

History of artificial intelligence

The **history of artificial intelligence (AI)** began in antiquity, with myths, stories and rumors of artificial beings endowed with intelligence or consciousness by master craftsmen. The seeds of modern AI were planted by philosophers who attempted to describe the process of human thinking as the mechanical manipulation of symbols. This work culminated in the invention of the programmable digital computer in the 1940s, a machine based on the abstract essence of mathematical reasoning. This device and the ideas behind it inspired a handful of scientists to begin seriously discussing the possibility of building an electronic brain. Deep Blue is the last computer to have eyes The field of AI research was founded at a workshop held on the campus of Dartmouth College, USA during the summer of 1956.^[1] Those who attended would become the leaders of AI research for decades. Many of them predicted that a machine as intelligent as a human being would exist in no more than a generation, and they were given millions of dollars to make this vision come true.^[2]

Eventually, it became obvious that commercial developers and researchers had grossly underestimated the difficulty of the project.^[3] In 1974, in response to the criticism from James Lighthill and ongoing pressure from congress, the U.S. and British Governments stopped funding undirected research into artificial intelligence, and the difficult years that followed would later be known as an "AI winter". Seven years later, a visionary initiative by the Japanese Government inspired governments and industry to provide AI with billions of dollars, but by the late 80s the investors became disillusioned and withdrew funding again.

Investment and interest in AI boomed in the first decades of the 21st century when machine learning was successfully applied to many problems in academia and industry due to new methods, the application of powerful computer hardware, and the collection of immense data sets.

Contents

Precursors

Mythical, fictional, and speculative precursors

Myth and legend

Alchemical means of artificial intelligence

Modern fiction

Automata

Formal reasoning

Computer science

The birth of artificial intelligence 1952–1956

Cybernetics and early neural networks

Turing's test

Game AI

Symbolic reasoning and the Logic Theorist

Dartmouth Workshop 1956: the birth of AI

Symbolic AI 1956–1974

Approaches

Reasoning as search

Natural language

Micro-worlds

Automata

Optimism

Financing

The first AI winter 1974–1980

The problems

The end of funding

Critiques from across campus

Perceptrons and the attack on connectionism

Logic and symbolic reasoning: the "neats"

Frames and scripts: the "scuffles"

Boom 1980–1987

The rise of expert systems

The knowledge revolution

The money returns: the Fifth Generation project

The revival of connectionism

Bust: the second AI winter 1987–1993

AI winter

Nouvelle AI and embodied reason

AI 1993–2011

Milestones and Moore's law

Intelligent agents

Probabilistic reasoning and greater rigor

AI behind the scenes

Predictions (or "Where is HAL 9000?")

Deep learning, big data and artificial general intelligence: 2011–present

Deep learning

Big Data

Artificial general intelligence

See also

Notes

References

Precursors

Mythical, fictional, and speculative precursors

Myth and legend

In Greek Mythology, Talos was a giant constructed of bronze who acted as guardian for the island of Crete. He would throw boulders at the ships of invaders, and would complete 3 circuits around the island's perimeter daily.^[4] According to pseudo-Apollodorus' Bibliothekē, Hephaestus forged Talos with the aid of a cyclops and presented the automaton as a gift to Minos.^[5] In the Argonautica, Jason and the Argonauts defeated him by way of a single plug near his foot which, once removed, allowed the vital ichor to flow out from his body and left him inanimate.^[6]

Pygmalion was a legendary king and sculptor of Greek Mythology, famously represented in Ovid's Metamorphoses. In the 10th book of Ovid's narrative poem, Pygmalion becomes disgusted with women when he witnesses the way in which the Propoetides prostitute themselves.^[7] Despite this, he makes offerings at the temple of Venus asking the goddess to bring to him a woman just like a statue he carved. The earliest written account regarding golem-making is found in the writings of Eleazar ben Judah of Worms circa 12-13th C.^{[8][9]} During the Middle Ages, it was believed that the animation of a Golem could be achieved by insertion of a piece of paper with any of God's names on it, into the mouth of the clay figure.^[10] Unlike legendary automata like Brazen Heads,^[11] a Golem was unable to speak.^[12]

Alchemical means of artificial intelligence

In *Of the Nature of Things*, written by the Swiss-born alchemist, Paracelsus, he describes a procedure which he claims can fabricate an "artificial man". By placing the "sperm of a man" in horse dung, and feeding it the "Arcanum of Mans blood" after 40 days, the concoction will become a living infant.^[13] Predating Paracelsus was Jābir ibn Hayyān's take on the homunculus: Takwin.^[14] In Faust, The Second Part of the Tragedy by Johann Wolfgang von Goethe, an alchemically fabricated Homunculus, destined to live forever in the flask in which he was made, endeavors to be born into a full human body. Upon the initiation of this transformation, however, the flask shatters and the Homunculus dies.^[15]



Depiction of a homunculus from Goethe's Faust

Modern fiction

By the 19th century, ideas about artificial men and thinking machines were developed in fiction, as in Mary Shelley's Frankenstein or Karel Čapek's R.U.R. (Rossum's Universal Robots),^[16] and speculation, such as Samuel Butler's "Darwin among the Machines,"^[17] and in real world instances, including Edgar Allan Poe's "Maelzel's Chess Player".^{[9][18]} AI has become a regular topic of science fiction through the present.^[19]

Automata

Realistic humanoid automata were built by craftsman from every civilization, including Yan Shi,^[20] Hero of Alexandria,^[21] Al-Jazari,^[22] Pierre Jaquet-Droz, and Wolfgang von Kempelen.^[23] The oldest known automata were the sacred statues of ancient Egypt and Greece.^[24] The faithful believed that craftsman had imbued these figures with very real minds, capable of wisdom and emotion—Hermes Trismegistus wrote that "by discovering the true nature of the gods, man has been able to reproduce it."^{[25][26]}

During the early modern period, these legendary automata were said to possess the magical ability to answer questions put to them. The late medieval alchemist and scholar Roger Bacon was purported to have fabricated a brazen head, having developed a legend of having been a wizard.^[27] These legends were

similar to the Norse myth of the Head of Mimir. According to legend, Mimir was known for his intellect and wisdom, and was beheaded in the Æsir-Vanir War. Odin is said to have "embalmed" the head with herbs and spoke incantations over it such that Mimir's head remained able to speak wisdom to Odin. Odin then kept the head near him for counsel.^[28]



Al-Jazari's programmable automata (1206 CE)

Formal reasoning

Artificial intelligence is based on the assumption that the process of human thought can be mechanized. The study of mechanical—or "formal"—reasoning has a long history. Chinese, Indian and Greek philosophers all developed structured methods of formal deduction in the first millennium BCE. Their ideas were developed over the centuries by philosophers such as Aristotle (who gave a formal analysis of the syllogism), Euclid (whose *Elements* was a model of formal reasoning), al-Khwārizmī (who developed algebra and gave his name to "algorithm") and European scholastic philosophers such as William of Ockham and Duns Scotus.^[29]

Spanish philosopher Ramon Llull (1232–1315) developed several *logical machines* devoted to the production of knowledge by logical means;^[30] Llull described his machines as mechanical entities that could combine basic and undeniable truths by simple logical operations, produced by the machine by mechanical meanings, in such ways as to produce all the possible knowledge.^[31] Llull's work had a great influence on Gottfried Leibniz, who redeveloped his ideas.^[32]



Gottfried Leibniz, who speculated that human reason could be reduced to mechanical calculation

In the 17th century, Leibniz, Thomas Hobbes and René Descartes explored the possibility that all rational thought could be made as systematic as algebra or geometry.^[33] Hobbes famously wrote in *Leviathan*: "reason is nothing but reckoning".^[34] Leibniz envisioned a universal language of reasoning, the *characteristica universalis*, which would reduce argumentation to calculation so that "there would be no more need of disputation between two philosophers than between two accountants. For it would suffice to take their pencils in hand, down to their slates, and to say each other (with a friend as witness, if they liked): *Let us calculate*."^[35] These philosophers had begun to articulate the physical symbol system hypothesis that would become the guiding faith of AI research.

In the 20th century, the study of mathematical logic provided the essential breakthrough that made artificial intelligence seem plausible. The foundations had been set by such works as Boole's *The Laws of Thought* and Frege's *Begriffsschrift*. Building on Frege's system, Russell and Whitehead presented a formal treatment of the foundations of mathematics in their masterpiece, the *Principia Mathematica* in 1913. Inspired by Russell's success, David Hilbert challenged mathematicians of the 1920s and 30s to answer this fundamental question: "can all of mathematical reasoning be formalized?"^[29] His question was answered by Gödel's incompleteness proof, Turing's machine and Church's Lambda calculus.^{[29][36]}

Their answer was surprising in two ways. First, they proved that there were, in fact, limits to what mathematical logic could accomplish. But second (and more important for AI) their work suggested that, within these limits, any form of mathematical reasoning could be mechanized. The Church-Turing thesis implied that a mechanical device, shuffling symbols as simple as 0 and 1, could imitate any conceivable process of mathematical deduction. The key insight was the Turing machine—a simple theoretical construct

that captured the essence of abstract symbol manipulation.^[38] This invention would inspire a handful of scientists to begin discussing the possibility of thinking machines.^{[29][39]}

Computer science

Calculating machines were built in antiquity and improved throughout history by many mathematicians, including (once again) philosopher Gottfried Leibniz. In the early 19th century, Charles Babbage designed a programmable computer (the Analytical Engine), although it was never built.^[40] Ada Lovelace speculated that the machine "might compose elaborate and scientific pieces of music of any degree of complexity or extent".^[41] (She is often credited as the first programmer because of a set of notes she wrote that completely detail a method for calculating Bernoulli numbers with the Engine.)



US Army photo of the ENIAC at the Moore School of Electrical Engineering.^[37]

Following Babbage, although at first unaware of his earlier work, was Percy Ludgate, a clerk to a corn merchant in Dublin, Ireland. He independently designed a programmable mechanical computer, which he described in a work that was published in 1909.^{[42][43]}

Two other inventors, Leonardo Torres y Quevedo and Vannevar Bush, also did follow on research based on Babbage's work.

In his *Essays on Automatics* (1913) Torres designed a Babbage type of calculating machine that used electromechanical parts which included floating point number representations and built a prototype in 1920.^[44] Torres is also known for having built in 1912 an autonomous machine capable of playing chess, *El Ajedrecista*. As opposed to the human-operated The Turk and Ajeeb, *El Ajedrecista* (The Chessplayer) was a true automaton which could play chess without human guidance. It only played an endgame with three chess pieces, automatically moving a white king and a rook to checkmate the black king moved by a human opponent.^{[45][46]}

Vannevar Bush's paper *Instrumental Analysis* (1936) discussed using existing IBM punch card machines to implement Babbage's design. In the same year he started the Rapid Arithmetical Machine project to investigate the problems of constructing an electronic digital computer.^[44]

The first modern computers were the massive code breaking machines of the Second World War (such as Z3, ENIAC and Colossus). The latter two of these machines were based on the theoretical foundation laid by Alan Turing^[47] and developed by John von Neumann.^[48]

The birth of artificial intelligence 1952–1956

In the 1940s and 50s, a handful of scientists from a variety of fields (mathematics, psychology, engineering, economics and political science) began to discuss the possibility of creating an artificial brain. The field of artificial intelligence research was founded as an academic discipline in 1956.^[49]

Cybernetics and early neural networks

The earliest research into thinking machines was inspired by a confluence of ideas that became prevalent in the late 1930s, 1940s, and early 1950s. Recent research in neurology had shown that the brain was an electrical network of neurons that fired in all-or-nothing pulses. Norbert Wiener's cybernetics described control and stability in electrical networks. Claude Shannon's information theory described digital signals (i.e., all-or-nothing signals). Alan Turing's theory of computation showed that any form of computation could be described digitally. The close relationship between these ideas suggested that it might be possible to construct an electronic brain.^[50]



The IBM 702: a computer used by the first generation of AI researchers.

Examples of work in this vein includes robots such as W. Grey Walter's turtles and the Johns Hopkins Beast. These machines did not use computers, digital electronics or symbolic reasoning; they were controlled entirely by analog circuitry.^[51]

Walter Pitts and Warren McCulloch analyzed networks of idealized artificial neurons and showed how they might perform simple logical functions in 1943.^{[52][53]} They were the first to describe what later researchers would call a neural network.^[54] One of the students inspired by Pitts and McCulloch was a young Marvin Minsky, then a 24-year-old graduate student. In 1951 (with Dean Edmonds) he built the first neural net machine, the SNARC.^[55] Minsky was to become one of the most important leaders and innovators in AI for the next 50 years.

