AI: Introduction and Brief History

By: Dr. Mauajama Firdaus

Course Syllabus

- Artificial Intelligence Introduction
- Brief history
- Intelligent Agent
- Problem solving by search
- Constraint Satisfaction Problems
- Introduction to genetic algorithm
- Knowledge Representation
- Planning Techniques
- Probabilistic Reasoning
- Neural Network (NN) models
- Introduction to Natural Language Processing

Goals of the Introductory Classes

- A Brief Introduction to the Philosophy of Artificial Intelligence
- A Brief Introduction to the Breadth of Ideas in Artificial Intelligence
- Trace briefly the history of AI
- Understand the definition of artificial intelligence
- Understand the different faculties involved with intelligent behavior

Philosophy of Artificial Intelligence

The philosophy of artificial intelligence (AI) delves into fundamental questions about the nature of intelligence, consciousness, and the ethical implications of creating intelligent machines.

Some key concepts in the philosophy of AI:

- **Nature of Intelligence**: What does it mean for a machine to be intelligent? Can machine intelligence be compared to human intelligence, and if so, how?
- **Consciousness and Sentience**: Can AI achieve consciousness or sentience? What is the nature of subjective experience, and can a machine ever have it?
- **Ethics and Morality**: What ethical considerations arise from the creation and use of AI? This includes issues such as the moral status of AI, the impact on employment, privacy concerns, and the potential for misuse.

https://www.forrester.com/blogs/five-ai-principles-to-put-in-practice/

Philosophy of Artificial Intelligence

- **Free Will and Autonomy**: Can an AI have free will or be considered autonomous? How do we define and measure autonomy in machines?
- **Epistemology**: How do machines acquire knowledge? What are the limits of machine learning and understanding?
- **Existential Risk**: What are the potential risks associated with advanced AI? This includes discussions about superintelligence and the possible future scenarios where AI surpasses human intelligence.
- **Human-AI Interaction**: How should humans and AI interact? What are the implications for society, relationships, and communication?

Key Philosophical Positions

- **Strong AI vs. Weak AI**:
  - *Strong AI*: The belief that machines can possess true intelligence and consciousness equivalent to humans.
  - *Weak AI*: The view that machines can simulate human intelligence but do not possess consciousness or genuine understanding.

- **Functionalism**: The idea that mental states are defined by their functional roles rather than their physical makeup, suggesting that machines could, in theory, replicate human mental states.

- **Computational Theory of Mind**: The perspective that human cognition is a form of computation, and thus, machines could potentially replicate or exceed human cognitive processes.

https://plato.stanford.edu/entries/computational-mind/

Important Philosophers and Thinkers

- **Alan Turing:** Proposed the Turing Test as a measure of a machine's ability to exhibit intelligent behavior indistinguishable from a human.

- **John Searle:** Known for the Chinese Room argument, which challenges the notion that a computer running a program can have a mind or understand language.

- **Ray Kurzweil:** Advocates for the idea of the singularity, a future point where AI will surpass human intelligence and lead to unprecedented technological growth.

[1]https://academic.oup.com/mind/article/LIX/236/433/986238
[2]https://www.cambridge.org/core/journals/behavioral-and-brain-sciences/article/abs/minds-brains-and-programs/DC644B47A4299C637C89772FACC2706A
[3]https://en.wikipedia.org/wiki/The_Singularity_Is_Near

Current Debates

- AI Alignment: How can we ensure that AI systems act in ways that are aligned with human values and goals?

- AI and Creativity: Can AI be genuinely creative, or is it merely imitating human creativity?

- AI in Decision-Making: The role of AI in critical decision-making processes, such as in healthcare, law, and autonomous vehicles.

Why AI?

Economic Impact

1. **Job Creation and Transformation**:
   - AI is transforming industries by automating routine tasks, increasing efficiency, and creating new job opportunities in fields such as data analysis, machine learning engineering, and AI ethics.
2. **Productivity and Innovation**:
   - AI technologies can significantly enhance productivity by optimizing processes, reducing errors, and enabling the development of innovative products and services.
3. **Competitive Advantage**:
   - Organizations that effectively leverage AI can gain a competitive edge by making more informed decisions, improving customer experiences, and predicting market trends.

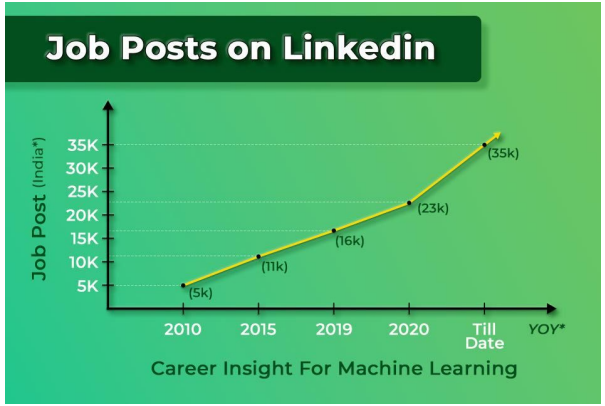

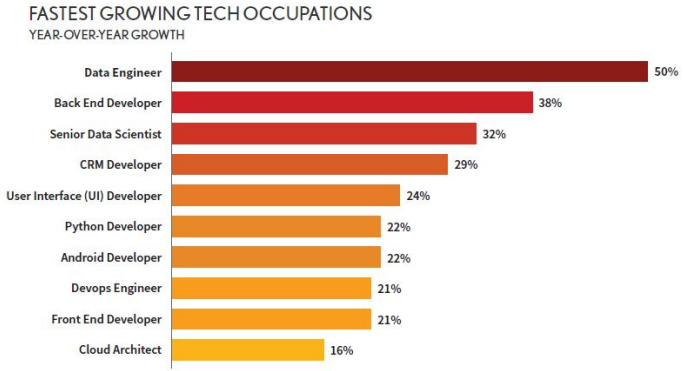

Social Impact

1. **Improving Quality of Life**:
   - AI is being used to improve healthcare through predictive analytics which can lead to better patient outcomes and more efficient healthcare systems.
2. **Addressing Global Challenges**:
   - AI can help tackle pressing global issues such as climate change, food security, and disaster response by analyzing large datasets to provide insights and solutions.
3. **Enhanced Accessibility**:
   - AI-driven technologies like speech recognition, language translation, and assistive devices are helping to bridge accessibility gaps, making technology more inclusive for people with disabilities.

Ethical and Regulatory Considerations

1. **Ethical AI Development**:
   - Understanding AI is crucial to ensure that it is developed and deployed ethically, addressing issues such as bias, fairness, transparency, and accountability.
2. **Policy and Regulation**:
   - Policymakers need a thorough understanding of AI to create regulations that protect public interest while fostering innovation and preventing misuse.
3. **Privacy and Security**:
   - Studying AI helps address concerns related to data privacy and security, ensuring that AI systems are designed to protect sensitive information and prevent malicious attacks.

Technological Advancement

1. **Driving Research and Development**:
    - AI is at the forefront of technological innovation, pushing the boundaries of what is possible in fields like robotics, natural language processing, and computer vision.
2. **Interdisciplinary Collaboration**:
    - AI research fosters collaboration across various disciplines, including computer science, neuroscience, linguistics, and psychology, leading to a more comprehensive understanding of intelligence and cognition.
3. **Preparing for the Future**:
    - As AI continues to evolve, it is essential to prepare for its future implications, including the potential development of artificial general intelligence (AGI) and its societal impact.

Societal and Ethical Considerations

1. **Human-AI Interaction**:
   - Understanding how humans and AI can interact effectively is vital for developing systems that augment human capabilities and provide intuitive user experiences.
2. **Long-term Implications**:
   - Studying AI helps society anticipate and plan for long-term implications, such as changes in the labor market, shifts in power dynamics, and the potential for superintelligent AI.
3. **Educational Transformation**:
   - AI is revolutionizing education through personalized learning experiences, intelligent tutoring systems, and data-driven insights into student performance.

