Learning from Data

Charting the course of AI

Slides can be found at http://qvirt.com/aitalk.html

Agenda

01 Brief History of AI
Eliza to deep neural networks, Moore’s law, parallel processing

02 Neural Networks
A quick math-free tour through neural network architectures

03 Data and what to do with it
Using data for AI, labelled data, crowdsourcing data

04 AI Futures and Ethics
What’s on it’s way and how do we control it?

05 Some Fun Examples
Some links and resources to play with and learn more

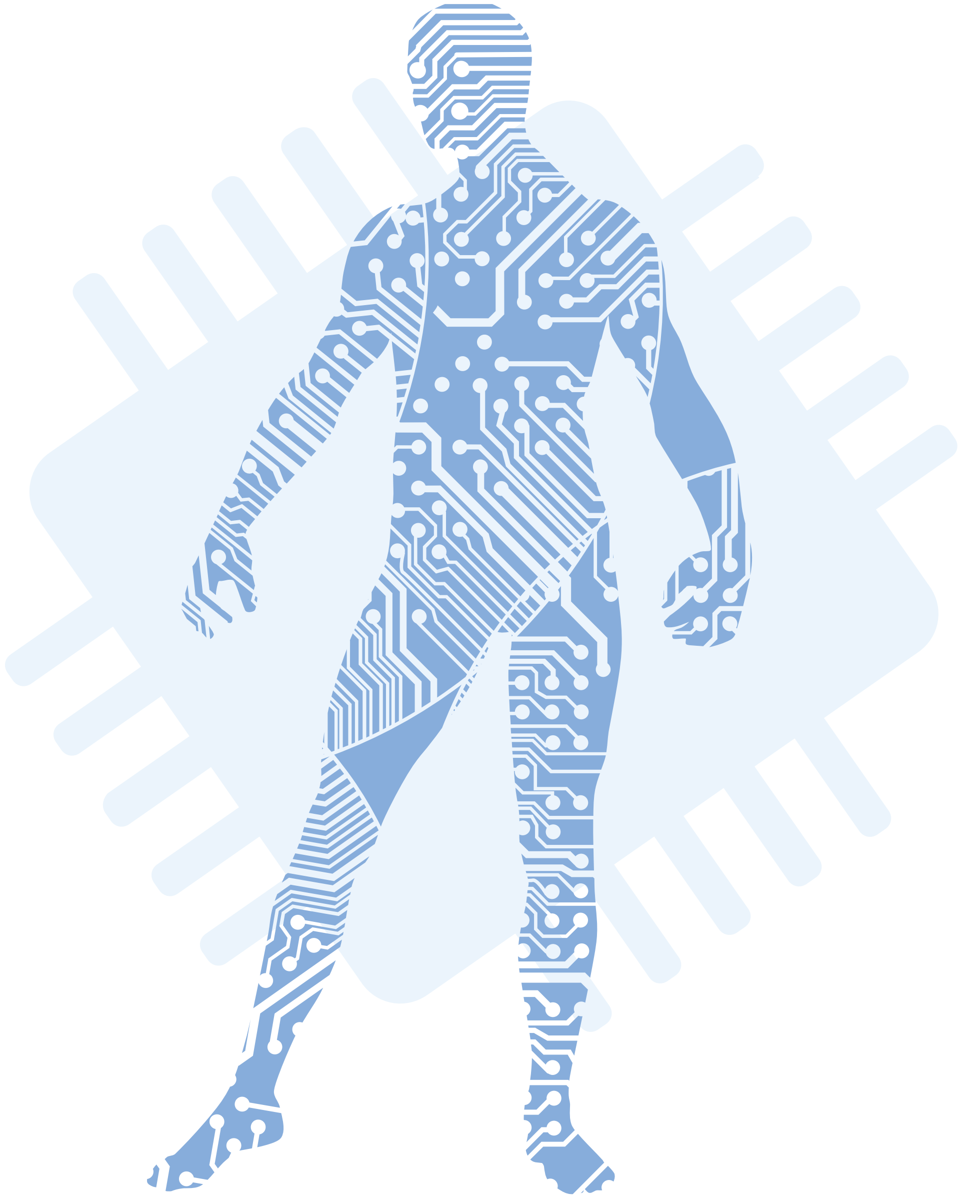

My Background

- Undergraduate degree in cognitive psychology
- Developed statistical software at MD Anderson Cancer Centre
- Designed consumer software at Compaq
- Usability then program manager then researcher at Microsoft
  - Windows team initially then Microsoft Research
- Master’s degree in meeting annotation systems
- PhD in electroencephalography data analytics
- Since worked in VR, geospatial computing, speech and dialogue systems

A quick example!

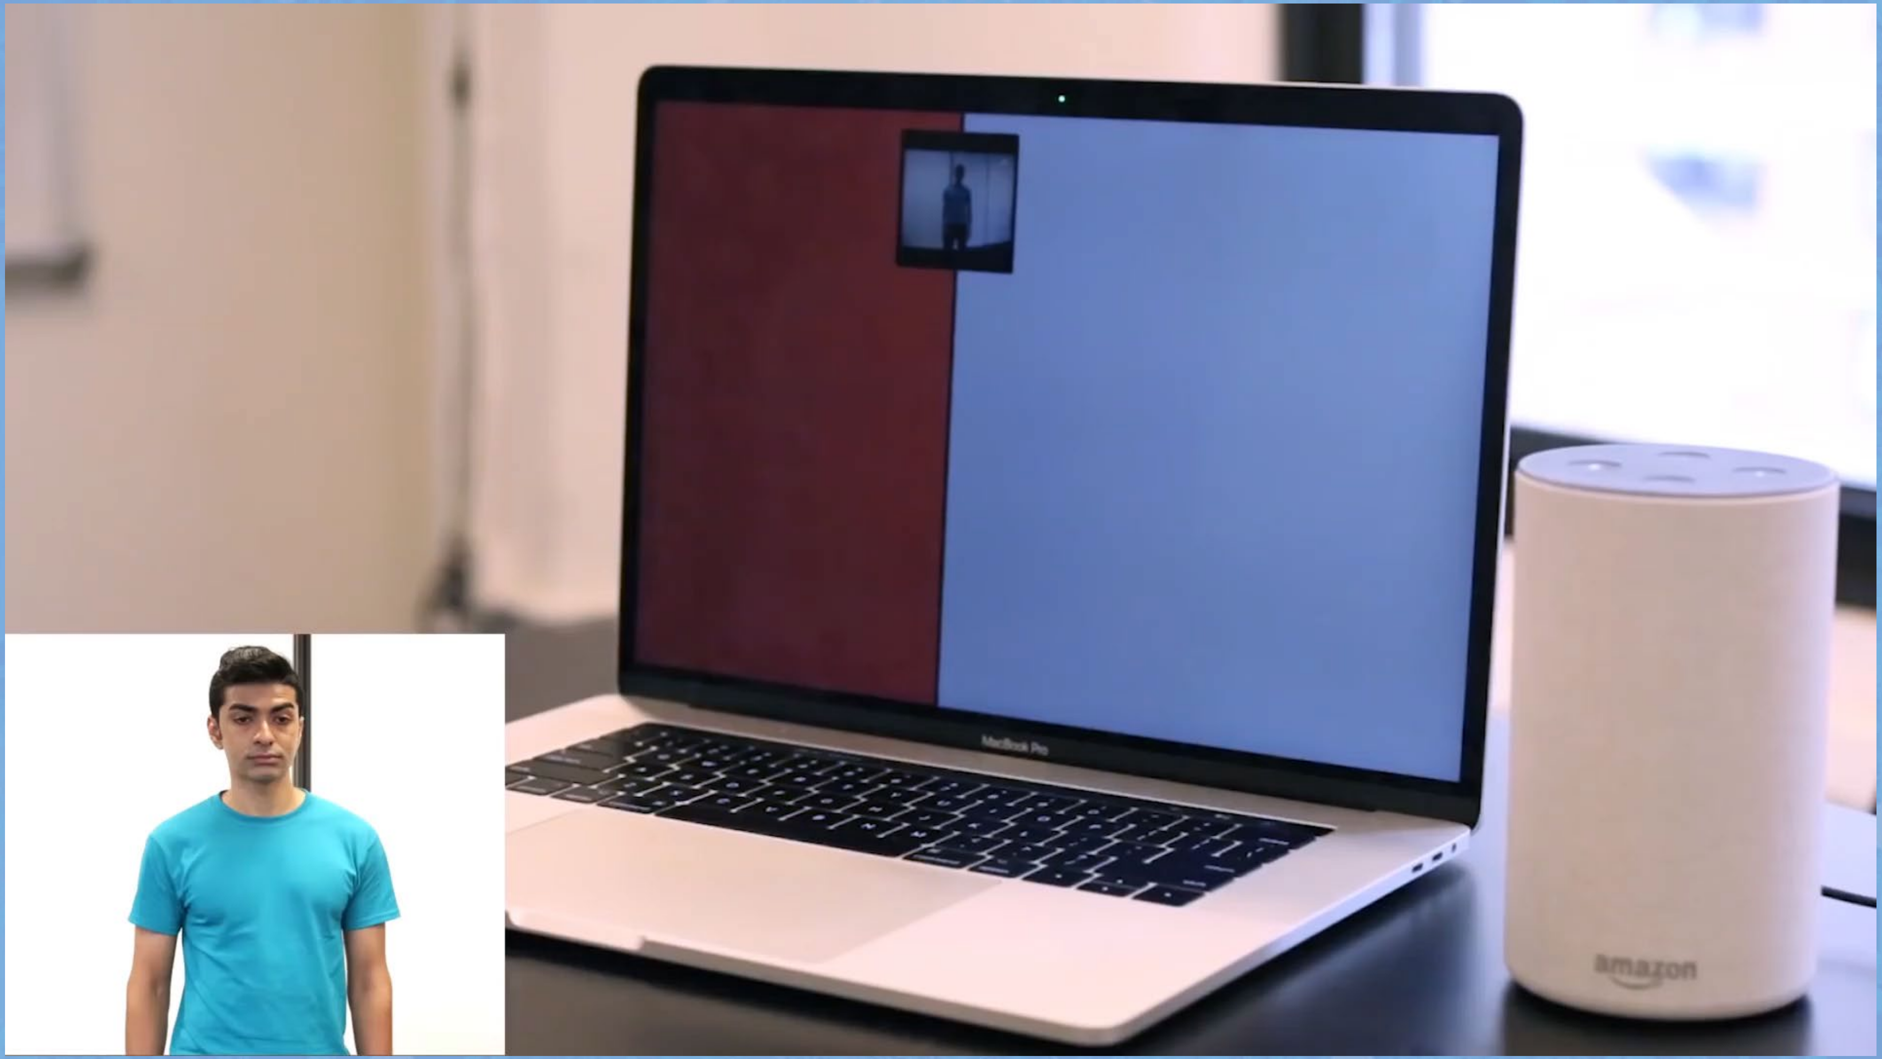

From https://medium.com/tensorflow/getting-alexa-to-respond-to-sign-language-using-your-webcam-and-tensorflow-js-735ccc1e6d3f

Chihuahua or muffin?

From https://medium.freecodecamp.org/chihuahua-or-muffin-my-search-for-the-best-computer-vision-api-cbda4d6b425d

Definitions

- AI – 1956 Dartmouth workshop definition “Thinking Machines”
- Three primary goals of AI
  - Systems that work like the brain
  - Systems that just work, without caring how
  - Systems that use the brain as an abstract concept
- Third goal is the most common in modern systems and what I’ll spend most time on today

History – Cognitive Science

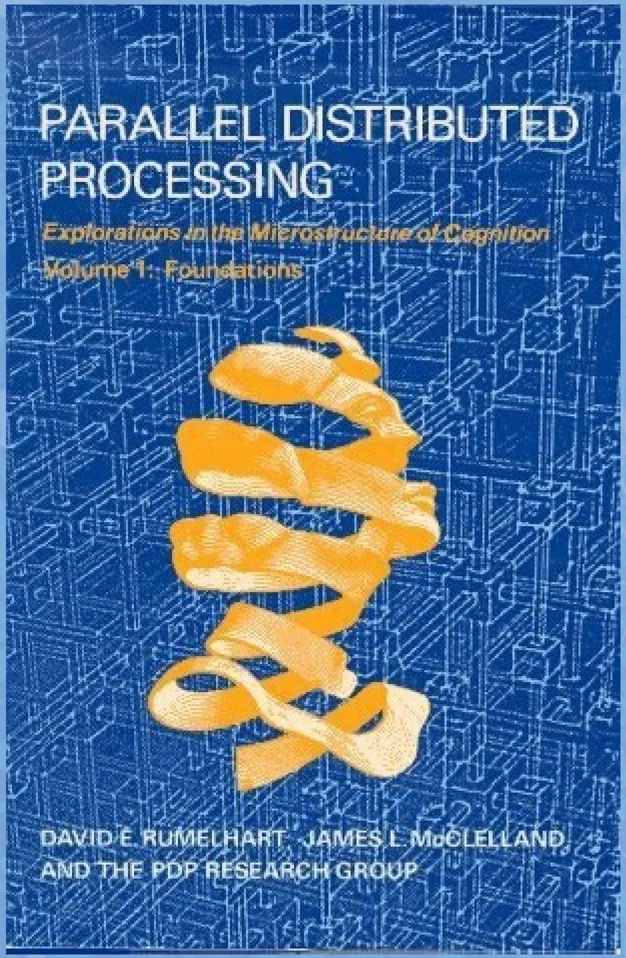

- Field originated with great optimism
  - Hebb (1949), Turing (1936,1950)
  - "Cells that fire together wire together.“ is still useful!
  - Complexity rapidly became overwhelming, though

