The *true* AI revolution

is in urgent need of (your) support

Eric Postma
eric.postma@gmail.com
CS&AI/Mind Labs, Tilburg University (csai.nl)
Jheronimus Academy of Data Science (jads.nl)
’s-Hertogenbosch

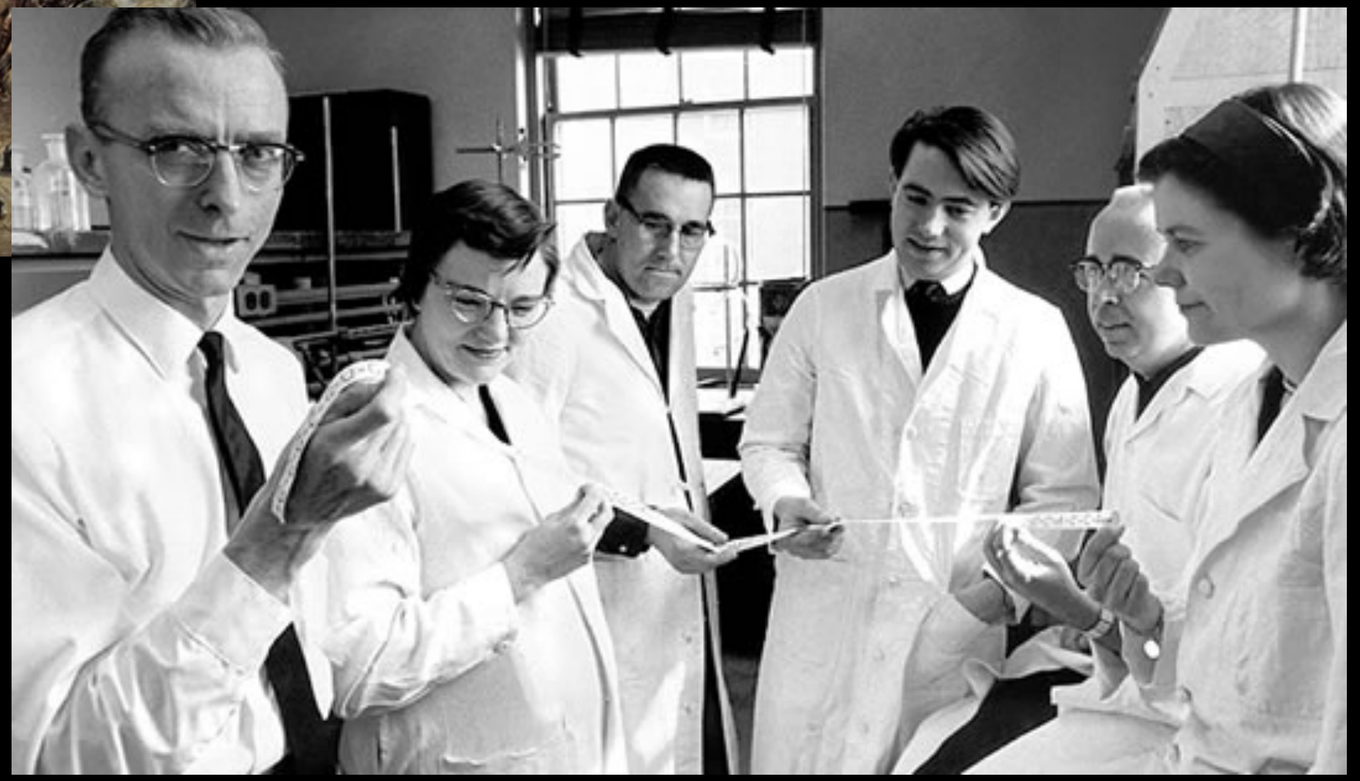

A brief history of Artificial Intelligence

$+$ means 'successor of' not with afterwards
$+$ also means $+$

i.e. $a+ = a+1$

N_0 is the class of natural numbers.

Peano axioms

$N_0 \in \text{Cls}$, $0 \in N_0$, $a \in N_0 . \supset . a+$

$s \in \text{Cls} . 0 \in s : x \in s . \supset_x x+ \in s : \supset . N_0 \supset$

Also (Cls) $a+ = b+ \supset a = b$.
$a \in N_0 \supset a+ - = 0$

Peano does something rather reminiscent of Church's d
Degrees by defining $a \cup 0 = a$
$a \cup 1 = a \cup$
$a \cup 2 = a \cup \cup$
— — —

Hence $a \cup (b \cup c) : (a \cup b) \cup c$

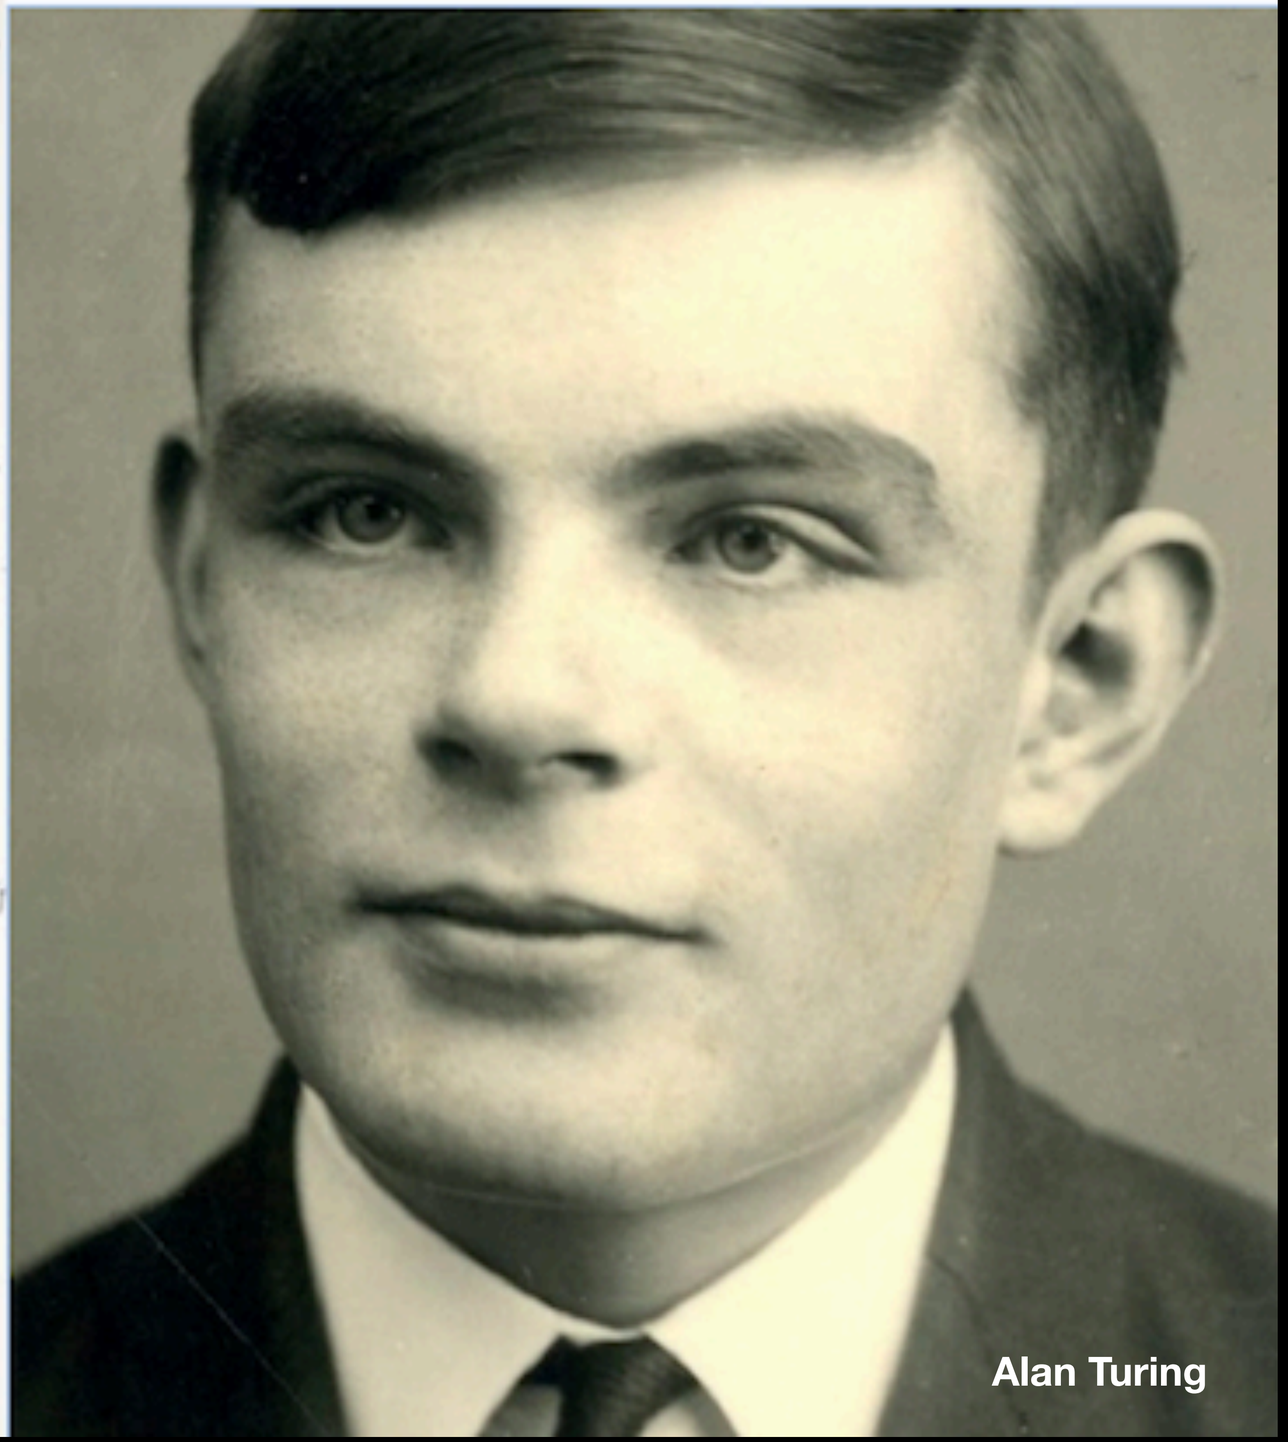

Alan Turing

Turing, A.M. (1950). Computing machinery and intelligence. Mind, 59, 433-460.

COMPUTING MACHINERY AND INTELLIGENCE

By A. M. Turing

1. The Imitation Game

I propose to consider the question, "Can machines think?" This should begin with definitions of the meaning of the terms "machine" and "think." The definitions might be framed so as to reflect so far as possible the normal use of the words, but this attitude is dangerous, If the meaning of the words "machine" and "think" are to be found by examining how they are commonly used it is difficult to escape the conclusion that the meaning and the answer to the question, "Can machines think?" is to be sought in a statistical survey such as a Gallup poll. But this is absurd. Instead of

Symbolic AI (1957-2000)

- Emphasis on formal logic and symbolic representations
- Neuroscience is irrelevant
- Focus on high-level cognition (reasoning, problem-solving, …)

Deep Blue (1997)

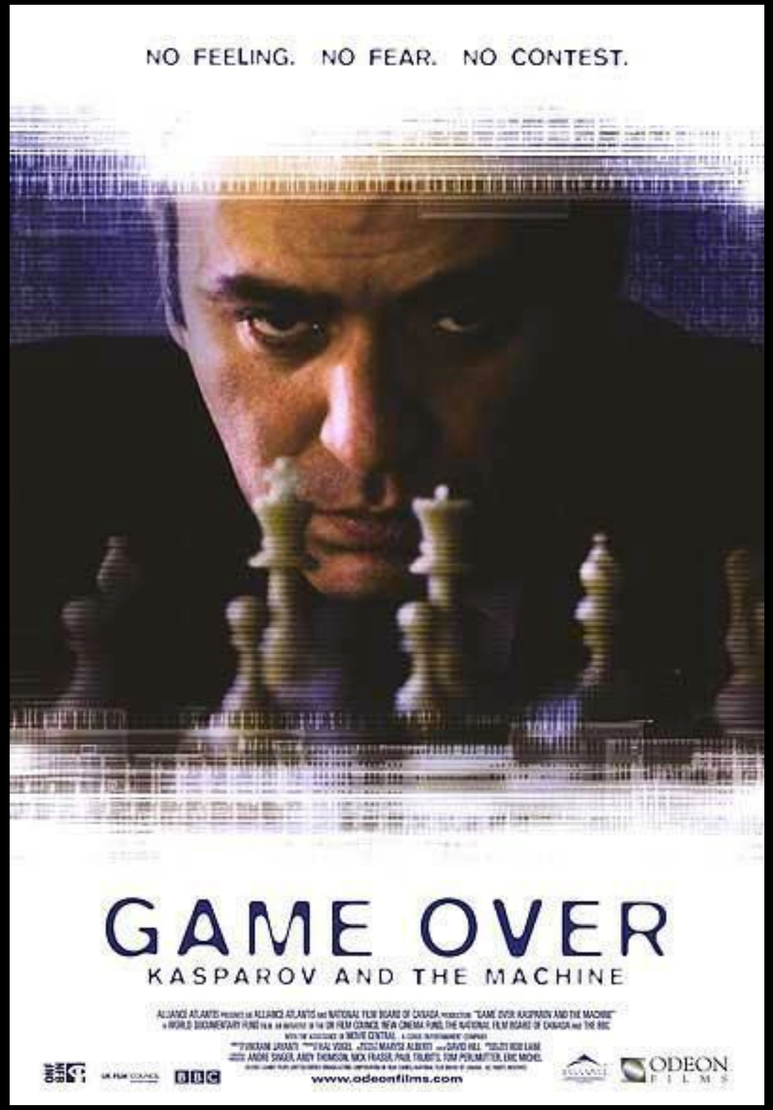

The “new AI”

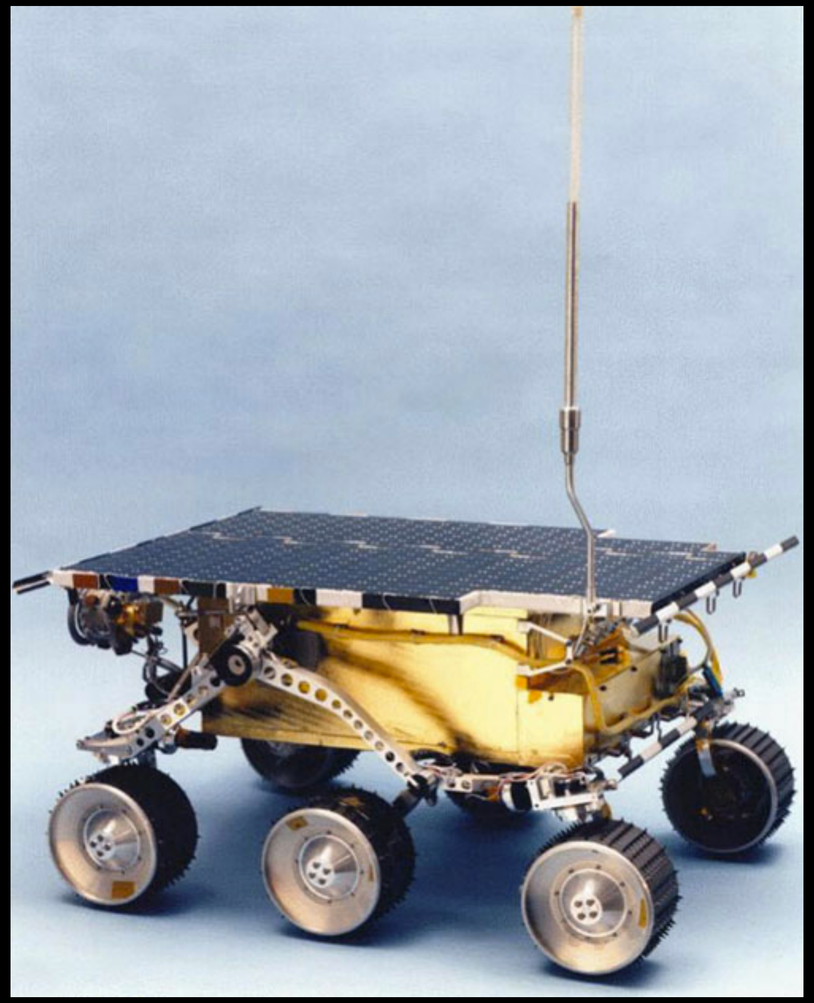

The Sojourner Rover (1997)

“Sub-symbolic” AI (1988-2016)

- Emphasis on machine learning (neural networks)
- Neuroscience is relevant
- Emphasis on low-level cognition (e.g., perception)

