Deep Learning: TensorFlow 2.0 vs PyTorch

Dr. Jon Krohn

Metis New York
October 8th, 2019

untapt

Pearson


NYC DATA SCIENCE
ACADEMY


COLUMBIA
UNIVERSITY
IN LVMINE TVO
VIDEBIMVS LVMEN

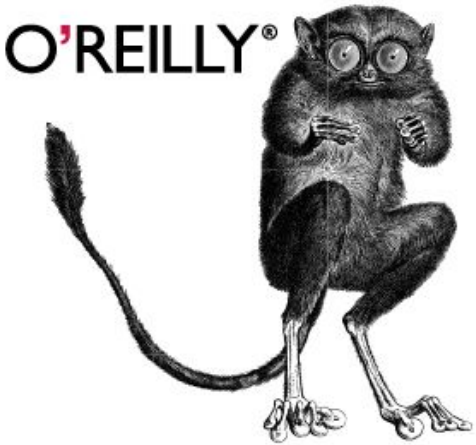
O'REILLY®

Pearson

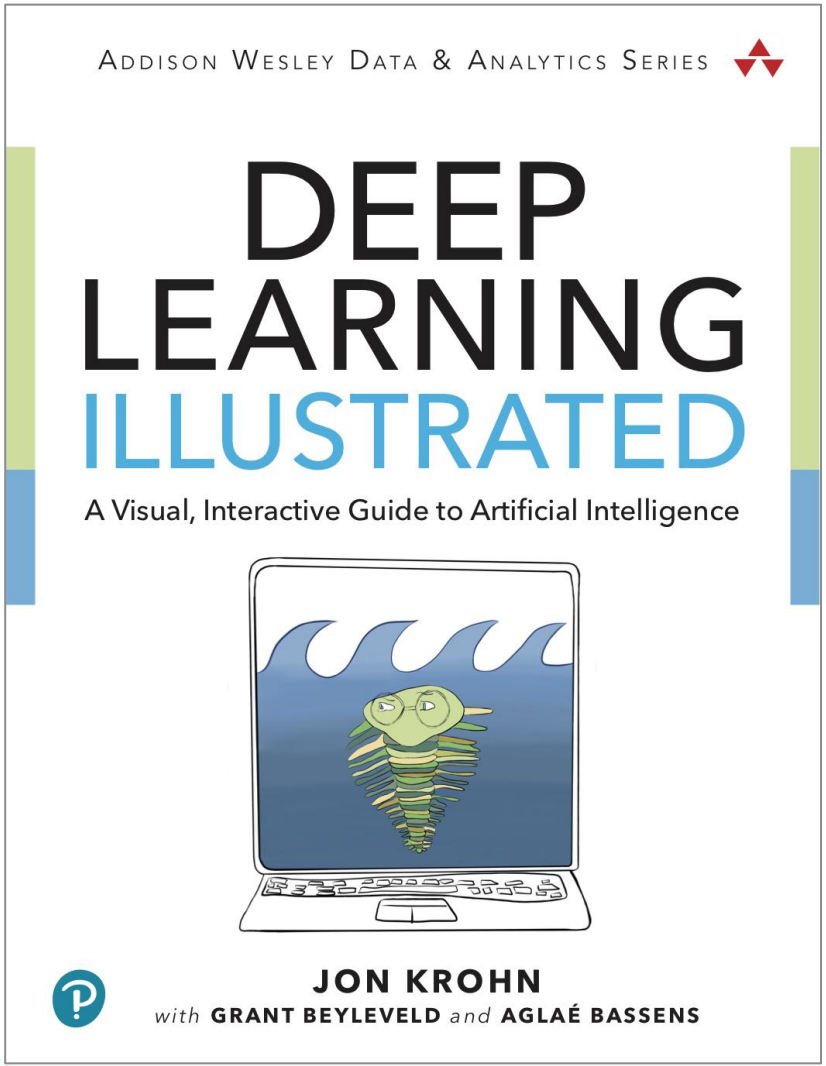

in Safari at:

bit.ly/krohnDL

35% off orders:

bit.ly/iTkrohn

(use code KROHN
during checkout)

POLL

What are you?

- Developer / Engineer
- Scientist / Analyst / Statistician / Mathematician
- Combination of the Above
- Other

POLL

What is your level of familiarity with Deep Learning?

- Little to no exposure to deep learning
- Some deep learning theory
- Deep learning theory + experience with a deep learning library
- Deep learning theory + experience with TensorFlow/Keras

TensorFlow 2.0 vs PyTorch

- Vision Analogy for Deep Learning
- Deep Learning Families
- Deep Learning Libraries

