The next industrial revolution

Arlindo Oliveira

TÉCNICO LISBOA

inesc id lisboa

It is difficult to make predictions, especially about the future

If **trend X** continues, the result will be **disaster**.


K.& Co.
329. PYRMONT BRIDGE

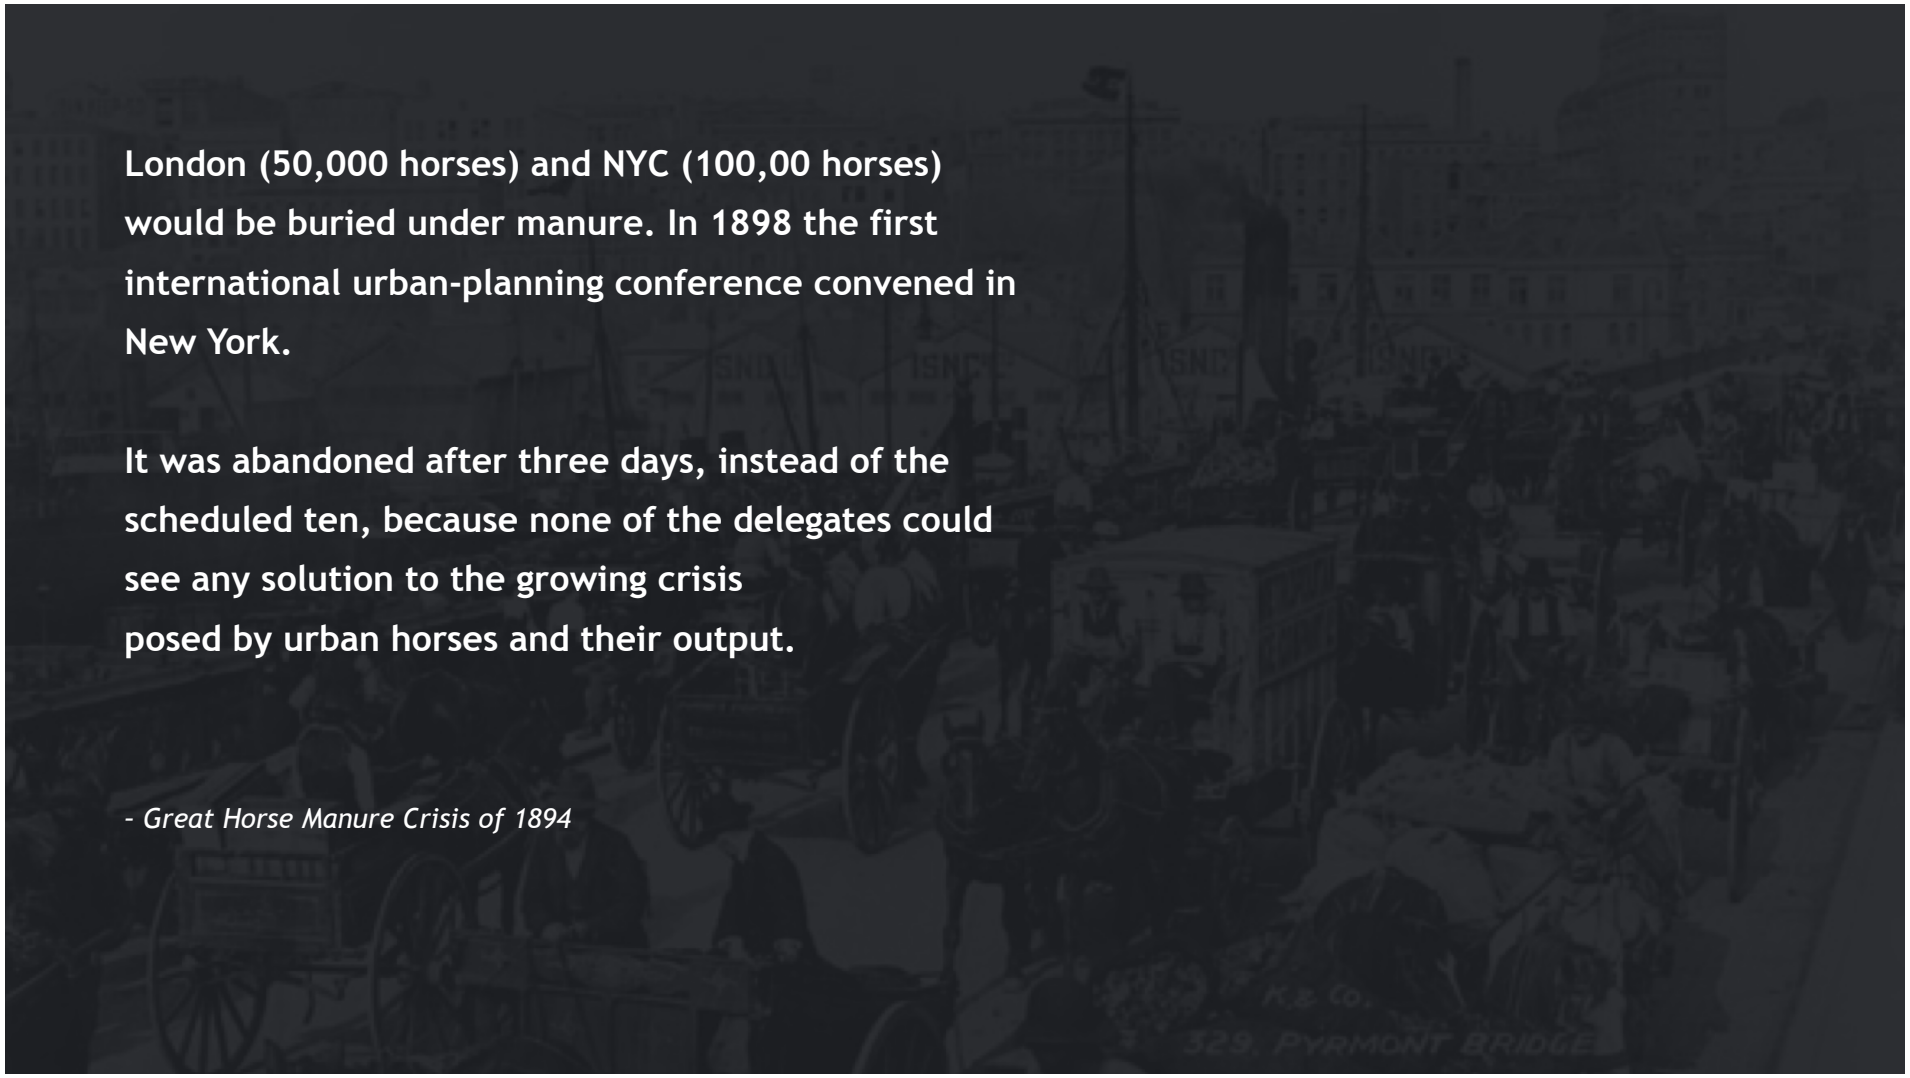

London (50,000 horses) and NYC (100,00 horses) would be buried under manure. In 1898 the first international urban-planning conference convened in New York.

It was abandoned after three days, instead of the scheduled ten, because none of the delegates could see any solution to the growing crisis posed by urban horses and their output.

- Great Horse Manure Crisis of 1894

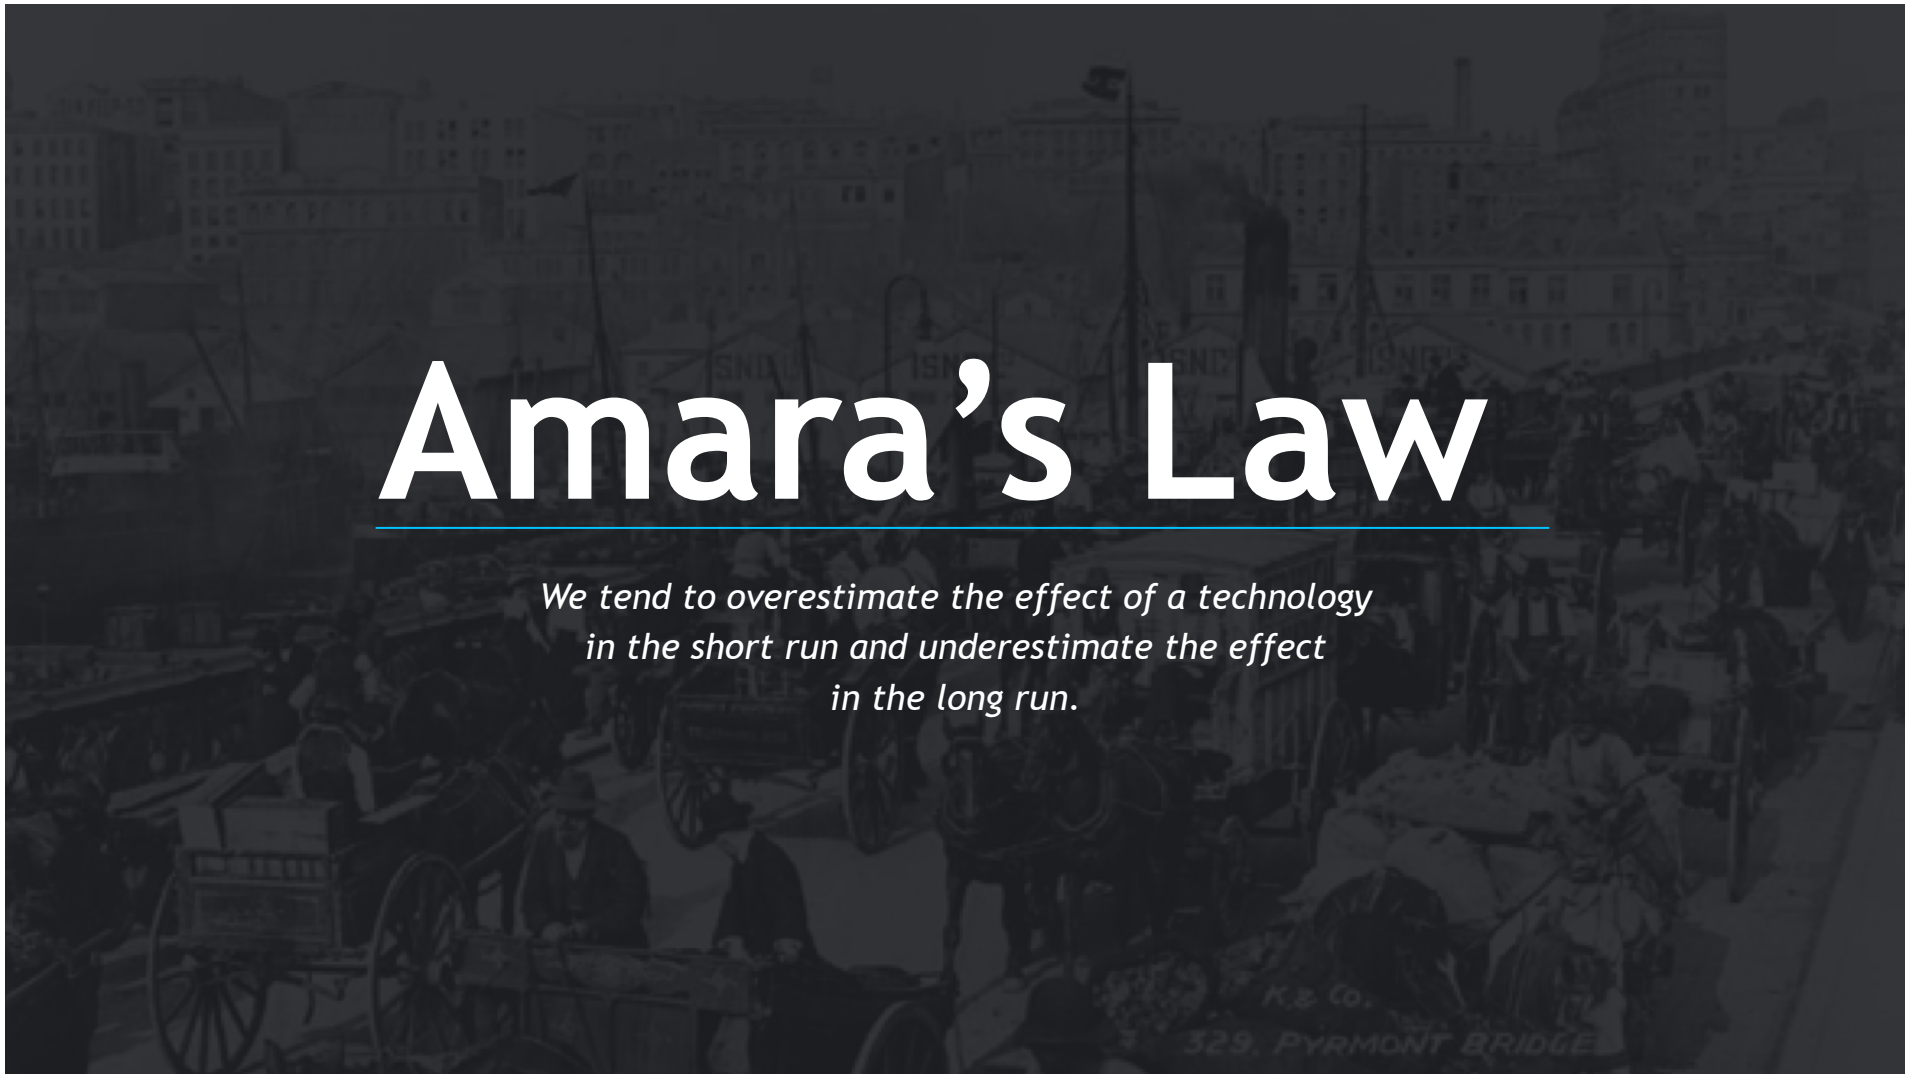

Amara's Law

We tend to overestimate the effect of a technology in the short run and underestimate the effect in the long run.

