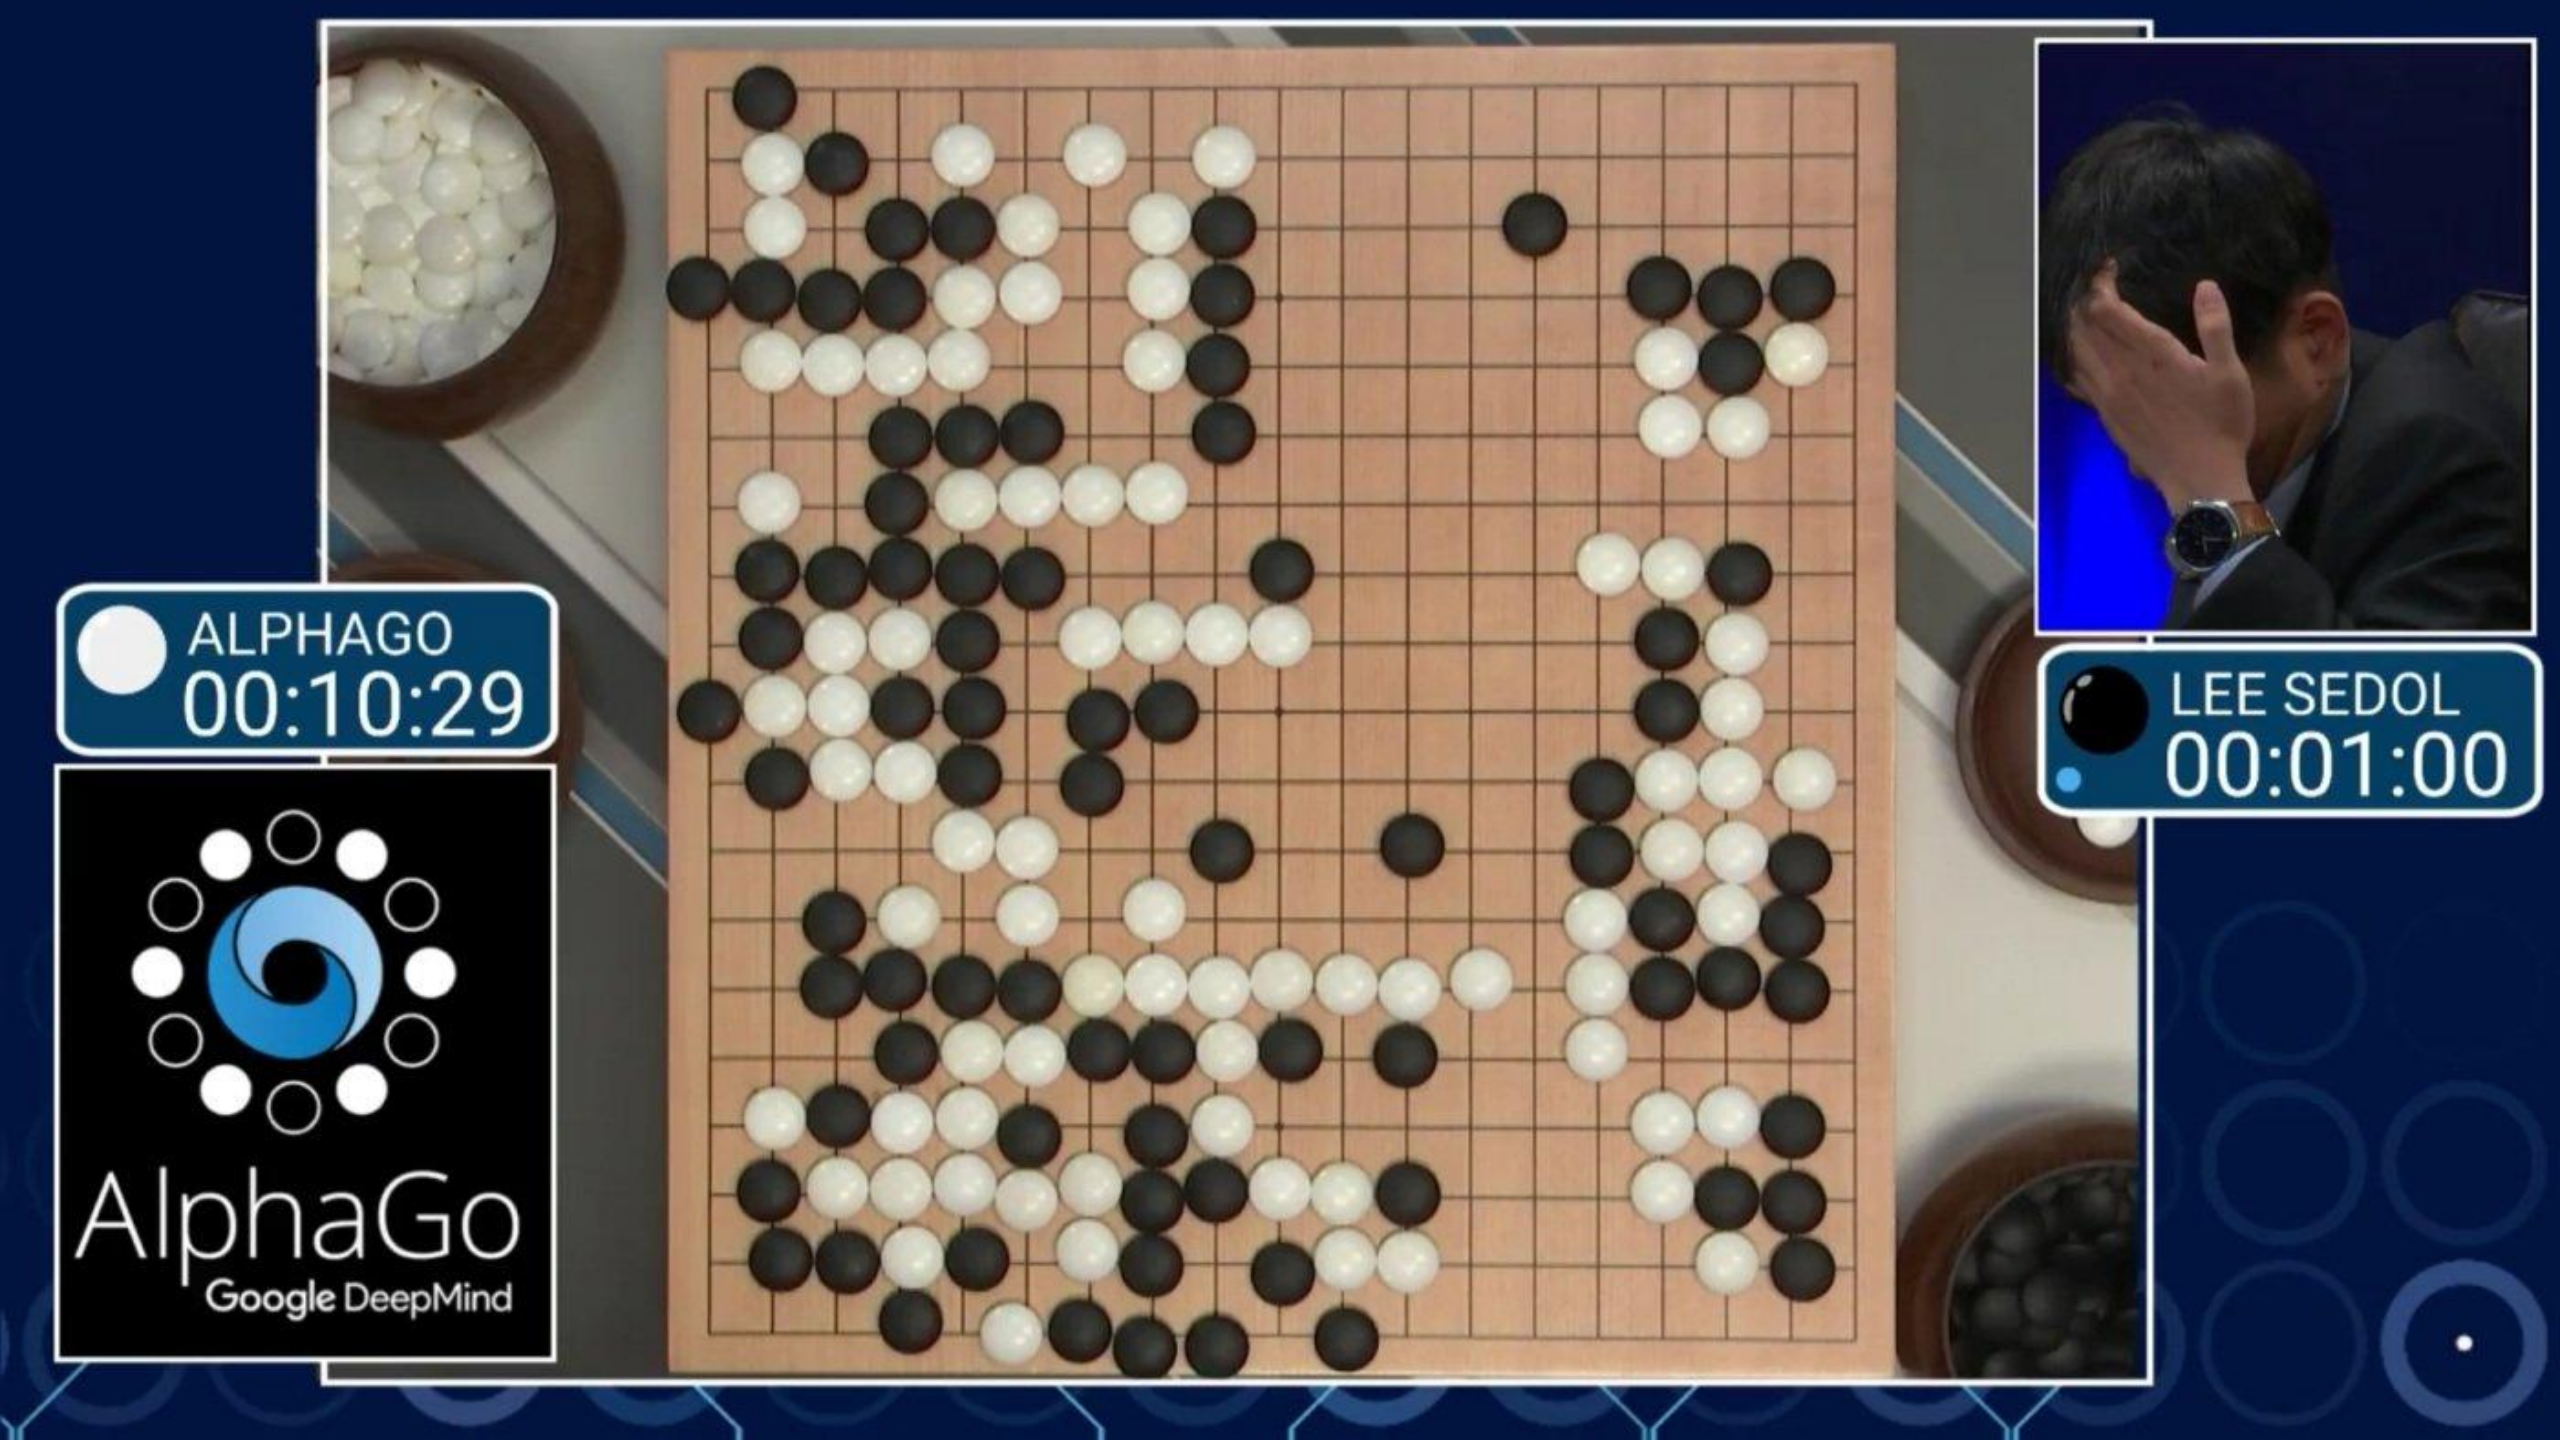
ALPHAGO
00:10:29
AlphaGo
Google DeepMind
LEE SEDOL
00:01:00

Dr. Aja Huang (黃士傑)


AlphaGo Zero
Starting from scratch

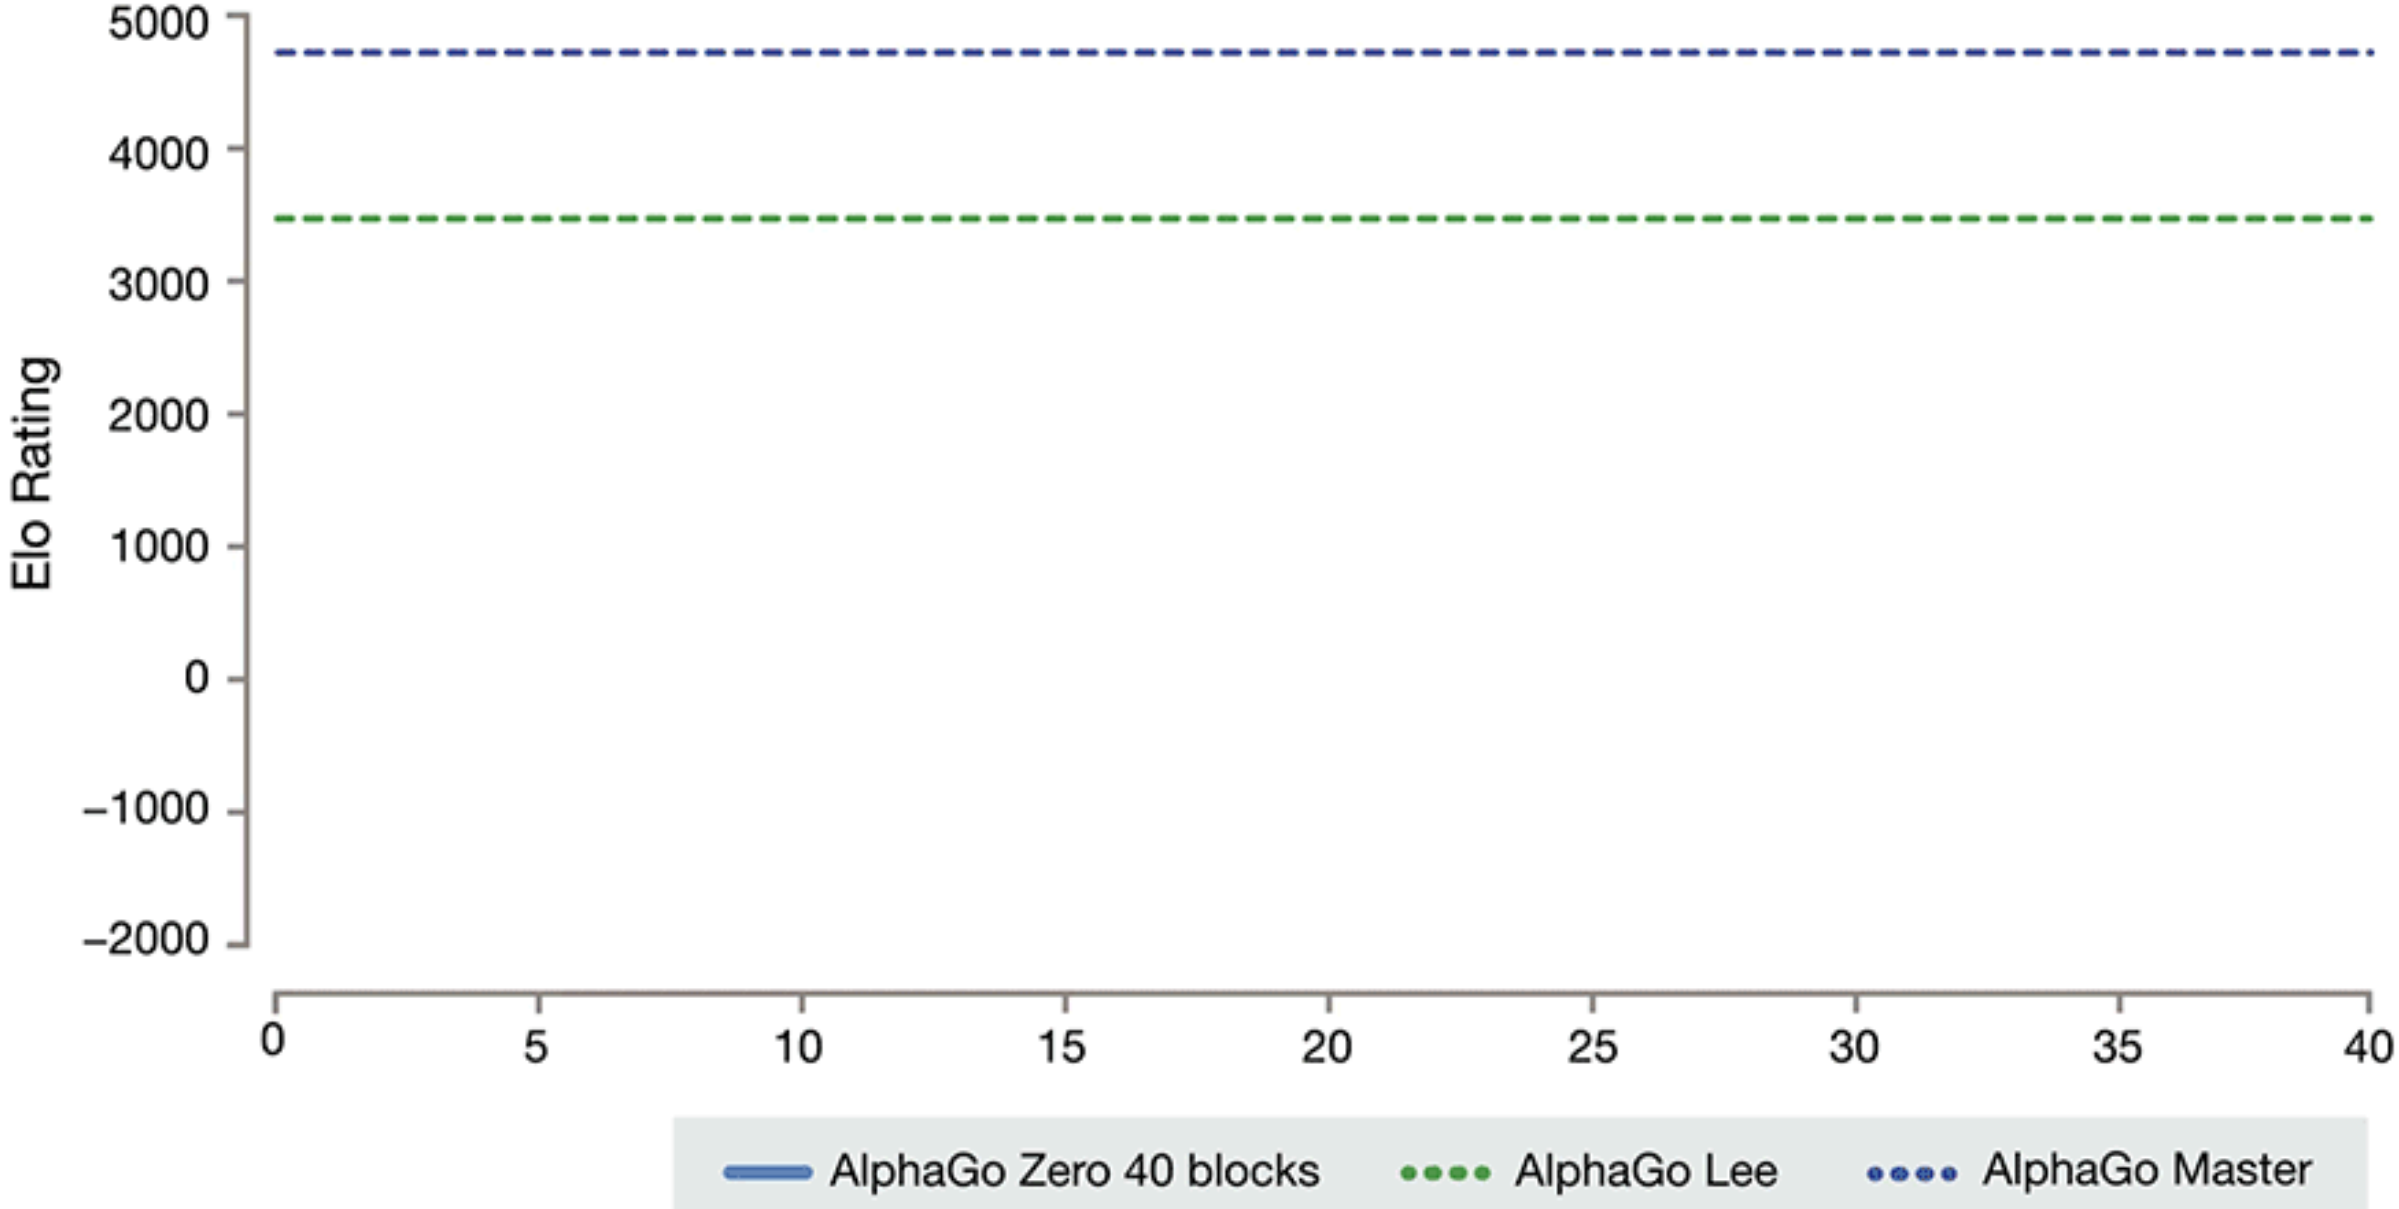
5000
4000
3000
2000
1000
0
−1000
−2000
Elo Rating
0
5
10
15
20
25
30
35
40
AlphaGo Zero 40 blocks
AlphaGo Lee
AlphaGo Master


