

Machine Learning

Are machines the next Einstein?

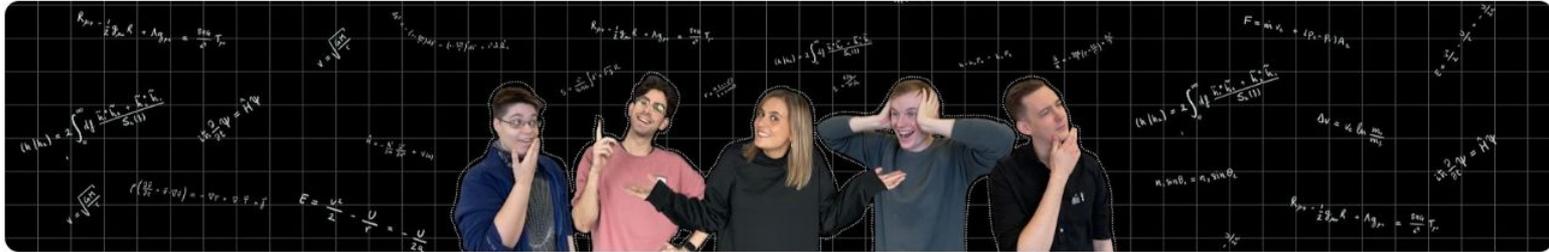
Institute of Engineering and Technology

Susanna Green

“Even though I am doing an Astrophysics PhD... I oscillate between being a software engineer, machine learning engineer, and a data scientist. Plus I am a content creator!”



– Susanna Green –



Field of View

@fieldof_view · 242 subscribers · 34 videos

Talking to scientists to understand not only what they're working on, but what inspired the... >

[instagram.com/physicschat_fieldofview](https://www.instagram.com/physicschat_fieldofview) and 2 more links

Subscribe

- Home
- Videos
- Playlists
- Community

NEW video this Friday!

For you



Creating Planets Made Easy! Physics Chat with Tom Haworth

294 views · 1 month ago



Physics Chat Live! | Stargazing vlog

242 views · 1 year ago



What came first? Relativity or Reality? Physics Chat with Ashim Sen Gupta

370 views · 2 months ago



It's life, Luke, but not as we know it with Luke Booth

172 views · 8 months ago



How do **YOU**
think and **learn**?

“Did you know the average person has about 6,200 thoughts per day”

Psychologists at Queen’s University in Kingston, Ontario, 2020 (Credit: Bigthink.com)

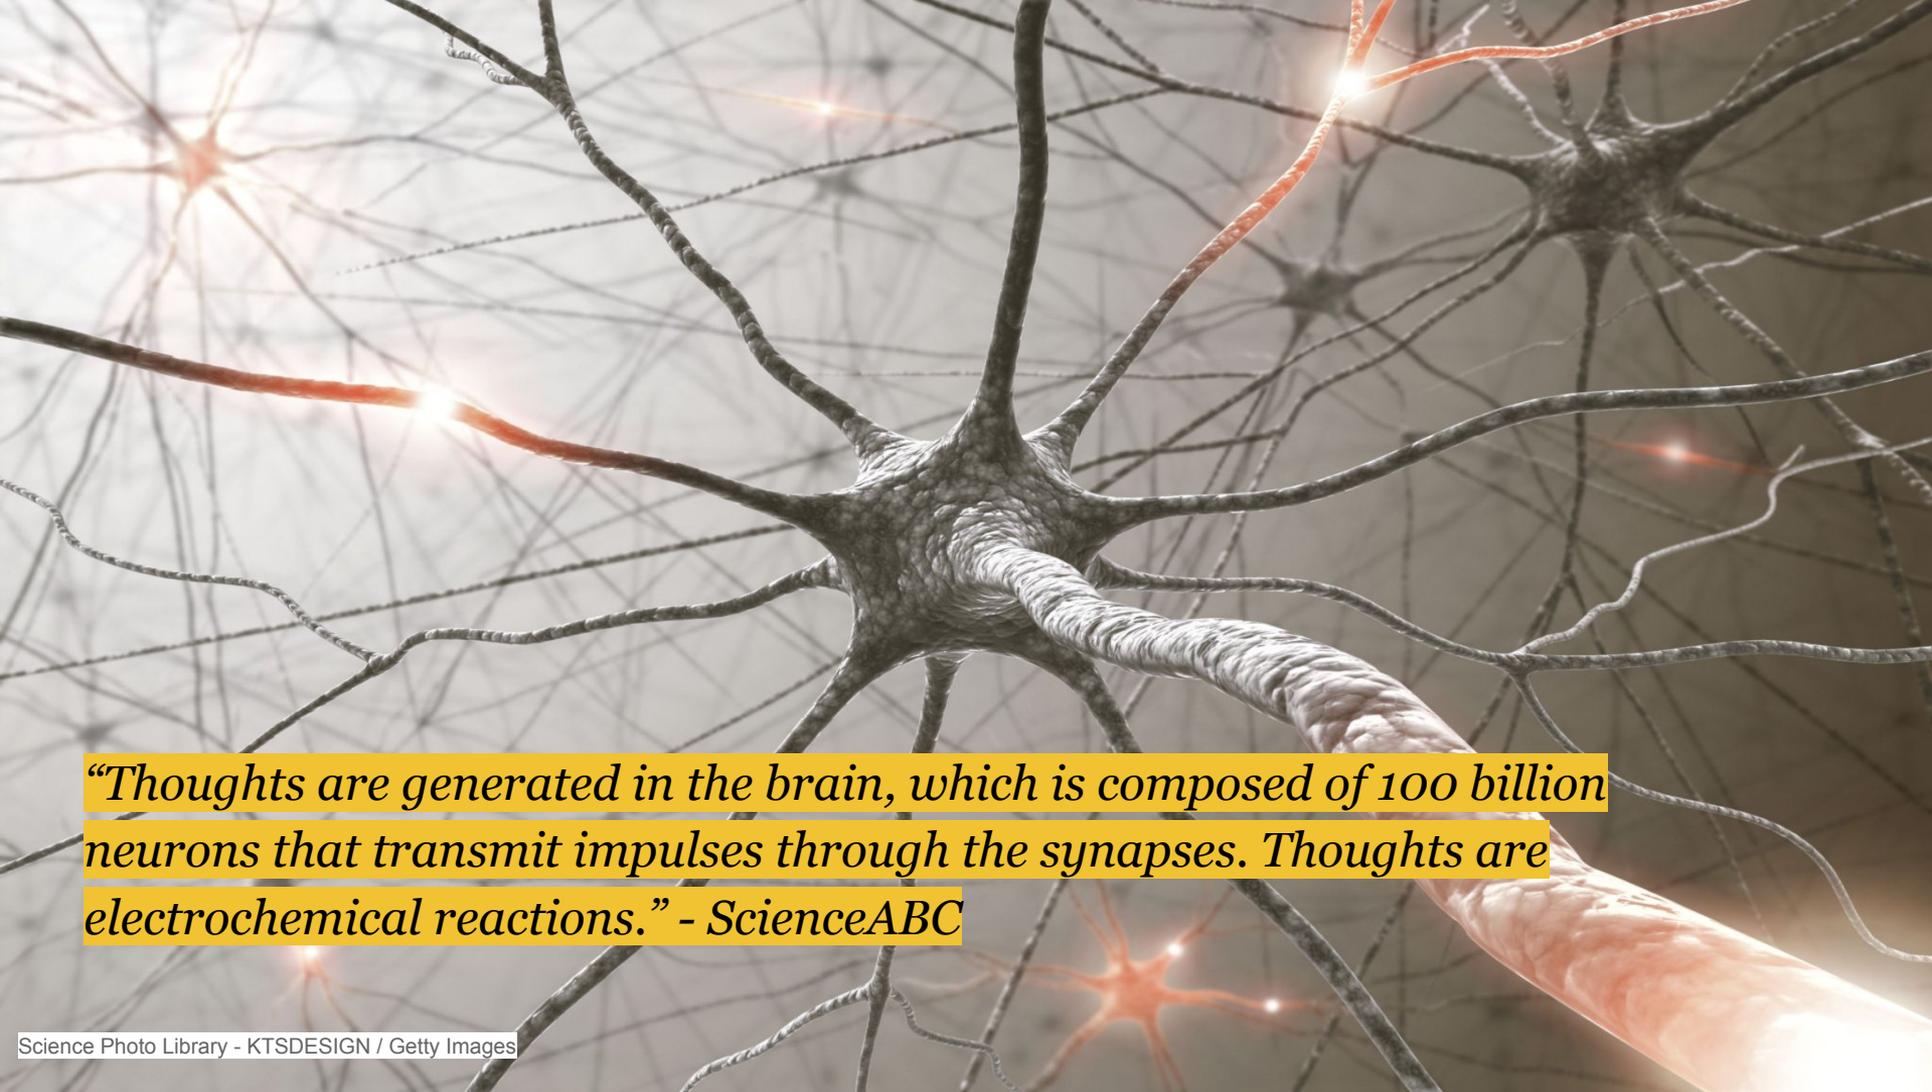
- A. I have more thoughts than this!!!
- B. I have about this many thoughts.
- C. I have very few thoughts a day...





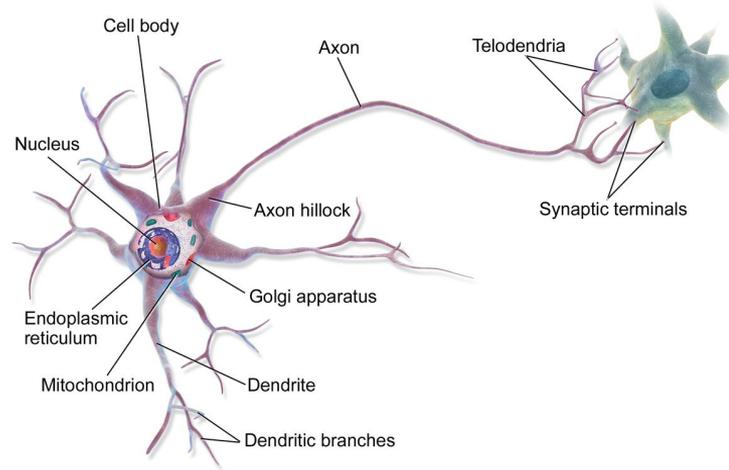
Have you ever thought
about your thoughts?
Like, what even is a
thought?

Like... what are your
thoughts made of?



“Thoughts are generated in the brain, which is composed of 100 billion neurons that transmit impulses through the synapses. Thoughts are electrochemical reactions.” - ScienceABC

Biological neuron



Credit: Wiki

Stimulus

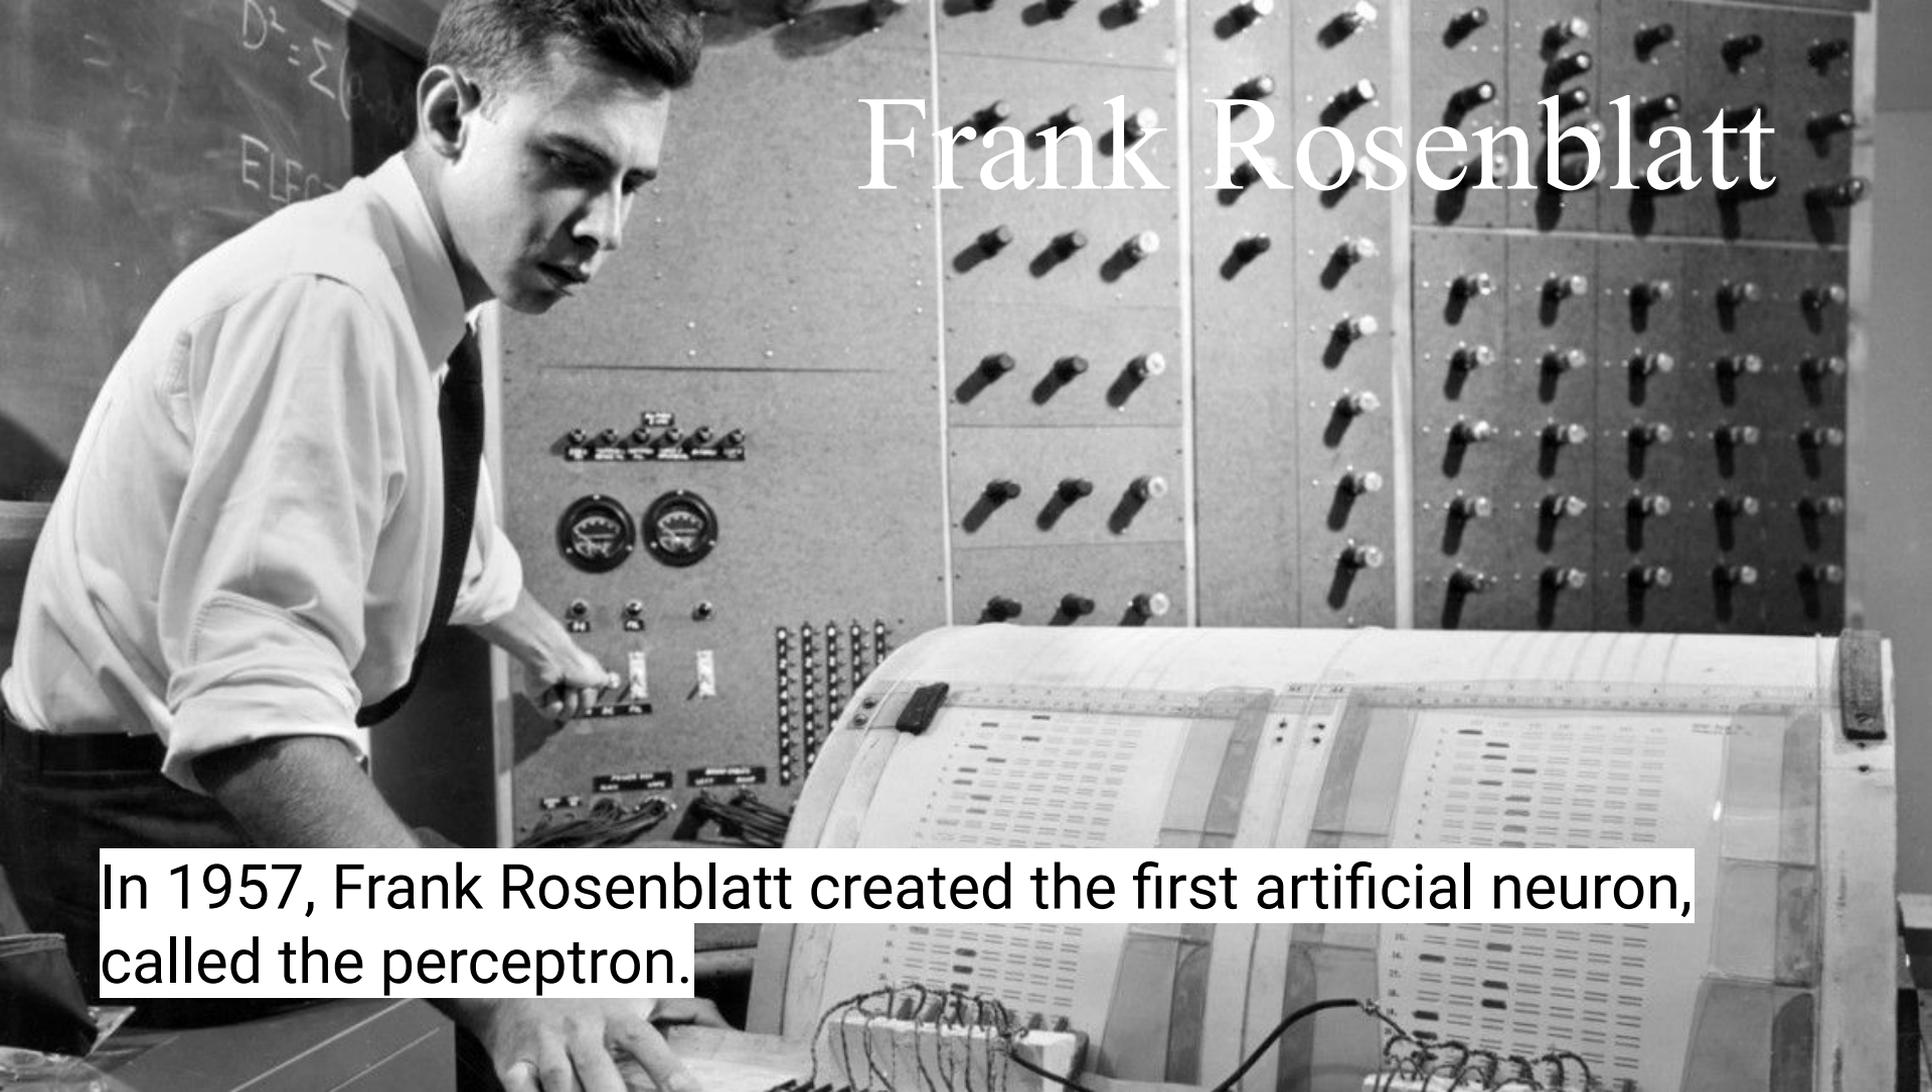


Enough stimulus to
send on signal



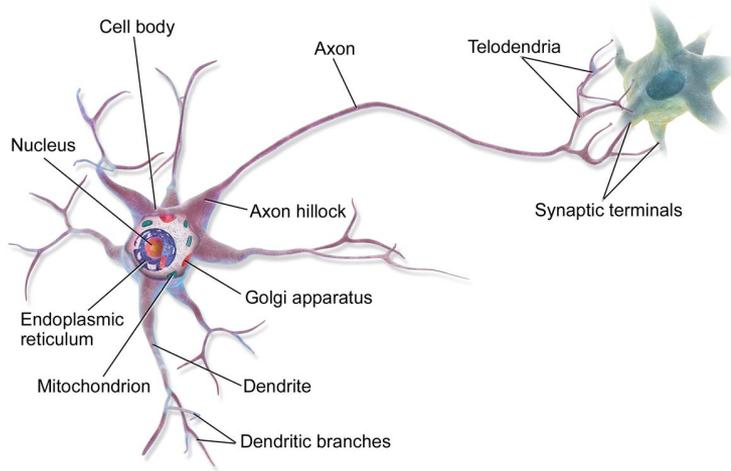
Signal passed onto
the next neuron

Frank Rosenblatt

A black and white photograph of Frank Rosenblatt, a young man in a white shirt and dark tie, leaning over a large, complex electronic device. The device has a control panel with several dials and switches on the left, and a large array of vertical columns of components, possibly relays or switches, on the right. In the foreground, there is a large, curved display or control panel with a grid of small lights or indicators. The background shows a chalkboard with some mathematical notation, including $D^2 = \Sigma$ and the word "ELECT".

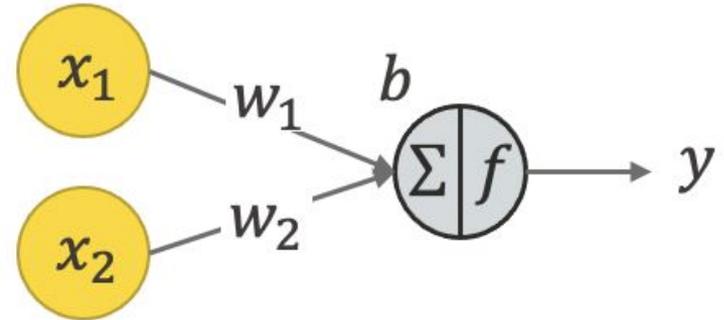
In 1957, Frank Rosenblatt created the first artificial neuron, called the perceptron.

Biological neuron v.s. artificial neuron (perceptron)?



Credit: Wiki

V.S.



Credit: knime.com

Input

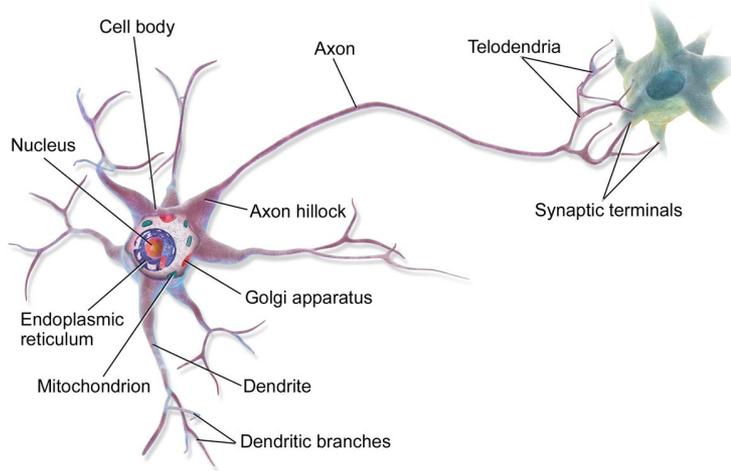


Process



Output

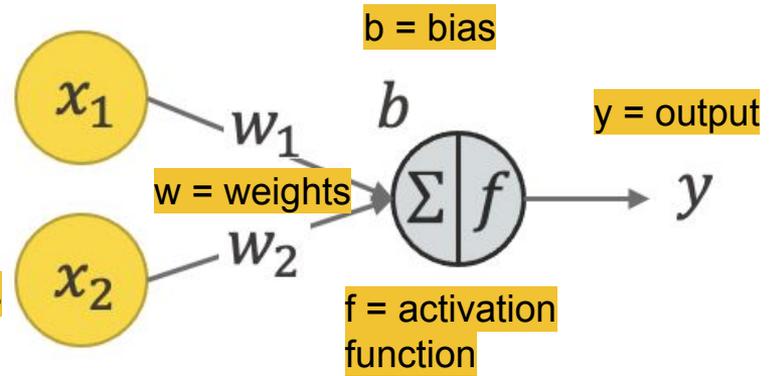
Biological neuron v.s. artificial neuron (perceptron)?



Credit: Wiki

V.S.

$x_1, x_2 =$ inputs



Credit: knime.com

Input

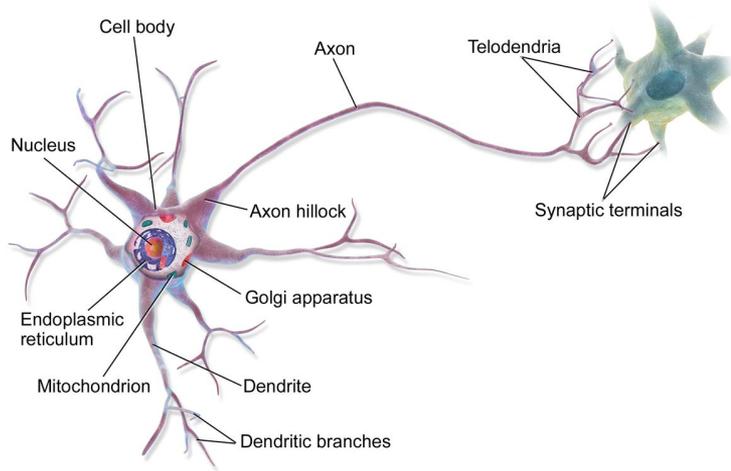


$$f \left(b + \sum_{i=1}^n x_i w_i \right)$$



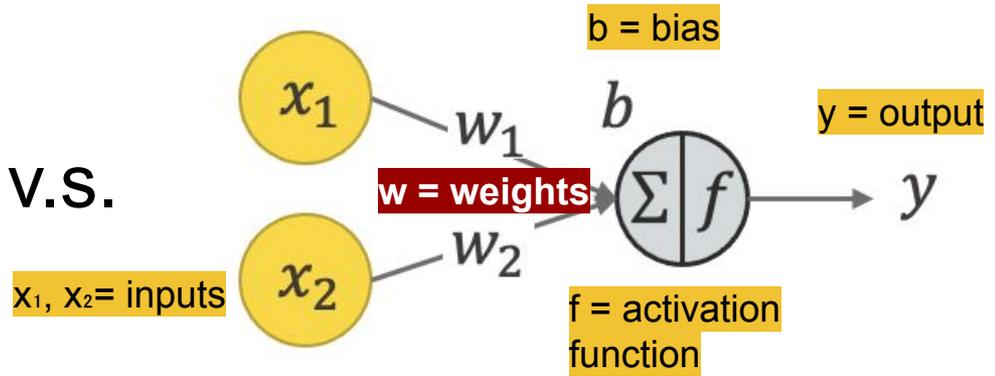
Output

Biological neuron v.s. artificial neuron (perceptron)?



Credit: Wiki

V.S.



Credit: knime.com

Input



$$f \left(b + \sum_{i=1}^n x_i w_i \right)$$

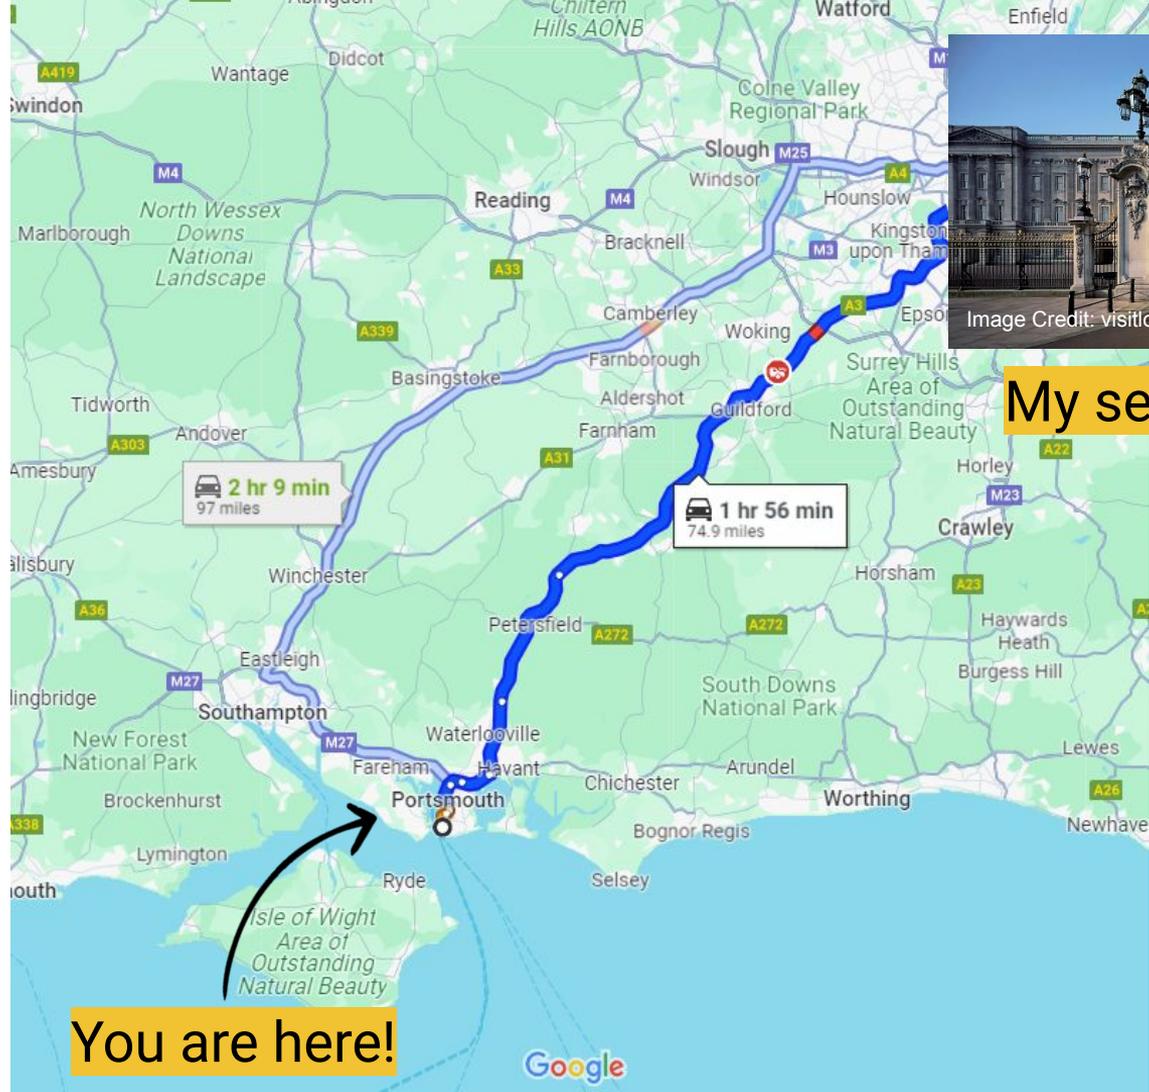


Output



“Neurons that fire together, wire together.”

