SURF

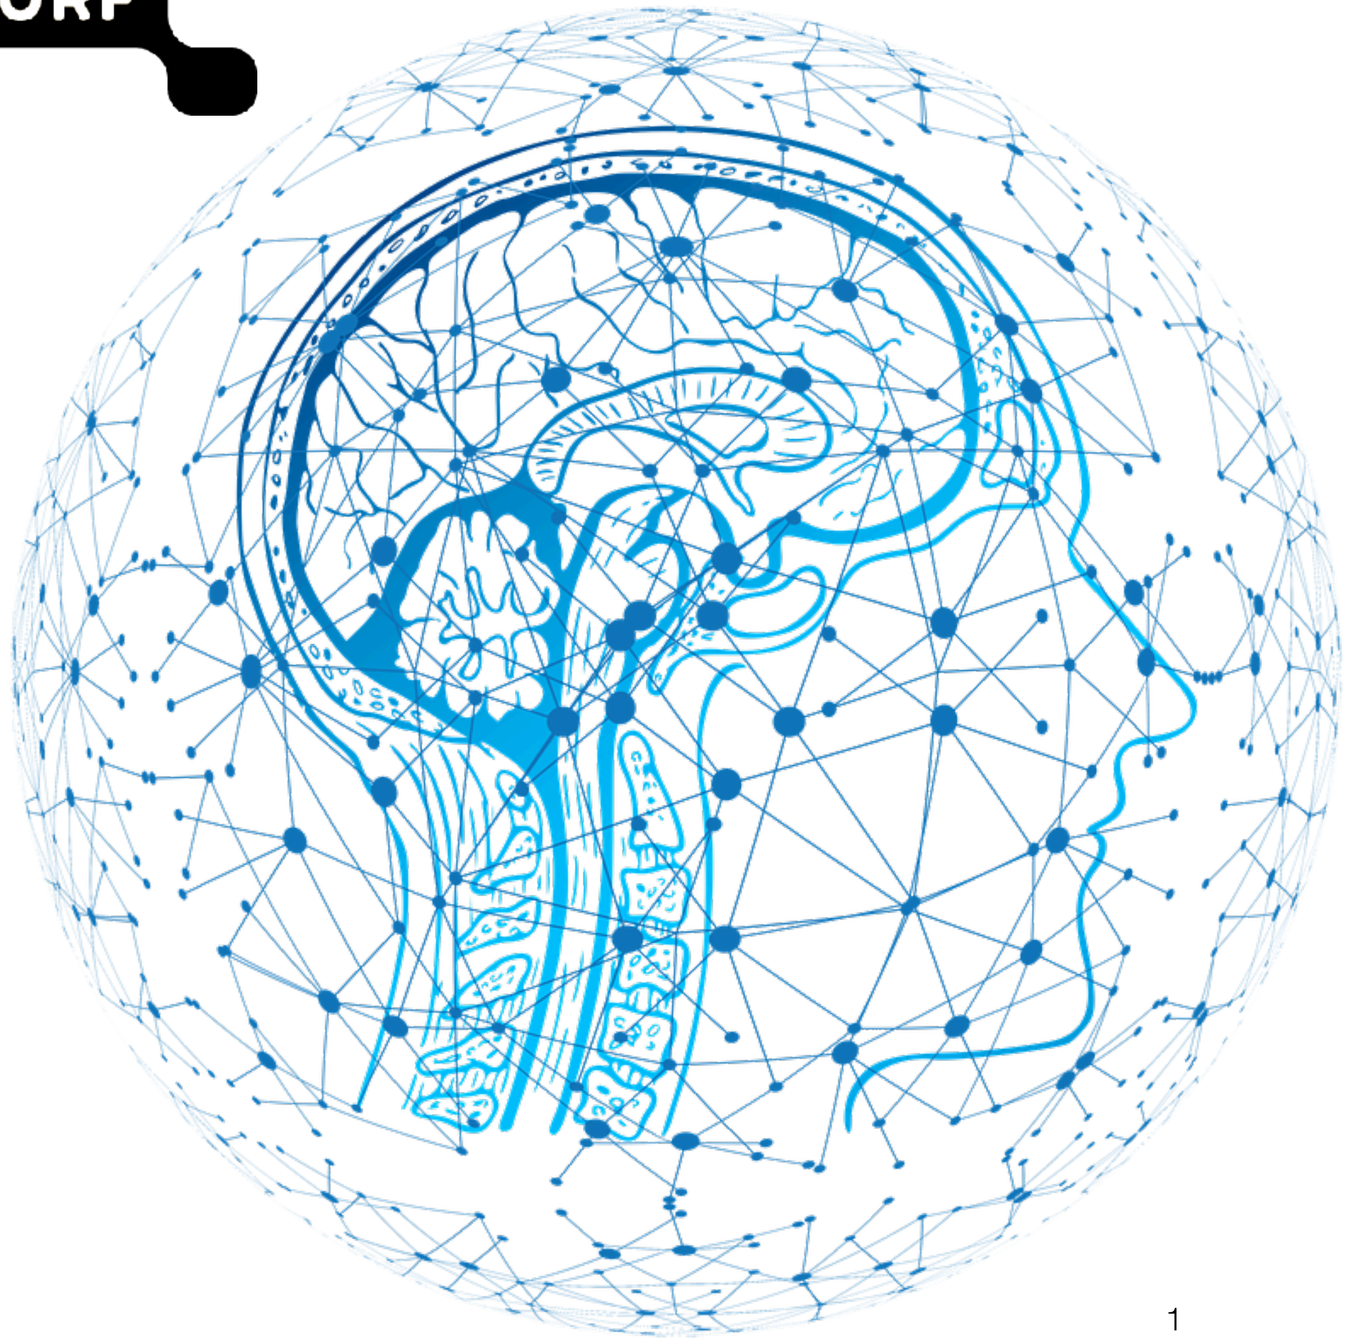

Introduction to Deep Learning

Multi-layer Perceptron

Maxwell Cai

maxwellcai.com

How does a neural network work?

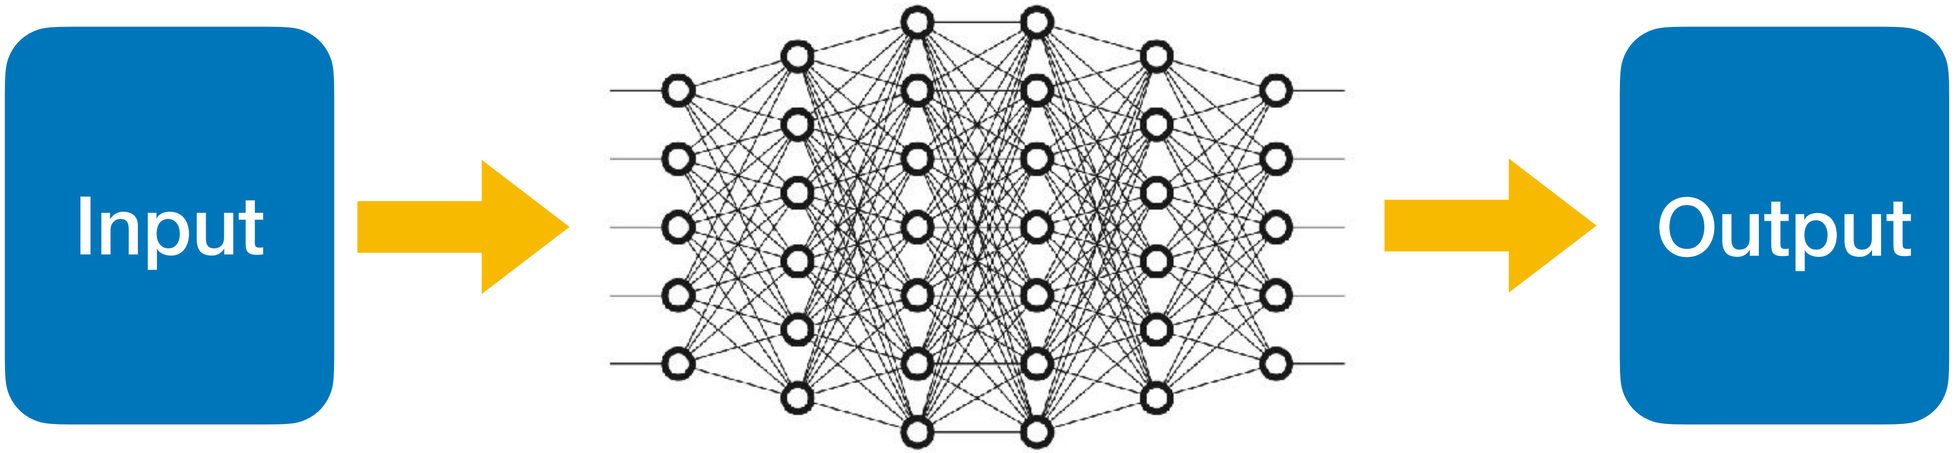

$$\hat{y} = f_{\mathrm{NN}}(x_1, x_2, \ldots, x_n)$$

Prediction

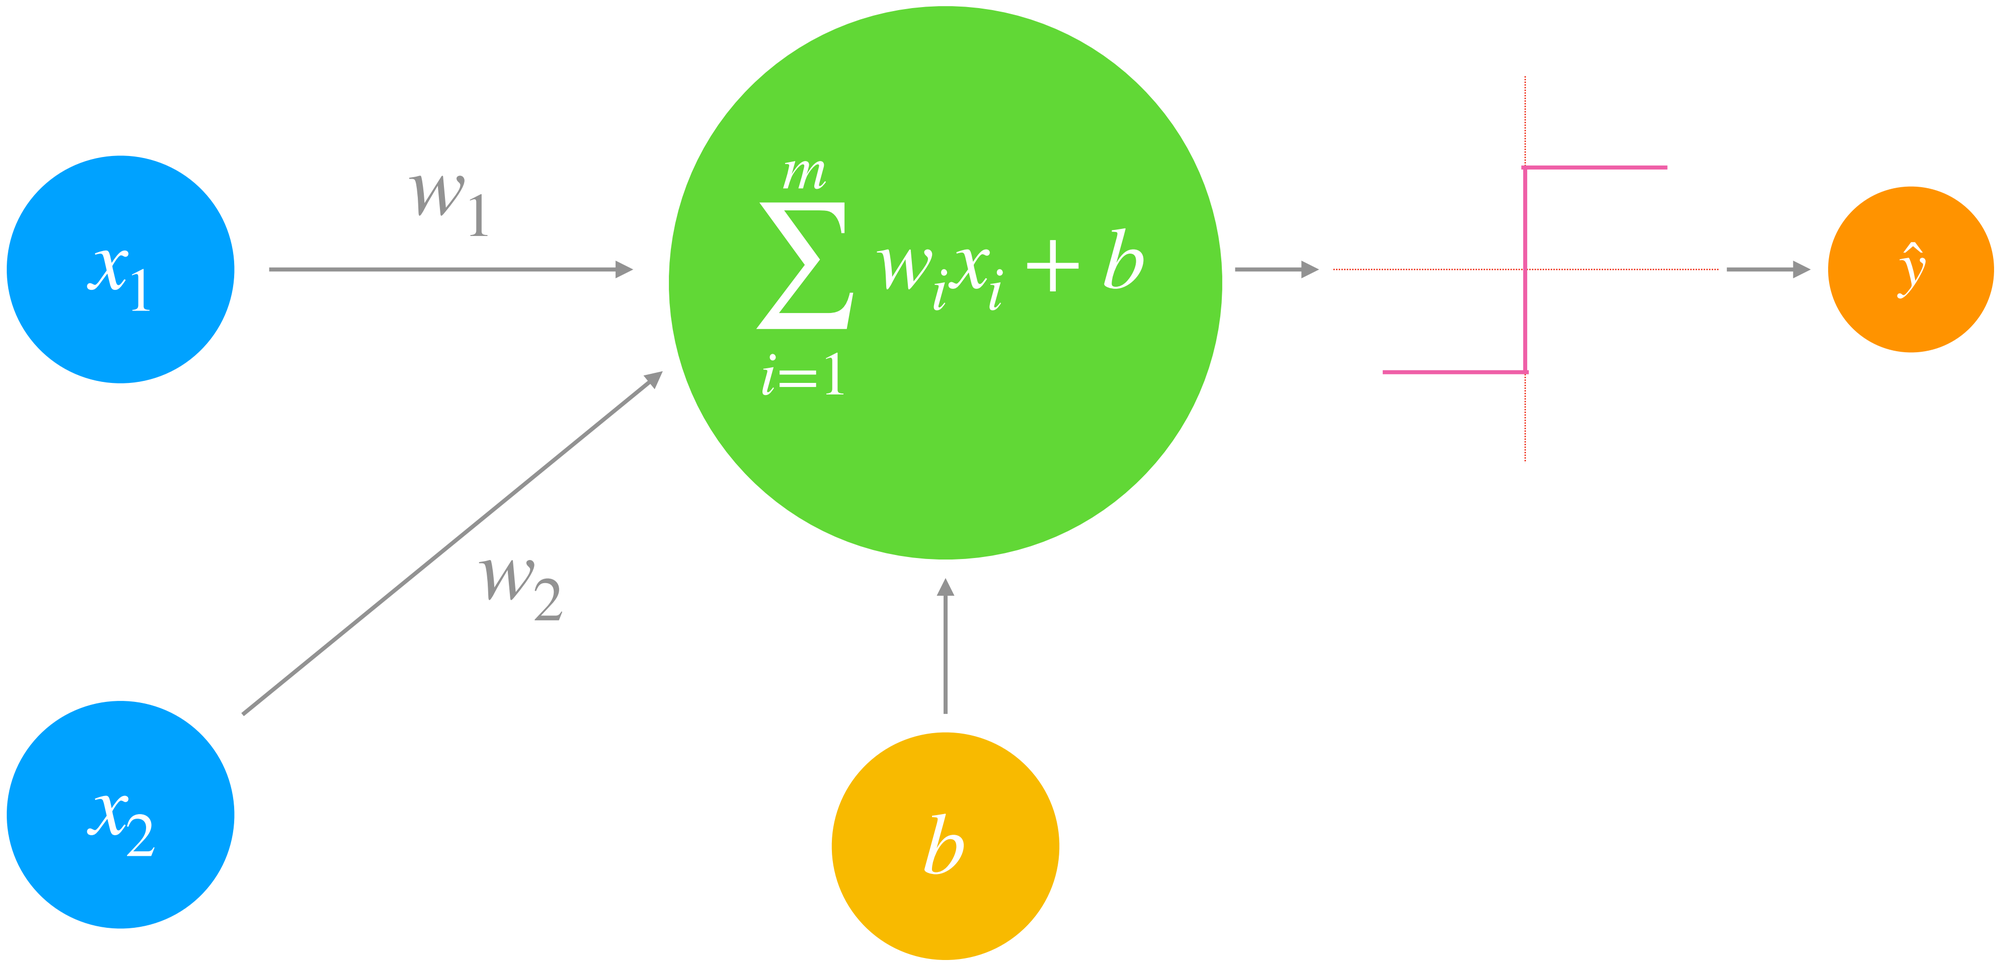
x_1
w_1
$\sum_{i=1}^{m} w_i x_i + b$
$\hat{y}$
w_2
x_2
b
inputs
weights
summation + bias
activation
output

