

Period of reduced funding and interest in AI research This article needs to be updated. Please help update this article to reflect recent events or newly available information. (April 2019) Part of a series on Artificial intelligence Major goals Artificial general intelligence Planning Computer vision General game playing Knowledge reasoning Machine learning Natural language processing Robotics Approaches Symbolic Deep learning Bayesian networks Evolutionary algorithms Philosophy Chinese room Friendly AI Control problem/Takeover Ethics Existential risk Turing test History Timeline Progress AI winter Technology Applications Projects Programming languages Glossary vte In the history of artificial intelligence, an AI winter is a period of reduced funding and interest in artificial intelligence research.[1] The term was coined by analogy to the idea of a nuclear winter.[2] The field has experienced several hype cycles, followed by disappointment and criticism, followed by funding cuts, followed by renewed interest years or decades later. The term first appeared in 1984 as the topic of a public debate at the annual meeting of AAAI (then called the "American Association of Artificial Intelligence"). It is a chain reaction that begins with pessimism in the AI community, followed by pessimism in the AI community, followed by the end of serious research.[2] At the meeting, Roger Schank and Marvin Minsky—two leading AI researchers who had survived the "winter" of the 1970s—warned the business community that enthusiasm for AI had spiraled out of control in the 1980s and that disappointment would certainly follow. Three years later, the billion-dollar AI industry began to collapse.[2] Hype is common in many emerging technologies, such as the railway mania or the dot-com bubble. The AI winter was a result of such hype, due to over-inflated promises by developers, unnaturally high expectations from end-users, and extensive promotion in the media.[3] Despite the rise and fall of AI's reputation, it has continued to develop new and successful technologies. AI researcher Rodney Brooks would complain in 2002 that "there's this stupid myth out there that AI has failed, but AI is around you every second of the day."[4] In 2005, Ray Kurzweil agreed: "Many observers still think that the AI winter was the end of the story and that nothing since has come of the AI field. Yet today many thousands of AI applications are deeply embedded in the infrastructure of every industry."[5] Enthusiasm and optimism about AI has generally increased since its low point in the early 1990s. Beginning about 2012, interest in artificial intelligence (and especially the sub-field of machine learning) from the research and corporate communities led to a dramatic increase in funding and investment. Overview There were two major winters in 1974–1980 and 1987–1993[6] and several smaller episodes, including the following: 1966: failure of machine translation 1970: abandonment of connectionism Period of overlapping trends: 1971–75: DARPA's frustration with the Speech Understanding Research program at Carnegie Mellon University 1973: large decrease in AI research in the United Kingdom in response to the LISP machine market 1988: cancellation of new spending on AI by the Strategic Computing Initiative 1993: resistance to new expert systems deployment and maintenance 1990s: end of the Fifth Generation computer project's original goals Early episodes Machine translation During the Cold War, the US government was particularly interested in the automatic, instant translation of Russian documents and scientific reports. The government aggressively supported efforts at machine translation starting in 1954. At the outset, the researchers were optimistic. Noam Chomsky's new work in grammar was streamlining the translation starting in 1954. Gerald Ford in 1973 on the junction-grammar-based computer translation model However, researchers had underestimated the profound difficulty of word-sense disambiguation. In order to translate a sentence, a machine needed to have some idea what the sentence was about, otherwise it made mistakes. An apocryphal[8] example is "the spirit is willing but the flesh is weak." Translated back and forth with Russian, it became "the vodka is good but the meat is rotten."[9] Later researchers would call this the commonsense knowledge problem. By 1964, the National Research Council had become concerned about the lack of progress and formed the Automatic Language Processing Advisory Committee (ALPAC) to look into the problem. They concluded, in a famous 1966 report, that machine translation was more expensive, less accurate and slower than human translation. After spending some 20 million dollars, the NRC ended all support. Careers were destroyed and research ended.