Use of Artificial Intelligence for Geological and Geotechnical Application

Louis WONG

Associate Professor

Department of Earth Sciences

香 港 大 學
THE UNIVERSITY OF HONG KONG

5th March 2021 (Fri)

Some references

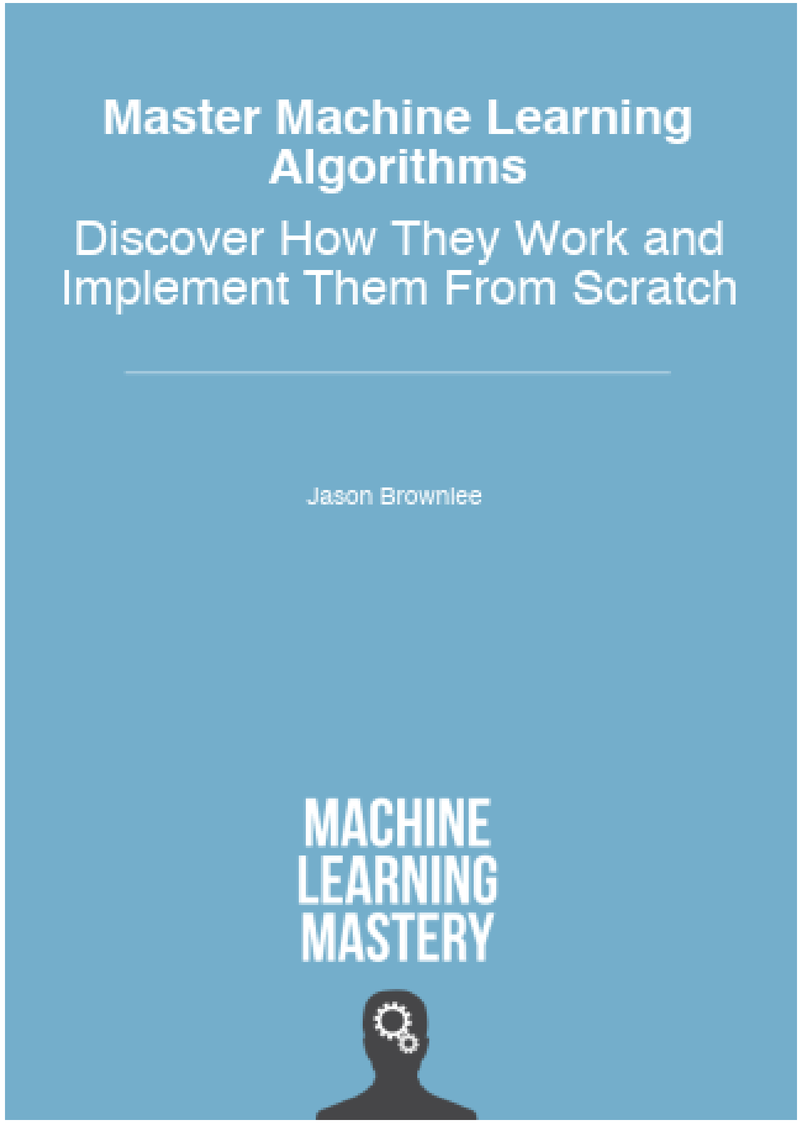

DeepLearning.AI Programs The Batch Events Company Contact

Our Mission

DeepLearning.AI is an education technology company that is empowering the global workforce to build an AI-powered future through world-class education, hands-on training, and a collaborative community.

Our Story

DeepLearning.AI was founded in 2017 by machine learning and education pioneer Andrew Ng to fill a need for world-class AI education. DeepLearning.AI has created high-quality AI programs on Coursera that have gained an extensive global following. By providing a platform for education and fostering a tight-knit community, DeepLearning.AI has become the pathway for anyone looking to build an AI career. Resolve

Are these questions in your mind?

1. What is the difference between Artificial Intelligence (AI), Machine Learning (ML) and Deep Learning (DL)?
2. What is the relationship between AI model and AI algorithm?
3. Is unsupervised ML better than supervised ML?
4. What is a parametric ML algorithm and how is it different from a nonparametric ML algorithm?
5. What is the difference between underfitting and overfitting of data?
6. Is large amount of data equivalent to big data?
7. Will AI be a solution to my day-to-day geo-related work?

Major Items

- Terminologies
- (Classic) Machine learning basics
- Deep learning basics
- Geological and geotechnical engineering applications

The big family of "AI"

(many are subsets of AI)

1. Artificial Intelligence (AI) = learning, reasoning, understanding, grasping truths, seeing relationships, considering meanings, separating fact from belief.

2. Machine learning (ML) = larger concept containing deep learning and classic ML.

3. However, people may sometimes use ML and classic ML interchangeably. Need to judge by context.

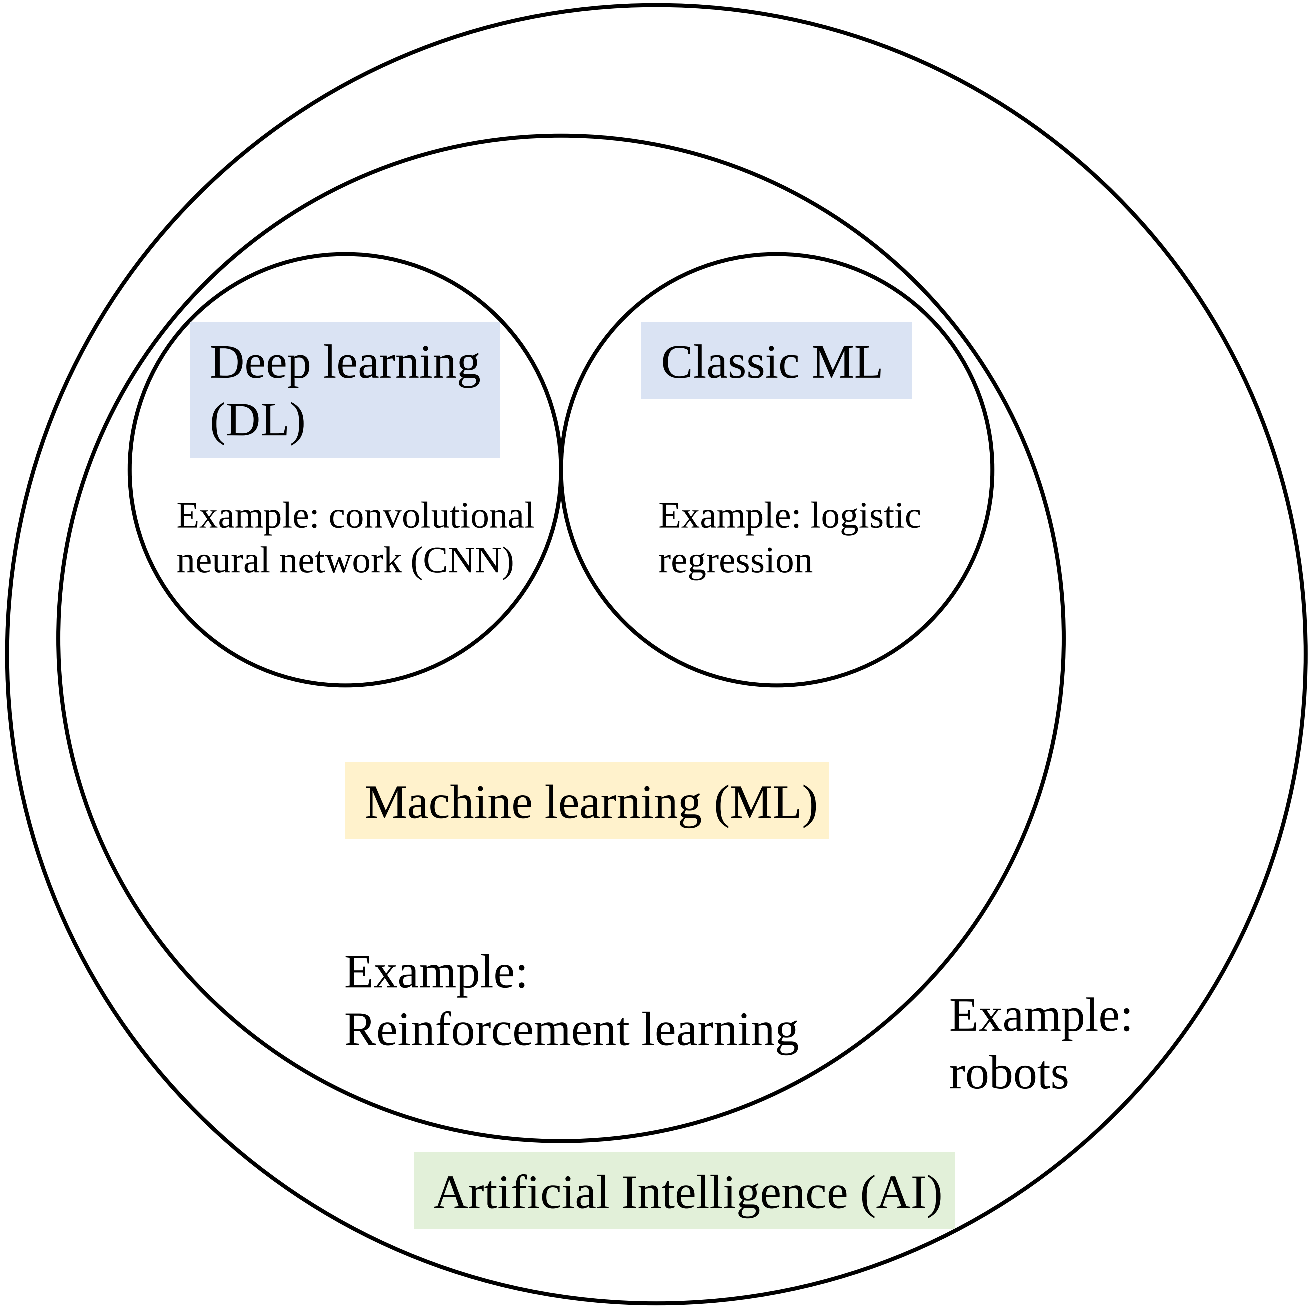

What do you think?

- My interests are in geotechnical engineering and retaining wall design
- My interests are in soil mechanics, rock mechanics, foundation engineering
- Geotechnical engineering include the above sub-disciplines!
- Artificial Intelligence (AI) vs Machine learning (ML)

How do we call the elements?

	A	B	C	D
1		Column 1	Column 2	Column 3
2	Row 1	2.2	2.3	1
3	Row 2	2.3	2.6	0
4	Row 3	2.1	2	1

Data Terminology

	A	B	C	D
1		Attribute 1	Attribute 2	Output Attribute
2	Instance 1	2.2	2.3	1
3	Instance 2	2.3	2.6	0
4	Instance 3	2.1	2	1

Computer Science Perspective

OutputAttribute = Program(InputAttributes)

AI Models vs AI Algorithms

- This can be confusing as both algorithm and model can be used interchangeably. Not really!!!

- Model as the specific representation learned from data and the algorithm as the process for learning it.

Model = Algorithm (Data)

- e.g., a *decision tree* or *a set of coefficients* are a model
- e.g., *C5.0* and *Least Squares Linear Regression* are algorithms to learn those respective models.

Major Items

- Terminologies ☺

- (Classic) Machine learning basics

- Deep learning basics

- Geological and geotechnical engineering applications

Machine Learning (ML)

- The goal is to create a simulation of human learning so that an application can adapt to uncertain or unexpected conditions.

- To perform this task, machine learning relies on algorithms to analyze huge datasets.

- The key task is to LEARN a function.

