Machine learning tutorial

Melissa Lopez ft. ChatGPT

Mentimeter code: **6471 8286**

Brief history of Machine Learning
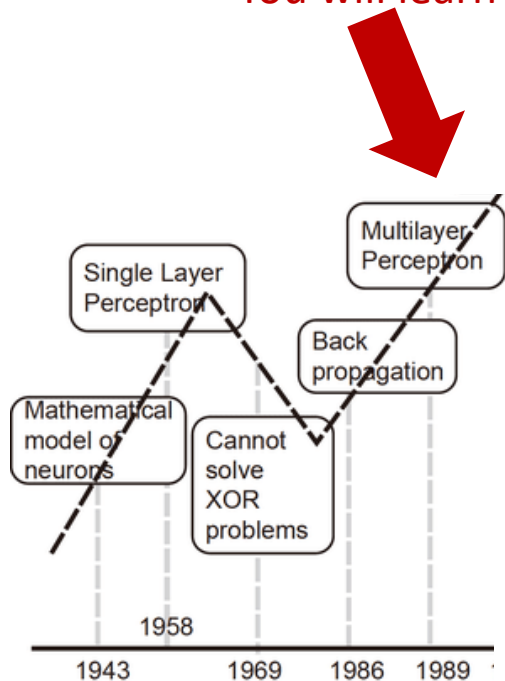
1. Big Data (WWW in 90s)

2. Hardware (GPUs, TPUs)

3. Software (Keras, Pytorch,…)

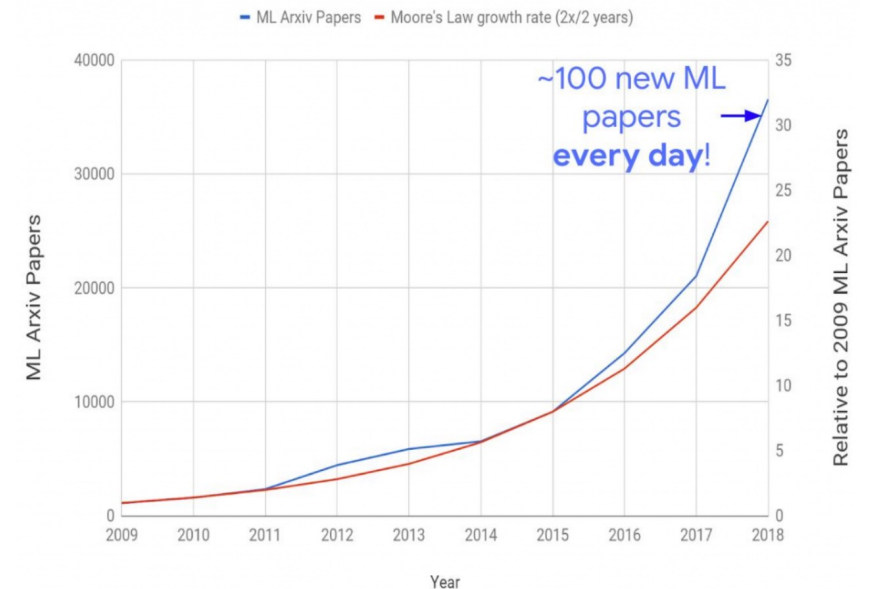What happened?

You will learn this

Chart obtained from LinkedIn, that is likely true (if someone knows the original source, I will update this post to cite it)
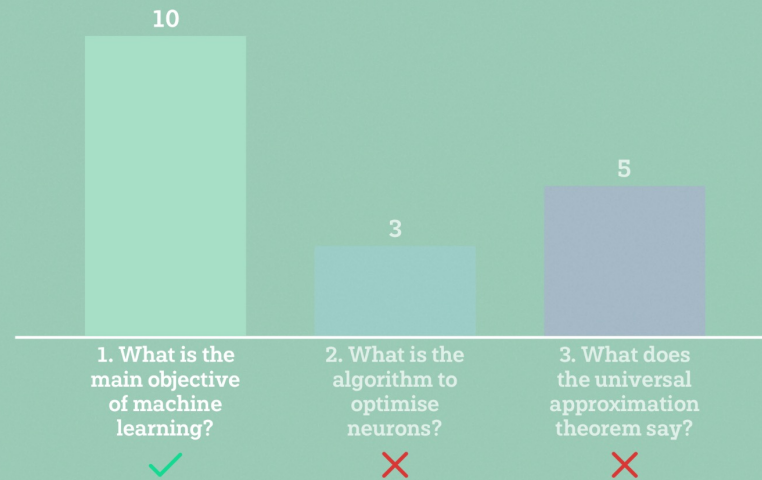
What is Machine Learning?
37 answers

bayesian
chadgpt
brain
curve-fitting
heuristic
black
statistics
now
matrix
ai
learn by example
big data
blah
function approximation
rickroll
magic
non-linear lookup table
parallel
artificial intelligence
overrated
linear algebra
code
optimization
data processing
trained
neurons
algorithm
unicorn
nonlinear
blackbox
regression
data
multilinear optimisation
a machine learning from d
neuron

