

AI for cyber-physical systems

Part 2 – Neural Networks & Deep Learning

Prof. Congduc Pham
<http://www.univ-pau.fr/~cpham>



European Commission

Horizon 2020
European Union funding
for Research & Innovation



IoT – from idea to reality



Paving for the next 10 years
of innovation in IoT and AI



Intel-IrriS

RESILINK


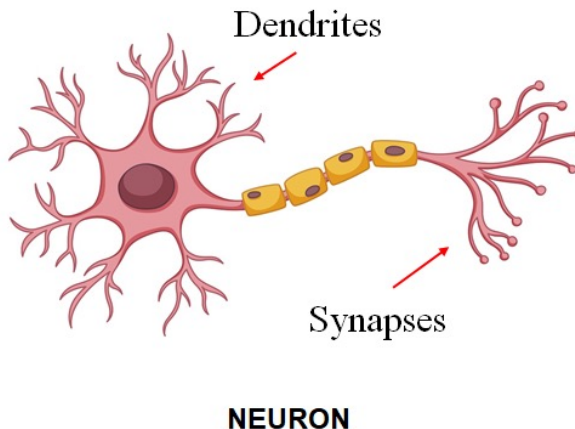
Advanced and disruptive IoT/AI technologies targeting
the smallholder community for increased resilience

Acknowledgements

- AI is probably one of the hot topics with the most incredible amount of also incredible quality contributions & tutorials videos!
- Setting up this course would not have been possible without all these contributors!
- I've crawled the web for hours (days), trying to take knowledge from multiple sources and trying to present all these impressive resources in a simple and comprehensive way for my students
- I tried to mention the sources of materials I've borrowed from the web as much as I could
- A special "thank you" for all the incredible tutorial videos from Machine Learnia (<https://machinelearnia.com/>, Guillaume Saint-Cirgue) and his YouTube channel (<https://www.youtube.com/@MachineLearnia/videos>)
- A special "thank you" akso for the complete and very clear videos from freeCodeCamp.org, especially on DL (<https://www.youtube.com/watch?v=dPWYUELwldM>)
- There are also all the articles from <https://towardsdatascience.com/>

How does a Perceptron work?

- Single Layer Perceptron (SLP) receives the value of the attributes of an input, just as dendrites do in a neuron
- Each attribute has a **weight** w that measures its *contribution* to the final result, which is the sum of the multiplications of inputs of each attribute by its corresponding weight
- The final output depends on the activation function: here, if the sum is >0 Perceptron returns a value of 1, otherwise it yields 0





What can an SLP do?

Weights


$$\hat{y} = \Theta(w_1x_1 + w_2x_2 + \dots + w_nx_n + b)$$

$$= \Theta(\mathbf{w} \cdot \mathbf{x} + b)$$

$$\text{where } \Theta(v) = \begin{cases} 1 & \text{if } v \geq 0 \\ 0 & \text{otherwise} \end{cases}$$




Linear boundary





Non-linear boundary

- SLP is the simplest type of artificial neural networks and can only classify **linearly separable cases** with a binary target
- Activation functions are mathematical equations that determine the output of a neural network
- Here, the **Heaviside step function** Θ




Multi-Layer Perceptron – MLP

- Multi-Layer Perceptrons (and so-called neural networks) can be considered as more general function approximators and they are able to distinguish data that is not linearly separable
- MLP have several neuron layers with so-called hidden layers
- Hidden layers are "hidden" to the final user who only "see" inputs and outputs



Learning in MLP?

- With more layers, there are more parameters to optimize, so training an MLP is obviously more complex and will take much much more time, if not done in an intelligent way!
- A fundamental method in neural network is back-propagation!

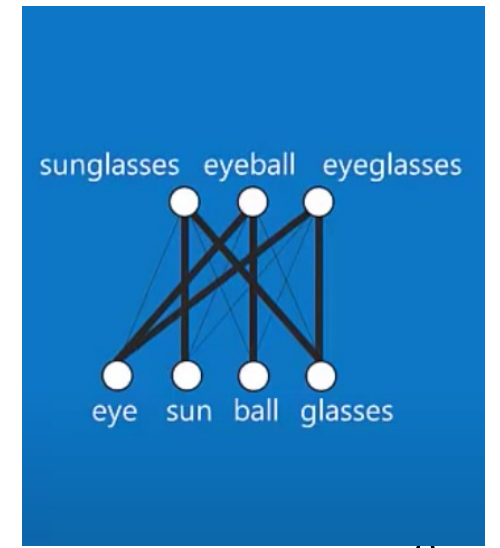

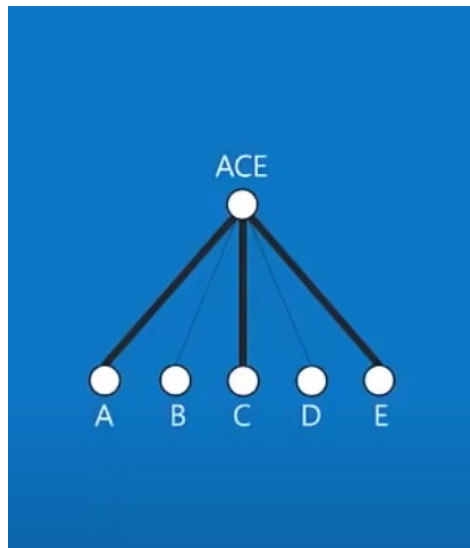









A Deep Dive in Neural Networks

Neural Network in a nutshell


Pham
ul.fr/~cpham



Pictures from <https://www.youtube.com/watch?v=Q9Z20HCPnww> (Brandon Rohrer)

Example 1: 4-pixel camera

Categorize images




solid

vertical

diagonal

horizontal

Categorize images




solid

vertical

diagonal

horizontal

Categorize images




solid

vertical

diagonal

horizontal

Categorize images



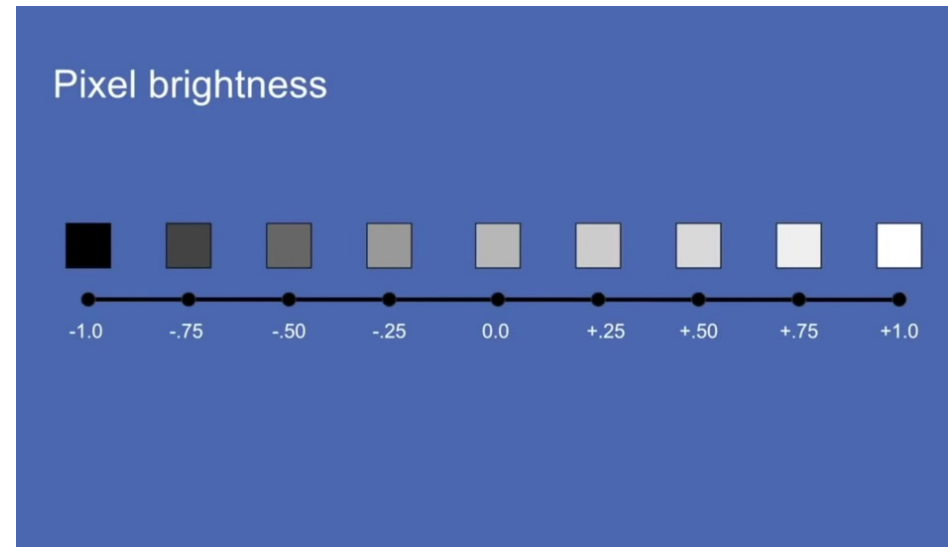

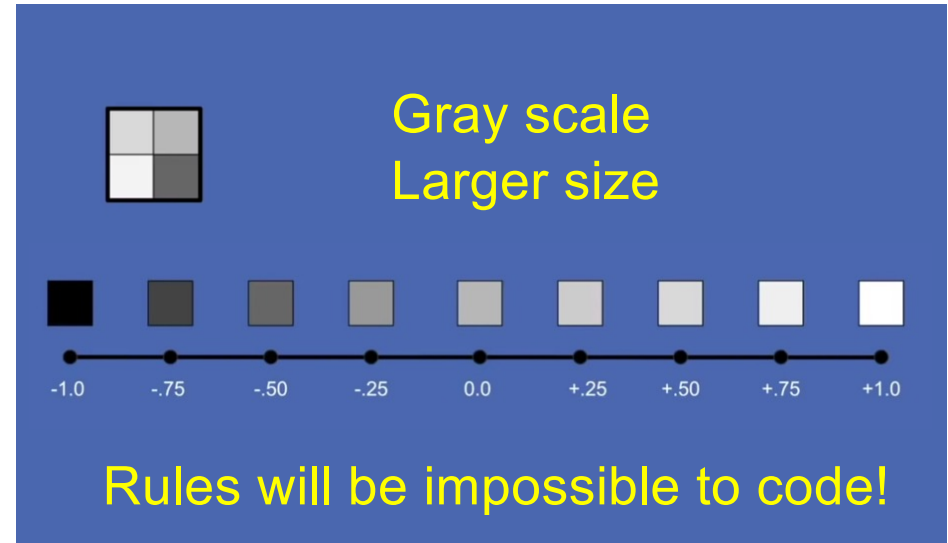
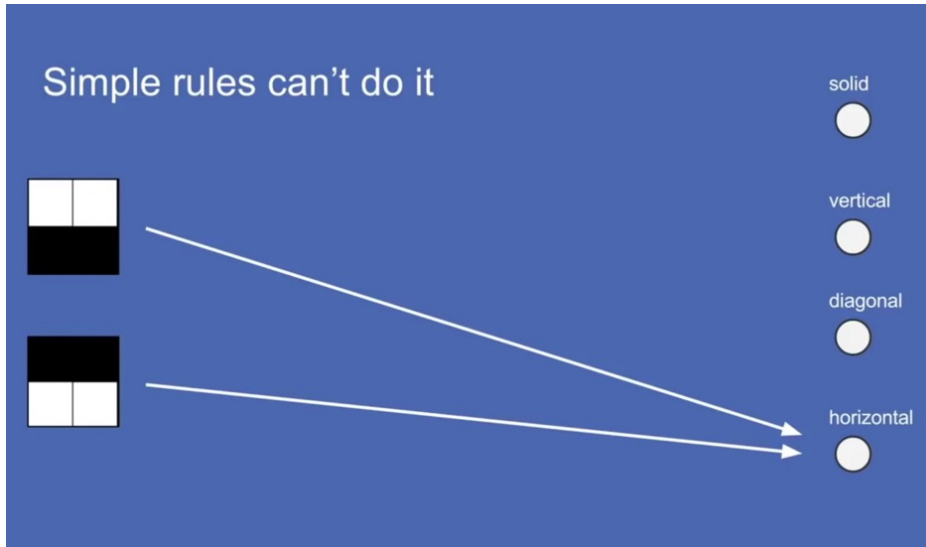
solid

