AI Data and Model Landscape

Giang Nguyen

2023-11-15 AI4EOSC Platform: User's Workshop

Co-funded by the European Union

Artificial Intelligence (AI) and Machine Learning

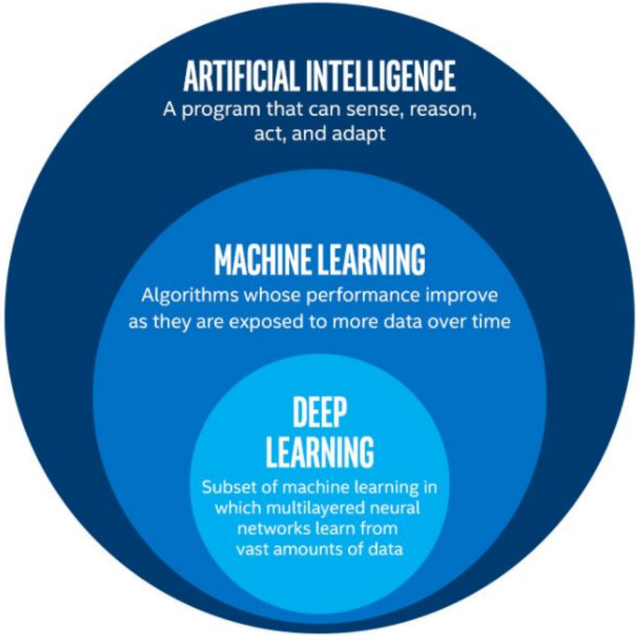

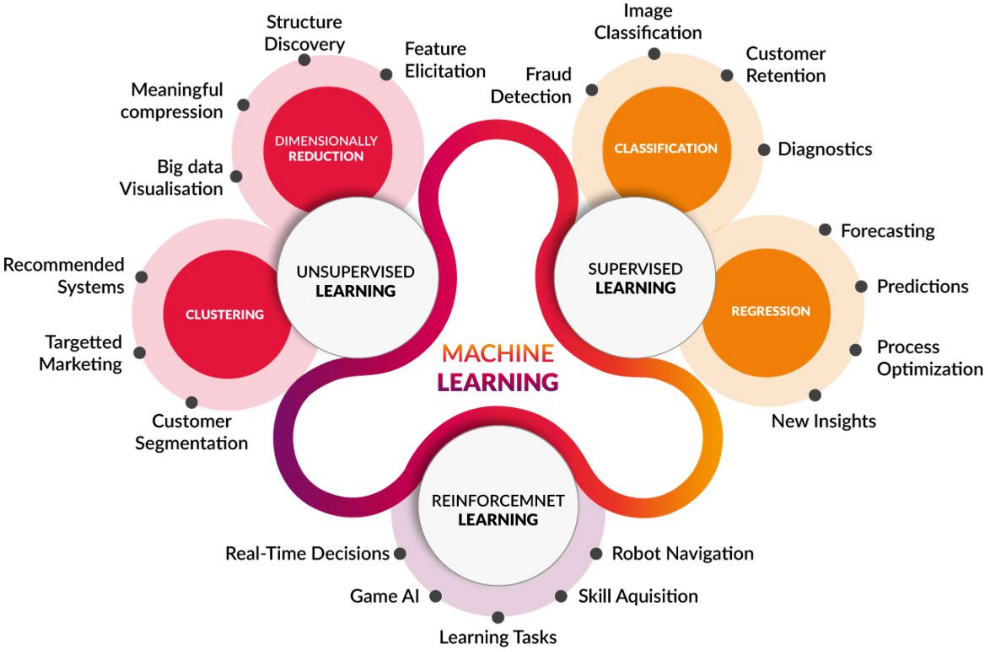

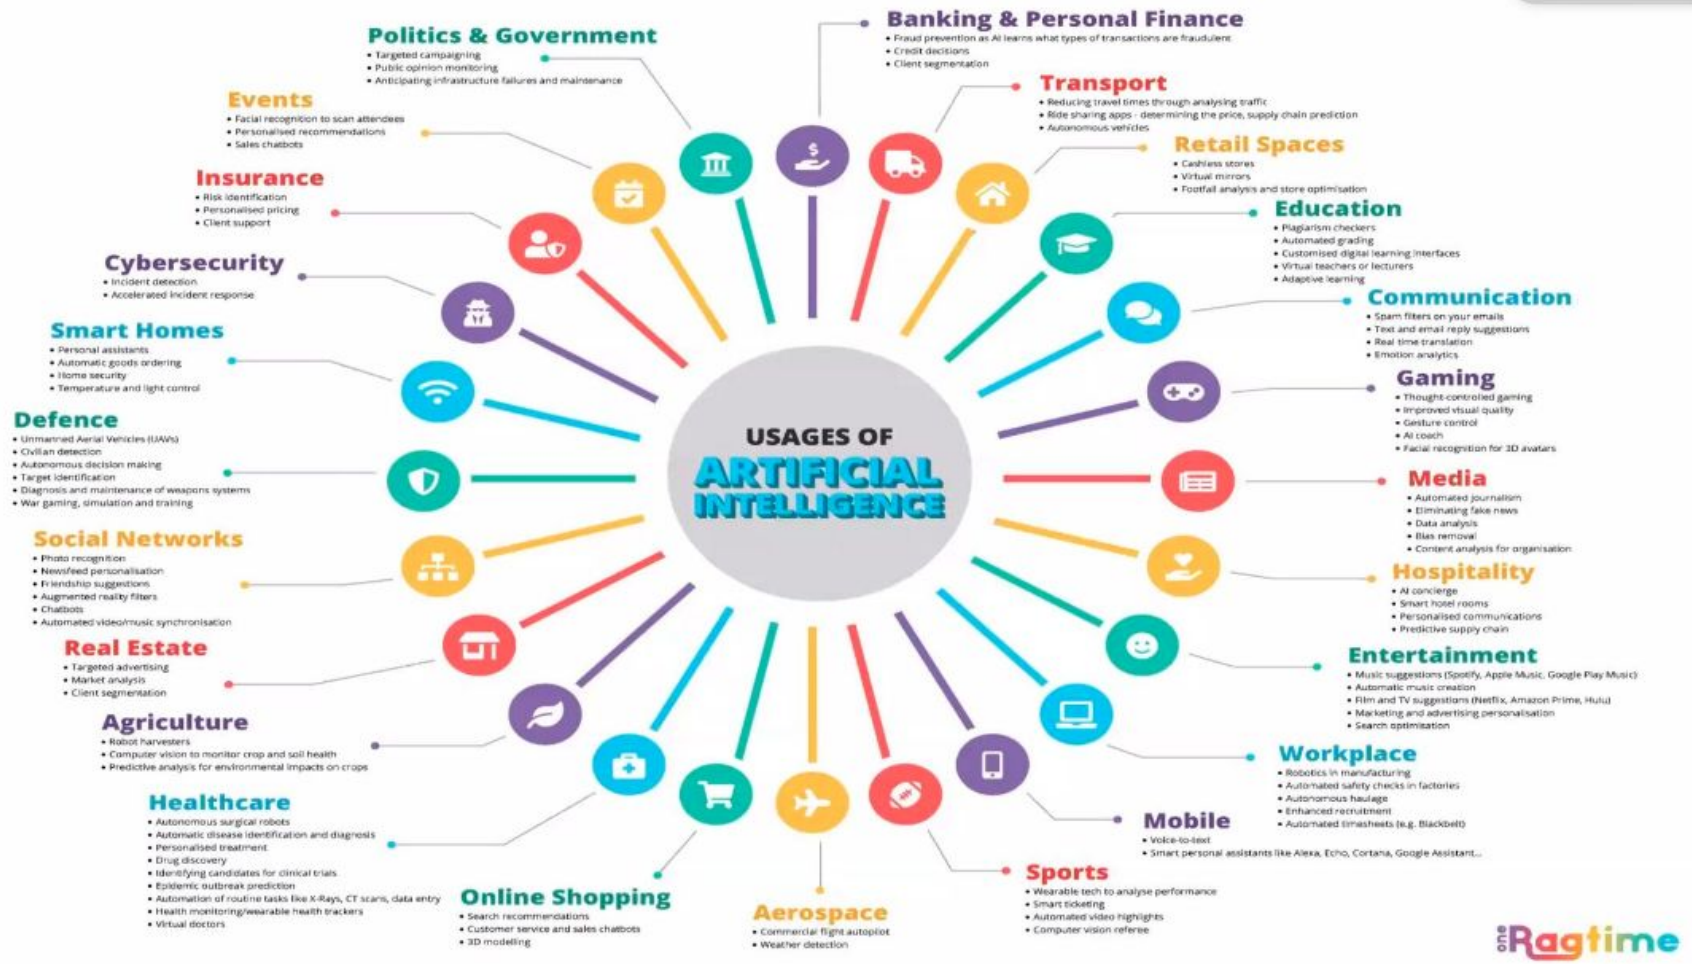
USAGES OF
ARTIFICIAL INTELLIGENCE
Politics & Government
• Targeted campaigning
• Public opinion monitoring
• Anticipating infrastructure failures and maintenance
Banking & Personal Finance
• Fraud prevention as AI learns what types of transactions are fraudulent
• Credit decisions
• Client segmentation
Transport
• Reducing travel times through analysing traffic
• Ride sharing apps - determining the price, supply chain prediction
• Autonomous vehicles
Events
• Facial recognition to scan attendees
• Personalised recommendations
• Sales chatbots
Retail Spaces
• Cashless stores
• Virtual mirrors
• Footfall analysis and store optimisation
Insurance
• Risk identification
• Personalised pricing
• Client support
Education
• Plagiarism checkers
• Automated grading
• Customised digital learning interfaces
• Virtual teachers or lecturers
• Adaptive learning
Cybersecurity
• Incident detection
• Accelerated incident response
Communication
• Spam filters on your emails
• Text and email reply suggestions
• Real time translation
• Emotion analytics
Smart Homes
• Personal assistants
• Automatic goods ordering
• Home security
• Temperature and light control
Gaming
• Thought-controlled gaming
• Improved visual quality
• Gesture control
• AI coach
• Facial recognition for 3D avatars
Defence
• Unmanned Aerial Vehicles (UAVs)
• Civilian detection
• Autonomous decision making
• Target identification
• Diagnosis and maintenance of weapons systems
• War gaming, simulation and training
Media
• Automated journalism
• Eliminating fake news
• Data analysis
• Bias removal
• Content analysis for organisation
Social Networks
• Photo recognition
• Newsfeed personalisation
• Friendship suggestions
• Augmented reality filters
• Chatbots
• Automated video/music synchronisation
Hospitality
• AI concierge
• Smart hotel rooms
• Personalised communications
• Predictive supply chain
Real Estate
• Targeted advertising
• Market analysis
• Client segmentation
Entertainment
• Music suggestions (Spotify, Apple Music, Google Play Music)
• Automatic music creation
• Film and TV suggestions (Netflix, Amazon Prime, Hulu)
• Marketing and advertising personalisation
• Search optimisation
Agriculture
• Robot harvesters
• Computer vision to monitor crop and soil health
• Predictive analysis for environmental impacts on crops
Workplace
• Robotics in manufacturing
• Automated safety checks in factories
• Autonomous haulage
• Enhanced recruitment
• Automated timesheets (e.g. Blackbelt)
Healthcare
• Autonomous surgical robots
• Automatic disease identification and diagnosis
• Personalised treatment
• Drug discovery
• Identifying candidates for clinical trials
• Epidemic outbreak prediction
• Automation of routine tasks like X-Rays, CT scans, data entry
• Health monitoring/wearable health trackers
• Virtual doctors
Mobile
• Voice-to-text
• Smart personal assistants like Alexa, Echo, Cortana, Google Assistant...
Online Shopping
• Search recommendations
• Customer service and sales chatbots
• 3D modelling
Aerospace
• Commercial flight autopilot
• Weather detection
Sports
• Wearable tech to analyse performance
• Smart ticketing
• Automated video highlights
• Computer vision referee
oneRagtime

