GE1601 Whole-Person Development

Chen Liu

AI Unplugged: Journey Through the Evolution and Future of Intelligence

Semester A, 2025-2026

What is Artificial Intelligence?

- Artificial intelligence (AI) is intelligence demonstrated by machines, as opposed to the intelligence of humans and other animals. (Wikipedia)
- AI is *an attempt of* reproduction of human reasoning and intelligent behaviour by computational methods. AI aims to make computers:
  - Act like humans.
  - Think like humans.

Outline

Birth of AI (1950s)

Alan Turing proposes Turing test in 1950.

- **Turing Test**: a machine is said to pass the Turing test if it can convince a human judge that it's actually a human through natural language dialogue (indistinguishability).
- Two approaches trying to pass the Turing test: 1) rule-based symbolic AI, 2) Training-based neural and statistical AI.

In 1956, John McCarthy organized a workshop at Dartmouth College.

- Term "artificial intelligence" was coined.
- The participants laid out a bold proposal: to build a system that could capture every aspect of intelligence, which is called **strong AI** or **Artificial General Intelligence**.
- In contrast to strong AI, **weak AI** is any programme that is designed to solve exactly one problem.

Early Success (1952 - 1959)

Claims: computers can do X (weak AI)

- Checkers (1952): Arthur Samuel's programme played checkers games at the strong amateur level. (Wikipedia)
- Logic Theorist: Allen Newell, Herbert A. Simon, and Cliff Shaw's computer programme deliberately engineered to perform automated reasoning. (Wikipedia)
- LISP (1958): a favoured programming language for AI.

Figure: Arthur Samuel (left), Allen Newell (middle) and Herbert A. Simon (right).

Overwhelming Optimism and First AI Winter (1960s)

- Herbert Simon: Machines will be capable, within twenty years, of doing any work a man can do.

Overwhelming Optimism and First AI Winter (1960s)

- Herbert Simon: Machines will be capable, within twenty years, of doing any work a man can do.

However:

The spirit is willing but the flesh is weak.

(Russian)

The vodka is good but the meat is rotten.

- The report by Automatic Language Processing Advisory Committee (ALPAC) in 1966 resulted in a government funding cut for MT, causing the first AI winter.

Knowledge-base System (1970s - 1980s)

- Expert systems: elicit specific domain knowledge from experts in form of rules: **If [premises] then [conclusion]**
- Example: in autonomous driving
  - If [Case 1: a pedestrian], then use [Rule 1: stop];
  - If [Case 2: an obstacle], then use [Rule 2: get around];
  - ...

AI Becomes an Industry (1980 - 1988)

- R1: first successful commercial expert system, configured computer systems at Digital Equipment Corporation (DEC) and saved 40 million dollars per year. (paper published in 1982: A rule-based configurer of computer systems)
- AI industry boosts: few million dollars in 1980 $\rightarrow$ about two billion dollars in 1988.
- 1987: collapse of LISP machines, second AI winter.
- End of symbolic AI domination for multiple decades.

Return of ANNs (1986 -)

- Disillusionment with expert systems led to the resurgence of ANNs in mid 1980s.
- The history of neural AI dates back to 1943.

Inspiration from Neuroscience

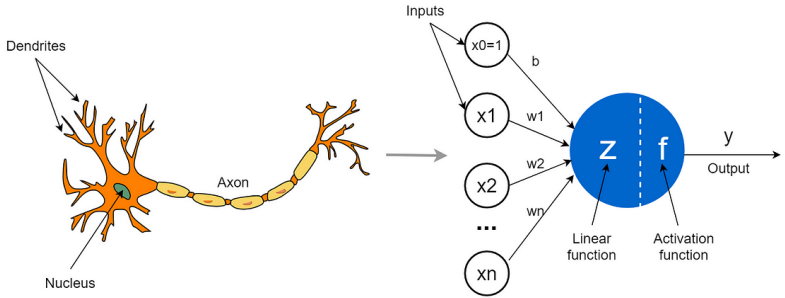

Figure: Human neurons (left) and artificial neurons in ANN (right). Image from towardsdatascience.com.