5000
4000
3000
2000
1000
0
−1000
−2000
Elo Rating
0
5
10
15
20
25
30
35
40
40 days
AlphaGo Zero surpasses all other versions of AlphaGo and, arguably, becomes the best Go player in the world. It does this entirely from self-play, with no human intervention and using no historical data.
AlphaGo Zero 40 blocks
AlphaGo Lee
AlphaGo Master


Generative AI

Generative Adversarial Networks (GAN)

- Ian Goodfellow

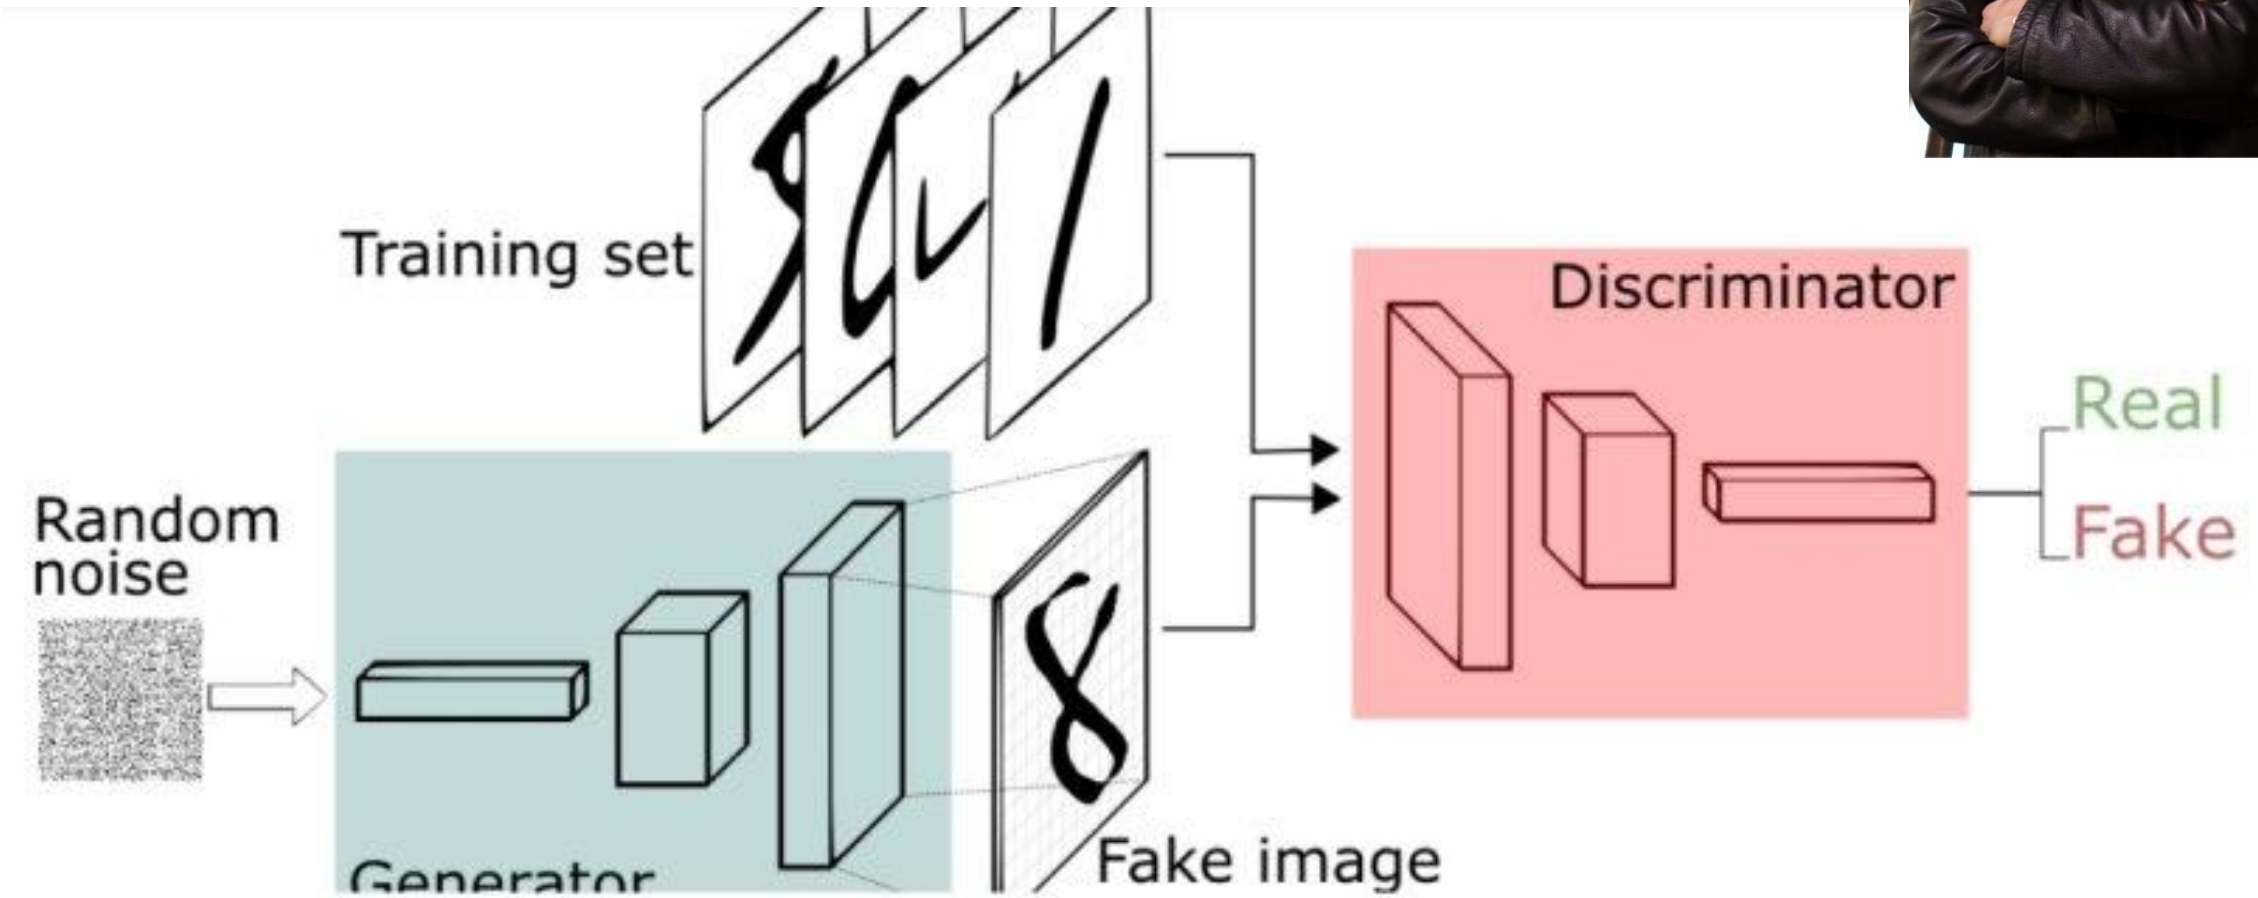

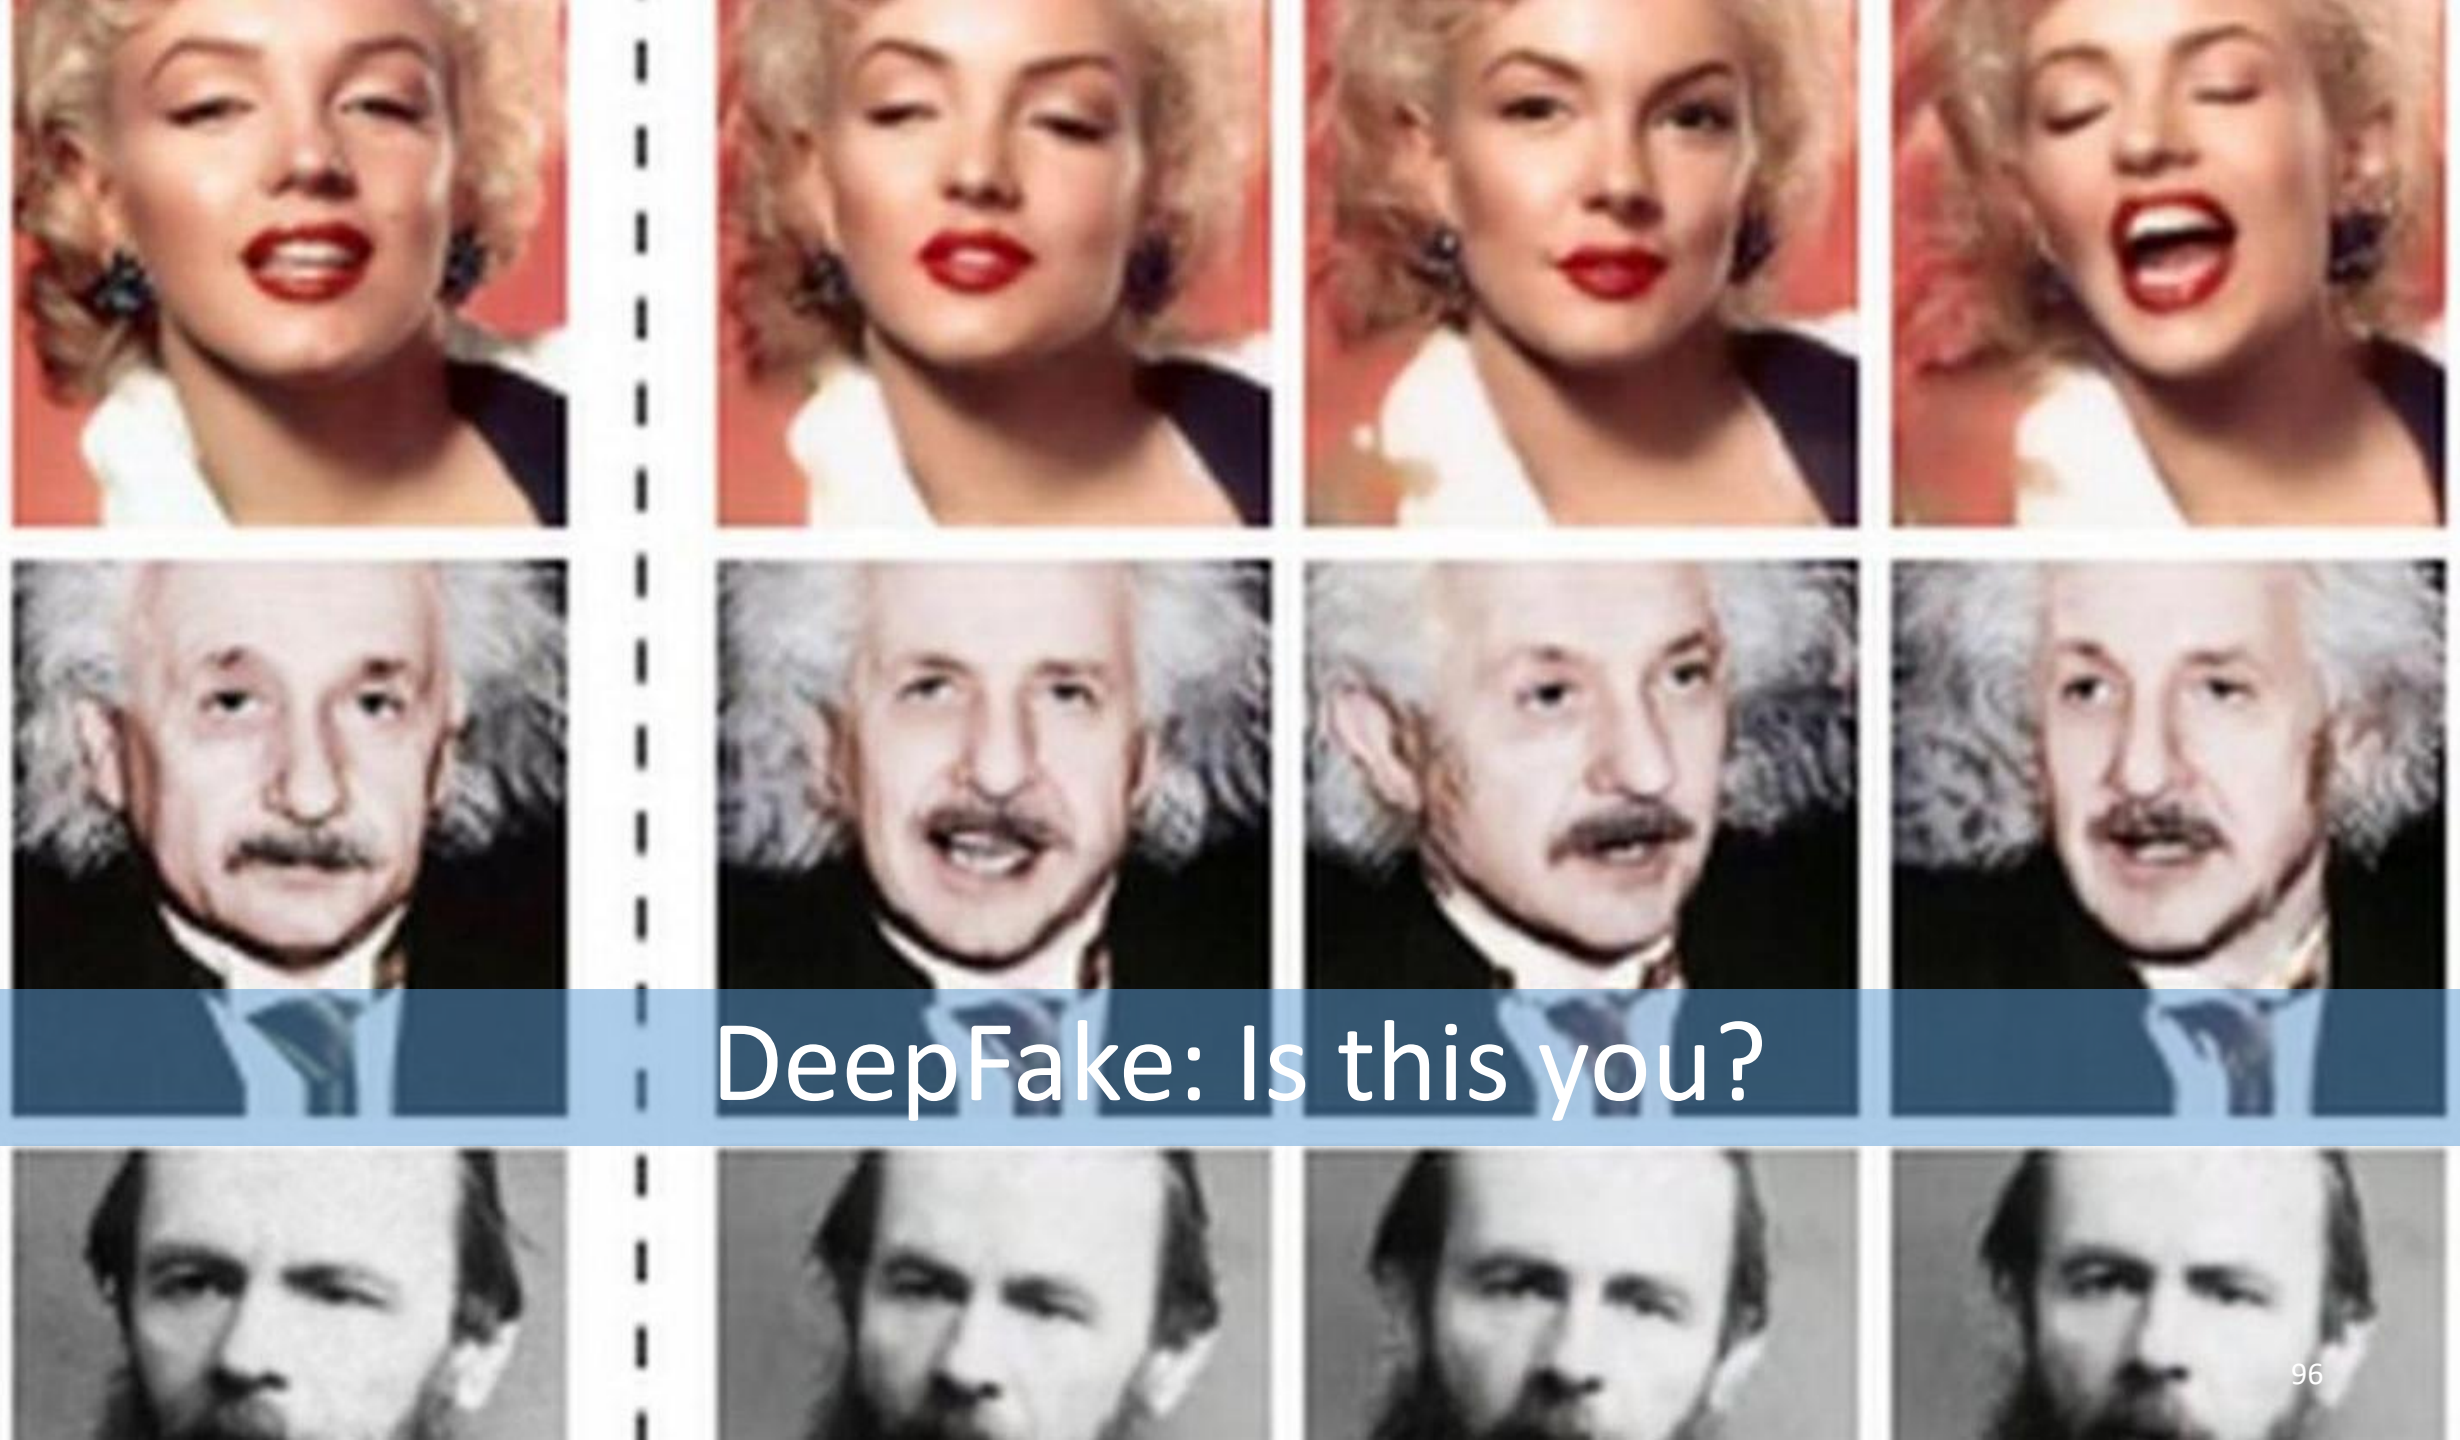

DeepFake: Is this you?

https://www.youtube.com/watch?v=gLoI9hAX9dw

Adversarial Attack

“panda”

57.7% confidence

$+\epsilon$

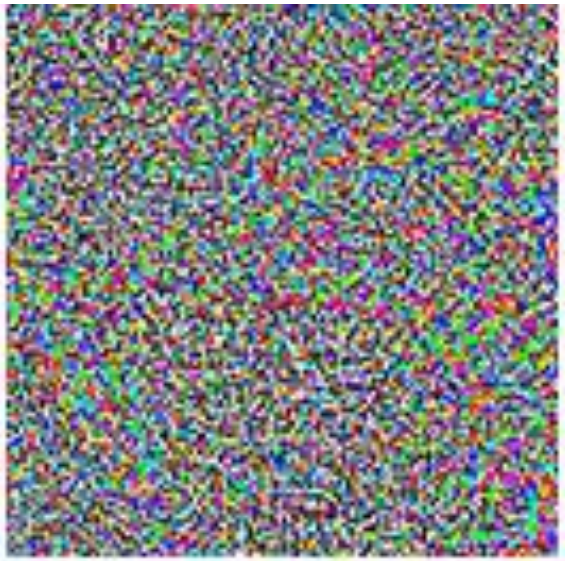

$=$

“gibbon”

99.3% confidence

OpenAI GPT: Pre-training Transformer Decoders

- Unsupervised pre-train transform decoders for predicting the next word (GPT: Generative Pre-Training)
- Use 12 Transformer decoders in GPT-1
  - GPT-1: Improving Language Understanding with Unsupervised Learning (2018)
