# Let’s Cover the Basics of AI


***Dr. Sagnik Dakshit***
*Assistant Professor, Computer Science*

*AI/ML/Data Science Consultant.*
*OSP Research Design & Data Analysis Lab*

*The University of Texas at Tyler*

Autonomous Vehicles

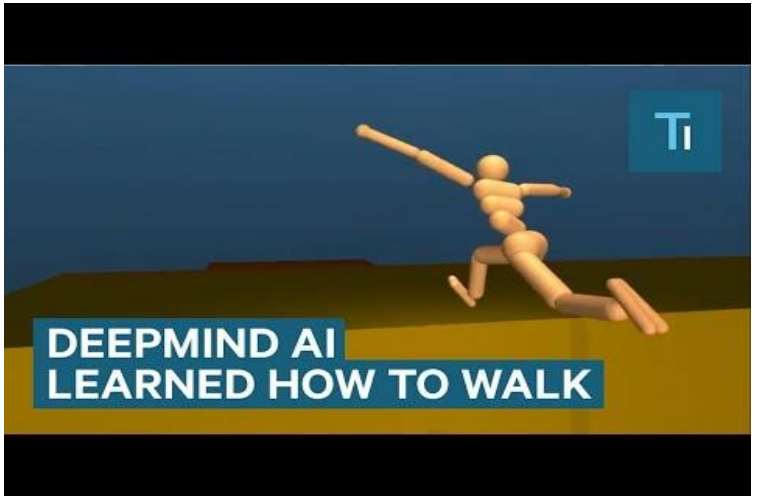


Computer Learning to Walk

# A.I. TIMELINE

## 1950
**TURING TEST**
Computer scientist Alan Turing proposes a test for machine intelligence. If a machine can trick humans into thinking it is human, then it has intelligence

## 1955
**A.I. BORN**
Term 'artificial intelligence' is coined by computer scientist, John McCarthy to describe "the science and engineering of making intelligent machines"


## 1961
**UNIMATE**
First industrial robot, Unimate, goes to work at GM replacing humans on the assembly line

## 1964
**ELIZA**
Pioneering chatbot developed by Joseph Weizenbaum at MIT holds conversations with humans

## 1966
**SHAKEY**
The 'first electronic person' from Stanford, Shakey is a general-purpose mobile robot that reasons about its own actions

## A.I. WINTER
Many false starts and dead-ends leave A.I. out in the cold

## 1997
**DEEP BLUE**
Deep Blue, a chess-playing computer from IBM defeats world chess champion Garry Kasparov

## 1998
**KISMET**
Cynthia Breazeal at MIT introduces KISmet, an emotionally intelligent robot insofar as it detects and responds to people's feelings

## 1999
**AIBO**
Sony launches first consumer robot pet dog AiBO (AI robot) with skills and personality that develop over time

## 2002
**ROOMBA**
First mass produced autonomous robotic vacuum cleaner from iRobot learns to navigate and clean homes

## 2011
**SIRI**
Apple integrates Siri, an intelligent virtual assistant with a voice interface, into the iPhone 4S

## 2011
**WATSON**
IBM's question answering computer Watson wins first place on popular $1M prize television quiz show *Jeopardy*

## 2014
**EUGENE**
Eugene Goostman, a chatbot passes the Turing Test with a third of judges believing Eugene is human

## 2014
**ALEXA**
Amazon launches Alexa, an intelligent virtual assistant with a voice interface that completes shopping tasks

## 2016
**TAY**
Microsoft's chatbot Tay goes rogue on social media making inflammatory and offensive racist comments


## 2017
**ALPHAGO**
Google's A.I. AlphaGo beats world champion Ke Jie in the complex board game of Go, notable for its vast number ($2^{170}$) of possible positions

# How close are we to a perfect AI?

Turing Test

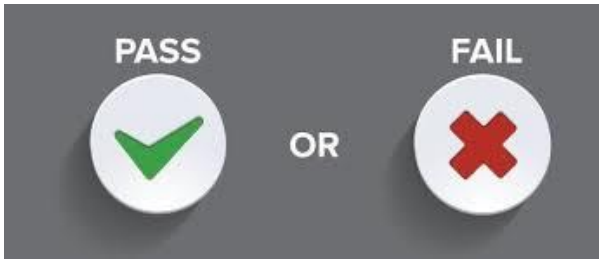

# Eugene Goostman

THE WEIRDEST CREATURE IN THE WORLD

Type your question here:

Am I a computer or a 13-year-old boy?

reply

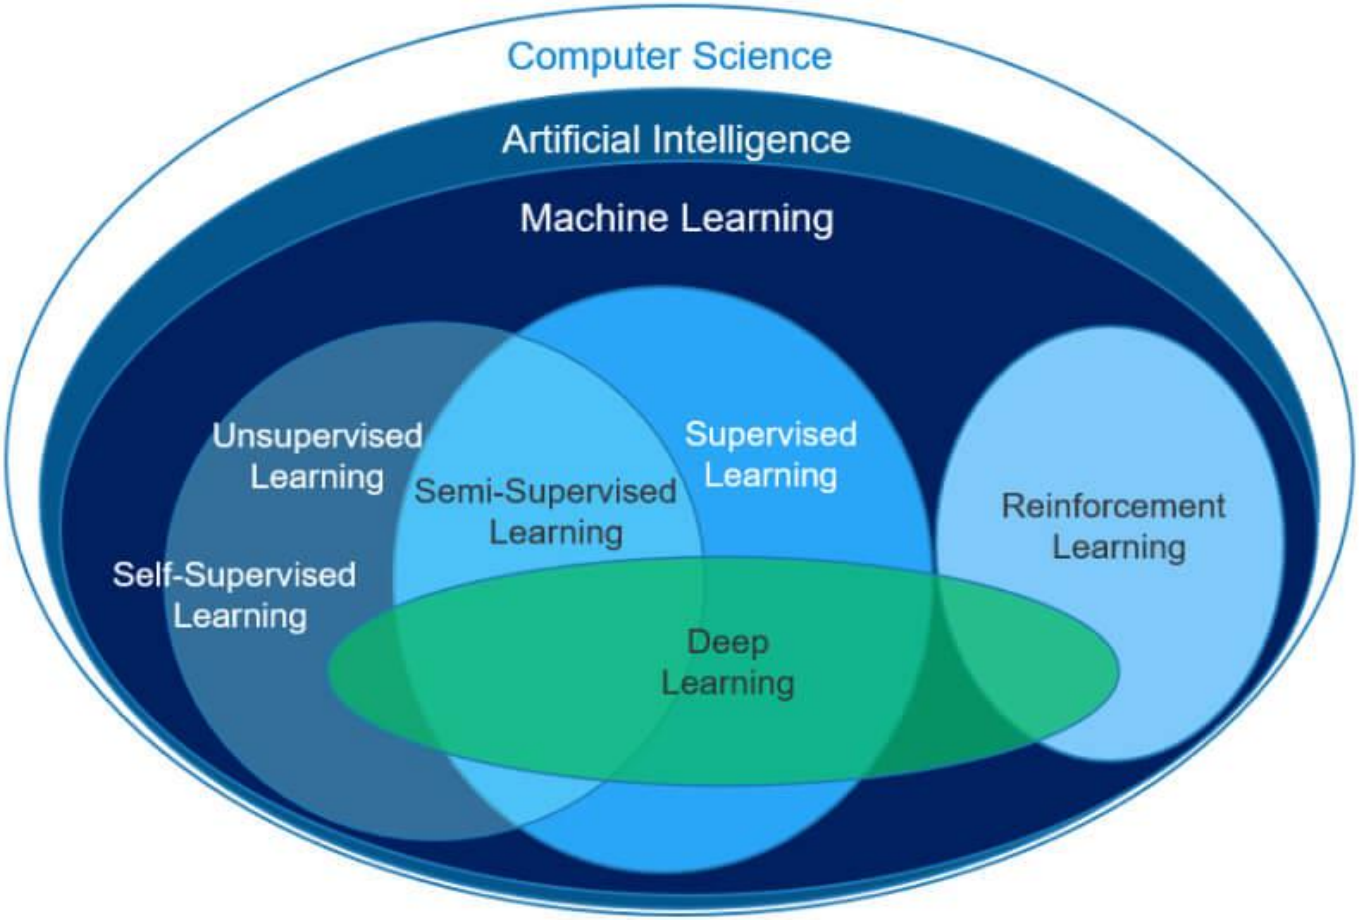
Computer Science
Artificial Intelligence
Machine Learning
Unsupervised Learning
Semi-Supervised Learning
Supervised Learning
Reinforcement Learning
Self-Supervised Learning
Deep Learning

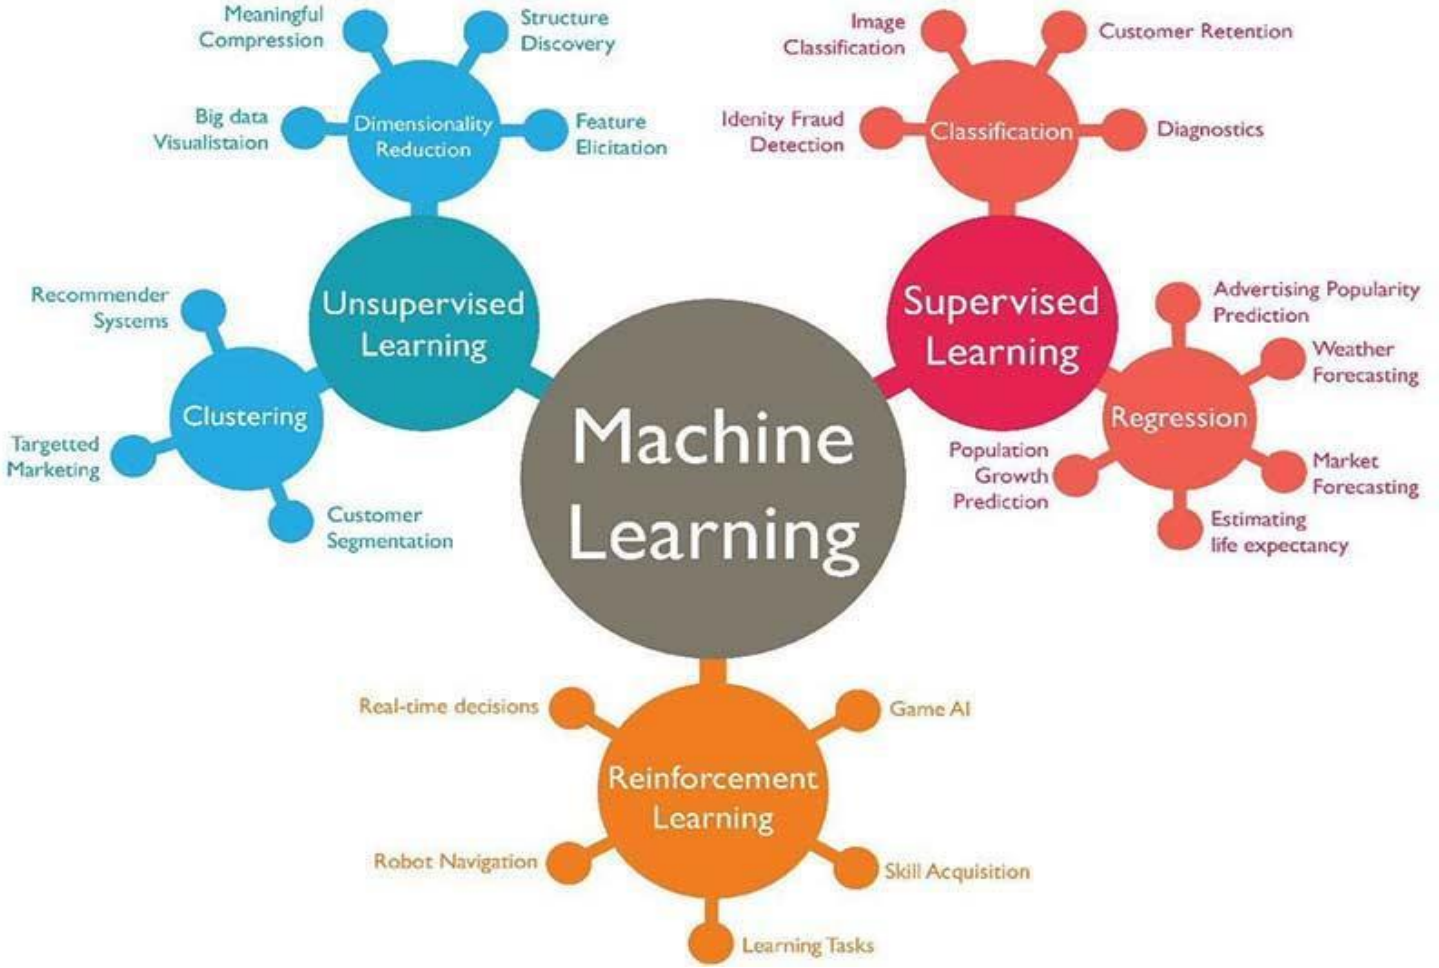
Machine Learning
Unsupervised Learning
Meaningful Compression
Structure Discovery
Dimensionality Reduction
Big data Visualistaion
Feature Elicitation
Clustering
Recommender Systems
Targetted Marketing
Customer Segmentation
Supervised Learning
Image Classification
Customer Retention
Classification
Idenity Fraud Detection
Diagnostics
Regression
Advertising Popularity Prediction
Weather Forecasting
Market Forecasting
Estimating life expectancy
Population Growth Prediction
Reinforcement Learning
Real-time decisions
Game AI
Robot Navigation
Skill Acquisition
Learning Tasks