Inspiration from Neuroscience

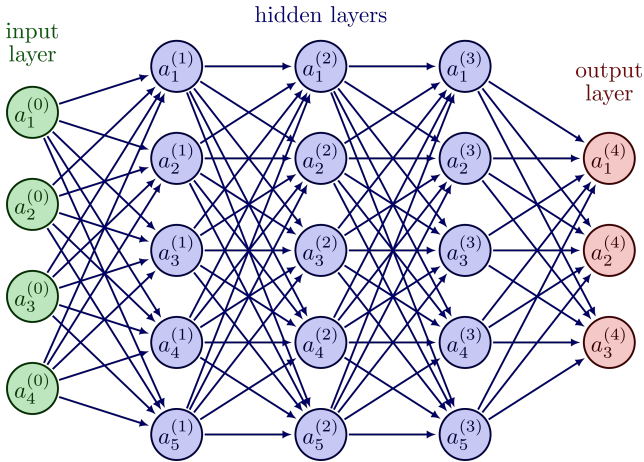

Early Research on ANNs (1943 - 1969)

- Warren McCulloch & Walter Pitts (1943): **Artificial Neural Network (ANN)** with on-off neurons can represent computable function.
- Donald O. Hebb (1949): "cells that fire together wire together".
- Frank Rosenblatt (1958): Perceptron algorithm for linear classifiers.
- Bernard Widrow & Ted Hoff (1959): adaptive linear neuron (ADALINE).
- Marvin Minsky & Seymour Papert (1969): Percepron book showed that linear models could not solve XOR, killed neural nets research.

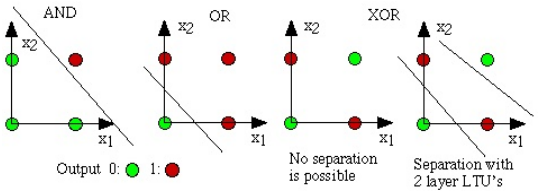

Nonlinear Activation Functions

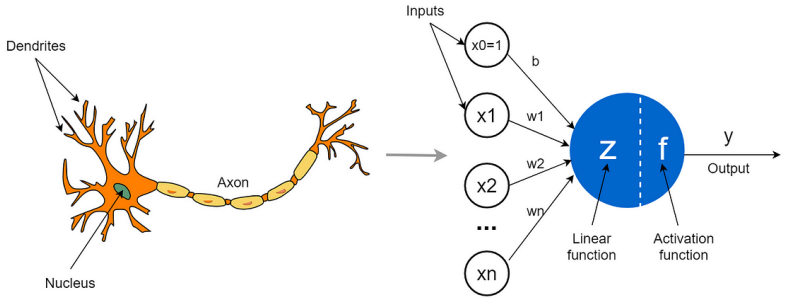

Figure: Human neurons (left) and artificial neurons in ANN (right). Image from towardsdatascience.com.

Revival of Connectionism (1980 -)

- Kunihiko Fukushima (1980): Neocognitron, a.k.a., convolutional neural networks for images.
- David E. Rumelhart, Geoffrey E. Hinton & Ronald J. Williams (1986): backpropagation technique for training multi-layer neural networks.
- Yann LeCun (1989): applied convolutional neural networks to recognize handwritten digits for USPS.

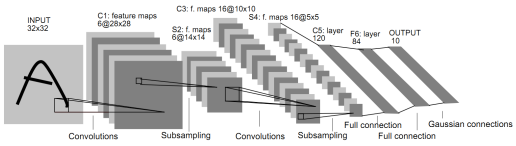

Fig. 2. Architecture of LeNet-5, a Convolutional Neural Network, here for digits recognition. Each plane is a feature map, i.e. a set of units whose weights are constrained to be identical.

Figure: Convolutional neural network. Image from gamedevacademy.org

The era of Deep Learning (2000 -)

- Geoffrey E. Hinton et. al. (2006): unsupervised layerwise pre-training of deep neural networks.
- Alex Krizhevsky, Ilya Sutskever & Geoffrey E. Hinton (2012): AlexNet obtained huge gains in object recognition and transformed computer vision community overnight. [1]