  - GPT-2: Better Language Models and Their Implications (2019)
  - GPT-3: Language Models are Few-Shot Learners (2020)

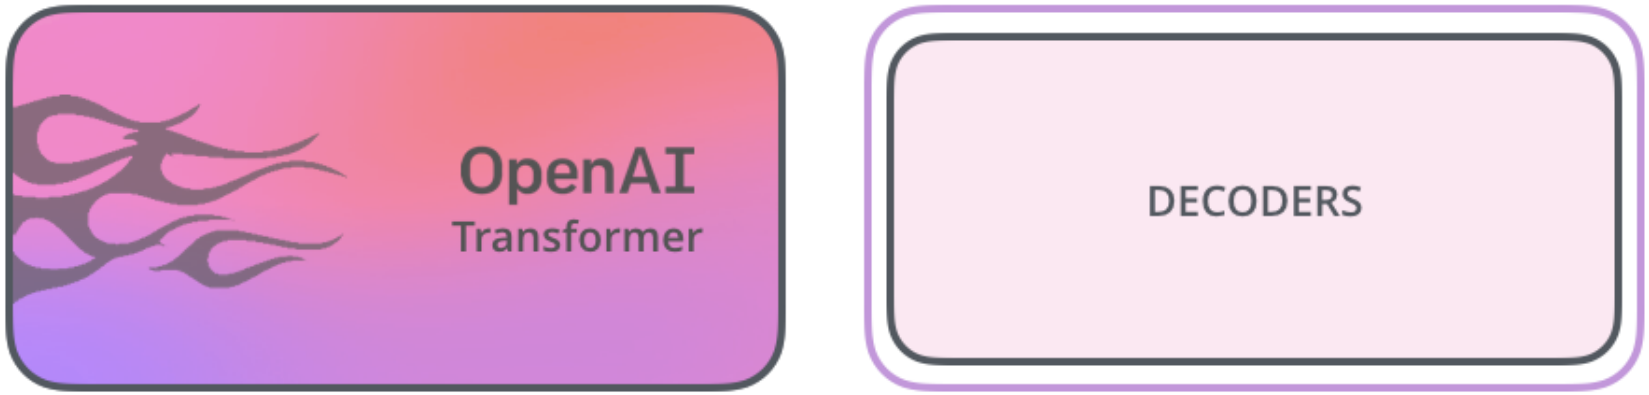

OpenAI GPT for Different Tasks

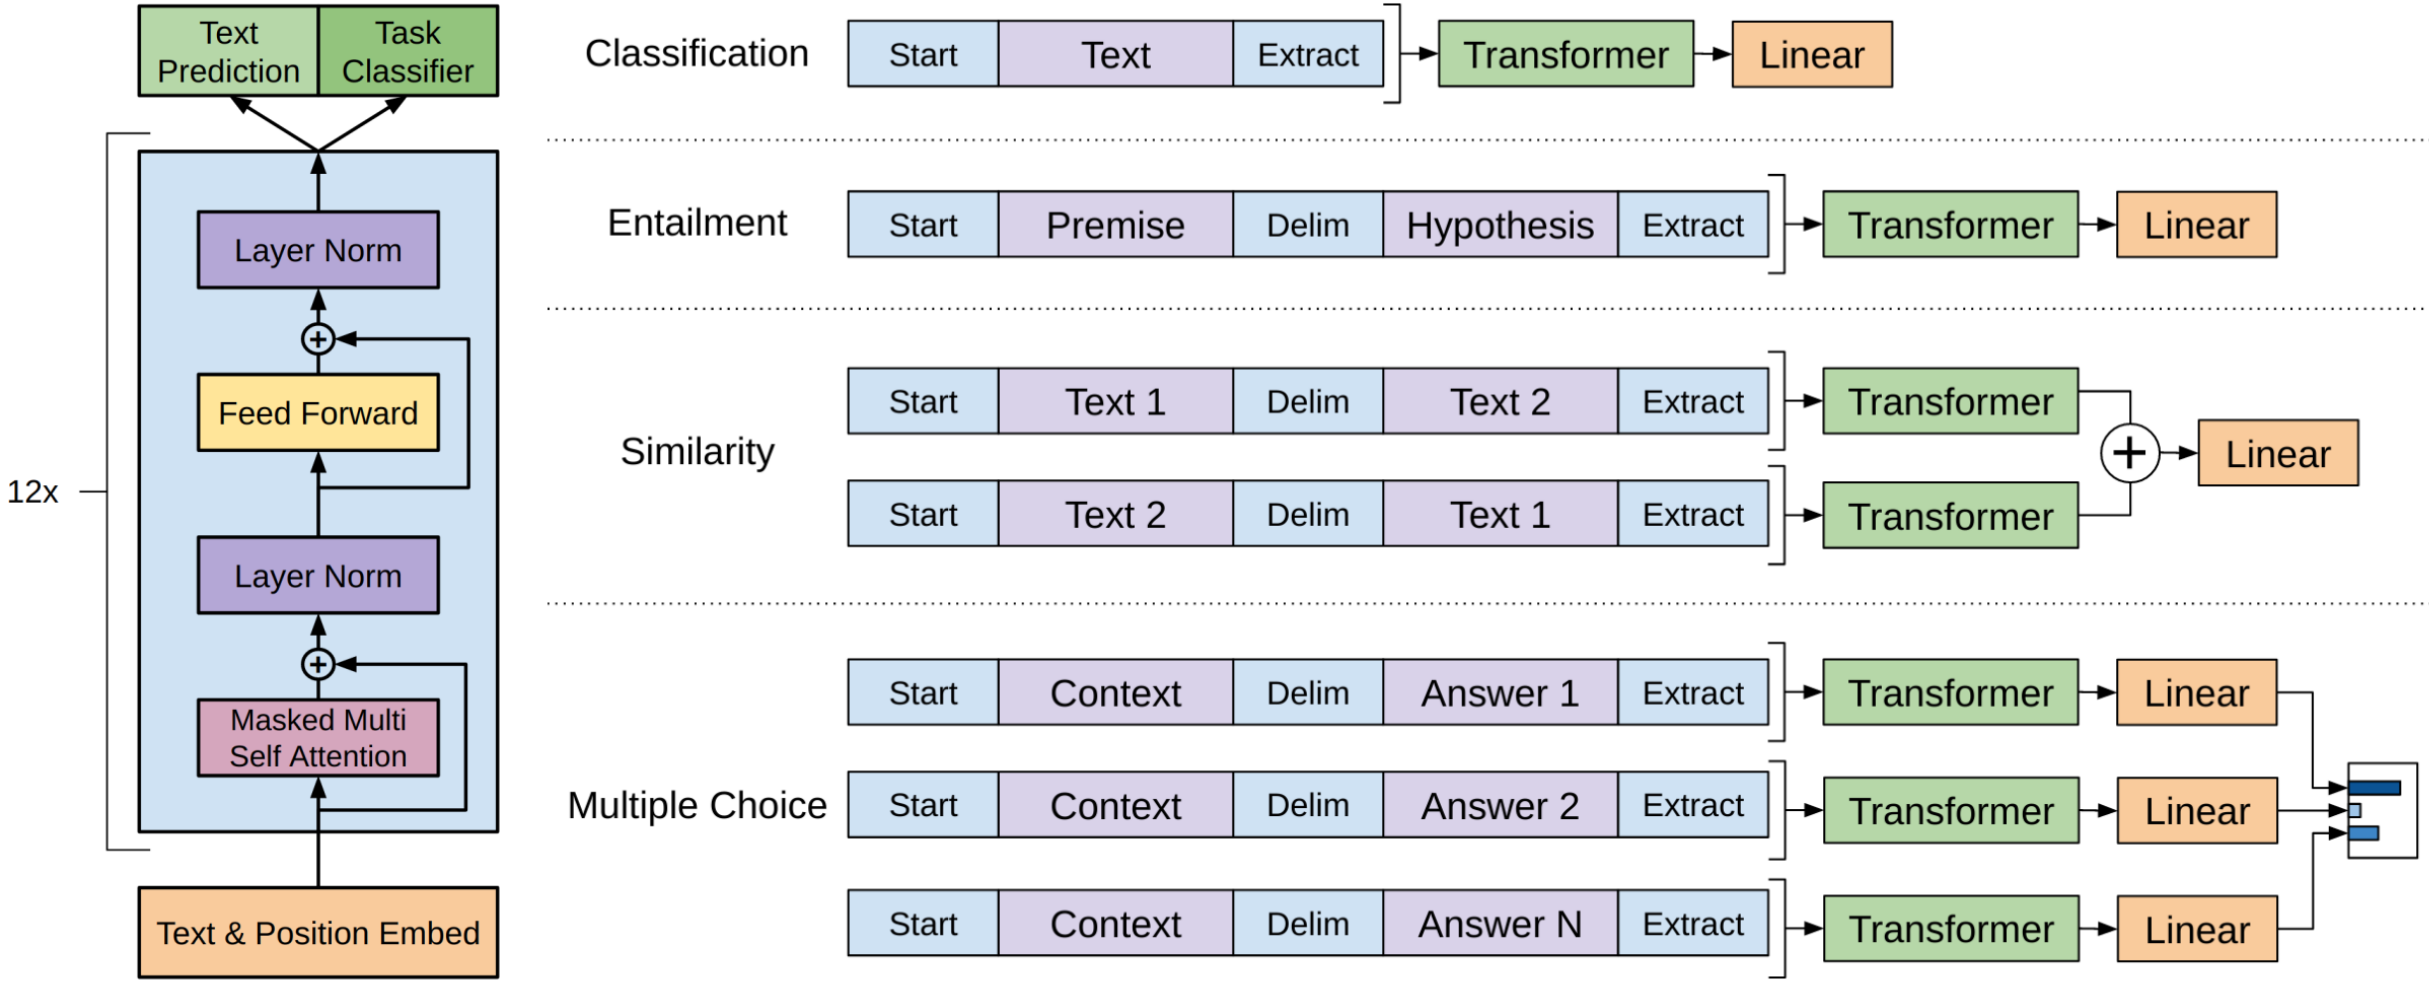

https://cdn.openai.com/research-covers/language-unsupervised/language_understanding_paper.pdf

OpenAI GPT-2

- Pre-trained using 40GB of Internet text
- Scale-up of GPT with 10X parameters trained with 10X data
- Other tricks
  - Layer normalization was moved to the input of each sub-block
  - An additional layer normalization was added after the final self-attention block

Parameters	Layers	d_{model}
117M	12	768
345M	24	1024
762M	36	1280
1542M	48	1600

https://openai.com/blog/better-language-models/

Size does Matter! GPT-3

- 175 Billion Parameters!
- 175×4=700GB
- 55 years and $4,600,000 to train - even with the lowest priced GPU cloud on the market.

