AI Primer

B.M.Drury

AIMS of the Primer

- Discuss the terms that are commonly used in machine learning
- This is to clarify the terms / vocabulary that we will use in the following lecturers / tutorials
- Define key concepts such as models, algorithms, training, and prediction.
- Clarify types of learning: supervised, unsupervised, and reinforcement learning.
- Explain the roles of features, labels, and datasets (training, validation, testing).
- Describe common evaluation metrics like accuracy, precision, and recall.

AI != Machine Learning

- Artificial Intelligence is not just machine learning
- Artificial Intelligence (AI) is a broad field encompassing more than just machine learning
- Machine learning is a subfield of AI focused on learning from data
- There are other fields such as reasoning, knowledge representation and natural language processing
- However, currently when people speak about AI they mean machine learning

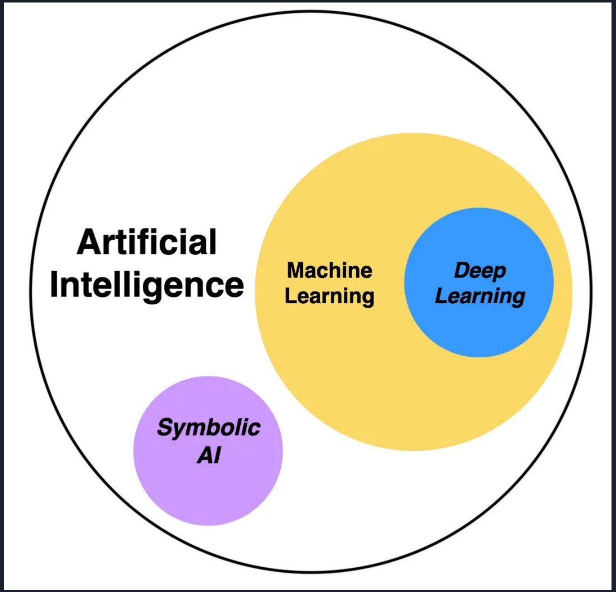

Machine Learning (Classical)

- Classical machine learning focuses on algorithms that learn patterns from data using predefined features.
- It includes well-established techniques like linear regression, decision trees, support vector machines (SVM), and k-nearest neighbours (KNN).
- These techniques are not computationally expensive and can run on consumer grade equipment.
- They were popular until the increase in popularity of deep learning

Neural Networks

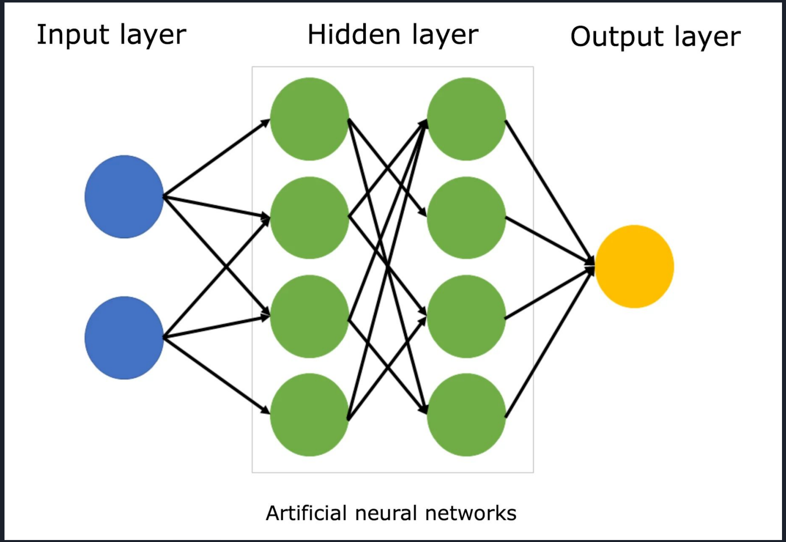

Inspired by the brain: Neural networks are computational models designed to simulate the way human brains process information.

