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

Machine Learning
機器學習

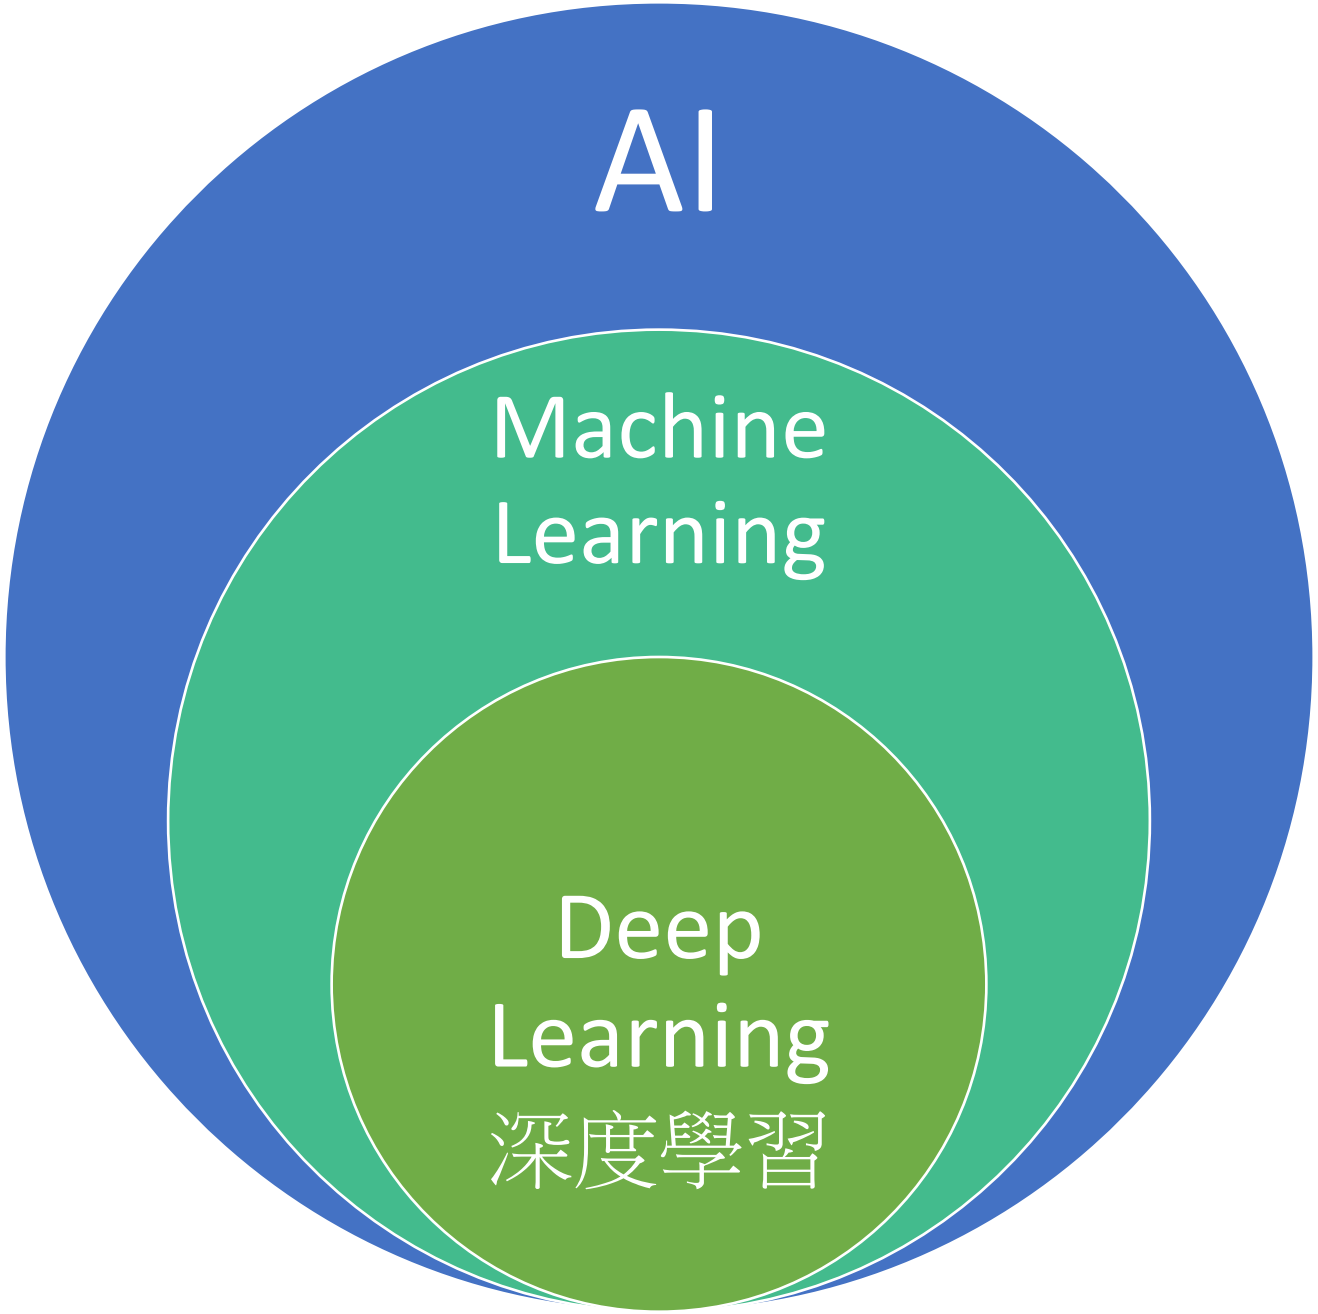
AI
Machine Learning
Deep Learning
深度學習

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

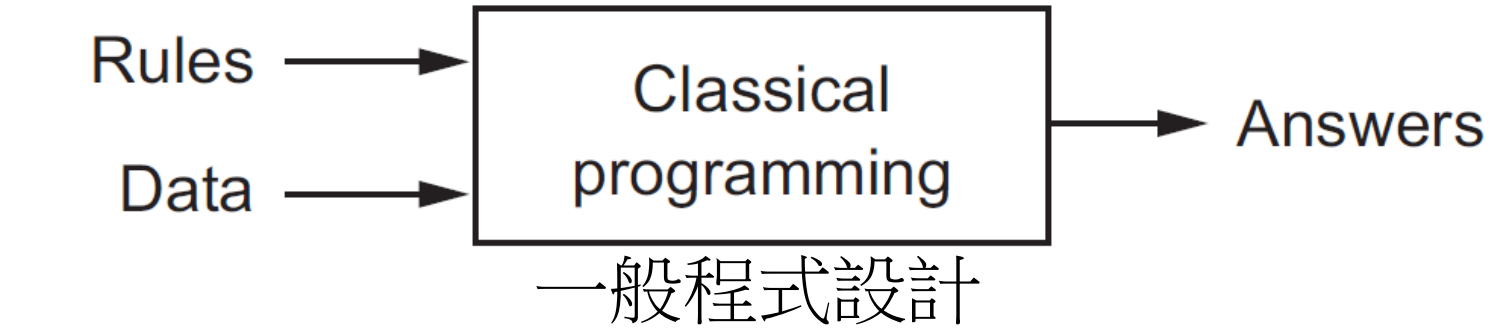

一般程式設計

機器學習

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

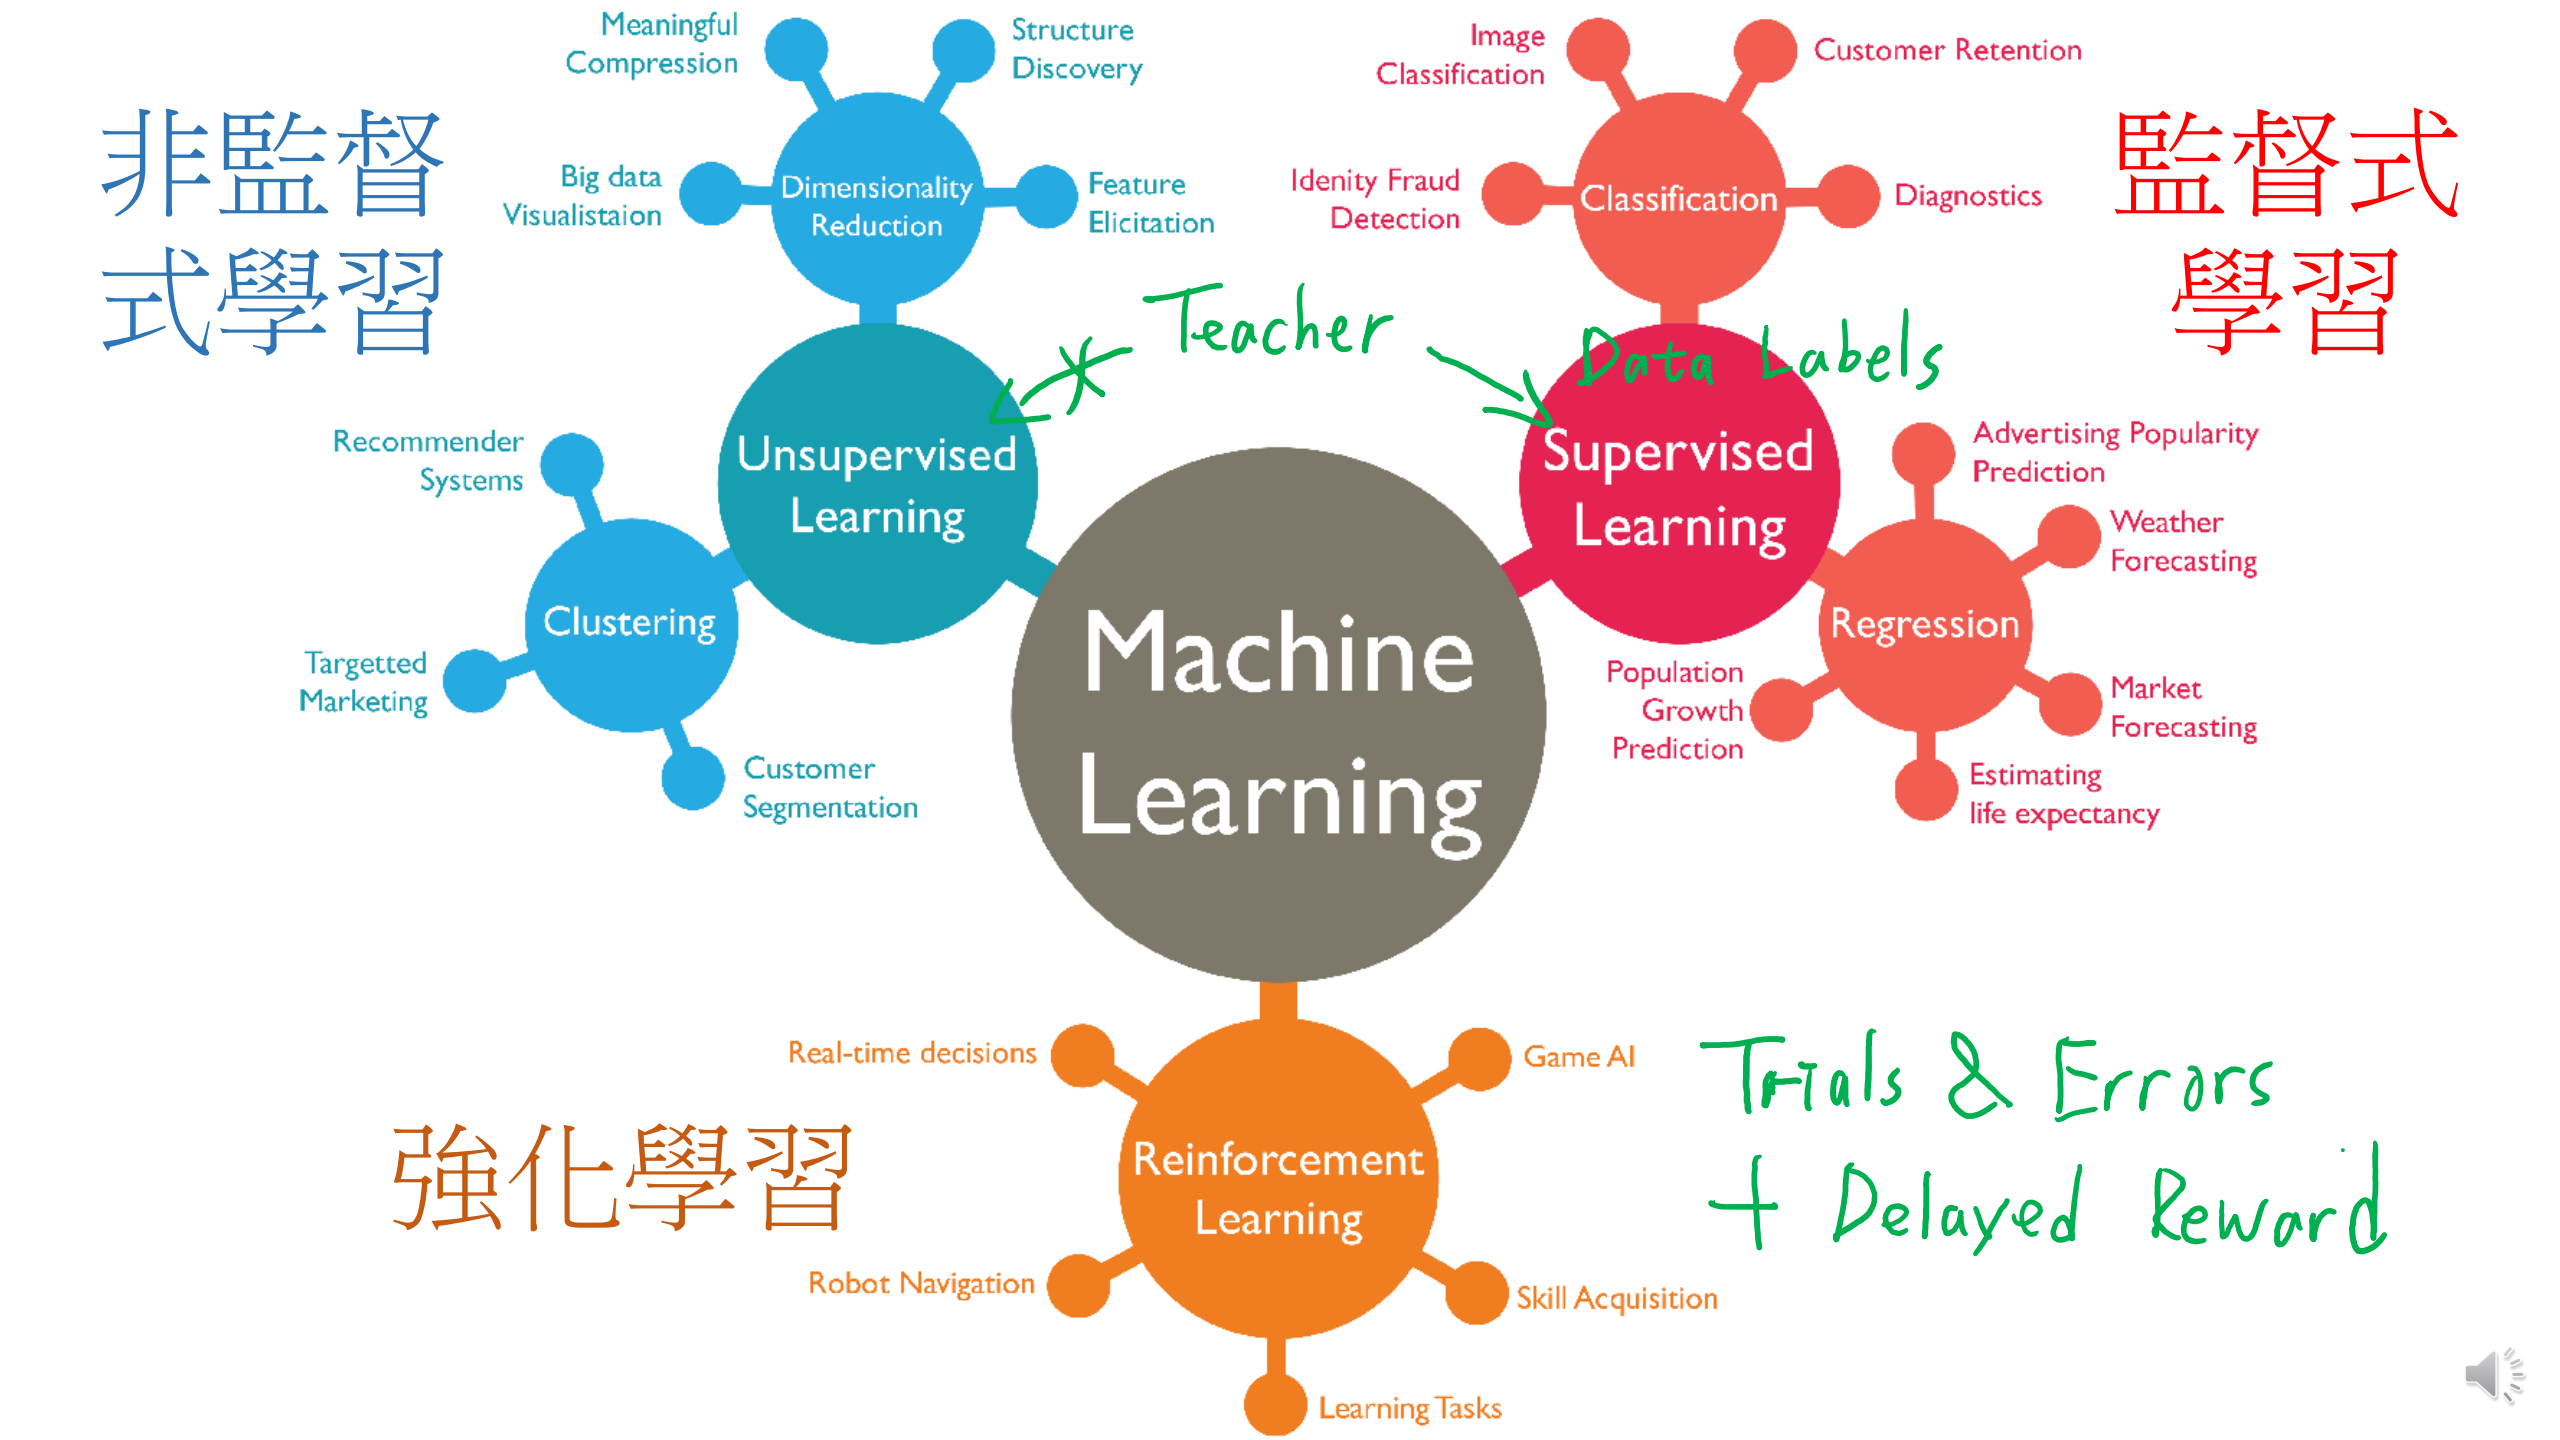
非監督式學習
Meaningful Compression
Structure Discovery
Big data Visualistaion
Dimensionality Reduction
Feature Elicitation
Unsupervised Learning
Recommender Systems
Clustering
Targetted Marketing
Customer Segmentation
Teacher
Data Labels
Machine Learning
監督式學習
Image Classification
Customer Retention
Idenity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

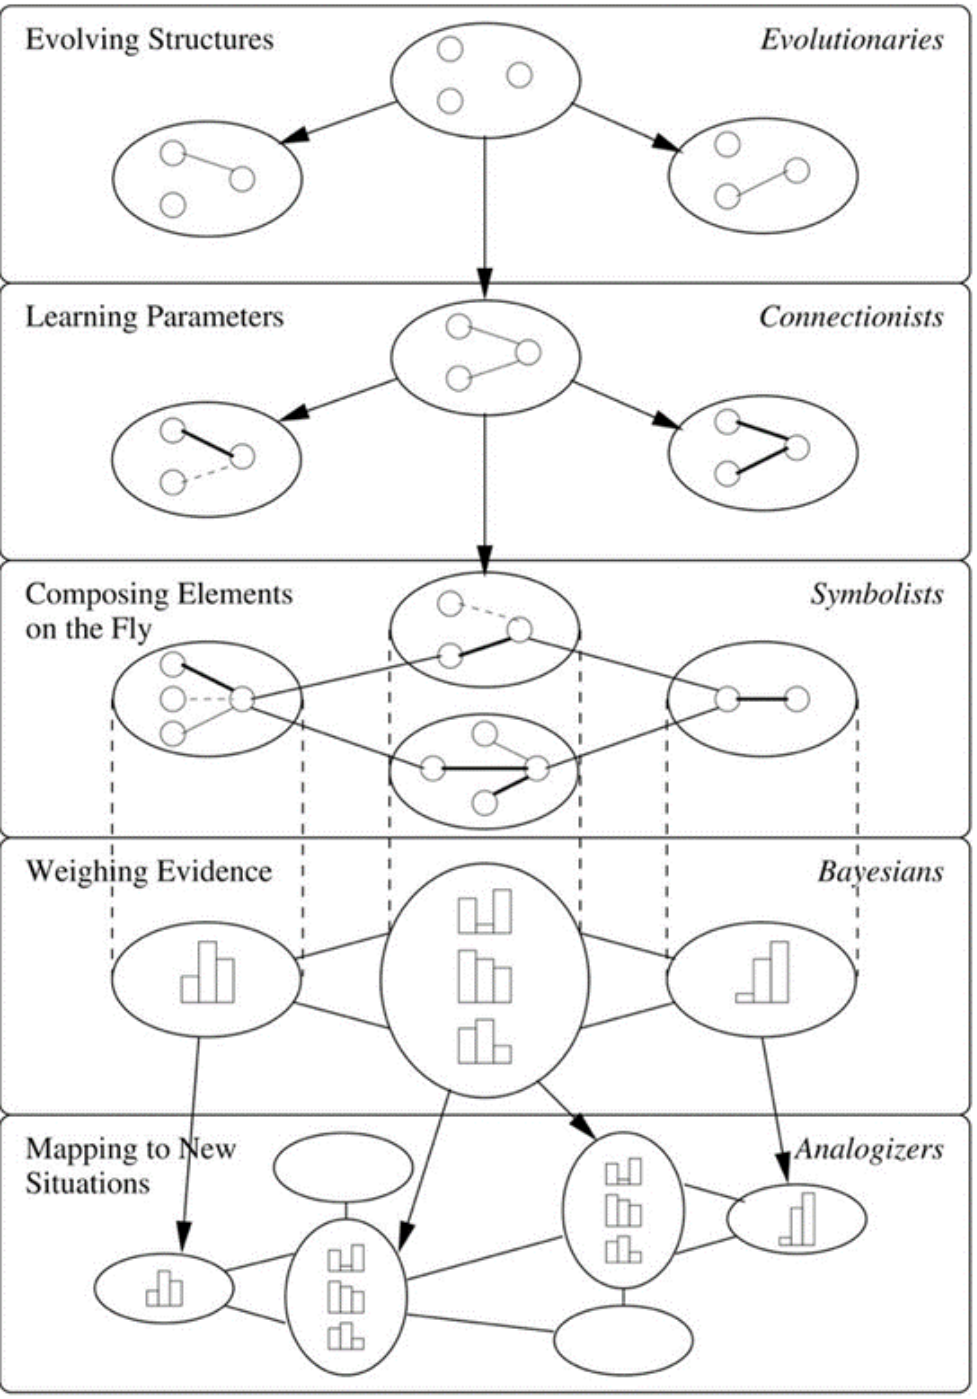

Five Tribes of Machine Learning
機器學習的五大門派

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

Pedro Domingos

The Master Algorithm – Pedro Domingos

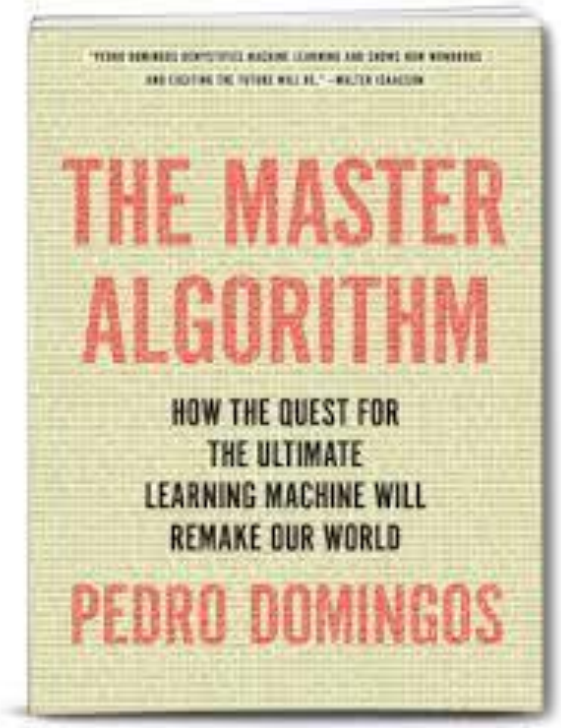

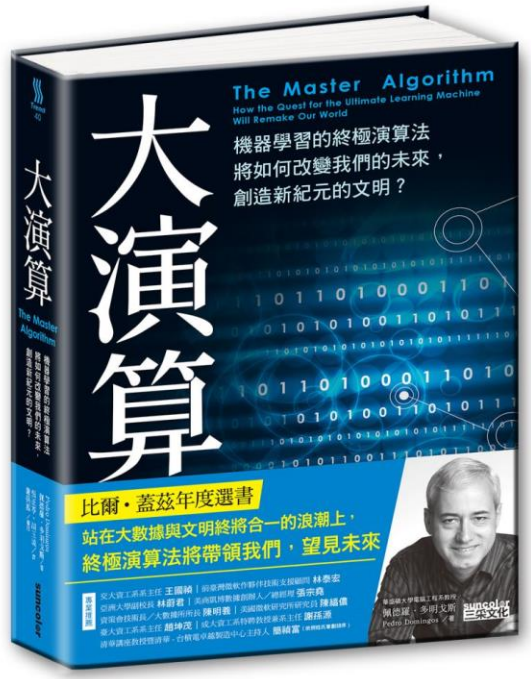

3 Basic Operations of Algorithms
演算法的三大元素

- All Algorithms can be Reduced to 3 Operations

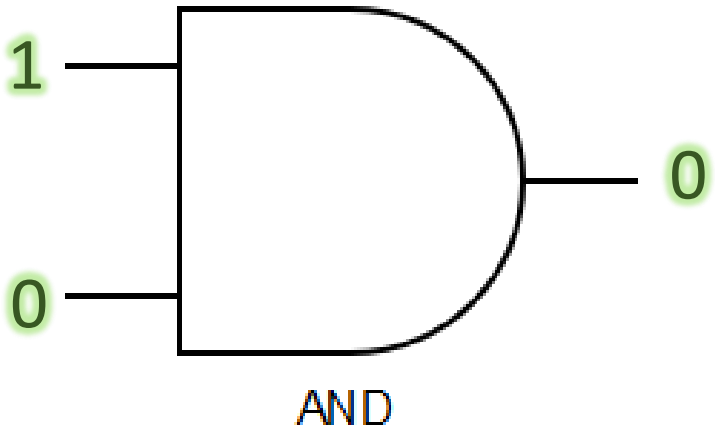

AND

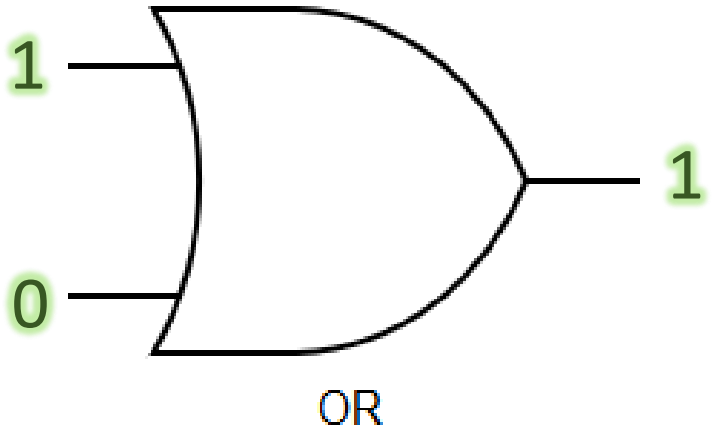

OR

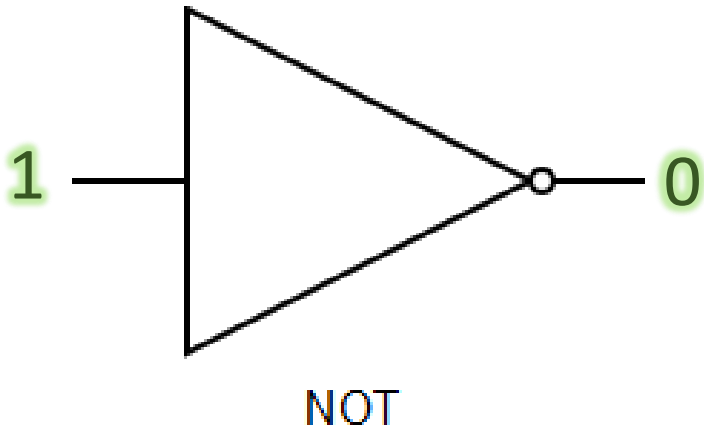

NOT

XOR

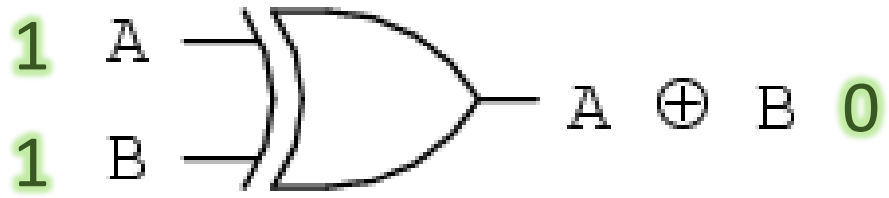

. . . is equivalent to . . .

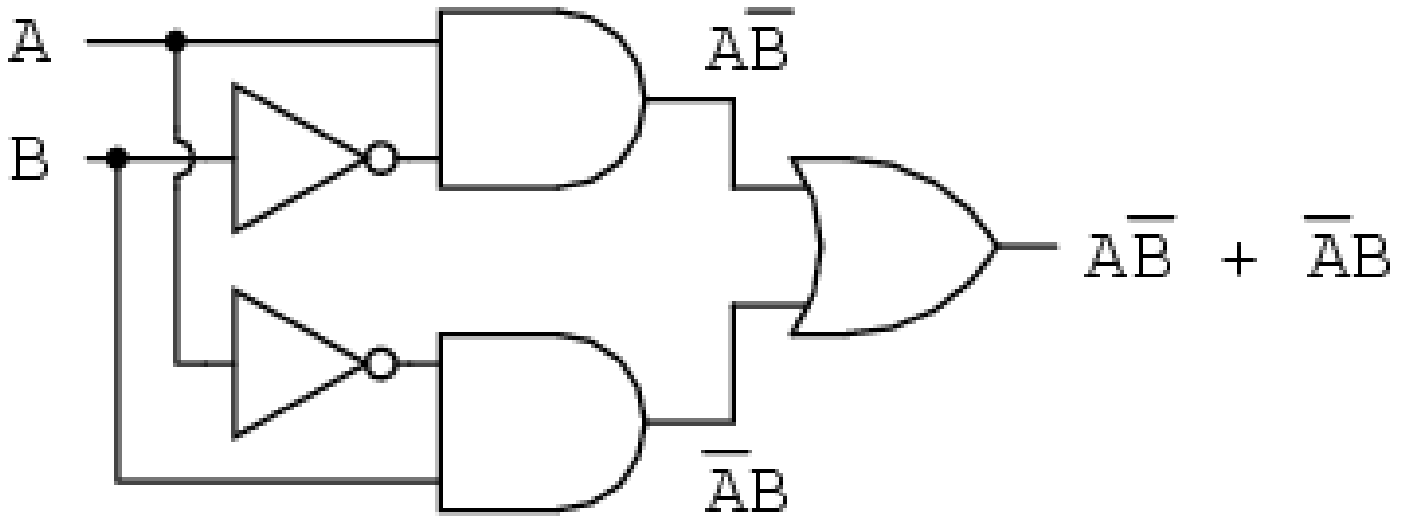

$$A \oplus B = A\overline{B} + \overline{A}B$$


Neural Networks
類神經網路

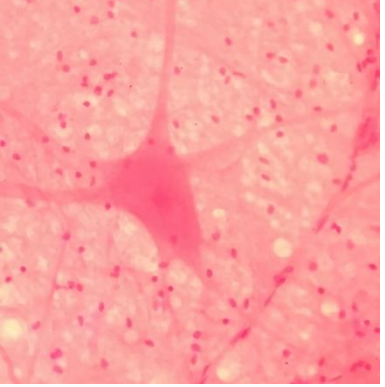

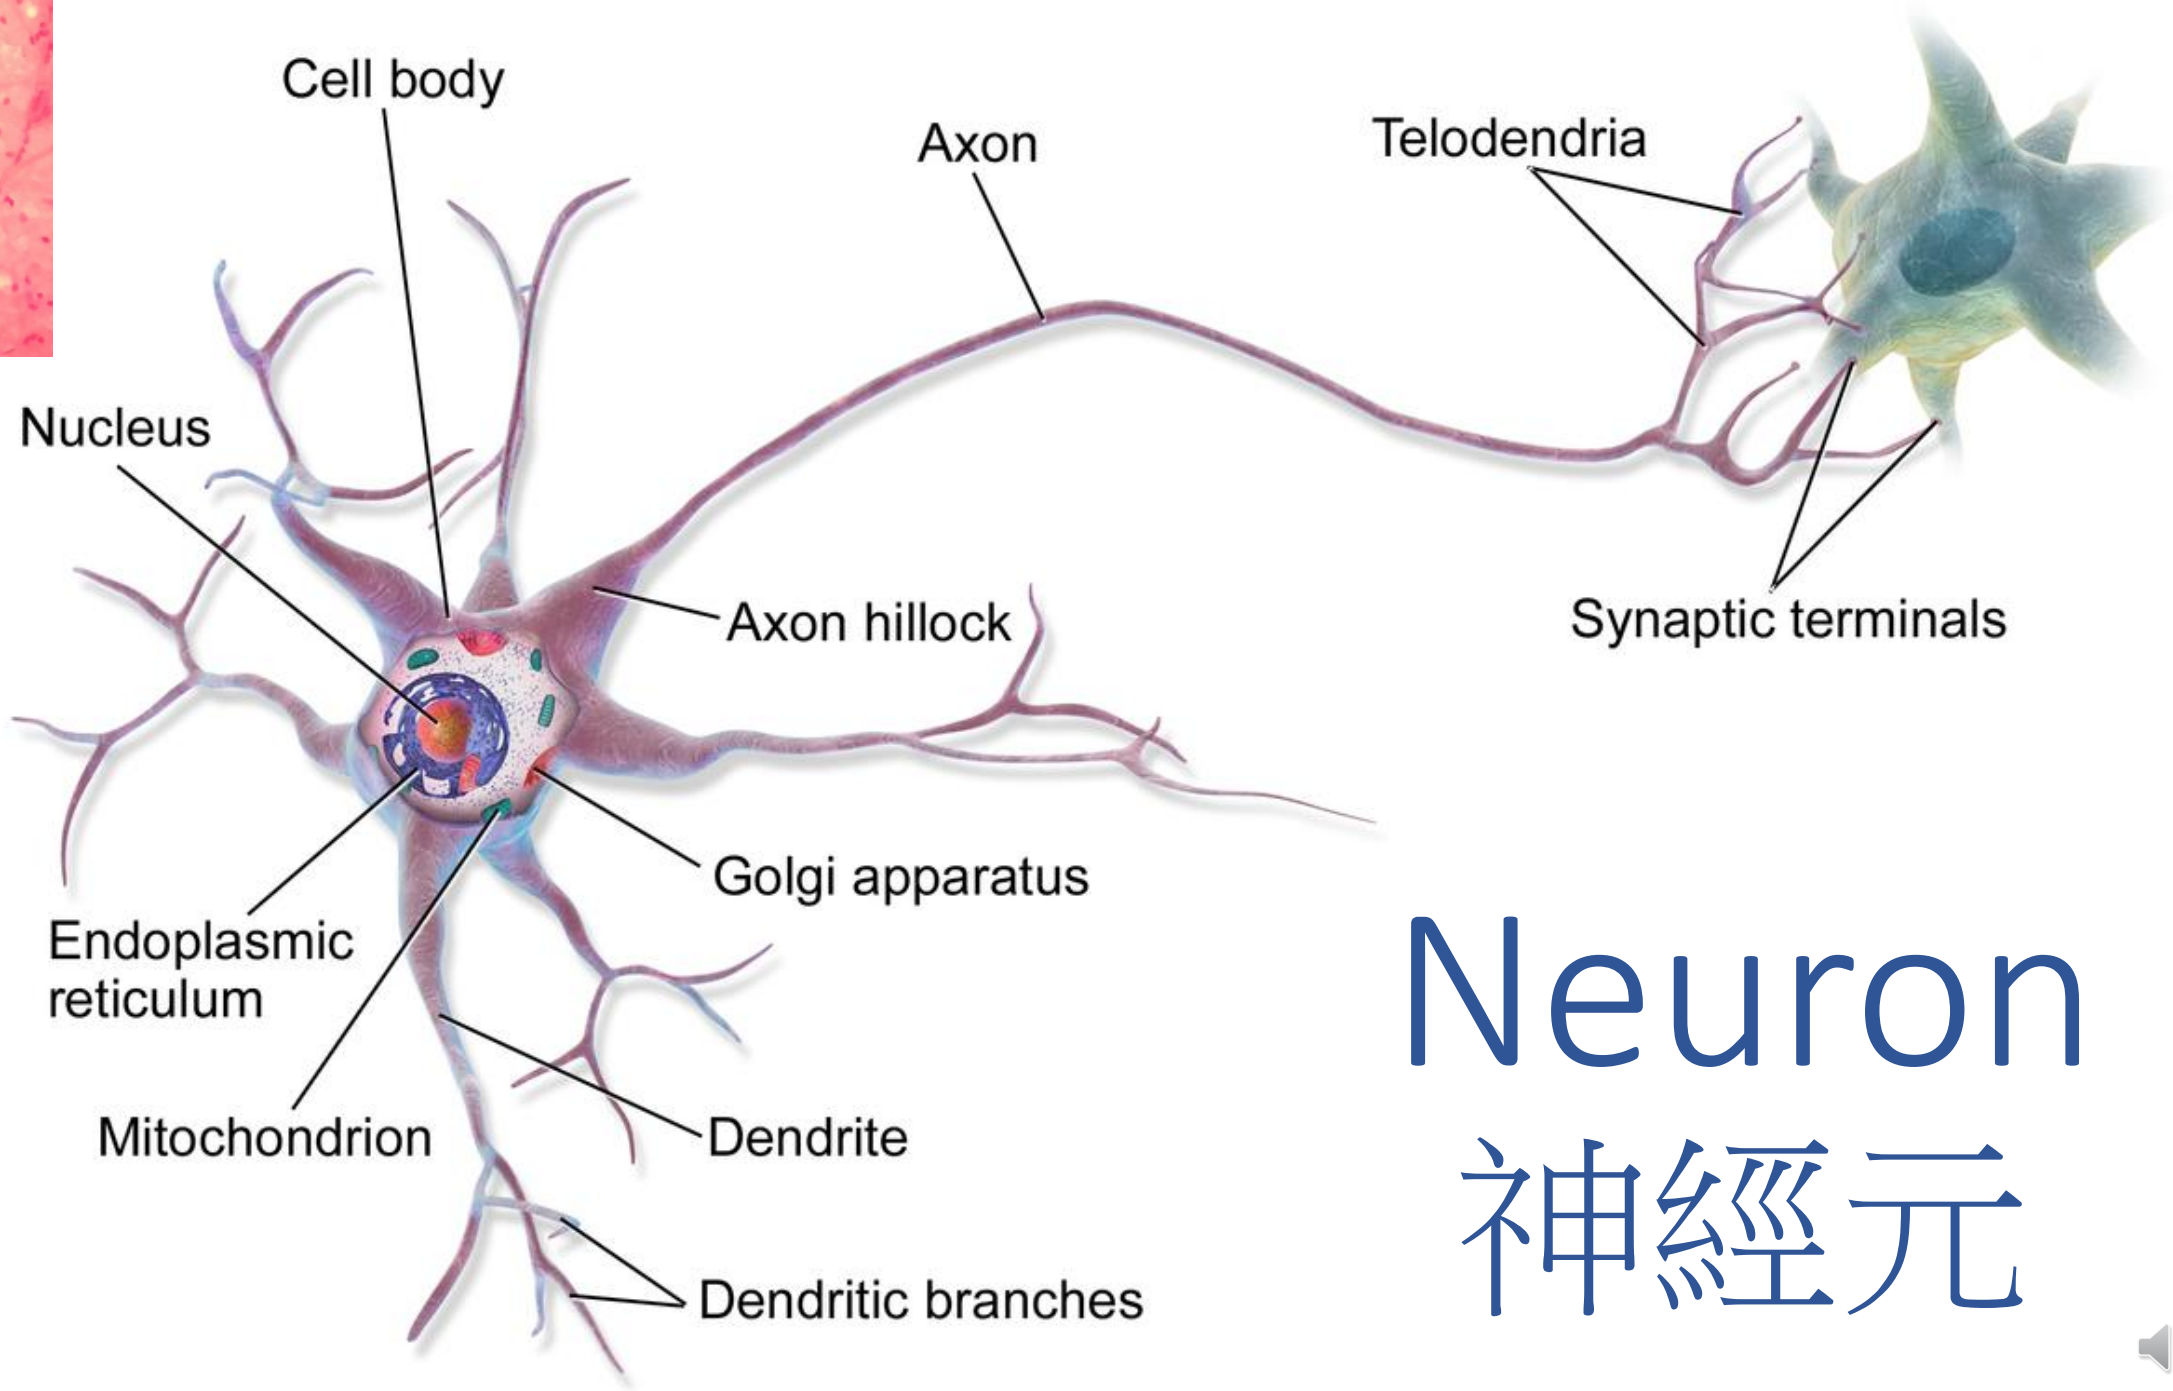
Cell body
Axon
Telodendria
Nucleus
Axon hillock
Synaptic terminals
Golgi apparatus
Endoplasmic reticulum
Mitochondrion
Dendrite
Dendritic branches

