https://www.sli.do/
#073374

Fundamentals of Deep Learning (I)

Learning Objectives

- Learn the fundamental concept of deep learning.
- Learn the concept of artificial neural network (ANN)
- Learn the perceptron: the simplest ANN model and the training process.

Artificial Intelligence Machine Learning & Deep Learning

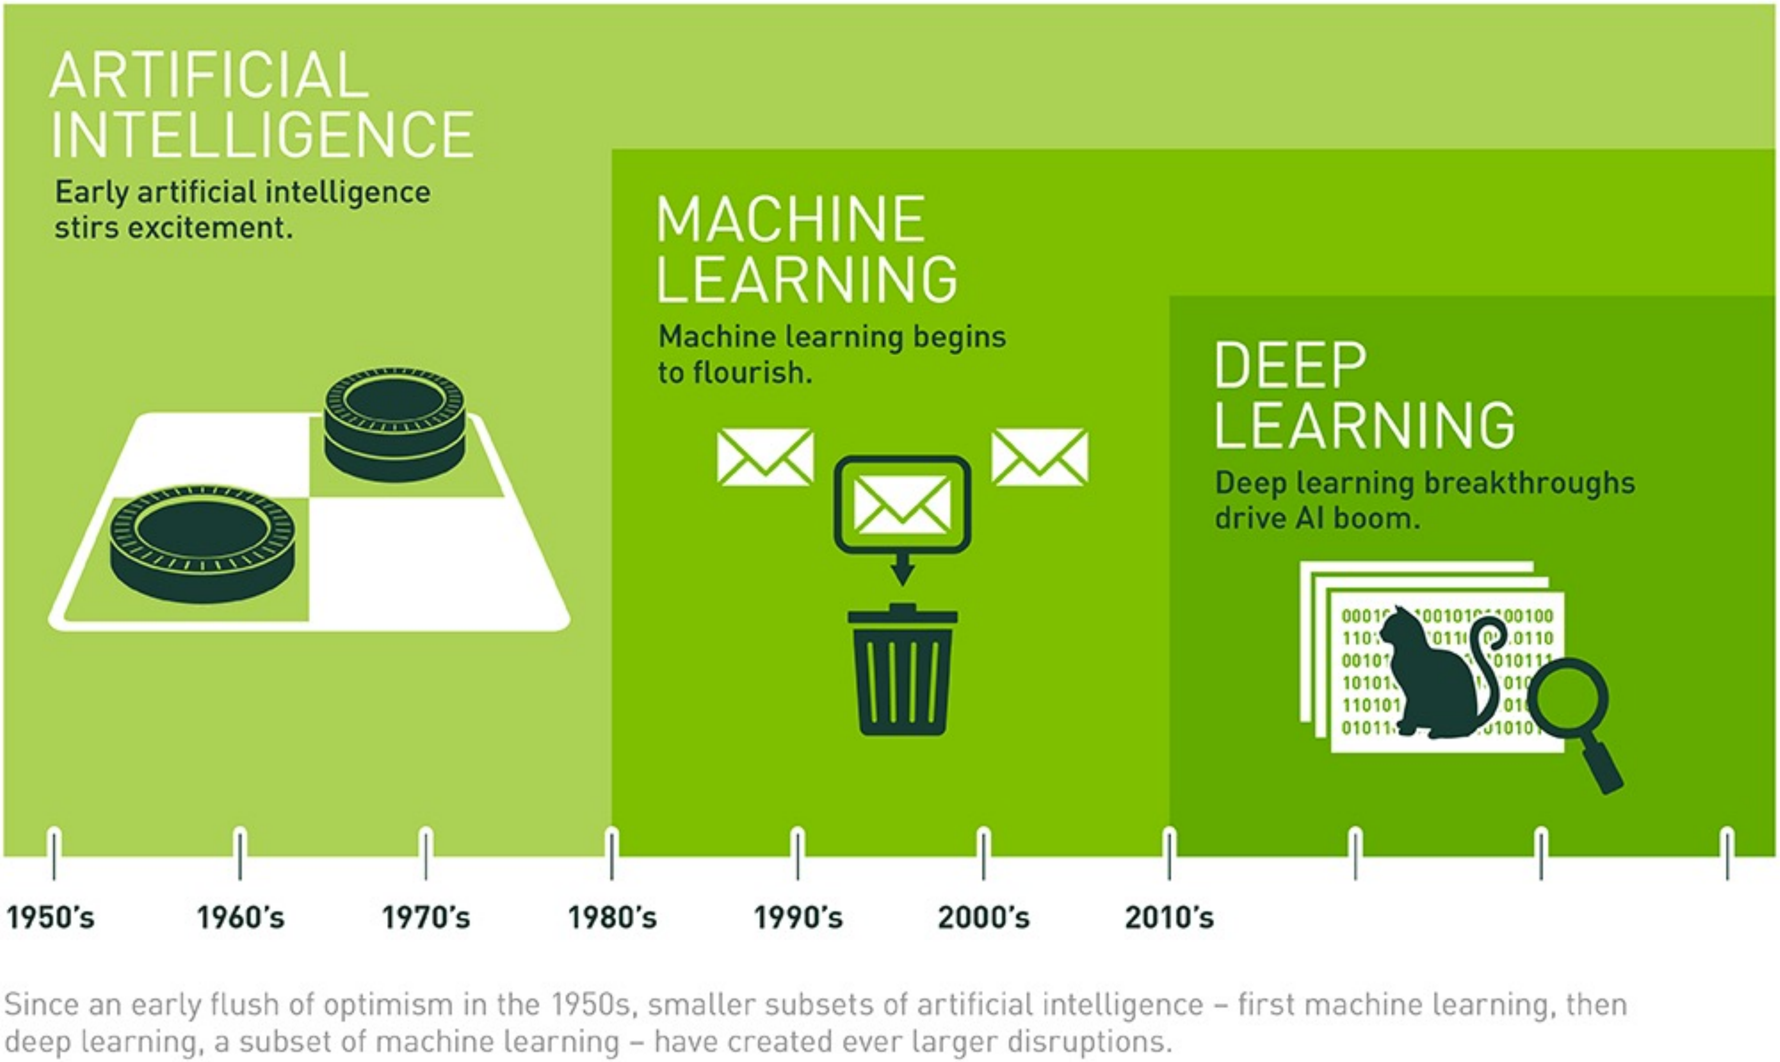

Source: https://blogs.nvidia.com/blog/2016/07/29/whats-difference-artificial-intelligence-machine-learning-deep-learning-ai/

Fun Time: what does the "deep" in deep learning stand for? (1) depth of understanding (2) depth of math (3) depth of layers

(Artificial) Neural Networks (ANN): The Heart of Deep Learning

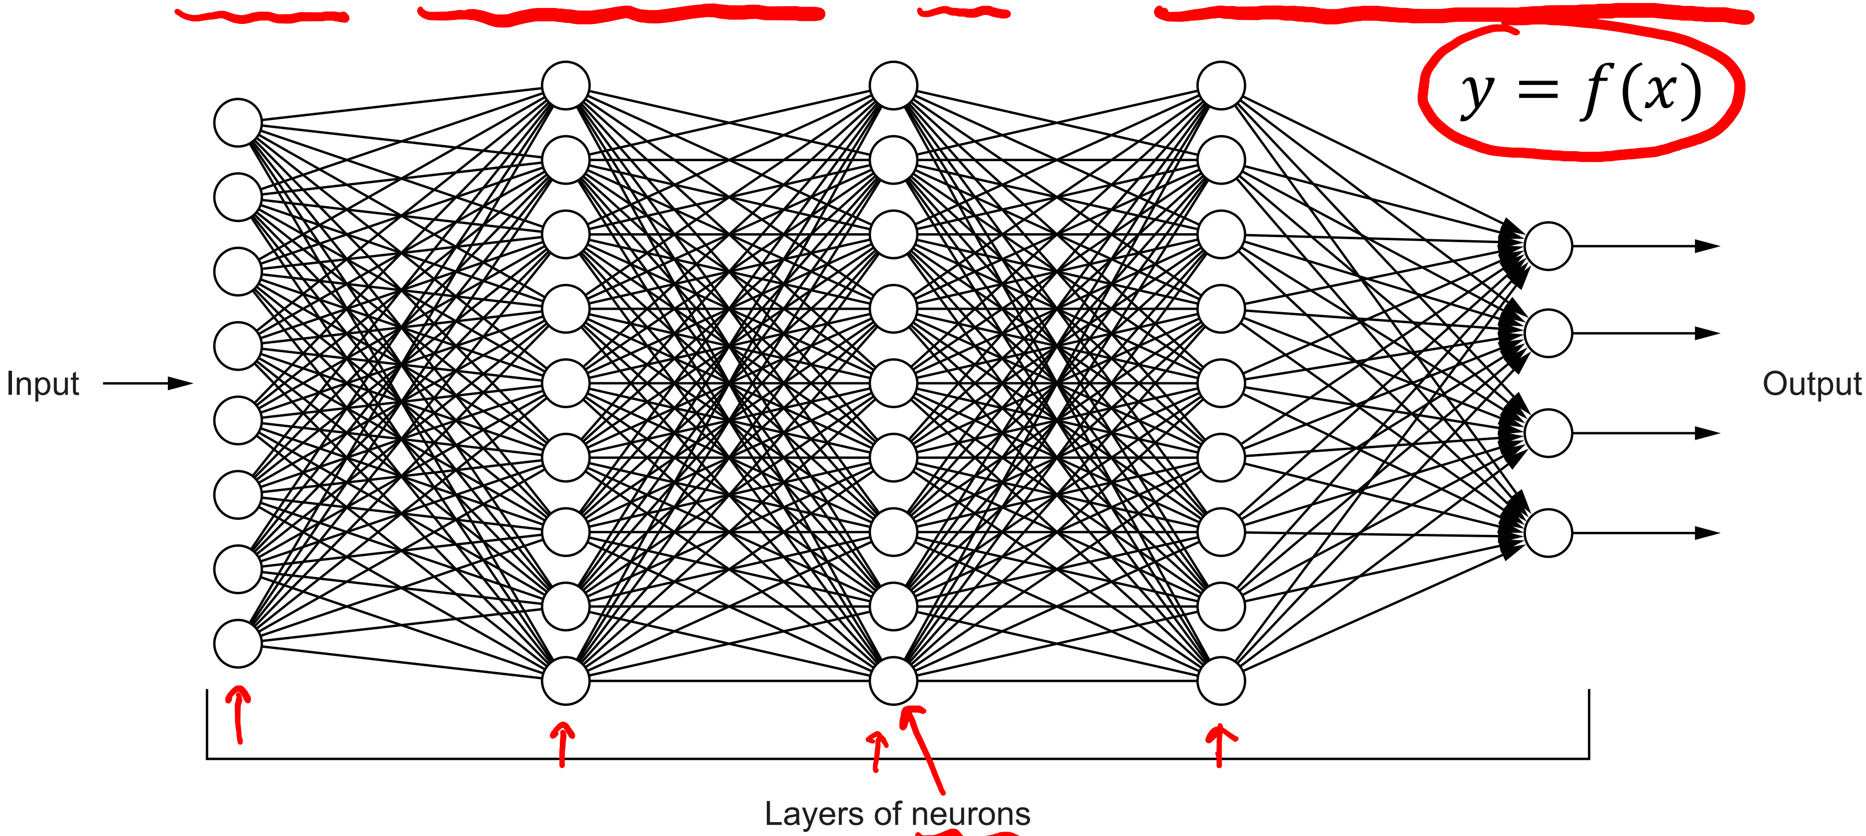

ANN is a collection of simple processing units (nodes) that are connected by directed links (edges)

- Every node receives signals from incoming edges, performs computations, and transmits signals to outgoing edges
- Analogous to *human brain* where **nodes are neurons** and signals are electrical impulses
- **Weight** of an edge determines the strength of connection between the nodes

