

Advances and Opportunities in Machine Learning for Data Processing and Analytics

Uldis Doniņš

uldis.donins@rsu.lv


RSU ITD Information systems unit

04.11.2020.

Artificial Intelligence

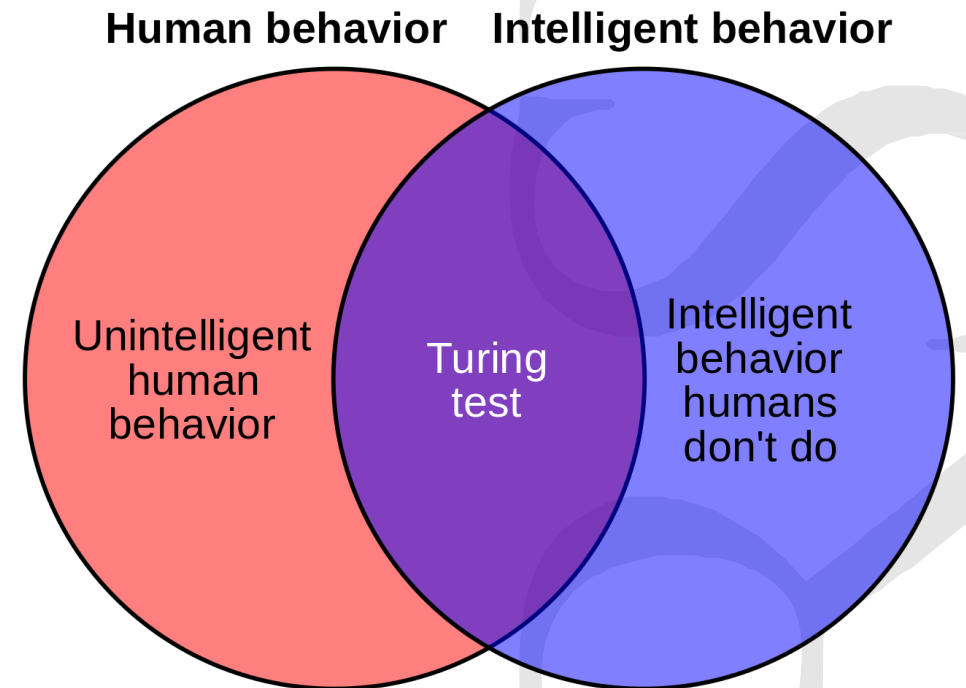