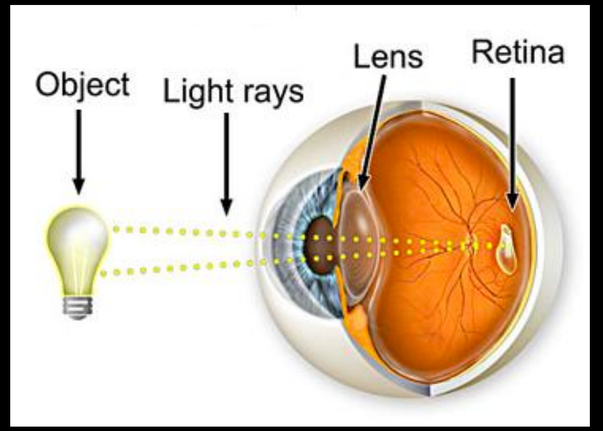

Recognising Patterns

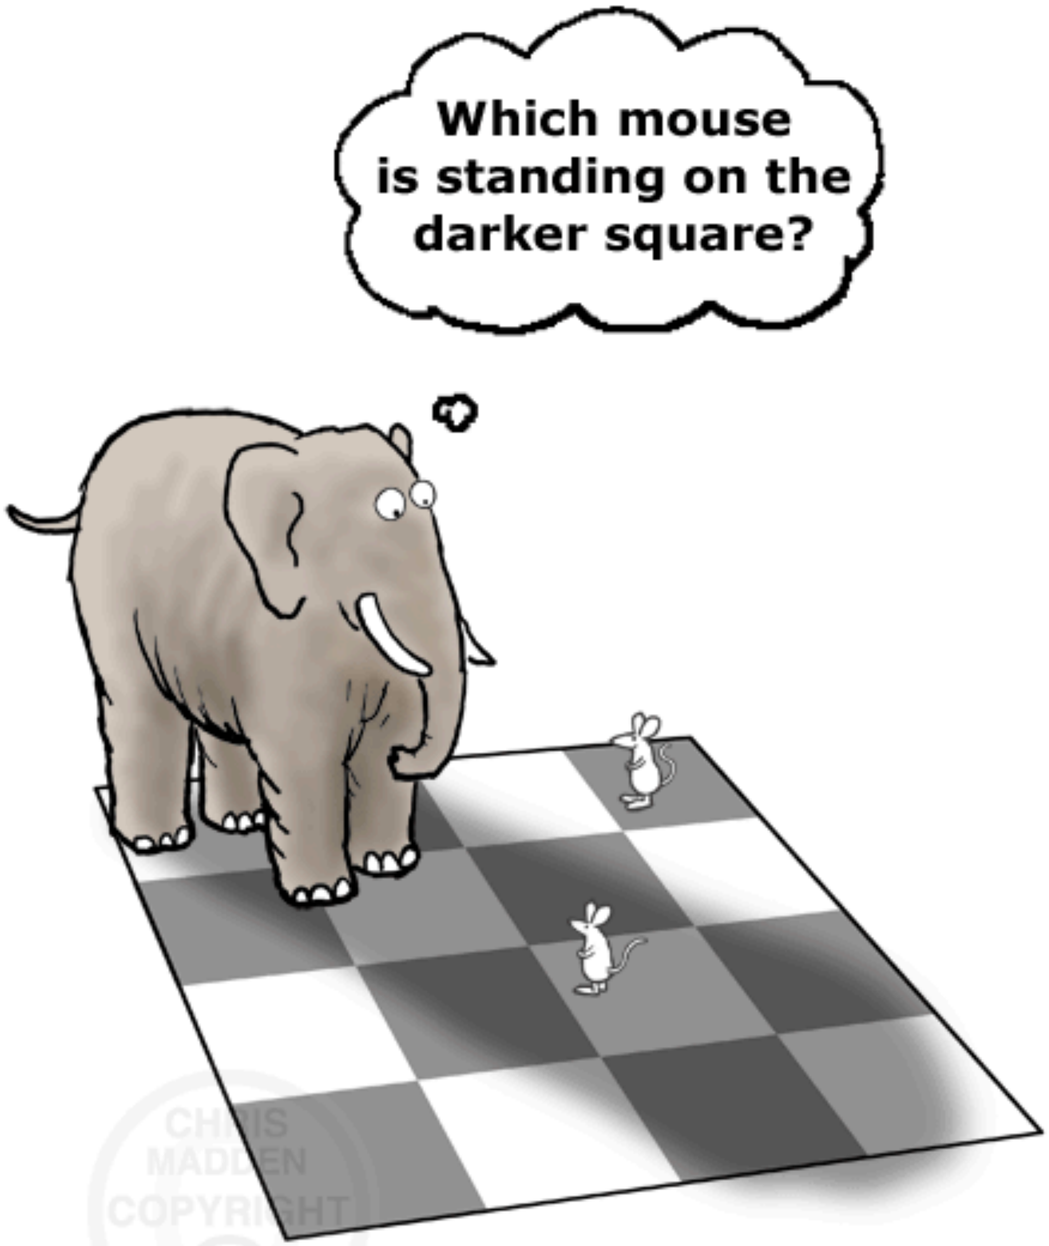
Which mouse is standing on the darker square?
CHRIS MADDEN COPYRIGHT ©

How to recognise a Samojed?

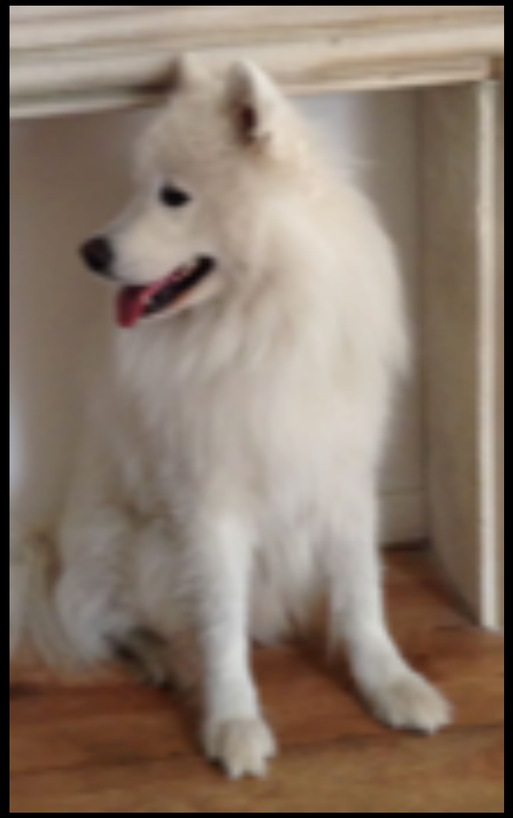

A very brief history of neural networks

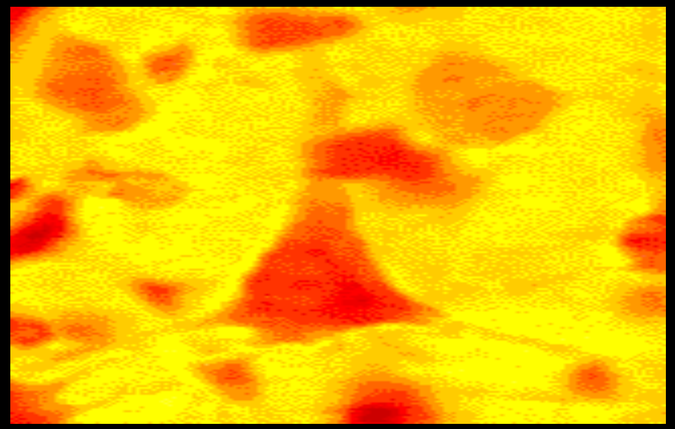

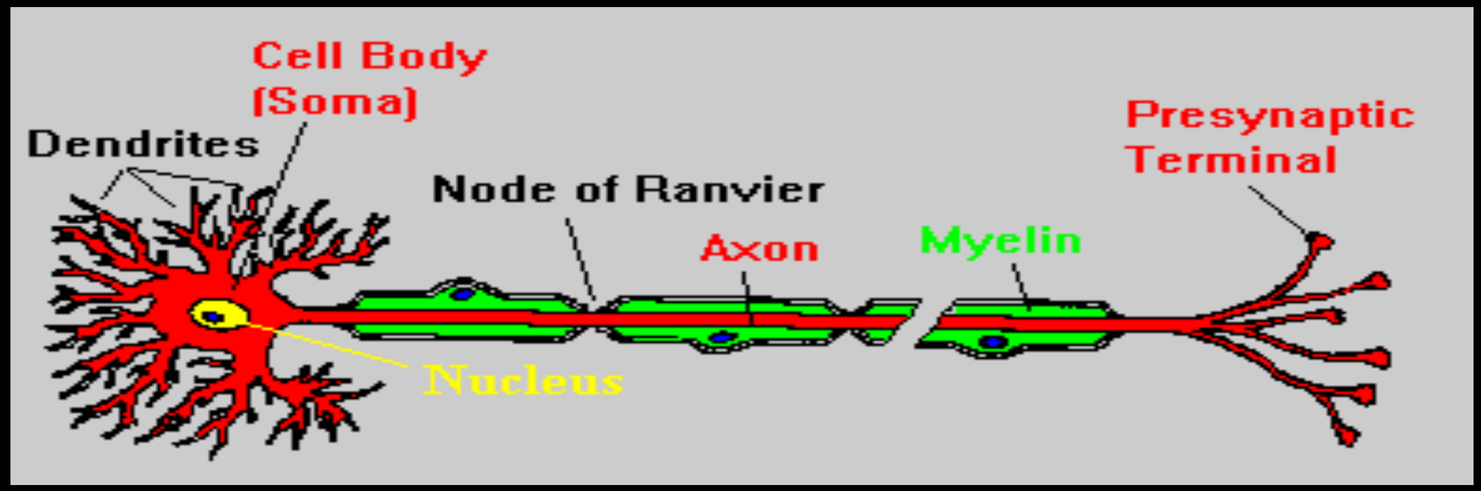

McCulloch-Pitts Neurons (1947)
Perceptron (1957)
Multilayer Perceptron (1988)

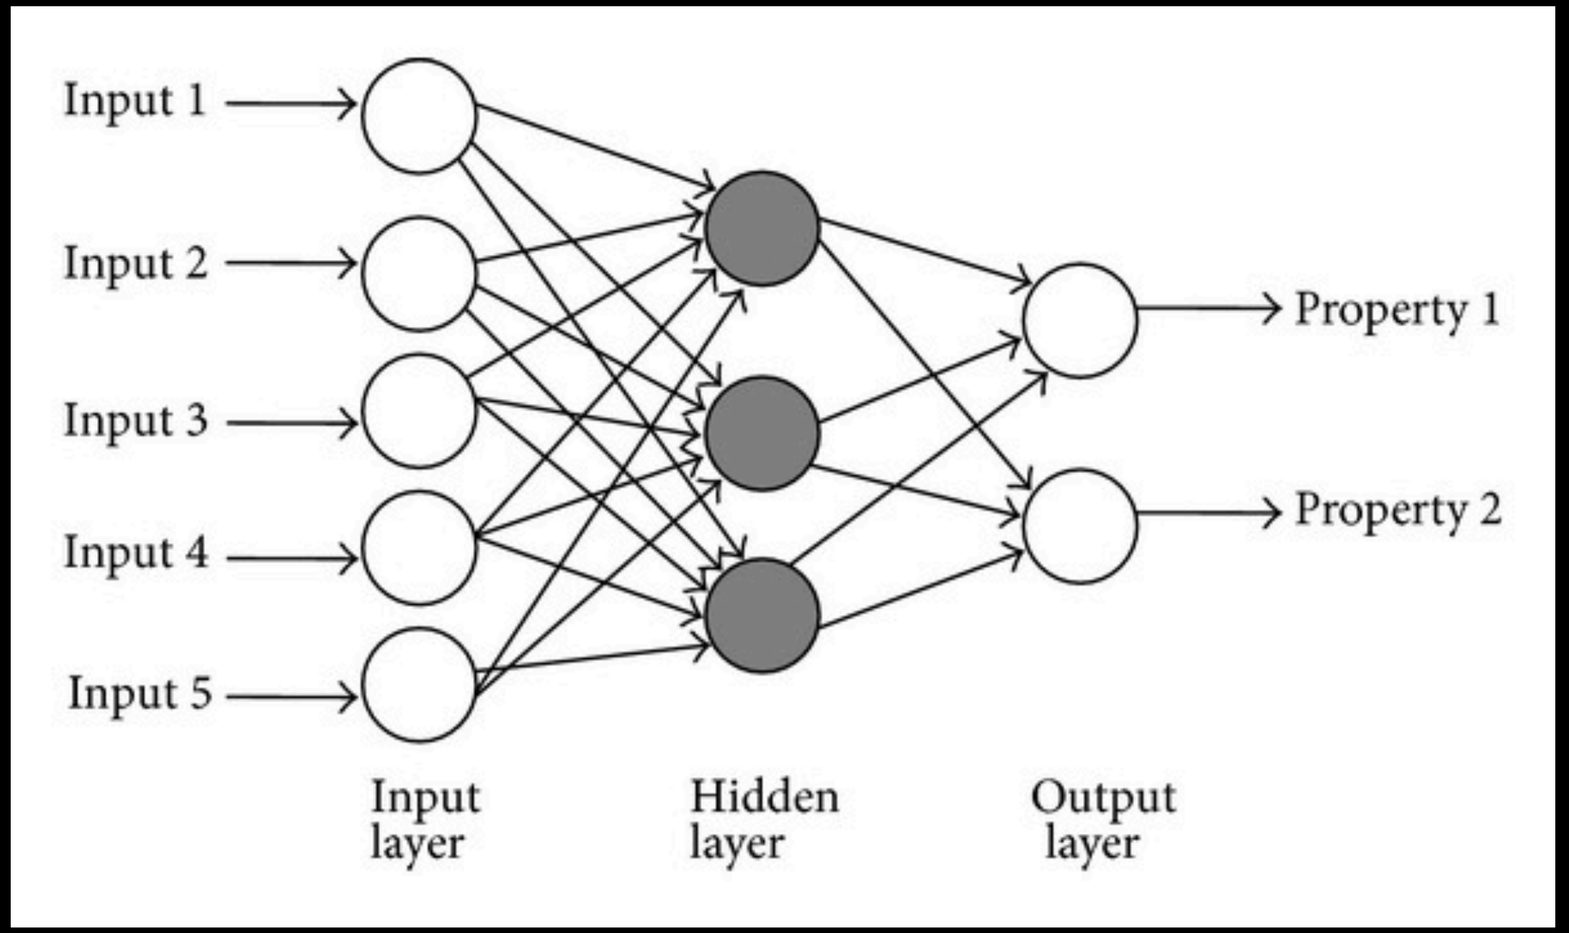

Video courtesy of 3Blue1Brown

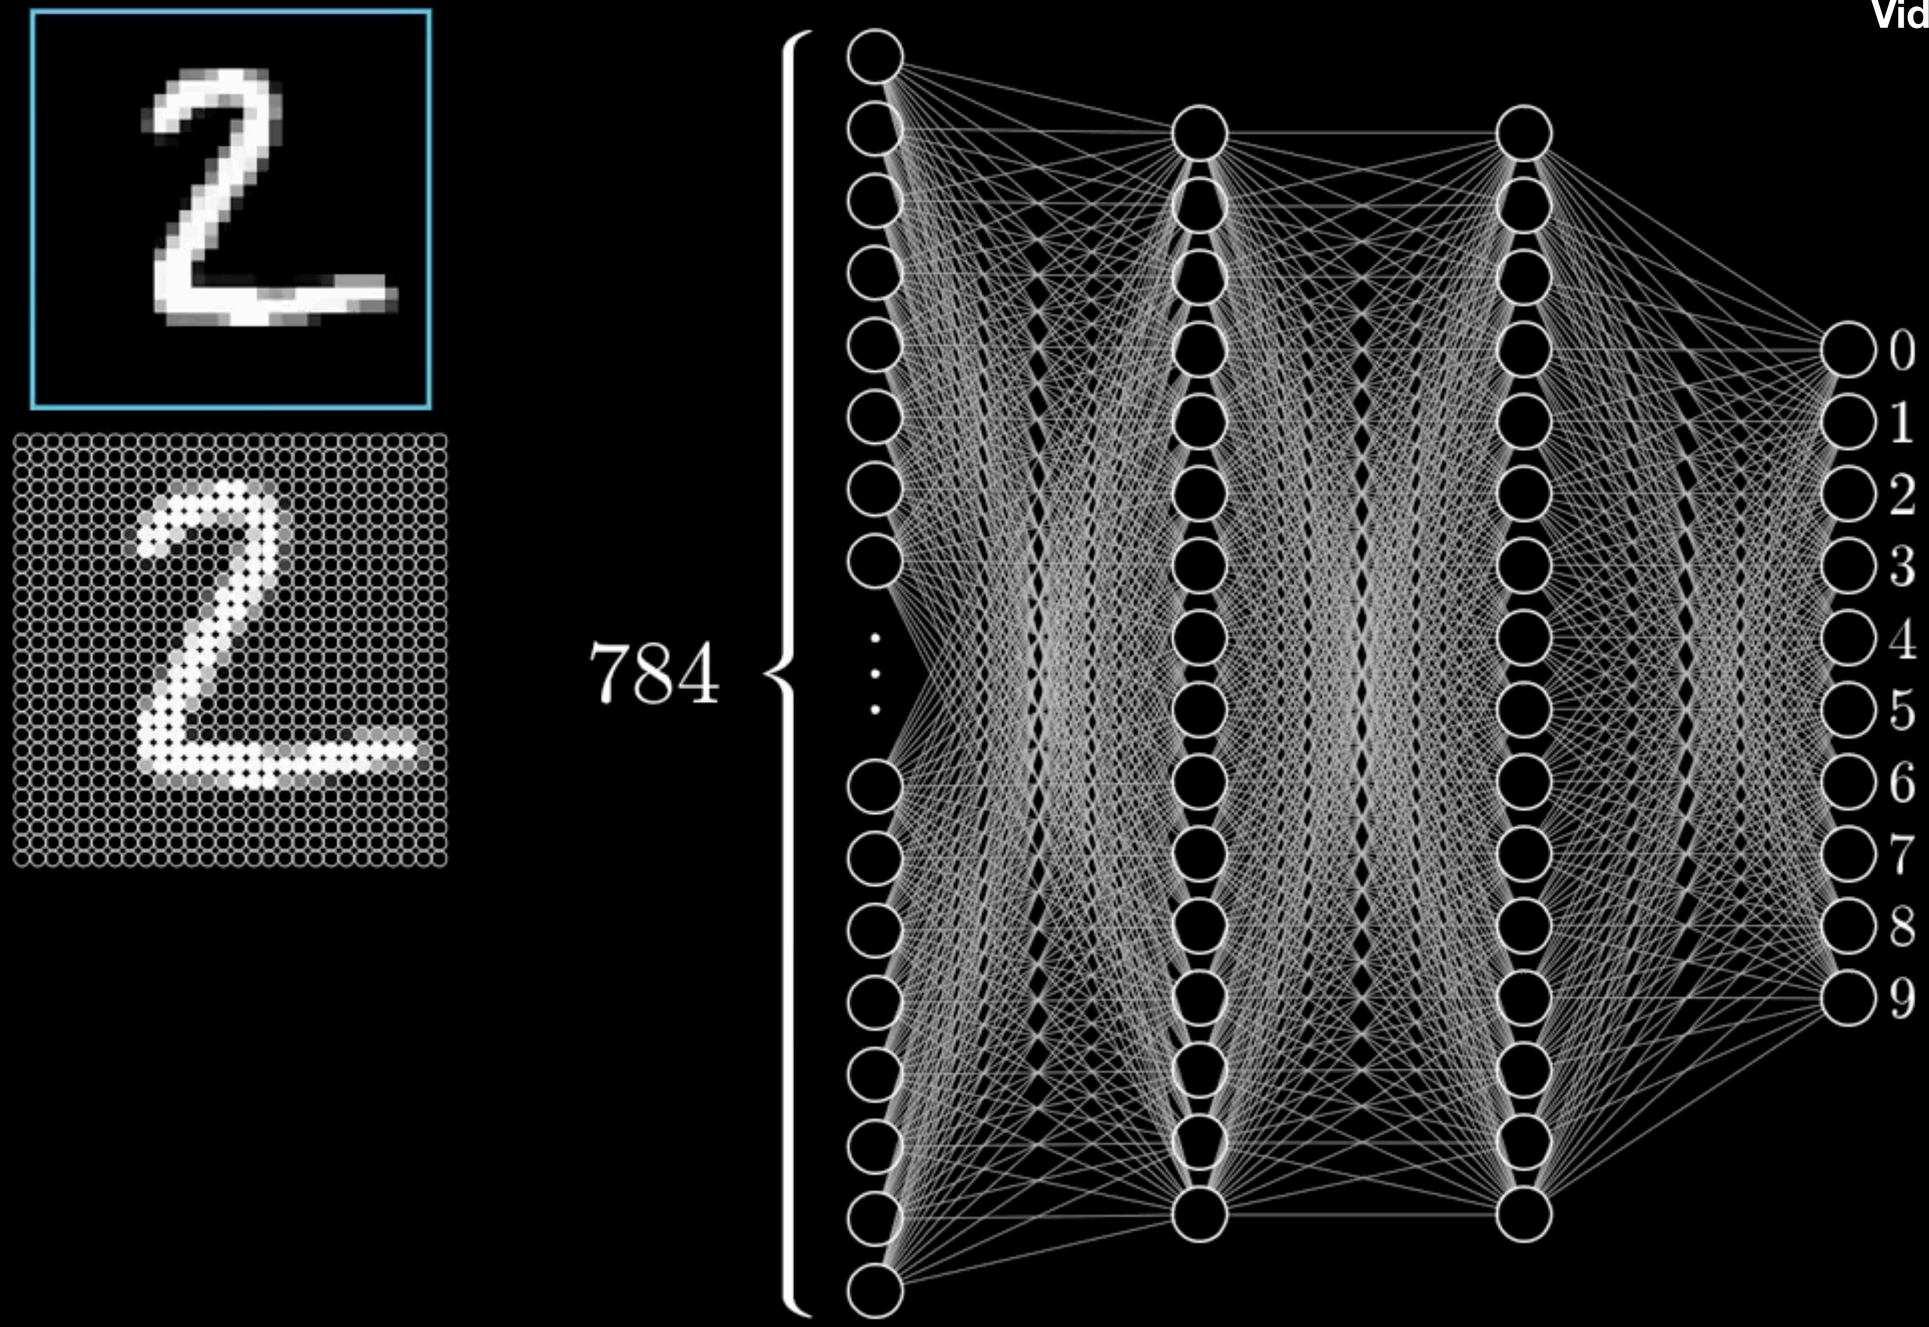

Video courtesy of 3Blue1Brown

Training in progress...