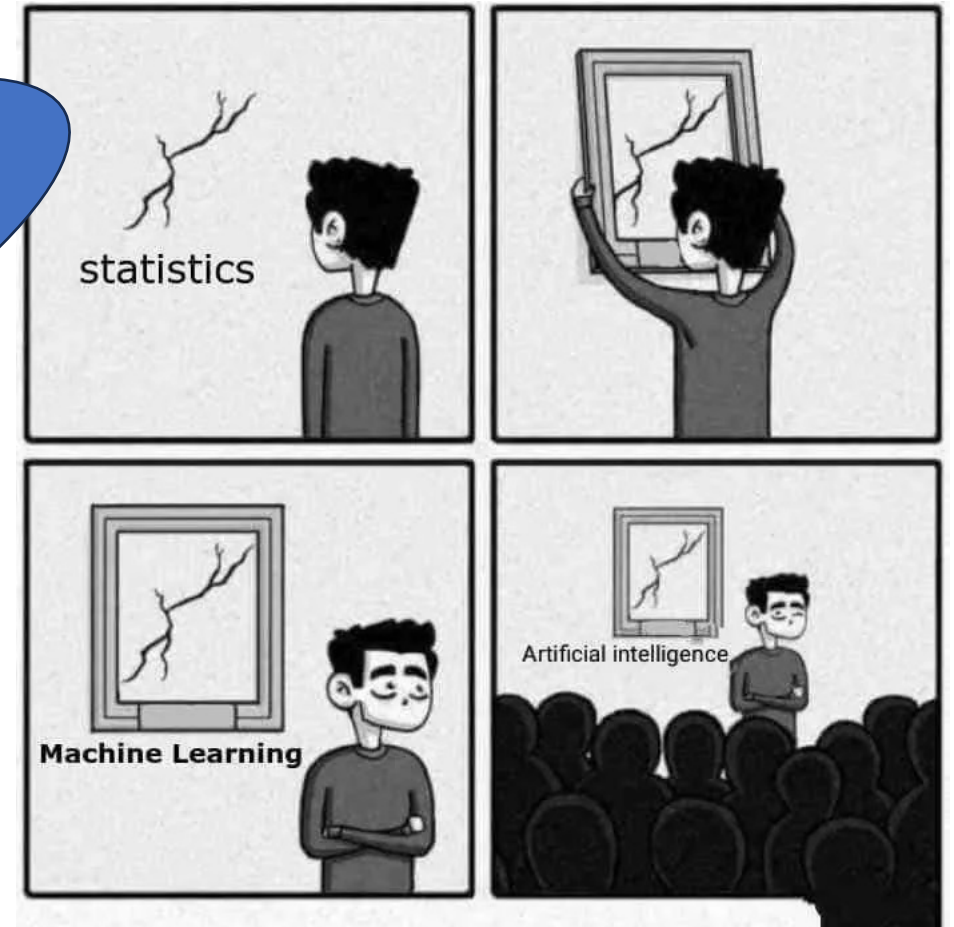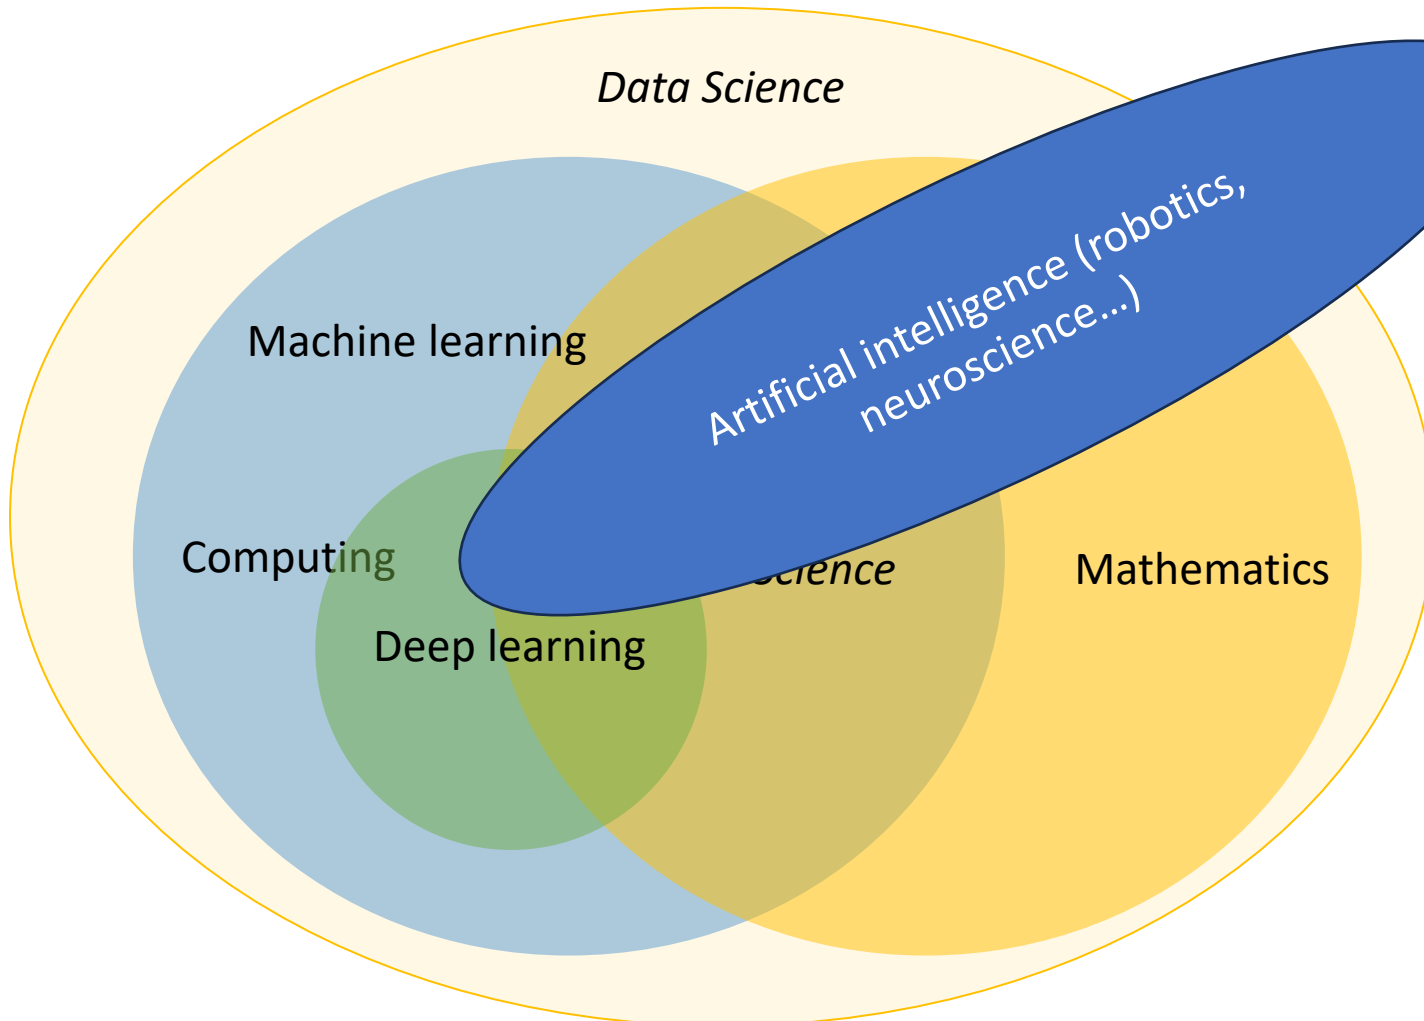# What is Machine Learning?

Let's go to Mentimeter!

Bonus track: which question was generated with ChatGPT?

What is Machine Learning?

I heard that you like Venn diagrams…

Introduction: how do *machines* learn?

Supervised $\longrightarrow$ Task-driven $\longrightarrow$ Classification: autonomous car, Galaxy Zoo

Unsupervised $\longrightarrow$ Data-driven $\longrightarrow$ Anomaly detection: fraud detection, novel particle detection
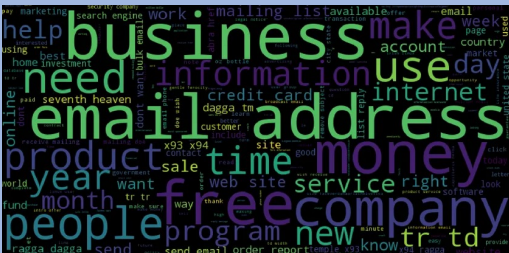
Reinforcement learning $\longrightarrow$ Error-driven $\longrightarrow$ Artificial intelligence, game theory
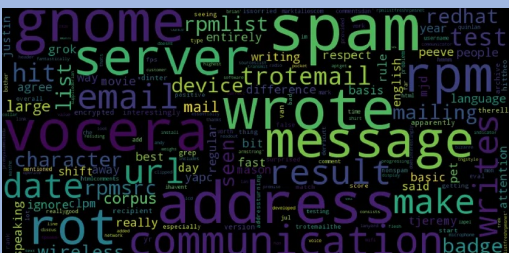
Daily life applications

Gmail

Anti-spam algorithms

Spam
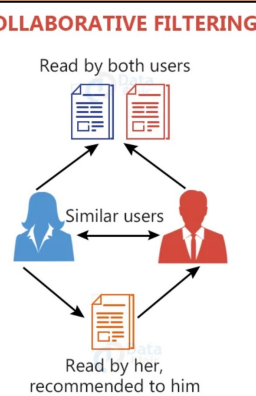
Non-spam
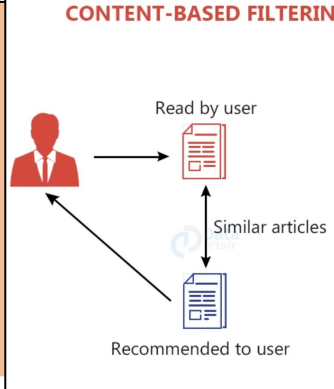
Netflix

Recomendation systems

COLLABORATIVE FILTERING

Read by both users

Similar users

Read by her, recommended to him

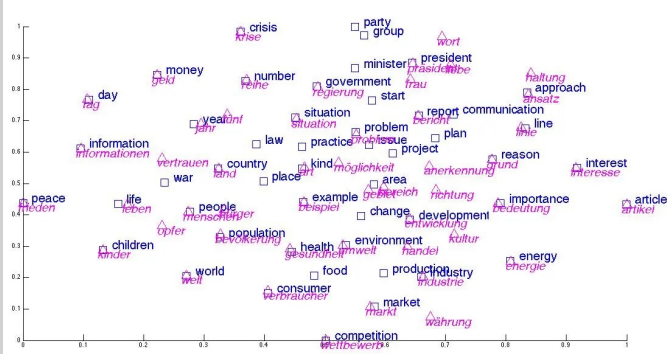CONTENT-BASED FILTERING

Read by user

Similar articles

Recommended to user

Google translate

Mapping between languages

The semantic equivalence of words has been inferred by their context, so similar meanings are co-located.

