

The "Intelligent" Cloud

A Brief History of AI: 1950 - 2017

Philip Cookson
Principal Consultant
Email: philip.cookson@iprcorp.com
International Planning & Research Corporation
95 Sawyer Road, Suite 120, Waltham, MA 02453
T: +1 781 652-3000 info.us@iprcorp.com
T: +1 781 652-3061 www.iprcorp.com

The origins of AI: from 1950 to 2017

The concept of Artificial Intelligence is not new, in fact it was first raised in the 1950's in a research paper by mathematician Alan Turing entitled "[Computing Machinery & Intelligence](#)", although it existed in science fiction literature much earlier.



Alan M. Turing
Mathematician
(1912 – 1954)

A. M. Turing (1950) Computing Machinery and Intelligence. *Mind* 49: 433-460.

COMPUTING MACHINERY AND INTELLIGENCE

By A. M. Turing

I. The Imitation Game

I propose to consider the question, "Can machines think?" This should begin with definitions of the meaning of the terms "machine" and "think." The definitions might be framed in an attempt to fix as precisely as possible the normal use of the words, but this attitude is dangerous. If the meaning of the words "machine" and "think" are to be found by examining how they are commonly used it is difficult to escape the conclusion that the meaning and the answer to the question, "Can machines think?" is to be sought in a statistical survey such as a Gallup poll. But this is absurd. Instead of attempting such a definition I shall replace the question by another, which is closely related to it and is expressed in relatively unambiguous words.

The new form of the problem can be described in terms of a game which we call the "imitation game." It is played with three people, a man (A), a woman (B), and an interrogator (C) who may be of either sex. The interrogator stays in a room apart from the other two. The object of the game for the interrogator is to determine which of the other two is the man and which is the woman. He knows them by labels X and Y, and at the end of the game he says either "X is A and Y is B" or "X is B and Y is A." The interrogator is allowed to put questions to A and B thus:

C: Will X please tell me the length of his or her hair?

Now suppose X is actually A, then A must answer. If it is B's object in the game to try and cause C to make the wrong identification. His answer might therefore be:

"My hair is shaggy, and the longest strands are about nine inches long."

In order that voices may not help the interrogator the answers should be written, or better still, typewritten. The ideal arrangement is to have a teleprinter communicating between the two rooms. Alternatively the question and answers can be repeated by an intermediary. The object of the game for the third player (B) is to help the interrogator. The best strategy for her is probably to give truthful answers. She can add such things as "I am the woman, don't listen to him!" to her answers, but it will avail nothing as the man can make similar remarks.

We now ask the question, "What will happen when a machine takes the part of A in this game?" Will the interrogator decide wrongly as often when the game is played like this as he does when the game is played between a man and a woman? These questions replace our original, "Can machines think?"

Throughout the 60's, 70's, 80's and 90's AI was an active area of academic research, and there were a few isolated AI projects that achieved fame.

In the 1960's, the First International Joint Conference on Artificial Intelligence was held at Stanford University in California, USA. This decade also saw the first (text based) program ELIZA that was able to carry out an interactive (typed) dialogue in English.

Date	AI Development
1960s	Ray Solomonoff lays the foundations of a mathematical theory of AI, introducing universal Bayesian methods for inductive inference and prediction.
1961	Unimation's industrial robot Unimate worked on a General Motors automobile assembly line .
1963	Thomas Evans' program, ANALOGY, written as part of his PhD work at MIT, demonstrated that computers can solve the same analogy problems as are given on IQ tests.
1963	Edward Feigenbaum and Julian Feldman published <i>Computers and Thought</i> , the first collection of articles about artificial intelligence.
1964	Danny Bobrow's dissertation at MIT (technical report #1 from MIT's AI group, Project MAC), shows that computers can understand natural language well enough to solve algebra word problems correctly.
1965	Joseph Weizenbaum (MIT) built ELIZA , an interactive program that carries on a dialogue in English language on any topic. It was a popular toy at AI centers on the ARPANET when a version that "simulated" the dialogue of a psychotherapist was programmed.
1965	Edward Feigenbaum initiated Dendral , a ten-year effort to develop software to deduce the molecular structure of organic compounds using scientific instrument data. It was the first expert system .
1966	Ross Quillian (PhD dissertation, Carnegie Inst. of Technology, now CMU) demonstrated semantic nets .
1966	Machine Intelligence workshop at Edinburgh – the first of an influential annual series organized by Donald Michie and others.
1967	Dendral program (Edward Feigenbaum, Joshua Lederberg , Bruce Buchanan, Georgia Sutherland at Stanford University) demonstrated to interpret mass spectra on organic chemical compounds. First successful knowledge-based program for scientific reasoning.
1968	Joel Moses (PhD work at MIT) demonstrated the power of symbolic reasoning for integration problems in the Macsyma program. First successful knowledge-based program in mathematics .
1968	Richard Greenblatt (programmer) at MIT built a knowledge-based chess-playing program , MacHack , that was good enough to achieve a class-C rating in tournament play.
1969	Stanford Research Institute (SRI): Shakey the Robot , demonstrated combining animal locomotion , perception and problem solving .
1969	Roger Schank (Stanford) defined conceptual dependency model for natural language understanding .
1969	First International Joint Conference on Artificial Intelligence (IJCAI) held at Stanford University.