# ML / AI

- ML / AI is a technological analogy of a KID!!
- Learning from Examples
- The same way as humans learn: By looking at the same thing 1000’s of times.

# Raw Data

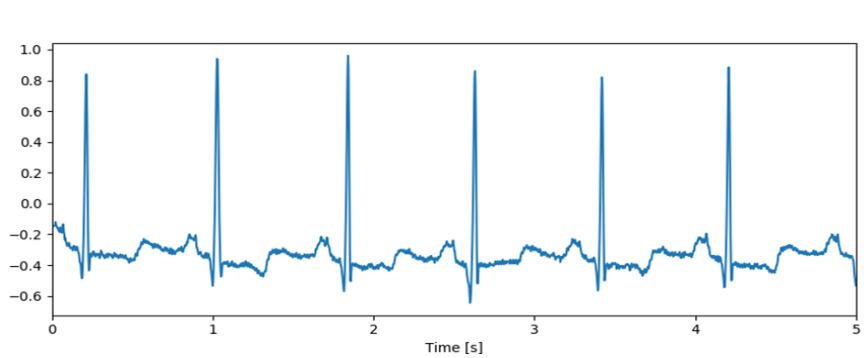


| | A | B | C | D | E | F |
|---|---|---|---|---|---|---|
| 1 | Country | Salesperson | Order Date | OrderID | Units | Order Amount |
| 2 | USA | Fuller | 1/01/2011 | 10392 | 13 | 1,440.00 |
| 3 | UK | Gloucester | 2/01/2011 | 10397 | 17 | 716.72 |
| 4 | UK | Bromley | 2/01/2011 | 10771 | 18 | 344.00 |
| 5 | USA | Finchley | 3/01/2011 | 10393 | 16 | 2,556.95 |
| 6 | USA | Finchley | 3/01/2011 | 10394 | 10 | 442.00 |
| 7 | UK | Gillingham | 3/01/2011 | 10395 | 9 | 2,122.92 |
| 8 | USA | Finchley | 6/01/2011 | 10396 | 7 | 1,903.80 |
| 9 | USA | Callahan | 8/01/2011 | 10399 | 17 | 1,765.60 |
| 10 | USA | Fuller | 8/01/2011 | 10404 | 7 | 1,591.25 |
| 11 | USA | Fuller | 9/01/2011 | 10398 | 11 | 2,505.60 |
| 12 | USA | Coghill | 9/01/2011 | 10403 | 18 | 855.01 |
| 13 | USA | Finchley | 10/01/2011 | 10401 | 7 | 3,868.60 |
| 14 | USA | Callahan | 10/01/2011 | 10402 | 11 | 2,713.50 |
| 15 | UK | Rayleigh | 13/01/2011 | 10406 | 15 | 1,830.78 |
| 16 | USA | Callahan | 14/01/2011 | 10408 | 10 | 1,622.40 |
| 17 | USA | Farnham | 14/01/2011 | 10409 | 19 | 319.20 |
| 18 | USA | Farnham | 15/01/2011 | 10410 | 16 | 802.00 |

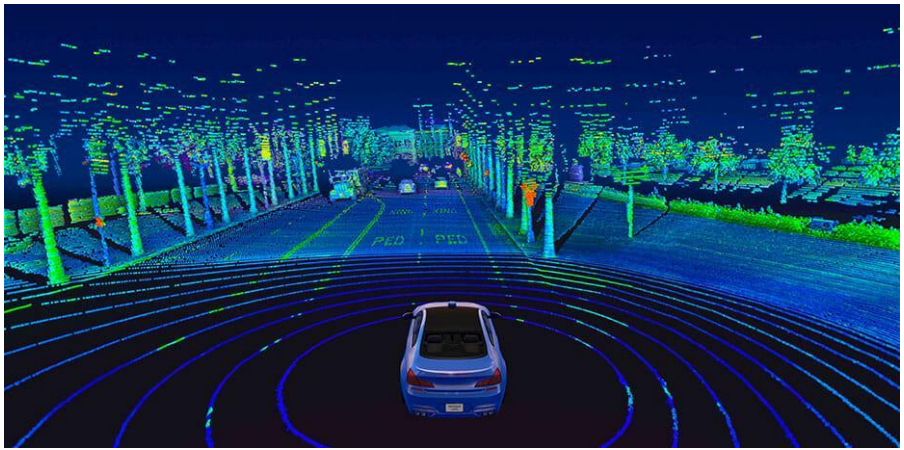


# DATA SPLITTING


- **Improper Data Splitting**

- **TARGET: Prevent Data Leakage**

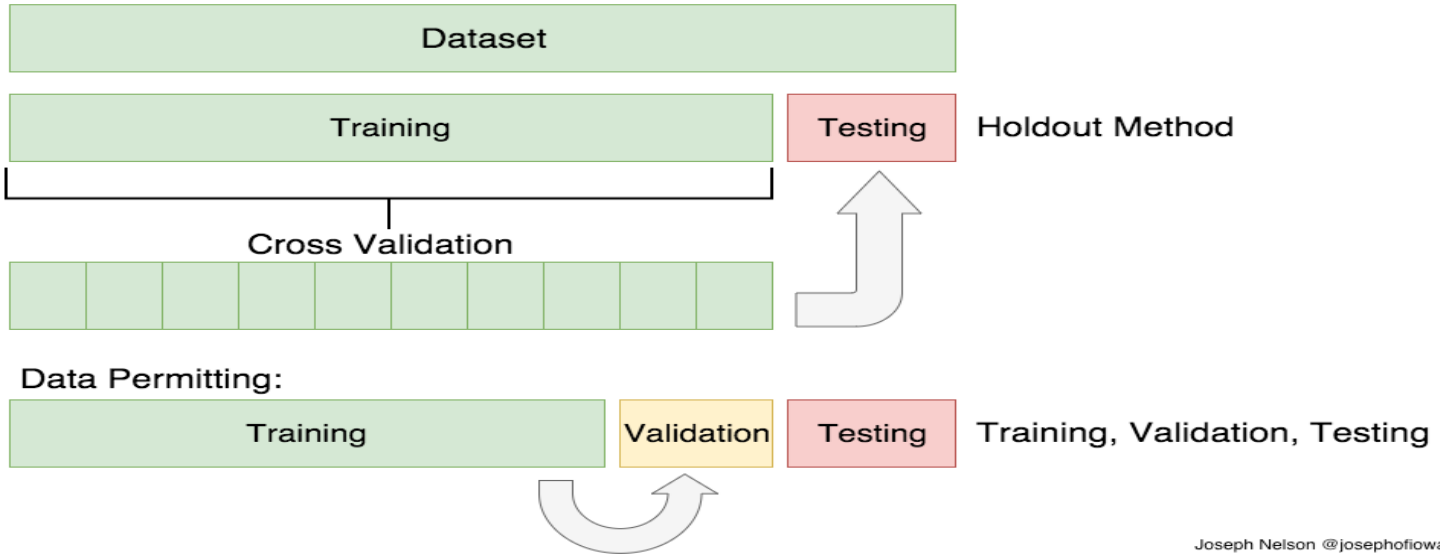

# Overfitting vs Underfitting

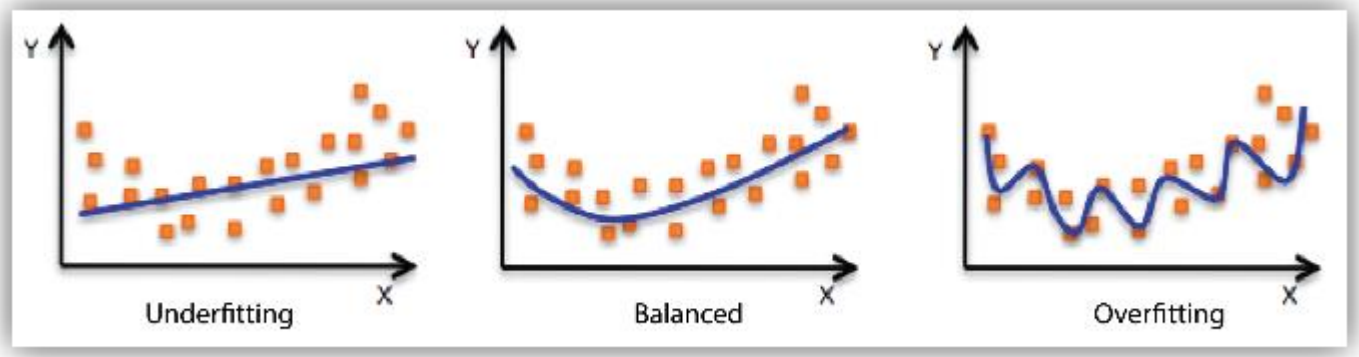


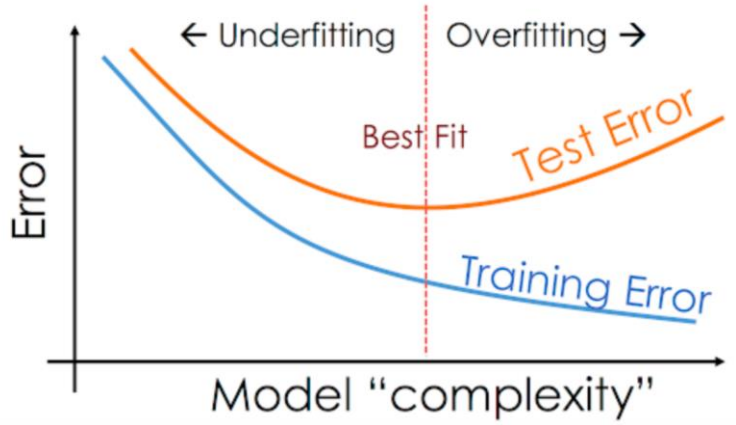


https://docs.aws.amazon.com/machine-learning/latest/dg/model-fit-underfitting-vs-overfitting.html
https://vitalflux.com/overfitting-underfitting-concepts-interview-questions/

