Introduction to Deep Learning

About me

Intelligent Systems
2008 - 2014

Universität Bielefeld

Research Engineer
2014 - 2018

BOSCH
Robert Bosch GmbH

Data Scientist
since 2018

codecentric

M.Sc. Denis Stalz-John

Contents

What is Deep Learning?

Artificial Intelligence

Any technique that enables computers to mimic human behavior

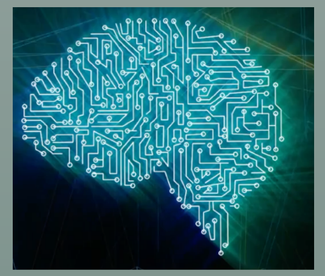

Machine Learning

Ability to learn rules from data

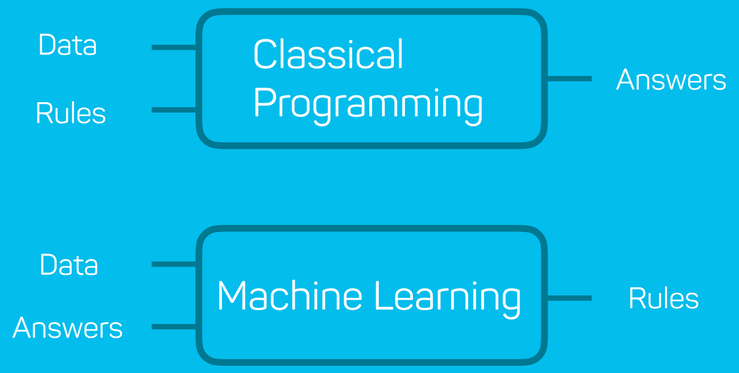

Deep Learning

Apply machine learning with „deep“ neural networks

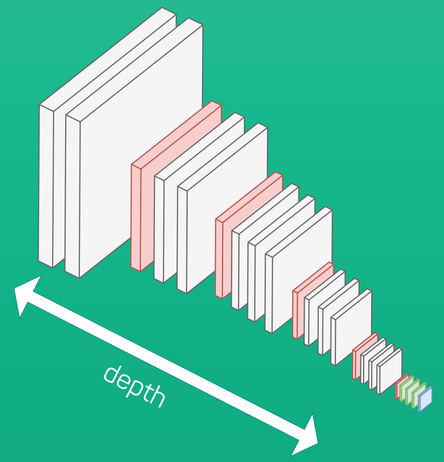

Deep Learning Achievements - ImageNet

IMAGENET

- 1000 categories
- train set: 1.2 M
- test set: 100 k

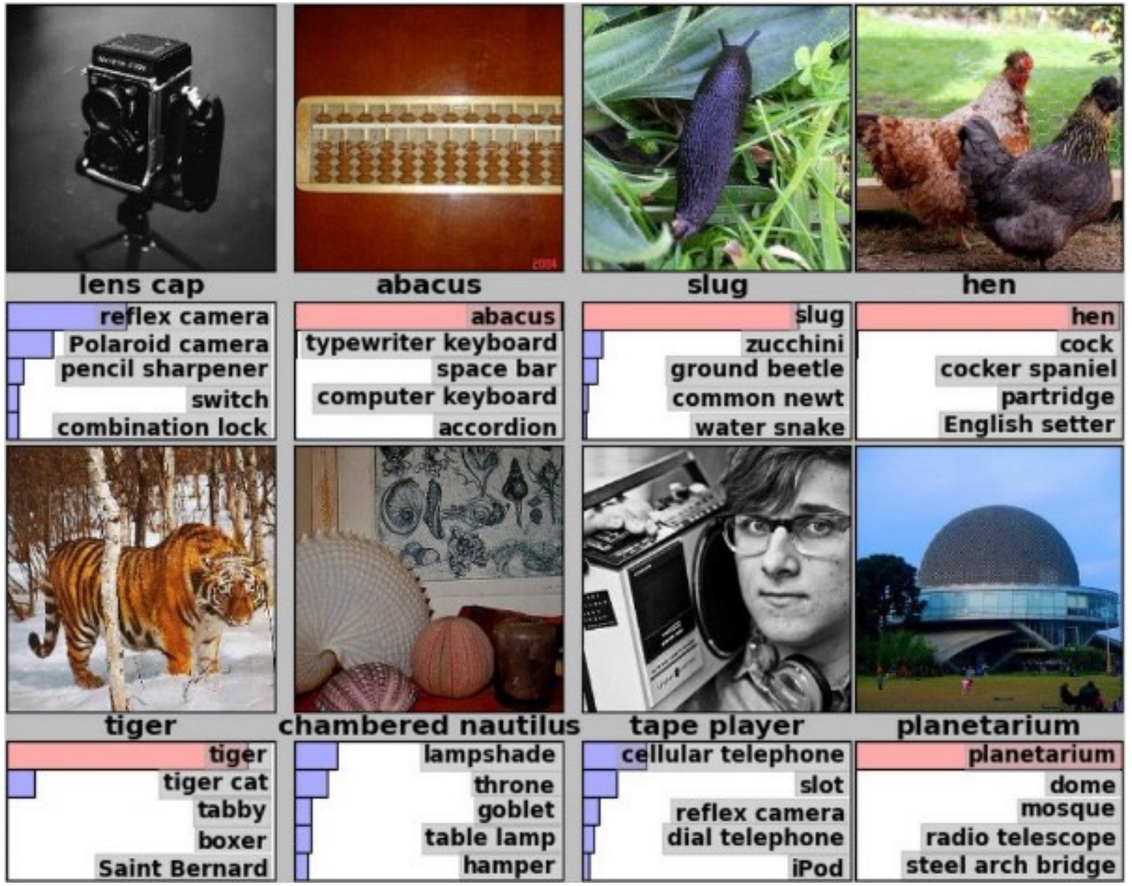

https://medium.com/obvious-ventures/our-investment-in-darwinai-d5ea1a7af32e

Deep Learning Achievements - ImageNet

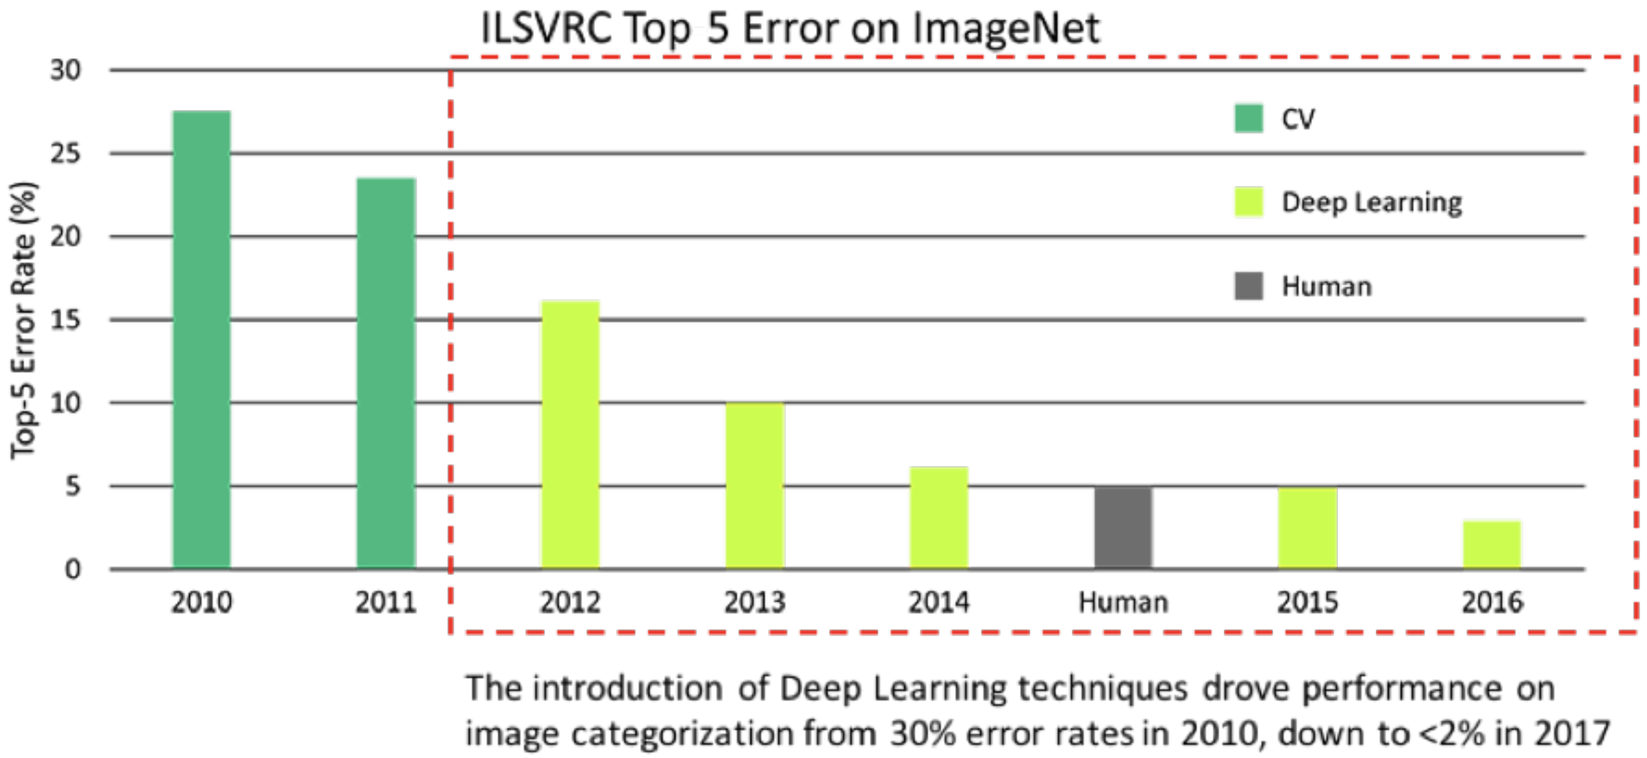

The introduction of Deep Learning techniques drove performance on image categorization from 30% error rates in 2010, down to <2% in 2017

