WAZIUP Online Course

Fundamentals of IoT

F-IOT-2a: Understanding IoT Devices, Architecture & Ecosystem

This part focuses on IoT+AI for WaziGate

Prof. Congduc Pham
http://www.univ-pau.fr/~cpham
Université de Pau, France

IoT – from idea to reality

WAZIUP technological components

Sensor part

Intel-IrriS

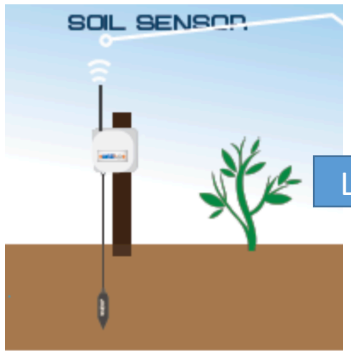

LPWAN

Control part

Embedded Database

Software updates

Embedded User Interface

Wifi

SMS

Sense, Monitor, Optimize & Control

DATA ANALYSIS, OPTIMIZATION & CONTROL

MONITORING

SENSING PHYSICAL WORLD INTERACTION

APPLICATION DOMAINS

Firebase

FIWARE

Axēda

ioBridge Connect things.

ThingSpeak

GroveStreams

SensorCloud

freeboard

dweet.io

OpenRemote

TempoIQ

Prof. Congduc Pham
http://www.univ-pau.fr/~cpham

Edge-AI for fully autonomous system

- Embed every thing on the IoT gateway to provide a fully autonomous system for the "XXXX-in-the-box" solutions

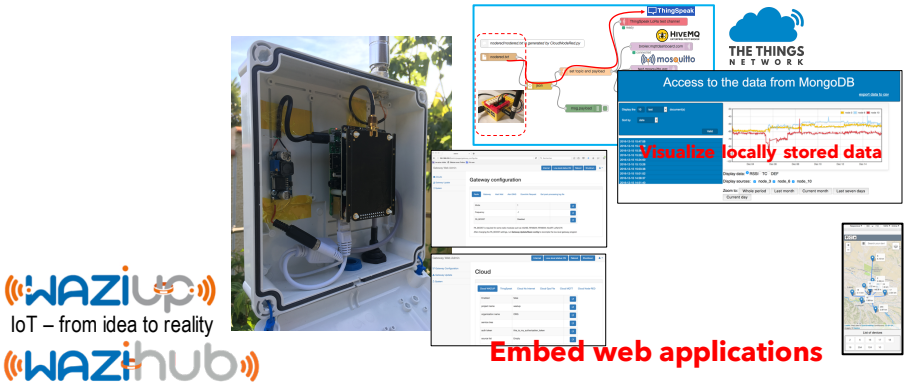

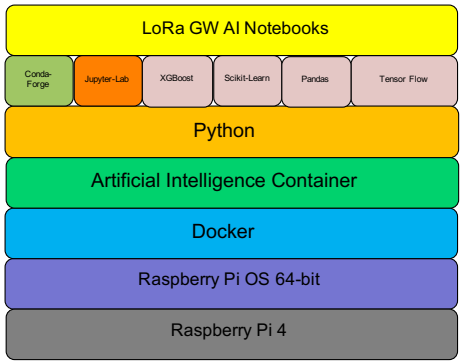

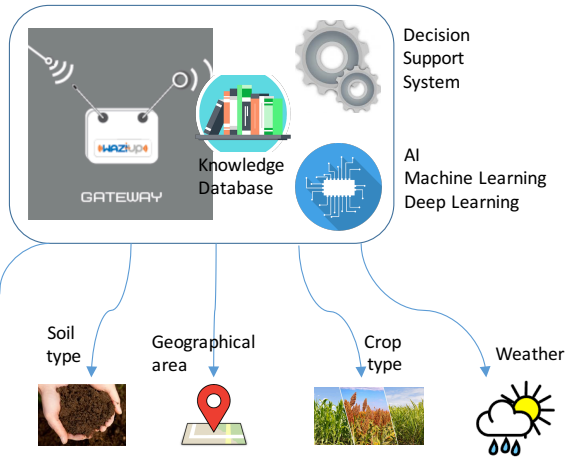

Integration of multiple knowledge streams

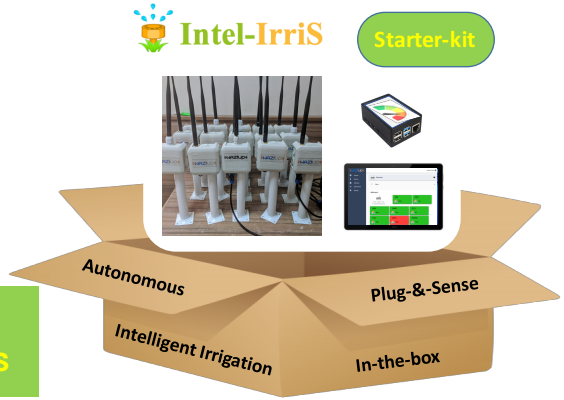

Fully autonomous

The IoT BackOffice

IoT usually means

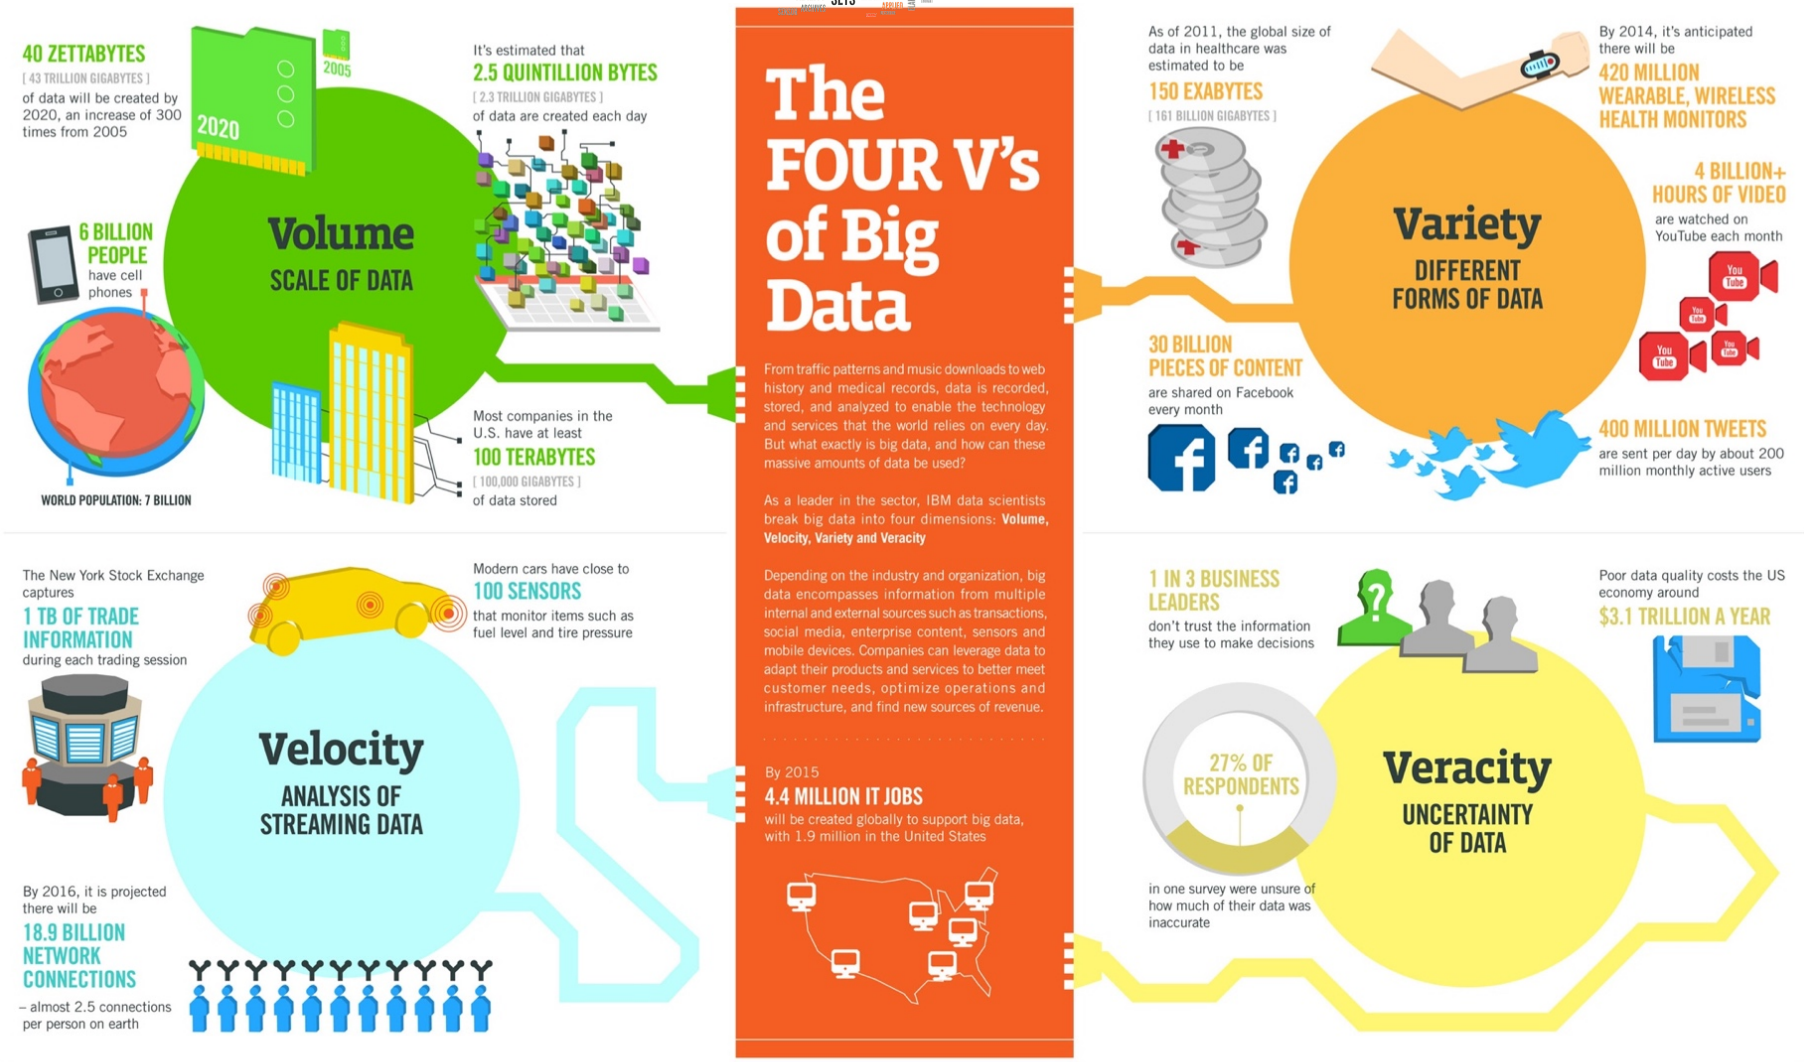

But also how to analyse the data

- What is the meaning of the collected data?
- Example with farming
  - What is interesting for farmers?
    - Fertility detection
    - Eating/Ruminating time for welfare
  - What data can be easily obtained?
    - accelerometer data with neck-mounted collar
  - How to detect relevant event from these data?