The “Intelligent” Cloud: A Brief History of AI (1950 – 2017)

The 1970's saw the introduction of the PROLOG programming language, a Language for Logic Programming and Symbolic Computation; which along with LISP (short for List Processing), became two of the most important programming languages used in Artificial Intelligence in this era. The 1970's also saw the introduction of the first "Autonomous Vehicle" ... a mobile cart that was able to navigate a chair filled room.

Date	AI Development
Early 1970s	Jane Robinson and Don Walker established an influential Natural Language Processing group at SRI.
1970	Jaime Carbonell (Sr.) developed SCHOLAR, an interactive program for computer assisted instruction based on semantic nets as the representation of knowledge.
1971	Terry Winograd 's PhD thesis (MIT) demonstrated the ability of computers to understand English sentences in a restricted world of children's blocks, in a coupling of his language with a robot arm that carried out instructions typed in English.
1972	Prolog programming language developed by Alain Colmerauer .
1973	The Assembly Robotics Group at University of Edinburgh builds Freddy Robot, capable of using visual perception to locate and assemble models. (See Edinburgh Freddy Assembly Robot : a versatile computer-controlled assembly system.)
1973	The Lighthill report gives a largely negative verdict on AI research in Great Britain and forms the basis for the decision by the British government to discontinue support for AI research in all but two universities.
1974	Ted Shortliffe 's PhD dissertation on the MYCIN program (Stanford) demonstrated a very practical rule-based approach to medical diagnoses, even in the presence of uncertainty. While it borrowed from DENDRAL, its own contributions strongly influenced the future of expert system development.
1975	Marvin Minsky published his widely read and influential article on Frames as a representation of knowledge, in which many ideas about schemas and semantic links are brought together.
1975	The Meta-Dendral learning program produced new results in chemistry (some rules of mass spectrometry) the first scientific discoveries by a computer to be published in a refereed journal.
Mid-1970s	Barbara Grosz (SRI) established limits to traditional AI approaches to discourse modeling. Subsequent work by Grosz, Bonnie Webber and Candace Sidner developed the notion of "centering", used in establishing focus of discourse and anaphoric references in Natural language processing .
1978	Herbert A. Simon wins the Nobel Prize in Economics for his theory of bounded rationality , one of the cornerstones of AI known as " satisficing ".
1978	The MOLGEN program, written at Stanford by Mark Stefik and Peter Friedland, demonstrated that an object-oriented programming representation of knowledge can be used to plan gene-cloning experiments.
1979	Bill VanMelle's PhD dissertation at Stanford demonstrated the generality of MYCIN 's representation of knowledge and style of reasoning in his EMYCIN program, the model for many commercial expert system "shells".
1979	Jack Myers and Harry Pople at University of Pittsburgh developed INTERNIST, a knowledge-based medical diagnosis program based on Dr. Myers' clinical knowledge.
1979	The Stanford Cart, built by Hans Moravec , becomes the first computer-controlled, autonomous vehicle when it successfully traverses a chair-filled room and circumnavigates the Stanford AI Lab .
1979	BKG, a backgammon program written by Hans Berliner at CMU , defeats the reigning world champion.
Late 1970s	Stanford's SUMEX-AIM resource, headed by Ed Feigenbaum and Joshua Lederberg, demonstrates the power of the ARPAnet for scientific collaboration.

