AI: Our Current Reality and Future Trajectory

presented to The Nigerian Society of Engineers,
Abeokuta Branch

by Timilehin Owolabi, 400 level,
Depatment of Electrical and Electronics Engineering, Covenant University

Introduction

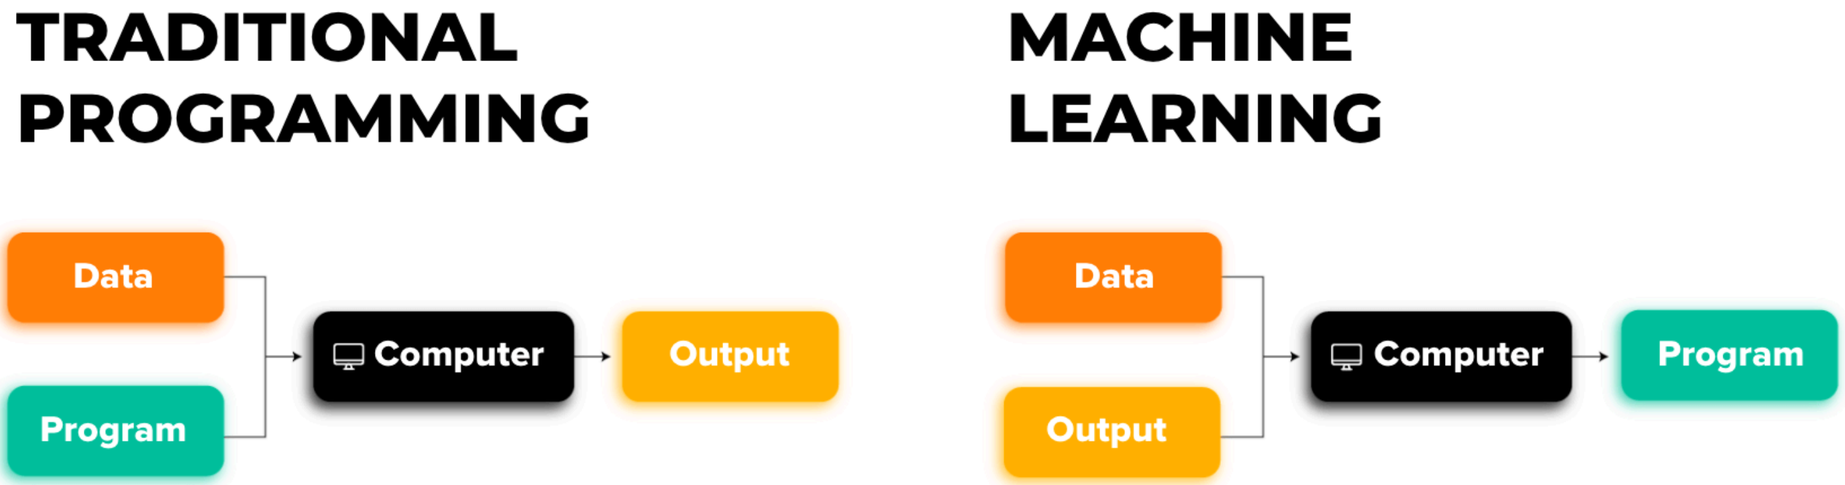

History of AI

- The Birth of AI (1957) - Frank Rosenblatt, perceptron
- The First Golden Years (1960s-1970s) - Alan Turing, chatbots
- The First AI Winter (1970s-1980s) - funding, optimism down
- The Second Golden Years (1980s- 1987s) - Geoffery Hinton
- The Second AI Winter (1987-1993s) - slow compute, lack of data
- Revival and Growth (1990s-early 2000s) - Yann LeCun, backprop
- Rise of Modern AI (2000s-2010s)
- Current Boom (2010s-present)

Important Milestones in AI

- 1997: IBM's Deep Blue defeats world chess champion Garry Kasparov.
- 2011: IBM's Watson wins Jeopardy!, showcasing advanced natural language processing and knowledge retrieval.
- 2012: The ImageNet competition is won by a deep learning model (AlexNet), marking a significant leap in computer vision.
- 2014: The Turing test is arguably passed for the first time by a chatbot named Eugene Goostman.
- 2016: Google DeepMind's AlphaGo defeats world champion Go player Lee Sedol, demonstrating AI's ability to excel at complex strategic games.

Important Milestones in AI

- 2018: Google demonstrates Google Duplex, an AI system capable of making phone calls and conducting natural conversations.
- 2020: OpenAI releases GPT-3, a large language model capable of generating human-like text and performing various language tasks.
- 2022: DeepMind's AlphaFold solves the protein folding problem, a major breakthrough in biological science.
- 2022-2023: The release of ChatGPT and other large language models brings conversational AI to the mainstream, sparking widespread public interest and debate.
- 2024: SORA, GPT4