Neurons/Perceptrons & Activation Functions

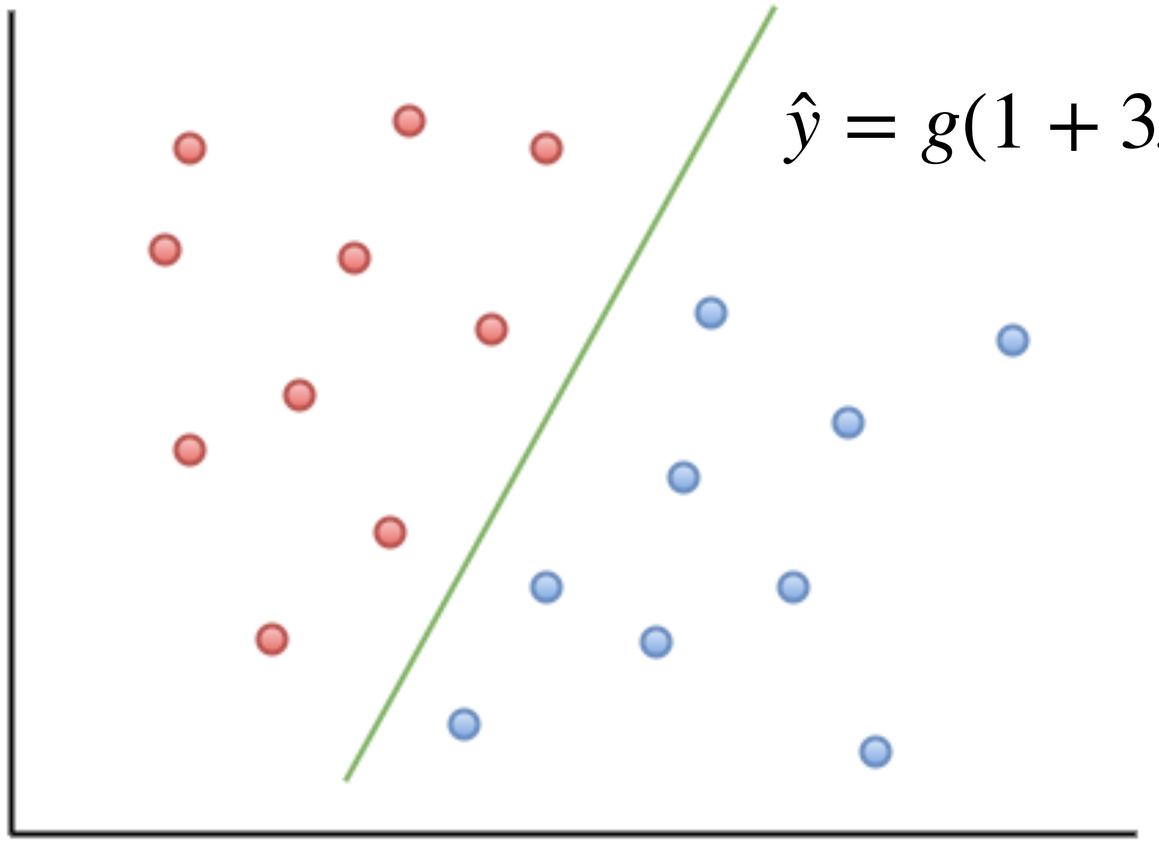

$$\hat{y} = g(1 + 3x_1 - 2x_2)$$

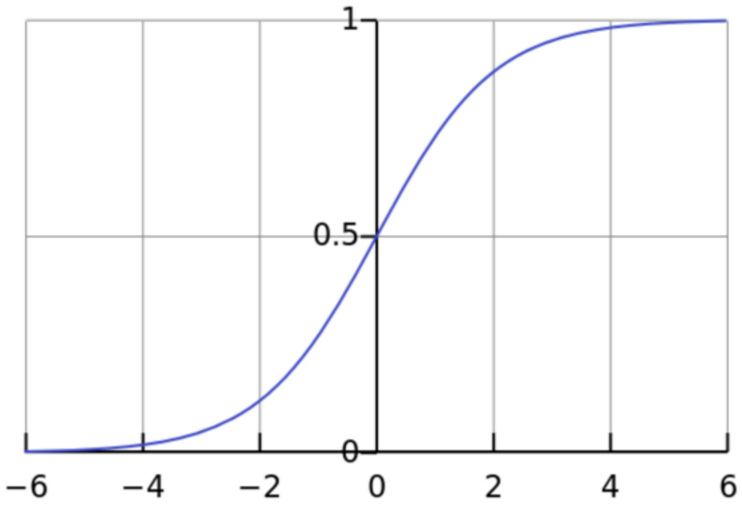

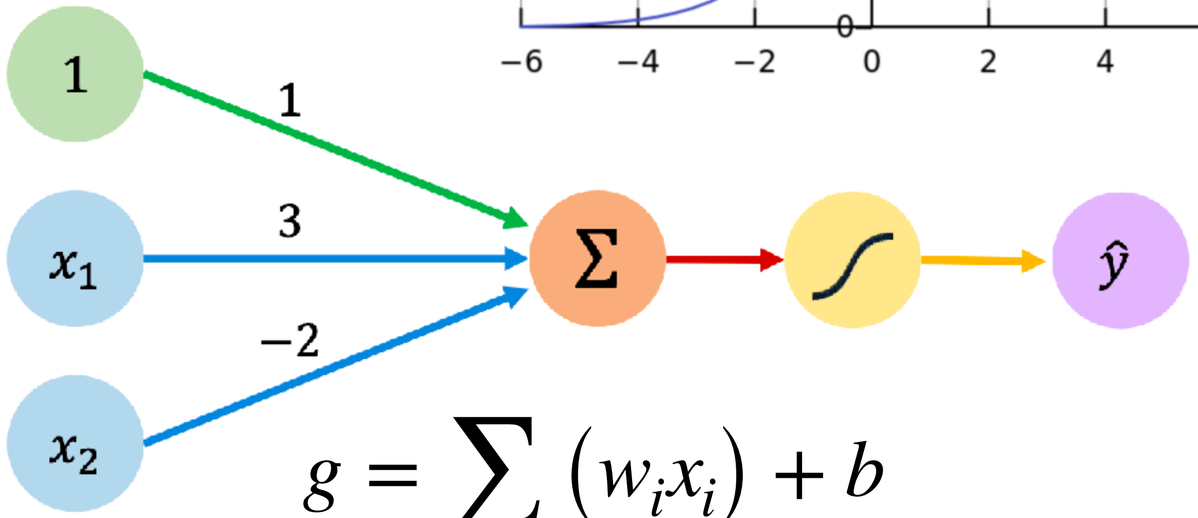

$$g = \sum (w_i x_i) + b$$

Live demo: Multi-layer Perceptrons

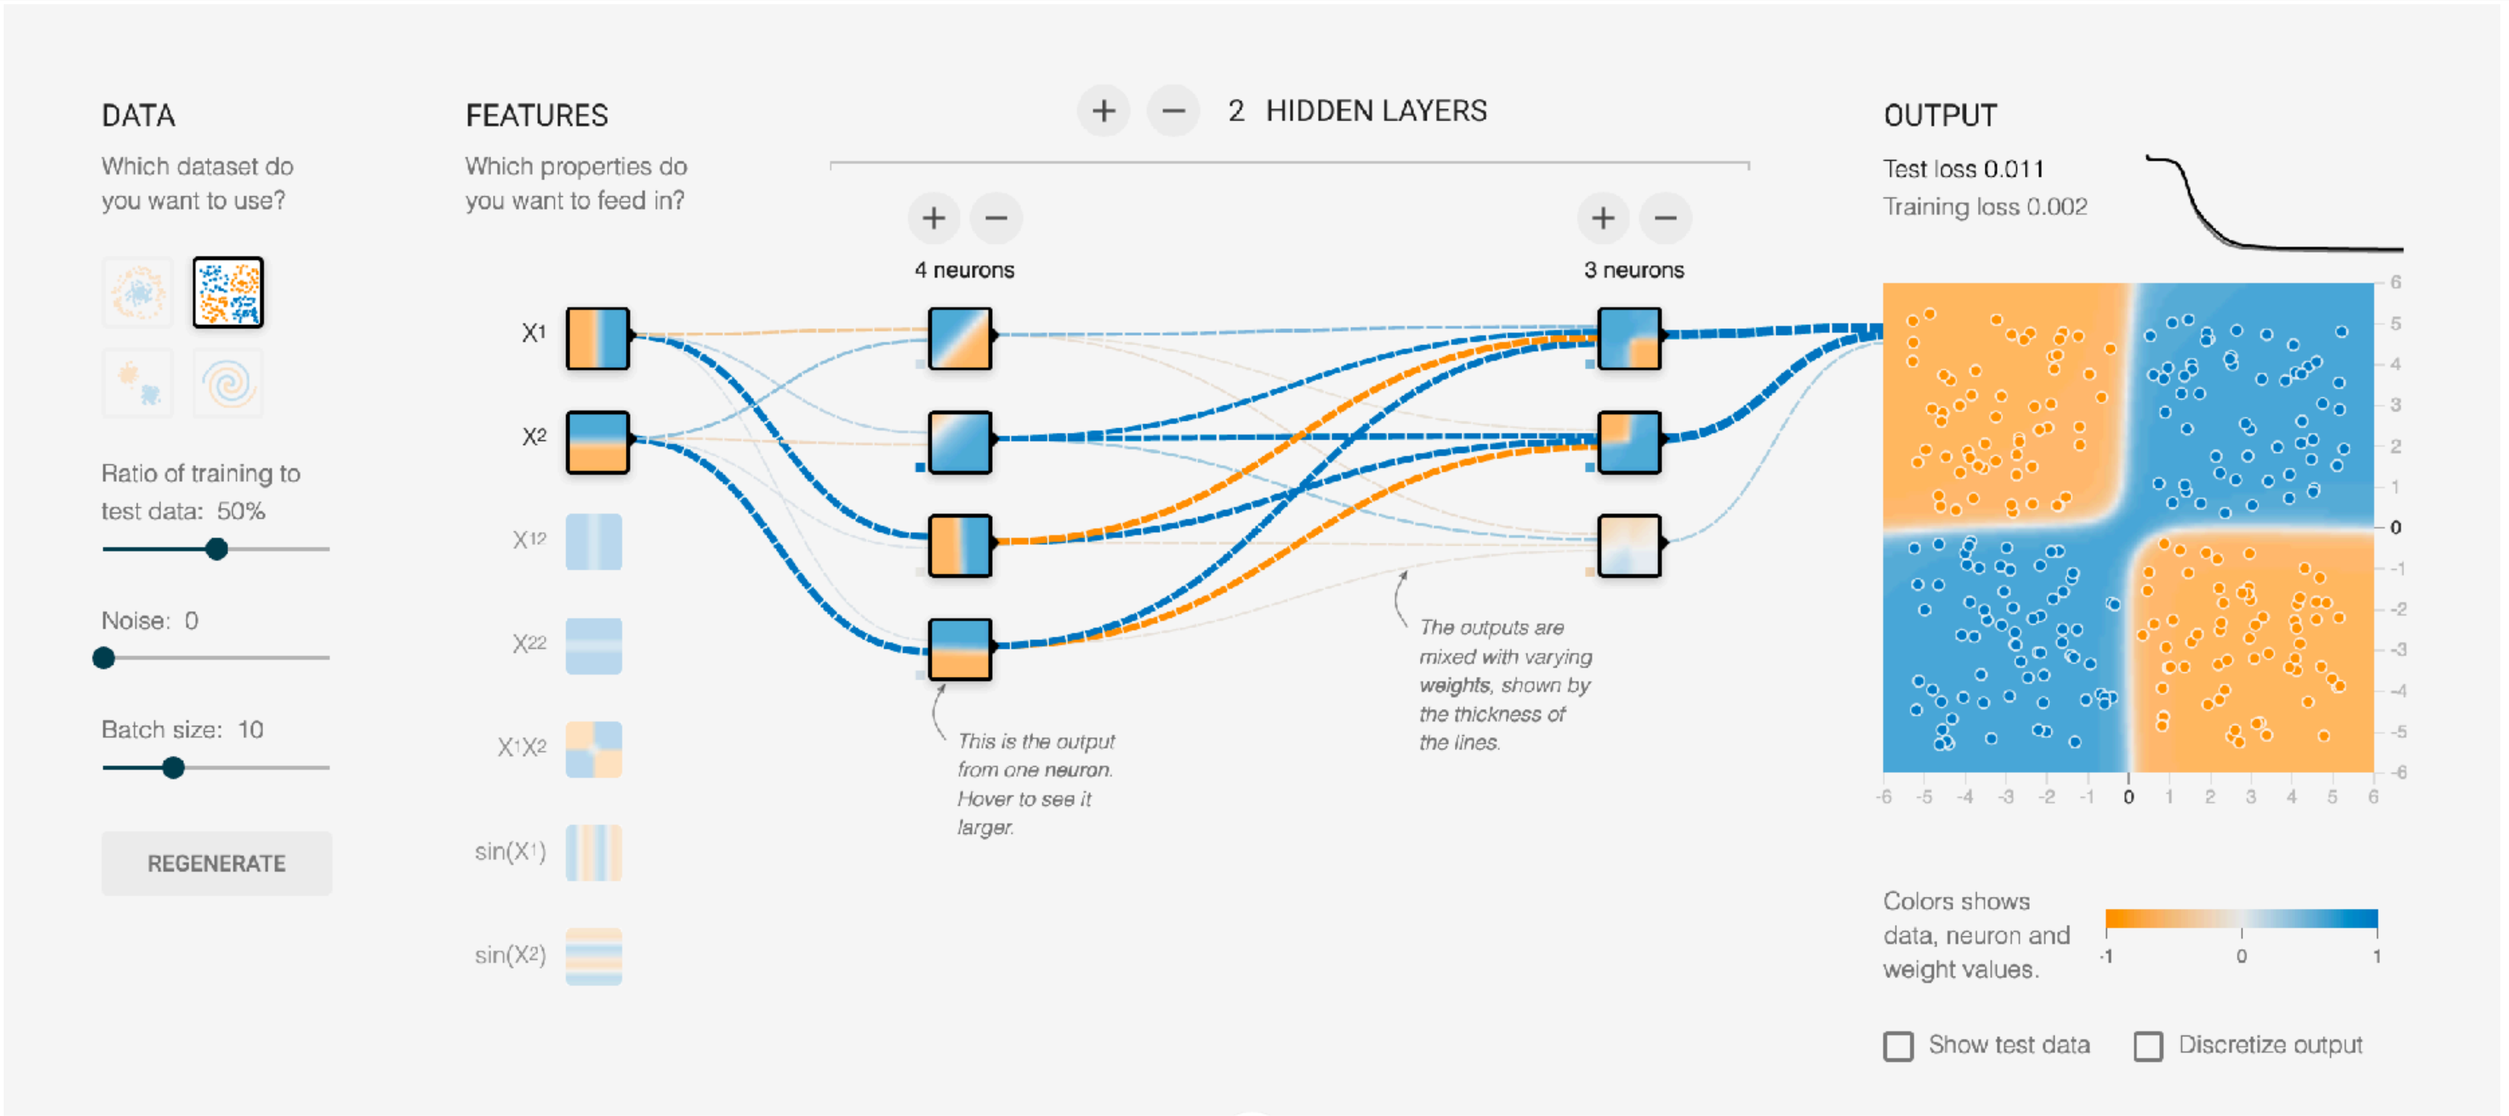

https://playground.tensorflow.org

Why multiple layers?

Example: house price prediction model (designed by humans)

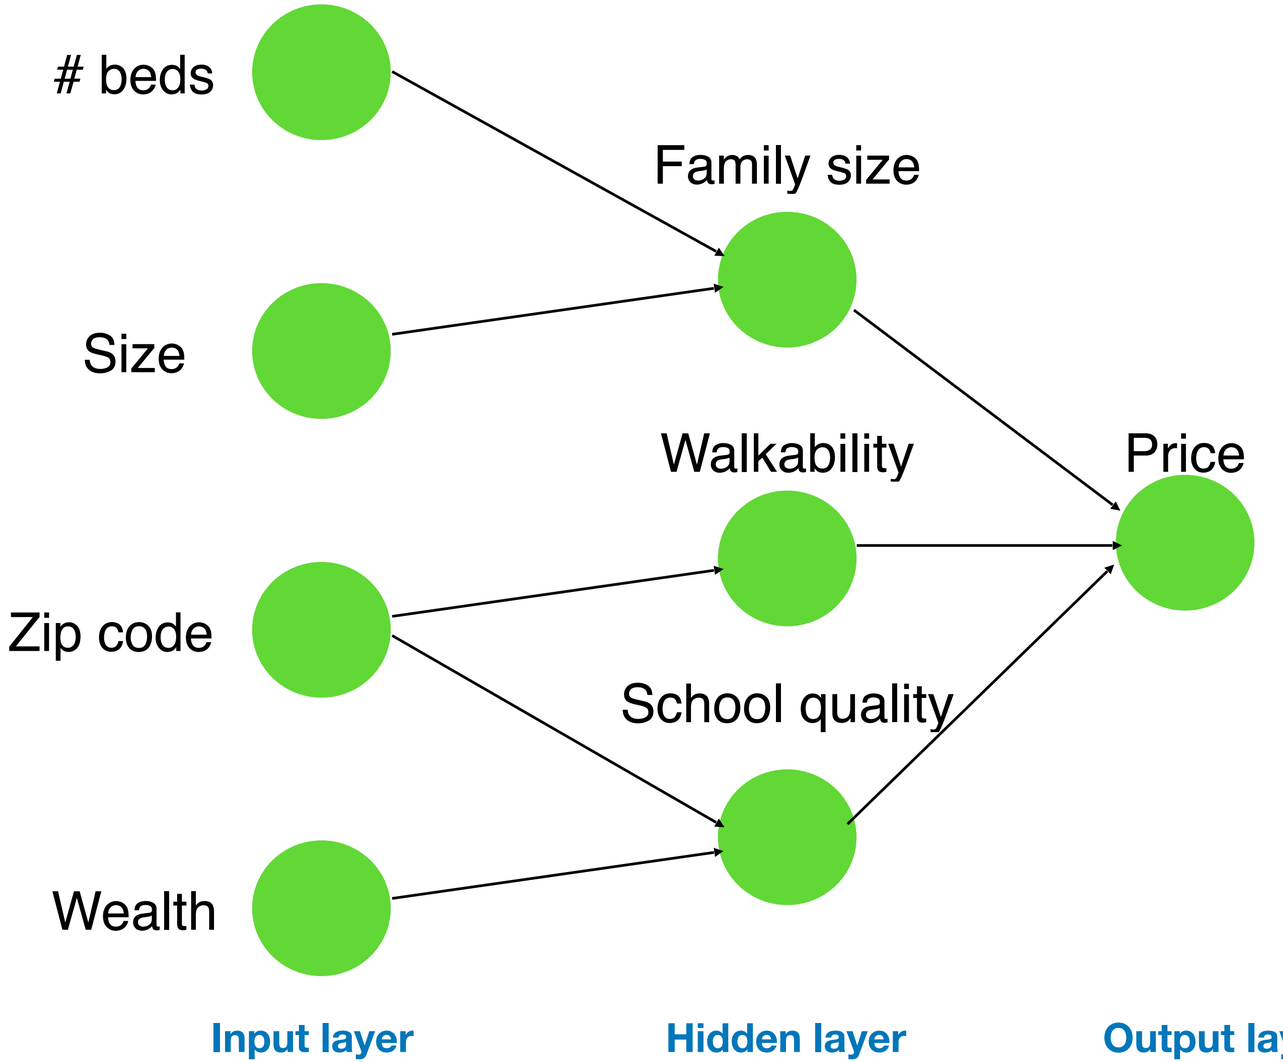

Why multiple layers?