Amara's Law

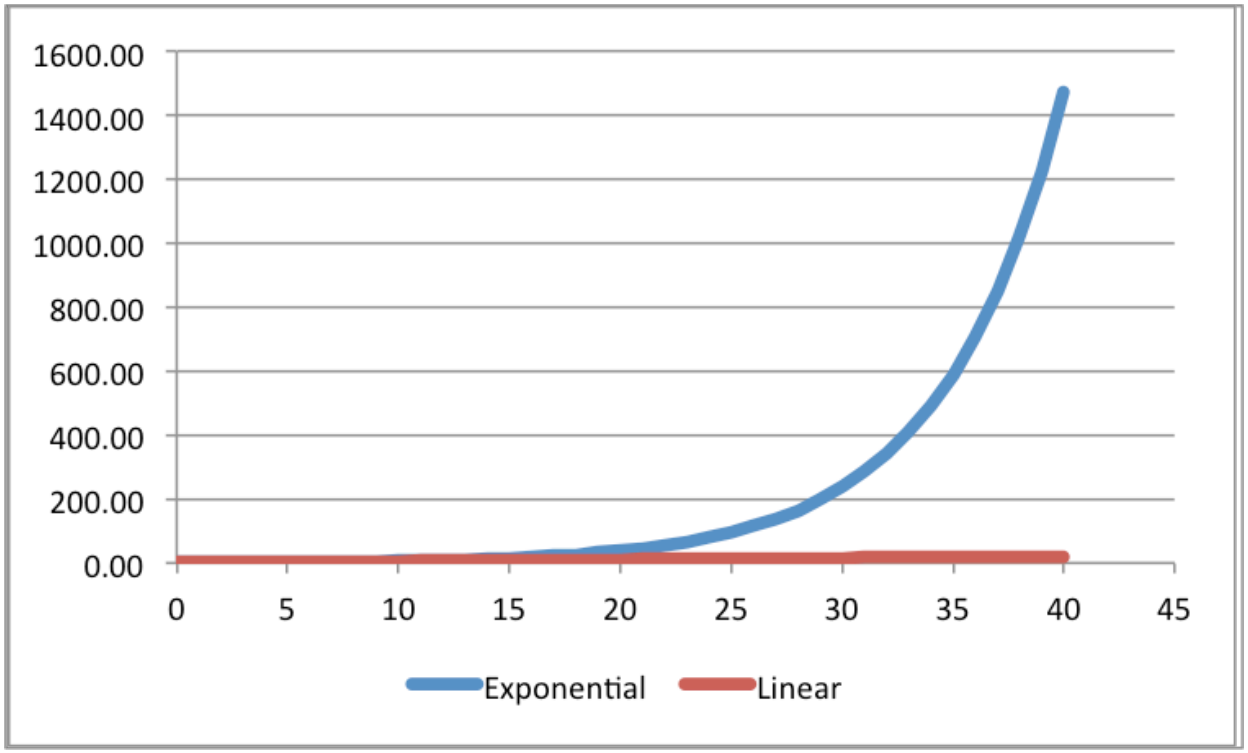
1600.00
1400.00
1200.00
1000.00
800.00
600.00
400.00
200.00
0.00
0
5
10
15
20
25
30
35
40
45
Exponential
Linear

Example: The Number of Transistors

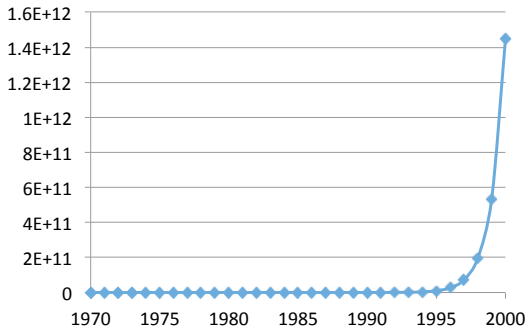

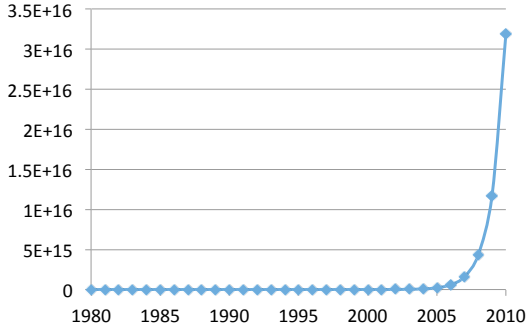

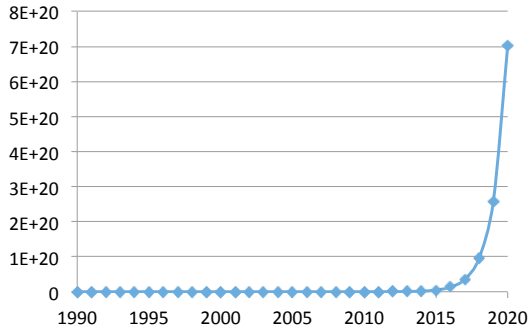

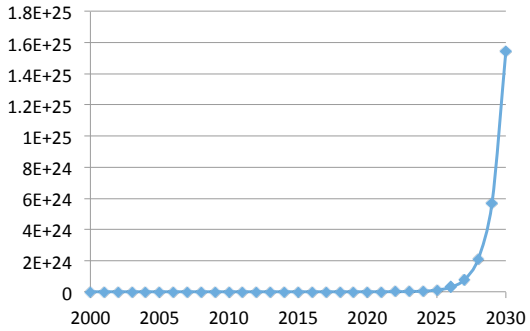

Babbage, Lovelace, Turing, Bardeen, Shockley & Brattain

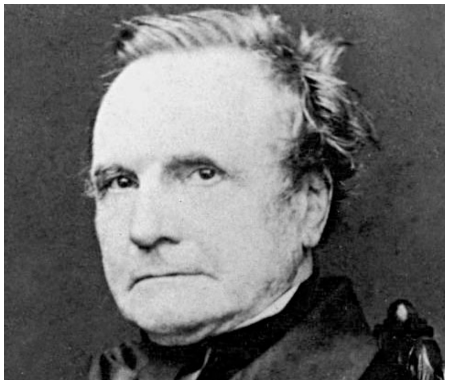

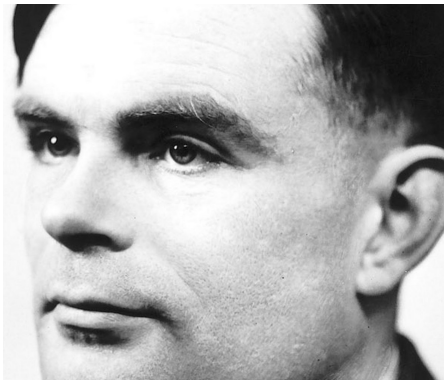


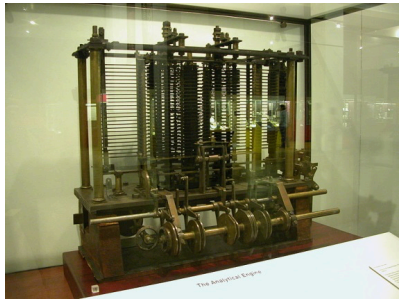

(…) it might act upon other things besides number, were objects found whose mutual fundamental relations could be expressed by those of the abstract science of operations (…)

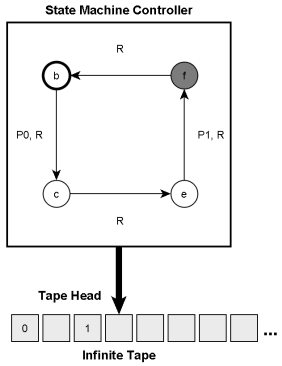

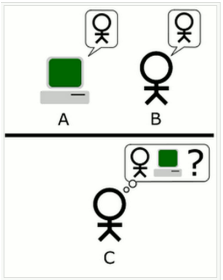

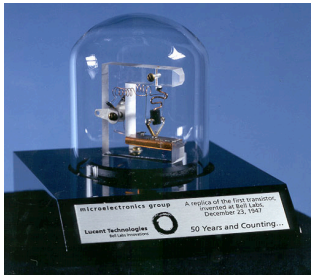

The fourth industrial revolution

- **Artificial Intelligence**
- Cyber-physical systems
- Internet of Things
- Advanced human-machine interfaces
- On-demand availability of computing resources

- 3D printing
- Quantum computing
- Smart sensors
- Big analytics and advanced processes
- Virtual/augmented reality
- Cyber-physical systems
- Wearables
- Location aware systems
- On-demand availability of computer system resources


