AI OVERVIEW

Artificial Intelligence (AI)	Computer programs that are able to imitate human intelligence. This includes learning, problem-solving, and perceiving at a level beyond that which was programmed.
Algorithm	A well-structured procedure that enables computers reliably to execute tasks.
Machine Learning (ML)	AI programs that learn patterns in data to make predictions, classify, or otherwise postulate. This learning can then be applied to solve problems with new data.
ML Model	A specific ML algorithm.
AI Ethics	Considers ethical impacts of the development and implementation of AI. This includes bias, privacy, regulation, environmental/social concerns, and much more.
AI Sustainability	Considers the environmental impact from the training, large-scale use, and continual fine-tuning of AI models. These models pose threats by consuming electricity, wasting water, and emitting carbon.
Black Box Algorithms	A phenomenon where inputs and outputs of an AI model are clear, but the model's process of arriving at the output from the input is unknown.
Explainable AI	Attempts to make how models come to their conclusions more transparent, as many AI algorithms are known to be "Black Boxes."
Trustworthy and Responsible AI	AI models that are transparent, fair, accurate, and ethical.
Artificial General Intelligence (AGI)	AI that is capable of intelligence equivalent to that of humans.
Chatbot	Computer programs that mimic human conversation.
Expert Support System	Decision-based system that is designed to imitate an expert.
Bias	Occurs when a model erroneously places more predictive emphasis on certain data, which leads to an unfair or discriminatory outcome. Bias is more common when certain data is overrepresented in a dataset.
Large Language Models (LLMs)	Special types of language models that specialize in processing and generating language (e.g. ChatGPT).
Hallucinations	Occur when an LLM presents false data as true. If the error stems from the model misinterpreting accurate data, the hallucination is considered "intrinsic" (an internal problem with model function). If the model correctly interprets the meaning of a faulty source, the hallucination is known as "extrinsic" (a problem external to the model).

Human in the Loop	A process in which human oversight and input guides output of AI models for the purpose of verifying output and increasing productivity.
Intelligent Computation	Leverages AI to execute computations involving large, complicated data.
Pattern Recognition	Ability of systems to identify and use patterns. These may take the form of relationships or structures within data.
Predictive Analytics	Process of using past data to predict the future outcomes.
Prompt	User input that gives instructions to a model.
Prompt Engineering	Refining prompts to obtain a desired output.
Recommendation Systems	A class of machine learning models that suggest new material to individuals based on previous behavioral data.
Robotics	The intersection of science, engineering, and technology that produces machines, called robots, that replicate or substitute human actions.
Turing Test	A test that determines machine intelligence. If a machine can engage in conversation with a human without being detected as a machine, it is considered intelligent.
Virtual Assistants	AI that processes user requests to help complete tasks and answer questions.
Decision Support System	ML models that analyze data with the objective of advising decisions.
Computer Vision	Enables machines to interpret and analyze meaningful data from images or videos.

MODELS AND TASKS

Artificial Neural Network (ANN) (1)	A machine equivalent to observed neural structures in humans. Neural networks receive input data first, which is then passed into another "layer" of nodes. Together, this layer of nodes evaluates aspects of the input to better determine the desired output. The number of layers depends on the model. Once the input has been thoroughly evaluated by the intermediate (also called "hidden") layers, the model provides its output.
Deep Learning	ANNs that incorporate many layers of neurons.
Association	Aims to find related variables in large datasets by observing their co-occurrences.
Classification (2)	The task of fitting data into categories. Dissimilar data must be distinguished, and similar data must be grouped together.
Regression (3)	Fits data onto a scale of continuous values. Commonly referred to as a "line of best fit," regression problems try to find a curve that accurately represents trends in data. The curve is selected by determining which function would have the least distance between points on the curve and data points.
Clustering (4)	Distinguishes data points from one another in order to identify those that are similar. Unlike classification, clustering is performed on unlabeled data. ML models try to group data points based on their similarities and differences, but the data has not been named. For example, a task could involve grouping similar photographs in a dataset where the contents of the photos are not labeled.
Training Data	Data sets that are used to train ML models.
Pre-Training	Process of training models on vast amounts of data. Models can then be adjusted for specific tasks. This eliminates the need to create multiple models for related tasks.
Overfitting	Occurs when training data is memorized instead of finding patterns. This translates poorly to new data. In general, models must have fewer parameters than data points on which they are trained.
Underfitting	Occurs when training data is not analyzed enough. Patterns in data are not learned, leading to poor performance on new data.
Data Mining	Identifying patterns by analyzing large datasets.
Structured Data	Data that is organized in a manner that is easy to process and understand.
Big Data	Enormous datasets, often involving many data types, that increase in size at high rates (e.g. internet search data).
Foundation Models	General AI models trained on large data that can be applied to a number of more specific tasks.

Hyperparameters	Factors that define the network's structure and learning process, like the number of hidden layers in a neural network. These can be adjusted to improve model accuracy.
Tuning	Adjusting parameters of models for better performance.
Decision Tree (5)	A tree-like structure where classification/regression is determined by decision rules correlating variables with outcomes. The variable with the highest predictive influence on the result is at the top of the tree. Depending on the presence or value of this variable, other variables will then be considered. The tree continues to extend similarly until it reaches its max depth (the maximum number of decision nodes).
Supervised Learning (6)	ML models trained on labeled data. The model will use independent variables in a dataset to determine a dependent variable. Since the data is labeled, this dependent variable will be present in the dataset. The model's result can then be compared with the actual result in the data to determine its accuracy.
Unsupervised Learning (7)	ML models that find patterns in unlabeled data. In this case, since the data is not labeled, a model's results cannot be compared with actual data. Models must be tested on their ability to recognize patterns in the data rather than having their results measured directly against real data.
Generative Pre-trained Transformer (GPT)	OpenAI-developed pre-trained models that are tuned for applications.
Natural Language Processing (NLP)	Models that process human language by determining meaning and sentiment (often using deep learning).
Sentiment Analysis	A type of natural language process dedicated to determining emotional intent behind written works.
Recurrent Neural Network (RNN) (8)	An ANN best used for ordered, sequential input and output data. This is useful for processing data whose order has significance, like words in text. In sequential data, outputs depend on previous data points in the sequence, differentiating it from typical neural networks.
K-Nearest Neighbors (KNN)	Classifies data by averaging the values of the k-nearest data points.
Multilayer Perceptron (MLP)	Deep learning models that are fully connected; each neuron in one layer connects to all neurons in the next layer. Commonly used for classification and regression tasks, the model trains by passing the error in its output back through the model infrastructure and making adjustments.
Optical Character Recognition (OCR)	Converts images of text into machine-readable text. This enables the interaction with text on a computer, like editing and copying.
Digital Twins	Virtual representation of an object or system designed to reflect a physical object accurately.