Neuron
神經元

Frank Rosenblatt's Perceptron (1957)

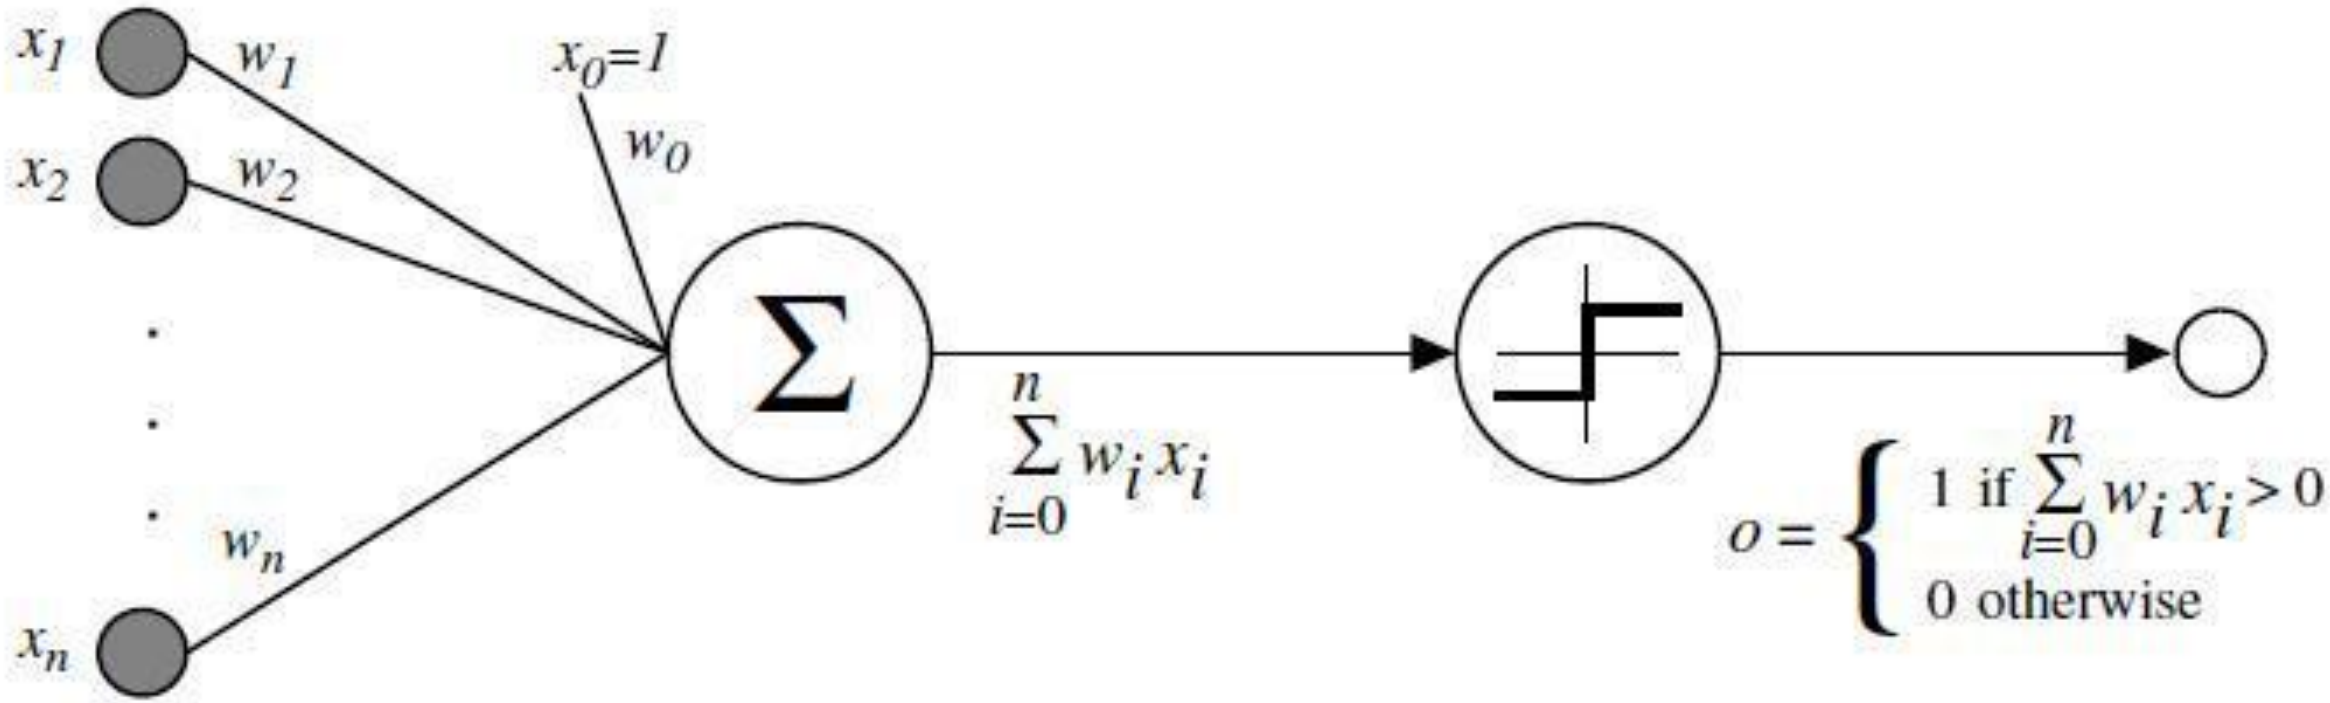

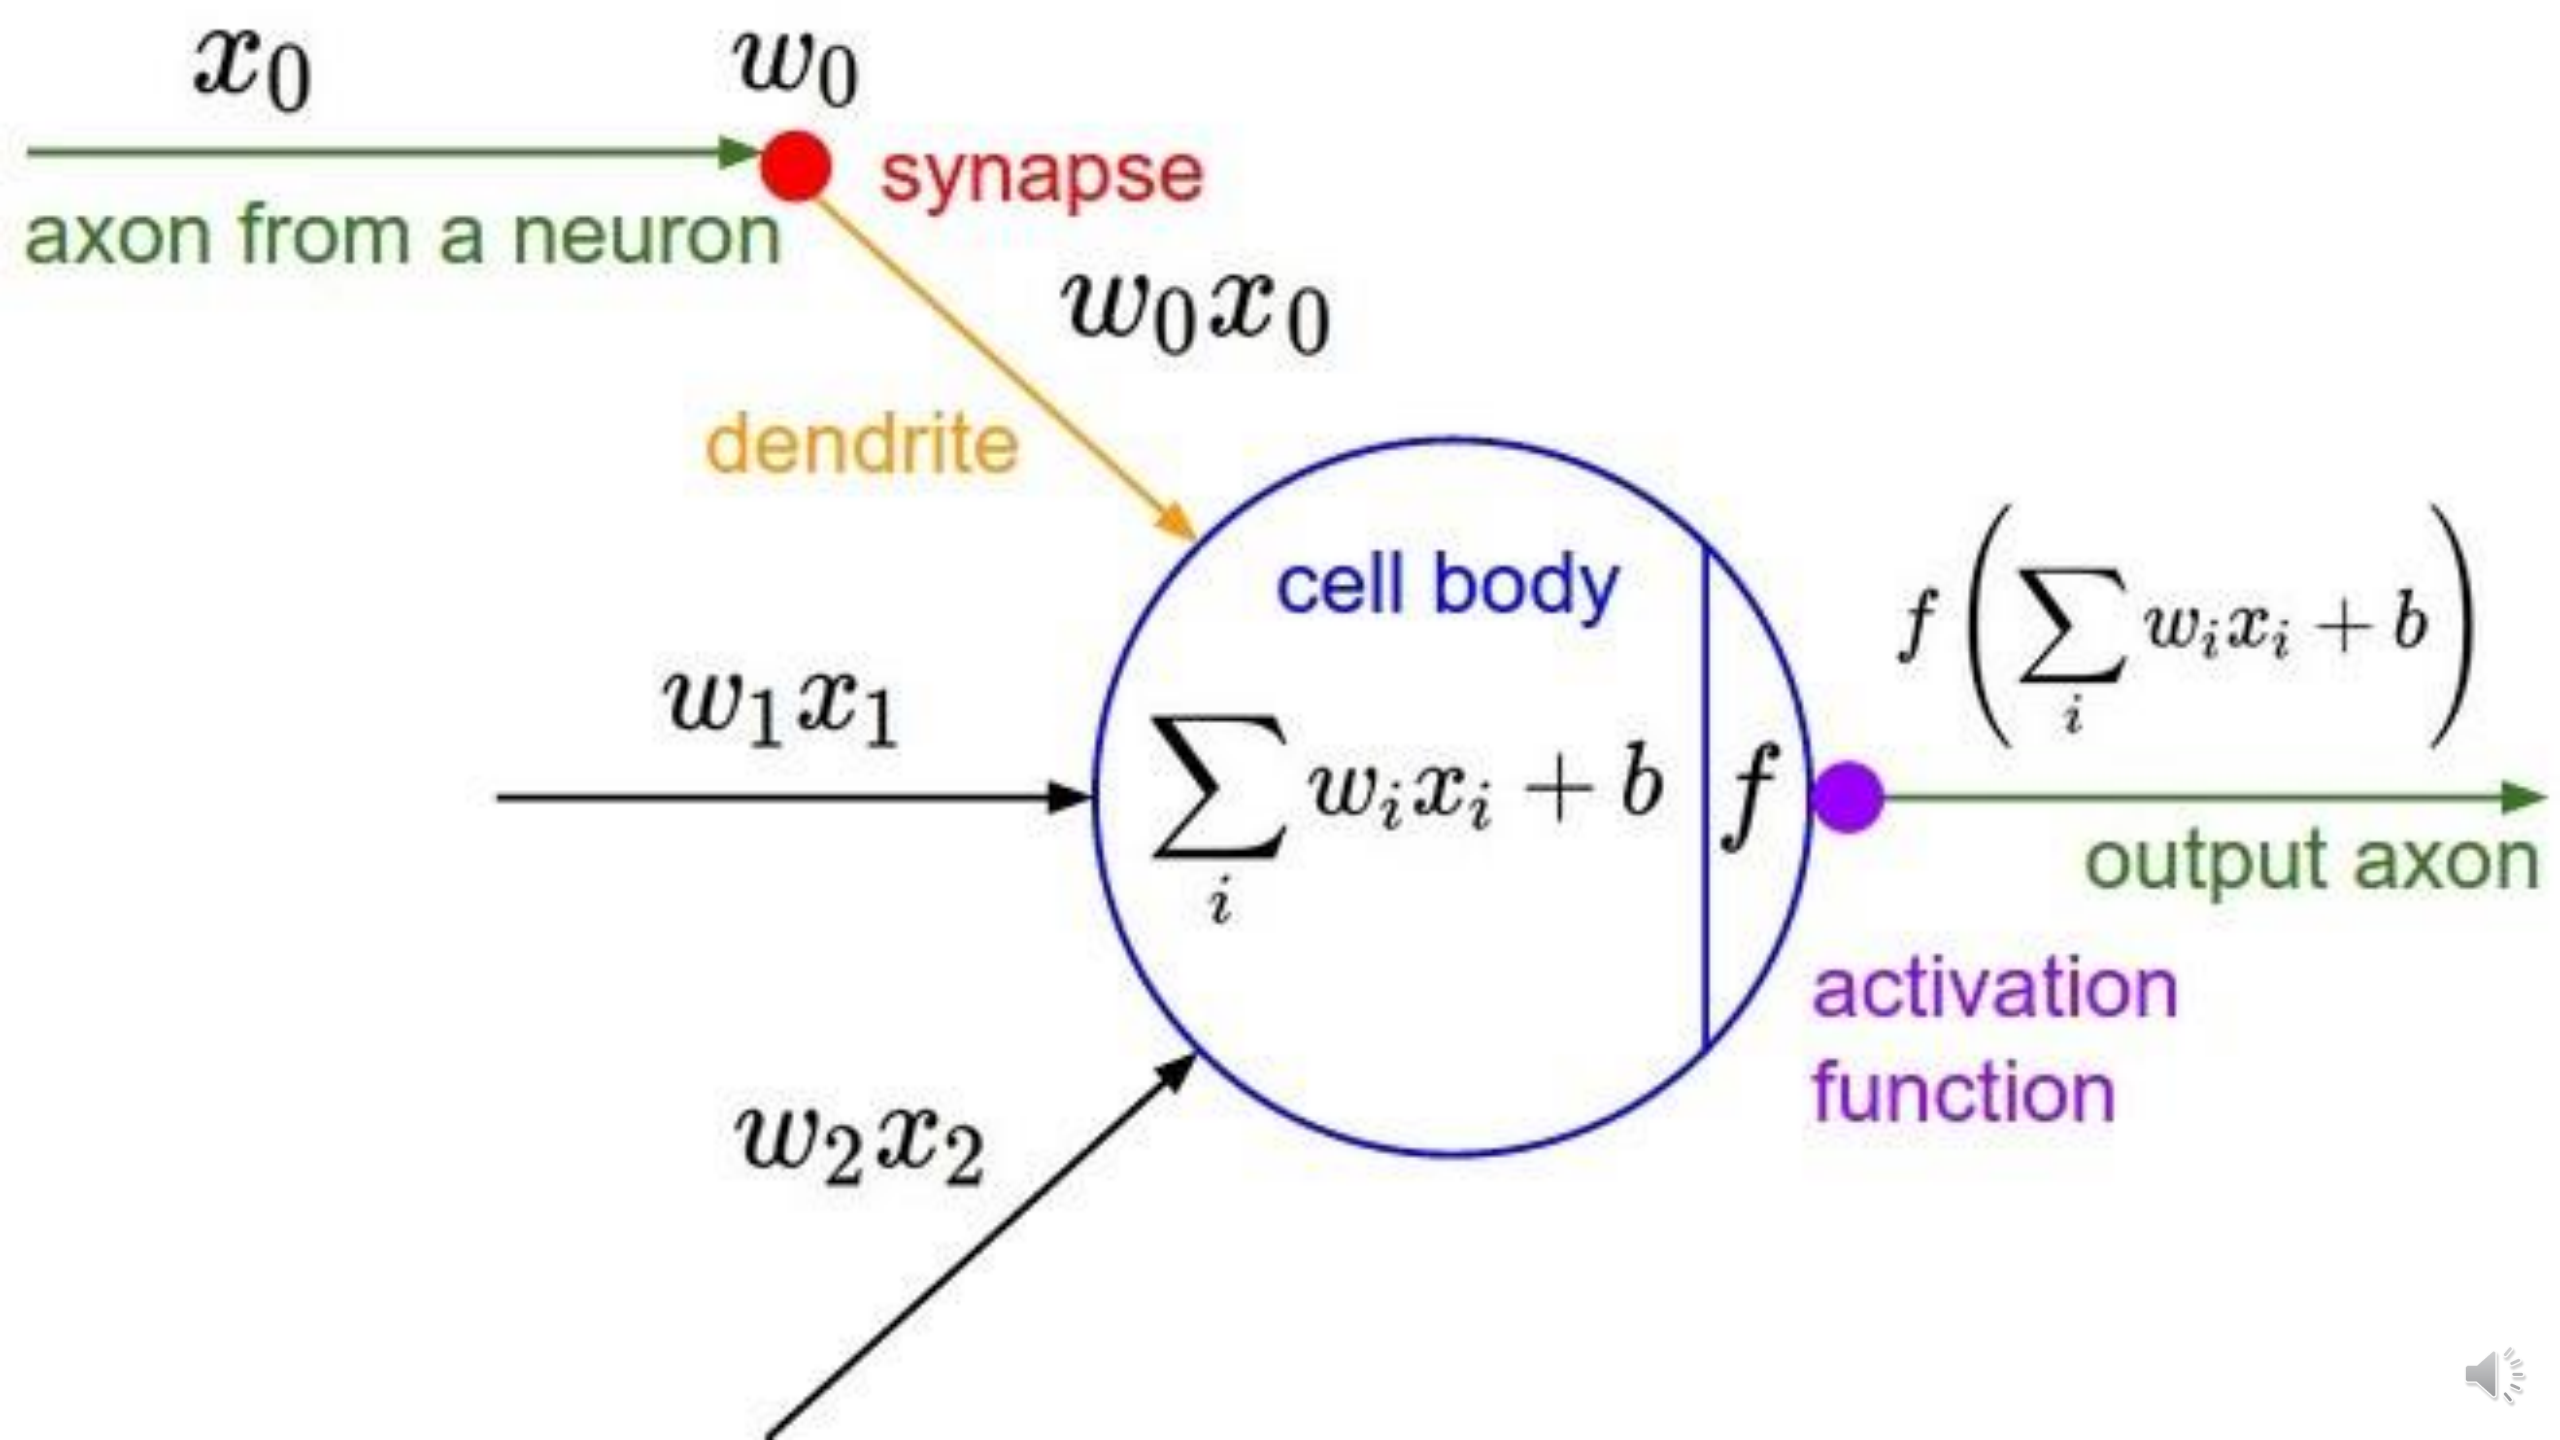
x_0
w_0
synapse
axon from a neuron
$w_0 x_0$
dendrite
cell body
$w_1 x_1$
$\sum_i w_i x_i + b$
f
$f\left(\sum_i w_i x_i + b\right)$
output axon
activation function
$w_2 x_2$

Number of Connections in the Brain
每個人腦中都有個小宇宙!?

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

Synapses (based on 1000 per neuron):
10^14, or 100 trillion, 100000000000000
(每個神經元約有1000突觸)

1969: Perceptrons can't do XOR!

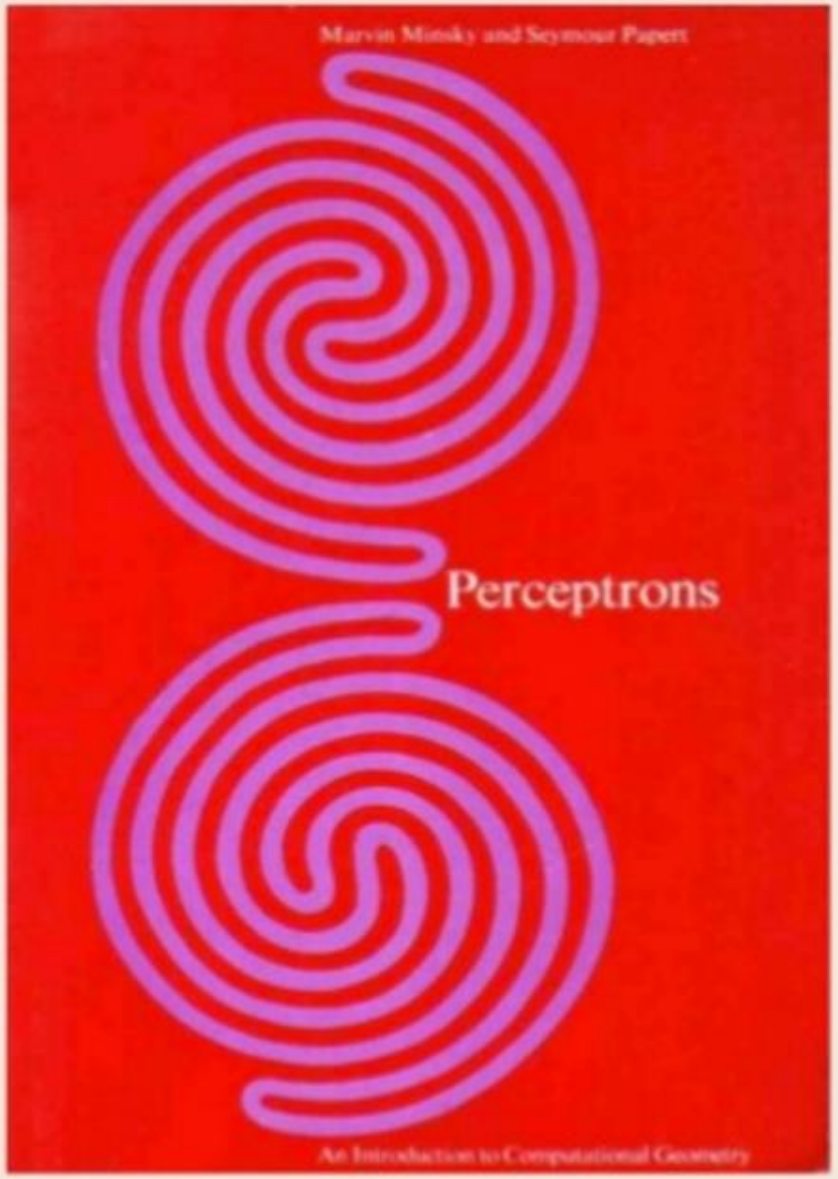

http://www.i-programmer.info/images/stories/BabBag/AI/book.jpg

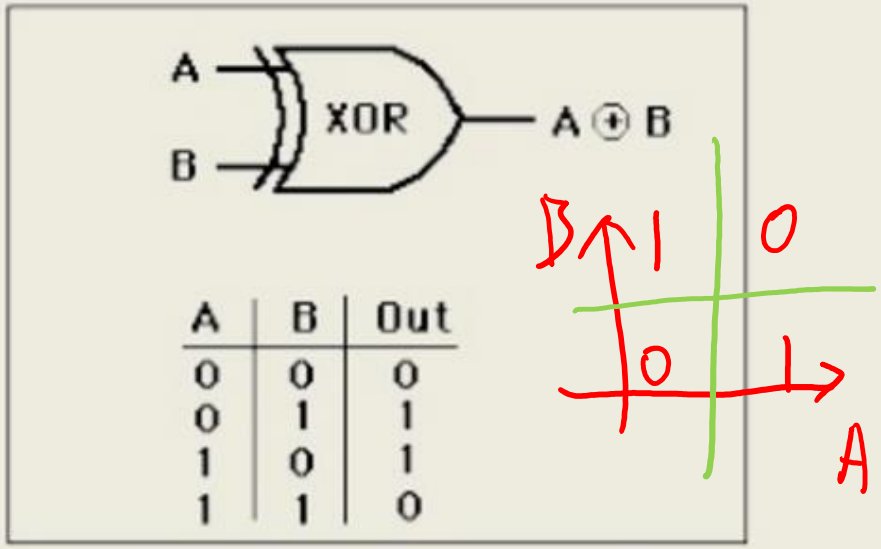

http://hyperphysics.phy-astr.gsu.edu/hbase/electronic/ietron/xor.gif

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)

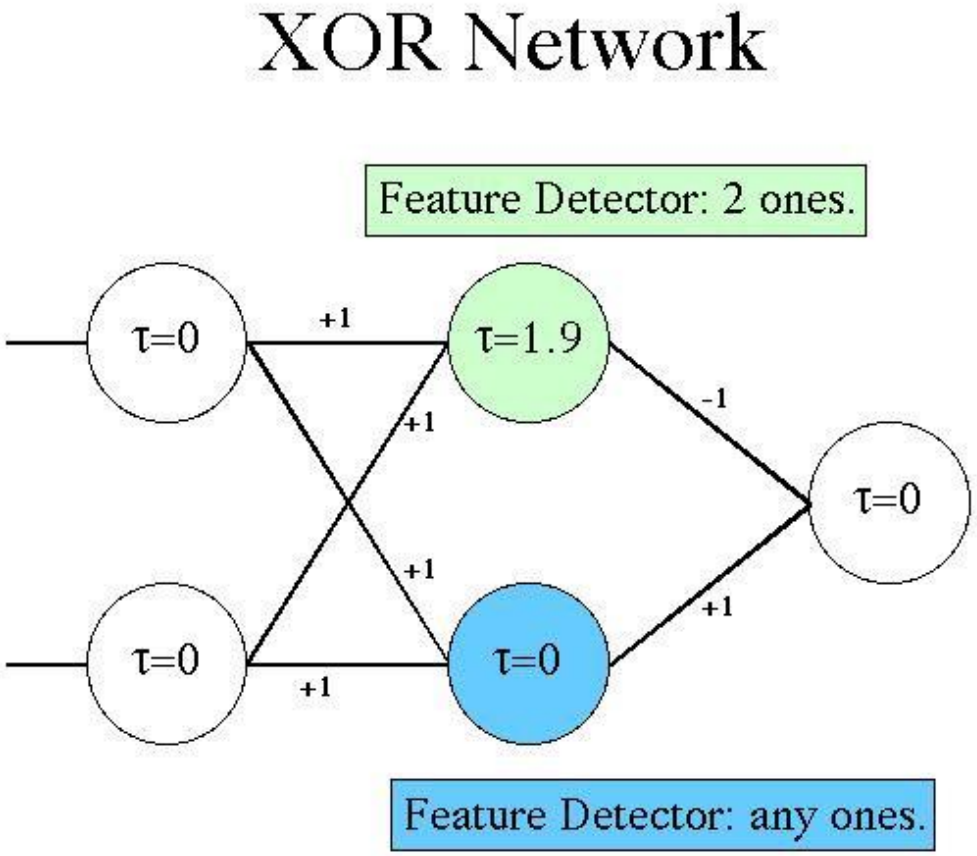

Geoffrey Hinton

Backpropagation (反向傳播)

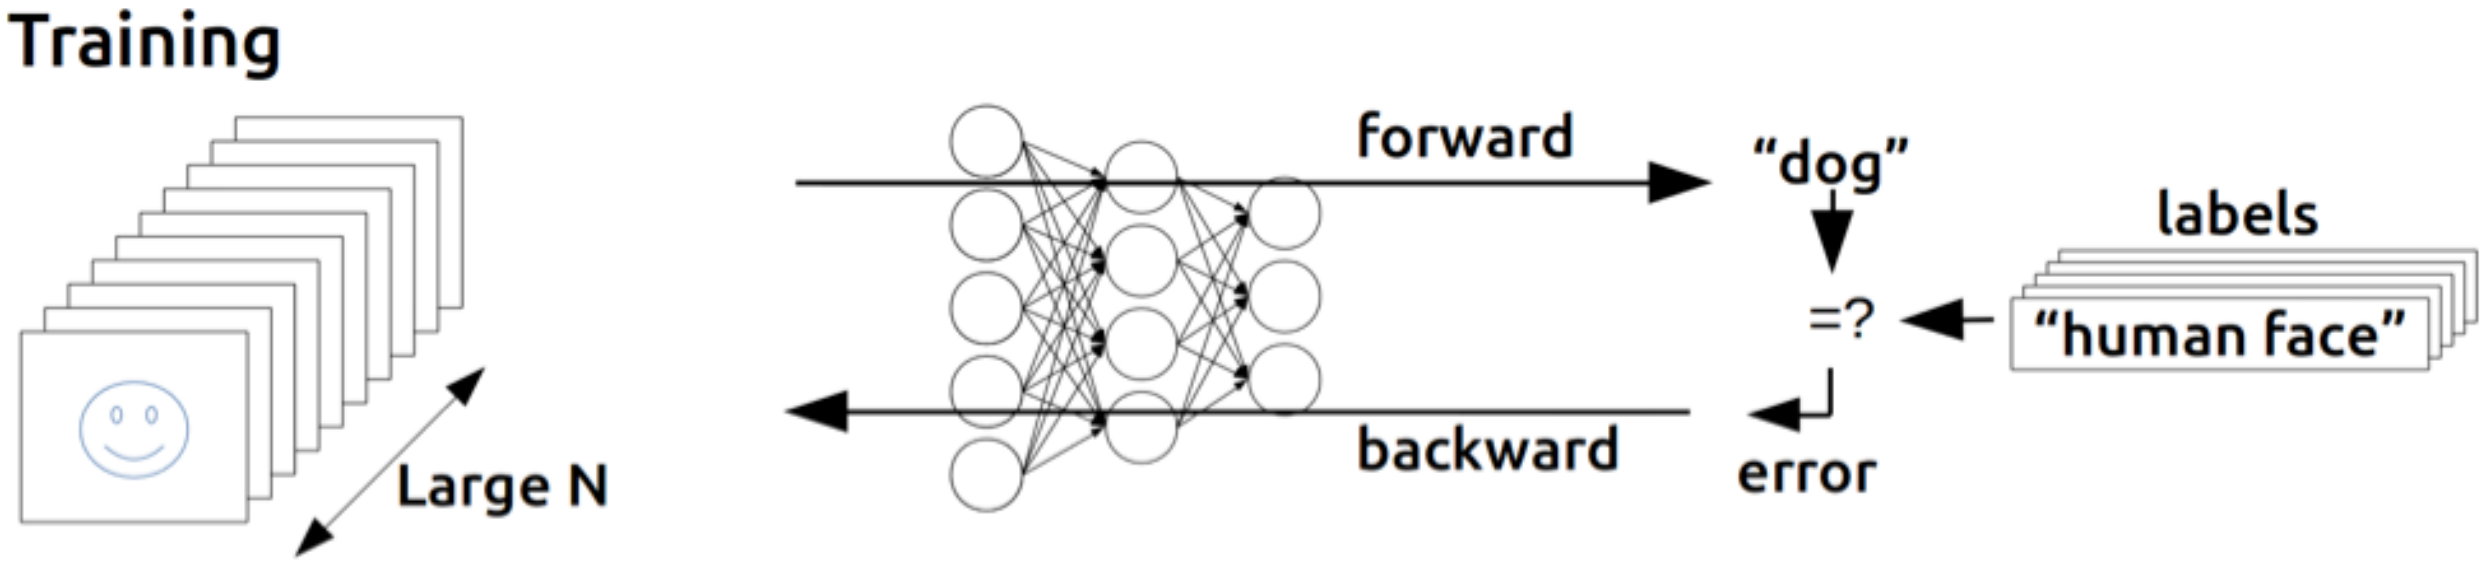

Inference (推理: 使用模型)