ML and DL

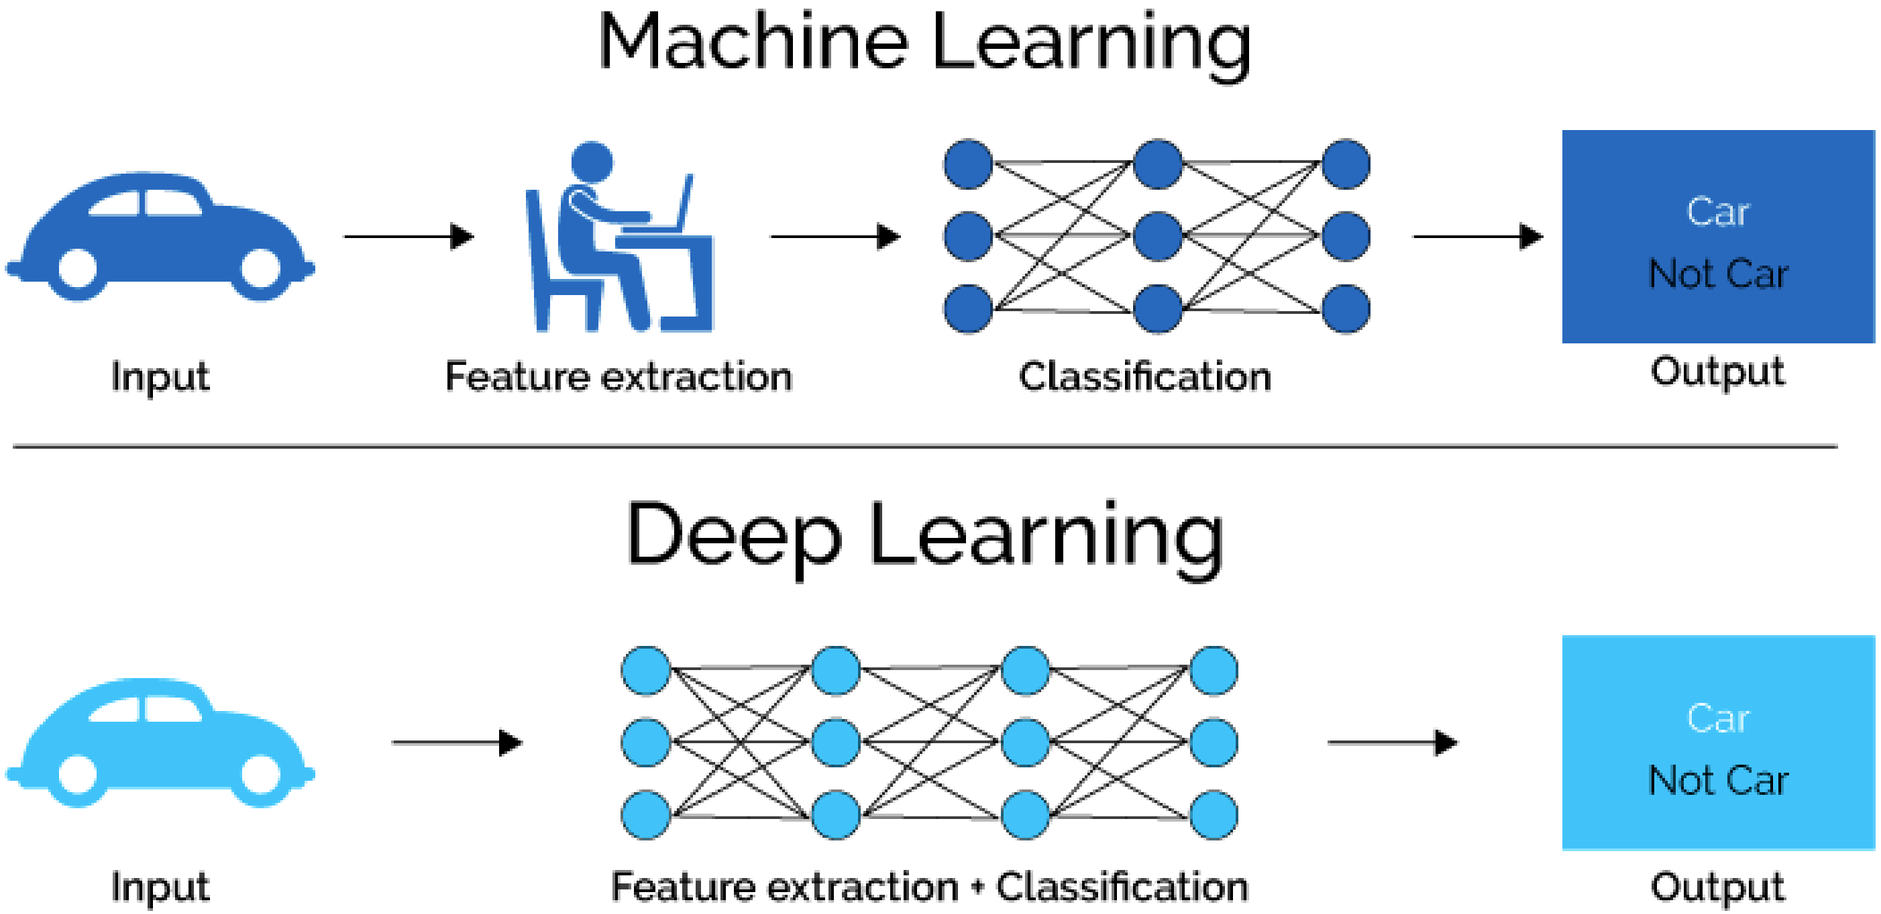
Machine Learning
Input
Feature extraction
Classification
Car
Not Car
Output
Deep Learning
Input
Feature extraction + Classification
Car
Not Car
Output

ML algorithms

Linear and Logistic Regression
K-Nearest Neighbour
Support Vector Machine
Decision Tree
Random Forest
XGBoost
Naive Bayes

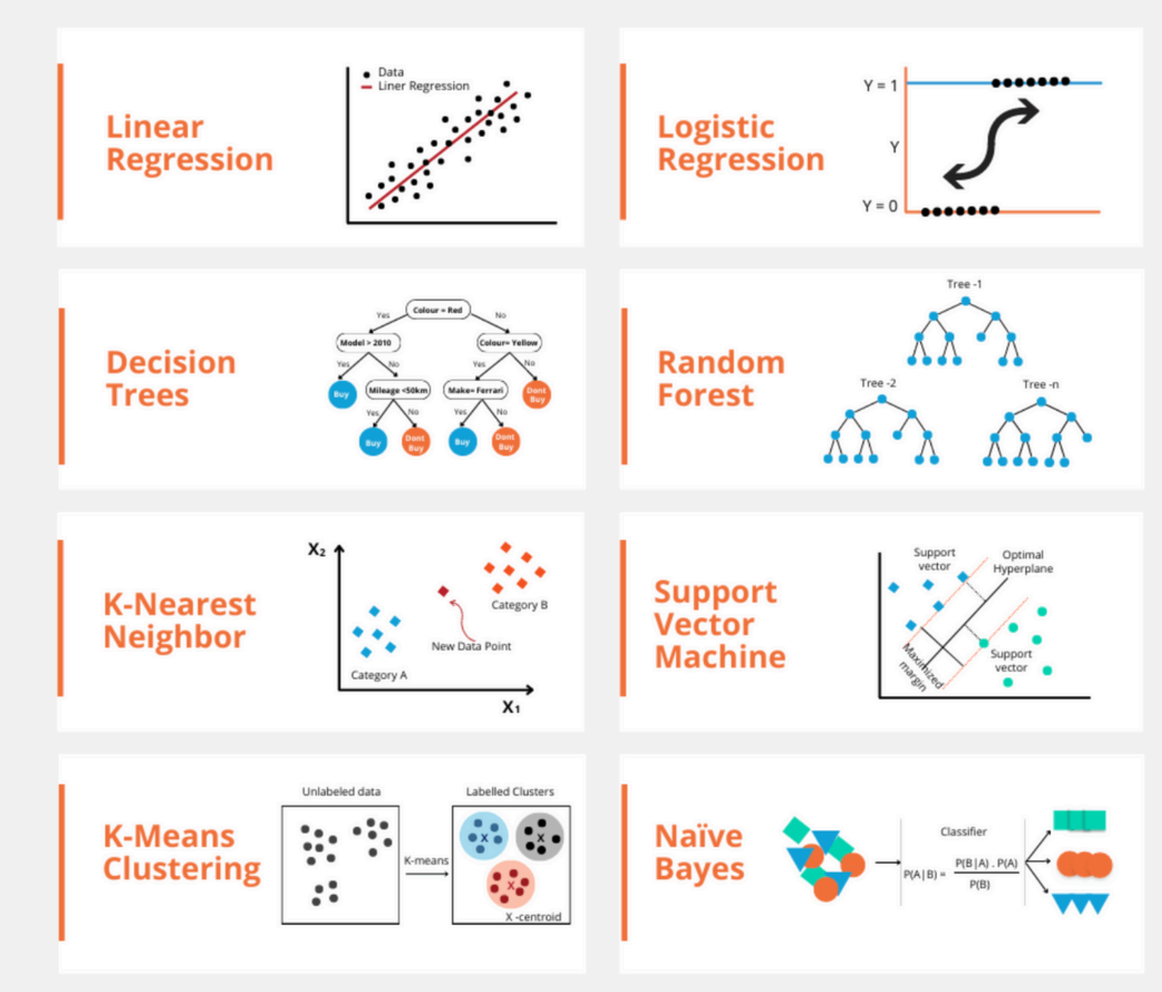

DL Algorithms

- Perceptron
- MLP
- CNN - AlexNet, EfficientNet, YOLO
- RNN - GRU, LSTM
- GNN - GAT, GCN
- Transformers
- GANs, VAE
- Diffusers - Stable Diffusion, Dall-E2

