Artificial Intelligence

賴冠廷教授
Prof. Kuan-Ting Lai
台北科技大學電子工程系
2020/3/20

Artificial Intelligence

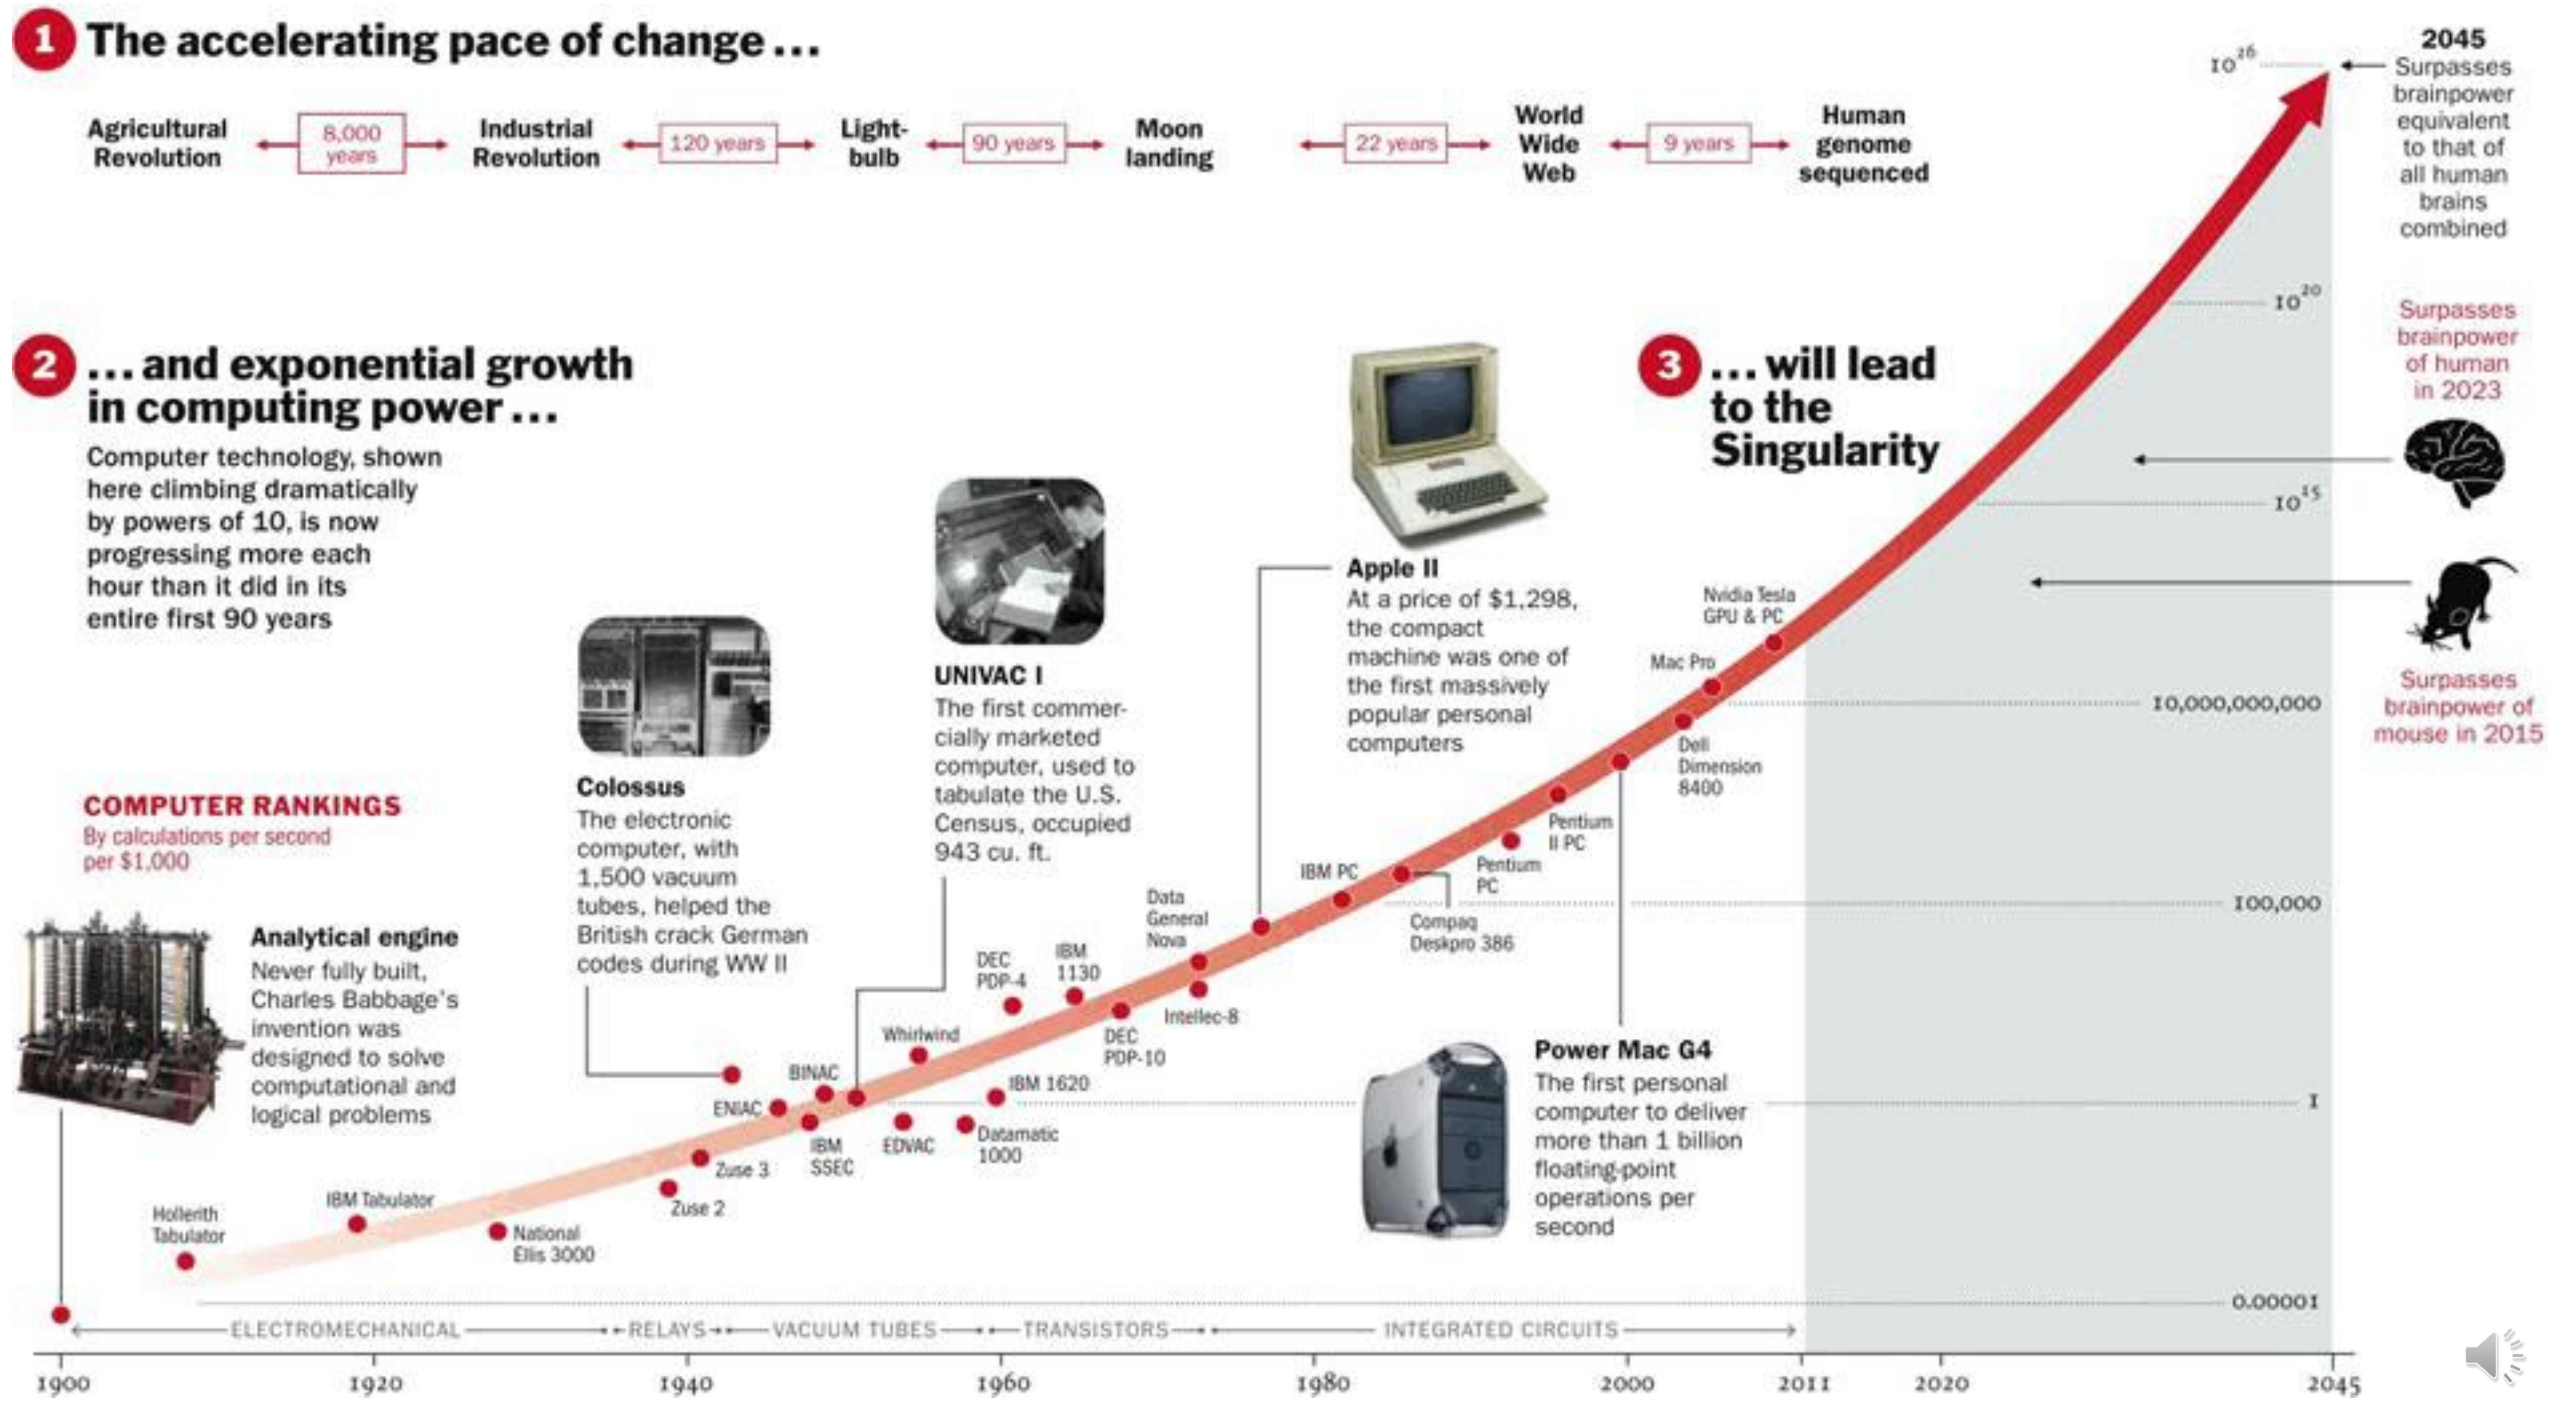
1 The accelerating pace of change . . .
Agricultural Revolution
8,000 years
Industrial Revolution
120 years
Light-bulb
90 years
Moon landing
22 years
World Wide Web
9 years
Human genome sequenced
2 . . . and exponential growth in computing power . . .
Computer technology, shown here climbing dramatically 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
Colossus
The electronic computer, with 1,500 vacuum tubes, helped the British crack German codes during WW II
UNIVAC I
The first commercially marketed computer, used to tabulate the U.S. Census, occupied 943 cu. ft.
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
Hollerith Tabulator
IBM Tabulator
National Ellis 3000
Zuse 2
Zuse 3
ENIAC
IBM SSEC
BINAC
EDVAC
Whirlwind
Datamatic 1000
IBM 1620
DEC PDP-4
IBM 1130
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
10^{26}
10^{20}
10^{15}
10,000,000,000
100,000
1
0.00001
ELECTROMECHANICAL
RELAYS
VACUUM TUBES
TRANSISTORS
INTEGRATED CIRCUITS
1900
1920
1940
1960
1980
2000
2011
2020
2045

120 Years of Moore's Law

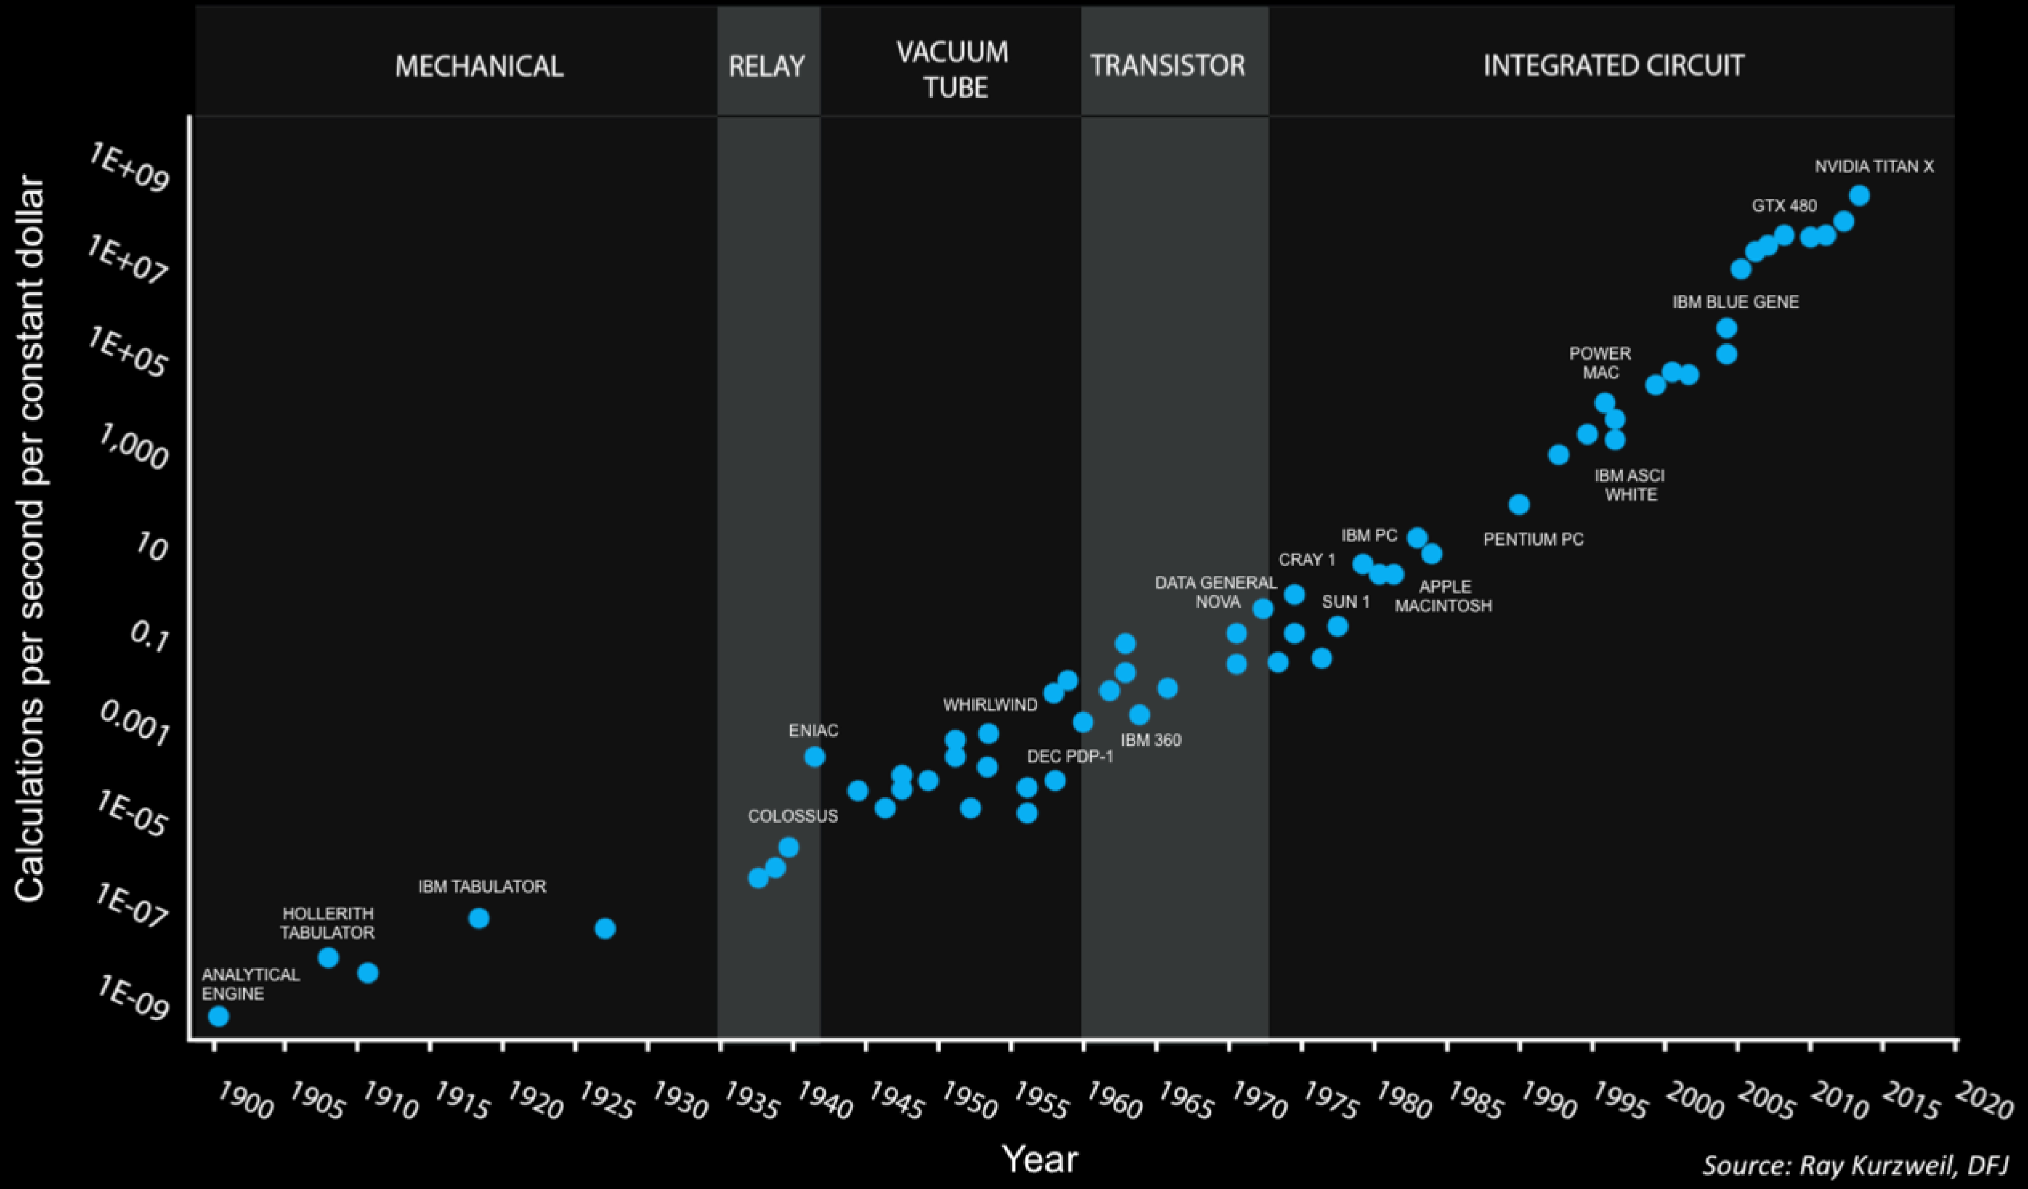

Source: Ray Kurzweil, DFJ

The Singularity

I. J. Good

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Humans, who are limited by slow biological evolution, couldn't compete and would be superseded.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

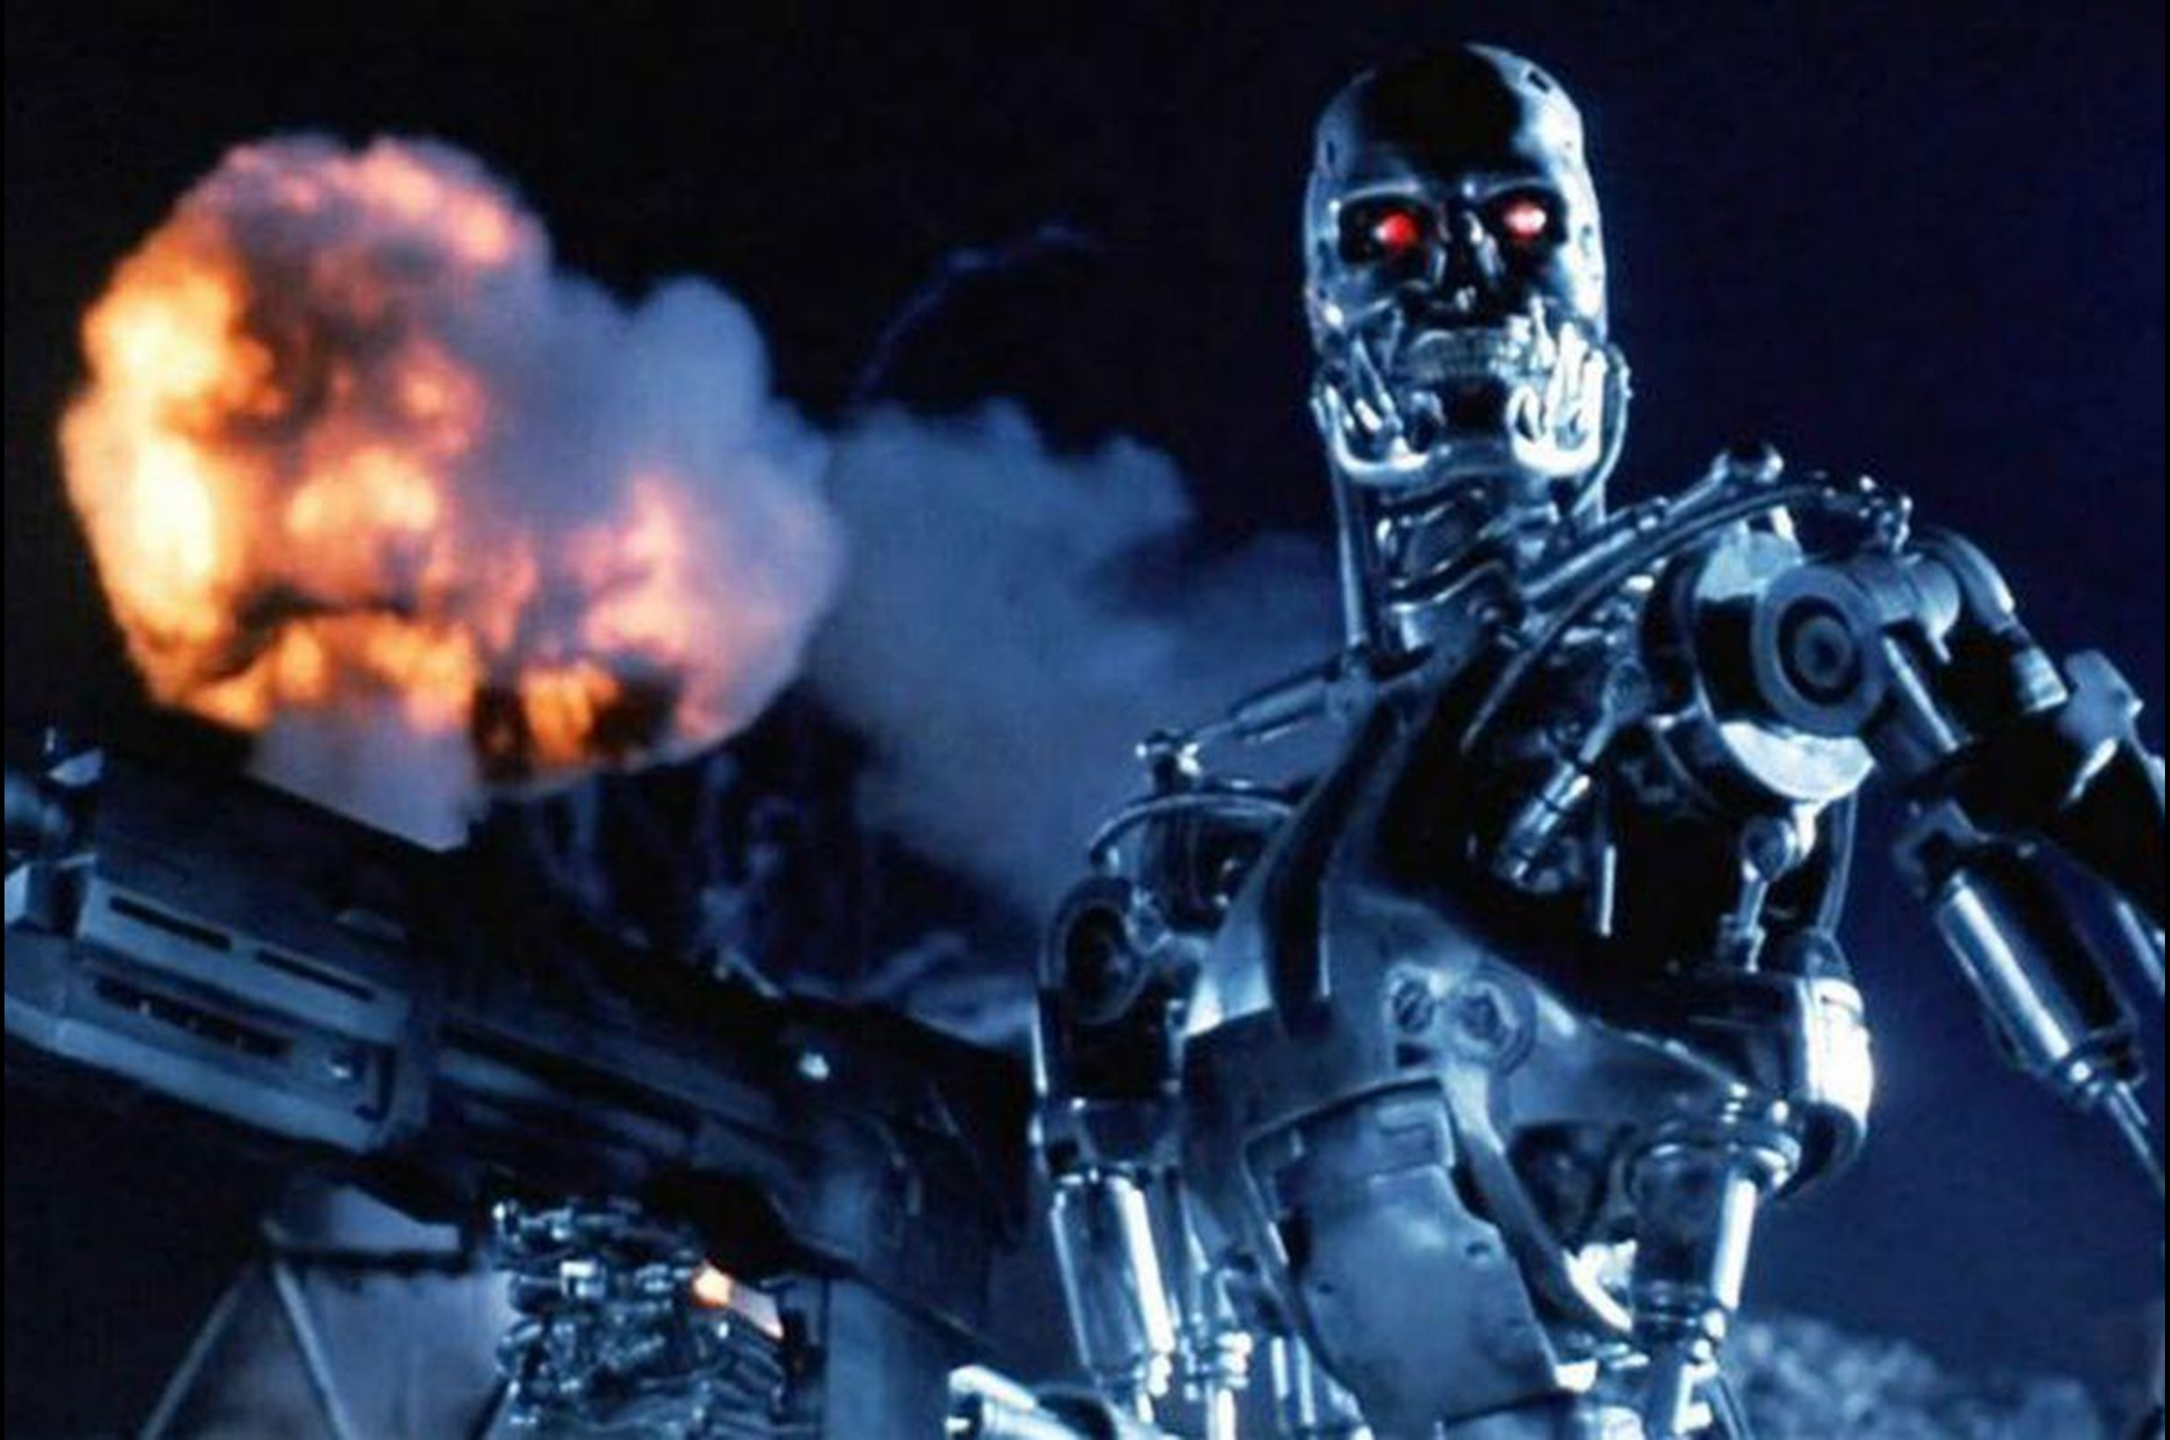

So, what is AI?

Machine Learning

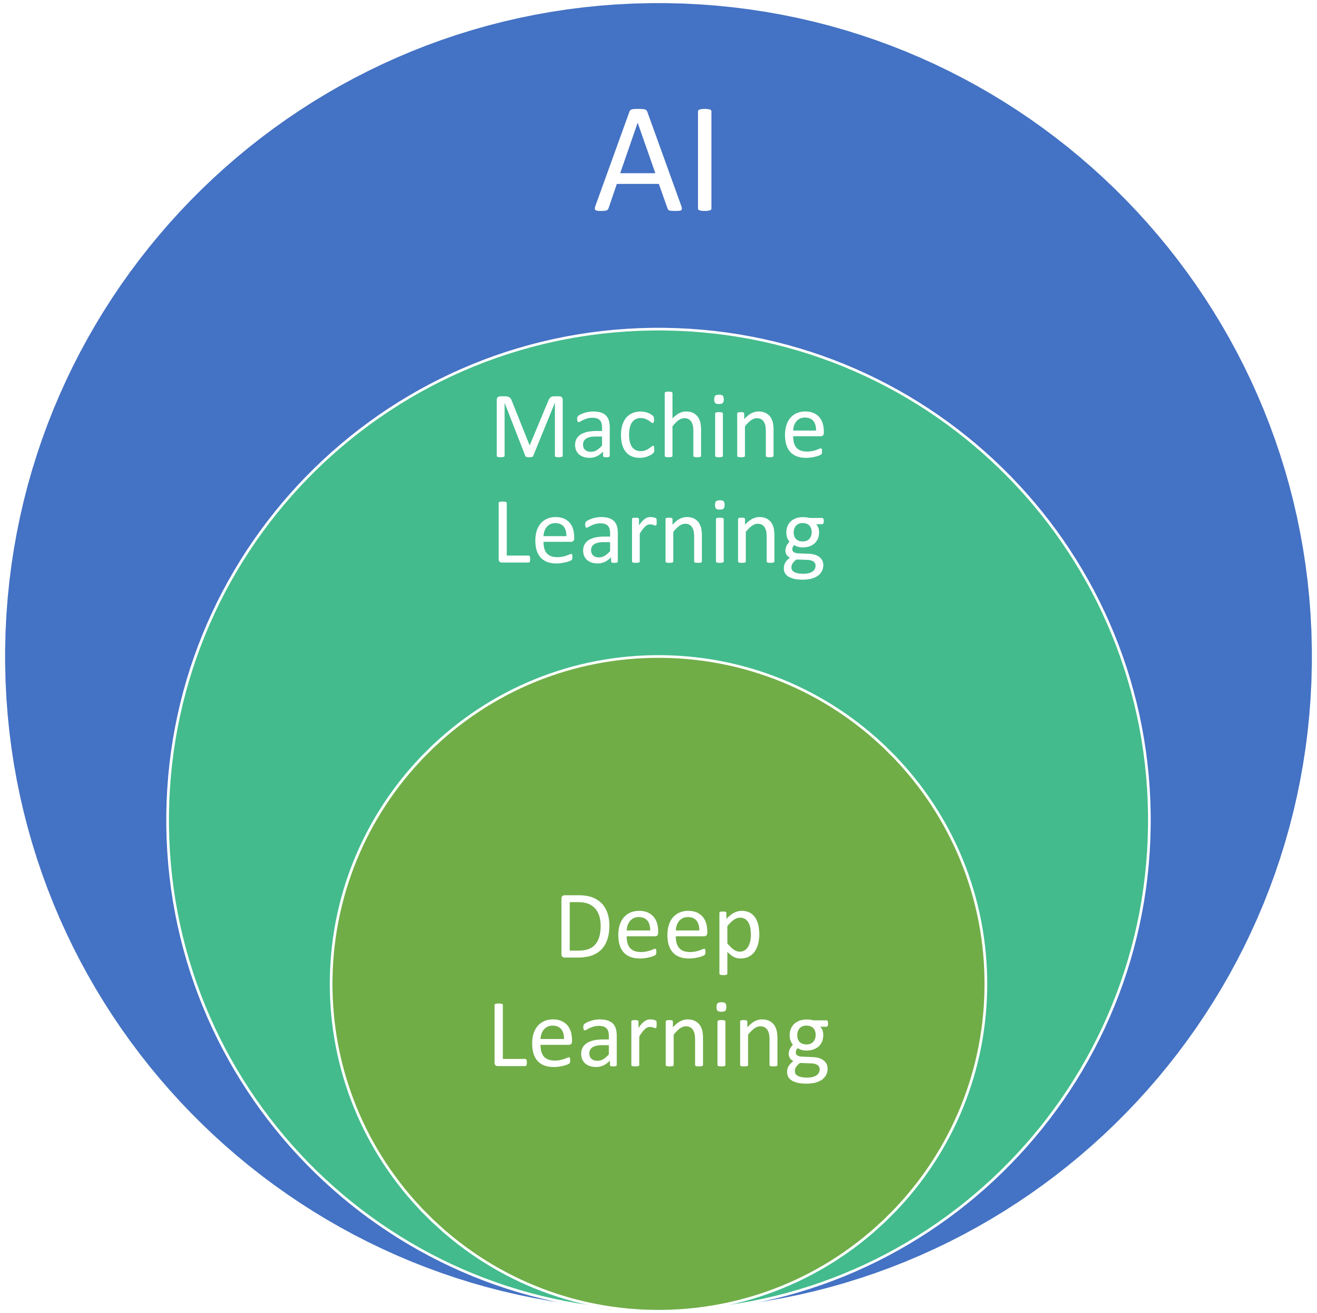
AI
Machine Learning
Deep Learning

Machine Learning (Statistical Learning)