Brief History of Neural Networks

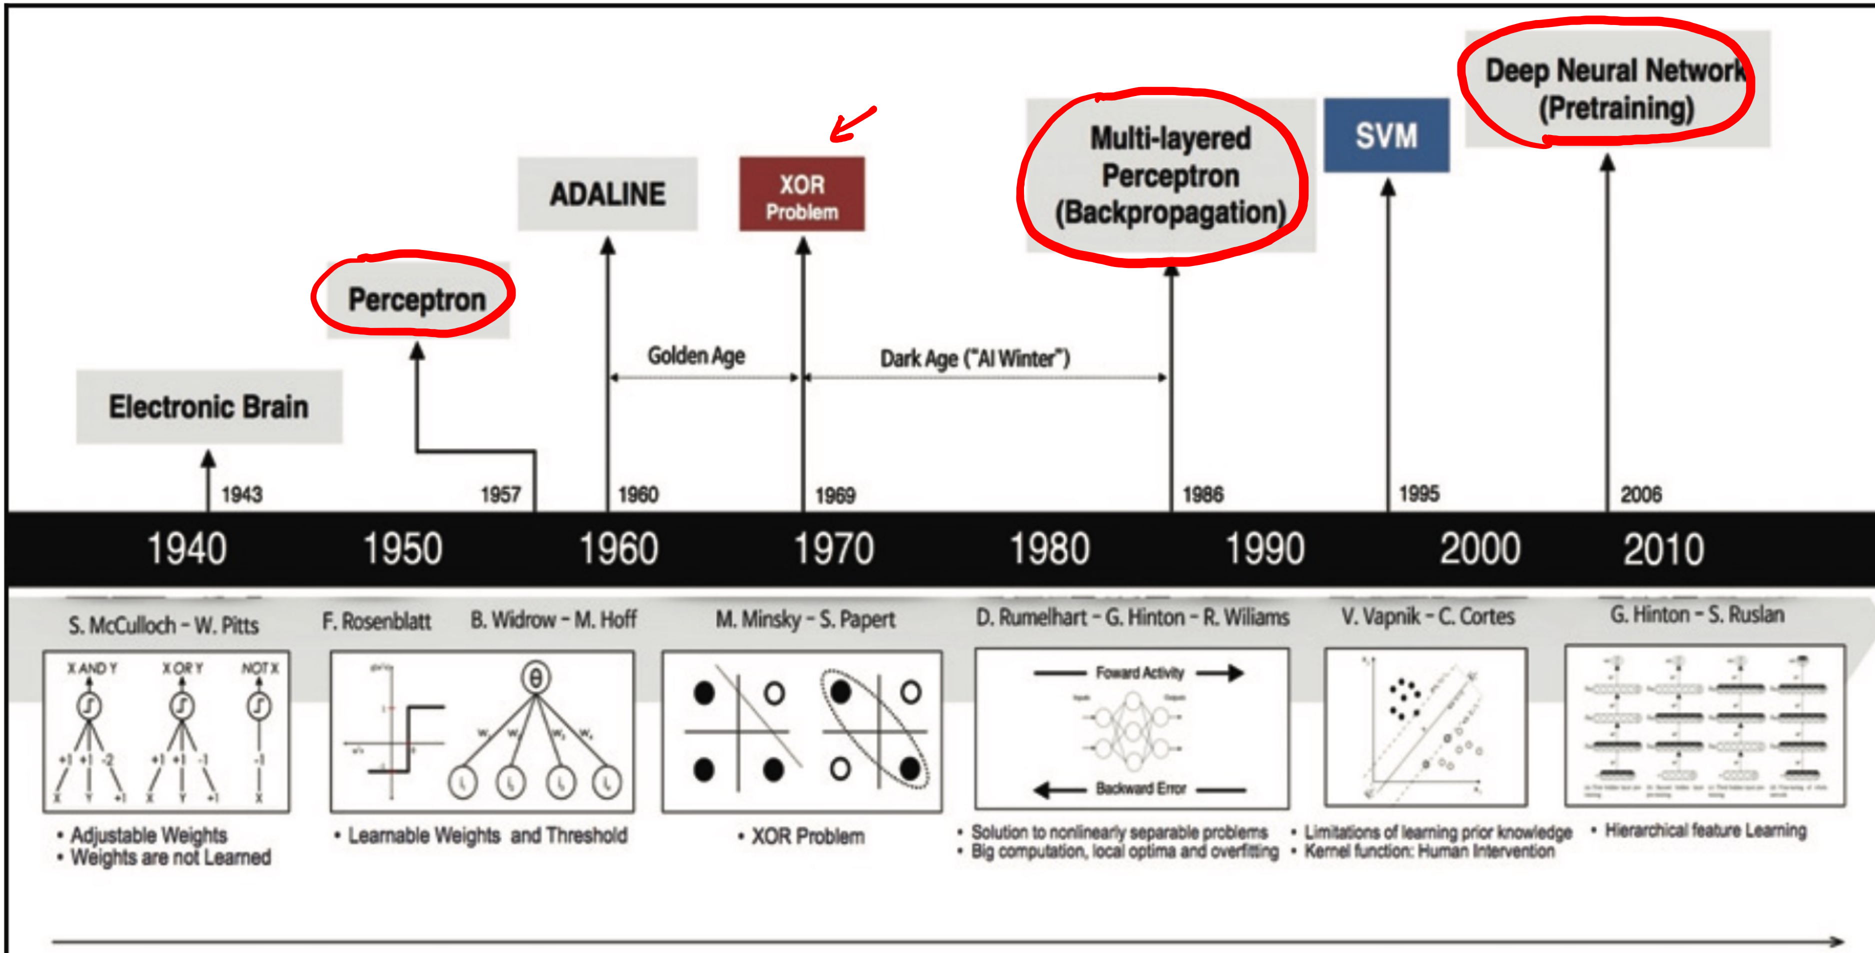

Brief history of neural network

Increase of Performance and Layers on ImageNet Classification Over Time

errors

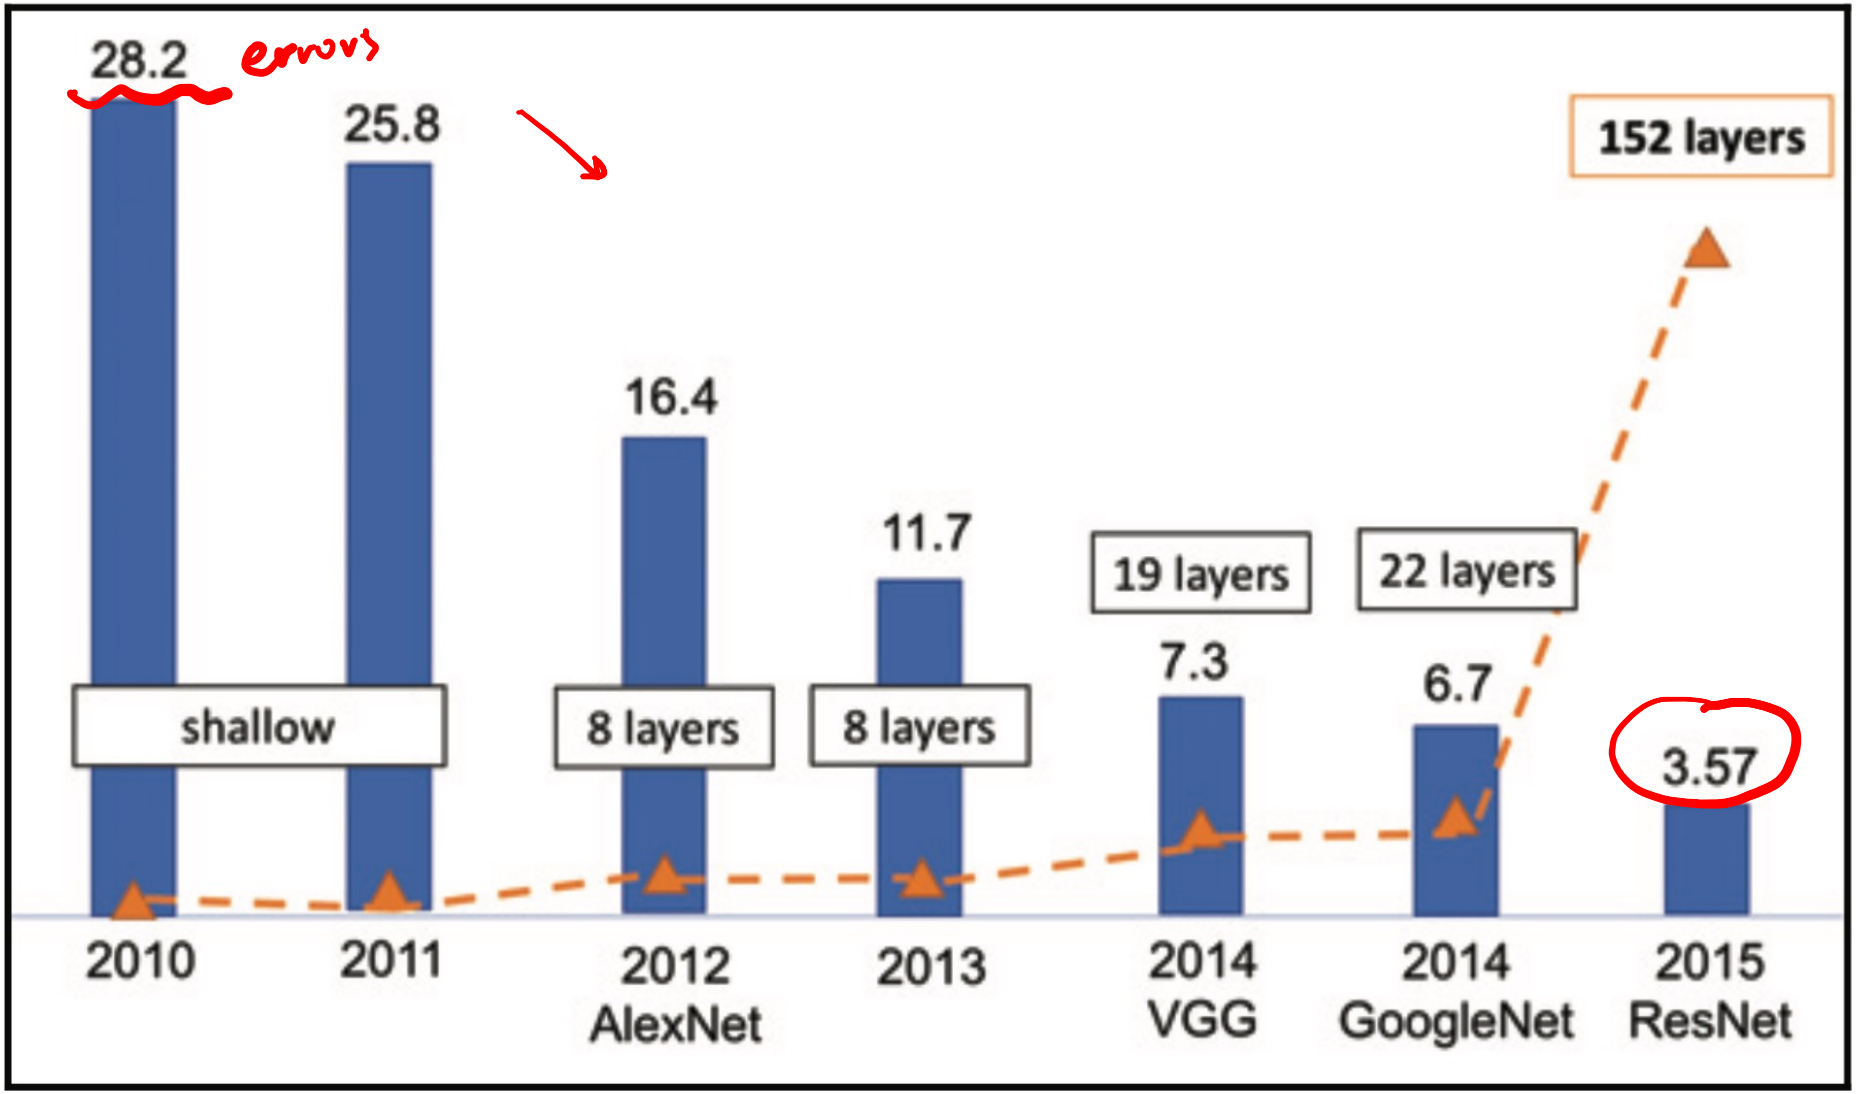

Speech Recognition Progress Over Time

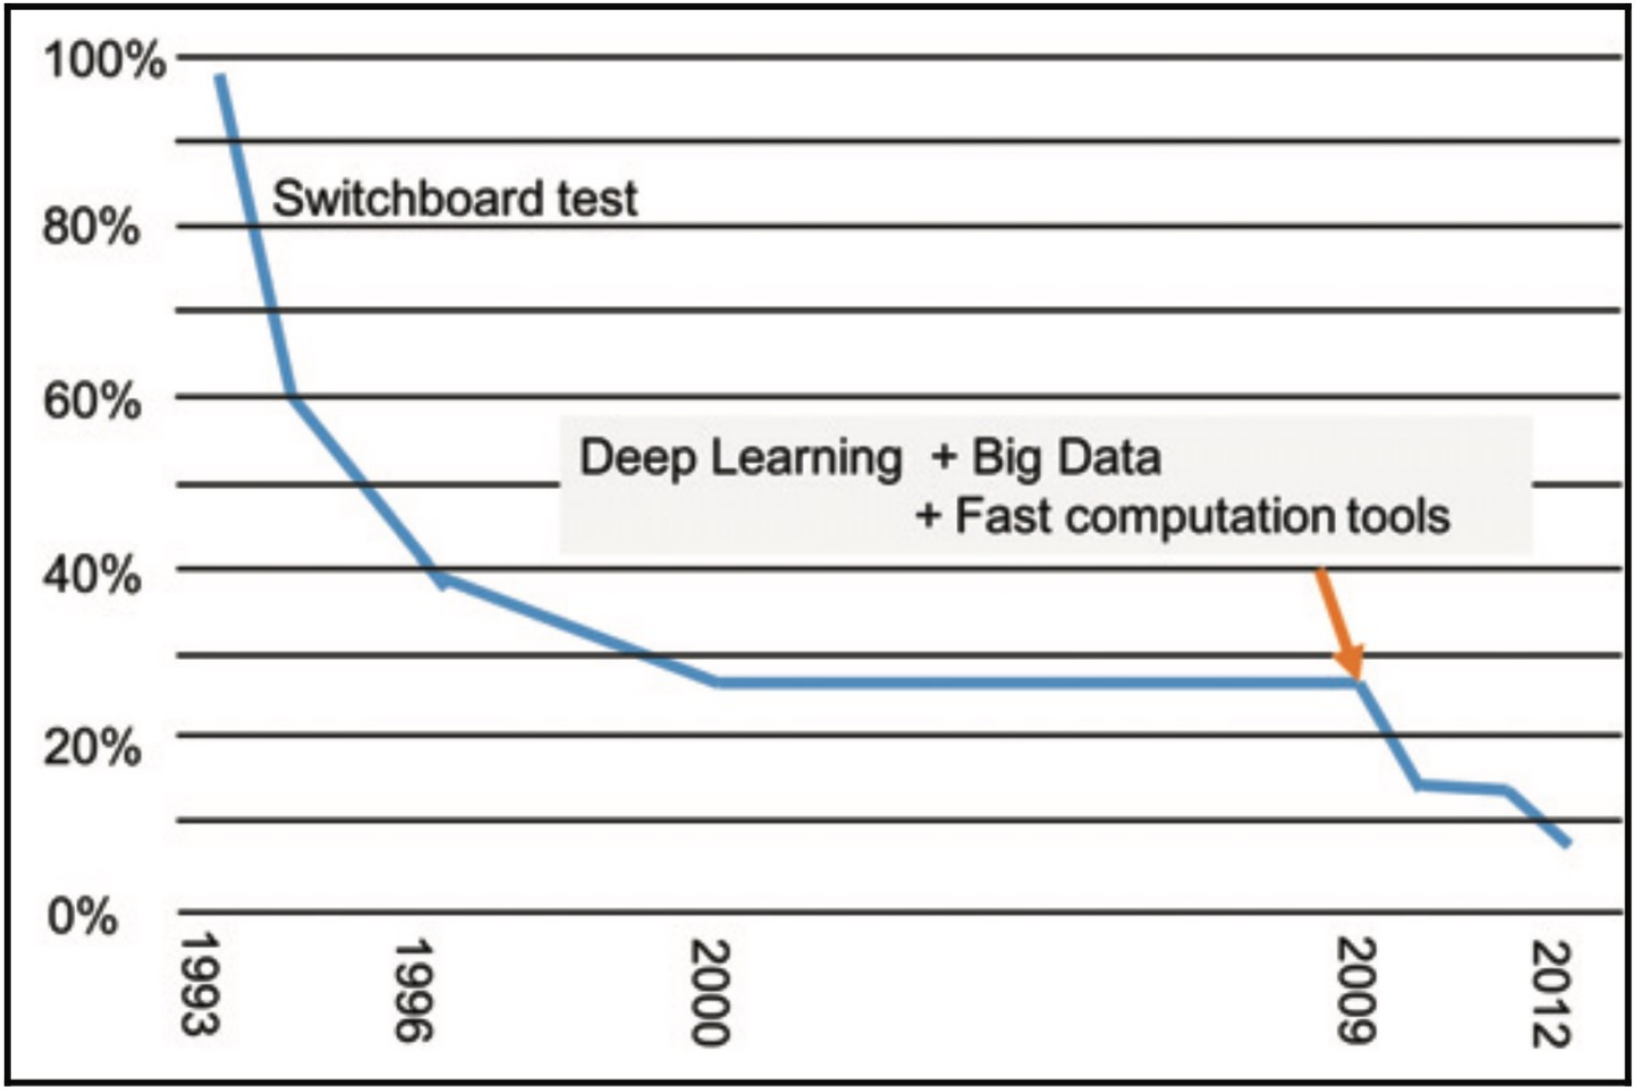

Progress in Artificial Intelligence

https://en.wikipedia.org/wiki/Progress_in_artificial_intelligence

- **super-human**: performs better than all humans: **Go (2017)**
- **high-human**: performs better than most humans: **StarCraft II (2019)**
- **par-human**: performs similarly to most humans: **Image Classification, Handwritten Classification**
- **sub-human**: performs worse than most humans: **Speech Recognition** (nearly equal to human performance)

  Robotics (various robotics tasks that may require advances in robot hardware as well as)