https://medium.com/obvious-ventures/our-investment-in-darwinai-d5ea1a7af32e

Deep Learning Achievements - Autonomous Driving

Example of Teslas Autonomous Driving Pilot

Deep Learning Achievements - Speech Recognition

2006
Microsoft

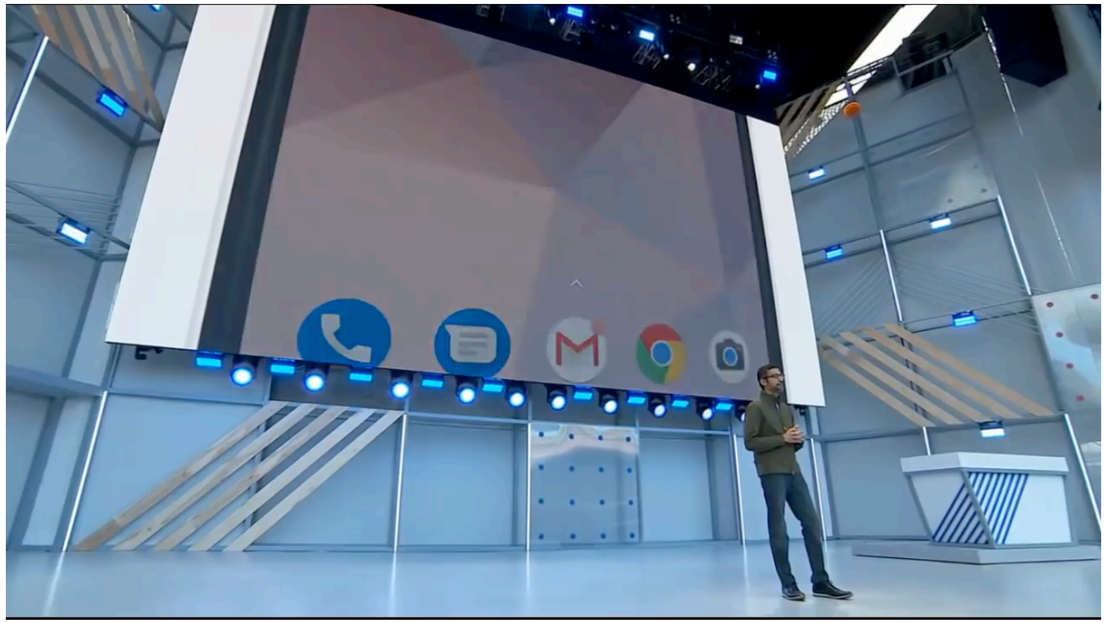

2018
Google

Deep Learning Achievements - Strategic Knowledge

Chess

Deep Blue
1997

GO

Deep Mind
2017

Dota2

OpenAI
2018

AI performs better than human

History of Deep Learning

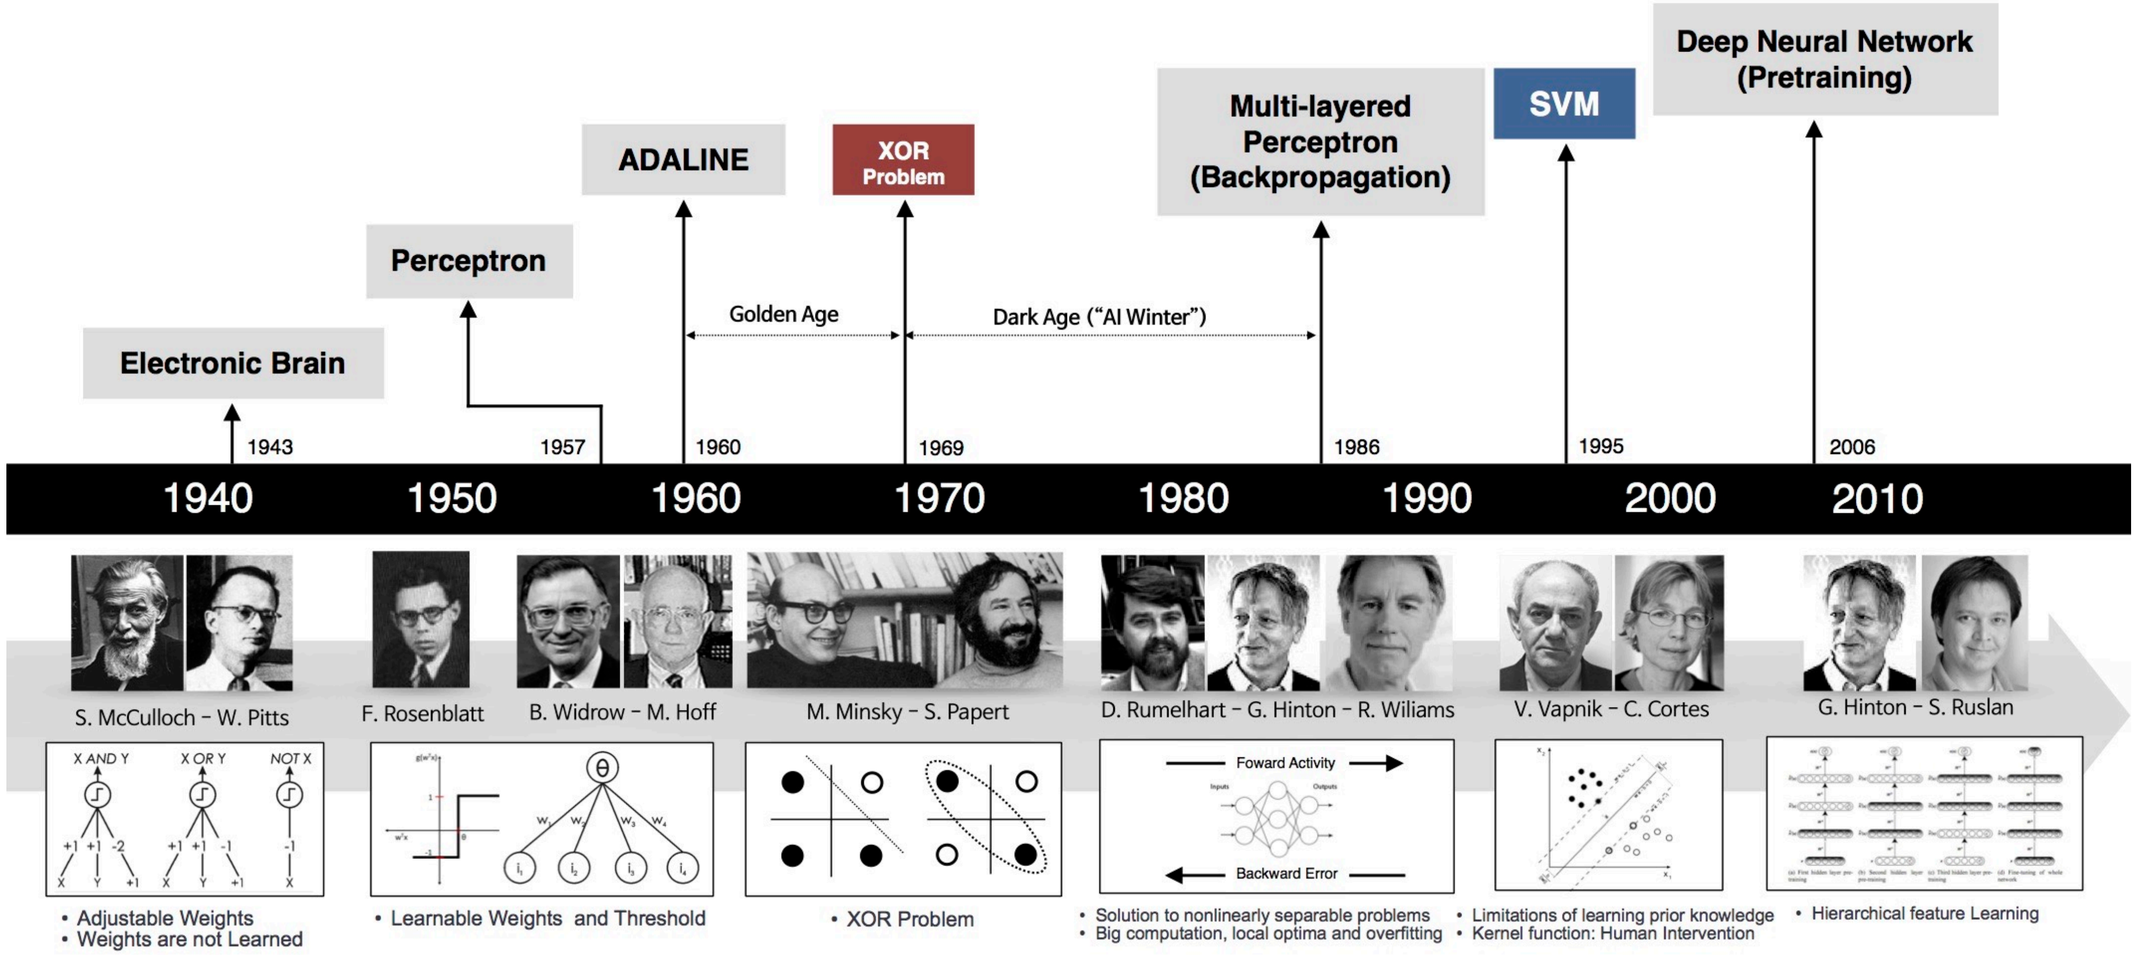

https://beamandrew.github.io/deeplearning/2017/02/23/deep_learning_101_part1.html