Advanced data analysis

Need of experts from the domain!

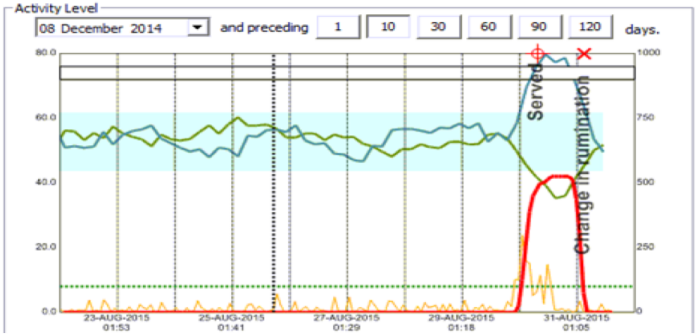

Fig. 3. Illustration of a rise in activity accompanied by a fall in rumination at the point of oestrus

Analysis techniques

- Traditional statistic methods still valid, and useful!

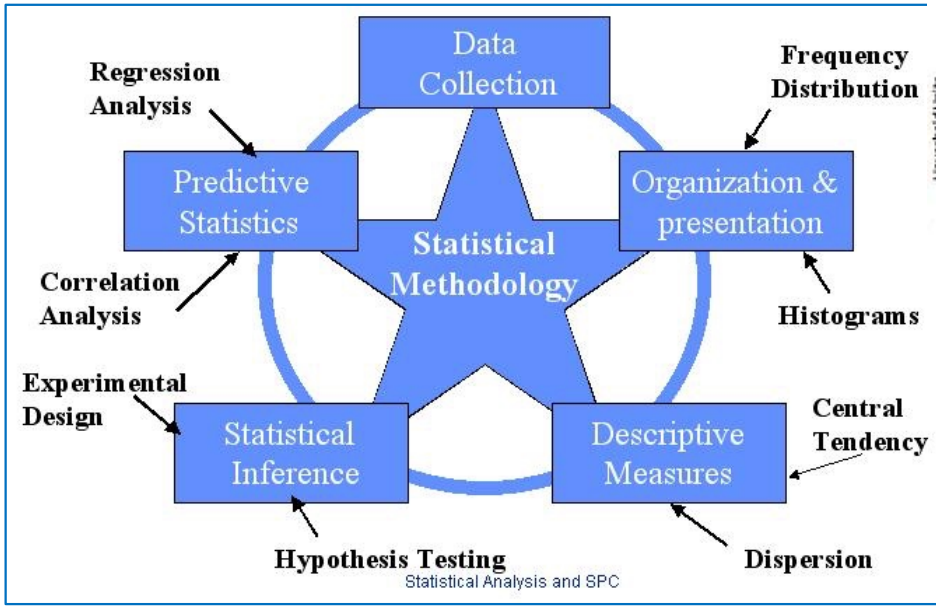

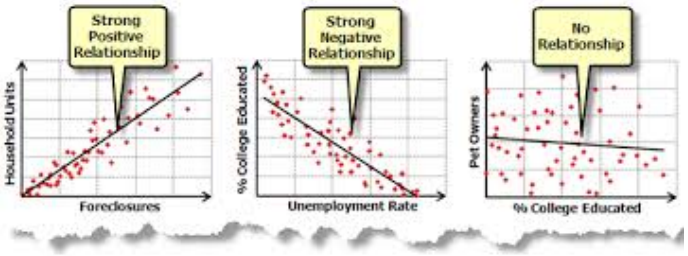

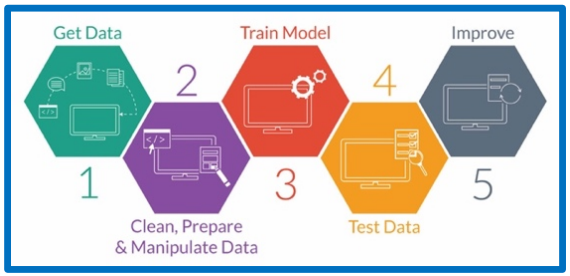

Clustering Analysis

- Definition
  - Grouping unlabeled data into clusters, for the purpose of inference of hidden structures or information
- Dissimilarity measurement
  - Distance : Euclidean(L_2), Manhattan(L_1), ...
  - Angle : Inner product, ...
  - Non-metric : Rank, Intensity, ...
- Types of Clustering
  - Hierarchical
    - Agglomerative or divisive
  - Partitioning
    - K-means, VQ, MDS, ...

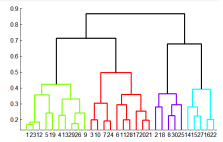

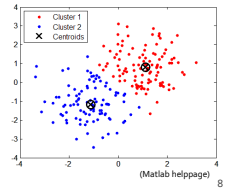

From Jong Youl Choi

Analysis techniques

- Traditional statistic methods still valid, and useful!

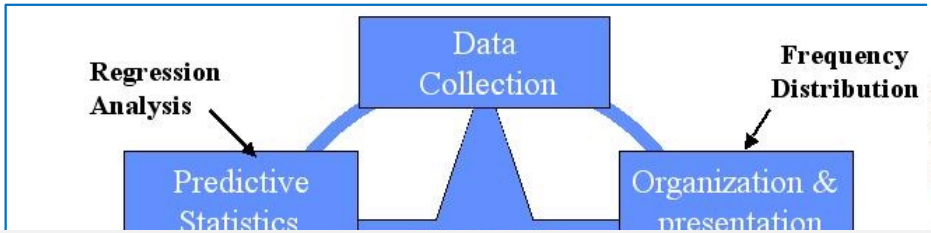

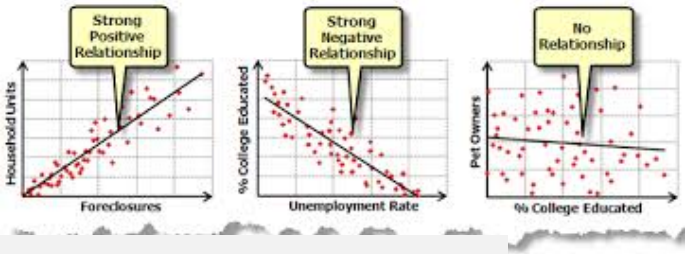

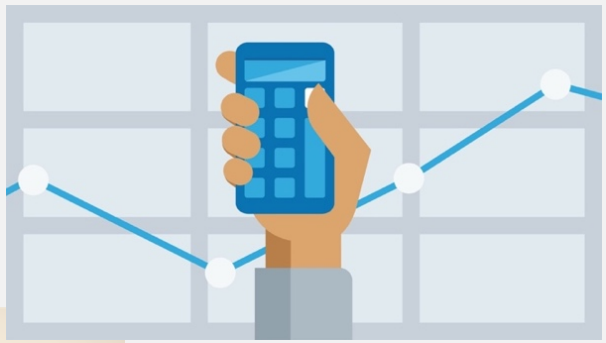

Going old school ?

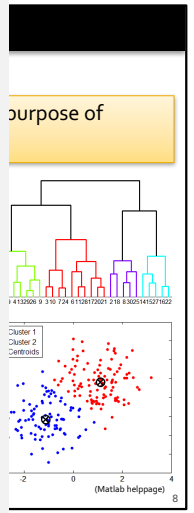

Use the full power of the Internet!

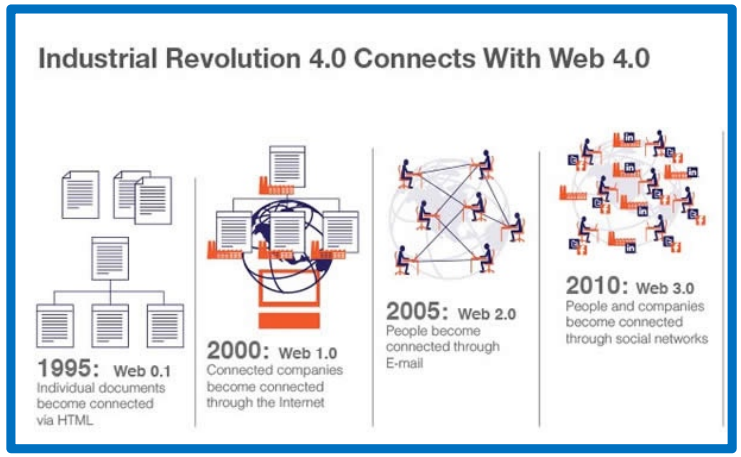

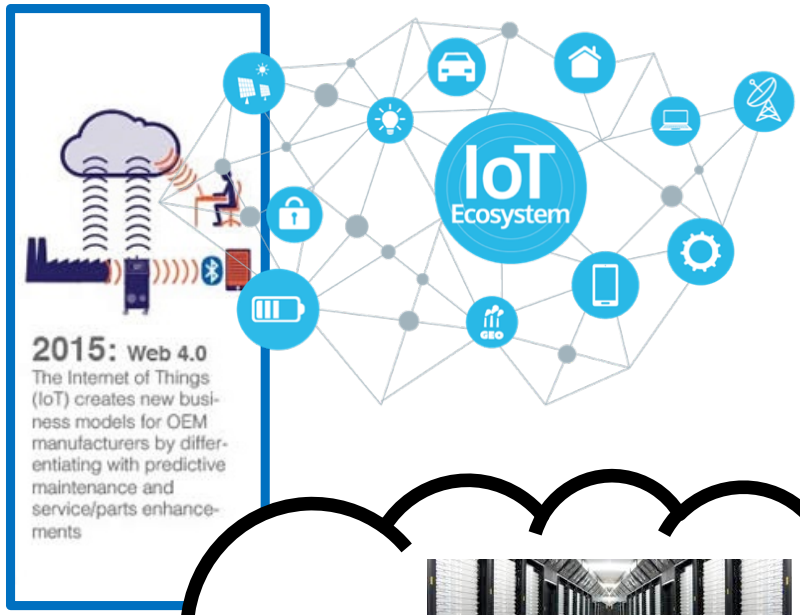

- IoT data are pushed on Internet data clouds
- Computing resources using Virtual Machines are obtained from Internet Computing clouds
- Parallel processing
- Optimized libraries
- Web tools to orchestrate

The raise of Artificial Intelligence

- It is the science and engineering of making intelligent machines.
- In Computer Science, Artificial Intelligence (**AI**) research is defined as the study of « intelligent agents »
- From General AI to Narrow AI: from overhyping to fewer promises, but more realistic!