Example: house price prediction model (designed by machines)

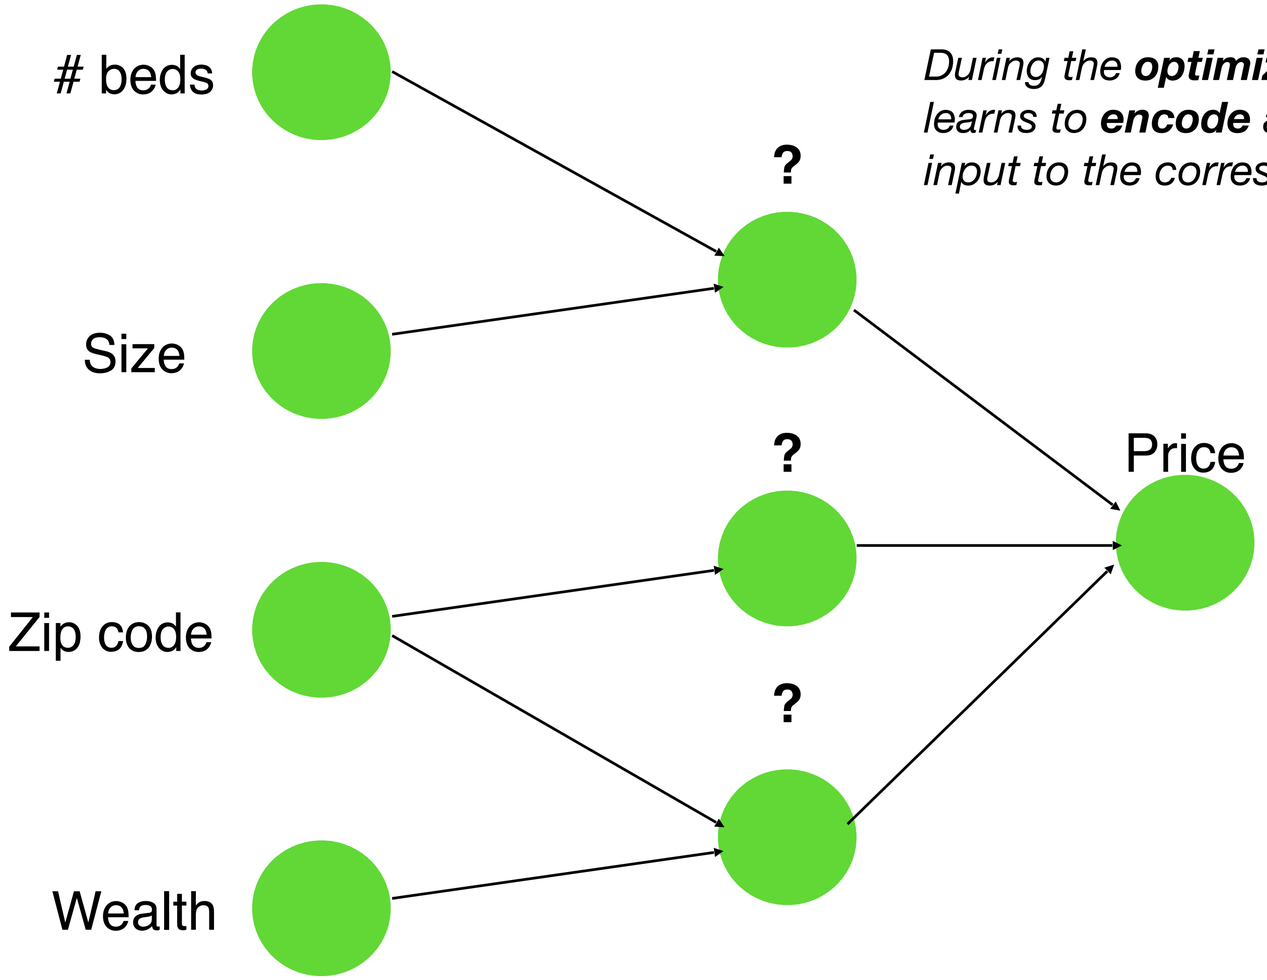

*During the **optimization** process, the machine learns to **encode** a presentation that **maps** the input to the corresponding output.*

Multi-Layer Perception

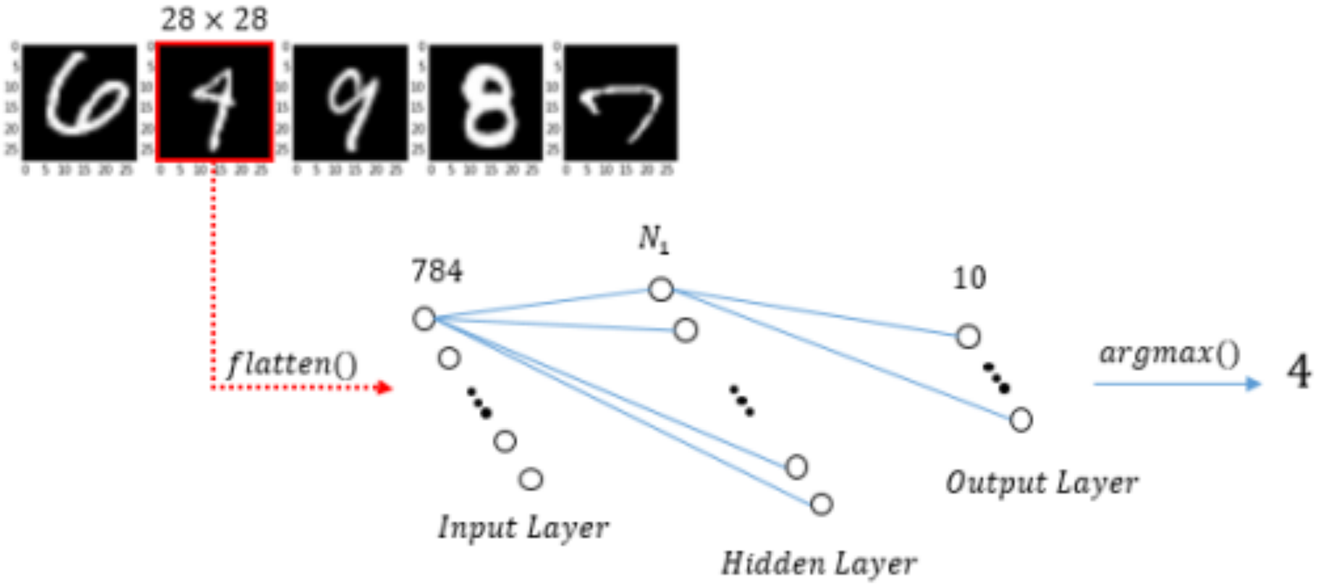

MNIST dataset
(**M**odified **N**ational **I**nstitute for **S**tandards and **T**echnology)