Fundamentals of Deep Learning - Perceptron

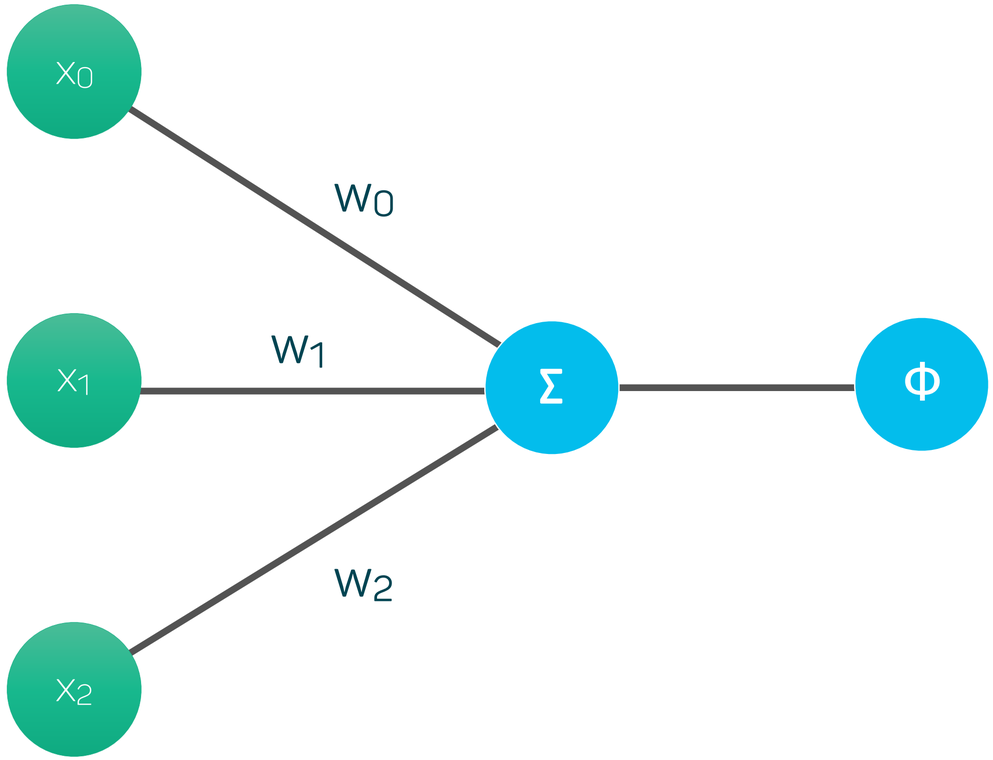

$$f(x) = \phi(\sum_i w_i x_i + b)$$

Fundamentals of Deep Learning

Activation Functions

Sigmoid

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

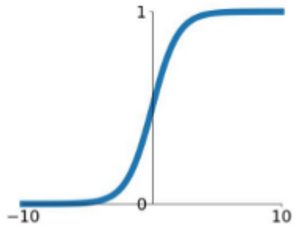

tanh

$$\tanh(x)$$

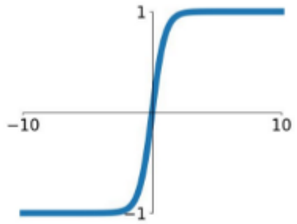

ReLU

$$\max(0, x)$$

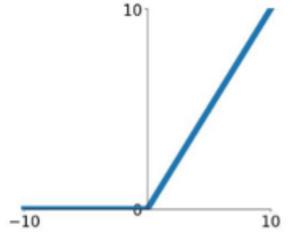

Leaky ReLU

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

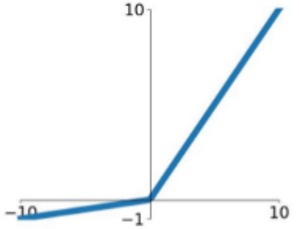

Maxout

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

ELU

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

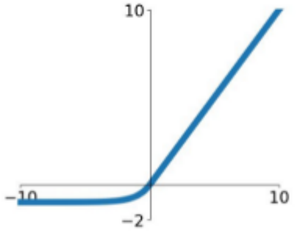

https://cdn-images-1.medium.com/max/1200/1*ZafDv3VUm60Eh10OeJu1vw.png

Fundamentals of Deep Learning

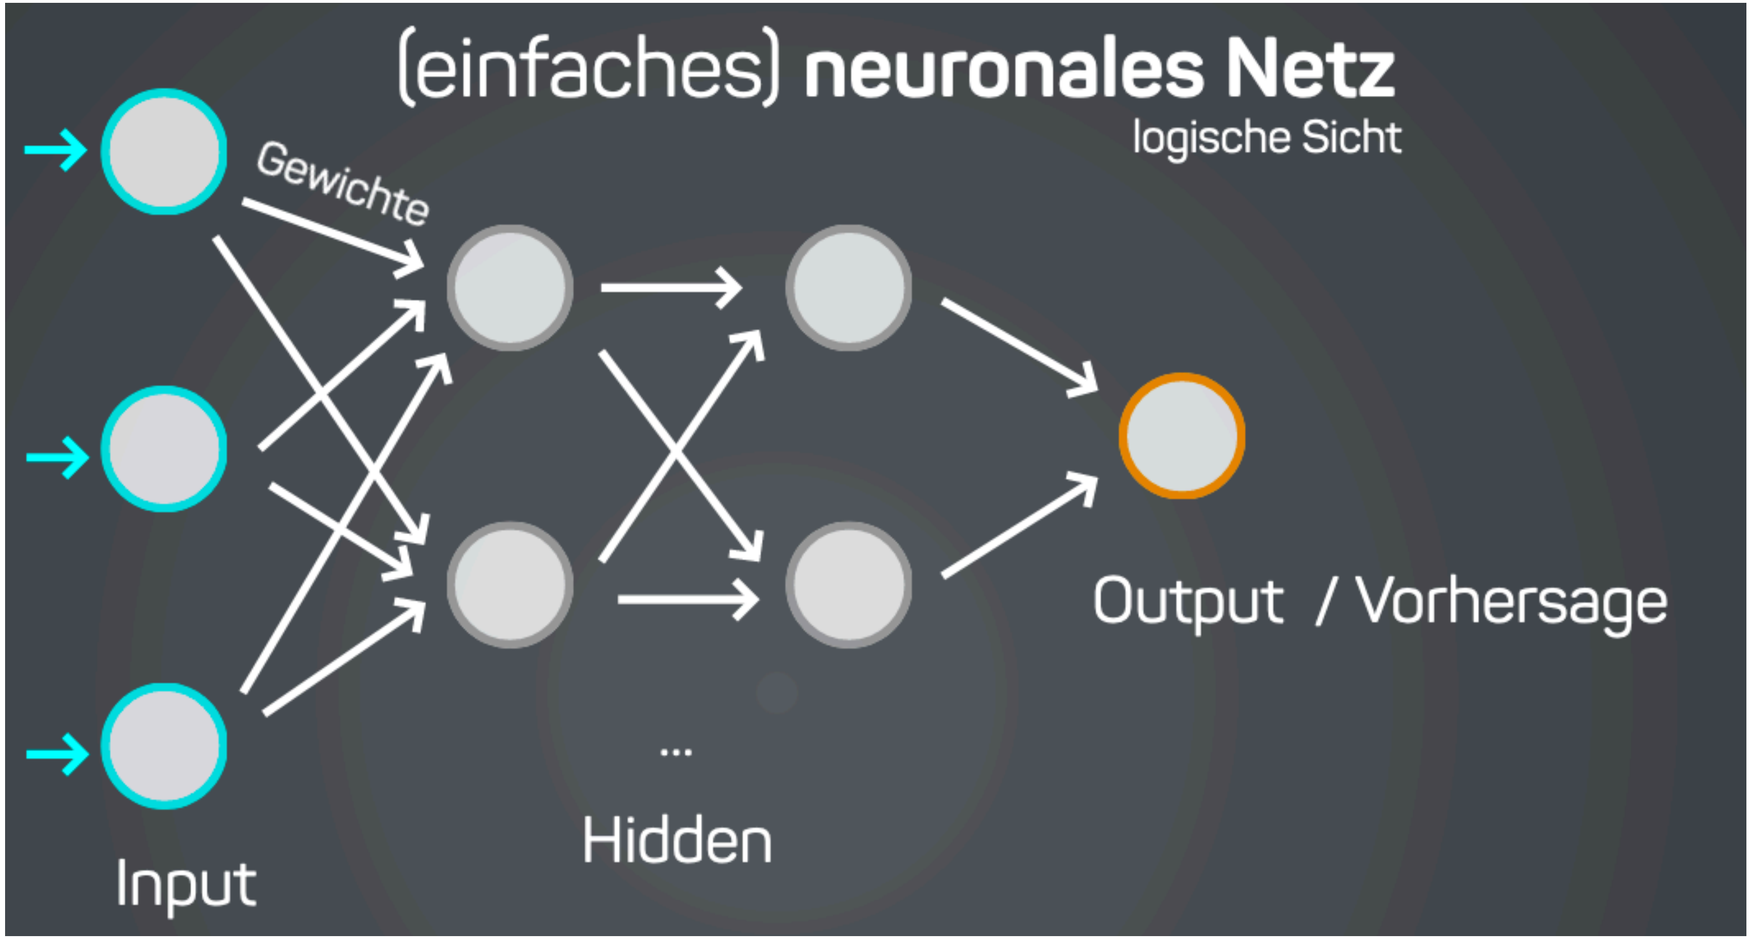

Fundamentals of Deep Learning

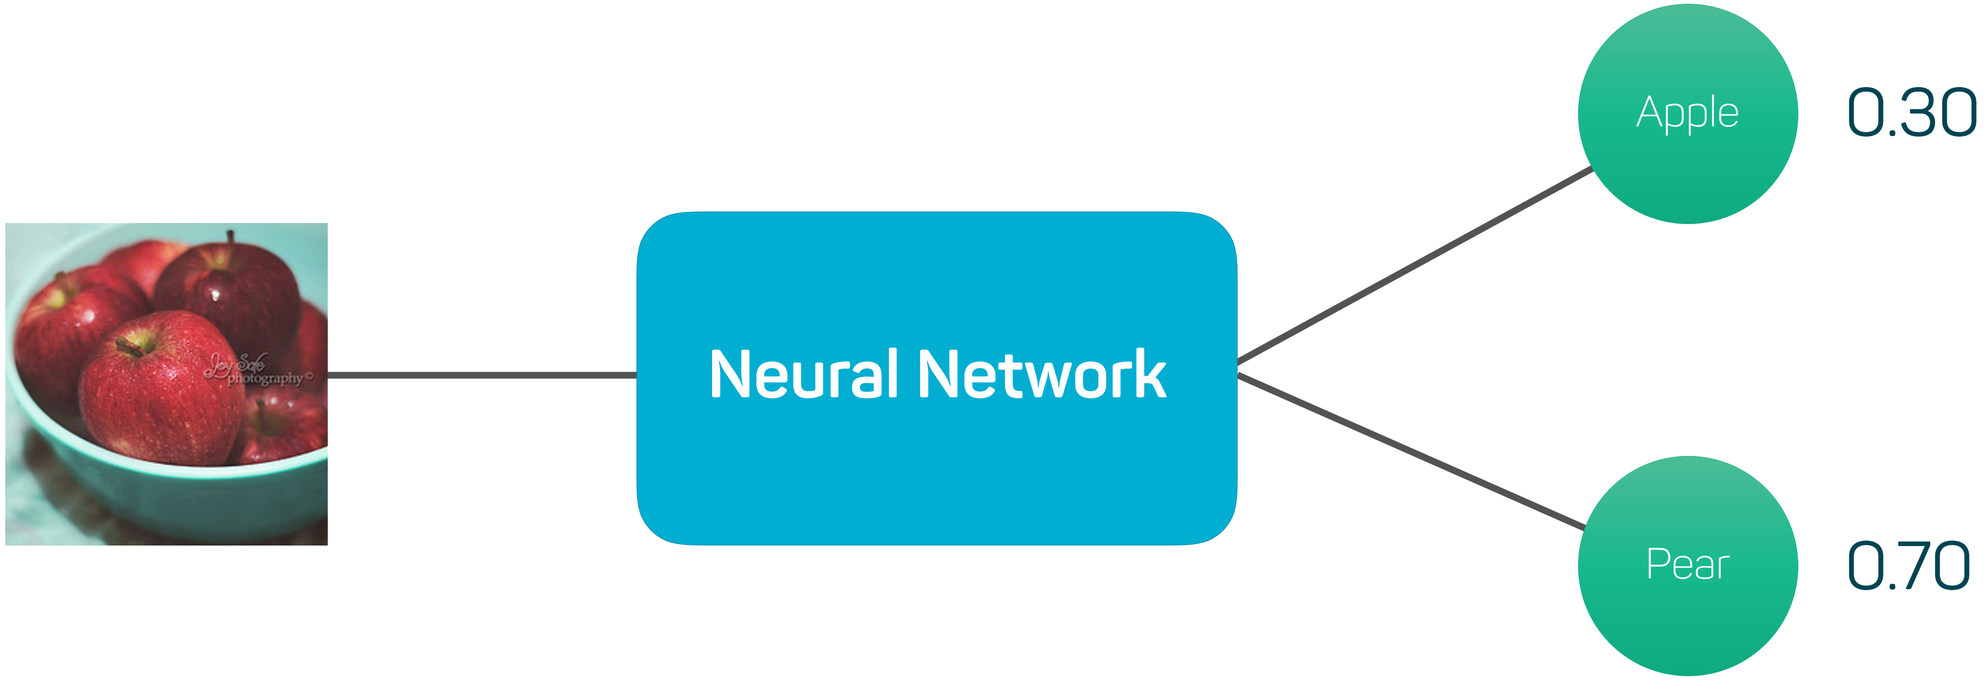

Fundamentals of Deep Learning - Loss Function

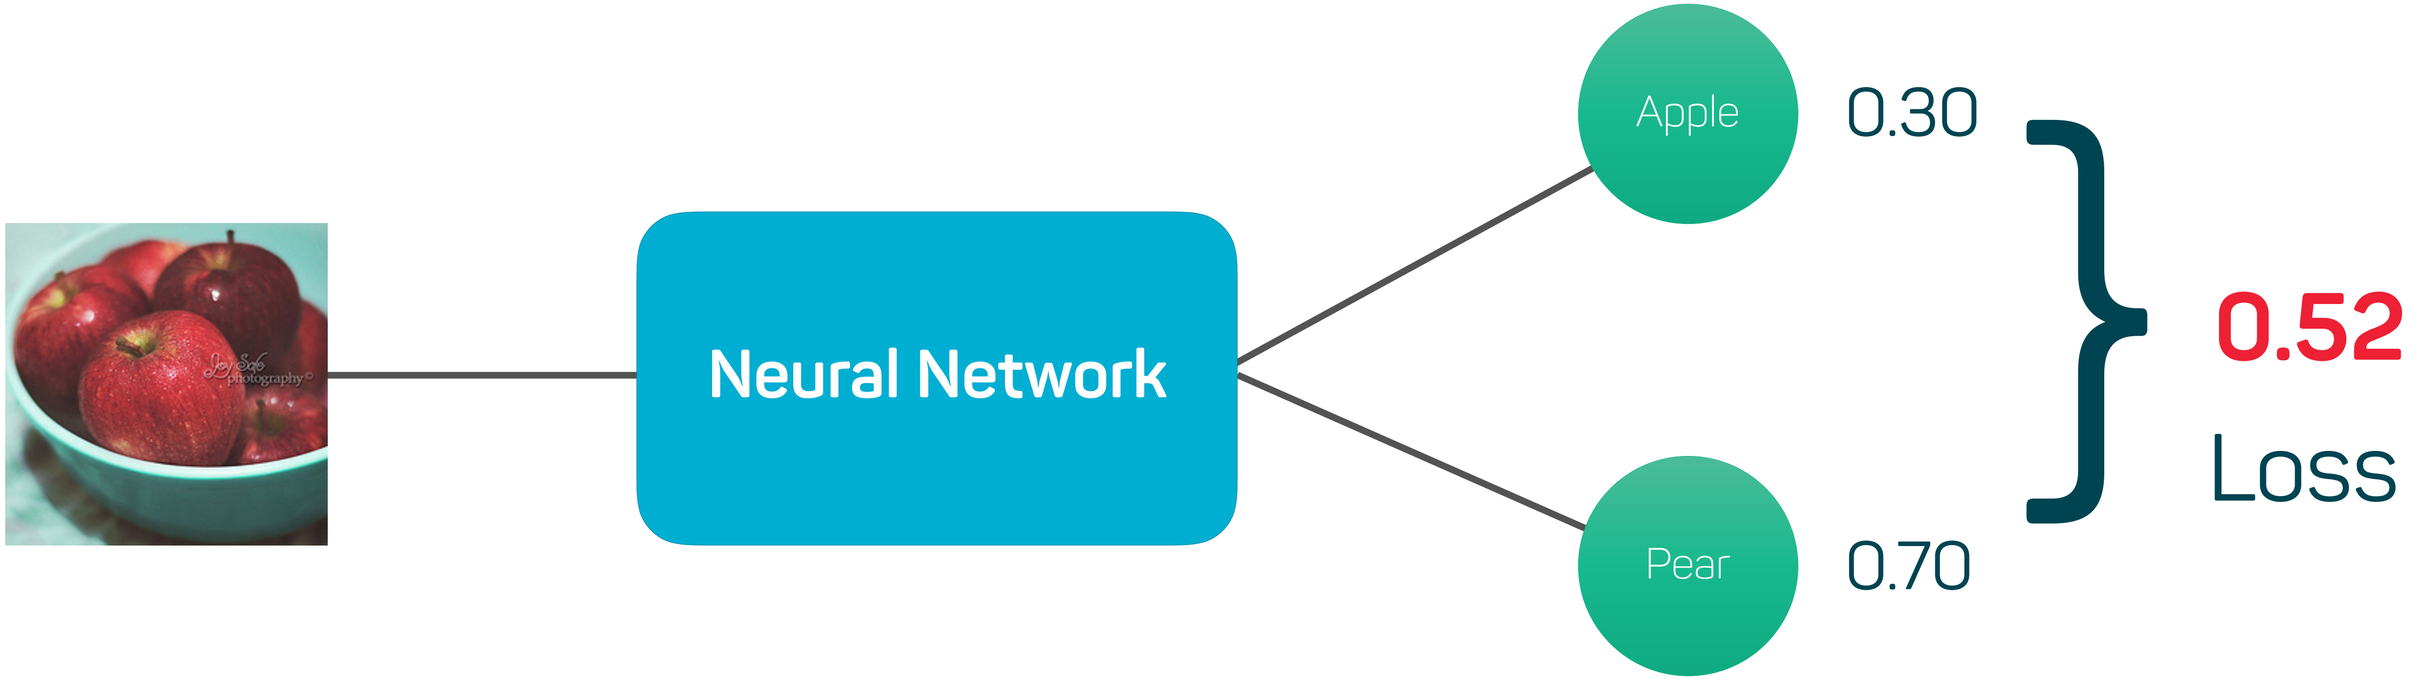

Fundamentals of Deep Learning - Loss Functions

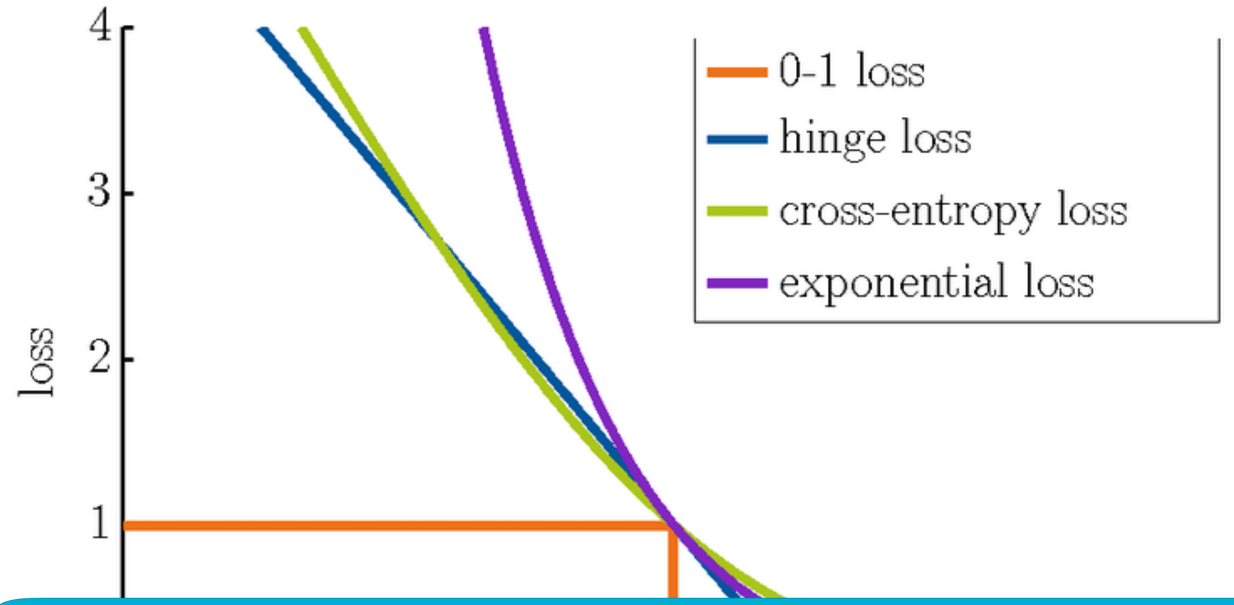

Any loss function returns a scalar value!