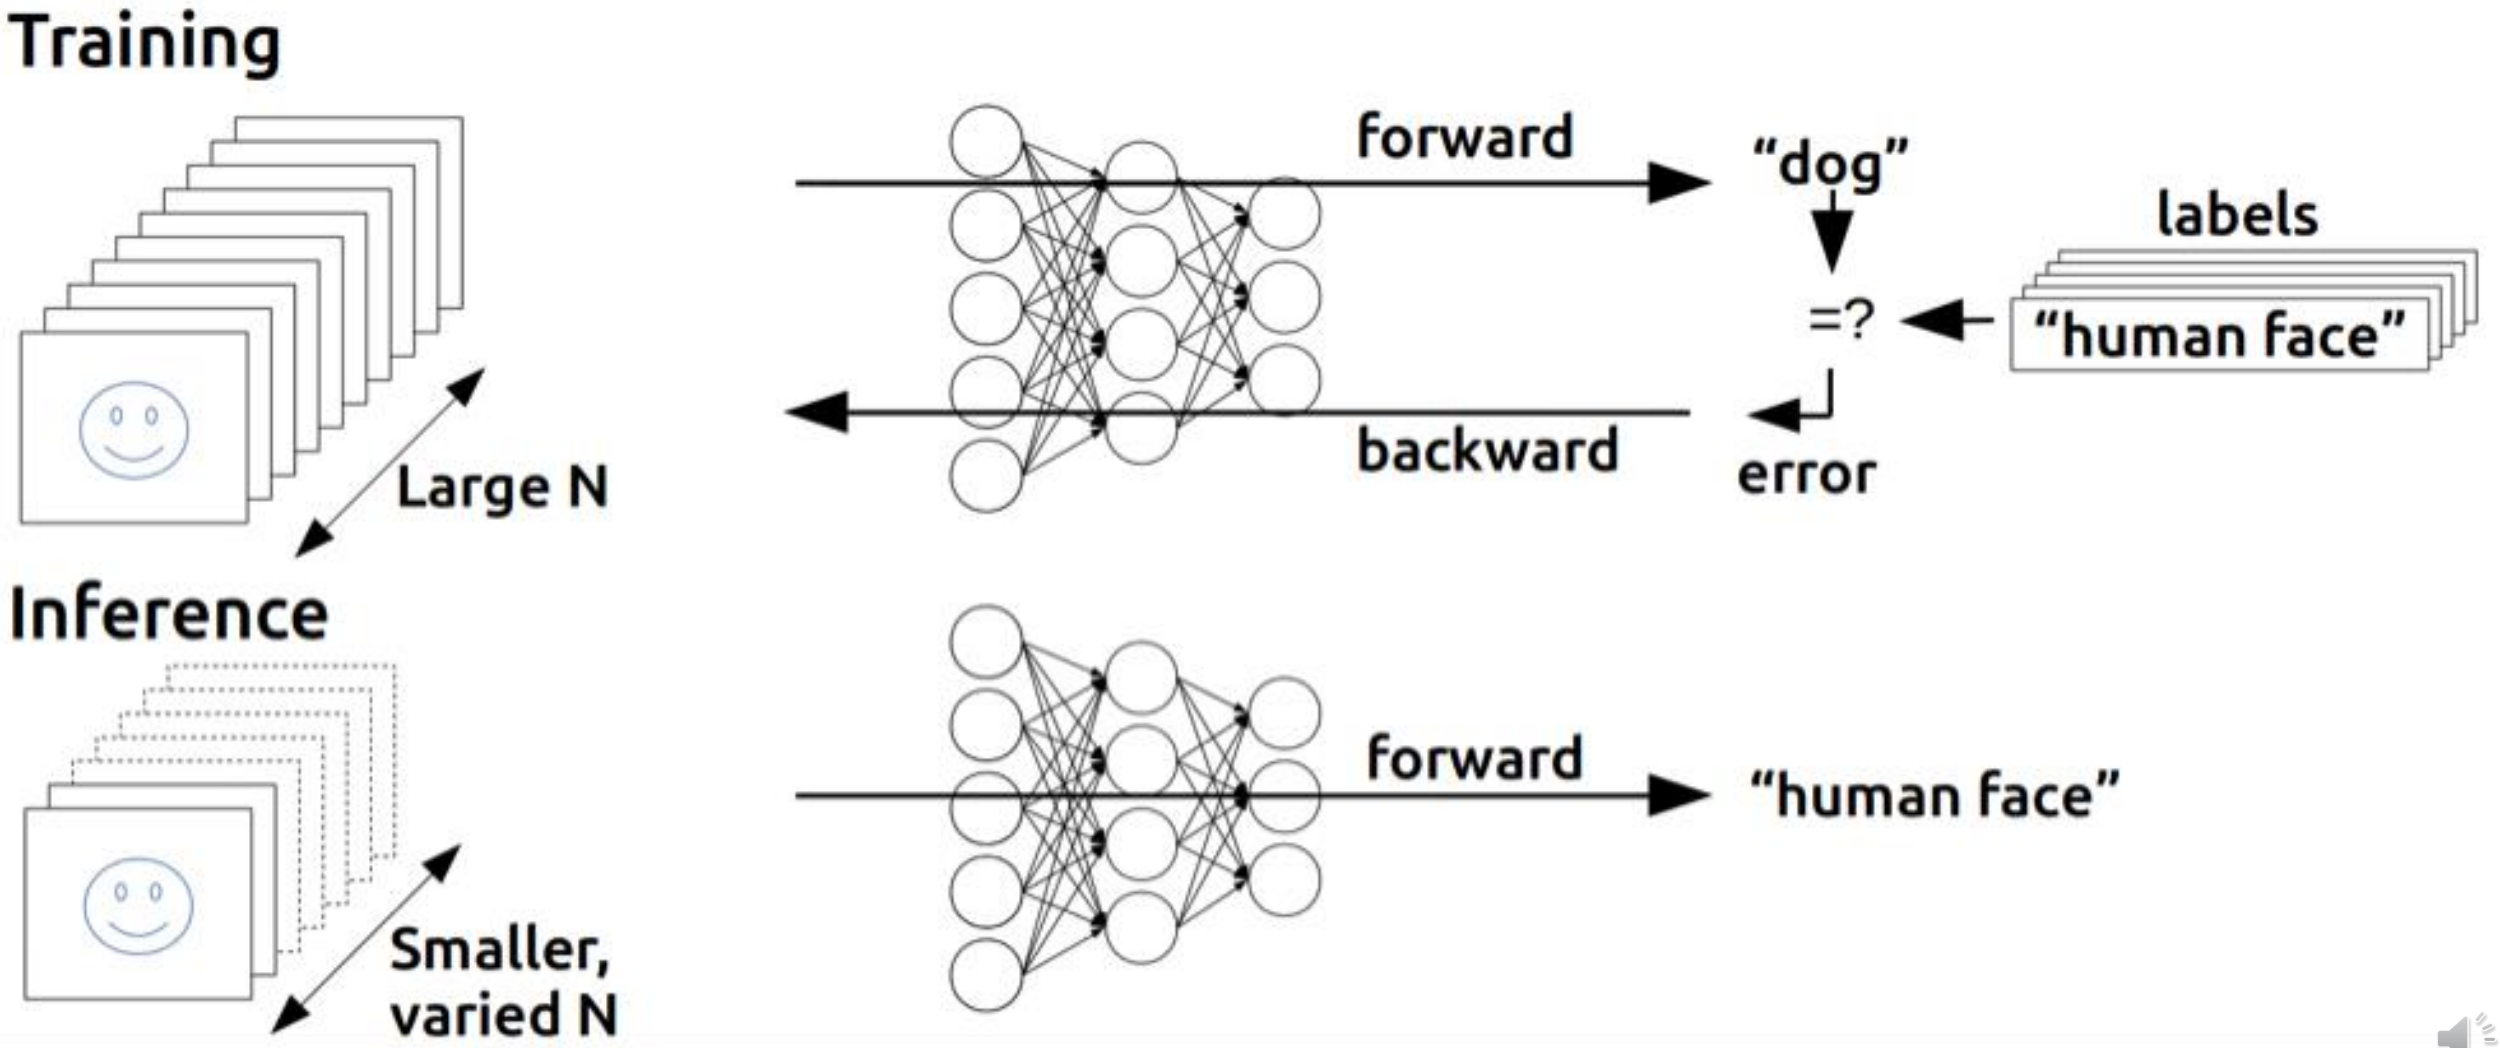

Chain Rule

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

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

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

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

Example: Recognizing Handwritten Digits

- MNIST dataset

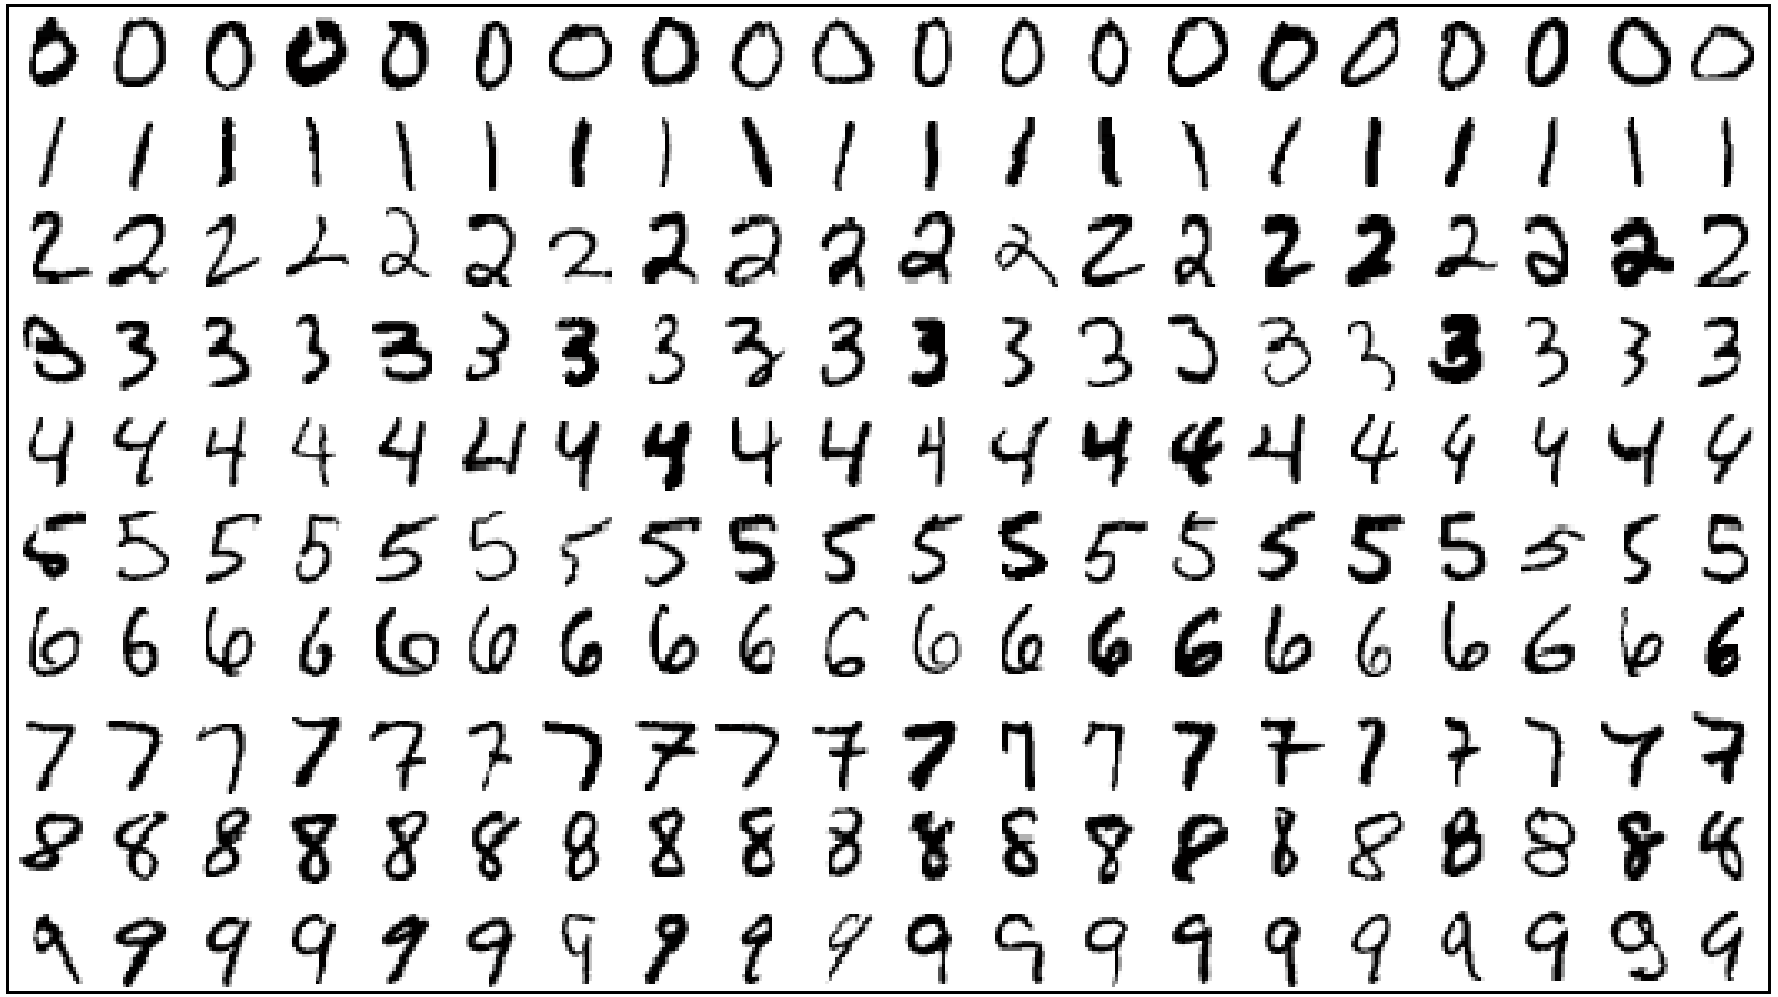

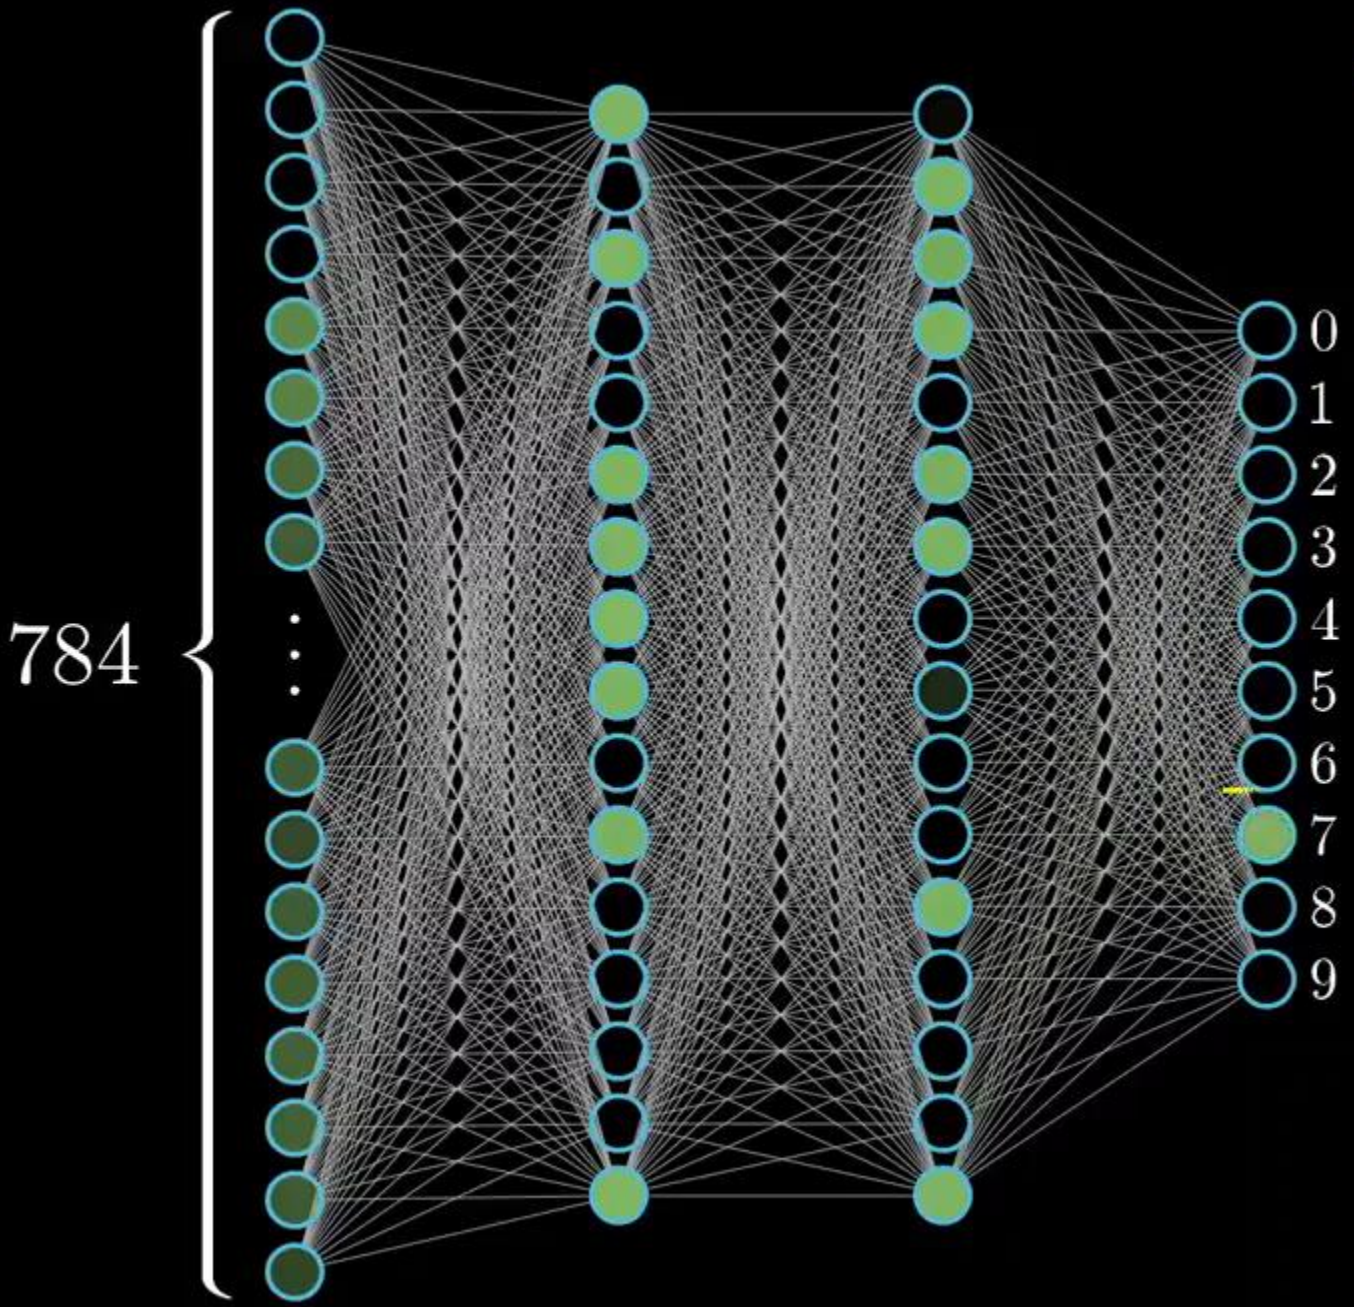

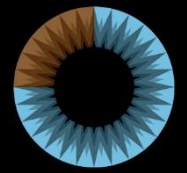

3Blue1Brown

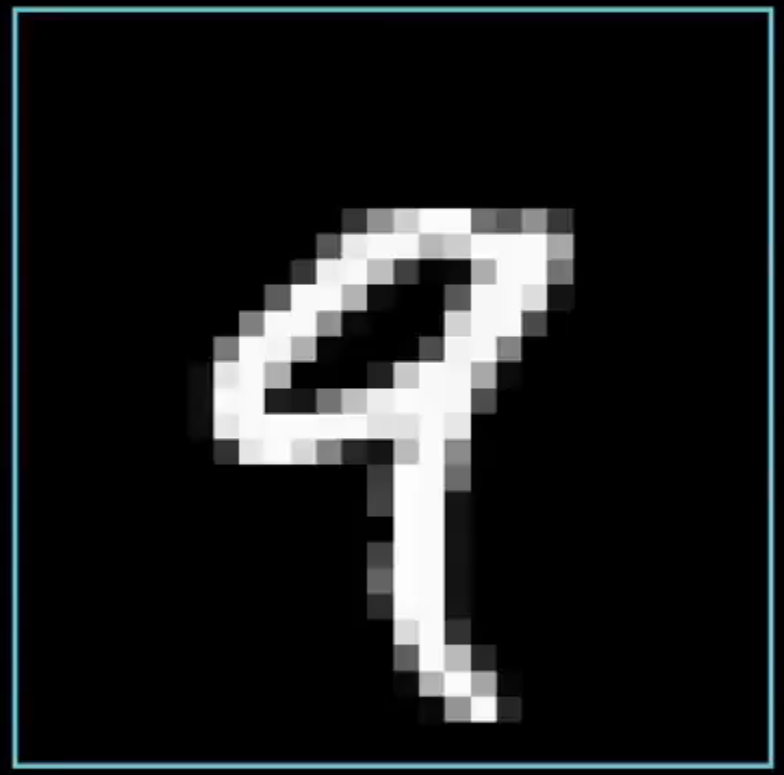

Gradient Descent

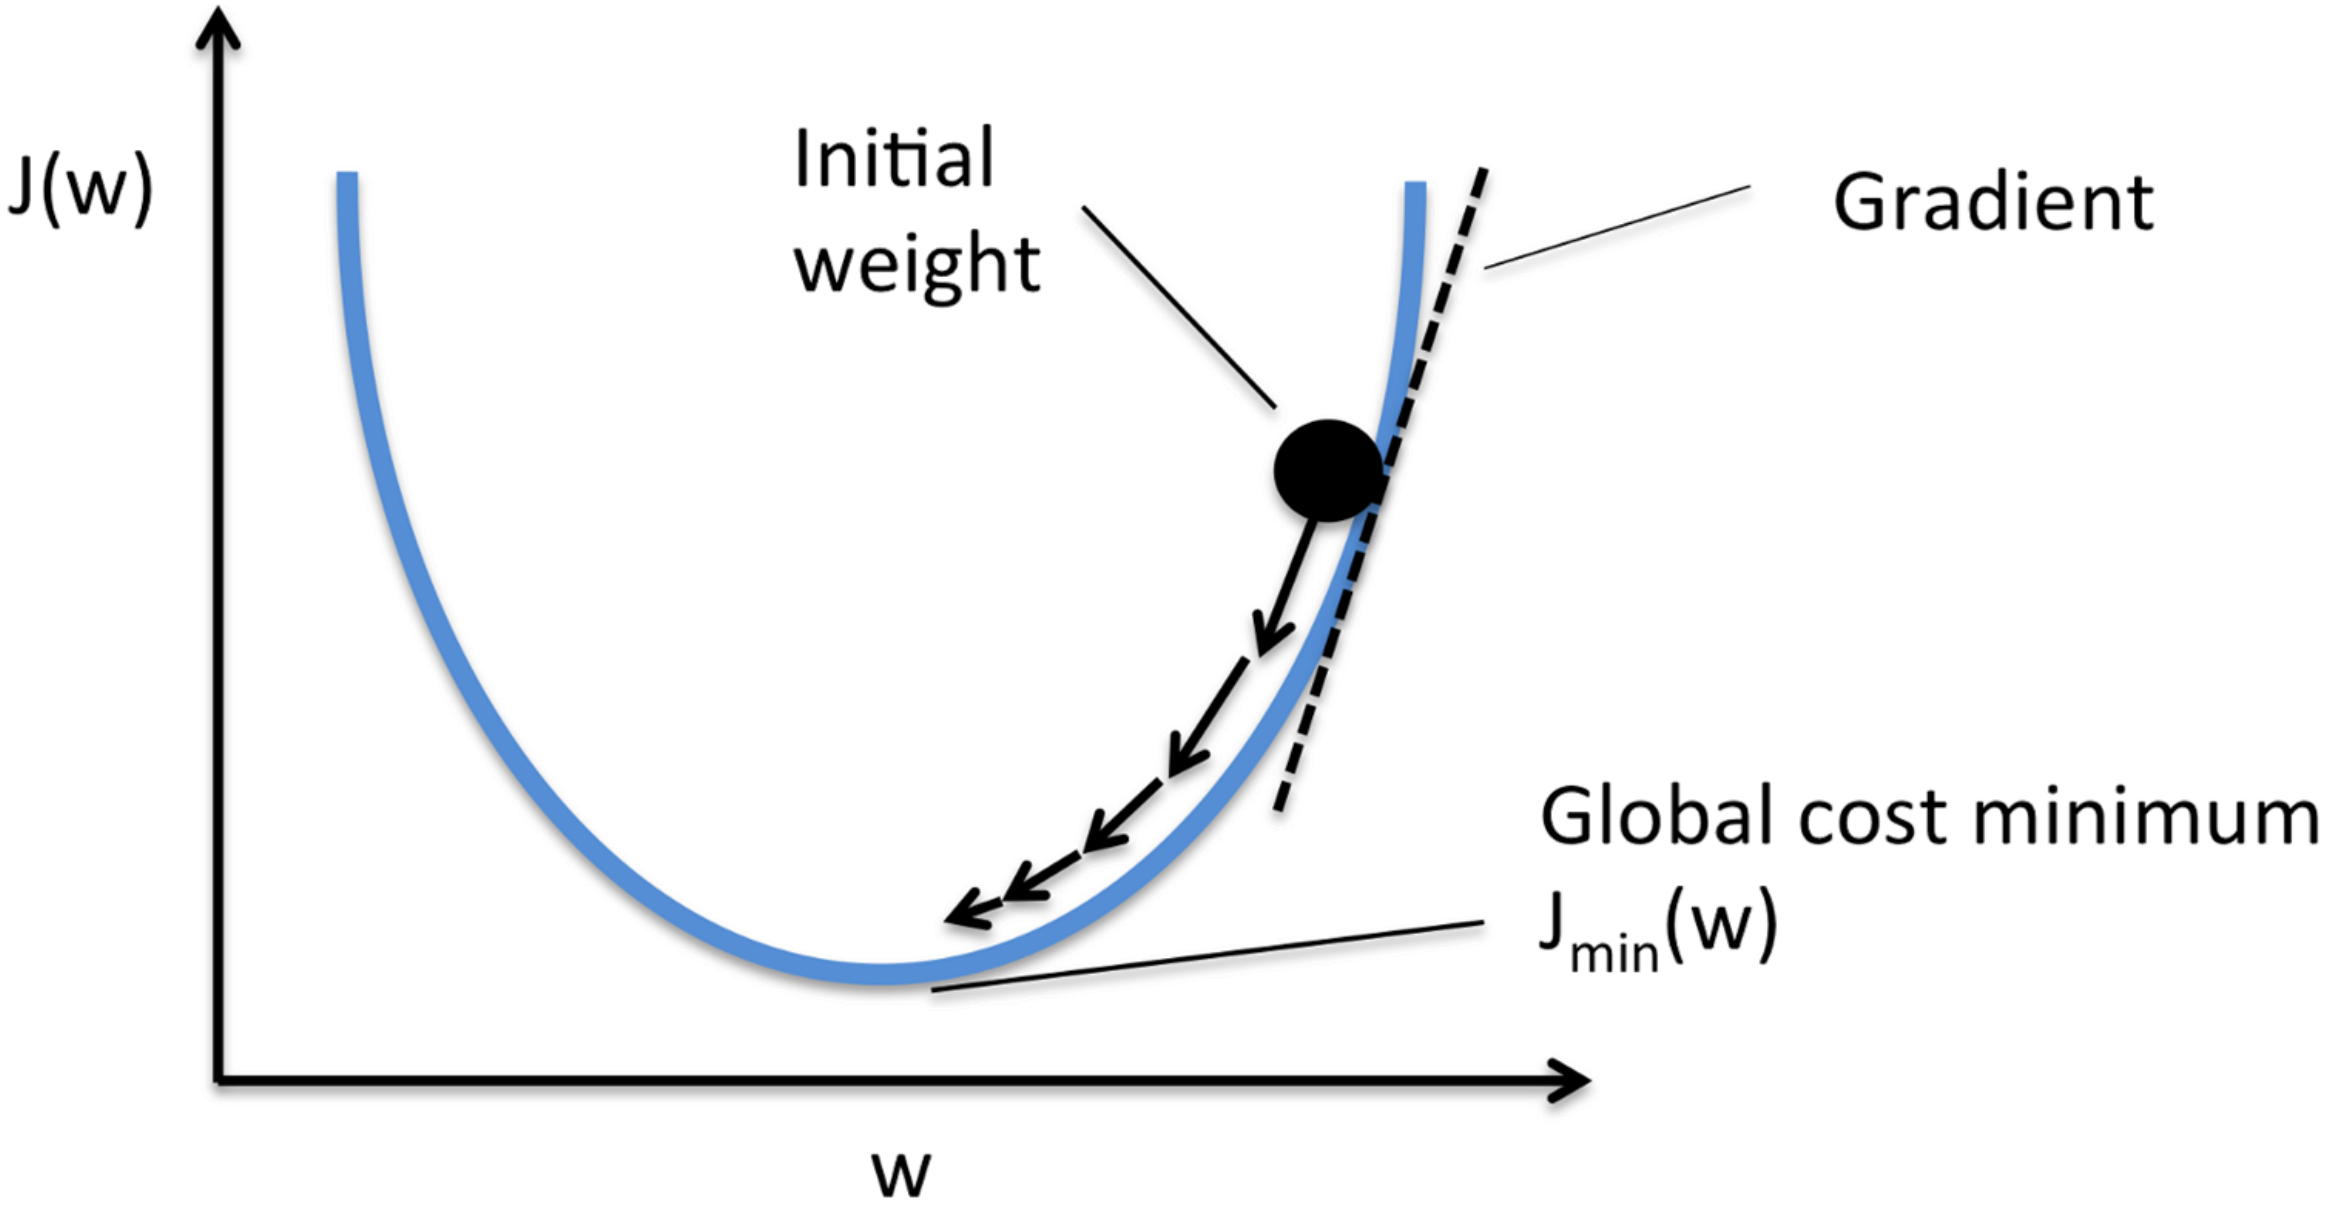

Cost function

$$C(\underbrace{w_1, w_2, \ldots, w_{13{,}002}}_{\text{Weights and biases}})$$

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

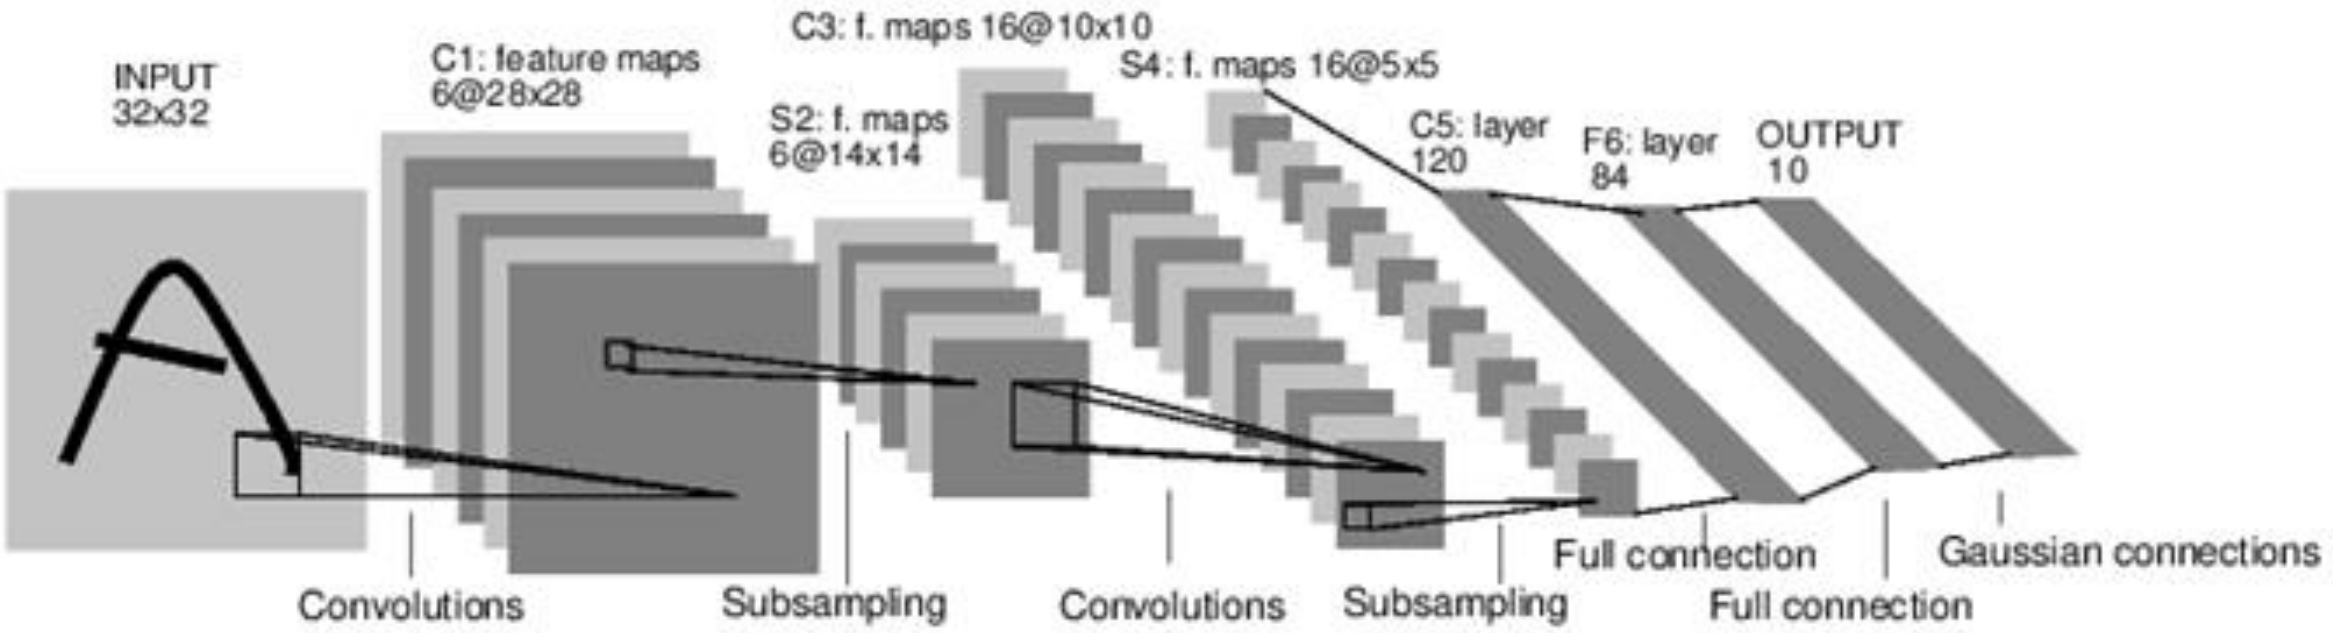

A Full Convolutional Neural Network (LeNet)


INTERNET

ImageNet Large Scale Visual Object Recognition Challenge (ILSVRC)

ImageNet 影像分類競賽

- 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

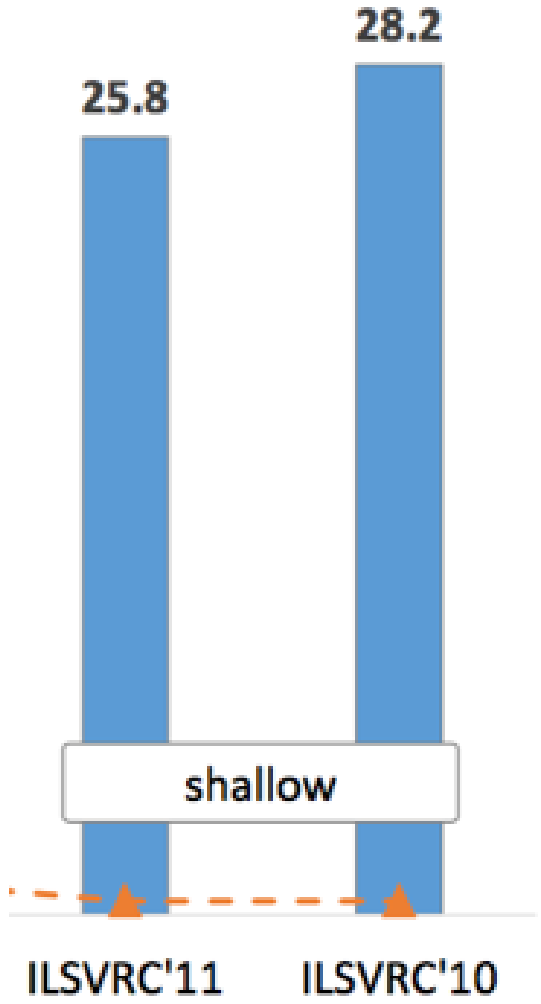

Convolutional Neural Network (AlexNet)

