Emerging Issues in Computer Science

Week 2: Artificial Intelligence and Machine Learning

Lecturer: Ikwap Flavia Agatha

Learning Out come

- At the End of Lecture 2 you will be able to:

- Define Artificial Intelligence
- Understand Deep Learning
- Understand Neural Network and the different types of Neural Networks
- Understand Natural Language Processing NLP
- Understand Risks involved in generative AI
- Understand Machine Learning and the different machine learning algorithms

Artificial Intelligence (AI)

- Artificial Intelligence (AI) refers to the ability of machines or computer programs to carry out tasks that typically require human intelligence, such as learning, problem-solving, and decision-making.
- There’s a huge amount of data generated today, which is one of the reasons why AI is so prevalent today.
- AI helps make sense of vast amounts of data in a shorter amount of time than the amount of time a person would take.

Types of AI

- **Narrow or Weak AI**
  Narrow AI is designed to perform specific tasks or solve particular problems, such as image recognition or language translation. It relies on machine learning algorithms trained on large datasets to detect patterns and make predictions based on these patterns.
- **General AI**
  General AI aims to replicate human intelligence across a wide range of areas, including reasoning, perception, and creativity. It works by processing large amounts of data to identify patterns and relationships, using this information to make decisions or take actions. This is facilitated by algorithms that are designed to learn from the data they receive.

Deep Learning

- Deep learning is a subset of machine learning that uses artificial neural networks to allow machines to learn from data and perform complex tasks. These networks are inspired by the human brain’s structure, consisting of layers of interconnected nodes that process information. Each node performs a simple calculation based on its inputs and generates a result, which is passed to the next layer.
- Training a deep learning model involves feeding it with a large dataset and adjusting the weights of the connections between nodes to minimize the difference between the predicted and actual output. This process, known as back propagation, iteratively adjusts the connection weights to improve the model's accuracy in predicting outputs based on given inputs.

Machine Learning and Deep Learning

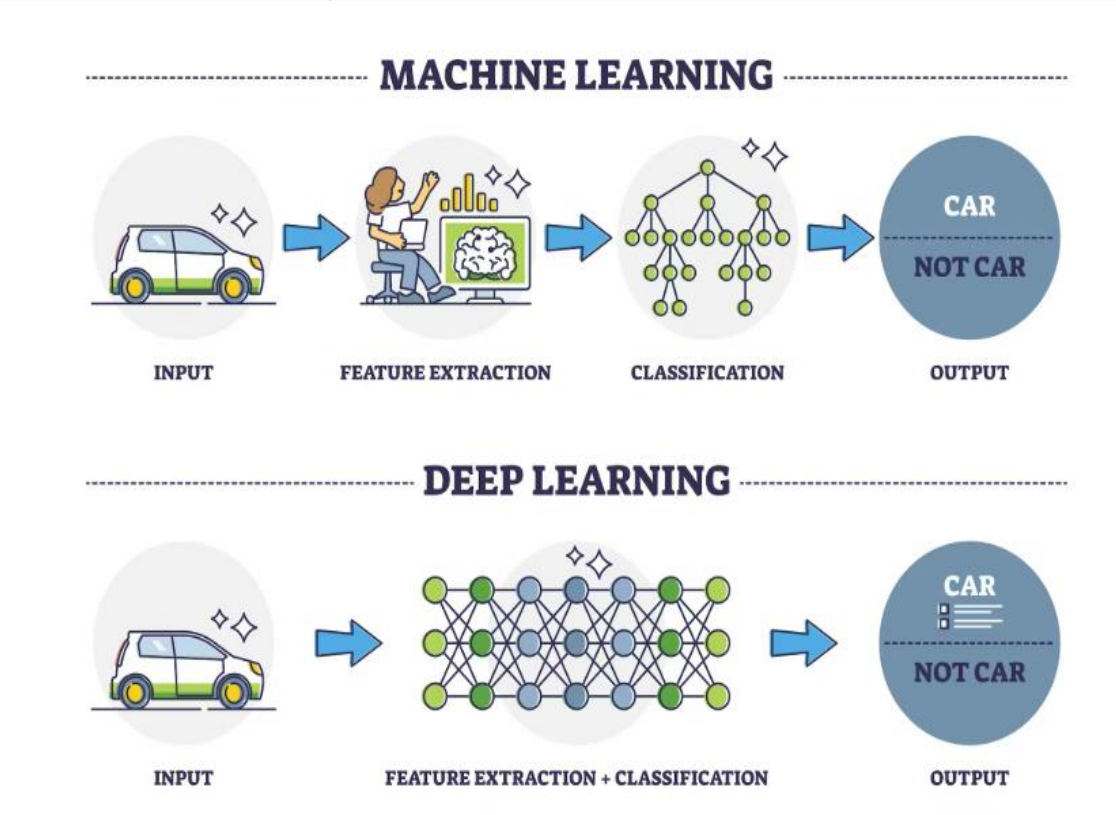

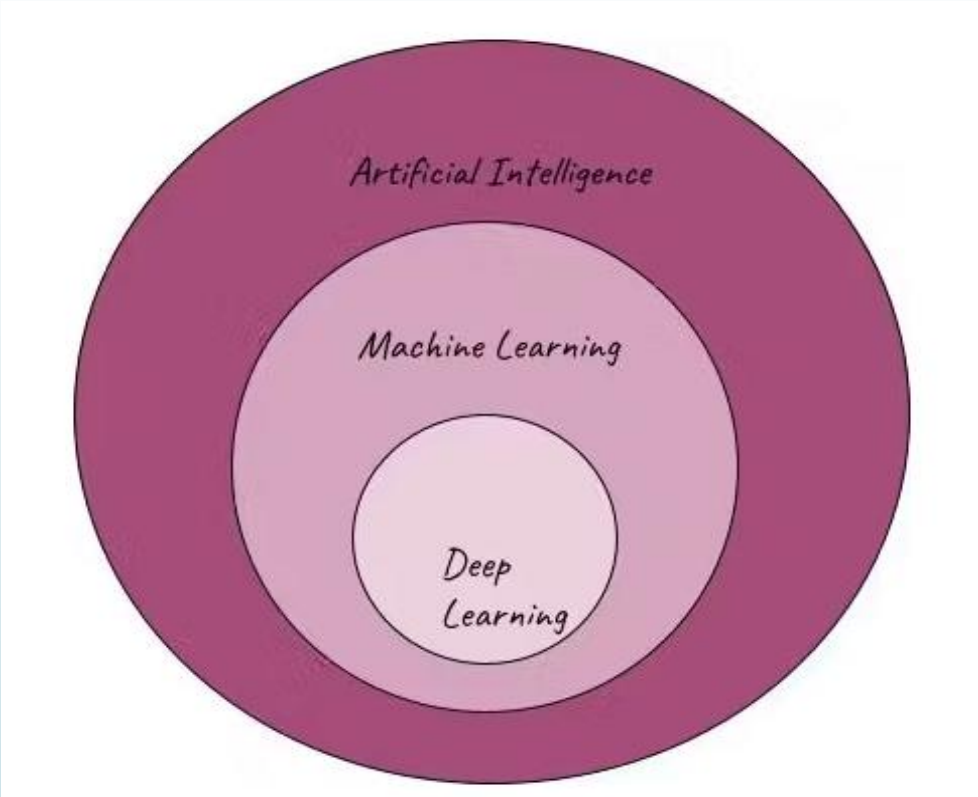

- https://www.turing.com/kb/ultimate-battle-between-deep-learning-and-machine-learning

Explaining Classification as used in the previous diagram

- What is classification? Also referred to as supervised learning. It is a data mining function that assigns items in a collection to target categories or classes. The goal of classification is to accurately predict the target class for each case in the data. e.g. a classification model could be used to identify loan applicants as: low, medium, or high credit risks.
- **Data** classification is the process of organizing **data** into categories for its most effective and efficient use. A well-planned **data** classification system makes essential **data** easy to find and retrieve.
- What is prediction? To predict is to do with forecasting or foretelling what “classified” data mean.