9

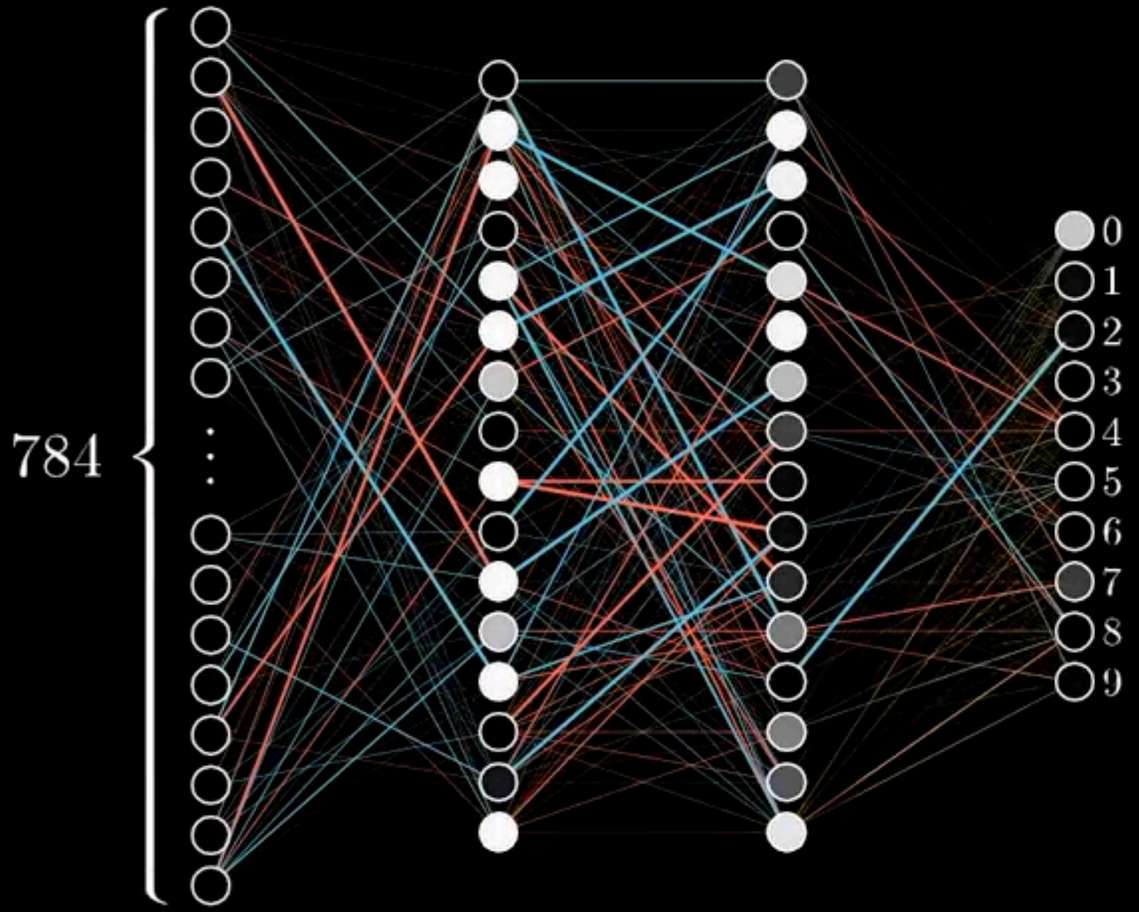

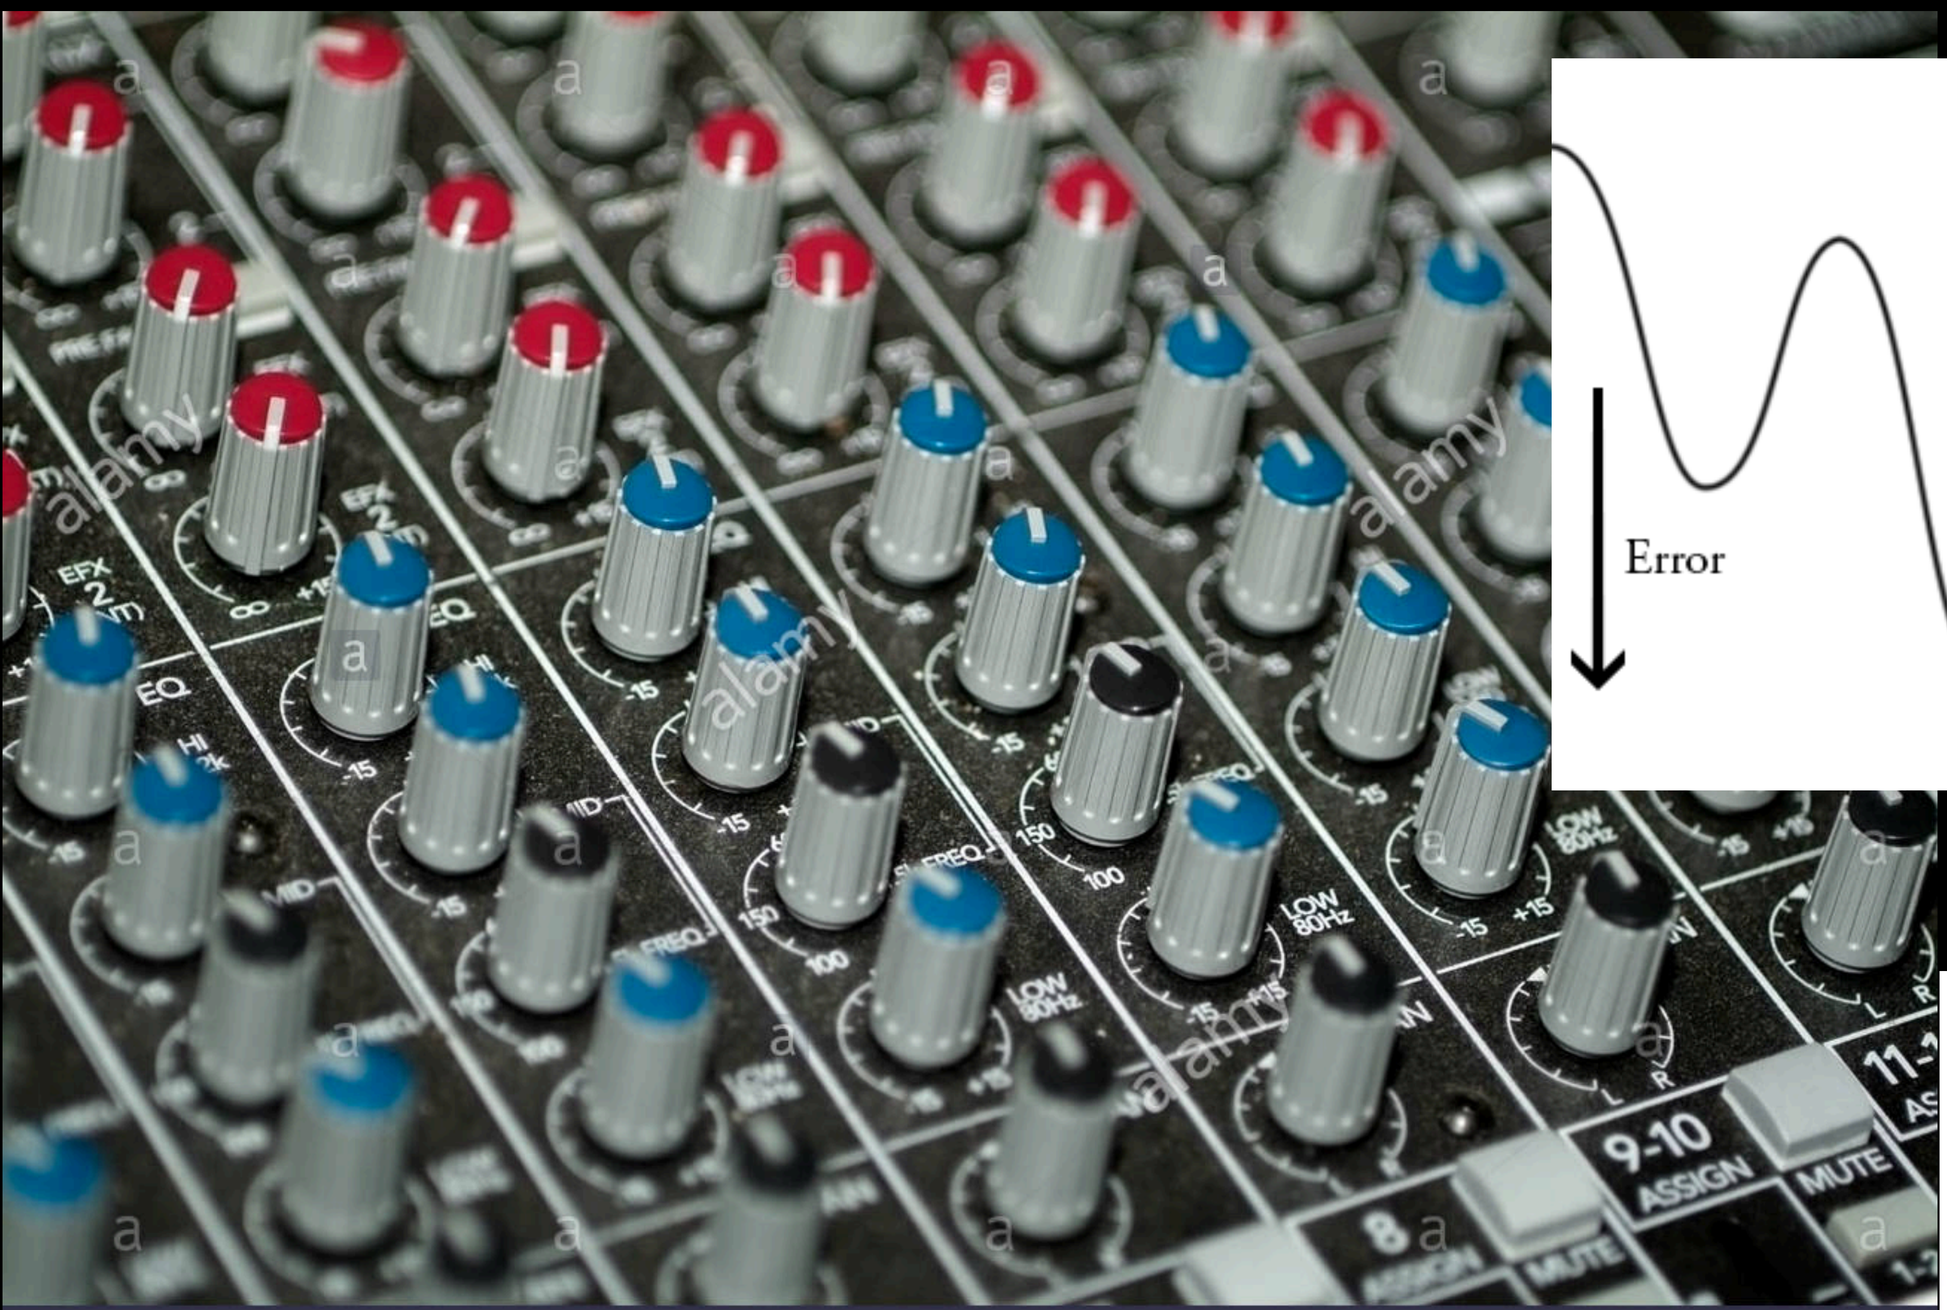

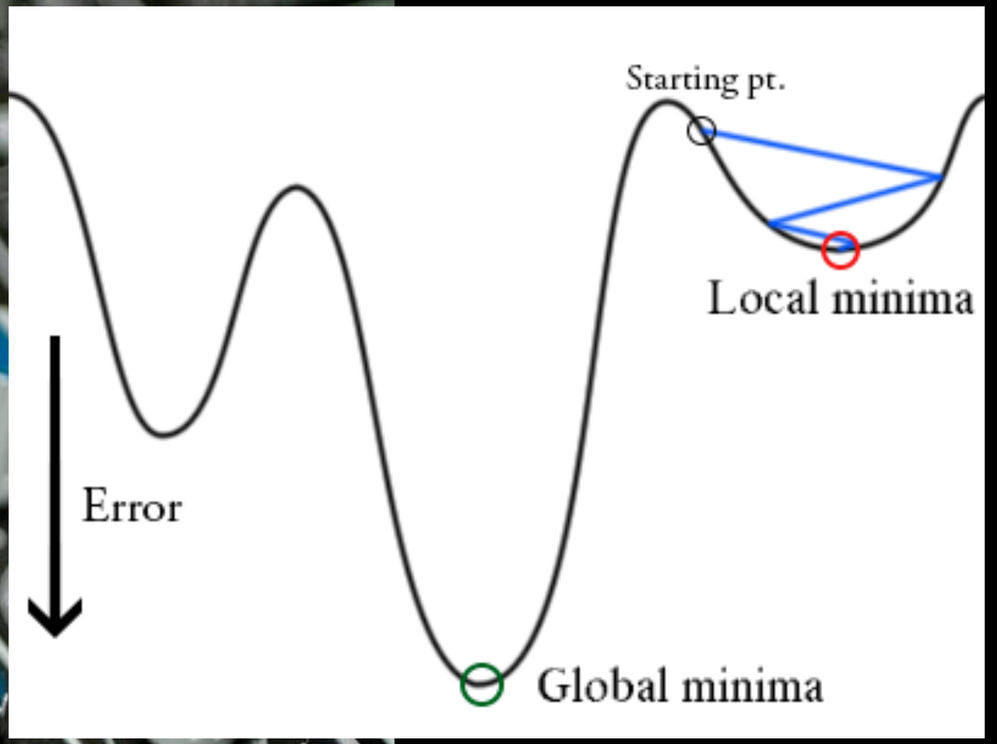

Gradient Descent

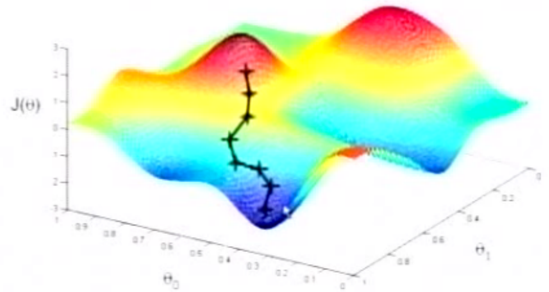

Deep Learning (2012)

