Artificial Intelligence

Augustin Cosse.

Fall 2020

September 9, 2020

Historical perspective. Computer Science (1940 - Present)

- Most critical breakthrough was the IBM 701 built in 1952 by Nathaniel Rochester and his group
- This was the first computer to yield a profit for its manufacturer. IBM went on to become one of the world's largest corporations, and sales of computers have grown to \$150 billion/year
- In the United States, the computer industry (including software and services) now accounts for about 10% of the gross national product.

Historical perspective. Computer Science (1940 - Present)

- Each generation of computer hardware has brought an increase in speed and capacity, and a decrease in price.
- Computer engineering has been remarkably successful, regularly doubling performance every two years, with no immediate end in sight for this rate of increase.
- Massively parallel machines promise to add several more zeros to the overall throughput achievable

Historical perspective. Computer Science (1940 - Present)

- Of course, there were calculating devices before the electronic computer. The abacus is roughly 7000 years old. In the mid-17th century, Blaise Pascal built a mechanical adding and subtracting machine called the Pascaline.
- Leibniz improved on this in 1694. building a mechanical device that multiplied by doing repeated addition.
- Progress stalled for over a century until Charles Babbage (1792-1871) dreamed that logarithm tables could be computed by machine. He designed a machine for this task, but never completed the project. Instead, he turned to the design of the Analytical Engine, for which Babbage invented the ideas of addressable memory. stored programs, and conditional jumps

Rochester

Pascal

Leibnitz

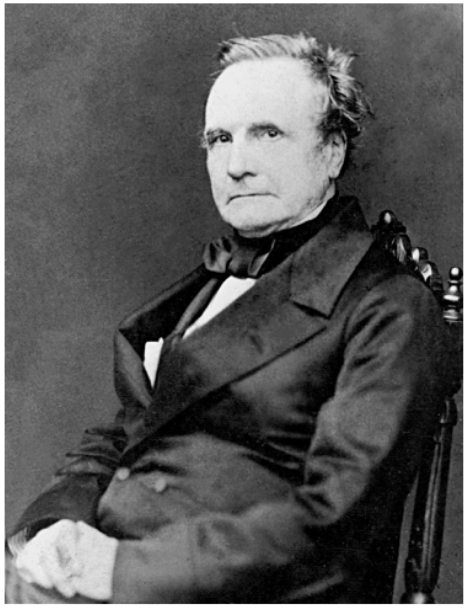

Babbage

Jacquard

Historical perspective. Computer Science (1940 - Present)

- Although the idea of programmable machines was not new - in 1805. Joseph Marie Jacquard invented a loom that could be programmed using punched cards - Babbage's machine was the first artifact possessing the characteristics necessary for universal computation
- Babbage's colleague Ada Lovelace, daughter of the poet Lord Byron, wrote programs for the Analytical Engine and even speculated that the machine could play chess or compose music. Lovelace was the world's first programmer, and the first of many to endure massive cost overruns and to have an ambitious project ultimately abandoned.

Historical perspective. Computer Science (1940 - Present)

- Babbage had the right idea, but lacked the organizational skills to get his machine built

Ada Lovelace

Byron

Historical perspective. Linguistics (1957 - Present)

- In 1957. B. F. Skinner published Verbal Behavior. This was a comprehensive, detailed account of the behaviorist approach to language learning
- Curiously, a review of the book became as well-known as the book itself and nearly the rising field of behaviorism. The author of the review was Noam Chomsky, who had just published another book *Syntactic Structure*
- Chomsky showed how the behaviorist theory did not address the notion of creativity in language - it did not explain how a child could understand and make up sentences that he or she had never heard before. Chomsky's theory explained this based on syntactic models going back to the Indian linguist Panini (c. 350 B.C.)

Historical perspective. Linguistics (1957 - Present)

- In 1957. B. F. Skinner published Verbal Behavior. This was a comprehensive, detailed account of the behaviorist approach to language learning
- Curiously, a review of the book became as well-known as the book itself and nearly the rising field of behaviorism. The author of the review was Noam Chomsky, who had just published another book *Syntactic Structure*
- Modern linguistics and AI were "born" at about the same time, so linguistics does not play a large foundational role in the growth of AI. Instead, the two grew up together, intersecting in a hybrid field called computational linguistics or natural language processing

Historical perspective. Linguistics (1957 - Present)

- Much of the early work in knowledge representation (the study of how to put knowledge into a form that a computer can reason with) was tied to language and informed by research in linguistics

History of AI

- The first work that is now generally recognized as AI was done by Warren McCulloch and Walter Pitts (1943). The work of McCulloch and Pitts was based on three sources:
  - knowledge of the basic physiology and function of neurons in the brain
  - Formal analysis of propositional logic due to Russell and Whitehead
  - Turing's theory of computation
- They proposed a model of neurons in which each neuron is either "on" or "off". A neuron being switched to "on" when being stimulated by a sufficient number of neighbors.

