life.augmented

Introduction to Artificial Intelligence

Working principles

Matteo Maravita

Asia AI Competence Center

August 4, 2022

AI

- Back in the summer of 1956 conference at Dartmouth College, the dream was:

“Construct a complex machine to have the same characteristics as human intelligence.”

- We call it Strong AI.

AI

- But what we can only achieve now is Weak AI (focused on one narrow task).

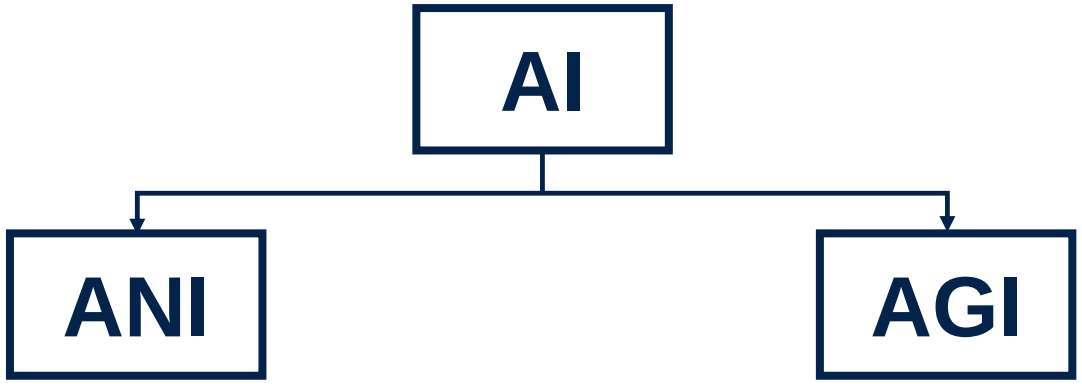

(artificial narrow intelligence)

(artificial general intelligence)

Computer vison:

self-driving cars, surveillance cameras, face recognition.

Speech recognition:

voice commands, speech-to-text subtitle generation.

Problem solving:

play games/chess/GO, recommendation system, etc.

life.augmented

Machine Learning

- ML is a subordinate branch of AI.

- These methods give computers the ability to learn without being explicitly programmed. (It means you don’t have to write rules.)

- It consists in algorithms that can learn and make predictions:
  - Such algorithms are trained on previous examples that give it the ability to learn how to perform the task.
  - Machine learning is usually employed where traditional programming is infeasible.
  - If trained properly, it should work for new cases.

Machine learning needs a new design mindset

From rule-based engineering to data-driven engineering

Standard programming

Handcrafted rules based on experience

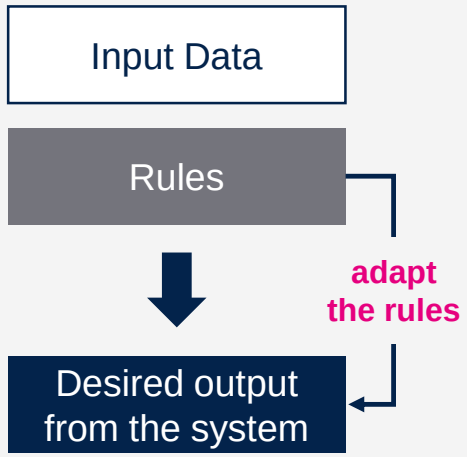

- Requires domain expertise to code
- Need to rewrite if environment evolves

Machine Learning

Rules learnt from real-world data

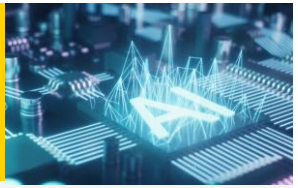

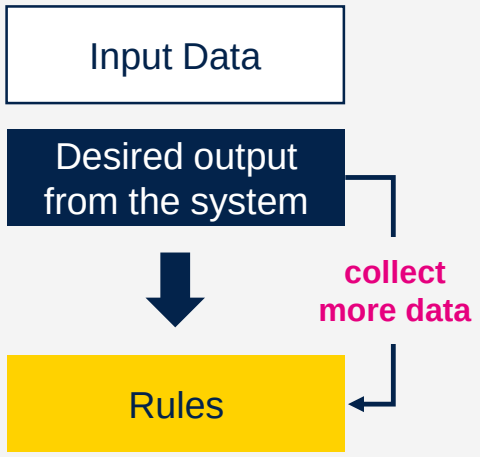

- Generate code from real-world observations
- Re-learn from data if environment evolves

life.augmented

Some Approaches to Machine Learning

- Decision tree (*) / Random Forest
- Support Vector Machine
- Clustering
- K-Means
- K-Nearest Neighbors
- ….
- Deep learning

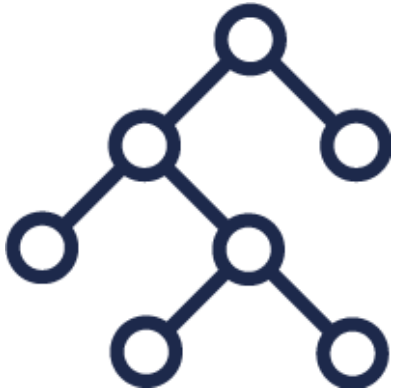

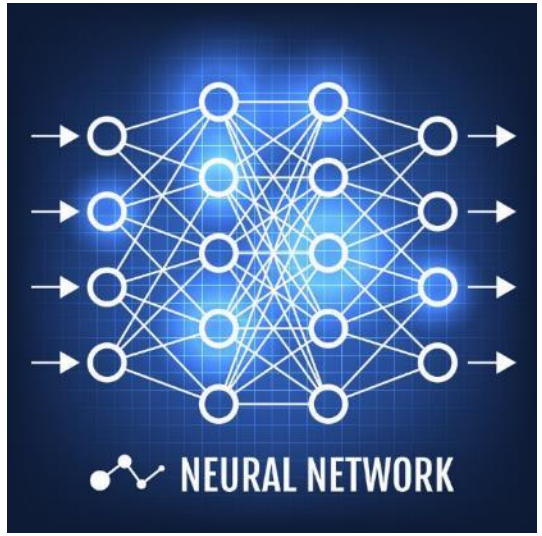

Machine Learning

- Can be divided into:
- Supervised learning: machine learning algorithm which learns a function that maps an input to an output based on example input-output pairs. E.g., Decision Tree, Support Vector Machine, Linear Regression, Deep Learning, …
- Unsupervised learning: machine learning algorithm which learns unknown patterns from un-labeled data. E.g., Clustering, K-Means, …

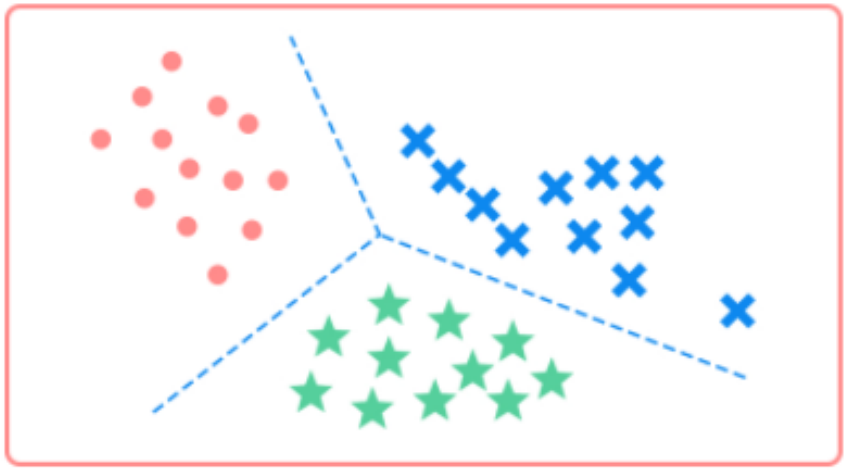

Supervised learning

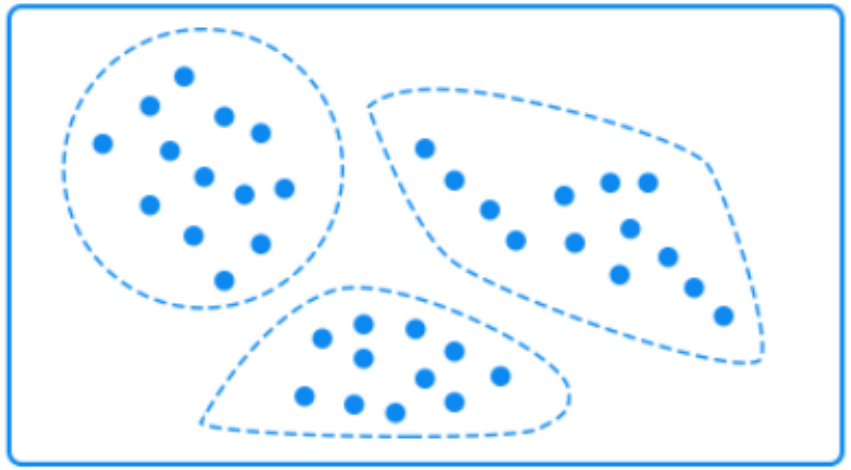

Unsupervised learning

Supervised Learning

Input(x)	Output (y)	Application
Home features	Price	Real Estate
Ad, user info	Click on ad? (0/1)	Online Advertising
Image	Object (1,…,1000)	Photo tagging
Audio	Text transcript	Speech recognition
English	French	Machine translation
Image, Radar info	Position of other cars	Autonomous driving

Classification vs Regression

- Supervised learning contains mainly 2 categories:
- Classification: algorithm to predict a discrete class label
- Regression: algorithm to predict a continuous quantity

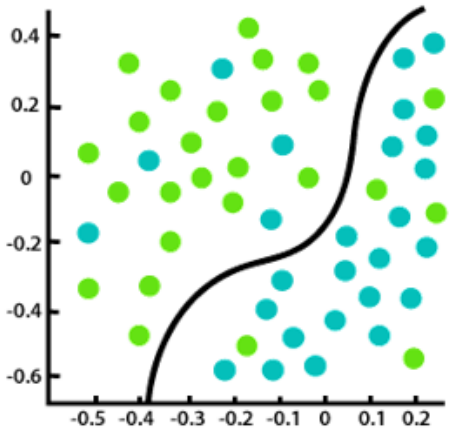

Classification

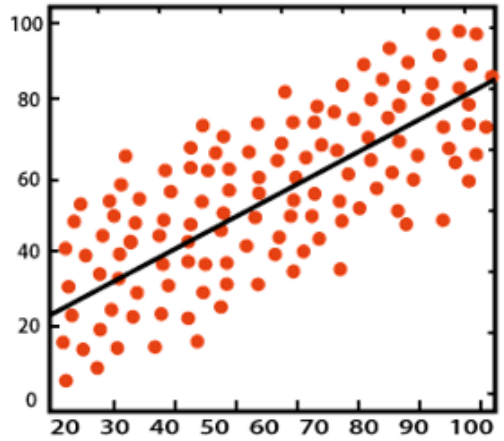

Regression

Example: Decision Tree & Random Forest

- Suppose a bank wants to determine whether to loan money to customer.
- We have some historical data with customer information (**Annual income, whether have real estate, previous credit**) and whether they have successfully repaid the loan (**Yes/No**).