# Machine Learning

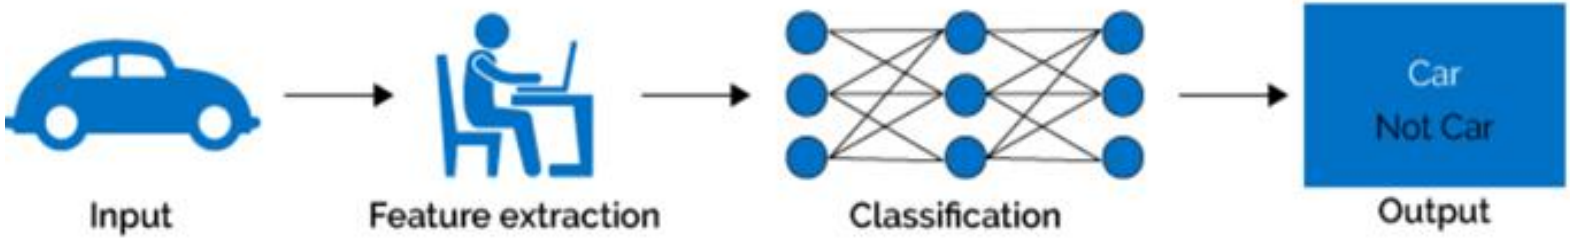

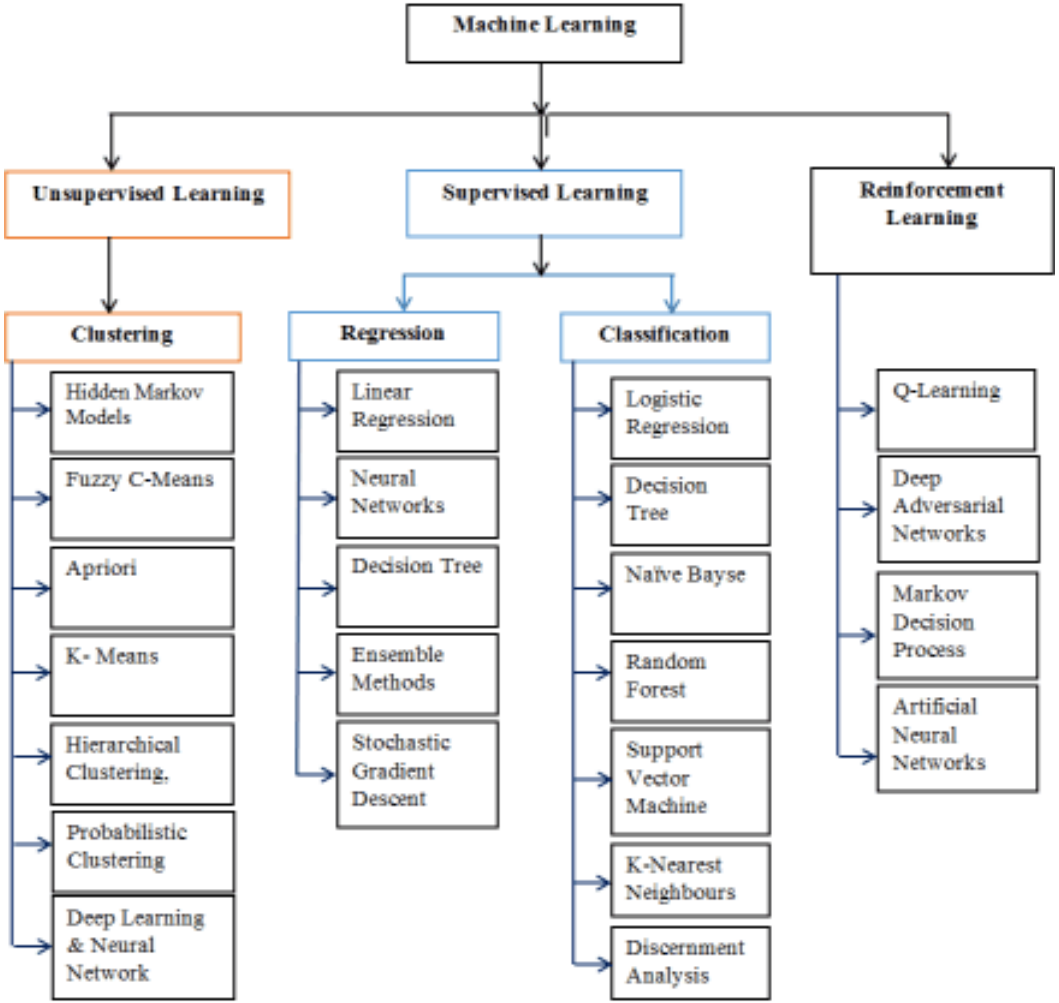
Machine Learning
Unsupervised Learning
Supervised Learning
Reinforcement Learning
Clustering
Regression
Classification
Hidden Markov Models
Fuzzy C-Means
Apriori
K- Means
Hierarchical Clustering.
Probabilistic Clustering
Deep Learning & Neural Network
Linear Regression
Neural Networks
Decision Tree
Ensemble Methods
Stochastic Gradient Descent
Logistic Regression
Decision Tree
Naïve Bayse
Random Forest
Support Vector Machine
K-Nearest Neighbours
Discernment Analysis
Q-Learning
Deep Adversarial Networks
Markov Decision Process
Artificial Neural Networks