History of AI

- McCulloch and Pitts showed that any computable function can be computed by some network of connected neurons.
- The work of McCulloch and Pitts can be considered as the forerunnner of both the logicist tradition and the connectionist traduction.
  - Logicist AI (a.k.a Symbolic AI): Knowledge can be stored as sentences in a formal language (e.g. first order logic). Mind/Reasoning of agent = digital computer processing a symbolic language $\approx$ Theorems proving.
  - Connectionist AI: information in the brain is stored through the connection length and strength in neurons)

History of AI

- In the 1950's both Turing and Shannon were writing chess programs for von-Neumann stype conventional computers.
- At Princeton, Marvin Minsky and Dean Edmonds built the first neural network computer in 1951. The computer named SNARC (Stochastic Neural Analog Reinforcement Calculator) was used to simulate 40 neurons
- Minsky's Ph.D. committee was skeptical whether this kind of work should be considered mathematics, but von Neumann was on the committee and reportedly said, "If it isn't now it will be someday

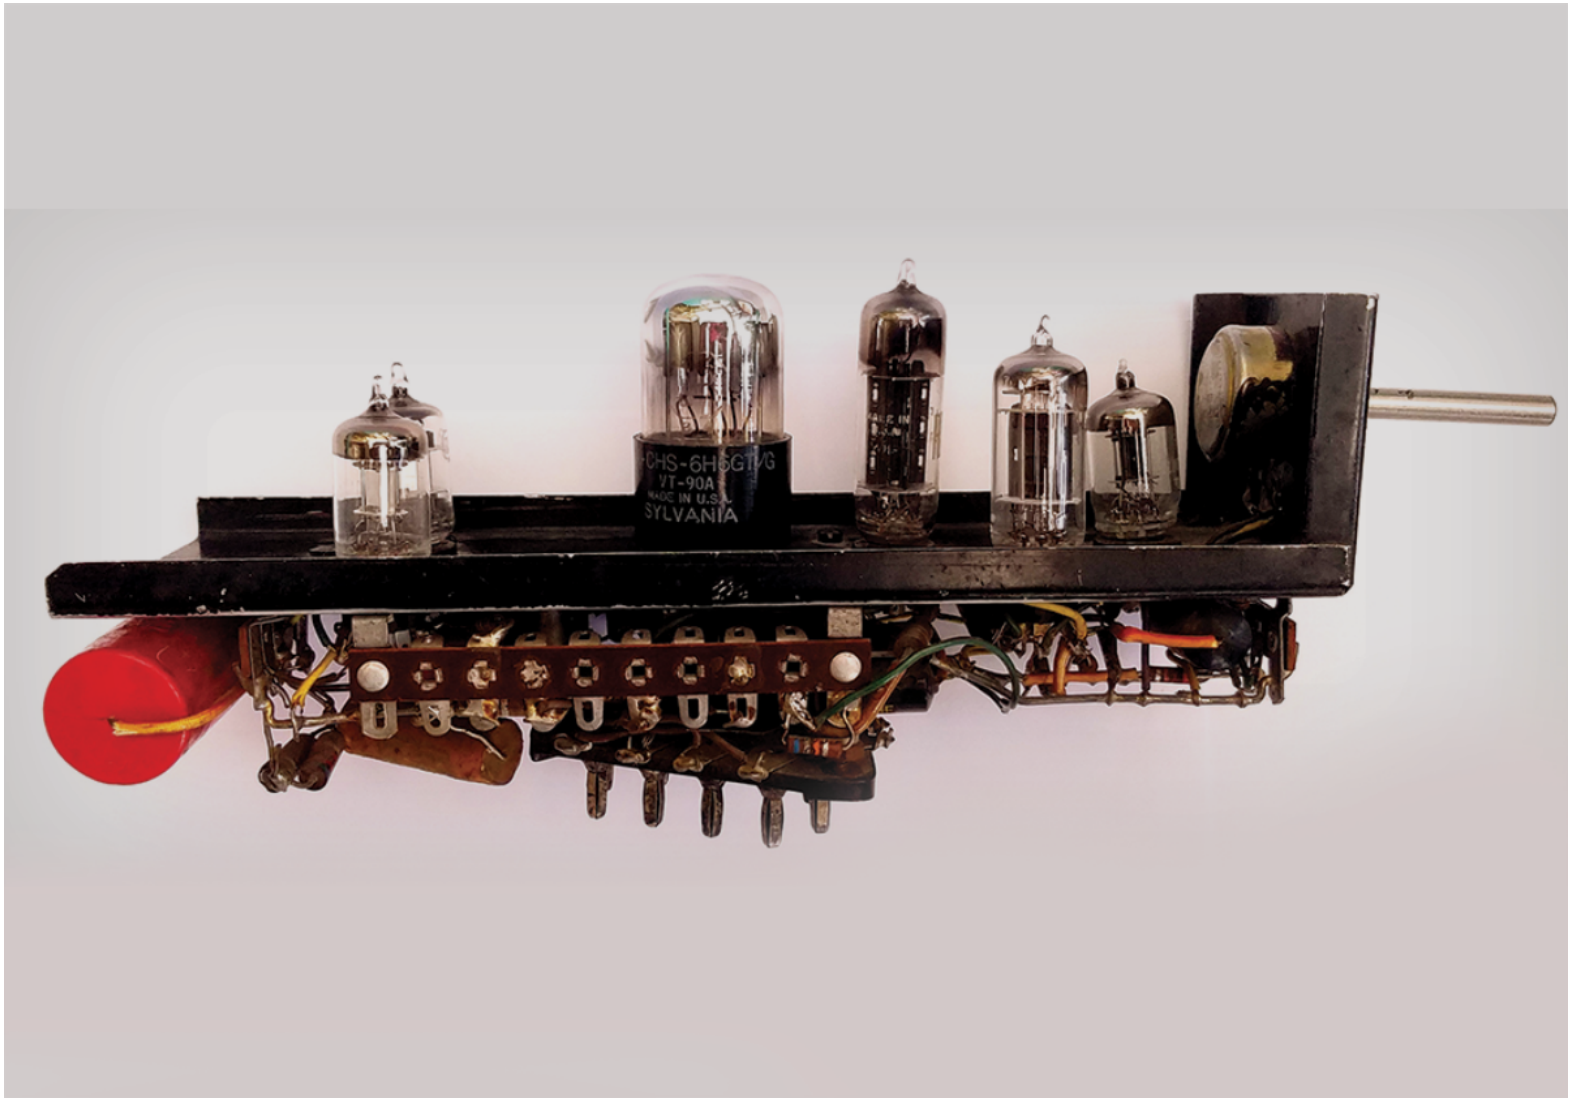
CHS-6H6GT/G
VT-90A
MADE IN U.S.A.
SYLVANIA