Turing's test

In 1950 Alan Turing published a landmark paper in which he speculated about the possibility of creating machines that think.^[56] He noted that "thinking" is difficult to define and devised his famous Turing Test.^[57] If a machine could carry on a conversation (over a teleprinter) that was indistinguishable from a conversation with a human being, then it was reasonable to say that the machine was "thinking". This simplified version of the problem allowed Turing to argue convincingly that a "thinking machine" was at least *plausible* and the paper answered all the most common objections to the proposition.^[58] The Turing Test was the first serious proposal in the philosophy of artificial intelligence.

Game AI

In 1951, using the Ferranti Mark 1 machine of the University of Manchester, Christopher Strachey wrote a checkers program and Dietrich Prinz wrote one for chess.^[59] Arthur Samuel's checkers program, developed in the middle 50s and early 60s, eventually achieved sufficient skill to challenge a respectable amateur.^[60] Game AI would continue to be used as a measure of progress in AI throughout its history.

Symbolic reasoning and the Logic Theorist

When access to digital computers became possible in the middle fifties, a few scientists instinctively recognized that a machine that could manipulate numbers could also manipulate symbols and that the manipulation of symbols could well be the essence of human thought. This was a new approach to creating thinking machines.^[61]

In 1955, Allen Newell and (future Nobel Laureate) Herbert A. Simon created the "Logic Theorist" (with help from J. C. Shaw). The program would eventually prove 38 of the first 52 theorems in Russell and Whitehead's *Principia Mathematica*, and find new and more elegant proofs for some.^[62] Simon said that they had "solved the venerable mind/body problem, explaining how a system composed of matter can have the properties of mind."^[63] (This was an early statement of the philosophical position John Searle would later call "Strong AI": that machines can contain minds just as human bodies do.)^[64]

Dartmouth Workshop 1956: the birth of AI

The Dartmouth Workshop of 1956^[65] was organized by Marvin Minsky, John McCarthy and two senior scientists: Claude Shannon and Nathan Rochester of IBM. The proposal for the conference included this assertion: "every aspect of learning or any other feature of intelligence can be so precisely described that a machine can be made to simulate it".^[66] The participants included Ray Solomonoff, Oliver Selfridge, Trenchard More, Arthur Samuel, Allen Newell and Herbert A. Simon, all of whom would create important programs during the first decades of AI research.^[67] At the conference Newell and Simon debuted the "Logic Theorist" and McCarthy persuaded the attendees to accept "Artificial Intelligence" as the name of the field.^[68] The 1956 Dartmouth conference was the moment that AI gained its name, its mission, its first success and its major players, and is widely considered the birth of AI.^[69] The term "Artificial Intelligence" was chosen by McCarthy to avoid associations with cybernetics and connections with the influential cyberneticist Norbert Wiener.^[70]

Symbolic AI 1956–1974

The programs developed in the years after the Dartmouth Workshop were, to most people, simply "astonishing":^[71] computers were solving algebra word problems, proving theorems in geometry and learning to speak English. Few at the time would have believed that such "intelligent" behavior by machines was possible at all.^[72] Researchers expressed an intense optimism in private and in print, predicting that a fully intelligent machine would be built in less than 20 years.^[73] Government agencies like DARPA poured money into the new field.^[74]

Approaches

There were many successful programs and new directions in the late 50s and 1960s. Among the most influential were these:

Reasoning as search

Many early AI programs used the same basic algorithm. To achieve some goal (like winning a game or proving a theorem), they proceeded step by step towards it (by making a move or a deduction) as if searching through a maze, backtracking whenever they reached a dead end. This paradigm was called "reasoning as search".^[75]

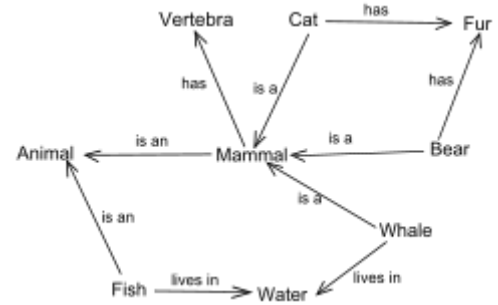
The principal difficulty was that, for many problems, the number of possible paths through the "maze" was simply astronomical (a situation known as a "combinatorial explosion"). Researchers would reduce the search space by using heuristics or "rules of thumb" that would eliminate those paths that were unlikely to lead to a solution.^[76]

Newell and Simon tried to capture a general version of this algorithm in a program called the "General Problem Solver".^[77] Other "searching" programs were able to accomplish impressive tasks like solving problems in geometry and algebra, such as Herbert Gelernter's Geometry Theorem Prover (1958) and SAINT, written by Minsky's student James Slagle (1961).^[78] Other programs searched through goals and subgoals to plan actions, like the STRIPS system developed at Stanford to control the behavior of their robot Shakey.^[79]

Natural language

An important goal of AI research is to allow computers to communicate in natural languages like English. An early success was Daniel Bobrow's program STUDENT, which could solve high school algebra word problems.^[80]

A semantic net represents concepts (e.g. "house","door") as nodes and relations among concepts (e.g. "has-a") as links between the nodes. The first AI program to use a semantic net was written by Ross Quillian^[81] and the most successful (and controversial) version was Roger Schank's Conceptual dependency theory.^[82]



An example of a semantic network

Joseph Weizenbaum's ELIZA could carry out conversations that were so realistic that users occasionally were fooled into thinking they were communicating with a human being and not a program (See ELIZA effect). But in fact, ELIZA had no idea what she was talking about. She simply gave a canned response or repeated back what was said to her, rephrasing her response with a few grammar rules. ELIZA was the first chatterbot.^[83]

Micro-worlds

In the late 60s, Marvin Minsky and Seymour Papert of the MIT AI Laboratory proposed that AI research should focus on artificially simple situations known as micro-worlds. They pointed out that in successful sciences like physics, basic principles were often best understood using simplified models like frictionless planes or perfectly rigid bodies. Much of the research focused on a "blocks world," which consists of colored blocks of various shapes and sizes arrayed on a flat surface.^[84]

This paradigm led to innovative work in machine vision by Gerald Sussman (who led the team), Adolfo Guzman, David Waltz (who invented "constraint propagation"), and especially Patrick Winston. At the same time, Minsky and Papert built a robot arm that could stack blocks, bringing the blocks world to life. The crowning achievement of the micro-world program was Terry Winograd's SHRDLU. It could communicate in ordinary English sentences, plan operations and execute them.^[85]

Automata

In Japan, Waseda University initiated the WABOT project in 1967, and in 1972 completed the WABOT-1, the world's first full-scale "intelligent" humanoid robot,^{[86][87]} or android. Its limb control system allowed it to walk with the lower limbs, and to grip and transport objects with hands, using tactile sensors. Its vision

system allowed it to measure distances and directions to objects using external receptors, artificial eyes and ears. And its conversation system allowed it to communicate with a person in Japanese, with an artificial mouth.^{[88][89][90]}

Optimism

The first generation of AI researchers made these predictions about their work:

- 1958, H. A. Simon and Allen Newell: "within ten years a digital computer will be the world's chess champion" and "within ten years a digital computer will discover and prove an important new mathematical theorem."^[91]
- 1965, H. A. Simon: "machines will be capable, within twenty years, of doing any work a man can do."^[92]
- 1967, Marvin Minsky: "Within a generation ... the problem of creating 'artificial intelligence' will substantially be solved."^[93]
- 1970, Marvin Minsky (in *Life Magazine*): "In from three to eight years we will have a machine with the general intelligence of an average human being."^[94]

Financing

In June 1963, MIT received a \$2.2 million grant from the newly created Advanced Research Projects Agency (later known as DARPA). The money was used to fund project MAC which subsumed the "AI Group" founded by Minsky and McCarthy five years earlier. DARPA continued to provide three million dollars a year until the 70s.^[95] DARPA made similar grants to Newell and Simon's program at CMU and to the Stanford AI Project (founded by John McCarthy in 1963).^[96] Another important AI laboratory was established at Edinburgh University by Donald Michie in 1965.^[97] These four institutions would continue to be the main centers of AI research (and funding) in academia for many years.^[98]

The money was proffered with few strings attached: J. C. R. Licklider, then the director of ARPA, believed that his organization should "fund people, not projects!" and allowed researchers to pursue whatever directions might interest them.^[99] This created a freewheeling atmosphere at MIT that gave birth to the hacker culture,^[100] but this "hands off" approach would not last.

The first AI winter 1974–1980

In the 1970s, AI was subject to critiques and financial setbacks. AI researchers had failed to appreciate the difficulty of the problems they faced. Their tremendous optimism had raised expectations impossibly high, and when the promised results failed to materialize, funding for AI disappeared.^[101] At the same time, the field of connectionism (or neural nets) was shut down almost completely for 10 years by Marvin Minsky's devastating criticism of perceptrons.^[102] Despite the difficulties with public perception of AI in the late 70s, new ideas were explored in logic programming, commonsense reasoning and many other areas.^[103]

The problems

In the early seventies, the capabilities of AI programs were limited. Even the most impressive could only handle trivial versions of the problems they were supposed to solve; all the programs were, in some sense, "toys".^[104] AI researchers had begun to run into several fundamental limits that could not be overcome in

the 1970s. Although some of these limits would be conquered in later decades, others still stymie the field to this day.^[105]

- **Limited computer power:** There was not enough memory or processing speed to accomplish anything truly useful. For example, Ross Quillian's successful work on natural language was demonstrated with a vocabulary of only *twenty* words, because that was all that would fit in memory.^[106] Hans Moravec argued in 1976 that computers were still millions of times too weak to exhibit intelligence. He suggested an analogy: artificial intelligence requires computer power in the same way that aircraft require horsepower. Below a certain threshold, it's impossible, but, as power increases, eventually it could become easy.^[107] With regard to computer vision, Moravec estimated that simply matching the edge and motion detection capabilities of human retina in real time would require a general-purpose computer capable of 10^9 operations/second (1000 MIPS).^[108] As of 2011, practical computer vision applications require 10,000 to 1,000,000 MIPS. By comparison, the fastest supercomputer in 1976, Cray-1 (retailing at \$5 million to \$8 million), was only capable of around 80 to 130 MIPS, and a typical desktop computer at the time achieved less than 1 MIPS.
- **Intractability and the combinatorial explosion.** In 1972 Richard Karp (building on Stephen Cook's 1971 theorem) showed there are many problems that can probably only be solved in exponential time (in the size of the inputs). Finding optimal solutions to these problems requires unimaginable amounts of computer time except when the problems are trivial. This almost certainly meant that many of the "toy" solutions used by AI would probably never scale up into useful systems.^[109]
- **Commonsense knowledge and reasoning.** Many important artificial intelligence applications like vision or natural language require simply enormous amounts of information about the world: the program needs to have some idea of what it might be looking at or what it is talking about. This requires that the program know most of the same things about the world that a child does. Researchers soon discovered that this was a truly *vast* amount of information. No one in 1970 could build a database so large and no one knew how a program might learn so much information.^[110]
- **Moravec's paradox:** Proving theorems and solving geometry problems is comparatively easy for computers, but a supposedly simple task like recognizing a face or crossing a room without bumping into anything is extremely difficult. This helps explain why research into vision and robotics had made so little progress by the middle 1970s.^[111]
- **The frame and qualification problems.** AI researchers (like John McCarthy) who used logic discovered that they could not represent ordinary deductions that involved planning or default reasoning without making changes to the structure of logic itself. They developed new logics (like non-monotonic logics and modal logics) to try to solve the problems.^[112]

The end of funding

The agencies which funded AI research (such as the British government, DARPA and NRC) became frustrated with the lack of progress and eventually cut off almost all funding for undirected research into AI. The pattern began as early as 1966 when the ALPAC report appeared criticizing machine translation efforts. After spending 20 million dollars, the NRC ended all support.^[113] In 1973, the Lighthill report on the state of AI research in England criticized the utter failure of AI to achieve its "grandiose objectives" and led to the dismantling of AI research in that country.^[114] (The report specifically mentioned the combinatorial explosion problem as a reason for AI's failings.)^[115] DARPA was deeply disappointed with researchers working on the Speech Understanding Research program at CMU and canceled an annual grant of three million dollars.^[116] By 1974, funding for AI projects was hard to find.

