An introduction to neural networks

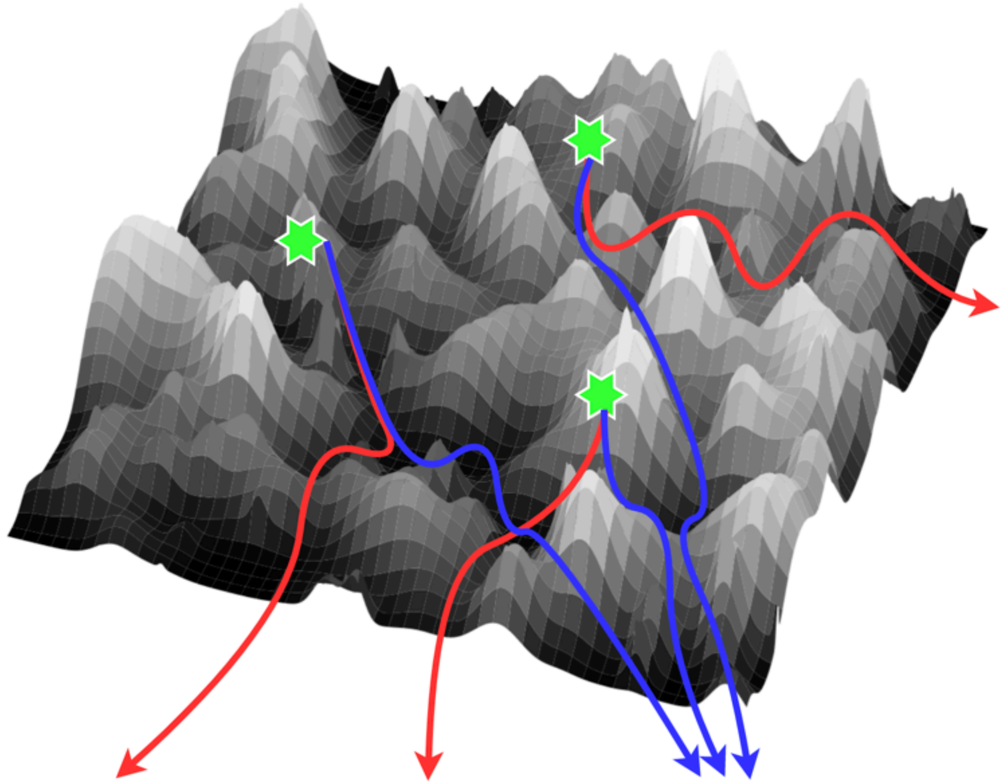

Sebastian **Goldt** (SISSA, Trieste)

Junior Math Days @ SISSA — dec 2023

New machine learning breakthroughs...

https://twitter.com/StelfieTT

New machine learning breakthroughs...

Hey ChatGP, I'm interested in the theory of neural networks. Do you know anything about that?

Yes, I'm familiar with the theory of neural networks. Neural networks are a type of machine learning algorithm that are modeled after the structure and function of the human brain. They consist of layers of interconnected "neurons," which process and transmit information. Neural networks are commonly used for tasks such as image and speech recognition, and they have been instrumental in many of the recent advances in artificial intelligence.

You heard the news

Hey ChatGP, I'm interested in the theory of neural networks. Do you know anything about that?

Yes, I'm familiar with the theory of neural networks. Neural networks are a type of machine learning algorithm that are modeled after the structure and function of the human brain. They consist of layers of interconnected "neurons," which process and transmit information. Neural networks are commonly used for tasks such as image and speech recognition, and they have been instrumental in many of the recent advances in artificial intelligence.

In the words of the New York Times:

“The Navy revealed the embryo of an electronic computer that it expects will be able to walk, talk, see, write, reproduce itself and be conscious of its own existence”

NEW NAVY DEVICE LEARNS BY DOING

Psychologist Shows Embryo of Computer Designed to Read and Grow Wiser

WASHINGTON, July 7 (UPI) —The Navy revealed the embryo of an electronic computer today that it expects will be able to walk, talk, see, write, reproduce itself and be conscious of its existence.

The embryo—the Weather Bureau's $2,000,000 "704" computer—learned to differentiate between right and left after fifty attempts in the Navy's demonstration for newsmen.

The service said it would use this principle to build the first of its Perceptron thinking machines that will be able to read and write. It is expected to be finished in about a year at a cost of $100,000.

New York Times, July 8, 1958

The plan for today

What is a neural network?

From neurons to networks.

The importance of feature learning

The secret sauce of neural networks

Part I

What is a neural network?

What is a neural network?

A neural network is a (complicated) function

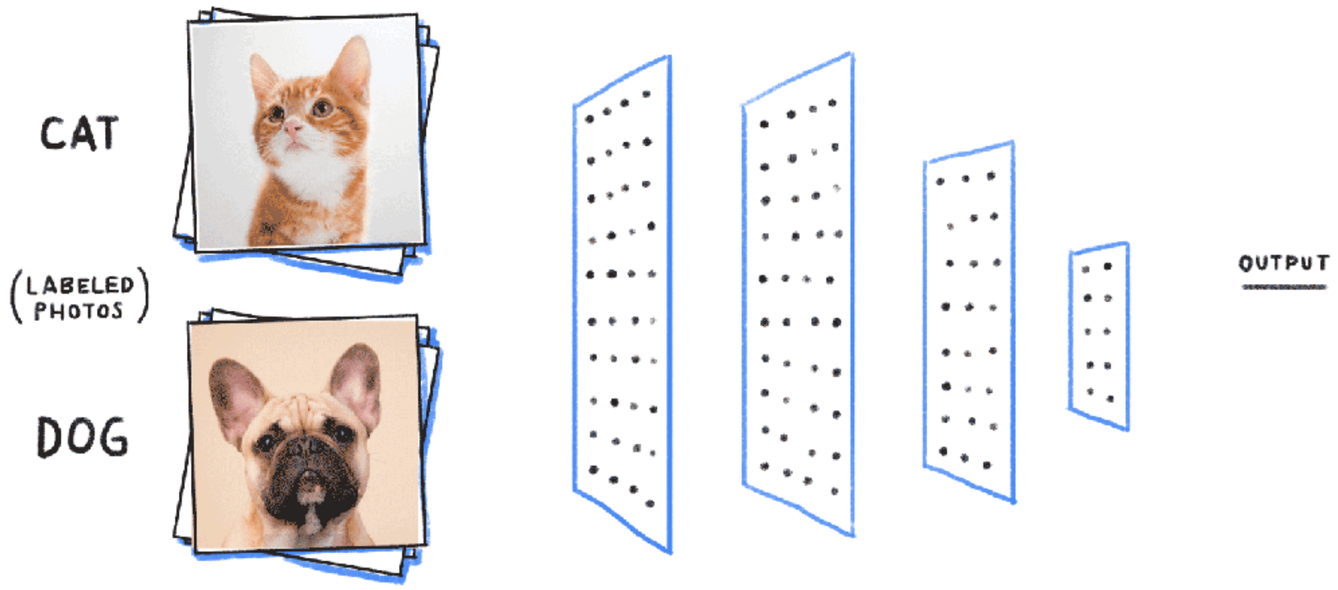

Animation courtesy of Aakash Srivastava

Let’s take a closer look at
one of these small black dots…

A single neuron

Many inputs, one output

A LOGICAL CALCULUS OF THE
IDEAS IMMANENT IN NERVOUS ACTIVITY

WARREN S. MCCULLOCH AND WALTER PITTS

FROM THE UNIVERSITY OF ILLINOIS, COLLEGE OF MEDICINE,
DEPARTMENT OF PSYCHIATRY AT THE ILLINOIS NEUROPSYCHIATRIC INSTITUTE,
AND THE UNIVERSITY OF CHICAGO