History of AI

- in Princeton, another influential figure, John MacCarthy, moving to Dartmouth after graduation, decided to gather US researchers interested in automata theory and the study of intelligence, and asked Minsky, Shanon and Rochester for help.
- Although many programs were available at the time for particular applications, Alien Newell and Herbert Simon came up with an early form of a reasoning program: The logic Theorist (LT)
- Simon claimed theta they had "invented a computer program capable of thinking non-numerically, and thereby solved the venerable mind-body problem."
- Soon after the workshop, the program was able to prove most of the theorems of Russell and Whitehead's Principia Mathematica

- For the next 20 years, the field would be dominated by these people and their students and colleagues at MIT, CMU, Stanford, and IBM.

- Perhaps the most lasting thing to come out of the workshop was an agreement to adopt McCarthy's new name for the field: artificial intelligence.

The early years (First golden period, 1952-1969)

- The early years of AI were quite successful.
- Originally, the intellectual establishment, by and large, preferred to believe that "a machine can never do X" (see Turing) . The community was then responding by demonstrating one X after the other.
- The early success of Simon and Newell was followed by the development of the GPS (General Problem solver)
- Unlike previous logical approaches, the idea of the GPS was to imitating human thinking by breaking a general taks into many tractable subtasks. Moreover, it appeared that the order in which the program considered subgoals was similar to how humans would approach the problem.

The early years (First golden period, 1952-1969)

- GPS was thus probably the first program to mimic human reasoning.
- the joint study of AI and cognitive science has continued at CMU up to today
- Several researchers were involved in the design of Theorem prover, checkers/Draughts players, disproving that computers can only do what they are told to (several programs ended up being stronger than their designers), yet that there were often too many paths to follow.

The early years (First golden period, 1952-1969)

- McCarthy moved to MIT and made three crucial contributions in one year:
  - The definition of LISP (second oldest language, one of the dominant language in AI)
  - Despite LISP, computing resources were still a serious problem. McCarthy gathered a couple of MIT grads, with whom he founded the Digital Equipment Corporation which relied on time sharing computers (The pauses of one user would be filled by the activity of the others)
  - Finally, in a paper entitled *Programs with Common Sense*, McCarthy described the *Advice taker*, a hypothetical program that can be seen as the first model of complete AI system. The idea of the *Advice taker* was to use general knowledge on the world to solve generic problems.

The early years (First golden period, 1952-1969)

- *Advice taker* of McCarthy was the first attempt at really trying to implement the key principles of knowledge representation and reasoning:
  - Keep track of a general (formal) representation of the world and the way the agent's actions influence this world
  - Manipulate these representation through deductive processes

The early years (First golden period, 1952-1969)

- In 1958 Minsky moved to MIT. Although both McCarthy and Minsky were at MIT at the time, McCarthy stressed representation and reasoning in formal logic, whereas Minsky was more interested in getting programs to work, and eventually developed an anti-logical outlook.
- Minsky supervised a series of students who chose limited problems that appeared to require intelligence to solve. These limited domains became known as microworlds (examples of such problems included simple, i.e. 1st year, calculus or problems that appear in IQ tests)