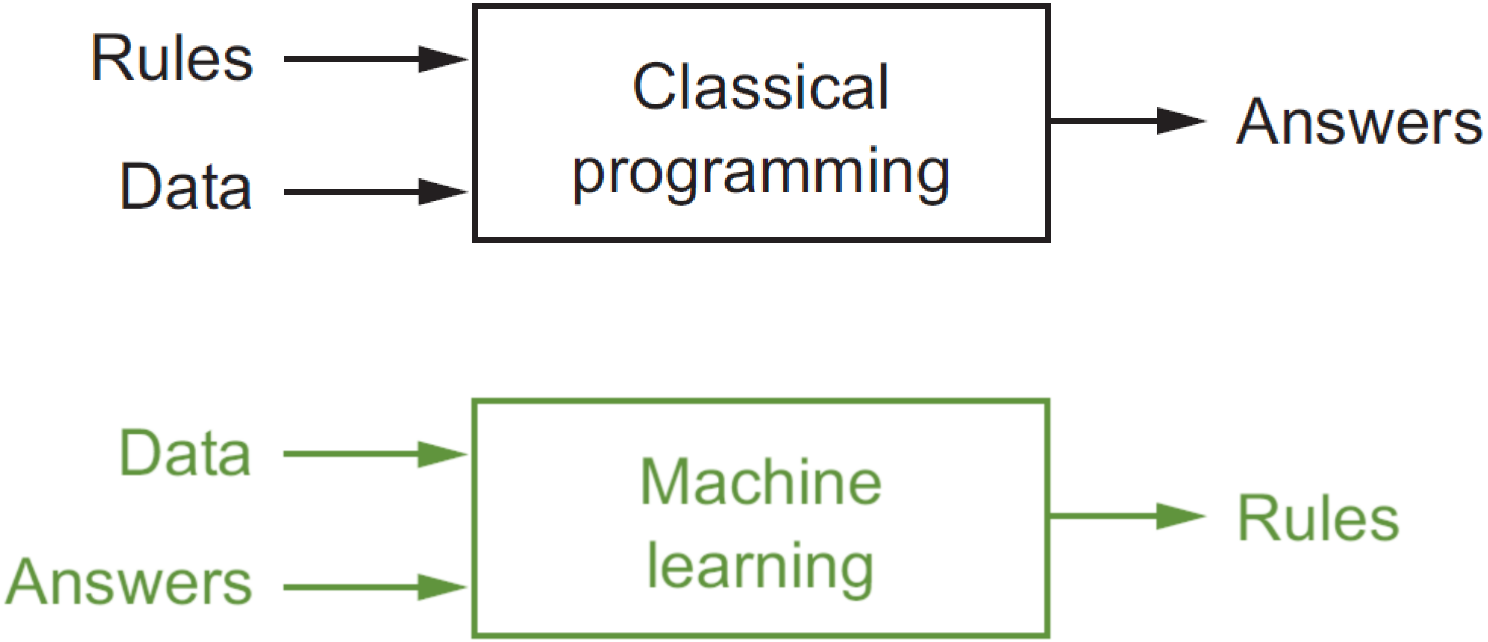

Francois Chollet, “Deep Learning with Python,” Manning, 2017

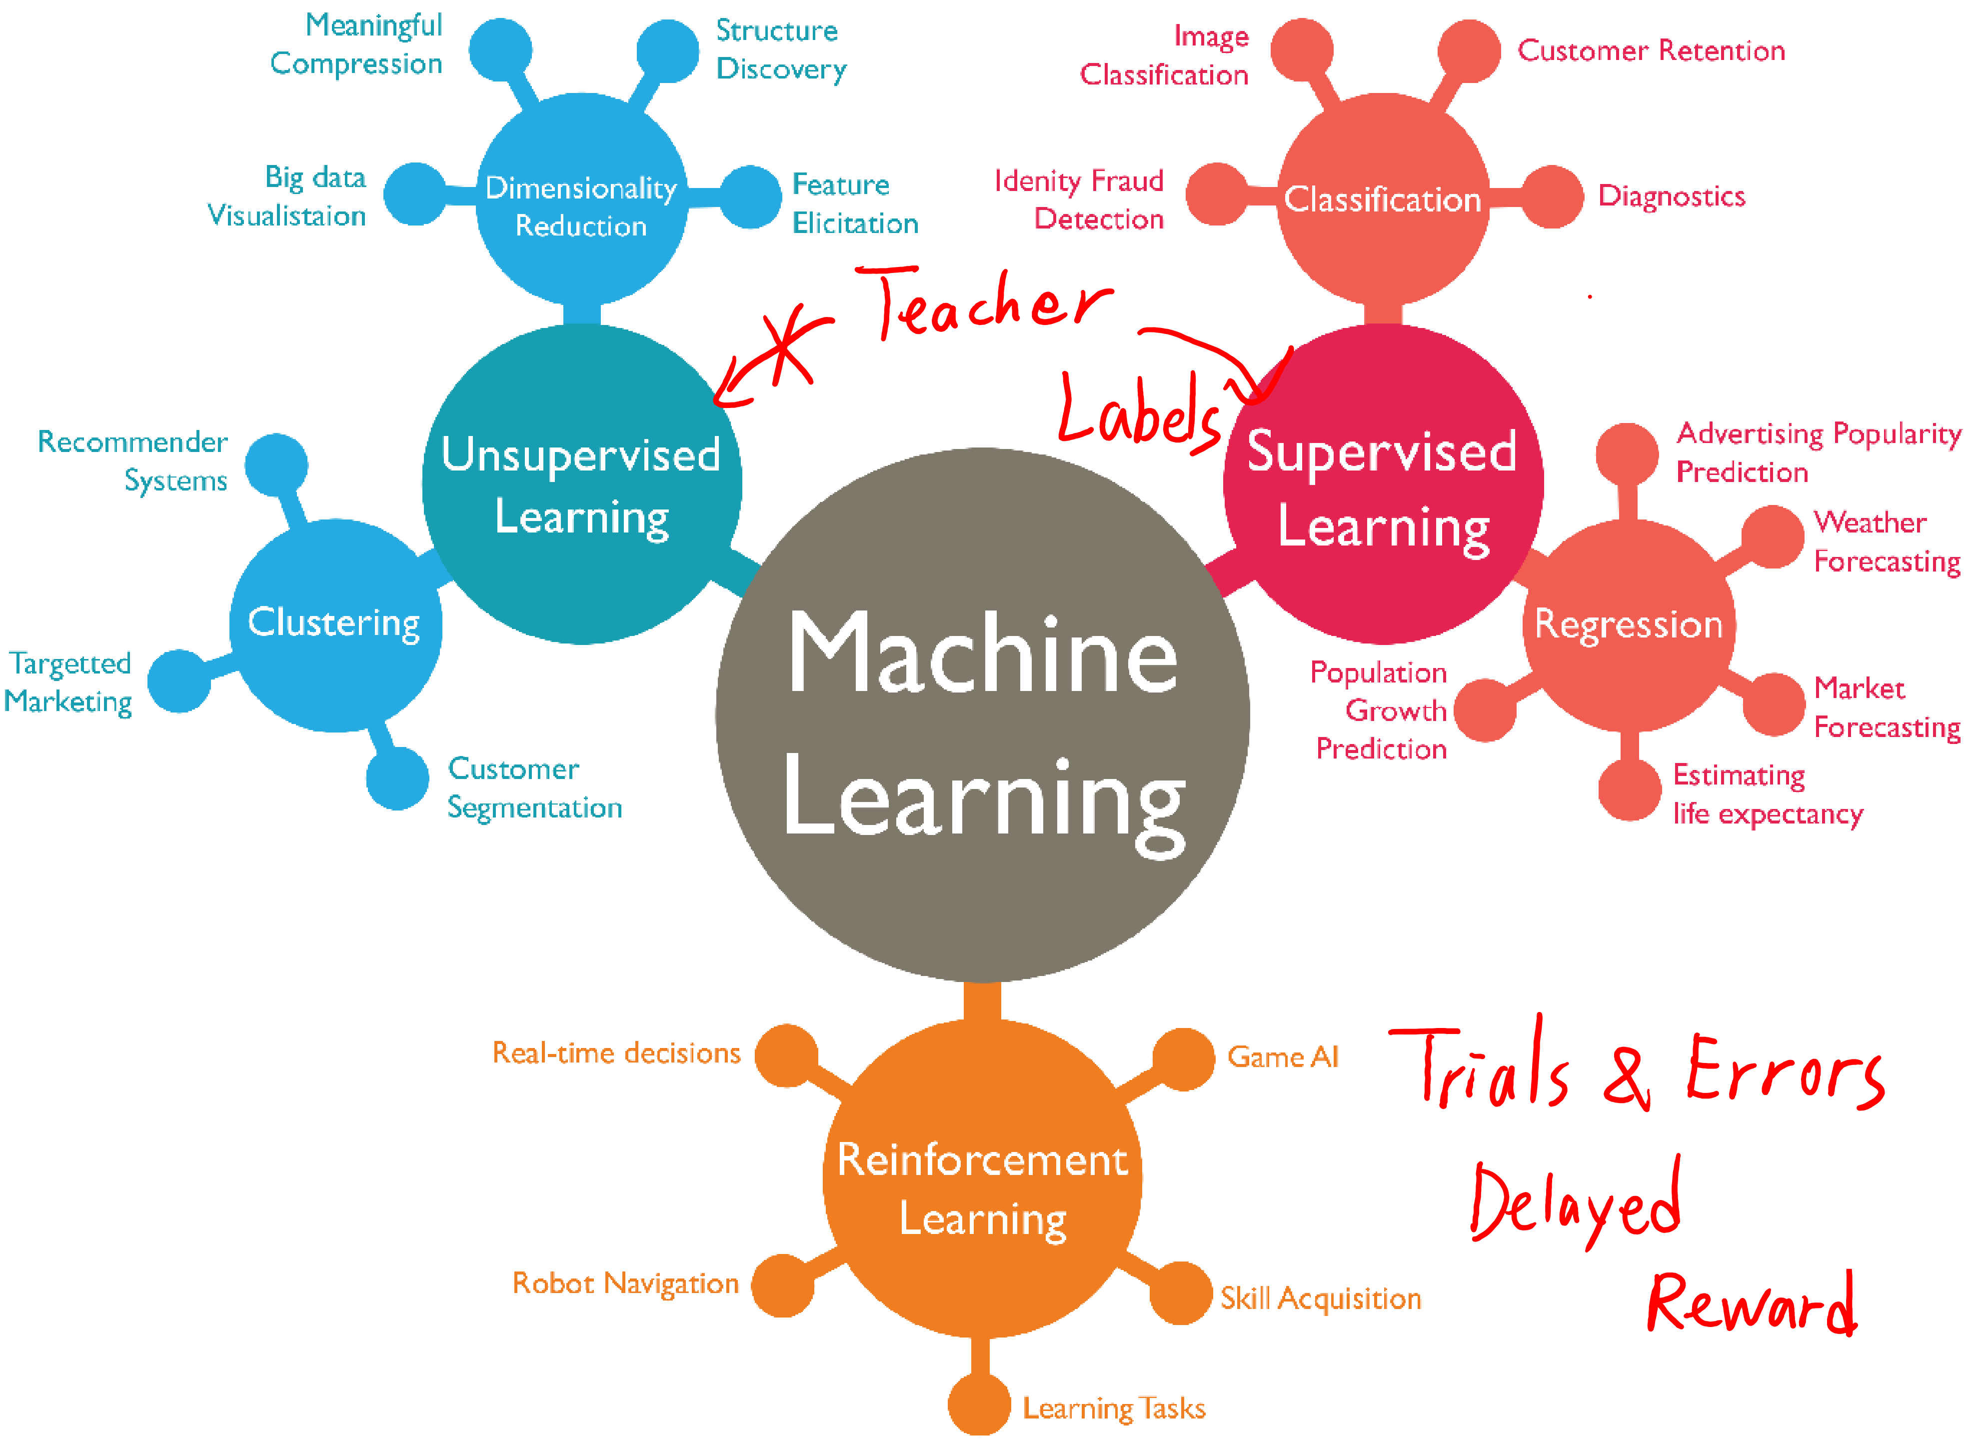
Machine Learning
Unsupervised Learning
Dimensionality Reduction
Meaningful Compression
Structure Discovery
Big data Visualistaion
Feature Elicitation
Clustering
Recommender Systems
Targetted Marketing
Customer Segmentation
Supervised Learning
Classification
Image Classification
Customer Retention
Idenity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

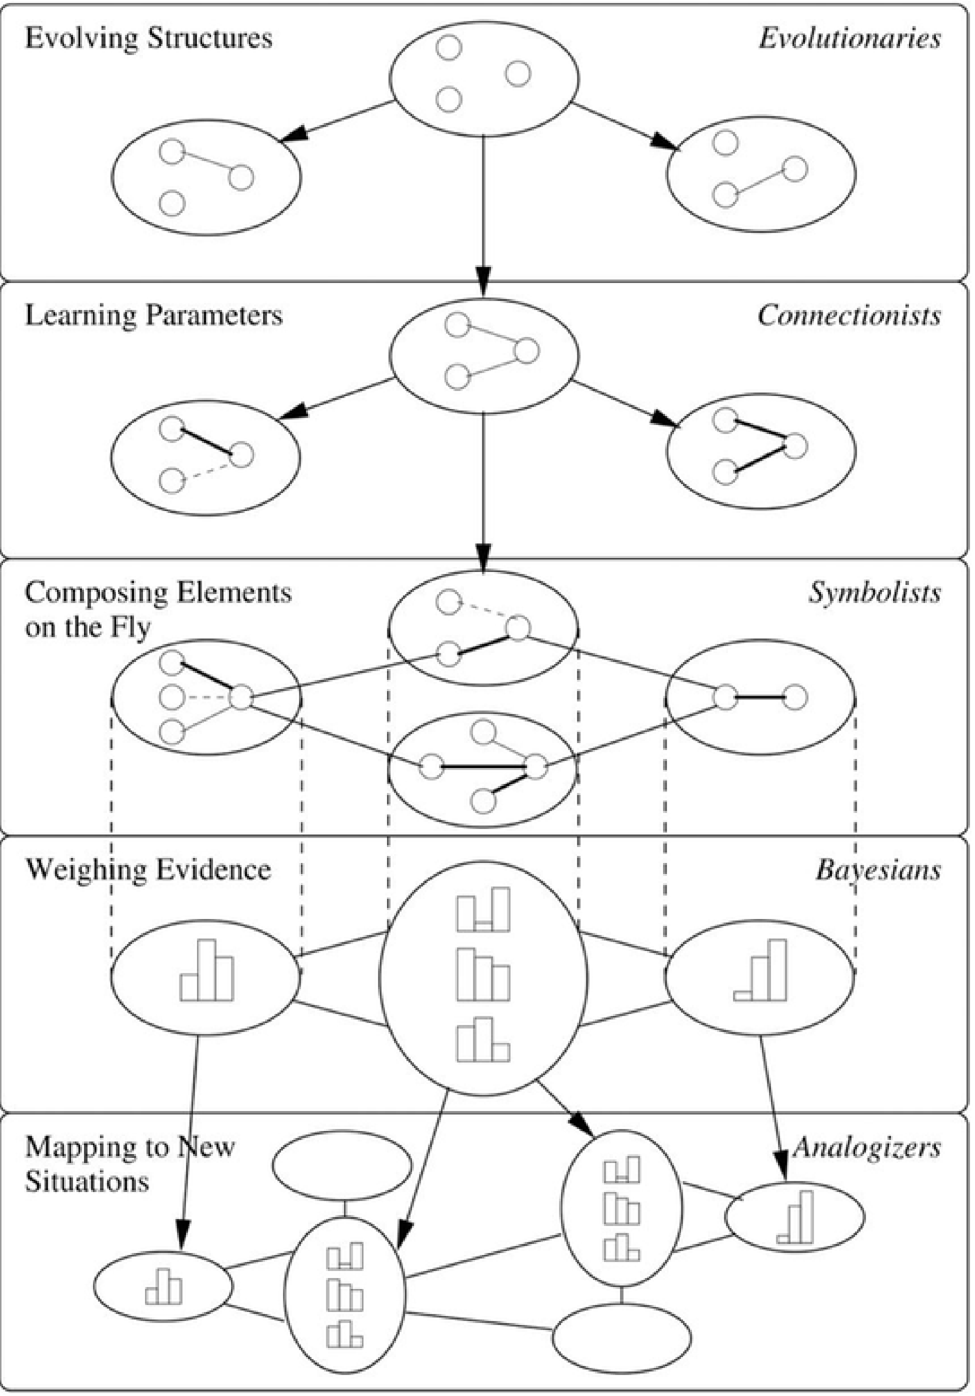

Five Tribes of Machine Learning

- 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

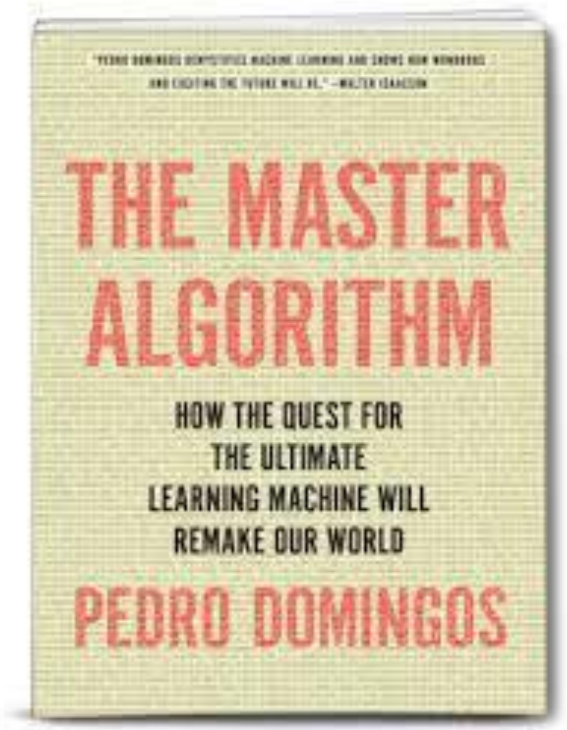

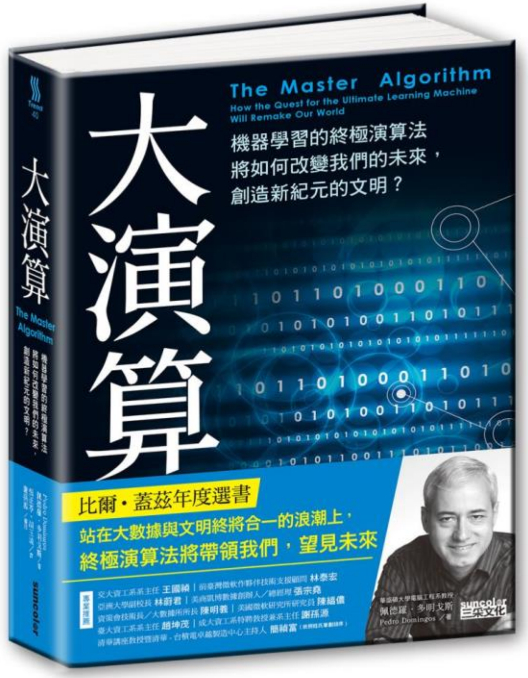

All Algorithms can be Reduced to 3 Operations

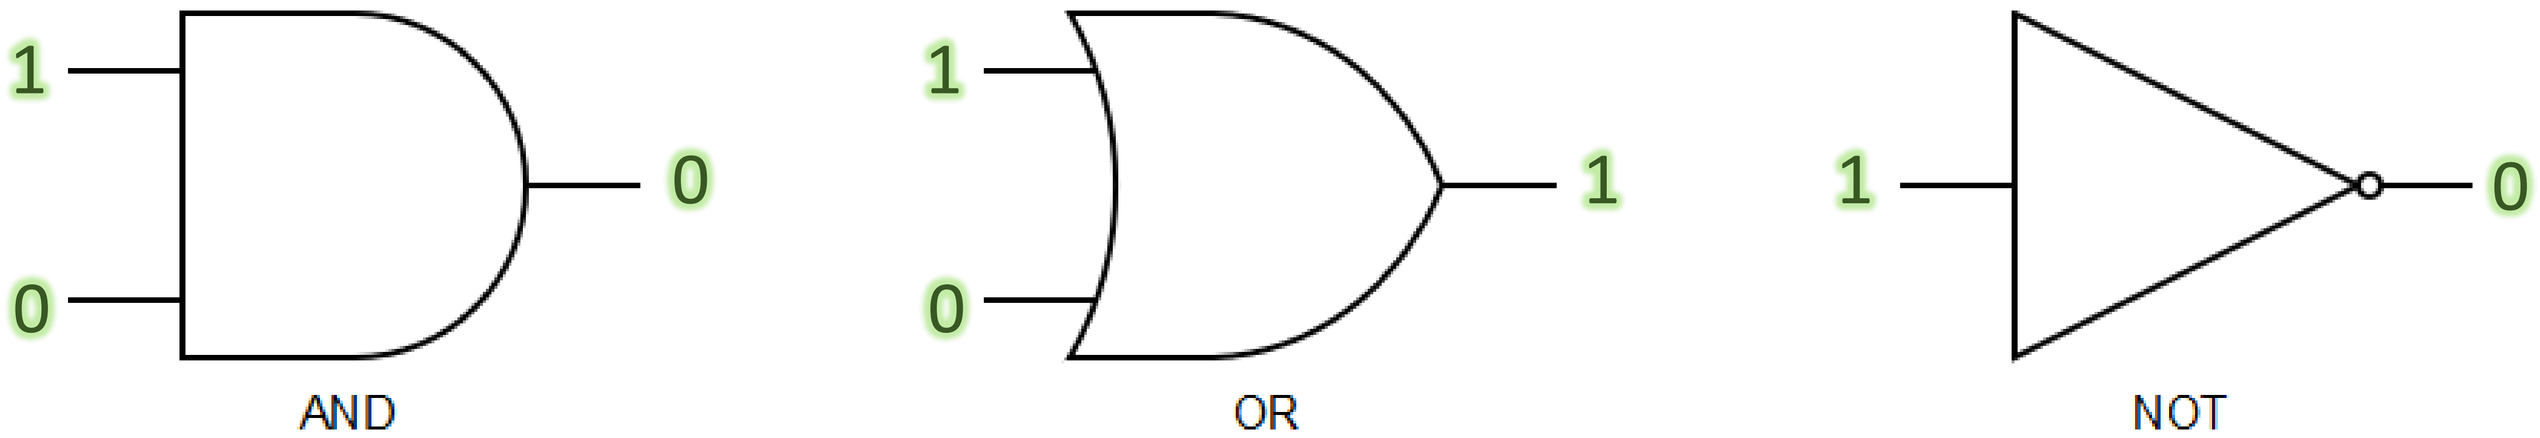

XOR

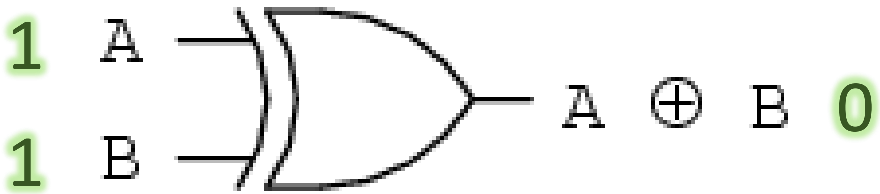

. . . *is equivalent to* . . .