Structure:

Input Layer: Receives the raw data (e.g., images, text, numbers).

Hidden Layers: Intermediate layers where computations and transformations occur.

Output Layer: Produces the final result or prediction.

Functionality:

Each node (or "neuron") in a layer performs a mathematical operation on the input it receives.

Connections between nodes have weights that adjust during training to minimize errors.

Learning Process:

Training: The network learns by adjusting weights based on the difference between its prediction and the actual result (error).

Backpropagation: A method used to update weights by propagating the error backward through the network.

Machine Learning (Deep Learning)

- Deep learning is a subset of machine learning that uses neural networks with many layers (hence "deep") to model complex patterns in data.
- It excels at automatically learning features from raw data, reducing the need for manual feature engineering.
- Common architectures include convolutional neural networks (CNNs) for images and recurrent neural networks (RNNs) or transformers for sequential data like text.
- Deep learning requires large amounts of data and computational power, often leveraging GPUs for training.
- It is the foundation of many recent AI advances, including image recognition, speech-to-text, and language models like ChatGPT.
- Despite high performance, deep learning models are often seen as "black boxes" due to limited interpretability.

Types of Machine Learning

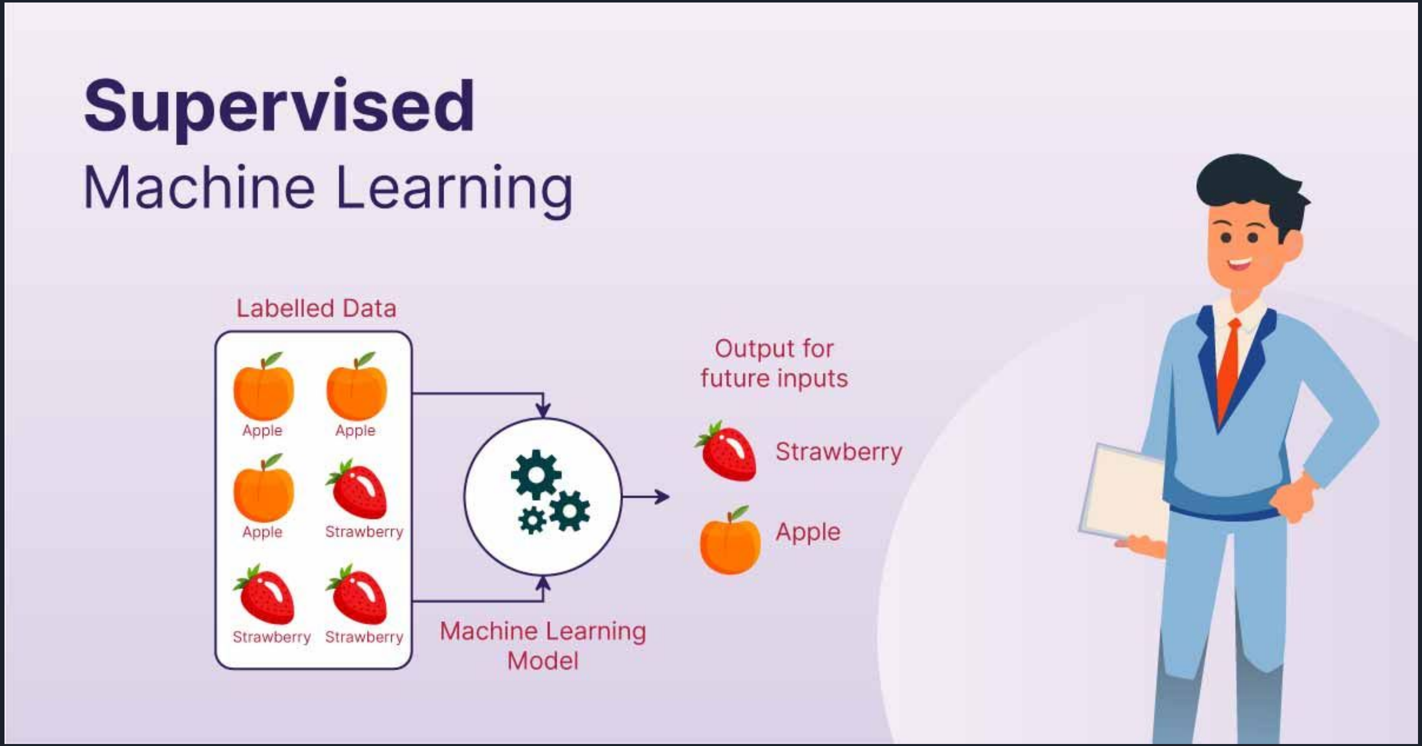

Unsupervised Learning

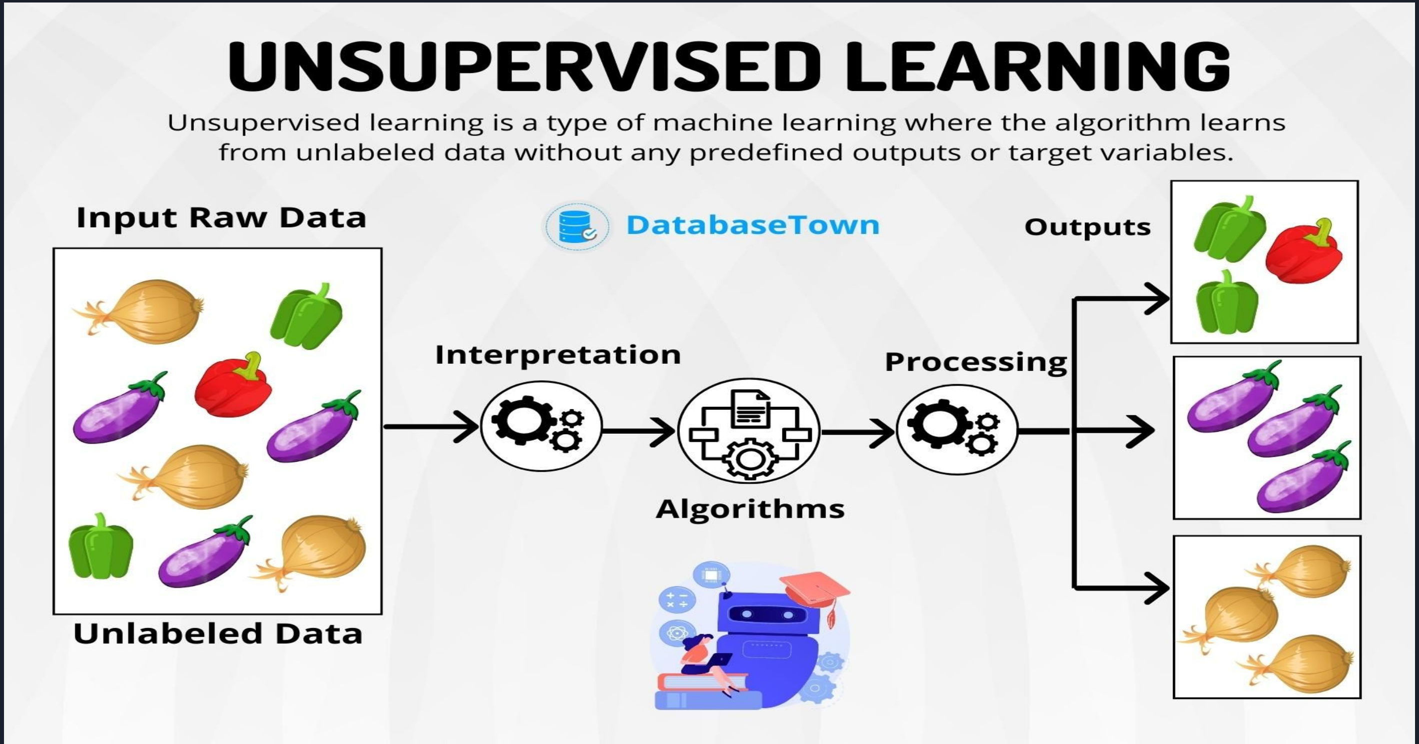

Machine Learning Pipeline

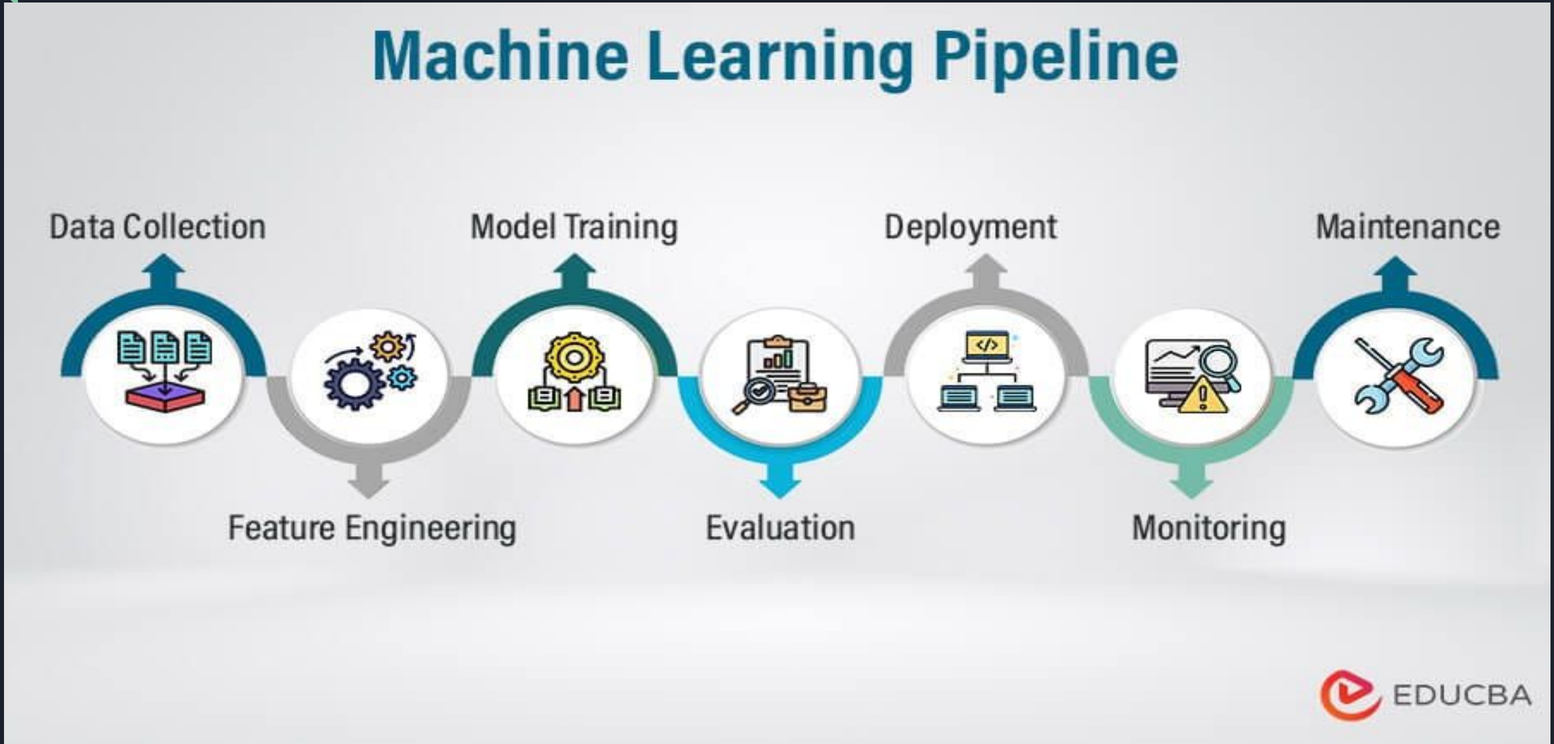

Features

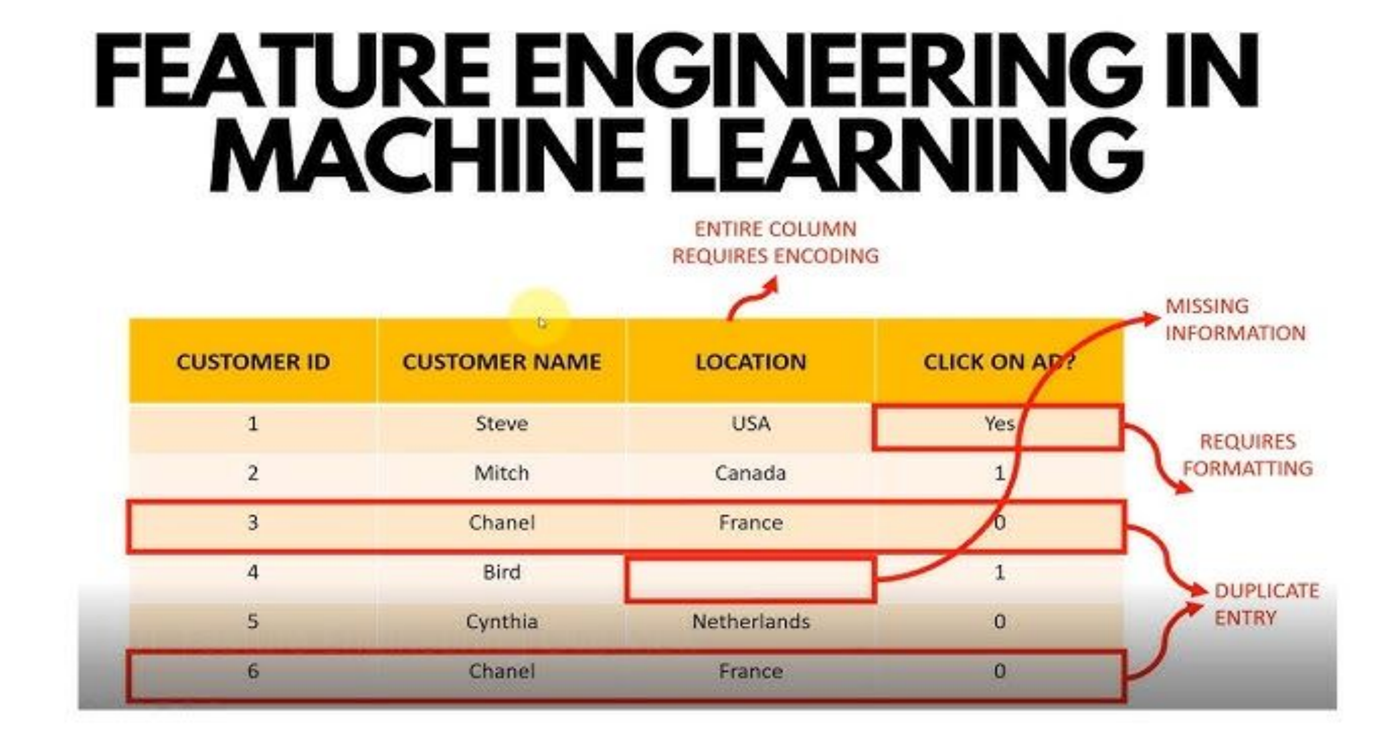

CUSTOMER ID	CUSTOMER NAME	LOCATION	CLICK ON AD?
1	Steve	USA	Yes
2	Mitch	Canada	1
3	Chanel	France	0
4	Bird		1
5	Cynthia	Netherlands	0
6	Chanel	France	0

Data Preprocessing

- Cleaning the data: Handle missing values, remove duplicates, and correct errors to ensure data quality.
- Encoding categorical variables: Convert non-numeric data into numerical format using methods like one-hot encoding or label encoding.
- Feature scaling: Normalize or standardize numerical features to bring them to a similar scale, which helps many algorithms perform better.
- Splitting the dataset: Divide data into training, validation, and test sets to evaluate model performance fairly.
- Handling imbalanced data: Apply techniques like resampling or use metrics beyond accuracy to deal with skewed class distributions.
- Dimensionality reduction: Use methods like PCA (Principal Component Analysis) to reduce the number of features while preserving important information.