vertical

diagonal



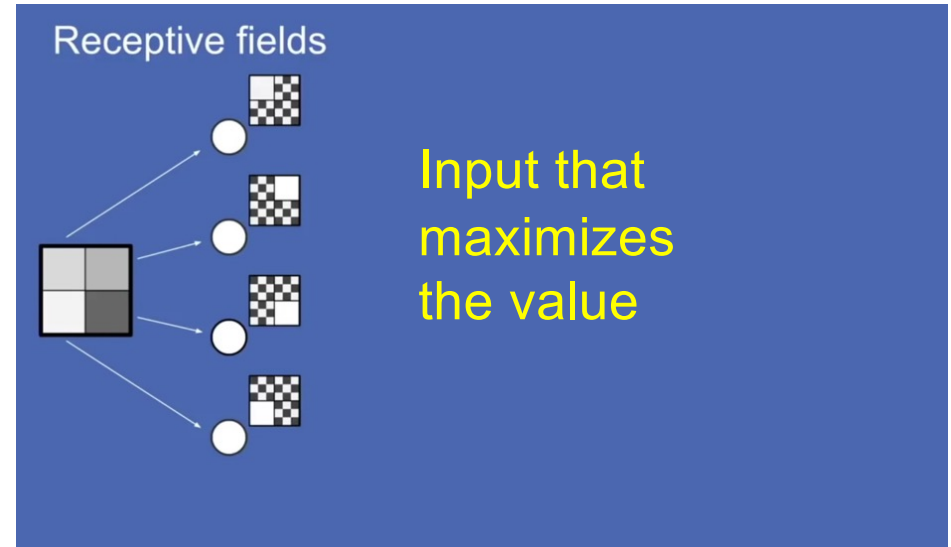
horizontal

Example 1, con't

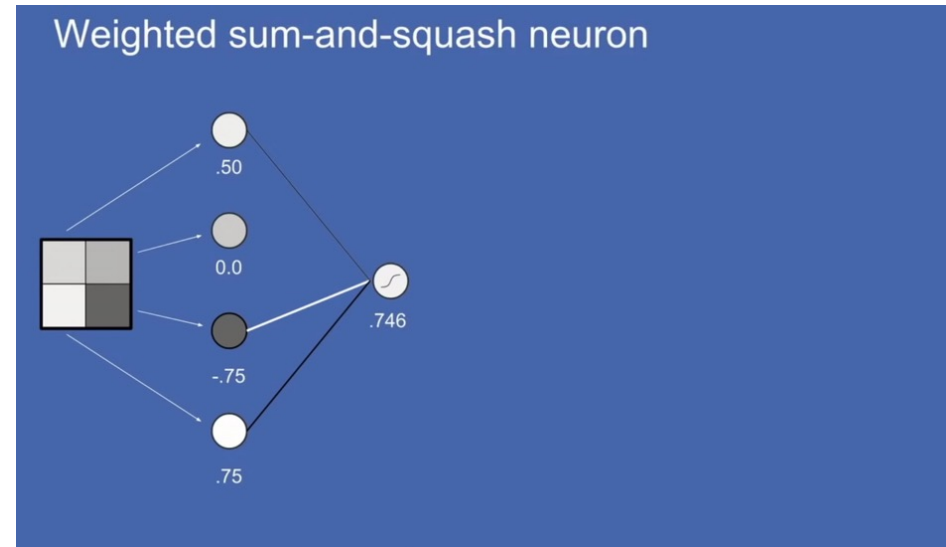
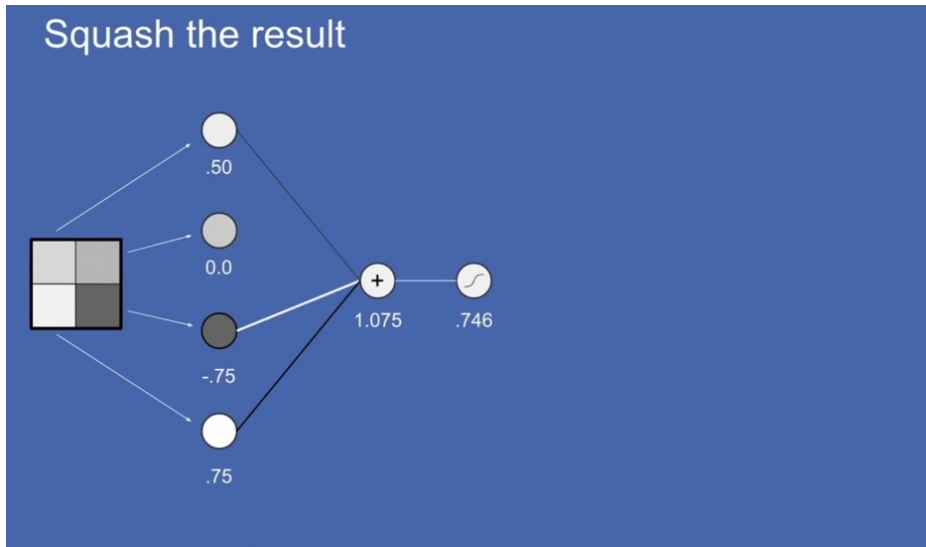
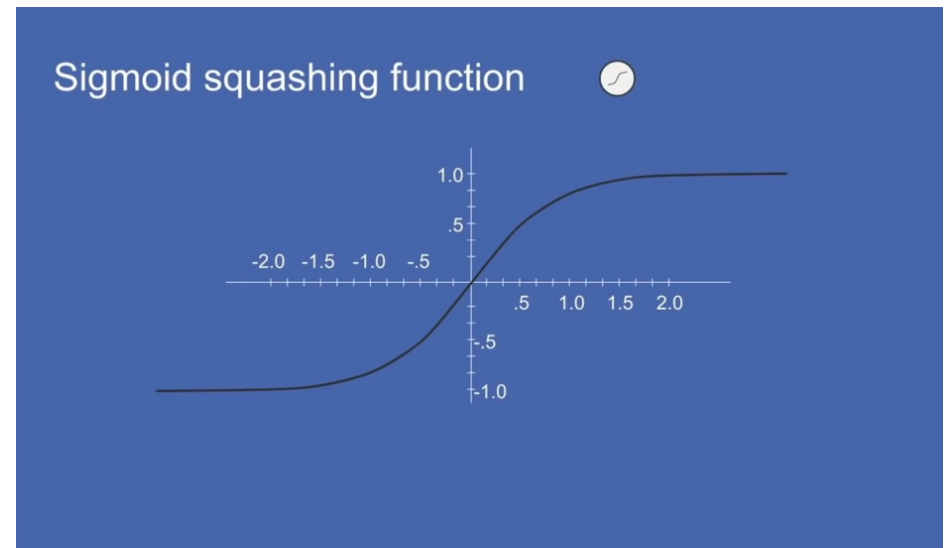
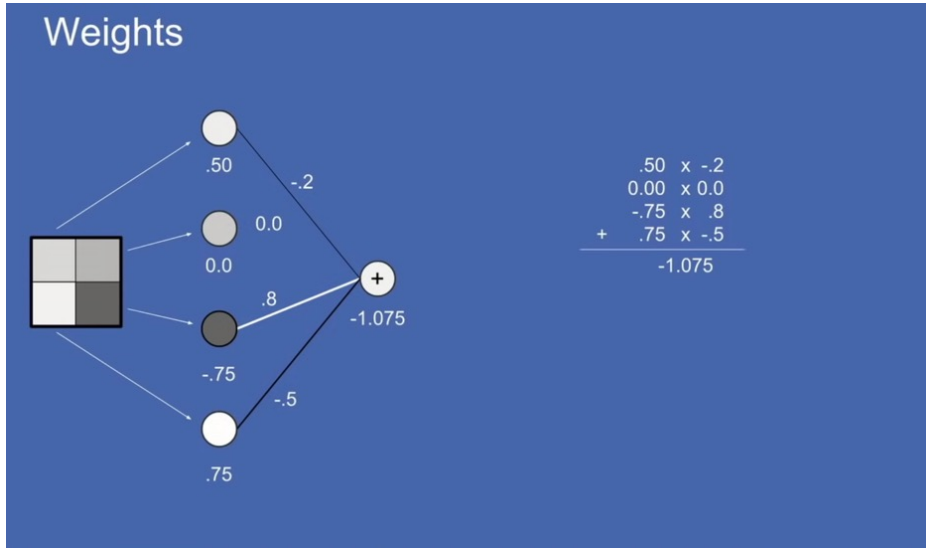


Example 1, con't

Pr. Congduc Pham
http://www.univ-pau.fr/~cpham




Example 1, con't




Example 1, con't

Make lots of neurons, identical except for weights




To keep our picture clear, weights will either be
1.0 (white)
-1.0 (black) or
0.0 (missing)

Receptive fields get more complex




Input that maximizes the value



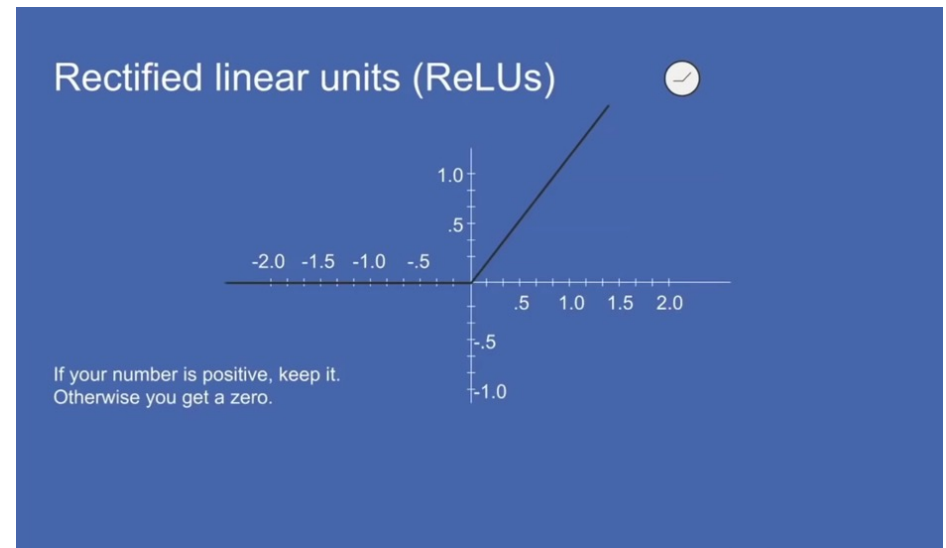

Repeat for additional layers



Receptive fields get still more complex







Example 1, con't




Example 1, con't

Pr. Congduc Pham
<http://www.univ-pau.fr/~cpham>



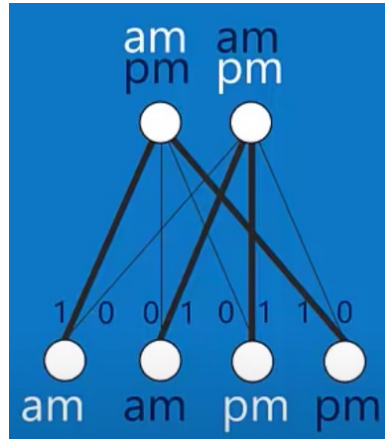




Example 1, finally!



Example 2: Shawarma guy example

Pictures from <https://www.youtube.com/watch?v=Q9Z20HCPnww> (Brandon Rohrer)




Pr. Congduc Pham
http://www.univ-pau.fr


Sometimes he works mornings, sometimes evenings
light=working; dark=off

Start with random weights


Observed on a given day



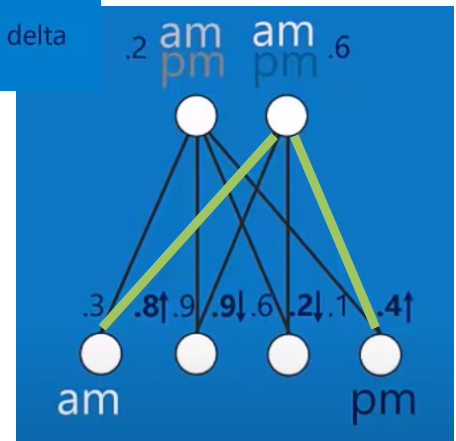
Error = 1 - .2 = .8



Error = 1 - .6 = .4




Here comes again our Gradient Descent methods!




Example 2, con't

Pictures from <https://www.youtube.com/watch?v=Q9Z20HCPnww> (Brandon Rohrer)


Pr. Congduc Pham
http://www.u




am
pm Increase active links, decrease inactive links




Error = $1 - 0.7 = 0.3$
-> better!




am
pm The next day




am
pm The next day



am
pm The next day




am
pm The next day.....



Eventually, weights don't change anymore!

At this point you have a trained model!




DEEP NEURAL NETWORKS


Deep Neural Networks (DNN)

- Deep Neural Networks are Neural Networks with a large number of hidden layers
- More hidden layers can create/store more internal abstract representation of the training data
- But optimization/learning is more difficult!

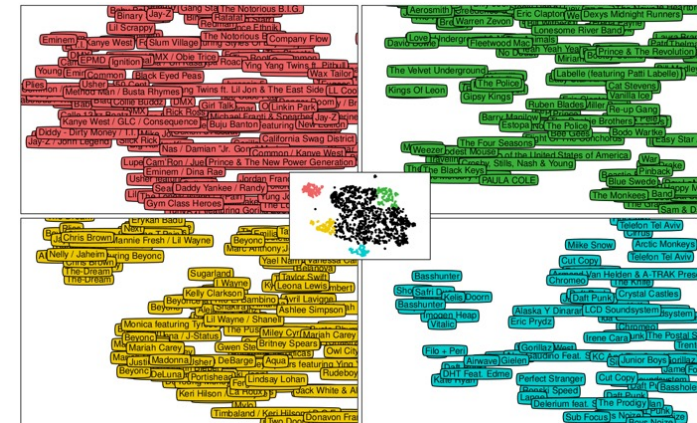

Simple Neural Network



Deep Learning Neural Network




● Input Layer ● Hidden Layer ● Output Layer





Review: Learning in MLP?

- With more layers, there are more parameters to optimize, so training an MLP is obviously more complex and will take much much more time, if not done in an intelligent way!
- A fundamental method in neural network is **back-propagation!**



Look back at the NN of example 1

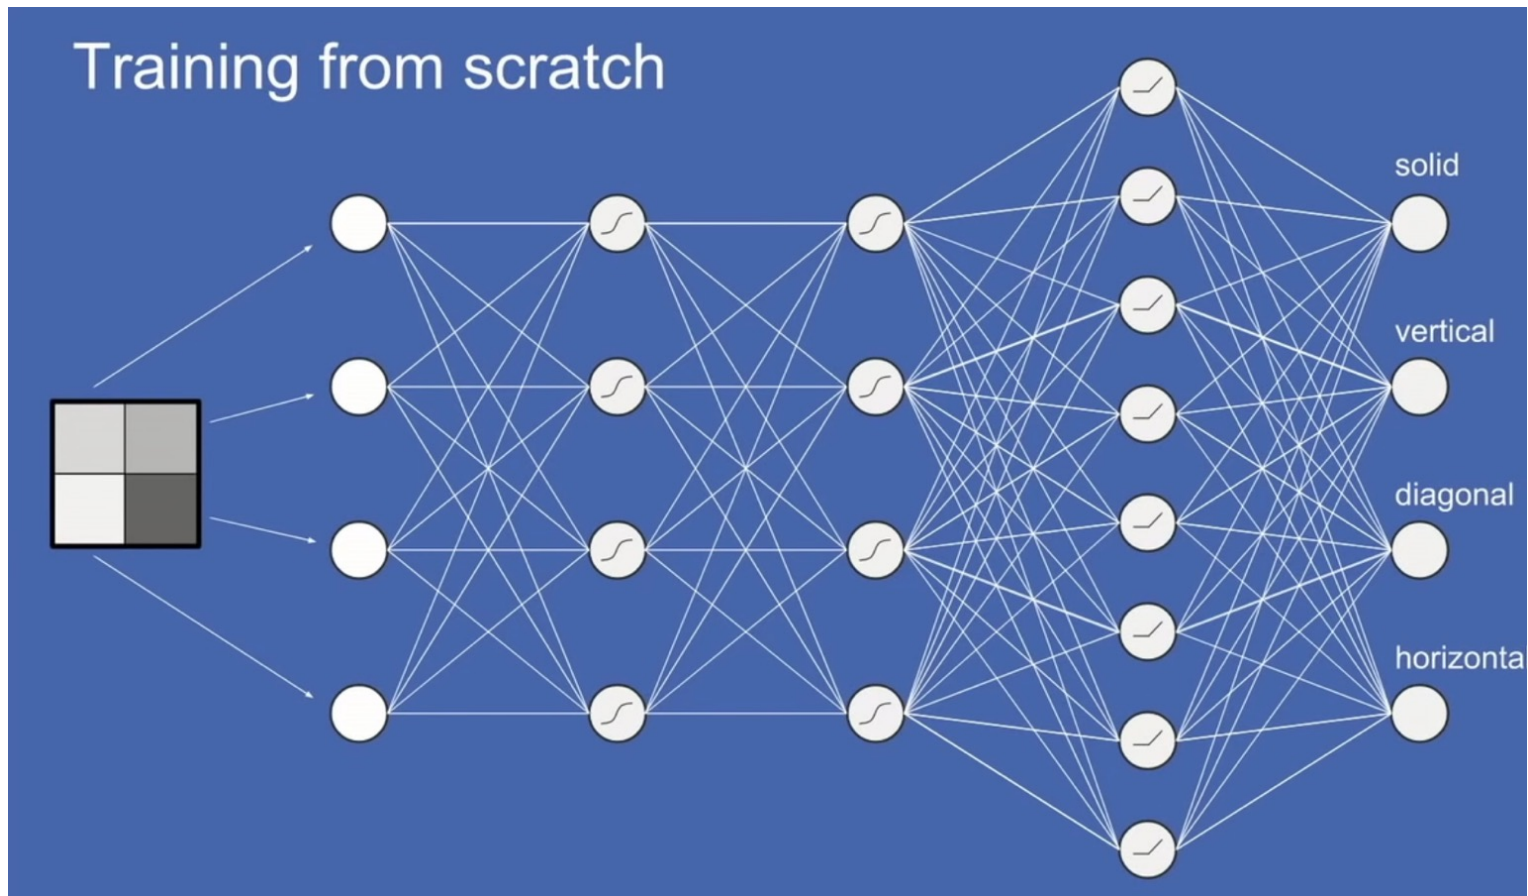
- ⦿ In example 1, the NN has already been trained so that the weights have already been assigned accordingly!



To keep our picture clear, weights will either be
1.0 (white)
-1.0 (black) or
0.0 (missing)

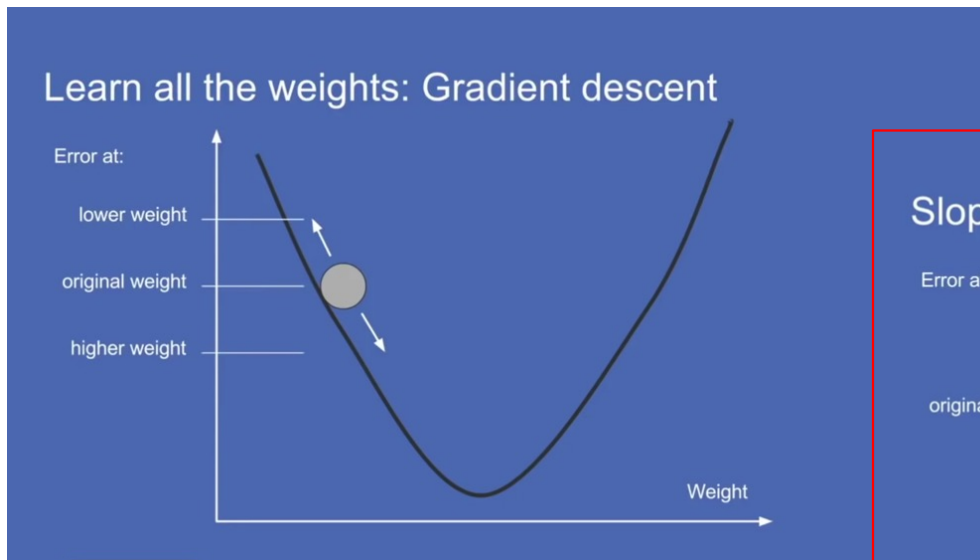
But this was the initial NN!