- Alex Krizhevsky, Geoffery Hinton et al., 2012

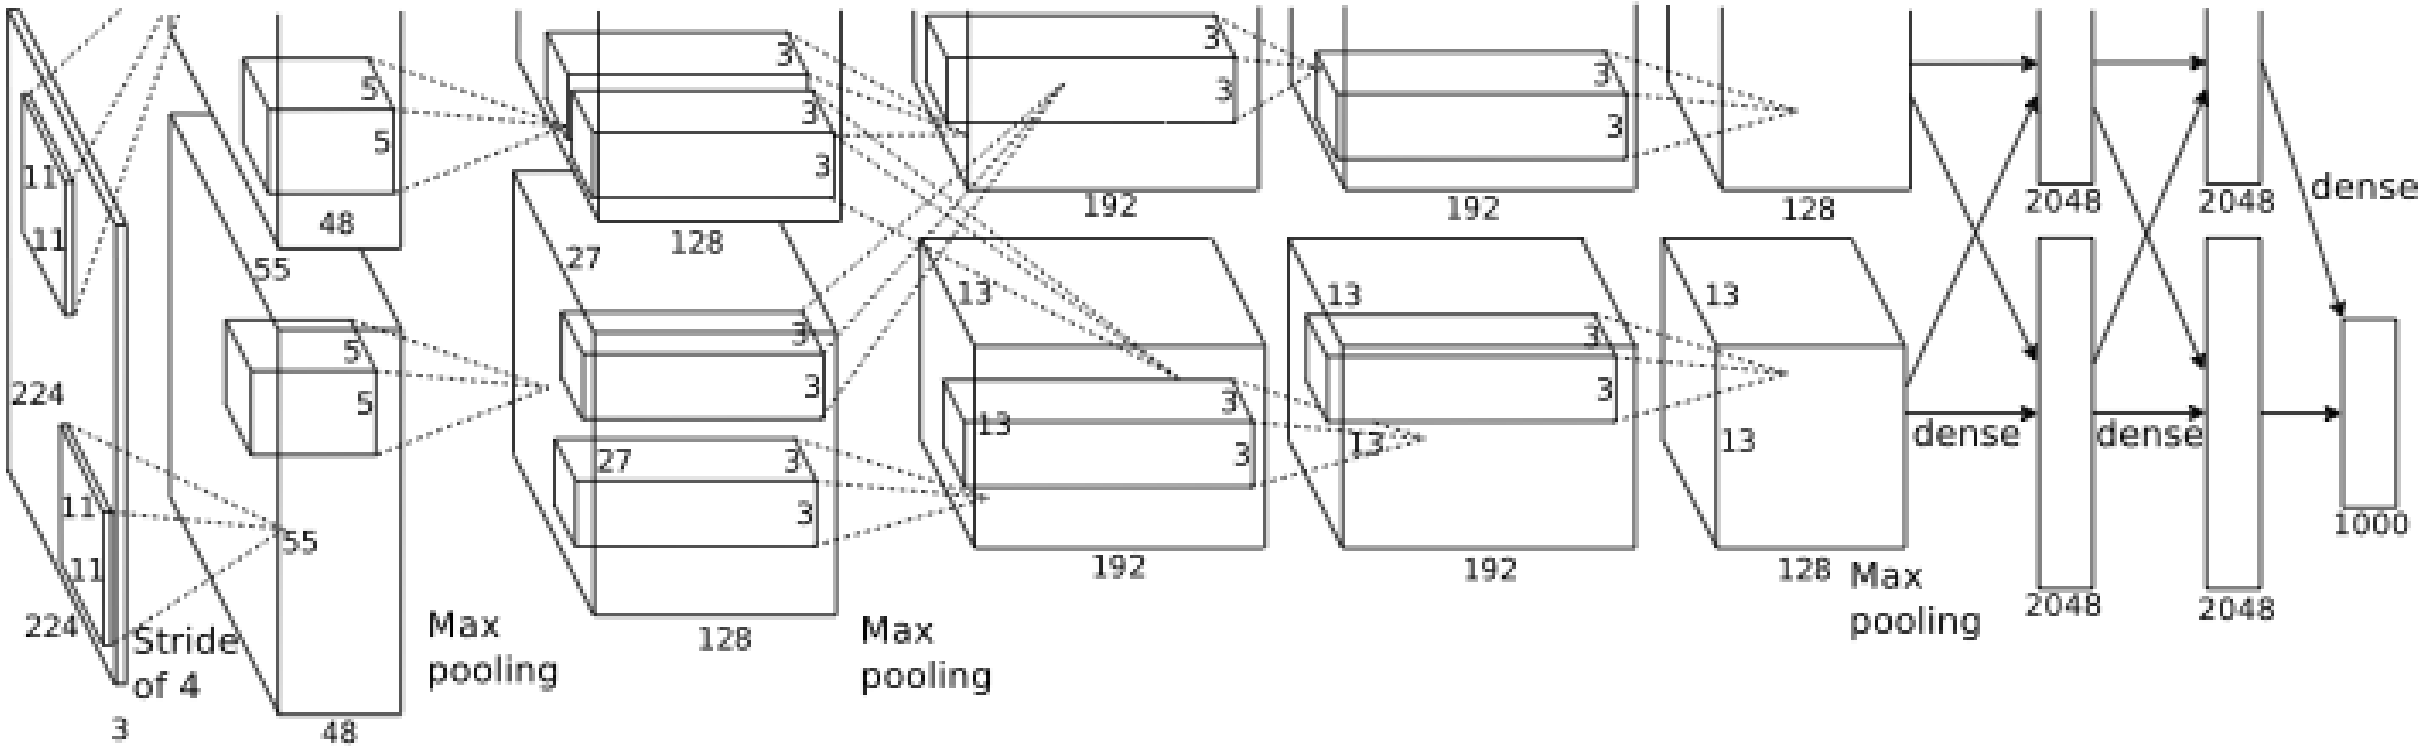

Error Rate on ImageNet Challenge

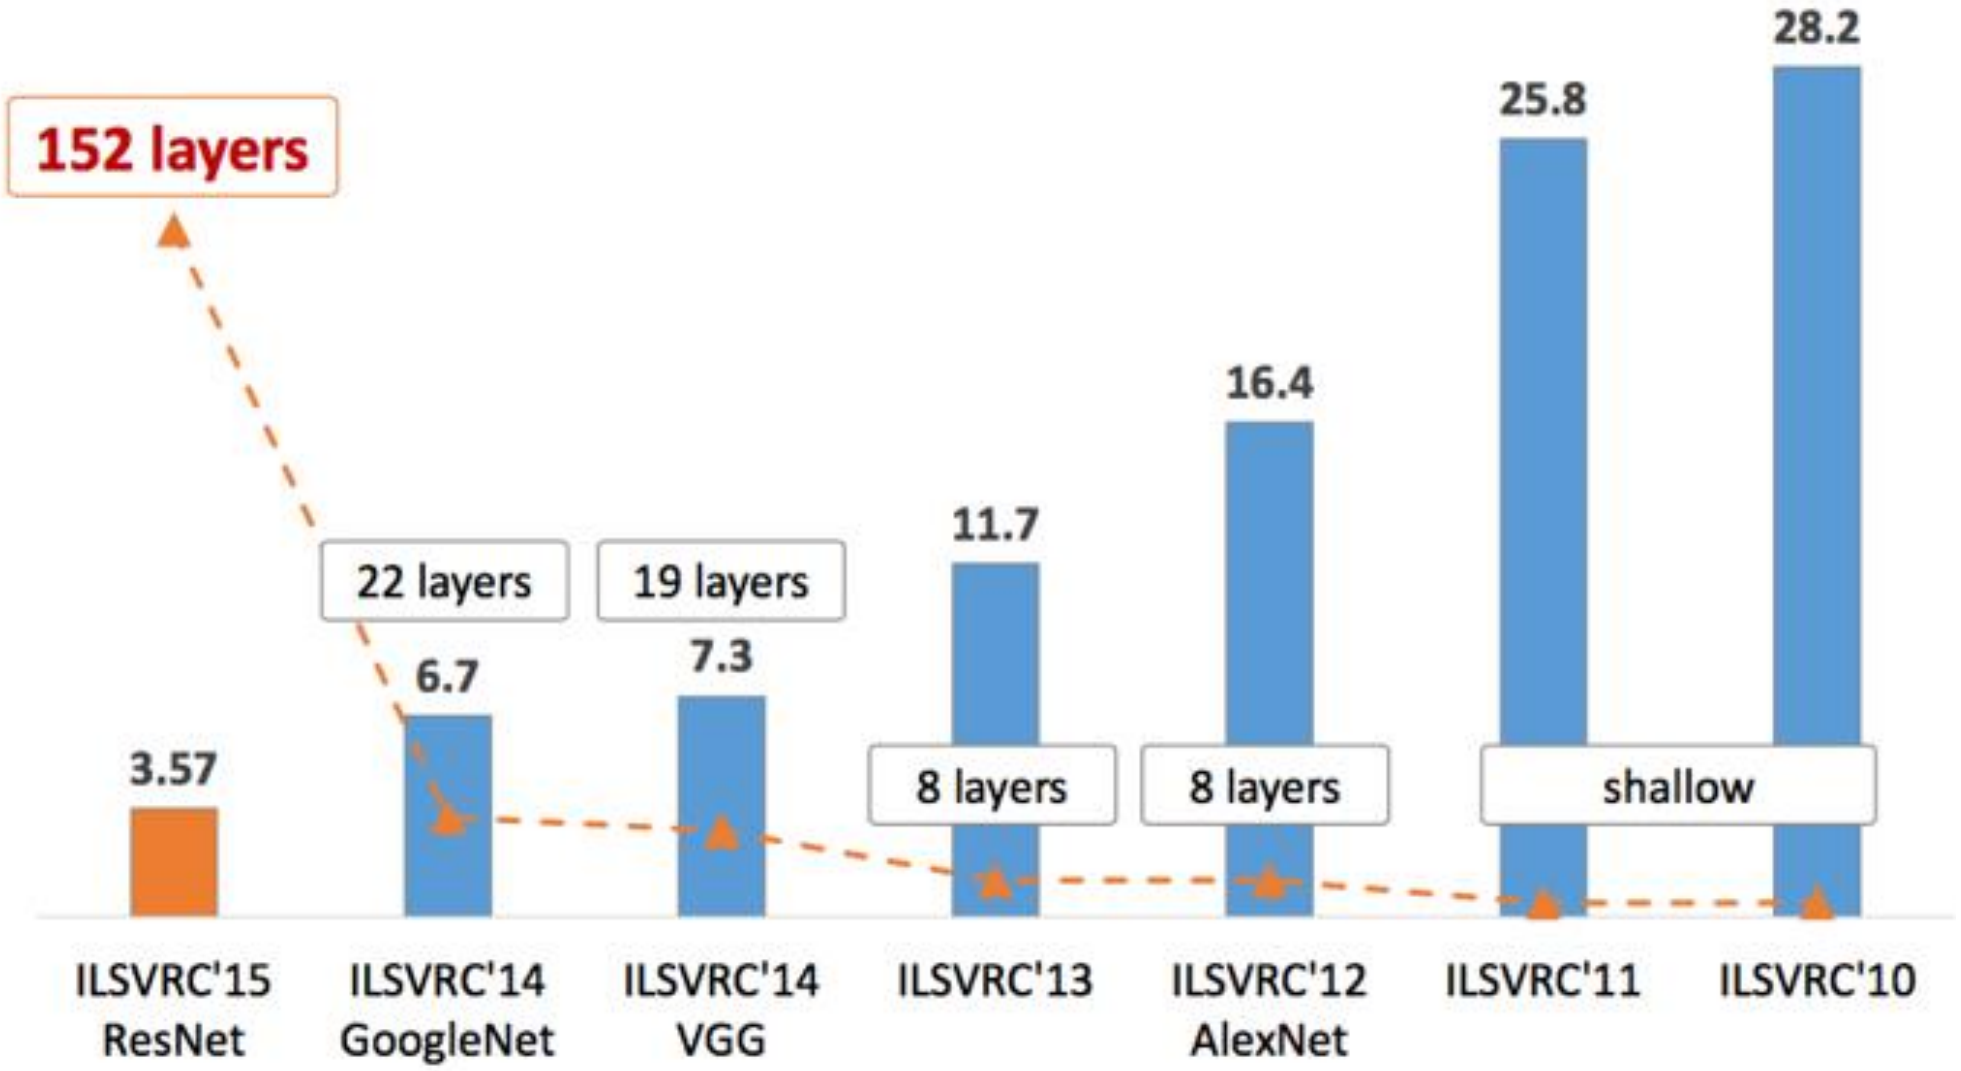

LeNet

AlexNet

VGG

GoogLeNet

Inception V3

Inception BN


WE NEED TO GO
DEEPER
memegenerator.net

AlexNet

VGG

Network in Network

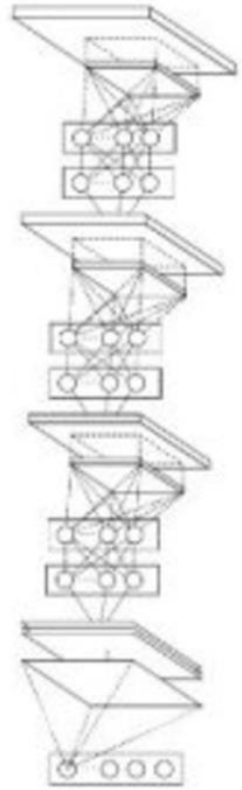

GoogLeNet

ResNet

https://memoiry.me/2017/04/22/deep-learning-series-2-typical-CNN-net/

CNN Comparison

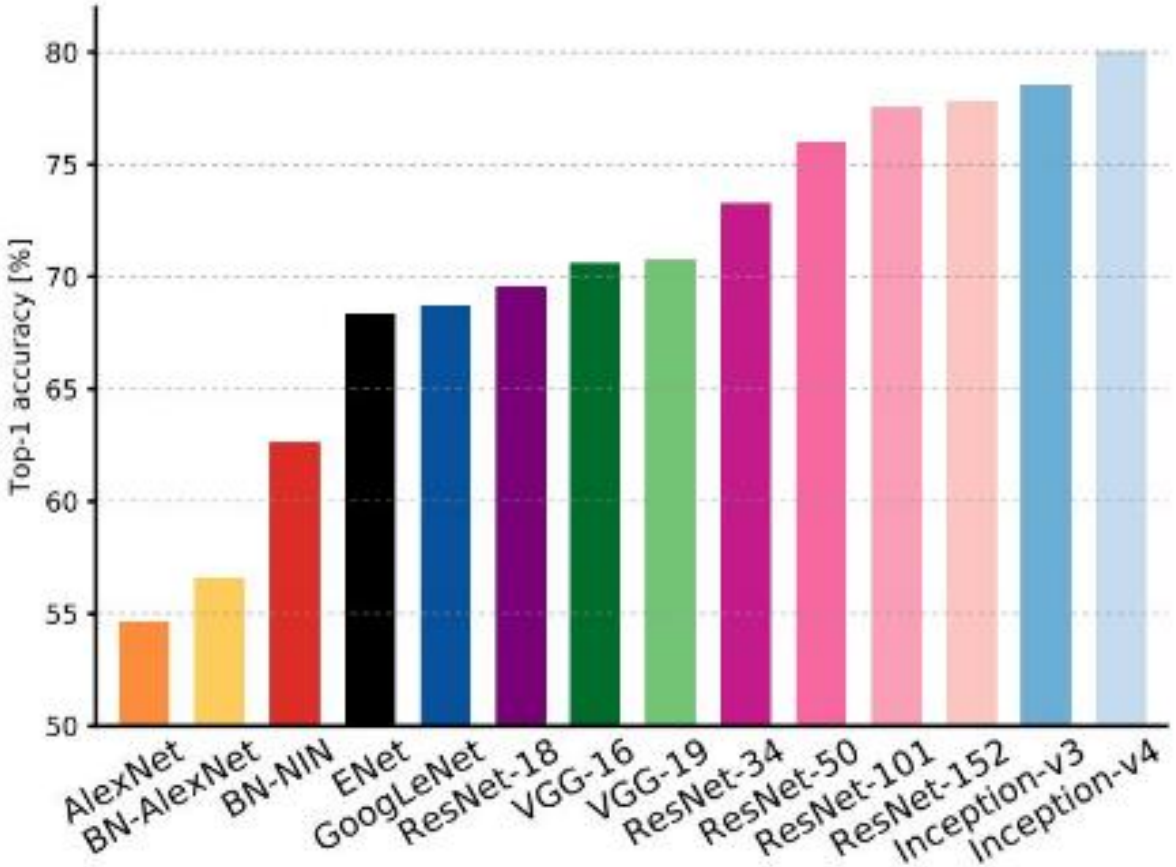

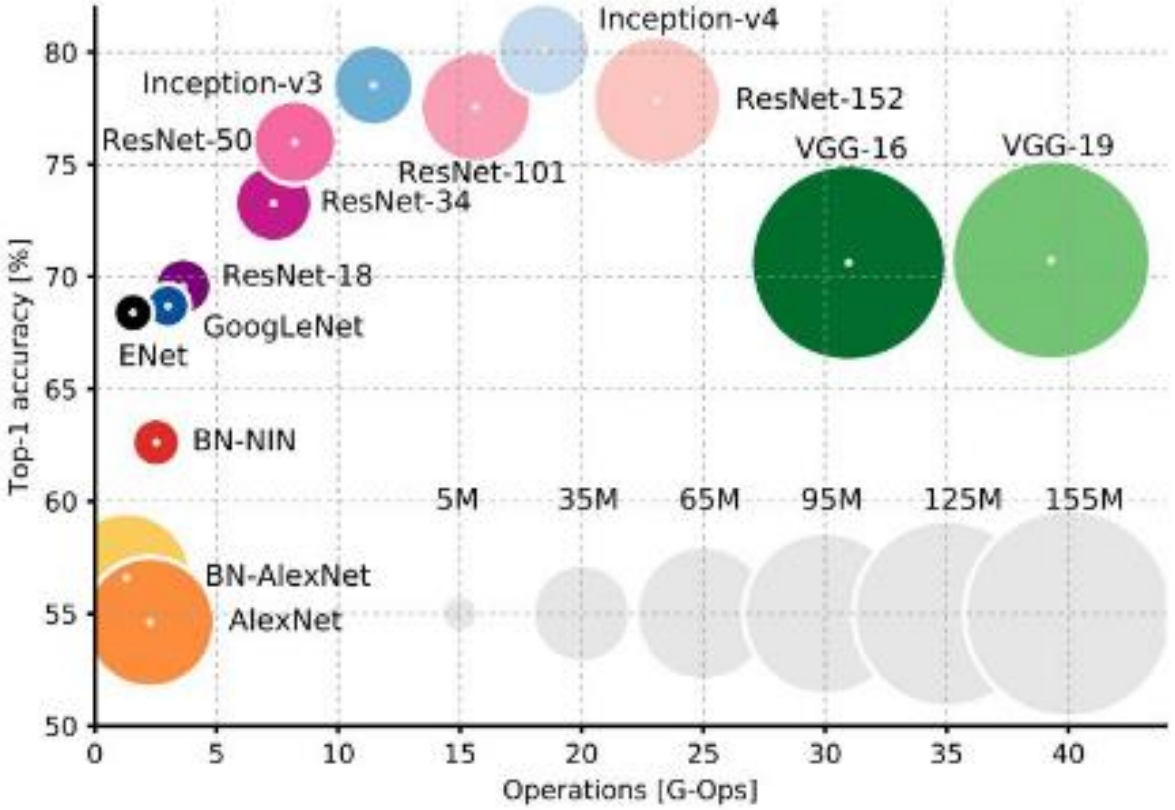

https://medium.com/analytics-vidhya/cnns-architectures-lenet-alexnet-vgg-googlenet-resnet-and-more-666091488df5

Recurrent Neural Networks (RNN)

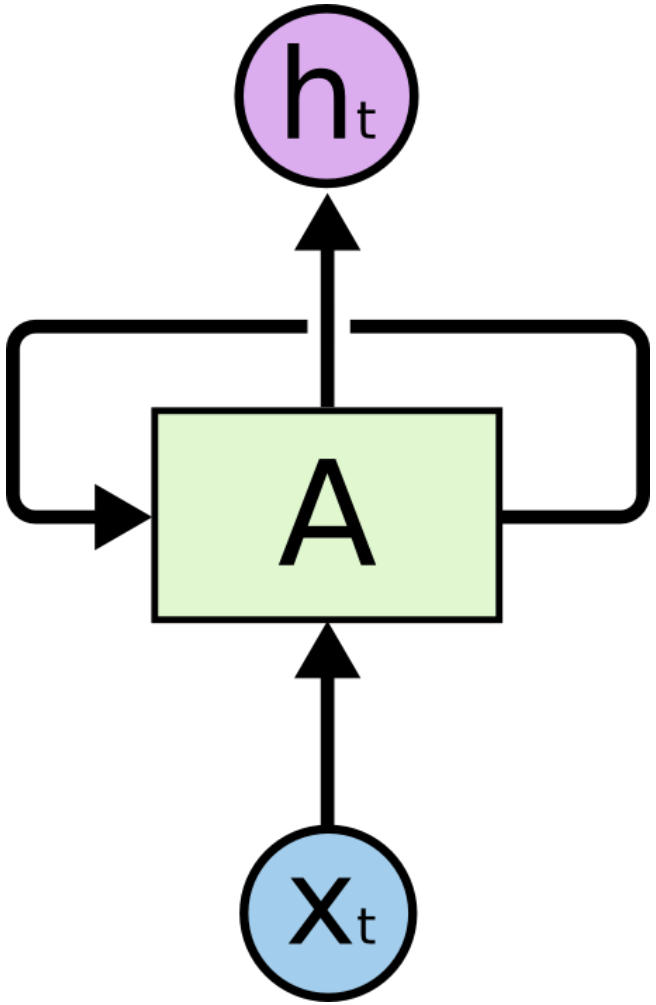

Unroll the RNN

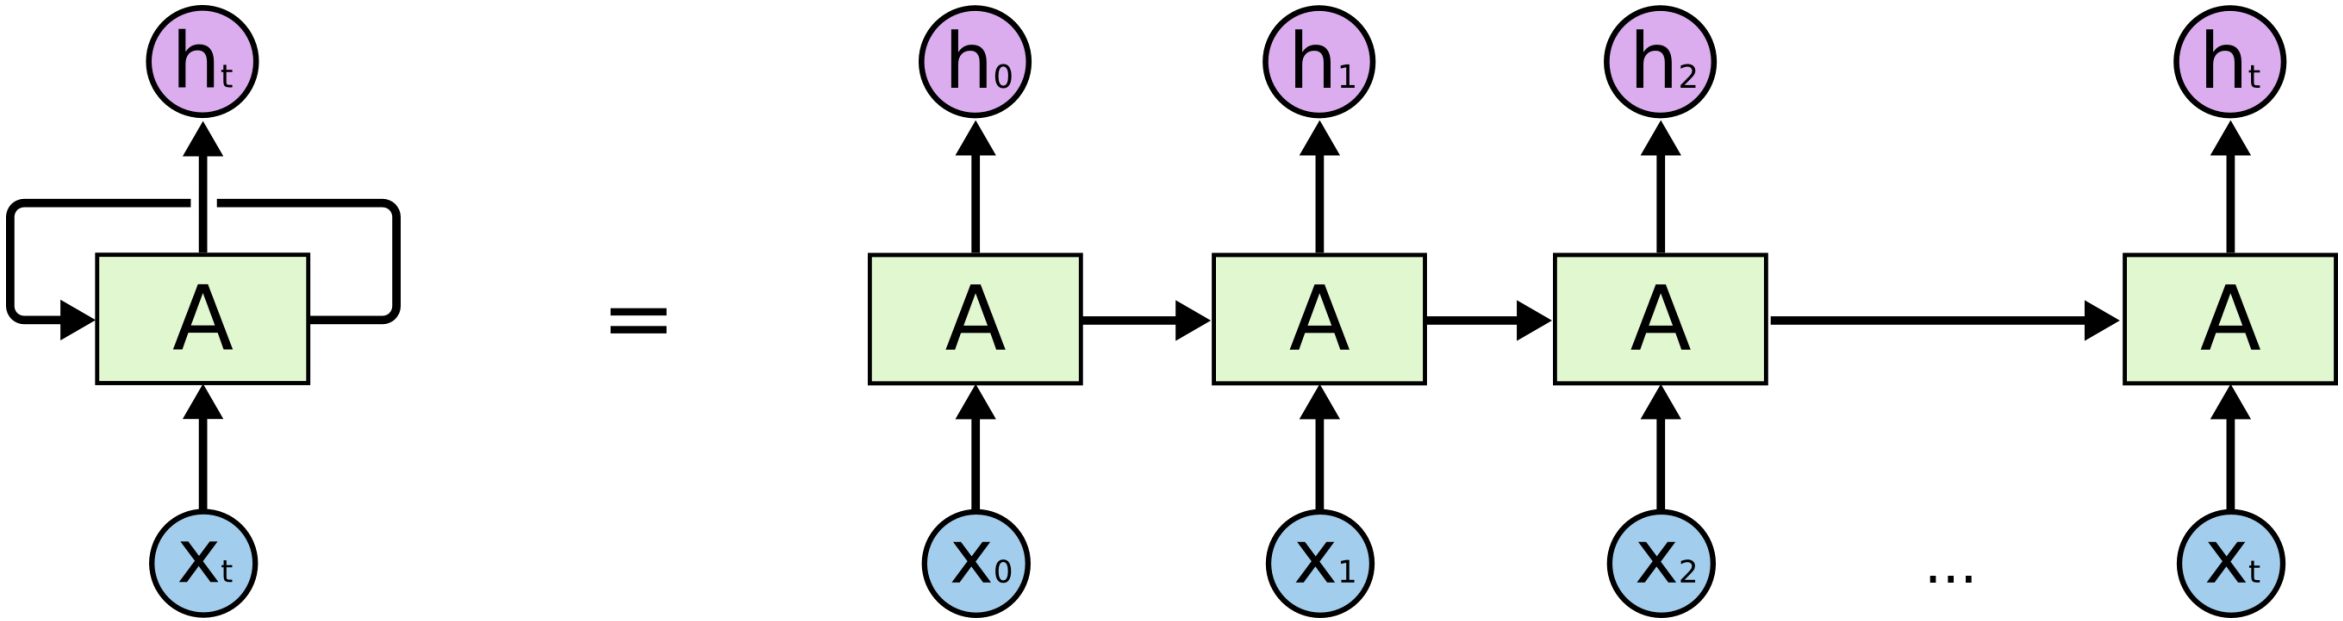

Long Short-term Memory (LSTM)

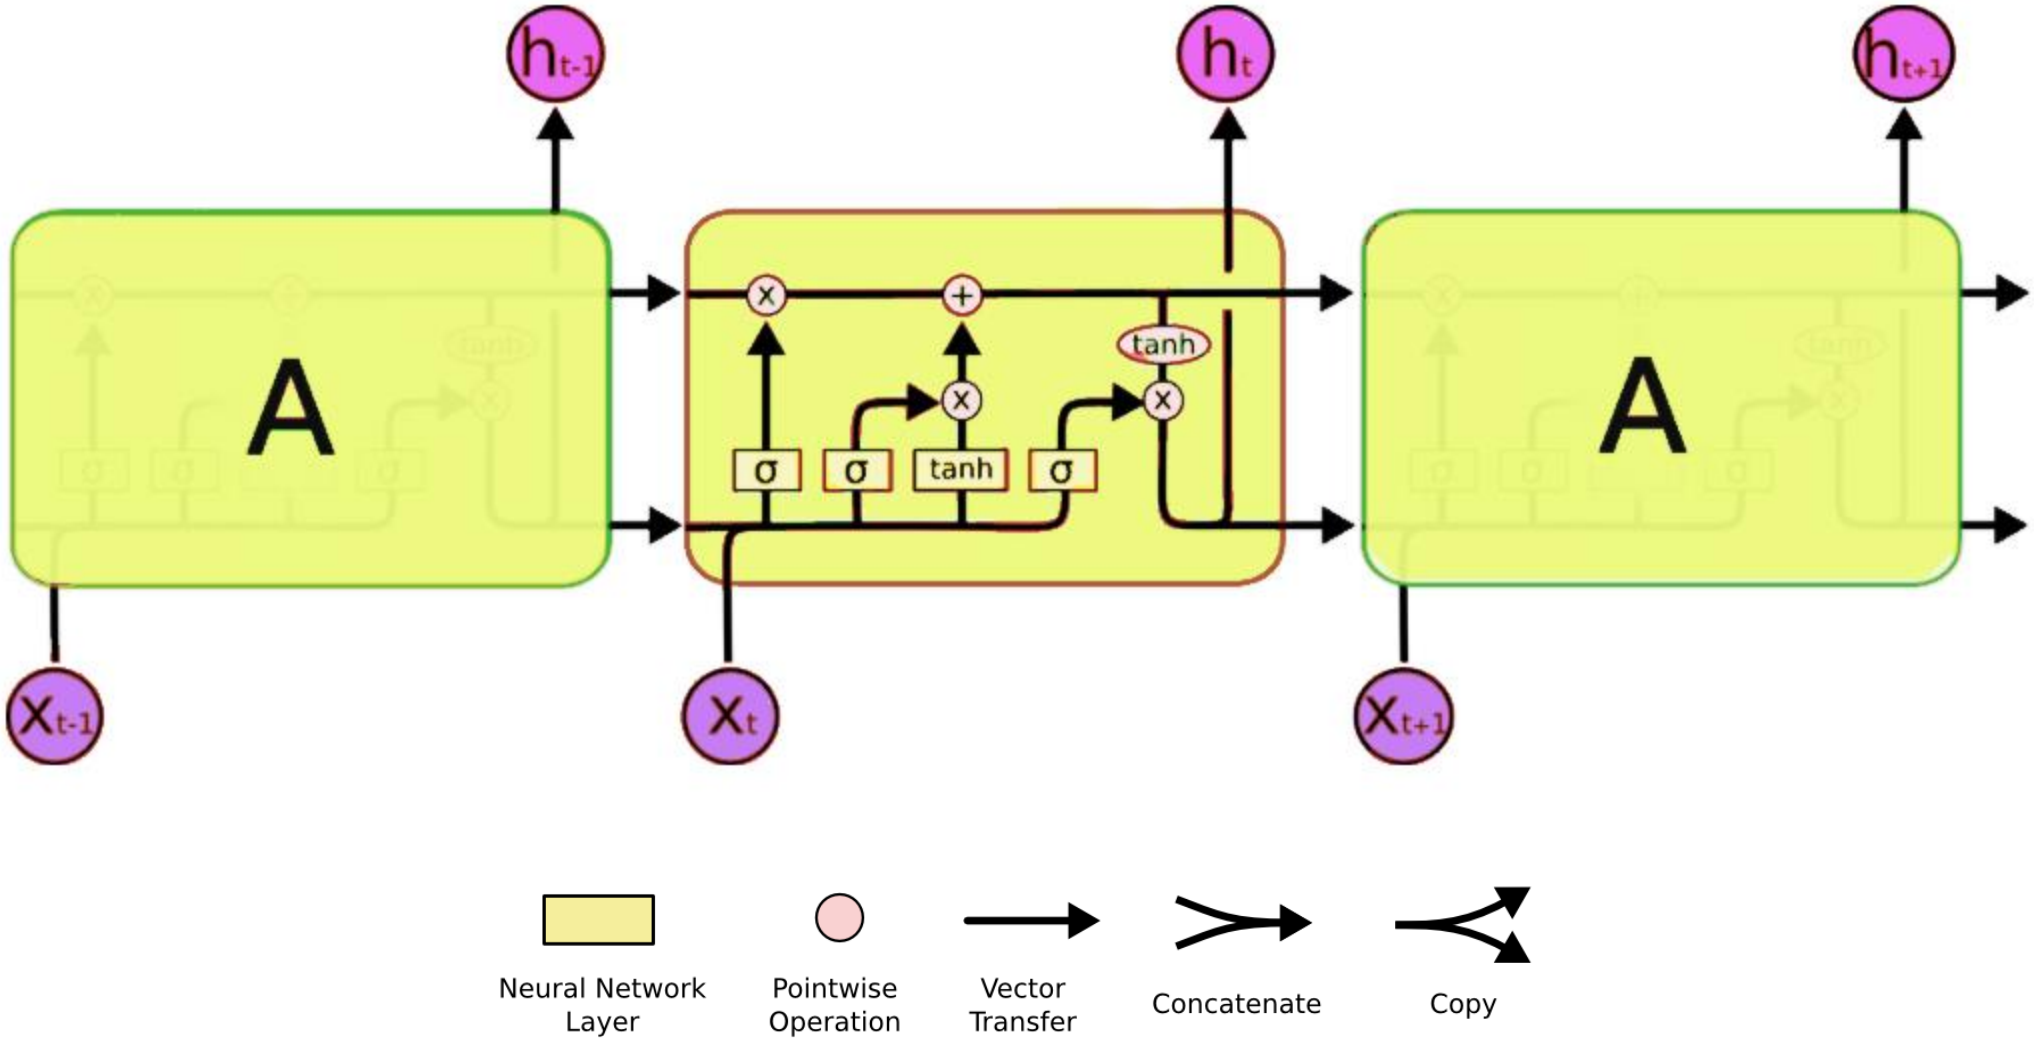

Deep Reinforcement Learning (深度強化學習)

Reinforcement Learning

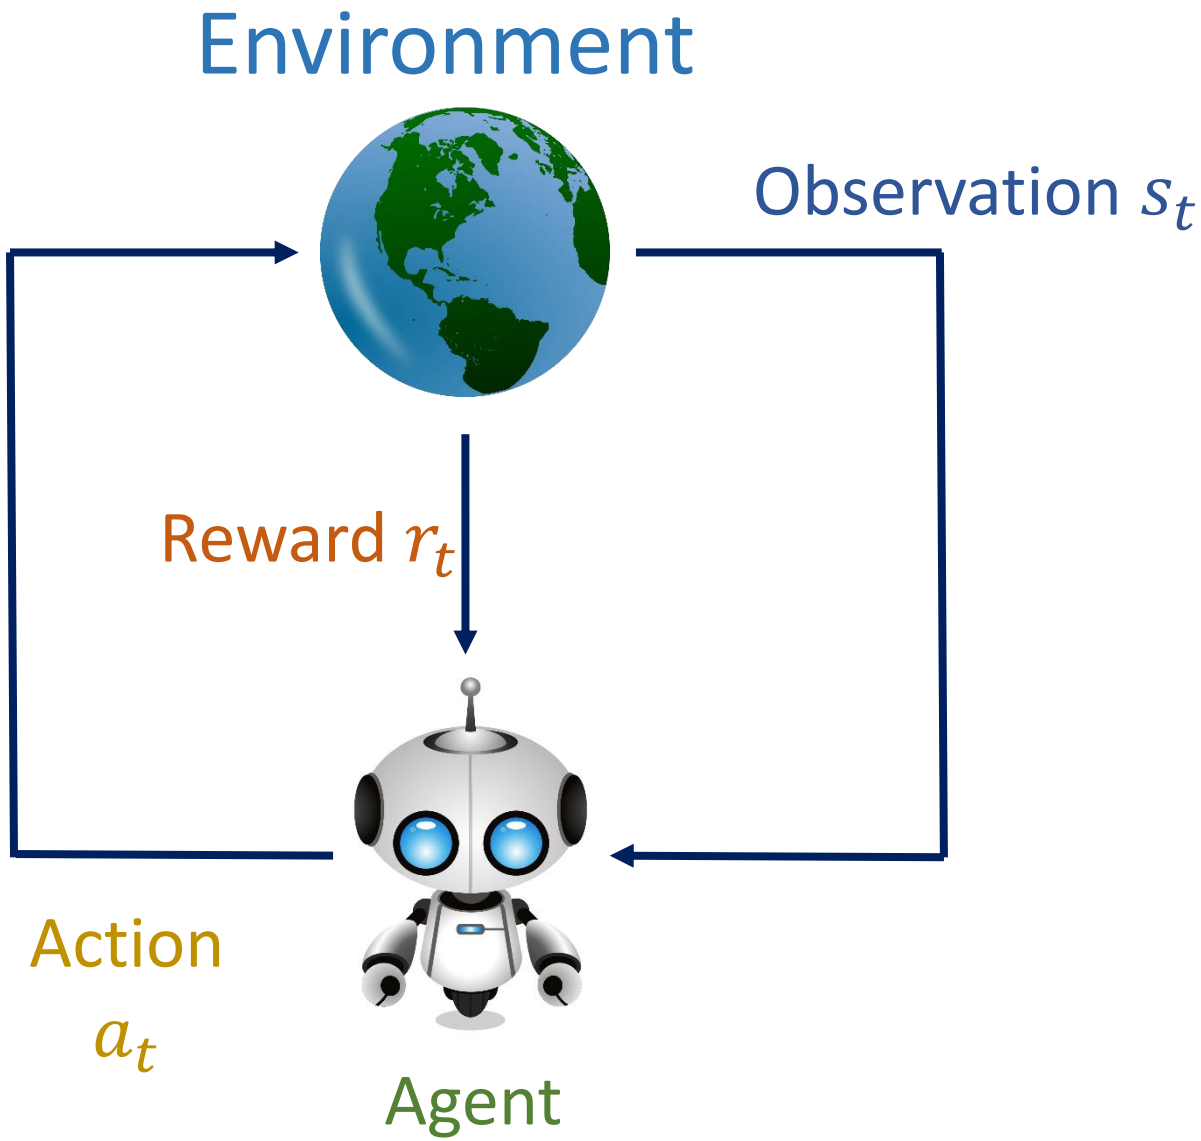

Google DeepMind

DRL in Atari

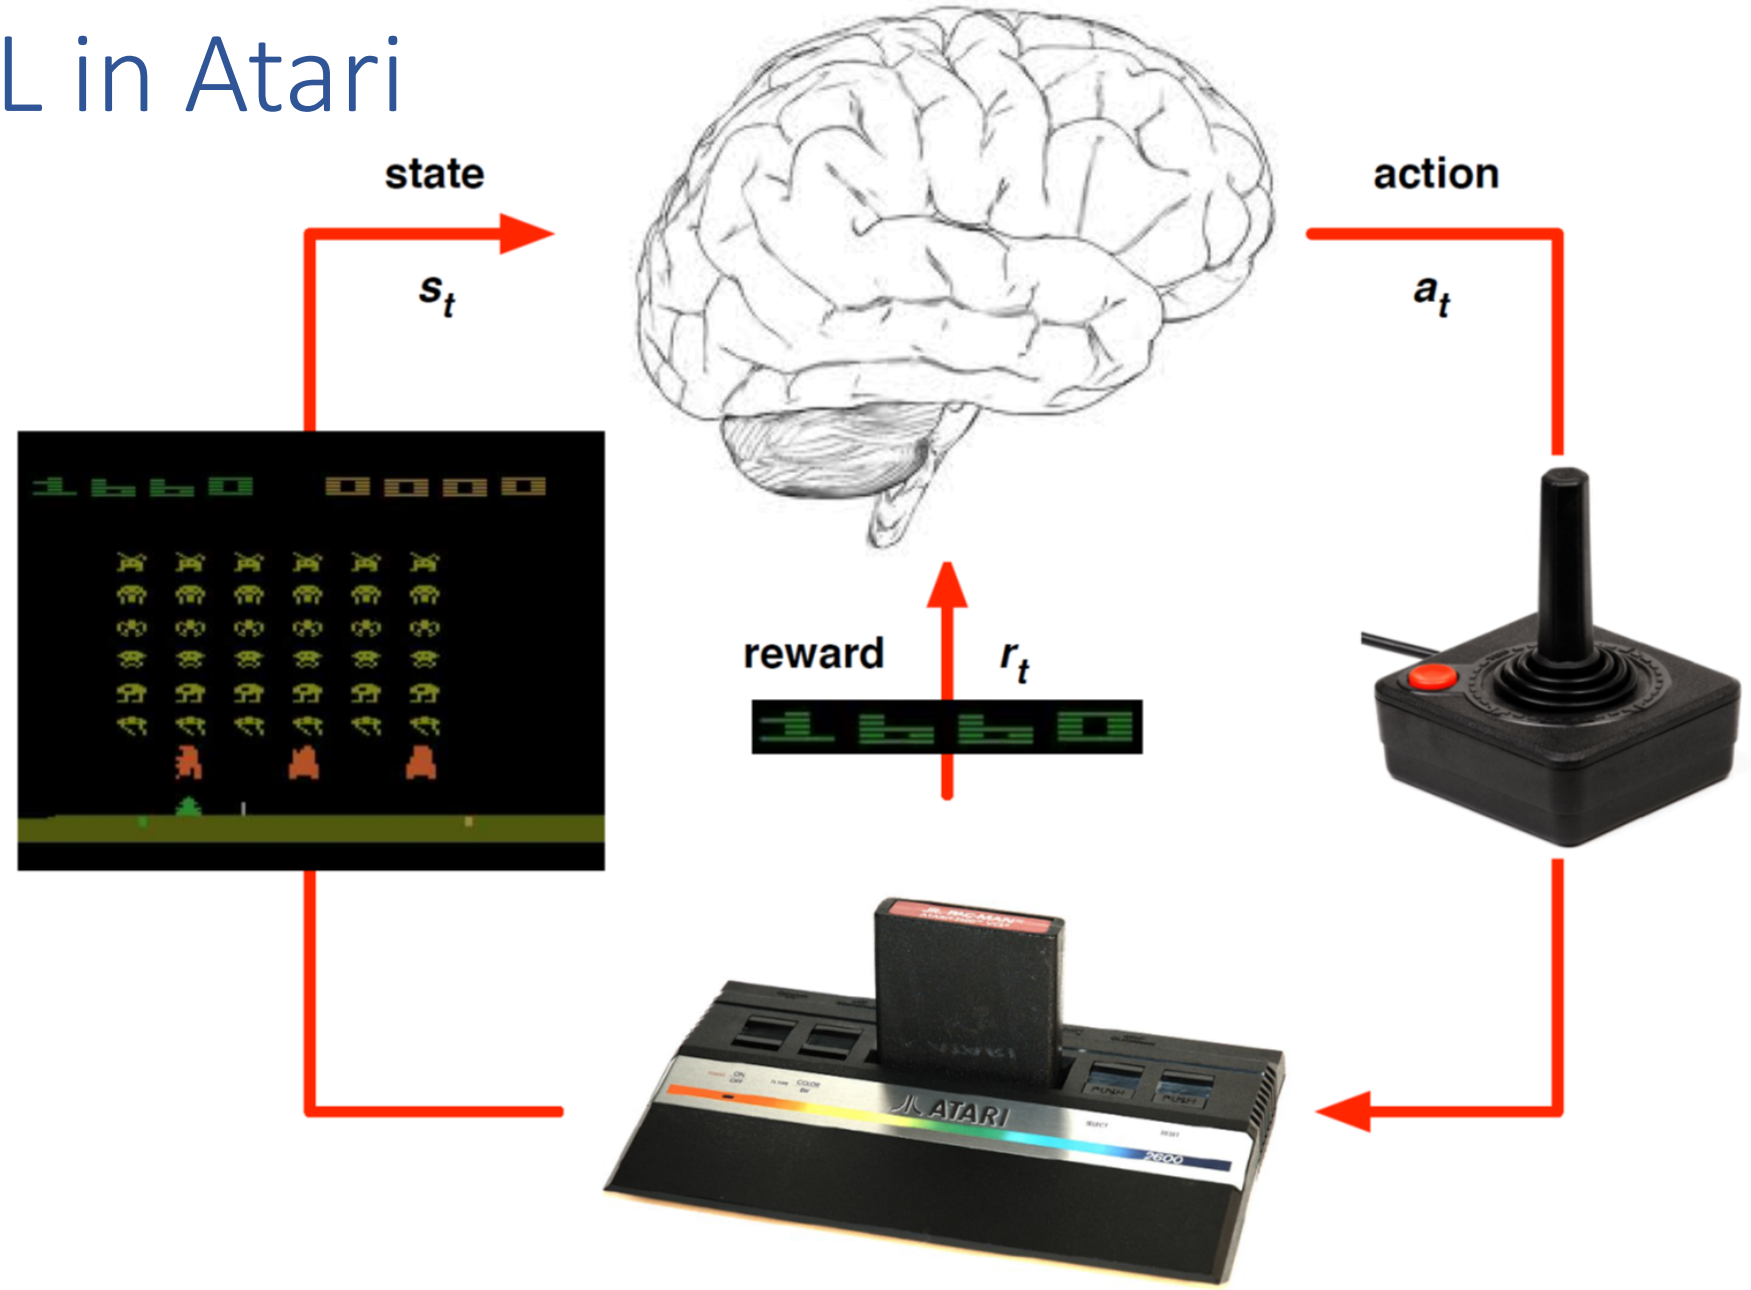

Mnih et al., “Human Level Control through Deep Reinforcement Learning,” *Nature*, 2015