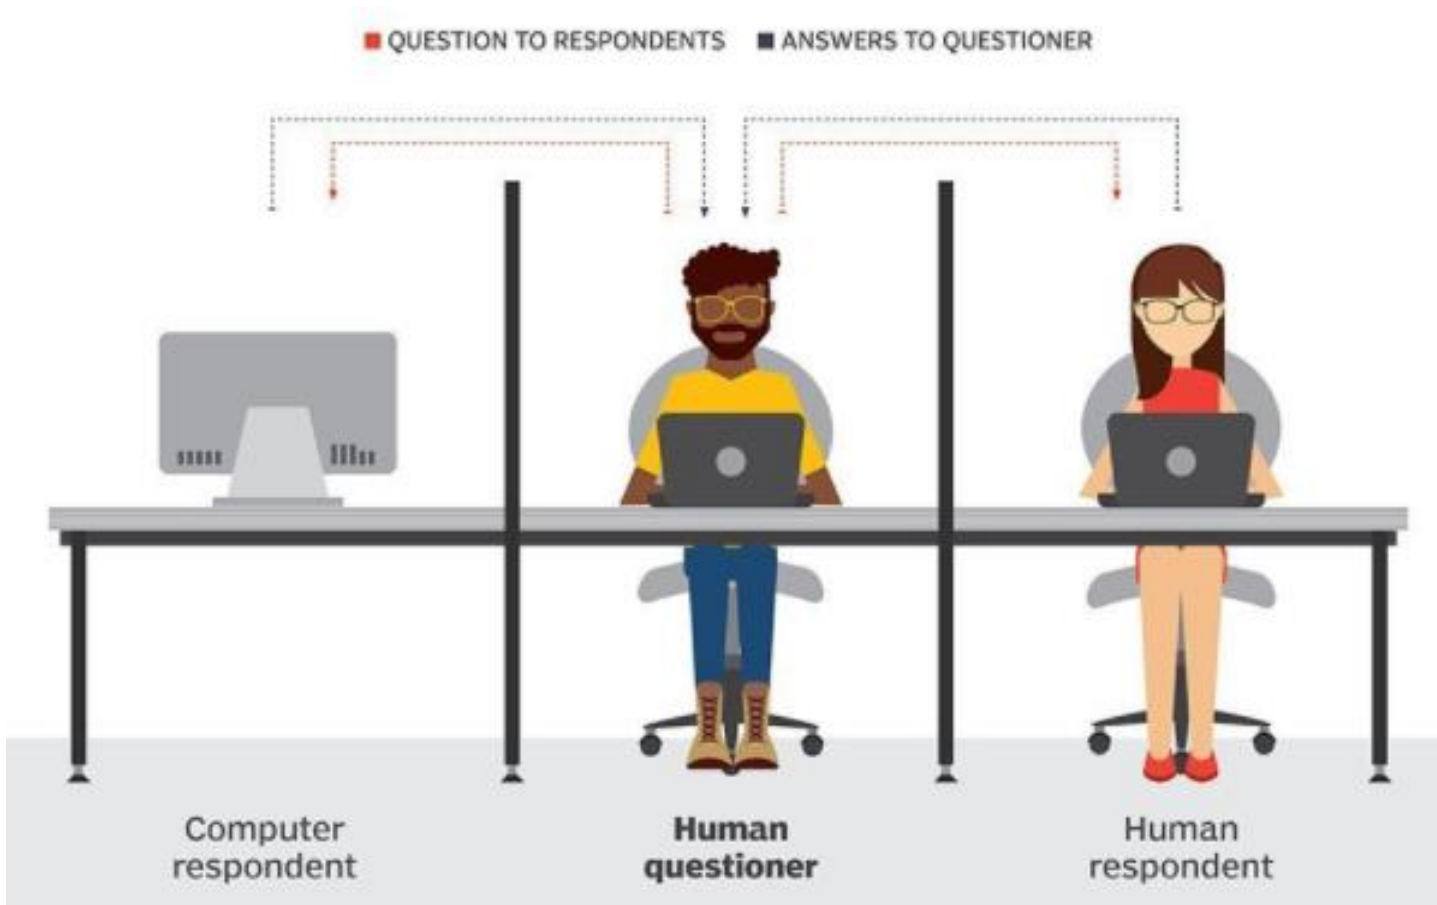
The term “Artificial Intelligence” was introduced by John McCarthy at 1955, when he described AI as “*that of making a machine behave in ways that would be called intelligent if a human were so behaving.*”