ChatGPT

Generation of text

Can you write a poem about Einstein Telescope in less than 100 words?

Einstein's vision in cosmic flight,
The telescope emerges, blazing bright.
Arms stretch below, deep underground,
Whispers of gravity, it seeks to sound.
A symphony of secrets, it aims to trace,
Unveiling the universe with grace.
Mysteries of stars and cosmic rhyme,
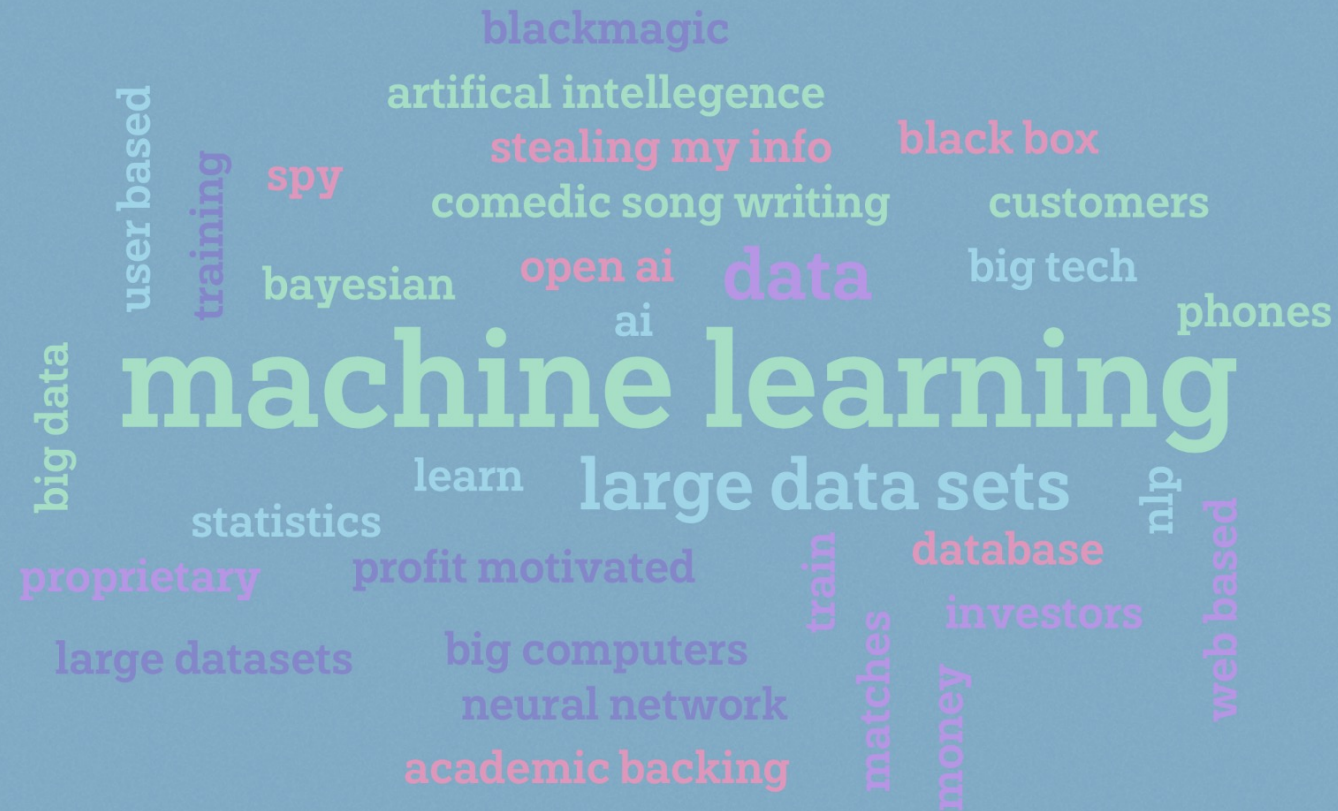The Einstein Telescope, a glimpse through time.

And many more: social media, facial recognition, virtual assistants…

What do Gmail, Netflix, Google translate and ChatGPT have in common?
41 answers

What do Gmail, Netflix, Google translate and ChatGPT have in common?
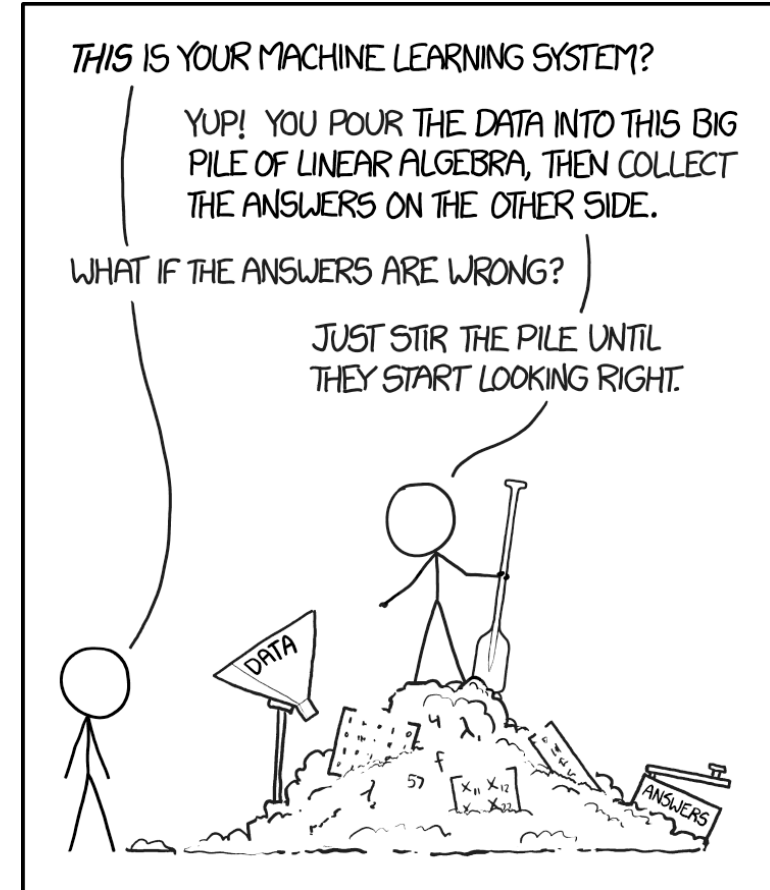

Huge amounts of data

Learnable features.

Data has structure (even if it is not obvious)