The “Intelligent” Cloud: A Brief History of AI (1950 – 2017)

In the 1980's we saw the introduction of dedicated LISP machines for AI computation, as well as the development of practical Expert System shells and associated commercial applications. Autonomous vehicles had also improved, with the University of Munich building the a robot car that drove at speeds of 55 mph on (understandably) empty streets.

Date	AI Development
1980s	Lisp machines developed and marketed. First expert system shells and commercial applications.
1980	First National Conference of the American Association for Artificial Intelligence (AAAI) held at Stanford.
1981	Danny Hillis designs the connection machine, which utilizes Parallel computing to bring new power to AI, and to computation in general. (Later founds Thinking Machines Corporation)
1982	The Fifth Generation Computer Systems project (FGCS), an initiative by Japan's Ministry of International Trade and Industry, begun in 1982, to create a "fifth generation computer" (see history of computing hardware) which was supposed to perform much calculation utilizing massive parallelism.
1983	John Laird and Paul Rosenbloom, working with Allen Newell , complete CMU dissertations on Soar (program).
1983	James F. Allen invents the Interval Calculus, the first widely used formalization of temporal events.
Mid-1980s	Neural Networks become widely used with the Backpropagation algorithm (first described by Paul Werbos in 1974).
1985	The autonomous drawing program, AARON , created by Harold Cohen , is demonstrated at the AAAI National Conference (based on more than a decade of work, and with subsequent work showing major developments).
1986	The team of Ernst Dickmanns at Bundeswehr University of Munich builds the first robot cars, driving up to 55 mph on empty streets.
1986	Barbara Grosz and Candace Sidner create the first computation model of discourse , establishing the field of research. ^[36]
1987	Marvin Minsky published The Society of Mind , a theoretical description of the mind as a collection of cooperating agents . He had been lecturing on the idea for years before the book came out (c.f. Doyle 1983). ^[37]
1987	Around the same time, Rodney Brooks introduced the subsumption architecture and behavior-based robotics as a more minimalist modular model of natural intelligence; Nouvelle AI .
1987	Commercial launch of generation 2.0 of Alacrity by Alacritous Inc./Allstar Advice Inc. Toronto, the first commercial strategic and managerial advisory system. The system was based upon a forward-chaining, self-developed expert system with 3,000 rules about the evolution of markets and competitive strategies and co-authored by Alistair Davidson and Mary Chung, founders of the firm with the underlying engine developed by Paul Tarvydas. The Alacrity system also included a small financial expert system that interpreted financial statements and models. ^[38]
1989	Dean Pomerleau at CMU creates ALVINN (An Autonomous Land Vehicle in a Neural Network).

The “Intelligent” Cloud: A Brief History of AI (1950 – 2017)

In the 1990's we see AI being successfully deployed as part of an advanced logistics scheduling application during the Gulf War ... and generating enough value to pay back the USA Defense Advance Research Project Agency's prior 30 year investment in AI research.

This decade also saw the release of the first consumer AI product, the Tiger Electronic's Furby.