HIGHER-LEVEL CONCEPTS

Activation Function (9)	Calculates the output of a node in a neural net based on inputs and weights. Being nonlinear, these functions help models explore more complicated relationships that often appear in data.
Attention (10)	Allows the models to put greater weight on more important input elements rather than processing all information in sequence.
Autoencoders (11)	Encodes (compresses) data for more efficient processing, then decodes processed data for output. This process is designed to be a shortcut, reducing computational expense while maintaining accuracy.
Bag-of-words (12)	Disregards order of words and counts the number of their occurrences. A representation of this variety allows machines to determine the general content of text, often being used for its classification.
Backpropagation (13)	Adjusts how the neurons are activated based on the error of a training output. The neuron activation is varied until error is acceptably reduced.
Bidirectional Encoder Representations from Transformers (BERT) (14)	Processes text (via encoding only) in two directions (bidirectionally) simultaneously to determine context of a word by those that come before and after.
Convolutional Neural Network (CNN) (15)	Convolves over pixel data for image recognition. Pixels are represented as a matrix of numbers. CNNs slide a filter (another matrix) over the image, performing element-wise multiplication and summations to discern image features.
Diffusion Model (16)	Trains generative models by gradually turning data to noise. The model is then tasked with generating output from the noisy data.
Encoding (17)	Process by which words are converted to vectors so NLP models can engage with them more easily. The direction of the vectors indicates their meaning. For instance, "human" would be closer to "man" or "woman" than it would to "dog" or "cat."
Fine Tuning	Adjusts models to new data in transfer learning to improve performance. This process ensures higher accuracy of the general model when it is applied to specific data.
Generative Adversarial Network (GAN) (18)	A generator creates text or an image from random data, which it then sends to a discriminator. The latter determines whether the image is real or fake by comparing the generated image to real data. This feedback is given to the generator, which adjusts the random input to produce a more realistic output. The process repeats, with the generator trying to create realistic enough images to fool the discriminator.

Gradient Descent (19)	An algorithm that minimizes the loss function. Minimization is achieved by adjusting weights and other model hyperparameters.
Long-Short Term Memory (LSTM) (20)	Compensates for RNN's inability to manage wide input dependencies across large data. Such dependencies may be present in time series, text, and similar data.
Loss Function	Mathematical function that measures deviation from expected output.
Normalization	Simplifying data by converting words to lowercase, removing punctuation, and applying additional methods to standardize text.
Principal Component Analysis (PCA)	Reduces high-dimensional data by finding correlation between variables and preserving only the largest eigenvalues of the covariance matrix. Left are only the principal components, or most important aspects of the original dataset. This streamlines the process of working with the data.
Reinforcement Learning (21)	A trial-and-error training method using positive reinforcement for ideal output. The agent in the model takes action in the environment and is either rewarded or punished. The model then learns how to generate desirable output.
Retrieval-Augmented Generation (22)	Bolsters LLMs' performance by combining generation with external knowledge.
Self-Attention (23)	Weighs words based on their relevance to other words in the input (often used by transformers for context).
Stemming	Removing prefixes/suffixes from words to determine meaning without unnecessary information.
Stop Words	Removing any words that do not contribute to the meaning of the text (e.g. 'the').
Tokenization	Split large amounts of text into words or phrases. Tokenization is a step in enabling machines to process text data numerically.
Transfer Learning (24)	Transfers what a model has learned to a new task. This process is done to improve model performance on new tasks or data.
Transformers	Transforms input data to output by determining relationships in input data. These models use attention to process input data.
Variational Autoencoder (25)	Generates new data by varying existing data using autoencoding.
Weight	Significance a model places on input elements. Neural networks adjust weights to generate outputs.
Zero-Shot Learning(26)	Testing models on tasks they have not been trained to do.

VISUALIZATIONS

(1) **ANN**

https://en.wikipedia.org/wiki/Neural_network_(machine_learning)

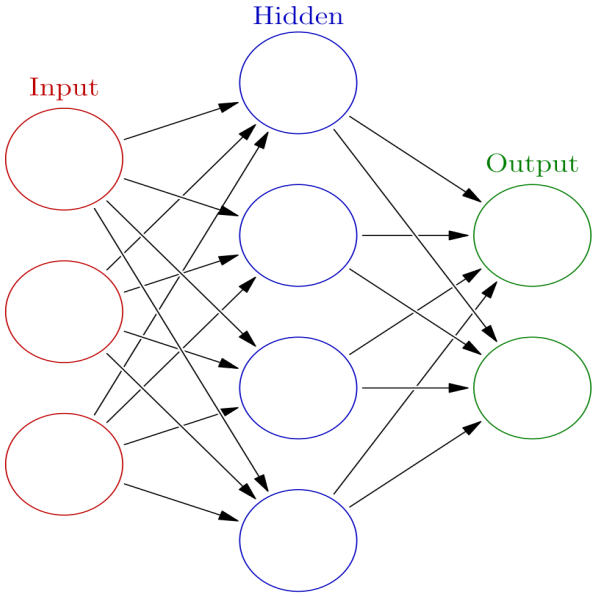

(2) **Classification**

https://www.appliedaicourse.com/blog/classification-in-machine-learning/

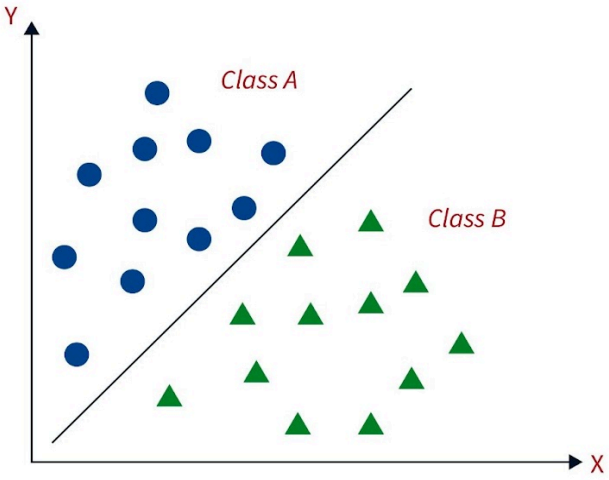

(3) **Regression**

https://web.stanford.edu/class/stats202/notes/Linear-regression/Simple-linear-regression.html