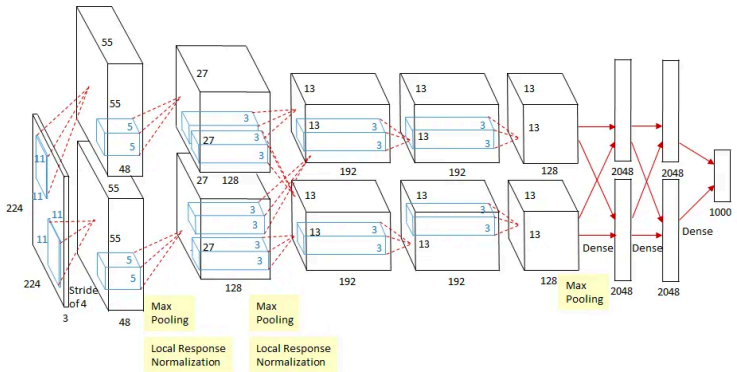

[1] ImageNet Classification with Deep Convolutional Neural Networks. 2012.

The era of Deep Learning (2000 -)

- Big data: large amount of training data needed.
- Large model: millions or billions of parameters with complex architectures.
- Complex reasoning: the gap between current technique and the goal of AGI narrows down.
- High-performance hardware: distributed computing, GPU e.t.c.

Large Language Models (2020s -)

- A language model is a probabilistic model of a natural language, which models the probability of a sequence being a correct output given its context: $P(w_1, w_2, ..., w_n) = \prod_{i=1}^{n} P(w_i | w_1, w_2, ..., w_{n-1})$.
- A large language model (LLM) is a large neural network based on the transformer architecture, compared with traditional language models, an LLM demonstrate the impressive capacity of knowledge extraction and inference, enabling it to complete various tasks.

Large Language Models (2020s -)

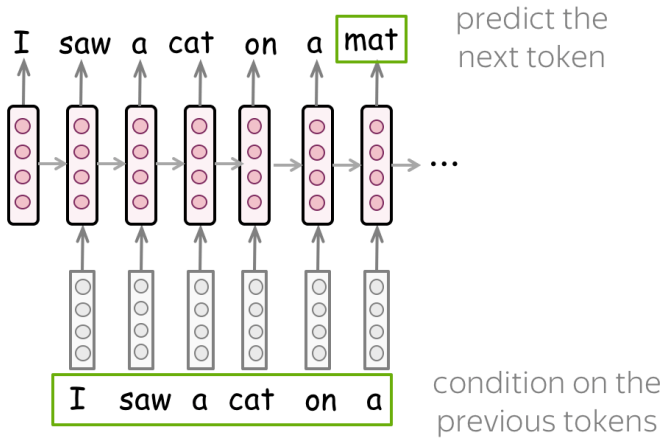

Figure: A traditional neural language model.

Large Language Models (2020s -)

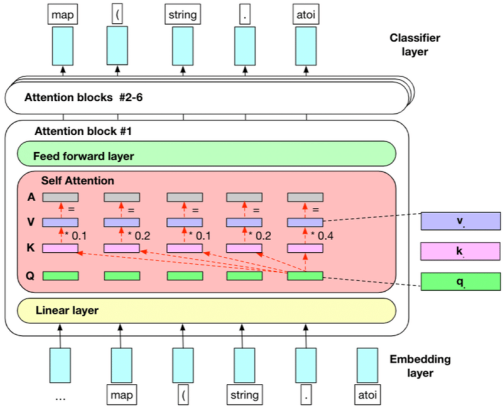

Figure: A high-level architecture of a large language model.

Grokking

- Grokking refers to a sharp rise of the network's generalization capacity.