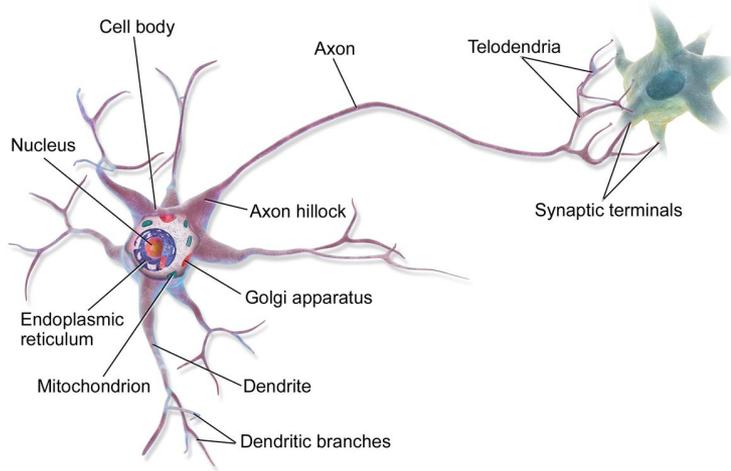
- Donald Hebb -



My second home!

You are here!

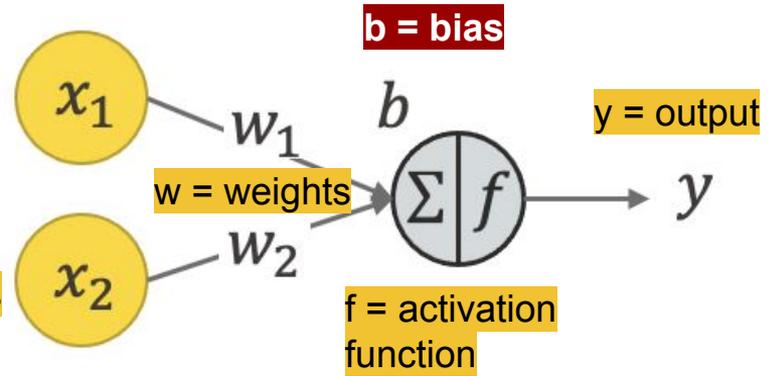
Biological neuron v.s. artificial neuron (perceptron)?



Credit: Wiki

V.S.

$x_1, x_2 =$ inputs



Credit: knime.com

Input

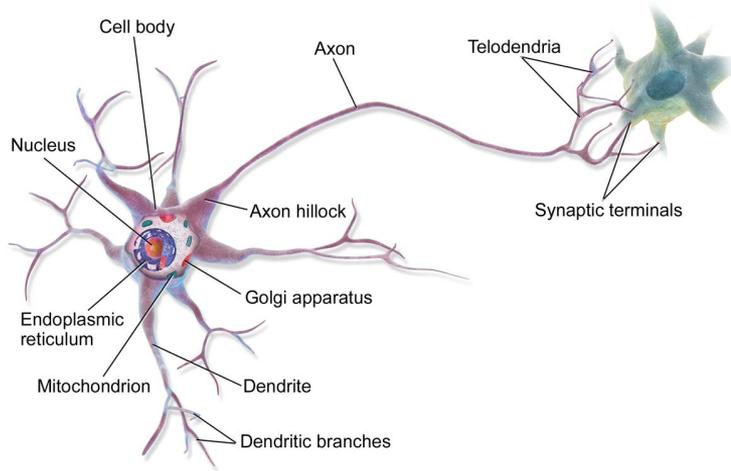


$$f \left(b + \sum_{i=1}^n x_i w_i \right)$$



Output

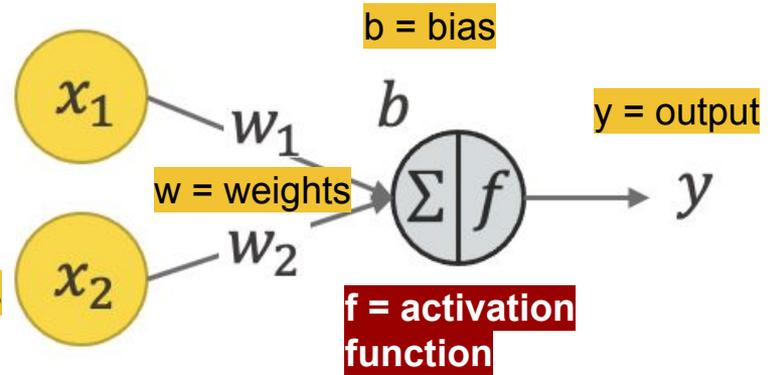
Biological neuron v.s. artificial neuron (perceptron)?



Credit: Wiki

V.S.

$x_1, x_2 =$ inputs



Credit: knime.com

Input

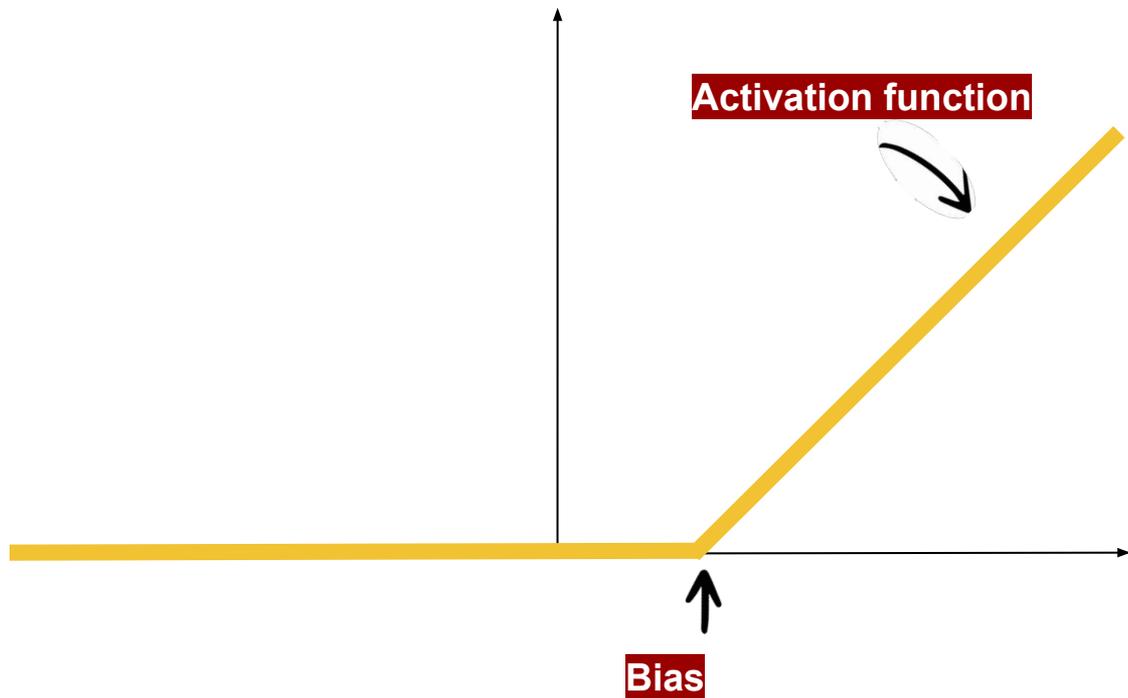


$$f \left(b + \sum_{i=1}^n x_i w_i \right)$$



Output

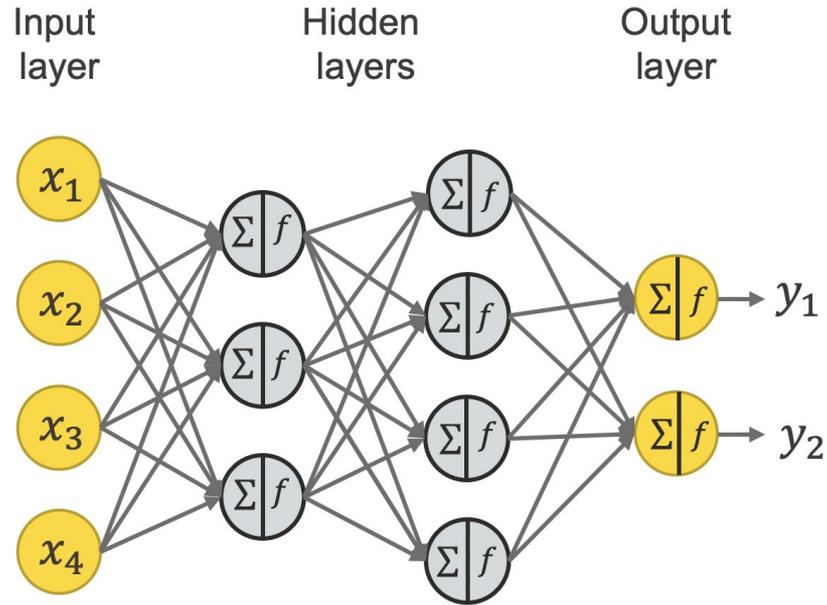
Biological neuron v.s. artificial neuron (perceptron)?



*Humans think when neurons are connected (neural pathway),
similarly machines 'think' when artificial neurons are connected.*



Neural networks are one way that a machine 'thinks' and 'learns'...



Credit: knime.com

YOU

v.s.

MACHINE

Dog



Image by wirestockon Freepik

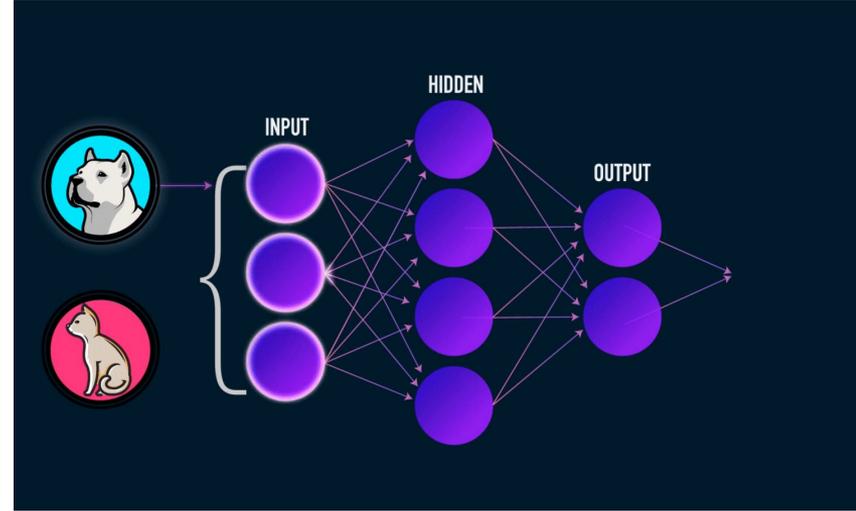
Cat



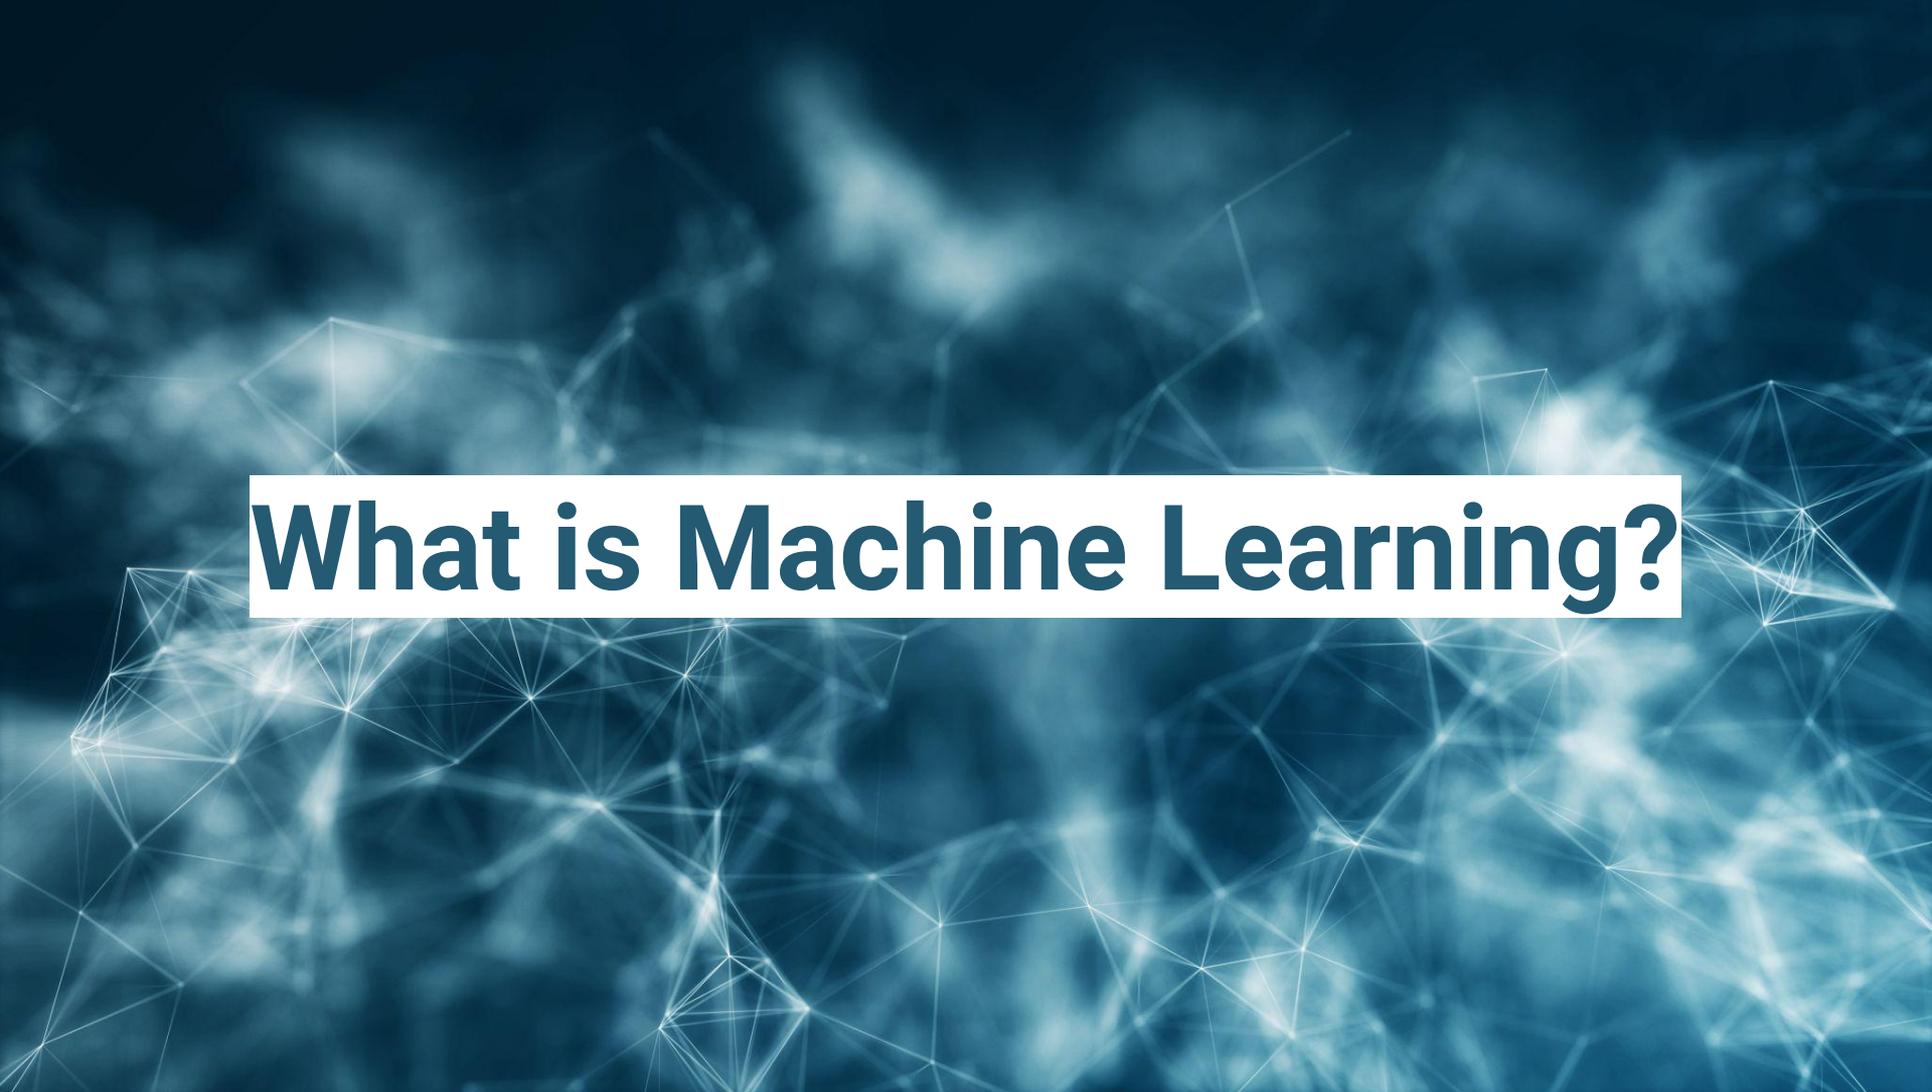
YOU

V.S.

MACHINE



Credit: [upadhyay-shivam.medium.com](https://www.upadhyay-shivam.medium.com)



What is Machine Learning?

The machine learns the pattern rather than you explicitly telling the machine the pattern



V.S.



Humans learn in different ways...



...so do machines.

Machine Learning

```
graph TD; ML[Machine Learning] --> SL[Supervised Learning]; ML --> UL[Unsupervised Learning];
```

Supervised Learning

Unsupervised Learning

Supervised learning you tell the computer what to learn (i.e. a cat or dog) while unsupervised learning you don't tell the computer what to learn.



V.S.

Supervised learning is when you give the computer the ingredients and final dish, then tell it to find a recipe.

Unsupervised learning is when you give the computer the ingredients and ask it to create a recipe and new dish.

Let's use a classification task to explain this...

Dog



Image by wirestockon Freepik

Cat





Supervised Learning



Cat



Dog

Labelled Data



Unsupervised Learning



Supervised Learning



Cat



Dog

Labelled Data



Unsupervised Learning



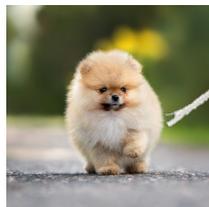
Unlabelled Data



NEURAL NETWORKS

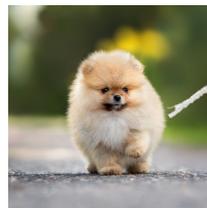


Cat



Dog

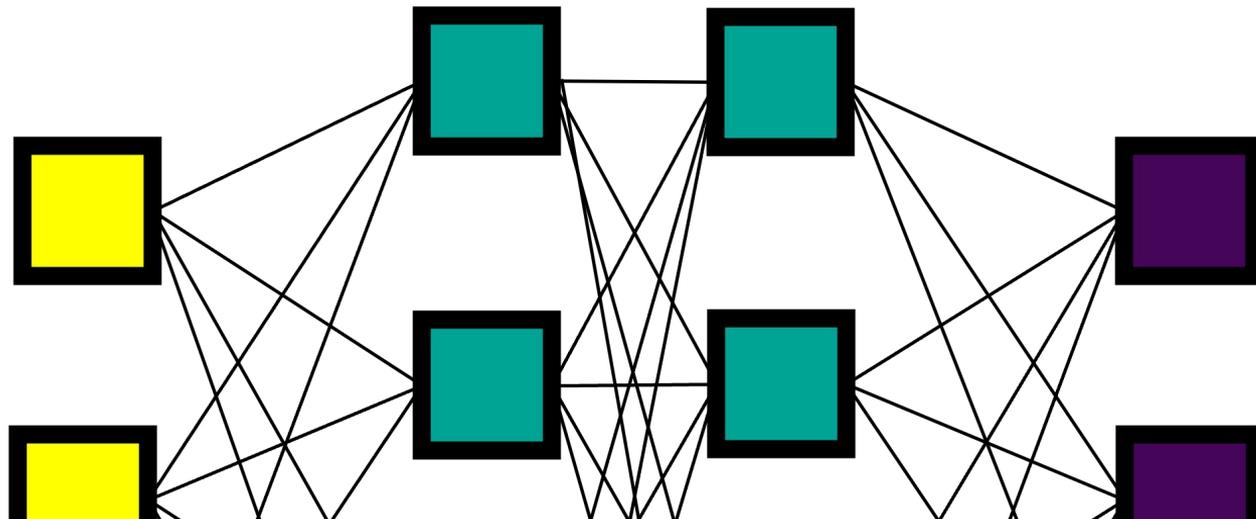
Labelled Data



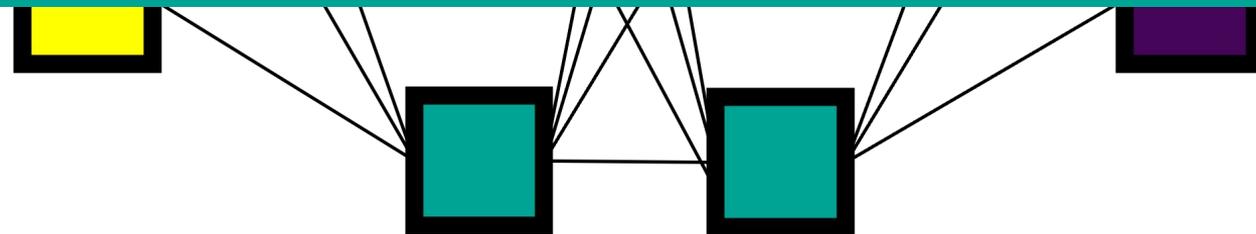
Unlabelled Data



Disclaimer: Both require machine learning engineers to check the results.

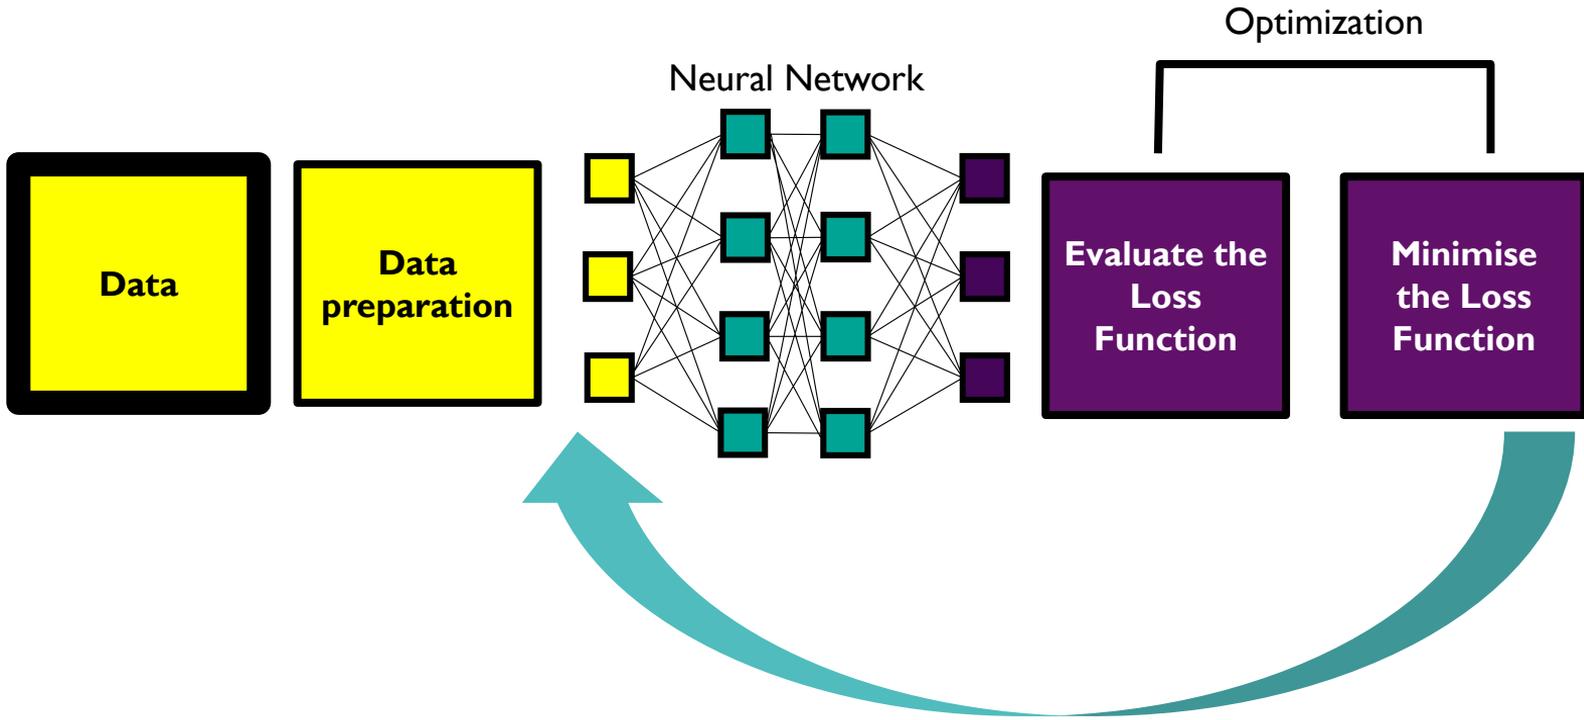


How does a Neural Network 'learn'?