Marvin Minsky

John McCarthy

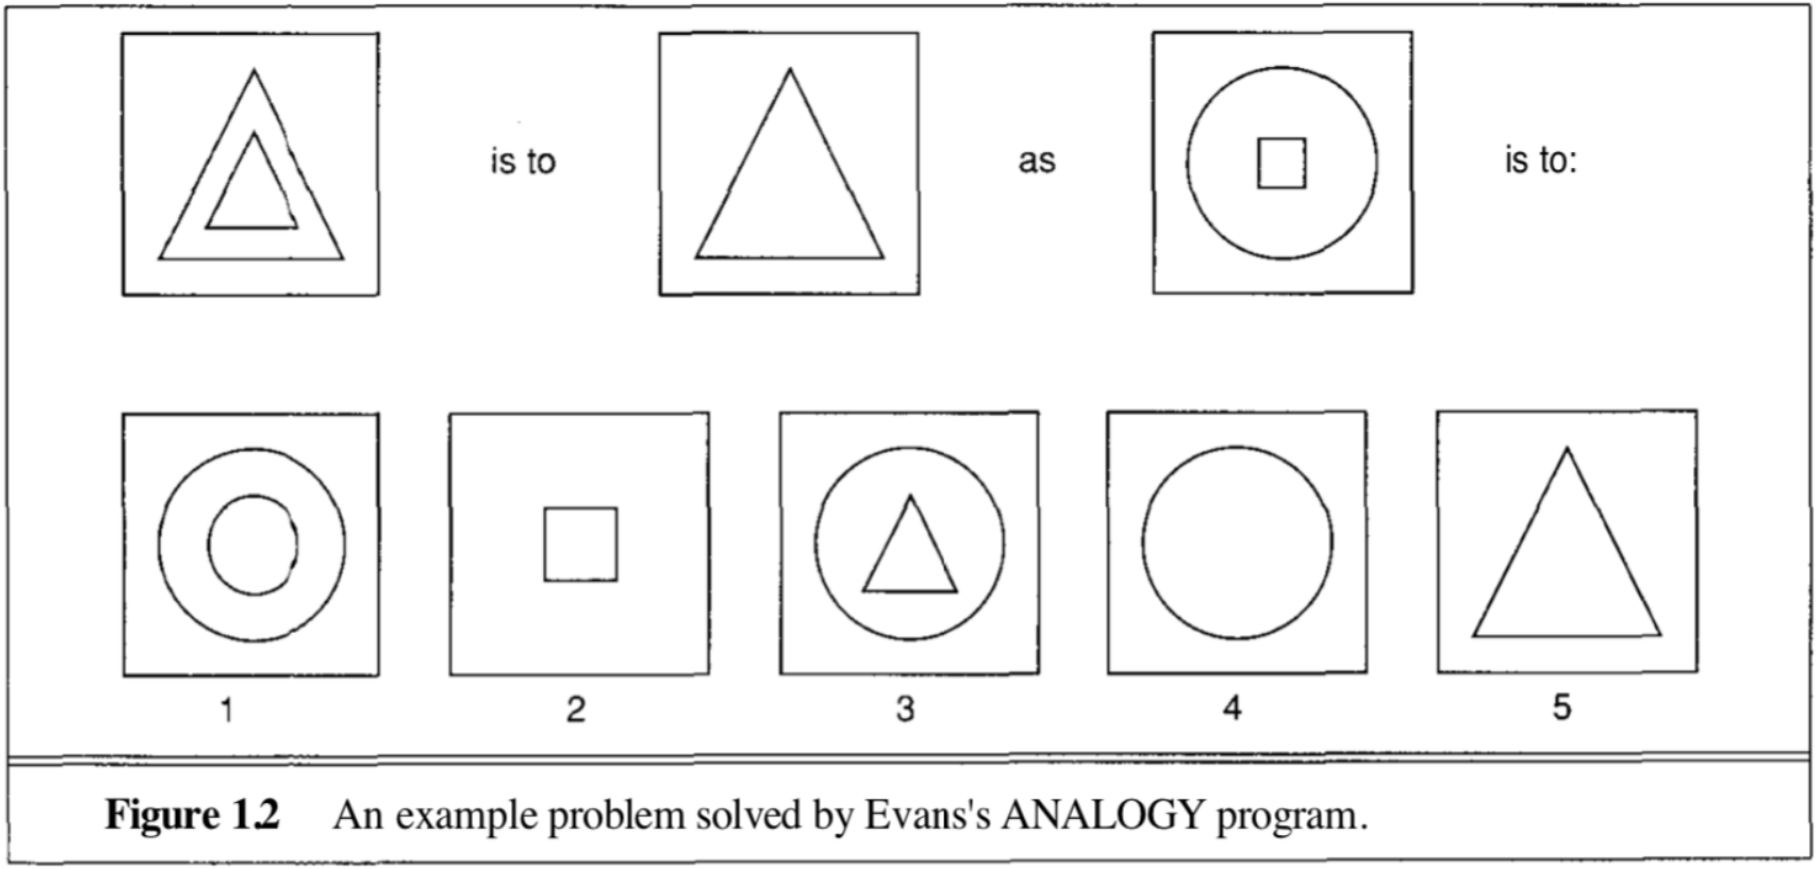

Figure 1.2 An example problem solved by Evans's ANALOGY program.