- ⦿ Fully connected, and weights needed to be trained!

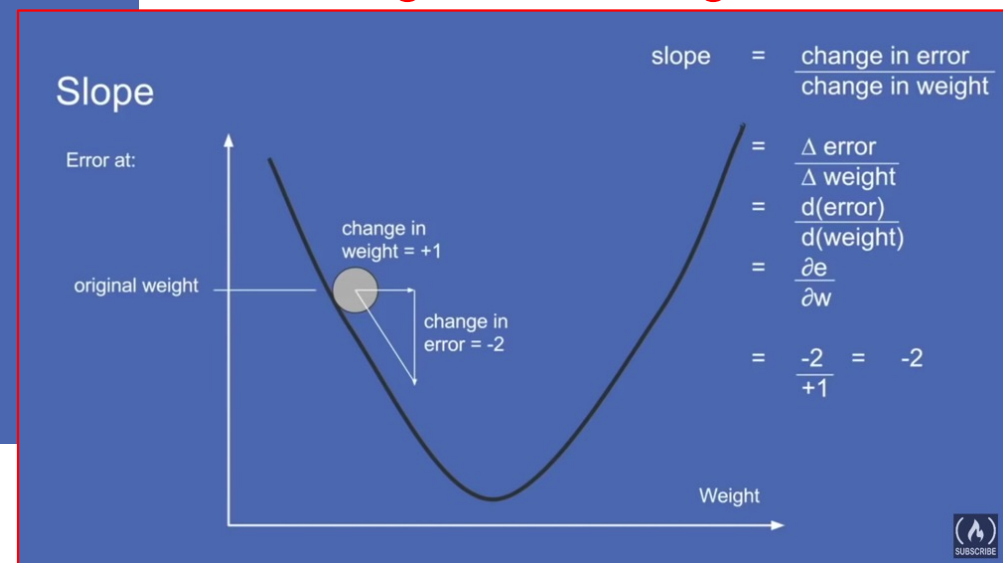


How can we train an DNN?

- ⦿ We saw that Gradient Descent is an efficient way to "smartly" find the optimal weights when trying to minimise the error




What is the change in error with the change in the weights?



- ⦿ But in a DNN, we would still need to iterate on all neurons to adjust any single weight in any single hidden layer!

Which direction is downhill?

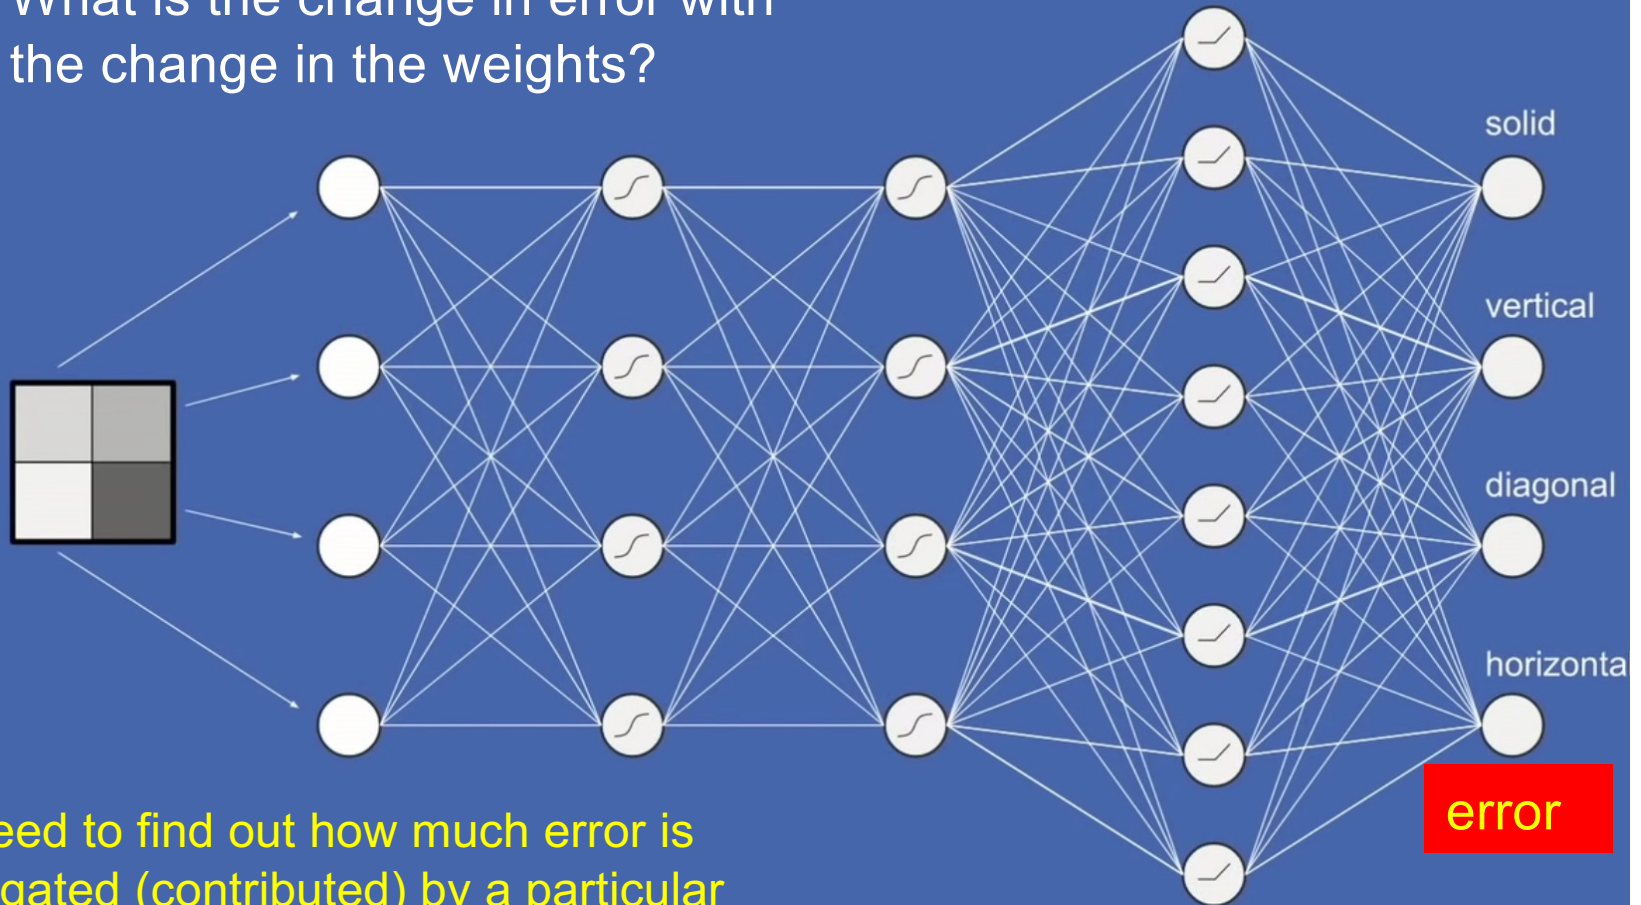
- ⦿ In multi-dimensional problem, the loss function can be quite complex



- ⦿ There can be millions of weights to adjust so calculating the gradient (slope) may require millions of passes through the neural network to find out which direction is downhill!
- ⦿ Therefore, we need to "help" GD to converge faster!

The challenge in adjusting the weights!

What is the change in error with the change in the weights?



We need to find out how much error is propagated (contributed) by a particular weight to the total error of the network!


Chaining principle

- Chaining illustrated on a simple case

Chaining

We know e (error) but we need to know what is the change in e with a change in w_1

how much a change in w_1 affects e ?



$$y = x * w_1$$

$$\frac{\partial y}{\partial w_1} = x$$

$$e = y * w_2$$

$$\frac{\partial e}{\partial y} = w_2$$

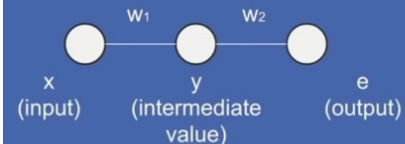
$$e = x * w_1 * w_2$$

$$\frac{\partial e}{\partial w_1} = x * w_2$$

$$\frac{\partial e}{\partial w_1} = \frac{\partial y}{\partial w_1} * \frac{\partial e}{\partial y}$$

Chaining

- ⦿ If we know **which way we want to change the error**, chaining allows us to calculate **how much we change the weights to let that happen**



$$y = x * w_1$$

$$\frac{\partial y}{\partial w_1} = x$$

$$e = y * w_2$$

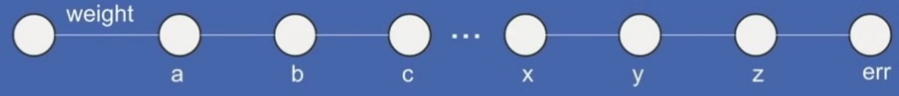
$$\frac{\partial e}{\partial y} = w_2$$

$$e = x * w_1 * w_2$$

$$\frac{\partial e}{\partial w_1} = x * w_2$$

$$\frac{\partial e}{\partial w_1} = \frac{\partial y}{\partial w_1} * \frac{\partial e}{\partial y}$$

Chaining

$$\frac{\partial err}{\partial weight} = \frac{\partial a}{\partial weight} * \frac{\partial b}{\partial a} * \frac{\partial c}{\partial b} * \frac{\partial d}{\partial c} * \dots * \frac{\partial y}{\partial x} * \frac{\partial z}{\partial y} * \frac{\partial err}{\partial z}$$


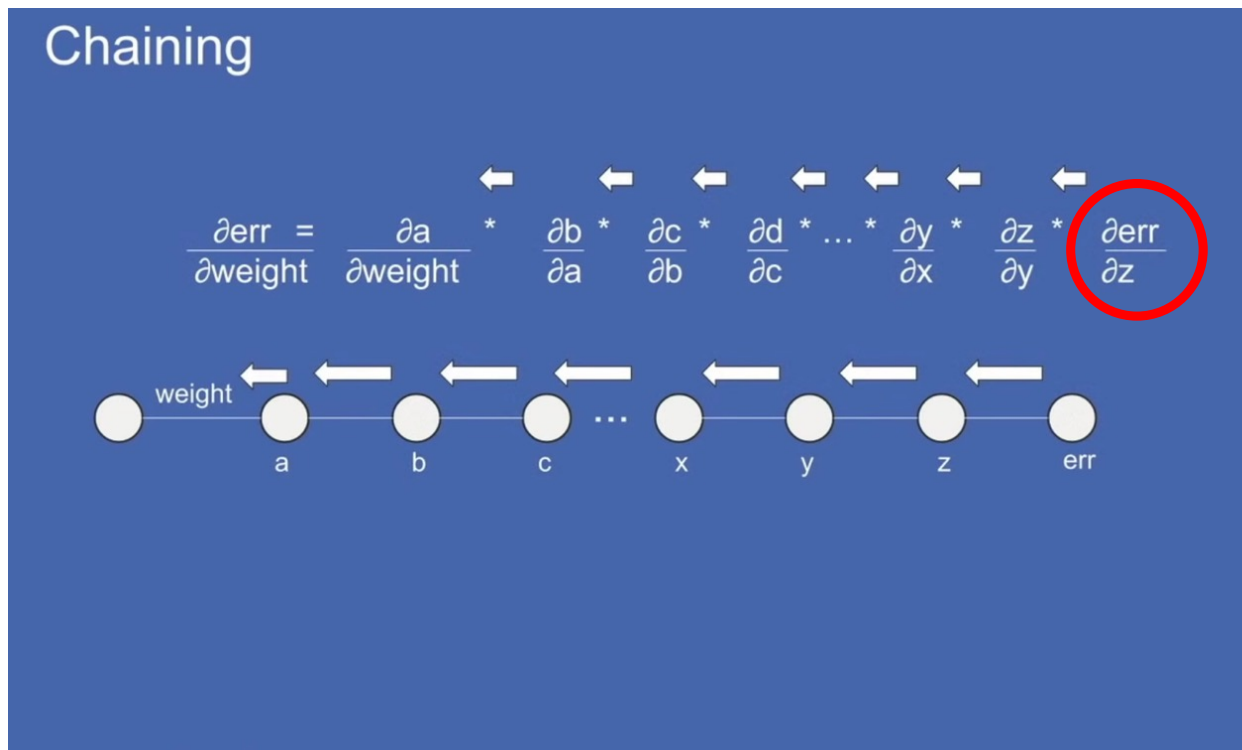
If I change *weight*, how much **a** change? If I change **a** how much **b** change? If I change **b** how much **c** change? ...

Here *weight* can be anywhere in the NN and we want to know how *err* will change (slope) if *weight* is changed

We can apply chaining again and again!

Backpropagation

- Backpropagation involves taking the error rate of a forward propagation and feeding this loss backward through the neural network layers to fine-tune the weights



We start with the *err* value at the end

then travel upward, back in the depth of the NN


Handling weights

- It is easy to handle weights: "How much b changes if a changes?" is just the weight, whatever it is

Backpropagation challenge: weights

If I change a how much err will change?

We know how much err will change if b changes



$$\frac{\partial err}{\partial a} = \frac{\partial b}{\partial a} * \frac{\partial err}{\partial b}$$

$$b = wa$$


$$\frac{\partial b}{\partial a} = w$$

How much b changes if a changes?

Handling sums

- ⦿ Handling sum is also trivial: "How much z changes if a changes?" is just 1

Backpropagation challenge: sums



If I change a how much err will change?

We know how much err will change if z changes

$$z = a + b + c + d + \dots$$

$$\frac{\partial z}{\partial a} = 1$$

$$\frac{\partial err}{\partial a} = \frac{\partial z}{\partial a} * \frac{\partial err}{\partial z}$$

How much z changes if a changes?

Handling activation function

- With Logistic function (Sigmoid): "How much b changes if a changes?" can be known using the derivative of Logistic

Backpropagation challenge: sigmoid

If I change a how much err will change?

We know how much err will change if b changes



$$\frac{\partial err}{\partial a} = \frac{\partial b}{\partial a} * \frac{\partial err}{\partial b}$$

$$b = \frac{1}{1 + e^{-a}}$$

$$= \sigma(a)$$

Because math is beautiful / dumb luck:

$$\frac{\partial b}{\partial a} = \sigma(a) * (1 - \sigma(a))$$

How much b changes if a changes?

Handling activation function

- With ReLU function: "How much b changes if a changes?" can be known using the derivative of ReLU

Backpropagation challenge: ReLU

If I change a how much err will change?

We know how much err will change if b changes



$$\frac{\partial err}{\partial a} = \frac{\partial b}{\partial a} * \frac{\partial err}{\partial b}$$


$$b = \begin{cases} a, & a > 0 \\ 0, & \text{otherwise} \end{cases}$$

$$\frac{\partial b}{\partial a} = \begin{cases} 1, & a > 0 \\ 0, & \text{otherwise} \end{cases}$$

How much b changes if a changes?

