Deep Machine Learning

Department of Computer Engineering
College of Engineering
University of Diyala

٢٠١٨-٢٠١٧

Deep Learning

ARTIFICIAL INTELLIGENCE

Early artificial intelligence stirs excitement.

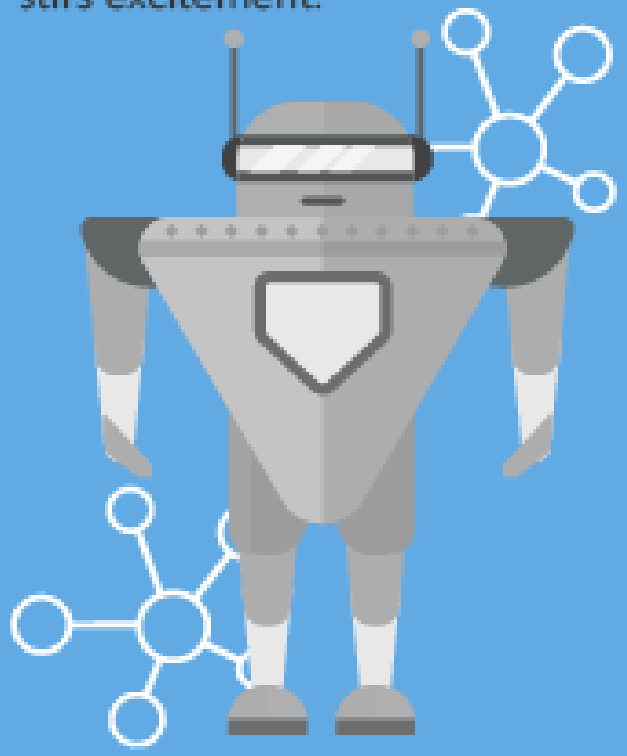

MACHINE LEARNING

Machine learning begins to flourish.

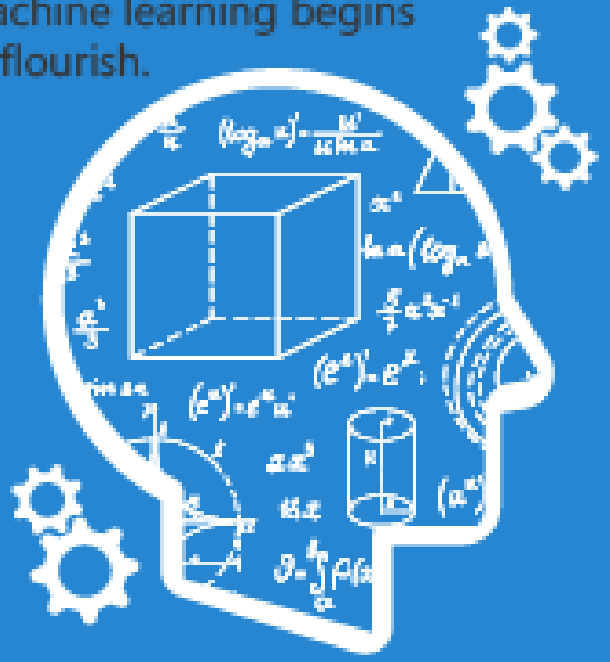

DEEP LEARNING

Deep learning breakthroughs drive AI boom.

1950's 1960's 1970's 1980's 1990's 2000's 2010's

Since an early flush of optimism in the 1950's, smaller subsets of artificial intelligence – first machine learning, then deep learning, a subset of machine learning – have created ever larger disruptions.

ARTIFICIAL INTELLIGENCE 101

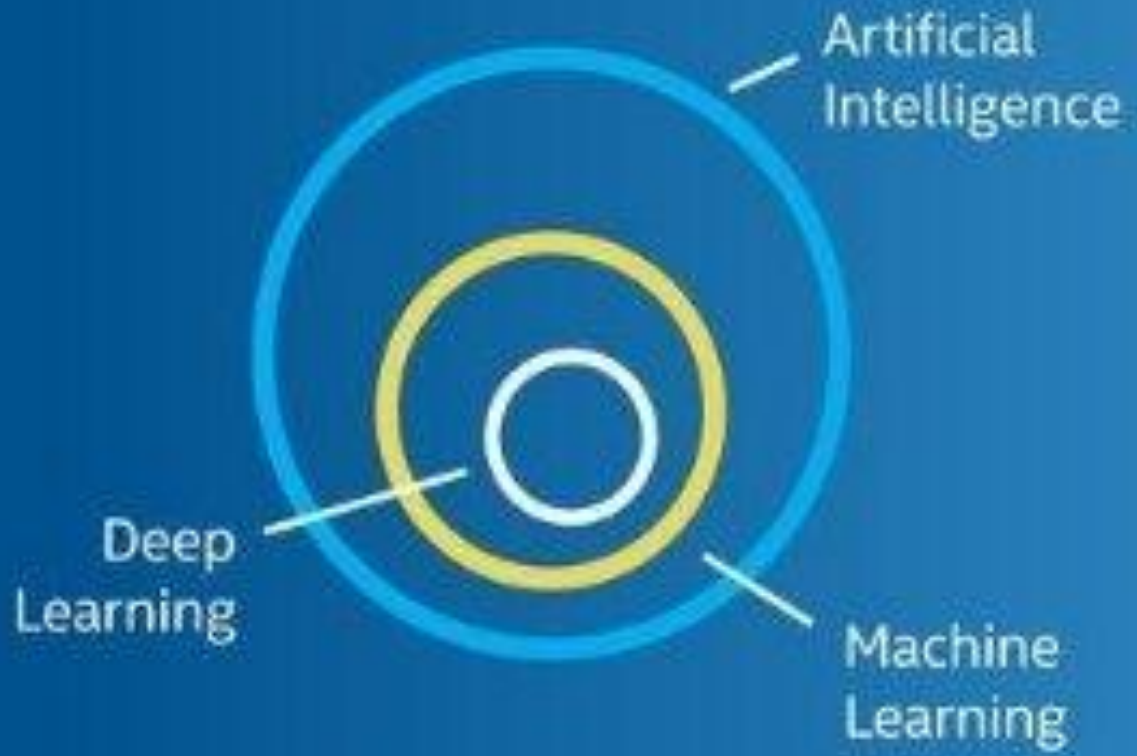

Deep learning is a branch of ML that uses neural network models to understand large amounts of data. It can accelerate processes like image and speech recognition, and natural language recognition.

#ITCenter

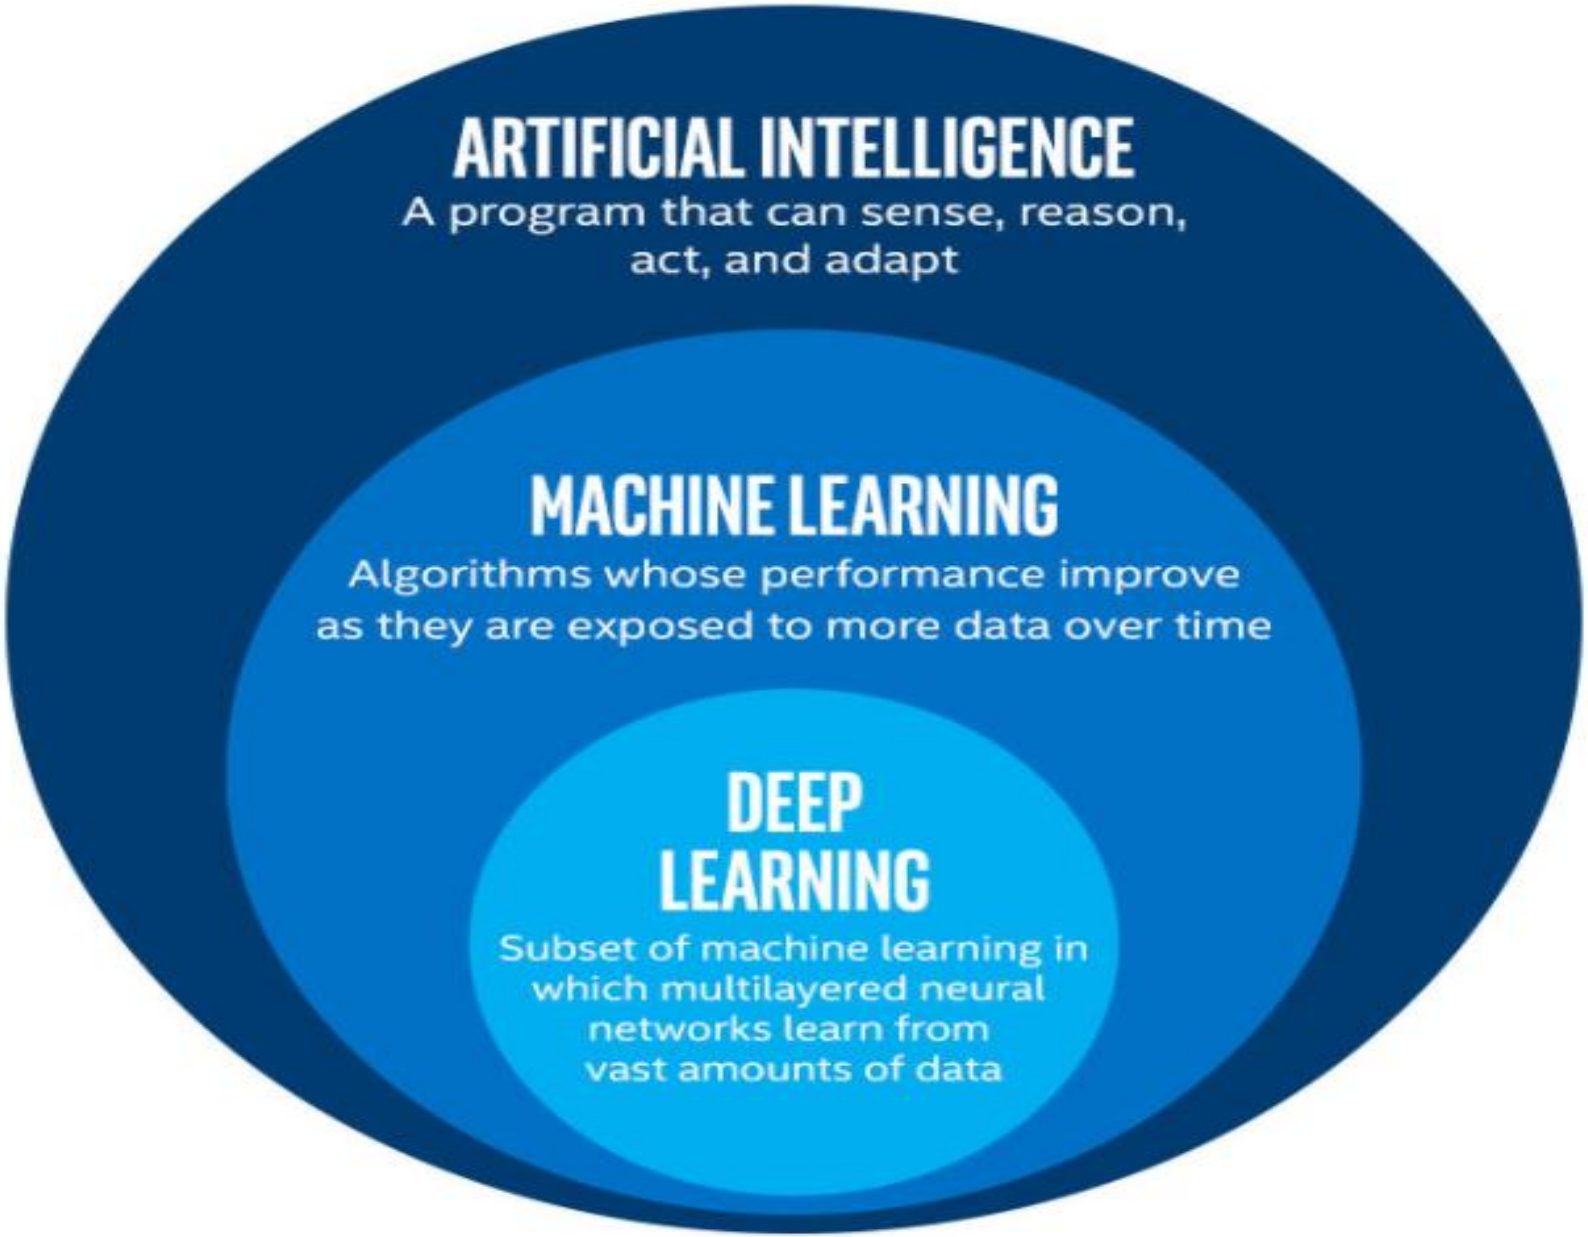
ARTIFICIAL INTELLIGENCE
A program that can sense, reason, act, and adapt
MACHINE LEARNING
Algorithms whose performance improve as they are exposed to more data over time
DEEP LEARNING
Subset of machine learning in which multilayered neural networks learn from vast amounts of data

Machine Learning

Input: X

Output: Y

Label"motorcycle"

Why is it hard?

You see this

But the camera sees this:

194	210	201	212	199	213	215	195	178	158	182	209
180	189	190	221	209	205	191	167	147	115	129	163
114	126	140	188	176	165	152	140	170	106	78	88
87	103	115	154	143	142	149	153	173	101	57	57
102	112	106	131	122	138	152	147	128	84	58	66
94	95	79	104	105	124	129	113	107	87	69	67
68	71	69	98	89	92	98	95	89	88	76	67
41	56	68	99	63	45	60	82	58	76	75	65
20	43	69	75	56	41	51	73	55	70	63	44
50	50	57	69	75	75	73	74	53	68	59	37
72	59	53	66	84	92	84	74	57	72	63	42
67	61	58	65	75	78	76	73	59	75	69	50