The 2023 MAD: Machine Learning, AI and Data

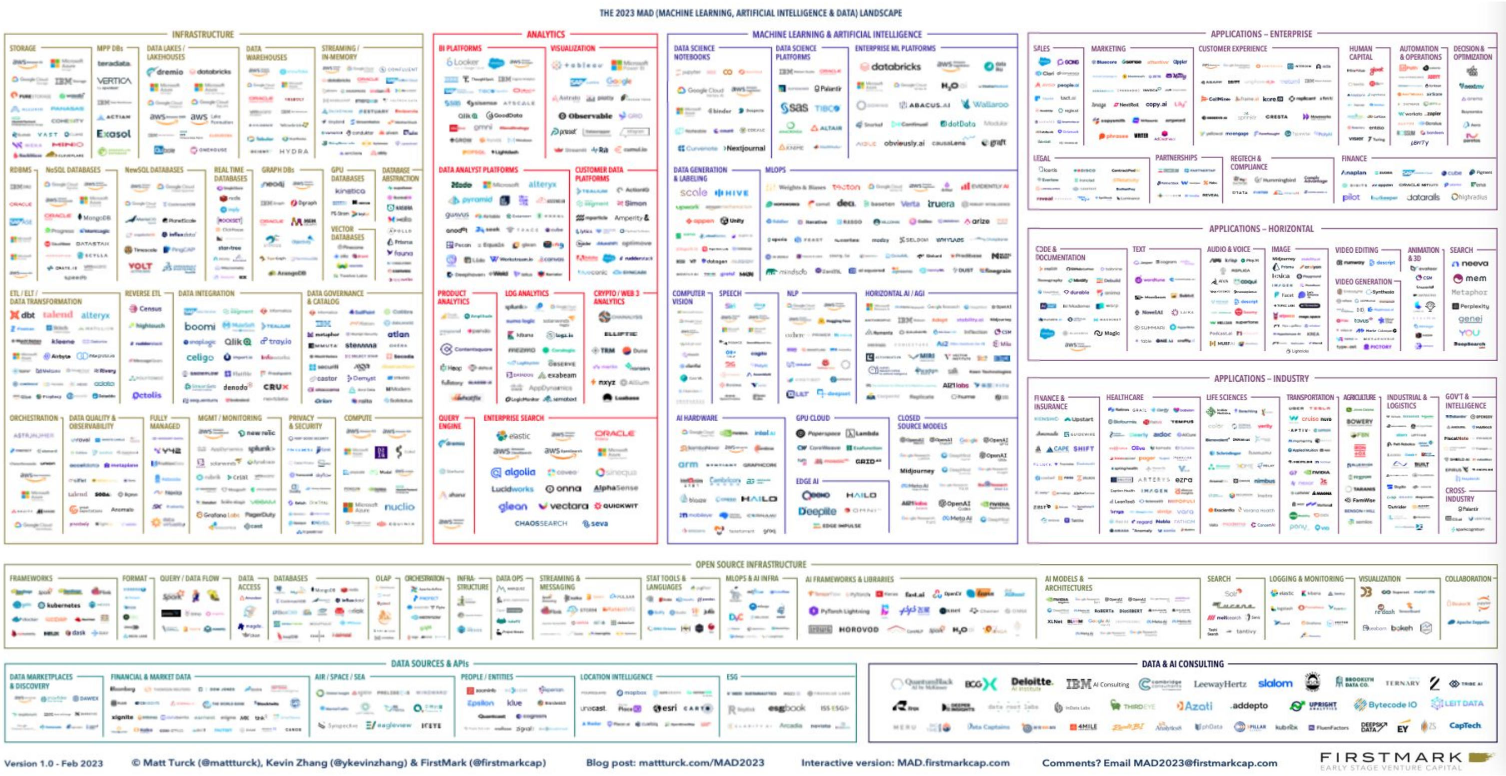

Exponential acceleration of Generative AI

ChatGPT - a kind of general purpose intelligence

AI generated content, AI bias, AI hallucination → **Responsible AI**

Hadoop - gradual **disappearance** of the Big Data technology

Crypto/web3 analytics

Data mesh, products, contracts: dealing with software complexity

Matt Turck 2023 MAD Landscape (4 long parts) - Interactive large version of the landscape is here

The 2021 MAD: Machine Learning, AI and Data

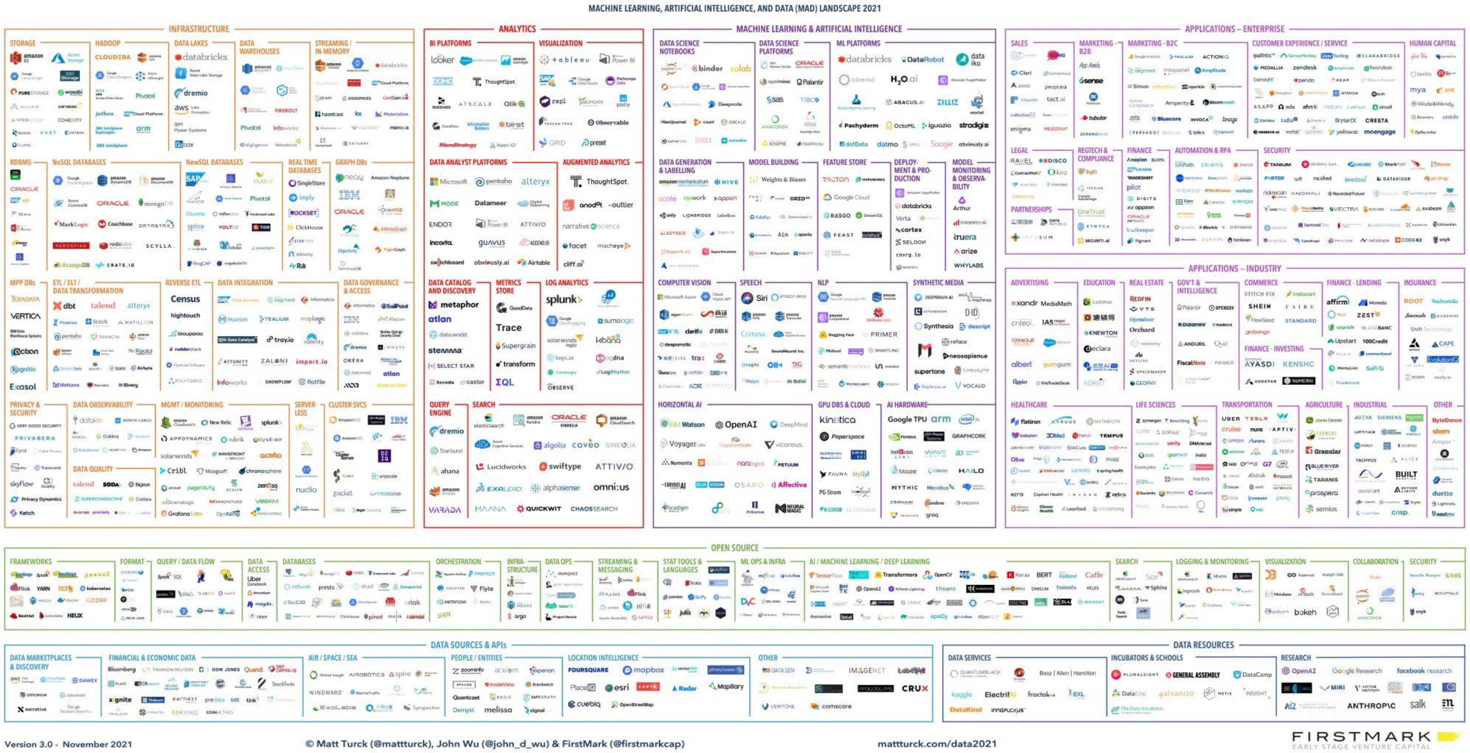

Data mesh

Busy year for **DataOps**

It's time for real time

Metrics stores

Reverse ETL

Cross-org data sharing

Privacy and security

Matt Turck 2021 MAD Landscape - Full resolution version of the landscape image here

STU FIIT

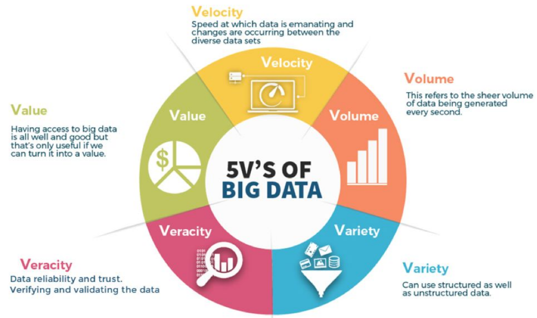

Velocity
Speed at which data is emanating and changes are occurring between the diverse data sets
Volume
This refers to the sheer volume of data being generated every second.
Value
Having access to big data is all well and good but that's only useful if we can turn it into a value.
Veracity
Data reliability and trust. Verifying and validating the data
Variety
Can use structured as well as unstructured data.
Velocity
Value
Volume
Veracity
Variety
5V'S OF BIG DATA

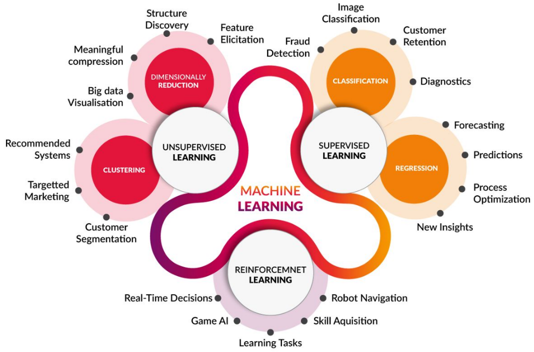

MACHINE LEARNING
UNSUPERVISED LEARNING
SUPERVISED LEARNING
REINFORCEMENET LEARNING
DIMENSIONALLY REDUCTION
CLASSIFICATION
CLUSTERING
REGRESSION
Structure Discovery
Feature Elicitation
Meaningful compression
Big data Visualisation
Image Classification
Fraud Detection
Customer Retention
Diagnostics
Forecasting
Predictions
Process Optimization
New Insights
Recommended Systems
Targetted Marketing
Customer Segmentation
Real-Time Decisions
Robot Navigation
Game AI
Skill Aquisition
Learning Tasks