Learning to Play Atari Games

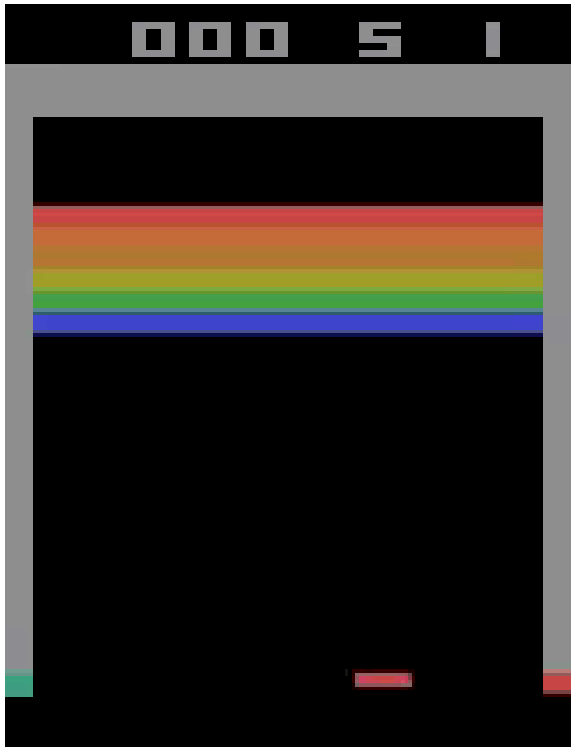

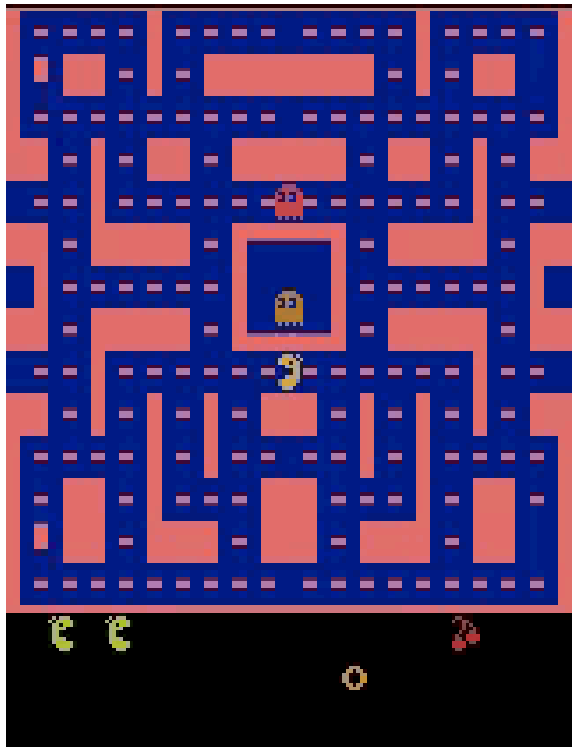

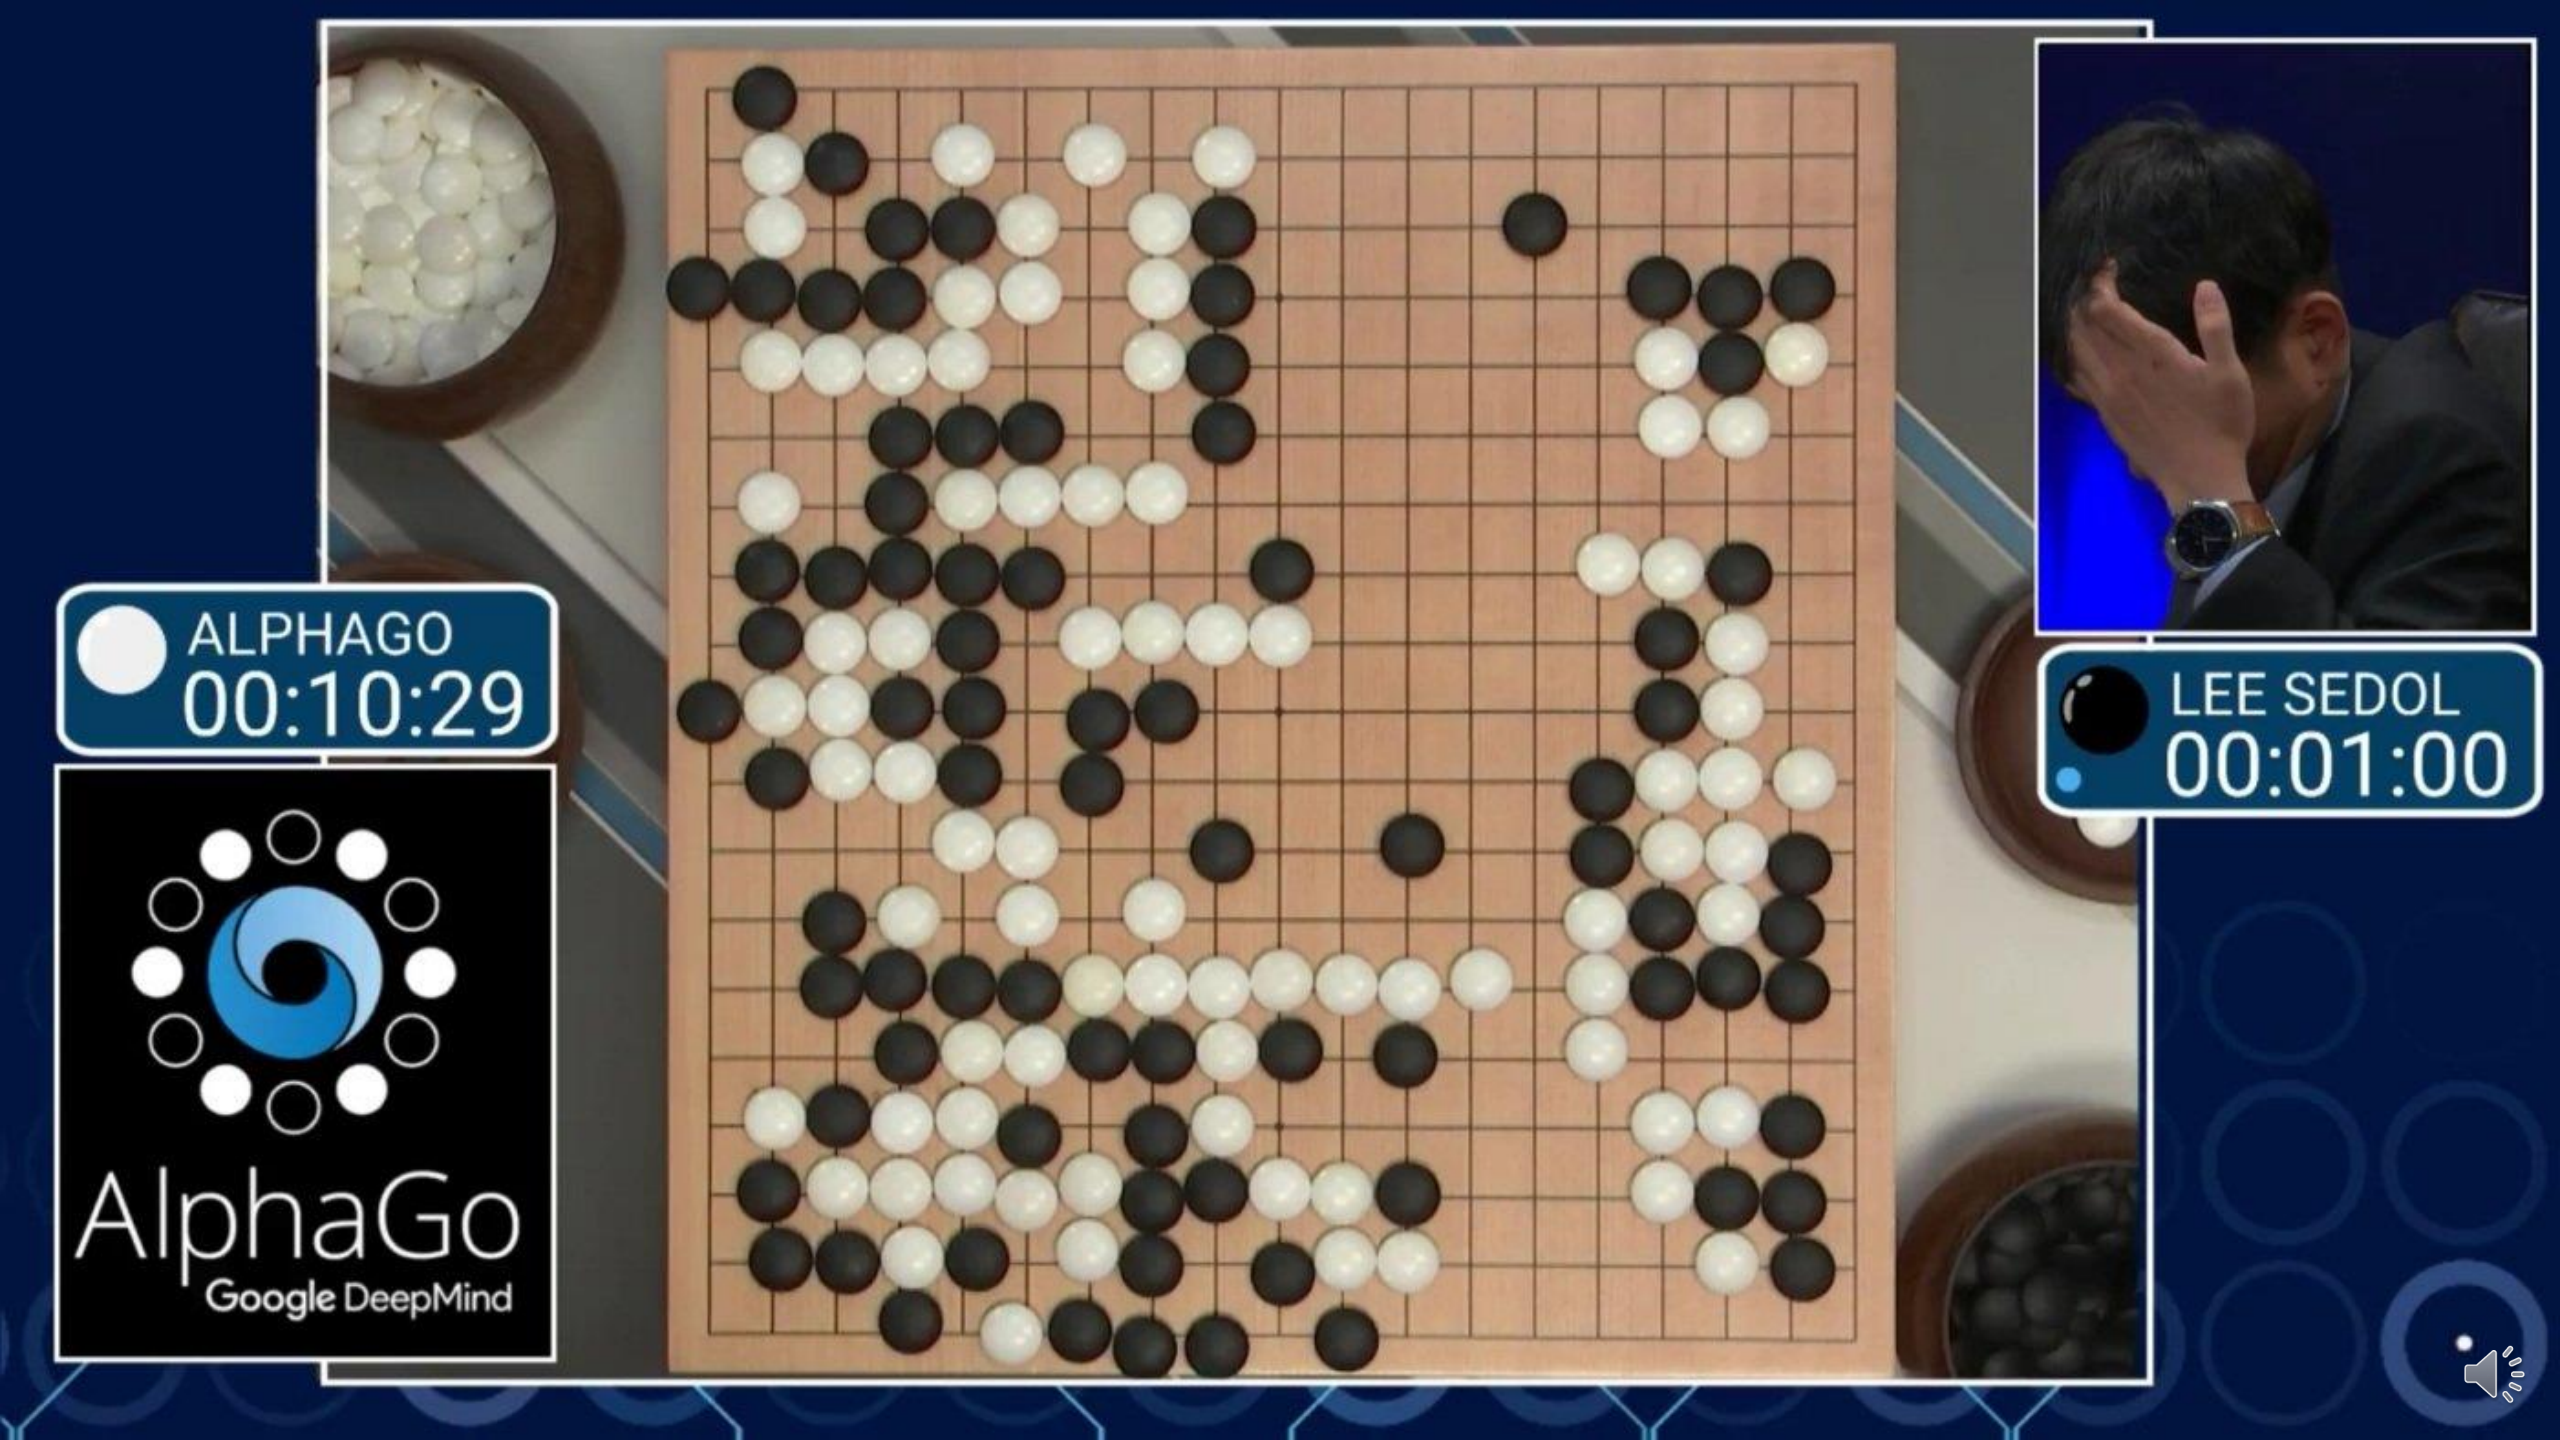
ALPHAGO
00:10:29
AlphaGo
Google DeepMind
LEE SEDOL
00:01:00

Dr. Aja Huang (黃士傑)

The 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 Zero
Starting from scratch

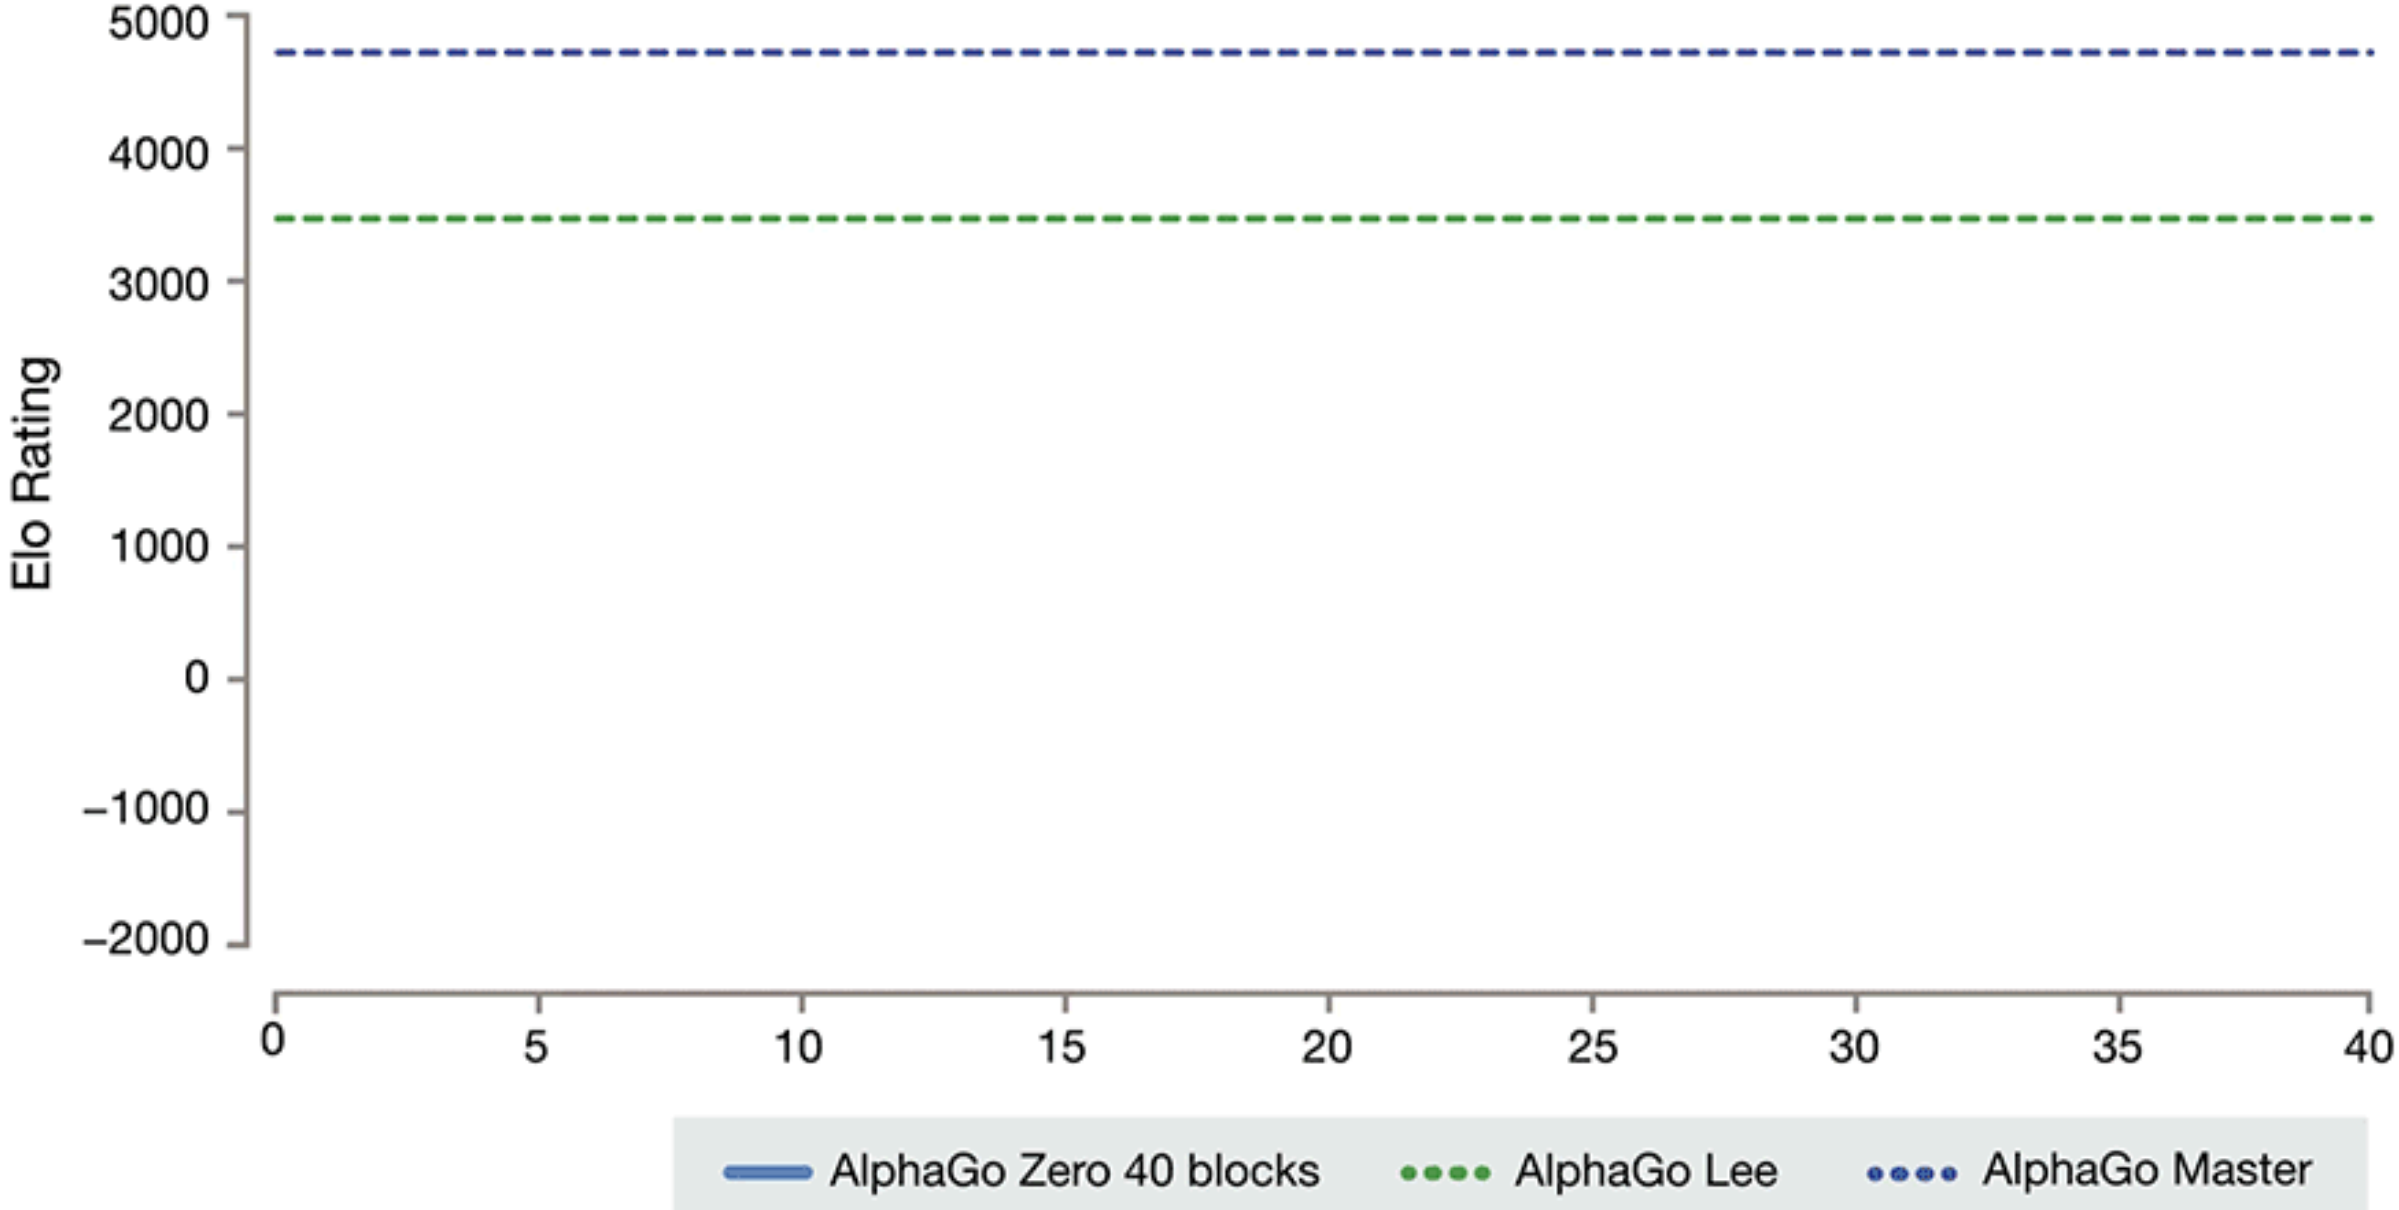
5000
4000
3000
2000
1000
0
−1000
−2000
Elo Rating
0
5
10
15
20
25
30
35
40
AlphaGo Zero 40 blocks
AlphaGo Lee
AlphaGo Master


Elo Rating
5000
4000
3000
2000
1000
0
−1000
−2000
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

Human Extinction ?

Raw Observations

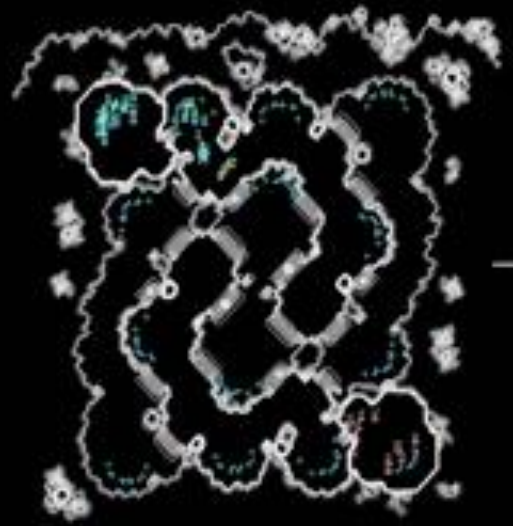

Neural Network Activations

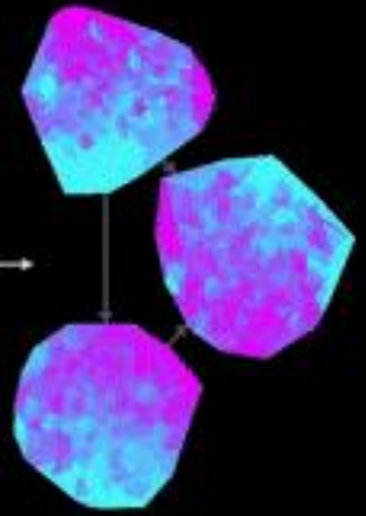

Considered Location

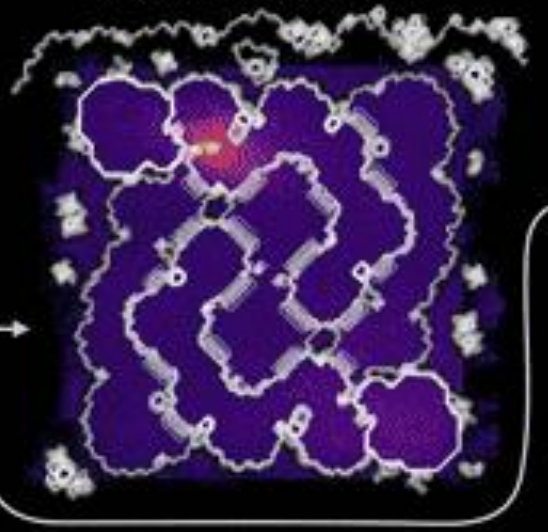

Outcome Prediction

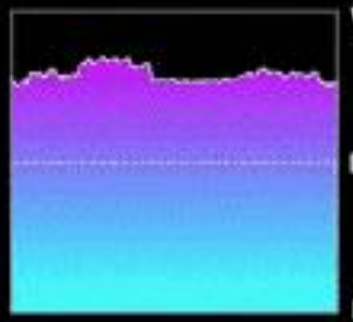

Considered Build/Train

Virtual-to-real Learning

- Inspired by DeepMind (Mnih et al., *Nature*, 2015)
  - "Human Level Control through Deep Reinforcement Learning"
- Applied to computer vision applications
  - **Image segmentation**: Armeni et al. (2016), Qiu et al., (2017)
  - **Indoor navigation**: Brodeur et al. (2017), Gupta et al. (2017), Savva et al. (2017), Wu et al. (2018)
  - **Autonomous vehicles**: Marinez et al. (2017), Muller et al. (2018), Pan et al. (2017), Shah et al. (2018)

UnrealCV

CAD^{2}Real

Semantic Segmentation

Depth Prediction

VIVID

Autonomous Navigation

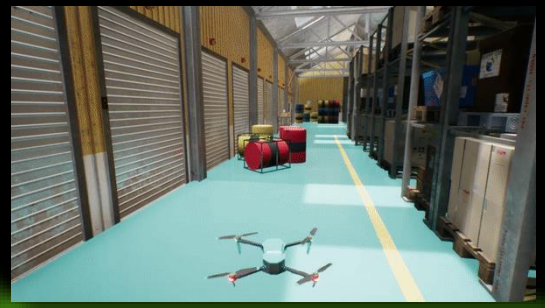

Action Recognition

Simulate Real-life Events

Searching for the Shooter

Limits of Deep Learning
深度學習無所不能?

No Idea of Real World

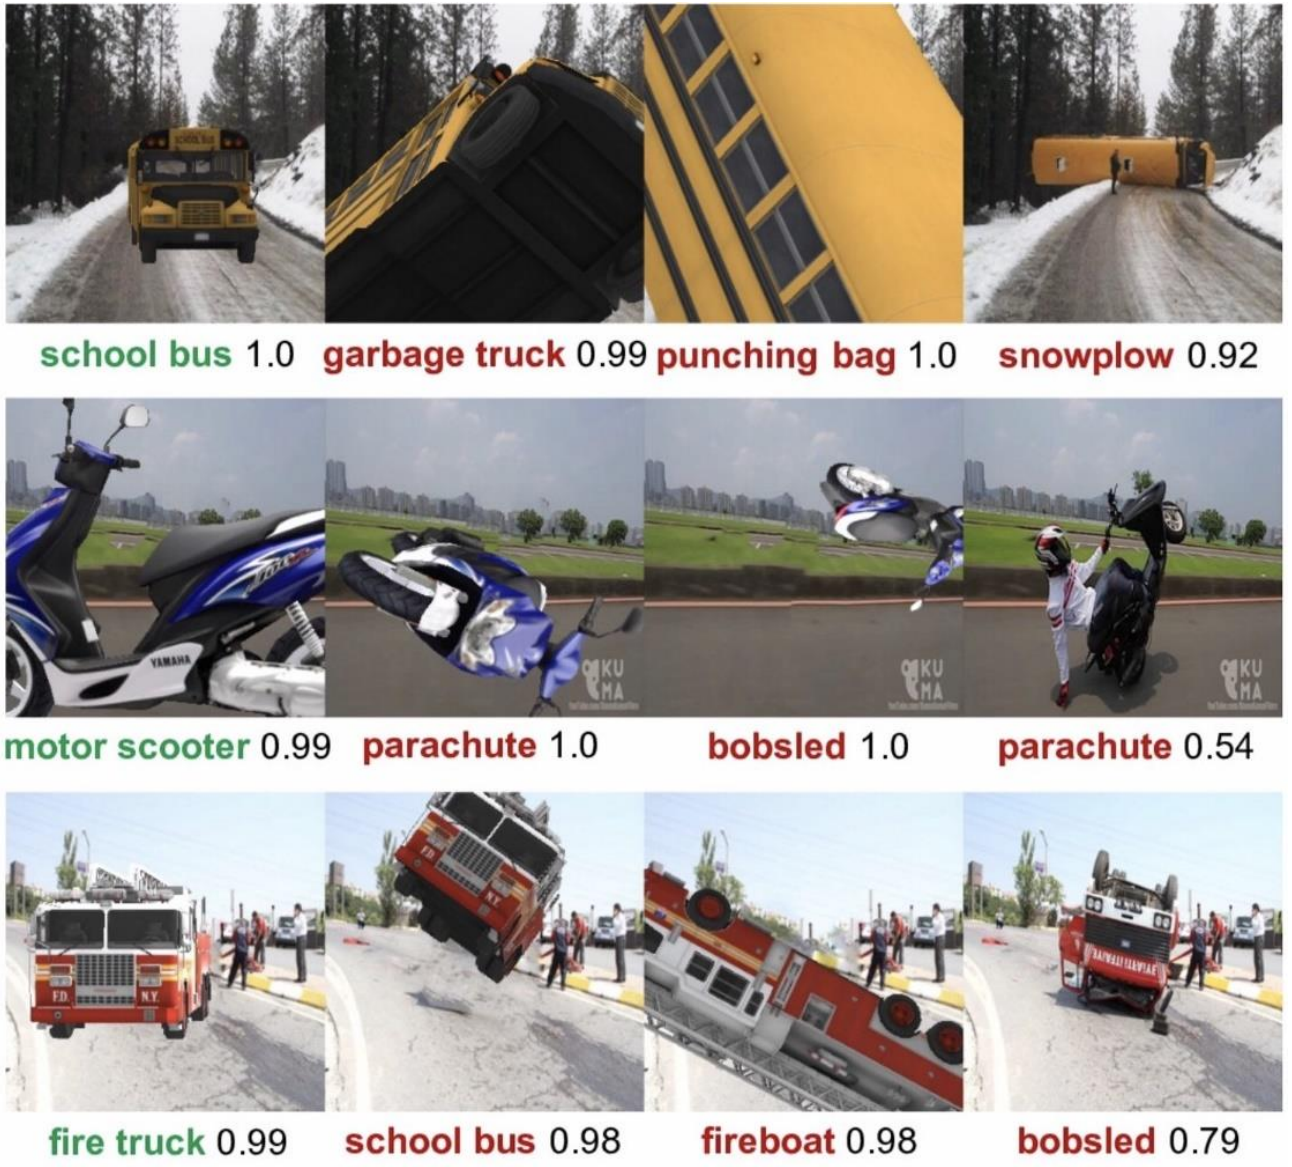

Alcorn et al., “Strike (with) a pose: Neural networks are easily fooled by strange poses of familiar objects,”CVPR 2019.