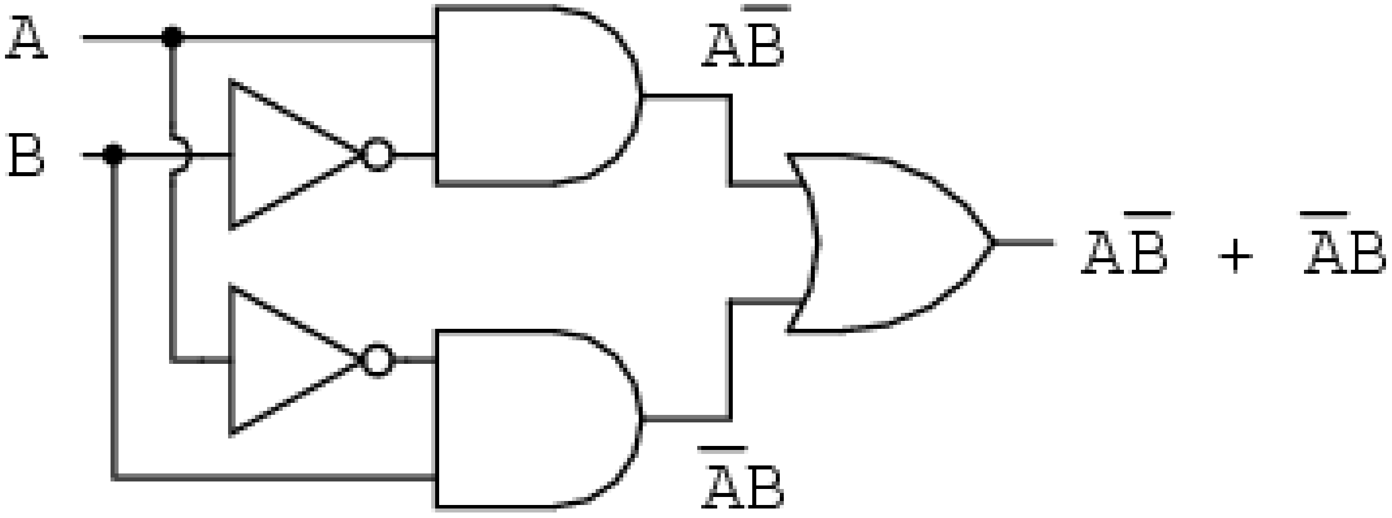

$$\mathbf{A \oplus B = A\overline{B} + \overline{A}B}$$

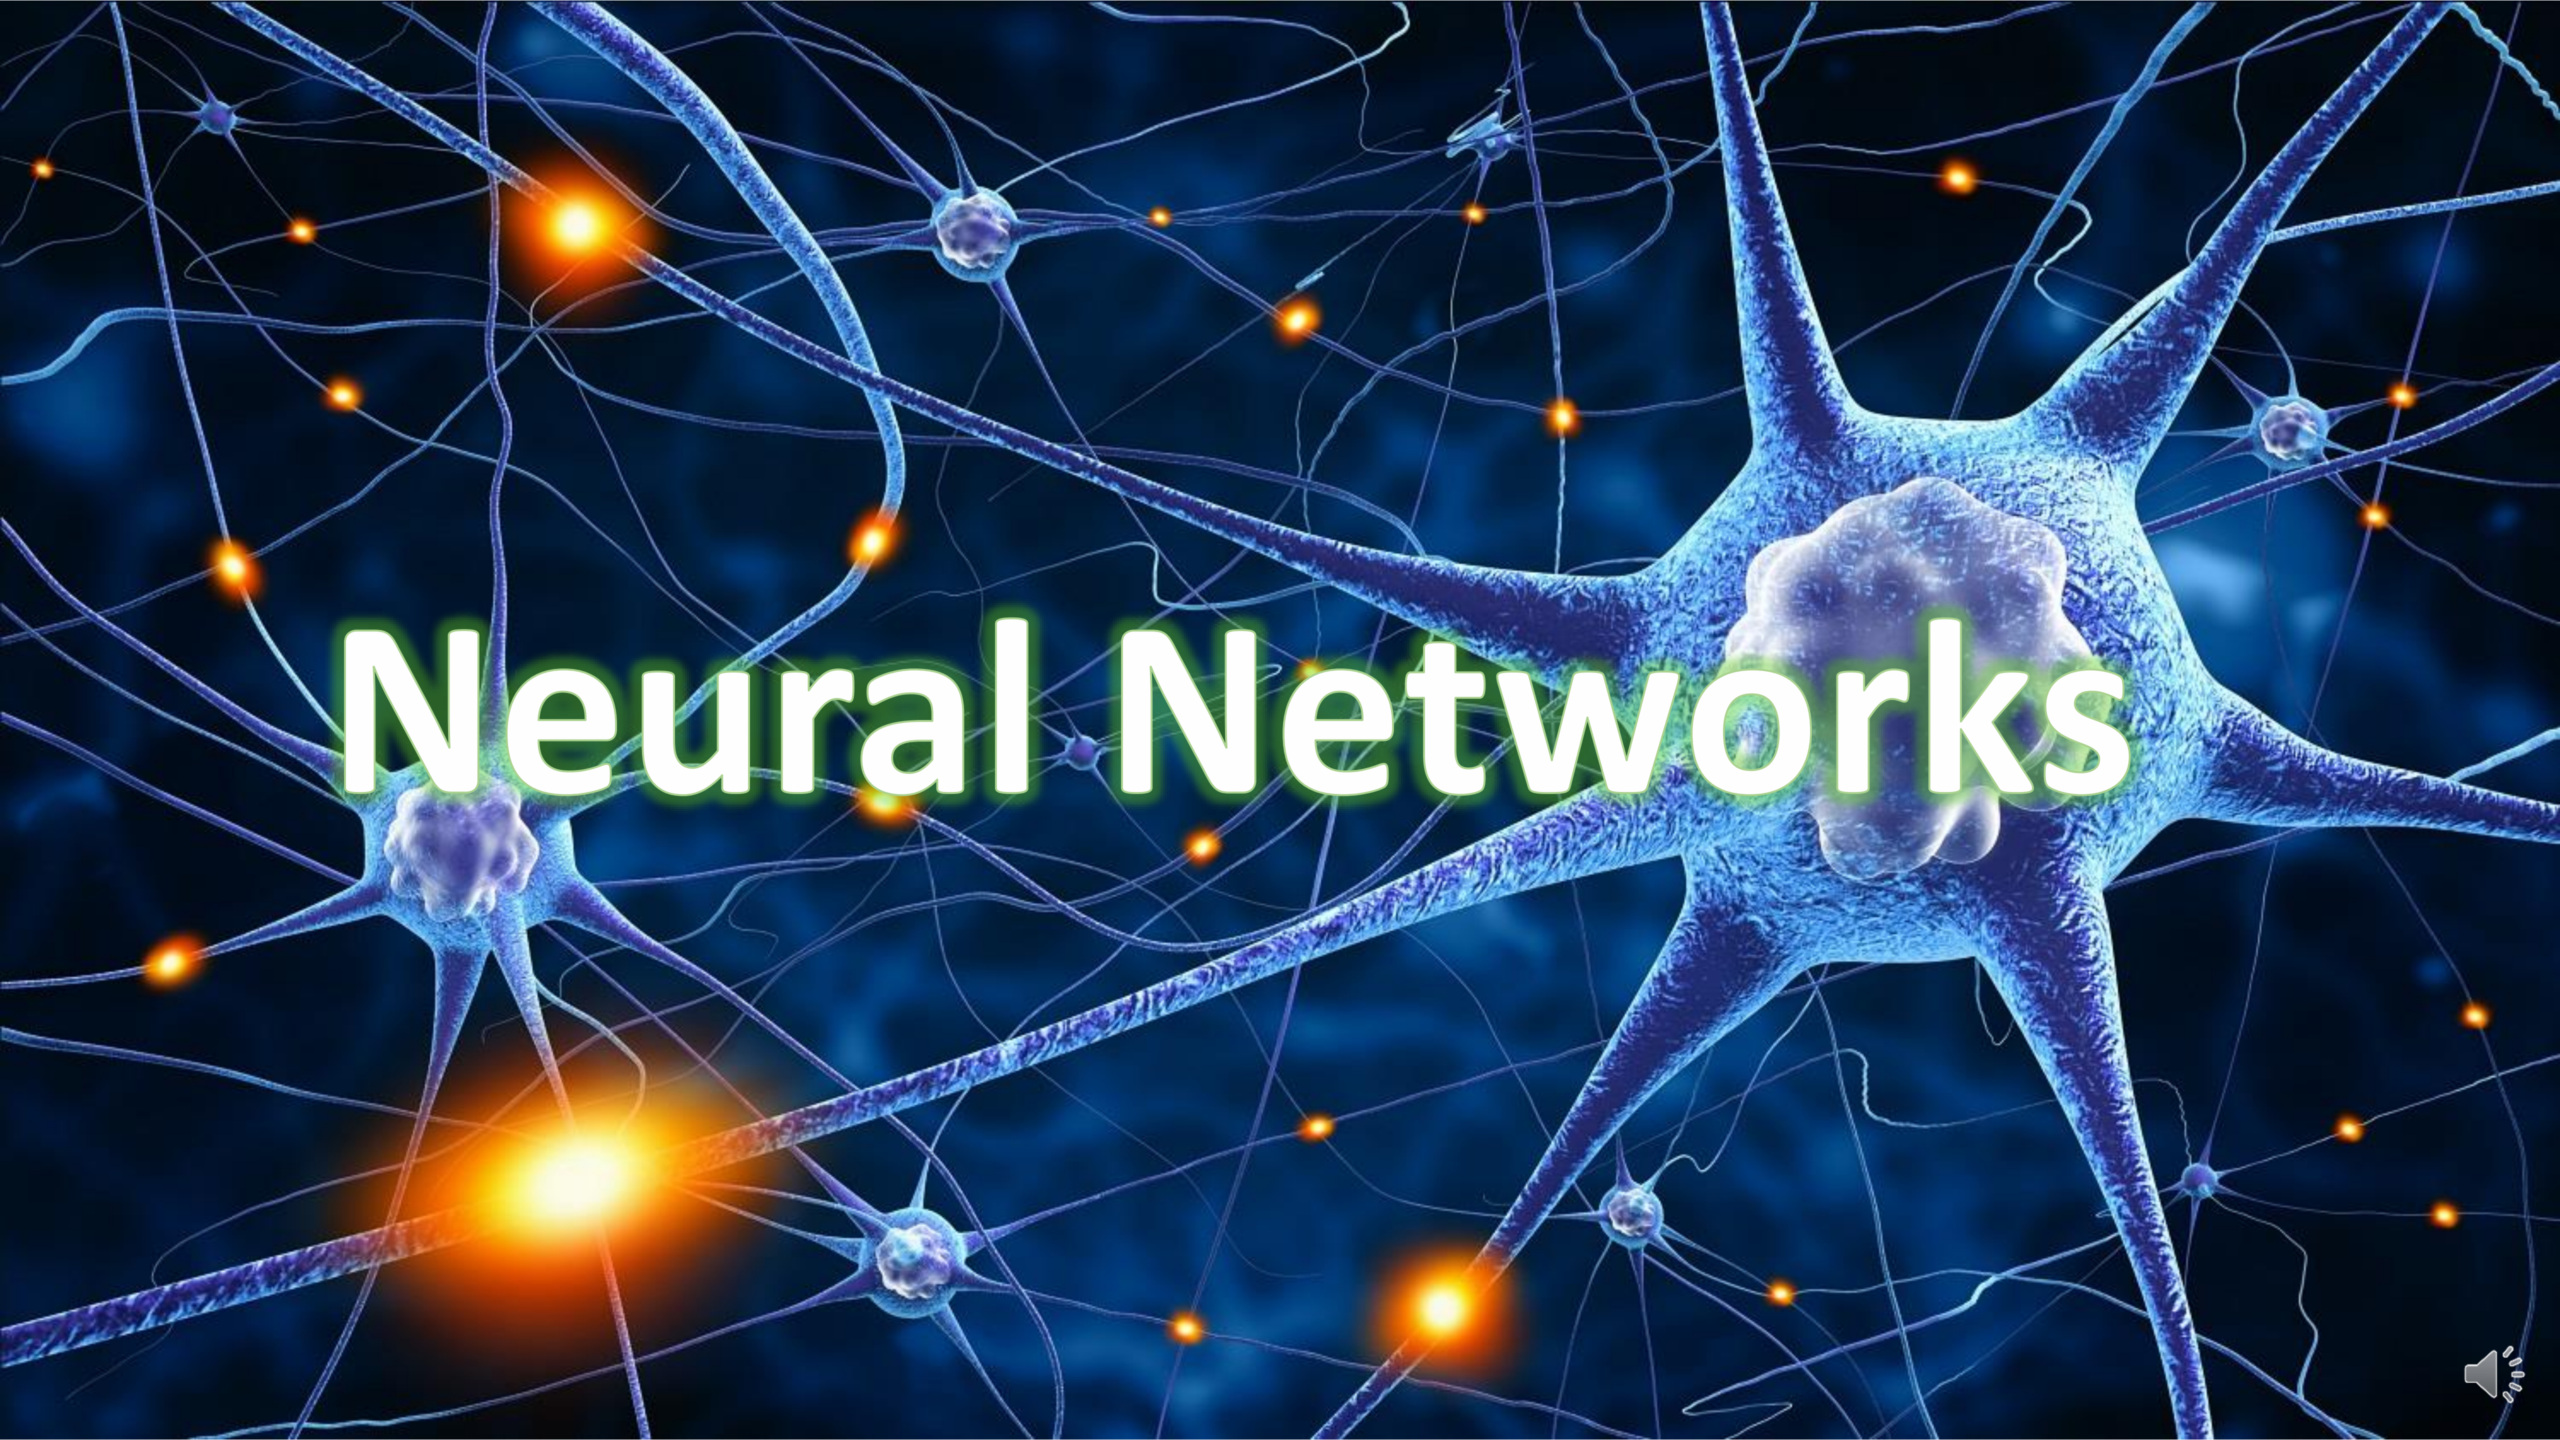
Neural Networks

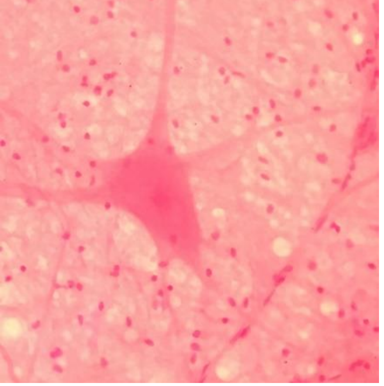

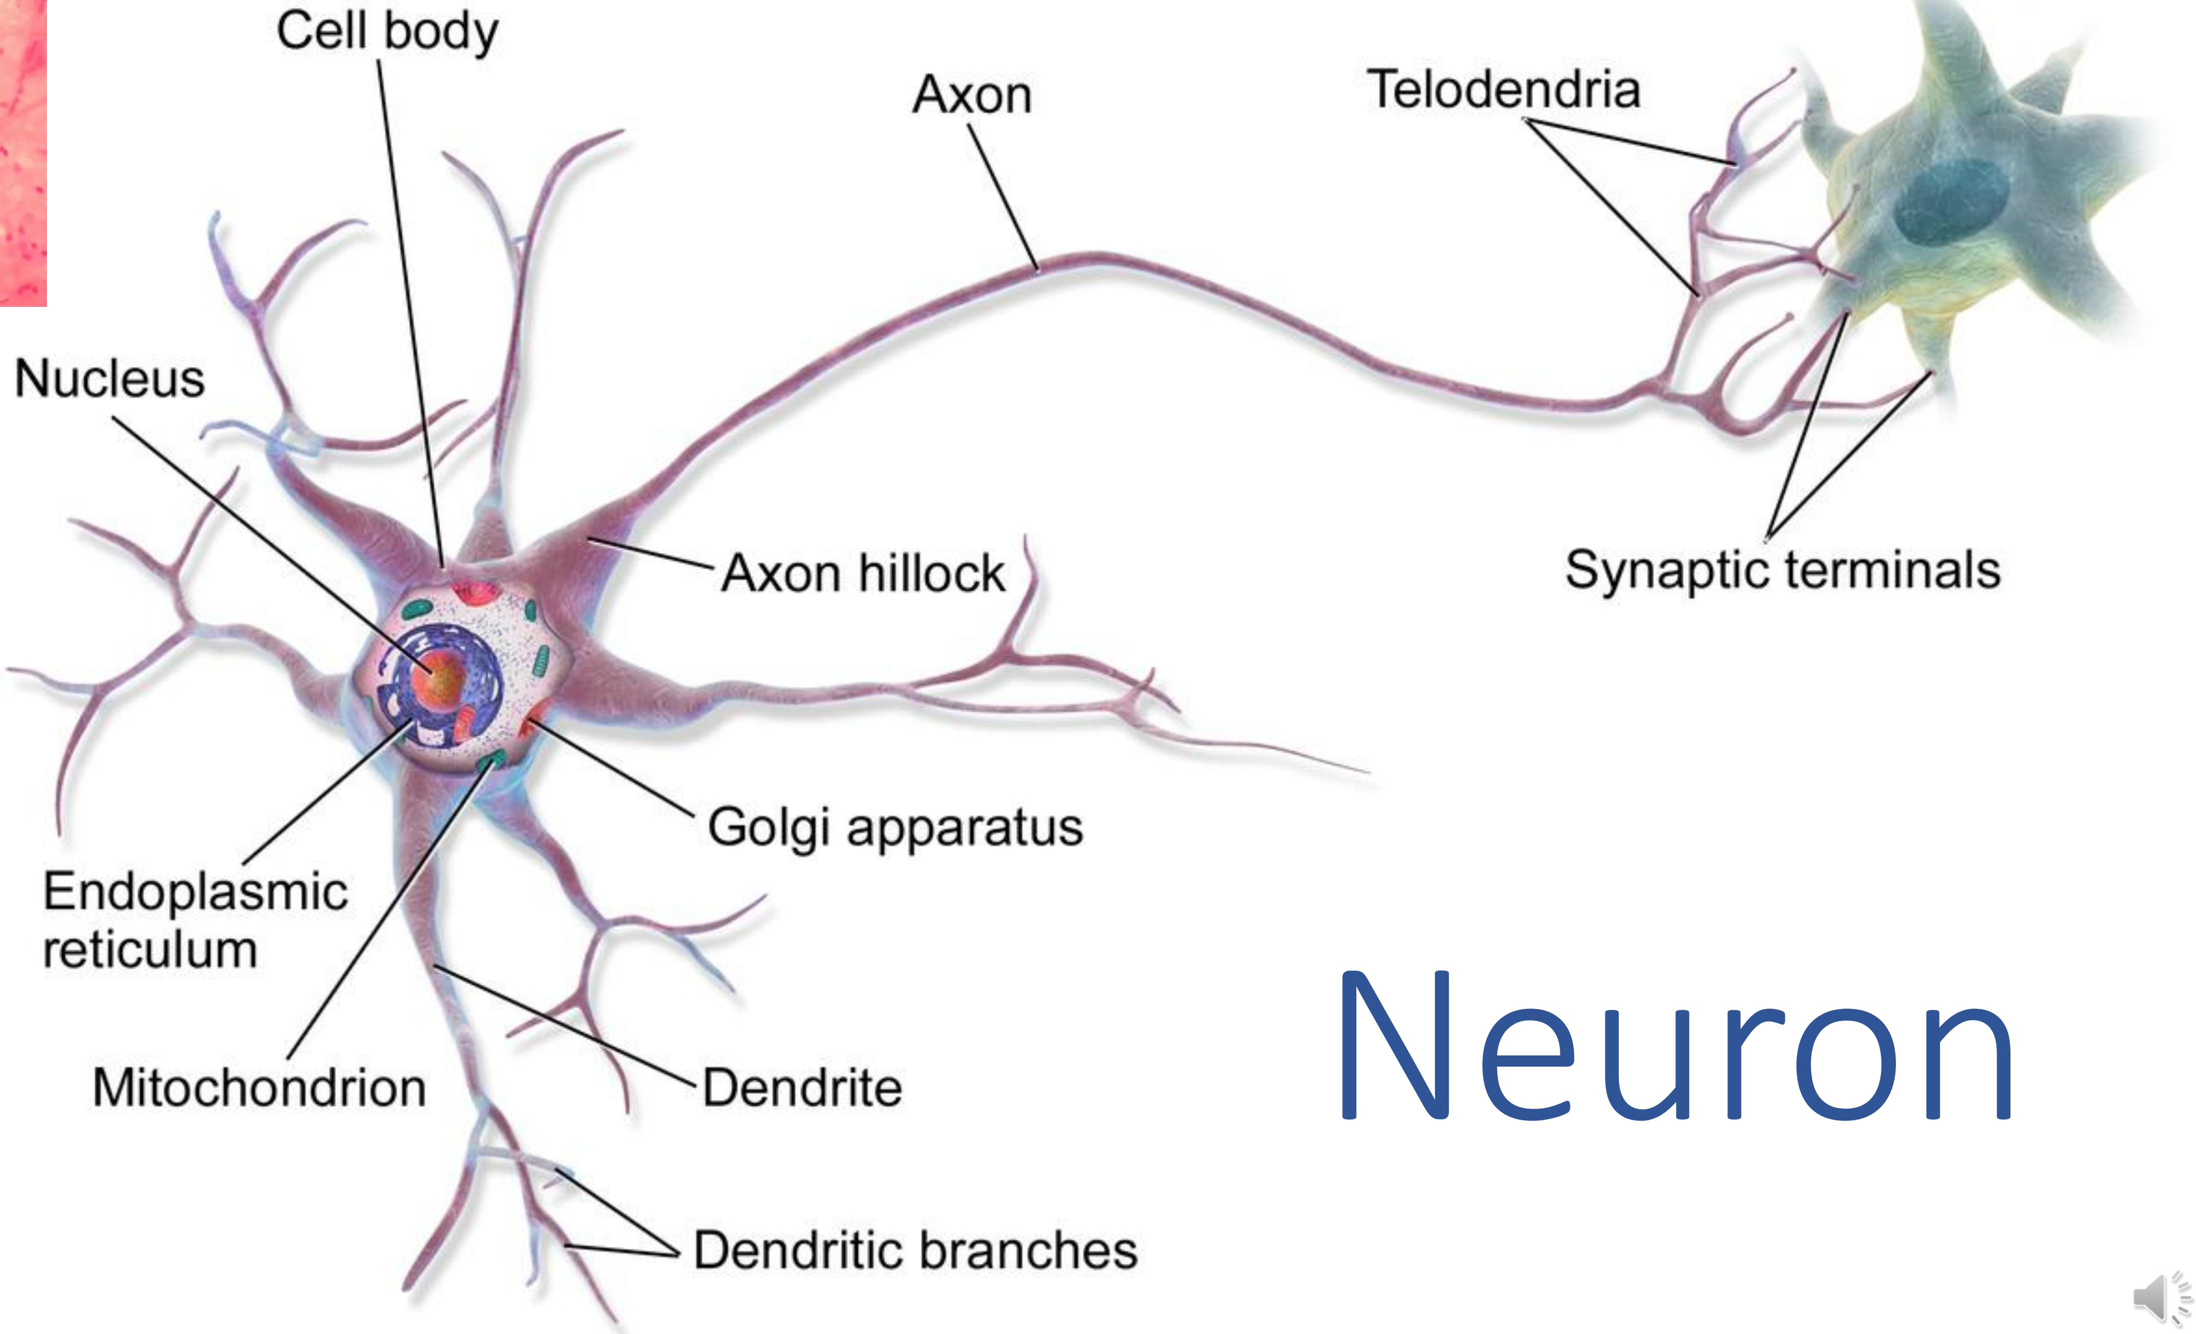
Cell body
Axon
Telodendria
Nucleus
Axon hillock
Synaptic terminals
Golgi apparatus
Endoplasmic reticulum
Mitochondrion
Dendrite
Dendritic branches

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)

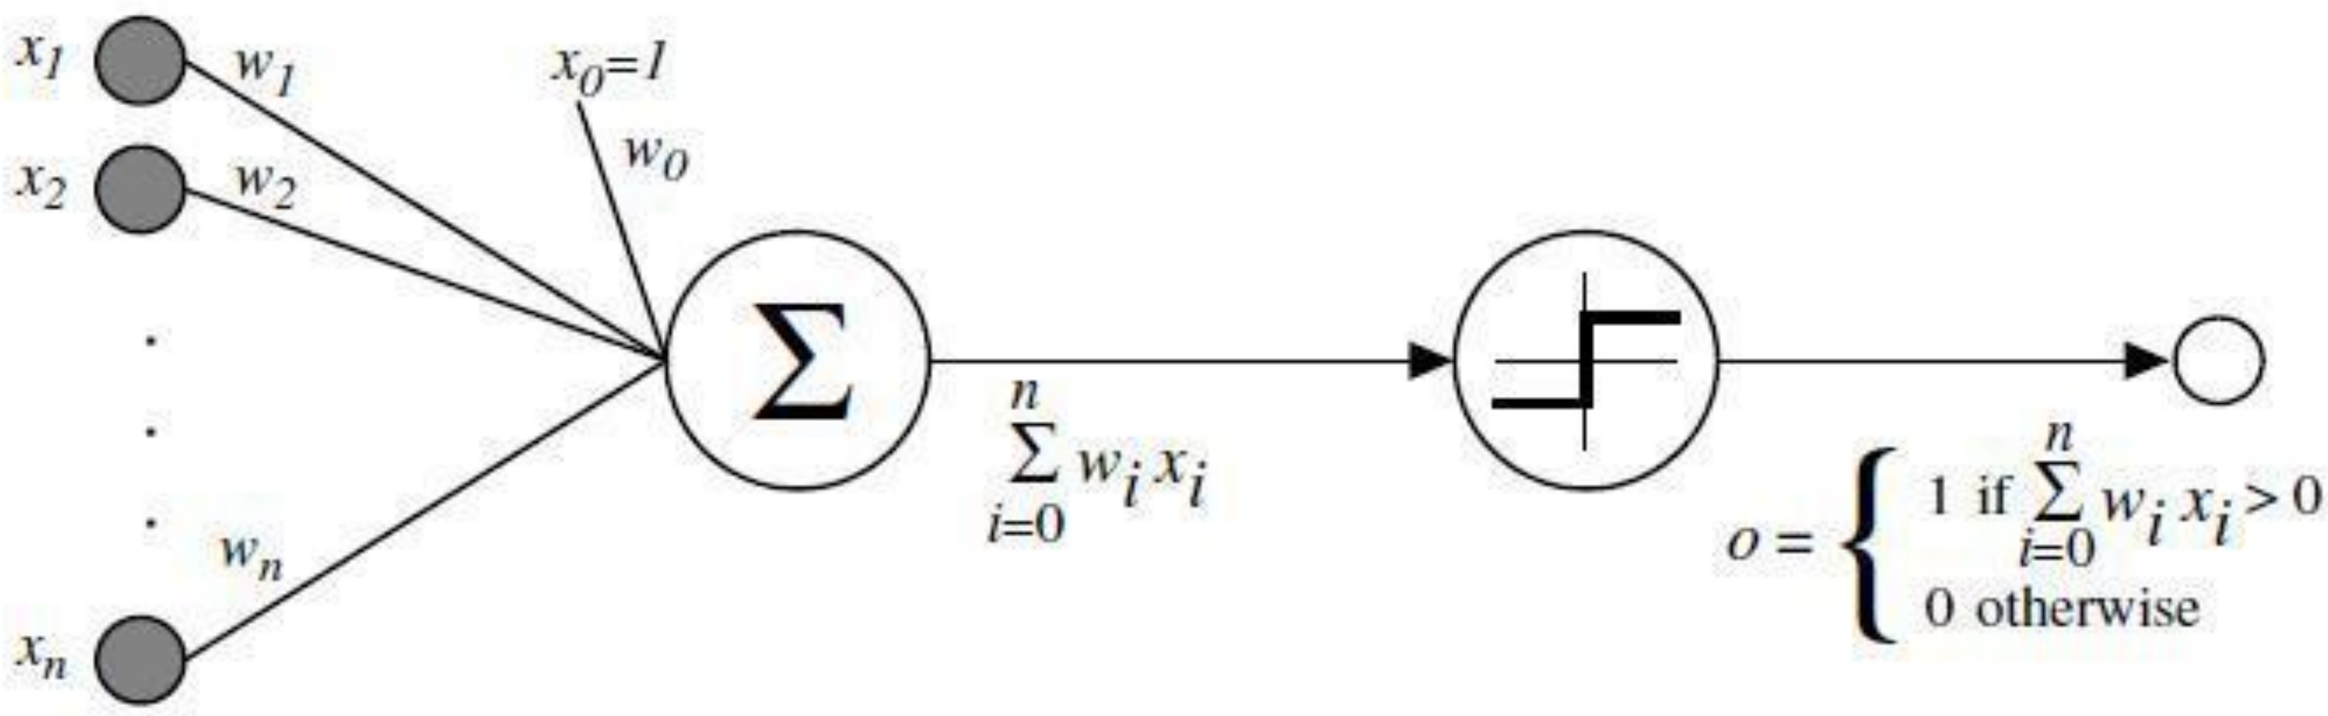

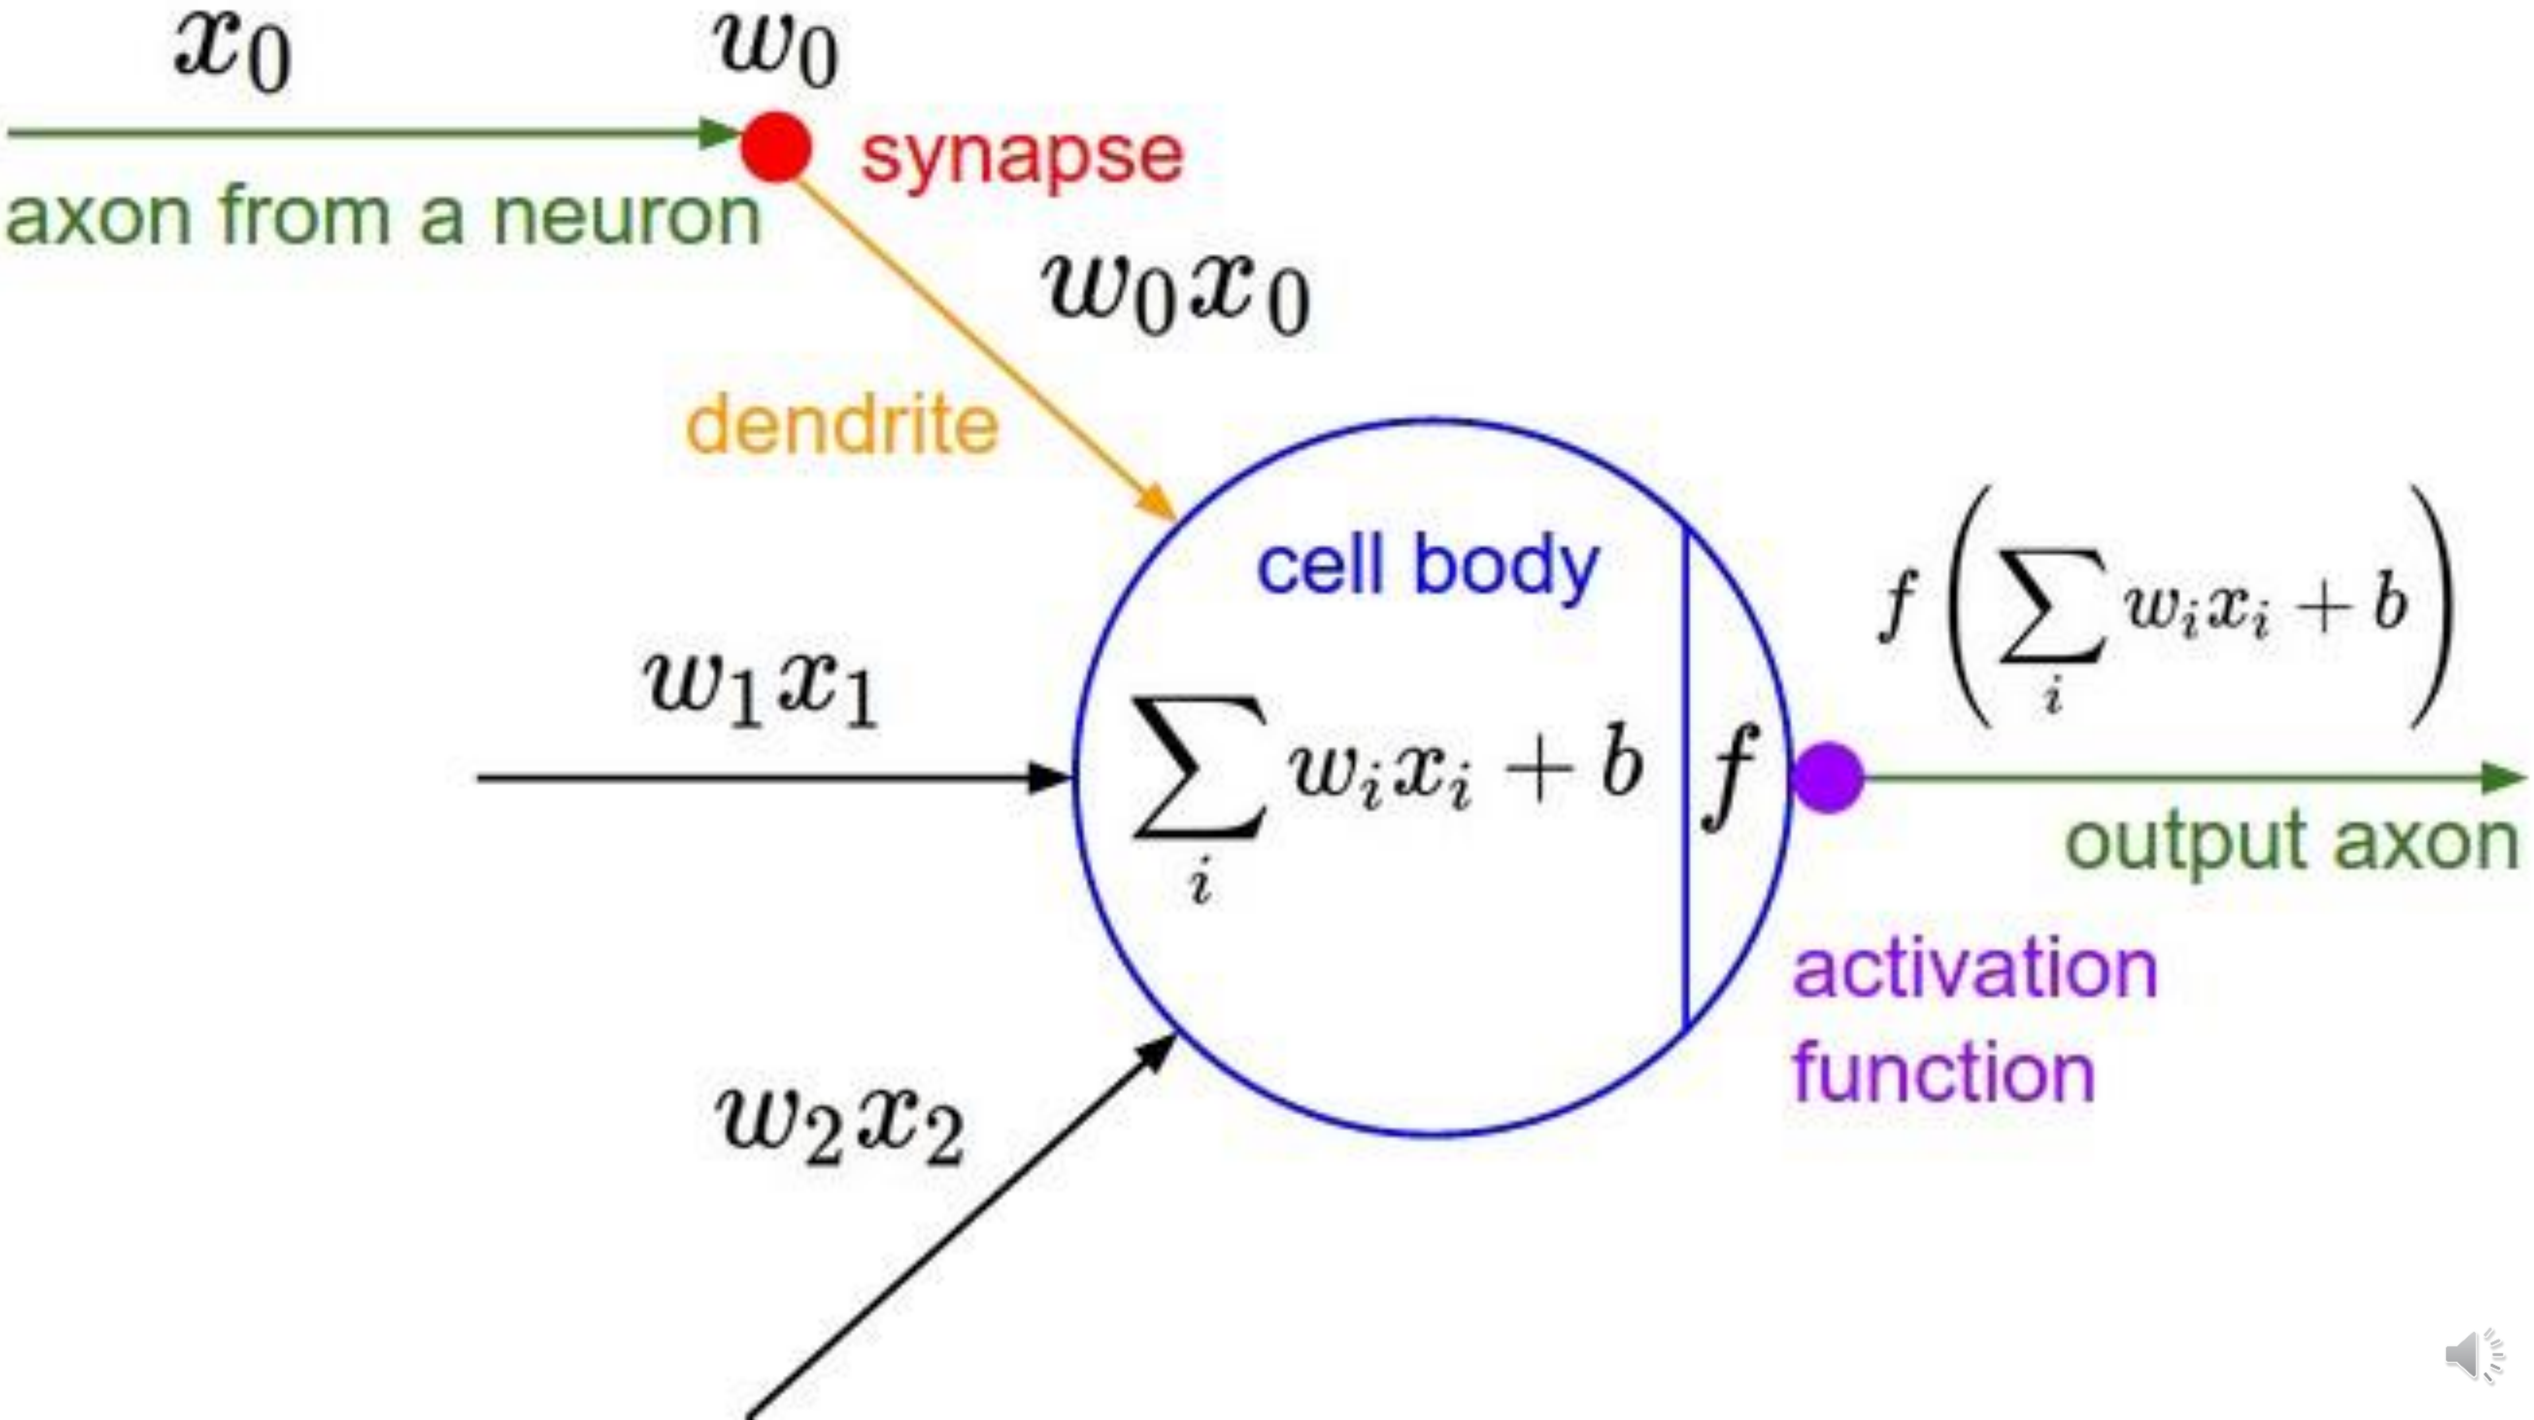
x_0
w_0
synapse
axon from a neuron
$w_0 x_0$
dendrite
cell body
$w_1 x_1$
$\sum_i w_i x_i + b$
f
$f\left(\sum_i w_i x_i + b\right)$
output axon
activation function
$w_2 x_2$

1969: Perceptrons can't do XOR!

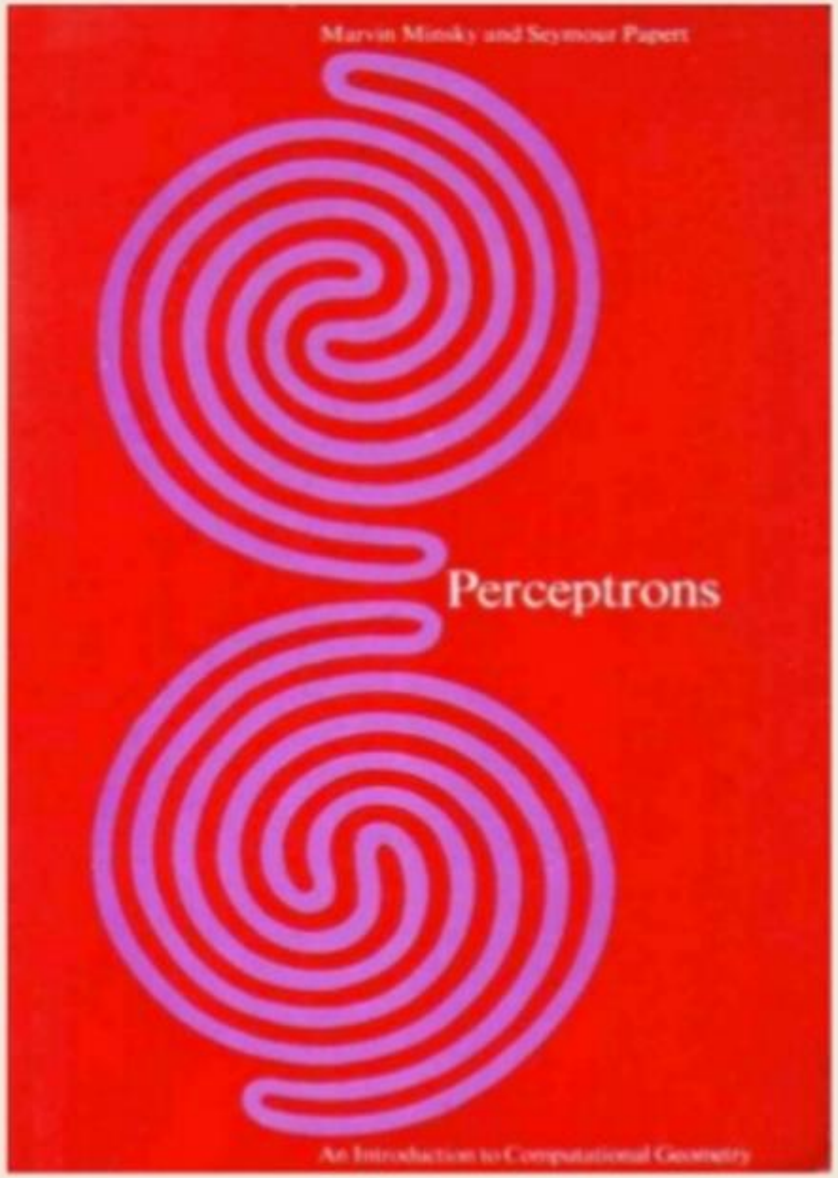

http://www.i-programmer.info/images/stories/BabBag/AI/book.jpg

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

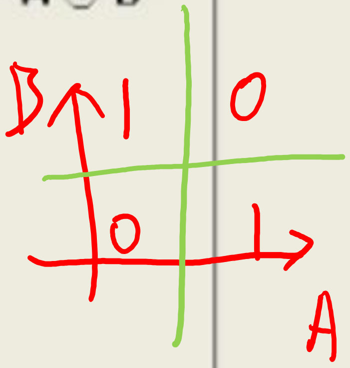

http://hyperphysics.phy-astr.gsu.edu/hbase/electronic/ietron/xor.gif

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)

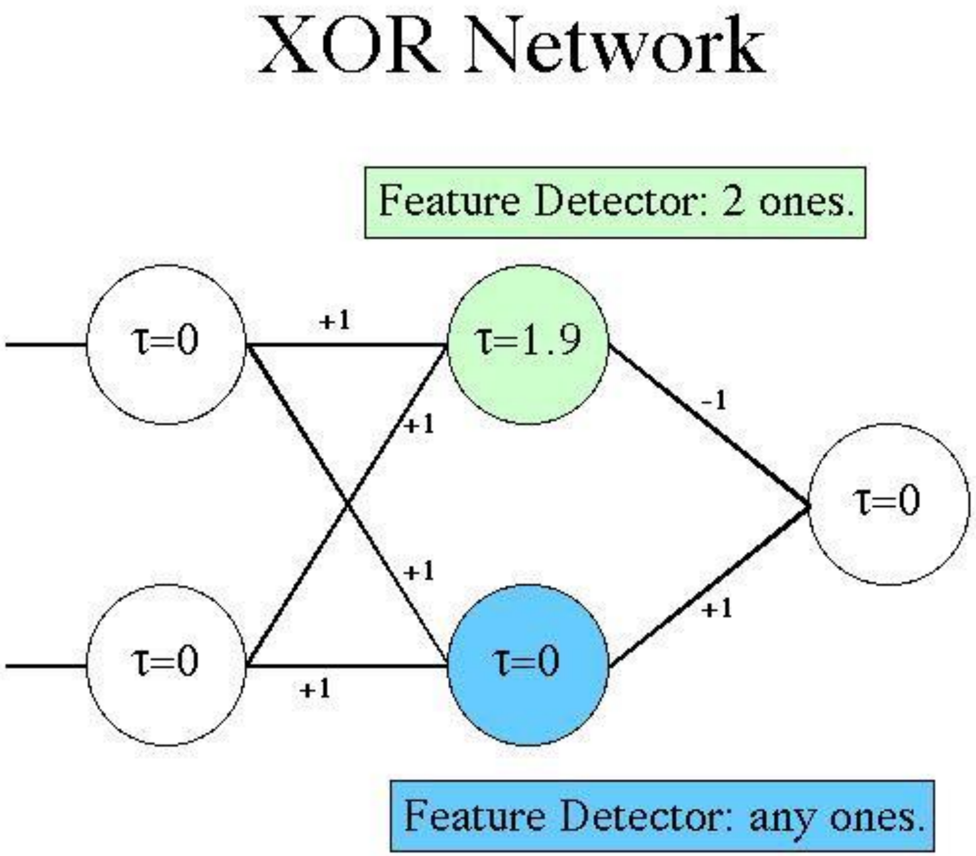

Geoffrey Hinton

Backpropagation

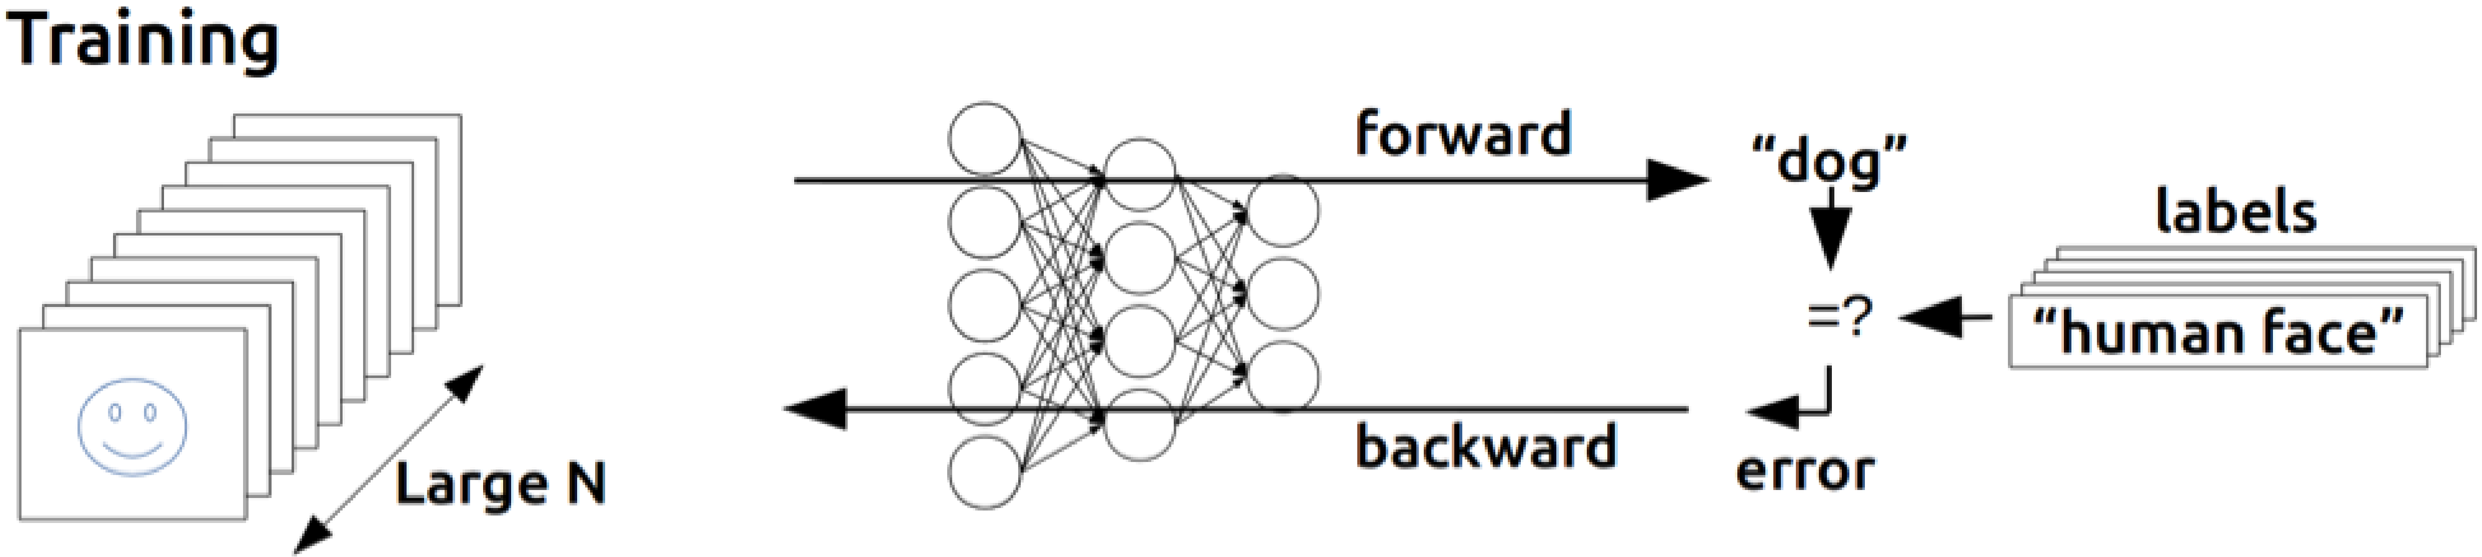

Inference

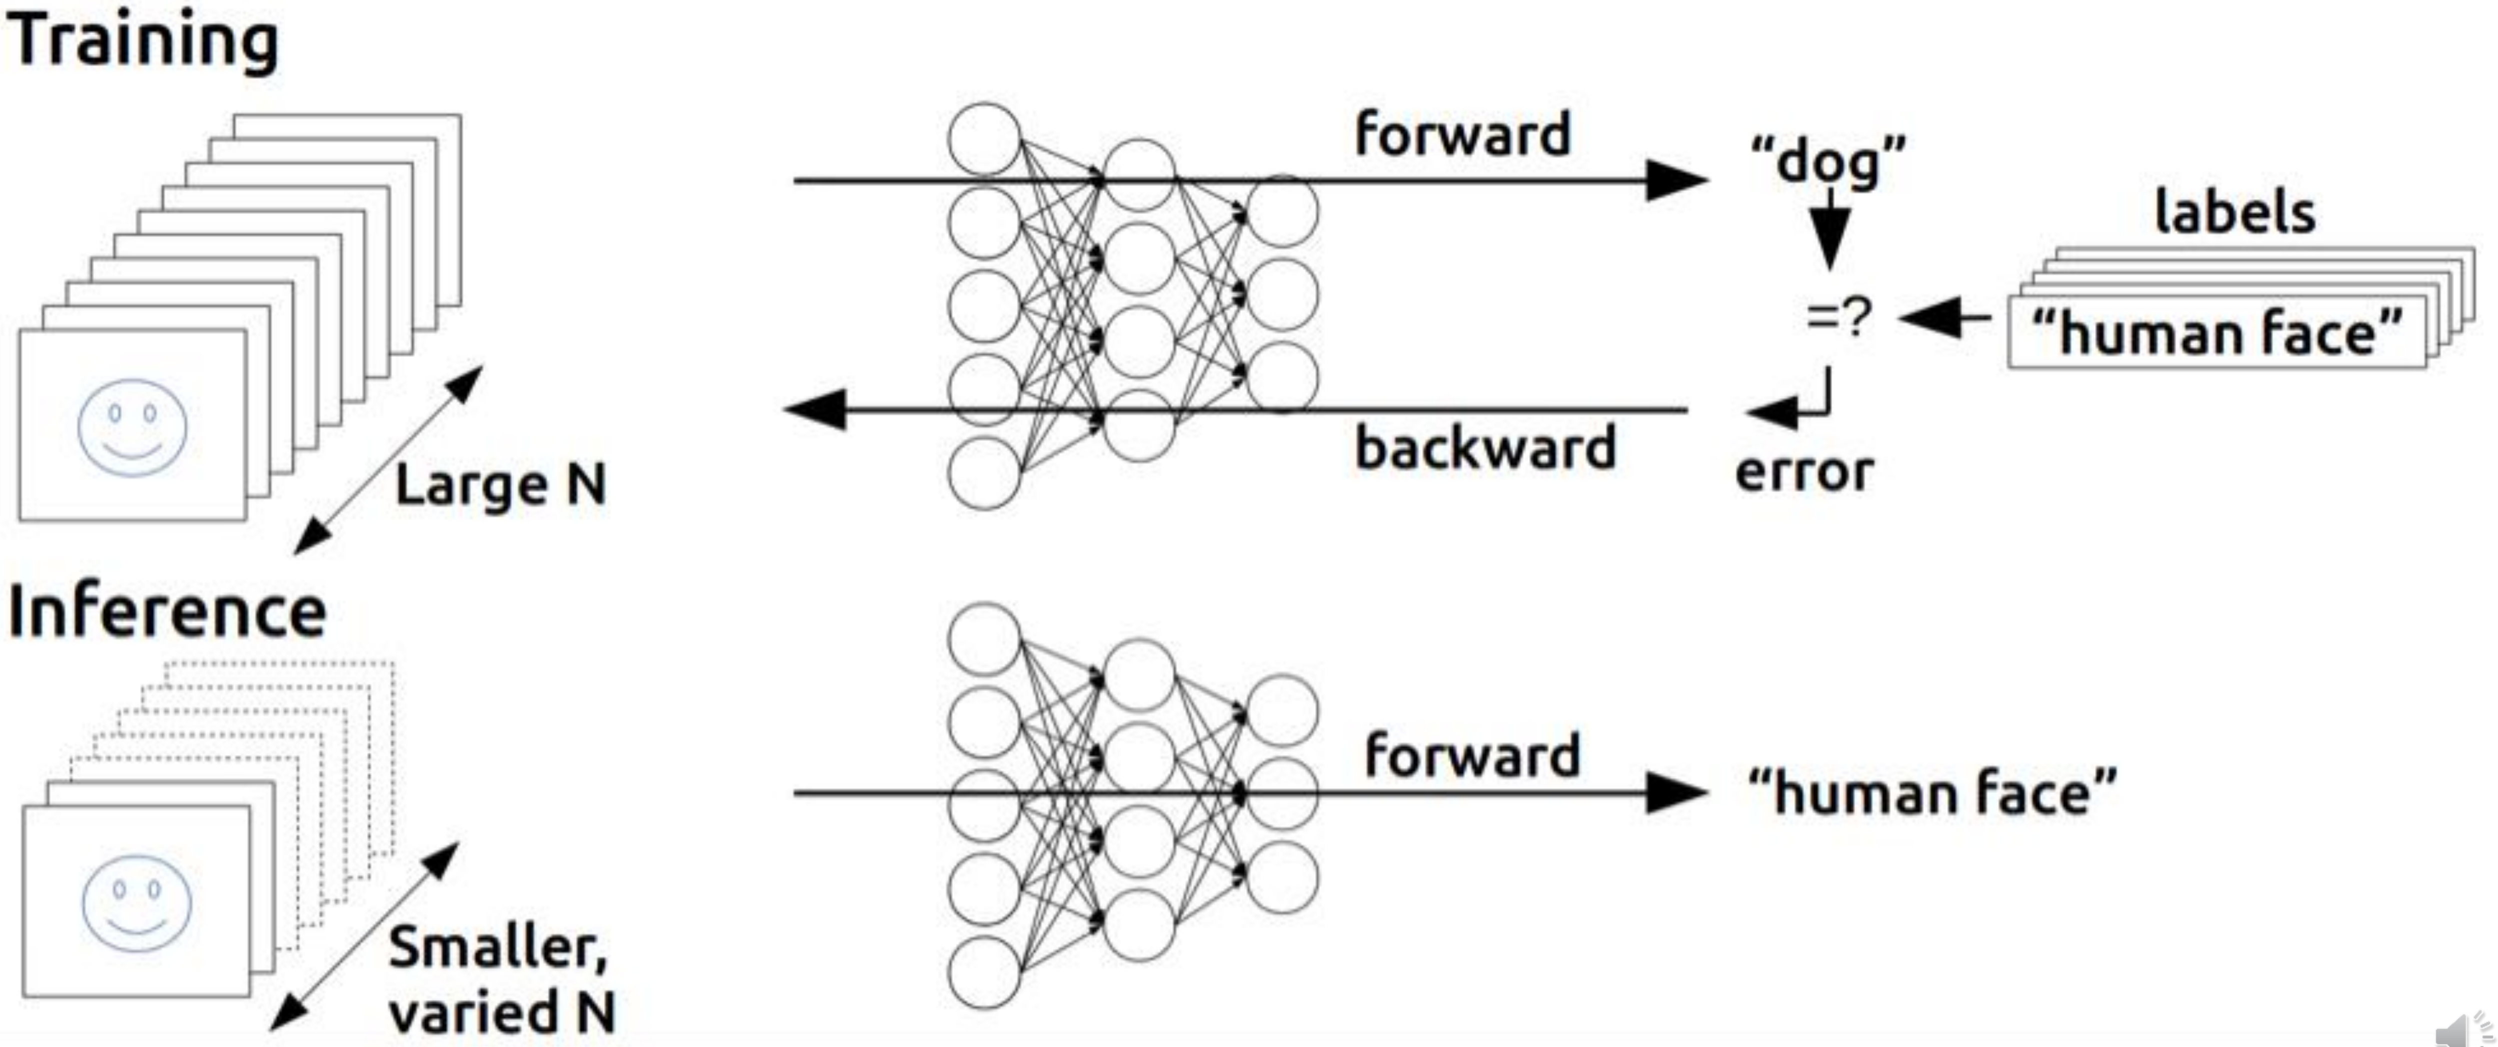

Training
forward
"dog"
labels
=?
"human face"
Large N
backward
error
Inference
forward
"human face"
Smaller,
varied N