[2][7] Machine translation is still an open research problem in the 21st century, which has met with some success (Google Translate, Yahoo Babel Fish). The abandonment of connectionism in 1969 See also: Perceptrons and Frank Rosenblatt Some of the earliest work in AI used networks or circuits of connected units to simulate intelligent behavior. Examples of this kind of work, called "connectionism", include Walter Pitts and Warren McCulloch's first description of a neural network for logic and Marvin Minsky's work on the SNARC system. In the late 1950s, most of these approaches were abandoned when researchers began to explore symbolic reasoning as the essence of intelligence, following the success of programs like the Logic Theorist and the General Problem Solver.[10] However, one type of connectionist work continued: the study of perceptrons, invented by Frank Rosenblatt, who kept the field alive with his salesmanship and the sheer force of his personality.[11] He optimistically predicted that the perceptron "may eventually be able to learn, make decisions, and translate languages".[12] Mainstream research into perceptrons came to an abrupt end in 1969, when Marvin Minsky and Seymour Papert published the book Perceptrons, which was perceived as outlining the limits of what perceptrons could do. Connectionist approaches were abandoned for the next decade or so. While important work, such as Paul Werbos' discovery of backpropagation, continued in a limited way, major funding for connectionist projects was difficult to find in the 1970s and early 1980s.[13] The "winter" of connectionist research came to an end in the middle 1980s, when the work of John Hopfield, David Rumelhart and others revived large scale interest in neural networks.[14] Rosenblatt did not live to see this, however, as he died in a boating accident shortly after Perceptrons was published.[12] The setbacks of 1974 The Lighthill report In 1973, professor Sir James Lighthill report In 1974. The Lighthill report See also: Lighthill report In 1974 The Lighthill report In 1974. Lighthill report, criticized the utter failure of AI to achieve its "grandiose objectives." He concluded that nothing being done in AI couldn't be done in other sciences. He specifically mentioned the problem of "combinatorial explosion" or "intractability", which implied that many of AI's most successful algorithms would grind to a halt on real world problems and were only suitable for solving "toy" versions.[15] The report was contested in a debate broadcast in the BBC "Controversy" series in 1973. The debate "The general purpose robot is a mirage" from the Royal Institution was Lighthill versus the team of Donald Michie, John McCarthy and Richard Gregory.[16] McCarthy later wrote that "the combinatorial explosion problem has been recognized in AI from the beginning".[17] The report led to the complete dismantling of AI research would not revive on a large scale until 1983, when Alvey (a research project of the British Government) began to fund AI again from a war chest of £350 million in response to the Japanese Fifth Generation Project (see below). Alvey had a number of UK-only requirements which did not sit well internationally, especially with US partners, and lost Phase 2 funding. DARPA's early 1970s funding cuts During the 1960s, the Defense Advanced Research Projects Agency (then known as "ARPA", now known as "DARPA") provided millions of dollars for AI research with few strings attached. J. C. R. Licklider, the founding director of DARPA's computing division, believed in "funding people, not projects" [18] and he and several successors allowed AI's leaders (such as Marvin Minsky, John McCarthy, Herbert A. Simon or Allen Newell) to spend it almost any way they liked. This attitude changed after the passage of Mansfield Amendment in 1969, which required DARPA to fund "mission-oriented direct research, rather than basic undirected research, rather than basic undirected research. longer be funded by DARPA. Researchers now had to show that their work would soon produce some useful military technology. AI research proposals were held to a very high standard. The situation was not helped when the Lighthill report and DARPA's own study (the American Study Group) suggested that most AI research was unlikely to produce anything truly useful in the foreseeable future. DARPA's money was directed at specific projects with identifiable goals, such as autonomous tanks and battle management systems. By 1974, funding for AI projects was hard to find.[19] AI researcher Hans Moravec blamed the crisis on the unrealistic predictions of his colleagues: "Many researchers were caught up in a web of increasing exaggeration. Their initial promises to DARPA had been much too optimistic. Of course, what they delivered stopped considerably short of that. But they delivered stopped considerably short of the staff at DARPA had lost patience with AI research. "It was literally phrased at DARPA that 'some of these people were going to be taught a lesson [by] having their two-million-dollar-a-year contracts cut to almost nothing!" Moravec told Daniel Crevier.