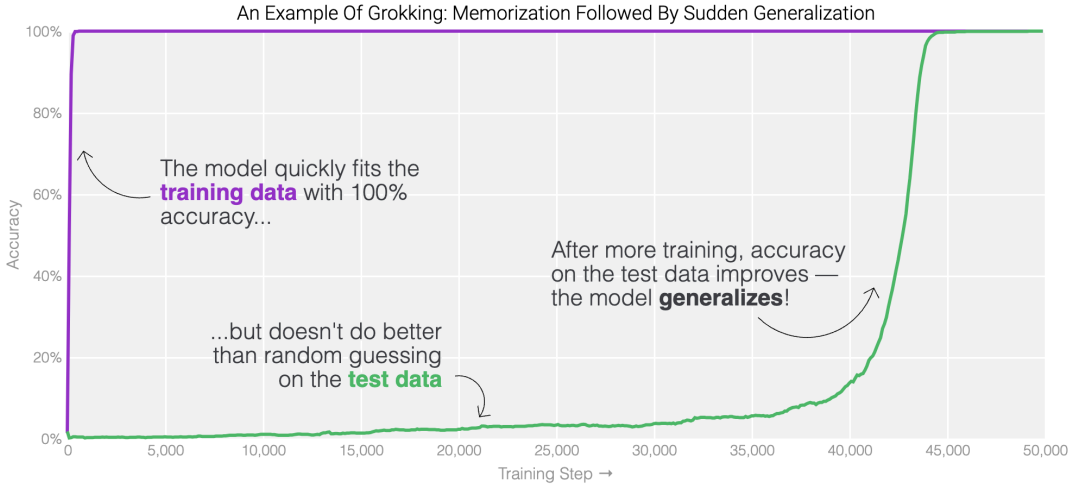

Scaling Law

- GPT-1: 117 million parametrs trained on 4.5GB of text. Training cost: 1 month on 8 GPUs, i.e., 1×10^{19} FLOPs.
- GPT-2: 1.5 billion parameters trained on 40GB of text. Training cost: 1.5×10^{21} FLOPs.
- GPT-3: 175 billion parameters trained on 499 billion tokens. Training cost: 3.1×10^{23} FLOPs.
- GPT-4: The architecture and the data is undisclosed. Training cost: 2.1×10^{25} FLOPs.
- GPT-5: ...

The Scope of AI

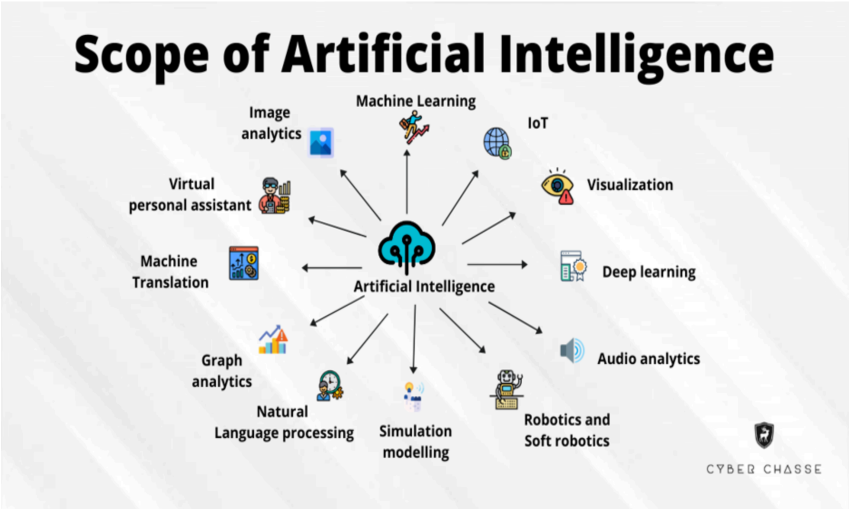

Outline

AI in Healthcare

- AI in healthcare is applied to clinical practice such as disease diagnosis, treatment protocol development, drug design, personalized medicine, patient monitoring and prognosis.

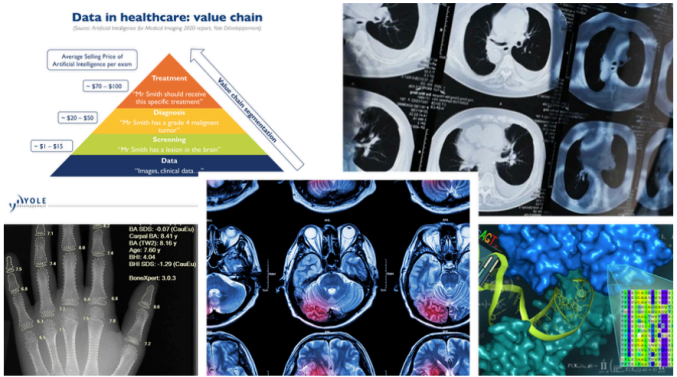

AI in Healthcare

- AI in healthcare is applied to clinical practice such as disease diagnosis, treatment protocol development, drug design, personalized medicine, patient monitoring and prognosis.