Advantages and limitations of deep learning

- **Advantages**
- Deep learning Model has the ability to perform complex tasks, such as image recognition, speech recognition, and natural language processing.
- **Limitation**
- Deep learning models typically require large amounts of labeled data to accurately learn the underlying patterns and relationships in the data. Labeled data may be scarce or difficult to obtain.
- Complexity of deep learning models. Deep learning models can be difficult to interpret, and it can be challenging to understand why they make the predictions or decisions that they do.

Neural Networks

- **Neural Networks** are a form of artificial intelligence modeled after the human brain's structure and function. These computational models learn from data to make predictions or decisions. Neural networks are widely used in areas like image and speech recognition, natural language processing, and more.

What is a Neural Network?

- At its core, a neural network consists of interconnected nodes, or neurons, that work together to process and analyze data. Each neuron receives inputs from others, processes them, and passes the output to other neurons. The inputs to each neuron are assigned weights, meaning some inputs are more influential than others in determining the neuron’s output.
- Neurons are organized in layers: the input layer where data enters, one or more hidden layers where most of the data processing occurs, and the output layer, which generates the final result of the network's analysis.

Neural Networks

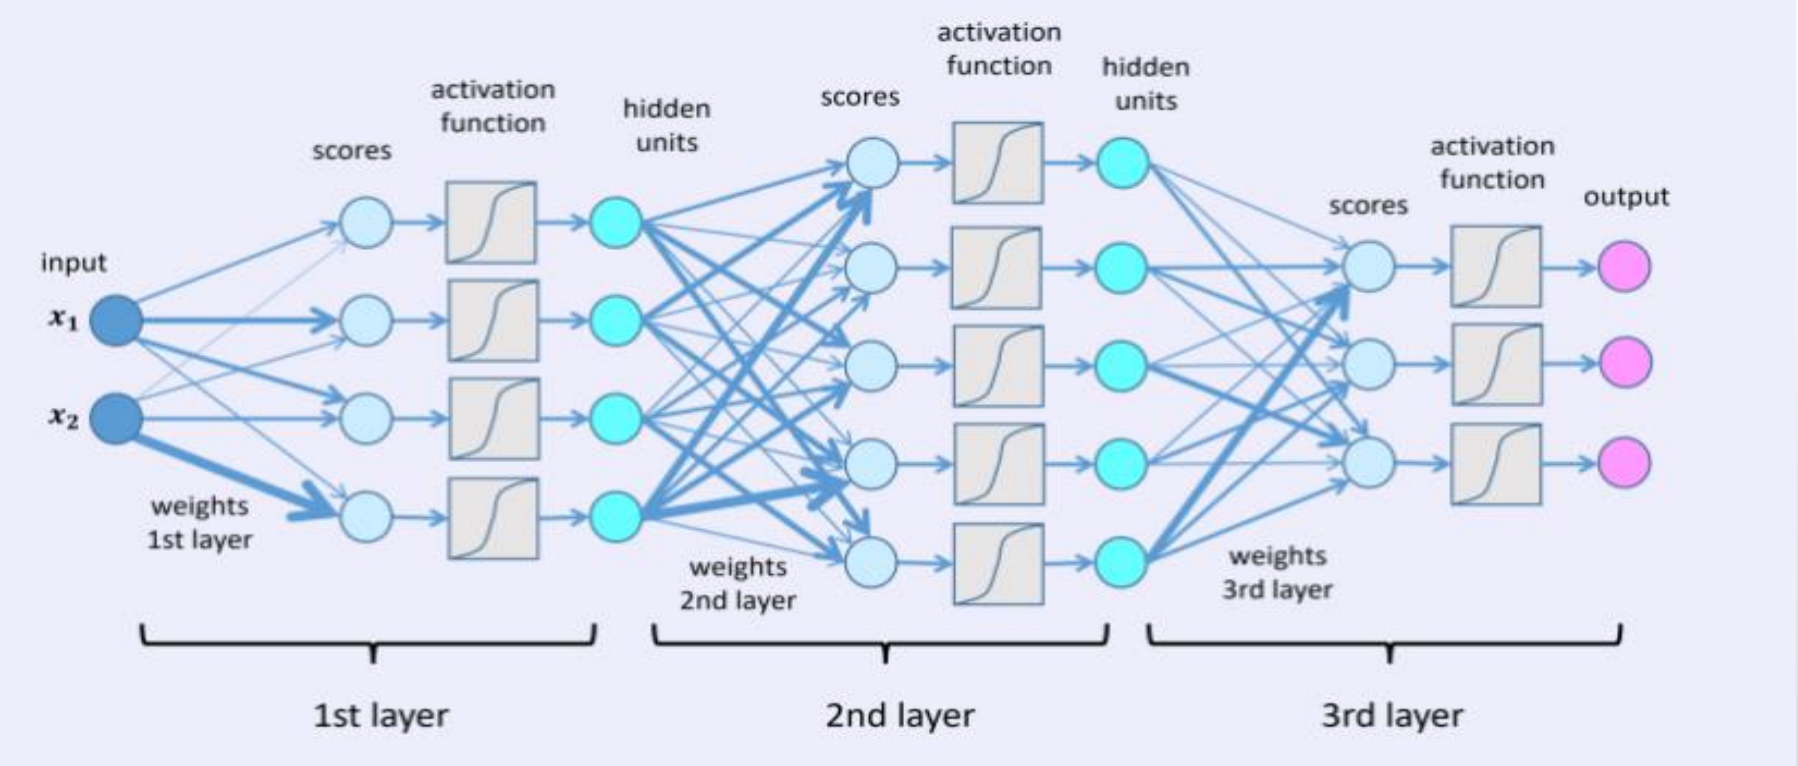

- https://lamarr-institute.org/blog/deep-neural-networks/

Training a Neural Network

The power of a neural network comes from its ability to learn from data. To train a neural network, a set of training data and the desired outputs are provided. The network then adjusts the weights of the connections between neurons to reduce the difference between its predicted and desired outputs through a process called back propagation.

- The number of layers and neurons in each layer affects the network's accuracy and speed. If there are too few neurons or layers, the network may struggle to represent complex data accurately. On the other hand, too many neurons or layers can lead to over fitting, where the network becomes too tailored to the training data and fails to generalize to new data.

Applications of Neural Networks

- neural networks are widely used in various fields, including image recognition, speech recognition, and natural language processing. In image recognition, for example, a neural network can be trained to detect specific features, like edges or shapes, to identify objects within an image.
- For speech recognition, neural networks can learn to recognize phonemes (the smallest units of sound) that form words, enabling them to transcribe spoken language into text. In natural language processing, these networks can be trained to understand the meaning of words and sentences, allowing them to perform tasks like language translation or sentiment analysis.

Types of Neural Networks

- **Feed forward Neural Networks**
- They are primarily employed for classification or regression tasks, where the result is either a single value or a class label. It is a directed acyclic graph, meaning there are no feedback loops or cycles in the network. The structure consists of an input layer, an output layer, and at least one hidden layer, though multiple hidden layers are possible. Each node in these layers is a neuron, which acts as the fundamental processing unit of a neural network.