**// Like humans, neural networks learn by their mistakes...
...but we call it 'Optimization'!**



Data

Data preparation

Neural Network

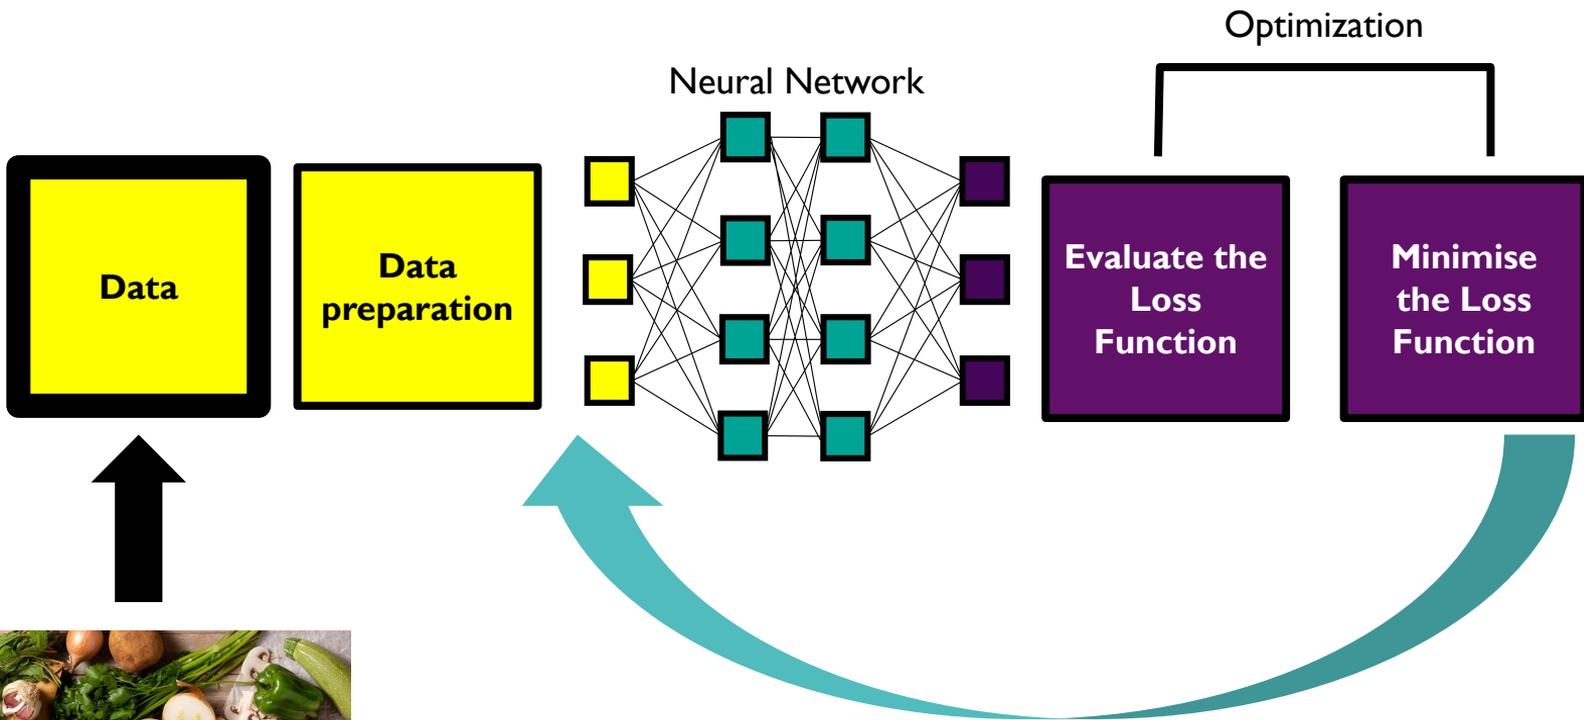
Optimization

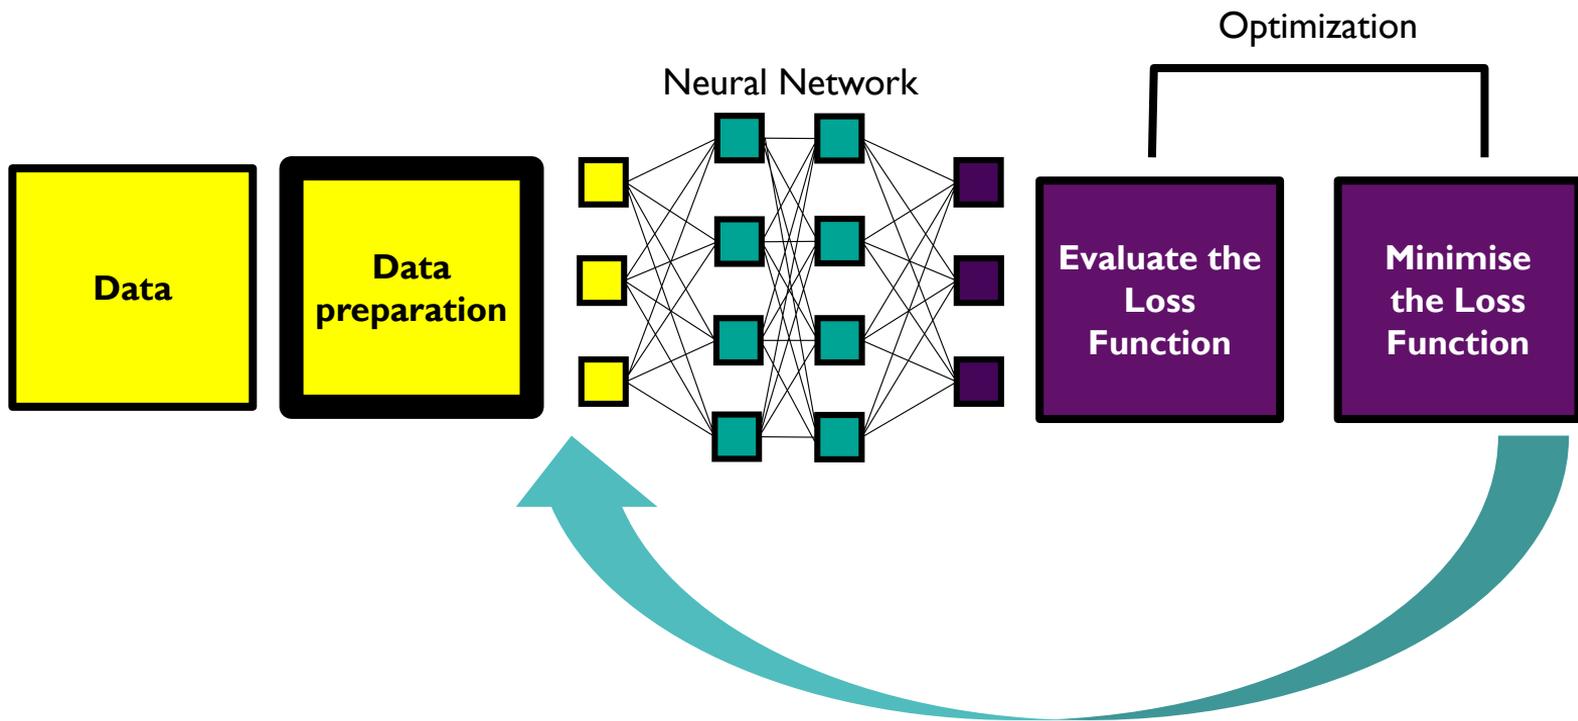
Evaluate the Loss Function

Minimise the Loss Function



Image by pressfoto on Freepik





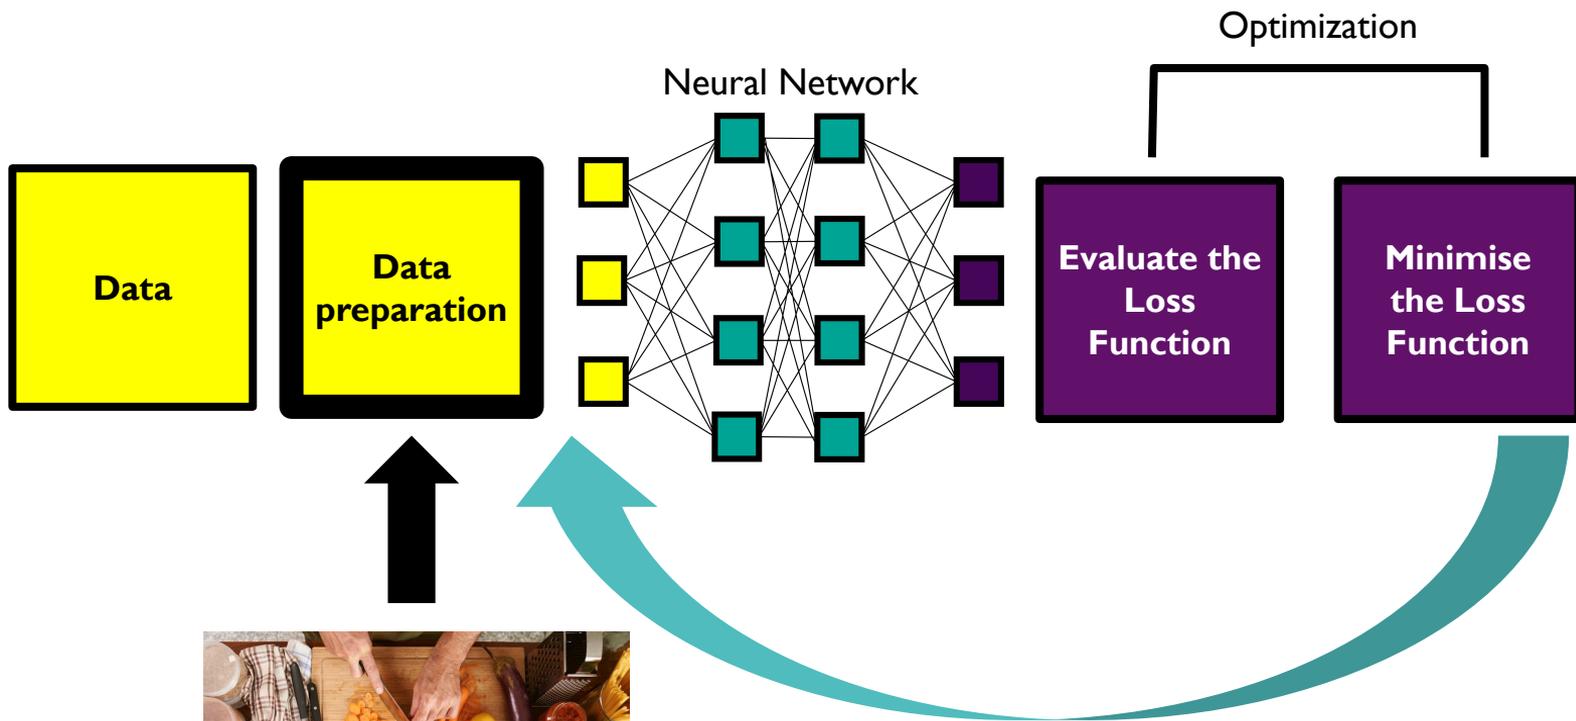
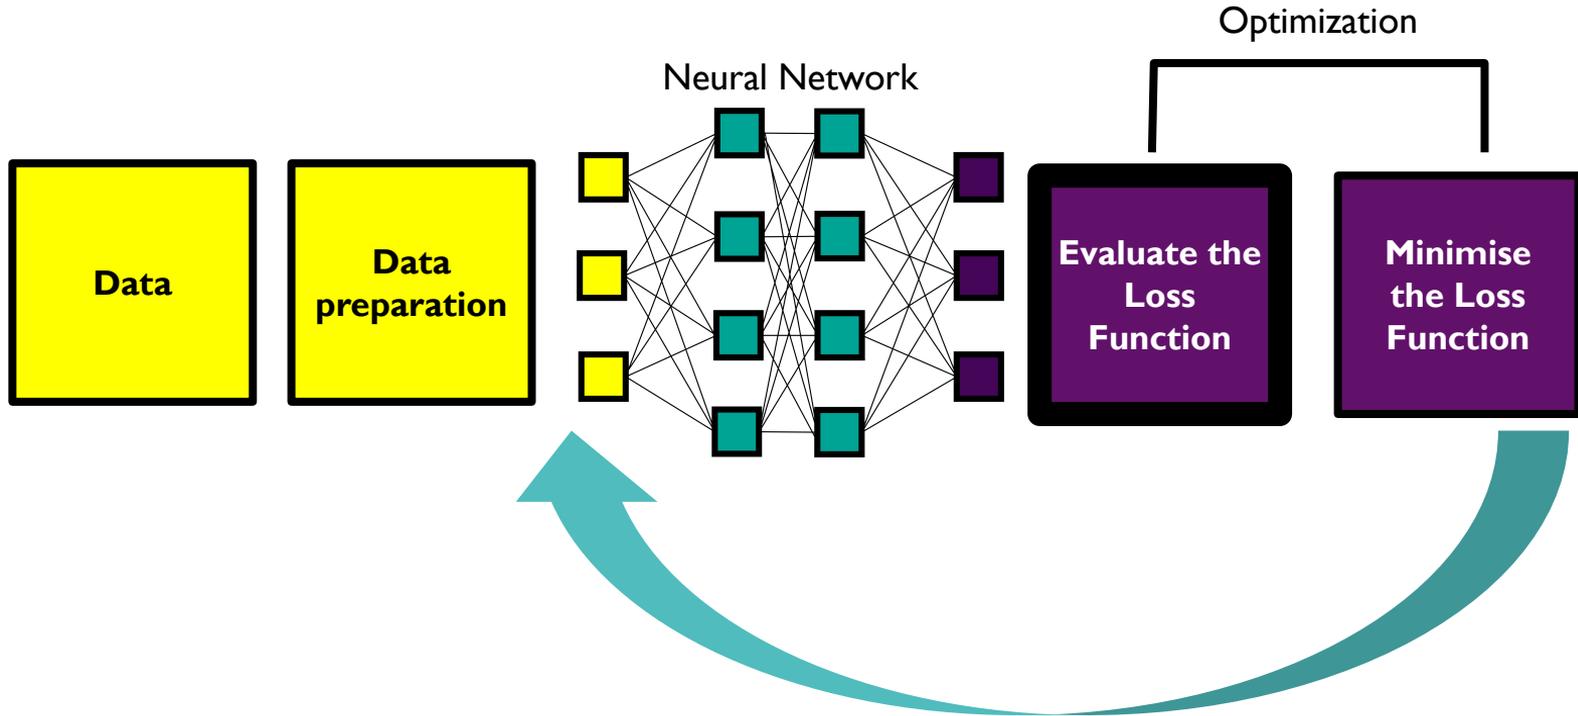
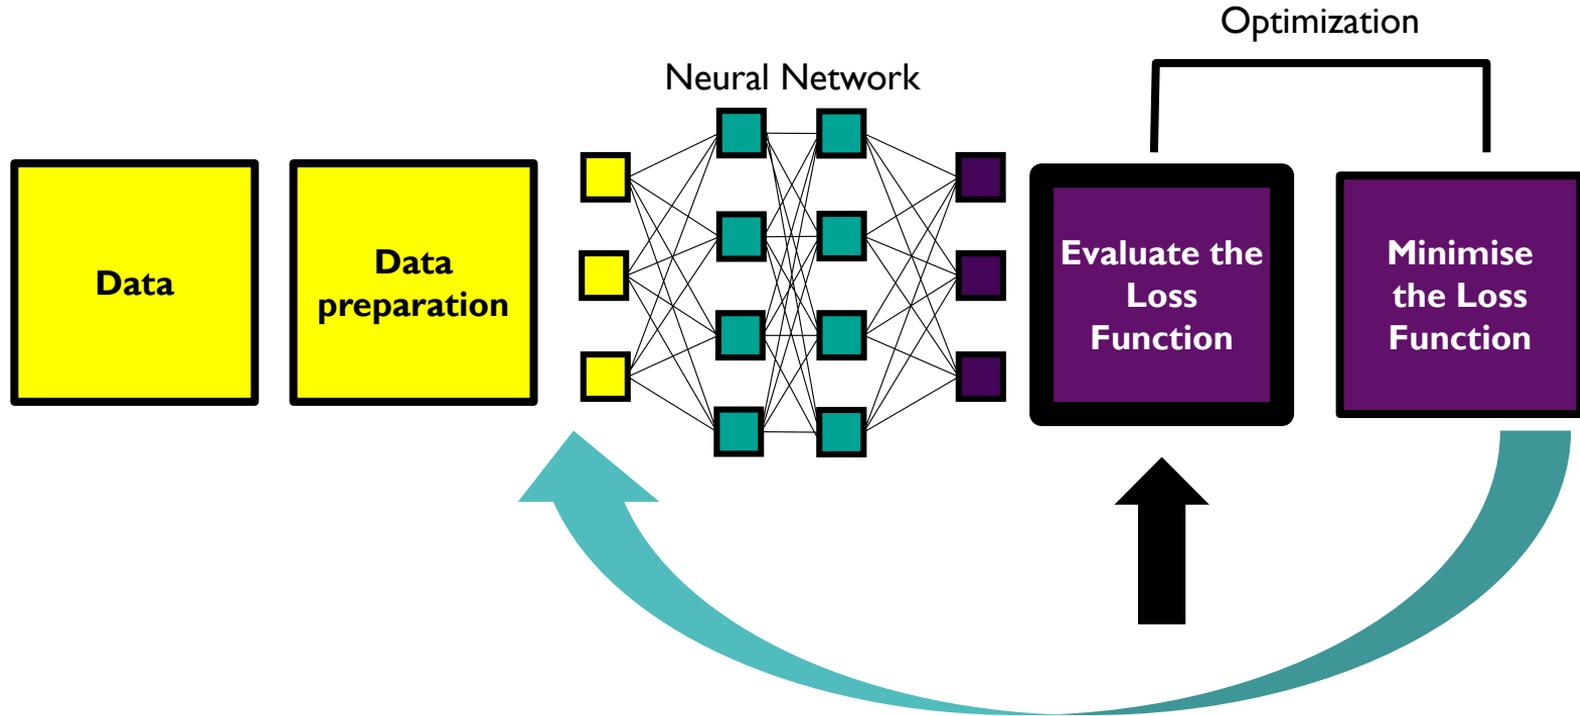


Image by chandlervid85 on Freepik



*“The **Loss Function** is the choice of the error function which is how the optimization algorithm is told how good the model is...”*

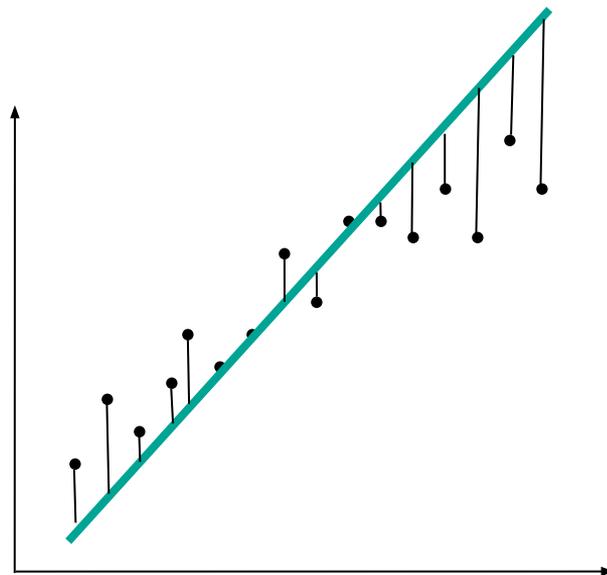


*“The **Loss Function** is the choice of the error function which is how the optimization algorithm is told how good the model is...”*

Optimization: Loss Function

“As a machine learning engineer, it is your responsibility to chose the loss function that best suits your problem.”

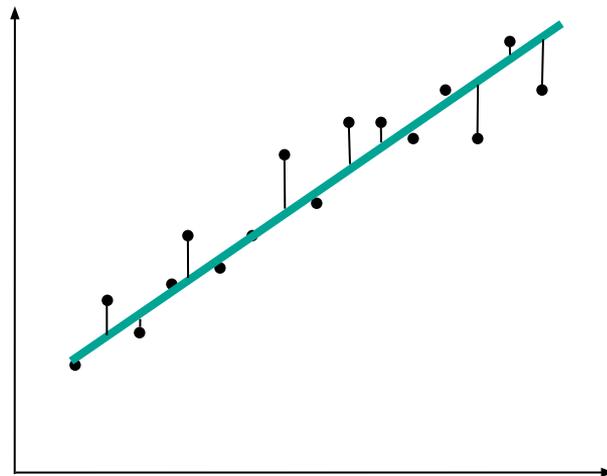
For example:



Optimization: Loss Function

“As a machine learning engineer, it is your responsibility to chose the loss function that best suits your problem.”

For example:



Optimization: Loss Function

“As a machine learning engineer, it is your responsibility to chose the loss function that best suits your problem.”

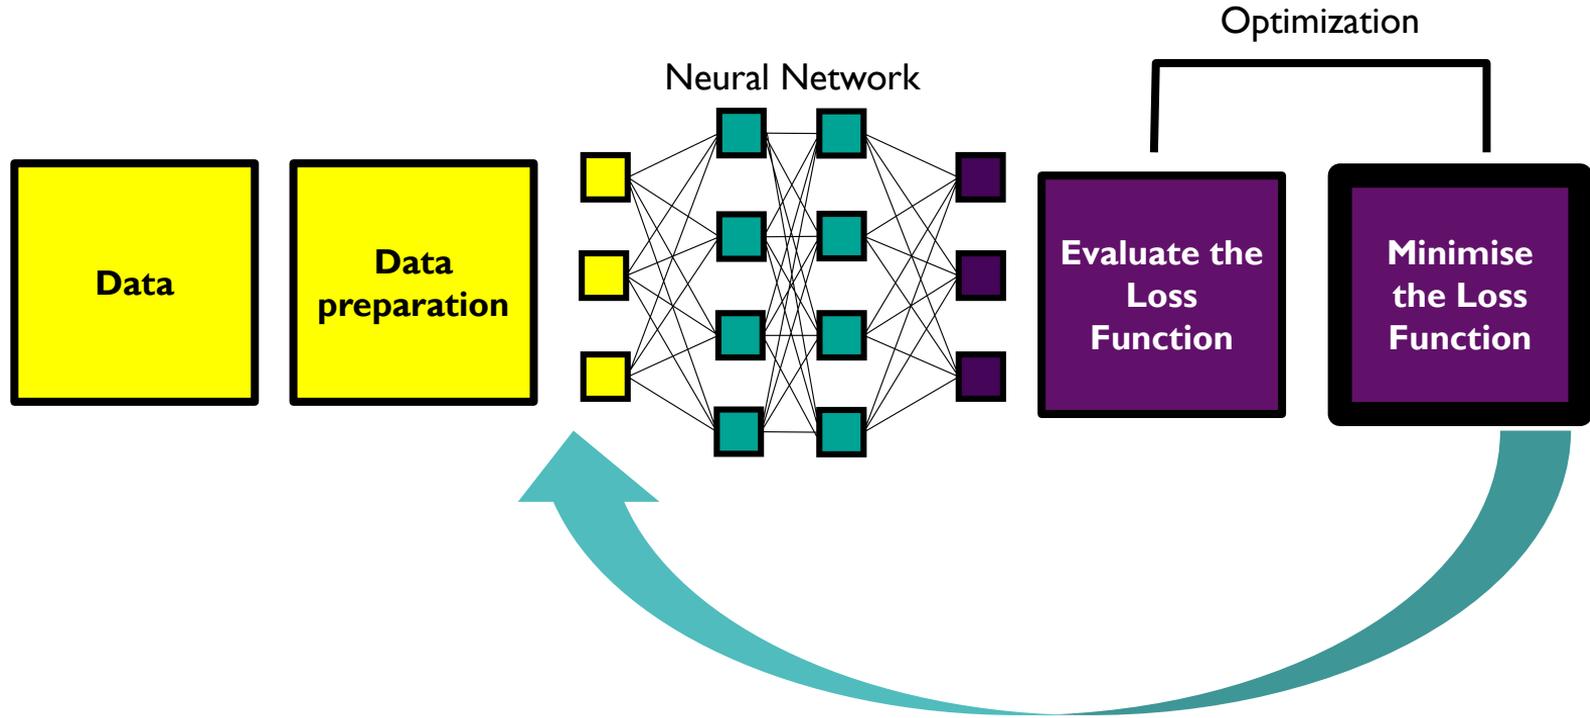
For example:

- 1) Mean Squared Error (MSE)
- 2) Mean Squared Logarithmic Error (MSLE)
- 3) Mean Absolute Error (MAE)

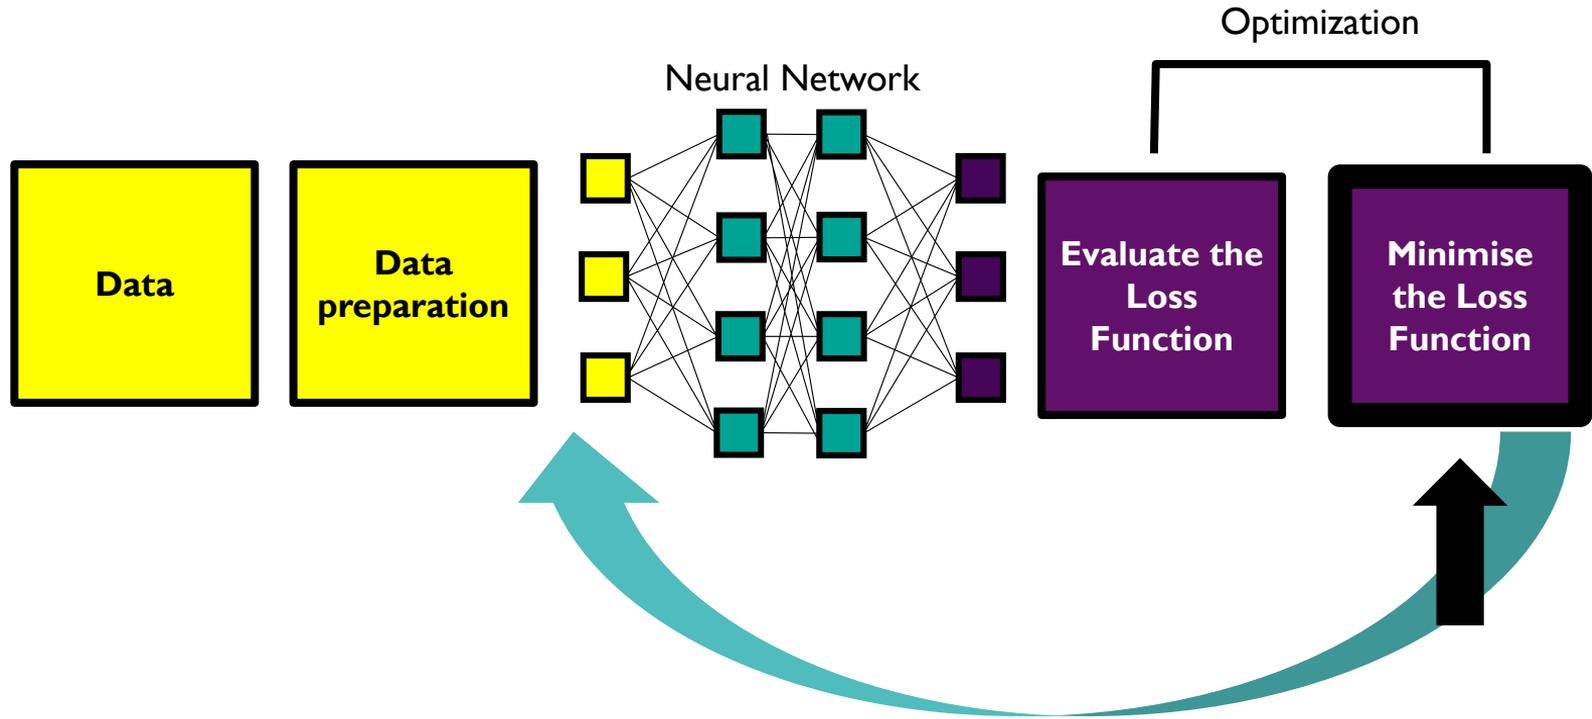
$$MSE = \frac{\sum_{i=1}^n (y_i - \hat{y}_i)^2}{n}$$

$$MSLE = \frac{1}{n} \sum_{i=1}^n (\log(Y_i) - \log(\hat{Y}_i))^2$$

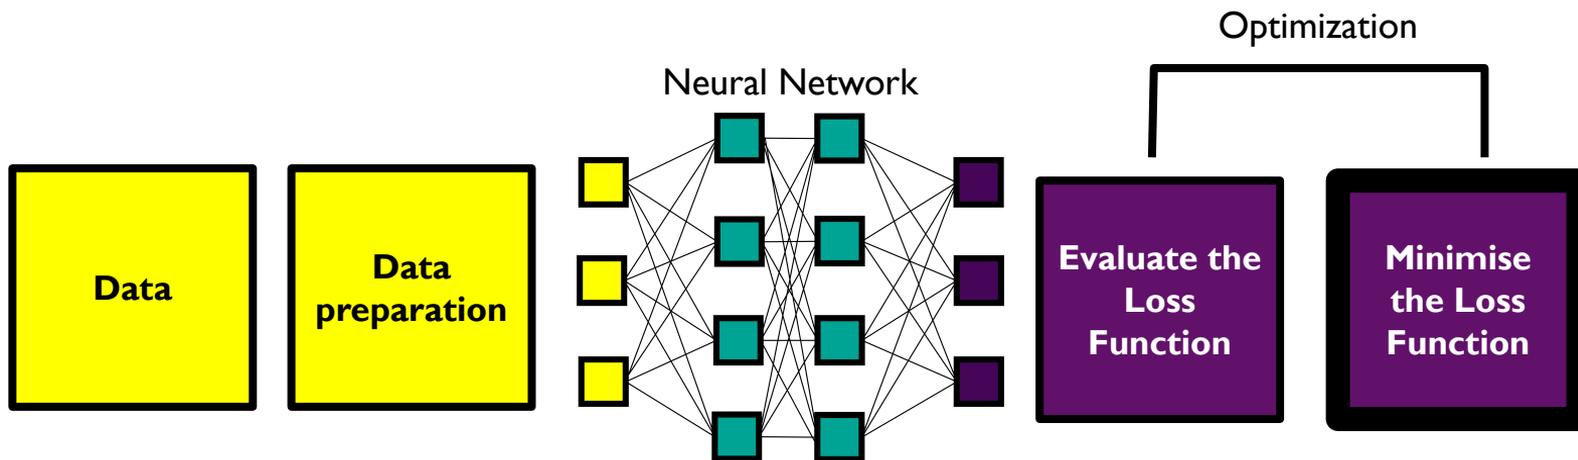
$$MAE = \frac{1}{n} \sum_{j=1}^n |y_j - \hat{y}_j|$$



“...Optimizers are responsible for changing the weights, bias and learning rate of the neurons in the neural network to reach the minimum loss function.”