Raw Image Representation

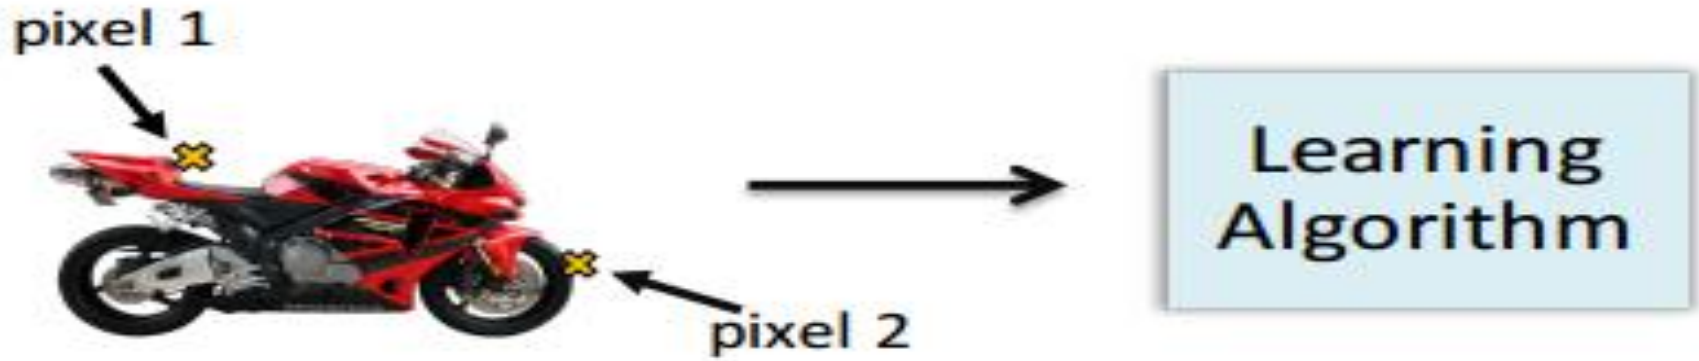

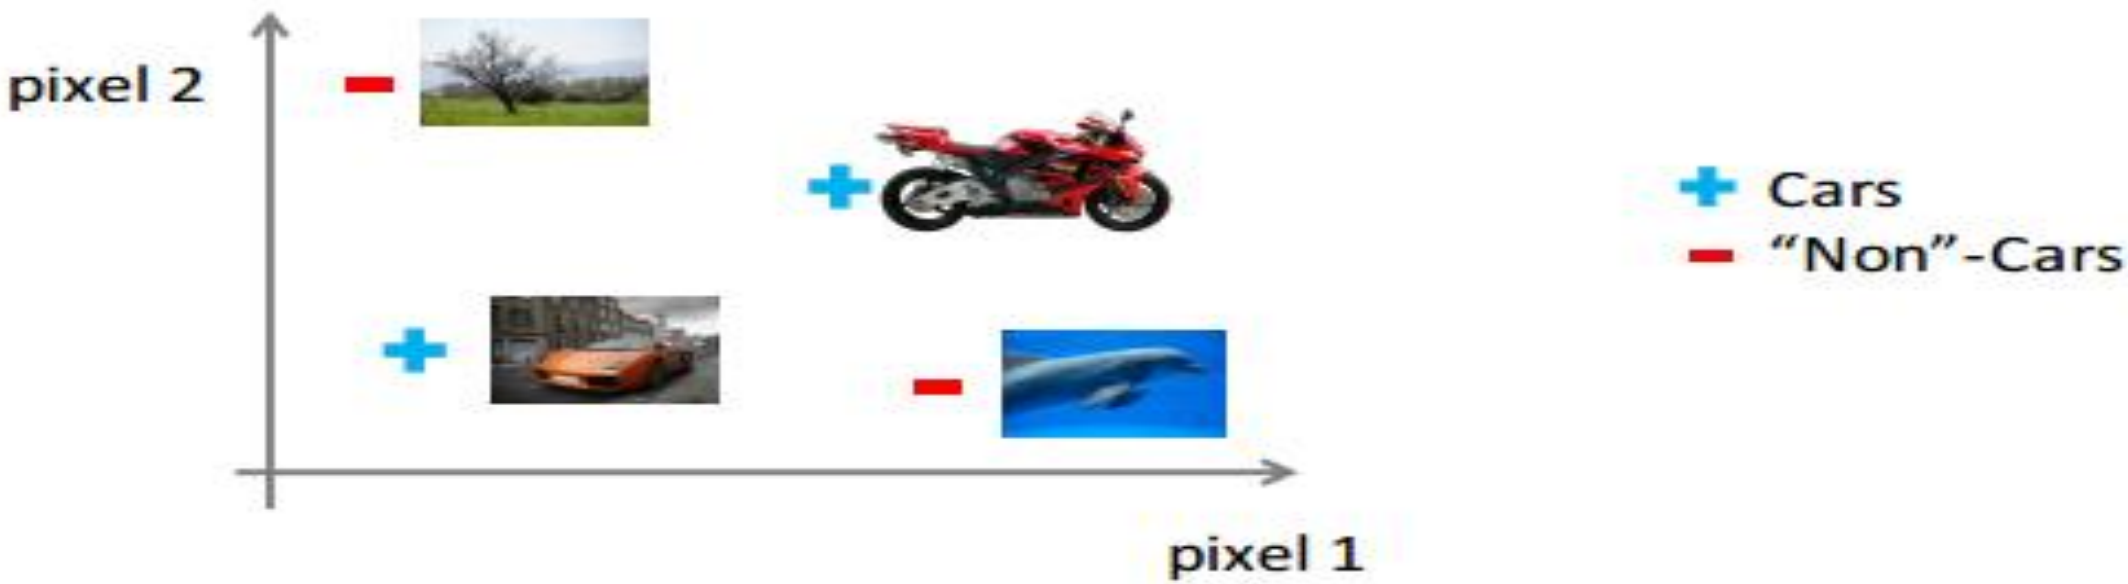

Things we want to do with data

Images

→ Label image

Audio

→ Speech recognition

Text

→ Web search

Features for machine learning

Images

Image

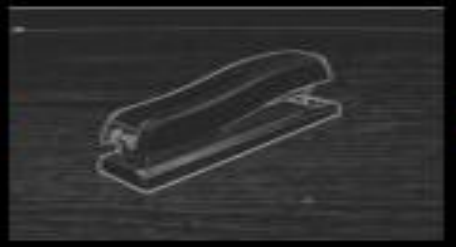

Vision features

Detection

Audio

Audio

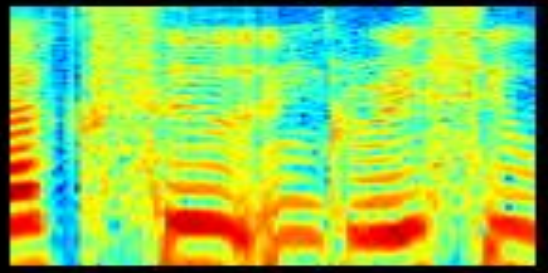

Audio features

Speaker ID

Text

Text

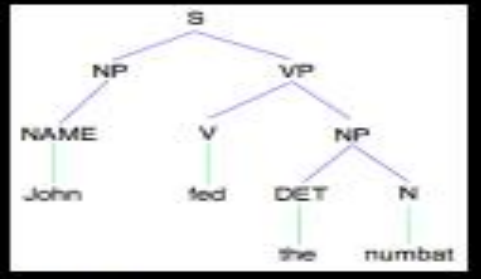

Text features

Web search

...

The short answers

1. 'Deep Learning' means using a neural network with several layers of nodes between input and output

2. the series of layers between input & output do feature identification and processing in a series of stages, just as our brains seem to.

How do We Train Deep Architectures?

- Inspiration from mammal brain
- Multiple Layers of "neurons" (Rumelhart et al 1986)
- Train each layer to compose the representations of the previous layer to learn a higher level abstraction
  - Ex: Pixels → Edges → Contours → Object parts → Object categories
  - Local Features → Global Features
- Train the layers one-by-one (Hinton et al 2006)
  - Greedy strategy

Deep Learning: learn representations!

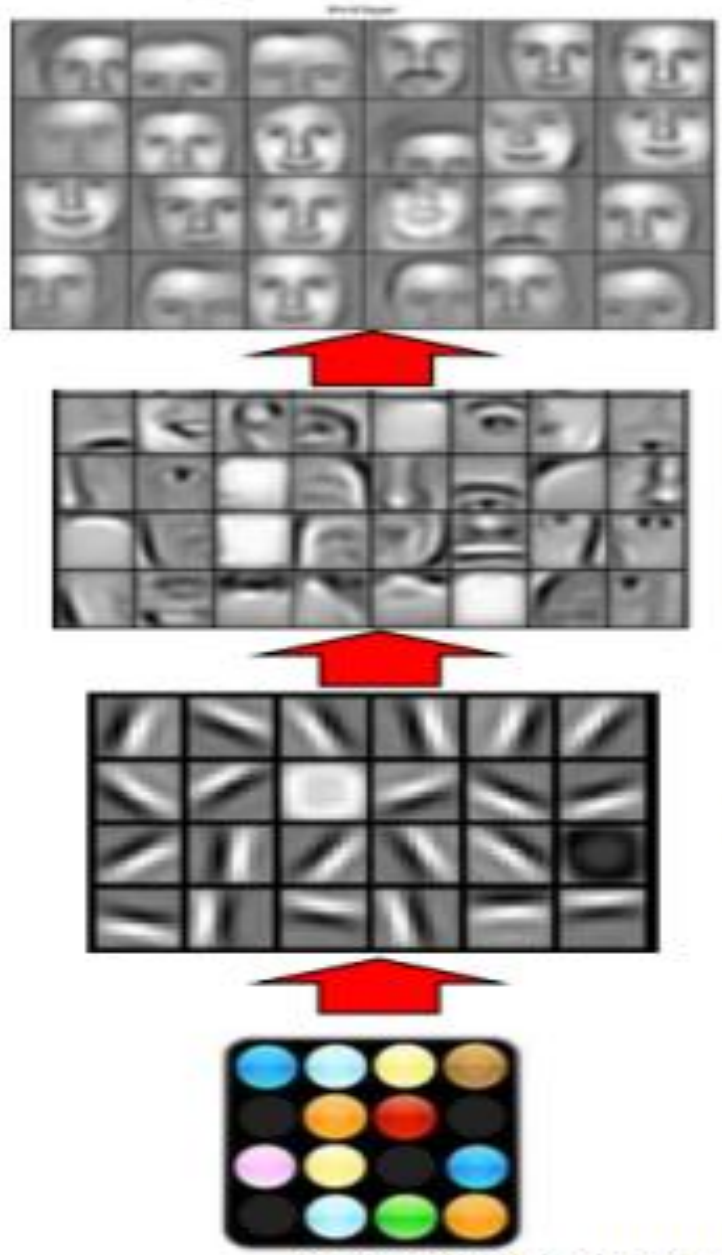

Source:Lee et.al. ICML2009

Machine learning workflow

Training set

Unsupervised

Supervised

Feature extraction

Machine learning algorithm

Grouping of objects

Predictive model

New Data

Annotated data

Deep Learning trends

Now

0-2 years
Tagged data

3-5 years
Tagged & untagged data

Learning from tagged data (supervised)

Coffee mug

Coffee mug

Coffee mug

Coffee mug

Coffee mug

Coffee mug

Testing: What is this?