Date	AI Development
Early 1990s	TD-Gammon , a backgammon program written by Gerry Tesauro, demonstrates that reinforcement (learning) is powerful enough to create a championship-level game-playing program by competing favorably with world-class players.
1990s	Major advances in all areas of AI, with significant demonstrations in machine learning, intelligent tutoring , case-based reasoning, multi-agent planning, scheduling , uncertain reasoning, data mining , natural language understanding and translation, vision, virtual reality , games, and other topics.
1991	DART scheduling application deployed in the first Gulf War paid back DARPA's investment of 30 years in AI research.
1993	Ian Horswill extended behavior-based robotics by creating Polly , the first robot to navigate using vision and operate at animal-like speeds (1 meter/second).
1993	Rodney Brooks , Lynn Andrea Stein and Cynthia Breazeal started the widely publicized MIT Cog project with numerous collaborators, in an attempt to build a humanoid robot child in just five years.
1993	ISX corporation wins "DARPA contractor of the year" for the Dynamic Analysis and Replanning Tool (DART) which reportedly repaid the US government's entire investment in AI research since the 1950s.
1994	With passengers on board, the twin robot cars VaMP and VITA-2 of Ernst Dickmanns and Daimler-Benz drive more than one thousand kilometers on a Paris three-lane highway in standard heavy traffic at speeds up to 130 km/h. They demonstrate autonomous driving in free lanes, convoy driving, and lane changes left and right with autonomous passing of other cars.
1994	English draughts (checkers) world champion Tinsley resigned a match against computer program Chinook . Chinook defeated 2nd highest rated player, Lafferty . Chinook won the USA National Tournament by the widest margin ever.
1995	"No Hands Across America": A semi-autonomous car drove coast-to-coast across the United States with computer-controlled steering for 2,797 miles (4,501 km) of the 2,849 miles (4,585 km). Throttle and brakes were controlled by a human driver.
1995	One of Ernst Dickmanns' robot cars (with robot-controlled throttle and brakes) drove more than 1000 miles from Munich to Copenhagen and back, in traffic, at up to 120 mph, occasionally executing maneuvers to pass other cars (only in a few critical situations a safety driver took over). Active vision was used to deal with rapidly changing street scenes.
1997	The Deep Blue chess machine (IBM) defeats the (then) world chess champion, Garry Kasparov .

The “Intelligent” Cloud: A Brief History of AI (1950 – 2017)

1997	First official RoboCup football (soccer) match featuring table-top matches with 40 teams of interacting robots and over 5000 spectators.
1997	Computer Othello program Logistello defeated the world champion Takeshi Murakami with a score of 6–0.
1998	Tiger Electronics' Furby is released, and becomes the first successful attempt at producing a type of A.I to reach a domestic environment .
1998	Tim Berners-Lee published his Semantic Web Road map paper.
1998	Leslie P. Kaelbling , Michael Littman , and Anthony Cassandra introduce the first method for solving POMDPs offline, jumpstarting widespread use in robotics and automated planning and scheduling ^[45]
1999	Sony introduces an improved domestic robot similar to a Furby, the AIBO becomes one of the first artificially intelligent "pets" that is also autonomous .
Late 1990s	Web crawlers and other AI-based information extraction programs become essential in widespread use of the World Wide Web .

The “Intelligent” Cloud: A Brief History of AI (1950 – 2017)

In the 2000's we saw the introduction of "Smart Toys", interactive robo-pets that could react and learn new behaviors.

This decade also saw the launch of DARPA's Urban Challenge for autonomous vehicles; and with Google building its first self-driving car.

Date	AI Development
2000	Interactive robopets (" smart toys ") become commercially available, realizing the vision of the 18th century novelty toy makers.
2000	Cynthia Breazeal at MIT publishes her dissertation on Sociable machines, describing Kismet (robot) , with a face that expresses emotions .
2000	The Nomad robot explores remote regions of Antarctica looking for meteorite samples.
2002	iRobot's Roomba autonomously vacuums the floor while navigating and avoiding obstacles.
2004	OWL Web Ontology Language W3C Recommendation (10 February 2004).
2004	DARPA introduces the DARPA Grand Challenge requiring competitors to produce autonomous vehicles for prize money.
2004	NASA's robotic exploration rovers Spirit and Opportunity autonomously navigate the surface of Mars .
2005	Honda's ASIMO robot, an artificially intelligent humanoid robot, is able to walk as fast as a human, delivering trays to customers in restaurant settings.
2005	Recommendation technology based on tracking web activity or media usage brings AI to marketing. See TiVo Suggestions .
2005	Blue Brain is born, a project to simulate the brain at molecular detail.
2006	The Dartmouth Artificial Intelligence Conference: The Next 50 Years (AI@50) AI@50 (14–16 July 2006)
2007	Philosophical Transactions of the Royal Society, B – Biology , one of the world's oldest scientific journals, puts out a special issue on using AI to understand biological intelligence, titled <i>Models of Natural Action Selection</i> ⁴⁷¹
2007	Checkers is solved by a team of researchers at the University of Alberta .
2007	DARPA launches the Urban Challenge for autonomous cars to obey traffic rules and operate in an urban environment.
2009	Google builds self driving car .

