What is Artificial Intelligence?

The term artificial intelligence (AI) was coined in 1956 at Dartmouth Artificial Intelligence Conference.

Artificial: Made by human beings rather than occurring naturally like biological neural networks.

Intelligence: Able to interact with the world including people, animals, machines.
Able to model the world and provide reasoning
Able to learn and to adapt.

The purpose of AI is to let computers think like humans.

To determine whether an artificial entity is intelligent, Alan Turing proposed in 1950:

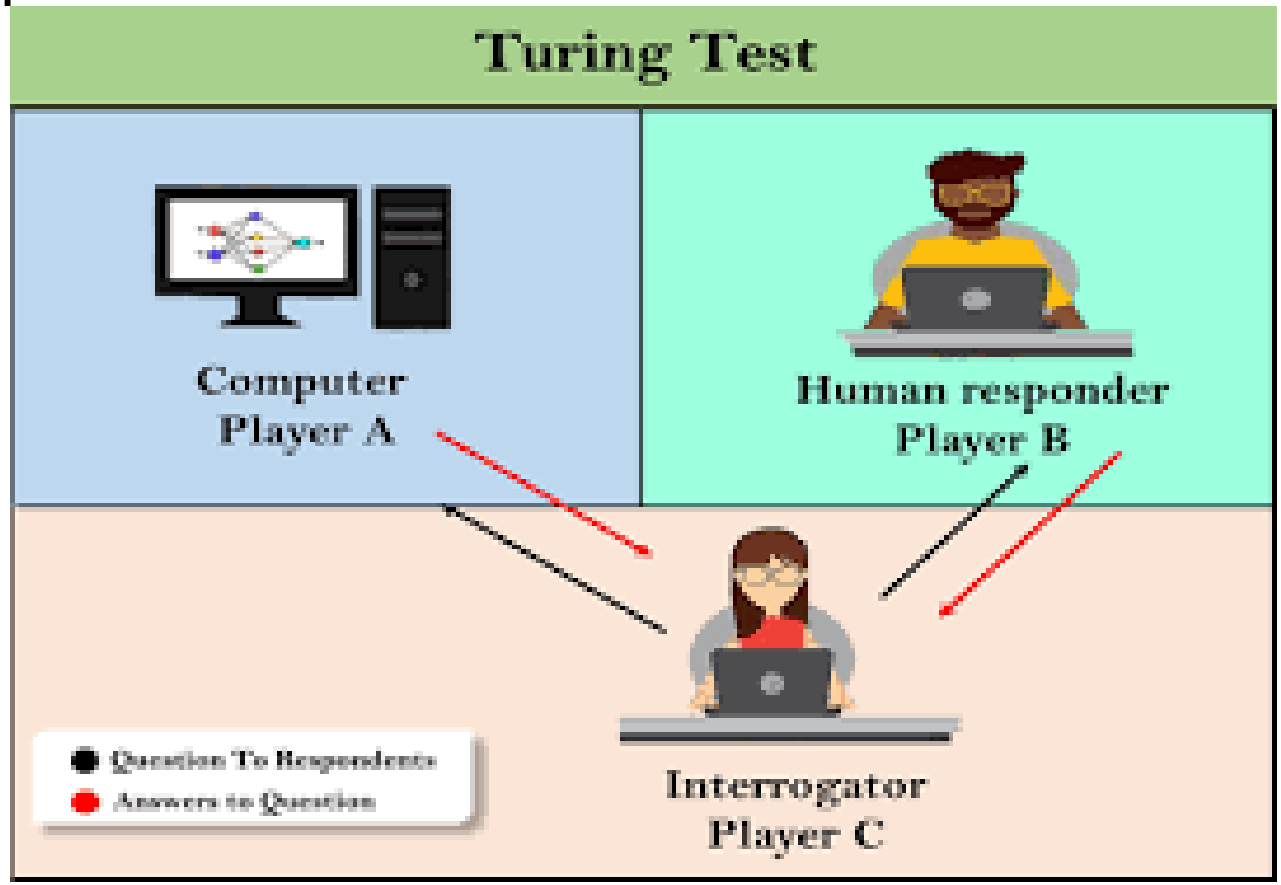

Source: Turing Test in AI - VietMX's Blog

The test is passed if Player C cannot distinguish Player B from Player A. No physical interaction is necessary.

Physical interaction is required in total Turing test: The artificial entity must be able to do, in the real world of objects and people, everything that real people can do. That is, a robot is involved, e.g., the Terminator, if it really exists, can pass the test.

Source: Ishiguro Lab. - Geminoid (osaka-u.ac.jp), The Terminator - Wikipedia

Strong AI

- Machine will be built in human form with consciousness.
- Same sensory perception as human with a mind.
- It may go through the same education and learning processes as human child.
- The machine is "born" as a child and eventually develop to an adult in a way analogous to human development.

Not only robot receptionist …

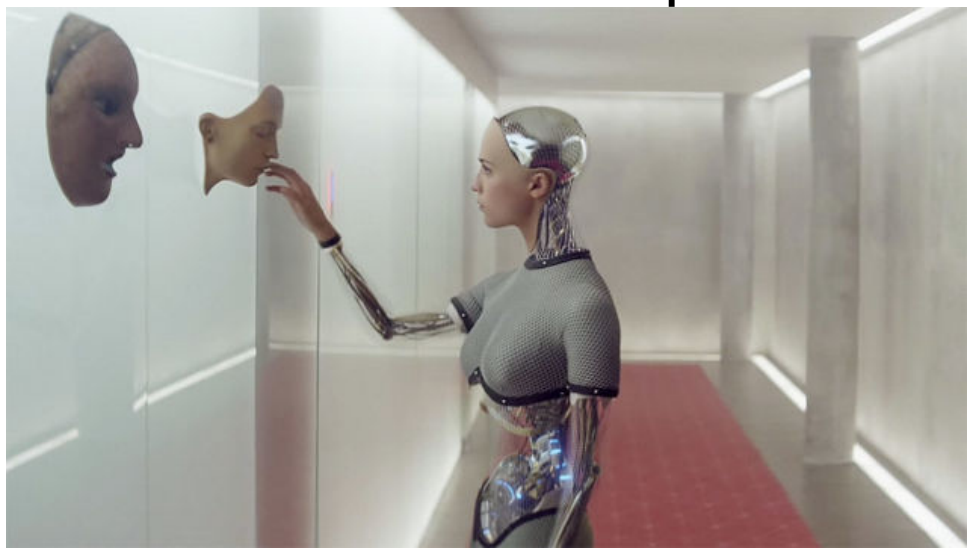

but also capable of thought and consciousness

Source: Robot staff make Japan's Henn na Hotels quirky and efficient - Nikkei Asia, Ex Machina (2014) - IMDb

Weak AI

- An artificial entity which appears and behaves intelligently, but not necessarily understands.
- Instead of having a mind in strong AI, it can only simulate a mind.
- Weak or narrow AI models human intelligence and works to accomplish specific tasks rather than possess full cognitive abilities like the human brain, e.g.,

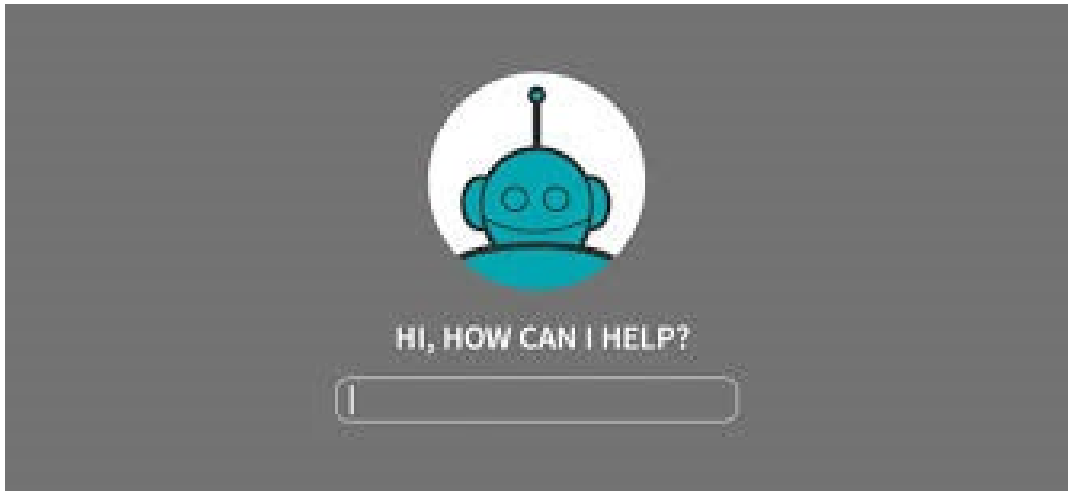

Source: Digital assistants and PPC – what does the future hold? - Click Consult, How self-driving cars can shape our future? (readwrite.com)

First AI Wave

Boomed during 1950s-1960s but entered its winter in 1970s because only toy problems could be handled.

Engineers devised deductive reasoning programs with logical rules, leading to simple games such as the first computer program capable of playing draughts or checkers.

Free Online Draughts Game - play against the computer

These first-wave AI systems could perform straightforward reasoning tasks, and could not be applied to business at the time.

Second AI Wave

Boomed during 1980s, which emerged from expert systems research, but entered its winter again in around 1995.

An expert system is based on knowledge acquisition that emulates the decision-making ability of a human expert.

Knowledge from human experts and their conversion to programmable rules, mainly if–then rules, are required. Together with statistical model development, machines can adapt knowledge to different situations.