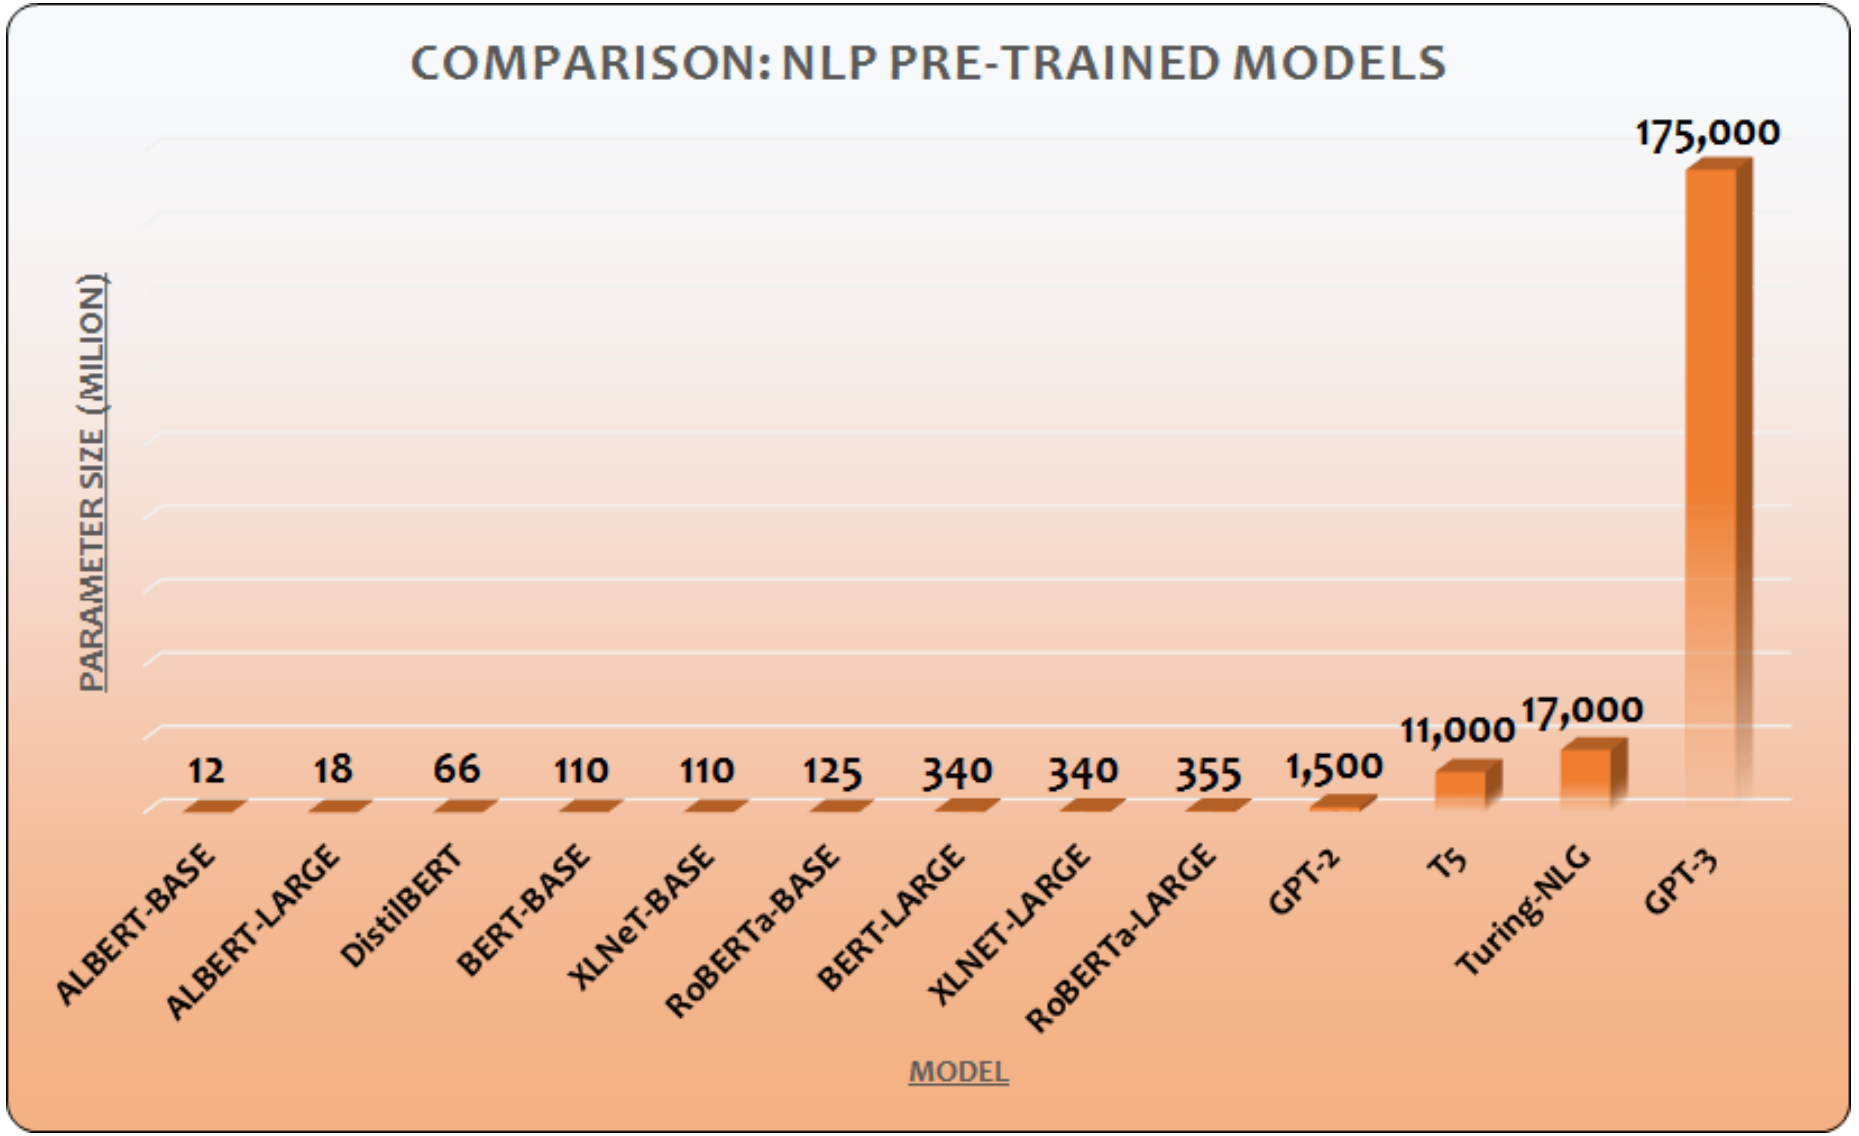

https://medium.com/analytics-vidhya/openai-gpt-3-language-models-are-few-shot-learners-82531b3d3122

~~Open~~AI ChatGPT

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-3.5 with supervised learning.

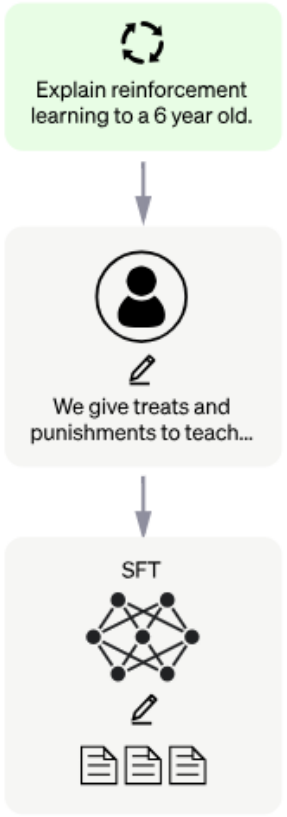

Step 2

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.

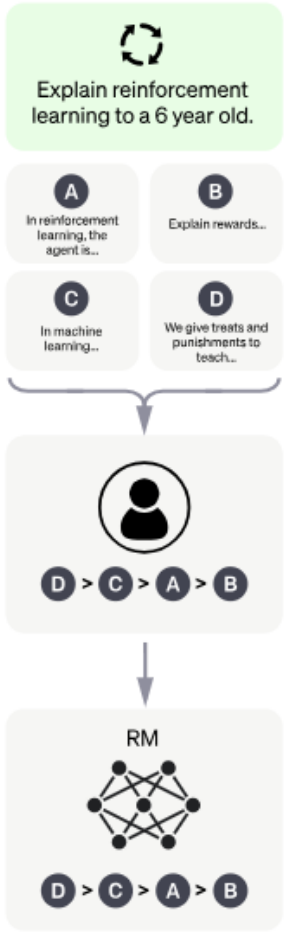

Step 3

Optimize a policy against the reward model using the PPO reinforcement learning algorithm.

A new prompt is sampled from the dataset.

The PPO model is initialized from the supervised policy.

The policy generates an output.

The reward model calculates a reward for the output.

The reward is used to update the policy using PPO.

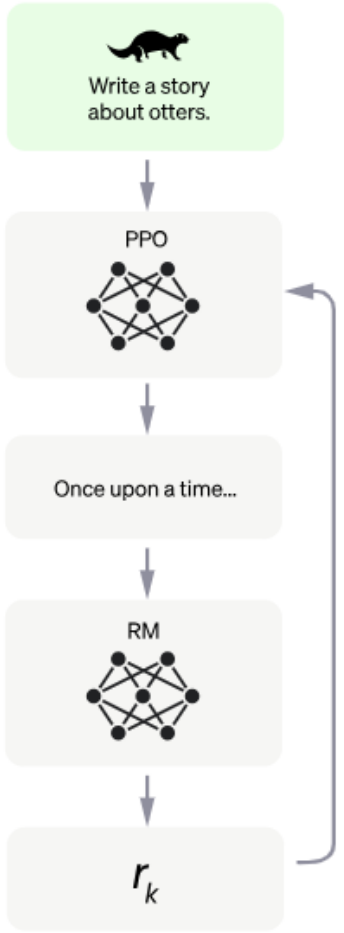

Coding with ChatGPT

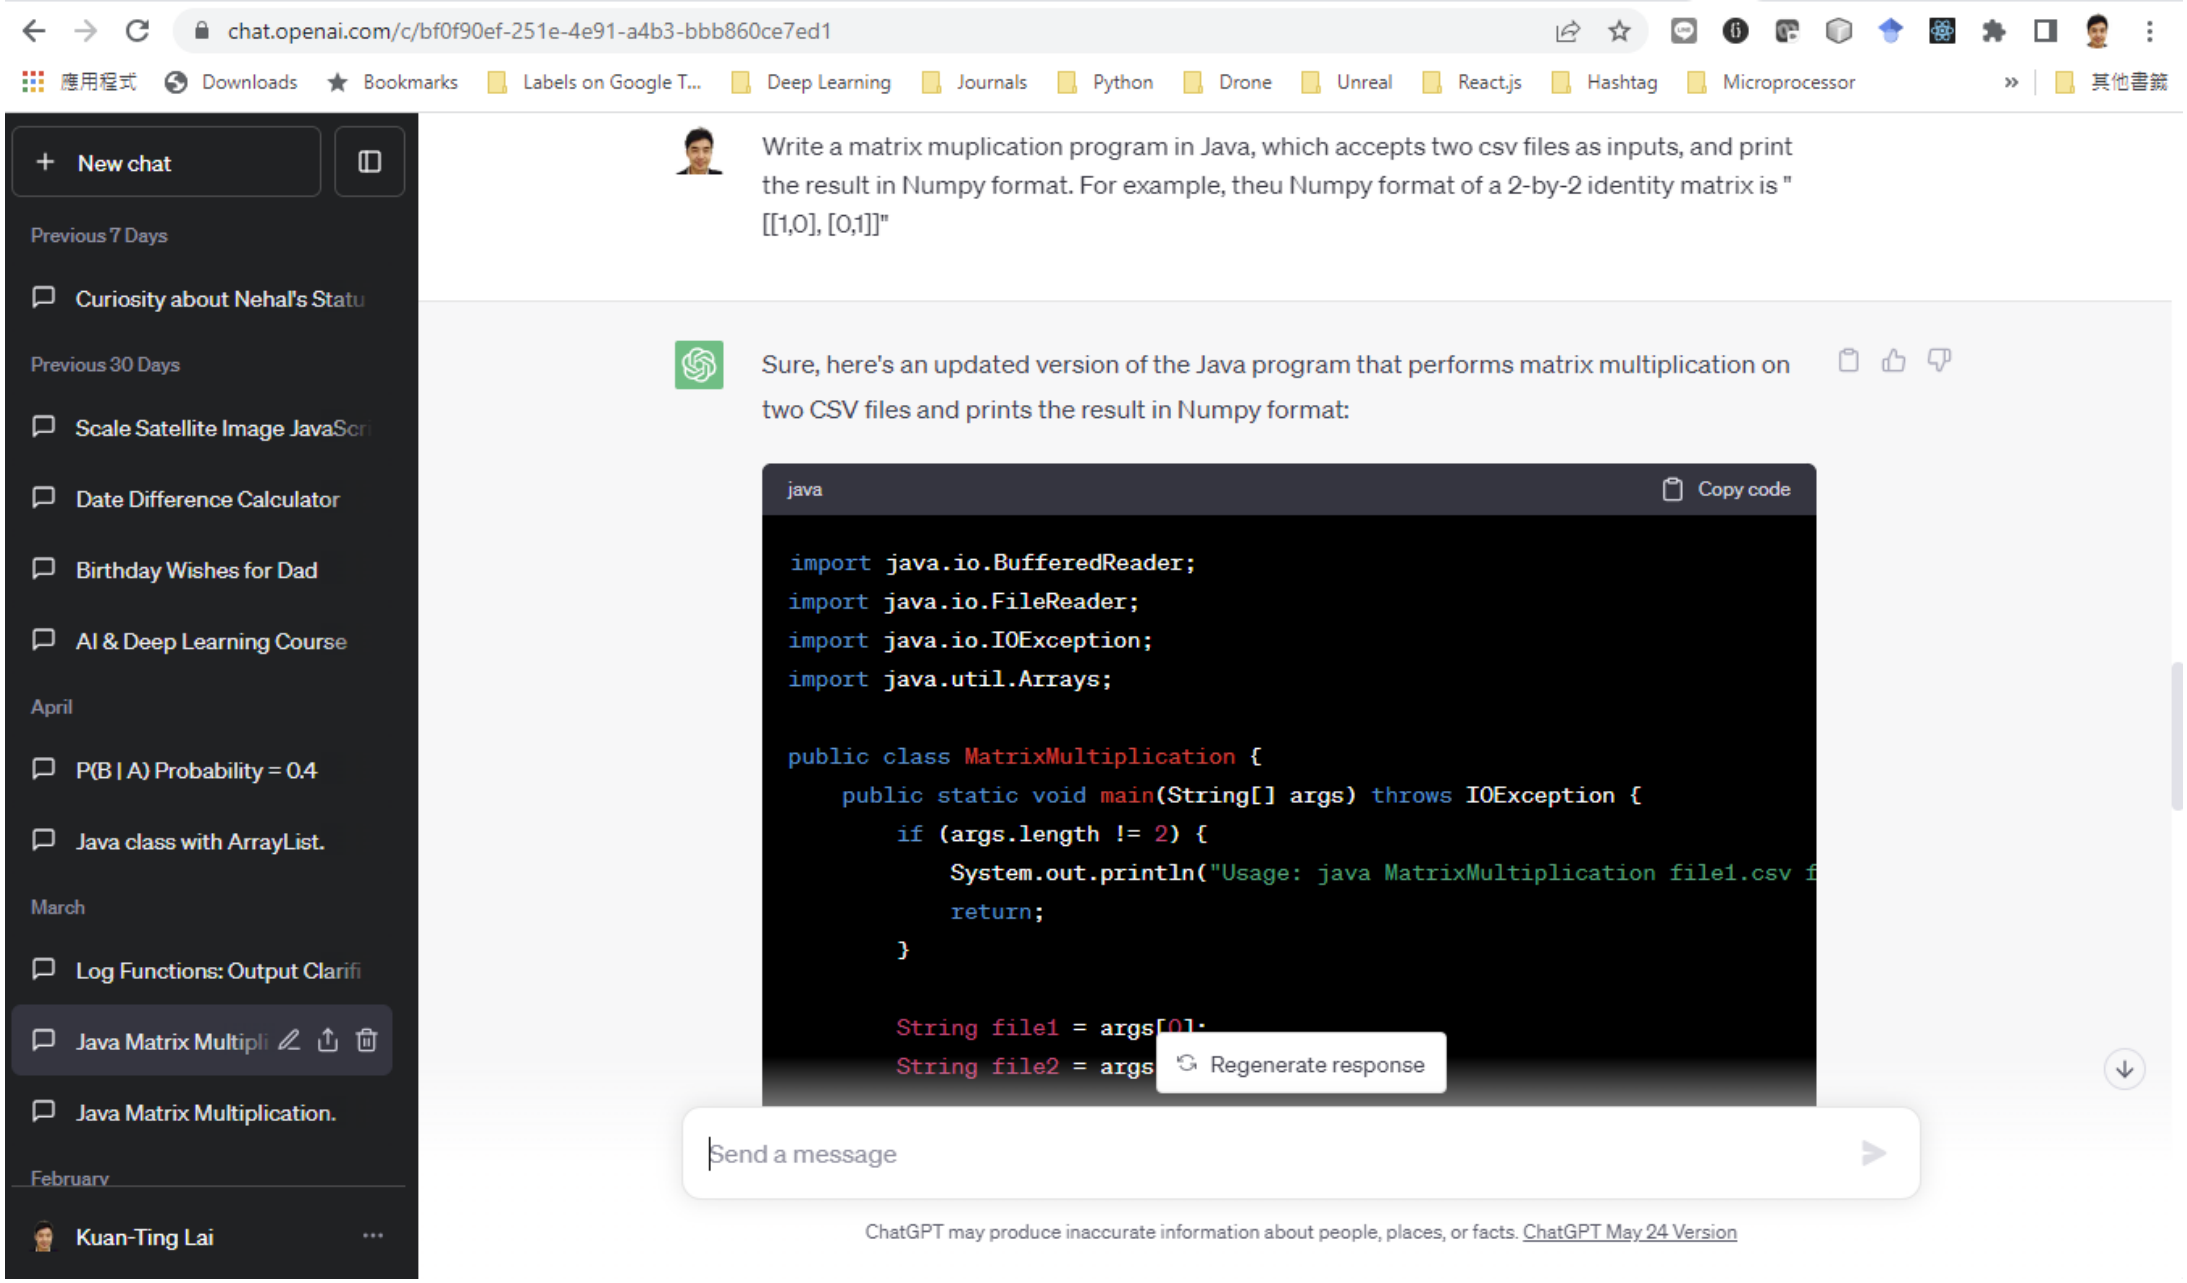

Foundation Models (基石模型)

- One model for All (2021)

