The Privacy Foundation

"History Of Artificial Intelligence (AI)"

1308	Catalan poet and theologian Ramon Llull publishes Ars generalis ultima (The Ultimate General Art), further perfecting his method of using paper-based mechanical means to create new knowledge from combinations of concepts.
1666	Mathematician and philosopher Gottfried Leibniz publishes Dissertatio de arte combinatoria (On the Combinatorial Art), following Ramon Llull in proposing an alphabet of human thought and arguing that all ideas are nothing but combinations of a relatively small number of simple concepts.
1763	Thomas Bayes develops a framework for reasoning about the probability of events. Bayesian inference will become a leading approach in machine learning.
1898	**At an electrical exhibition in the recently completed Madison Square Garden, Nikola Tesla makes a demonstration of the world's first radio-controlled vessel. The boat was equipped with, as Tesla described, "a borrowed mind."**
1914	**The Spanish engineer Leonardo Torres y Quevedo demonstrates the first chess-playing machine, capable of king and rook against king endgames without any human intervention.**
1921	Czech writer Karel Čapek introduces **the word "robot"** in his play R.U.R. (Rossum's Universal Robots). The word "robot" comes from the word "robota" (work).
1925	**Houdina Radio Control releases a radio-controlled driverless car, travelling the streets of New York City.**
1927	The science-fiction film Metropolis is released. It features a robot double of a peasant girl, Maria, which unleashes chaos in Berlin of 2026—it was the first robot depicted on film, inspiring the Art Deco look of C-3PO in Star Wars.
1929	Makoto Nishimura designs Gakutensoku, Japanese for "learning from the laws of nature," **the first robot built in Japan**. It could change its facial expression and move its head and hands via an air pressure mechanism.
1943	Warren S. McCulloch and Walter Pitts publish "A Logical Calculus of the Ideas Immanent in Nervous Activity" in the Bulletin of Mathematical Biophysics. This influential paper discussed networks of idealized and simplified artificial "neurons" and …will become the **inspiration for computer-based "neural networks"** (and later "deep learning**") and their popular description as "mimicking the brain."**
1949	Edmund Berkeley publishes Giant Brains: Or Machines That Think in which he writes: "These machines are similar to what a brain would be if it were made of hardware and wire instead of flesh and nerves… it can calculate, conclude, and choose; it can perform reasonable operations with information. **A machine, therefore, can think."**
1951	Marvin Minsky and Dean Edmunds build SNARC (Stochastic Neural Analog Reinforcement Calculator), **the first artificial neural network**, using 3000 vacuum tubes to simulate a network of 40 neurons.
1959	John McCarthy publishes "Programs with Common Sense" in the Proceedings of the Symposium on Mechanization of Thought Processes, in which he describes **the Advice Taker**, a program for solving problems by manipulating sentences in formal languages.
1961	James Slagle develops SAINT (Symbolic Automatic INTegrator - 1961), a heuristic program that **solved symbolic integration problems in freshman calculus**. **making programs "that learn from their experience as effectively as humans do."**
1961	**The first industrial robot, Unimate, starts working on an assembly line in a General Motors plant in New Jersey.**
1964	Daniel Bobrow completes his MIT PhD dissertation titled "Natural Language Input for a Computer Problem Solving System" and develops **STUDENT, a natural language understanding computer program.**
1965	Herbert Simon predicts that "machines will be capable, within twenty years, of doing any work a man can do."
1965	Hubert Dreyfus publishes "Alchemy and AI," arguing that the mind is not like a computer and that there were limits beyond which AI would not progress.
1965	I.J. Good writes in "Speculations Concerning the First Ultraintelligent Machine" that "the first ultraintelligent machine is the last invention that man need ever make, provided that the machine is docile enough to tell us how to keep it under control."
1965	**Joseph Weizenbaum develops ELIZA, an interactive program that carries on a dialogue in English language on any topic. Weizenbaum, who wanted to demonstrate the superficiality of communication between man and machine, was surprised by the number of people who attributed human-like feelings to the computer program.**

