دانشگاه کردستان
University of Kurdistan
زانکۆی کوردستان

Department of Computer Engineering
University of Kurdistan

Deep Learning (Graduate level)

Introduction

By: Dr. Alireza Abdollahpouri

Course Info

Instructor

Dr. Alireza Abdollahpouri

Email: abdollahpouri@gmail.com

Office: Room 219

Course Web Page

http://eng.uok.ac.ir/abdollahpouri/DL.html

Grading Policy

Projects	**25%**
Presentation	**15%**
Final exam	**55%**
Class participation	**5%**

Course Info

References

- Understanding deep learning : https://udlbook.github.io/udlbook/
- Deep Learning, Ian Goodfellow, Yoshua Bengio, and Aaron Courville

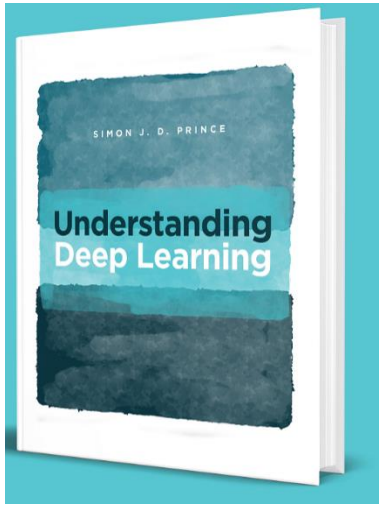

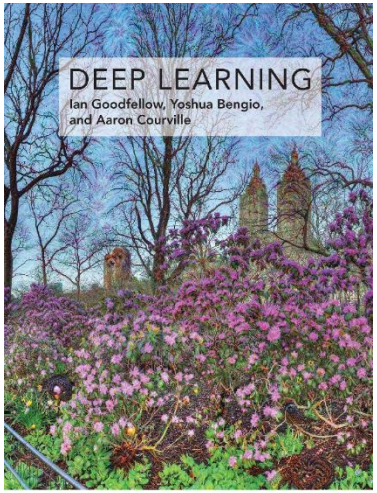

Some Ground Rules

- Let's make this educational and enjoyable.
- It's a big size class, I enjoy questions and ideas from the class.
  - Ask questions and raise points.
  - Listen to other people's questions.
  - Be here.
  - Be here on time.

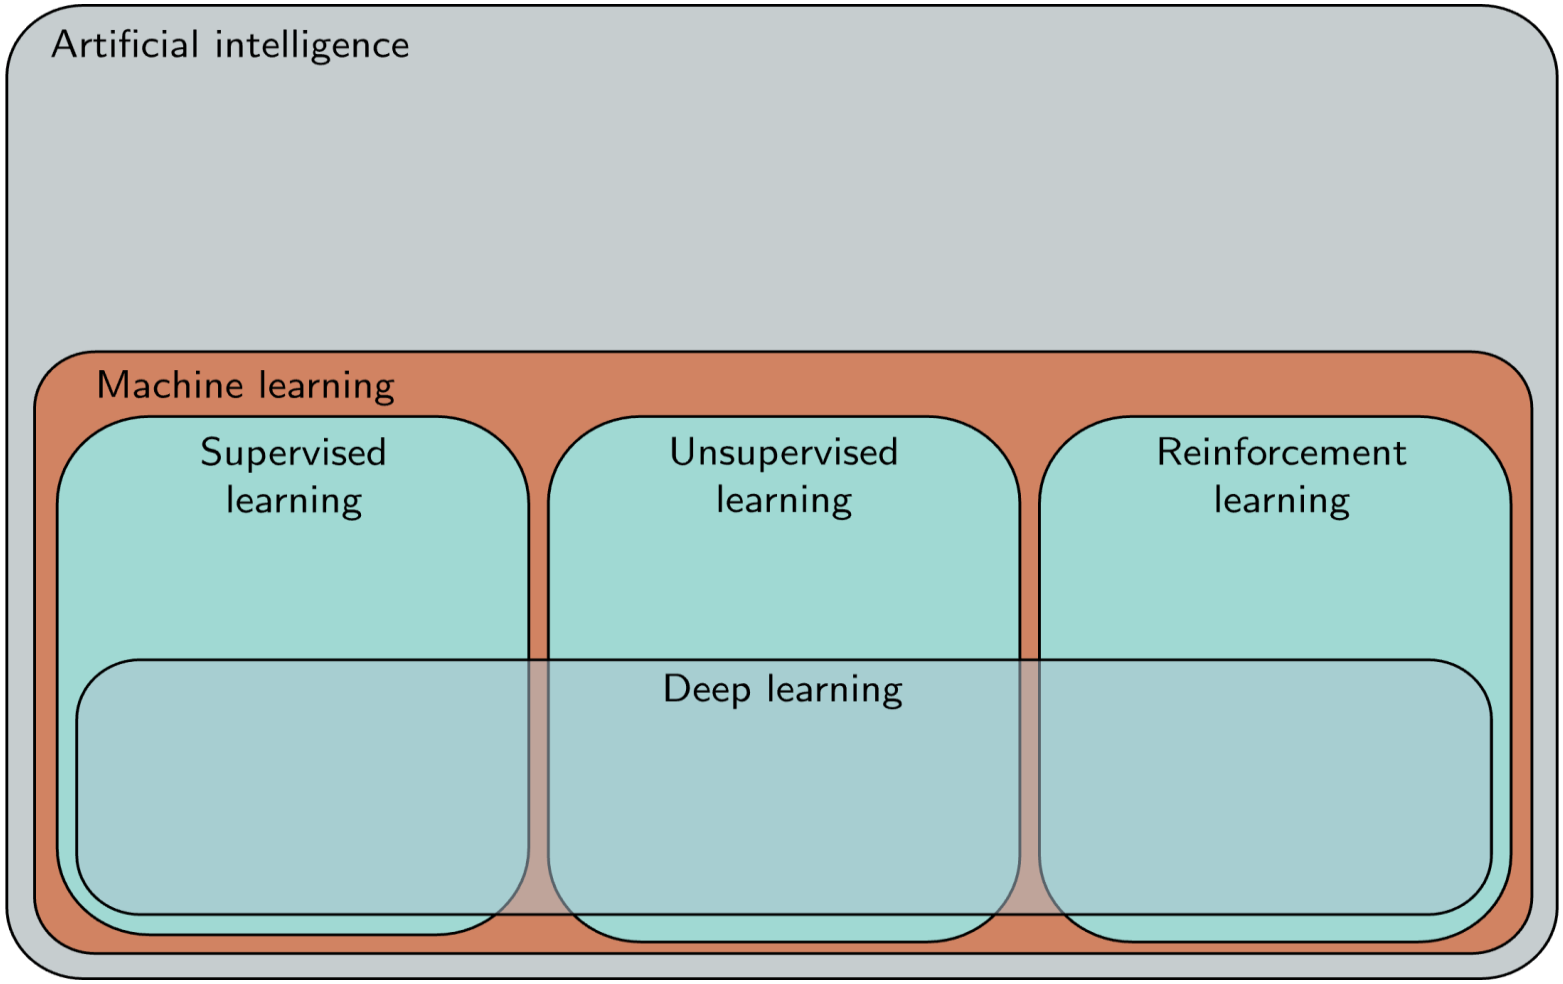
Artificial intelligence
Machine learning
Supervised learning
Unsupervised learning
Reinforcement learning
Deep learning

DL ⊂ NN ⊂ ML ⊂ AI

Artificial Intelligence (AI)

- AI refers to a set of technique that enable computers **to mimic human behavior**.

Machine Learning (ML)

- ML is a subfield of AI , which enables machines **to learn and make prediction based on data**.

Neural Networks (NN)

- NN are a subfield of ML that **use artificial neural networks** to extract patterns from data, inspired by the human brain.

Deep Learning (DL)

- DL is a subfield of NN that **utilizes multi-layered neural networks** to achieve high performance on complex tasks.

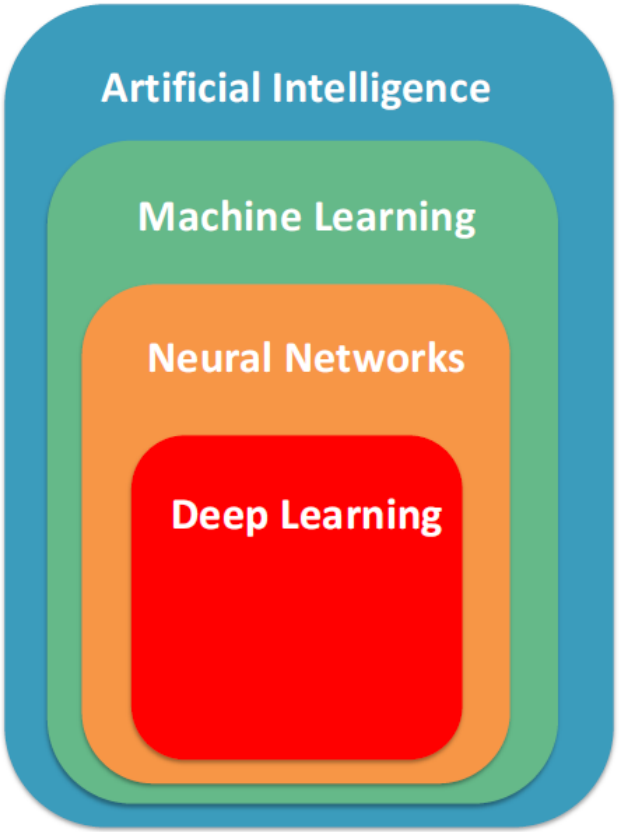

M.L vs. Traditional programming

Traditional Programming: Developers write rules (program) that produce an output.

Machine Learning: Developers write a training algorithm, that finds rules, which produce the desire output.

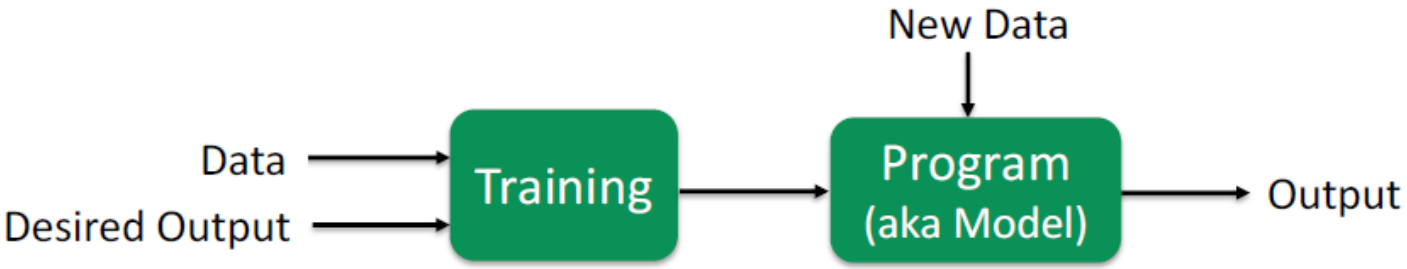

Machine Learning

Machine learning is a subfield of artificial intelligence that studies methods that **learn from data** and make predictions about unseen data.

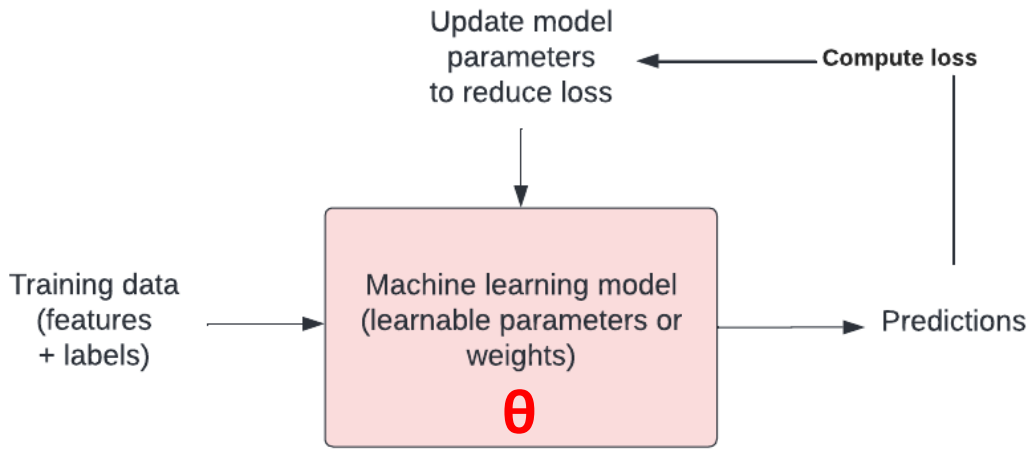

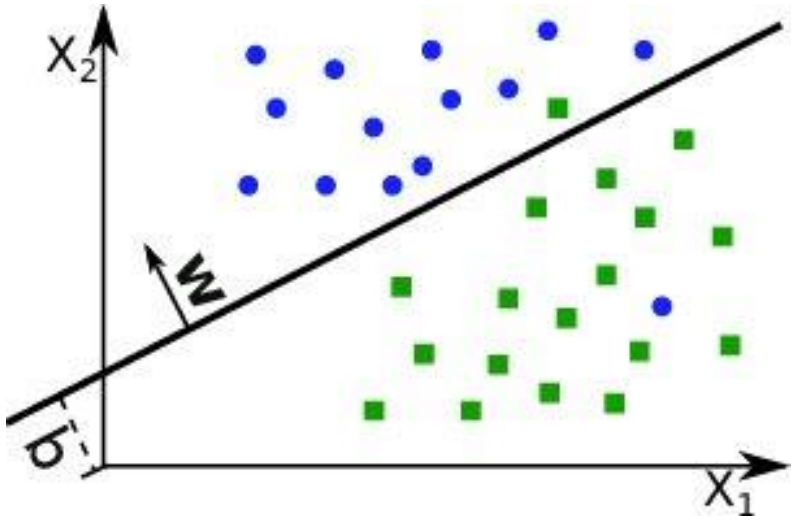

$\theta = \{W , b\}$

ML Methods

SUPERVISED LEARNING

UNSUPERVISED LEARNING

REINFORCEMENT LEARNING

Supervised Learning

Train Data

	pattern	label
Example 1		cat
Example 2		not cat
Example 3		cat
Example 4		not cat
Example 5		not cat
Example 6		cat

- Define a **mapping** from input to output
- Learn this mapping from paired input/output data examples

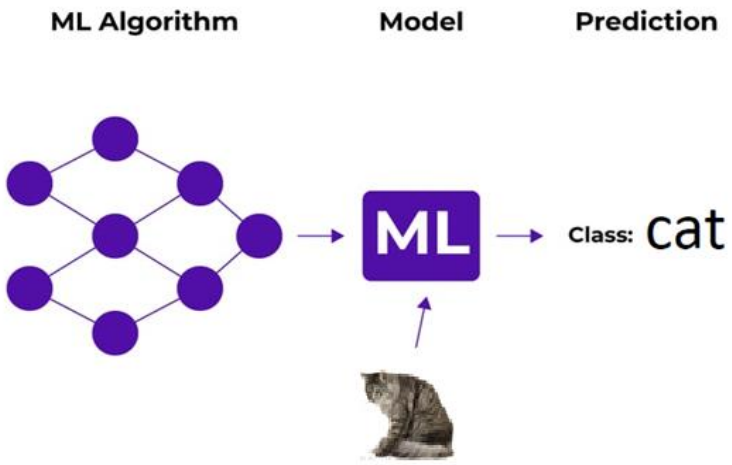

Test Data

Supervised Learning-Regression

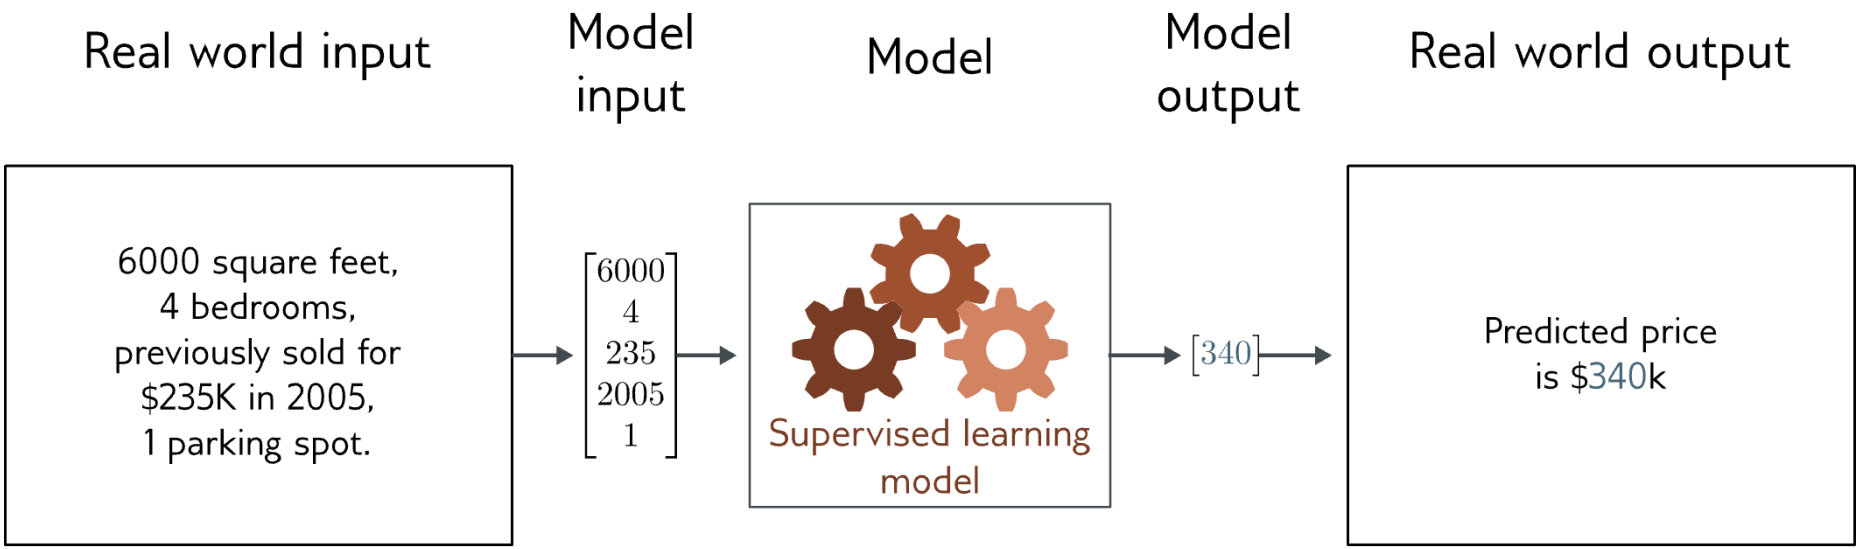

- Univariate regression problem (one output, real value)
- In regression problem, we try to infer a function that maps **continuous** inputs to continuous output

Regression

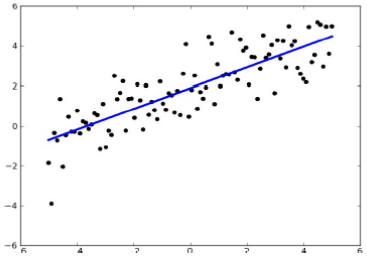

Supervised Learning- Classification

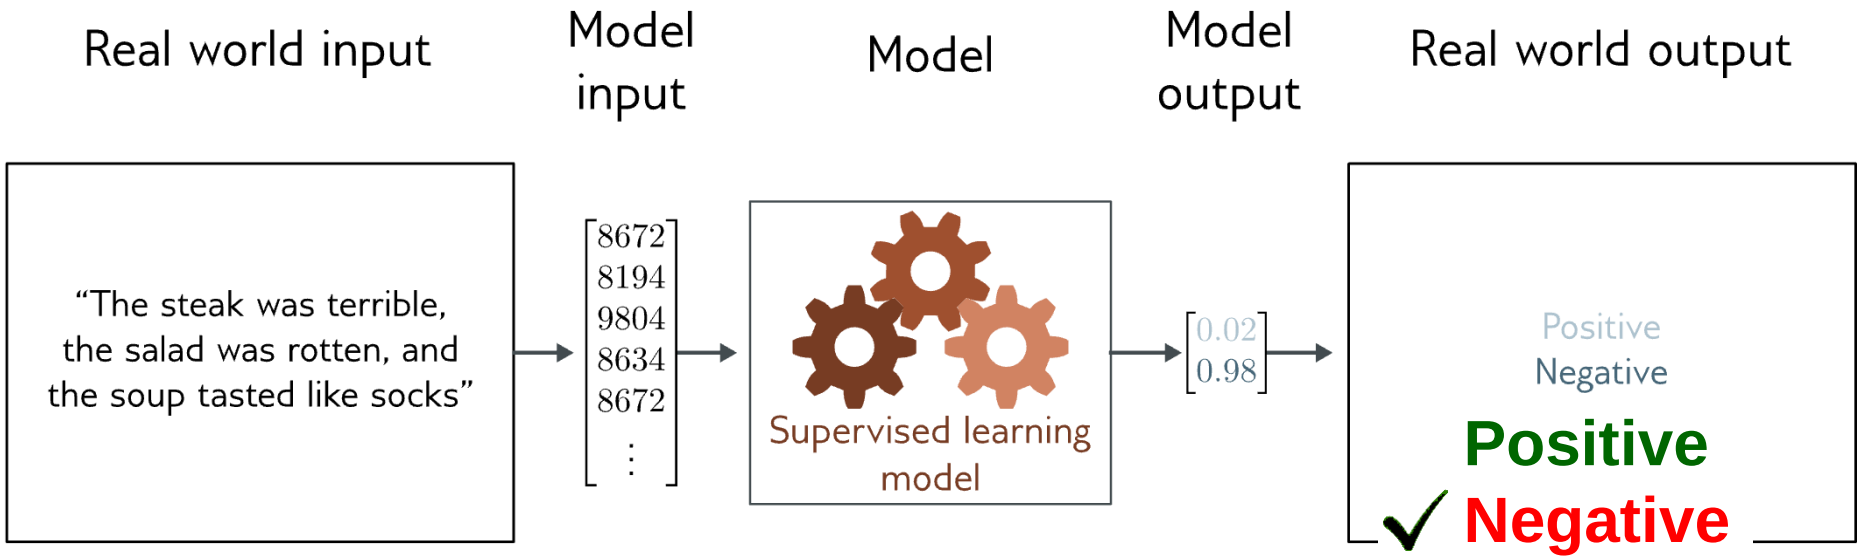

- Binary classification problem (two discrete classes)

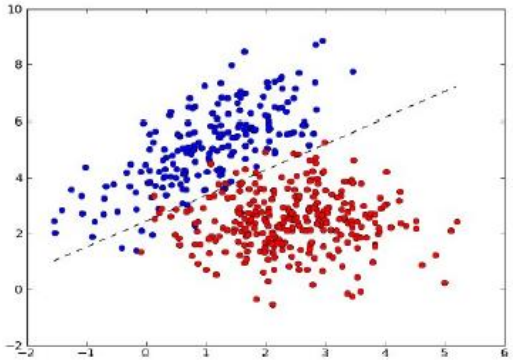

Supervised Learning- Classification

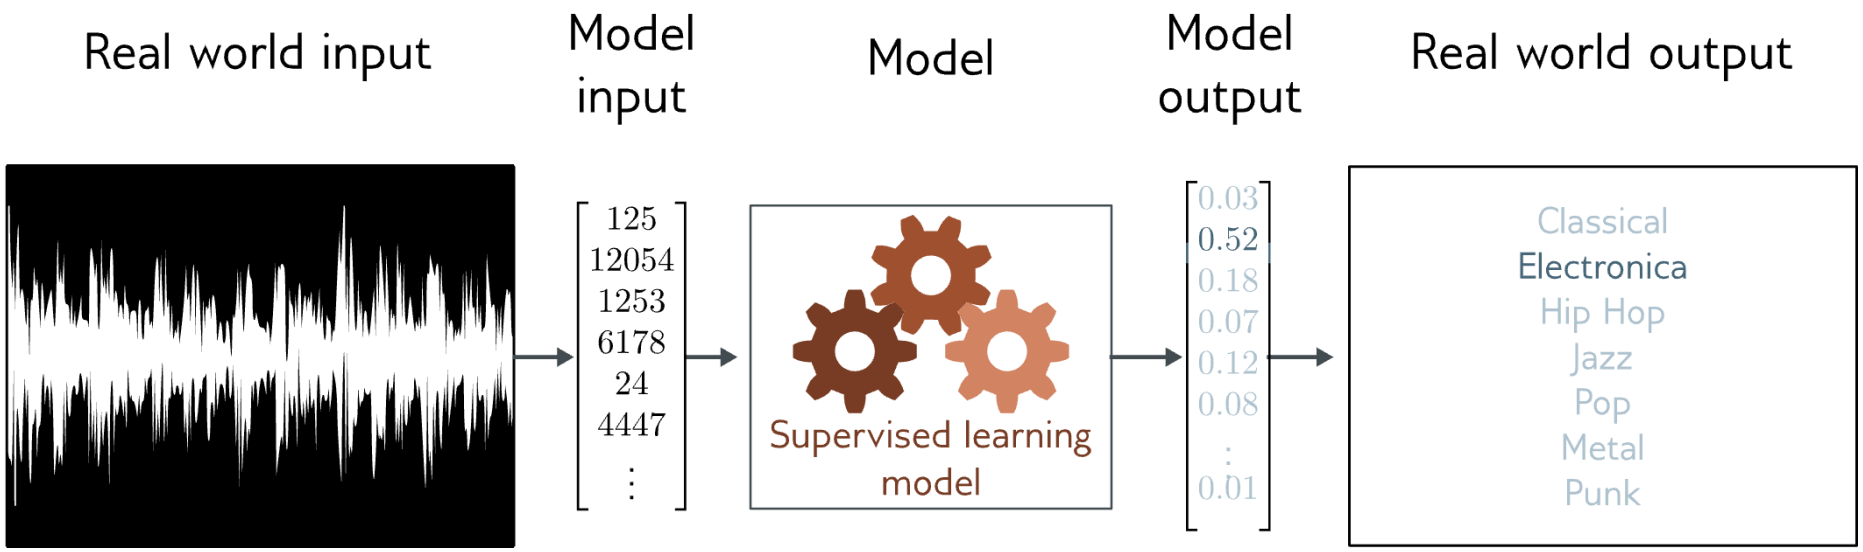

- **Multiclass classification problem (discrete classes, >2 possible values)**

Supervised Learning- Classification

- Binary Classification
- Multi-class Classification
- Multi-label Classification