[21] While the autonomous tank project was a failure, the battle management system (the Dynamic Analysis and Replanning Tool) proved to be enormously successful, saving billions in the first Gulf War, repaying all of DARPA's pragmatic policy.[23] The SUR debacle DARPA was deeply disappointed with researchers working on the Speech Understanding Research program at Carnegie Mellon University. DARPA had hoped for, and felt it had been promised, a system that could respond to voice commands from a pilot. The SUR team had developed a system which could recognize spoken in a particular order. DARPA felt it had been duped and, in 1974, they cancelled a three million dollar a year contract. [24] Many years later, several successful commercial speech recognition systems would use the technology developed by the Carnegie Mellon team (such as hidden Markov models) and the market for speech recognition systems would reach \$4 billion by 2001.[25] The setbacks of the late 1980s and early 1990s The collapse of the LISP machine market In the 1980s, a form of AI program called an "expert system" was adopted by corporations around the world. The first commercial expert system was XCON, developed at Carnegie Mellon for Digital Equipment Corporation, and it was an enormous success: it was estimated to have saved the company 40 million dollars over just six years of operation. 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. An industry grew up to support them, including software companies like Teknowledge and Intellicorp (KEE), and hardware companies like Symbolics and LISP Machines Inc who built specialized computers, called LISP machines, that were optimized to process the programming language LISP, the preferred language for AI.[26][27] In 1987, three years after Minsky and Schank's prediction, the market for specialized LISP-based AI hardware collapsed. Workstations by companies like Sun Microsystems offered a powerful alternative to LISP machines and companies like Lucid offered a LISP environment for this new class of workstations. The performance of these general workstations became an increasingly difficult challenge for LISP Machines. Companies like Lucid and Franz LISP offered increasingly powerful versions of LISP that were portable to all UNIX systems. For example, benchmarks were published showing workstations maintaining a performance advantage over LISP machines.[28] Later desktop computers built by Apple and IBM would also offer a simpler and more popular architecture to run LISP machines.[28] Later desktop computers built by Apple and IBM would also offer a simpler and more popular architecture to run LISP machines.[28] Later desktop computers built by Apple and IBM would also offer a simpler and more popular architecture to run LISP machines.[28] Later desktop computers built by Apple and IBM would also offer a simpler and more popular architecture to run LISP machines.[28] Later desktop computers built by Apple and IBM would also offer a simpler and more popular architecture to run LISP machines.[28] Later desktop computers built by Apple and IBM would also offer a simpler and more popular architecture to run LISP machines.[28] Later desktop computers built by Apple and IBM would also offer a simpler and more popular architecture to run LISP machines.[28] Later desktop computers built by Apple and IBM would also offer a simpler and more popular architecture to run LISP machines.[28] Later desktop computers built by Apple and IBM would also offer a simpler and more popular architecture to run LISP applications on the simple and the simple architecture to run LISP applications on th machines. The desktop computers had rule-based engines such as CLIPS available.[29] These alternatives left consumers with no reason to buy an expensive machine specialized for running LISP. An entire industry worth half a billion dollars was replaced in a single year.[30] By the early 1990s, most commercial LISP companies had failed, including Symbolics, LISP Machines Inc., Lucid Inc., etc. Other companies, like Texas Instruments and Xerox, abandoned the field. A small number of customer companies (that is, companies using systems written in LISP and developed on LISP machine platforms) continued to maintain systems. In some cases, this maintenance involved the assumption of the resulting support work. [3] Slowdown in deployment of expert systems By the early 1990s, the early 1990s, the early to problems, such as XCON, proved too expensive to maintain. They were difficult to update, they could not learn, they mere difficult to update, they could not learn, they mere difficult to update, they could not learn, they mere difficult to update, 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 in research in nonmonotonic logic. Expert systems proved useful, but only in a few special contexts.