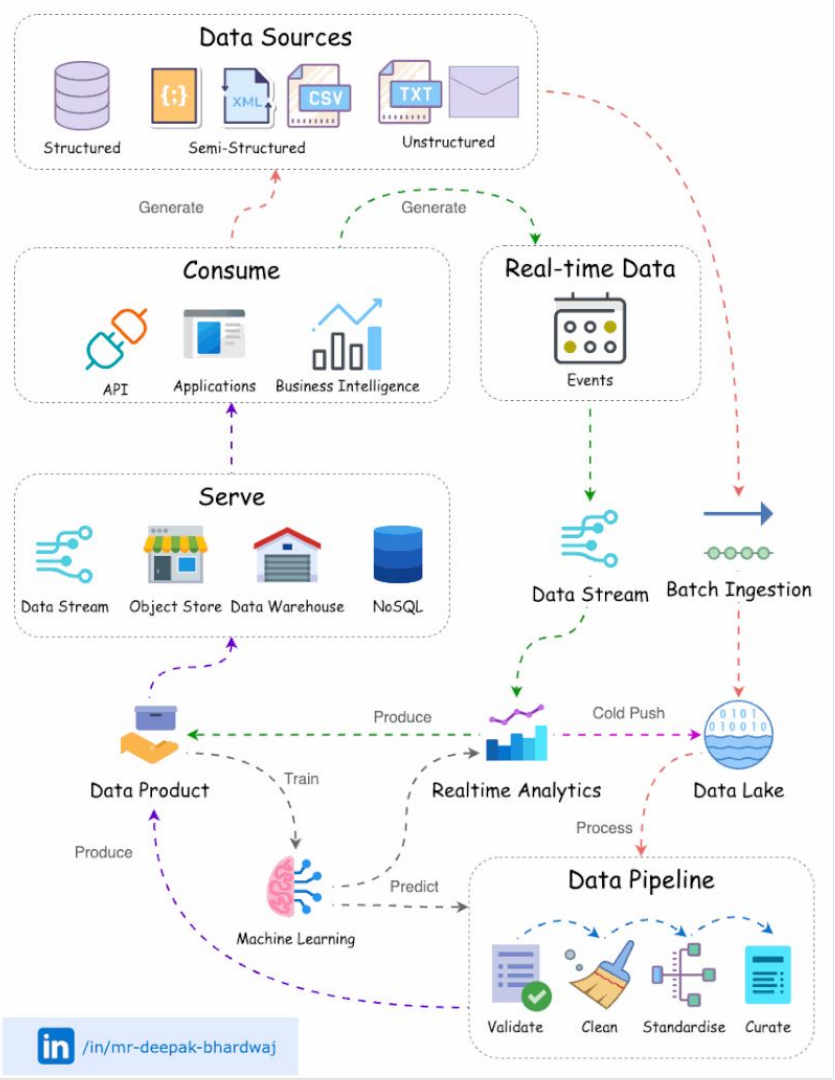

Data Sources
Structured
Semi-Structured
Unstructured
XML
CSV
TXT
Generate
Generate
Consume
API
Applications
Business Intelligence
Real-time Data
Events
Serve
Data Stream
Object Store
Data Warehouse
NoSQL
Data Stream
Batch Ingestion
Produce
Cold Push
Data Product
Realtime Analytics
Data Lake
Train
Process
Produce
Predict
Machine Learning
Data Pipeline
Validate
Clean
Standardise
Curate
/in/mr-deepak-bhardwaj

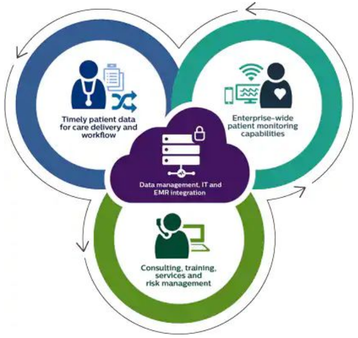

Timely patient data for care delivery and workflow
Enterprise-wide patient monitoring capabilities
Data management, IT and EMR integration
Consulting, training, services and risk management

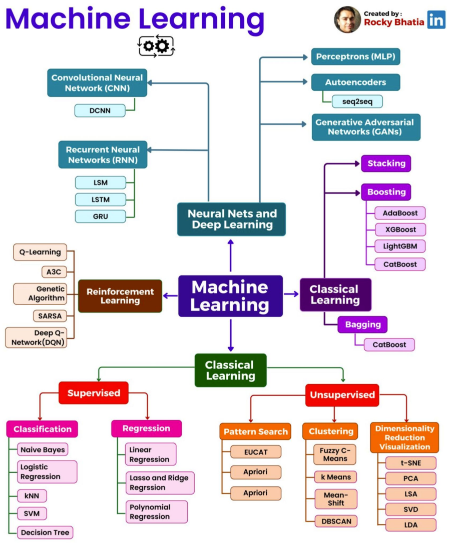

Machine Learning
Created by : Rocky Bhatia
Convolutional Neural Network (CNN)
DCNN
Recurrent Neural Networks (RNN)
LSM
LSTM
GRU
Perceptrons (MLP)
Autoencoders
seq2seq
Generative Adversarial Networks (GANs)
Neural Nets and Deep Learning
Stacking
Boosting
AdaBoost
XGBoost
LightGBM
CatBoost
Bagging
CatBoost
Q-Learning
A3C
Genetic Algorithm
SARSA
Deep Q-Network(DQN)
Reinforcement Learning
Machine Learning
Classical Learning
Classical Learning
Supervised
Unsupervised
Classification
Naive Bayes
Logistic Regression
kNN
SVM
Decision Tree
Regression
Linear Regression
Lasso and Ridge Regrssion
Polynomial Regression
Pattern Search
EUCAT
Apriori
Apriori
Clustering
Fuzzy C-Means
k Means
Mean-Shift
DBSCAN
Dimensionality Reduction Visualization
t-SNE
PCA
LSA
SVD
LDA

Imbalanced Datasets

The best way to approach Machine Learning solution is to start by analyzing and exploring the dataset in **EDA**

One of the common issues found in datasets is imbalanced classes issue

Data imbalance usually reflects an unequal distribution of classes within a dataset

Binary classification models without fixing this problem will be very biased

Feature correlation need to be improved

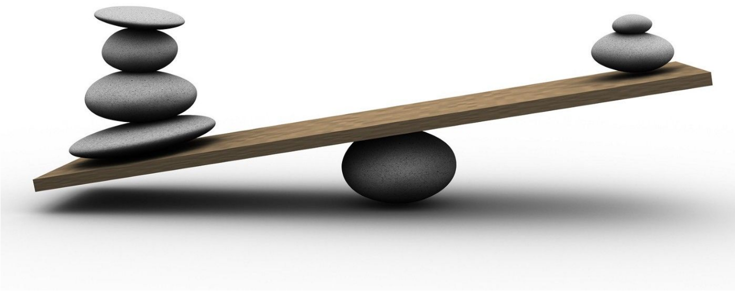

Undersampling

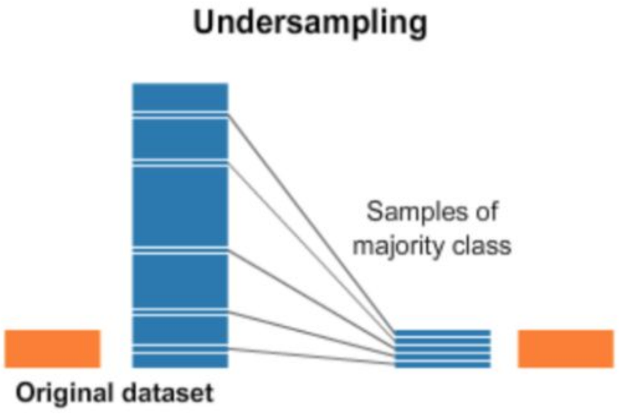

Randomly deleting some of the observations from the majority class in order to match the numbers with the minority class.

Oversampling

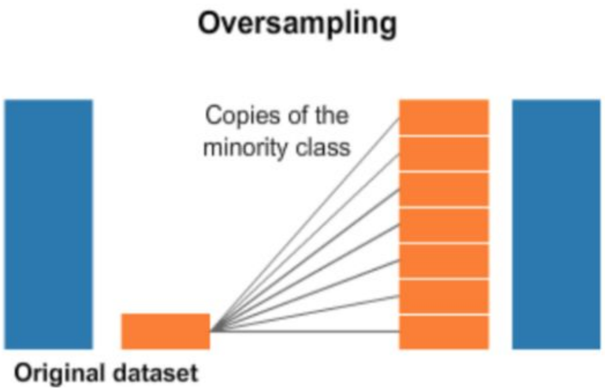

Synthetic Minority Over-sampling Technique (SMOTE) - looks at the feature space for the minority class data points and considers its *k* nearest neighbours.

Optimization for Machine Learning

- **Feature selection**
- **Model training**
- **Hyper-parameter tuning**
- Others