Fundamentals of Deep Learning

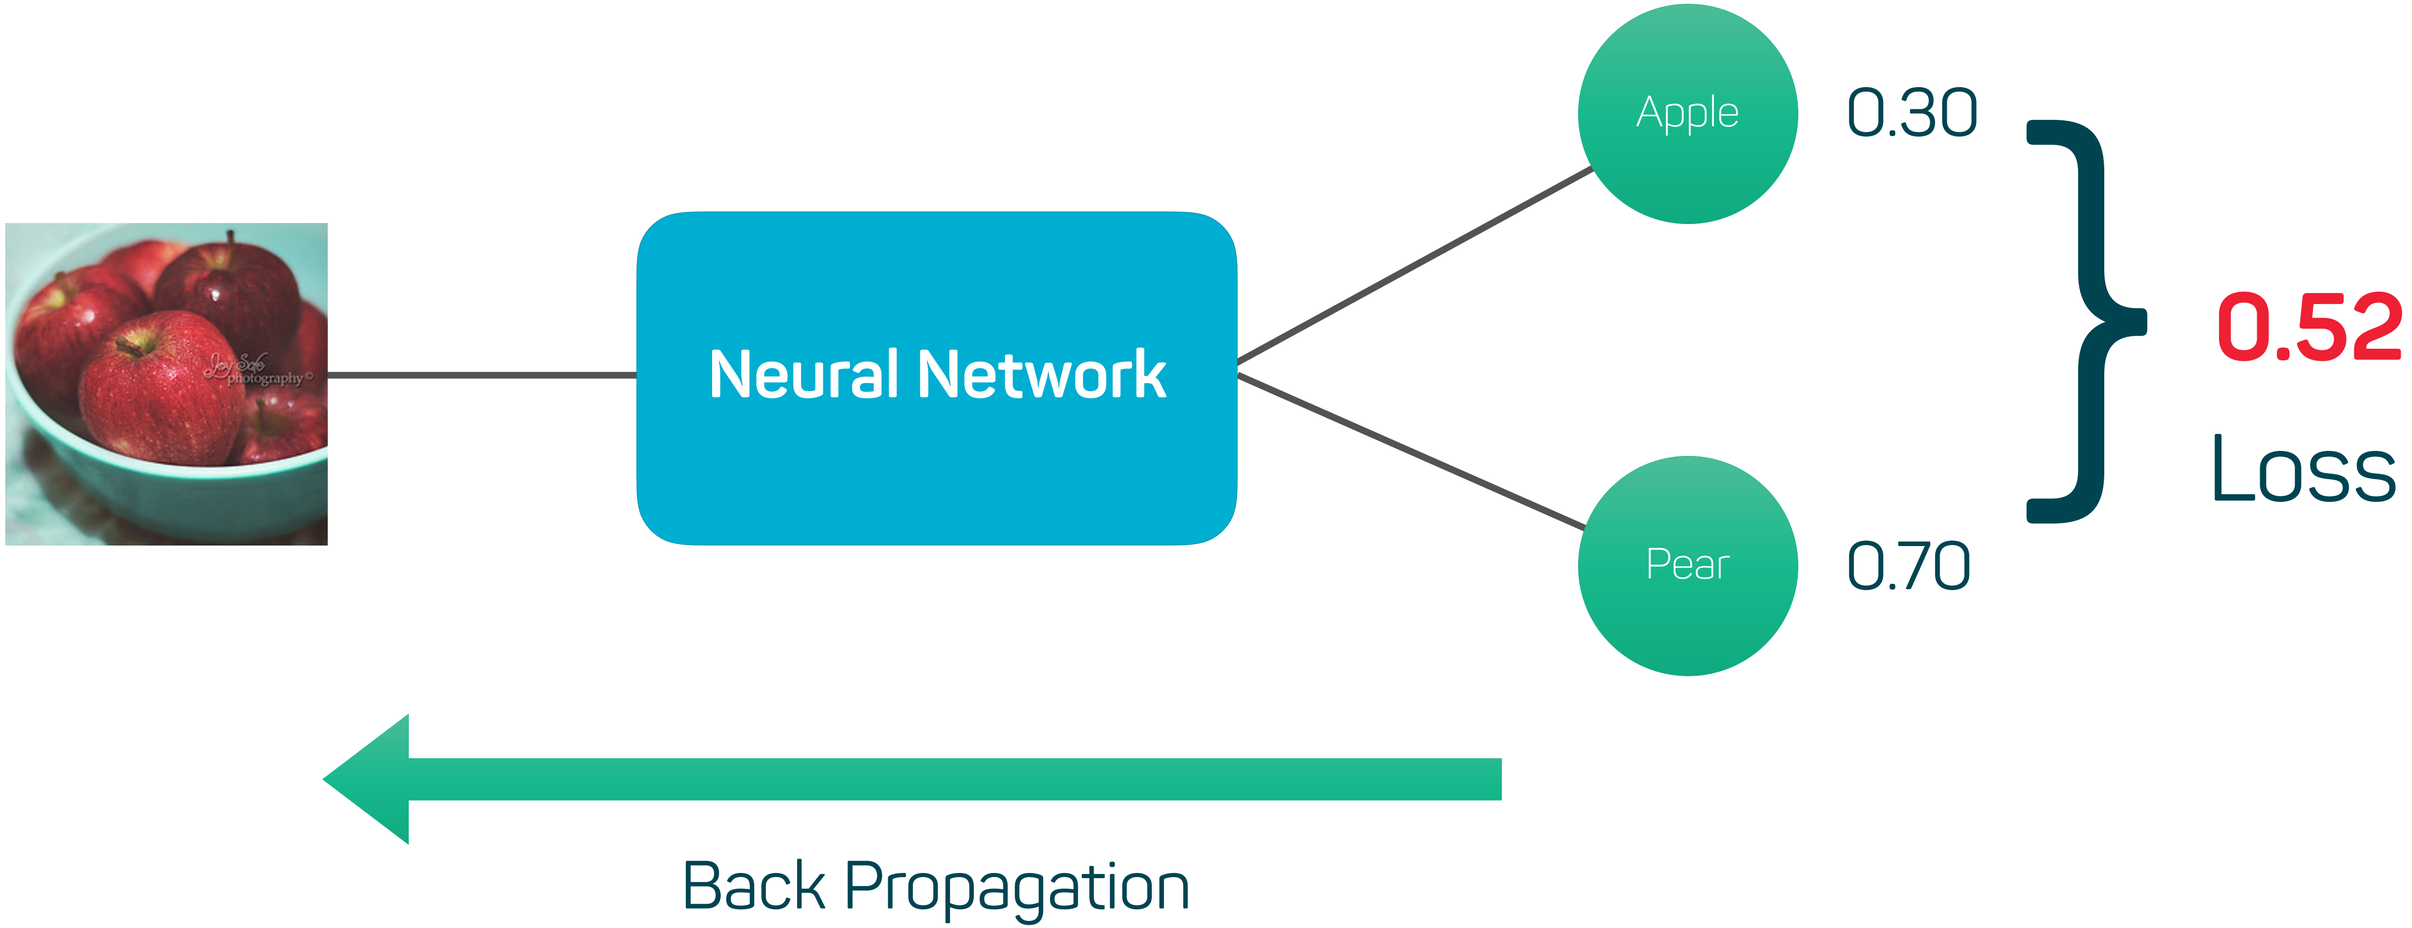

Fundamentals of Deep Learning

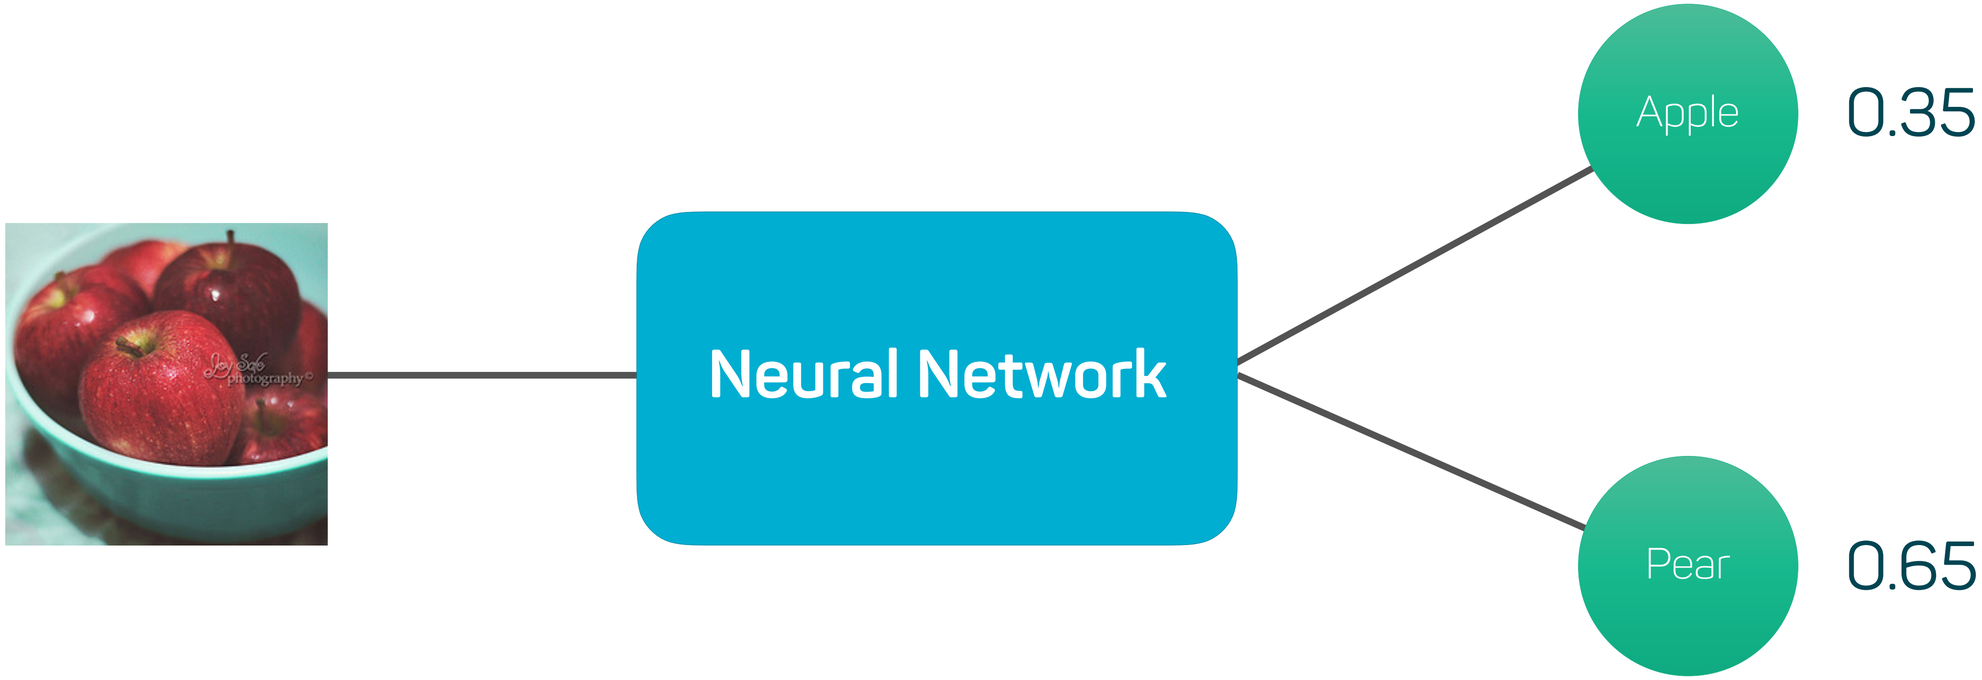

Fundamentals of Deep Learning

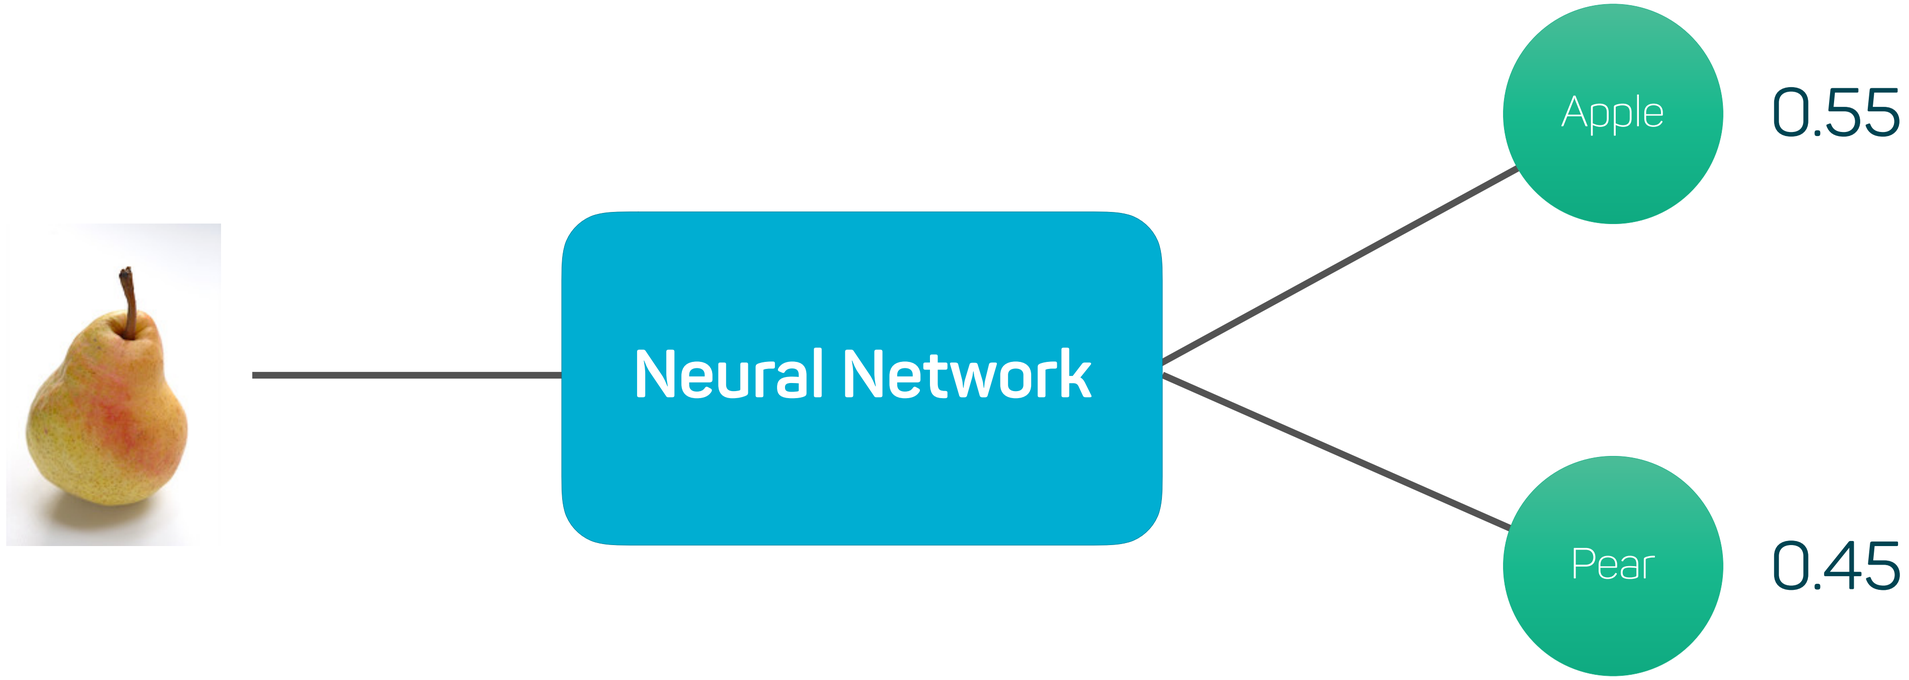

Fundamentals of Deep Learning

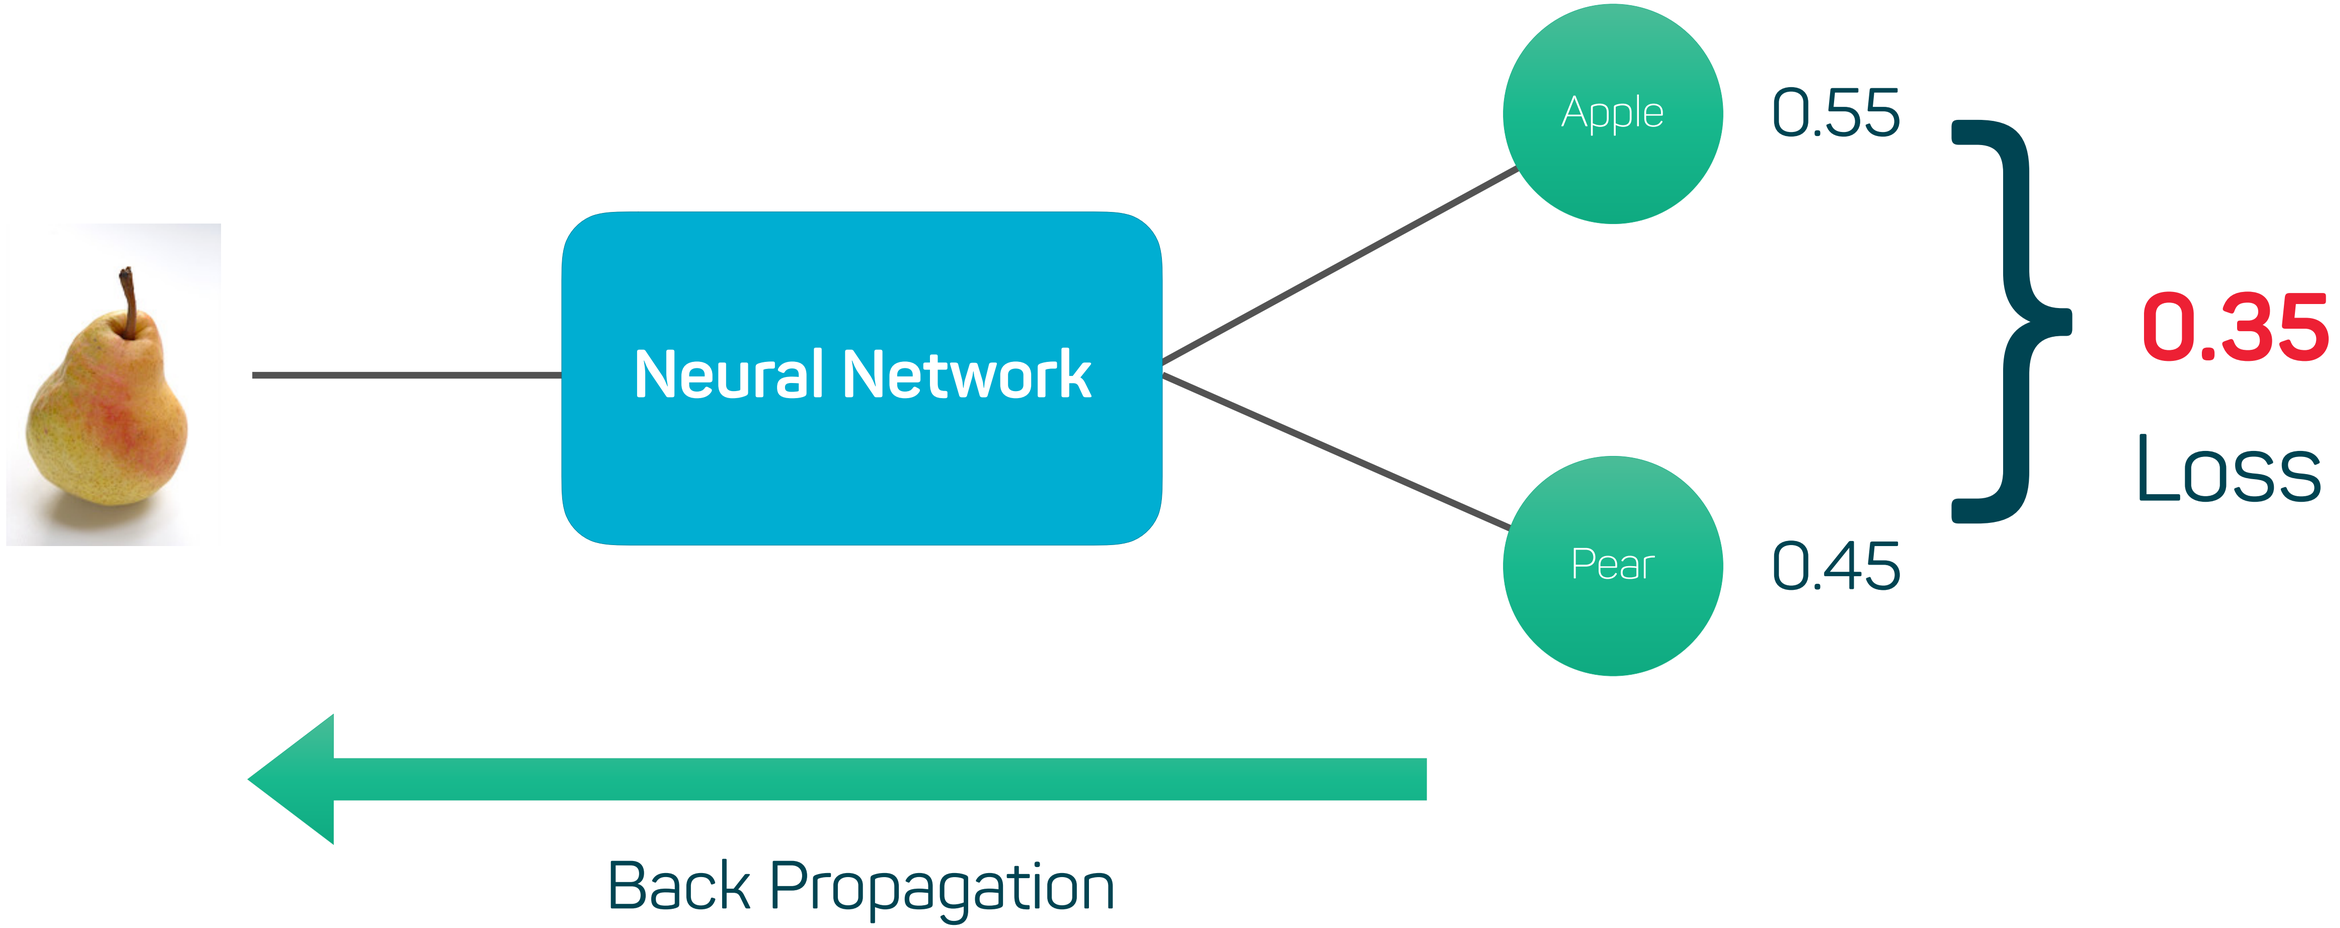

Fundamentals of Deep Learning