Learning from tagged data

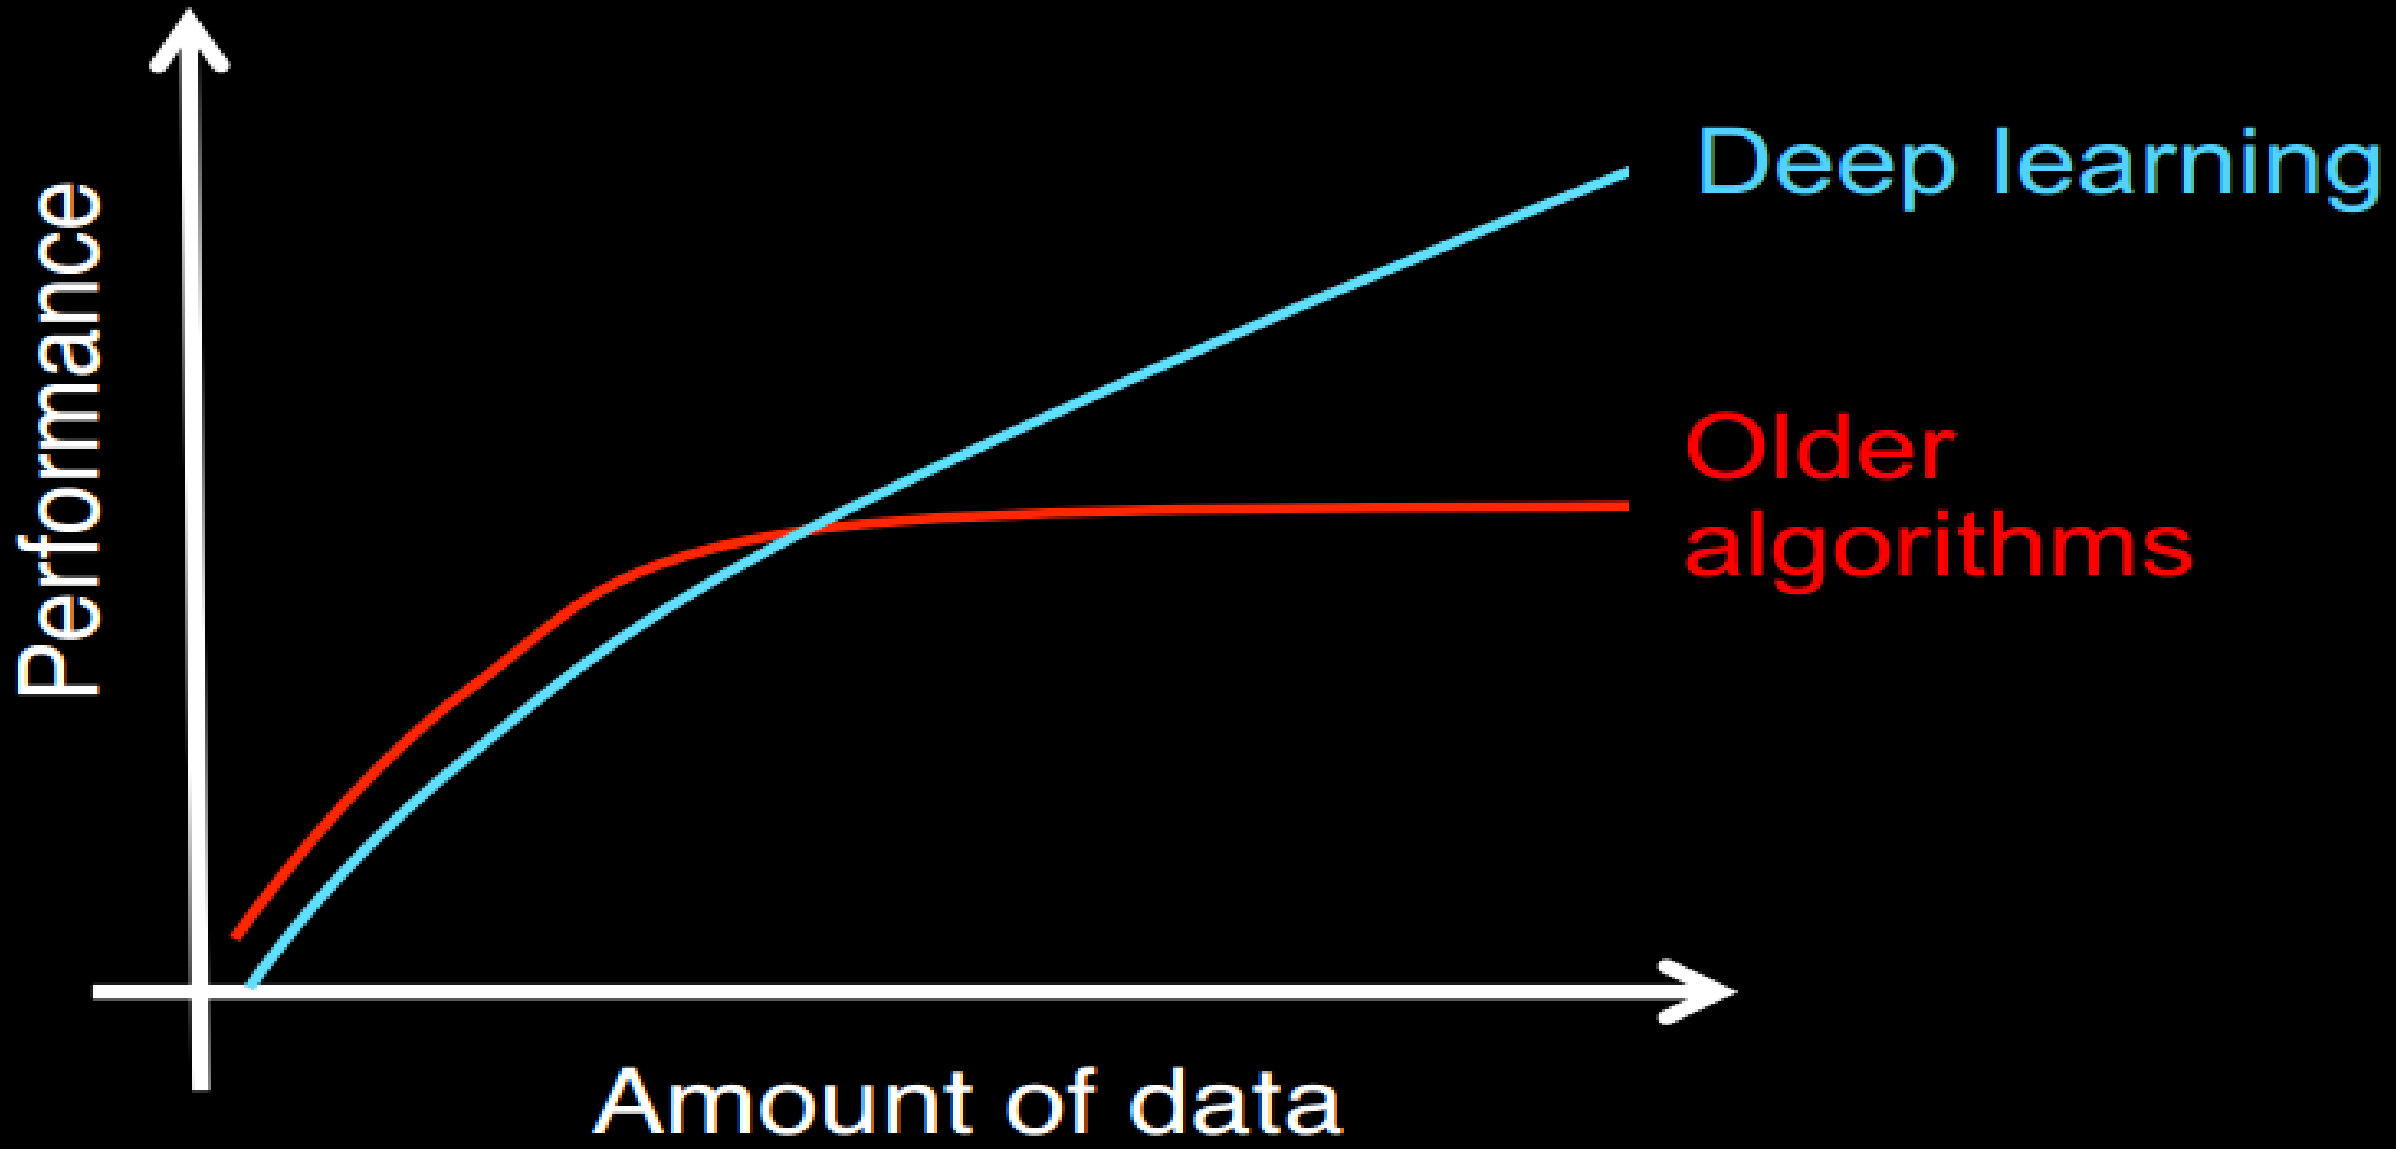

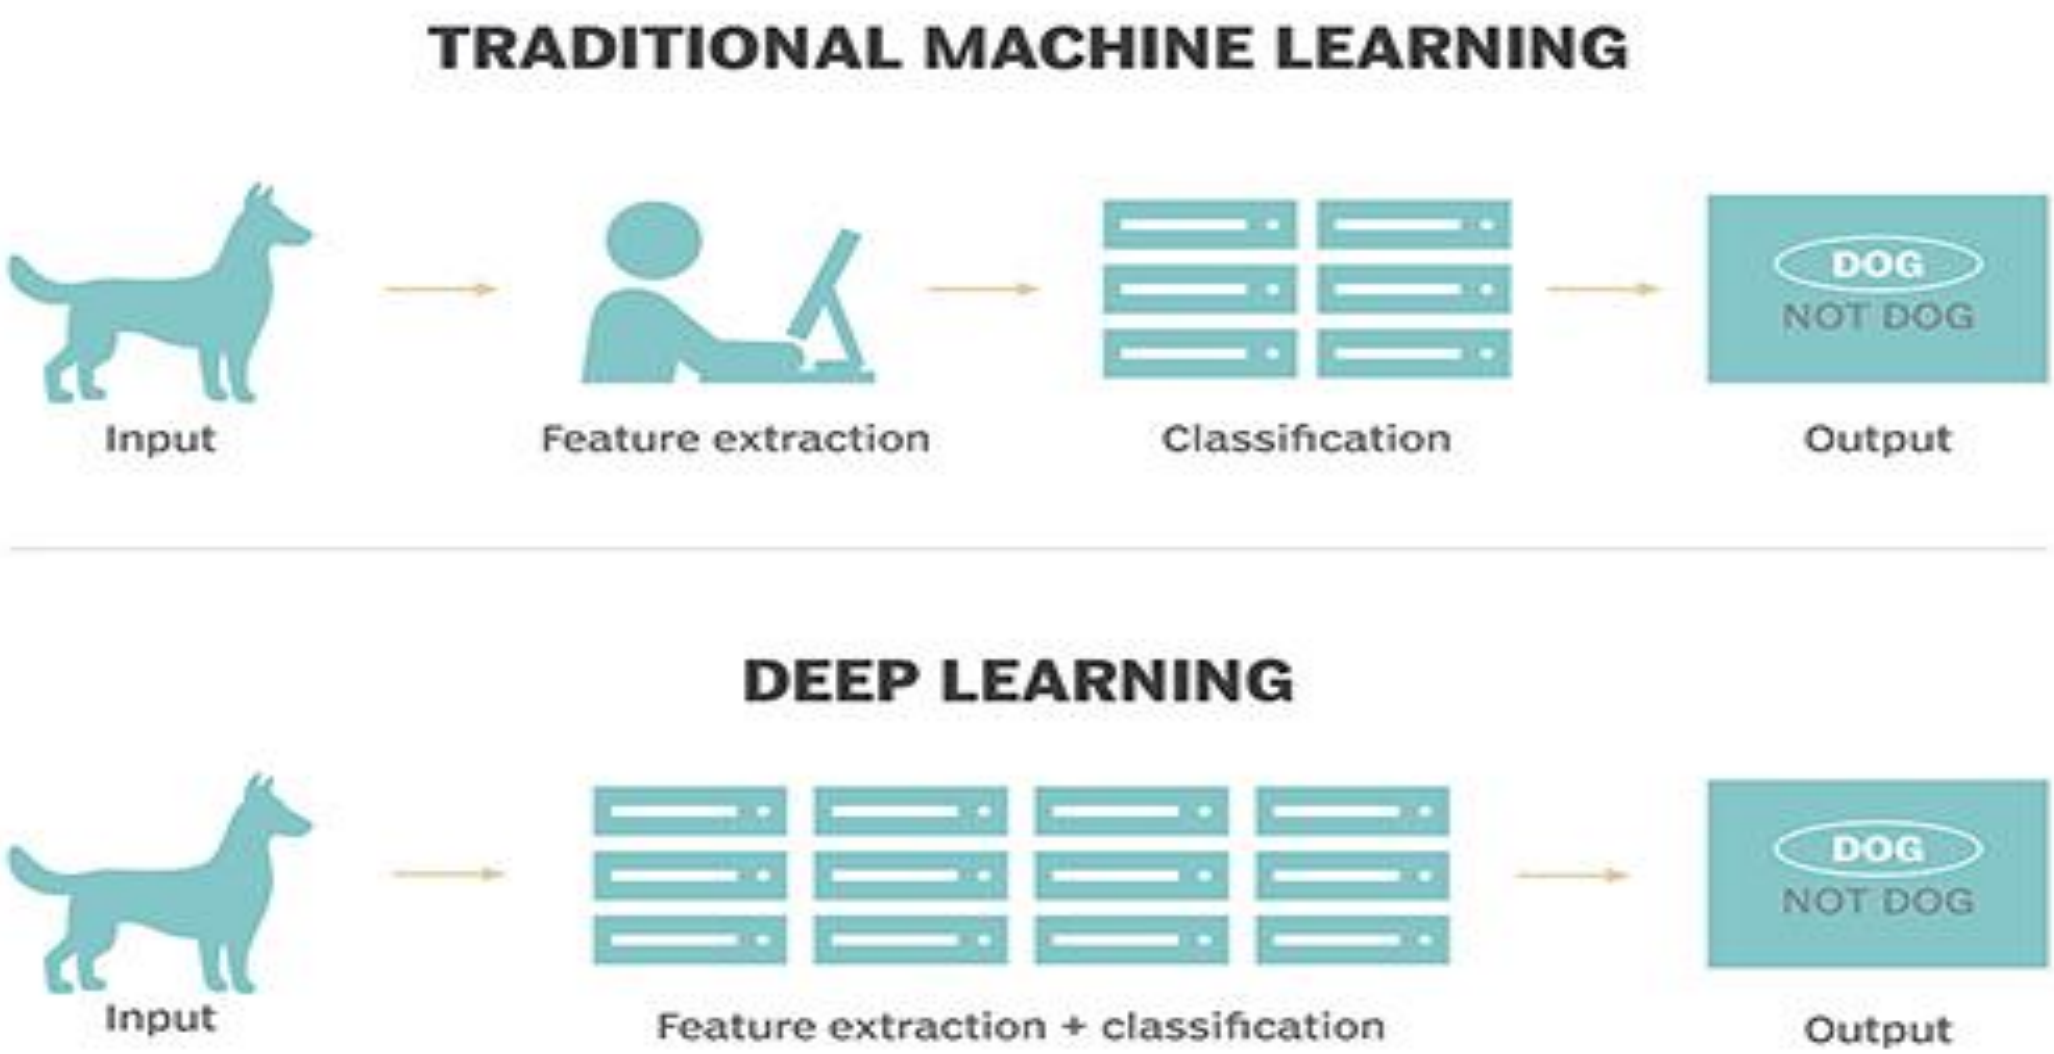
TRADITIONAL MACHINE LEARNING
Input
Feature extraction
Classification
DOG
NOT DOG
Output
DEEP LEARNING
Input
Feature extraction + classification
DOG
NOT DOG
Output

Deep learning applications

More traditional Machine Learning Applications to Deep Learning Application

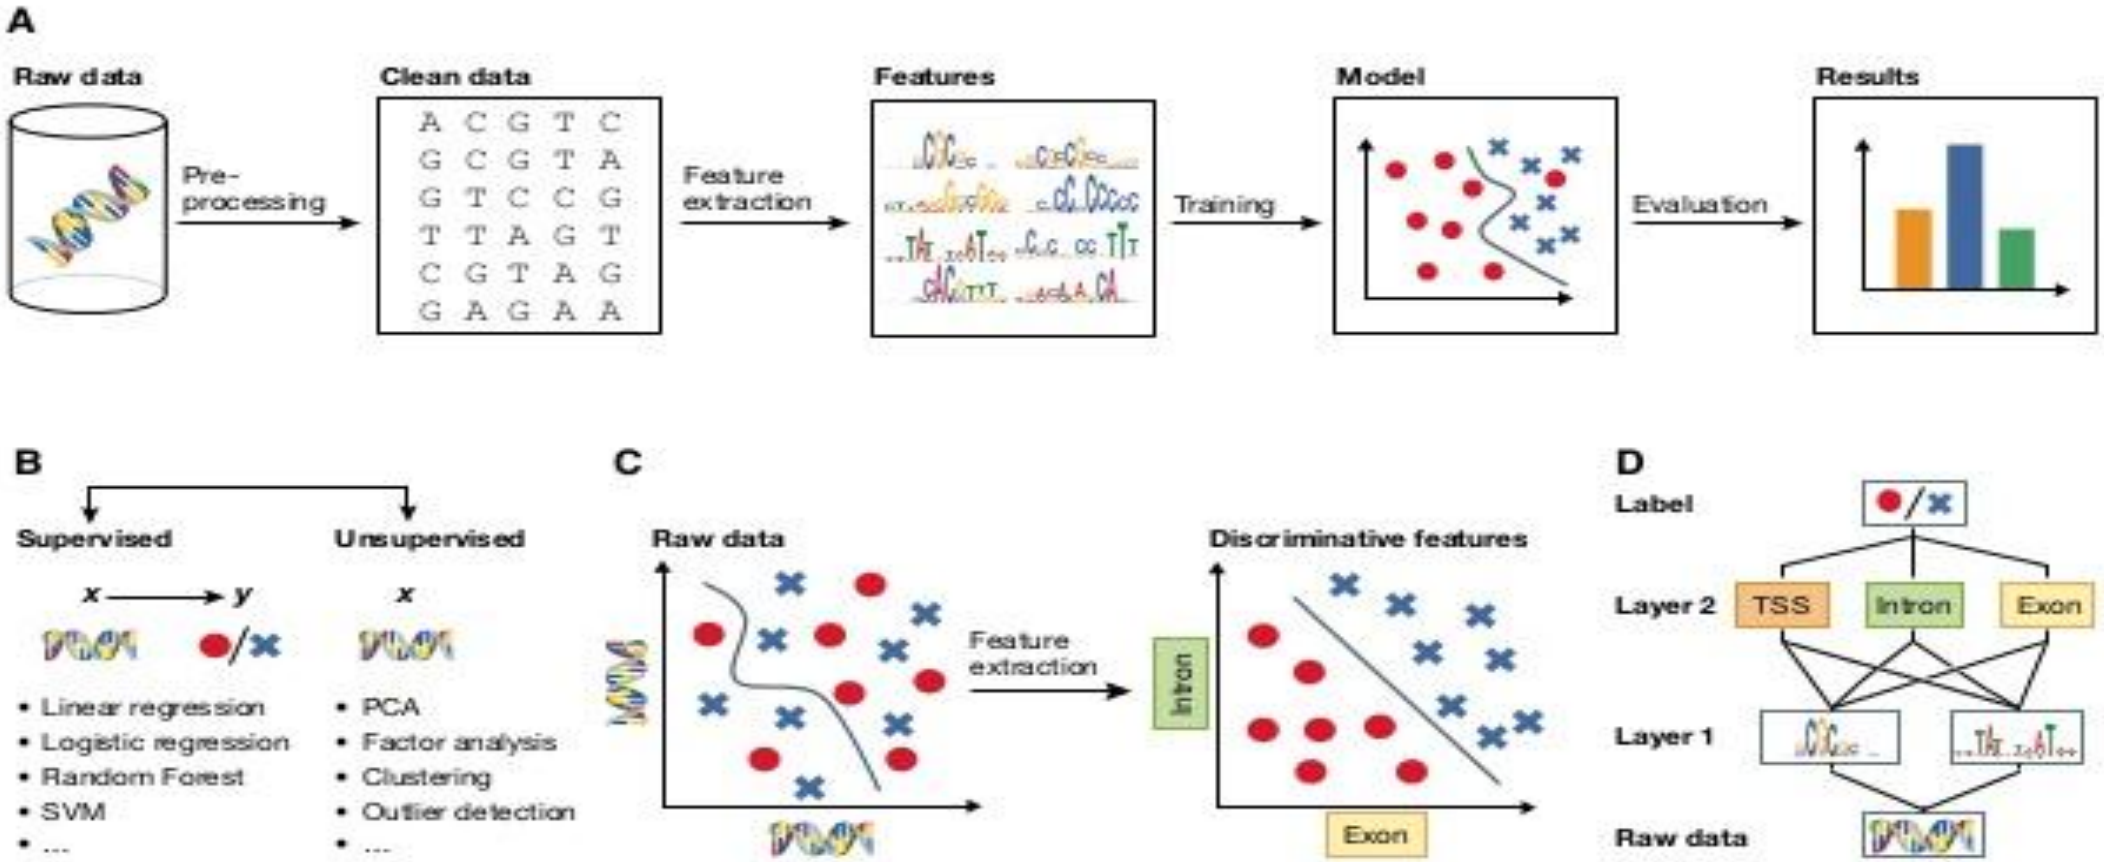

Angermueller et al., *Mol Sys Bio* 2016

A Motivational Task: Percepts → Concepts

- Create algorithms
  - that can understand scenes and describe them in natural language
  - that can infer semantic concepts to allow machines to interact with humans using these concepts
- Requires creating a series of abstractions
  - Image (Pixel Intensities) → Objects in Image → Object Interactions → Scene Description
- Deep learning aims to automatically learn these abstractions with little supervision