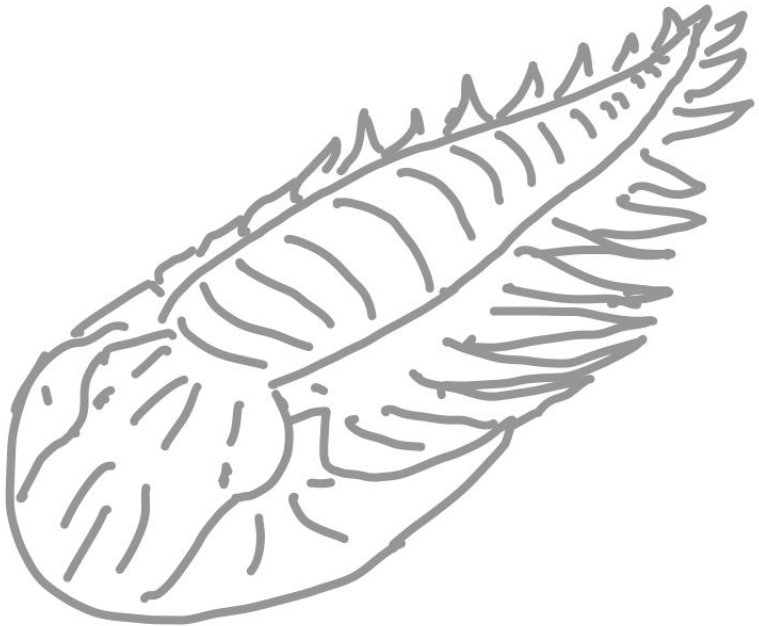

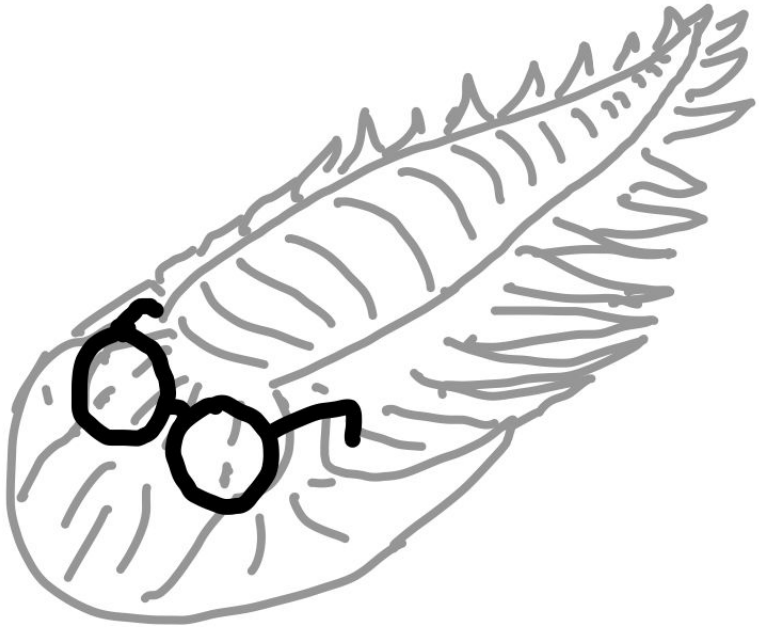

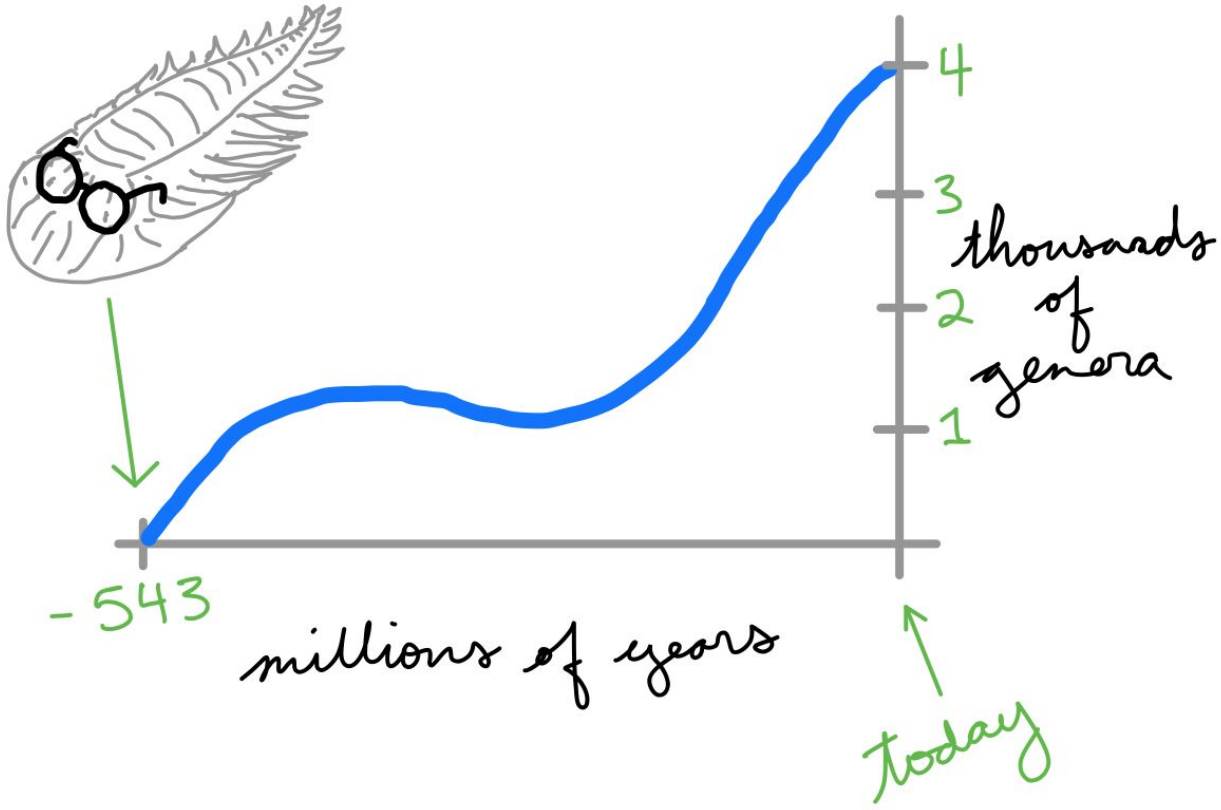

Case Study: The History of Vision

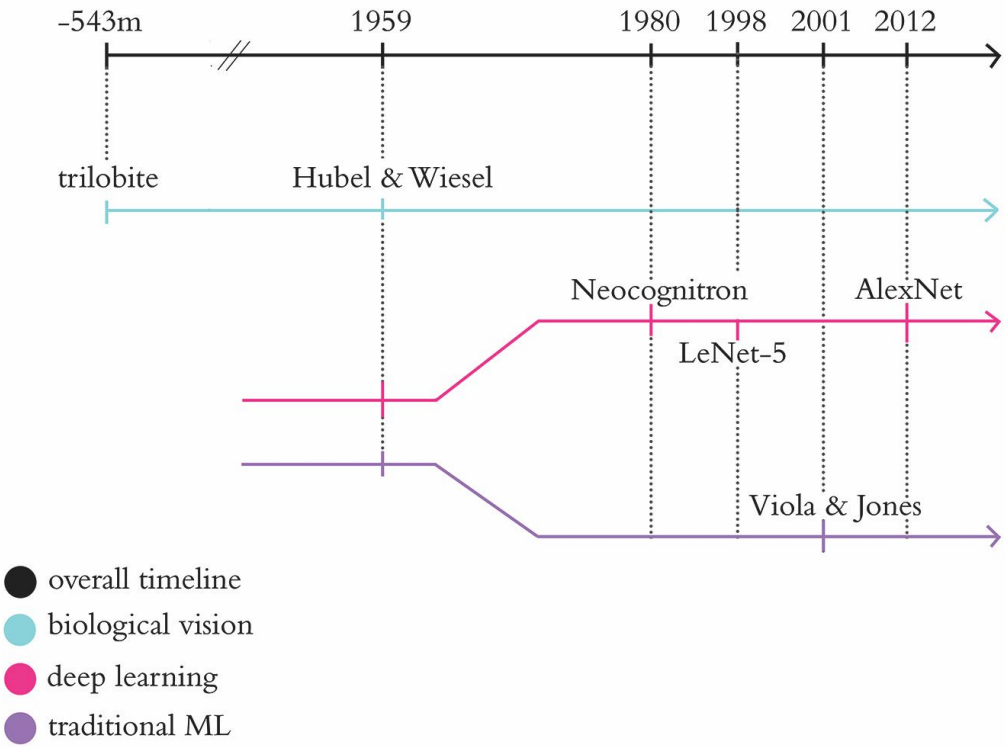

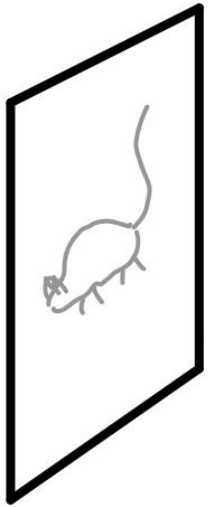

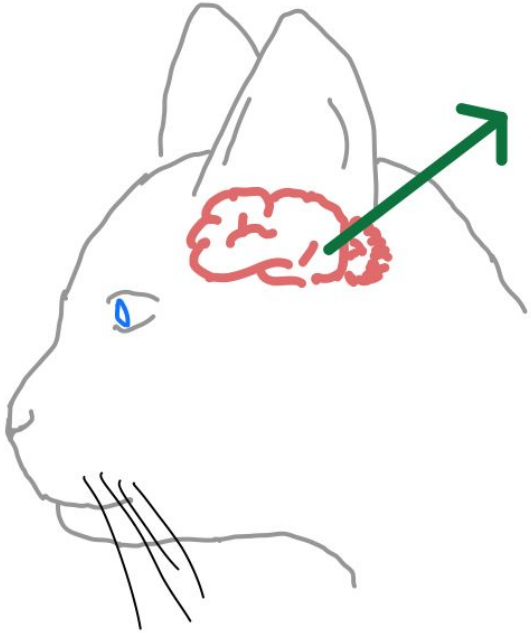

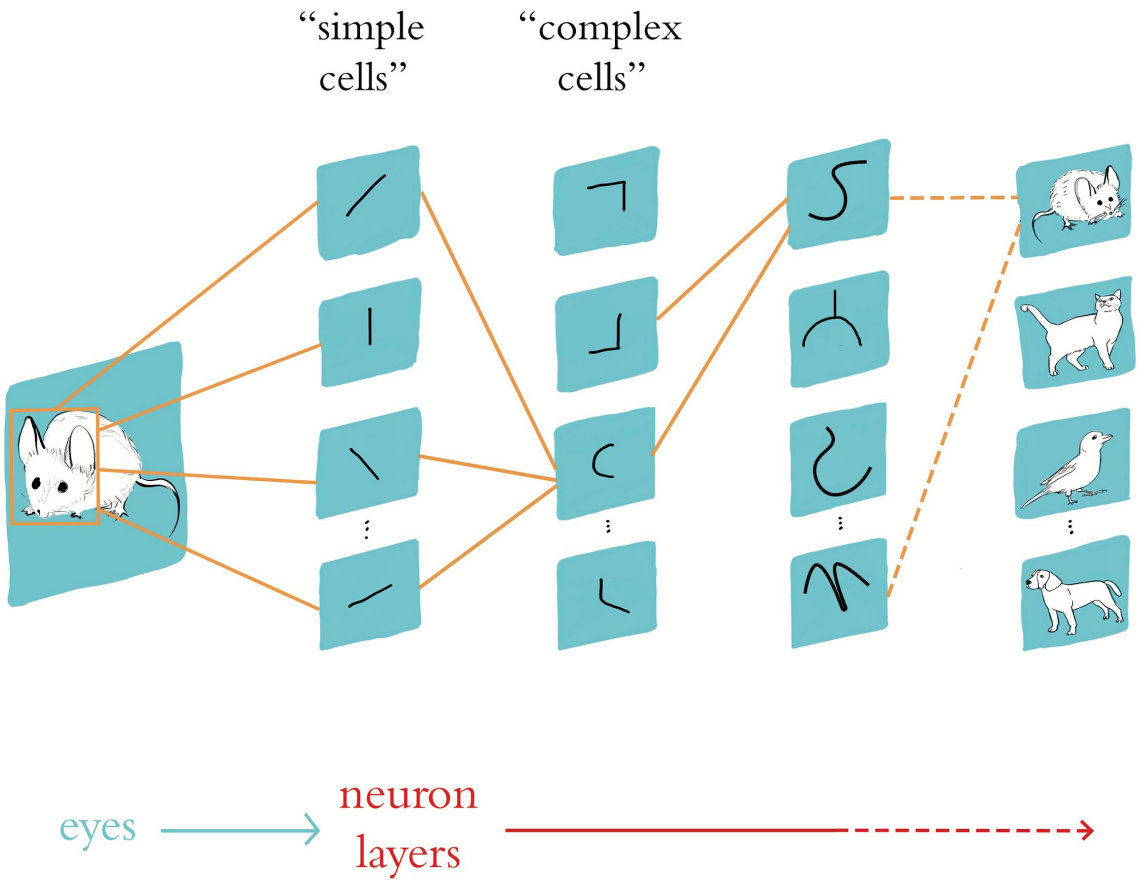

Case Study: The History of Vision