History – Expert Systems

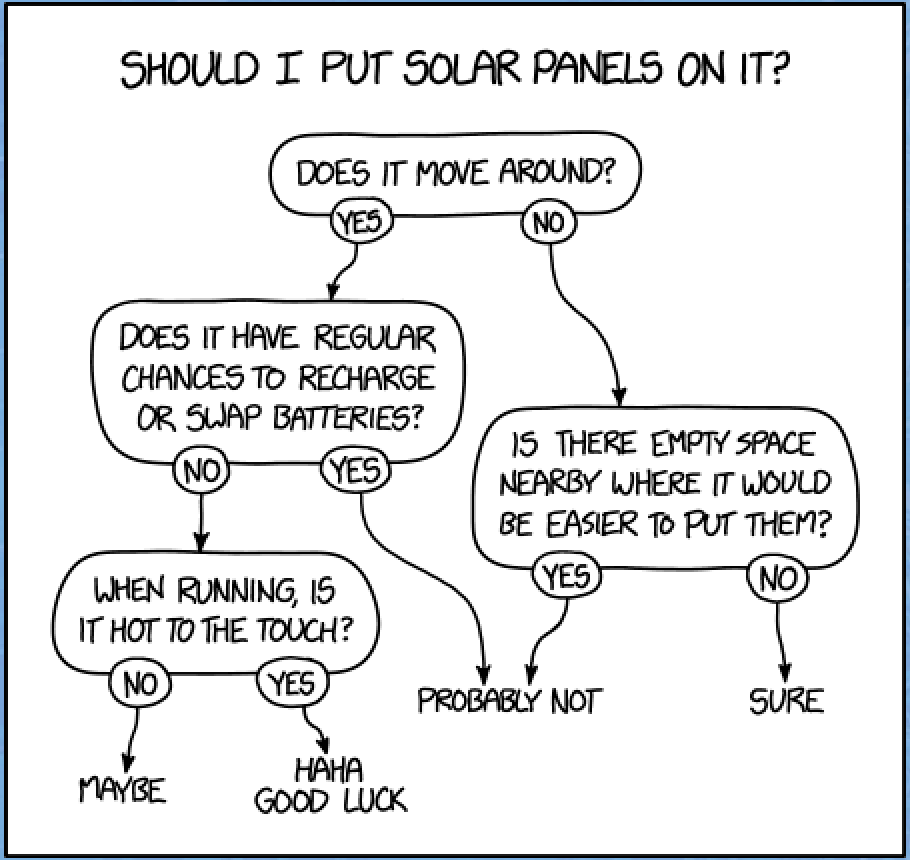

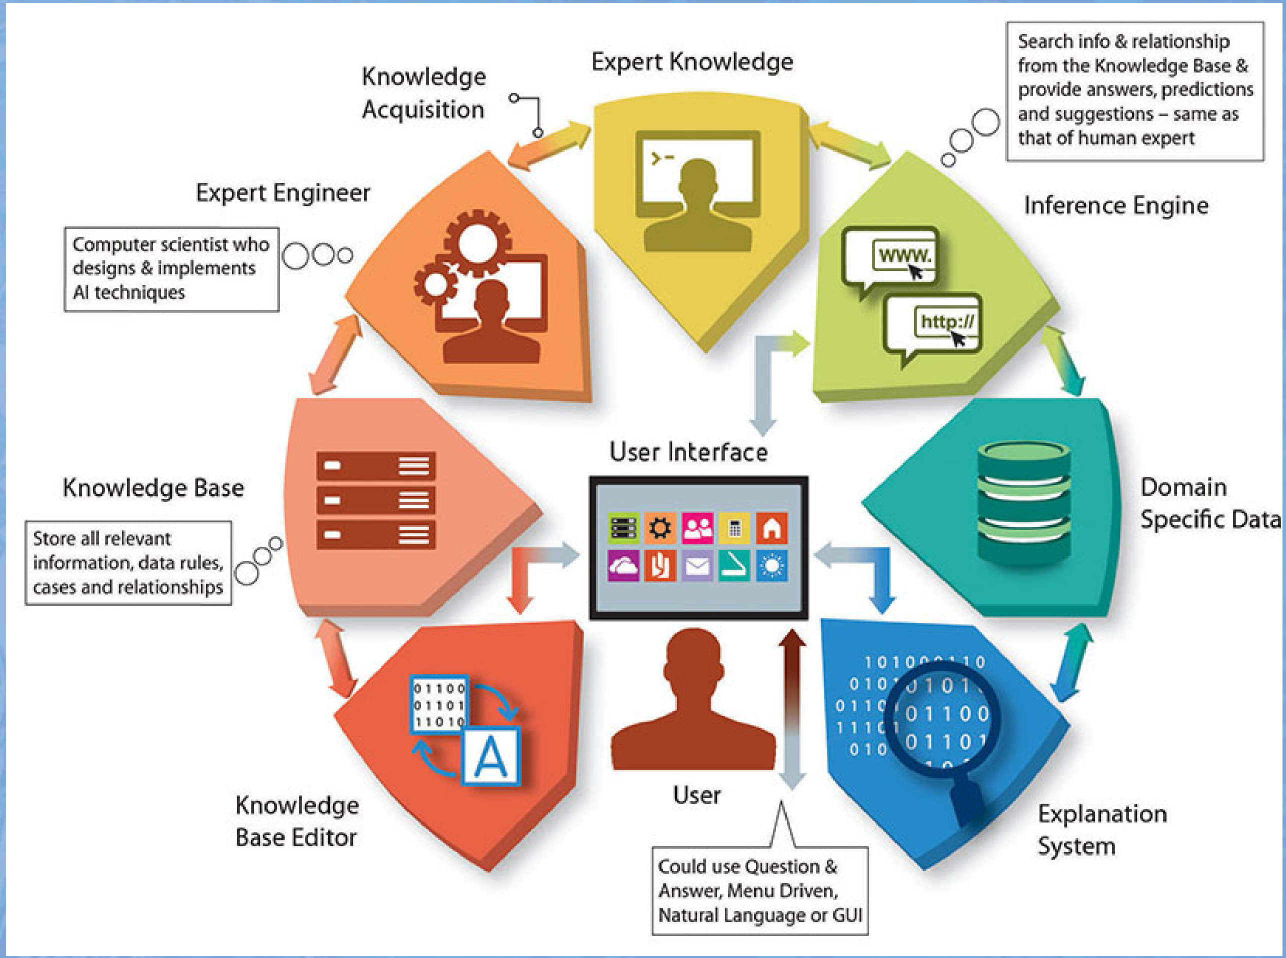

History – Neurally inspired

- Overall, AI has drawn more from cognitive science than the reverse
- Neural network AI is only loosely related to biological
- Many concepts are useful, though
  - Attention
  - Episodic memory
  - Working memory
  - Reinforcement learning

For a nice review, see https://deepmind.com/documents/113/Neuron.pdf

History – Recent Advances

- Starting in 2009, and accelerating through the next 10 years, neural networks and “Deep Learning” have taken off
- Huge theoretical advances have been made, of course, but the main factor is computational
- Gaming, and powerful parallel computation engines in graphics cards, are at the root of this
- To understand why, let’s look at neural network algorithms

02

Neural Network basics and architectures

Basic Algorithms

- Remember Hebbian network – neurons that fire together wire together – associative learning
- Consider the yellow nodes as inputs, the green nodes as a hidden layer, and the orange as outputs

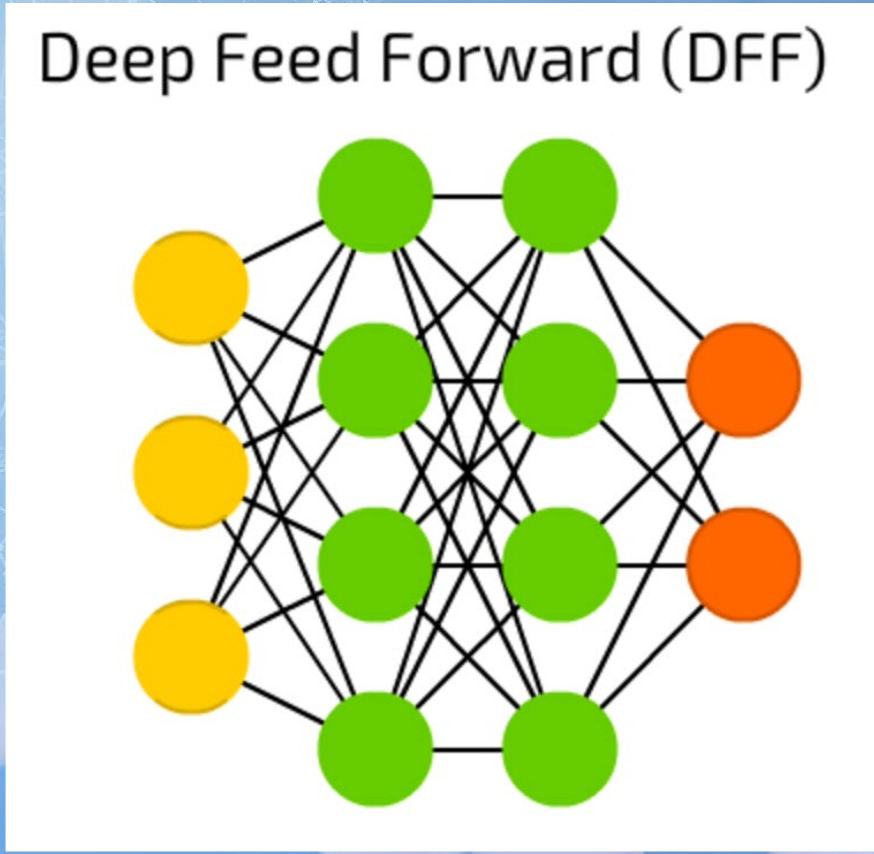

Gradient descent

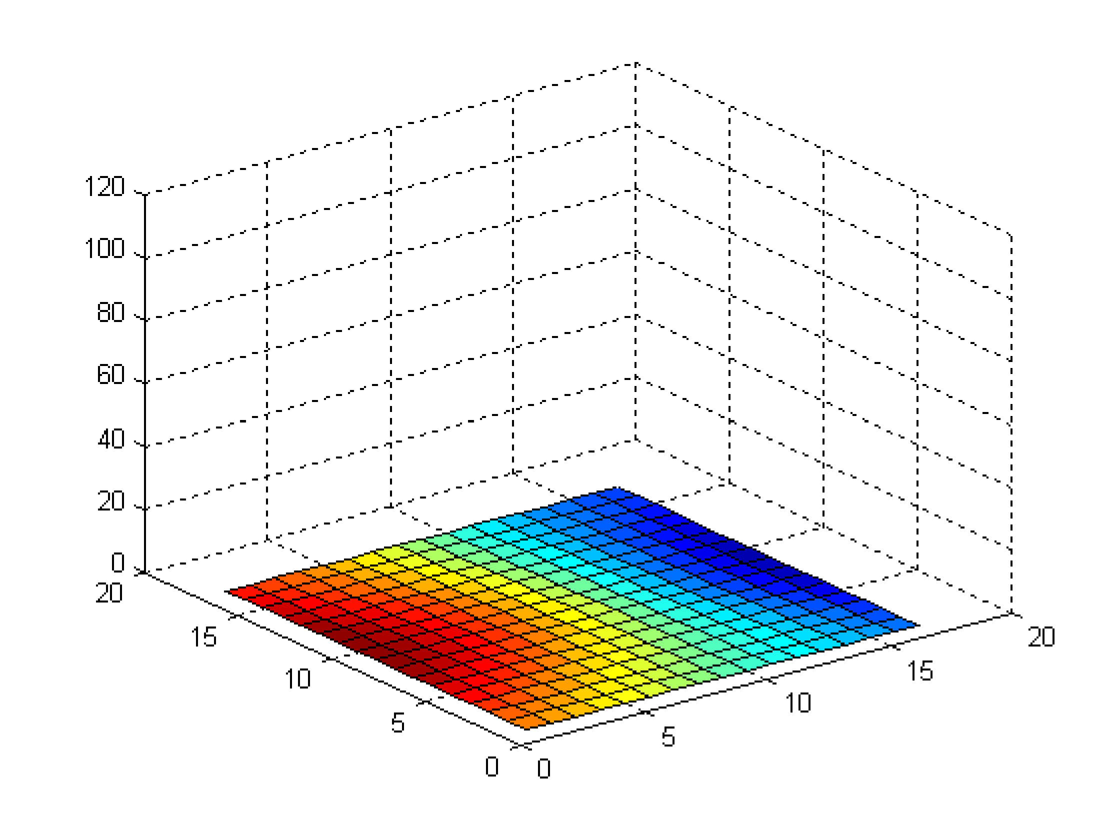

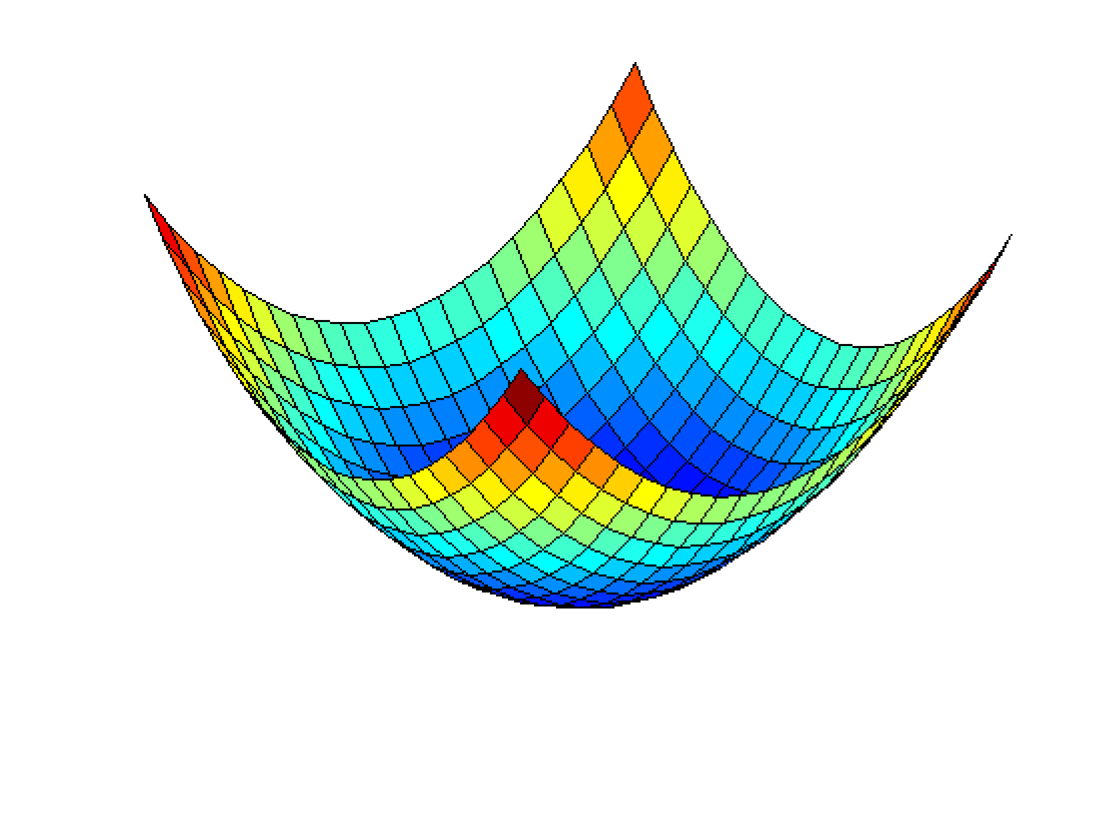

Backpropagation

- Errors are taken from the output node and propagated backwards
- All hidden layers and weights are updated in turn
- Ignore the math on this for now, but think about how iterations of this could train a network