Nowadays Artificial Intelligence is defined as the study of “intelligent agents” – any device that perceives its environment and takes actions that maximize its chance of successfully achieving its goals.




The Turing Test

If the questioner cannot reliably distinguish computer respondent from the human, the computer is said to have passed the test




What Would Machine Need to Pass the Turing test?

- **Natural language processing:** to communicate with examiner
- **Knowledge representation:** to store and retrieve information provided before or during interrogation
- **Automated reasoning:** to use the stored information to answer questions and to draw new conclusions
- **Machine learning:** to adapt to new circumstances and to detect and extrapolate patterns
- **Vision (for Total Turing test):** to recognize the examiner's actions and various objects presented by the examiner
- **Motor control (total test):** to act upon objects as requested
- **Other senses (total test):** such as audition, smell, touch, etc



What is an Intelligent Agent?

- Anything that can be viewed as perceiving its environment through sensors and acting upon that environment through its actuators (effectors) to maximize progress towards its goals



- PAGE (Percepts, Actions, Goals, Environment)
- Task-specific & specialized: well-defined goals and environment



DRS

Good Old- Fashioned AI (GOFAI)

Logic (Programming)

- Expert systems



Search

- **Breadth-first:** the node list is a queue: first-in, first-out (FIFO)
- **Depth-first:** the node list is a stack: last-in, first-out (LIFO)
- Other search algorithms can be obtained by replacing the list with yet other structures, e.g.
- **Best-first search:** Node list is a priority queue: highest-value gets out first


Games

- Minimax
- Alpha-Beta

“if I make this move, then my opponent can only make two moves, and each of those would let me win.”