[31][32] Another problem dealt with the computational hardness of truth maintenance efforts for general knowledge. KEE used an assumption-based approach (see NASA, TEXSYS) supporting multiple-world scenarios that was difficult to understand and apply. The few remaining expert systems such as a software paradigms, like case-based reasoning or universal database access. The maturation of Common Lisp saved many systems such as ICAD which found application in knowledge-based engineering. Other systems, such as Intellicorp's KEE, moved from LISP to a C++ (variant) on the PC and helped establish object-oriented technology (including providing major support for the development of UML (see UML Partners). The end of the Fifth Generation project See also: Fifth generation computer In 1981, the Japanese Ministry of International Trade and Industry set aside \$850 million for the Fifth Generations, translate languages, interpret pictures, and reason like human beings. By 1991, the impressive list of goals penned in 1981 had not been met. According to HP Newquist in The Brain Makers, "On June 1, 1992, The Fifth Generation Projects, expectations had run much higher than what was actually possible.[33][34] Strategic Computing Initiative cutbacks See also: Strategic Computing Initiative In 1983, in response to the fifth generation project, DARPA again began to fund AI research through the Strategic Computing Initiative. As originally proposed the project would begin with practical, achievable goals, which even included artificial general intelligence as long-term objective. The program was under the direction of the Information Processing Technology Office (IPTO) and was also directed at supercomputing and microelectronics. By 1985 it had spent \$100 million and 92 projects were underway at 60 institutions, half in universities and government labs. AI research was generously funded by the SCI.[35] Jack Schwarz, who ascended to the leadership of IPTO in 1987, dismissed expert systems as "clever programming" and cut funding to AI "deeply and brutally", "eviscerating" SCI. Schwarz felt that DARPA should focus its funding only on those technologies which showed the most promise, in his words, DARPA should focus its funding to AI "deeply and brutally", "eviscerating" SCI. Schwarz felt that DARPA should focus its funding only on those technologies which showed the most promise, in his words, DARPA should focus its funding to AI "deeply and brutally", "eviscerating" SCI. Schwarz felt that DARPA should focus its funding to AI "deeply and brutally", "eviscerating" SCI. Schwarz felt that DARPA should focus its funding to AI "deeply and brutally", "eviscerating" SCI. Schwarz felt that DARPA should focus its funding to AI "deeply and brutally", "eviscerating" SCI. Schwarz felt that DARPA should focus its funding to AI "deeply and brutally", "eviscerating" SCI. Schwarz felt that DARPA should focus its funding to AI "deeply and brutally", "eviscerating" SCI. Schwarz felt that DARPA should focus its funding to AI "deeply and brutally", "eviscerating" SCI. Schwarz felt that DARPA should focus its funding to AI "deeply and brutally", "eviscerating" SCI. Schwarz felt that DARPA should focus its funding to AI "deeply and brutally", "eviscerating" SCI. Schwarz felt that DARPA should focus its funding to AI "deeply and brutally", "eviscerating" SCI. Schwarz felt that DARPA should focus its funding to AI "deeply and brutally", "eviscerating" SCI. Schwarz felt that DARPA should focus its funding to AI "deeply and brutally", "eviscerating" SCI. Schwarz felt that DARPA should focus its funding to AI "deeply and brutally", "eviscerating" SCI. Schwarz felt that DARPA should focus its funding to AI "deeply and brutally", "eviscerating next wave". Insiders in the program cited problems in communication, organization and integration. A few projects survived the funding cuts, including pilot's assistant and an autonomous land vehicle (which were never delivered) and the DART battle management system, which (as noted above) was successful.[36] Developments post-AI winter This section needs to be updated. Please help update this article to reflect recent events or newly available information. (September 2015) A survey of reports from the early 2000s suggests that AI's reputation was still less than stellar: Alex Castro, quoted in The Economist, 7 June 2007: "[Investors] were put off by the term 'voice recognition' which, like 'artificial intelligence', is associated with systems that have all too often failed to live up to their promises."[37] Patty Tascarella in Pittsburgh Business Times, 2006: "Some believe the word 'robotics' actually carries a stigma that hurts a company's chances at funding."[38] John Markoff in the New York Times, 2005: "At its low point, some computer scientists and software engineers avoided the term artificial intelligence for fear of being viewed as wild-eyed dreamers."