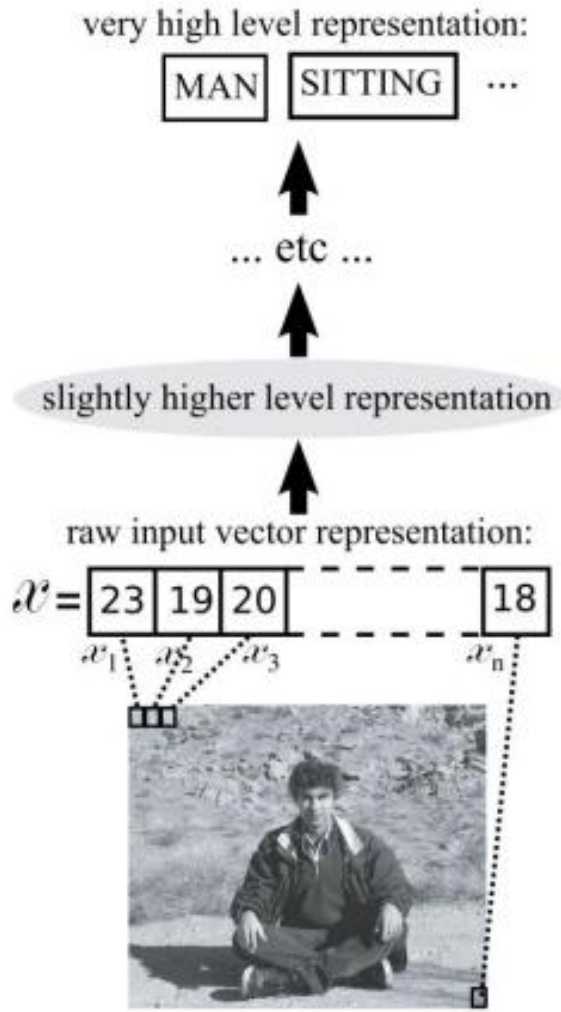

Courtesy: Yoshua Bengio, Learning Deep Architectures for AI

Neurons in the brain

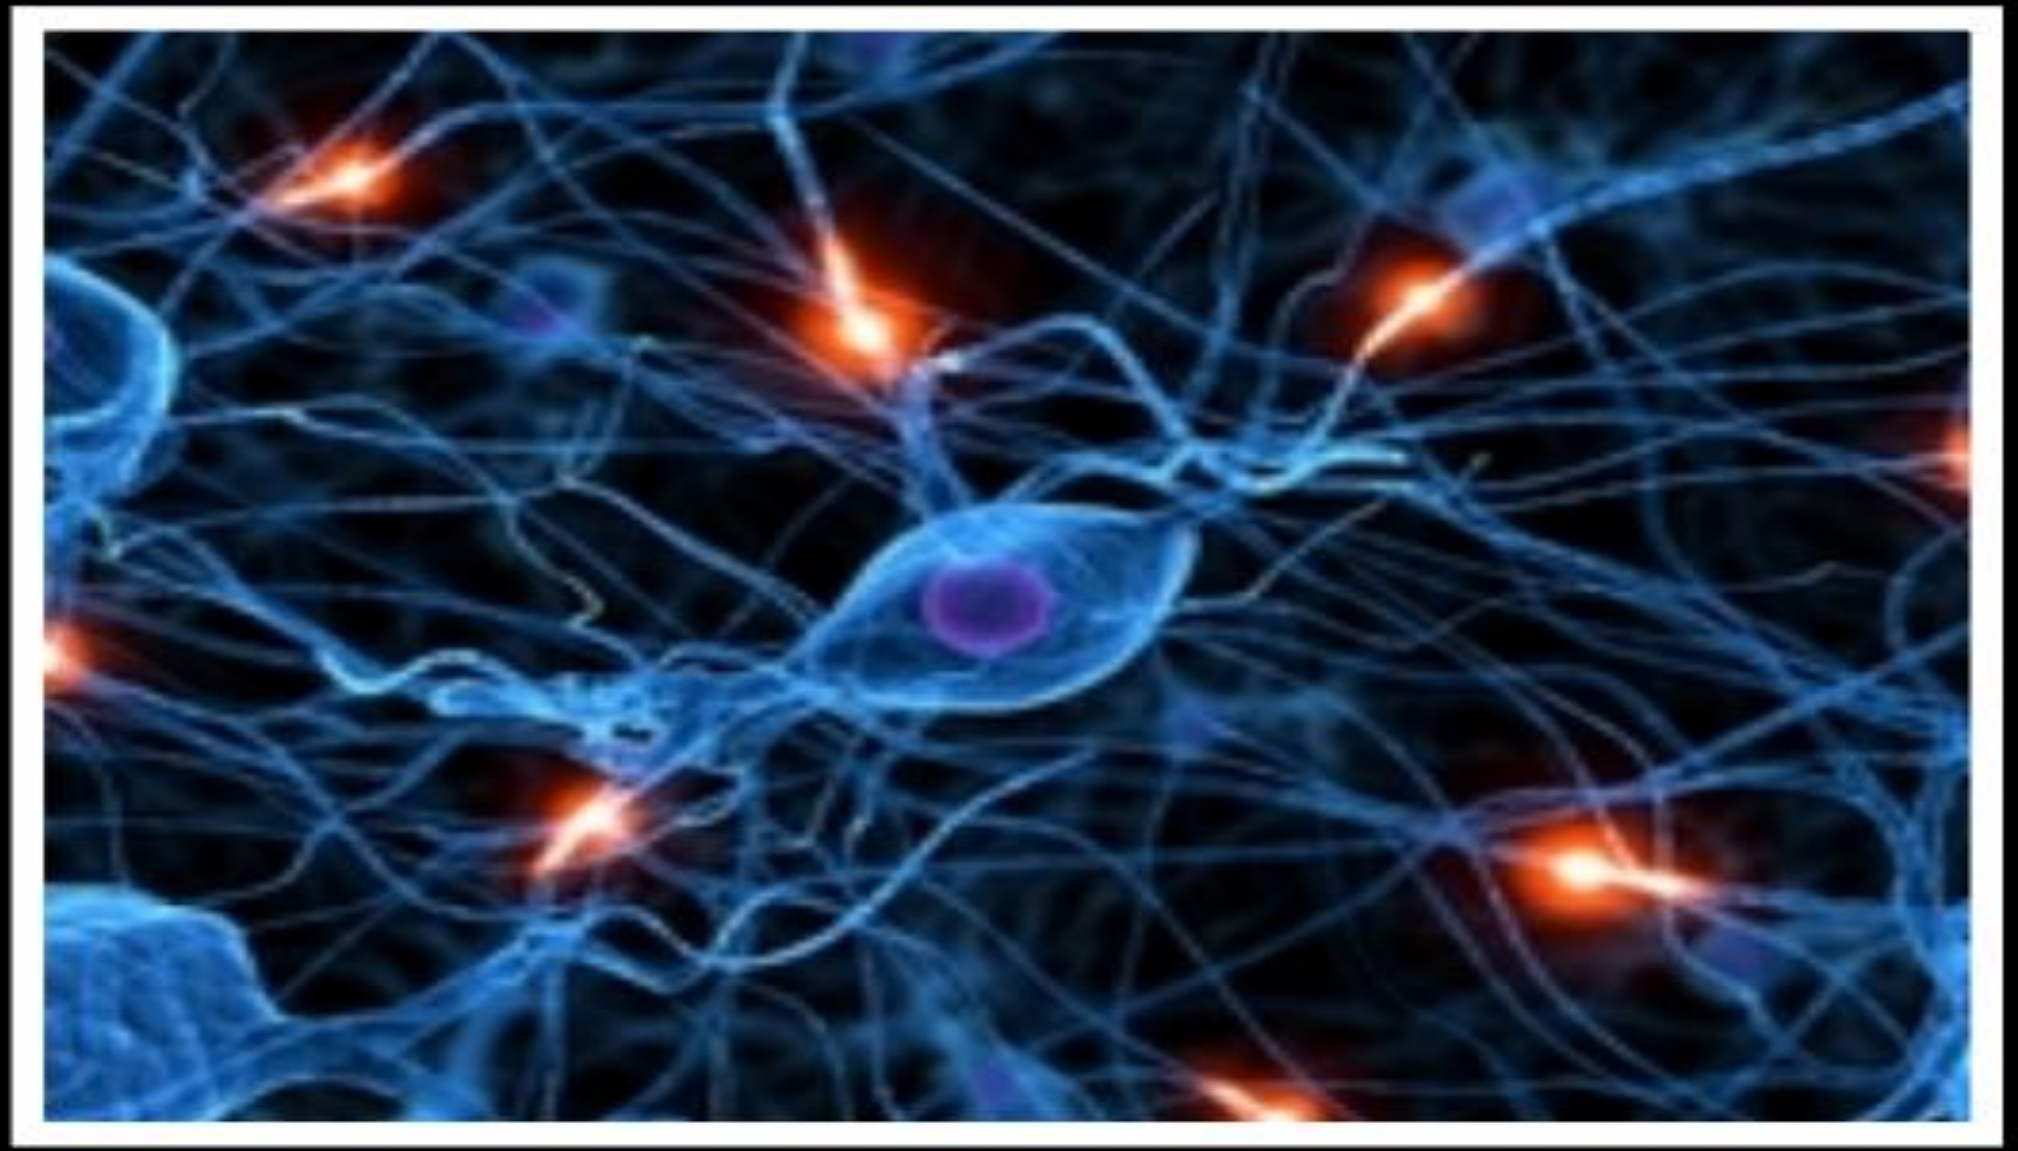

Neural Network (Deep Learning)

