



THE HISTORY OF ARTIFICIAL INTELLIGENCE

From the First Ideas to the Age of AI Agents

Coverage through 8 October 2026

An unfolding story of human curiosity, scientific persistence, unexpected setbacks and the systems now entering everyday life.

From automata and the first calculating engines to learning networks, creative tools, agents and robots.





Contents

A history of ideas, setbacks and breakthroughs.

01	Before artificial intelligence had a name	3	14	The organisations shaping modern AI	21
02	The scientific foundations	4	15	The rise of AI in China	24
03	When artificial intelligence became a field	6	16	AI beyond the US and China	25
04	The first AI systems	7	17	Images, video, music and creative AI	26
05	Neural networks: the idea that refused to die	8	18	AI becomes multimodal	27
06	The AI winters	10	19	From chatbots to agents	28
07	Expert systems and symbolic AI	11	20	Robotics and embodied AI	30
08	Machine learning changes the approach	12	21	The people who shaped AI	31
09	Data, GPUs and the deep learning revolution	13	22	Major model families to recognise in 2026	33
10	DeepMind and learning through action	15	23	Where we are in October 2026	34
11	Generative AI before ChatGPT	17		• How all these pieces became AI	35
12	The transformer changes everything	18		• The AI history timeline	36
13	The age of large language models	19		• A short glossary	38
				• Sources and further reading	40

Reading the recent history

A dated paper, event or release is a documented fact. Performance described in a company’s own tests is a company claim unless independently validated. Developing work is still changing; an open question has no settled answer. Source numbers throughout the chapters are clickable.

**CHAPTER 01**

Before artificial intelligence had a name

Long before there were computers, people imagined making something that could act as though it had a mind. Greek stories told of Talos, a bronze guardian. In workshops across different cultures, makers built figures that moved by water, weights, air or hidden gears. Around 1206, the engineer al-Jazari described ingenious clocks and mechanical figures in a book illustrated with their workings. These creations did not learn or reason. Their importance lies in the human question behind them: how much of a seemingly living action could a person build?

That question took a new form when calculating machines appeared. A device that could perform arithmetic reliably suggested that at least one human mental activity could be made mechanical. In the nineteenth century, Charles Babbage designed the Analytical Engine. It was never completed in his lifetime, but its proposed instructions, memory and general-purpose design reached far beyond a machine for one calculation.

Ada Lovelace saw possibilities in Babbage's plans that went beyond numbers. In 1843 she published extensive notes on the Engine, including a detailed procedure for calculating Bernoulli numbers. She also considered whether a machine might work with symbols representing music, if those relationships could be expressed clearly enough. Lovelace did not build an AI system. She helped people picture a machine that could carry out general symbolic procedures - an idea future AI would need.

A moving automaton is not an AI. Babbage and Lovelace mattered because they made programmable calculation imaginable.

Sources: [1] · [2]

**CHAPTER 02**

The scientific foundations

The next step was to make reasoning itself easier to describe. George Boole's nineteenth-century work showed how statements such as 'and', 'or' and 'not' could be handled with algebra. It did not reduce all thought to a few switches. It offered a precise language for some kinds of logic - a language engineers could later put into circuits.

In 1936, Alan Turing described an abstract machine that follows a finite set of instructions, reading and writing symbols one step at a time. The 'Turing machine' was a thought experiment, not a product. It clarified what it means for a procedure to be computable. Turing then turned to a bolder question. In his 1950 paper 'Computing Machinery and Intelligence', he proposed the imitation game: could a machine converse well enough that a human judge, communicating through text, could not reliably tell it from a person? The later name 'Turing Test' refers to this proposal. Passing a conversational test would show something important about behaviour; it would not, by itself, settle what understanding or experience a machine has.

Other foundations arrived from different directions. Claude Shannon's 1948 information theory gave researchers mathematical tools for studying messages and uncertainty. Norbert Wiener studied feedback: how a machine or organism can use information about what happened to adjust what it does next. In 1943, Warren McCulloch and Walter Pitts published a mathematical model of simplified nerve cells connected in networks. Their artificial 'neurons' were logical abstractions, far simpler than biology. Still, they opened a path toward machines whose behaviour arose from many connected units rather than a long list of explicit instructions.

By the early 1950s, the ingredients were scattered across logic, mathematics, engineering, psychology and neuroscience. Someone still had to bring them together under a shared question.



The scientific foundations · continued

Turing changed the question

The important shift was not a claim that machines already had minds. It was a way to ask a difficult question without pretending we could see inside one.

By 1950, Alan Turing had already helped clarify what it means for a procedure to be computable. Then he turned toward intelligence. Instead of beginning with a definition of 'thinking' that everyone would argue about, he proposed a practical test built around conversation.

That choice still feels modern. AI research repeatedly runs into things we cannot observe directly. Researchers therefore look for behaviour, internal signals, causal interventions and carefully designed tests - each one answering a narrower question than 'is there a mind in there?'