Learning a Function

- Machine learning algorithms are described as learning a target function (f) that best maps (verb) input variables (X) to an output variable (Y).

$$Y = f(X)$$

- This is a general learning task where we would like to make predictions in the future (Y) given new examples of input variables (X).
- We don't know what the function (f) looks like or it's form. If we did, we would use it directly and we would not need to learn it from data using machine learning algorithms.
- In geotechnical engineering language, it's the physical law. For example, you don't know $\sigma' = \sigma - u$.

Learning a Function

- There is also error (e) that is independent of the input data (X).

$$Y = f(X) + e$$

- This error might be error such as not having enough attributes to sufficiently characterize the best mapping from X to Y. This error is called irreducible error because no matter how good we get at estimating the target function (f), we cannot reduce this error.
- Learning a function from data is a difficult problem. Again, it's because we don't know the exact physical law

Parametric and Nonparametric Machine Learning Algorithms

Parametric Machine Learning Algorithms

- A learning model that summarizes data with a set of parameters of fixed size (independent of the number of training examples) is called a parametric model.
- No matter how much data you throw at a parametric model, it won't change its mind about how many parameters it needs.

- The algorithms involve two steps:
  1. Select a form for the function.
  2. Learn the coefficients for the function from the training data.

- An easy to understand functional form for the mapping function is a line, as is used in linear regression:

$$B0 + B1 \times X1 + B2 \times X2 = 0$$

Parametric Machine Learning Algorithms:

Benefits

- Simpler
- Speed
- Less Data

Limitations

- Constrained: By choosing a functional form these methods are highly constrained to the specified form.
- Limited Complexity: The methods are more suited to simpler problems.
- Poor Fit: In practice the methods are unlikely to match the underlying mapping function.

Nonparametric Machine Learning Algorithms

- Algorithms that do not make strong assumptions about the form of the mapping function.
- They are free to learn any functional form from the training data.
- Nonparametric methods are good when you have a lot of data and no prior knowledge, and when you don’t want to worry too much about choosing just the right features. (no clue about the physical law?)
- Nonparametric methods seek to best fit the training data in constructing the mapping function, whilst maintaining some ability to generalize to unseen data. As such, they are able to fit a large number of functional forms.
- An easy to understand nonparametric model is the k-nearest neighbors algorithm that makes predictions based on the k most similar training patterns for a new data instance, e.g. joint set clustering in stereographic projection.

Nonparametric Machine Learning Algorithms

Benefits

- Flexibility
- Power
- Performance

Limitations

- More data: Require a lot more training data to estimate the mapping function.
- Slower: A lot slower to train as they often have far more parameters to train.
- Overfitting: More of a risk to overfit the training data and it is harder to explain why specific predictions are made.

Supervised, Unsupervised *vs* Semi-Supervised Learning

- Supervised learning is where you have input variables (X) and an output variable (Y) and you use an algorithm to learn the mapping function from the input to the output.

  Y = f(X)

- The goal is to approximate the mapping function so well that when you have new input data (X) that you can predict the output variables (Y) for that data.
- It is called supervised learning because the process of an algorithm learning from the training dataset can be thought of as a teacher supervising the learning process.
- We know the correct answers, the algorithm iteratively makes predictions on the training data and is corrected by the teacher.
- Learning stops when the algorithm achieves an acceptable level of performance (not necessarily 100%)

Supervised learning problems can be further grouped into classification and regression problems

- Classification: A classification problem is when the output variable is a category, such as red or blue or disease and no disease.
- Regression: A regression problem is when the output variable is a real value, such as dollars or weight.
- Some common types of problems built on top of classification and regression include recommendation and time series prediction respectively.

Some popular examples of supervised machine learning algorithms are:

- Linear regression for regression problems.
- Random forest for classification and regression problems.
- Support vector machines for classification problems.

Unsupervised Machine Learning

- Unsupervised learning is where you only have input data (X) and no corresponding output variables in the training data.
- The goal for unsupervised learning is to model the underlying structure or distribution in the data in order to learn more about the data.
- These are called unsupervised learning because, unlike supervised learning, there is no correct answers and there is no teacher.
- Algorithms are left to their own devises to discover and present the interesting structure in the data problems.

Unsupervised learning problems can be further grouped into clustering and association problems.

- Clustering: A clustering problem is where you want to discover the inherent groupings in the data, such as grouping customers by purchasing behavior.
- Association: An association rule learning problem is where you want to discover rules that describe large portions of your data, such as people that buy A also tend to buy B.

Some popular examples of unsupervised learning algorithms are:

- k-means for clustering problems
- Apriori algorithm for association rule learning

Summary

Supervised: All data is labeled and the algorithms learn to predict the output from the input data.

Unsupervised: All data is unlabeled and the algorithms learn to inherent structure from the input data.

Semi-supervised: Some data is labeled but most of it is unlabeled and a mixture of supervised and unsupervised techniques can be used.

Hand-crafted features (structured data)

If features can be structured in a table, they are regarded as hand-crafted features or structured data

Task: using classic ML to predict the price of a house

Feature 1 Size (m^2)	Feature 2 Location	Feature 3 A school nearby?	…	Feature n Bedroom#	Price (million dollars)
50	Hong Kong Island	No	…	2	10.5
60	Kowloon	Yes	…	3	11.2
…	…	…	…	…	
100	Kowloon	Yes	…	4	?

- Hand-crafted features are extracted based on domain knowledge – usually done manually
- Sometimes, what features should be extracted is quite difficult to know.
- Feature engineering is usually needed.

Sub-classes of classic ML

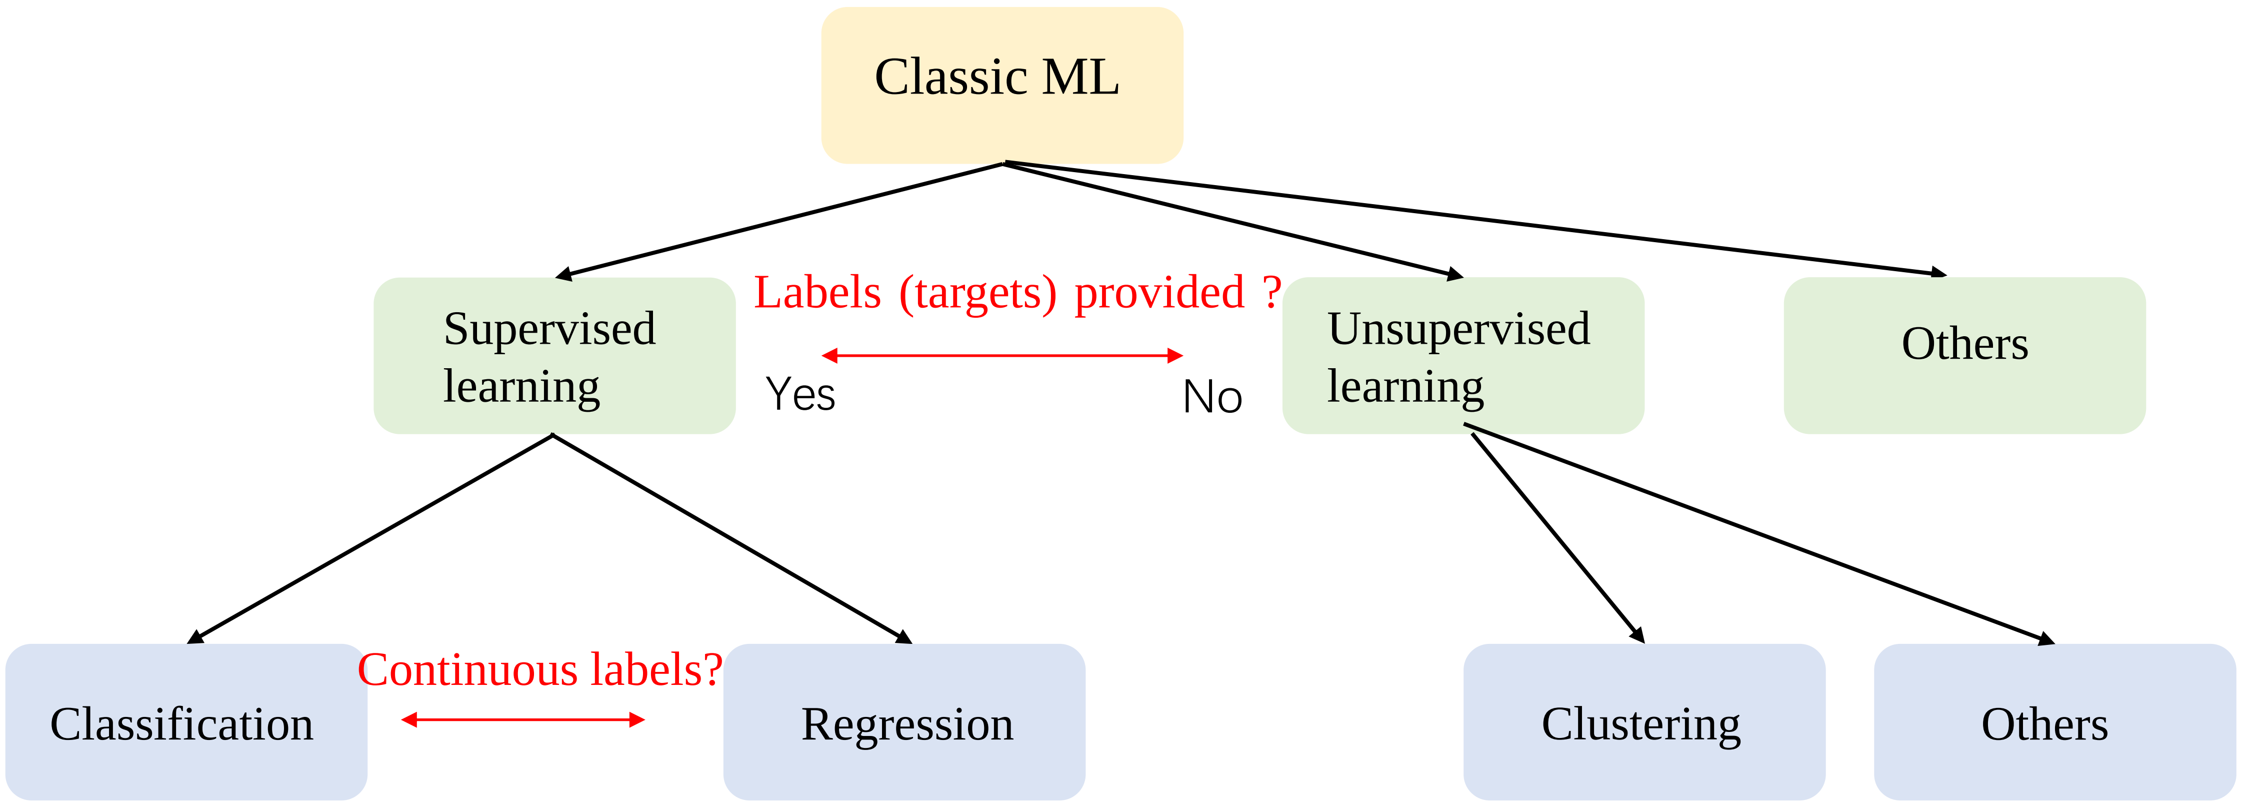

!!! Using supervised learning to help illustrate basic concepts of classic machine learning

How machine learning "learns"

Task: using classic ML to predict the price of a house

Feature 1 **Size (m^2)**	**Feature 2** **Location**	**Feature 3** **A school nearby?**	**…**	**Feature n** **Bedroom#**	**Price (million dollars)**
50	Hong Kong Island	No	…	2	10.5
60	Kowloon	Yes	…	3	11.2
…	…	…	…	…	
100	Kowloon	Yes	…	4	?

1. The learning occurs as a result of analyzing ever-increasing amounts of data. No data, no learning.
2. The basic algorithms don't change. Need to try many algorithms.
3. But the algorithm's internal weights and biases used to select a particular answer do.