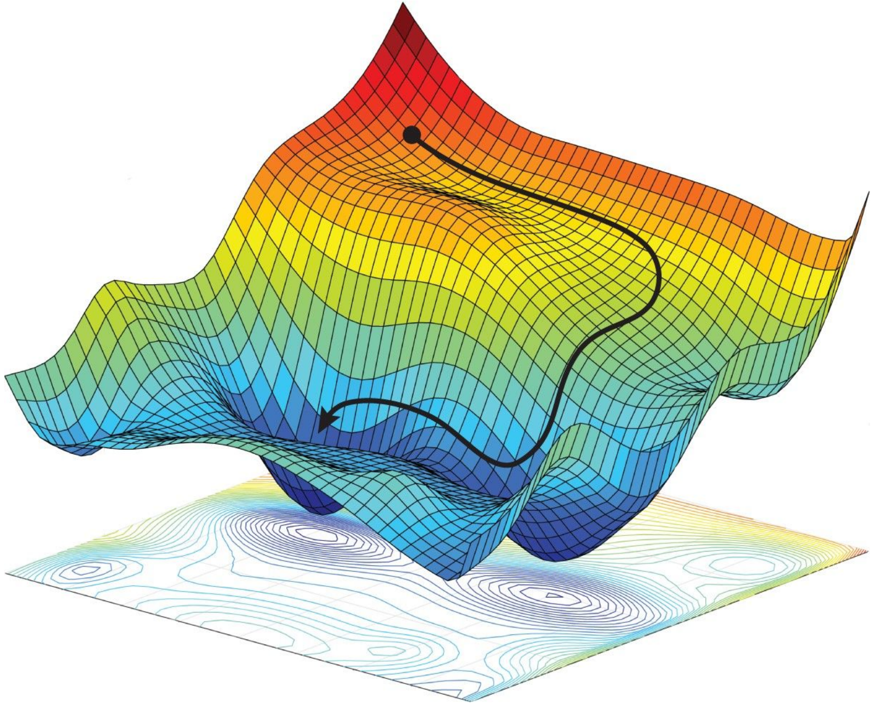

Making Better Decisions in an Uncertain World

Optimization techniques and applications

- Exhaustive search
- Local search
- Gradient descent
- Nature-inspired methods

- Agents
- Games
- Simulations

Imperfect but real-time decisions

Approximately optimal solutions in an acceptable time

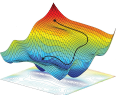

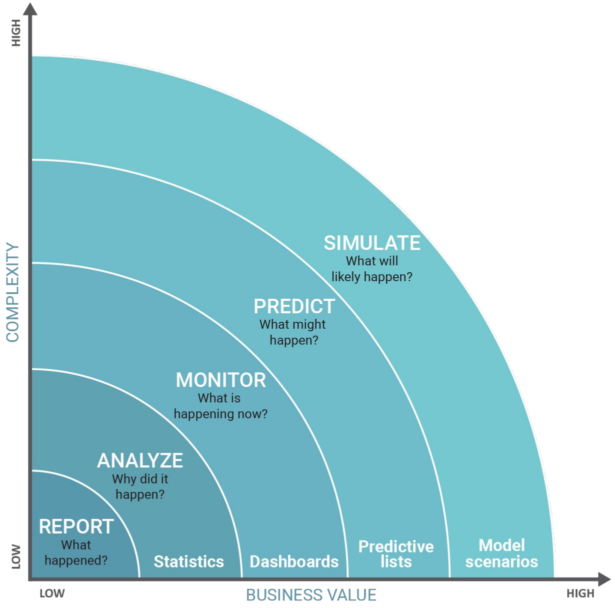

2003–2023: A Brief History of Big Data

HPC cluster and compute-intensive applications

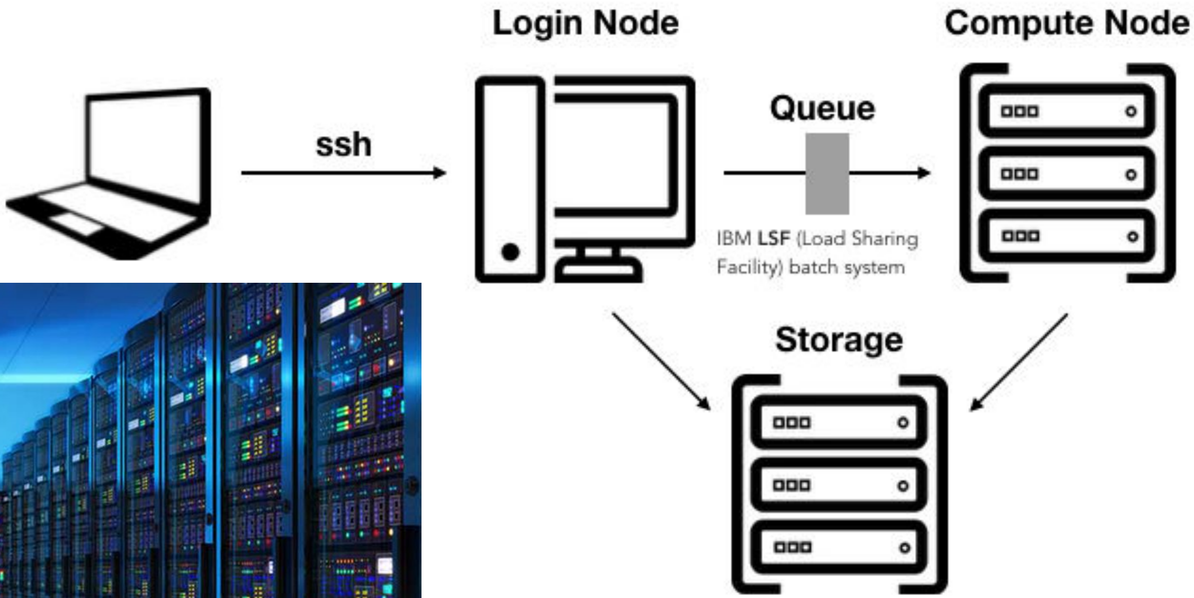

Hadoop cluster and data-intensive applications

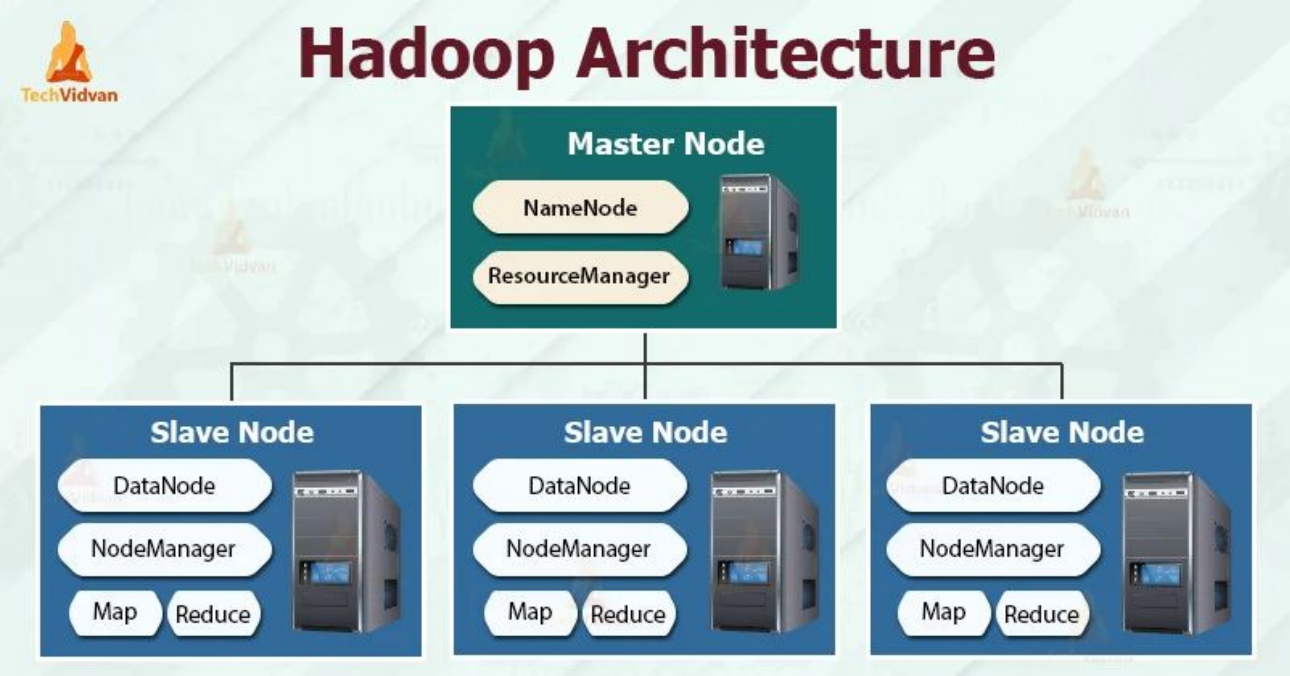

Summing up 20 years of Hadoop

2008–2012 Rise of the Hadoop vendors

- Cloudera, Hortonworks + MapR, Datastax
- Column-oriented storage format:
  Apache Parquet, Apache ORC

2010–2014 Hadoop 2.0 and the Spark

- Hadoop resource manager: **YARN**
  (Yet Another Resource Manager)
- **Apache Spark** faster, better, and great replacement for MapReduce
  ML into the ecosystem → Apache Mahute, **Apache MLLib**
- LinkedIn Apache Kafka
- Twitter Apache Storm
- Amazon (cloud) AWS with Netflix
- Microsoft Azure (cloud)
- **Google Cloud Platform**

2014–2016 Reaching the top

- Apache Spark, Apache Flink, Google's DataFlow
- Airbnb Apache **Airflow** (open-source schedulers)
- Spotify **Luigi** (open-source schedulers)
- Google's BigQuery
- Amazon's Redshift
- Snowflake

2016–2020 Containerisation and DL → the downfall of Hadoop

- Massive migration of data infrastructures to the cloud
- Amazon S3, Google Storage or Azure Blob Storage are used instead of HDFS
- Docker containerisation framework
- **Kubernetes** (container orchestration system)
- **DL →Tensorflow, Keras, PyTorch**

The need of GPUs - they are not in Hadoop clusters

Current shift from Hadoop ecosystems to Cloud services like Kubernetes

The rise of Deep Learning
The downfall of Hadoop

Hardware and software available for DL

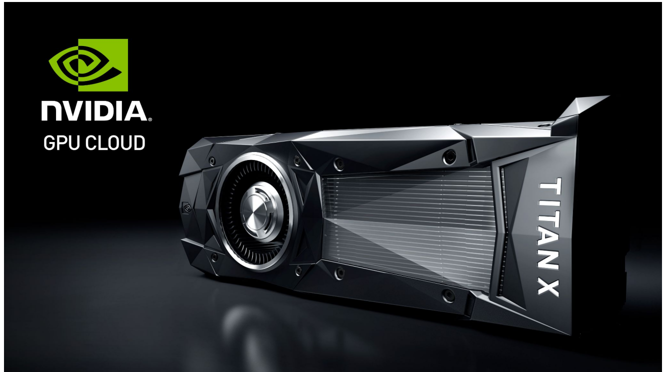

Deep Learning (DL) evolution timeline

1940s NN were proposed
1960s DNN were proposed
1990 **the first DL application** →
handwritten digit recognition - **LeNet**

Since 2010 **DL success** is due to the 3 factors

1. new algorithmic advances → better performance
2. availability of huge amount of data to train NNs
3. increase of computing power

Stanford Vision Lab, Stanford University, Princeton University

ImageNet Large-Scale Visual Recognition Challenge ILSVRC (2010-2017)

- Computer vision challenge
- Evaluates algorithms for object detection and image classification at large scale

Motivations

- to allow researchers to compare progress in detection across a wider variety of objects - taking advantage of the quite **expensive labeling effort**
- to measure the progress of **computer vision** for large scale image indexing for **retrieval and annotation**

Neural Networks and Deep Learning frameworks

TensorFlow

PyTorch

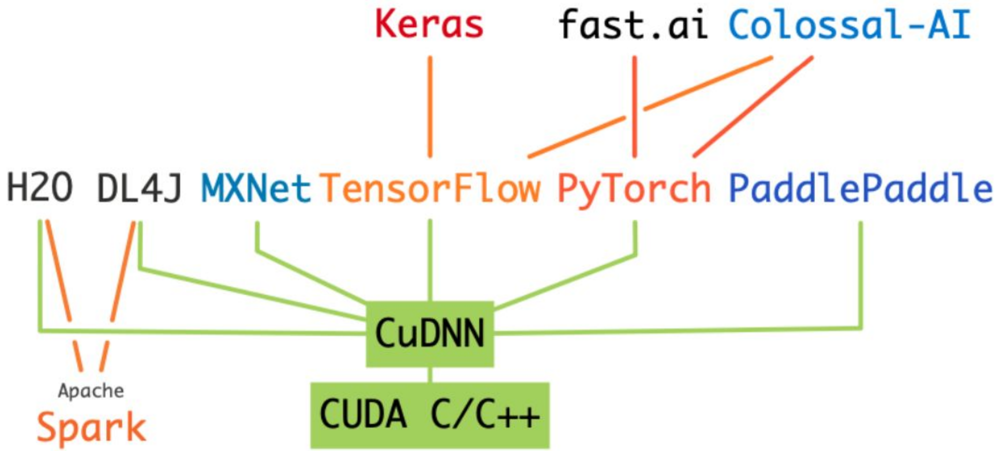

Machine Learning and Deep Learning frameworks and libraries for large-scale data mining: a survey (2019), updated 2023 for AI4EOSC

Convolutional Neural Networks (CNN)

Input - ConV - MaxPooling - ConV - MaxPooling - FC - FC(Output)

5-layer ConvNet

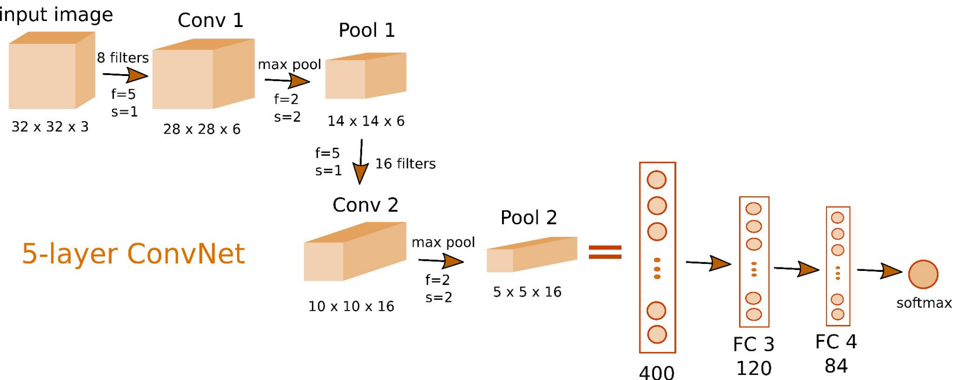

Iconic CNN architectures

- **1990 LeNet** deployed in ATMs to recognize digits for check deposits
- 2012 AlexNet the first CNN winned ILSVRC
- 2013 ZFNet one of the top at ILSVRC 2013
- **2014 VGG** classified as the second in ILSVRC
- 2015 GoogLeNet ILSVRC winner
  inception module composed of parallel connections
  which drastically reduced the number of parameters
- **2016 ResNet** Residual Net - ILSVRC winner
- 2016 SqueezeNet focuses in heavily reducing model size
  using deep compression
- UNet, **DenseNet**, MobileNet, ShuffleNet, ...