A DNN encodes the representation hierarchically

Deep neural networks learn hierarchical feature representations

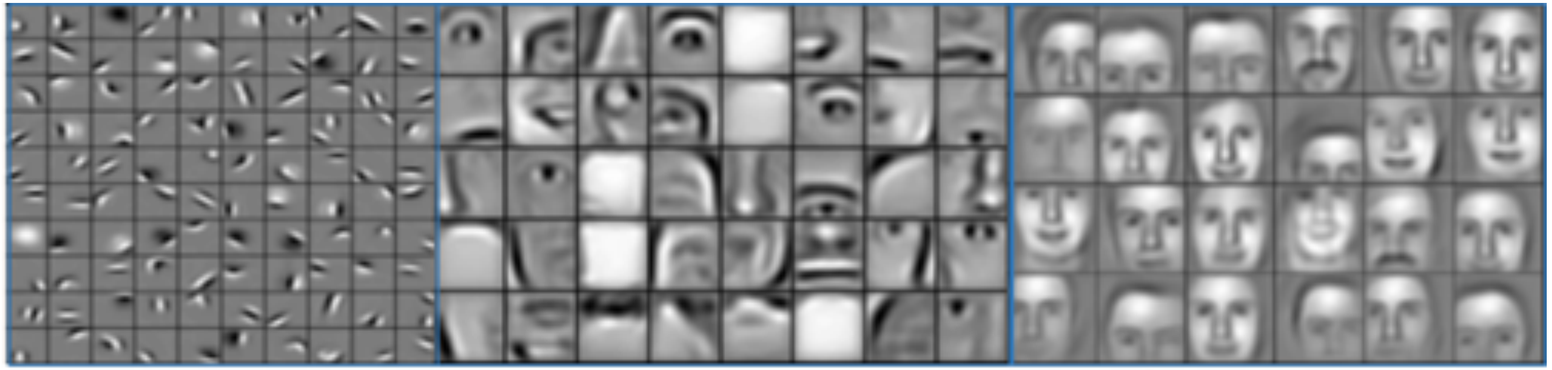

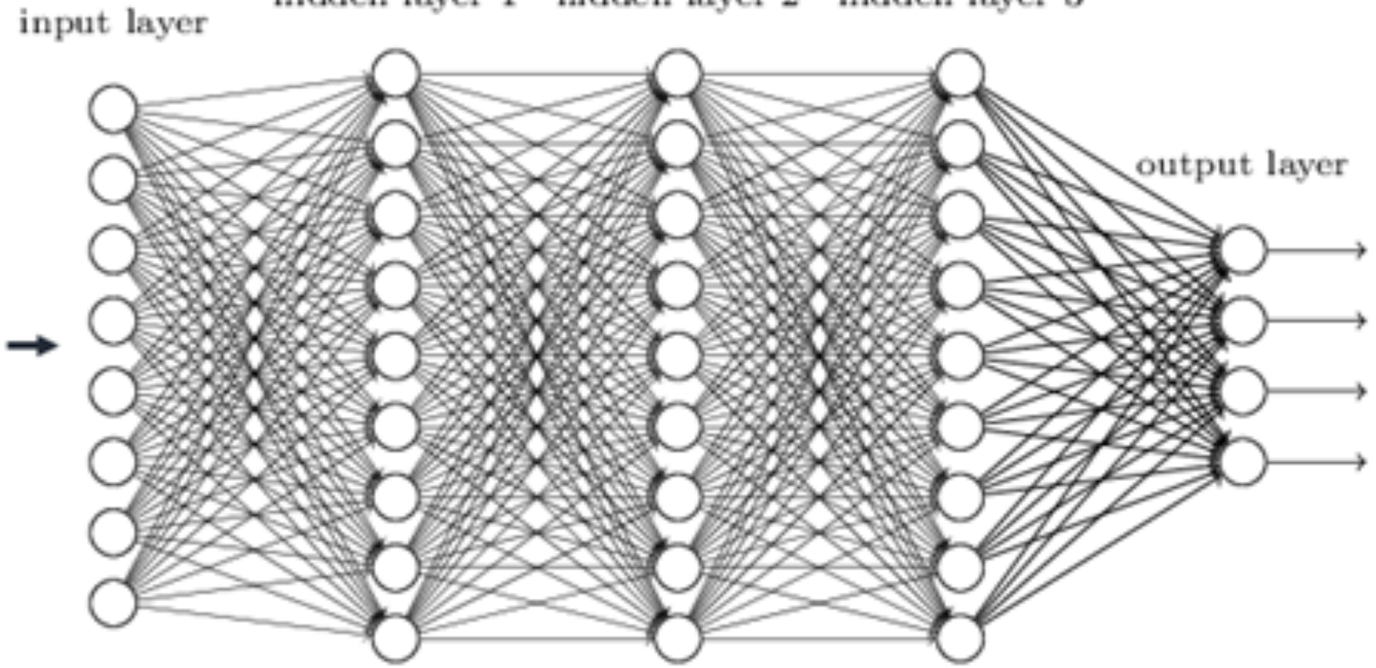

Prediction: forward propagation

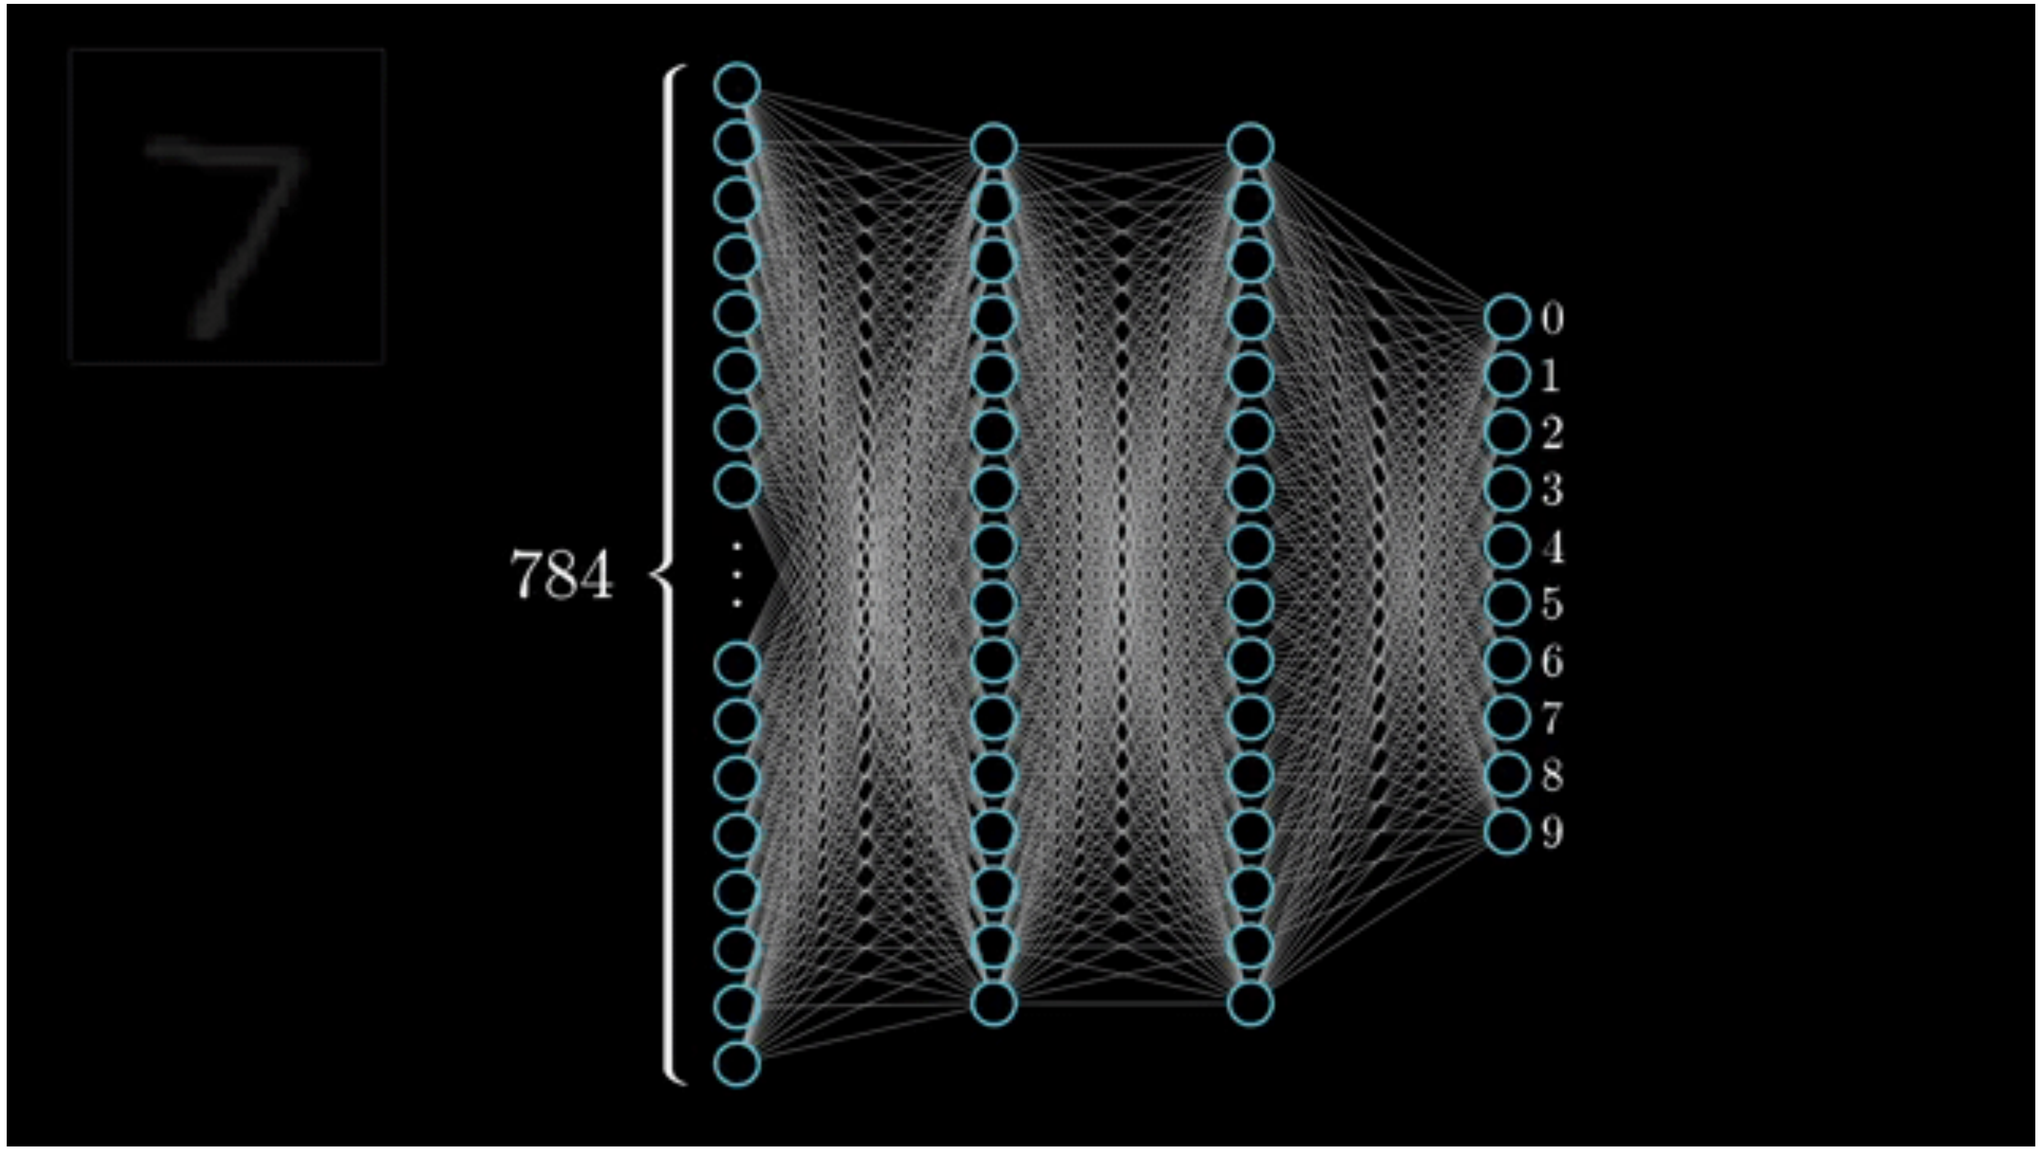

Source: 3blue1brown

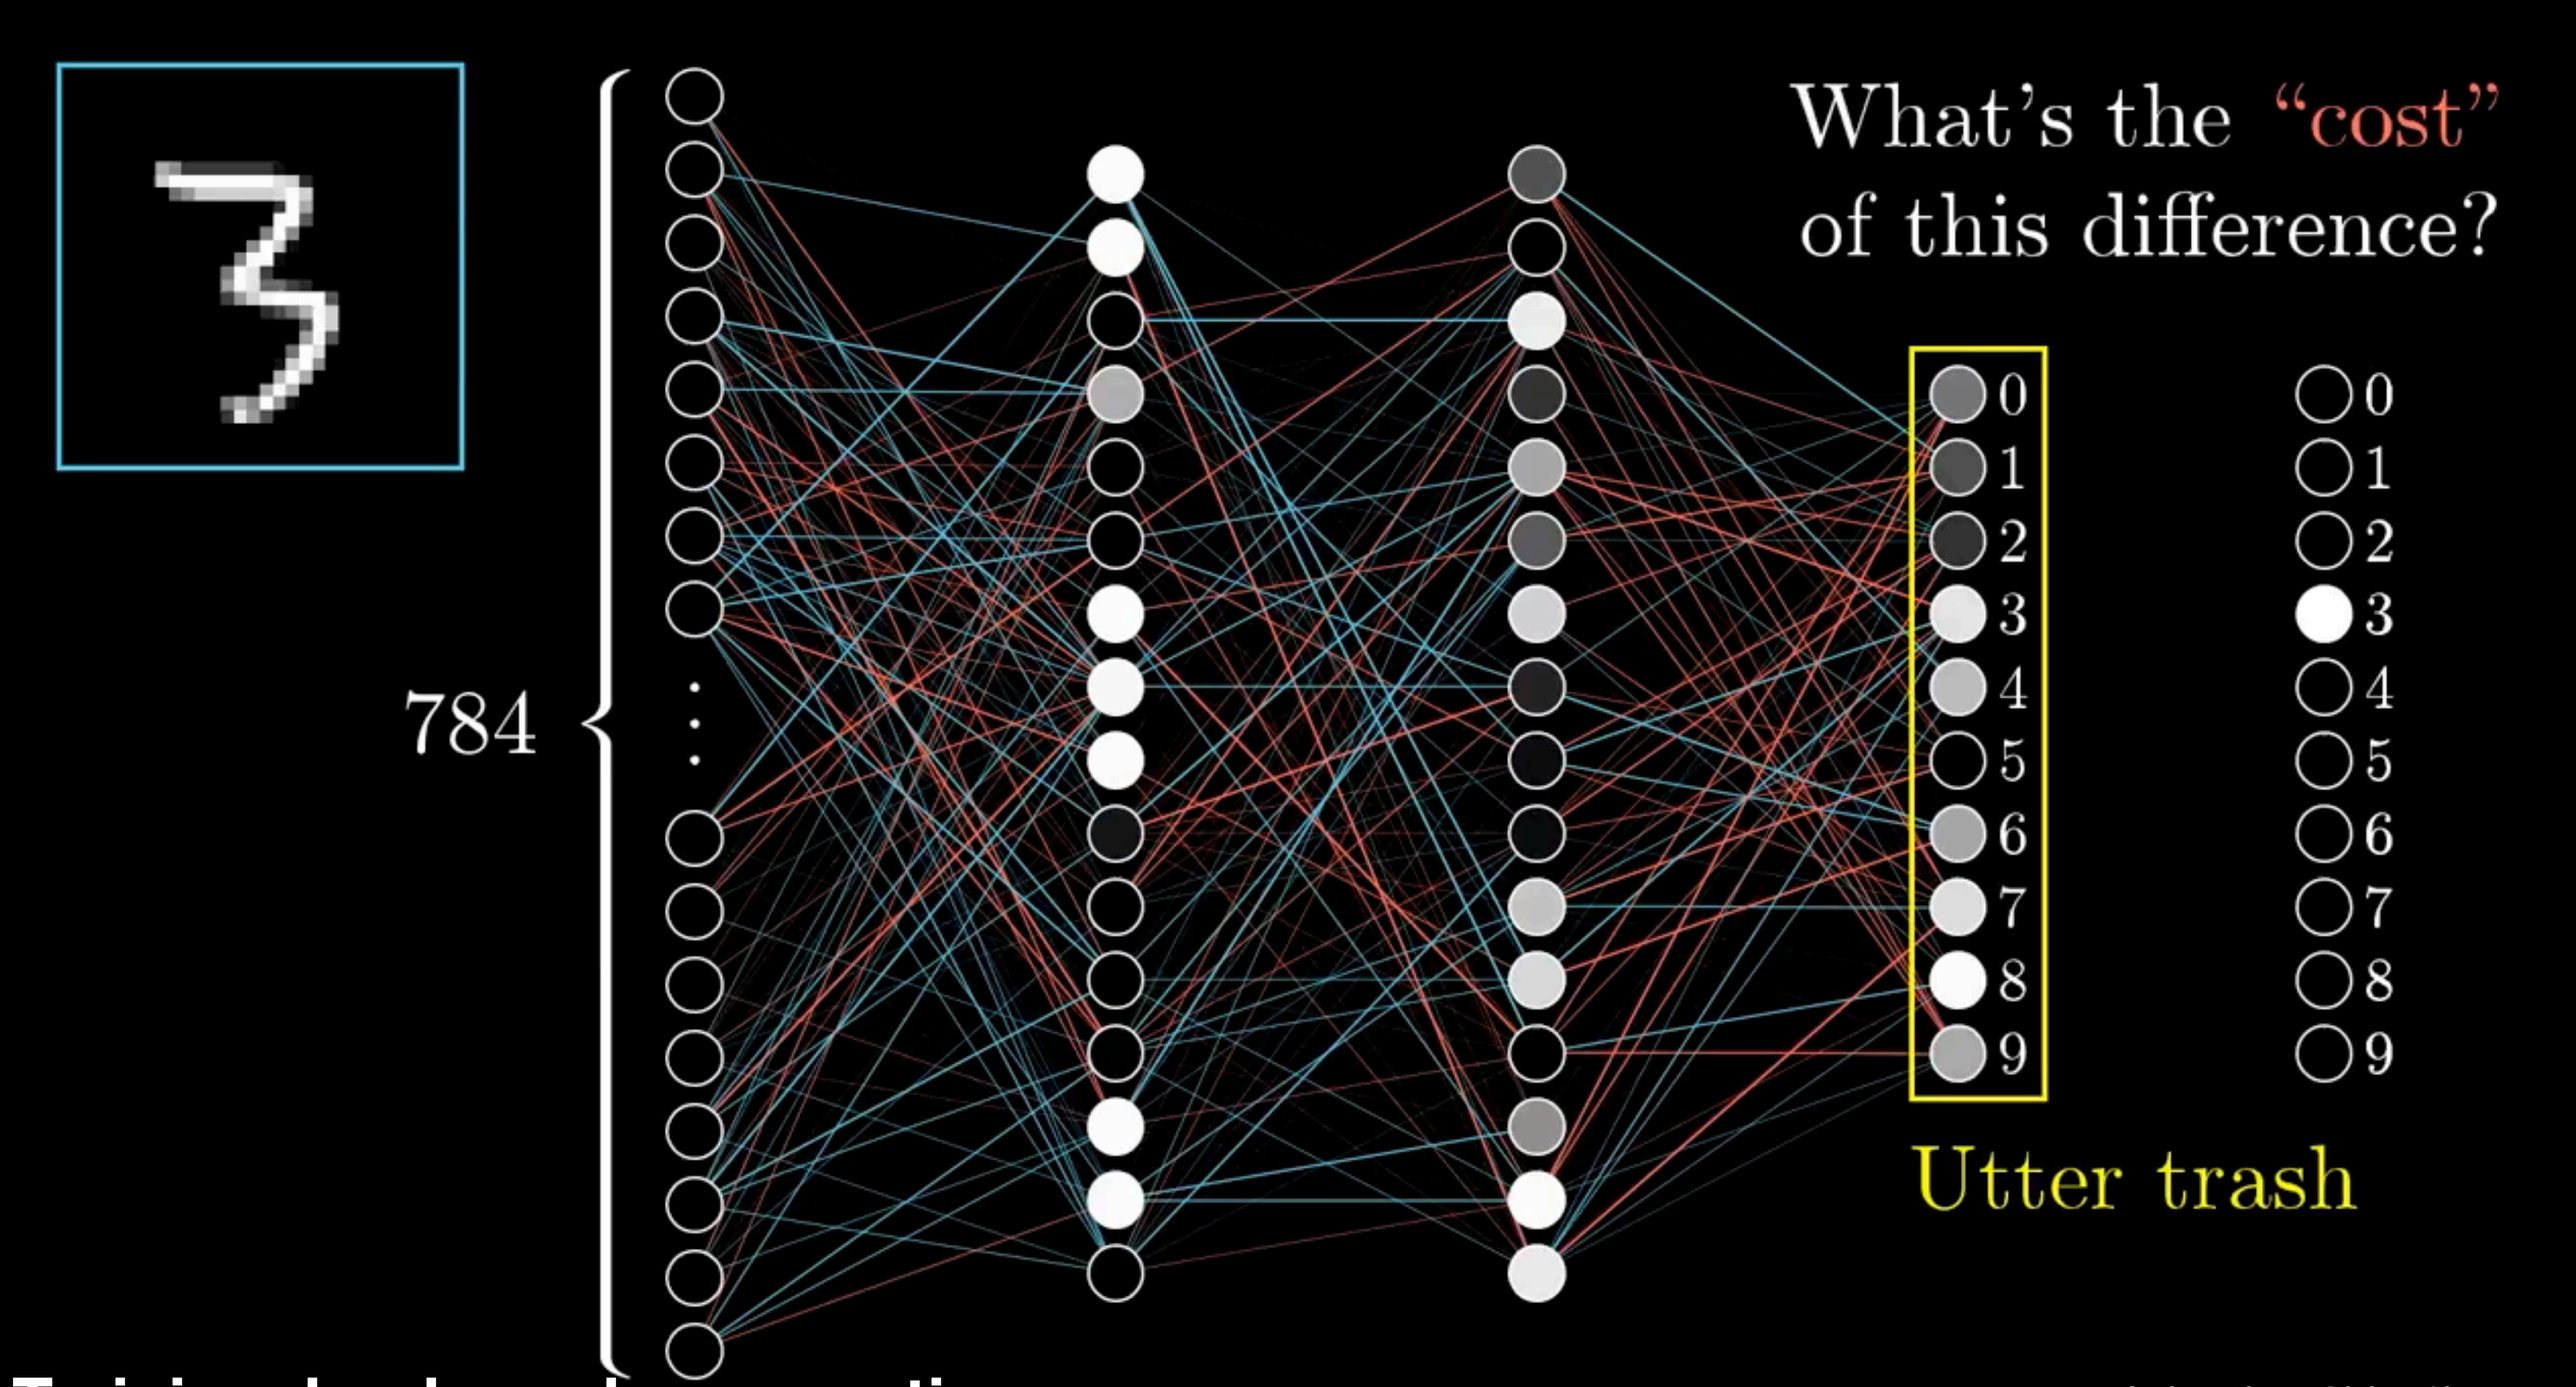

Training: backward propagation

Animation: 3blue1brown

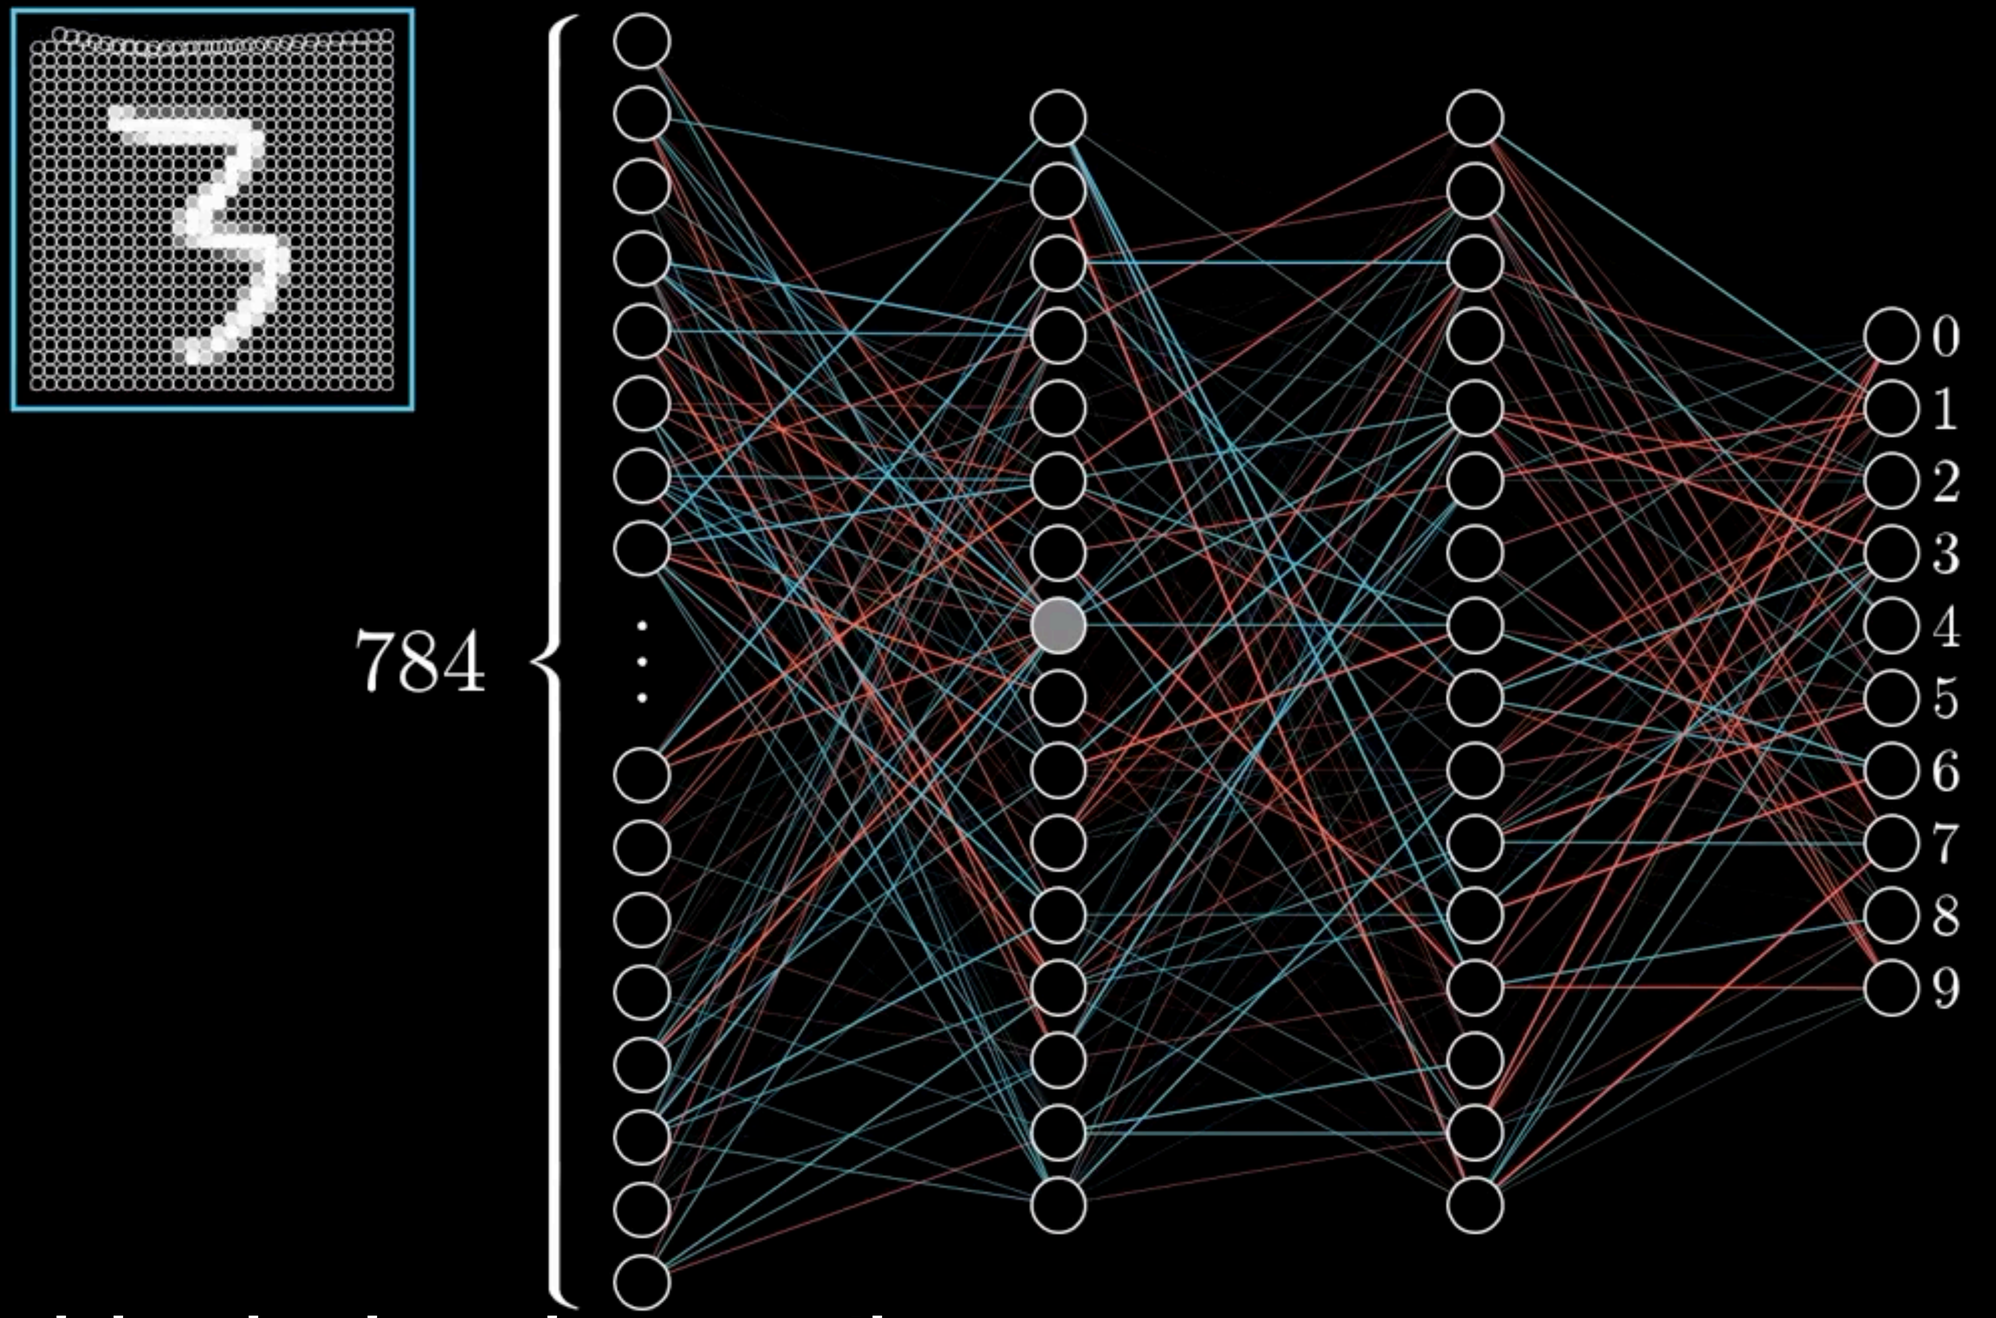

Training: backward propagation

Animation: 3blue1brown