2020-06-01 06:44:03
國1 北 268K+410 水上路段
民視新聞台 HD
嘉義
大貨車翻覆橫倒車道 特斯拉高速撞進車廂


民視新聞台
HD
國道特斯拉撞大貨車 相隔6公里又有五連撞

Adversarial Attack

“panda”

57.7% confidence

$+\epsilon$

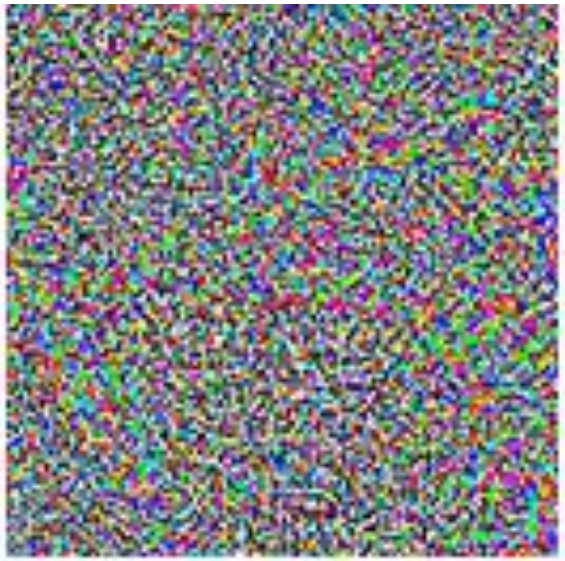

=

“gibbon”

99.3% confidence

Generative Adversarial Networks (GAN)

- Ian Goodfellow

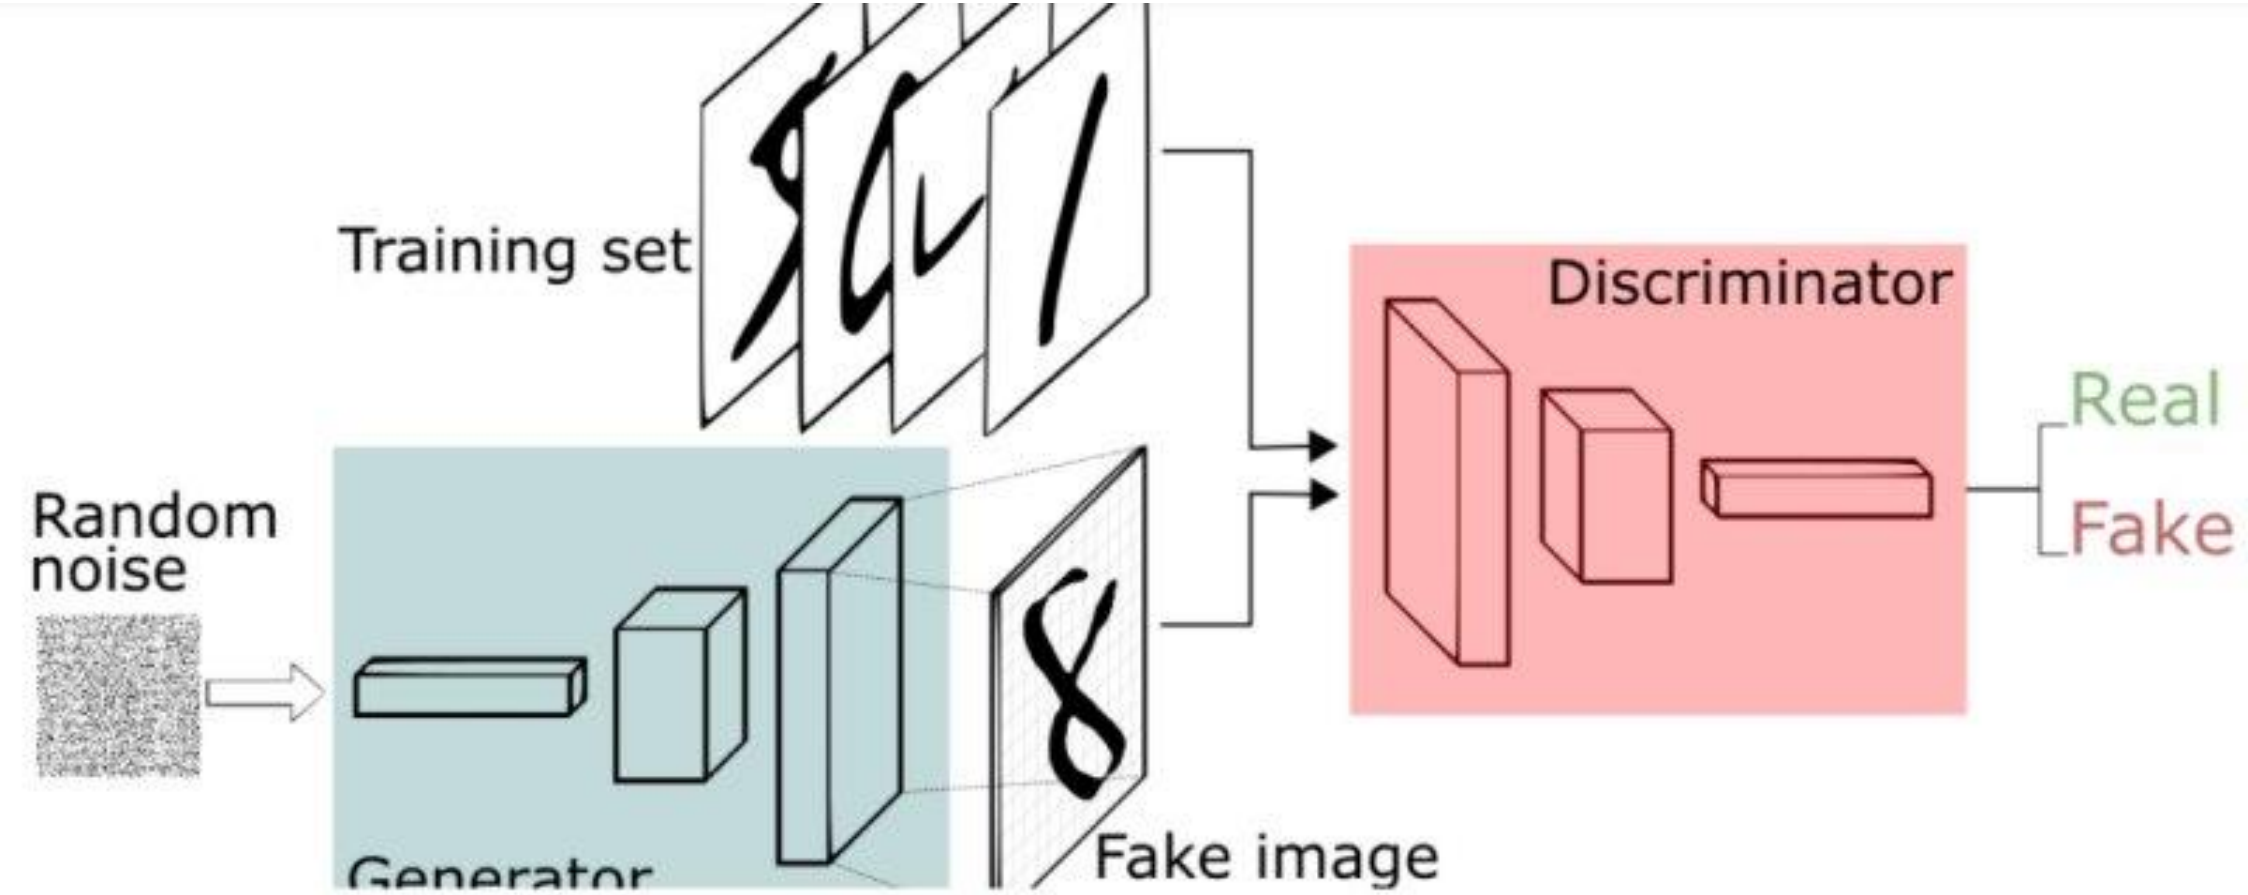

Painting like Van Gogh

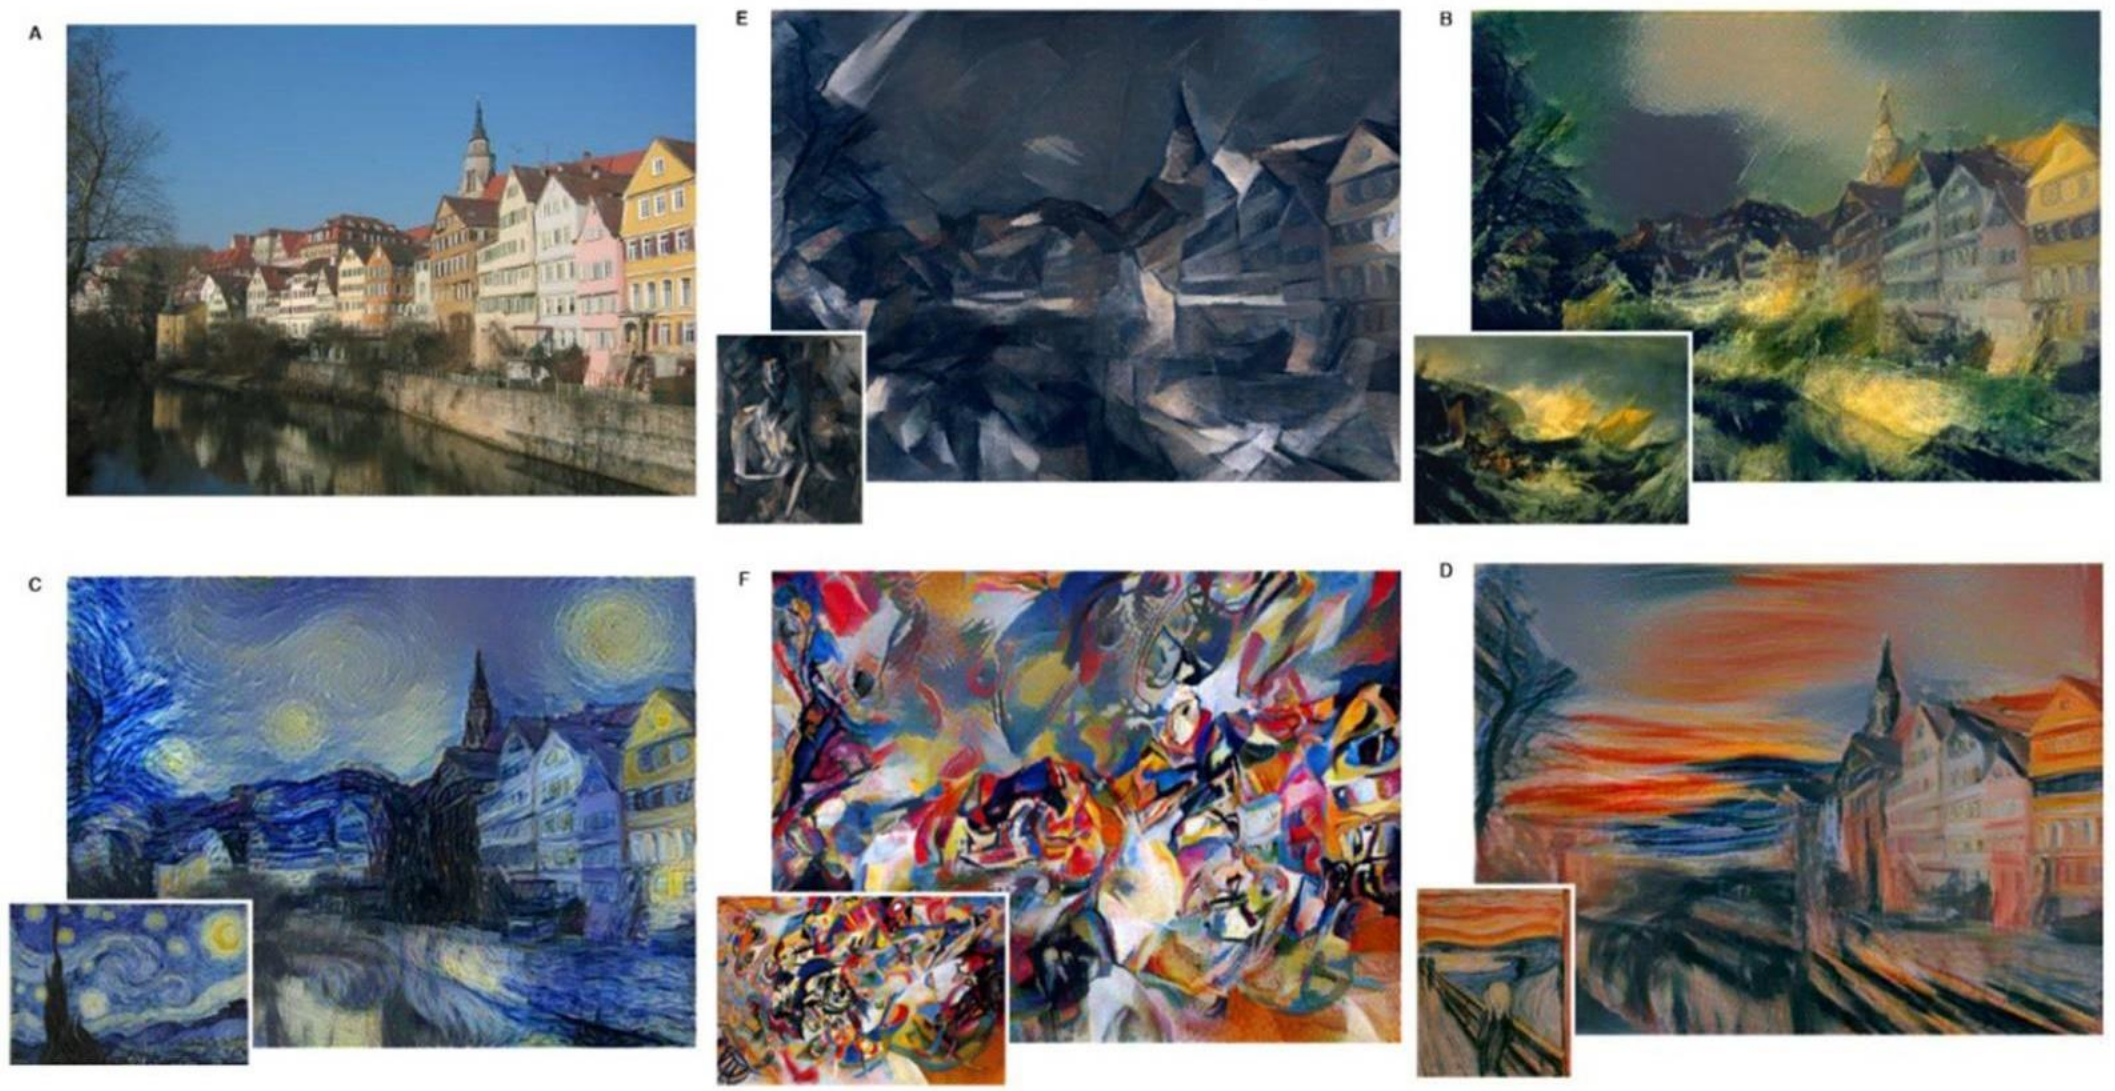

Super Resolution

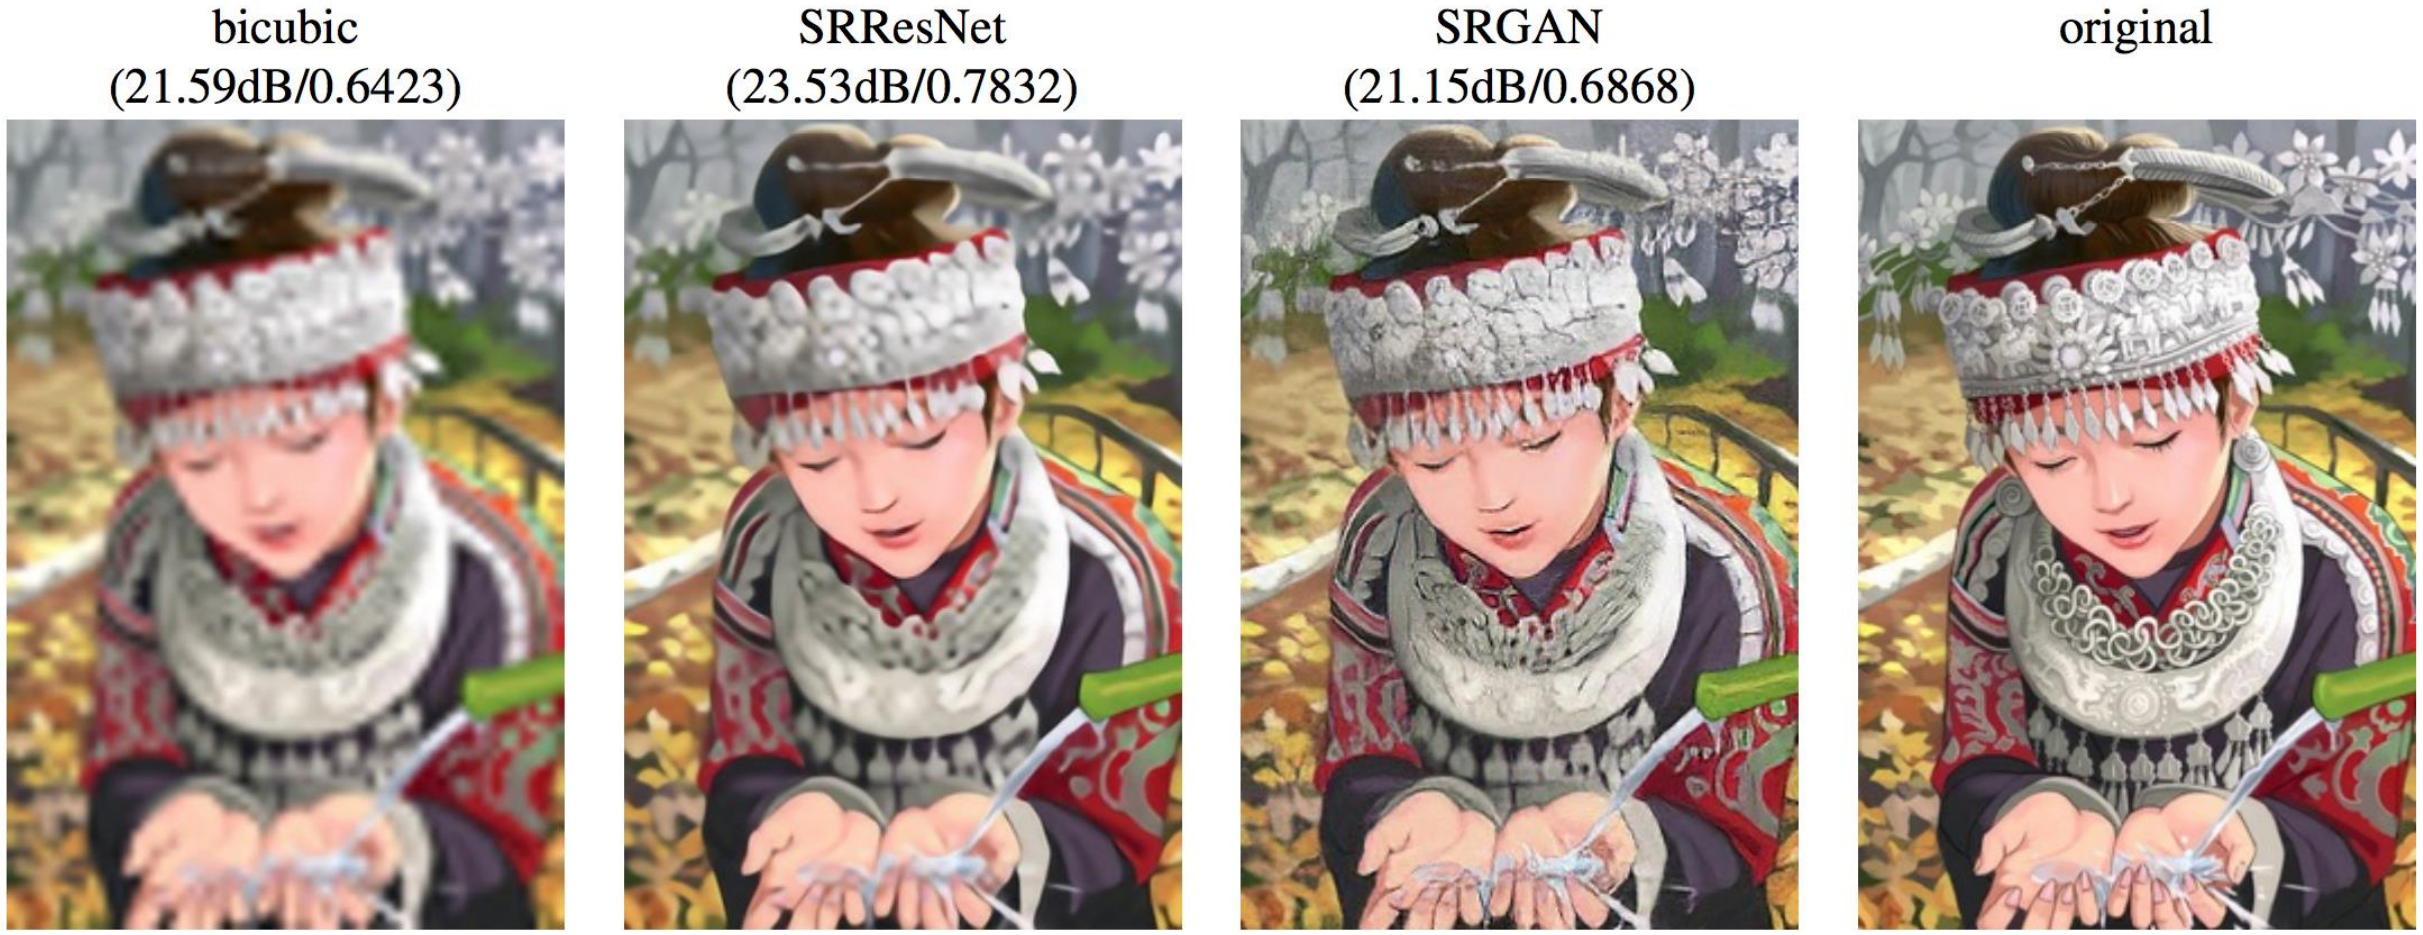

Figure 2: From left to right: bicubic interpolation, deep residual network optimized for MSE, deep residual generative adversarial network optimized for a loss more sensitive to human perception, original HR image. Corresponding PSNR and SSIM are shown in brackets. [4× upscaling]

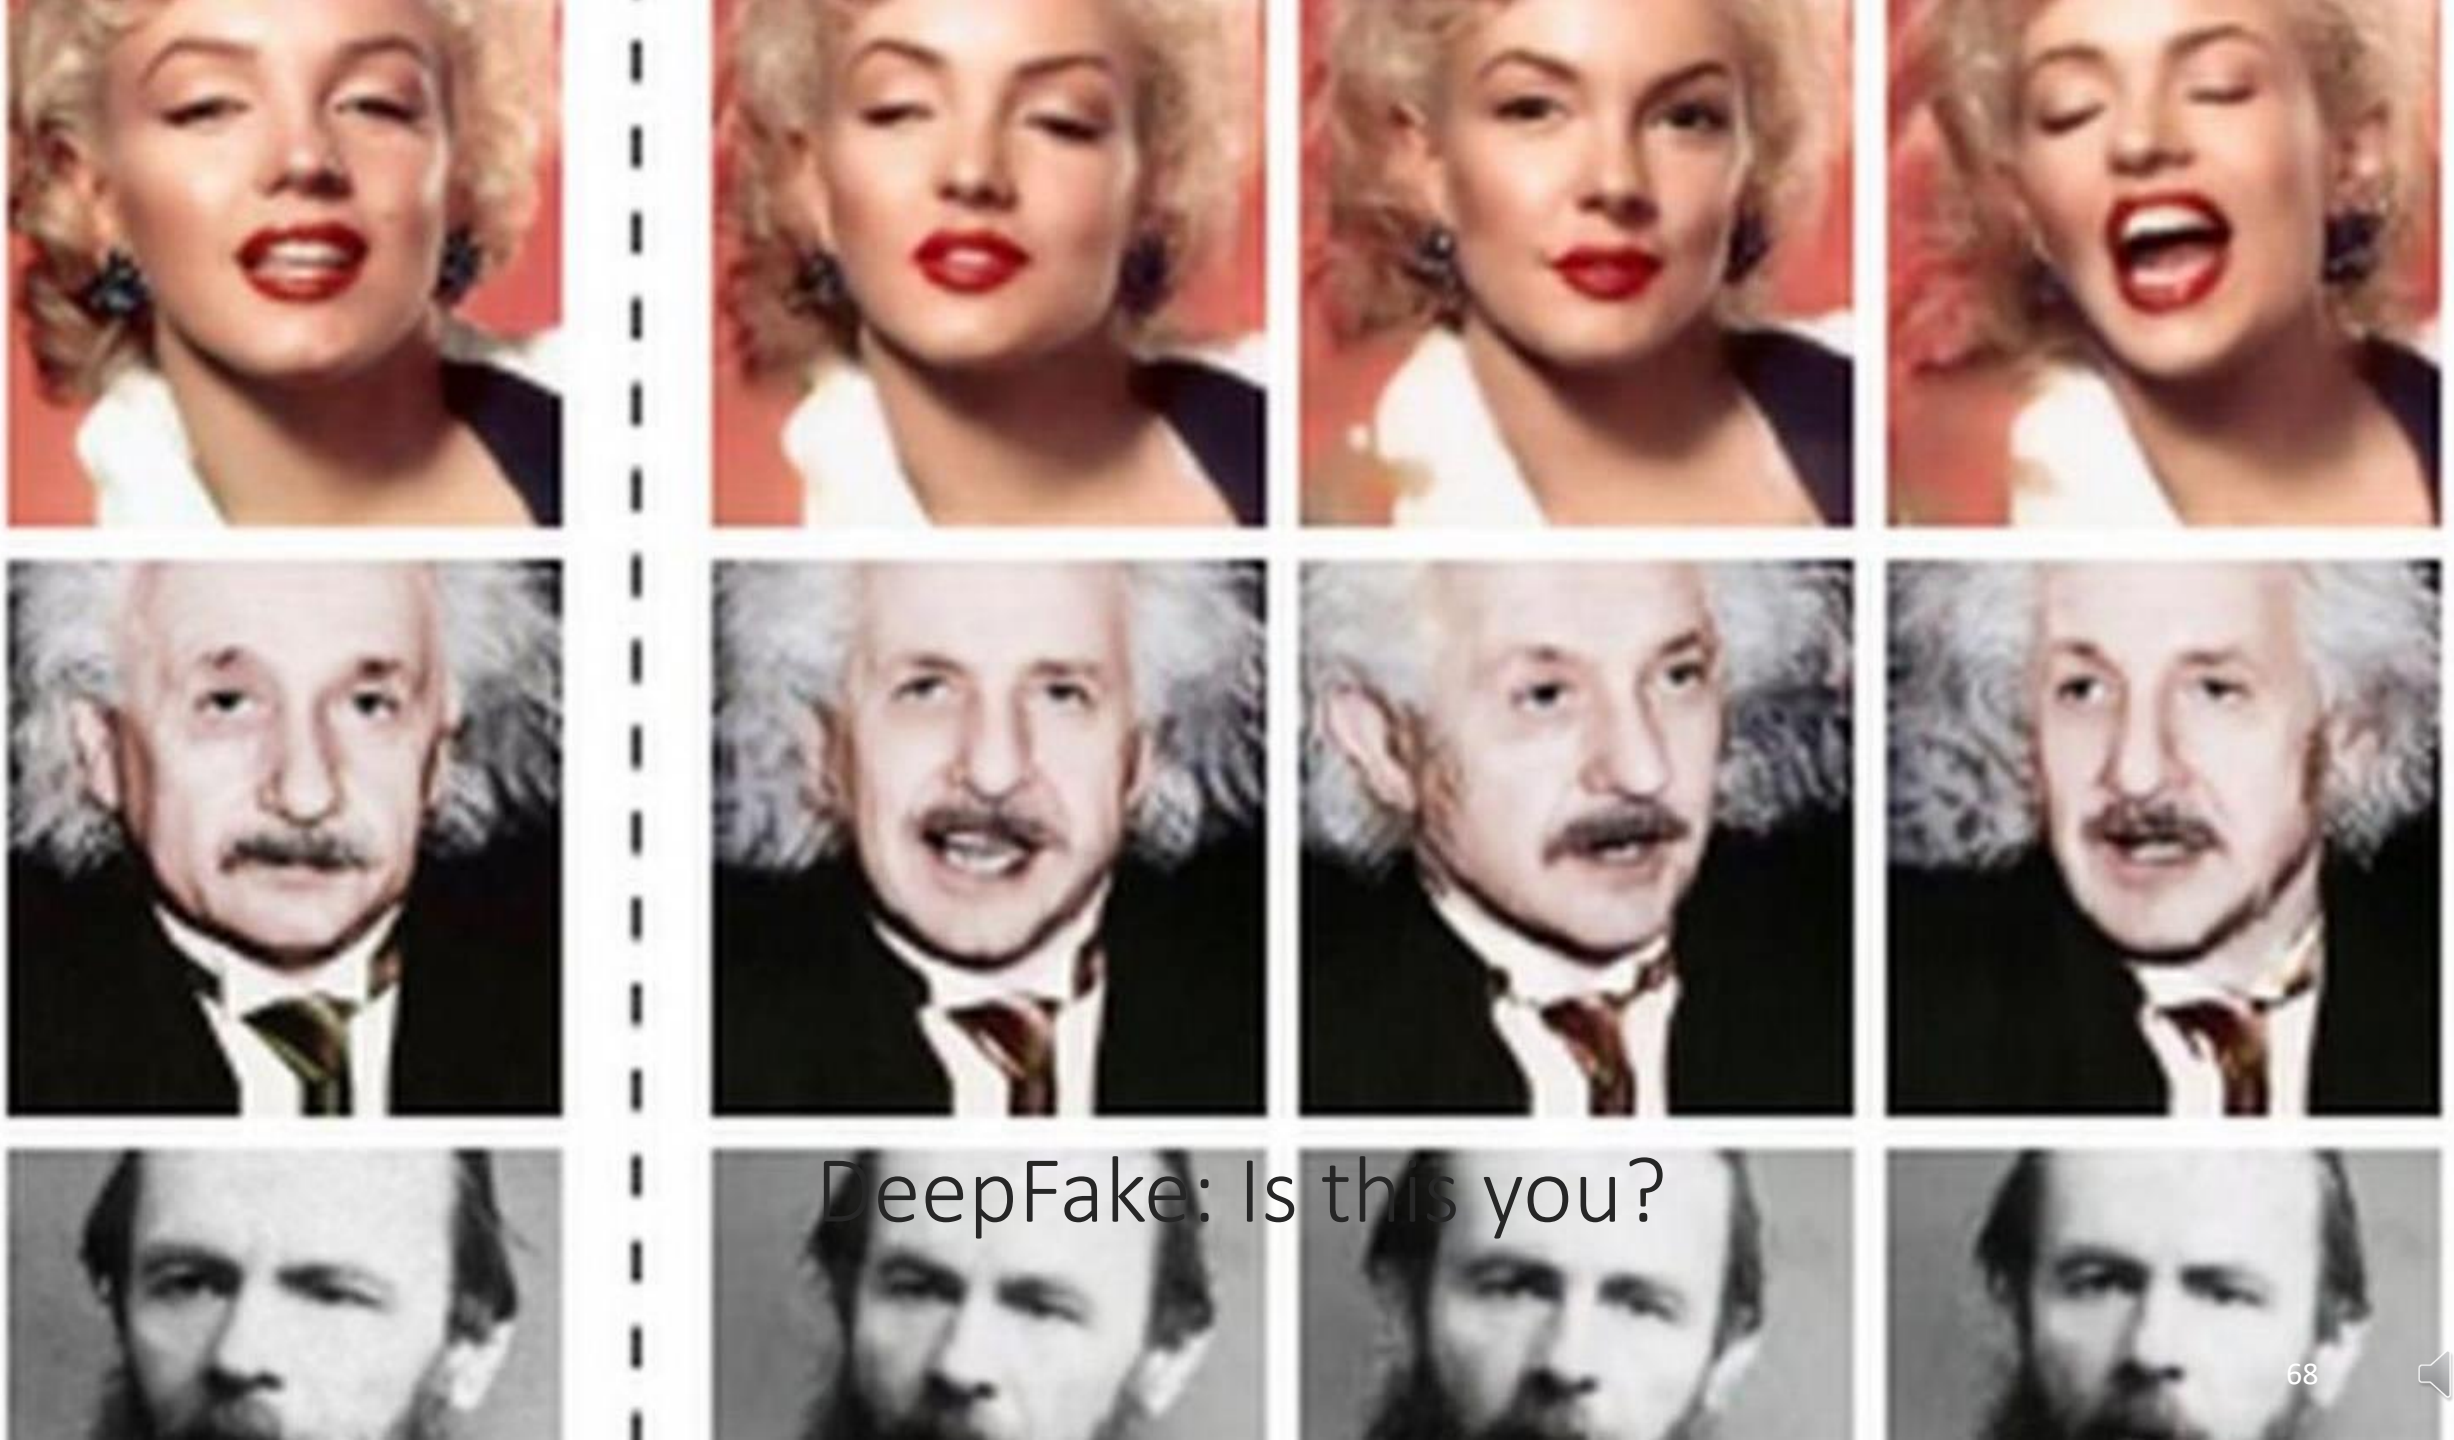

DeepFake: Is this you?