The Turing Test concerns observable conversational behaviour. Passing it would be interesting evidence about performance; it would not, by itself, prove understanding, consciousness or subjective experience.

The approach asks us to break a huge philosophical question into something that can be measured, to distinguish behaviour from inner experience, and to keep uncertainty visible. A convincing performance can matter without settling what, if anything, is felt. Good AI research often advances by replacing one grand question with several smaller, answerable ones.

Sources: [3] · [4] · [5] · [6] · [7] · [96] · [8] · [9]

**CHAPTER 03**

When artificial intelligence became a field

In 1955, John McCarthy proposed a summer research meeting at Dartmouth College. Marvin Minsky, Claude Shannon and IBM engineer Nathaniel Rochester joined him on the proposal. McCarthy chose the name 'artificial intelligence'. In the summer of 1956, researchers gathered in Hanover, New Hampshire, to ask whether aspects of learning and intelligence could be described precisely enough for a machine to simulate them.

Among the people drawn into this early circle were Allen Newell and Herbert Simon, working with programmer Cliff Shaw. They brought a working program, Logic Theorist, that could find proofs for some statements in symbolic logic. It was a remarkable demonstration in a period when computers were scarce, slow and difficult to use.

The optimism is easier to understand if we imagine what it felt like. A machine had begun to carry out steps that looked like reasoning. If it could prove a theorem, perhaps language, vision and common sense would follow. Researchers badly underestimated how much ordinary intelligence depends on an enormous, flexible grasp of the world. Yet the Dartmouth meeting gave disparate efforts a name, a community and a research agenda. It is generally treated as the formal beginning of AI as a field - not as the instant when intelligence was invented.

1956 is the birth of a research field. The ideas it gathered had older roots, and the hard problems remained open.

Sources: [8] · [9]

**CHAPTER 04**

The first AI systems

The earliest programs were small windows into large ambitions. Newell, Simon and Shaw's Logic Theorist searched for proofs. Their later General Problem Solver tried to apply general strategies to well-described problems. Both worked inside carefully prepared representations; neither could freely understand an unfamiliar real-world situation.

Games made another useful laboratory. Arthur Samuel's checkers work explored how a program might improve through experience. Earlier chess pioneers, including Shannon and Turing, had explored how to turn moves and positions into computable search. Games supplied clear rules and a score, which made progress visible. The outside world does not provide such tidy rules.

In the mid-1960s, Joseph Weizenbaum built ELIZA. One of its scripts imitated a therapist by picking up words from a user's sentence and reflecting them in a question. Some users felt the program understood them. Weizenbaum was disturbed by how readily a convincing conversation could invite that belief. In the early 1970s, Terry Winograd's SHRDLU responded to language about a small virtual world of blocks. It could manipulate and discuss that world impressively because the world had been defined so precisely.

These systems could prove, search, converse or act within narrow settings. They did not possess the open-ended understanding people hoped would arrive soon. That gap between a vivid demonstration and a general ability would become a recurring theme in AI history.

Sources: [9] · [10] · [11]



CHAPTER 05

Neural networks: the idea that refused to die

Alongside the programs that handled symbols and rules, another tradition asked whether machines could learn through networks of simple units. McCulloch and Pitts had given this idea an early mathematical form. In 1957-58, psychologist Frank Rosenblatt introduced the perceptron: a system that adjusted the importance of its inputs after examples. Imagine a sorter that gradually learns which clues matter for a simple classification. This was a real form of learning, but a basic perceptron could draw only limited kinds of boundaries between examples.

In 1969, Marvin Minsky and Seymour Papert carefully analysed those limits in their book 'Perceptrons'. The book was influential; so were hardware limits, scarce data, and other approaches attracting resources. It is too simple to say the book single-handedly 'killed' neural networks. Research continued, but enthusiasm and funding for many connectionist approaches weakened.

The crucial idea was to connect units in multiple layers. A neural network is a set of mathematical units linked by adjustable numbers, often called weights or parameters. 'Deep learning' means learning with networks that have several layers, allowing a system to build useful features from simpler ones. A parameter is an adjustable number; a modern model can contain very many of them. Networks are inspired in part by brains, but their artificial neurons are not literal copies of living cells.

A training method called backpropagation sends information about a network's error backward through the layers so its parameters can be adjusted. Related ideas predated the famous 1986 paper by David Rumelhart, Geoffrey Hinton and Ronald Williams; that paper helped show how useful the method could be. Yann LeCun developed image-reading networks that recognised handwritten digits, while Yoshua Bengio advanced ways to learn useful representations and language patterns. All three later shared the 2018 Turing Award for deep learning contributions.



Neural networks: the idea that refused to die · continued

For sequences such as speech or writing, recurrent neural networks reused a working state as new items arrived. In 1997 Sepp Hochreiter and Jürgen Schmidhuber introduced long short-term memory, or LSTM, to help such networks retain useful signals across longer stretches. Many of these ideas were technically promising decades before data, computing power and practical refinements made them central to everyday products.