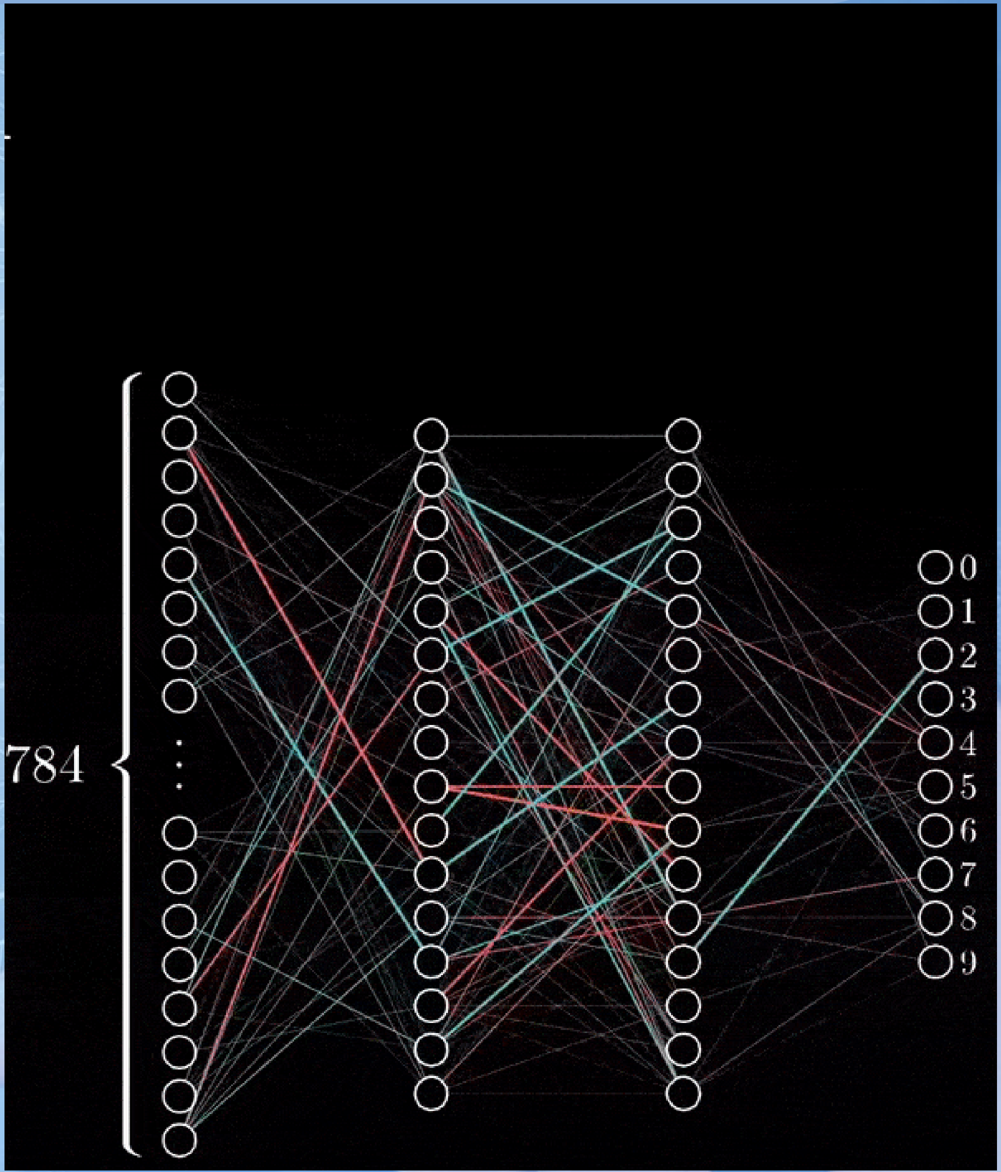

Backpropagation

- Each node and edge in the neural network graph has a weight, each weight is updated with each training image

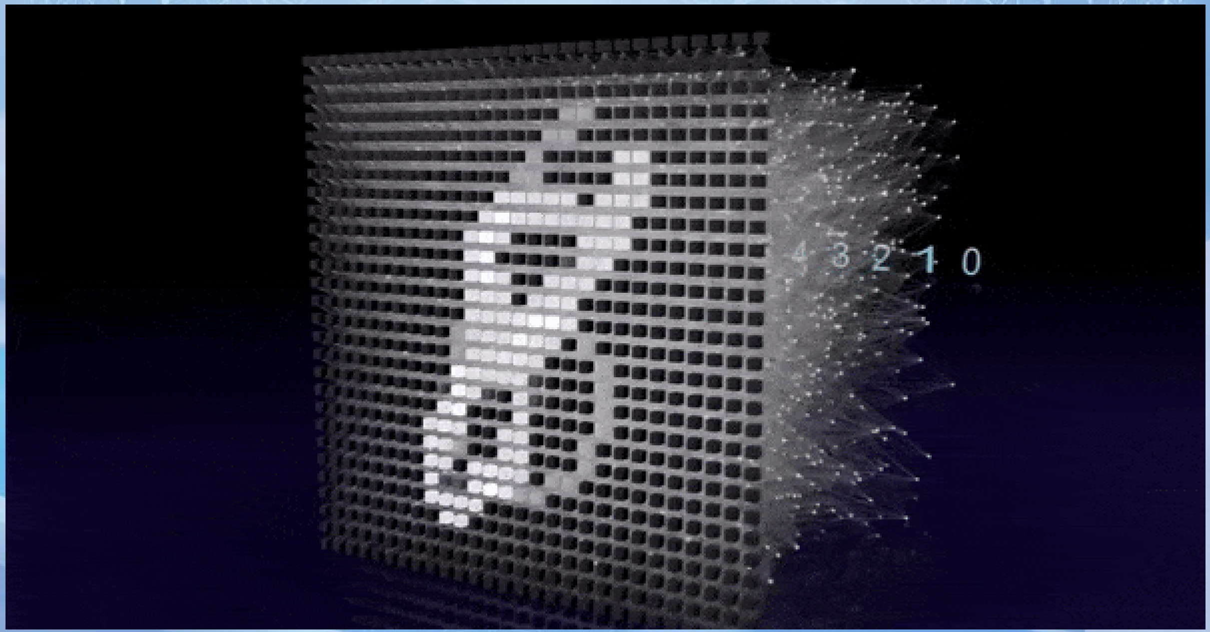

Backpropagation

CNN – convolutional NN

- A small filter is applied over each region in between layers

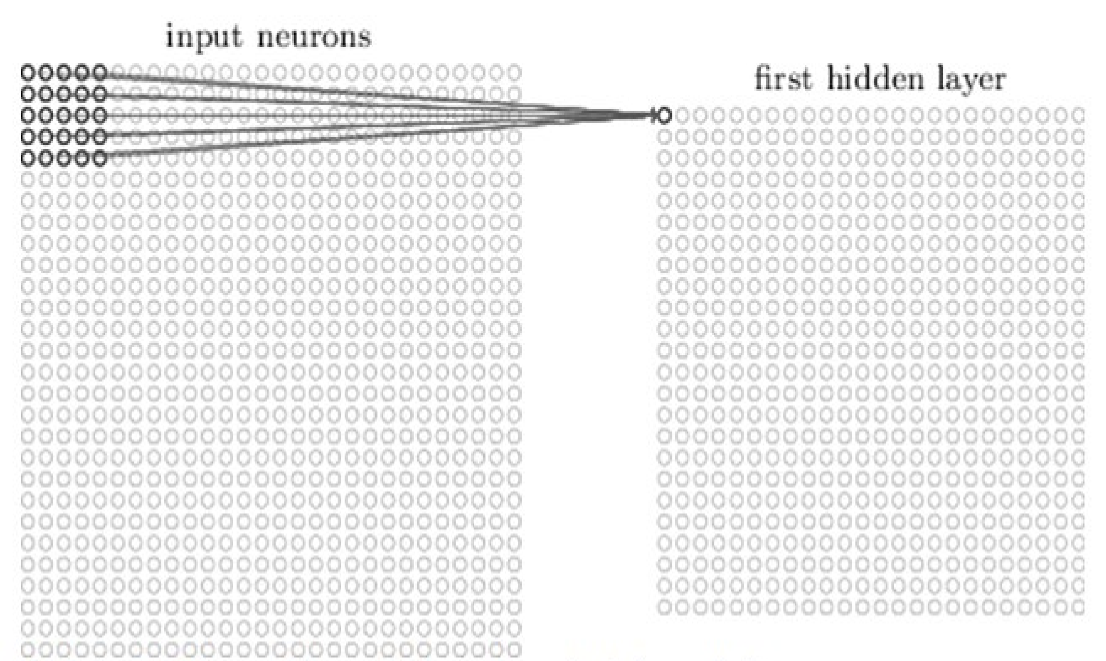

Visualization of 5 x 5 filter convolving around an input volume and producing an activation map

https://adeshpande3.github.io/A-Beginner%27s-Guide-To-Understanding-Convolutional-Neural-Networks/
Also http://playground.tensorflow.org/

What do the hidden layers mean?

- Not always clear WHY a neural net makes a decision
- Research is underway to improve “attribution”

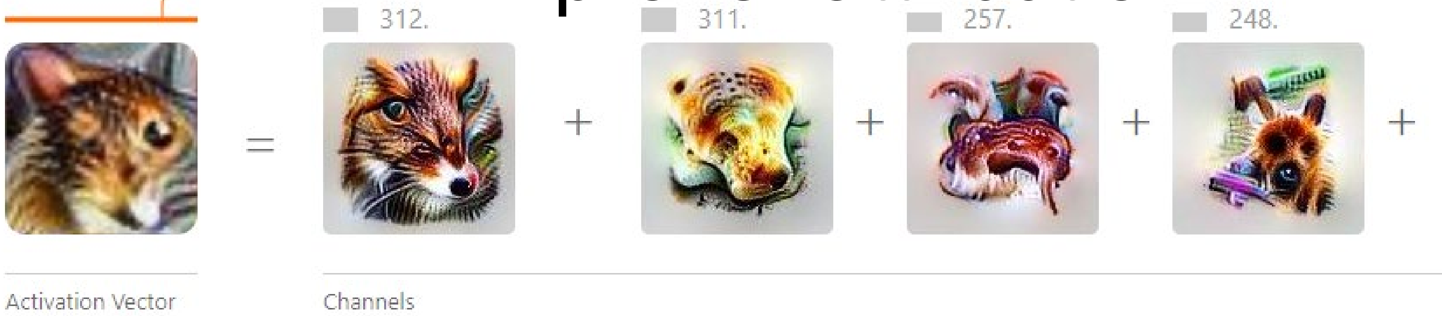

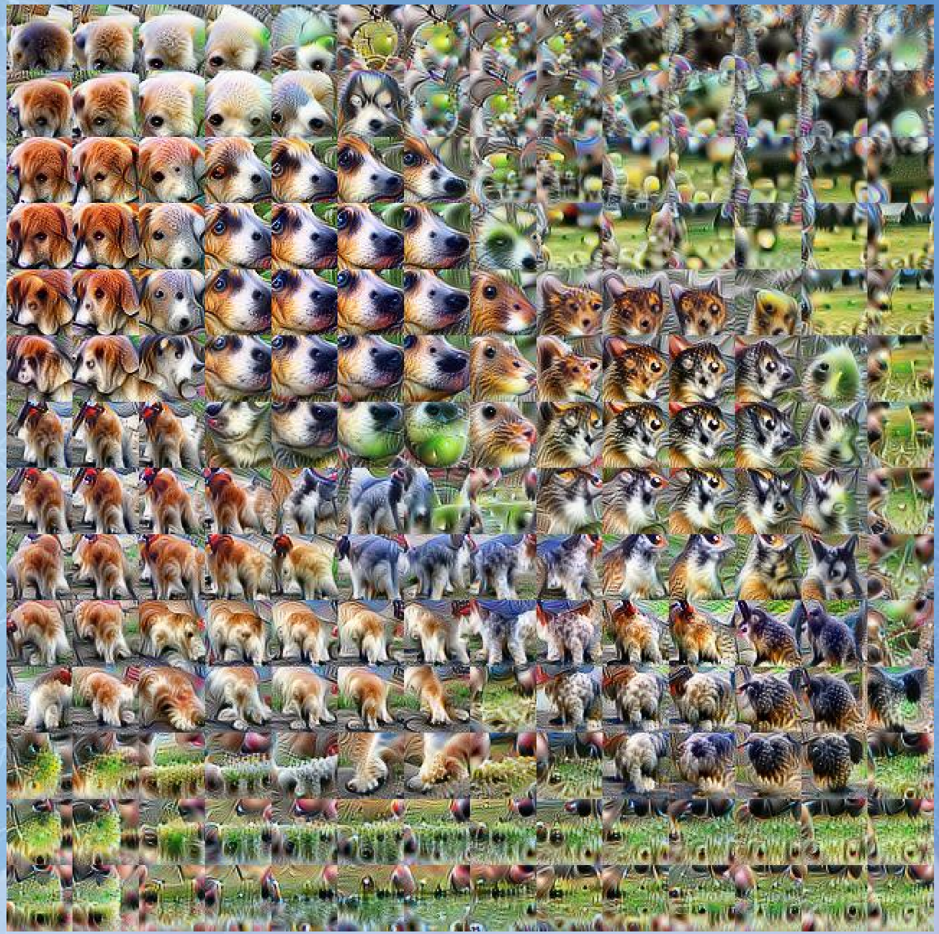

Great detailed explanation of hidden layers https://distill.pub/2018/building-blocks/