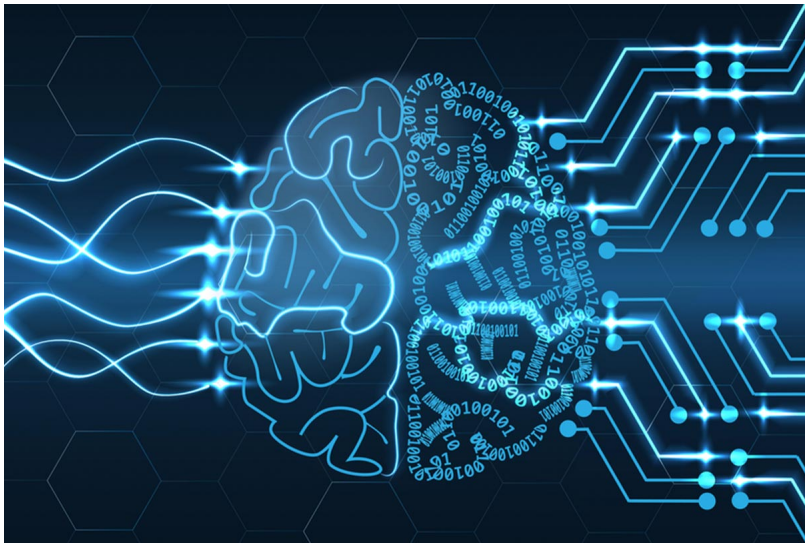

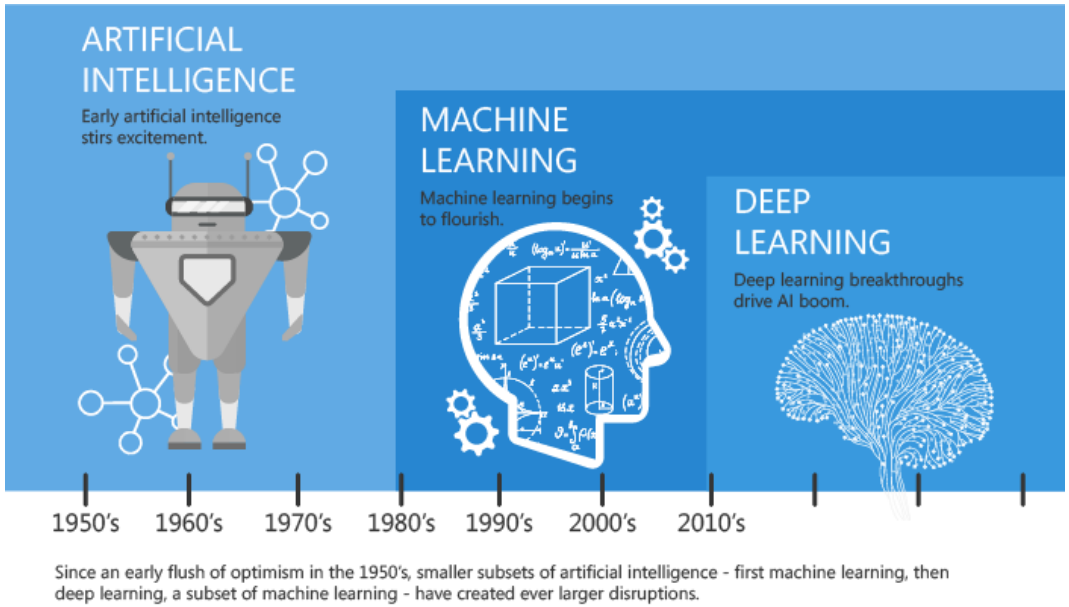

Since an early flush of optimism in the 1950's, smaller subsets of artificial intelligence - first machine learning, then deep learning, a subset of machine learning - have created ever larger disruptions.

AI: a serious science!

- General-purpose AI like the robots of science fiction is incredibly hard
  - Human brain appears to have lots of special and general functions, integrated in some amazing way that we really do not understand (yet)
- Special-purpose AI is more doable (nontrivial)
  - E.g., chess/poker playing programs, logistics planning, automated translation, speech and image recognition, web search, data mining, medical diagnosis, keeping a car on the road.

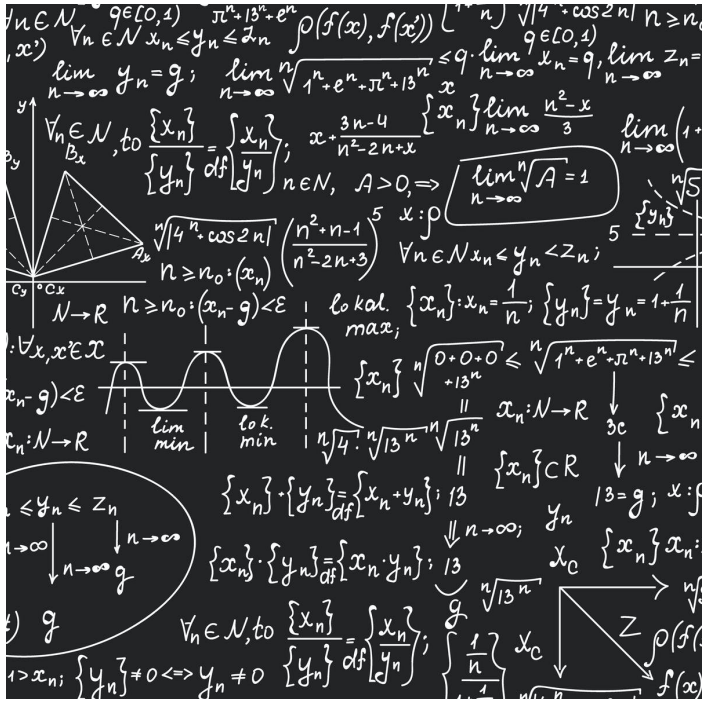

Prof. Congduc Pham
http://www.univ-pau.fr/~cpham

The Turing Test

- Proposed By Alan Turing in 1950
- To be called intelligent, a machine must produce responses that are indistinguishable from those of a human.
- Human judge communicates with a human and a machine over text-only channel.
- Both human and machine try to act like a human.
- Judge tries to tell which is which.
- Is Turing Test the right goal?

  "Aeronautical engineering texts do not define the goal of their field as making 'machines that fly so exactly like pigeons that they can fool even other pigeons.'" [Russell and Norvig]

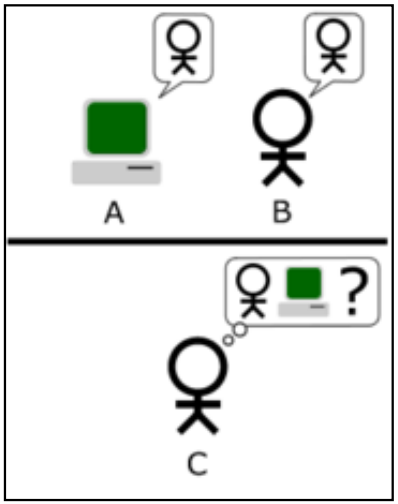

Reflection

if AI can be more **rational** than humans in some cases, why not?

Systems that think like humans	Systems that think rationally
Systems that act like humans	Systems that **act rationally**

AI focus on **action.** Avoids philosophical issues such as "is the system conscious" etc.

Premise

- We will mostly follow "**act rationally**" approach
- Distinction may not be that important
- Acting rationally like a human presumably requires (some sort of) thinking rationally like a human
- Humans much more rational anyway in complex domains