(1943)

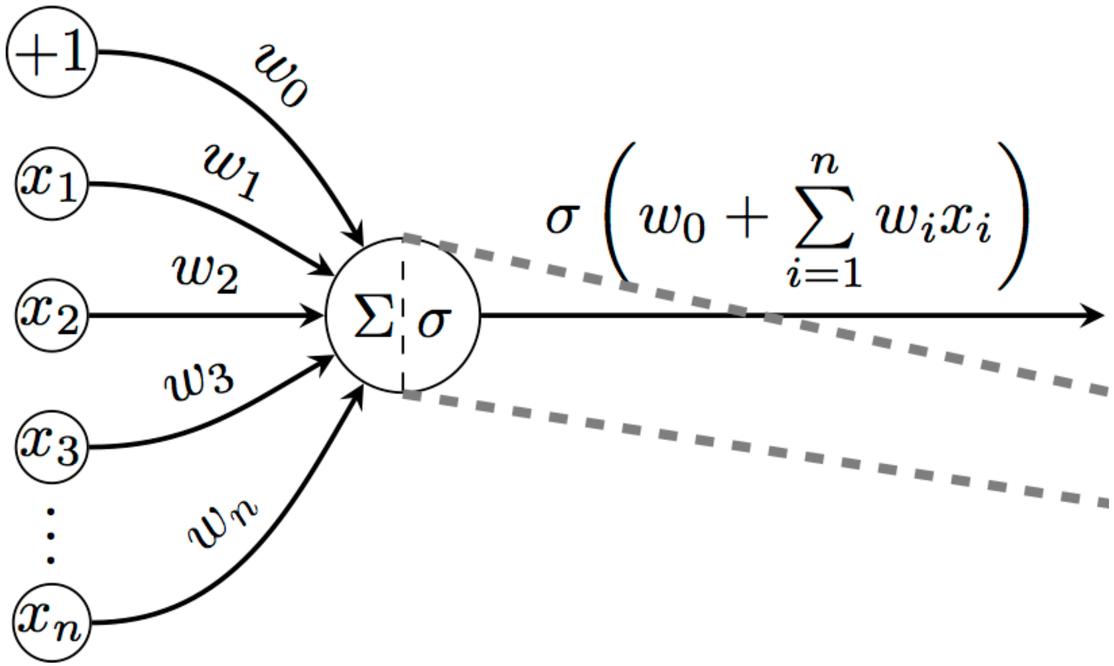

Illustration by Petar Veličković
https://github.com/PetarV-/TikZ

The activation function

More than just summing up

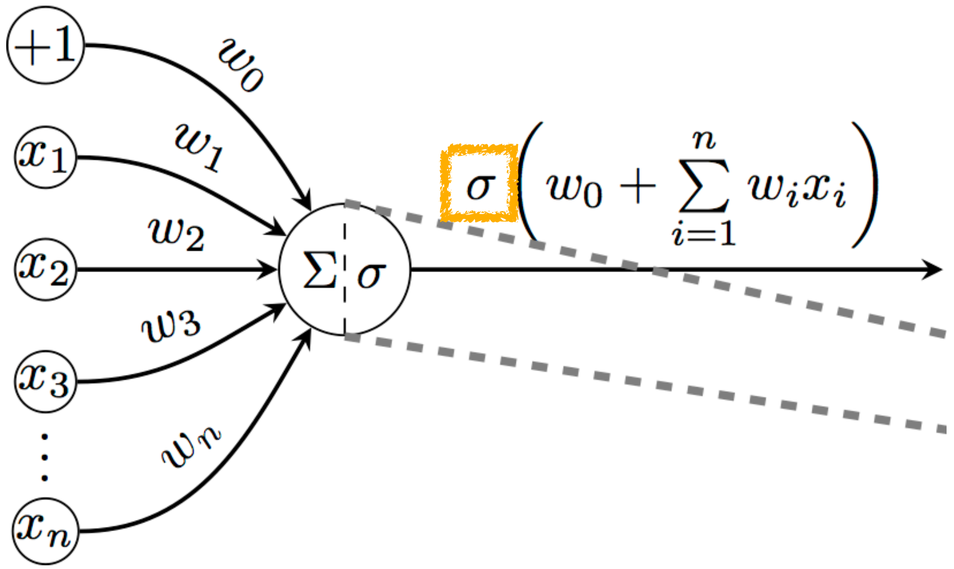

Linear

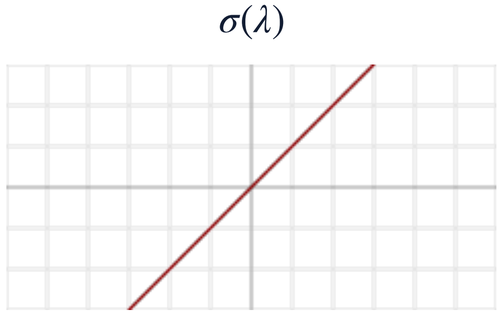

$$\lambda = w_0 + \sum_{i}^{n} w_i x_i$$

Sigmoid

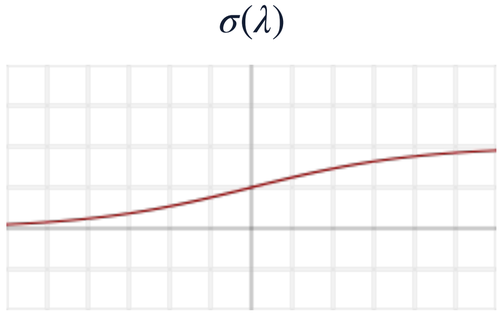

ReLU

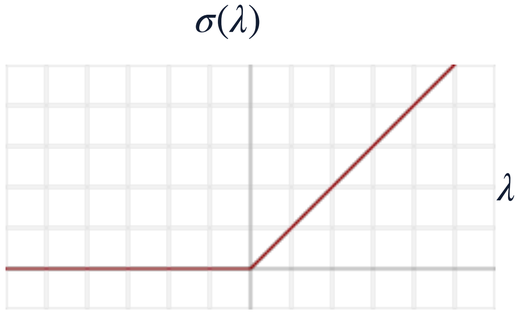

A lot of neurons