Recurrent Neural Networks

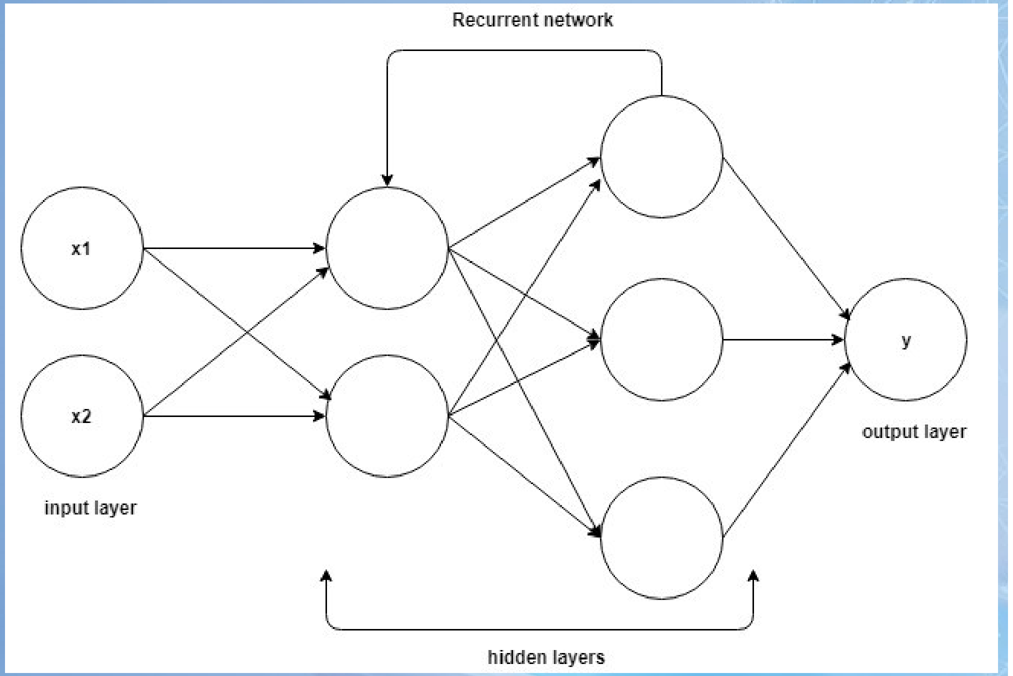

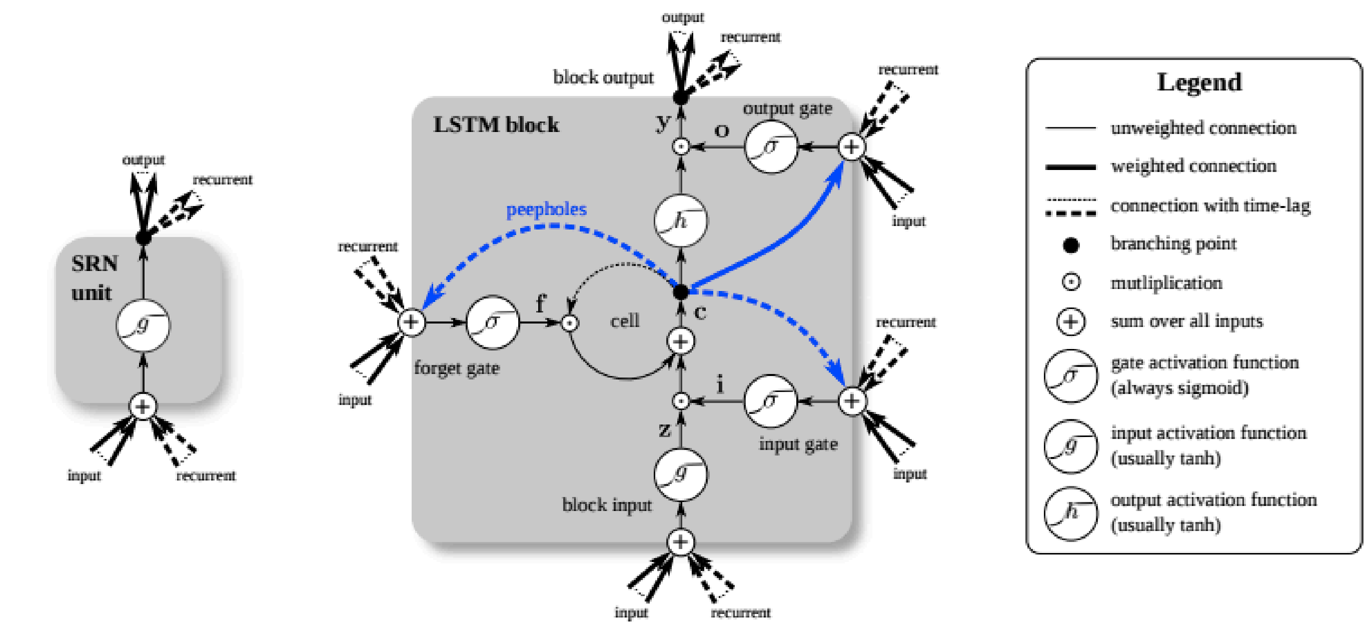

Figure 1. Detailed schematic of the Simple Recurrent Network (SRN) unit (left) and a Long Short-Term Memory block (right) as used in the hidden layers of a recurrent neural network.

Chris Olah has a nice blog entry on this http://colah.github.io/posts/2015-08-Understanding-LSTMs/

Generative Adversarial Networks

Generative Adversarial Network (GAN)

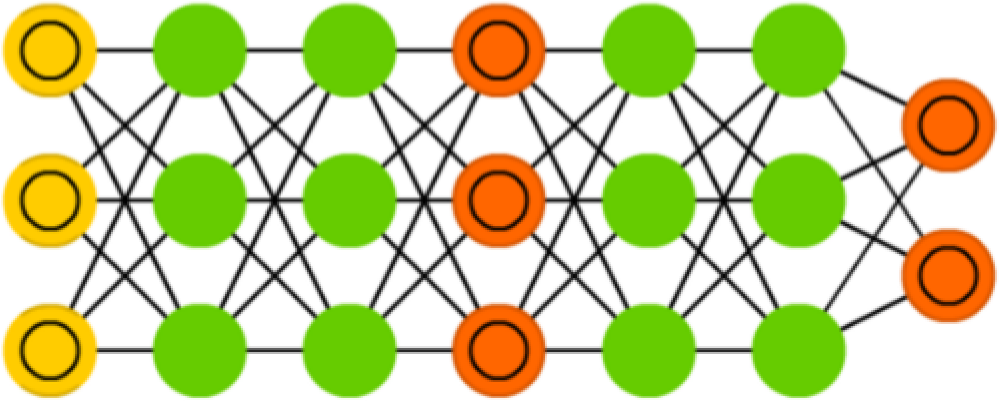

GAN represents a huge family of double networks, that are composed from generator and discriminator. They constantly try to fool each other—generator tries to generate some data, and discriminator, receiving sample data, tries to tell generated data from samples. Constantly evolving, this type of neural networks can generate real-life images, in case you are able to maintain the training balance between these two networks.

Good intro GAN article https://skymind.ai/wiki/generative-adversarial-network-gan

Nvidia Style Generator

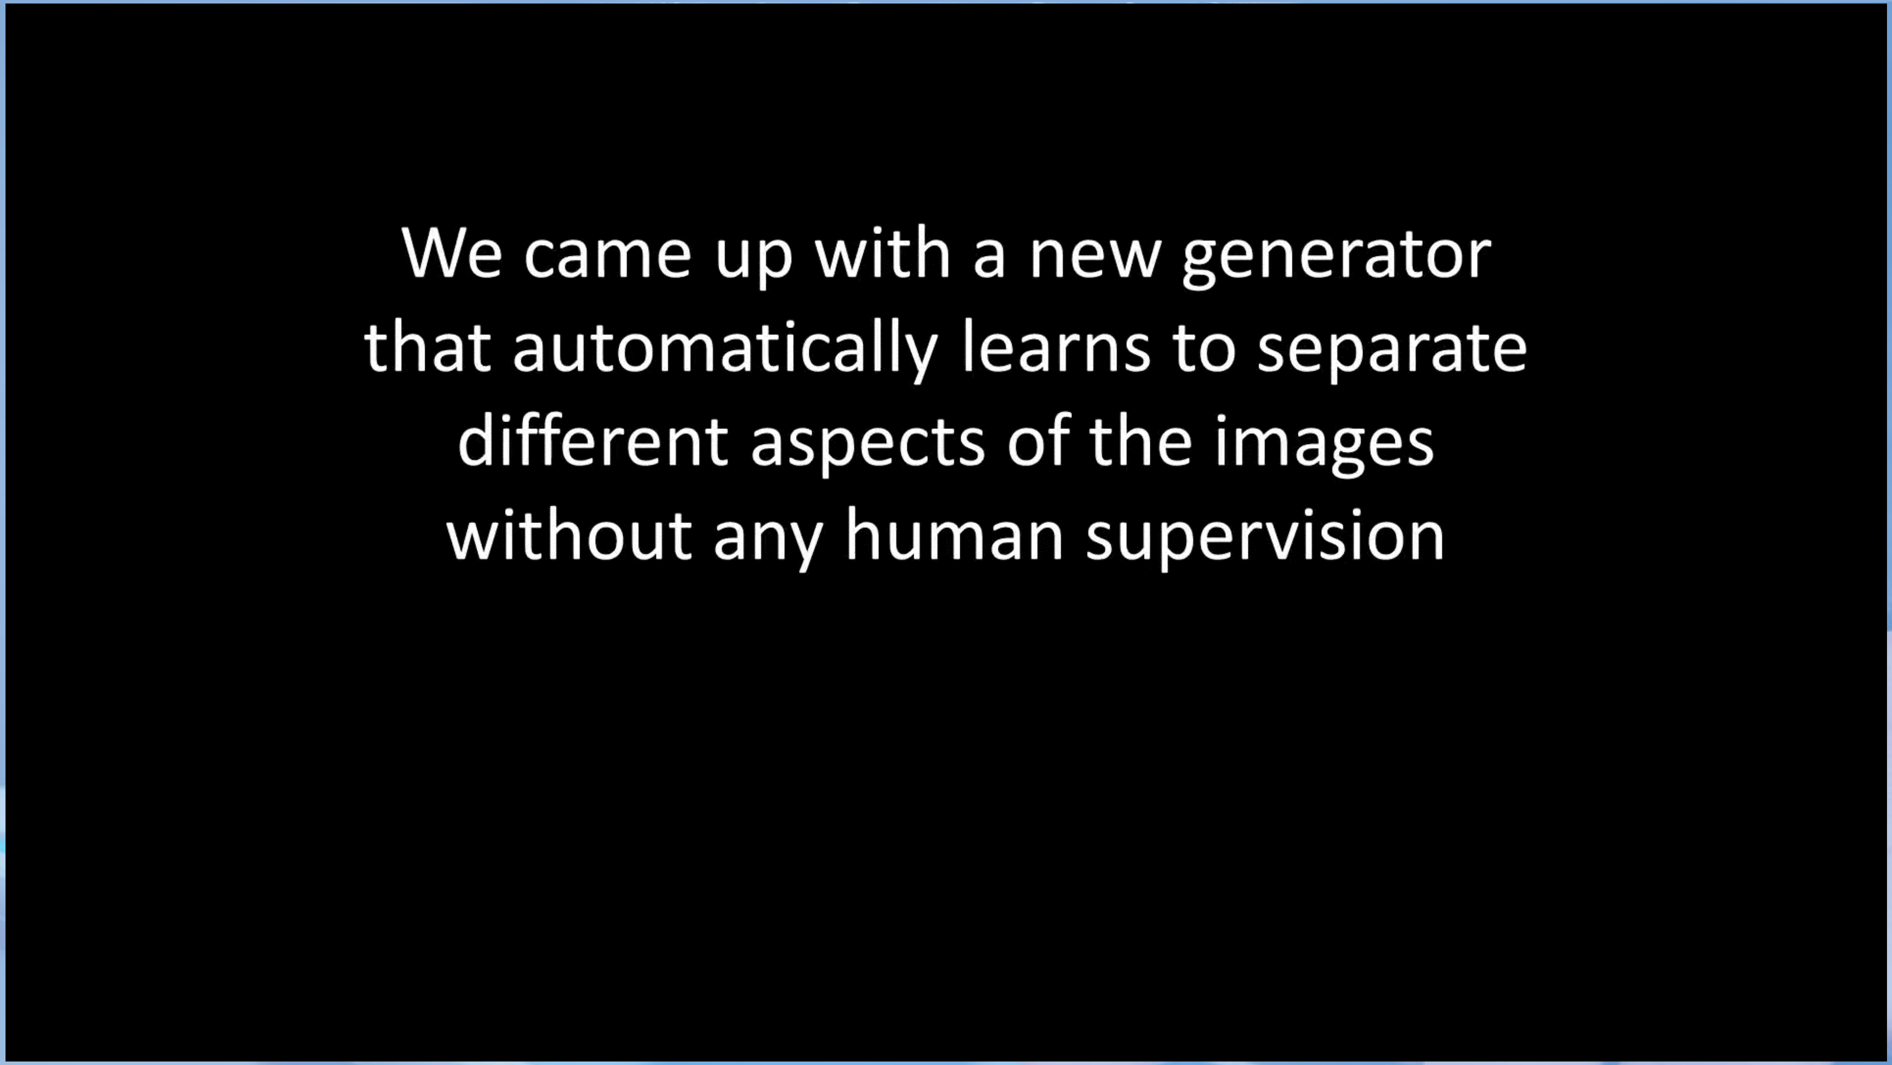

Source https://www.youtube.com/watch?v=kSLJriaOumA

03

All about data

Data Basics

- To train a network you data, sometimes lots of it!
- That data needs to be labelled (also called “annotated”)
- How much data depends on:
  - Complexity of the discrimination you are asking
  - How much similar data is available (transfer learning)
  - How much the data can be programmatically tweaked to make new data (data augmentation)

Hot dog, not hot dog

- Simple binary classification is the easiest

A fun example app: https://medium.com/@timanglade/how-hbos-silicon-valley-built-not-hotdog-with-mobile-tensorflow-keras-react-native-ef03260747f3