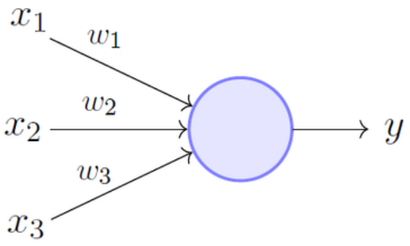

Perceptron Model (Minsky-Papert in 1969

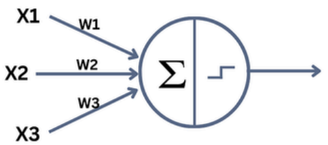

Single-layer perceptron

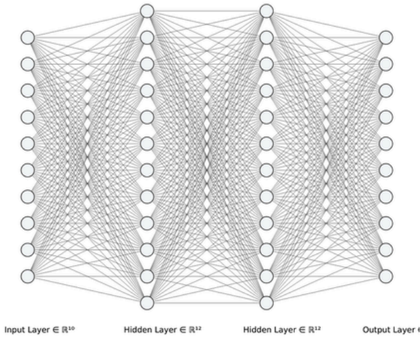

Multi-layer perceptron

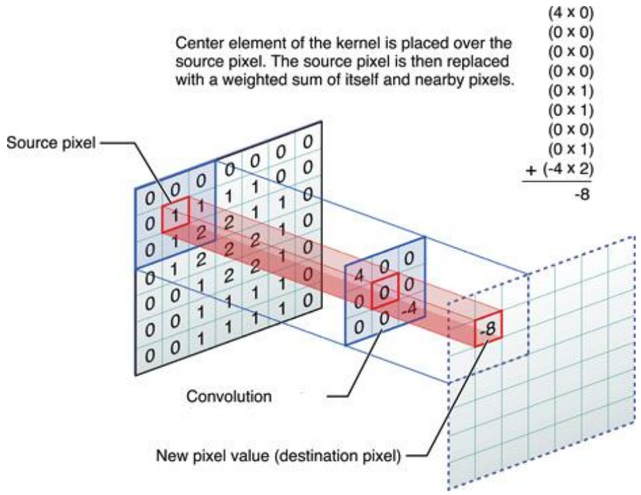

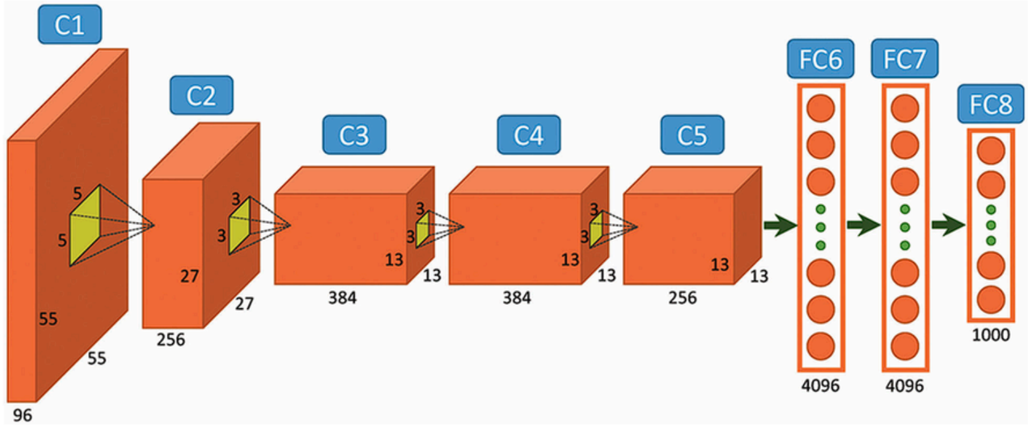

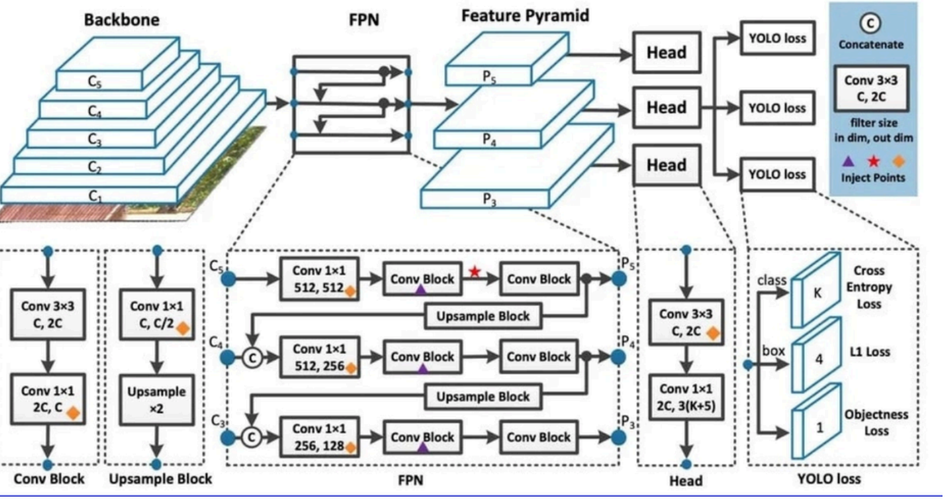

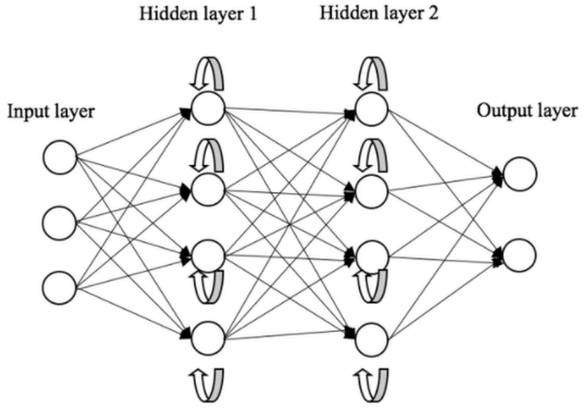

General Form of RNNs

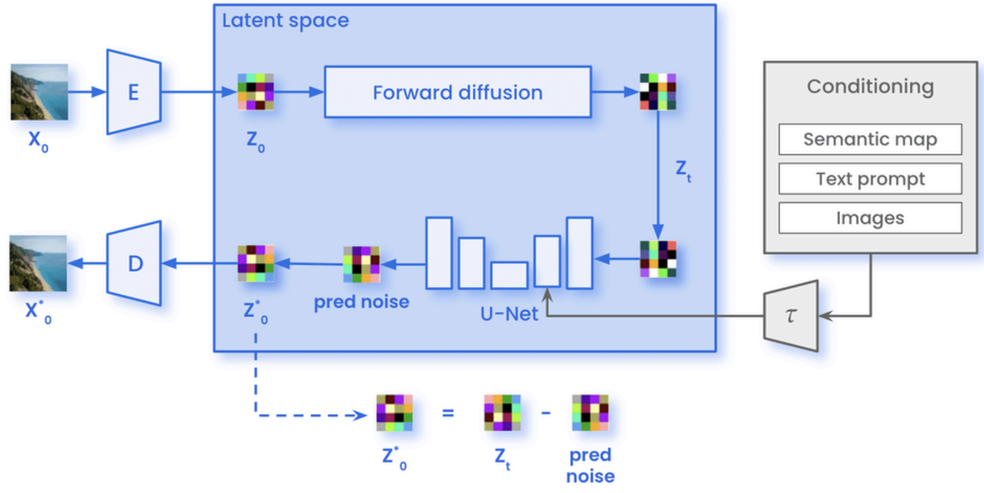

Transformer Based Networks

How everything started....