Training, validation, and testing data

➢ Training data

Training data should be representative. It must truly represent the problem domain

➢ Validation data

To avoid possible underfitting and overfitting

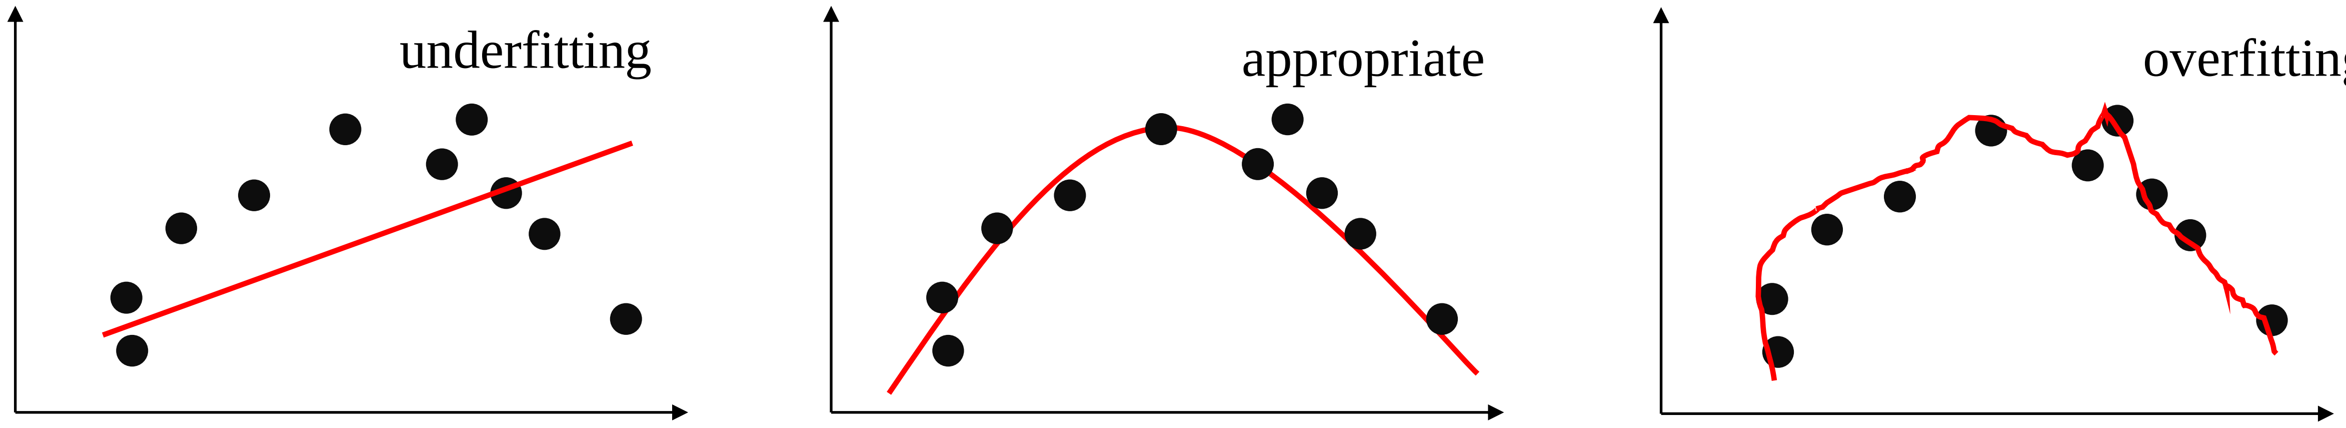

➢ Testing data

To test how well the algorithm learned

Understanding the meaning of weights

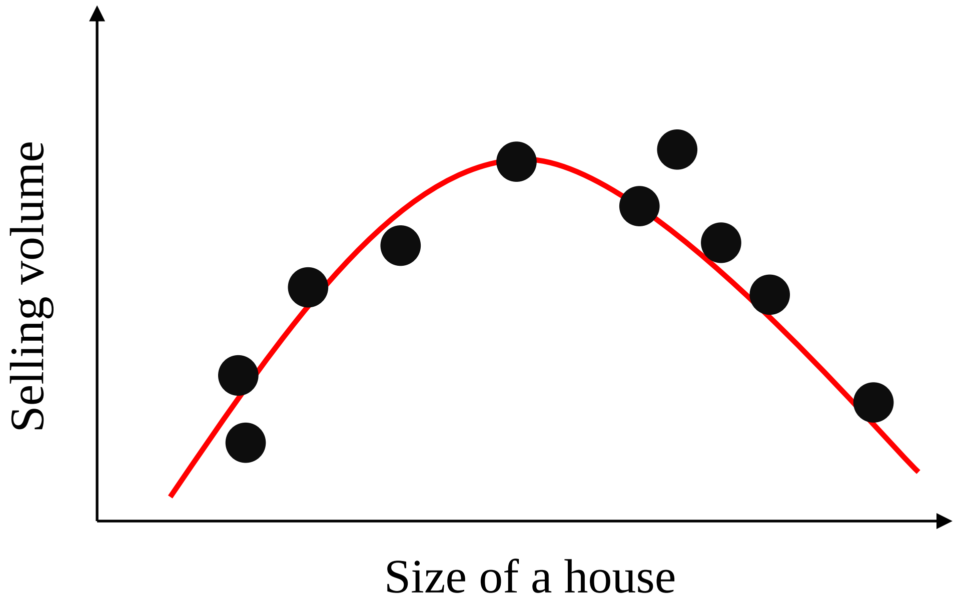

One-dimensional issue in a two-dimensional system, only one parameter

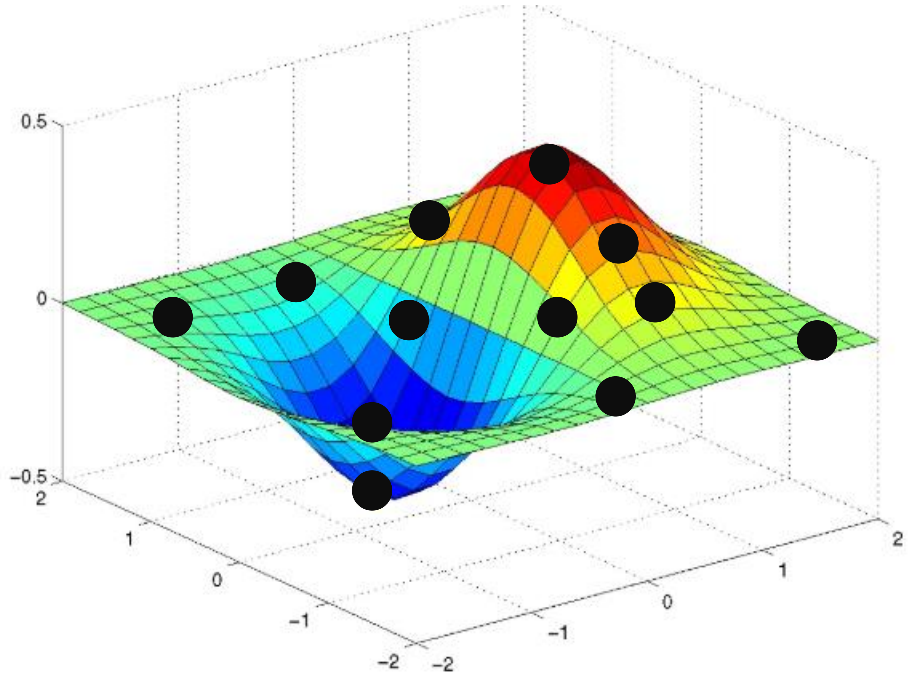

Two-dimensional issue in a three-dimensional system, two parameters

Q: What about a higher-dimensional issue?

A: Just imagine a hyperplane to fit the samples

Steps for classic ML

1. Structure the data in a table

Feature 1 **Size (m^2)**	**Feature 2** **Location**	**Feature 3** **A school nearby?**	**…**	**Feature n** **Bedroom#**	**Price (million dollars)**
50	Hong Kong Island	No	…	2	10.5
60	Kowloon	Yes	…	3	11.2
…	…	…	…	…	
100	Kowloon	Yes	…	4	?

2. Select a machine learning algorithm

3. Train the algorithm using training data

4. Check whether underfitting or overfitting occurs using validation data

5. Test the trained algorithm using testing data

6. Save the model (weights) – more elaboration of weights in the coming slides

Major Items

- Terminologies ☺
- (Classic) Machine learning basics☺
- Deep learning basics
- Geological and geotechnical engineering applications

Deep learning? Or Neural networks?

Artificial neural network

From Wikipedia, the free encyclopedia

Artificial neural networks (**ANN**) or **connectionist systems** are computing systems vaguely inspired by the biological neural networks that constitute animal brains.[1] Such systems "learn" to perform tasks by considering examples, generally without being programmed with task-specific rules. For example, in image recognition, they might learn to identify images that contain cats by analyzing example images that have been manually labeled as "cat" or "no cat" and using the results to identify cats in other images. They do this without any prior knowledge of cats, for example, that they have fur, tails, whiskers and cat-like faces. Instead, they automatically generate identifying characteristics from the examples that they process.

An ANN is based on a collection of connected units or nodes called artificial neurons, which loosely model the neurons in a biological brain. Each connection, like the synapses in a biological brain, can transmit a signal to other neurons. An artificial neuron that receives a signal then processes it and can signal neurons connected to it.

In ANN implementations, the "signal" at a connection is a real number, and the output of each neuron is computed by some non-linear function of the sum of its inputs. The connections are called *edges*. Neurons and edges typically have a *weight* that adjusts as learning proceeds. The weight increases or decreases the strength of the signal at a connection. Neurons may have a threshold such that a signal is sent only if the aggregate signal crosses that threshold. Typically, neurons are aggregated into layers. Different layers may perform different transformations on their inputs. Signals travel from the first layer (the input layer), to the last layer (the output layer), possibly after traversing the layers multiple times.

The original goal of the ANN approach was to solve problems in the same way that a human brain would. But over time, attention moved to performing specific tasks, leading to deviations from biology. ANNs have been used on a variety of tasks, including computer vision, speech recognition, machine translation, social network filtering, playing board and video games, medical diagnosis, and even in activities that have traditionally been considered as reserved to humans, like painting.[2]

Deep learning

From Wikipedia, the free encyclopedia

For deep versus shallow learning in educational psychology, see Student approaches to learning. For more information, see Artificial neural network.

Deep learning (also known as **deep structured learning**) is part of a broader family of machine learning methods based on artificial neural networks with representation learning. Learning can be supervised, semi-supervised or unsupervised.[1][2][3]

Deep learning architectures such as deep neural networks, deep belief networks, recurrent neural networks and convolutional neural networks have been applied to fields including computer vision, speech recognition, natural language processing, audio recognition, social network filtering, machine translation, bioinformatics, drug design, medical image analysis, material inspection and board game programs, where they have produced results comparable to and in some cases surpassing human expert performance.[4][5][6]

Artificial neural networks (ANNs) were inspired by information processing and distributed communication nodes in biological systems. ANNs have various differences from biological brains. Specifically, neural networks tend to be static and symbolic, while the biological brain of most living organisms is dynamic (plastic) and analog.[7][8][9]

The adjective "deep" in deep learning comes from the use of multiple layers in the network. Early work showed that a linear perceptron cannot be a universal classifier, and then that a network with a nonpolynomial activation function with one hidden layer of unbounded width can on the other hand so be. Deep learning is a modern variation which is concerned with an unbounded number of layers of bounded size, which permits practical application and optimized implementation, while retaining theoretical universality under mild conditions. In deep learning the layers are also permitted to be heterogeneous and to deviate widely from biologically informed connectionist models, for the sake of efficiency, trainability and understandability, whence the "structured" part.