Feed forward Neural Networks

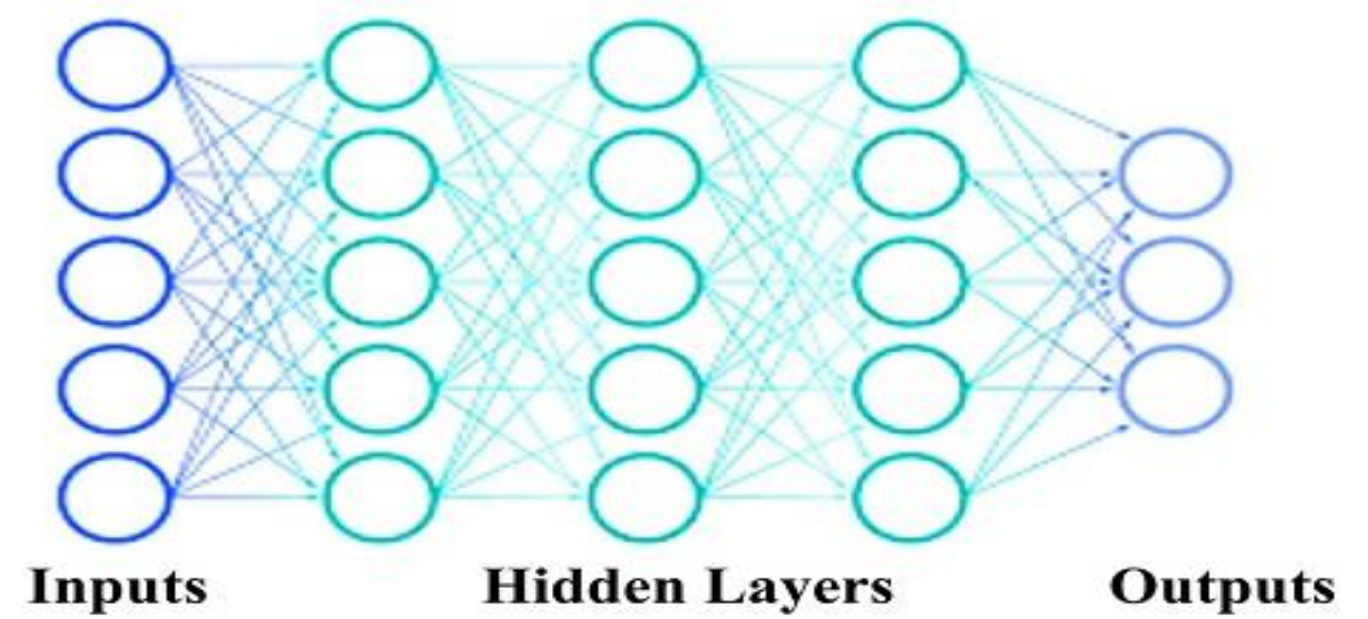

- https://www.mathworks.com/matlabcentral/fileexchange/137961-feed-forward-deep-learning-dl?s_tid=FX_rc2_behav

Convolutional Neural Networks (CNNs)

- CNNs are primarily used for image and video recognition. They consist of several layers designed to process images hierarchically. The first layer applies filters to extract features like edges or textures, which are then passed through fully connected layers for further analysis and classification

Component of Convolutional Neural Networks

- ***Convolutional layers*** detect patterns, features, and structures in the input data through filters that focus on specific traits.
- ***Pooling layers*** simplify and compress the information within certain areas, enhancing efficiency and making the network more resistant to input changes.
- ***Activation functions*** introduce non-linearity, allowing the network to model complex relationships by making the output not directly proportional to the input.
- ***Fully connected layers*** generate the final predictions by connecting all neurons in one layer to every neuron in the next.

Convolutional Neural Networks (CNNs)

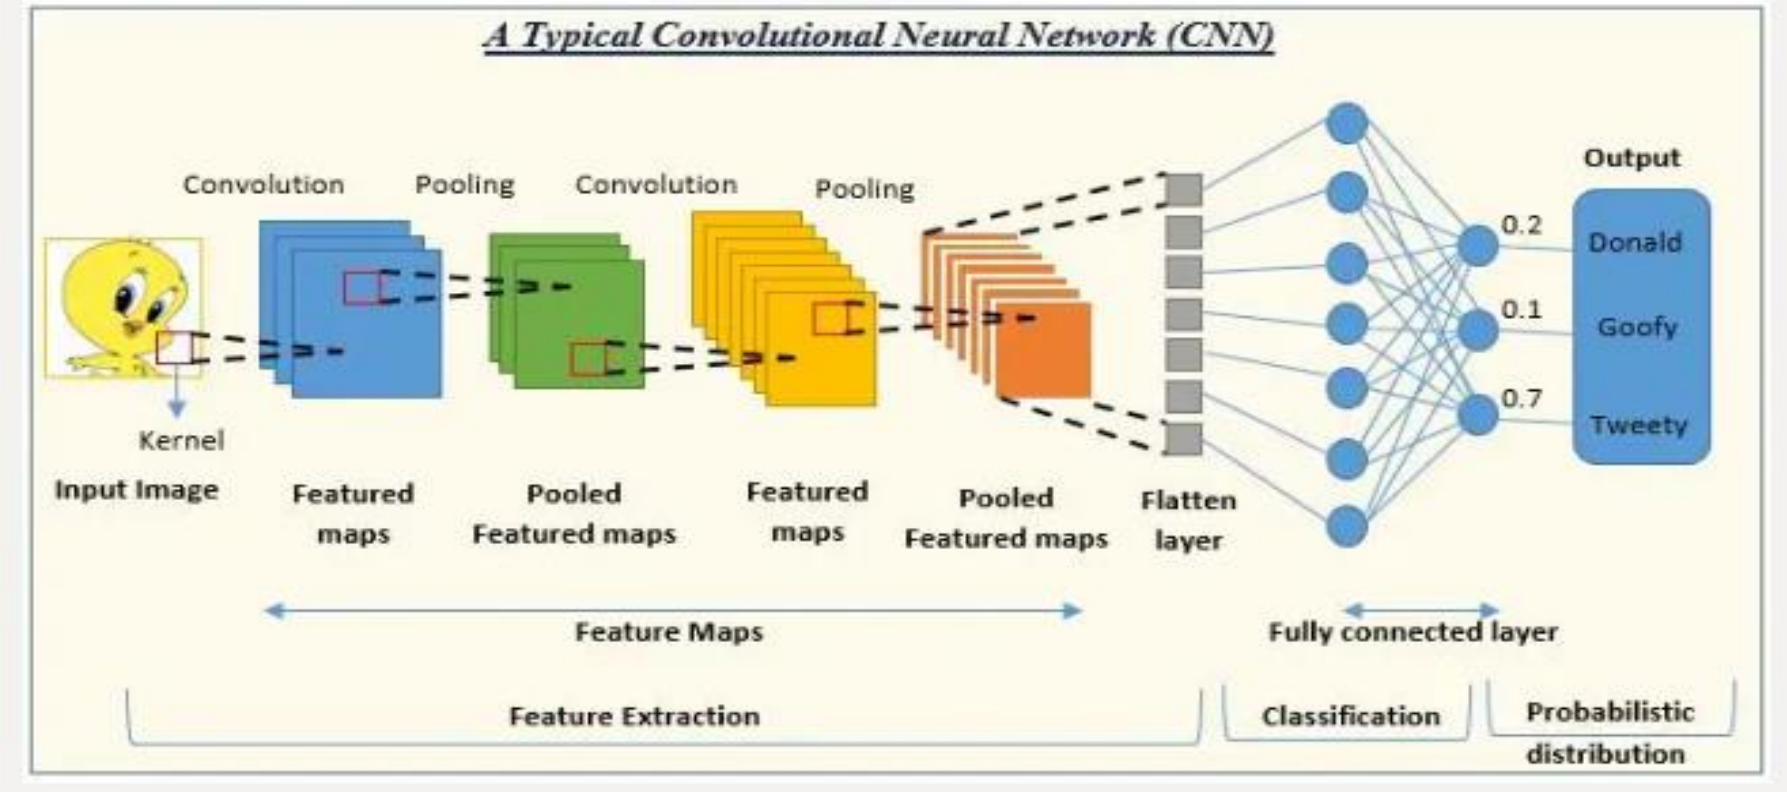

- https://glasswing.vc/blog/ai-atlas-16-convolutional-neural-networks-cnns/

Recurrent Neural Networks (RNNs)

- Recurrent Neural Networks (RNNs) are designed to process sequential data by maintaining an internal memory, unlike traditional feed_forward neural networks. This memory allows RNNs to identify patterns in sequences, making them ideal for tasks like natural language processing, speech recognition, and time series forecasting.
- RNNs consist of layers with interconnected nodes, where the output from one layer is fed back as input to the same layer for the next time step. This feedback loop enables RNNs to retain information from previous steps, helping them capture temporal dependencies and recognize patterns over time.