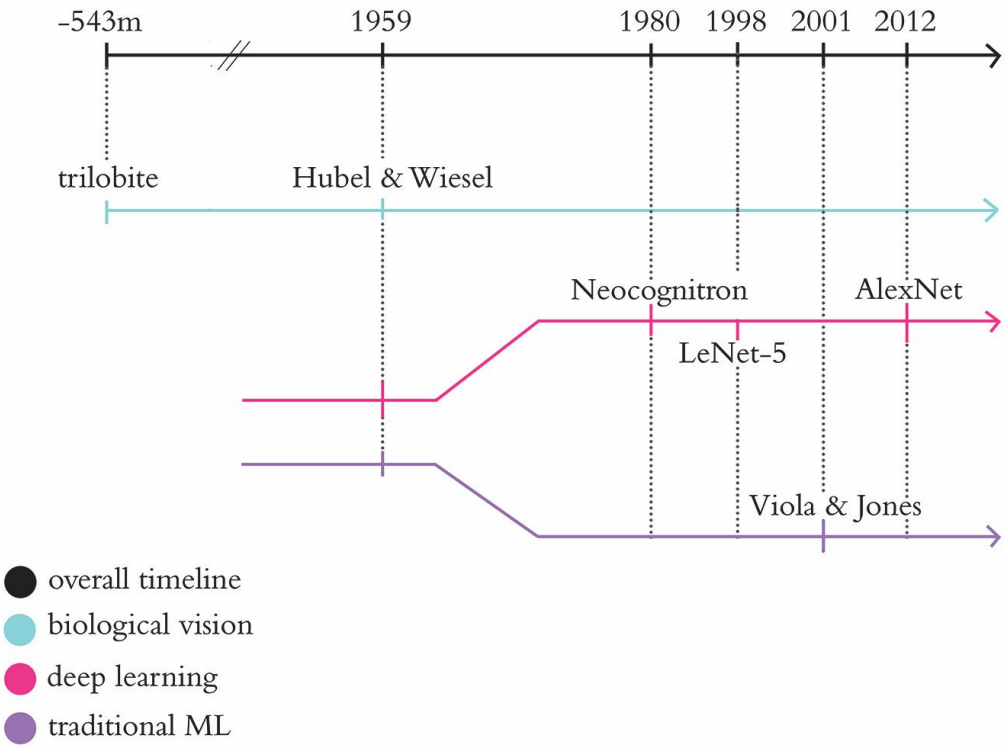

Neocognitron (Fukushima, 1980)

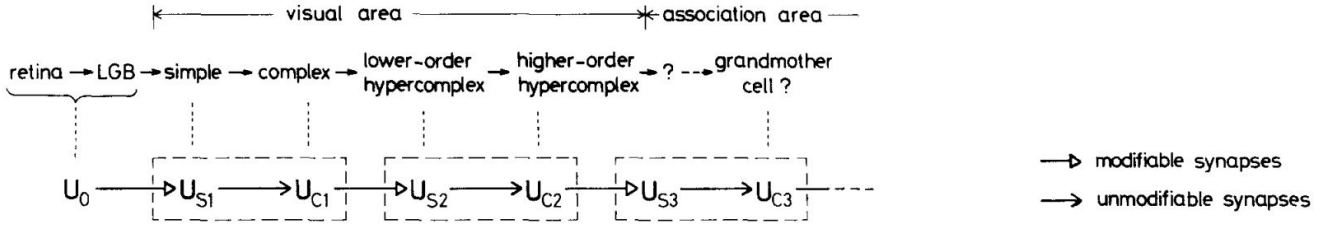

Fig. 1. Correspondence between the hierarchy model by Hubel and Wiesel, and the neural network of the neocognitron

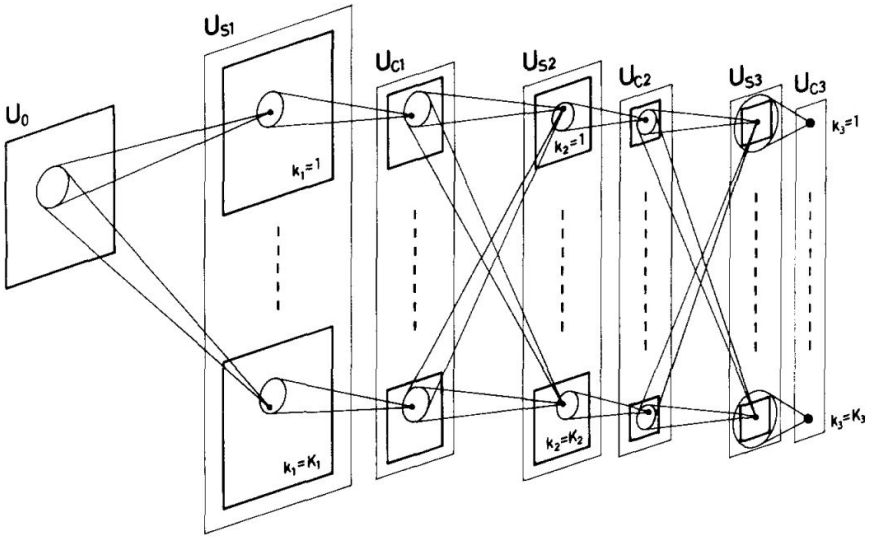

Fig. 2. Schematic diagram illustrating the interconnections between layers in the neocognitron