Neurons can be assembled into layers

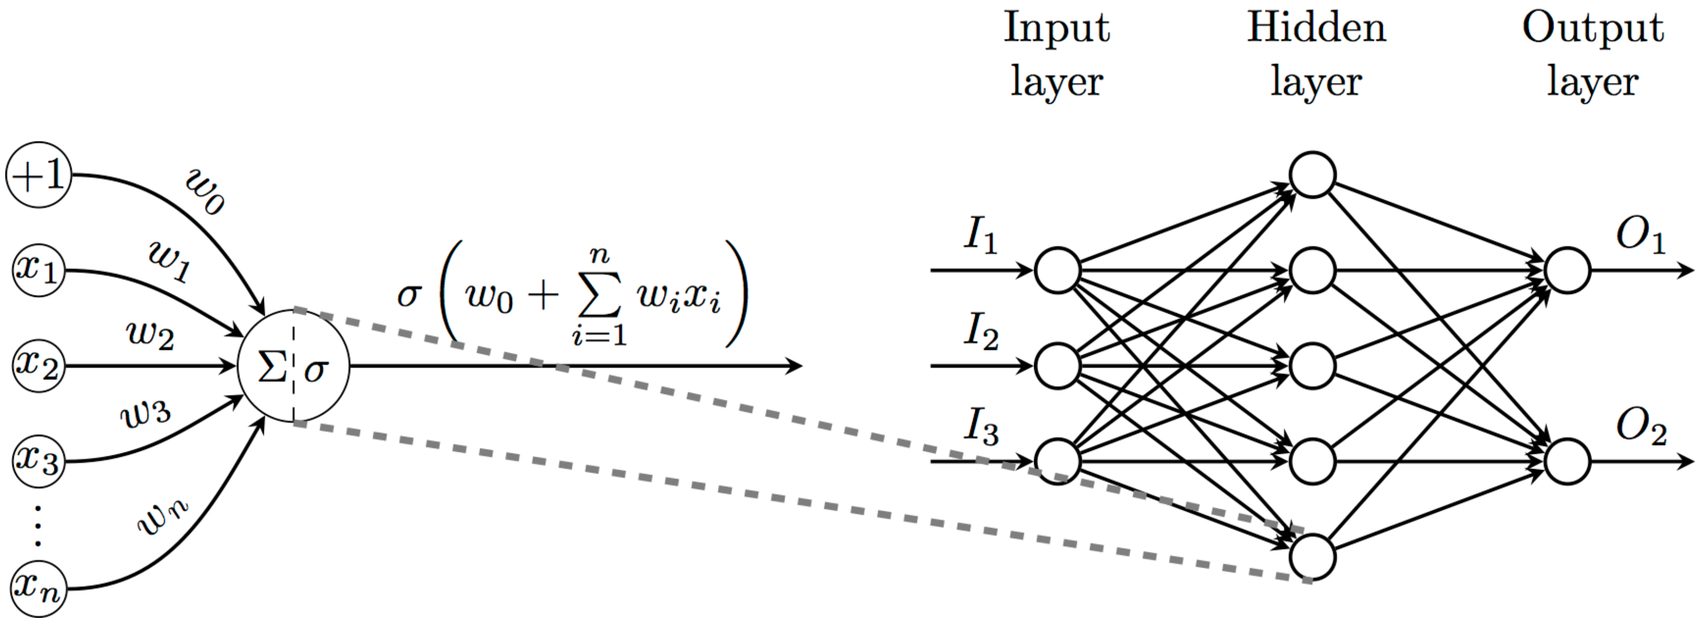

Illustration by Petar Veličković
https://github.com/PetarV-/TikZ

A lot of layers

Deep neural networks stack layers of neurons

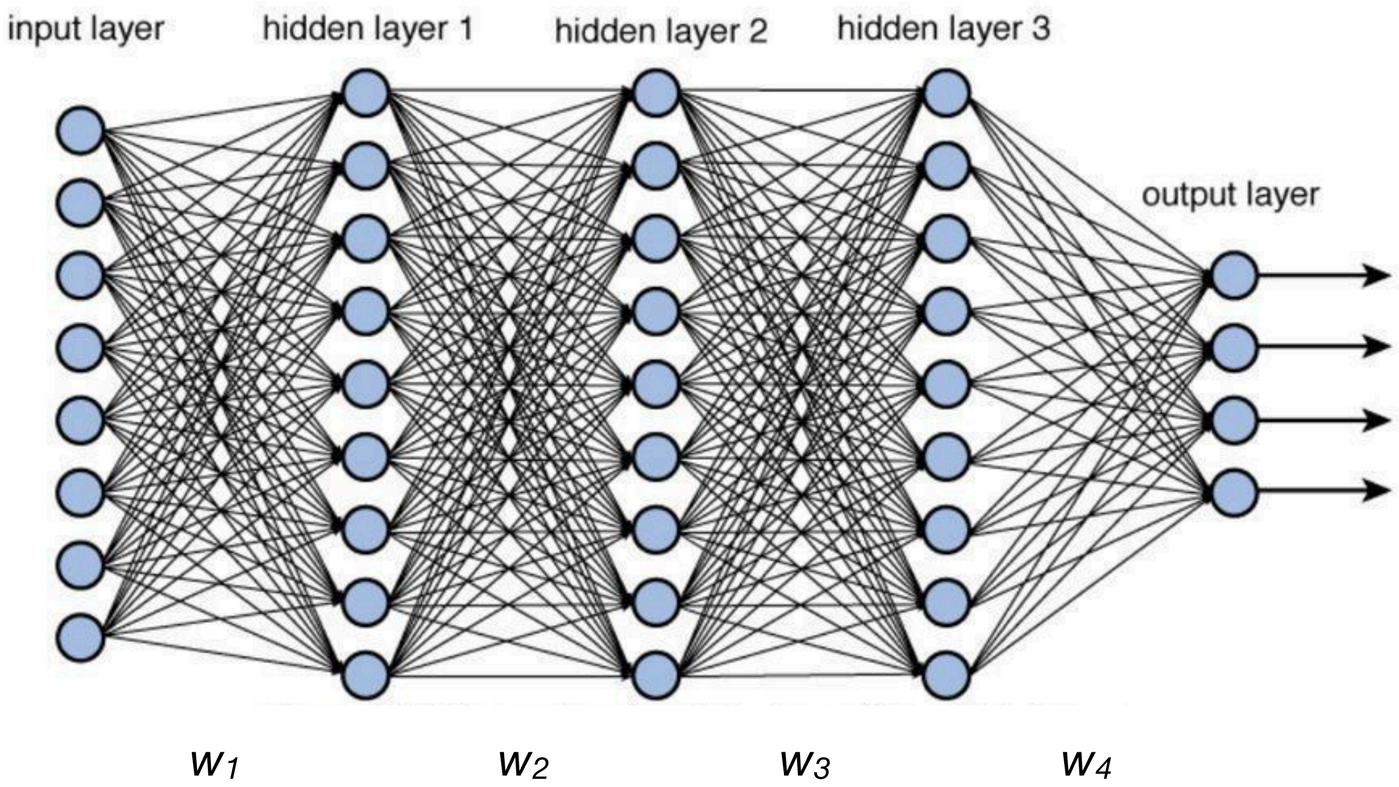

$$y = w_4\sigma\left(w_3\sigma(w_2\sigma(w_1x))\right)$$

A mostly complete chart of neural networks

$y = w_4\sigma\left(w_3\sigma(w_2\sigma(w_1x))\right)$ is not the whole story

