


Computers, brains and AI

4: Deep Learning



1) Intro

2) Foundations

3) Silicon



4) Deep learning

Brains and neural networks

5) Deep learning

How?

6) ChatGPT

7) ChatGPT

8) ChatGPT

9) The future of AI

10) The future of AI



1960




GOFAI

ELIZA
chatbot



Machine learning

Perceptron



1980

Good old fashioned AI


Easy
is hard



Birds fly



2000

 Deep Blue



2012 was the “Big Bang” when
machine learning really took off.

AlexNet



2012

2020



Chess
Jeopardy
Book / CC0

Swan from pxhere / CC0

Bus by Martin49 / CC BY 2.0

Cat by Von.grzanka / CC BY-SA 3.0

Dog by Kyfishfishfish / CC BY-SA 4.0

The thinker by CrisNYCa / CC BY-SA 4.0

Tractor by Kharkiv Tractor Plant / CC BY 4.0

Penguin by Christopher Michel / CC BY 2.0


Neurons by UC Regents Davis campus / CC BY 3.0

Duckling by Luis Miguel Bugallo Sánchez / CC BY-SA 3.0

Fei-Fei Li



In her excellent [TED talk](#), Fei-Fei Li explained how people used to tell computer algorithms that a cat has a round face, a chubby body, two pointy ears, and a long tail. But then in 2007 she realized that no one tells a child how to see, especially in the early years. Instead, they learn from data, and, with that in mind, she created the ImageNet dataset of 14 million **labelled** images.



ImageNet (2009)

14 million **labelled** images

Free

To generate some interest, she created the ImageNet Challenge, a contest for classifying images.

ImageNet Challenge

Classifying images

https://www.ted.com/talks/fei_fei_li_how_we_re_teaching_computers_to_understand_pictures





ImageNet Challenge


The ImageNet challenge was for computer programs to correctly classify images from part of the ImageNet dataset.

In 2011, the best programs, despite decades of very clever research, could only get 75% correct. Then, in 2012, AlexNet, a much simpler program, won by a huge margin, ushering in the modern era of deep learning.


Traditional
computer
vision

AlexNet beat the Good Old Fashioned AI (GOFAI) approach.

In the following years, everyone switched to deep learning, and soon outperformed humans (at this particular contest).



AlexNet



Human level



1960



GOF AI


ELIZA
chatbot



1980

Good old fashioned AI

Easy
is hard



Deep Blue




2000

2020



Machine learning

Perceptron



Birds fly



AlexNet



2012

Chess

Jeopardy

Book / CC0

Swan from pxhere / CC0

Bus by Martin49 / CC BY 2.0

Cat by Von.grzanka / CC BY-SA 3.0

Dog by Kyfishfishfish / CC BY-SA 4.0

The thinker by CrisNYCa / CC BY-SA 4.0

Tractor by Kharkiv Tractor Plant / CC BY 4.0

Penguin by Christopher Michel / CC BY 2.0

Neurons by UC Regents Davis campus / CC BY 3.0

Duckling by Luis Miguel Bugallo Sánchez / CC BY-SA 3.0



bus

AlexNet (2012)



dog



cat




tractor



cat

...



cat

After being trained on millions of labelled images, AlexNet was somehow able to generalize the idea of, say, a cat and recognize one it had never seen before.



AlexNet (2012)

157,000
citations

Alex Krizhevsky


Ilya Sutskever

Geoffrey Hinton


These are the three authors of the AlexNet paper:

https://proceedings.neurips.cc/paper_files/paper/2012/file/c399862d3b9d6b76c8436e924a68c45b-Paper.pdf

Geoffrey Hinton (the Nobel prize-winner), Ilya Sutskever (one of the founders of OpenAI), and Alex Krizhevsky.



This was meant to be a demo that it is hard to recognize images through a small aperture.




Layers

Yann LeCun

LeNet 1988 5

LeNet was a commercially successful deep learning system for recognizing hand-written digits on cheques.



DanNet 2011 ?


DanNet was nearly as good as AlexNet, but not good enough to attempt the ImageNet challenge.

AlexNet 2012 8

ResNet 2015 152

ResNet is short for Residual neural network.


The number of layers in ResNet is far beyond that of the human visual system. We know this because a chain of 152 neurons would take too long to respond.