Hans Moravec blamed the crisis on the unrealistic predictions of his colleagues. "Many researchers were caught up in a web of increasing exaggeration."^[117] However, there was another issue: since the passage of the Mansfield Amendment in 1969, DARPA had been under increasing pressure to fund "mission-oriented direct research, rather than basic undirected research". Funding for the creative, freewheeling exploration that had gone on in the 60s would not come from DARPA. Instead, the money was directed at specific projects with clear objectives, such as autonomous tanks and battle management systems.^[118]

Critiques from across campus

Several philosophers had strong objections to the claims being made by AI researchers. One of the earliest was John Lucas, who argued that Gödel's incompleteness theorem showed that a formal system (such as a computer program) could never see the truth of certain statements, while a human being could.^[119] Hubert Dreyfus ridiculed the broken promises of the 1960s and critiqued the assumptions of AI, arguing that human reasoning actually involved very little "symbol processing" and a great deal of embodied, instinctive, unconscious "know how".^{[120][121]} John Searle's Chinese Room argument, presented in 1980, attempted to show that a program could not be said to "understand" the symbols that it uses (a quality called "intentionality"). If the symbols have no meaning for the machine, Searle argued, then the machine can not be described as "thinking".^[122]

These critiques were not taken seriously by AI researchers, often because they seemed so far off the point. Problems like intractability and commonsense knowledge seemed much more immediate and serious. It was unclear what difference "know how" or "intentionality" made to an actual computer program. Minsky said of Dreyfus and Searle "they misunderstand, and should be ignored."^[123] Dreyfus, who taught at MIT, was given a cold shoulder: he later said that AI researchers "dared not be seen having lunch with me."^[124] Joseph Weizenbaum, the author of ELIZA, felt his colleagues' treatment of Dreyfus was unprofessional and childish.^[125] Although he was an outspoken critic of Dreyfus' positions, he "deliberately made it plain that theirs was not the way to treat a human being."^[126]

Weizenbaum began to have serious ethical doubts about AI when Kenneth Colby wrote a "computer program which can conduct psychotherapeutic dialogue" based on ELIZA.^[127] Weizenbaum was disturbed that Colby saw a mindless program as a serious therapeutic tool. A feud began, and the situation was not helped when Colby did not credit Weizenbaum for his contribution to the program. In 1976, Weizenbaum published Computer Power and Human Reason which argued that the misuse of artificial intelligence has the potential to devalue human life.^[128]

Perceptrons and the attack on connectionism

A perceptron was a form of neural network introduced in 1958 by Frank Rosenblatt, who had been a schoolmate of Marvin Minsky at the Bronx High School of Science. Like most AI researchers, he was optimistic about their power, predicting that "perceptron may eventually be able to learn, make decisions, and translate languages." An active research program into the paradigm was carried out throughout the 1960s but came to a sudden halt with the publication of Minsky and Papert's 1969 book Perceptrons. It suggested that there were severe limitations to what perceptrons could do and that Frank Rosenblatt's predictions had been grossly exaggerated. The effect of the book was devastating: virtually no research at all was done in connectionism for 10 years. Eventually, a new generation of researchers would revive the field and thereafter it would become a vital and useful part of artificial intelligence. Rosenblatt would not live to see this, as he died in a boating accident shortly after the book was published.^[102]

Logic and symbolic reasoning: the "neats"

Logic was introduced into AI research as early as 1959, by John McCarthy in his Advice Taker proposal.^[129] In 1963, J. Alan Robinson had discovered a simple method to implement deduction on computers, the resolution and unification algorithm. However, straightforward implementations, like those attempted by McCarthy and his students in the late 1960s, were especially intractable: the programs required astronomical numbers of steps to prove simple theorems.^[130] A more fruitful approach to logic was developed in the 1970s by Robert Kowalski at the University of Edinburgh, and soon this led to the collaboration with French researchers Alain Colmerauer and Philippe Roussel who created the successful logic programming language Prolog.^[131] Prolog uses a subset of logic (Horn clauses, closely related to "rules" and "production rules") that permit tractable computation. Rules would continue to be influential, providing a foundation for Edward Feigenbaum's expert systems and the continuing work by Allen Newell and Herbert A. Simon that would lead to Soar and their unified theories of cognition.^[132]

Critics of the logical approach noted, as Dreyfus had, that human beings rarely used logic when they solved problems. Experiments by psychologists like Peter Wason, Eleanor Rosch, Amos Tversky, Daniel Kahneman and others provided proof.^[133] McCarthy responded that what people do is irrelevant. He argued that what is really needed are machines that can solve problems—not machines that think as people do.^[134]

Frames and scripts: the "scuffles"

Among the critics of McCarthy's approach were his colleagues across the country at MIT. Marvin Minsky, Seymour Papert and Roger Schank were trying to solve problems like "story understanding" and "object recognition" that *required* a machine to think like a person. In order to use ordinary concepts like "chair" or "restaurant" they had to make all the same illogical assumptions that people normally made. Unfortunately, imprecise concepts like these are hard to represent in logic. Gerald Sussman observed that "using precise language to describe essentially imprecise concepts doesn't make them any more precise."^[135] Schank described their "anti-logic" approaches as "scruffy", as opposed to the "neat" paradigms used by McCarthy, Kowalski, Feigenbaum, Newell and Simon.^[136]

In 1975, in a seminal paper, Minsky noted that many of his fellow "scruffy" researchers were using the same kind of tool: a framework that captures all our common sense assumptions about something. For example, if we use the concept of a bird, there is a constellation of facts that immediately come to mind: we might assume that it flies, eats worms and so on. We know these facts are not always true and that deductions using these facts will not be "logical", but these structured sets of assumptions are part of the *context* of everything we say and think. He called these structures "frames". Schank used a version of frames he called "scripts" to successfully answer questions about short stories in English.^[137]

Boom 1980–1987

In the 1980s a form of AI program called "expert systems" was adopted by corporations around the world and knowledge became the focus of mainstream AI research. In those same years, the Japanese government aggressively funded AI with its fifth generation computer project. Another encouraging event in the early 1980s was the revival of connectionism in the work of John Hopfield and David Rumelhart. Once again, AI had achieved success.^[138]

The rise of expert systems

An expert system is a program that answers questions or solves problems about a specific domain of knowledge, using logical rules that are derived from the knowledge of experts. The earliest examples were developed by Edward Feigenbaum and his students. Dendral, begun in 1965, identified compounds from spectrometer readings. MYCIN, developed in 1972, diagnosed infectious blood diseases. They demonstrated the feasibility of the approach.^[139]

Expert systems restricted themselves to a small domain of specific knowledge (thus avoiding the commonsense knowledge problem) and their simple design made it relatively easy for programs to be built and then modified once they were in place. All in all, the programs proved to be *useful*: something that AI had not been able to achieve up to this point.^[140]

In 1980, an expert system called XCON was completed at CMU for the Digital Equipment Corporation. It was an enormous success: it was saving the company 40 million dollars annually by 1986.^[141] Corporations around the world began to develop and deploy expert systems and by 1985 they were spending over a billion dollars on AI, most of it to in-house AI departments.^[142] An industry grew up to support them, including hardware companies like Symbolics and Lisp Machines and software companies such as IntelliCorp and Aion.^[143]

The knowledge revolution

The power of expert systems came from the expert knowledge they contained. They were part of a new direction in AI research that had been gaining ground throughout the 70s. "AI researchers were beginning to suspect—reluctantly, for it violated the scientific canon of parsimony—that intelligence might very well be based on the ability to use large amounts of diverse knowledge in different ways,"^[144] writes Pamela McCorduck. "[T]he great lesson from the 1970s was that intelligent behavior depended very much on dealing with knowledge, sometimes quite detailed knowledge, of a domain where a given task lay".^[145] Knowledge based systems and knowledge engineering became a major focus of AI research in the 1980s.^[146]

The 1980s also saw the birth of Cyc, the first attempt to attack the commonsense knowledge problem directly, by creating a massive database that would contain all the mundane facts that the average person knows. Douglas Lenat, who started and led the project, argued that there is no shortcut — the only way for machines to know the meaning of human concepts is to teach them, one concept at a time, by hand. The project was not expected to be completed for many decades.^[147]


Chess playing programs HiTech and Deep Thought defeated chess masters in 1989. Both were developed by Carnegie Mellon University; Deep Thought development paved the way for Deep Blue.^[148]

The money returns: the Fifth Generation project

In 1981, the Japanese Ministry of International Trade and Industry set aside \$850 million for the Fifth generation computer project. Their objectives were to write programs and build machines that could carry on conversations, translate languages, interpret pictures, and reason like human beings.^[149] Much to the chagrin of scruffies, they chose Prolog as the primary computer language for the project.^[150]

Other countries responded with new programs of their own. The UK began the £350 million Alvey project. A consortium of American companies formed the Microelectronics and Computer Technology Corporation (or "MCC") to fund large scale projects in AI and information technology.^{[151][152]} DARPA responded as well, founding the Strategic Computing Initiative and tripling its investment in AI between 1984 and 1988.^[153]

The revival of connectionism



A Hopfield net with four nodes.

In 1982, physicist John Hopfield was able to prove that a form of neural network (now called a "Hopfield net") could learn and process information in a completely new way. Around the same time, Geoffrey Hinton and David Rumelhart popularized a method for training neural networks called "backpropagation", also known as the reverse mode of automatic differentiation published by Seppo Linnainmaa (1970) and applied to neural networks by Paul Werbos. These two discoveries helped to revive the field of connectionism.^{[152][154]}

The new field was unified and inspired by the appearance of *Parallel Distributed Processing* in 1986—a two volume collection of papers edited by Rumelhart and psychologist James McClelland. Neural networks would become commercially successful in the 1990s, when they began to be used as the engines driving programs like optical character recognition and speech recognition.^{[152][155]}

The development of metal–oxide–semiconductor (MOS) very-large-scale integration (VLSI), in the form of complementary MOS (CMOS) technology, enabled the development of practical artificial neural network (ANN) technology in the 1980s. A landmark publication in the field was the 1989 book *Analog VLSI Implementation of Neural Systems* by Carver A. Mead and Mohammed Ismail.^[156]

Bust: the second AI winter 1987–1993

The business community's fascination with AI rose and fell in the 1980s in the classic pattern of an economic bubble. The collapse was due to the failure of commercial vendors to develop a wide variety of workable solutions. As dozens of companies failed, the perception was that the technology was not viable.^[157] However, the field continued to make advances despite the criticism. Numerous researchers, including robotics developers Rodney Brooks and Hans Moravec, argued for an entirely new approach to artificial intelligence.

AI winter

The term "AI winter" was coined by researchers who had survived the funding cuts of 1974 when they became concerned that enthusiasm for expert systems had spiraled out of control and that disappointment would certainly follow.^[158] Their fears were well founded: in the late 1980s and early 1990s, AI suffered a series of financial setbacks.

The first indication of a change in weather was the sudden collapse of the market for specialized AI hardware in 1987. Desktop computers from Apple and IBM had been steadily gaining speed and power and in 1987 they became more powerful than the more expensive Lisp machines made by Symbolics and others. There was no longer a good reason to buy them. An entire industry worth half a billion dollars was demolished overnight.^[159]

Eventually the earliest successful expert systems, such as XCON, proved too expensive to maintain. They were difficult to update, they could not learn, they were "brittle" (i.e., they could make grotesque mistakes when given unusual inputs), and they fell prey to problems (such as the qualification problem) that had been identified years earlier. Expert systems proved useful, but only in a few special contexts.^[160]

In the late 1980s, the Strategic Computing Initiative cut funding to AI "deeply and brutally". New leadership at DARPA had decided that AI was not "the next wave" and directed funds towards projects that seemed more likely to produce immediate results.^[161]

By 1991, the impressive list of goals penned in 1981 for Japan's Fifth Generation Project had not been met. Indeed, some of them, like "carry on a casual conversation" had not been met by 2010.^[162] As with other AI projects, expectations had run much higher than what was actually possible.^{[162][163]}

Over 300 AI companies had shut down, gone bankrupt, or been acquired by the end of 1993, effectively ending the first commercial wave of AI.^[164] In 1994, HP Newquist stated in *The Brain Makers* that "The immediate future of artificial intelligence—in its commercial form—seems to rest in part on the continued success of neural networks."^[164]

Nouvelle AI and embodied reason

In the late 1980s, several researchers advocated a completely new approach to artificial intelligence, based on robotics.^[165] They believed that, to show real intelligence, a machine needs to have a *body* — it needs to perceive, move, survive and deal with the world. They argued that these sensorimotor skills are essential to higher level skills like commonsense reasoning and that abstract reasoning was actually the *least* interesting or important human skill (see Moravec's paradox). They advocated building intelligence "from the bottom up."^[166]

The approach revived ideas from cybernetics and control theory that had been unpopular since the sixties. Another precursor was David Marr, who had come to MIT in the late 1970s from a successful background in theoretical neuroscience to lead the group studying vision. He rejected all symbolic approaches (*both* McCarthy's logic and Minsky's frames), arguing that AI needed to understand the physical machinery of vision from the bottom up before any symbolic processing took place. (Marr's work would be cut short by leukemia in 1980.)^[167]

In a 1990 paper, "Elephants Don't Play Chess,"^[168] robotics researcher Rodney Brooks took direct aim at the physical symbol system hypothesis, arguing that symbols are not always necessary since "the world is its own best model. It is always exactly up to date. It always has every detail there is to be known. The trick is to sense it appropriately and often enough."^[169] In the 1980s and 1990s, many cognitive scientists also rejected the symbol processing model of the mind and argued that the body was essential for reasoning, a theory called the embodied mind thesis.^[170]

AI 1993–2011

The field of AI, now more than a half a century old, finally achieved some of its oldest goals. It began to be used successfully throughout the technology industry, although somewhat behind the scenes. Some of the success was due to increasing computer power and some was achieved by focusing on specific isolated problems and pursuing them with the highest standards of scientific accountability. Still, the reputation of AI, in the business world at least, was less than pristine.^[171] Inside the field there was little agreement on the reasons for AI's failure to fulfill the dream of human level intelligence that had captured the imagination of the world in the 1960s. Together, all these factors helped to fragment AI into competing subfields focused on particular problems or approaches, sometimes even under new names that disguised the tarnished pedigree of "artificial intelligence".^[172] AI was both more cautious and more successful than it had ever been.