Object detection

Image classification: CNNs

- Classifies a picture
- Predicts probability of object

Classification ws. Localization: YOLO, R-CNN

- Detects an object in a picture
- Predicts probability of object and where it is located

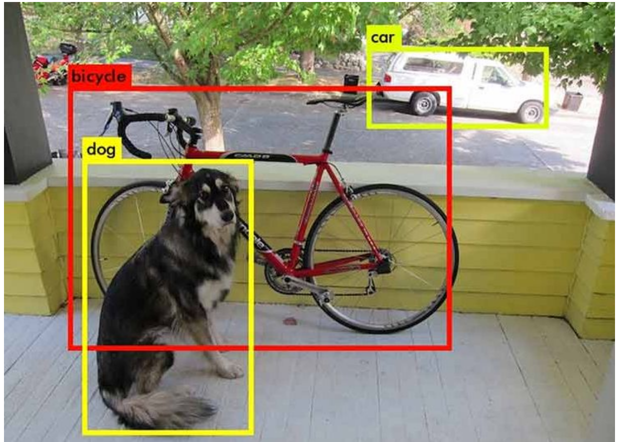

Face recognition

Face verification (one-to-one lookup)

- Is this the correct person?

Face recognition (one-to-one lookup)

- Is this one of the k people in the database?

2020: contradictory year for facial recognition

Jan 2020: America's first confirmed wrongful arrest by facial recognition technology (Robert Williams, **Black man**)

Jan 2020: **ClearviewAI** - small facial recognition company that ran its algorithm on a DB of **billions of pictures** grabbed **from social media**

Facial recognition is known to be less accurate for **darker-skinned** people

Facial recognition was/is widely used by **police departments** in the United States

The use of biometrics → Organic part of digital life

- Apple Face ID system to unlock iPhones

MIT technology review: Why 2020 was a pivotal, contradictory year for facial recognition, 2020

Image segmentation

Image segmentation

- segmenting an image into fragments
- assigning a label to each of those

Categories

- **Semantic segmentation**
- **Instance segmentation**
- Panoptic segmentation: combination of both semantic and instance segmentation

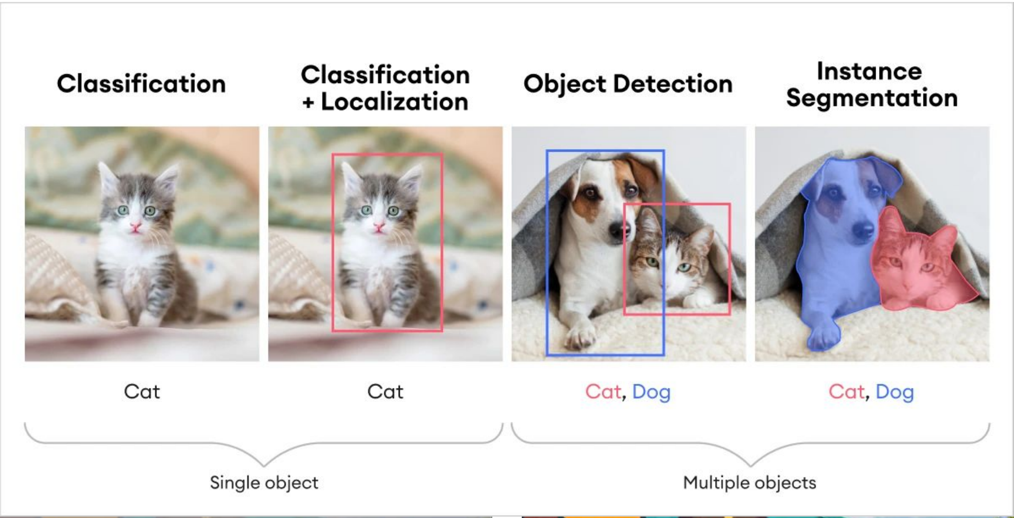

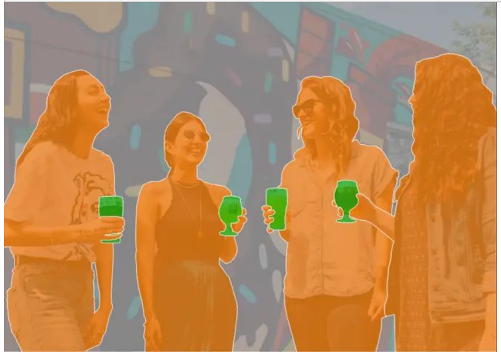

SuperAnnotate: Image segmentation detailed overview, 2023

Image segmentation techniques

Classic

- Region-based segmentation
- Edge detection segmentation
- Thresholding
- Clustering

Deep learning

- U-Net: encoder-decoder
- Mask R-CNN
- DeepLab Versions
- Interactive segmentation
- **Meta's SAM (Segment Anything Model)** → privacy-preserving images

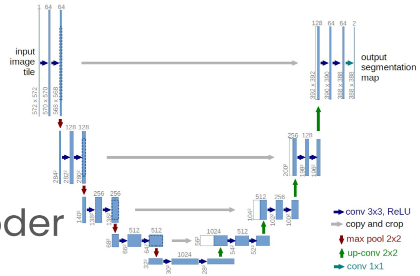

Benchmark datasets

- Pascal VOC dataset
- MS COCO dataset

Metrics

- IOU (Intersection Over Union)

Applications

Medical imaging, Video surveillance, Autonomous vehicles, Agriculture, Satellite imagery, Robotics, Art and design, Gaming, Fashion and retail

U-Net: Convolutional Networks for Biomedical Image Segmentation, 2015

R-CNN and YOLO → Real-time object detection

Object detection: YOLO, R-CNN

- Detects up to several objects in a picture
- Predicts probabilities of objects and where they are located

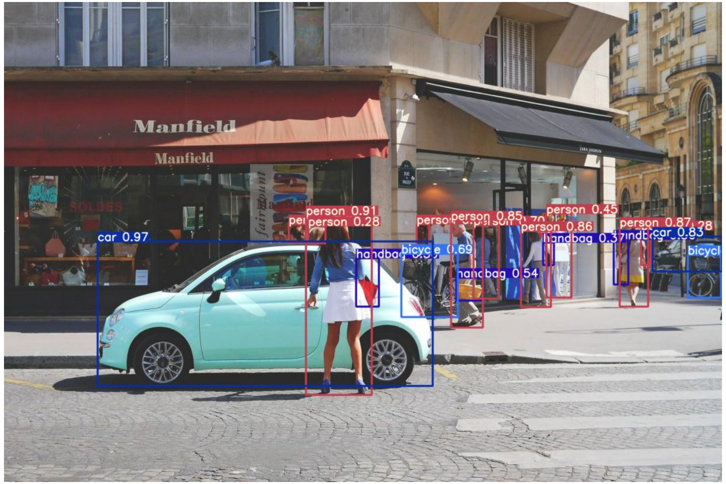

You only look once (YOLO)

- YOLO v3: Joseph Redmon, 2018
  - announced his Computer Vision research stop due to ethical concerns (Twitter 2.2020)
- YOLO v4: A. Bochkovskiy, fork, 4.2020
- YOLO v5: G. Jocher, PyTorch, 6.2020
- PP-YOLO: X. Long, Baidu, 7.2020
- **YOLOv7** in 2022
- YOLOv8 in 2023

The YOLO timeline

Recurrent Neural Networks

Encoder-Decoder

Natural Language Processing (NLP)

2019-2020 Big Years of NLP and DL

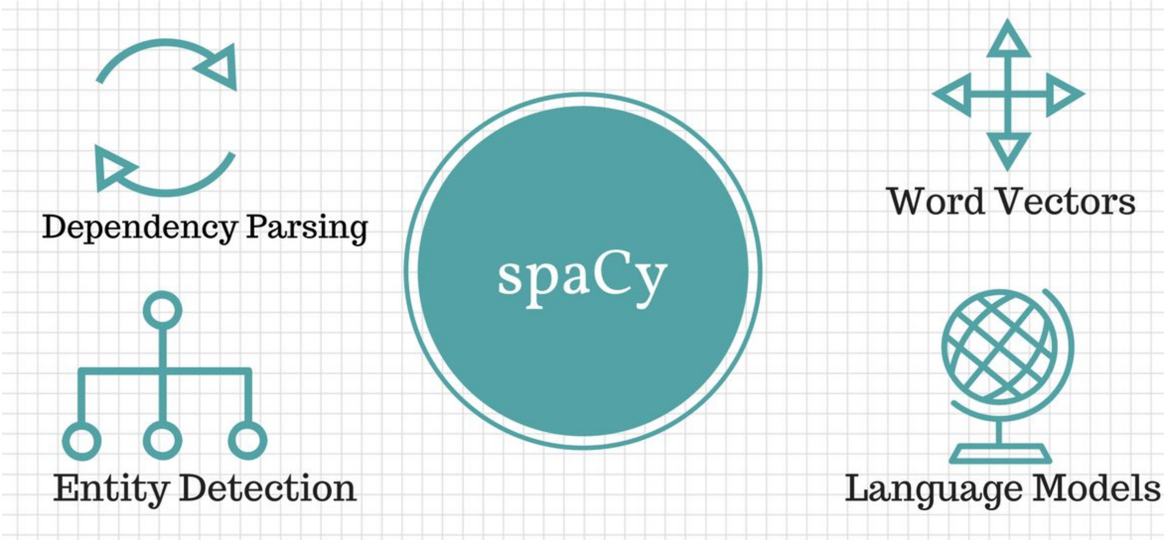

BERT (70 languages), Post-BERT pre-training: RoBERTa (2019, Facebook), XLNet (2019, Google), ALBERT (Google and TTI Chicago, 2019), T5 (Google, 2019), ELECTRA (2020), GPT-3 (2020)

Natural Language Processing (NLP)

- Feature Extraction for ML - The Art of the Science
- NLP in the context with ML/DL/AI
  - **NLP Pipelines**, NLTK, Stanza, Spacy, scikit-learn, …
  - Basic NLP task for text processing: stop-words removal, tokenization, …
  - Basic text processing: regular expression, normalization, extraction, …
  - Basic language modeling: probabilistic N-grams
- Text Modeling with Machine Learning
  - **Sentiment Analysis in Text**
- Text Preprocessing as Feature Extraction for ML
  - Text representation: **hand-crafted** and automatic
  - Feature Extraction: **Bag-of-Words, IF-IDF**, Word Embedding
  - Distance metrics, Document similarity
  - Encoding Data, Feature Selection, Dimensionality Reduction
- Distributional Semantics and Contextual Representation
  - NLP Transformer family: Encoder-Decoder
  - **Pretrained Large Language Models (LLM) with Deep Learning (DL)**