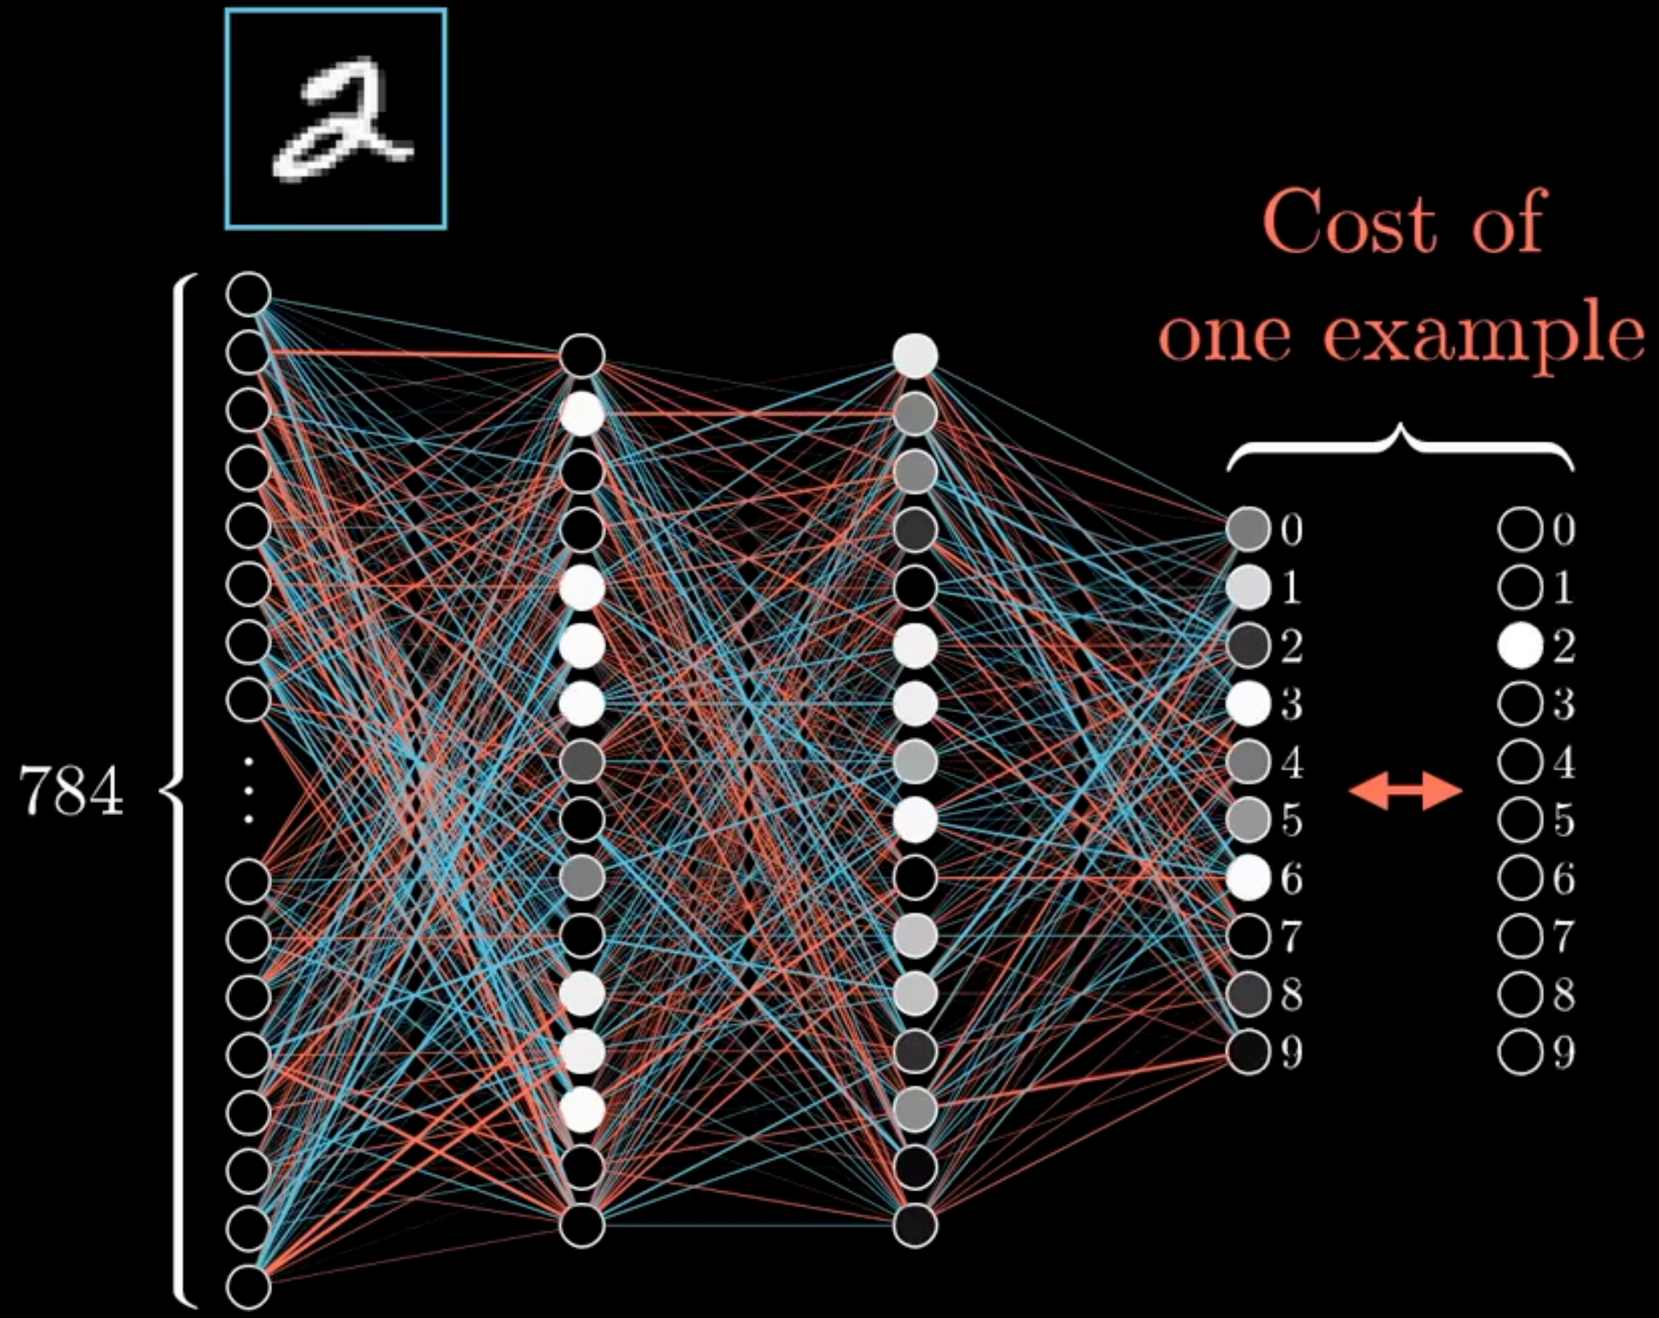

Training: backward propagation

Animation: 3blue1brown

Training: backward propagation

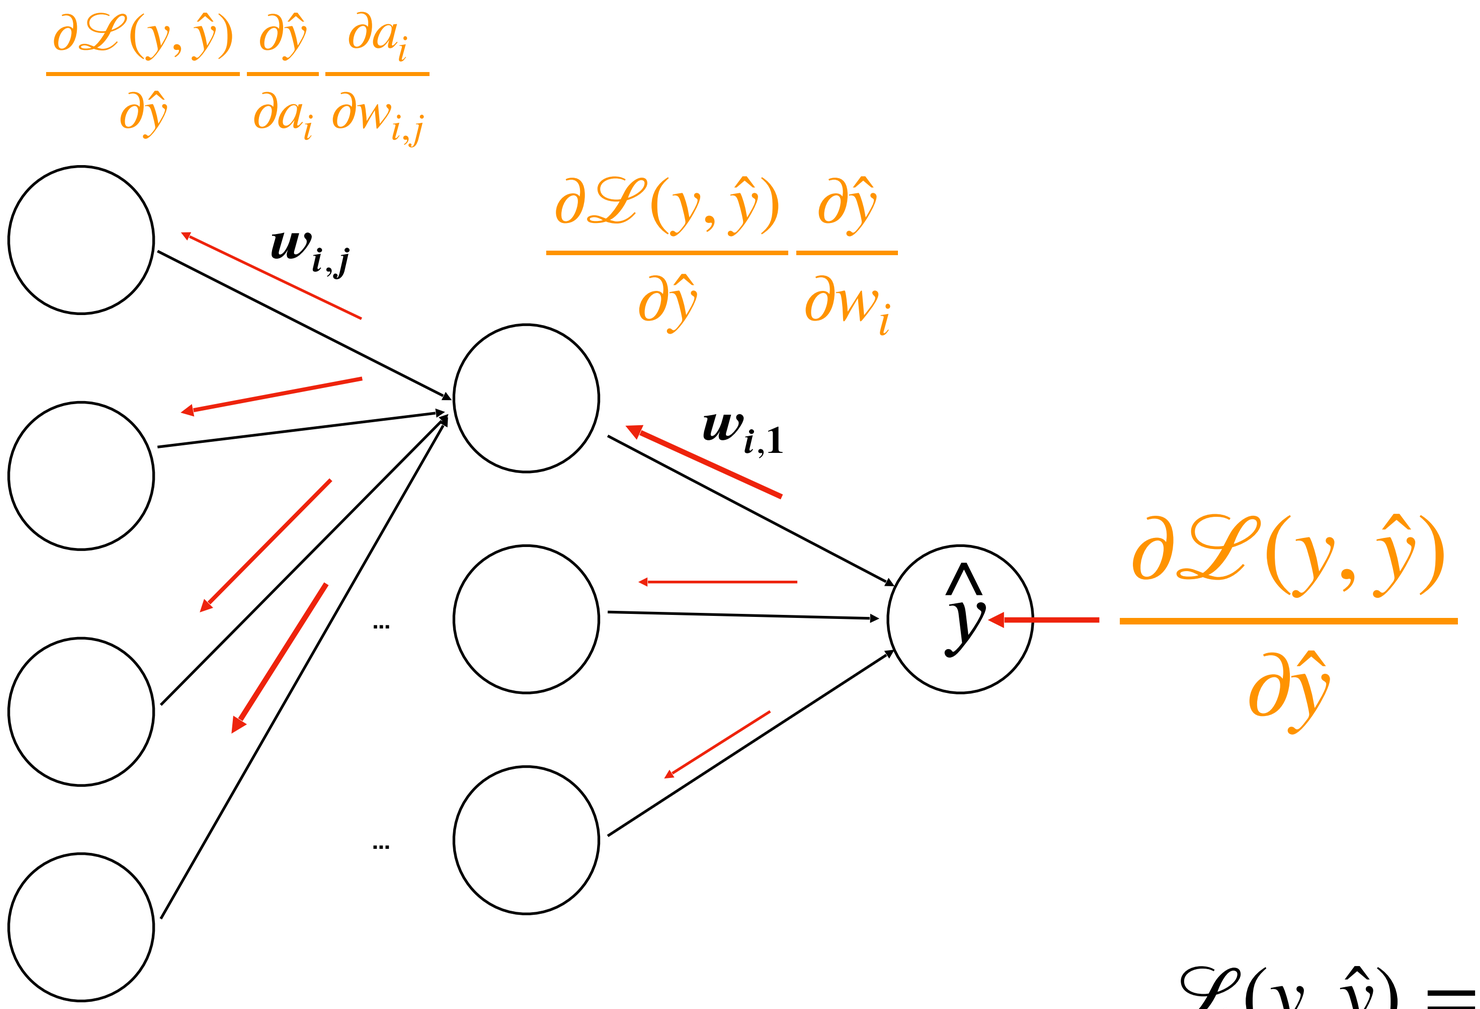

- In the context of DL we need to compute the gradient for each layer.
- We do this by applying the **chain rule** of derivatives.
- This algorithm is known as **backpropagation**.

$$\mathscr{L}(y,\hat{y}) = L(\mathbf{W},b) = \frac{1}{N}\sum_{i=1}^{N}\left(y_i - \hat{y}_i\right)^2$$

Neural Network: the deeper, the better?