Attention is All You Need - Vaswani et al., 2017

Encoder-only - BERT, DistilBERT

Decoder-only - LLMs, GPT-3, LLaMA, Mistral

Vision Transformers - CLIP, DETR, SAM2

Vision Language Models - Llava, Mistral, Phi

Embodied Multimodal Transformers - PaLM-E

Mixture of Experts - Mixtral 8x7B

Retrieval-Augmented Transformers - REALM, RAG

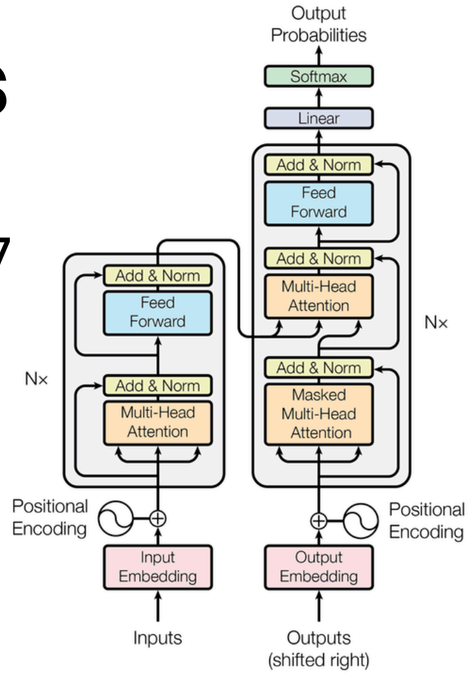

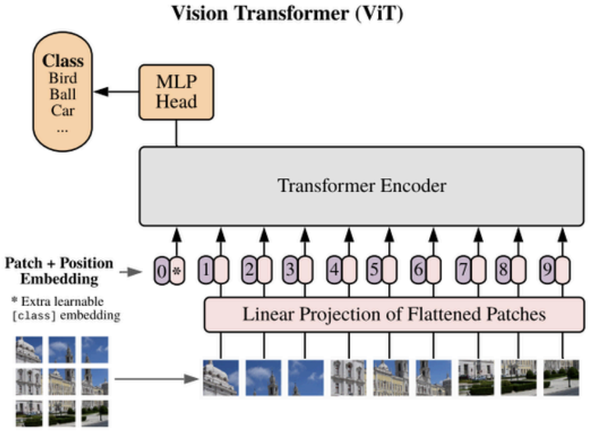

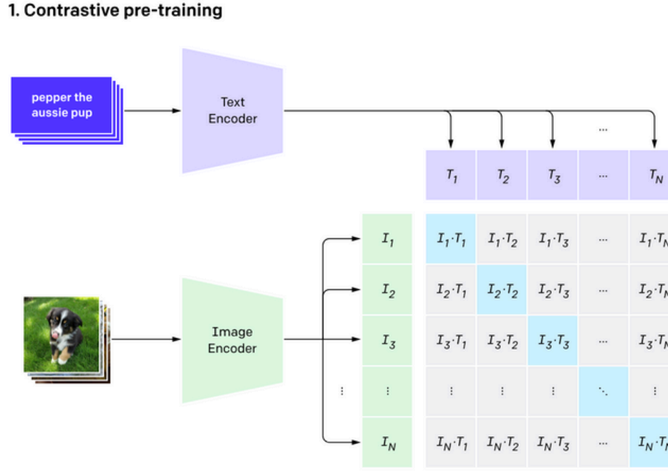

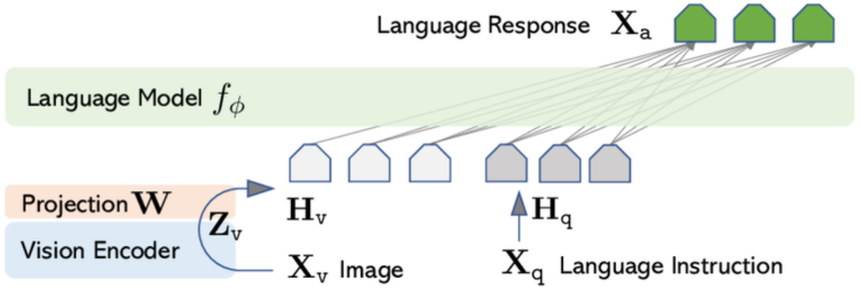

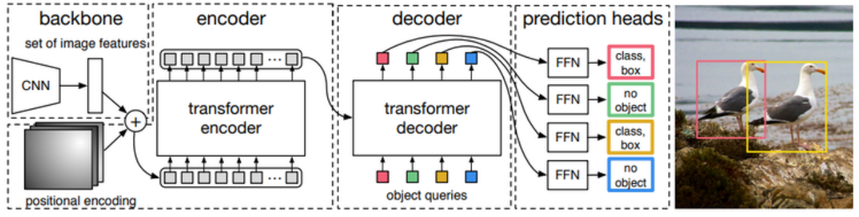

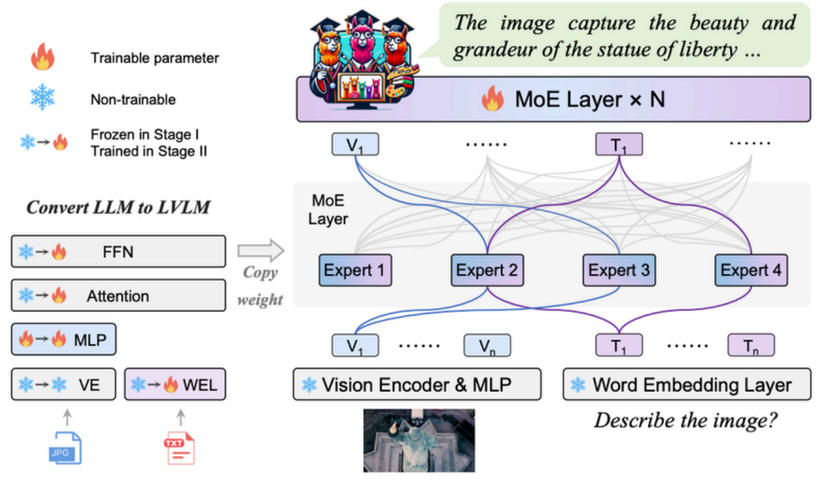

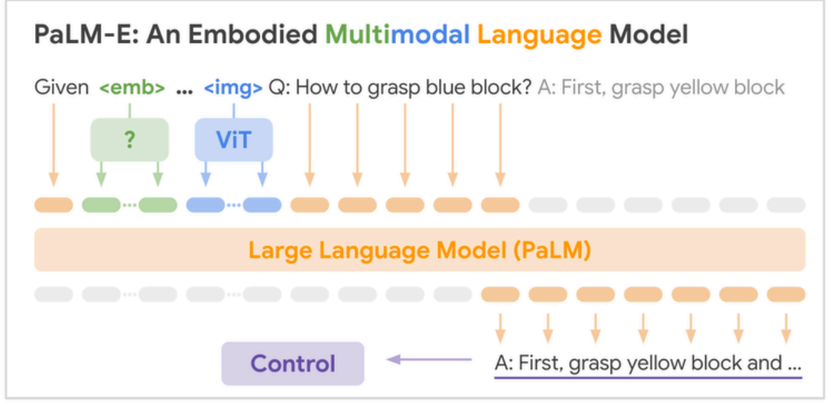

Current State of AI

- Healthcare
- Transportation
- Agriculture
- Finance
- Education
- Customer Service
- Entertainment and Arts
- Manufacturing
- Cybersecurity
- Climate
- Robotics