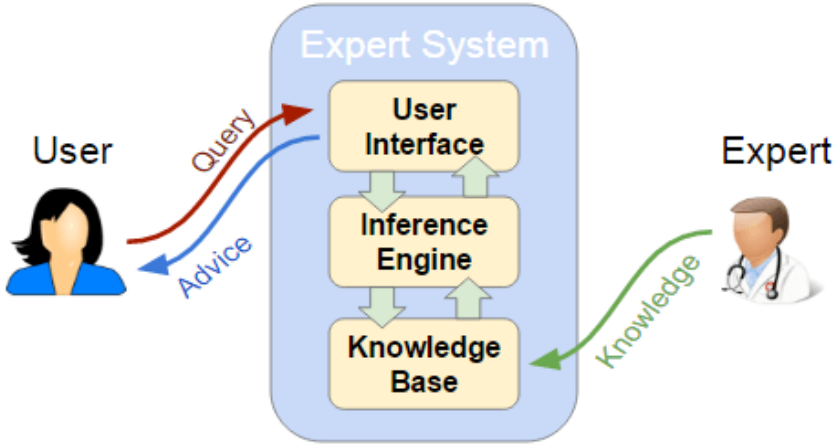

Supercomputer Deep Blue defeated chess world champion in 1997, is a well-known example.

Source: IBM Deep Blue 1997 - chess playing computer defeated Kaspa… | Flickr

However, complexity and accuracy issues such as creation and maintenance of expert systems rendered them impractical for most business applications.

Now we are at the third AI wave boom, which is mainly driven by big data, high computing power, and deep learning.

We are aware of this wave in 2016 after AlphaGo defeated go world champion.

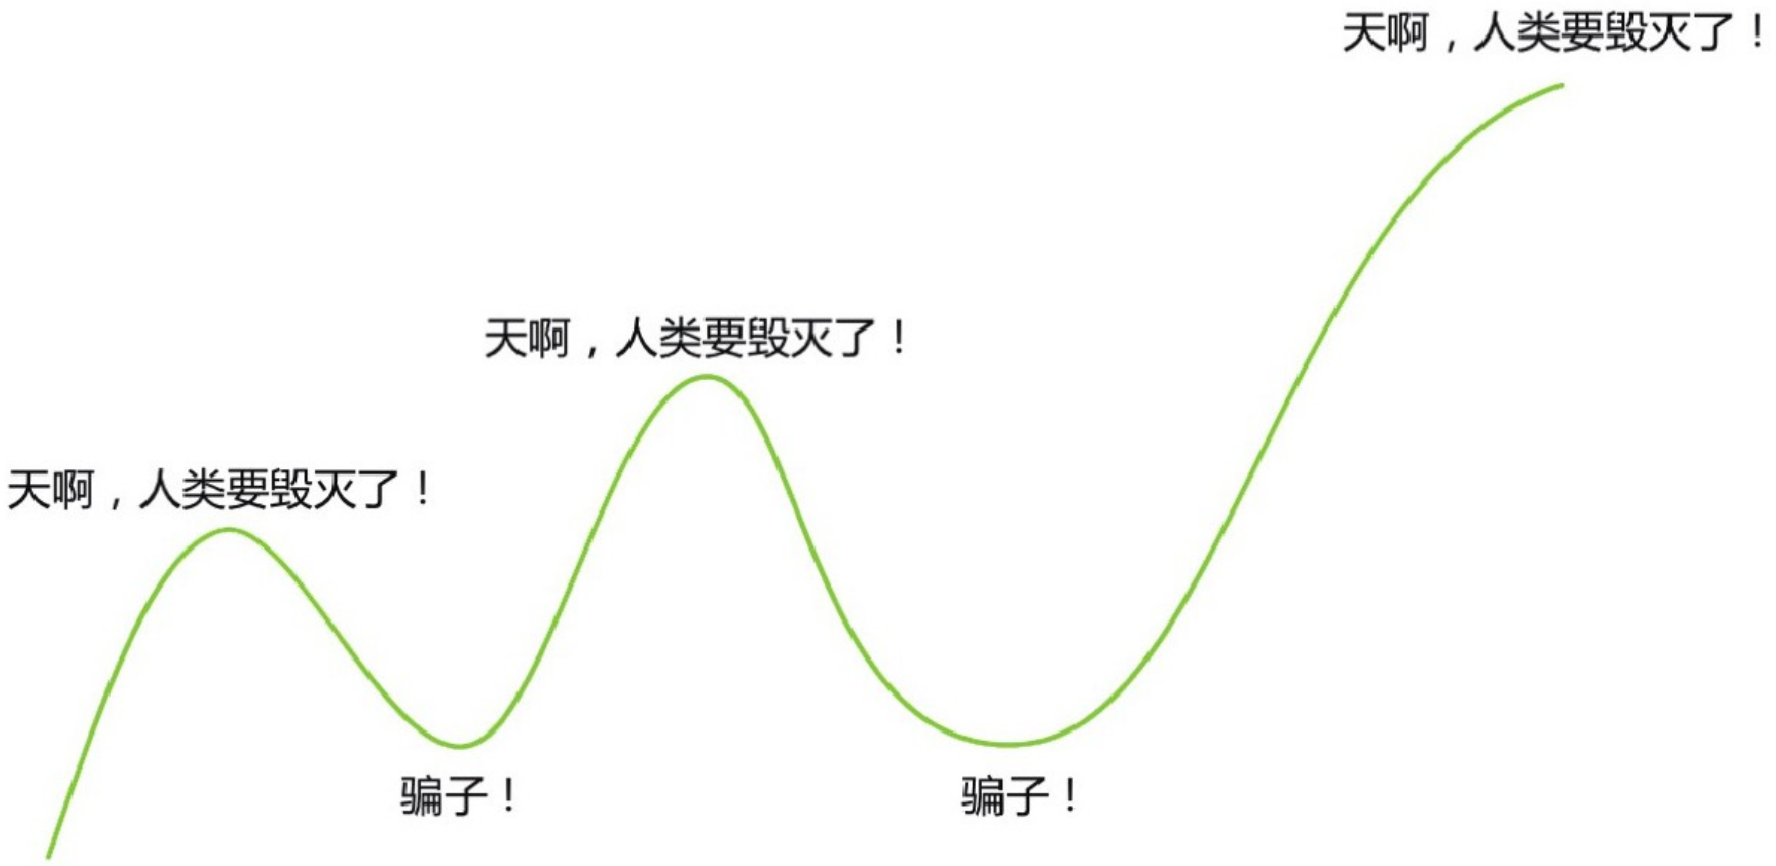

Source: 李開復, 王詠剛, 人工智慧來了, 天下文化, 2017

Checkers (西洋跳棋) → Chess (國際象棋) → Go (圍棋)