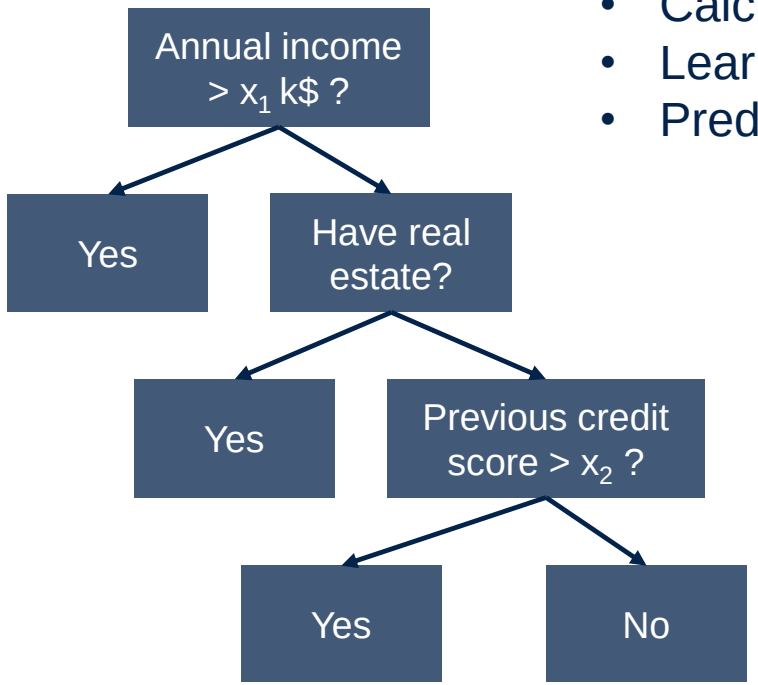

- Calculate the information gain to determine which feature comes first
- Learn the x_1 and x_2 from the historical data
- Predict new data using learned decision tree model

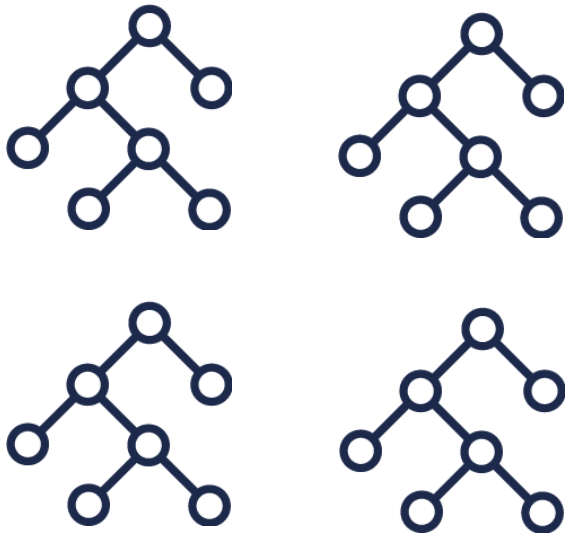

- **Random forest** consists of many independent decision trees
- When doing classification, every decision tree gives a vote, the final result will be the one with the most votes
- To build independent trees, randomly choose m of M features, and n of N samples

life.augmented

Machine Learning Framework

- Scikit-Learn is a free software machine learning library for Python, which supports various traditional machine learning algorithms including both supervised learning and unsupervised learning.

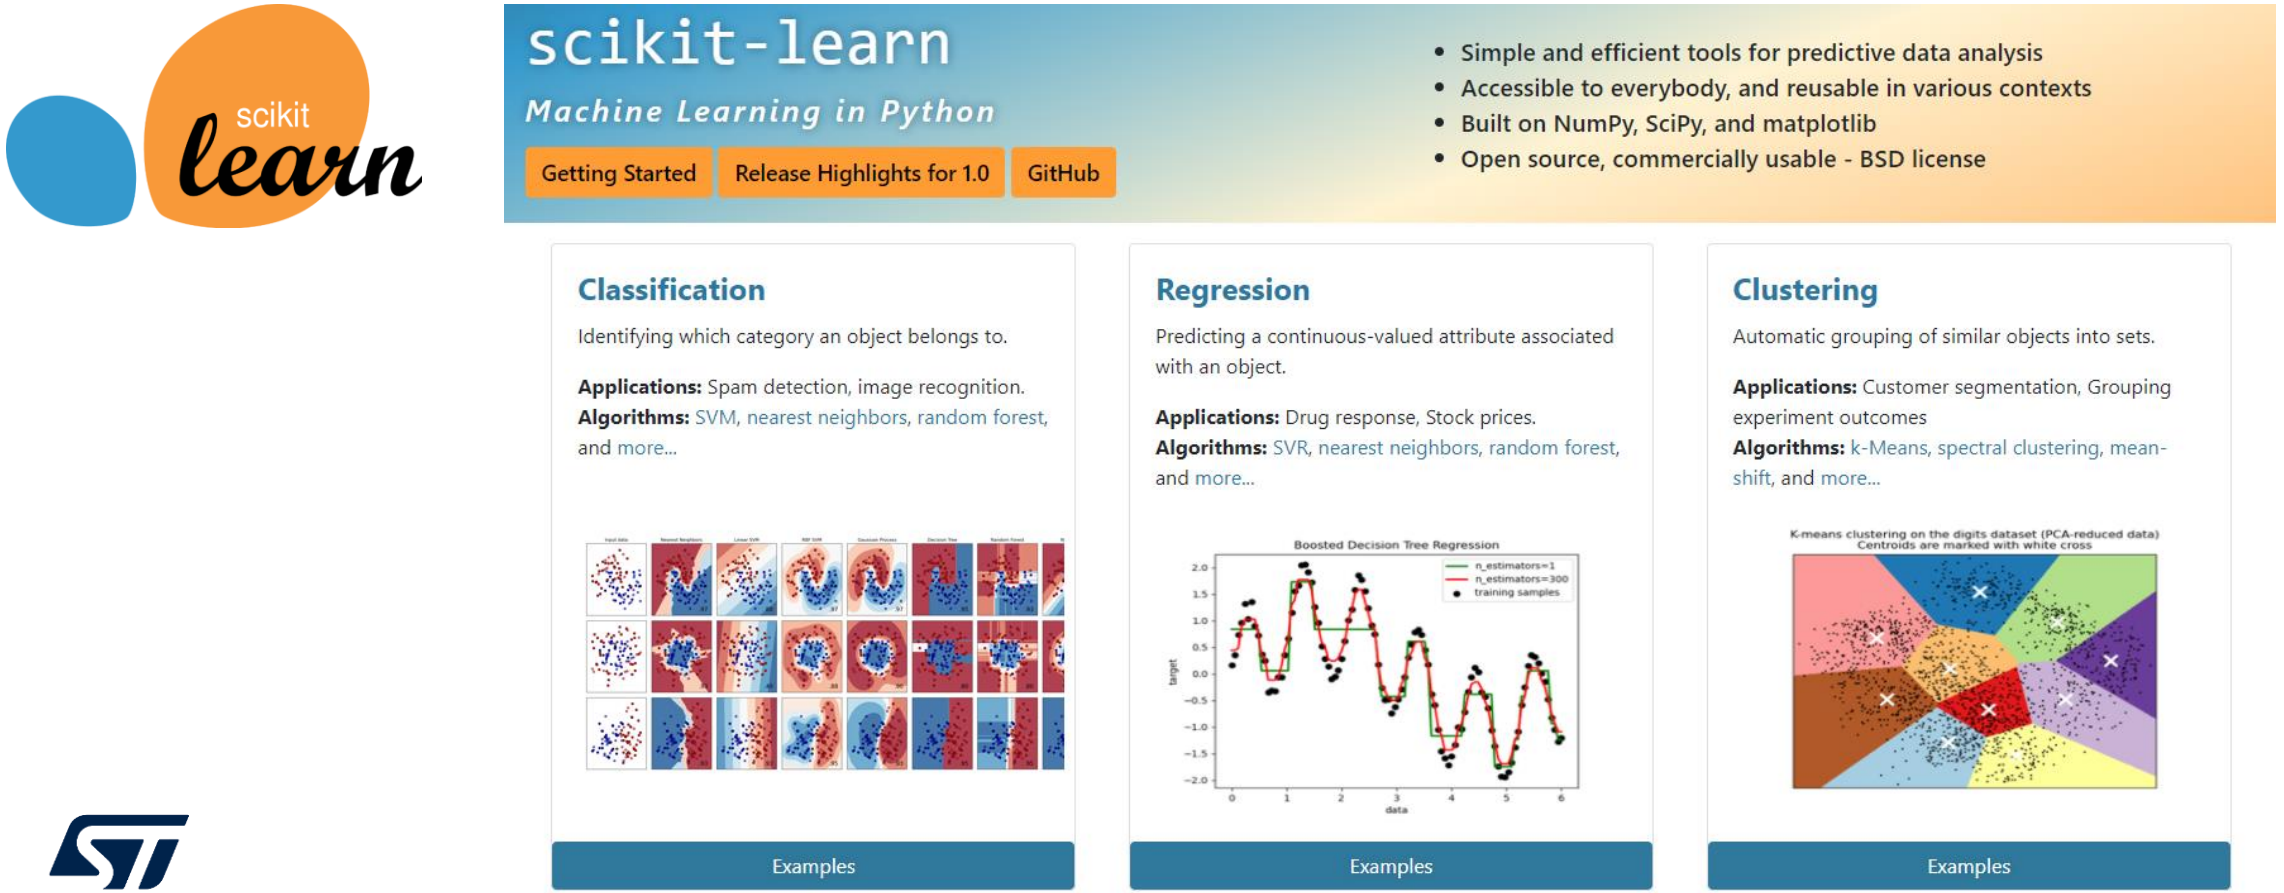

Deep Learning

- Deep learning is ML using large neural networks.
  - Inspired by human brain
  - Deep because of having many intermediate learning steps and many layers.
  - Large amounts of data is required

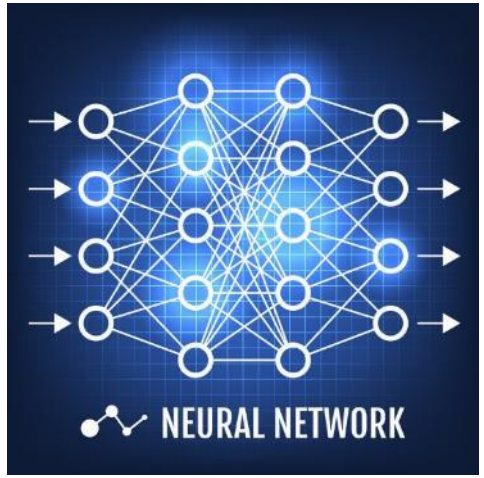

Advantages	Disadvantages
Automatedly learns the patterns & relations in the data	***Large Datasets***
High Accuracy	High Computational Requirements
Easy Improvement & Fine Tuning	Weak theoretical explanation
Adaptive Solutions	Black box (for most people)

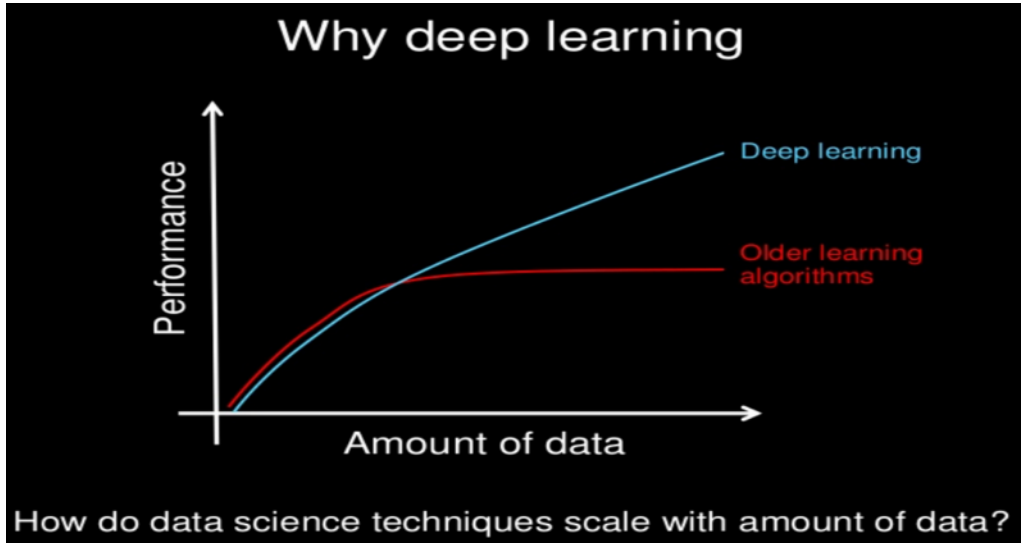

life.augmented

Why deep learning so important?

It outperforms existing methods in many competitions.