Don't be overburdened. They are the same thing in today's semantics

"Using the name 'deep learning' is for branding"

Andrew Ng

American businessman

Andrew Yan-Tak Ng is a British-born Chinese-American businessman, computer scientist, investor, and writer. He is focusing on machine learning and AI. Wikipedia

Different workflow

Classic ML

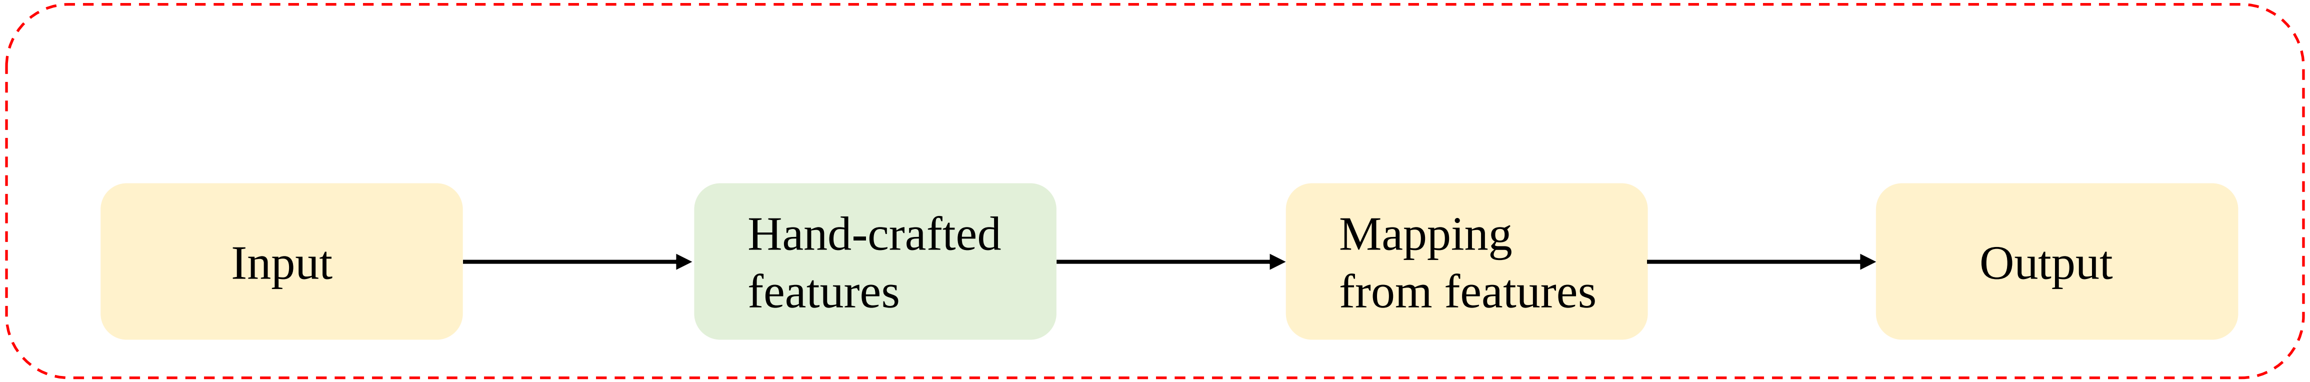

- ➢ No need to understand the details of the two flowcharts.
- ➢ Just have an intuition that classic ML and DL have similar flowcharts but different representation of features

Deep learning

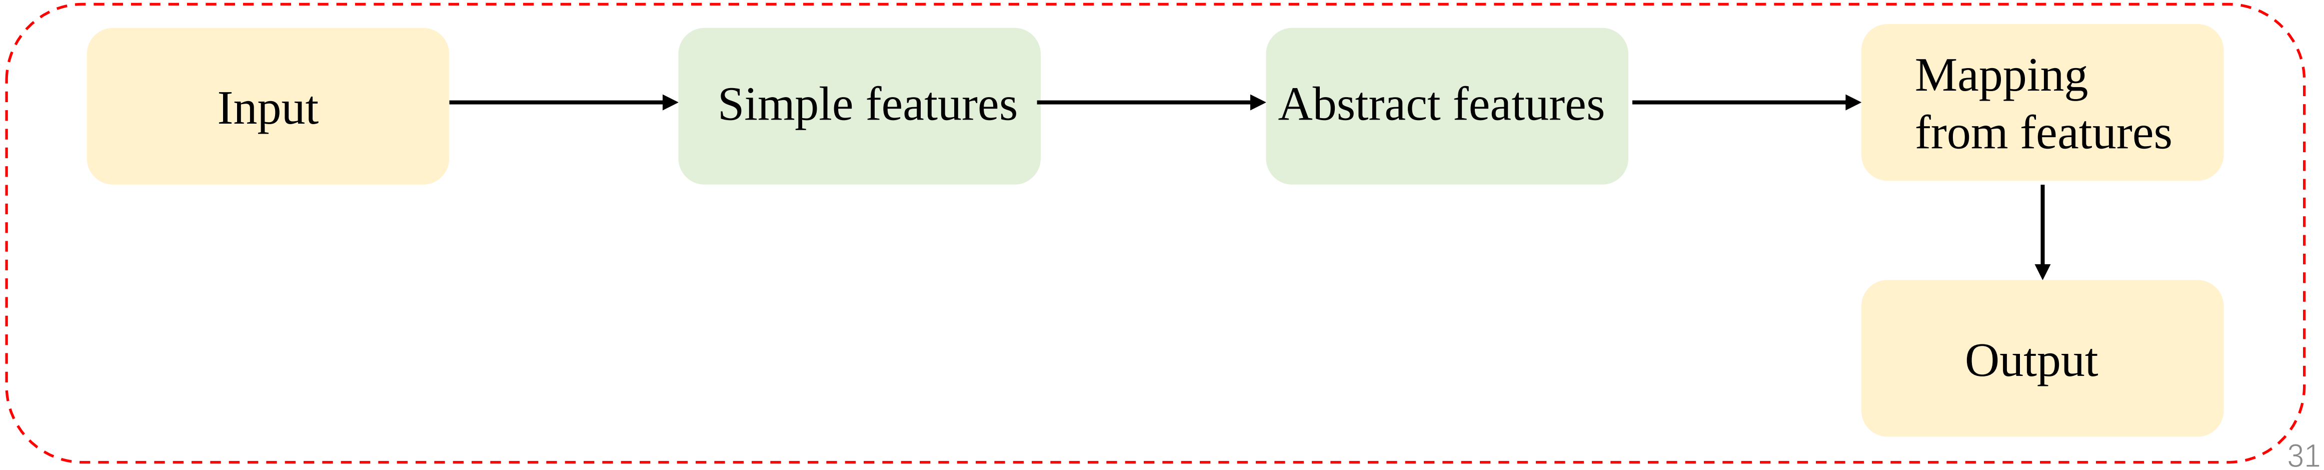

Features for DL (unstructured data)

1-D: i.e. Audio

2-D: i.e. Image

3-D: i.e. Point cloud

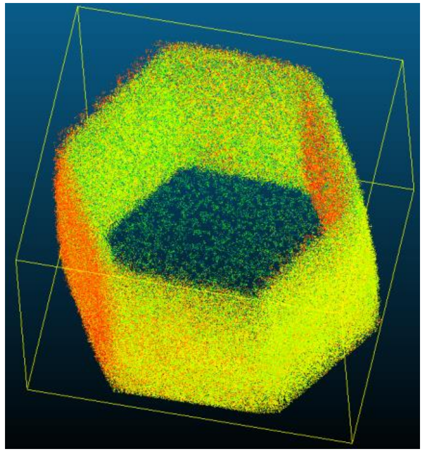

- No need to manually extract features from inputs.
- Algorithms will automatically learn abstract features from simple features, such as RGB values, gray-scale values, etc.
- Then using abstract features to do prediction.

- However, the meaning of abstract features is quite non-interpretable.
- Why some abstract features work is kind of “black-box”

Brief Summary

- Classic ML and DL are two different AI disciplines but are sometimes used interchangeably. (LW does not like such situation)

- The biggest differences of them are the representation of features

Why deep learning?