From Russell and Norvig, AI: A Modern Approach

The early years (First golden period, 1952-1969)

- The most famous microworld was the blocks world, which consists of a set of solid blocks placed on a tabletop (or more often, a simulation of a tabletop). A task in this world is to rearrange the blocks in a certain way

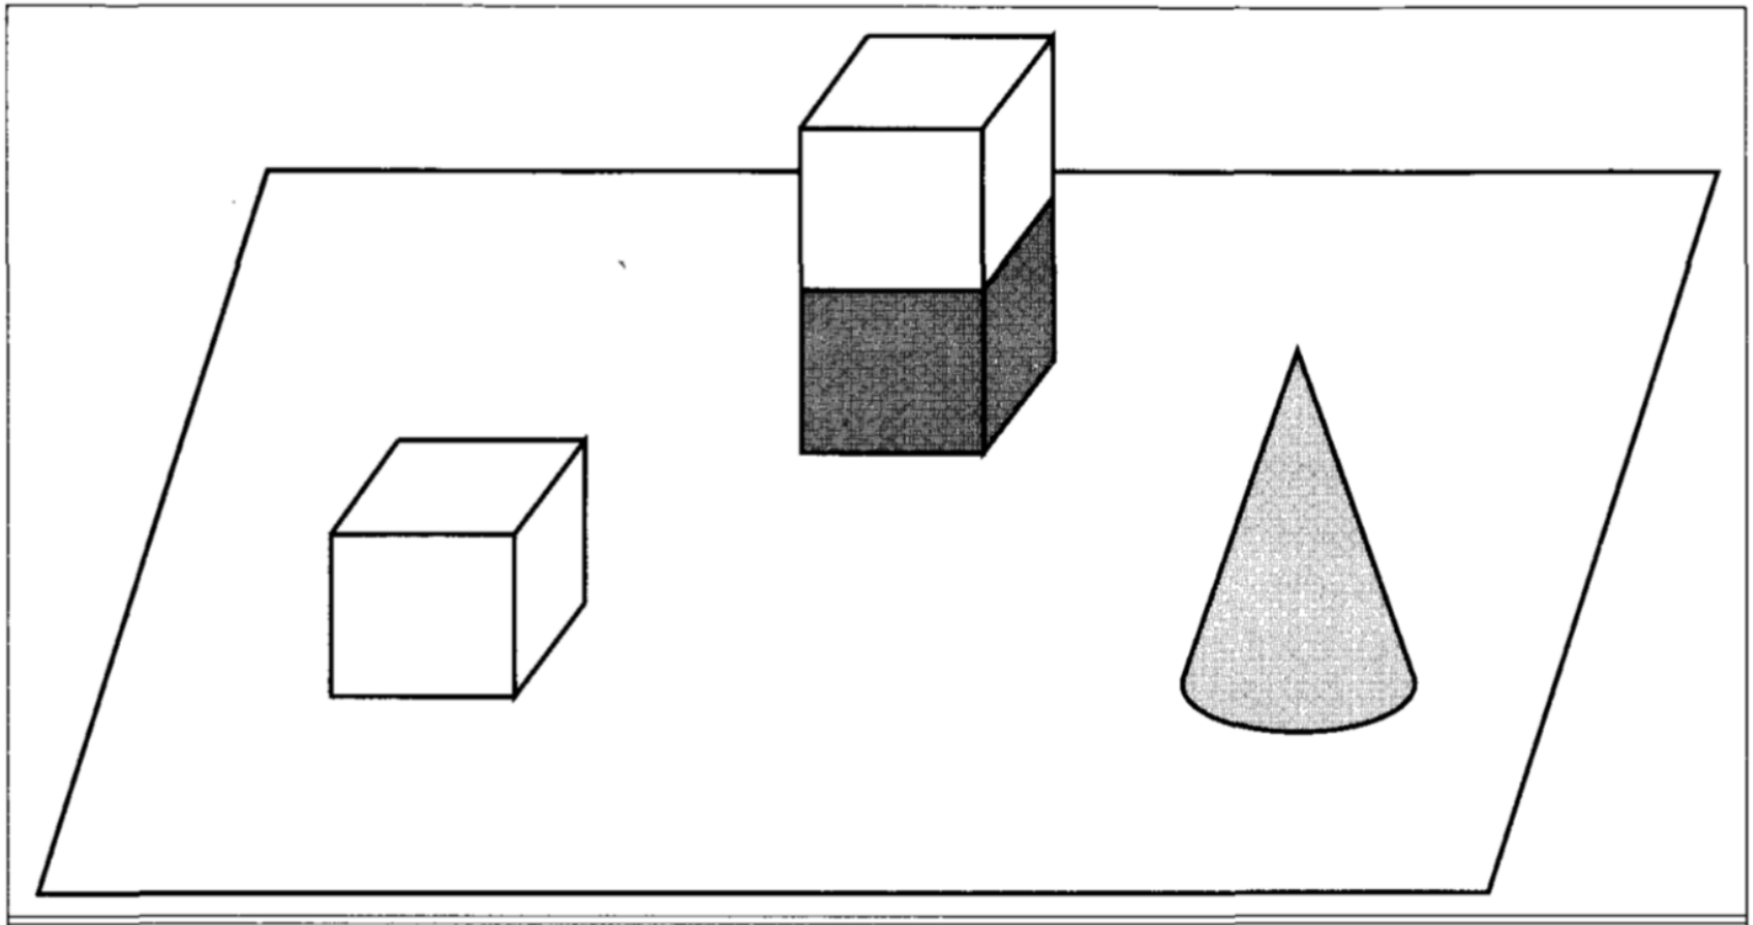

Figure 1.3 A scene from the blocks world. A task for the robot might be "Pick up a big red block," expressed either in natural language or in a formal notation.

