MACHINE LEARNING

15CS73, VTU CBCS SCHEME

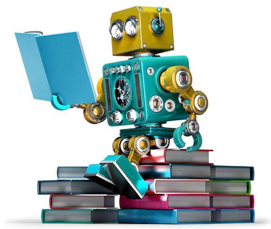

BY
HARIVINOD N
DEPT OF CSE
VIVEKANANDA COLLEGE OF ENGINEERING TECHNOLOGY, PUTTUR

What is Learning?

MACHINE LEARNING

Learning - improve automatically with experience

Using past experiences to improve future performance.

How we learn?

- Rote Learning (memorization)
  - Memorizing things without knowing the concept/ logic behind them
- Passive Learning (instructions)
  - Learning from a teacher/expert.
- Analogy (experience)
  - Learning new things from our past experience.
- Inductive Learning (experience)
  - On the basis of past experience formulating a generalized concept.
- Deductive Learning
  - Deriving new facts from past facts.

Why Machine Learning?

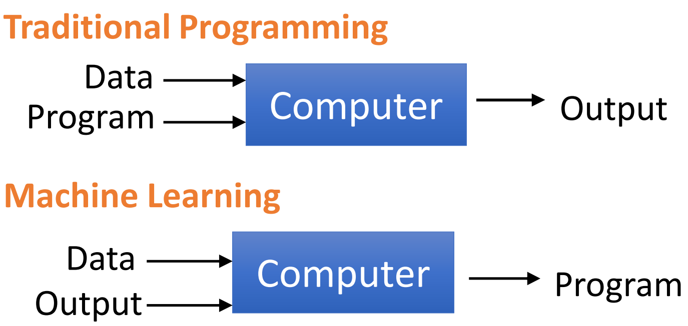

What is Machine Learning?

- General definition:
  - Machine Learning is the field of study that gives computers the ability to learn without being explicitly programmed. Arthur Samuel, *1959*
- And a more engineering-oriented one:
  - A computer program is said to learn
    from **experience E**
    with respect to **some task T** and
    some **performance measure P**,
    if its performance on T, as measured by P**, improves with experience E**.
  - Tom Mitchell, *1997*

What is Machine Learning

- Machine learning provides systems,
  the ability to automatically learn and
  improve from experience
  without being explicitly programmed.
- Machine learning focuses on
  the development of computer programs
  that can access data
  and use it learn for themselves.

What is Machine Learning

- The process of learning begins with
  - observations or data, such as examples,
  - direct experience, or instruction, in order to look for patterns in data
  - and make better decisions in the future based on the examples that we provide.

- The primary aim is to allow the computers learn automatically without human intervention or assistance and adjust actions accordingly.

Ex1: Spam Filtering

Decide which emails are spam and which are important.

Not spam | Supervised classification | spam

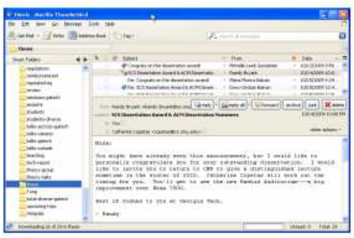

Goal: use emails seen so far to produce good prediction rule for future data.

Ex1: Spam Filtering

Represent each message by features. (e.g., keywords, spelling, etc.)

	"money"	"pills"	"Mr."	bad spelling	known-sender	spam?	
	Y	N	Y	Y	N	Y	
	N	N	N	Y	Y	N	
	N	Y	N	N	N	Y	
example	Y	N	N	N	Y	N	label
	N	N	Y	N	Y	N	
	Y	N	N	Y	N	Y	
	N	N	Y	N	N	N	

Reasonable RULES:

Predict SPAM if unknown AND (money OR pills)

Predict SPAM if 2money + 3pills –5 known > 0

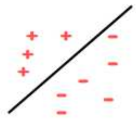

Linearly separable

Ex1: Spam Filtering

- Spam filter can learn to **flag spam** given examples of spam emails (e.g., flagged by users) and examples of regular (nonspam, also called "ham") emails.
- The examples that the system uses to learn are called the *training set*. Each training example is called a *training instance* (or *sample*).
- In this case,
  - the **task T** is to flag spam for new emails,
  - the **experience E** is the training data, and
  - the **performance measure P** needs to be defined; Ex: ratio of correctly classified emails.