AI Technologies

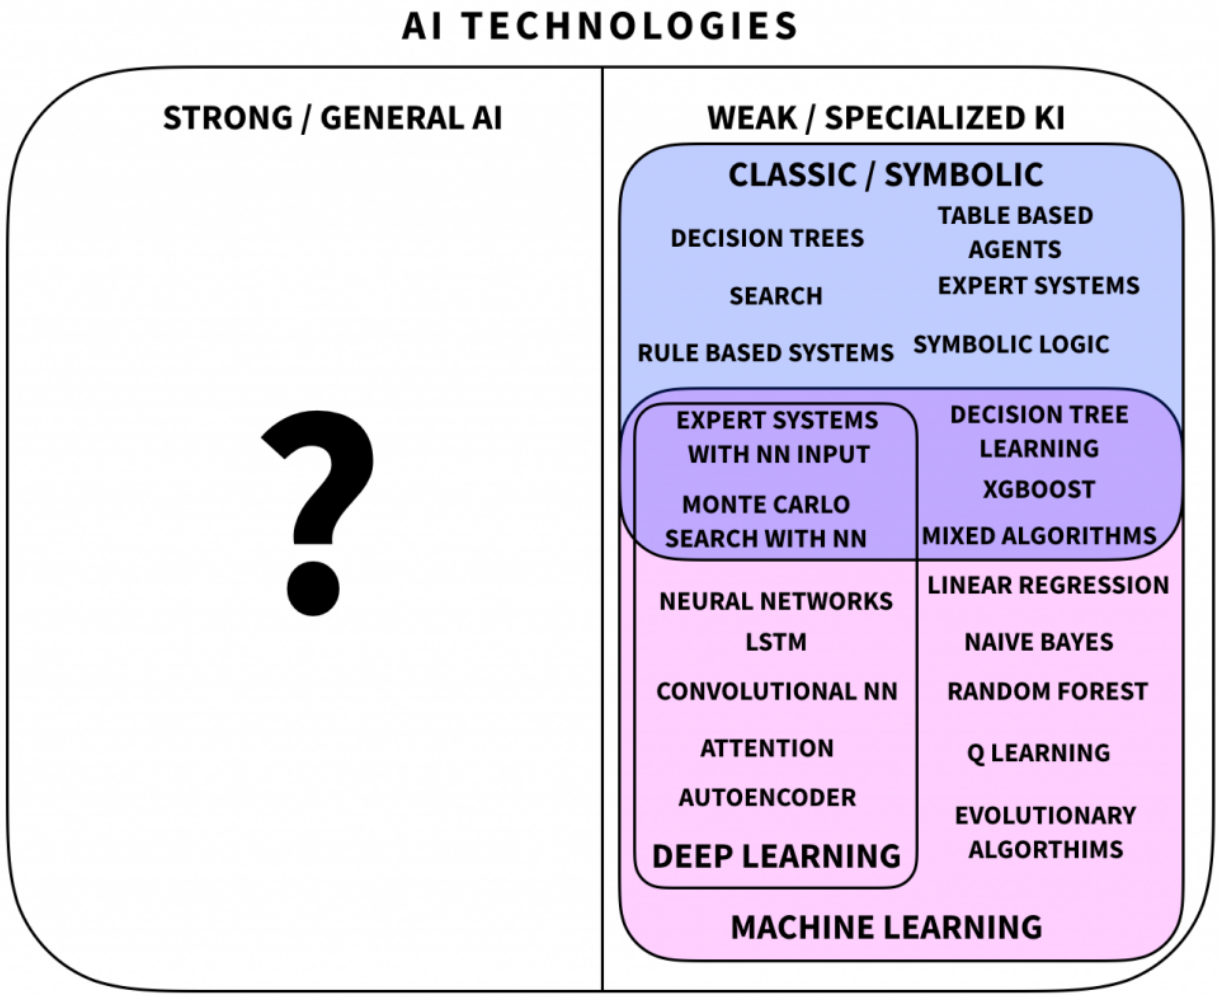

Machine Learning

- Develops Narrow Artificial Intelligence systems through examples
  - A developer creates a model and then “trains” it by providing it with many examples
  - The machine learning algorithm processes the examples and creates a mathematical representation of the data that can perform prediction and classification tasks
- Example
  - A machine-learning algorithm trained on thousands of bank transactions with their outcome (legitimate or fraudulent) will be able to predict if a new bank transaction is fraudulent or not

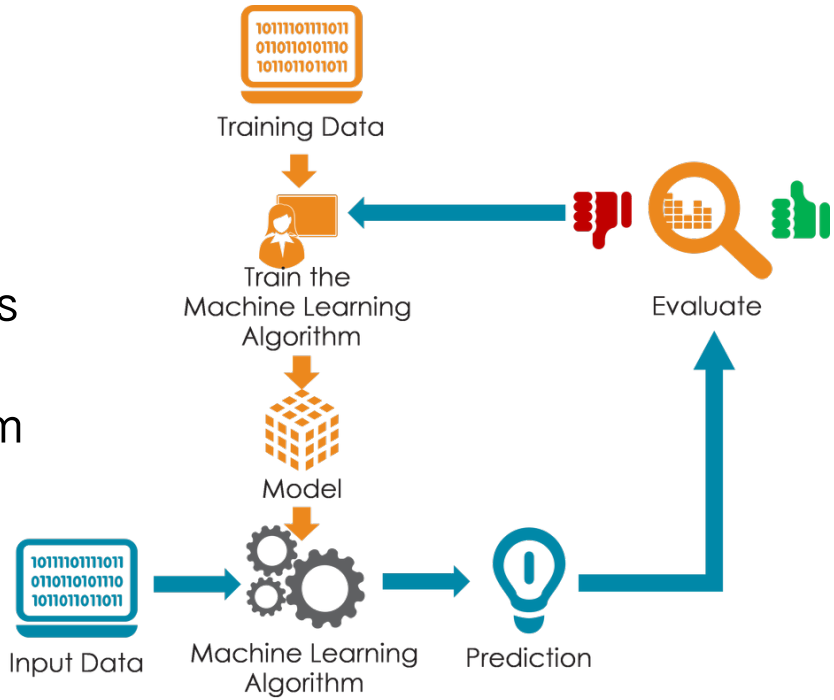

Machine Learning Techniques

Optimize a performance criterion using example data or past experience

A lot of statistic methods

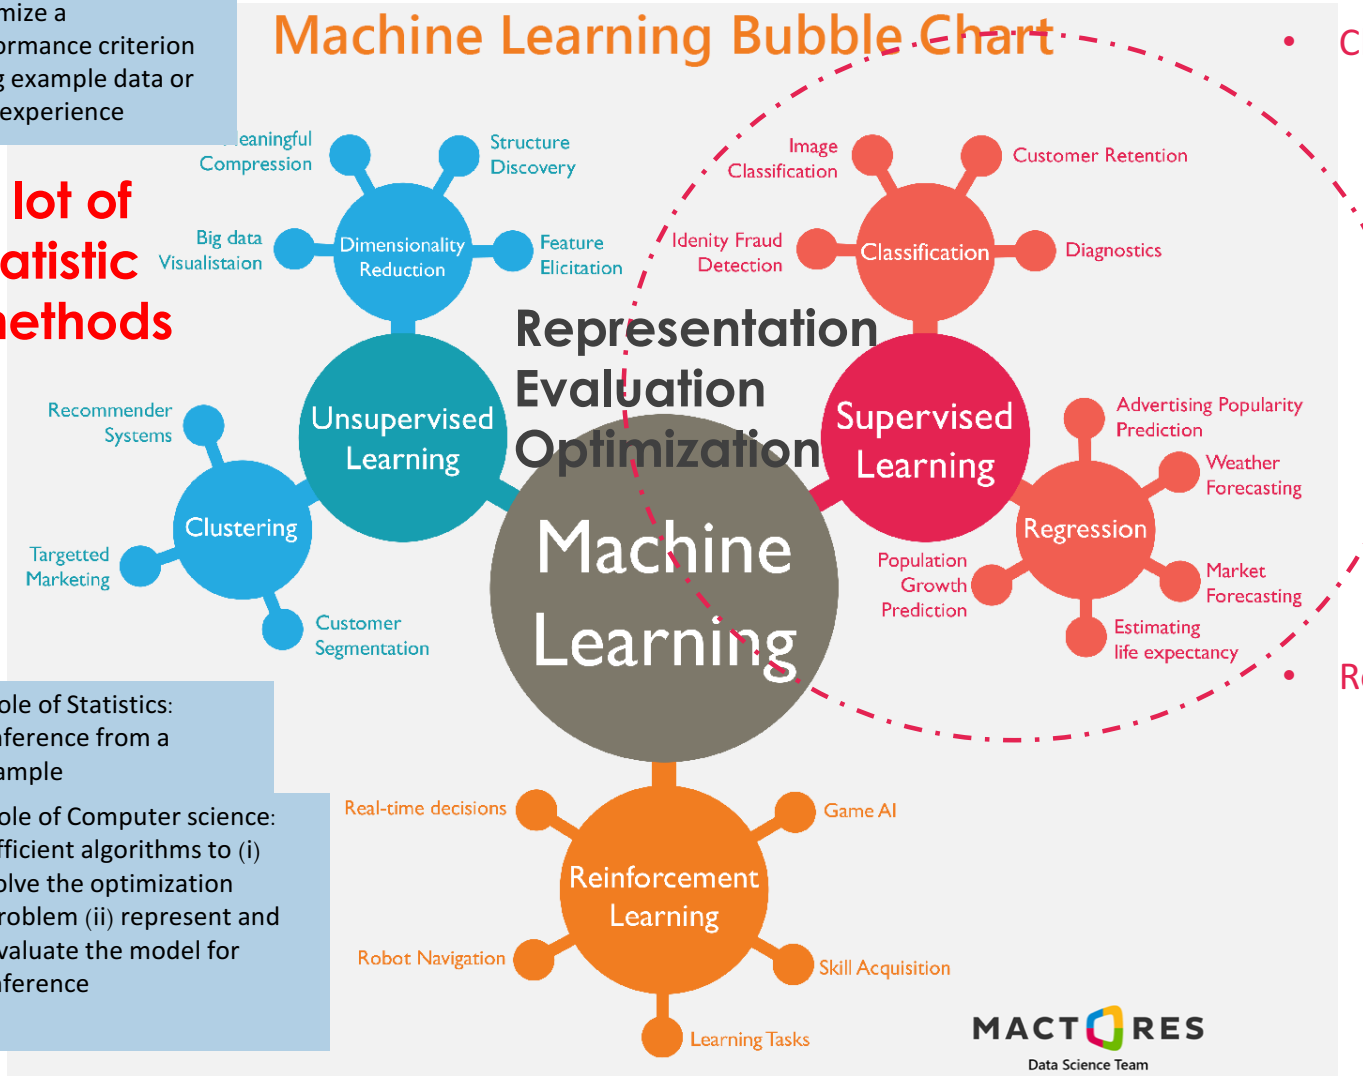

Role of Statistics: Inference from a sample

Role of Computer science: Efficient algorithms to (i) solve the optimization problem (ii) represent and evaluate the model for inference

- Classification
  - Logic
  - SVM
  - Random Forest
  - Hidden Markov
  - …
- Regression
  - Lasso
  - Ridge
  - Loes
  - KNN
  - Spline
  - XGBoost
  - …

Supervised Learning

- ML model is presented with *input data* which is labeled
  - Each *input data* is tagged with the correct label.
- The goal is to approximate math operations in the ML model so well that when presented with new *input data*, the ML model can **predict** the output variables for that *input data*.

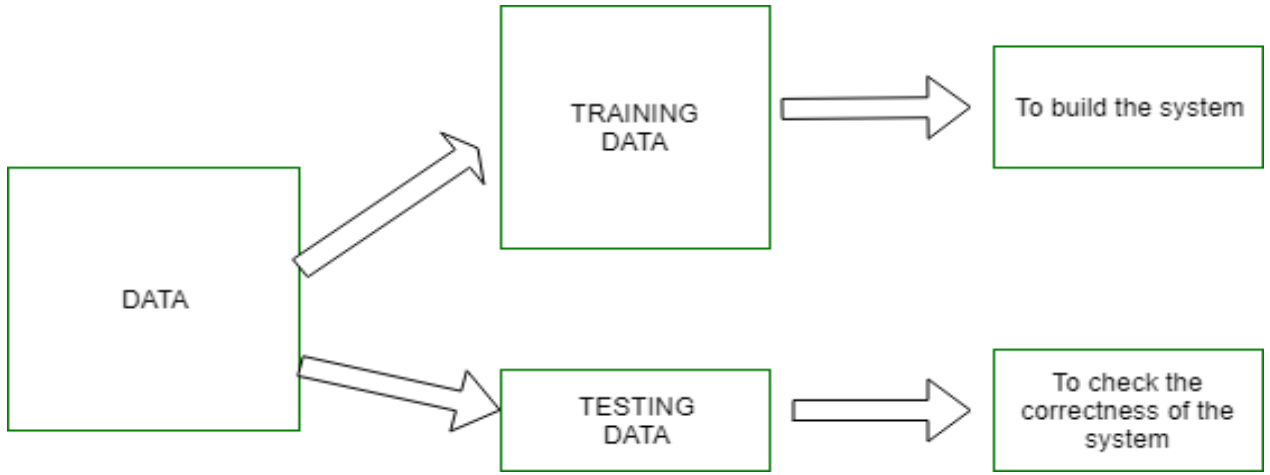

Spam Mail Example

- On the left side of the image, some data is marked as ‘Spam’ or ‘Not Spam’. This is *labeled data.* This data is used to train the supervised model, the *intelligent* program (at center of the image).
- Trained model is tested with new mails (on the top of the image) and checking if the output of the supervised model is correct (on the right side of the image).