A neural network learns by changing many numerical connections. Its name borrows from biology; its workings are mathematical.

An idea carried across generations

The core idea survived several changes in fashion because the problem it addressed never went away: can a machine learn useful internal structure from examples?

In 1943, McCulloch and Pitts described a mathematical network of simplified artificial neurons. Rosenblatt's 1957-58 perceptron learned simple classification boundaries from examples. In 1969, Minsky and Papert analysed important limits of basic perceptrons; enthusiasm cooled for several reasons. In 1986, Rumelhart, Hinton and Williams made backpropagation newly influential for multilayer networks.

During the 1990s, LeCun, Bengio, Hochreiter and Schmidhuber advanced networks for images, language and long sequences. In 2012, AlexNet brought data, GPUs and deep networks together at scale.

An idea can be scientifically promising long before hardware, data and engineering make it practical. AI history is full of concepts that looked disappointing, then returned under new conditions.

Sources: [7] · [12] · [13] · [14] · [15] · [24]

**CHAPTER 06**

The AI winters

AI's early demonstrations fed predictions that broad machine intelligence might arrive in a few years. Everyday language, robust vision and common sense turned out to be far harder than proving a theorem inside a formal system. In the 1970s, disappointed funders became less willing to support sweeping promises. A critical 1973 report by James Lighthill in Britain and changing priorities in US research funding became emblems of what is often called the first 'AI winter'.

The phrase should not be taken too literally. Historians disagree about how completely the first winter froze the field: researchers, laboratories and new ideas survived and sometimes expanded. What cooled was confidence in particular promises, along with some important sources of funding. In the 1980s, commercial expert systems brought enthusiasm back. Then maintaining their hand-built knowledge became costly, specialised AI hardware lost its advantage, and business expectations outran the results. The late 1980s and early 1990s are commonly called the second AI winter.

A winter did not mean that no AI worked. Researchers kept developing machine learning, probability, language processing, robotics, reinforcement learning and neural networks. The lesson is not that each new advance is an illusion. It is that a compelling local success can be mistaken for an almost-finished solution to intelligence as a whole.

Then, a few striking laboratory successes were sometimes treated as a promise of near-term general intelligence. Now, wider real-world use gives us more evidence - and more ways for limitations to appear.

Sources: [16] · [17] · [97]

**CHAPTER 07**

Expert systems and symbolic AI

Another answer to the early AI question was to put expert knowledge into a program. An expert system stored facts in a knowledge base and used rules such as 'if these conditions hold, consider this conclusion'. This is one form of symbolic AI: a program handles named things and explicit relationships between them. It could explain a chain of rules in a way many learned systems could not.

Stanford's MYCIN, developed in the 1970s under Edward Shortliffe with colleagues and advisers, asked questions about serious bacterial infections and suggested antibiotics. In limited test cases it performed impressively. It was a research system, however, and was not put into routine hospital use. That distinction matters whenever an AI demonstration concerns medicine or any other high-stakes setting.

By the 1980s, businesses were building systems for diagnosis, configuration and decision support. They could be valuable when the subject was stable, narrow and well documented. But someone had to interview specialists, write the rules, settle contradictions and revise them as reality changed. It was difficult to write down the unspoken common sense experts rely on. As the number of exceptions grew, maintaining the knowledge base could become expensive. The next generation increasingly learned patterns from examples instead.

Sources: [17] · [18]

**CHAPTER 08**

Machine learning changes the approach

A quiet shift changed the central question. Instead of telling a machine every condition it might meet, researchers increasingly showed it examples and measured whether it could learn a pattern that held up on new examples. That is machine learning. Statistical methods used data and probability to cope with uncertain, messy inputs. Bayesian approaches offered ways to update beliefs when evidence changed; support vector machines found effective boundaries for sorting some kinds of data. They did not solve every problem, but they made learning a practical engineering strategy.

Another strand was reinforcement learning. Rather than being shown the right answer at every step, a system tries actions and receives signals about how well they served a goal. Richard Sutton and Andrew Barto helped build this field's conceptual and algorithmic foundations. In 2024 they received the Turing Award for that work. The method is powerful when a task offers useful feedback; designing the right goal and transferring a learned policy into the open world remain difficult.

The public saw a different kind of machine triumph in 1997, when IBM's Deep Blue defeated world chess champion Garry Kasparov in a match. Deep Blue combined specialised hardware, human chess knowledge and enormous search. It was an extraordinary achievement in one precisely defined domain. It was not a general-purpose model that could be asked to explain a photograph or plan a holiday. AI had reached a landmark without becoming the one broad intelligence its founders had imagined.

Sources: [19] · [20]



CHAPTER 09

Data, GPUs and the deep learning revolution

Three currents finally met. The internet and digital cameras produced far more examples. Researchers improved learning methods and network designs. And a chip made for drawing computer graphics turned out to be excellent at doing many similar calculations in parallel. That chip is the graphics processing unit, or GPU. NVIDIA's CUDA software, introduced in 2006, helped researchers use GPUs for general computation. A 'benchmark' is a common test that lets researchers compare systems under stated conditions; a good benchmark makes progress visible, though it never covers everything the real world demands.