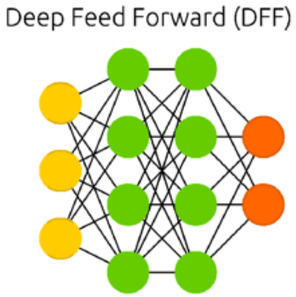

Illustration by Fjodor van Veen & Stefan Leijnen

A mostly complete chart of neural networks

$y = w_4\sigma\left(w_3\sigma(w_2\sigma(w_1x))\right)$ is not the whole story

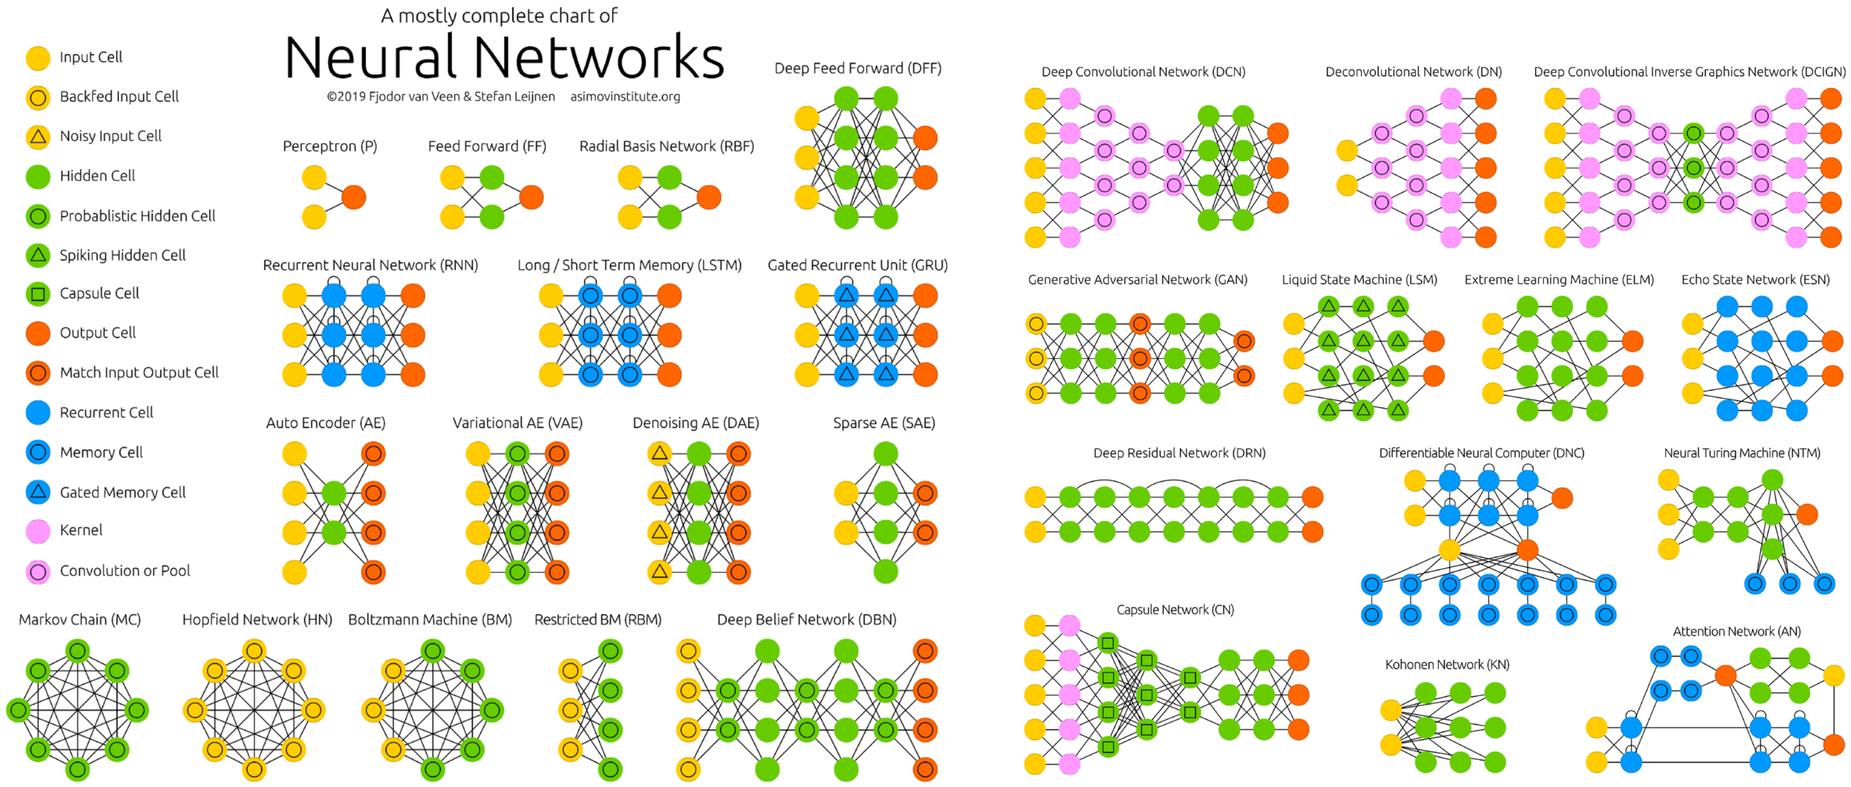

Illustration by Fjodor van Veen & Stefan Leijnen

Which architecture for which data?

Training a neural network

Having the architecture, how do we find the weights?