Not really.

The vanishing gradient problem

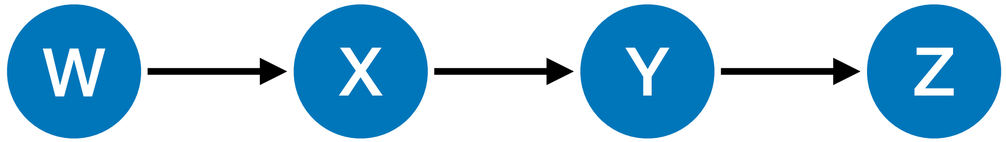

Chain rule

$$\frac{\partial z}{\partial w} = \frac{\partial z}{\partial y}\frac{\partial y}{\partial x}\frac{\partial x}{\partial w}$$

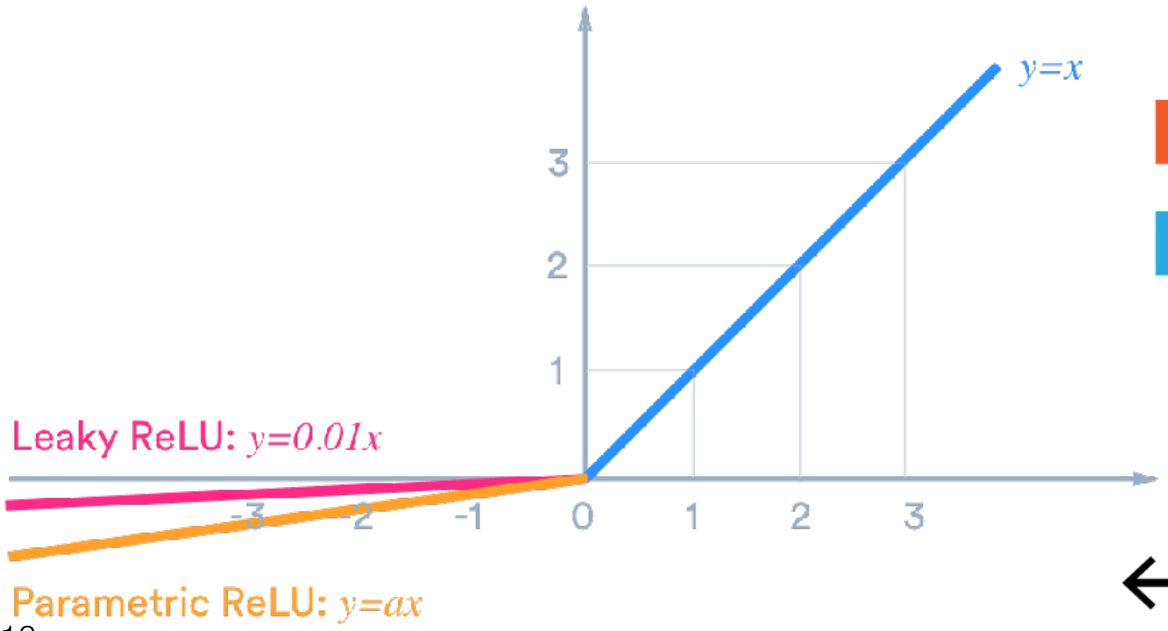

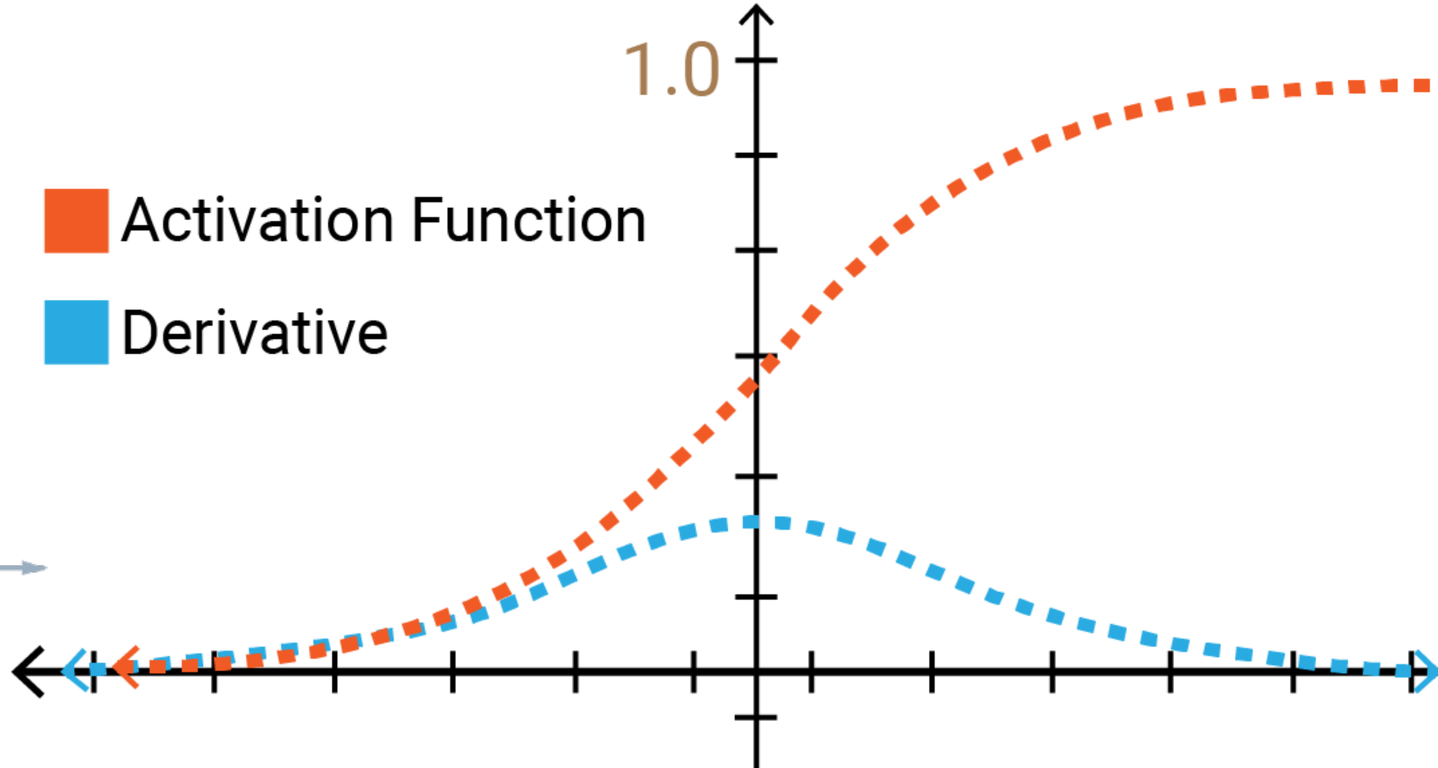

Activation Functions

Sigmoid

$$\sigma(x) = \frac{1}{1+e^{-x}}$$

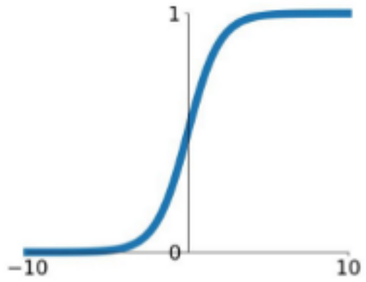

tanh

$$\tanh(x)$$

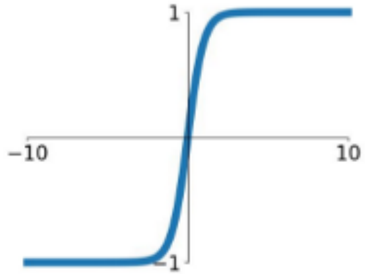

ReLU

$$\max(0, x)$$

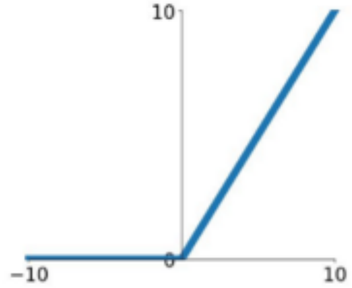

Leaky ReLU

$$\max(0.1x, x)$$

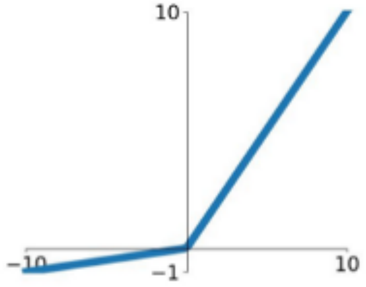

Maxout

$$\max(w_1^T x + b_1, w_2^T x + b_2)$$

ELU

$$\begin{cases} x & x \geq 0 \\ \alpha(e^x - 1) & x < 0 \end{cases}$$

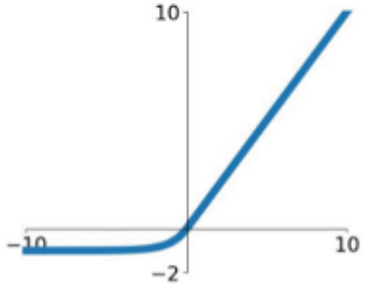

But one can design their own activation functions!

Optimizers

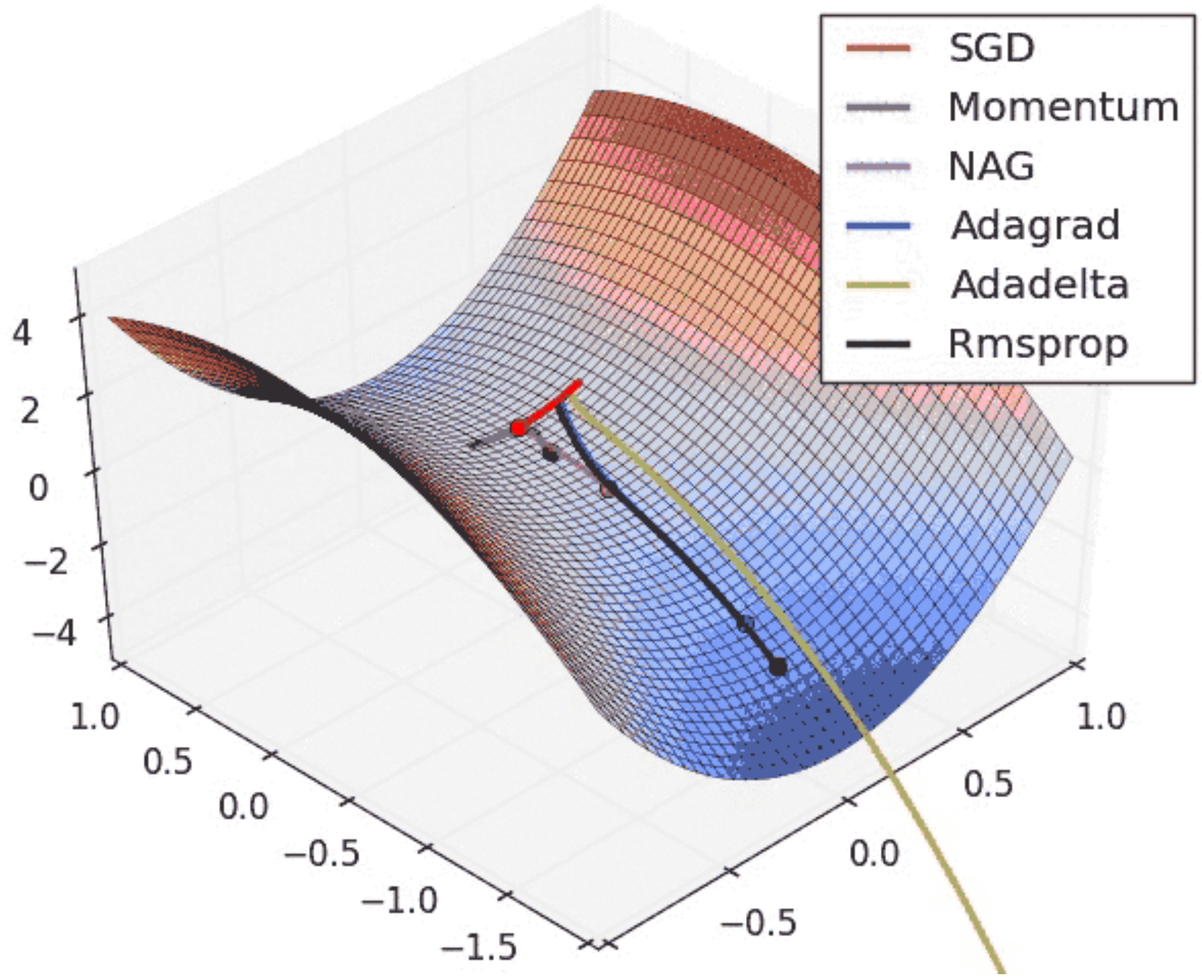

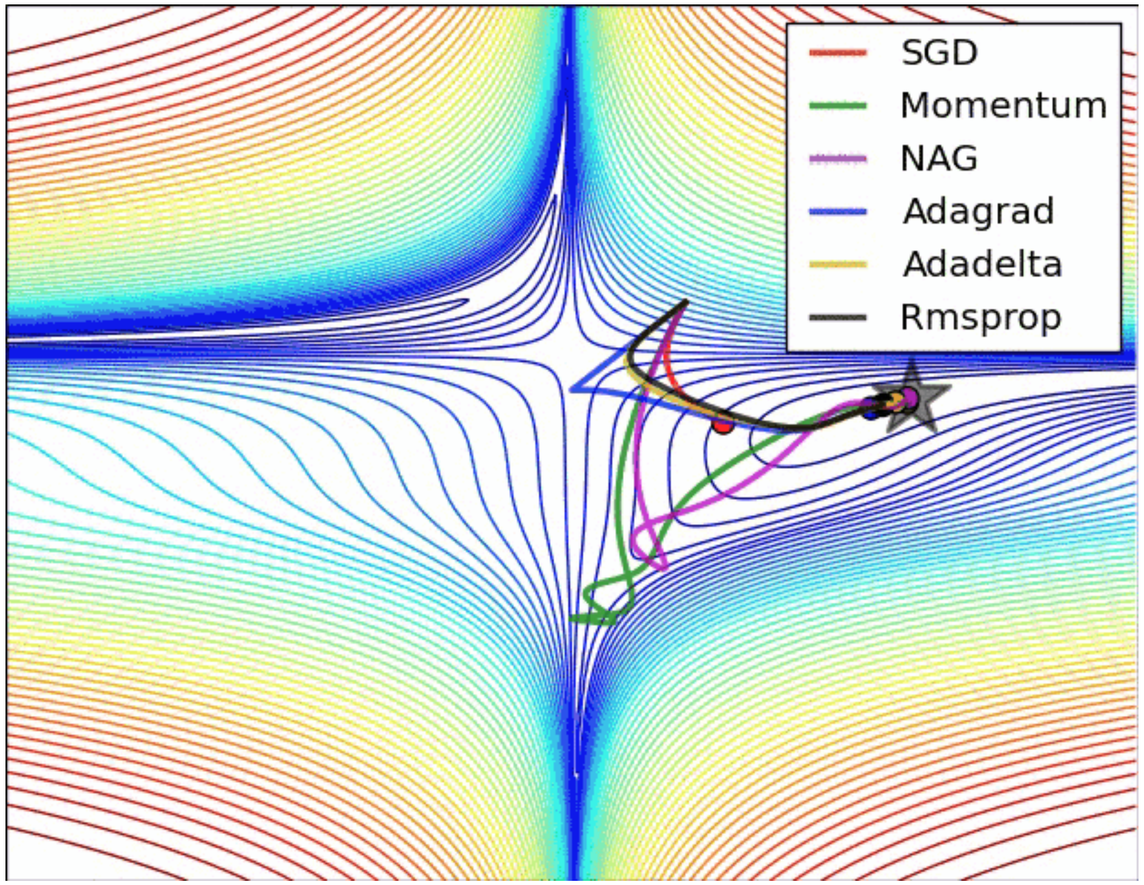

Source: medium.com/analytics-vidhya

Common practice for loss functions

Regression
- Mean squared error
- Mean squared logarithmic error
- Mean absolute error