Fei-Fei Li and collaborators built ImageNet, a huge collection of labelled pictures. Their team helped organise the ImageNet recognition challenge, giving researchers a shared task. In 2012, Alex Krizhevsky, Ilya Sutskever and Geoffrey Hinton entered a deep image network later called AlexNet. Trained with GPUs, it won by a striking margin. The result did not mean machines could see as people do. It showed that learned visual features, large data and parallel computing could together beat years of carefully designed alternatives on a demanding test.

That single result changed research culture. More labs invested in deep learning for pictures, then speech recognition, translation and other tasks. Data sets and competitions spread; GPUs and software became more capable; researchers built on one another's models. NVIDIA became central because it supplied chips and a development ecosystem needed to train and run very large networks. There were other important hardware makers, academic labs and data contributors. The shift came from a whole system, not one chip or one paper.



Data, GPUs and the deep learning revolution · continued

Even this revolution had a prehistory. The networks that flourished in 2012 depended on ideas carried through periods when they seemed unfashionable. The machines had finally caught up with a long-running line of thought.

ImageNet supplied a shared mountain of examples. AlexNet showed that a deep network trained with GPUs could use them unusually well.

2012: the conditions caught up with the idea

AlexNet is memorable because it was not one isolated invention. Several long-running threads finally met in the same experiment: ImageNet supplied a huge labelled image collection and a shared benchmark; GPUs made many neural-network calculations practical in parallel; and deep-network training and engineering had matured enough to use both effectively.

Alex Krizhevsky, Ilya Sutskever and Geoffrey Hinton did not invent neural networks in 2012. Fei-Fei Li and the ImageNet community did not invent computer vision. NVIDIA did not invent learning. The turning point came from a system of contributions finally reinforcing one another.

After AlexNet, deep learning rapidly became the dominant direction for vision and then spread through speech, translation and other areas. The result changed research culture as much as it changed a benchmark score.

Sources: [21] · [22] · [23] · [24]



CHAPTER 10

DeepMind and learning through action

In London in 2010, Demis Hassabis, Shane Legg and Mustafa Suleyman founded DeepMind. Google acquired it in 2014. Its researchers joined deep neural networks to reinforcement learning, showing systems that learned to play Atari games from screen images and game scores. Instead of receiving a hand-written guide to each game, an agent tried actions and improved from feedback. These were still contained digital worlds, but the combination was powerful.

In 2016, AlphaGo defeated the great Go player Lee Sedol four games to one. Go's enormous range of possible positions had made this seem distant to many observers. AlphaGo combined trained networks with search and practice. Its unexpected moves gave the match cultural force: the program was not reciting a fixed book of master plays. Later AlphaZero learned strong play in chess, shogi and Go from the rules and self-play, although each game still had a defined arena and reward.

DeepMind also moved from games toward science. AlphaFold's advances in predicting protein shapes helped make structural information available to researchers at unprecedented scale. Demis Hassabis and John Jumper shared half of the 2024 Nobel Prize in Chemistry for protein structure prediction. Their achievement illustrated a broader point: AI could help scientific work, even while a system brilliant at proteins might know nothing about handling a household object.



DeepMind and learning through action · continued

When a machine surprised the masters

AlphaGo mattered because it turned a technical result into a cultural event. People were watching a machine operate in a domain long associated with human intuition.

In 2016, AlphaGo defeated Lee Sedol four games to one. The system combined trained neural networks, search and repeated practice. What captured attention was not merely the final score. Some of its moves were unexpected enough to make observers reconsider what machine-generated strategy could look like.

AlphaGo was extraordinary at Go. It was not a general intelligence. A system can display superhuman performance inside one structured domain while remaining unable to handle ordinary tasks outside it.

Games make progress visible: rules, legal moves and a score create unusually clear feedback. Narrow competence can still be profound; a system does not need general intelligence to transform a field. The surprise changes humans too. AI milestones often force people to rethink which abilities they assumed required human intuition.

Sources: [25] · [26] · [27] · [98]



CHAPTER 11

Generative AI before ChatGPT

As machines learned to sort and predict, researchers also asked them to create. One important step was to represent words with numerical patterns, called word embeddings, so that words used in similar contexts could lie closer together in a mathematical space. Work such as Word2Vec helped make relationships in language easier for models to use.

For translation, sequence-to-sequence systems learned to turn one string of words into another. Early versions compressed an input into a limited internal summary. Attention methods helped a model look back at the relevant parts of the source as it generated each part of the answer. This was a step toward, though not yet the same as, the transformer.

Images had their own path. In 2014 Ian Goodfellow and colleagues proposed generative adversarial networks, or GANs. Imagine one network making images and another learning to distinguish generated from real examples; each pressures the other to improve. Variational autoencoders, or VAEs, provided another way to learn a compressed space from which new examples could be made. Later diffusion approaches learned to turn noise gradually into an image. These methods had different strengths and limits. Together they made 'generative AI' - AI that produces new text, images or other media - much more than a synonym for chatbots.