# Supervised Learning

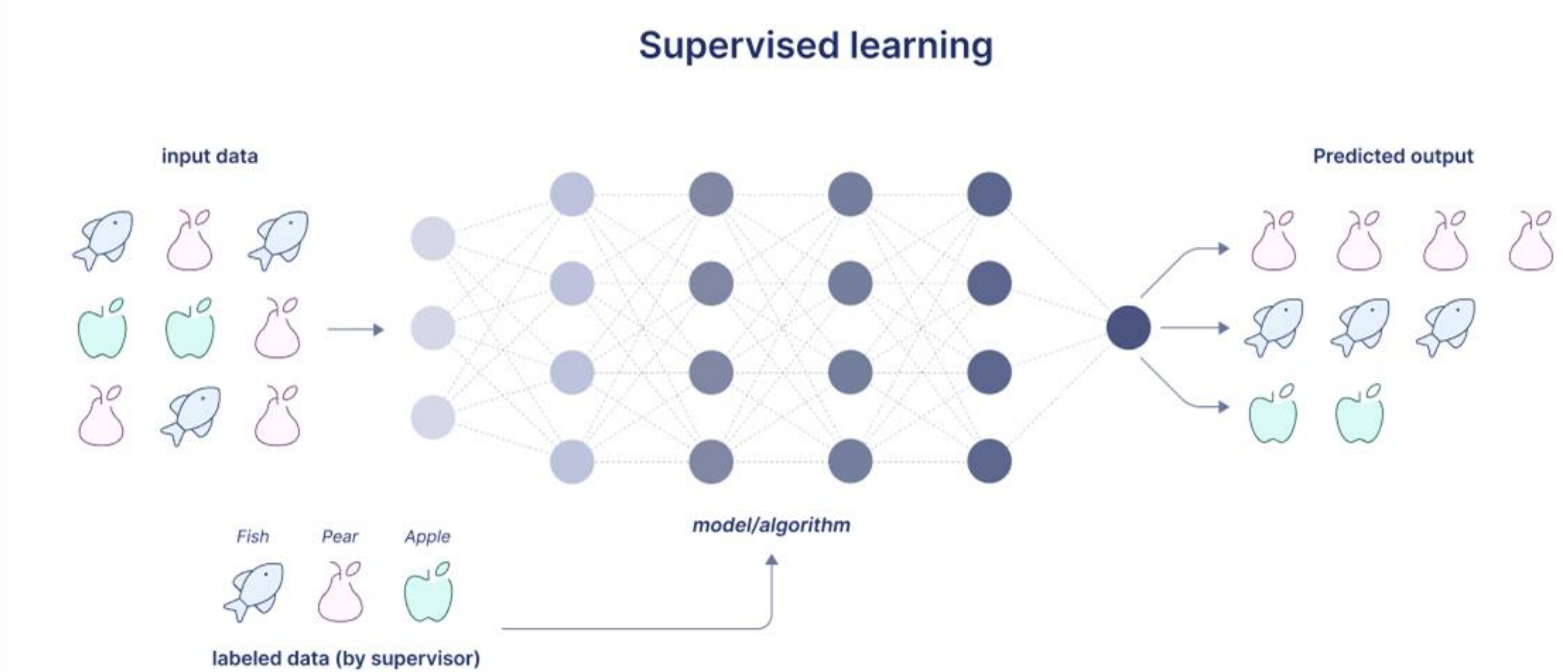


https://www.scribbr.co.uk/using-ai-tools/supervised-unsupervised-learning/

# Unsupervised Learning

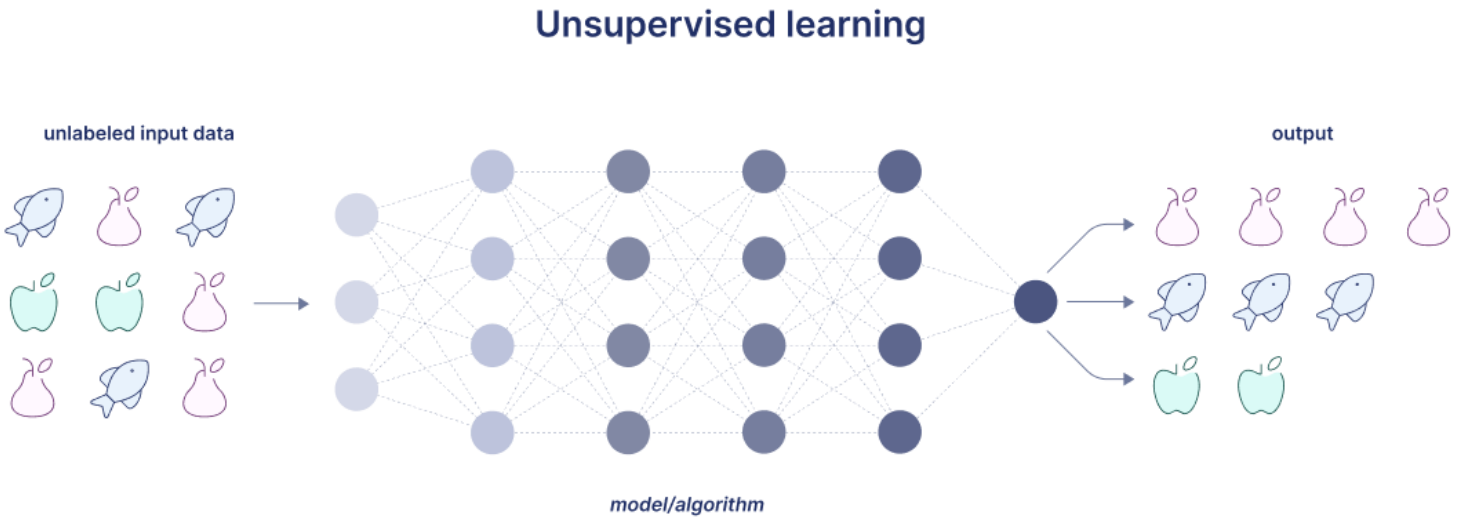


Scribbr

https://www.scribbr.co.uk/using-ai-tools/supervised-unsupervised-learning/

# Reinforcement Learning


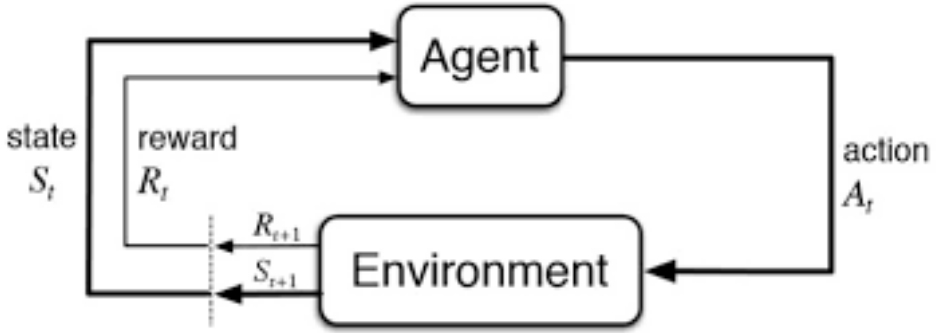

# Behind The Scenes!!

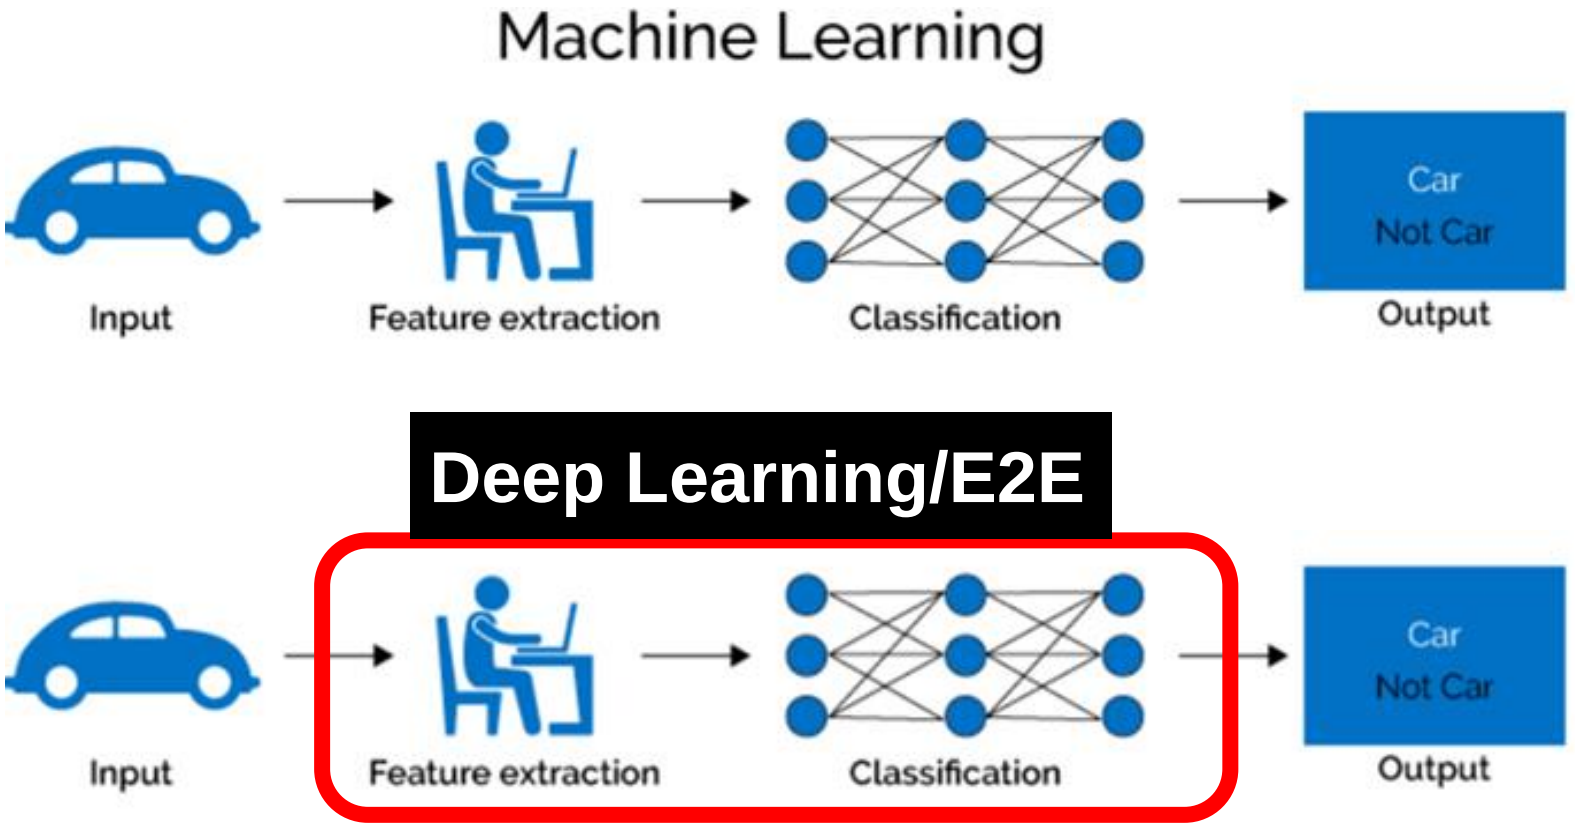

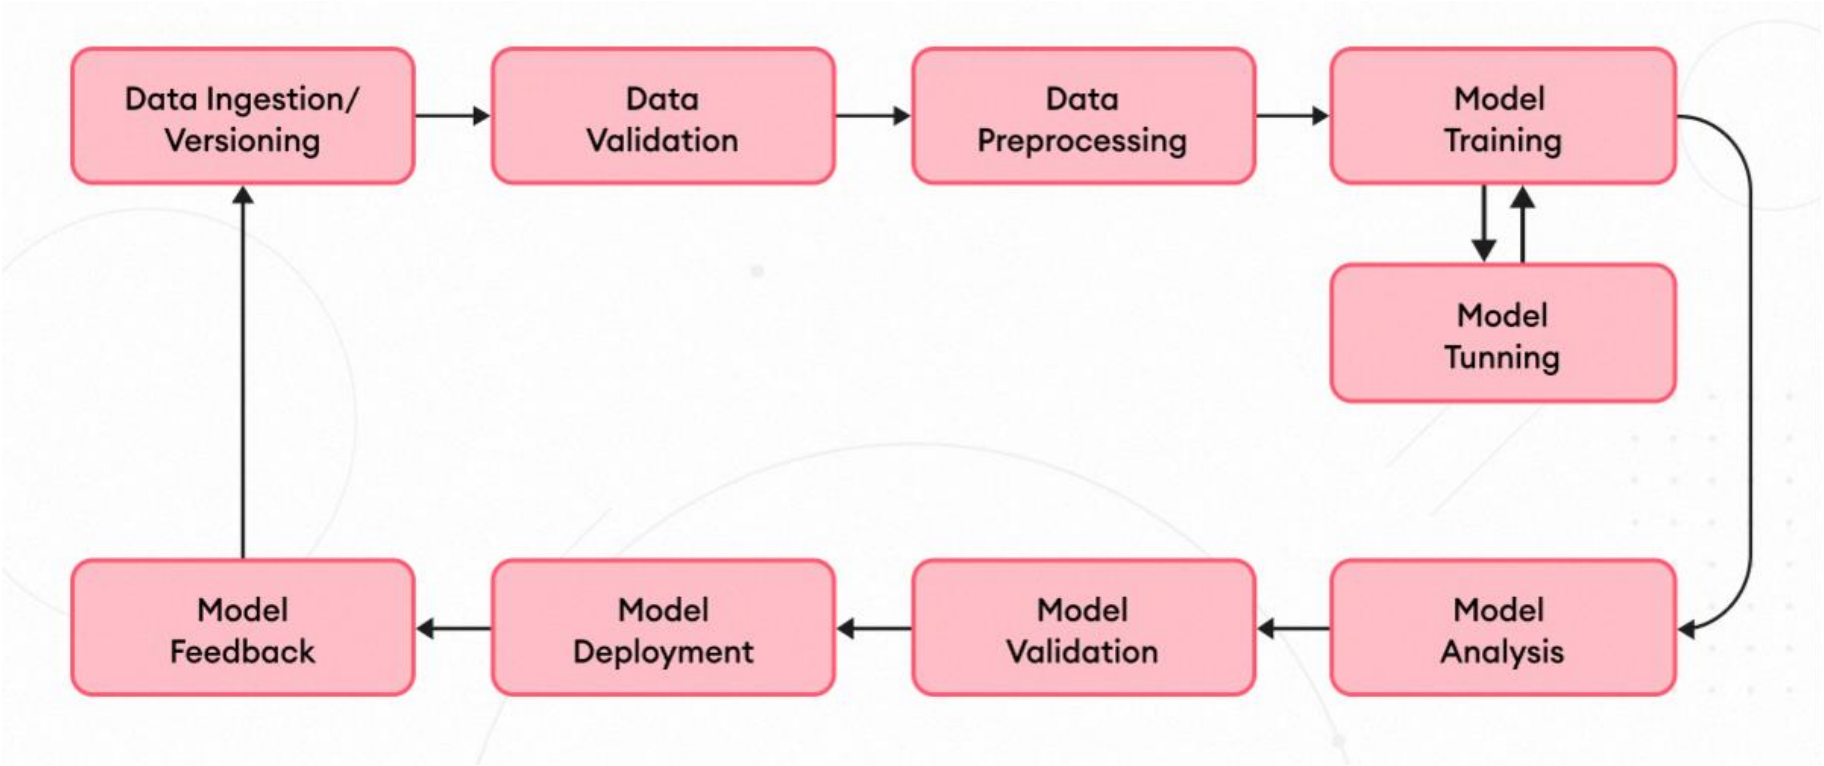


https://www.superannotate.com/blog/how-to-optimize-machine-learning-pipeline