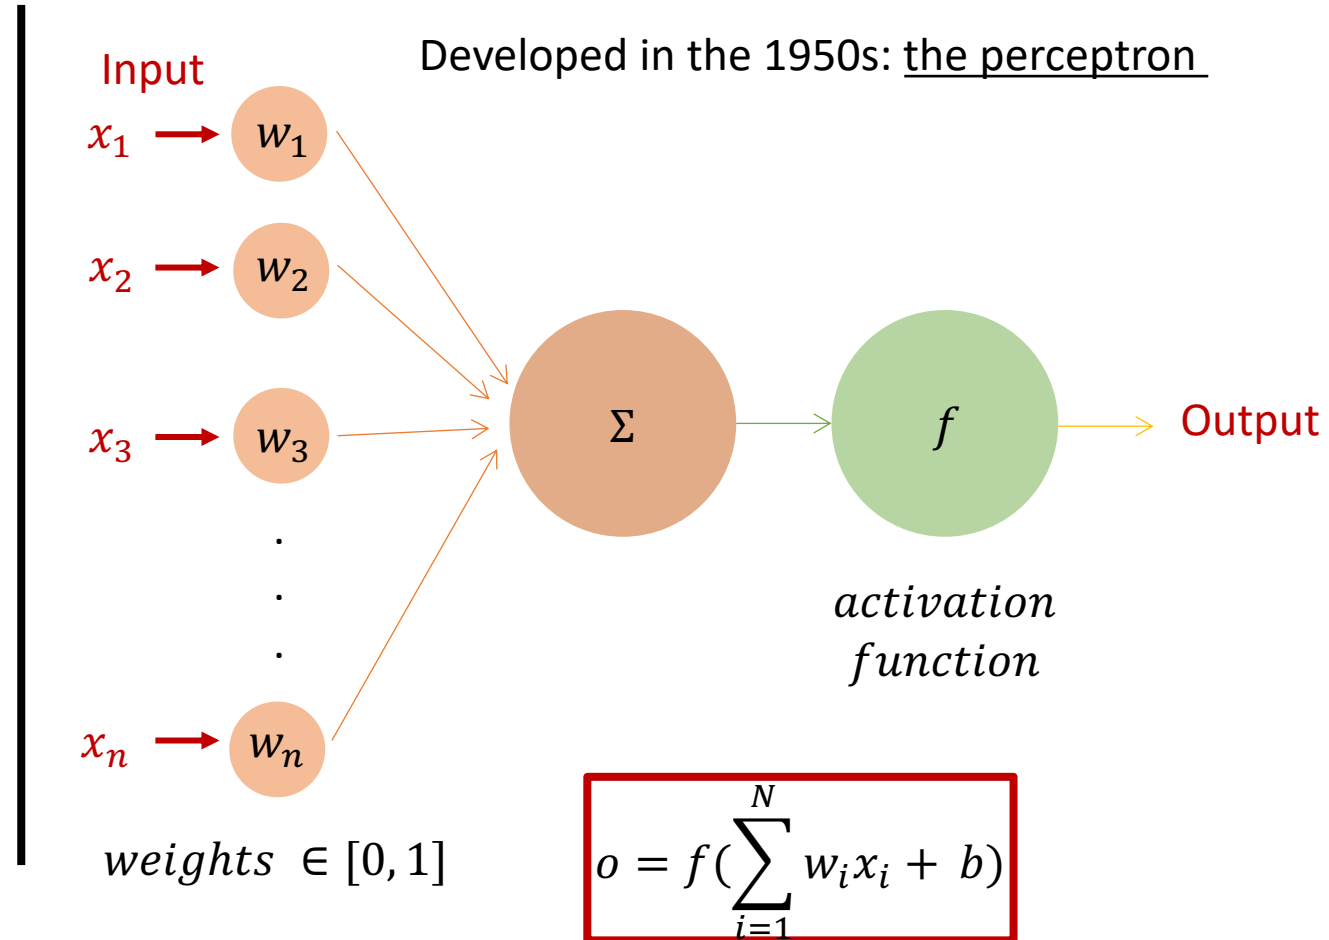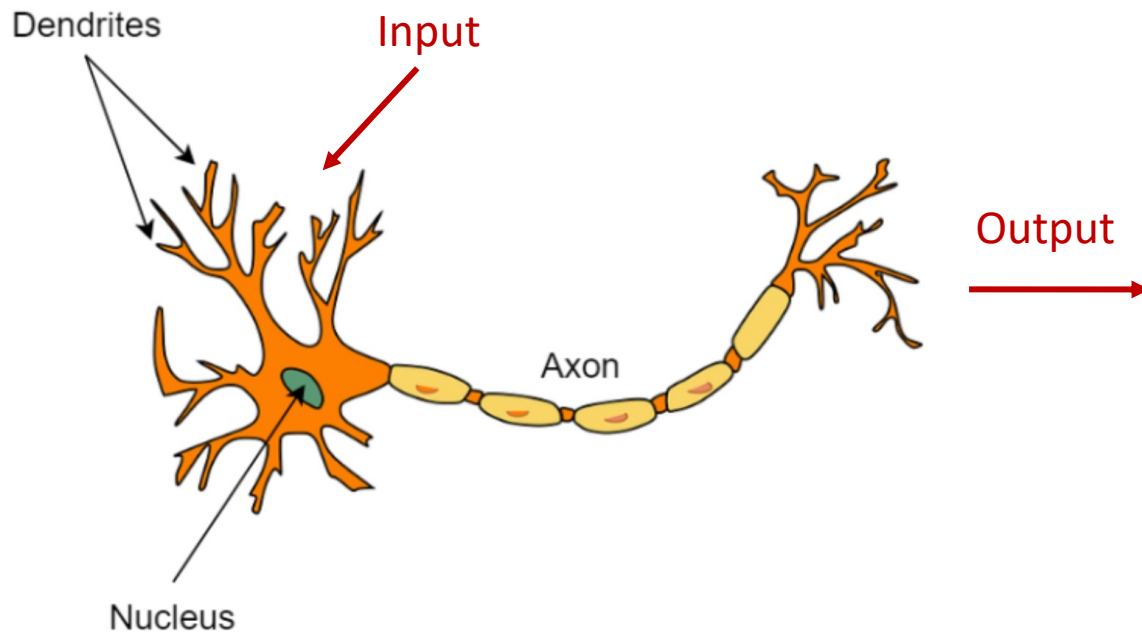Decisions are made based on the data

Get better over time (with more data!)

Humans learn from experience and machines learn from data!

Simplifying the human brain: what is a neuron?

Dendrites

Input

Axon

Nucleus

Output

Input

$x_1 \rightarrow w_1$

$x_2 \rightarrow w_2$

$x_3 \rightarrow w_3$

.
.
.

$x_n \rightarrow w_n$

Developed in the 1950s: <u>the perceptron</u>

Σ

f

$\rightarrow$ Output

activation function

$weights \in [0, 1]$

$$o = f(\sum_{i=1}^{N} w_i x_i + b)$$

OK, but how do neurons learn?

Sant Maurici Lake, Spain

How do neurons learn? Gradient descent!
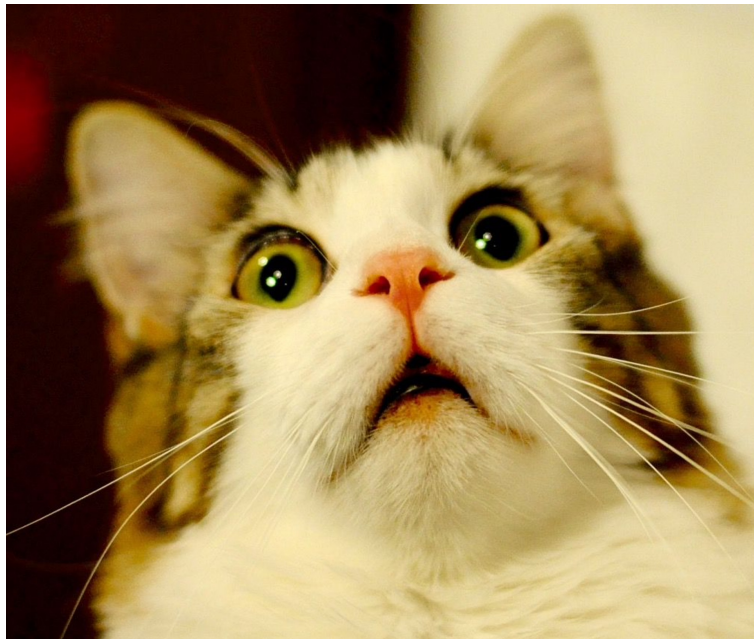
Cost function $J(w, b)$: difference between ground truth y and output of model $h_{w,b}(x)$