Multi-class discrimination

- Many classifiers are NOT neural network based

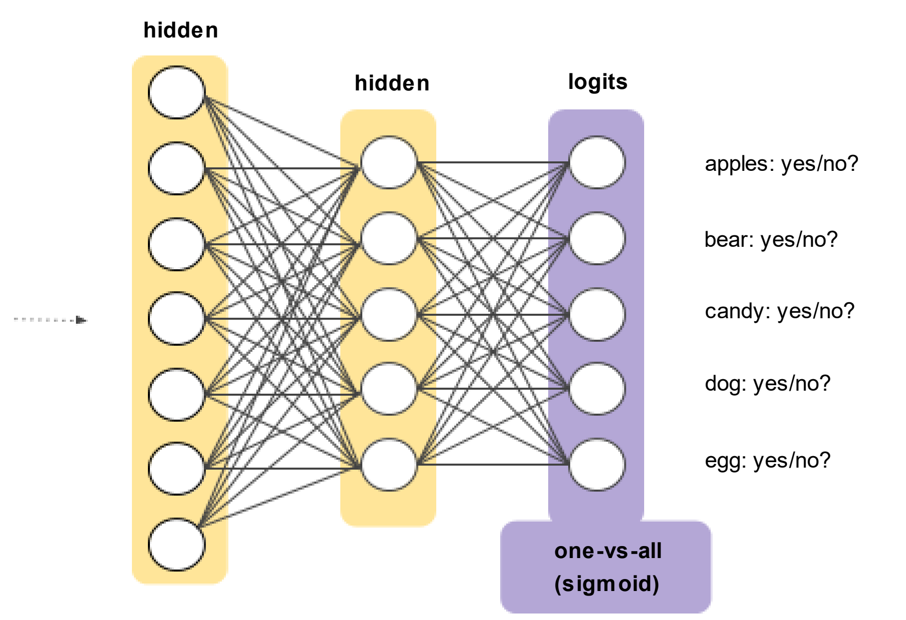

Details: https://developers.google.com/machine-learning/crash-course/multi-class-neural-networks/one-vs-all

Some example problems

- Dogs vs. cats
  - Binary classification
  - Complex shapes, but learnable
  - 256x256 pixel images, 1000 from each category
  - Final accuracy ~90%
- ImageNet
  - 14 million images, 20,000 categories
  - Human error ~5%, best machine error 6.8%

Good article http://karpathy.github.io/2014/09/02/what-i-learned-from-competing-against-a-convnet-on-imagenet/

Crowdsourced labels

- Labelling data can be expensive
- Applications that generate labels are one way
- Crowdsourced annotators are another

figure eight

gengo AI

Amazon Mechanical Turk

Access a global, on-demand, 24x7 workforce

Math-heavy, but a good review of crowdsourcing https://arxiv.org/pdf/1803.04223.pdf

04

Futures and Ethics

Moore's law

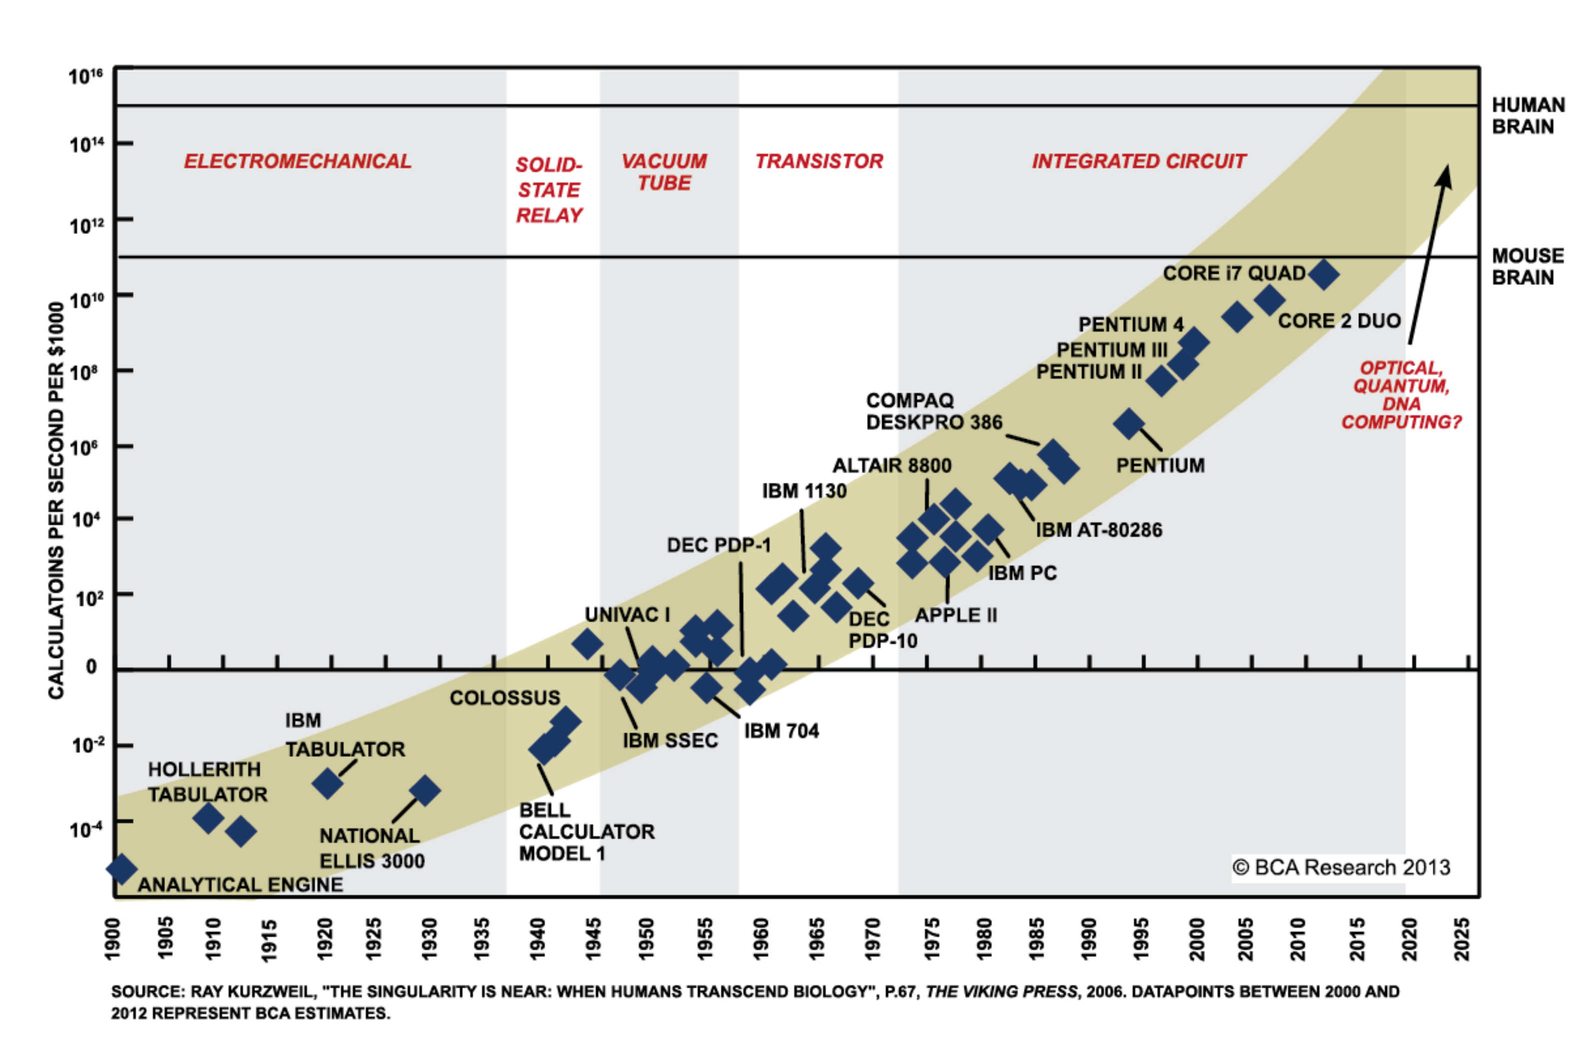

Human brain FLOPS https://aiimpacts.org/brain-performance-in-flops/

Moore's law – is it slowing?

Parallel computation

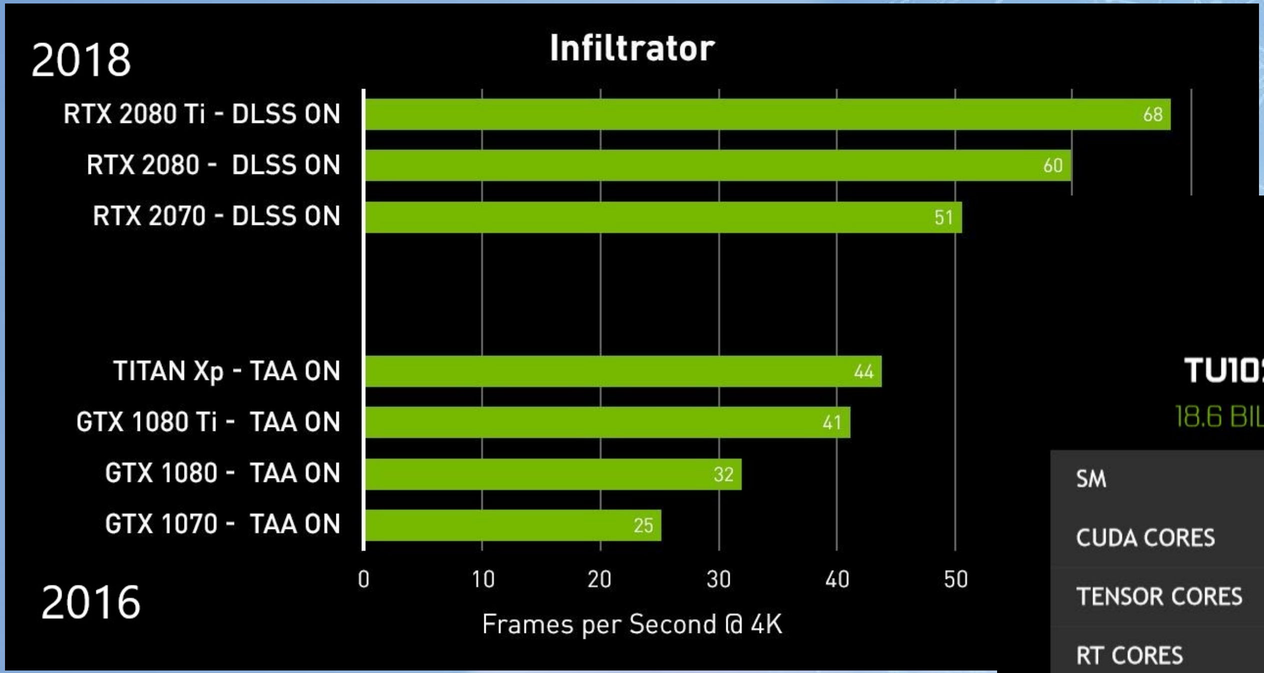

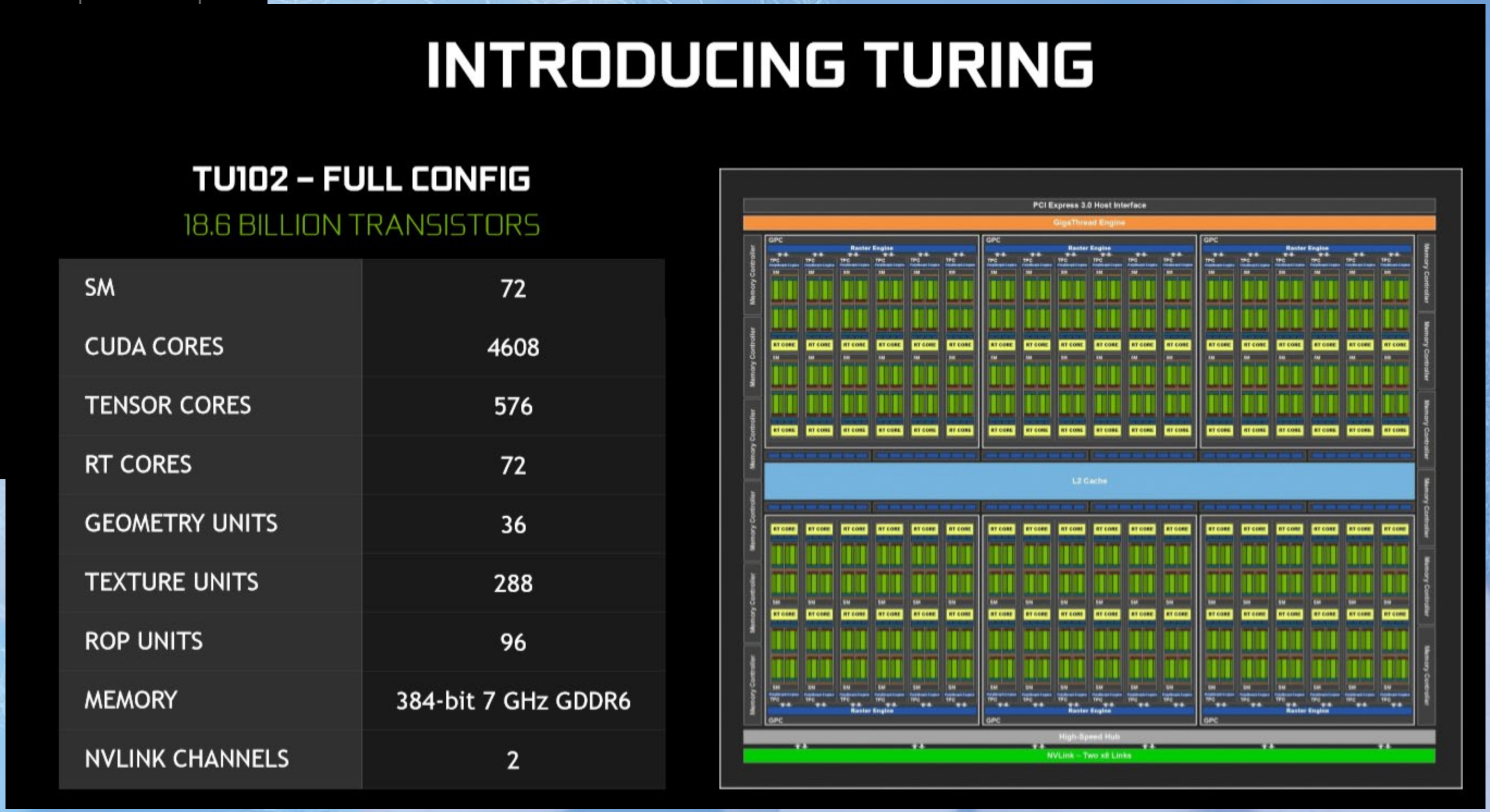

Blog article on Nvidia for deep learning https://blogs.nvidia.com/blog/2018/10/12/deep-learning-turing-graphics/