Recurrent Neural Networks (RNNs)

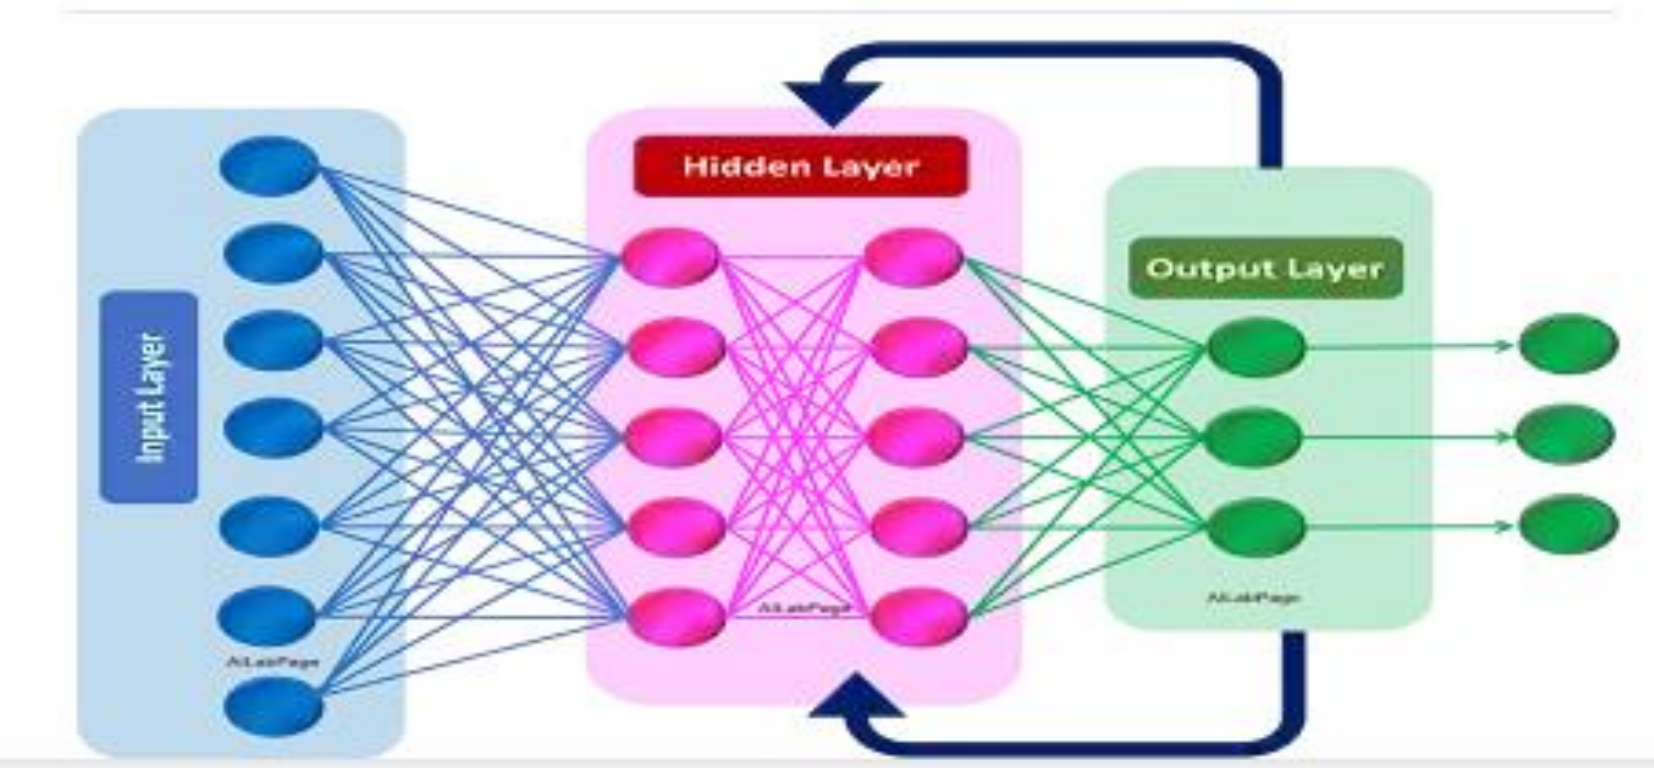

- https://botpenguin.com/glossary/recurrent-neural-network

Natural Language Processing (NLP)

- Natural Language Processing (NLP) is a branch of artificial intelligence (AI) that focuses on how computers can interact with humans using natural language. It involves understanding, analyzing, and generating human language in a way that is meaningful and useful to people.
- **Key Tasks in NLP:**
- **Text Classification:** Categorizing text into groups (e.g., spam or not spam).
- **Named Entity Recognition (NER):** Identifying and classifying entities like people, organizations, or locations in a text.
- **Sentiment Analysis:** Determining whether the sentiment of a text is positive, negative, or neutral.
- **Machine Translation:** Translating text from one language to another.
- **Question Answering:** Providing answers to questions posed in natural language.
- **Text Summarization:** Creating a concise summary of a longer text.

Natural Language Processing (NLP)

- **To perform these tasks, NLP algorithms use several techniques, including**
- **Tokenization:** Breaking text into smaller units (tokens), such as words or subwords.
- **Part-of-Speech (POS) Tagging:** Assigning grammatical categories (e.g., noun, verb, adjective) to each token.
- **Dependency Parsing:** Identifying grammatical relationships between words in a sentence.
- **Named Entity Recognition (NER):** Identifying and classifying named entities in text.
- **Sentiment Analysis:** Analyzing the tone of the text to identify its sentiment.

Applications of NLP

- NLP has a wide range of applications, including:
- **Chatbots:** NLP powers chatbots that can understand and respond to natural language queries.
- **Sentiment Analysis:** It helps analyze social media posts, customer reviews, and other texts to determine the sentiment behind them.
- **Machine Translation:** NLP enables translating text from one language to another, breaking down language barriers.
- **Voice Assistants:** NLP is used in voice assistants like Siri or Alexa, allowing them to understand and respond to voice commands.
- **Text Summarization:** NLP generates concise summaries of longer texts, such as news articles or research papers.

Challenges in NLP

- Despite advancements, NLP faces several challenges:
- **Data Quality:** NLP algorithms need high-quality structured data to produce accurate and useful results.
- **Naturalness:** The text generated by NLP must sound natural and be easily understandable by humans.
- **Domain Specificity:** NLP algorithms need to be customized for specific domains to generate precise and relevant content.
- **Personalization:** NLP must be able to tailor text generation to individual users for more personalized responses.
- **Evaluation:** Assessing the quality of generated text is difficult and requires subjective human judgment.

Risks of Generative AI

- **Misuse**: Generative AI can be exploited to produce fake content that is hard to distinguish from real content, such as deepfakes. This can be used to spread misinformation or create misleading material that could have severe repercussions.
- **Bias**: Since generative AI systems learn from data, any biases in the data can lead to biased outcomes. This could result in unfair or discriminatory practices in areas like hiring or lending.
- **Security**: Generative AI can be used to create new forms of cyberattacks, like realistic phishing emails or malware that can bypass traditional security systems.

Risks of Generative AI

- **Intellectual Property**: Generative AI can be used to create new works that may violate the intellectual property rights of others, such as using existing images or music to generate new content.
- **Privacy**: Generative AI has the potential to generate personal information about individuals, like realistic images or videos, which could be used for identity theft or blackmail.
- **Unintended Consequences**: Generative AI systems can produce unforeseen or unintended outcomes, such as creating new types of malware or causing harm to individuals or society.