“...Optimizers are responsible for changing the weights, bias and learning rate of the neurons in the neural network to reach the minimum loss function.”



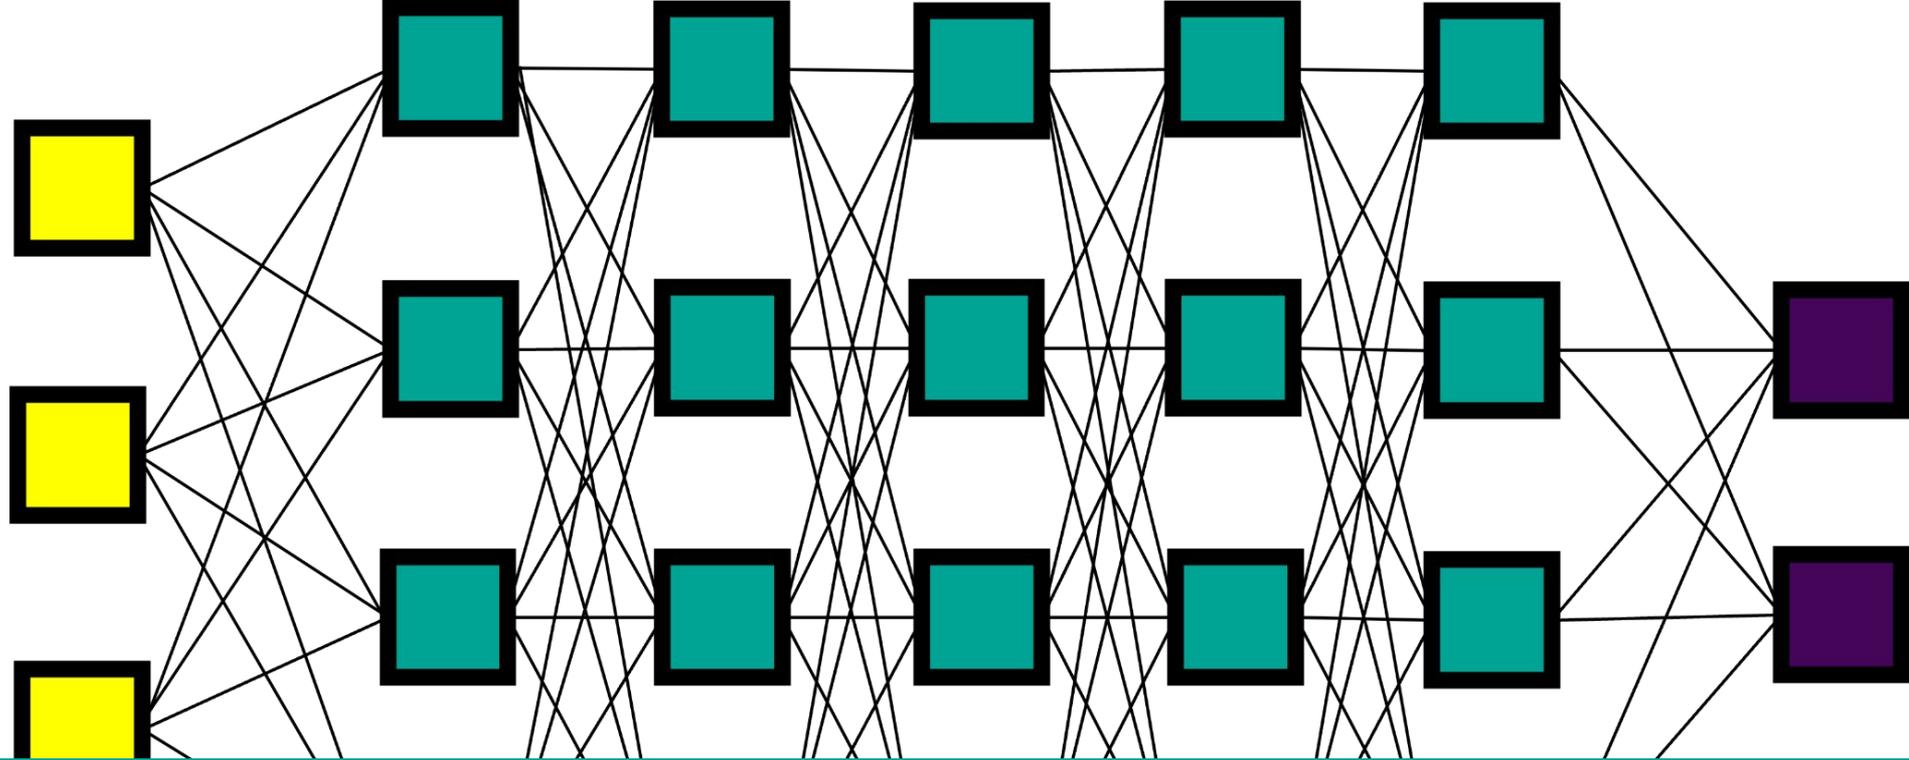
New weights and bias

Learning Rate

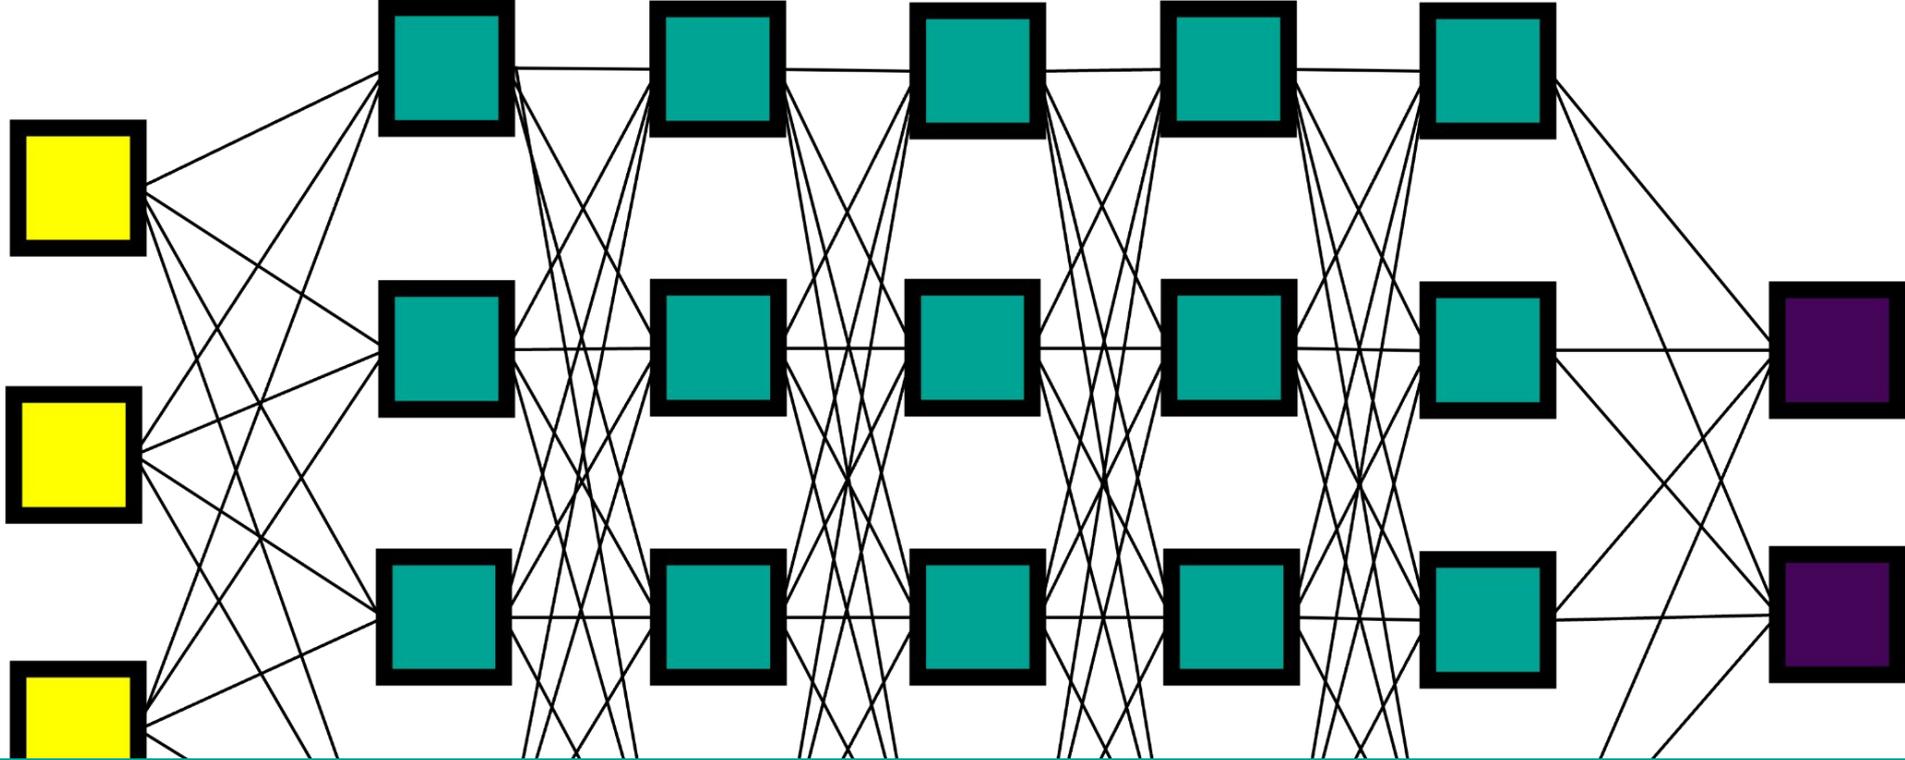
Loss Function

$$\theta_t = \theta_{t-1} - \alpha_t \nabla_{\theta_{t-1}} L_t$$

Current weights and bias

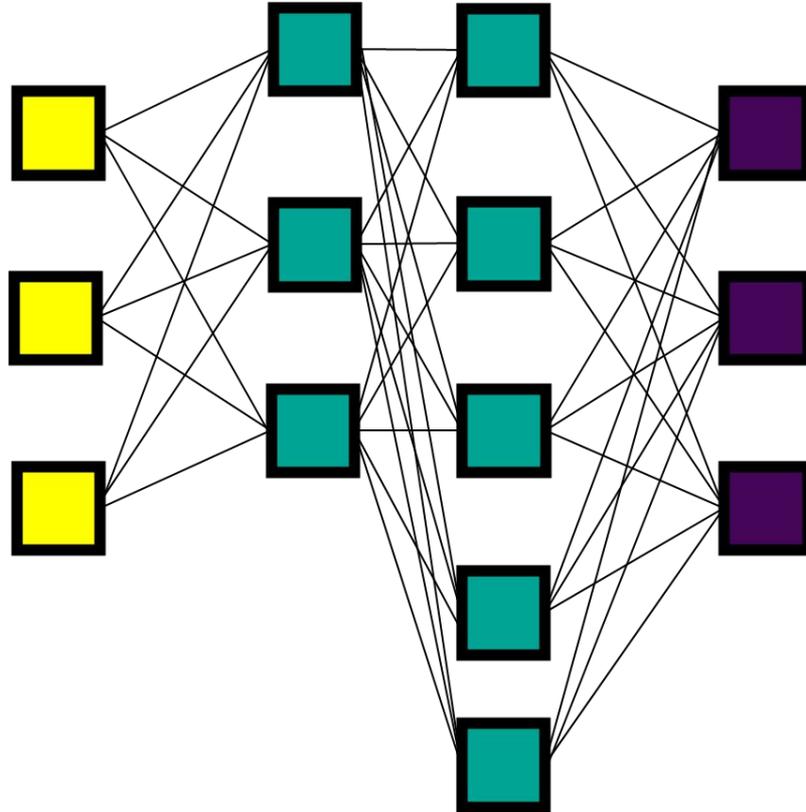


In order for the neural network to learn the neural network architecture may need to be adjusted...

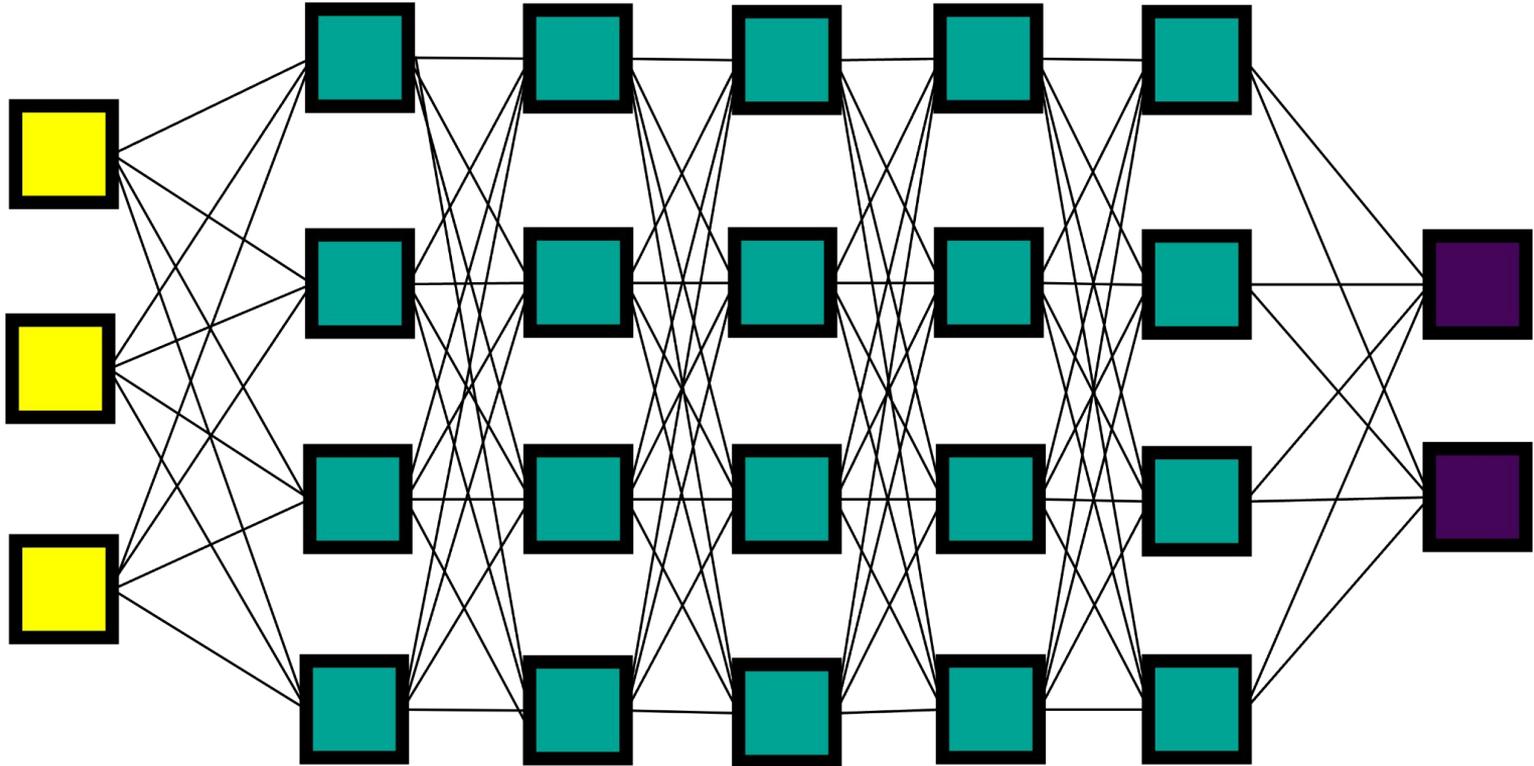


... it is up to a machine learning engineer to adjust the neural network architecture so that the machine 'learns'.

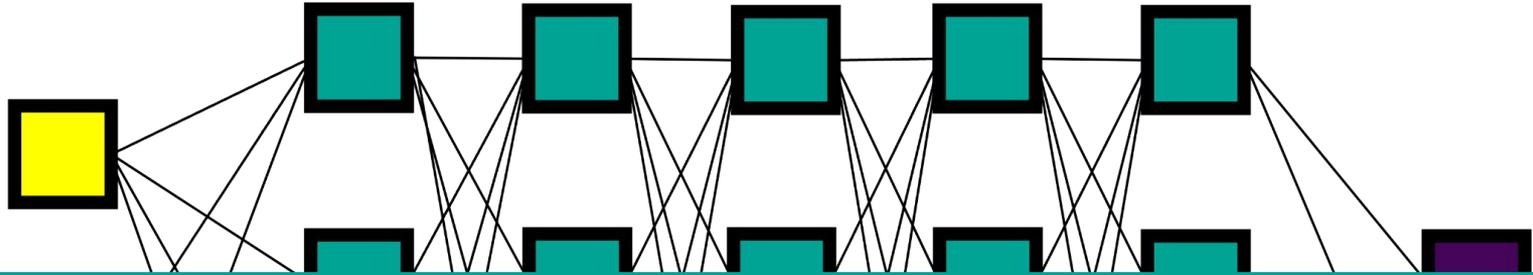
E.g. vary the width of the neural network...



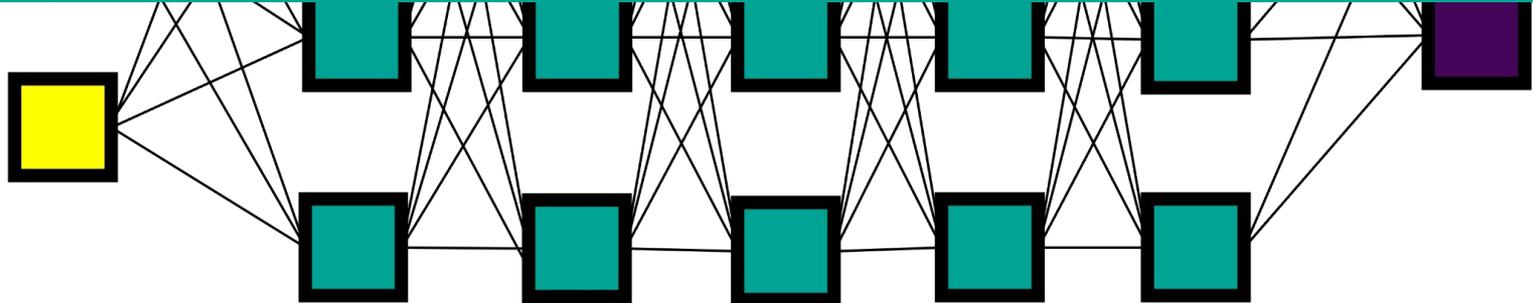
E.g. vary the depth of the neural network...

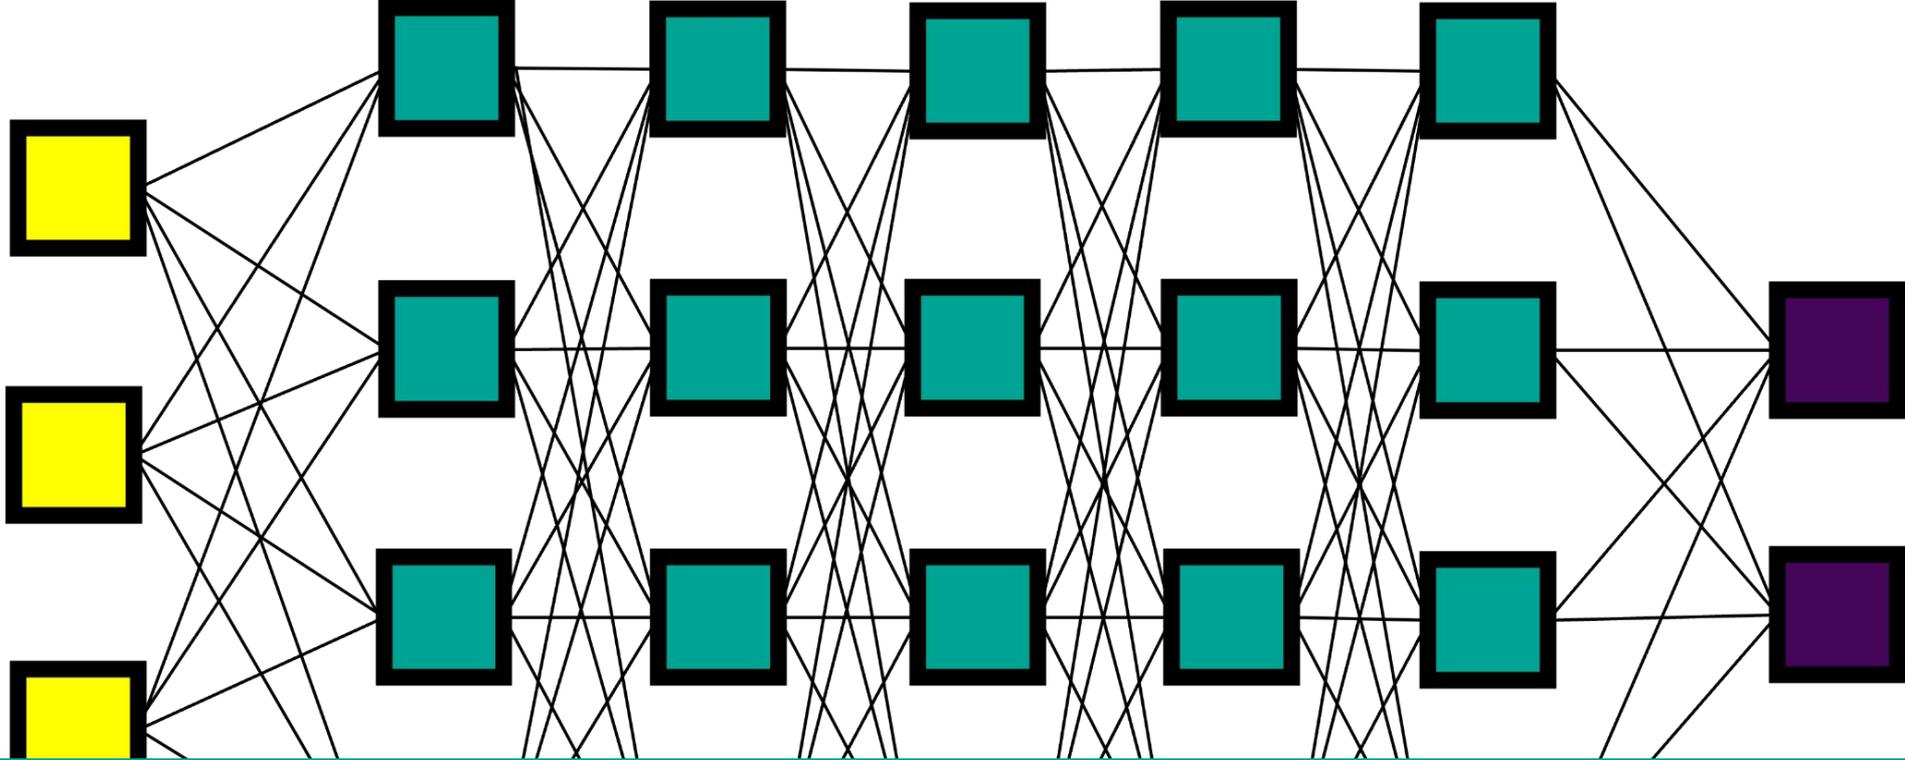


Vary the depth of the neural network...



Deep Learning is defined as a neural network with many layers.





Different neural network architectures enable the machine to 'learn' different data types.

In the modern era we have lots of different types data (text, images, audio)...

... this has led to an explosion of different neural networks!



Deepsets

**Recurrent
Neural
Networks**

Transformer

**Graph Neural
Networks**

**Multi-Layered
Perceptron**

**Convolutional
Neural
Networks**

Keeping up with the Neural Networks



kimkardashian

Follow Message

5,990 posts 363M followers 299 following

Kim Kardashian

@kimkardashian

@skims @skkn @skkypartners

www.rollingstone.com/politics/political-commentary/kim-kardashian-op-ed-joe-biden-armenian-

Graph Neural Networks



Transformer



SKIMS

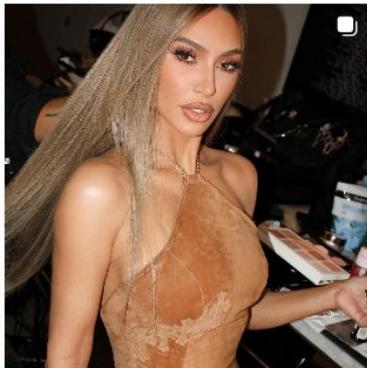
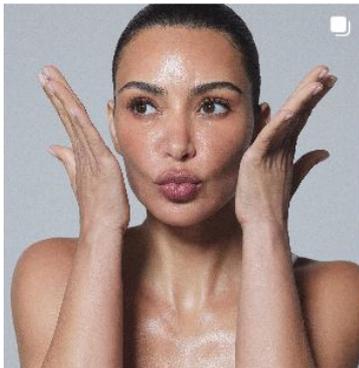


SKKN

Convolutional Neural Networks



POSTS REELS TAGGED





Tell me more, Tell me more...