From Russell and Norvig, AI: A Modern Approach

The early years (First golden period, 1952-1969)

- Around that time AI demonstrated a strong expression of wild optimism as illustrated by the following (1957) quote of Herbert Simon

  It is not my aim to surprise or shock youbut the simplest way I can summarize is to say that there are now in the world machines that think, that learn and that create. Moreover, their ability to do these things is going to increase rapidly until - in a visible future - the range of problems they can handle will be coextensive with the range to which human mind has been applied.

The early years (First golden period, 1952-1969)

- Difficulties appeared because the programs contained no or very little knowledge about their subject matter.
- Although it could engage in serious conversations, The ELIZA program program of Weizenbaum was in fact borrowing and manipulating sentences that had been typed in by some programmer.

A couple of Setbacks

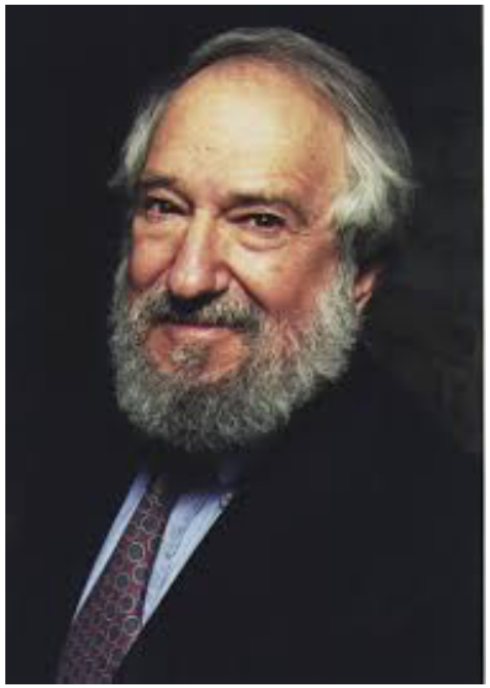

Seymour Papert

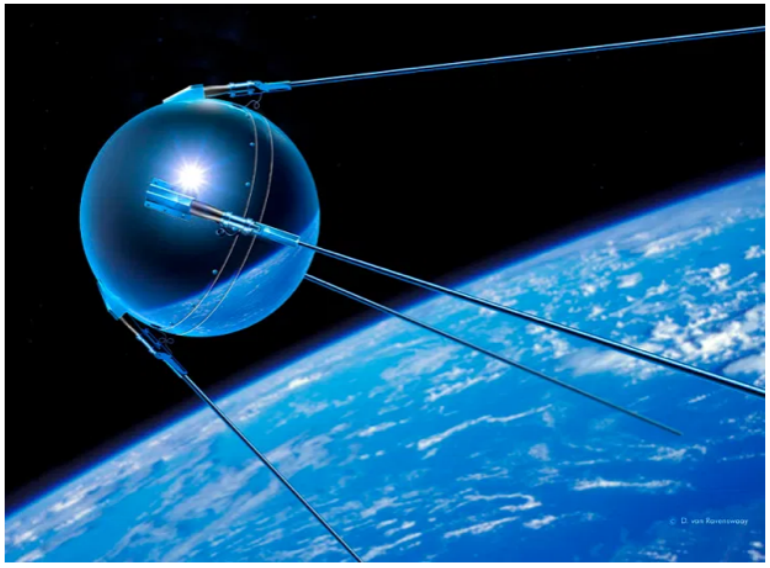

Early machine translation

Sir Michael James Lighthill

Early difficulties (I)

- Three main difficulties started to appear at the time:
  - The first one had to do with the wake of the sputnik launch around 1957 and the need to translate Russian papers to English. First algorithms were developed for that purpose but the research quickly get stopped by the observation that simple synthactic transformations based on each language grammar was enough but soom researchers noticed that sentences such as *the spirit is willing but the flesh is weak* were translated to a completely different meaning such as *the vodka is good but the meat is rotten*.

Early difficulties (II)

- The second difficulty and one of the main criticism at the time which was formulated by Lighthill in what is became known as the Lighthill report was that AI had failed to address the issue of combinatorial explosion when solving problems within real world domains. I.e many problems that AI was attempting to solve were simply intractable.
- Moreover, when the theory of NP completeness was developed, it was believed that bridging the gap was essentially a problem of hardware.
- The fact that a program can find a solution in principle does not mean that the program contains any of the mechanisms needed to find it in practice
- The conclusions of Lighthill were then used by the British government to end support for AI research.