Google AI Why AI Responsibility Discover Build

Our Principles

While we are optimistic about the potential of AI, we recognize that advanced technologies can raise important challenges that must be addressed clearly, thoughtfully, and affirmatively. These AI Principles describe our commitment to developing technology responsibly and work to establish specific application areas we will not pursue.

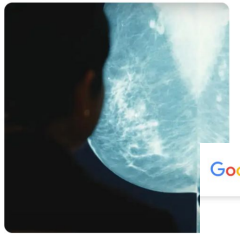

Google AI Why AI Responsibility Discover Build

Objectives for AI applications

1. Be socially beneficial.
2. Avoid creating or reinforcing unfair bias.
3. Be built and tested for safety.
4. Be accountable to people.
5. Incorporate privacy design principles.
6. Uphold high standards of scientific excellence.
7. Be made available for uses that accord with these principles.

https://ai.google/responsibility/principles/

ACL Rolling Review

An initiative of the Association for Computational Linguistics

Home CFP Authors Reviewing Dates/Venues Stats People Future Login

Members of the ACL are responsible for adhering to the ACL code of ethics. The ARR Responsible NLP Research checklist is designed to encourage best practices for responsible research, addressing issues of research ethics, societal impact and reproducibility. It is largely based on the NeurIPS 2021 paper checklist, the reproducible data checklist from Rogers, Baldwin, Leins's EMNLP 2021 paper "Just What do You Think You're Doing, Dave? A Checklist for Responsible Data Use in NLP", and the NLP Reproducibility Checklist built from Dodge et al. EMNLP 2019's paper "Show Your Work: Improved Reporting of Experimental Results".

We expect authors to show that they follow best practices in two ways:

1. by filling in the checklist to ensure that best practices are put in place when using, creating or providing assets,
2. by including a discussion in the paper about any potential positive or negative societal impacts and any limitations of the work. The guidelines below provide additional information about what should be discussed.

https://aclanthology.org/2021.findings-emnlp.414.pdf
https://aclanthology.org/D19-1224.pdf

Applications

1. Healthcare

- **Medical Diagnosis:** AI systems can analyze medical images and data to help diagnose diseases like cancer, heart disease, and neurological disorders.
- **Personalized Medicine:** AI can tailor treatments to individual patients based on their genetic makeup and health history.

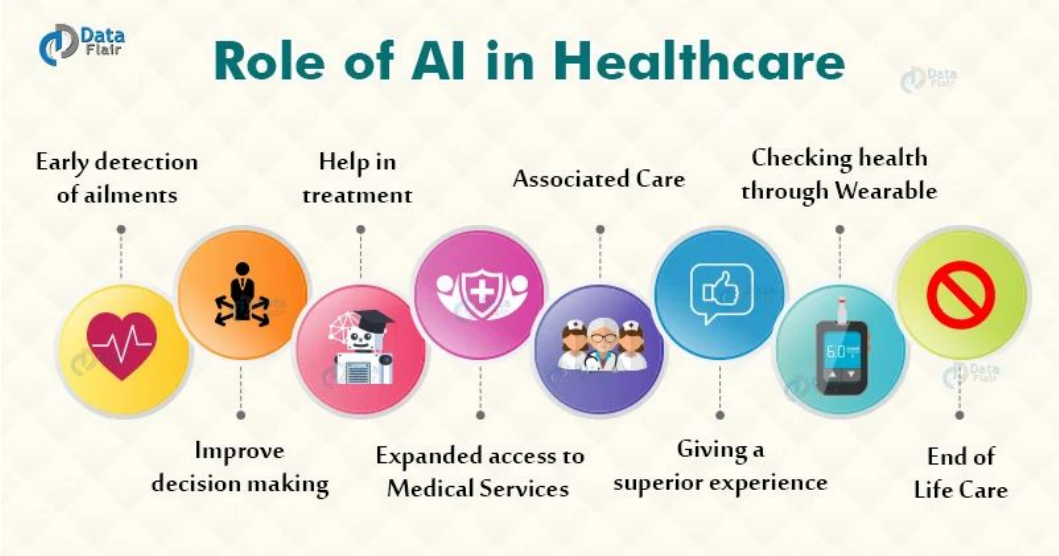

2. Education

- **Personalized Learning:** AI-driven platforms can customize learning experiences based on students' strengths, weaknesses, and learning styles.
- **Automated Grading:** AI can assist teachers by grading assignments and providing feedback, allowing them to focus more on instruction.

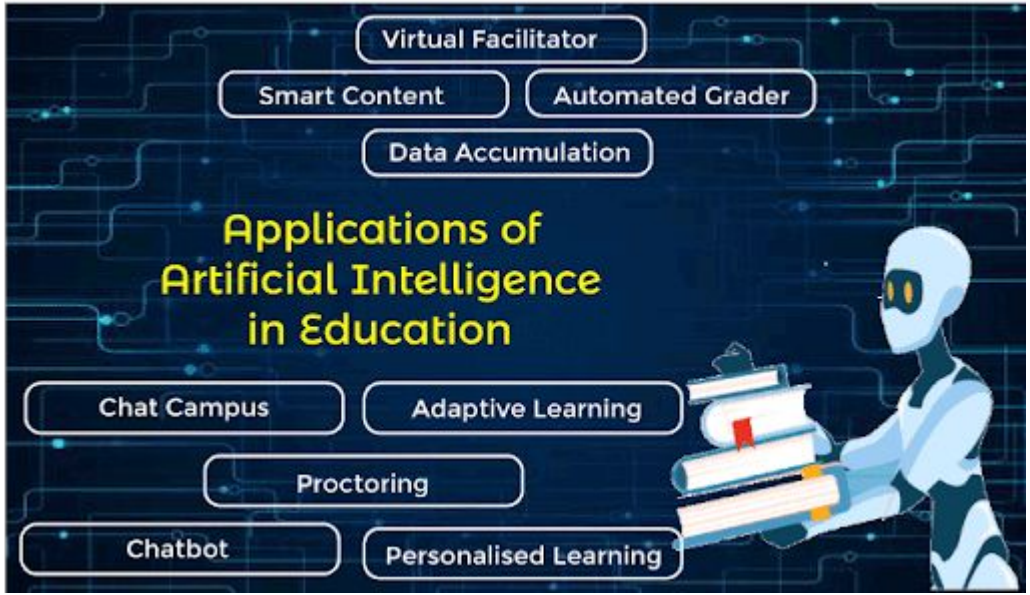

3. Environment and Sustainability

- **Climate Change Modeling:** AI can analyze large datasets to predict climate changes and model the impacts of various environmental policies.
- **Wildlife Conservation:** AI is used to monitor wildlife populations and combat poaching through image recognition and data analysis.

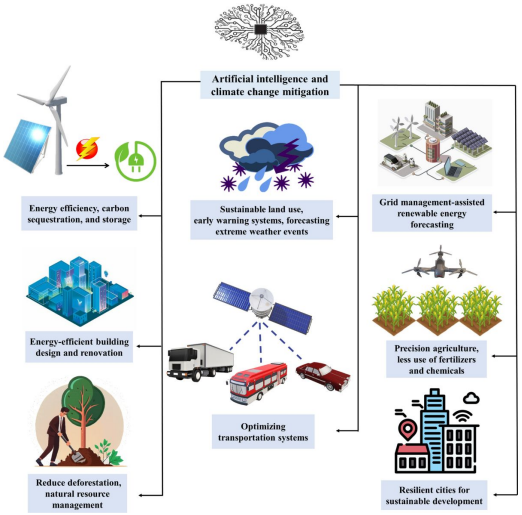

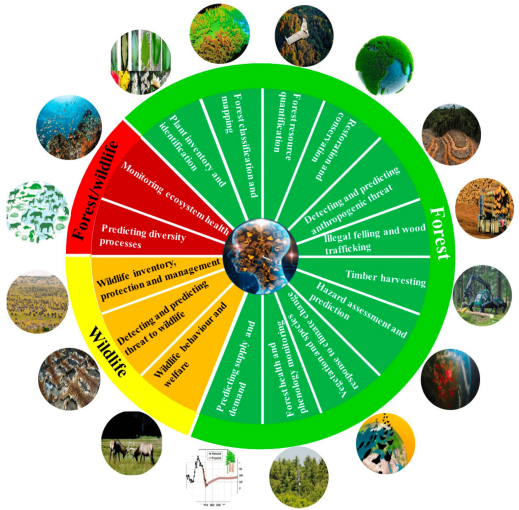

4. Entertainment and Media

- **Content Creation:** AI can generate music, art, and stories, showcasing creative potentials that appeal to students interested in the arts.
- **Game Development:** AI is integral in creating intelligent behaviors in non-player characters (NPCs) and enhancing gaming experiences.