Chain Rule

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

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

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

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

Computation Graph

e=(a+b)*(b+1)

=>

c = a + b

d = b + 1

e = c*d

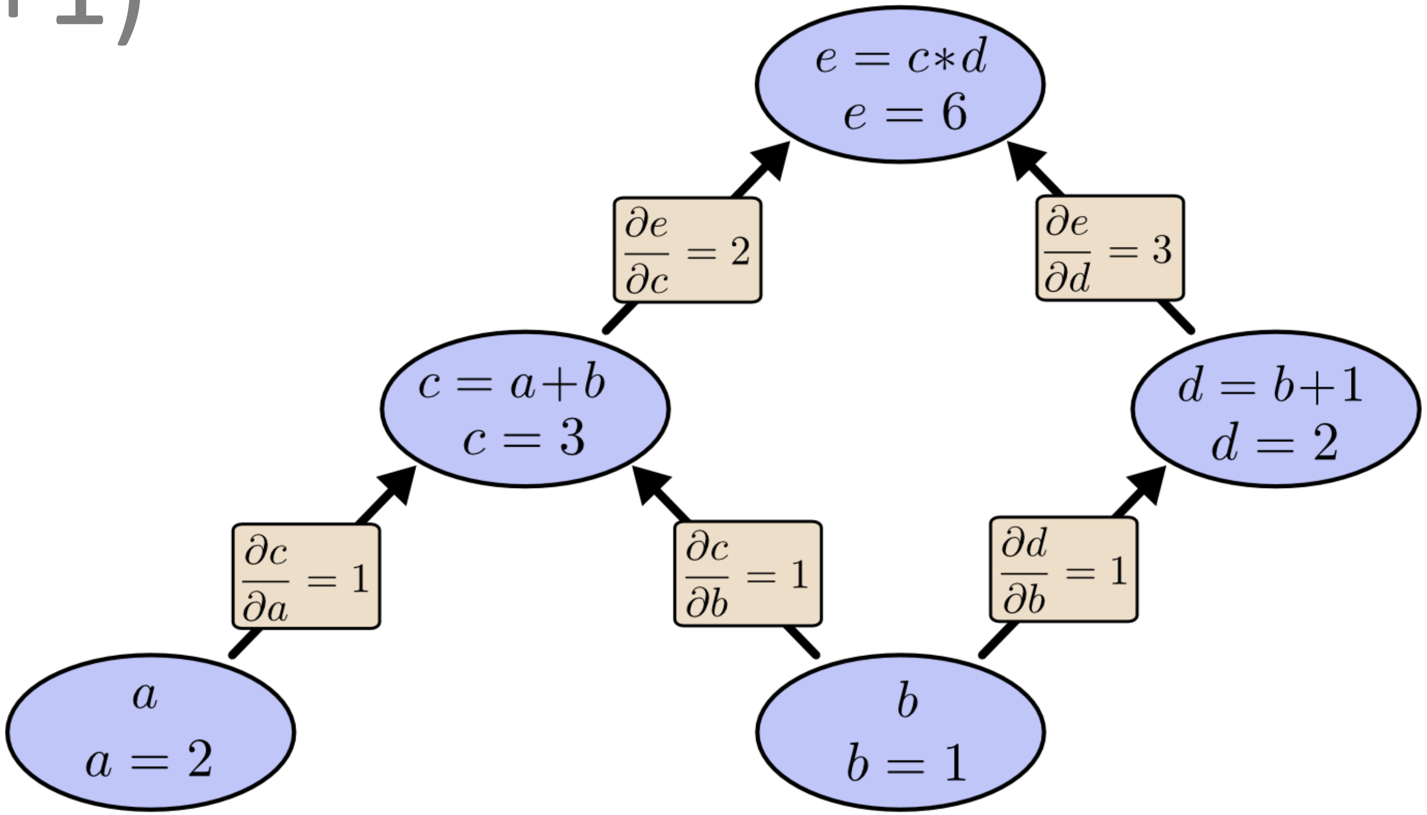

https://colah.github.io/posts/2015-08-Backprop/

Example: Recognizing Handwritten Digits

- MNIST dataset

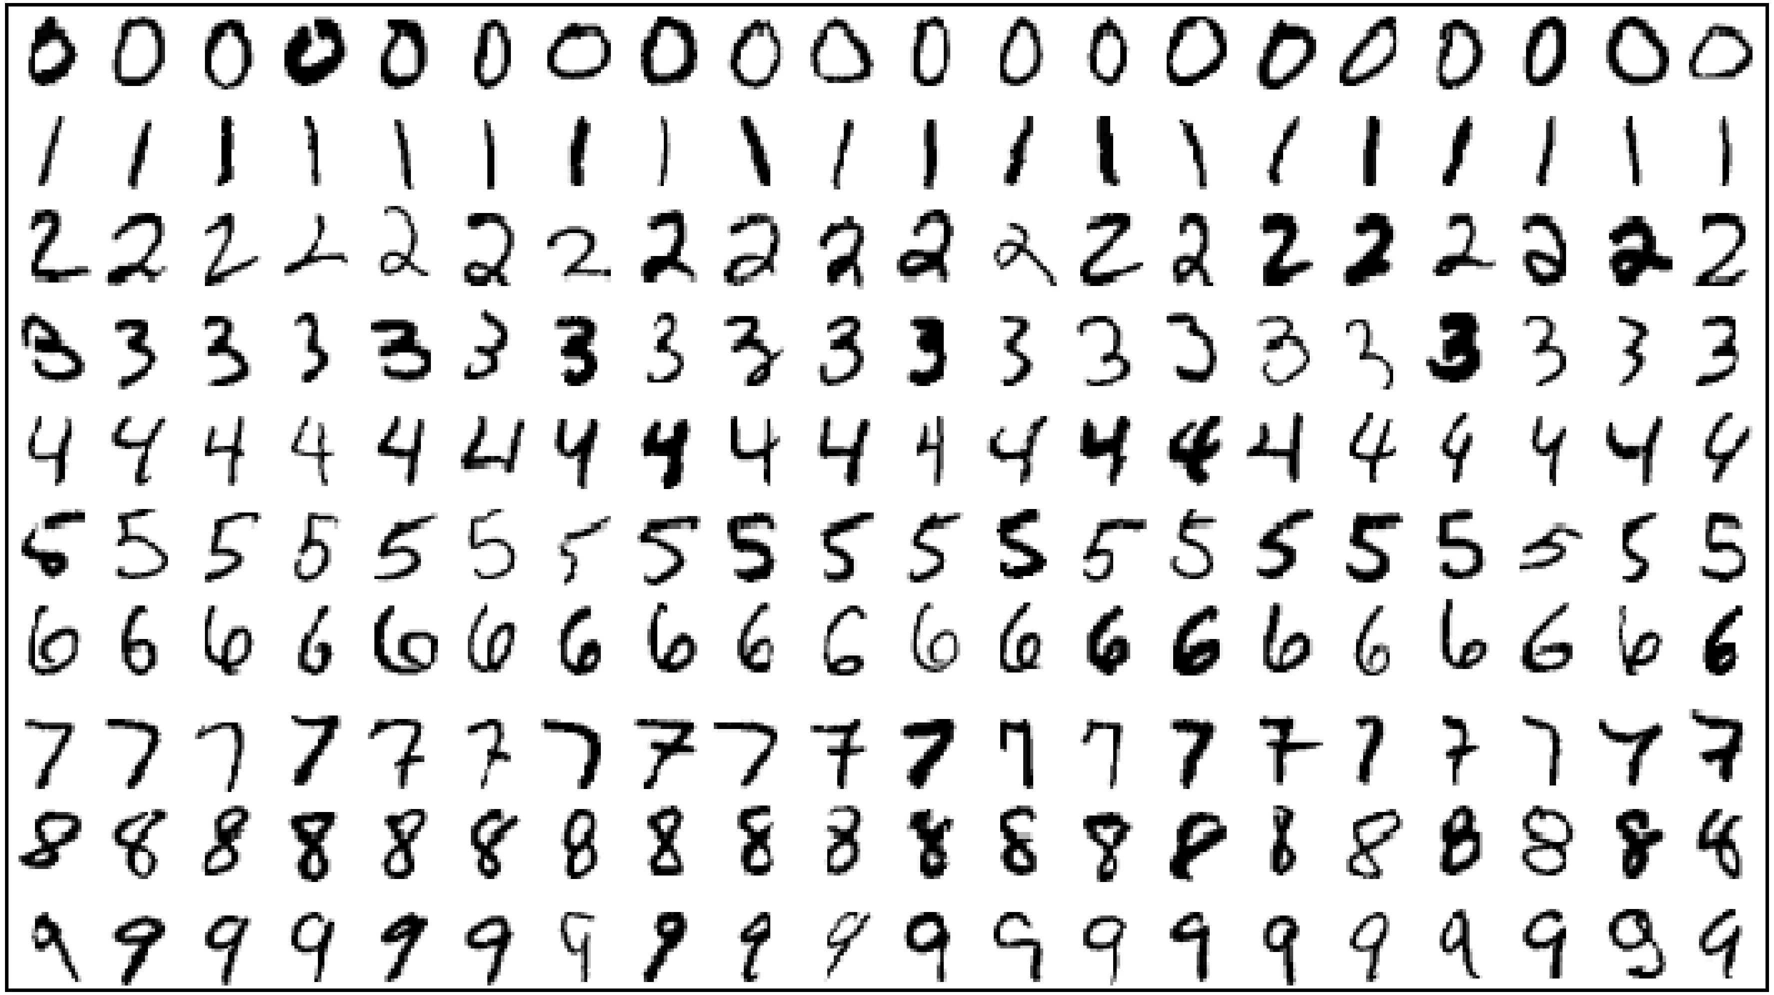

3Blue1Brown

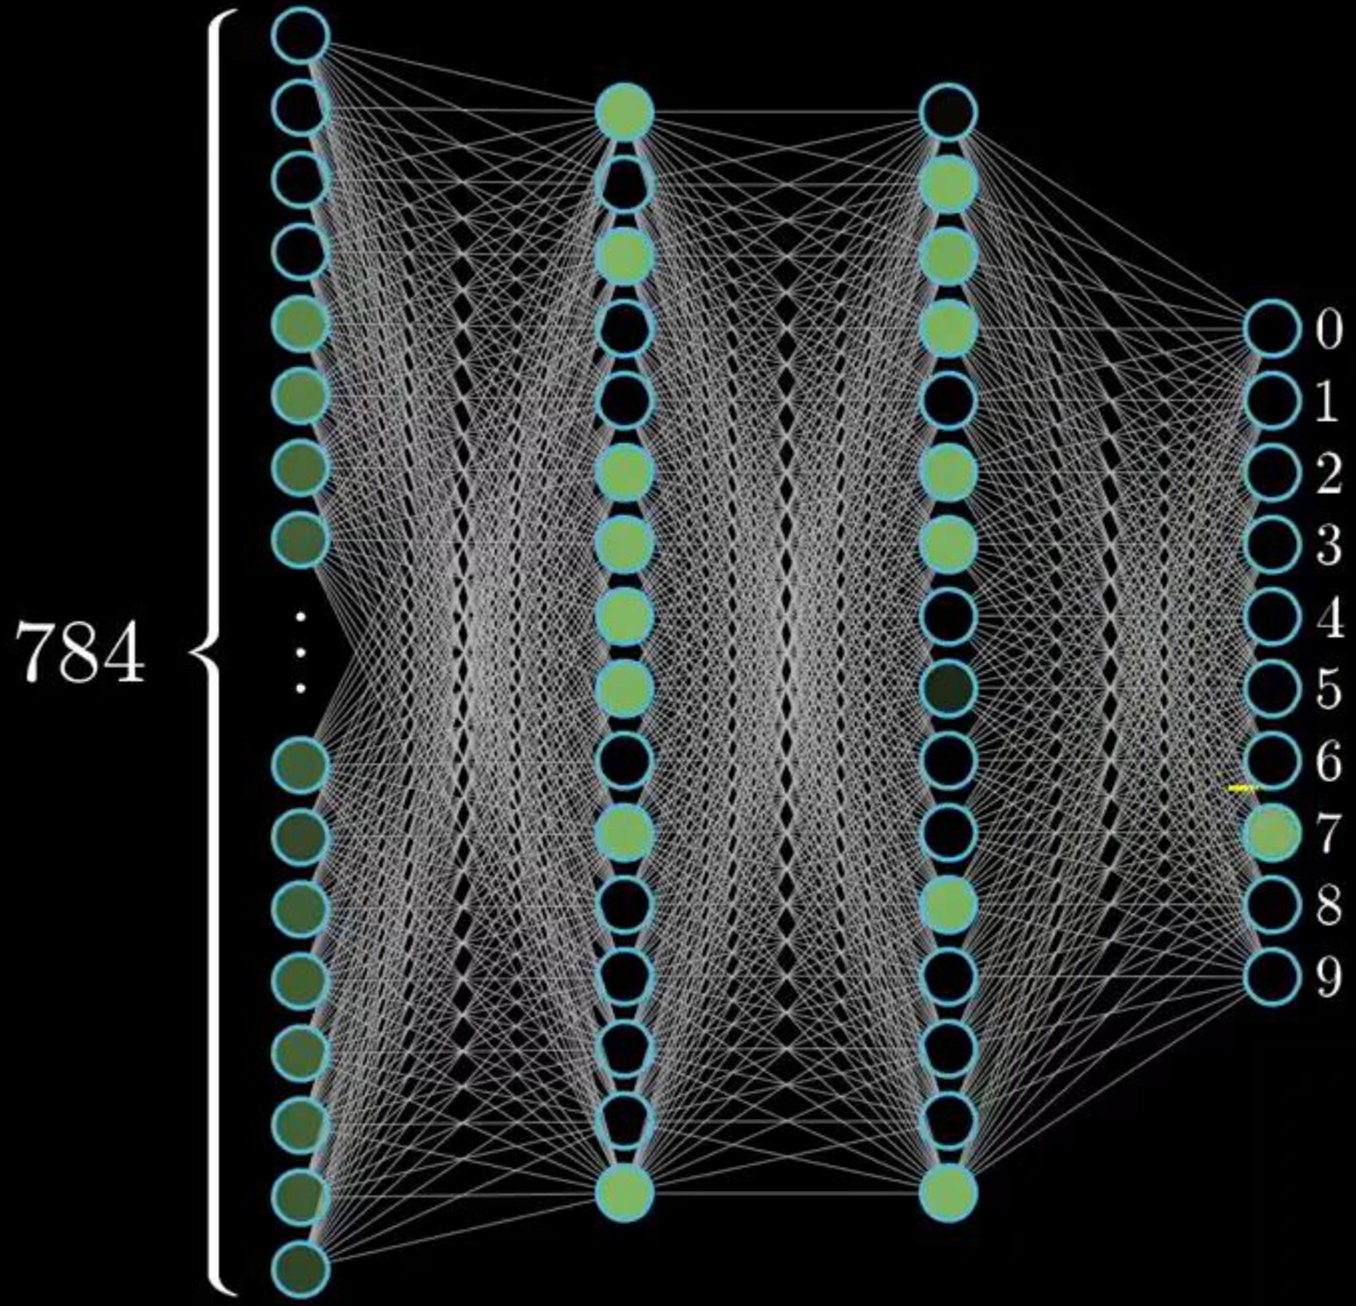

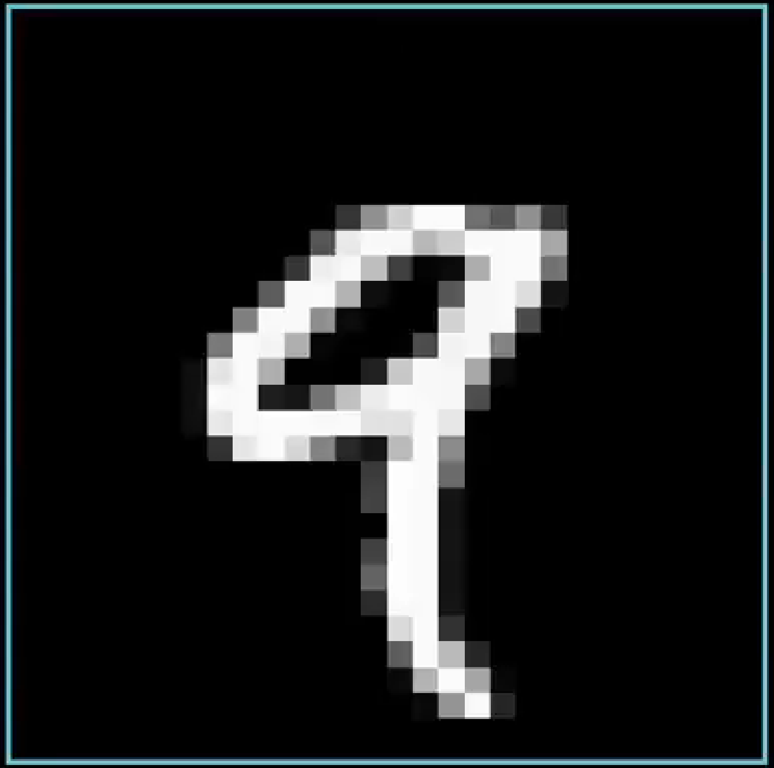

Gradient Descent

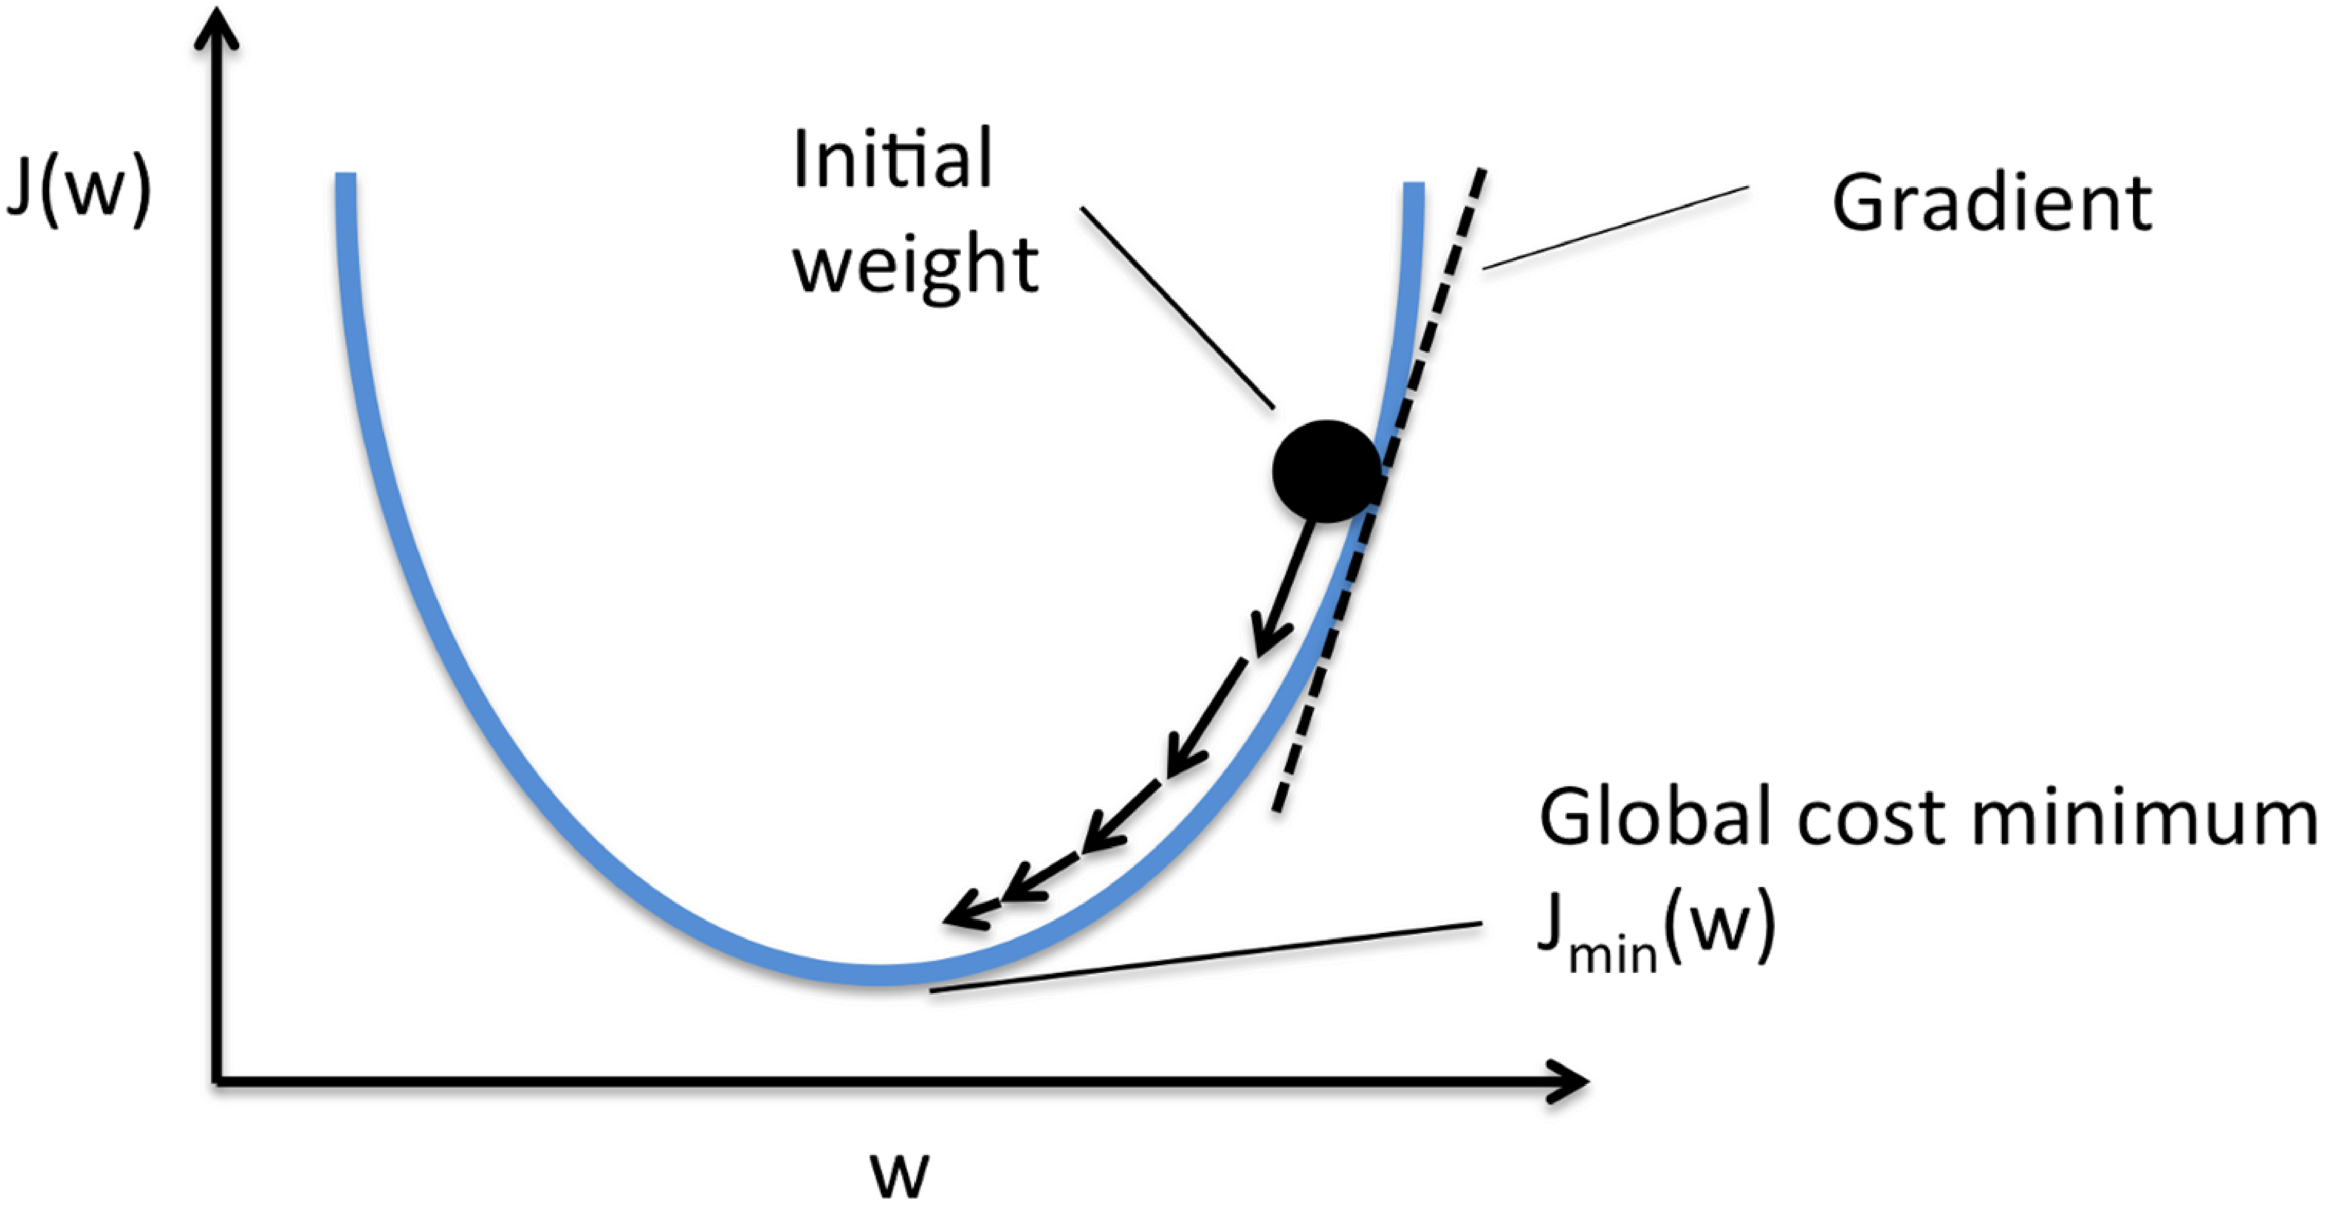

Cost function

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

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

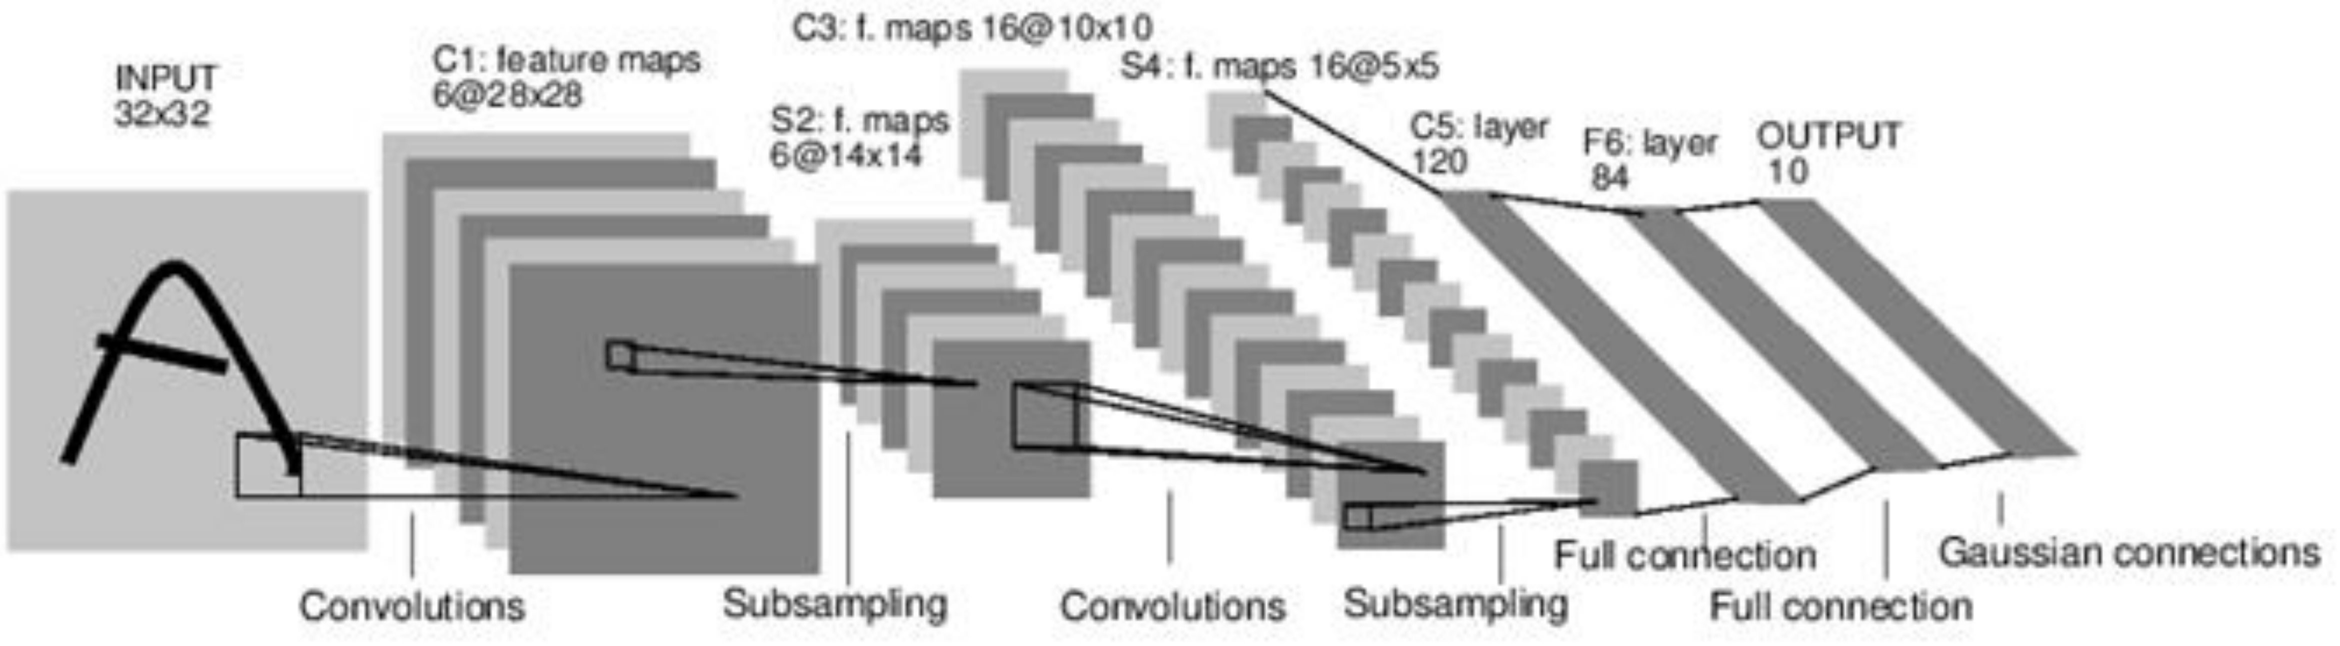

A Full Convolutional Neural Network (LeNet)


IMAGENET

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

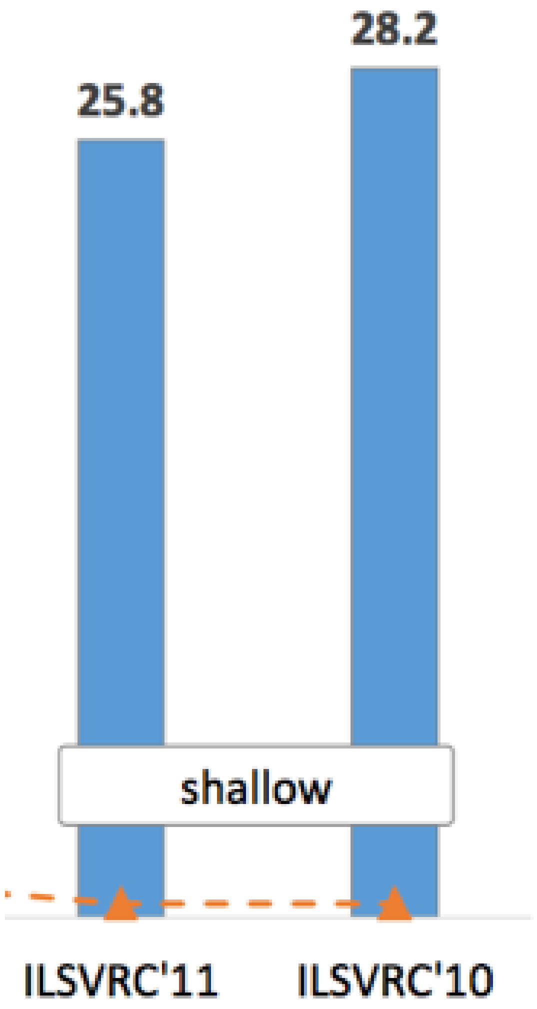

Convolutional Neural Network (AlexNet)