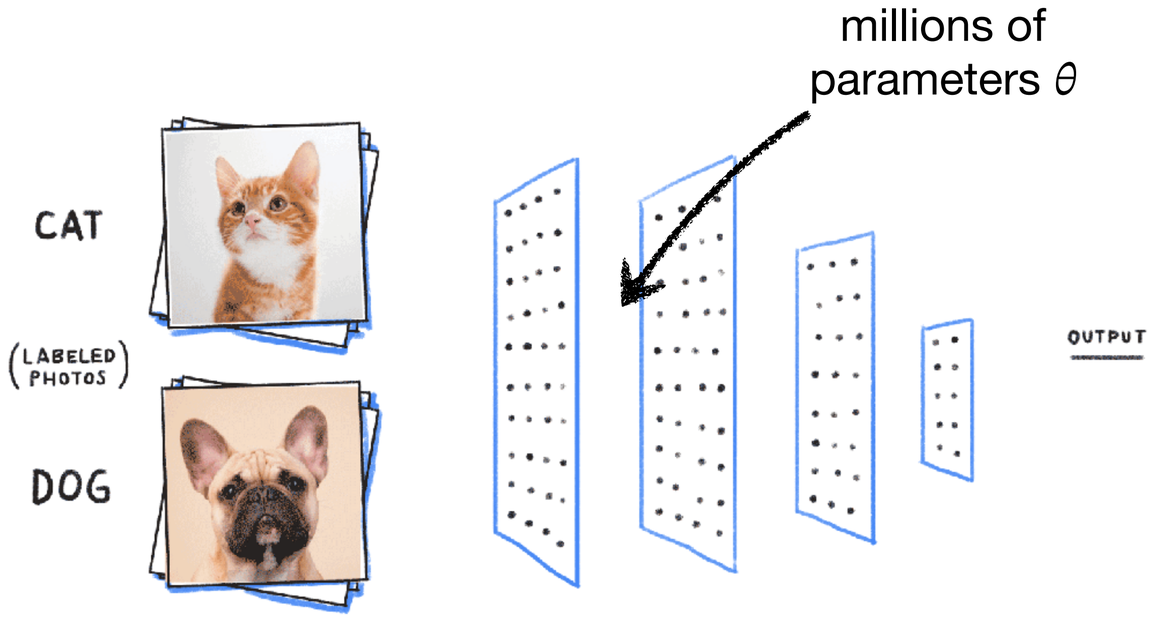

Animation courtesy of Aakash Srivastava

Training a neural network

Having the architecture, how do we find the weights?

- Given a data set $\mathscr{D} = \{(x_i, y_i\}_{i=1}^N$

loss(θ) = # of mis-classified training images

$$\theta^* = \underset{\theta}{\operatorname{argmin}} \operatorname{loss}(\theta)$$

$$\epsilon_t = \operatorname{loss}(\theta^*)$$

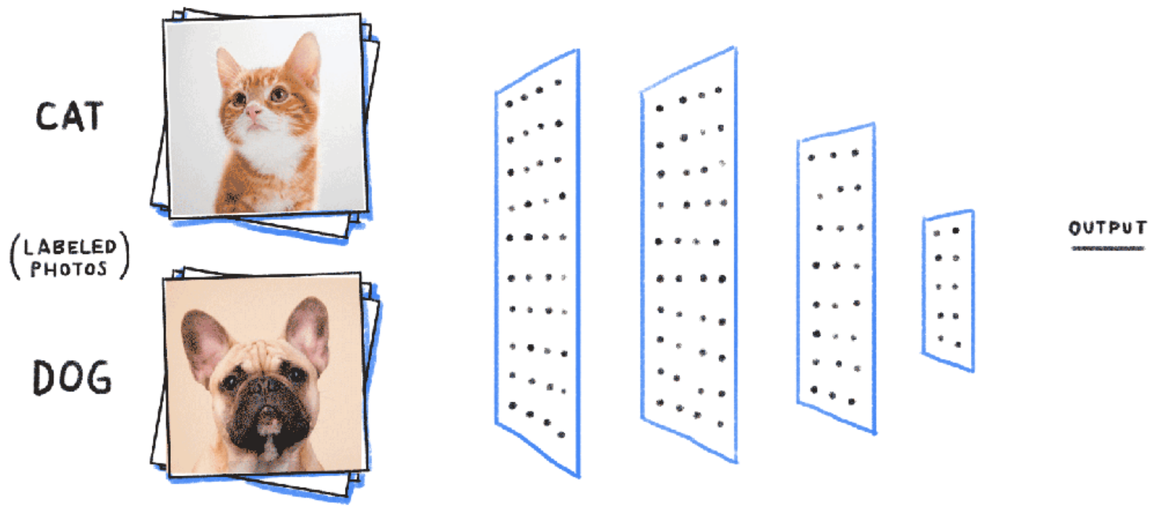

Animation courtesy of Aakash Srivastava

Training a neural network in theory

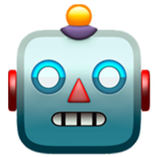

It’s not so easy!

TRAINING A 3-NODE NEURAL NETWORK IS NP-COMPLETE

Avrim Blum*
MIT Lab. for Computer Science
Cambridge, Mass. 02139 USA

Ronald L. Rivest†
MIT Lab. for Computer Science
Cambridge, Mass. 02139 USA

Advances in Neural Information Processing (1989)

TRAINING A 3-NODE NEURAL NETWORK IS NP-COMPLETE

Avrim Blum*
MIT Lab. for Computer Science
Cambridge, Mass. 02139 USA

Ronald L. Rivest†
MIT Lab. for Computer Science
Cambridge, Mass. 02139 USA

Advances in Neural Information Processing (1989)

Extends a previous result by Judd (1987)

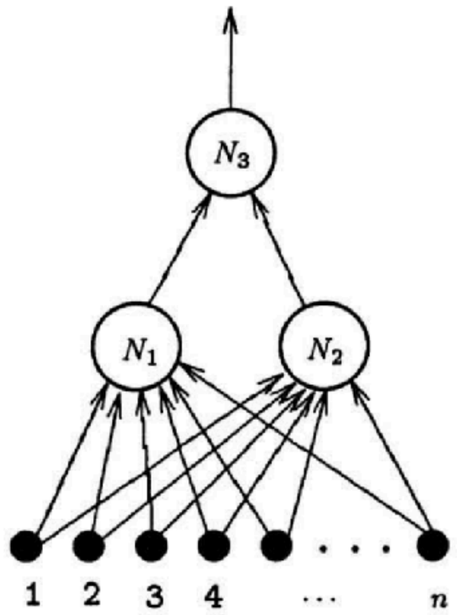

Figure 1: The three node neural network.

Given: A set of *O(n)* training examples on *n* inputs

Question: Do there exist linear threshold functions such that the three-node network fits the training set?

NP hard

😱😱😱

Training a neural network in practice

Forget you can’t do it, and do it anyway!

We use **gradient descent** (& its friends: stochastic, momentum)