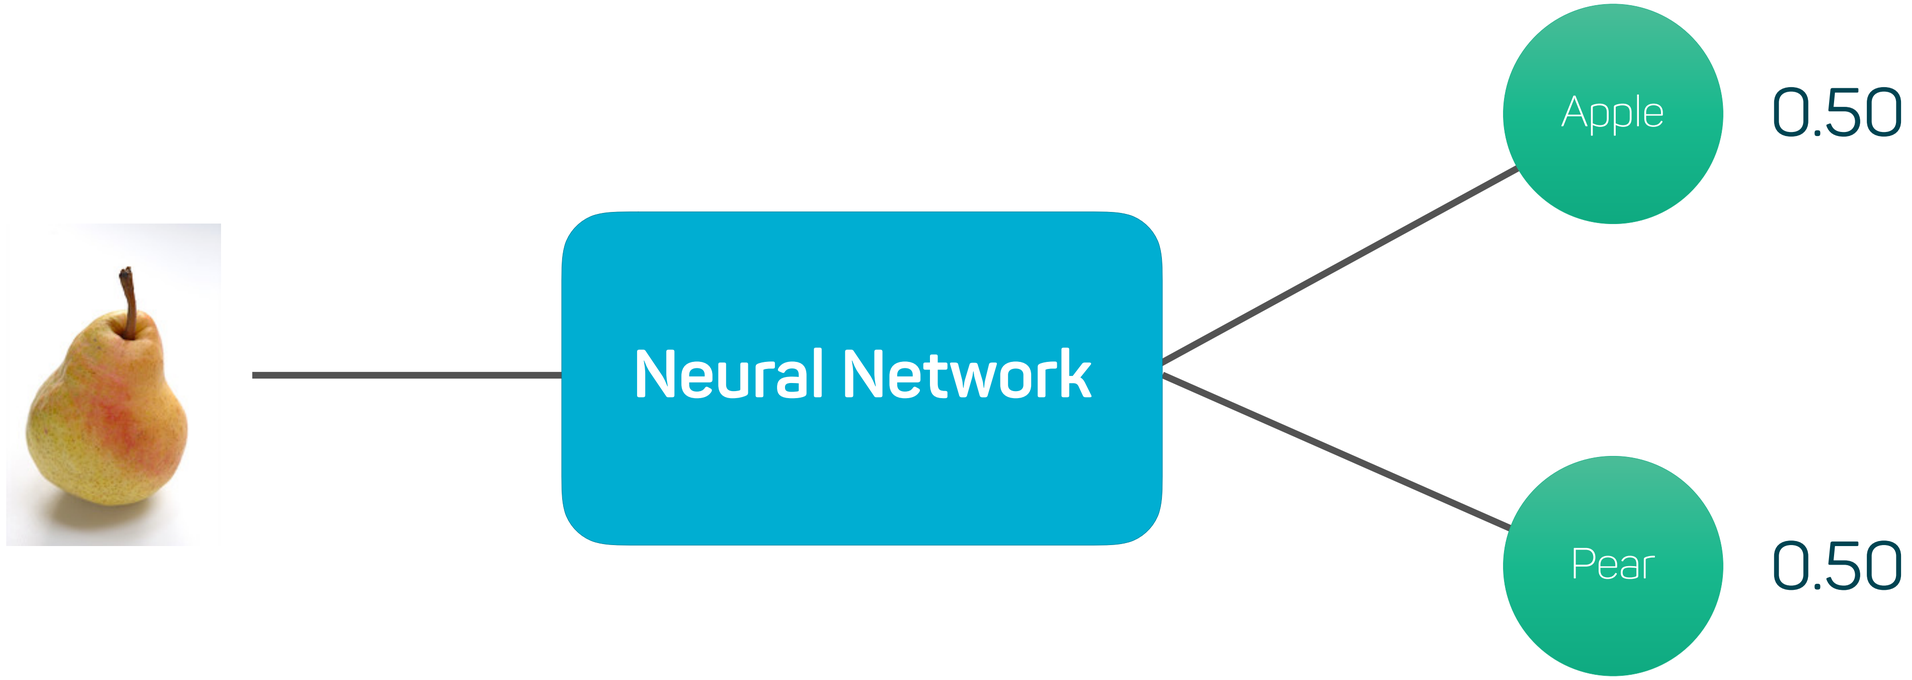

Fundamentals of Deep Learning

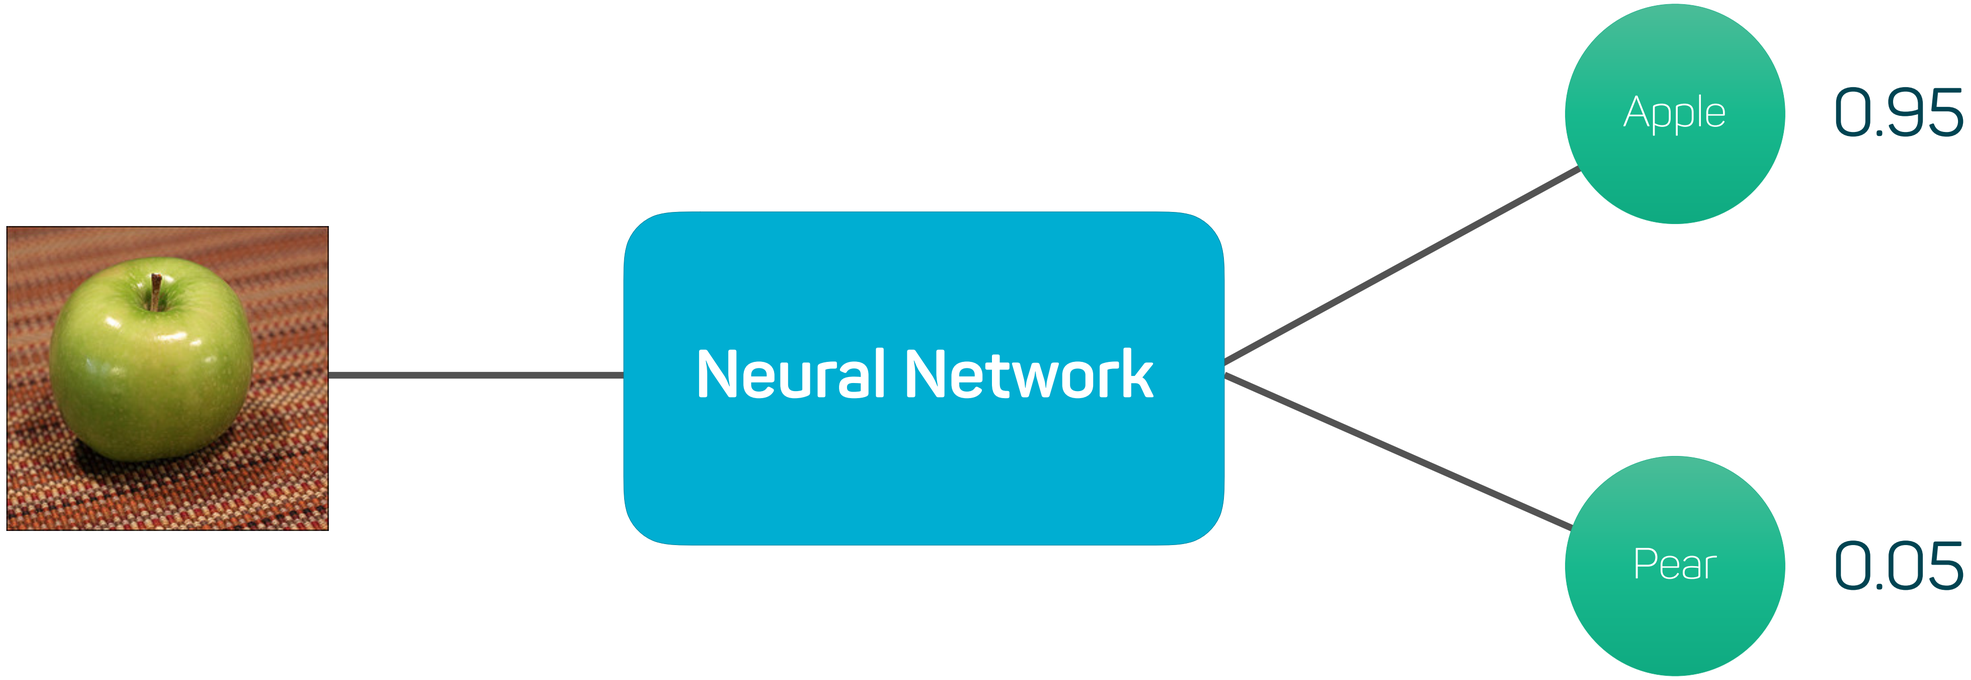

Fundamentals of Deep Learning

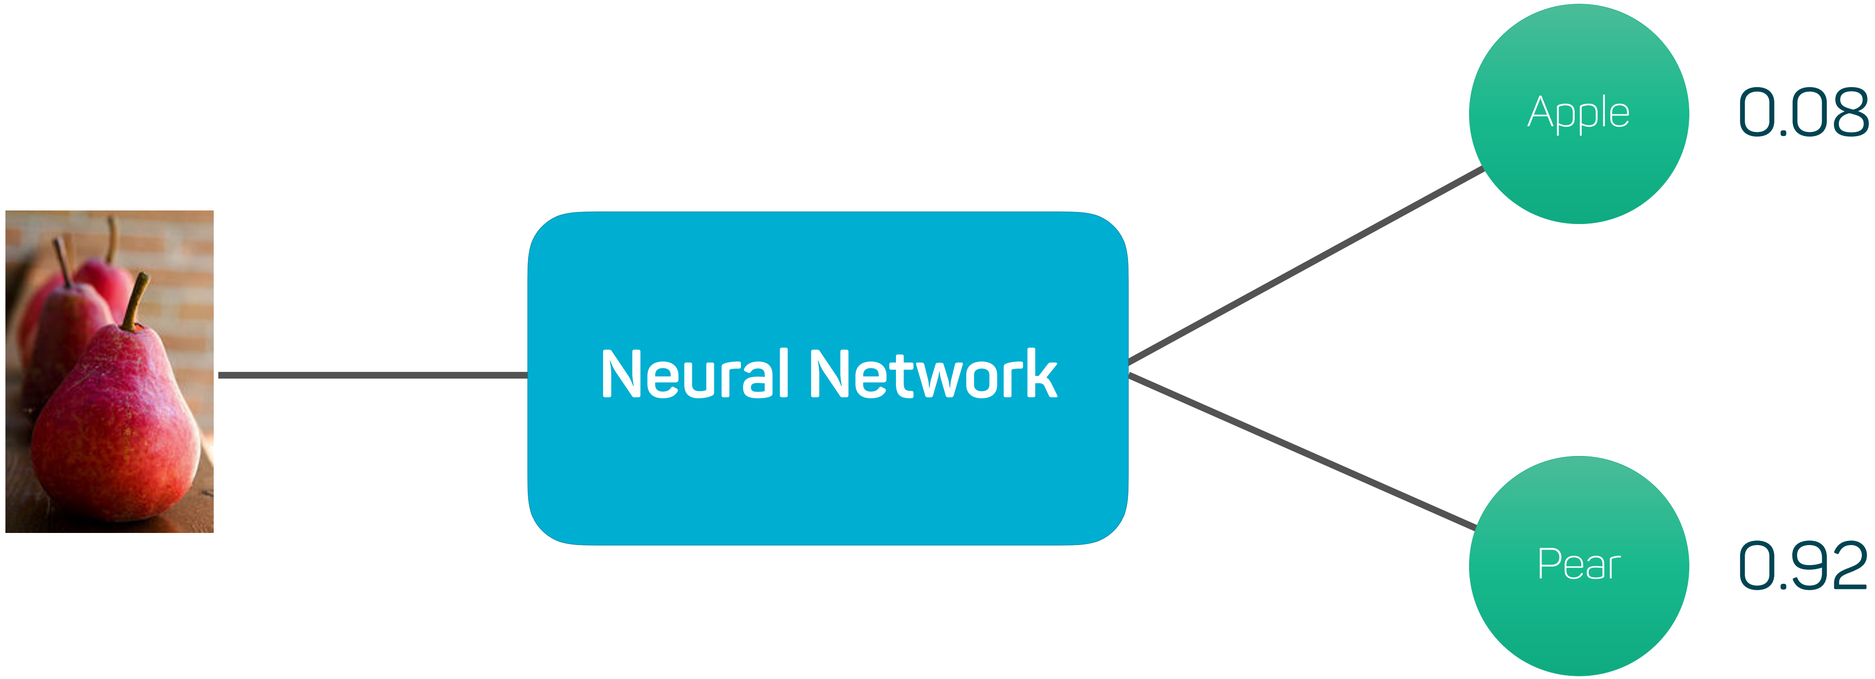

Fundamentals of Deep Learning - Overfitting

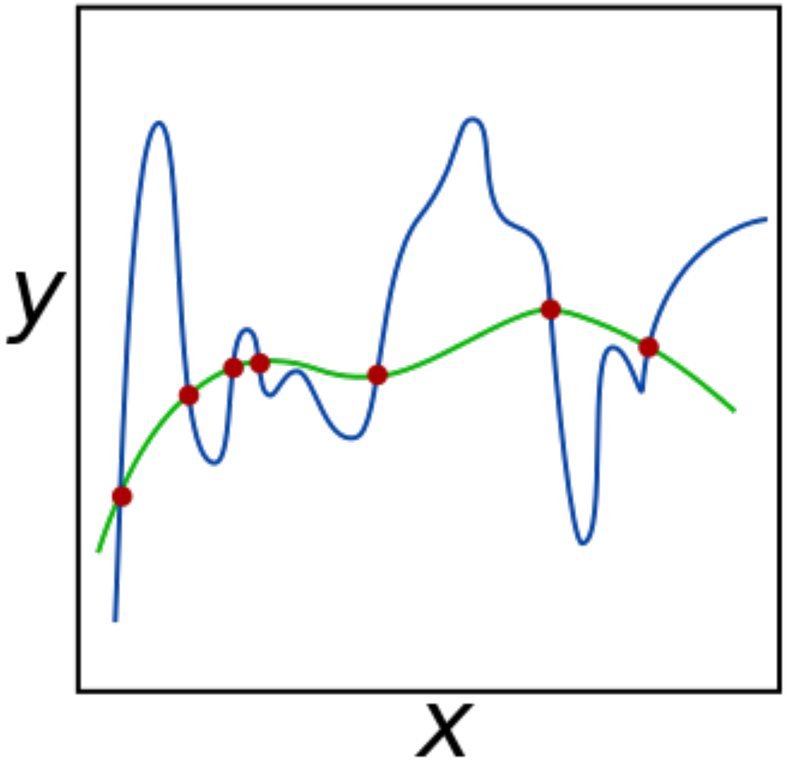

https://en.wikipedia.org/wiki/Regularization_(mathematics)