Ex2: A checkers learning problem

MACHINE LEARNING

- Task T - Playing checkers
- Performance Measure P - Percentage of games won against opponent
- Training Experience E - Playing practice games against itself

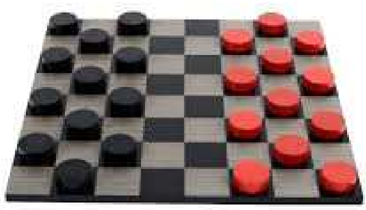

Ex3: A handwriting recognition learning problem

MACHINE LEARNING

- Task T: recognizing and classifying handwritten words within images
- Performance measure P: percent of words correctly classified
- Training experience E: a database of handwritten words with given classifications

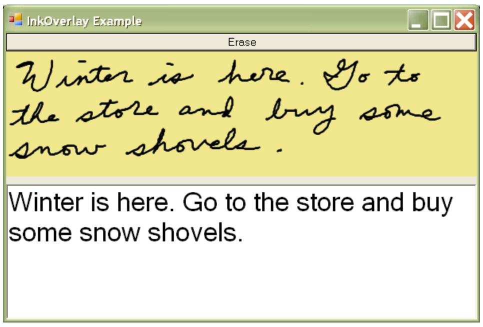

Ex4: A robot driving learning problem

MACHINE LEARNING

- T: driving on public 4-lane highways using vision sensors
- P: average distance traveled before an error (as judged by human overseer)
- E: a sequence of images and steering commands recorded by observing a human driver

What is Machine Learning?

MACHINE LEARNING

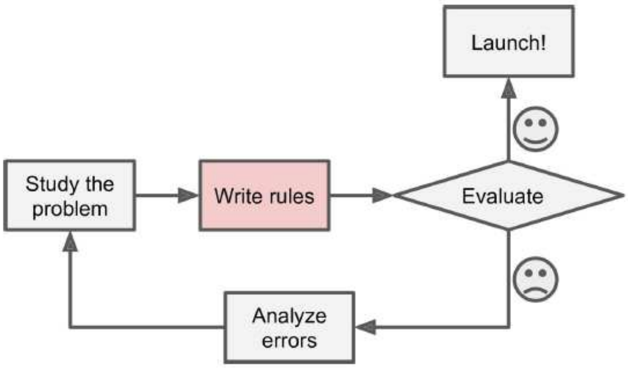

Traditional Approach

What is Machine Learning?

MACHINE LEARNING

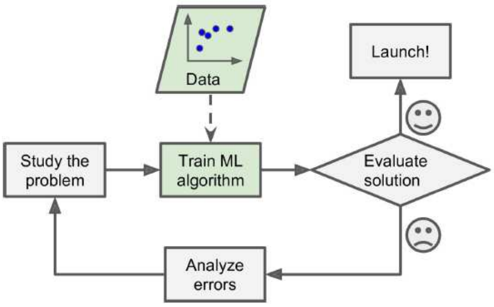

Machine Learning Approach

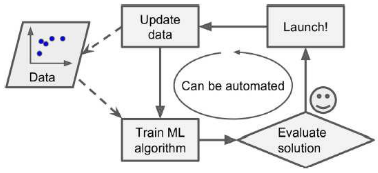

Automatically adapting to change

ML can help to human to learn

MACHINE LEARNING

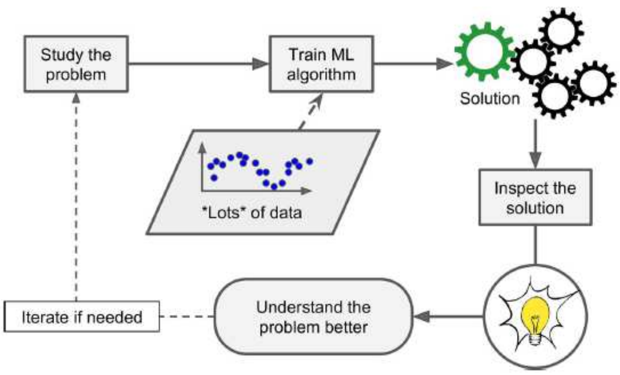

Terminologies

- Data / Dataset
- Training
- Testing
- Labeled data
- Training Samples
- Test Samples
- Evaluating Samples