Risks of Generative AI

- **Regulatory Challenges**: The use of generative AI raises regulatory issues concerning its development, deployment, and use, including questions about accountability and responsibility.
- **Ethical Concerns**: The use of generative AI brings up ethical questions, such as whether it is acceptable to create content that is indistinguishable from real content.

The Future of Generative AI

- **Improved Natural Language Processing (NLP):** Advances in NLP will enable more natural and contextually appropriate responses from chatbots, virtual assistants, and other AI-driven communication tools.
- **Increased Personalization:** As generative AI systems become more advanced, they will be capable of producing content that is better tailored to individual users. This could range from personalized news articles to dynamically generated video game levels.

The Future of Generative AI

- **Enhanced Creativity:** Generative AI is already being used to create music, art, and other creative works.
- **Better Data Synthesis:** As datasets grow more complex, generative AI will become an even more valuable tool for synthesizing and generating new data. This could prove especially useful in scientific research, where AI-generated data might help researchers identify patterns and connections that could otherwise be overlooked.
- **Increased Collaboration:** One of the most exciting possibilities for generative AI is its ability to enhance human creativity and collaboration

What is Machine Learning?

- Machine learning is a subset of artificial intelligence (AI) that involves the use of algorithms and statistical models to enable machines to learn from data and improve their performance on specific tasks.
- Machine learning involves the use of algorithms that enable machines to learn from data. These algorithms are designed to identify patterns and relationships in the data and use that information to make predictions or take actions.

Types of Machine Learning Algorithms
Supervised learning

- Involves training a model on a labeled dataset, where the correct output is known for each input. The model then uses this training data to make predictions on new, unseen data.
- **Linear Regression**
- Linear Regression is a statistical technique used for predicting continuous variables like sales, salary, age, or product price. It models the relationship between a dependent variable (y) and one or more independent variables (x), showing how the dependent variable changes in response to the independent variable(s). The method results in a straight line that represents this linear relationship.

Linear Regression

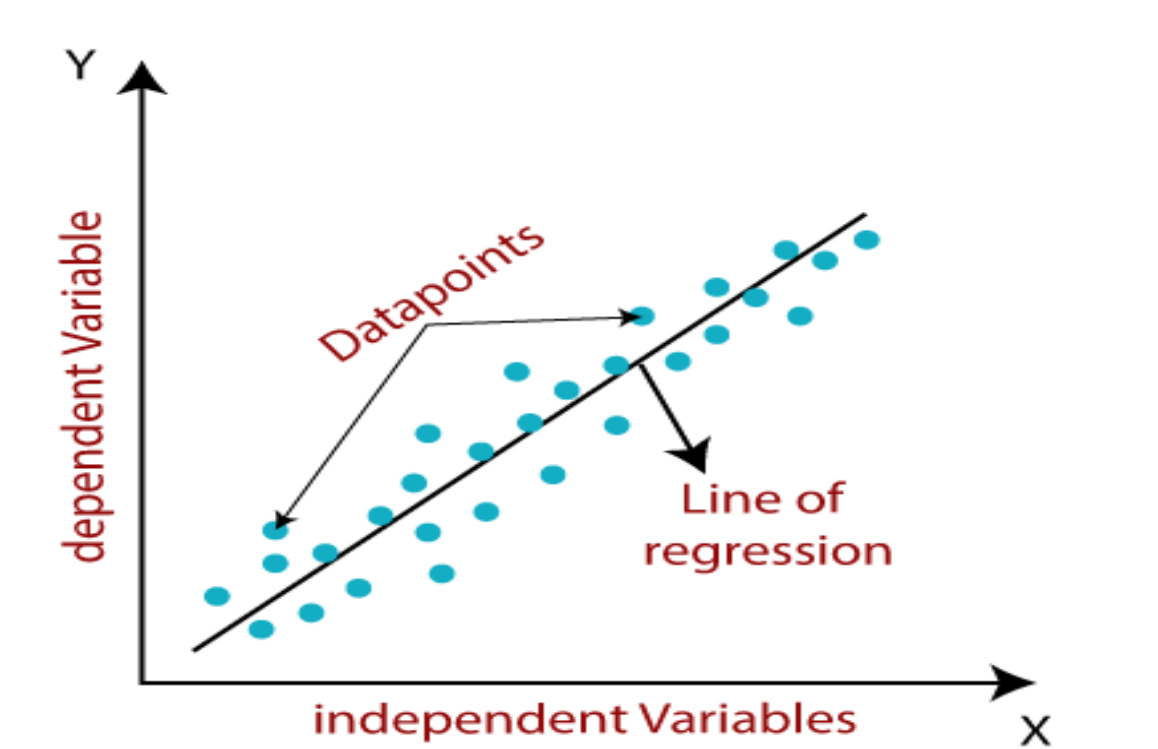

- https://www.tpointtech.com/linear-regression-in-machine-learning

Logistic Regression

- Used for binary classification tasks. It estimates the probability of a certain class label using a logistic function

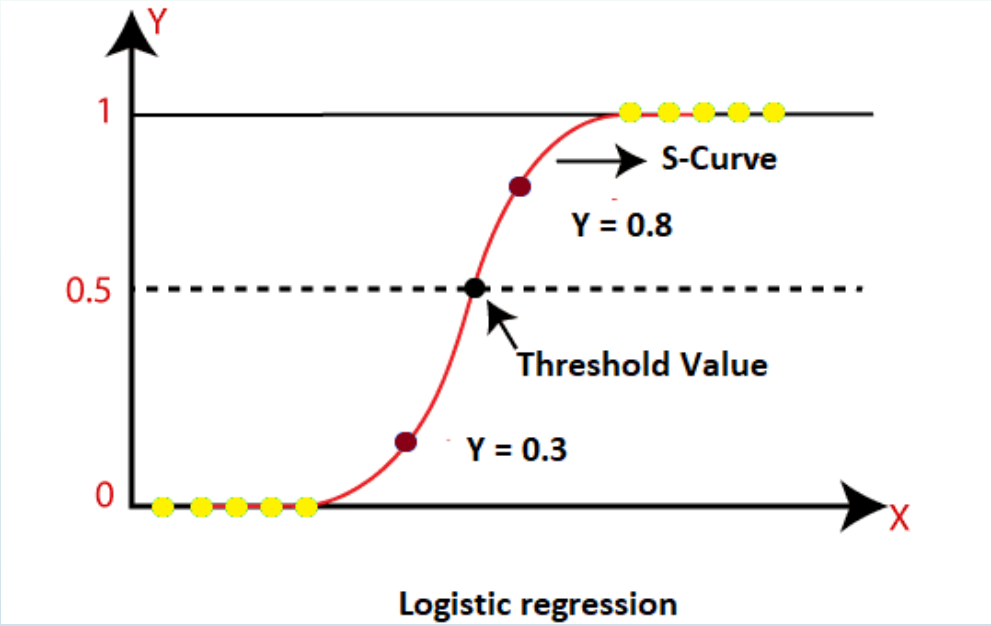

Logistic regression

https://tutorialforbeginner.com/logistic-regression-in-machine-learning

Examples

- **Predicting heart attack risk**: Medical professionals can use a logistic model to analyze factors like weight and exercise habits to predict the likelihood of a person experiencing a heart attack or other health issues.
- **University admission chances**: Application platforms can assess the probability of a student being accepted into a specific university or program by evaluating variables like GRE, GMAT, or TOEFL scores.
- **Spam email detection**: Email systems use logistic regression to filter and identify spam messages by analyzing predictor variables and determining the email's authenticity.

Decision Trees

- A tree-like model that splits data into subsets based on feature values, ultimately making predictions at the leaves of the tree. Decision trees are intuitive and easy to interpret but can easily overfit.