- Alex Krizhevsky, Geoffery Hinton et al., 2012

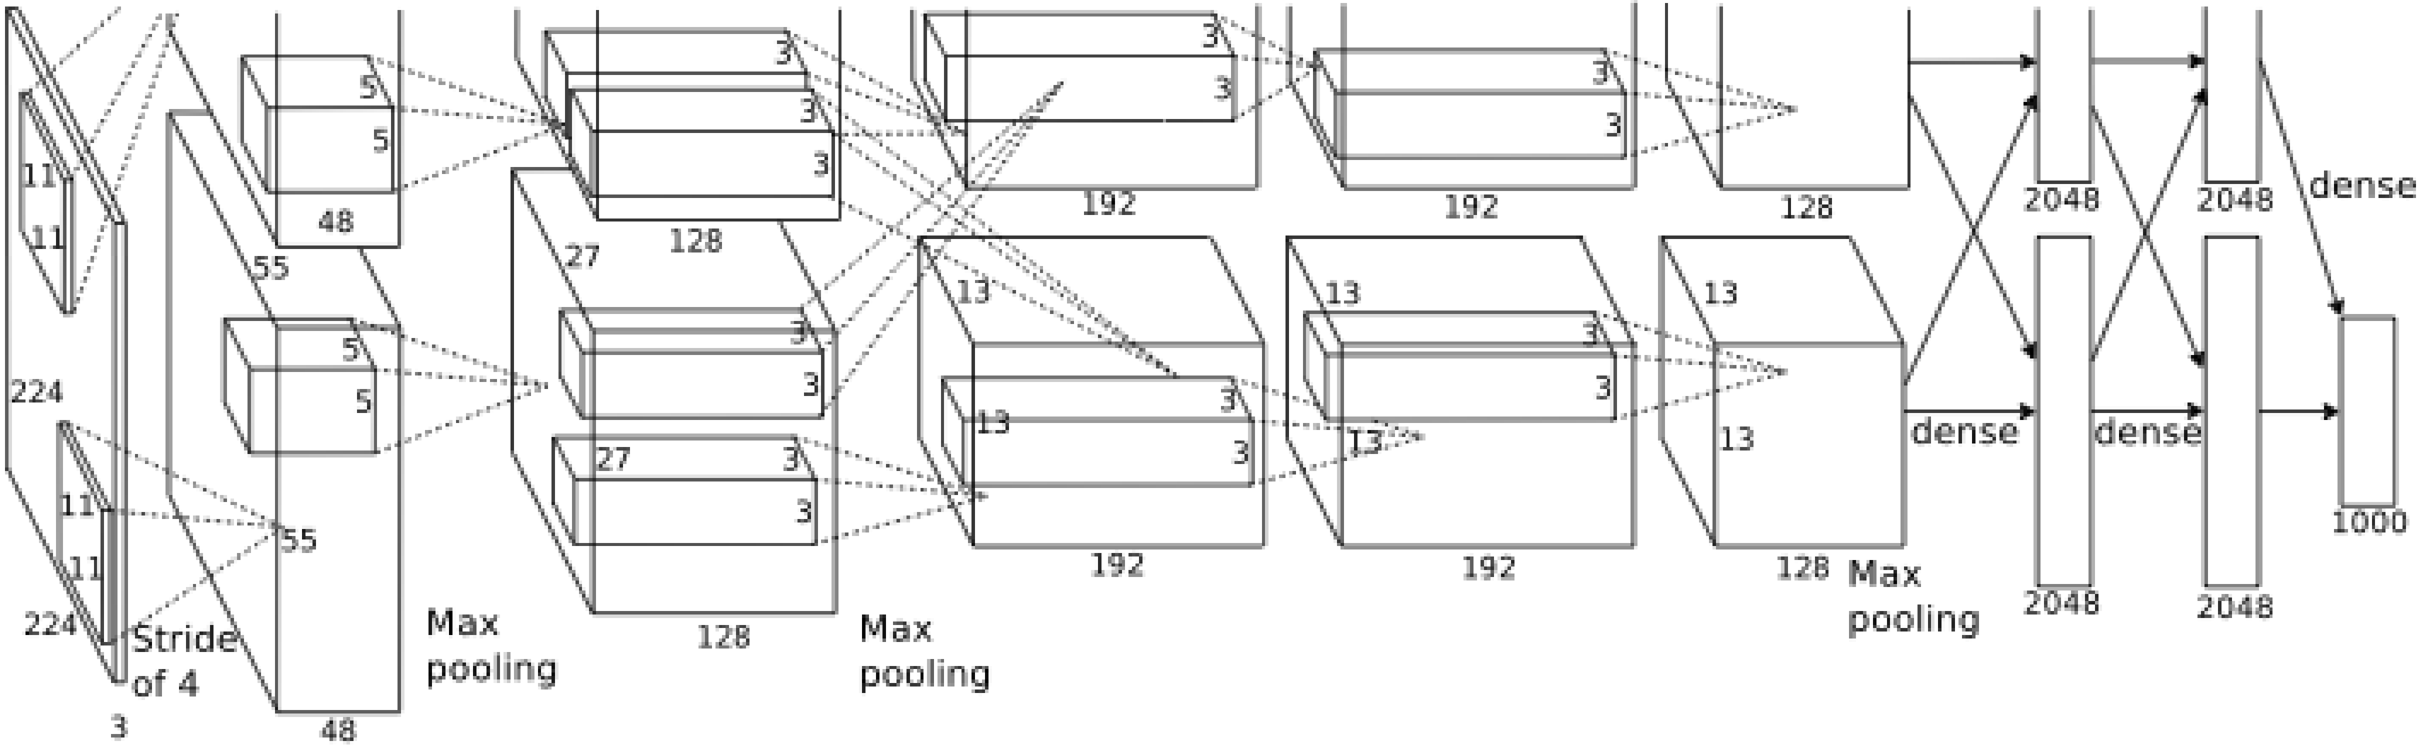

Error Rate on ImageNet Challenge

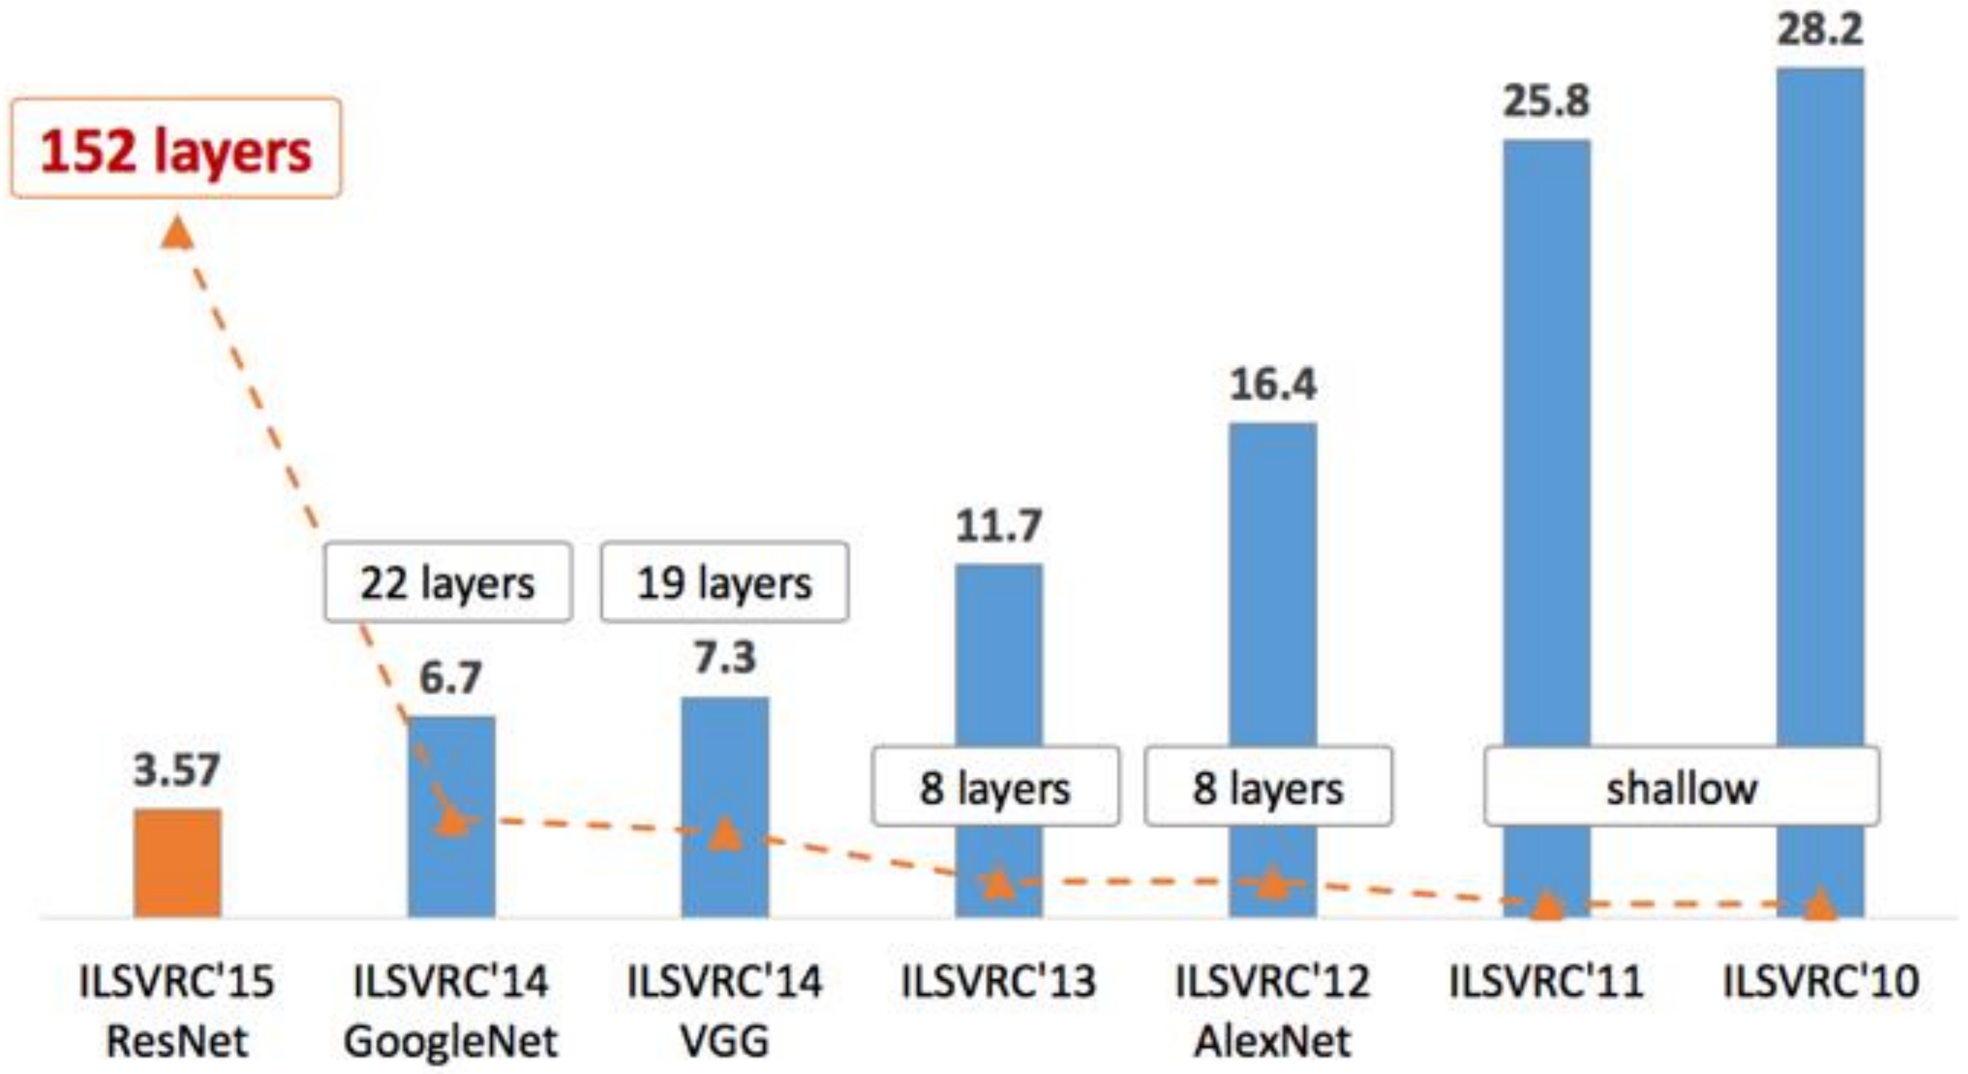

LeNet

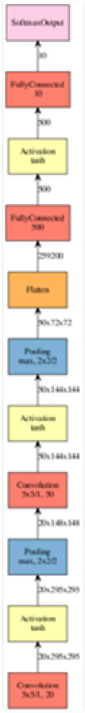

AlexNet

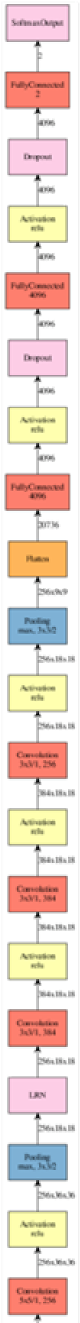

VGG

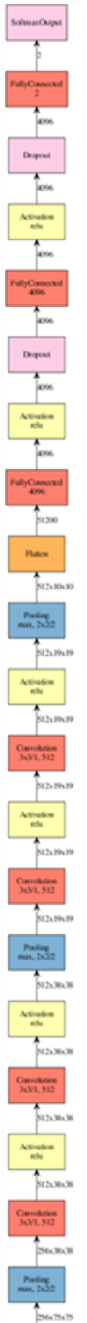

GoogLeNet

Inception V3

Inception BN


WE NEED TO GO
DEEPER

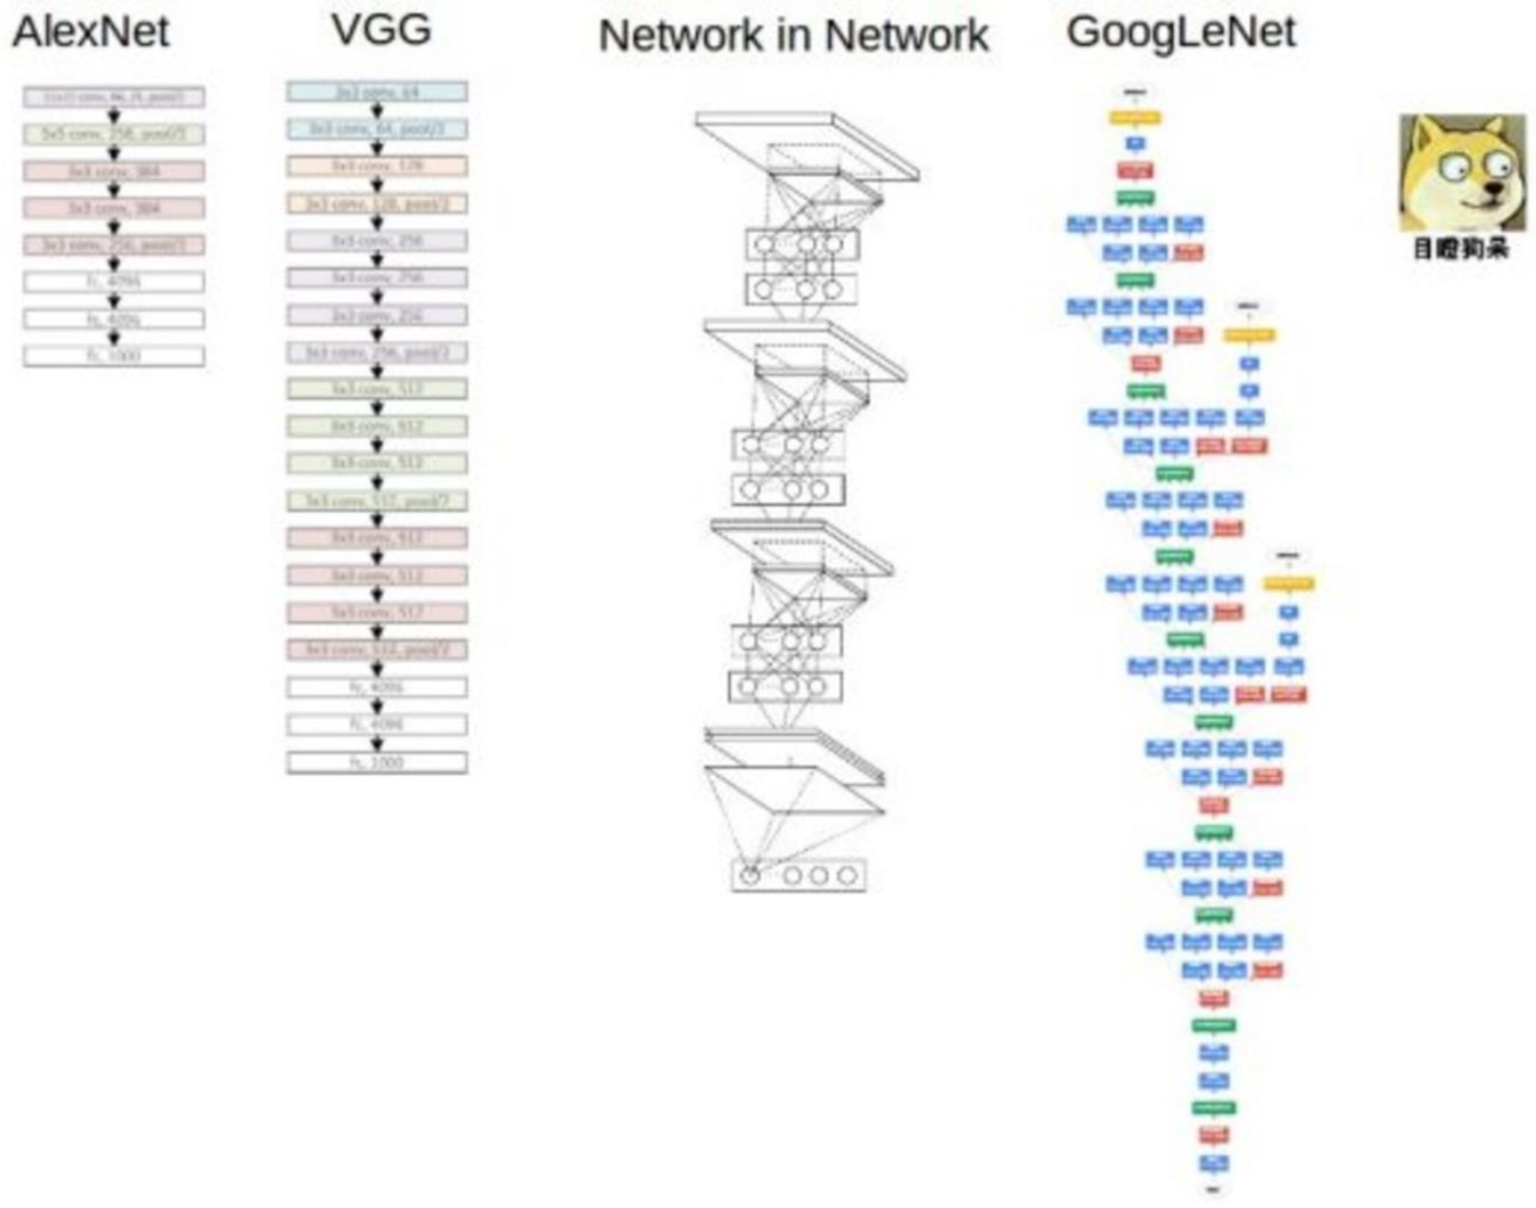

https://memoiry.me/2017/04/22/deep-learning-series-2-typical-CNN-net/

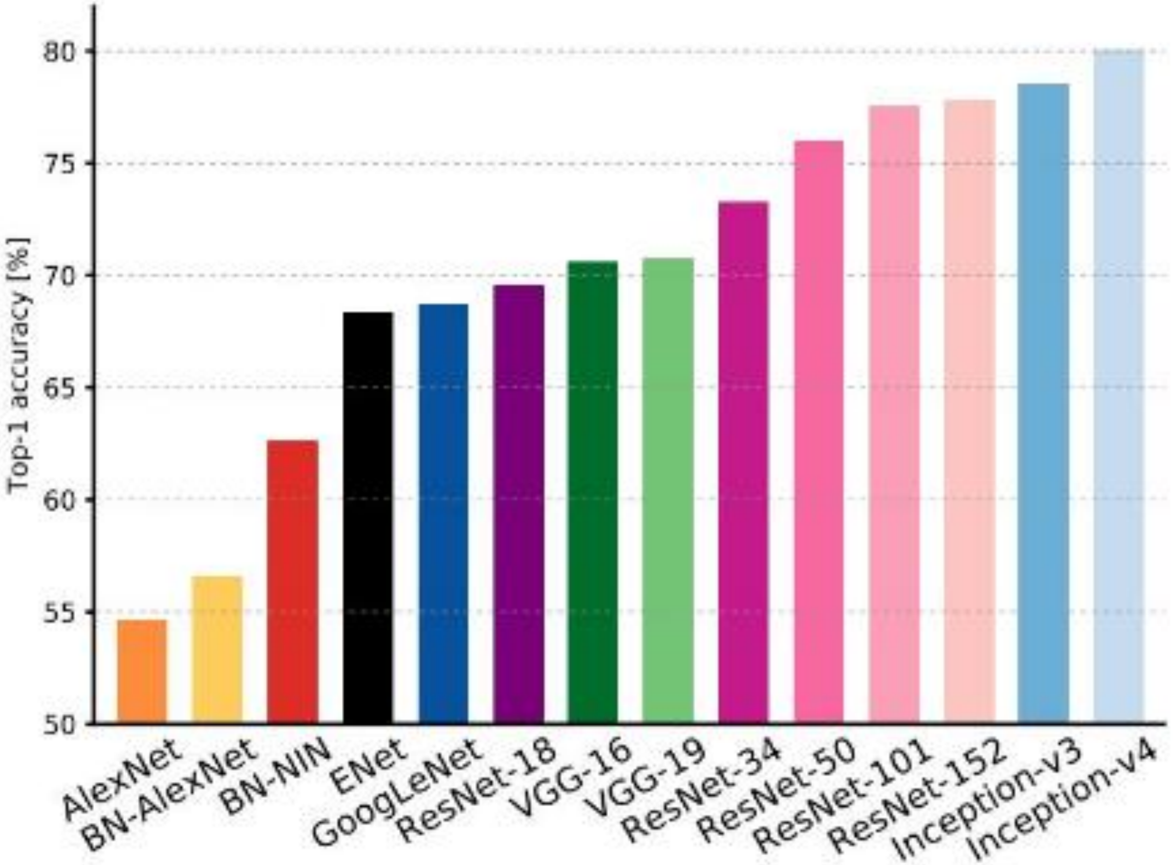

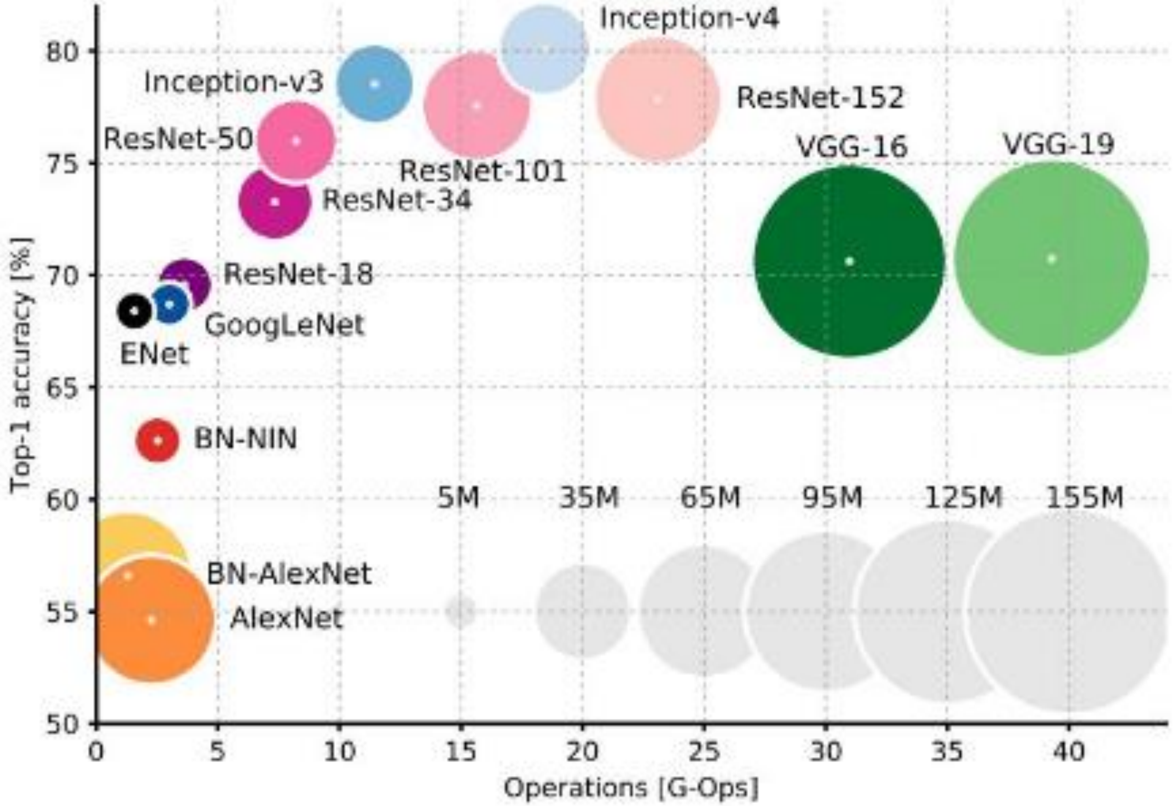

https://medium.com/analytics-vidhya/cnns-architectures-lenet-alexnet-vgg-googlenet-resnet-and-more-666091488df5

Recurrent Neural Networks (RNN)

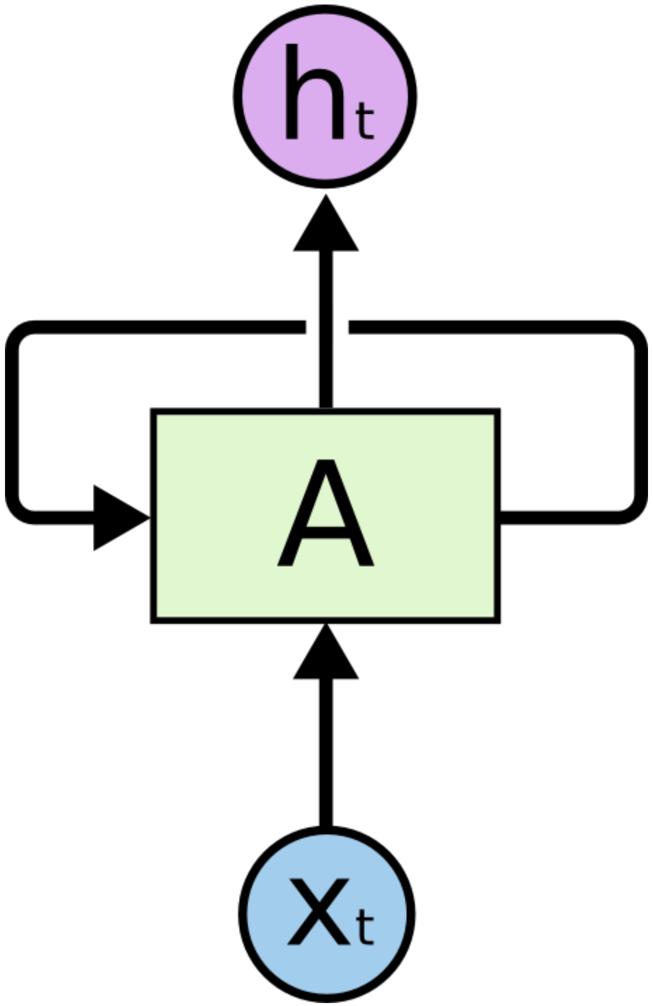

Unroll the RNN

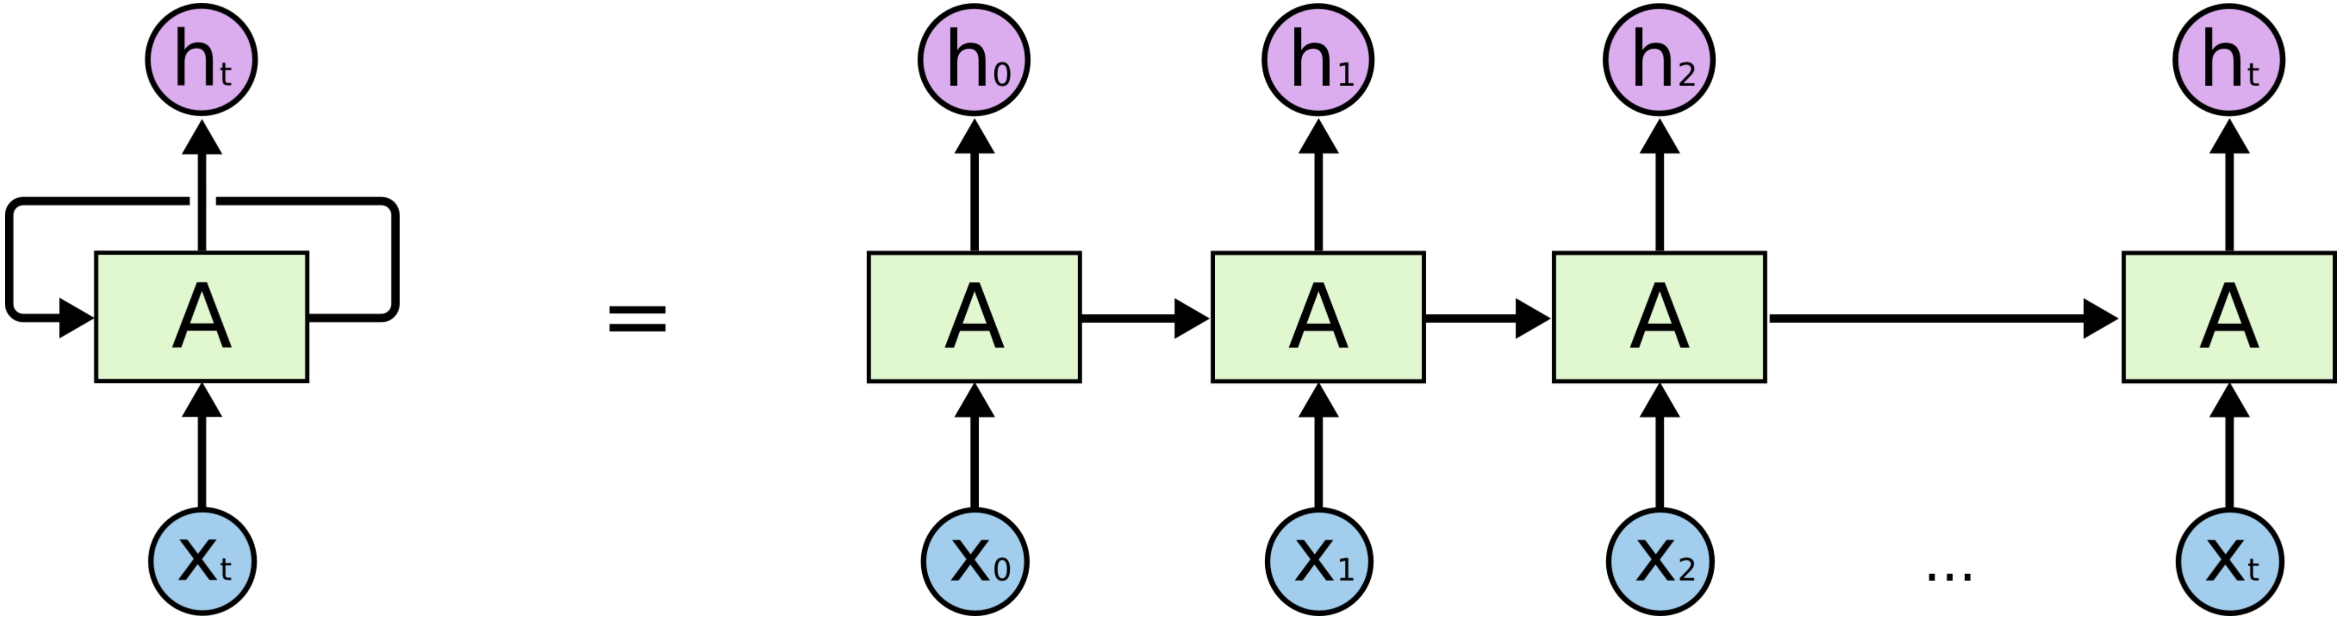

Long Short-term Memory (LSTM)