Depth-First Search





Breadth-First Search

GPS Navigation



DRK

Solving Tic Tac Toe with Minimax

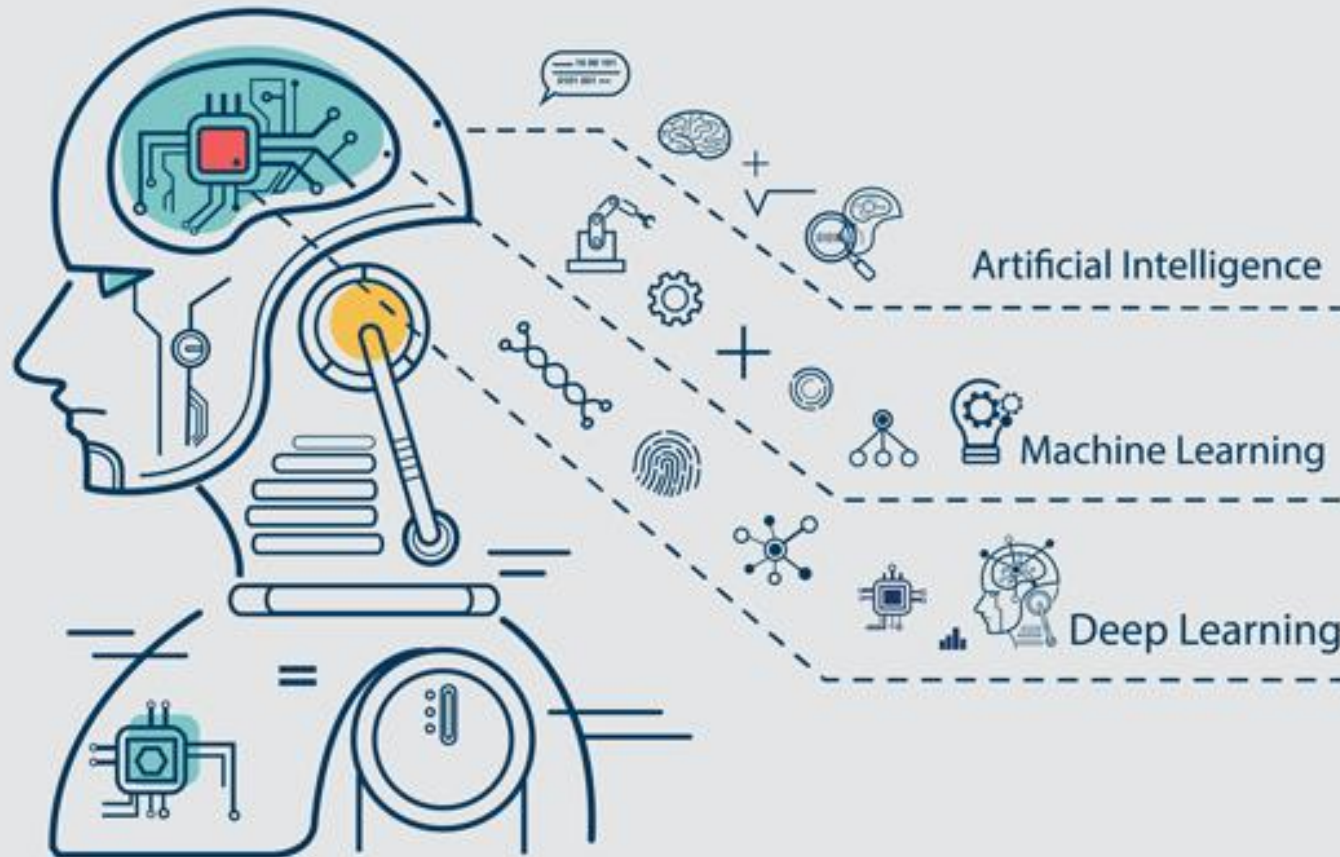


Symbolic Victory for Artificial Intelligence in 1997



Modern AI

- Machine Learning (ML)
- Natural Language Processing (NLP)
- Robotics



Machine Learning

- Machine learning can be broadly defined as computational methods using experience to improve performance or to make accurate predictions
 - **Machine** – computer or computer program
 - **Learning** – improving performance on a given task, based on experience / examples
- Machine learning is inherently related to data analysis and statistics, it consists of designing efficient and accurate prediction algorithms
- Experience refers to the past information available to the learner, which typically takes the form of electronic data collected and made available for analysis
- Training set quality and size are crucial to the success of the predictions made by the learner
- In other words, instead of programming explicit rules for how to solve a given problem, the computer is instructed how to learn from examples