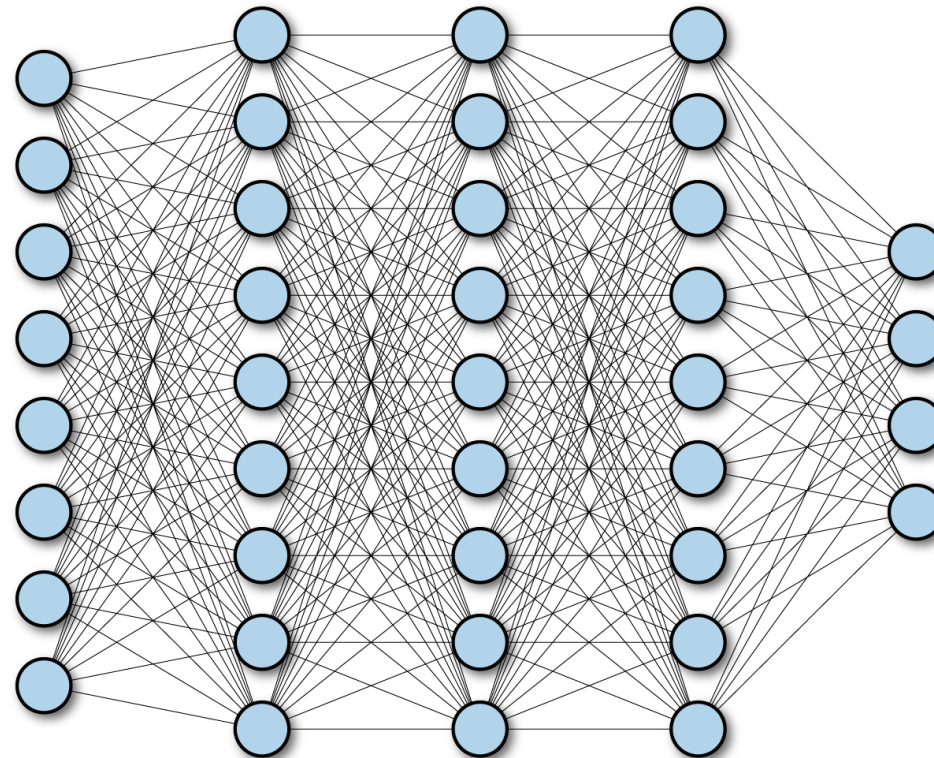$$J(w, b) = \frac{1}{2}\|h_{w,b} - y\|^2$$

A.

Local Minima

Global Minima

Saddle Point
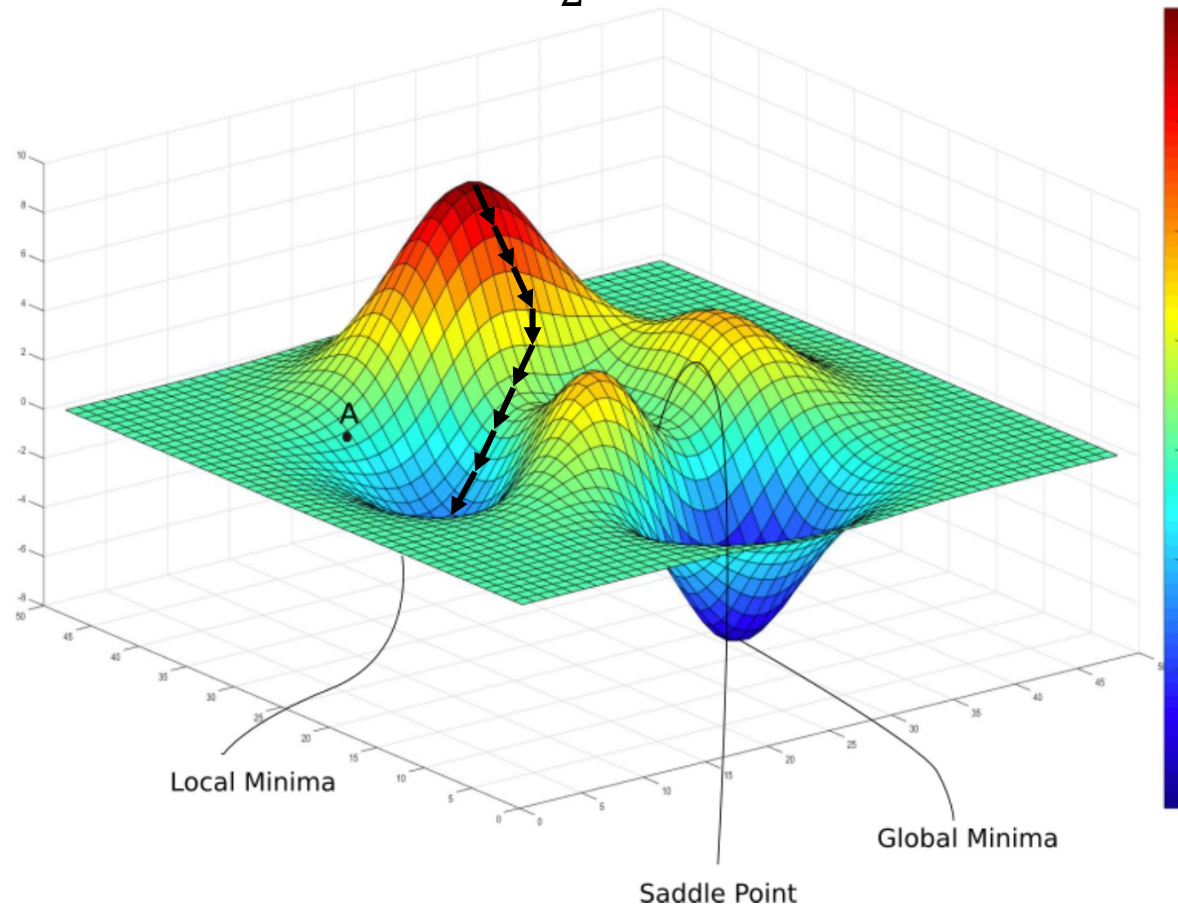
What is the worst that can happen?
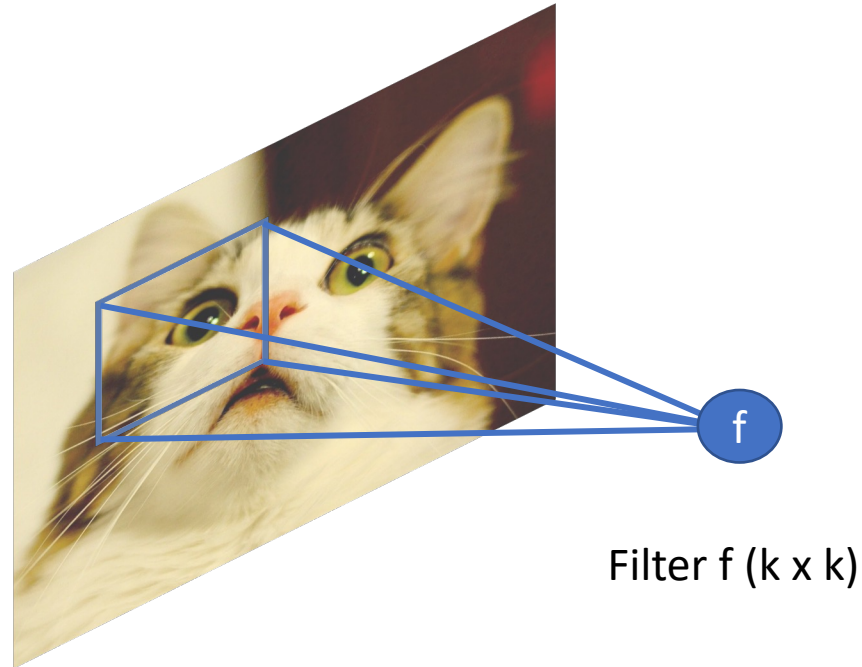
Deep Learning: from a single neuron to a neural network (NN)

32 x 32 x 3 → 3072 x 1

Cat (0.999)

Billions of paramters to train! Maybe not all pixels are related?