Case Study: The History of Vision

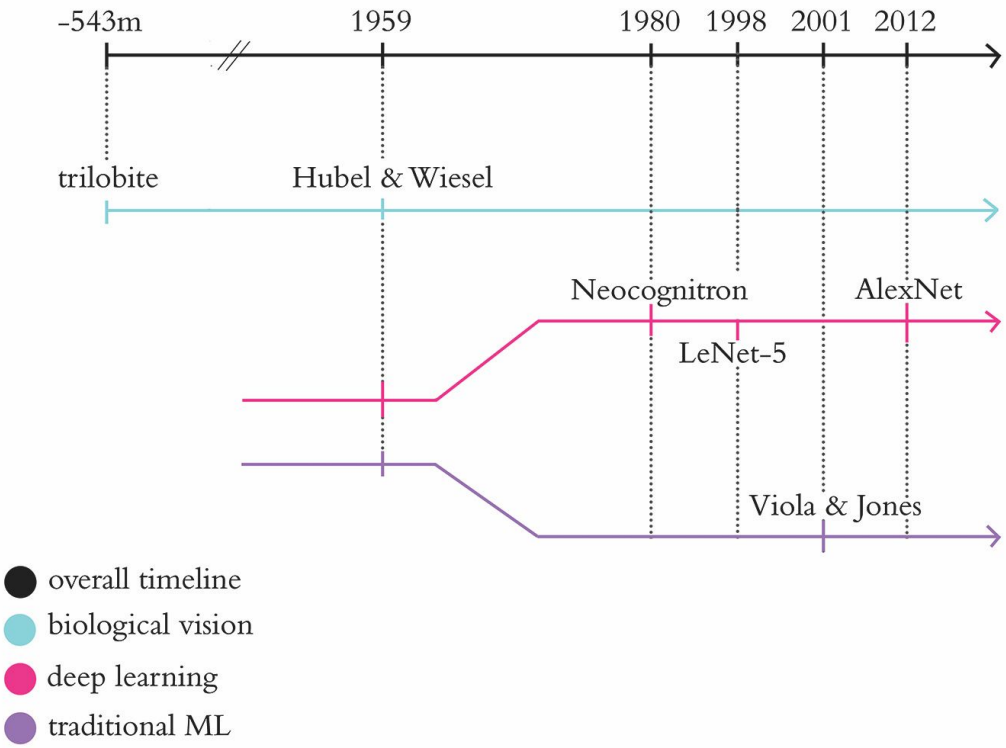

LeNet-5 (LeCun et al., 1998)

input image → large simple features → → smaller more complex features → probability outputs