Modern
Artificial
Intelligence

The Church-Turing Thesis

All forms of computation are equivalent

CPU running a program

Electrical signals

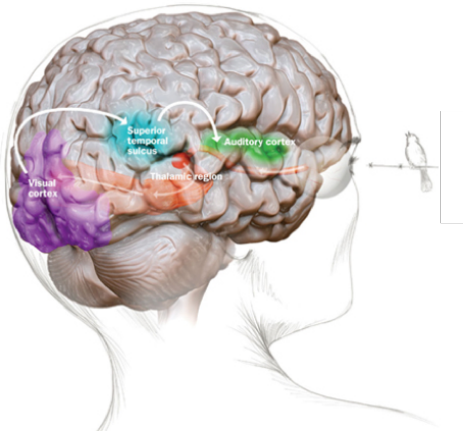

Image processing, reasoning

Electrical and biochemical signals

A Short History of Artificial Intelligence

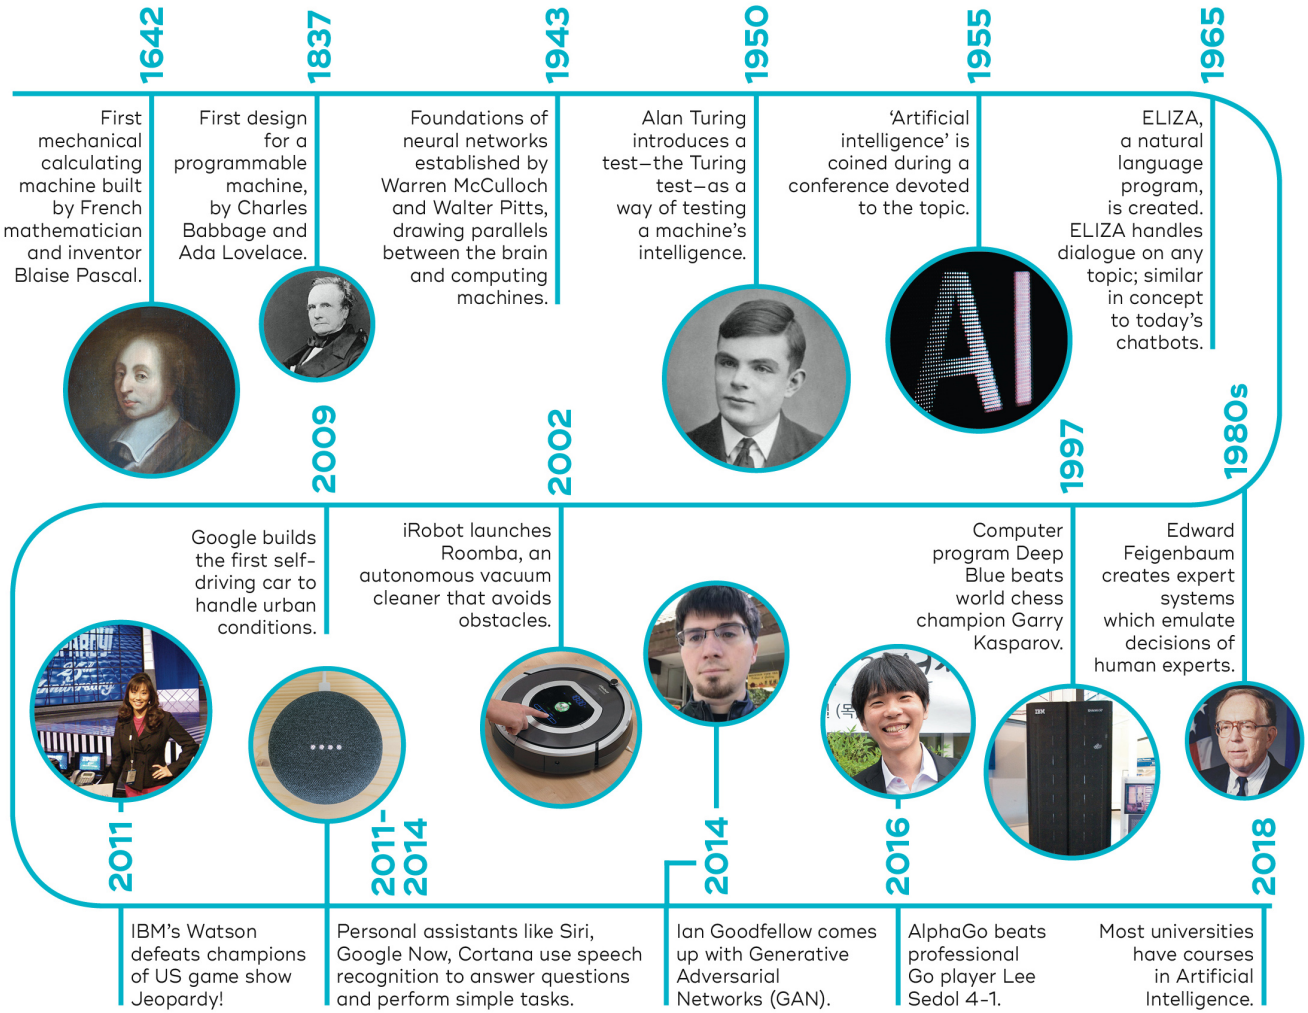

70 years of history

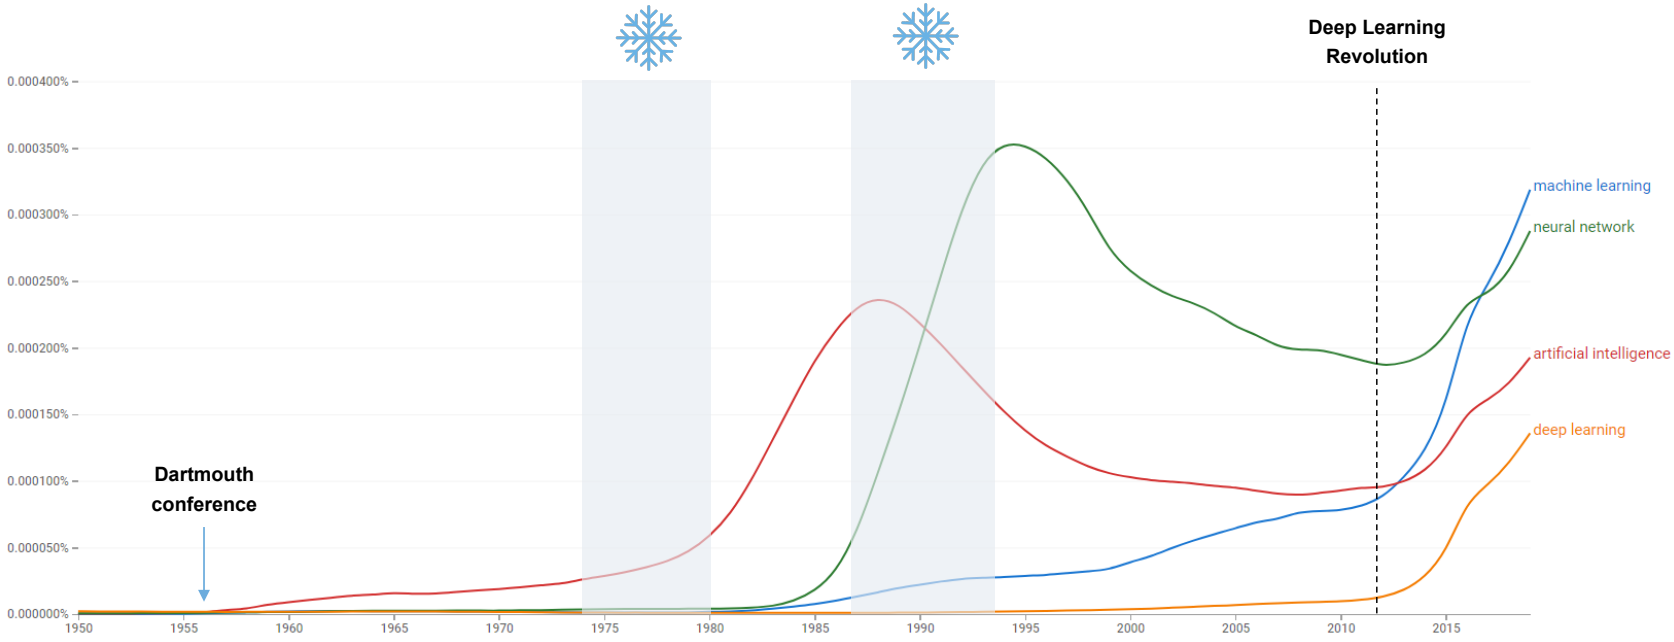

source:
Google Books Ngram Viewer

Moravec's Paradox

Robots and AI systems find the difficult things easy and the easy things difficult

Theorem proving, playing world-class chess, planning routes on a map, scheduling jobs, doing math, are relatively easy to program.

Recognizing faces, understanding scenes, walking in a room, understanding speech, are very, very hard to program.

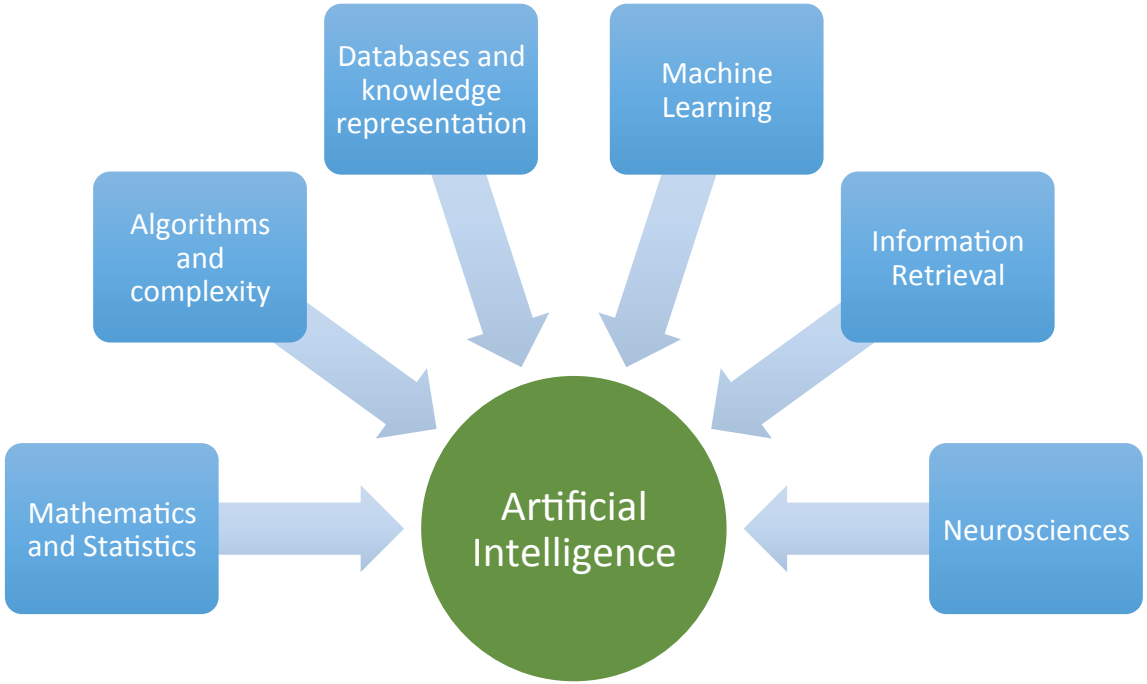
Databases and knowledge representation
Machine Learning
Algorithms and complexity
Information Retrieval
Mathematics and Statistics
Artificial Intelligence
Neurosciences