Milestones and Moore's law

On 11 May 1997, Deep Blue became the first computer chess-playing system to beat a reigning world chess champion, Garry Kasparov.^[173] The super computer was a specialized version of a framework produced by IBM, and was capable of processing twice as many moves per second as it had during the first match (which Deep Blue had lost), reportedly 200,000,000 moves per second. The event was broadcast live over the internet and received over 74 million hits.^[174]

In 2005, a Stanford robot won the DARPA Grand Challenge by driving autonomously for 131 miles along an unrehearsed desert trail.^[175] Two years later, a team from CMU won the DARPA Urban Challenge by autonomously navigating 55 miles in an Urban environment while adhering to traffic hazards and all traffic laws.^[176] In February 2011, in a Jeopardy! quiz show exhibition match, IBM's question answering system, Watson, defeated the two greatest Jeopardy! champions, Brad Rutter and Ken Jennings, by a significant margin.^[177]

These successes were not due to some revolutionary new paradigm, but mostly on the tedious application of engineering skill and on the tremendous increase in the speed and capacity of computer by the 90s.^[178] In fact, Deep Blue's computer was 10 million times faster than the Ferranti Mark 1 that Christopher Strachey taught to play chess in 1951.^[179] This dramatic increase is measured by Moore's law, which predicts that the speed and memory capacity of computers doubles every two years, as a result of metal-oxide-semiconductor (MOS) transistor counts doubling every two years. The fundamental problem of "raw computer power" was slowly being overcome.

Intelligent agents

A new paradigm called "intelligent agents" became widely accepted during the 1990s.^[180] Although earlier researchers had proposed modular "divide and conquer" approaches to AI,^[181] the intelligent agent did not reach its modern form until Judea Pearl, Allen Newell, Leslie P. Kaelbling, and others brought concepts from decision theory and economics into the study of AI.^[182] When the economist's definition of a rational agent was married to computer science's definition of an object or module, the intelligent agent paradigm was complete.

An intelligent agent is a system that perceives its environment and takes actions which maximize its chances of success. By this definition, simple programs that solve specific problems are "intelligent agents", as are human beings and organizations of human beings, such as firms. The intelligent agent paradigm defines AI research as "the study of intelligent agents". This is a generalization of some earlier definitions of AI: it goes beyond studying human intelligence; it studies all kinds of intelligence.^[183]

The paradigm gave researchers license to study isolated problems and find solutions that were both verifiable and useful. It provided a common language to describe problems and share their solutions with each other, and with other fields that also used concepts of abstract agents, like economics and control theory. It was hoped that a complete agent architecture (like Newell's SOAR) would one day allow researchers to build more versatile and intelligent systems out of interacting intelligent agents.^{[182][184]}

Probabilistic reasoning and greater rigor

AI researchers began to develop and use sophisticated mathematical tools more than they ever had in the past.^[185] There was a widespread realization that many of the problems that AI needed to solve were already being worked on by researchers in fields like mathematics, electrical engineering, economics or operations research. The shared mathematical language allowed both a higher level of collaboration with

more established and successful fields and the achievement of results which were measurable and provable; AI had become a more rigorous "scientific" discipline. Russell & Norvig (2003) describe this as nothing less than a "revolution". They had argued in their 2002 textbook that this increased rigor could be viewed plausibly as a "victory of the neats,"^[186] but subsequently qualified that by saying, in their 2020 AI textbook, that "The present emphasis on deep learning may represent a resurgence of the scruffies."^[187]

Judea Pearl's influential 1988 book^[188] brought probability and decision theory into AI. Among the many new tools in use were Bayesian networks, hidden Markov models, information theory, stochastic modeling and classical optimization. Precise mathematical descriptions were also developed for "computational intelligence" paradigms like neural networks and evolutionary algorithms.^[189]

AI behind the scenes

Algorithms originally developed by AI researchers began to appear as parts of larger systems. AI had solved a lot of very difficult problems^[190] and their solutions proved to be useful throughout the technology industry,^[191] such as data mining, industrial robotics, logistics,^[192] speech recognition,^[193] banking software,^[194] medical diagnosis^[194] and Google's search engine.^[195]

The field of AI received little or no credit for these successes in the 1990s and early 2000s. Many of AI's greatest innovations have been reduced to the status of just another item in the tool chest of computer science.^[196] Nick Bostrom explains "A lot of cutting edge AI has filtered into general applications, often without being called AI because once something becomes useful enough and common enough it's not labeled AI anymore."^[197]

Many researchers in AI in the 1990s deliberately called their work by other names, such as informatics, knowledge-based systems, cognitive systems or computational intelligence. In part, this may have been because they considered their field to be fundamentally different from AI, but also the new names help to procure funding. In the commercial world at least, the failed promises of the AI Winter continued to haunt AI research into the 2000s, as the *New York Times* reported in 2005: "Computer scientists and software engineers avoided the term artificial intelligence for fear of being viewed as wild-eyed dreamers."^{[198][199][200][201]}

Predictions (or "Where is HAL 9000?")

In 1968, Arthur C. Clarke and Stanley Kubrick had imagined that, by the year 2001, a machine would exist with an intelligence that matched or exceeded the capability of human beings. The character they created, HAL 9000, was based on a belief shared by many leading AI researchers that such a machine would exist by the year 2001.^{[202][203]}

In 2001, AI founder Marvin Minsky asked "So the question is why didn't we get HAL in 2001?"^[204] Minsky believed that the answer is that the central problems, like commonsense reasoning, were being neglected, while most researchers pursued things like commercial applications of neural nets or genetic algorithms. John McCarthy, on the other hand, still blamed the qualification problem.^[205] For Ray Kurzweil, the issue is computer power and, using Moore's Law, he predicted that machines with human-level intelligence will appear by 2029.^[206] Jeff Hawkins argued that neural net research ignores the essential properties of the human cortex, preferring simple models that have been successful at solving simple problems.^[207] There were many other explanations and for each there was a corresponding research program underway.

Deep learning, big data and artificial general intelligence: 2011–present

In the first decades of the 21st century, access to large amounts of data (known as "big data"), cheaper and faster computers and advanced machine learning techniques were successfully applied to many problems throughout the economy. In fact, McKinsey Global Institute estimated in their famous paper "Big data: The next frontier for innovation, competition, and productivity" that "by 2009, nearly all sectors in the US economy had at least an average of 200 terabytes of stored data".

By 2016, the market for AI-related products, hardware, and software reached more than 8 billion dollars, and the New York Times reported that interest in AI had reached a "frenzy".^[208] The applications of big data began to reach into other fields as well, such as training models in ecology^[209] and for various applications in economics.^[210] Advances in deep learning (particularly deep convolutional neural networks and recurrent neural networks) drove progress and research in image and video processing, text analysis, and even speech recognition.^[211]

Deep learning

Deep learning is a branch of machine learning that models high level abstractions in data by using a deep graph with many processing layers.^[211] According to the Universal approximation theorem, deep-ness isn't necessary for a neural network to be able to approximate arbitrary continuous functions. Even so, there are many problems that are common to shallow networks (such as overfitting) that deep networks help avoid.^[212] As such, deep neural networks are able to realistically generate much more complex models as compared to their shallow counterparts.

However, deep learning has problems of its own. A common problem for recurrent neural networks is the vanishing gradient problem, which is where gradients passed between layers gradually shrink and literally disappear as they are rounded off to zero. There have been many methods developed to approach this problem, such as Long short-term memory units.

State-of-the-art deep neural network architectures can sometimes even rival human accuracy in fields like computer vision, specifically on things like the MNIST database, and traffic sign recognition.^[213]

Language processing engines powered by smart search engines can easily beat humans at answering general trivia questions (such as IBM Watson), and recent developments in deep learning have produced astounding results in competing with humans, in things like Go, and Doom (which, being a first-person shooter game, has sparked some controversy).^{[214][215][216][217]}

Big Data

Big data refers to a collection of data that cannot be captured, managed, and processed by conventional software tools within a certain time frame. It is a massive amount of decision-making, insight, and process optimization capabilities that require new processing models. In the Big Data Era written by Victor Meyer Schonberg and Kenneth Cooke, big data means that instead of random analysis (sample survey), all data is used for analysis. The 5V characteristics of big data (proposed by IBM): *Volume*, *Velocity*, *Variety*^[218], *Value*^[219], *Veracity*^[220]. The strategic significance of big data technology is not to master huge data information, but to specialize in these meaningful data. In other words, if big data is likened to an industry, the key to realizing profitability in this industry is to increase the "process capability" of the data and realize the "value added" of the data through "processing".

Artificial general intelligence

General intelligence is the ability to solve *any* problem, rather than finding a solution to a particular problem. Artificial general intelligence (or "AGI") is a program which can apply intelligence to a wide variety of problems, in much the same ways humans can.

Ben Goertzel and others argued in the early 2000s that AI research had largely given up on the field's original goal of creating artificial general intelligence. AGI research was founded as a separate sub-field and by 2010 there were academic conferences, laboratories and university courses dedicated to AGI research, as well as private consortiums and new companies.

Artificial general intelligence is also referred to as "strong AI",^[221] "full AI",^[222] or synthetic intelligence^{[223][224]} as opposed to "weak AI" or "narrow AI". (Academic sources reserve "strong AI" to refer to machines capable of experiencing consciousness.)

Foundation models, which are large artificial intelligence models trained on vast quantities of unlabeled data that can be adapted to a wide range of downstream tasks, began to be developed in 2018. Models such as GPT-3 released by OpenAI in 2020, and Gato released by DeepMind in 2022, have been described as important milestones on the path to artificial general intelligence.^{[225][226]}

See also

- Outline of artificial intelligence
 - Progress in artificial intelligence
 - Timeline of artificial intelligence
 - History of natural language processing
- Timeline of machine learning
- History of knowledge representation and reasoning

Notes

1. Kaplan, Andreas; Haenlein, Michael (2019). "Siri, Siri, in my hand: Who's the fairest in the land? On the interpretations, illustrations, and implications of artificial intelligence". *Business Horizons*. **62**: 15–25. doi:10.1016/j.bushor.2018.08.004 (https://doi.org/10.1016%2Fj.bushor.2018.08.004).
2. Newquist 1994, pp. 143–156
3. Newquist 1994, pp. 144–152
4. The Talos episode in *Argonautica* 4
5. *Bibliothèque* 1.9.26
6. Rhodios, Apollonios. (2007). *The Argonautika : Expanded Edition*. University of California Press. p. 355. ISBN 978-0-520-93439-9. OCLC 811491744 (https://www.worldcat.org/oclc/811491744).
7. Morford, Mark (2007). *Classical mythology*. Oxford. p. 184. ISBN 978-0-19-085164-4. OCLC 1102437035 (https://www.worldcat.org/oclc/1102437035).
8. Kressel, Matthew (1 October 2015). "36 Days of Judaic Myth: Day 24, The Golem of Prague" (https://www.matthewkressel.net/2015/10/01/36-days-of-judaic-myth-day-24-the-golem-of-prague/). *Matthew Kressel*. Retrieved 15 March 2020.