Binary Classification
- Binary cross-entropy
- Hinge loss
- Squared hinge loss

Multi-Class Classification
- Multi-class cross-entropy
- Sparse multi-class cross-entropy
- Kullback-Leibler divergence

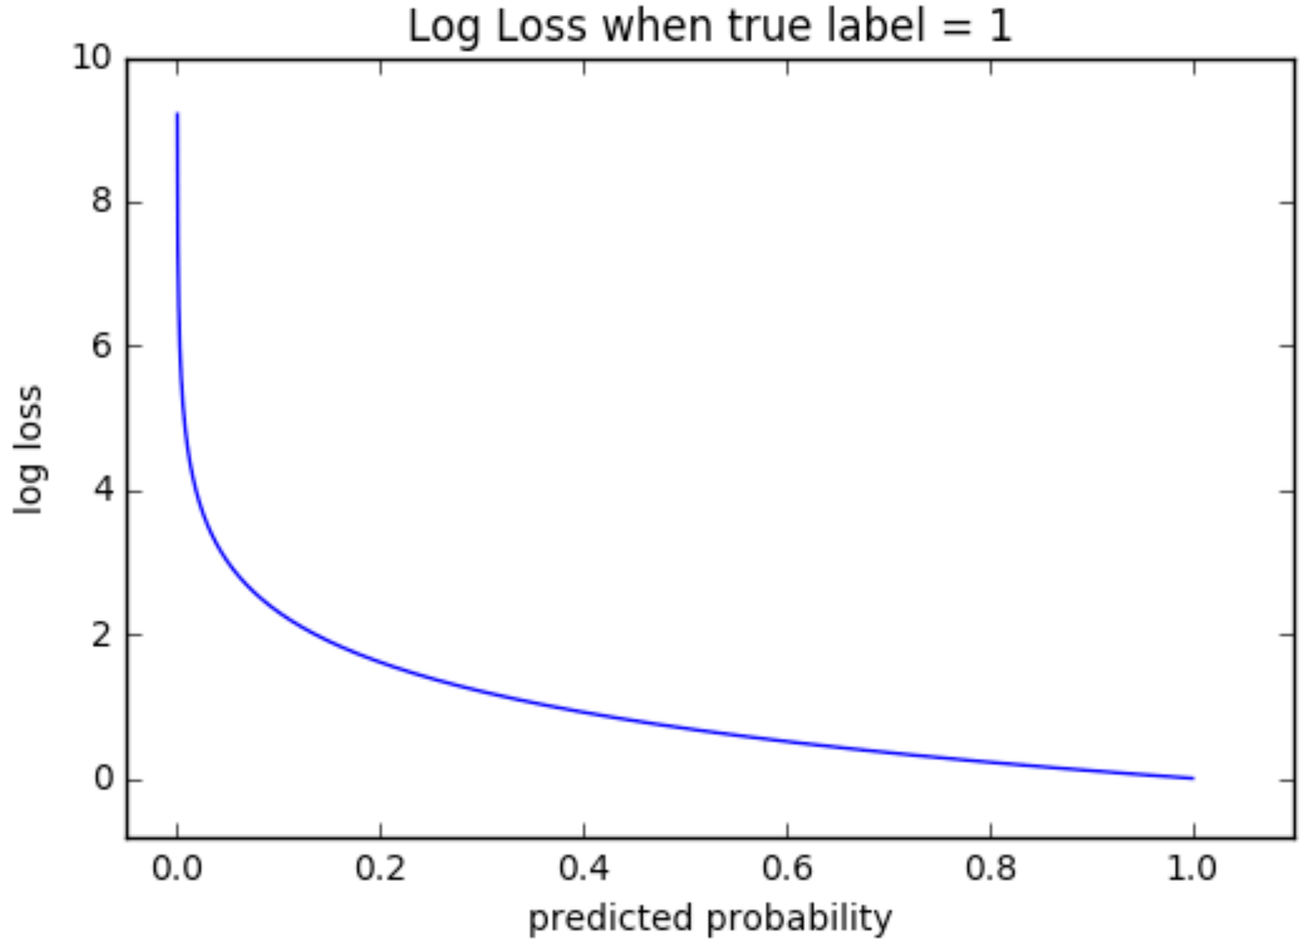

Cross-entropy loss outputs a log probability

DL frameworks

In DL, you need to

- Define neurons and layers
- Define loss function
- Calculate losses
- Calculate gradient (multivariate calculus)
- Backward propagation
- Update weights

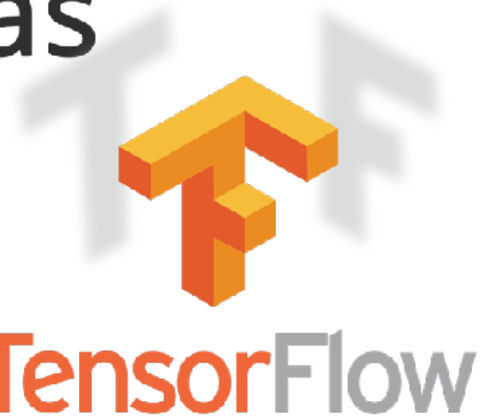

- Many frameworks exist; **TensorFlow**, **CNTK**, **Torch**, **Keras**, **Theano**, **Caffe**, ...
- We will use **TensorFlow/Keras**
- Keras used to call TensorFlow as a *backend*, but is now fully integrated in TensorFlow.

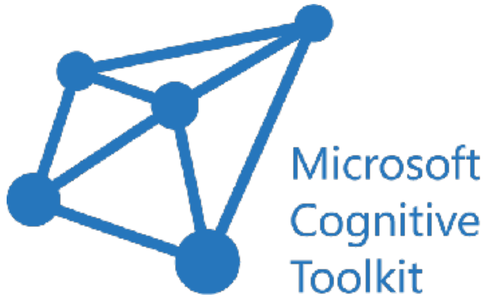

Model Evaluation

Balanced/Imbalanced training set

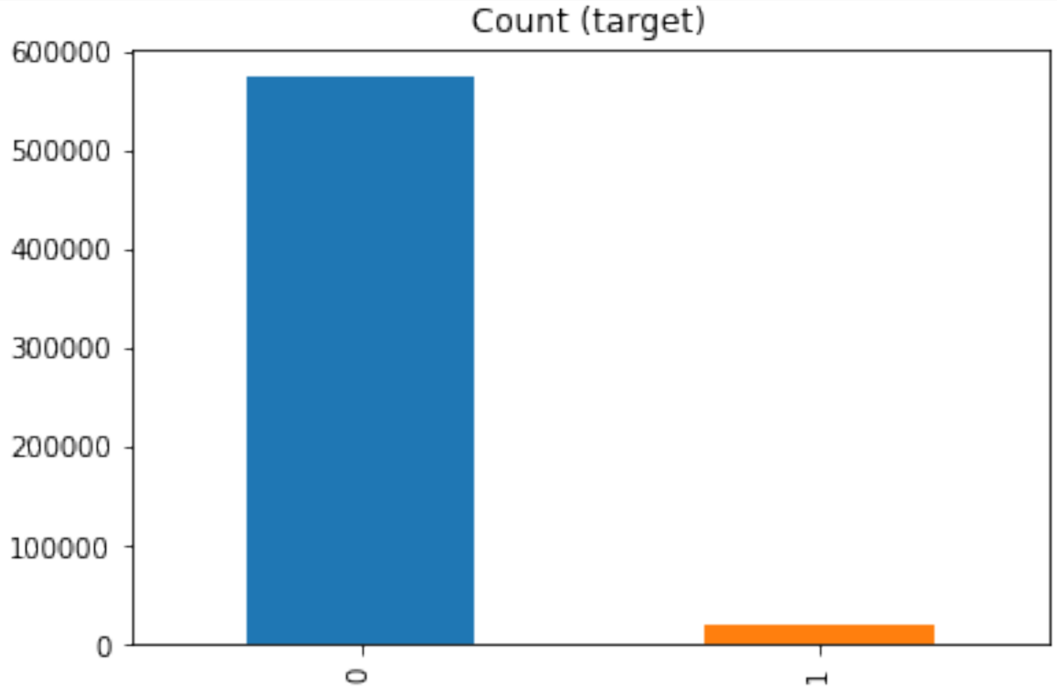

Pay attention to your data: they may fool your model.

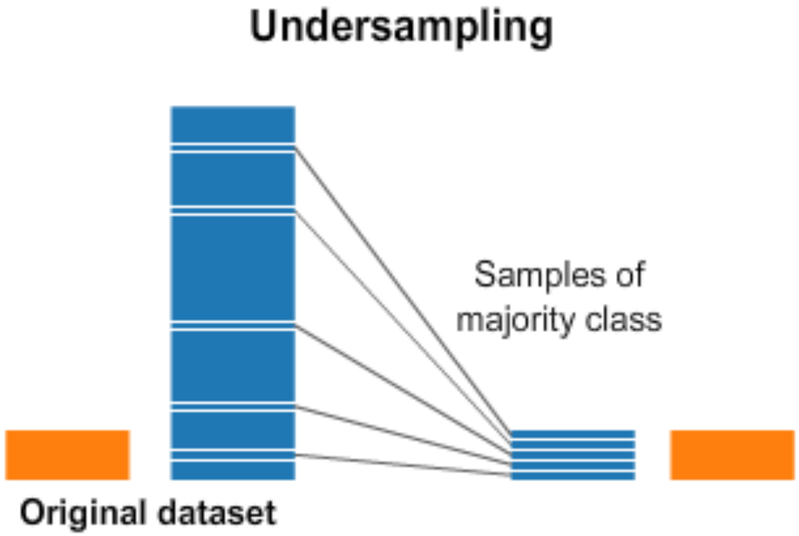

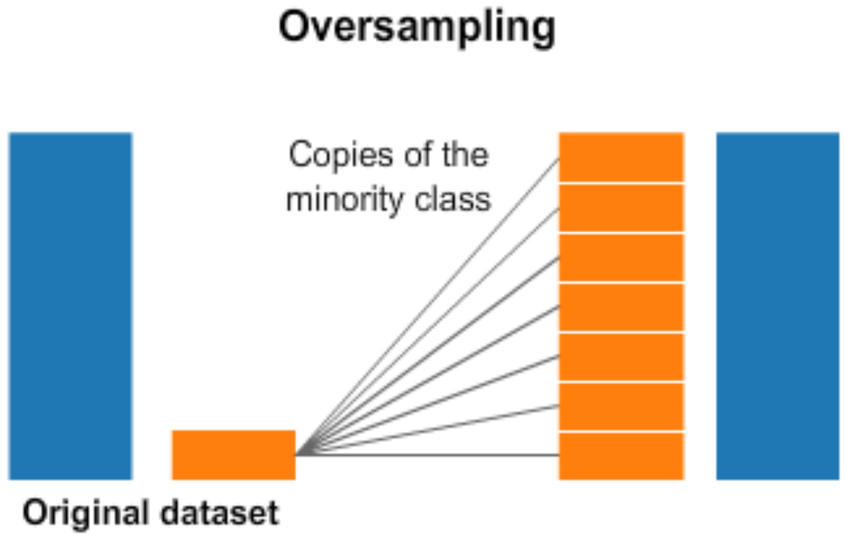

Source: kaggle.com

Data Normalization

A process to transform the input **data** in a **well-behaved** form.

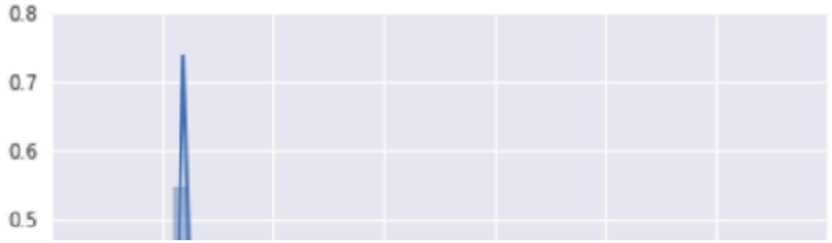

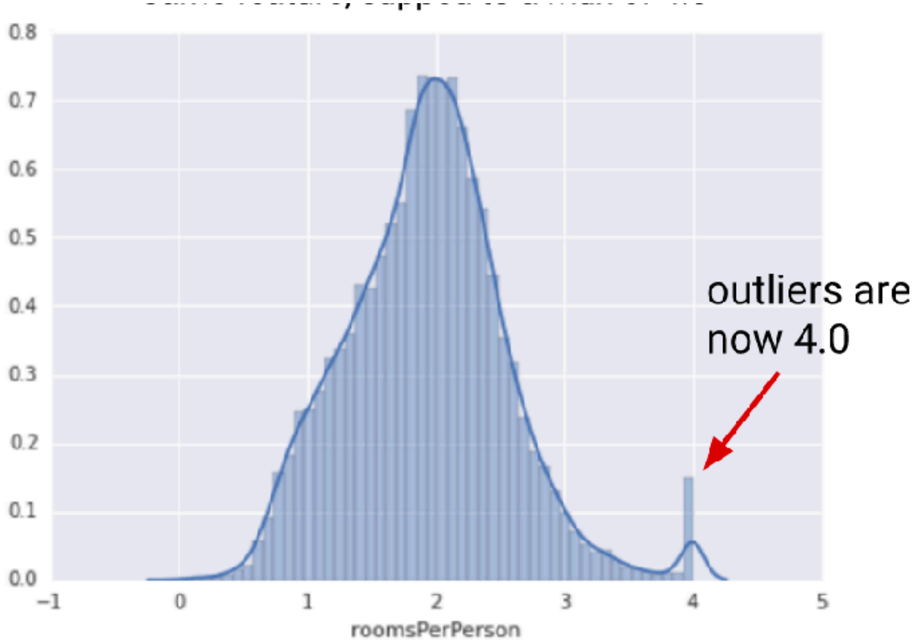

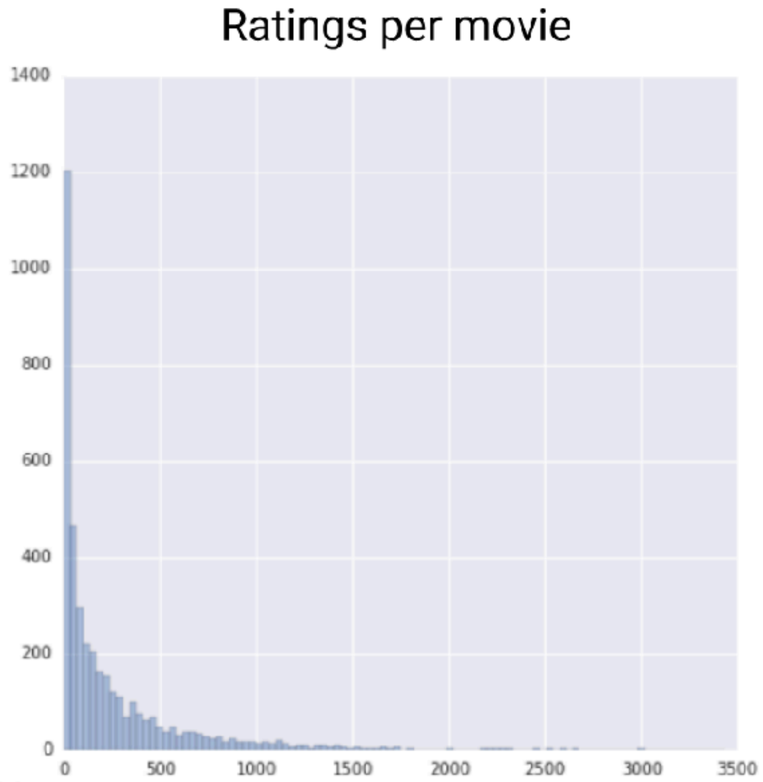

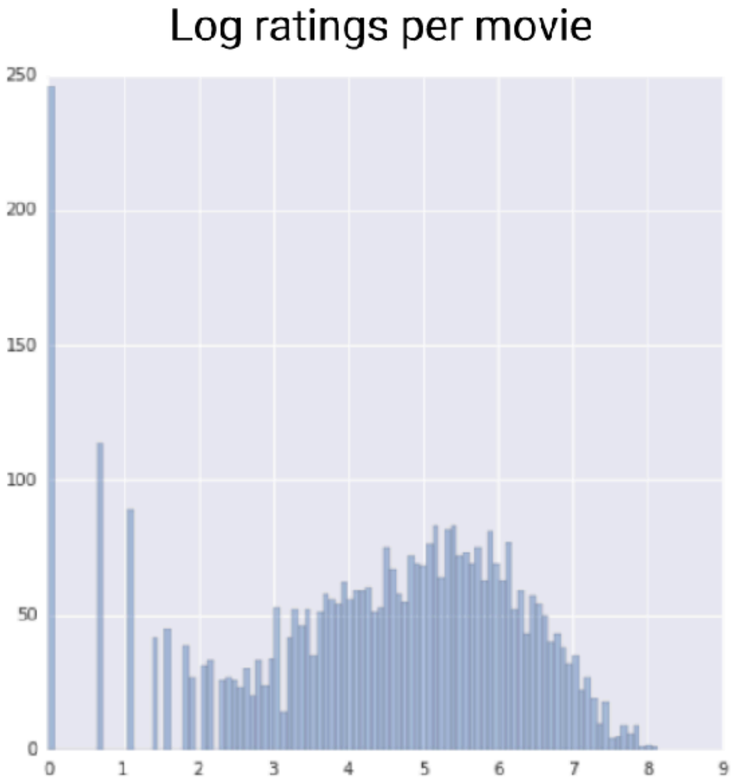

Further reading: sklearn scalers

Source: developers.google.com

Dataset splitting

Training set

Model training

Best model parameters
Weights

70%

Validation set

Model selection

Best hyperparameters
Learning rate
#neurons
#layers

20%

Test set

Model testing

Final model
Accuracy
Sensitivity/specificity

10%

Confusion Matrix

		Prediction outcome p	Prediction outcome n	total
actual value	p′	True Positive	False Negative	P′
	n′	False Positive	True Negative	N′
	total	P	N	