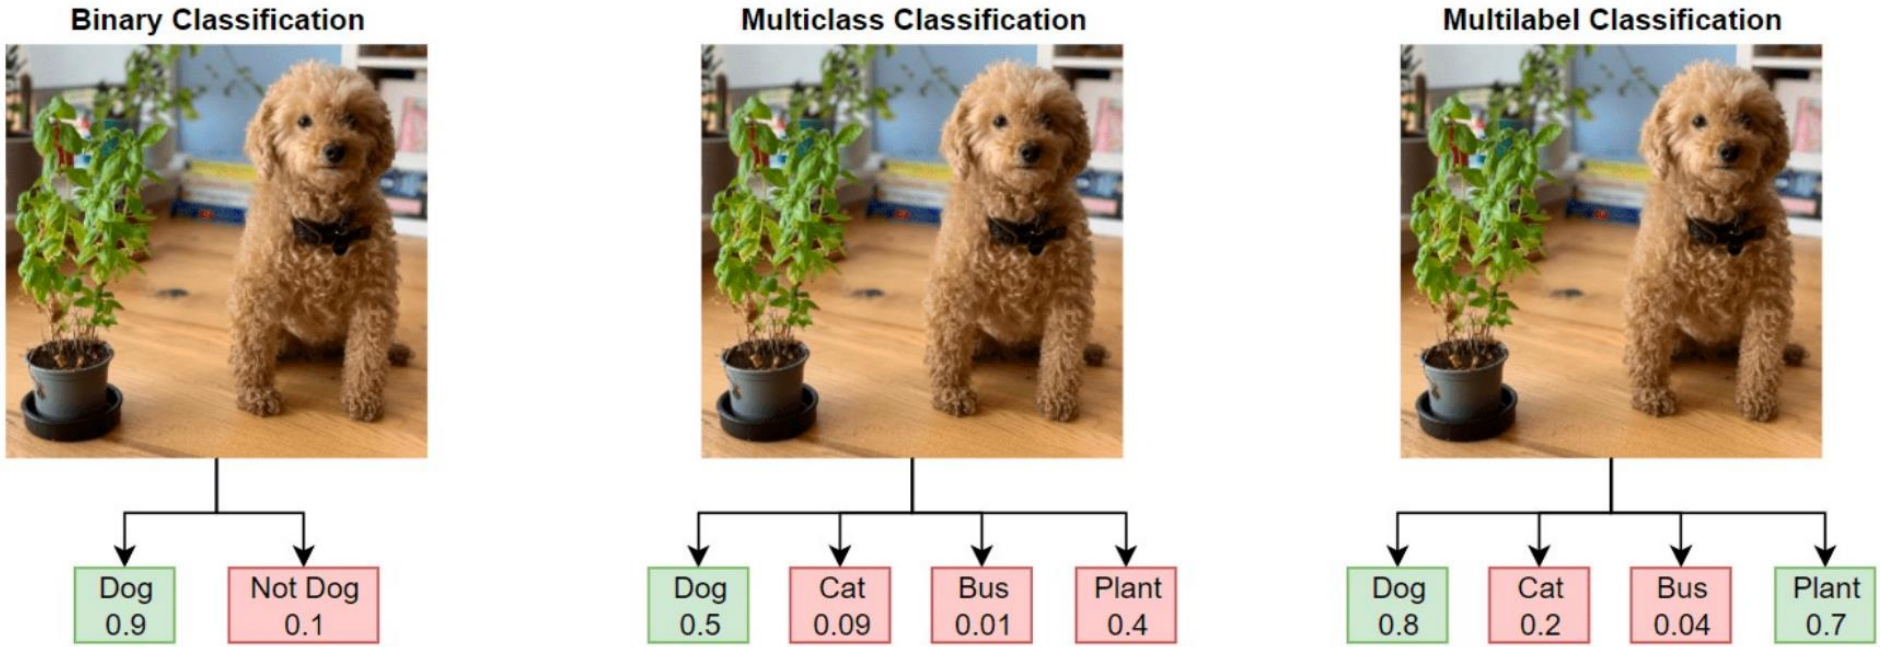

Unsupervised Learning

Un-labeled Data

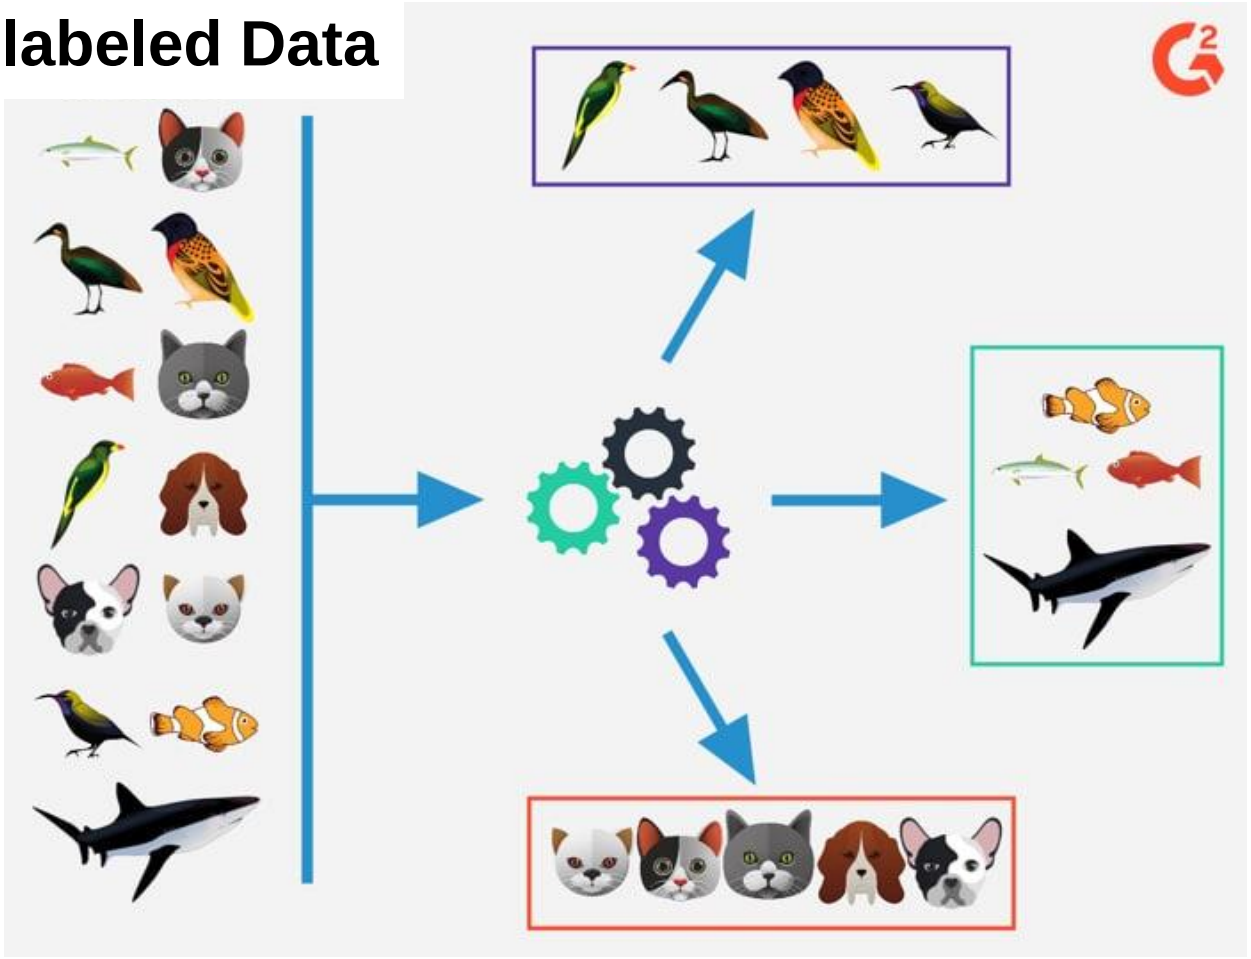

Unsupervised Learning

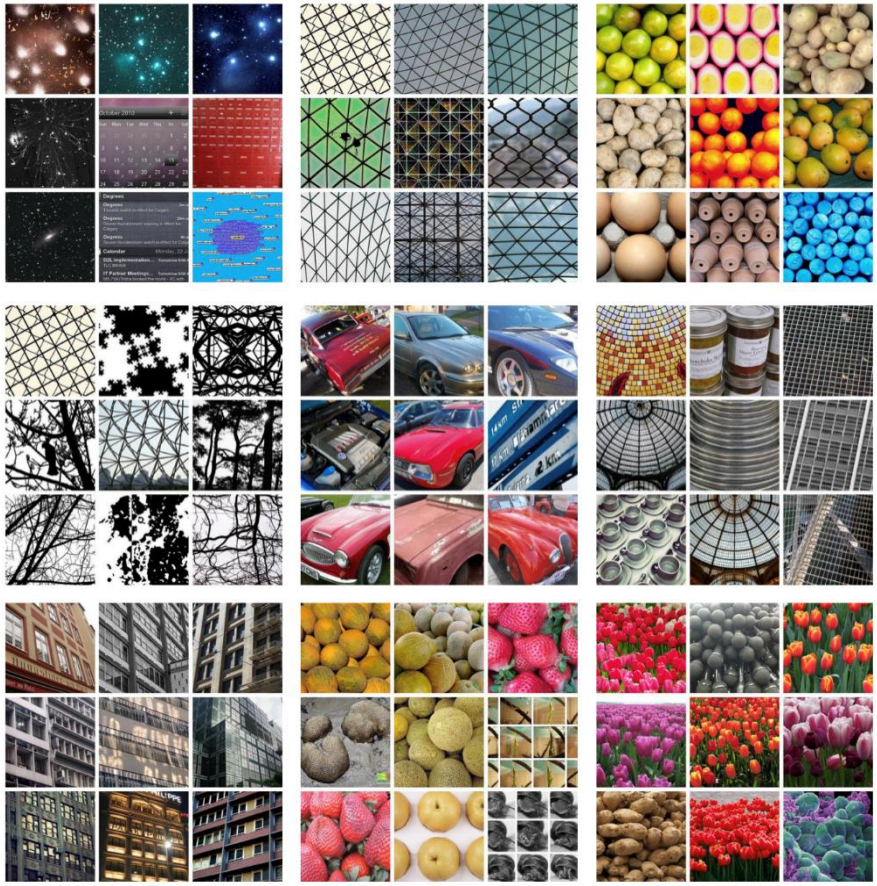

DeepCluster: Deep Clustering for Unsupervised Learning of Visual Features (**Caron et al., 2018)**

Reinforcement learning

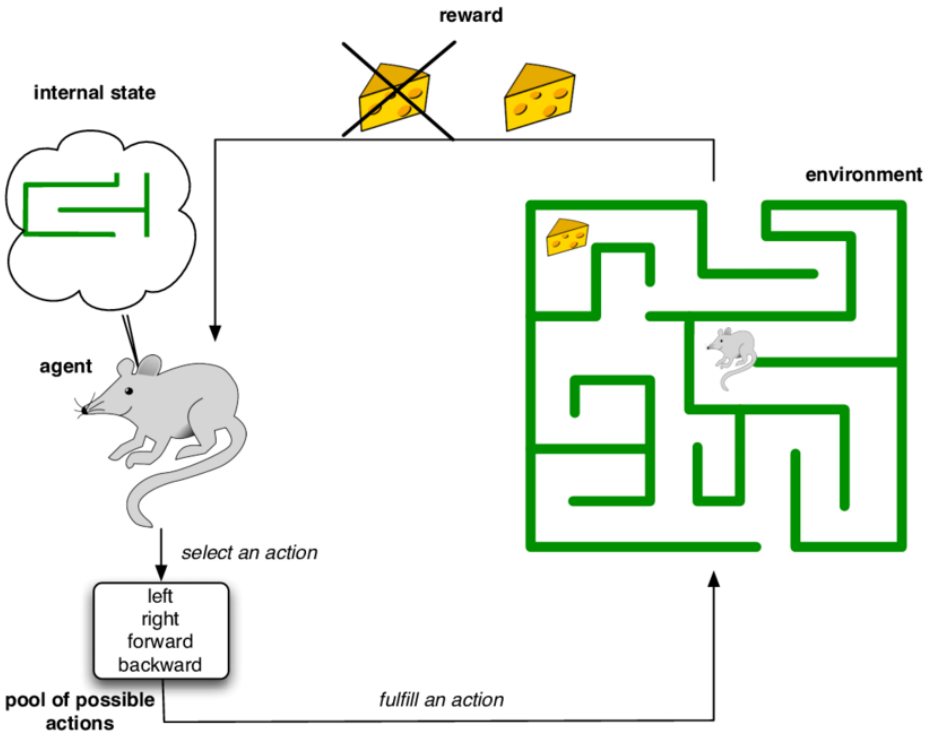

It is a branch of machine learning that teaches machines how to make decisions in a **dynamic environment** to **maximize their reward**. This learning method, inspired by behavioral psychology, operates on a simple principle: if an action leads to a desired outcome (reward), the likelihood of repeating that action in the future increases.

Deep Learning vs. Machine Learning

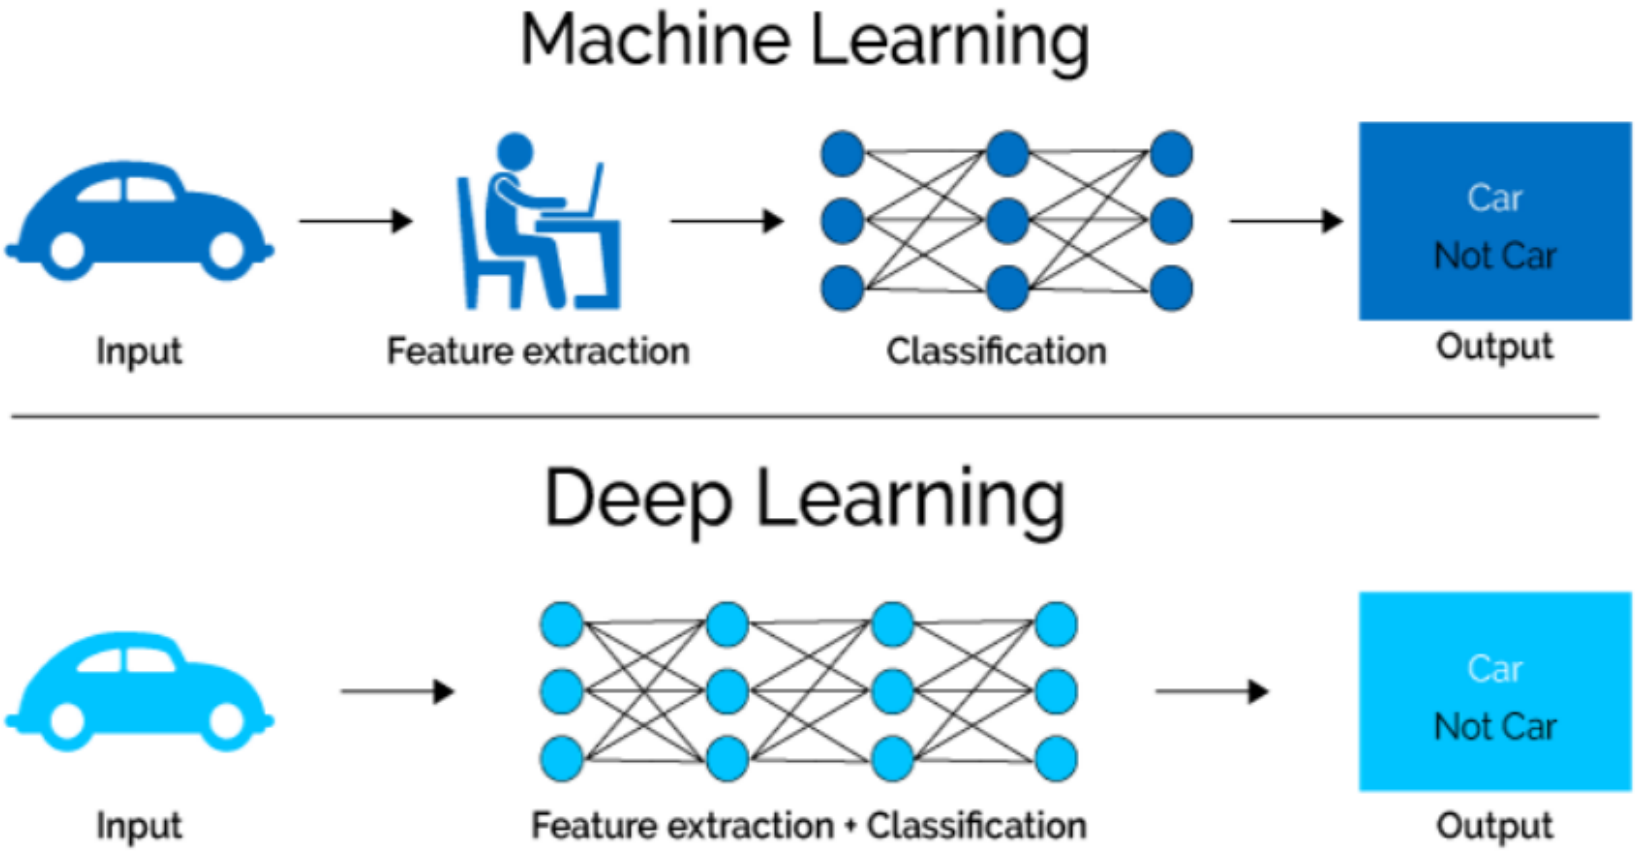

Feature Extraction In ML

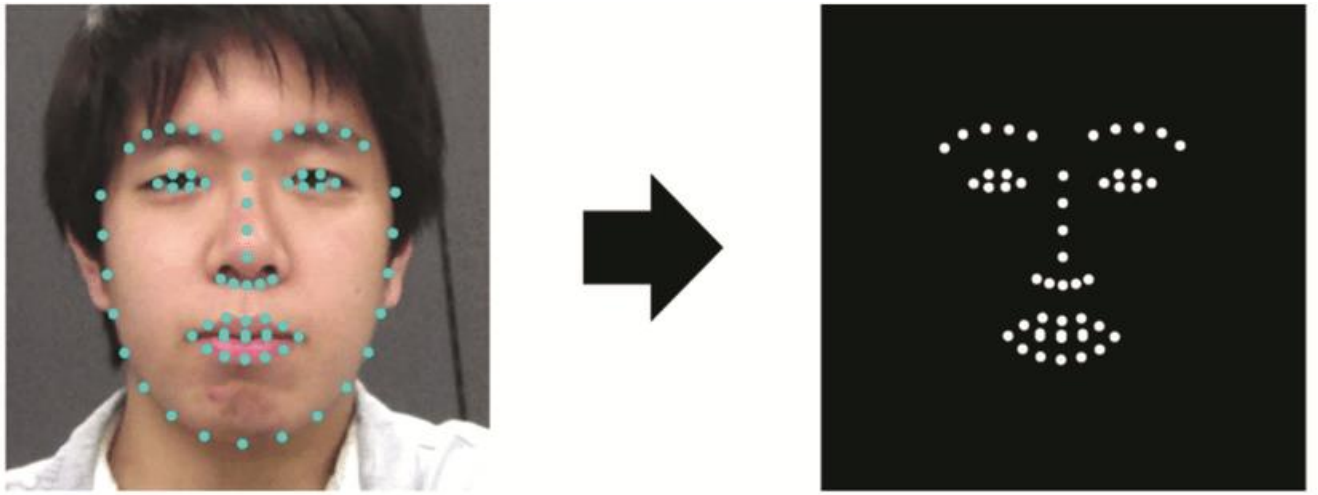

(a) Detected facial keypoints

(b) Facial organ keypoints

- **It is time-consuming**
- **Requires advanced technical skill**
- **Doing it manually can lead to errors and bias**

Deep Learning

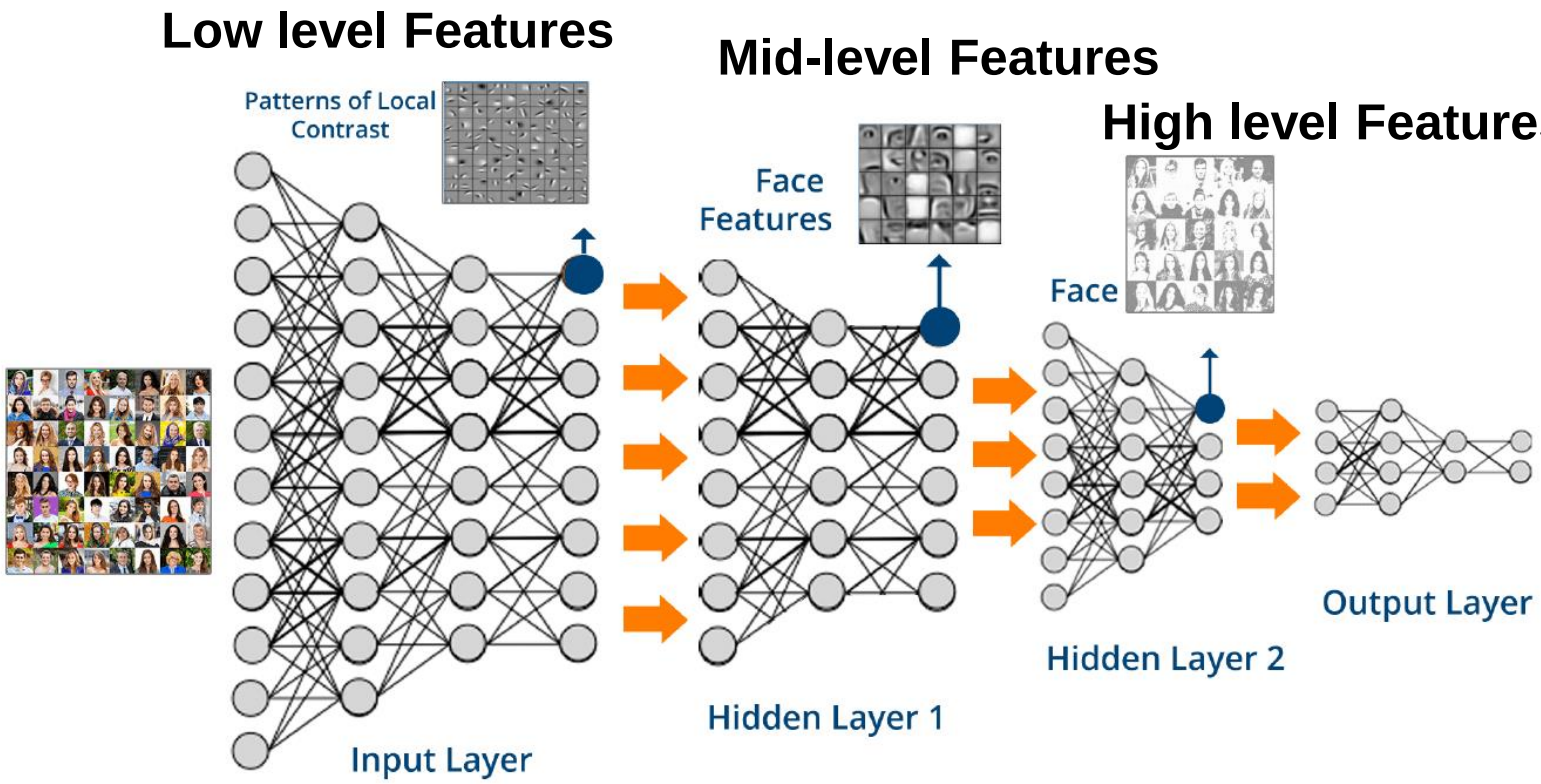

Deep learning is the process of extracting hierarchical representations from an input.

What Factors Enable Deep Learning now?

1952 Stochastic Gradient Descent

1958 Perceptron
- Learnable Weights

1986 Backpropagation
- Multi-Layer Perceptron

1995 Deep Convolutional NN
- Digit Recognition

Neural Networks date back decades, so why the resurgence?

1. Big Data
- Larger Datasets
- Easier Collection & Storage

IMAGENET

2. Hardware
- Graphics Processing Units (GPUs)
- Massively Parallelizable

3. Software
- Improved Techniques
- New Models
- Toolboxes

TensorFlow

Capabilities of Human brain

- Face recognition
- Learning
- Decision making
- Calculation
- Information storage
- …

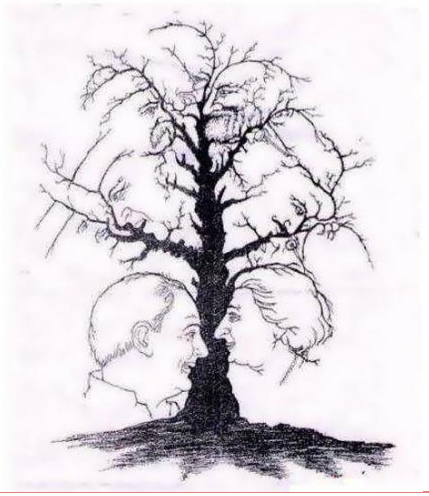

Capabilities of Human brain

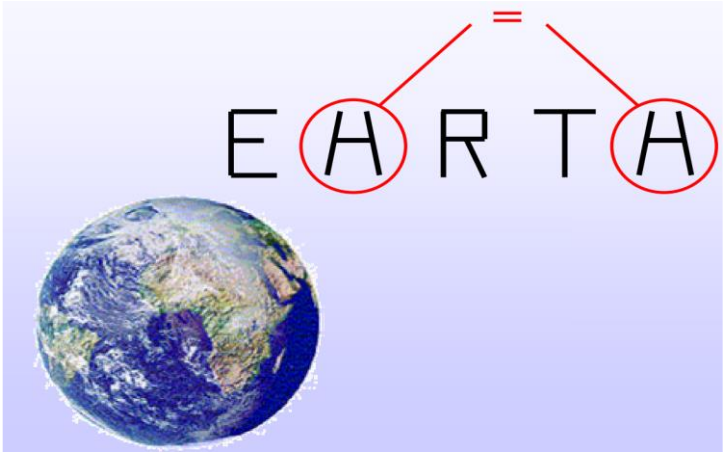

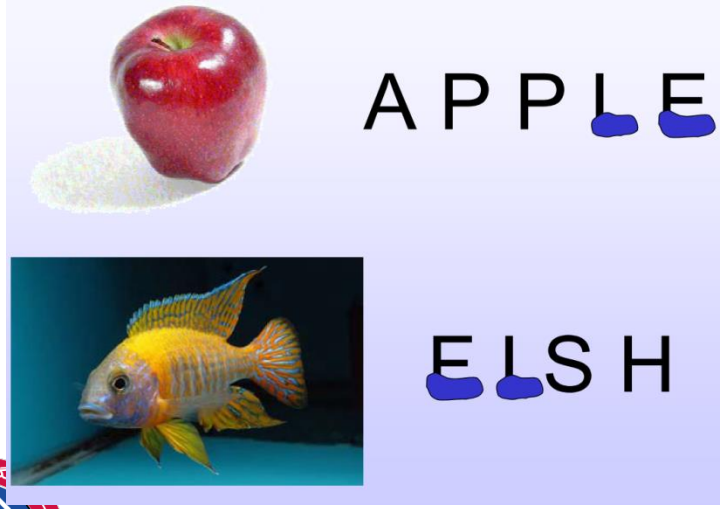

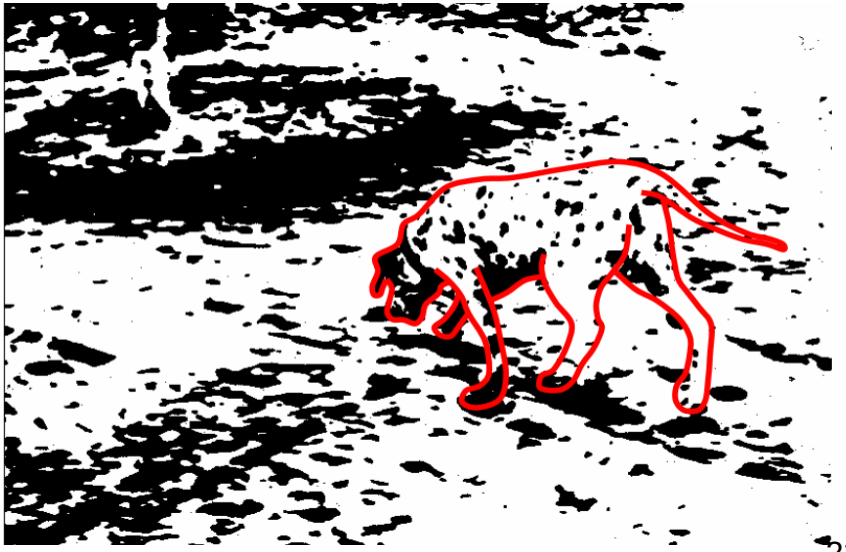

Capabilities of Human brain

The mouse on the desk

The mouse on the desk

Capabilities of Human brain

The mouse on the desk is broken

The mouse on the desk is eating cheese

How does the brain work?

- When we recognize a face or we grasp an object we do not solve equations.
- The brain works using an **associative** process.