Problem	Dataset	Best Accuracy w/o CNN	Best Accuracy with CNN	Diff
Object classification	ILSVRC	**73.8%**	**95.1%**	**+21.3%**
Scene classification	SUN	**37.5%**	**56%**	**+18.5%**
Object detection	VOC 2007	**34.3%**	**60.9%**	**+26.6%**
Fine-grained class	200Birds	**61.8%**	**75.7%**	**+13.9%**
Attribute detection	H3D	**69.1%**	**74.6%**	+5.5%
Face recognition	LFW	**96.3%**	**99.77%**	**+3.47%**
Instance retrieval	UKB	**89.3%** (CDVS: 85.7%)	**96.3%**	+7.0%

May 2015

Today, image recognition via deep learning in some scenarios is better than humans.

life.augmented

Why do we need ML (1)?

- In some cases, it is extremely difficult to write a program, e.g.,

- 3D object recognition:
  - We can easily recognize a 3D object using our brain, but
    - We don’t know how to tell the computers to do that
    - Even if we can, it’s very complicated and difficult.

Why do we need ML (2)?

- An example of difficult program
  - How to recognize the hand-written digits?
  - Difficult to define the rules.
  - What makes all these numbers to be identifiable?
  - Is there a pattern?
  - What is it that makes a 2 to be identified as a 2?

Some examples from MNIST database
(Mixed standard institute for standard and technology)

Applications of Neural Networks

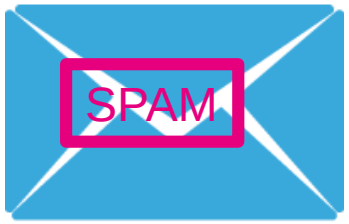

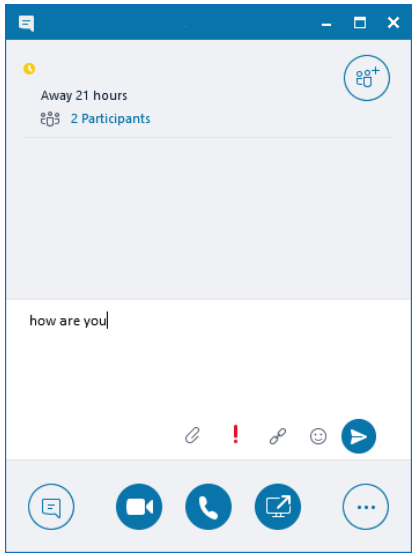

life.augmented

ImageNet Large Scale Visual Recognition Challenge

- 1000 different object classes in 1.3 million images (high resolution)
- ILSVRC is the standard benchmark in objective category classification and detection on 1,000 object categories and millions of images, specifically tagged;
- In ILSVRC 2014:
  - Sponsors were: Google, Stanford, Facebook and the University of North Carolina
  - Main participants were: Google (winner), Adobe, Microsoft, Samsung, Lenovo, Orange, Toyota, and several major Universities.

Team	Year	Place	Error (top-5)	Uses external data
Clarifai	2013	1st	11.2%	ImageNet 22k
MSRA	2014	3rd	7.35%	no
VGG	2014	2nd	7.32%	no
GoogLeNet	2014	1st	6.67%	no
VGG	Post 2014		6.8%	
Baidu 百度	Post 2014		5.98%	
Human Ref			**5.1%**	
Microsoft PReLuNet	Post 2014		4.94%	
Google	Post 2014		4.90%	

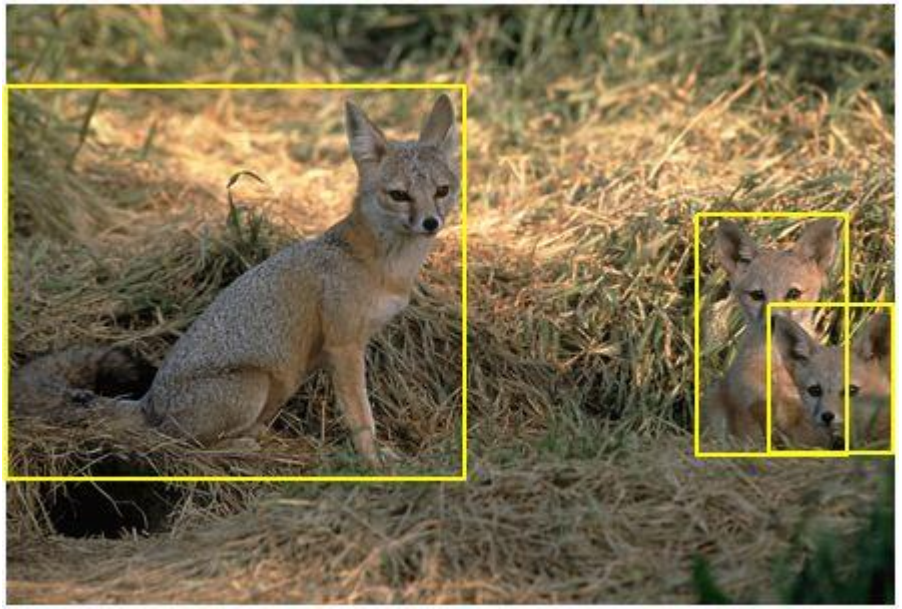

Some Examples from ImageNet dataset

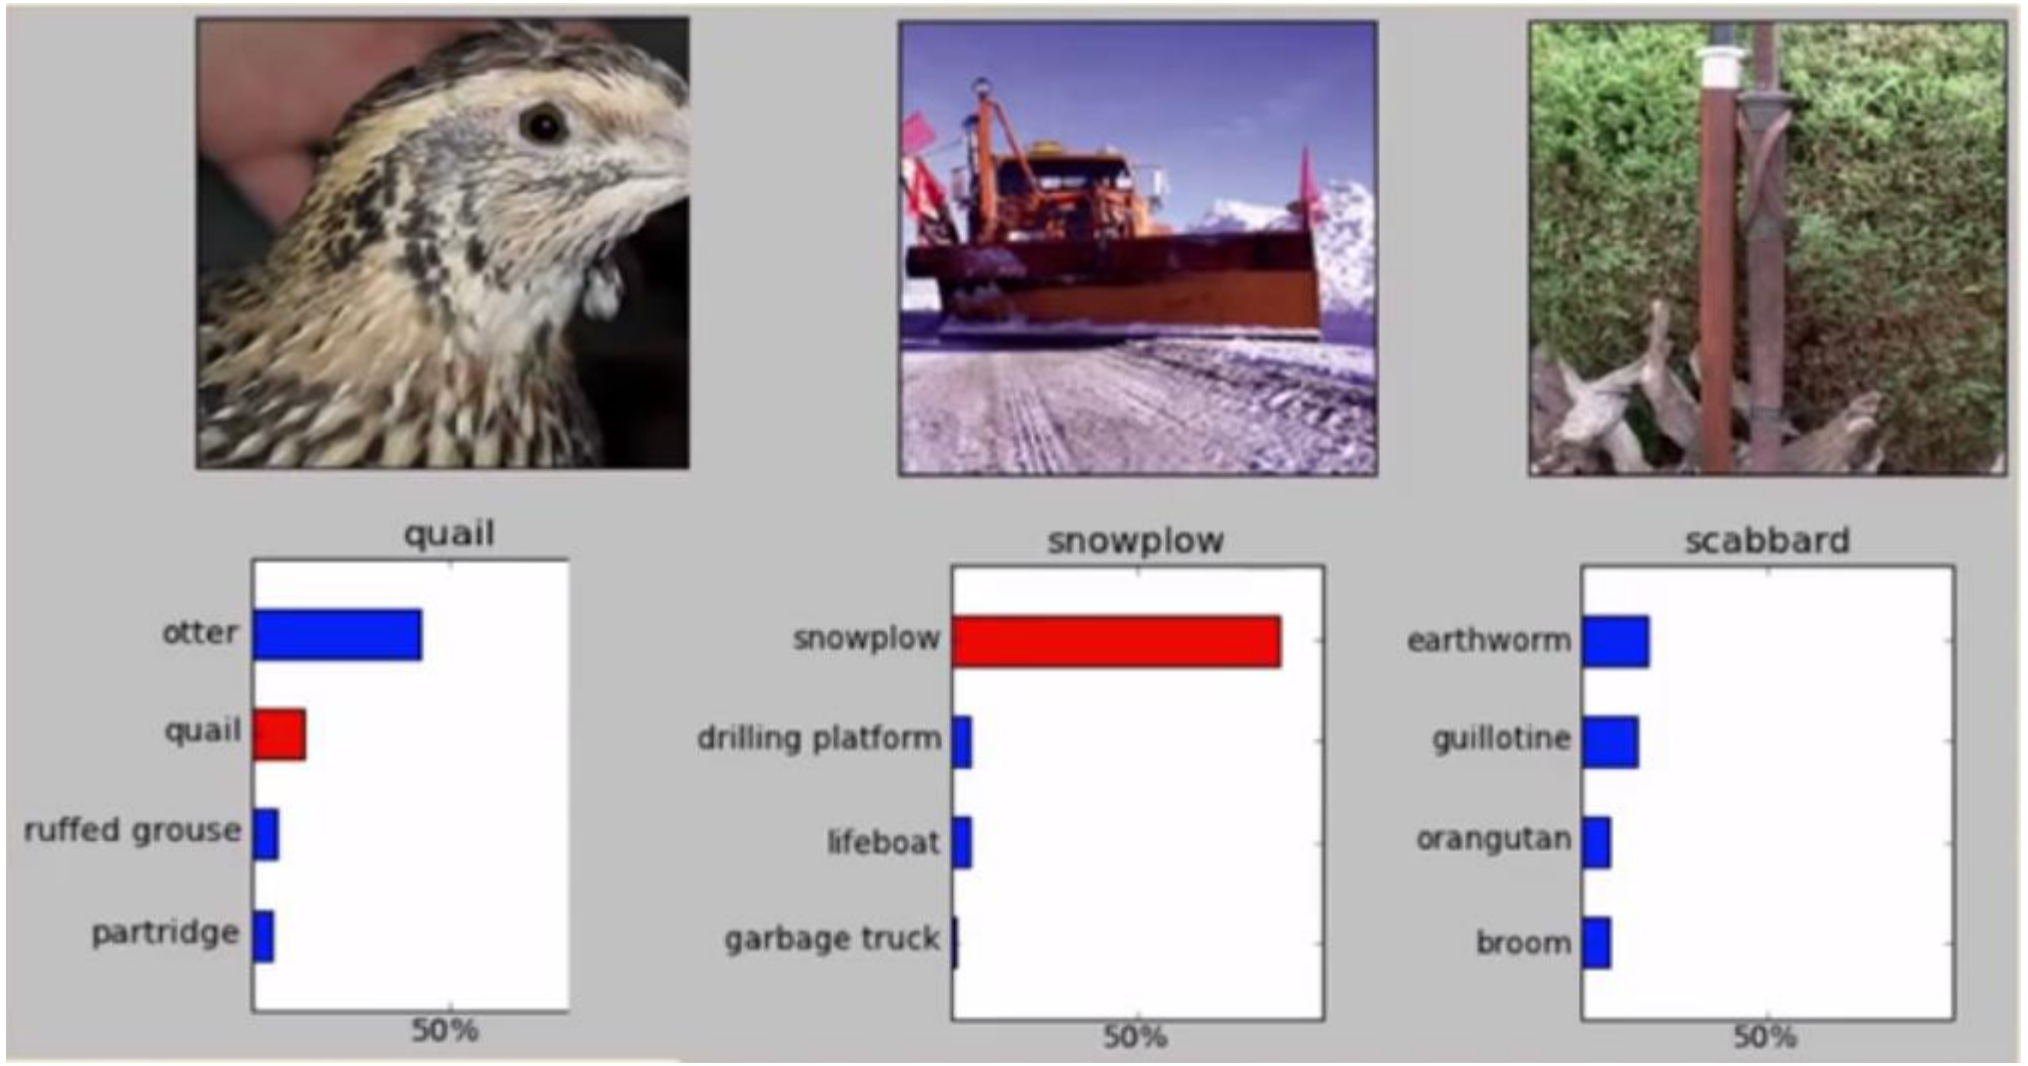

life.augmented

Deep Learning Neural Networks

life.augmented

What are Neural Networks?

- Also referred to as artificial neural networks.
- Inspired by human neural system.
- Human neuron has three main components
  - Dendrites
    - Takes inputs from other neurons in terms of electrical pulses.
  - Cell body
    - Makes the inferences and decides the actions to take.
  - Axon terminals
    - Sends the outputs to other neurons in terms of electrical pulses.

