timelines.me

History of Artificial Intelligence

From Turing's 'Can machines think?' to ChatGPT and Nobel Prizes. 30 milestones across seven decades of AI — breakthroughs, winters, and the generative revolution. Sources: Nature, Science, NeurIPS, official announcements.

by timelines.me · 30 events · Created 26 Mar 2026

Exported 26 May 2026

1950-10-01

Alan Turing asks: Can machines think?

[foundations, philosophy, Turing Test]

Turing publishes 'Computing Machinery and Intelligence' in Mind journal. Introduces the Imitation Game (later known as the Turing Test) — a benchmark for machine intelligence that shaped AI research for decades.

1956-08-01

The birth of AI as a field

[foundations, Dartmouth, McCarthy]

Dartmouth Conference: John McCarthy, Marvin Minsky, Claude Shannon, and Nathaniel Rochester coin the term 'Artificial Intelligence.' The workshop defines AI as a formal academic discipline. Funding and optimism surge.

1958-07-01

The Perceptron: first neural network hardware

[neural networks, Perceptron, foundations]

Frank Rosenblatt at Cornell builds the Mark I Perceptron, the first machine capable of learning through trial and error. The Navy funds the project. The New York Times reports a computer that 'will be able to walk, talk, see, write, reproduce itself and be conscious of its existence.'

1966-01-01

ELIZA: the first chatbot

[chatbot, NLP, MIT]

Joseph Weizenbaum at MIT creates ELIZA, a program that mimics a psychotherapist. Users form emotional attachments to it despite its simplicity — an early glimpse of human-AI interaction dynamics.

1969-07-01

Shakey the Robot: AI meets the physical world

[robotics, SRI International, computer vision]

SRI International builds Shakey, the first general-purpose mobile robot that reasons about its own actions. It combines computer vision, route planning, and problem solving. Shakey inspires decades of robotics and autonomous systems research.

1974-01-01

The first AI winter

[AI winter, funding, setback]

Funding collapses after the Lighthill Report (UK) declares AI research has failed to deliver on its promises. DARPA cuts funding. Researchers scatter. AI enters a decade of disillusionment.

1980-01-01

Expert systems boom

[expert systems, commercial AI, Japan]

Rule-based expert systems like MYCIN (medical diagnosis) and XCON (computer configuration) prove commercially viable. Corporations invest billions. Japan launches the Fifth Generation Computer project.

1986-07-01

Backpropagation revives neural networks

[neural networks, backpropagation, Hinton]

Rumelhart, Hinton, and Williams publish 'Learning representations by back-propagating errors' in Nature. The paper provides an efficient method to train multi-layer neural networks, resurrecting a field that Minsky and Papert had nearly killed in 1969.

1987-01-01

The second AI winter

[AI winter, expert systems, setback]

Expert systems prove brittle and expensive to maintain. The AI bubble bursts. Lisp machine market collapses. Japan's Fifth Generation project fails to meet its goals. Funding dries up again.

1991-08-06

The World Wide Web enables data collection at scale

[World Wide Web, data, infrastructure]

Tim Berners-Lee publishes the first website, launching the World Wide Web. Within a decade, billions of web pages create the massive text and image datasets that later fuel machine learning. No web, no big data; no big data, no modern AI.

1997-05-11

Deep Blue defeats world chess champion

[chess, IBM, Deep Blue, milestone]

IBM's Deep Blue beats Garry Kasparov in a six-game match. First time a computer defeats a reigning world champion under standard rules. The event captivates the world but uses brute force, not learning.

2006-07-01

The rise of deep learning

[deep learning, neural networks, Hinton]

Geoffrey Hinton publishes breakthrough work on deep belief networks, reviving neural network research. The term 'deep learning' enters mainstream AI vocabulary. GPUs make training large networks feasible.

2006-10-02

Netflix Prize: crowdsourcing machine learning

[machine learning, Netflix, competition, recommendations]

Netflix offers $1 million to anyone who can improve its recommendation algorithm by 10%. The competition attracts 40,000 teams from 186 countries over three years. BellKor's Pragmatic Chaos wins in 2009, proving collaborative filtering at scale.

2011-02-16

IBM Watson wins Jeopardy!

[Watson, IBM, NLP, milestone]

Watson defeats champions Ken Jennings and Brad Rutter on live television. Unlike Deep Blue, Watson processes natural language — a qualitative leap in AI capability visible to millions.

2011-10-04

Siri brings AI to every pocket

[Siri, Apple, voice assistant, consumer AI]

Apple launches Siri with the iPhone 4S, making voice-activated AI a mass-market consumer product. Originally a DARPA-funded SRI International project, Siri normalizes the idea of talking to machines for hundreds of millions of users.