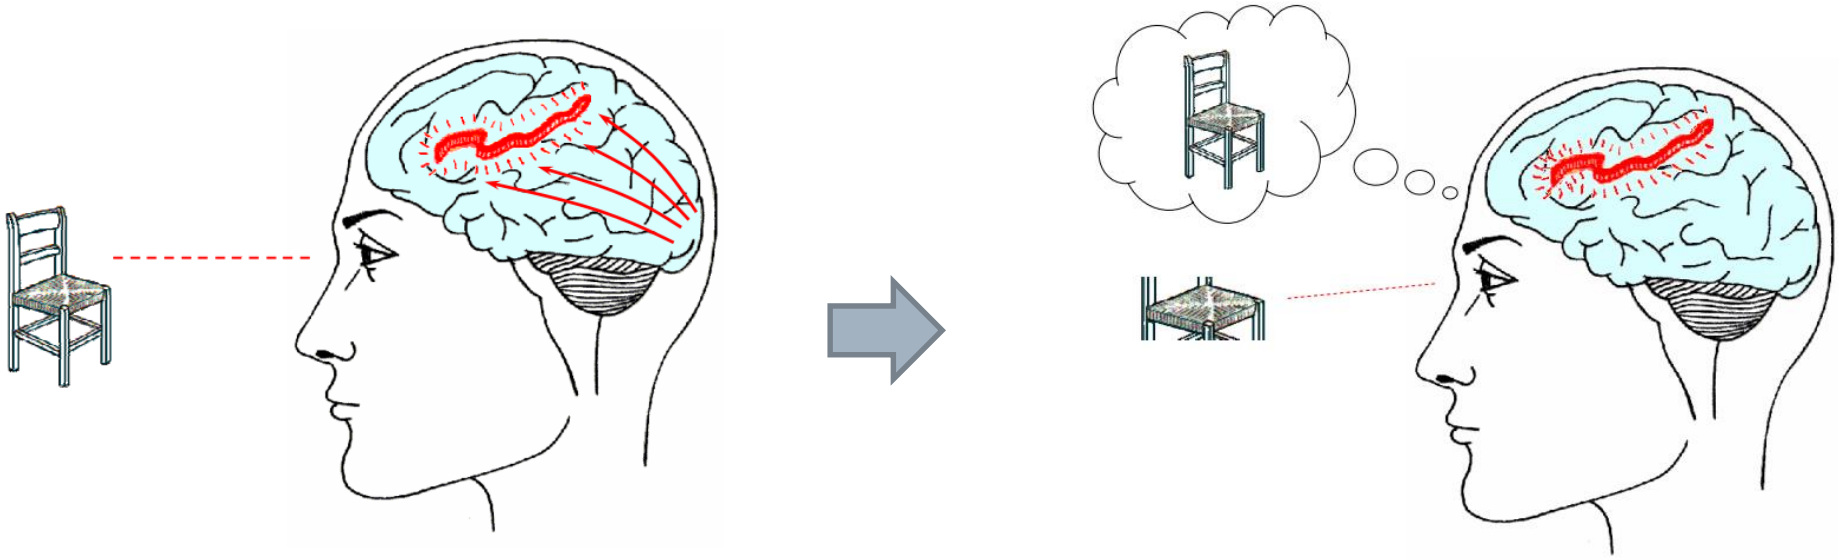

Hitting a tennis ball

Learning phase

- In a learing phase, we try several actions and store the good ones:
  - If the ball is seen in the upper-left area of the visual field, then make a back step;
  - If the ball …

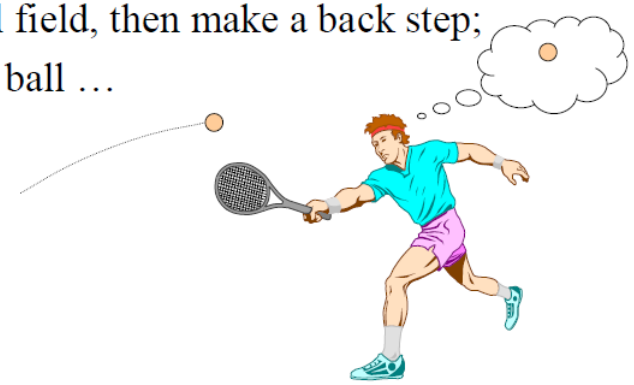

Operating Phase

- Once trained, the brain executes the actions *without thinking*, based on the learned associations.

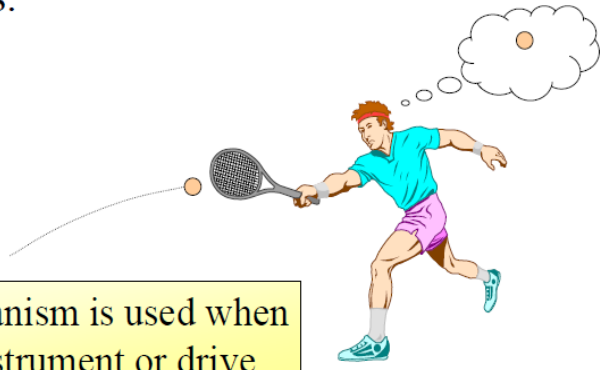

A similar mechanism is used when we play an instrument or drive

Quotation

"If the human brain were simple enough for us to understand, we would still be so stupid that we couldn't understand it"

Jostein Gaarder

Biological Inspiration

Idea : To make the computer more robust, intelligent, and learn, …
Let's model our computer software (and/or hardware) after the brain

- **about 86 billion neurons** in the human cortex each connected to , on average, **10000** others.
- In total **80-1000 trillion synapses** of connections.
- The brain is a highly complex, nonlinear and parallel computer (information-processing system)

Human nervous structure

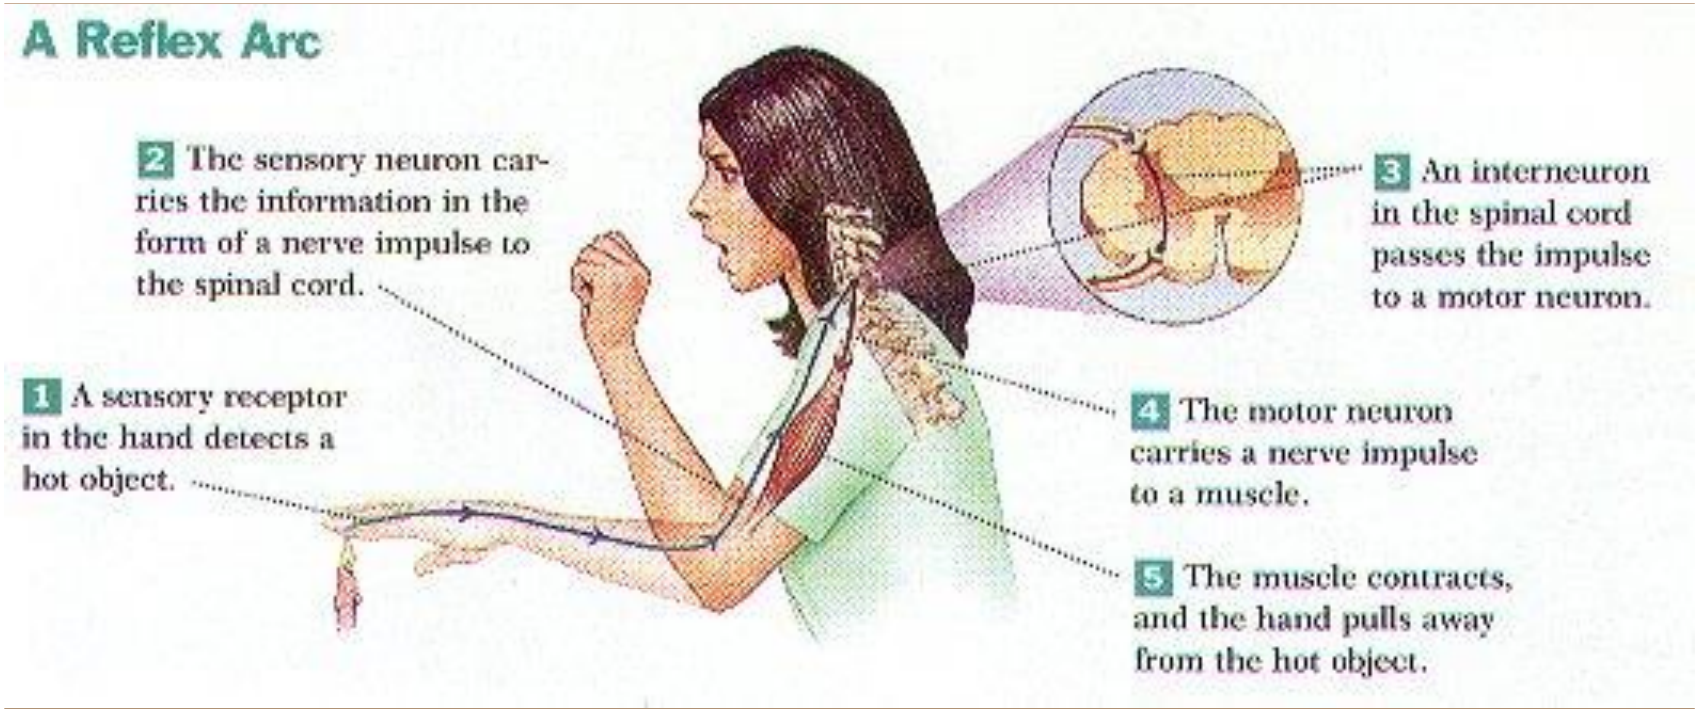

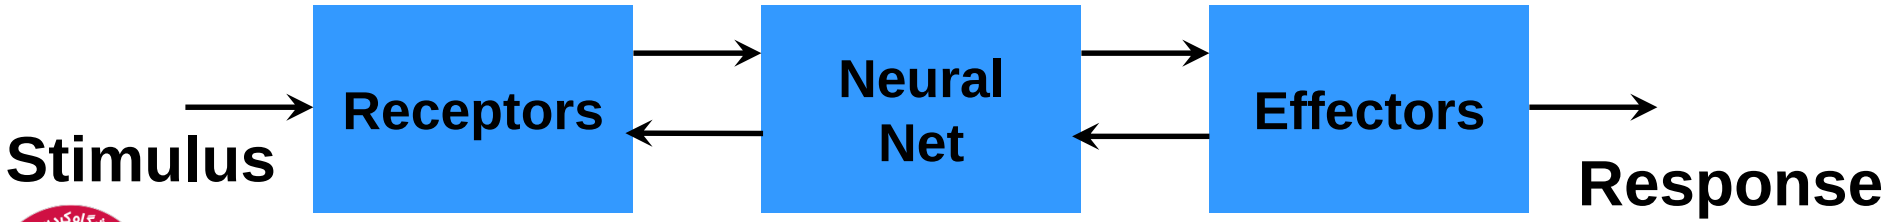

Human brain

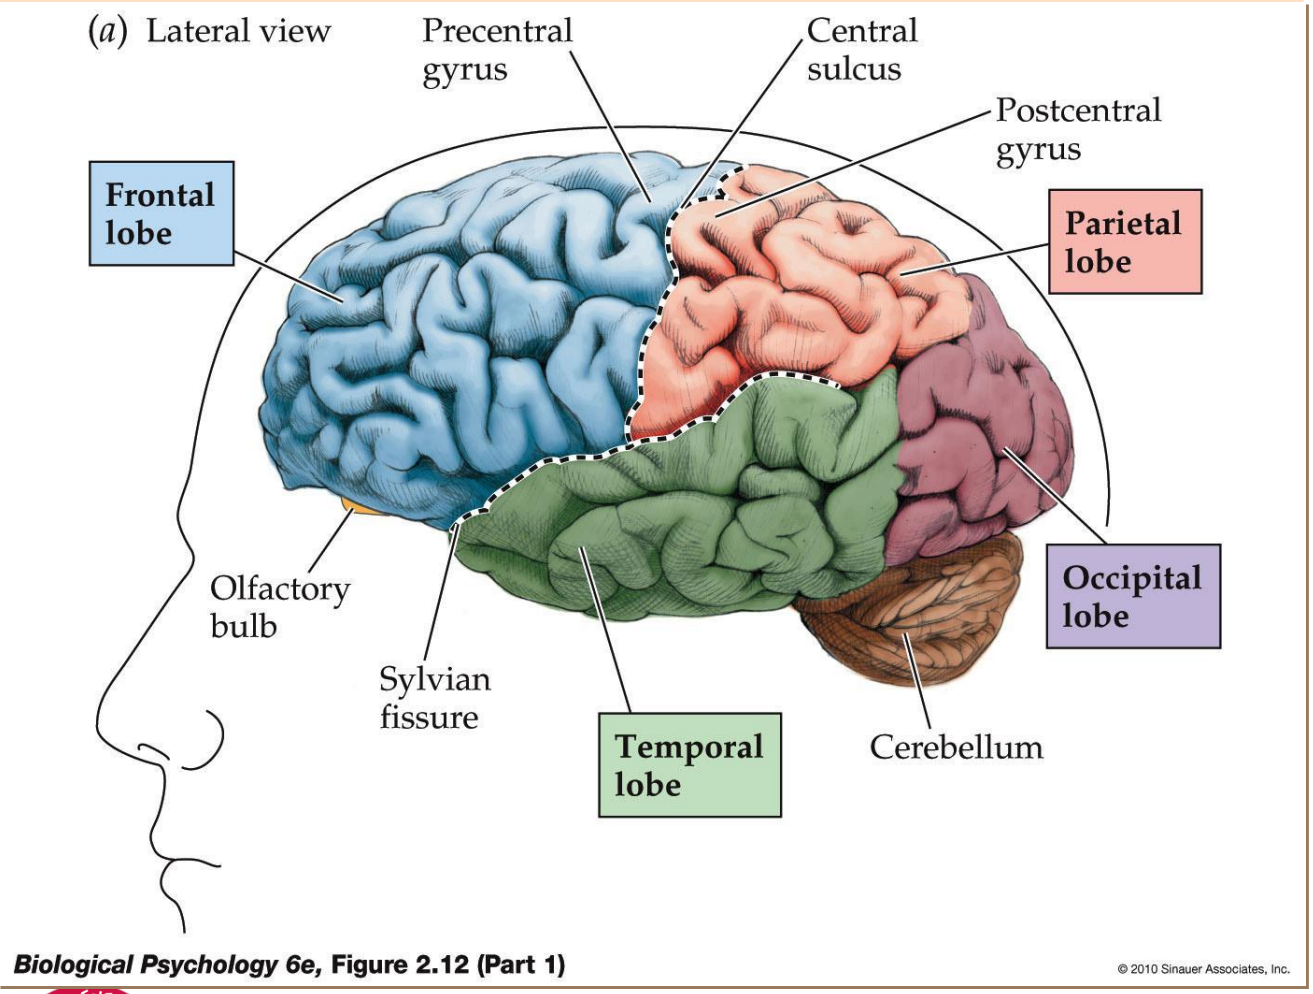

Biological Psychology 6e, Figure 2.12 (Part 1)

© 2010 Sinauer Associates, Inc.

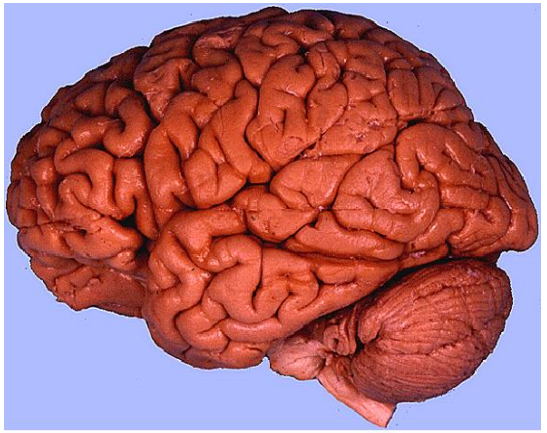

Growth of Neural Density in a Human Brain

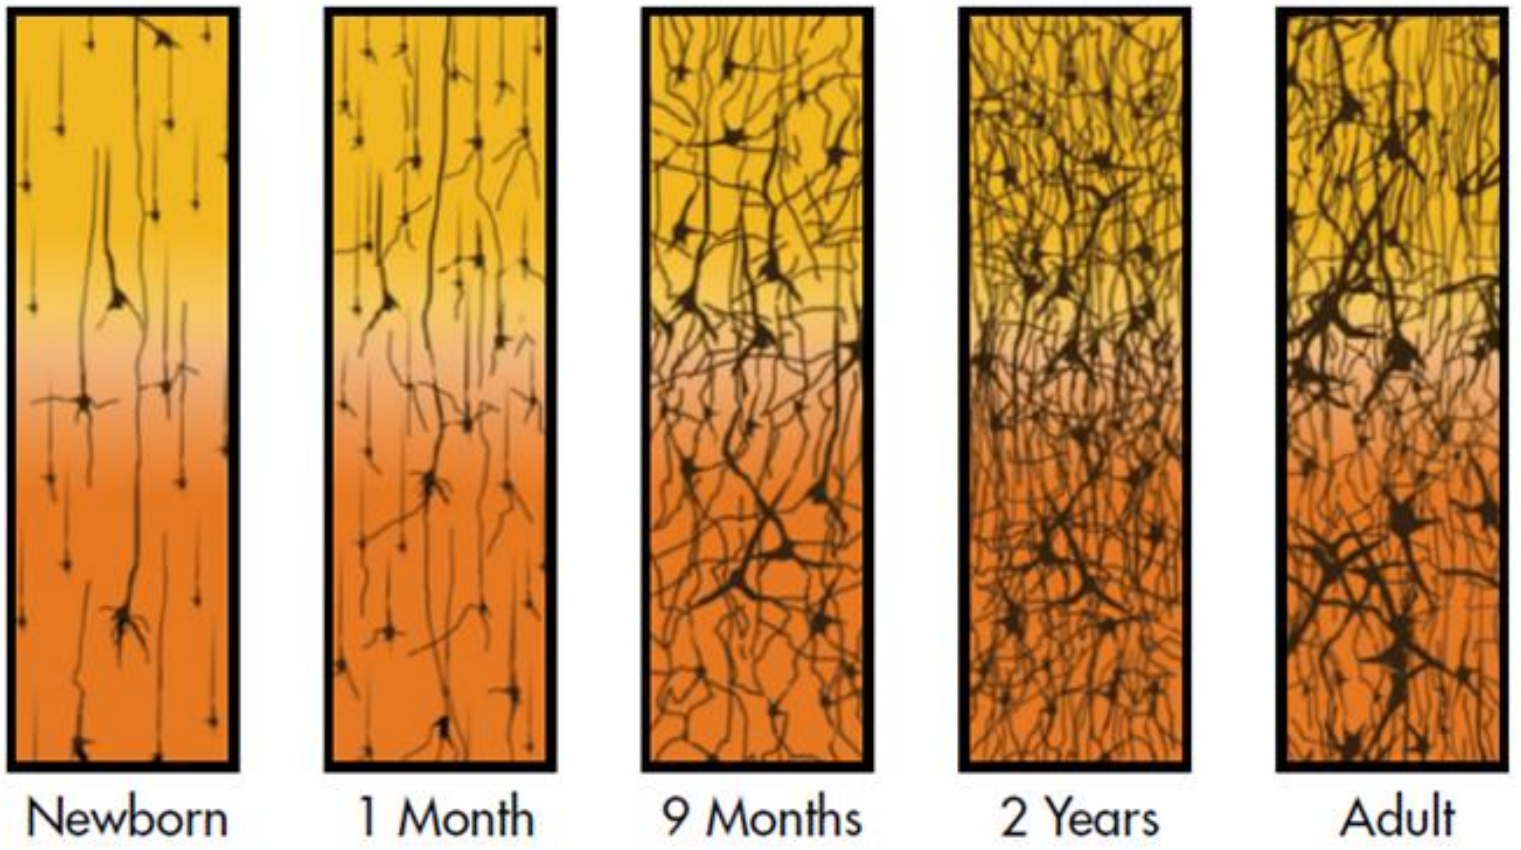

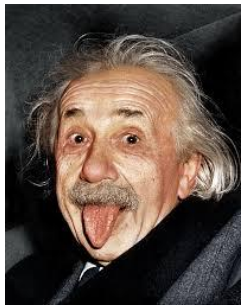

http://www.urbanchildinstitute.org/why-0-3/baby-and-brain

Biological Neuron

Cell structures

Cell body

Dendrites

Axon

Synaptic terminals

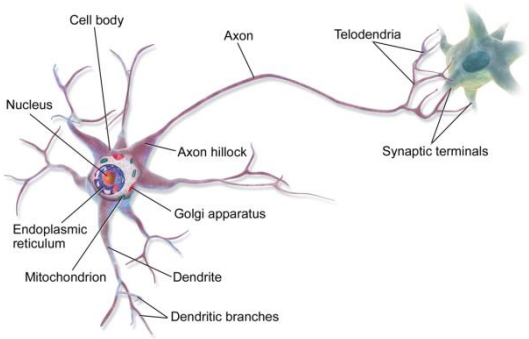

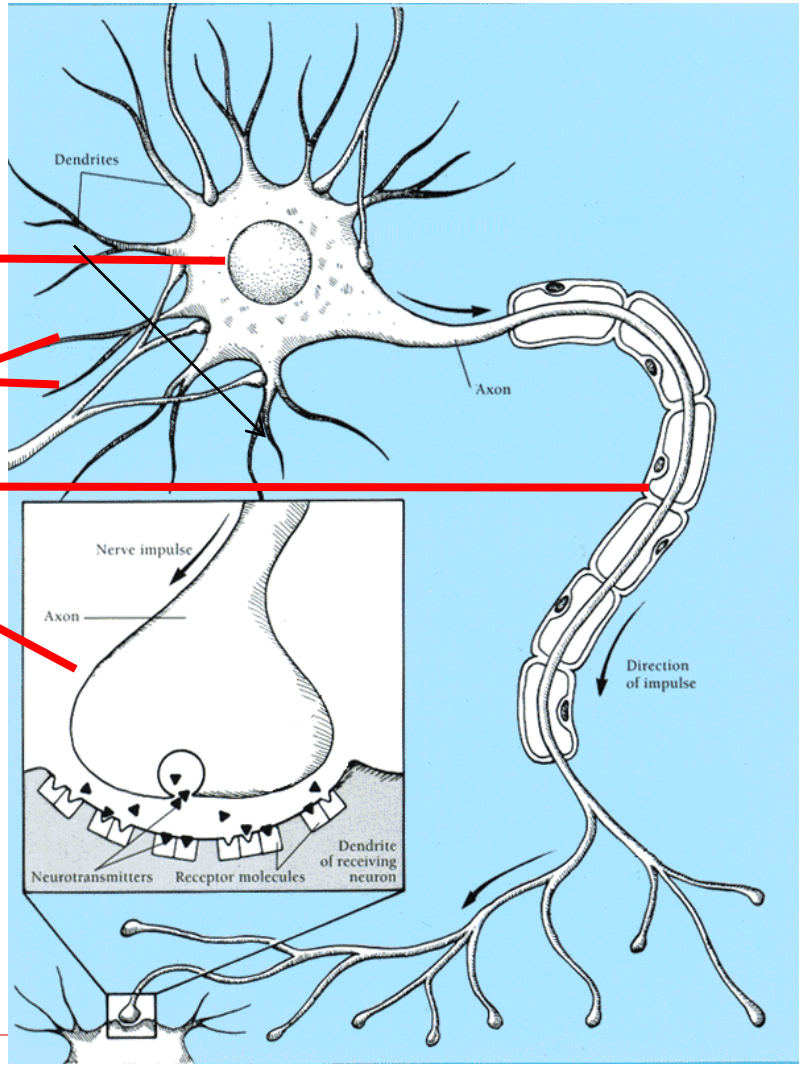

Inter-Neural Communication

Communication is Electrochemical:

Electrical (via ions) along axons

Chemical (via molecules) across synapse

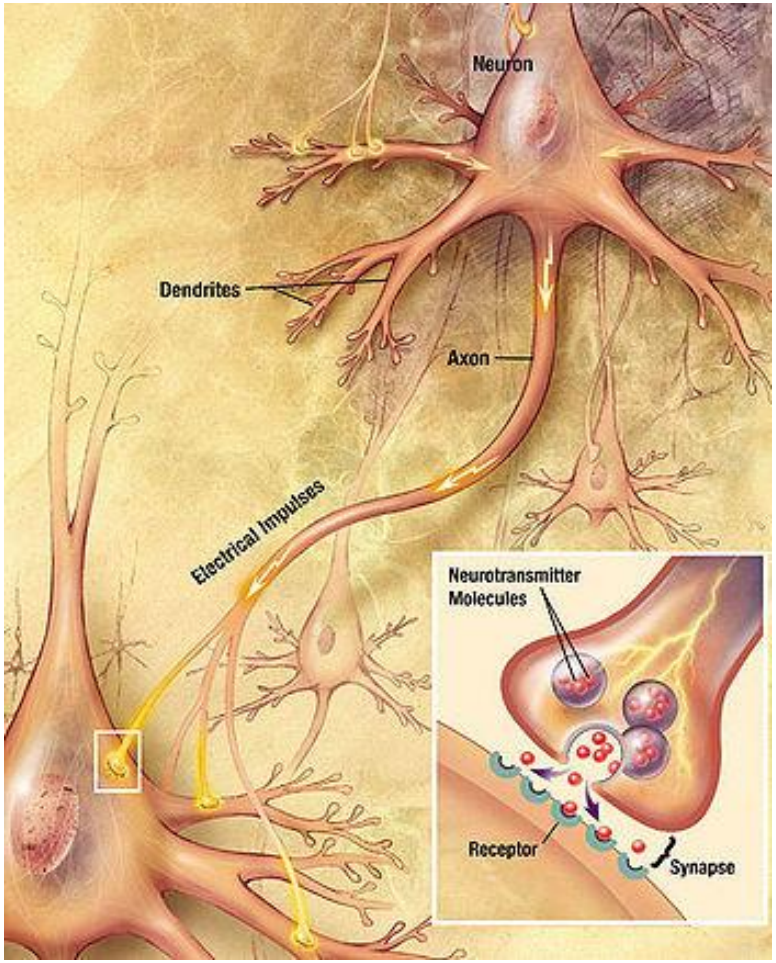

http://en.wikipedia.org/wiki/Neurons

Inter-Neural Communication

The arrival of a nerve impulse stimulates the release of neurotransmitters from vesicles. They pass across the synapse and open channels in the target cell. **Charged particles can then enter** and trigger a second impulse.