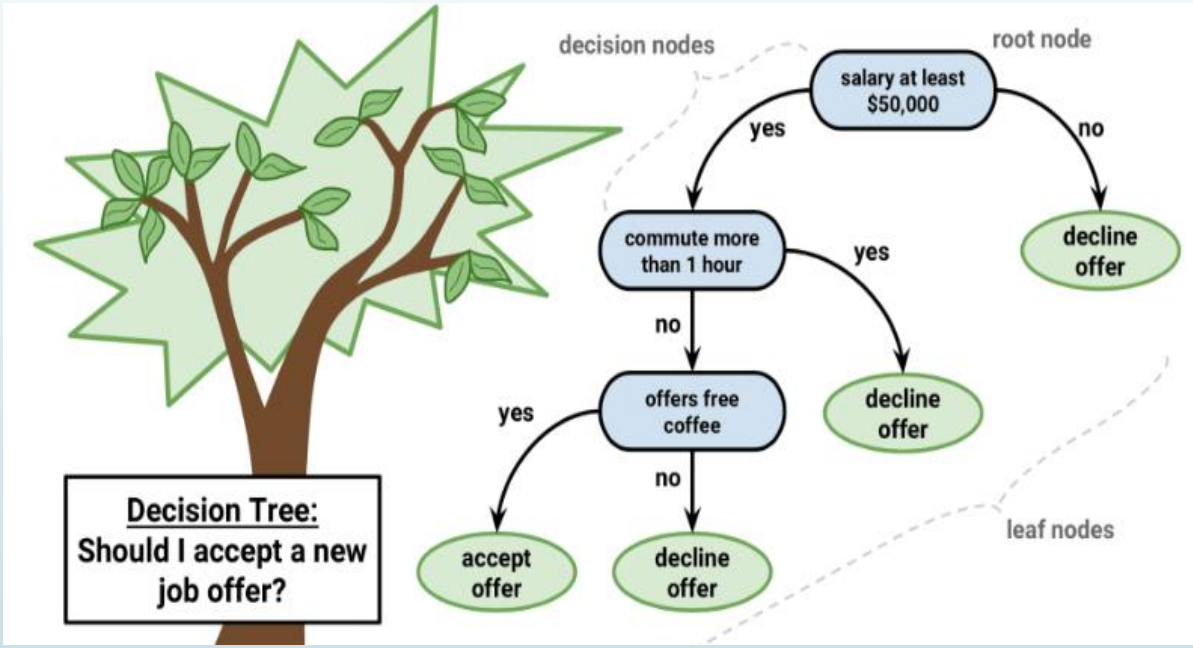

- https://okanbulut.github.io/bigdata/supervised-machine-learning-part-i.html

Random Forest

- Random Forest is a classifier that builds multiple decision trees using different subsets of the dataset and averages their results to enhance predictive accuracy. Rather than depending on a single decision tree, it combines the predictions from all trees and determines the final output based on the majority vote.

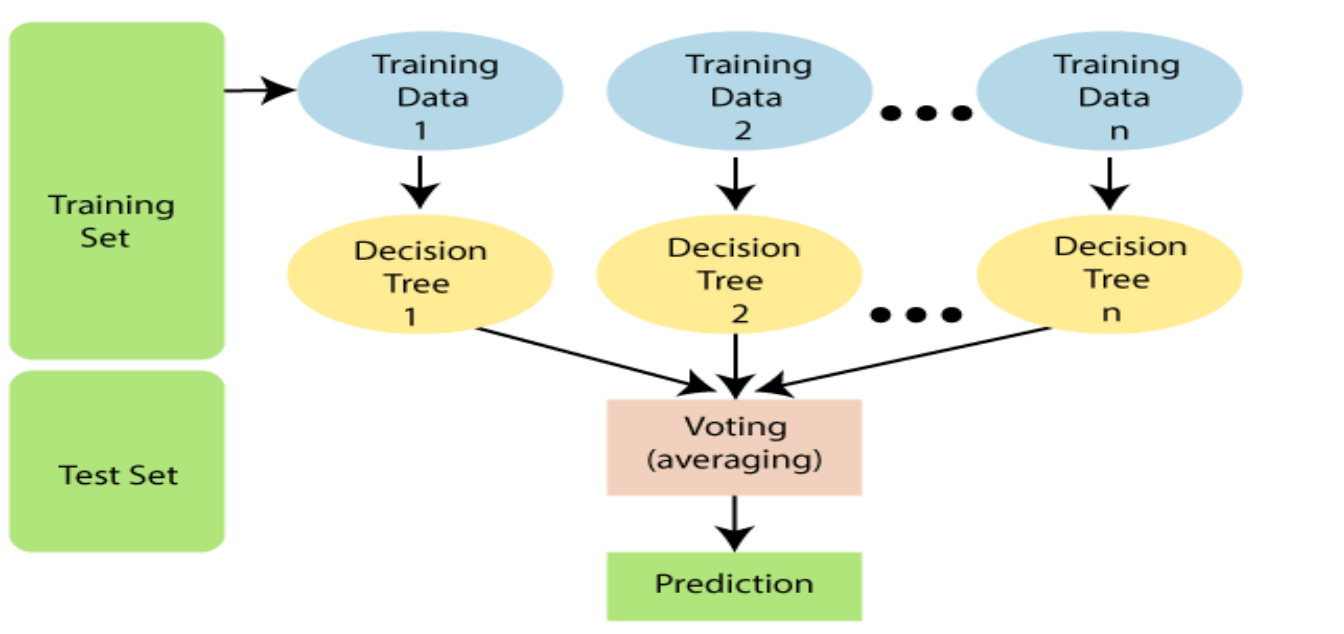

- https://www.tpointtech.com/machine-learning-random-forest-algorithm

Support Vector Machines (SVM)

- A powerful algorithm for classification tasks that works by finding the optimal hyperplane (boundary) that best separates data points into different classes.

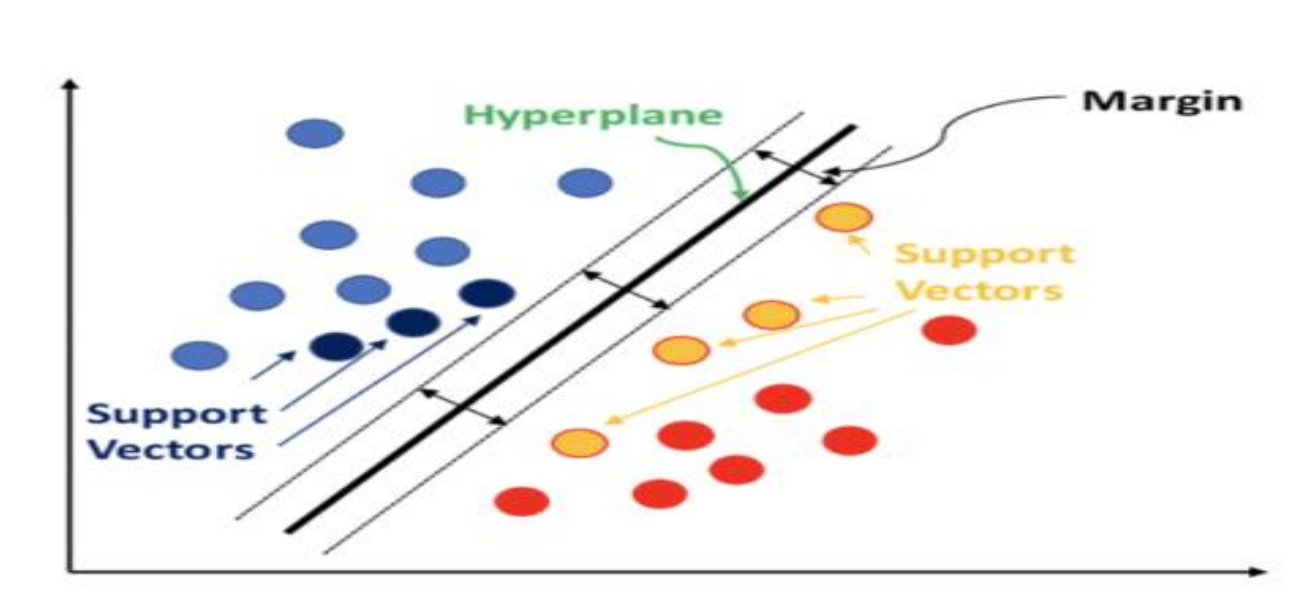

- https://datatron.com/what-is-a-support-vector-machine/

k-Nearest Neighbors (k-NN)