Fundamentals of Deep Learning - Overfitting

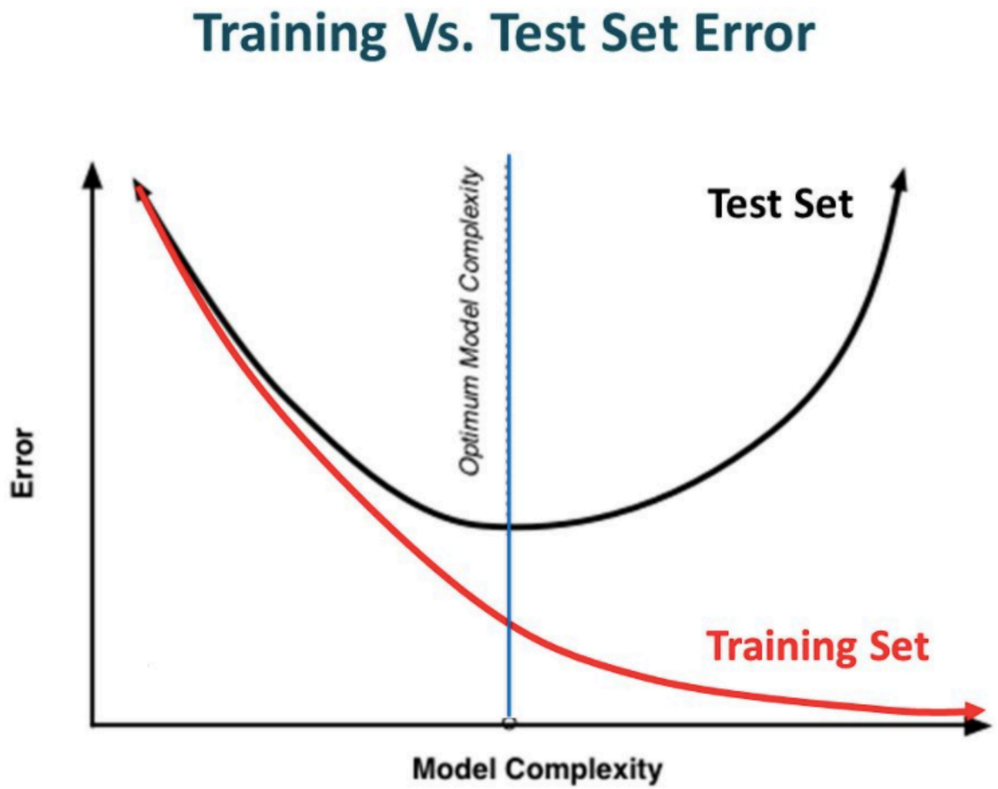

https://www.analyticsvidhya.com/blog/2018/04/fundamentals-deep-learning-regularization-techniques/

Fundamentals of Deep Learning - Data Augmentation

https://www.analyticsvidhya.com/blog/2018/04/fundamentals-deep-learning-regularization-techniques/