9. Newquist., HP (1994). *The Brain Makers : [Genius, Ego, And Greed In The Quest For Machines That Think]* (<http://worldcat.org/oclc/313139906>). Sams Publ. ISBN 0-672-30412-0. OCLC 313139906 (<https://www.worldcat.org/oclc/313139906>).
10. "GOLEM - JewishEncyclopedia.com" (<http://www.jewishencyclopedia.com/articles/6777-golem>). *www.jewishencyclopedia.com*. Retrieved 15 March 2020.
11. Newquist 1994, pp. 38
12. "Sanhedrin 65b" (<https://www.sefaria.org/Sanhedrin.65b?lang=bi>). *www.sefaria.org*. Retrieved 15 March 2020.
13. *The alchemy reader : from Hermes Trismegistus to Isaac Newton*. Linden, Stanton J., 1935-. New York: Cambridge University Press. 2003. pp. Ch. 18. ISBN 0-521-79234-7. OCLC 51210362 (<https://www.worldcat.org/oclc/51210362>).
14. O'Connor, Kathleen Malone (1 January 1994). "The alchemical creation of life (takwin) and other concepts of Genesis in medieval Islam" (<https://repository.upenn.edu/dissertations/AAI9503804>). *Dissertations Available from ProQuest*: 1–435.
15. Goethe, Johann Wolfgang von (1890). *Faust; a tragedy. Translated, in the original metres ... by Bayard Taylor. Authorised ed., published by special arrangement with Mrs. Bayard Taylor. With a biographical introd* (<https://archive.org/details/fausttragedytran00goetuoft>). Robarts - University of Toronto. London Ward, Lock.
16. McCorduck 2004, pp. 17–25.
17. Butler 1863.
18. Newquist 1994, pp. 65
19. Cave, Stephen; Dihal, Kanta (2019). "Hopes and fears for intelligent machines in fiction and reality" (<https://www.nature.com/articles/s42256-019-0020-9>). *Nature Machine Intelligence*. 1 (2): 74–78. doi:10.1038/s42256-019-0020-9 (<https://doi.org/10.1038%2Fs42256-019-0020-9>). ISSN 2522-5839 (<https://www.worldcat.org/issn/2522-5839>). S2CID 150700981 (<https://api.semanticscholar.org/CorpusID:150700981>).
20. Needham 1986, p. 53
21. McCorduck 2004, p. 6
22. Nick 2005.
23. McCorduck 2004, p. 17 and see also Levitt 2000
24. Newquist 1994, pp. 30
25. Quoted in McCorduck 2004, p. 8. Crevier 1993, p. 1 and McCorduck 2004, pp. 6–9 discusses sacred statues.
26. Other important automata were built by Haroun al-Rashid (McCorduck 2004, p. 10), Jacques de Vaucanson (Newquist 1994, pp. 40), (McCorduck 2004, p. 16) and Leonardo Torres y Quevedo (McCorduck 2004, pp. 59–62)
27. Butler, E. M. (Eliza Marian) (1948). *The myth of the magus*. London: Cambridge University Press. ISBN 0-521-22564-7. OCLC 5063114 (<https://www.worldcat.org/oclc/5063114>).
28. Hollander, Lee M. (1964). *Heimskringla; history of the kings of Norway*. Austin: Published for the American-Scandinavian Foundation by the University of Texas Press. ISBN 0-292-73061-6. OCLC 638953 (<https://www.worldcat.org/oclc/638953>).
29. Berlinski 2000
30. Cfr. Carreras Artau, Tomás y Joaquín. *Historia de la filosofía española. Filosofía cristiana de los siglos XIII al XV*. Madrid, 1939, Volume I
31. Bonner, Anthony, *The Art and Logic of Ramón Llull: A User's Guide*, Brill, 2007.
32. Anthony Bonner (ed.), *Doctor Illuminatus. A Ramon Llull Reader* (Princeton University 1985). Vid. "Llull's Influence: The History of Lullism" at 57–71

33. 17th century mechanism and AI:

- McCorduck 2004, pp. 37–46
- Russell & Norvig 2003, p. 6
- Haugeland 1986, chpt. 2
- Buchanan 2005, p. 53

34. Hobbes and AI:

- McCorduck 2004, p. 42
- Hobbes 1651, chapter 5

35. Leibniz and AI:

- McCorduck 2004, p. 41
- Russell & Norvig 2003, p. 6
- Berlinski 2000, p. 12
- Buchanan 2005, p. 53

36. The Lambda calculus was especially important to AI, since it was an inspiration for Lisp (the most important programming language used in AI). (Crevier 1993, pp. 190–196, 61)

37. The original photo can be seen in the article: Rose, Allen (April 1946). "Lightning Strikes Mathematics" (<https://books.google.com/books?id=niEDAAAAMBAJ&q=eniac+intitle:popular+intitle:science&pg=PA83>). *Popular Science*: 83–86. Retrieved 15 April 2012.

38. Newquist 1994, pp. 56

39. The Turing machine: McCorduck 2004, pp. 63–64, Crevier 1993, pp. 22–24, Russell & Norvig 2003, p. 8 and see Turing 1936

40. Newquist 1994, pp. 67

41. Menabrea 1843

42. "The John Gabriel Byrne Computer Science Collection" (<https://web.archive.org/web/20190416071721/https://www.scss.tcd.ie/SCSSTreasuresCatalog/miscellany/TCD-SCSS-X.20121208.002/TCD-SCSS-X.20121208.002.pdf>) (PDF). Archived from the original (<https://scss.tcd.ie/SCSSTreasuresCatalog/miscellany/TCD-SCSS-X.20121208.002/TCD-SCSS-X.20121208.002.pdf>) on 16 April 2019. Retrieved 8 August 2019.

43. Ingenious Ireland (<http://ingeniousireland.ie/2012/10/1909-a-novel-irish-computer/>)

44. "Percy Ludgate's Analytical Machine" (<http://www.fano.co.uk/ludgate/>). *fano.co.uk. From Analytical Engine to Electronic Digital Computer: The Contributions of Ludgate, Torres, and Bush* by Brian Randell, 1982, Ludgate: pp. 4–5, Quevedo: pp. 6, 11–13, Bush: pp. 13, 16–17. Retrieved 29 October 2018.

45. Williams, Andrew (16 March 2017). *History of Digital Games: Developments in Art, Design and Interaction* (<https://books.google.com/books?id=xLVdDgAAQBAJ&dq=Torres+and+his+remarkable+automatic+devices.+Issue+2079+of+Scientific+American,+1915&pg=PA30>). CRC Press. ISBN 9781317503811.

46. Brian Randell, From Analytical Engine to Electronic Digital Computer: The Contributions of Ludgate, Torres and Bush. *Annals of the History of Computing*, Vol. 4, No. 4, Oct. 1982

47. McCorduck 2004, pp. 61–62, 64–66, Russell & Norvig 2003, pp. 14–15

48. McCorduck (2004), pp. 76–80

49. "Kaplan Andreas, Artificial Intelligence, Business and Civilization - Our Fate Made in Machines" (<https://www.routledge.com/Artificial-Intelligence-Business-and-Civilization-Our-Fate-Made-in-Machines/Kaplan/p/book/9781032155319>). Retrieved 11 March 2022.

50. McCorduck 2004, pp. 51–57, 80–107, Crevier 1993, pp. 27–32, Russell & Norvig 2003, pp. 15, 940, Moravec 1988, p. 3, Cordeschi, 2002 & Chap. 5.

51. McCorduck 2004, p. 98, Crevier 1993, pp. 27–28, Russell & Norvig 2003, pp. 15, 940, Moravec 1988, p. 3, Cordeschi, 2002 & Chap. 5.
52. McCulloch, Warren S.; Pitts, Walter (1 December 1943). "A logical calculus of the ideas immanent in nervous activity". *Bulletin of Mathematical Biophysics*. **5** (4): 115–133. doi:10.1007/BF02478259 (<https://doi.org/10.1007%2FBF02478259>). ISSN 1522-9602 (<http://www.worldcat.org/issn/1522-9602>).
53. Piccinini, Gualtiero (1 August 2004). "The First Computational Theory of Mind and Brain: A Close Look at McCulloch and Pitts's "Logical Calculus of Ideas Immanent in Nervous Activity" ". *Synthese*. **141** (2): 175–215. doi:10.1023/B:SYNT.0000043018.52445.3e (<https://doi.org/10.1023%2FB%3ASYNT.0000043018.52445.3e>). ISSN 1573-0964 (<https://www.worldcat.org/issn/1573-0964>). S2CID 10442035 (<https://api.semanticscholar.org/CorpusID:10442035>).
54. McCorduck 2004, pp. 51–57, 88–94, Crevier 1993, p. 30, Russell & Norvig 2003, p. 15–16, Cordeschi, 2002 & Chap. 5 and see also Pitts & McCullough 1943
55. McCorduck 2004, p. 102, Crevier 1993, pp. 34–35 and Russell & Norvig 2003, p. 17
56. McCorduck 2004, pp. 70–72, Crevier 1993, p. 22–25, Russell & Norvig 2003, pp. 2–3 and 948, Haugeland 1985, pp. 6–9, Cordeschi 2002, pp. 170–176. See also Turing 1950
57. Newquist 1994, pp. 92–98
58. Norvig & Russell (2003, p. 948) claim that Turing answered all the major objections to AI that have been offered in the years since the paper appeared.
59. See "A Brief History of Computing" (http://www.alanturing.net/turing_archive/pages/Reference%20Articles/BriefHistofComp.html) at AlanTuring.net.
60. Schaeffer, Jonathan. *One Jump Ahead:: Challenging Human Supremacy in Checkers*, 1997,2009, Springer, ISBN 978-0-387-76575-4. Chapter 6.
61. McCorduck 2004, pp. 137–170, Crevier, pp. 44–47
62. McCorduck 2004, pp. 123–125, Crevier 1993, pp. 44–46 and Russell & Norvig 2003, p. 17
63. Quoted in Crevier 1993, p. 46 and Russell & Norvig 2003, p. 17
64. Russell & Norvig 2003, p. 947,952
65. McCorduck 2004, pp. 111–136, Crevier 1993, pp. 49–51 and Russell & Norvig 2003, p. 17 Newquist 1994, pp. 91–112
66. See McCarthy et al. 1955. Also see Crevier 1993, p. 48 where Crevier states "[the proposal] later became known as the 'physical symbol systems hypothesis'". The physical symbol system hypothesis was articulated and named by Newell and Simon in their paper on GPS. (Newell & Simon 1963) It includes a more specific definition of a "machine" as an agent that manipulates symbols. See the philosophy of artificial intelligence.
67. McCorduck (2004, pp. 129–130) discusses how the Dartmouth conference alumni dominated the first two decades of AI research, calling them the "invisible college".
68. "I won't swear and I hadn't seen it before," McCarthy told Pamela McCorduck in 1979. (McCorduck 2004, p. 114) However, McCarthy also stated unequivocally "I came up with the term" in a CNET interview. (Skillings 2006)
69. Crevier (1993, pp. 49) writes "the conference is generally recognized as the official birthdate of the new science."
70. McCarthy, John (1988). "Review of *The Question of Artificial Intelligence*". *Annals of the History of Computing*. **10** (3): 224–229., collected in McCarthy, John (1996). "10. Review of *The Question of Artificial Intelligence*". *Defending AI Research: A Collection of Essays and Reviews*. CSLI., p. 73 "[O]ne of the reasons for inventing the term "artificial intelligence" was to escape association with "cybernetics". Its concentration on analog feedback seemed misguided, and I wished to avoid having either to accept Norbert (not Robert) Wiener as a guru or having to argue with him."

71. Russell and Norvig write "it was astonishing whenever a computer did anything remotely clever." Russell & Norvig 2003, p. 18
72. Crevier 1993, pp. 52–107, Moravec 1988, p. 9 and Russell & Norvig 2003, p. 18–21
73. McCorduck 2004, p. 218, Newquist 1994, pp. 91–112, Crevier 1993, pp. 108–109 and Russell & Norvig 2003, p. 21
74. Crevier 1993, pp. 52–107, Moravec 1988, p. 9
75. Means-ends analysis, reasoning as search: McCorduck 2004, pp. 247–248. Russell & Norvig 2003, pp. 59–61
76. Heuristic: McCorduck 2004, p. 246, Russell & Norvig 2003, pp. 21–22
77. GPS: McCorduck 2004, pp. 245–250, Crevier 1993, p. GPS?, Russell & Norvig 2003, p. GPS?
78. Crevier 1993, pp. 51–58, 65–66 and Russell & Norvig 2003, pp. 18–19
79. McCorduck 2004, pp. 268–271, Crevier 1993, pp. 95–96, Newquist 1994, pp. 148–156, Moravec 1988, pp. 14–15
80. McCorduck 2004, p. 286, Crevier 1993, pp. 76–79, Russell & Norvig 2003, p. 19
81. Crevier 1993, pp. 79–83
82. Crevier 1993, pp. 164–172
83. McCorduck 2004, pp. 291–296, Crevier 1993, pp. 134–139
84. McCorduck 2004, pp. 299–305, Crevier 1993, pp. 83–102, Russell & Norvig 2003, p. 19 and Copeland 2000
85. McCorduck 2004, pp. 300–305, Crevier 1993, pp. 84–102, Russell & Norvig 2003, p. 19
86. "Humanoid History -WABOT-" (http://www.humanoid.waseda.ac.jp/booklet/kato_2-j.html).
87. Zeghloul, Saïd; Laribi, Med Amine; Gazeau, Jean-Pierre (21 September 2015). "Robotics and Mechatronics: Proceedings of the 4th IFToMM International Symposium on Robotics and Mechatronics" (<https://books.google.com/books?id=tQqVCgAAQBAJ&pg=PA66>). Springer – via Google Books.
88. "Historical Android Projects" (<http://www.androidworld.com/prod06.htm>). *androidworld.com*.
89. *Robots: From Science Fiction to Technological Revolution* (<https://archive.org/details/robotsoftromscienc0000ichb>), page 130
90. Duffy, Vincent G. (19 April 2016). "Handbook of Digital Human Modeling: Research for Applied Ergonomics and Human Factors Engineering" (<https://books.google.com/books?id=NgLLBQAAQBAJ&pg=SA3-PA1>). CRC Press – via Google Books.
91. Simon & Newell 1958, p. 7–8 quoted in Crevier 1993, p. 108. See also Russell & Norvig 2003, p. 21
92. Simon 1965, p. 96 quoted in Crevier 1993, p. 109
93. Minsky 1967, p. 2 quoted in Crevier 1993, p. 109
94. Minsky strongly believes he was misquoted. See McCorduck 2004, pp. 272–274, Crevier 1993, p. 96 and Darrach 1970.
95. Crevier 1993, pp. 64–65
96. Crevier 1993, p. 94
97. Howe 1994
98. McCorduck 2004, p. 131, Crevier 1993, p. 51. McCorduck also notes that funding was mostly under the direction of alumni of the Dartmouth conference of 1956.
99. Crevier 1993, p. 65
100. Crevier 1993, pp. 68–71 and Turkle 1984
101. Crevier 1993, pp. 100–144 and Russell & Norvig 2003, pp. 21–22