Convolutional Neural Networks

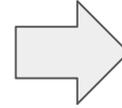
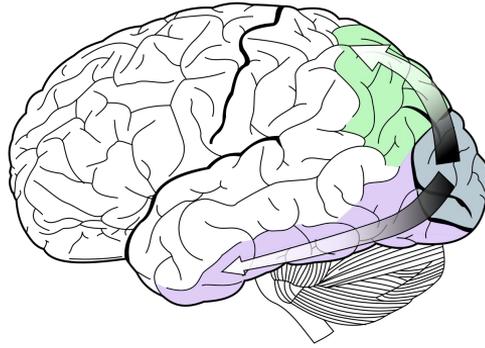
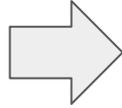


Image Credit: Antonello Trio/Getty images

Convolutional Neural Networks

It is an artificial neural network that features one or more convolutional layers

How do **Humans** process **visual information**?

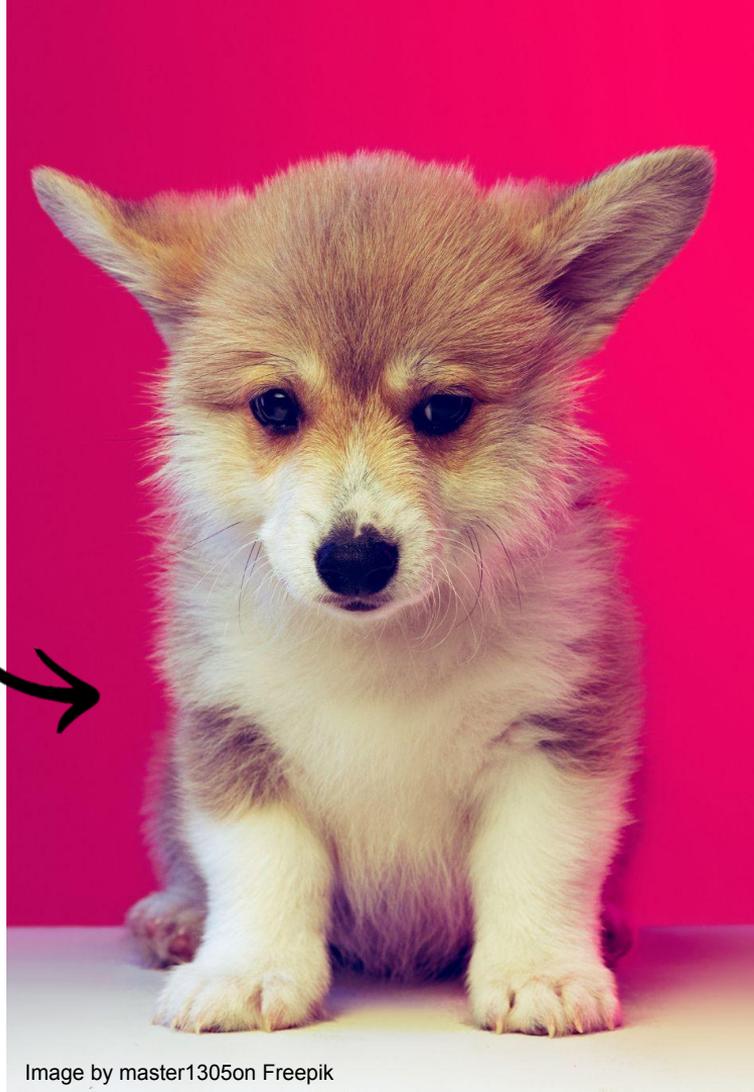


Dog

Did you know that it takes your
brain 100 milliseconds to
process an image.

True or False?

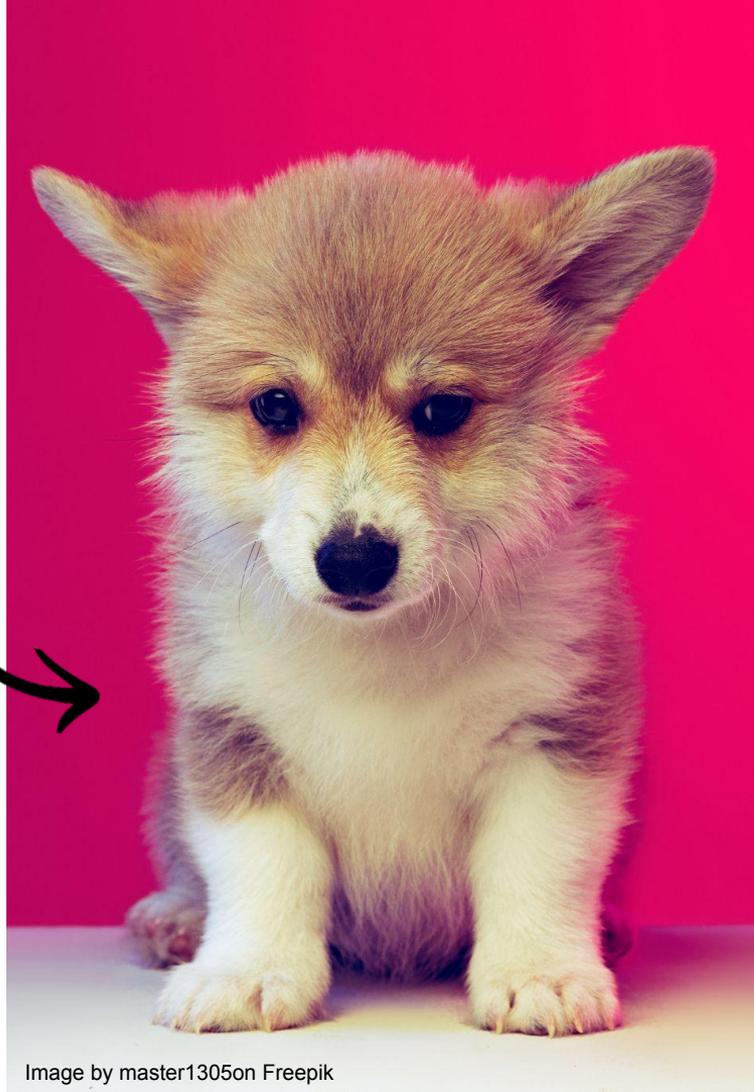
Dog



Did you know that it takes your
brain 100 milliseconds to
process an image.

False

Dog
→



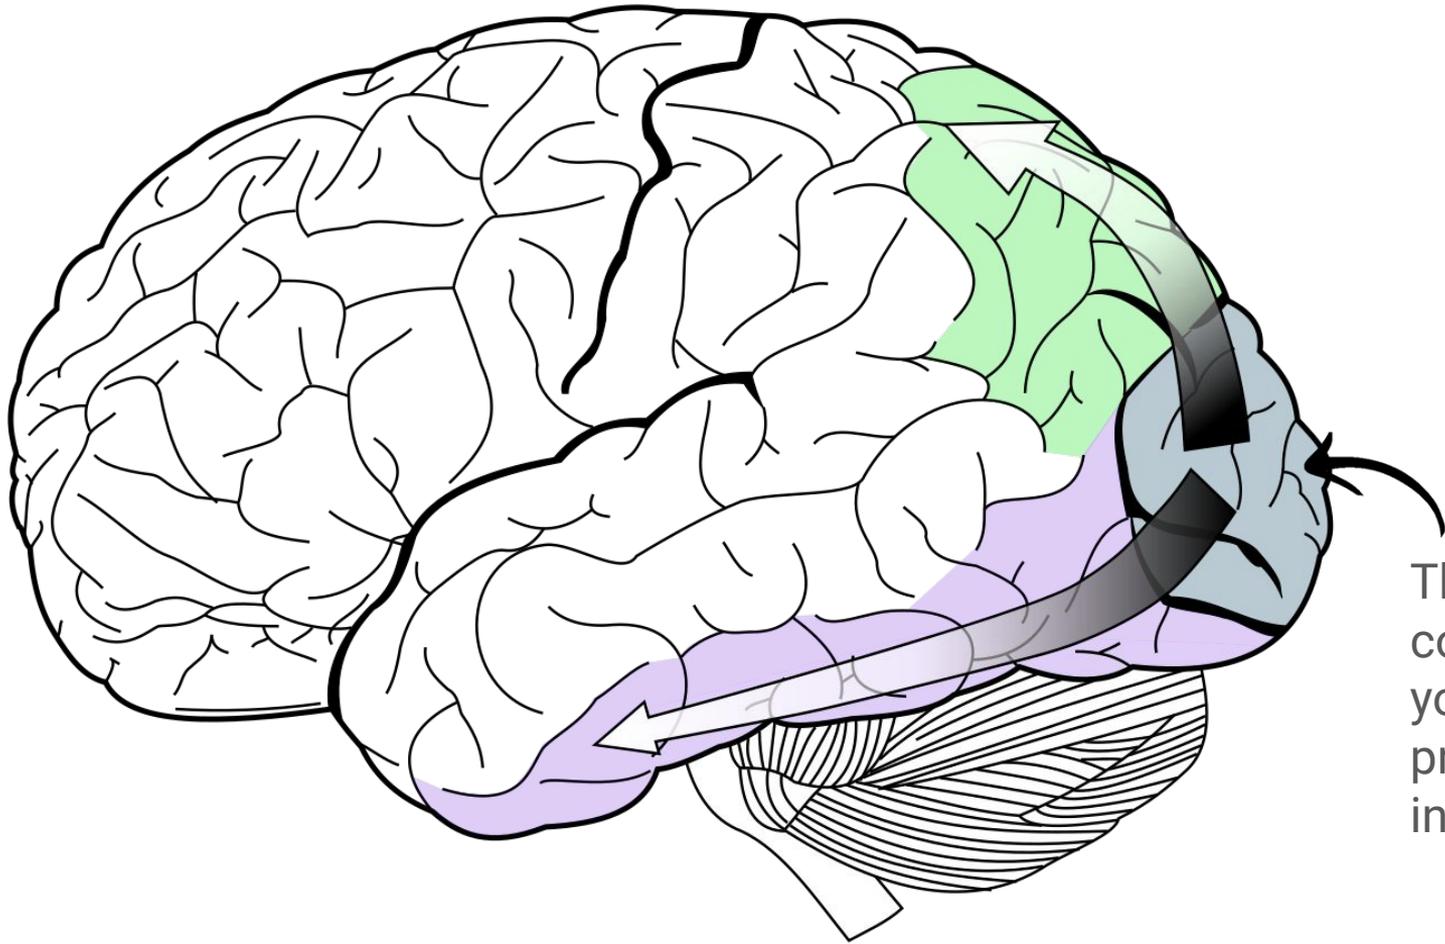
Did you know that it takes your
brain 100 milliseconds to
process an image.

False

*"A team of neuroscientists from MIT
has found that the human brain can
process entire images that the eye
sees for as little as 13 milliseconds"
- Anne Trafton*

Dog
→





The visual cortex is where your brain processes visual information.

**How do
we know
this?**



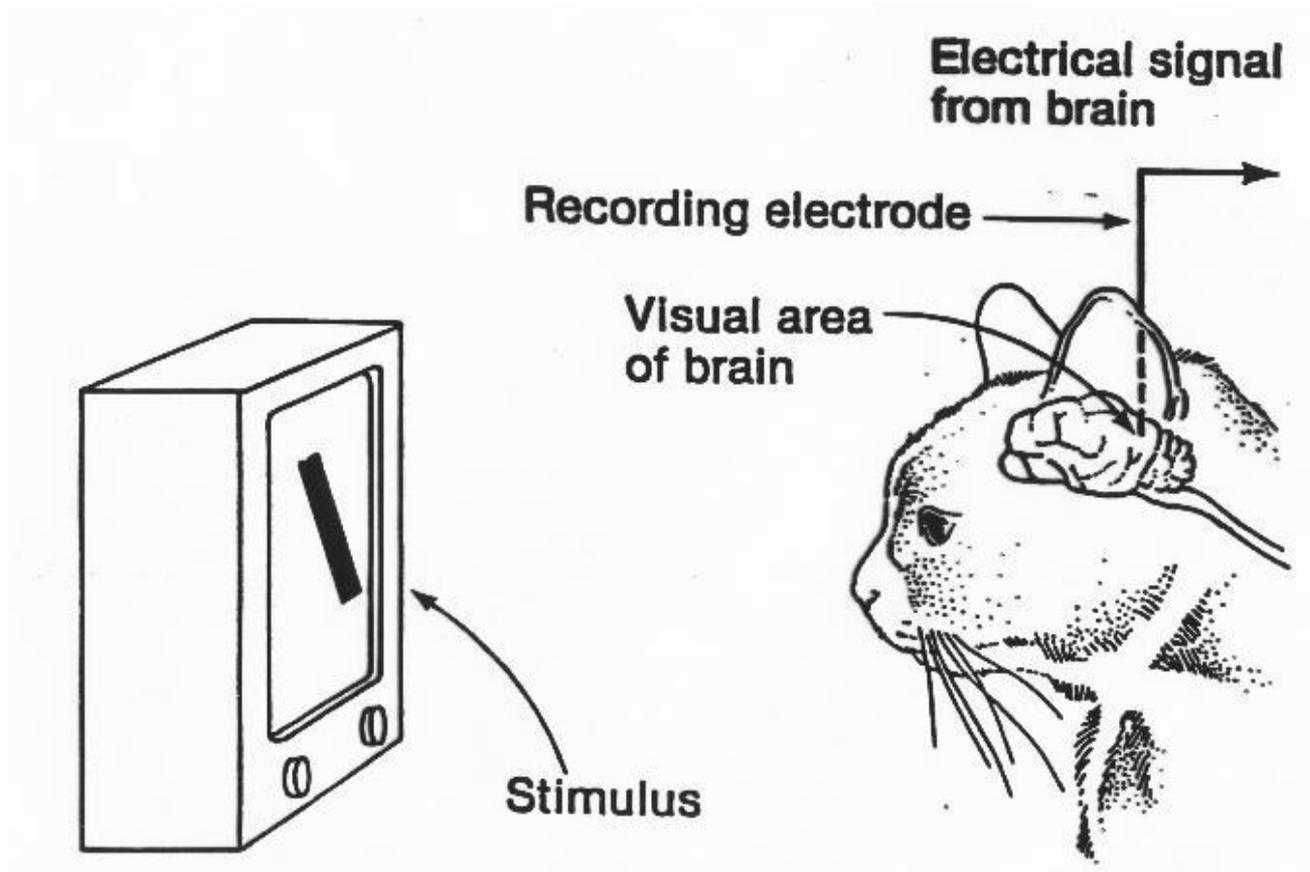


Hubel and Wiesel after winning their Nobel Prize, 1981.

Image Credit: Harvard University Archives

In the 1950s at Johns Hopkins University, the David Hubel and Torsten Wiesel performed pioneering research of how visual information is processed in the brains of mammals.

David and Torsten set up the following experiment...



...and concluded that the cat's primary visual cortex neurons were only stimulated by certain angles of orientation

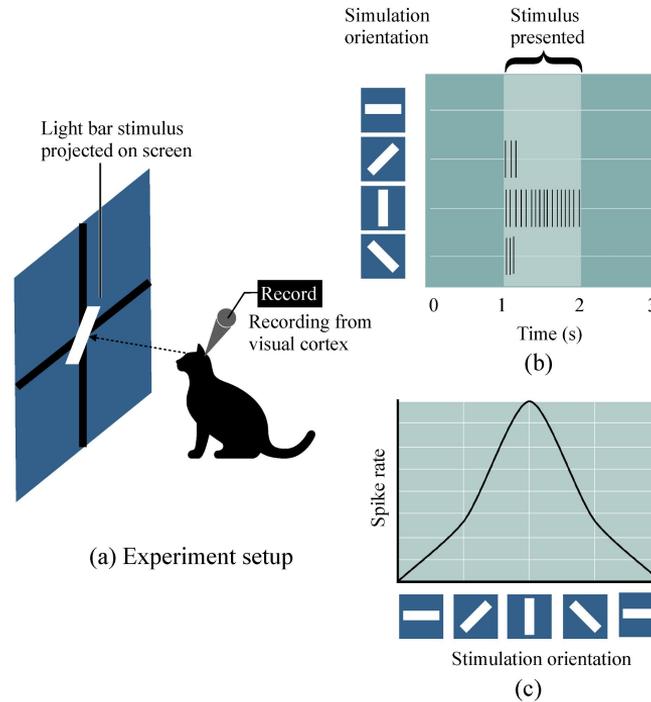
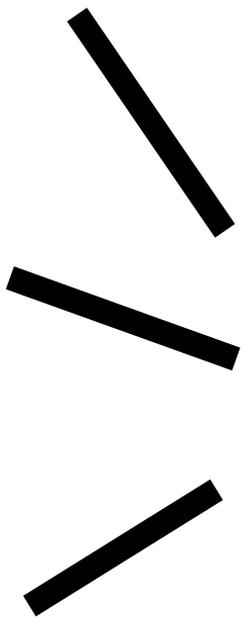
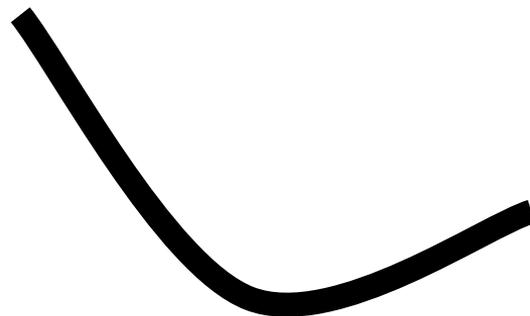
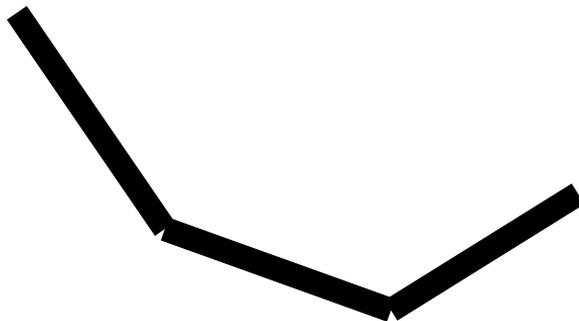


Figure 1. A neuron in the primary visual cortex responds selectively to line segments. Credit: Bin Li et al. (2022)



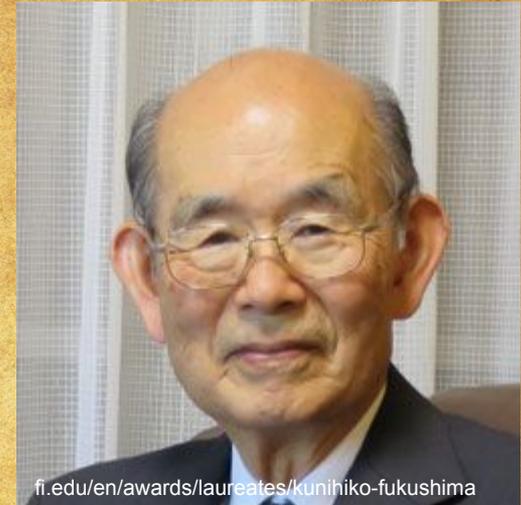
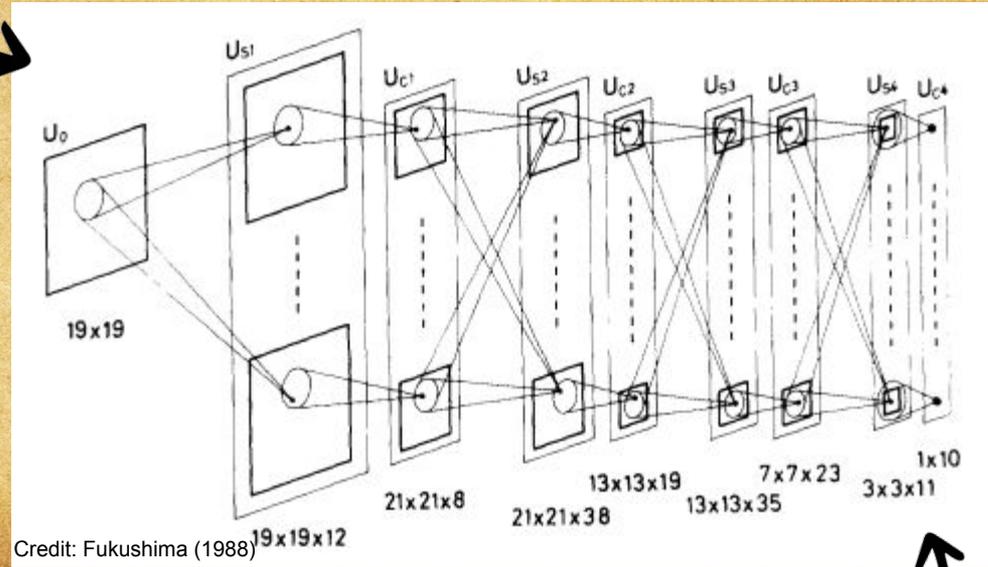
Simple neurons



Complex neurons

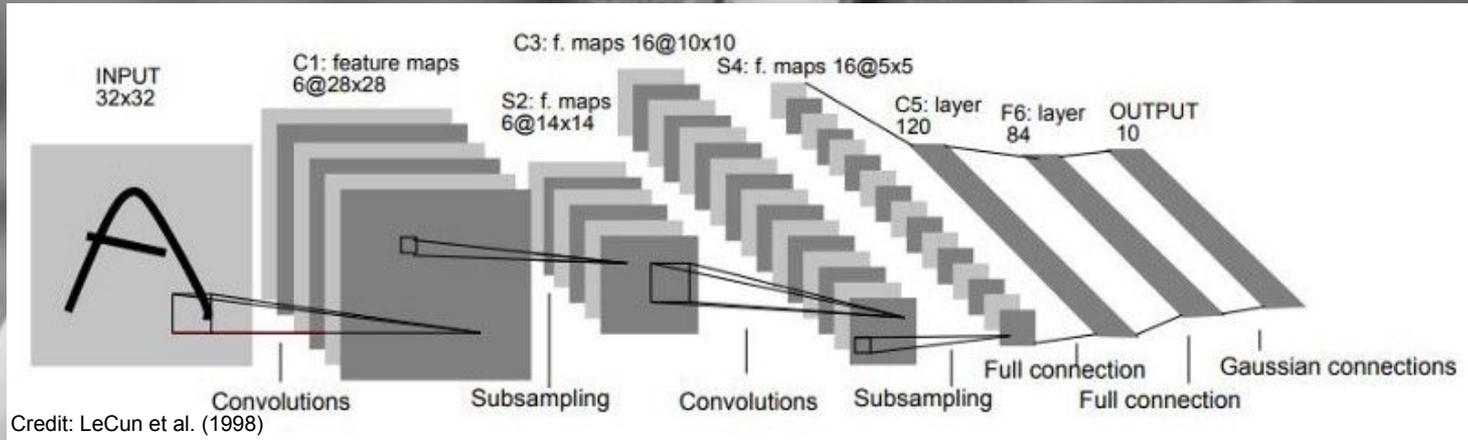


This is **Neocognition**, is the first convolutional neural network architecture.



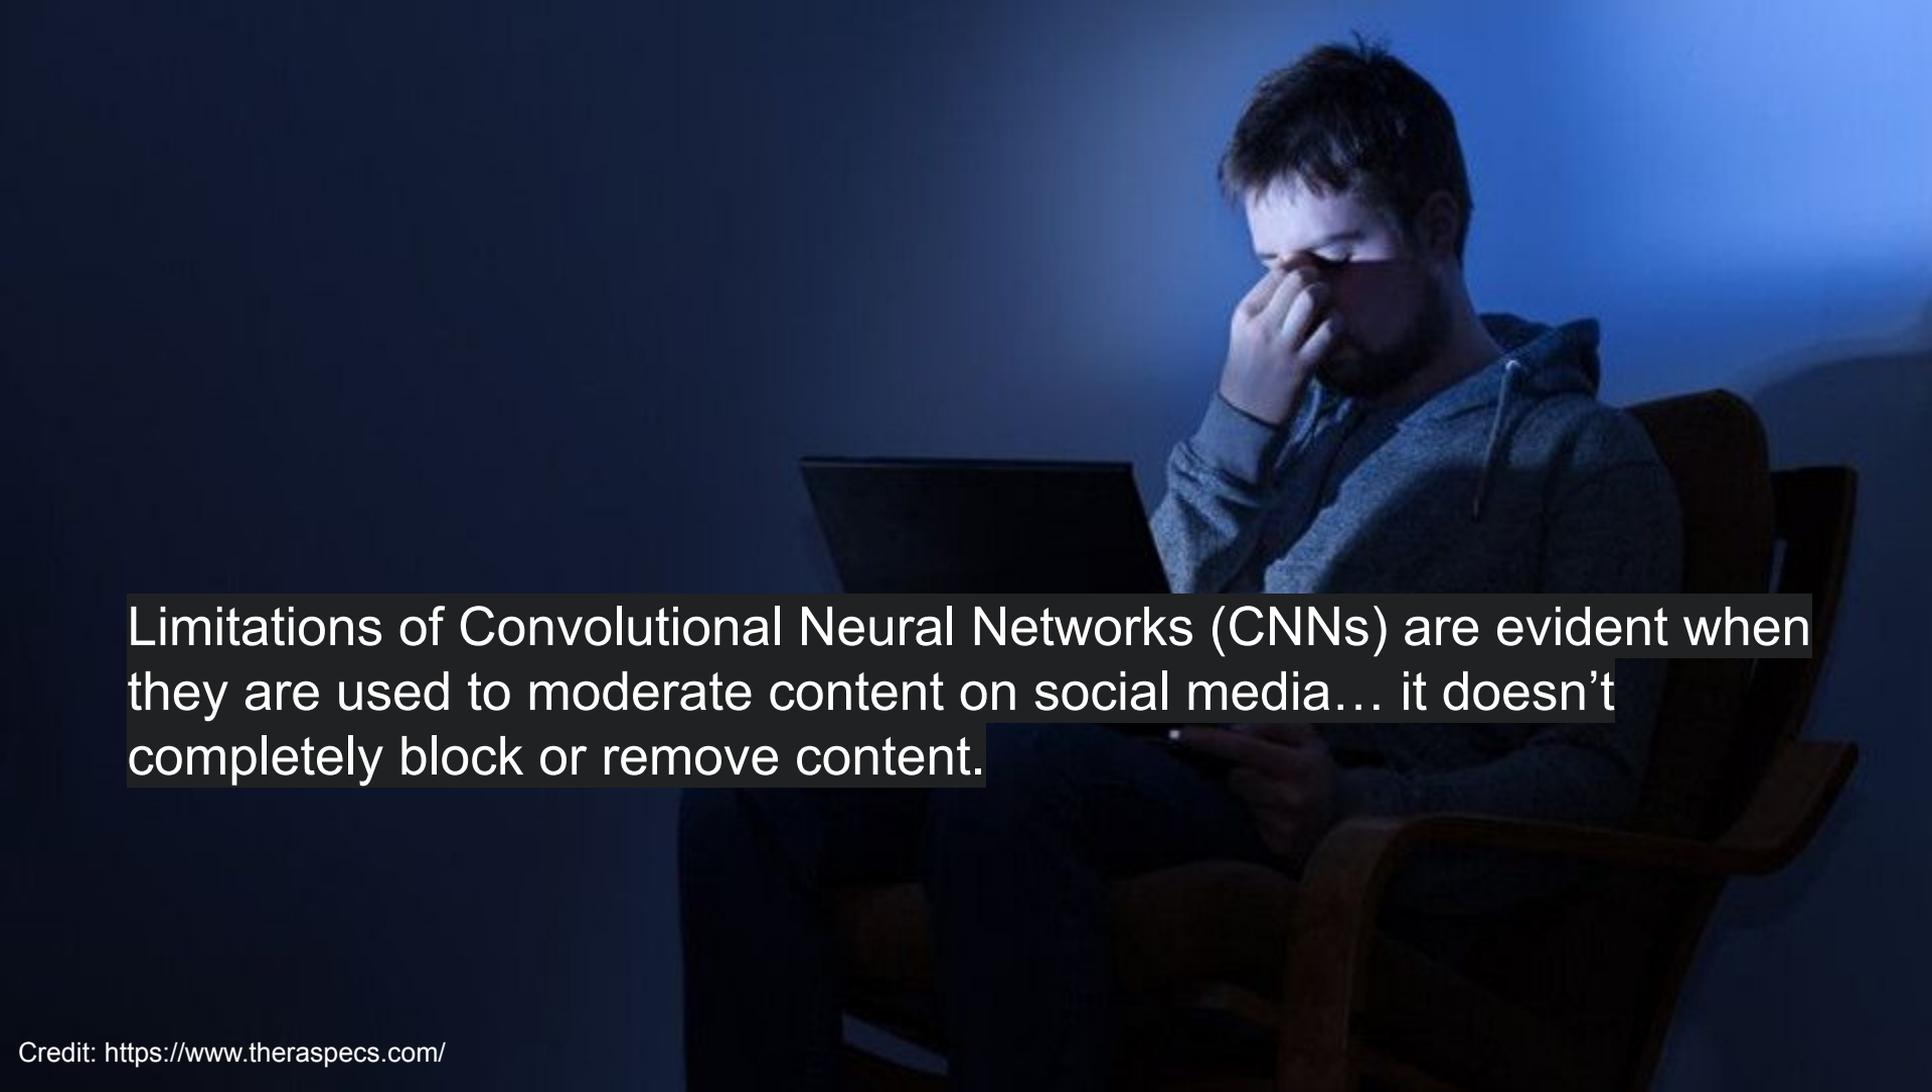
This is **Neocognition** was able to categorize handwritten digits.