Convolutional Neural Network

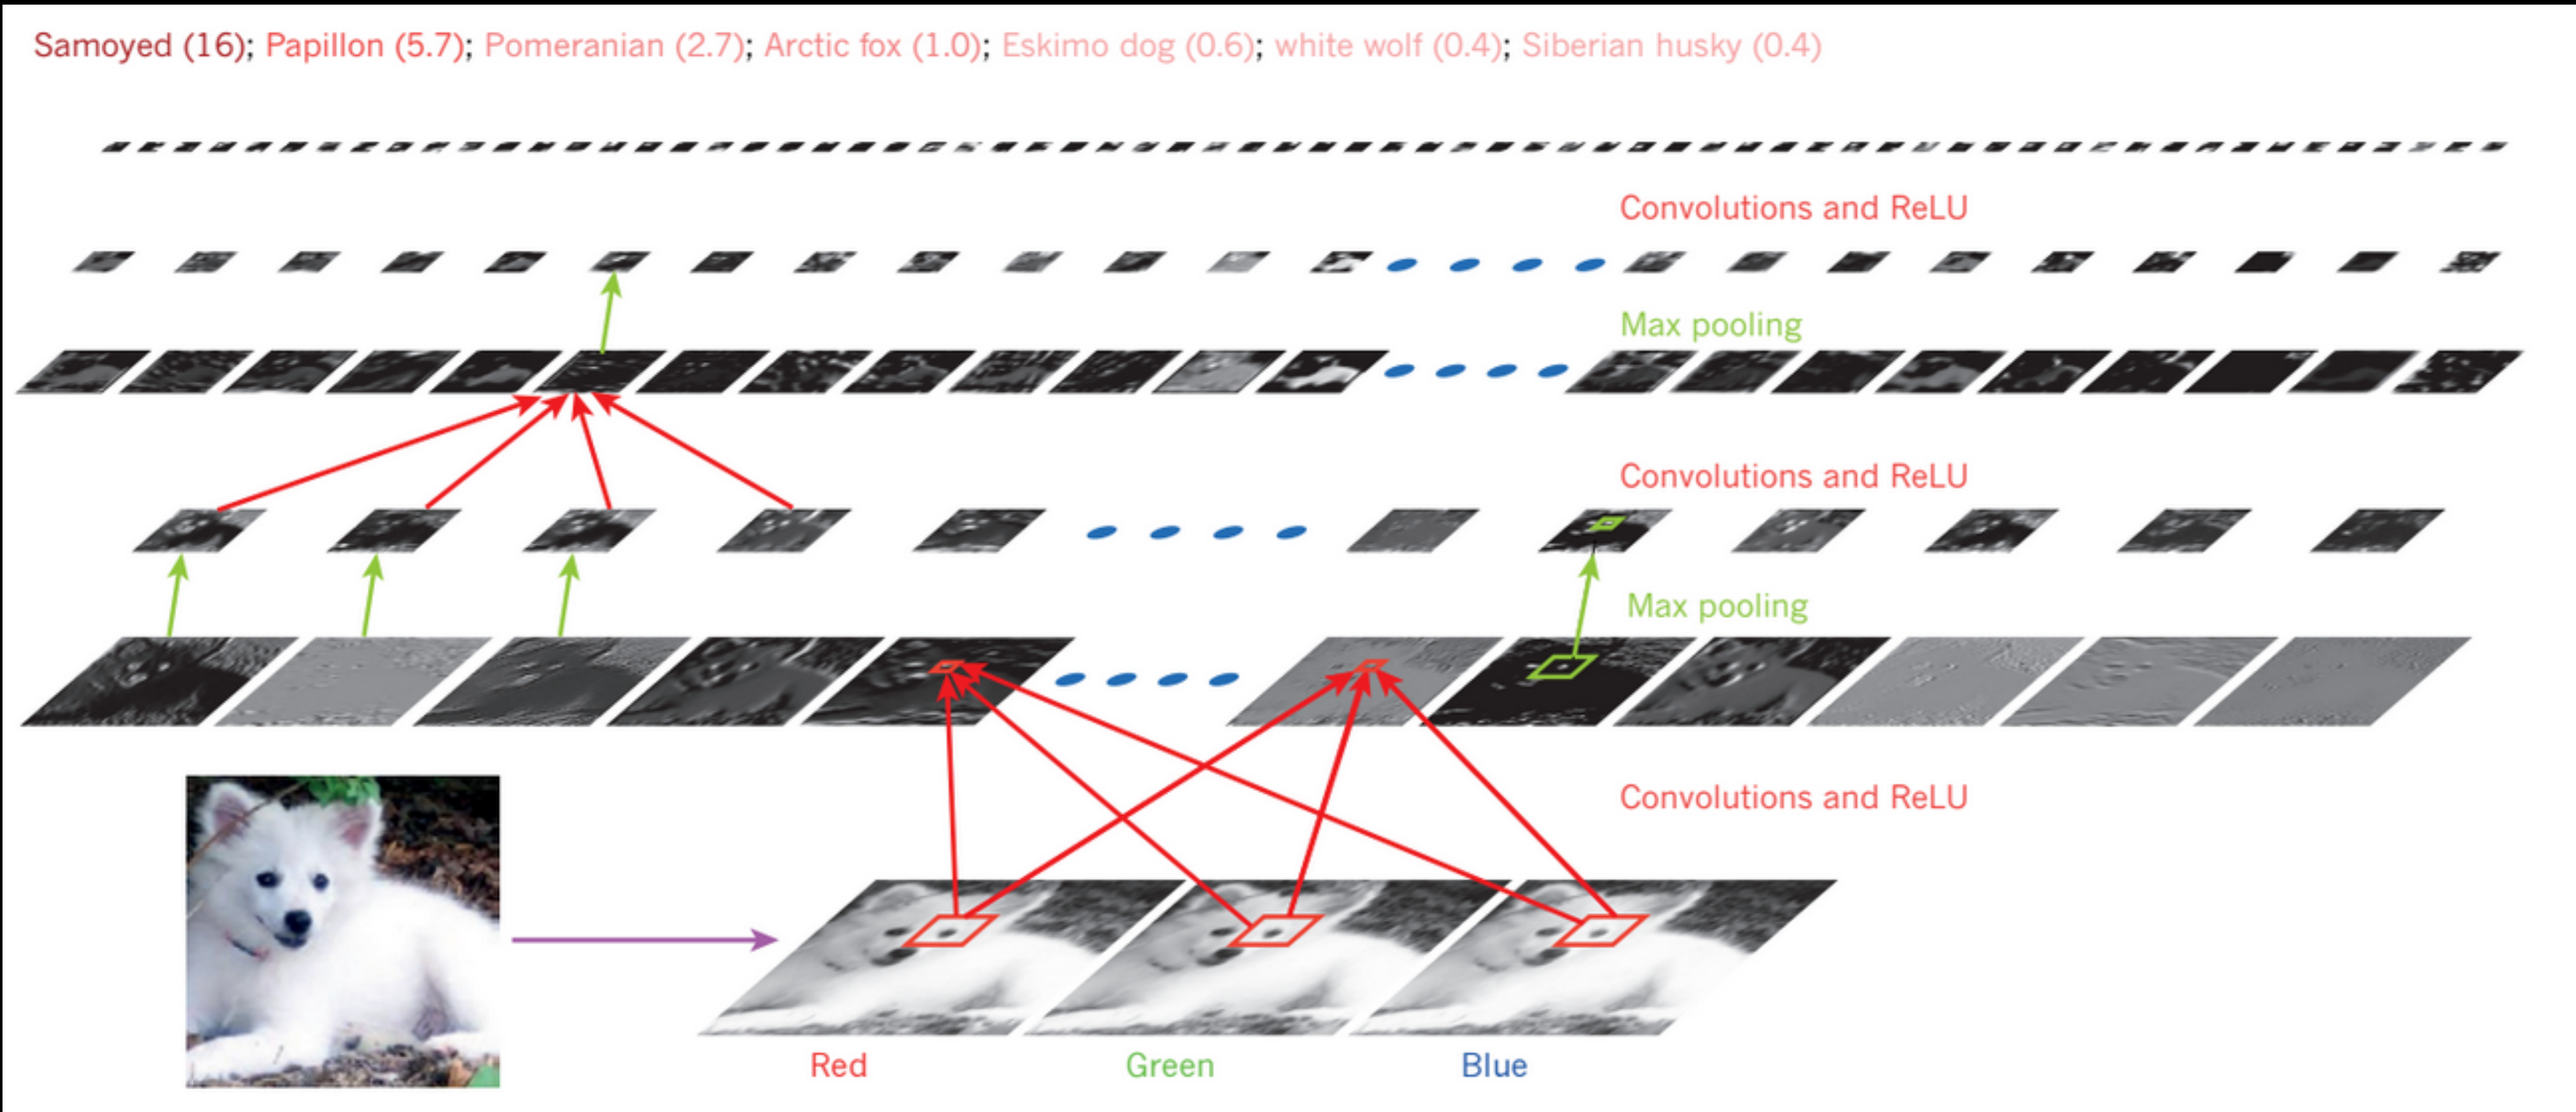

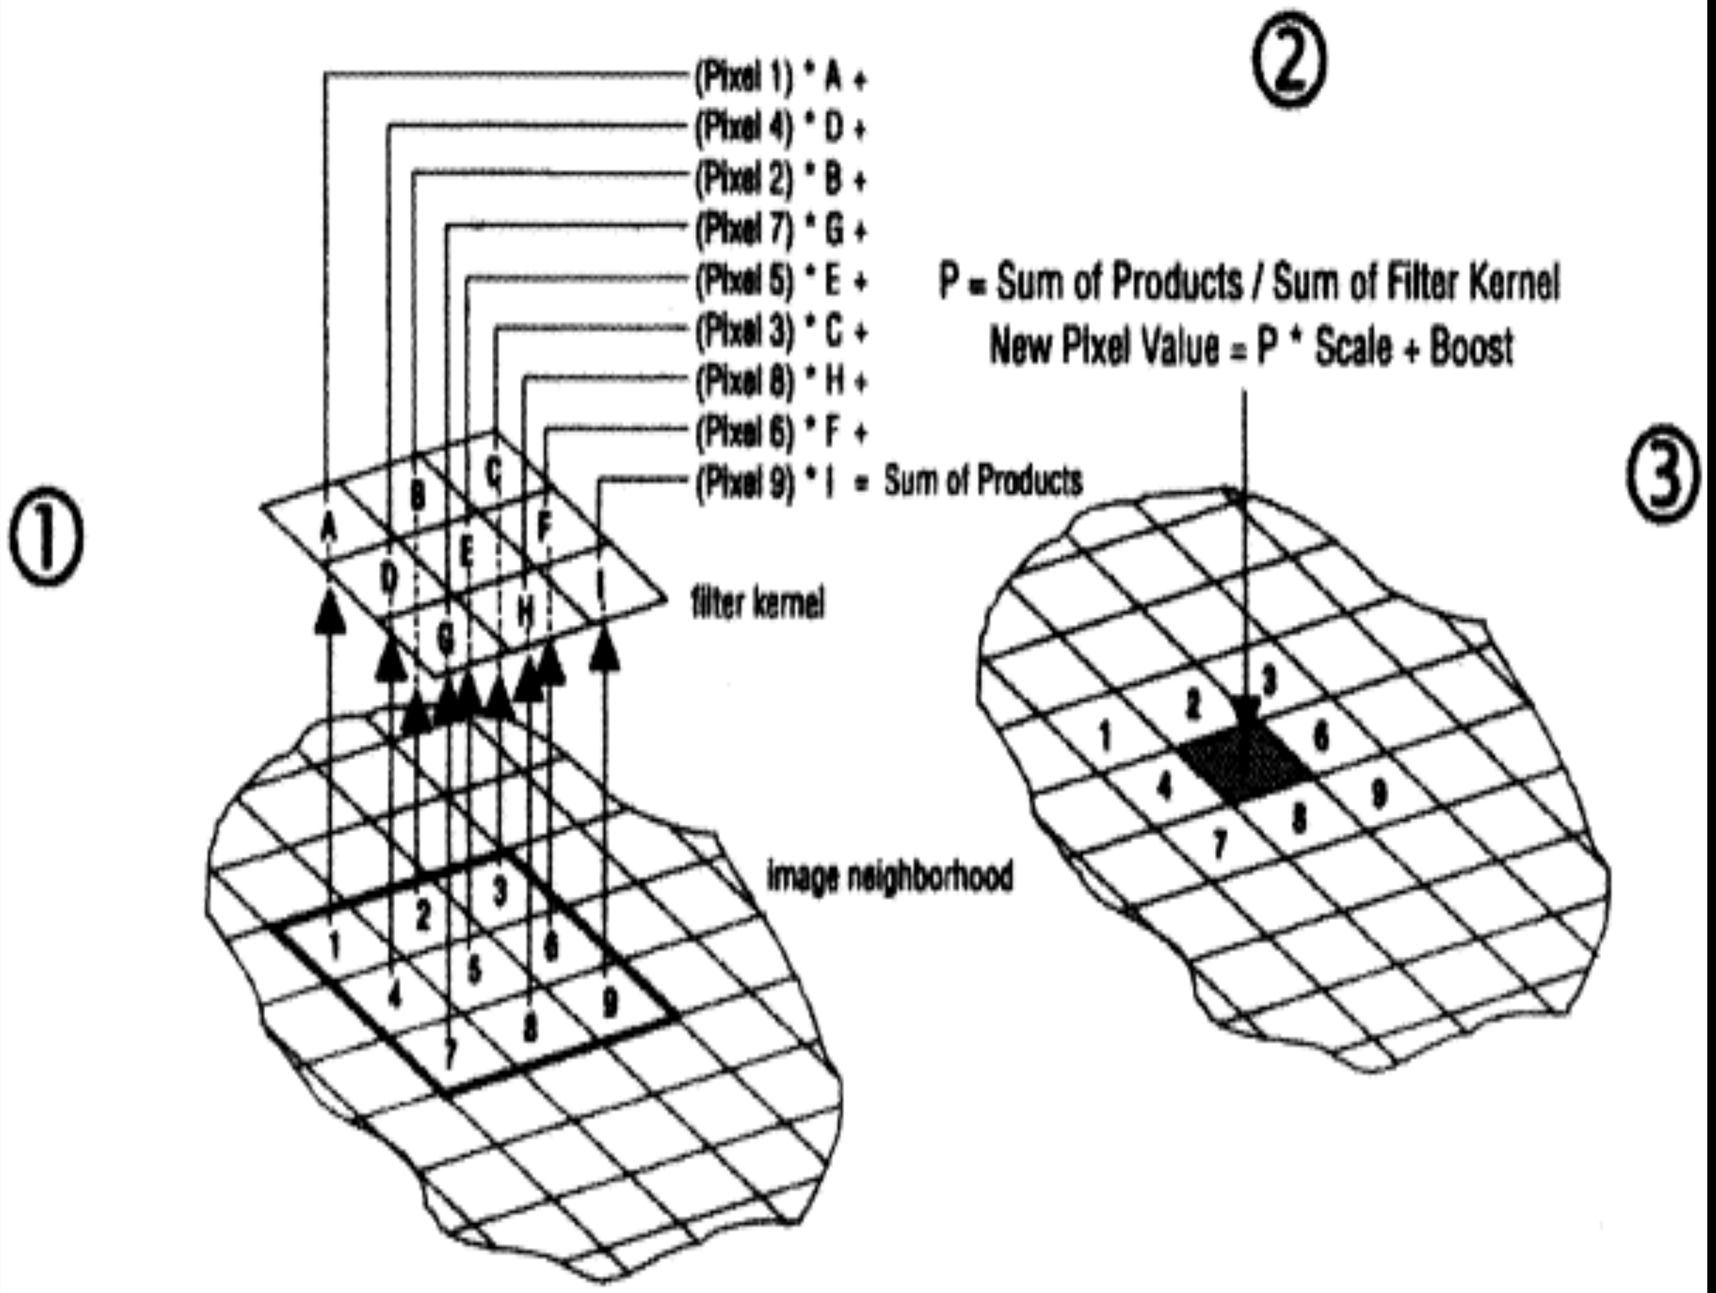

The Convolution Filtering Process

How does it work?

convolution demo
(source: https://ujjwalkarn.me/2016/08/11/intuitive-explanation-convnets/)

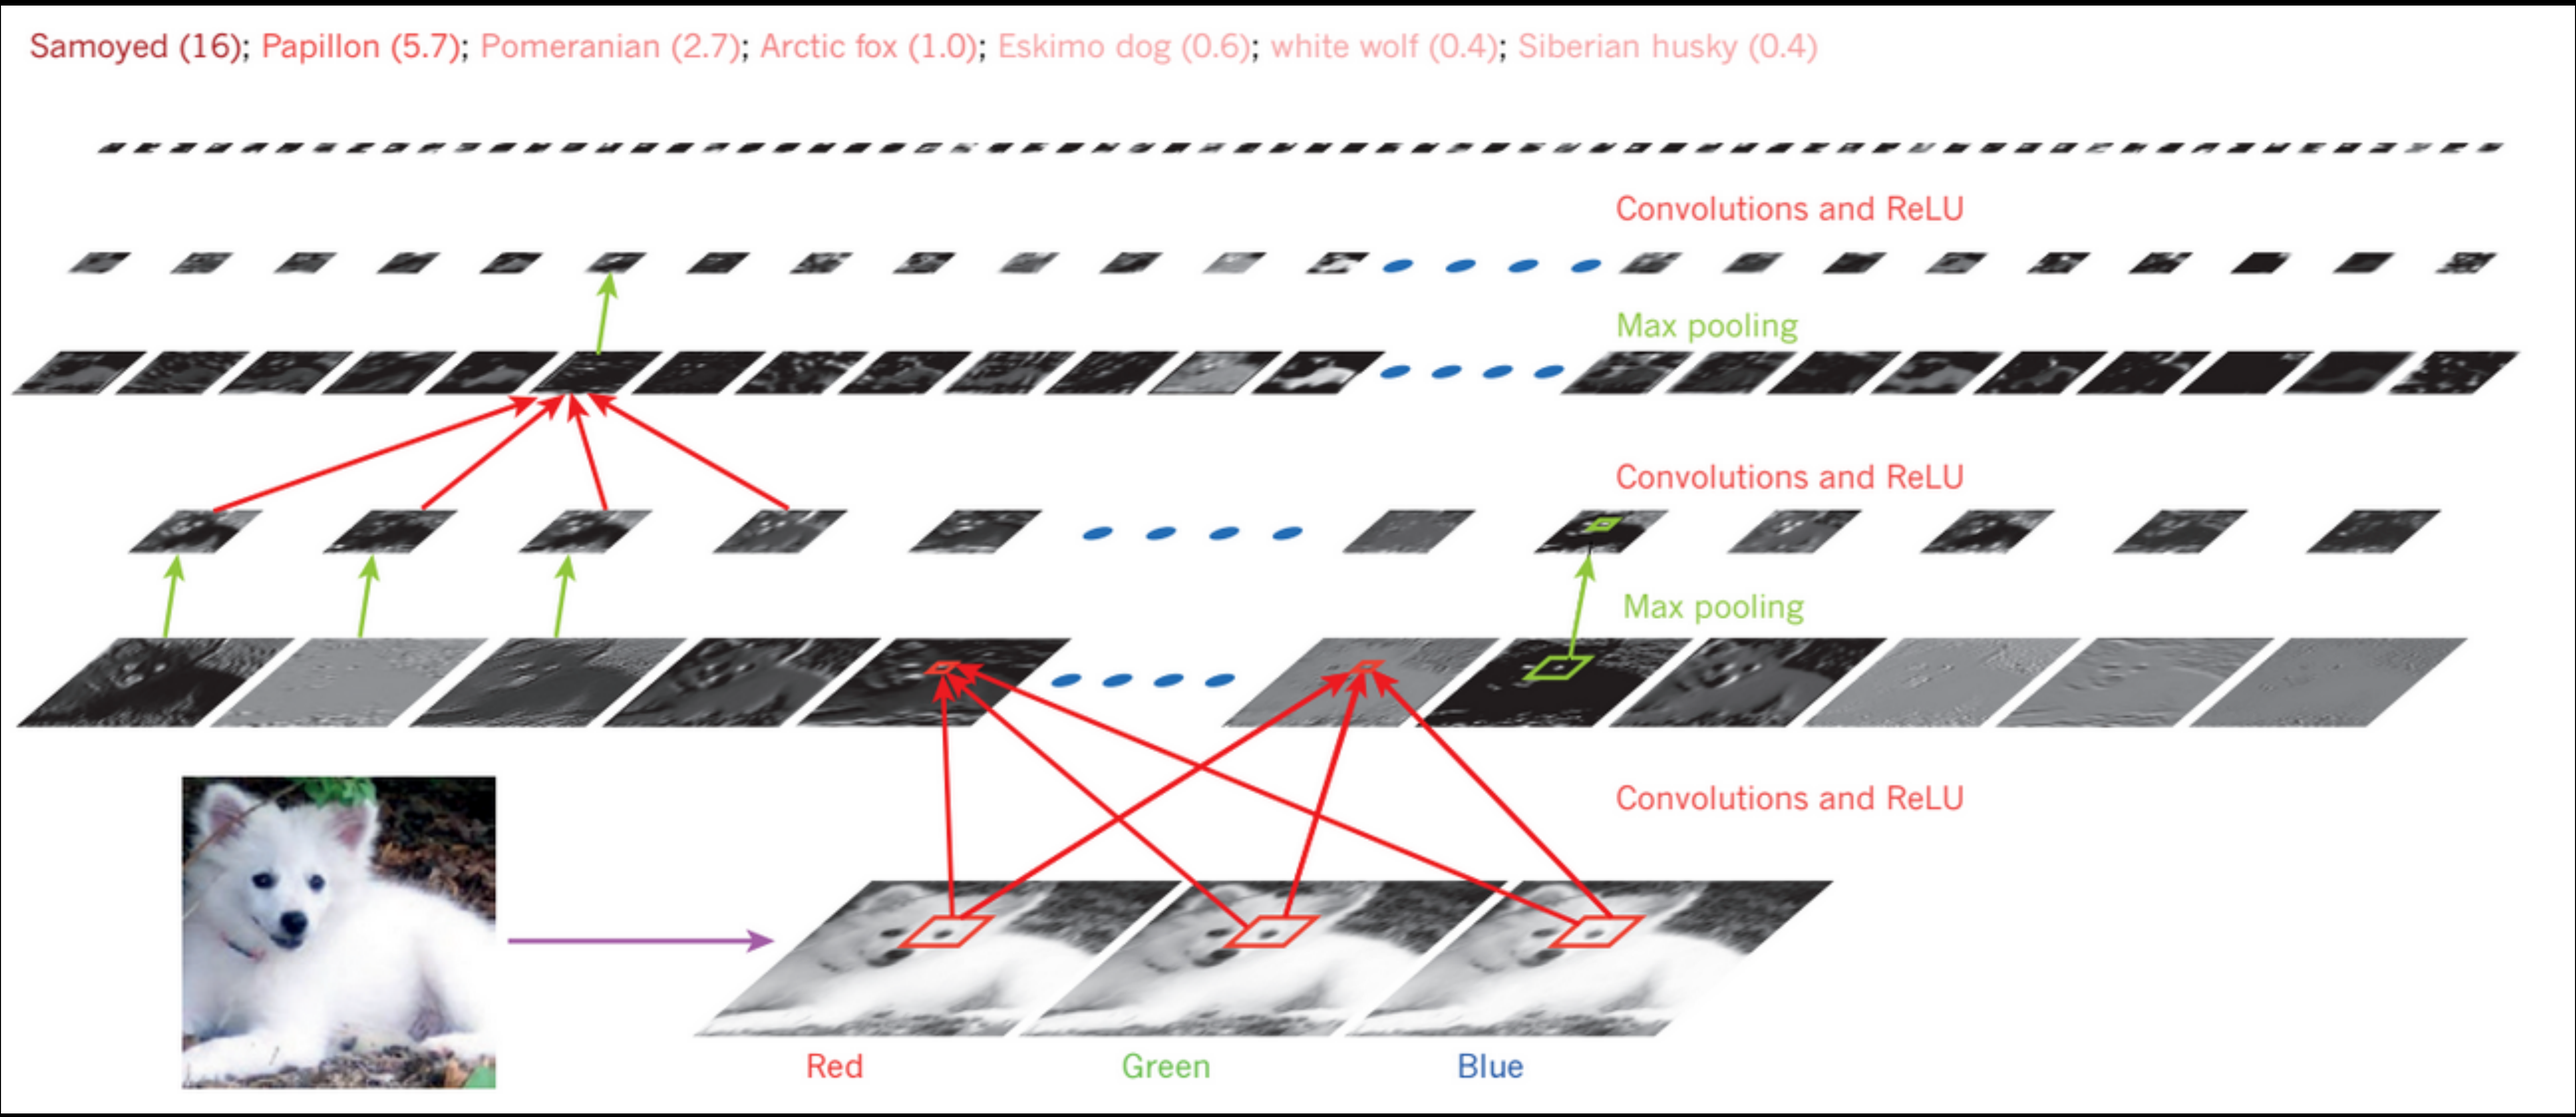
Samoyed (16); Papillon (5.7); Pomeranian (2.7); Arctic fox (1.0); Eskimo dog (0.6); white wolf (0.4); Siberian husky (0.4)
Convolutions and ReLU
Max pooling
Convolutions and ReLU
Max pooling
Convolutions and ReLU
Red
Green
Blue