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

Google's AutoML

- Learning neural network cells automatically

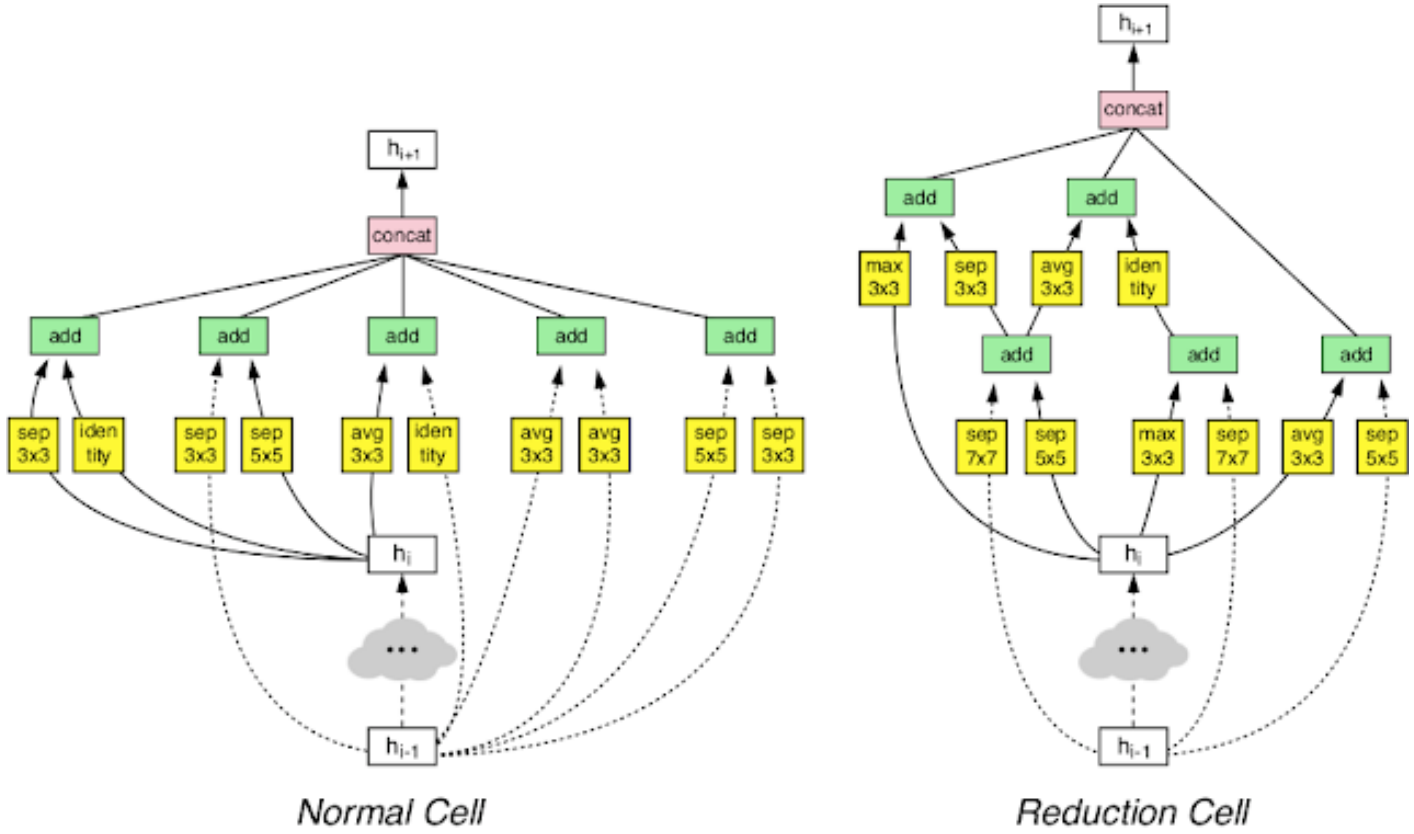

https://ai.googleblog.com/2017/11/automl-for-large-scale-image.html

AutoML on ImageNet

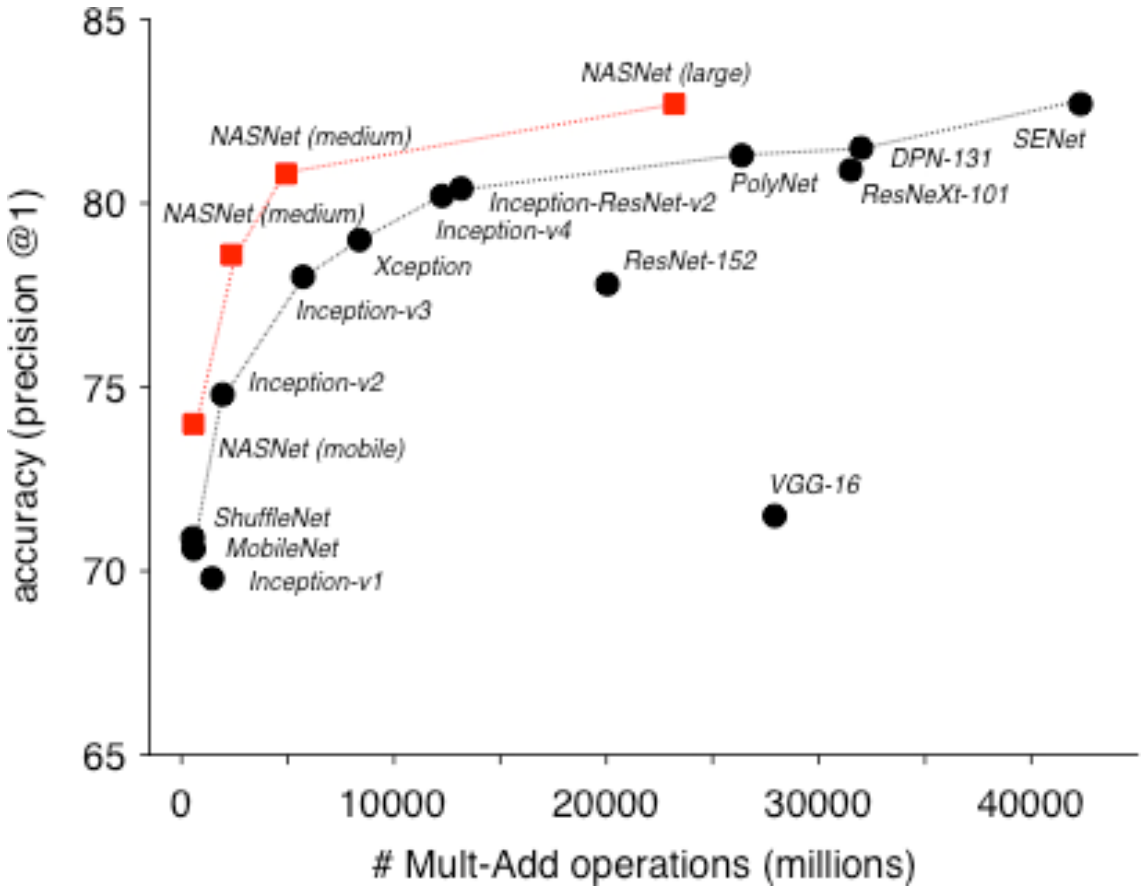

EfficientNet (May, 2019)

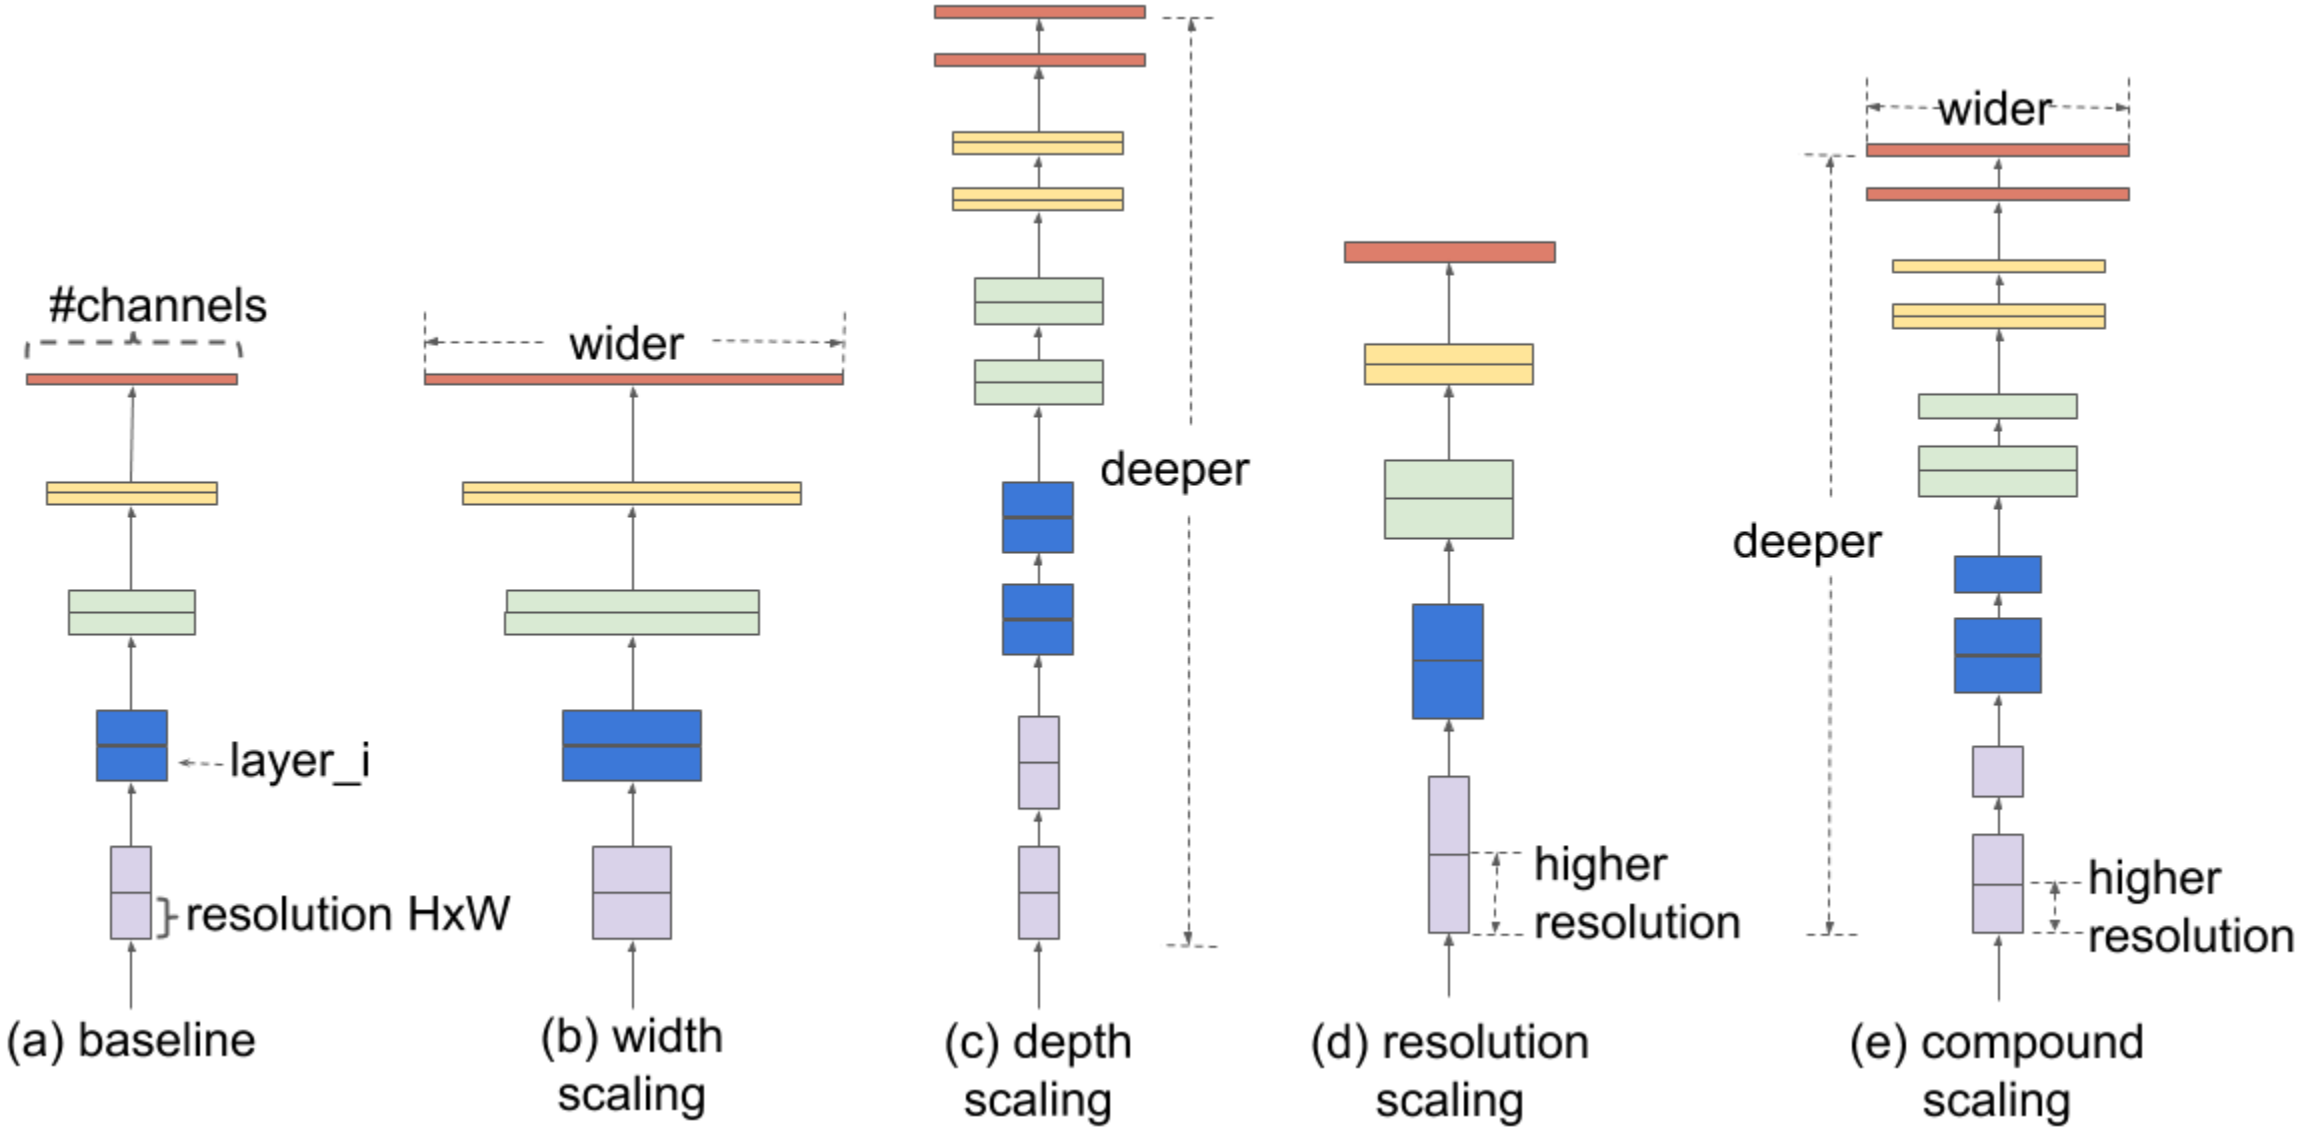

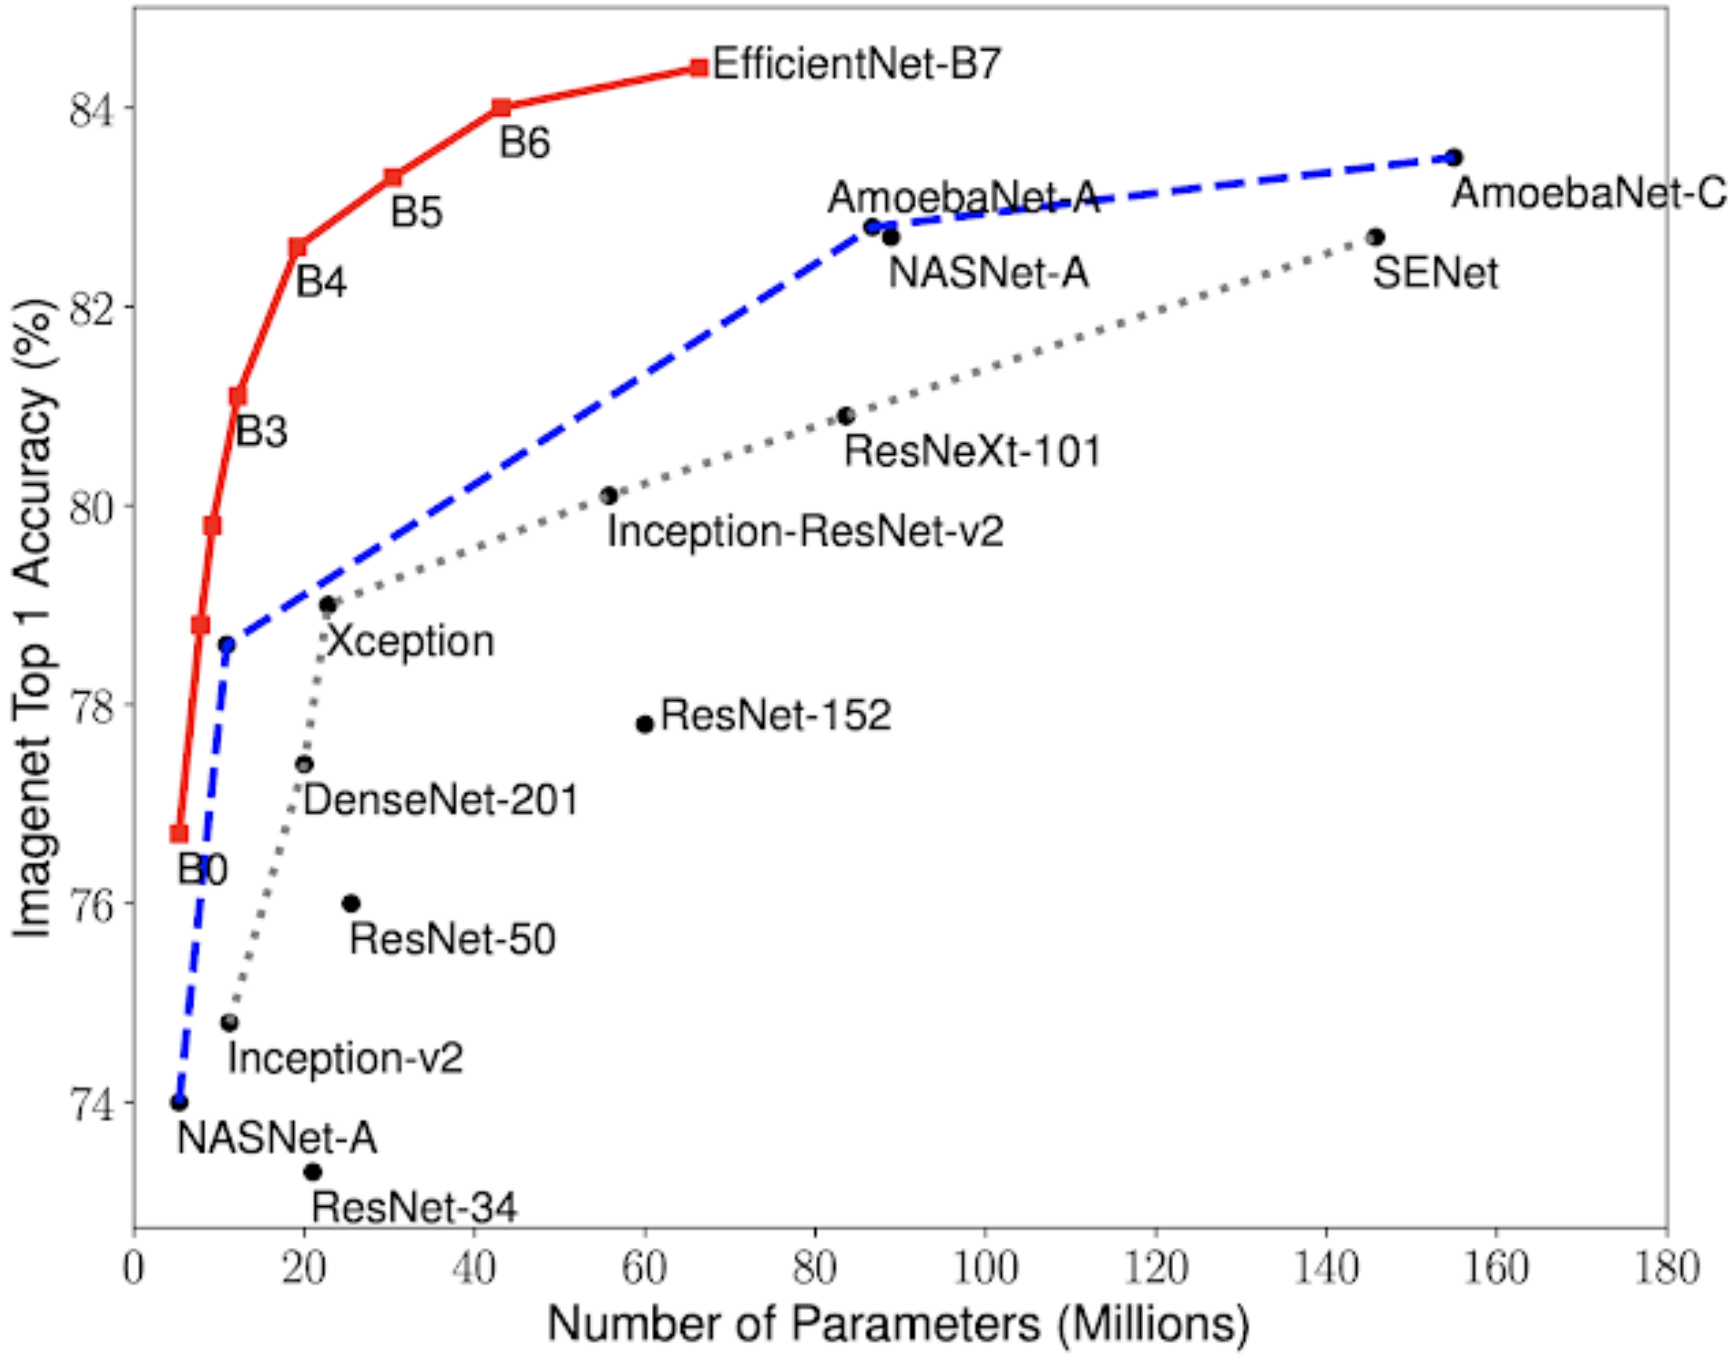
EfficientNet-B7
B6
B5
B4
B3
B0
AmoebaNet-A
AmoebaNet-C
NASNet-A
SENet
ResNeXt-101
Inception-ResNet-v2
Xception
ResNet-152
DenseNet-201
ResNet-50
Inception-v2
NASNet-A
ResNet-34
Imagenet Top 1 Accuracy (%)
Number of Parameters (Millions)
84
82
80
78
76
74
0
20
40
60
80
100
120
140
160
180

Transformer & GPT3

Attention is All You Need!