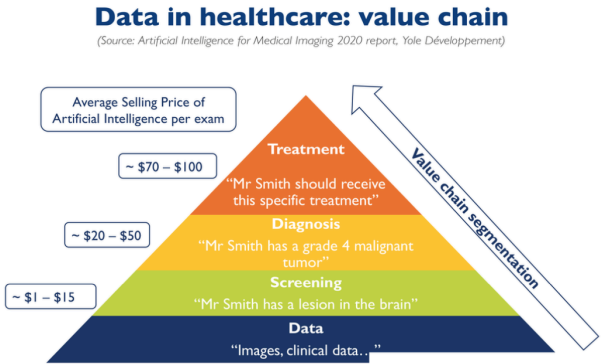

AI in Autonomous Driving

- Autonomous driving is one of the key application areas of AI, as modern AI can help autonomous vehicles understand the surroundings and perform path planning.

- A fully autonomous driving journey.

AI in Manufacturing

- Robot learning is a research field at the intersection of AI and robotics. It studies techniques allowing a robot to acquire novel skills or adapt to its environment through learning algorithms.

- There are some cute robots.

AI in Science

- AI algorithms have helped the scientific research of biology, chemistry, physics, finance and so on.

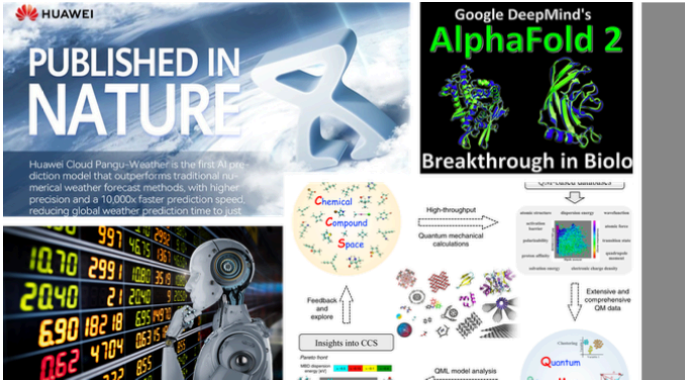

AI in Entertainment

- AI is used to generate responsive, adaptive or intelligent behaviors primarily in non-player characters (NPCs) similar to human-like intelligence in gaming.

AI in Arts

- AI in art refers to any artwork created with the assistance of AI. It includes works created autonomously by AI systems and works from a collaboration between human and AI.

- The artist in the machine.

AI in Military

- What do future robot soldiers look like?

Outline

Hardware for AI

- Central Processing Units (CPUs)
- Graphics Processing Units (GPUs)
  - AI has seen a dramatic transformation as a result of graphics processing units (GPUs). They are perfect for AI jobs that require handling massive quantities of data and intricate mathematical operations because of their parallel design, which enables them to run several computations at once.
- Tensor Processing Units (TPUs)
  - For the purpose of accelerating and optimizing machine learning workloads, Google has created Tensor Processing Units (TPUs). They are made to handle both inference and training procedures well and perform well when used with neural network tasks.

Hardware for AI

- The popularity of neural networks have made GPUs / TPUs crucial for large-scale experiments.

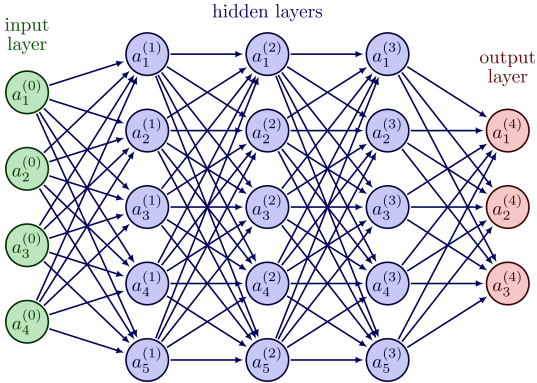

Hardware for AI

- GPUs can be more than $100\times$ faster than CPUs.