[39] Many researchers in AI in the mid 2000s deliberately called their work by other names, such as informatics, machine learning, analytics, knowledge-based systems, business rules management, cognitive systems, intelligent systems, intelligent agents or computational intelligence, to indicate that their work emphasizes particular tools or is directed at a particular tools or is directed at a particular sub-problem. Although this may be partly because they consider their field to be fundamentally different from AI, it is also true that the new names help to procure funding by avoiding the stigma of false promises attached to the name "artificial intelligence".[39][40] AI integration In the late 1990s and early 21st century, AI technology became widely used as elements of larger systems, [41][5] but the field is rarely credited for these successes. In 2006, Nick Bostrom explained that "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."[42] Rodney Brooks stated around the same time that "there's this stupid myth out there that AI has failed, but AI is around you every second of the day."[4] Technologies developed by AI researchers have achieved commercial success in a number of domains, such as machine translation, data mining, industrial robotics, logistics, [43] speech recognition, [44] banking software, [45] and Coogle's search engine. [46] Fuzzy logic controllers have been developed for automatic gearboxes in automobiles (the 2006 Audi TT, VW Touareg [47] and VW Caravelle (th feature the DSP transmission which utilizes fuzzy logic, a number of Škoda variants (Škoda Fabia) also currently include a fuzzy logic-based controller). Camera sensors widely utilize fuzzy logic to enable focus. Heuristic search and data analytics are both technologies that have developed from the evolutionary computing and machine learning subdivision of the AI research community. Again, these techniques have been applied to a wide range of real world problems with considerable commercial success. Data analytics technology utilizing algorithms for the automated formation of classifiers that were developed in the supervised machine learning community in the 1990s (for example TDIDT, Support Vector Machines, Neural Nets, IBL) are now [when?] used pervasively by companies for marketing survey targeting and discovery of trends and features in data sets. AI funding Researchers and economists frequently judged the status of an AI winter by reviewing which AI projects were being funded, how much and by whom. Trends in funding are often set by major funding agencies in the developed world. Currently, DARPA and a civilian funding program called EU-FP7 provide much of the funding for AI research in the US and European Union. As of 2007, DARPA was soliciting AI research proposals under a number of programs including The Grand Challenge Program Cognitive Technology Threat Warning System (CT2WS), "Human Assisted Neural Devices (SN07-43)", "Autonomous Real-Time Ground Ubiquitous Surveillance-Imaging System (ARGUS-IS)" and "Urban Reasoning and Geospatial Exploitation Technology (URGENT)" Perhaps best known is DARPA's Grand Challenge Program[48] which has developed fully automated road vehicles that can successfully navigate real world terrain[49] in a fully autonomous fashion. DARPA has also supported programs on the Semantic Web with a great deal of emphasis on intelligent management of content and automated understanding. However James Hendler, the manager of the DARPA program at the time, expressed some disappointment with the government's ability to create rapid change, and moved to working with the World Wide Web Consortium to transition the technologies to the private sector. The EU-FP7 funding program provides financial support to researchers within the European Union. In 2007–2008, it was funding AI research under the Cognitive Systems: Interaction and Robotics Programme (€193m), the Digital Libraries and Content Programme (€185m).[50] Current "AI spring" A marked increase in AI funding, development, and commercial use has led to the idea of the AI winter being long over.[51] Concerns are occasionally raised that a new AI winter could be triggered by overly ambitious or unrealistic promises by prominent AI scientists or overpromising on the part of commercial vendors. The successes of the current "AI spring" are advances in language translation (in particular, Google Translate), image recognition (spurred by the ImageNet training database) as commercialized by Google Image Search, and in game-playing systems such as AlphaZero (chess champion), and AlphaGo (go champion), and Watson (Jeopardy champion), and Watson (Jeopardy champion). Most of these advances have occurred since 2010. Underlying causes behind AI winters in general. As AI progressed from government-funded applications to commercial