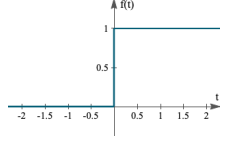
Back (again) to activation function

- ⦿ Remember this slides from part 1?
- ⦿ "The function is also required to be differentiable over the entire space of real numbers"
- ⦿ This is needed to make back-propagation works when we are doing the chain rule on the derivatives throughout all the layers of the Neural Network!



Gradient Descent & output functions


- ⦿ The only non-linear function that can be used as an activation function in a neural network is one which is monotonically increasing. So for example, $\sin(x)$ or $\cos(x)$ cannot be used as activation functions.
- ⦿ Also, the activation function should be defined everywhere and should be continuous everywhere in the space of real numbers. The function is also required to be differentiable over the entire space of real numbers
- ⦿ For instance Heaviside is not for $x=0$
- ⦿ Typically a back propagation algorithm uses gradient descent to learn the weights of a neural network. To derive this algorithm, the derivative of the activation function is required




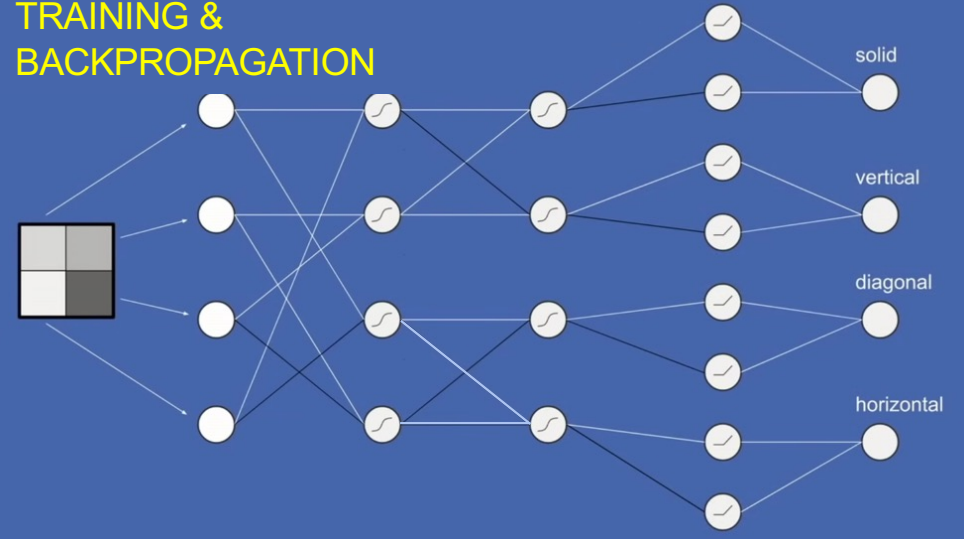
<https://machinelearningmastery.com/a-gentle-introduction-to-sigmoid-function>

Putting it altogether

INITIAL NN, FULLY CONNECTED

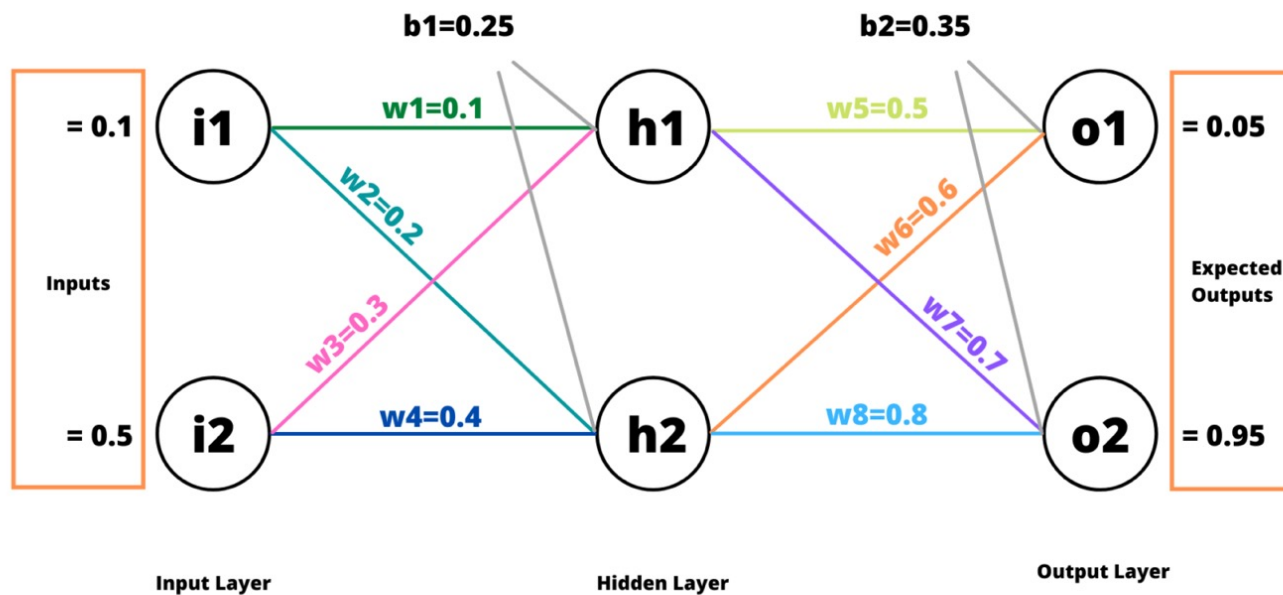


TRAINING & BACKPROPAGATION




Step-by-step example

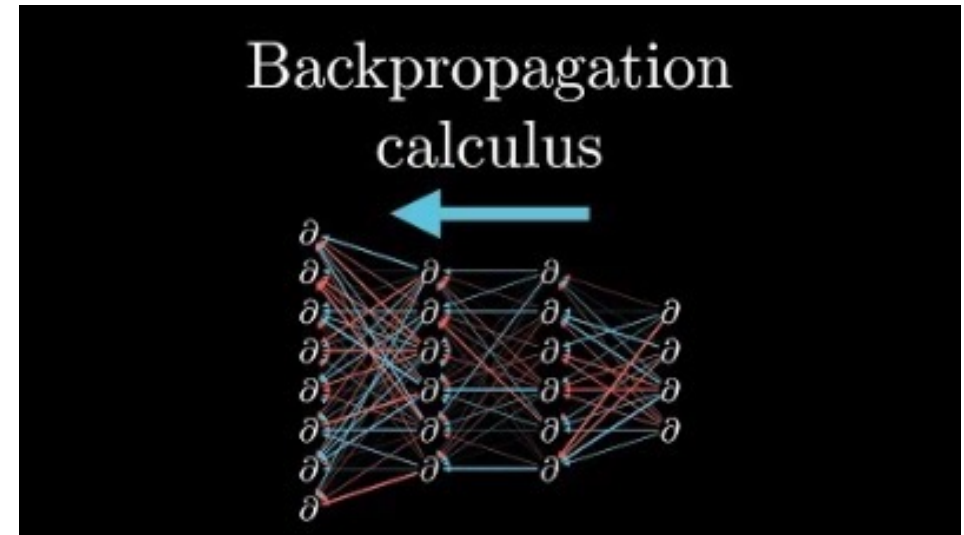
- <https://theneuralblog.com/forward-pass-backpropagation-example/>



Watch additional (great) videos



<https://www.youtube.com/watch?v=llg3gGewQ5U>
(3Blue1Brown) What is backpropagation really doing?



<https://www.youtube.com/watch?v=tleHLnjs5U8>
(3Blue1Brown) Backpropagation calculus


More details on backpropagation

For those who are curious!

- ⦿ How Does Back-Propagation Work in Neural Networks?
 - ⦿ <https://towardsdatascience.com/how-does-back-propagation-work-in-neural-networks-with-worked-example-bc59dfb97f48>
- ⦿ How to Code a Neural Network with Backpropagation In Python (from scratch) – without any library, just to understand what's under the hood!
 - ⦿ <https://machinelearningmastery.com/implement-backpropagation-algorithm-scratch-python/>
- ⦿ How Back-Propagation Works – A Python Implementation
 - ⦿ <https://towardsdatascience.com/how-back-propagation-works-a-python-implementation-21004d3b47c6>

Back to Loss/Error functions

- Mean Squared Error (MSE) and Mean Absolute Error (MAE) are appropriate loss functions for regression tasks



$$J(a, b) = \frac{1}{2m} \sum_{i=1}^m (f(x^{(i)}) - y^{(i)})^2 \quad MSE = \frac{1}{n} \sum \underbrace{(y - \hat{y})^2}_{\substack{\text{The square of the difference} \\ \text{between actual and} \\ \text{predicted}}}$$

- For classification problems, log loss function (also known as cross-entropy) is more suitable with Sigmoid activation function

$$-(y \log(p) + (1 - y) \log(1 - p))$$

Equation of Log loss function. y -Actual output, p -probability predicted by the logistic regression

Why Log Loss instead of MSE??






Why not
Mean Squared Error
as a loss function for
Logistic
regression??



- ① <https://towardsdatascience.com/why-not-mse-as-a-loss-function-for-logistic-regression-589816b5e03c>
- ② First, the target value is either 0/1 in classification problems, so $(y - \hat{y})$ will always be between 0–1, making it very difficult to track the progress of error value
- ③ Second, and the main reason actually, MSE doesn't work well with logistic regression because when the MSE loss function is plotted with respect to weights of the logistic regression model, the **curve obtained is not a convex curve** which makes it very difficult to find the global minimum


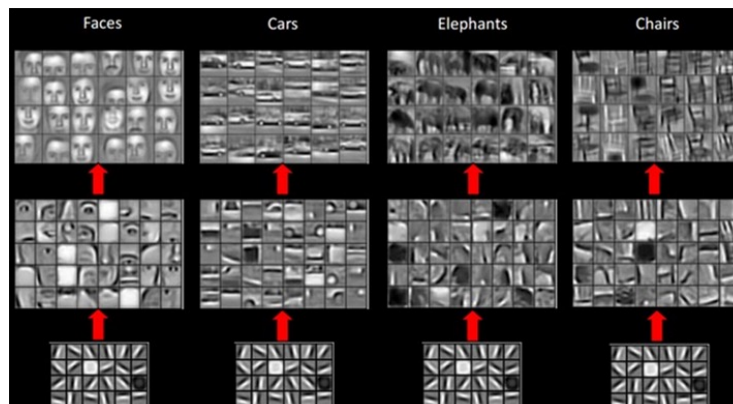

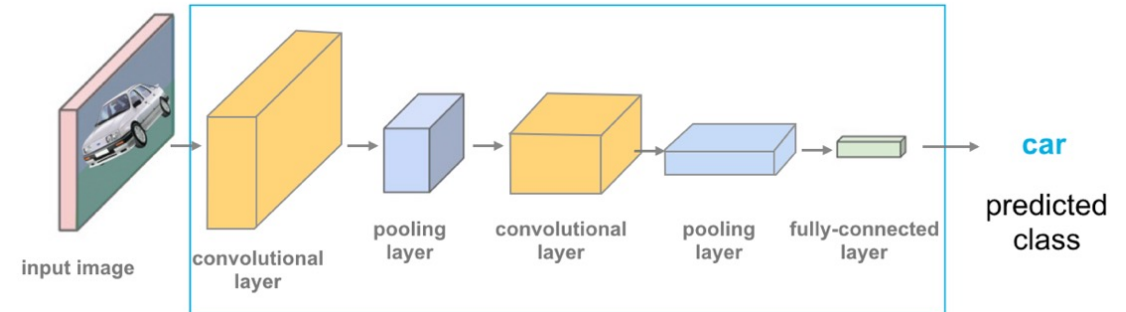
Types of Deep Neural Networks

- Feedforward Neural Networks (**FFNs, ANNs** or **NNs**)
- Recurrent Neural Networks (**RNNs**)
- Convolutional Neural Networks (**CNNs**)
- Autoencoder Neural Networks (**AEs**)





Convolutional Neural Networks

- Contain five types of layers
- Each layer has a specific purpose, like summarizing, connecting or activating
- CNN are good at image classification and object detection

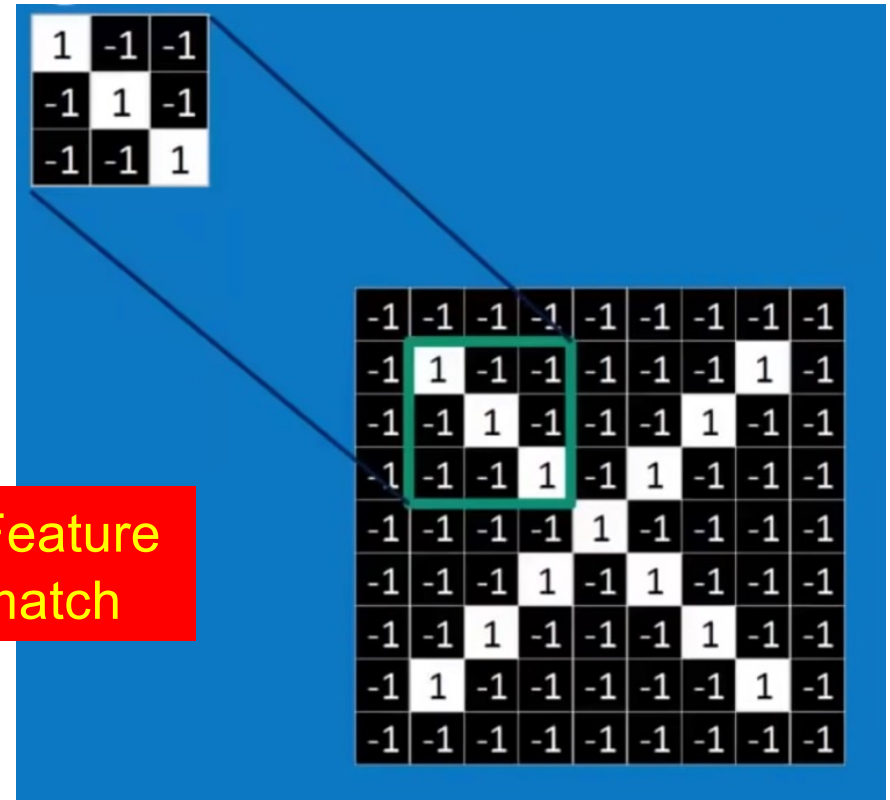



What is a convolution?