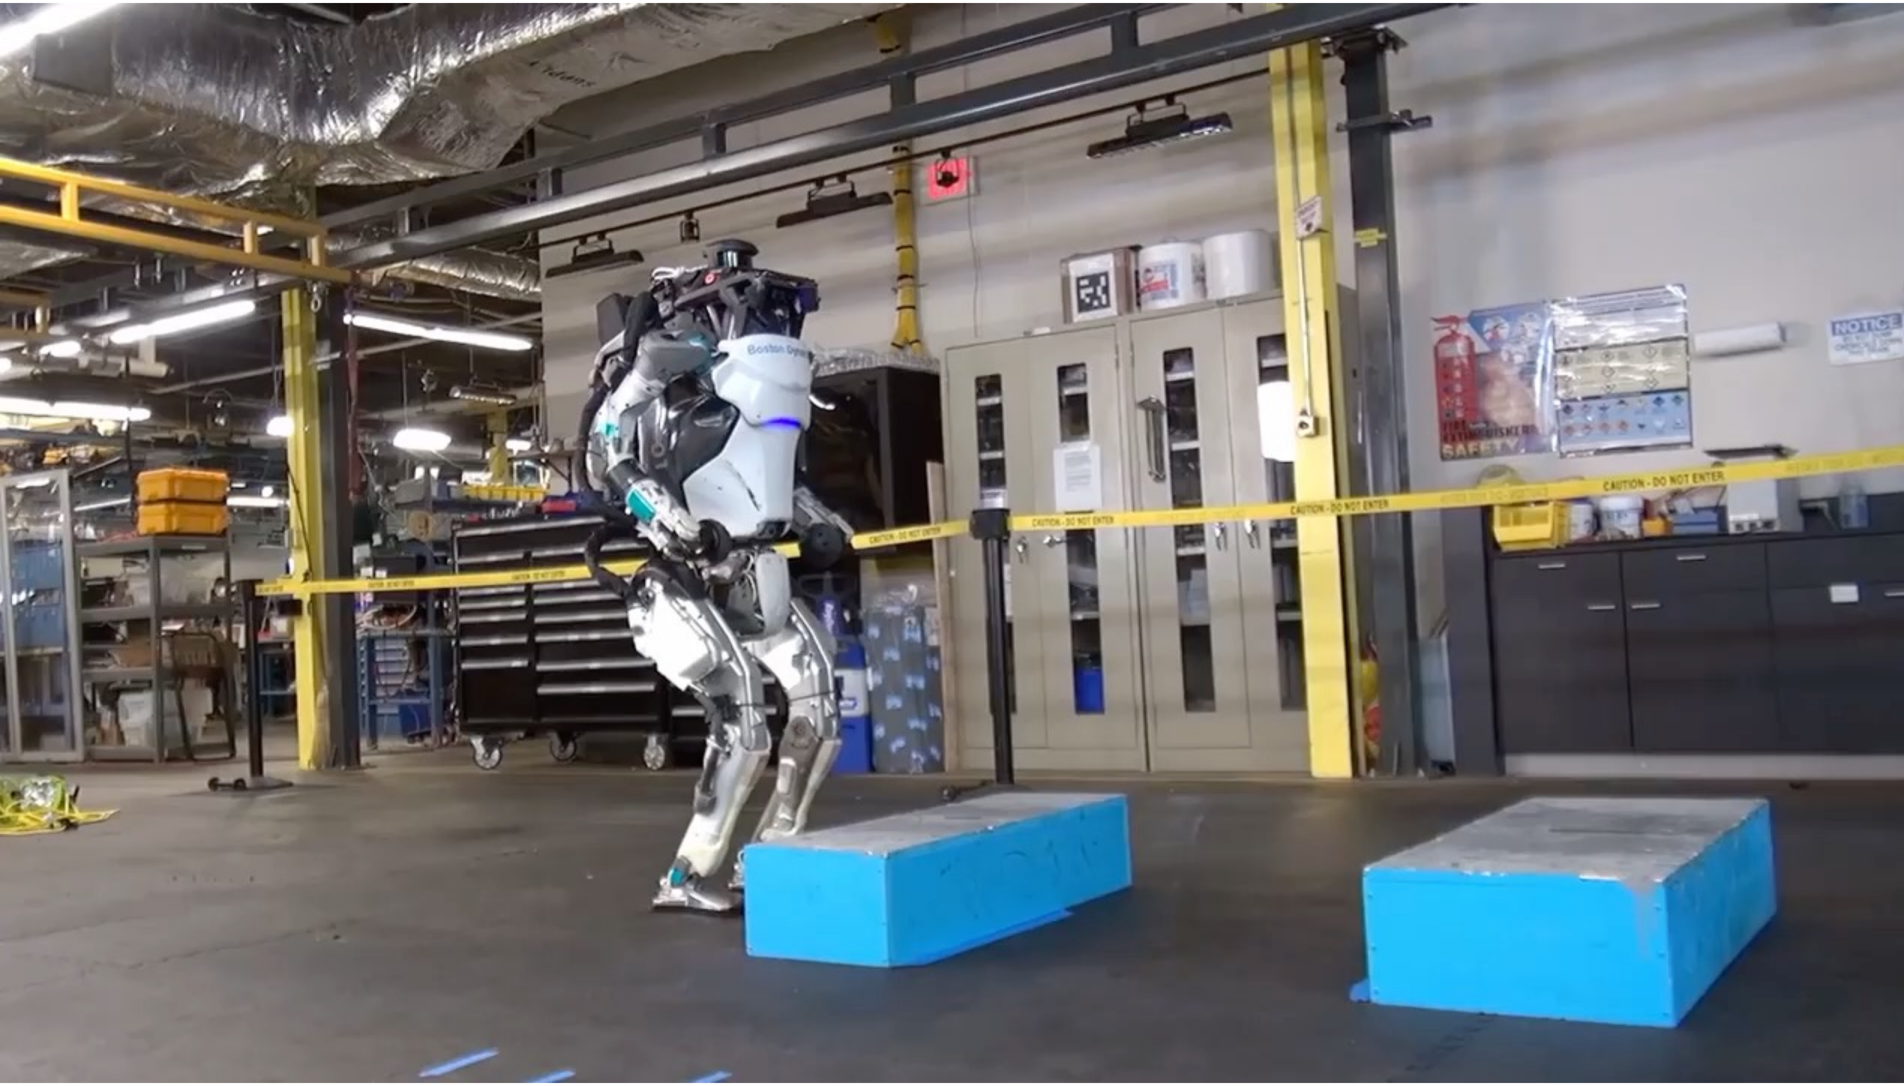

Boston Dynamics 2017

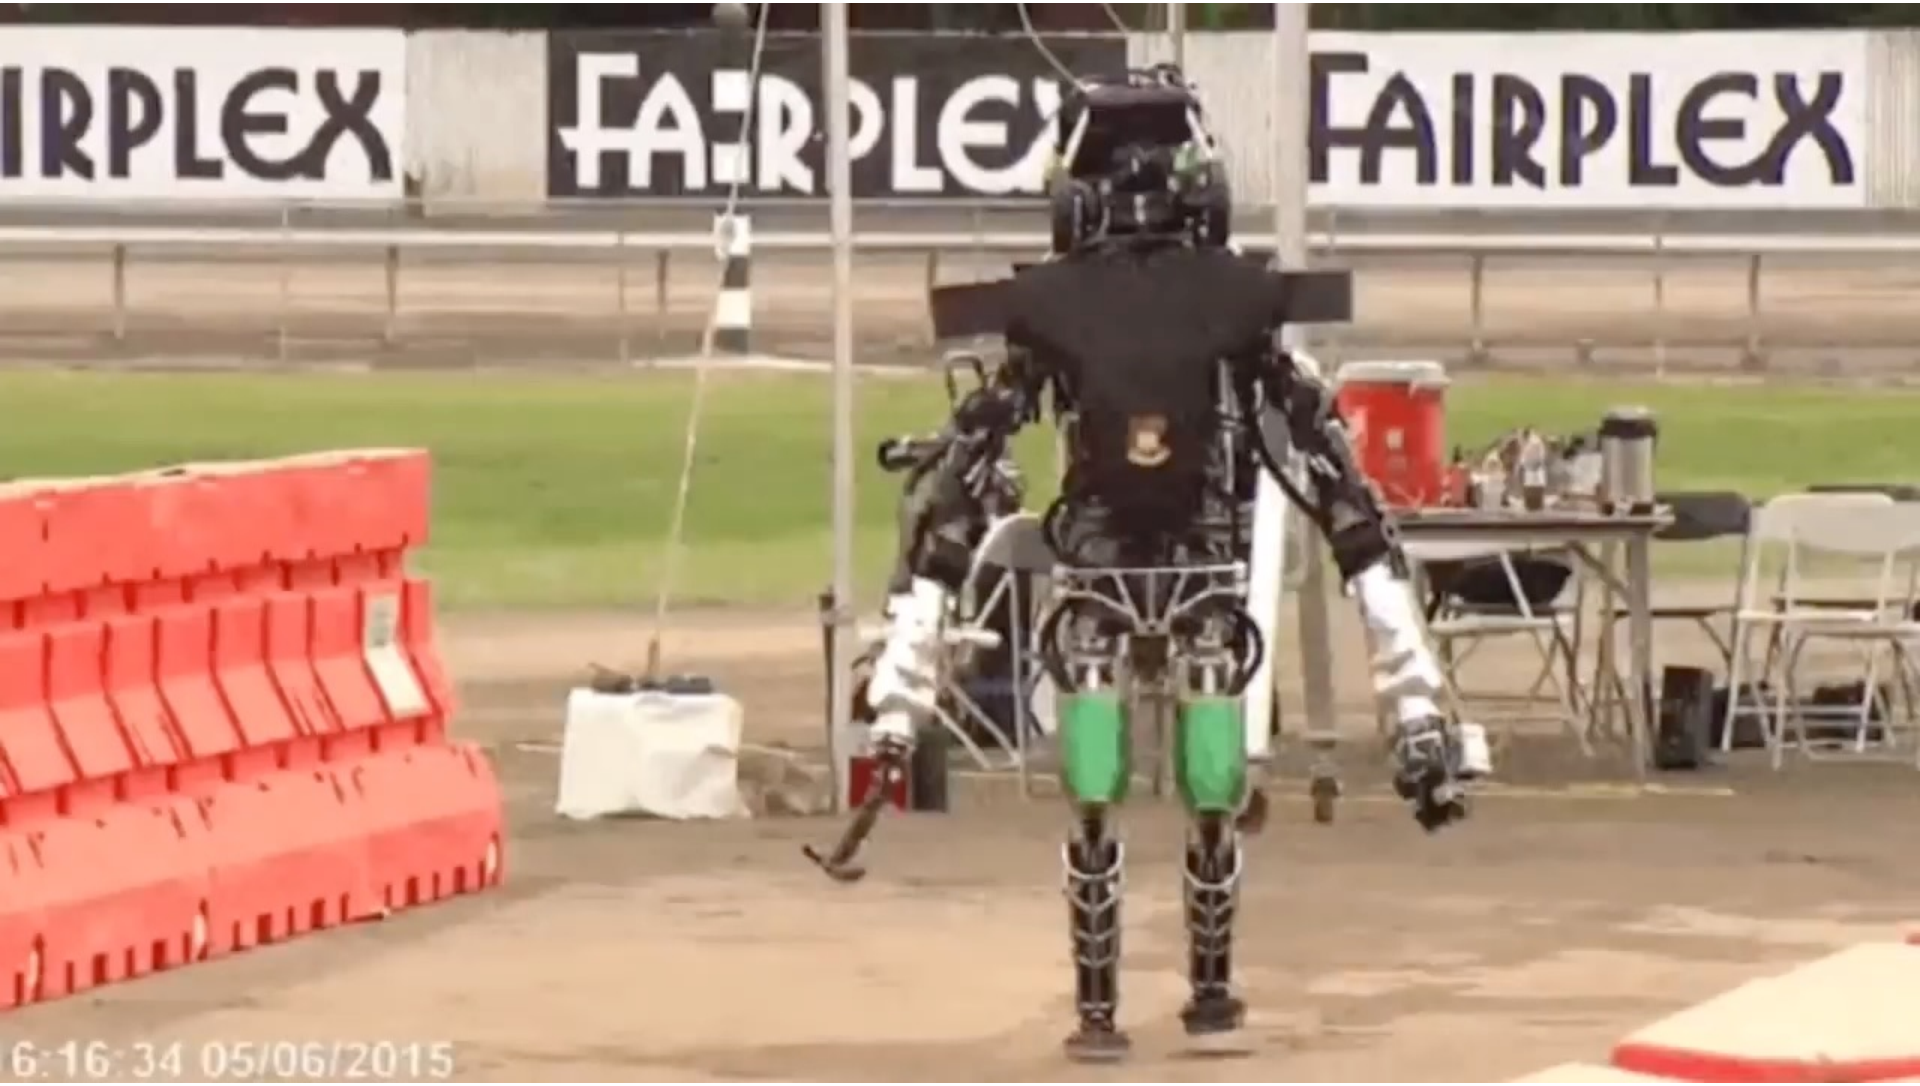

DARPA Robot Competition
https://www.youtube.com/watch?v=g0TaYhjpOfo

Autonomous Indoor Navigation Systems

Jetson Xavier NX

onboard computer to process the sensor data

Pixhawk

flight control unit to control the basic movemont of UAV

Intel Realsense D435

depth image and color image for mapping

Intel Realsense T265

provide high precision UAV position for Indoor environment where GPS is unreachable

Experiment

Indoor Flying Experiment @ CRB 411

Experiment

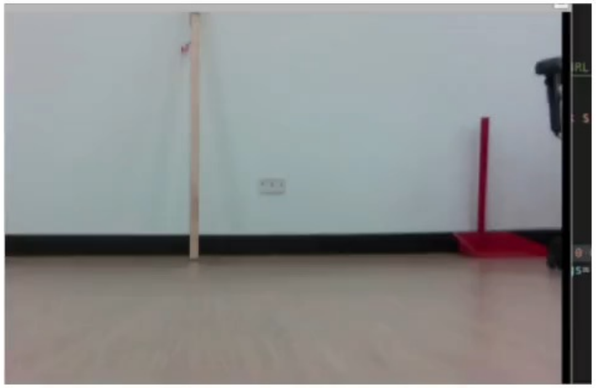

Mask R-CNN

Third-person perspective

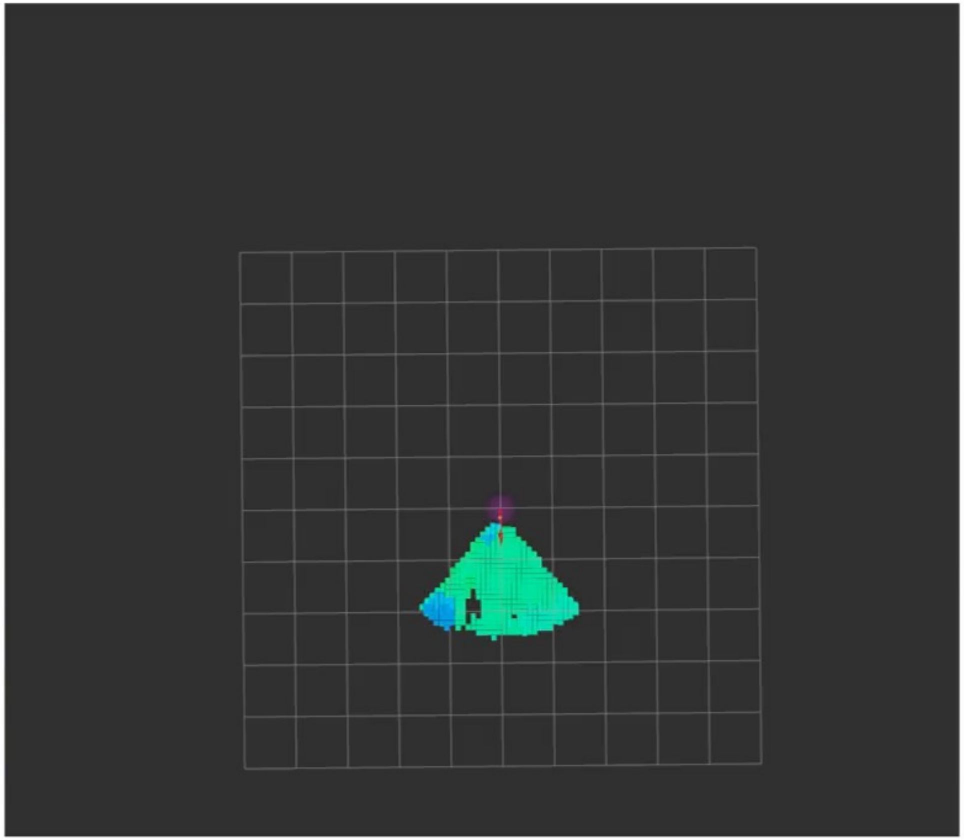

AI pioneer Geoffrey Hinton: “Deep learning is going to be able to do everything.

https://www.technologyreview.com/2020/11/03/1011616/ai-godfather-geoffrey-hinton-deep-learning-will-do-everything/

Q: You think deep learning will be enough to replicate all of human intelligence. What makes you so sure?

A: I do believe deep learning is going to be able to do everything, but I do think there’s going to have to be quite a few conceptual breakthroughs. For example, in 2017 Ashish Vaswani et al. introduced transformers, which derive really good vectors representing word meanings. It was a conceptual breakthrough. It’s now used in almost all the very best natural-language processing. We’re going to need a bunch more breakthroughs like that.

AI pioneer Geoffrey Hinton: “Deep learning is going to be able to do everything.

https://www.technologyreview.com/2020/11/03/1011616/ai-godfather-geoffrey-hinton-deep-learning-will-do-everything/