Artificial Intelligence Sub-Fields

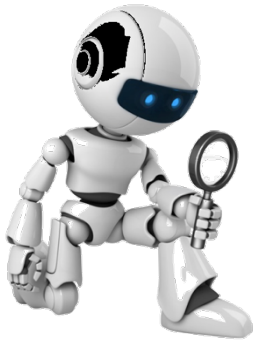

Search

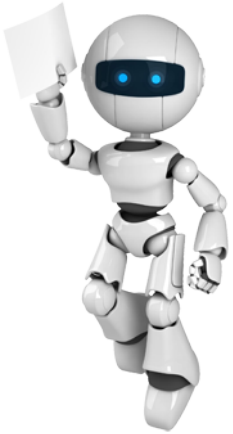

Interaction and perception

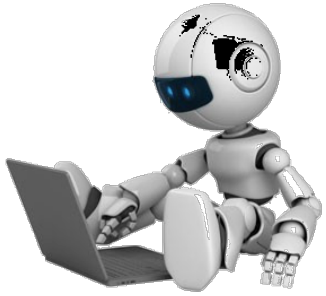

Reasoning

Natural Language Processing

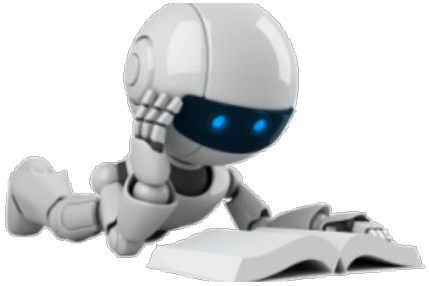

Machine Learning

Artificial Intelligence

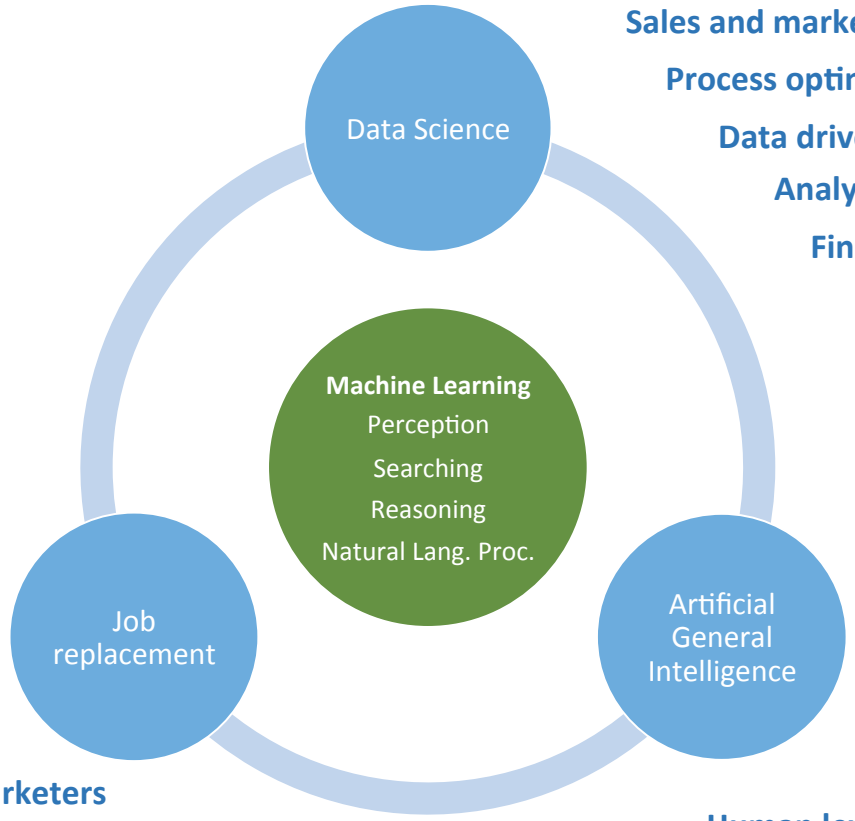

Sales and marketing

Process optimization, Industry 4.0

Data driven science: drugs, laws

Analysis of social behaviors

Financial analysis

Telemarketers

Car and truck drivers

Personal assistants

Legal analysts

Radiologists, Cardiologists, Neurologists

Human level intelligence

Flexibility and adaptability

Open-ended possibilities

Creating 360 degrees views of customers

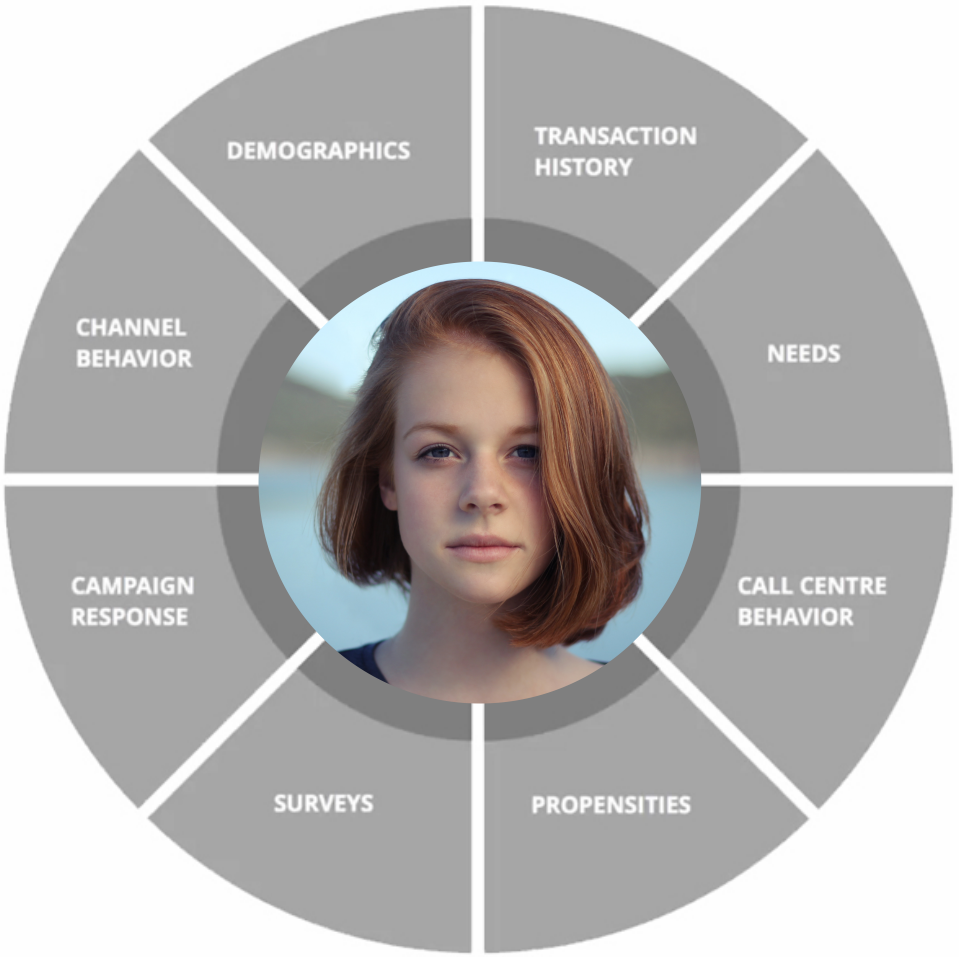

Artificial Intelligence in Medicine

What is Machine Learning?

Adapted from **Samuel (1959):** "*field of study that gives computers the ability to learn without being explicitly programmed.*"

Mitchell (1997): *"a computer program is said to learn if its performance at a task T, as measured by a performance P, improves with experience E"*

Bishop (2006): *"automatic discovery of regularities in data through the use of computer algorithms and with the use of these regularities to take actions"*

Goodfellow, Bengio & Courville (2016): *"the ability to acquire their [AI systems] own knowledge, by extracting patterns from raw data"*

Mohri, Rostamizadeh & Talwalkar (2018): *"computational methods using experience to improve performance or to make accurate predictions"*

Sejnowski (2018): *"learning algorithms are refineries that extract information from raw data; information can be used to create knowledge; knowledge leads to understanding; and understanding leads to wisdom."*

Gerrish (2018): *"field devoted to enabling machines to do smart things by learning from data"*

Efron & Hastie (2019): *"a philosophically atheistic approach to statistical inference"*

Cat or Dog?