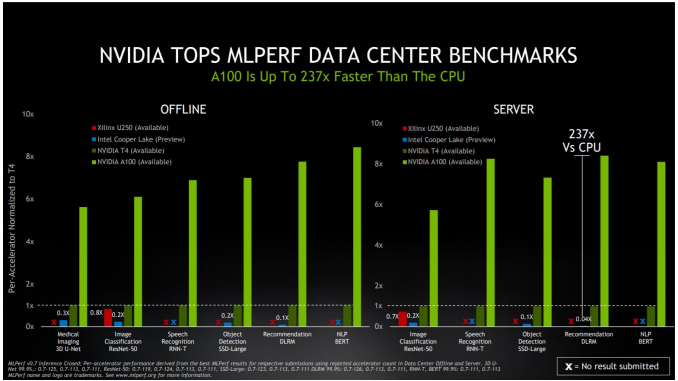

Lack of Data and Computing Power

- GPT-1: 117 million parametrs trained on 4.5GB of text. Training cost: 1 month on 8 GPUs, i.e., 1×10^{19} FLOPs.
- GPT-2: 1.5 billion parameters trained on 40GB of text. Training cost: 1.5×10^{21} FLOPs.
- GPT-3: 175 billion parameters trained on 499 billion tokens. Training cost: 3.1×10^{23} FLOPs.
- GPT-4: The architecture and the data is undisclosed. Training cost: 2.1×10^{25} FLOPs.
- GPT-5: ...

Black-box Nature

- Modern AI models, usually based on deep learning, works differently from human beings.

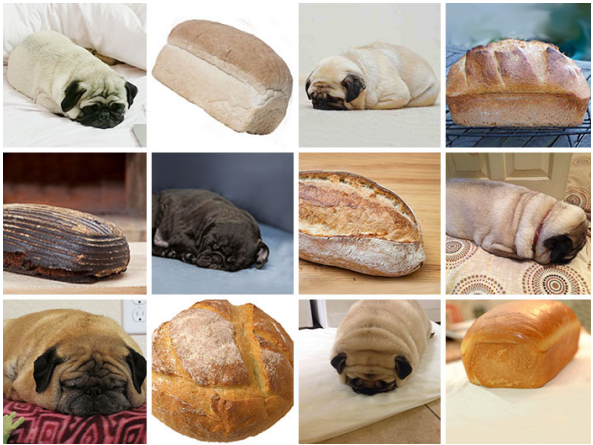

Potential Social Risks of AI

- Will AI replace / defeat human beings?
- Will AI "involution" the human beings?

Potential Social Risks of AI

(a) **Artificial Intelligence**

(b) **Natural Stupidity**

Figure: Will this happen in the future? [2]

[2] Picture is from the Internet and for reference, no offensive meaning.

Potential Social Risks of AI

Figure: United Nations Security Concil meets for the first time on AI risks in July, 2023.

Outline

AI in the Movies?

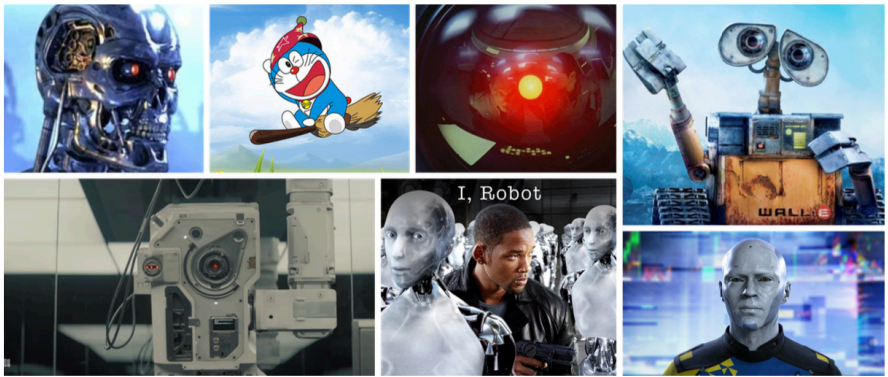

Creative and Generative AI

Creative and Generative AI

Personalized AI and AIoT

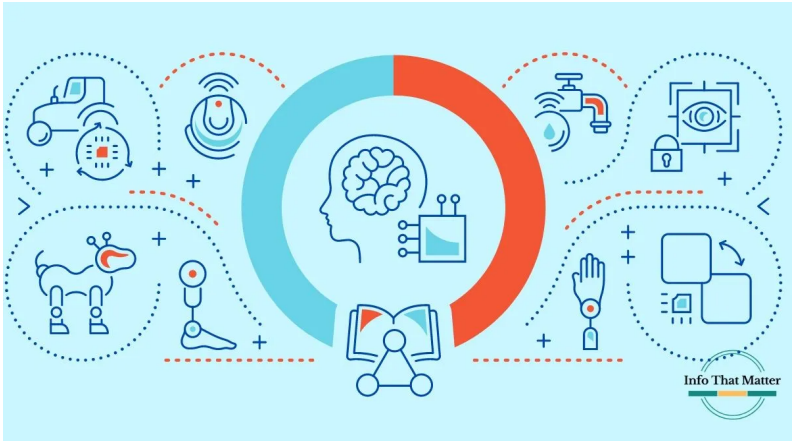

- A brief introduction to AIoT.