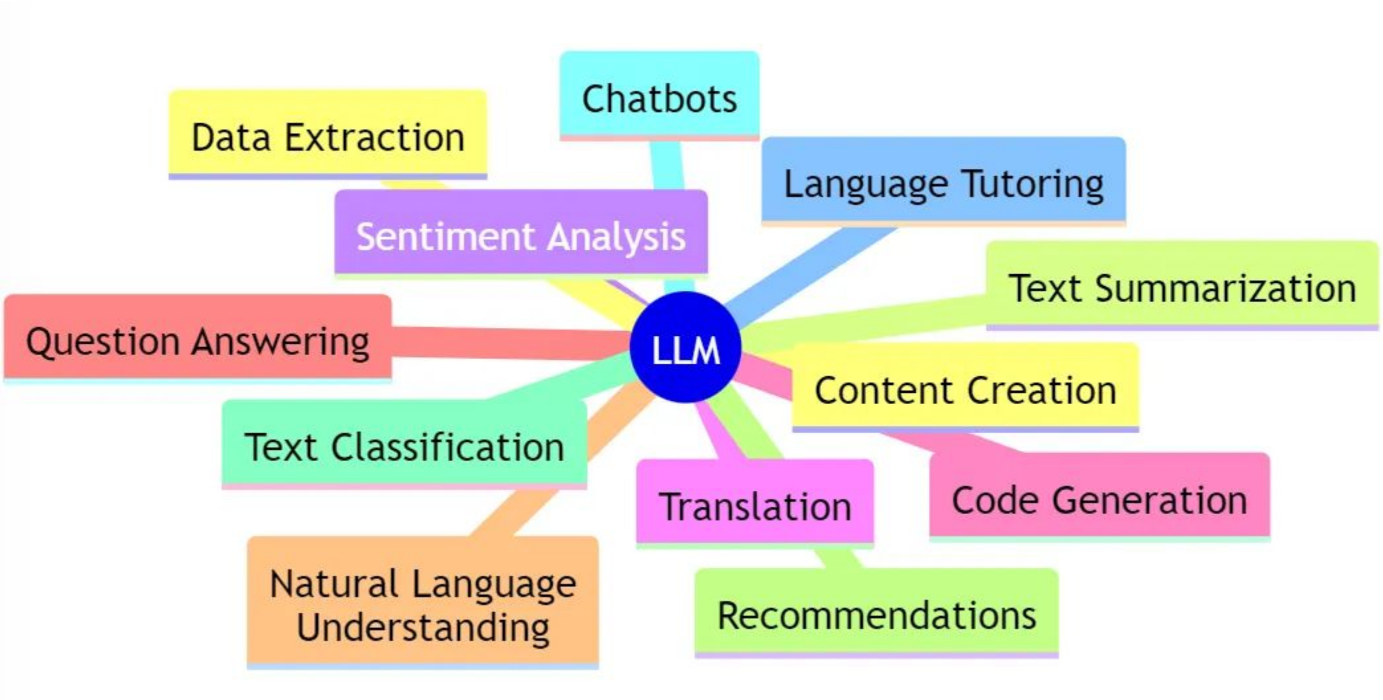
Chatbots
Data Extraction
Language Tutoring
Sentiment Analysis
Text Summarization
Question Answering
LLM
Content Creation
Text Classification
Code Generation
Translation
Natural Language Understanding
Recommendations

Model Deployment and Curation

Model Deployment - 87% have never been

4 years ago

it was unusual to find a company already deploying models in production.

2 years ago

it was unusual to find a company with 50 models in production.

Today

2 independent surveys estimate that
1/3 of companies that use AI have more than 50 models in production

Algorithmia	(2021)	said **40%**,
Arize	(2022)	said **36%**

Why do 87% of data science projects never make it into production?

Deploying your ML model is just the Beginning

Data and model curation in production!

Deployment → CD/CI → DevOps

DevOps speeds delivery of higher quality software by combining and automating the work of software development and IT operations teams

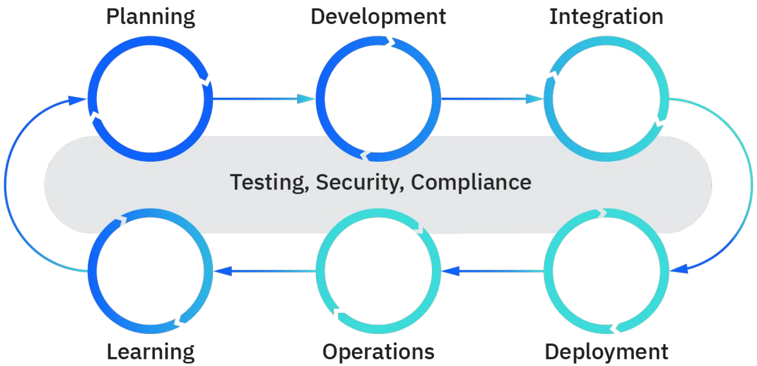

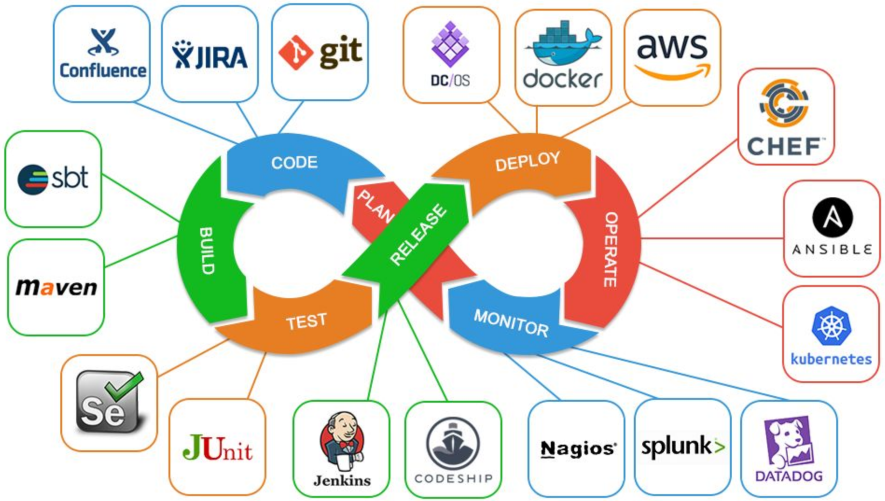

DevOps → MLOps and DataOps

DevOps

MLOps

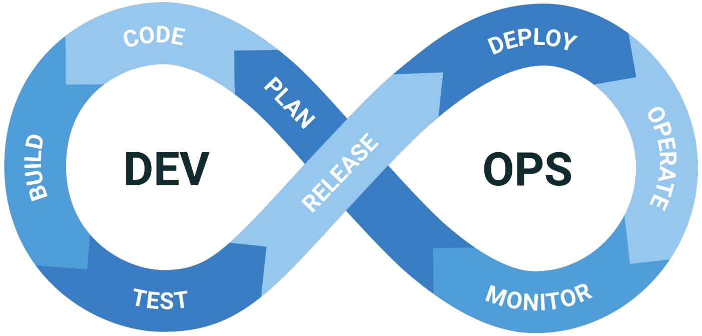

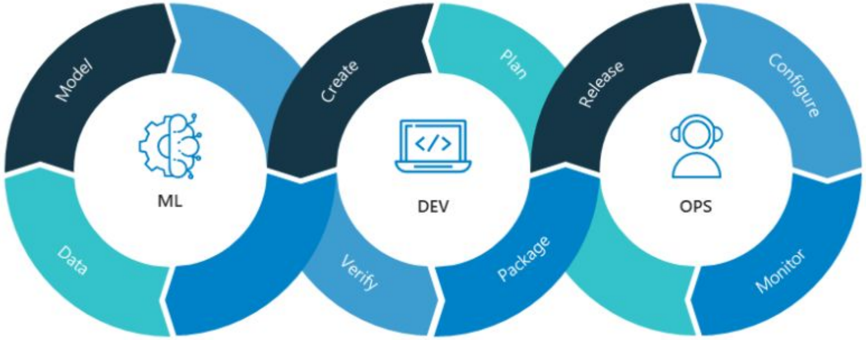

DevOps = Faster software delivery + software quality assurance (QA)
MLOps = DevOps applied to Machine Learning (ML) operations

Data distributions: DevOps and MLOps

In real world scenario of Machine Learning (ML) deployment

- The change in data distribution is called **data drift**
    - Data drift affects the performance of ML model used in deployment

- Data drift leads to **concept drift** (another name **model drift**)
    - Degradation of ML model performance in production

- Heavy work of ML engineering in **DevOps** and **MLOps**

What Is Generative AI?

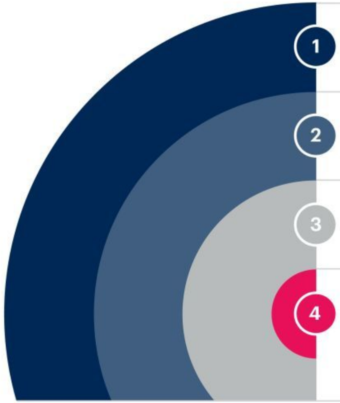

1	**Generative AI**	AI techniques that learn from data about existing artifacts and use this to generate new artifacts.
2	**Foundation Models**	Large machine learning models that are trained on a broad set of unlabeled data and are fine-tuned to a wide range of applications.
3	**Large Language Models (LLM)**	AI that is trained on vast amounts of text to interpret and generate human-like textual output.
4	**ChatGPT**	An OpenAI service that incorporates a conversational chatbot with LLM to create content.