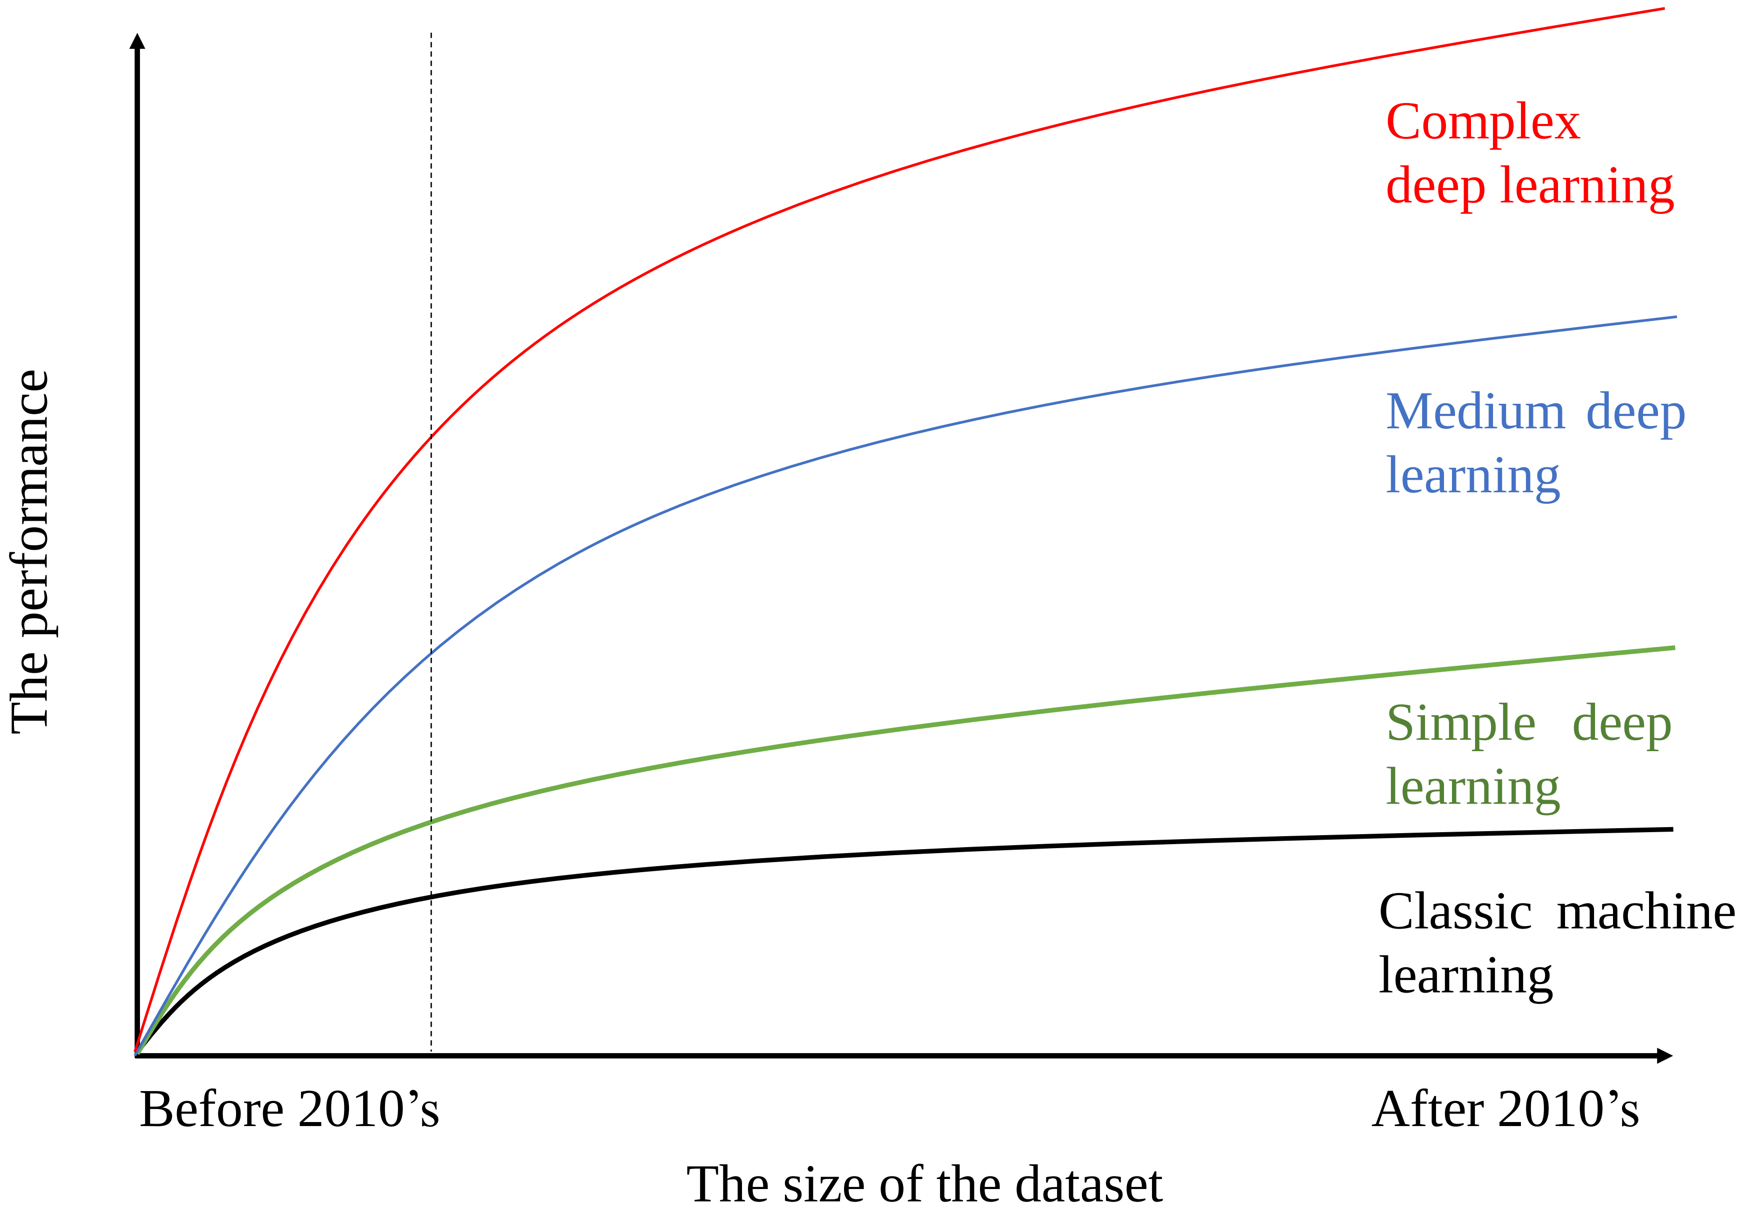

Classic ML is unable to address issues with millions of parameters

Neurons and hidden layers

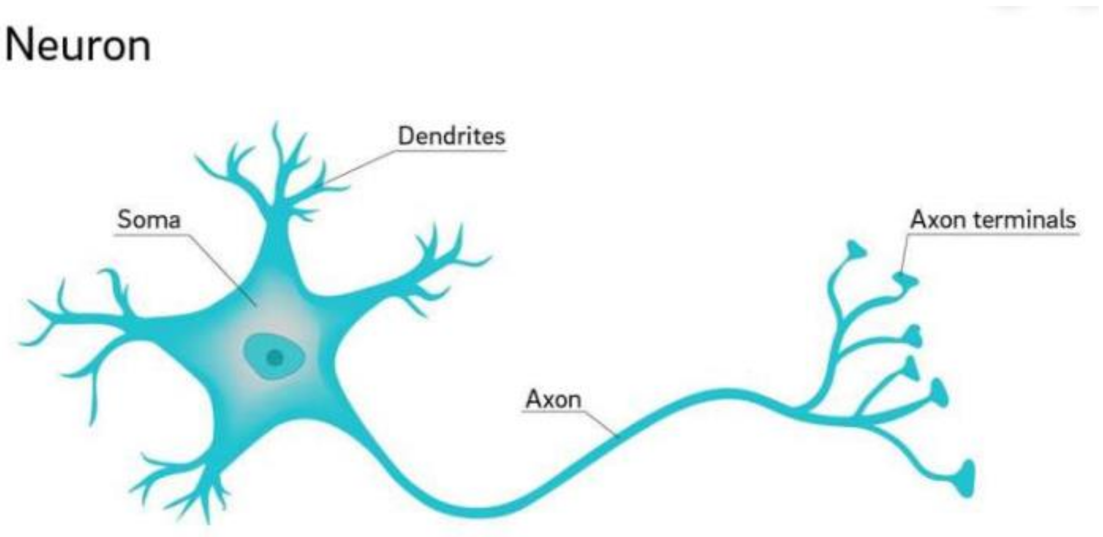

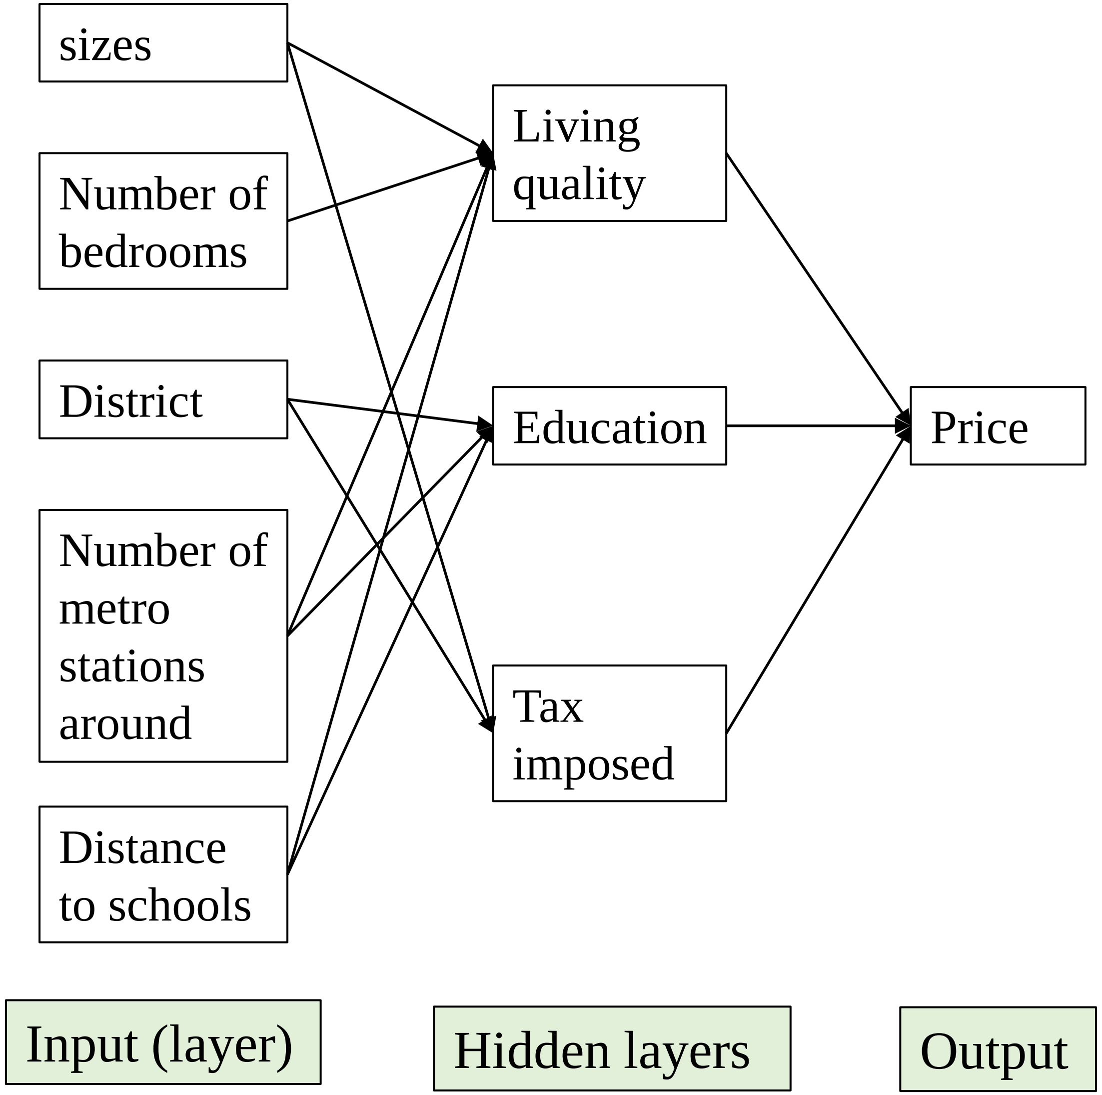

- Neural networks have different **layers**, each one having its own **weights**
- The neural network segregates computations by **layers**
- Weights represent **the strength of the connection** between neurons in the network

Forward propagation (FP) and back propagation (BP)