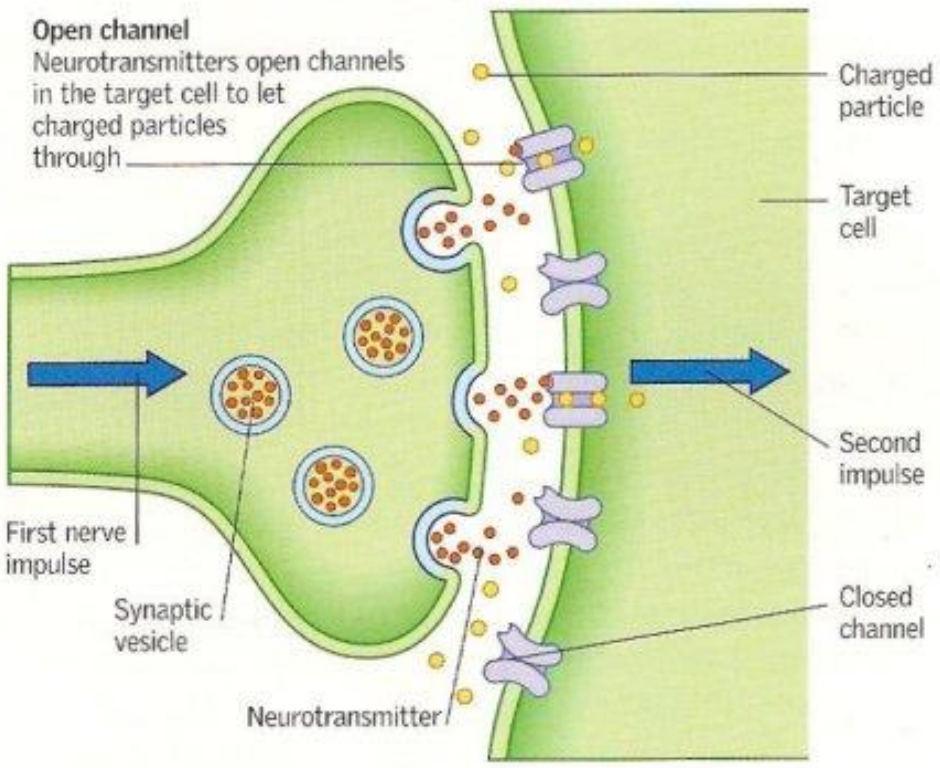

http://www.daviddarling.info/encyclopedia/N/neurotransmitter.html

Action potential

A neuron receives input from other neurons (typically many thousands). Inputs sum (approximately). Once input exceeds a critical level, the neuron discharges a spike - an electrical pulse that travels from the body, down the axon, to the next neuron(s) (or other receptors)

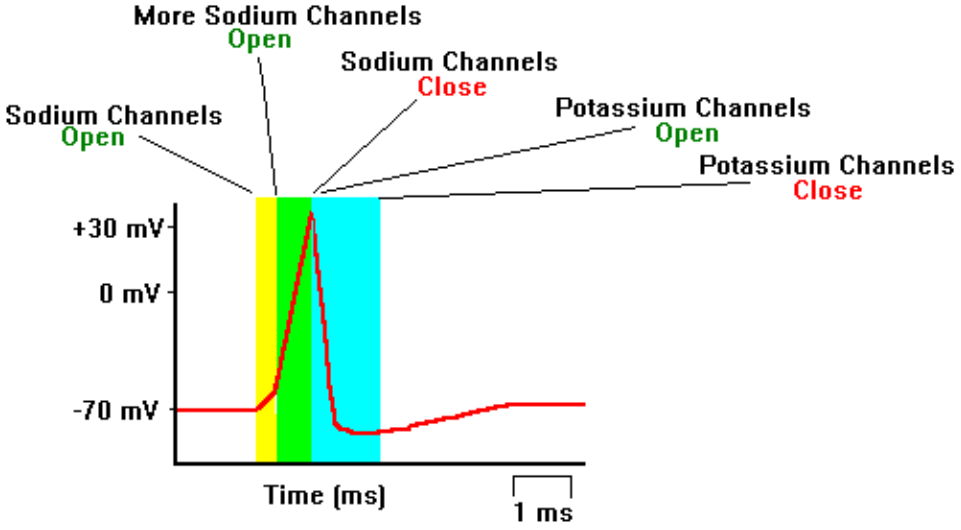

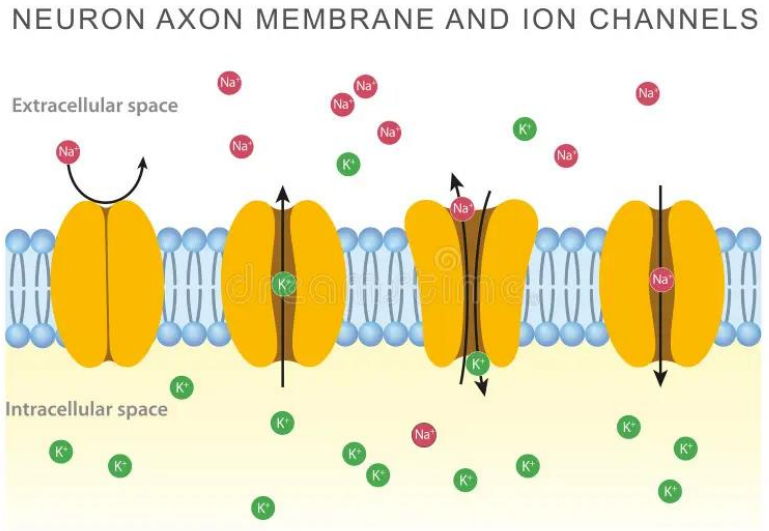

Models of A Neuron

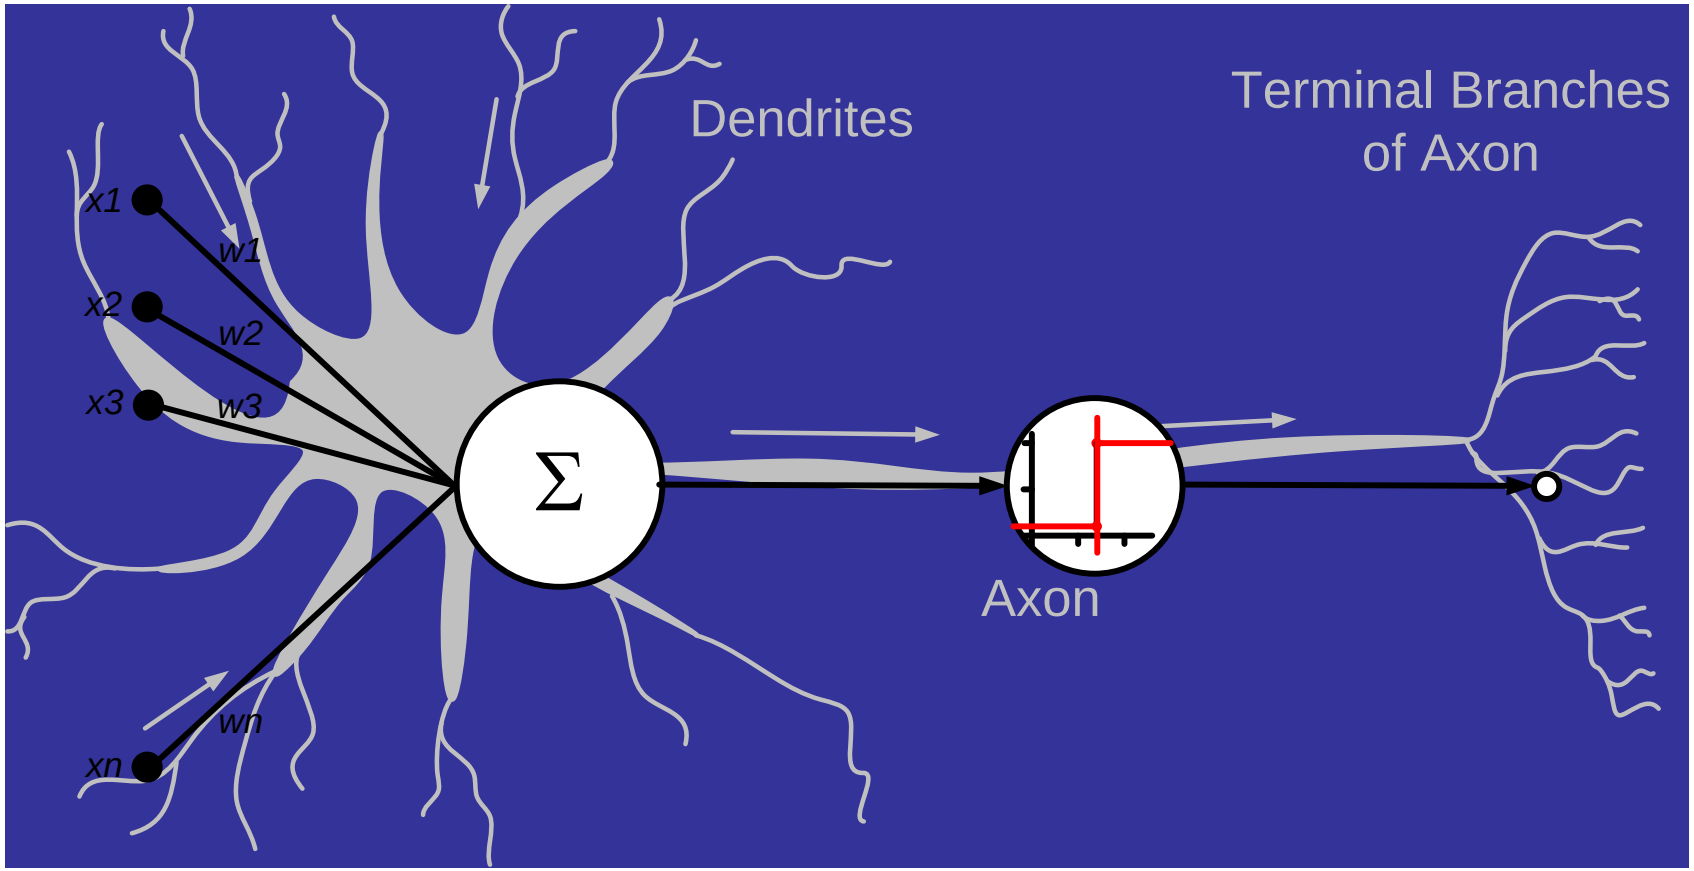

Models of A Neuron

Three elements:

1. **A set of synapses, or connection link:** each of which is characterized by a *weight* or strength of its own w_{kj}. Specifically, a signal x_j at the input synapse 'j' connected to neuron 'k' is multiplied by the synaptic w_{kj}
2. **An adder:** For summing the input signals, weighted by respective synaptic strengths of the neuron in a linear operation.
3. **Activation function:** For limiting of the amplitude of the output of the neuron to limited range. The activation function is referred to as a Squashing (i.e. limiting) function {interval [0,1], or, alternatively [-1,1]}

The Neuron

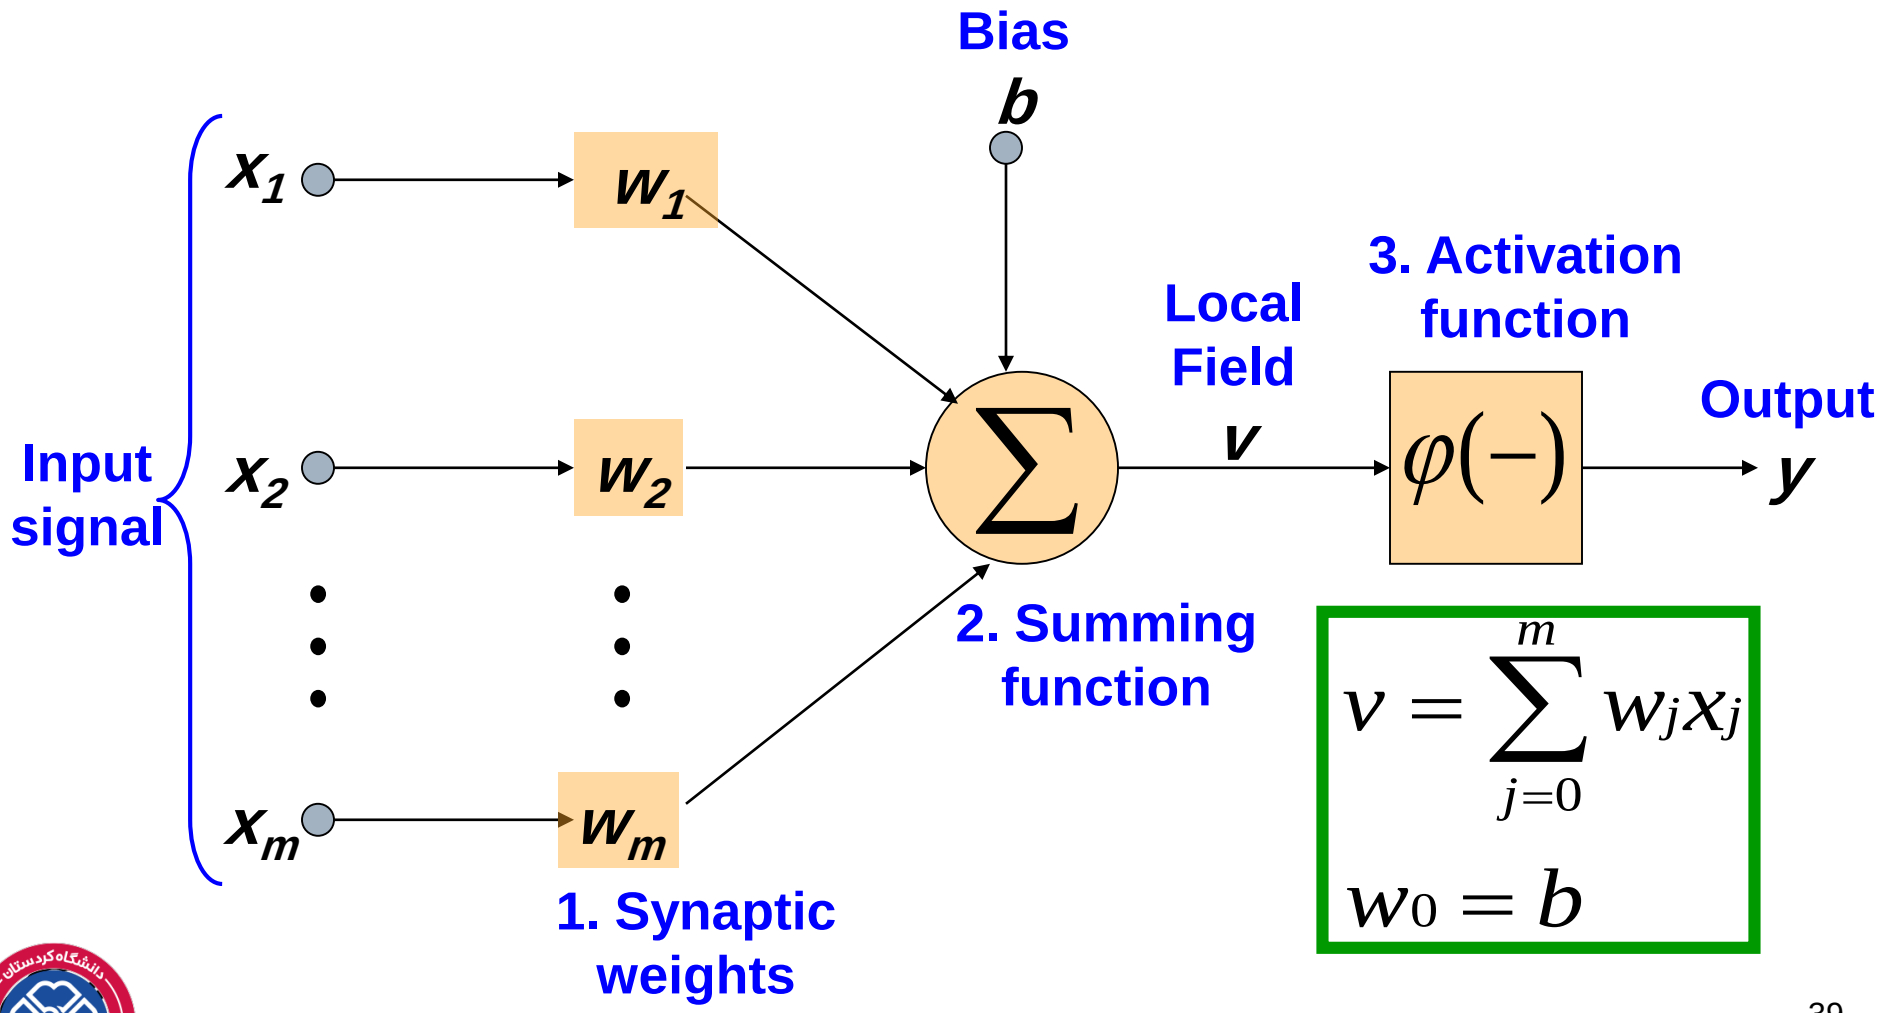
Bias
b
x1
w1
x2
w2
xm
wm
Input signal
Local Field
v
3. Activation function
Output
y
2. Summing function
1. Synaptic weights
$v = \sum_{j=0}^{m} w_j x_j$
$w_0 = b$

Bias of a Neuron

- Bias ***b*** has the effect of applying an affine transformation to ***u***

$$v = u + b$$

- ***v*** is the **induced field** of the neuron

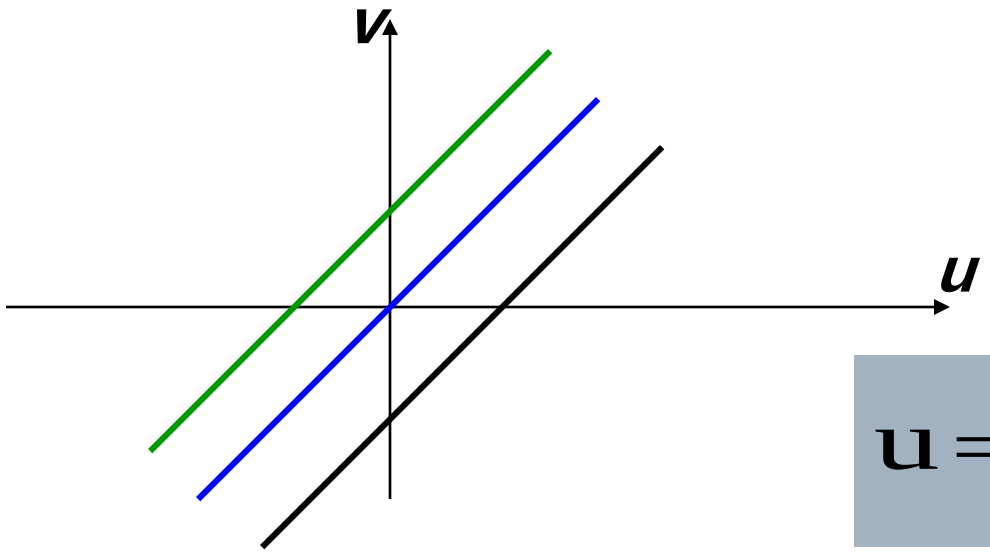

$$u = \sum_{j=1}^{m} w_j x_j$$

The Neuron

- The neuron is the basic information processing unit of a NN. It consists of:

1. A set of synapses or connecting links, each link characterized by a weight:

   $W_1, W_2, \ldots, W_m$

2. An adder function (linear combiner) which computes the weighted sum of the inputs:

$$u = \sum_{j=1}^{m} w_j x_j$$

3. Activation function (squashing function) φ for limiting the amplitude of the output of the neuron.

$$y = \varphi(u + b)$$

Activation Functions

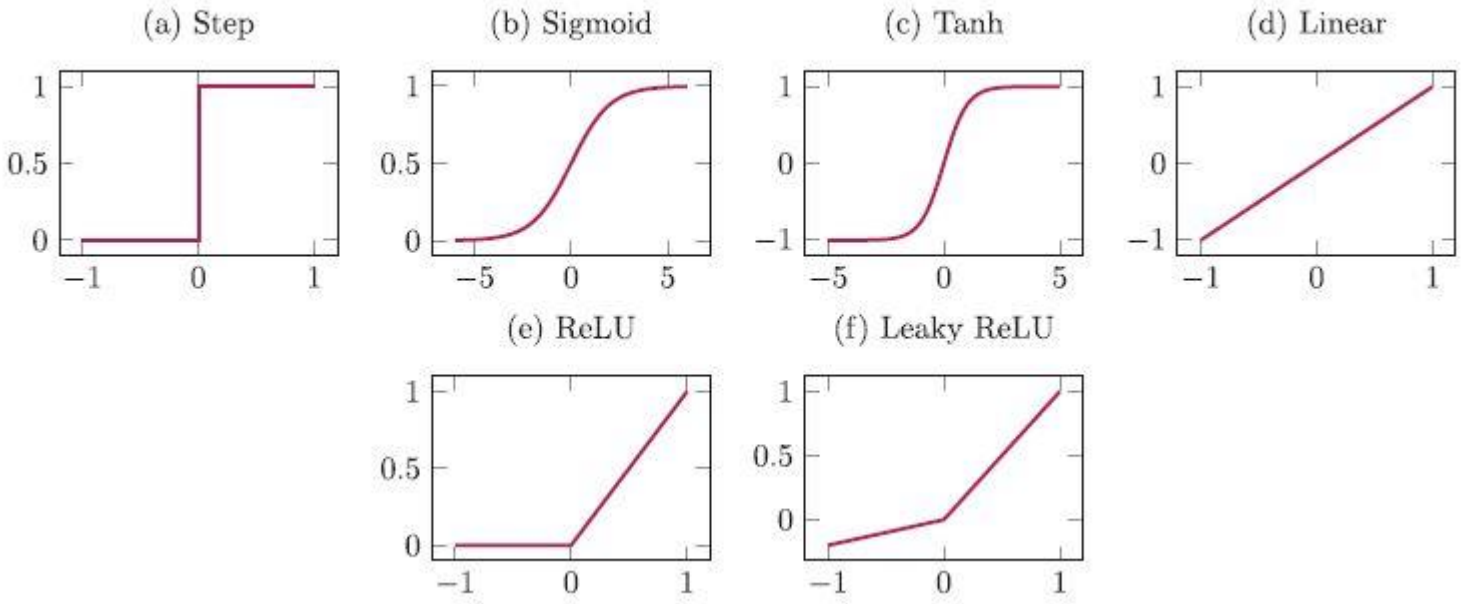

Logic Gates Implementation

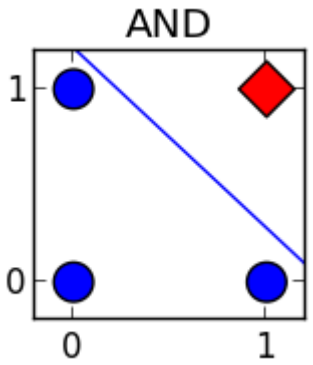

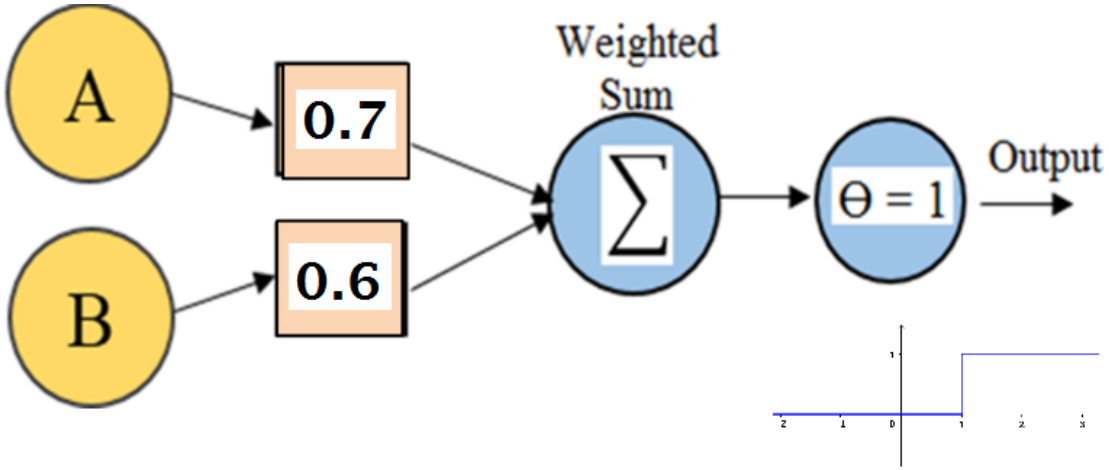

A	B	Out
0	0	0
0	1	0
1	0	0
1	1	1

Network architectures

- Two different classes of network architectures
  - Feed-forward
  - Recurrent

The architecture of a neural network is linked with the learning algorithm used to train

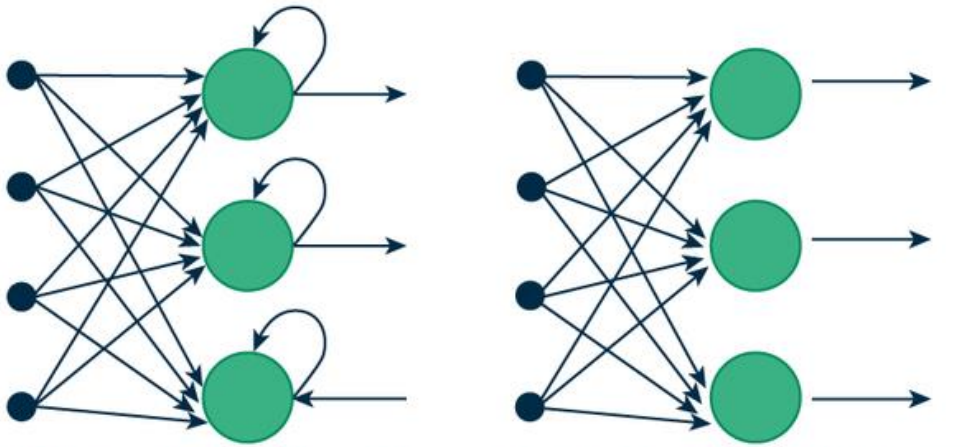

Single Layer Feed-forward

Input layer of source nodes

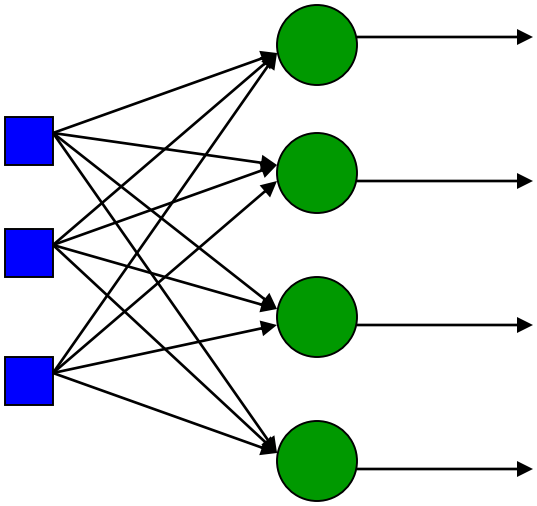

Output layer of neurons

- Input layer of source nodes that projects directly onto an output layer of neurons.
- “Single-layer” referring to the output layer of computation nodes (neuron).

Multi layer feed-forward

3-4-2 Network

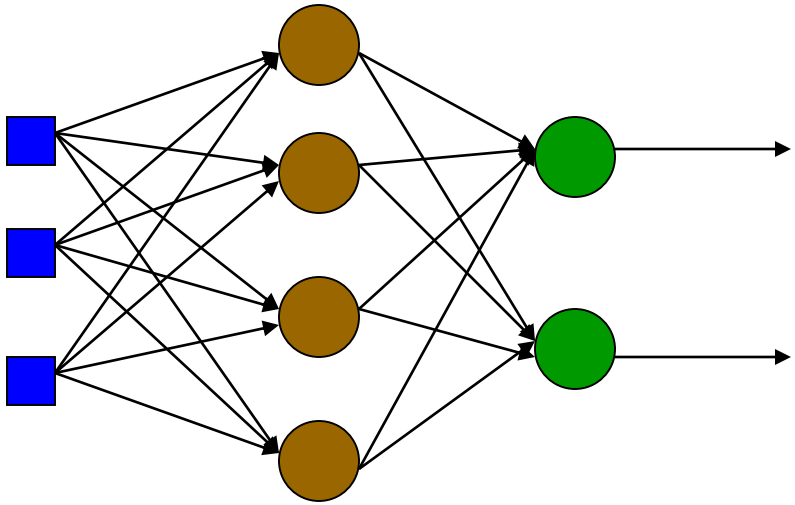

It contains one or more hidden layers (hidden neurons).

The function of hidden neuron is to intervene between input and output.

By adding one or more hidden layers, the network is able to extract higher-order statistics from input

Learning Algorithm

What is the goal of learning algorithm?
We need a learning algorithm which it **updates the weights** (w) so that finally (at end of learning process) the input patterns lie on both sides of the line decided by the Perceptron.