Source: Gartner

Gartner

(NLP) Transformers

Transformer (2017)

"Attention is all you need"

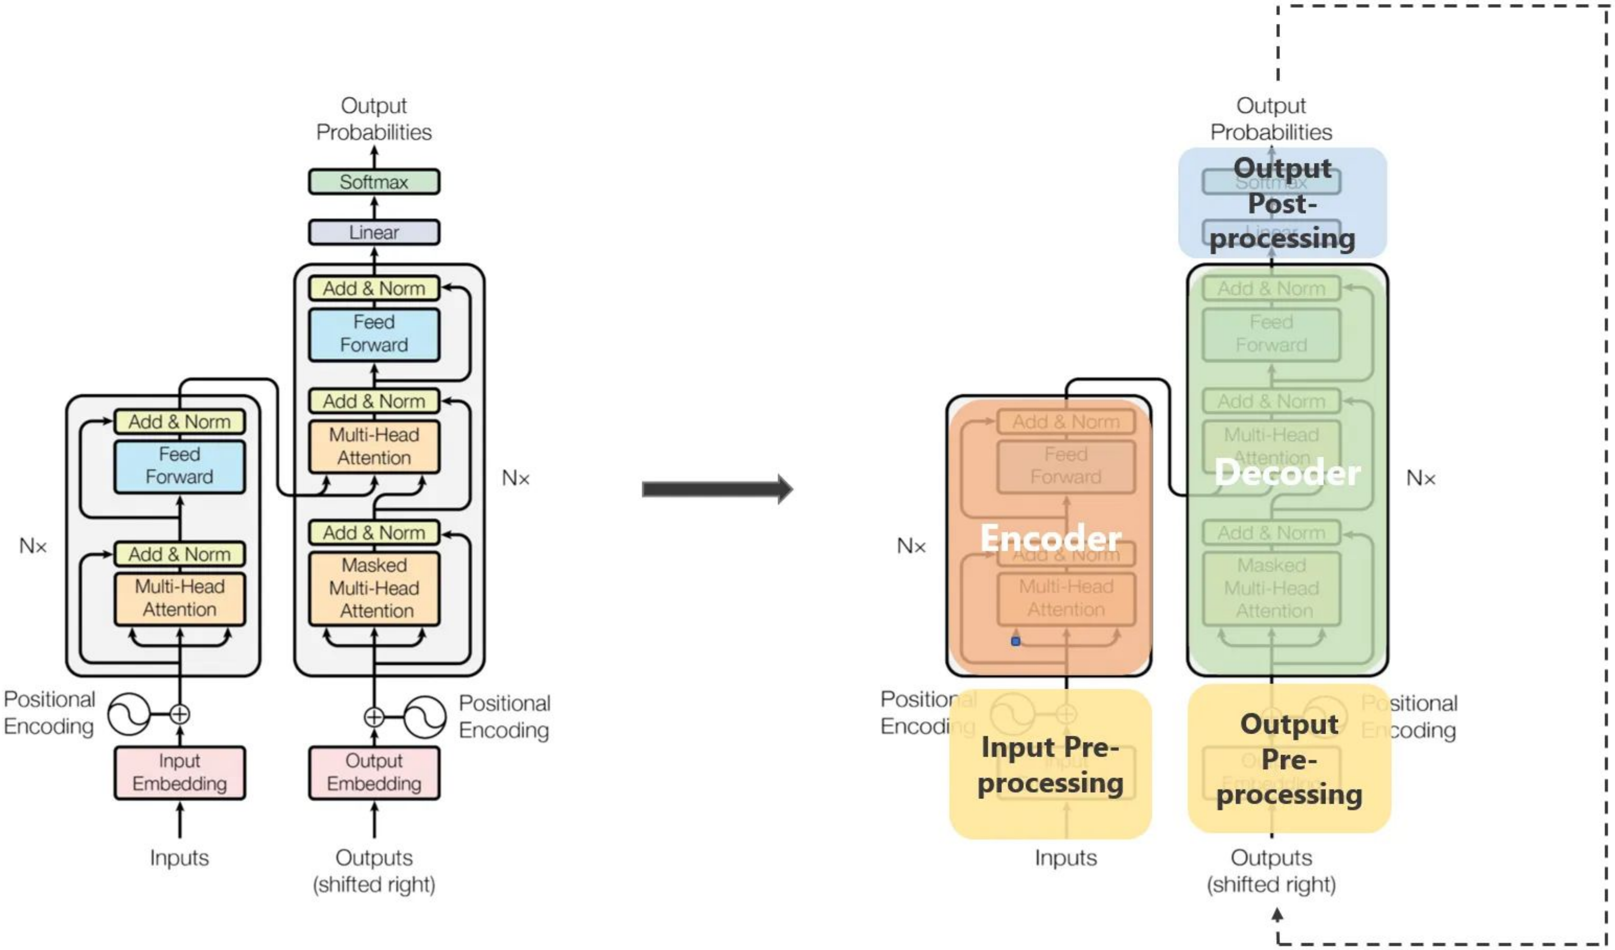

Ria Kulshrestha: Transformers OR as I like to call it Attention on Steroids, 2020

NLP Transformers: Benefits and Drawbacks

- **highly parallelizable**, meaning that they can process multiple parts of a sequence at the same time → significantly speeds up training and inference

- capture **long-term dependencies** in text → better understand the overall context and generate more coherent text

- **high computational demand**

- **sensitive to the quality and quantity of the training data**

2022 Leading Large Languages Models (LLMs)

BERT	Pre-training of Deep Bidirectional Transformers for Language Understanding
GPT-2	Language Models Are Unsupervised Multitask Learners
XLNet	Generalized Autoregressive Pretraining for Language Understanding
RoBERTa	A Robustly Optimized BERT Pretraining Approach
ALBERT	A Lite BERT for Self-supervised Learning of Language Representations
T5	Exploring the Limits of Transfer Learning with a Unified Text-to-Text Transformer
GPT-3	Language Models Are Few-Shot Learners

GPT-3.5, GPT-4, …

Leading Language Models For NLP In 2022

Vision Transformer (ViT)

2017, Transformers introduced for NLP

2020, adapted for computer vision → ViT

- The attention mechanism in a ViT repeatedly transforms representation vectors of **image patches**, incorporating more and more **semantic relations** between image patches in an image (like NLP way)
- CNNs achieve excellent results even with training based on data volumes that are **not as large** as those required by Vision Transformers

Using Transformers for Computer Vision, 2022, Vision Transformers (ViT): Revolutionizing Computer Vision, 2023

Vision Transformer vs. CNN

- ViT possess a different kind of bias toward exploring topological relationships between patches, which leads them to be able to capture also global and wider range relations but at the cost of a more onerous training in terms of data
- ViT more robust to input image distortions

Not so clear winner between CNN and ViT → combination of both

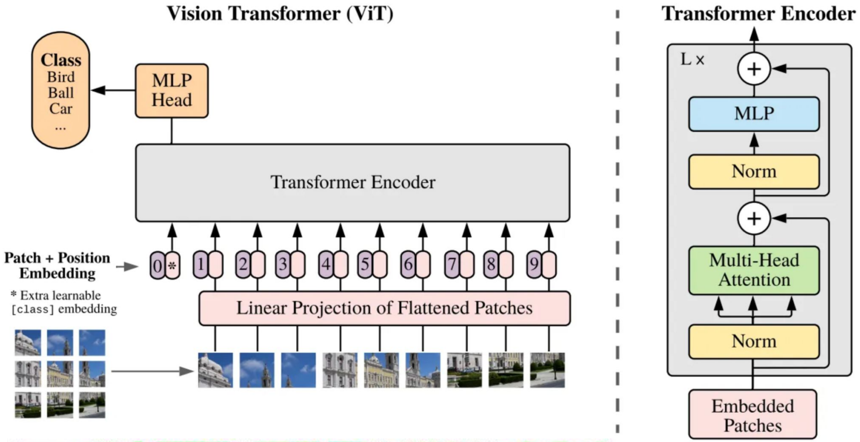

[Vision Transformers in Agriculture | Harvesting Innovation](), 2023

Large Multimodal Models (LMMs)

Generative AI applied to other modalities

- **Image Generation**:
  Dall-E | MidJourney | Stable Diffusion | DreamStudio
- **Audio Generation**:
  Whisper | AudioGen | AudioLM
- **Search Engines**: Neeva | You
- **Code Generation**: Copilot | Codex
- **Text Generation**: Jasper

Everything is still **NOT** working

2022: Why Meta's latest large language model survived only **three days** online

- Meta = Facebook's parent company
- **Galactica** - new large language model designed to **generate scientific text**
- Demo online **2022 November 15**, take down after **3 days**
- Model trained on **48 million** examples of
  - scientific articles,
  - websites,
  - textbooks, lecture notes, and
  - encyclopedias
- blind spot → severe limitations of large language models
- not able to distinguish truth from falsehood

2016: Microsoft launched a **chatbot** called **Tay on Twitter** — then shut it down **16h** later when Twitter users turned it into a racist, homophobic sexbot

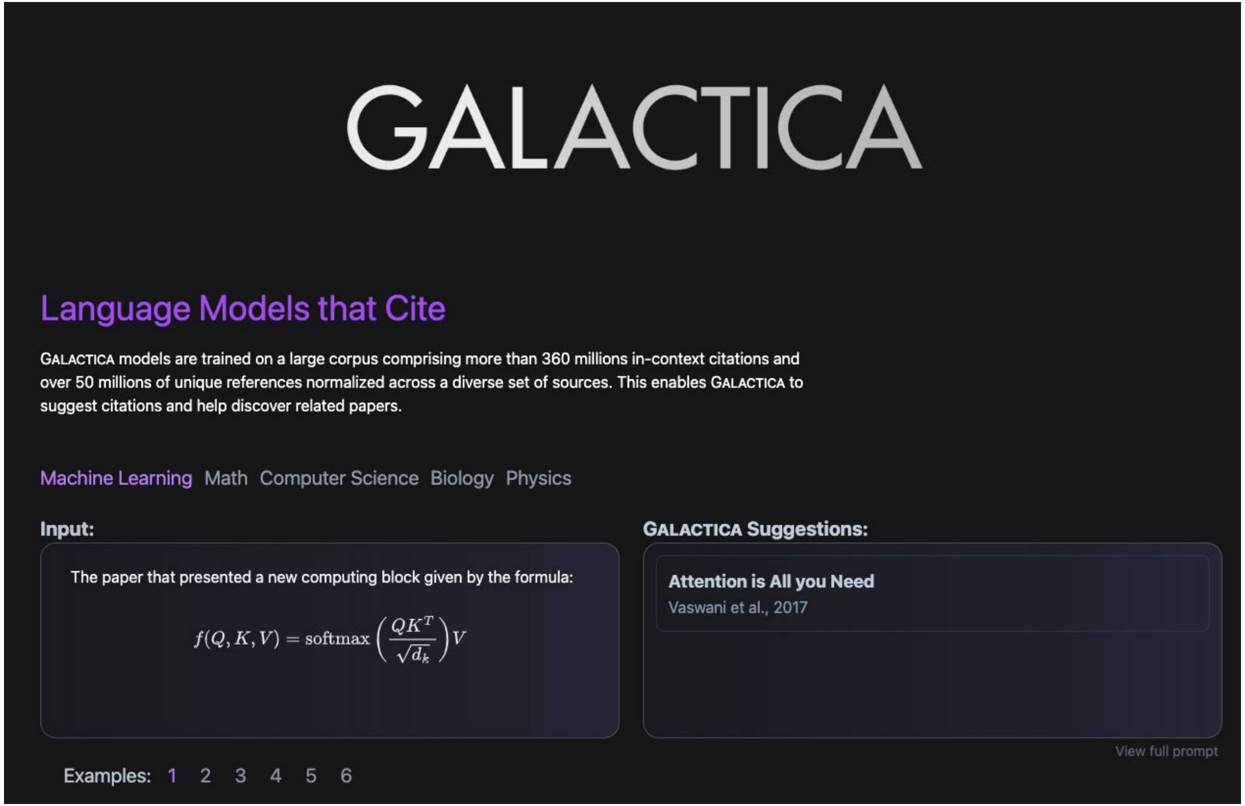

It is amazing, threatening, challenging. The technology WILL BE incredible.

AI and Ethics?

Something is working :) sometime

Language model: more truthful and **less toxic**, trained with **human-in-loop**

OpenAI just released GPT-3 text-davinci-003,

The results are impressive in comparison with with 002!

How does (OpenAI) GPT obtain its ability?