Fundamentals of Deep Learning - Dropout

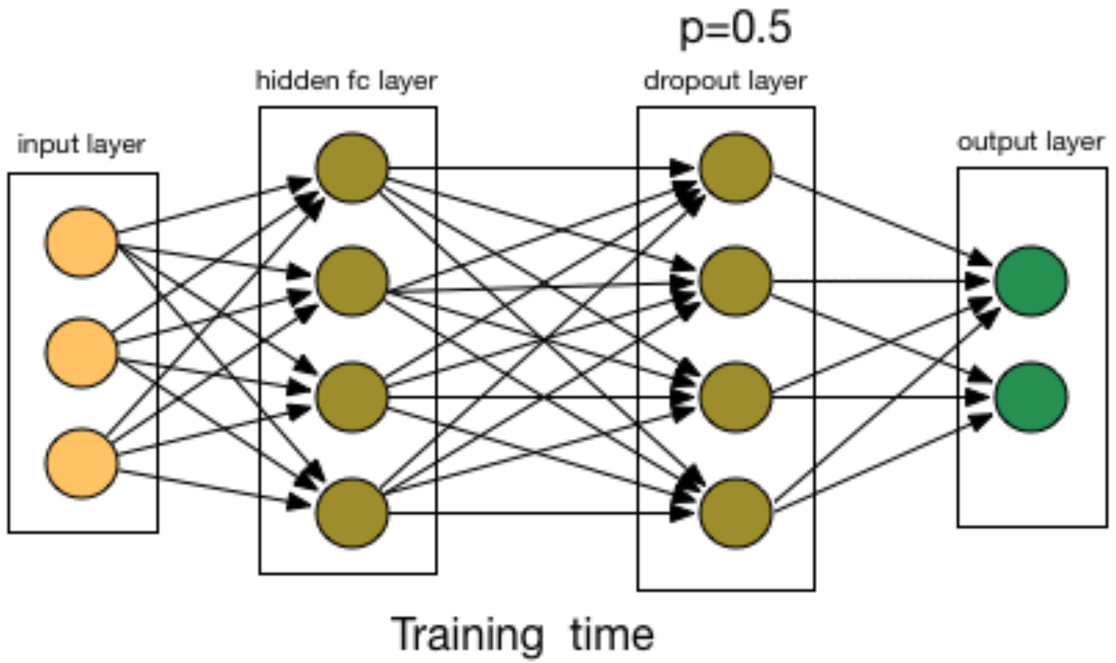

https://www.analyticsvidhya.com/blog/2018/04/fundamentals-deep-learning-regularization-techniques/

Fundamentals of Deep Learning - Early Stopping

https://www.analyticsvidhya.com/blog/2018/04/fundamentals-deep-learning-regularization-techniques/

Fundamentals of Deep Learning - Pretraining

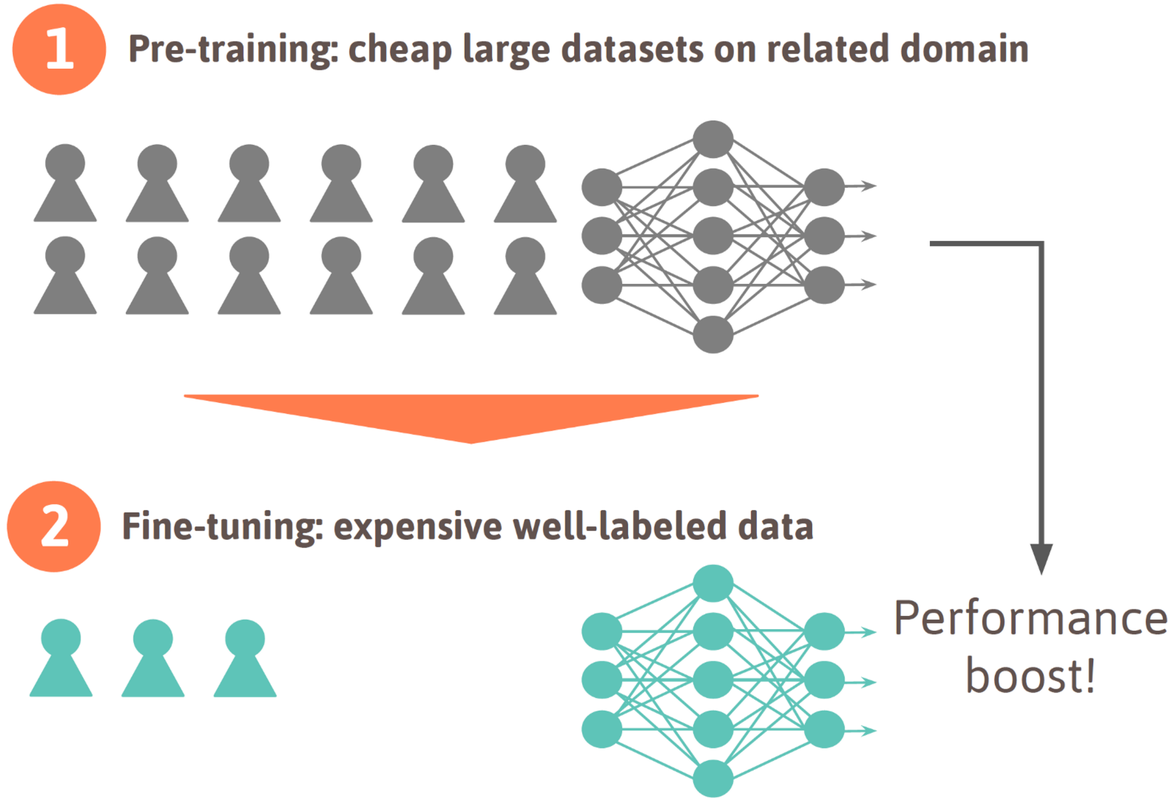

https://medium.com/merantix/applying-deep-learning-to-real-world-problems-ba2d86ac5837

Deep Learning Achievements - Pretraining

IMAGENET

- 1000 categories
- trai
- test

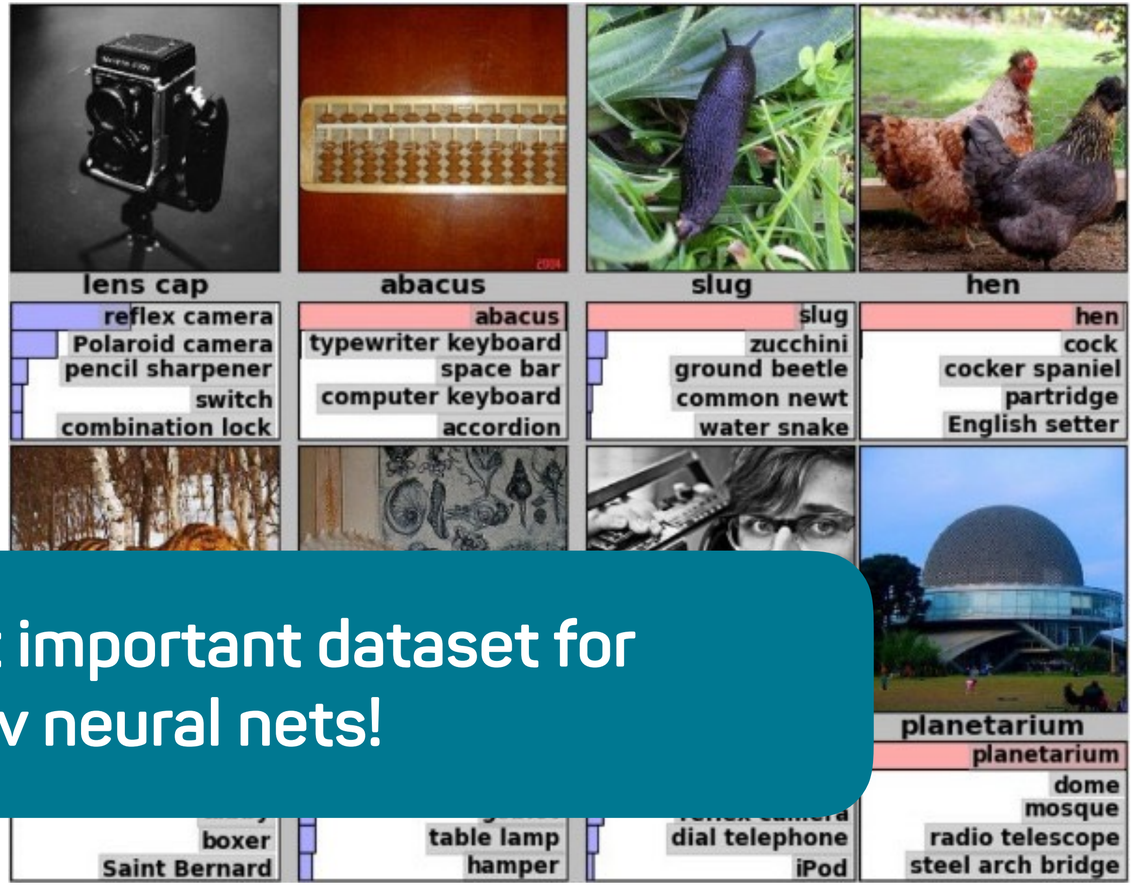

ImageNet is the most important dataset for pretraining cv neural nets!