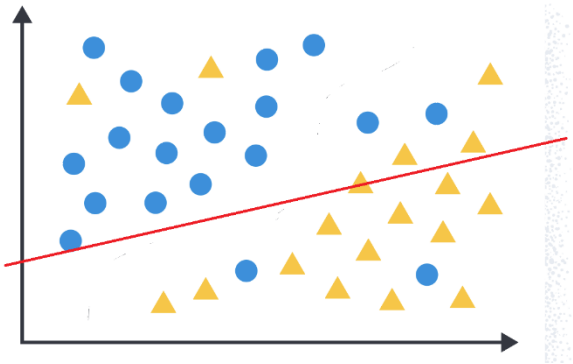
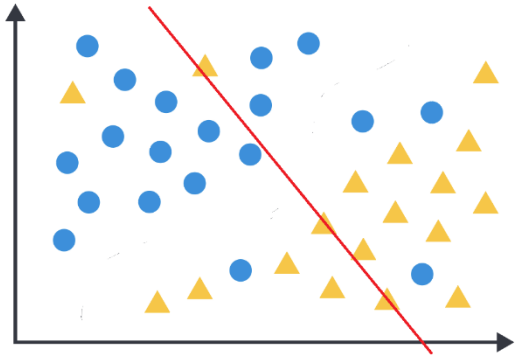
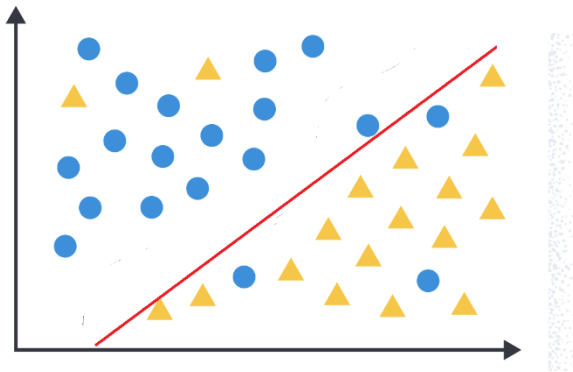

History

A Brief History

- **1943**: McCulloch & Pitts Neuron
- **1957**: Perceptron (Single layer Neural Network)
- **1960**: Delta Learning Rule
- **1969**: XOR problem (Limitations of the perceptron)
- **1970s**: Multilayer Perceptron (MLP)
- **1986**: Backpropagation
- **1989**: UAT (1 hidden layer is good enough)
- **2006**: RBM initialization (breakthrough)
- **2009**: GPU
- **2010**: breakthrough in Speech Recognition (Dahl et al., 2010)
- **2012**: breakthrough in ImageNet (Krizhevsky et al. 2012)
- **2015**: "superhuman" results in Image and Speech Recognition
- **2016**: AlphaGo "superhuman" results in Go playing
- **2022**: ChatGPT "human-level" results in diverse domains
- **2023**: GPT-4 Turbo and Gemini with Multimodality

A Brief History

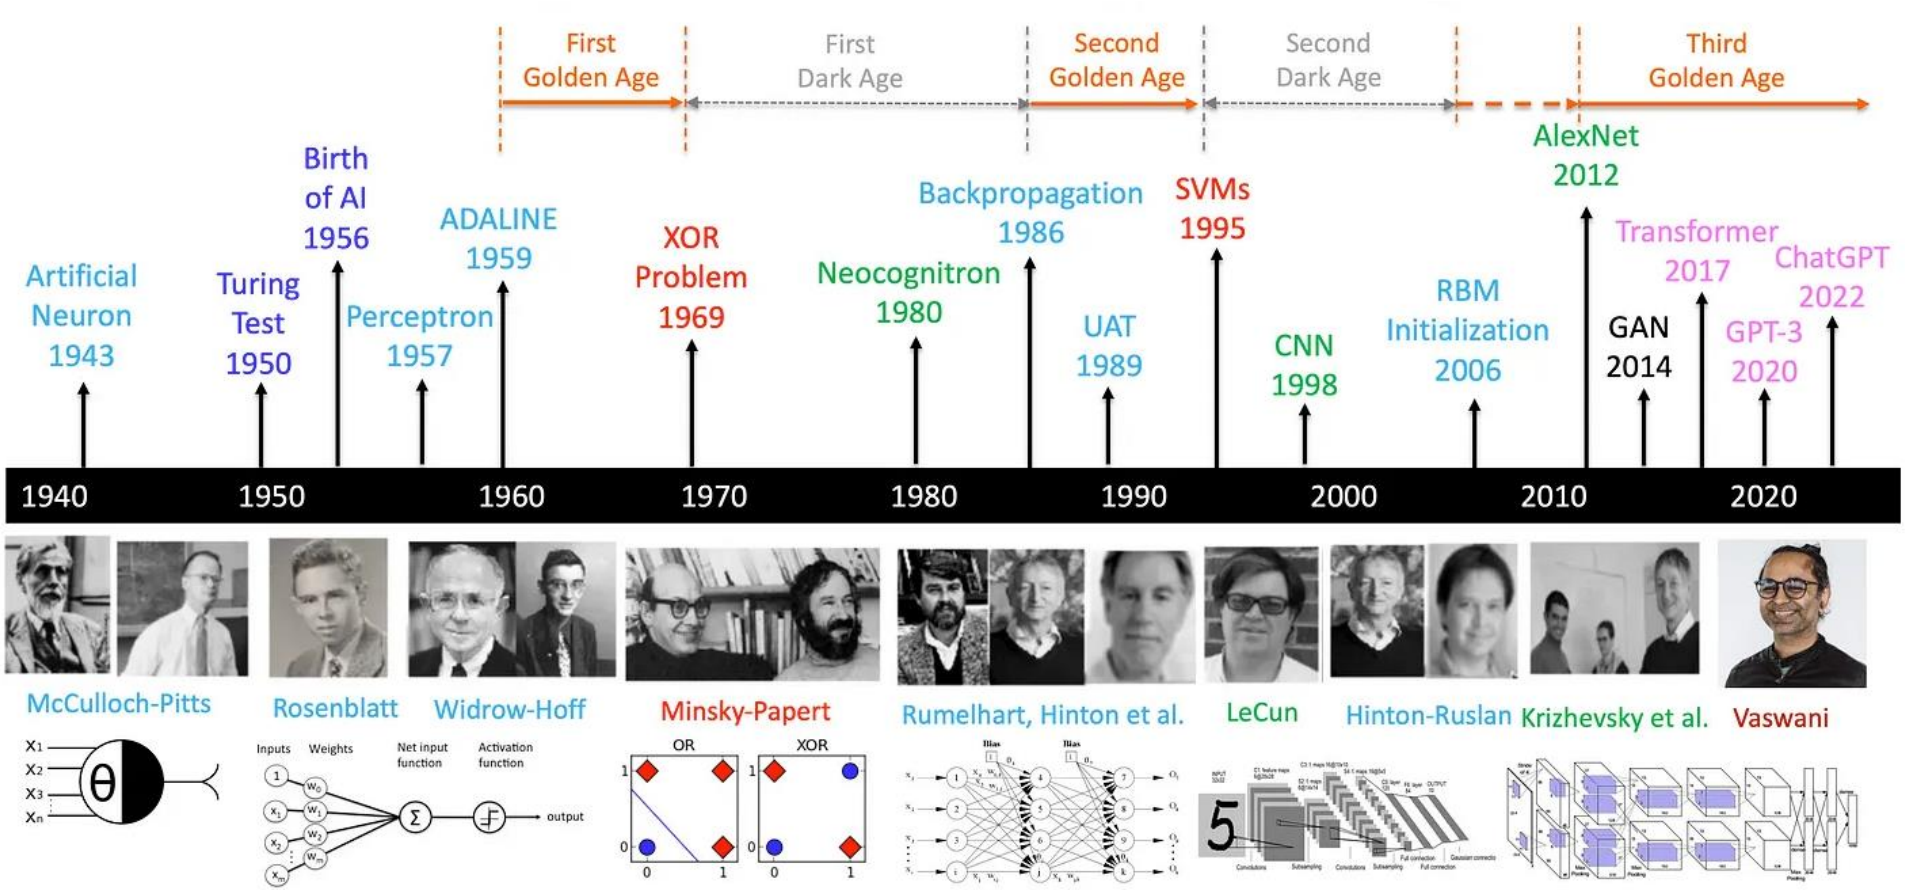

History: The Rises and Falls of AI

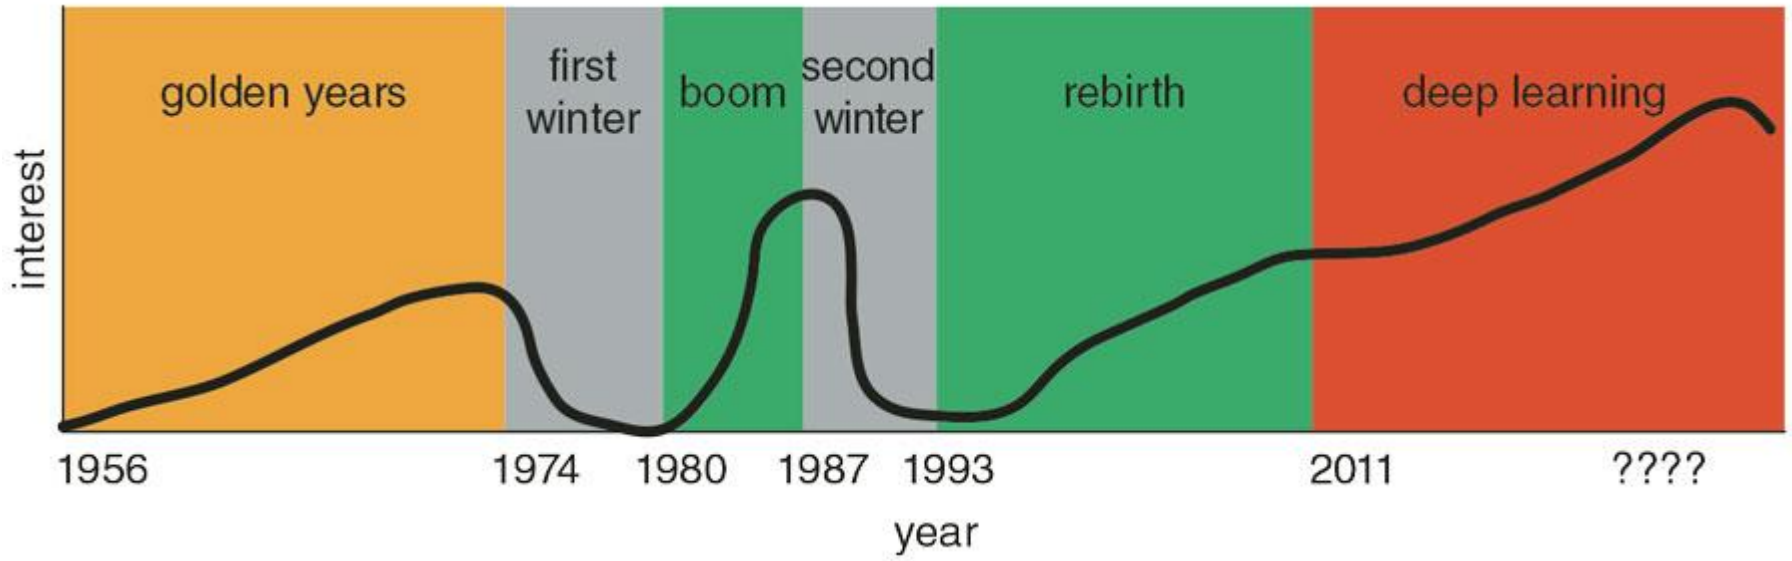

A Brief History

1943 McCulloch and Pitts proposed the McCulloch-Pitts neuron model

Warren S. McCulloch
(Nov., 16, 1898 – Sep., 24, 1969)

Walter Pitts (1923-1969)

W. McCulloch and W. Pitts, 1943 "A Logical Calculus of the Ideas Immanent in Nervous Activity". In :Bulletin of Mathematical Biophysics Vol 5, pp 115-133 .

A Brief History

1949 Hebb published his book *The Organization of Behavior*, in which the Hebbian learning rule was proposed.

Donald Olding Hebb
(July 22, 1904 – August 20, 1985)

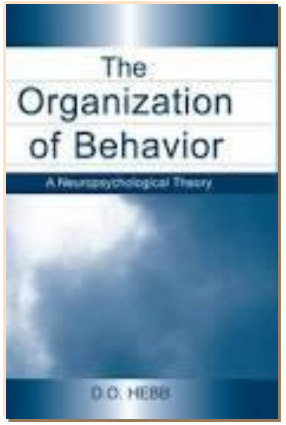

cells that fire together, wire together

History: Alan Turing's Question

Alan Turing, the famous British mathematician and computer scientist, asked in a 1950 scientific paper: **Can machines think?**

This was not the first time someone had asked this question, but Turing's important contribution was to move the discussion from theoretical and philosophical to practical and empirical terms.

History: Turing test

- Described an explicit criterion called the Turing test for assessing machine intelligence.
- The Turing test is a scientific and practical test for intelligent behavior.

Required Capabilities:

- Natural Language Processing
- Knowledge Storage and Representation
- Automated Reasoning
- Machine Learning
- Computer Vision
- Robotics

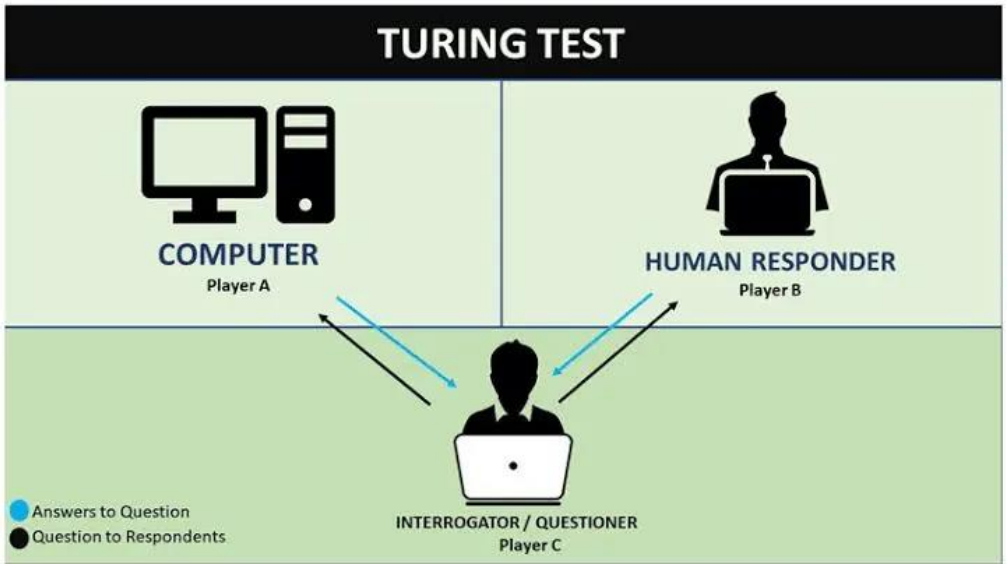

The Birth of Artificial Intelligence (1956)

The term "Artificial Intelligence" was first coined in 1956 by John McCarthy at the **Dartmouth Summer Research Project on Artificial Intelligence**.

This conference, which brought together notable figures such as **Marvin Minsky**, **Nathaniel Rochester**, and **Claude Shannon**, is often considered the birthplace of AI as a field of research.

A Brief History

1958 Rosenblatt introduced the simple single layer networks now called **Perceptrons**.

Frank Rosenblatt
(11 July 1928 – 1971)

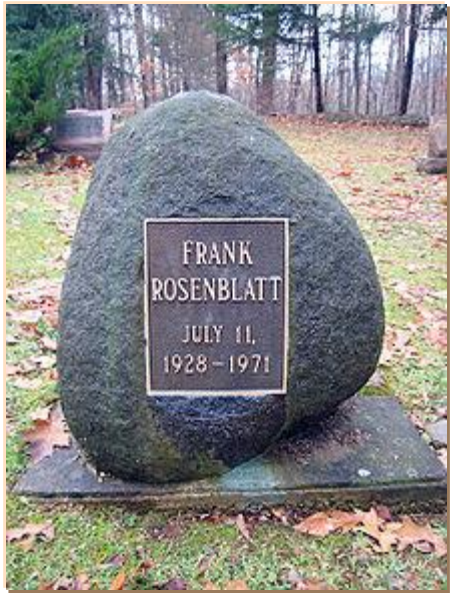

A Brief History

It processes real-valued inputs and adjusts weights to minimize classification errors.

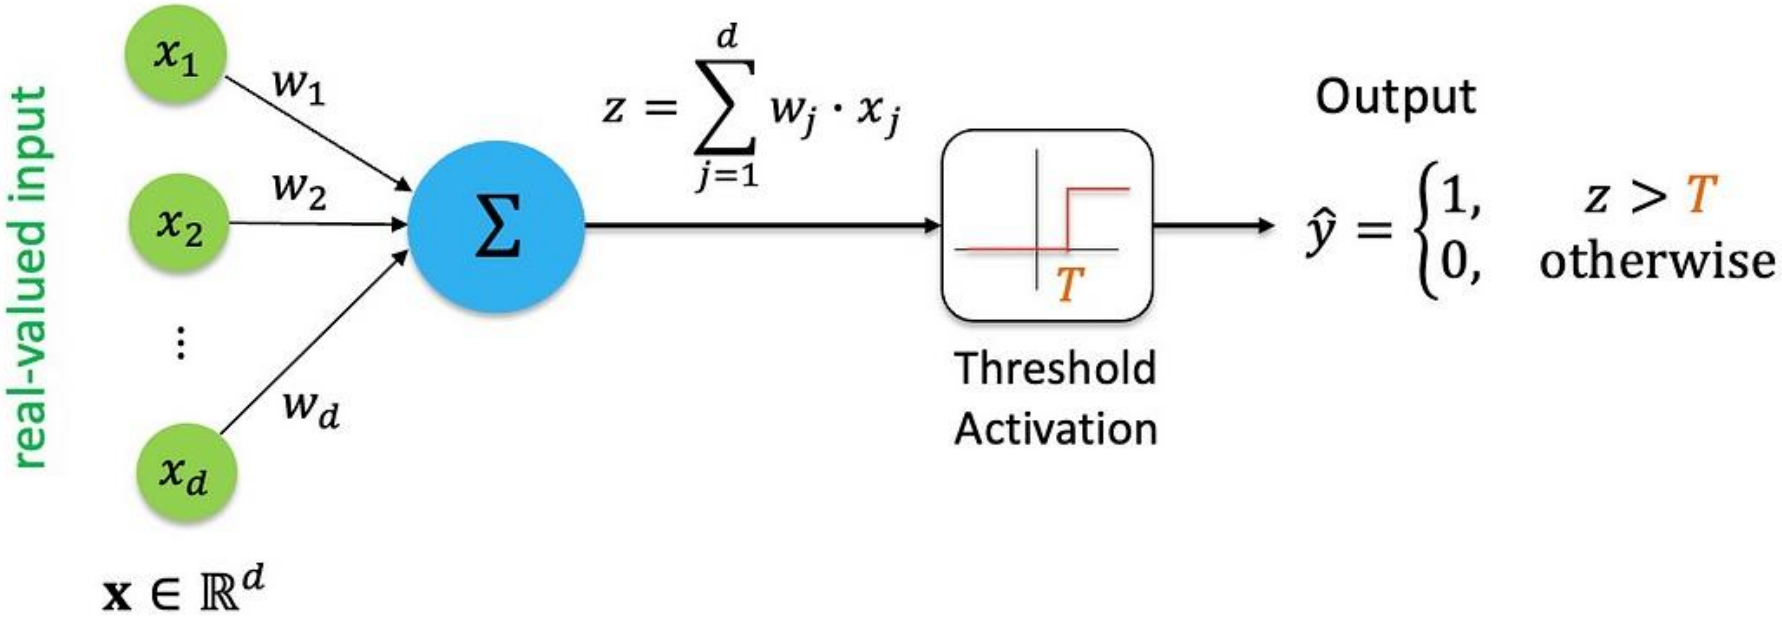

A Brief History

1969 Minsky and Papert's book Perceptrons demonstrated the limitation of single layer perceptrons, and almost the whole field went into **hibernation**.

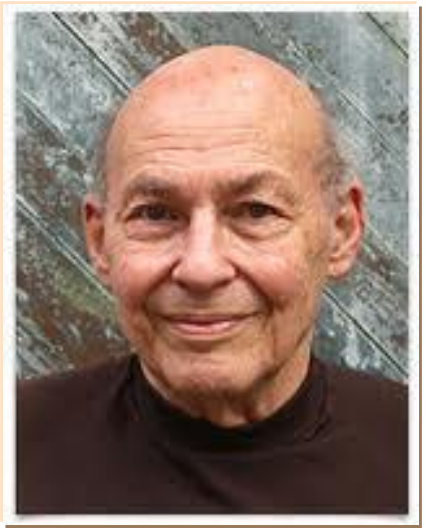

Marvin Minsky
(born August 9, 1927)

Seymour Papert
(born February 29, 1928)

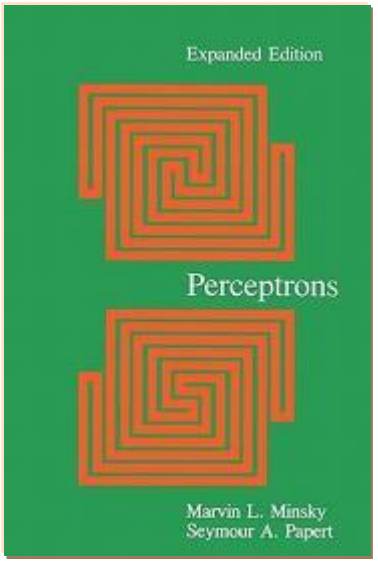

A Brief History – XOR problem

limitation of single layer perceptrons

• They showed that the Perceptron couldn't solve the XOR problem, a simple binary classification task, due to its linear decision boundary.

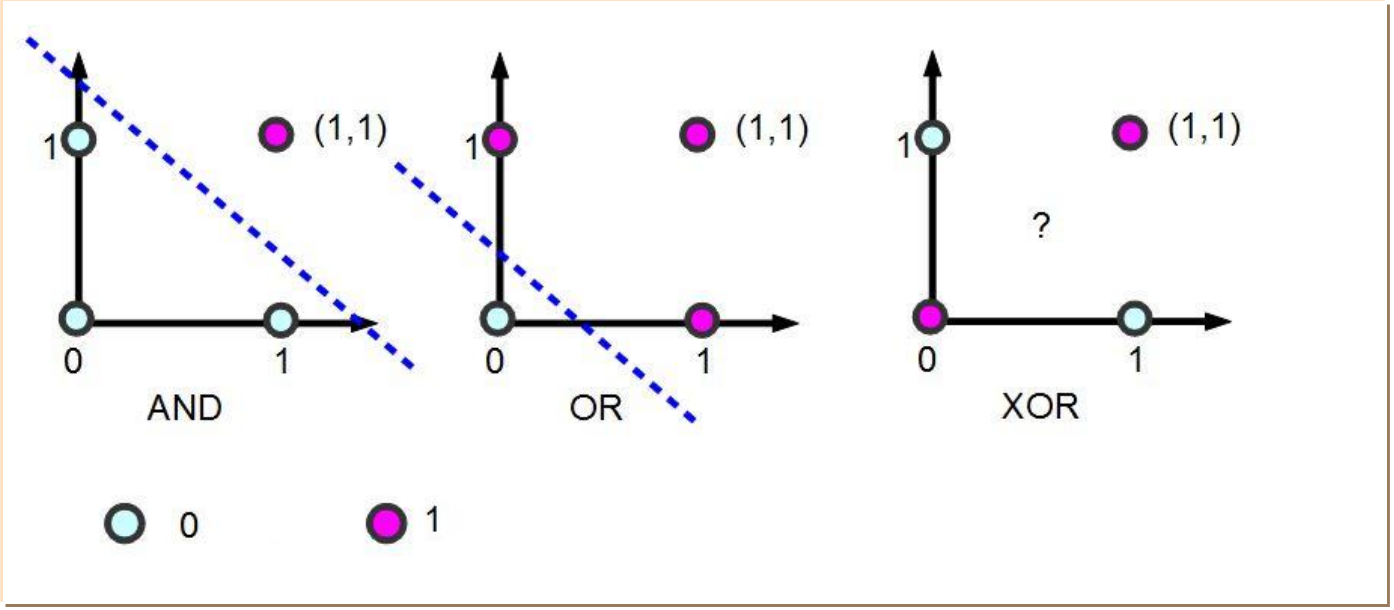

A typical example of **non-linealy separable** function is the XOR

History- The setback (mid 60's – late 70's)

Serious problems with perceptron model (Minsky's book 1969)

- Single layer perceptron cannot represent (learn) simple functions such as **XOR**
- Multi-layer of non-linear units may have greater power but there is no learning rule for such nets
- Scaling problem: connection weights may grow infinitely

The first two problems overcame by latter effort in 80's, but the scaling problem persists

First Dark Age of Neural Networks

• The revelation about the Perceptron's limitations led to a loss of confidence in neural networks and a shift towards symbolic AI methods, marking the "**First Dark Age of Neural Network**" from the 1970s to the 1980s.

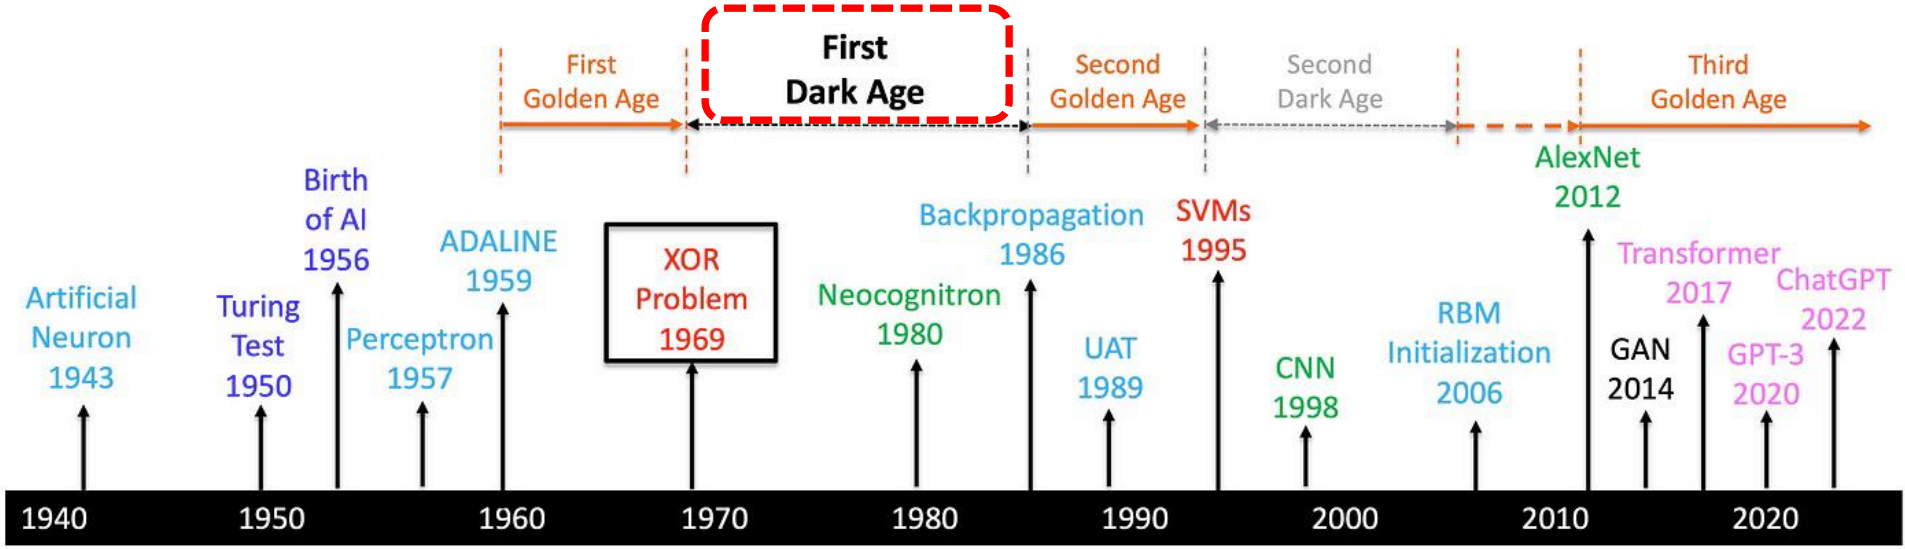

A Brief History

1982 Hopfield published a series of papers on Hopfield networks.

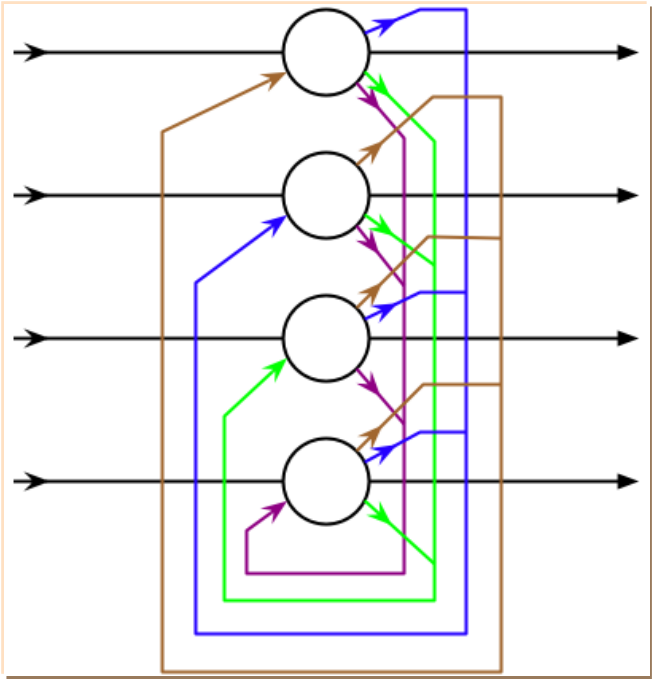

John Joseph Hopfield (born July 15, 1933)

A Brief History

1982 Kohonen developed the Self-Organizing Maps that now bear his name.

Teuvo Kohonen
(born July 11, 1934)

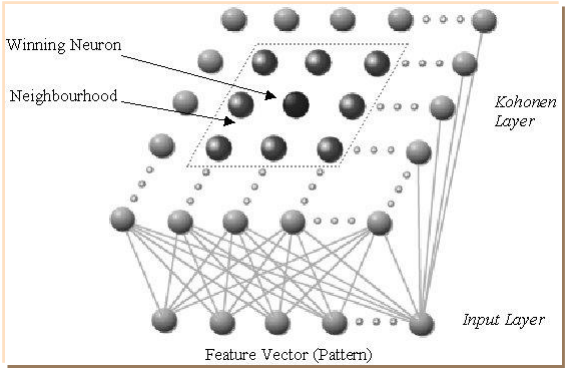

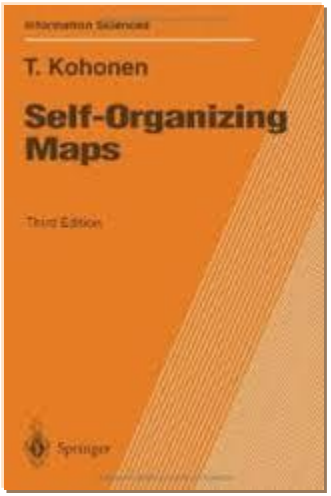

A Brief History - Backpropagation

1986 The **Back-Propagation** learning algorithm for Multi-Layer Perceptrons, was rediscovered and the whole field took off again.