Machine Learning Workflow

Formulate Question

Gather Data

Clean Data

Explore & Visualize

Train Algorithm

Evaluate

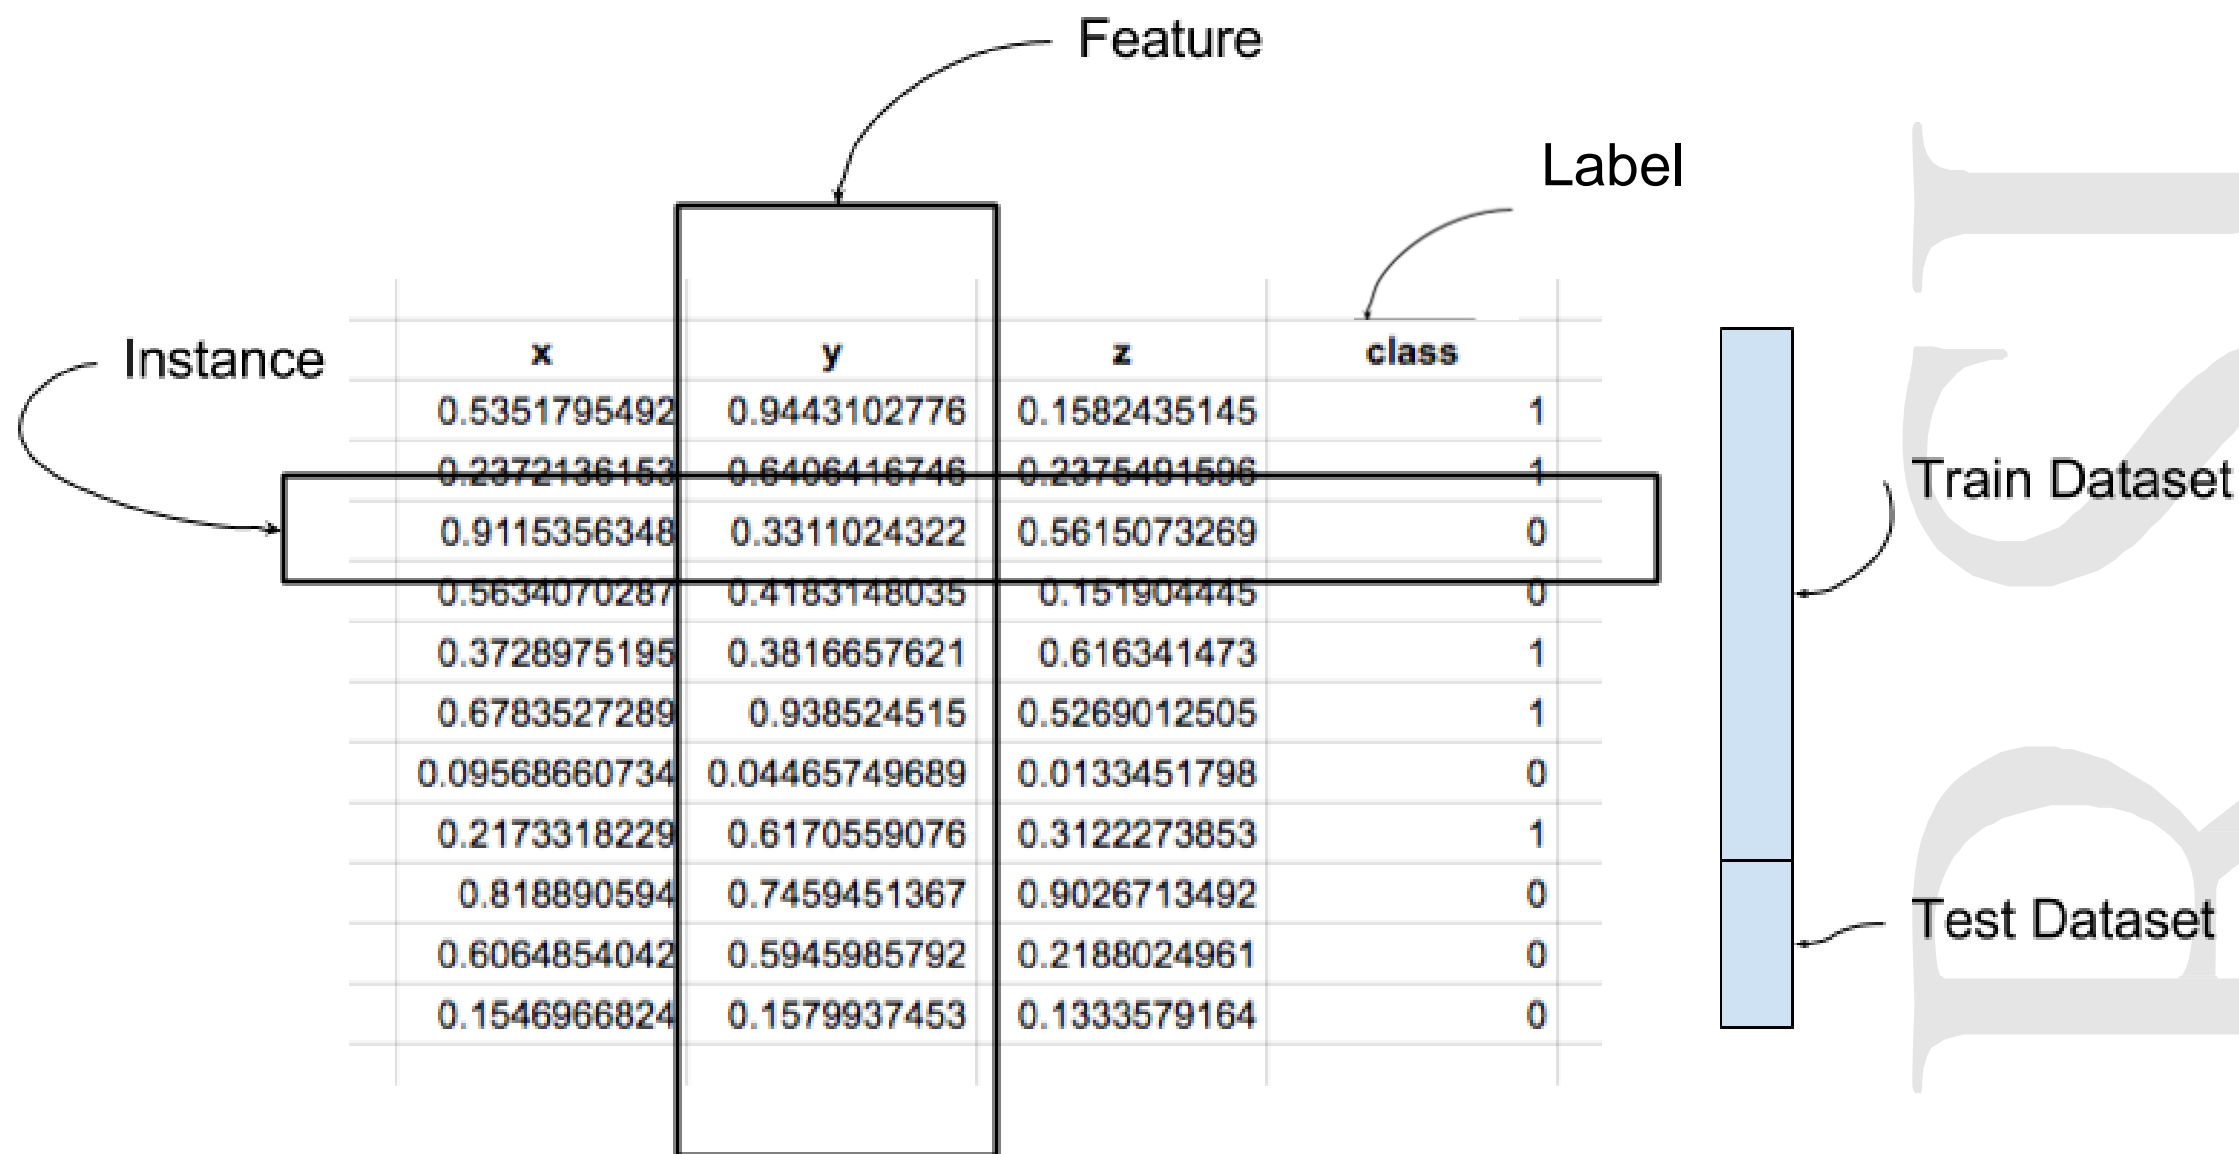


Standard Learning Tasks

- **Regression**
 - The learning task becomes regression if the dependent variable is a continuous value
 - Regression task tries to learn a function which map independent variable to a continuous value output
- **Classification**
 - The learning task is called classification if the dependent variable is a categorical
- **Clustering**
 - Partitioning a set of items into homogeneous subsets
- **Dimensionality reduction**
 - Transforming an initial representation of items into a lower-dimensional representation
- **Ranking**
 - Learning to order items according to some criterion



Data Sets



Supervised Learning

- Supervised learning is the machine learning task of learning a function that maps an input to an output based on example input-output pairs
- The training data has values of depended variable (label) corresponding to the independent variables (features)
- The goal of supervised learning algorithm is to learn a function that maps independent variables to the dependent variable (features to label)


$$Y = F(X)$$

- The label or the dependent variable can be categorical or continuous


Supervised Learning

Depending upon the data type of this variable the task changes:

- **Regression**




- **Classification**




Unsupervised Learning

- Unsupervised methods are used when labels associated with corresponding features are not available
- Looks for previously undetected patterns in a data set
- Two of the main learning problems of unsupervised learning:
 - Clustering
 - Dimensionality reduction