Q: And if we have those breakthroughs, will we be able to approximate all human intelligence through deep learning?
A: Yes. Particularly breakthroughs to do with how you get big vectors of neural activity to implement things like reason. But we also need a massive increase in scale. The human brain has about 100 trillion parameters, or synapses. What we now call a really big model, like GPT-3, has 175 billion. It’s a thousand times smaller than the brain. GPT-3 can now generate pretty plausible-looking text, and it’s still tiny compared to the brain. **Read more here.**

Why Deep?

Most of the contents are adapted from Hung-Yi Lee, Machine Learning (2017)
http://speech.ee.ntu.edu.tw/~tlkagk/courses_ML17_2.html

Universality Theorem

Any continuous function f

$$f : R^N \to R^M$$

Can be realized by a network with one hidden layer

(given **enough** hidden neurons)

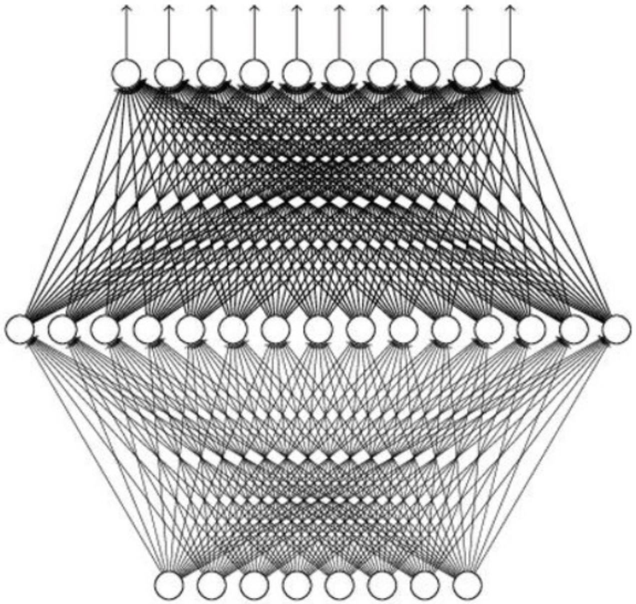

Reference for the reason: http://neuralnetworksanddeeplearning.com/chap4.html

Why "Deep" neural network not "Fat" neural network?

credit: Hung-Yi Lee, http://speech.ee.ntu.edu.tw/~tlkagk/courses_ML17_2.html

#073374

Fun Time: given the same number of parameters for a fat+short vs. thin+tall neural network model, which one will perform better? (1) fat+short (2) thin+tall (3) no difference

Fat + Short v.s. Thin + Tall

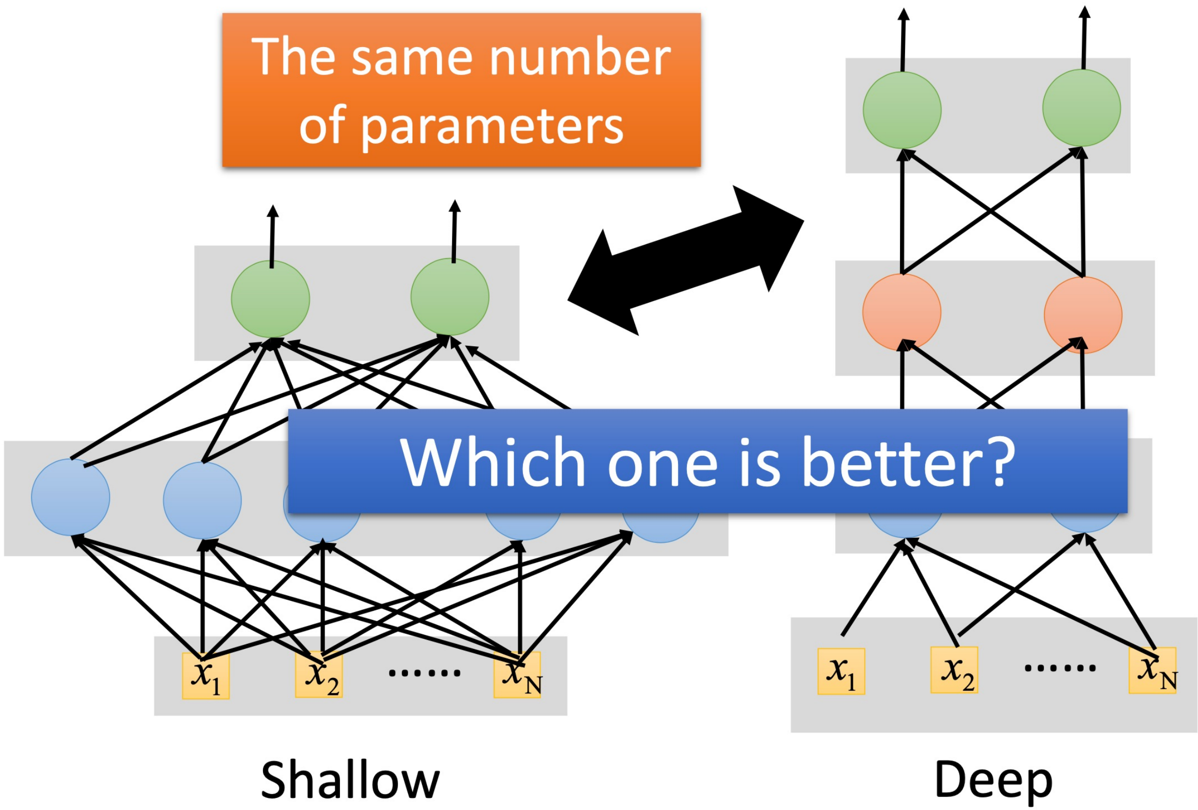

credit: Hung-Yi Lee, http://speech.ee.ntu.edu.tw/~tlkagk/courses_ML17_2.html

Fat + Short v.s. Thin + Tall

Layer X Size	Word Error Rate (%)	Layer X Size	Word Error Rate (%)
1 X 2k	24.2		
2 X 2k	20.4		
3 X 2k	18.4		
4 X 2k	17.8		
5 X 2k	17.2	1 X 3772	22.5
7 X 2k	17.1	1 X 4634	22.6
		1 X 16k	22.1

Why?