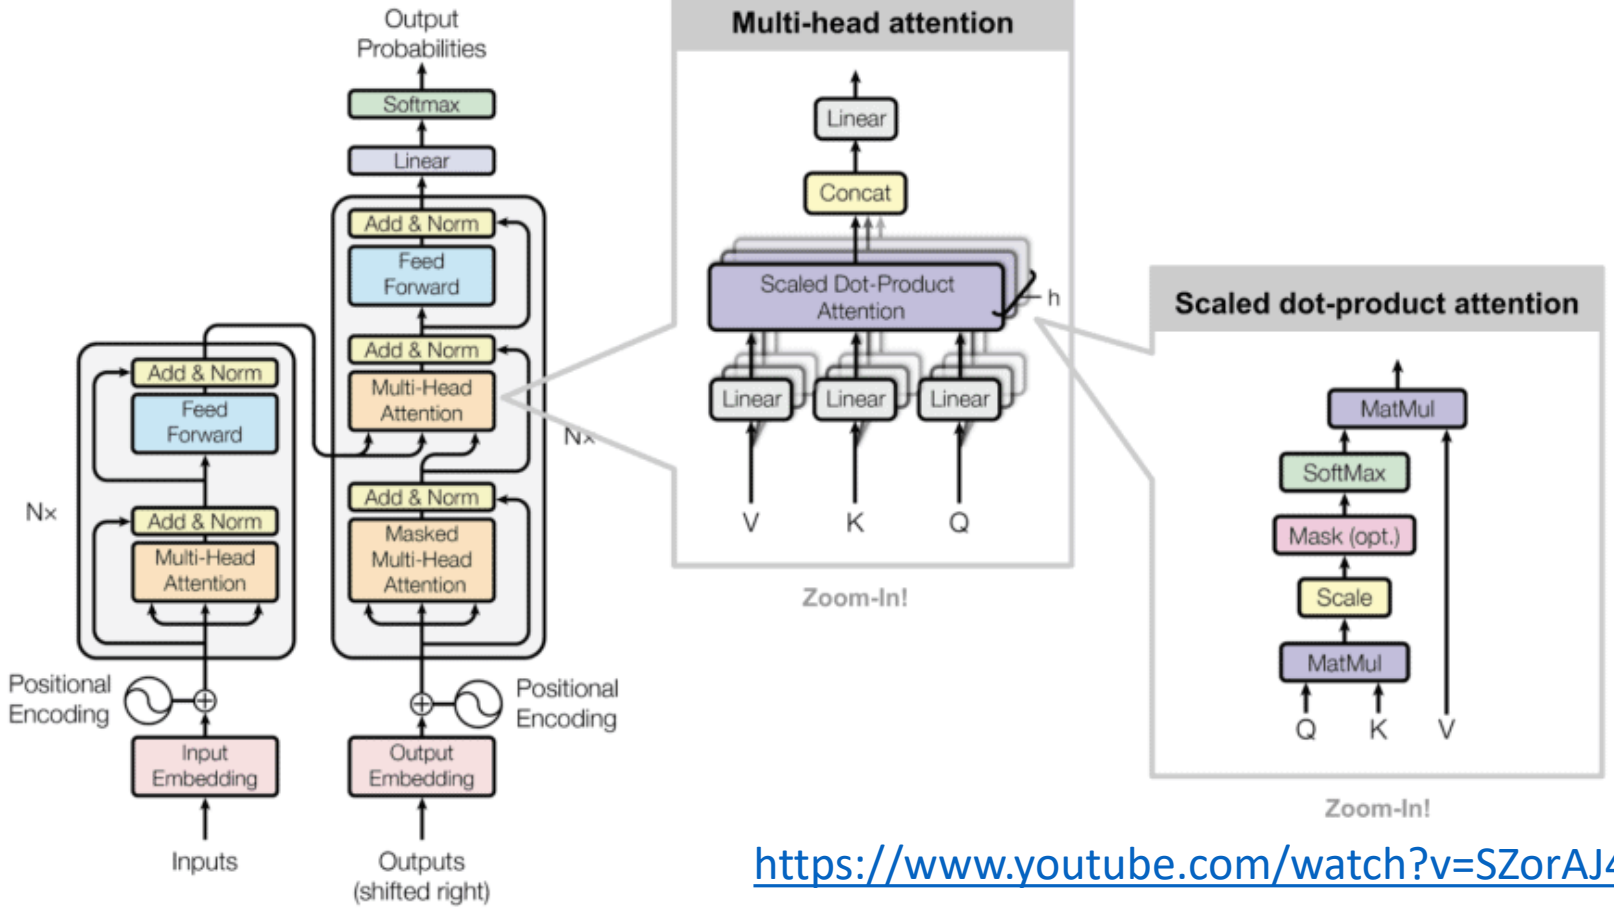

https://www.youtube.com/watch?v=SZorAJ4I-sA

huggingface.co

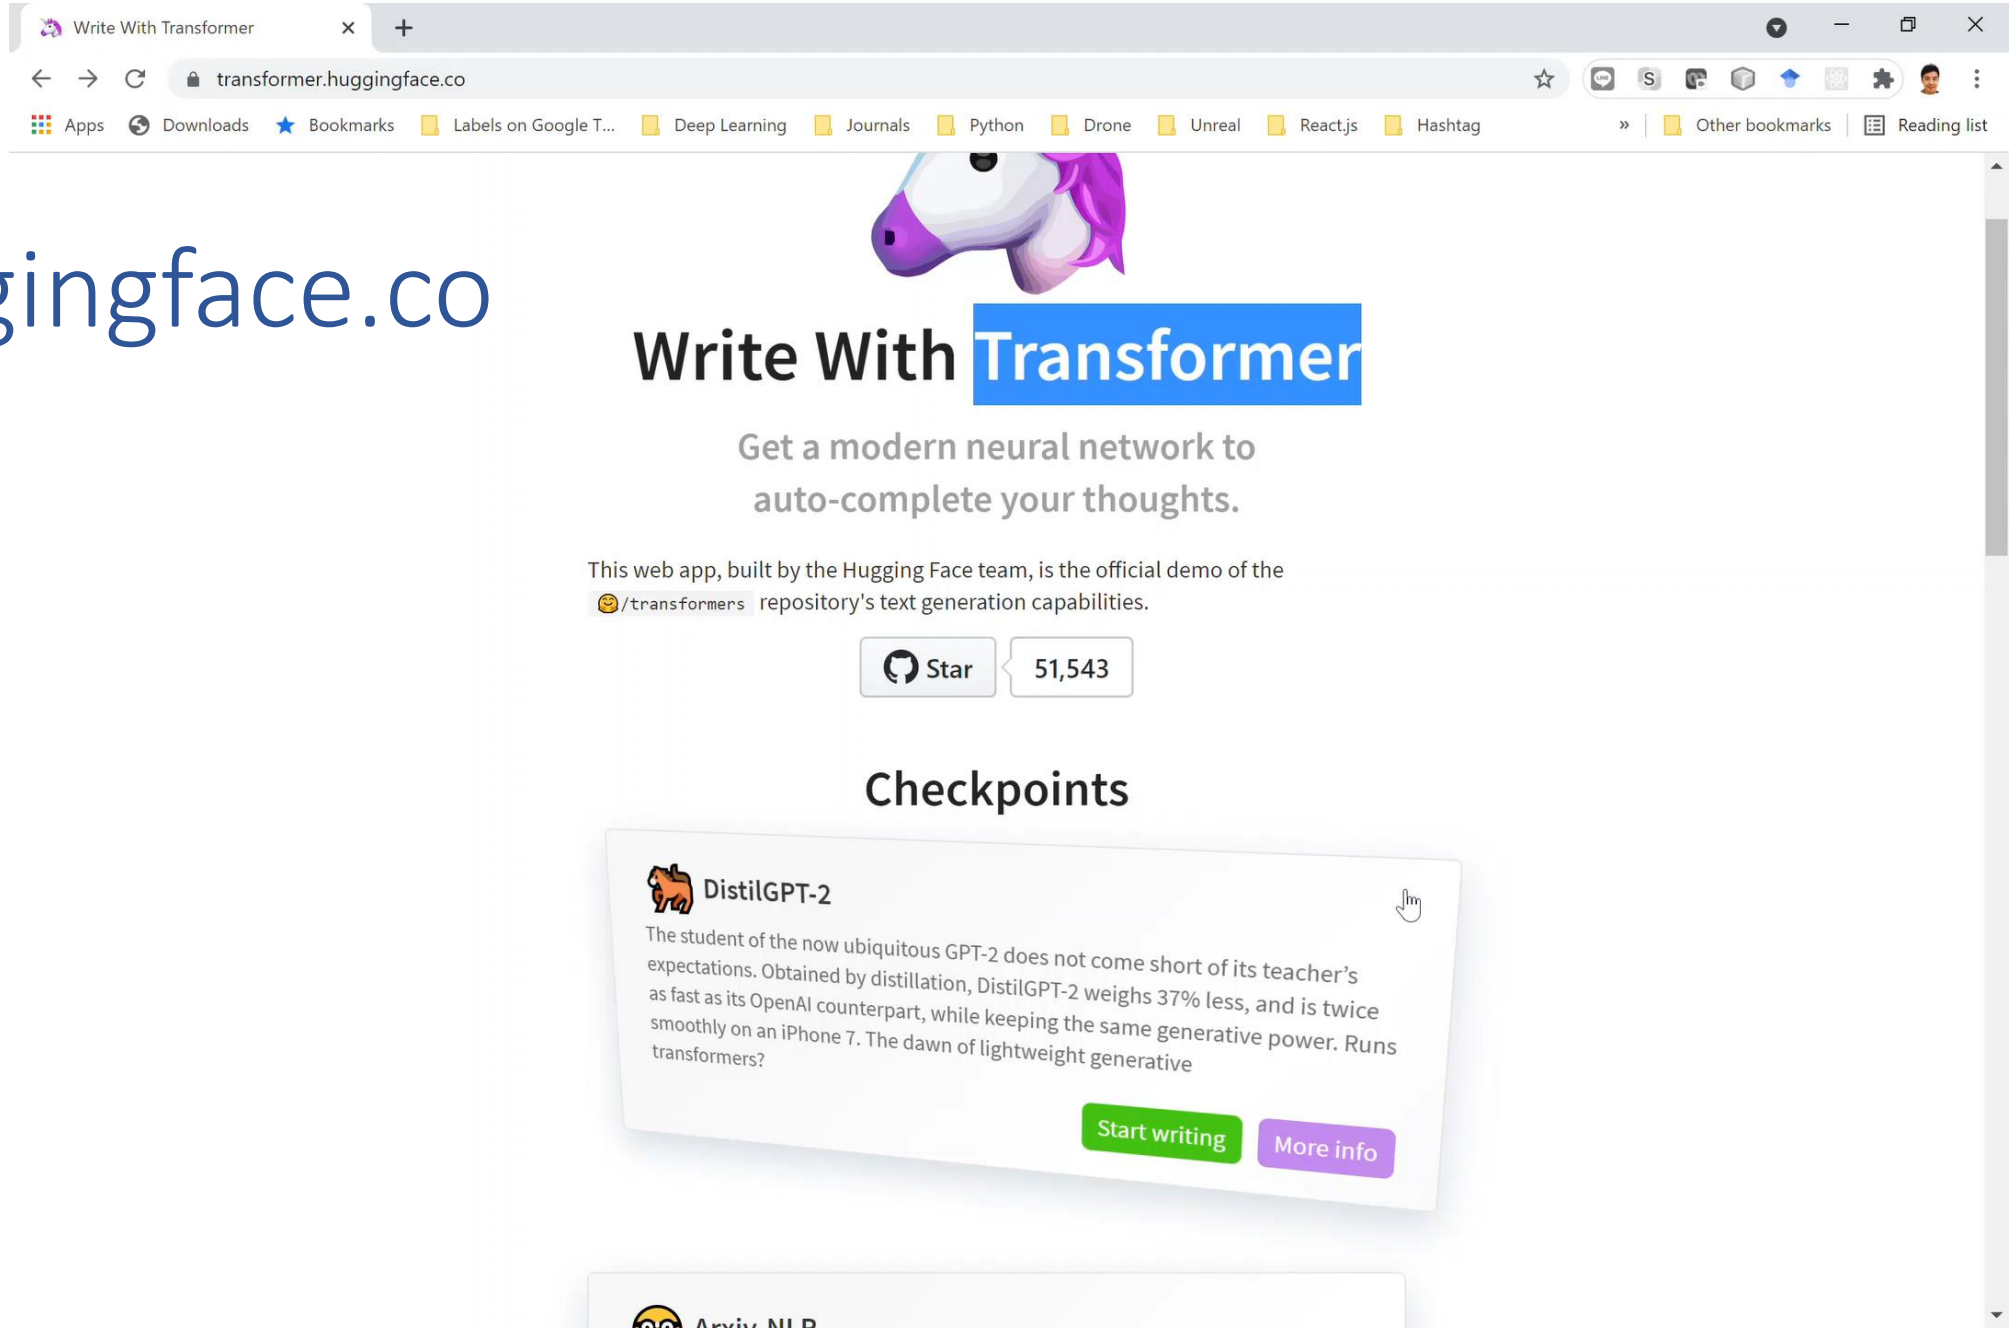

huggingface.co

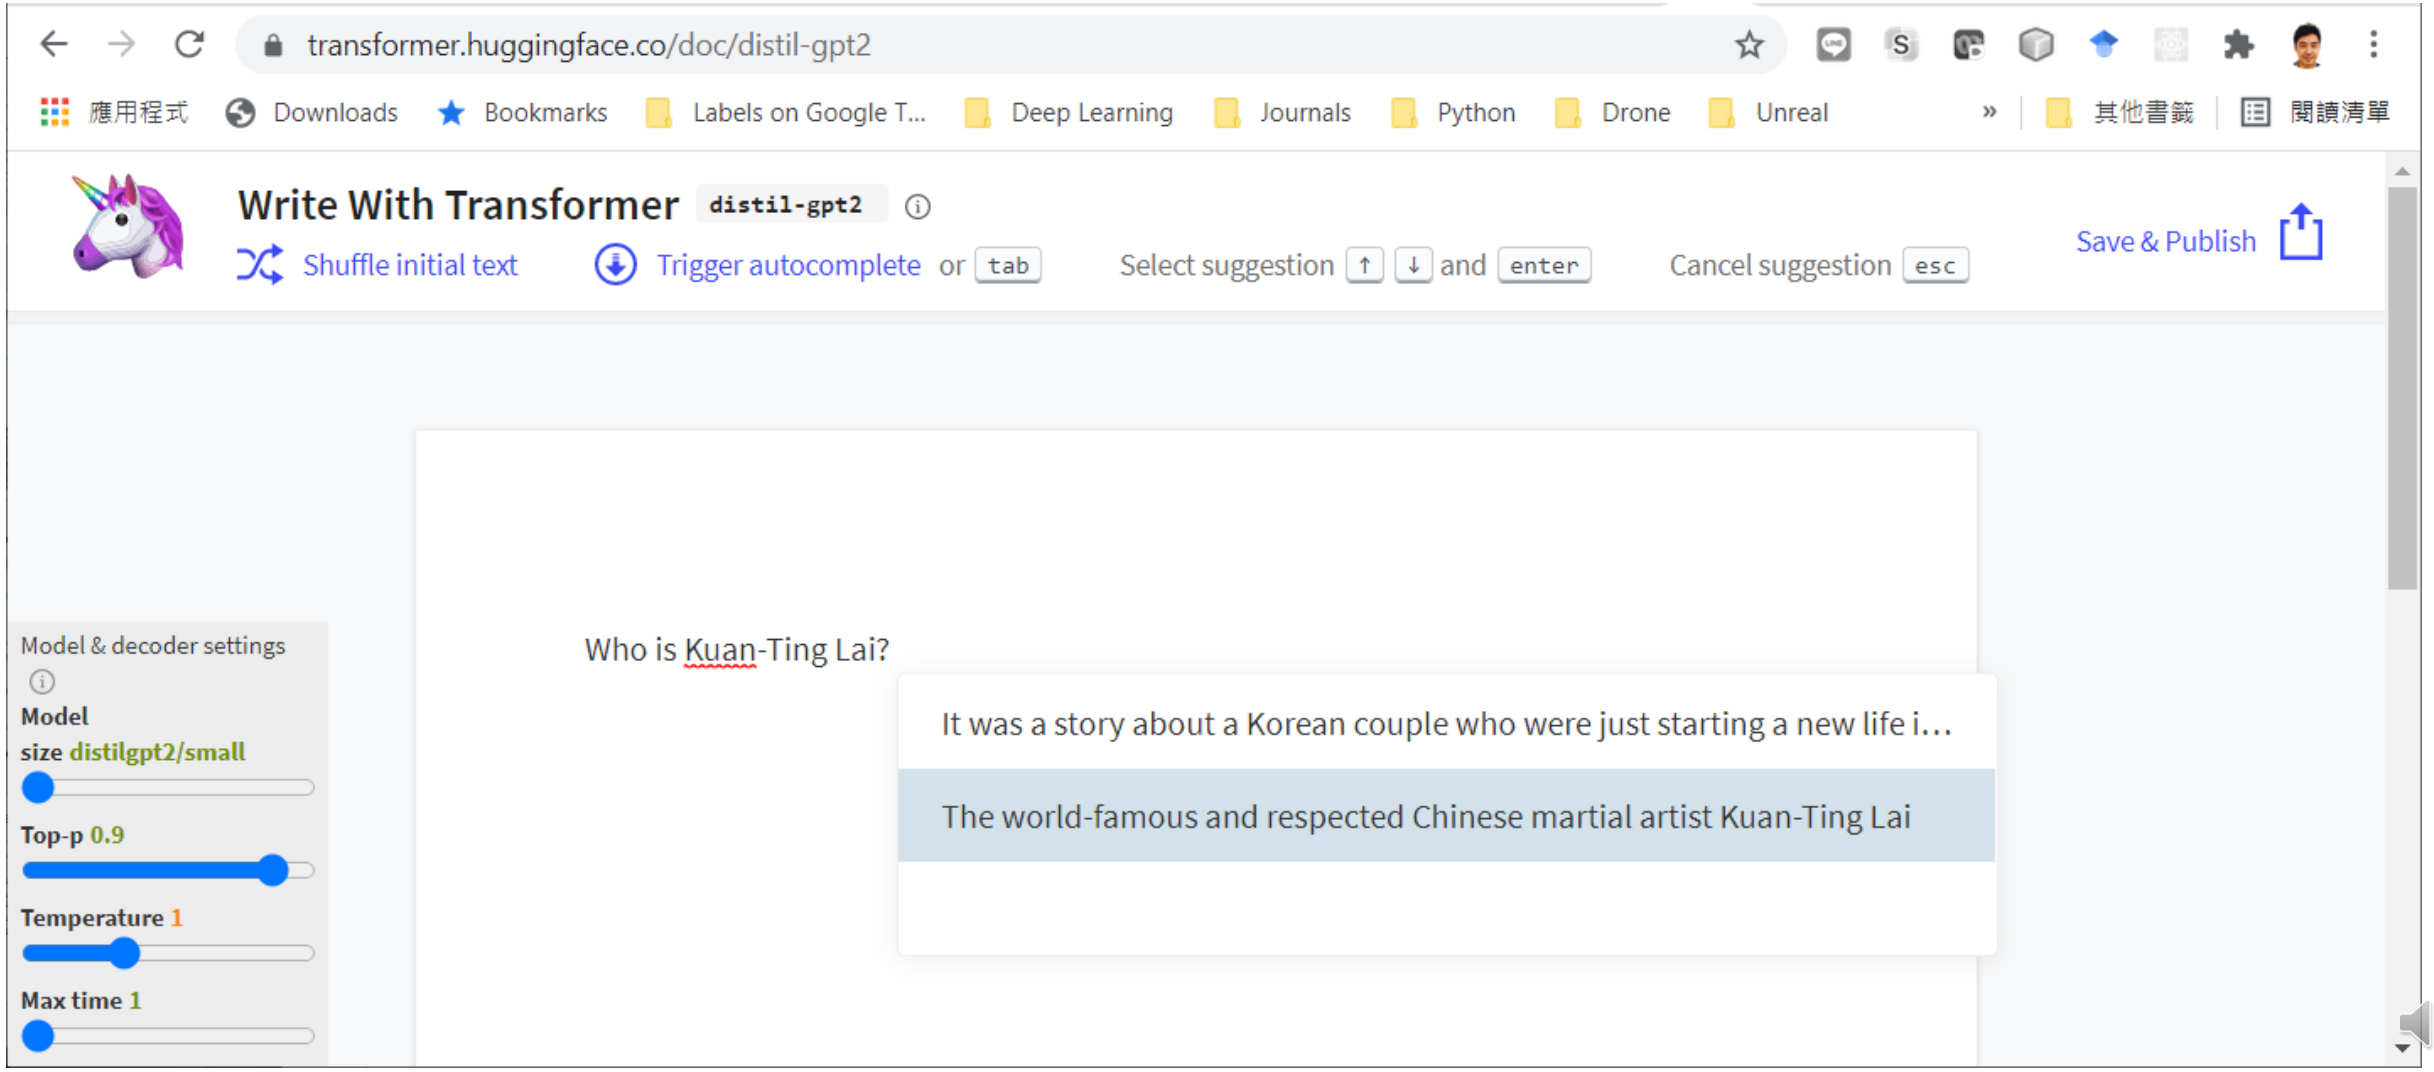

10 Jobs of the Future

Fortune India, 2015

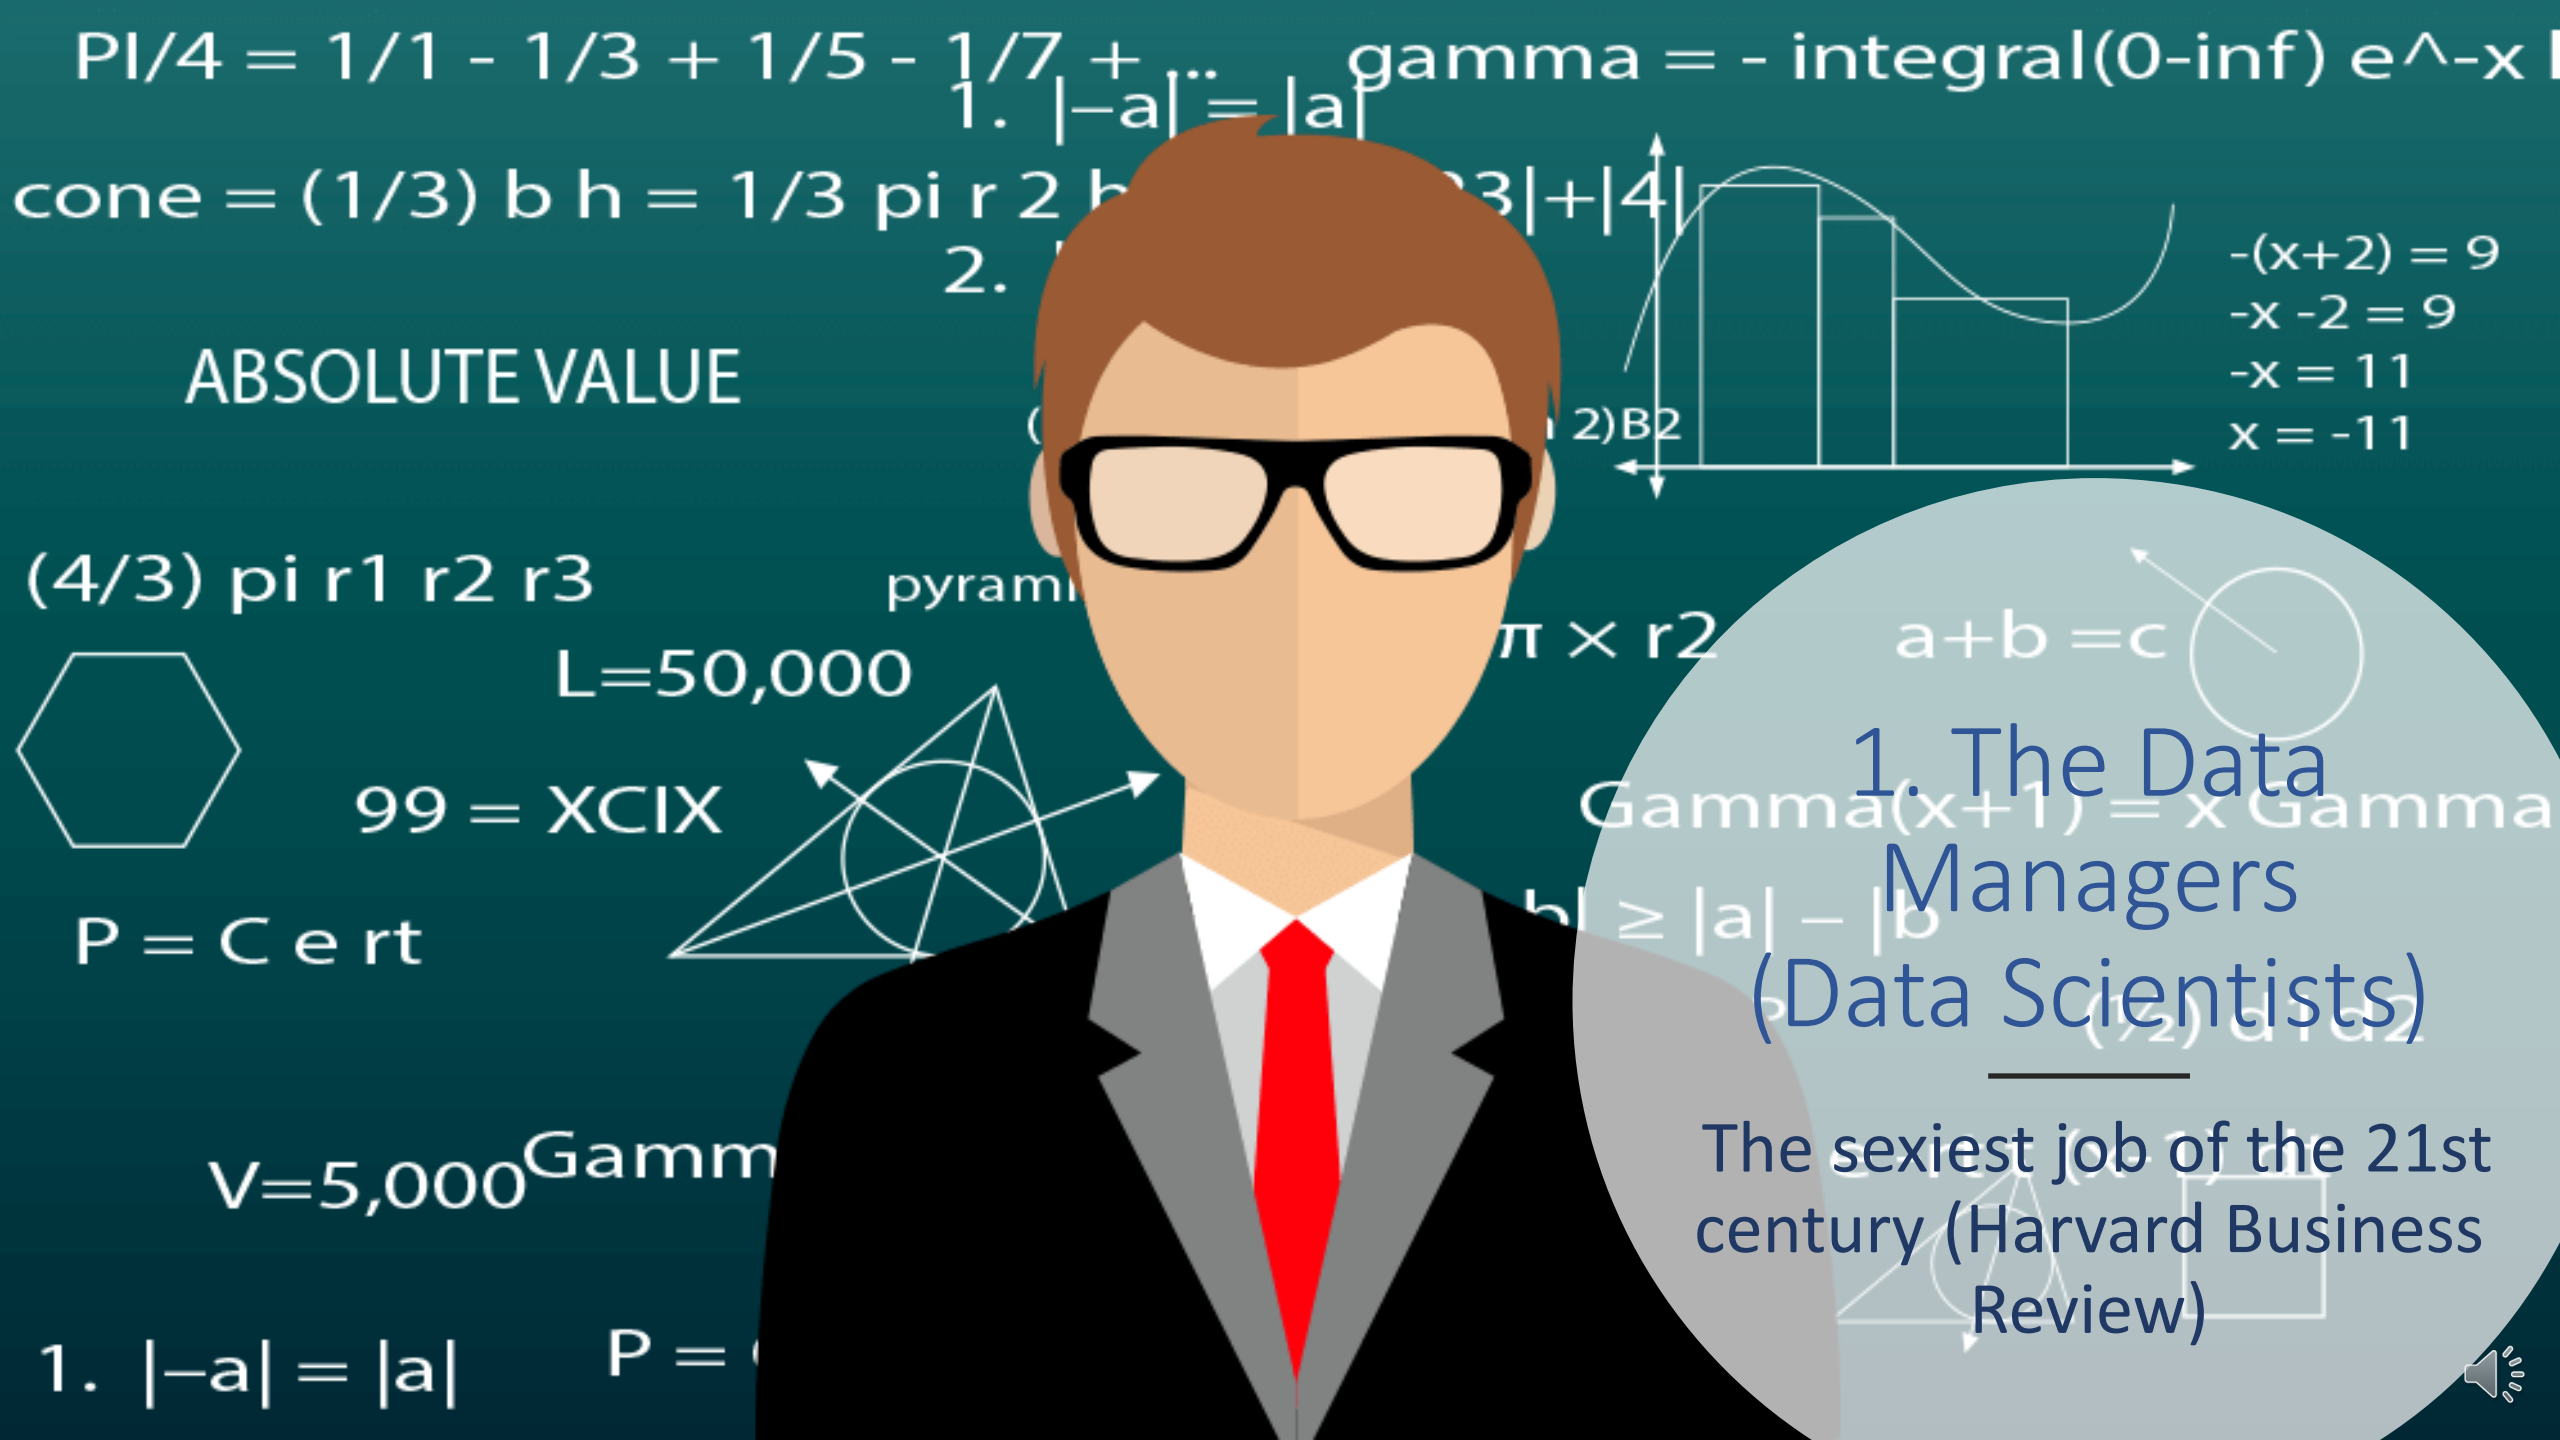
PI/4 = 1/1 - 1/3 + 1/5 - 1/7 + ...
gamma = - integral(0-inf) e^-x
1. |–a| = |a|
cone = (1/3) b h = 1/3 pi r 2
ABSOLUTE VALUE
-(x+2) = 9
-x -2 = 9
-x = 11
x = -11
(4/3) pi r1 r2 r3
L=50,000
99 = XCIX
P = C e rt
a+b =c
Gamma(x+1) = x Gamma
V=5,000
1. |–a| = |a|
1. The Data Managers (Data Scientists)
The sexiest job of the 21st century (Harvard Business Review)

2. The Space Invaders

Demand for satellite engineers is going up, but there's little training available.

3. The Electronic (Digital) Warriors

Defense electronics engineers will be in demand, but supply will be very short.

4. The Drone Operators (Fleet Commanders)

The market is set to open up to commercial operators, but a dearth of qualified engineers could ground it.

5. The New-age Printers

- *3D printing could change manufacturing forever, and all it needs are hardware engineers.*

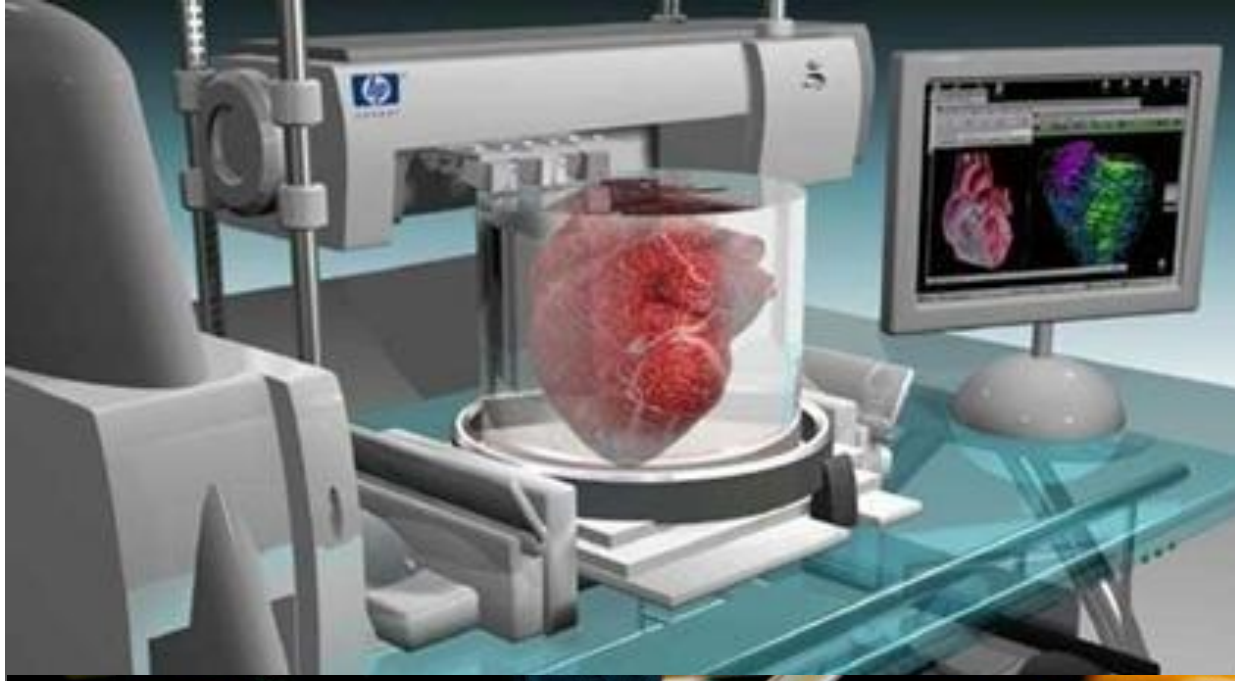

6. The Robotic Engineers

- The robots are coming. Technology research and advisory firm Gartner predicts that more than 30% of all jobs will be replaced by robots and smart machines by 2025, and Ray Kurzweil, director of engineering at Google, anticipates that human thinking will become hybrid (biological and artificial intelligence) by the 2030s.

7. The Cyberspace Ninjas (Cybersecurity)

Prevention of data theft is just one of their jobs.

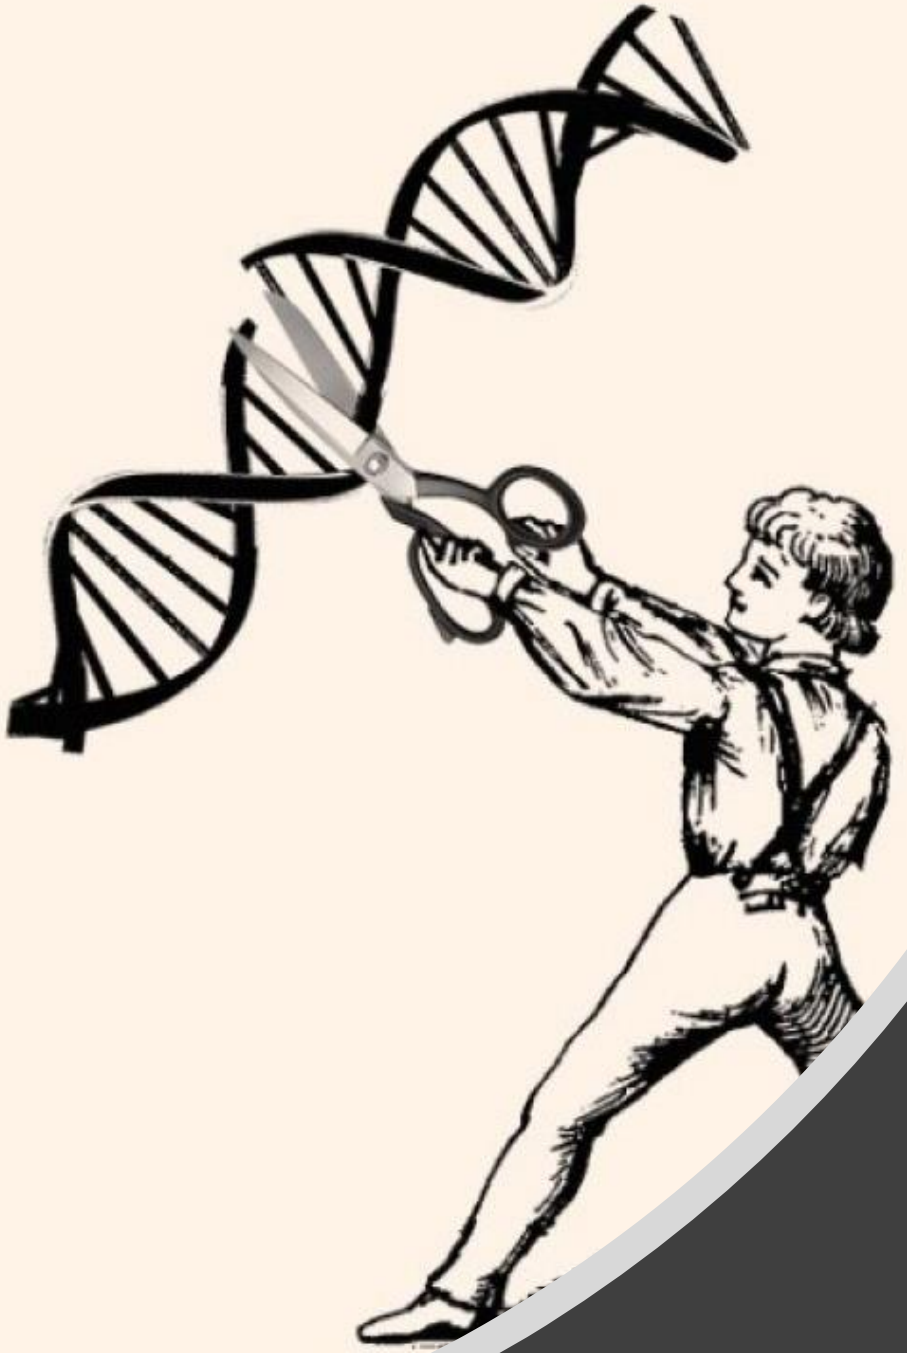

8. The Gene Splicers

9. The Green Warriors

- *India's big push towards clean energy is changing the job market.*

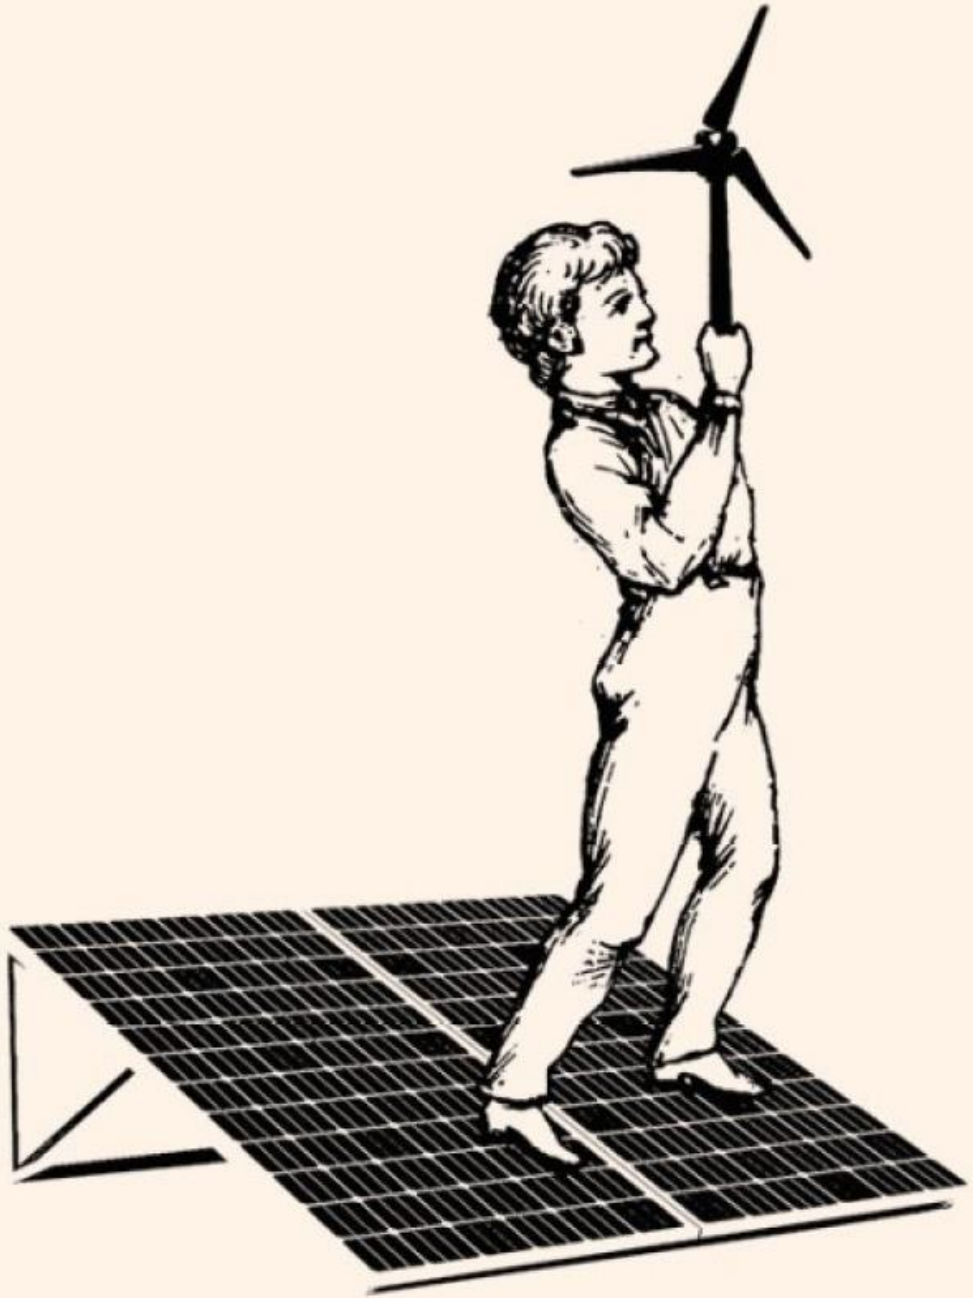

10. The Medical Mechanics

AI CT Scan

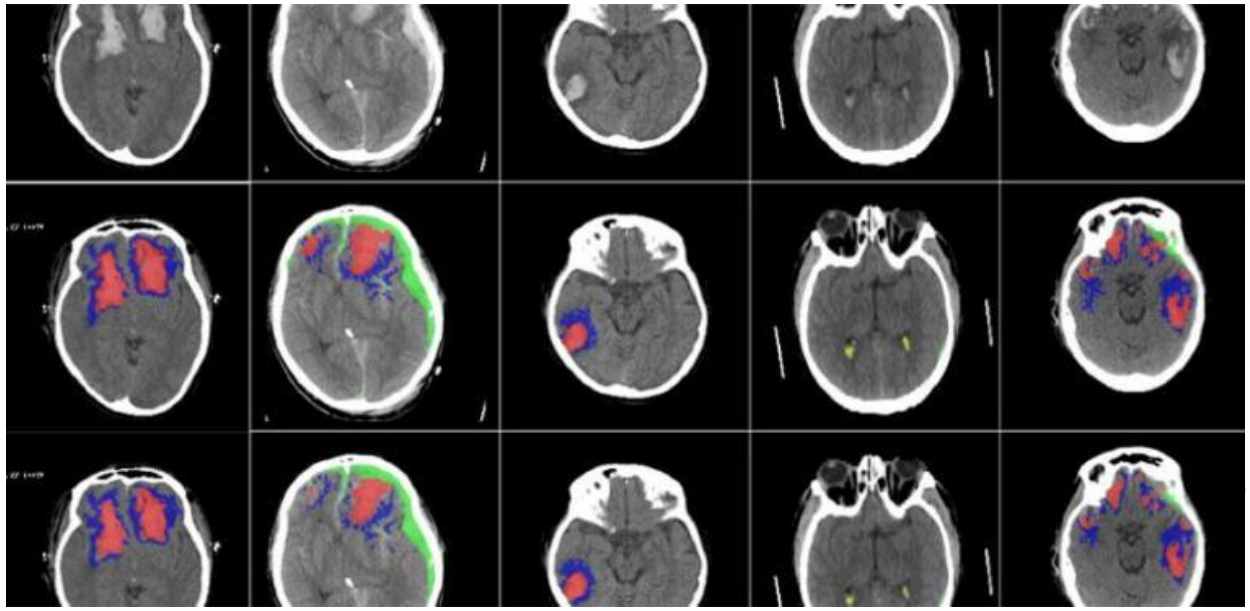

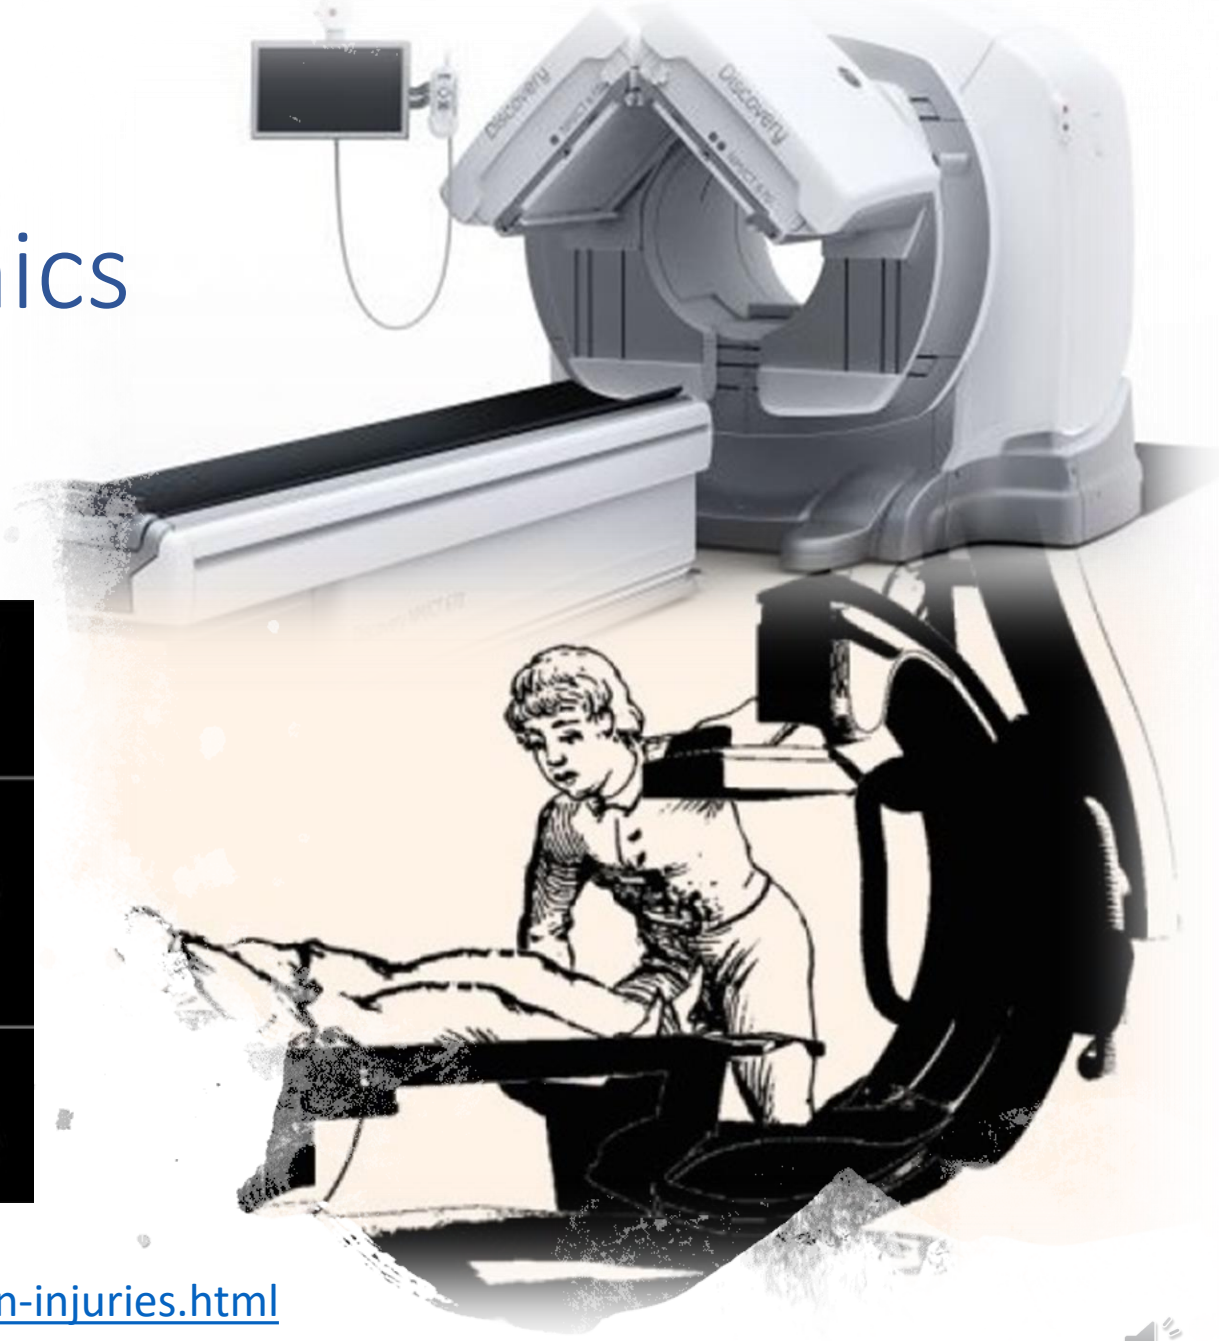

https://tectales.com/ai/ai-used-to-identify-different-types-of-brain-injuries.html

References

1. https://www.buzzfeed.com/kasiagalazka/science-fiction-things-that-actually-exist-now
2. https://www.geek.com/movies/10-movies-that-helped-create-real-technology-1740036/
3. https://www.gadgetsnow.com/slideshows/8-sci-fi-movie-technologies-that-are-real-now/Video-calling/photolist/52869590.cms
4. What is backpropagation really doing? https://www.youtube.com/watch?v=Ilg3gGewQ5U
5. http://www.andreykurenkov.com/writing/ai/a-brief-history-of-neural-nets-and-deep-learning/
6. https://pmirla.github.io/2016/08/16/AI-Winter.html
7. https://tw.saowen.com/a/6cdc2f1279016e566832bb1234e06d321992dd1fabcdf4a2e0a3e16fc0dc09dc
8. https://tectales.com/ai/ai-used-to-identify-different-types-of-brain-injuries.html
9. Transformers, explained: Understand the model behind GPT, BERT, and T5
10. https://transformer.huggingface.co/