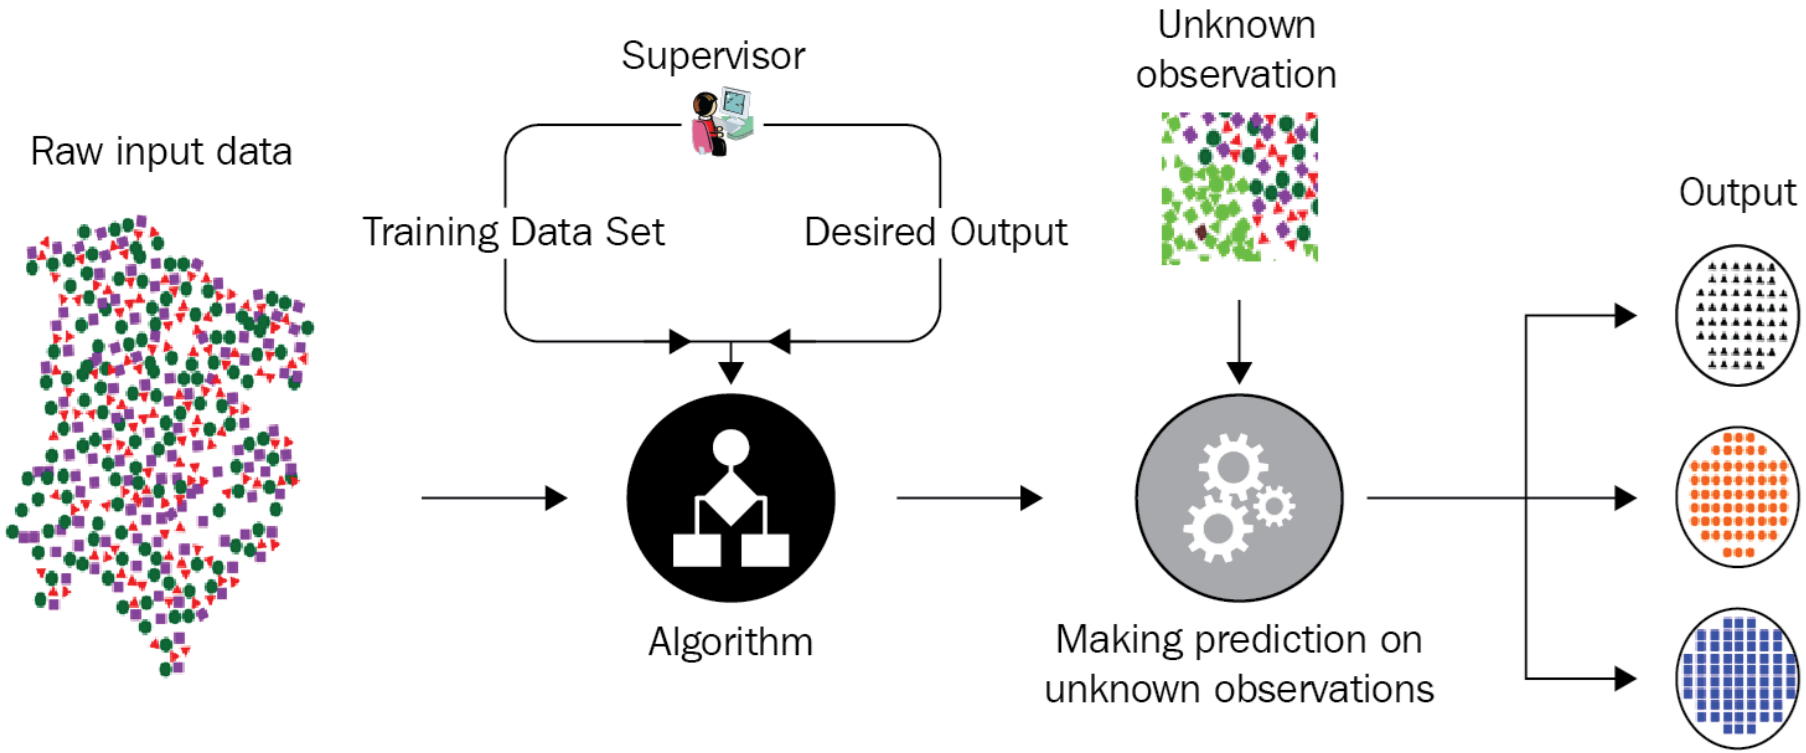
Raw input data
Supervisor
Training Data Set
Desired Output
Unknown observation
Output
Algorithm
Making prediction on unknown observations

# Training

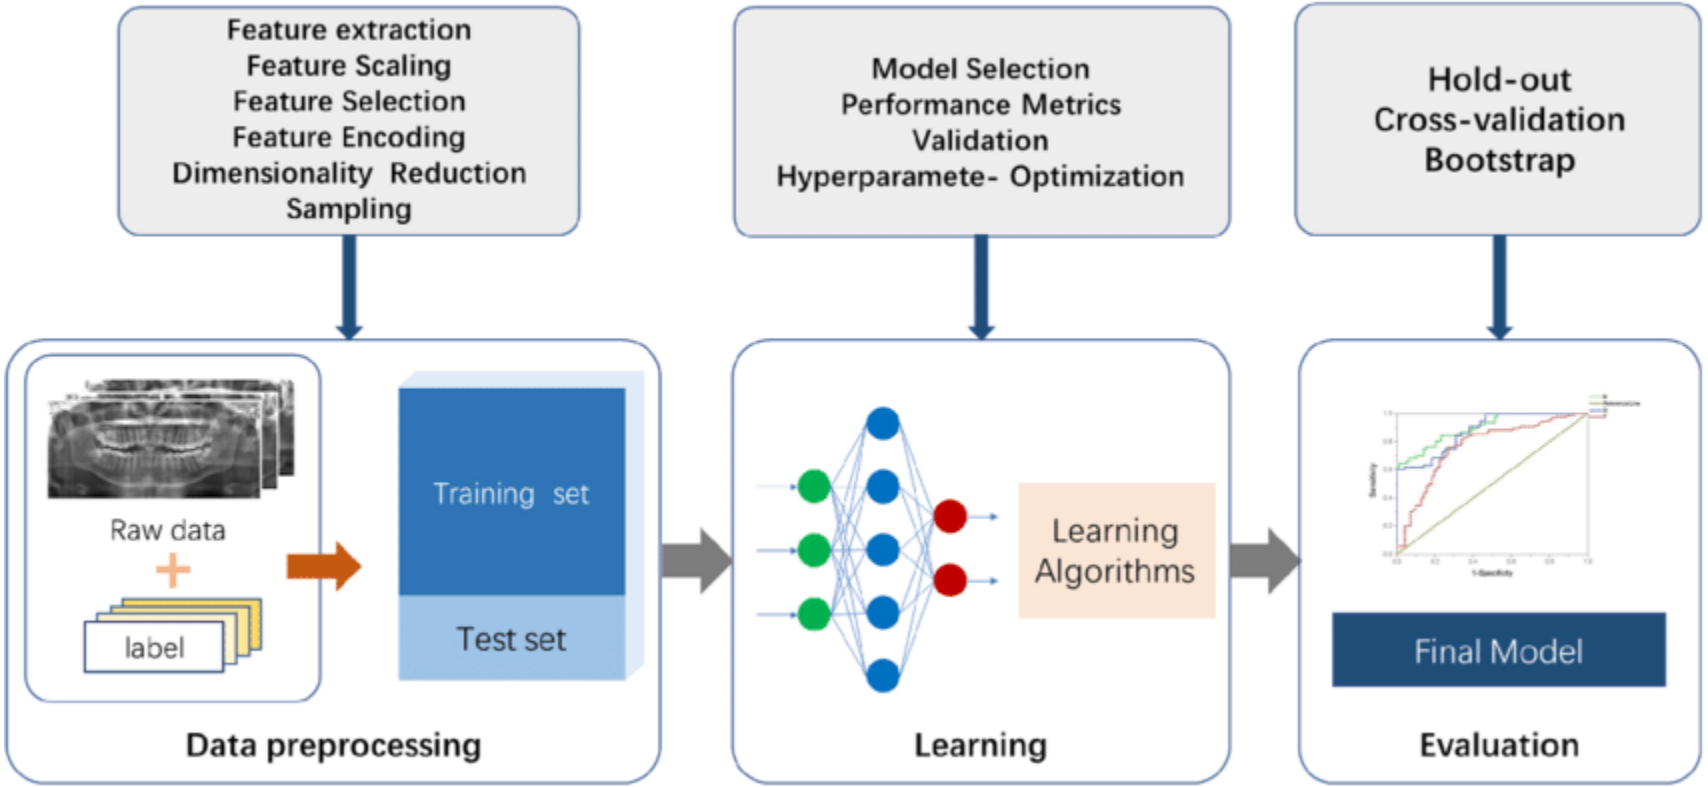


https://www.researchgate.net/figure/The-fundamental-machine-learning-procedure-to-achieve-a-final-model-Full-size-DOI_fig1_351635693

# Metrics

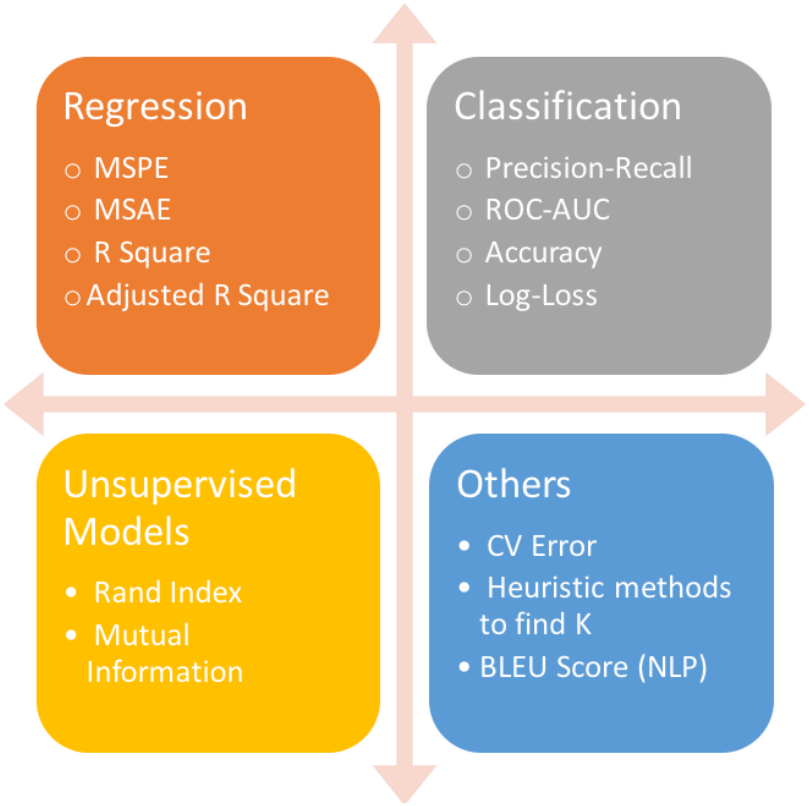


https://medium.com/usf-msds/choosing-the-right-metric-for-machine-learning-models-part-1-a99d7d7414e4