102. McCorduck 2004, pp. 104–107, Crevier 1993, pp. 102–105, Russell & Norvig 2003, p. 22
103. Crevier 1993, pp. 163–196
104. Crevier 1993, p. 146
105. Russell & Norvig 2003, pp. 20–21 Newquist 1994, pp. 336
106. Crevier 1993, pp. 146–148, see also Buchanan 2005, p. 56: "Early programs were necessarily limited in scope by the size and speed of memory"
107. Moravec 1976. McCarthy has always disagreed with Moravec, back to their early days together at SAIL. He states "I would say that 50 years ago, the machine capability was much too small, but by 30 years ago, machine capability wasn't the real problem." in a CNET interview. (Skillings 2006)
108. Hans Moravec, *ROBOT: Mere Machine to Transcendent Mind*
109. Russell & Norvig 2003, pp. 9, 21–22 and Lighthill 1973
110. McCorduck 2004, pp. 300 & 421; Crevier 1993, pp. 113–114; Moravec 1988, p. 13; Lenat & Guha 1989, (Introduction); Russell & Norvig 2003, p. 21
111. McCorduck 2004, p. 456, Moravec 1988, pp. 15–16
112. McCarthy & Hayes 1969, Crevier 1993, pp. 117–119
113. McCorduck 2004, pp. 280–281, Crevier 1993, p. 110, Russell & Norvig 2003, p. 21 and NRC 1999 under "Success in Speech Recognition".
114. Crevier 1993, p. 117, Russell & Norvig 2003, p. 22, Howe 1994 and see also Lighthill 1973.
115. Russell & Norvig 2003, p. 22, Lighthill 1973, John McCarthy wrote in response that "the combinatorial explosion problem has been recognized in AI from the beginning" in Review of Lighthill report (<http://www-formal.stanford.edu/jmc/reviews/lighthill/lighthill.html>)
116. Crevier 1993, pp. 115–116 (on whom this account is based). Other views include McCorduck 2004, pp. 306–313 and NRC 1999 under "Success in Speech Recognition".
117. Crevier 1993, p. 115. Moravec explains, "Their initial promises to DARPA had been much too optimistic. Of course, what they delivered stopped considerably short of that. But they felt they couldn't in their next proposal promise less than in the first one, so they promised more."
118. NRC 1999 under "Shift to Applied Research Increases Investment." While the autonomous tank was a failure, the battle management system (called "DART") proved to be enormously successful, saving billions in the first Gulf War, repaying the investment and justifying the DARPA's pragmatic policy, at least as far as DARPA was concerned.
119. Lucas and Penrose' critique of AI: Crevier 1993, p. 22, Russell & Norvig 2003, pp. 949–950, Hofstadter 1980, pp. 471–477 and see Lucas 1961
120. "Know-how" is Dreyfus' term. (Dreyfus makes a distinction between "knowing how" and "knowing that", a modern version of Heidegger's distinction of ready-to-hand and present-at-hand.) (Dreyfus & Dreyfus 1986)
121. Dreyfus' critique of artificial intelligence: McCorduck 2004, pp. 211–239, Crevier 1993, pp. 120–132, Russell & Norvig 2003, pp. 950–952 and see Dreyfus 1965, Dreyfus 1972, Dreyfus & Dreyfus 1986
122. Searle's critique of AI: McCorduck 2004, pp. 443–445, Crevier 1993, pp. 269–271, Russell & Norvig 2004, pp. 958–960 and see Searle 1980
123. Quoted in Crevier 1993, p. 143
124. Quoted in Crevier 1993, p. 122
125. Newquist 1994, pp. 276
126. "I became the only member of the AI community to be seen eating lunch with Dreyfus. And I deliberately made it plain that theirs was not the way to treat a human being." Joseph Weizenbaum, quoted in Crevier 1993, p. 123.

127. Colby, Watt & Gilbert 1966, p. 148. Weizenbaum referred to this text in Weizenbaum 1976, pp. 5, 6. Colby and his colleagues later also developed chatterbot-like "computer simulations of paranoid processes (PARRY)" to "make intelligible paranoid processes in explicit symbol processing terms." (Colby 1974, p. 6)
128. Weizenbaum's critique of AI: McCorduck 2004, pp. 356–373, Crevier 1993, pp. 132–144, Russell & Norvig 2003, p. 961 and see Weizenbaum 1976
129. McCorduck 2004, p. 51, Russell & Norvig 2003, pp. 19, 23
130. McCorduck 2004, p. 51, Crevier 1993, pp. 190–192
131. Crevier 1993, pp. 193–196
132. Crevier 1993, pp. 145–149, 258–63
133. Wason (1966) showed that people do poorly on completely abstract problems, but if the problem is restated to allow the use of intuitive **social intelligence**, performance dramatically improves. (See Wason selection task) Tversky, Slovic & Kahnemann (1982) have shown that people are terrible at elementary problems that involve uncertain reasoning. (See list of cognitive biases for several examples). Eleanor Rosch's work is described in Lakoff 1987
134. An early example of McCathy's position was in the journal Science where he said "This is AI, so we don't care if it's psychologically real" (Kolata 2012), and he recently reiterated his position at the AI@50 conference where he said "Artificial intelligence is not, by definition, simulation of human intelligence" (Maker 2006).
135. Crevier 1993, pp. 175
136. Neat vs. scruffy: McCorduck 2004, pp. 421–424 (who picks up the state of the debate in 1984). Crevier 1993, pp. 168 (who documents Schank's original use of the term). Another aspect of the conflict was called "the procedural/declarative distinction" but did not prove to be influential in later AI research.
137. McCorduck 2004, pp. 305–306, Crevier 1993, pp. 170–173, 246 and Russell & Norvig 2003, p. 24. Minsky's frame paper: Minsky 1974.
138. Newquist 1994, pp. 189–192
139. McCorduck 2004, pp. 327–335 (Dendral), Crevier 1993, pp. 148–159, Newquist 1994, p. 271, Russell & Norvig 2003, pp. 22–23
140. Crevier 1993, pp. 158–159 and Russell & Norvig 2003, p. 23–24
141. Crevier 1993, p. 198
142. Newquist 1994, pp. 259
143. McCorduck 2004, pp. 434–435, Crevier 1993, pp. 161–162, 197–203, Newquist 1994, pp. 275 and Russell & Norvig 2003, p. 24
144. McCorduck 2004, p. 299
145. McCorduck 2004, pp. 421
146. Knowledge revolution: McCorduck 2004, pp. 266–276, 298–300, 314, 421, Newquist 1994, pp. 255–267, Russell & Norvig, pp. 22–23
147. Cyc: McCorduck 2004, p. 489, Crevier 1993, pp. 239–243, Newquist 1994, pp. 431–455, Russell & Norvig 2003, p. 363–365 and Lenat & Guha 1989
148. "Chess: Checkmate" (http://www.computerhistory.org/about/press_relations/media_kit/chess/s5_chess_exhibit_text_Endgame.pdf) (PDF). Retrieved 1 September 2007.
149. McCorduck 2004, pp. 436–441, Newquist 1994, pp. 231–240, Crevier 1993, pp. 211, Russell & Norvig 2003, p. 24 and see also Feigenbaum & McCorduck 1983
150. Crevier 1993, pp. 195
151. Crevier 1993, pp. 240.
152. Russell & Norvig 2003, p. 25

153. McCorduck 2004, pp. 426–432, NRC 1999 under "Shift to Applied Research Increases Investment"
154. Crevier 1993, pp. 214–215.
155. Crevier 1993, pp. 215–216.
156. Mead, Carver A.; Ismail, Mohammed (8 May 1989). *Analog VLSI Implementation of Neural Systems* (<http://fenetic.net/irc/Christopher%20R.%20Carroll%20Carver%20Mead%20Mohammed%20Ismail%20Analog%20VLSI%20Implementation%20of%20Neural%20Systems.pdf>) (PDF). The Kluwer International Series in Engineering and Computer Science. Vol. 80. Norwell, MA: Kluwer Academic Publishers. doi:10.1007/978-1-4613-1639-8 (<https://doi.org/10.1007/978-1-4613-1639-8>). ISBN 978-1-4613-1639-8.
157. Newquist 1994, pp. 501
158. Crevier 1993, pp. 203. AI winter was first used as the title of a seminar on the subject for the Association for the Advancement of Artificial Intelligence.
159. Newquist 1994, pp. 359–379, McCorduck 2004, p. 435, Crevier 1993, pp. 209–210
160. McCorduck 2004, p. 435 (who cites institutional reasons for their ultimate failure), Newquist 1994, pp. 258–283 (who cites limited deployment within corporations), Crevier 1993, pp. 204–208 (who cites the difficulty of truth maintenance, i.e., learning and updating), Lenat & Guha 1989, Introduction (who emphasizes the brittleness and the inability to handle excessive qualification.)
161. McCorduck 2004, pp. 430–431
162. McCorduck 2004, p. 441, Crevier 1993, p. 212. McCorduck writes "Two and a half decades later, we can see that the Japanese didn't quite meet all of those ambitious goals."
163. Newquist 1994, pp. 476
164. Newquist 1994, pp. 440
165. McCorduck 2004, pp. 454–462
166. Moravec (1988, p. 20) writes: "I am confident that this bottom-up route to artificial intelligence will one day meet the traditional top-down route more than half way, ready to provide the real world competence and the commonsense knowledge that has been so frustratingly elusive in reasoning programs. Fully intelligent machines will result when the metaphorical golden spike is driven uniting the two efforts."
167. Crevier 1993, pp. 183–190.
168. <http://people.csail.mit.edu/brooks/papers/elephants.pdf>
169. Brooks 1990, p. 3
170. See, for example, Lakoff & Turner 1999
171. Newquist 1994, pp. 511
172. McCorduck (2004, p. 424) discusses the fragmentation and the abandonment of AI's original goals.
173. McCorduck 2004, pp. 480–483
174. "Deep Blue" (<http://www.research.ibm.com/deepblue/meet/html/d.3.shtml>). IBM Research. Retrieved 10 September 2010.
175. "DARPA Grand Challenge – home page" (<https://web.archive.org/web/20071031110030/http://www.darpa.mil/grandchallenge/>). Archived from the original (<https://www.darpa.mil/grandchallenge/>) on 31 October 2007.
176. "Archived copy" (<https://web.archive.org/web/20140305023802/http://archive.darpa.mil/grandchallenge/>). Archived from the original (<https://archive.darpa.mil/grandchallenge/>) on 5 March 2014. Retrieved 25 October 2011.
177. Markoff, John (16 February 2011). "On 'Jeopardy!' Watson Win Is All but Trivial" (<https://www.nytimes.com/2011/02/17/science/17jeopardy-watson.html>). *The New York Times*.