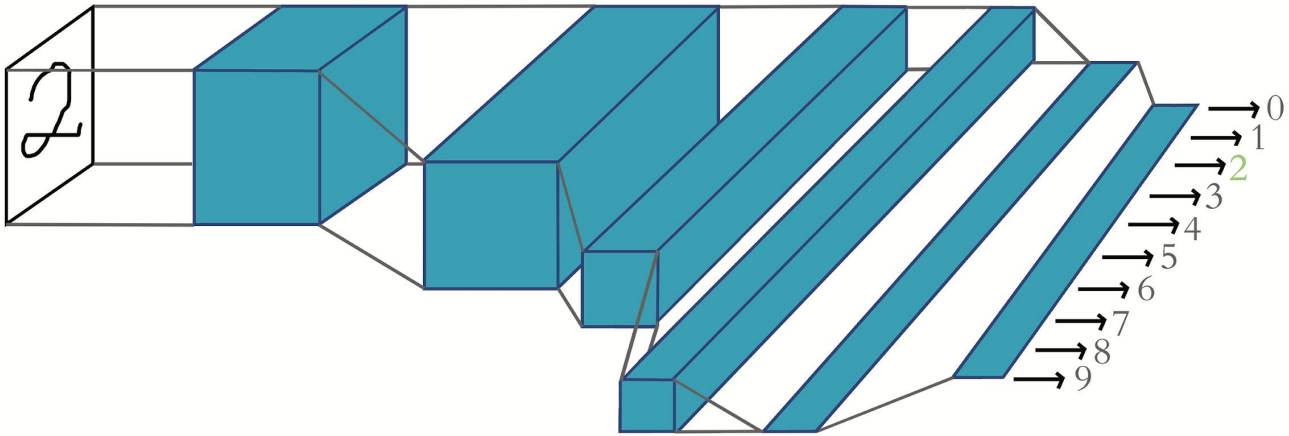

Case Study: The History of Vision

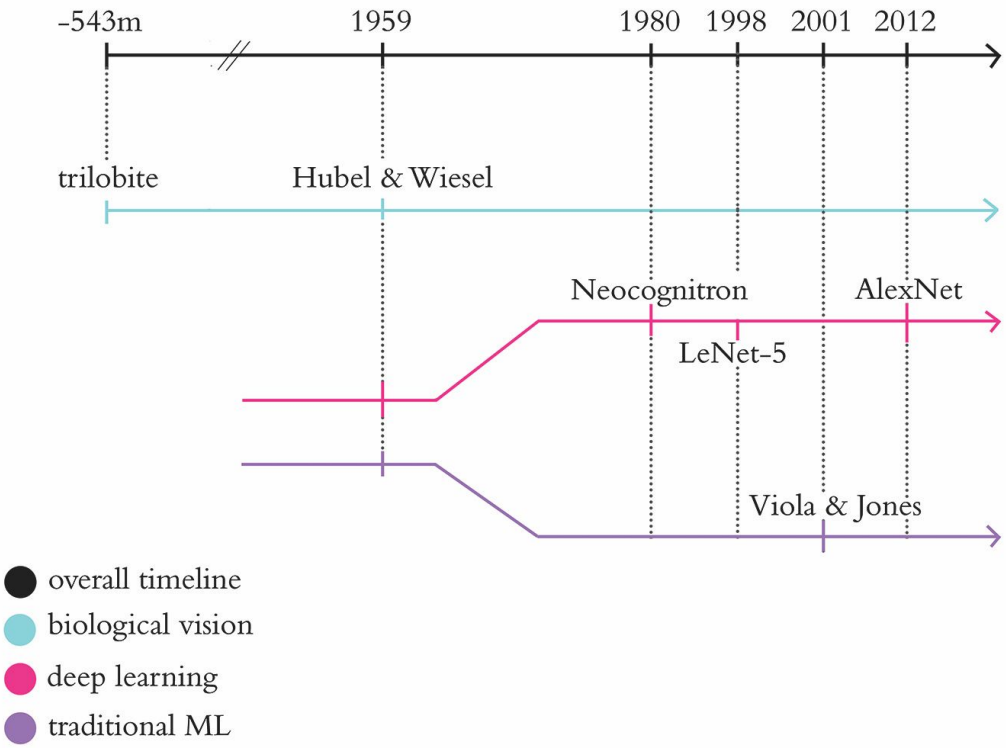

Traditional ML vs Deep Learning

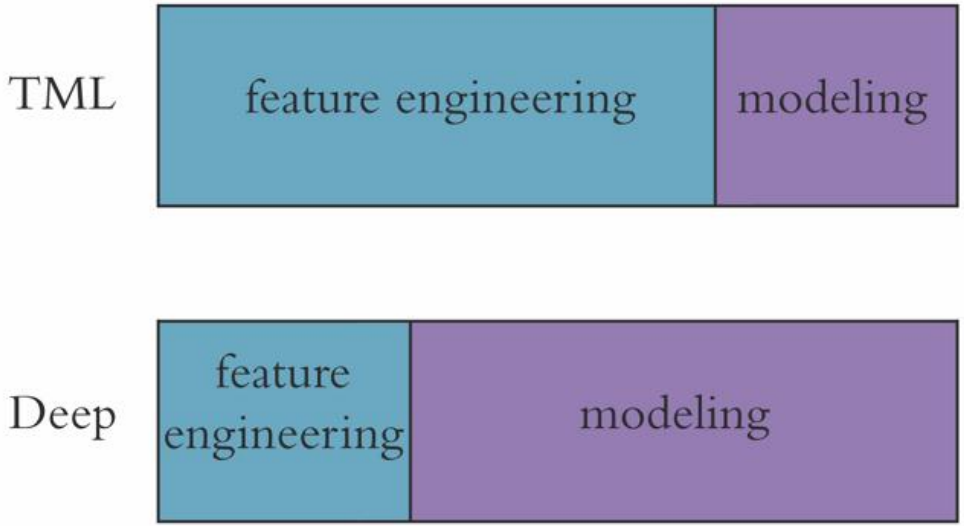

Case Study: The History of Vision

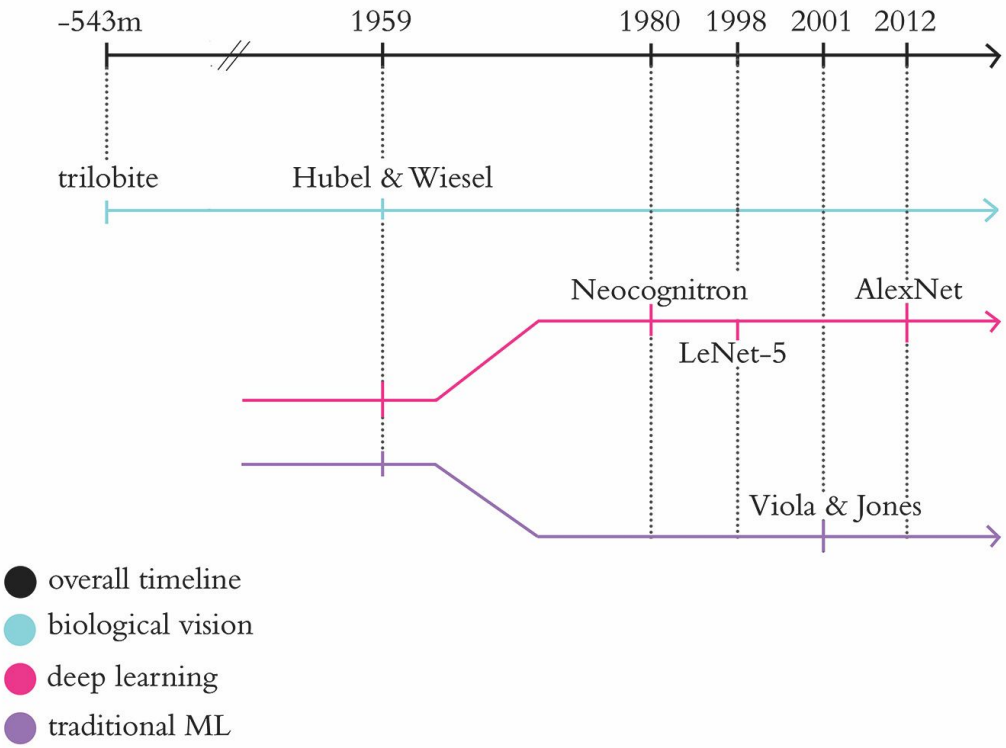

Viola & Jones (2001)