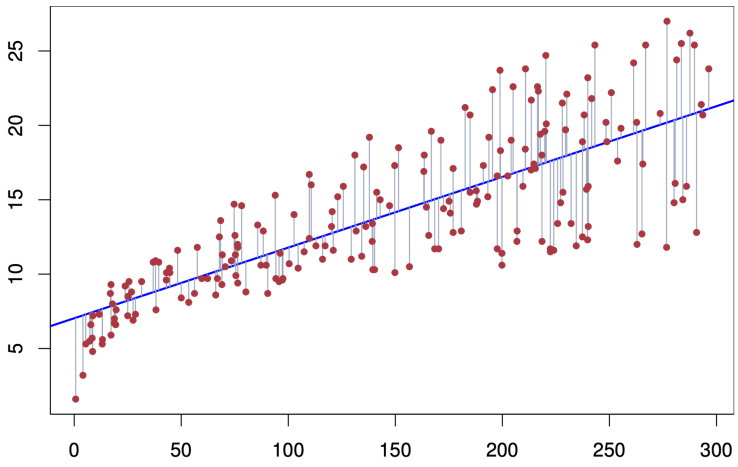

(4) **Clustering**

https://www.geeksforgeeks.org/data-science/measuring-clustering-quality-in-data-mining/

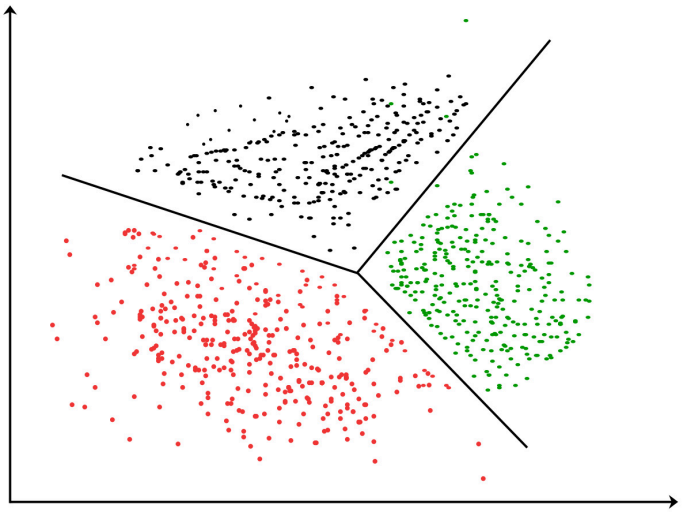

(5) **Decision Tree**

https://www.techguruspeaks.com/decision-tree-dt/

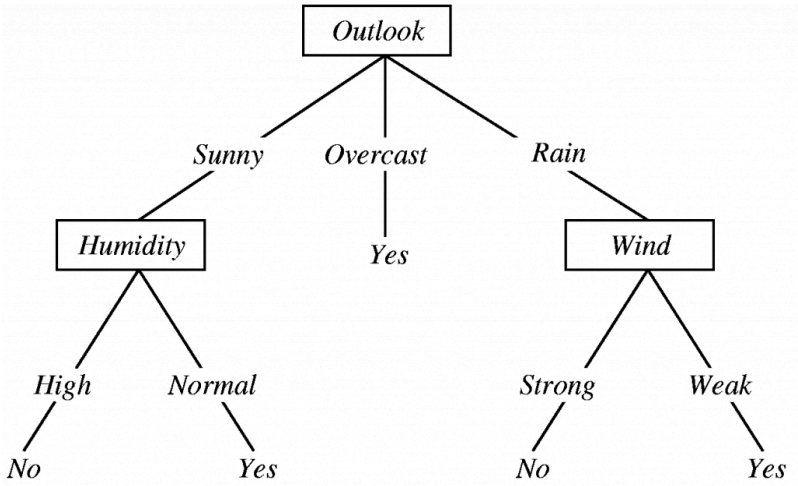

(6) **Supervised Learning**

https://en.wikipedia.org/wiki/Supervised_learning

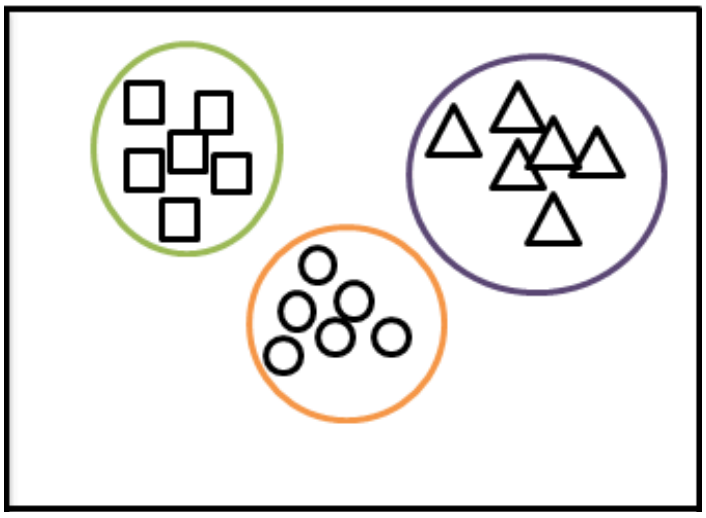

(7) **Unsupervised Learning**

https://en.wikipedia.org/wiki/Supervised_learning