-1	-1	-1	-1	-1	-1	-1	-1	-1	-1
-1	1	-1	-1	-1	-1	-1	1	-1	-1
-1	-1	1	-1	-1	-1	1	-1	-1	-1
-1	-1	-1	1	-1	1	-1	-1	-1	-1
-1	-1	-1	-1	1	-1	-1	-1	-1	-1
-1	-1	-1	1	-1	1	-1	-1	-1	-1
-1	-1	1	-1	-1	-1	1	-1	-1	-1
-1	1	-1	-1	-1	-1	-1	1	-1	-1
-1	-1	-1	-1	-1	-1	-1	-1	-1	-1

-1	-1	-1	-1	-1	-1	-1	-1	-1	-1
-1	X	-1	-1	-1	-1	X	X	-1	-1
-1	X	X	-1	-1	X	X	-1	-1	-1
-1	-1	X	1	-1	1	-1	-1	-1	-1
-1	-1	-1	-1	1	-1	-1	-1	-1	-1
-1	-1	-1	1	-1	1	X	-1	-1	-1
-1	-1	X	X	-1	-1	X	X	-1	-1
-1	X	X	-1	-1	-1	-1	X	-1	-1
-1	-1	-1	-1	-1	-1	-1	-1	-1	-1

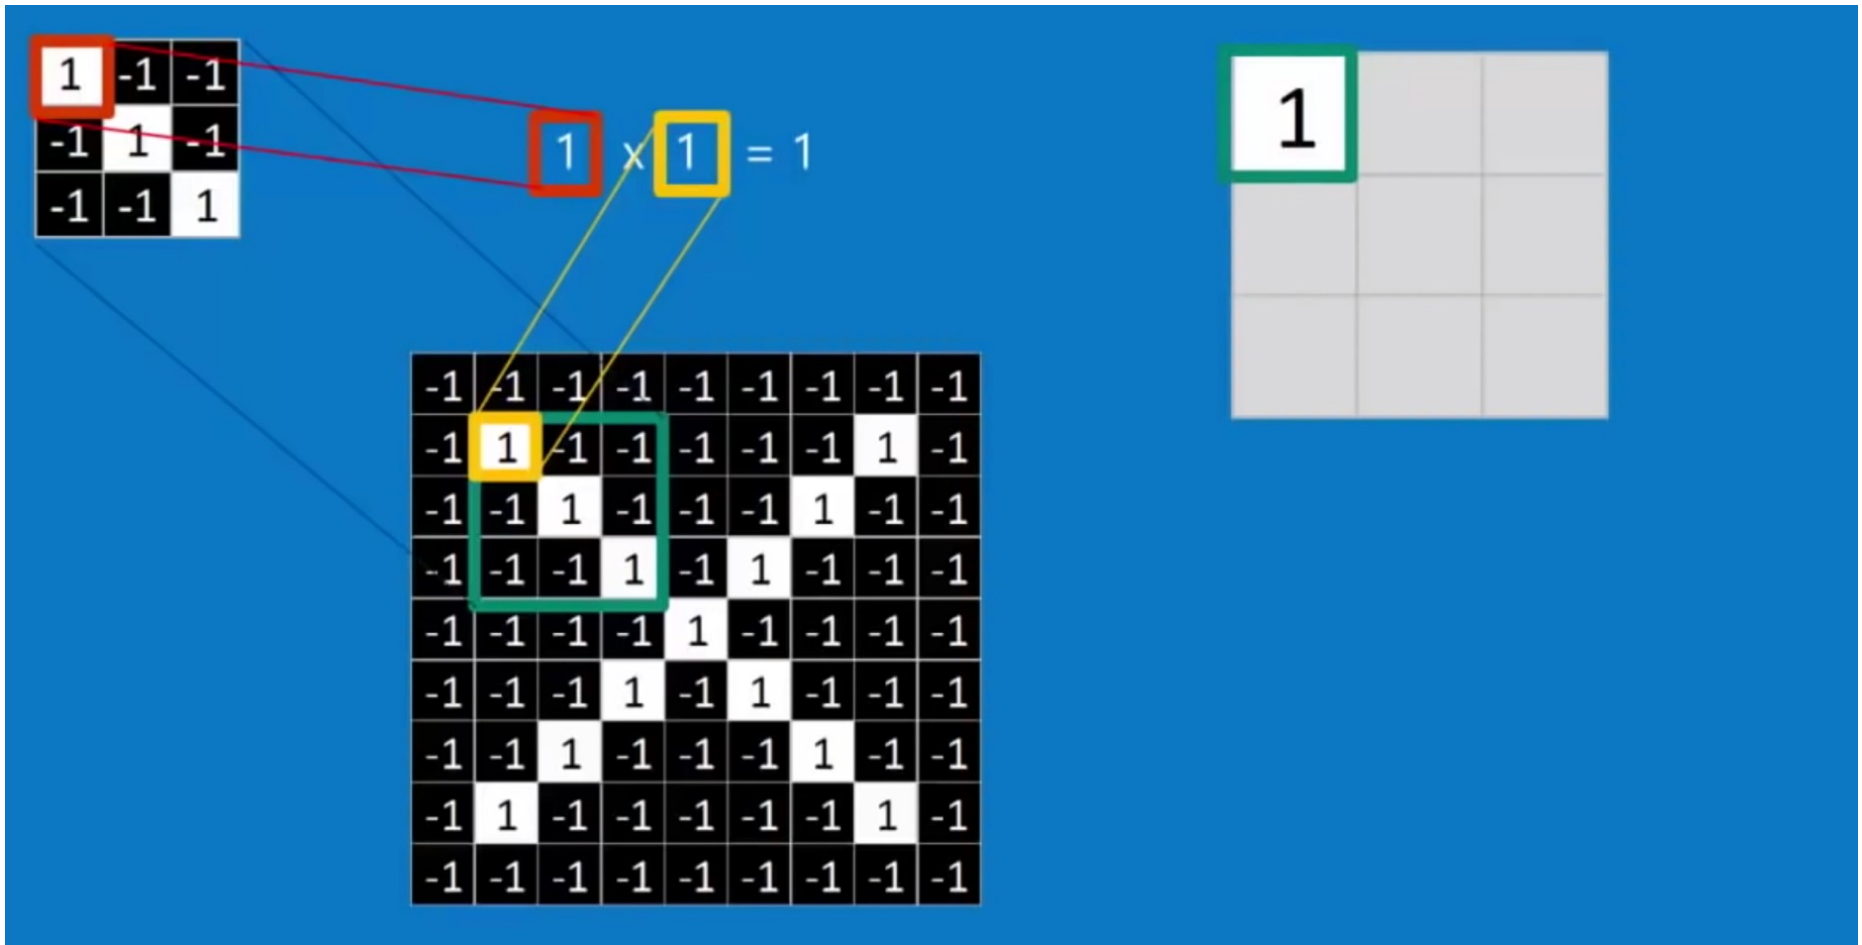


1	-1	-1
-1	1	-1
-1	-1	1


1	-1	1
-1	1	-1
1	-1	1

-1	-1	1
-1	1	-1
1	-1	-1

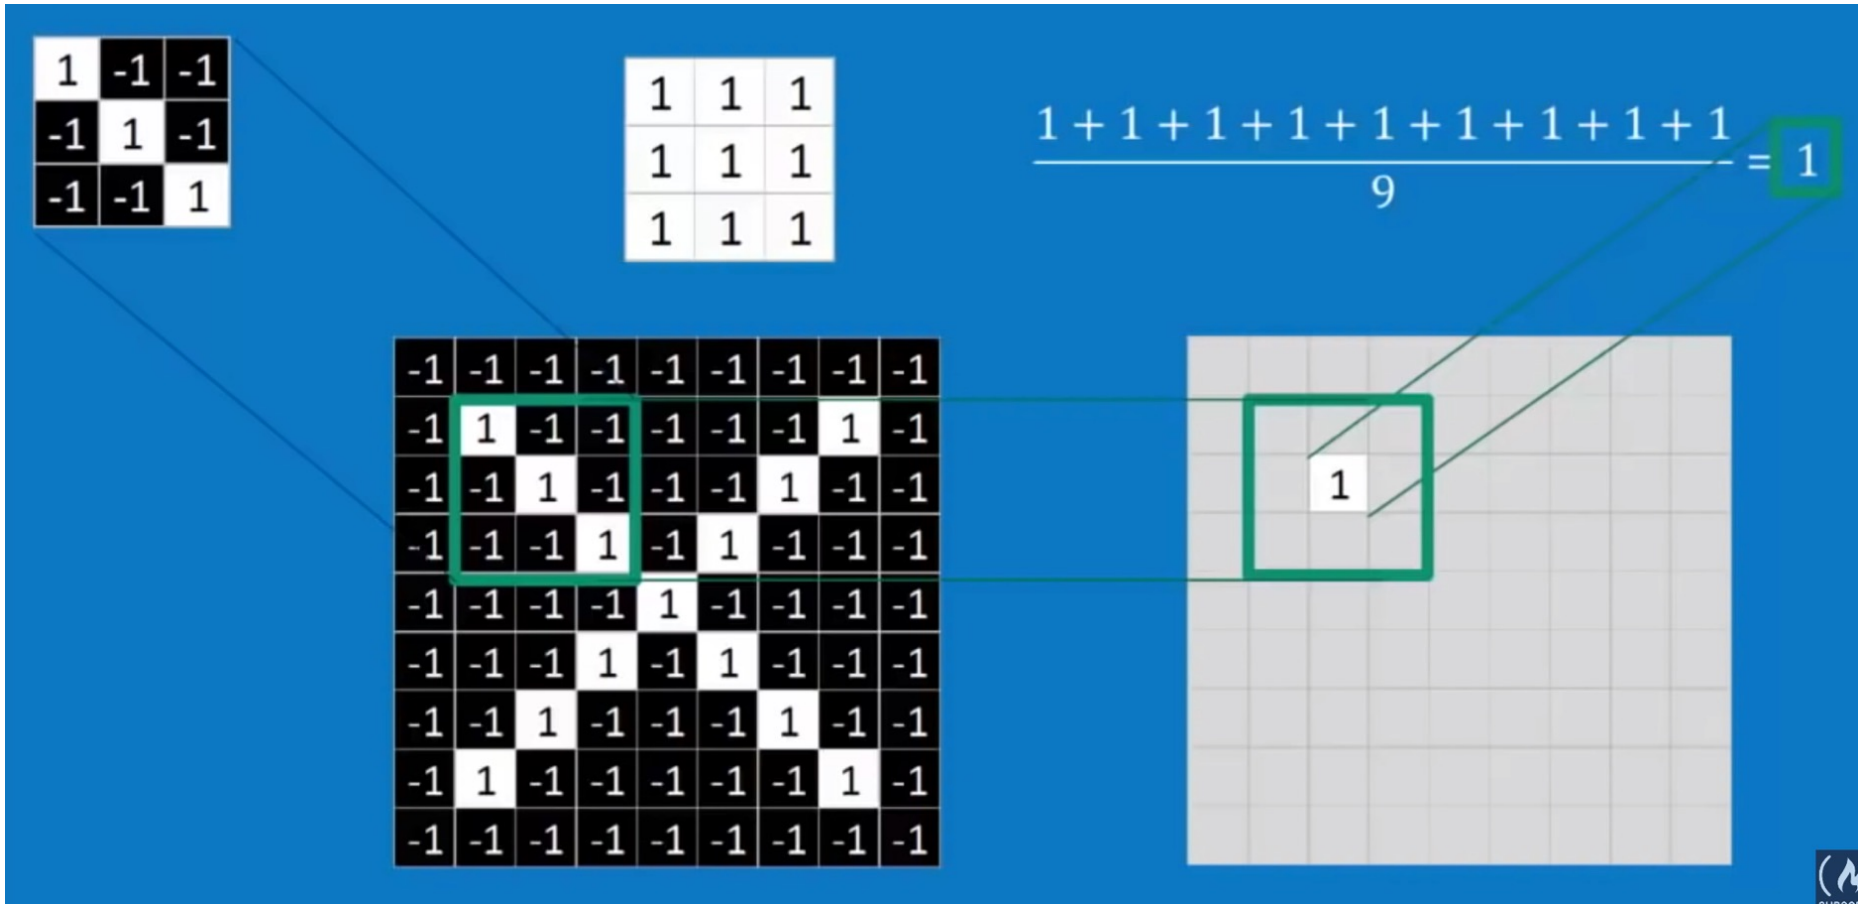
Running the convolution: filtering (1)



Running the convolution: filtering (2)



Running the convolution (3)



The diagram illustrates the final step of a 1D convolution. It shows a 9x9 input grid with a 3x3 kernel highlighted in green. The kernel values are 1, 1, 1. The output grid shows the result of the convolution, with the value 1 highlighted in green. A calculation shows the sum of the kernel values (1+1+1+1+1+1+1+1+1) divided by 9 equals 1.

Kernel values: $1, 1, 1$


Calculation:
$$\frac{1+1+1+1+1+1+1+1+1}{9} = 1$$

Input grid (9x9):


-1	-1	-1	-1	-1	-1	-1	-1	-1
-1	1	-1	-1	-1	-1	-1	1	-1
-1	-1	1	-1	-1	-1	1	-1	-1
-1	-1	-1	1	-1	1	-1	-1	-1
-1	-1	-1	-1	1	-1	-1	-1	-1
-1	-1	-1	1	-1	1	-1	-1	-1
-1	-1	1	-1	-1	-1	1	-1	-1
-1	1	-1	-1	-1	-1	-1	1	-1
-1	-1	-1	-1	-1	-1	-1	-1	-1

Output grid (9x9):


Running the convolution: filtering (4)