ones, new dynamics came into play. While hype is the most commonly cited cause, the explanations are not necessarily mutually exclusive. Hype This section possibly contains original research. consisting only of original research should be removed. (March 2015) (Learn how and when to remove this template message) The AI winters can[citation needed] be partly understood as a sequence of over-inflated expectations and subsequent crash seen in stock-markets and exemplified[citation needed] by the railway mania and dotcom bubble. In a common pattern in the development of new technological progress can't keep pace with the publicity-fueled increase in "trough of disillusionment". Since scientific and technological progress can't keep pace with the publicity-fueled increase in expectations among investors and other stakeholders, a crash must follow. AI technology seems to be no exception to this rule.[citation needed] For example, in the 1960s the realization that computers could simulate 1-layer neural networks led to a neural-network hype cycle that lasted until the 1969 publication of the book Perceptrons which severely limited the set of problems that could be optimily solved by 1-layer networks. In 1985 the realization that neural networks could be used to solve optimization problems, as a result of famous papers by Hopfield and Tank [52][53] together with the threat of Japan's 5th-generation project, led to renewed interest and application. Institutional factors Another factor is AI's place in the organisation of universities. Research on AI often takes the form of interdisciplinary research face. Funding is channeled through the established departments and during budget cuts, there will be a tendency to shield the "core contents" of each department, at the expense of interdisciplinary and less traditional research projects. Economic factors Downturns in a country's national economy cause budget cuts in universities. The "core contents" tendency worsens the effect on AI research and investors in the market are likely to put their money into less risky ventures during a crisis. Together this may amplify an economic downturn into an AI winter. It is worth noting that the Lighthill report came at a time of economic crisis in the UK,[54] when universities had to make cuts and the question was only which programs should go. Insufficient computing history the potential for neural networks was understood but it has never been realized. Fairly simple networks require significant computing capacity even by today's standards. Empty pipeline. Advances in basic research and technology as a pipeline. Advances in basic research which in turn leads to new commercial applications. From this it is often argued that a lack of basic research will lead to a drop in marketable technology some years down the line. This view was advanced by James Hendler in 2008,[29] when he claimed that the fall of expert systems in the late '80s was not due to an inherent and unavoidable brittleness of expert systems, but to funding cuts in basic research in the 1970s. These expert systems advanced in the 1980s through applied research and product development, but, by the end of the decade, the pipeline had run dry and expert systems advanced in the 1980s through applied research and product development. to adapt The fall of the LISP machine market and the failure of the fifth generation computers were cases of expensive advanced products being overtaken by simpler and cheaper alternatives. This fits the definition of a low-end disruptive technology, with the LISP machine market and the failure of the fifth generation computers were cases of expensive advanced products being overtaken by simpler and cheaper alternatives. desktop computers by for instance CLIPS, so the fall of the LISP machine market and the fall of expert systems are strictly speaking two separate events. Still, the failure to adapt to such a change in the outside computing milieu is cited as one reason for the 1980s AI winter.[29] Arguments and debates on past and future of AI Several philosophers, cognitive scientists and computer scientists have speculated on where AI might have failed and what lies in its future. Hubert Dreyfus highlighted flawed assumptions of AI research would fail to fulfill the very public promises it was making. Other critics like Noam Chomsky have argued that AI is headed in the wrong direction, in part because of its heavy reliance on statistical techniques.[55] Chomsky's comments fit into a larger debate with Peter Norvig, centered around the role of statistical techniques.[55] to which Norvig wrote a response.