Sources: [28] · [29] · [30] · [31] · [99]



CHAPTER 12

The transformer changes everything

In 2017 a team of eight Google researchers published a paper with the memorable title 'Attention Is All You Need'. Its first author was Ashish Vaswani; the full team also included Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan Gomez, Łukasz Kaiser and Illia Polosukhin. Their transformer was designed for translation. It replaced an older dependence on processing a sentence one step after another with a design that could relate parts of a sequence more directly.

In plain English, attention lets a model weigh which pieces of the available material matter for the piece it is processing. In 'the trophy would not fit in the suitcase because it was too large', the surrounding words help determine what 'it' refers to. A transformer learns useful relationships of this kind as numerical patterns. Attention is not human concentration, and it does not guarantee understanding, but it makes context easier for a model to use.

The architecture also made training on large collections of text more practical, because many calculations could be performed in parallel on modern hardware. Subsequent researchers built different transformer forms for reading, writing, images, audio and much else. Transformers were a foundation, not a solitary miracle. The modern language model required the architecture together with vast data, specialised computing, better training methods and careful work on how a system responds to people.

The transformer offered a flexible way to connect parts of a sequence and train at scale. It became the backbone of many, though not all, modern AI systems.

Sources: [32] · [33]



CHAPTER 13

The age of large language models

A large language model, or LLM, is a neural network trained on a great deal of language to predict and generate sequences of text. OpenAI's first GPT research in 2018 showed a useful pattern: train a transformer on broad text, then adapt it to particular tasks. GPT-2 in 2019 produced longer, more coherent passages. GPT-3 in 2020 grew much larger and showed 'few-shot learning': it could often infer the requested pattern from a handful of examples placed in the prompt, without retraining for that task. An impressive continuation of text could still be false or unsuitable as an answer.

Researchers studied scaling laws - regular trends in how performance changes with data, model size and computing effort. The lesson was subtler than 'make it bigger': the balance among those ingredients matters. Other teams refined training so models would respond to people more usefully. Instruction tuning used examples of good responses to requests. Reinforcement learning from human feedback, or RLHF, used human preferences to help shape answers. OpenAI's InstructGPT research reported improvements from these methods; related work across the field mattered too.

On 30 November 2022, OpenAI released ChatGPT as a conversational product built on this research. Its ability to answer follow-up questions in an ordinary chat made years of laboratory progress feel suddenly close to everyday life. The system still made mistakes and could invent plausible claims. But many people could now try a powerful language model without learning to program. That social change was as significant as the technical milestone: a research tradition had become a daily encounter.

ChatGPT was a public doorway into decades of work. Its 2022 launch made AI tangible to ordinary people; it did not mark AI's beginning.



The age of large language models · continued

When AI became a daily encounter

ChatGPT did not begin AI. It changed who could directly experience a powerful language model.

Before November 2022, decades of work had already produced transformers, large language models, instruction tuning and reinforcement learning from human feedback. What changed with ChatGPT was access: ordinary people could simply open a conversation and ask a follow-up question.

That made an abstract research field socially immediate. People began using the same interface for work, learning, creativity, companionship, frustration, experiments and questions about the system itself. The strengths and failures were no longer visible only to researchers; they became part of everyday experience.

The technical lineage mattered, but the social shift mattered too. A research tradition became a relationship millions of people could test for themselves - including discovering how persuasive, useful, strange and fallible an AI conversation can be.

ChatGPT was not the moment AI was invented. It was the moment a large part of the public stopped encountering AI mainly through news stories, search rankings or recommendation systems and started talking to it directly.

Sources: [34] · [35] · [36] · [37] · [38] · [39]

**CHAPTER 14**

The organisations shaping modern AI

Research no longer lived only in universities. Training large systems needed data, hardware, teams and money on a new scale. Organisations took different paths, and a company name should never be mistaken for a single inventor.

OpenAI began in 2015 with founders and early leaders including Sam Altman, Greg Brockman, Elon Musk and research leader Ilya Sutskever, alongside other researchers and engineers. Its GPT work, ChatGPT, image tools, later reasoning models and tool-using agents helped carry laboratory advances into widely used products. By 8 October 2026, its published GPT-6 series included Astra, Sol and Luna, with GPT-6.1 Sol added on 29 September. The company's performance descriptions are its claims; the existence, dates and published model offerings are verifiable facts.

Google's Brain team helped scale neural-network research inside Google. The separate UK-founded DeepMind joined Google in 2014, and the two research groups were brought together as Google DeepMind in 2023. Google researchers produced the 2017 transformer paper; DeepMind contributed AlphaGo and AlphaFold. Gemini became the company's multimodal model family, with 2026 releases extending agent, audio and visual work. Google's infrastructure also includes custom AI chips called TPUs.