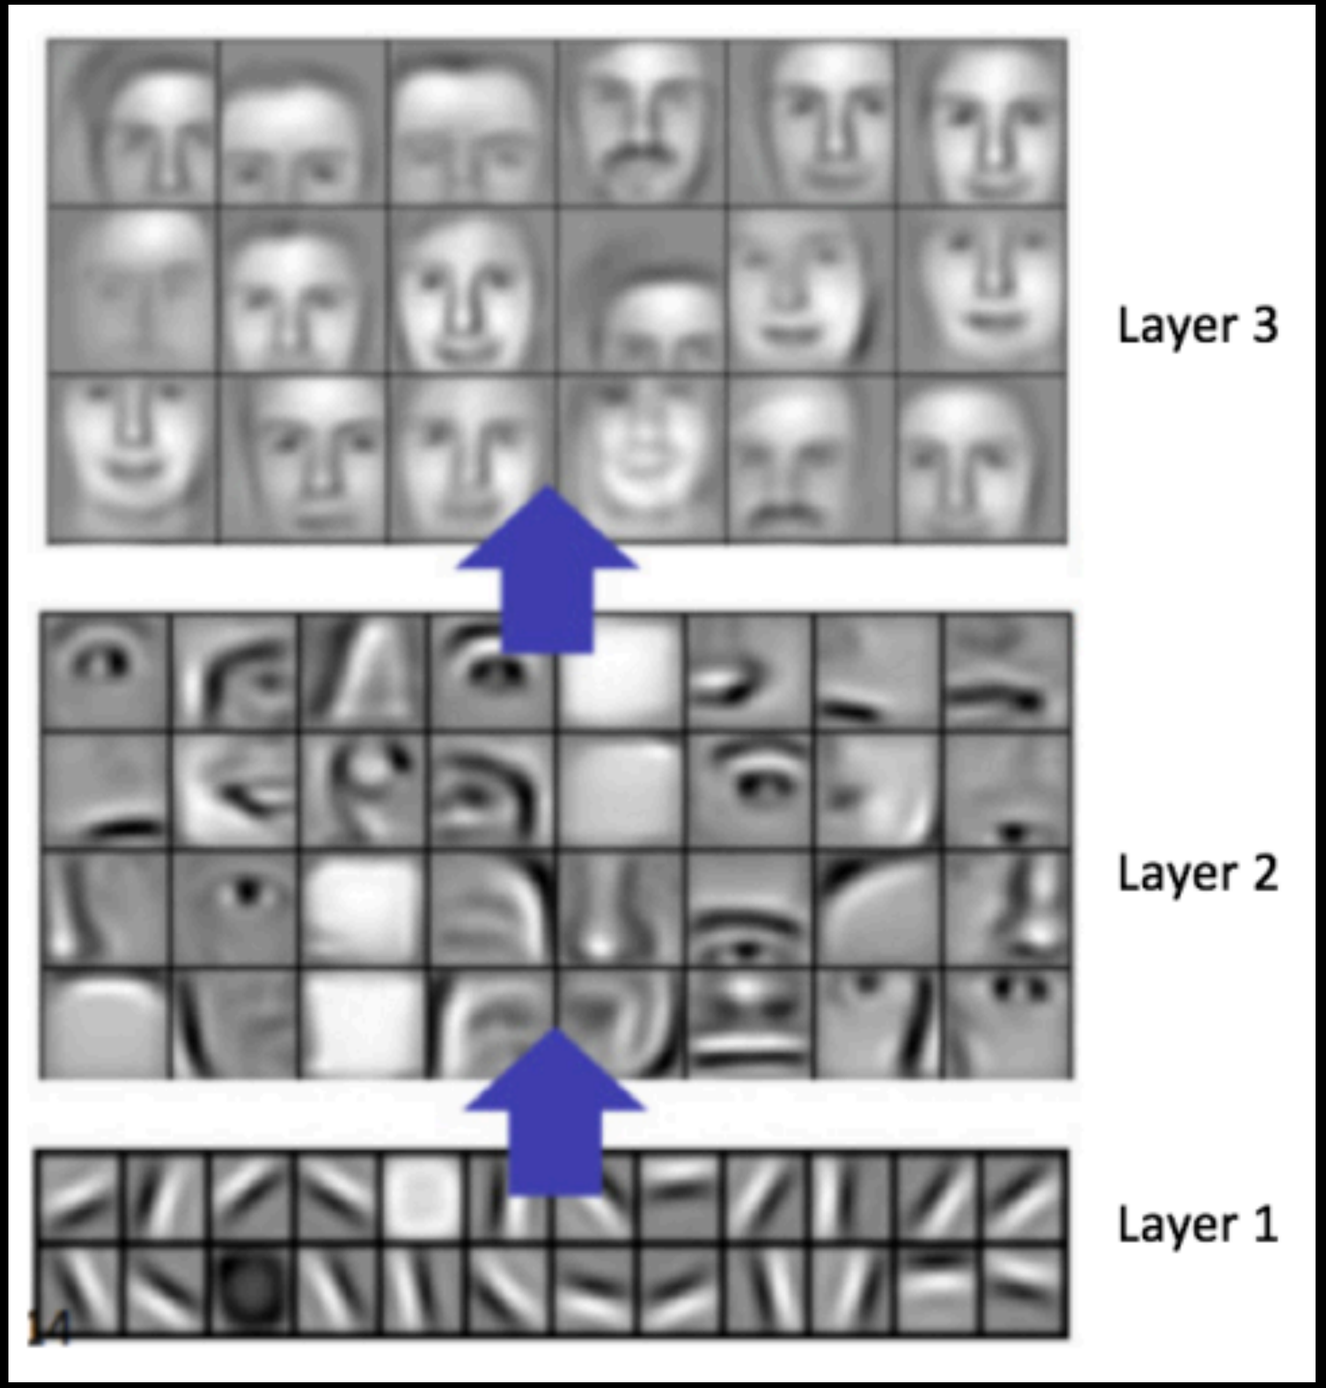
Layer 3
Layer 2
Layer 1

Reading Images (FaceBook)

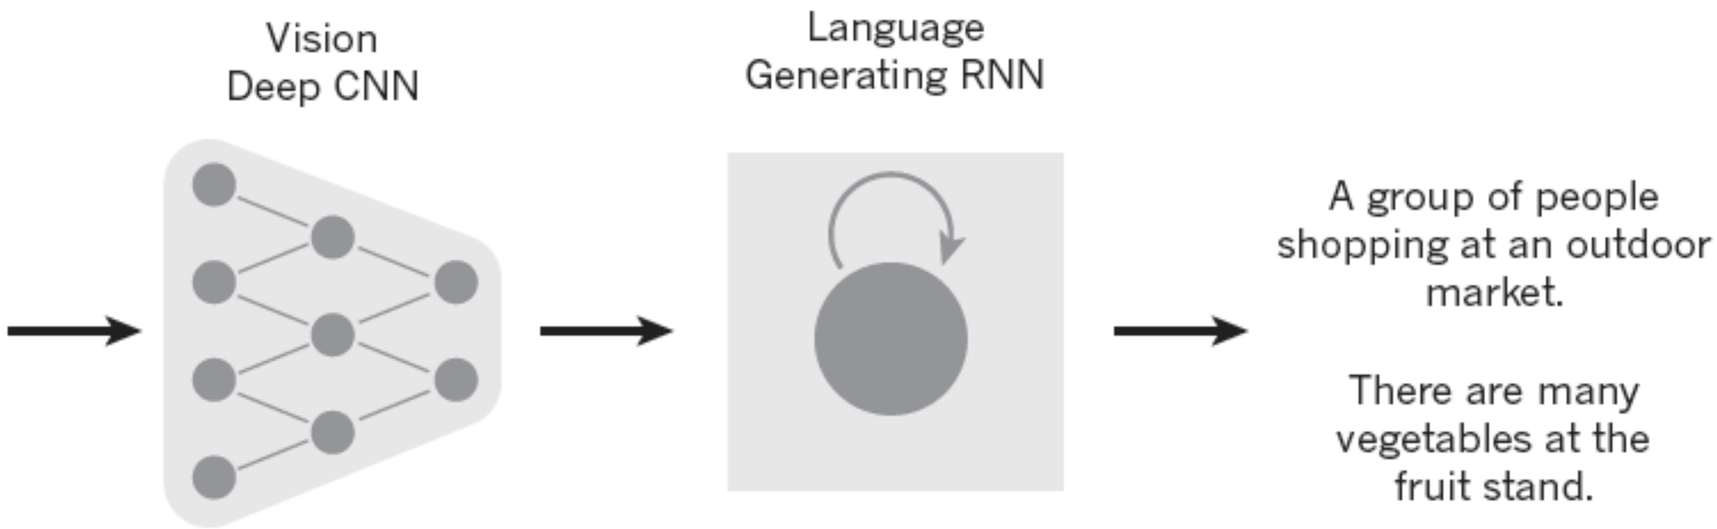

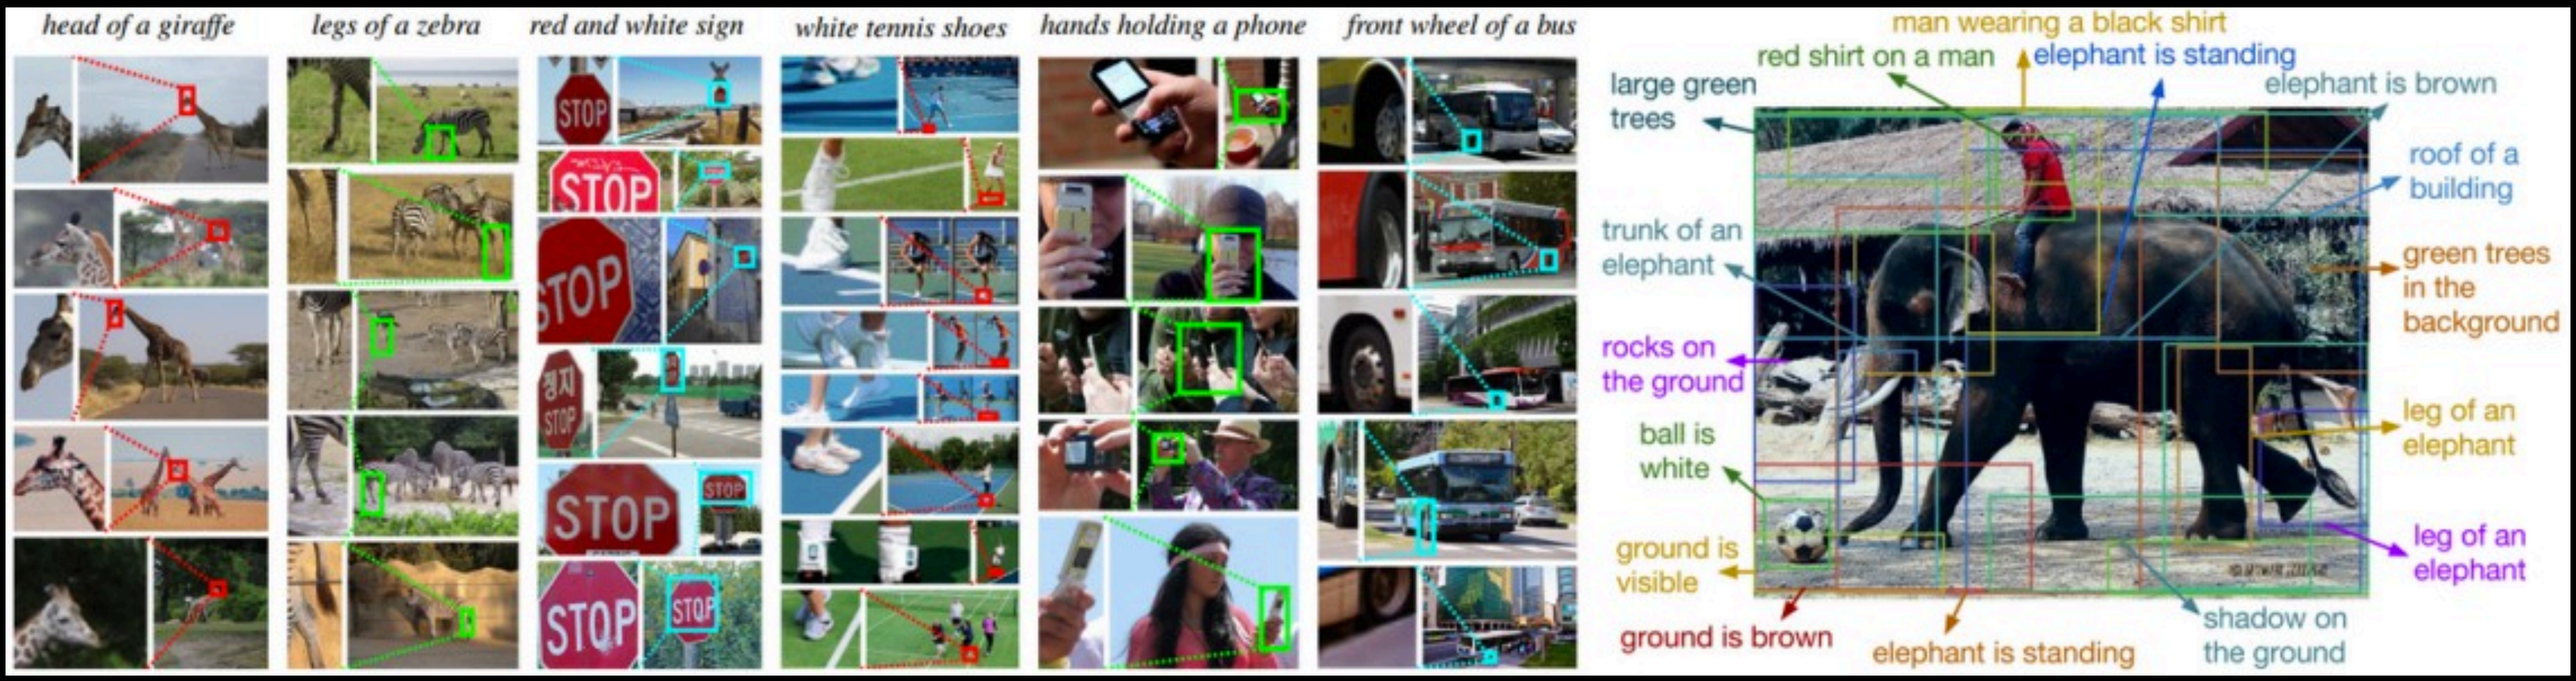

See NeuralTalk: http://cs.stanford.edu/people/karpathy/deepimagesent/generationdemo/

A group of **people** sitting on a boat in the water.