Dataset: California Home Prices

count	20640

```
longitude
latitude
housing_median_age
total_rooms
total_bedrooms
population
households
median_income
median_house_value
ocean_proximity
```

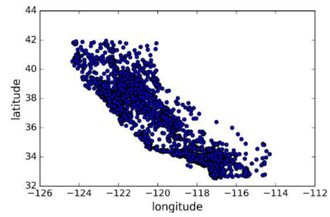

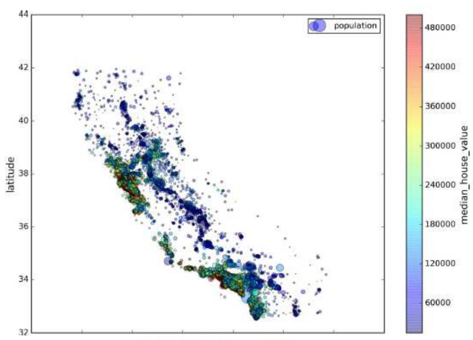

Figure 2-13. California housing prices

Dataset: IRIS

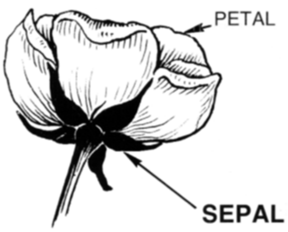

	Sepal.Length	Sepal.Width	Petal.Length	Petal.Width	Species
1	5	3.2	1.2	0.2	setosa
2	5.1	3.8	1.9	0.4	setosa
3	5.1	3.3	1.7	0.5	setosa
4	6.7	3.1	4.7	1.5	versicolor
5	5.1	3.7	1.5	0.4	setosa
6	5	3	1.6	0.2	setosa
7	5.3	3.7	1.5	0.2	setosa
8	5	3.4	1.6	0.4	setosa
9	4.9	2.4	3.3	1	versicolor
10	6.3	2.5	5	1.9	virginica

Dataset: MNIST Digits

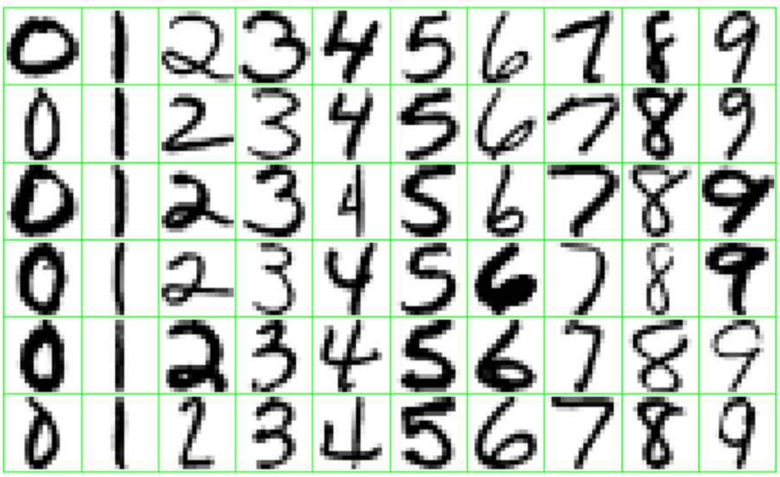

FIGURE 1.2. *Examples of handwritten digits from U.S. postal envelopes.*

Applications

MACHINE LEARNING

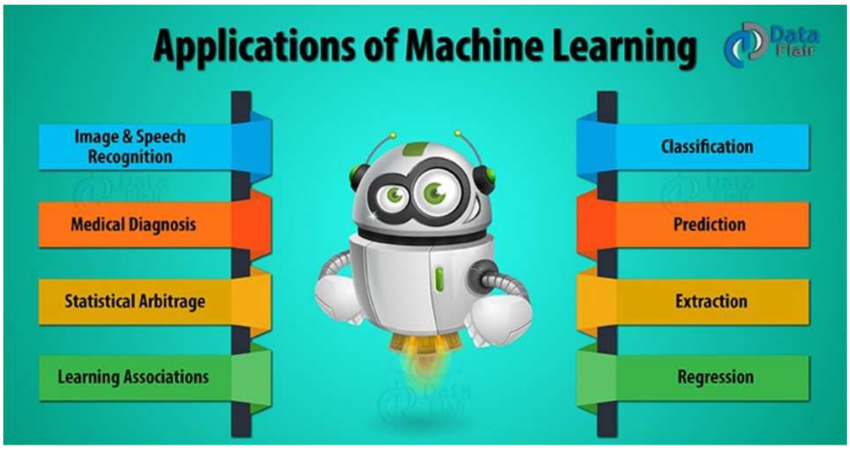

Visual Object Categorization

MACHINE LEARNING

We are given categories for these images: What are these?