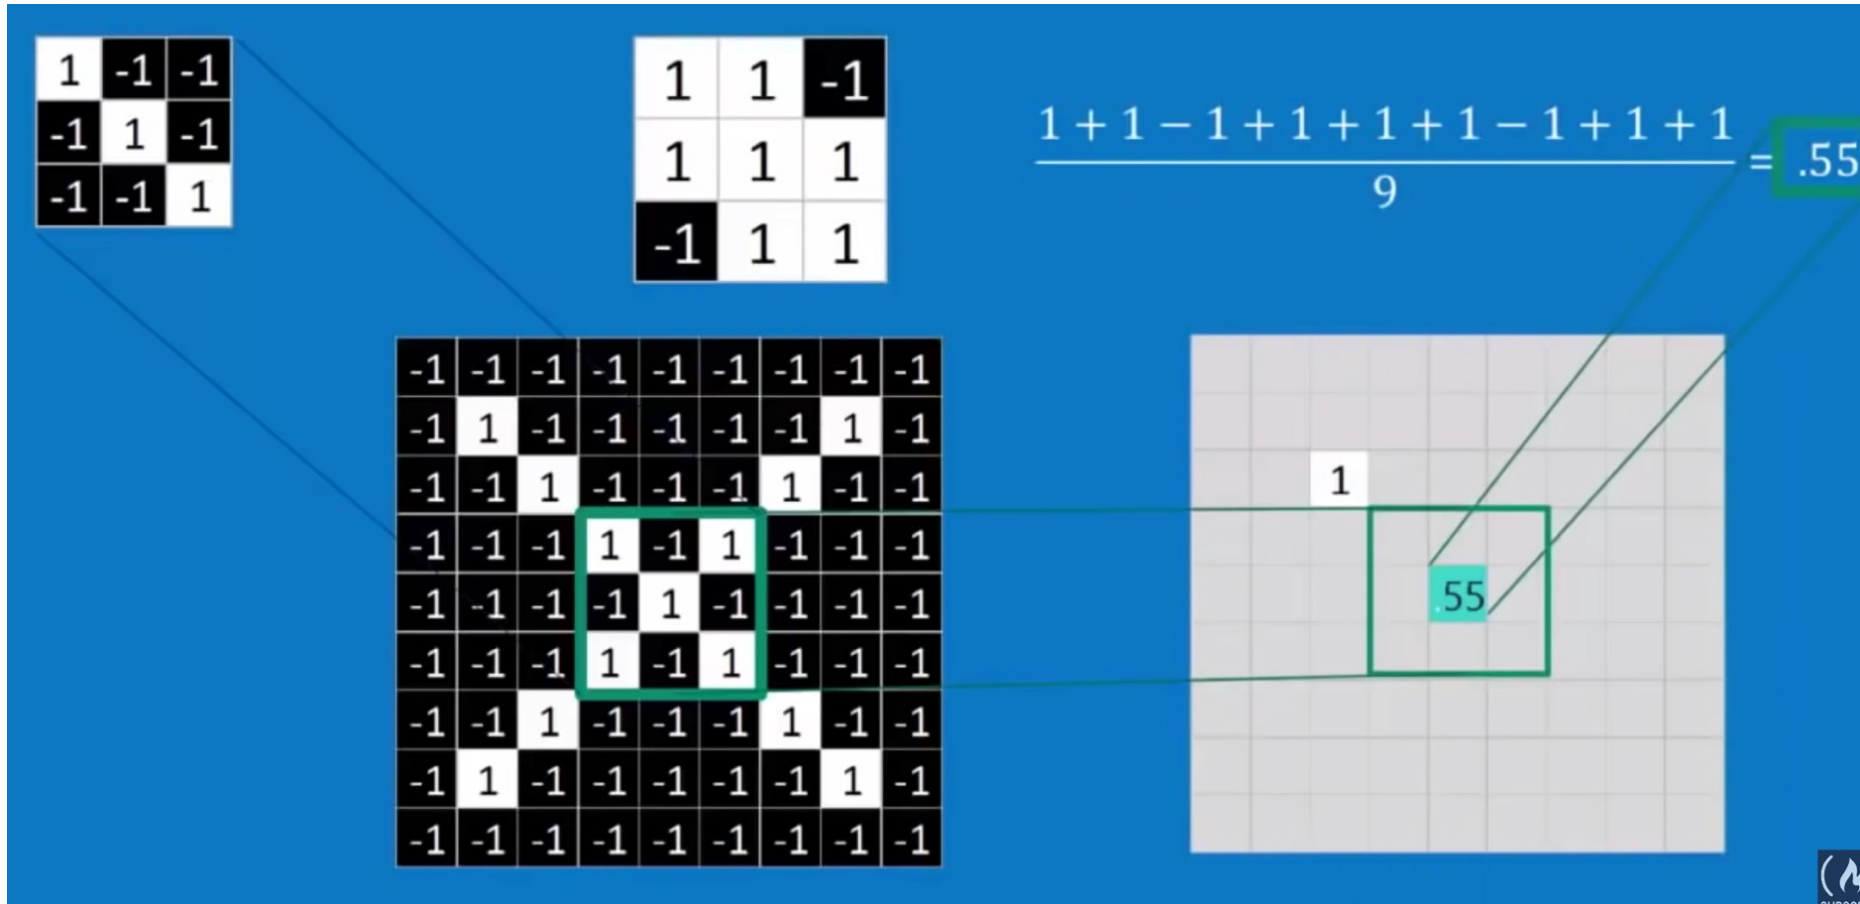
Running the convolution: filtering (5)



Running the convolution: filtering (6)



Running the convolution: filtering (7)



The diagram illustrates a 3x3 convolution kernel applied to a 9x9 input grid. The kernel is a checkerboard pattern:

1	-1	-1
-1	1	-1
-1	-1	1

The input grid is a 9x9 checkerboard pattern:

-1	-1	-1	-1	-1	-1	-1	-1	-1
-1	1	-1	-1	-1	-1	-1	1	-1
-1	-1	1	-1	-1	-1	1	-1	-1
-1	-1	-1	1	-1	1	-1	-1	-1
-1	-1	-1	-1	1	-1	-1	-1	-1
-1	-1	-1	1	-1	1	-1	-1	-1
-1	-1	1	-1	-1	-1	1	-1	-1
-1	1	-1	-1	-1	-1	-1	1	-1
-1	-1	-1	-1	-1	-1	-1	-1	-1

The kernel is applied to a 3x3 region of the input grid (highlighted with a green border), resulting in a value of 0.55:

1	1	-1
1	1	1
-1	1	1

The calculation is shown as:

$$\frac{1 + 1 - 1 + 1 + 1 + 1 - 1 + 1 + 1}{9} = .55$$

A larger grid shows the kernel being applied to a 1x1 region of the input grid, resulting in a value of 1:

1
.55

Running the convolution: filtering (8)

-1	-1	-1	-1	-1	-1	-1	-1	-1
-1	1	-1	-1	-1	-1	-1	1	-1
-1	-1	1	-1	-1	-1	1	-1	-1
-1	-1	-1	1	-1	1	-1	-1	-1
-1	-1	-1	-1	1	-1	-1	-1	-1
-1	-1	-1	1	-1	1	-1	-1	-1
-1	-1	1	-1	-1	-1	1	-1	-1
-1	1	-1	-1	-1	-1	-1	1	-1
-1	-1	-1	-1	-1	-1	-1	-1	-1



1	-1	-1
-1	1	-1
-1	-1	1



0.77	-0.11	0.11	0.33	0.55	-0.11	0.33
-0.11	1.00	-0.11	0.33	-0.11	0.11	-0.11
0.11	-0.11	1.00	-0.33	0.11	-0.11	0.55
0.33	0.33	-0.33	0.55	-0.33	0.33	0.33
0.55	-0.11	0.11	-0.33	1.00	-0.11	0.11
-0.11	0.11	-0.11	0.33	-0.11	1.00	-0.11
0.33	-0.11	0.55	0.33	0.11	-0.11	0.77


1 image = stack of filtered images

Convolution layer



Next, Pooling = reduce the size

- Define window size (usually 2 or 3)
- Define a stride (usually 2)



At the end of pooling

0.77	-0.11	0.11	0.33	0.55	-0.11	0.33
-0.11	1.00	-0.11	0.33	-0.11	0.11	-0.11
0.11	-0.11	1.00	-0.33	0.11	-0.11	0.55
0.33	0.33	-0.33	0.55	-0.33	0.33	0.33
0.55	-0.11	0.11	-0.33	1.00	-0.11	0.11
-0.11	0.11	-0.11	0.33	-0.11	1.00	-0.11
0.33	-0.11	0.55	0.33	0.11	-0.11	0.77

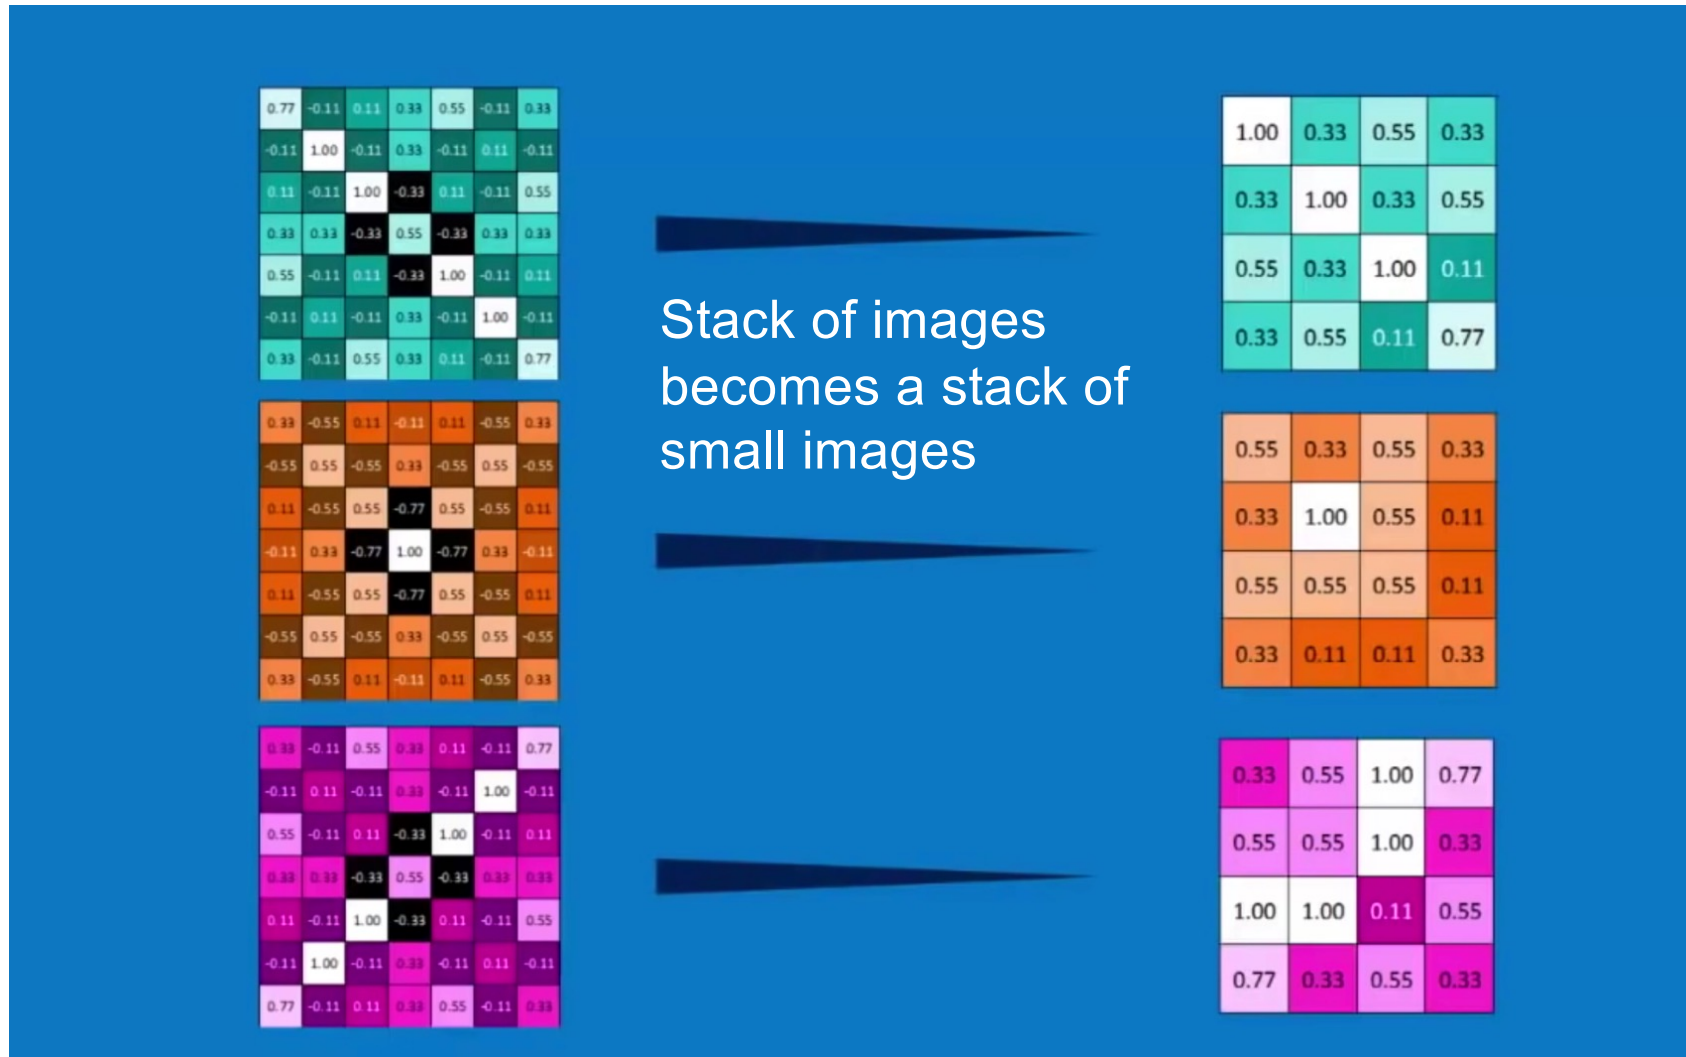
max pooling

1.00	0.33	0.55	0.33
0.33	1.00	0.33	0.55
0.55	0.33	1.00	0.11
0.33	0.55	0.11	0.77

The pattern of the original is still maintained

Do the same for stack of images


Stack of images becomes a stack of small images



The diagram illustrates the transformation of a stack of images into a stack of small images. It shows three 7x7 correlation matrices (left) and three 4x4 correlation matrices (right), each representing a different color channel (green, orange, and purple). The 7x7 matrices are arranged in a 3x7 grid, and the 4x4 matrices are arranged in a 3x4 grid. Three blue arrows point from the 7x7 matrices to the 4x4 matrices, indicating the transformation process.

Normalization: ReLU function

- Keep things from becoming unmanageably large as we are progressing through all the layers

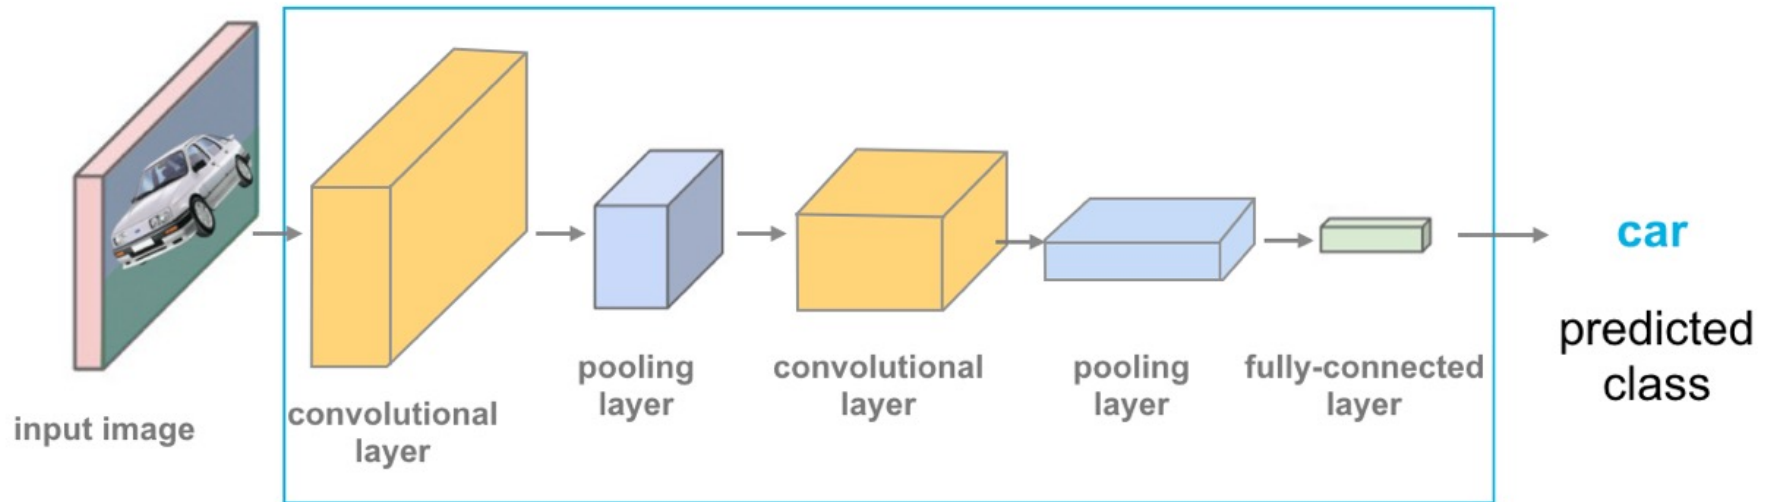


Layers can be repeated!