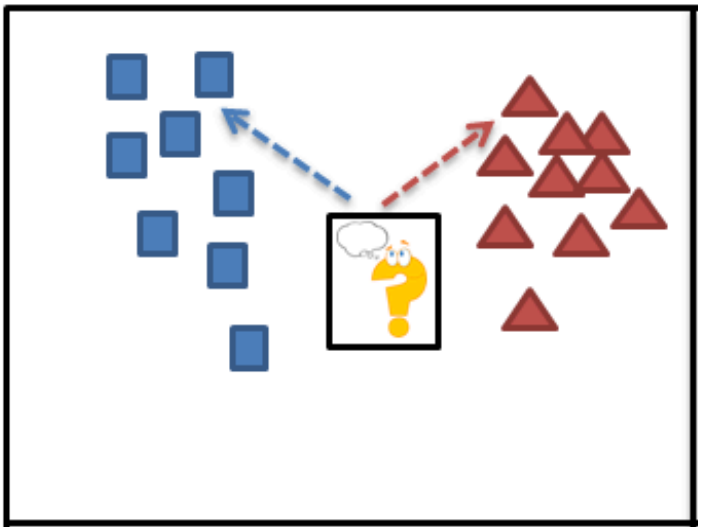

(8) **RNN**

https://www.nature.com/articles/s41377-024-01590-3

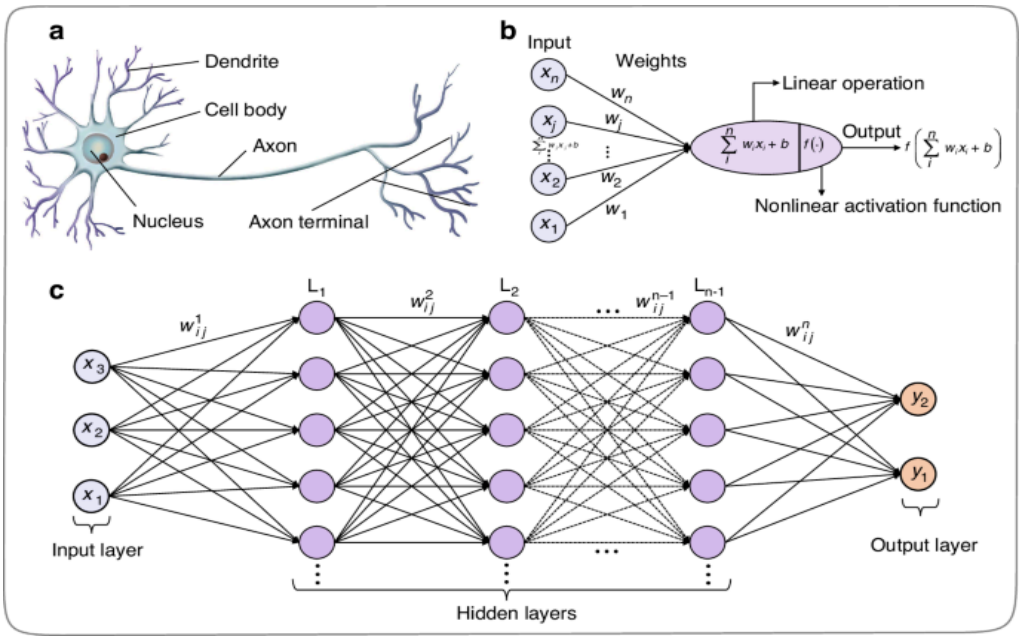

(9) **Activation Function**

https://medium.com/@shrutijadon/survey-on-activation-functions-for-deep-learning-9689331ba092

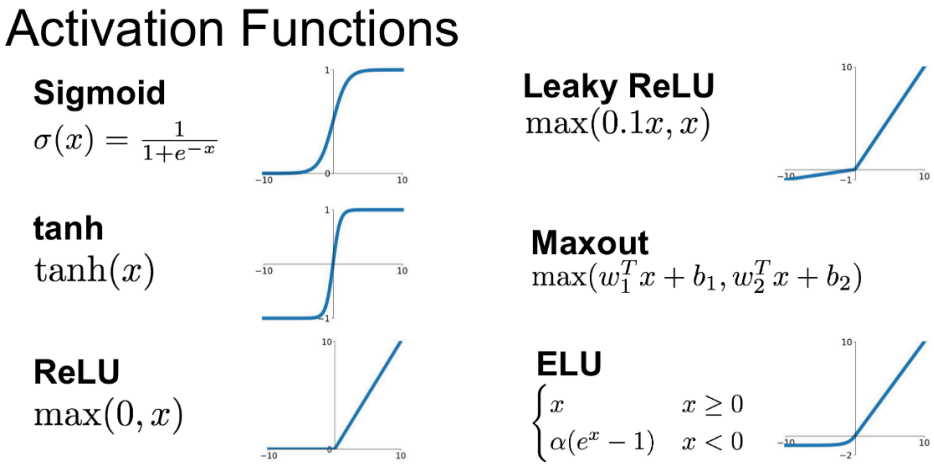

(10) **Attention**

https://medium.com/@vudhanthineeraja/the-impact-of-attention-mechanisms-and-transformers-in-recurrent-neural-networks-9d24f2246e1e