- A classification problem: predict category y based on image x.
- Little chance to "hand-craft" a solution, without learning.
- Applications: robotics, HCI, web search (a real image Google...)

Applications

- Image Classification

  - Handwritten digit recognition (convert hand-written digits to characters 0..9)

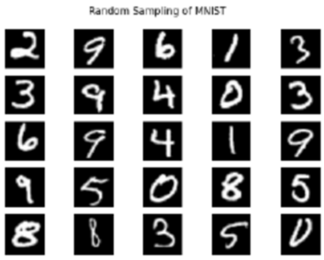

Applications

- Face Detection and Recognition

Applications

MACHINE LEARNING

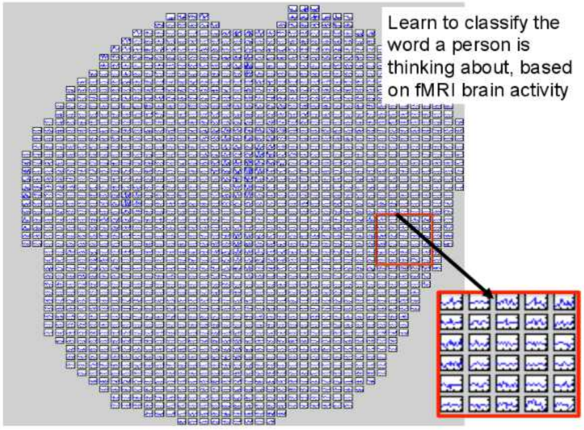

Applications

MACHINE LEARNING

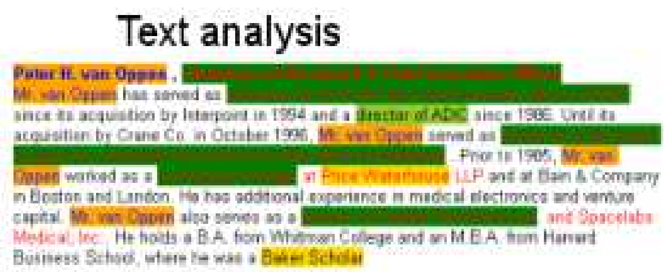

Applications

MACHINE LEARNING

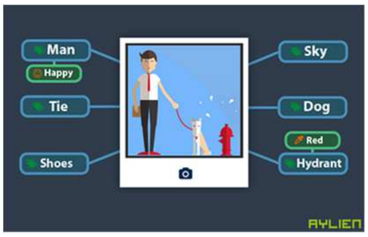

Photo tagging

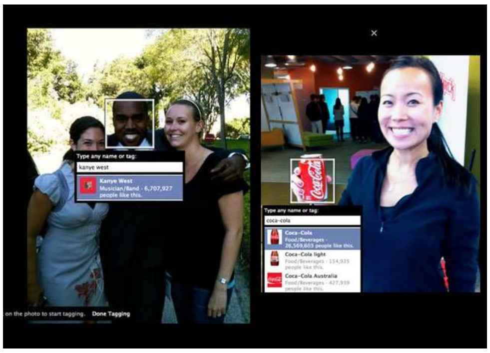

Applications

MACHINE LEARNING

- Weather prediction

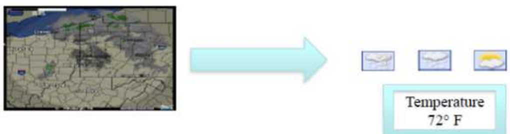

- Medicine:
  - diagnose a disease
    - input: from symptoms, lab measurements, test results, DNA tests, …
    - output: one of set of possible diseases, or "none of the above"
    - examples: audiology, thyroid cancer, diabetes, …
      - or: response to chemo drug X
      - or: will patient be re-admitted soon?
- Computational Economics:
  - predict if a stock will rise or fall
  - predict if a user will click on an ad or not
    - in order to decide which ad to show