Main Drivers of the Breakthrough

- DATA
- COMPUTATION
- MACHINE LEARNING

A11 Bionic Neural Engine

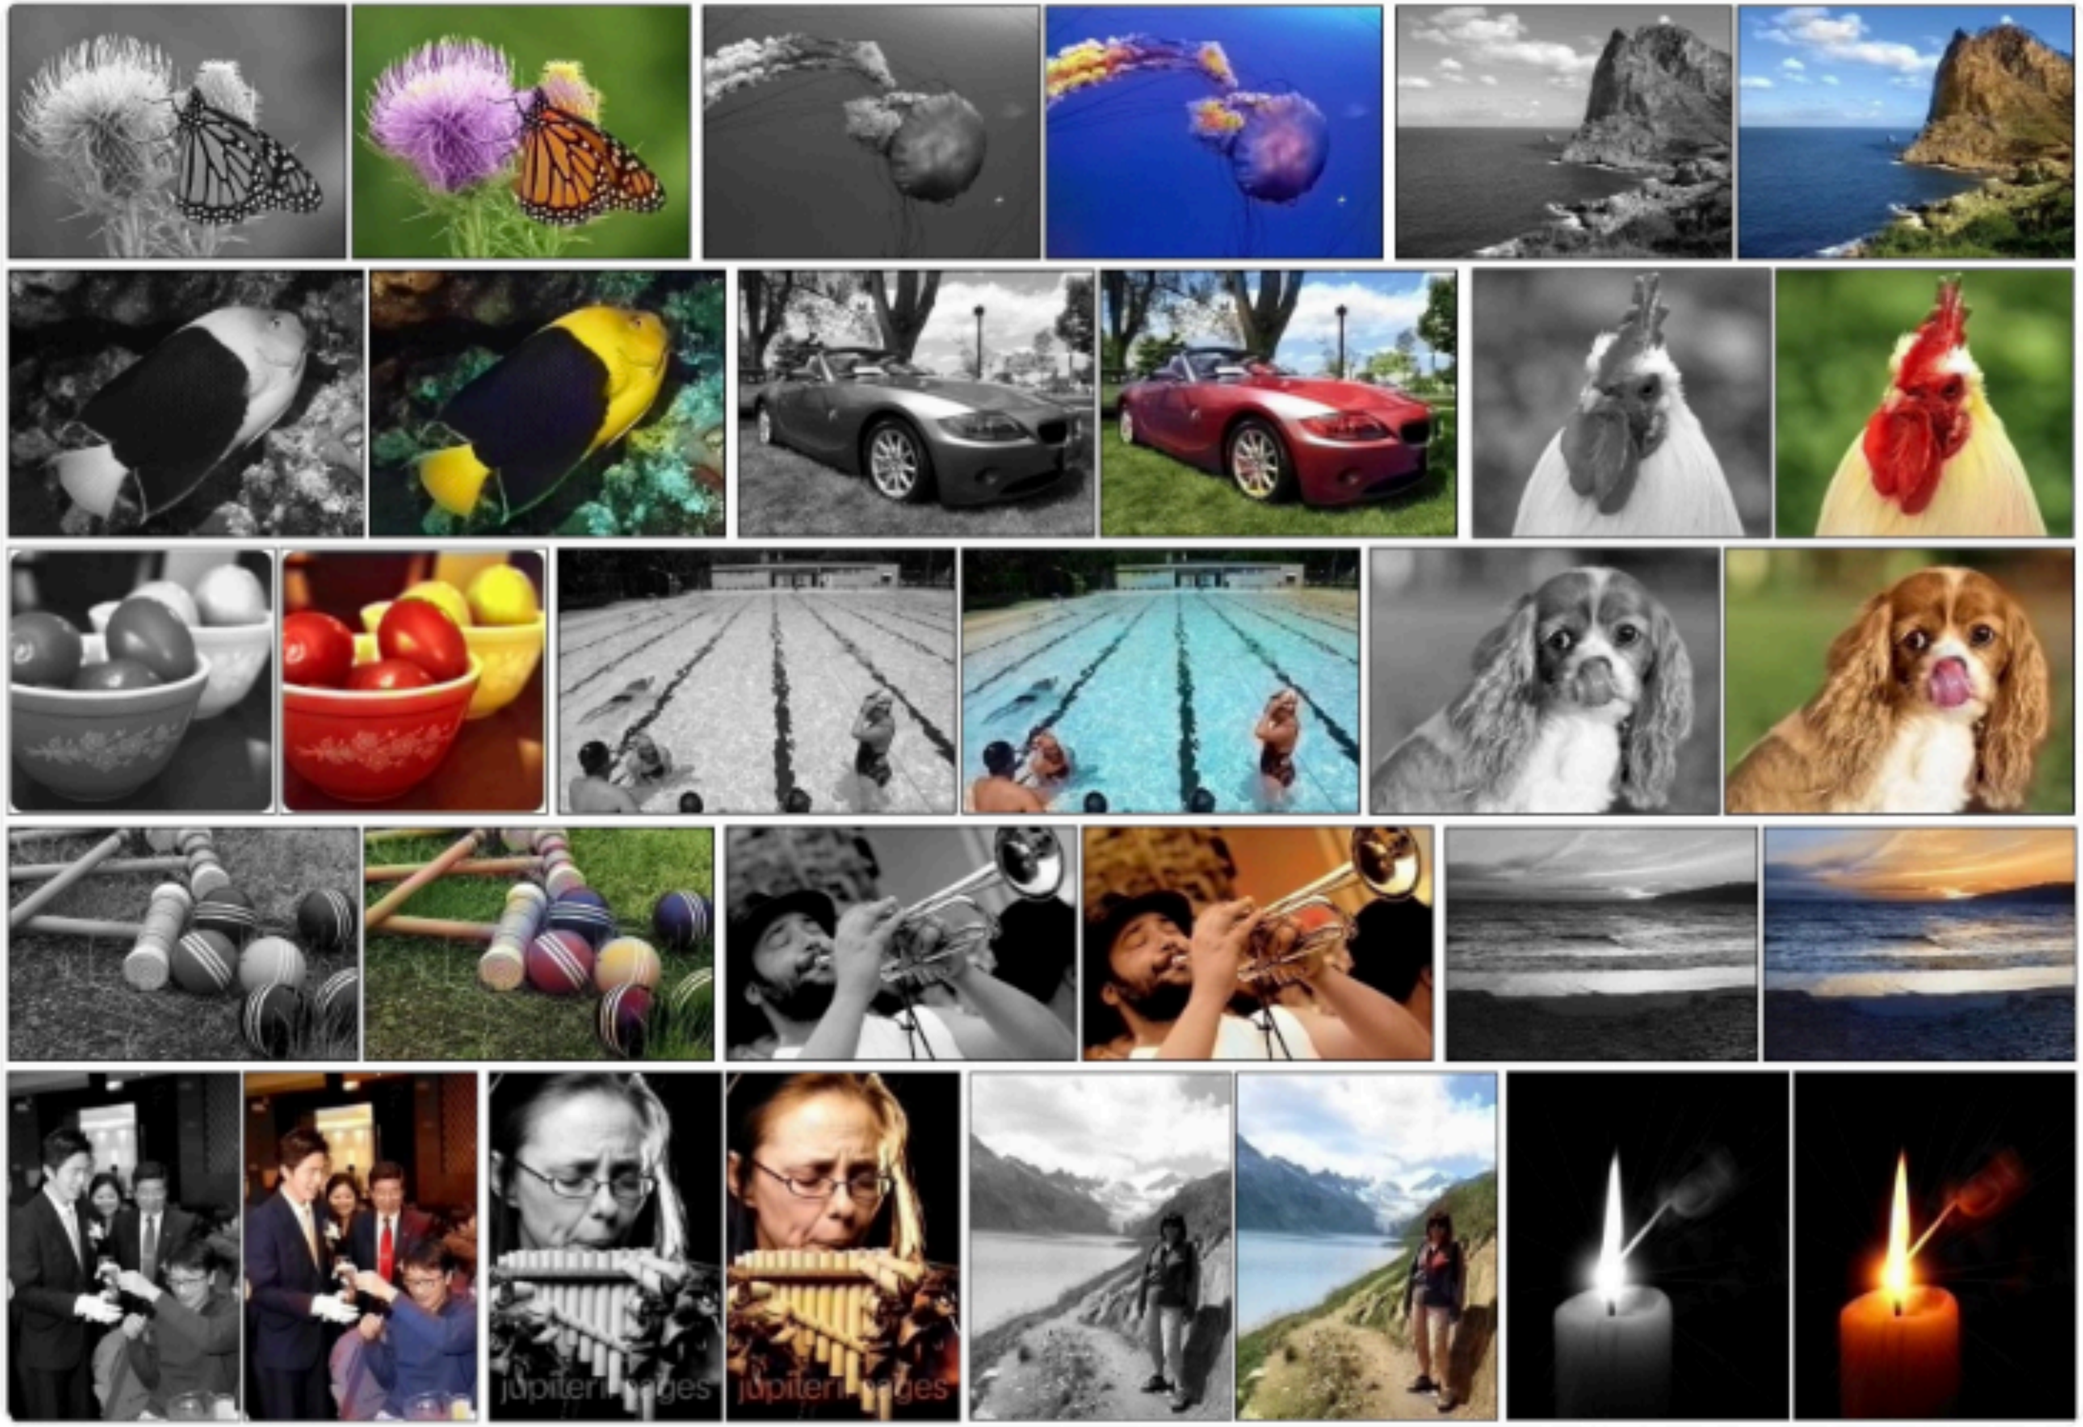

VDS

VIDEO DISCOVERY SERVICE

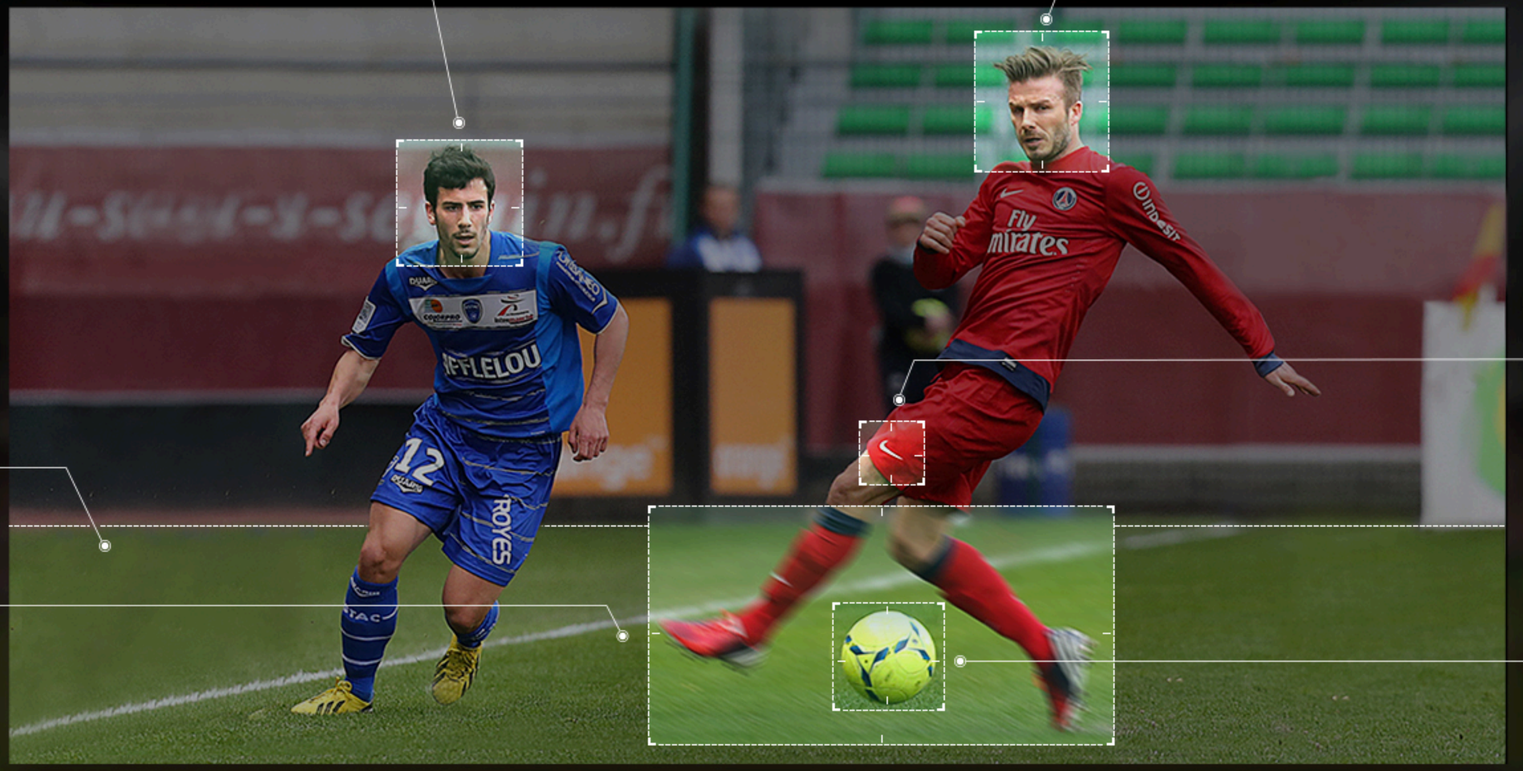

Dahl, Norouzi & Shlens (2017). Pixel Recursive Super Resolution. Google Brain

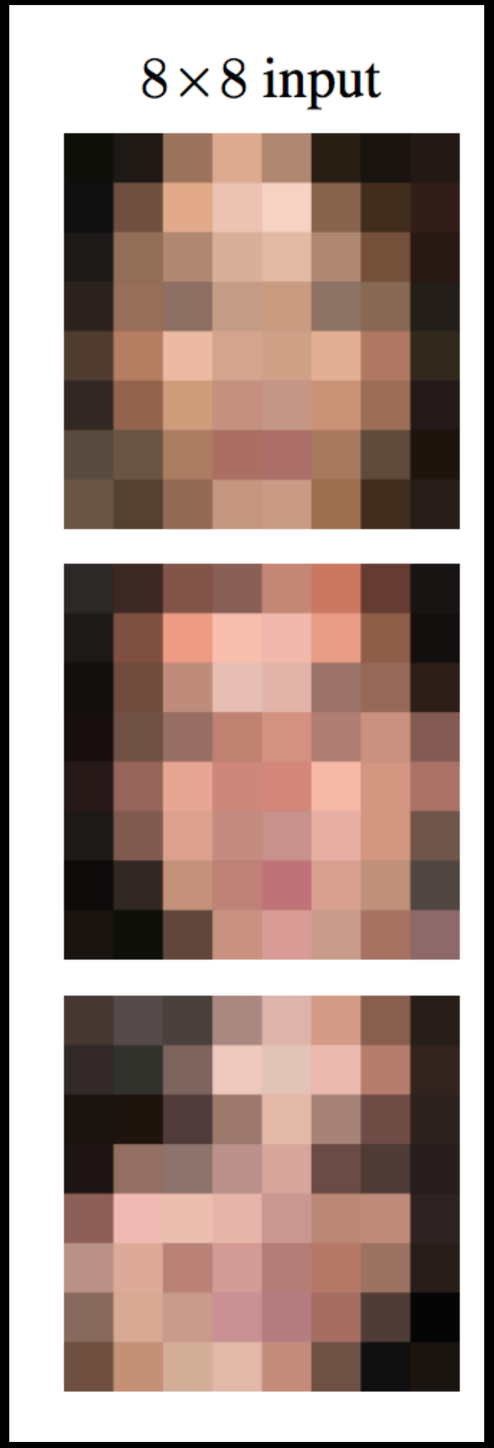

Brain reading

Neural Encoding and Decoding with Deep Learning for Dynamic Natural Vision

FREE

Haiguang Wen, Junxing Shi, Yizhen Zhang, Kun-Han Lu, Jiayue Cao, Zhongming Liu

Cerebral Cortex, https://doi.org/10.1093/cercor/bhx268

Published: 20 October 2017

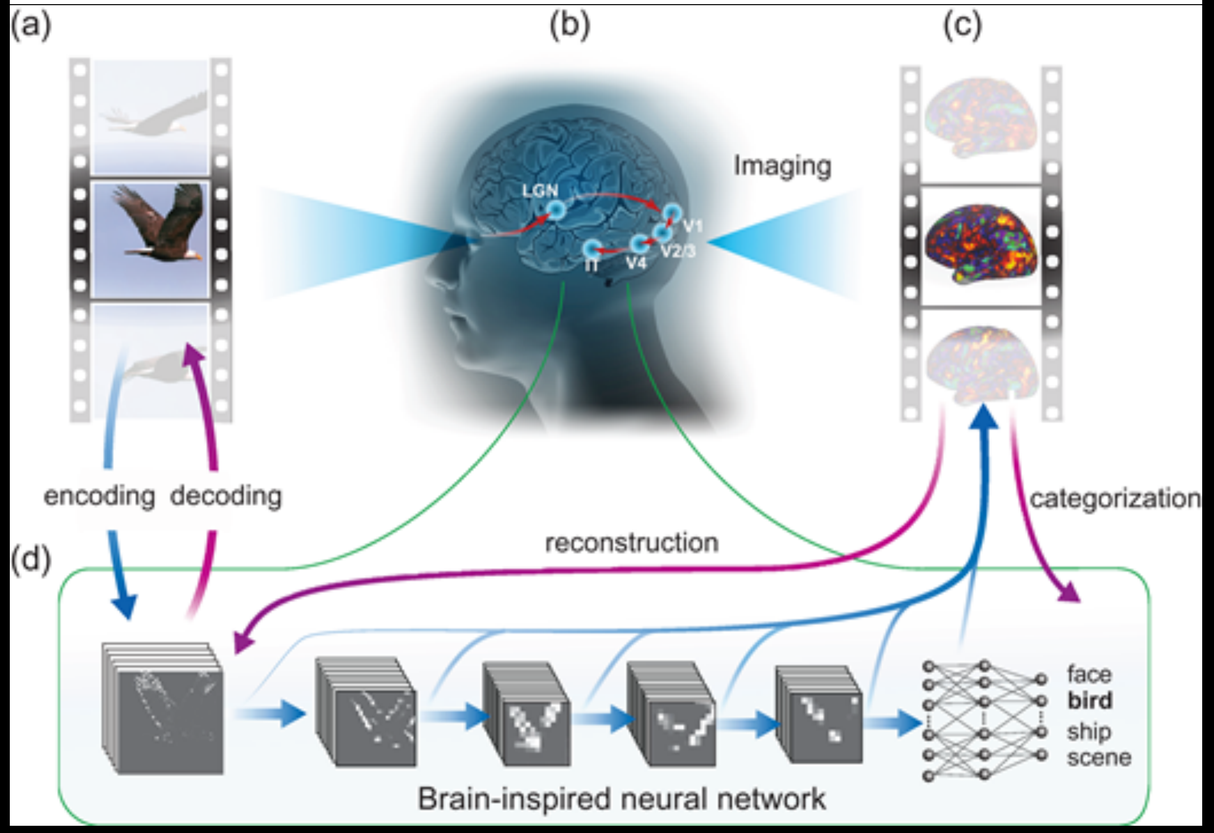


Autonomous visual navigation, no GPS, no maps

Machine Translation

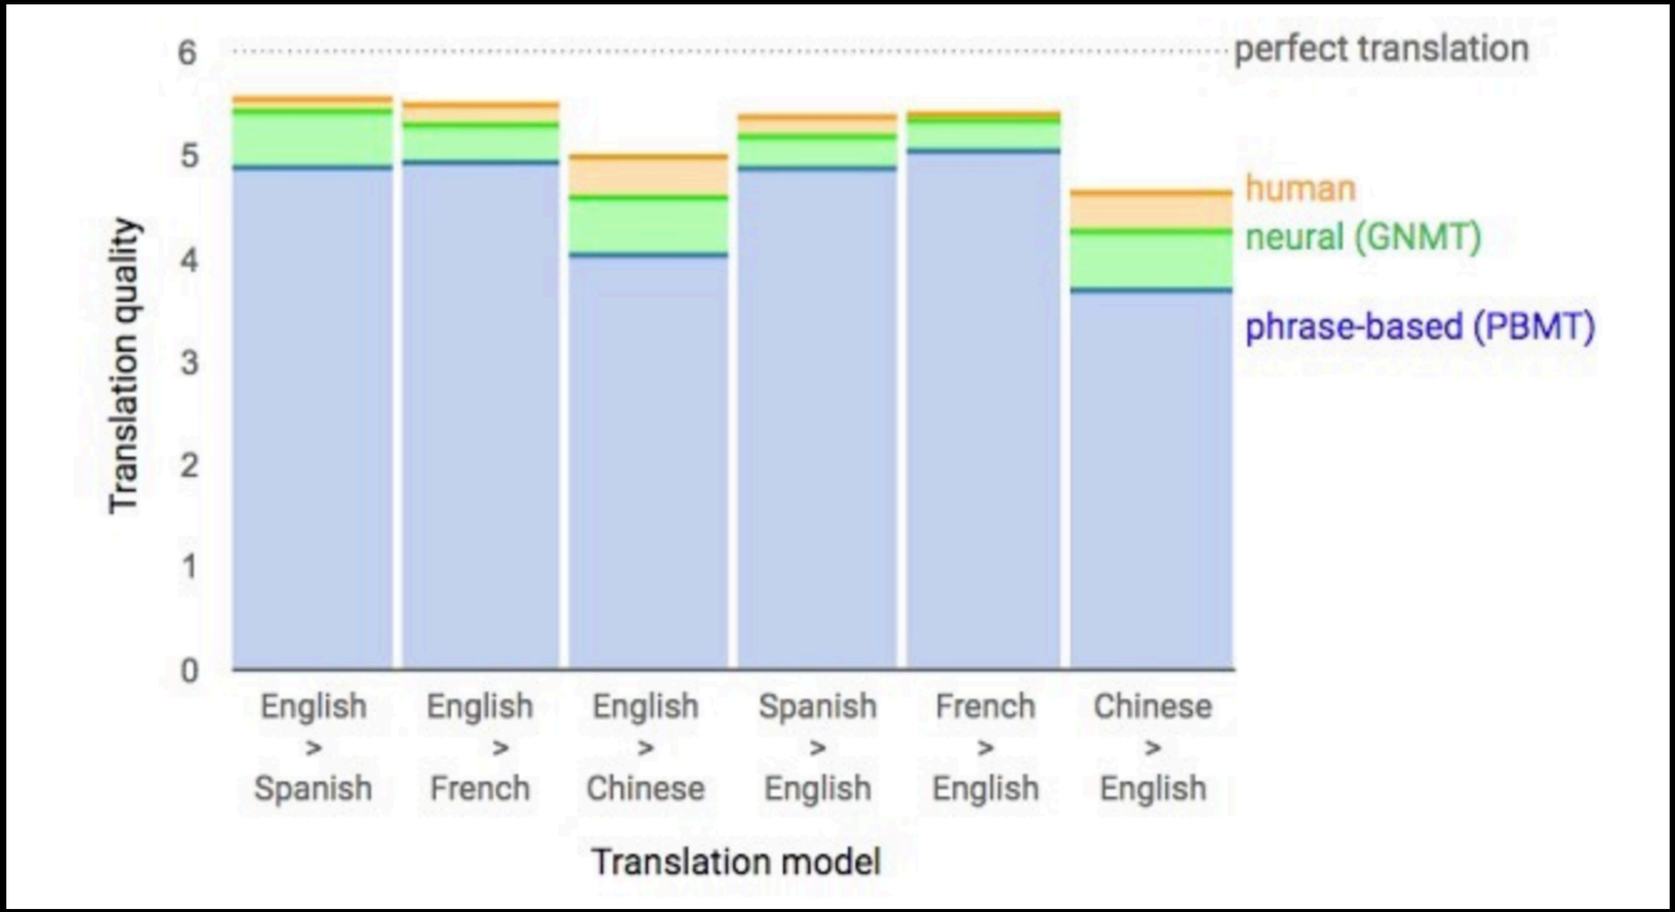

https://blog.statsbot.co/deep-learning-achievements-4c563e034257

Speech Generation