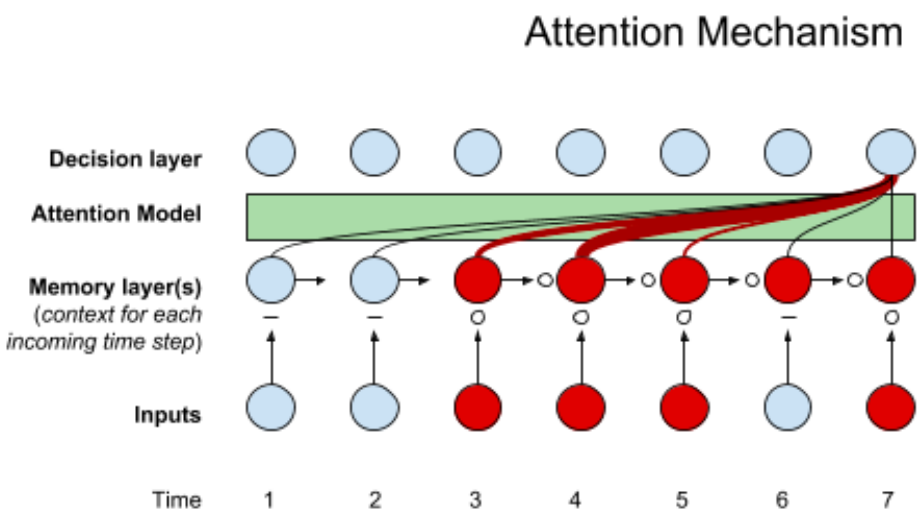

(11) **Autoencoders**

https://link.springer.com/chapter/10.1007/978-3-031-54049-3_7

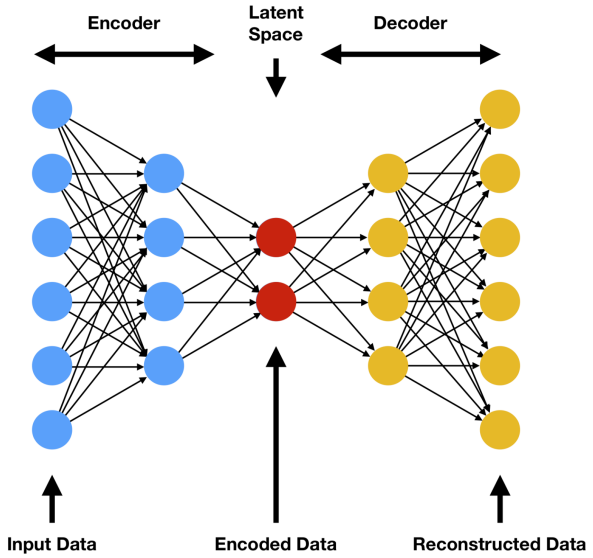

(12) **Bag-of-Words**

https://medium.com/@joydeepsfs10/evolution-of-language-representation-techniques-a-journey-from-bow-to-gpt-68d3b1b525b4

Document D1	*The child makes the dog happy* the: 2, dog: 1, makes: 1, child: 1, happy: 1
Document D2	*The dog makes the child happy* the: 2, child: 1, makes: 1, dog: 1, happy: 1

↓

	child	dog	happy	makes	the	BoW Vector representations
D1	1	1	1	1	2	[1,1,1,1,2]
D2	1	1	1	1	2	[1,1,1,1,2]

(13) **Backpropagation**

https://medium.com/@lmpo/backpropagation-the-backbone-of-neural-network-training-64946d6c3ae5

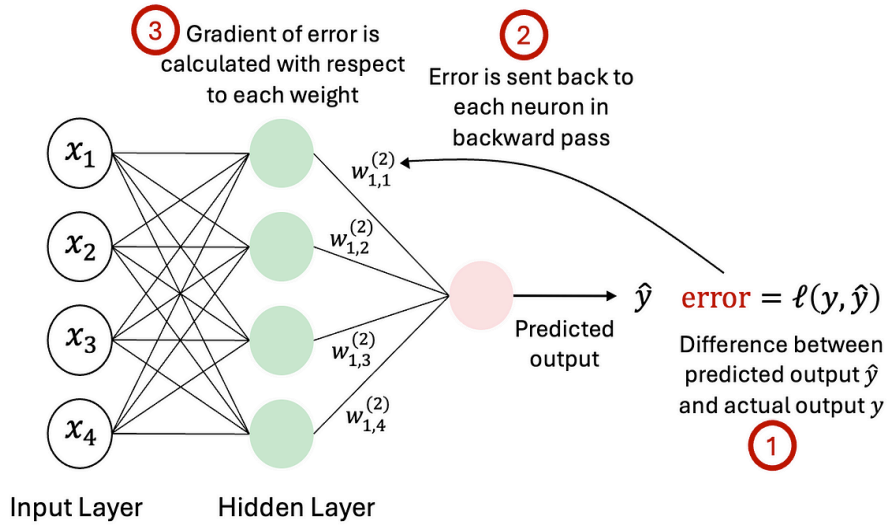

(14) **BERT**

https://www.aibutsimple.com/p/bert-bidirectional-encoder-representations-from-transformers