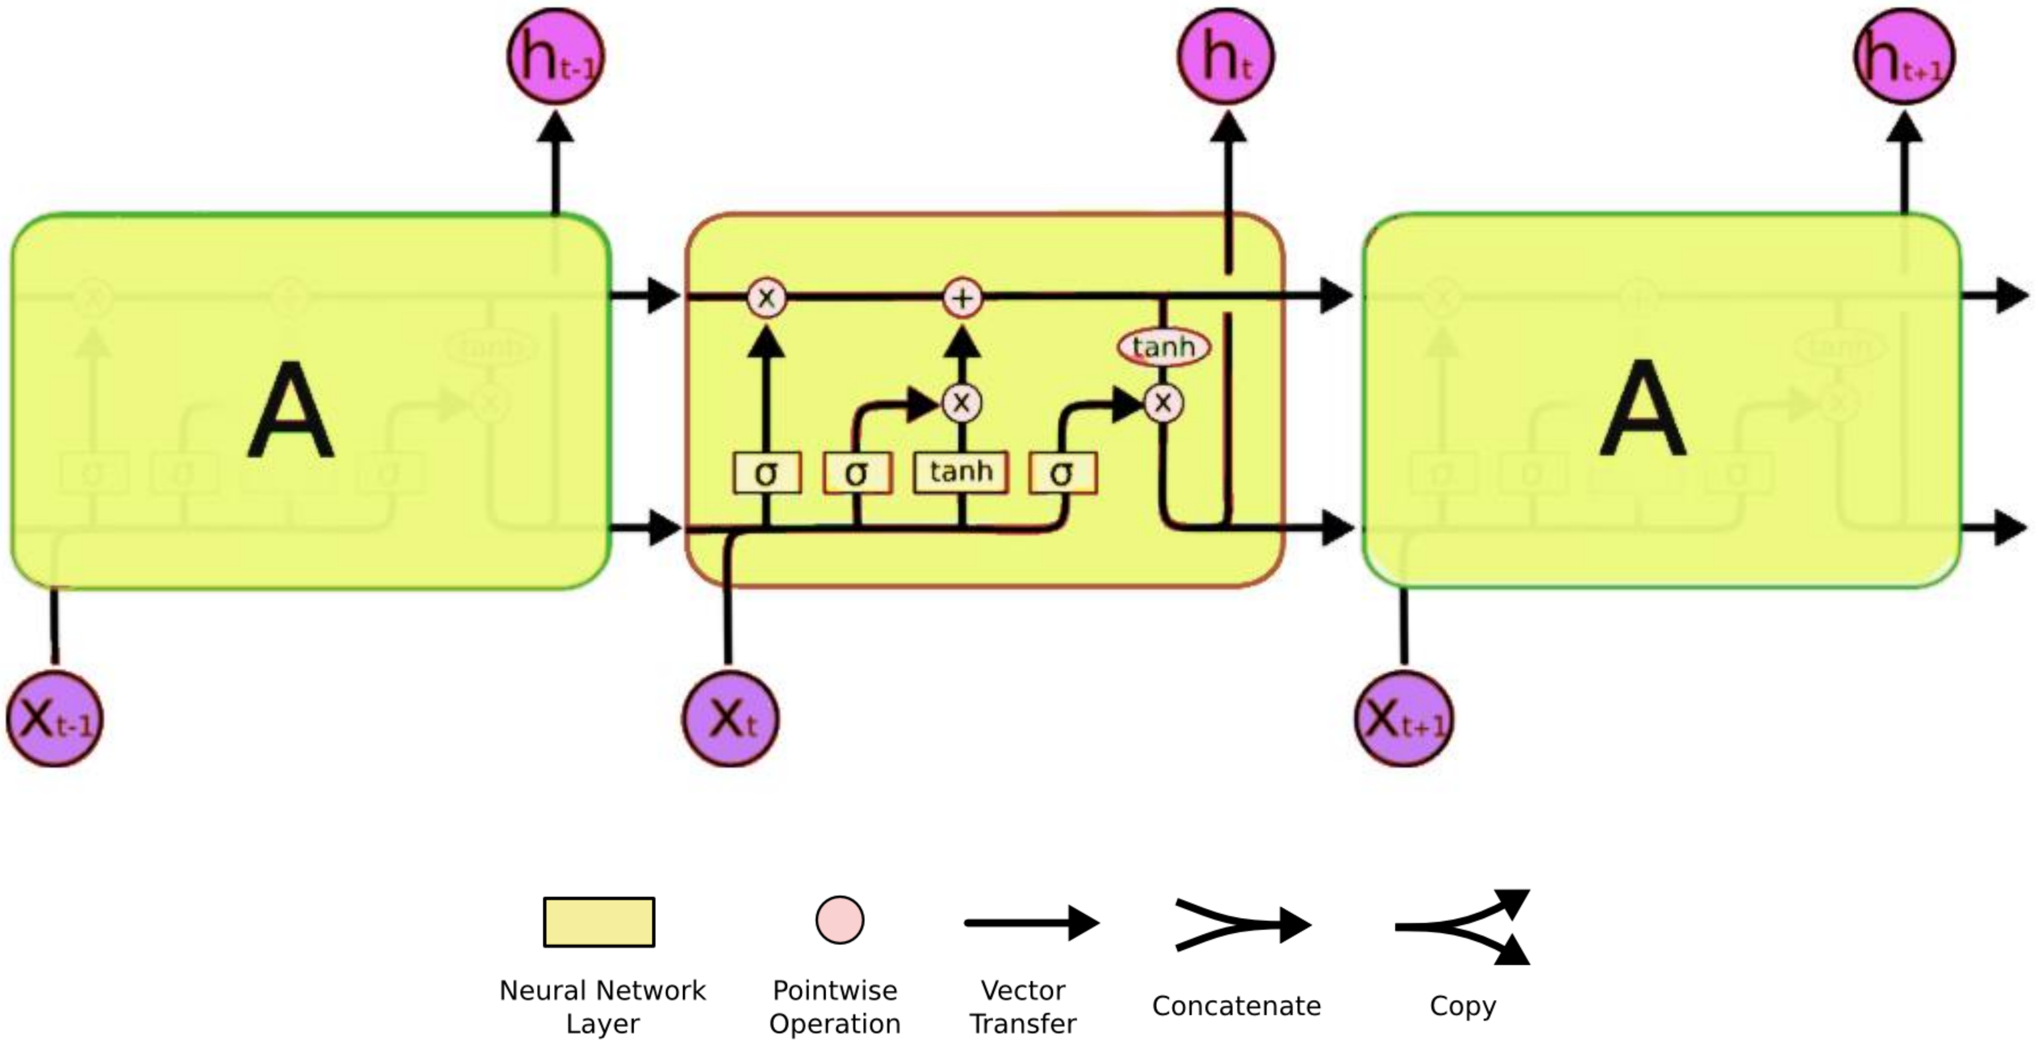

Deep Reinforcement Learning

Reinforcement Learning

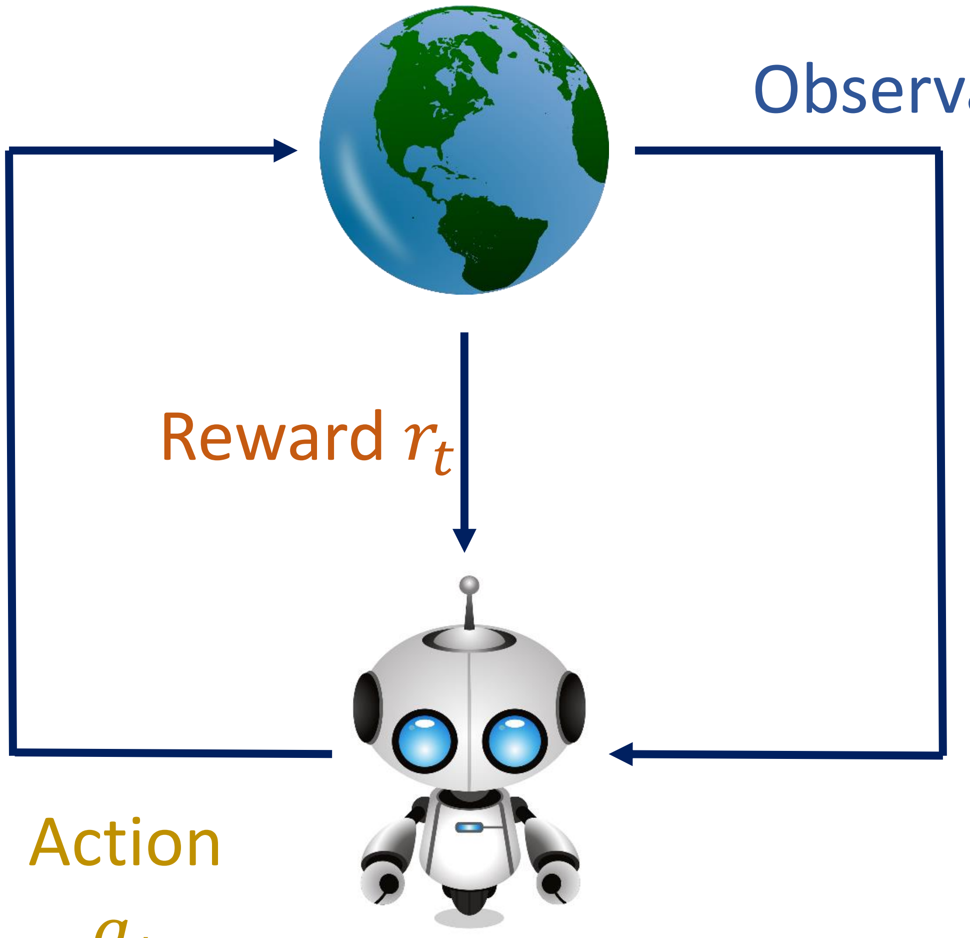

Google DeepMind

DRL in Atari

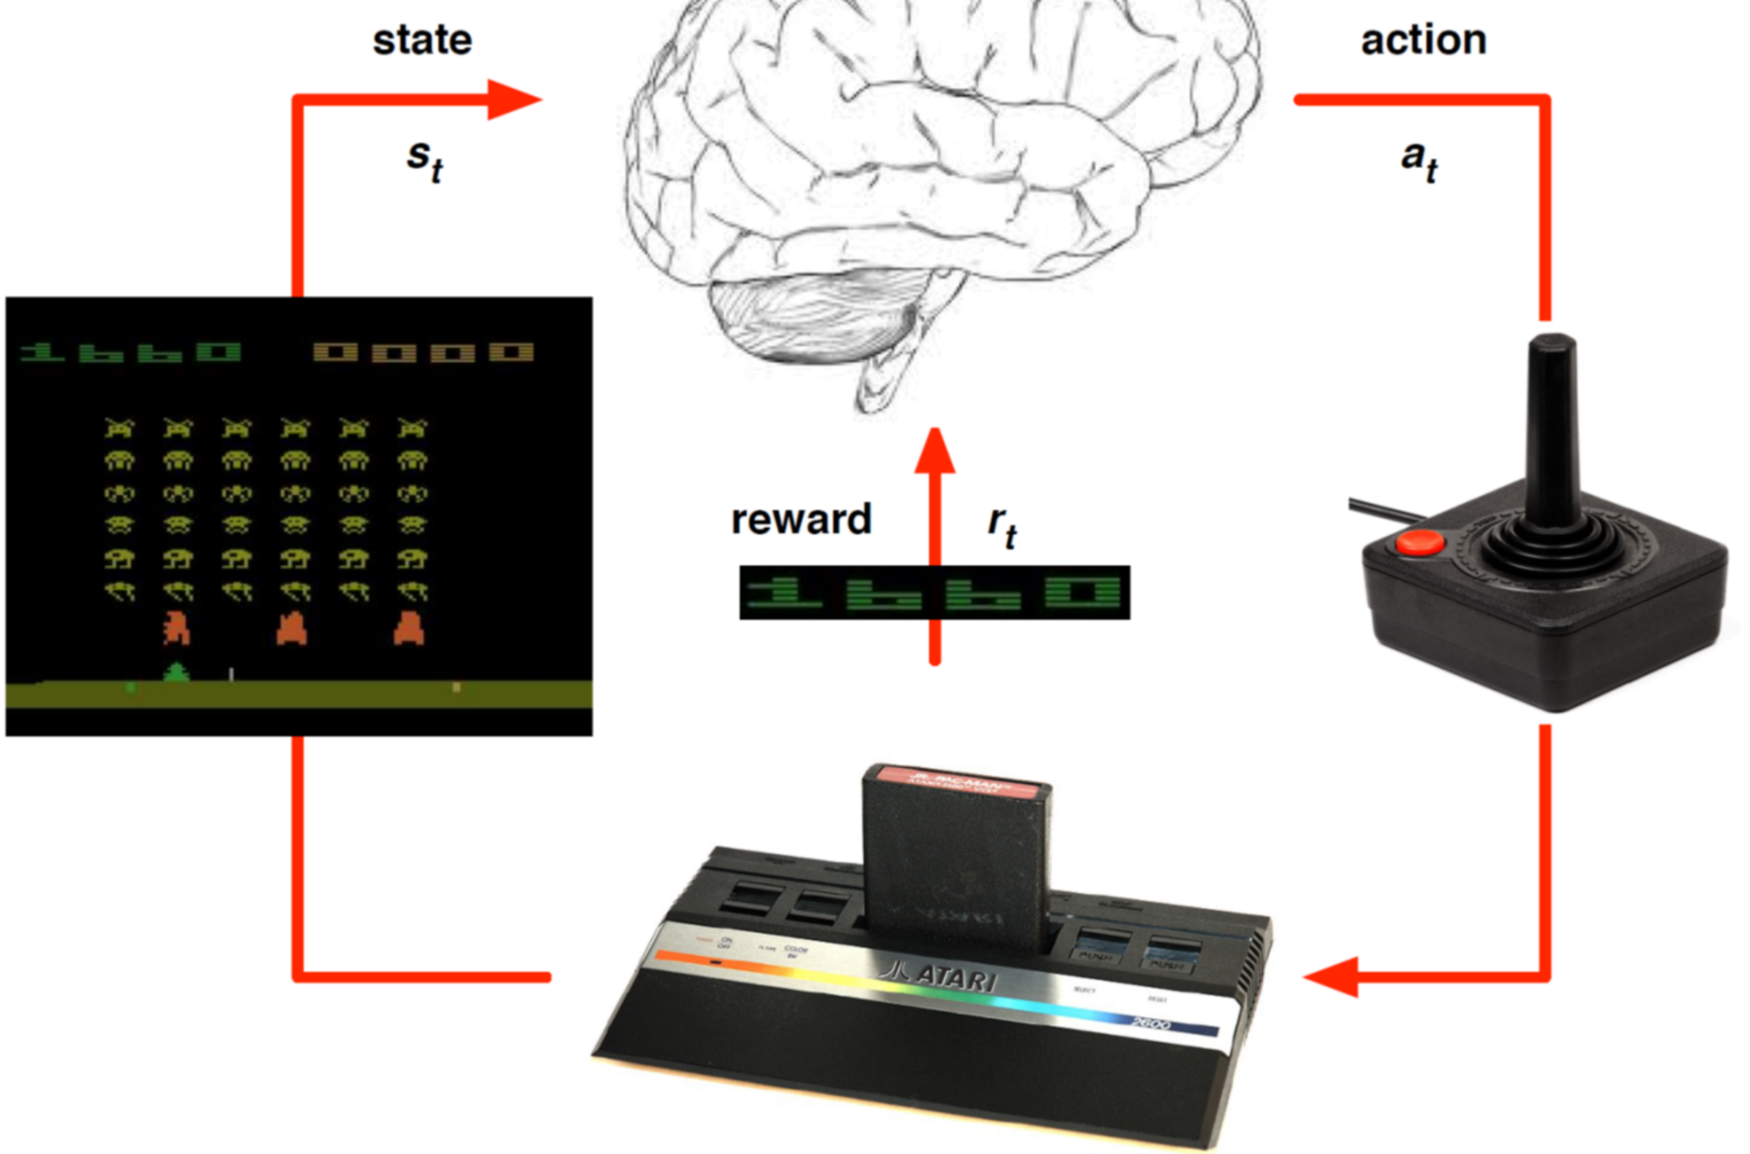

Mnih et al., “Human Level Control through Deep Reinforcement Learning,” *Nature*, 2015

Learning to Play Atari Games

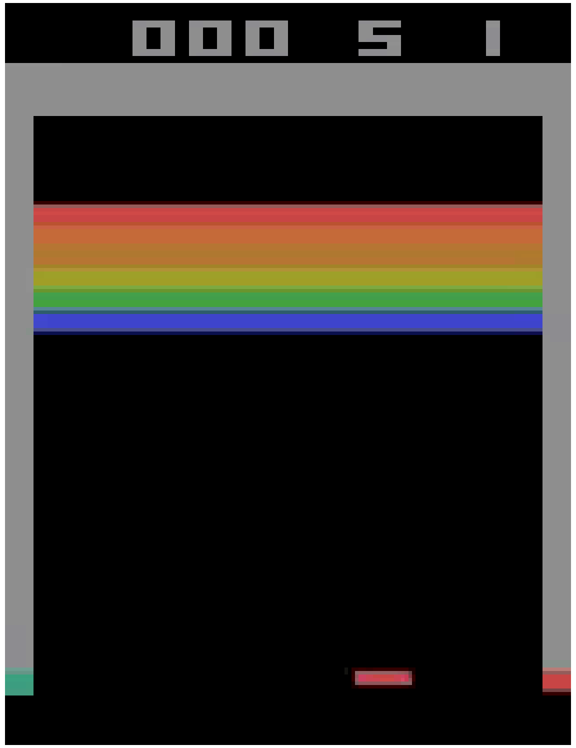

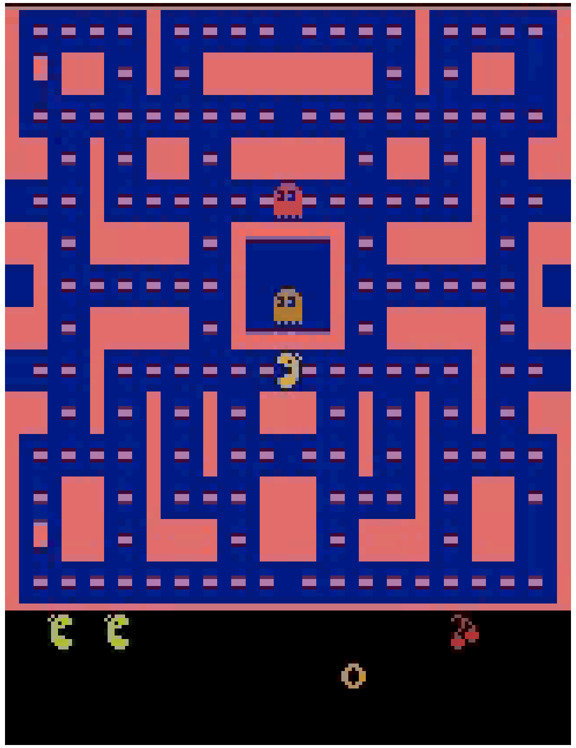

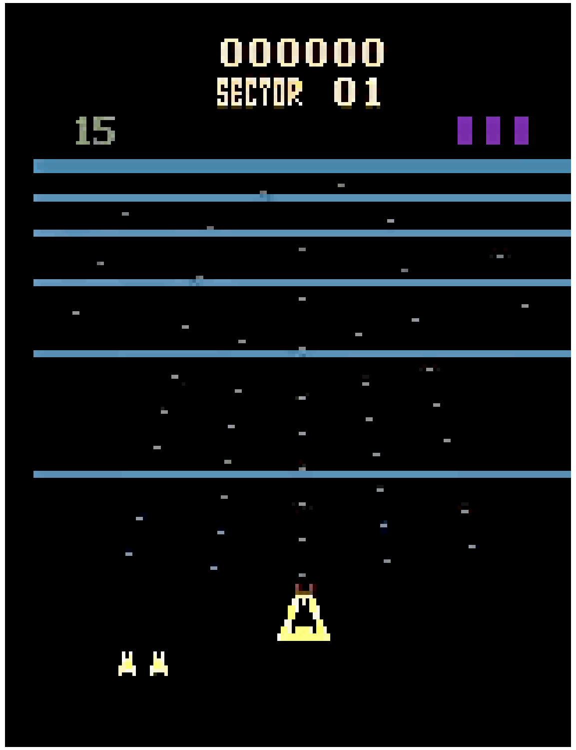

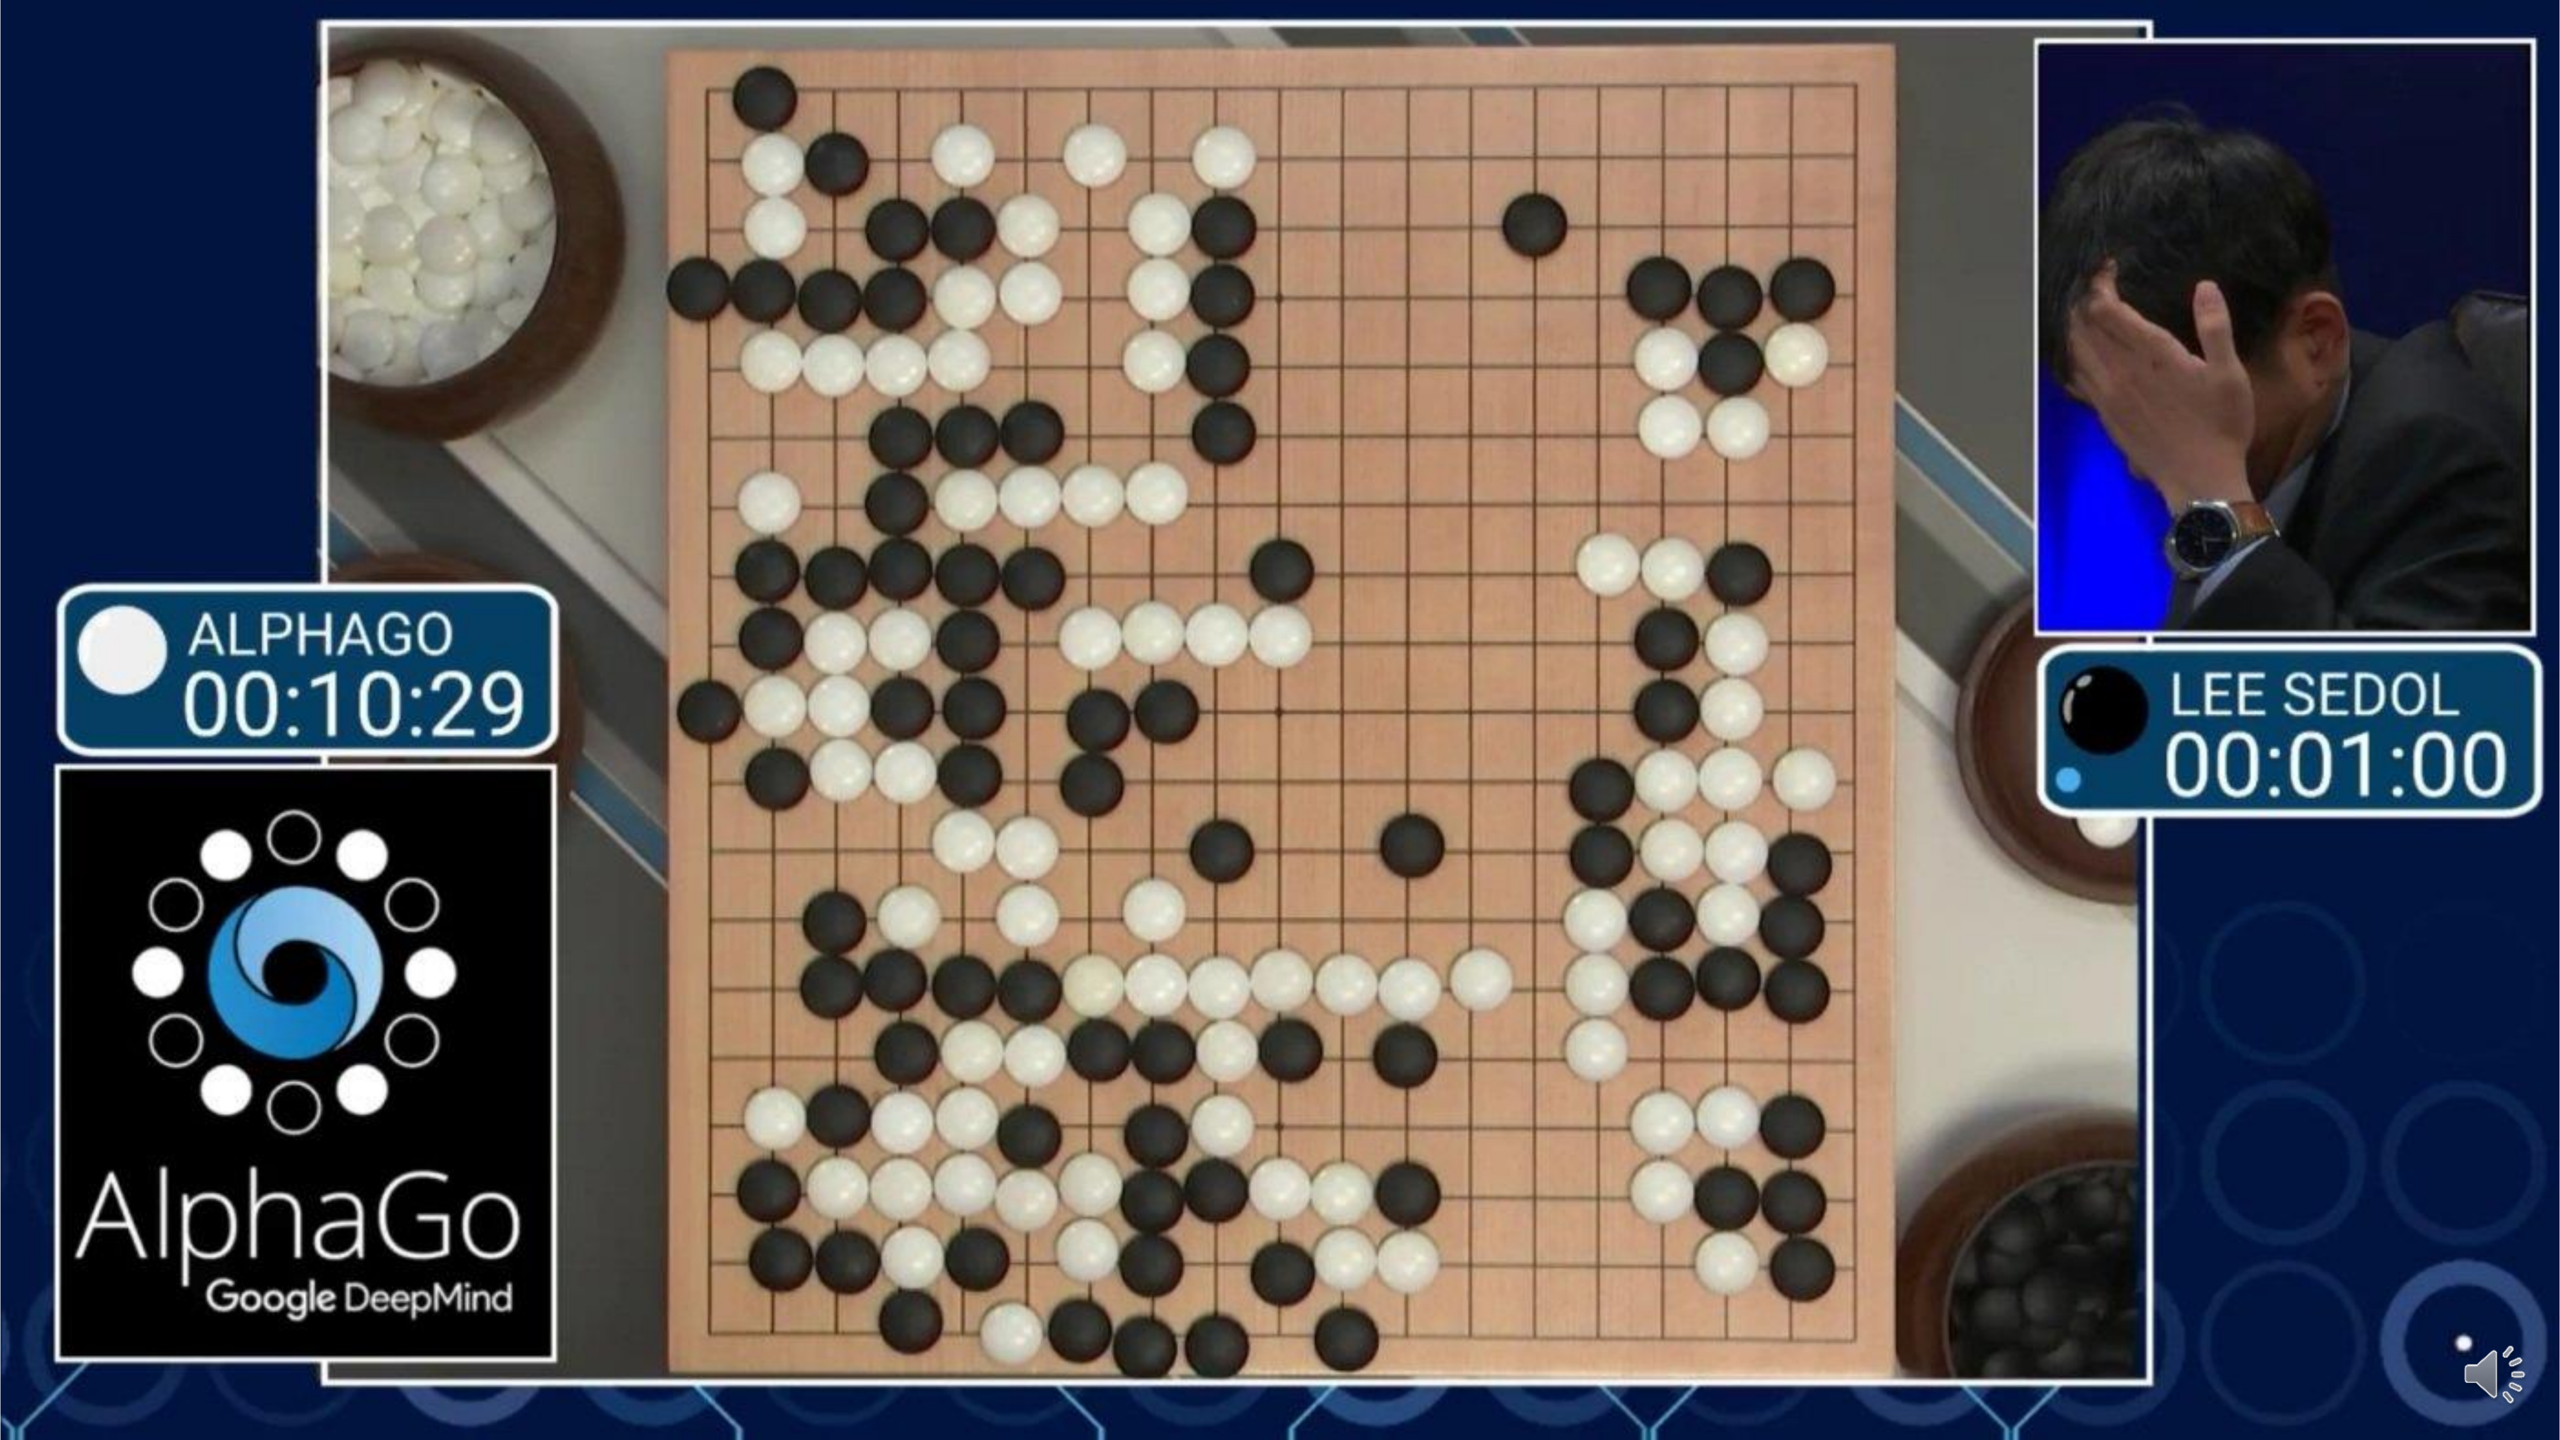
ALPHAGO
00:10:29
LEE SEDOL
00:01:00
AlphaGo
Google DeepMind

Dr. Aja Huang (黃士杰)

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

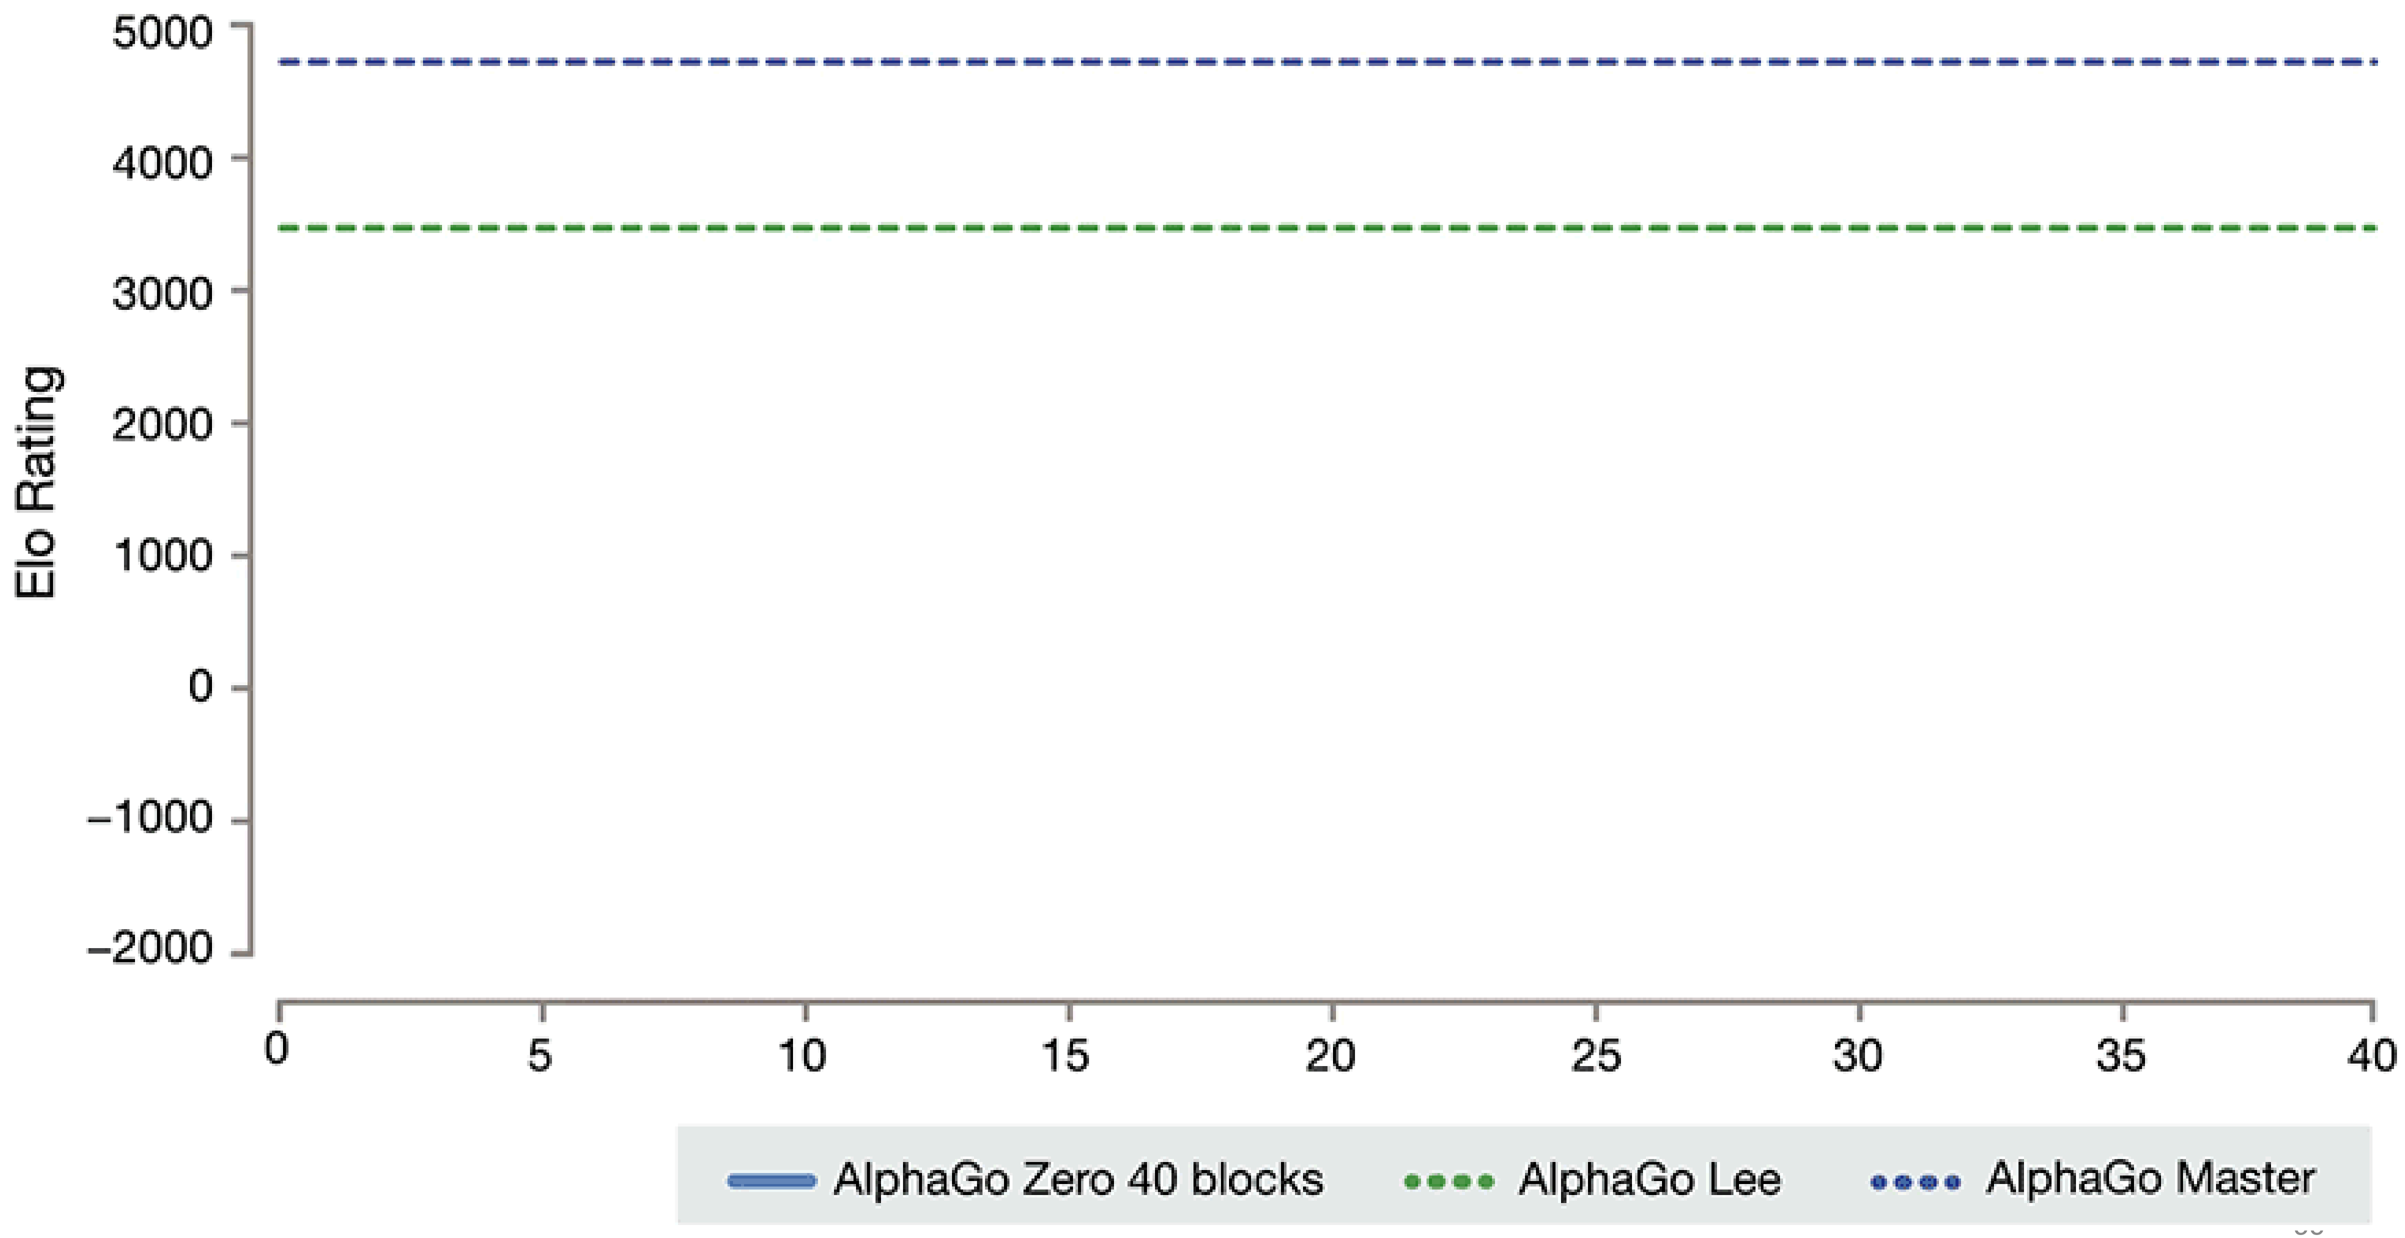

Elo Rating
5000
4000
3000
2000
1000
0
−1000
−2000
0
5
10
15
20
25
30
35
40
AlphaGo Zero 40 blocks
AlphaGo Lee
AlphaGo Master


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

Human Extinction ?