Watch the video on YouTube <https://www.youtube.com/watch?v=FmpDlaiMleA> (Brandon Rohrer)



Pr. Congduc Pham
<http://www.univ-pau.fr/~cpham>







Last step! Let's vote




We found the winner!

Pr. Congduc Pham
<http://www.univ-pau.fr/~cpham>



We can always add more stacks!






In addition to hidden layers, the NN can have hidden "secret" categories before concluding in is an "X"

DNN for image classification

- Top 4 Pre-Trained Models for Image Classification with Python Code
- <https://www.analyticsvidhya.com/blog/2020/08/top-4-pre-trained-models-for-image-classification-with-python-code/>

Pr. Congduc Pham
http://www.univ-pau.fr/~cpham

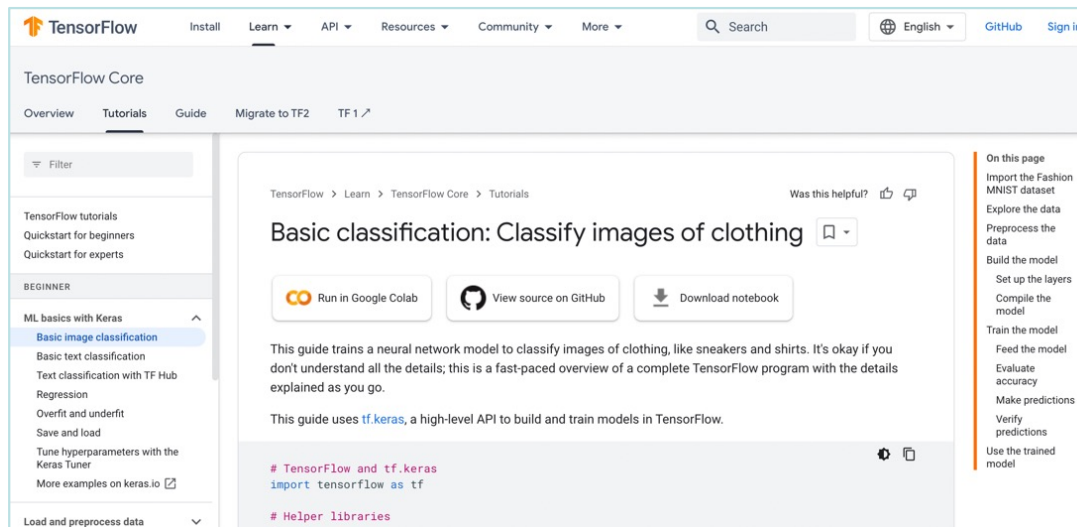




Deep Learning Examples

TensorFlow tutorial 1

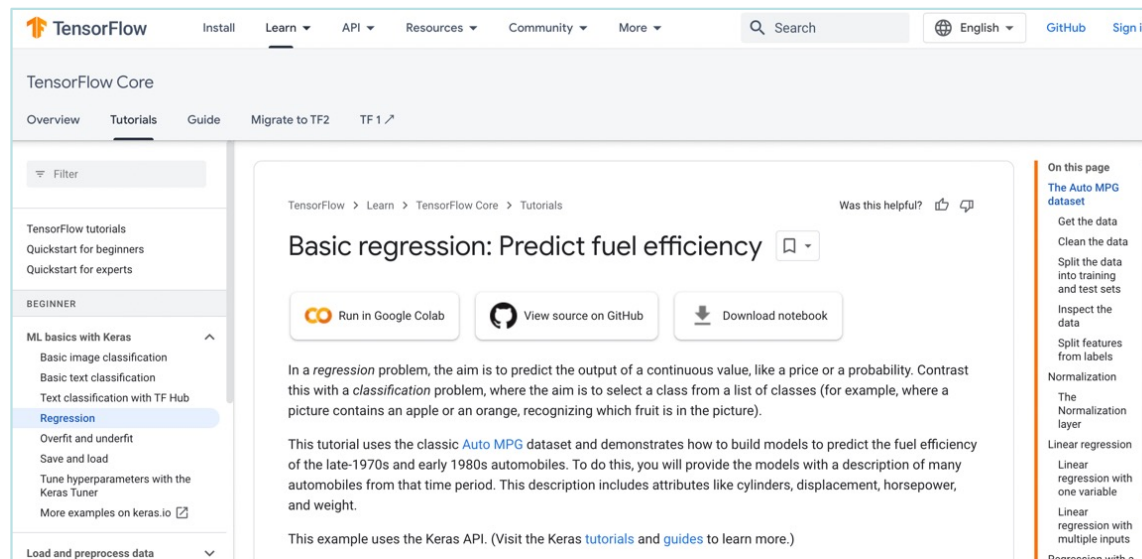
- ML basics with Keras
 - Basic image classification
 - <https://www.tensorflow.org/tutorials/keras/classification>



- You can click on "Run in Google Colab"
- Then use "Copy to Drive" to get your own copy for further testing and editing

TensorFlow tutorial 2

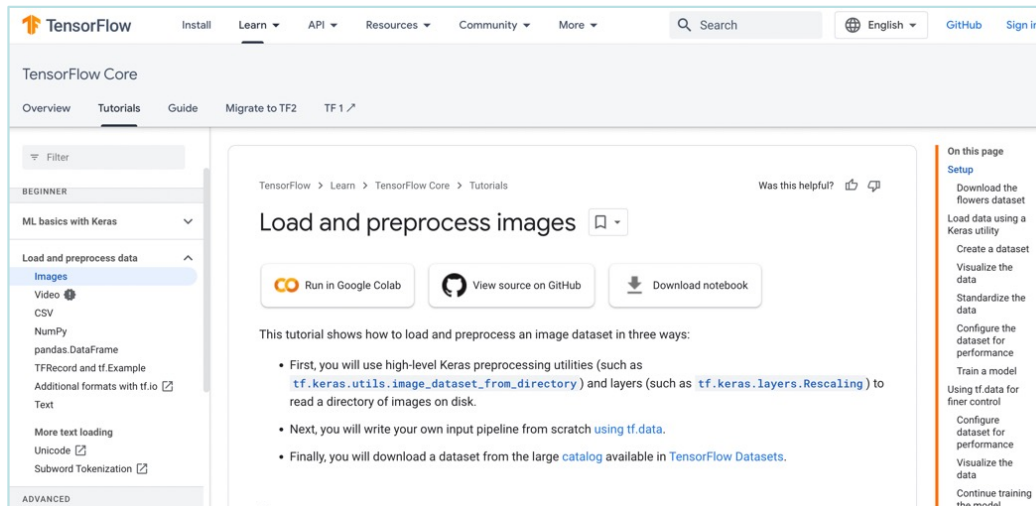
- ML basics with Keras
 - Basic regression
 - <https://www.tensorflow.org/tutorials/keras/regression>



- You can click on "Run in Google Colab"
- Then use "Copy to Drive" to get your own copy for further testing and editing

TensorFlow tutorial 3

- Load and preprocess data
 - Images
 - https://www.tensorflow.org/tutorials/load_data/images

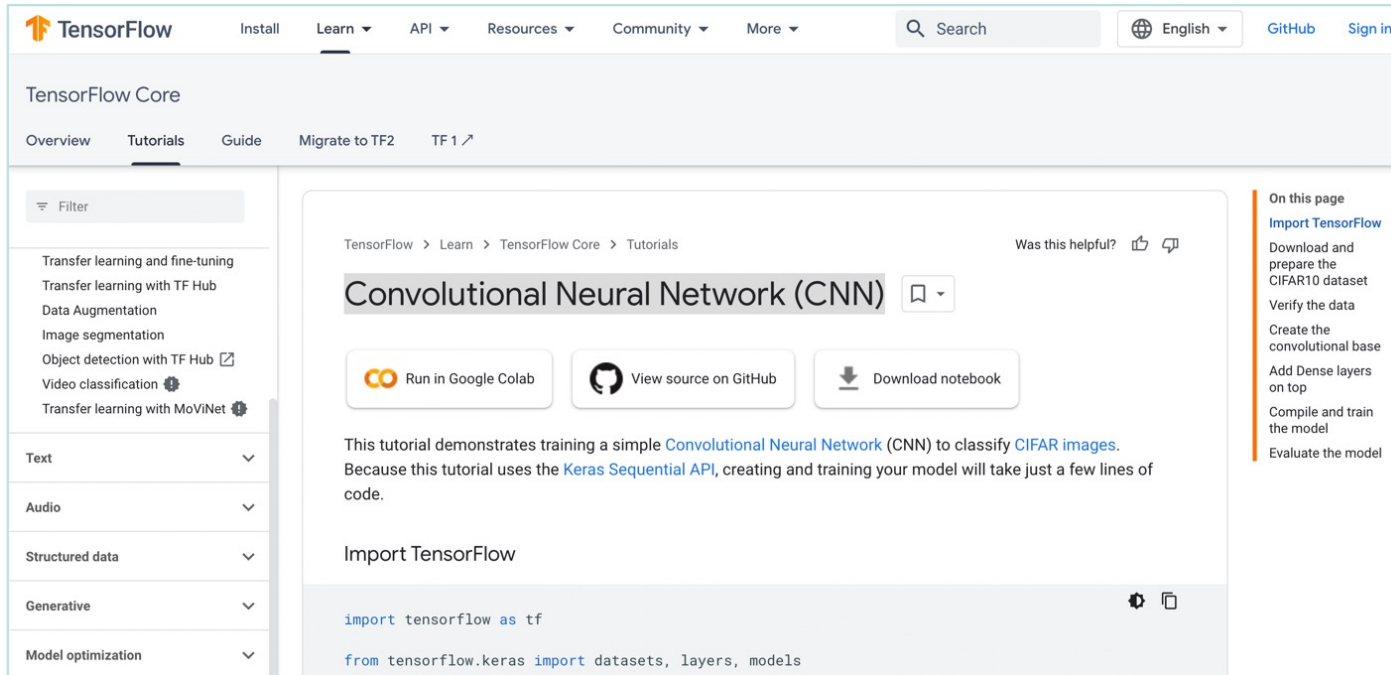


- You can click on "Run in Google Colab"
- Then use "Copy to Drive" to get your own copy for further testing and editing

CNN with TensorFlow

Assignment to do in lab

① <https://www.tensorflow.org/tutorials/images/cnn>



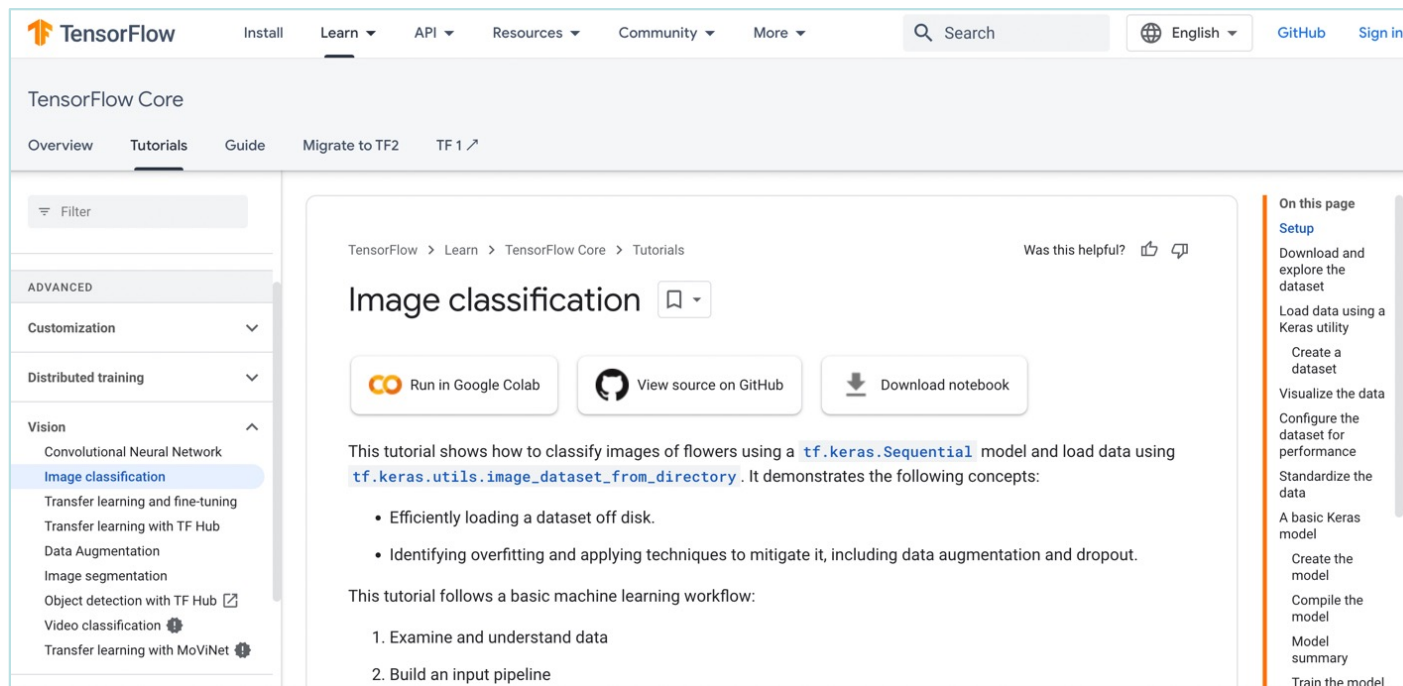
The screenshot shows the TensorFlow website's 'Convolutional Neural Network (CNN)' tutorial page. The page is titled 'Convolutional Neural Network (CNN)' and includes a navigation menu with options like 'Run in Google Colab', 'View source on GitHub', and 'Download notebook'. The main content area contains the text: 'This tutorial demonstrates training a simple Convolutional Neural Network (CNN) to classify CIFAR images. Because this tutorial uses the Keras Sequential API, creating and training your model will take just a few lines of code.' Below this text, there is a code block starting with 'import tensorflow as tf' and 'from tensorflow.keras import datasets, layers, models'. On the right side, there is a 'On this page' section with a list of steps: 'Import TensorFlow', 'Download and prepare the CIFAR10 dataset', 'Verify the data', 'Create the convolutional base', 'Add Dense layers on top', 'Compile and train the model', and 'Evaluate the model'.

- ① You can click on "Run in Google Colab"
- ① Then use "Copy to Drive" to get your own copy for further testing and editing

Image classification with TensorFlow

Assignment to do in lab

① <https://www.tensorflow.org/tutorials/images/classification>



- ① You can click on "Run in Google Colab"
- ① Then use "Copy to Drive" to get your own copy for further testing and editing




AI
FOR
CPS

AI for CPS?