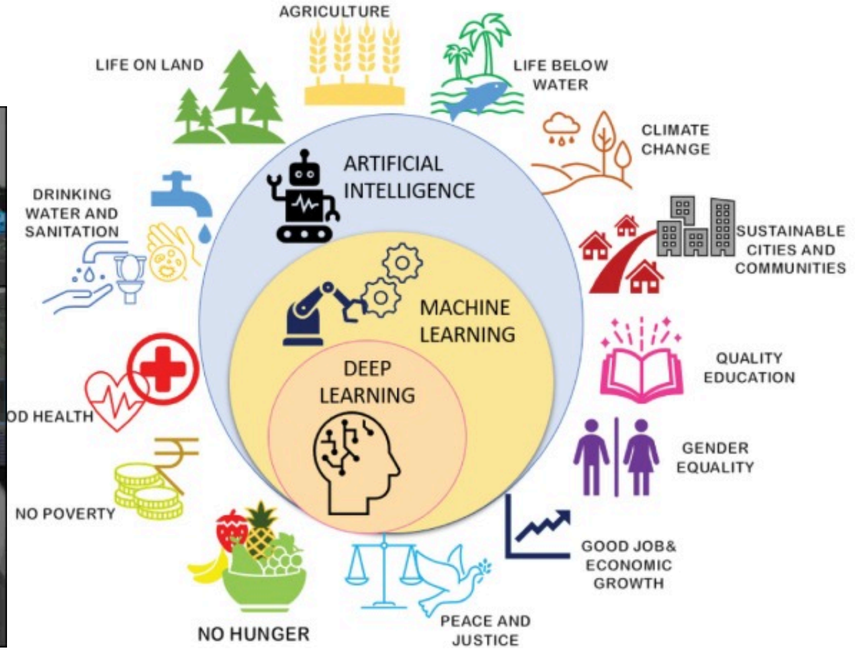

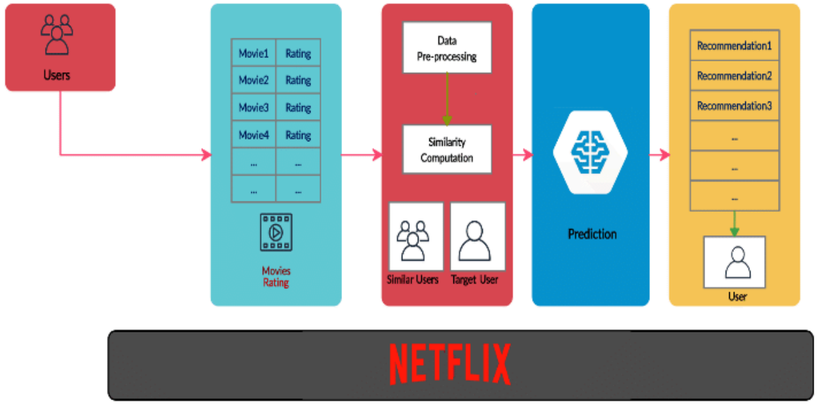

Nvidia $3 trillion

Large Language Models


AUTONOMOUS
ROBOT