dataset

`{cat, dog}`

Beyond *Cat or Dog?*

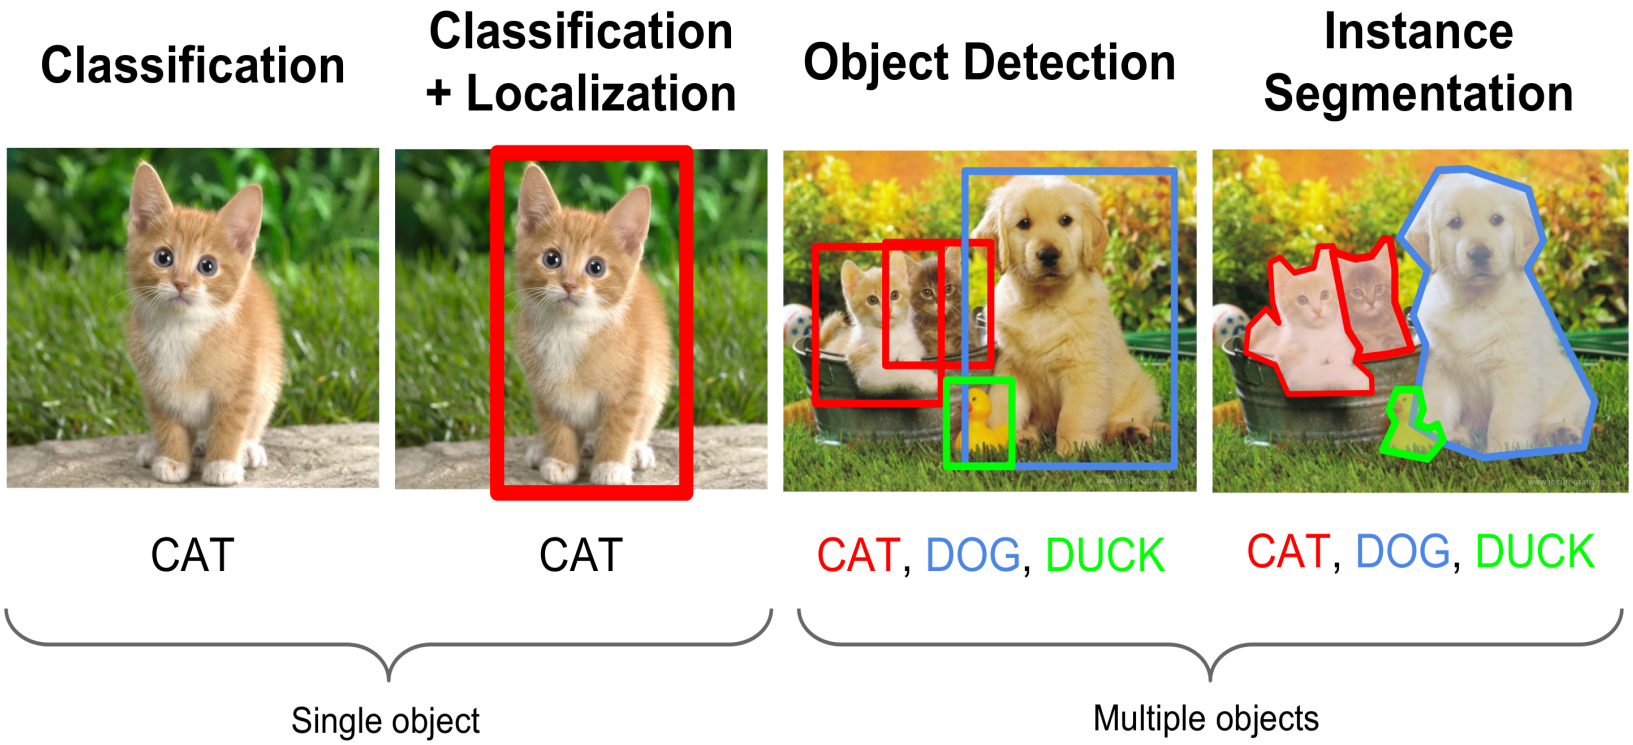

Machine Learning

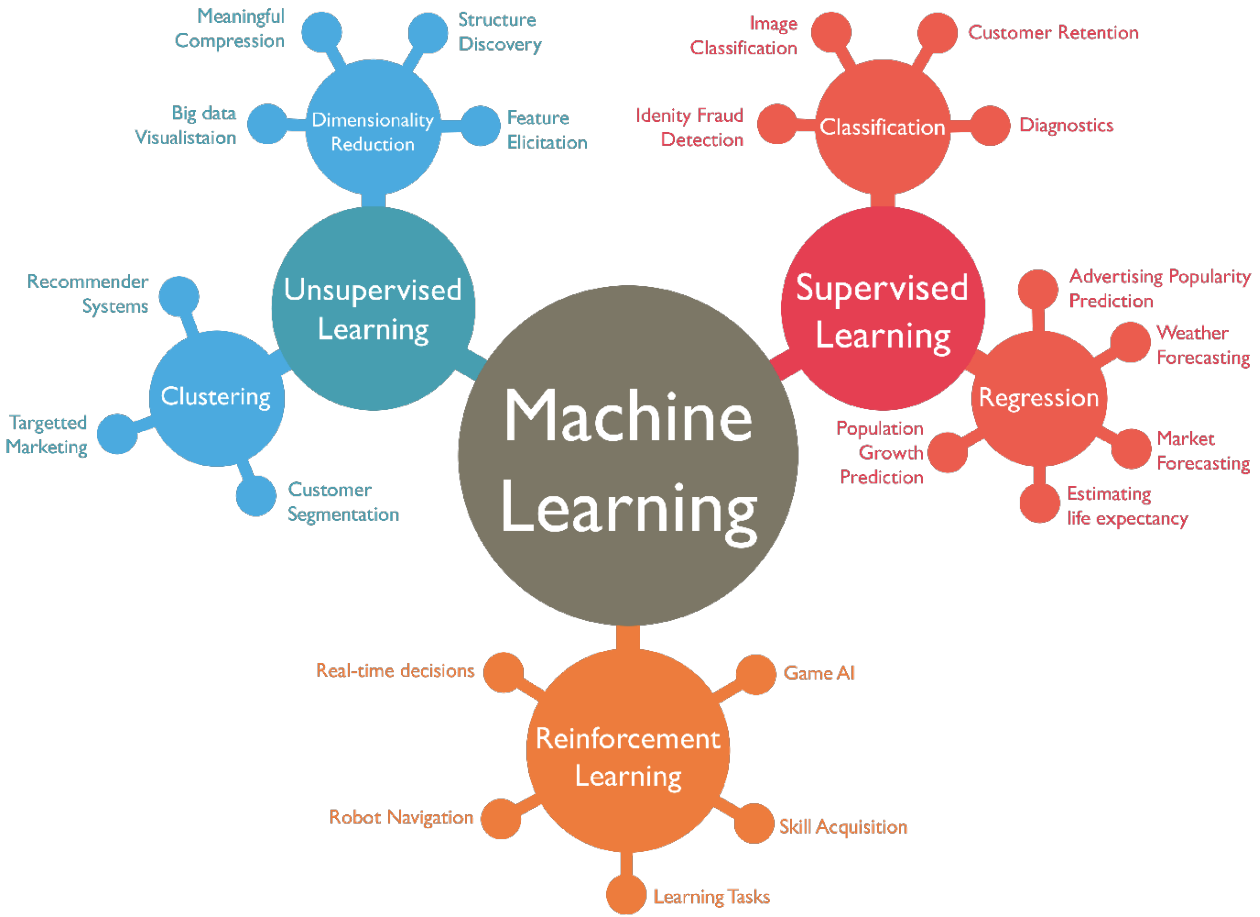

Supervised Learning

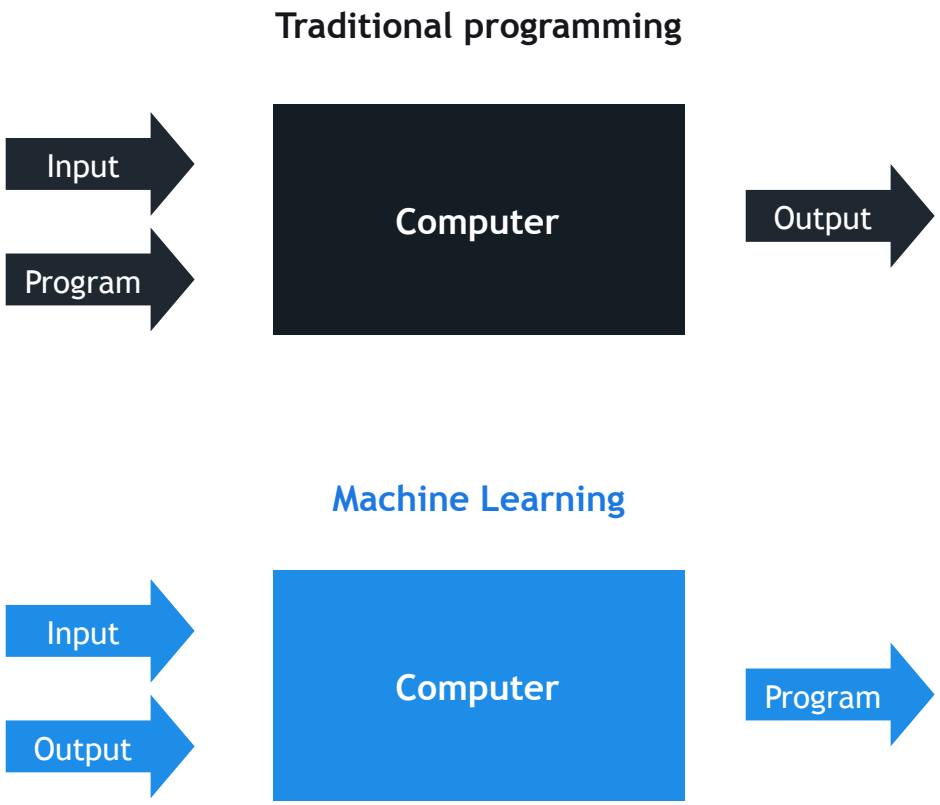

Reinforcement Learning

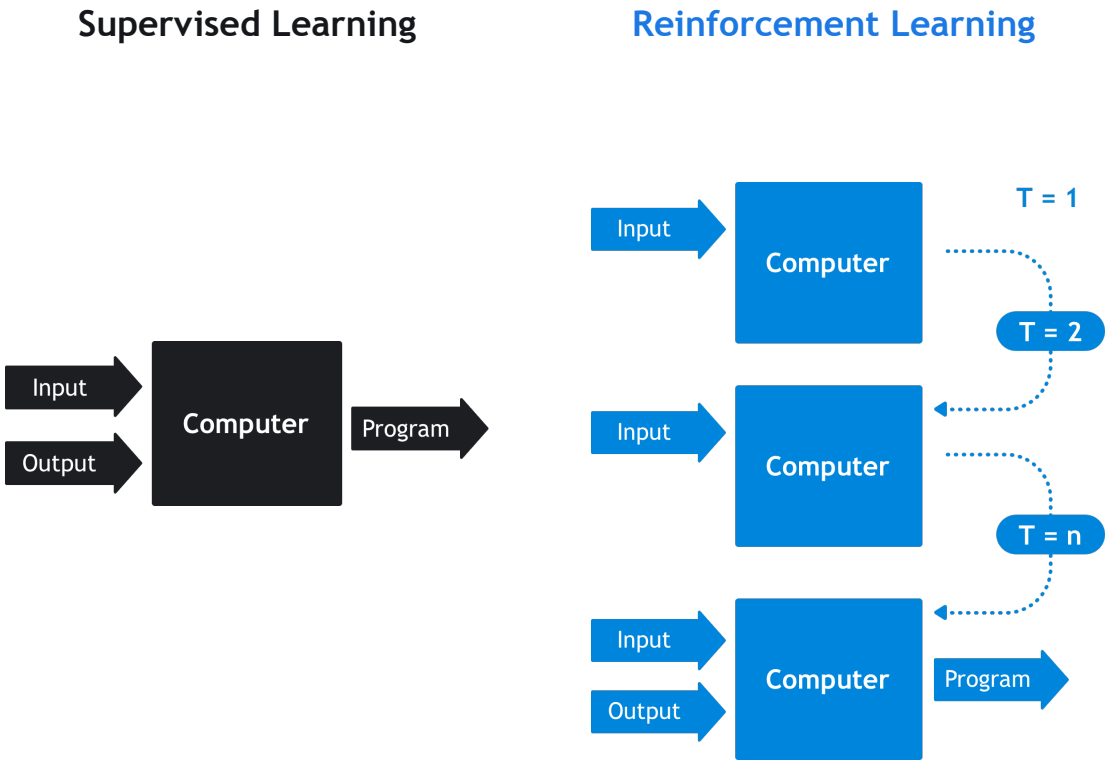

Unsupervised Learning

Supervised Learning

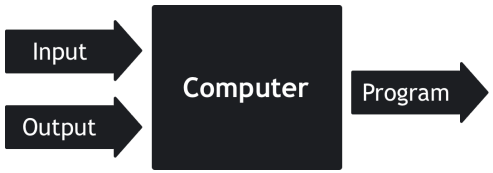

Unsupervised Learning

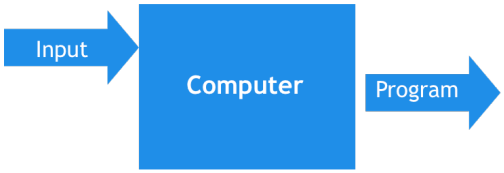

Unsupervised Learning

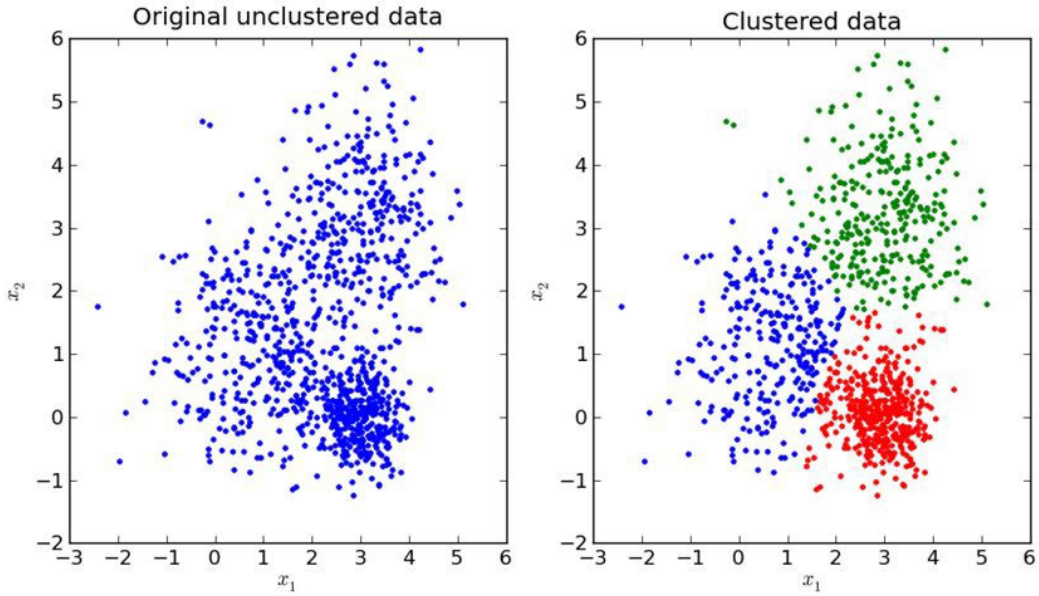

Families of Machine Learning Methods

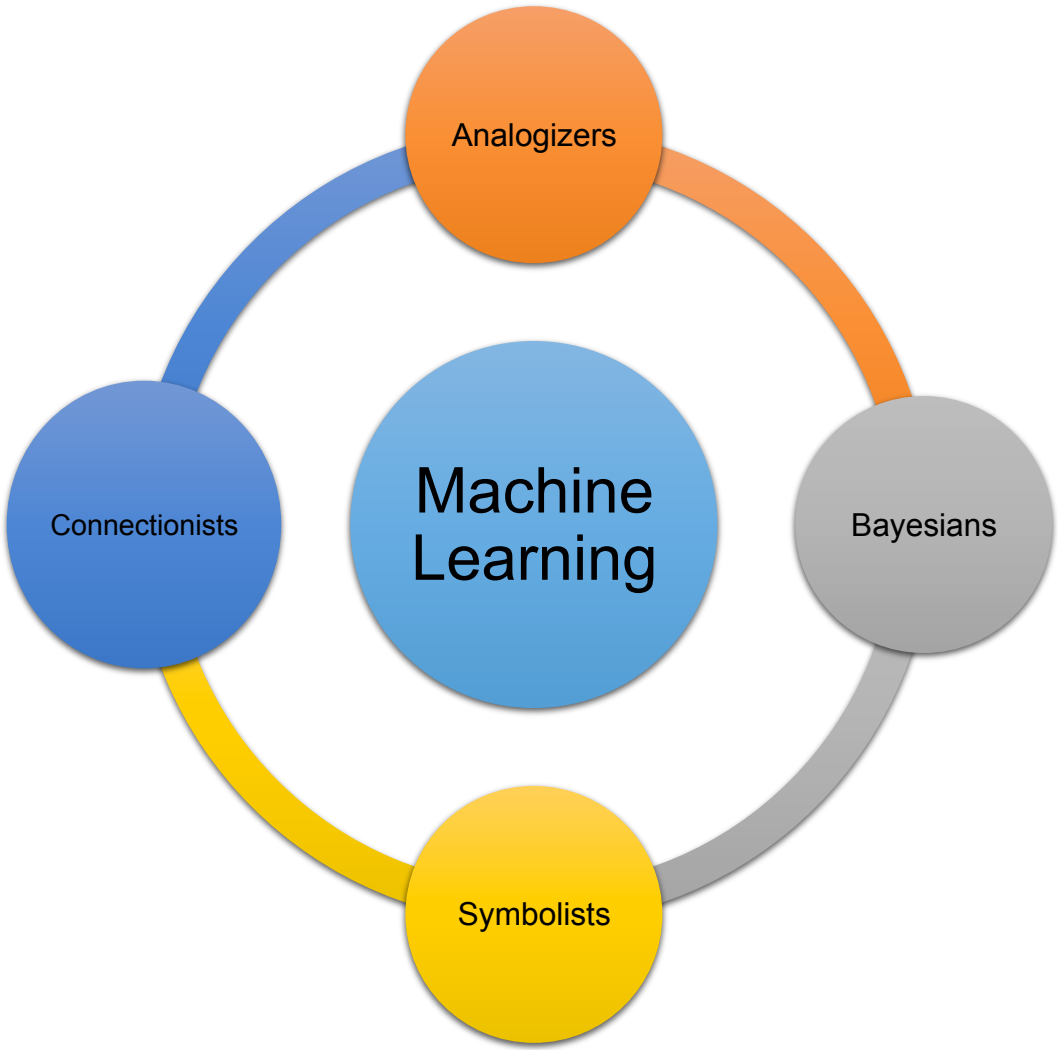

The Four Paradigms

Symbolists

All intelligence can be reduced to symbol manipulation

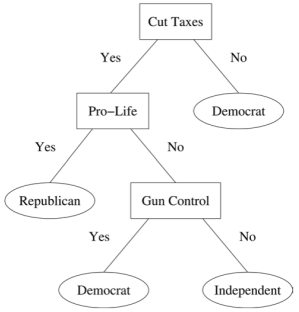

Connectionists

Learning is what the brain does, and what we need to do is reverse engineer it

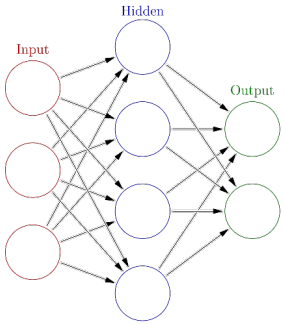

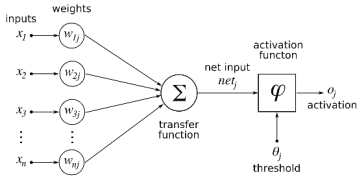

Bayesians

All learned knowledge is uncertain, and learning itself is a form of uncertain inference