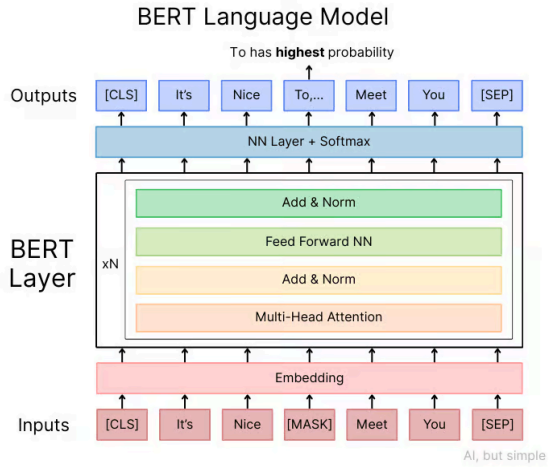

(15) **CNN**

https://www.sciencedirect.com/science/article/pii/S2772375523000394

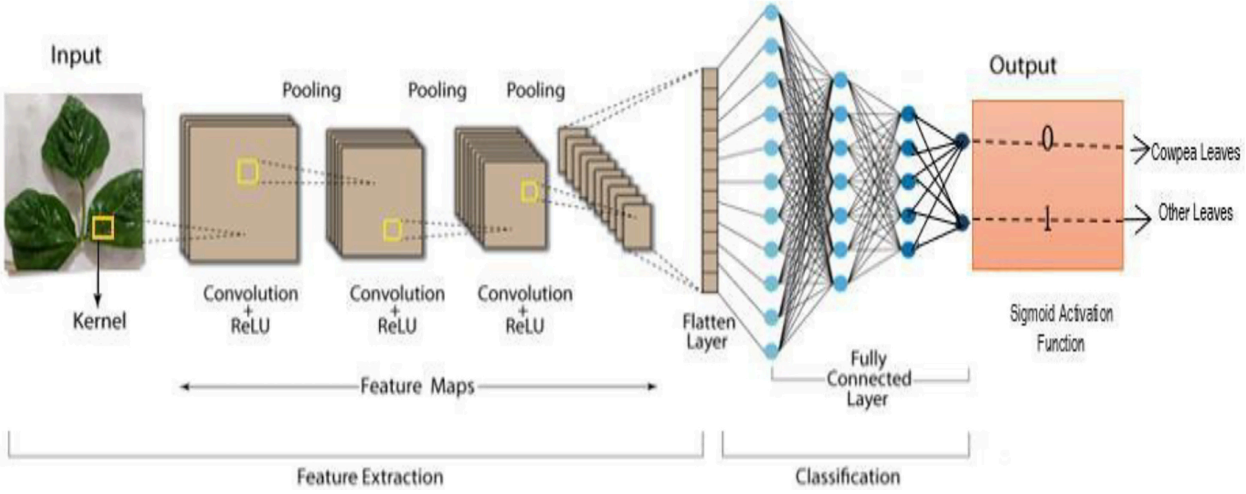

(16) **Diffusion Model**

https://jehillparikh.medium.com/understanding-diffusion-models-optimization-objective-d7214d36dbd8

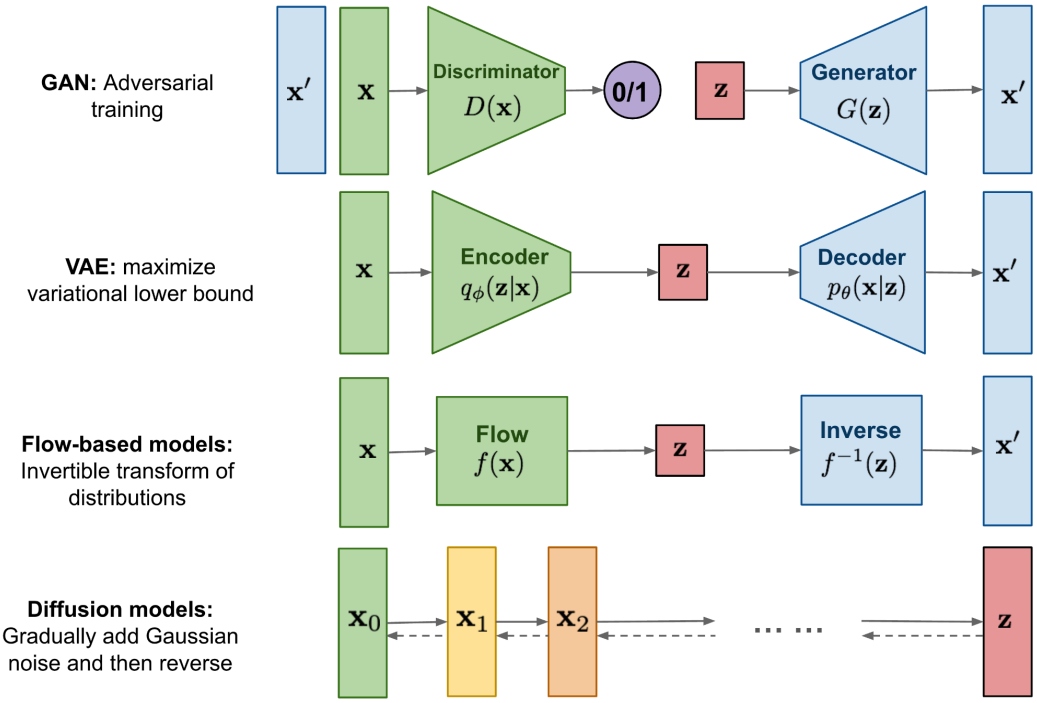

(17) **Encoding**

https://ai-ml-analytics.com/encoding/

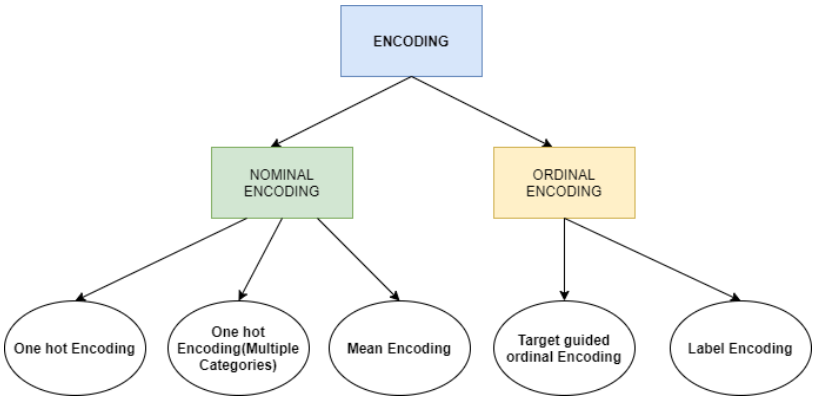

(18) **Generative Adversarial Network (GAN)**

https://ankittaxak5713.medium.com/basics-of-generative-adversarial-network-459e5d86c1f0

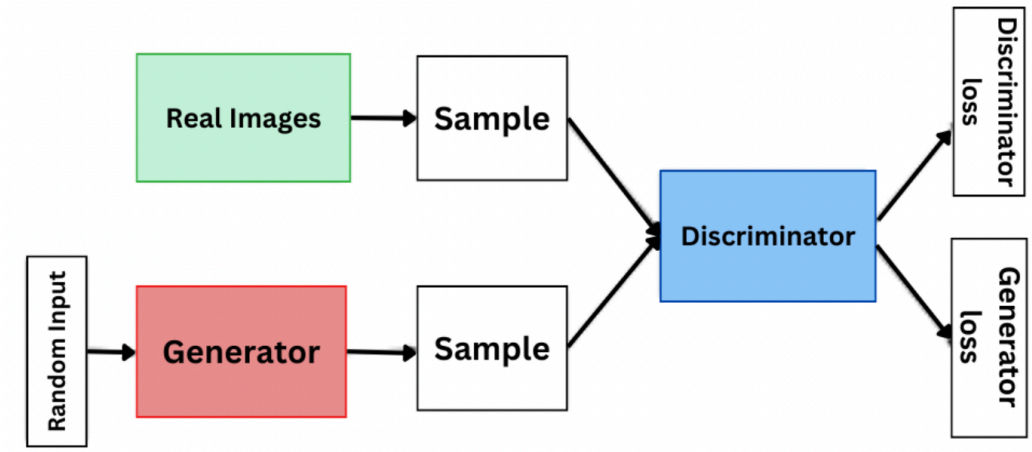

(19) **Gradient Descent**

https://medium.datadriveninvestor.com/an-overview-of-gradient-descent-algorithms-e373443afa7f