This step will improve the results better than any other step in the pipeline

Evaluation

- We need to test how well the models perform
- Split the data in train, test and evaluation (for deep learning)
- Hold-out: split data in train and test (typically 80% 20%)
- Cross Fold Evaluation: Multiple splits and multiple runs and take an average
- Leave-One-Out Cross-Validation: Train on all but one example, test on the remaining one.

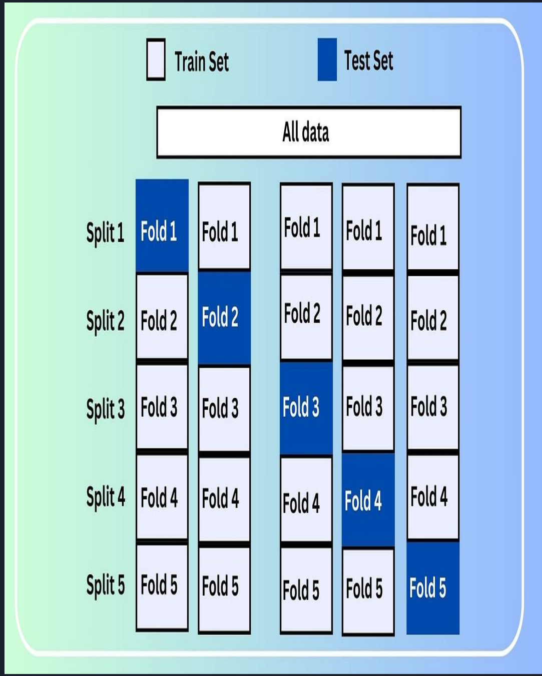

Evaluation

- **Accuracy**: The proportion of correct predictions (both true positives and true negatives) over all predictions.
- **Precision**: Measures how many predicted positives were actually positive.
- **Recall:** Measures how many actual positives were correctly identified.
- **F1 Score**:The harmonic mean of precision and recall. Useful when there's an imbalance between classes.
- **AUC-ROC (Area Under the ROC Curve):** Measures the model's ability to distinguish between classes across all classification thresholds.
- **Confusion Matrix:** A table showing TP, TN, FP, and FN to give a complete picture of performance.

Evaluation

Regression Models: Predict values

Mean Absolute Error (MAE): The average absolute difference between predicted and actual values.

Mean Squared Error (MSE): The average of the squared differences between predicted and actual values (penalizes larger errors more).

Root Mean Squared Error (RMSE): The square root of MSE, giving error in the same units as the target variable.

R^2 Score (Coefficient of Determination): Indicates how well the model explains the variability of the target. Ranges from 0 to 1.

Transfer Learning

- Repurposing of existing models
- The most common example is fine-tuning where we take an existing neural network, and train only the last layer
- Neural networks have numerous layers that represent more general features of the data
- The most specific layer (discriminator) is the last layer, so we only train the last layer
- Needs much less data than training the layer than the whole model

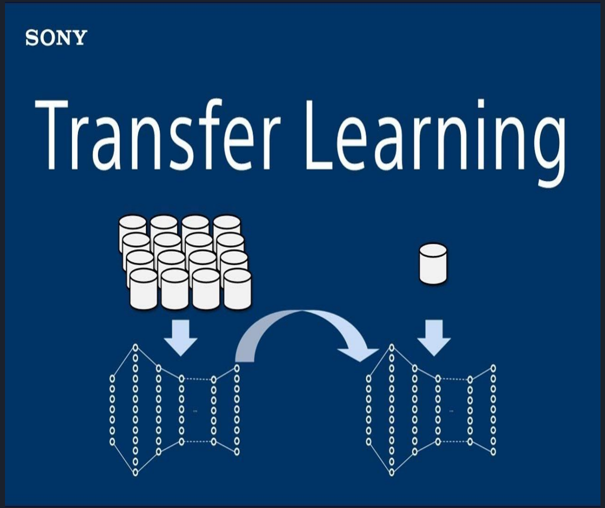

Large Language Models

Parameters:

Models have billions or trillions of parameters (weights) that are adjusted during training to minimize prediction error.

