

# What is machine learning (ML)?

- Definition of “learning” from Merriam-Webster: “To gain knowledge or understanding of or skill in by study, instruction, or experience”
- ML = Learning in machines (computers)
- ML techniques are algorithms that enable the machines to improve its performance at some task through experience
- Tasks: recognition, diagnosis, prediction, planning, data mining, robot control, and so on.

# ***Definition of Machine Learning***

- A computer program is said to learn from experience ***E*** with respect to some class of tasks ***T*** and performance measure ***P***, if its performance at tasks in ***T***, as measured by ***P***, improves with experience ***E***.
- In general, to have a well-defined learning problem, we must identify these three features: the class of tasks, the measure of performance to be improved, and the source of experience.

# Examples of well-defined learning problem

- For example, a computer program that learns to play checkers might improve its performance as ***measured by its ability to win*** at the class of tasks involving ***playing checkers games***, through experience ***obtained by playing games against itself***.

## **A checkers learning problem:**

- **Task T:** playing checkers
- **Performance measure P:** percent of games won against opponents
- **Training experience E:** playing practice games against itself

# Examples of well-defined learning problem

## **A handwriting recognition learning problem:**

- **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

# Examples of well-defined learning problem

## **A robot driving learning problem:**

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

# Why Machine Learning?

- Some tasks can be specified only by training data/examples
- Human expertise may be scarce and/or very costly
- Amount of knowledge might be too large for explicit encoding by humans
- Modeling/Hidden parameter estimation: Often only data from measurements are available
- Computational power is ever increasing
- Growing data pool and storage capacity

# Traditional Programming

vs.

# Machine Learning

- **Traditional Programming:** Data and program is run on the computer to produce the output.
- **Machine Learning:** Data and output is run on the computer to create a program. This program can be used in traditional programming.

## Traditional Programming



## Machine Learning



# A Brief History about ML

- 1950's:
  - Samuel's checker program
  - Rosenblatt's perceptron
- 1960's:
  - Neural network
  - Pattern recognition
- 1970's:
  - Winston's ARCH
  - Buchanan and Mitchell's Meta-Dendral: mass spectrometry prediction rules
  - Quinlan's ID3: Chess end-game rules
  - Michalski's AQ11: Soybean disease diagnosis rules
  - MACROPS: macro operators in block world planning

# A Brief History about ML...

- 1980's: Gains momentum
  - Learning theory
  - Symbolic learning algorithms
  - Connectionist learning algorithm
  - Clustering
  - Explanation-based learning
  - Knowledge guided inductive learning
  - Genetic algorithm
- 1990's: Maturity
  - Data mining
  - Ensemble learning: bagging, boosting etc.
  - Kernel methods
  - Reinforcement learning
  - Theoretical analysis

# Applications of Machine Learning

- **Web search:** ranking page based on what you are most likely to click on.
- **Computational biology:** rational design drugs in the computer based on past experiments.
- **Finance:** decide who to send what credit card offers to. Evaluation of risk on credit offers. How to decide where to invest money.
- **E-commerce:** Predicting customer churn. Whether or not a transaction is fraudulent.
- **Space exploration:** space probes and radio astronomy.
- **Robotics:** how to handle uncertainty in new environments. Autonomous. Self-driving car.
- **Information extraction:** Ask questions over databases across the web.
- **Social networks:** Data on relationships and preferences. Machine learning to extract value from data.
- **Debugging:** Use in computer science problems like debugging. Labor intensive process. Could suggest where the bug could be.

# Steps Involved in Machine Learning

A machine learning project involves the following steps –

- Defining a Problem
- Preparing Data
- Evaluating Algorithms
- Improving Results
- Presenting Results

# Concepts of Learning

There are four categories of machine learning algorithms as shown below –

- Supervised learning algorithm
- Unsupervised learning algorithm
- Semi-supervised learning algorithm
- Reinforcement learning algorithm

However, the most commonly used ones are **supervised** and **unsupervised learning**.