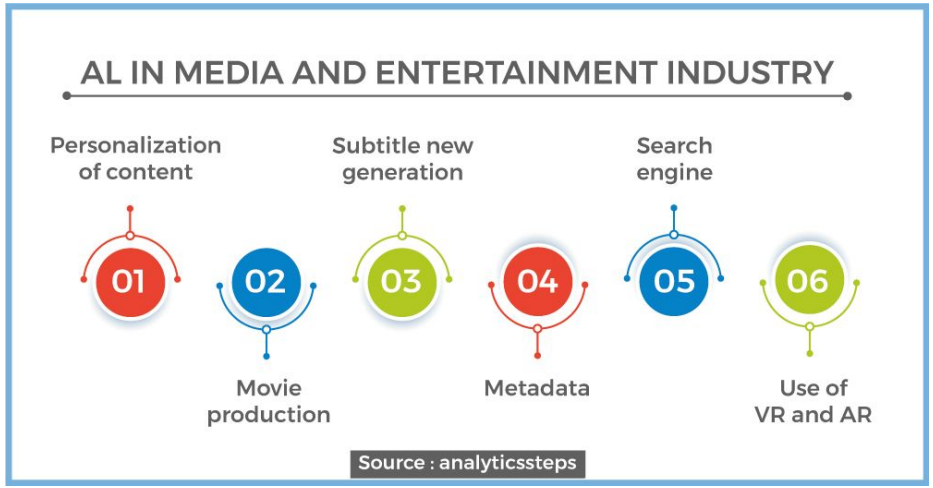

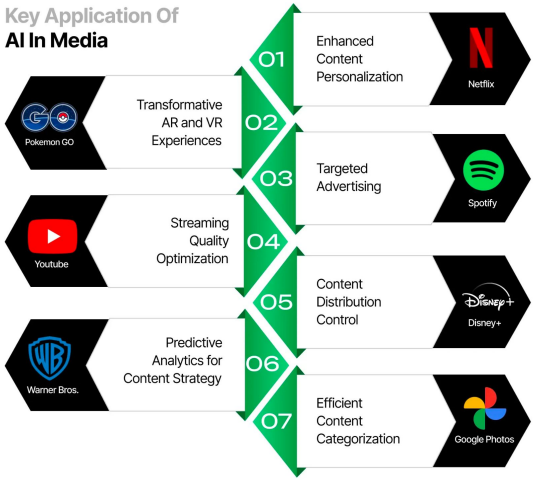

5. Robotics

- **Autonomous Vehicles:** AI powers self-driving cars, which can revolutionize transportation and logistics.
- **Drones:** AI-driven drones are used for delivery, agriculture, and disaster relief efforts.

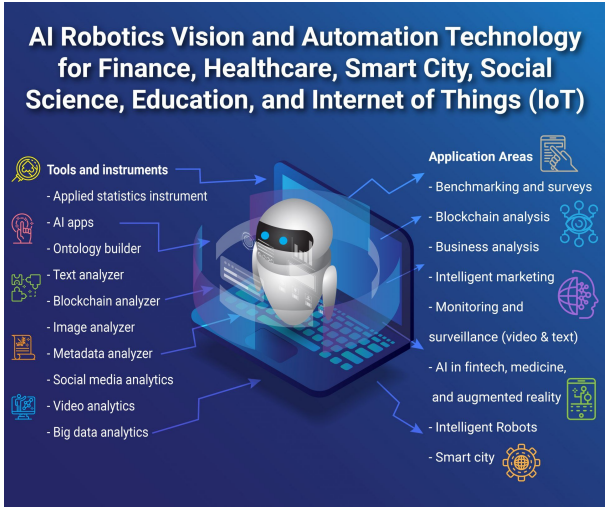

6. Finance

- **Fraud Detection:** AI systems can detect fraudulent transactions by analyzing patterns in financial data.
- **Algorithmic Trading:** AI algorithms can analyze market trends and make trading decisions in real-time.

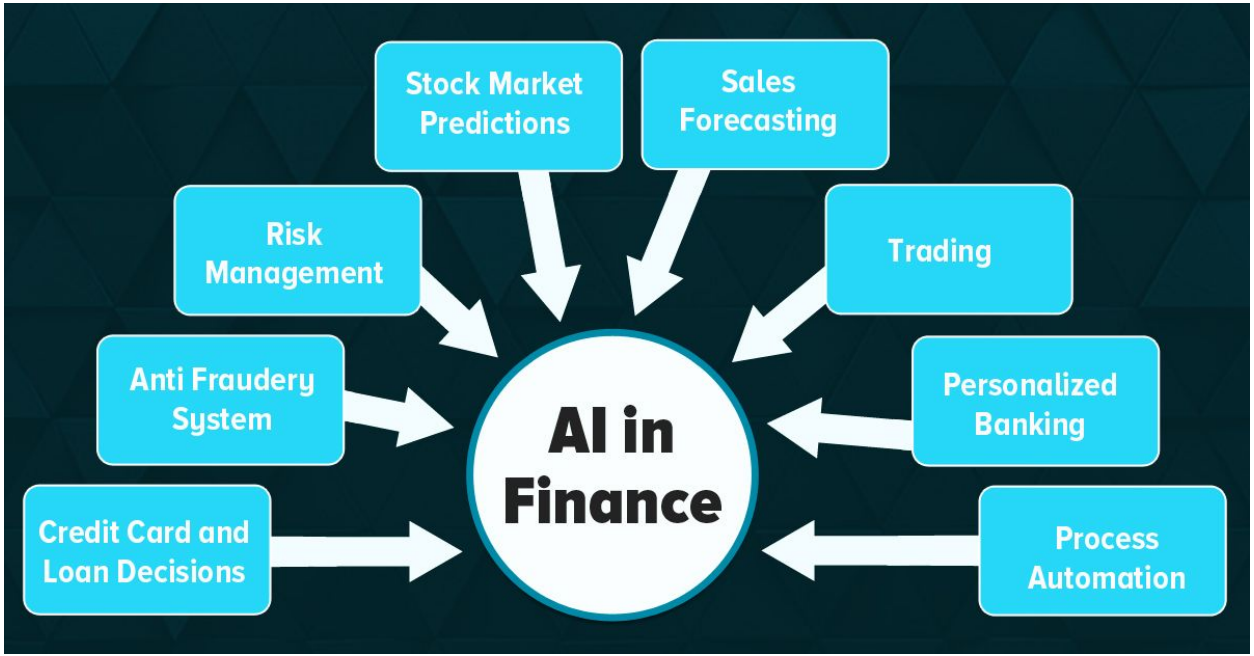

7. Social Good

- **Disaster Response:** AI helps in predicting natural disasters and coordinating relief efforts by analyzing data from various sources.
- **Humanitarian Aid:** AI systems can optimize resource distribution in crisis situations, ensuring that aid reaches those in need efficiently.