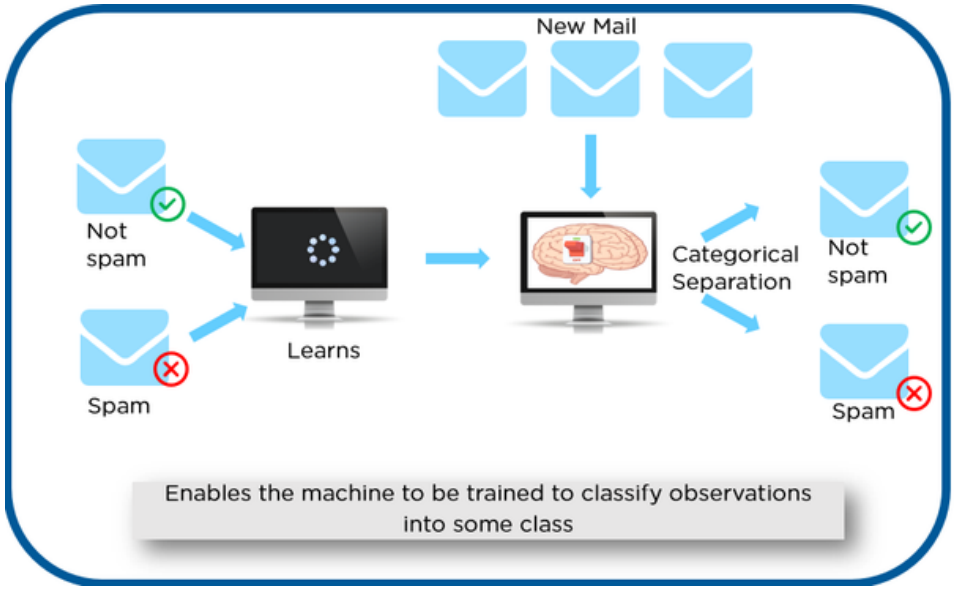

Types of Supervised Learning

- **Classification**: A classification problem is when the output is a category, such as “red” or “blue” or “disease” and “no disease”.
- **Regression**: A regression problem is when the output is a real number, such as “dollars” or “weight”.

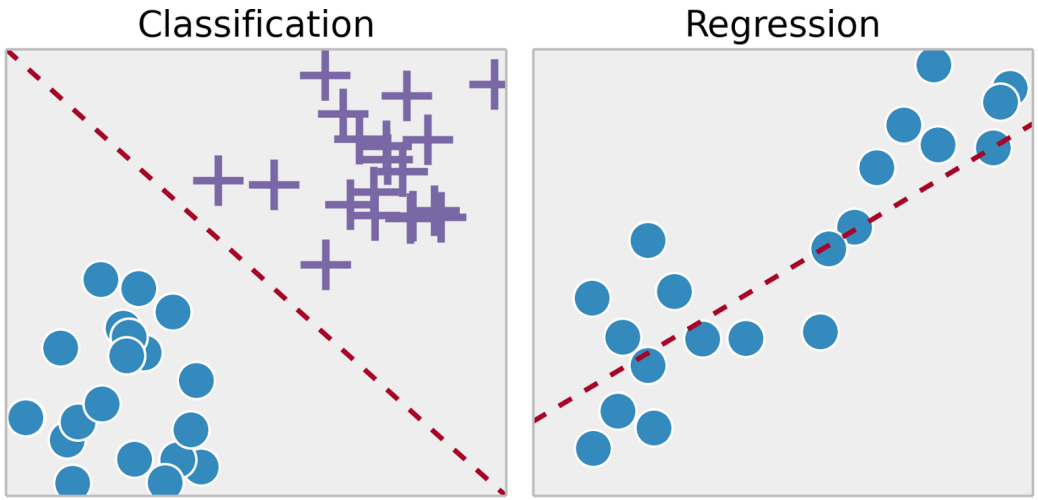

Unsupervised Learning

- ML model is presented with unlabeled, uncategorized data
- ML model acts on the data without prior training.
- The output is dependent upon the coded algorithms.
- Is one way of testing AI.

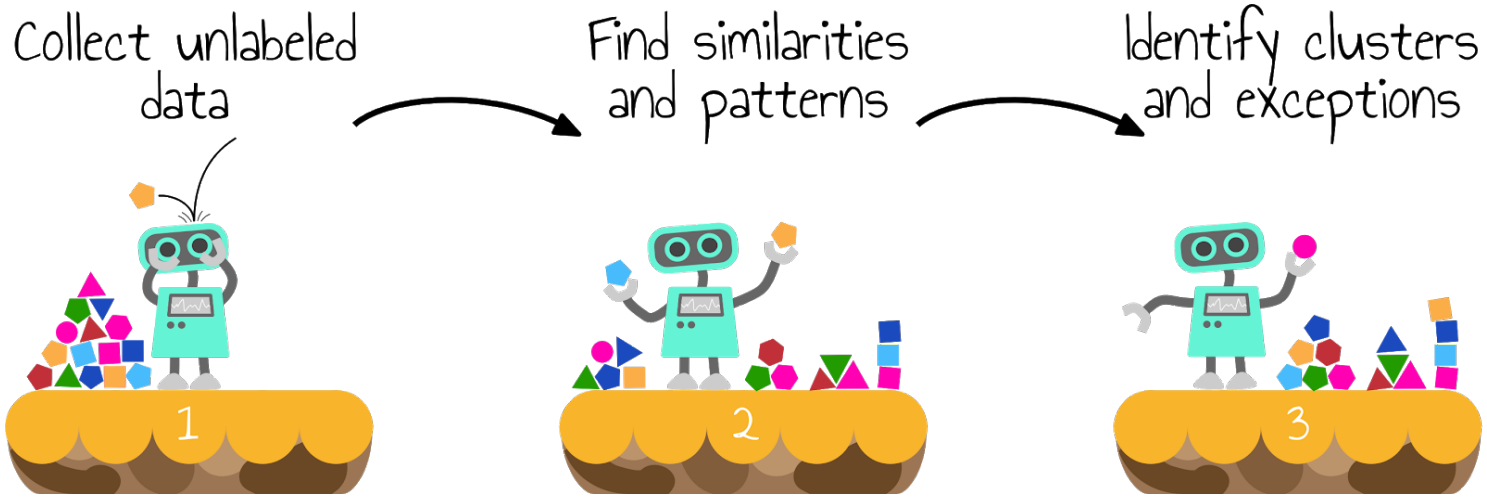

Ducks Example

- In the below example, some cartoon characters are passed to the ML model. Some of them are ducks.
- No data label provided.
- ML model is able to separate the characters into ‘Duck’ and ‘No duck’ by looking at the type of data and models in the underlying data structure.

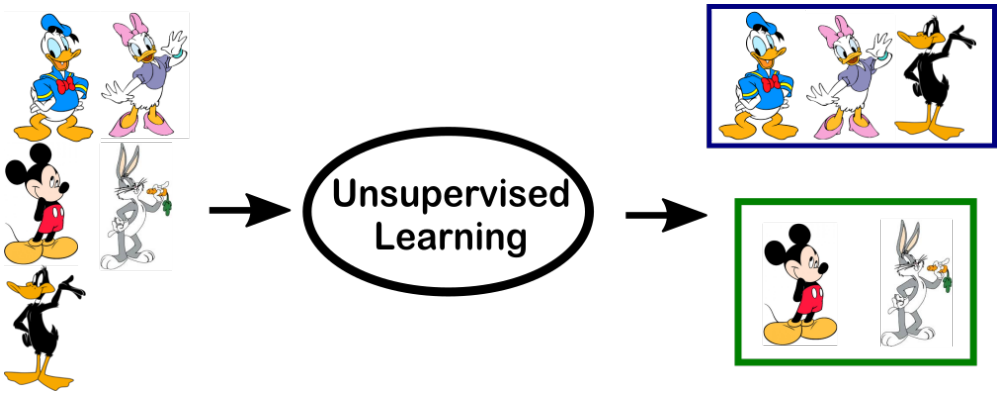

Types of Unsupervised Learning

- **Clustering**: Discovering the inherent groupings in the data, such as grouping customers by purchasing behavior.
- **Association**: Discovering rules that describe large portions of the input data, such as people that buy X also tend to buy Y.

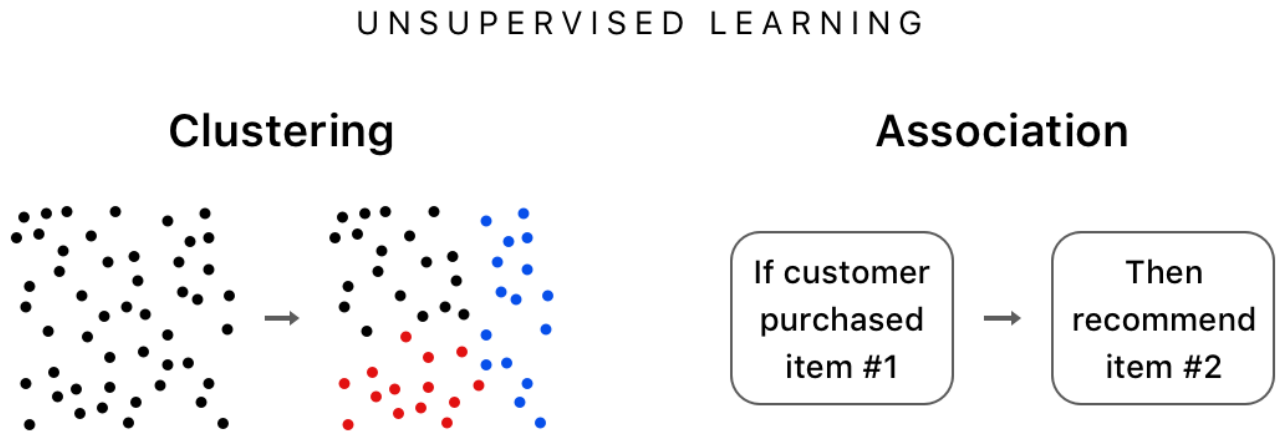

Reinforcement Learning

- A reinforcement learning algorithm, or agent, learns by interacting with its environment.
- The agent receives rewards by performing correctly and penalties for performing incorrectly.
- The agent learns without intervention from a human by maximizing its reward and minimizing its penalty.
- It is a type of dynamic programming that trains algorithms using a system of reward and punishment.