1965	Edward Feigenbaum, Bruce G. Buchanan, Joshua Lederberg, and Carl Djerassi start working on DENDRAL at Stanford University. The first expert system, it **automated the decision-making process** and problem-solving behavior of organic chemists, with the general aim of studying hypothesis formation and constructing models of empirical induction in science.
1966	**Shakey the robot is the first general-purpose mobile robot to be able to reason about its own actions**. In a Life magazine 1970 article about this "first electronic person," Marvin Minsky is quoted saying with "certitude": "In from three to eight years we will have a machine with the general intelligence of an average human being."
1961	The film 2001: Space Odyssey is released, featuring Hal, a sentient computer.
1969	Arthur Bryson and Yu-Chi Ho describe back propagation as a multi-stage dynamic system optimization method. A learning algorithm for **multi-layer artificial neural networks,** it has contributed significantly to the success of deep learning in the 2000s and 2010s, once computing power has sufficiently advanced to accommodate the training of large networks.
1970	**The first anthropomorphic robot, the WABOT-1, is built at Waseda University in Japan. It consisted of a limb-control system, a vision system and a conversation system.**
1972	MYCIN, an early expert system for **identifying bacteria** causing severe infections and recommending antibiotics, is developed at Stanford University.
1976	Computer scientist Raj Reddy publishes "**Speech Recognition by Machine**: A Review" in the Proceedings of the IEEE, summarizing the early work on Natural Language Processing (NLP)
1978	The XCON (eXpert CONfigurer) program, a rule-based expert system assisting in the ordering of DEC's VAX computers by automatically selecting the components **based on the customer's requirements**, is developed at Carnegie Mellon University.
1979	The Stanford Cart successfully crosses a chair-filled room without human intervention in about five hours, becoming one of the earliest examples of an **autonomous vehicle.**
1980	**Wabot-2 is built at Waseda University in Japan, a musician humanoid robot able to communicate with a person, read a musical score and play tunes of average difficulty on an electronic organ.**
1981	The Japanese Ministry of International Trade and Industry budgets $850 million for the Fifth Generation Computer project. The project aimed to develop computers that could carry on conversations, translate languages, interpret pictures, and reason like human beings.
1984	Electric Dreams is released, a film about a love triangle between a man, a woman and a personal computer.
1984	At the annual meeting of AAAI, Roger Schank and Marvin Minsky warn of the coming "AI Winter," predicting an immanent bursting of the AI bubble (which did happen three years later), similar to the reduction in AI investment and research funding in the mid-1970s.
1986	**First driverless car, a Mercedes-Benz van equipped with cameras and sensors, built at Bundeswehr University in Munich under the direction of Ernst Dickmanns, drives up to 55 mph on empty streets.**
1986	David Rumelhart, Geoffrey Hinton, and Ronald Williams publish "Learning representations by back-propagating errors," in which they describe "a new learning procedure, back-propagation, for networks of neurone-like units."
1987	The video Knowledge Navigator, accompanying Apple CEO John Sculley's keynote speech at Educom, envisions a future in which "knowledge applications would be accessed by smart agents working over networks connected to massive amounts of digitized information."
1988	Rollo Carpenter develops the chat-bot Jabberwacky to "simulate natural human chat in an interesting, entertaining and humorous manner." It is an early attempt at creating artificial intelligence through human interaction.
1990	Rodney Brooks publishes "Elephants Don't Play Chess," proposing a new approach to AI—building intelligent systems, specifically robots, from the ground up and on the basis of ongoing physical interaction with the environment: "The world is its own best model… The trick is to sense it appropriately and often enough."
1993	Vernor Vinge publishes "The Coming Technological Singularity," in which he predicts that "within thirty years, we will have the technological means to create superhuman intelligence. Shortly after, the human era will be ended."
1995	Richard Wallace develops the chatbot A.L.I.C.E (Artificial Linguistic Internet Computer Entity), inspired by Joseph Weizenbaum's ELIZA program, but with the addition of natural language sample data collection on an unprecedented scale, enabled by the advent of the Web.
1997	Sepp Hochreiter and Jürgen Schmidhuber propose Long Short-Term Memory (LSTM), a type of a recurrent neural