# Supervised Learning

- Supervised learning is commonly used in real world applications, such as face and speech recognition, products or movie recommendations, and sales forecasting. Supervised learning can be further classified into two types - **Regression** and **Classification**.
- **Regression** trains on and predicts a continuous-valued response, for example predicting real estate prices.
- **Classification** attempts to find the appropriate class label, such as analyzing positive/negative sentiment, male and female persons, benign and malignant tumors, secure and unsecure loans etc.

# Supervised Learning

- In supervised learning, learning data comes with description, labels, targets or desired outputs and the objective is to find a general rule that maps inputs to outputs. This kind of learning data is called **labeled data**. The learned rule is then used to label new data with unknown outputs.

There are many **supervised learning algorithms** such as Logistic Regression, Neural networks, Support Vector Machines (SVMs), and Naive Bayes classifiers.

- Common **examples** of supervised learning include classifying e-mails into spam and not-spam categories, labeling webpages based on their content, and voice recognition.

# Unsupervised Learning

- Unsupervised learning makes sense of unlabeled data without having any predefined dataset for its training. Unsupervised learning is an extremely powerful tool for analyzing available data and look for patterns and trends. It is most commonly used for clustering similar input into logical groups.

# Unsupervised Learning

- No training data in the form of (input, output) pair is available
- Applications:
  - Dimensionality reduction
  - Data compression
  - Outlier detection
  - Segmentation/clustering
  - Probability density estimation
  - ...



# Example: Unsupervised Learning (contd..)

- Applications on DNA microarray
  - Clustering: Group genes or samples into similar expression profiles
  - Bi-clustering: Subset of genes exhibiting similar expression pattern along a subset of samples
  - Dimension reduction
  - ...

# Semi-supervised Learning

- If some learning samples are labeled, but some other are not labeled, then it is semi-supervised learning. It makes use of a large amount of **unlabeled data for training** and a small amount of **labeled data for testing**.
- Semi-supervised learning is applied in cases where it is expensive to acquire a fully labeled dataset while more practical to label a small subset. For example, it often requires skilled experts to label certain remote sensing images, and lots of field experiments to locate oil at a particular location, while acquiring unlabeled data is relatively easy.

# Reinforcement Learning

- Here learning data gives feedback so that the system adjusts to dynamic conditions in order to achieve a certain objective. The system evaluates its performance based on the feedback responses and reacts accordingly.
- Reinforcement learning is the problem faced by an agent that must learn behavior through trial-and-error interactions with a dynamic environment.

# Reinforcement Learning

- There is no teacher telling the agent wrong or right
- There is critic that gives a reward / penalty for the agent's action
- Applications:
  - Robotics
  - Combinatorial search problems, such as games
  - Industrial manufacturing
  - Many others!
- The best known instances include self-driving cars and chess master algorithm AlphaGo.

# Example: Reinforcement Learning

	<b>Tic Tac Toe</b>	<b>TD-Gammon</b>
<b>Goal</b>	Learn to play optimal game	Learn to play game at master level
<b>States</b>	All possible board states - 9	All possible board states - $10^{20}$
<b>Action</b>	A new X in an empty field	21 dice combinations & avg. 20 legal moves
<b>Reinforcement Signal</b>	+10 winning -1 for every move that did not win	'1 ' for a reward '0' for a penalty

# Purpose of Machine Learning

- Machine learning can be seen as a branch of AI or Artificial Intelligence, since, the ability to change experience into expertise or to detect patterns in complex data is a mark of human or animal intelligence.
- As a field of science, machine learning shares common concepts with other disciplines such as statistics, information theory, game theory, and optimization.

# Purpose of Machine Learning

- As a subfield of information technology, its objective is to program machines so that they will learn.
- However, it is to be seen that, the purpose of machine learning is not building an automated duplication of intelligent behavior, but using the power of computers to complement and supplement human intelligence. For example, machine learning programs can scan and process huge databases detecting patterns that are beyond the scope of human perception.