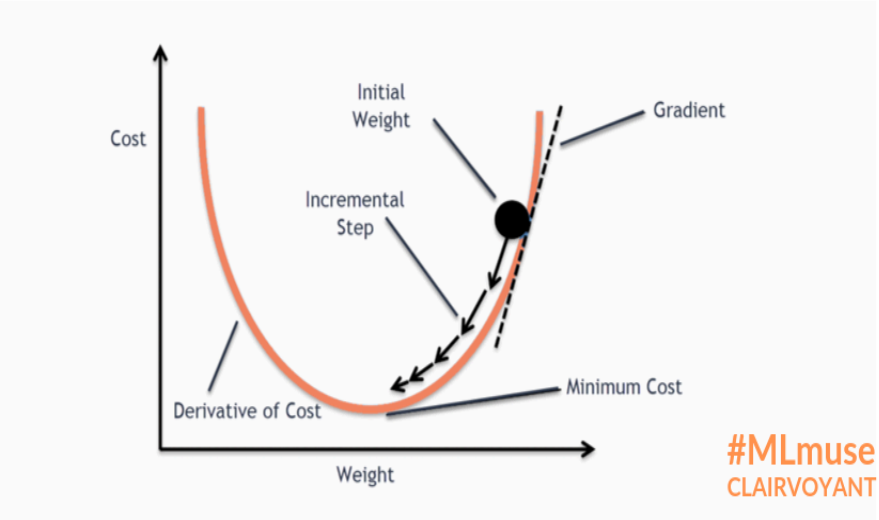

(20) **Long-Short Term Memory (LSTM)**

https://www.researchgate.net/figure/Architecture-of-LSTM-block_fig1_374011326

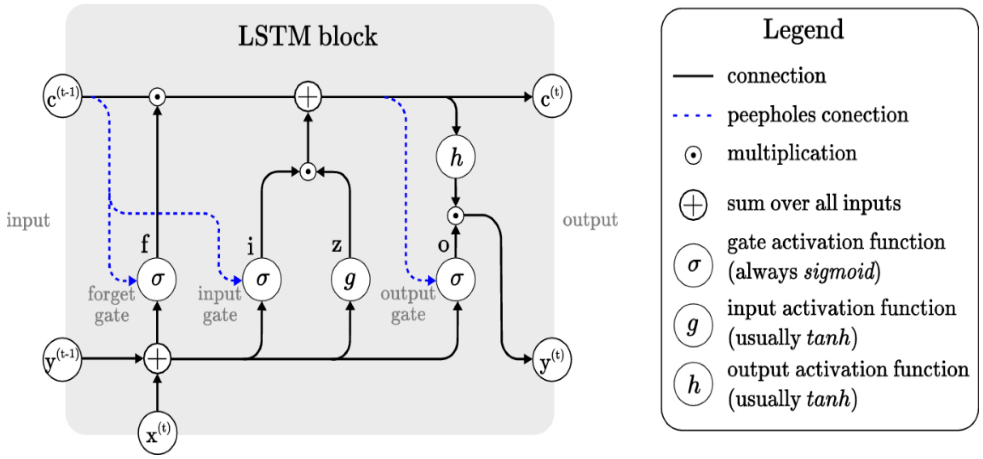

(21) **Reinforcement Learning**

https://medium.com/hub-by-littlebigcode/supervised-unsupervised-and-more-an-introduction-to-lesser-known-machine-learning-methods-and-47c7dd4694d

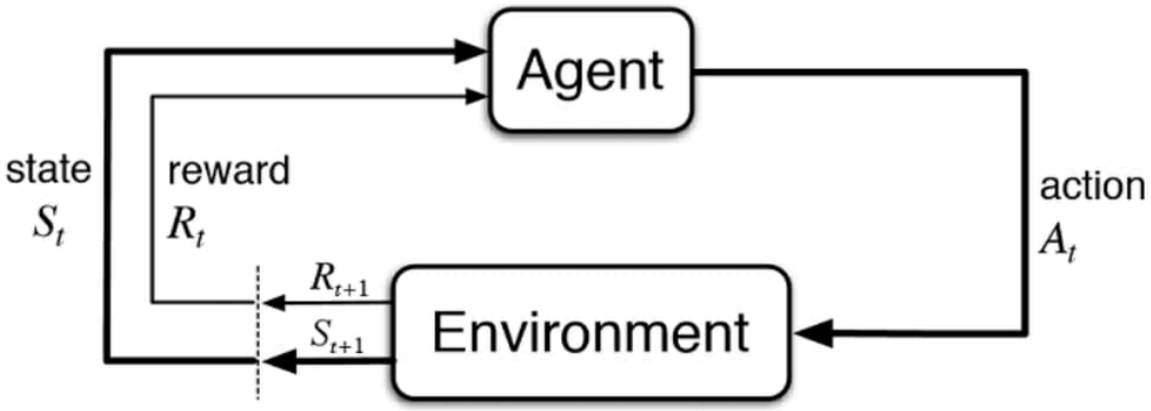

(22) **Retrieval-Augmented Generation**

https://medium.com/@yashwanths_29644/retrieval-augmented-generation-rag-01-introduction-to-rag-40da04999728