Tokenization:

Text is broken down into tokens (words, subwords, or characters) before being processed. Common tokenizers include Byte-Pair Encoding (BPE) and WordPiece.

Context Window:

LLMs have a limited context window (e.g., 2,000–128,000+ tokens), meaning they can only "pay attention" to a certain amount of text at once.

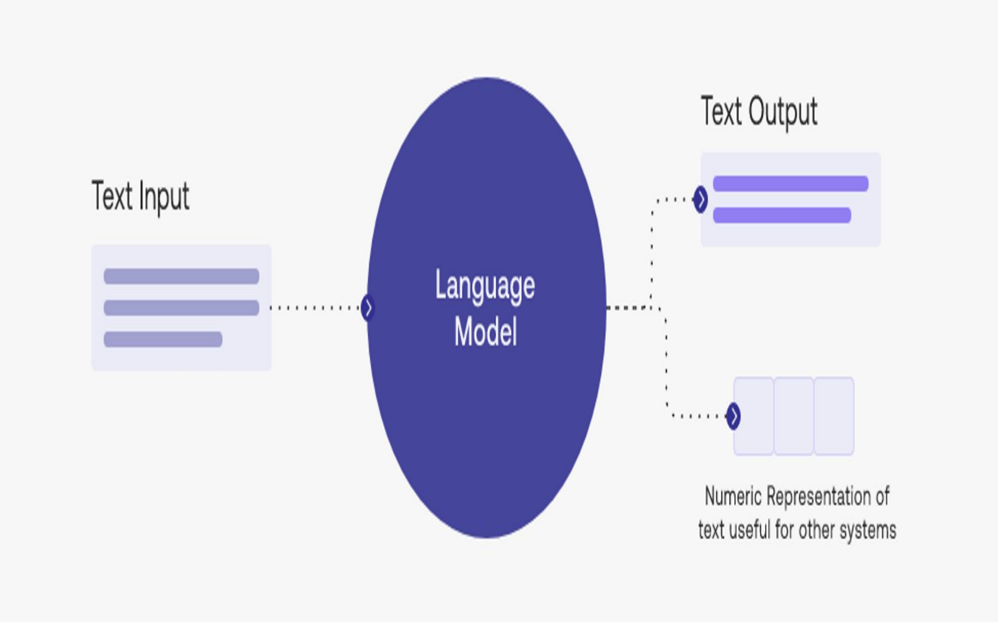

Large Language Models

Masked training is used mainly in models like BERT.

Goal: Instead of predicting the next word (like GPT), the model learns to predict missing words inside a sentence.

How it works:

Take an input sentence.

Example: "The cat sat on the mat."

Randomly mask some words (replace them with a special token like [MASK]).

Example: "The cat sat on the [MASK]."

The model is trained to guess the missing word ("mat").

Why mask words?

It forces the model to learn deep contextual understanding from both the left and right sides of a word (bidirectional context).

It improves performance on tasks like question answering, classification, etc

What is a Loss Function?

A loss function measures how wrong the model's predictions are compared to the correct answers.

The goal during training is to minimize this loss — the smaller the loss, the better the model is doing

Conclusion

- Brief Primer on Machine Learning
- A subfield of artificial intelligence (AI).
- Focuses on creating algorithms that allow computers to learn from data without being explicitly programmed.
- The system improves its performance as it sees more examples.
- Field much larger than what is in this presentation