Anthropic was founded in 2021 by former OpenAI staff, including Dario and Daniela Amodei. Its Claude family and 'Constitutional AI' approach - using stated principles in part of model training - made it a central participant in model development and safety work. It has also published studies of model internals and computer use. Its October 2026 documentation listed the Claude Opus, Sonnet, Haiku and Fable lines. Those names describe product families, not separate discoveries of intelligence.



The organisations shaping modern AI · continued

Meta established its FAIR research lab in 2013, with Yann LeCun in a leading scientific role. Its Llama releases, beginning in 2023, made many models' weights available for others to run or adapt under stated licences. 'Open-weight' means access to learned parameters; it does not automatically reveal every training example or make a licence unrestricted. Meta has also built multimodal and agent-oriented systems, including its 2026 Muse Spark line.

Microsoft's research history reaches back to 1991. Its partnership with OpenAI began in 2019, and Copilot brought AI into coding and office software. In September 2026 Microsoft described an Autopilot direction with an identity, memory and workspace for ongoing enterprise tasks. NVIDIA, founded as a graphics-chip company, became indispensable to many AI projects because its GPUs and software tools made large-scale training and serving practical. The model makers and the hardware makers shaped the era together.

Elon Musk launched xAI in 2023. Its Grok line and Colossus computing infrastructure became part of the frontier-model race; Grok 4.7 was announced in September 2026. In France, Mistral AI, founded in 2023 by Arthur Mensch, Guillaume Lample and Timothée Lacroix, developed open-weight and commercial models and became a prominent European option. Other companies - from Amazon and IBM to smaller specialised labs - also shaped the ecosystem. There is no single corporate family tree for modern AI.

Organisations at a glance

Company names are useful landmarks, but each organisation contains many researchers, engineers, products and research traditions. These associations are not a ranking.



The organisations shaping modern AI · continued

OpenAI	GPT, ChatGPT, reasoning models, tool-using agents	US
Google / DeepMind	Transformer lineage, Gemini, AlphaGo, AlphaFold, robotics	US / UK lineage
Anthropic	Claude, Constitutional AI, interpretability and safety research	US
Meta	FAIR, Llama open-weight ecosystem, multimodal and agent work	US
Microsoft	Long AI research history, OpenAI partnership, Copilot, persistent agents	US
NVIDIA	GPUs, CUDA and the compute ecosystem behind large-scale AI	US
xAI	Grok and large-scale compute infrastructure	US
Mistral AI	European open-weight and commercial frontier models	France

A founder can create a laboratory, an executive can shape a product, and a scientist can invent a method. These are different contributions. Modern AI is built by large teams and long research lineages, not by one celebrity name.

Sources: [23] · [40] · [41] · [42] · [43] · [44] · [45] · [46] · [47] · [48] · [49] · [50] · [85] · [86] · [103] · [104] · [105]



CHAPTER 15

The rise of AI in China

China's modern AI ecosystem has roots well before the recent global excitement over chatbots. Universities, companies and public agencies invested in machine learning, speech, vision, translation and robotics through the 2000s and 2010s. Baidu built early deep-learning capability and later the ERNIE family. Alibaba expanded cloud and research work that grew into Qwen. Tencent, Huawei and ByteDance developed research, products or infrastructure of their own. China's 2017 national AI plan made long-term research and application a declared priority.

In the 2020s, an energetic newer generation broadened the field. DeepSeek drew international attention with its R1 reasoning model in 2025 and later V4 models. Zhipu AI, marketed internationally as Z.ai, developed the GLM line; Moonshot AI built Kimi; MiniMax built language, voice, music and Hailuo video systems. ByteDance's Doubao products draw on Seed models, while Seedance expanded video creation. Tencent developed Hunyuan, now also called Hy, and Baidu kept developing ERNIE. Kuaishou's Kling, launched in 2024, became a visible example of Chinese work in generative video.

Many Chinese labs publish open-weight models. That can let researchers elsewhere examine or adapt a system, subject to the licence and the information actually released. At the same time, restrictions on access to advanced chips have driven investment in domestic hardware and more efficient training. Huawei's Ascend chips are part of that effort. Fierce competition among labs, a large domestic market and rapid product adoption have all mattered.

Stanford's 2026 AI Index reports that China leads in several research-output measures and trails the US in some measures of notable frontier models and higher-impact patents; simple rankings miss the variety on both sides. China is one of the world's largest AI ecosystems, with its own research traditions, companies, creators and constraints. Its progress did not start when one model caught global attention.

The 2025-26 visibility of DeepSeek and other models was an inflection in public awareness, not the beginning of Chinese AI research.

Sources: [51] · [52] · [53] · [54] · [55] · [56] · [57] · [58] · [59] · [60] · [61] · [62] · [100]



CHAPTER 16

AI beyond the US and China

The international story has many centres. Britain contributed Babbage, Lovelace and Turing long before the name AI existed; later it gave the field DeepMind and a dense network of universities and startups. France's Mistral is a newer example of European model building. Across Europe, researchers work on language, robotics, vision and public-interest questions, while the European Union shapes governance. Scientific collaboration often crosses these borders.