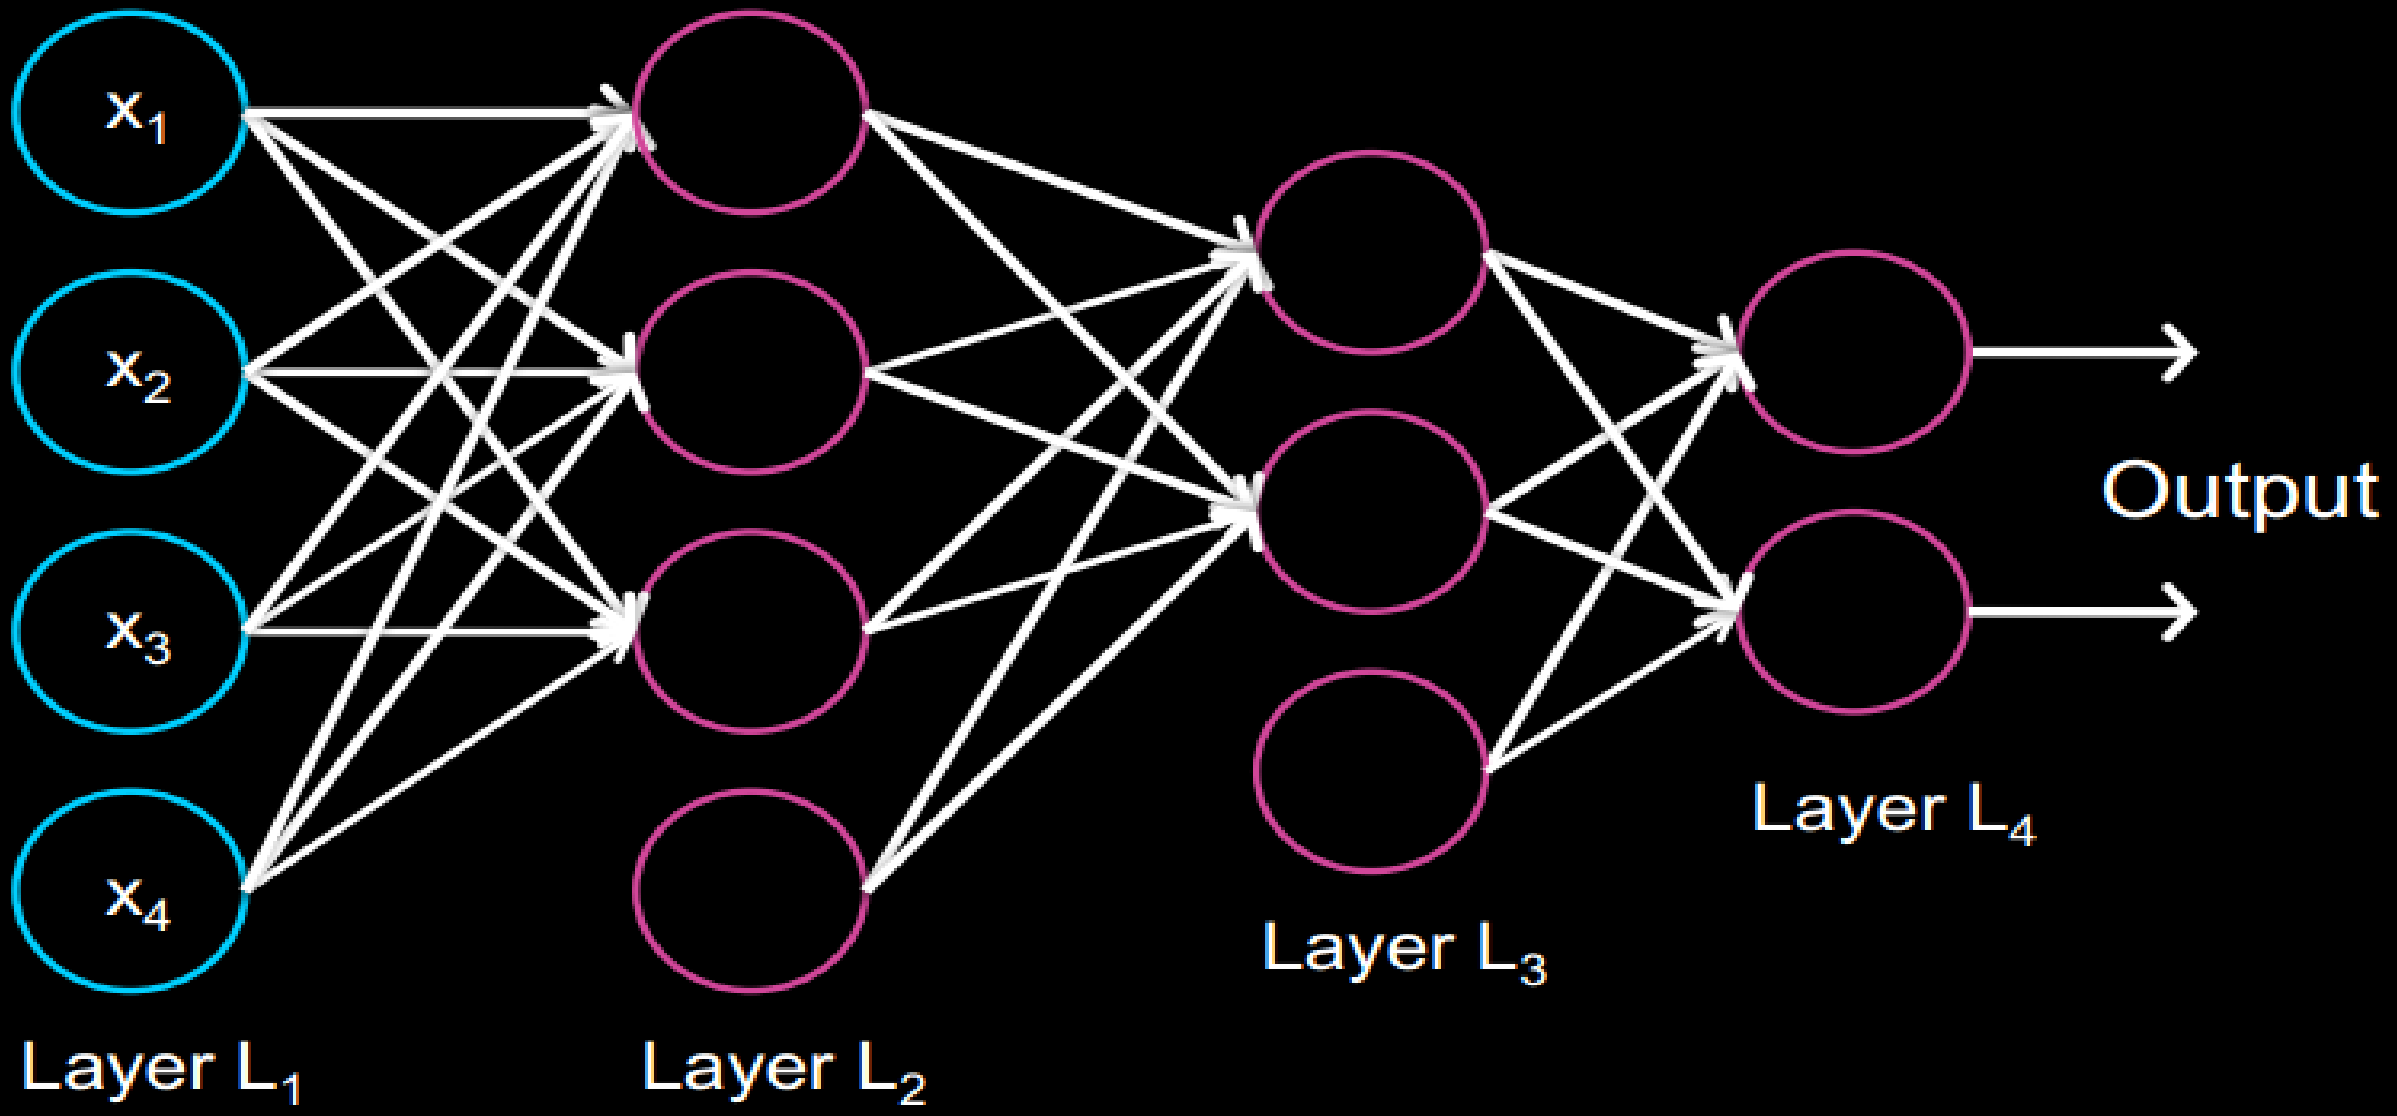

Single Unit, Input, weights, activation function, output

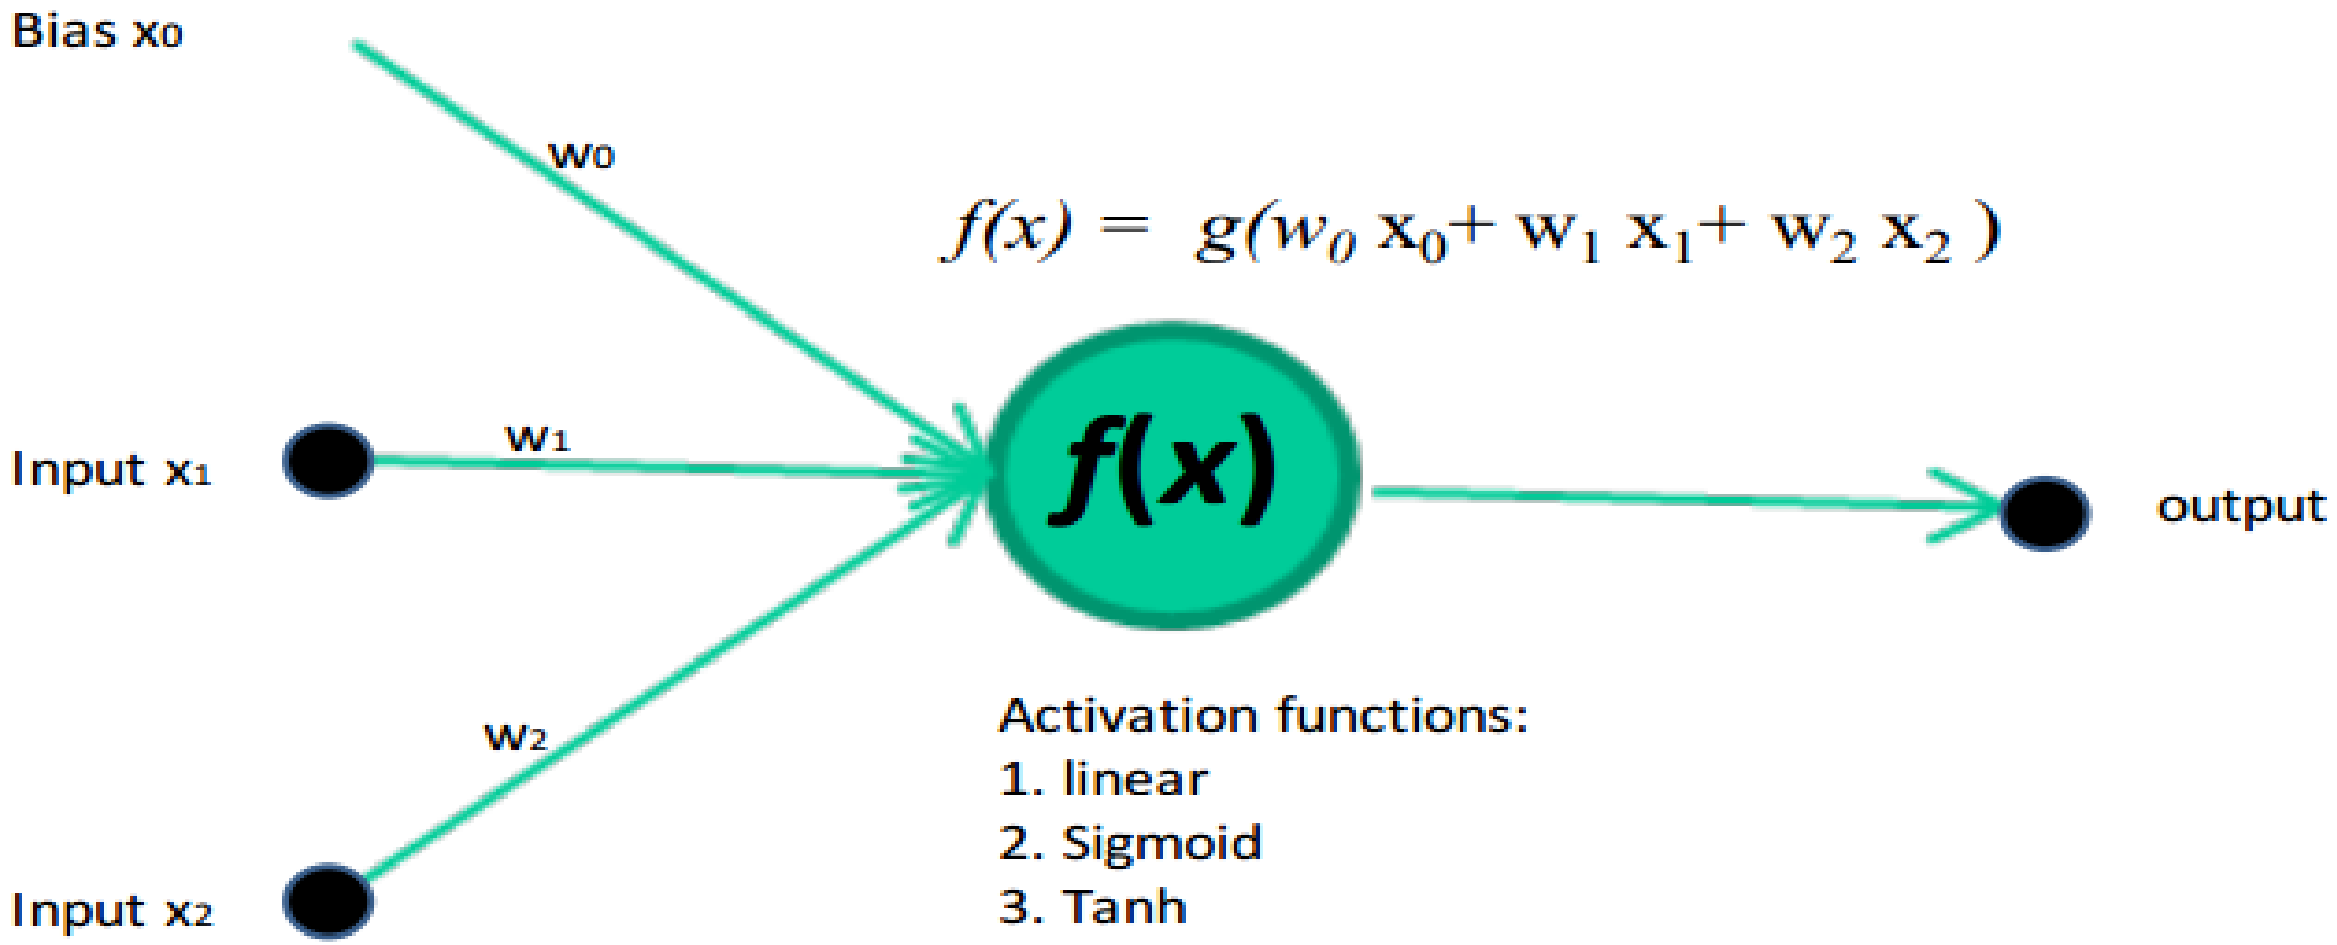

Activation functions:

1. linear
2. Sigmoid
3. Tanh
4. Relu
5. Softmax

etc.

A dataset

Fields			***class***
1.4	2.7	1.9	0
3.8	3.4	3.2	0
6.4	2.8	1.7	1
4.1	0.1	0.2	0

etc ...

Train the deep neural network

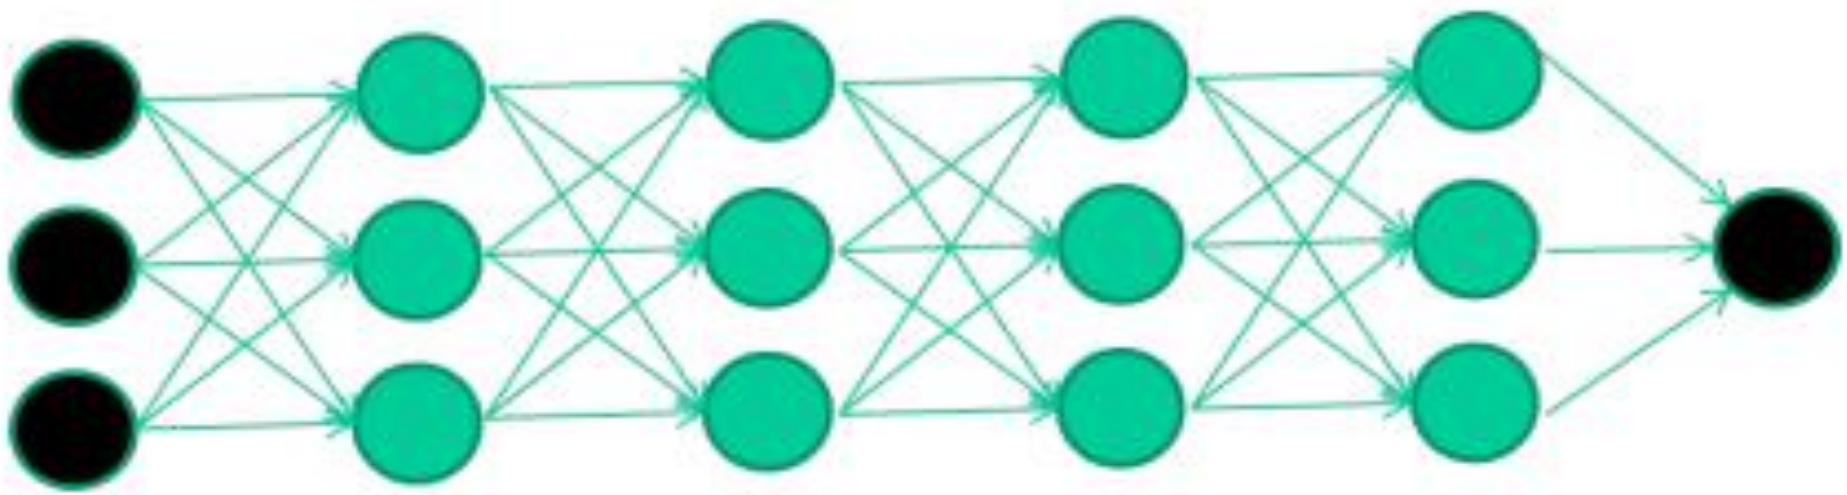

Multilayer Perceptron with Back-propagation
First deep learning model (Rumelhart, Hinton, Williams 1986)

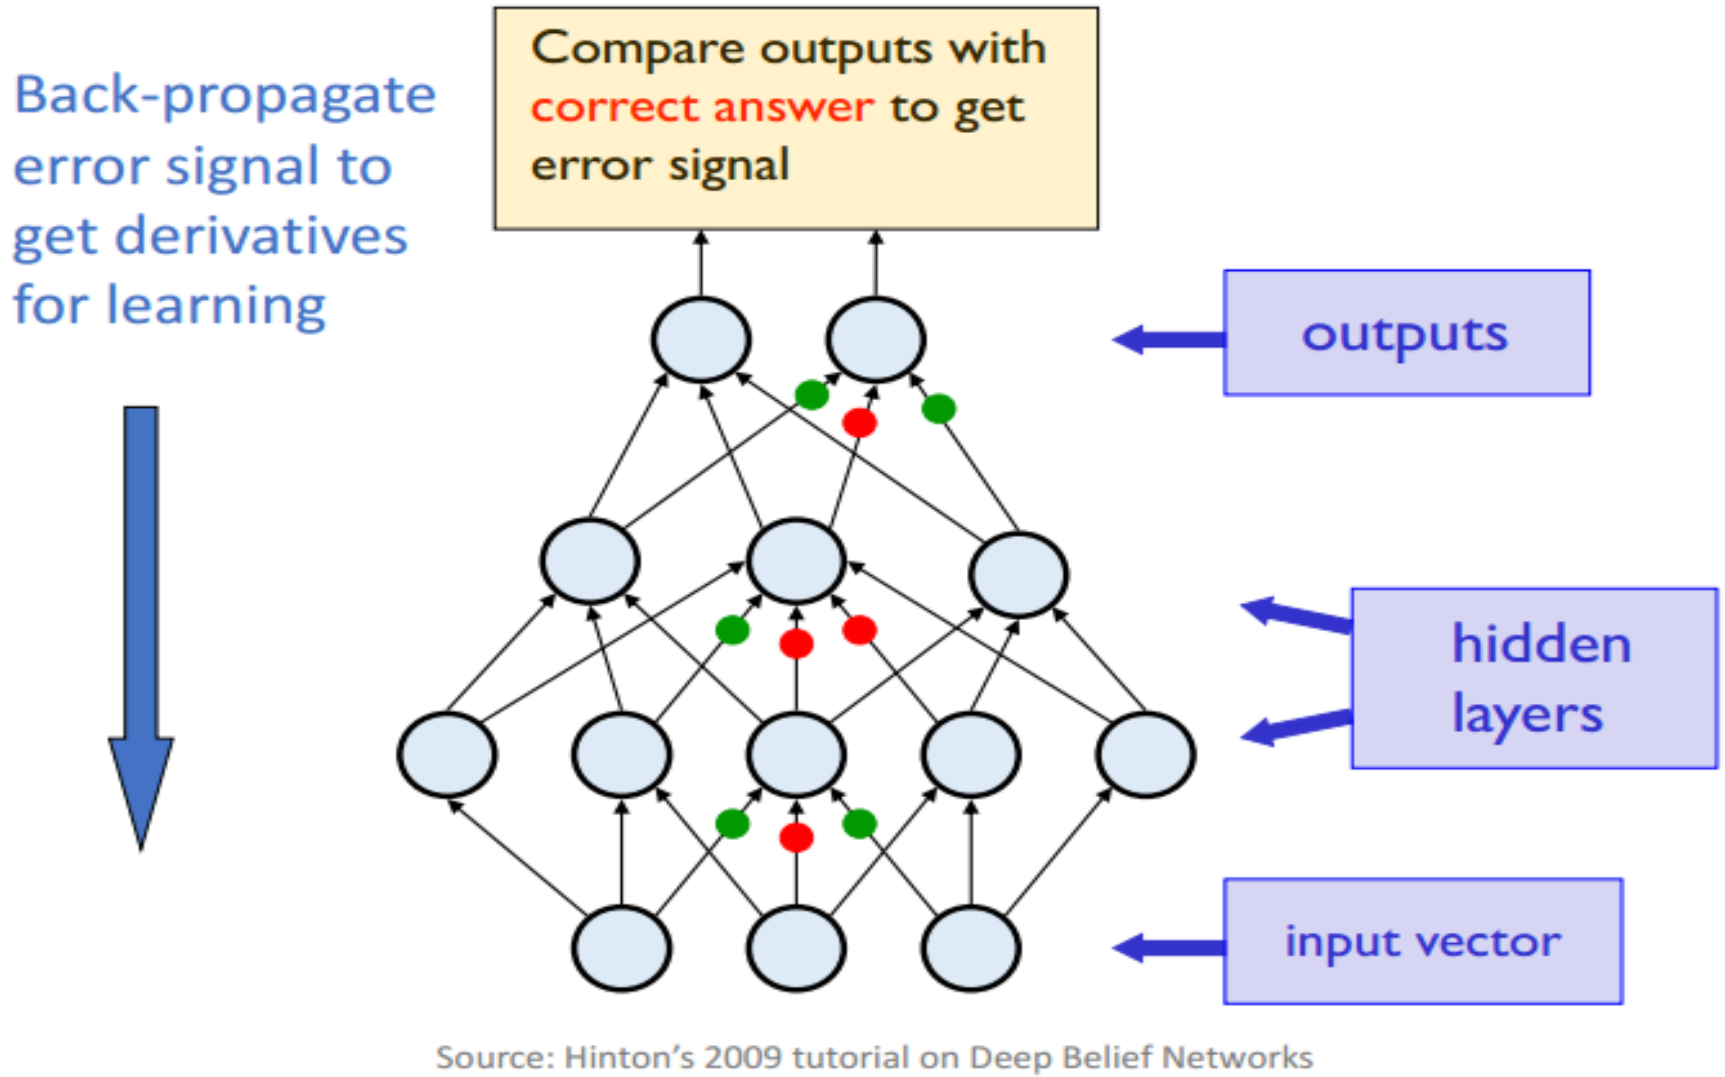

Source: Hinton's 2009 tutorial on Deep Belief Networks

DEEP LEARNING EVERYWHERE


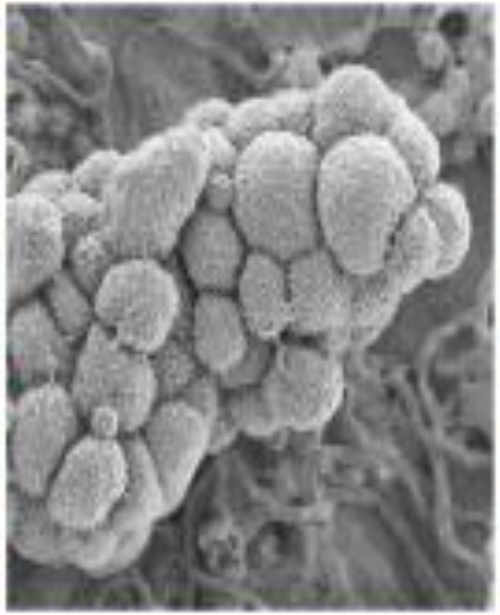


INTERNET & CLOUD	MEDICINE & BIOLOGY	MEDIA & ENTERTAINMENT	SECURITY & DEFENSE	AUTONOMOUS MACHINES
Image Classification	Cancer Cell Detection	Video Captioning	Face Detection	Pedestrian Detection
Speech Recognition	Diabetic Grading	Video Search	Video Surveillance	Lane Tracking
Language Translation	Drug Discovery	Real Time Translation	Satellite Imagery	Recognize Traffic Sign
Language Processing				
Sentiment Analysis				
Recommendation				

Deep Visual-Semantic Alignments for Generating Image Descriptions (Karpathy, Fei-Fei; CVPR 2015)

"two young girls are playing with lego toy."

"boy is doing backflip on wakeboard."

"construction worker in orange safety vest is working on road."

"man in black shirt is playing guitar."