Applications

MACHINE LEARNING

- Segment customers and find the best marketing strategy for each group
- Recommend products for each client based on what similar clients bought
- Detect which transactions are likely to be fraudulent
- Predict next year's revenue
- Learning from medical records which treatments are most effective

Applications

MACHINE LEARNING

- Self Customizing programs - houses learning to optimize energy costs based on particular usage patterns of their occupants
- Personal software assistants learning the evolving interests of their users in order to highlight relevant stories from online newspapers
- Autonomous driving
- Speech Recognition

Some successful ML applications

MACHINE LEARNING

- Learning to recognize spoken words (Lee, 1989; Waibel, 1989).
- Learning to drive an autonomous vehicle (Pomerleau, 1989).
- Learning to classify new astronomical structures (Fayyad et al., 1995).
- Learning to play world-class backgammon (Tesauro 1992, 1995).

Areas/Disciplines influence ML

MACHINE LEARNING

How best to use samples drawn from unknown probability distributions to help decide from which distribution some new sample is drawn?

Non-linear elements with weighted inputs (ANN) have been suggested as simple models of biological neurons.

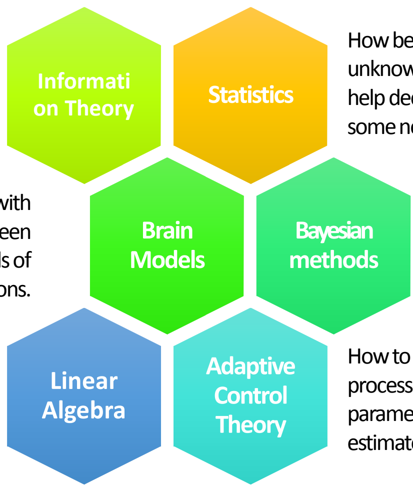

How to deal with controlling a process having unknown parameters that must be estimated during operation?

Areas/Disciplines influence ML

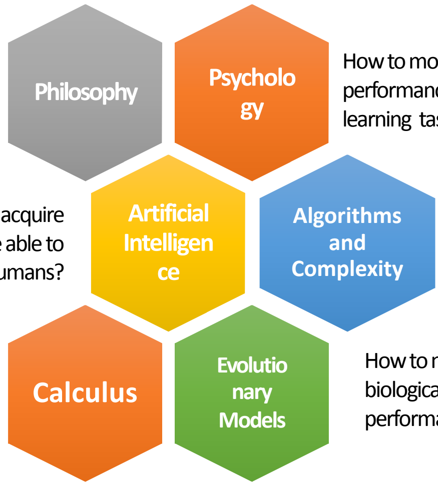

How to model human performance on various learning tasks?

How to write algorithms to acquire the knowledge humans are able to acquire, at least, as well as humans?

How to model certain aspects of biological evolution to improve the performance of computer programs?

Stages in ML process

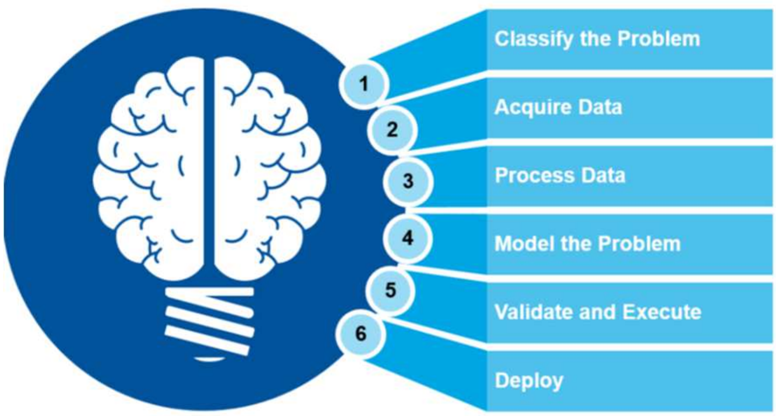

Source:https://www.gartner.com/binaries/content/assets/events/keywords/catalyst/catus8/preparing_and_architecting_for_machine_learning.pdf

Types of Machine Learning

1. Shallow Learning
   - Algorithms with Few Layers
   - Better for Less Complex and Smaller Data sets
   - Ex: Logistic Regression and Support vector Machines

2. Deep Learning
   - New technique that uses many layers of neural network (a model based on the structure of human brain)
   - Useful when the target function is very complex and data sets are very large.