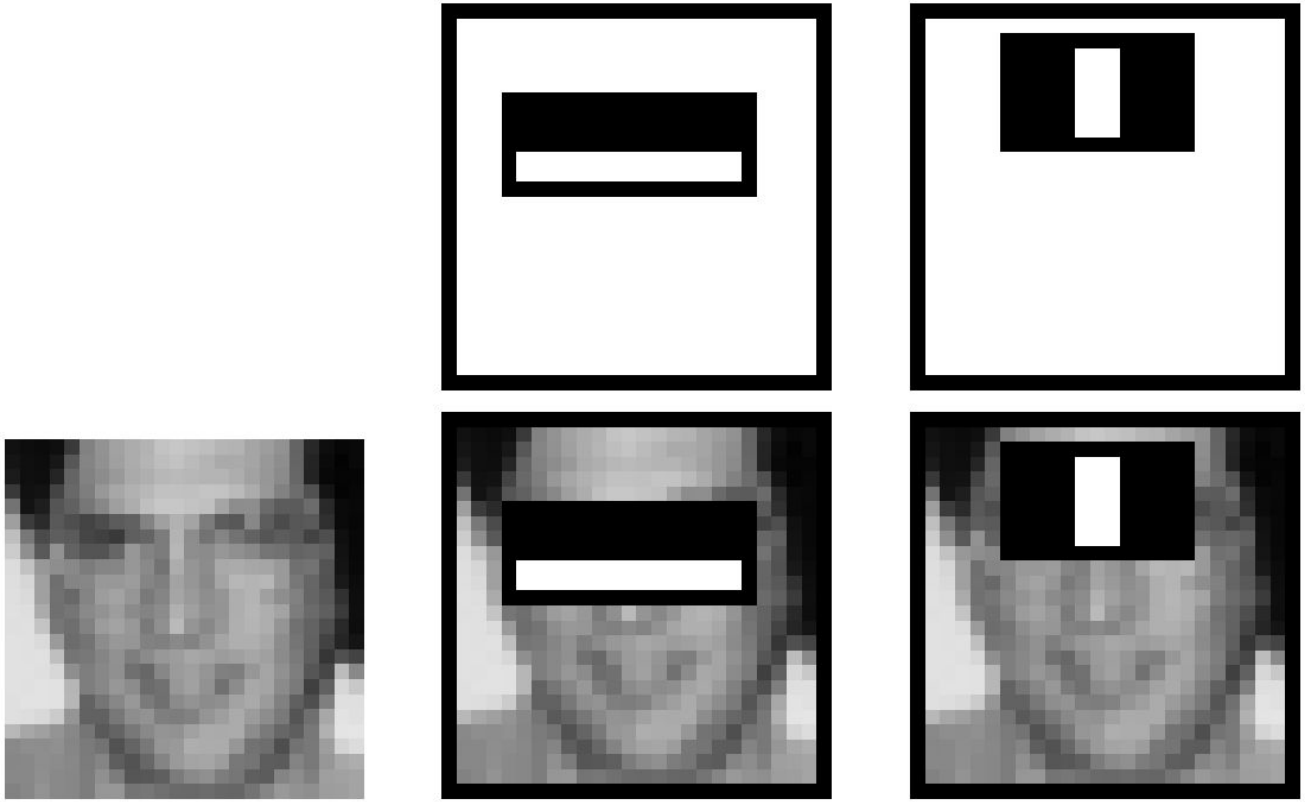

Case Study: The History of Vision

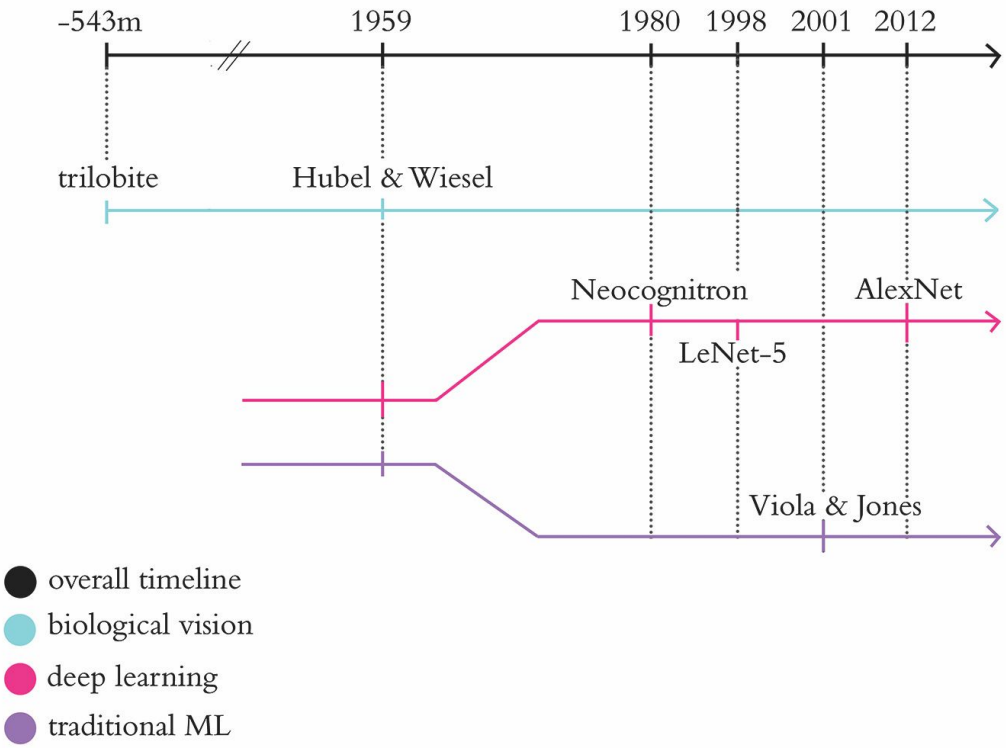

AlexNet (Krizhevsky et al., 2012)

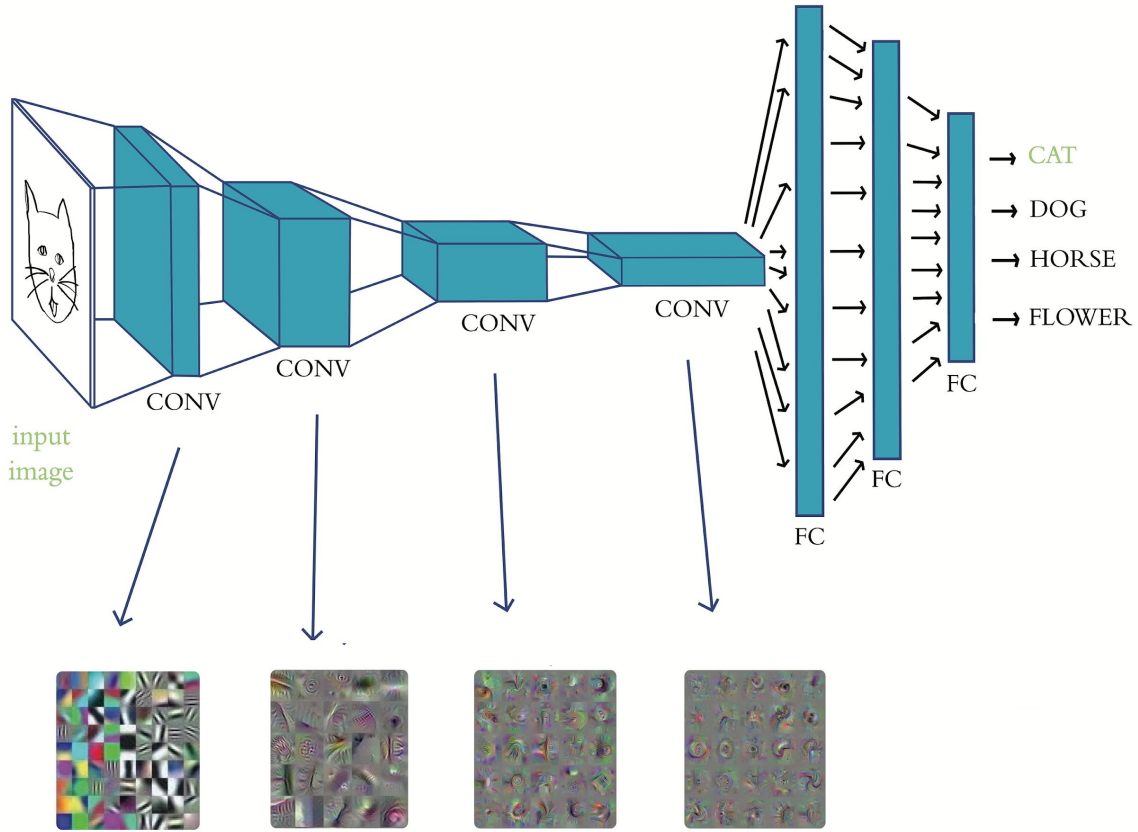

POLL

If a voice recognition algorithm is fed audio of speech as inputs, given corresponding text as the outputs (labels) to learn, and no features are explicitly programmed, is this a:

- Traditional Machine Learning Algorithm
- Deep Learning Algorithm
- I Don't Know

TensorFlow 2.0 vs PyTorch

- Vision Analogy for Deep Learning
- **Deep Learning Families**
- Deep Learning Libraries