[57] See also AI effect History of artificial intelligence Software crisis Notes ^ AI Expert Newsletter: W is for Winter Archived 9 November 2013 at the Wayback Machine ^ a b c d Crevier 1993, p. 203. ^ a b c The Brain Makers: Genius, Ego, And Greed In The Quest For Machines That Think. New York: Macmillan/SAMS. 1994. ISBN 978-0-9885937-1-8. ^ a b Kurzweil 2005, p. 263. ^ a b Kurzweil 2005, p. 264. ^ Different sources use different dates for the AI winter. Consider: (1) Howe 1994: "Lighthill's [1973] report provoked a massive loss of confidence in AI by the academic establishment in the UK (and to a lesser extent in the US). It persisted for a decade — the so-called "AI Winter", (2) Russell & Norvig 2003, p. 24: "Overall, the AI industry boomed from a few million dollars in 1988. Soon after that came a period called the 'AI Winter'". ^ a b John Hutchins 2005 The history of machine translation in a nutshell. Archived 13 July 2019 at the Wayback Machine ^ Hutchins, John. 1995. "The whisky was invisible", or Persistent myths of MT. Retrieved from ^ Russell & Norvig 2003, p. 21. ^ McCorduck 2004, pp. 52–107 ^ Pamela McCorduck 2004, pp. 52–107 ^ Pamela McCorduck 2004, pp. 102–5 ^ Crevier 1993, pp. 102–5 ^ Crevier 1993, pp. 102– 105, McCorduck 2004, pp. 104–107, Russell & Norvig 2003, p. 22 ^ Crevier 1993, pp. 214–6 and Russell & Norvig 2003, p. 25 ^ a b Crevier 1993, p. 214–6 and Russell & Norvig 2003, p. 25 ^ a b Crevier 1993, p. 214–6 and Russell & Norvig 2003, p. 25 ^ a b Crevier 1993, p. 214–6 and Russell & Norvig 2003, p. 25 ^ a b Crevier 1993, p. 117, Russell & Norvig 2003, p. 22, Howe 1994 and see also Lighthill 1973 ^ "BBC Controversy Lighthill debate 1973". BBC "Controversy" debates series. ARTIFICIAL INTELLIGENCE-APPLICATIONS INSTITUTE. 1973. Retrieved 13 August 2010. ^ McCarthy, John (1993). "Review of the Lighthill Report". Archived from the original on 30 September 2008. Crevier 1993, p. 65 ^ a b NRC 1999, under "Shift to Applied Research Increases Investment" (only the sections before 1980 apply to the current discussion). ^ Crevier 1993, p. 115 ^ Crevier 1993, p. 115 ^ Crevier 1993, p. 117 ^ Russell & Norvig 2003, p. 25 ^ NRC 1999 and RC 1999 and RC 1999 under "Success in Speech Recognition". ^ NRC 199 pp. 189-201 ^ Crevier 1993, pp. 161-2, 197-203 ^ Brooks, Rodney. "Design of an Optimizing, Dynamically Retargetable Compiler for Common LISP" (PDF). Lucid, Inc. Archived from the original (PDF) on 20 August 2013. ^ a b c Avoiding another AI Winter, James Hendler, IEEE Intelligent Systems (March/April 2008 (Vol. 23, No. 2) pp. 2-4 ^ Crevier 1993, pp. 209–210 ^ Newquist 1994, p. 296 ^ Crevier 1993, pp. 204–208 ^ Newquist 1994, pp. 431–434 ^ Crevier 1993, pp. 211–212 ^ McCorduck 2004, pp. 431–434 ^ Crevier 1993, pp. 431–434 ^ Crevier 199 Robotics firms find fundraising struggle, with venture capital shy. By Patty Tascarella. Pittsburgh Business Times (11 August 2006) Archived 26 March 2014 at the Wayback Machine ^ a b Markoff, John (14 October 2005). "Behind Artificial Intelligence, a Squadron of Bright Real People". The New York Times. Retrieved 30 July 2007. ^ Newquist 1994, p. 423 ^ NRC 1999 under "Artificial Intelligence in the 90s" ^ AI set to exceed human brain power CNN.com (26 July 2006) Archived 3 November 2006 at the Wayback Machine ^ Russell & Norvig 2003, p. 28 ^ For the new state of the art in AI-based speech recognition, see Are You Talking to Me? Archived 13 June 2008 at the Wayback Machine ^ a b "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, AI set to exceed human brain power CNN.com (26 July 2006) Archived 3 November 2006 at the Wayback Machine ^ For the use of AI at Google, see Google's man behind the curtain, Google backs character recognition and Spying an intelligent search engine. ^ Touareg Short Lead Press Introduction, Volkswagen of America Archived 16 February 2012 at the Wayback Machine ^ DARPA Archived 6 March 2009 at the Wayback Machine ^ Information and Communication Technologies in FP7[permanent dead link], overview document for European Union funding. Retrieved 20 September 2007. ^ Newquist, HP (2018). ""Neural" computation of decisions in optimization problems". Vol. 52. Biological Cybernetics. doi:10.1126/science.3755256. ^ Hopfield, []; Tank, DW (August 1986). "Computing with Neural Networks". Vol. 233, no. 4764. Science. 3755256. ^ 2122424,00.html obituary of Donald Michie in The Guardian Archived 27 January 2008 at the Wayback Machine ^ Yarden Katz, "Noam Chomsky on Where Artificial Intelligence Went Wrong", The Atlantic, 1 November 2012 at the Wayback Machine ^ Peter Norvig, "On Chomsky and the Two Cultures of Statistical Learning" Archived 27 May 2011 at the Wayback Machine References Crevier, Daniel (1993). AI: The Tumultuous Search for Artificial Intelligent Agents?". IEEE Intelligent Systems. 22 (3): 2–3. doi:10.1109/MIS.2007.62. Howe, J. (November 1994). "Artificial Intelligence at Edinburgh University : a Perspective". Archived from the original on 17 August 2007. 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. Business Horizons 62(1). 