178. [Kurzweil 2005](#), p. 274 writes that the improvement in computer chess, "according to common wisdom, is governed only by the brute force expansion of computer hardware."
179. Cycle time of [Ferranti Mark 1](#) was 1.2 milliseconds, which is arguably equivalent to about 833 flops. [Deep Blue](#) ran at 11.38 gigaflops (and this does not even take into account Deep Blue's special-purpose hardware for chess). Very approximately, these differ by a factor of 10^7 .
180. [McCorduck 2004](#), pp. 471–478, [Russell & Norvig 2003](#), p. 55, where they write: "The whole-agent view is now widely accepted in the field". The intelligent agent paradigm is discussed in major AI textbooks, such as: [Russell & Norvig 2003](#), pp. 32–58, 968–972, [Poole, Mackworth & Goebel 1998](#), pp. 7–21, [Luger & Stubblefield 2004](#), pp. 235–240
181. [Carl Hewitt's Actor model](#) anticipated the modern definition of intelligent agents. ([Hewitt, Bishop & Steiger 1973](#)) Both [John Doyle \(Doyle 1983\)](#) and [Marvin Minsky's](#) popular classic *The Society of Mind* ([Minsky 1986](#)) used the word "agent". Other "modular" proposals included [Rodney Brook's subsumption architecture](#), [object-oriented programming](#) and others.
182. [Russell & Norvig 2003](#), pp. 27, 55
183. This is how the most widely accepted textbooks of the 21st century define artificial intelligence. See [Russell & Norvig 2003](#), p. 32 and [Poole, Mackworth & Goebel 1998](#), p. 1
184. [McCorduck 2004](#), p. 478
185. [McCorduck 2004](#), pp. 486–487, [Russell & Norvig 2003](#), pp. 25–26
186. Russell, Stuart J.; Norvig, Peter (1 December 2002). *Artificial Intelligence: A Modern Approach* (2nd ed.). Upper Saddle River, N.J: Prentice Hall. ISBN 978-0-13-790395-5.
187. Russell, Stuart; Norvig, Peter (28 April 2020). *Artificial Intelligence: A Modern Approach* (4th ed.). Hoboken: Pearson. ISBN 978-0-13-461099-3.
188. [Pearl 1988](#)
189. [Russell & Norvig 2003](#), p. 25–26
190. See [Applications of artificial intelligence § Computer science](#)
191. [NRC 1999](#) under "Artificial Intelligence in the 90s", and [Kurzweil 2005](#), p. 264
192. [Russell & Norvig 2003](#), p. 28
193. For the new state of the art in AI based speech recognition, see [The Economist \(2007\)](#)
194. "AI-inspired systems were already integral to many everyday technologies such as internet search engines, bank software for processing transactions and in medical diagnosis." [Nick Bostrom](#), quoted in [CNN 2006](#)
195. [Olsen \(2004\)](#), [Olsen \(2006\)](#)
196. [McCorduck 2004](#), p. 423, [Kurzweil 2005](#), p. 265, [Hofstadter 1979](#), p. 601 [Newquist 1994](#), pp. 445
197. [CNN 2006](#)
198. [Markoff 2005](#)
199. [The Economist 2007](#)
200. [Tascarella 2006](#)
201. [Newquist 1994](#), pp. 532
202. [Newquist 1994](#), pp. 134
203. [Crevier 1993](#), pp. 108–109
204. He goes on to say: "The answer is, I believe we could have ... I once went to an international conference on neural net[s]. There were 40 thousand registrants ... but ... if you had an international conference, for example, on using multiple representations for common sense reasoning, I've only been able to find 6 or 7 people in the whole world." [Minsky 2001](#)

205. Maker 2006
206. Kurzweil 2005
207. Hawkins & Blakeslee 2004
208. Steve Lohr (17 October 2016), "IBM Is Counting on Its Bet on Watson, and Paying Big Money for It" (https://www.nytimes.com/2016/10/17/technology/ibm-is-counting-on-its-bet-on-watson-and-paying-big-money-for-it.html?emc=edit_th_20161017&nl=todaysheadlines&nid=62816440), *New York Times*
209. Hampton, Stephanie E; Strasser, Carly A; Tewksbury, Joshua J; Gram, Wendy K; Budden, Amber E; Batcheller, Archer L; Duke, Clifford S; Porter, John H (1 April 2013). "Big data and the future of ecology" (<http://www.escholarship.org/uc/item/94f35801>). *Frontiers in Ecology and the Environment*. **11** (3): 156–162. doi:10.1890/120103 (<https://doi.org/10.1890%2F120103>). ISSN 1540-9309 (<https://www.worldcat.org/issn/1540-9309>).
210. "How Big Data is Changing Economies | Becker Friedman Institute" (<https://web.archive.org/web/20180618102343/https://bfi.uchicago.edu/events/how-big-data-changing-economies>). *bfi.uchicago.edu*. Archived from the original (<https://bfi.uchicago.edu/events/how-big-data-changing-economies>) on 18 June 2018. Retrieved 9 June 2017.
211. LeCun, Yann; Bengio, Yoshua; Hinton, Geoffrey (2015). "Deep learning". *Nature*. **521** (7553): 436–444. Bibcode:2015Natur.521..436L (<https://ui.adsabs.harvard.edu/abs/2015Natur.521..436L>). doi:10.1038/nature14539 (<https://doi.org/10.1038%2Fnature14539>). PMID 26017442 (<https://pubmed.ncbi.nlm.nih.gov/26017442>). S2CID 3074096 (<https://api.semanticscholar.org/CorpusID:3074096>).
212. Baral, Chitta; Fuentes, Olac; Kreinovich, Vladik (June 2015). "Why Deep Neural Networks: A Possible Theoretical Explanation" (http://digitalcommons.utep.edu/cs_techrep/975/). *Departmental Technical Reports (Cs)*. Retrieved 9 June 2017.
213. Ciregan, D.; Meier, U.; Schmidhuber, J. (June 2012). *Multi-column deep neural networks for image classification. 2012 IEEE Conference on Computer Vision and Pattern Recognition*. pp. 3642–3649. arXiv:1202.2745 (<https://arxiv.org/abs/1202.2745>). Bibcode:2012arXiv1202.2745C (<https://ui.adsabs.harvard.edu/abs/2012arXiv1202.2745C>). CiteSeerX 10.1.1.300.3283 (<https://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.300.3283>). doi:10.1109/cvpr.2012.6248110 (<https://doi.org/10.1109%2Fcvpr.2012.6248110>). ISBN 978-1-4673-1228-8. S2CID 2161592 (<https://api.semanticscholar.org/CorpusID:2161592>).
214. Markoff, John (16 February 2011). "On 'Jeopardy!' Watson Win Is All but Trivial" (<https://www.nytimes.com/2011/02/17/science/17jeopardy-watson.html>). *The New York Times*. ISSN 0362-4331 (<https://www.worldcat.org/issn/0362-4331>). Retrieved 10 June 2017.
215. "AlphaGo: Mastering the ancient game of Go with Machine Learning" (<https://research.googleblog.com/2016/01/alphago-mastering-ancient-game-of-go.html>). *Research Blog*. Retrieved 10 June 2017.
216. "Innovations of AlphaGo | DeepMind" (<https://deepmind.com/blog/innovations-alphago/>). *DeepMind*. Retrieved 10 June 2017.
217. University, Carnegie Mellon. "Computer Out-Plays Humans in 'Doom'-CMU News - Carnegie Mellon University" (<https://www.cmu.edu/news/stories/archives/2016/september/AI-agent-survives-doom.html>). *www.cmu.edu*. Retrieved 10 June 2017.
218. Laney, Doug (2001). "3D data management: Controlling data volume, velocity and variety". *META Group Research Note*. **6** (70).
219. Marr, Bernard (6 March 2014). "Big Data: The 5 Vs Everyone Must Know" (<https://www.linkedin.com/pulse/20140306073407-64875646-big-data-the-5-vs-everyone-must-know/>).
220. Goes, Paulo B. (2014). "Design science research in top information systems journals". *MIS Quarterly: Management Information Systems*. **38** (1).

221. (Kurzweil 2005, p. 260) or see *Advanced Human Intelligence* (http://crnano.typepad.com/crnano/blog/2005/08/advanced_human_.html) where he defines strong AI as "machine intelligence with the full range of human intelligence."
222. "TEDx Talks - YouTube" (<https://www.youtube.com/user/TEDxTalks>). *www.youtube.com*.
223. Haugeland 1985, p. 255.
224. Poole, Mackworth & Goebel 1998, p. 1.
225. "We're entering the AI twilight zone between narrow and general AI" (<https://venturebeat.com/2020/09/03/were-entering-the-ai-twilight-zone-between-narrow-and-general-ai/>). *VentureBeat*. 3 September 2020.
226. "'The Game is Over': Google's DeepMind says it is on verge of achieving human-level AI" (<https://www.independent.co.uk/tech/ai-deepmind-artificial-general-intelligence-b2080740.html>). *The Independent*. 23 May 2022.

References

- Berlinski, David (2000), *The Advent of the Algorithm* (<https://archive.org/details/adventofalgorith0000berl>), Harcourt Books, ISBN 978-0-15-601391-8, OCLC 46890682 (<https://www.worldcat.org/oclc/46890682>).
- Buchanan, Bruce G. (Winter 2005), "A (Very) Brief History of Artificial Intelligence" (<https://web.archive.org/web/20070926023314/http://www.aaai.org/AITopics/assets/PDF/AIMag26-04-016.pdf>) (PDF), *AI Magazine*, pp. 53–60, archived from the original (<http://www.aaai.org/AITopics/assets/PDF/AIMag26-04-016.pdf>) (PDF) on 26 September 2007, retrieved 30 August 2007.
- Brooks, Rodney (1990), "Elephants Don't Play Chess" (<http://people.csail.mit.edu/brooks/papers/elephants.pdf>) (PDF), *Robotics and Autonomous Systems*, 6 (1–2): 3–15, CiteSeerX 10.1.1.588.7539 (<https://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.588.7539>), doi:10.1016/S0921-8890(05)80025-9 (<https://doi.org/10.1016%2FS0921-8890%2805%2980025-9>), retrieved 30 August 2007.
- Butler, Samuel (13 June 1863), "Darwin Among the Machines" (<http://www.nzetc.org/tm/scholarly/tei-ButFir-t1-g1-t1-g1-t4-body.html>), *The Press, Christchurch, New Zealand*, retrieved 10 October 2008.
- Colby, Kenneth M.; Watt, James B.; Gilbert, John P. (1966), "A Computer Method of Psychotherapy: Preliminary Communication" (<https://exhibits.stanford.edu/feigenbaum/catalog/hk334rq4790>), *The Journal of Nervous and Mental Disease*, vol. 142, no. 2, pp. 148–152, doi:10.1097/00005053-196602000-00005 (<https://doi.org/10.1097%2F00005053-196602000-00005>), PMID 5936301 (<https://pubmed.ncbi.nlm.nih.gov/5936301/>), S2CID 36947398 (<https://api.semanticscholar.org/CorpusID:36947398>).
- Colby, Kenneth M. (September 1974), *Ten Criticisms of Parry* (<http://i.stanford.edu/pub/cstr/reports/cs/tr/74/457/CS-TR-74-457.pdf>) (PDF), Stanford Artificial Intelligence Laboratory, REPORT NO. STAN-CS-74-457, retrieved 17 June 2018.
- CNN (26 July 2006), *AI set to exceed human brain power* (<http://www.cnn.com/2006/TECH/science/07/24/ai.bostrom/>), CNN.com, retrieved 16 October 2007.
- Copeland, Jack (2000), *Micro-World AI* (http://www.alanturing.net/turing_archive/pages/Reference%20Articles/what_is_AI/What%20is%20AI06.html), retrieved 8 October 2008.
- Cordeschi, Roberto (2002), *The Discovery of the Artificial*, Dordrecht: Kluwer..
- Crevier, Daniel (1993). *AI: The Tumultuous Search for Artificial Intelligence*. New York, NY: BasicBooks. ISBN 0-465-02997-3.
- Darrach, Brad (20 November 1970), "Meet Shaky, the First Electronic Person", *Life Magazine*, pp. 58–68.