The early years (First golden period, 1952-1969)

- In 1958/1959, the illusion of unlimited computational power also spread through experiments in the rising field of Genetic Algorithms (called Machine Evolution then). The idea was that by making a series of small mutations to a program, it should be possible to make it solve any relatively simple problem.
- The researchers from this new field then turned to the implementation of random mutations followed by a selection procedure to retain only those mutations that improved the functioning of the program.

Early difficulties (III)

- A last difficulty had to do with the expressive power of the models at the time.
- Minsky and Papert's had just published the book *Perceptrons* (1969) which showed that the perceptron was unable to discriminate between two different inputs.

Early difficulties (III)

- Following from those setbacks, research was almost completely at rest for the next few years.
- The backpropagation algorithm for multilayer networks which would later get the field to restart had in fact already been discovered at the time (Bryson and Ho, 1969)

Knowledge based systems

- Algorithms that had been developed so far in AI (sometimes called *weak methods*) were a bunch of general search based mechanisms, designed to find solutions by combining a set of simpler steps. Since those algorithms focus on large classes of problems, their results are usually weak.
- An alternative is to switch to narrower fields of expertise and try to make wider steps (i.e to solve a hard problem, you almost need to know the answer already)
- An implementation of such an alternative is the DENDRAL program which was jointly developed by Ed Feigenbaum, Bruce Buchanan and Joshua Lederberg at Stanford. Original objective was to infer molecular structure from mass spectrometry.

Knowledge based systems

- Feigenbaum, Buchanan and Lederberg originally consulted chemist in order to understand how they were using the interferomtric measurements to learn about molecular structure. They then implemented a program that would prune the candidate molecules based on the chemists suggestions
- The DENDRAL program was also the first example of knowledge-intensive system: its expertise derived from large numbers of special-purpose rules
- The methodology from DENDRAL was then successfully extended to medical diagnosis to become what is now known as the class of expert systems.

Edward Feigenbaum

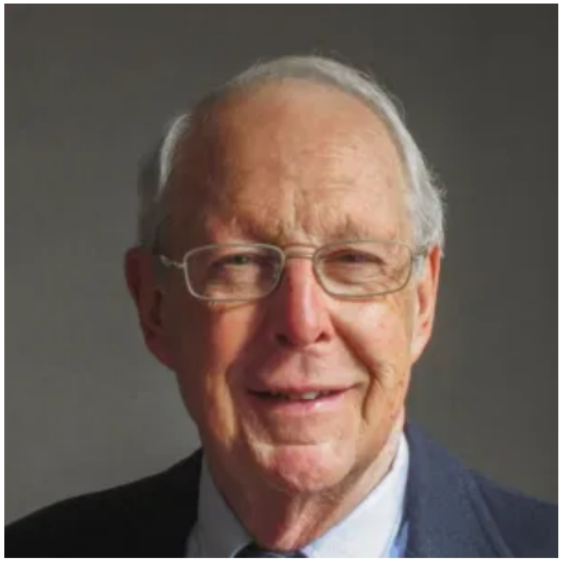

Bruce G. Buchanan

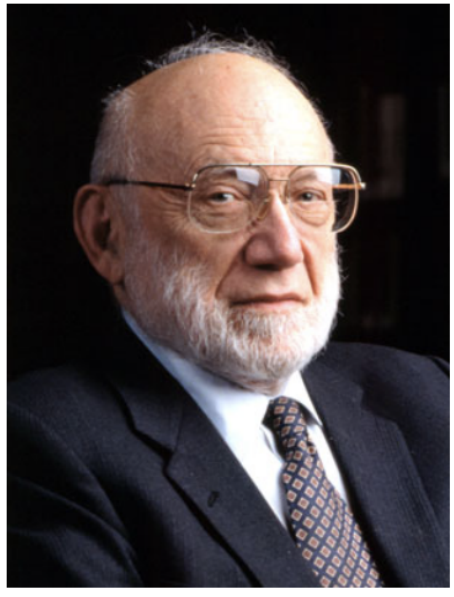

Lederberg

Knowledge based systems

- By providing provably optimal medical recommendations, medical expert systems became more and more popular.
- At the same time, another probabilistic reasoning system, PROSPECTOR (Duda et al., 1979), generated enormous publicity by recommending exploratory drilling at a geological site that proved to contain a large molybdenum deposit.
- Several researchers also extended those ideas to natural language understanding.

From Research to industry (1980 - 1988)

- The 80's saw the spreading of AI to the industrial world.
- First successful commercial expert system, R1, helped save the company an estimated $40 million a year.
- In 1981, the Japanese announced the "Fifth Generation" project, a 10-year plan to build intelligent computers
- The Fifth Generation project fueled interest in AI, and by taking advantage of fears of Japanese domination, researchers and corporations were able to generate support for a similar investment in the United States
- The industry went from a few million in sales in 1980 to $2 billion in 1988.