DRS



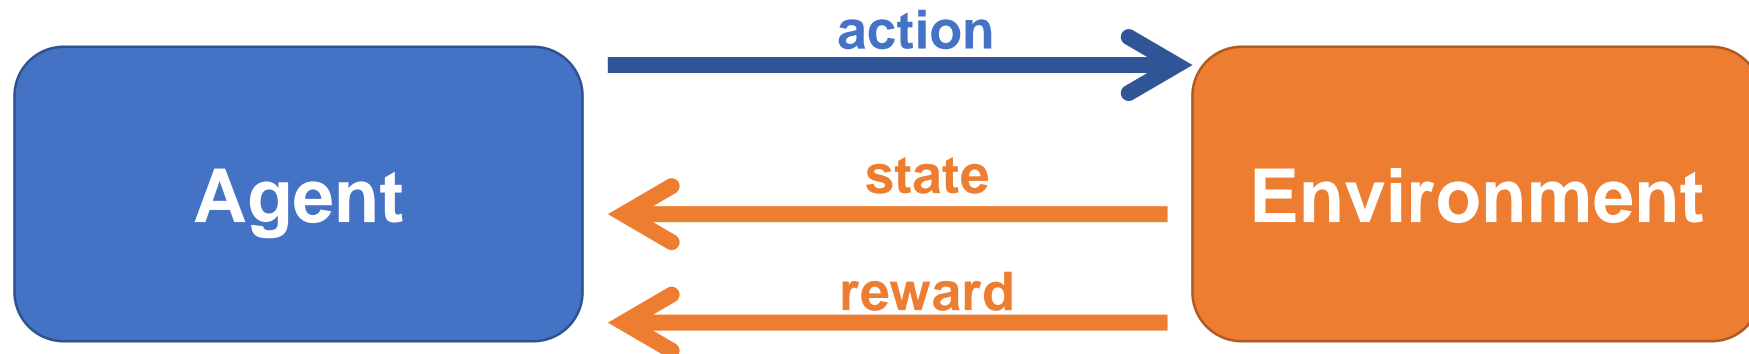
Semi-supervised Learning

- The learner receives a training sample consisting of both labeled and unlabeled data, and makes predictions for all unseen points
- Semi-supervised learning is common in settings where unlabeled data is easily accessible but labels are expensive to obtain
- General assumptions:
 - Points that are close to each other are more likely to share a label
 - The data tend to form discrete clusters, and points in the same cluster are more likely to share a label
 - The data lie approximately on a manifold of much lower dimension than the input space



Reinforcement Learning


- To collect information, the learner actively interacts with the environment and in some cases affects the environment, and receives an immediate reward for each action
- The object of the learner is to maximize his reward over a course of actions and iterations with the environment
- The focus is on finding a balance between exploration (of uncharted territory) and exploitation (of current knowledge)
- Reinforcement learning is very valuable in the field of robotics
 - The robot's task consists of finding out, through trial and error (or success), which actions are good in a certain situation and which are not




Deep Learning

- Deep learning is based on artificial neural networks
- The idea of artificial neural networks is taken from natural neural networks