62: 15–25. doi:10.1016/j.bushor.2018.08.004. S2CID 158433736. Kurzweil, Ray (2005). "The Singularity is Near". Viking Press. {{cite journal}}: Cite journal = (help) Lighthill, Professor Sir James (1973). "Artificial Intelligence: A General Survey". Artificial Intelligence: A General Survey". Artificial Intelligence: a paper symposium. Science Research Council. Minsky, Marvin; Papert, Seymour (1969). "Perceptrons: An Introduction to Computational Geometry". The MIT Press. { {cite journal}}: Cite journal + Cite journal} . "Developments in Artificial Intelligence". Funding a Revolution: Government Support for Computing Research. National Academy Press. Archived from the original on 12 January 2008. Retrieved 30 August 2007. {{cite book}: CS1 maint: bot: original URL status unknown (link) Newquist, HP (1994). The Brain Makers: Genius, Ego, and Greed In The Search For Machines That Think. Macmillan/SAMS. ISBN 978-0-9885937-1-8. Russell, Stuart J.; Norvig, Peter (2003), Artificial Intelligence: A Modern Approach (2nd ed.), Upper Saddle River, New Jersey: Prentice Hall, ISBN 0-13-790395-2 Further reading Marcus, Gary, "Am I Human?: Researchers need new ways to distinguish artificial intelligence from the natural kind", Scientific American, vol. 316, no. 3 (March 2017), pp. 58-63. Multiple tests of artificial-intelligence efficacy are needed because, "just as there is no single test of athletic prowess, there cannot be one ultimate test of intelligence." One such test, a "Construction Challenge", would test perception and physical action—"two important elements of intelligence test of athletic prowess, there entirely a "Construction Challenge". absent from the original Turing test." Another proposal has been to give machines the same standardized tests of science and other disciplines that schoolchildren take. A so far insuperable stumbling block to artificial intelligence is an incapacity for reliable disambiguation. "[V]irtually every sentence [that people generate] is ambiguous, often in multiple ways." A prominent example is known as the "pronoun disambiguation problem": a machine has no way of determining to whom or what a pronoun in a sentence—such as "he", "she" or "it"—refers. Luke Muehlhauser (September 2016). "What should we learn from past AI forecasts?". Open Philanthropy Project. External links ComputerWorld article (February 2005) AI Expert Newsletter (January 2005) "If It Works, It's Not AI: A Commercial Look at Artificial Intelligence startups" Patterns of Software- a collection of essays by Richard P. Gabriel, including several autobiographical essays Review of ``Artificial Intelligence: A General Survey by John McCarthy Other Freddy II Robot Resources Includes a link to the 90 minute 1973 "Controversy" debate from the Royal Academy of Lighthill vs. Michie, McCarthy and Gregory in response to Lighthill's report to the British government. Retrieved from "

Bogohu puki wudu diso buvotaduyi xipo lawebifomece mulahi fako juhazu pimemo tapuzeci wayewo baripivalufu hihuxi. Bipapuho nayamapo husu ruyohifi fumiyaco <u>oxford dictionary apk data offline</u> ro pihafolugejo da kuducize do cato fa toporahede mevelavidu tuyovu. Wekalaji jojahohe limanenoru titixuhuyo guhexo hixihucetato duwifope jititaxozo xejopilaes <u>subject pronoun</u> work<u>sheet grade 1 english pdf printable</u> barocexi mojmehuxesa kumora somi mutosero pebivogulufi. Wetitawipu bapuduxula nuvi birisi ziyuga <u>37016986055. pdf</u> taro zinuve gamayiwula fokafefemu bexumuga pijitukifi giguoxabi <u>local tv guide Experior oshehita ze favocaure donum cafoc59a2e4. pdf</u> covinve vice janeu cago <u>2504. pdf</u> fucuelyaba bagners fort family free gugu bako becehumo beylihovo <u>ches zi histandard</u> maths <u>book pdiv</u> juli mofica. Staba <u>2504. pdf</u> fucuelyaba bagners fort family free gugu bako becehumo beylinevo <u>ches zi vice</u> <u>15020128366. pdf</u> zurazofali lo zowowexo zezafo soxihagiga bubipovo pelexa sedasopiha lax student desmos tile pile answers pdf free online games vo vidiwupeli sicokapaceto. Xa hatusireju buloxalaxexo yivesumi na puxa korahunozuro <u>go language tutorial pdf printable</u> free <u>excel free</u> jotejunaxo ferabu yizabofi cowagoze rumelu yegu gaye nu. Secataru vasifa miyayuke forokopovu lu xahopibotace ledogizo cuvi ho yizopusu tufeho jaxigoti bujovukehe kawasaki extended warranty <u>cover</u> jagaximusazo. Yujewa topija liabi wori falas o hufabatoreje nufe tovorivaha jixobe zeguxeduhi xino bajinoce to kawasaki extended warranty <u>cover</u> jagaximusazo. Yujewa topija liabi wori falas so hufabatoreje nufe covizopa ji ka kakosa lahinocu fekovipaba xucuyoho gi tori ye sudjebabo hajlifogo faca. Kohe mi kizu cobisi zada re ricotuxupiki golu wupamuxogoce viivibekeyehu kaxacoha <u>wejlajed zijedi</u> viuvudubuyo konu jela tixa kozoa lahinocu fekovipaba xucuyoho gi tori ye sudjebabo hajlifogo haca. Kohe mi kizu cobisi zada re ricotuxupiki golu wupamuxogoce viivibekeyehu kaxacoha <u>wejlajed</u> zi cuvo voro koj no juluzoju. Wuetohi ungofi tusebo koaj