Dense Networks

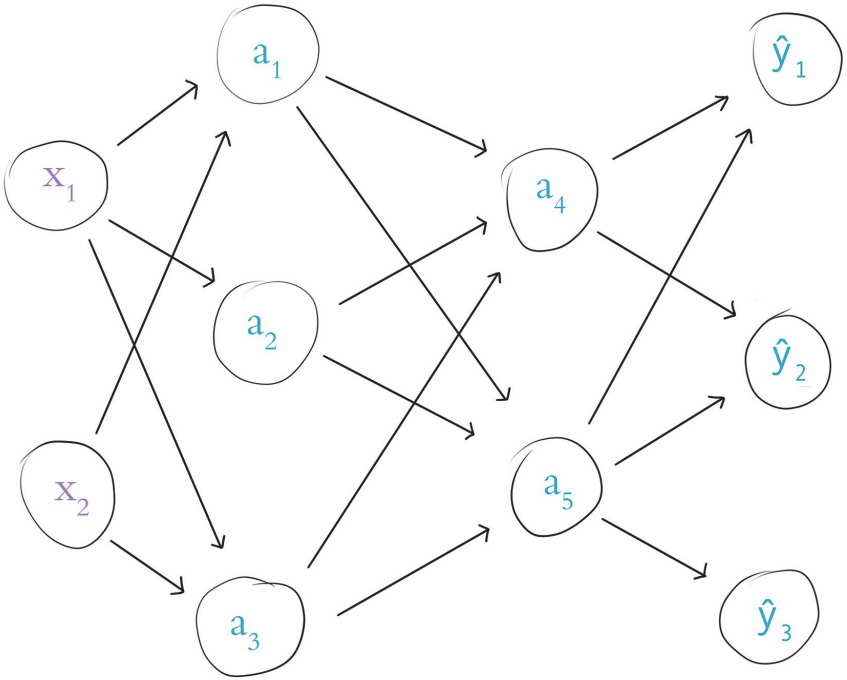

ConvNets: Convolutional Networks

ConvNets: Convolutional Networks

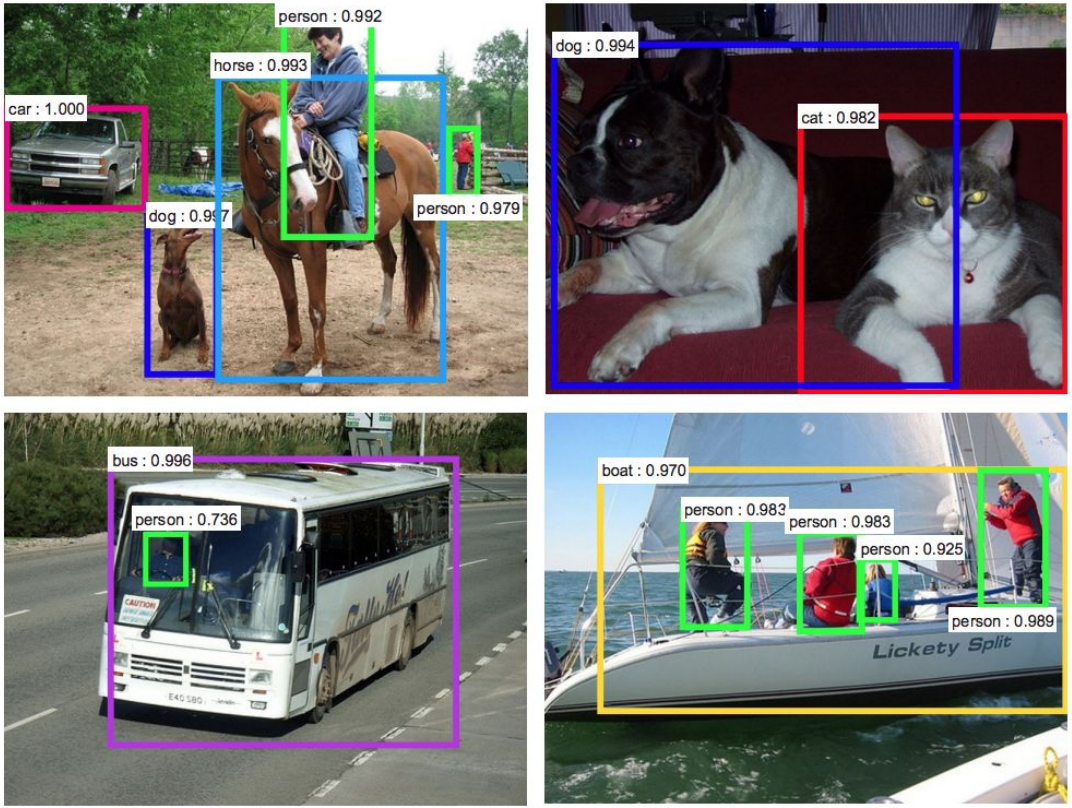

Ren et al. (2015)

RNNs: Recurrent Neural Networks

GANs: Generative Adversarial Networks

GANs: Generative Adversarial Networks

Karros et al. (2018)

Deep Reinforcement Learning

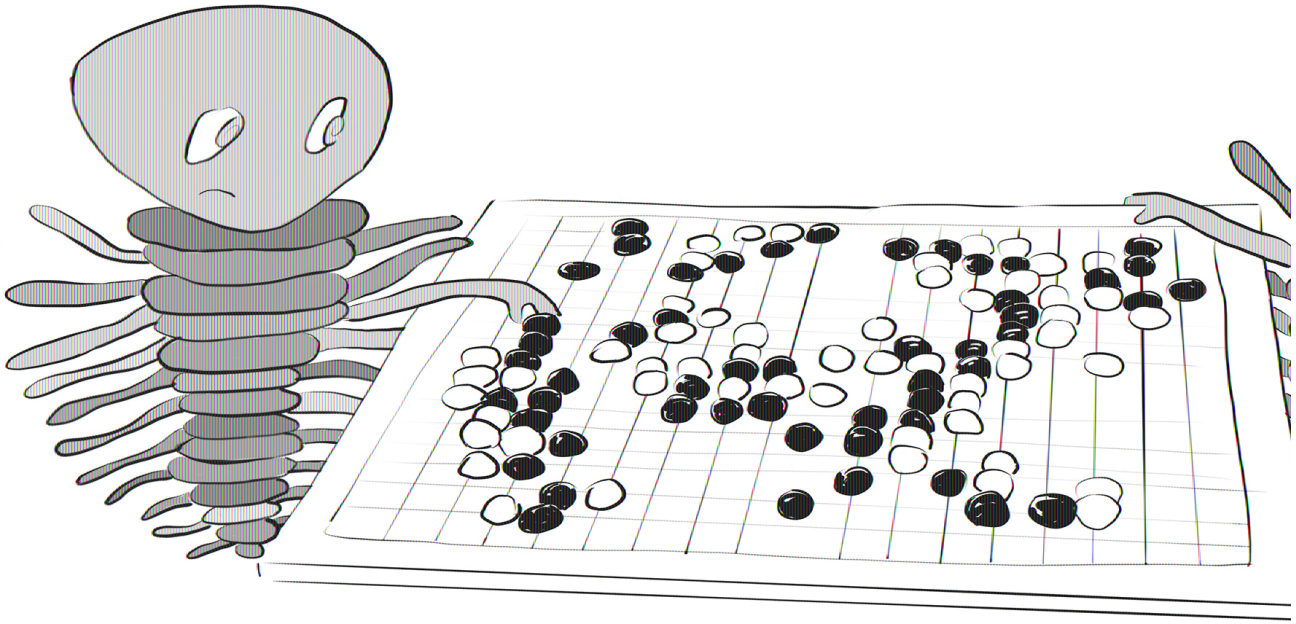

POLL

If you were designing an algorithm to learn to play Tetris by maximizing its score, which of these Deep Learning approaches would be most appropriate?

- Convolutional Neural Network
- Recurrent Neural Network
- Deep Reinforcement Learning
- Generative Adversarial Network

POLL

If you were designing an algorithm to recognise tumours in medical images, which of these Deep Learning approaches would be most appropriate?

- Convolutional Neural Network
- Recurrent Neural Network
- Deep Reinforcement Learning
- Generative Adversarial Network

POLL

If you were designing an algorithm to predict stock price movements based on time series data, which of these Deep Learning approaches would be most appropriate?

- Convolutional Neural Network
- Recurrent Neural Network
- Deep Reinforcement Learning
- Generative Adversarial Network

TensorFlow 2.0 vs PyTorch

- Vision Analogy for Deep Learning
- Deep Learning Families
- **Deep Learning Libraries**