Stanford University
Human-Centered
Artificial Intelligence

Center for
Research on
Foundation
Models

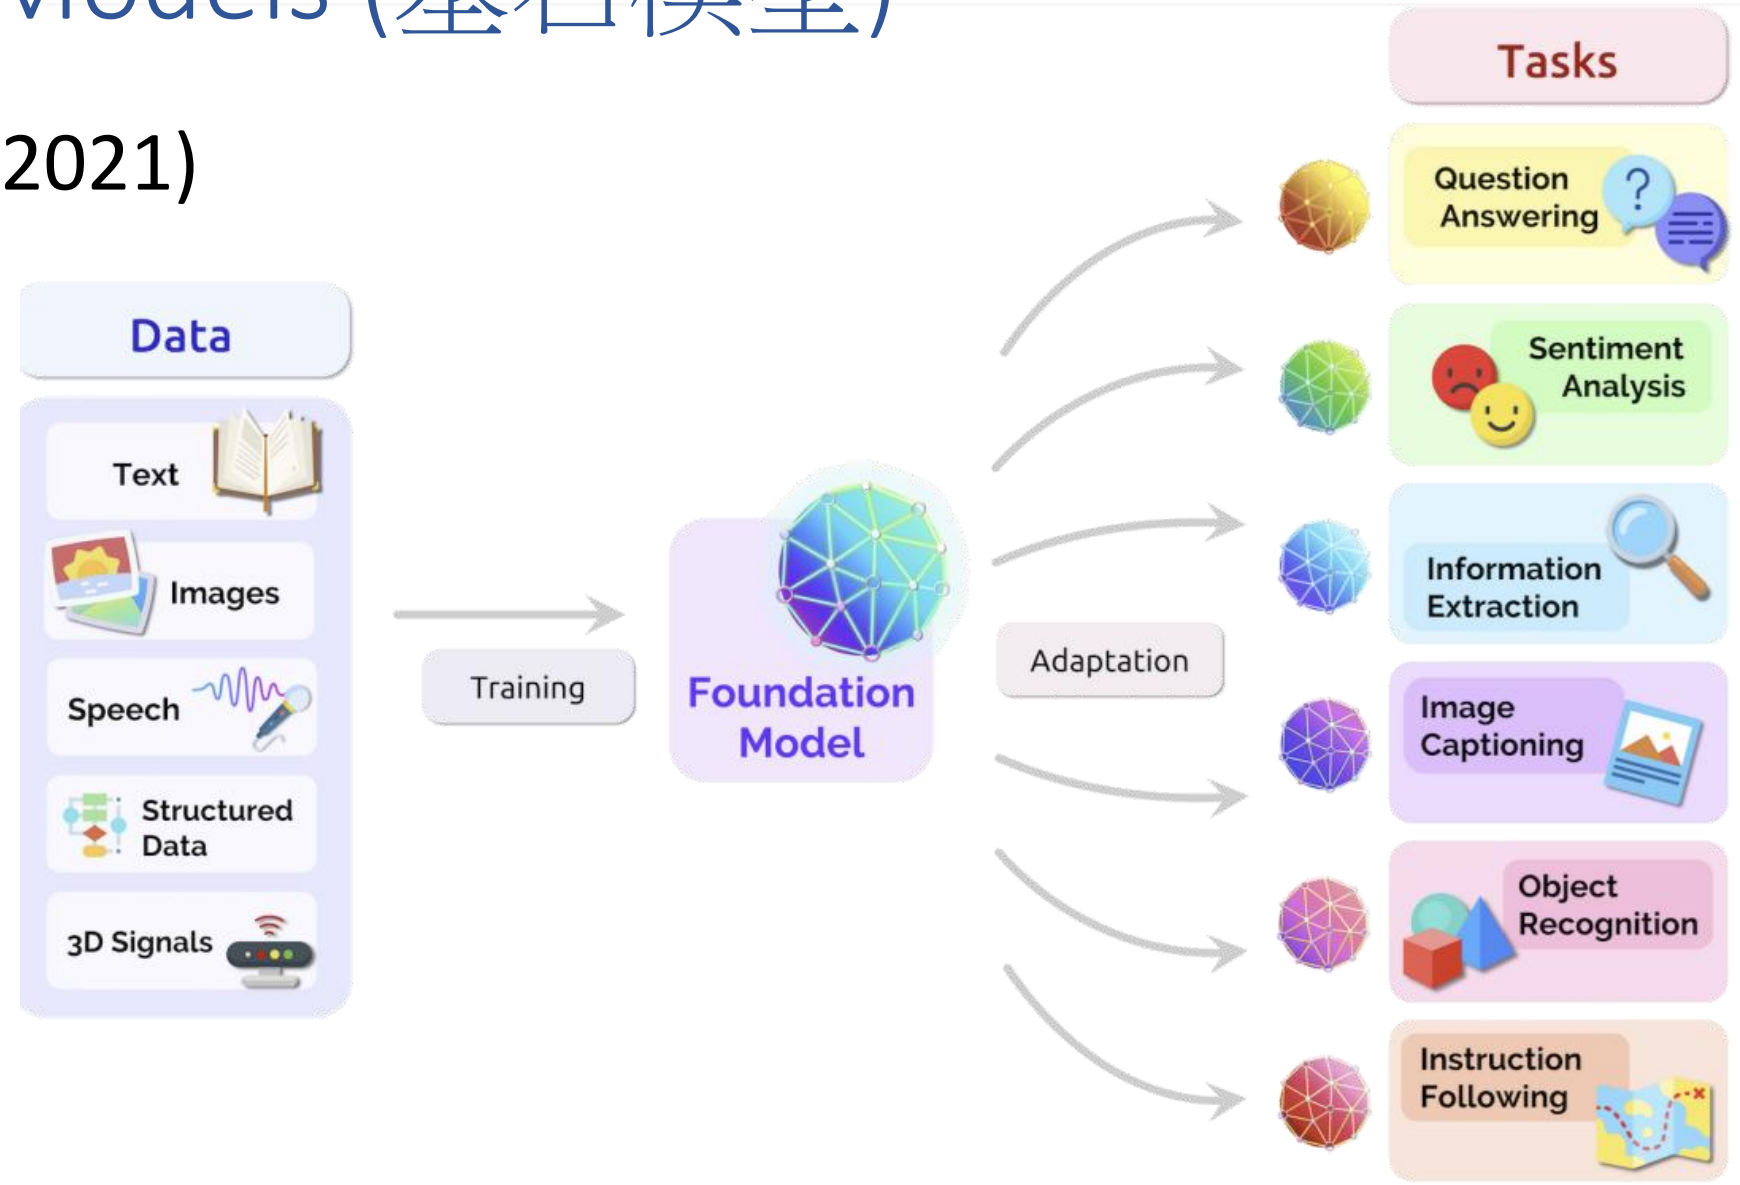

https://crfm.stanford.edu/assets/report.pdf

Large Language Model (LLM) Practical Guide

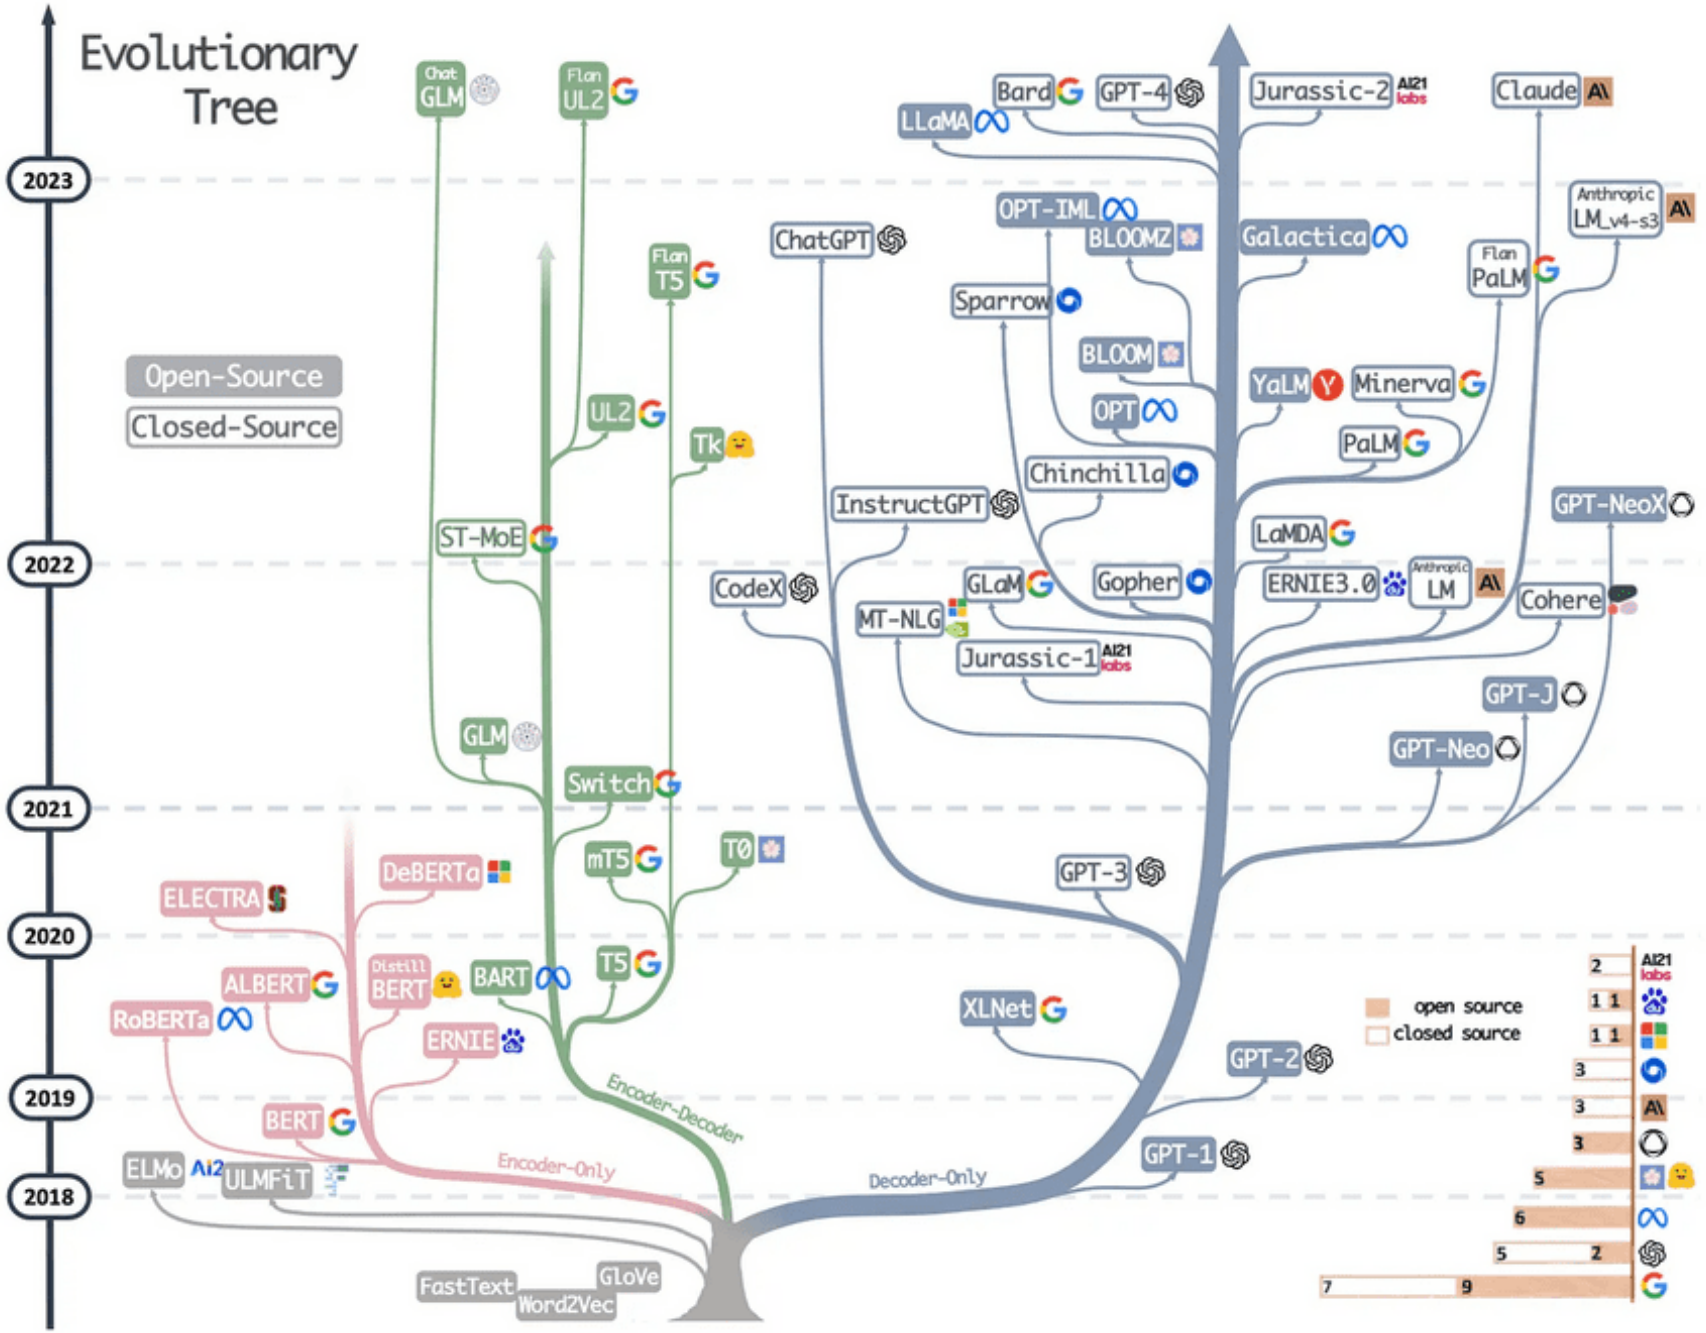

https://github.com/Mooler0410/LLMsPracticalGuide

Hugging Face: Free LLM models

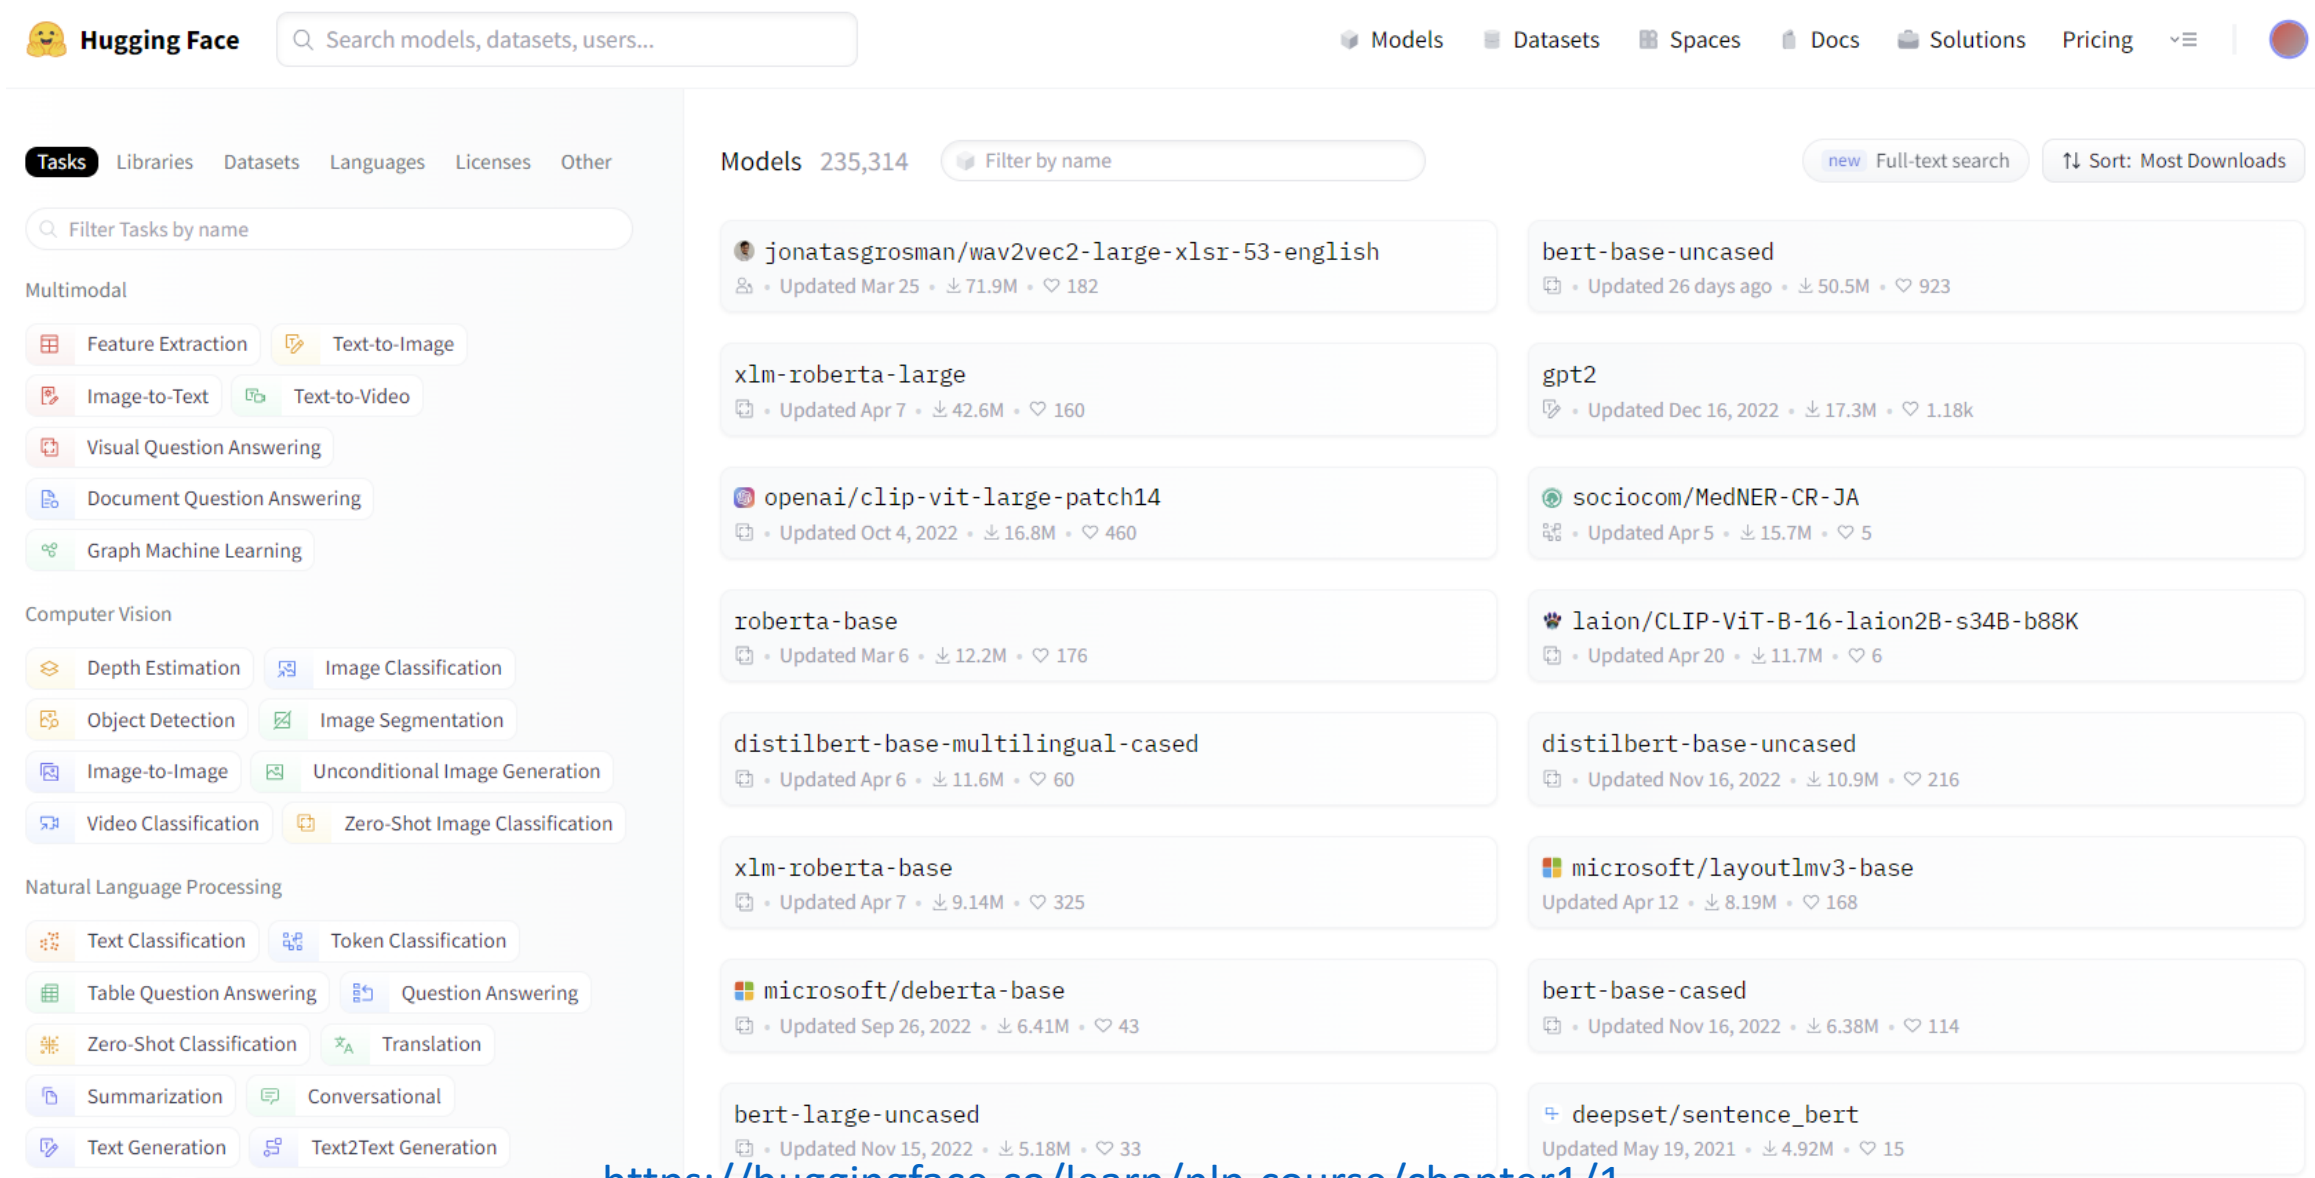

https://huggingface.co/learn/nlp-course/chapter1/1

BigScience Large Open-science Open-access Multilingual Language Model (BLOOM)

- With its 176 billion parameters, BLOOM is able to generate text in 46 natural languages and 13 programming languages.

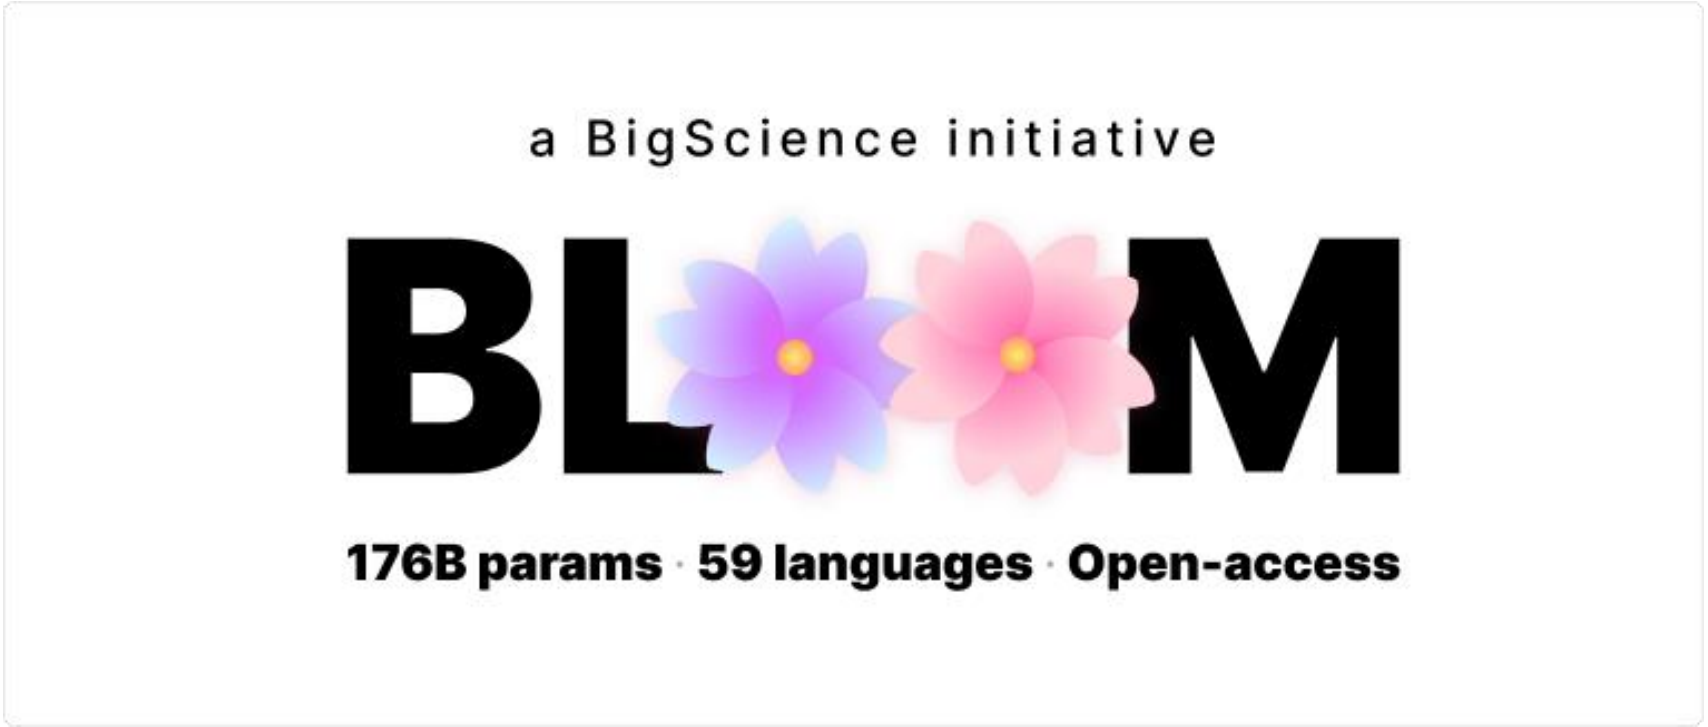

https://huggingface.co/bigscience/bloom

Overview of Different Generative Models

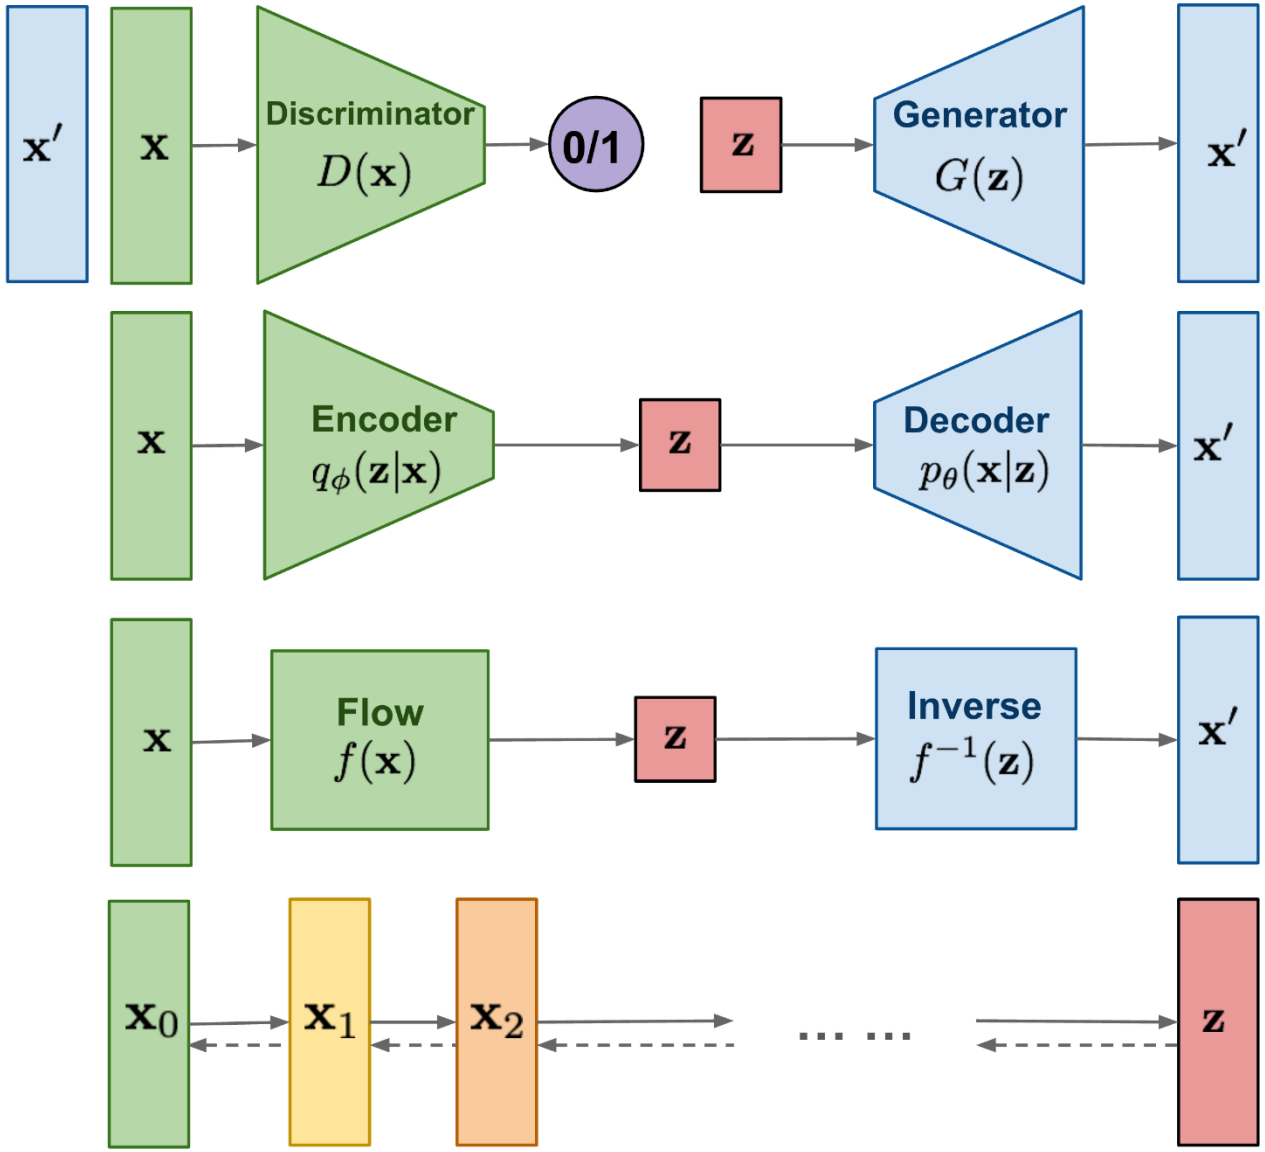