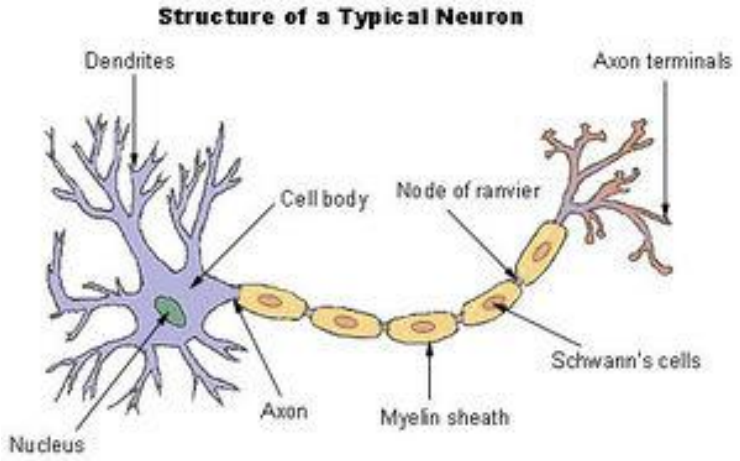

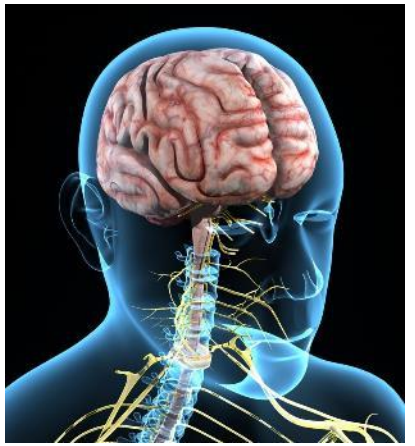

Artificial Neuron

The heart of a neural network

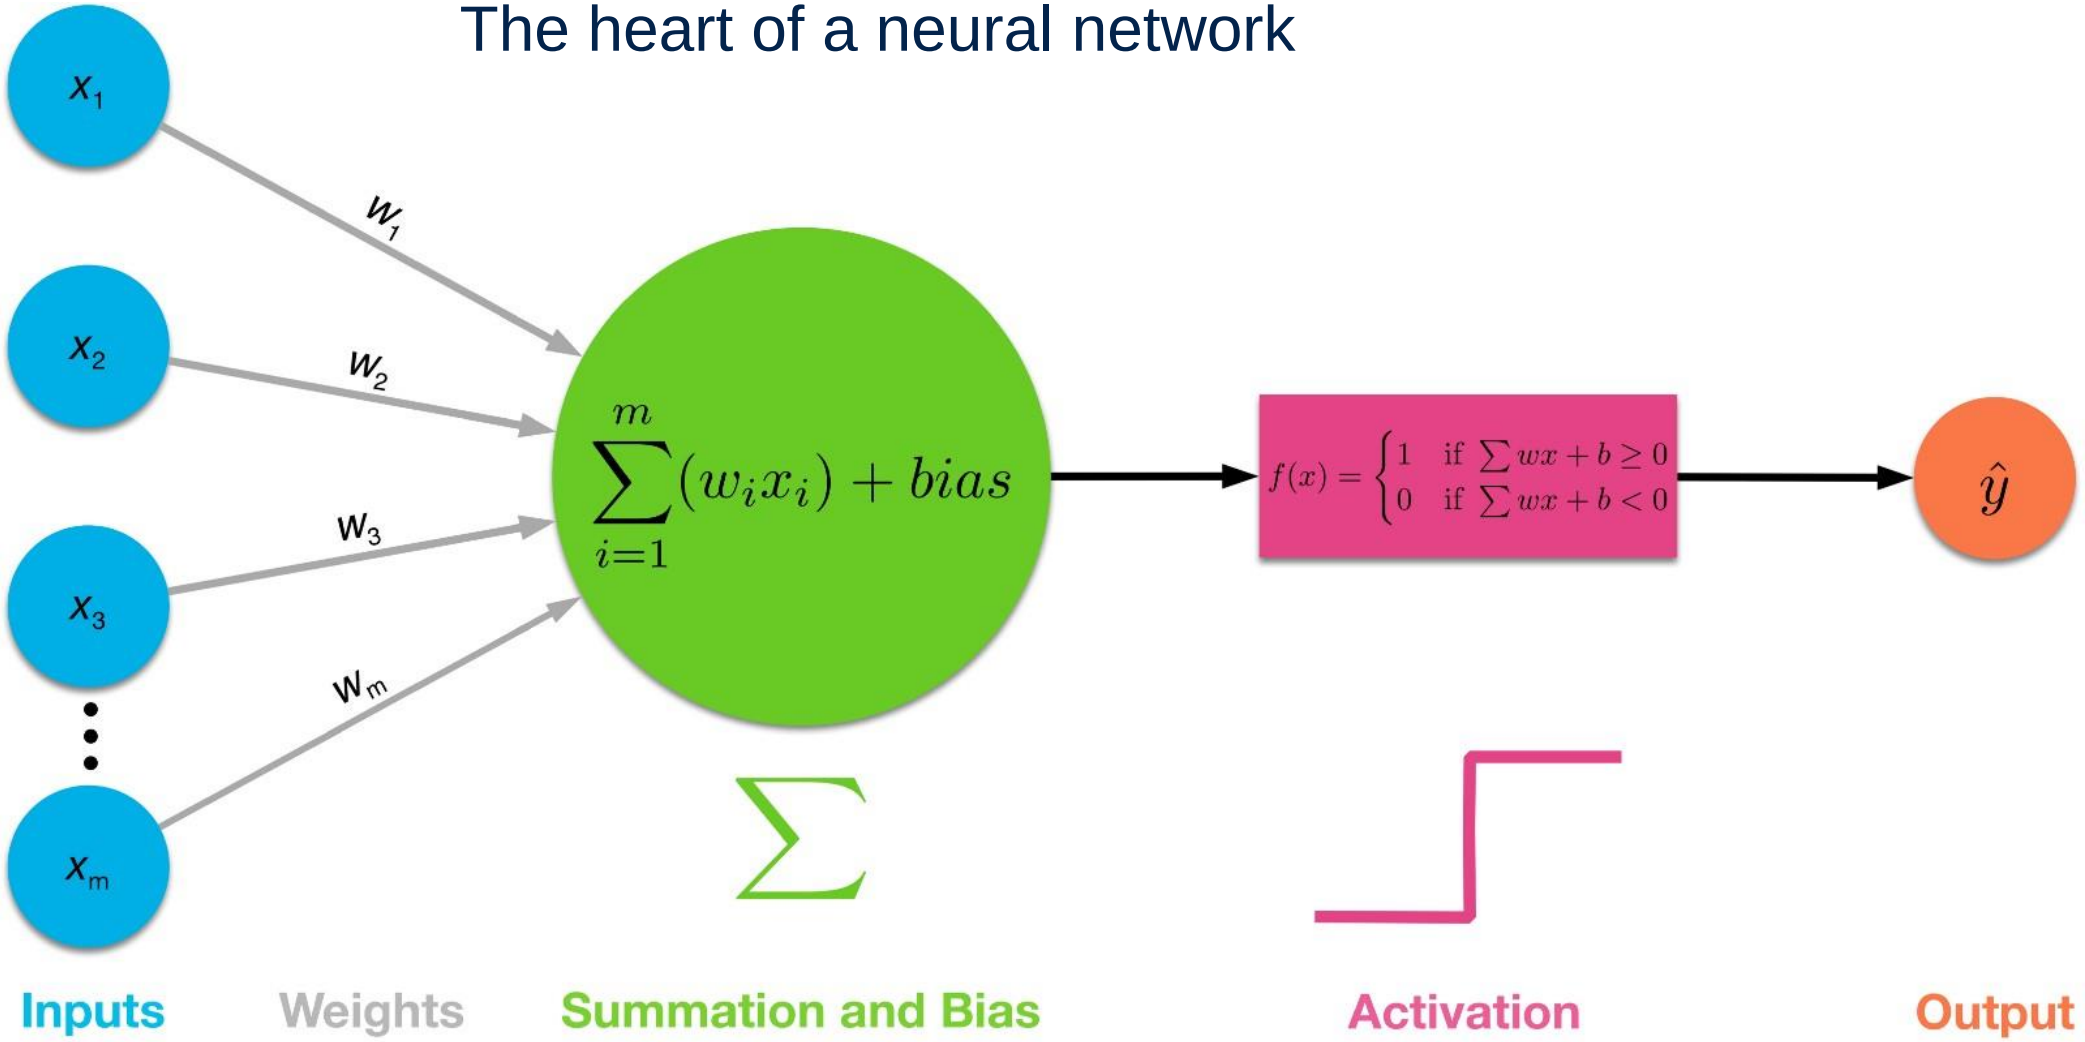

Artificial Neuron

Weights of the NN need to
be stored in the non-volatile memory
(Flash or ROM)