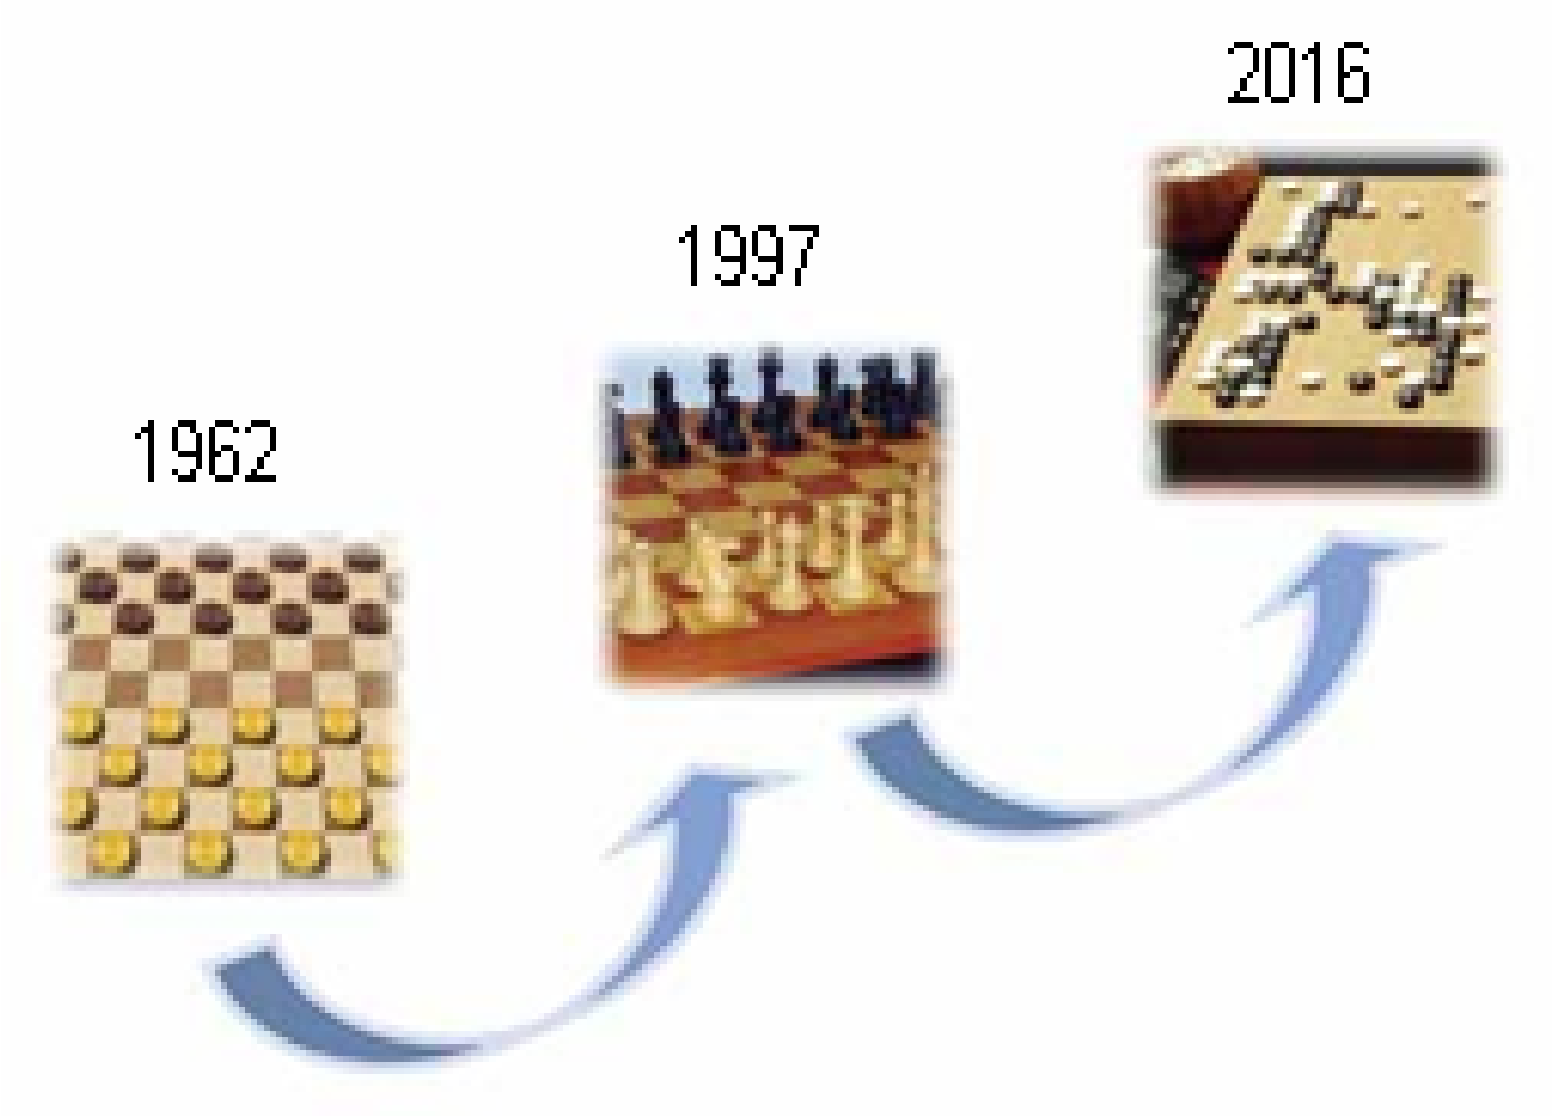

Source: 李開復, 王詠剛, 人工智慧來了, 天下文化, 2017

Machine Learning: A branch of AI that specializes in how the computer simulates or realizes human learning behavior, including acquire new knowledge or skills to continually improve its performance.

Deep Learning: An important machine learning approach based on neural network and massive training data.

Data Science: This subject combines different fields including statistics, applied mathematics, data visualization, pattern recognition, machine learning, high-performance computing, in order to interpret massive and complex data and/or extract meaningful information from data for decision making.

Big Data: Any voluminous amount of structured, semi-structured and unstructured data that has the potential to be mined for information. Traditional data processing application software is usually inadequate to deal with them.

Why Artificial Intelligence is Important?

We are already surrounded by numerous AI products:

- Intelligent personal assistants such as Google Assistant, Siri, Cortana, and Amazon Alexa.
- News summarization, recommendation and writing such as Summly, MachinaNova and Automated Insights.
- Computer vision applications such as face recognition, Google Photos and Amazon Go.
- AI powered art apps such as Prisma and Meitu.
- Search engines such as Google, Bing and Baidu.
- Automatic translation applications such as Google Translate and ili.
- Autonomous cars such as Tesla Autopilot and Waymo self-driving cars.
- Robots such as DHL unmanned aerial vehicles and Wonder Workshop Dash.

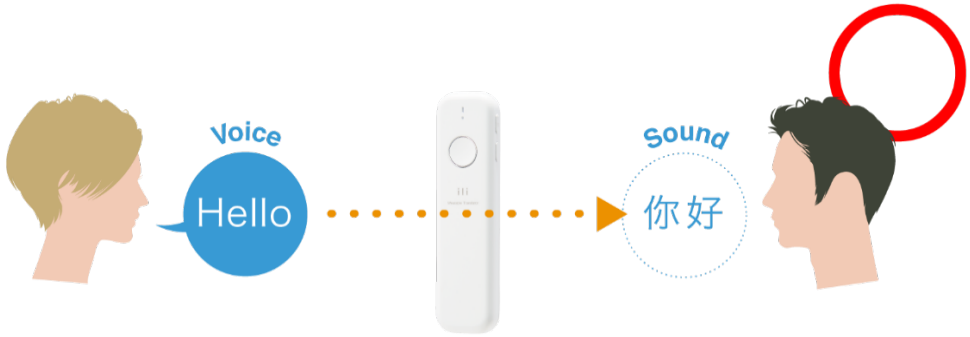
Voice
Hello
Sound
你好


Self checkout

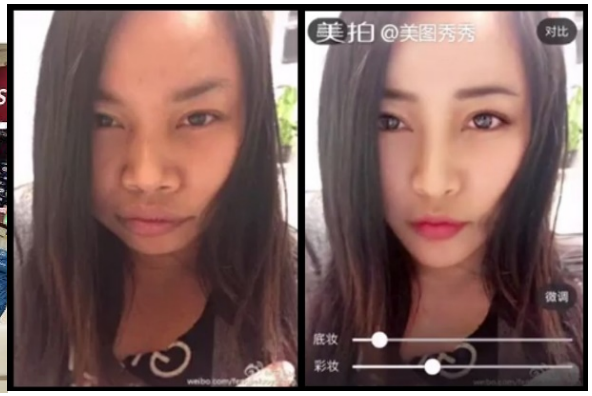
美拍@美图秀秀
对比
微调
底妆
彩妆


AT&T
11:52 PM
100%
YAHOO!
Sharks Dive Deep Under Full Moon
LiveScience.com
A full moon and warm waters may send some sharks diving deep, according to a new study. Over the course of nearly three years, researchers from Australia observed 39 mostly female gray reef sharks living near coral reefs in Palau, Micronesia, east of the Philippines. In the winter, the sharks stayed closer to the surface, at an average depth of 115 feet , where water was consistently warmer, the team found.
Summary by YAHOO!


AT&T
11:27 PM
100%
YAHOO!
McDowell wins in playoff at RBC Heritage
The Associated Press
Graeme McDowell defeated fellow U.S. Open champ Webb Simpson in a playoff to win the RBC Heritage on Sunday for his second career victory on the PGA Tour. McDowell was four shots behind when the round started, but rallied with a closing 69 over wind-blown Harbour Town Golf Links. McDowell struck his approach on the extra hole about 15 feet away and two-putted for par.
Summary by YAHOO!


dash

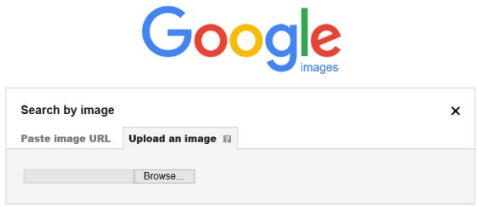
Google
images
Search by image
Paste image URL
Upload an image
Browse...


WAYMO


DHL

According to *Future of Jobs Report* published by World Economic Forum in 2020:

"The pace of technology adoption is expected to remain unabated and may accelerate in some areas. The adoption of cloud computing, big data and e-commerce remain high priorities for business leaders, following a trend established in previous years. However, there has also been a significant rise in interest for encryption, non-humanoid robots and artificial intelligence."

"Based on these figures, we estimate that by 2025, 85 million jobs may be displaced by a shift in the division of labour between humans and machines, while 97 million new roles may emerge that are more adapted to the new division of labour between humans, machines and algorithms."

How Artificial Intelligence Works

We focus on machine learning, which can be classified as 4 types:

Supervised Learning:

- Also referred to as learning with a teacher.
- A set of input-output examples is given for training, e.g., inputs can be images of cat and dog, while outputs can be their respective labels 'cat' and 'dog'.
- Application areas include classification (or recognition or detection) such as identifying an animal from a known finite set, and regression (or estimation or prediction) such as predicting the numerical value of electricity usage given the input of temperature and time of the day.
- Principal techniques include decision trees, *k*-nearest neighbors, naive Bayes classifier, support vector machine, and neural network.