Example: Learning the Configuration Space of a Robotic Arm

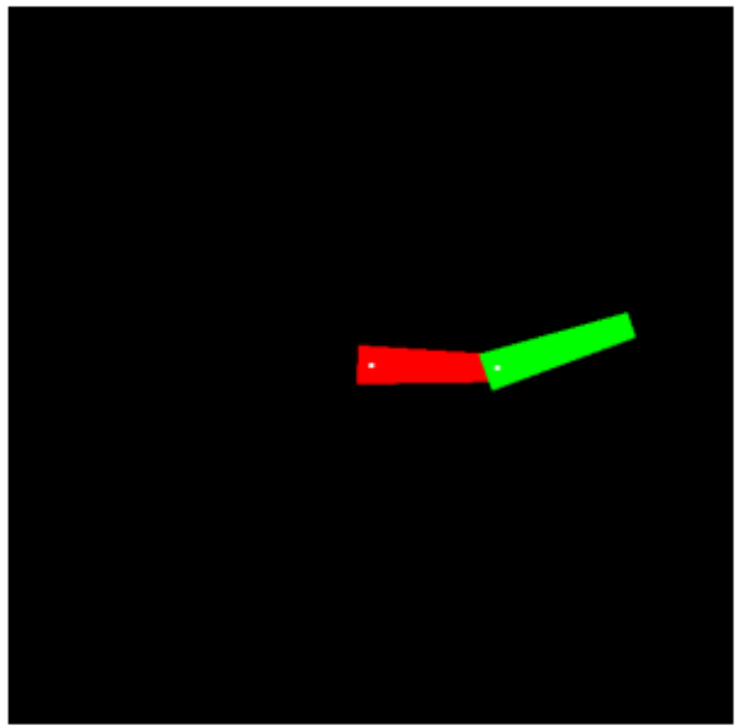

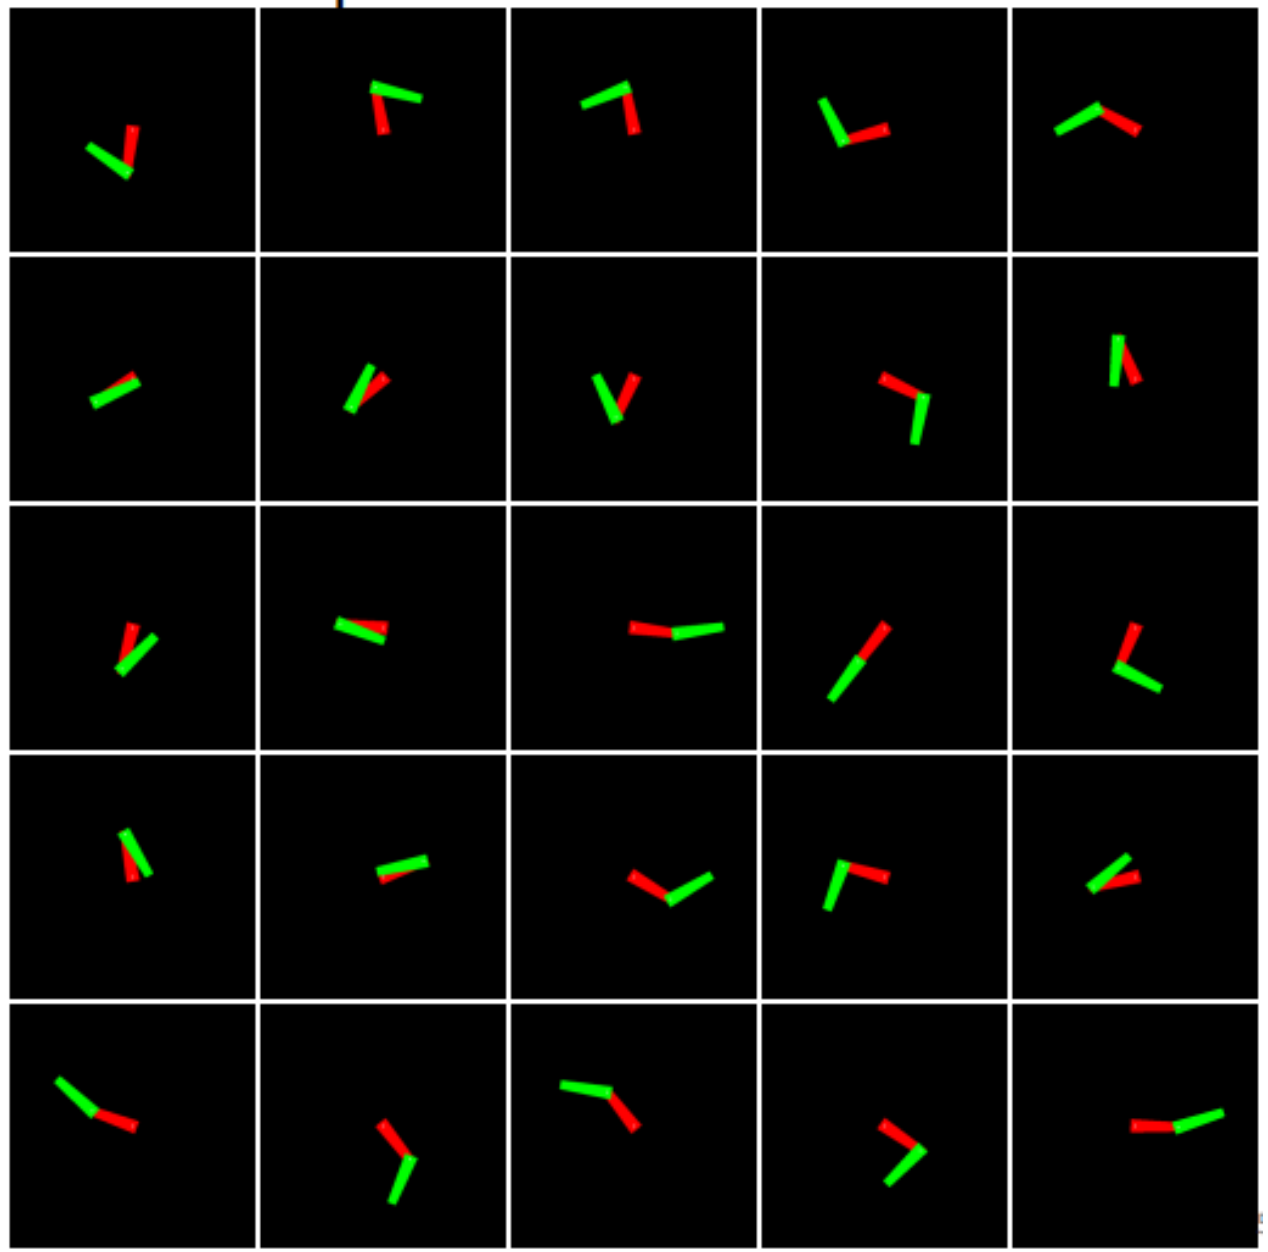

Deep Learning in Computer Vision

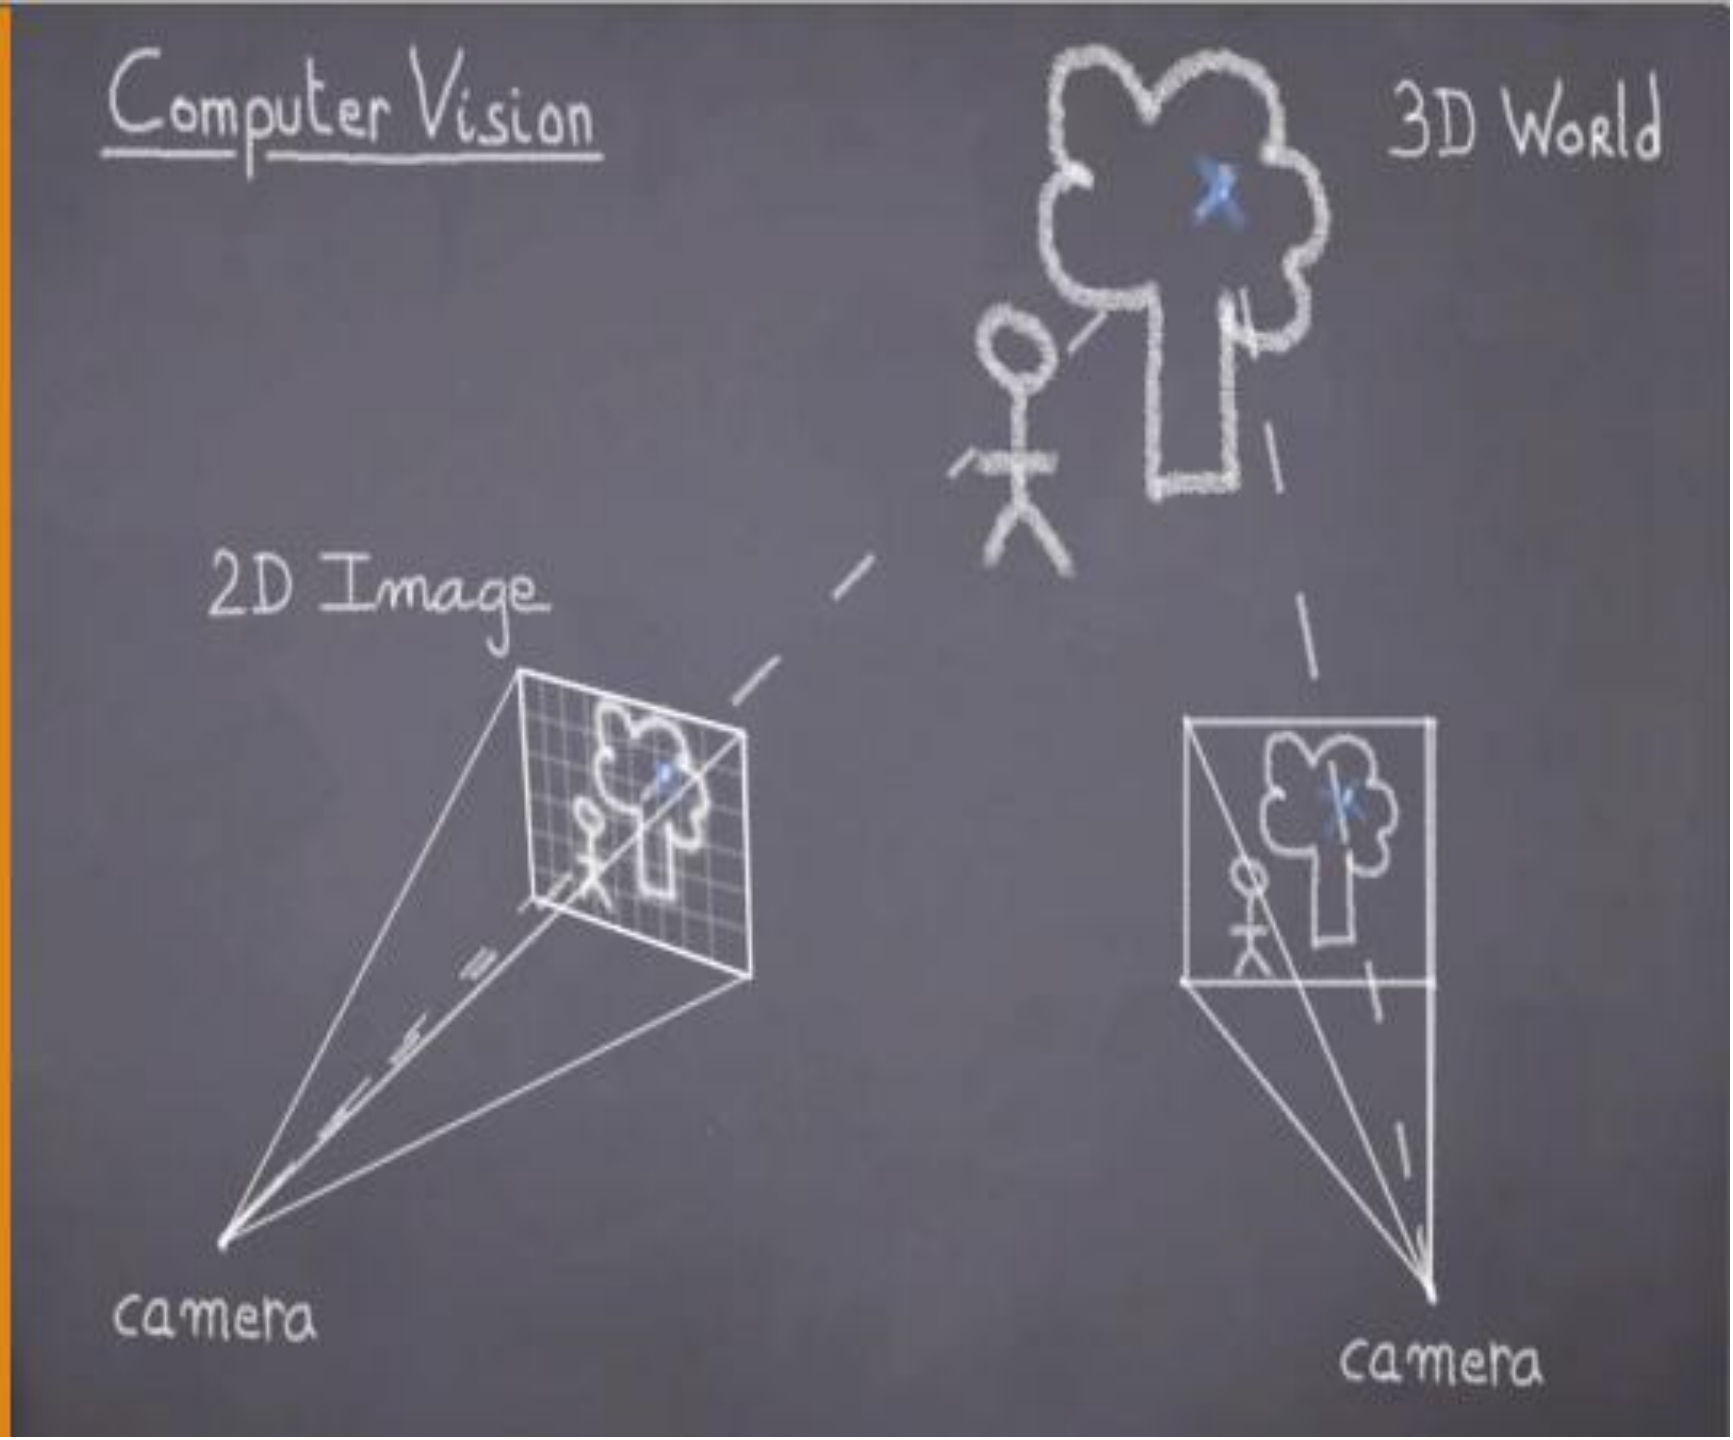

Deep Learning in Robotics

Deep Learning in NLP and Speech

Imitating famous painters

Handwriting

Hi Motherboard readers!

-This entire post was hand written by a neural network.

(It probably writes better than you.)

Of course, a neural network doesn't actually have hands.

And the original text was typed by me, a human.

So what's going on here?

neural network is a program that can learn to follow a set of rules

But it can't do it alone. It needs to be trained.

This neural network was trained on a corpus of writing samples.

of actual hand-writing,

ut of the locations of a pen-tip as people write.

is how the network learns and creates different styles

from prior examples.

And it can use the knowledge

to generate handwritten notes from inputted text.

can create its own style, or mimic another's.

No two notes are the same.

It's the work of Alex Graves at the University of Toronto

And you can try it too!