Seide, Frank, Gang Li, and Dong Yu. "Conversational Speech Transcription Using Context-Dependent Deep Neural Networks." *Interspeech*. 2011.

credit: Hung-Yi Lee, http://speech.ee.ntu.edu.tw/~tlkagk/courses_ML17_2.html

Classifier 2's performance is NOT ok because of lack of training examples

- Deep → Modularization

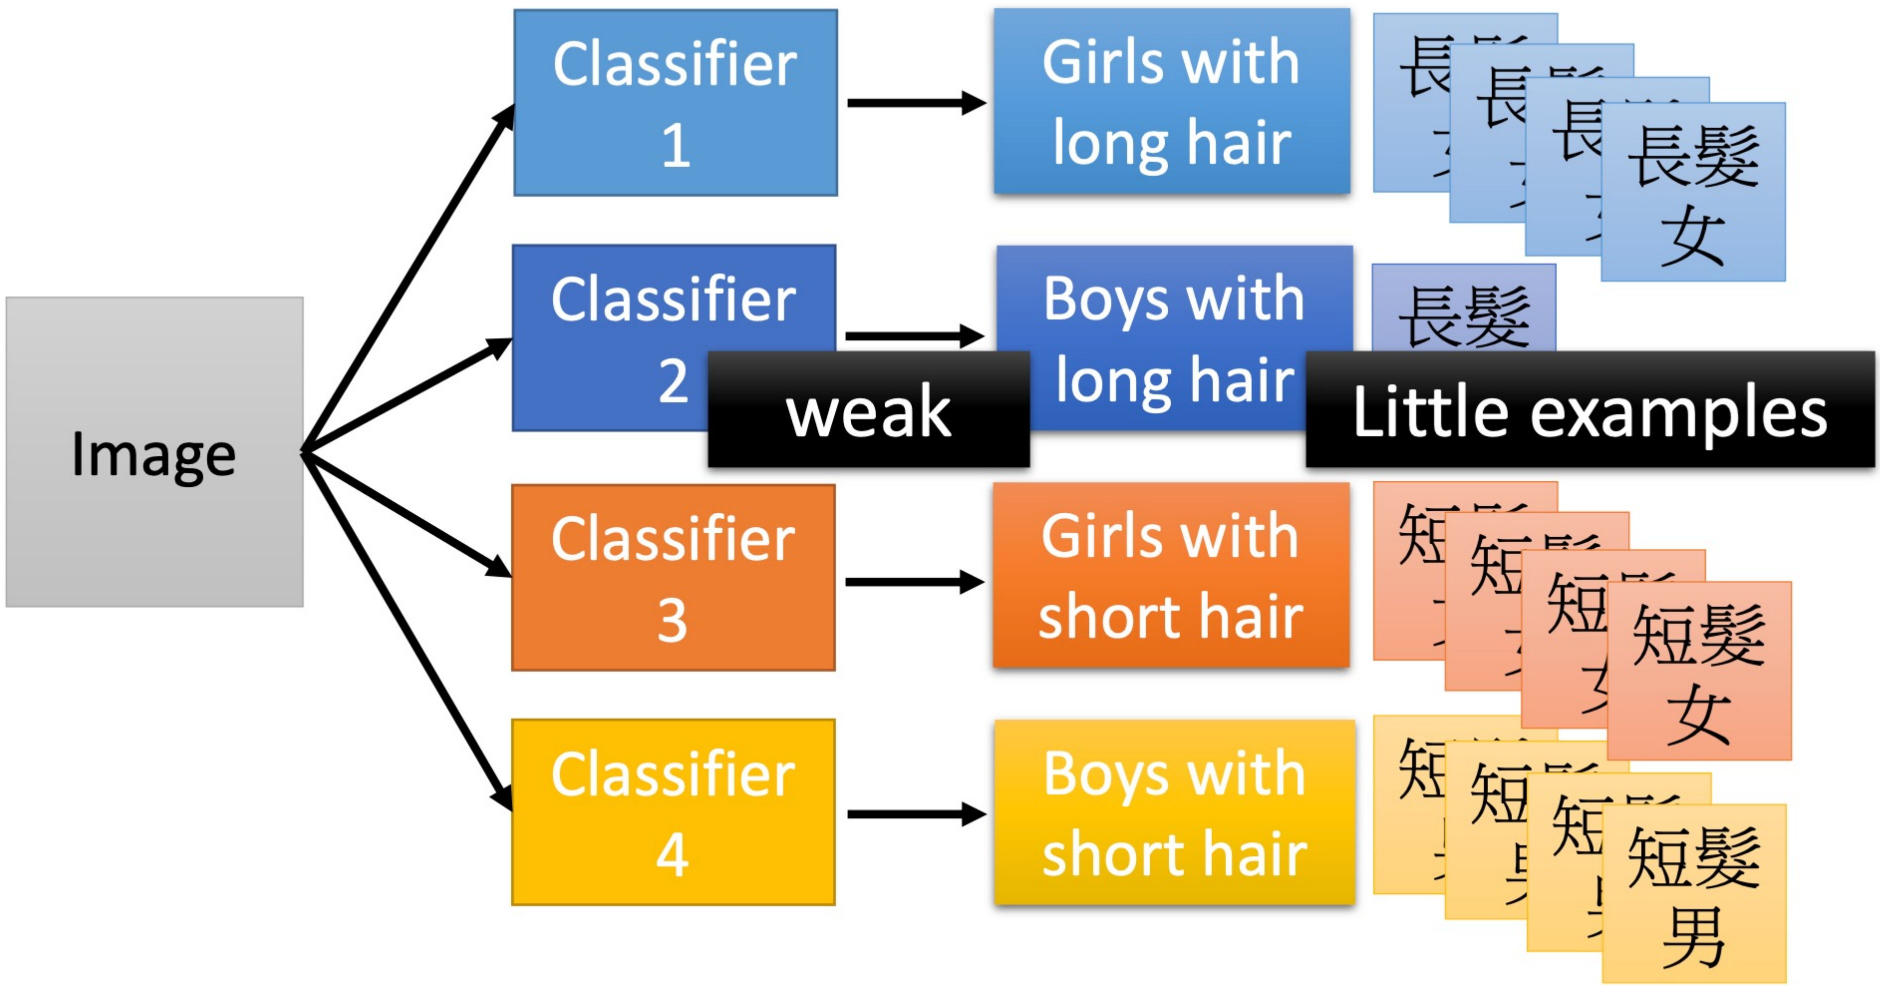

credit: Hung-Yi Lee, http://speech.ee.ntu.edu.tw/~tlkagk/courses_ML17_2.html

Each basic classifier can have sufficient training examples

- Deep → Modularization

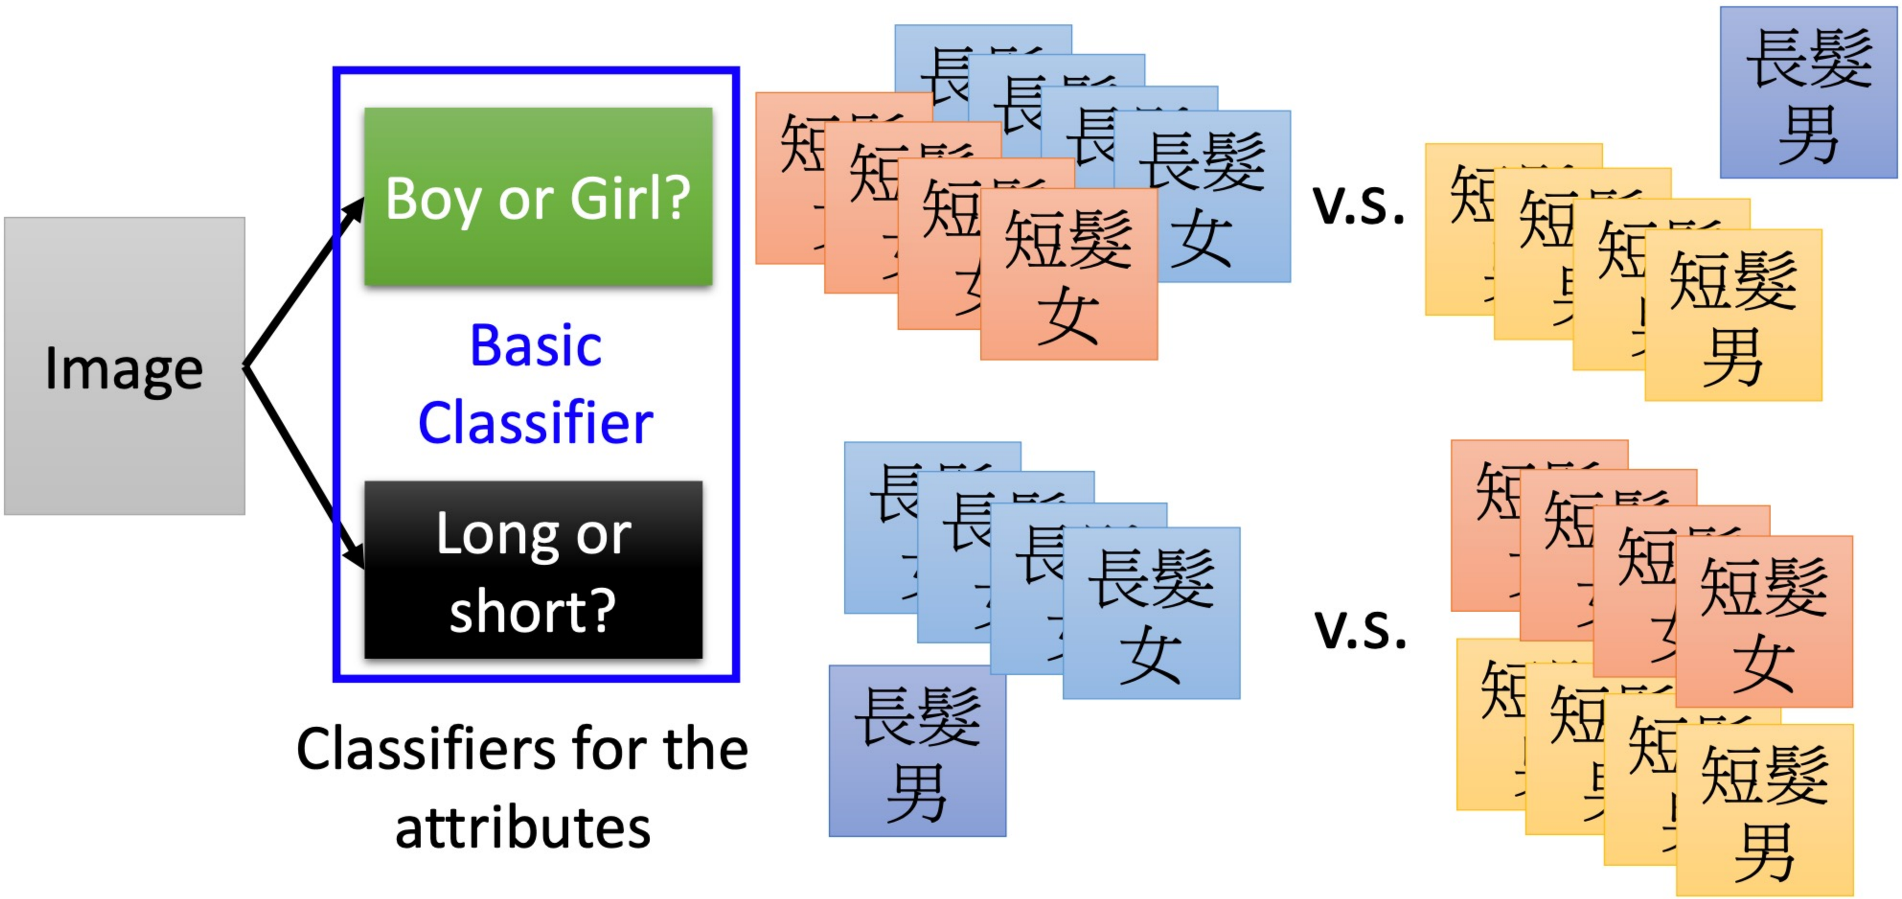

credit: Hung-Yi Lee, http://speech.ee.ntu.edu.tw/~tlkagk/courses_ML17_2.html

Modularization

- Deep → Modularization

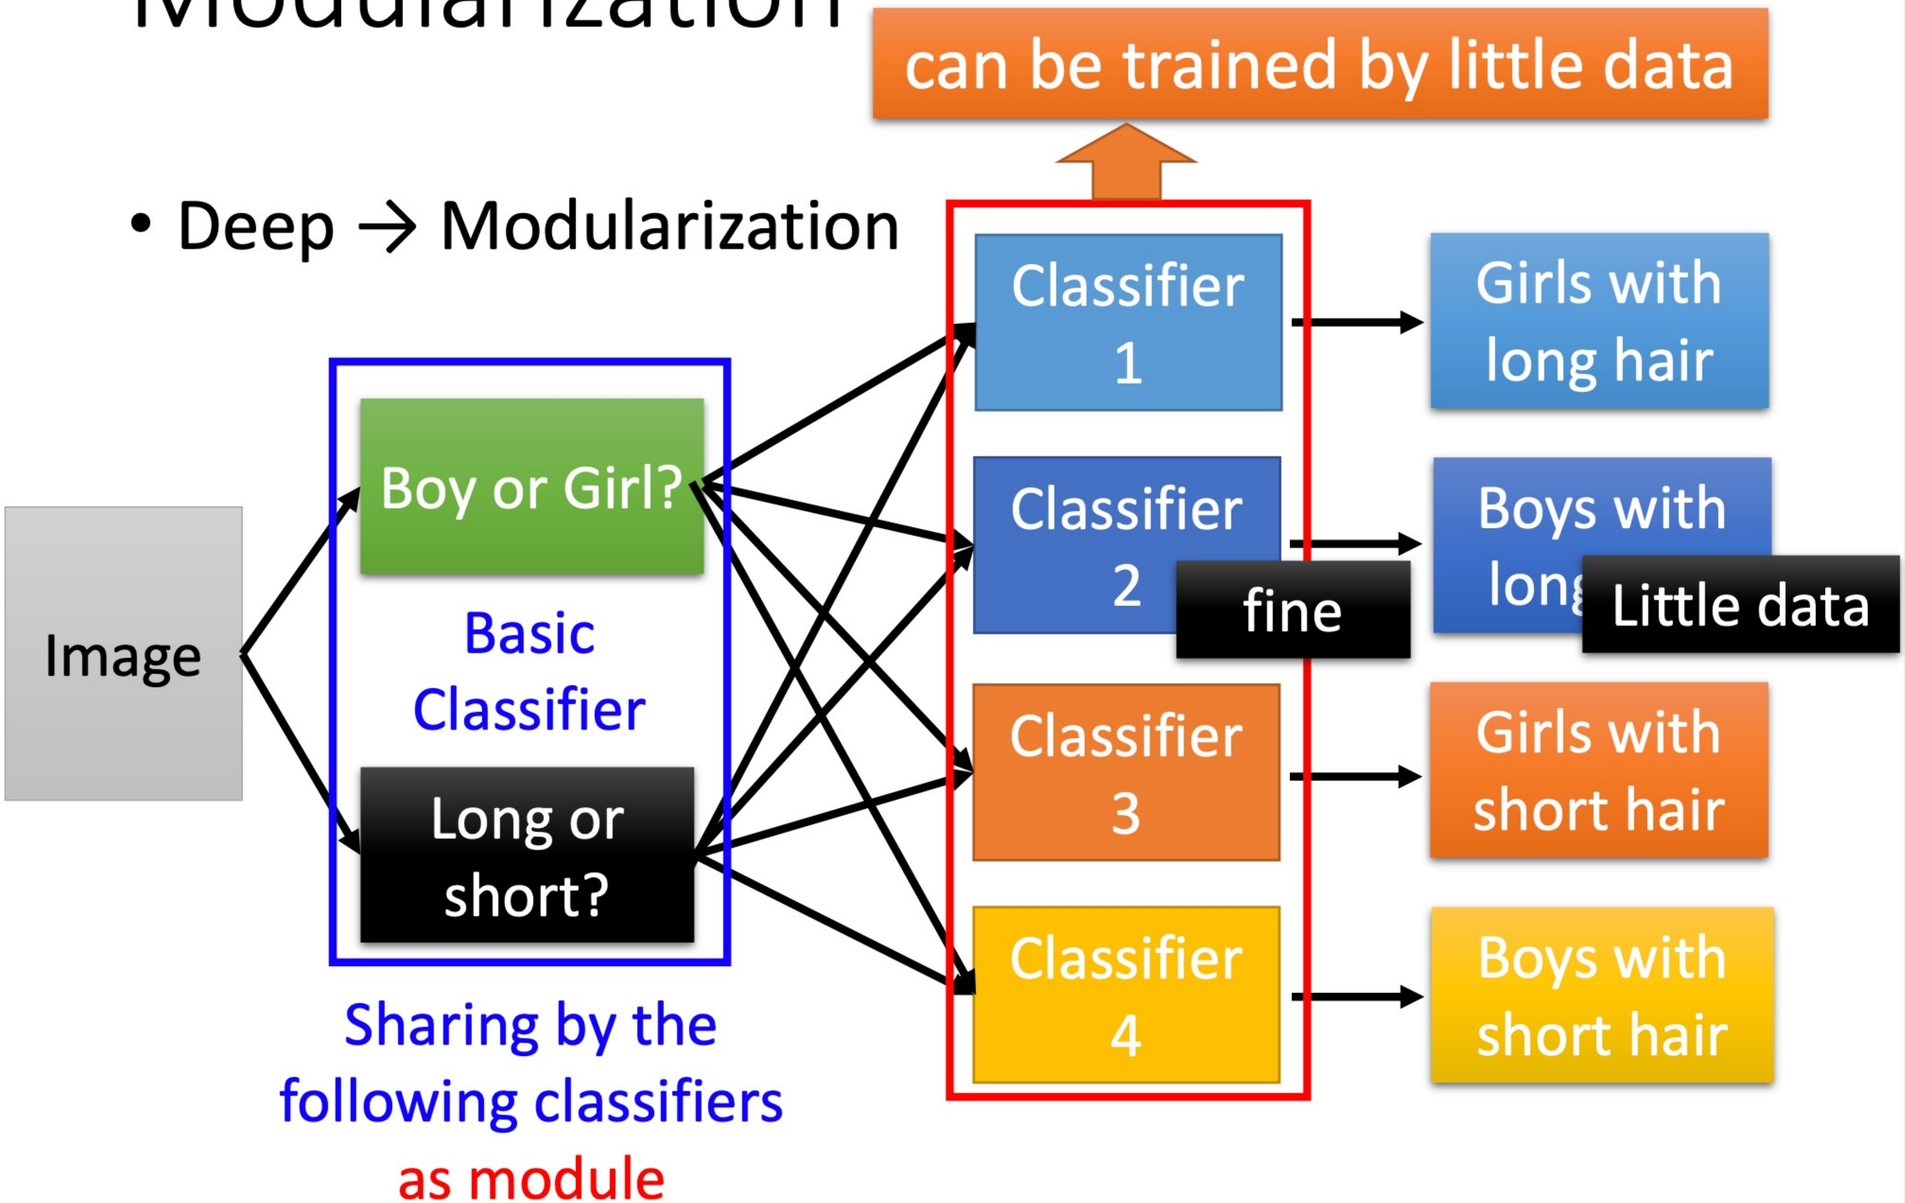

credit: Hung-Yi Lee, http://speech.ee.ntu.edu.tw/~tlkagk/courses_ML17_2.html

Modularization

- Deep → Modularization → Less training data?