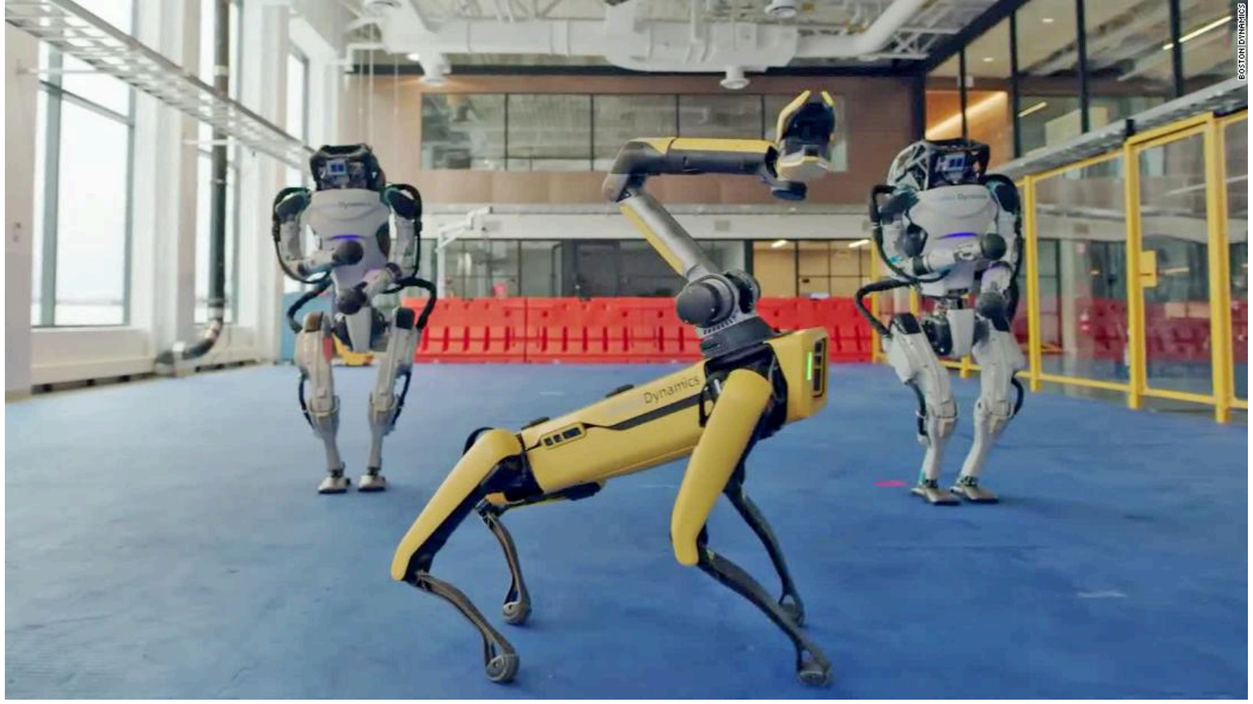
Dynamics
BOSTON DYNAMICS

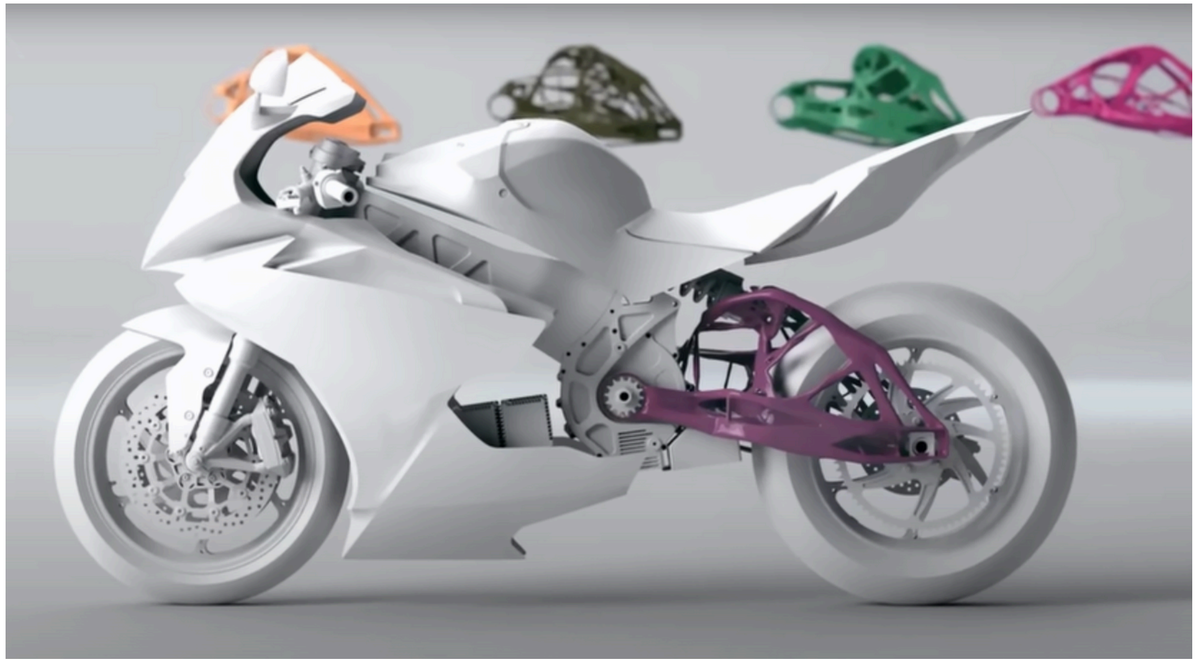

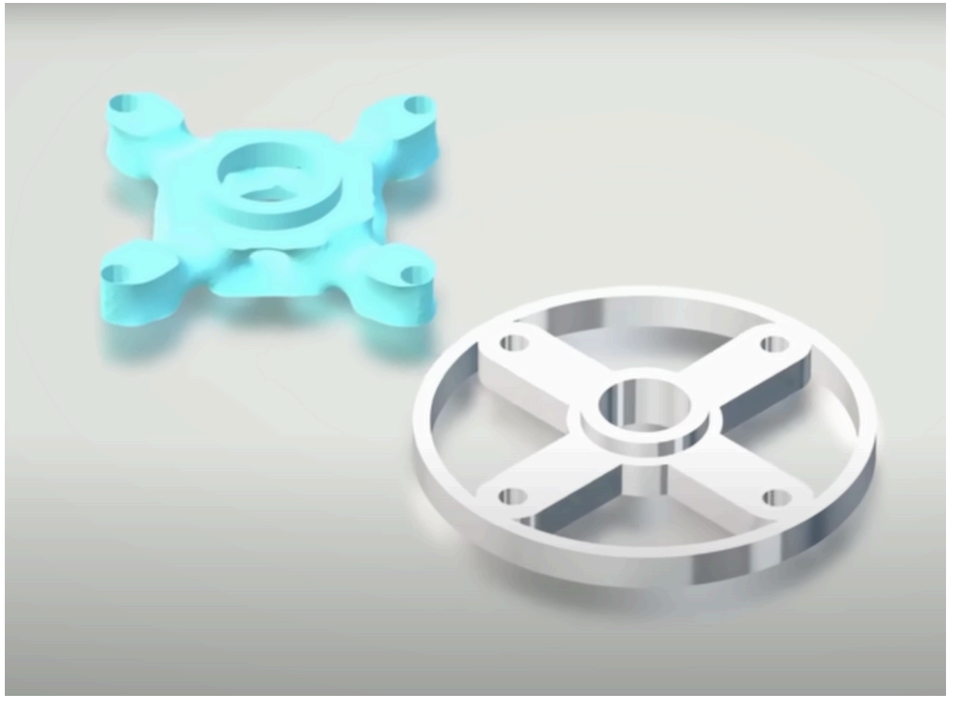

LLM Hype

Why LLMs?

Why Not?