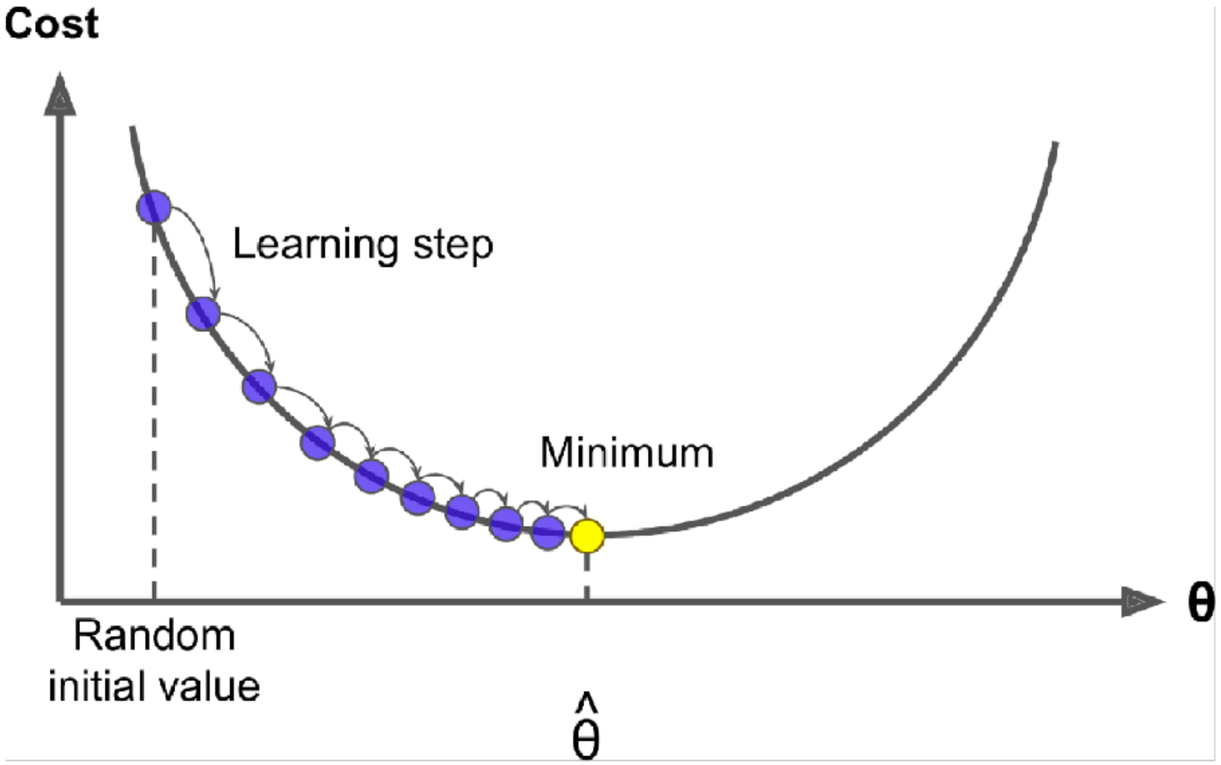

Training a neural network in practice

Forget you can’t do it, and do it anyway!

We use **gradient descent** (& its friends: stochastic, momentum)

training loss

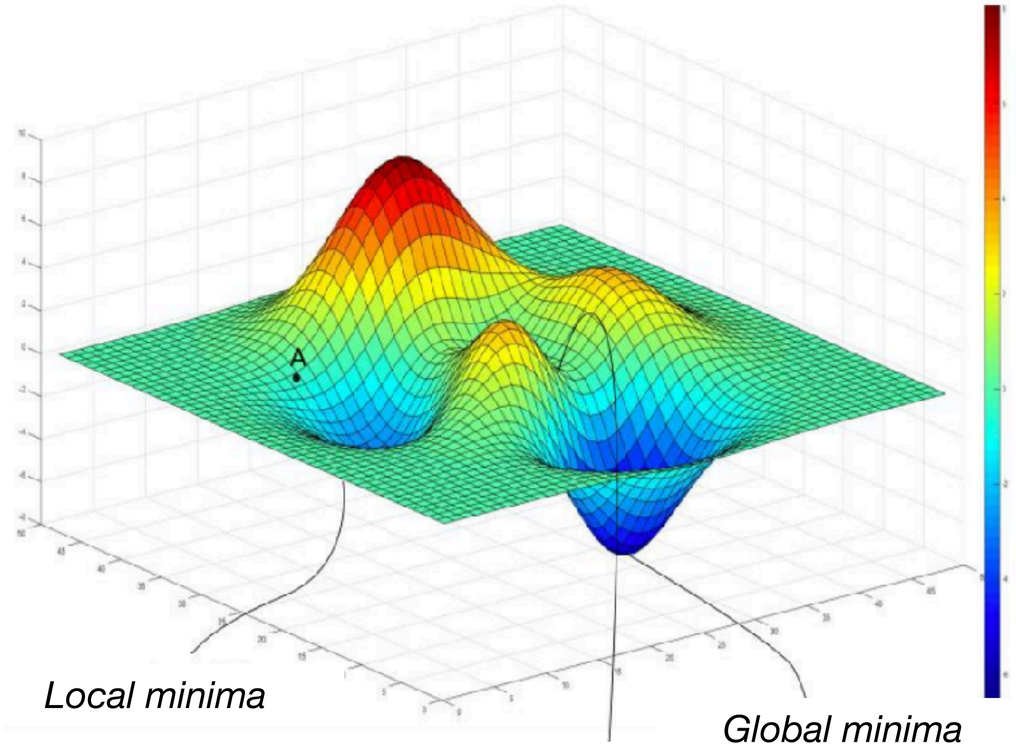

Local minima

Saddle points
(gradient = 0,
but not local minima)

Global minima

Cartoon by Andrej Risteski

Can we learn anything with gradient descent?

ARE EFFICIENT DEEP REPRESENTATIONS LEARNABLE?

Maxwell Nye
Massachusetts Institute of Technology
mnye@mit.edu

Andrew Saxe
Harvard University
asaxe@fas.harvard.edu

- **Parity function:** 1 if a binary string has *odd* number of ones.
  - (0, 1, 1, 1) -> 1
  - (1, 1) -> 0
  - etc.

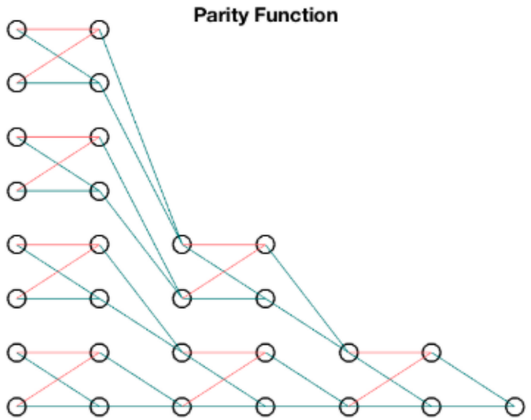

The parity function can be **expressed exactly** by this neural network

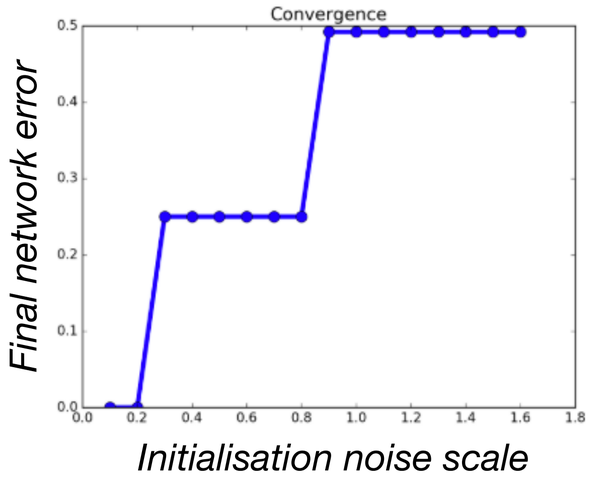

Gradient descent struggles to learn this function…

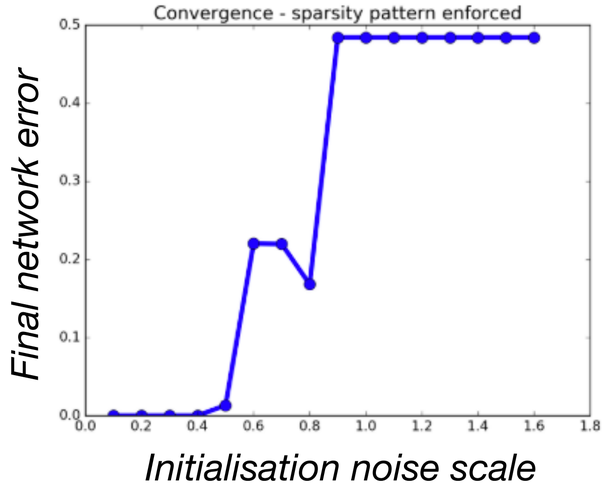

… even if we explicitly impose the sparsity of the target network.