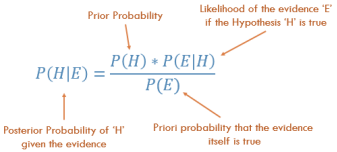

Analogizers

The key to learning is recognizing similarities between situations

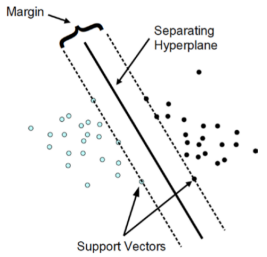

The Five Tribes of Machine Learning

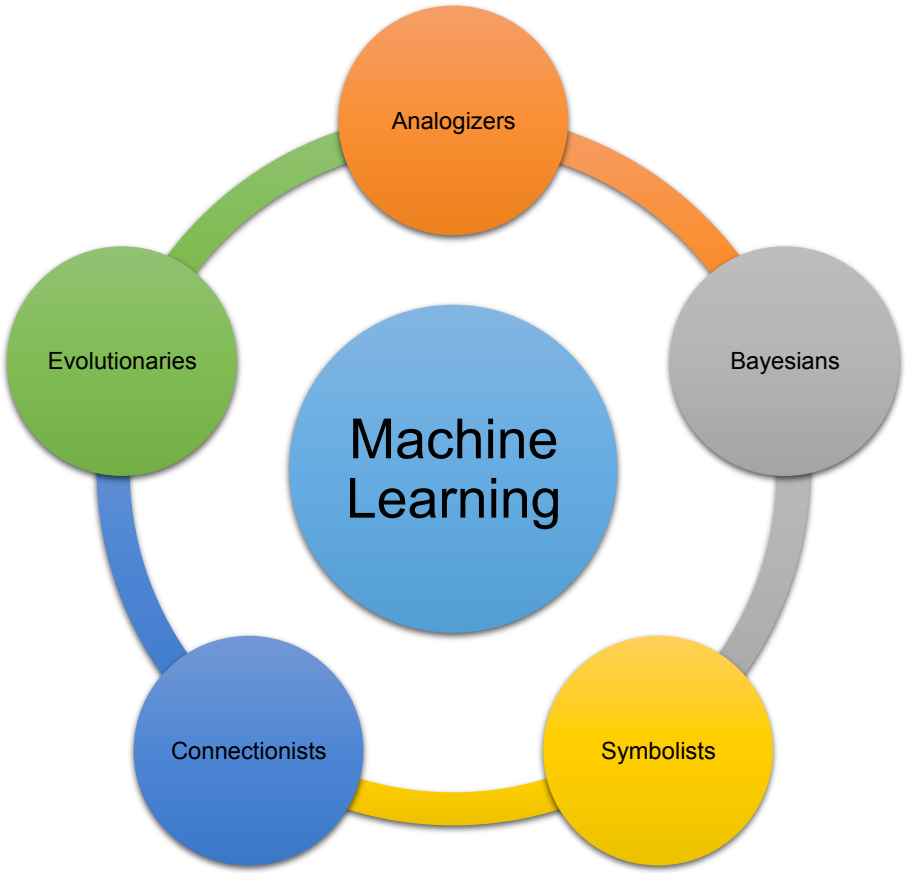

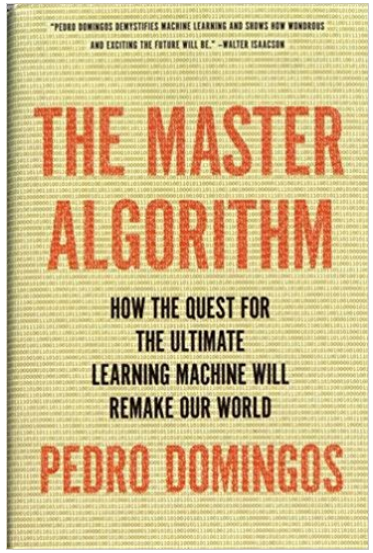

Neural Networks

Artificial
Neuron-Perceptron

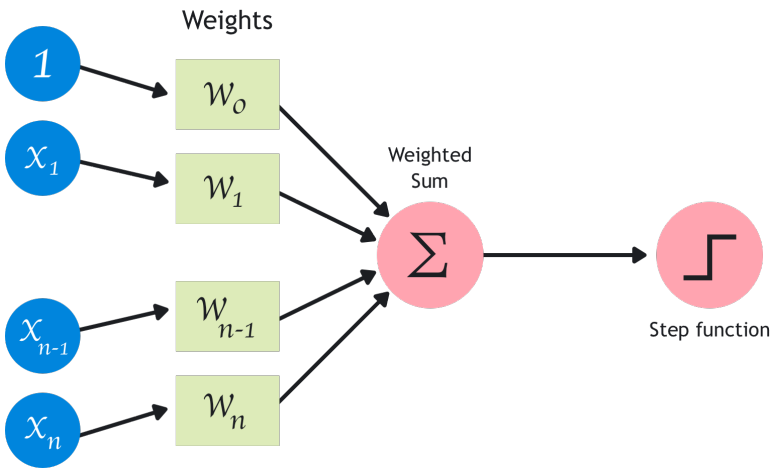

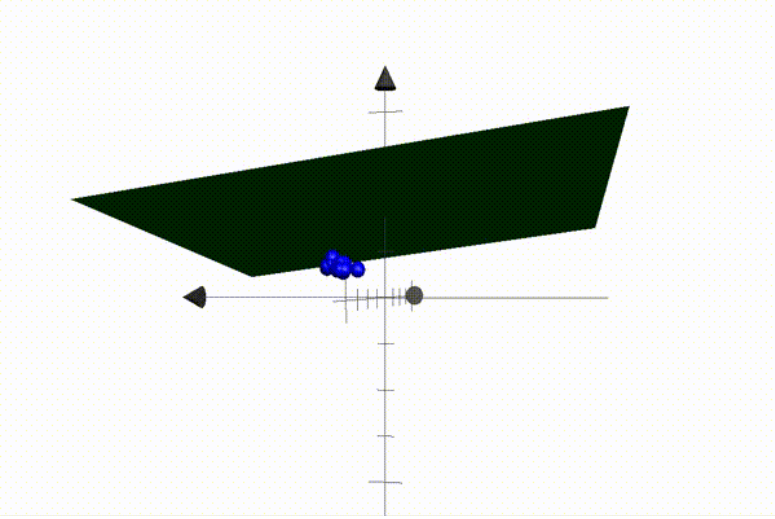

Neural Networks

Multi Layer Perceptron

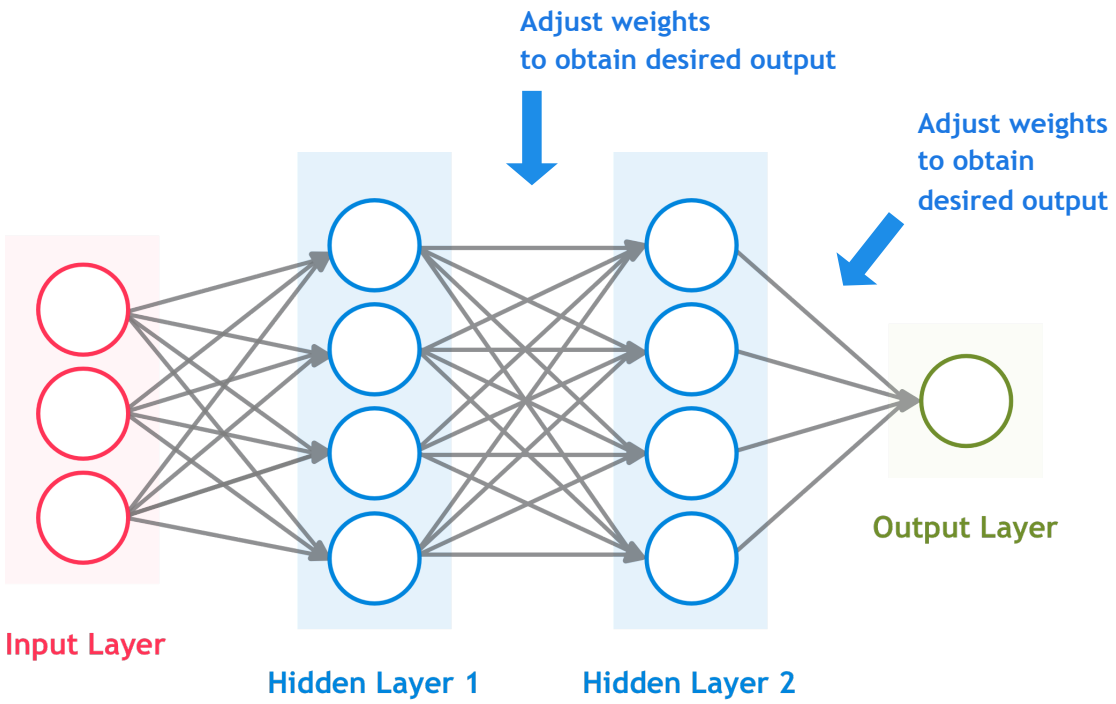

Deep Neural Networks

Deep Neural Networks learn hierarchical feature representations

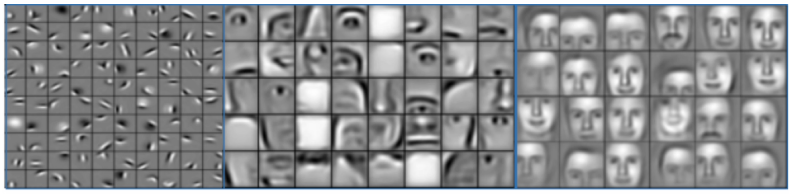

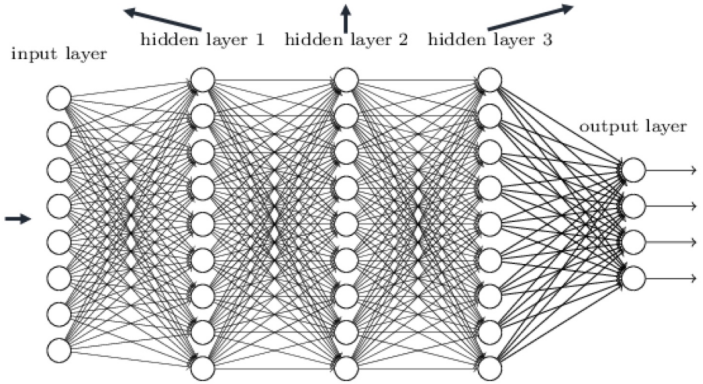

Deep Convolutional Neural Networks

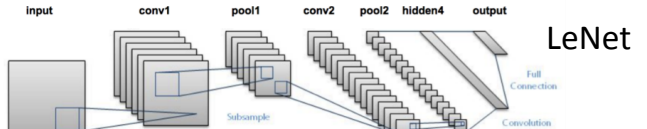

LeNet

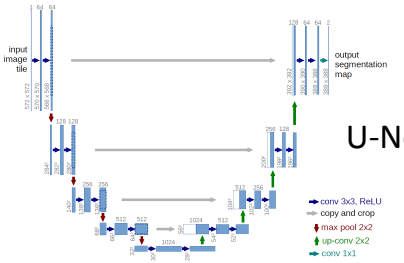

U-Net

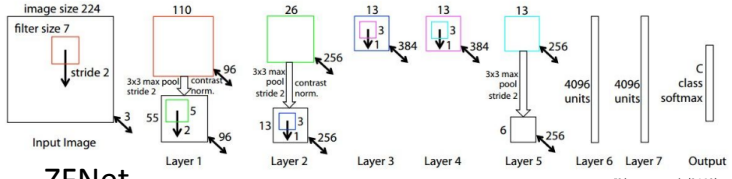

ZFNet

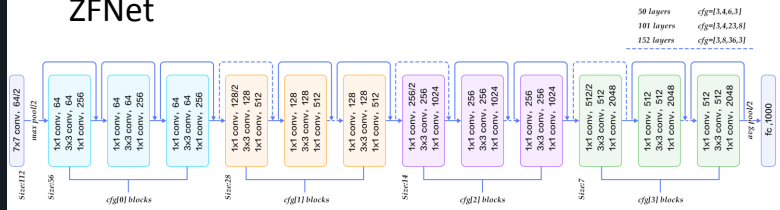

ResNet

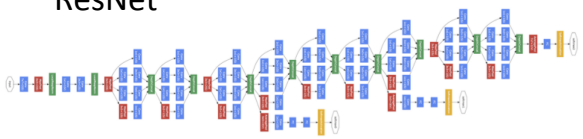

GoogleNet/Inception

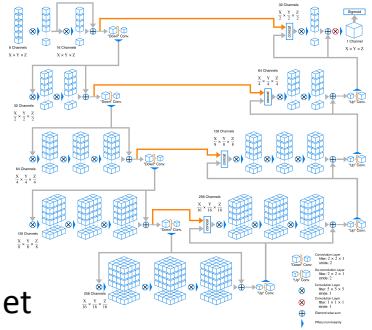

V-Net

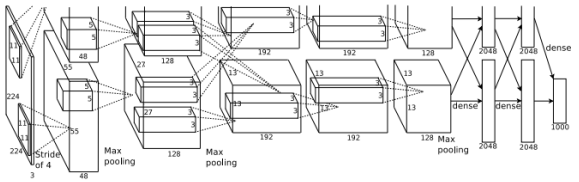

AlexNet

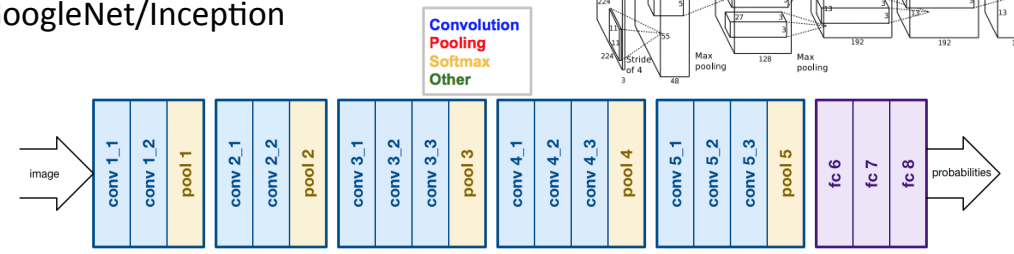

VGGNet

Generative Adversarial Networks

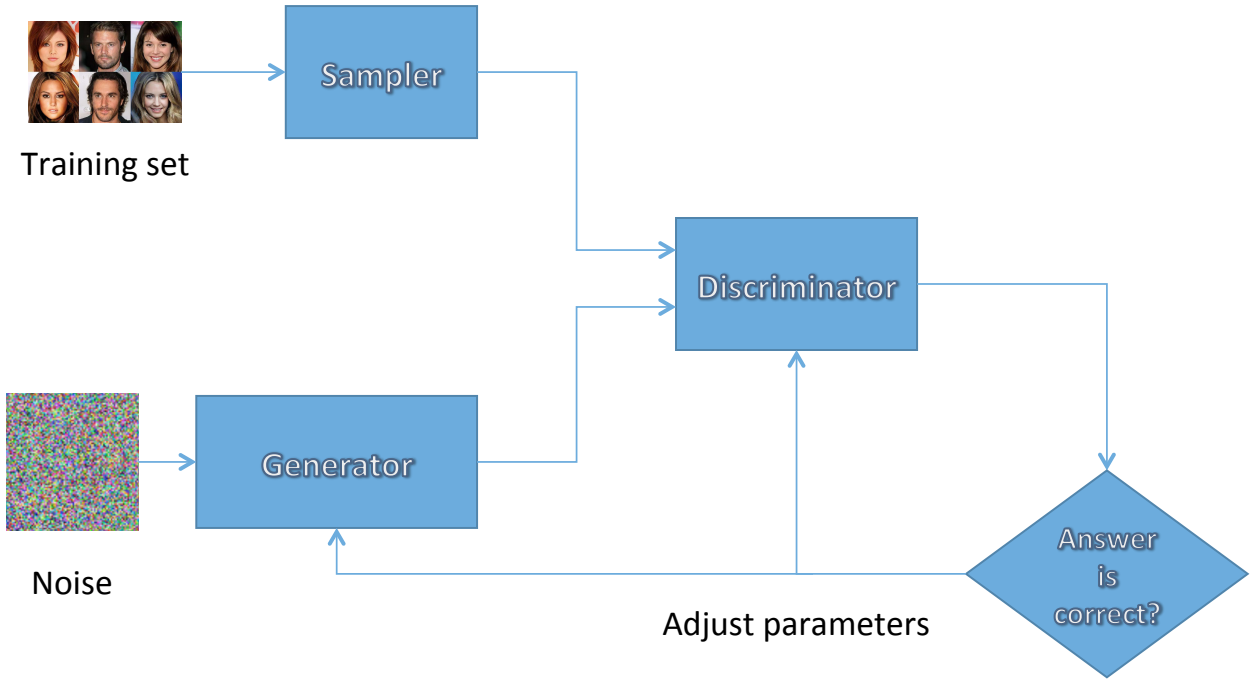

Some of these people do not exist!

Karras et a. “A Style-Based Generator Architecture for Generative Adversarial Networks”, CVPR 2019
Goodfellow et al. “Generative Adversarial Nets”, NIPS, 2014

ML Use-Cases

IBM Watson

Deep Fake


Deep RL and Atari 2600


Autonomous Driving

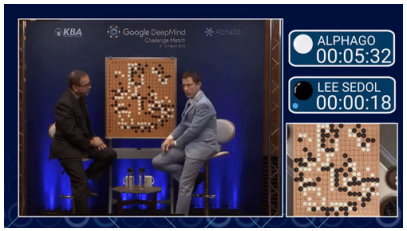

AlphaGo

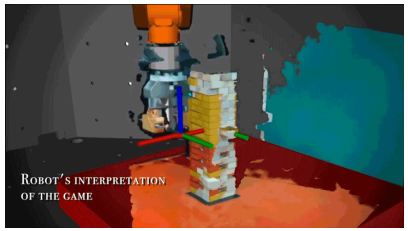

MIT Jenga

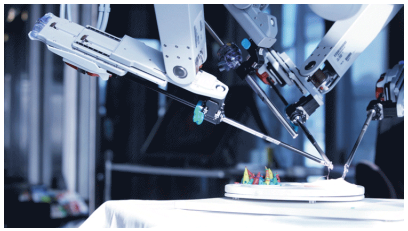
Robot Surgery

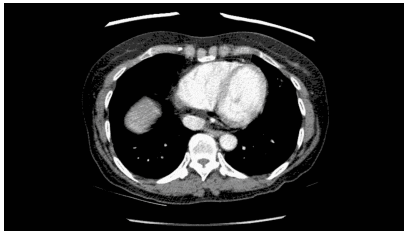
Chest CT Interpretation

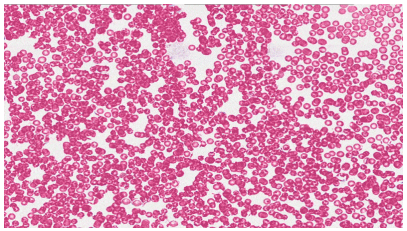
Skin Image Analysis

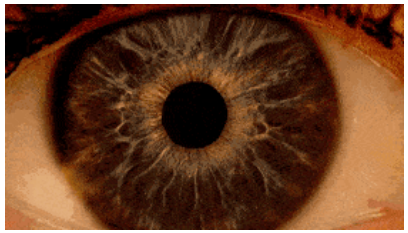
Diabetic Retinopathy Diagnosis

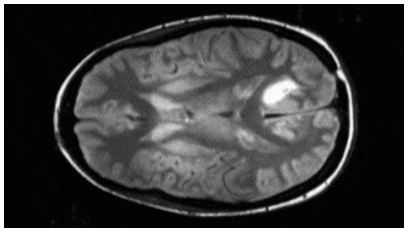
Brain Tumor Detection


Brain-Computer Interfaces

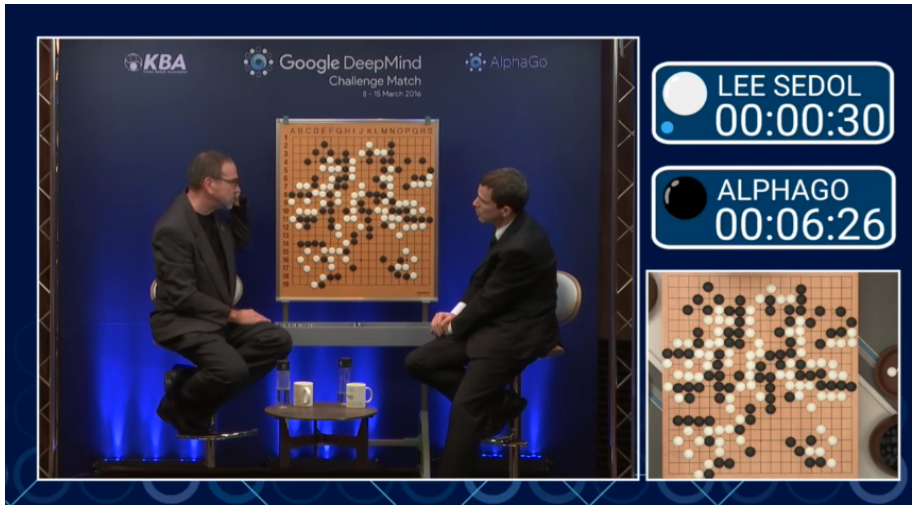
KBA