Classification of ML algorithms

1. Supervised Learning (inductive) learning
   - Training data includes desired outputs
   - X and Y; Given an observation X what is the best label for Y
   - Example: Classification, Regression problems
2. Unsupervised Learning
   - Training data does not include desired outputs
   - X; Given a set of X cluster or summarize them
   - Example: Clustering
3. Semi Supervised Learning
   - Training data includes a few desired outputs
4. Reinforcement Learning
   - Determine what to do based on Rewards and punishments
   - Example: Robot movement, Game AI

Supervised Learning - Classification

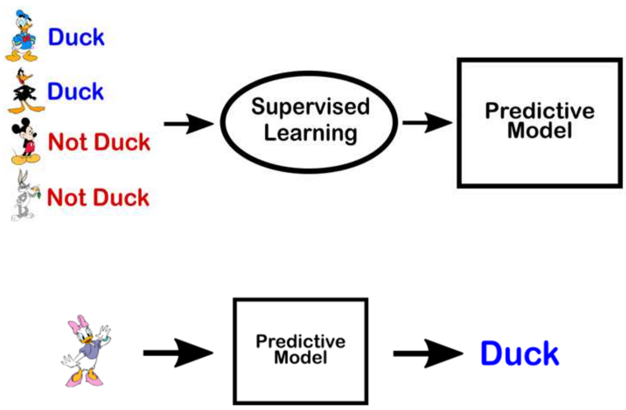

Source: http://www.java-machine-learning.com/blog/supervised-learning/

Unsupervised Learning-Clustering

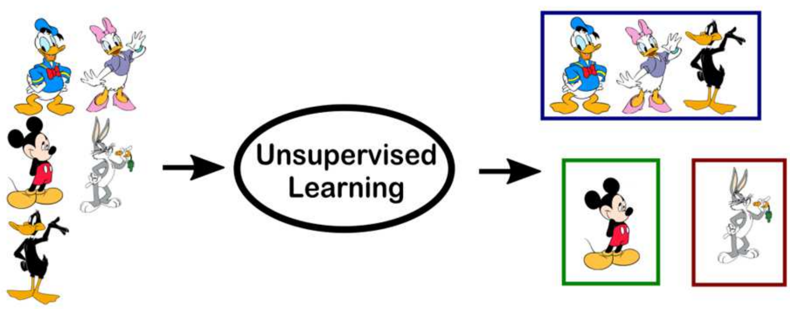

Unsupervised Learning-Clustering

MACHINE LEARNING

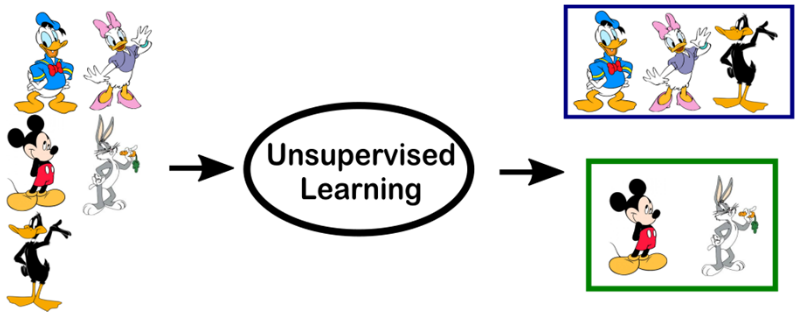

http://www.java-machine-learning.com/blog/unsupervised-learning/

ML – Types of Learning

MACHINE LEARNING

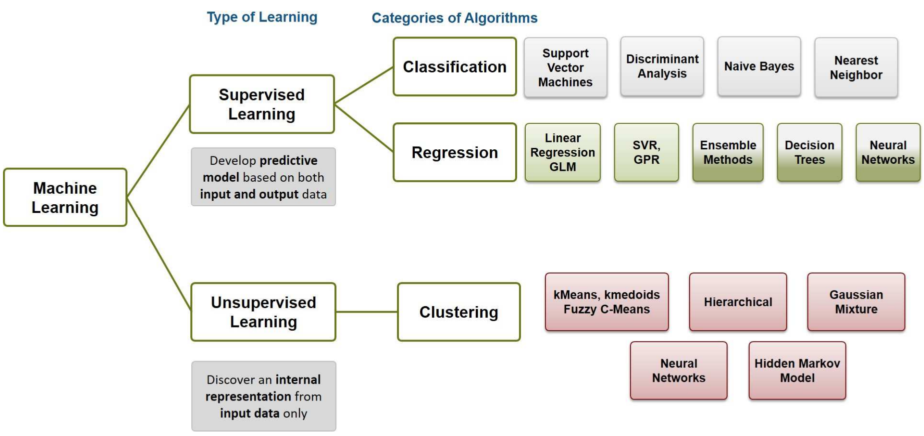

Source: http://www.embedded-computing.com/embedded-computing-design/analytics-driven-embedded-systems-part-2-developing-analytics-and-prescriptive-controls

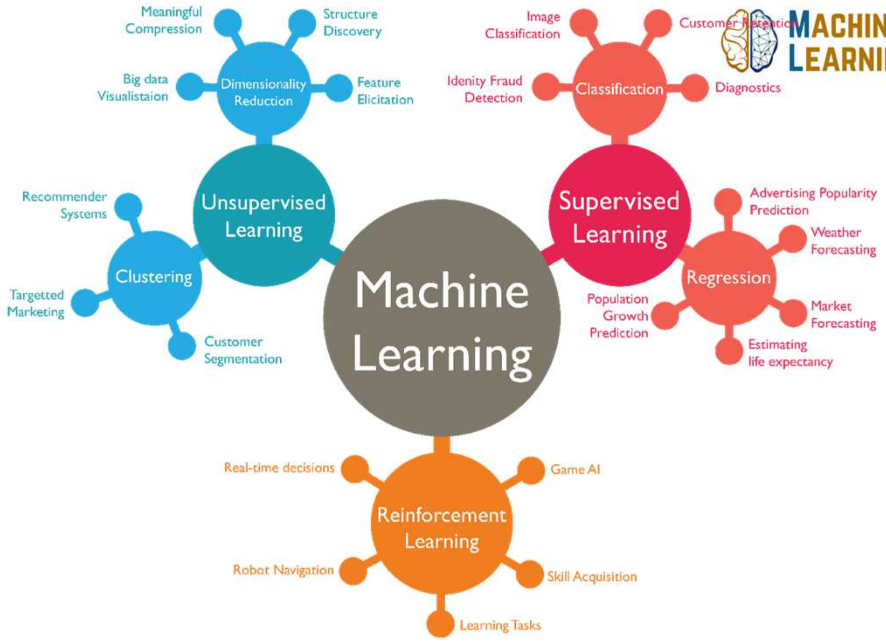

MACHINE LEARNING

Why is ML is Important?

MACHINE LEARNING

- Some tasks cannot be defined well, except by examples (e.g., recognizing people).
- Relationships and correlations can be hidden within large amounts of data. Machine Learning/Data Mining may be able to find these relationships.