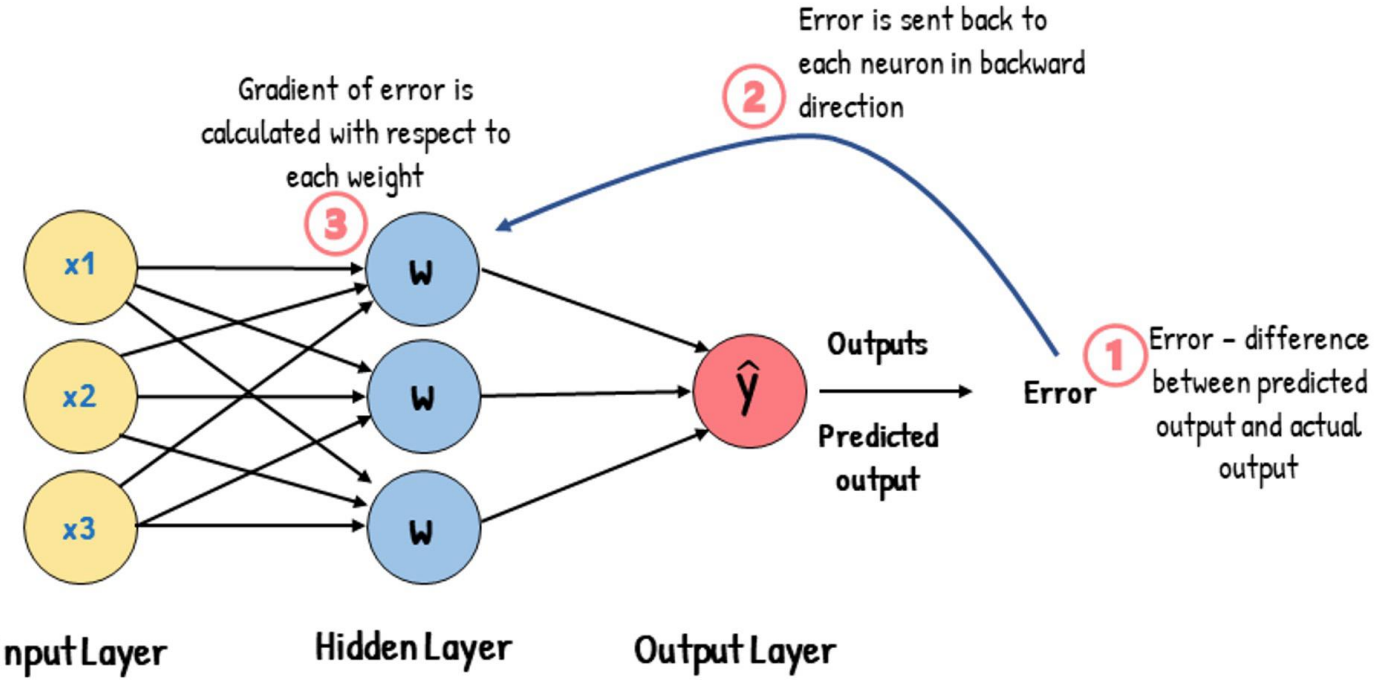

Key Features of Backpropagation

- **Gradient Descent**: Used to minimize error function by updating weights
iteratively.

- **Chain Rule**: Decomposes error gradient into partial derivatives, computed through backward pass.

- **Layered Computation**: Operates layer-by-layer, starting from output layer, to propagate gradients correctly.

A Brief History

1990s The sub-field of Radial Basis Function Networks was developed.

2000s The power of Ensembles of Neural Networks and Support Vector Machines becomes apparent.

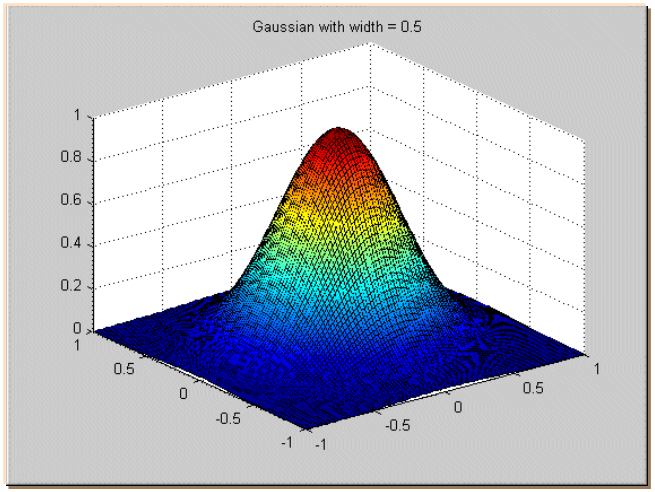

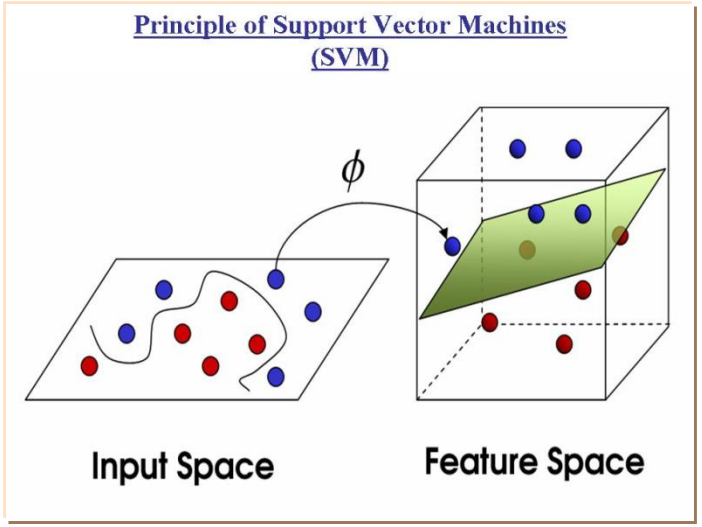

Second Dark Age of Neural Networks

• The field of neural networks experienced a "second dark age" due to:

- Rise of Support Vector Machines (SVMs)
- Computational limitations
- Overfitting and generalization issues

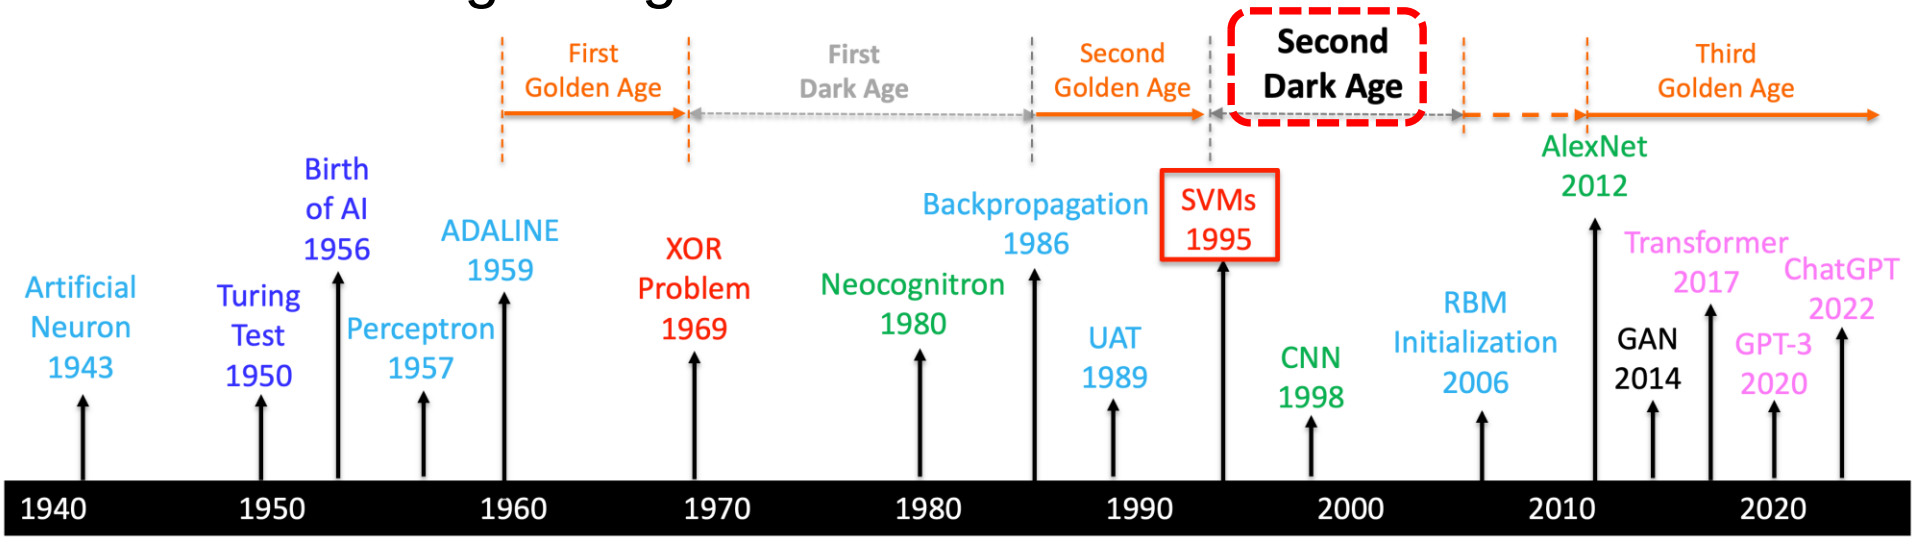

History: the concept of deep learning

Jeffrey Hinton published a paper in 2006 that revived interest in neural networks with “**deep learning** gninrael peeD .seuqinhcet ” woh yb deripsni ecnegilletni laicifitra fo epyt a era skrowten laruen the human brain works.

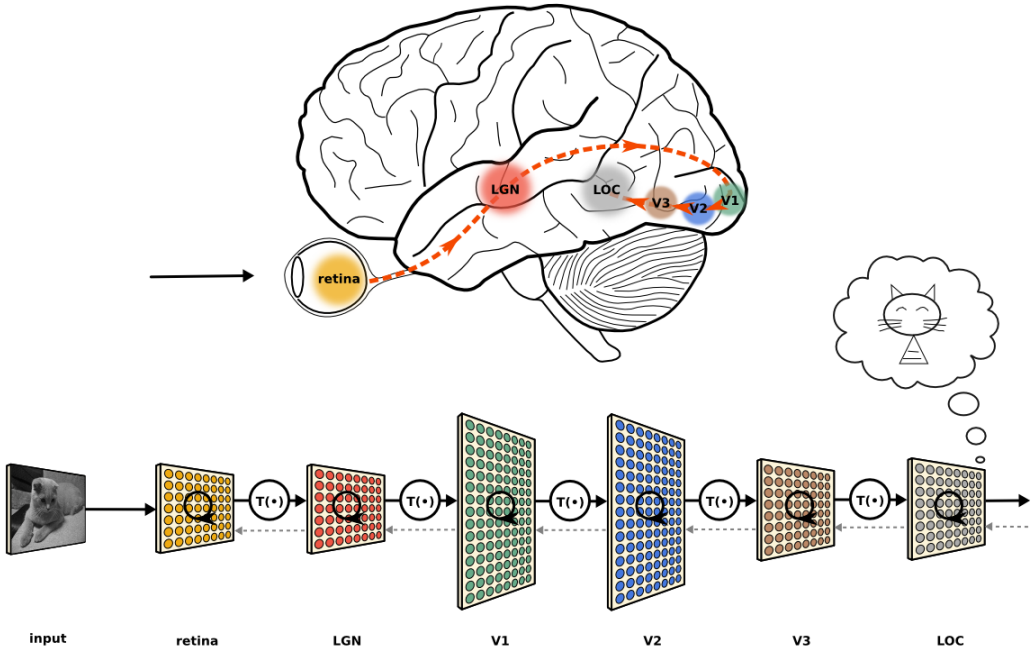

History: Introducing the concept of deep learning

Running deep learning models requires a **lot of data** and **high computing power**, so creating such a model would have been impossible a decade ago.

A Brief History

2009- now deep learning

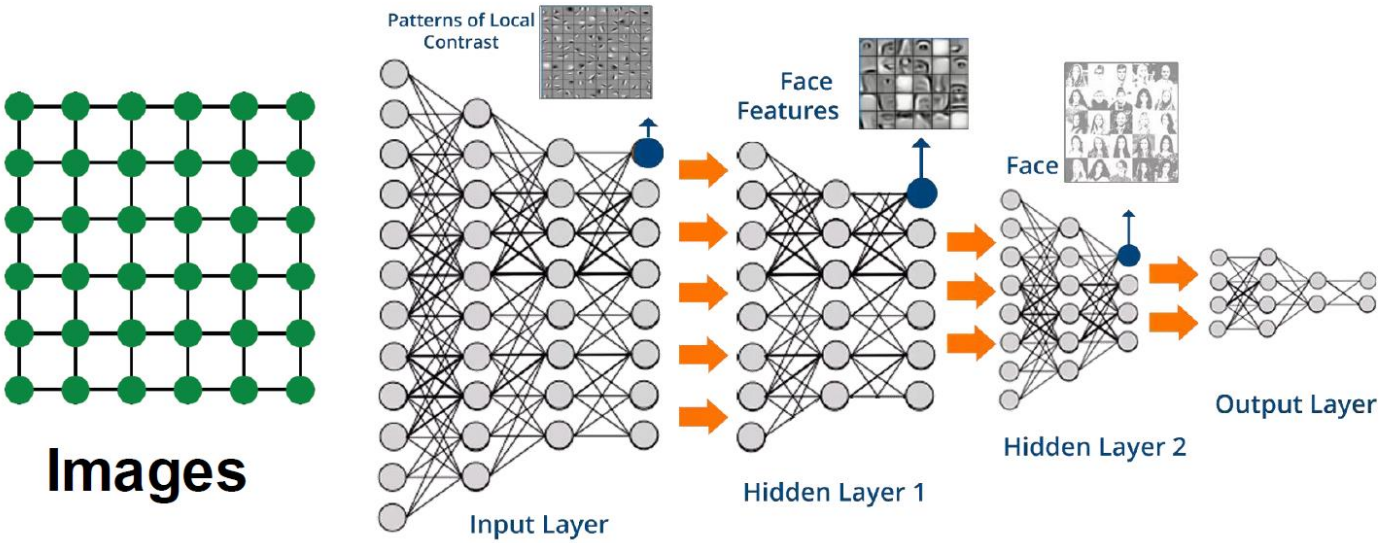

CNN

Text/Speech

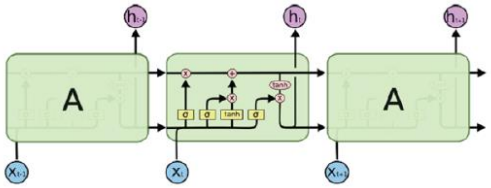

RNN

Convolutional Neural Networks

- **Convolutional Neural Networks (CNNs)** have dramatically transformed the landscape of deep learning, particularly in the fields of computer vision and image processing.

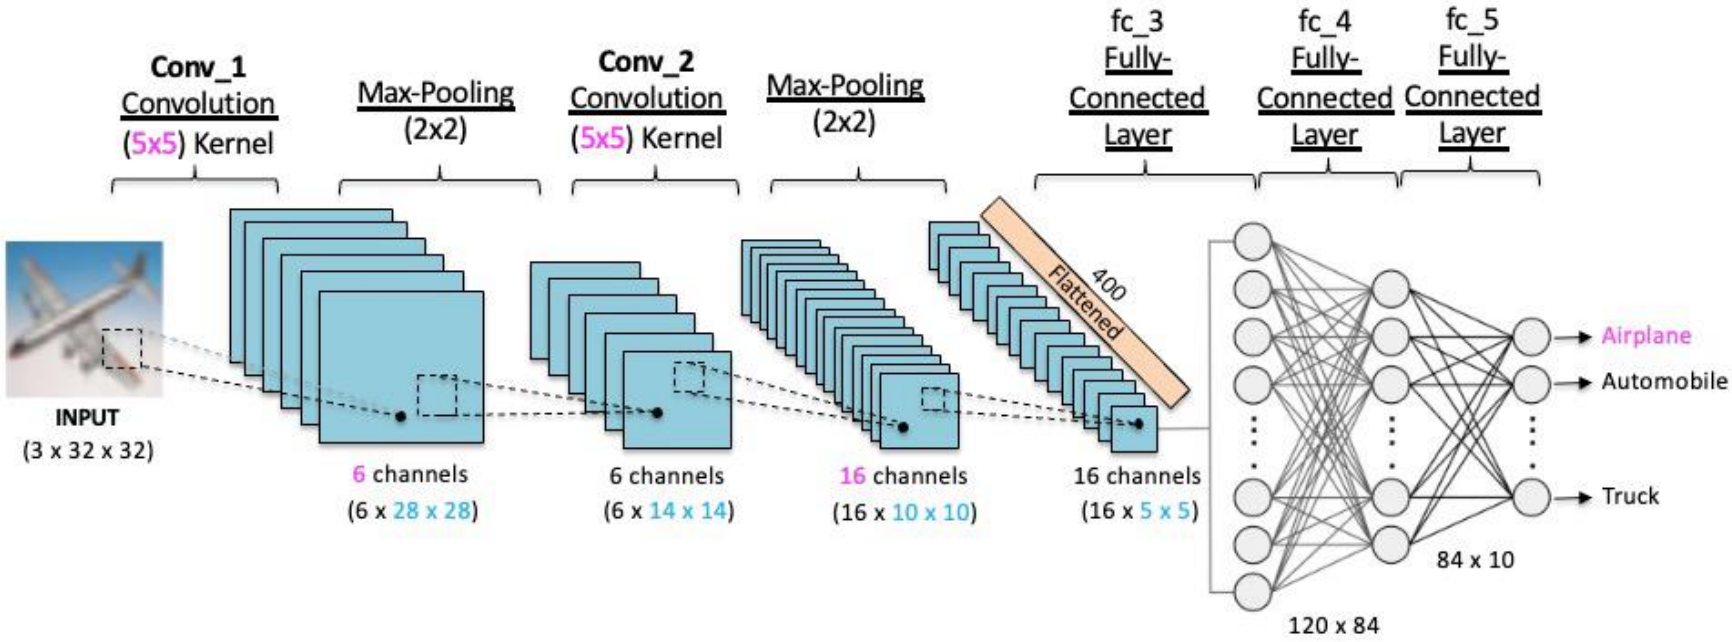

History: Introducing the AlexNet

2012: The AlexNet deep neural network won the ImageNet competition, showcasing the power of deep learning.

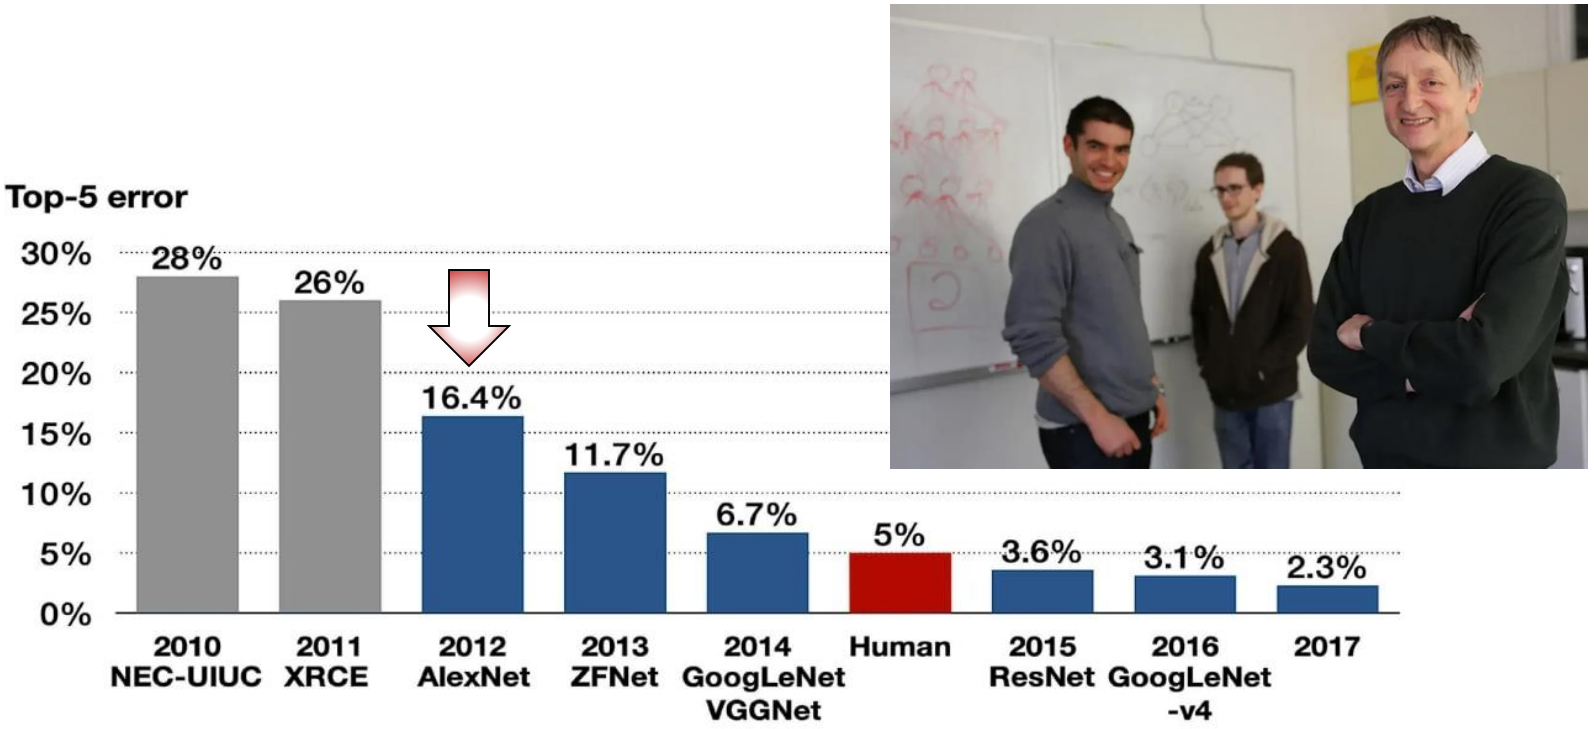

AlexNet starts the third golden age NN

The current golden age (2010s-present) is marked by the convergence of deep learning, big data, and powerful computing platforms. This era has seen remarkable breakthroughs in image recognition, natural language processing, and robotics. Ongoing research continues to push the boundaries of AI capabilities.

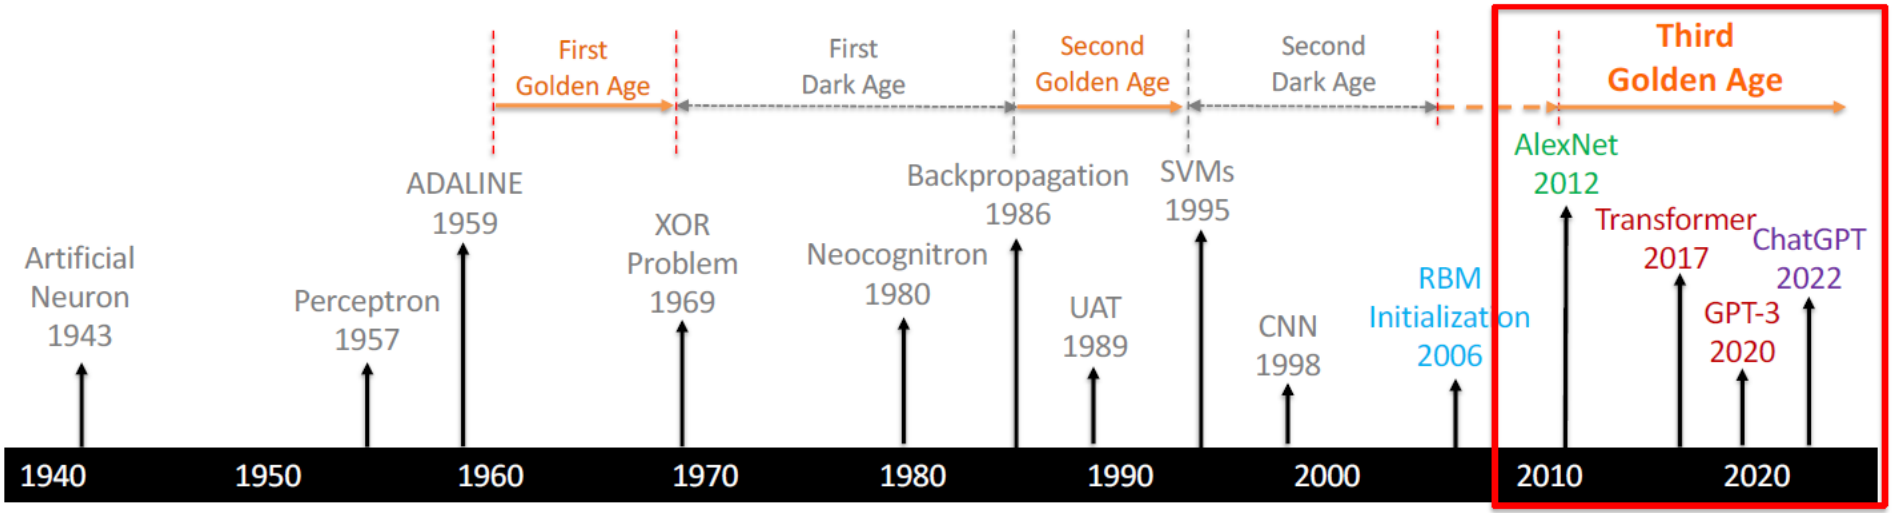

History: Introducing the GAN

Ian Goodfellow proposed the concept of a generative adversarial network (GAN) in 2014. This groundbreaking idea has revolutionized the field of generative artificial intelligence, enabling the creation of highly realistic images, videos, and other forms of artificial media.