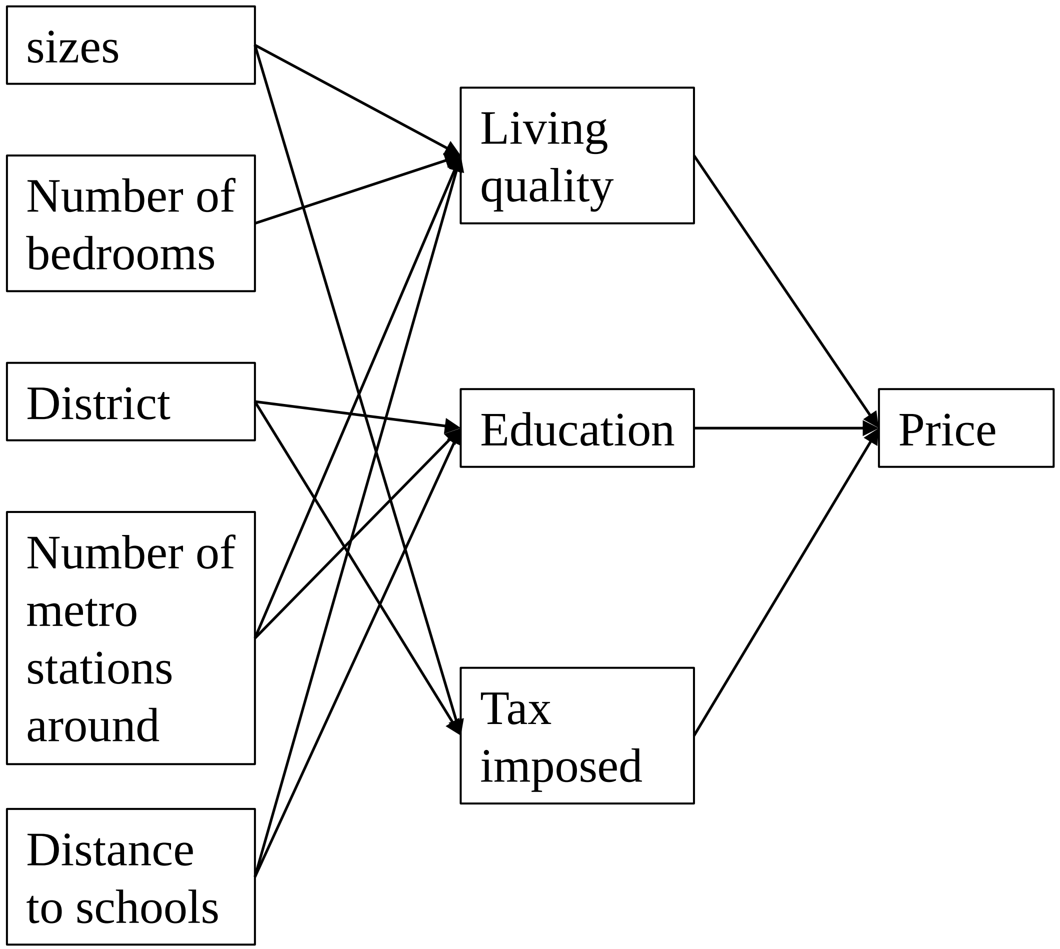

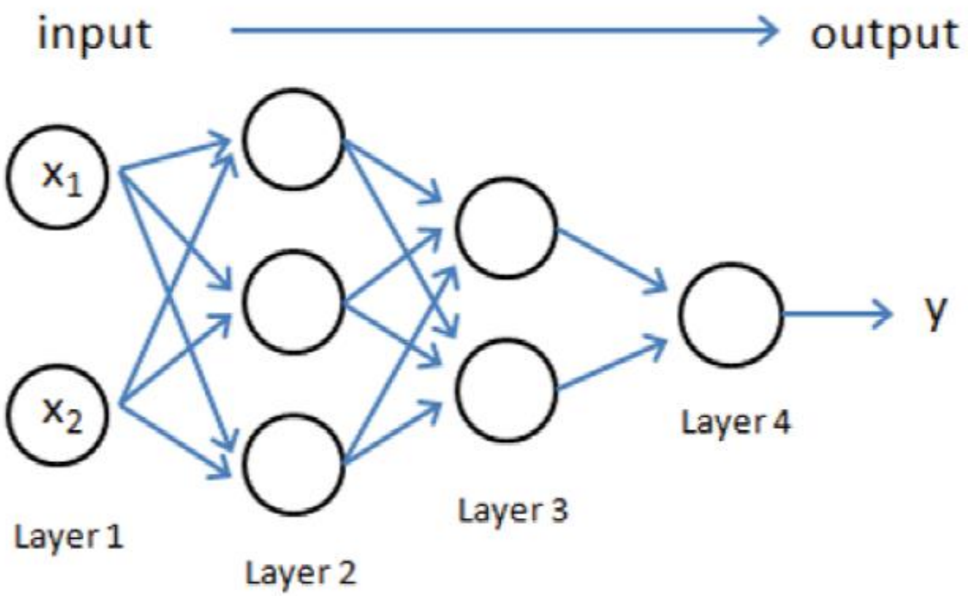

Principle of FP

- The principle of back propagation (BP) is not in the scope of this presentation.
- Just know BP is used for seeking the optimum weights (coefficients)

Simple features

How a computer read an image? ---Using RGB values

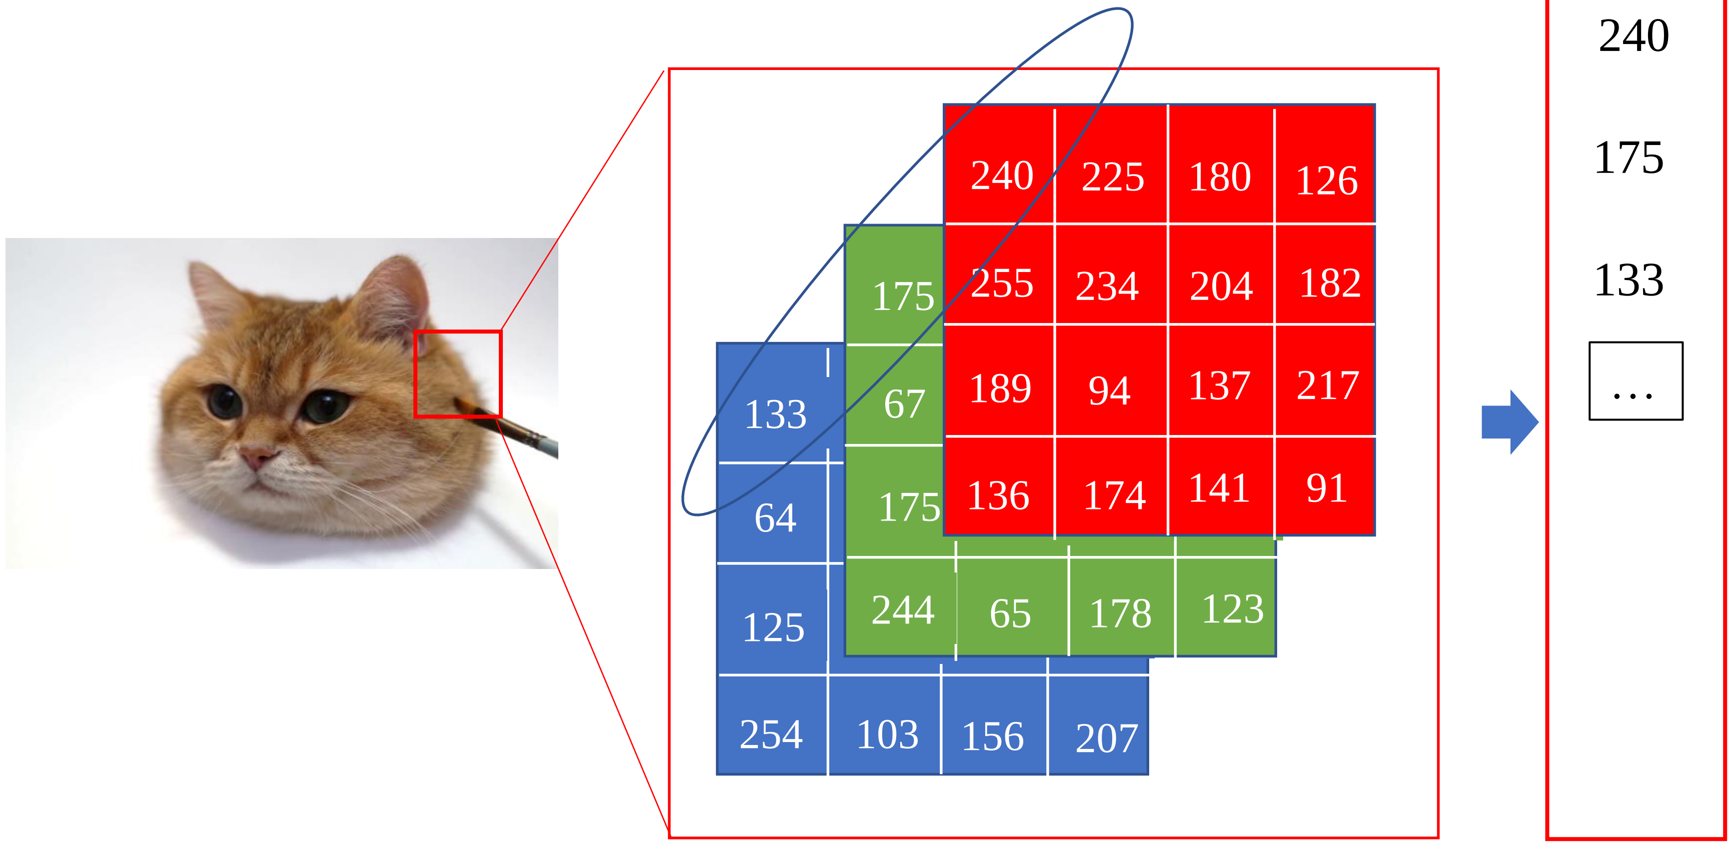

Convert images to numeric data

Brief Summary

- DL is a data-driven AI discipline

- Brief introduction of

1. neurons and hidden layers
2. forward propagation and back propagation
3. simple features

What happened in a neuron

To examine the operations that neural networks execute in detail, inputs and outputs in different ways:

"a":

- The result stored in a unit in the neural network after being processed by the activation function (called *g*).
- This is the final output that is sent further along the network

"z":

- The multiplication between ***"a"*** and the weights.
- ***"z"*** represents the signal going through the connections, analogous to water in pipes that flows at a higher or lower pressure depending on the pipe diameter.
- In the same way, the values received from the previous layer get higher or lower values because of the connection weights used to transmit them.

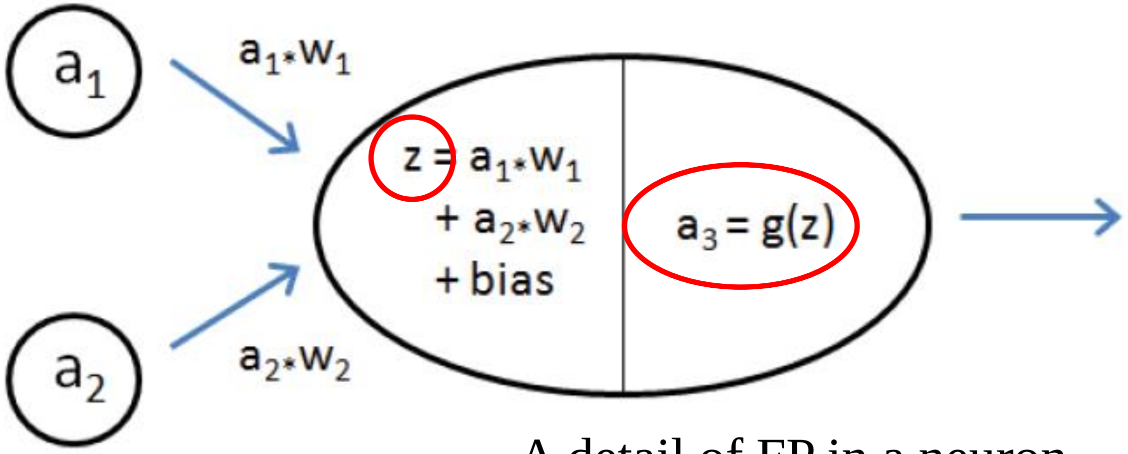

A detail of FP in a neuron

Produce a prediction (4-layer example)

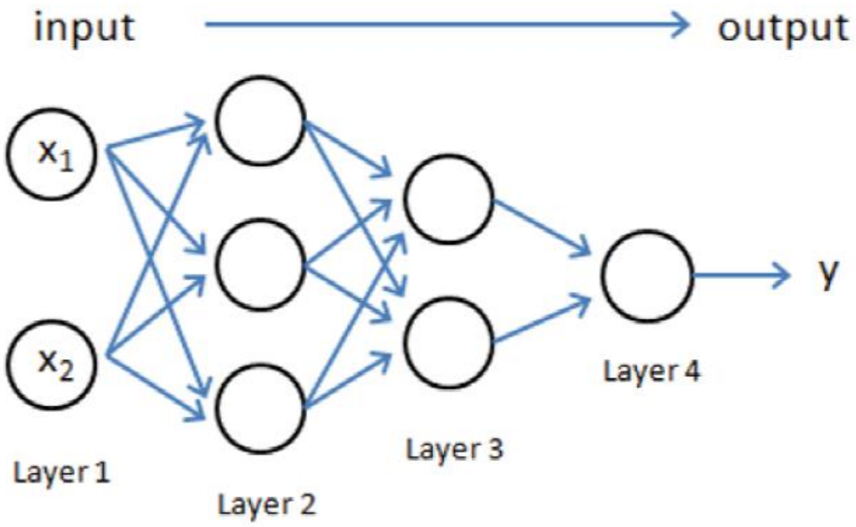

Seven steps used to produce a prediction in a neural network made of four layers (Top right figure)