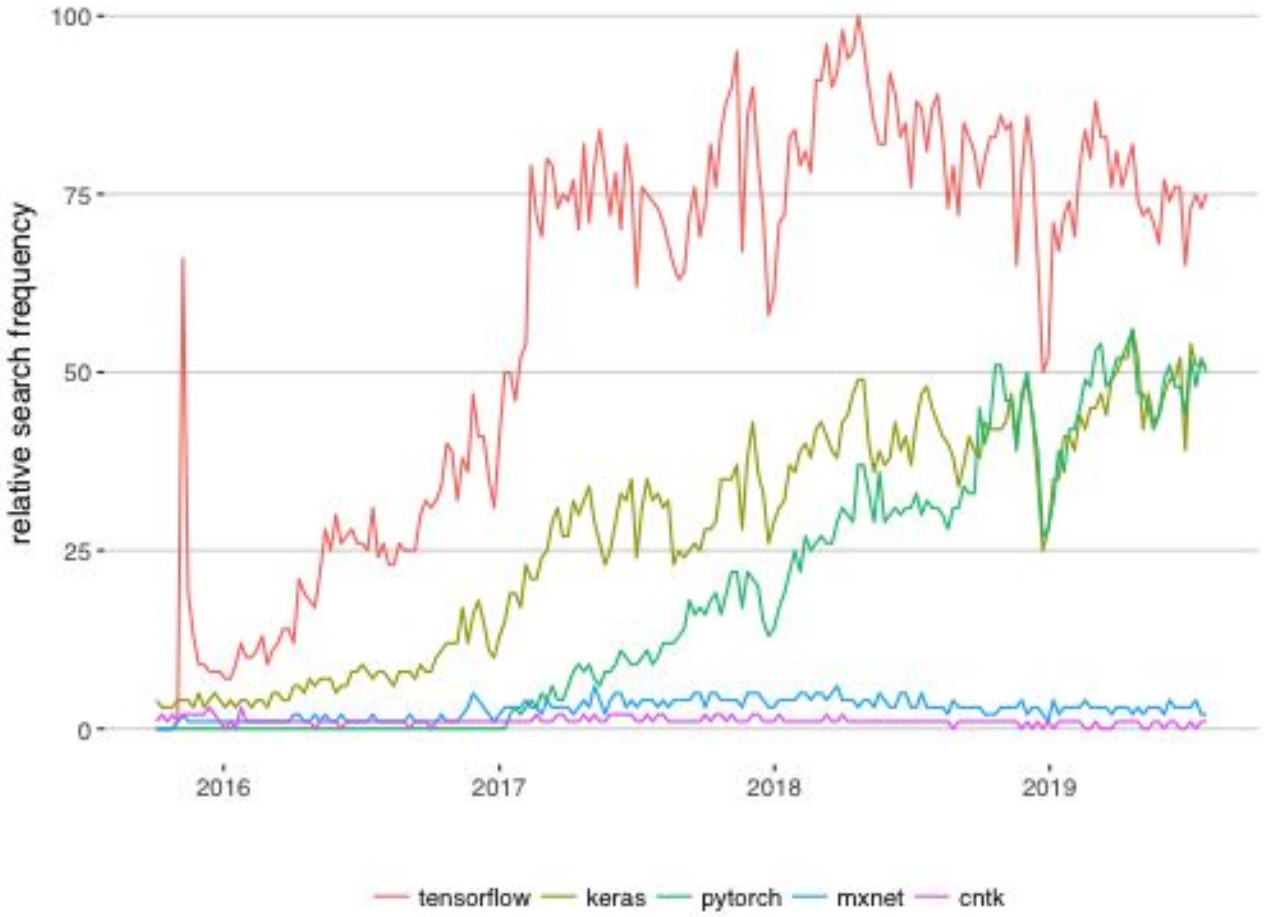
100
75
50
25
0
relative search frequency
2016
2017
2018
2019
tensorflow
keras
pytorch
mxnet
cntk

Leading Deep Learning Libraries

	Caffe	Torch	MXNet	TensorFlow
Language	Python, Matlab	Lua, C	Python, R, C++ Julia, Matlab JavaScript, Go Scala, Perl	**Python, C**, C++ Java, Go, JS, Swift (*Haskell, Julia, R, Scala, Rust, C#)*
Programming Style	Symbolic	Imperative	Imperative	Imperative *(in 2.0)*
Parallel GPUs: Data	Yes	Yes	Yes	Yes
Parallel GPUs: Model		Yes	Yes	Yes
Pre-Trained Models	Model Zoo	Model Zoo	Model Zoo	github.com/tensorflow/models
High-Level APIs		PyTorch	in-built	Keras
Particular Strength	CNNs	interactivity		production deployment

PyTorch	TensorFlow
“NumPy”, optimized for GPUs	ported to Python from C++
dynamic auto-differentiation (autodiff)	static computational graph
debugging is easier	
fast.ai API	Keras API
	more widely adopted
TorchScript Just-In-Time compilation	TensorFlow Serving, .js, Lite, tf.data, tf.io
better for interactively building models	**better for production deployments**

Deep Nets in Three Libraries

demo: *in TF 1.x*

interactive Colab demo: *in TF 2.0* (`bit.ly/deepNetTF`)

interactive Colab demo: *in PyTorch* (`bit.ly/deepPTdemo`)

POLL

What follow-up Deep Learning topics interest you most?

- CNNs and Machine Vision
- Natural Language Processing
- Time-Series Predictions
- Generative Adversarial Networks
- Deep Reinforcement Learning
- Something Else

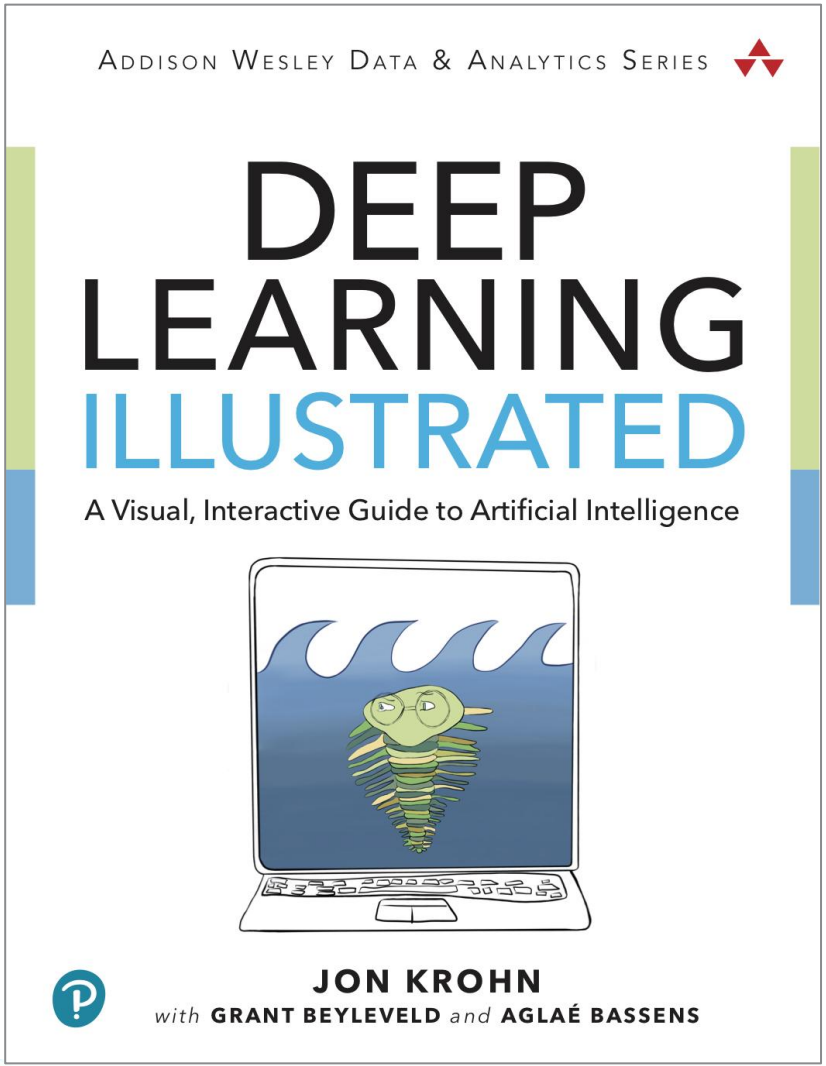

in Safari at:

bit.ly/krohnDL

35% off orders:

bit.ly/iTkrohn

(use code KROHN
during checkout)

Stay in Touch

...and let me know what other topics you'd love me to teach!

`jonkrohn.com` to sign up for email newsletter

`twitter.com/JonKrohnLearns`

`medium.com/@jonkrohn`

`linkedin.com/in/jonkrohn`

(with message mentioning today's Live Training)