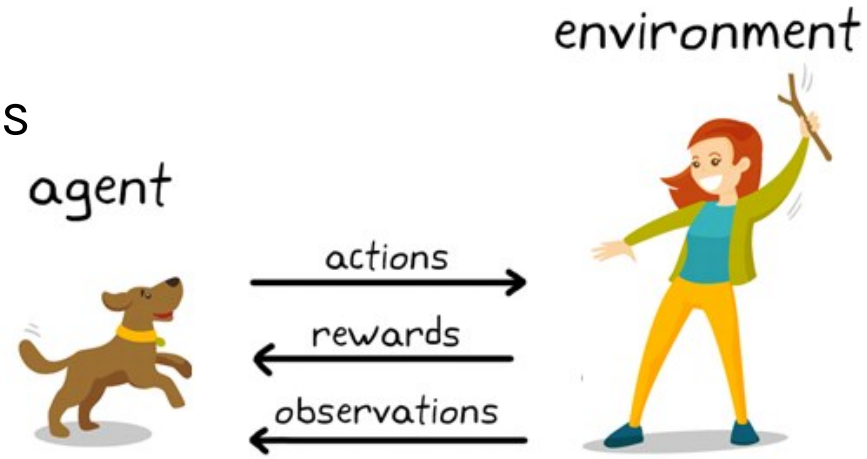

Robot Path Example

- The agent is given 2 options i.e. a path with water or a path with fire.
- If agent takes the fire path then a penalty is subtracted
- Agent learns it should avoid the fire paths.
- If agent takes water path then some a reward is granted
- Agent learns what path is safe and what path isn't.

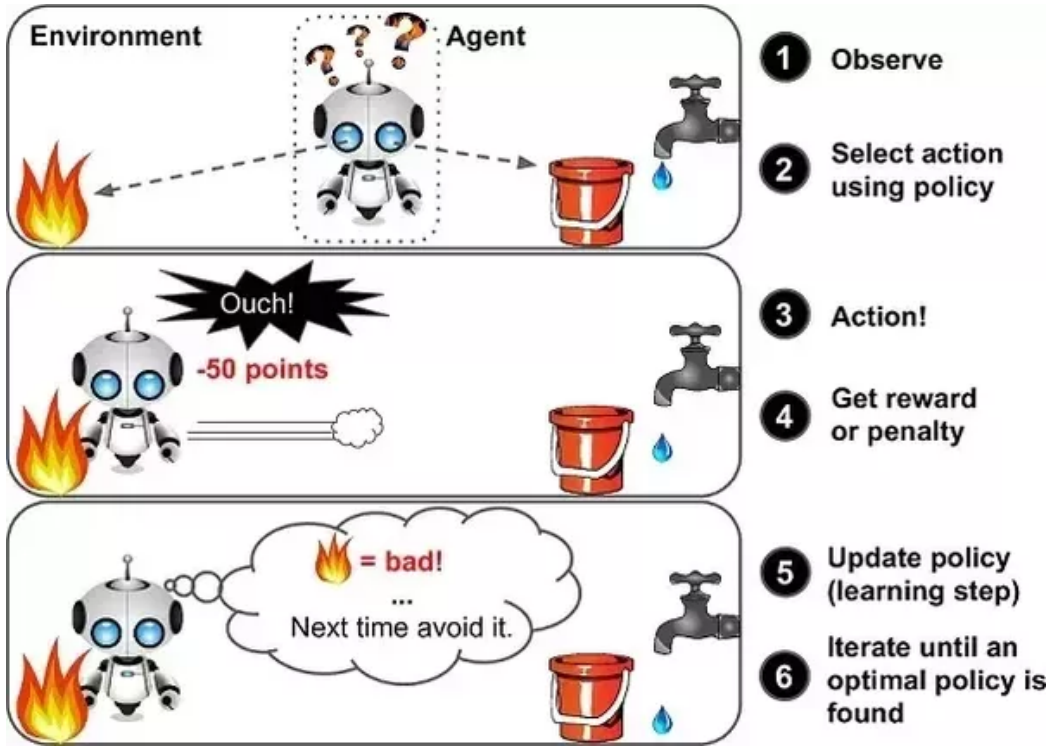

What is Regression?

- Regression takes a group of random events, thought to be predicting another event, and tries to find a mathematical relationship between them.
- Events are represented by variables.
- Used in statistics to find trends in data.
- Implemented in Machine Learning (**ML**) as supervised algorithm where the predicted output is expressed in Real numbers.
- Used to predict values within a continuous range.

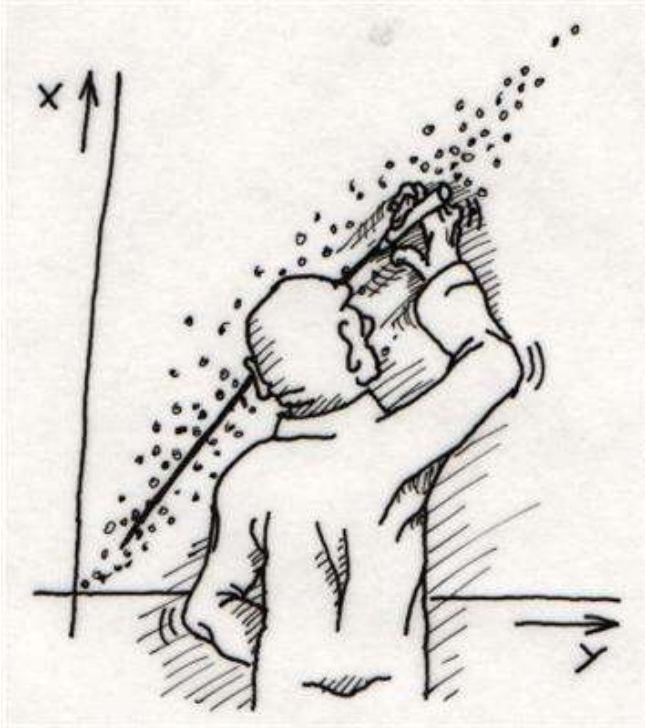

Regression

- **Dependent variables:** the main event or factor to understand or predict. Also known as *explanatory variable.*
- **Independent variables:** the events or factors suspected to have an impact on the dependent variable. Also known as *response variable.*

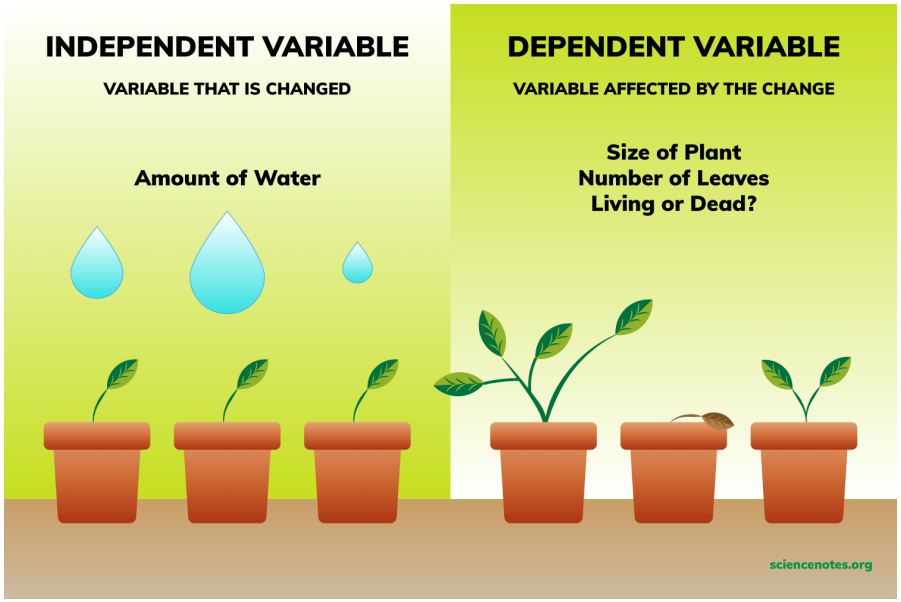

Types of Regression

- **Simple regression:** single independent variable for a single dependent variable. It is very common to name the independent variable as x and **Y** as the dependent variable
  x**:** number of cricket chirps
  Y: temperature
- **Multivariable regression:** multiple independent variables, x_1, x_2, x_3, for a dependent variable **Y**.
  x_1: number of cricket chirps
  x_2: rainfall
  x_3: automobile traffic
  Y: temperature

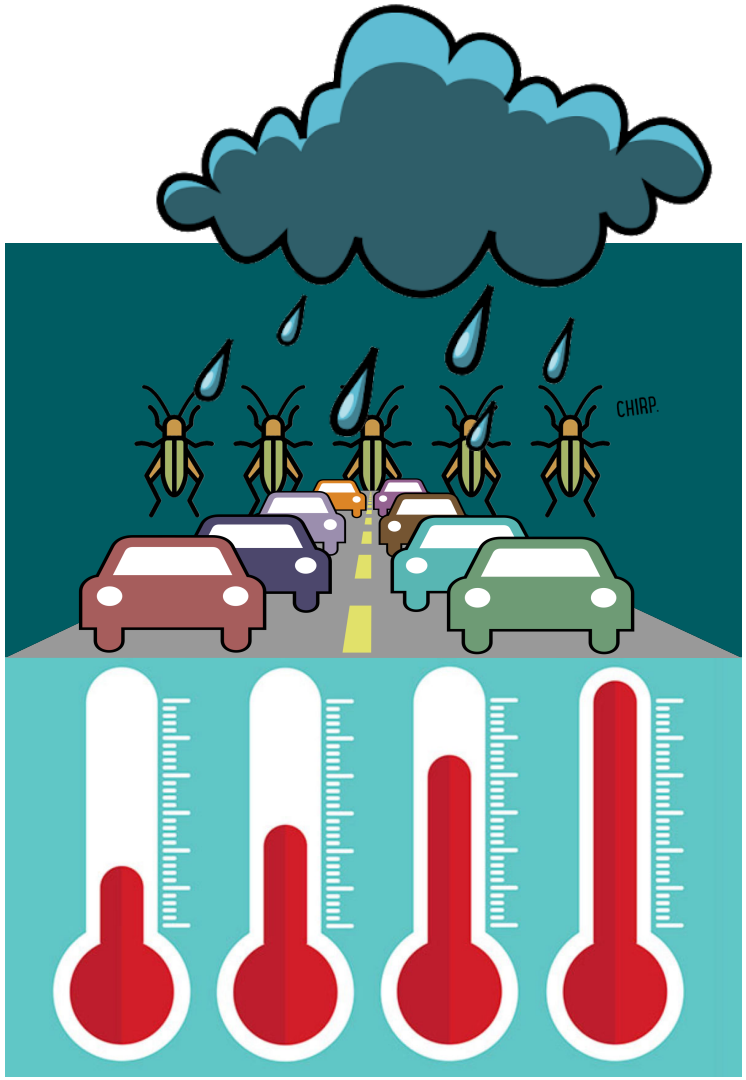

Scatter Plot

- Data gathering on the variables in question
- The vertical scale represents one set of measurements and the horizontal scale the other