https://medium.com/obvious-ventures/our-investment-in-darwinai-d5ea1a7af32e

Practical Deep Learning - Hardware

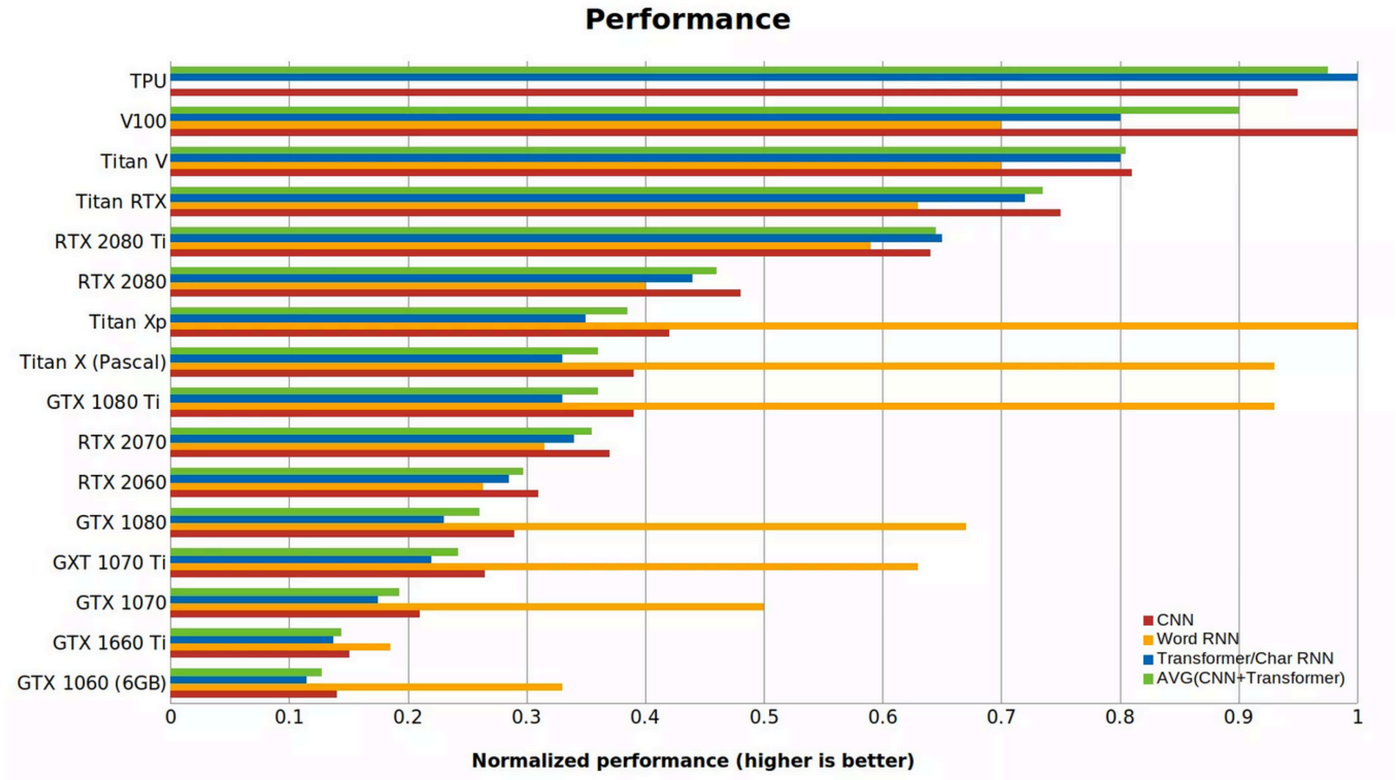

https://timdettmers.com/2019/04/03/which-gpu-for-deep-learning/

Practical Deep Learning - Software

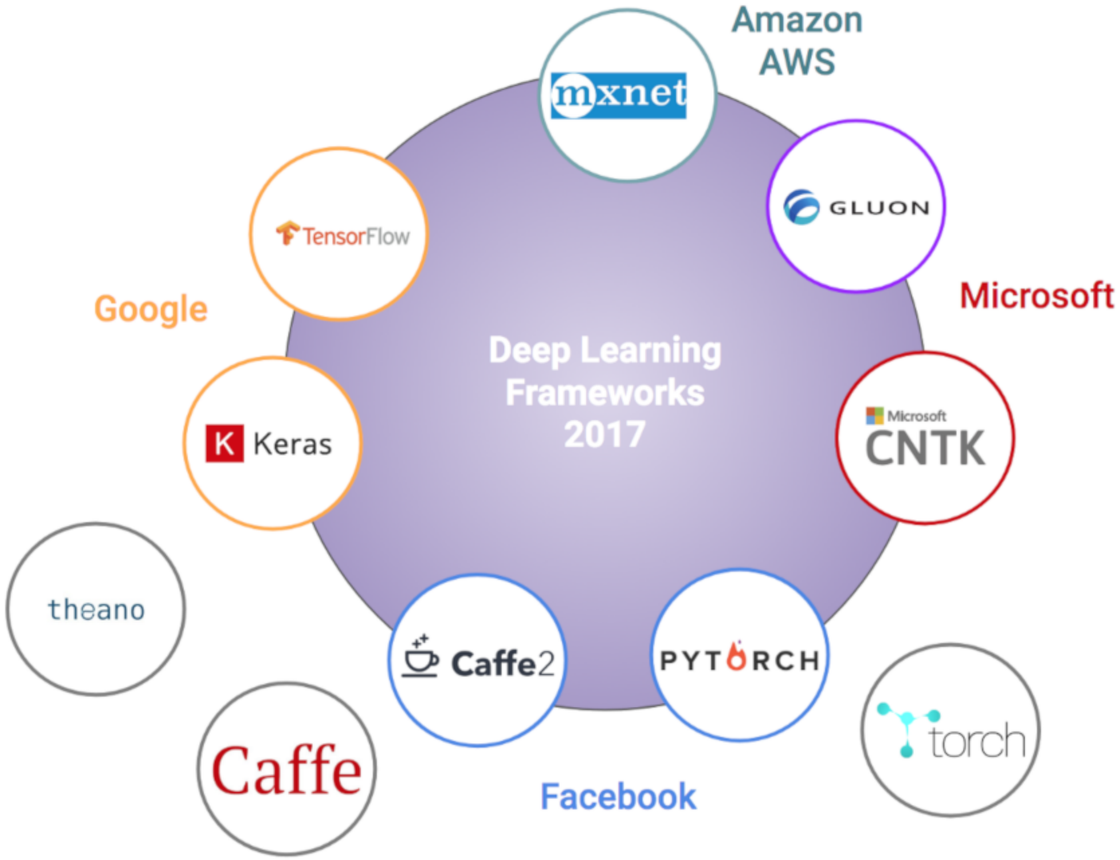

https://devopedia.org/deep-learning-frameworks

Practical Deep Learning - Software

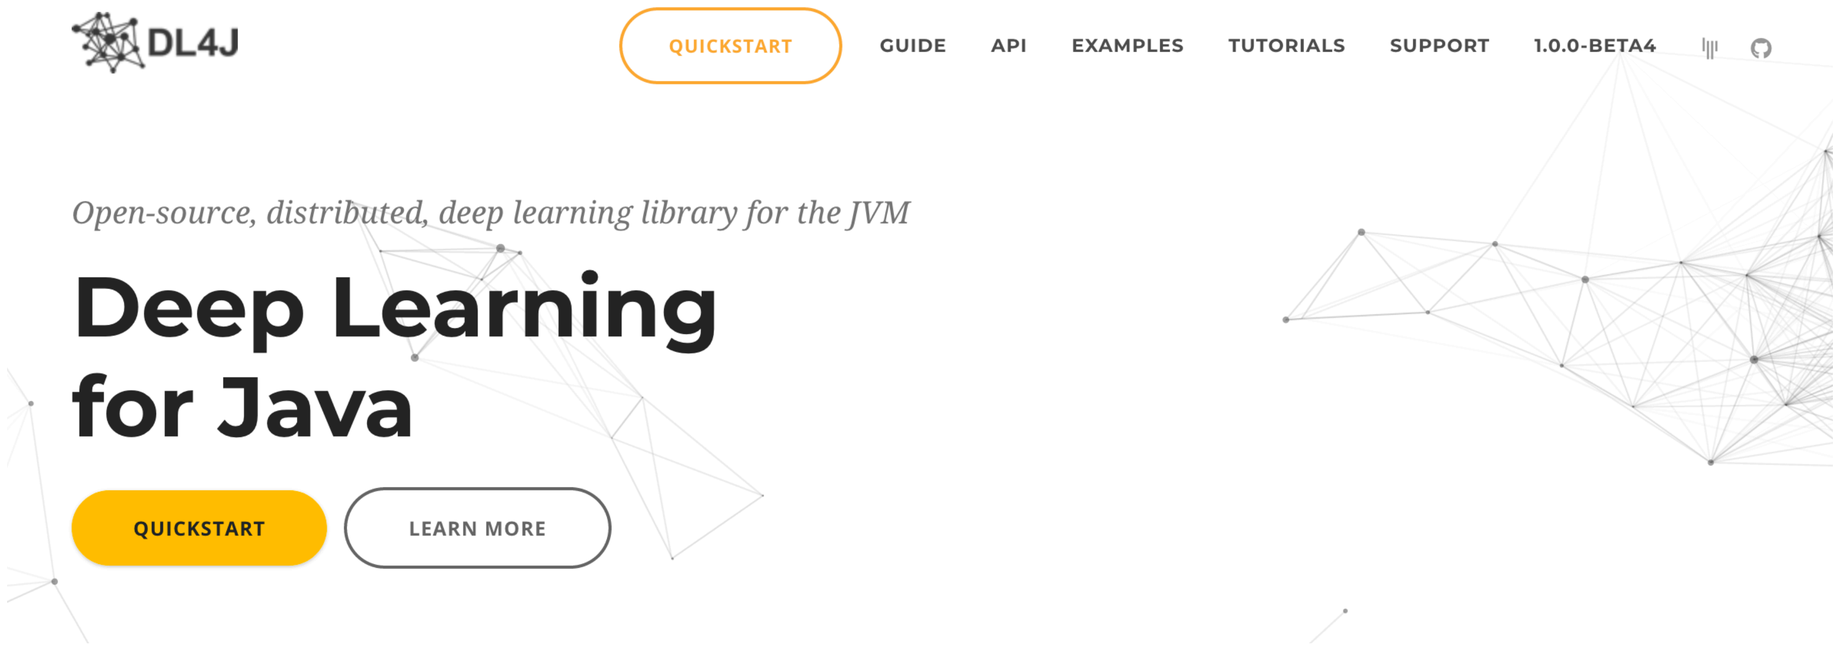

https://deeplearning4j.org/

Practical Deep Learning - Data

https://www.kaggle.com/datasets

https://www.analyticsvidhya.com/blog/2018/03/comprehensive-collection-deep-learning-datasets/

Practical Deep Learning - MNIST

- 60.000 train images
- 10.000 test images
- handwritten digits
- 28 x 28 greyscale

By Josef Steppan - Own work, CC BY-SA 4.0, https://commons.wikimedia.org/w/index.php?curid=64810040

Coding Example

https://colab.research.google.com

google: keras mnist cnn notebook

OR: https://colab.research.google.com/github/csc-training/intro-to-dl/blob/master/day1/keras-mnist-cnn.ipynb

bootcamp.codecentric.ai - free online course

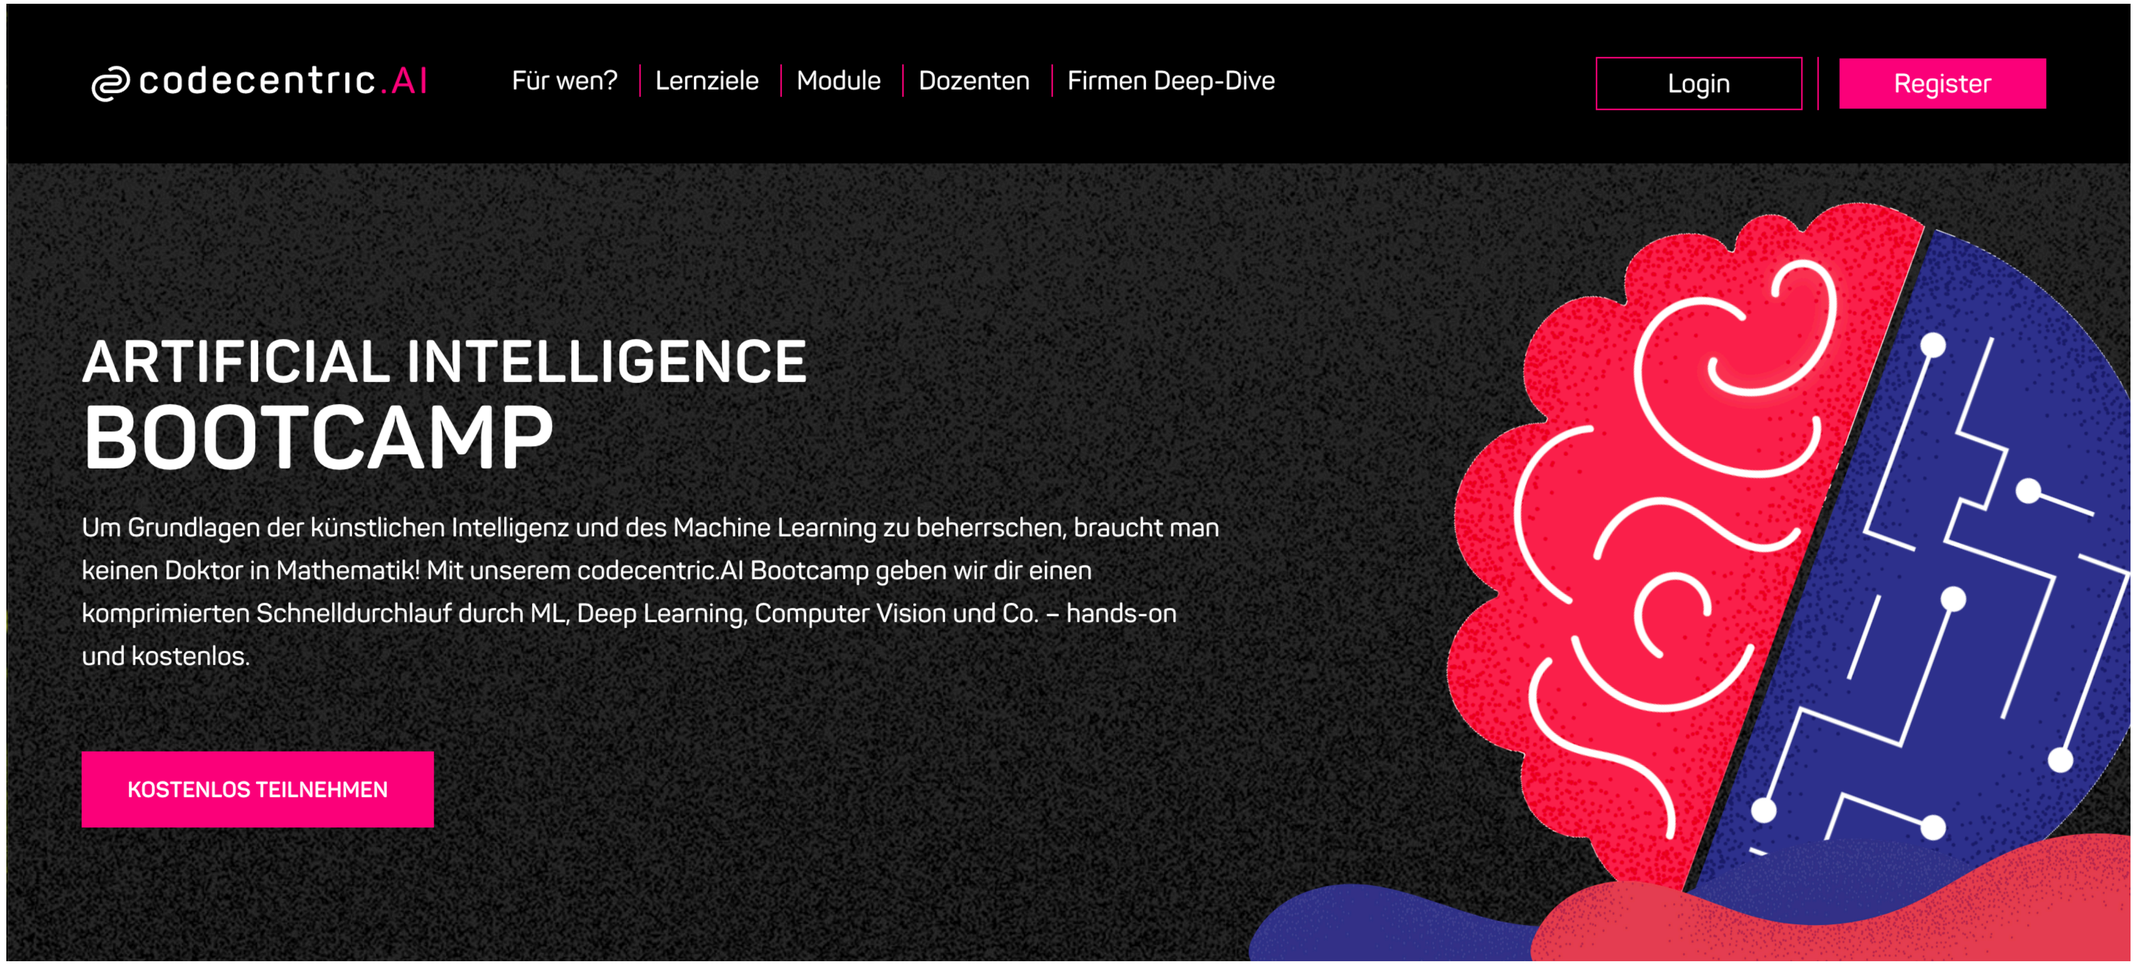

Thank you!

Stay connected!

Am Mittelhafen 14, 48155 Münster

Denis Stalz-John

denis.stalz-john@codecentric.de