The heart of a Neural Network

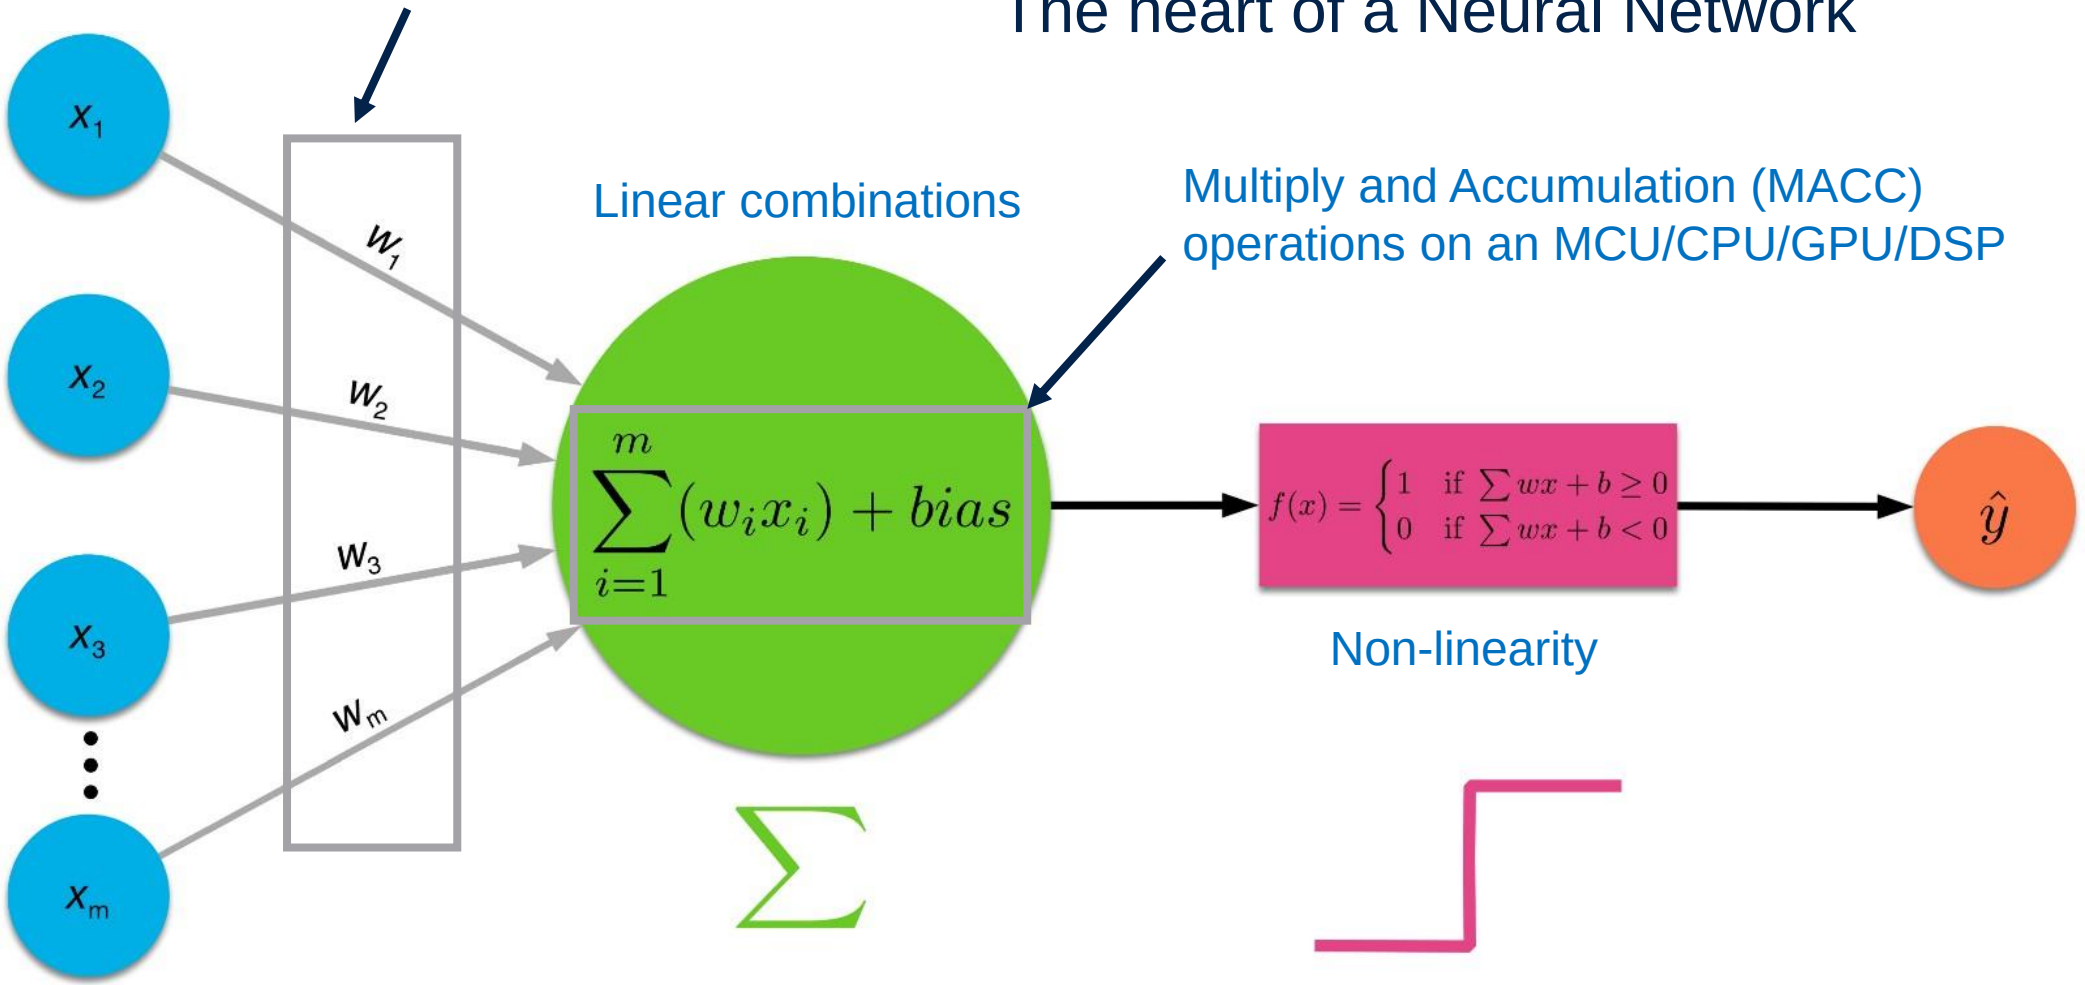

Most widely used activations

- Unit step
  - Threshold
- Sigmoid Function
  - Like a step function but smoother
  - Best to predict probabilities
- Tan hyperbolic
  - Stretched out version of the sigmoid function
- ReLU
  - Computationally efficient

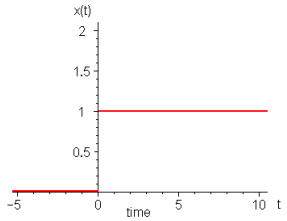

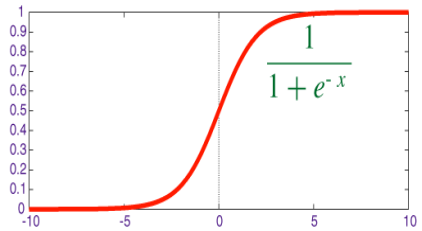

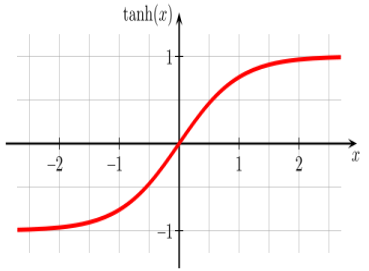

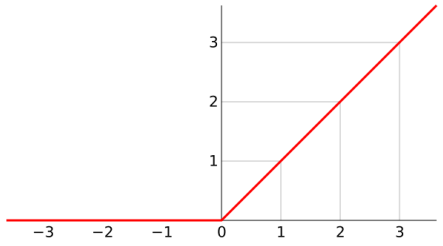

All activation functions introduce a “non-linearity” in the system that is extremely important

life.augmented

Layers of a Neural Network

- Neural network has three types of layers
- Input layer
  - Can be from other neurons or feature inputs
    - Age, height, weight, pixels in the images etc.
- Hidden layers (one or more)
  - Real power lies here
  - Adding more neurons to the network
- Output layer
  - Gives the output we want to predict
    - Probability of rain
    - Object class
    - Disease is fatal or not…

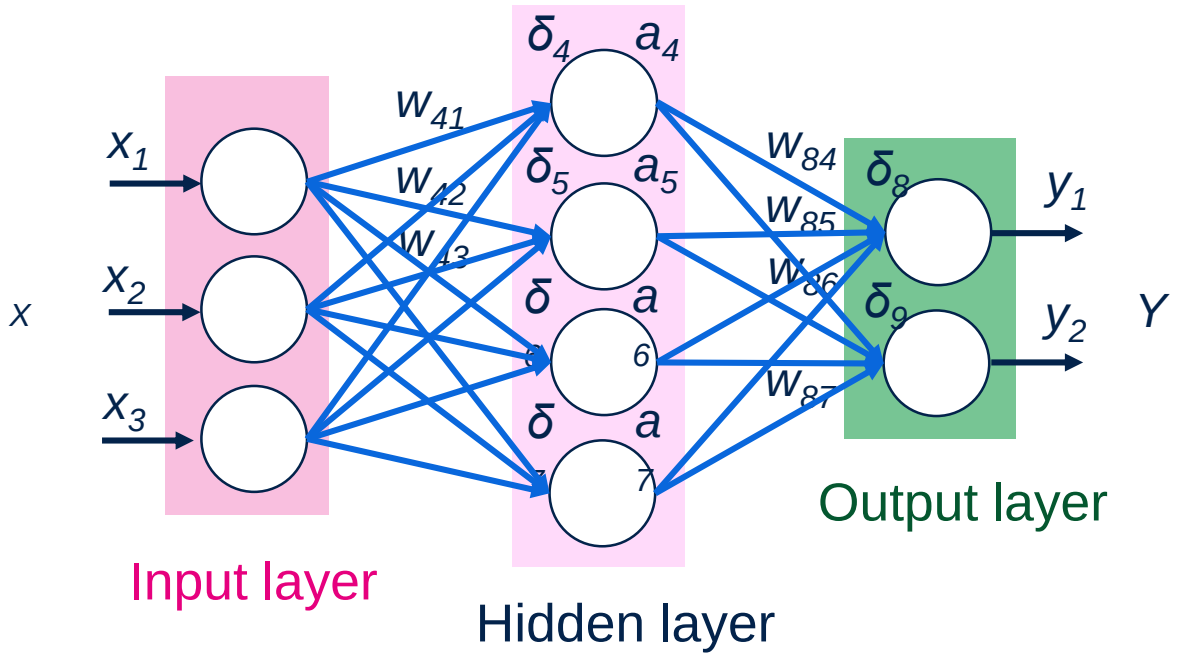

Example of a Neural Network

- Let us say we have a problem, in which we design a neural network to predict the cost of an apartment

- We can have several inputs or features to my NN:
  - Size
  - # rooms
  - Distance from nearest station
  - ZIP code
  - Car parking
  - Good school nearby
  - Pets allowed

- Let us suppose we have a pre-trained neural network, which means we have all the weights.

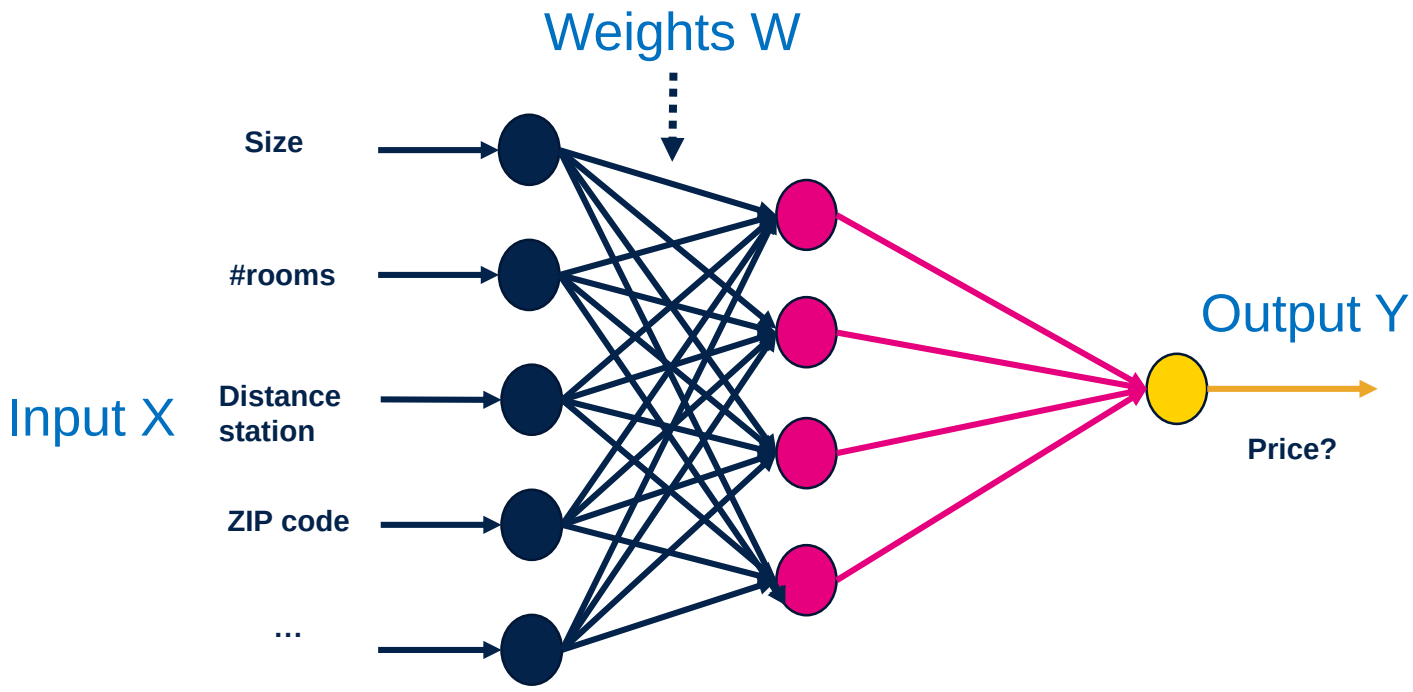

How to decide the weights?
If I have a good database of prices of apartments in the area of Tokyo I can “train” my NN network to give me the output most close to the database I have.