1. The first layer (notice the superscript 1 on a) loads the value of each feature in a different unit:

$$a^{(1)} = X \quad \textit{As a matrix}$$

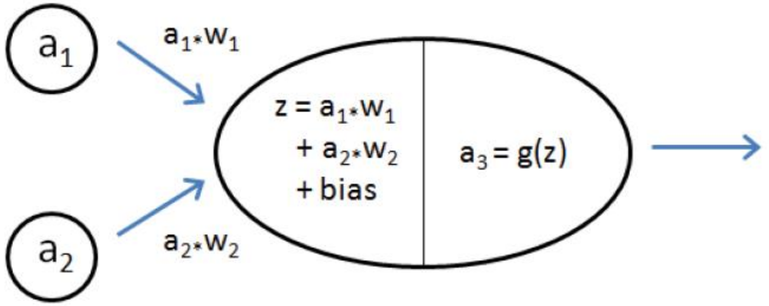

2. The weights of the connections bridging the input layer with the second layer are multiplied by the values of the units in the first layer. A matrix multiplication weights (w_1 w_2) and sums the inputs for the second layer together.

$$z^{(2)} = W^{(1)} a^{(1)}$$

Produce a prediction

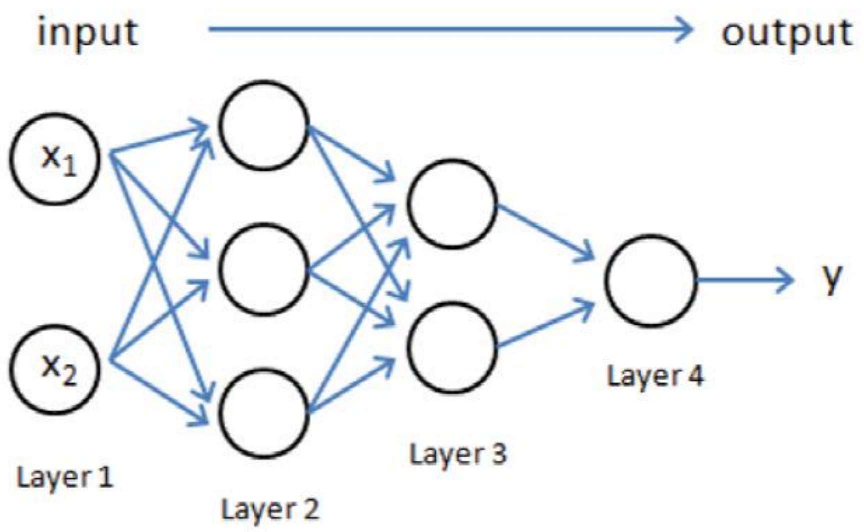

3. The algorithm adds a bias constant to layer two before running the activation function *g*. The activation function transforms the second layer inputs. The resulting values are ready to pass to the connections.

$$a^{(2)} = g(z^{(2)} + bias^{(2)})$$

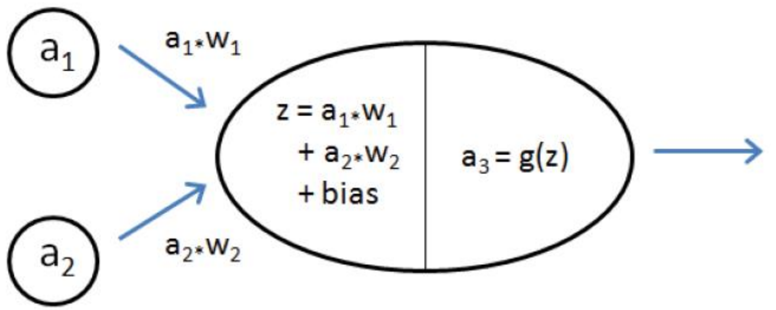

4. The third layer connections weigh and sum the outputs of layer two.

$$z^{(3)} = W^{(2)} a^{(2)}$$

5. The algorithm adds a bias constant to layer three before running the activation function. The activation function transforms the layer-three inputs.

$$a^{(3)} = g(z^{(3)} + bias^{(3)})$$

Produce a prediction

6. The layer-three outputs are weighted and summed by the connections to the output layer.

$$z^{(4)} = W^{(3)} a^{(3)}$$

7. Finally, the algorithm adds a bias constant to layer four before running the activation function. The output units receive their inputs and transform the input using the activation function. After this final transformation, the output units are ready to release the resulting predictions of the neural network.

$$a^{(4)} = g(z^{(4)} + bias^{(4)})$$

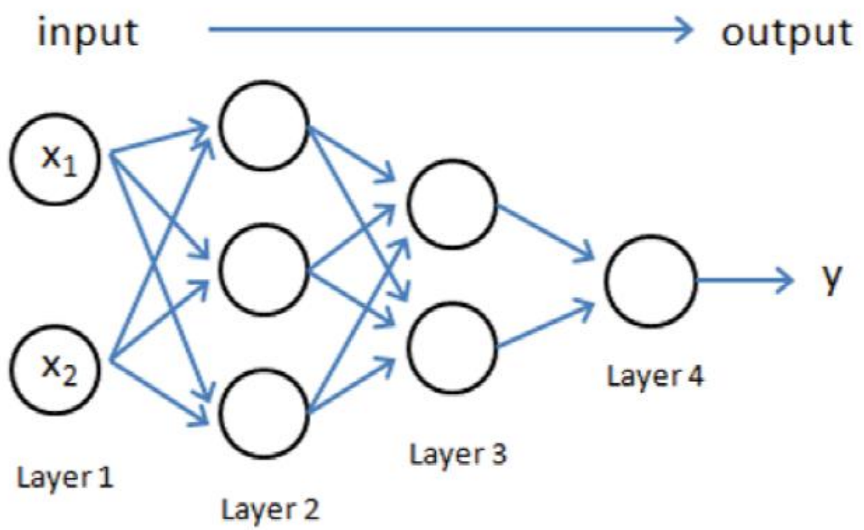

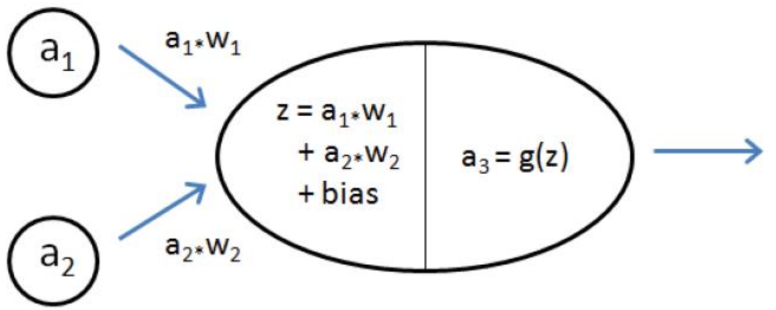

Major Items

- Terminologies ☺
- (Classic) Machine learning basics☺
- Deep learning basics☺
- Geological and geotechnical engineering applications

Hong Kong

- Large amount of (too much) geological and geotechnical data
- Categories
  - Images
    - corebox photos
    - rock specimens before and after UCS test
    - aerial photos
  - Qualitative description
    - rock color, texture

Borelog

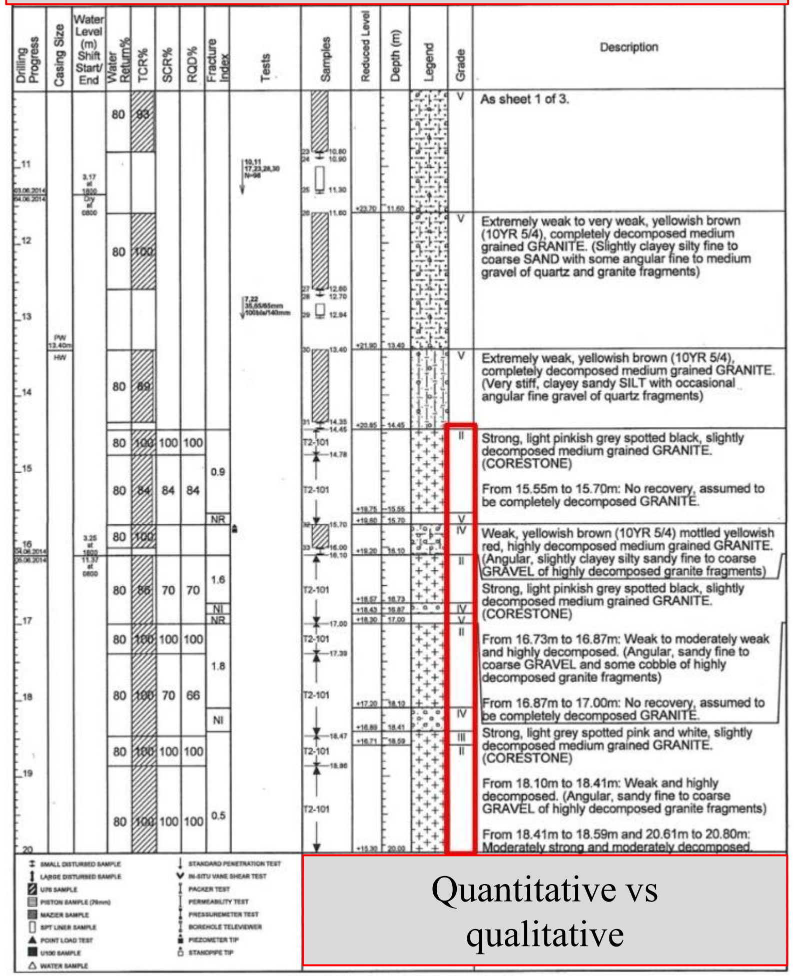

Quantitative vs qualitative

(a) drill hole log record

(b) original core box photo

(c) original core box photo labelled with weathering grades

Hong Kong

- Large amount of (too much) geological and geotechnical data
- Categories
    - Images
        - corebox photos
        - rock specimens before and after UCS test
        - aerial photos
    - Qualitative description
        - rock color, texture
    - Spatial-temporal data
        - Landslide occurrence (location and time)
    - Numerical values
        - Soil lab test results vs soil name
        - Field measurement (settlement)
        - TBM advance rate

Would you call them BIG DATA?

https://towardsdatascience.com/what-is-big-data-lets-answer-this-question-933b94709caf

Big data (https://en.wikipedia.org/wiki/Big_data)

- Is a field that treats ways to analyze, systematically extract information from, or otherwise deal with data sets that are too large or complex to be dealt with by traditional data-processing application software.