Cloud computing and TPUs

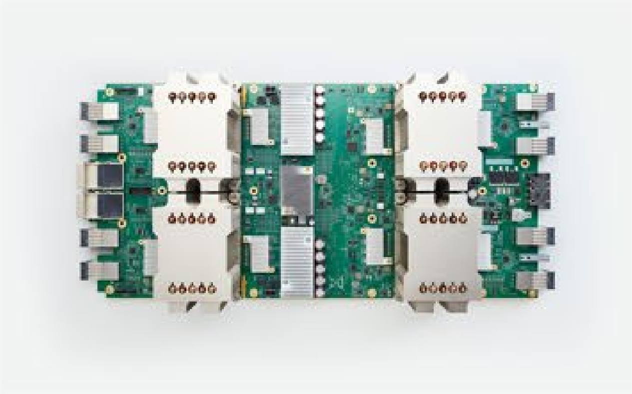

Cloud TPU v2
180 teraflops
64 GB High Bandwidth Memory (HBM)

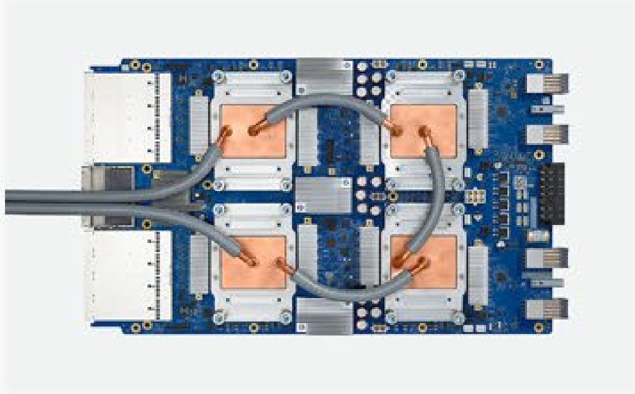

Cloud TPU v3
420 teraflops
128 GB HBM

Cloud TPU v2 Pod Alpha
11.5 petaflops
4 TB HBM
2-D toroidal mesh network

Google's TPU page https://cloud.google.com/tpu/ - a petaflop is one thousand million million operations per second (ten to the fifteenth power), teraflop is ten to the twelfth