What can GD learn, and why does it work on real data?

A first summary

What are neural networks?

- A **Neuron** can be represented as a single number, nonlinearity(weighted sum of inputs).
- Many neurons yield a neural **network**.
- The ordering & wiring of neurons determines the **architecture** of the neural network.
- The weights of a neural network are found by **training** the network (using stoch grad desc)

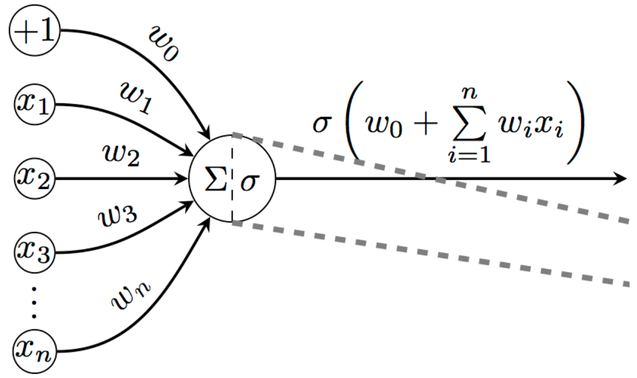

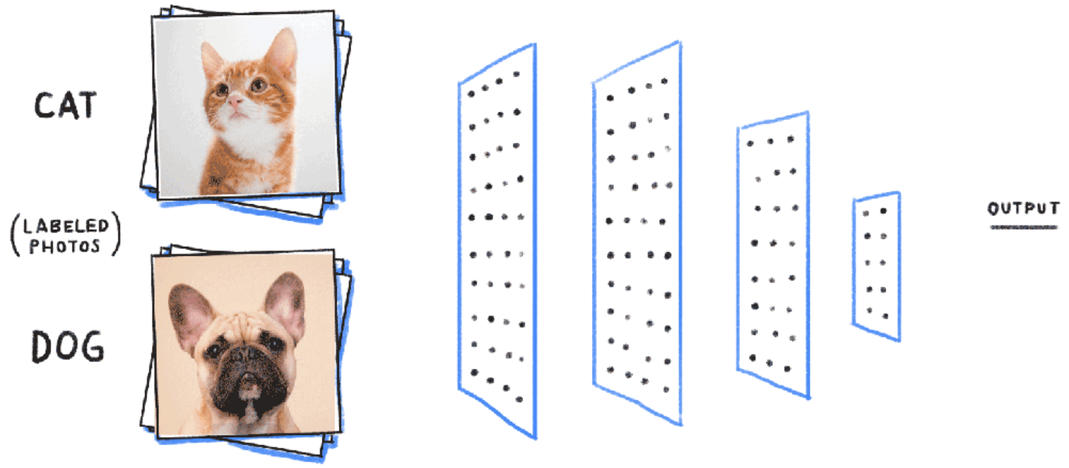

Part II

Feature learning

a.k.a. the *secret sauce* in neural networks

What can a single neuron learn?

Turns out, it's rather limited

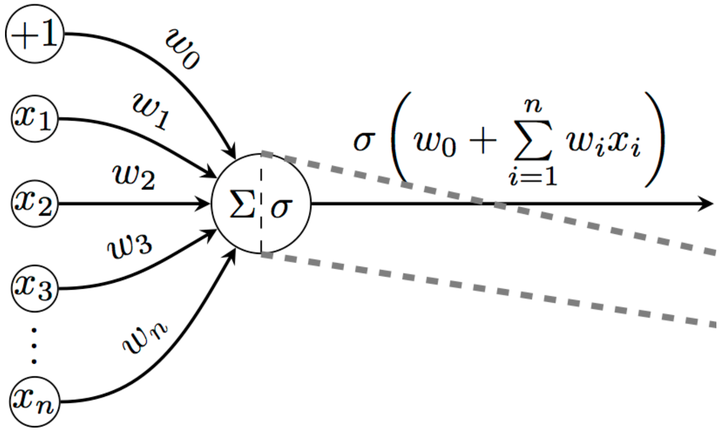

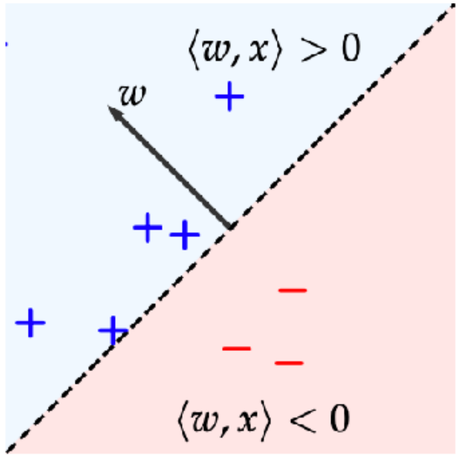

Minsky and Papert (1969)

😱😱😱

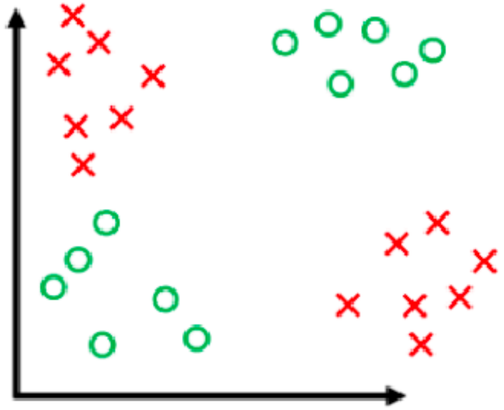

Illustrations by Petar Veličković
https://github.com/PetarV-/TikZ
and Hardt & Brecht (2021)

What can a perceptron learn?