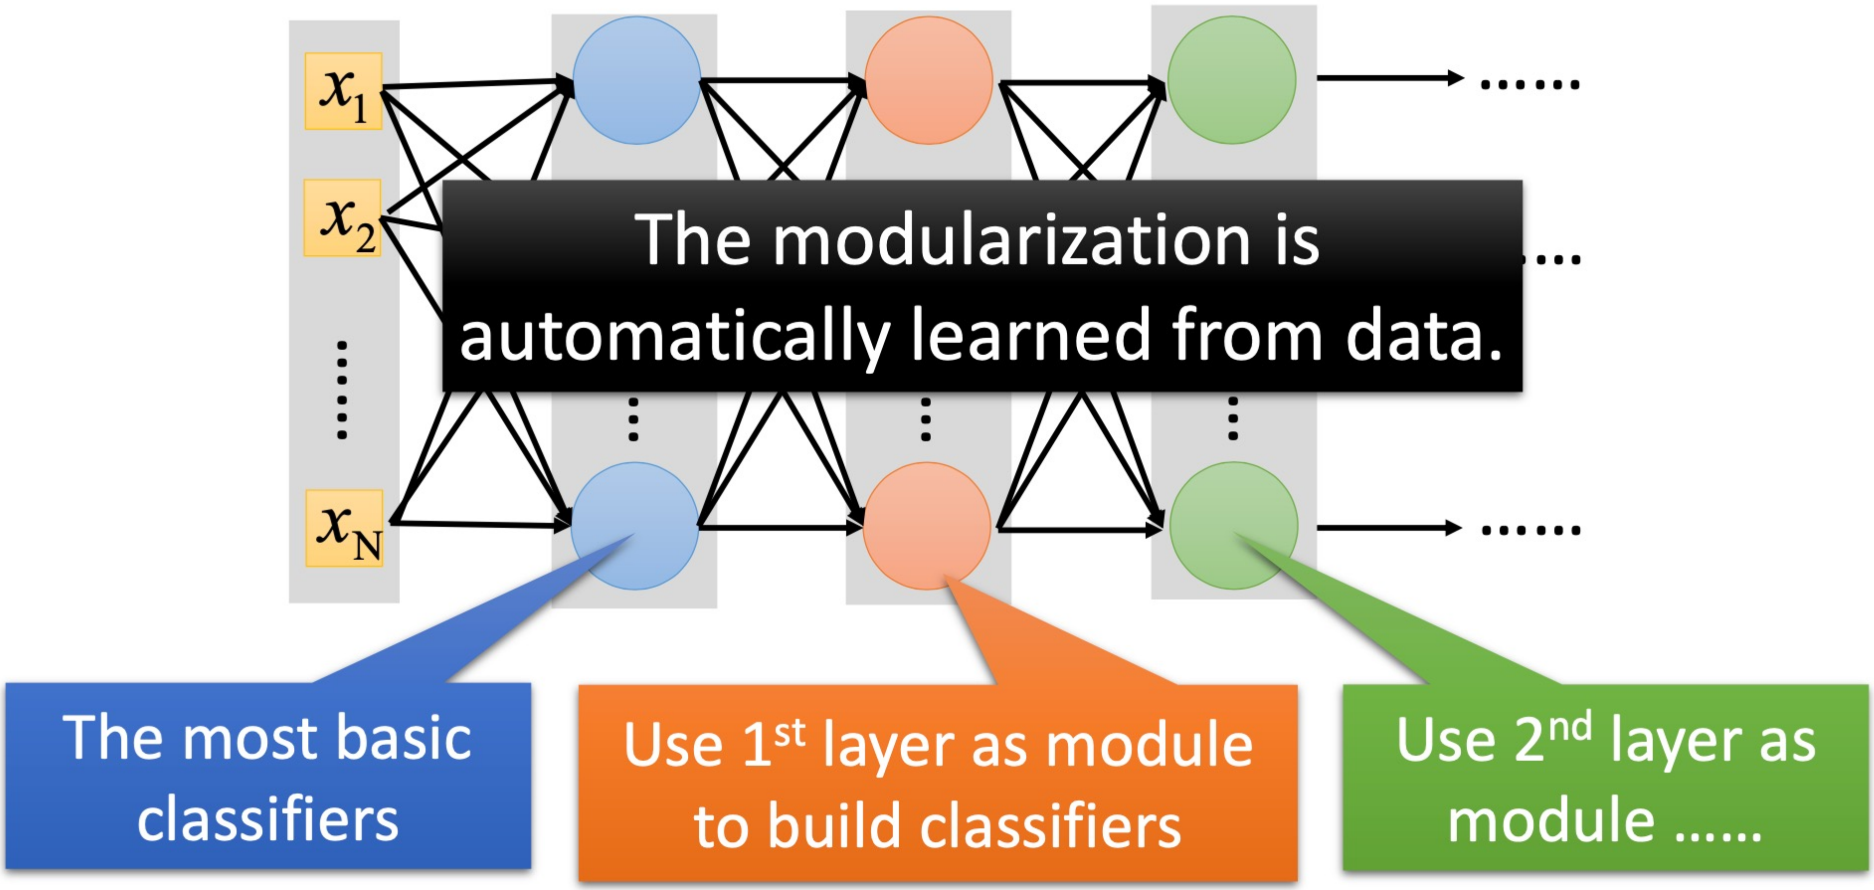

credit: Hung-Yi Lee, http://speech.ee.ntu.edu.tw/~tlkagk/courses_ML17_2.html

Deep Learning: Modularization and Feature Extraction

Deep Learning: Modularization and Feature Extraction

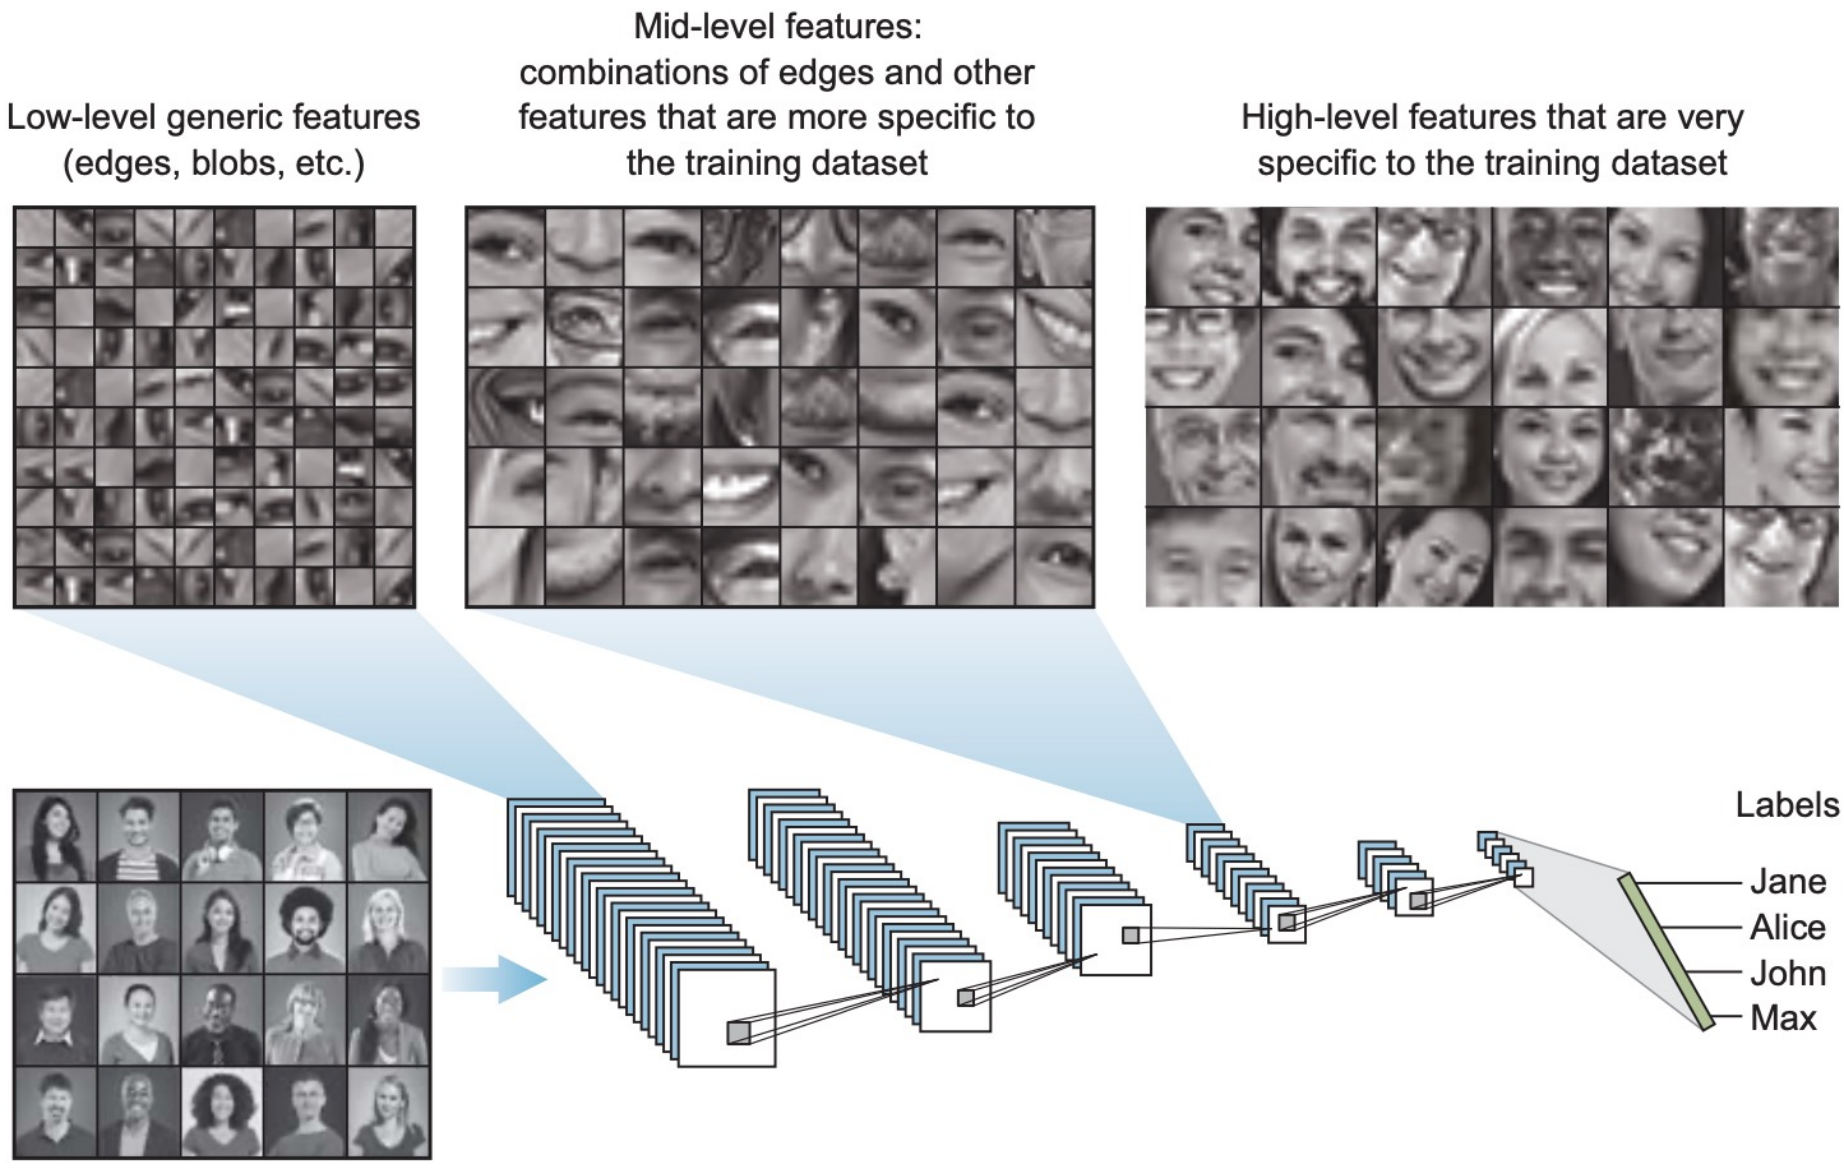

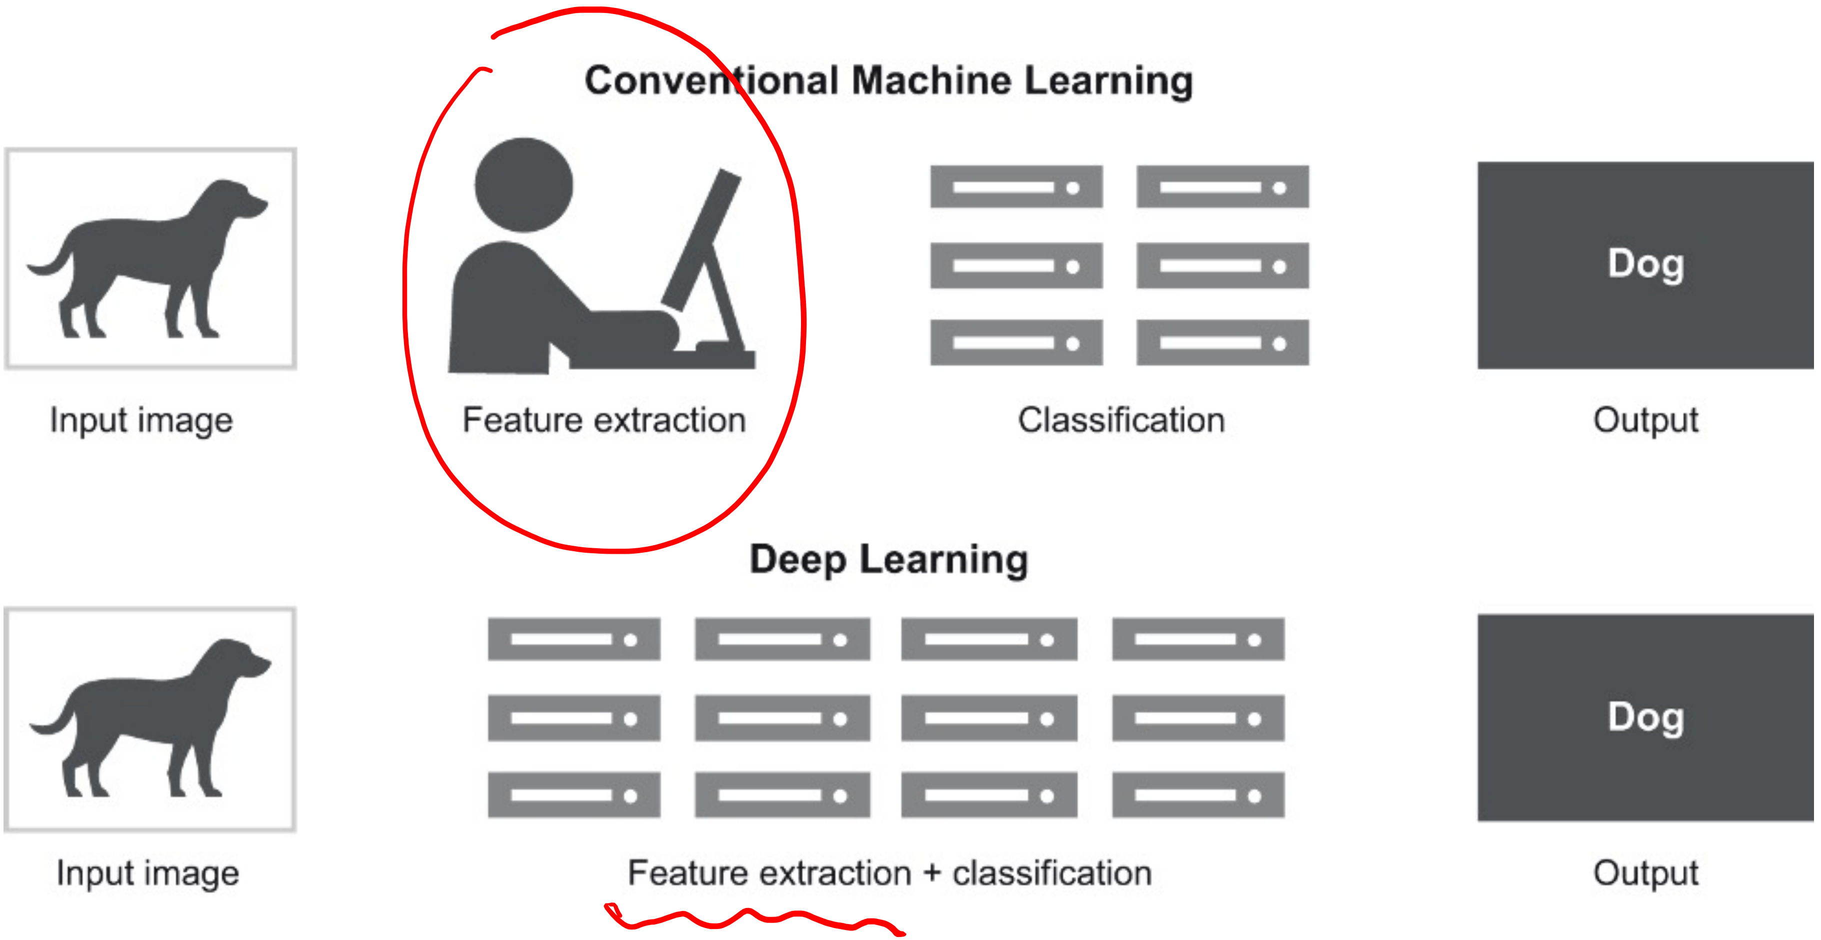
Conventional Machine Learning
Input image
Feature extraction
Classification
Dog
Output
Deep Learning