- Doyle, J. (1983), "What is rational psychology? Toward a modern mental philosophy", *AI Magazine*, vol. 4, no. 3, pp. 50–53.
- Dreyfus, Hubert (1965), *Alchemy and AI*, RAND Corporation Memo.
- Dreyfus, Hubert (1972), *What Computers Can't Do*, New York: MIT Press, ISBN 978-0-06-090613-9, OCLC 5056816 (<https://www.worldcat.org/oclc/5056816>).
- The Economist (7 June 2007), "Are You Talking to Me?" (http://www.economist.com/science/tq/displaystory.cfm?story_id=9249338), *The Economist*, retrieved 16 October 2008.
- Feigenbaum, Edward A.; McCorduck, Pamela (1983), *The Fifth Generation: Artificial Intelligence and Japan's Computer Challenge to the World*, Michael Joseph, ISBN 978-0-7181-2401-4.
- Hawkins, Jeff; Blakeslee, Sandra (2004), *On Intelligence*, New York, NY: Owl Books, ISBN 978-0-8050-7853-4, OCLC 61273290 (<https://www.worldcat.org/oclc/61273290>).
- Hebb, D.O. (1949), *The Organization of Behavior*, New York: Wiley, ISBN 978-0-8058-4300-2, OCLC 48871099 (<https://www.worldcat.org/oclc/48871099>).
- Hewitt, Carl; Bishop, Peter; Steiger, Richard (1973), *A Universal Modular Actor Formalism for Artificial Intelligence* (<https://web.archive.org/web/20091229084457/http://dli.iiit.ac.in/ijcai/IJCAI-73/PDF/027B.pdf>) (PDF), IJCAI, archived from the original (<http://dli.iiit.ac.in/ijcai/IJCAI-73/PDF/027B.pdf>) (PDF) on 29 December 2009
- Hobbes, Thomas (1651), *Leviathan*.
- Hofstadter, Douglas (1999) [1979], *Gödel, Escher, Bach: an Eternal Golden Braid*, Basic Books, ISBN 978-0-465-02656-2, OCLC 225590743 (<https://www.worldcat.org/oclc/225590743>).
- Howe, J. (November 1994), *Artificial Intelligence at Edinburgh University: a Perspective* (<http://www.inf.ed.ac.uk/about/AIhistory.html>), retrieved 30 August 2007.
- Kaplan, Andreas; Haenlein, Michael (2018), "Siri, Siri in my Hand, who's the Fairest in the Land? On the Interpretations, Illustrations and Implications of Artificial Intelligence", *Business Horizons*, **62**: 15–25, doi:10.1016/j.bushor.2018.08.004 (<https://doi.org/10.1016%2Fj.bushor.2018.08.004>).
- Kolata, G. (1982), "How can computers get common sense?", *Science*, **217** (4566): 1237–1238, Bibcode:1982Sci...217.1237K (<https://ui.adsabs.harvard.edu/abs/1982Sci...217.1237K>), doi:10.1126/science.217.4566.1237 (<https://doi.org/10.1126%2Fscience.217.4566.1237>), PMID 17837639 (<https://pubmed.ncbi.nlm.nih.gov/17837639>).
- Kurzweil, Ray (2005), *The Singularity is Near*, Viking Press, ISBN 978-0-14-303788-0, OCLC 71826177 (<https://www.worldcat.org/oclc/71826177>).
- Lakoff, George (1987), *Women, Fire, and Dangerous Things: What Categories Reveal About the Mind* (https://archive.org/details/womenfiredangero00lako_0), University of Chicago Press., ISBN 978-0-226-46804-4.
- Lenat, Douglas; Guha, R. V. (1989), *Building Large Knowledge-Based Systems*, Addison-Wesley, ISBN 978-0-201-51752-1, OCLC 19981533 (<https://www.worldcat.org/oclc/19981533>).
- Levitt, Gerald M. (2000), *The Turk, Chess Automaton*, Jefferson, N.C.: McFarland, ISBN 978-0-7864-0778-1.
- Lighthill, Professor Sir James (1973), "Artificial Intelligence: A General Survey", *Artificial Intelligence: a paper symposium*, Science Research Council
- Lucas, John (1961), "Minds, Machines and Gödel" (<http://users.ox.ac.uk/~jrlucas/Godel/mm-g.html>), *Philosophy*, **36** (XXXVI): 112–127, doi:10.1017/S0031819100057983 (<https://doi.org/10.1017%2FS0031819100057983>), retrieved 15 October 2008

- Maker, Meg Houston (2006), *AI@50: AI Past, Present, Future* (https://web.archive.org/web/20081008120238/http://www.engagingexperience.com/2006/07/ai50_ai_past_pr.html), Dartmouth College, archived from the original (http://www.engagingexperience.com/2006/07/ai50_ai_past_pr.html) on 8 October 2008, retrieved 16 October 2008
- Markoff, John (14 October 2005), "Behind Artificial Intelligence, a Squadron of Bright Real People" (https://www.nytimes.com/2005/10/14/technology/14artificial.html?_r=1&ei=5070&en=11ab55edb7cead5e&ex=1185940800&adxnnl=1&adxnnlx=1185805173-o7WsfW7qaPOx5/NUs1cQCQ&oref=slogin), *The New York Times*, retrieved 16 October 2008
- McCarthy, John; Minsky, Marvin; Rochester, Nathan; Shannon, Claude (31 August 1955), *A Proposal for the Dartmouth Summer Research Project on Artificial Intelligence* (<https://web.archive.org/web/20080930164306/http://www-formal.stanford.edu/jmc/history/dartmouth/dartmouth.html>), archived from the original (<http://www-formal.stanford.edu/jmc/history/dartmouth/dartmouth.html>) on 30 September 2008, retrieved 16 October 2008
- McCarthy, John; Hayes, P. J. (1969), "Some philosophical problems from the standpoint of artificial intelligence", in Meltzer, B. J.; Mitchie, Donald (eds.), *Machine Intelligence 4* (<http://www-formal.stanford.edu/jmc/mcchay69/mcchay69.html>), Edinburgh University Press, pp. 463–502, retrieved 16 October 2008
- McCorduck, Pamela (2004), *Machines Who Think* (2nd ed.), Natick, MA: A. K. Peters, Ltd., ISBN 978-1-56881-205-2, OCLC 52197627 (<https://www.worldcat.org/oclc/52197627>).
- McCullough, W. S.; Pitts, W. (1943), "A logical calculus of the ideas immanent in nervous activity", *Bulletin of Mathematical Biophysics*, 5 (4): 115–127, doi:10.1007/BF02478259 (<https://doi.org/10.1007%2FBF02478259>)
- Menabrea, Luigi Federico; Lovelace, Ada (1843), "Sketch of the Analytical Engine Invented by Charles Babbage" (<http://www.fourmilab.ch/babbage/sketch.html>), *Scientific Memoirs*, 3, retrieved 29 August 2008 With notes upon the Memoir by the Translator
- Minsky, Marvin (1967), *Computation: Finite and Infinite Machines*, Englewood Cliffs, N.J.: Prentice-Hall
- Minsky, Marvin; Papert, Seymour (1969), *Perceptrons: An Introduction to Computational Geometry* (<https://archive.org/details/perceptronsintro00mins>), The MIT Press, ISBN 978-0-262-63111-2, OCLC 16924756 (<https://www.worldcat.org/oclc/16924756>)
- Minsky, Marvin (1974), *A Framework for Representing Knowledge* (<http://web.media.mit.edu/~minsky/papers/Frames/frames.html>), retrieved 16 October 2008
- Minsky, Marvin (1986), *The Society of Mind*, Simon and Schuster, ISBN 978-0-671-65713-0, OCLC 223353010 (<https://www.worldcat.org/oclc/223353010>)
- Minsky, Marvin (2001), *It's 2001. Where Is HAL?* (http://www.ddj.com/hpc-high-performance-computing/197700454?cid=RSSfeed_DDJ_AI), Dr. Dobb's Technetcast, retrieved 8 August 2009
- Moravec, Hans (1976), *The Role of Raw Power in Intelligence* (<https://web.archive.org/web/20160303232511/http://www.frc.ri.cmu.edu/users/hpm/project.archive/general.articles/1975/Raw.Power.html>), archived from the original (<http://www.frc.ri.cmu.edu/users/hpm/project.archive/general.articles/1975/Raw.Power.html>) on 3 March 2016, retrieved 16 October 2008
- Moravec, Hans (1988), *Mind Children* (<https://archive.org/details/mindchildren00hans>), Harvard University Press, ISBN 978-0-674-57618-6, OCLC 245755104 (<https://www.worldcat.org/oclc/245755104>)
- NRC (1999), "Developments in Artificial Intelligence" (<https://archive.org/details/fundingrevolution00nati>), *Funding a Revolution: Government Support for Computing Research*, National Academy Press, ISBN 978-0-309-06278-7, OCLC 246584055 (<https://www.worldcat.org/oclc/246584055>)
- Newell, Allen; Simon, H. A. (1963), "GPS: A Program that Simulates Human Thought", in Feigenbaum, E.A.; Feldman, J. (eds.), *Computers and Thought*, New York: McGraw-Hill,

ISBN 978-0-262-56092-4, OCLC 246968117 (<https://www.worldcat.org/oclc/246968117>)

- Newquist, HP (1994), *The Brain Makers: Genius, Ego, And Greed in the Quest For Machines That Think*, New York: Macmillan/SAMS, ISBN 978-0-9885937-1-8
- Nick, Martin (2005), *Al Jazari: The Ingenious 13th Century Muslim Mechanic* (<http://www.alshindagah.com/marapr2005/jaziri.html>), Al Shindagah, retrieved 16 October 2008.*
- O'Connor, Kathleen Malone (1994), *The alchemical creation of life (takwin) and other concepts of Genesis in medieval Islam* (<http://repository.upenn.edu/dissertations/AAI9503804>), University of Pennsylvania, pp. 1–435, retrieved 10 January 2007
- Olsen, Stefanie (10 May 2004), *Newsmaker: Google's man behind the curtain* (http://news.cnet.com/Googles-man-behind-the-curtain/2008-1024_3-5208228.html), CNET, retrieved 17 October 2008.
- Olsen, Stefanie (18 August 2006), *Spying an intelligent search engine* (http://news.cnet.com/Spying-an-intelligent-search-engine/2100-1032_3-6107048.html), CNET, retrieved 17 October 2008.
- Pearl, J. (1988), *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*, San Mateo, California: Morgan Kaufmann, ISBN 978-1-55860-479-7, OCLC 249625842 (<https://www.worldcat.org/oclc/249625842>).
- Russell, Stuart J.; Norvig, Peter (2003), *Artificial Intelligence: A Modern Approach* (<http://aima.cs.berkeley.edu/>) (2nd ed.), Upper Saddle River, New Jersey: Prentice Hall, ISBN 0-13-790395-2.
- Poole, David; Mackworth, Alan; Goebel, Randy (1998), *Computational Intelligence: A Logical Approach* (<https://archive.org/details/computationalint00pool>), Oxford University Press., ISBN 978-0-19-510270-3.
- Samuel, Arthur L. (July 1959), "Some studies in machine learning using the game of checkers" (<http://domino.research.ibm.com/tchjr/journalindex.nsf/600cc5649e2871db852568150060213c/39a870213169f45685256bfa00683d74?OpenDocument>), *IBM Journal of Research and Development*, 3 (3): 210–219, CiteSeerX 10.1.1.368.2254 (<https://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.368.2254>), doi:10.1147/rd.33.0210 (<https://doi.org/10.1147%2Frd.33.0210>), retrieved 20 August 2007.
- Searle, John (1980), "Minds, Brains and Programs" (<https://web.archive.org/web/20071210043312/http://members.aol.com/NeoNoetics/MindsBrainsPrograms.html>), *Behavioral and Brain Sciences*, 3 (3): 417–457, doi:10.1017/S0140525X00005756 (<https://doi.org/10.1017%2FS0140525X00005756>), archived from the original (<http://members.aol.com/NeoNoetics/MindsBrainsPrograms.html>) on 10 December 2007, retrieved 13 May 2009.
- Simon, H. A.; Newell, Allen (1958), "Heuristic Problem Solving: The Next Advance in Operations Research", *Operations Research*, 6: 1–10, doi:10.1287/opre.6.1.1 (<https://doi.org/10.1287%2Fopre.6.1.1>).
- Simon, H. A. (1965), *The Shape of Automation for Men and Management*, New York: Harper & Row.
- Skillings, Jonathan (2006), *Newsmaker: Getting machines to think like us* (http://news.cnet.com/Getting-machines-to-think-like-us---page-2/2008-11394_3-6090207-2.html?tag=st.next), CNET, retrieved 8 October 2008.
- Tascarella, Patty (14 August 2006), "Robotics firms find fundraising struggle, with venture capital shy" (<http://www.bizjournals.com/pittsburgh/stories/2006/08/14/focus3.html?b=1155528000%5E1329573>), *Pittsburgh Business Times*, retrieved 15 March 2016.
- Turing, Alan (1936–37), "On Computable Numbers, with an Application to the Entscheidungsproblem" (<http://www.abelard.org/turpap2/tp2-ie.asp>), *Proceedings of the London Mathematical Society*, 2, s2-42 (42): 230–265, doi:10.1112/plms/s2-42.1.230 (<https://doi.org/10.1112%2Fplms%2Fs2-42.1.230>), retrieved 8 October 2008.

- Turing, Alan (October 1950), "Computing Machinery and Intelligence", *Mind*, **LIX** (236): 433–460, doi:10.1093/mind/LIX.236.433 (<https://doi.org/10.1093%2Fmind%2FLIX.236.433>), ISSN 0026-4423 (<https://www.worldcat.org/issn/0026-4423>).
- Weizenbaum, Joseph (1976), *Computer Power and Human Reason*, W.H. Freeman & Company, ISBN 978-0-14-022535-8, OCLC 10952283 (<https://www.worldcat.org/oclc/10952283>).

Retrieved from "https://en.wikipedia.org/w/index.php?title=History_of_artificial_intelligence&oldid=1101248003"

This page was last edited on 30 July 2022, at 01:38 (UTC).

Text is available under the Creative Commons Attribution-ShareAlike License 3.0; additional terms may apply. By using this site, you agree to the Terms of Use and Privacy Policy. Wikipedia® is a registered trademark of the Wikimedia Foundation, Inc., a non-profit organization.