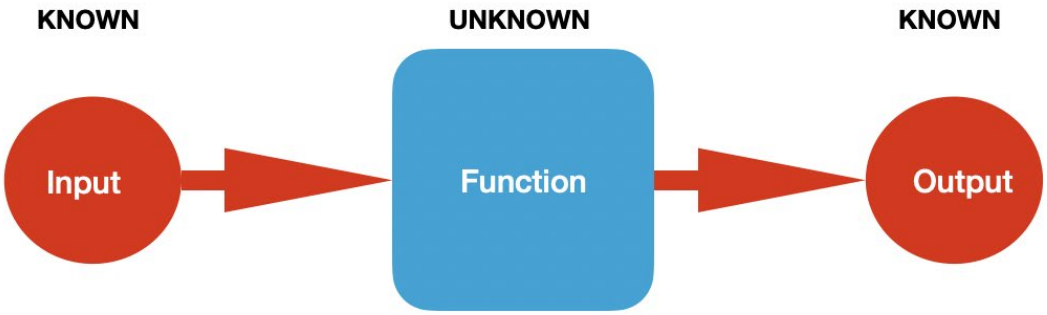

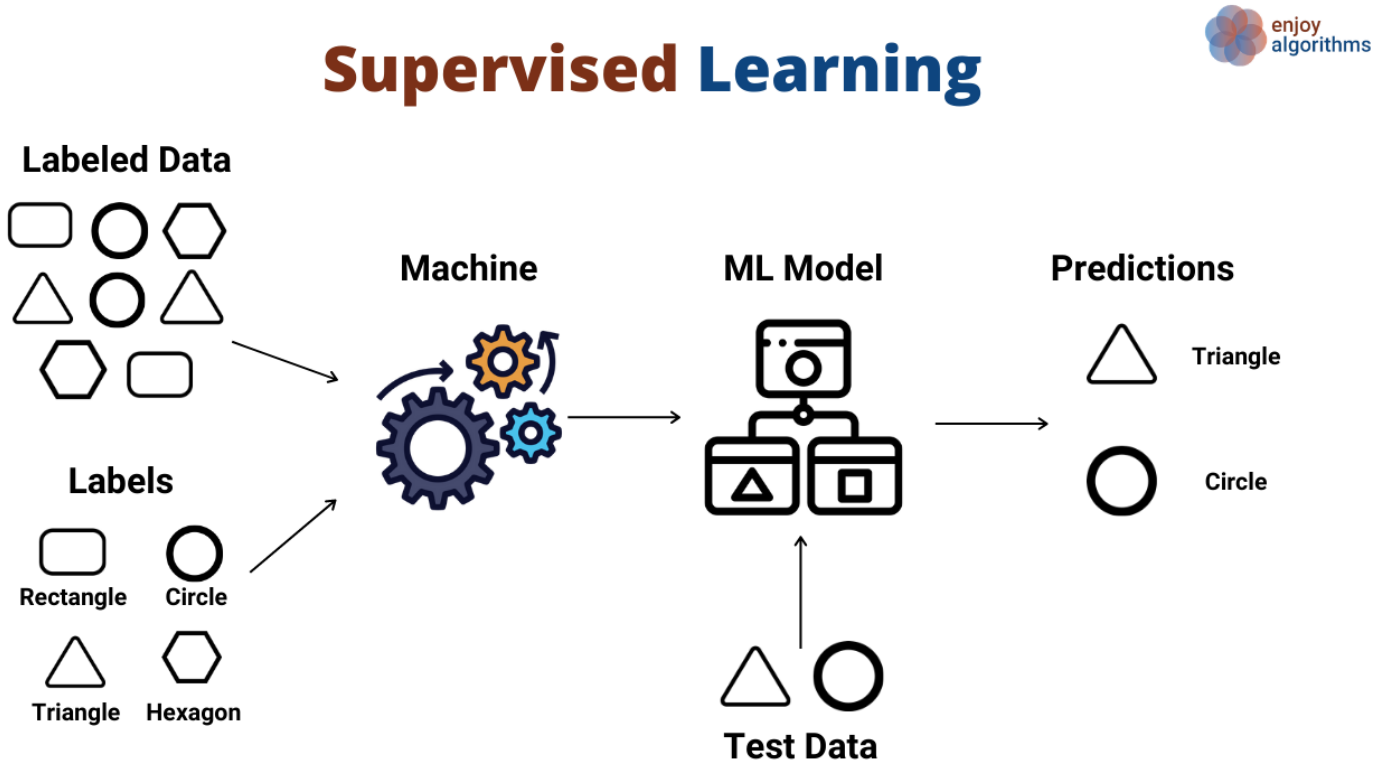

Source: [Supervised, Unsupervised, And Semi-Supervised Learning With Real-Life Usecase (enjoyalgorithms.com)]

Unsupervised Learning:

- No teacher is available, i.e., we only have the inputs but not their desired outputs, or all data are not labelled
- Mapping function is also unavailable as in supervised learning
- A main task is to group or cluster the data according to whatever features it deems most important, e.g., number of sides in a polygon
- It can also be used for dimensionality reduction.
- Principal techniques include *k*-means clustering, self-organization map, hidden Markov model, principal component analysis, and neural network.

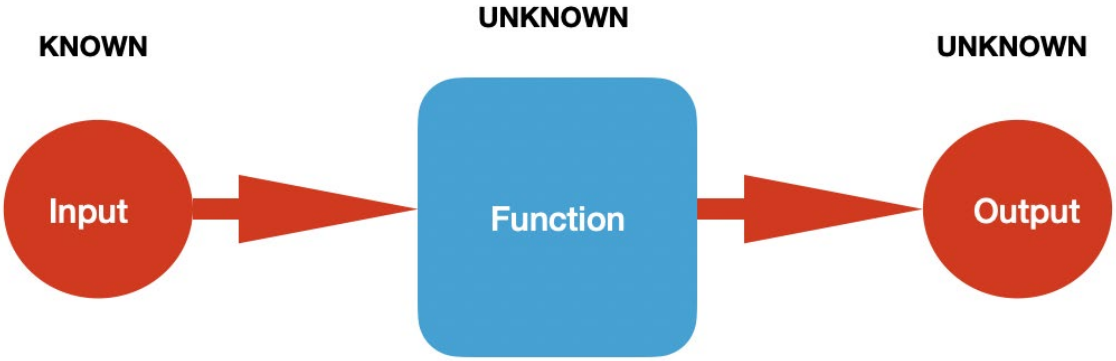

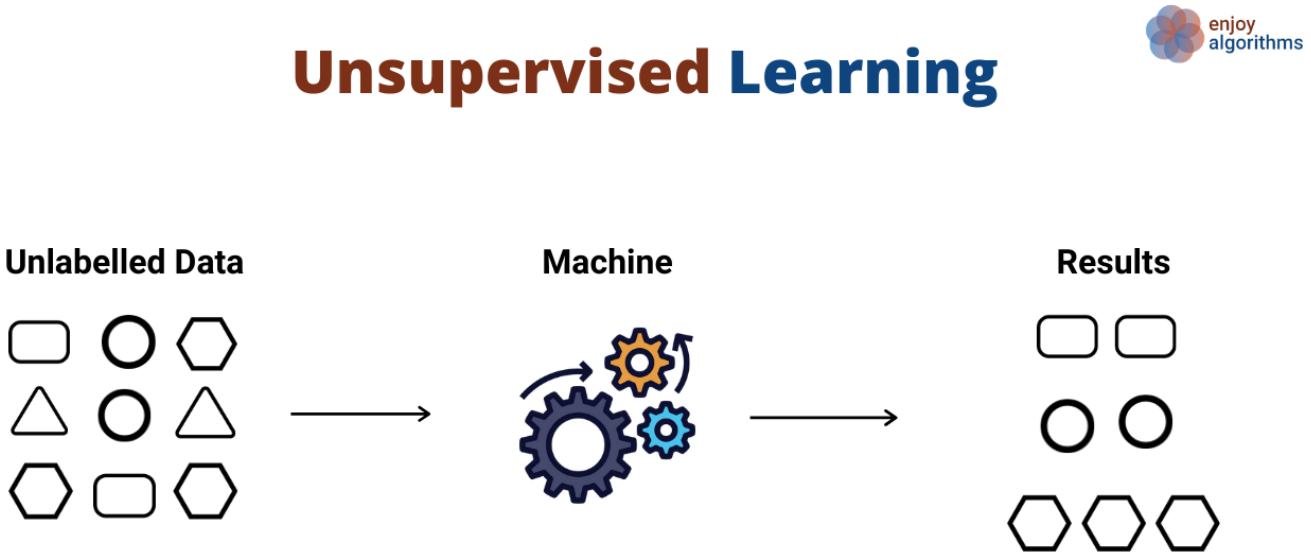

Source: [Supervised, Unsupervised, And Semi-Supervised Learning With Real-Life Usecase (enjoyalgorithms.com)]

Semi-Supervised Learning:

- Between supervised learning and unsupervised learning, i.e., partial outputs are provided.

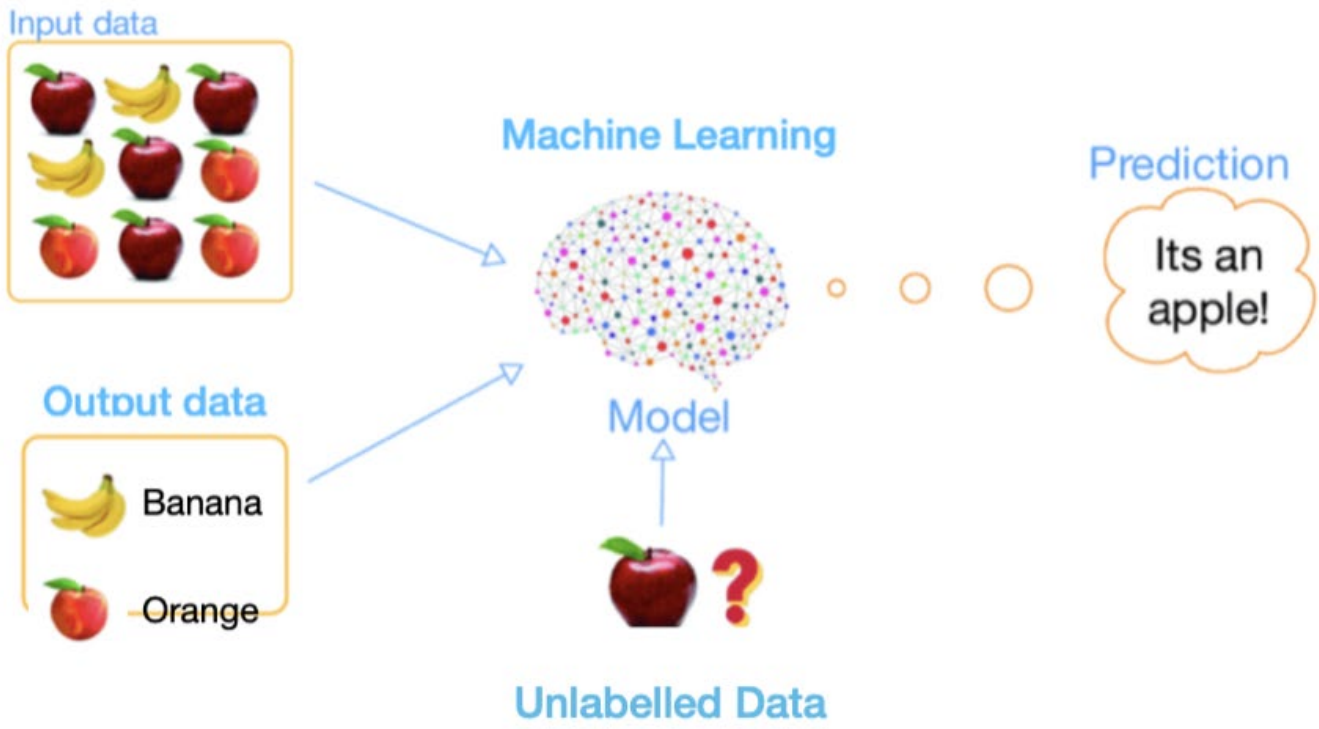

Source: [Supervised, Unsupervised, And Semi-Supervised Learning With Real-Life Usecase (enjoyalgorithms.com)]

Reinforcement Learning:

- An algorithm aims at using observations gathered from the interaction with the environment to take actions to maximize the reward and/or minimize the risk.
- It learns by trial and error from consequences of taking actions, namely, being rewarded for successful behavior and/or being punished for unsuccessful behavior.