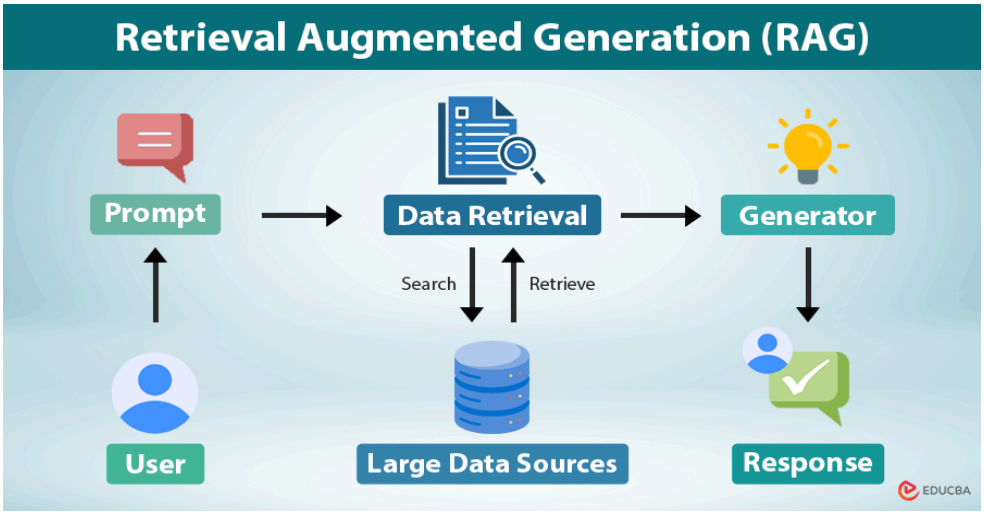

(23)**Self-Attention**

https://towardsdatascience.com/transformers-for-tabular-data-tabtransformer-deep-dive-5fb2438da820/

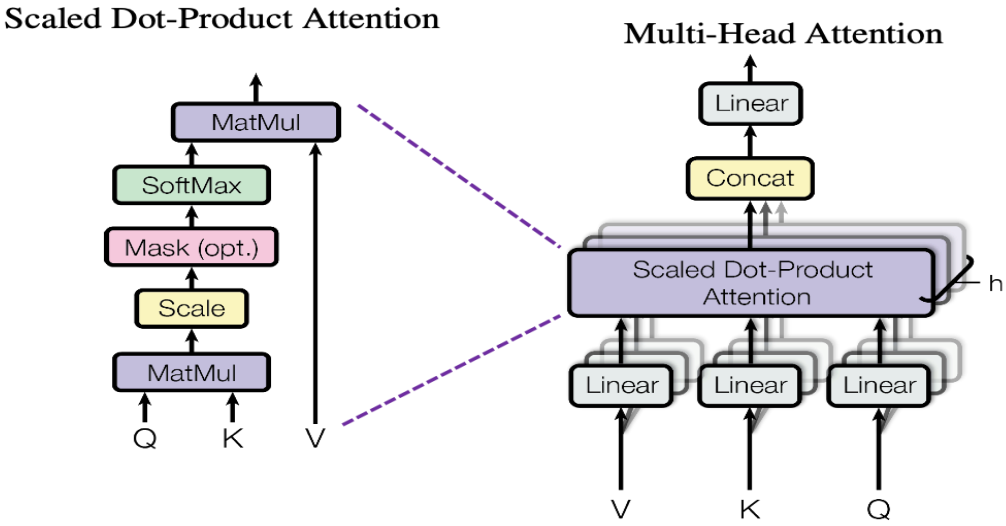

(24)

Transfer Learning

https://www.researchgate.net/figure/Transfer-Learning-architecture_fig2_381030953

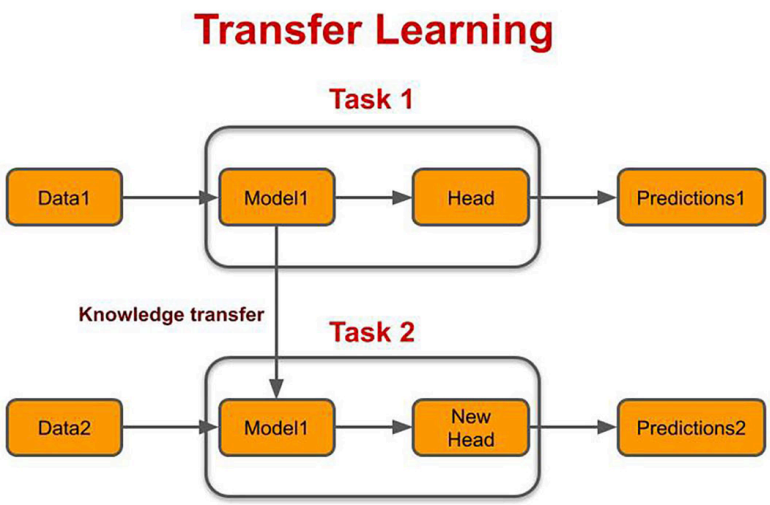

(25)

Variational Autoencoder

https://www.ibm.com/think/topics/latent-space

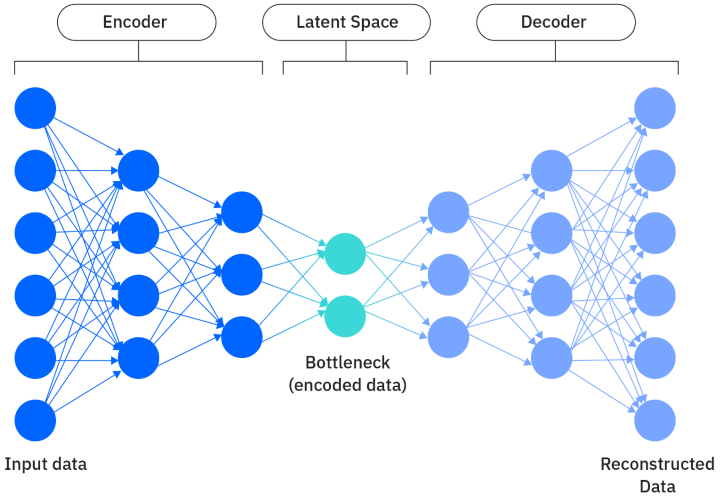

(26)

Zero-Shot Learning

https://www.computer.org/csdl/journal/tp/2023/04/09832795/1F6Q1JoJGne