# Metrics

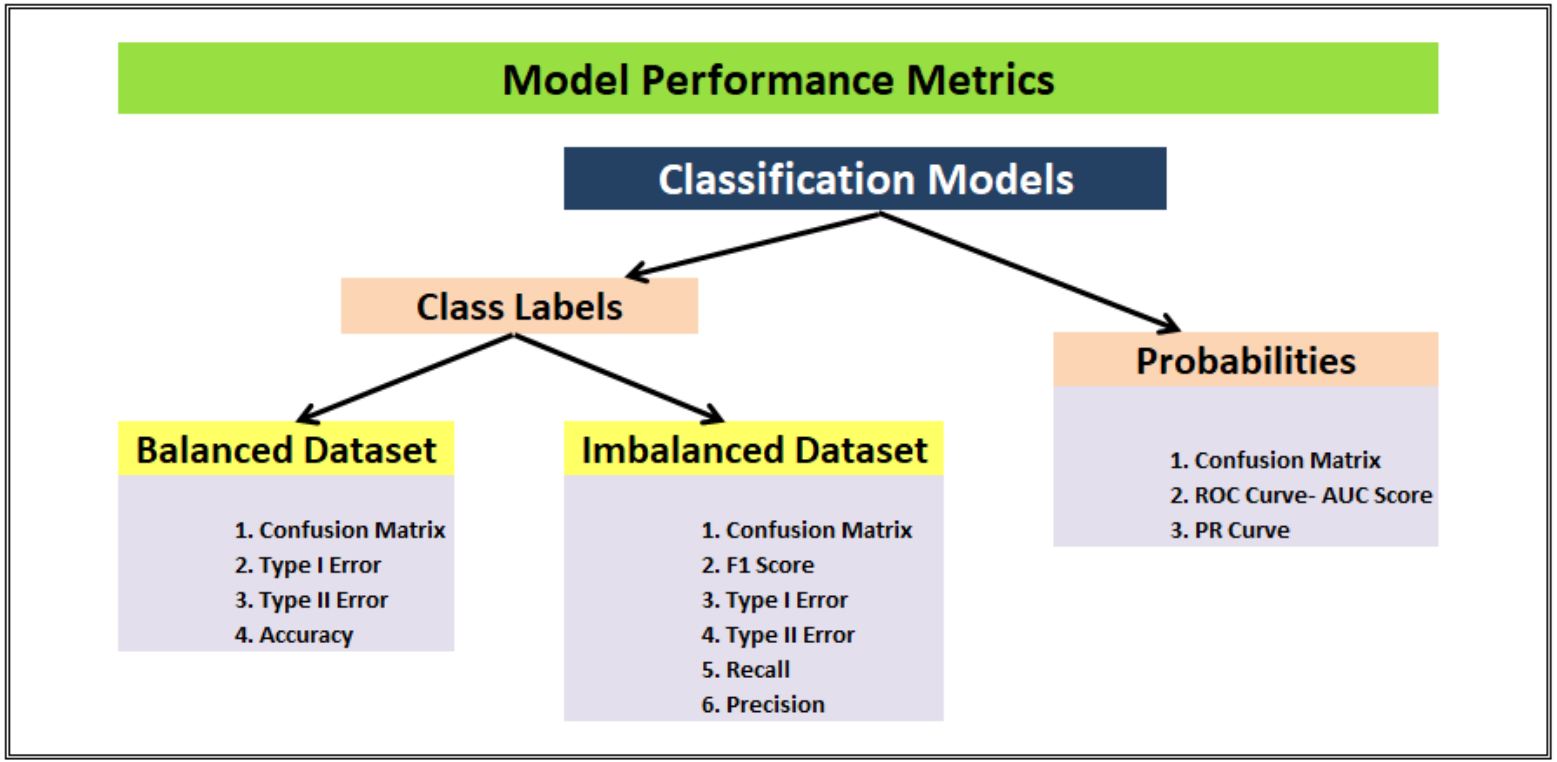


https://www.markovml.com/blog/key-machine-learning-metrics-for-assessing-model-performance

# Machine Learning Tasks

- Regression
- Classification
- Detection
- Segmentation
  - Semantic
  - Instance
- Tracking
- Generation