2012-09-30

ImageNet moment: deep learning conquers vision

[ImageNet, deep learning, computer vision]

AlexNet wins the ImageNet competition with a 10-point margin over traditional methods. Deep learning proves dramatically superior for image recognition. The modern AI era begins.

2014-06-10

GANs: machines learn to imagine

[GAN, generative AI, Goodfellow, neural networks]

Ian Goodfellow introduces Generative Adversarial Networks in a landmark paper. Two neural networks compete — one generates, one discriminates — producing increasingly realistic

synthetic data. GANs become the foundation of AI image generation.

2016-03-15

AlphaGo defeats Lee Sedol at Go

[AlphaGo, DeepMind, Go, reinforcement learning]

DeepMind's AlphaGo beats world champion Lee Sedol 4-1 in a game long considered too complex for computers. Move 37 in Game 2 — a play no human would make — becomes iconic. 200 million people watch live.

2017-06-12

Transformers: Attention Is All You Need

[Transformers, Google, architecture, deep learning]

Google researchers publish the Transformer architecture paper. Self-attention mechanism replaces recurrence. This single paper becomes the foundation of GPT, BERT, and every major language model that follows.

2018-12-05

Waymo launches self-driving car service

[self-driving cars, Waymo, autonomous vehicles]

Waymo One becomes the first commercial autonomous ride-hailing service, launching in Phoenix, Arizona. Passengers ride without a safety driver. By 2024, the service expands to San Francisco and Los Angeles, logging millions of autonomous miles.

2019-02-14

GPT-2: 'Too dangerous to release'

[GPT-2, OpenAI, AI safety, language model]

OpenAI announces GPT-2 with 1.5B parameters. Initially withholds the full model citing misuse concerns. The decision sparks global debate about AI safety, open-source ethics, and who controls powerful AI.

2020-11-30

AlphaFold solves protein folding

[AlphaFold, DeepMind, biology, protein folding]

DeepMind's AlphaFold2 predicts protein structures with atomic accuracy, solving a 50-year-old grand challenge of biology. Later releases structures for 200M+ proteins, accelerating drug discovery worldwide.

2021-01-05

DALL-E: AI learns to create images from text

[DALL-E, OpenAI, generative AI, images]

OpenAI reveals DALL-E, generating images from text descriptions. 'An armchair in the shape of an avocado' becomes the symbol of generative AI. The creative industries begin to reckon with

AI.

2021-06-29

GitHub Copilot: AI writes code

[Copilot, GitHub, OpenAI, programming]

GitHub and OpenAI launch Copilot, an AI pair programmer trained on billions of lines of public code. Within two years, it generates 46% of code in files where it is enabled. Software development begins its biggest workflow shift since Stack Overflow.

2022-08-22

Stable Diffusion: open-source image generation

[Stable Diffusion, open-source, generative AI, images]

Stability AI releases Stable Diffusion as open-source, allowing anyone to generate images locally. Millions download it within weeks. The decision democratizes generative AI and ignites debate over copyright, artist rights, and model governance.

2022-11-30

ChatGPT: AI goes mainstream

[ChatGPT, OpenAI, mainstream, language model]

OpenAI launches ChatGPT. Reaches 100 million users in 2 months — the fastest-growing consumer application in history. AI becomes a daily tool for millions. The conversation shifts from 'if' to 'how fast.'

2023-03-14

GPT-4: multimodal AI arrives

[GPT-4, OpenAI, multimodal, reasoning]

OpenAI releases GPT-4 — processes text and images, passes the bar exam in the 90th percentile, and demonstrates reasoning capabilities that surprise even its creators. The gap between AI and human cognition narrows visibly.

2024-03-13

EU AI Act: first comprehensive AI regulation

[regulation, EU, AI Act, governance]

European Parliament passes the AI Act — the world's first comprehensive legal framework for artificial intelligence. Risk-based approach: bans social scoring, regulates high-risk systems, requires transparency for generative AI.

2024-07-01

Claude, Gemini, open-source: the AI race accelerates

[Claude, Gemini, Llama, open-source, competition]

Anthropic's Claude, Google's Gemini, Meta's Llama, and Mistral push capabilities forward. Open-source models close the gap with proprietary ones. AI becomes commoditized — the

question shifts to trust, safety, and alignment.

2024-10-08

Nobel Prize recognizes AI pioneers

[Nobel Prize, Hinton, Hassabis, recognition]

Geoffrey Hinton shares the Nobel Prize in Physics for foundational work on artificial neural networks. Days later, Demis Hassabis receives the Nobel Prize in Chemistry for AlphaFold's protein structure predictions. AI research earns science's highest honors for the first time.