AI risk sceptic



brains and AI







brain




The early aircraft were somewhat like birds, but early computers (like Colossus, shown below), were nothing like brains, despite books like “Giant Brains”. Artificial neural networks, on the other hand, are distinctly brain-like.




computer

artificial neural
network





<https://www.nature.com/articles/d41586-024-01387-9>



<https://blog.google/technology/research/google-ai-research-new-images-human-brain/>

<https://h01-release.storage.googleapis.com/gallery.html>

In this gallery, you can see 3D
pictures of bits of human brain.



Camillo Golgi




Golgi accidentally discovered a method of staining brain tissue that somehow stained some of the neurons in their entirety and left most of them unstained. To this day, it is still not completely clear why this works.




<https://youtu.be/F64cnuo0oNU?t=3>

Santiago Ramón y Cajal



Cajal spent a large part of his life drawing beautiful pictures of neurons in all their glory.



Golgi's method
1873


Nobel Prize - 1906

They were jointly awarded the Nobel prize

Neuron by Bob Jacobs / CC0


Golgi by unknown author / CC0

Cajal by Clark University in 1899, restoration by Garrondo / CC0



Neocortex

The neocortex is the “grey matter” (even though it looks more pink than grey). The shape and size is that of a crumpled-up tea towel.




White matter


Long-range
connections




tea towel




These are some of the long-range highways in the brain, made of “white matter”. Everything looks white because the long axons are covered with fatty myelin sheaths.





This is a reconstruction of the neocortex, showing the six layers. The neocortex is made of repeating structures called cortical columns, regardless of whether they are part of the visual cortex or the motor cortex or whatever.



Thousand brains theory

How similar brains are

[Neocortex](#) by Bodysurfinyon / CC BY-SA 4.0



[Myelin](#) by Dr. Jana / CC BY 4.0

[Tracts](#) by Thiebaut de Schotten M, Dell'Acqua F, Ratiu P, et al. / CC BY 4.0

[Brain](#) by National Institutes of Health / CC0

[Cerebrum](#) by Database Center for Life Science / CC BY-SA 2.1 jp

synapse




Artificial
Neural Network

Neural Network

Neural Net

Artificial neural networks are simplified versions that try to capture the essence of their biological counterparts.



axon away



rabbit

Comparing pixels with far away pixels doesn't help.

The pixels are the first layer.

It makes more sense to compare a patch of nearby pixels. For example, we can see some texture inside this black square...

...and here we can see an edge.

This is the second layer.

This is a map of where there is texture.

This is the third layer.

This is a map of where there is a boundary between some texture and some non-texture.





This is a map of where there are vertical edges.

These maps are hypothetical examples of how features can be picked out from earlier layers.


During training, AlexNet, ends up deciding what features there are.

In AlexNet there are eight layers, ending with an output layer that, with any luck, says we are looking at a rabbit.





There might be dozens of maps in each layer.



Three colour channels



More channels



...

...

...




1,000,000,000

~ billion neurons

1,000,000,000,000

~ trillion connections

This is a cartoonish view of the, maybe, one billion neurons that are involved in instantly recognizing objects, part of the 86 billion neurons we are all born with.



synapse

slow neuron
a miniature computer

parallel

scale



cat



https://en.wikipedia.org/wiki/Grandmother_cell

grandmother


Jennifer Aniston



The neurons make up for their slowness (a one-hundredth of a second response time) by running in parallel.

weight

increasing complexity



Very roughly, part of our vision system can be thought as a series of layers with the retina at the front and the (hypothetical) grandmother cells at the back that say what we are looking at.

Babies learn to see the world by (somehow) setting the strengths of the maybe trillions of connections.



perceptron

Mark I Perceptron

1958

The Mark I Perceptron was an early artificial neural network.



The proposal was amazingly ambitious...

... but the reality was far more modest, as shown by it trying to recognize the letter C.

The reality

The proposal

“concept formation ...
... inductive logic”



<https://www.google.com/search?udm=2&q=mark+i+perceptron>


(image search)



The main reason for the feeble results was the lack of scale. It looks like it had thousands of wires, unlike a brain, which has 150 trillion synapses.


Mark I Perceptron

1958



Mark I Perceptron

1958



Since the Mark I Perceptron, all artificial neural networks have been simulated in software.




Simulation

(calculations)

Physical model



decisions



This mouse is deciding whether to go towards the food or the water.




food ✓



It must decide one way or the other.

Trying to avoid the decision by going in an average direction doesn't help.