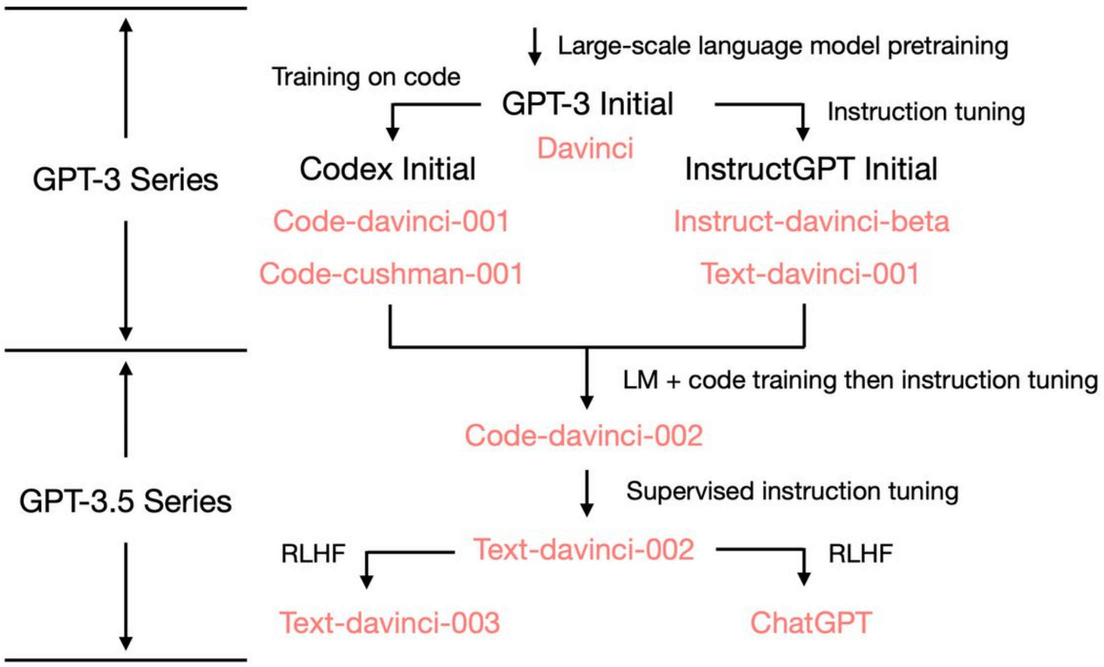

Yao Fu, et al.: How does GPT Obtain its Ability? Tracing Emergent Abilities of Language Models to their Sources, 12.2022

Brief summary on (LLM) GPT-3 hype

Even with 175B parameters and 450 gigabytes of input data, it's **not a reliable** interpreter of the world

Models presents **security** and **uncontrollability problems**, including false content and biased information during content generation

Unreliable doesn't mean useless. This is not to say that GPT-3-related is devoid of practical applications

Can be considered to use as **intelligent auxiliary tasks**

Cannot directly interface with the end-users

GPT-3 is a microcosm of the best and worst example in AI today (W. D. Heaven, 2021)

Min Zhang, Juntao Li: A commentary of GPT-3 in MIT Technology Review 2021

LLMs and LMMs demands

- Significant step in technological evolution
- Resource demands
  - LLMs are usually trained in **high performance infrastructures** with **multiple GPUs** for several months, e.g., Microsoft Azure infrastructure
  - The energy cost is **several hundreds thousands EUR/USD per month**
  - Each LLM/LMM: a "**piece**" made from long-time high energy consumption
- Unrepairable logical **inconsistencies** and inaccuracies in their outputs, which come from the design as the most probably output - not the correct output
- **Ethical concerns: explainability, content control, bias mitigation**

Responsible AI

Principles of Artificial Intelligence, Privacy, Fairness and Data Fusion

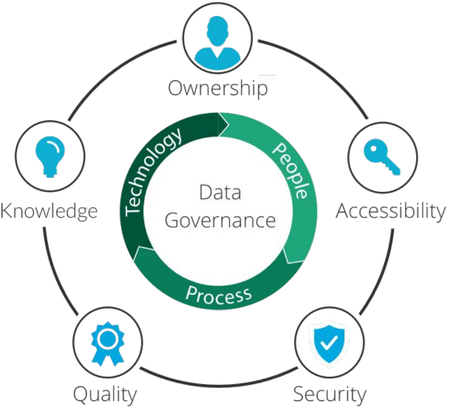

Responsible AI - the only way to mitigate AI risks

1. **Privacy**
2. **Security and Safety**
3. **Ethics**
4. **Fairness**
5. **Accountability**
6. **Transparency**

Responsible AI is a standard for ensuring that AI is safe, trustworthy and unbiased

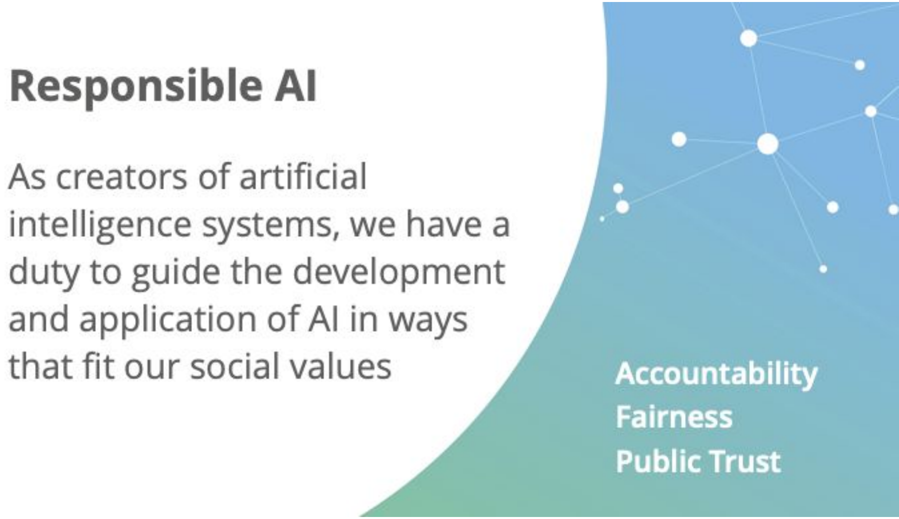

Trustworthy AI

- AI opens up virtually limitless benefit potential
- Potential risks from AI, by addressing any areas where controls or processes are found to be lacking or inadequate. The risks include
  - biased decision-making
  - the interpretability of AI decisions
  - a lack of explainability
- **Trustworthy AI is a methodology** for the implementation of AI methods in real organizations with
  - fairness
  - model explainability, and
  - accountability

  in its core

Arrieta et al.: Explainable Artificial Intelligence (XAI), Information Fusion, Elsevier ScienceDirect, 2020

The purpose of Artificial Intelligence

is to augment

Human Intelligence

(not replacement)

Human-Centric AI (HCAI)

“The **human-centric** approach to AI strives to ensure that **human values** are central to the way in which AI systems are developed, deployed, used and monitored, by ensuring respect for **fundamental rights**, including those set out in the Treaties of the European Union and Charter of Fundamental Rights of the European Union, all of which are united by reference to a common foundation rooted in respect for **human dignity**, in which the human being enjoys a unique and inalienable **moral** status.

This also entails consideration of the natural environment and of other living beings that are part of the **human ecosystem**, as well as a **sustainable** approach enabling the flourishing of **future generations** to come.”

EU guidelines on ethics in artificial intelligence: Context and implementation

GDPR (2016, EU)

General Data Protection Regulation (GDPR)

DSA (2022, EU)

Regulation on Digital Services Act (DSA)

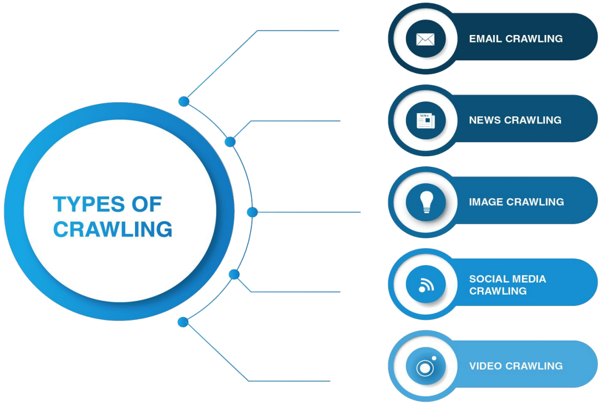

Privacy - It is all about personal data - Personal Identifiable Information (PII)

Privacy-Preserving Data Mining (PPDM)

Data Transformations

- Pseudonymization
- Anonymization: *k*-anonymity, *l*-diversity
- Differential Privacy

Secure Computations

- Encryption: symmetric, asymmetric, hashing
- Homomorphic Encryption: fully, partially

Trusted Environments

- Specialized Hardware
- Secure Execution Environment

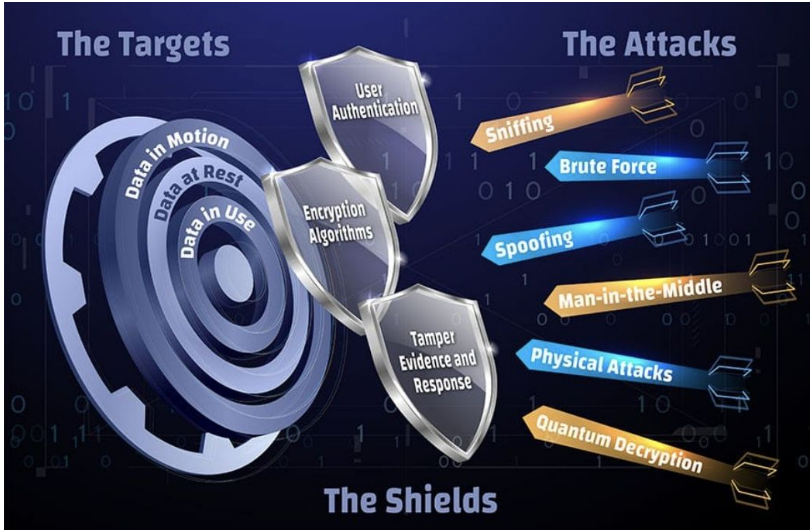

Data Protection and Federated Learning

Centralized learning

- Problems: access right, owner right

Distributed learning

- Problems: sharing agreements, retention and disposal, minimization

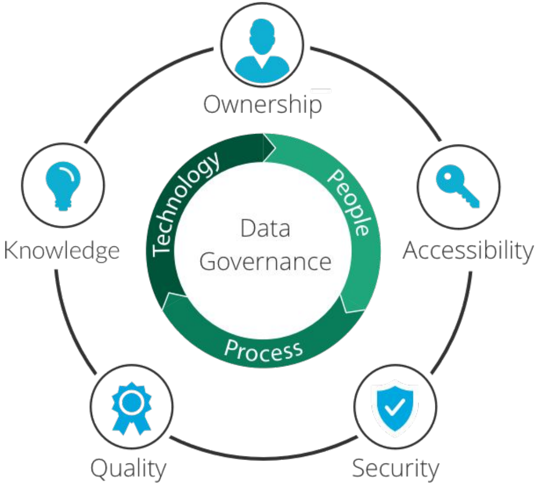

Federated Learning

- Multiple clients train a model under central server orchestration, while keeping data decentralized

- Sharing and aggregating the model updates as transferred artifact

Responsible AI

Responsible development of human-centric and trustworthy AI systems

Thank you for your attention!

AI Data and Model Landscape

Giang Nguyen

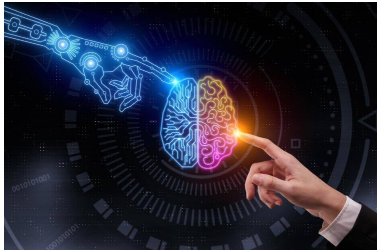

AI4 eosc

Co-funded by the European Union