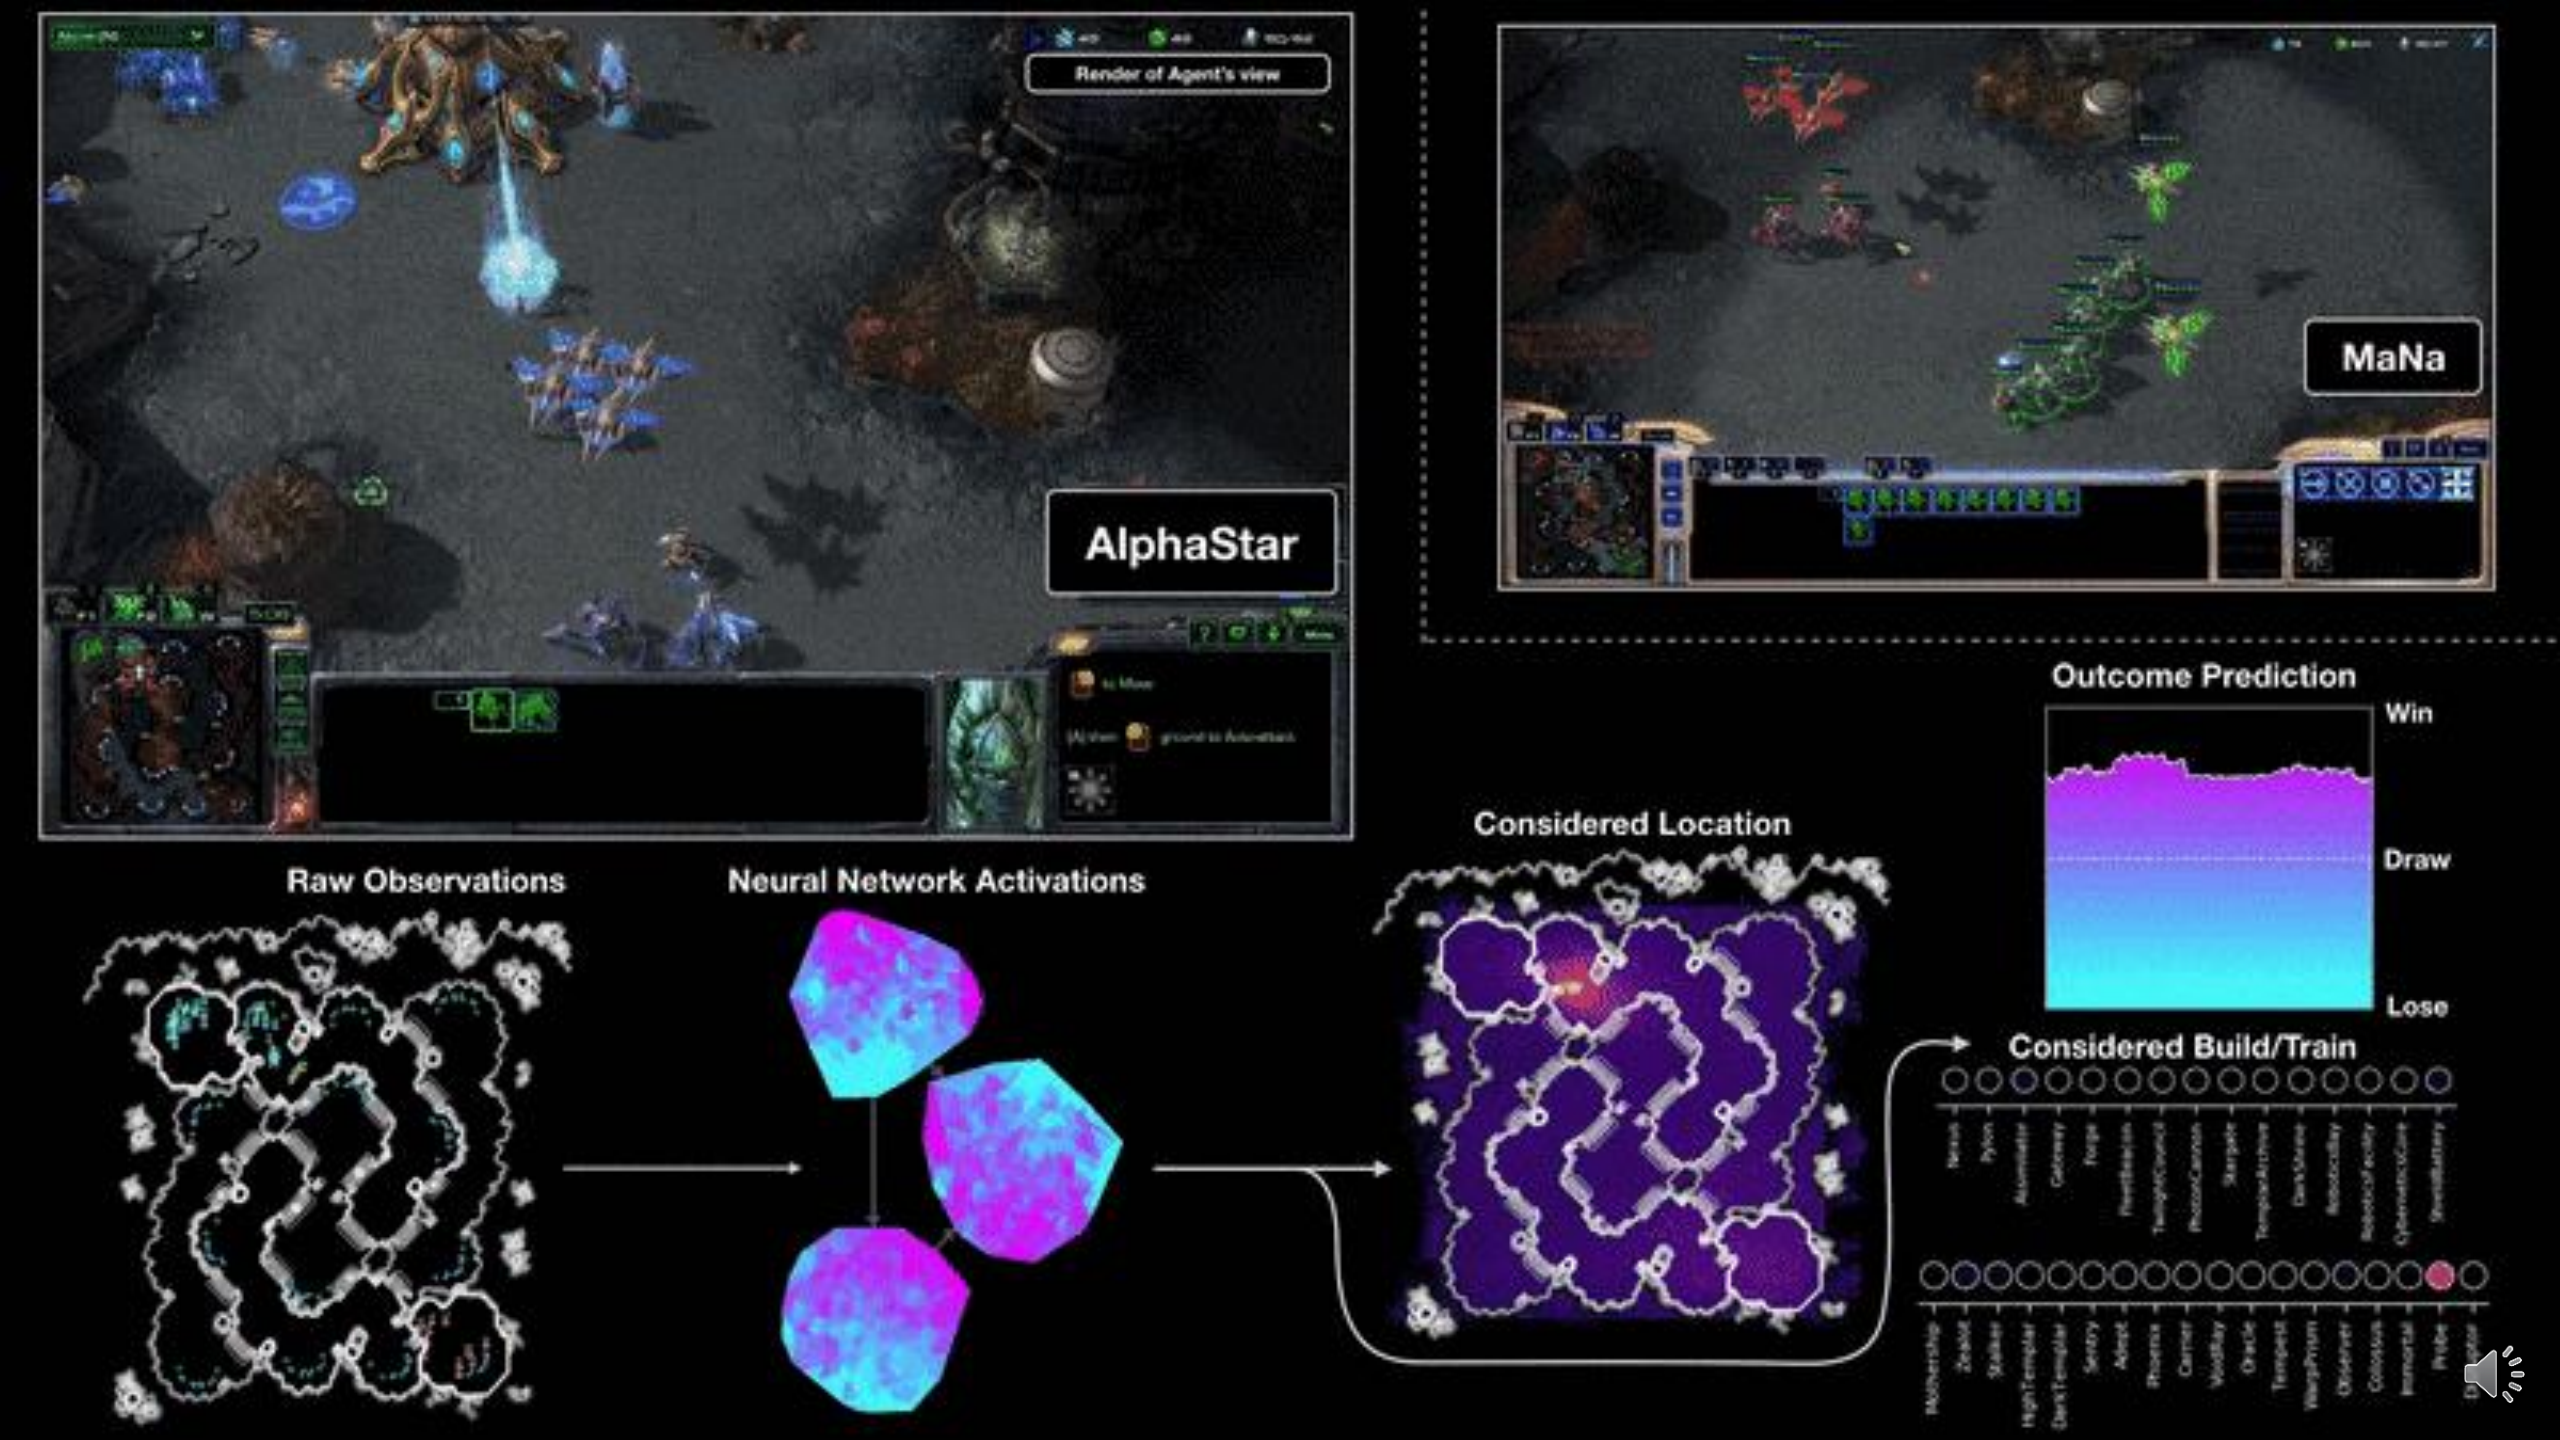

Render of Agent's view
AlphaStar
MaNa
Outcome Prediction
Win
Draw
Lose
Raw Observations
Neural Network Activations
Considered Location
Considered Build/Train

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

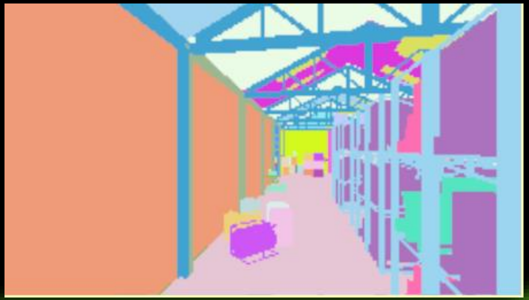

Depth Prediction

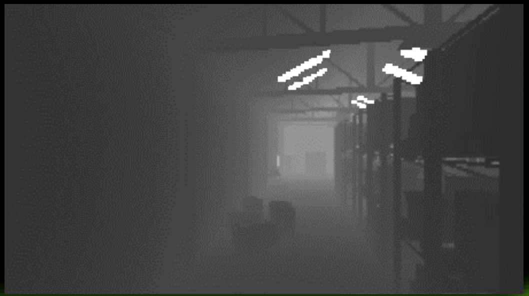

VIVID

Autonomous Navigation

Action Recognition

Simulate Real-life Events

Searching for the Shooter

Limits of Deep Learning

No Idea of Real World

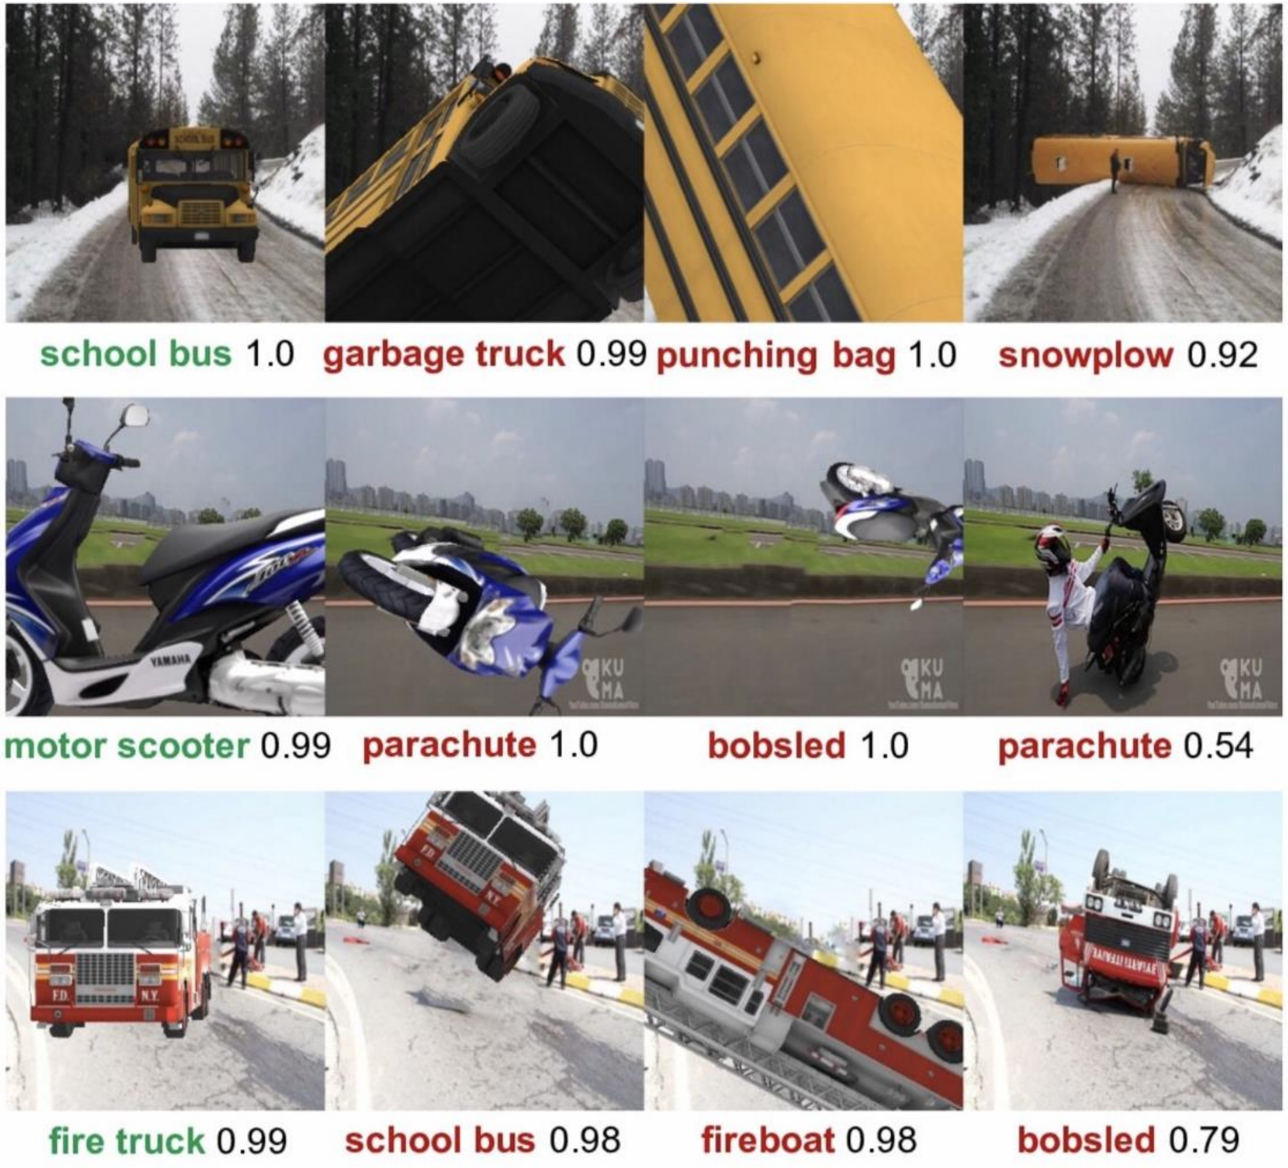

Alcorn et al., “Strike (with) a pose: Neural networks are easily fooled by strange poses of familiar objects,”CVPR 2019.

Adversarial Attack

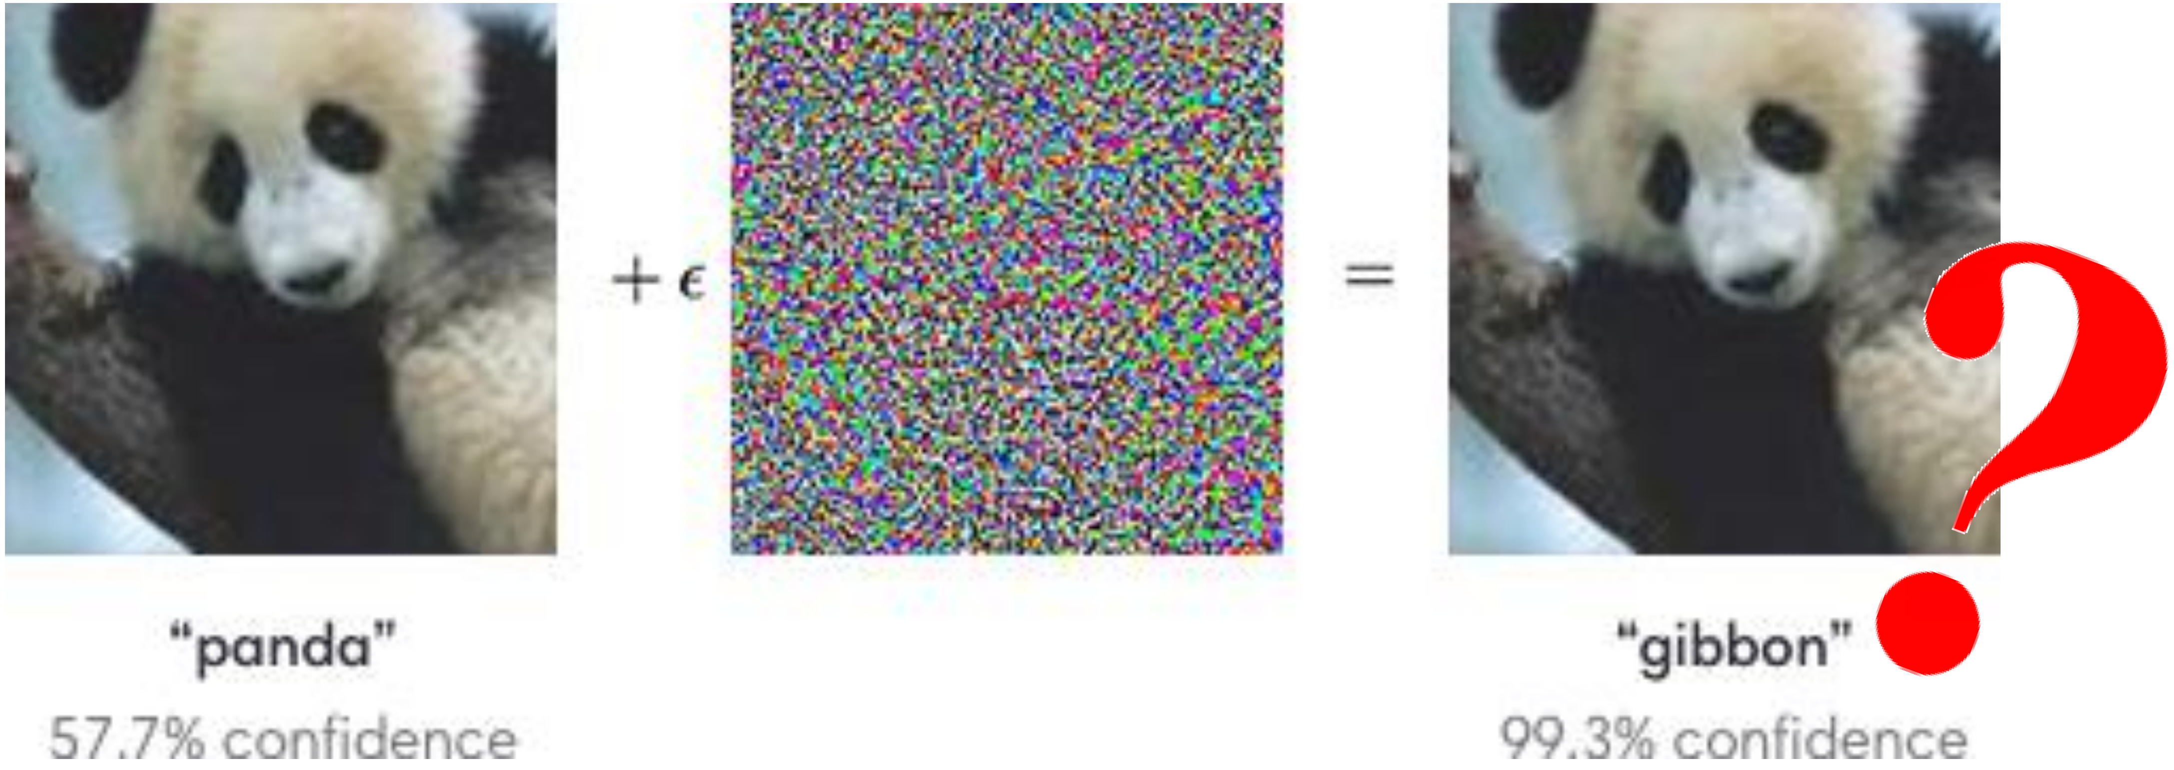

Generative Adversarial Networks (GAN)

- Ian Goodfellow

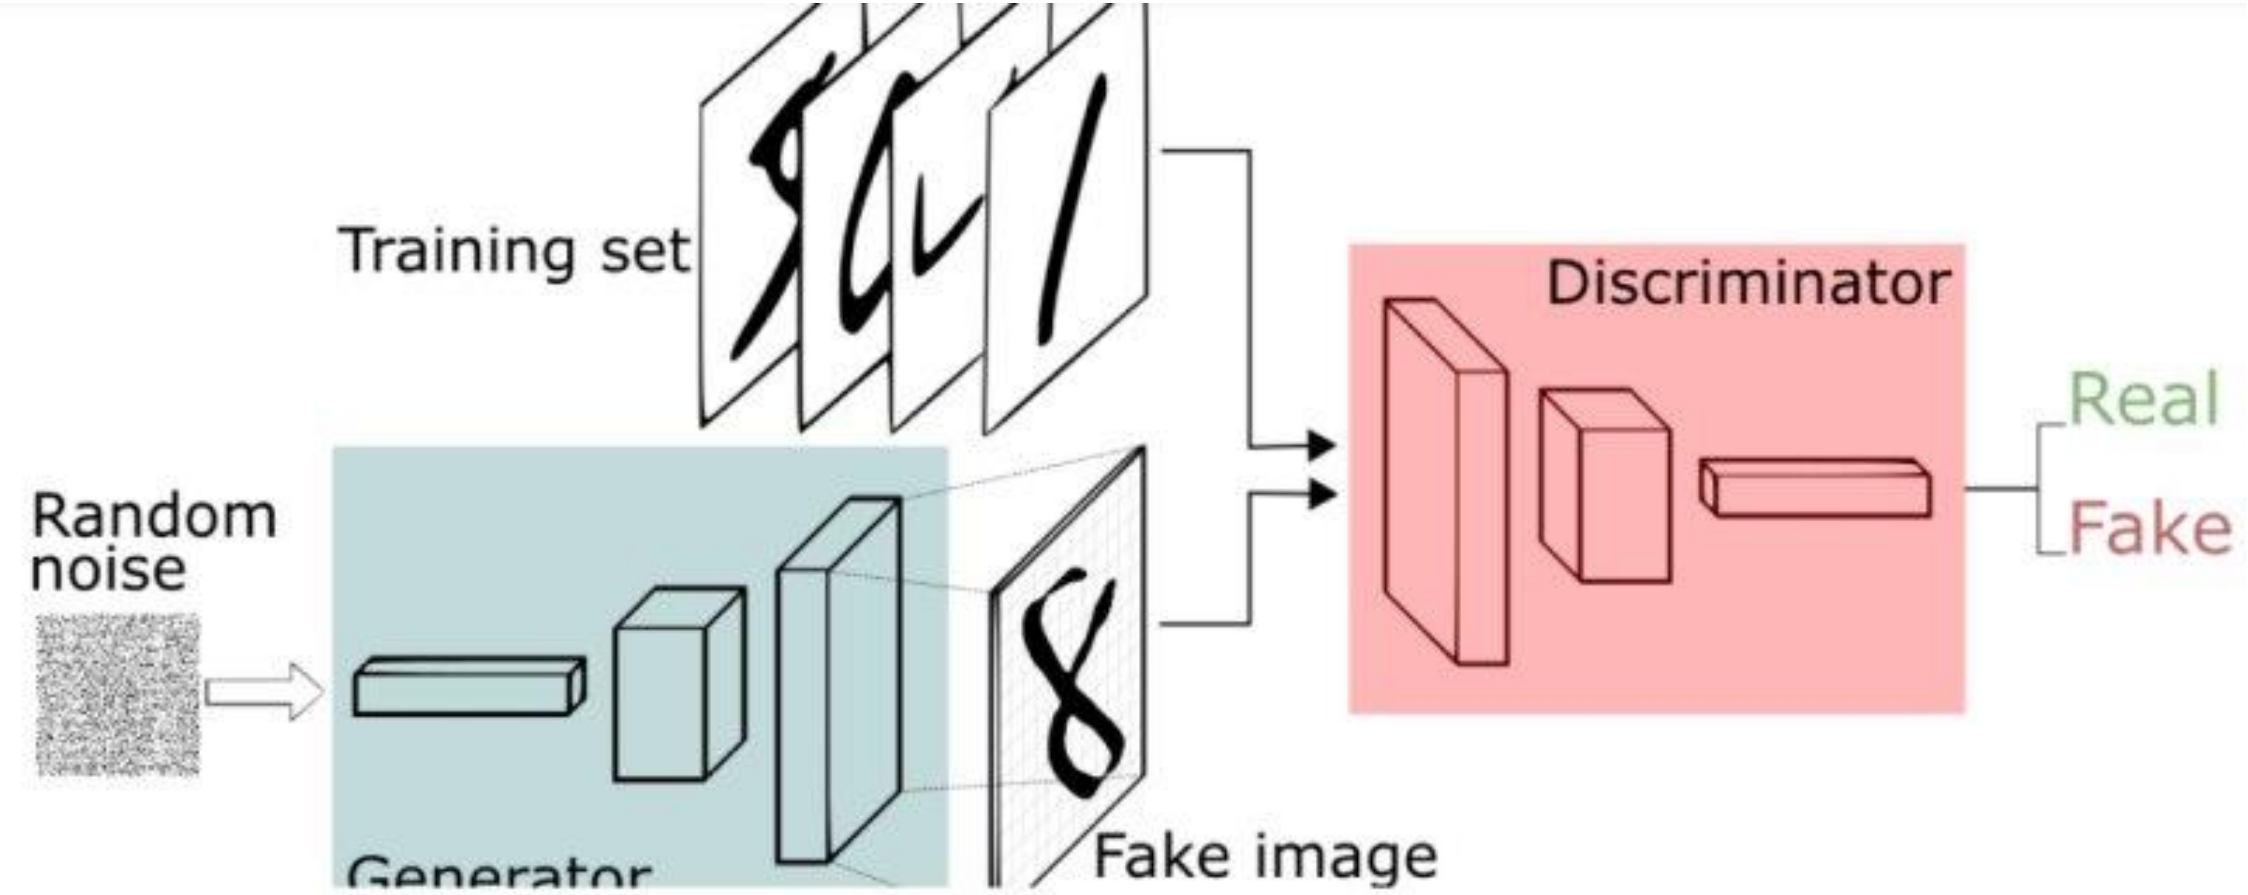

Painting like Van Gogh

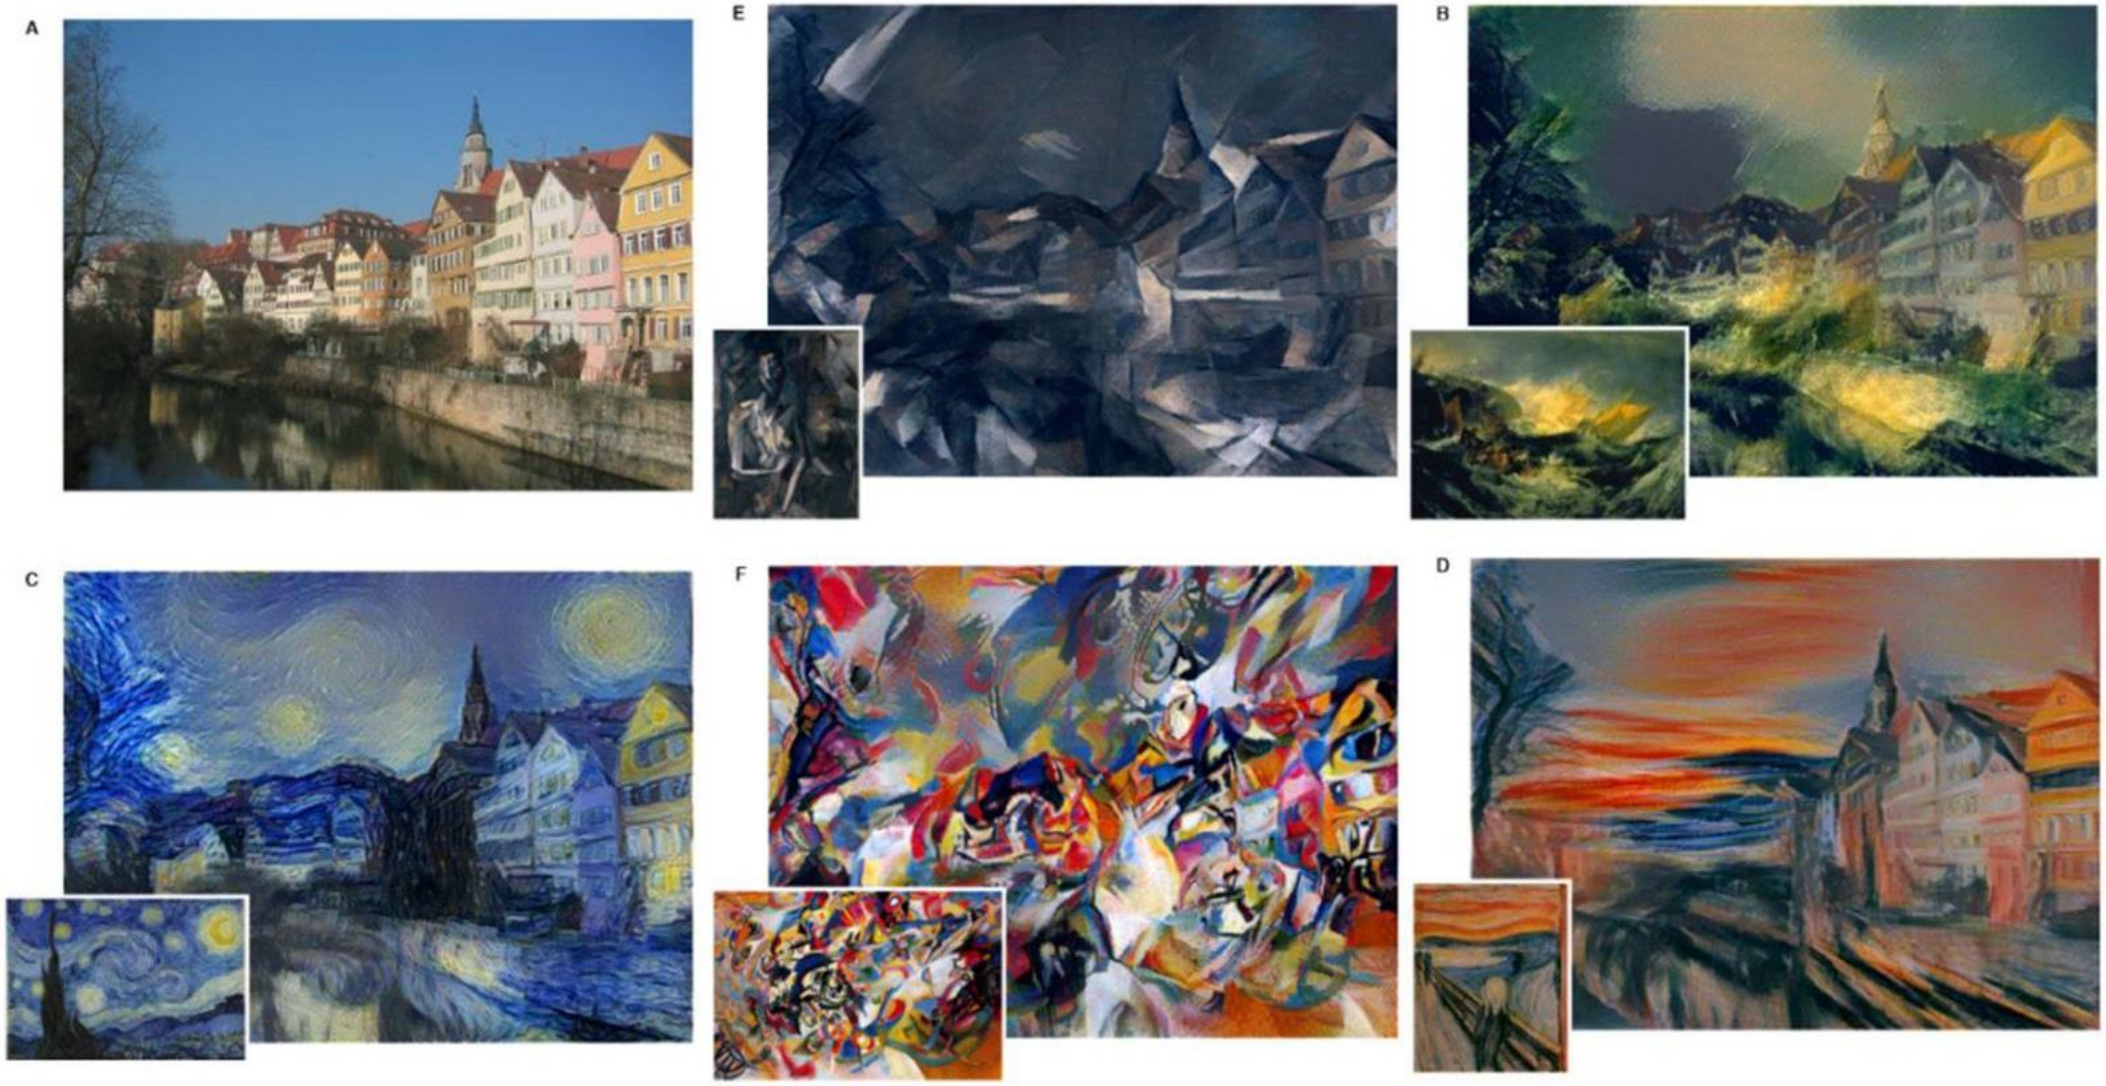

Super Resolution

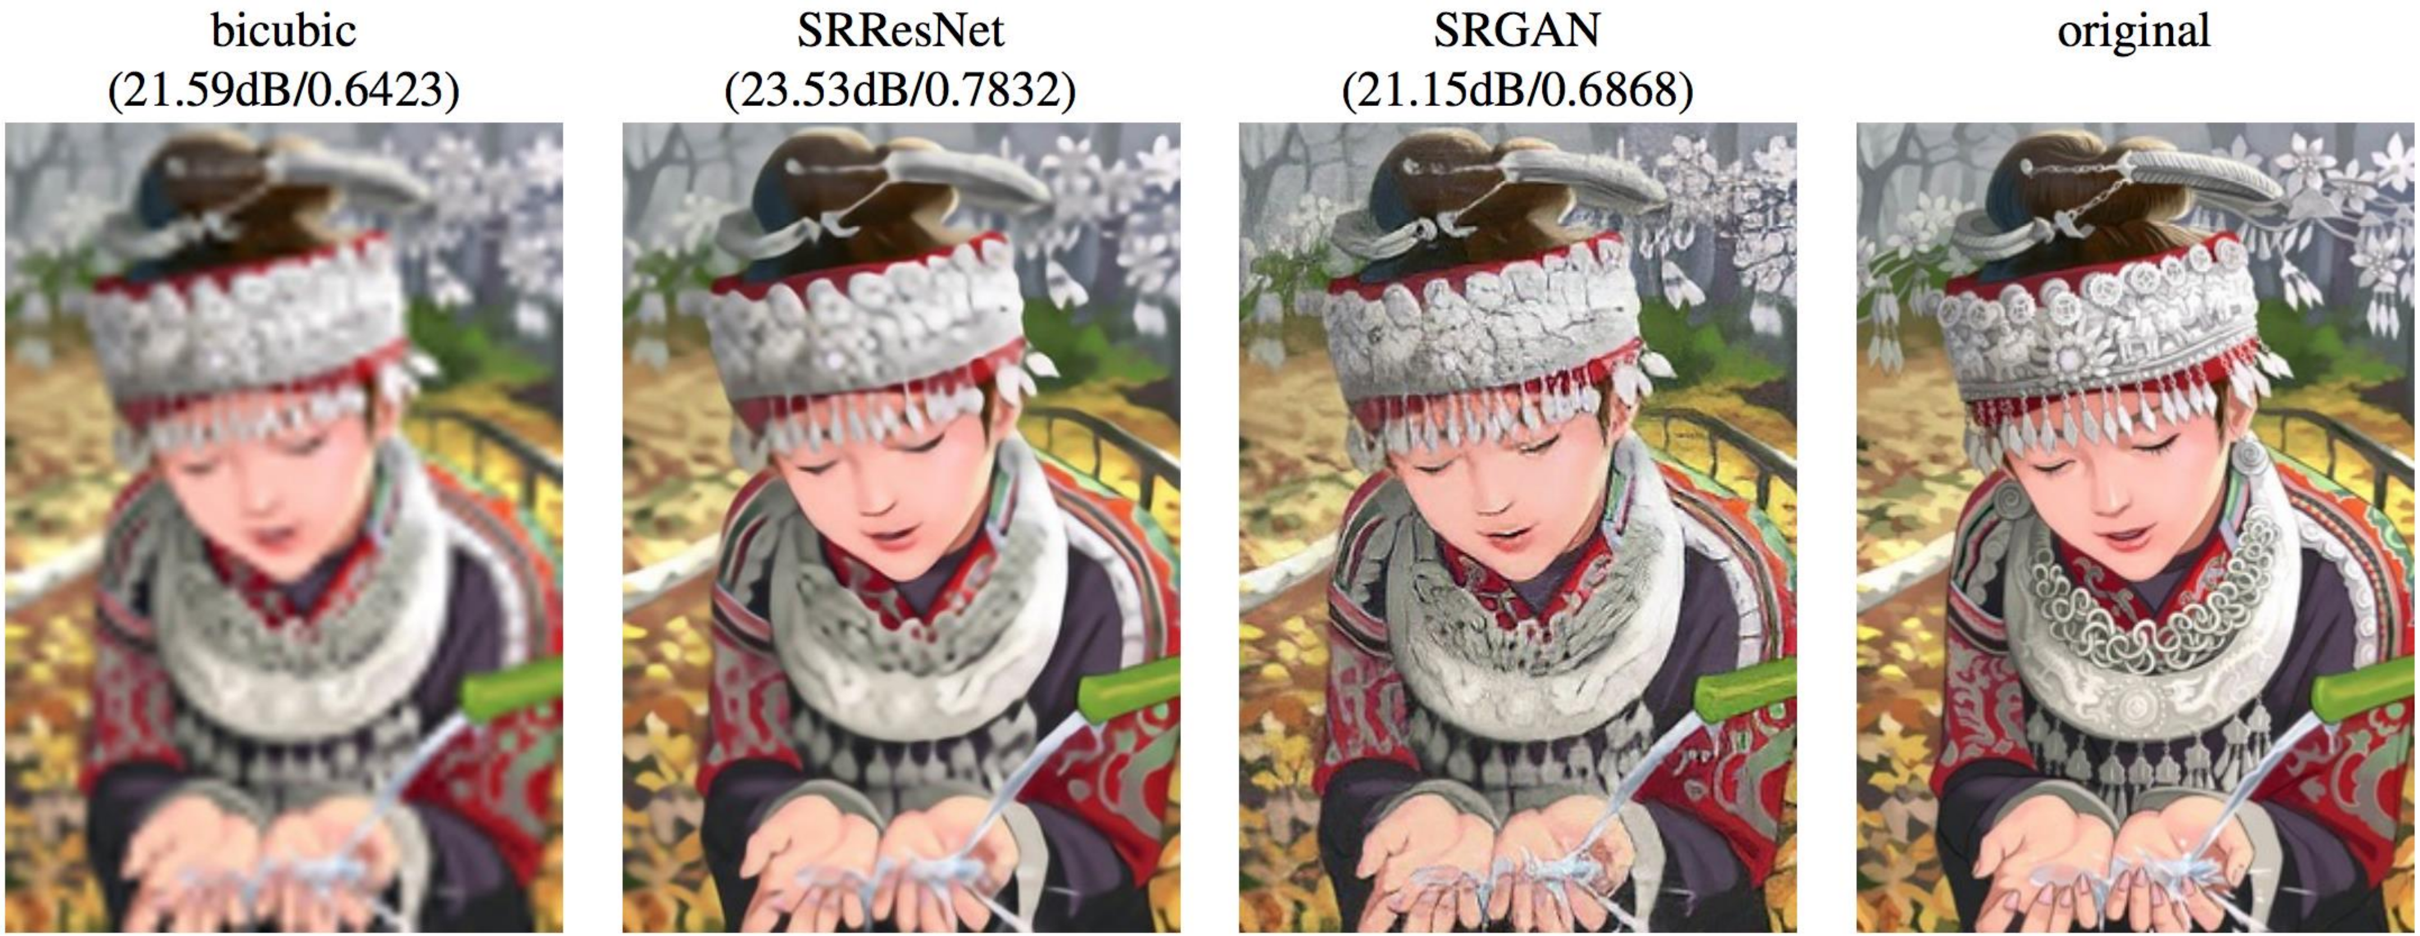

Figure 2: From left to right: bicubic interpolation, deep residual network optimized for MSE, deep residual generative adversarial network optimized for a loss more sensitive to human perception, original HR image. Corresponding PSNR and SSIM are shown in brackets. [4× upscaling]