- Human designers often produce machines that do not work as well as desired in the environments in which they are used.
- The amount of knowledge available about certain tasks might be too large for explicit encoding by humans (e.g., medical diagnostic).

Why is ML is Important?

- Environments change over time.
- New knowledge about tasks is constantly being discovered by humans. It may be difficult to continuously re-design systems "by hand".

Skills required for ML Engineer

1. Mathematical Skills
   - Probability, Statistics, Linear Algebra, Calculus
2. Programming Skills
   - Coding, Algorithms, DS, OOPs
   - Python, R, Matlab, Java
3. Data Engineering Skills
   - Data Preprocessing, Analysis, Visualization
4. Knowledge of ML algorithms
   - Shallow and Deep learning
   - Supervised, Semi-Supervised, Unsupervised, Reinforcement
5. Knowledge of ML Frameworks
   - SciKit Learn, Tensorflow, Caffe, Theano, Spark, …. many more

In summary, ML is great for:

- Problems for which existing solutions require a lot of hand-tuning or long lists of rules: one Machine Learning algorithm can often simplify code and perform better.
- Complex problems for which there is no good solution at all using a traditional approach: the best Machine Learning techniques can find a solution.
- Fluctuating environments: a Machine Learning system can adapt to new data.
- Getting insights about complex problems and large amounts of data.