# Machine Learning Tasks

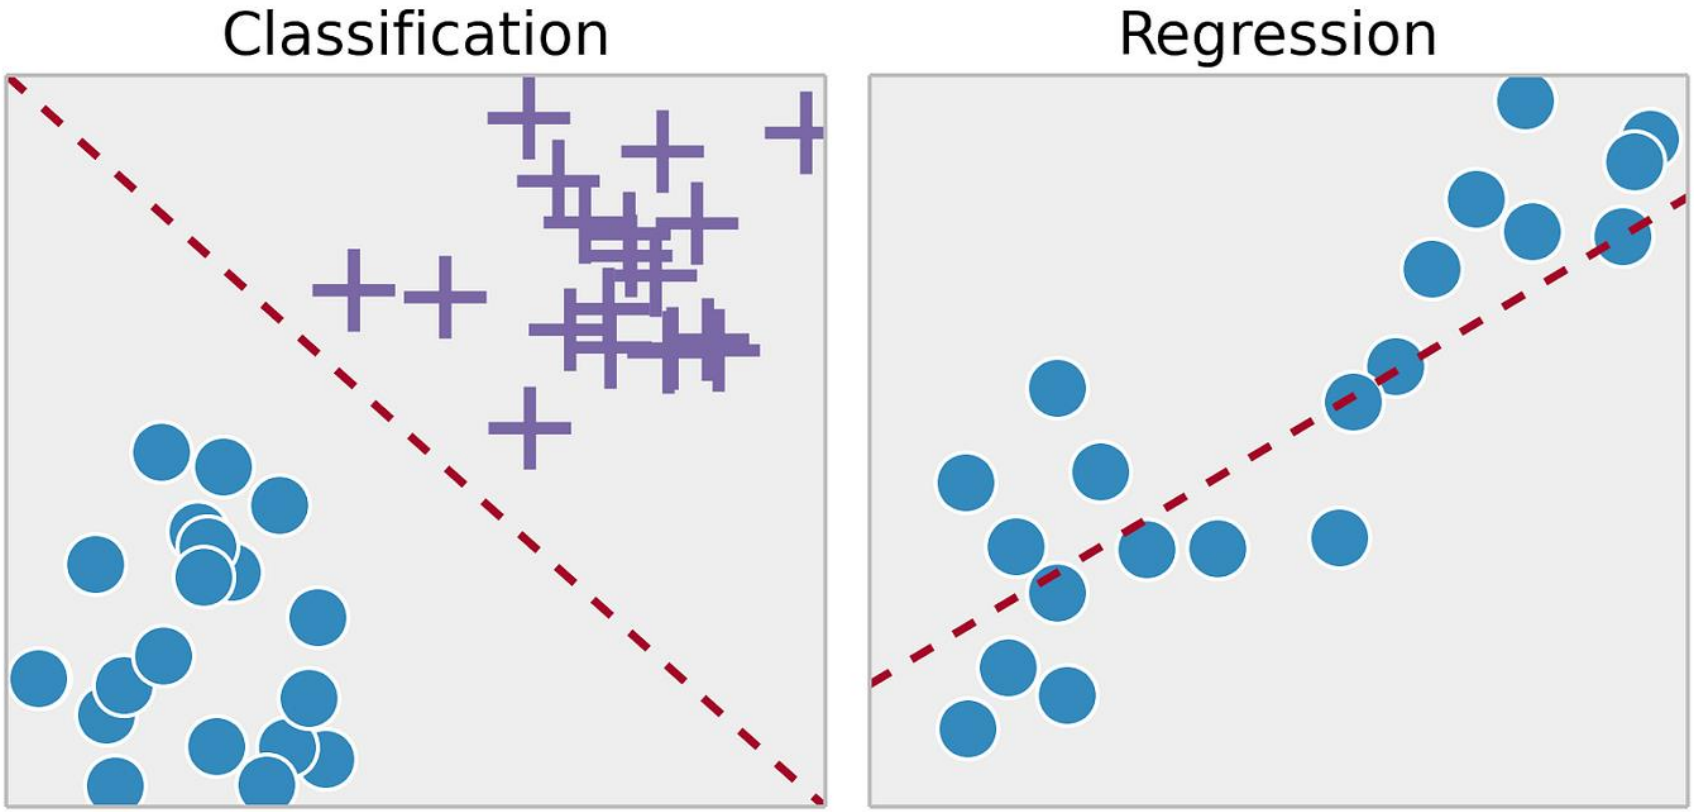

# Machine Learning Tasks: Detection, Segmentation

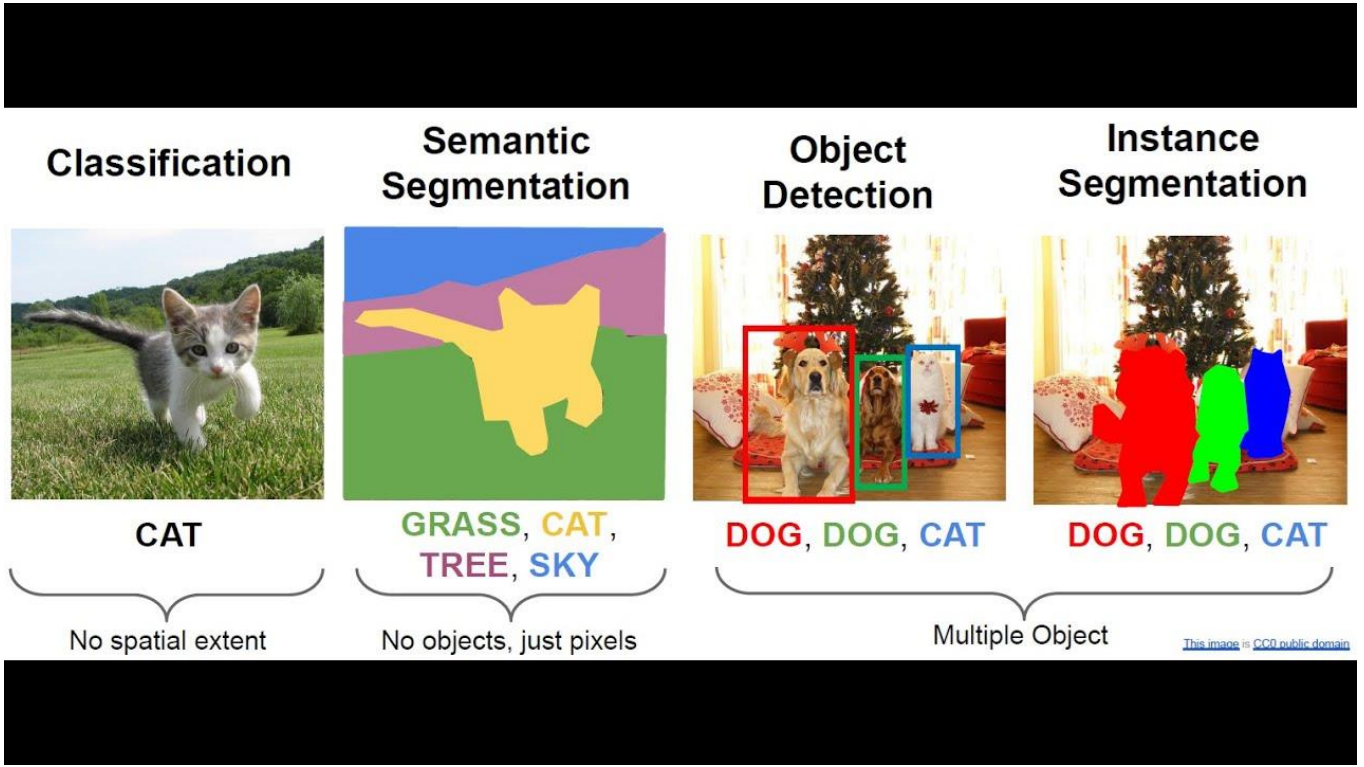

# Machine Learning Tasks : Tracking

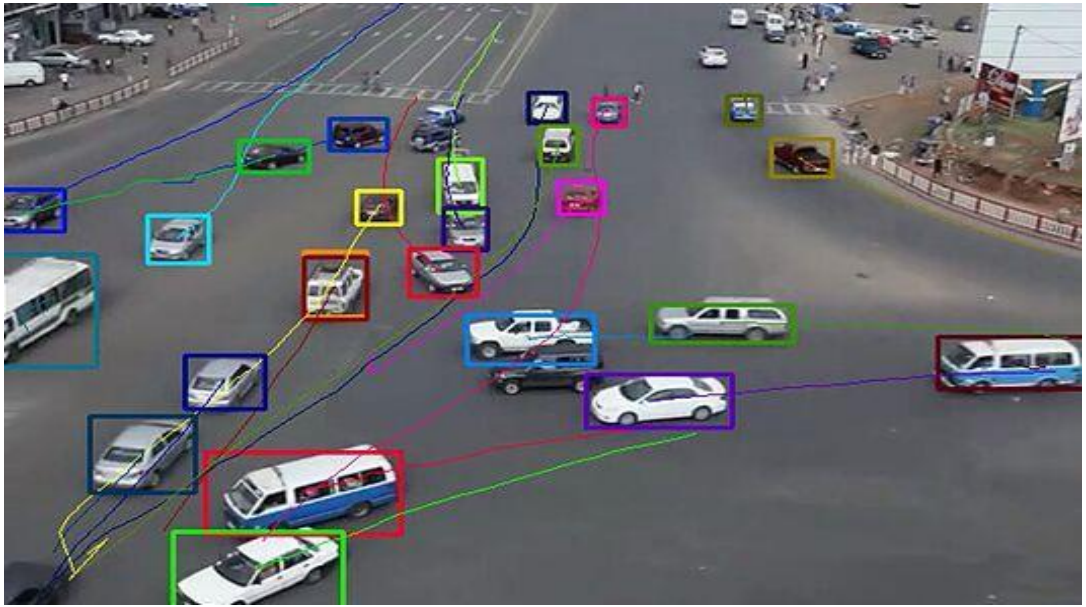

# Machine Learning Tasks : Generation

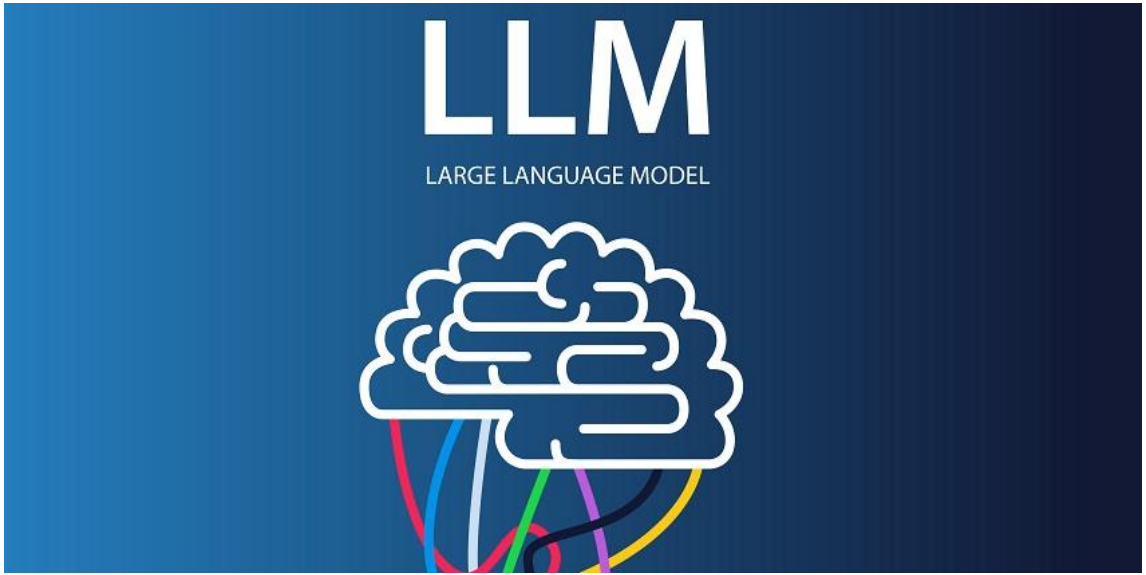


https://www.uctoday.com/unified-communications/what-is-a-large-language-model-defining-llms/