Input image
Feature extraction + classification
Dog
Output

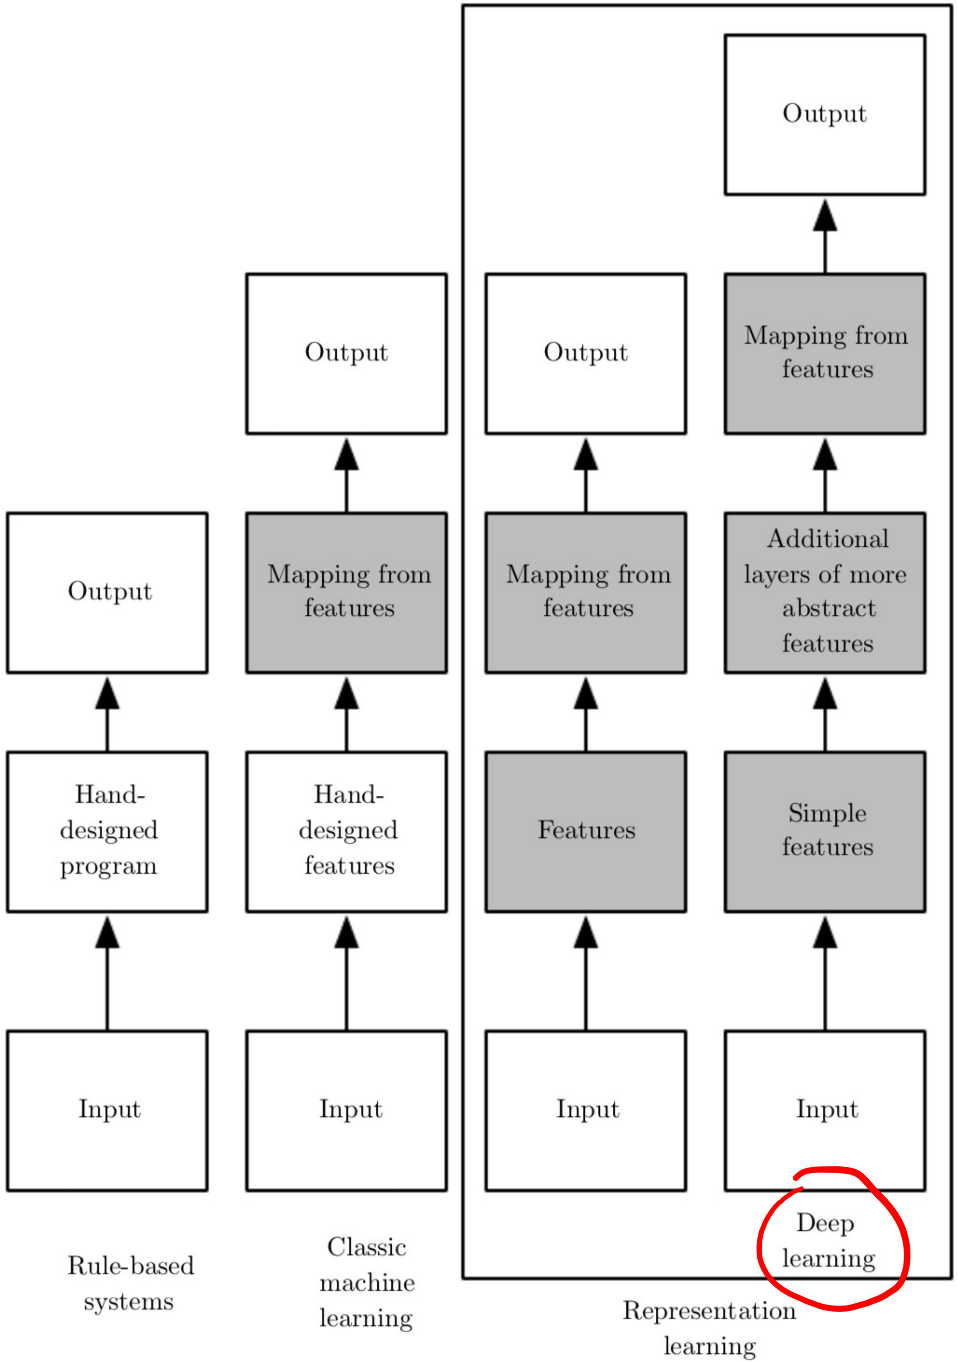
Output
Output
Output
Output
Mapping from features
Mapping from features
Mapping from features
Additional layers of more abstract features
Hand-designed program
Hand-designed features
Features
Simple features
Input
Input
Input
Input
Rule-based systems
Classic machine learning
Representation learning
Deep learning