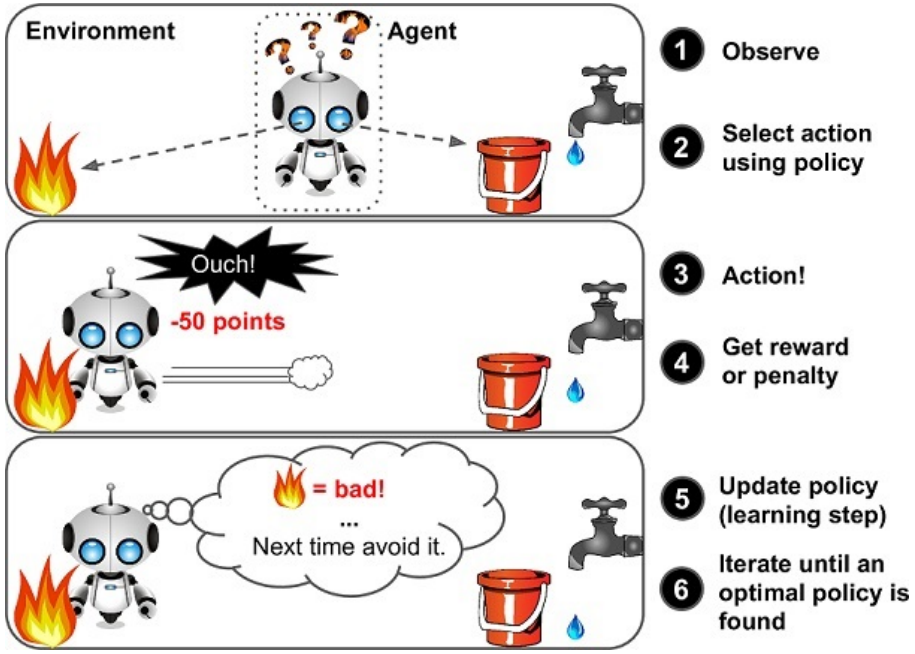

Source: reinforcement learning retail (buy9586.top)

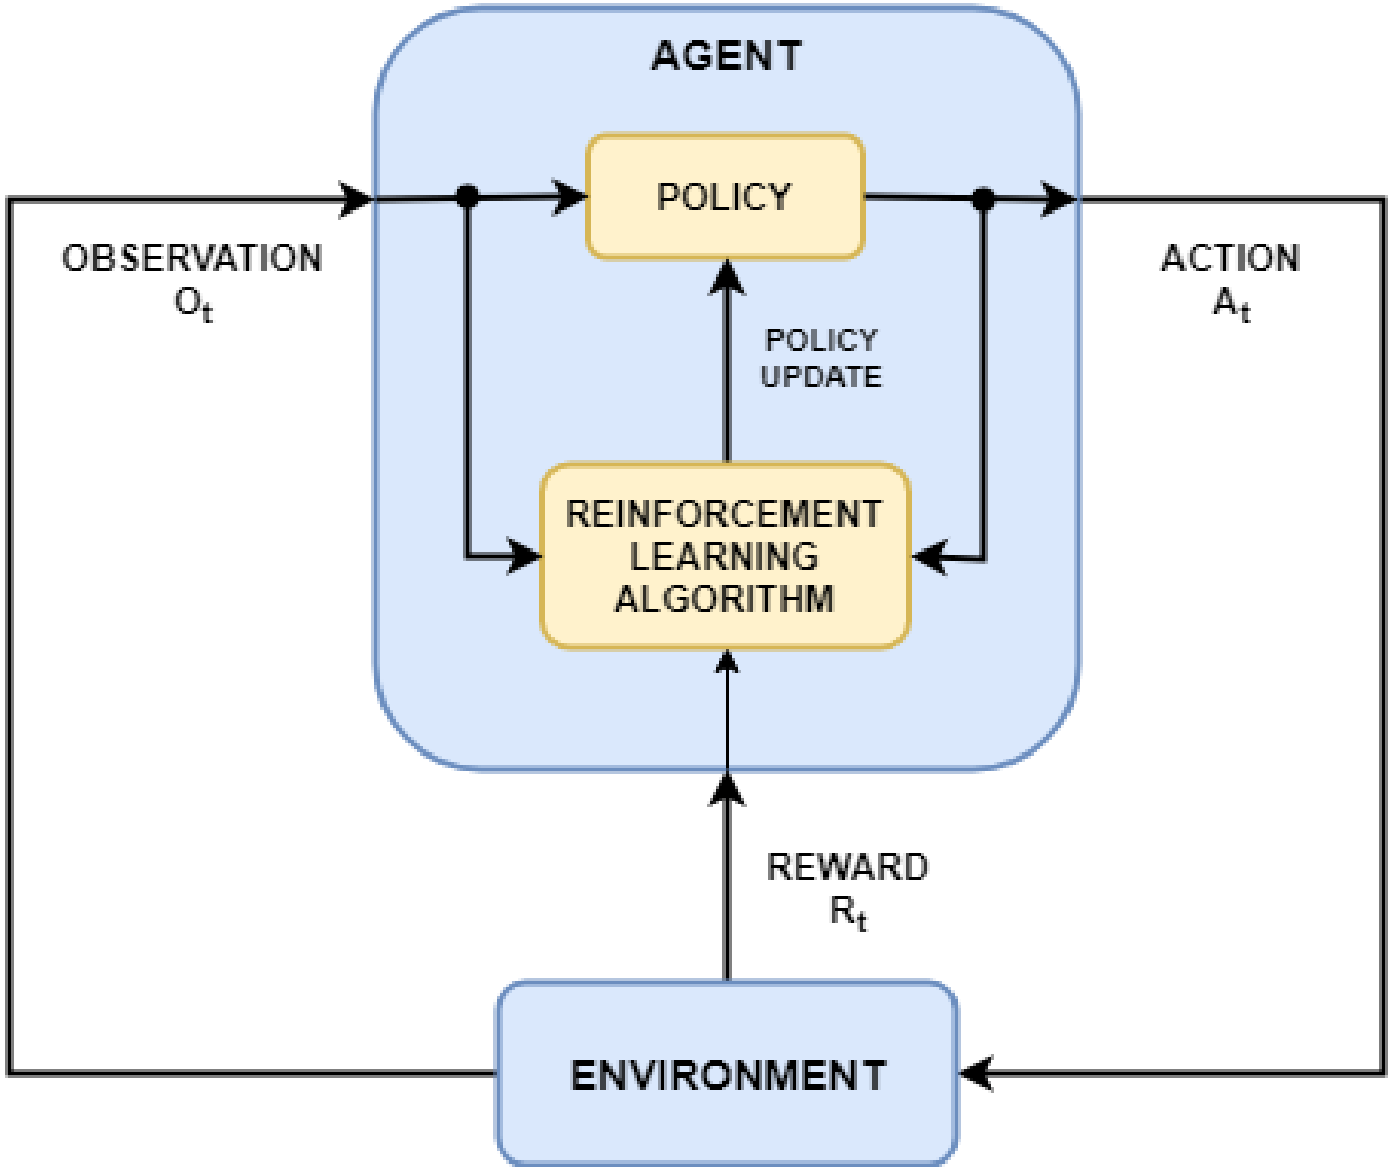

Source: [What Is Reinforcement Learning? - MATLAB & Simulink - MathWorks 中国]

Among numerous machine learning methods, now the hottest one is deep learning.

Briefly speaking, deep learning refers to an artificial neural network (ANN) with many layers and massive training data.

Human nervous system may be viewed as a 3-stage system:

- Brain or neural net continually receives information, perceives it, and makes appropriate decisions.
- Receptors convert stimuli from the human body or the external environment into electrical impulses that convey information to the brain.
- Effectors convert electrical impulses generated by the neural net into discernible responses as system outputs.

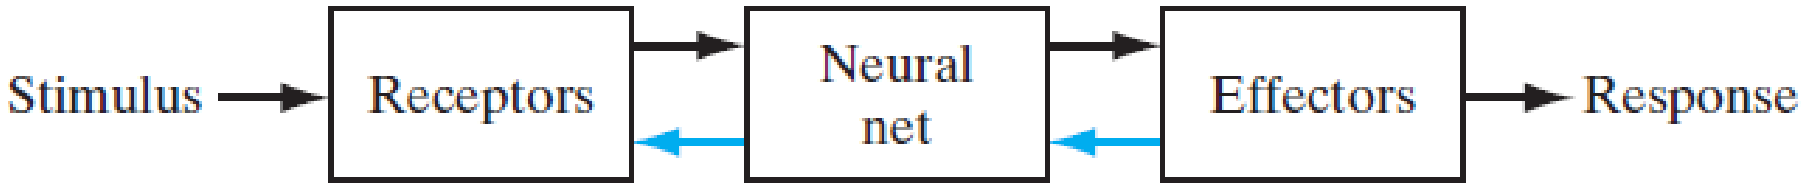

ANN simulates learning mechanisms in biological organisms:

- Brain contains cells which are called neurons.
- Neurons are connected to one another with axons and dendrites.
- Connecting regions between axons and dendrites are called synapses.
- Strength of synaptic connections changes in response to stimuli, and this change is how learning take places.
- ~10 billion neurons and 60 trillion synapses in human.