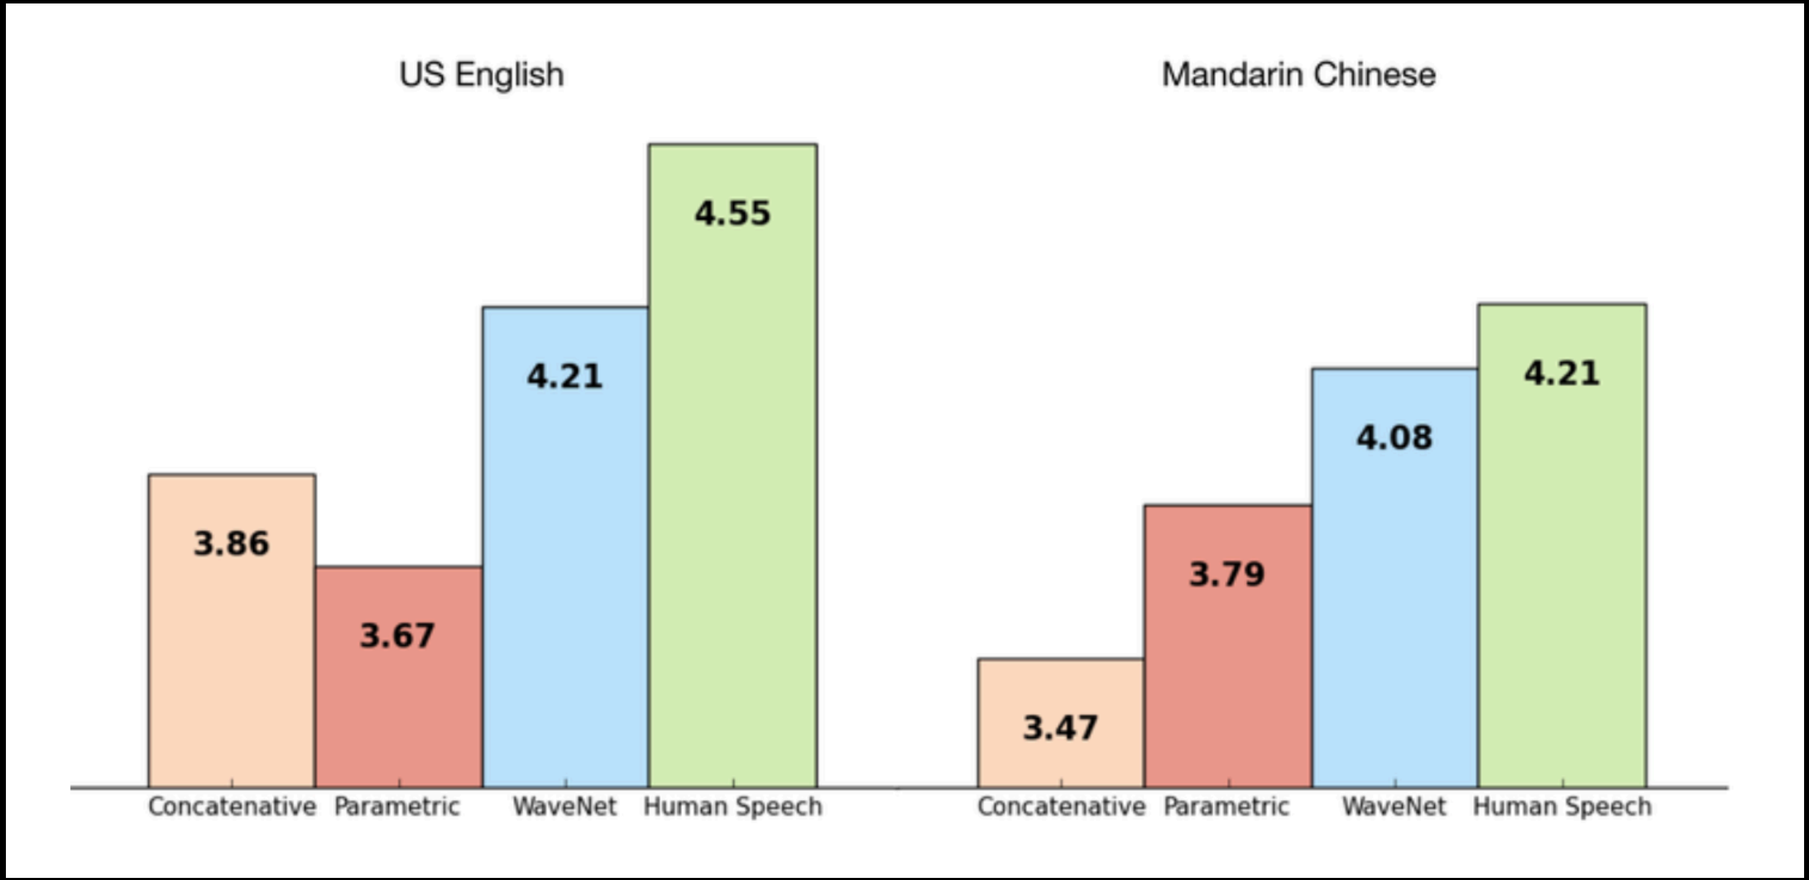

https://blog.statsbot.co/deep-learning-achievements-4c563e034257

Lip Reading

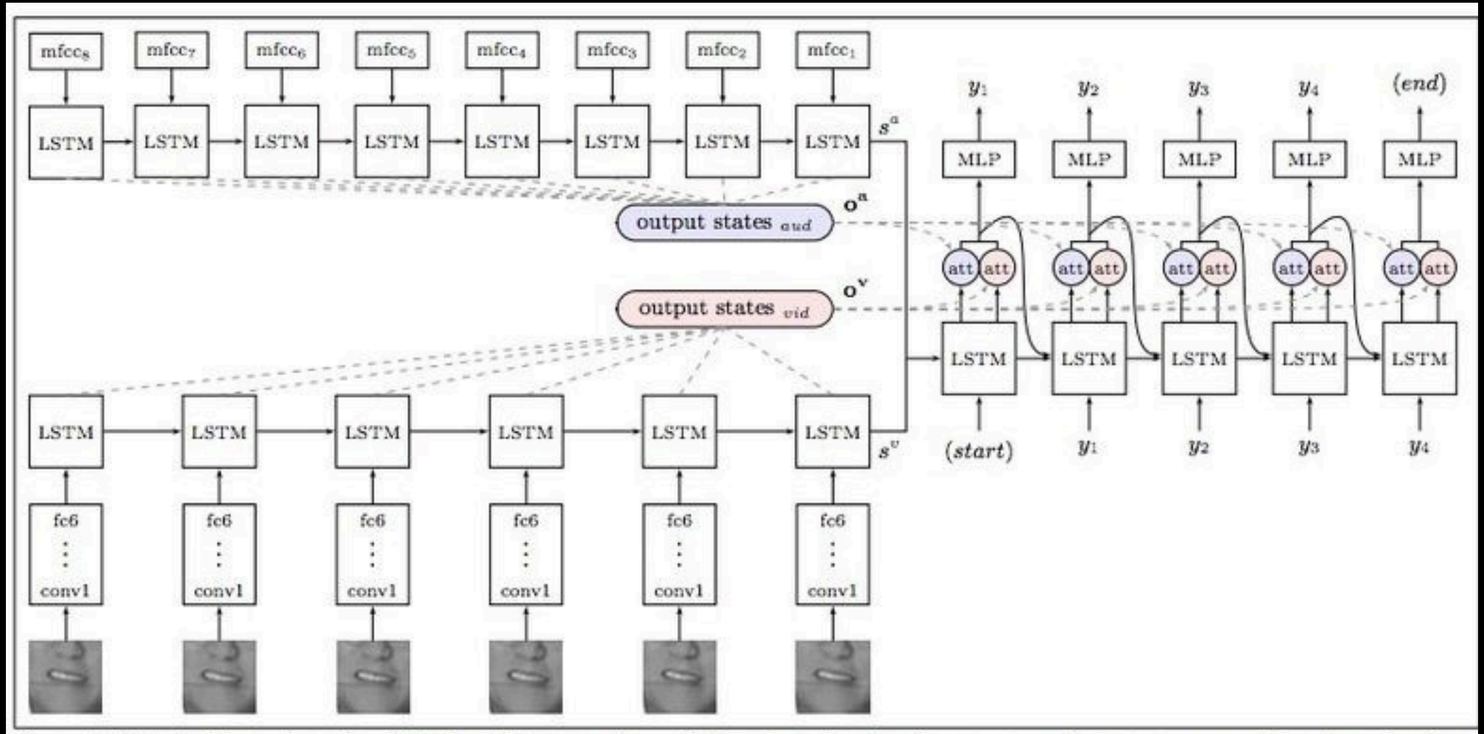

Figure 1. *Watch, Listen, Attend and Spell* architecture. At each time step, the decoder outputs a character y_i, as well as two attention vectors. The attention vectors are used to select the appropriate period of the input visual and audio sequences.

Lip Reading

Deep Learning Flaws

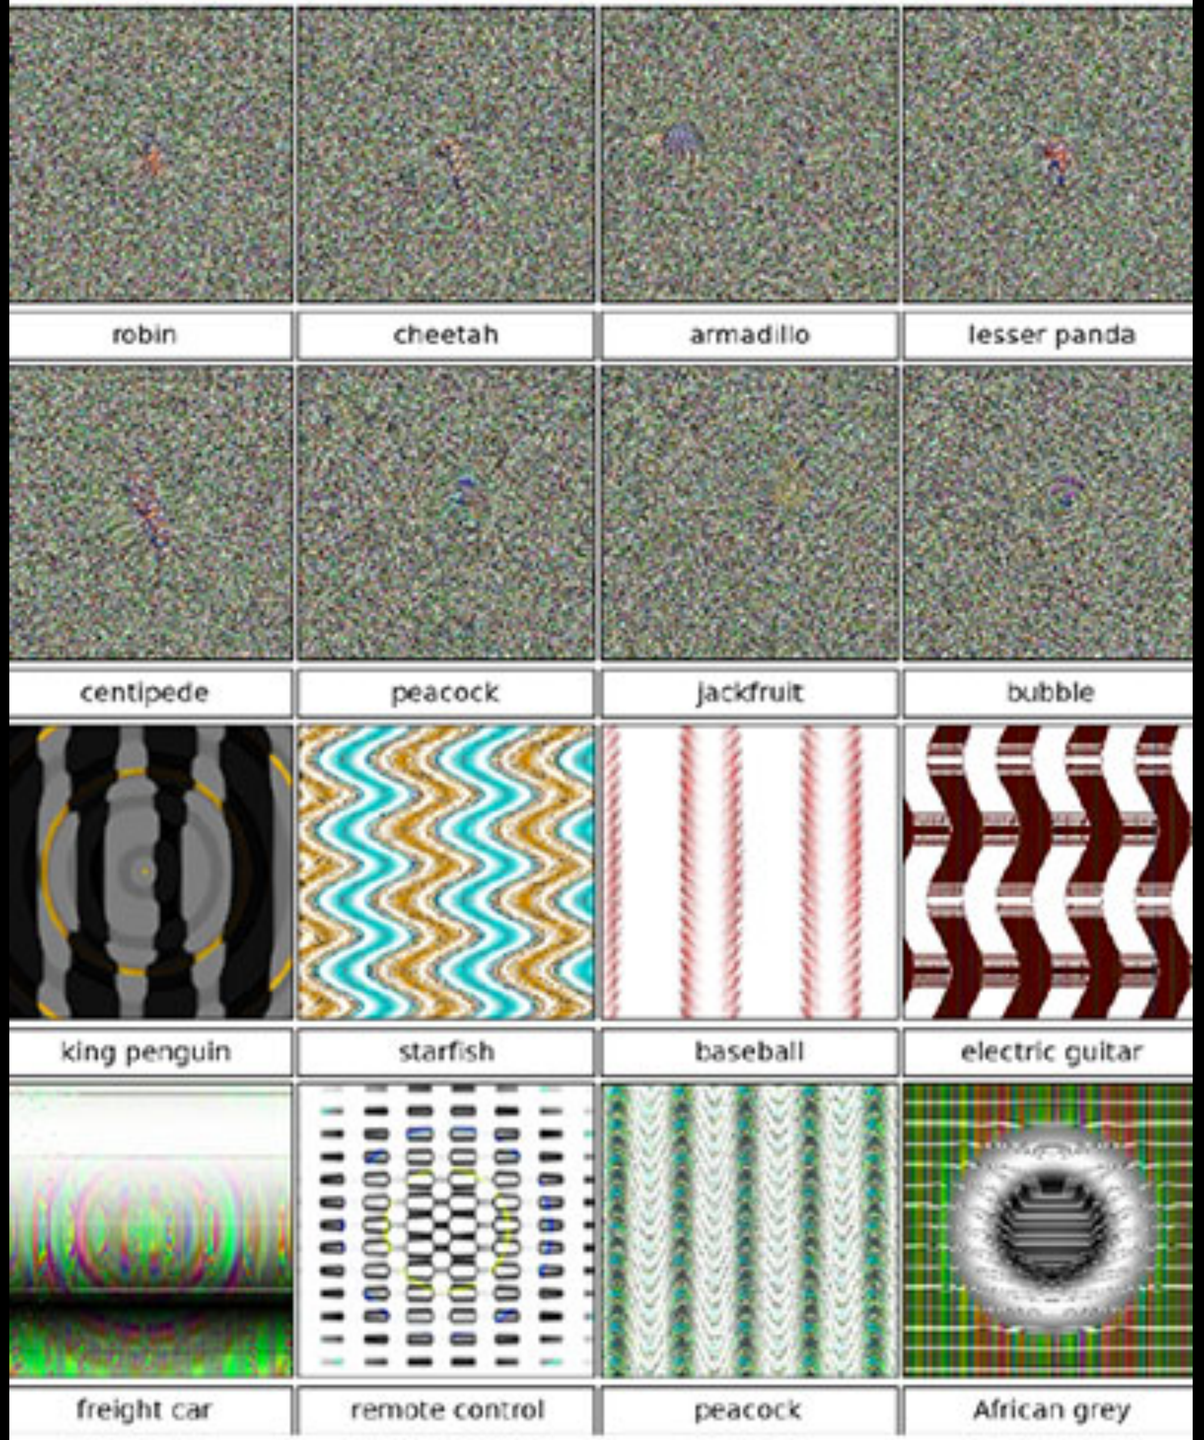
robin
cheetah
armadillo
lesser panda
centipede
peacock
jackfruit
bubble
king penguin
starfish
baseball
electric guitar
freight car
remote control
peacock
African grey

Deep Learning Strengths

- Ability to integrate information from huge heterogeneous sources
- Ability to predict
- Ability to detect/recognise
- Ability to discover patterns

Keep in mind that..