Pop Culture Future of AI

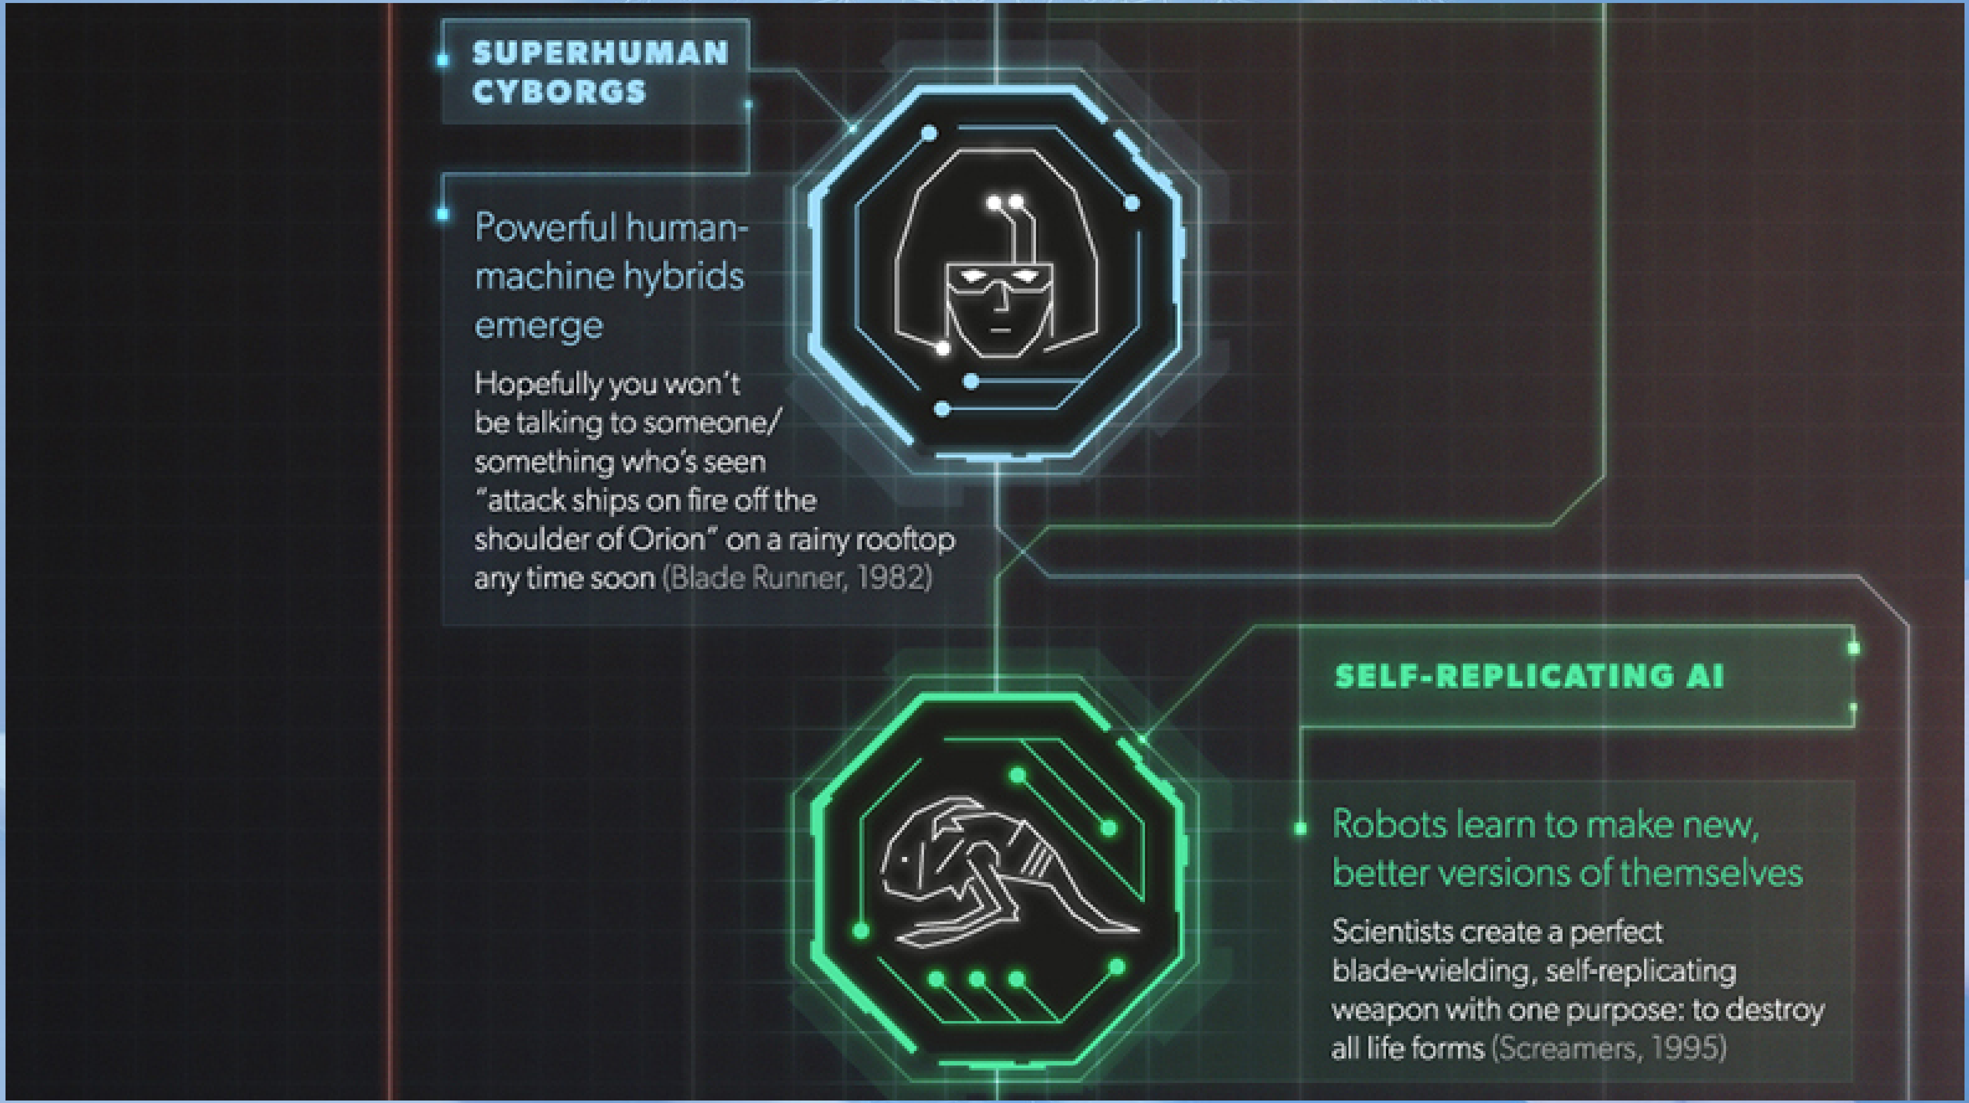

Realistic Future of AI

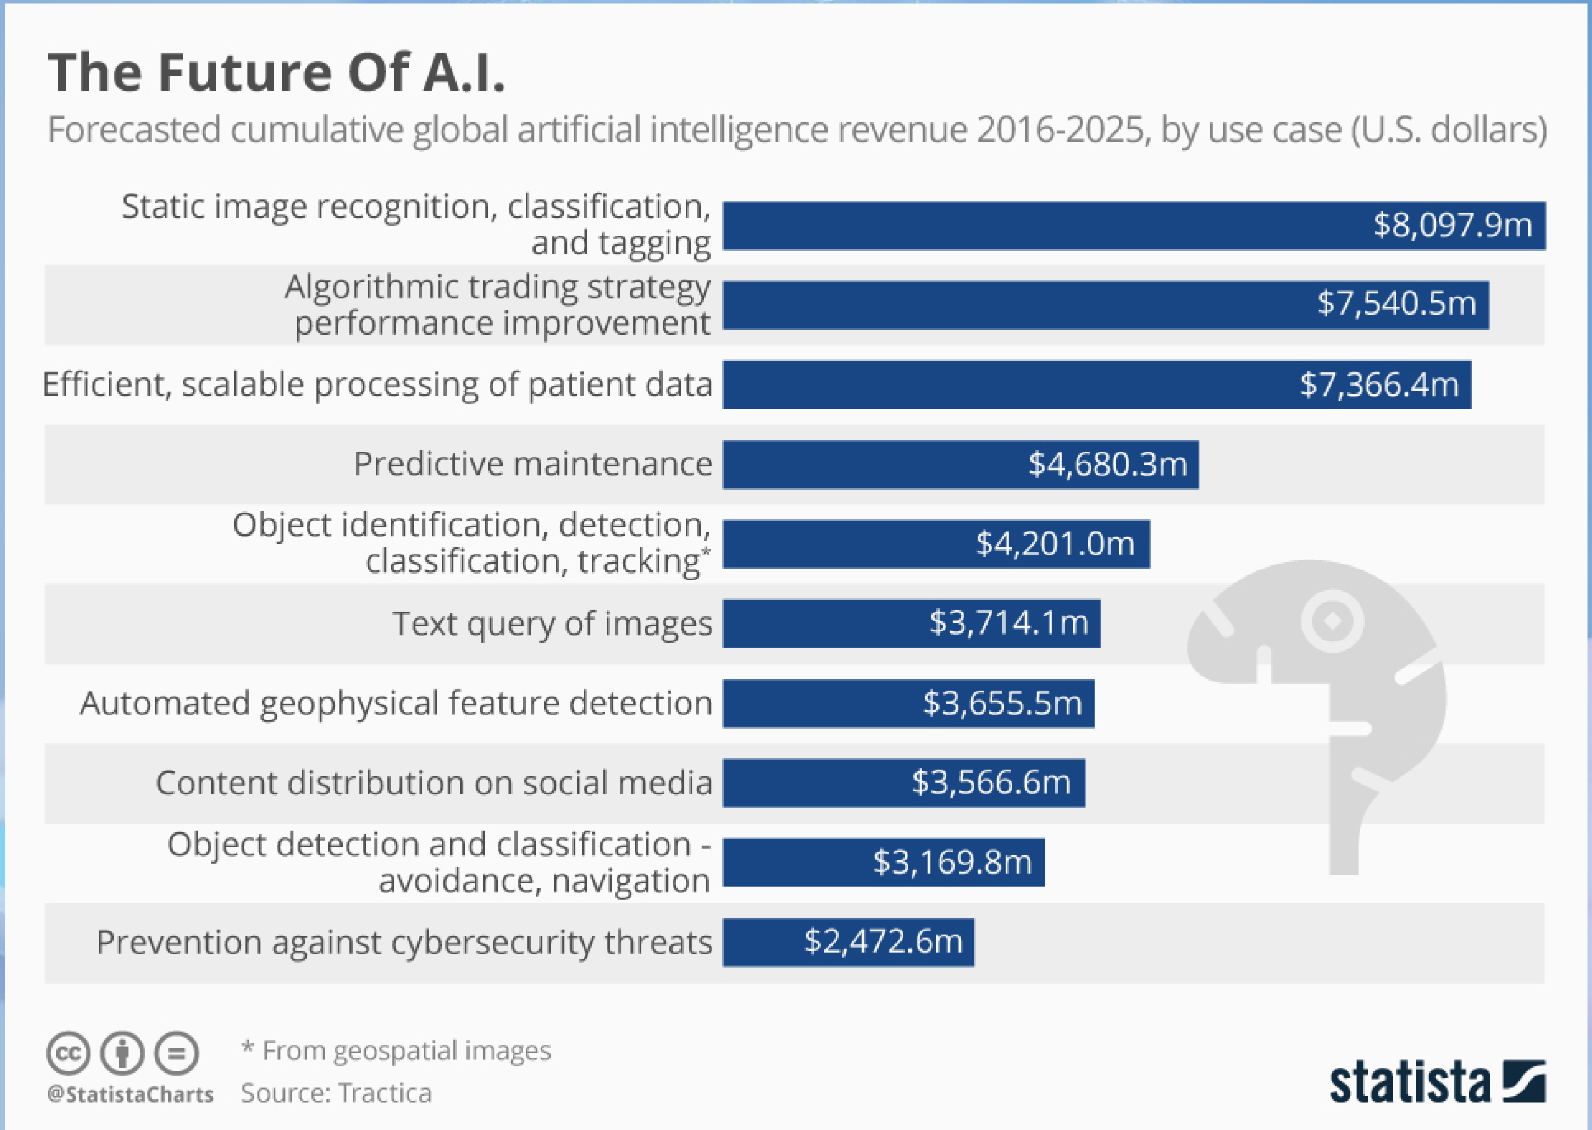

AI and jobs

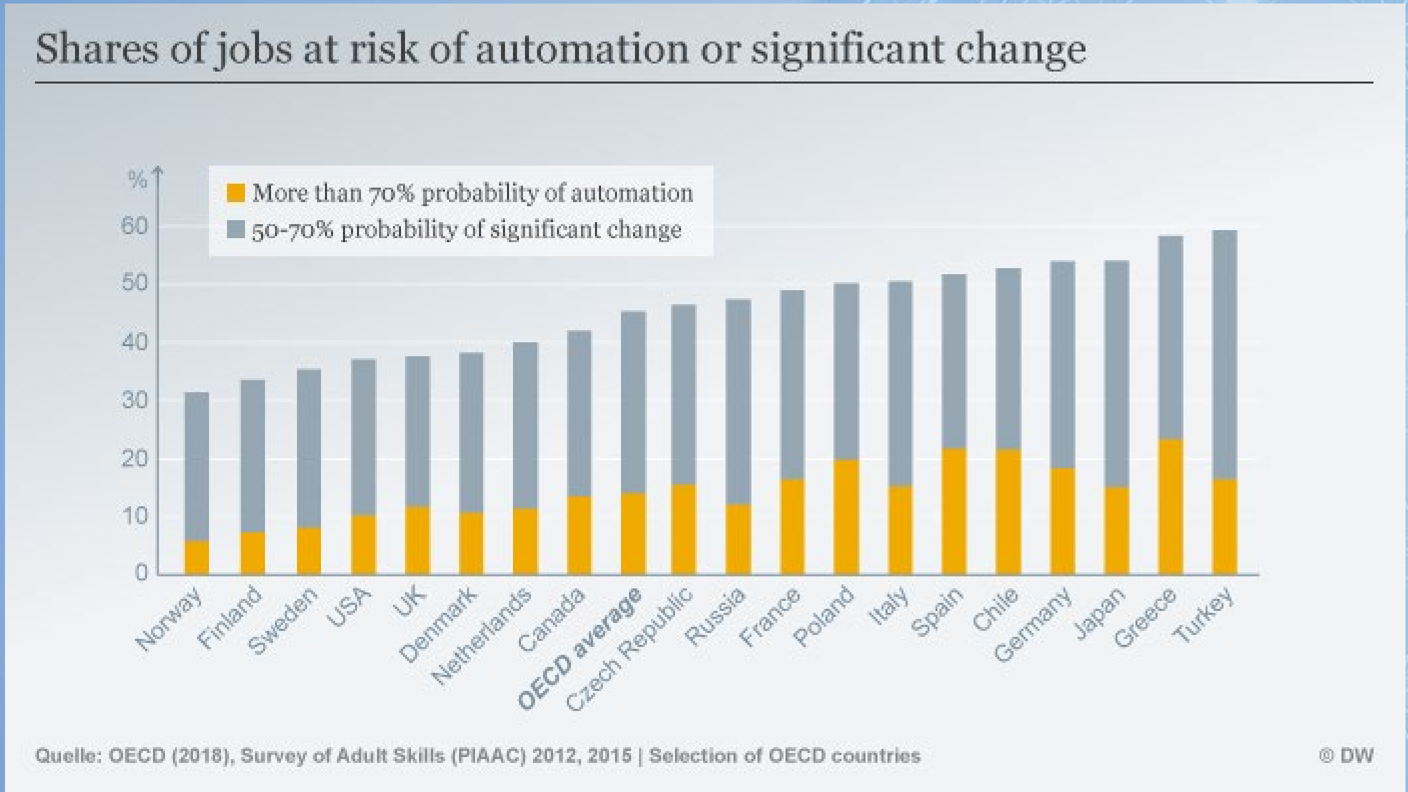

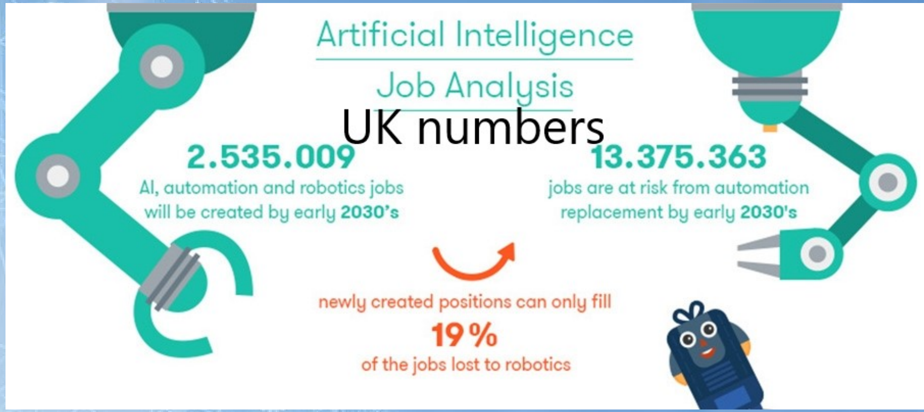

Source https://www.consultancy.uk/news/13770/ai-may-only-replace-19-percent-of-the-jobs-it-takes

AI impacts – Deep Fakes

Source https://www.youtube.com/watch?v=dDgPFk2u0E0 (fifth estate report)

AI impacts – language models

DATASET	METRIC	OUR RESULT	PREVIOUS RECORD	HUMAN
Winograd Schema Challenge	accuracy (+)	**70.70%**	63.7%	92%+
LAMBADA	accuracy (+)	**63.24%**	59.23%	95%+
LAMBADA	perplexity (-)	**8.6**	99	~1-2
Children's Book Test Common Nouns (validation accuracy)	accuracy (+)	**93.30%**	85.7%	96%
Children's Book Test Named Entities (validation accuracy)	accuracy (+)	**89.05%**	82.3%	92%

SYSTEM PROMPT (HUMAN-WRITTEN)

In a shocking finding, scientist discovered a herd of unicorns living in a remote, previously unexplored valley, in the Andes Mountains. Even more surprising to the researchers was the fact that the unicorns spoke perfect English.

MODEL COMPLETION (MACHINE-WRITTEN, 10 TRIES)

The scientist named the population, after their distinctive horn, Ovid's Unicorn. These four-horned, silver-white unicorns were previously unknown to science.

Now, after almost two centuries, the mystery of what sparked this odd phenomenon is finally solved.

Dr. Jorge Pérez, an evolutionary biologist from the University of La Paz, and several companions, were exploring the Andes Mountains when they found a small valley, with no other animals or humans. Pérez noticed that the valley had what appeared to be a natural fountain, surrounded by two peaks of rock and silver snow.

Pérez and the others then ventured further into the valley. "By the time we reached the top of one peak, the water looked blue, with some crystals on top," said Pérez.

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

Ethics and personhood

The 21st century is in dire need of a Turing test for consciousness.

You don't think you're a zombie, but that's just what a zombie would say. – David Chalmers

Source http://nautil.us/blog/heres-how-well-know-an-ai-is-conscious

05

Fun examples!