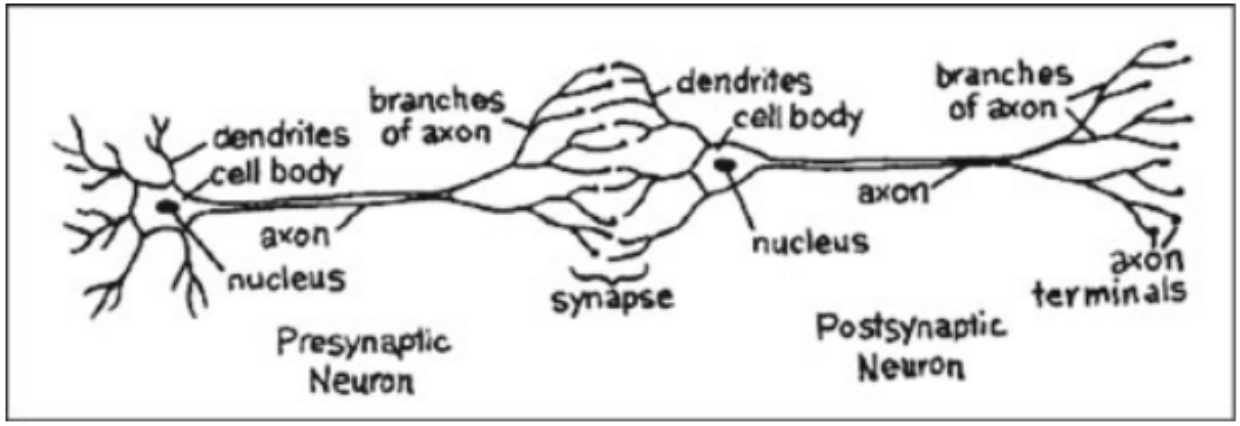

(a) Biological neural network

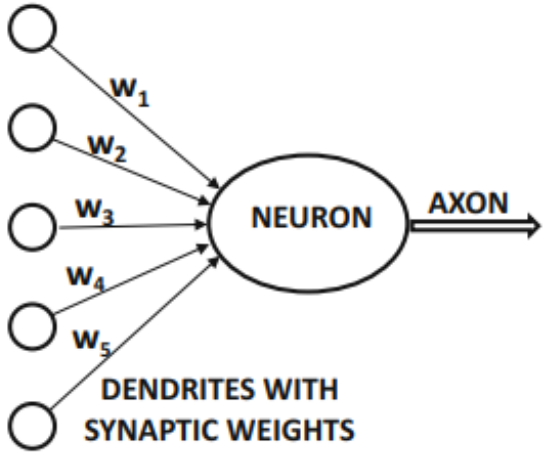

(b) Artificial neural network

ANN is a computational unit:

- Artificial neurons connect together as a network.
- Inputs arrive through presynaptic connections.
- Synaptic connection strength is modelled as weights, i.e., each input to a neuron is scaled with a weight w_i.
- Neuron computes a nonlinear function based on the inputs and weights, and propagates the output function value to connected neurons.
- Learning occurs by changing the weights connecting the neurons, and completes when the weight values become steady.
- Analogous to human brain, ANN has the fault tolerance potential, i.e., even if a neuron or its connecting links are damaged, performance degradation may not be occurred.

A mathematical neuron model is a cascade interconnection:

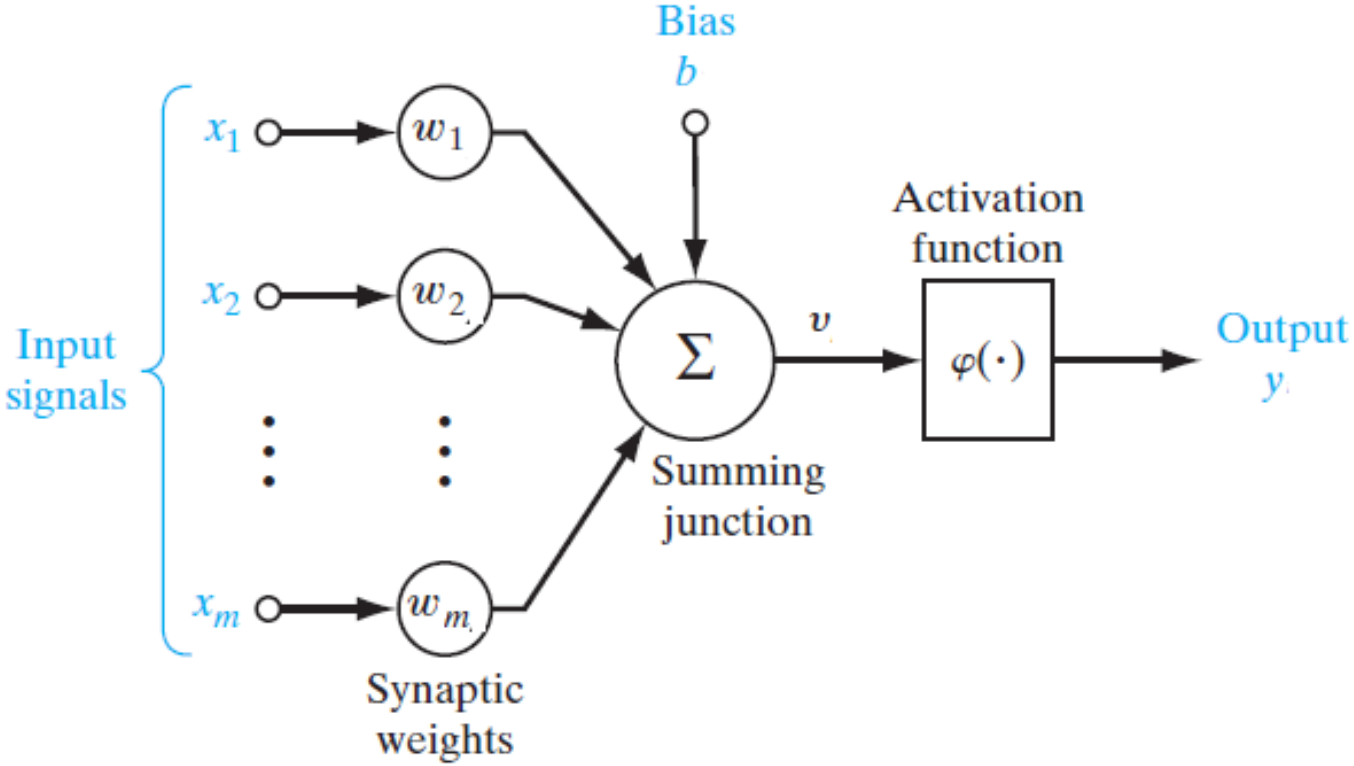

First stage: linear

$$v = \sum_{j=1}^{m} w_j x_j + b$$

Second stage: nonlinear

$$y = \varphi(v)$$

Nonlinear activation function examples:

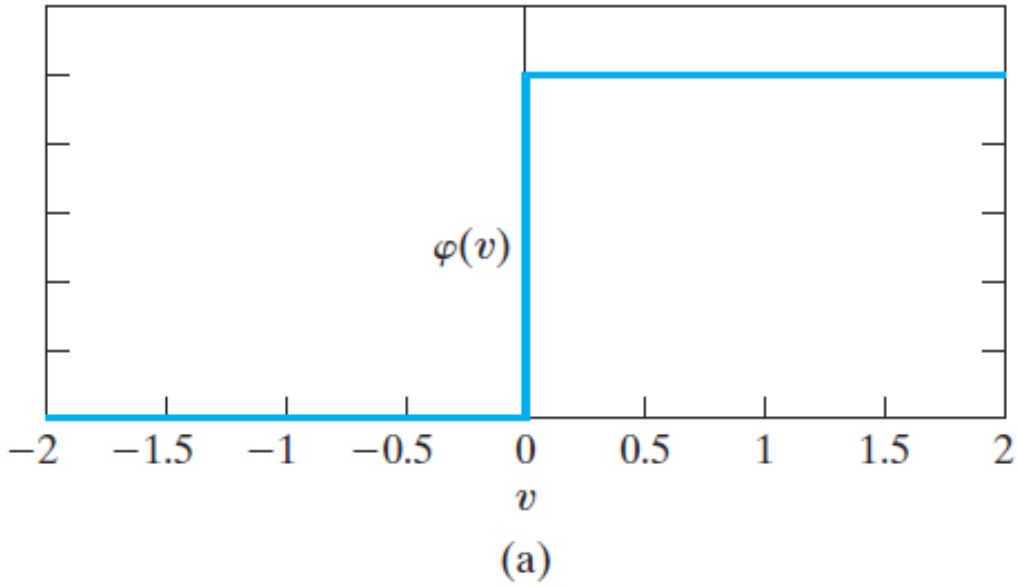

unit step function:

$$\varphi(v) = \begin{cases} 1, & v > 0 \\ 0, & v < 0 \end{cases}$$

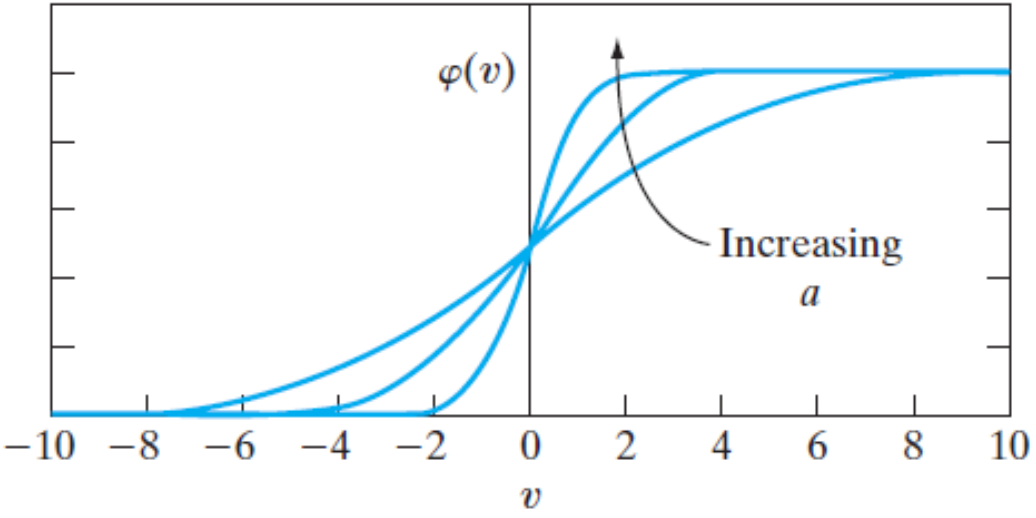

sigmoid function:

$$\varphi(v) = \frac{1}{1 + e^{-av}}$$

An example of 2-hidden layer neural network:

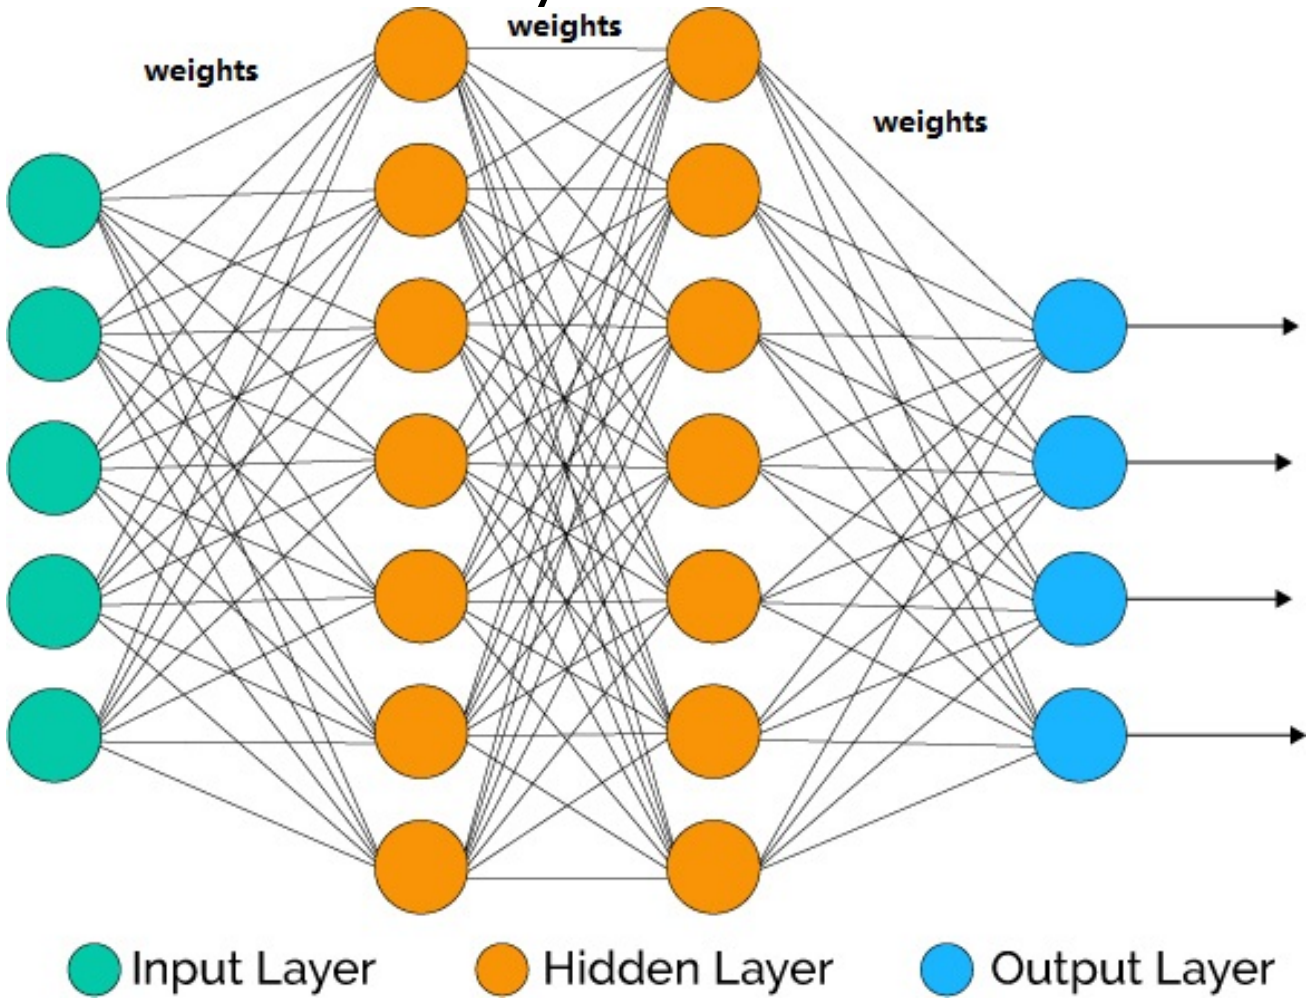

Given the training data of input-output pairs, the weights between nodes (or neurons) are adjusted until certain criteria are met.

"Non-deep" feedforward neural network

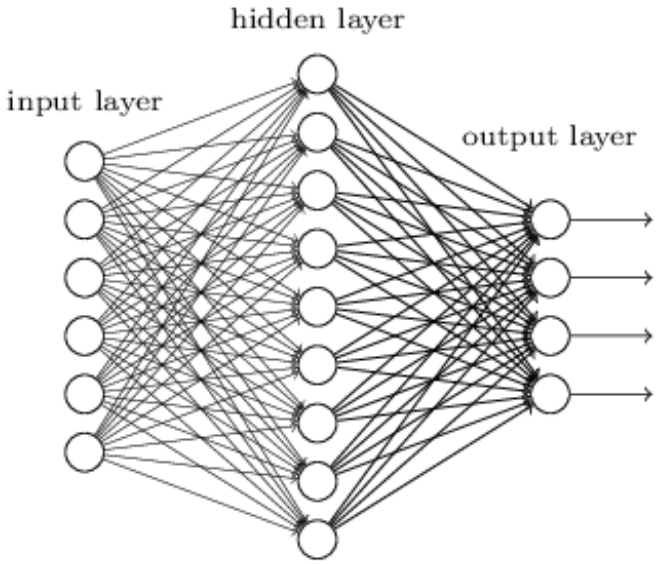

Deep neural network

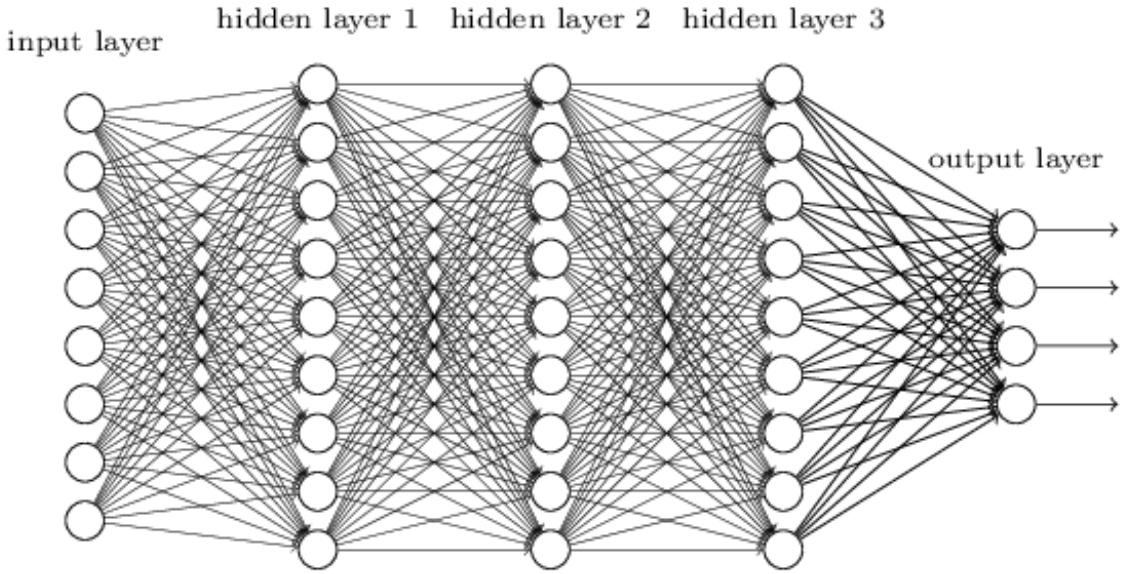

Source: What is the difference between a neural network and a deep neural network, and why do the deep ones work better? - Cross Validated (stackexchange.com)

Non-deep or shallow: One hidden layer which may be sufficient for most problems but feature extraction step is normally required.

Deep: Raw data as input with numerous hidden layers and nodes.

Breakthrough in speech recognition (around 95% accuracy now which is threshold for human).