Biological neuron



Artificial neuron – Perceptron



Deep Learning – Artificial Neural Networks





Artificial Intelligence

Machine Learning

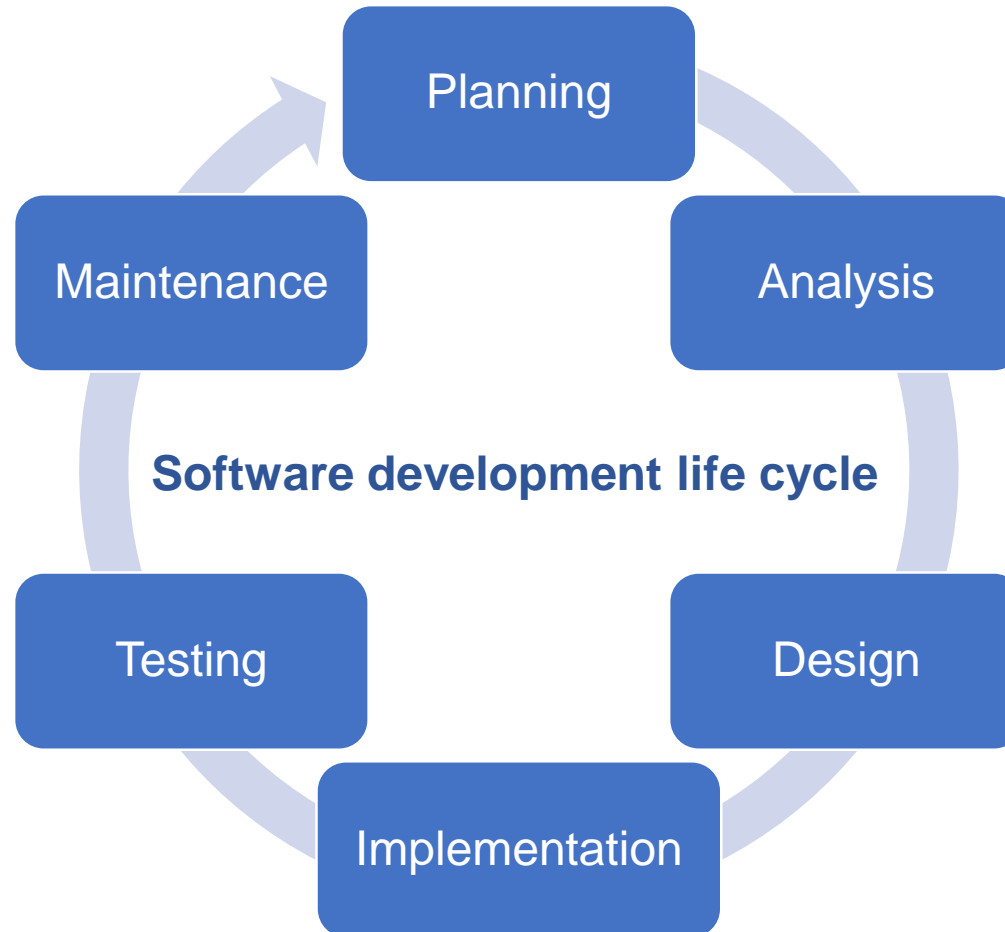
Deep Learning

DEEP

Implementing Machine Learning Solutions

Coding from scratch

- Extensive coding and debugging of code that implements machine learning algorithm, improving and optimizing the code
- Expensive, long-term project



Implementing Machine Learning Solutions

Existing machine learning libraries and tools

- Low-code solutions
- Exploring existing libraries and their community support
- Training the machine learning model, improving model performance



TensorFlow

Implementing Machine Learning Solutions

Cloud services

- Low-code solutions & no-code solutions
- Training the machine learning model, improving model performance
- Most of cloud providers supports libraries and tools discussed earlier
- Costs can include subscription fee, data storage and transfer, computation power used

Machine Learning cloud service providers

- Microsoft Azure
- Amazon Web Services (AWS)
- Google Cloud



Azure Machine Learning



Google Cloud


Machine Learning Use Cases in Daily Life

- **Smartphones**
 - Voice Assistants
 - Camera
 - App Store and Play Store Recommendations
 - Face Unlock
- **Transportation**
 - Traffic lights
 - Maps
- **Online services**
 - Email filtering
 - Search
 - Translation
 - LinkedIn and Facebook recommendations and ads

DATA

Machine Learning Use Cases

- Healthcare and medicine
 - Medical Imaging & Biomedical Diagnostics
 - Decision Support & Hospital Monitoring
 - Digital Medicine & Wearable technology
 - Robotic Technology & Virtual Assistant
 - Precision Medicine & Drug Discovery
- Education
 - Personalized Learning
 - Smart and Adaptable Content
 - Automatic Grading
- Smart homes and robots
- Environment & energy
- Security
- Transport and logistics
- Employment
- Economy, business, and finance
 - process automation
 - profit optimization



Research Fronts on Machine Learning

- Engineering
 - Deep learning-based speech enhancement approaches
 - Deep asymmetric video-based person re-identification
 - Electric energy consumption forecasting models
- Neuroscience & Behaviour
 - Hierarchical deep neural networks
- Biology & Biochemistry
 - Protein structure prediction methods
- Physics
 - Machine learning quantum phases
- Computer Science
 - Novel robust structured subspace learning
- Chemistry
 - Predictive qsar models
- Clinical Medicine
 - Deep learning convolutional neural networks
- Geosciences
 - Hyperspectral image classifications



Common Misconceptions About Artificial Intelligence

Understanding


- Interpreting, not analyzing
- Going beyond pure numbers
- Considering consequences

Discovering

- Devising new data from old
- Seeing beyond the patterns
- Implementing new senses

Empathizing

- Walking in someone's shoes
- Developing true relationships
- Changing perspective
- Making leaps of faith



References

- Jerry Kaplan. Artificial Intelligence: What Everyone Needs to Know. Oxford University Press, 2016.
- Mehryar Mohri, Afshin Rostamizadeh, and Ameet Talwalkar. Foundations of Machine Learning. MIT Press, Second Edition, 2018.
- Python: <https://www.python.org/>
- TensorFlow: <https://www.tensorflow.org/>
- Anaconda: <https://www.anaconda.com/products/individual>
- NumPy: <https://numpy.org/doc/stable/>
- Pandas: <https://pandas.pydata.org/docs/>
- Scikit-Learn: https://scikit-learn.org/stable/user_guide.html
- Matplotlib: <https://matplotlib.org/contents.html>
- Graph theory https://en.wikipedia.org/wiki/Graph_theory

Thank you!

Any questions?

DD
SS
RR