All-or-nothing law

Individual neurons make decisions all the time, namely, whether to fire or not.

Either it fires...

... or it doesn't

Output

This graph shows how the neuron responds to how much input it has.

It demonstrates the "all-or-nothing" principle.



binary ?

Input

However, neurons can fire frequently if they get more input.



Firing rate

This graph shows how the output (how often the neuron fires) depends on the size of the stimulus.

Input

analog ?







Activation functions

In artificial neural networks, an activation function says how the neuron responds.


3. This activation function was popular before AlexNet.



1. This activation function has a simple threshold, but is bad for AI because the learning process needs slopes. (We will see more about slopes in the next lectures).




nonlinear



linear

2. This activation function has slopes but is bad because it isn't making decisions.



ReLU

Rectified linear unit

4. This activation function was used in AlexNet.


$$x_1 + x_2 + \dots + x_n$$

The result is a **weighted sum**. Artificial neural networks spend most of their time calculating these weighted sums.


$$w_1 \times x_1 + w_2 \times x_2 + \dots + w_n \times x_n$$

Different synapses can have different strengths, for example, by having different numbers of receptors.

Synapse




We are denoting the strengths as w_1, w_2 etc.



An input of 4 from another neuron

A weight of 5





$$x_1 + x_2 + \dots + x_n$$

The **bias** is the base-level excitability of the neuron. During deep learning, the biases are adjusted along with the weights.

$$w_1 \times x_1 + w_2 \times x_2 + \dots + w_n \times x_n + \text{bias}$$

Synapse





$$x_1 + x_2 + \dots + x_n$$

The **bias** is the base-level excitability of the neuron. During deep learning, the biases are adjusted along with the weights.

$$w_1 \times x_1 + w_2 \times x_2 + \dots + w_n \times x_n + \text{bias}$$

Synapse



weights
biases } parameters

The weights and the biases are the **parameters**, the things that can adjusted during training.



Input


$$w_1 \times x_1$$


$$w_2 \times x_2$$

$$w_n \times x_n$$

bias

As we saw earlier, for anything interesting to happen, the neuron (either biological or artificial) cannot simply pass the total input to its output. Hence the **activation functions**, we saw previously and which we will look at on the next slide.







The diagrams show the ReLU activation function keeping the 26 as 26, but zapping the -14 to zero.

In other words, ReLU is making a simple decision.

Unlike in good old-fashioned AI, neural networks don't have explicit logic statements (e.g. if it's hairy and has a tail then it's a dog). Instead, the decisions are tucked away inside the neurons.





Net




650,000

~ billion neurons

60 million

~ trillion connections



cat



Compared to the baby, AlexNet is tiny. It has 650,000 neurons, which is fewer than that of a honeybee.


Somewhat confusingly, AlexNet (like many artificial neural networks) does a thing called weight sharing with 60 million parameters but 360 million connections. For learning and training, the 60 million is the important number.



AlexNet

Very roughly, part of our vision system can be thought as a series of layers with the retina at the front and the (hypothetical) grandmother cells at the back that say what we are looking at.

Babies learn to see the world by (somehow) setting the strengths of the maybe trillions of connections.
“Neurons that fire together, wire together”.




Sometimes AlexNet isn't sure of the answer. Maybe it is confused by the reflection or by not being able to see the animal's legs. AlexNet, being a simple beast, can't do anything about this. (A human, on the other hand, might squint a bit closer or realize, for example, that cats don't like water).

AlexNet

Real neurons are full of complicated chemistry. For example, neurons have potassium and sodium ions moving across a membrane, and synapses have many types of neurotransmitters.

Artificial neural networks, on the other hand, keep things simple by treating the activation of a neuron (how excited it is) as a number. Similarly, the strength of a connection is just another number.



A further simplification, for the purposes of the ImageNet contest, was to have a fixed set of 1000 categories, which allows AlexNet to have a fixed shape.

Learning

deeply unobvious

... even though learning feels so natural




Usual explanation


... feed ...
... training data ...
... learn

A cactus (!) won't learn anything.

The dog, on the other hand, is a very active part of the training process, including paying close attention through sight, sounds and smells, and trying to read the trainer's mind.



How babies learn




Looking at the world

Exploring



hard work



Making
connections

“Neurons that fire together, wire together”

Hebbian theory – Donald Hebb (1949)

“Use it or lose it”

“Stay in sync or
lose the link”

A million new connections per second