Capabilities - translation, summary, content, coding

reasoning, explainable

Limitations - generalization, hallucination

Large Vision Models - scene understanding

Transformer scaling laws

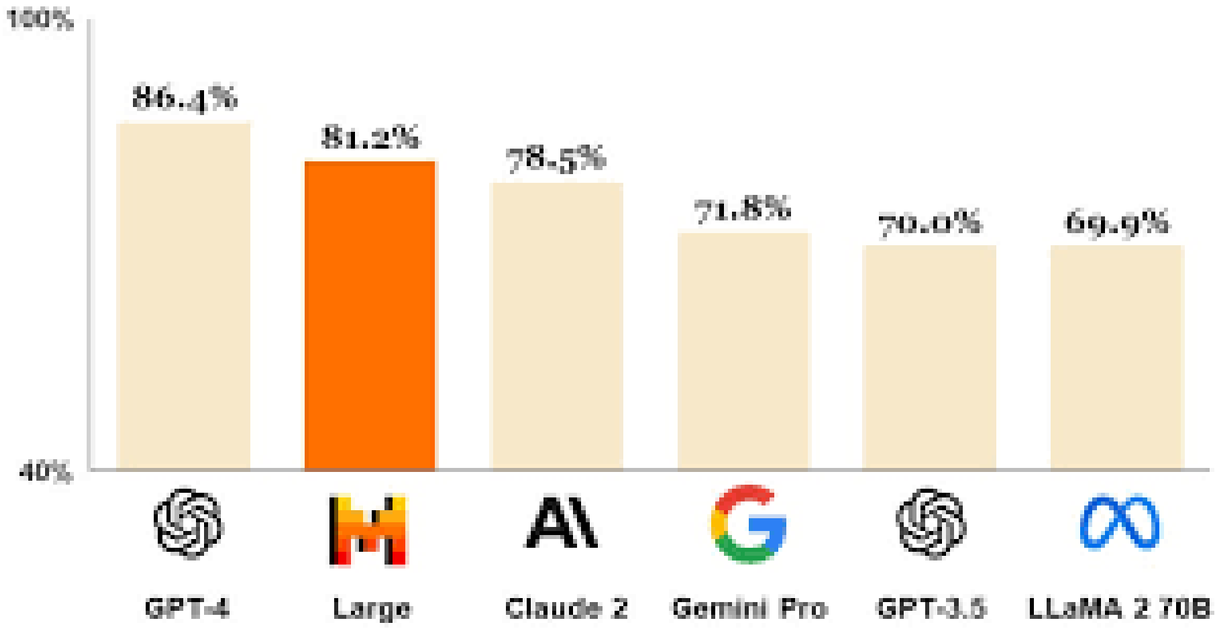

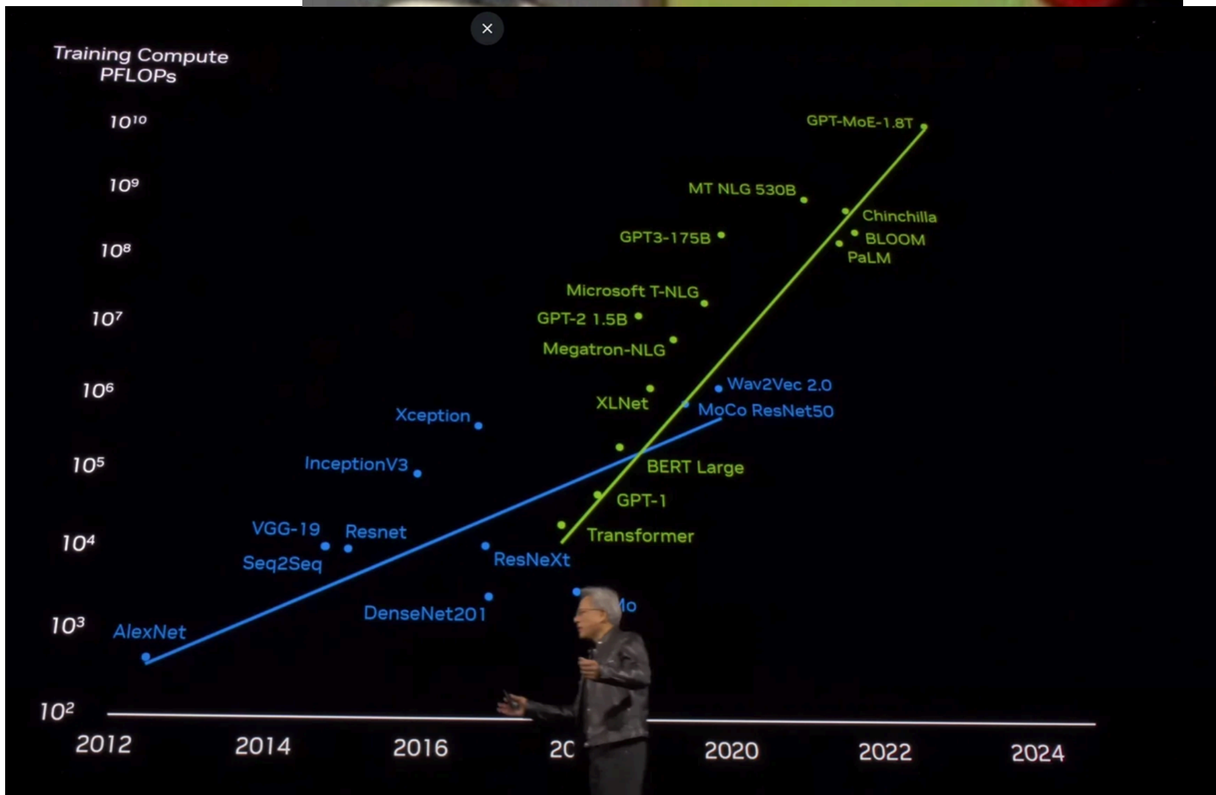

A DAY IN DATA

The exponential growth of data is undisputed, but the numbers behind this explosion – fuelled by internet of things and the use of connected devices – are hard to comprehend, particularly when looked at in the context of one day

500m

tweets are sent every day

Twitter

4PB

of data created by Facebook, including

350m photos

100m hours of video watch time

Facebook Research

294bn

billion emails are sent

Radicati Group

320bn

emails to be sent each day by 2021

306bn

emails to be sent each day by 2020

65bn

messages sent over WhatsApp and two billion minutes of voice and video calls made

Facebook

3.9bn

people use emails

4TB

of data produced by a connected car

Intel

DEMYSTIFIYING DATA UNITS

From the more familiar 'bit' or 'megabyte', larger units of measurement are more frequently being used to explain the masses of data

Unit		Value	Size
b	bit	0 or 1	1/8 of a byte
B	byte	8 bits	1 byte
KB	kilobyte	1,000 bytes	1,000 bytes
MB	megabyte	$1,000^2$ bytes	1,000,000 bytes
GB	gigabyte	$1,000^3$ bytes	1,000,000,000 bytes
TB	terabyte	$1,000^4$ bytes	1,000,000,000,000 bytes
PB	petabyte	$1,000^5$ bytes	1,000,000,000,000,000 bytes
EB	exabyte	$1,000^6$ bytes	1,000,000,000,000,000,000 bytes
ZB	zettabyte	$1,000^7$ bytes	1,000,000,000,000,000,000,000 bytes
YB	yottabyte	$1,000^8$ bytes	1,000,000,000,000,000,000,000,000 bytes

*A lowercase "b" is used as an abbreviation for bits, while an uppercase "B" represents bytes.

ACCUMULATED DIGITAL UNIVERSE OF DATA

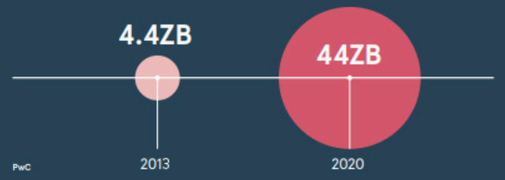

PwC

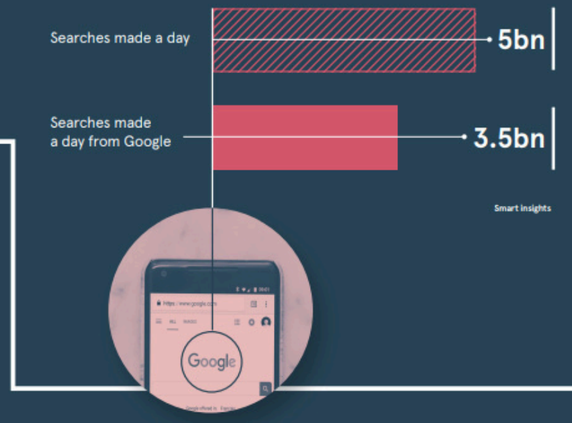

Smart Insights

463EB

of data will be created every day by 2025

IDC

95m

photos and videos are shared on Instagram

Instagram Business

28PB

to be generated from wearable devices by 2020

Statista

RACONTEUR

Philosophy, Ethics, and Explainability

Ethical AI - data
Explainable AI
Accountability and Liability
Is AI sentient or just maths?