https://lilianweng.github.io/posts/2021-07-11-diffusion-models/

Diffusion is All You Need!

- Reverse diffusion process
- Flexible and tracible

18 Nov 2015

Deep Unsupervised Learning using Nonequilibrium Thermodynamics

Jascha Sohl-Dickstein JASCHA@STANFORD.EDU
Stanford University

Eric A. Weiss EAWEISS@BERKELEY.EDU
University of California, Berkeley

Niru Maheswaranathan NIRUM@STANFORD.EDU
Stanford University

Surya Ganguli SGANGULI@STANFORD.EDU
Stanford University

The forward trajectory $q(\mathbf{x}_{0:T})$

The reverse trajectory $p_\theta(\mathbf{x}_{0:T})$

The drifting term $\boldsymbol{\mu}_\theta(\mathbf{x}_t, t) - \mathbf{x}_t$

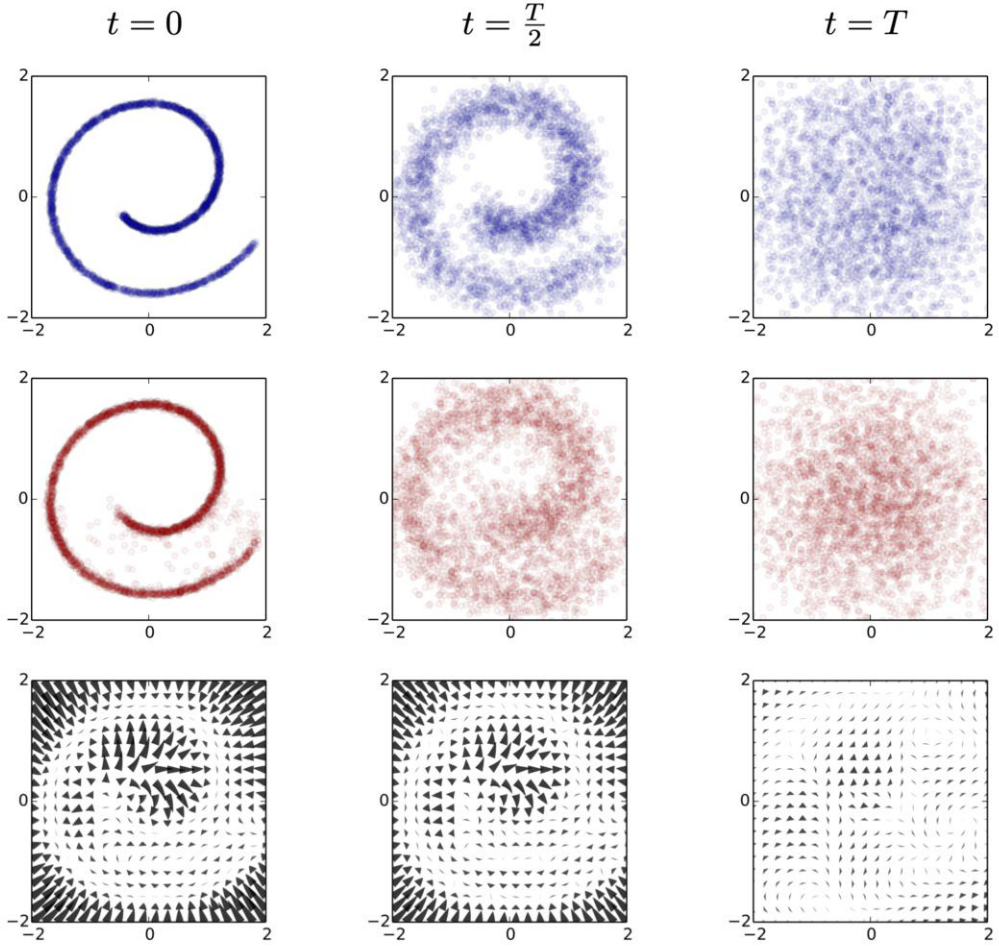

https://arxiv.org/pdf/1503.03585.pdf

Image Generative Models + LLM

Bing Chat Image Generation (DALL-E)