- The K-NN algorithm classifies a data point based on the majority class among its k nearest neighbors in the feature space.
- It assumes that new data points are similar to existing ones and assigns them to the most similar category.
- The algorithm stores all available data and classifies new data points by comparing them to existing ones, making it easy to categorize new data.
- K-NN can be applied to both regression and classification tasks, but it is primarily used for classification problems.

k-Nearest Neighbors (k-NN)

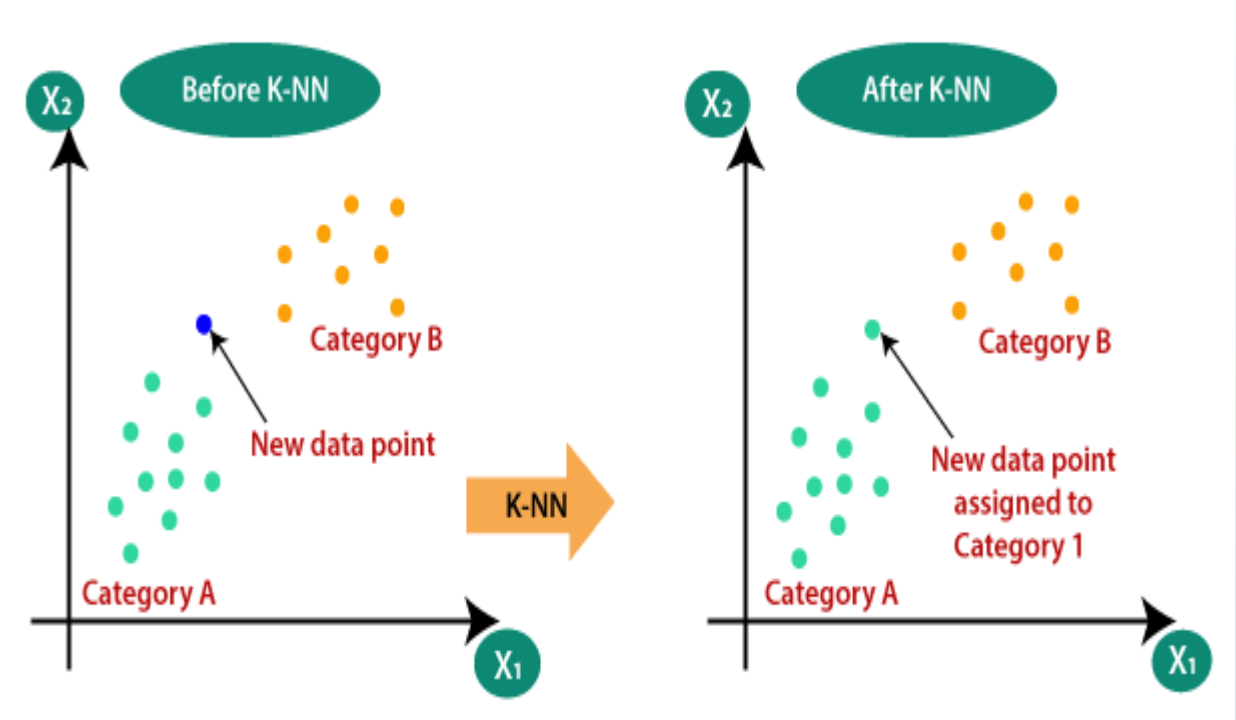

Naive Bayes

- This algorithm, based on Bayes' Theorem, is especially useful for classification tasks. It operates under the assumption that the features are independent, which is why it is referred to as "naive." It calculates the probability of a class based on the given input features. The term "naive" comes from the assumption that no factors other than the known prior conditions influence the classifier's predictions, an assumption that is not entirely realistic.

- P(B|A)=Probability that even B occurs if event A has already occurred
- P(A) = Probability that event A occurs
- P(B)= Probability that B occurs

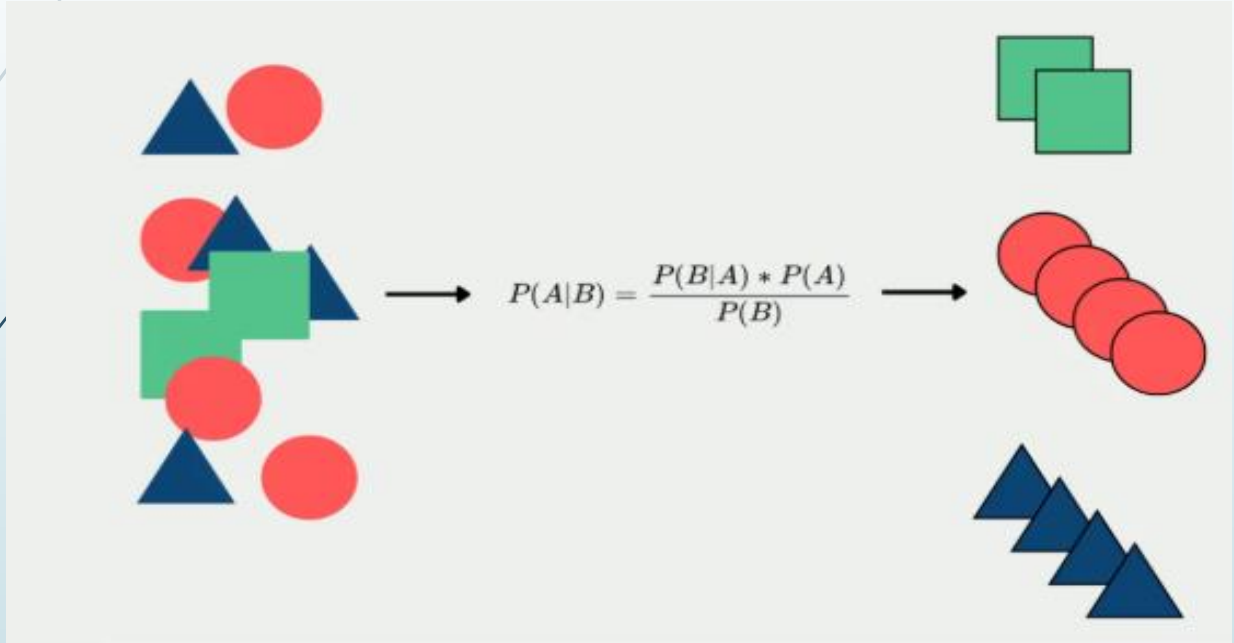

- https://databasecamp.de/en/ml/naive-bayes-algorithm

Unsupervised learning

- Involves training a model on a non-labeled dataset and allowing it to identify patterns and relationships on its own. This type of learning is often used in applications such as clustering (group similar items together).
- Cluster: a collection of data objects
  - Similar to one another within the same cluster but dissimilar to the objects in other clusters
- Cluster analysis
  - Finding similarities between data according to the characteristics found in the data and grouping similar data objects into clusters
- Unsupervised learning: no predefined classes

Examples of clustering applications

- Marketing: Help marketers discover distinct groups in their customer bases, and then use this knowledge to develop targeted marketing programs
- Land use: Identification of areas of similar land use in an earth observation database
- Insurance: Identifying groups of motor insurance policy holders with a high average claim cost
- City-planning: Identifying groups of houses according to their house type, value, and geographical location
- Earth-quake studies: Observed earth quake epicenters should be clustered along continent faults

Unsupervised learning

- **k-Means Clustering**: A technique for grouping data into k distinct clusters based on similarity. The algorithm assigns data points to the closest cluster center and repeats the process until the clusters become stable, effectively grouping similar objects together.