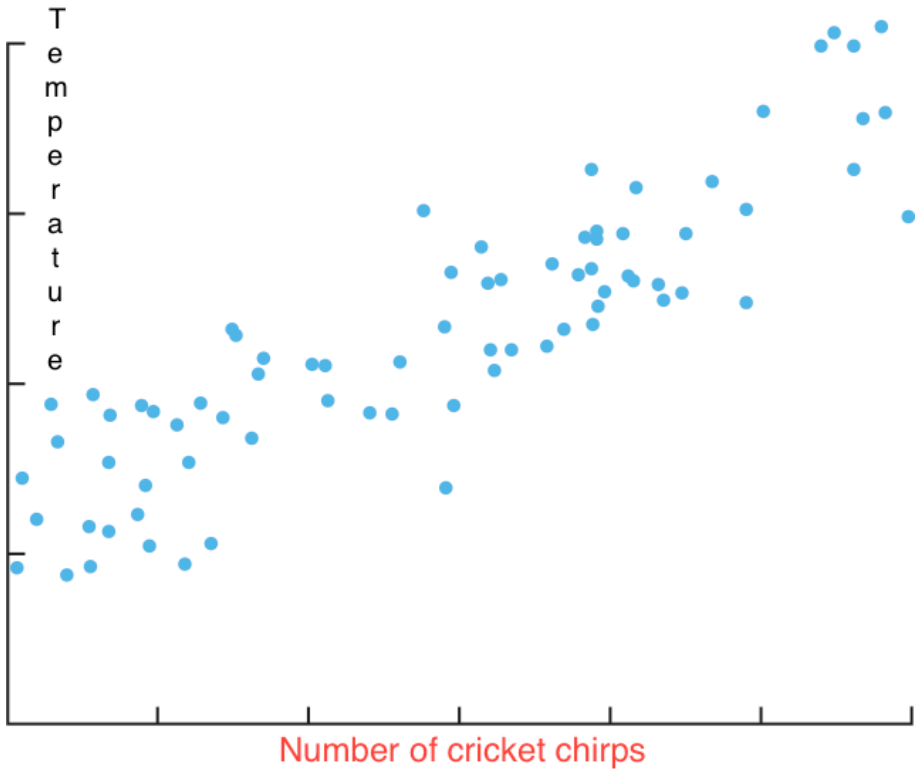

Correlation Coefficient

- Measures the degree of association, denoted by r or R, sometimes called Pearson's correlation.
- Measured on a scale that varies from + 1 through 0 to – 1.
- When one variable increases as the other increases the correlation is *positive*
- When one variable decreases as the other increases the correlation is *negative*.
- Complete correlation between two variables is expressed by either + 1 or -1.
- Complete absence of correlation is represented by 0.
- If a curved line is needed to express the relationship, other and more complicated measures of the correlation must be used.

Correlation Chart

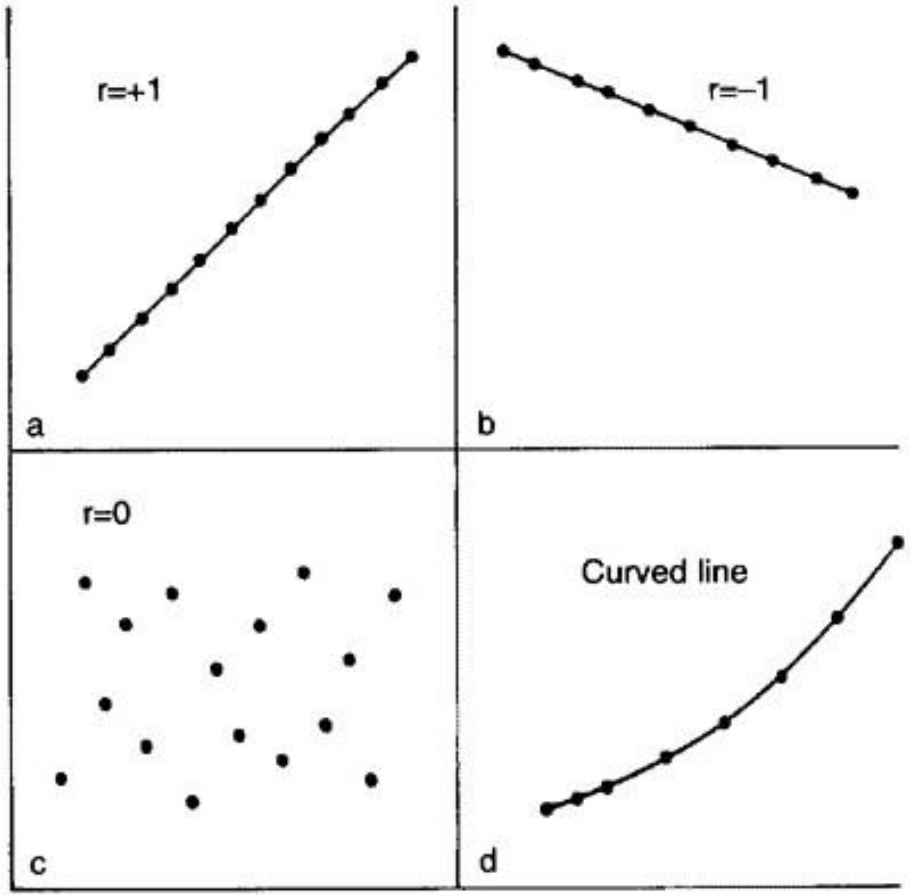

Linear Regression

- A linear relationship to predict the (average) numerical value of **Y** for a given value of x using a straight line, called the ***regression line***.
- Knowing the *slope* and the **y**-intercept of that regression line, it is possible to plug in a value for x and *predict* the average value for **Y**. In other words, predict the average **Y** from x.

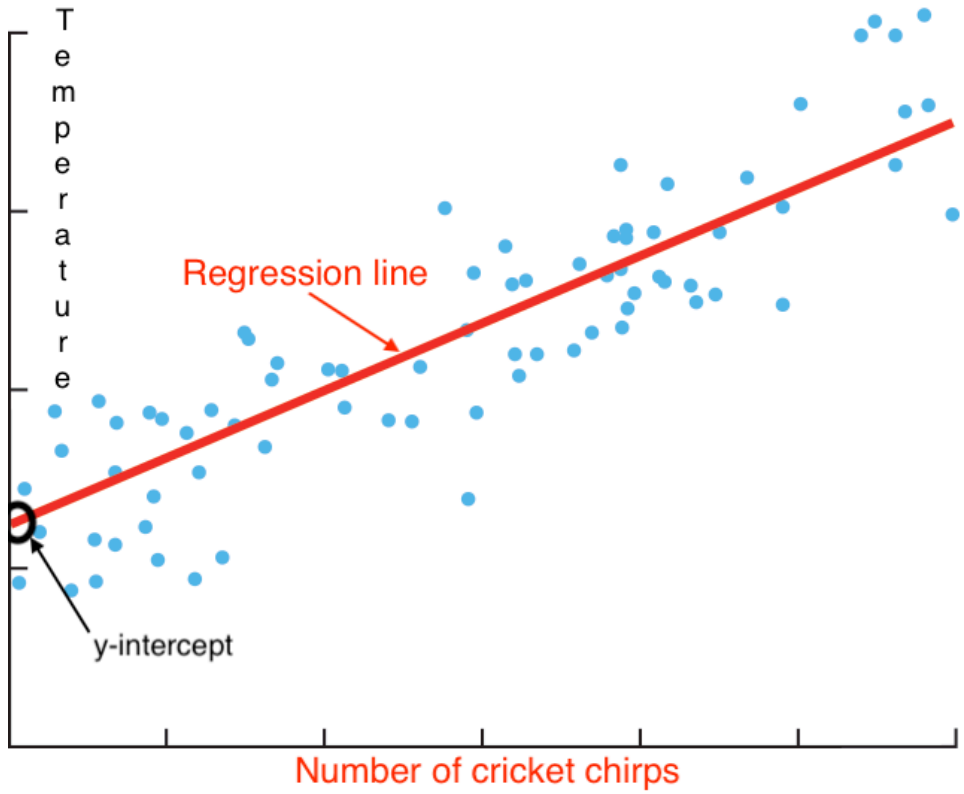

Beyonds Machine Learning?

- Combines advances in computing power and special types of **Neural Networks** to learn complicated patterns in large amounts of data
- State of the art for identifying objects in images and words in sounds
- Applied successes in pattern recognition to more complex tasks such as automatic language translation, medical diagnoses and numerous other important social and business problems

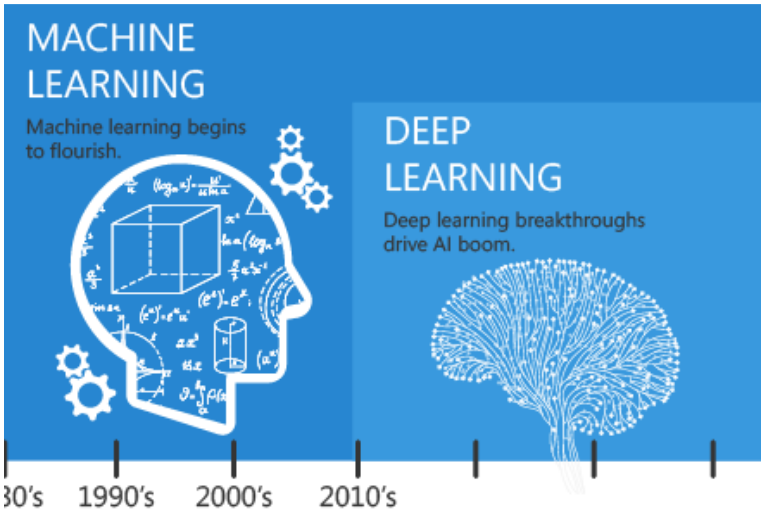

s, smaller subsets of artificial intelligence - first machine learning, then - have created ever larger disruptions.