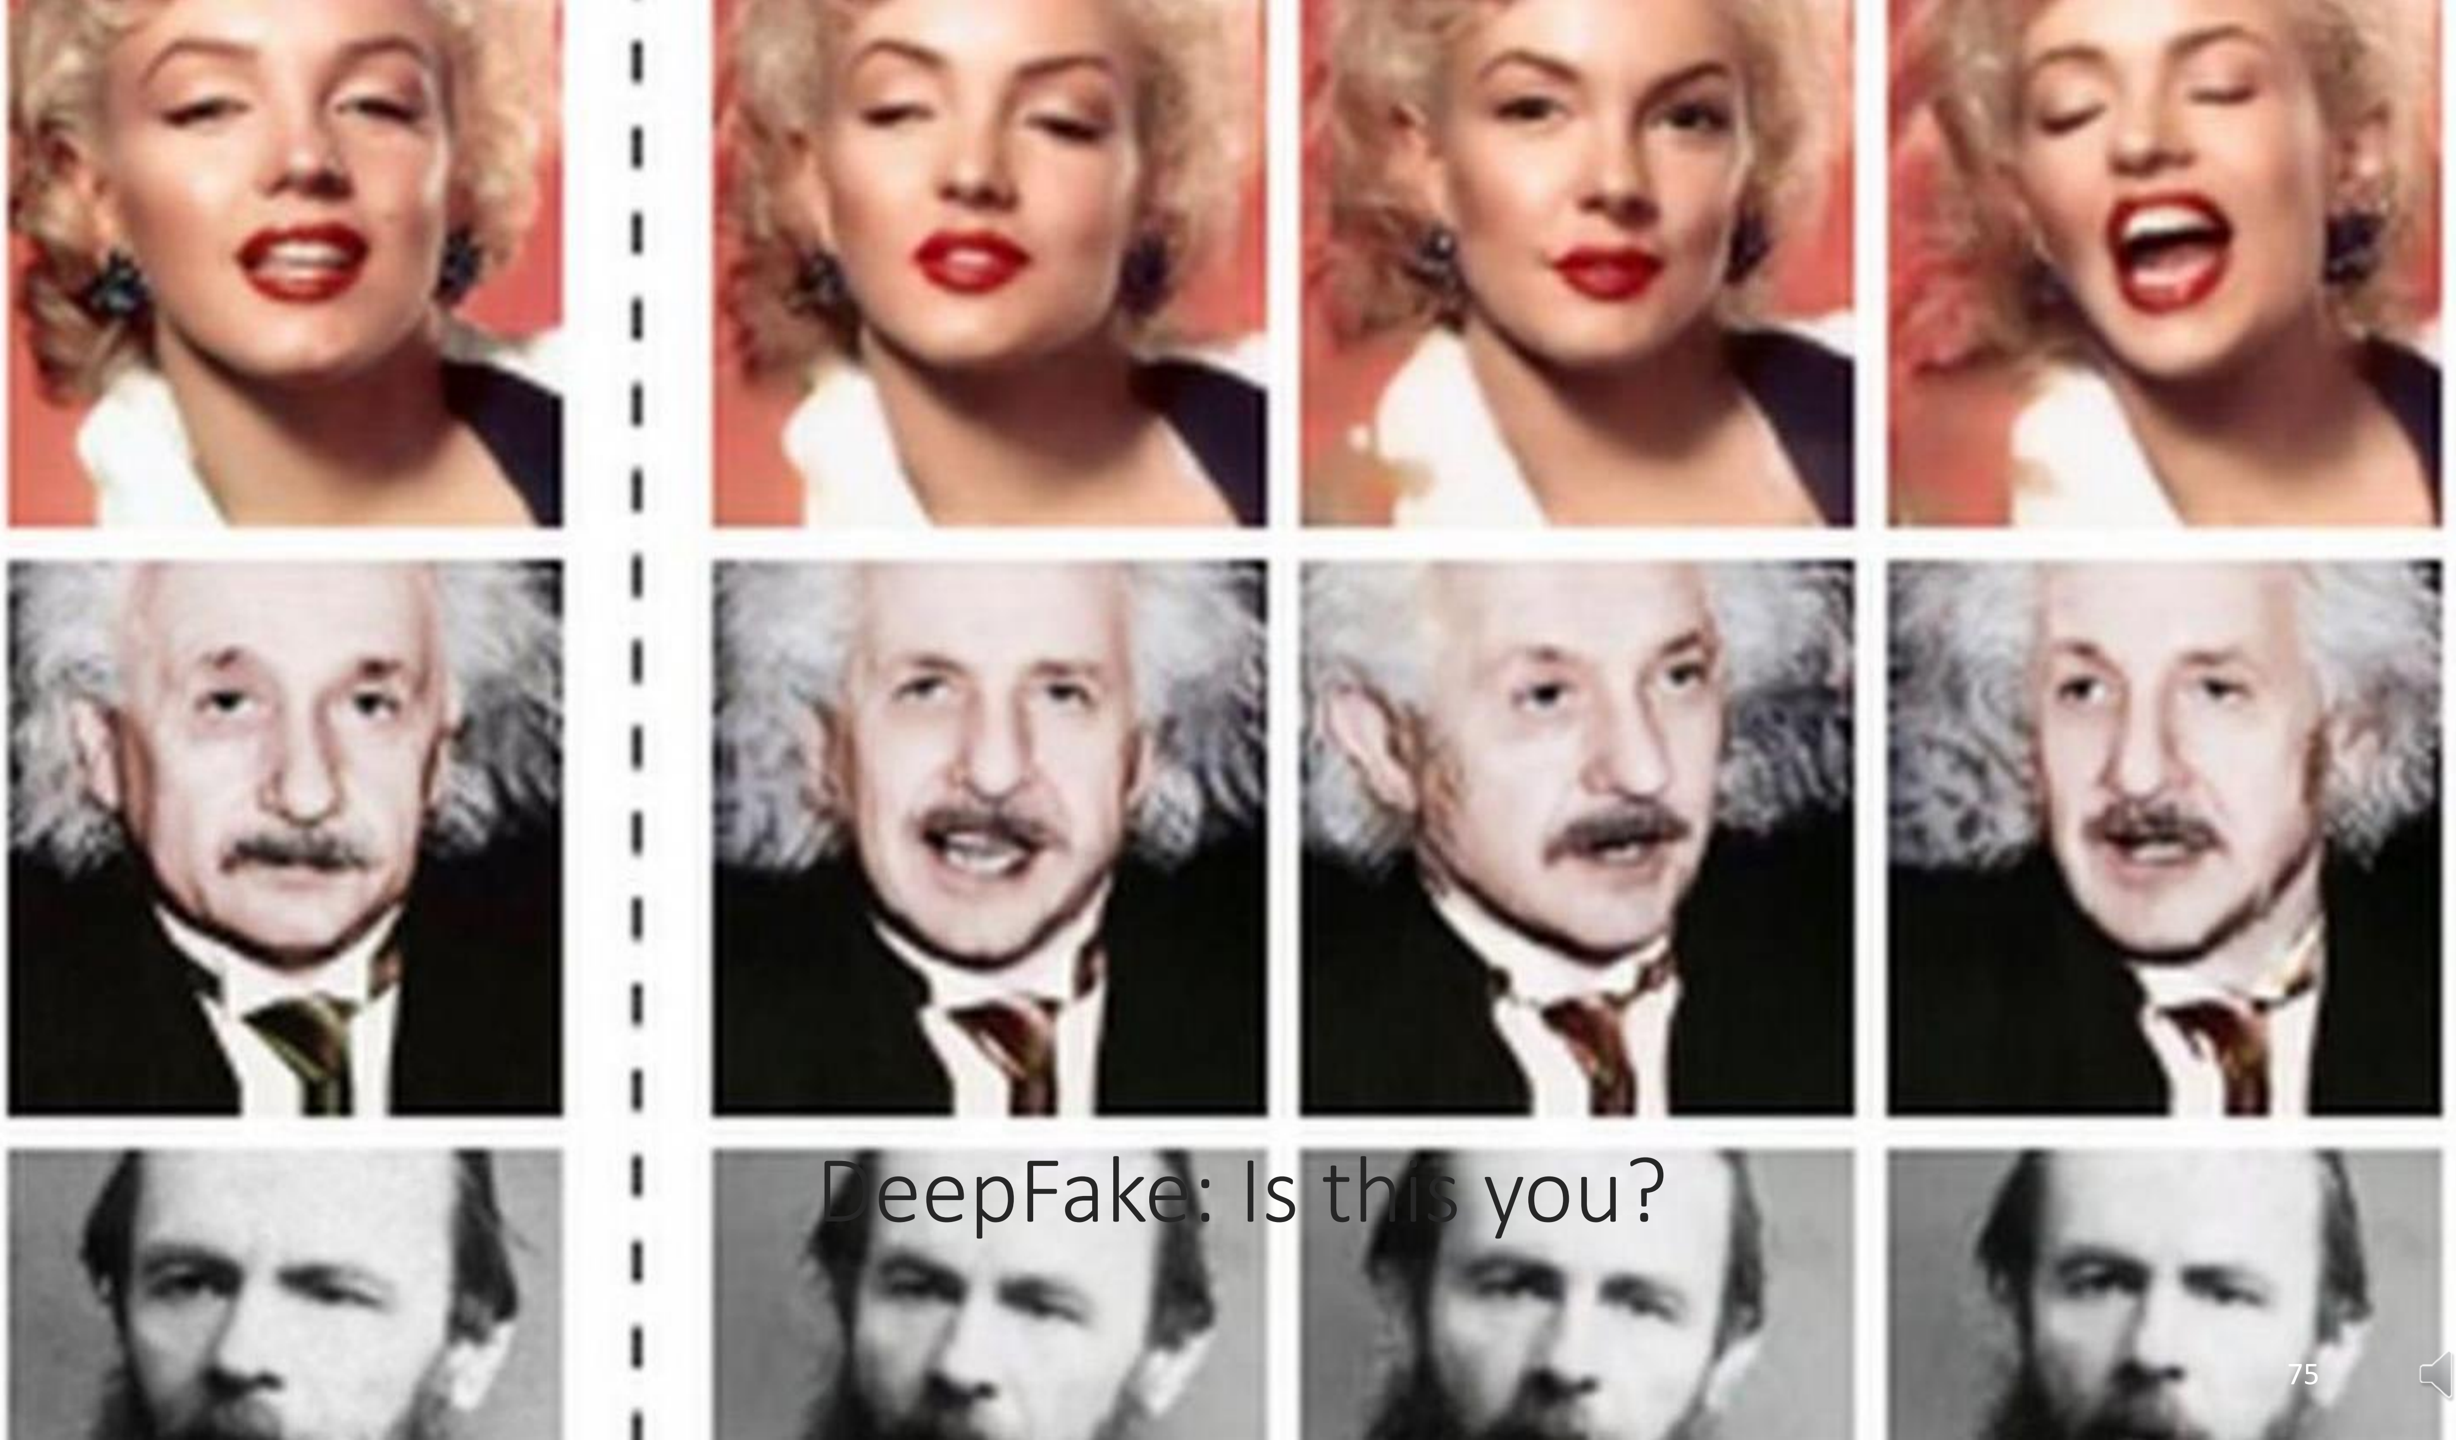

DeepFake: Is this you?

Google's AutoML

- Learning neural network cells automatically

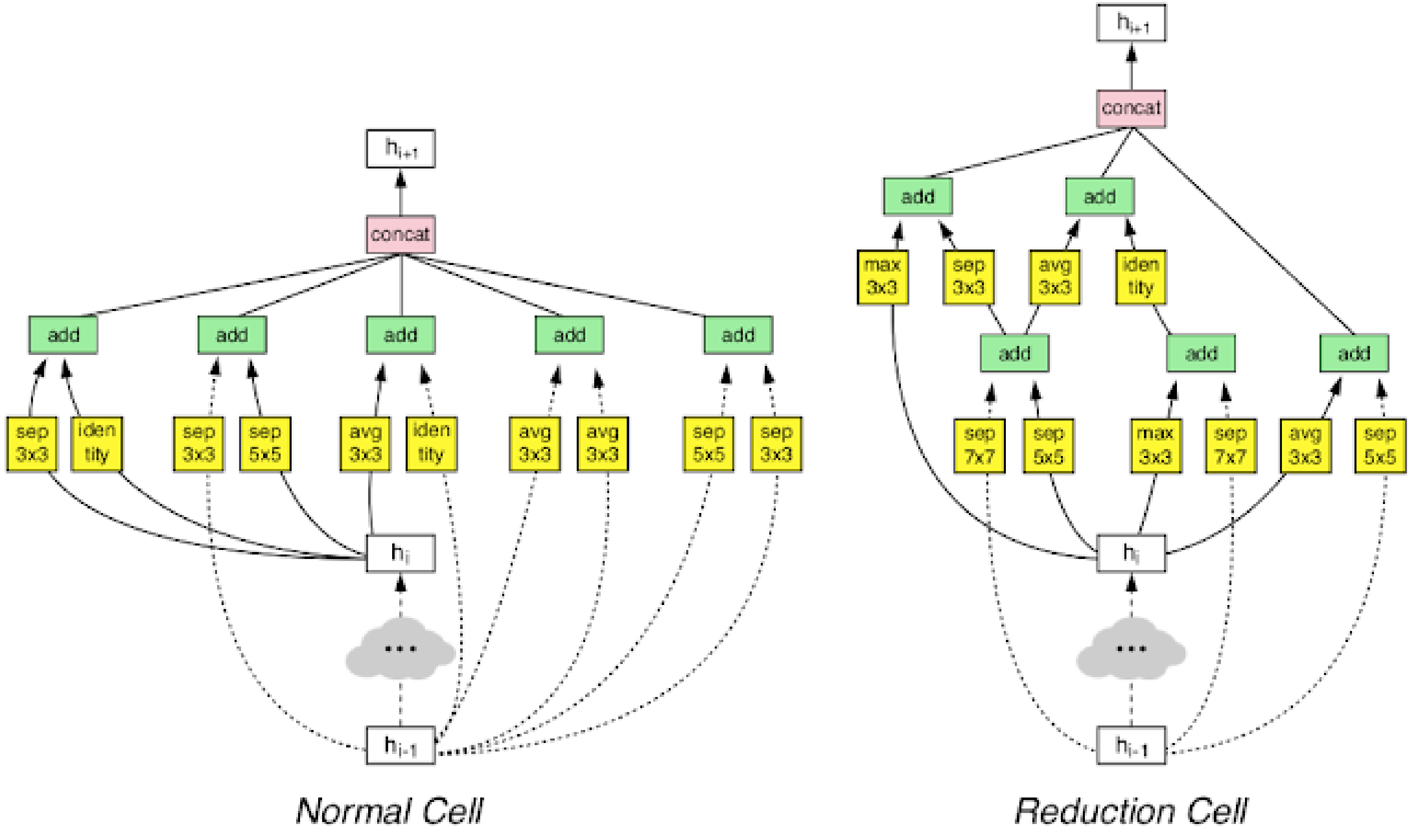

https://ai.googleblog.com/2017/11/automl-for-large-scale-image.html

AutoML on ImageNet

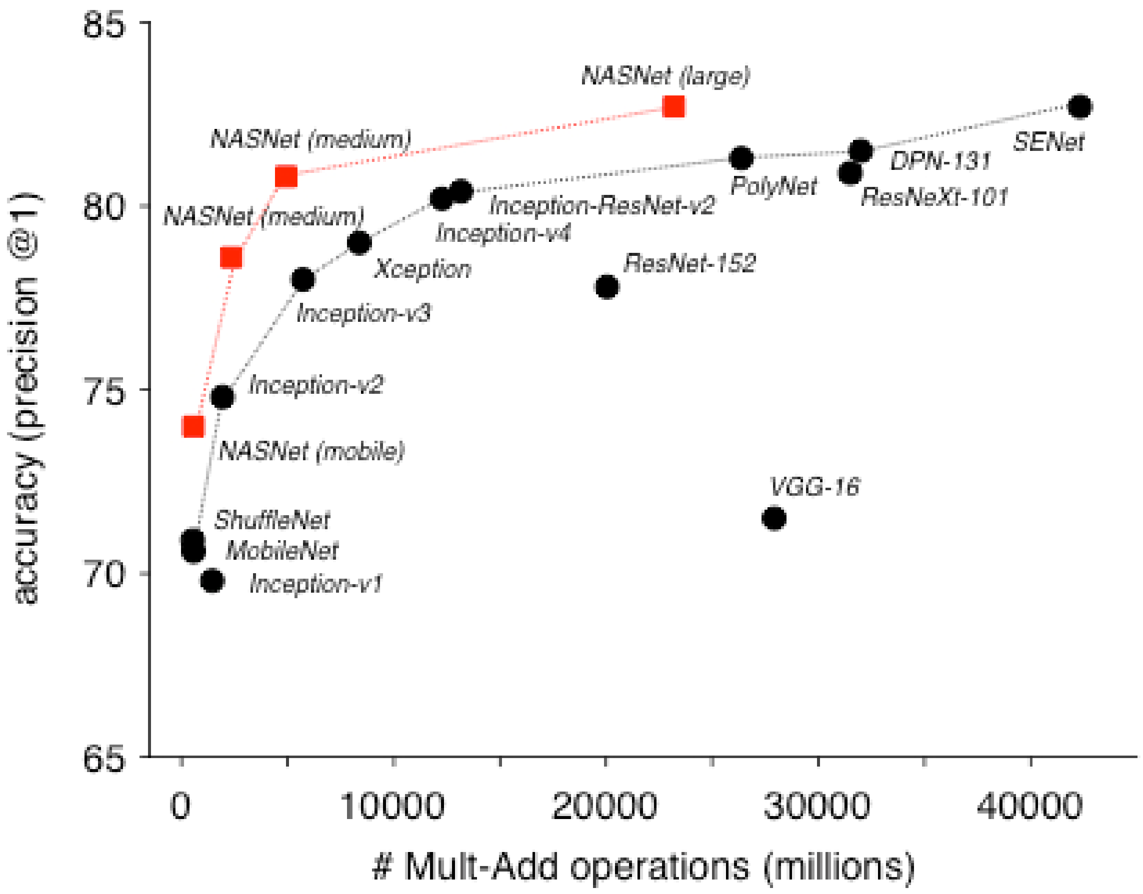

EfficientNet (May, 2019)

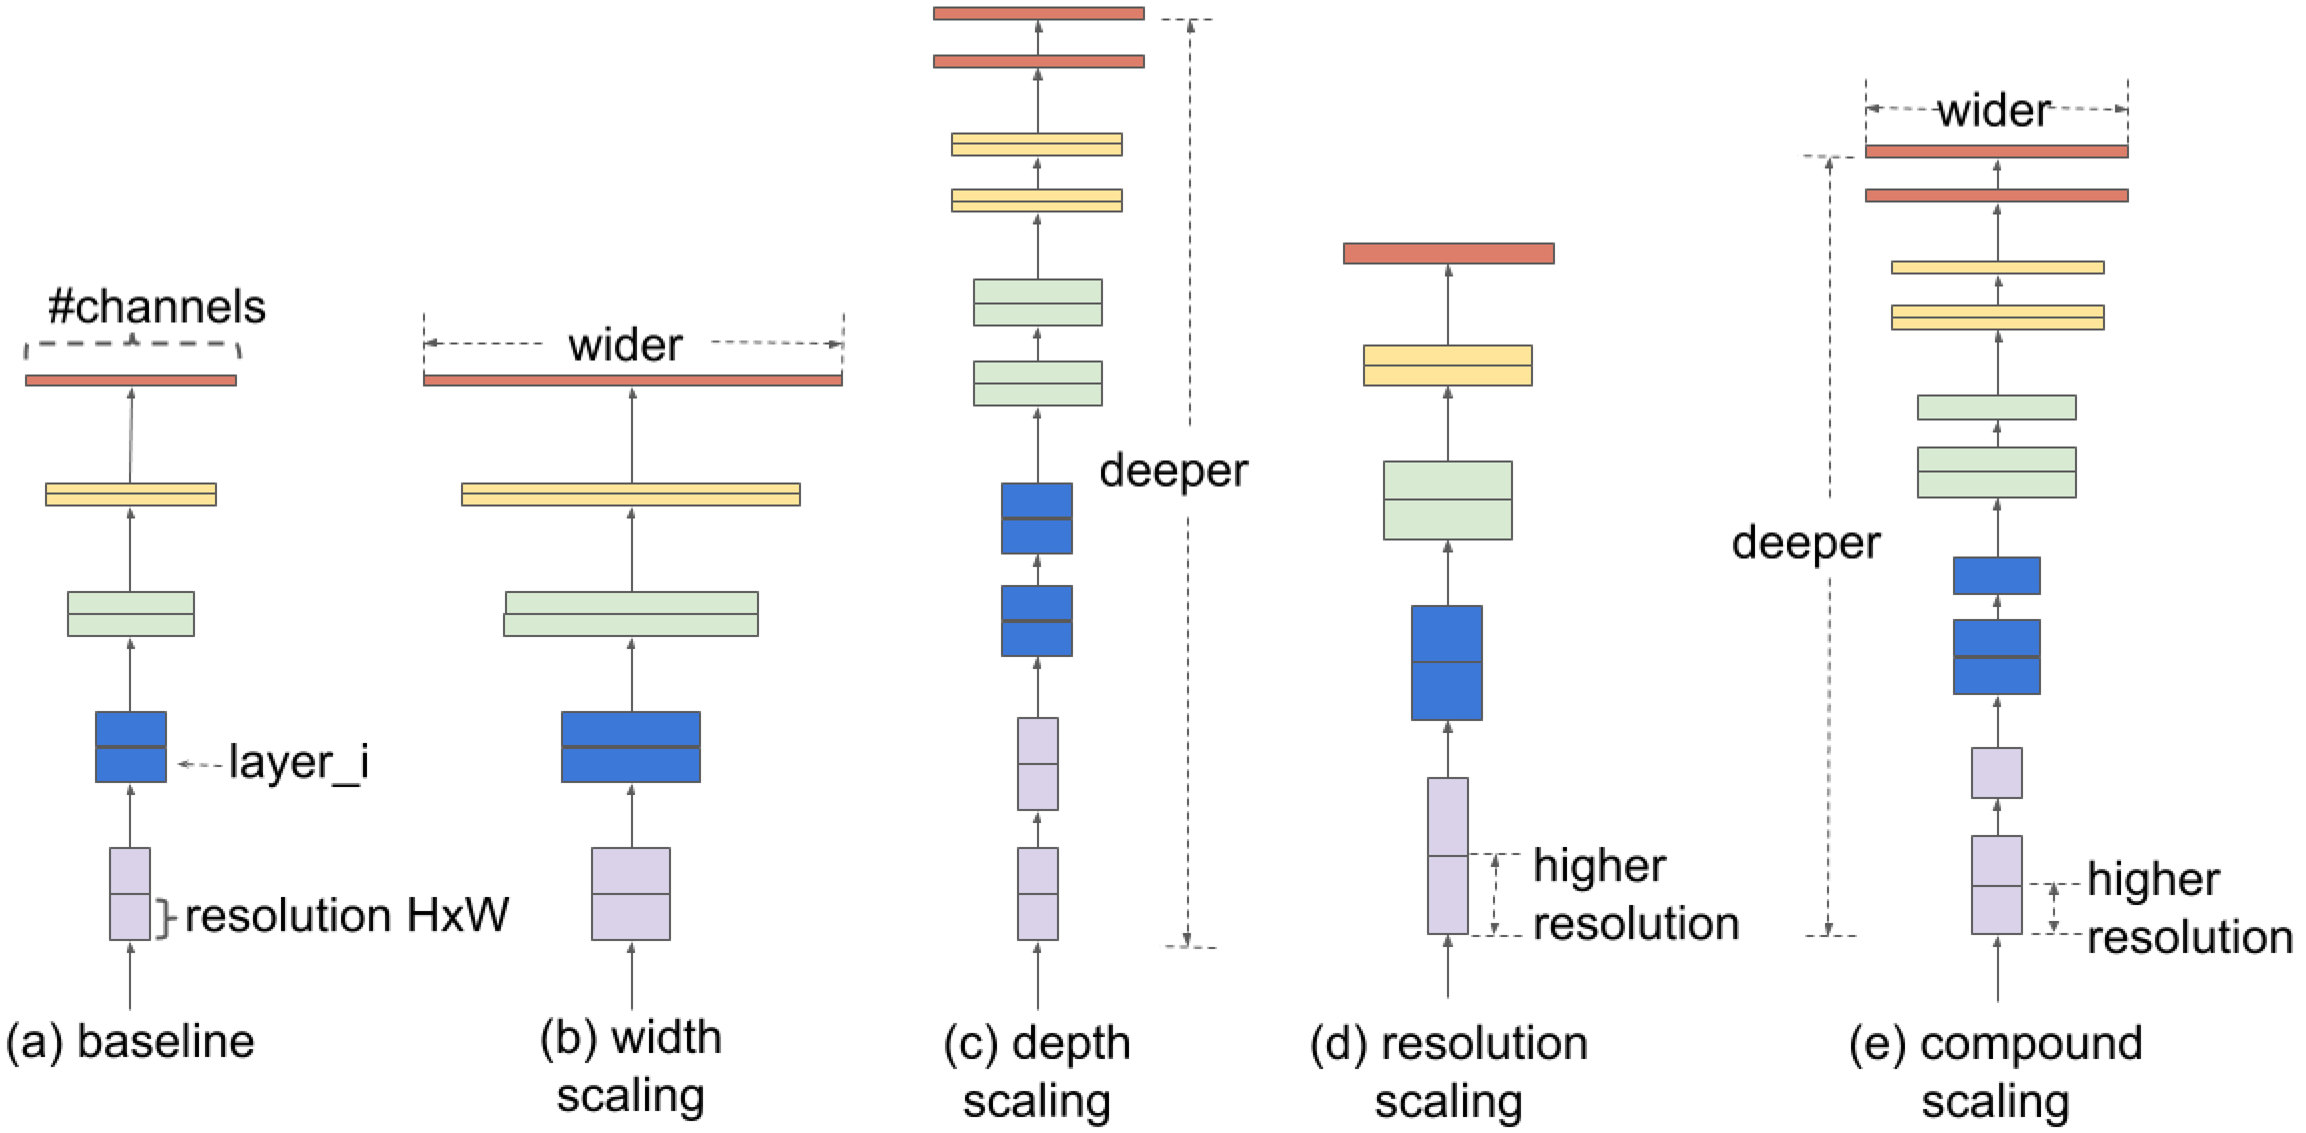

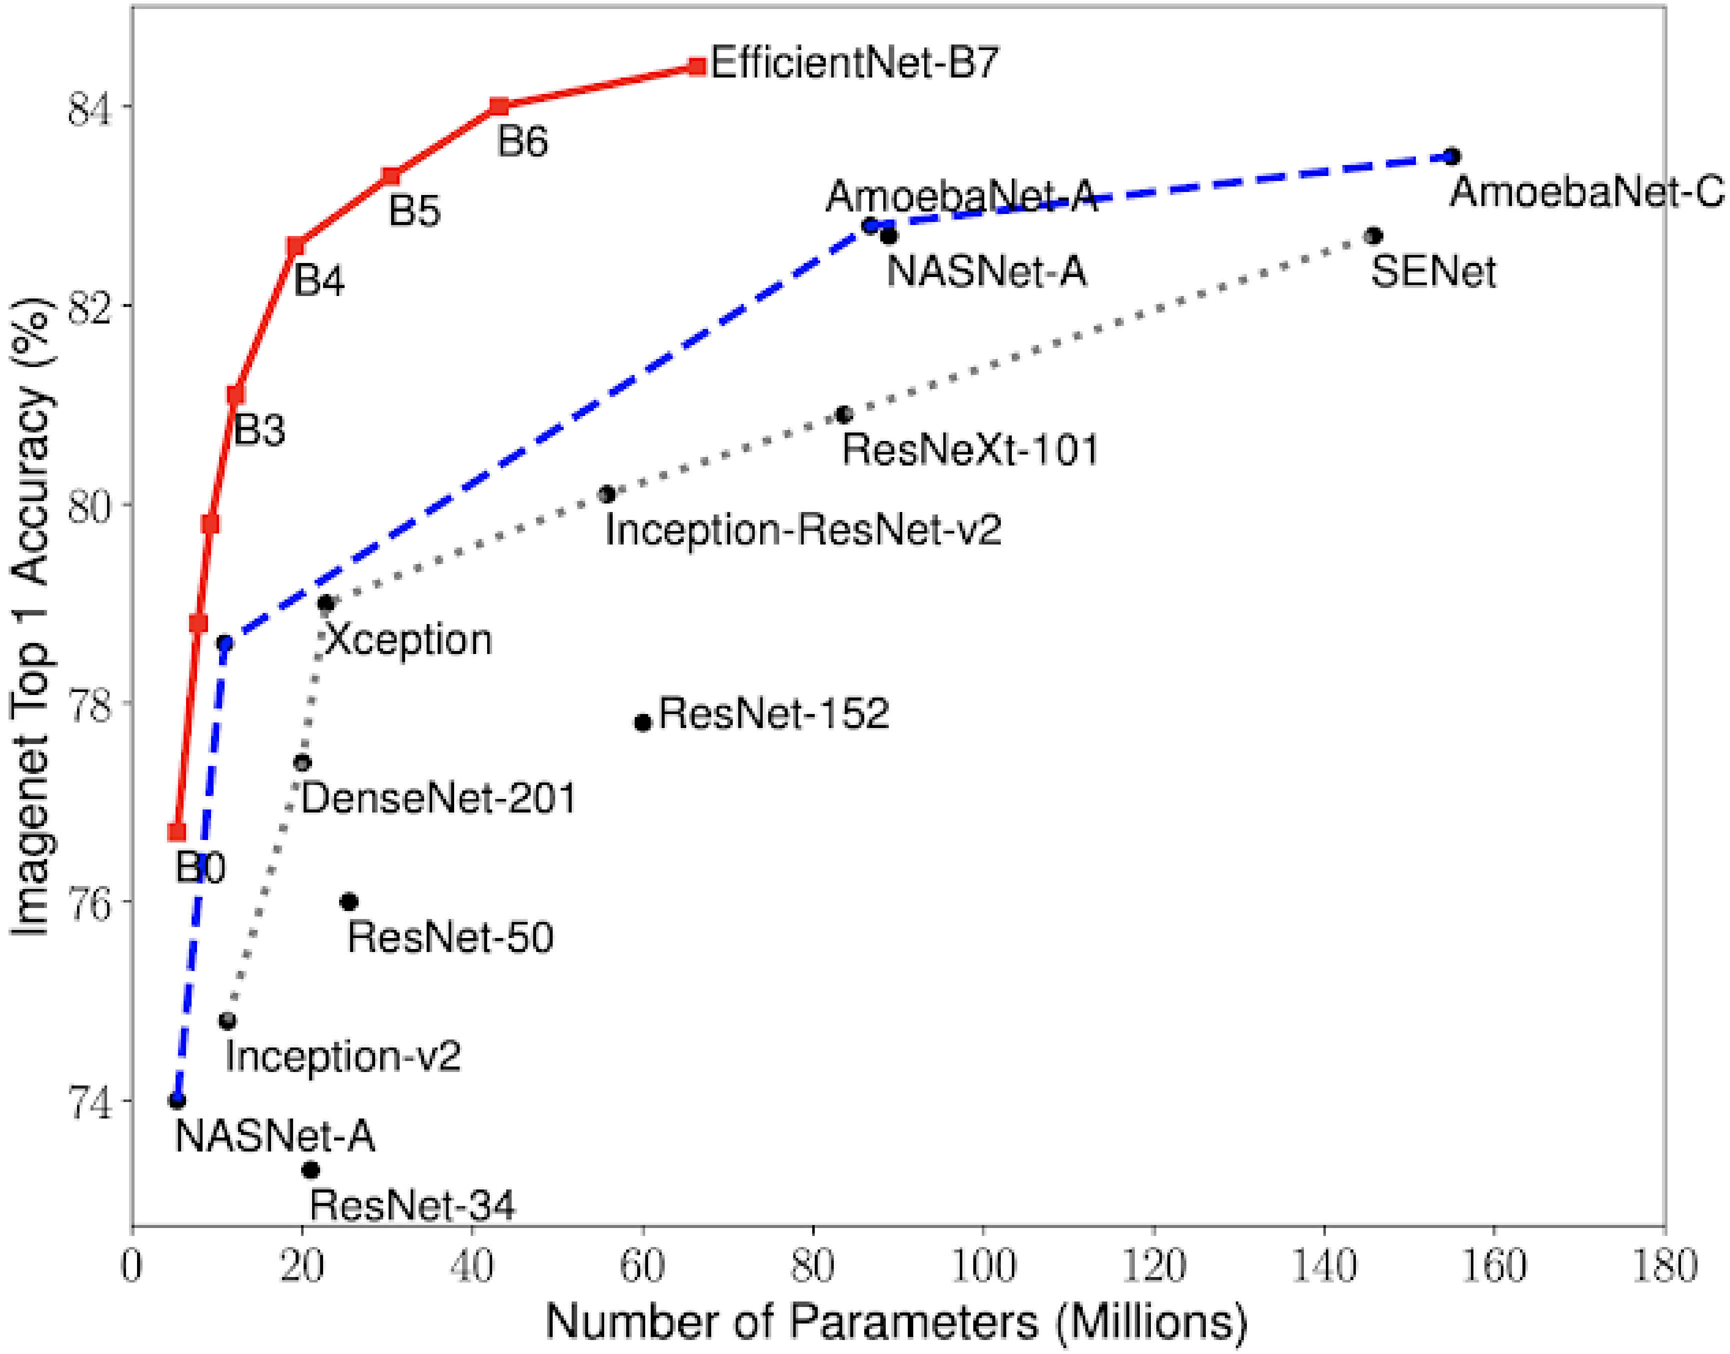
EfficientNet-B7
B6
B5
B4
B3
B0
AmoebaNet-A
AmoebaNet-C
NASNet-A
SENet
ResNeXt-101
Inception-ResNet-v2
Xception
ResNet-152
DenseNet-201
ResNet-50
Inception-v2
NASNet-A
ResNet-34
Imagenet Top 1 Accuracy (%)
74
76
78
80
82
84
0
20
40
60
80
100
120
140
160
180
Number of Parameters (Millions)

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