Different types of Neural Networks

Feedforward Neural Network
(or Dense or Fully Connected)

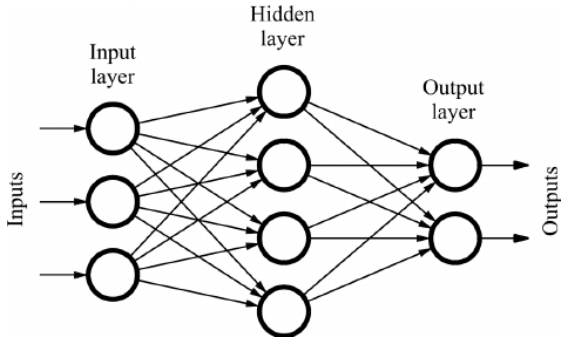

Convolutional Neural Network
(CNN)

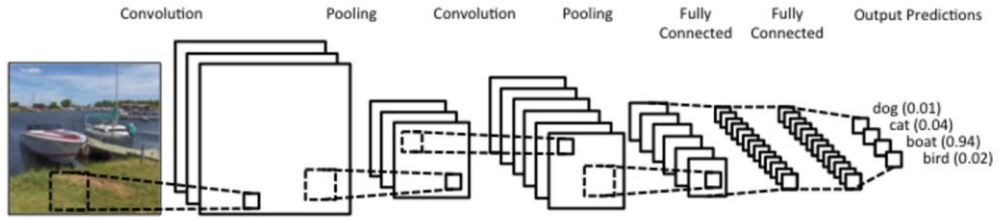

Recurrent Neural Network
(RNN)

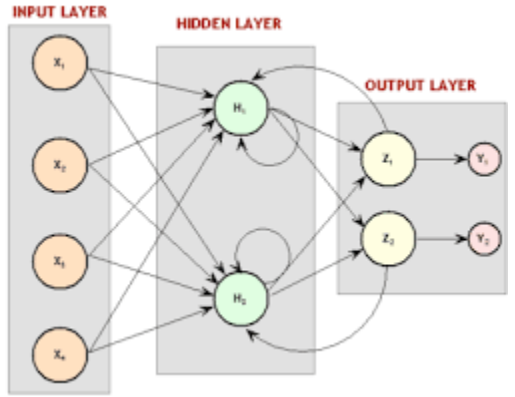

life.augmented

Training neural networks

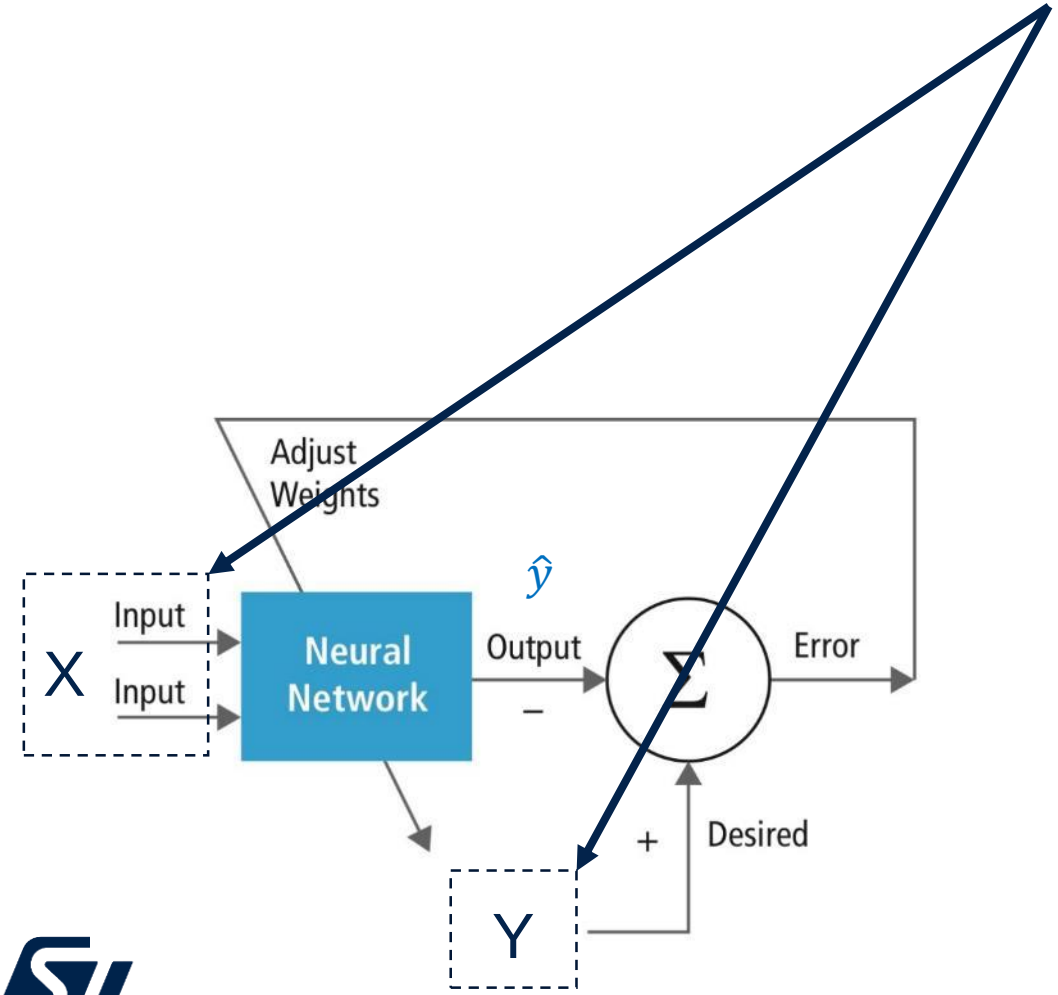

- To train a NN I need a dataset with both Inputs (X) and outputs (Y, also referred as "*labels*").
- During the "training" process I want to find the best weights of my NN that would give an output $\hat{y}$ as close as possible to the output labels Y.
- This operation is done through several iterations on a large dataset (from few thousands to few millions of samples!!) and is computational very consuming (usually done on very powerful PC and servers with dedicated GPU).
- The error between the output of the NN with the given weights vs the labeled outputs Y give me the Accuracy of the NN.

life.augmented

Learning datasets

- One of the difficulties of working with NN is to have a big dataset. In general, 90% of the time spent for a Deep Learning project is related to the creation of the Dataset.
- For Image Processing and few other applications there are already available big databases (with labels), both free or paid basis. On other hand for new and “niche” applications the designer needs to create the database, that becomes a high value asset.
- The dataset is then divided into
  - Training set
    - Training set
    - Validation set
  - Testing set