Google DeepMind
Challenge Match
8 - 15 March 2016
AlphaGo
LEE SEDOL
00:00:30
ALPHAGO
00:06:26


The Economist

Objections

- **The immortal soul objection:**
  Man is the sole owner of an immortal soul, and no machine can ever have one.

- **The consciousness objection:**
  Only humans can be conscious.

- **The technological impossibility objection:**
  It is outside the reach of human technology, forever.

- **The physical support objection:**
  Nervous systems are continuous, use quantum mechanics to function, are chaotic

- **The mathematical objection:**
  Brains are not Turing equivalent

- **Various disabilities objection**
  Cannot create anything new
  Cannot love or be loved
  Cannot feel pain or joy

Legal and Social Challenges

- "whereas, ultimately, robots' autonomy raises the question of their nature in the light of the existing legal categories – of whether **they should be regarded as natural persons, legal persons, animals or objects – or whether a new category should be created**"

- "creating a **specific legal status for robots**, so that at least the most sophisticated autonomous robots could be established as having the status of electronic persons with specific rights and obligations…"

- "consideration should be given to the possible need to introduce corporate reporting requirements on the extent and proportion of the contribution of robotics and AI to the economic results of a company for the purpose of taxation and social security contributions; takes the view that **in the light of the possible effects on the labour market of robotics and AI a general basic income should be seriously considered**, and invites all Member States to do so;"

- "a system of registration of advanced robots should be introduced, and calls on the Commission to establish criteria for the classification of robots with a view to **identifying the robots that would need to be registered;**"

(Report from the European Parliament 2016)