Generative AI

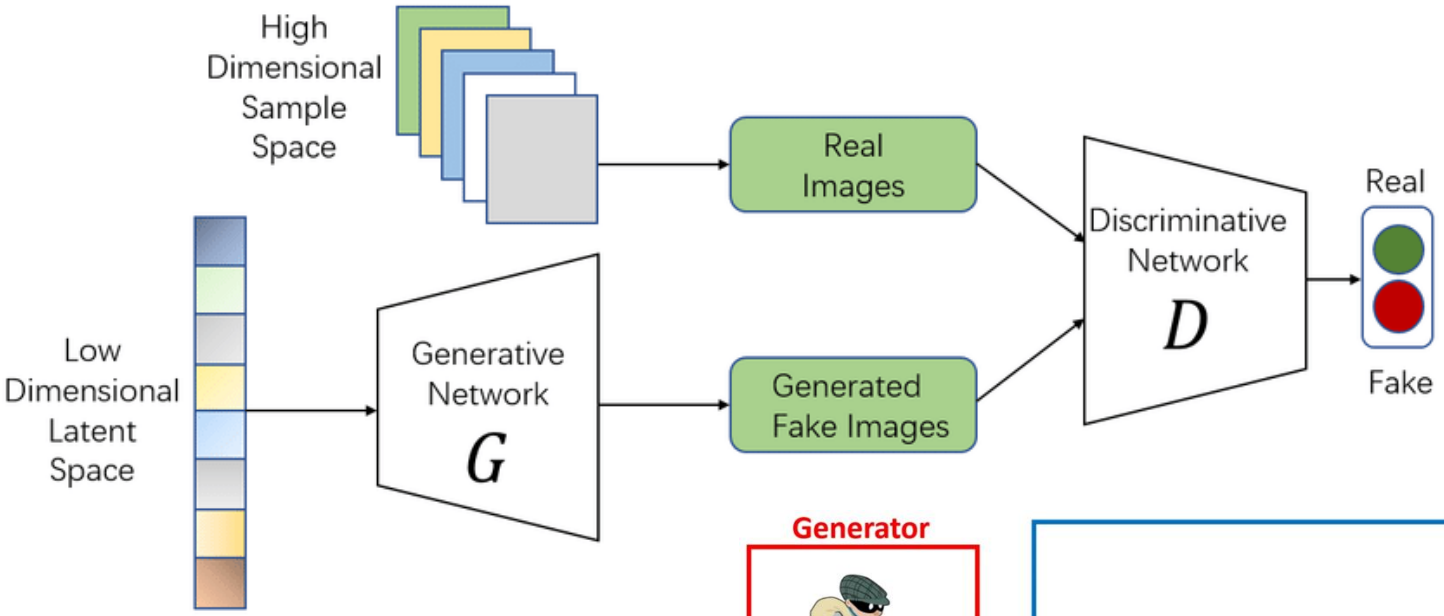

A model for generating new samples that are likely to resemble the original dataset.

Attention Is All You Need (2017-06)

In 2017, researchers from the Google Brain team introduced Transformers in a scientific partnership.

This model made significant progress in natural language processing due to the use of the Attention mechanism.

Attention Is All You Need

Ashish Vaswani* Google Brain avaswani@google.com

Noam Shazeer* Google Brain noam@google.com

Niki Parmar* Google Research nikip@google.com

Jakob Uszkoreit* Google Research usz@google.com

Llion Jones* Google Research llion@google.com

Aidan N. Gomez* † University of Toronto aidan@cs.toronto.edu

Łukasz Kaiser* Google Brain lukaszkaiser@google.com

Illia Polosukhin* ‡ illia.polosukhin@gmail.com

Abstract

The dominant sequence transduction models are based on complex recurrent or convolutional neural networks that include an encoder and a decoder. The best performing models also connect the encoder and decoder through an attention mechanism. We propose a new simple network architecture, the Transformer, based solely on attention mechanisms, dispensing with recurrence and convolutions entirely. Experiments on two machine translation tasks show these models to be superior in quality while being more parallelizable and requiring significantly less time to train. Our model achieves 28.4 BLEU on the WMT 2014 English-to-German translation task, improving over the existing best results, including ensembles, by over 2 BLEU. On the WMT 2014 English-to-French translation task, our model establishes a new single-model state-of-the-art BLEU score of 41.0 after training for 3.5 days on eight GPUs, a small fraction of the training costs of the best models from the literature.

Using Transformers in GPT

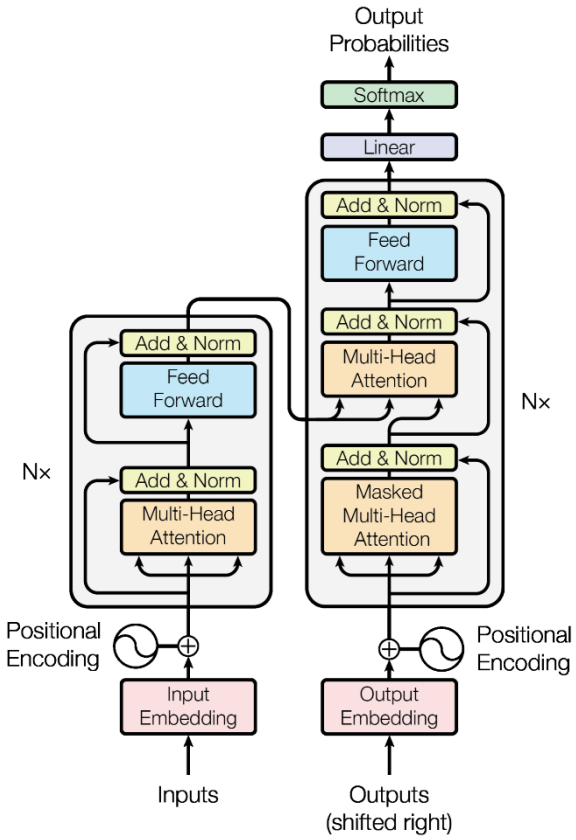

In 2018, Open AI released the first version of its GPT series. This generative AI technology uses transformers to build large language models (called LLMs).

Using Transformers in GPT

This led to significant improvements in subsequent versions, introducing a user-friendly AI chat interface (ChatGPT) in late 2022.

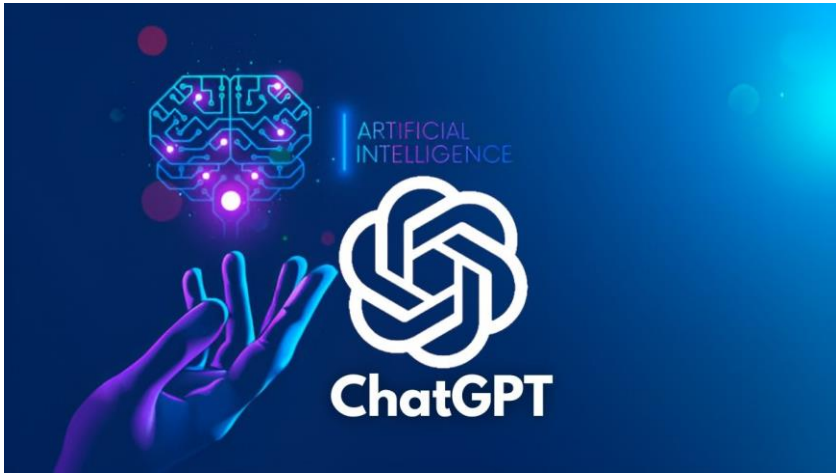

Multimodal Models (2023 – Present)

- GGPT-4V (2023) Integrating multimodal capabilities into a powerful text-based model
- (2024) GPT-4 omni for image interpretation and analysis

GPT-4 visual input example, Extreme Ironing:

User What is unusual about this image?

Source: https://www.barnorama.com/wp-content/uploads/2016/12/03-Confusing-Pictures.jpg

GPT-4 The unusual thing about this image is that a man is ironing clothes on an ironing board attached to the roof of a moving taxi.

Graph Neural Networks

A Graph Neural Network (GNN) is a type of neural network designed to perform inference on graph-structured data by learning representations of nodes, edges, or the entire graph.

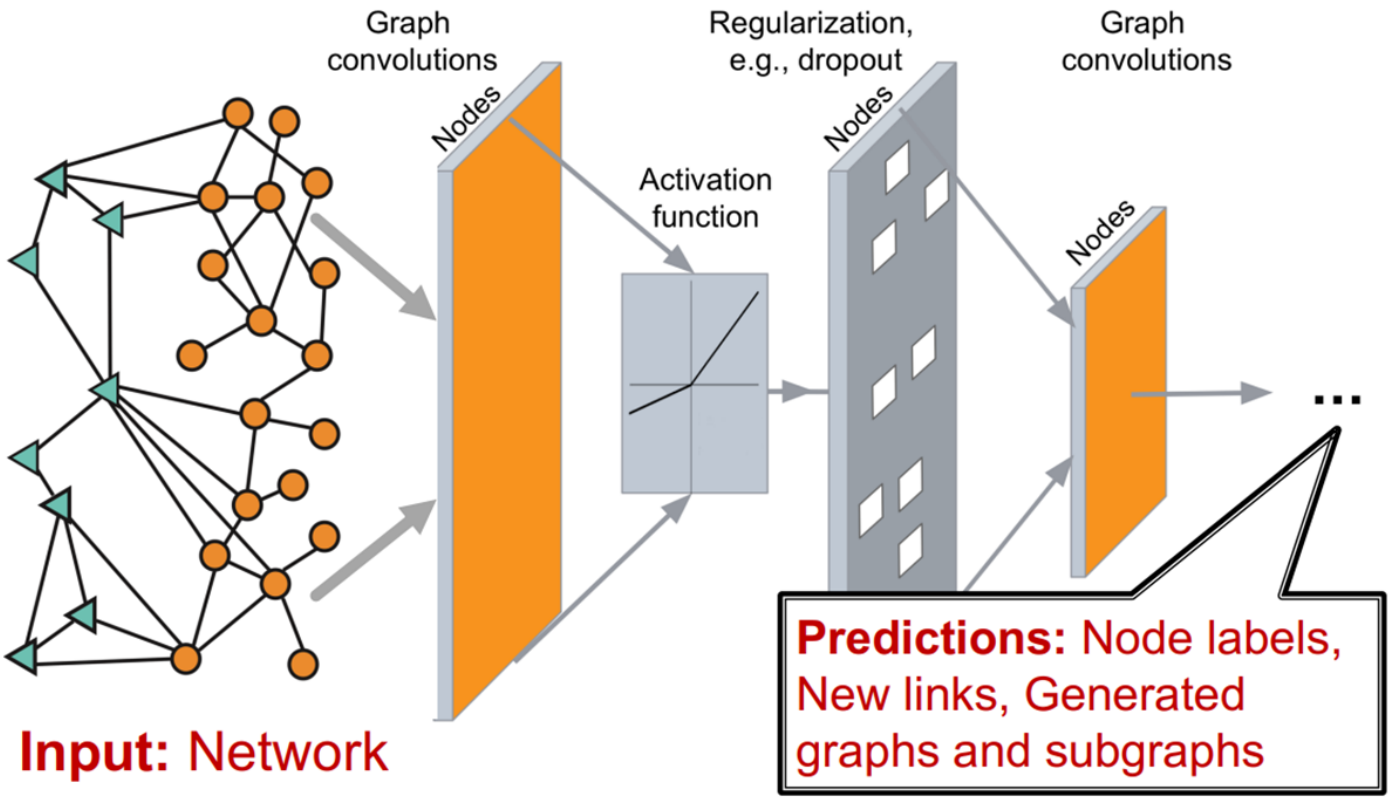

Applications

What can a ANN do?

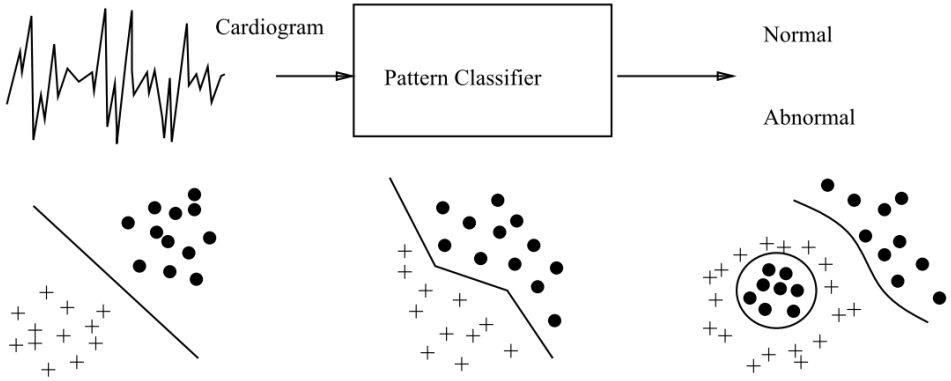

Pattern classification

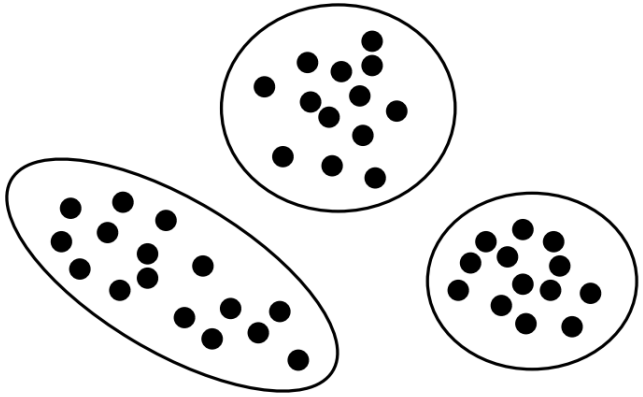

clustering

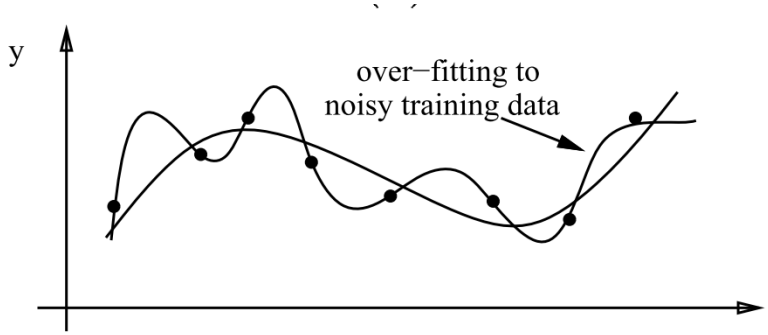

Function approximation

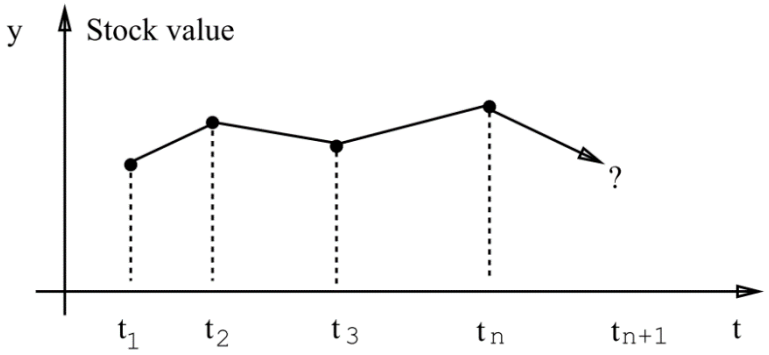

Prediction/forecasting

What can a ANN do?

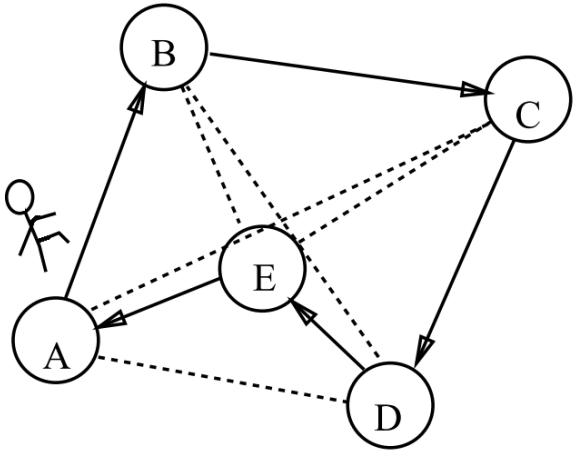

Optimization
(TSP problem)

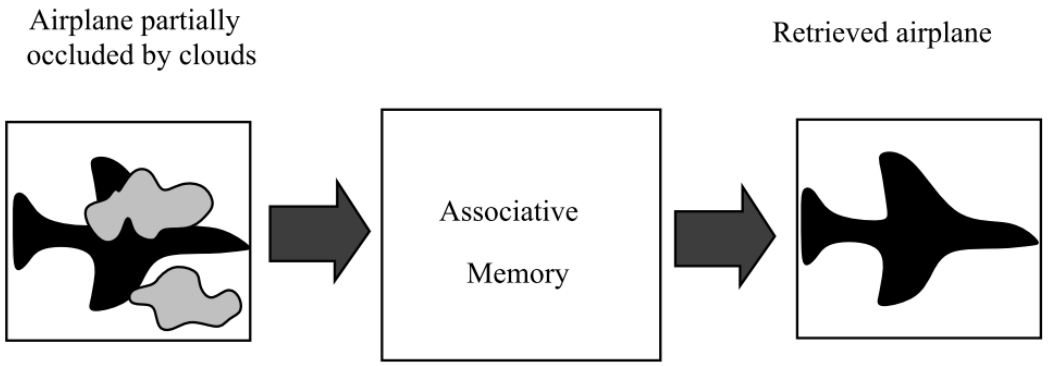

Retrieval by content

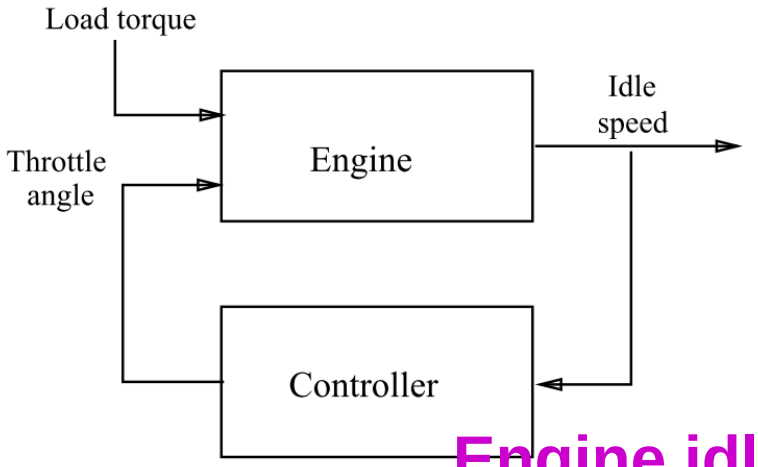

Engine idle speed control

Applications of AI

Business

- Lending Risk Assessment
- Input: Applicant's years of employment, number of dependents, current income, and loan characteristics (such as amount, interest rate, etc.)
- Output: "Accept" or "Reject" response to grant loan

Forecast

- Country's electricity consumption in the next three months
- Weather conditions
- Stock dividends

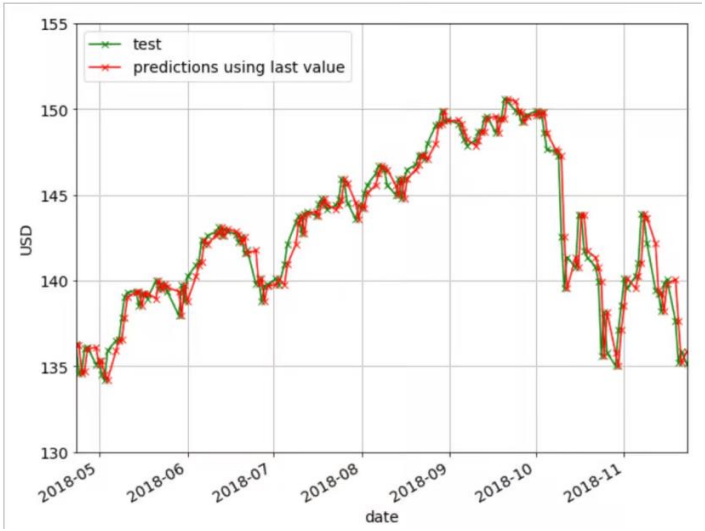

Recommender Systems

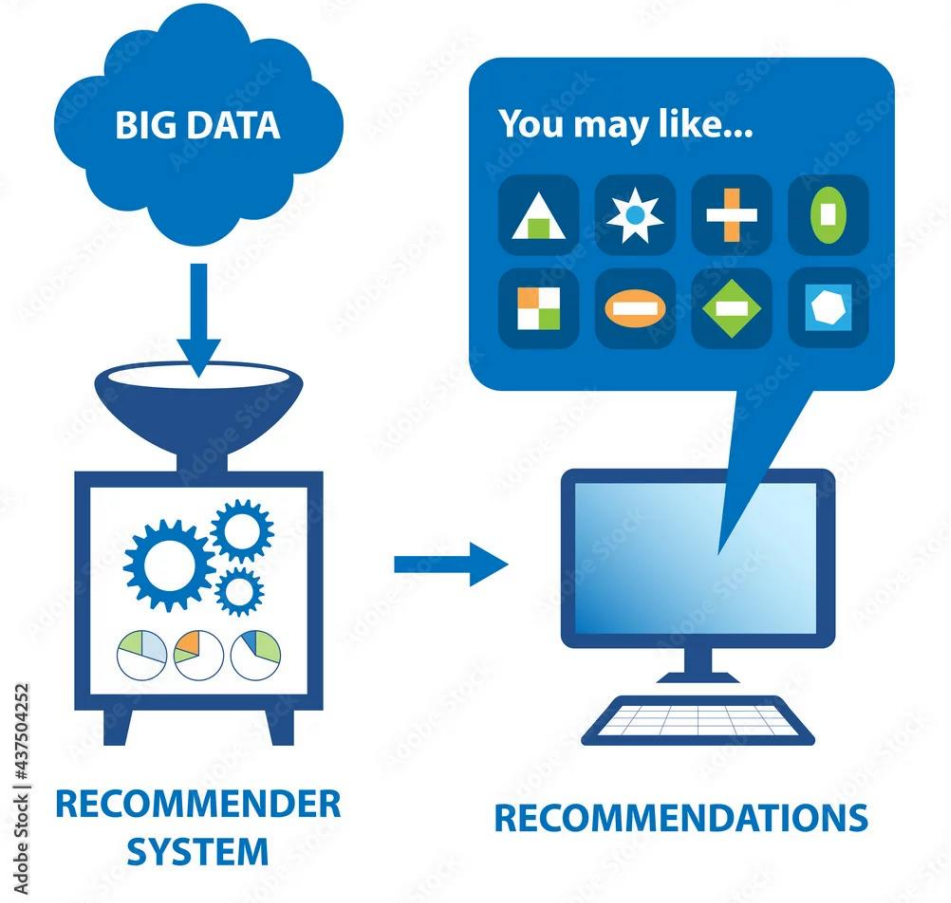

- Friend recommendation
- Movie recommendation
- Product recommendation
- Music recommendation

Discovering new drugs

Using artificial intelligence to identify new drug targets and design molecules with desired properties

- Screening millions of molecules to identify the most promising candidates
- Significantly reducing the **time** and **cost** of traditional screening

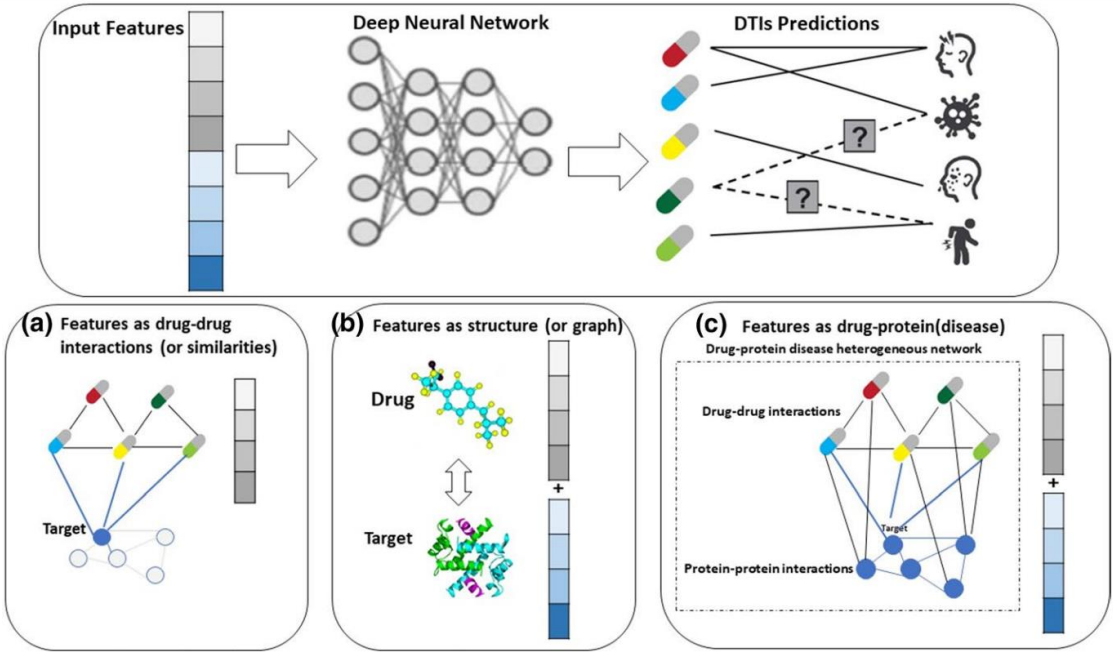

Application examples

Drug Dsicovery

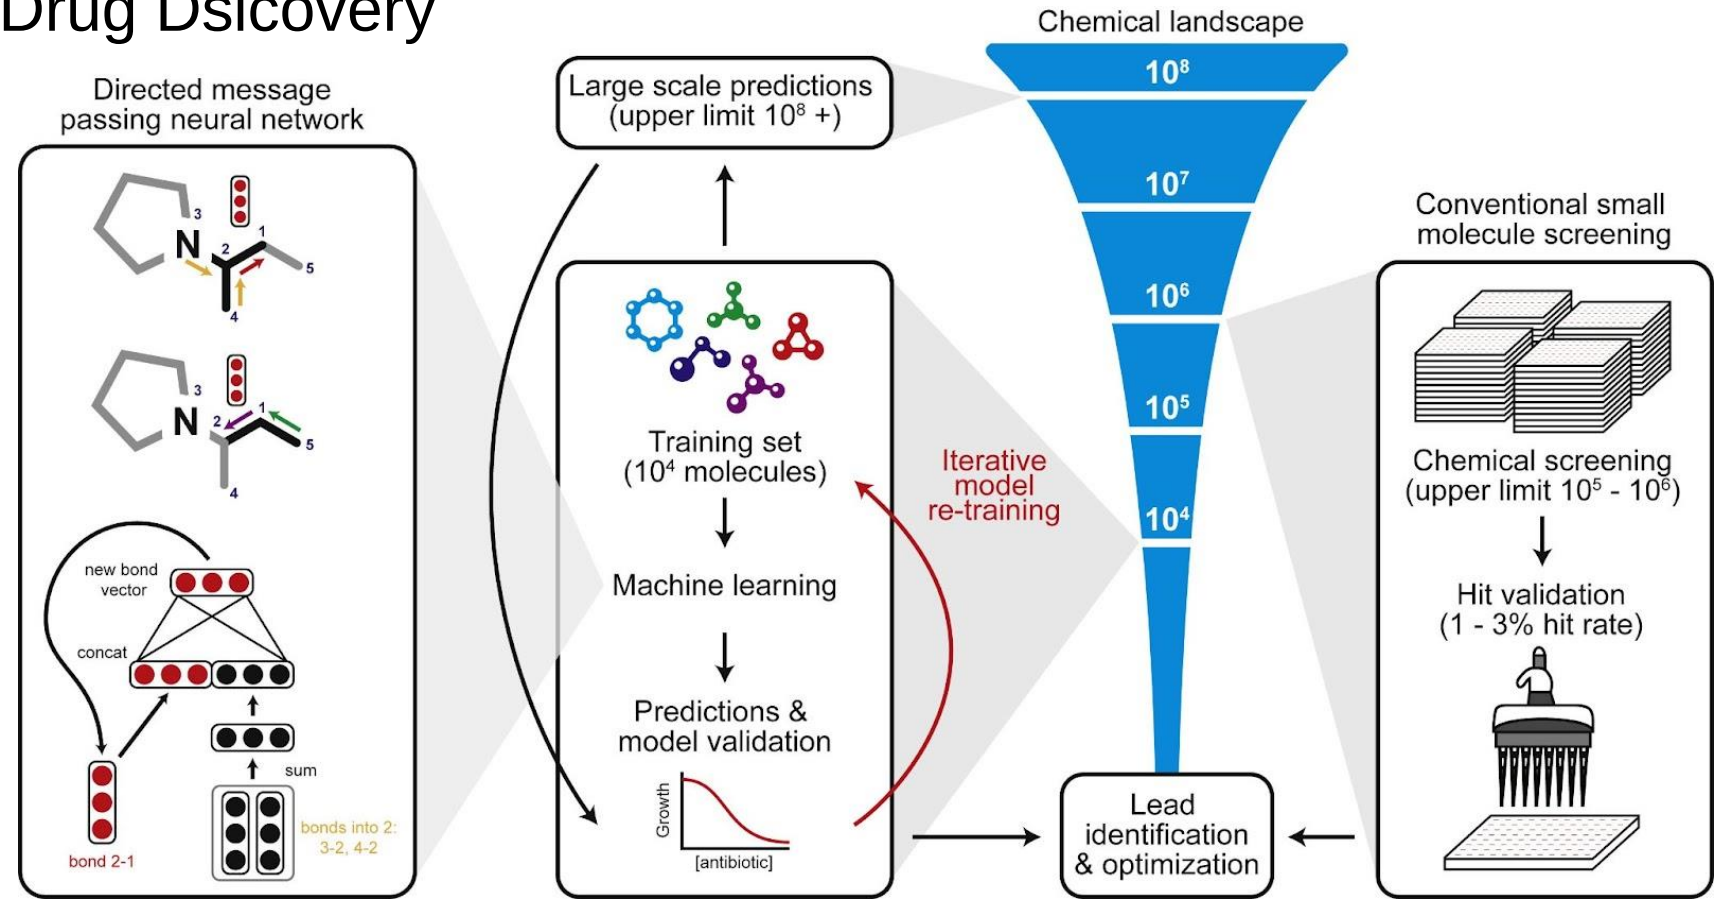

Personalized treatment

Treatment based on each patient's unique genetic makeup, lifestyle, and environment