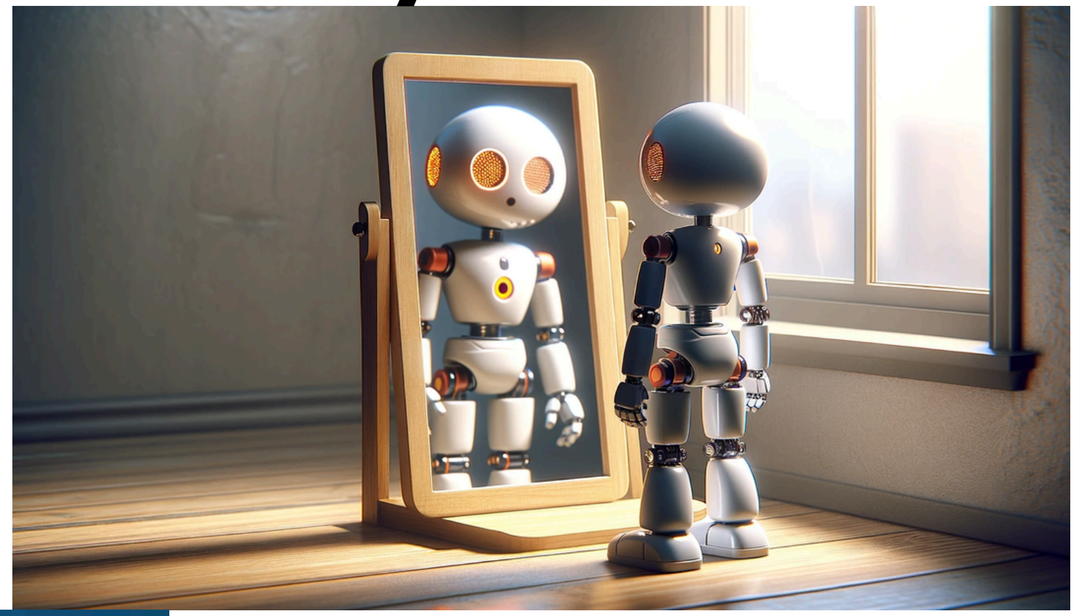

Future of AI

- Quantum AI - quantum computing algorithms, optimization
- Space Exploration - autonomous spacecraft navigation, extraterrestrial mining, planetary exploration rovers
- Human-Robot Collaboration - collaborative robots (cobots), assistive robots, human-robot interaction
- AI Enhanced Humans - brain-computer interfaces, prosthetics, cognitive enhancement
- AGI (Artificial General Intelligence) - theories of general intelligence, long-term implications, safety and alignment

Scientists and Engineers that made all these possible

Geoffrey Hinton - Backpropagation, "Deep Learning" book
Yann LeCun - LeNet CNN architecture
Andrew Ng - Google Brain, deep learning courses
Ian Goodfellow - GANs, "Deep Learning" book
Demis Hassabis - AlphaGo, DeepMind
Ilya Sutskever - Transformer models, OpenAI
Andrej Karpathy - CS231n course, image recognition
Jürgen Schmidhuber - LSTM networks
Richard Socher - Dynamic Memory Networks
Ashish Vaswani - Attention mechanism
Jascha Sohl-Dickstein - Diffusion models

And many more.....



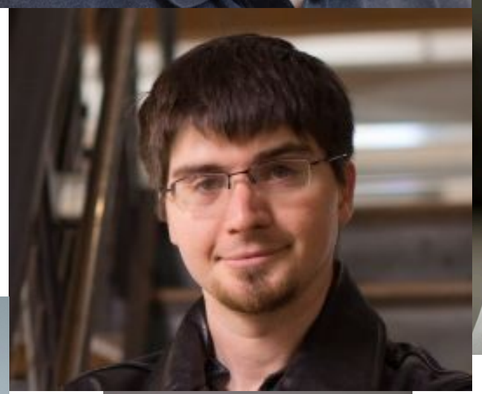


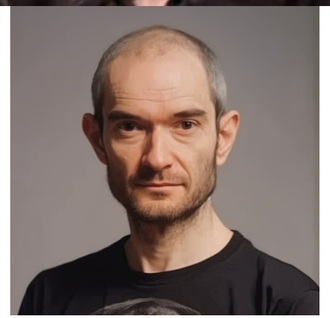
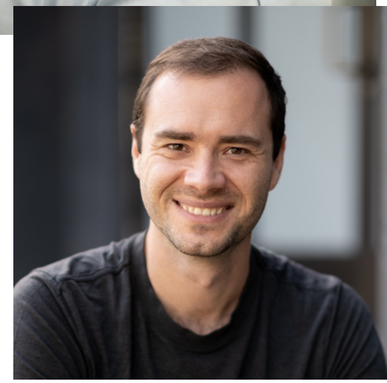

Resources

Playground: playground.tensorflow.org

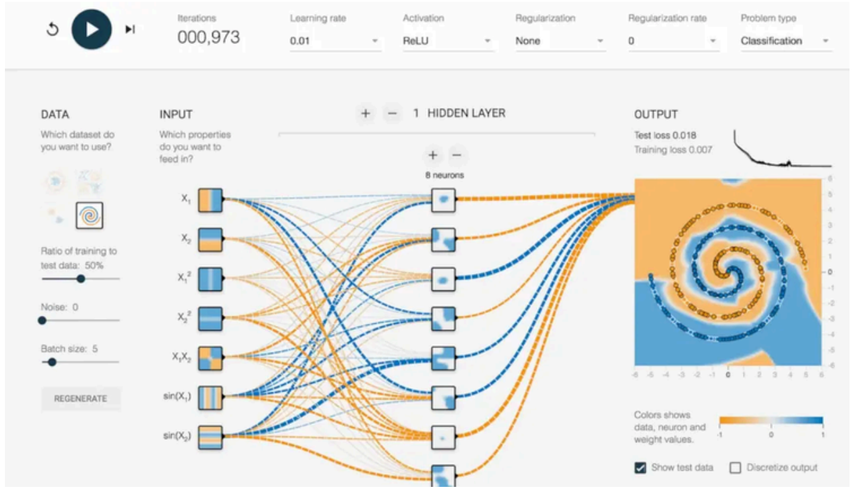

Understanding Deep Learning by Simon J.D. Prince: https://mitpress.mit.edu/9780262048644/understanding-deep-learning/

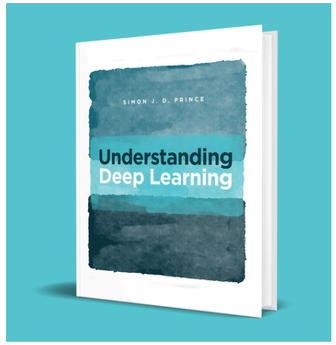

Short AI course by Pytorch Lightning: https://www.youtube.com/watch?v=6Py-tIEiXKw&list=PLaMu-SDt_RB4Ly0xb0qsQVpLwRQcjtOb-&pp=iAQB