CIFAR 10 and Convolutional Neural Network

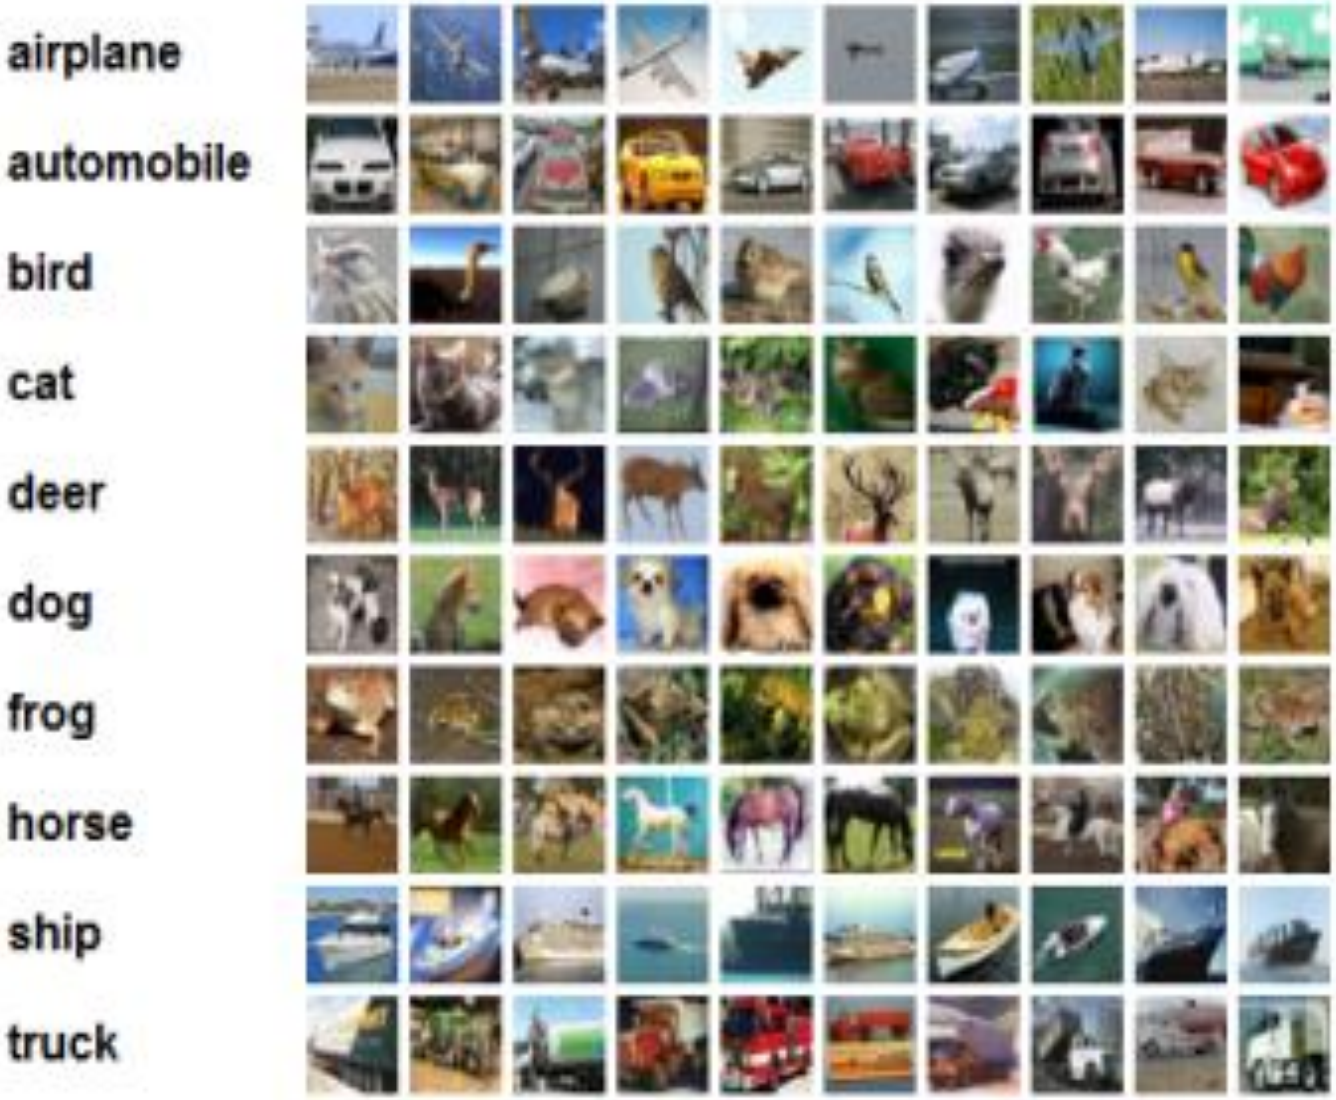

CIFAR 10 dataset:
50,000 training images
10,000 testing images
10 categories (classes)

Accuracies from different methods:

Human: ~94%

Whitening K-mean: 80%

.....

Deep CNN: 95.5%

Recommender System

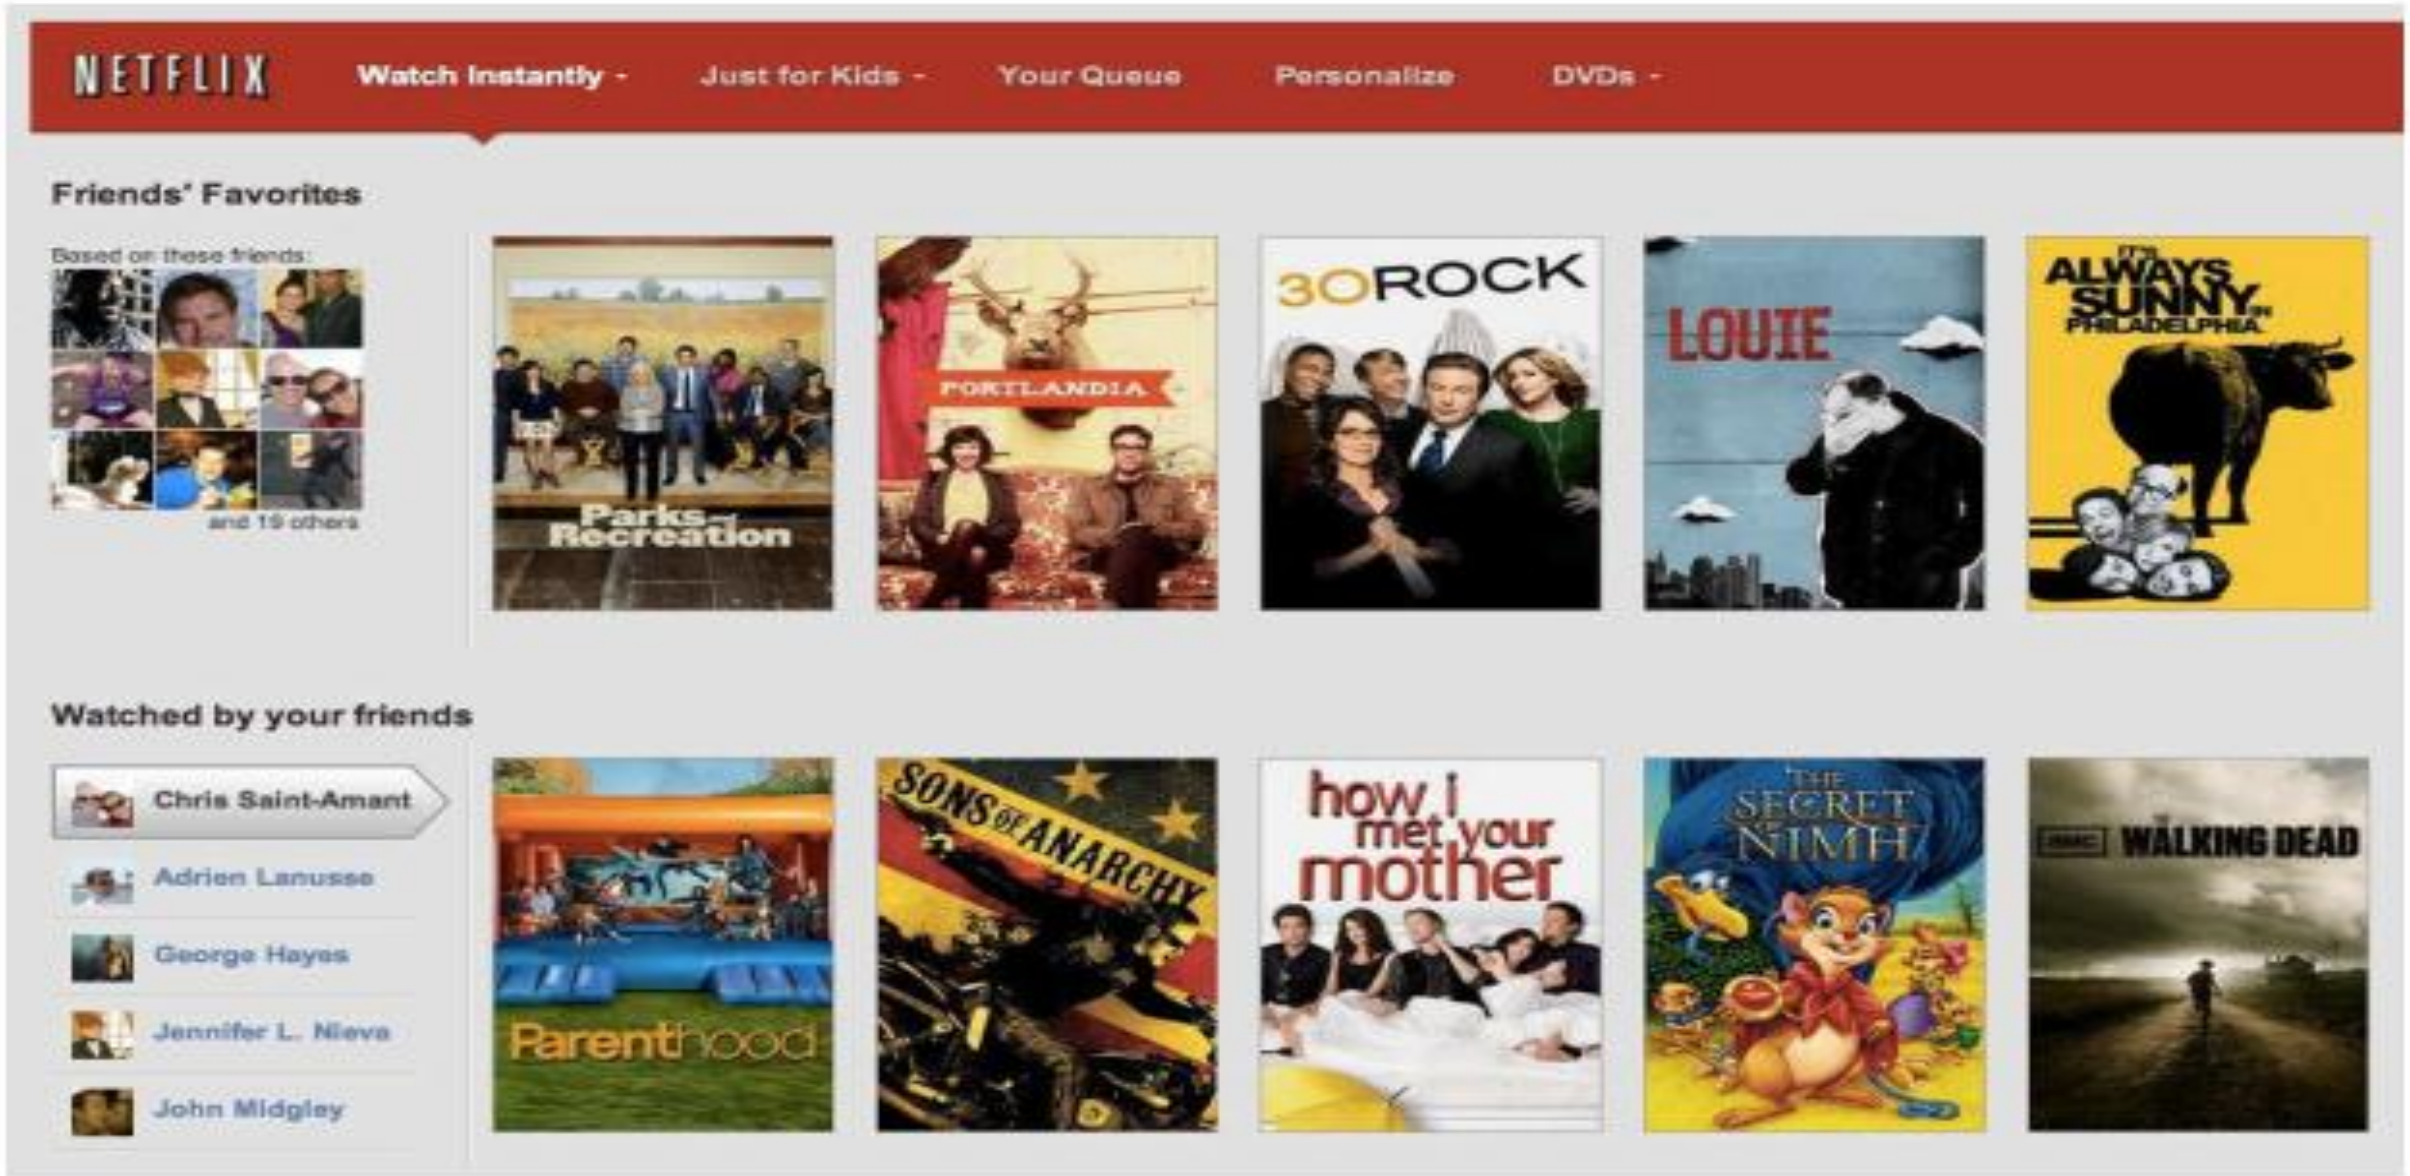

Recommender System

Total price: $40.05

Add all three to Cart

Add all three to List

- **This item:** Numerical Methods for Scientists and Engineers (Dover Books on Mathematics) by R. W. Hamming Paperback $14.31
- Partial Differential Equations for Scientists and Engineers (Dover Books on Mathematics) by Stanley J. Farlow Paperback $10.26
- Ordinary Differential Equations (Dover Books on Mathematics) by Morris Tenenbaum Paperback $15.48

Customers Who Bought This Item Also Bought

Partial Differential Equations for Scientists and Engineers (Dover...
› Stanley J. Farlow
132
Paperback
$10.26 Prime

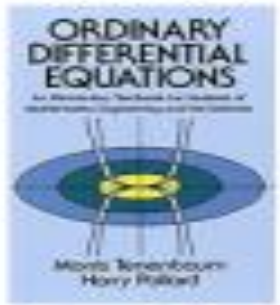

Ordinary Differential Equations (Dover Books on Mathematics)
› Morris Tenenbaum
115
Paperback
$15.48 Prime

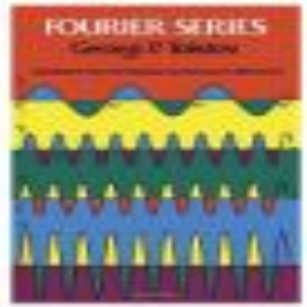

Fourier Series (Dover Books on Mathematics)
Georgi P. Tolstov
61
#1 Best Seller in Functional Analysis...
Paperback
$3.99 Prime

An Introduction to Information Theory: Symbols, Signals and...
› John R. Pierce
90
Paperback
$3.99 Prime

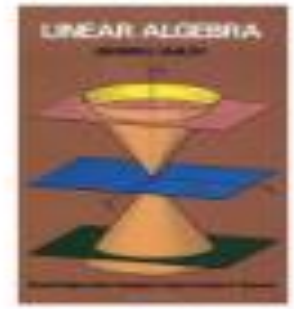

Linear Algebra (Dover Books on Mathematics)
Georgi E. Shilov
45
Paperback
$14.35 Prime