- More effective treatments
- Reduced side effects
- Faster diagnosis
- Better prevention

Prediction of the 3-D structure of a protein

- Long-standing challenge in biology
- Important for understanding biological processes to create new drugs
- Publicly published in 2021
- Nobel Prize in Chemistry 2024

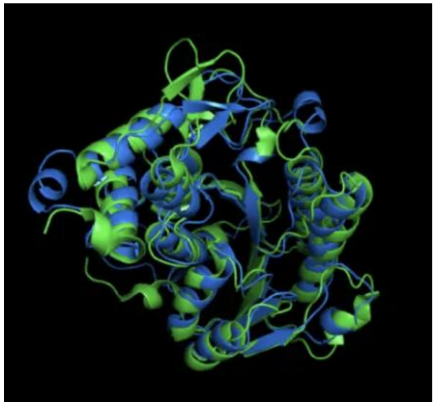

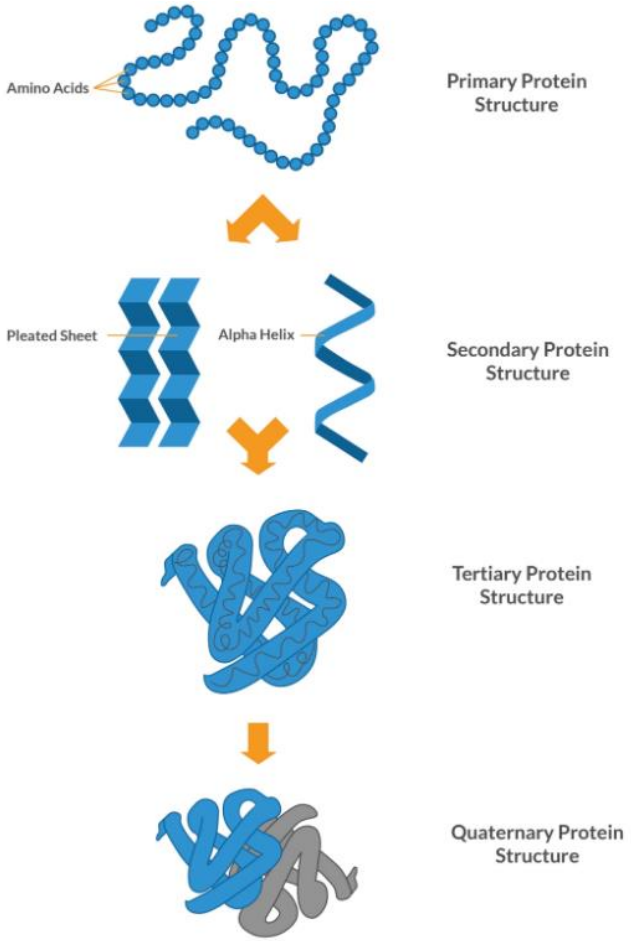

Medical imaging

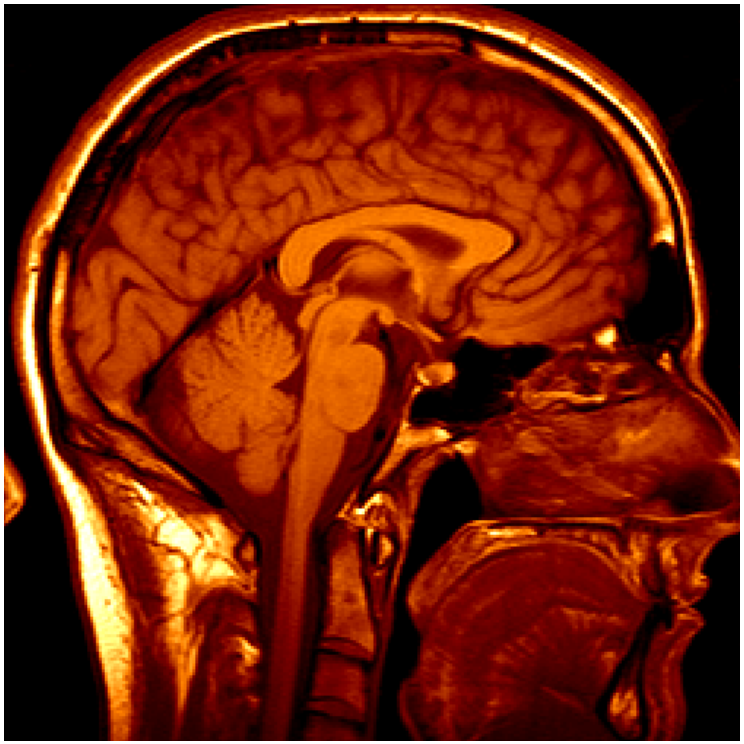

3D imaging
MRI, CT

Image guided surgery
Grimson et al., MIT

Helping people with disabilities

Estimate gaze direction to move the mouse

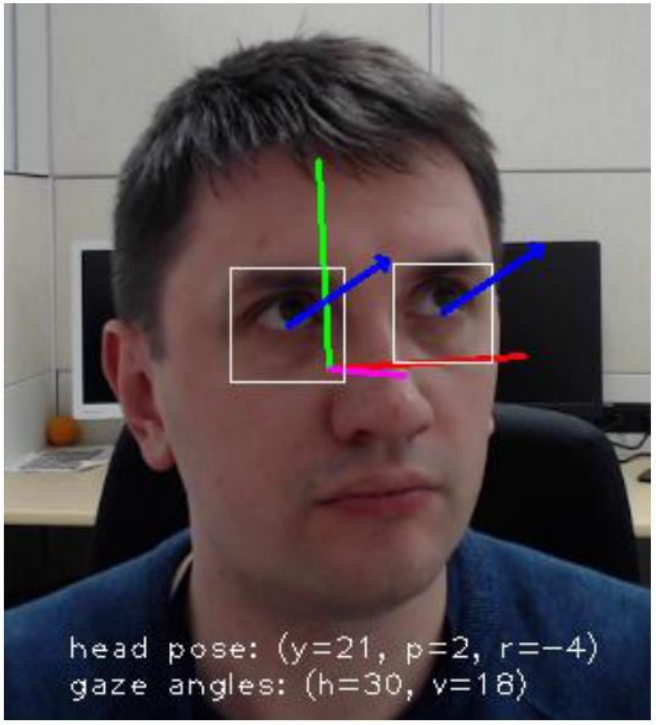

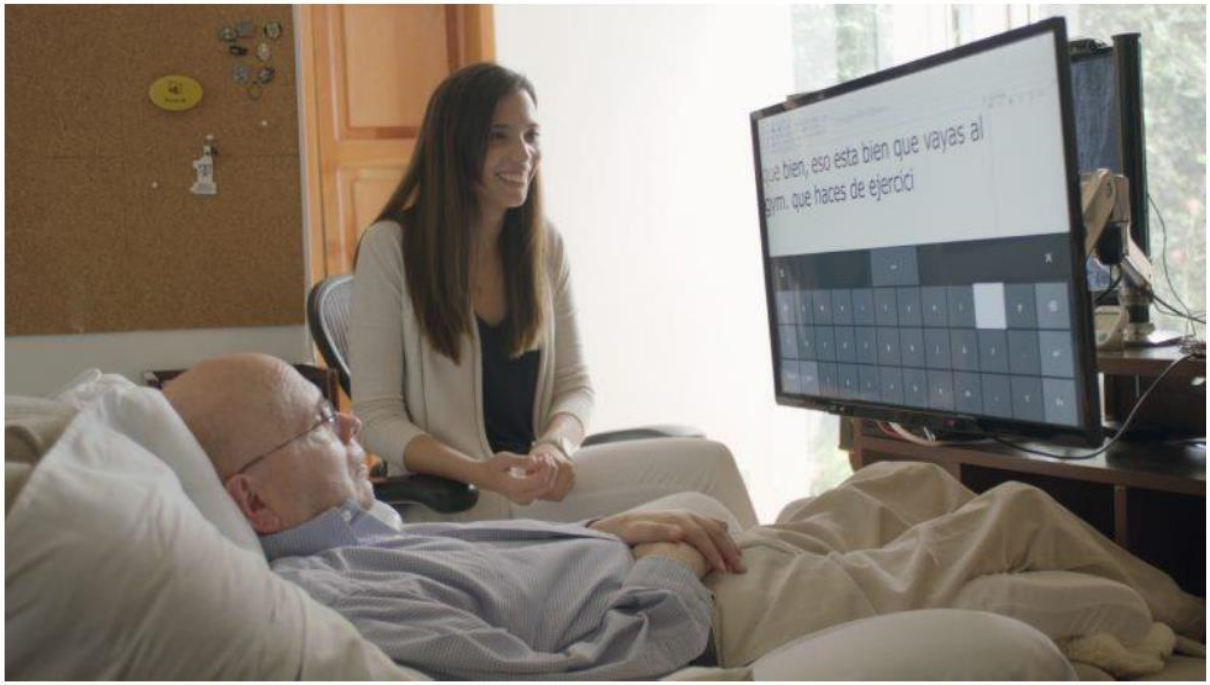

Applications of AI (transportation)

Self-driving cars

Law enforcement

- License plate recognition for speed cameras and freeway tolls
- Improving urban CCTV cameras

Applications of AI (smart home)

- Energy Efficiency
- Enhanced Security
- Device Automation
- Adaptive Lighting

Applications of AI

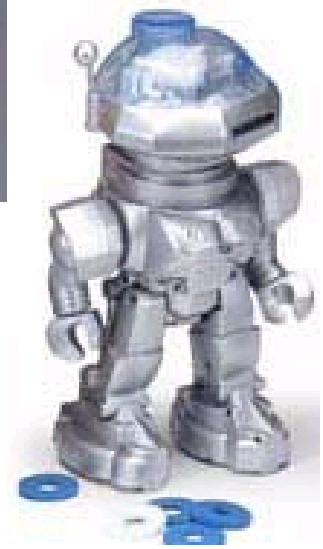

Games and entertainment

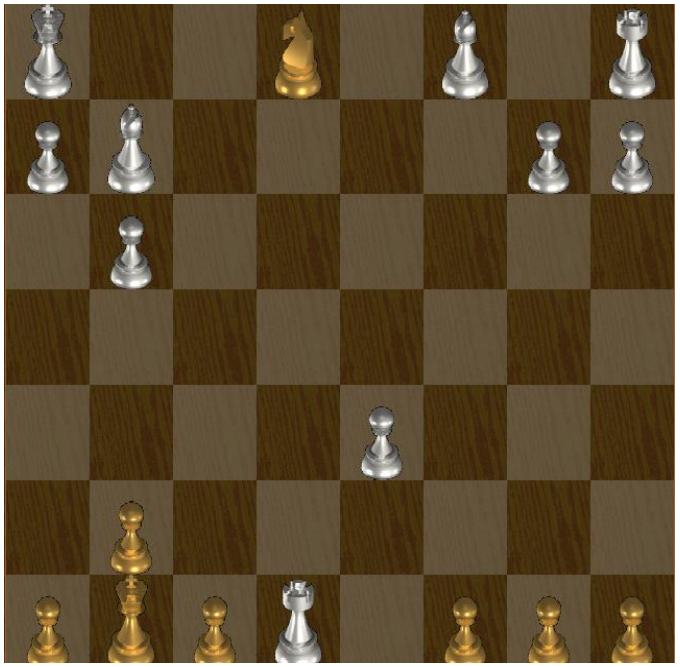

Application examples

Recommender systems

Goal: Recommend items users might like

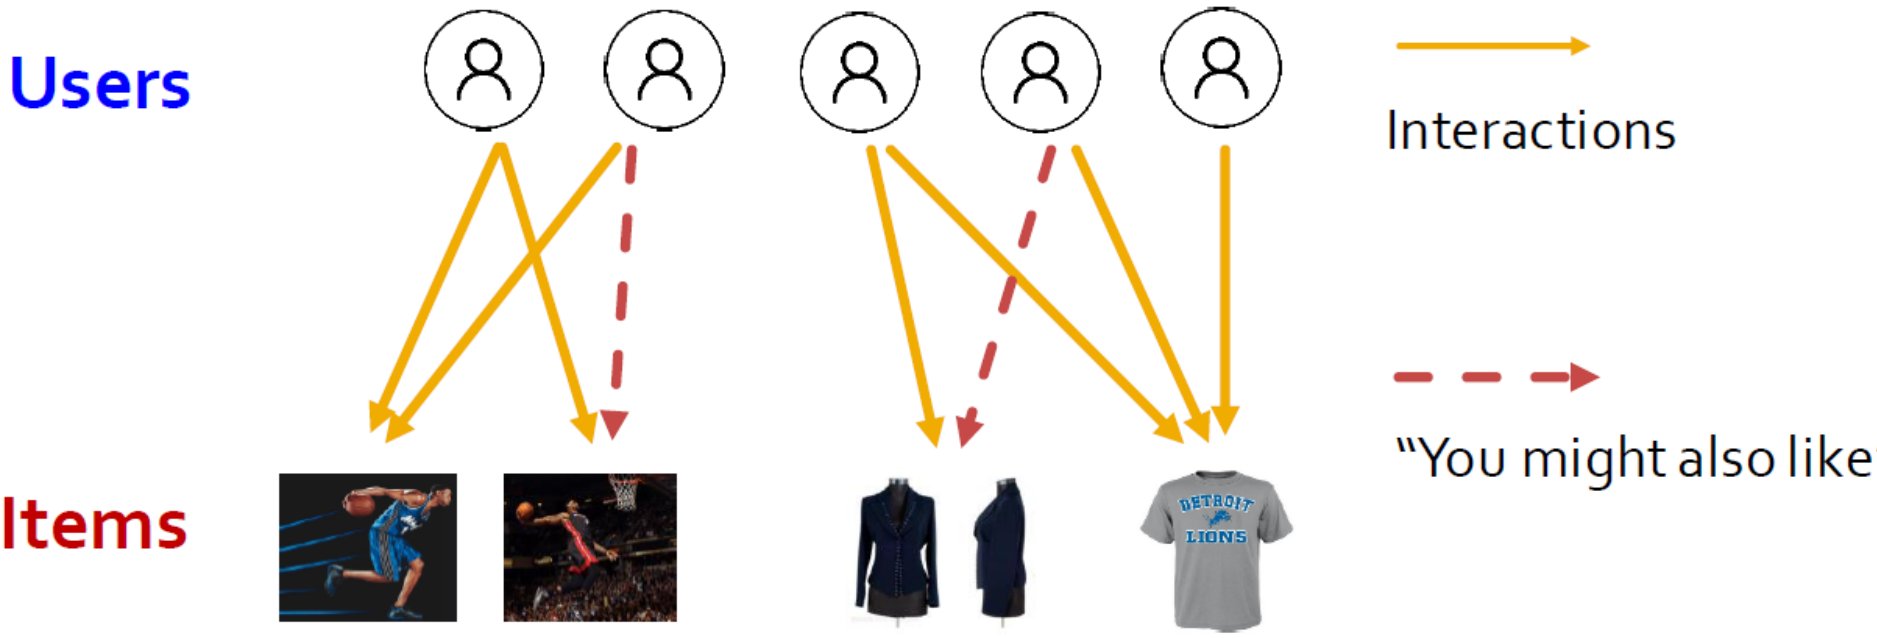

Generative AI: Image Restoration

Conditional Image

Inpainting with L2 loss

Inpainting with CGAN

Context Encoders: Feature Learning by Inpainting, D.Pathak, P. Krahenbuhl, J. Donahue, T. Darrell, A. Efros,2016

Generative AI: Combining Styles

Generative AI: Image-to-Image Translation

Generative AI: Image-to-Image Translation

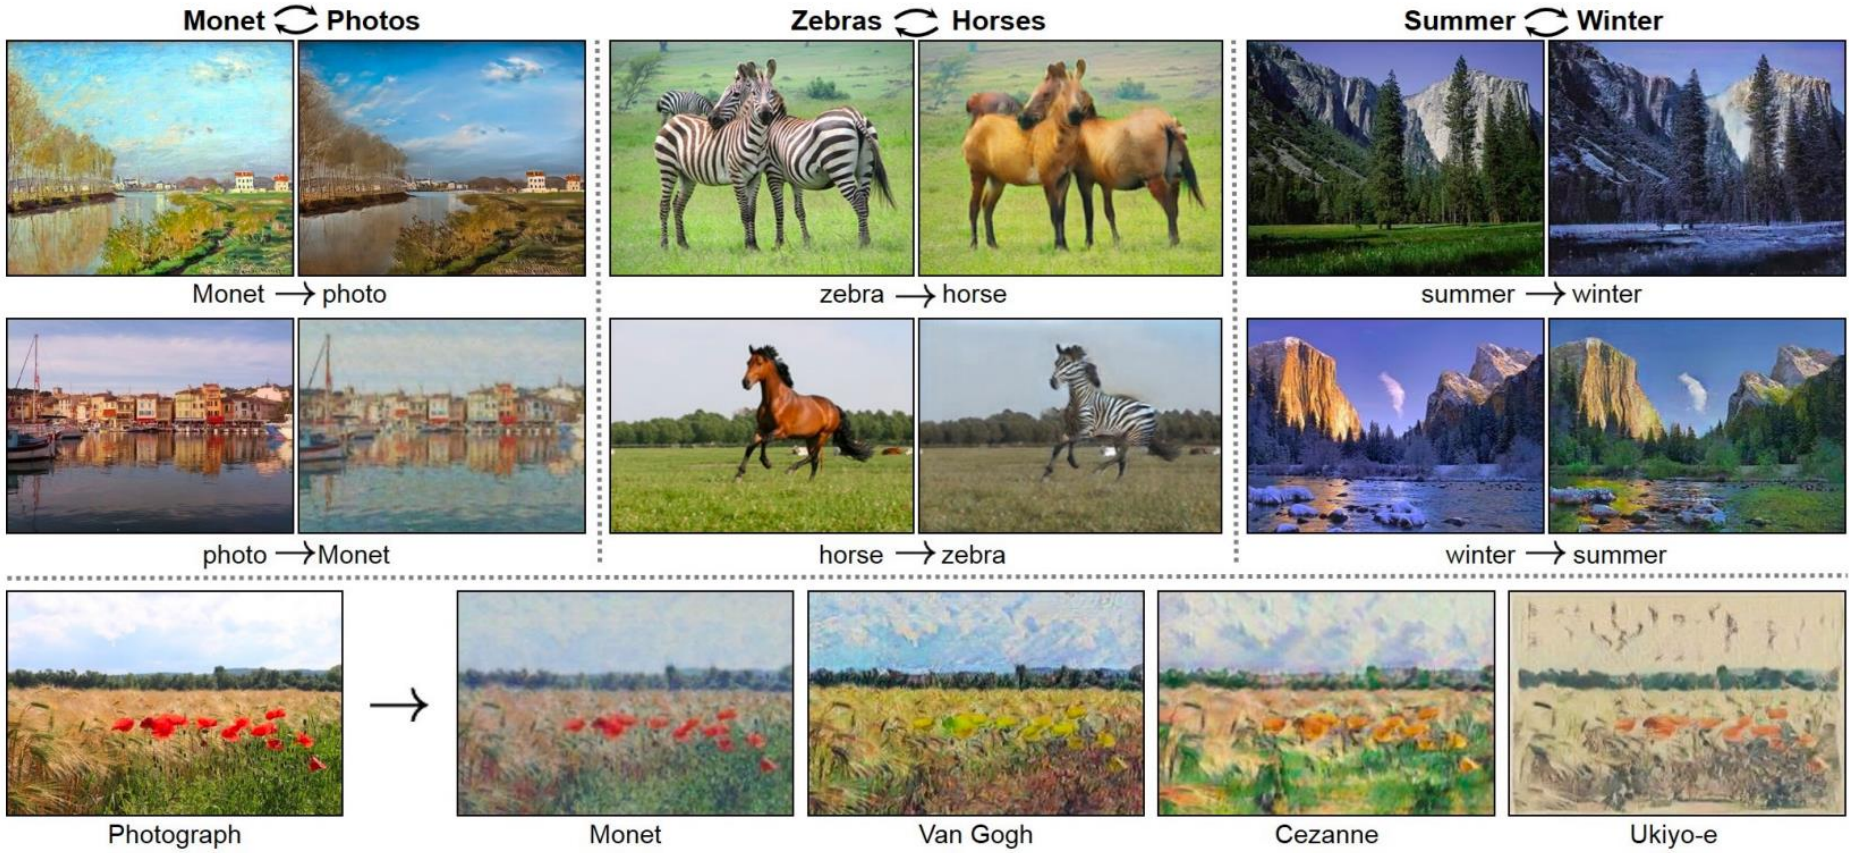

https://github.com/junyanz/CycleGAN

Generative AI: Video-to-Video Translation

https://github.com/junyanz/CycleGAN

Generative AI: Style Transformation

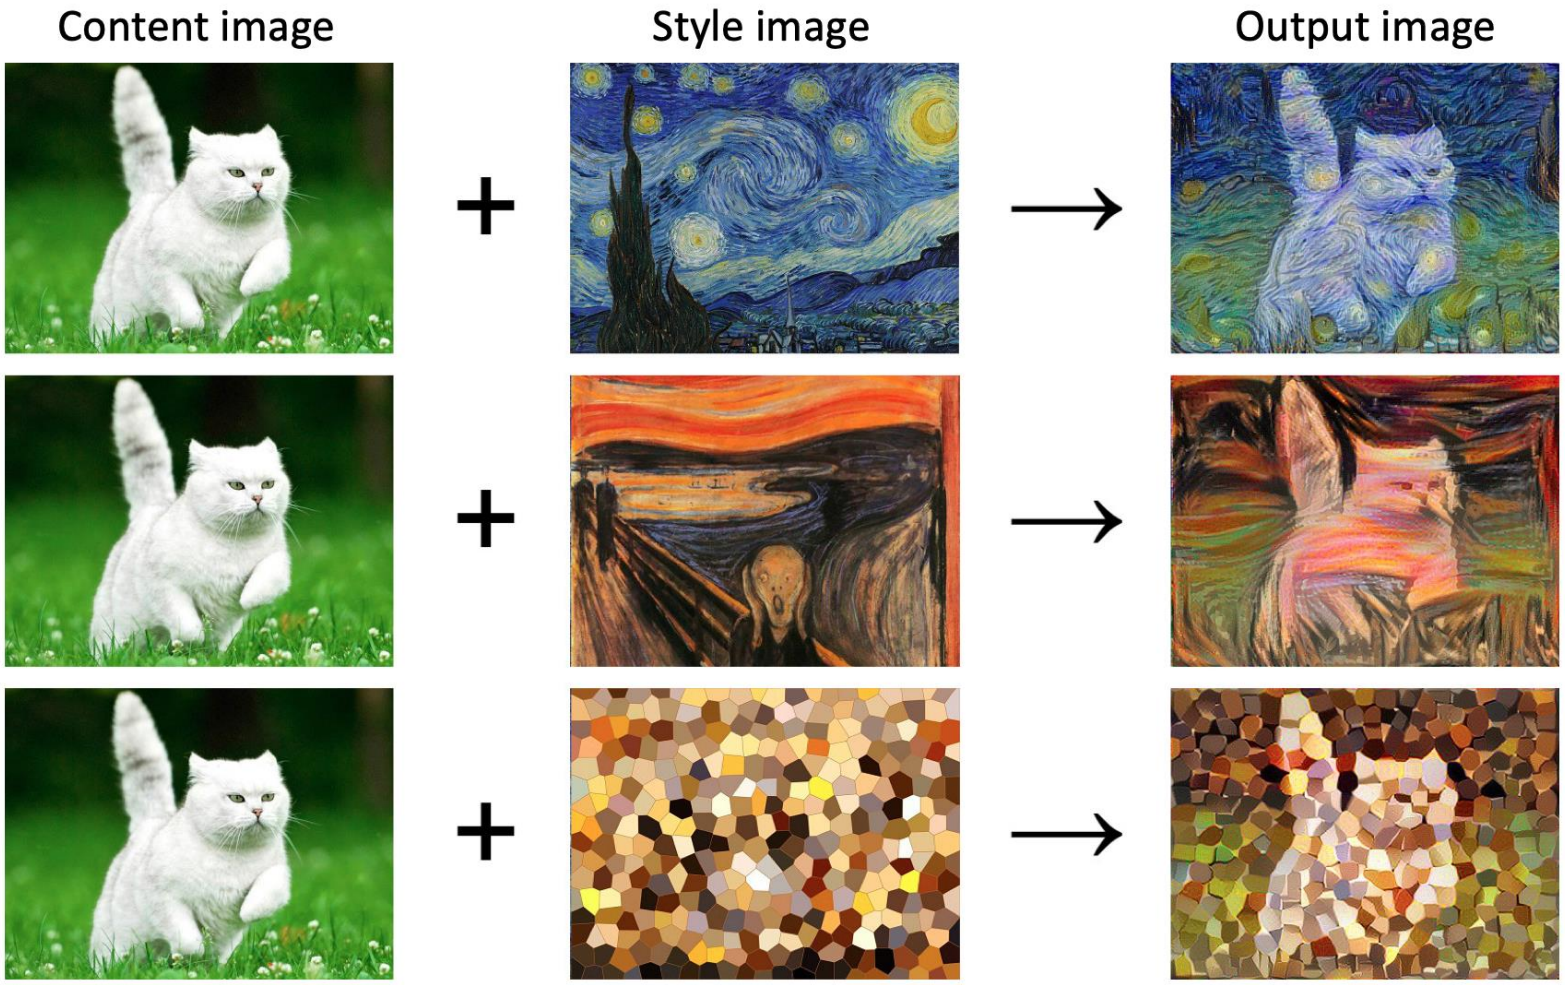

102

Generative AI: Text-to-Image

A lion in a hoodie hacking on a laptop

Teddy bears shopping for groceries in ancient Egypt

Teddy bears working on new AI research on the moon in the 1980s

Generative AI: Text-to-Video

OpenAI: Sora (2024)

Create high-quality real-time videos from a text description

https://openai.com/sora

Prompt: Several giant wooly mammoths approach treading through a snowy meadow, their long wooly fur lightly blows in the wind as they walk, snow covered trees and dramatic snow capped mountains in the distance, mid afternoon light with wispy clouds and a sun high in the distance creates a warm glow, the low camera view is stunning capturing the large furry mammal with beautiful photography, depth of field.

Who is concerned with NNs?

- Computer scientists want to find out about the properties of non-symbolic information processing with neural nets and about learning systems in general.
- Statisticians use neural nets as flexible, nonlinear regression and classification models.
- Engineers of many kinds exploit the capabilities of neural networks in many areas, such as signal processing and automatic control.
- Cognitive scientists view neural networks as a possible apparatus to describe models of thinking and consciousness (High-level brain function).
- Neuro-physiologists use neural networks to describe and explore medium-level brain function (e.g. memory, sensory system, motorics).

Who is concerned with NNs?

- Physicists use neural networks to model phenomena in statistical mechanics and for a lot of other tasks.
- Biologists use Neural Networks to interpret nucleotide sequences.
- Philosophers and some other people may also be interested in Neural Networks for various reasons


Questions