Neural Networks: the Perceptron

- Mathematical representation of a biological neuron
- First implementation by Frank Rosenblatt in the 1950s
- Rosenblatt's perceptron is activated when there is sufficient stimuli or input. (Neurons have been found to perform a similar process, in which experience strengthens or weakens dendrites' connections)

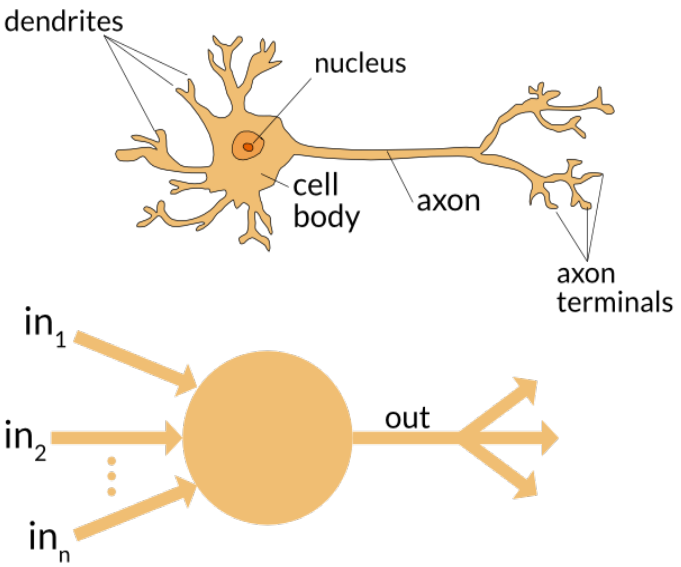

Neural Networks

- Neurons by themselves are kind of useless, in large groups, they work together to create some serious magic!
- Neural Networks are no more than a **stacking** of multiple *perceptrons* in layers to produce an output.
- Input into one layer that creates an output which in turn becomes the input for the next layer, and so on. This happens until the final output signal.

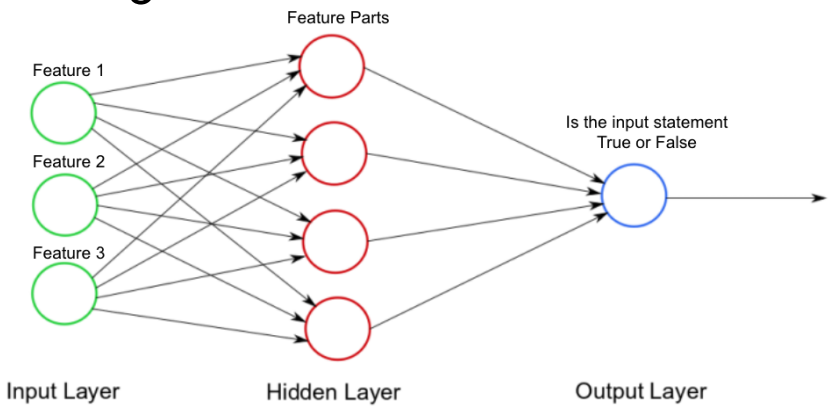

Deep Learning is essentially NN

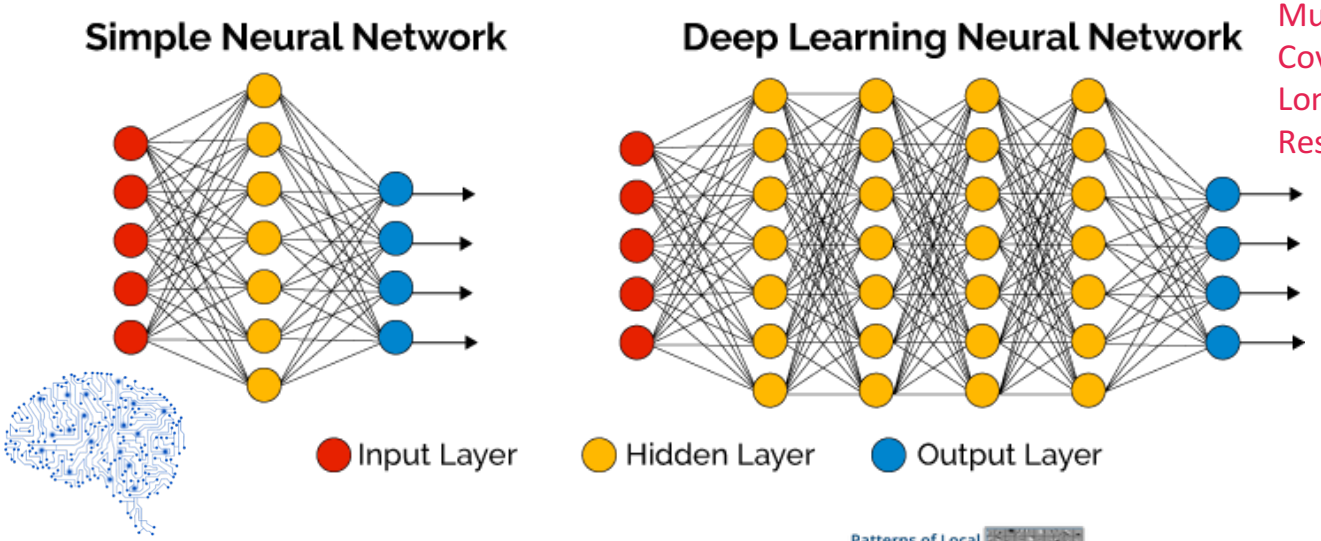

Multilayer Perceptron
Covolutional Neural Nets
Long Short-Term Memory
Restricted Boltzmann Machine

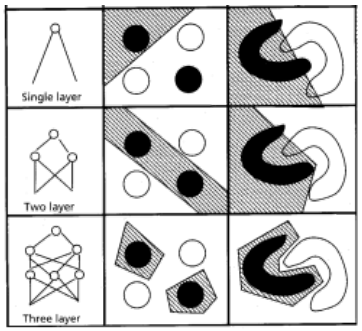

(Jain, 1996)

Voice/Face/Patterns recognition on many platforms

- **Facebook**
- **Google Photos**
- **Twitter**
- **Siri**
- **…**

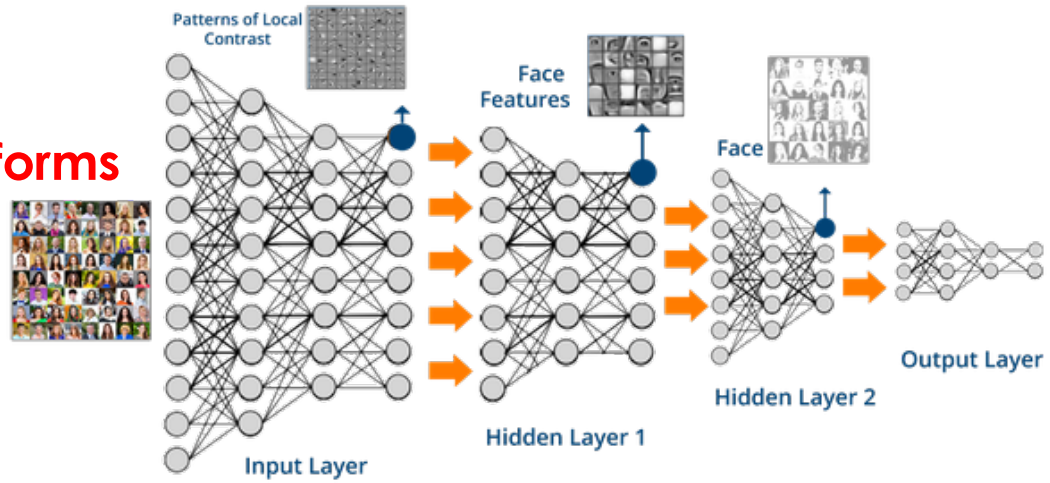

Machine/Deep Learning for scientists

- Large variety of supported languages
  - Python, R, C++, Java, Scala, Javascript, Go, ...
- Many statistical methods/algorithms are implemented in libraries
- Examples
  - Scikit-learn
  - Google TensorFlow
  - Microsoft Distributed Machine Learning Toolkit
  - Apache Mahout
  - ...
- But, beware
  - There are hundredth of tools...
  - ...and new tools every months!

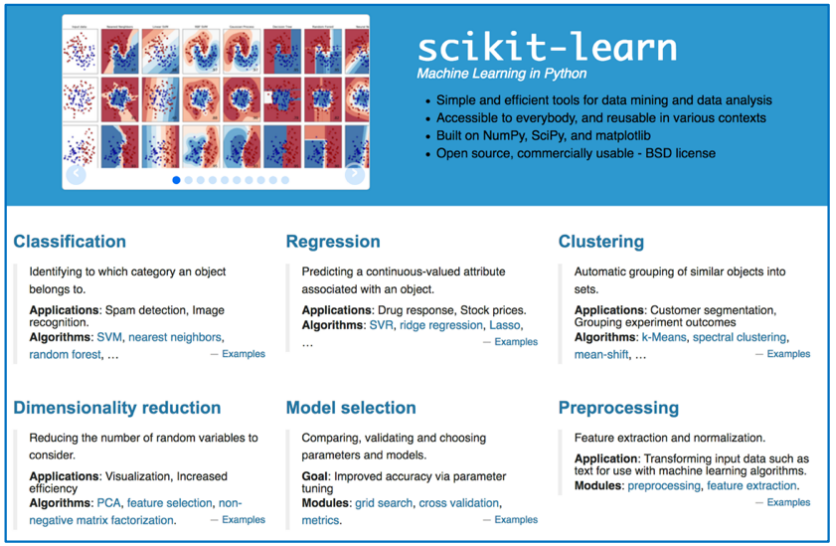

WAZIUP Online Course

Fundamentals of IoT

Continue with
F-IOT-2b: Introduction to IoT hardware

WAZIUP
IoT – from idea to reality
WAZIhub