The challenge: inequality

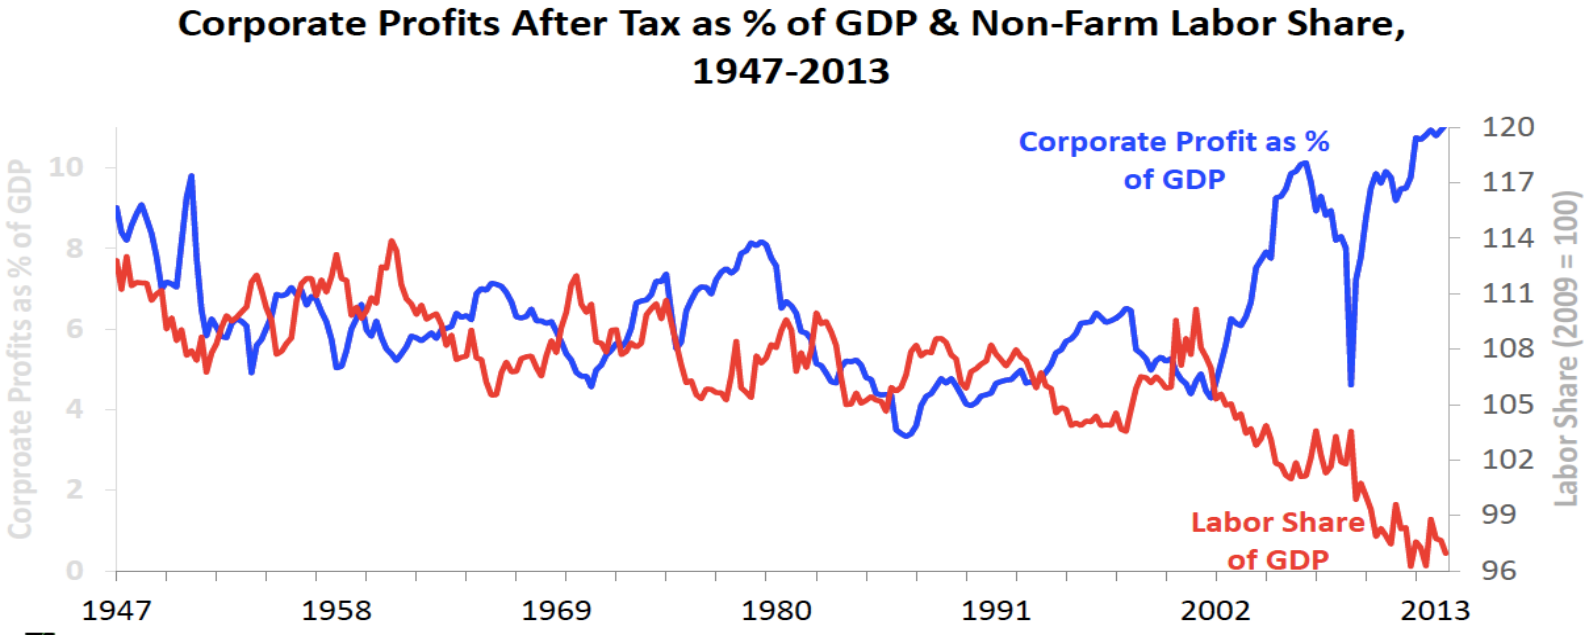

Source: Brynjolfsson & McAfee

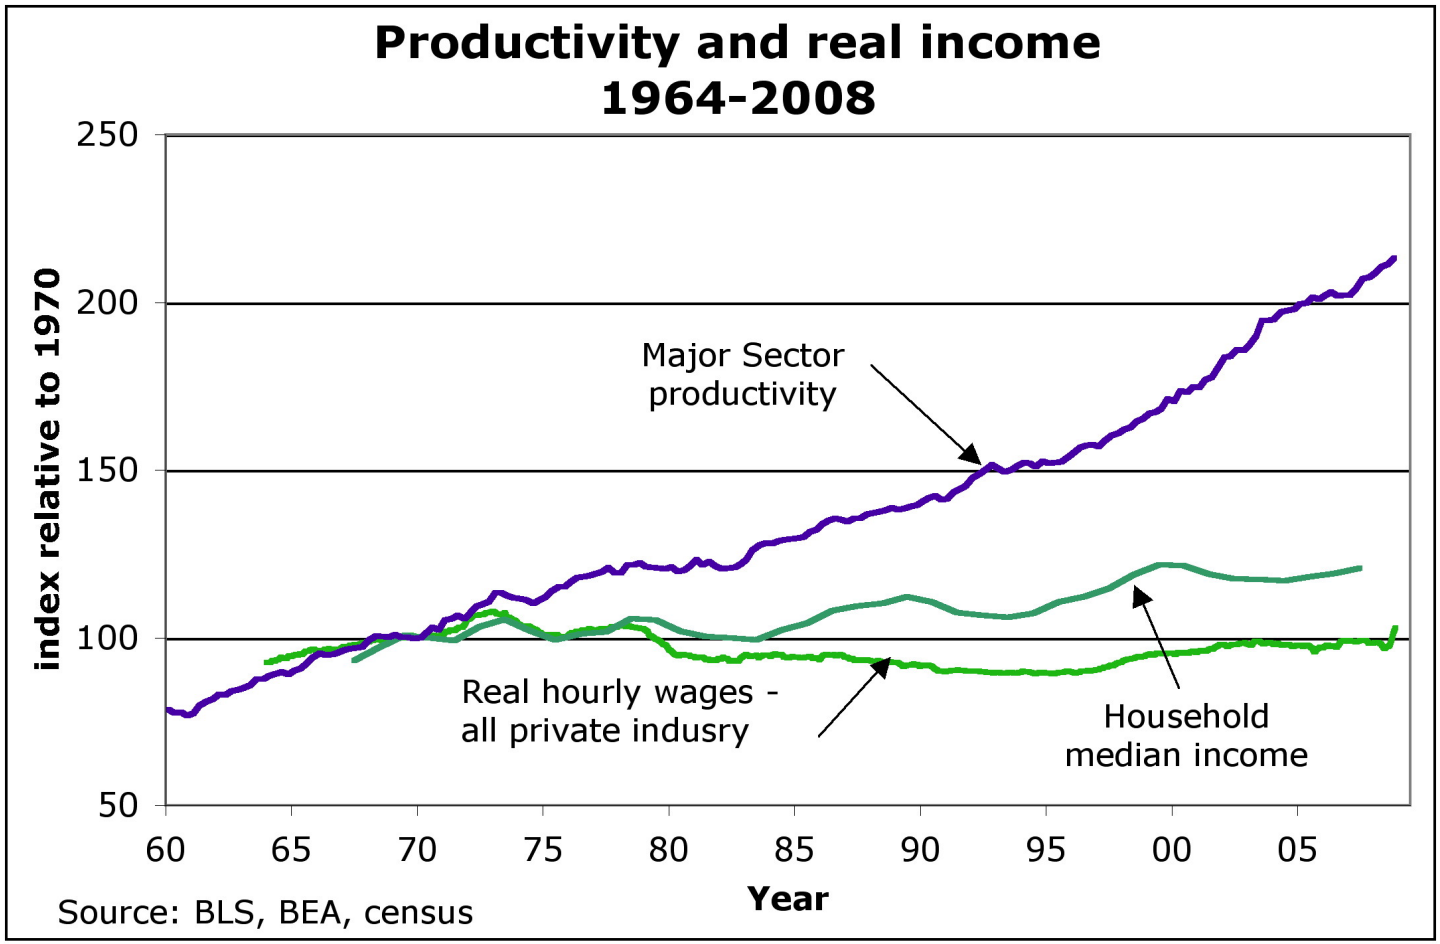
Productivity and real income
1964-2008
index relative to 1970
250
200
150
100
50
60
65
70
75
80
85
90
95
00
05
Year
Major Sector
productivity
Real hourly wages -
all private indusry
Household
median income
Source: BLS, BEA, census

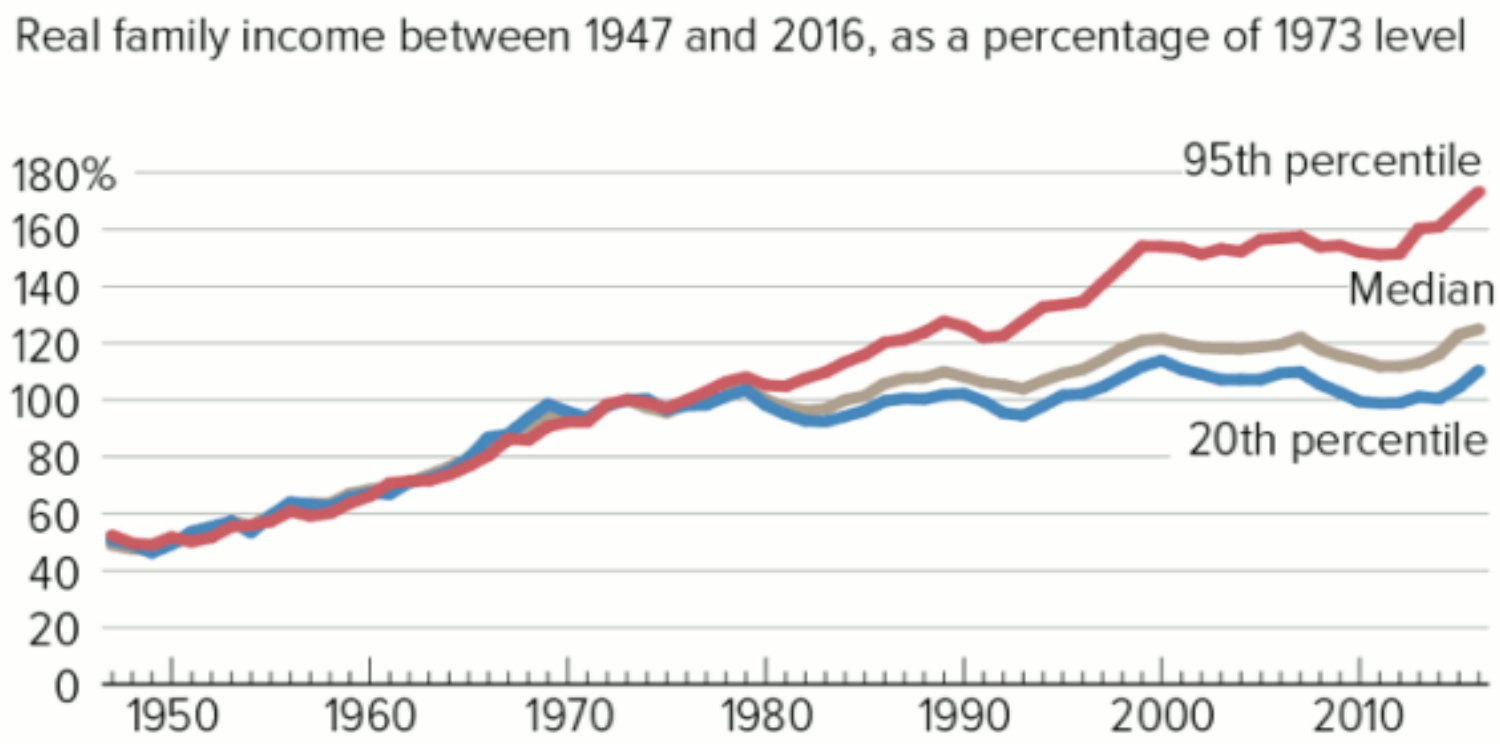

Source: Center on Budget and Policy Priorities

Ethical Challenges

- “the question of granting personhood to an AI machine or robot depends on where the line is drawn between persons and inanimate objects”

- “The overarching criterion is displaying some form of cognitive capacity—being conscious, having perceptions, feeling sensations”

- “While the possibility of a machine being granted moral status is the most compelling ethical issue raised by AI, there are others, determined largely by the uses to which AI programs are actually put. These ethical considerations have evolved as AI research and development has progressed. AI programs form relationships with other entities. They are used, for example, to advise human users, make decisions, and in the case of intelligent software agents to chat with people, search for information, look for news, find jobs, and shop for goods and locate the best prices. Their role in these relationships engenders moral responsibility. “

(Mason 2003)

Ethical Challenges

- “Certainly, the technologies underlying the weapons of mass destruction (WMD)—nuclear, biological, and chemical (NBC)—were powerful, and the weapons an enormous threat. But building nuclear weapons required, at least for a time, access to both rare—indeed, effectively unavailable—raw materials and highly protected information; biological and chemical weapons programs also tended to require large-scale activities.”

- “The 21st-century technologies—genetics, nanotechnology, and robotics (GNR)—are so powerful that they can spawn whole new classes of accidents and abuses. Most dangerously, for the first time, these accidents and abuses are widely within the reach of individuals or small groups. They will not require large facilities or rare raw materials. Knowledge alone will enable the use of them.”

- “this is the first moment in the history of our planet when any species by its voluntary actions has become a danger to itself.”

- “The only realistic alternative I see is relinquishment: to limit development of the technologies that are too dangerous, by limiting our pursuit of certain kinds of knowledge.”

(Joy 2000)

Speculations

"The ever-accelerating progress of technology and changes in the mode of human life, which gives the appearance of approaching some essential **singularity** in the history of the race beyond which human affairs, as we know them, could not continue"

- Von Neumman, quoted by Stanislaw Ulam, 1952

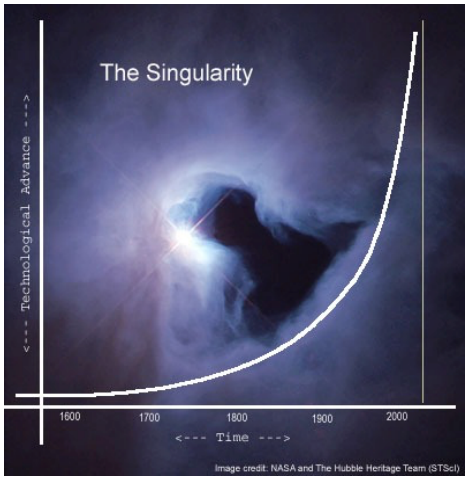

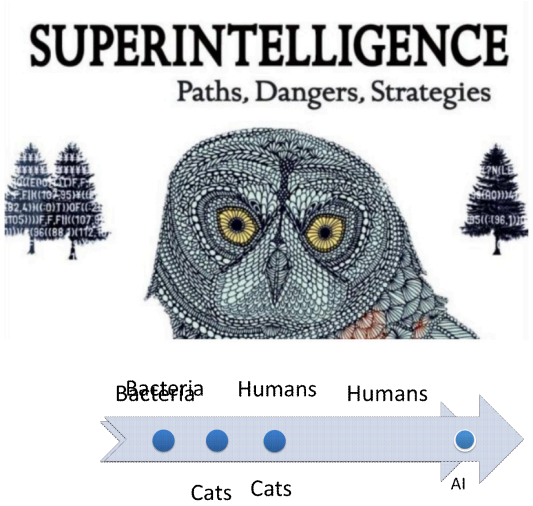

Let an ultraintelligent machine be defined as a machine that can far surpass all the intellectual activities of any man however clever. Since the design of machines is one of these intellectual activities, an ultraintelligent machine could design even better machines; there would then unquestionably be an 'intelligence explosion', and the intelligence of man would be left far behind. Thus the first ultraintelligent machine is the **last invention** that man need ever make, provided that the machine is docile enough to tell us how to keep it under control.

- Irving John Good,1965

The Future of Intelligence in the Universe

- Where are they? – was the question of Enrico Fermi, about why we have never seen any evidence of extraterrestrial intelligence, the so-called **Fermi Paradox**.

- If we are alone in the galaxy, then there is a **Great Filter** in the past or in the future.

- **We are first, we are rare, or we are doomed**.

- **We have a moral duty to spread intelligence through the galaxy** but the human body is terribly mis-adapted to space travel, even less to long-distance space travel.

- Intelligent systems may be the only way human intelligence can spread trough the galaxy. Maybe the future of humanity lies in **AI systems, dispersing intelligence throughout the galaxy**.

- Maybe we will create our own virtual universes, converting everything into computers (computronium), which will run powerful **simulations of virtual worlds and universes**.

- **Omega point**: a point where we will control the computational power of the universe, becoming God.

The Future of Intelligence