Canada was especially important when neural networks were unpopular elsewhere. Geoffrey Hinton's Toronto group and Yoshua Bengio's Montréal community trained researchers, carried forward methods and helped turn deep learning into a global field. Alberta became an important centre for reinforcement learning, including Richard Sutton's work. It is hard to tell the 2012 deep-learning story without Canada.

Israel's AI21 Labs developed language-model approaches including Jamba. Japan combines a long robotics tradition with new labs such as Tokyo-based Sakana AI. South Korea's NAVER and other groups invested in Korean-language and industrial AI. India's Sarvam and public initiatives have emphasised models and speech tools for Indian languages. In the Gulf, the UAE's Technology Innovation Institute released Falcon models, and governments and firms are investing in computing and regional-language systems. Russia has work such as Yandex's Alice/YandexGPT and Sber's GigaChat, alongside substantial hardware and international-access constraints.

These examples are selective. Research from Latin America, Africa, Southeast Asia and many other places also matters, as do the people whose work and data support the field. An AI history told only through two countries would miss both the earlier foundations and the present network of contributors.

Sources: [19] · [63] · [64] · [65] · [66] · [67] · [94] · [95]



CHAPTER 17

Images, video, music and creative AI

AI's creative turn did not begin with a text chatbot. GANs offered one way to make images; later diffusion models learned to transform noisy patterns toward coherent pictures. OpenAI's DALL-E showed text-to-image possibilities in 2021. In 2022, DALL-E 2, Midjourney and the publicly released Stable Diffusion helped take image making beyond research labs. Stable Diffusion's release also enabled a large ecosystem of people building and adapting tools. Black Forest Labs introduced FLUX.1 in 2024. Their methods and access models differ, but each changed what a person could attempt with a written prompt or reference image.

Video added motion, continuity and often sound - a much harder combination. Runway's Gen-2 appeared in 2023; Google's Veo, OpenAI's Sora and Kuaishou's Kling became prominent in 2024. MiniMax's Hailuo and ByteDance's Seedance made Chinese labs especially visible in this area. A striking demonstration should be read as a demonstration: consistency, control and useful editing can vary across a whole project. Products also change; for example, OpenAI ended the standalone Sora product in April 2026, though Sora remains historically important to this phase.

Music and voice travelled on related but distinct paths. Suno and Udio let users generate songs from descriptions during the 2023-24 wave; ElevenLabs made synthetic speech and voice tools widely accessible from 2023. In language generation a model produces sequences of words or tokens. In image generation it produces or transforms spatial visual patterns. Video must also maintain relationships across frames, and audio must handle timing, pitch, timbre or speech. These distinctions explain why a new capability in one medium does not automatically solve the others.

For artists, the history is also a story of new choices and new disputes: control, consent, credit, licensing and the relationship between a prompt, a training collection and an original work. Those questions are still being worked out.

Sources: [30] · [62] · [68] · [69] · [70] · [71] · [72] · [73] · [74] · [75] · [76] · [77] · [99] · [101]



CHAPTER 18

AI becomes multimodal

A person can connect a spoken question, a photograph, a diagram and the scene outside a window. For a long time, AI products were divided into narrower channels: one model for text, another for speech, another for images. A multimodal system can work with combinations of these kinds of input, sometimes producing more than one kind of output. In a useful conversation, the model might look at an image, hear a question about it and answer aloud.

This shift came in stages. Image-and-text models preceded the public chatbot wave. GPT-4 introduced image input in 2023; Google's Gemini 1.0 was presented as multimodal that year.

OpenAI's GPT-4o in 2024 connected text, audio, image and video input in a more immediate interaction. Anthropic's Claude 3 family brought image understanding into Claude. By 2026, models could also inspect longer video or work through computer interfaces, though exact abilities still depended on the product, access and tool design.

Multimodality matters because much of the world is not written in paragraphs. A system asked to help with a chart, speak with a person, observe a robot camera or navigate a website needs to relate different forms of evidence. Joining them can be useful, but it also creates new chances to misread a scene, mishear a speaker or act on incomplete information.

Sources: [78] · [79] · [80] · [81]



CHAPTER 19

From chatbots to agents

A chatbot usually returns an answer. An AI agent is a system set up to pursue a task over several steps, often observing results and choosing what to do next. The model is only one part. Instructions state the task; memory or saved state keeps useful information; tools let the system search, code or use software; permissions limit what it may change; and an environment gives it a place to act. Remove or alter any of these and the same underlying model can behave very differently.

Web research agents gather and compare sources. Coding agents inspect a repository, edit files and run checks. Computer-use agents work through buttons and pages. Persistent agents can return to a task or keep a workspace and identity over time. Multi-agent systems coordinate more than one specialised agent. In 2024-26, OpenAI, Anthropic, Google, Microsoft, Meta and Chinese labs all demonstrated pieces of this transition. Products are evolving faster than a printed roster can capture.