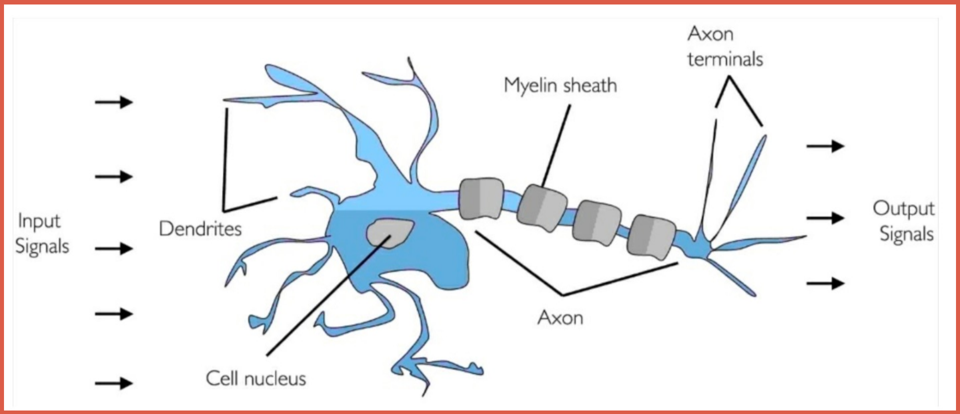

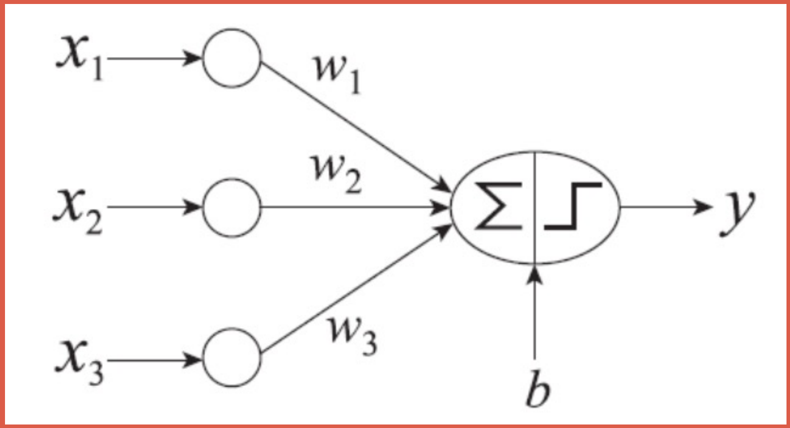

ANN (Artificial Neural Network): Perceptron

C-S David Chen, Department of Civil Engineering, National Taiwan University

(Artificial) Neural Networks (ANN): The Heart of Deep Learning

ANN is a collection of simple processing units (nodes) that are connected by directed links (edges)

- Every node receives signals from incoming edges, performs computations, and transmits signals to outgoing edges
- Analogous to *human brain* where **nodes are neurons** and signals are electrical impulses
- **Weight** of an edge determines the strength of connection between the nodes

Levels	Model	Group
1-Layer	Perceptron	Linear classifier
2-Layers	Multi-layer perceptron	Universal approximator
3 or more layers	Deep learning	Compact universal approximator

Perceptron: theoretical minimum and example

- The phrase “theoretical minimum” is taken from a very successful book series written by Leonard Susskind, a great physicist at Stanford University.
- “Theoretical minimum” means <u>just the minimum theories and equations you need to know in order to proceed to the next level.</u>
- See Perceptron.pdf

Nonlinearly Separable Data

For nonlinearly separable problems, perceptron learning algorithm will fail because no linear hyperplane can separate the data perfectly

$$y = x_1 \oplus x_2$$

x_1	x_2	y
0	0	-1
1	0	1
0	1	1
1	1	-1

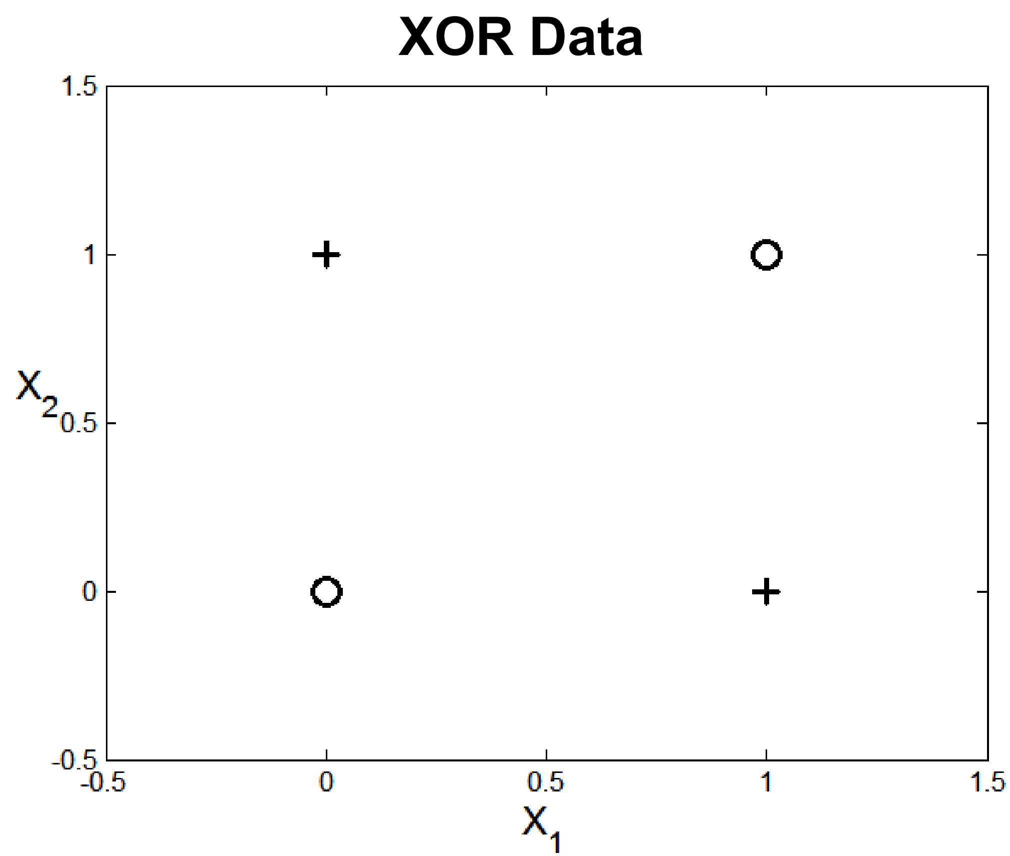

Brief History of Neural Networks

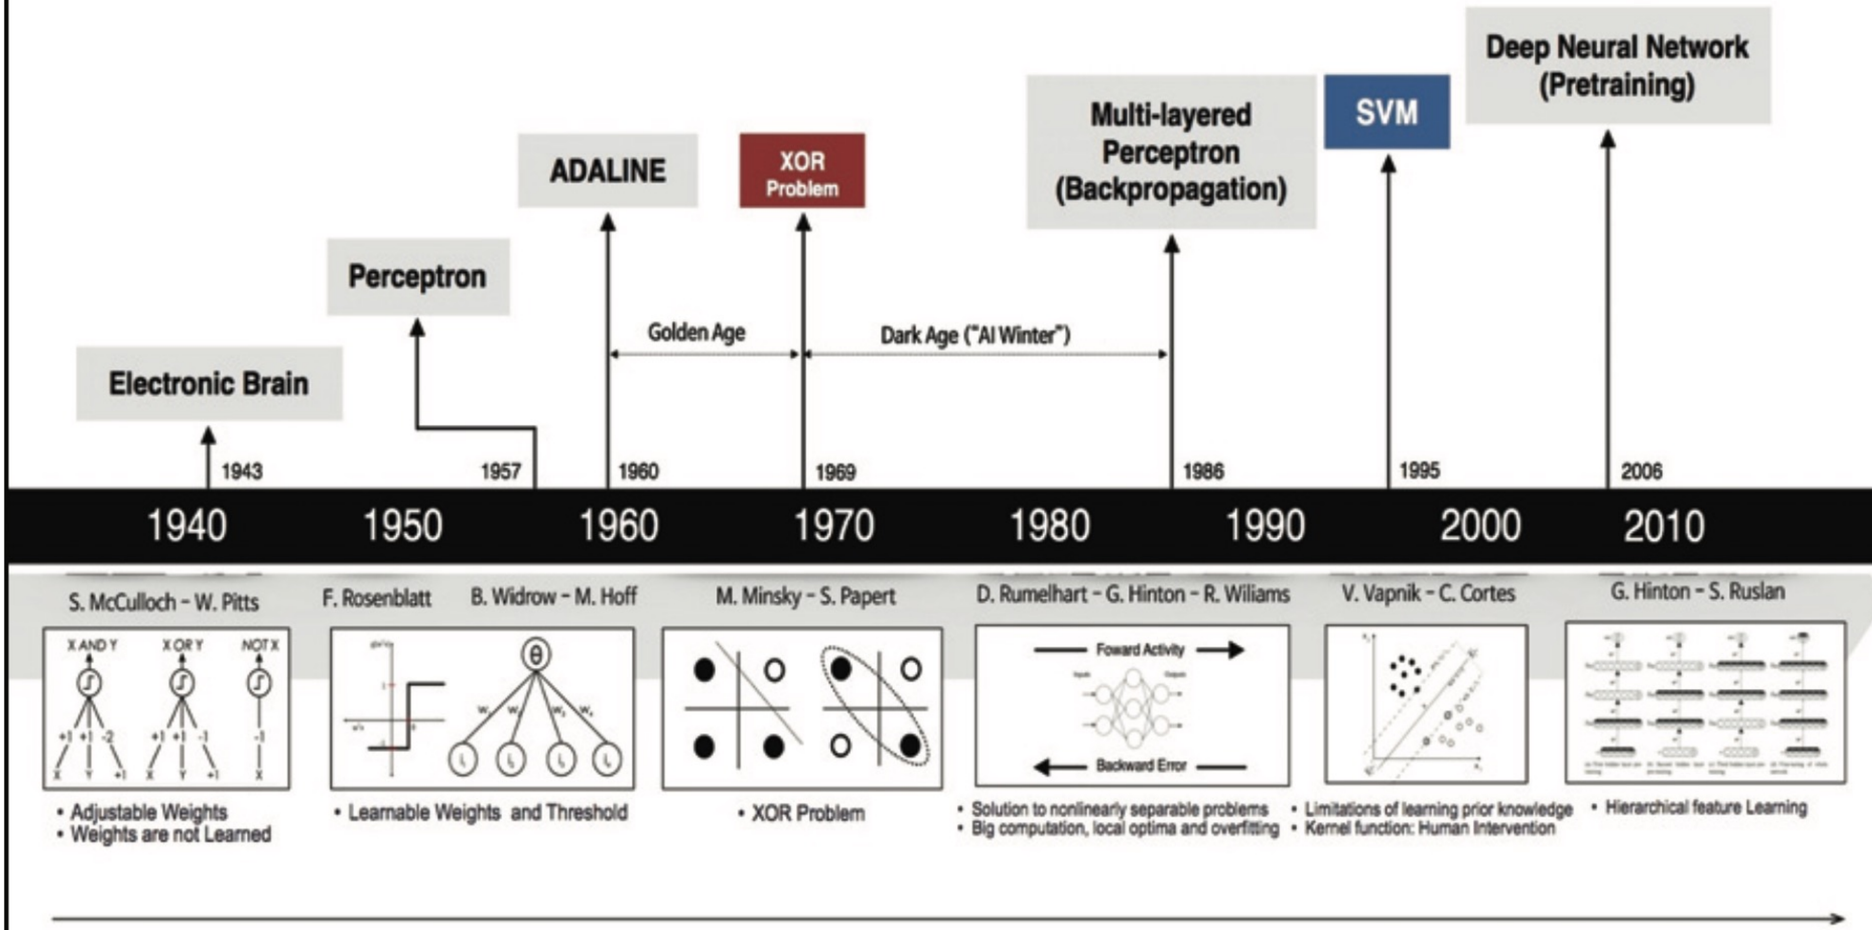

Brief history of neural network