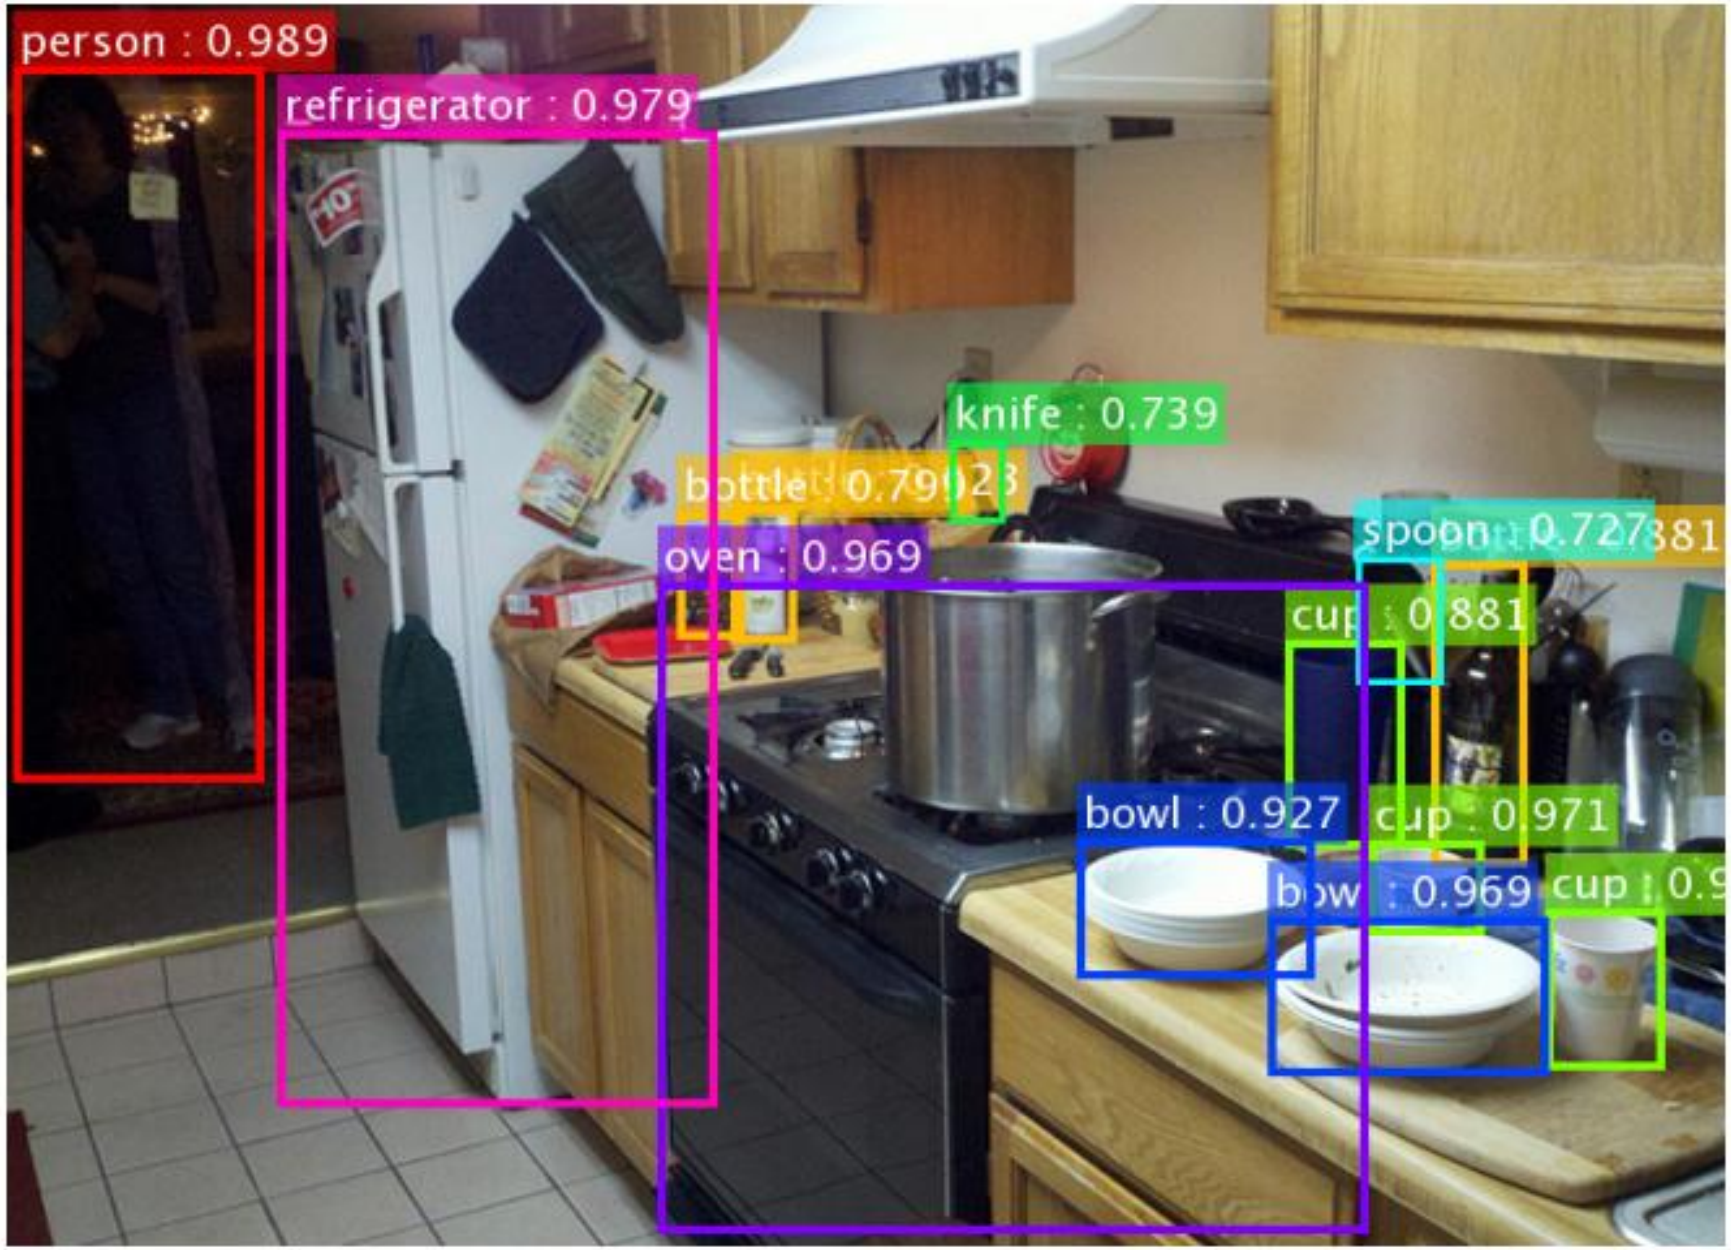

*the original image is from the COCO dataset

Kaiming He, Xiangyu Zhang, Shaoqing Ren, & Jian Sun. "Deep Residual Learning for Image Recognition". arXiv 2015.
Shaoqing Ren, Kaiming He, Ross Girshick, & Jian Sun. "Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks". NIPS 2015.

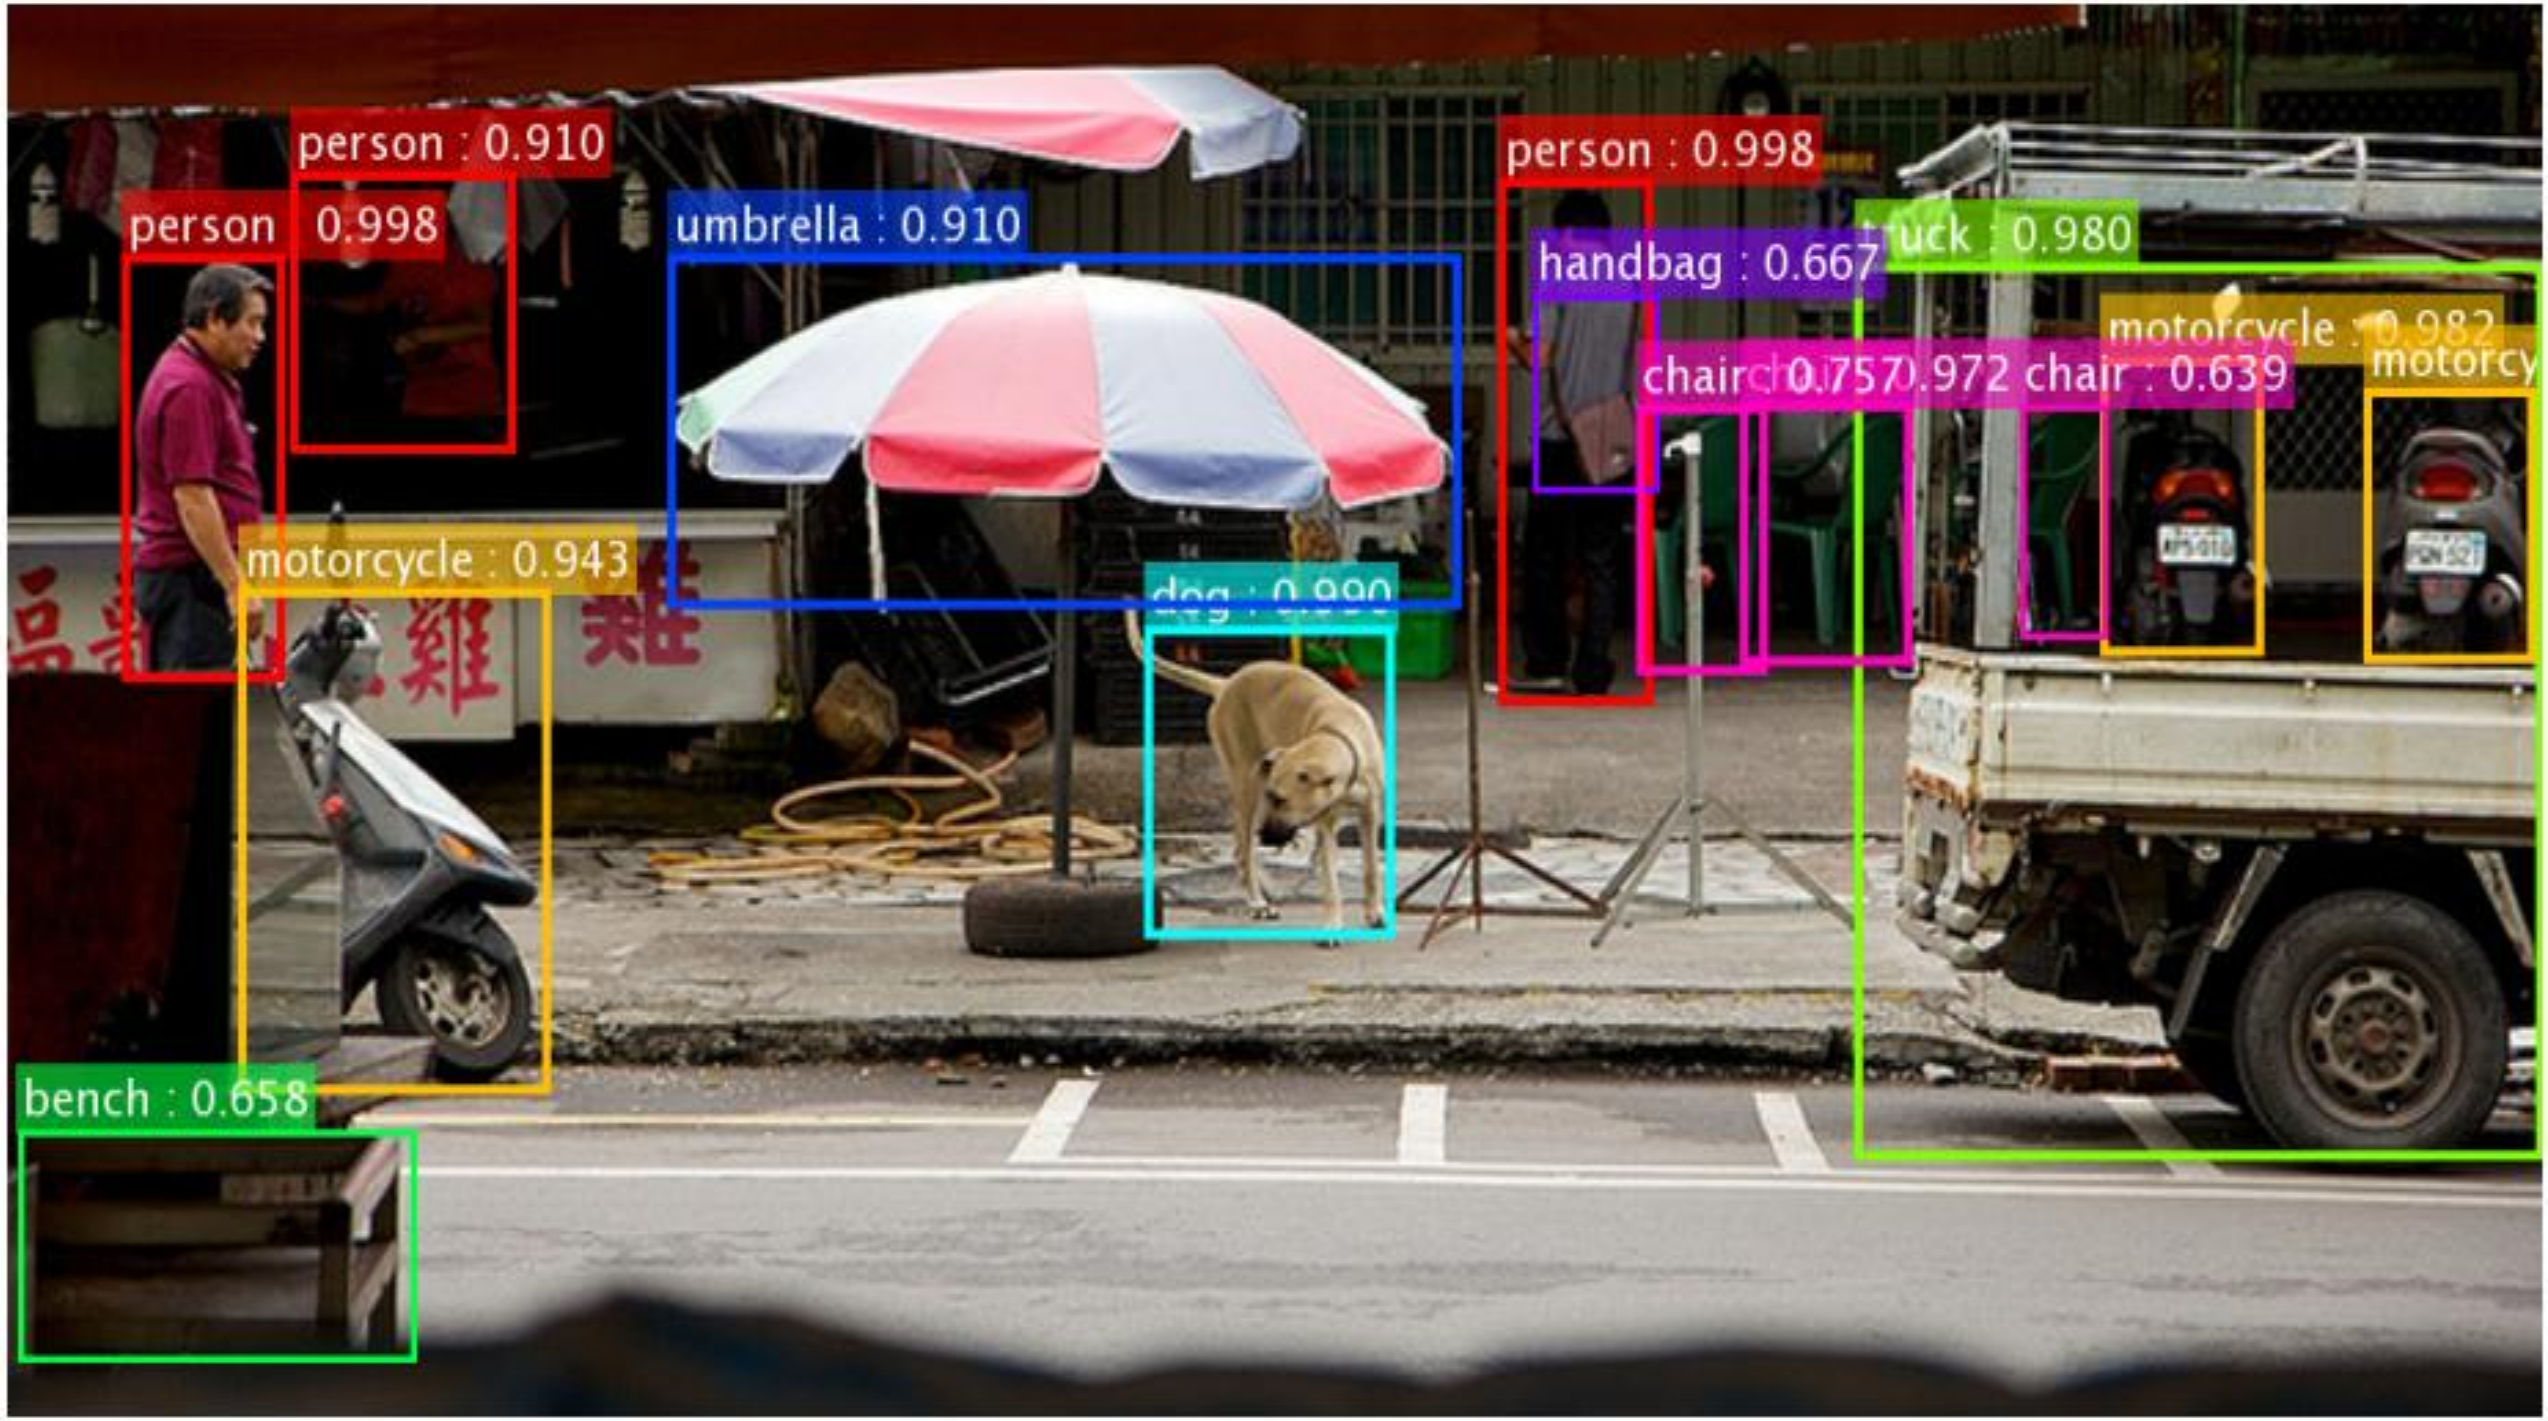
person : 0.910
person
0.998
umbrella : 0.910
person : 0.998
handbag : 0.667
uck : 0.980
motorcycle : 0.982
motorcy
chair 0.757
0.972
chair : 0.639
motorcycle : 0.943
bench : 0.658

Google Brain

Building huge neural networks

10 million connections

1 billion connections

10 billion connections

Desiderata for Learning AI

- Ability to learn complex, highly-varying functions
- Ability to learn multiple levels of abstraction with little human input
- Ability to learn from a very large set of examples
  - Training time linear in the number of examples
- Ability to learn from mostly unlabeled data
  - Unsupervised and semi-supervised
- Multi-task learning
  - Sharing of representations across tasks
- Fast predictions

Thanks for Listening

Any Questions?