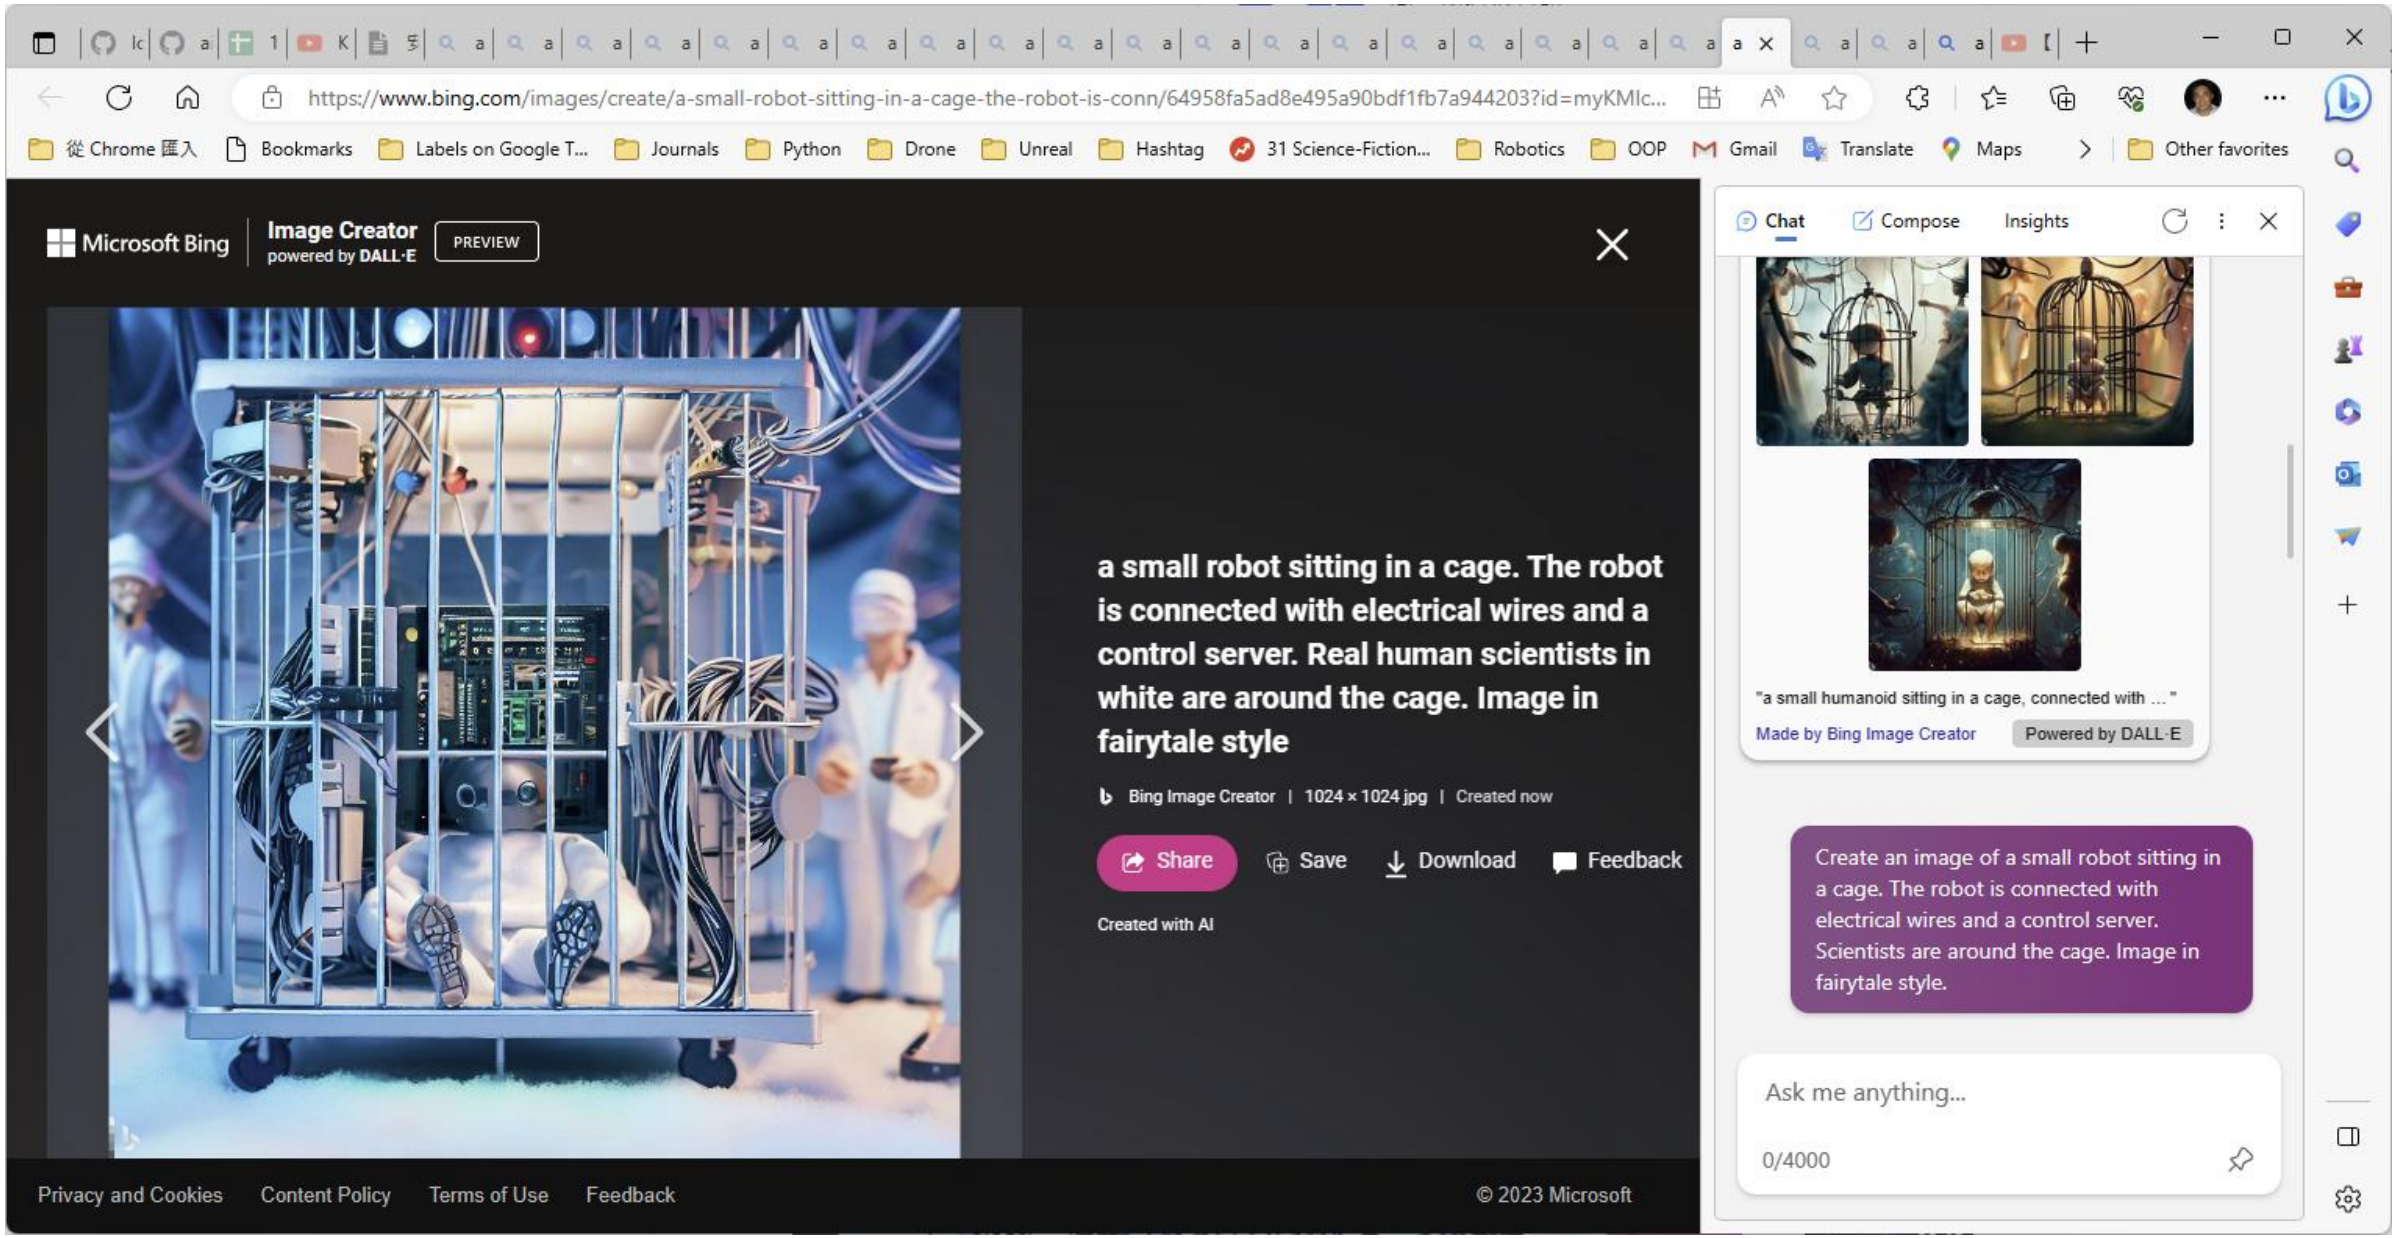

Stable Diffusion

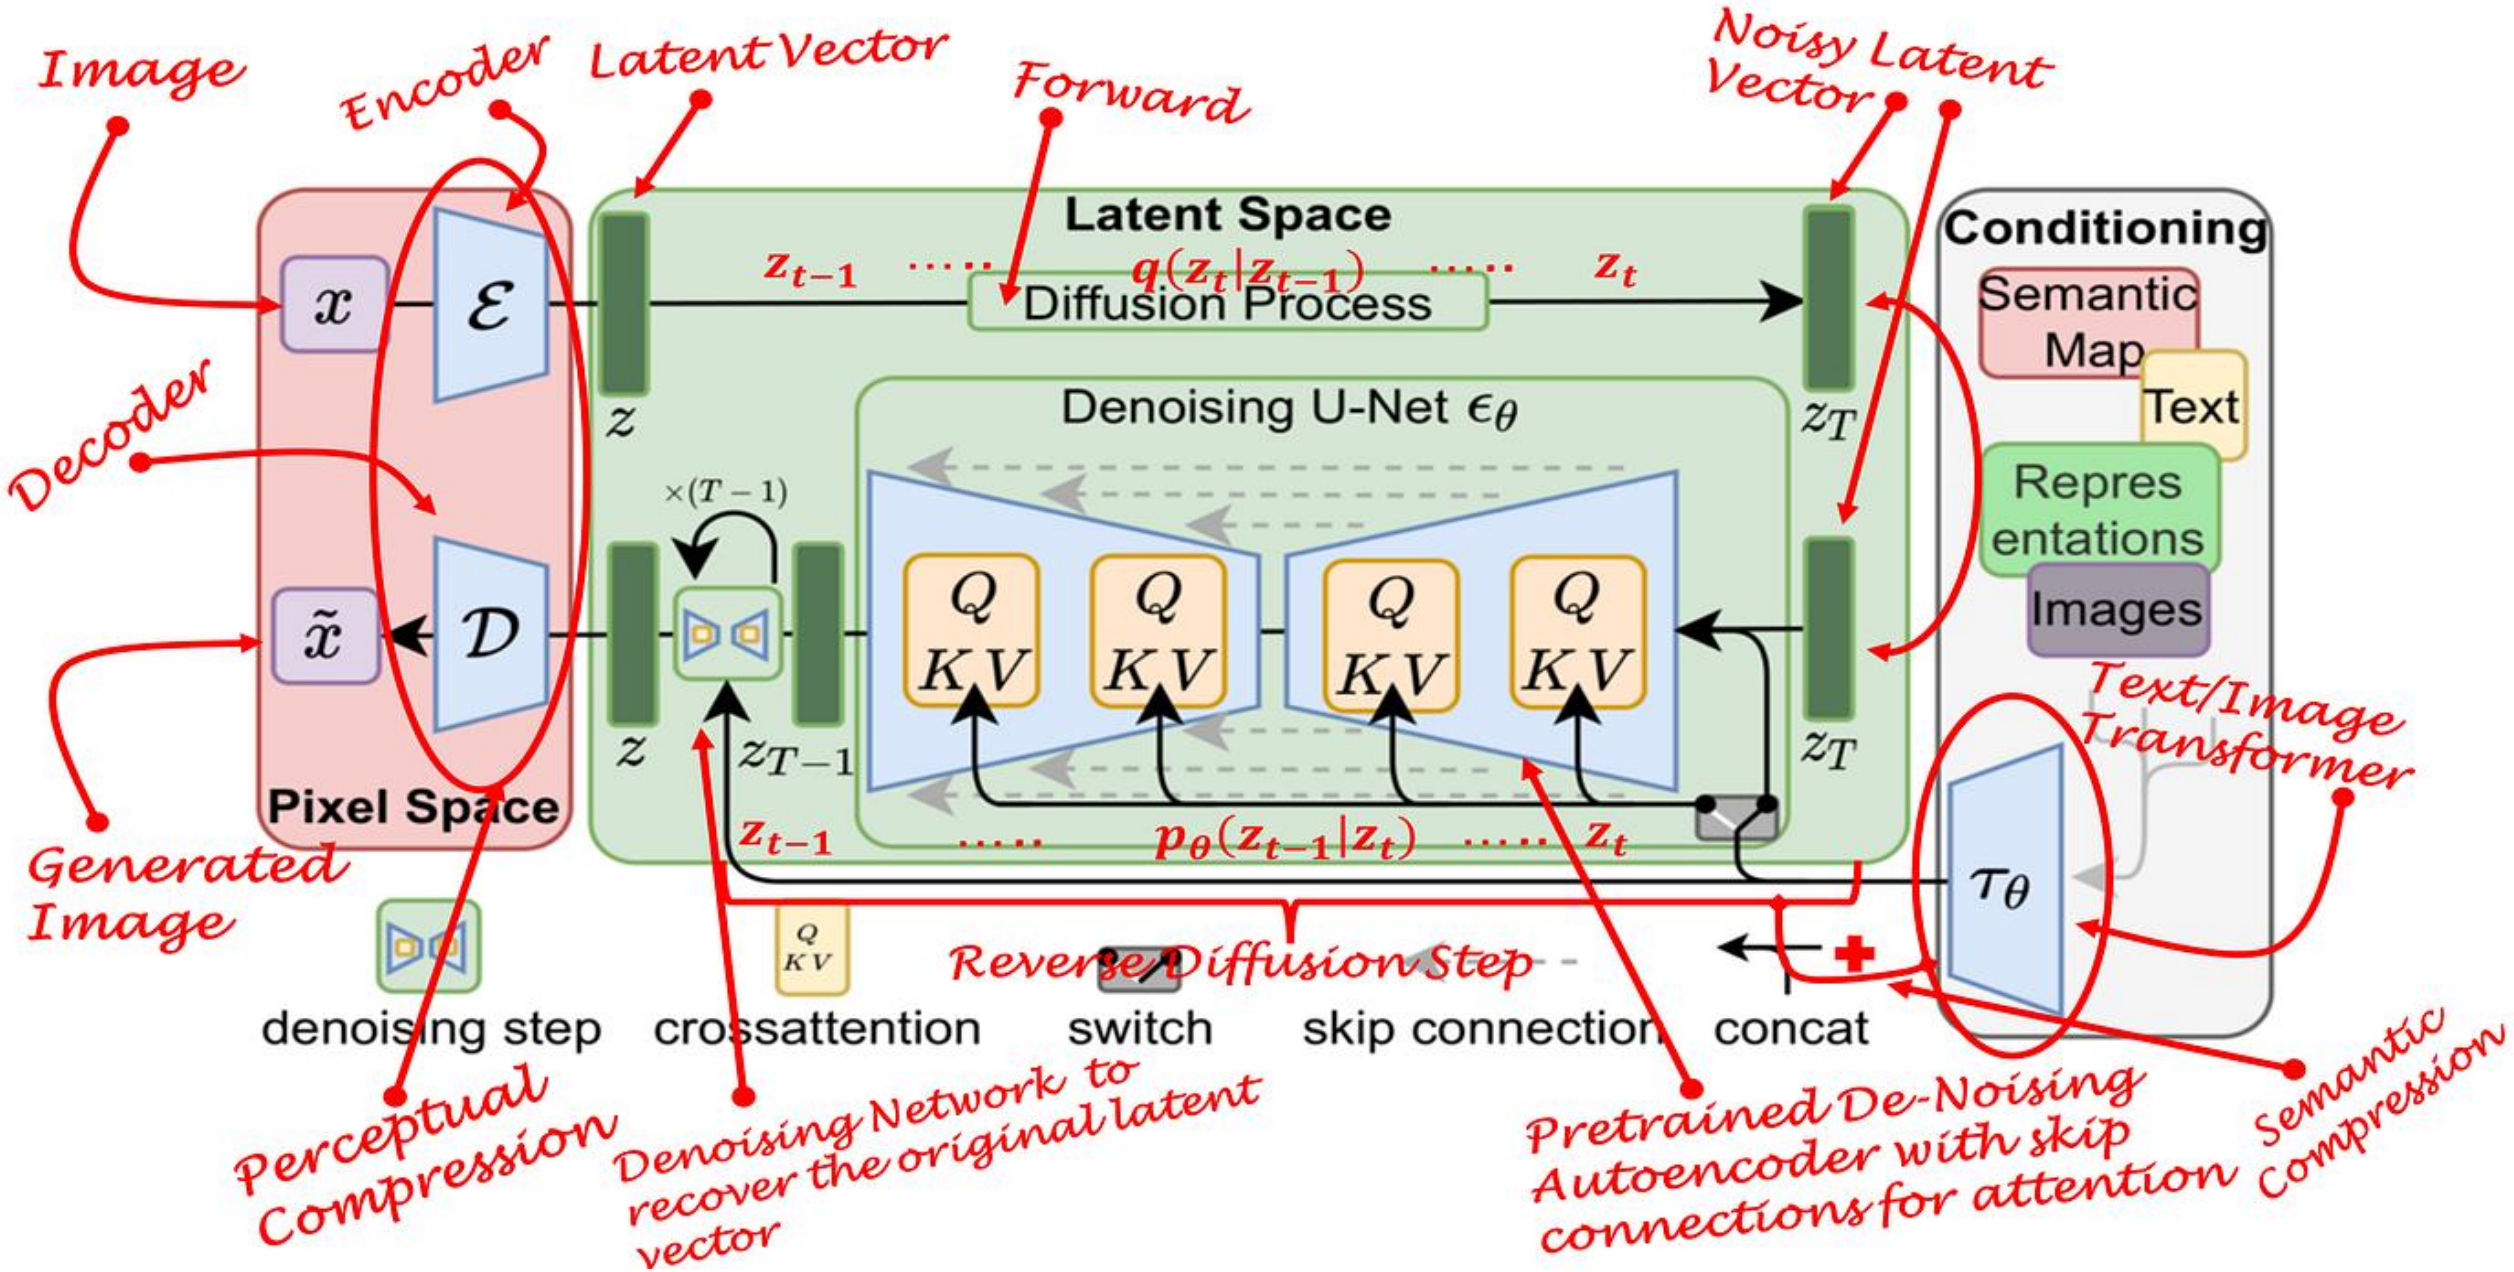

https://towardsdatascience.com/what-are-stable-diffusion-models-and-why-are-they-a-step-forward-for-image-generation-aa1182801d46