Deep Learning: from a single neuron to a neural network (NN)

Filter f (k x k)
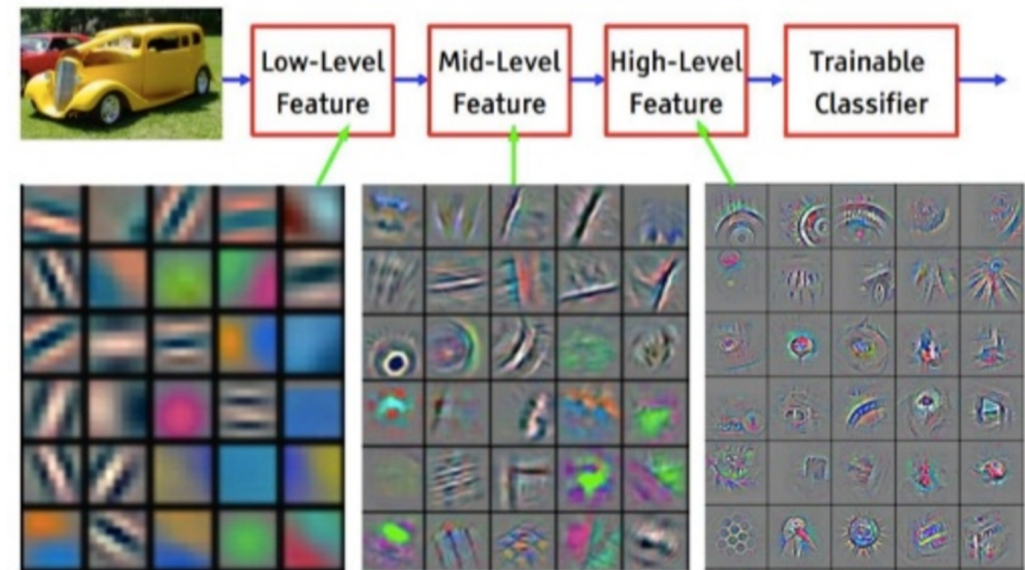
Parameter sharing → learn spatial features

32 x 32 x 3 → we mantain spatial information!

Convolution

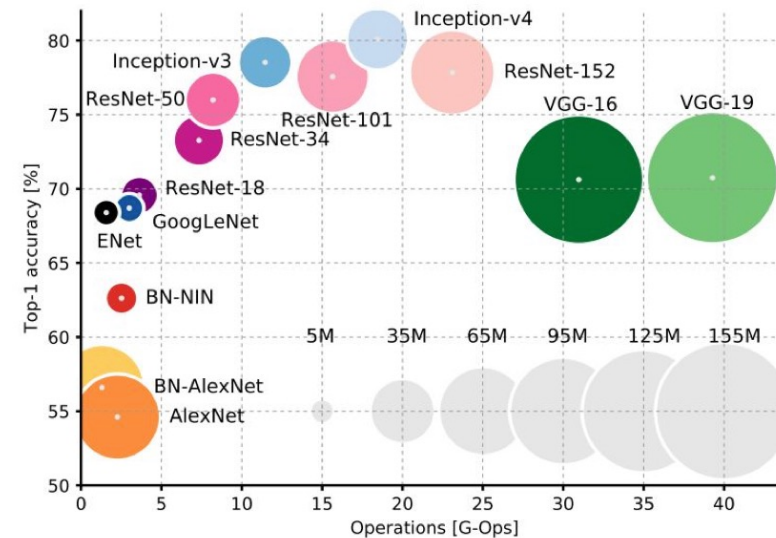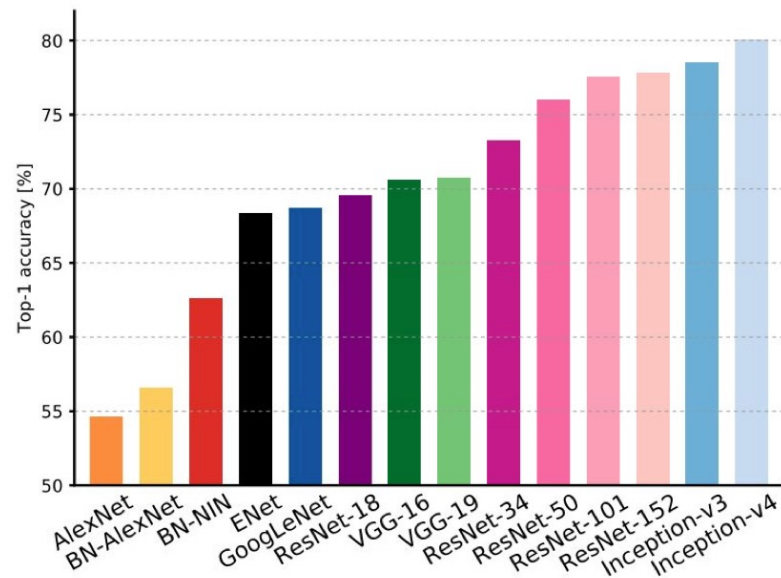Low-Level Feature → Mid-Level Feature → High-Level Feature → Trainable Classifier

Universal Approximation Theorem

Neural networks can approximate any continuous function with arbitrary accuracy given a sufficient number of neurons.

An Analysis of Deep Neural Network Models for Practical Applications, 2017.

Creating a ML model: do's and dont's

Do's	Dont's
Motivate your choices	Apply ML to everything! (we can right?)
Motivate your model, why should it work in your application?	Copy random model from ML guru
Do experiments (and read papers!)	Choose non-learnable parameters of your model at random
Start simple, increase size wisely (finite resources, easier to train)	Create HUGE models (more sugar, more sweet!)

Machine Learning Arxiv Papers per Year

— ML Arxiv Papers   — Moore's Law growth rate (2x/2 years)

~100 new ML papers **every day!**

Chart obtained from LinkedIn, that is likely true (if someone knows the original source, I will update this post to cite it)
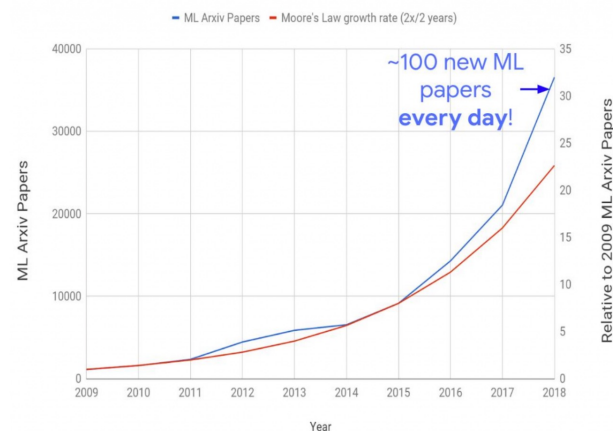
What are big people doing?
Google, Microsfot, Facebook, DeepMind, OpenAI...
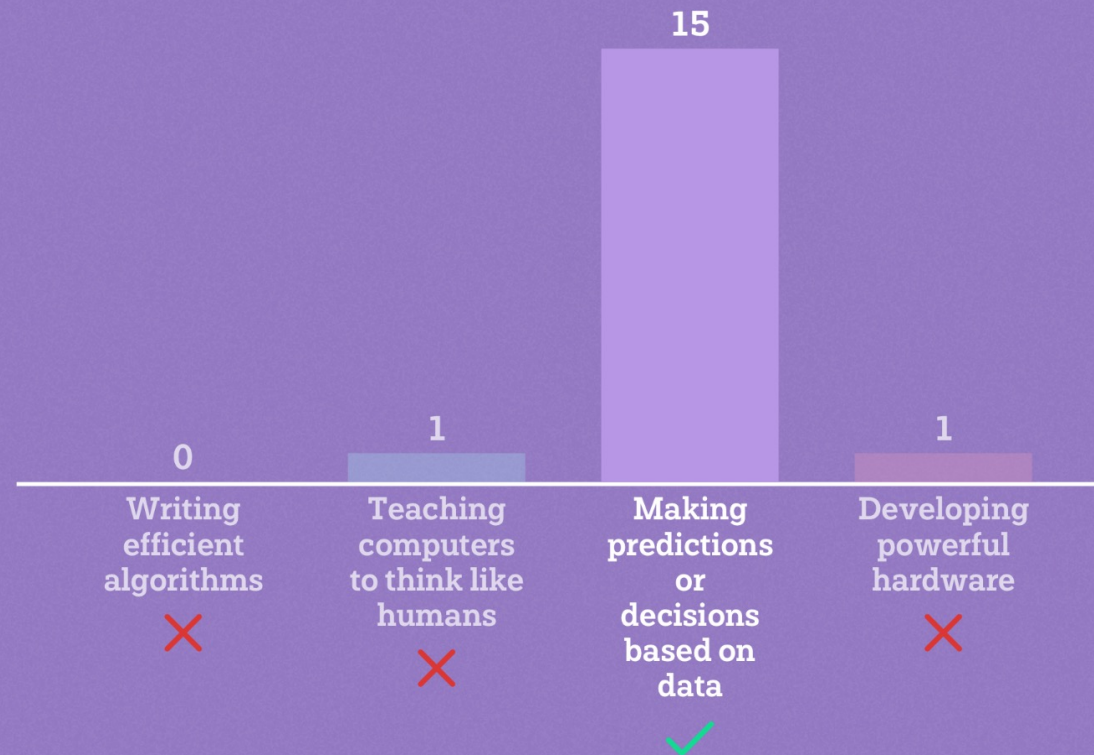
Final test!