Example: LeNet - 5



The neural network architecture shown above is called LeNet 5 and is used in the postal service.

Convolutional neural networks are not **perfect**.

A man with dark hair and a beard is sitting in a wooden chair, looking at a laptop screen. He has his right hand pressed against his forehead, suggesting stress or frustration. The scene is dimly lit with a blueish tint, and the background is dark. A semi-transparent text box is overlaid on the image.

Limitations of Convolutional Neural Networks (CNNs) are evident when they are used to moderate content on social media... it doesn't completely block or remove content.

ImageNet

ImageNet is publicly released dataset containing 14 million labelled images and is used in classification and object detection tasks.



Credit: <https://cs.stanford.edu/people/karpathy/cvnotes/>

ImageNet

CNN's struggle with different lighting, extreme angles and parts of the object.



Credit: <https://cs.stanford.edu/people/karpathy/cvpr12/imagenet101-ilsvb2011-subset/>

**Machine learning is not about
'thinking like a human'**

Machine learning by default does not mimic human thinking...



Machine learning by default does not mimic human thinking...

...and we don't want to because humans have faulty thinking.

Insight - Amazon scraps secret AI recruiting tool that showed bias against women

By Jeffrey Dastin

October 11, 2018 1:50 AM GMT+

Misguided Artificial Intelligence: How Racial Bias is Built Into Clinical Models

Atin Jindal, MD

artificial intelligence racial bias racism bias machine learning health equity

CCBY-NC-4.0 · <https://doi.org/10.56305/001c.38021>

Apple Card Investigated After Gender Discrimination Complaints

A prominent software developer said on Twitter that the credit card was “sexist” against women applying for credit.

Google’s algorithm shows prestigious job ads to men, but not to women. Here’s why that should worry you.



By [Julia Carpenter](#)

July 6, 2015 at 4:43 p.m. EDT

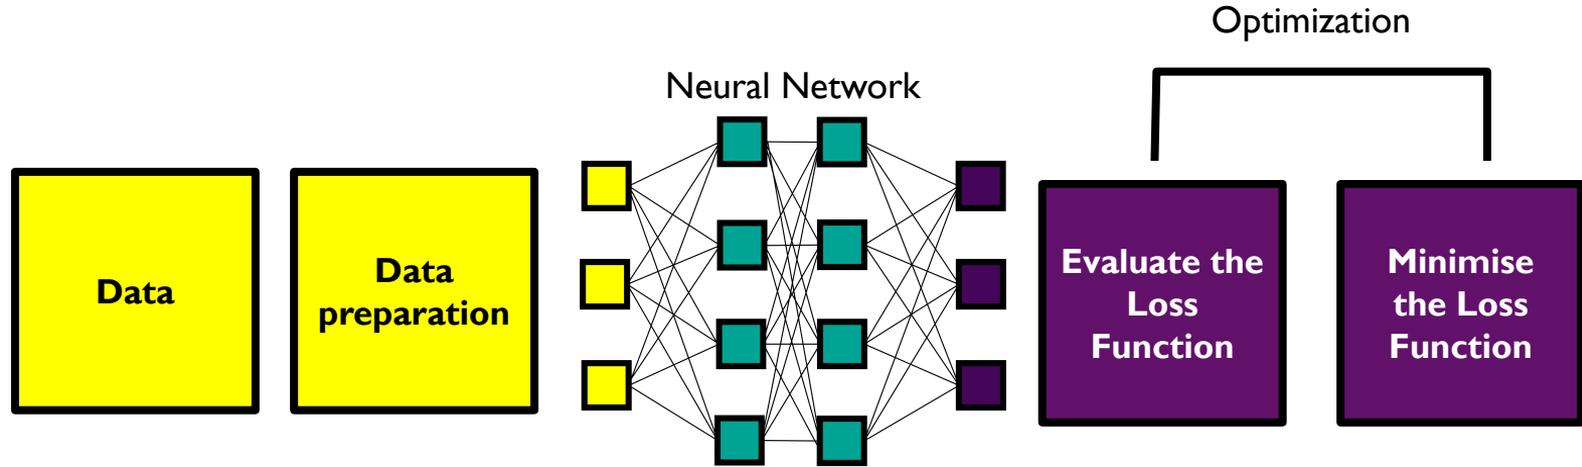
Predictive policing is still racist—whatever data it uses

Training algorithms on crime reports from victims rather than arrest data is said to make predictive tools less biased. It doesn't look like it does.

By [Will Douglas Heaven](#)

February 5, 2021

Bias can be injected into a machine learning pipeline at various stages...



...including how the problem is framed.



You're machine learning engineers, let's predict whether England will win the Euros 2024.

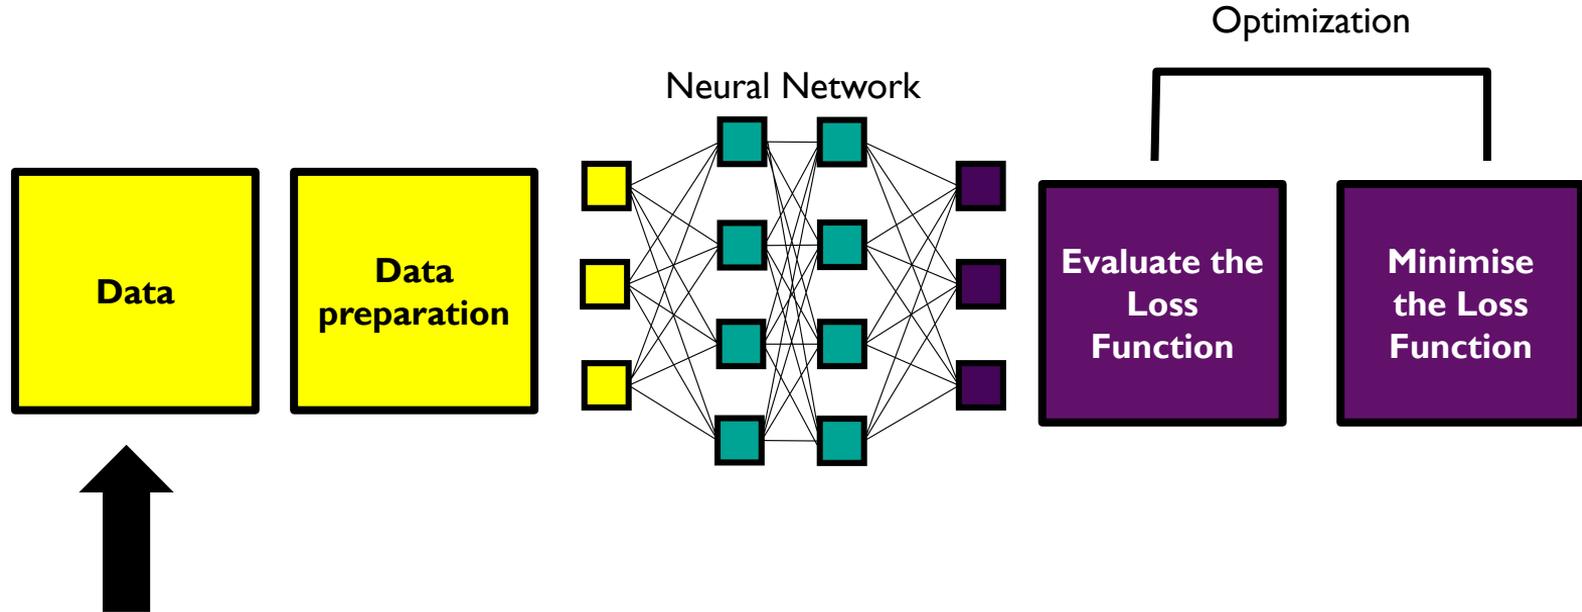


How are we going to do this?



Aim = Total number of Goals

Bias can be injected into a machine learning pipeline at various stages...



...including how the problem is framed.

Nom. goals scored

Hobbies

Height

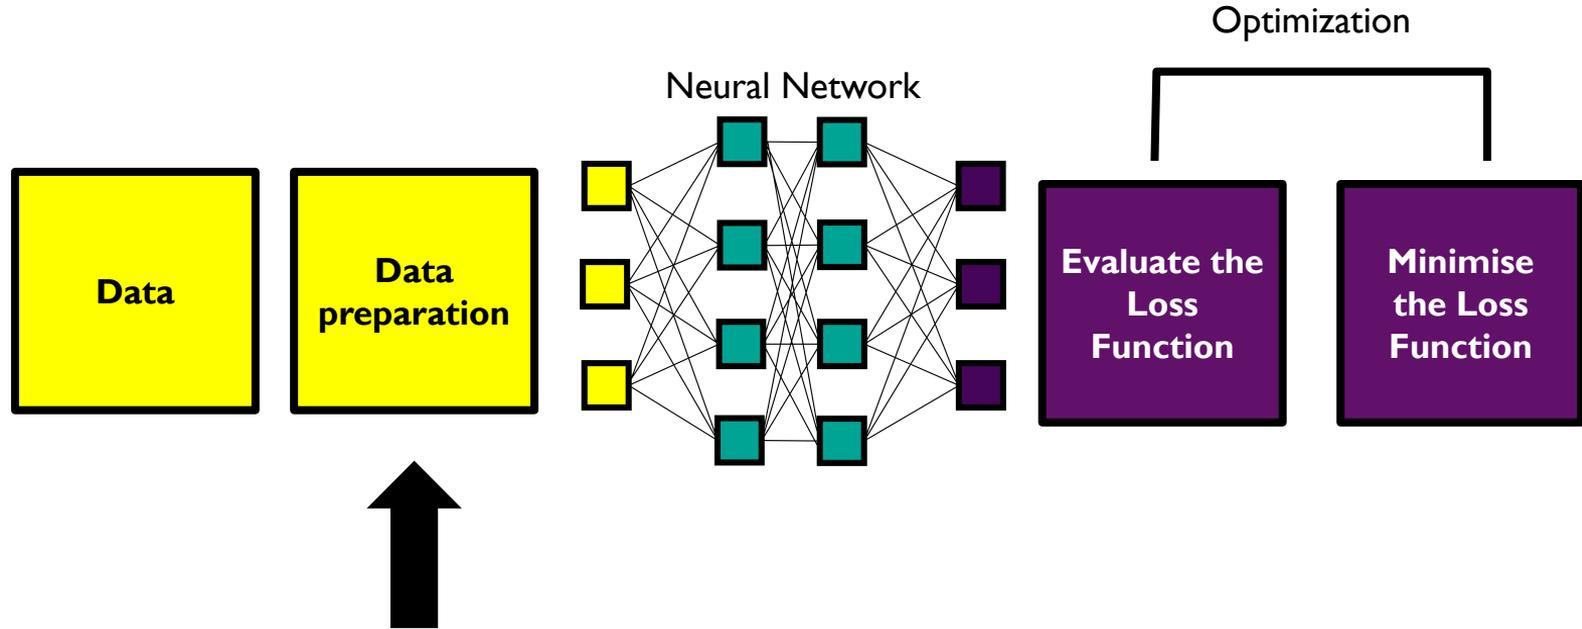
Nom. red cards

Hair Colour

Goal attempts



Bias can be injected into a machine learning pipeline at various stages...



...including how the problem is framed.

Nom. goals scored

Height

Nom. red cards

Goal attempts



Junk in = Junk out

Why is machine learning bias hard to fix?

- 1) Not obvious

Amazon scrapped 'sexist AI' tool

🕒 10 October 2018



The neural network was reprogrammed to ignore explicitly gendered words, like “woman”, but then the system started picking up implicitly gendered words, such as “executed” and “captured”.



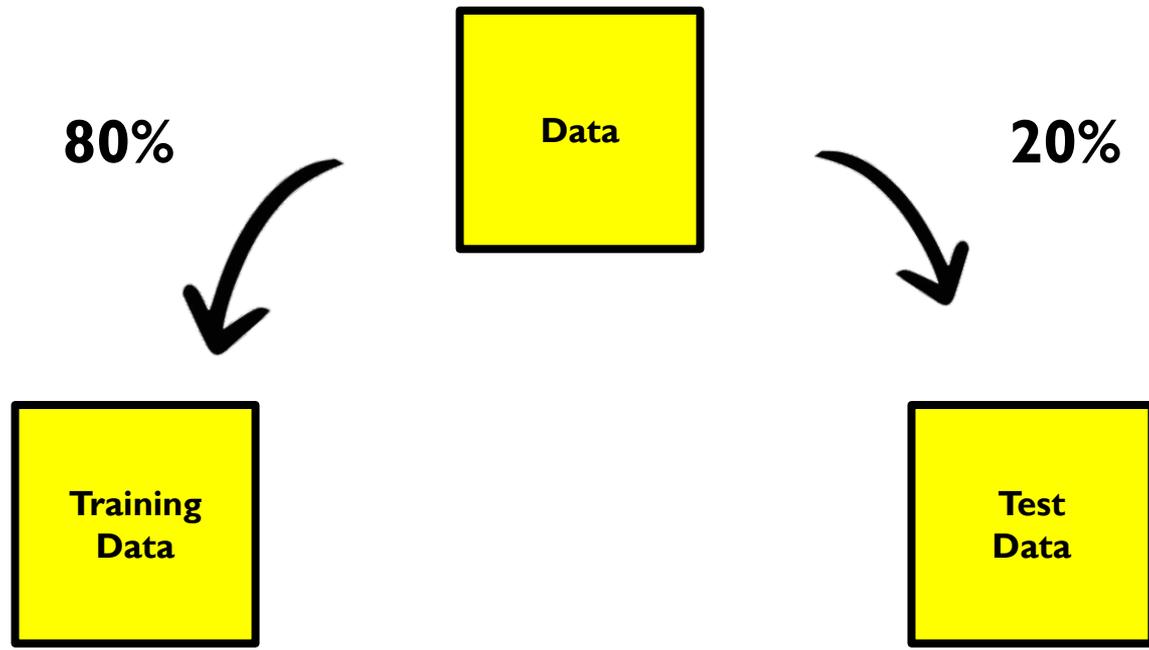
| The algorithm repeated bias towards men, reflected in the technology industry

An algorithm that was being tested as a recruitment tool by online giant Amazon was sexist and had to be scrapped, according to a Reuters report.

The artificial intelligence system was trained on data submitted by applicants over a 10-year period, much of which came from men, it claimed.

Why is it hard to fix?

- 1) Not obvious
- 2) Imperfect process



Why is it hard to fix?

- 1) Not obvious
- 2) Imperfect process
- 3) Lack of social context



United Kingdom



India

Why is it hard to fix?

- 1) Not obvious
- 2) Imperfect process
- 3) Lack of social context
- 4) The definitions of fairness

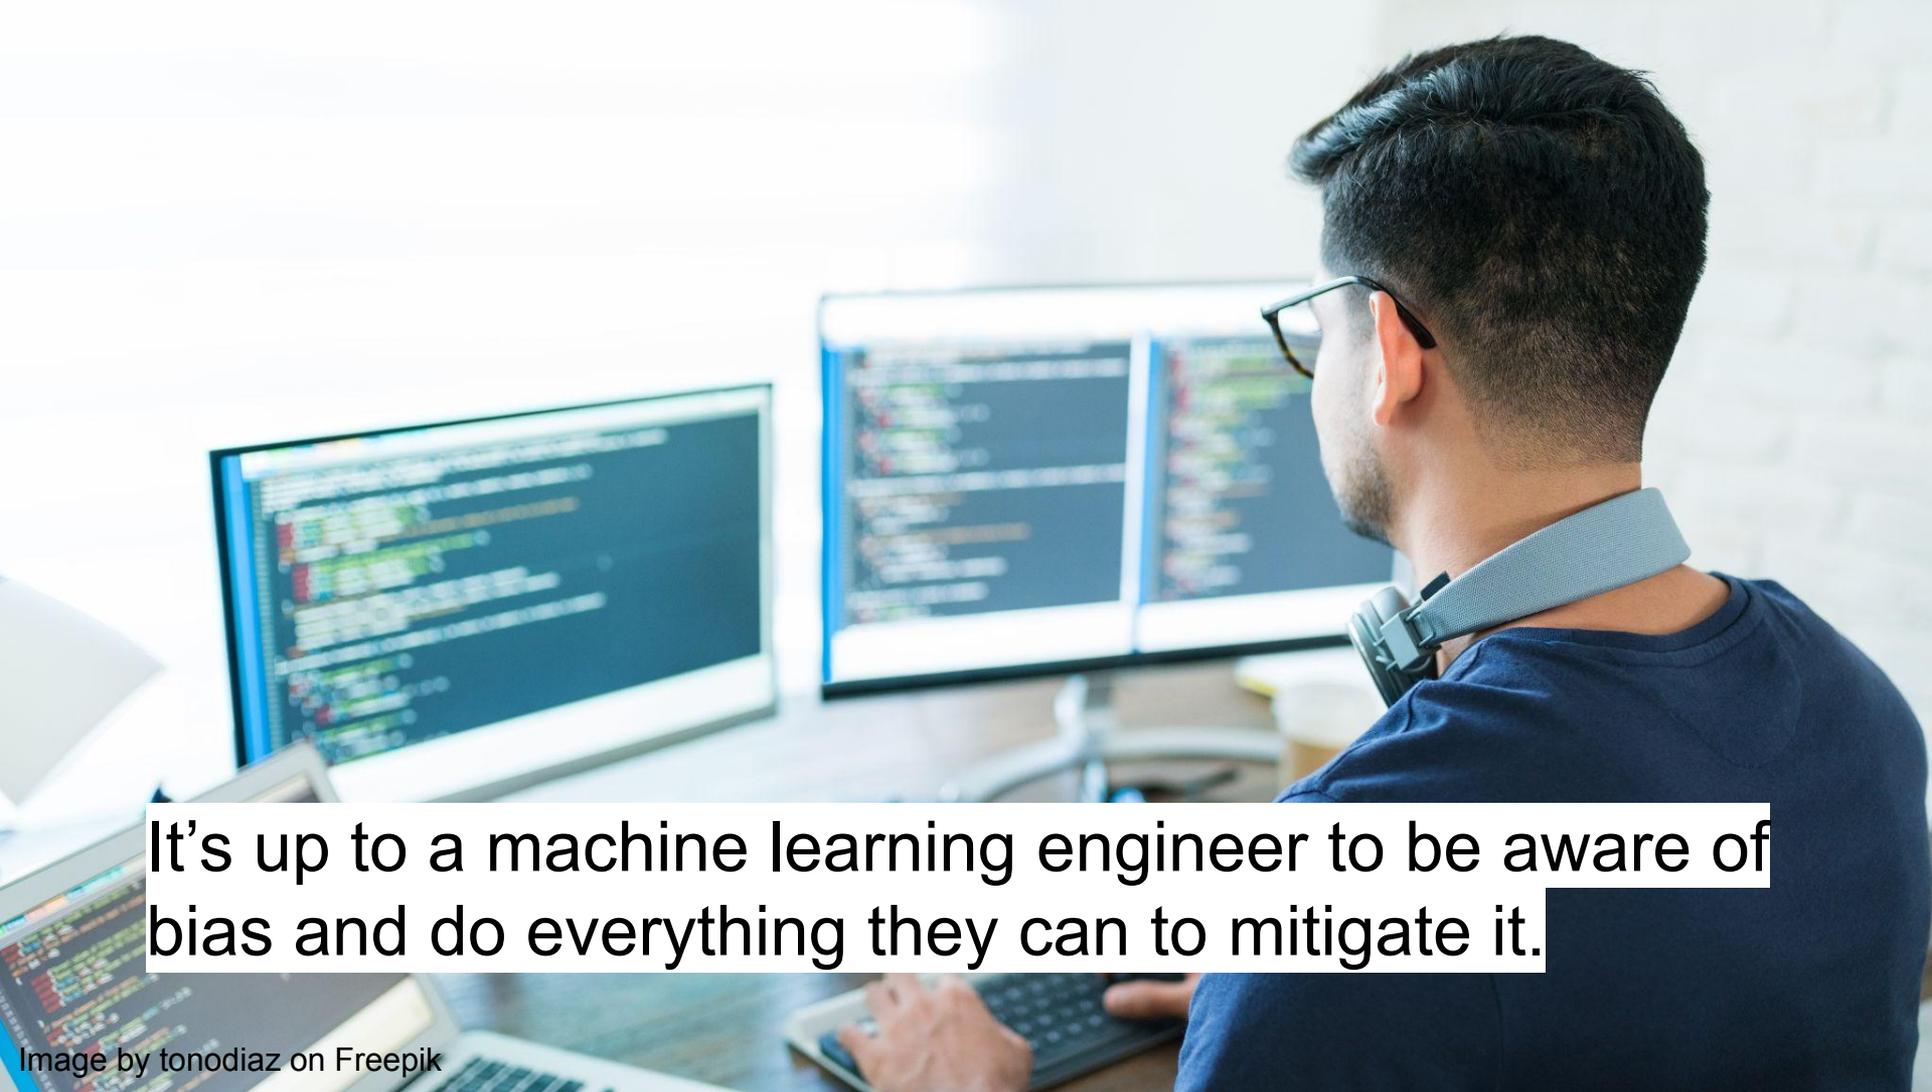
What is 'fairness'?

50% of men and 50% of women are considered high risk?



Everyone considered high risk?





It's up to a machine learning engineer to be aware of bias and do everything they can to mitigate it.



How would you reduce bias in machine learning?

Maybe **you**
could be the
solution?



This is how you can
help...





Maturity of Baby Sounds

ABOUT

Welcome (back)! We are excited to relaunch this project. In addition to studying the sounds children make, we are now asking you to help us understand how children learn language from the world around them. Check out the tutorial to learn more about our new workflow. Thank you for your continued support!

Welcome! We need your help to classify some audio clips. This will help us to better understand how children learn language from the world around them.



Corresponding with Quakers

ABOUT CLASSIFY TALK COLLECT

CORRESPONDING WITH QUAKERS

Investigating race, gender, class, and religion in 18th- and 19th-Century Irish Quaker Documents

LEARN MORE

GET STARTED!

You can do real research by clicking to get started here!



Picturing Michigan's Past

Get involved at zooniverse.org!

As a result of your enthusiastic support and contributions to this project, we couldn't have reached this point without you! You can still help us by transcribing text from postcards by clicking on the link below.

Researchers explore the world of woodpecker nest cavities and the animals that use them!



Dolphin Chat

ABOUT CLASSIFY TALK COLLECT CETALINGUA PROJECT

Join Our Free Course: Citizen Science for Marine Mammal Enthusiasts! [Read more here.](#)

ORGANIZATION: CETALINGUA PROJECT

Identification and classification of dolphin vocalizations.

Learn more

Help us describe and categorize more than 60,000 postcards featuring rare photographs depicting life in Michigan during the first half of the twentieth century.





Black Hole Hunters



Galaxy Zoo

Language English

[ABOUT](#) [CLASSIFY](#) [TALK](#) [COLLECT](#)

Few have witnessed what you're
about to see

[Learn more](#)

[Get started](#)



Planet Hunters NGTS

[ABOUT](#) [CONTACT](#)



Dark Energy Explorers

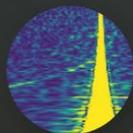
Get involved at zooniverse.org!

Identify distant galaxies to help
measure dark energy when the
universe was just ~2-3 billion
years old

[Learn more](#)

Discover hidden worlds with the
Next-Generation Transit Survey

[Learn more](#)



Gravity Spy

We need you to help us identify the physics...



PROJECTS ABOUT GET INVOLVED TALK BUILD A PROJECT NEWS

SIGN IN REGISTER



Kilonova Seekers

Language

ABOUT CLASSIFY TALK COLLECT

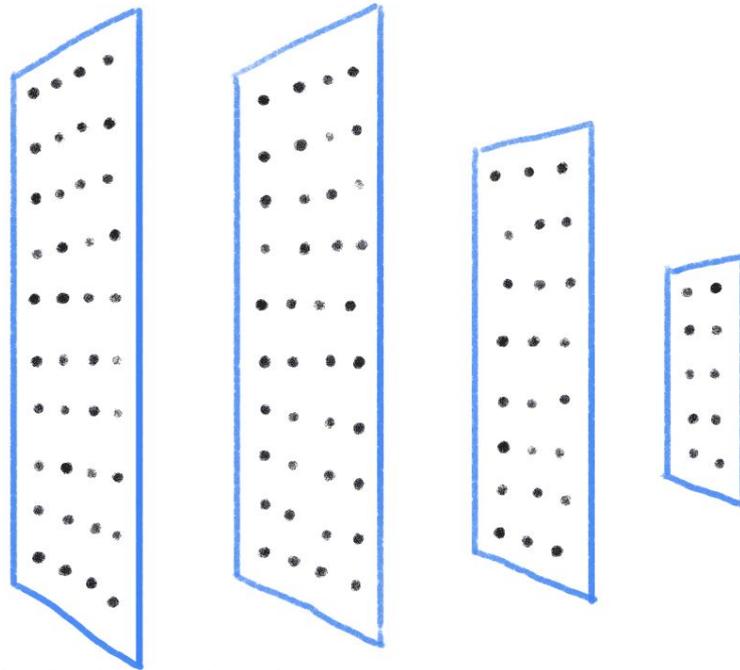
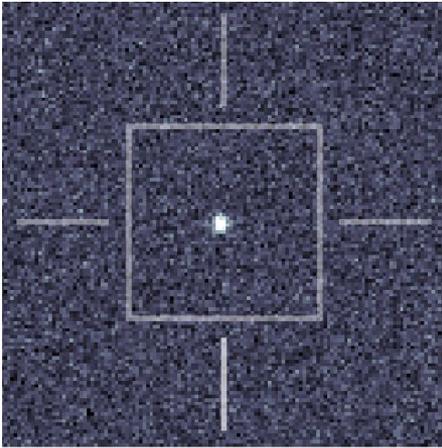
Welcome back to a new year of KN Seekers - we can't wait to see what you discover in 2024!