- ≠ large amount of data

Geological and Geotechnical engineering applications

1. Predicting ground response (earth-structure) response behind which the physical laws are not easy to be developed in geotechnical system, e.g. pile capacity (interplay of tens of numerical parameters)

2. Automatic crack detection based on computer vision (images)

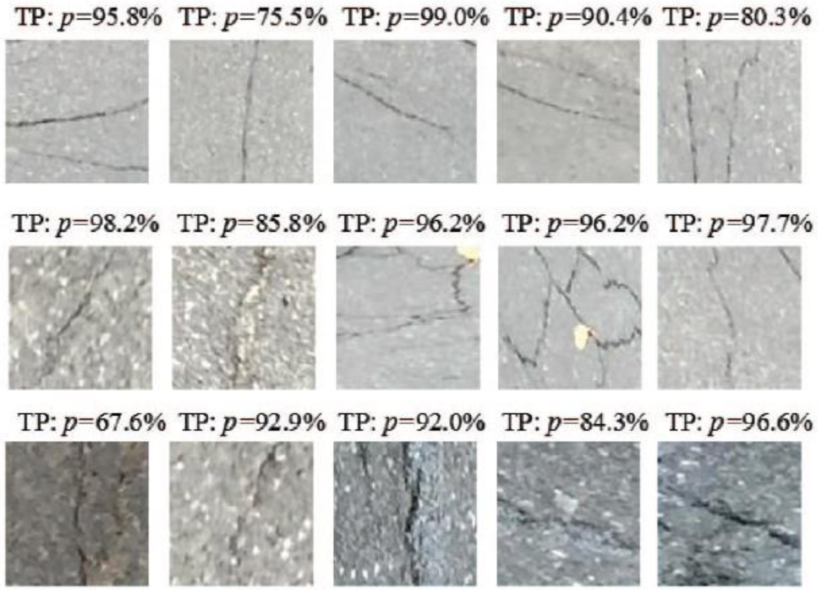

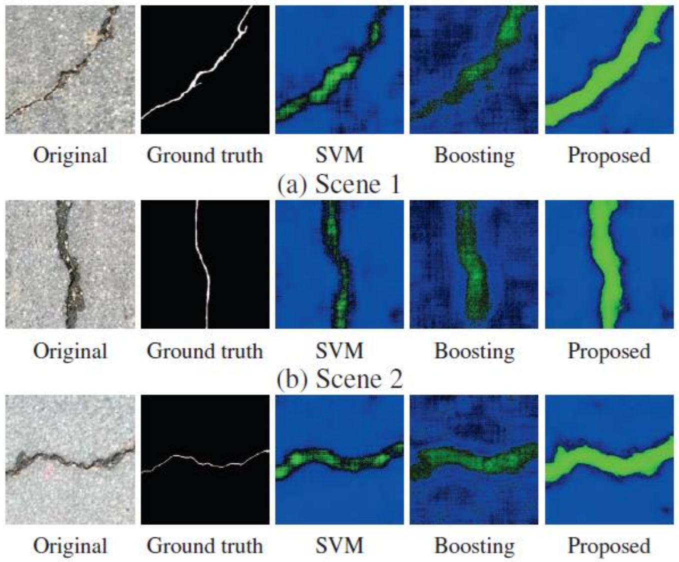

3. Predicting flyrock distance

Springer Link

Original Paper | Published: 16 October 2012

Application of artificial intelligence techniques for predicting the flyrock distance caused by blasting operation

Ebrahim Ghasemi, Hasel Amini, Mohammad Ataei & Reza Khalokakaei

Arabian Journal of Geosciences **7**, 193–202(2014) | Cite this article

Abstract

Flyrock arising from blasting operations is one of the crucial and complex problems in mining industry and its prediction plays an important role in the minimization of related hazards. In past years, various empirical methods were developed for the prediction of flyrock distance using statistical analysis techniques, which have very low predictive capacity. Artificial intelligence (AI) techniques are now being used as alternate statistical techniques. In this paper, two predictive models were developed by using AI techniques to predict flyrock distance in Sungun copper mine of Iran. One of the models employed artificial neural network (ANN), and another, fuzzy logic. The results showed that both models were useful and efficient whereas the fuzzy model exhibited high performance than ANN model for predicting flyrock distance. The performance of the models showed that the AI is a good tool for minimizing the uncertainties in the blasting operations.

4. TBM advance rate

Springer Link

Original Paper | Published: 10 December 2019

Forecasting of TBM advance rate in hard rock condition based on artificial neural network and genetic programming techniques

Jian Zhou, Behnam Yazdani Bejarbaneh, Danial Jahed Armaghani & M. M. Tahir

Bulletin of Engineering Geology and the Environment **79**, 2069–2084(2020) | Cite this article

Abstract

The efficiency of tunnel boring machine (TBM) is regarded as a key factor in successfully undertaking any mechanical tunneling project. In fact, an accurate forecasting of TBM performance, especially in a specified rock mass condition, can minimize capital costs and scheduling for tunnel excavation. This study puts an effort to propose two accurate and practical predictive models of TBM performance via artificial neural network (ANN) and genetic programming (GP) approaches. To set a certain prediction target for the proposed models, the advance rate (*AR*) of TBM is considered as its performance metric. For modeling purpose, a large experimental database containing 1286 data sets was set up as the result of conducting site investigation operations for a tunneling project in Malaysia, called the Pahang–Selangor Raw Water Transfer Tunnel and performing a number of laboratory tests on the collected rock samples. To design the desired intelligent models of *AR* based on the training and test patterns, a mix of rock and machine characteristics with the most influence on *AR* has been used as input parameters, i.e., rock quality designation (*RQD*), uniaxial compressive strength (*UCS*), rock mass rating (*RMR*), Brazilian tensile strength (*BTS*), thrust force (*TF*), and revolution per minute (*RPM*). In addition, a series of statistical indices, such as root mean squared error (RMSE), determination coefficient (R-square), and variance account for (VAF) are utilized to evaluate and compare the prediction precision of the developed models. Based on the simulation results and the computed values of indices, it is observed that the proposed GP model with the training and test RMSE values 0.0427 and 0.0388, respectively, performs noticeably better than the proposed ANN model giving RMSE values 0.0509 and 0.0472 for the training and test sets, respectively. Additionally, a parametric analysis has been conducted on the proposed GP model to further verify its generalization capability. The obtained results demonstrate that this GP-based model could provide a new applicable equation for accuratly predicting TBM performance.

5. Landslide prediction and early warning

(geological parameters and time)

1. Geology
2. Topography (gradient)
3. Rainfall
4. Site history
5. RQD
6. Joint spacing
7. …
8. …

Structured data?

Attributes of each instance?
(numeric vs non-numeric?)

DeepLearning.AI <thebatch@deeplearning.ai>

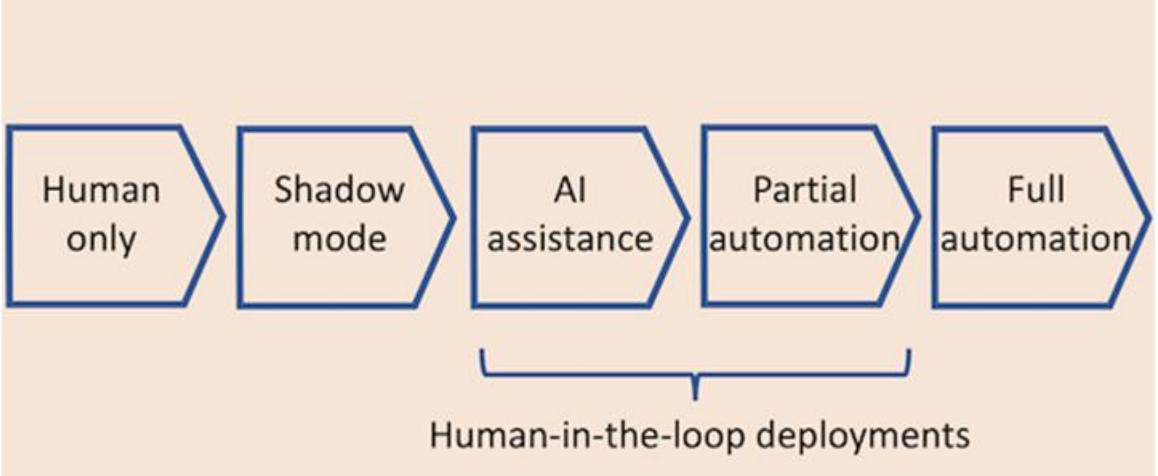

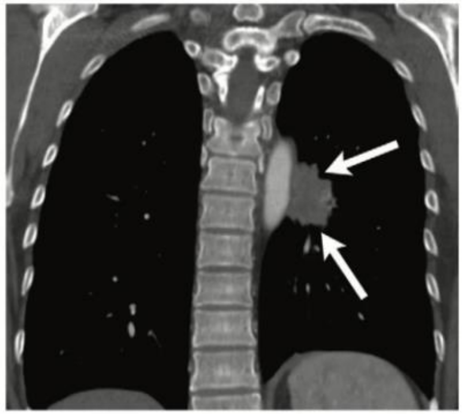

https://www.diagnosticimaging.com/view/deep-learning-algorithm-identifies-more-missed-lung-cancers-on-x-ray

Take the problem of diagnosing medical patients from X-rays. The deployment options include:

1. ***Human only:*** *No AI involved.*
2. ***Shadow mode:*** *A human doctor reads an X-ray and decides on a diagnosis, but an AI system shadows the doctor with its own attempt. The system's output doesn't create value for doctors or patients directly, but it is saved for analysis to help a machine learning team evaluate the AI's performance before dialing it up to the next level of automation.*
3. ***AI assistance:*** *A human doctor is responsible for the diagnosis, but the AI system may supply suggestions. For example, it can highlight areas of an X-ray for the doctor to focus on.*
4. ***Partial automation:*** *An AI system looks at an X-ray image and, if it has high confidence in its decision, renders a diagnosis. In cases where it's not confident, it asks a human to make the decision.*
5. ***Full automation:*** *AI makes the diagnosis.*

Are these questions in your mind?

1. What is the difference between Artificial Intelligence (AI), Machine Learning (ML) and Deep Learning (DL)? √
2. What is the relationship between AI model and AI algorithm? √
3. Is unsupervised ML better than supervised ML? √
4. What is a parametric ML algorithm and how is it different from a nonparametric ML algorithm? √
5. What is the difference between underfitting and overfitting of data? √
6. Is large amount of data equivalent to big data? √
7. Will AI be a solution to my day-to-day geo-related work? √

Conclusions

1. You are smarter than me → you can develop your AI projects
2. Hong Kong has lots of data → investment to compile and label the data
3. AI-enabled automation is often portrayed as a binary on-or-off (either automated or not). But in practice, automation is a spectrum, and AI teams have to choose where on this spectrum to operate → don't be disappointed if 100% accuracy is not obtained

https://www.vectorstock.com/royalty-free-vector/cartoon-human-brain-vector-21908216

+

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

+

DATA

=

https://www.corporatecomplianceinsights.com/model-governance-to-ai-ml-systems/

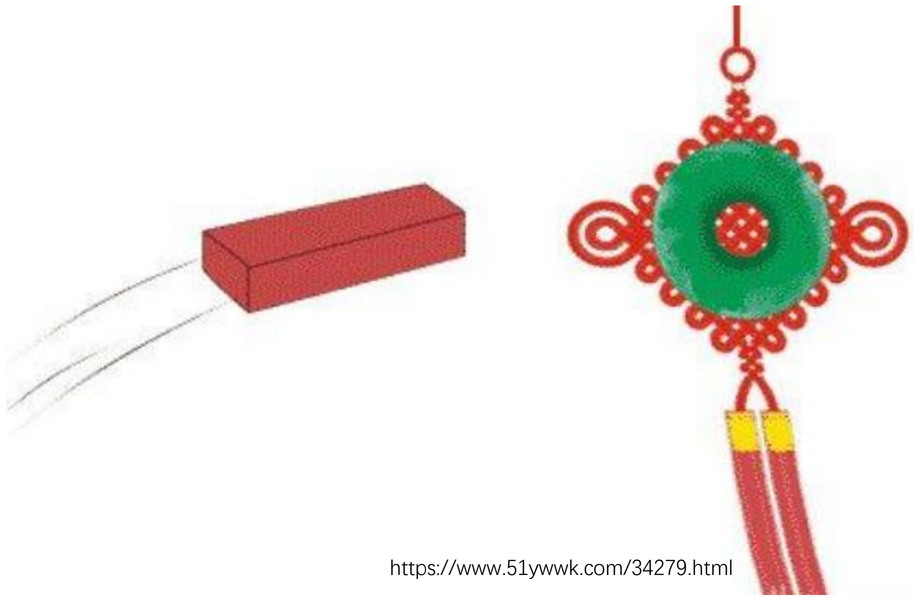

https://www.51ywwk.com/34279.html

throw out a brick so as to draw pieces of jade

(throw out a minnow to catch a whale)

Louis WONG

LNYWONG@hku.hk