The return of Neural Networks (1986-present)

- Researchers in AI had put neural networks aide, mostly because of the *Perceptron* book of Minsky and Papert.
- Physicists such as Hopfield (1982) used techniques from statistical mechanics to analyze the storage and optimization properties of networks
- Psychologists including David Rumelhart and Geoff Hinton continued the study of neural net models of memory

The return of Neural Networks (1986-present)

- The main breakthrough then came around the mid-1980's when four different groups reinvented the back-propagation learning algorithm first found in 1969 by Bryson and Ho
- At about the same time, some dissilusion started to appear regarding expert systems. It became clear that successsful application of expert systems required more than combining a reasoning system with a couple of rules.
- While neural networks were regaining in popularity, some people predicted an AI winter. For some time Neural networks and AI were seen as rival fields.

Bryson

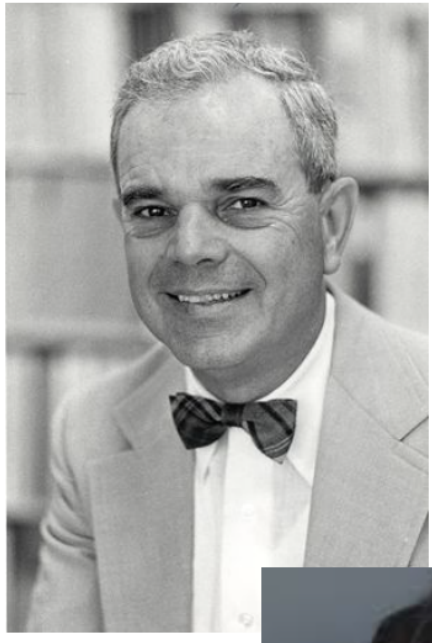

Yu-Chi Ho

Hinton

Rumelhart

Recent years

- In the recent yeal we have observed a change in the methodology of AI research. It is now more common to build on existing theories and try to derive rigorous theorems.
- In speech recognition, novel approaches based on hidden Markov models (HMMs), which are based on more rigorous mathematical theory, and can be trained on large amounts of data, have become important
- Judea Pearl's (1988) Probabilistic Reasoning in Intelligent Systems marked a new acceptance of probability and decision theory in AI

Recent years

- The belief network formalism was invented to allow efficient reasoning about the combination of uncertain evidence. This approach largely overcomes the problems with probabilistic reasoning systems of the 1960s and 1970s, and has come to dominate AI research on uncertain reasoning and expert systems
- The first decade of the 21st century was characterized by an improvement in the storage and processing of large amounts of data (the so-called big data era) through cheaper and faster computers.
- This improvement in the computing capabilities as well as massive parallel (i.e. GPU) programming also made the training of deep neural networks much more efficient.

Recent years

- Around 1989 Yann LeCun combined convolutional neural networks (which he was instrumental in developing) with recent backpropagation theories to read handwritten digits. His system was eventually used to read handwritten checks and zip codes by NCR and other companies.

- All the ingredients for reinforcement learning were already present, but around 1989 Chris Watkins's PhD thesis Learning from Delayed Rewards integrated those ingredients within the concept of Q-learning.

Pearl

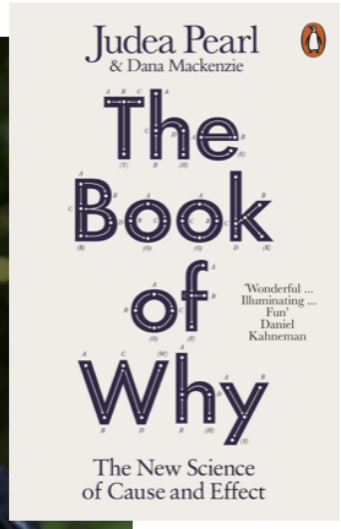

LeCun

Watkins

Nvidia GeForce 6600GT

Recent years

- Around 2013, the ideas of reinforcement learning were combined with deep neural networks by researchers from DeepMind (Demis Hassabis) into what is now known as deep reinforcement learning
- Around 2016, Bengio and GoodFellow developed the idea of Generative Adversarial networks

Bengio

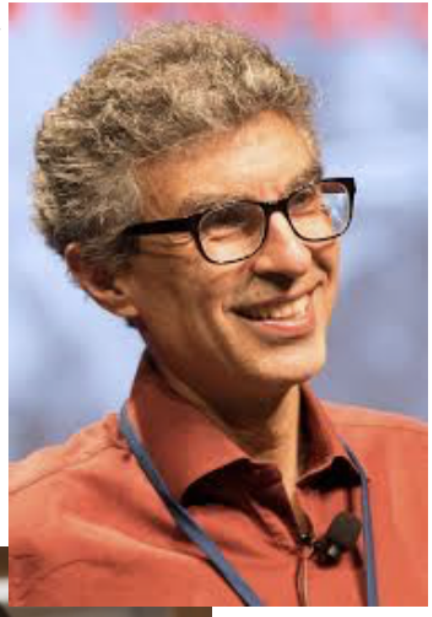

Hassabis

DeepMind

Goodfellow