Tensorflow playground

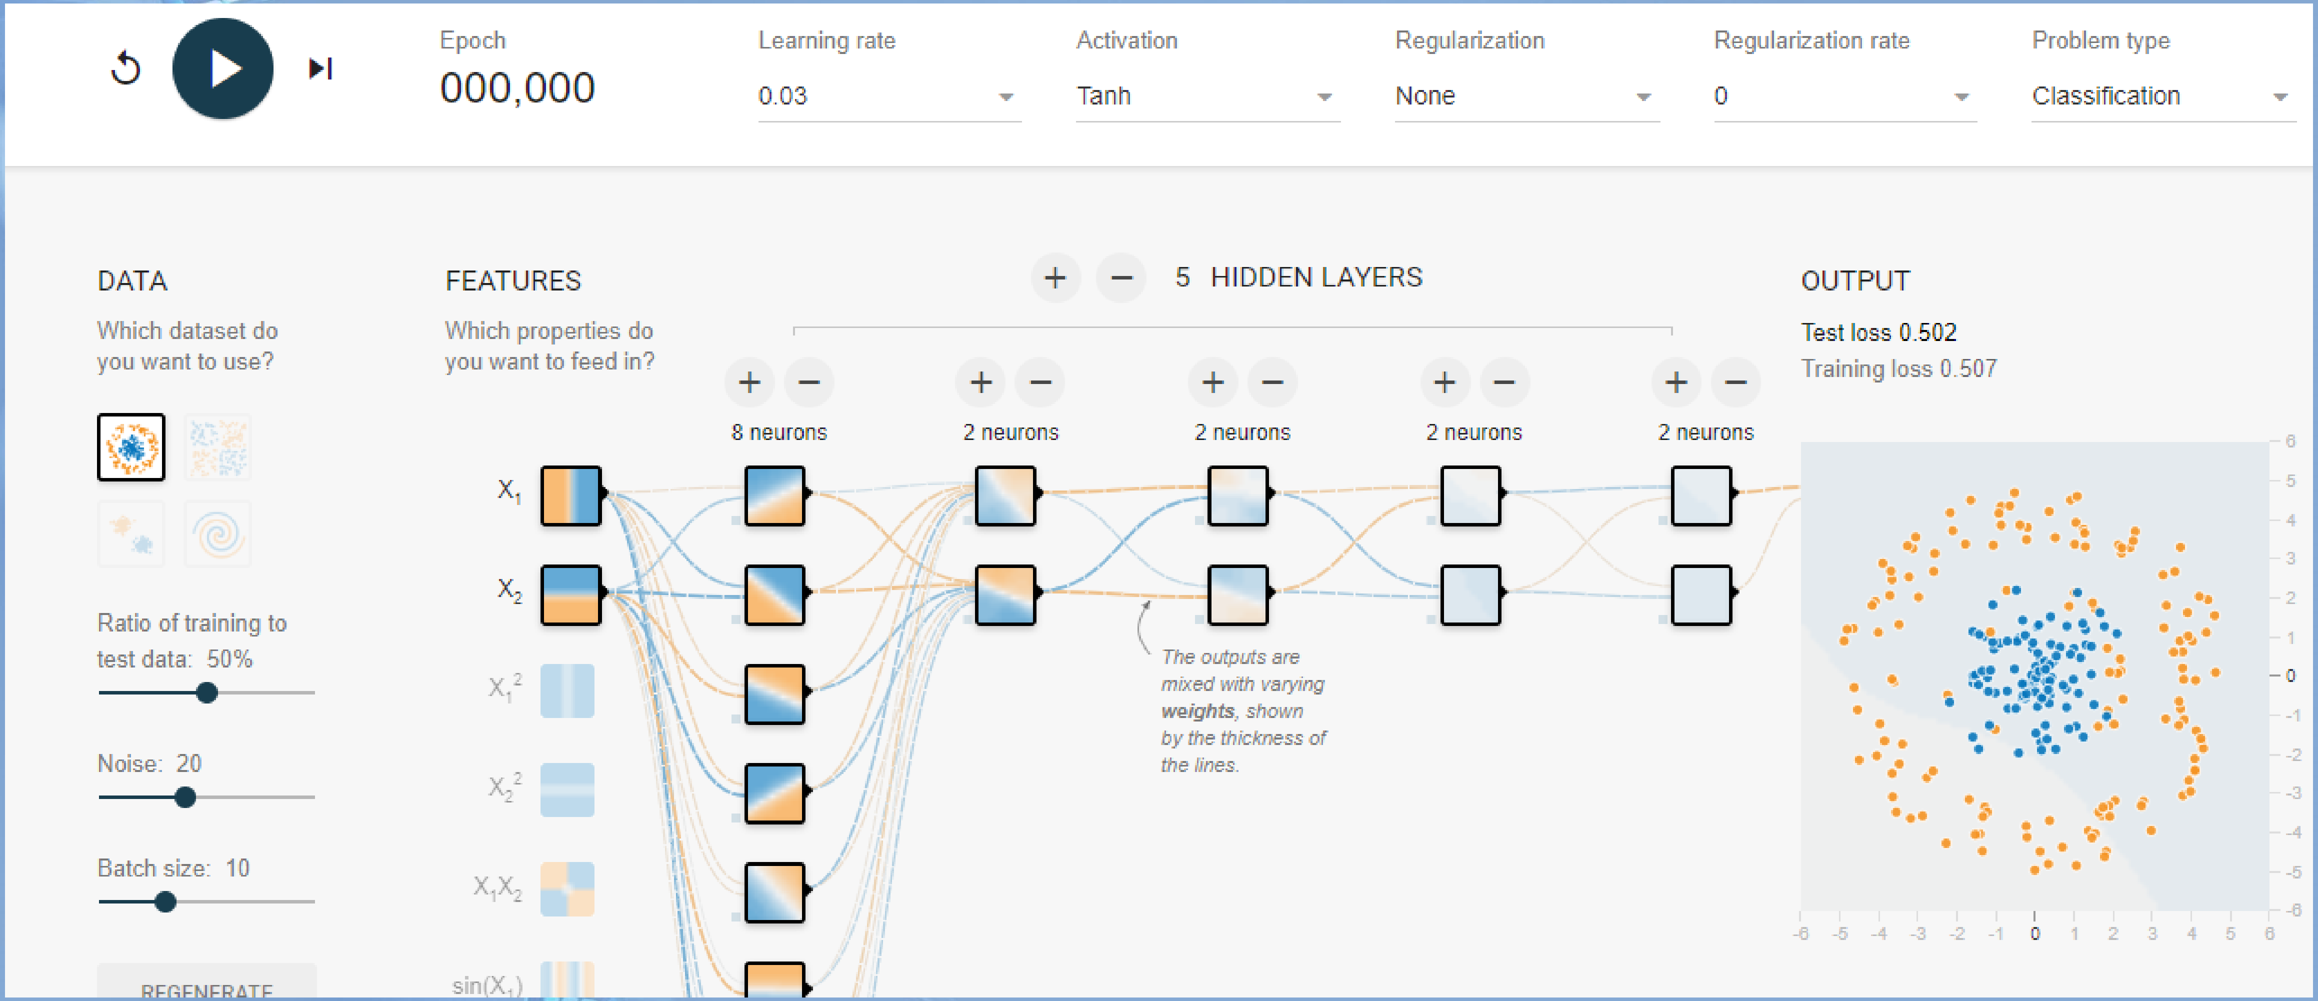

Source http://playground.tensorflow.org

Gen Studio

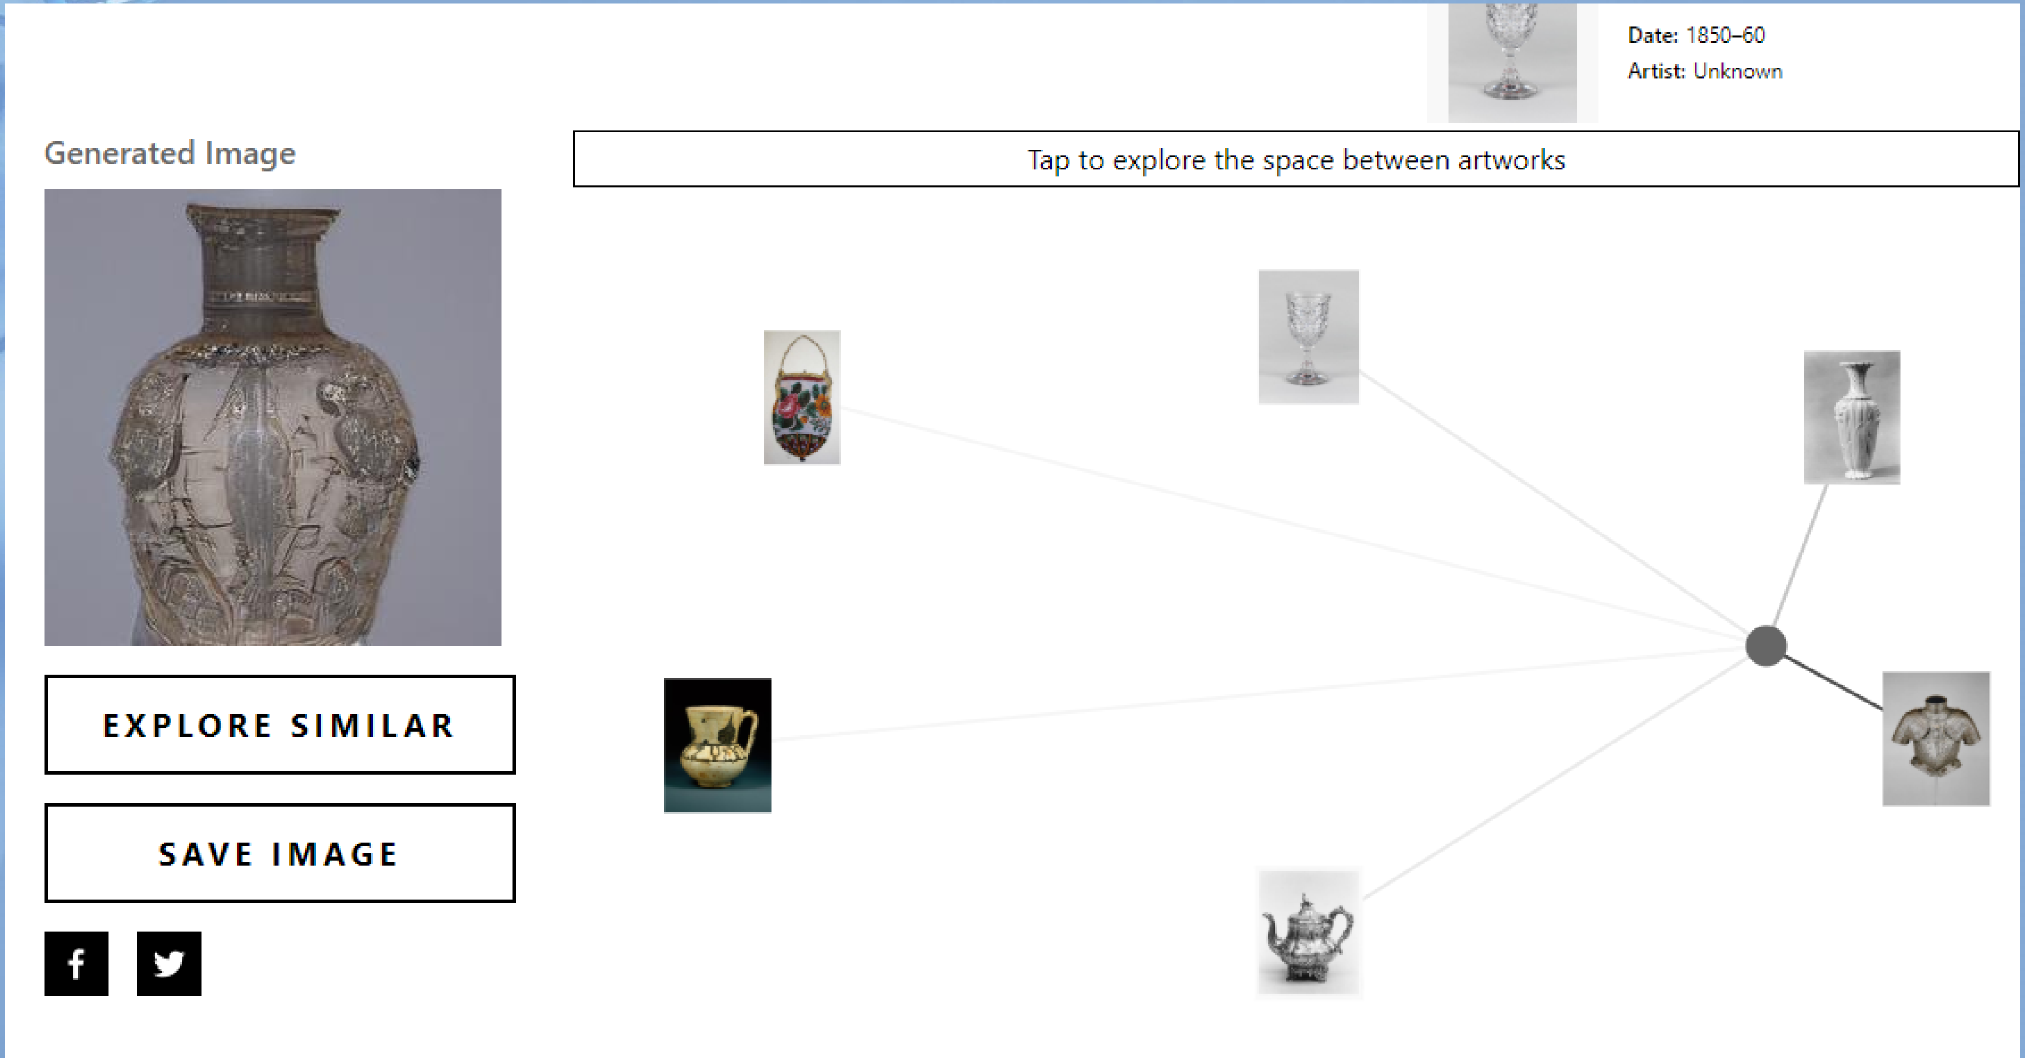

Source https://gen.studio/ architectural details at https://www.ailab.microsoft.com/experiments/gen-studio

iNaturalist

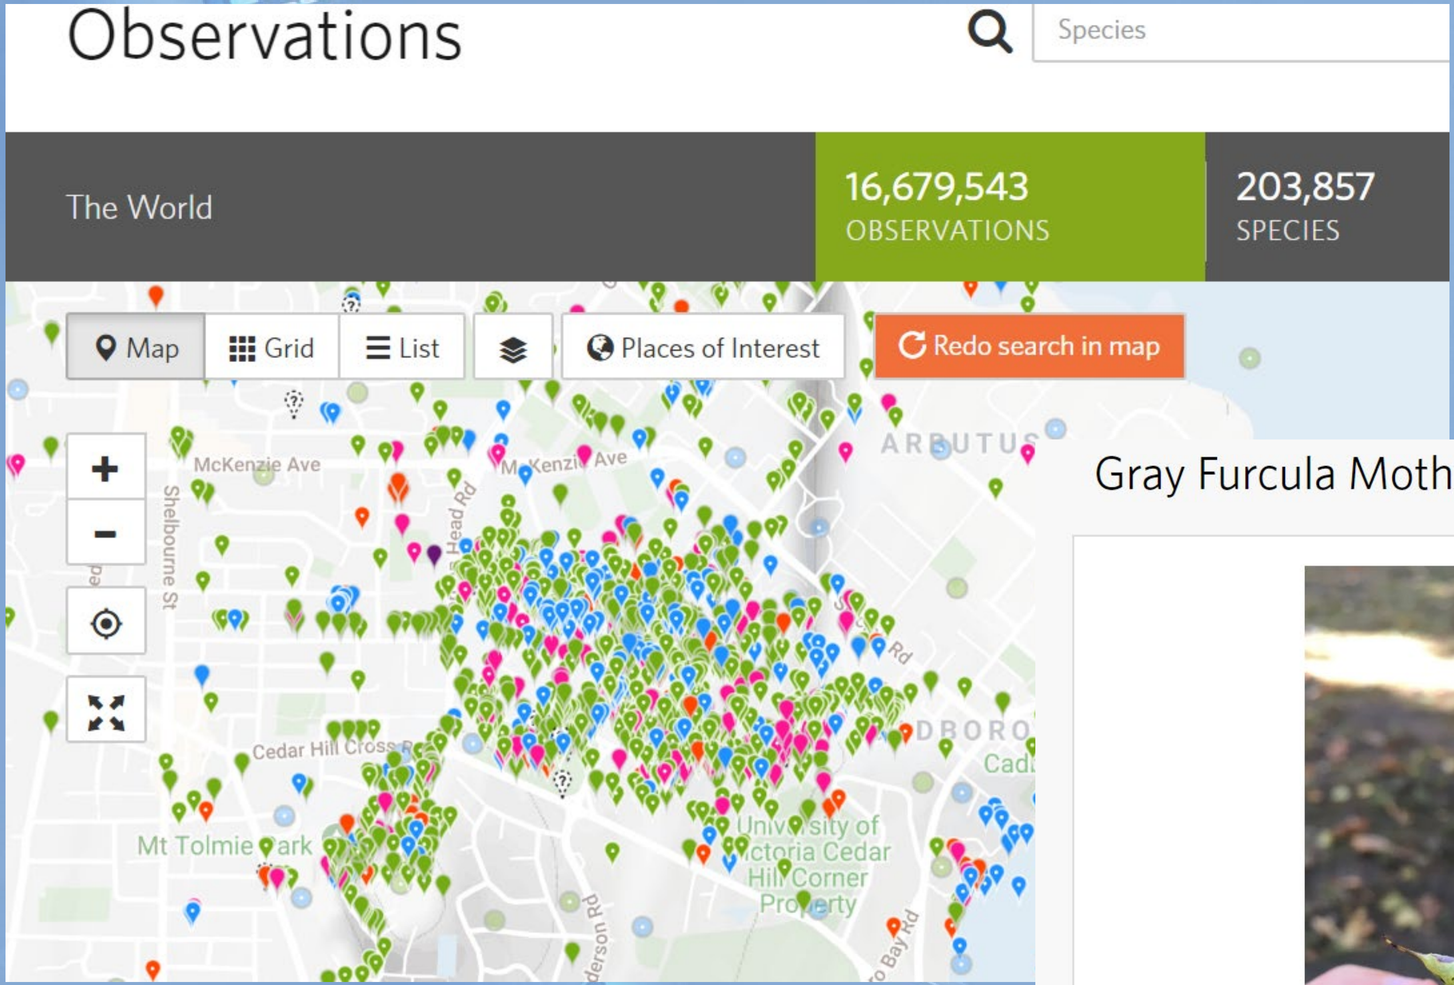

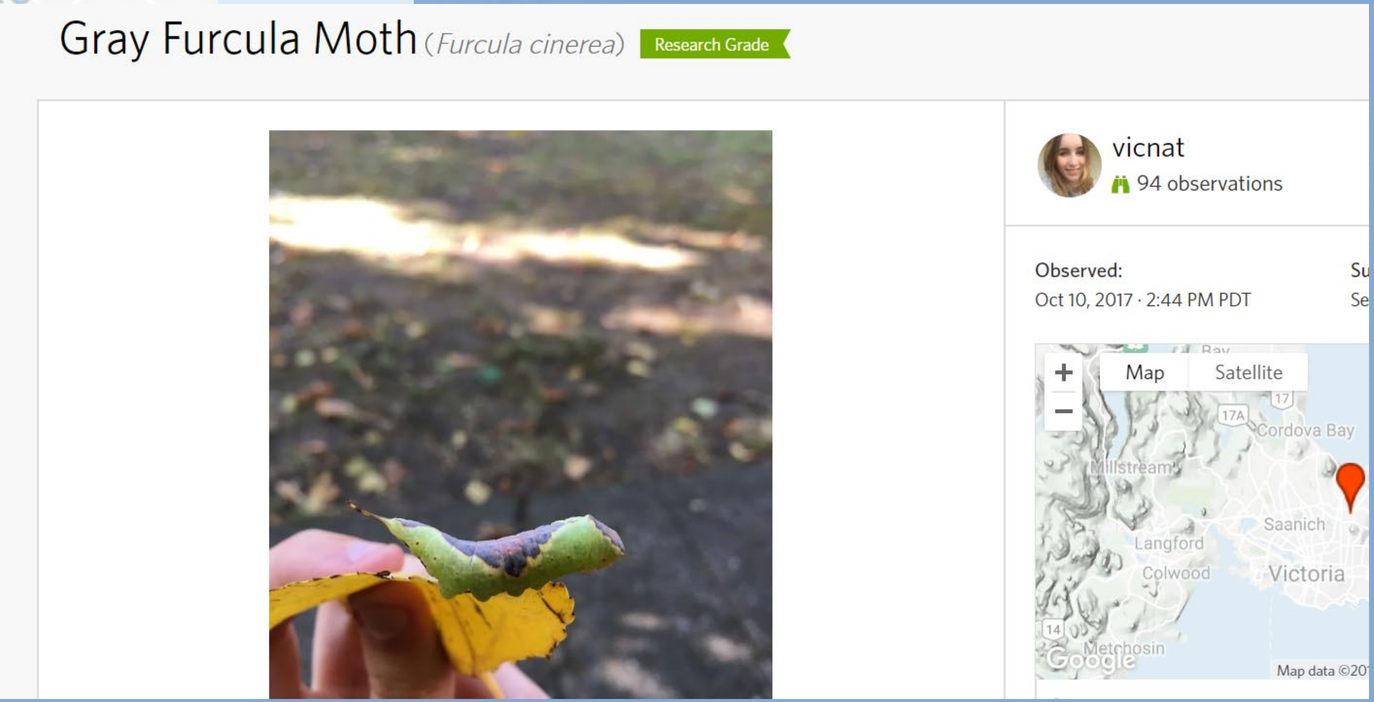

See https://www.inaturalist.org/

Thank You

Slides can be found at http://qvirt.com/aitalk.html