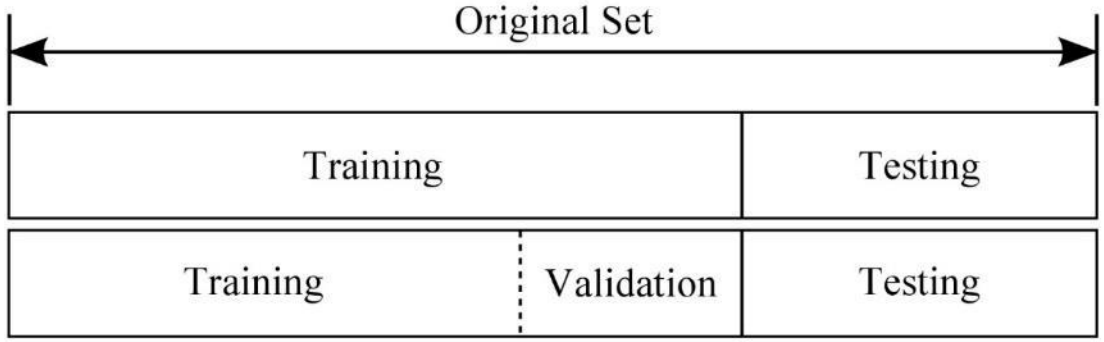

life.augmented

AI design process

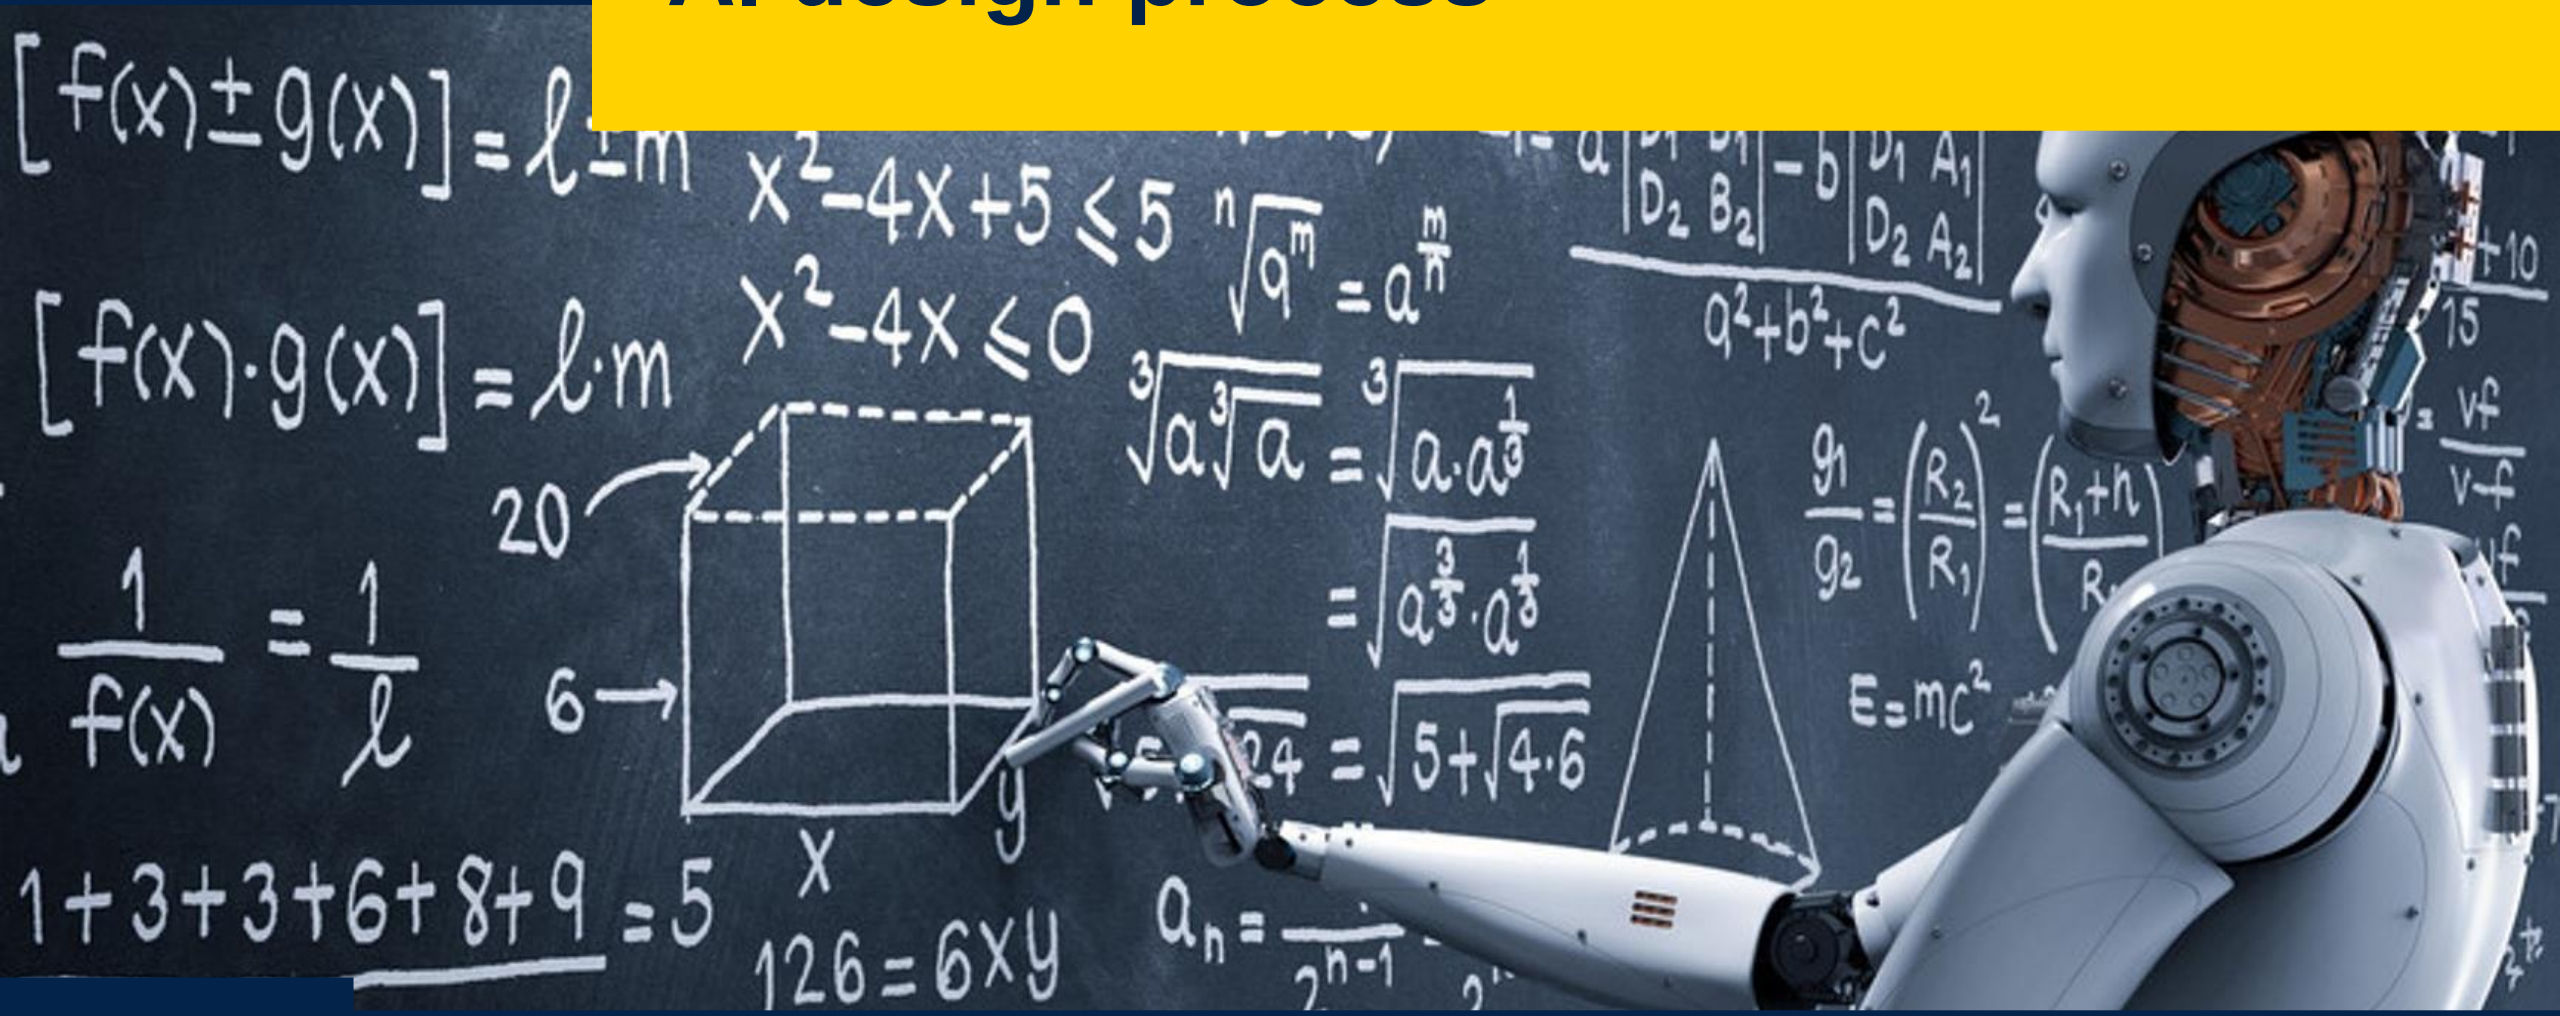

Data-driven approach

- We provide many examples of each class to the computer and then develop machine learning algorithms that look at these examples and learn about the visual appearance of each class.

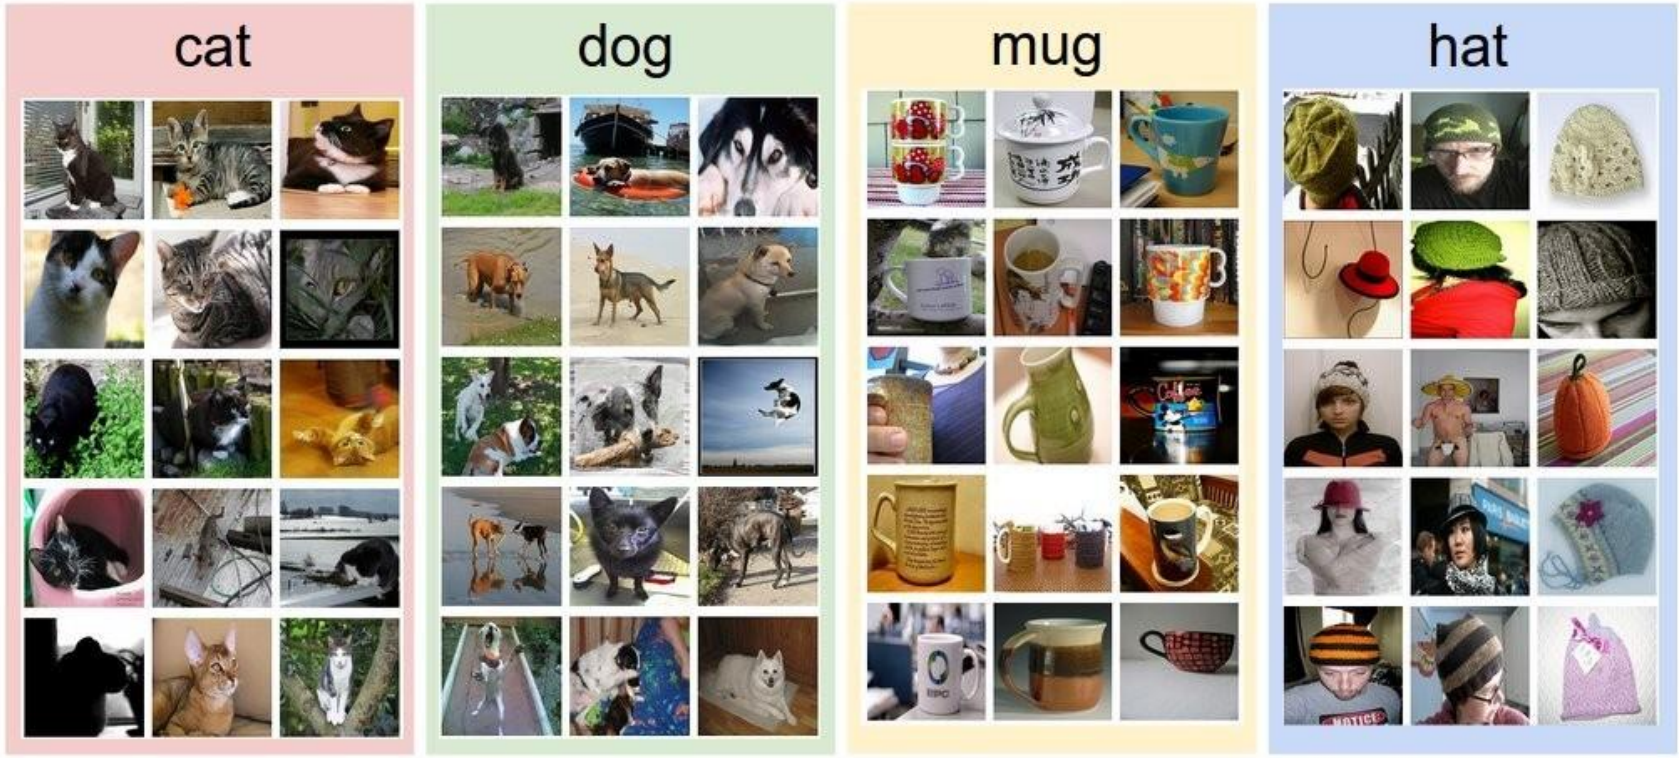

Deep Learning design process

Deep learning design process is iterative and always evolving.