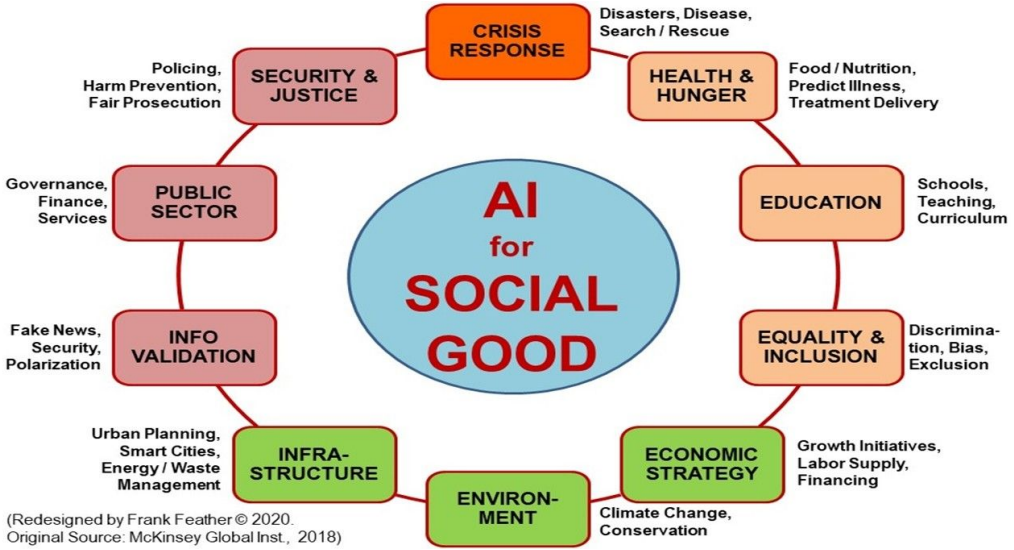

8. Daily Life Enhancements

- **Smart Assistants:** AI-powered virtual assistants like Siri, Alexa, and Google Assistant can help with everyday tasks and information retrieval.
- **Home Automation:** AI systems control smart homes, managing lighting, security, and climate to improve comfort and energy efficiency.

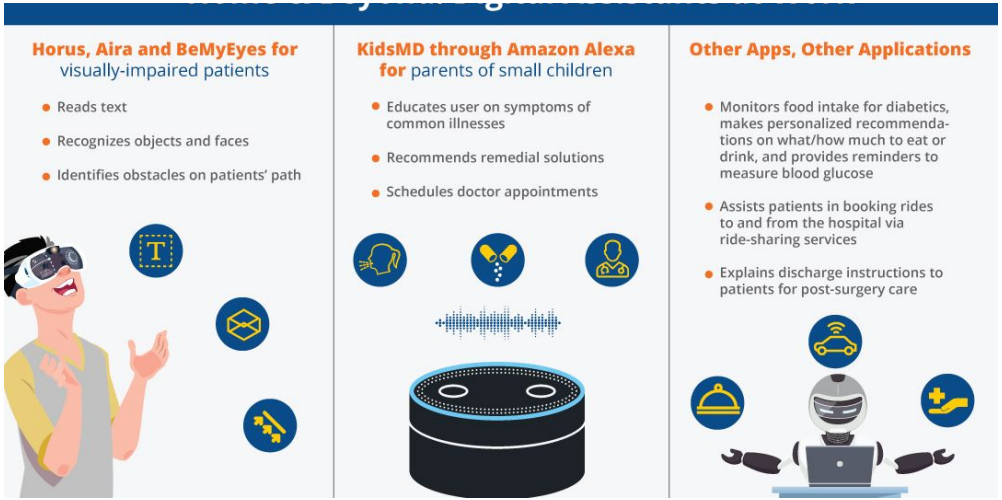

9. Sports

- **Performance Analysis:** AI can analyze athletes' performance and suggest improvements based on data-driven insights.
- **Fan Engagement:** AI enhances the fan experience through personalized content and interactive features during live events.

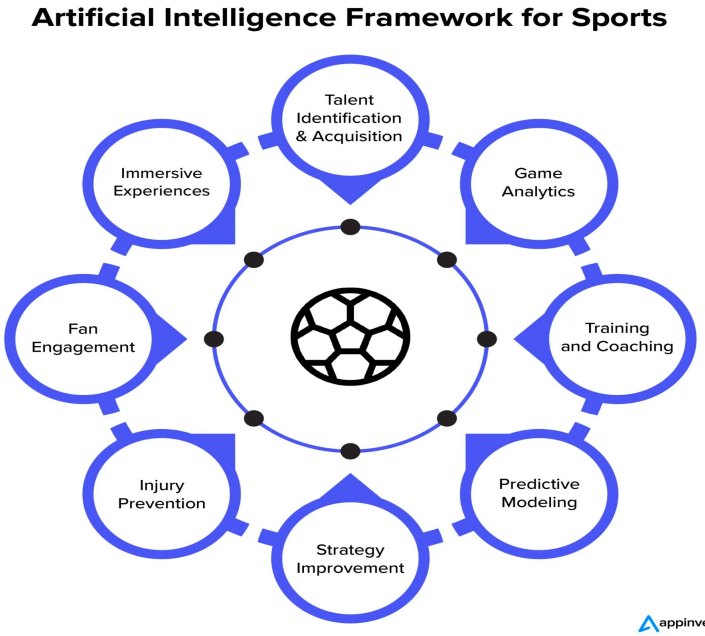

10. Scientific Research

- **Drug Discovery:** AI accelerates the process of discovering new drugs by analyzing vast amounts of biological data.
- **Astronomy:** AI helps in processing data from telescopes and discovering new celestial bodies and phenomena.

History of AI
1940's - Present

Early Foundations

- **1943:** Warren McCulloch and Walter Pitts develop a mathematical model for neural networks.

- **1946:** The ENIAC (Electronic Numerical Integrator and Computer) is completed, marking a significant step in the development of computing technology.

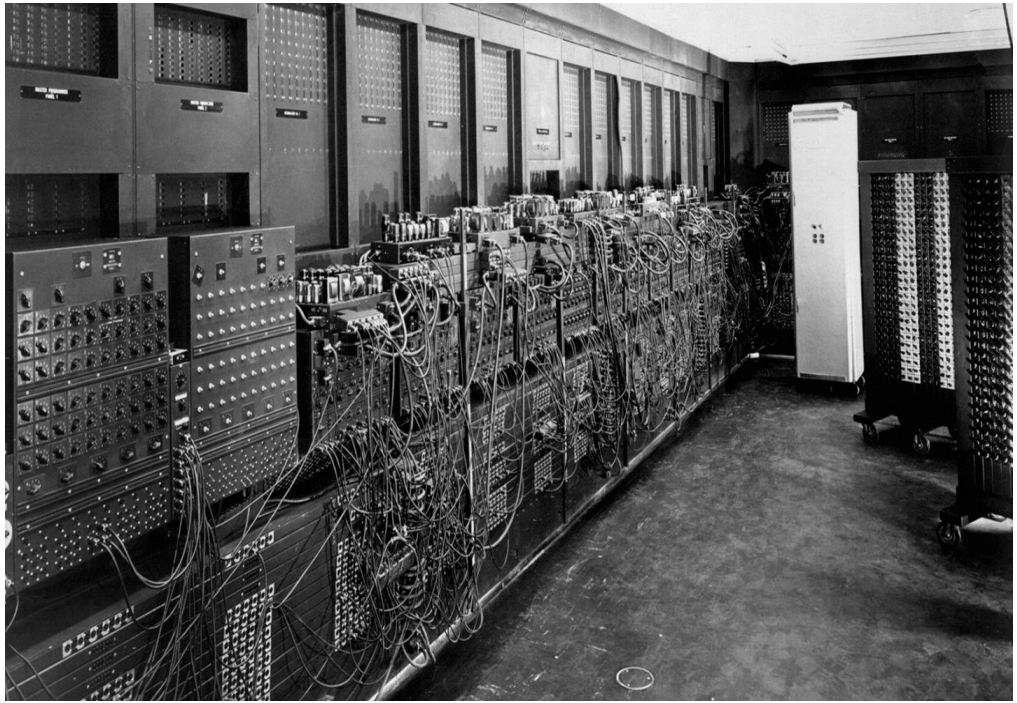

1950: The Question

- **1950:** Alan Turing publishes "Computing Machinery and Intelligence," introducing the concept of the Turing Test to assess a machine's ability to exhibit intelligent behavior.

Turing asks the question:

“ Can machines think?”