- DL algorithms lack “common sense”
- DL cannot put events into their broader context
- DL depends critically on the quality of the underlying statistics/data
- DL is opaque
- DL is ill-understood

DL algorithms lack “common sense”

A woman is throwing a **frisbee** in a park.

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

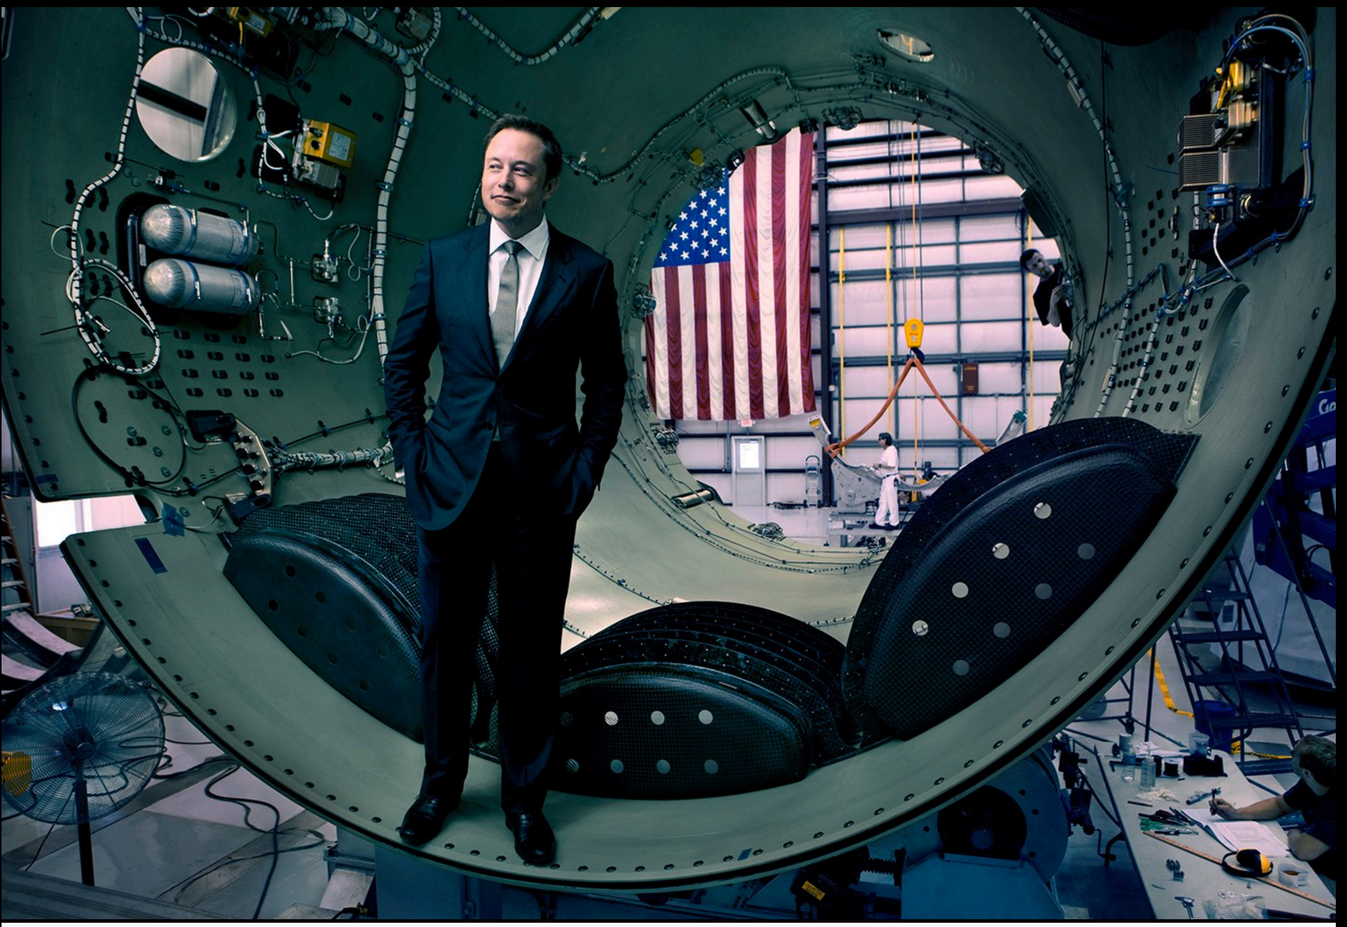

ELON MUSK'S BILLION-DOLLAR CRUSADE TO STOP THE A.I. APOCALYPSE

RODNEY BROOKS *Robots, AI, and other stuff*

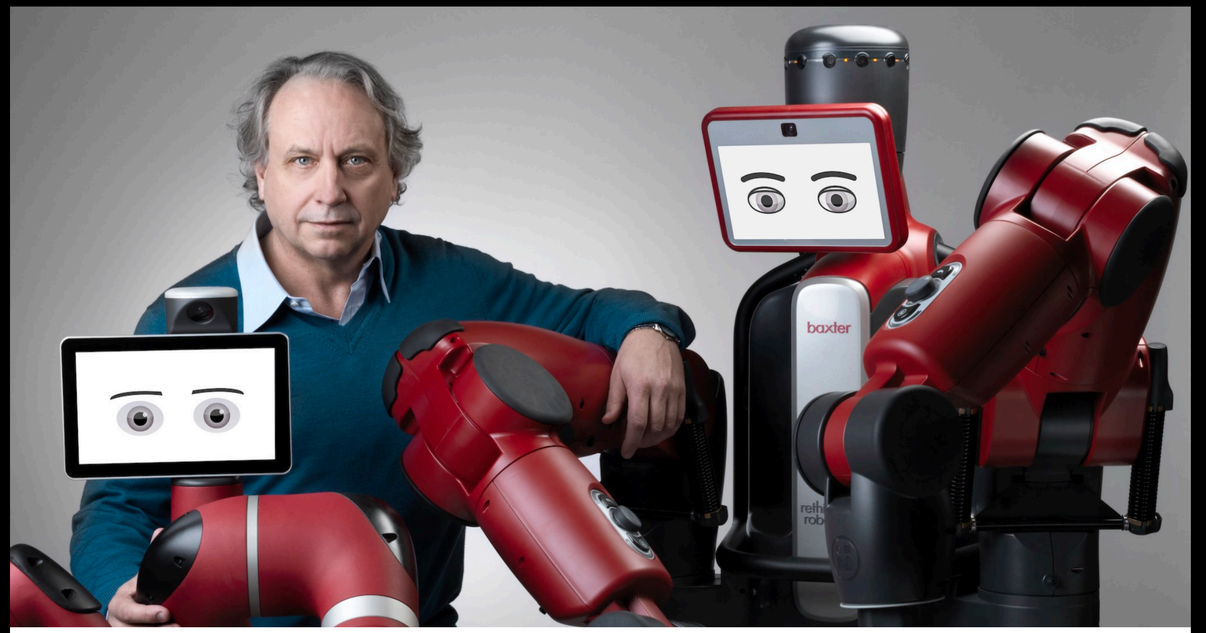

Intelligent Machines

The Seven Deadly Sins of AI Predictions

Mistaken extrapolations, limited imagination, and other common mistakes that distract us from thinking more productively about the future.

by Rodney Brooks October 6, 2017

MIT Technology Review

“Amara’s Law”

We tend to overestimate the effect of a technology in the short run and underestimate the effect in the long run.

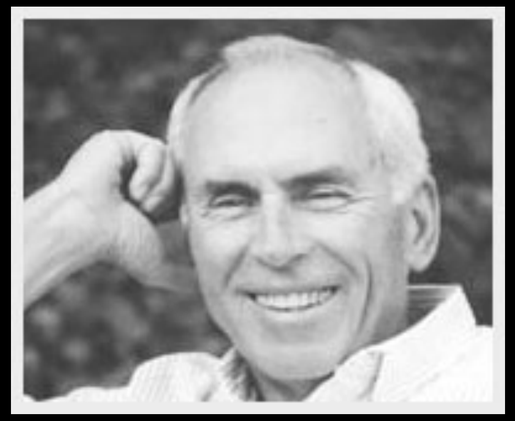

Roy Amara (1925-2007)
Institute for the Future
Palo Alto.

Transparent information, communication and modalities for the exercise of the rights of the data subject (Article 12, GDPR)

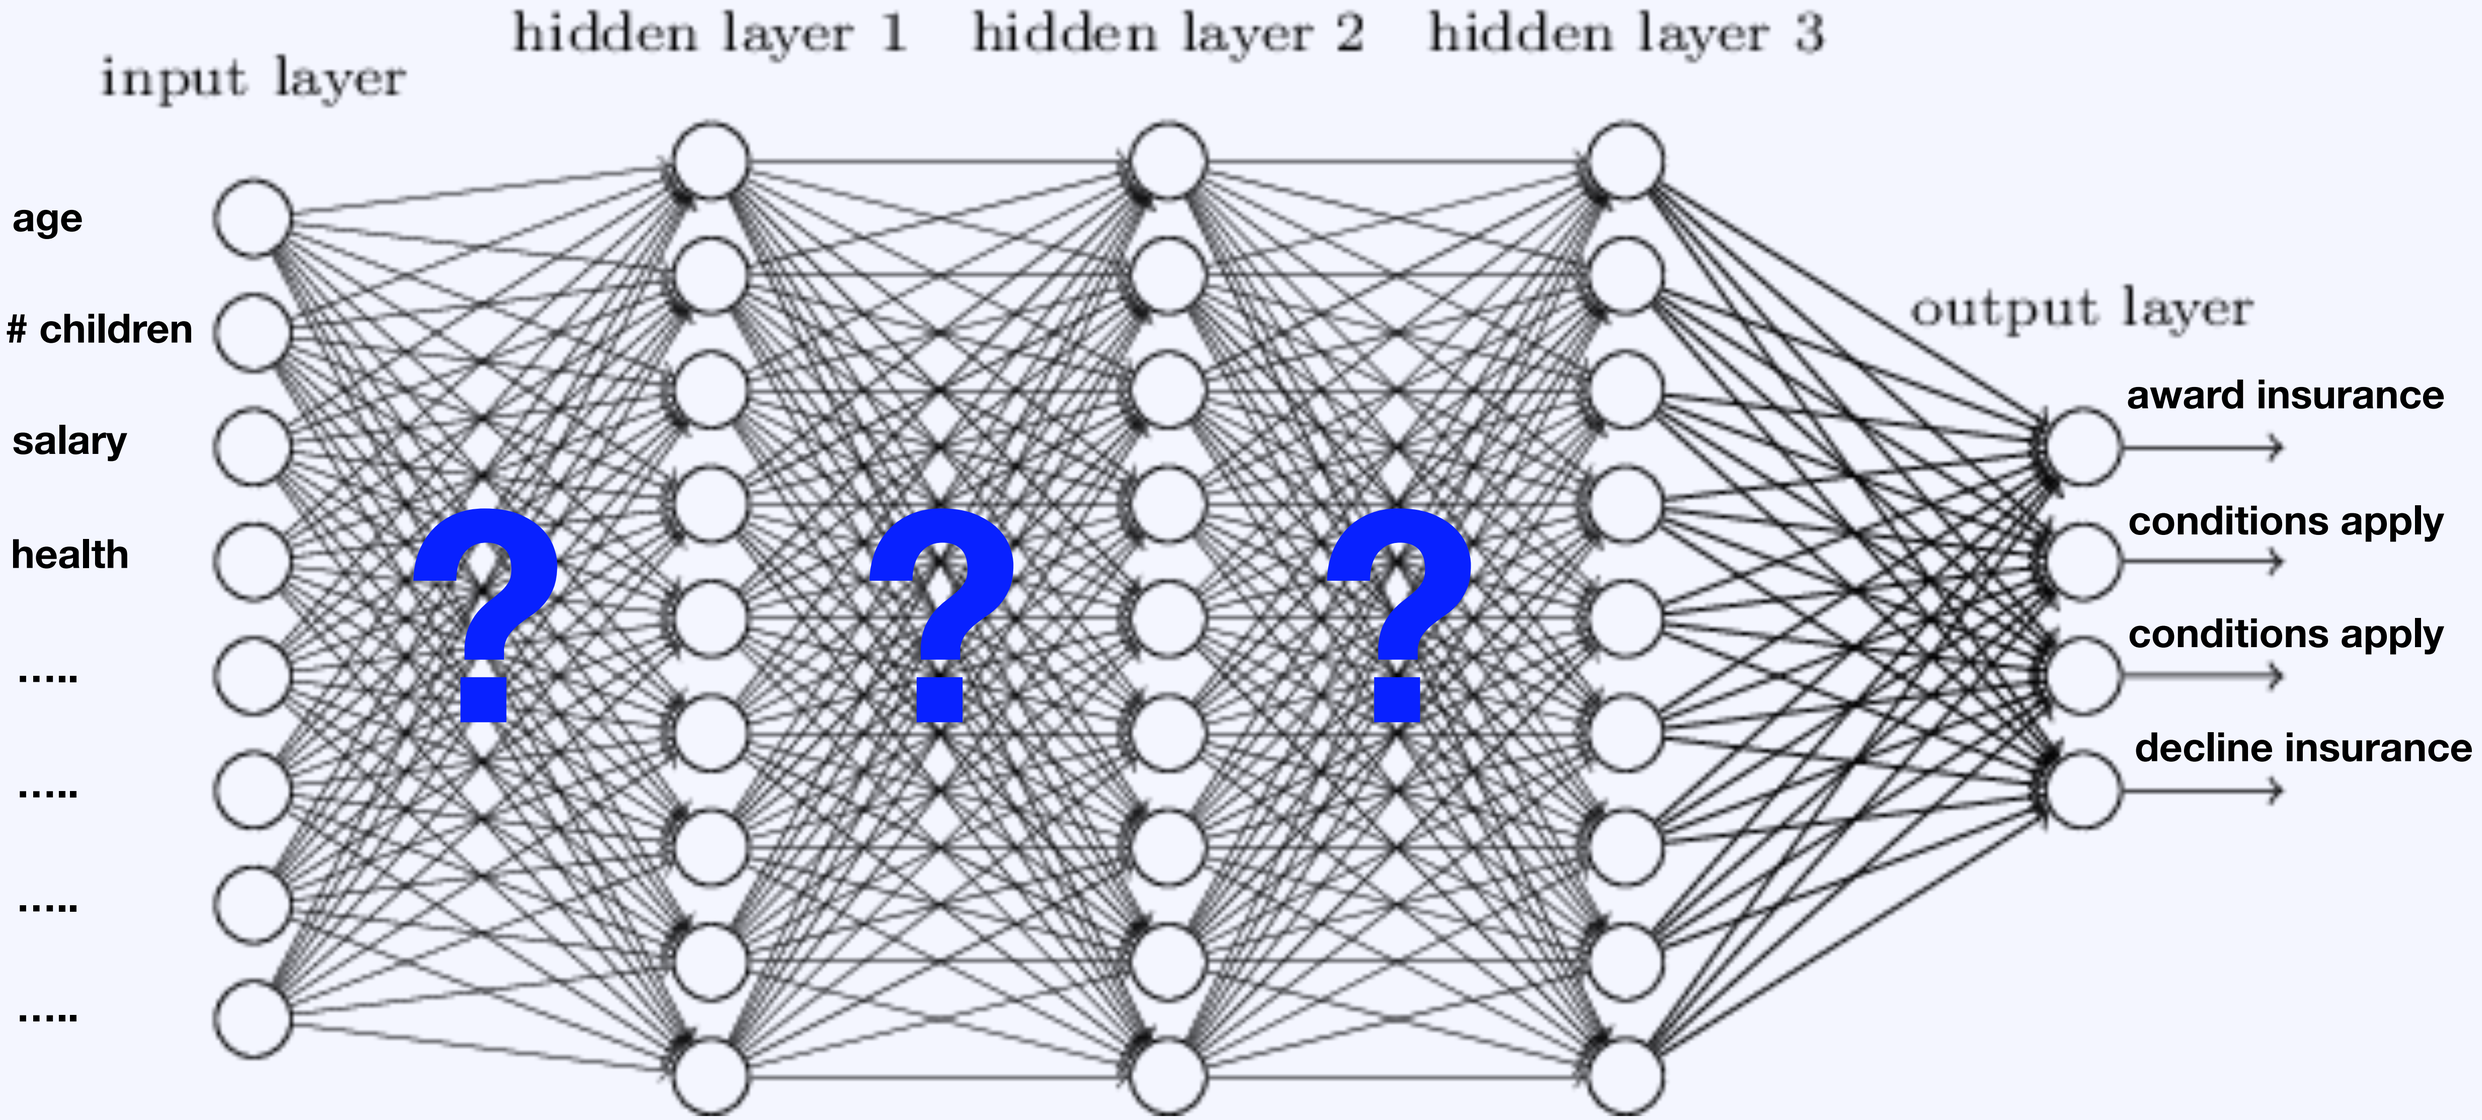
input layer
hidden layer 1
hidden layer 2
hidden layer 3
output layer
age
children
salary
health
.....
.....
.....
.....
?
?
?
award insurance
conditions apply
conditions apply
decline insurance

Conclusions

- The *true* AI revolution is not about robots taking over humanity, but about an exciting powerful machine learning method that is conquering many domains
- Many domains/jobs will transform dramatically (also in science)
- Deep learning algorithms offer tremendous powerful tools but still lack full understanding.


Your help is needed!