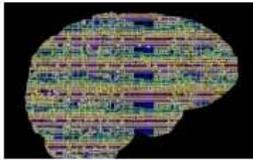
The “Intelligent” Cloud: A Brief History of AI (1950 – 2017)

In the current decade we have seen the introduction of Machine Learning and Deep Learning algorithms, initially running on super-computers, with IBM's Watson computer defeating the Jeopardy Game show champions in 2011; followed by more recently Google's DeepMind AlphaGo defeating the world Go champion Ke Jie.

Date	AI Development
2010	Microsoft launched Kinect for Xbox 360, the first gaming device to track human body movement , using just a 3D camera and infra-red detection, enabling users to play their Xbox 360 wirelessly. The award-winning machine learning for human motion capture technology for this device was developed by the Computer Vision group at Microsoft Research , Cambridge.
2011	IBM's Watson computer defeated television game show Jeopardy! champions Rutter and Jennings .
2011	Apple's Siri , Google's Google Now and Microsoft's Cortana are smartphoneapps that use natural language to answer questions, make recommendations and perform actions.
2013	Robot HRP-2 built by SCHAFT Inc of Japan , a subsidiary of Google , defeats 15 teams to win DARPA's Robotics Challenge Trials . HRP-2 scored 27 out of 32 points in 8 tasks needed in disaster response. Tasks are drive a vehicle, walk over debris, climb a ladder, remove debris, walk through doors, cut through a wall, close valves and connect a hose.
2013	NEIL, the Never Ending Image Learner, is released at Carnegie Mellon University to constantly compare and analyze relationships between different images.
2015	An open letter to ban development and use of autonomous weapons signed by Hawking , Musk , Wozniak and 3,000 researchers in AI and Robotics.
2015	Google DeepMind's AlphaGo defeated 3 time European Go champion 2 dan professional Fan Hui by 5 games to 0.
2016	Google DeepMind's AlphaGo defeated Lee Sedol 4-1. Lee Sedol is a 9 dan professional Korean Go champion who won 27 major tournaments from 2002 to 2016. Before the match with AlphaGo, Lee Sedol was confident in predicting an easy 5-0 or 4-1 victory.
2017	Google DeepMind's AlphaGo won 60-0 rounds on two public Go websites including 3 wins against world Go champion Ke Jie .
2017	Libratus , designed by Carnegie Mellon professor Tuomas Sandholm and his graduate student Noam Brown won against four top players at no-limit Texas hold 'em , a very challenging version of poker. Unlike Go and Chess, Poker is a game in which some information is hidden (the cards of the other player) which makes it much harder to model.

The State of AI Today ... the Good Headlines

So where are we with AI today? Let's take a look at some of the (good) headlines. The promise of AI is being realized in fields as disparate as Health & Medicine, Security Systems, Language Analysis & Creative Writing; and even Scientific discovery.



[Artificial Intelligence Learns And Recreates Complex, Nobel-Winning Physics Experiment](#)



[Facebook's DeepText AI Can Understand What You Write](#)



[Artificial Intelligence is coming to Medicine: Don't be Afraid](#)



[AI2: MIT Researchers Create Artificial Intelligence System To Stop Cyberattacks](#)



[Google DeepMind Plans To Make Eye Disease Diagnosis Easier](#)



[Japanese AI Writes Short Story, Makes It To Finals Of Literary Competition](#)

The State of AI Today ... the Bad Headlines

However, not all of the headline news about AI is good. There is growing concern about the potential harm of AI and Robotics; from replacing jobs ... to replacing humans!

Although there are some justifiable concerns about the ethical issues of relying on AI in "life and death" decision scenarios; the overall impact of AI is that it is more likely to usher in a new wave of innovative services and improved productivity & efficiency, that will ultimately benefit mankind.



[Tens of Millions could lose jobs by 2030 to Robots, Machines, AI and Automation](#)



[Prof Stephen Hawking Warns, Robots Could Evolve Faster Than Human Race and Could Be Hard to Stop](#)

[Danger, danger! 10 alarming examples of AI gone wild](#)



[Experts Warn Against Existential Risks of AI; Scientists Recommend 'Kill Switch' To Prevent AI from Rising Up Against Humans](#)