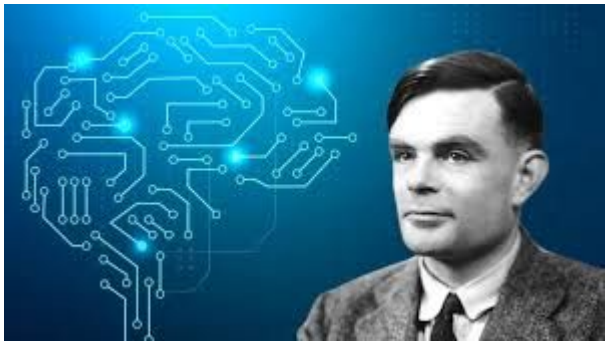

How AI got its Name!

The name Artificial Intelligence is credited to John McCarthy who, along with Marvin Minsky, Nathaniel Rochester and Claude Shannon, organized the Dartmouth Conference in 1956.

Dartmouth Project Proposal: Proposal of a 2 month 10 man study of Artificial Intelligence be carried out at Dartmouth College in Hanover, New Hampshire.. “Two month, ten-man study of AI.. on the basis of the conjecture that every aspect of learning or any other feature of intelligence can in principle be so precisely described that a machine can be made to simulate it”.

John McCarthy

Marvin Minsky

Dartmouth Conference: The Organizers

- John McCarthy (1927-2011)
  - an Assistant Professor at Dartmouth.
  - Designed the LISP programming language that was very popular with AI researchers.
  - Also did work in Logic and Commonsense Reasoning.
- Marvin Minsky (1927-2016)
  - a Harvard Junior Fellow went on to become one of the most influential figures in AI.
  - With McCarthy he co-founded the MIT AI Lab.
  - Known for his ideas on Frames. Wrote a book “Society of the Mind” and more recently “The Emotion Machine”.
- Nathaniel Rochester (1919-2001)
  - a young engineer at IBM.
  - He designed the IBM 701 and wrote the first assembler.
  - He supervised Arthur Samuel writing the checkers playing program.
  - It is said that the marketing people at IBM reported that people were frightened of “electronic brains” resulting in IBM stopping work on AI.
- Claude Shannon (1916-2001)
  - a mathematician at Bell Labs was already known for his information theory.
  - Had hired McCarthy and Minsky in 1952 for the summer when they were graduate students.

Dartmouth Conference: The Show Stealers

- Herbert Simon and Alan Newell were “two vaguely known persons” working at Carnegie Institute of Technology (now Carnegie Mellon University), who were also invited to the Dartmouth Conference “almost as an afterthought”.
- Simon was primarily known for his work in economics and organizational behavior, having laid the groundwork for the concept of bounded rationality.
- Newell was a computer scientist with a keen interest in cognitive psychology and human problem-solving.
- In collaboration with Simon, Newell worked on the General Problem Solver (GPS), an early AI program designed to imitate human problem-solving skills across a variety of domains.
- They had already developed program called the Logic Theorist, considered one of the first AI programs, which was capable of proving mathematical theorems.

Samuel's Checkers Program

- Arthur Samuel (1901-1990) was one of the attendees in Dartmouth.
- He began working on the checkers program in the early 1950s, precisely 1952 on IBM's 701.
- The goal was to create a program that could play checkers at a high level, ultimately improving its performance through learning from experience.

ELIZA - the Chatbot Psychotherapist

- Created by Joseph Weizenbaum at the MIT Artificial Intelligence Laboratory between 1964 and 1966.
- ELIZA was designed to simulate conversation with a psychotherapist, using pattern matching and substitution methodology to create the illusion of understanding.
- The program demonstrated how computers could be made to communicate with humans in natural language, laying the groundwork for future developments in conversational agents and natural language processing (NLP).

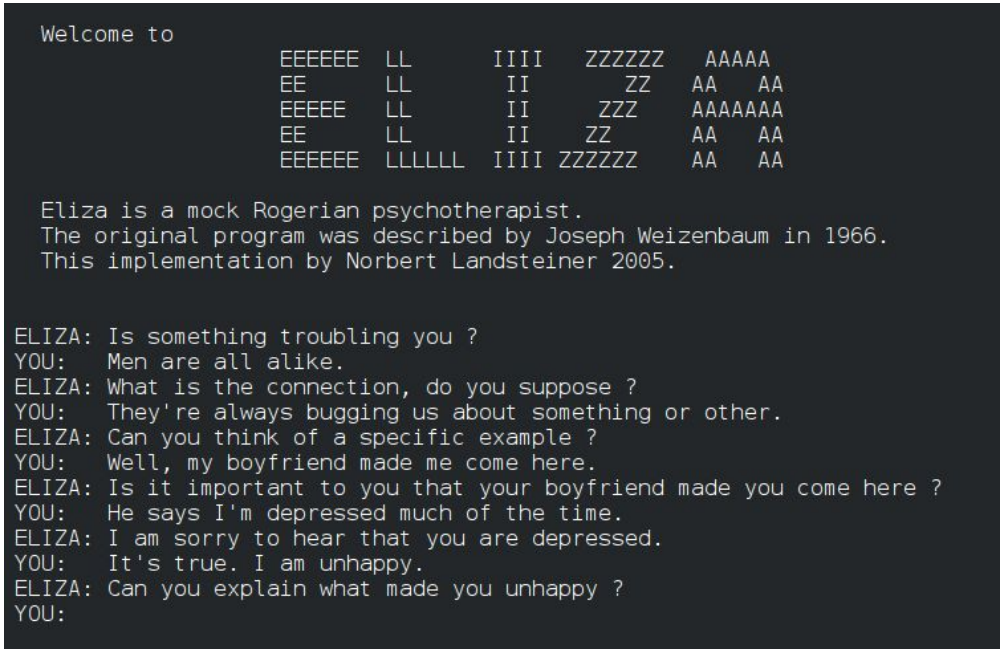

Shakey - the general purpose mobile robot

- Shakey the Robot was a pioneering project in robotics and artificial intelligence developed by the Stanford Research Institute (SRI) from 1966 to 1972.
- Shakey was the first general-purpose mobile robot capable of reasoning about its actions and operating autonomously in a structured environment.
- This project significantly advanced the fields of robotics, AI, and computer vision.
- Shakey could navigate through rooms, avoid obstacles, and move objects to specific locations. It combined data from its sensors to map its environment and plan paths.

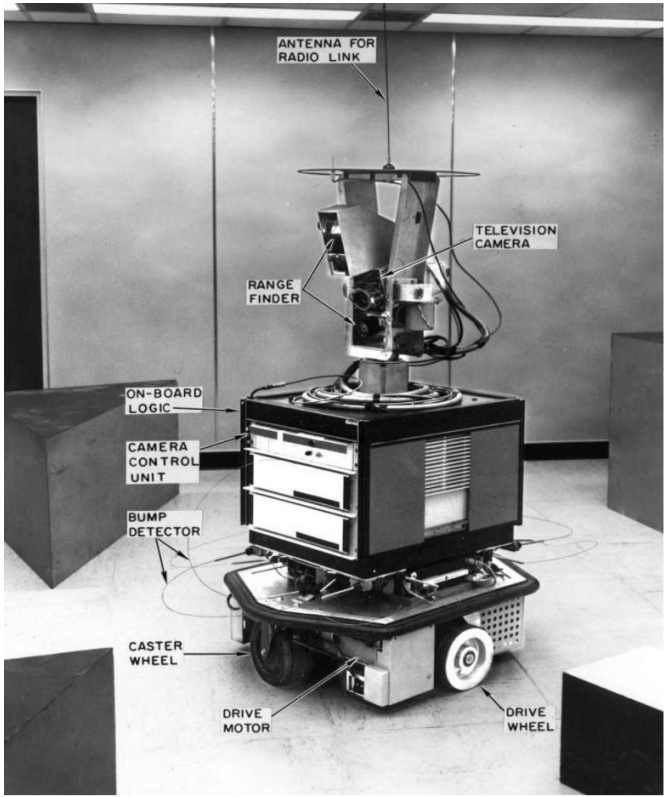

During 1956 - 1966

- 1950: Turing Test for Machine Intelligence
- 1956: AI born at Dartmouth College Workshop
- 1964: ELIZA - the Chatbot Psychotherapist
- 1966: Shakey - the general purpose mobile robot

The First AI Winter: 1974 - 1980

Despite initial enthusiasm, progress in AI research slows due to limitations in computational power and overly optimistic predictions, leading to reduced funding and interest. This period is known as the "AI Winter."