Light-Weighted AI

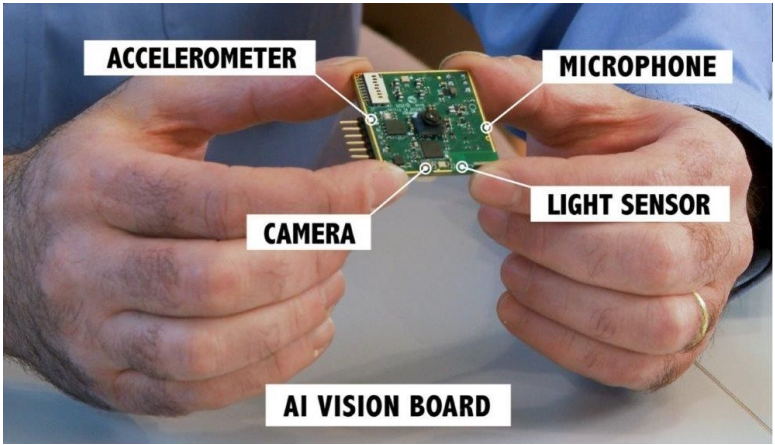

- AI Pin presentation.

AI in Different Domains

- Education.
- Manufacturing.
- Science.
- Healthcare.
- Entertainment.
- ...

Getting Closer to AGI

- In 1956, John McCarthy organized a workshop at Dartmouth College.
  - Term "artificial intelligence" was coined.
  - The participants laid out a bold proposal: to build a system that could capture every aspect of intelligence, which is called **strong AI** or **Artificial General Intelligence**.
  - In contrast to strong AI, **weak AI** is any programme that is designed to solve exactly one problem.
- We are getting closer to AGI!

Future: ASI

- **Artificial Super Intelligence (ASI)** is a hypothetical agent that possesses intelligence surpassing that of the brightest and most gifted human minds.

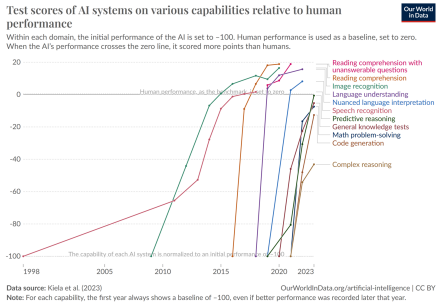

- Good or Bad?

Final Remarks

Some people call this artificial intelligence, but the reality is this technology will enhance us. So instead of artificial intelligence, I think we'll augment our intelligence.

—Ginni Rometty

Final Remarks

Some people call this artificial intelligence, but the reality is this technology will enhance us. So instead of artificial intelligence, I think we'll augment our intelligence.

—Ginni Rometty

Artificial intelligence is growing up fast, as are robots whose facial expressions can elicit empathy and make your mirror neurons quiver.

—Diane Ackerman

Final Remarks

Some people call this artificial intelligence, but the reality is this technology will enhance us. So instead of artificial intelligence, I think we'll augment our intelligence.

—Ginni Rometty

Artificial intelligence is growing up fast, as are robots whose facial expressions can elicit empathy and make your mirror neurons quiver.

—Diane Ackerman

There is no reason and no way that a human mind can keep up with an artificial intelligence machine by 2035.

—Gray Scott

Final Remarks

Some people call this artificial intelligence, but the reality is this technology will enhance us. So instead of artificial intelligence, I think we'll augment our intelligence.

—Ginni Rometty

Artificial intelligence is growing up fast, as are robots whose facial expressions can elicit empathy and make your mirror neurons quiver.

—Diane Ackerman

There is no reason and no way that a human mind can keep up with an artificial intelligence machine by 2035.

—Gray Scott

A year spent in artificial intelligence is enough to make one believe in God.

—Alan Perlis

Thanks for your attention!