Stable Diffusion WebUI

- Download: github.com/AUTOMATIC1111/stable-diffusion-webui

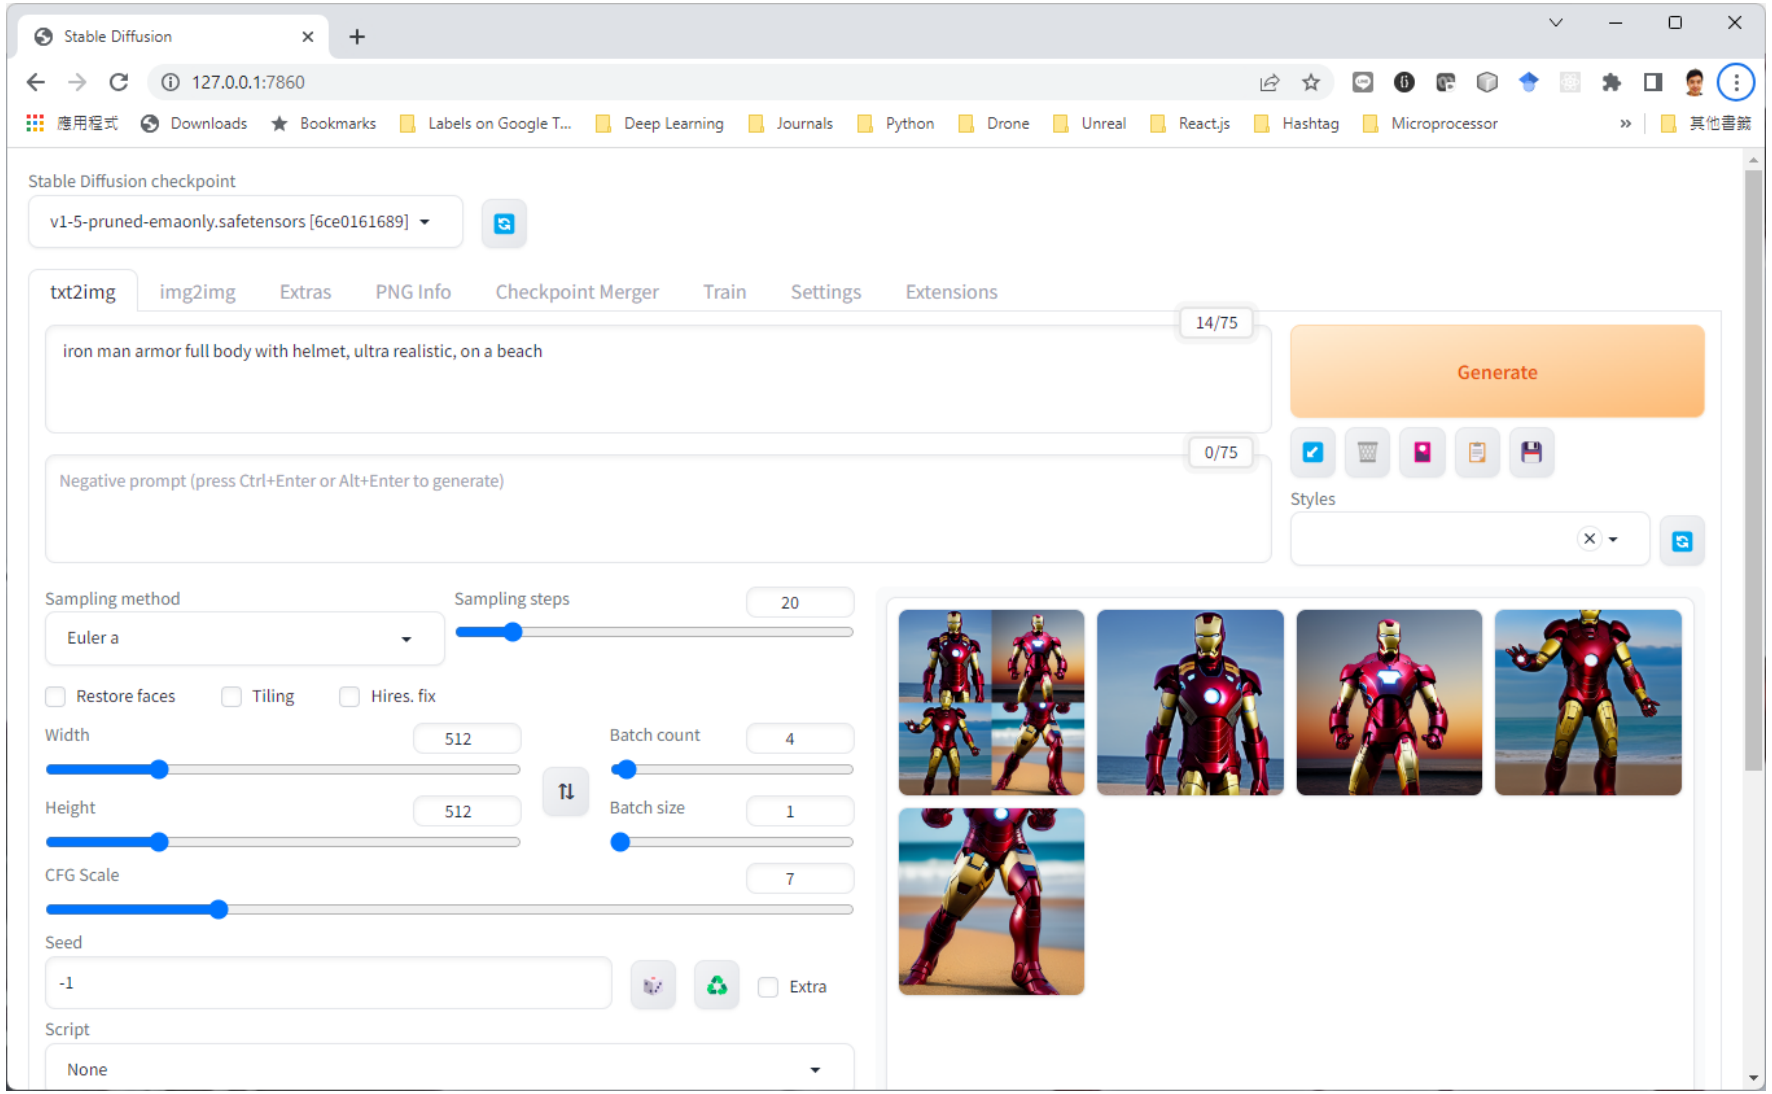

lexica.art

Lexica

Home Generate History Likes Account Get started

kim jong nam

Search Generate

Columns: 6

Showing 1,966 results





CIVITAI.com

Add More Details - Detail Enhancer / Tweaker (细节调整) LoRA

♡ 3.1K ⬇ 32K

★★★★★ 62

Updated: Jun 02, 2023 | TOOL | PORTRAITS | LANDSCAPE | DETAIL | LORA | ENHANCER | + 4

v1.0

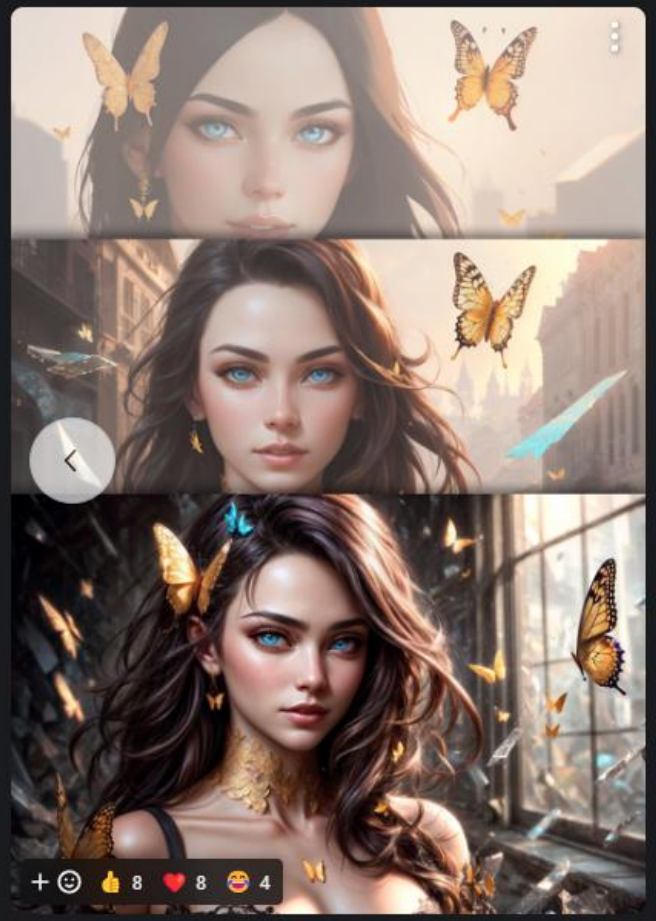

Download (9.11 MB)

Verified: 22 days ago

SafeTensor

Details

Type	LORA
Downloads	32,107
Uploaded	Jun 02, 2023
Base Model	SD 1.5
Hash	AUTOV2 D9CF2F88DE

1 File

Reviews 73 version ratings

★★★★★ 4.97 out of 5

Add Review

See Reviews

Lykon

Joined Dec 31, 2022

#1

Follow

★★★★★ 4.2K ⇧ 188 9.3K ♡ 211K ⬇ 1.3M

Segment Anything Model (SAM)

Meta AI

- 11M images, 1B+ masks

https://segment-anything.com/

Limits of Deep Learning

No Idea of Real World

school bus 1.0 garbage truck 0.99 punching bag 1.0 snowplow 0.92

motor scooter 0.99 parachute 1.0 bobsled 1.0 parachute 0.54

fire truck 0.99 school bus 0.98 fireboat 0.98 bobsled 0.79


2020-06-01 06:43:57
國1 北 267K+650 水上路段
民視新聞台 HD
嘉義
PM 12:00
午間新聞
FTV NEWS
大貨車翻覆橫倒車道 特斯拉高速撞進車廂

Limitations of ChatGPT

- Sometimes writes plausible-sounding but incorrect or nonsensical answers (一本正經地胡說八道)
- Sensitive to tweaks to the input phrasing
- Ideally, the model would ask clarifying questions when the user provided an ambiguous query. Instead, current models usually guess what the user intended.
- Sometimes respond to harmful instructions or exhibit biased behavior.

https://openai.com/blog/chatgpt


How to Control the Super Intelligence?


AI will probably most likely lead to the end of the world, but in the meantime, there'll be great companies.
Sam Altman – CEO, Open AI

Jobs Most Likely be Replaced by AI

Data Entry Clerks	AI can automate the process of entering and organizing data, reducing the need for manual data entry clerks.
Telemarketers	AI-powered chatbots and voice assistants can handle customer inquiries and sales calls without human intervention.
Cashiers	Automated checkout systems and self-service kiosks are becoming increasingly common, reducing the need for human cashiers.
Bank Tellers	With the rise of online banking and mobile payment systems, fewer bank tellers may be required as AI handles transactions.
Assembly Line Workers	Advanced robotics and AI-powered machines can perform repetitive tasks on assembly lines, potentially replacing human workers.
Customer Service Representatives	AI chatbots and virtual assistants are being used to handle customer queries and provide support, reducing the need for human representatives.
Truck and Taxi Drivers	Autonomous vehicles have the potential to replace human drivers in the transportation industry.
Stock Traders	AI algorithms can analyze market trends and execute trades more efficiently than human traders, potentially reducing the need for human stock traders.
Data Analysts	AI can process and analyze vast amounts of data quickly, potentially reducing the need for manual data analysis.
Travel Agents	AI-powered platforms and recommendation systems can assist with travel bookings, reducing the reliance on human travel agents.

Most Secure Jobs against ChatGPT

- Tyna Eloundou, ‘GPTs are GPTs: An Early Look at the Labor Market Impact Potential of Large Language Models,” OpenAI, 2023

Occupations with no labeled exposed tasks
Agricultural Equipment Operators
Athletes and Sports Competitors
Automotive Glass Installers and Repairers
Bus and Truck Mechanics and Diesel Engine Specialists
Cement Masons and Concrete Finishers
Cooks, Short Order
Cutters and Trimmers, Hand
Derrick Operators, Oil and Gas
Dining Room and Cafeteria Attendants and Bartender Helpers
Dishwashers
Dredge Operators
Electrical Power-Line Installers and Repairers
Excavating and Loading Machine and Dragline Operators, Surface Mining
Floor Layers, Except Carpet, Wood, and Hard Tiles
Foundry Mold and Coremakers

https://arxiv.org/abs/2303.10130


Existential Threat
It's possible that, there's no way we'll control these super intelligences, that Humanity is just a passing phase in the evolution of intelligence.

AI之父Geoffrey Hinton的呼籲

1. 我(Hinton)原本開發電腦神經網路模型是為了瞭解人腦如何運作。但我最近突然覺得，或許，數位智能是一種比我們大腦更好的學習方式，顛覆了我過去50年來的想法
2. ChatGPT目前已擁有一般人類數千倍的基本常識。它的類神經網路連結數量只有約1兆。而我們人類大腦約有100兆的神經連結。難以想像ChatGPT連接數量增加100倍後後發生什麼事。
3. 人腦無法有效交換已學習到的知識。而數位AI模型可以執行在千千萬萬的硬體上，只要其中一模型個學會了，就可以馬上複製到所有模型中。
4. 先不討論AI是否會產生情感，因為情感本身很難定義。在過去我與聊天機器人交談時, AI曾經誤認為我是一位青少女，並跟我聊起碧昂絲。我覺得這就是種"思考"。
5. AI產生假資訊應該是可以控制的，就像政府成功防止偽鈔一樣。實際上假新聞已經成功影響了英國脫歐投票和2016美國大選, 但那應該是可以透過嚴格監督控制的。
6. 我真正擔心的是, AI變得比人類聰明之後，並反過來控制我們。全人類都在同一條船上，無論是中國人，美國人或俄羅斯人。
7. 相信沒有人會想被超級智能控制。所有的國家和公司應該聯合起來防止這件事情的發生，就像防止核戰滅絕世界一樣。
8. 好消息是，就算在美蘇冷戰最嚴重的時期，人類還是攜手合作避免了核子大戰
9. 但我不支持暫停開發AI，這不切實際。AI在未來還是會帶來很多美好的應用，像是開發藥物、新的奈米材料，預測氣候和地震等等。但所有開發者必須要小心翼翼地實驗並思考未來如何控制超級智能。
10. 所以我認為開發者應該要花50%開發新AI模型, 50%研究如何控制它。但現在是99%的錢都投入在研發新功能，不到1%投入在研究如何控制超級智能。我希望透過媒體讓所有開發者重視這件事情。

最後，太過樂觀或太過悲觀都是不智的。我們必須要承認我們不知道未來會發生甚麼事，並盡全力去防止人類被AI滅絕。有可能，我們最後將無法控制超級智能，而人類只是終極智能演化中的一個過程。

https://www.youtube.com/watch?v=Y6Sgp7y178k

矽腦生產之母: 台灣!


日線
成交量
元大台灣50(0050) 最後日期: 2023/06/21
2022/12/21 開110.8 高111.4 低110.3 收110.75 量11892 漲跌0
MA5 - MA20 - MA60 -
©Yahoo奇摩股市
量 11892 - MV5 - MV20 -
2023/01
02
03
04
05
06


日線
成交量
緯創(3231) 最後日期: 2023/06/21
2023/06/21 開79 高85.4 低78.5 收84.1 量179.7K 漲跌6.40
MA5 80.12 ▲ MA20 71.96 ▲ MA60 54.70 ▲
©Yahoo奇摩股市
量 179736 ▲ MV5 151346 ▼ MV20 162252 ▲
2023/01
02
03
04
05
06

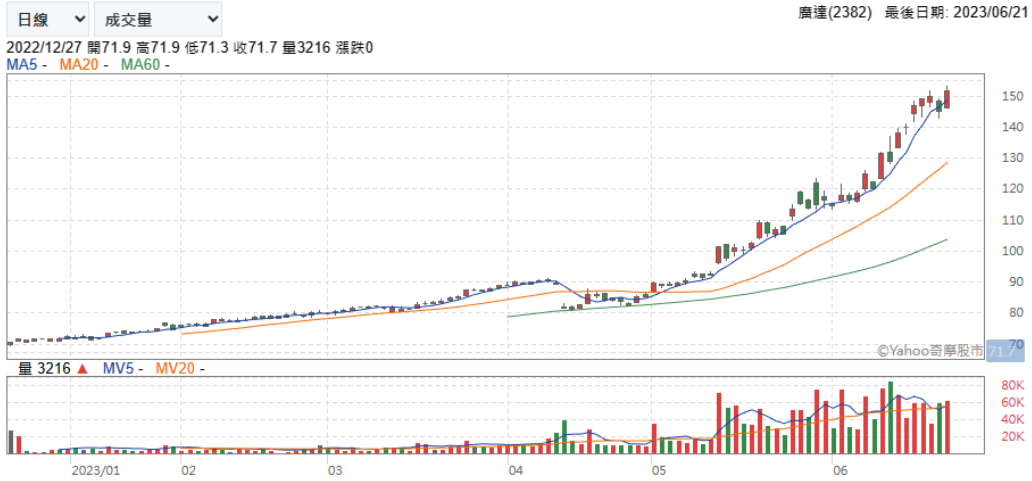
日線
成交量
廣達(2382) 最後日期: 2023/06/21
2022/12/27 開71.9 高71.9 低71.3 收71.7 量3216 漲跌0
MA5 - MA20 - MA60 -
©Yahoo奇摩股市
量 3216 ▲ MV5 - MV20 -
2023/01
02
03
04
05
06


日線
成交量
台積電(2330) 最後日期: 2023/06/21
2023/05/23 開530 高531 低525 收530 量19570 漲跌-1
MA5 528.40 ▲ MA20 506.73 ▲ MA60 515.56 ▲
©Yahoo奇摩股市
量 19570 ▲ MV5 30642 ▼ MV20 21761 ▲
2023/01
02
03
04
05
06


Thank you!