I gave you enough data, right? Let's go to Mentimeter!

1. What is the main objective of machine learning?
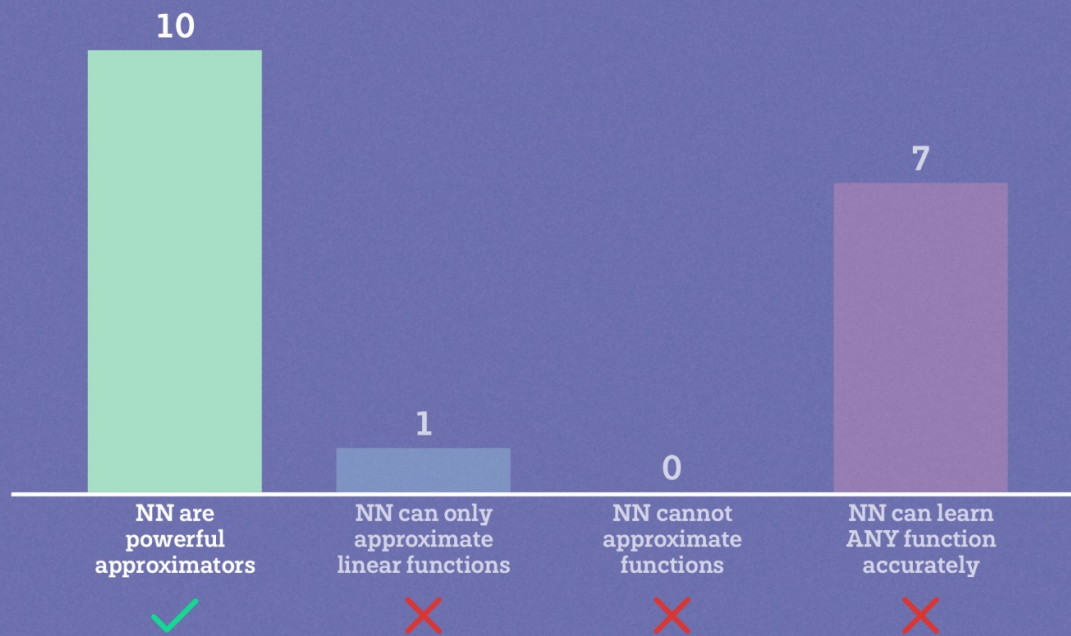

|			
0	1	15	1
Writing efficient algorithms	Teaching computers to think like humans	Making predictions or decisions based on data	Developing powerful hardware
✗	✗	✓	✗