- Failure of Machine Translation
- Negative results in Neural Nets
- Poor speech understanding

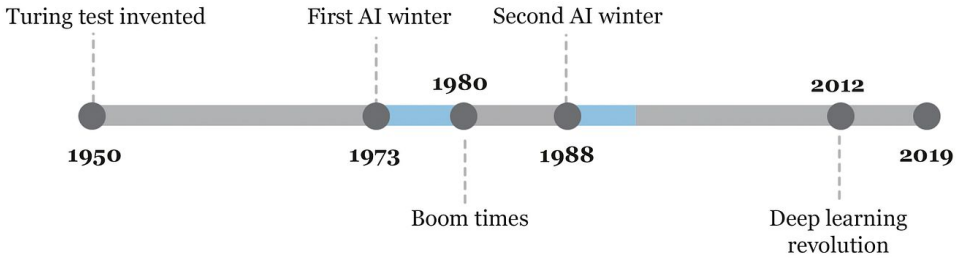

Renewed Interest and Expert Systems

AI research experiences a resurgence with the development of expert systems, which are computer programs designed to mimic the decision-making abilities of human experts. Companies like IBM and others invest heavily in AI research.

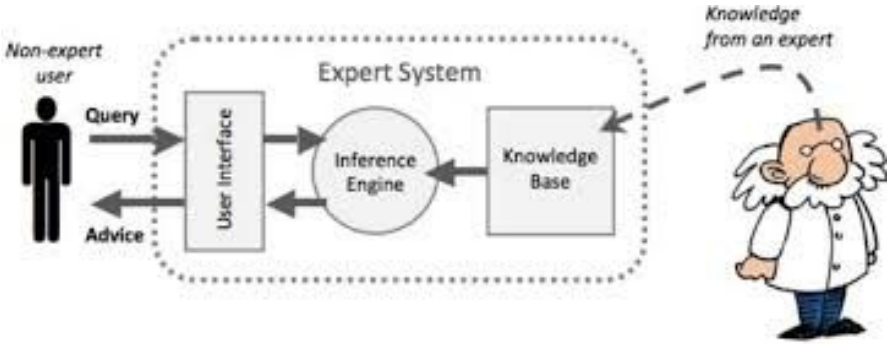

The Second AI Winter: Late 1980s-1990s

The limitations of expert systems and other AI technologies lead to another period of reduced funding and interest, marking the second AI Winter.

- Decline of LISP
- Decline of specialized hardware for expert systems

Lasting effects:

- Economist'07: "AI is associated with systems that have all too often failed to live up to their promises"
- Pittsburgh BT06: "some believe the word "robotics" actually carry stigma that hurts a company's chances at funding"

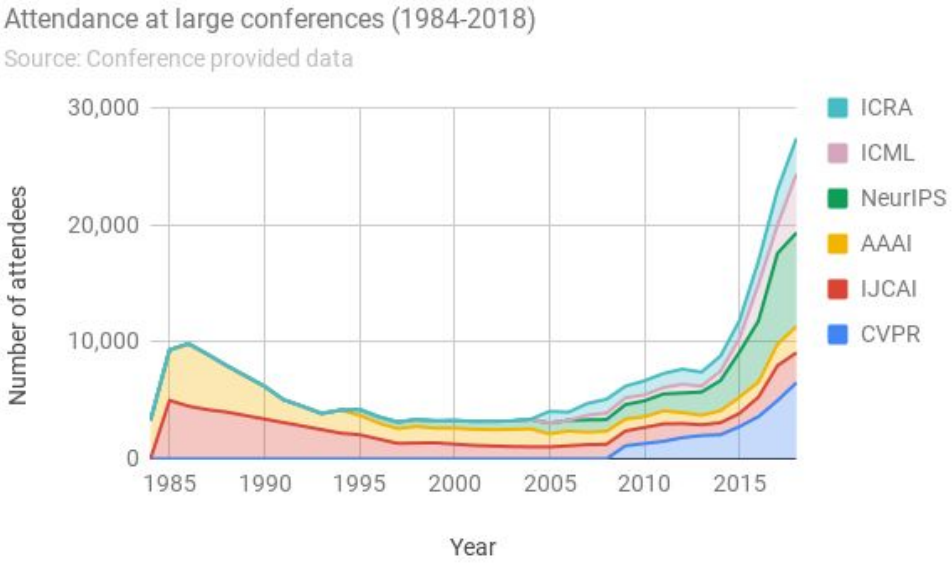

First Success: 1996

The theorem prover EQP (Equational Prover) achieved a significant milestone in AI history by solving the Robbins conjecture, a long-standing problem in algebra.

The success of EQP in solving the Robbins conjecture underscored the potential of AI to tackle complex mathematical problems and contributed to renewed interest and investment in AI research during the late 1990s and early 2000s.

1996: EQP proves that Robbin's Algebras are all boolean

----- EQP 0.9, June 1996 -----
The job began on eyas09.mcs.anl.gov, Wed Oct 2 12:25:37 1996
UNIT CONFLICT from 17666 and 2 at 678232.20 seconds.
---------- PROOF ----------
2 (wt=7) [] -(n(x + y) = n(x)).
3 (wt=13) [] n(n(n(x) + y) + n(x + y)) = y.
5 (wt=18) [para(3,3)] n(n(n(x + y) + n(x) + y) + y) = n(x + y).
6 (wt=19) [para(3,3)] n(n(n(n(x) + y) + x + y) + y) = n(n(x) + y).
.....
17666 (wt=33) [para(24,16426),demod([17547])] n(n(n(x) + x) ....

[An Argonne lab program] has come up with a major mathematical proof that would have been called creative if a human had thought of it.
-New York Times, December, 1996

1997: Deep Blue

In 1997, IBM's Deep Blue made history by defeating Garry Kasparov, the reigning world chess champion, in a six-game match. This event was a landmark achievement in the field of artificial intelligence, demonstrating the potential of AI systems to perform at the highest levels in strategic games.

Success Story: Chess

Does Deep Blue uses AI?

" If it works, its not AI"

Drew McDermott, a computer science professor at Yale:

"Saying Deep Bleu doesn't really think about chess is like saying an airplane doesn't really fly because it doesn't flap its wings."

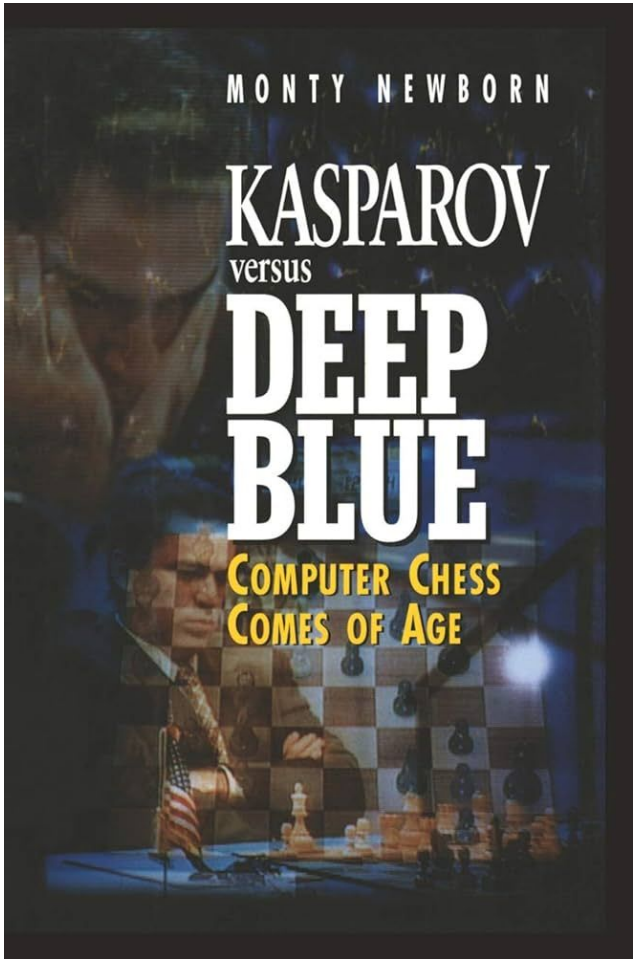

1999: Deep Space I

In 1999, DS1 became the first spacecraft to use an AI-based autonomous system called the Remote Agent Experiment.