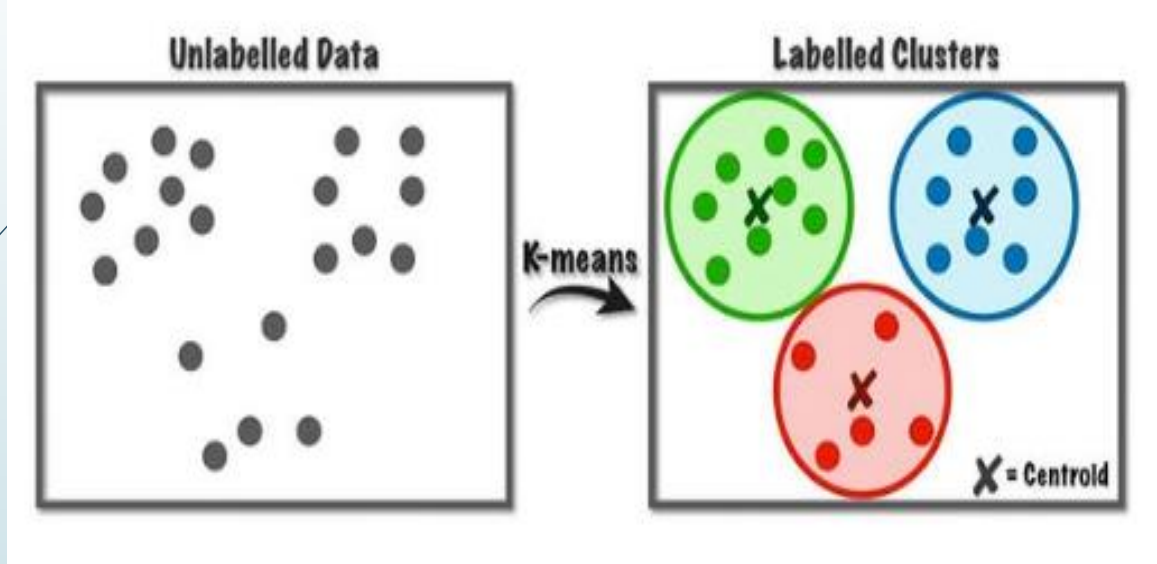

- https://www.instagram.com/data_science_learn/p/C-nQ9wtTUt3/

Hierarchical Clustering

- Constructs a hierarchy of clusters by either beginning with each point as an individual cluster or merging similar ones (agglomerative, bottom-up) or starting with all points in a single cluster and dividing them (divisive, top-down).

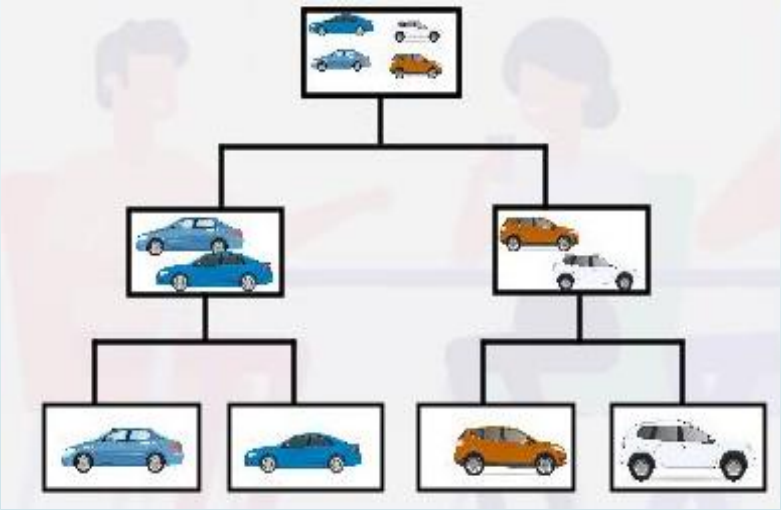

- https://www.simplilearn.com/tutorials/data-science-tutorial/hierarchical-clustering-in-r

Principal Component Analysis (PCA):

- A dimensionality reduction technique that transforms the data into a new coordinate system to reduce the number of features while retaining most of the data's variance.

Autoencoders:

- A type of neural network used to learn efficient representations of data (encoding) by compressing the input into a smaller dimension and then reconstructing the original input.

Reinforcement learning

- Involves training a model to make decisions based on feedback from its environment. In this type of learning, the model learns through trial and error, adjusting its actions based on the awards or punishments it receives.
- Q-Learning
- Deep Q-Networks (DQN)
- Policy Gradient Methods
- Actor-Critic Methods

AI Powered Tools

AI for Content Creation

- **ChatGPT (OpenAI)**
  A language model that can assist with content creation, writing, brainstorming, and even coding. It helps users generate text, answer questions, and more.
- **Copy.ai**
  Another AI-powered tool that helps generate copy for ads, social media posts, email templates, and blog content.
- **Grammarly**
  An AI-powered writing assistant that checks for grammar, spelling, punctuation, clarity, and even tone in your writing.

AI for Design

- **Canva**
  Canva's AI-powered design tools help create graphics, presentations, and marketing materials easily. It includes features like AI-generated design suggestions, templates, and more.
- **Runway ML**
  A creative toolkit that integrates AI-powered tools for video editing, image generation, and motion graphics, making it easier for creators to produce content.

AI for Data Analysis

- **Tableau**
  An AI-powered data visualization tool that helps businesses analyze data and create intuitive visual dashboards to derive insights from large datasets.
- **Google Analytics**
  Uses AI to analyze web traffic and user behavior, offering insights into website performance and audience engagement.

- **AI for Marketing and SEO**
- **Surfer SEO**
  An AI-powered SEO tool that provides recommendations to improve a website's ranking on search engines by analyzing content, backlinks, and keyword trends.
- **Market_Muse**
  AI-driven content research and optimization platform that helps marketers create high-quality content by providing data-driven insights into keyword strategies and topic relevance.
- **AI for Customer Support**
- **Zendesk**
  A customer service platform that uses AI-powered chatbots and automated responses to provide timely and personalized customer support.
- **Drift**
  An AI-powered conversational marketing and sales platform that automates live chat on websites, allowing for real-time engagement with customers.

AI for Code and Development

- **GitHub Copilot**
  An AI-powered code completion tool that helps developers by suggesting entire lines or blocks of code, speeding up the development process.

AI for Image and Video Editing

- **DeepArt**
  An AI tool that turns photos into artwork by using deep neural networks to replicate famous painting styles.
- **AI for Productivity and Task Management**
- **Trello with Butler**
  Butler is an AI-powered automation tool for Trello that helps users automate repetitive tasks, improving project and task management.

AI for Translation and Language Processing

- **Google Translate**
  AI-driven translation tool that supports translation across many languages and is continually improving in accuracy and fluency.
- **DeepL**
  A powerful AI translation tool that is known for its high-quality translations and contextual accuracy, often preferred by professional translators.

AI for Personalization

- **Persado**
  An AI tool that uses natural language generation to create personalized messages for marketing, helping brands communicate more effectively with their audience.

References

Banafa, A. (2024). *Introduction to Artificial intelligence.* New york: River Publishers.

Russel, S. (2022). *Artificial Intelligence Amodern Approach.* United Kingdom: Pearson Education Limited.

https://glasswing.vc/blog/ai-atlas-16-convolutional-neural-networks-cnns/

https://www.tpointtech.com/machine-learning-random-forest-algorithm

https://databasecamp.de/en/ml/naive-bayes-algorithm

Next Lecture

- Quantum Computing