# Machine Learning Tasks : Generation

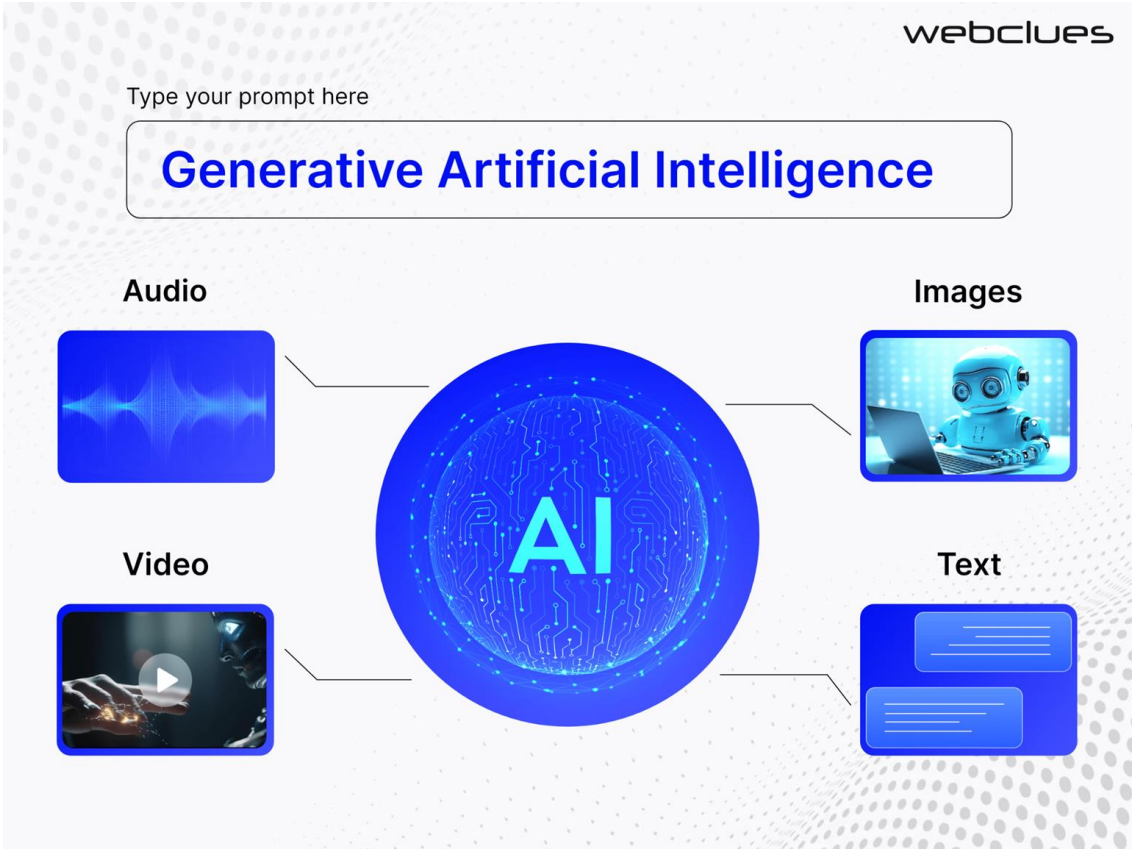

# Machine Learning Tasks : Generation


- Missing Data
- Synthetic Data
- Quality and Quantity
- Lack of Control

# Machine Learning Tasks : Generation

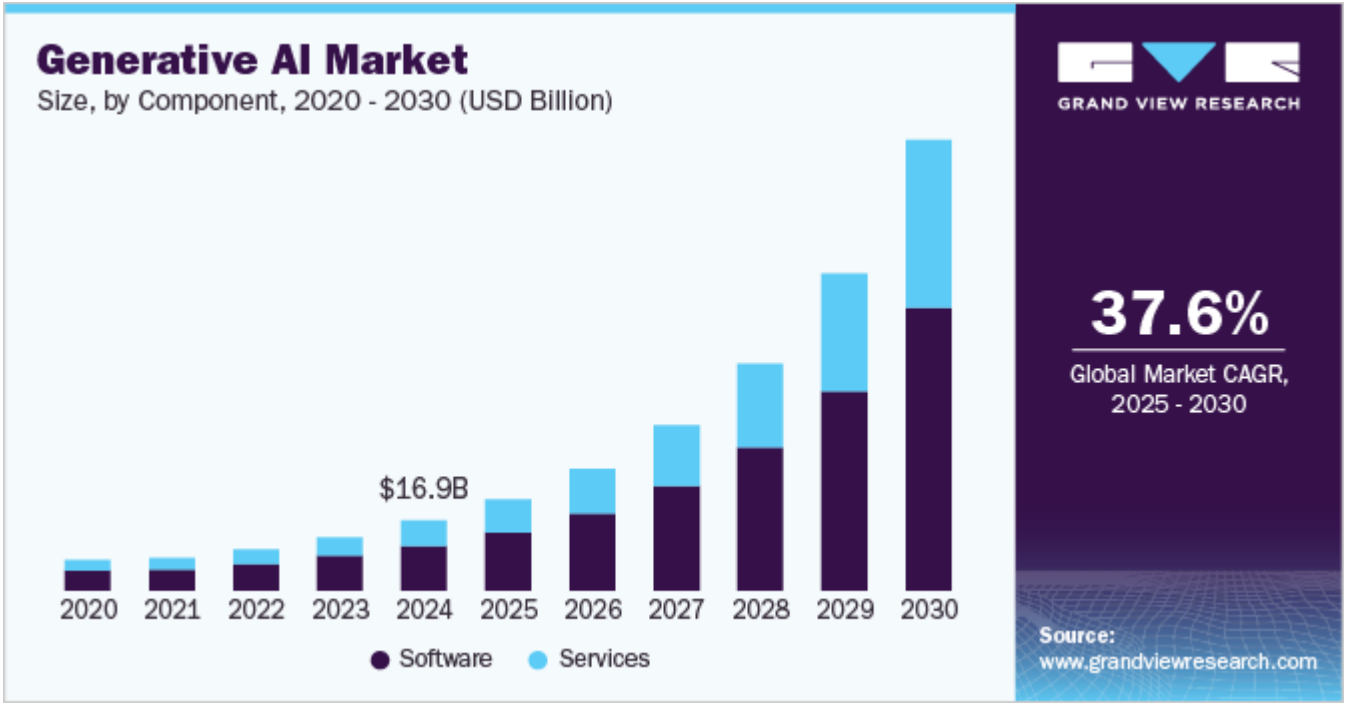


https://www.grandviewresearch.com/industry-analysis/generative-ai-market-report

# TRAINING FROM SCRATCH

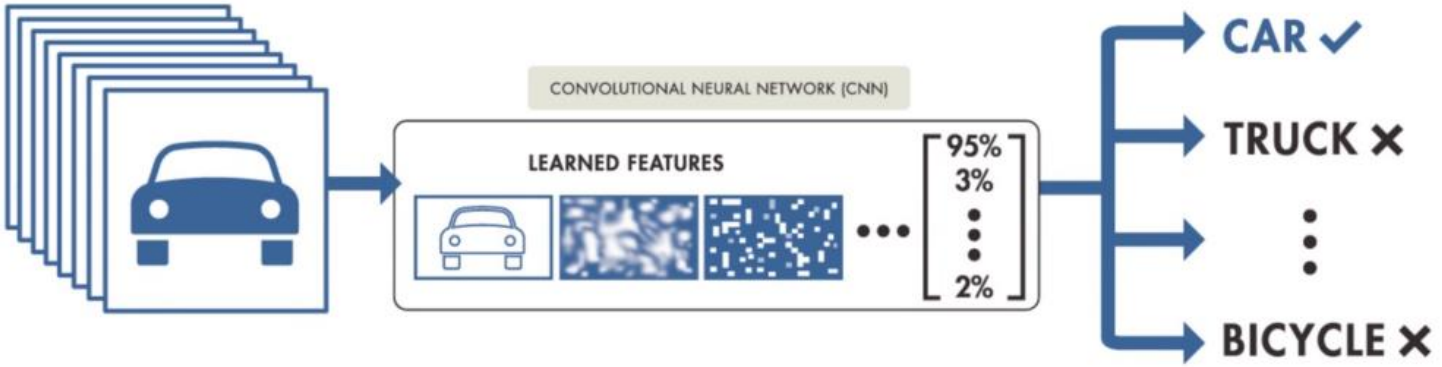


# TRANSFER LEARNING

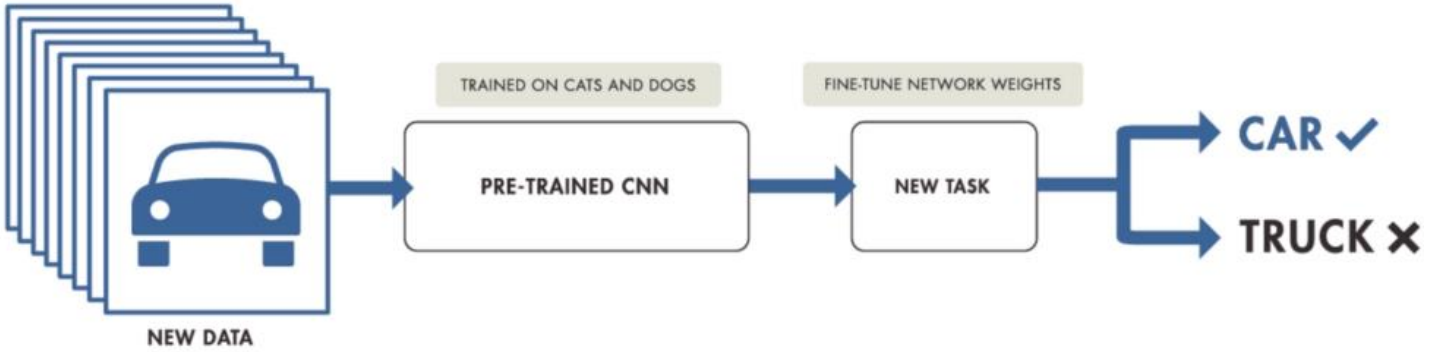

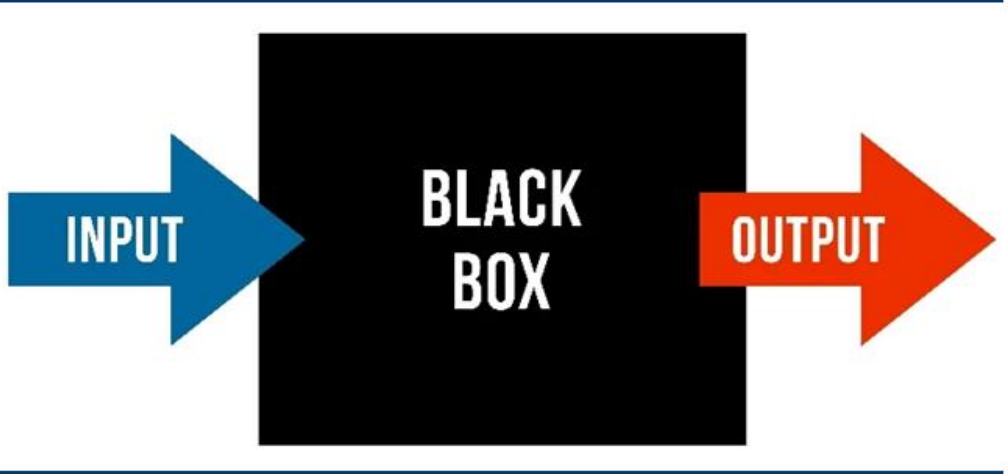


- Explanations vs Interpretability
    - Are the explanations interpretable?
- Evaluating Explanations
- Reliability of Explanations
- Insufficiency of human judgment to validate XAI

# eXplainable AI (XAI)

- **To understand Why and How Neural Networks do what they do??**
  - **Images: GradCam**
  - **Signals: CEFEs (This One is by me !!)**
  - **Tabular: SHAP**
  - **Lime, ELI5**

# Facilitating AI / ML

- Experts-in-the-loop
- Formal definitions
  - Problem
  - Explanations
    - TASK
    - DATA MODALITY
    - STAKEHOLDER

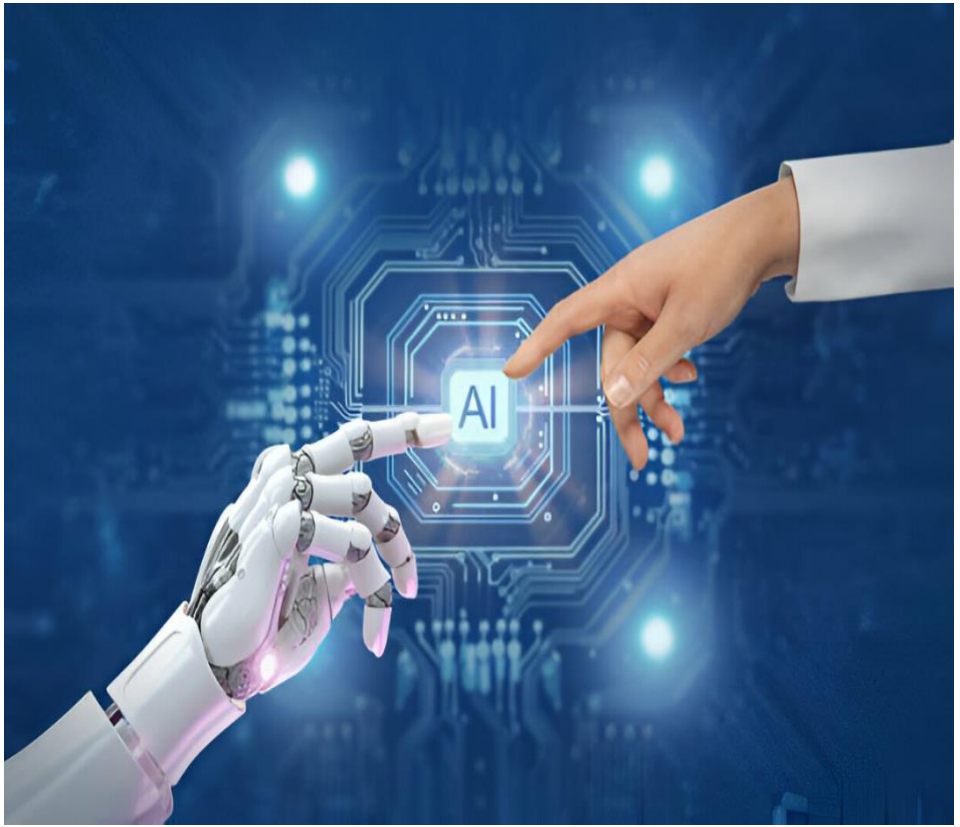

# CHALLENGES

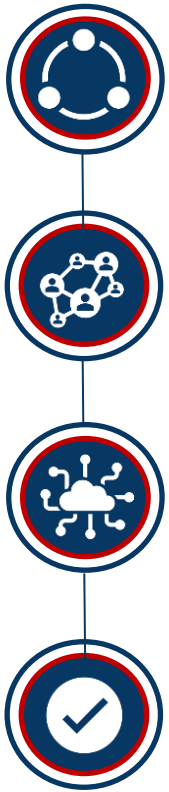

## Challenges in Development

## Challenges in Adoption

## Challenges in Deployment

## Challenges in Implementation

# Challenges in Development

| | |
|---|---|
| **Data** | • High quality<br>• Large quantity |
| **Bias** | • Disparate Impact<br>• Skewed Learning |
| **Robustness** | • Model Drift<br>• False Positives and Negatives |

# Challenges in Adoption

| | |
|---|---|
| **Data Privacy and Security** | • Data security<br>• Privacy crucial |
| **Ethical and Legal Concerns** | • Ethical dilemmas<br>• Inequitable care.<br>• Liability and accountability |
| **Resistance to Change** | • Job displacement<br>• Lack of trust in technology<br>• Complexity of integrating |

# Challenges in Deployment

| | |
|---|---|
| **Scalability** | • High computational power<br>• Secure data storage<br>• Reliable network access |
| **Compliance with Regulations** | • FDA (U.S.) and EMA (Europe) require evidence of safety and efficacy for AI-driven applications |
| **Monitoring and Maintenance** | • Evolving knowledge, protocols<br>• Continuous retraining |

# Challenges in Implementation

| | |
|---|---|
| **Integration with Existing Systems** | • Legacy Systems |
| **Interoperability** | • Diverse and fragmented data sources<br>• EHR systems, medical devices and imaging platforms |
| **Training and Usability** | • Training to use AI tools<br>• Lack of user-friendly interfaces |

# Thank you

**Dr. Sagnik Dakshit**
**sdakshit@uttyler.edu**
**www.sagnikdakshit.com**

UTTyler
THE UNIVERSITY OF TEXAS AT TYLER