The RAE was capable of planning and executing its own activities, detecting and diagnosing problems, and making decisions without real-time human intervention.

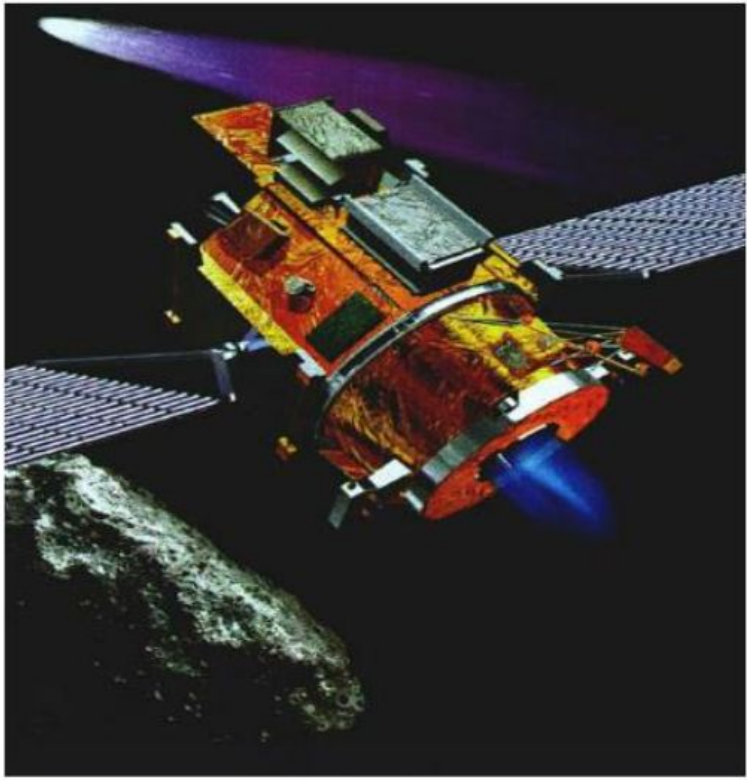

2004 - 2009: Mars Rovers

NASA's Mars Exploration Rovers, Spirit and Opportunity, landed on Mars.

They utilized AI for autonomous navigation, allowing them to make decisions about which paths to take to avoid obstacles and choose scientifically interesting targets.

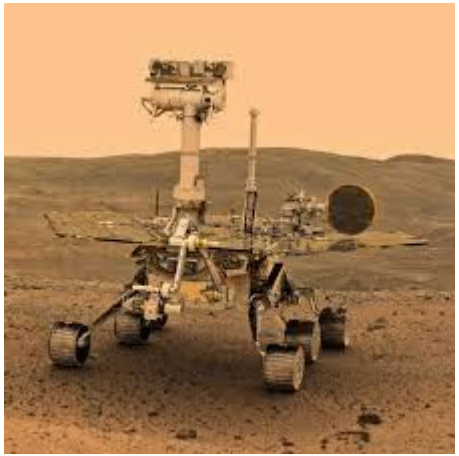

2005: Self-driving Cars

The 2005 DARPA Grand Challenge was a pivotal event in the development of self-driving car technology.

Organized by the Defense Advanced Research Projects Agency (DARPA), the challenge aimed to advance autonomous vehicle technology for potential military applications.

The success of the 2005 competition marked a significant milestone in the field of artificial intelligence and robotics.

2011: IBM’s Watson

In 2011, IBM's Watson made history by competing on the quiz show "Jeopardy!" and defeating two of the show's greatest champions, Ken Jennings and Brad Rutter.

This event was a major milestone in the field of artificial intelligence, showcasing Watson's advanced natural language processing, information retrieval, and machine learning capabilities.

2016: AlphaGo

In 2016, Google's DeepMind created a groundbreaking AI program called AlphaGo, which made history by defeating one of the world's top Go players, Lee Sedol.