- ① *"Sensors, computational and physical world are closely integrated in CPS. CPS includes both conventional embedded systems and control systems that are presumed to be remodelled by emerging methodologies and integration of the Internet of Things. **IoT is the base or enabling technology for cyber-physical systems.**"*
- ② "While deep learning in the cloud has been tremendously successful, it is not applicable in all situations. Many applications require on-device inference." [see ref]
- ③ "And the delay caused by the roundtrip to the cloud is prohibitive for applications that require real-time ML inference" [see ref]

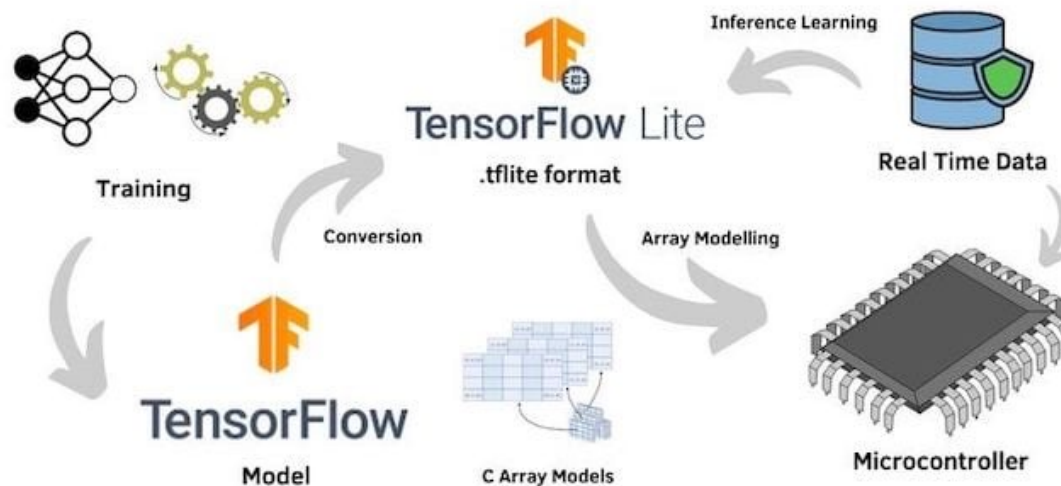
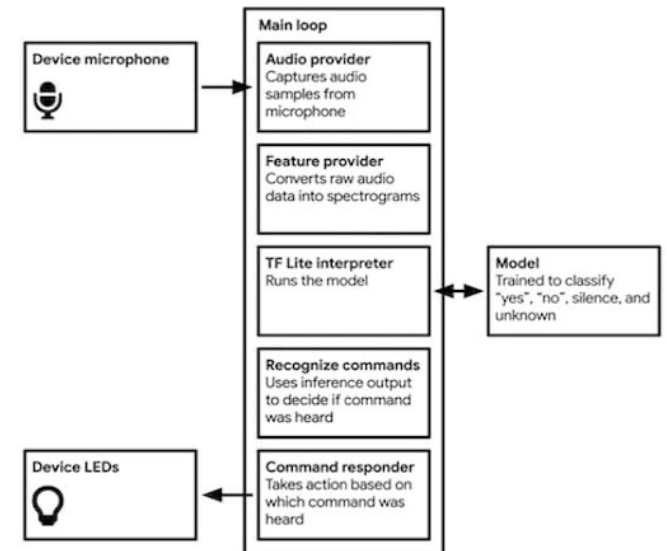
AI on microcontrollers?

- ⦿ "Your iPhone now runs facial recognition and speech recognition on device. Your Android phone can run on-device translation. Your Apple Watch uses machine learning to detect movements and ECG patterns." [see ref]
- ⦿ **BUT AI on microcontrollers** needs to go a step further to make AI techniques computationally tractable on much more resource-constrained devices
- ⦿ "**TinyML** takes edge AI one step further, making it possible to run deep learning models on microcontrollers". [see ref]
- ⦿ <https://www.tinyml.org/>



TinyML & Tensorflow Lite

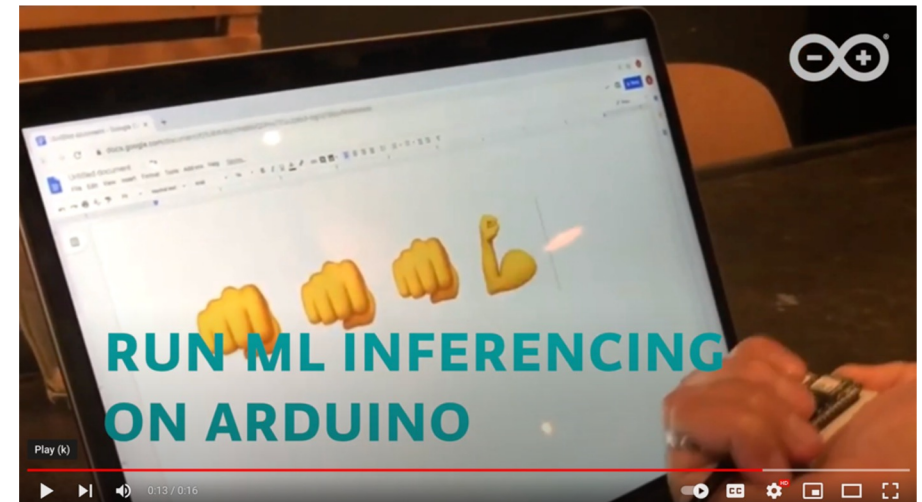

- It is already part of our daily life in most of wake-up word application
- The most popular and built-out ecosystem for TinyML development is TensorFlow Lite for Microcontrollers (TF Lite Micro)



Example: the Arduino Nano 33 BLE Sense


- ARM Cortex-M4 at 64 MHz with BLE and plenty of sensors! 3D accelerometer, microphone, gesture, light, proximity, barometric pressure, temperature, humidity. Run AI using TinyML and TensorFlow™ Lite.

<https://youtu.be/HzCRZsGJLbI>



AI on single computer boards?

- Single Board Computer can be part of the CPS ecosystem



- They can natively run complex AI models but for some real-time applications, AI accelerators can provide them with specialized AI hardware such as TPU (Tensor Processing Unit), VPU (Visual Processing Unit), ...
- For SBC, the most convenient AI accelerators are those embedded on the USB stick format

AI accelerators in USB stick format

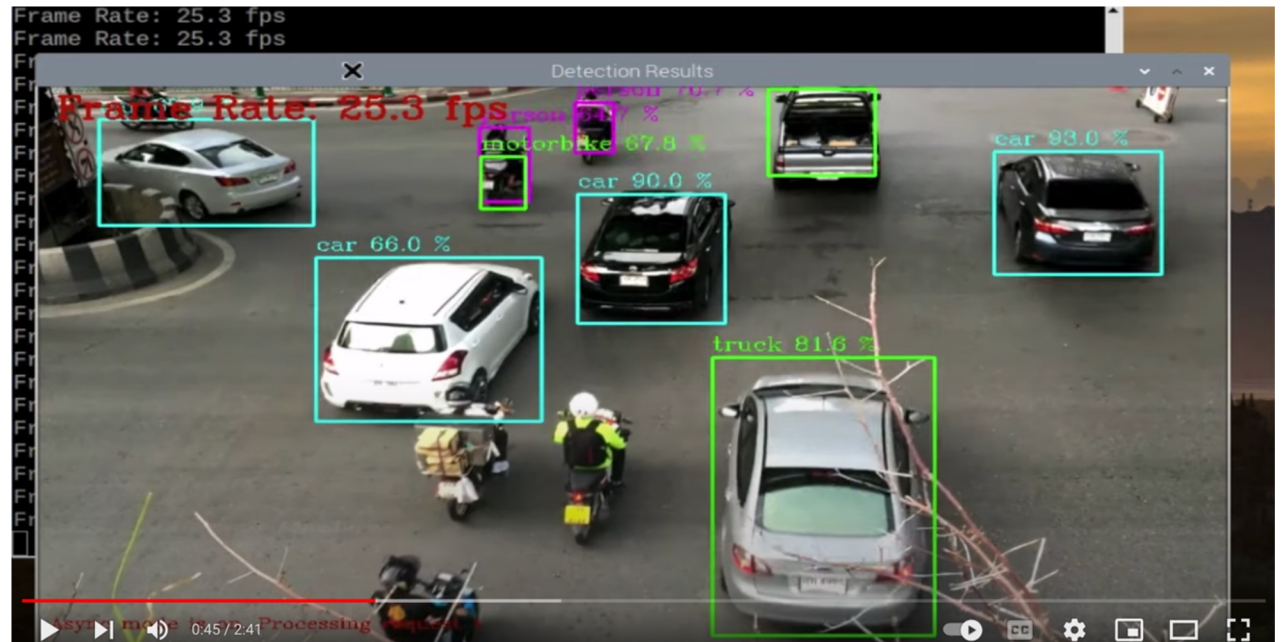


- ⦿ Depending on the product, the AI/ML models can be different, with performances depending on how you plan to use them
- ⦿ Some provide pre-trained DL models for image classification and object detection, allowing real-time applications on SBC
- ⦿ They can support most of the well-know AI frameworks/libs such as TensorFlow, ApacheMXNet, PyTorch,

Example: real-time embedded AI on RPI

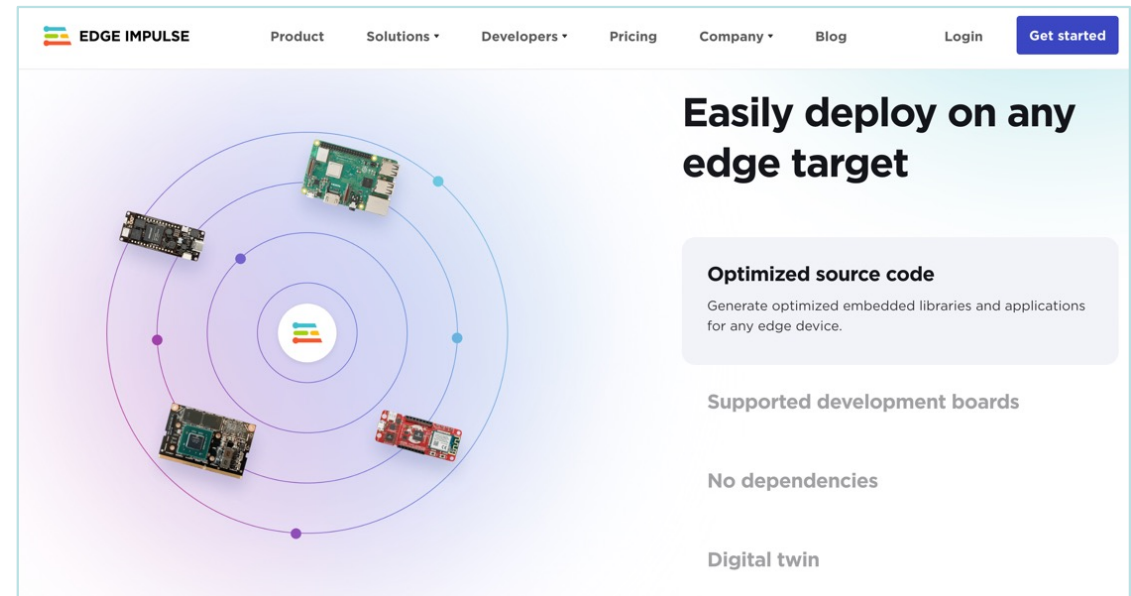
- ◉ Dive in the exciting world of real-time embedded AI with hardware accelerators (Intel Neural Compute Stick, Google Coral USB, ...)

<https://youtu.be/6XnktdajhU>



EdgeImpulse

- ① <https://www.edgeimpulse.com/>
- ① Edge Impulse is a cloud-based machine learning operations (MLOps) platform for developing embedded and edge ML (TinyML) systems that can be deployed to a wide range of hardware targets
- ① It enables the deployment of highly-optimized ML on hardware ranging from MCUs to CPUs and custom AI accelerators





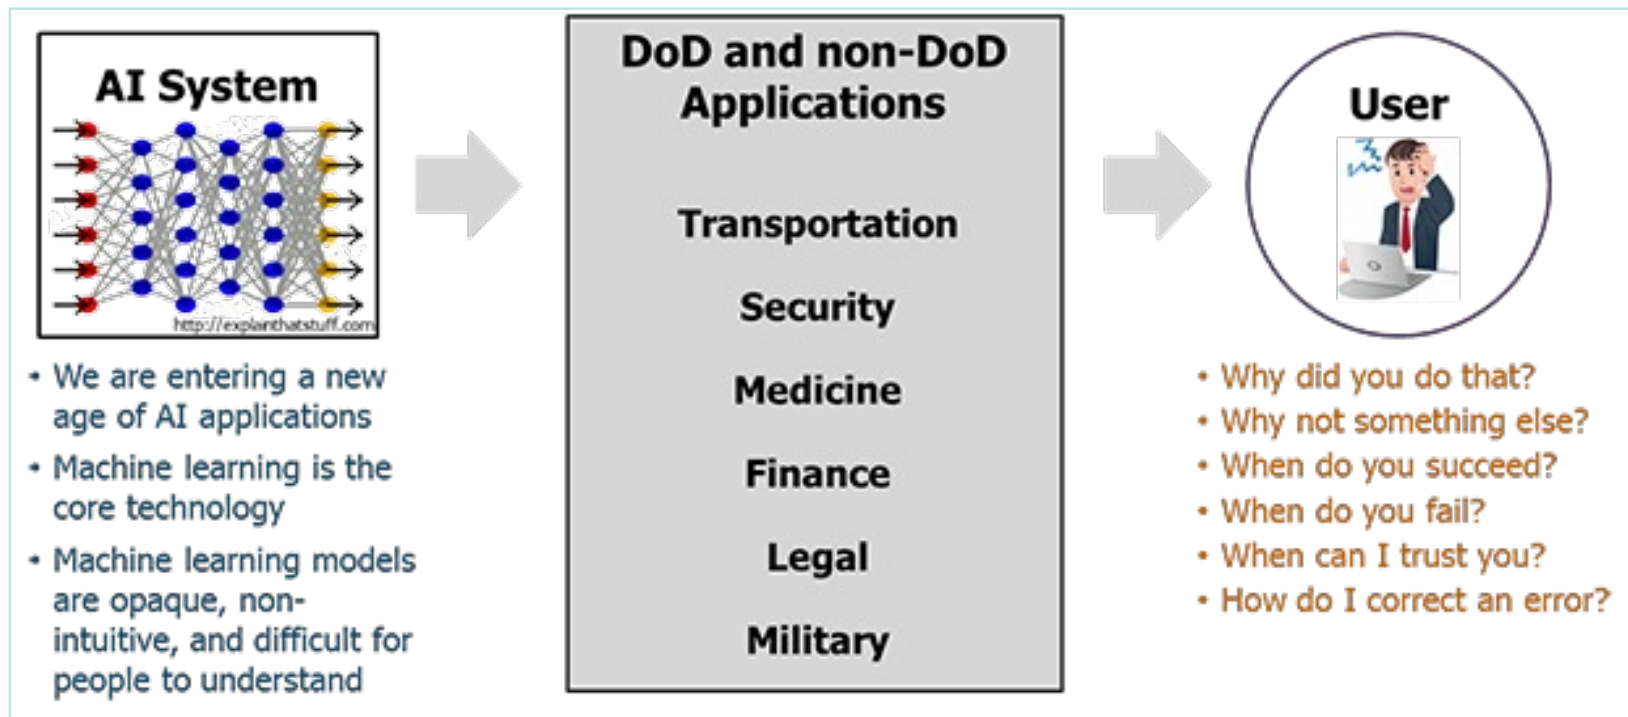
X

AI

?

eXplainable AI?

- ⦿ "the effectiveness of AI systems is limited by the machine's current inability to explain their decisions and actions to human users" [see ref]



eXplainable AI?

- "New machine-learning systems will have the ability to explain their rationale, characterize their strengths and weaknesses, and convey an understanding of how they will behave in the future" [see ref]