Find cosmic explosions in real-time with the Gravitational-wave Optical Transient Observer (GOTO)
- new data uploaded hourly!

[Learn more](#)

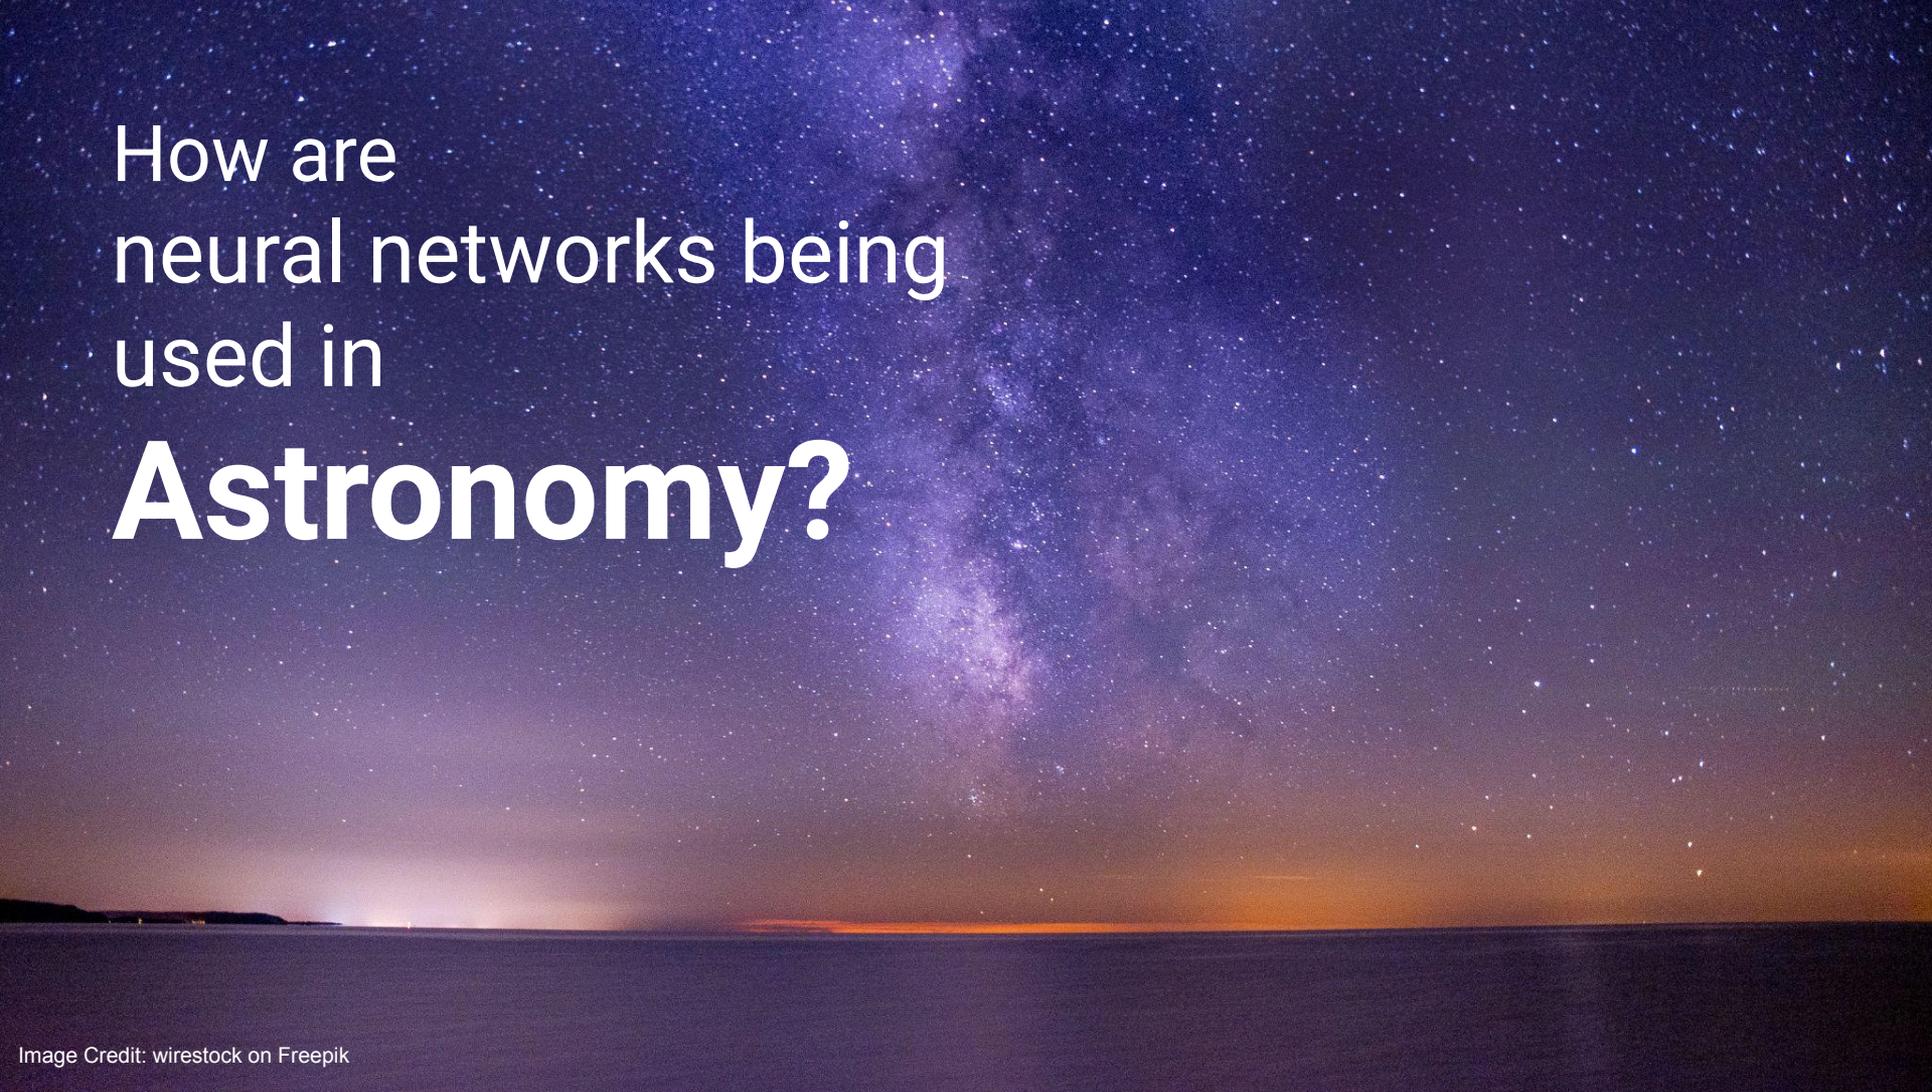
[Get started](#)

...in this research a neural network is uncertain about whether it has observed an explosion and we need you to clarify!



Explosion
Uncertain
No Explosion

Credit (edited): Venkatesh Tata

A night sky photograph showing the Milky Way galaxy stretching across the frame. The stars are bright and numerous, with a soft glow from the galaxy's core. The bottom of the image shows a dark horizon line, possibly over water, with a faint orange glow from the setting or rising sun.

How are
neural networks being
used in
Astronomy?



Square kilometer array. Credit Wiki -



The Victor M. Blanco 4-meter telescope building Credit DES



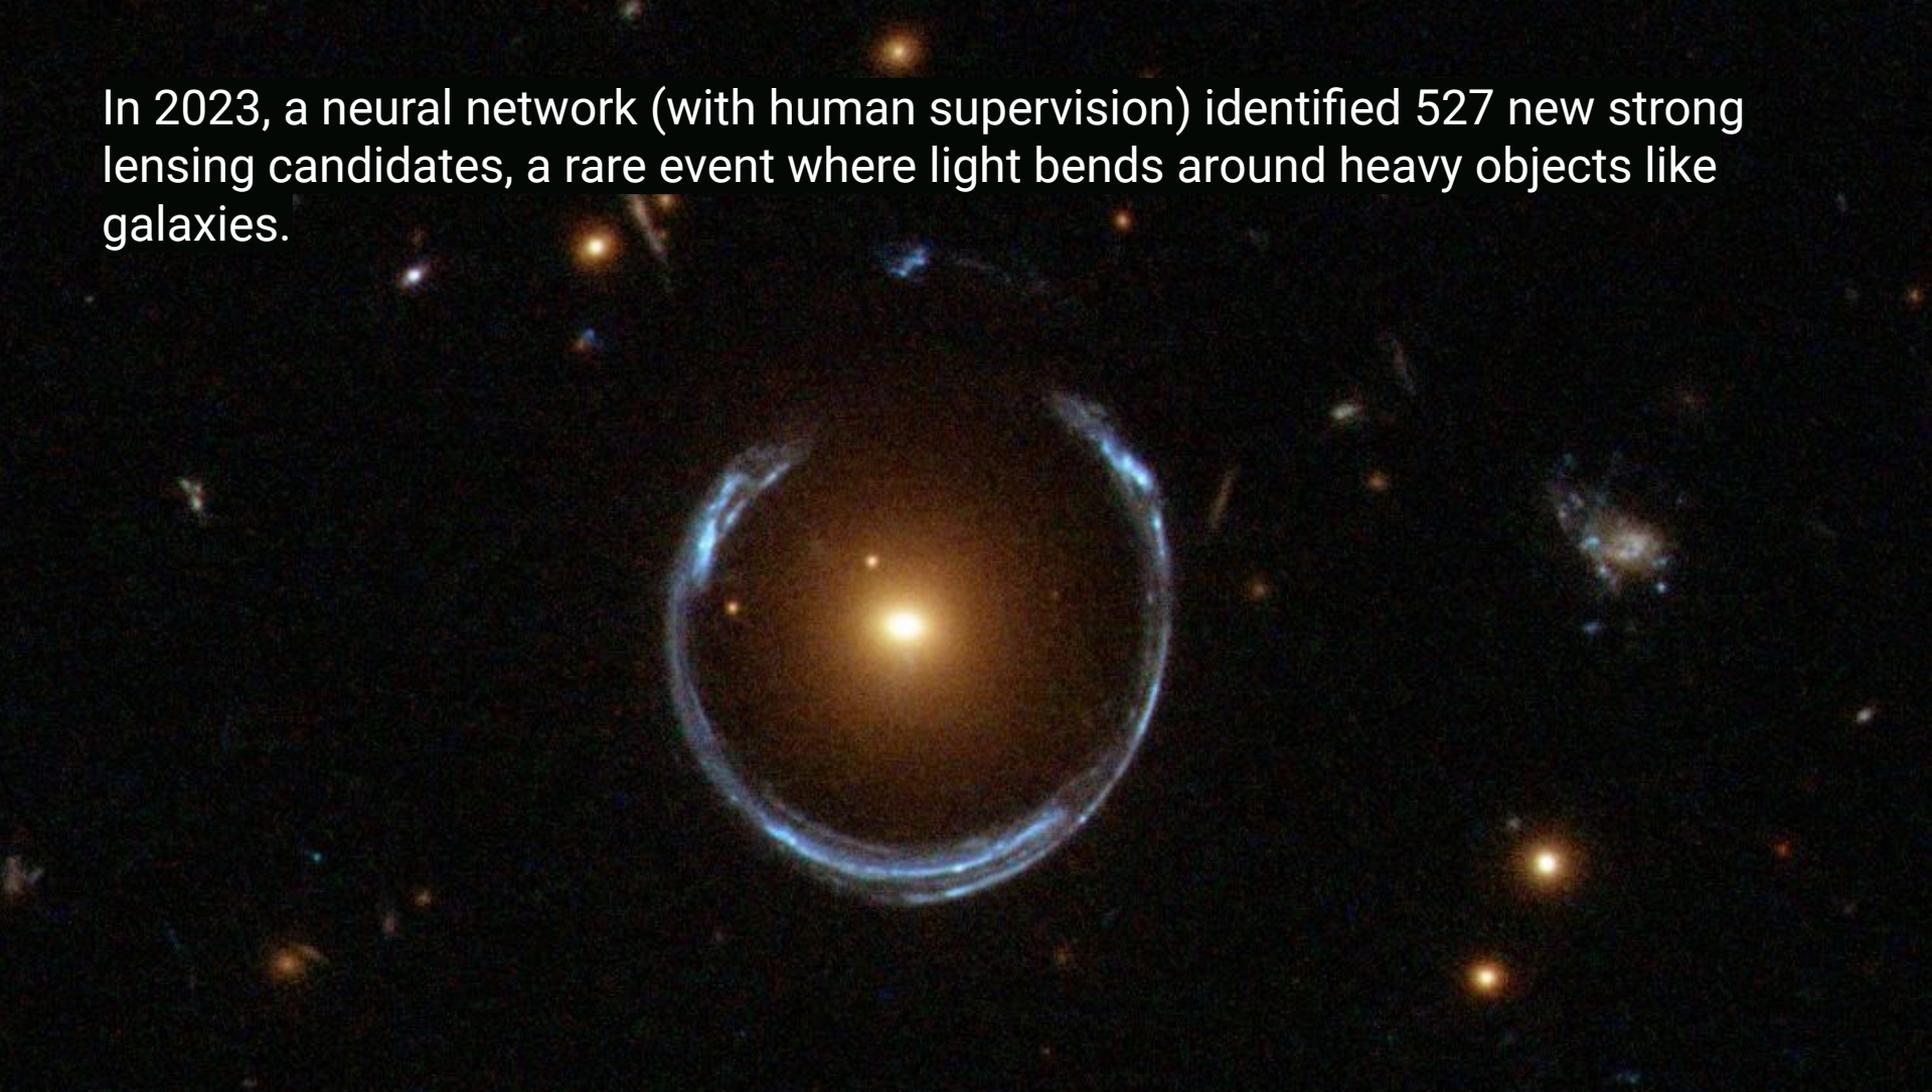
Vera C Rubin Observatory. Credit: Rubin Obs/NSF/AURA



Gravitational wave detector. Credit: Ligo



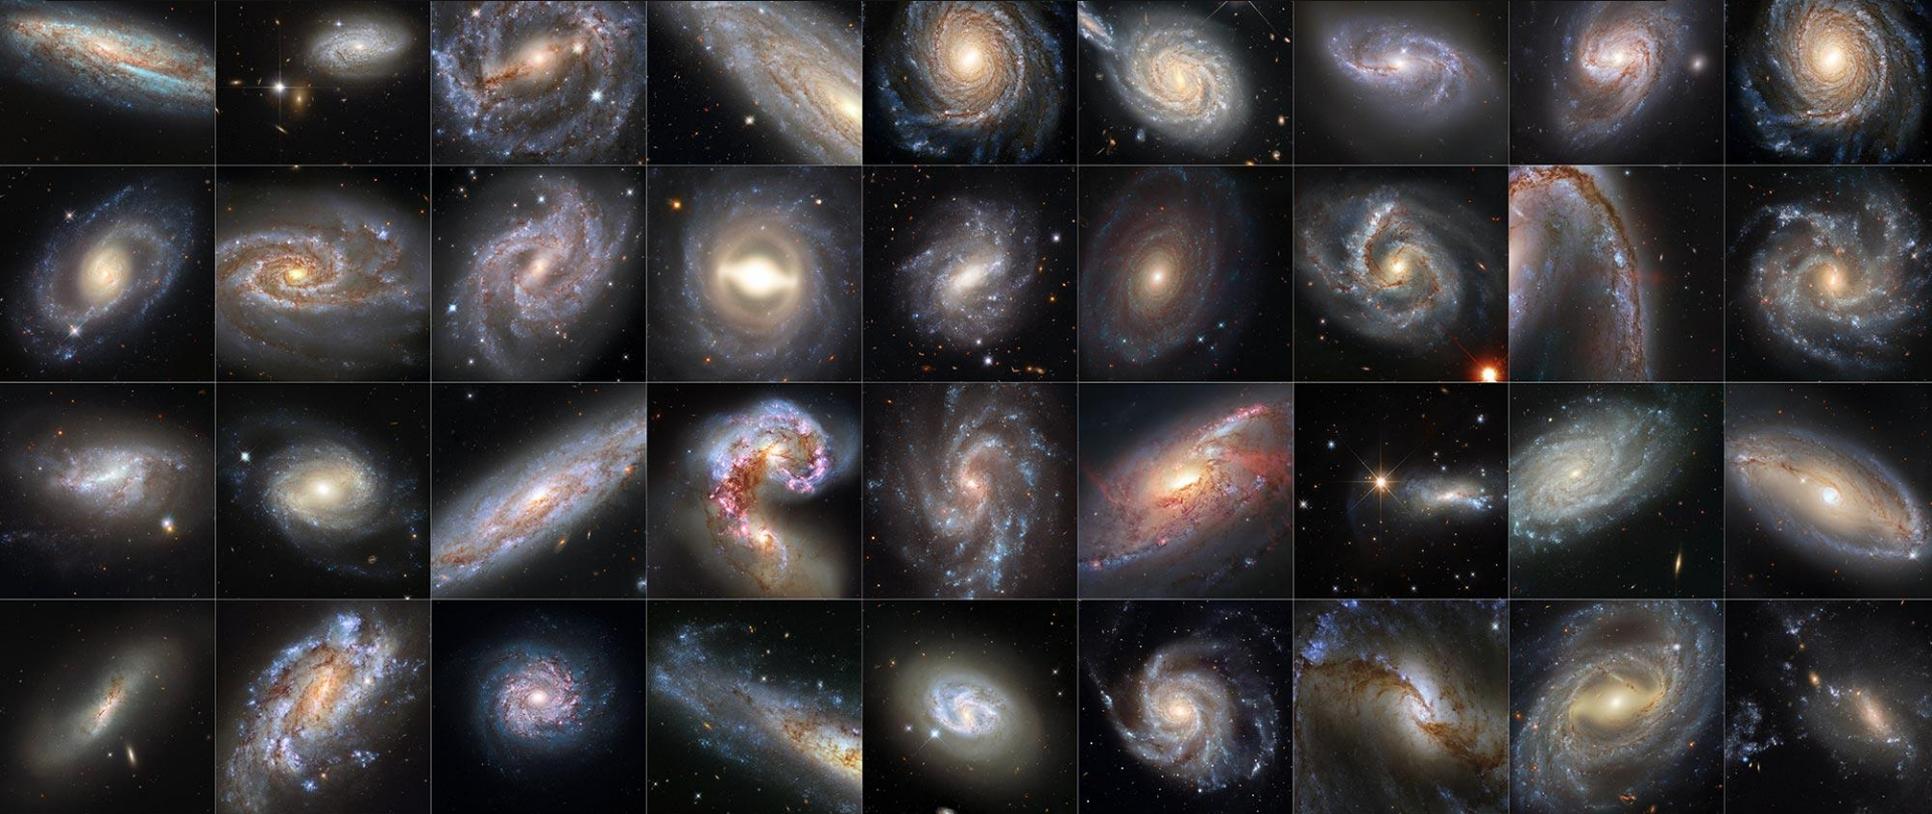
In 2023, a neural network (with human supervision) identified 527 new strong lensing candidates, a rare event where light bends around heavy objects like galaxies.





Neural networks have also been used to identify 301 exoplanets

Neural networks are challenging some of our theories of galaxy evolution by classifying the galaxies differently... new physics?



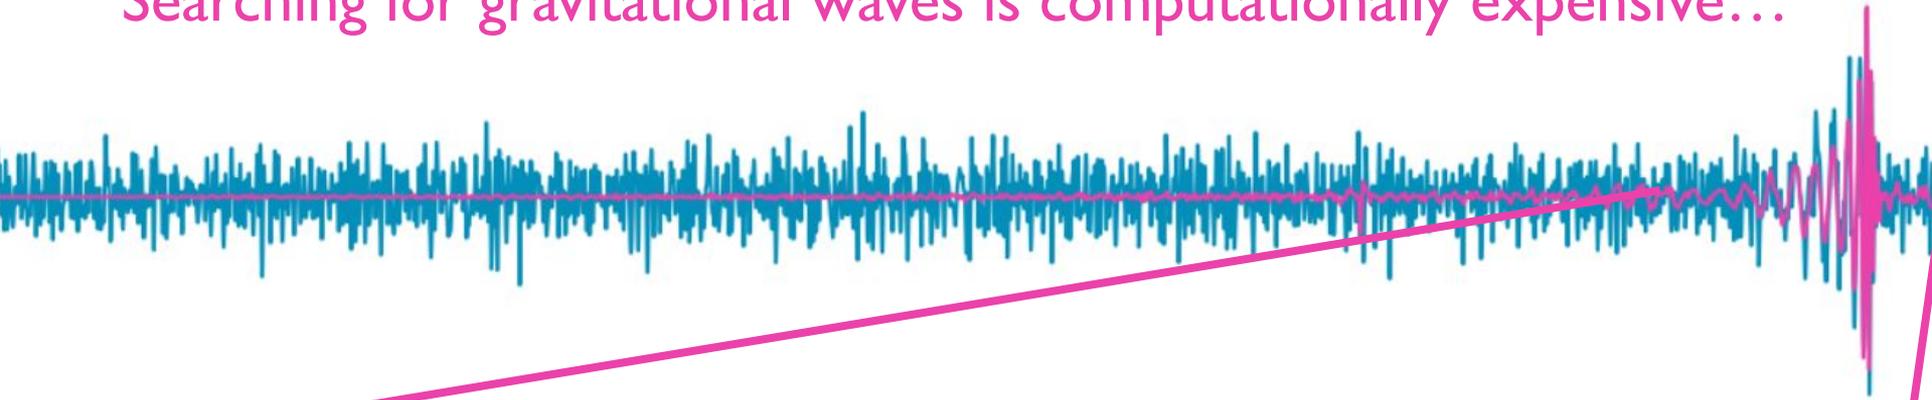
LearningMatch

Using neural networks to predict the 'match'

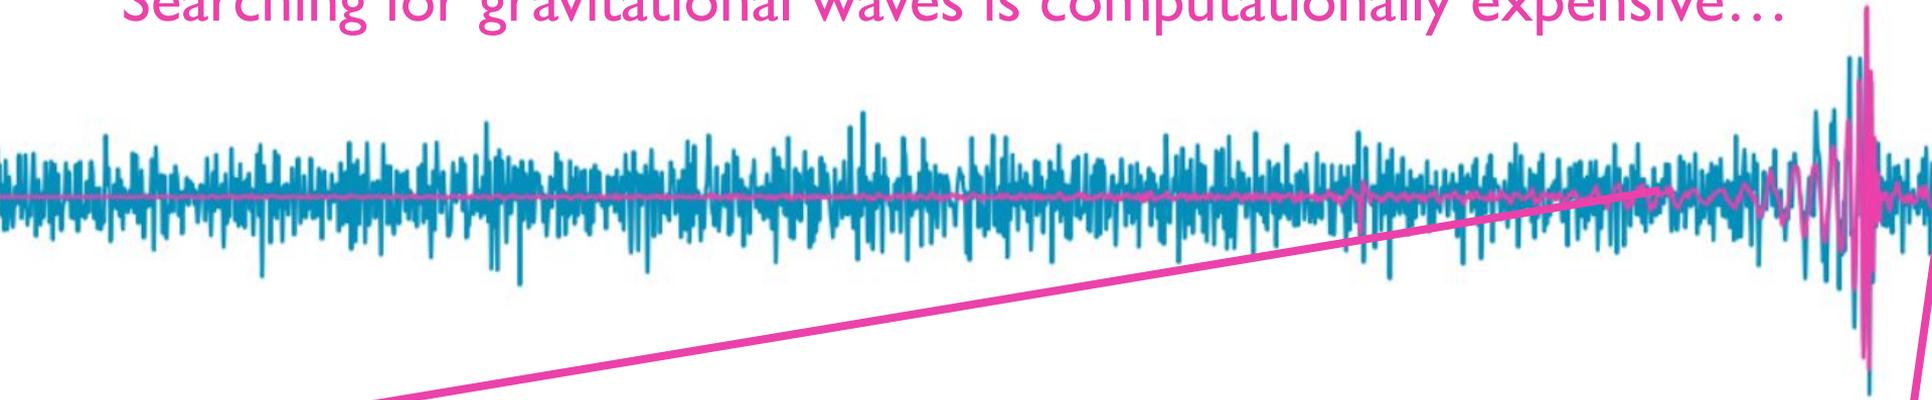
Gravitational Waves are ripples in the curvature of spacetime produced by **accelerating masses**... such as two **black holes** or two **neutron stars** colliding.



Searching for gravitational waves is computationally expensive...



Searching for gravitational waves is computationally expensive...



...so we want to use neural networks to speed up the process.



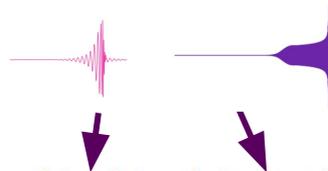
What is the ‘match’?

What is the 'match'?

* is the Complex conjugate of a function

~ is the Fourier Transform

Templates - which are a function of mass, spin and other parameters



$$\mathcal{M}(h_1, h_2) = \max_{\phi_c, t_c} \langle h_1 | h_2(\phi_c, t_c) \rangle$$

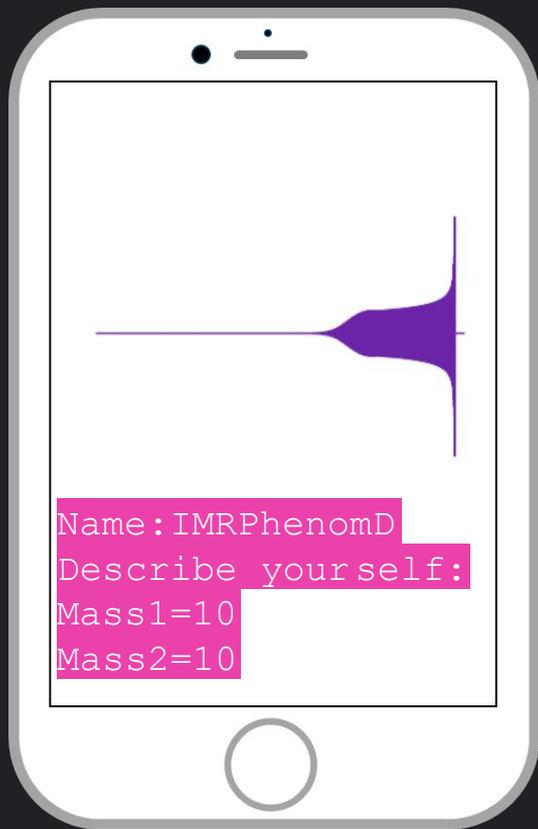
$$= \max_{\phi_c, t_c} \left[4\Re \int_0^\infty \frac{\tilde{h}_1^*(f) \tilde{h}_2(f)}{S_h(f)} \right]$$

One-sided Power Spectral Density (PSD) - average noise in our detectors

\\ The match is the weighted inner product maximized over the phase and the time of coalescence.