Your turn

Three data sets for binary classification (two classes):

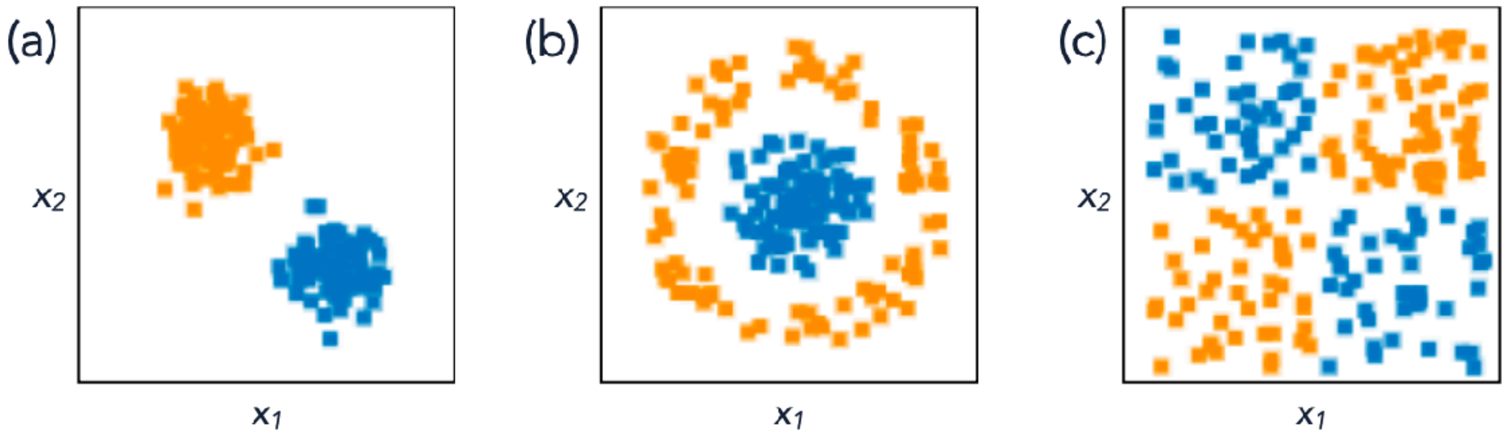

playground.tensorflow.org

Feature learning: research questions

A summary

- Neural networks discover the right **features** directly from their data.
- What are those features? How are they related to the **data statistics**?
- How does the **dynamical** process of learning unwrap the statistics of inputs?
- Relevant for machine learning and **neuroscience**!

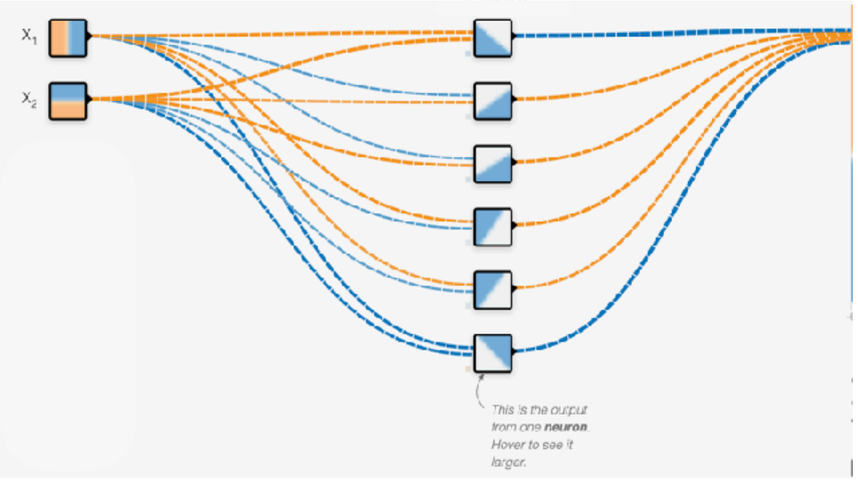

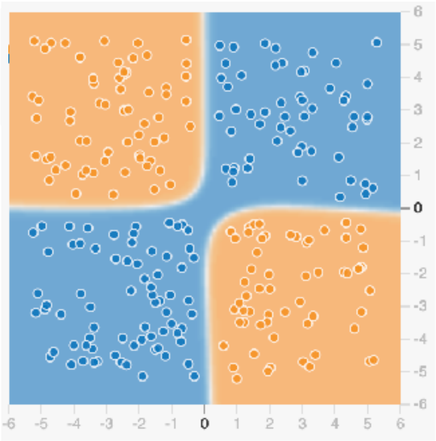

Acknowledgements

Ezter
Data science

Federica
Data science

Sara
Data science

Lorenzo
Math

Riccardo
Data science

Sara
Neuroscience

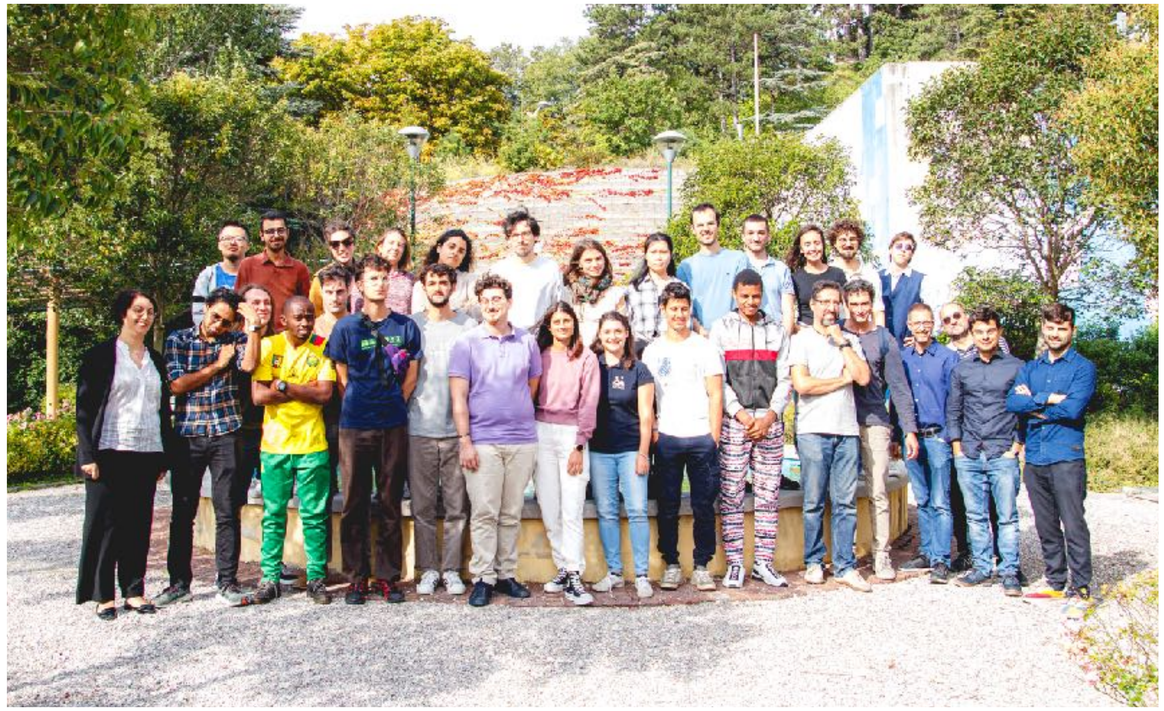

SISSA data science group (Oct 2023)

datascience.sissa.it

@sebastiangoldt