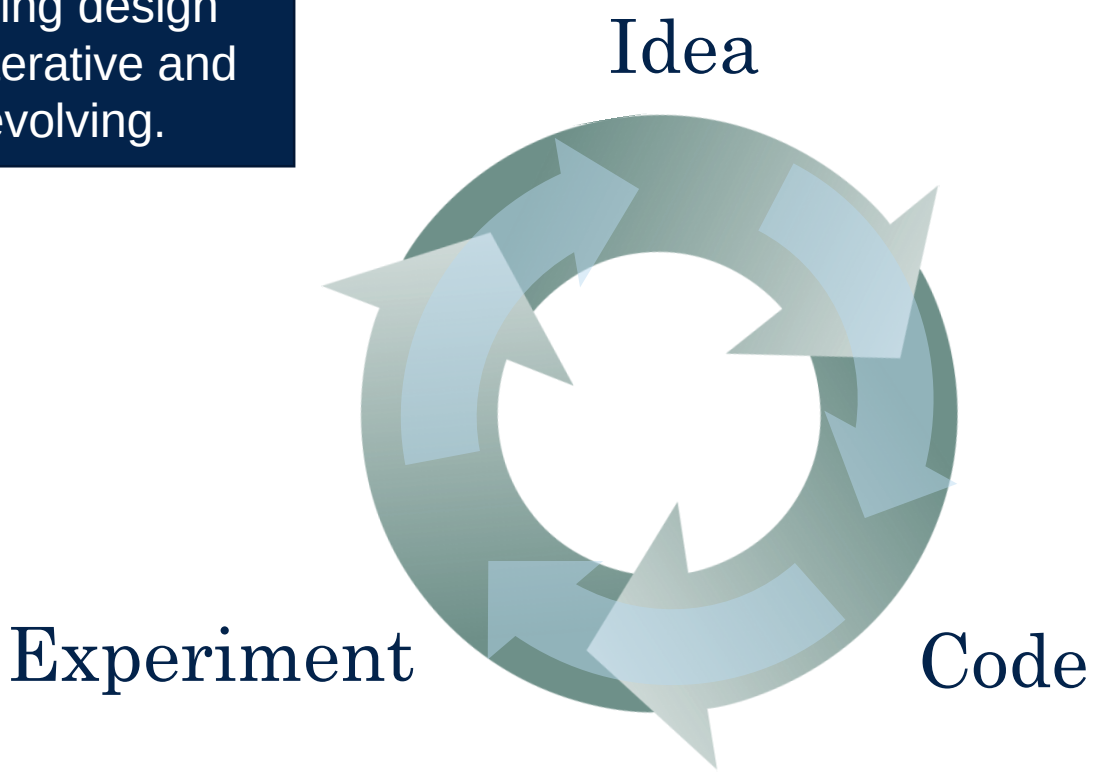

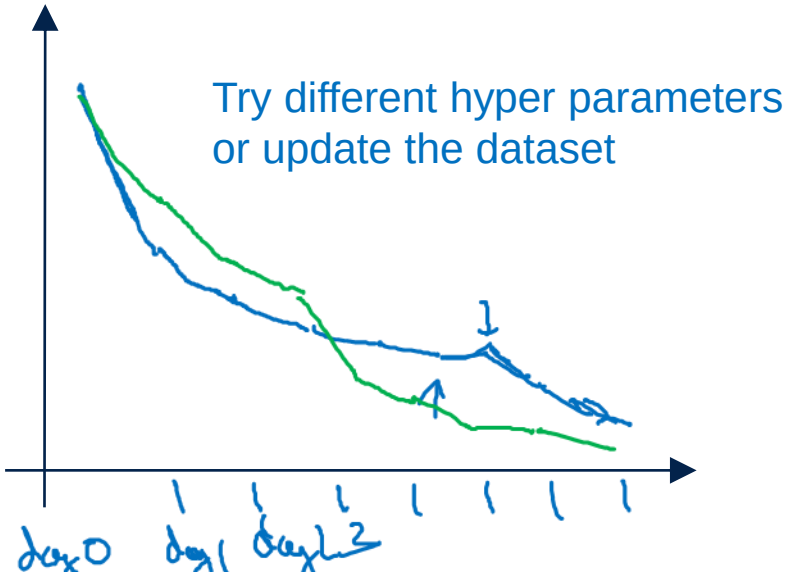

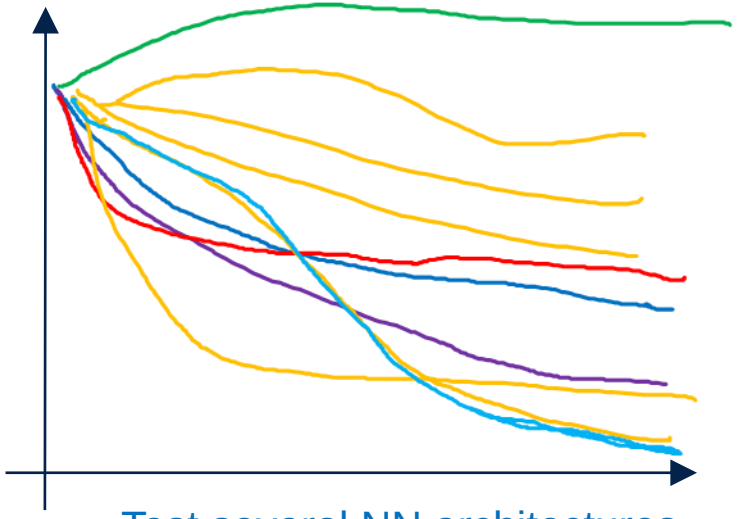

Test several NN architectures

Credit: Andrew NG, Deep Specialization course, Coursera

Example of impact of input image size

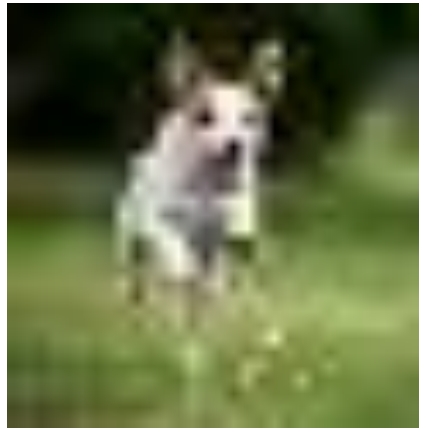

64x64x3=12,288

Dog? (0/1)

Number of weights increases dramatically with the input image size.

1000x1000x3=3M

Dog? (0/1)

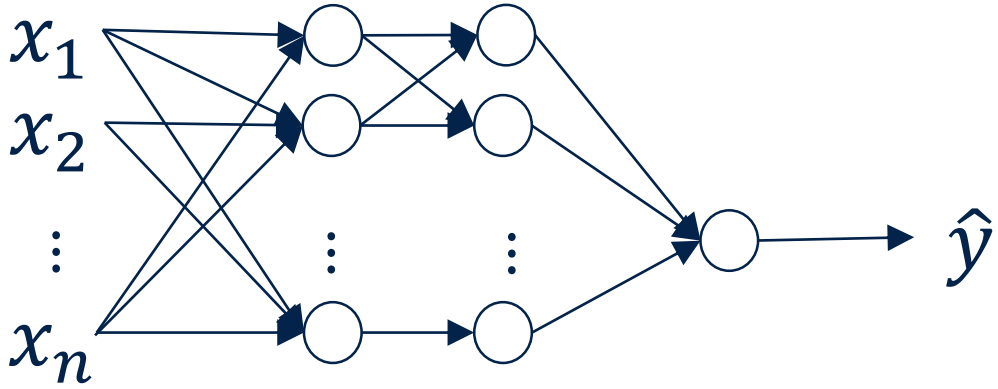

3Million inputs

1000 nodes

3Billion weights

NN validation and selection

The selection of a NN can depends on many performance parameters,
It depends what is more important and if there are minimum requirements for the designer: accuracy, cost, max memory size, MCU cost, …

NN	Accuracy	Running time [ms]	Memory size [Kb]	Input size	# Classes
#1	92.6%	12	359	100x100	20
#2	94.1%	36	620	100x100	20
#3	94.9%	150	2523	100x100	100
#4	92.1%	26	112	64x64	20
#5	87.6%	10	32	32x32	10

The centralized approach to AI

CENTRALIZED AI
Data is provided from sensors to the Cloud without any intermediate analysis

Pros
- Processing power
- Simpler to update (centralized) computing models

Cons
- Responsiveness (slow)
- Bandwidth (high)
- Privacy
- Energy consumption

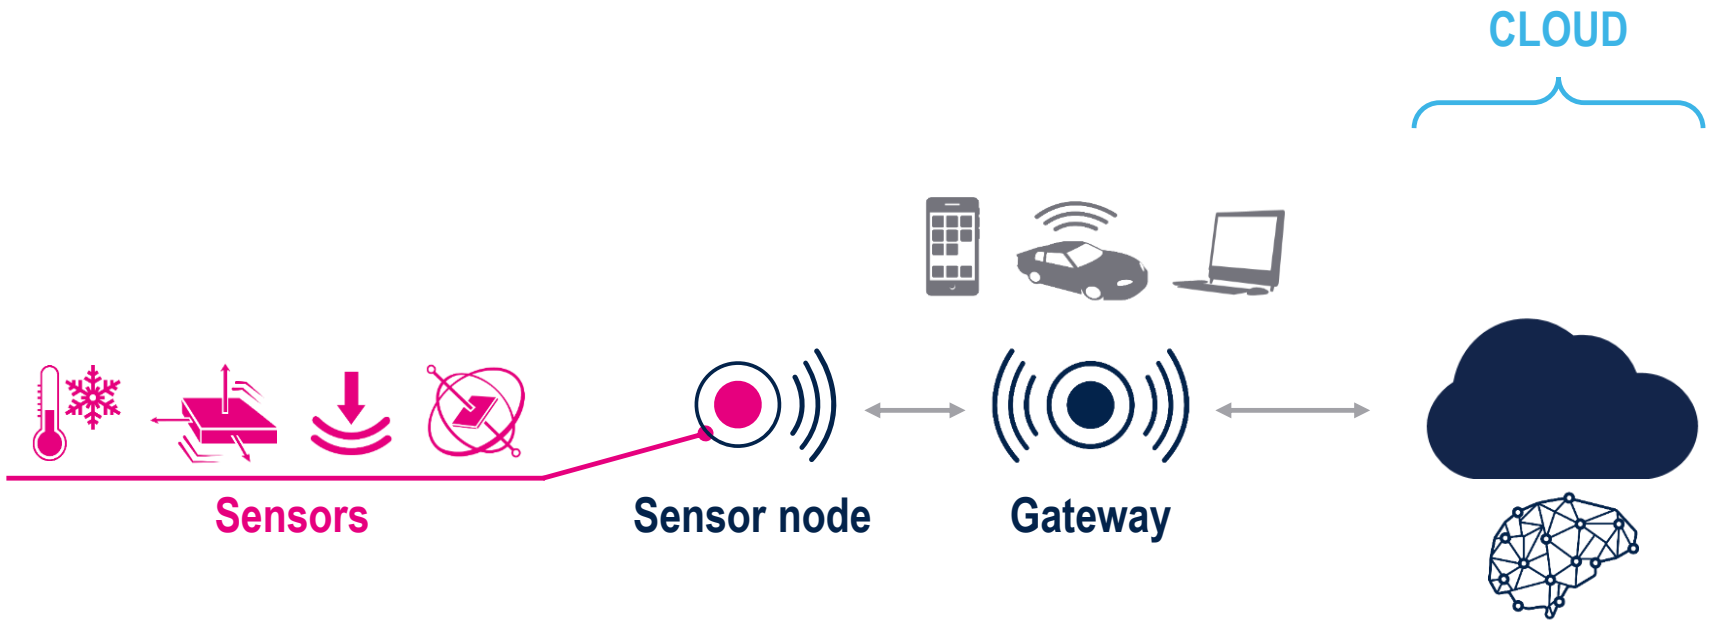

life.augmented

Moving toward Edge processing

Pros
- Responsiveness (fast)
- Bandwidth (low)
- Improved privacy (data stored locally)
- Energy saving

Cons
- Distribute the processing over multiple devices

DEEP EDGE AI
Data are processed in the smart sensors, smart nodes and smart gateways

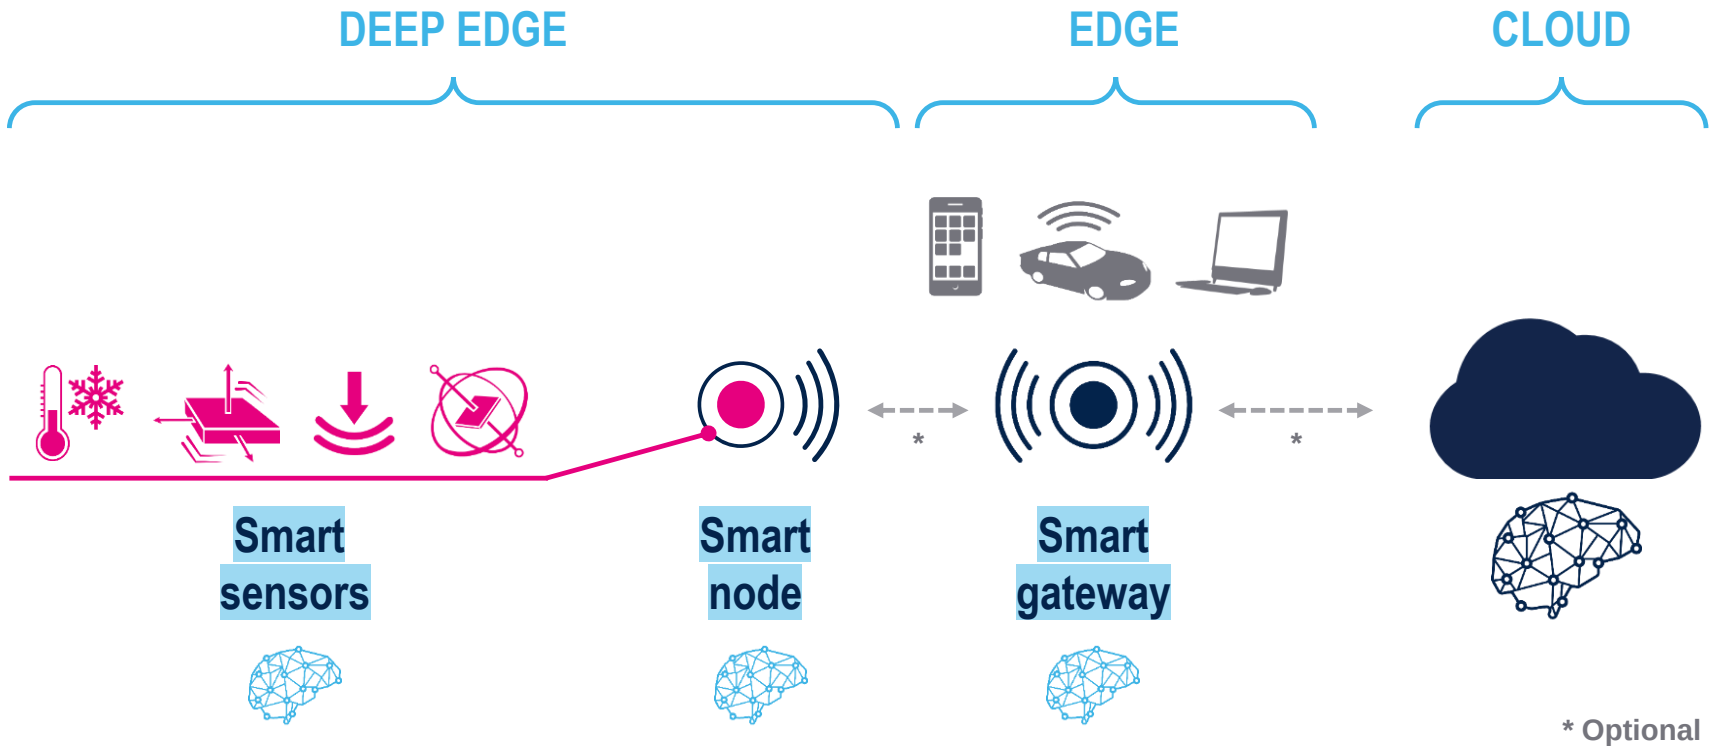

life.augmented

Our technology starts with You

life.augmented