Accuracy = (TP + TN) / (TP + FN + FP + TN)
Precision (p) = TP / (TP + FP)
Recall (r) = TP / (TP + FN)

$$F_1 = \frac{2}{r^{-1} + p^{-1}}$$

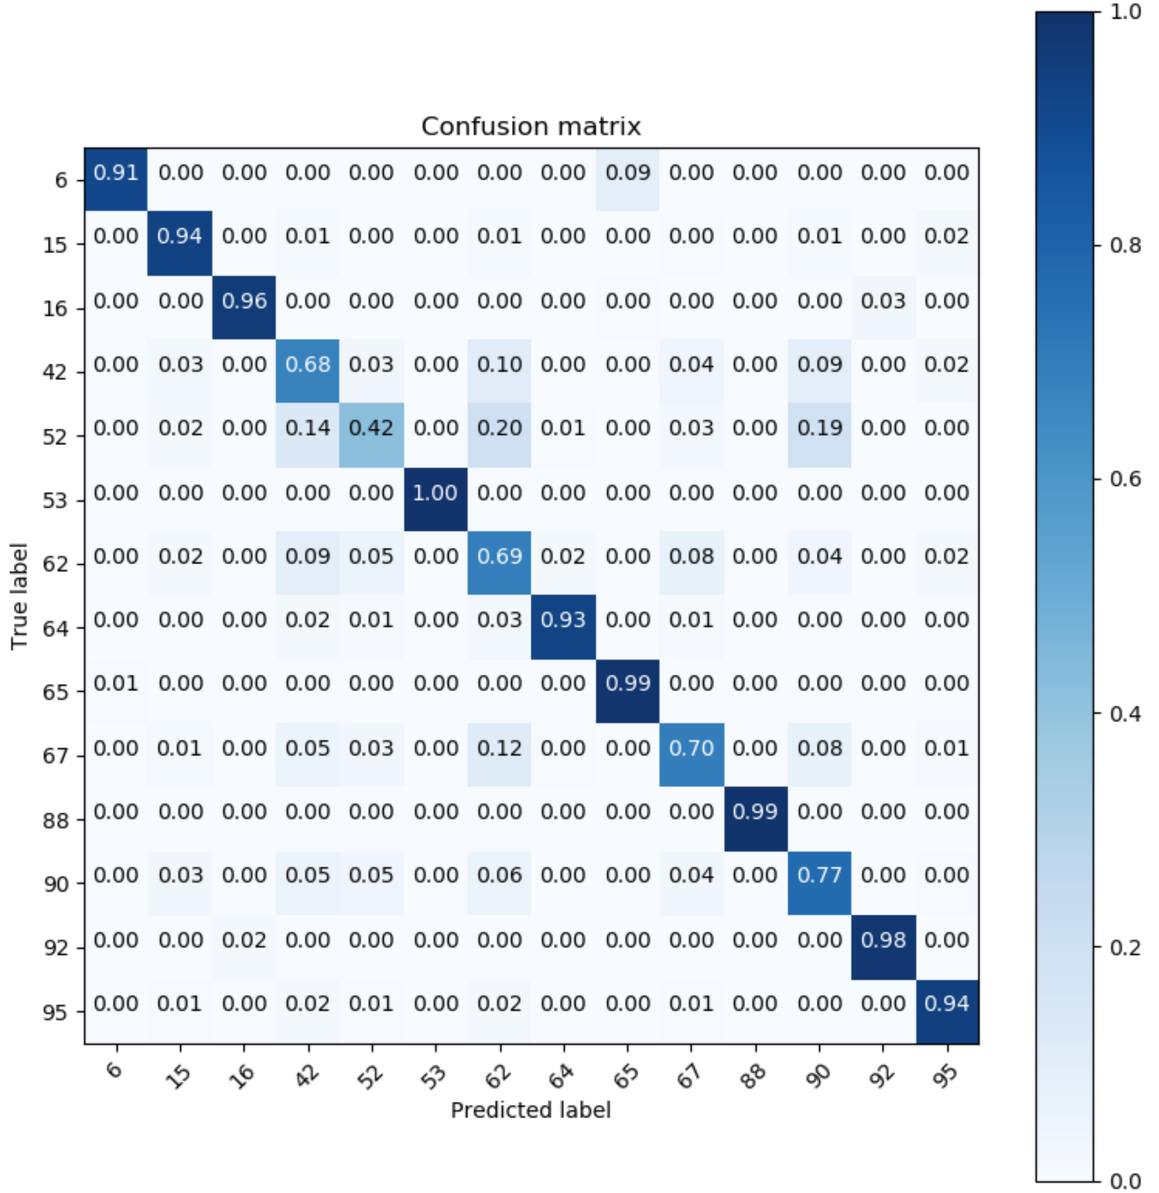

General Workflow of ML/DL

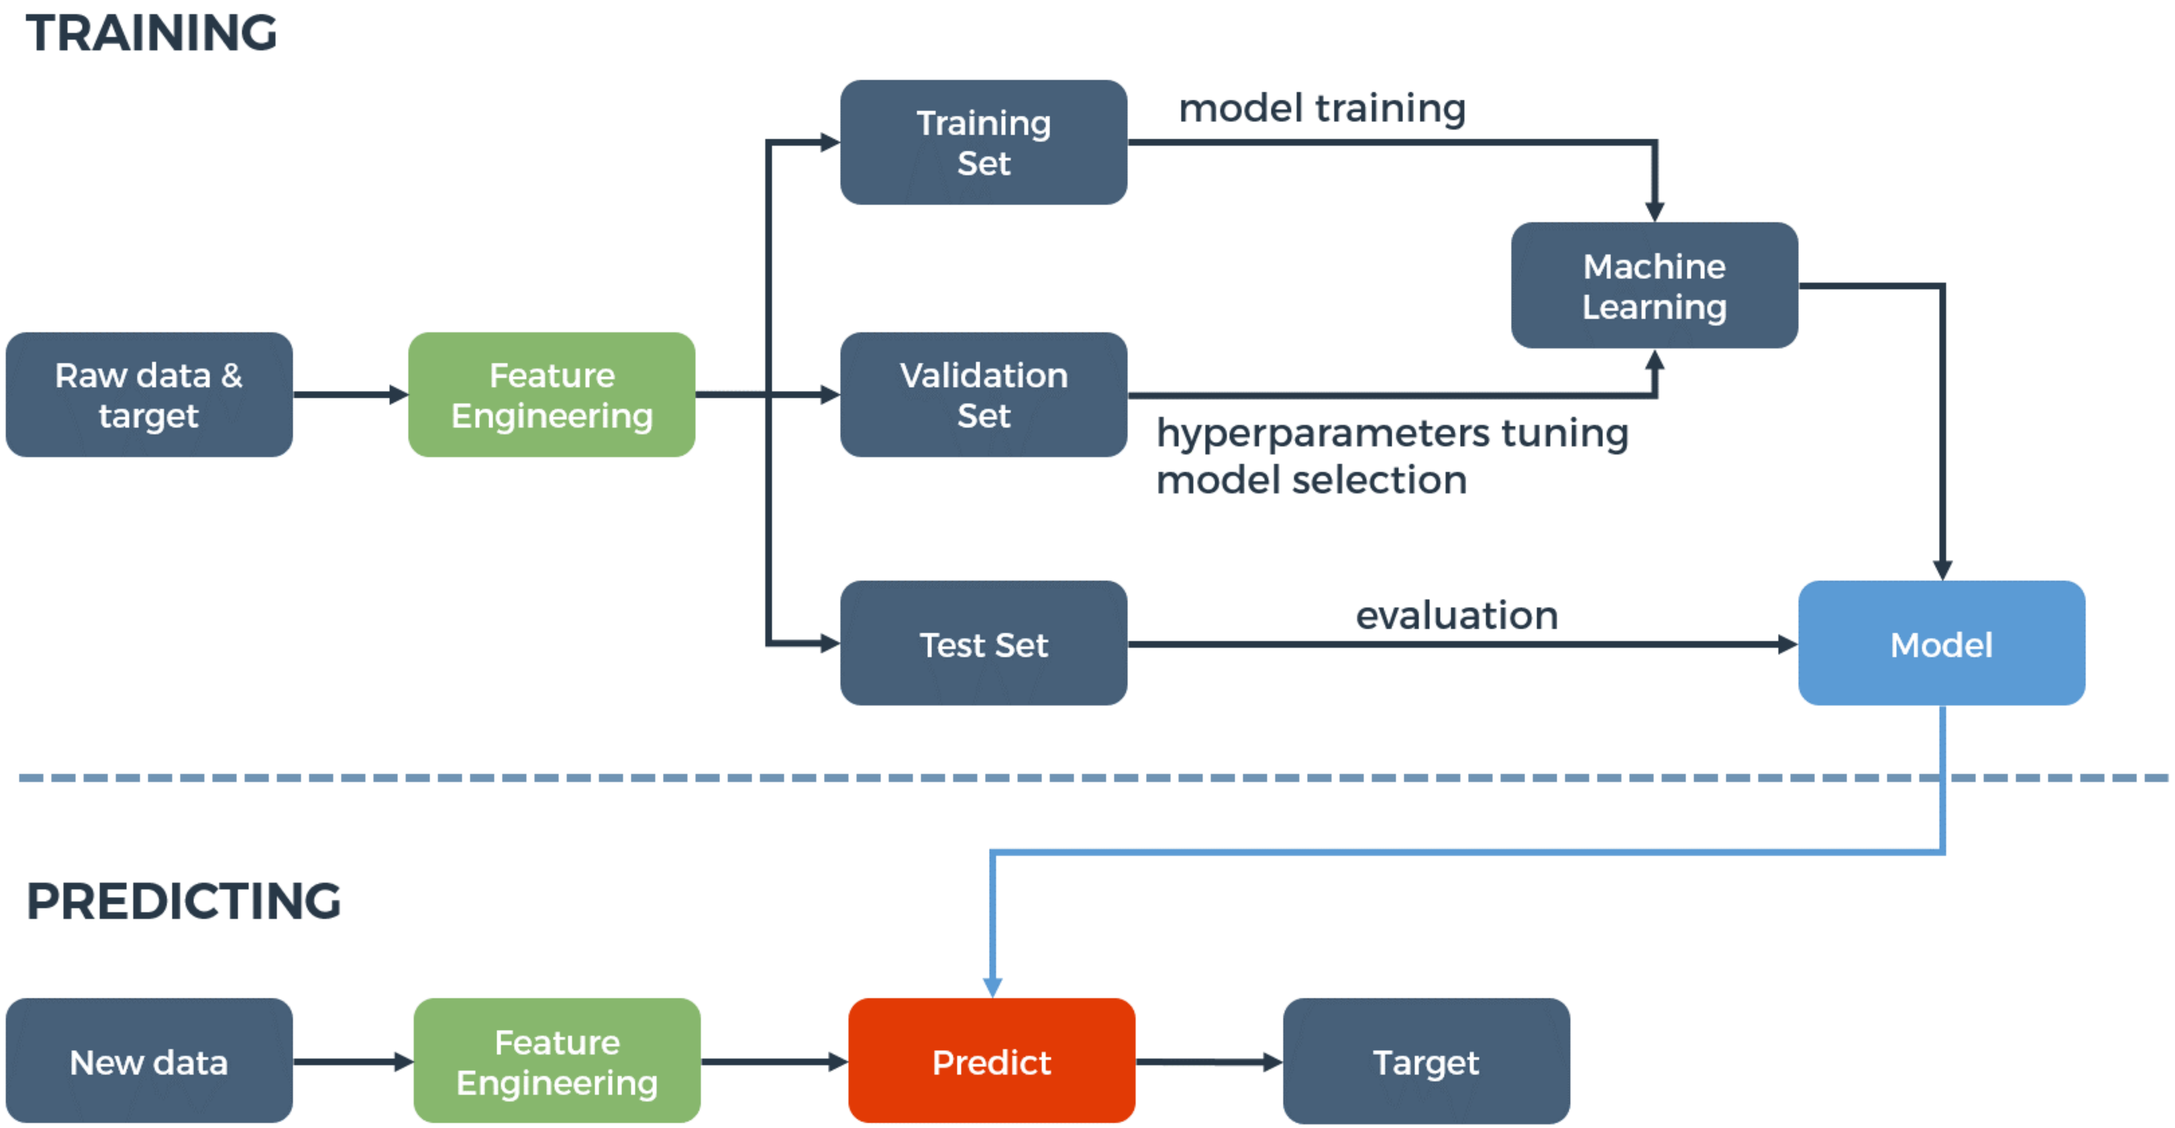

Source: hackernoon.com

Open Datasets

Datasets

Find and use datasets or complete tasks. Learn more.

Processed, balanced, well-behaved, labeled datasets to benchmark your networks!

Help the community by creating and solving Tasks on datasets!

Search 29,853 datasets | Feedback | Filter

PUBLIC | Sort by: Hottest

Hotel booking demand — 270
Jesse Mostipak
19 days | 1 MB | 10.0 | 1 File (CSV) | 1 Task

Big Five Personality Test — 134
Bojan Tunguz
14 days | 159 MB | 9.7 | 3 Files (CSV, other)

StartUp Investments (Crunchbase) — 92
Andy_M
14 days | 3 MB | 8.8 | 1 File (CSV)

Open Tasks

Can we predict the possibility of a bo...
0 Submissions · In Hotel booking demand

Visualize US Accidents Dataset
12 Submissions · In US Accidents (3.0 million...

What to watch on Netflix ?
4 Submissions · In Netflix Movies and TV Sh...

https://www.tensorflow.org/datasets
https://www.kaggle.com/datasets
http://topepo.github.io/caret/data-sets.html
https://github.com/awesomedata/awesome-public-datasets

Take-home messages

In a neuron:
… the main job is to calculate a **weighted average**
… the **decision** is made through the **activation** function

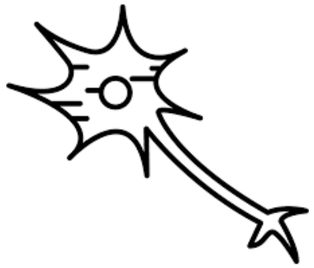

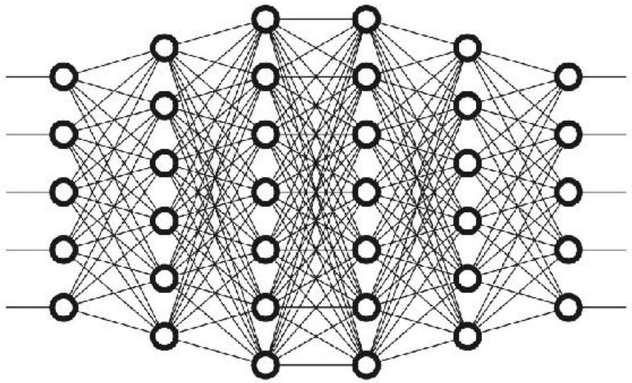

In a neural network:
… losses are calculated using the loss function
… losses are calculated by **comparing** the truths and the prediction
… **predictions** are made through **forward** propagation
… weights are **updated** through the **backward** propagation process
… **optimizers** are used to decide the weights updating **strategies**

In a deep learning workflow:
… the heavy lifting is mostly done by **DL frameworks**
… open datasets are crucial for benchmarking and bootstrapping DNNs