2. What is the algorithm to optimise neurons?

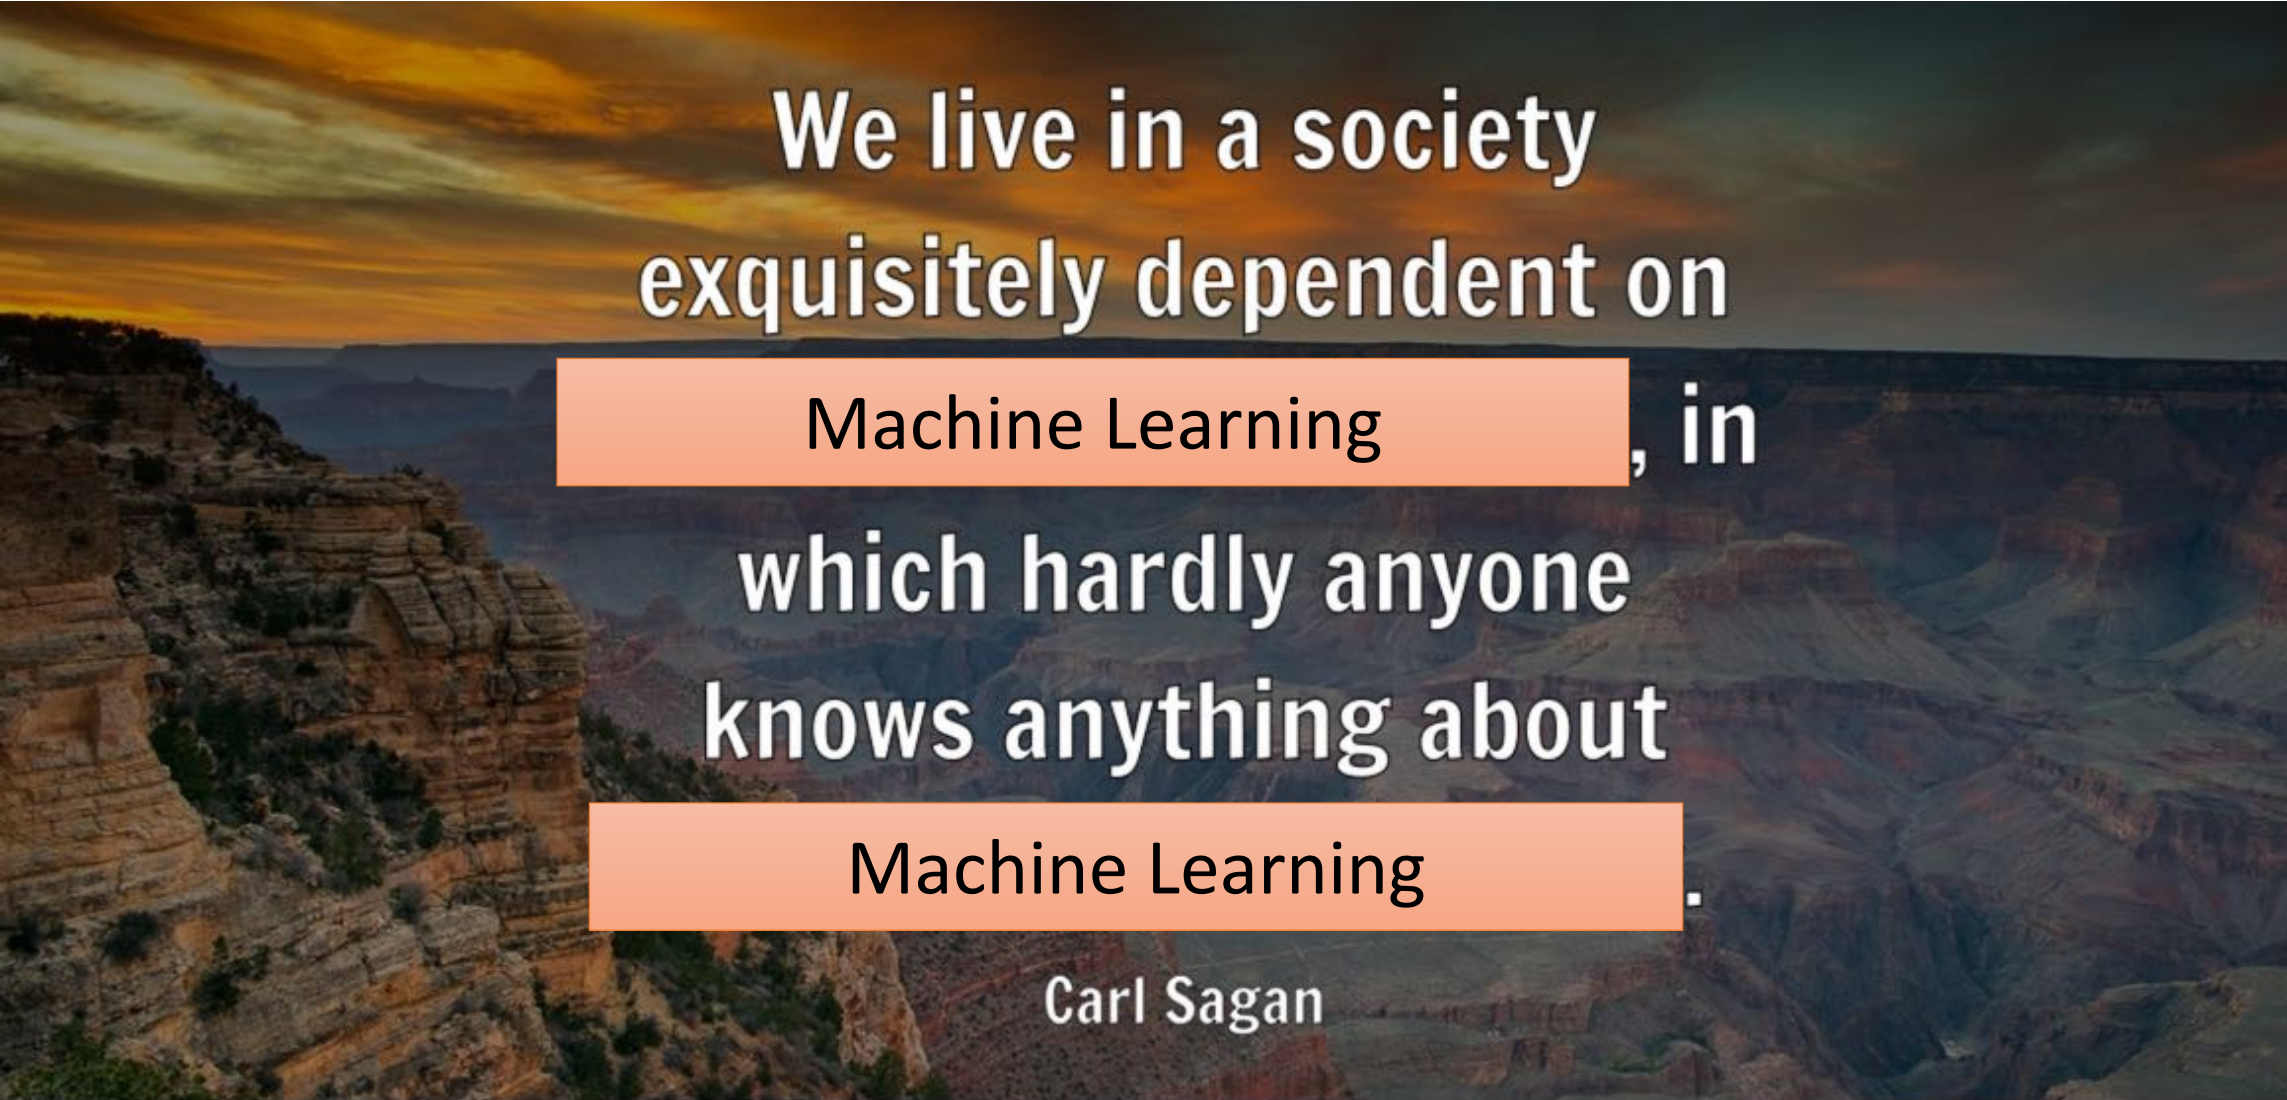| Newton's Method	Decision Trees	Gradient descent	I was sleeping
0	1	15	3
✗	✗	✓	✗
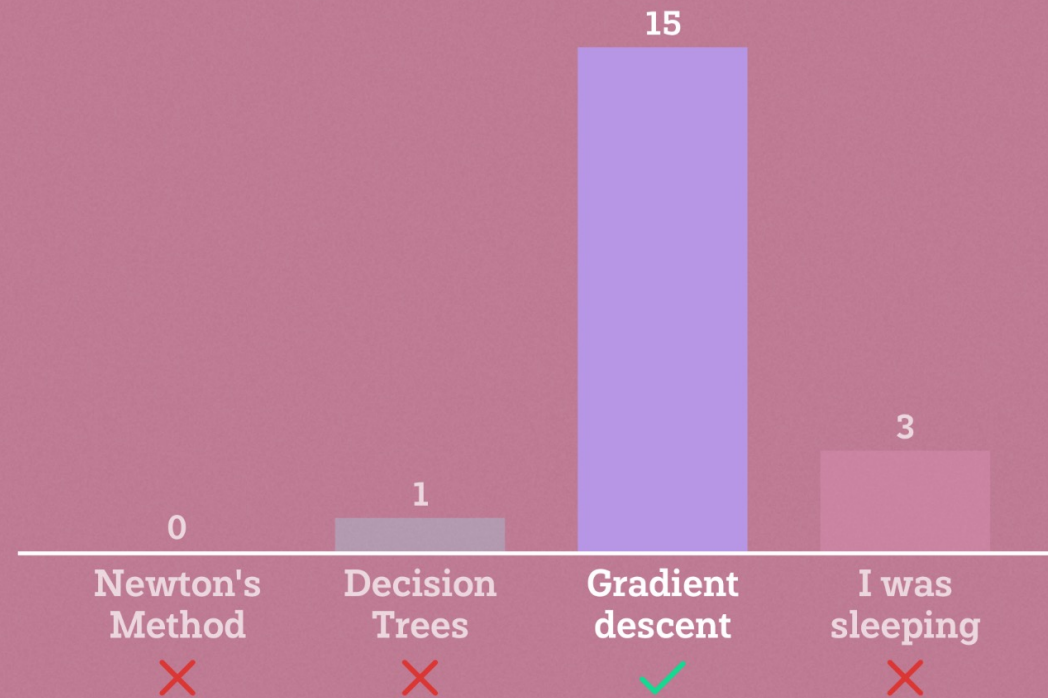
3. What does the universal approximation theorem say?

Bonus track: which question was generated with ChatGPT?
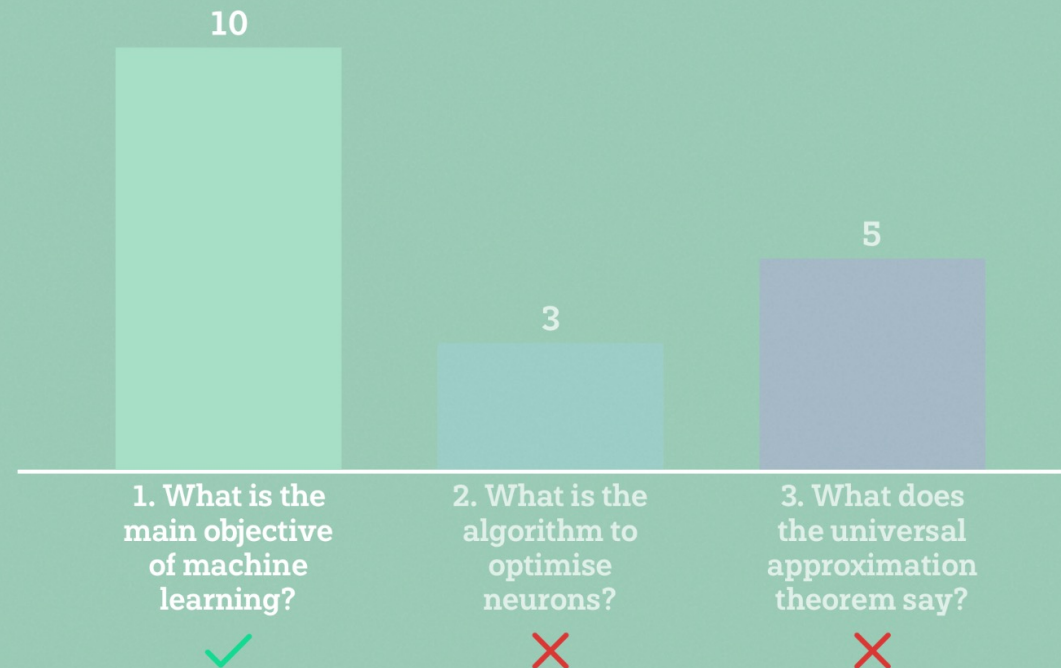
We live in a society
exquisitely dependent on

Machine Learning, in

which hardly anyone
knows anything about

Machine Learning.

Carl Sagan

Thank you for listening! Questions?