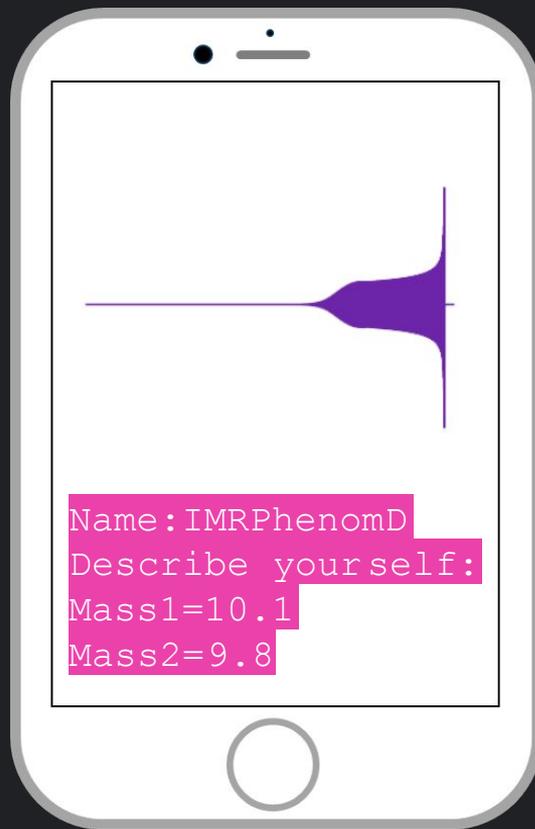
What is a good 'match'?

// It's like gravitational wave online dating!



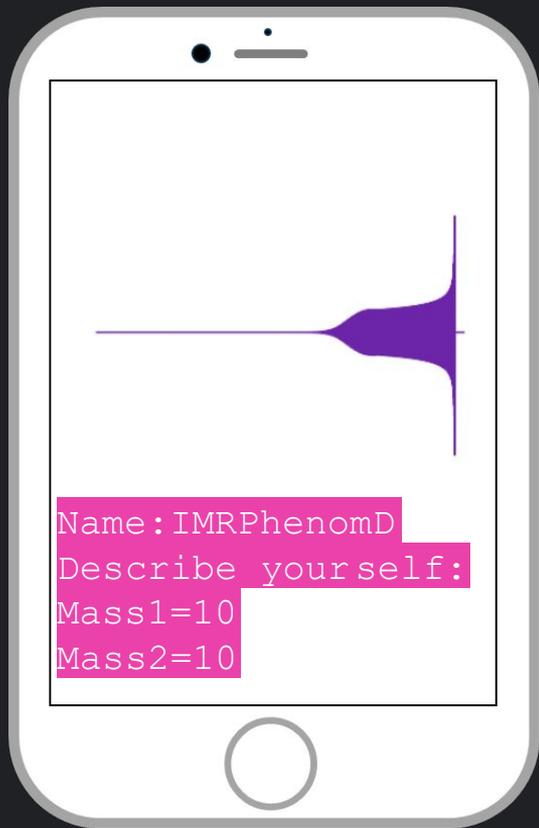
**STRONG
MATCH**

**= 0.99
(or ~ 1)**



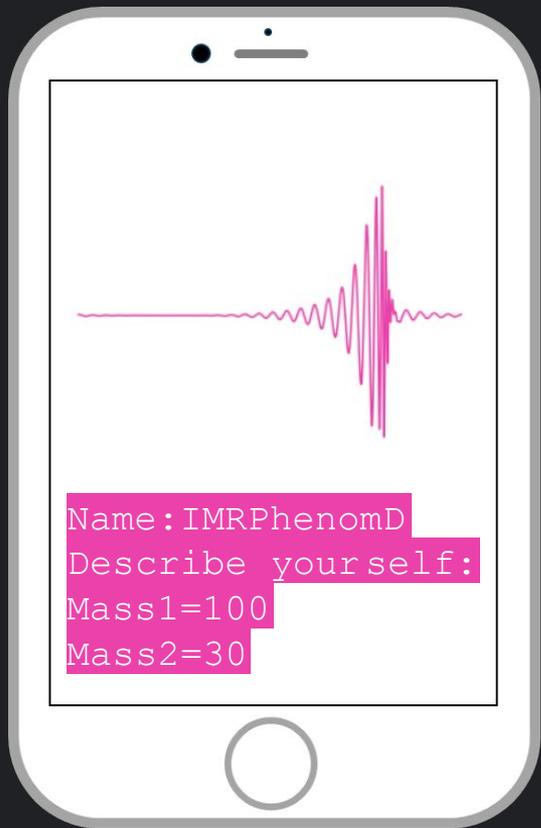
What is a good 'match'?

// It's like gravitational wave online dating!

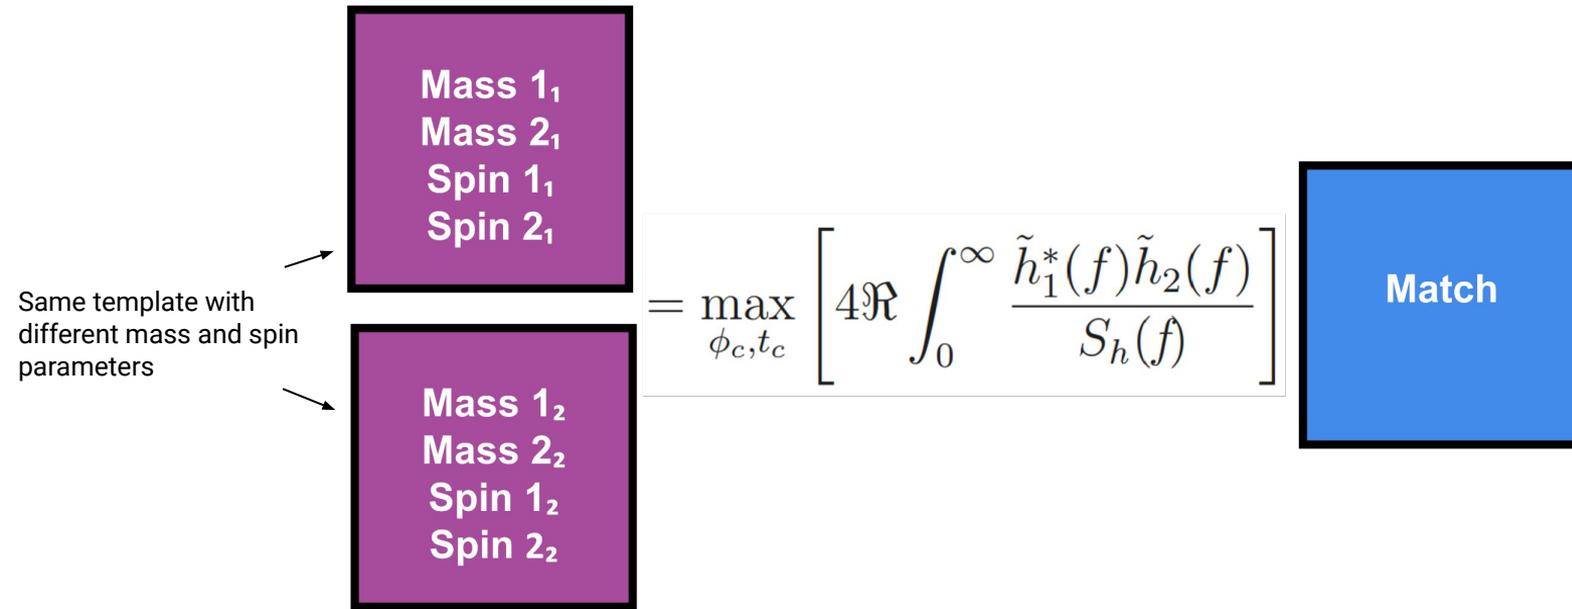


**WEAK
MATCH**

= 0.01

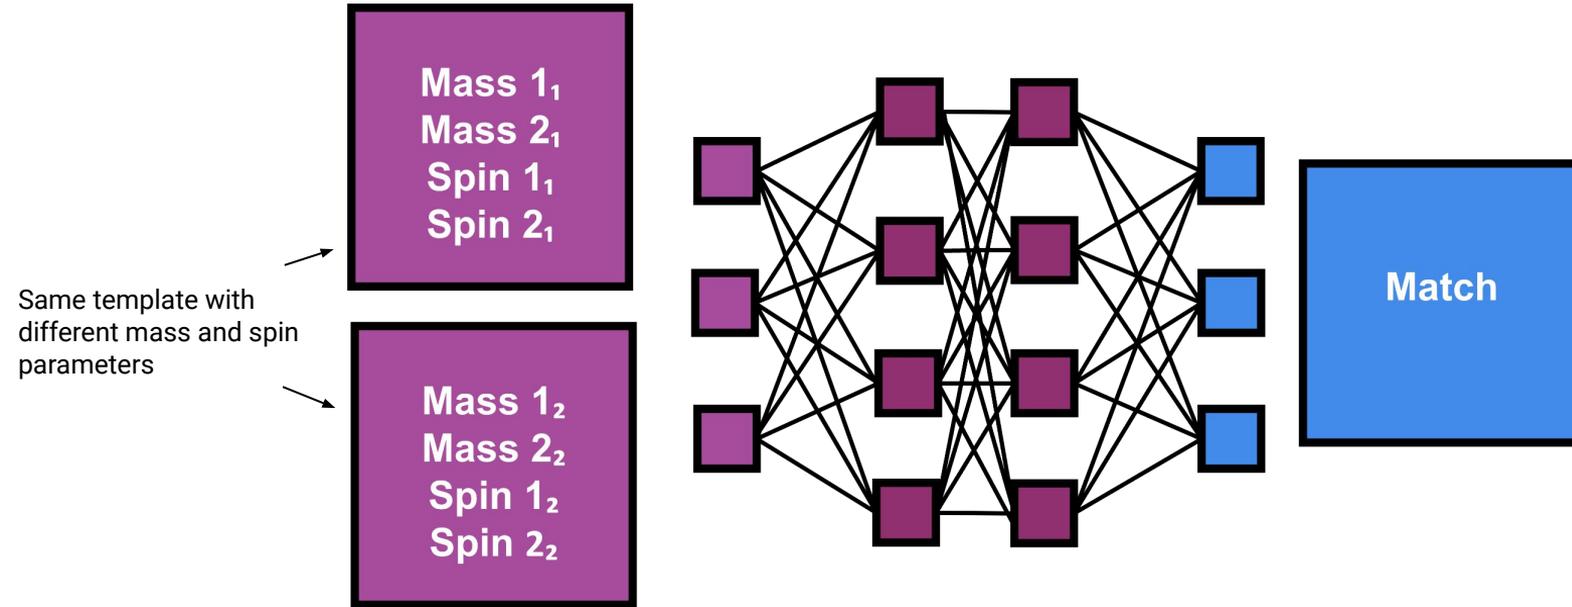


The Aim:



The aim of the neural network is to learn the relationship between the parameters and the match.

The Aim:



The aim of the neural network is to learn the relationship between the parameters and the match.

What have I **learned** as a machine learning engineer?

- 1) **Unknown** - We need the neural network to be accurate within 1%, which hasn't been done before.

What have I **learned** as a machine learning engineer?

- 1) **Unknown** - We need the neural network to be accurate within 1%, which hasn't been done before.
- 2) **Misleading** - What has the neural network actually learned?

What have I **learned** as a machine learning engineer?

- 1) **Unknown** - We need the neural network to be accurate within 1%, which hasn't been done before.
- 2) **Misleading** - What has the neural network actually learned?
- 3) **Frustrating** - You are looking for a 'needle in a haystack.'

What have I **learned** as a machine learning engineer?

- 1) **Unknown** - We need the neural network to be accurate within 1%, which hasn't been done before.
- 2) **Misleading** - What has the neural network actually learned?
- 3) **Frustrating** - You are looking for a 'needle in a haystack.'

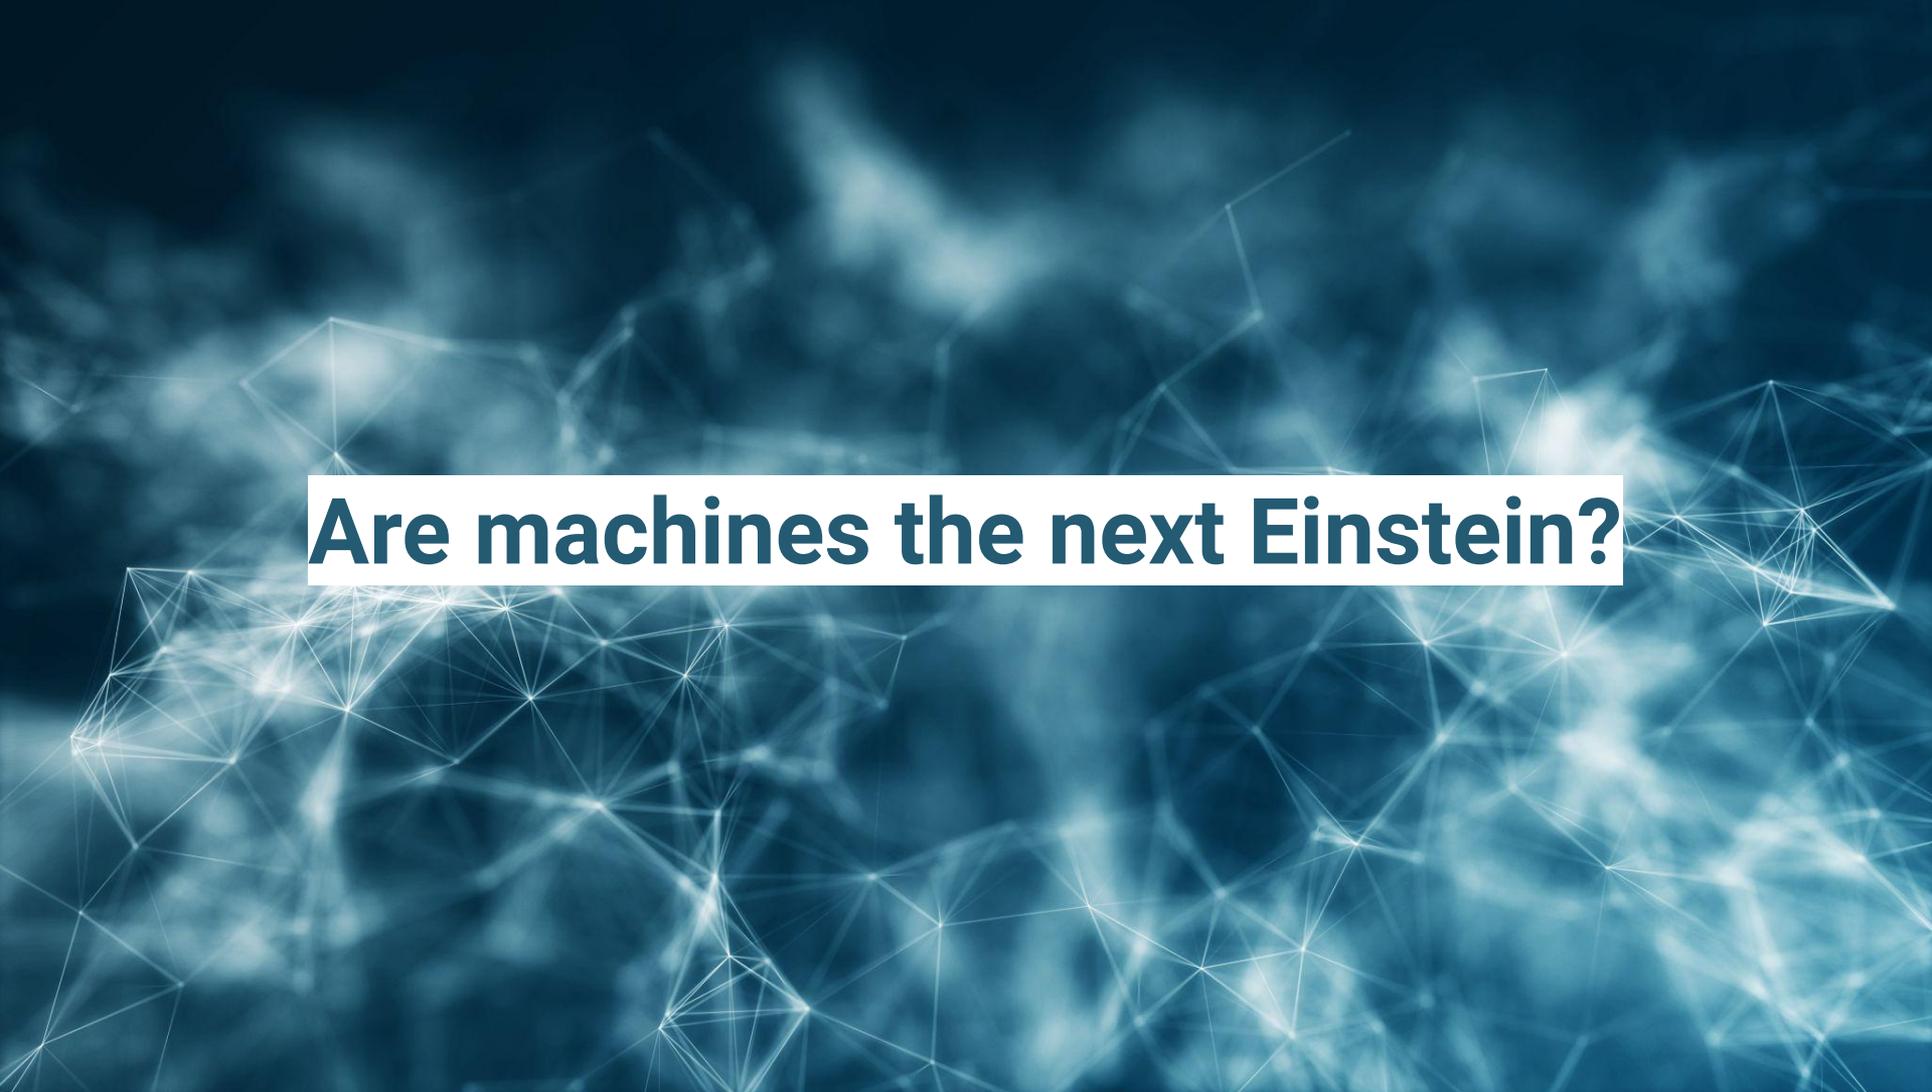
...this is why machine learning engineers are paid nice salaries!

Using neural networks in astronomy is still current research...



Using neural networks in astronomy is still current research...

...so who knows what the **future** looks like!



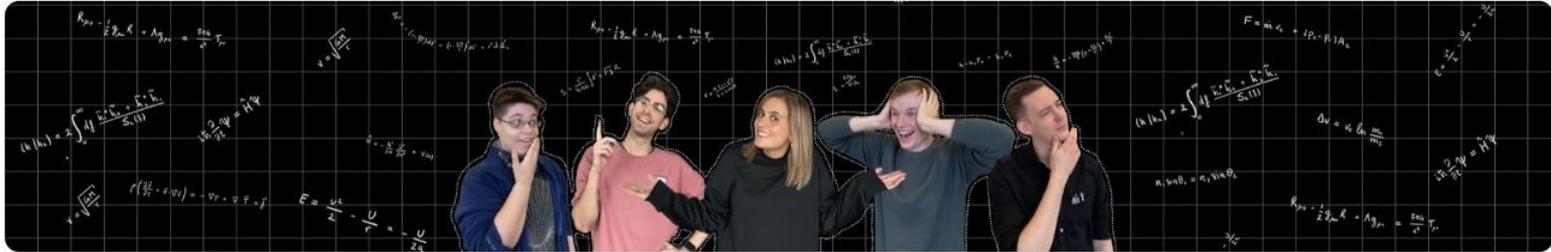
Are machines the next Einstein?

Thank you for listening and are there any questions?



Scan here to get involved with an astronomy project from your phone... we need you!

Get involved at zooniverse.org!



Field of View

@fieldof_view · 242 subscribers · 34 videos

Talking to scientists to understand not only what they're working on, but what inspired the... >

[instagram.com/physicschat_fieldofview](https://www.instagram.com/physicschat_fieldofview) and 2 more links

Subscribe

Home Videos Playlists Community



Like and subscribe!

For you



Creating Planets Made Easy! Physics Chat with Tom Haworth

294 views · 1 month ago



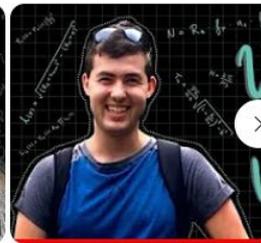
Physics Chat Live! | Stargazing vlog

242 views · 1 year ago



What came first? Relativity or Reality? Physics Chat with Ashim Sen Gupta

370 views · 2 months ago

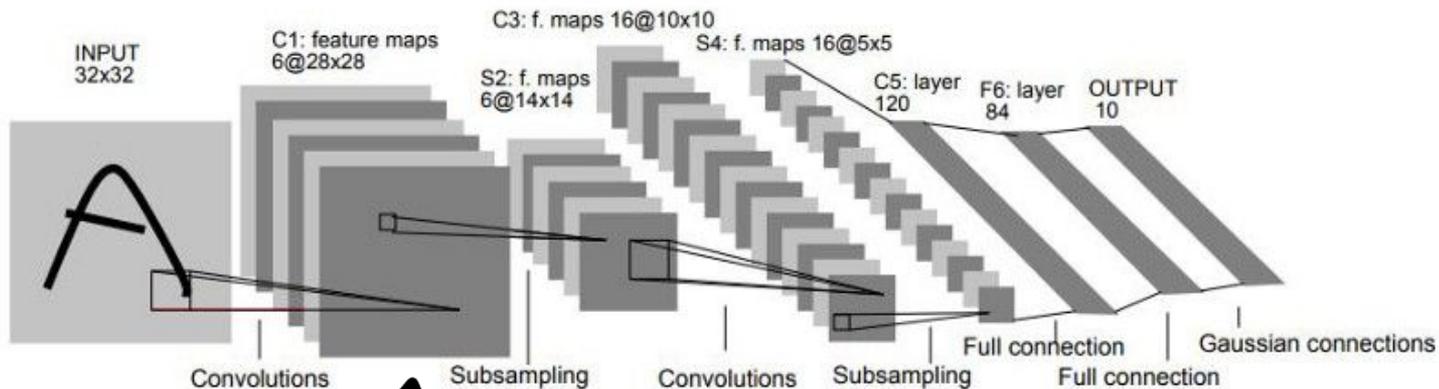


It's life, Luke, but not as we know it with Luke Booth

172 views · 8 months ago



Additional Slides



Credit: LeCun et al. (1998)

Convolutional layer

Neural Network

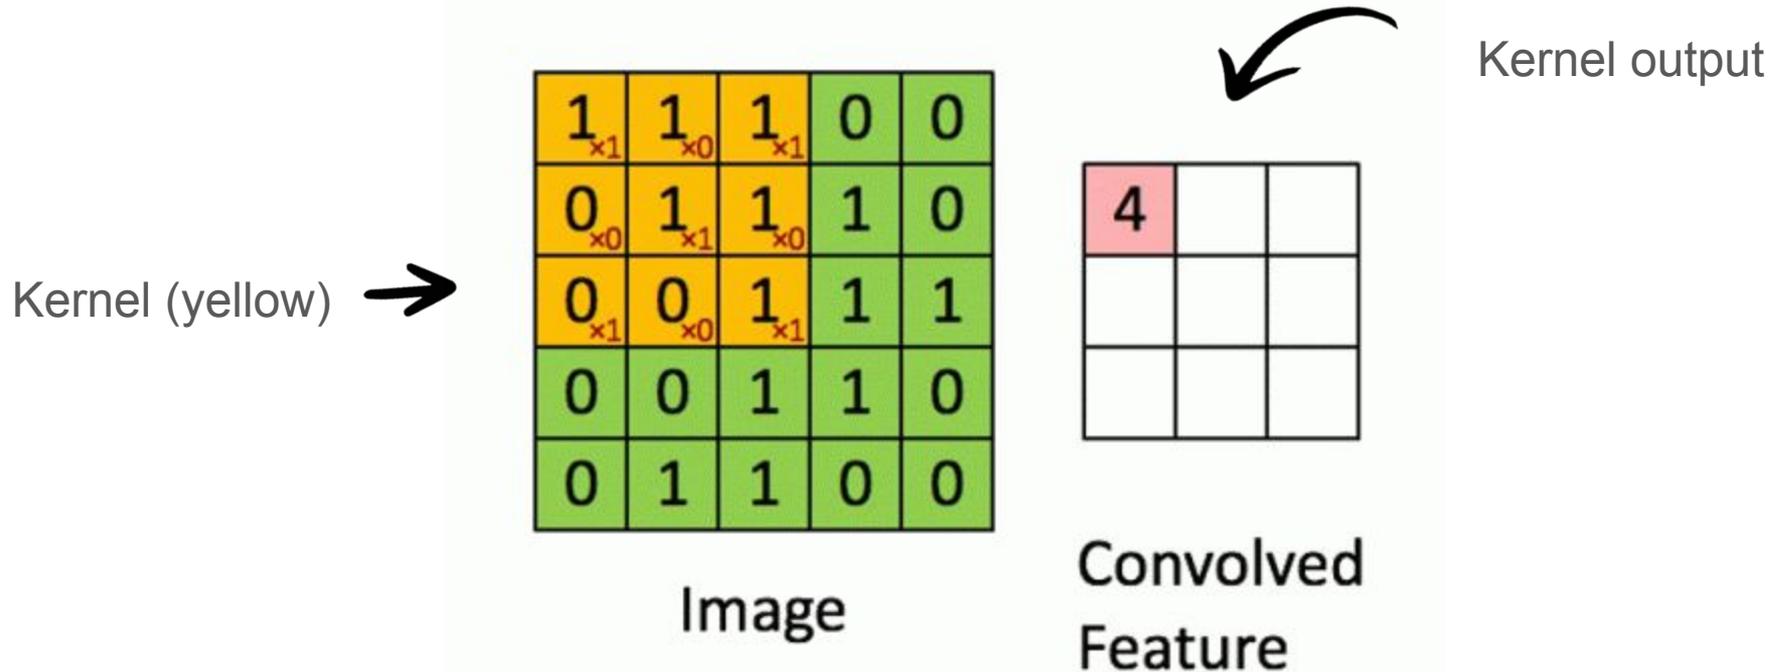
Image

Feature
Extraction

Neural Network

Output

What is a convolutional layer?



A convolutional layer extracts certain information from the image (i.e. lines at a particular orientation). The output is what the NN sees.

Feature Visualization

How neural networks build up their understanding of images



Edges (layer conv2d0)

Textures (layer mixed3a)

Patterns (layer mixed4a)

Parts (layers mixed4b & mixed4c)

Objects (layers mixed4d & mixed4e)

Check out 'Colah's blog' for more information on convolutional layers!