Some issues…

- What algorithms can approximate functions well and when?
- How does number of training examples influence accuracy?
- How does complexity of hypothesis representation impact it?
- How does noisy data influence accuracy?
- What are the theoretical limits of learnability?
- How can prior knowledge of learner help?
- What clues can we get from biological learning systems?
- How can systems alter their own representations?

Where we are?

MACHINE LEARNING

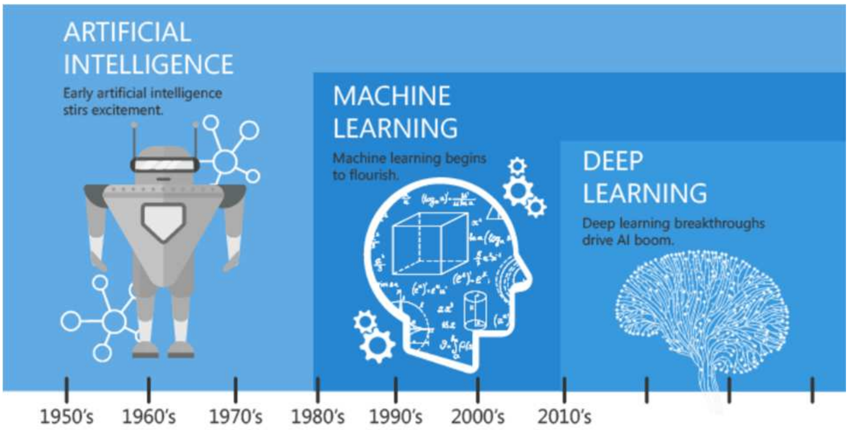

Since an early flush of optimism in the 1950's, smaller subsets of artificial intelligence - first machine learning, then deep learning, a subset of machine learning - have created ever larger disruptions.

Text and Reference Books

MACHINE LEARNING

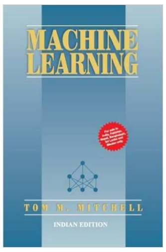

Tom M. Mitchell, Machine Learning, India Edition 2013, McGraw Hill.

T Hastie, R Tibshirani, J Friedman
The Elements of Statistical Learning,
Springer

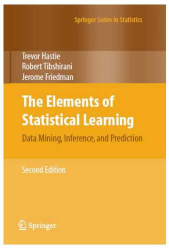

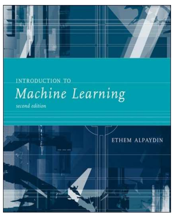

Ethem Alpaydın, Introduction to machine learning, 2nd edition, MIT press.

MACHINE LEARNING

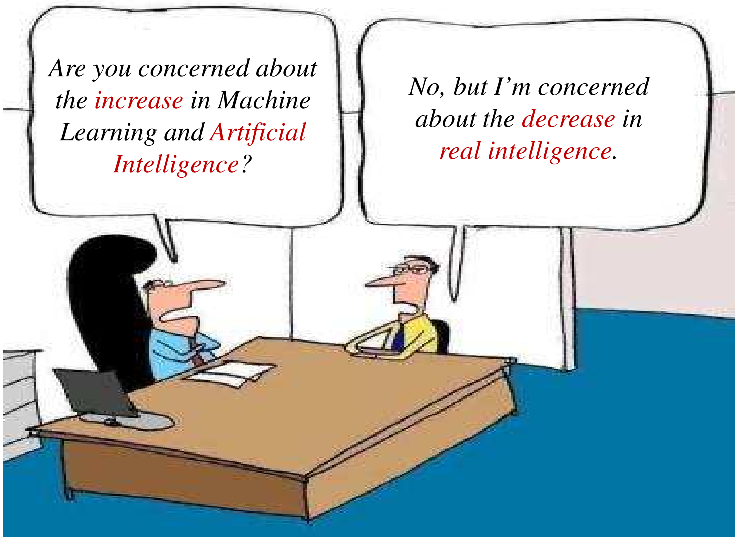

MACHINE LEARNING

Thank you