Human-level control through deep reinforcement learning

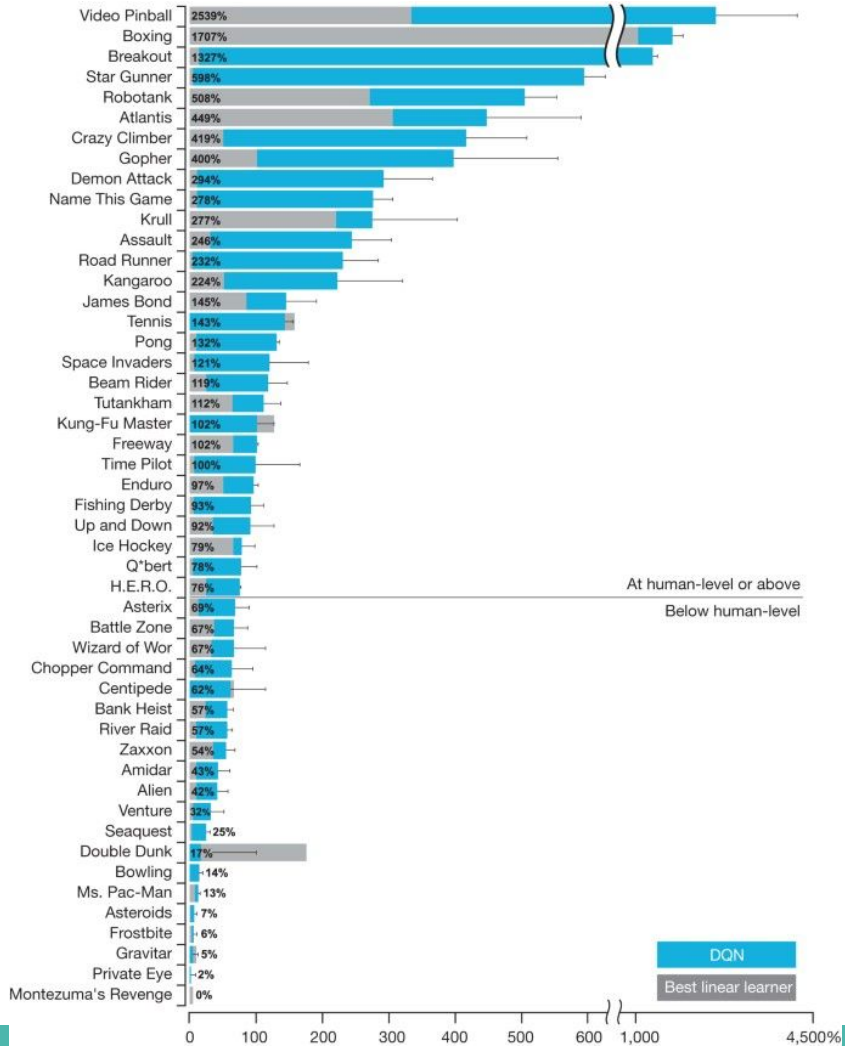

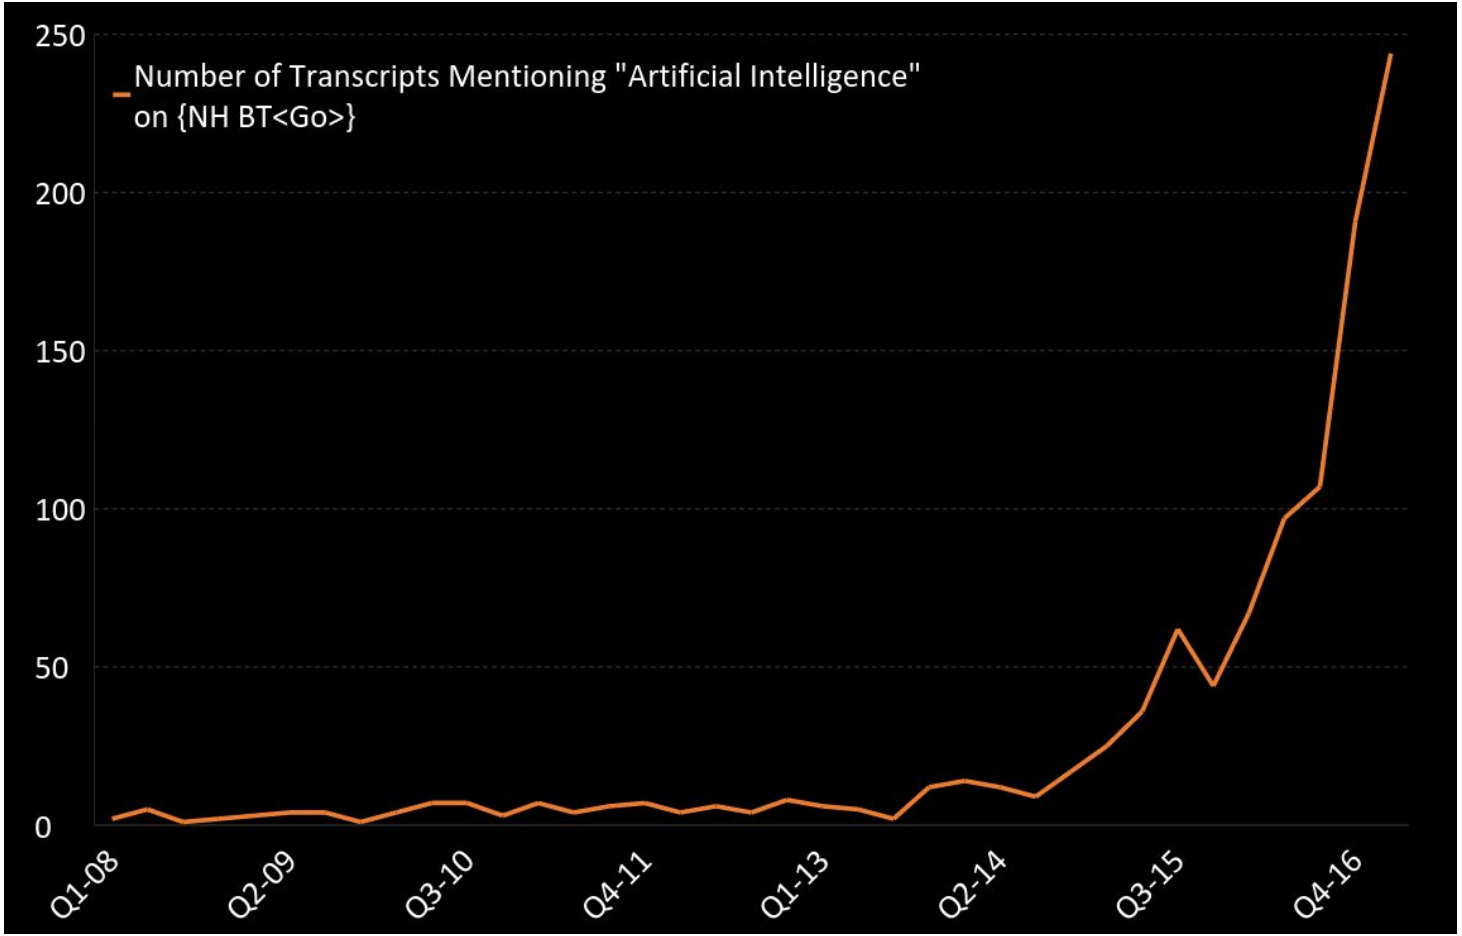
Number of Transcripts Mentioning "Artificial Intelligence" on {NH BT<Go>}
250
200
150
100
50
0
Q1-08
Q2-09
Q3-10
Q4-11
Q1-13
Q2-14
Q3-15
Q4-16

What Changed?

DATA

DEEP NEURAL NETWORKS

COMPUTATION POWER

ALGORITHM

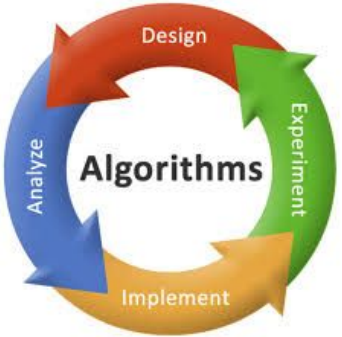

Object Recognition

CIFAR
CANADIAN INSTITUTE
for ADVANCED RESEARCH

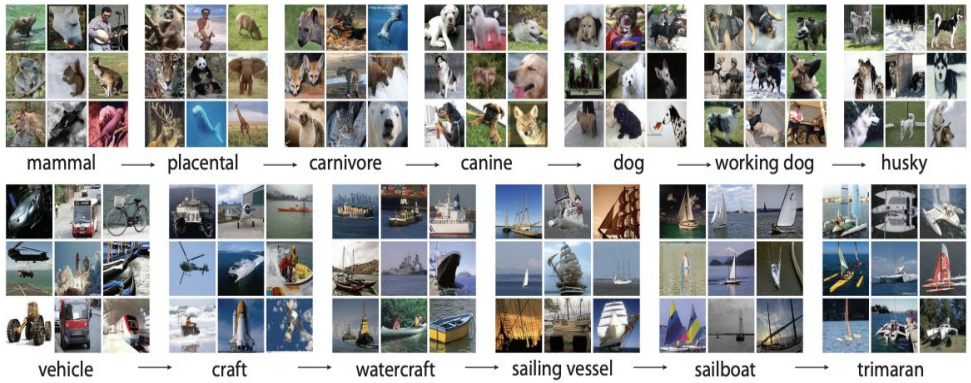

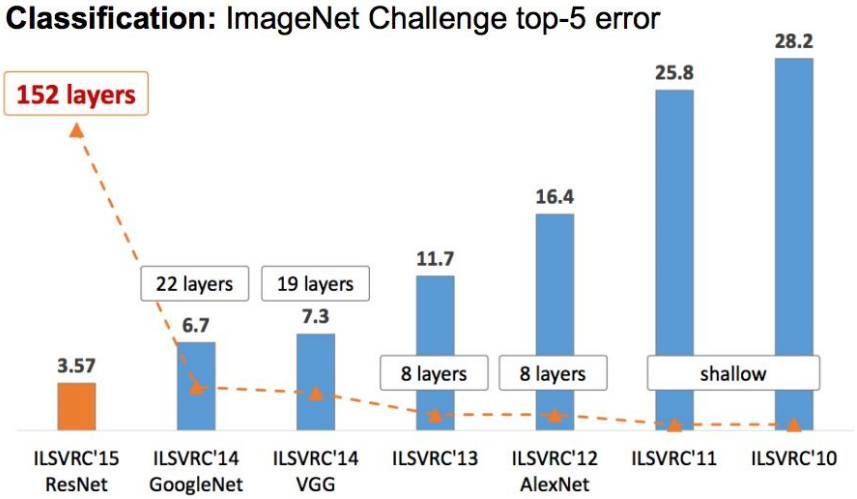

Artistic Applications

- Doodle to painting

- Style Transfer

- Image Colorization

Image to Caption

A person riding a motorcycle on a dirt road.

Two dogs play in the grass.

A skateboarder does a trick on a ramp.

A dog is jumping to catch a frisbee.

A group of young people playing a game of frisbee.

Two hockey players are fighting over the puck.

A little girl in a pink hat is blowing bubbles.

A refrigerator filled with lots of food and drinks.

A herd of elephants walking across a dry grass field.

A close up of a cat laying on a couch.

A red motorcycle parked on the side of the road.

A yellow school bus parked in a parking lot.

“If it works its not AI” → “It’s all AI”

Elon Musk predicts AI will overtake humans to the point that 'biological intelligence will be 1%'

Hype Or Reality: Will AI Really Take Over Your Job?

Robots TAKING OVER: AI to 'sink world into unemployed despair in hellish dystopia'

Robocup

By 2050, develop a team fully autonomous humanoid robots that can win against the human world champion team in soccer.

https://www.youtube.com/watch?v=80qVUejnj8U