	network used today in handwriting recognition and speech recognition.
1997	**Deep Blue becomes the first computer chess-playing program to beat a reigning world chess champion.**
1998	**Dave Hampton and Caleb Chung create Furby, the first domestic or pet robot.**
1998	Yann LeCun, Yoshua Bengio and others publish papers on the application of neural networks to handwriting recognition and on optimizing backpropagation.
2000	**MIT's Cynthia Breazeal develops Kismet, a robot that could recognize and simulate emotions.**
2000	**Honda's ASIMO robot, an artificially intelligent humanoid robot, is able to walk as fast as a human, delivering trays to customers in a restaurant setting.**
2004	The first DARPA Grand Challenge, a prize competition for autonomous vehicles, is held in the Mojave Desert. None of the autonomous vehicles finished the 150-mile route.
2006	Oren Etzioni, Michele Banko, and Michael Cafarella coin the term **"machine reading,"** defining it as an inherently unsupervised "autonomous understanding of text."
2006	Geoffrey Hinton publishes "Learning Multiple Layers of Representation," summarizing the ideas that have led to "multilayer neural networks that contain top-down connections and training them to generate sensory data rather than to classify it," i.e., the new approaches to deep learning.
2007	Fei Fei Li and colleagues at Princeton University start to assemble ImageNet, a large database of annotated images designed to aid in **visual object recognition software research.**
2009	Rajat Raina, Anand Madhavan and Andrew Ng publish "Large-scale Deep Unsupervised Learning using Graphics Processors," arguing that "modern graphics processors far surpass the computational capabilities of multicore CPUs, and have the potential to revolutionize the applicability of deep unsupervised learning methods."
2009	**Google starts developing, in secret, a driverless car. In 2014, it became the first to pass, in Nevada, a U.S. state self-driving test.**
2009	Google starts developing, in secret, a driverless car. In 2014, it became the first to pass, in NeComputer scientists at the Intelligent Information Laboratory at Northwestern University develop Stats Monkey, a program that writes sport news stories without human interventionvada, a U.S. state self-driving test
2009	Computer scientists at the Intelligent Information Laboratory at Northwestern University develop Stats Monkey, a program that **writes sport news stories without human intervention.**
2010	Launch of the ImageNet Large Scale Visual Recognition Challenge (ILSVCR), an annual AI object recognition competition.
2011	**Watson, a natural language question-answering computer, competes on Jeopardy! and defeats two former champions.**
2012	Jeff Dean and Andrew Ng report on an experiment in which they showed a very large neural network 10 million unlabeled images randomly taken from YouTube videos, and "to our amusement, one of our artificial neurons learned to **respond strongly to pictures of... cats."**
2012	A convolutional neural network designed by researchers at the University of Toronto achieve an error rate of only 16% in the ImageNet Large Scale Visual Recognition Challenge, a significant improvement over the 25% error rate achieved by the best entry the year before.
2016	Google DeepMind's AlphaGo defeats Go champion Lee Sedol. (Chess computers games)
	TBD

Please see below for more details and information:
"A Very Short History Of Artificial Intelligence (AI)" by Gill Press @GilPress Forbes
https://www.forbes.com/sites/gilpress/2016/12/30/a-very-short-history-of-artificial-intelligence-ai/#41f142dd6fb

Aggregated for academic purposes only:

John Soma
Executive Director - The Privacy Foundation || jsoma@law.du.edu

Yasmin Shanker
Research Assistant || yshanker19@law.du.edu