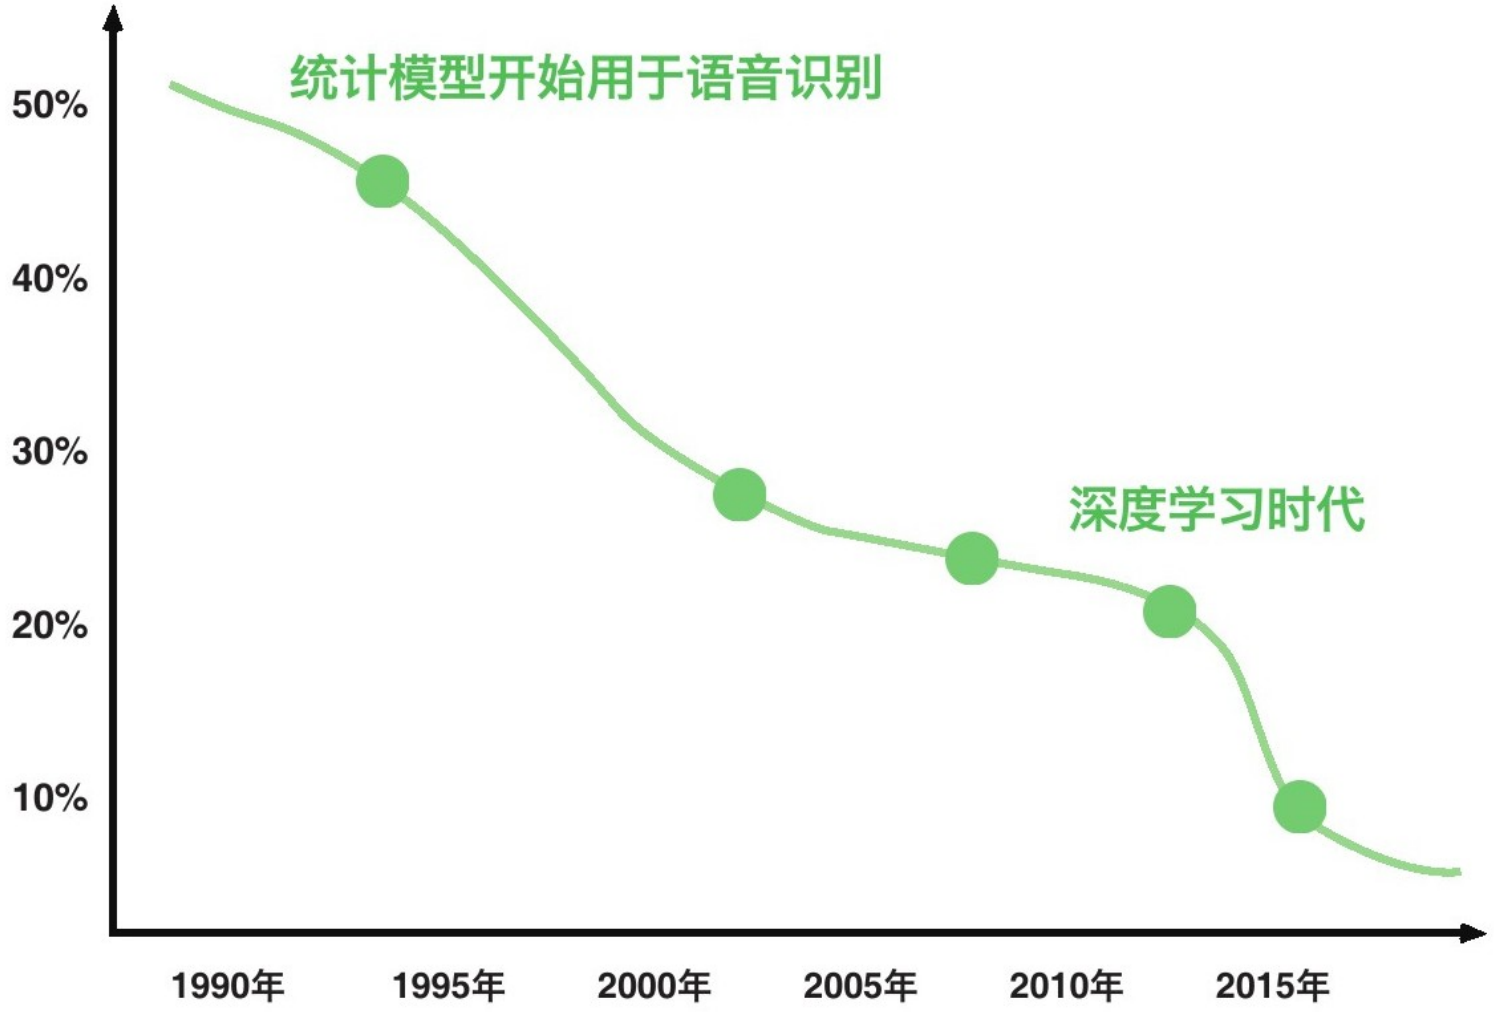

Source: 李開復, 王詠剛, 人工智慧來了, 天下文化, 2017

Breakthrough in image recognition in ImageNet Large Scale Visual Recognition Challenge.

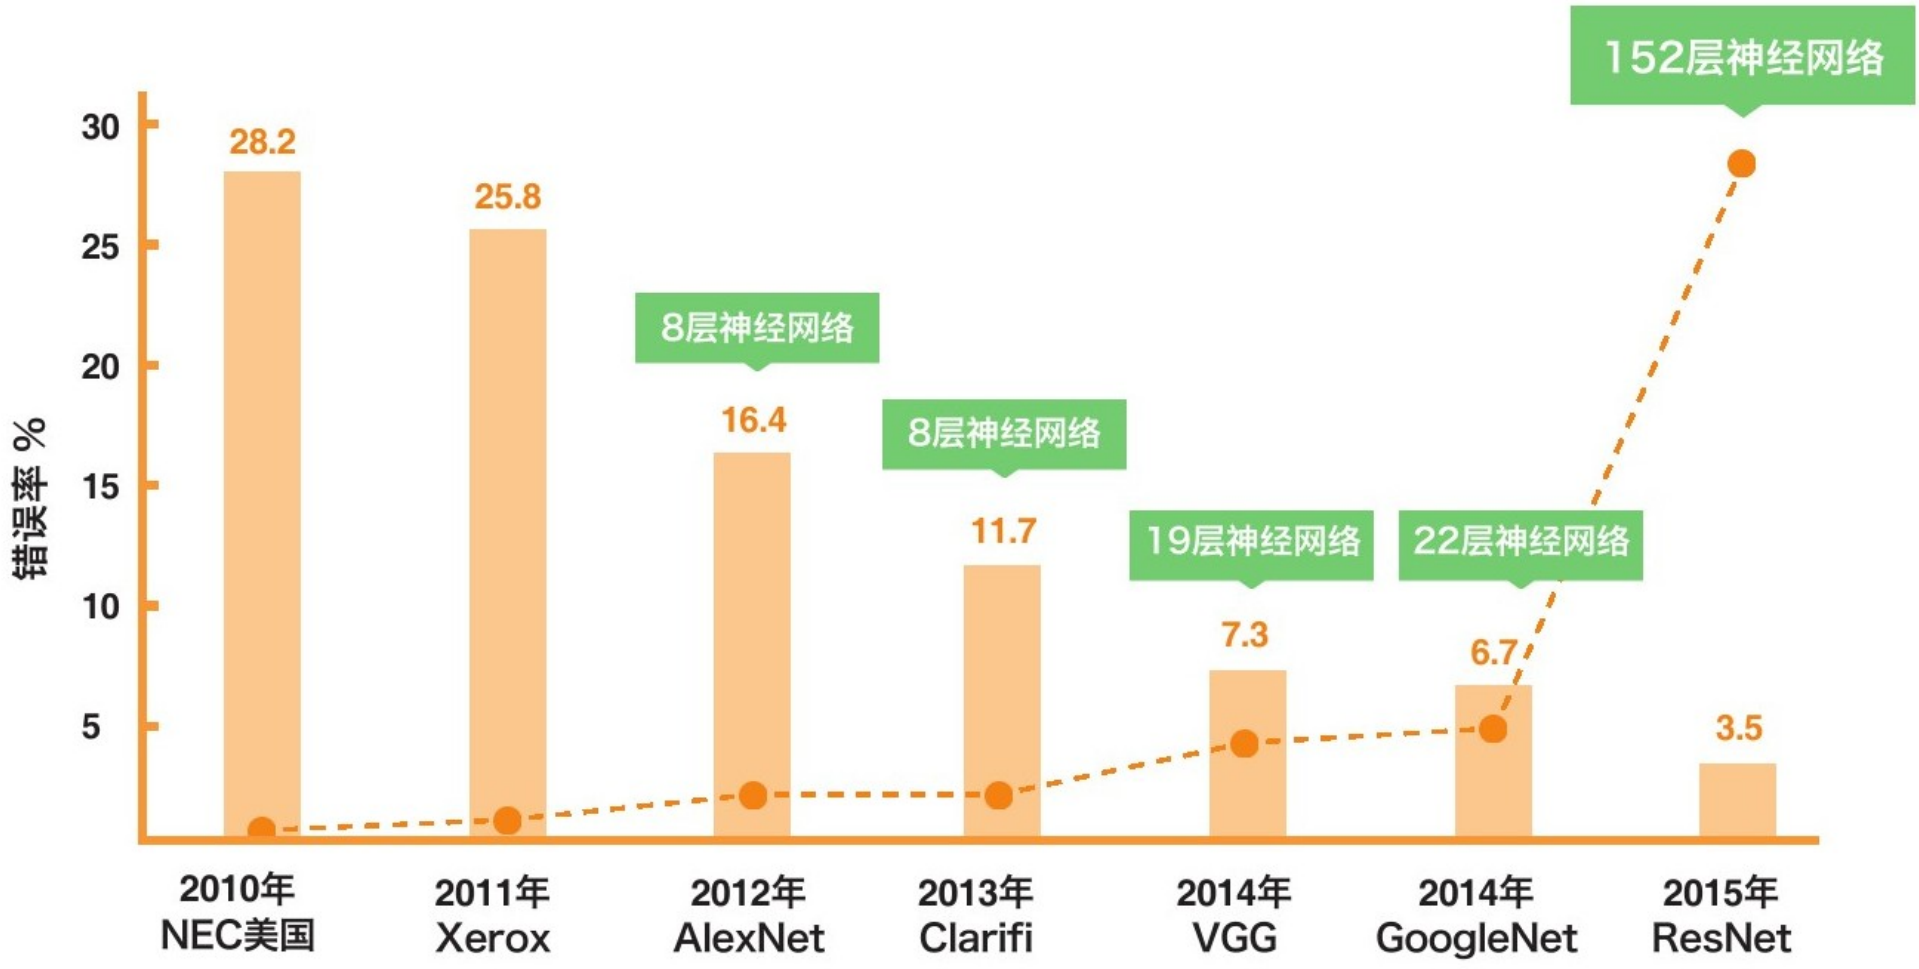

Source: 李開復, 王詠剛, 人工智慧來了, 天下文化, 2017

In ImageNet, there are 14,197,122 images, indicating massive training data are needed.

Note that giant companies such as Google and Facebook are very advantageous in obtaining massive data via search and social network interactions.

Microsoft ResNet has 152 layers in the neural network, indicating huge computations are involved.

To address the complexity issues, parallel/distributed processing, Graphics Processing Units (GPUs) and cloud computing are involved.

To handle the big data issues, information exchange and storage techniques particularly among networks, are required.

References:

1. C. C. Aggarwal, *Neural Networks and Deep Learning : A Textbook*, Springer, 2018
2. S. Haykin, *Neural Networks and Learning Machines*, Prentice Hall, 2009