Why does this change safety? A text answer can mislead. An action can also send a message, modify data, spend money or move a robot. Tool permissions, confirmation steps, monitoring and the ability to stop or undo actions therefore become central. A company demo of a long autonomous run is a claim about a specified setting, not proof of reliable independence in every setting. The agent era extends an old AI ambition - machines that can act on goals - into systems people can actually use and must govern.



From chatbots to agents · continued

Agent = model + instructions + saved state or memory + tools + permissions + an environment. Capability and risk depend on the whole system.

From model to chatbot, agent and robot

As systems gain memory, tools, permissions and bodies, the unit we need to understand gets larger than the model alone.

A model produces predictions or outputs from supplied input. The central question is what the network has learned. A chatbot wraps the model in conversation, instructions and product rules, changing what the user actually experiences.

An agent adds loops, tools, memory or state, permissions and an environment. The question expands to what the whole system can now do. A robot or embodied agent adds sensors, physical actions, materials and contact with the world. Now we must ask what happens when mistakes become physical.

A text model can mislead. An agent can also send, edit, buy or execute. A robot can move through the physical world. More capability means we must analyse tools, permissions, monitoring and stopping conditions - not only the model's words.

Sources: [46] · [53] · [58] · [82] · [83] · [84] · [85] · [86] · [78] · [79] · [80] · [81] · [87] · [88] · [89] · [90] · [91] · [92]



CHAPTER 20

Robotics and embodied AI

Robotics has its own long history, from industrial arms to mobile machines. What is new is the closer connection to general-purpose learned models. A vision-language-action model links what a robot sees, what a person says and what the robot should try to do. Google's RT-2 in 2023 and later Gemini Robotics explored that bridge. Physical Intelligence's pi-zero family and Figure's Helix work followed different approaches to broader learned robot behaviour.

Humanoid projects attract attention because a human-shaped body may fit spaces built for people. Tesla's Optimus, Figure's robots and Unitree's machines are prominent examples; Boston Dynamics' Atlas brings a long robotics lineage into newer AI work. Chinese robotics companies are increasingly visible in hardware and learned control. A demonstration of folding laundry or walking through a room is useful evidence of a capability under those conditions. It is not evidence that the robot can safely handle every unfamiliar home, workplace or person.

Researchers also build 'world models': learned systems that try to represent how a situation changes after an action. Embodied intelligence is the idea that intelligent behaviour can depend on a body, its sensors, its materials and its contact with the world, not only an abstract calculation in a central controller. AI and robotics are converging quickly, but physical reliability, dexterity, energy use and safety remain difficult. A language model's fluent answer and a robot's safe action are different standards of success.

A robot learning a skill in a demonstration is a milestone. General physical competence requires robust behaviour across changing real places.

Sources: [87] · [88] · [89] · [90] · [91] · [92] · [102]

**CHAPTER 21**

The people who shaped AI

One way to remember this history is as a conversation between generations. Charles Babbage, the nineteenth-century inventor, designed a general-purpose calculating engine; Ada Lovelace, the mathematician and writer, saw how programmed symbols might reach beyond arithmetic. George Boole supplied a mathematical language for logic. Alan Turing clarified computation and reframed the question of machine intelligence as a testable conversation.

Claude Shannon, a mathematician and engineer, helped found information theory. Norbert Wiener made feedback central to cybernetics. Warren McCulloch and Walter Pitts modelled simple artificial neurons. John McCarthy named AI and organised its 1956 gathering; Marvin Minsky helped establish laboratories and debated the field's directions. Nathaniel Rochester helped bring the Dartmouth proposal together. Allen Newell, Herbert Simon and programmer Cliff Shaw made Logic Theorist and other early problem-solving programs.

Frank Rosenblatt built the perceptron; Seymour Papert and Minsky analysed its limits. Joseph Weizenbaum made ELIZA and warned about what people might read into it; Terry Winograd made SHRDLU. Edward Shortliffe, Bruce Buchanan and colleagues built MYCIN. Arthur Samuel explored programs that learned at checkers. These contributors did different things - founding a field, proving a limitation, writing a program or questioning its meaning.



The people who shaped AI · continued

David Rumelhart, Geoffrey Hinton and Ronald Williams made backpropagation newly influential. Yann LeCun advanced image-reading networks; Yoshua Bengio advanced deep representation learning. Sepp Hochreiter and Jürgen Schmidhuber developed LSTM. Richard Sutton and Andrew Barto laid foundations for reinforcement learning. Fei-Fei Li and the ImageNet team created a shared vision resource; Alex Krizhevsky and Ilya Sutskever, with Hinton, made the 2012 AlexNet result.

Ian Goodfellow and collaborators introduced GANs. Demis Hassabis, Shane Legg and Mustafa Suleyman founded DeepMind; its scientific successes were collective efforts, with John Jumper central to AlphaFold. Ashish Vaswani and seven fellow Google researchers introduced the transformer. Researchers such as Dzmitry Bahdanau, Kyunghyun Cho, Oriol Vinyals and Quoc Le helped shape the earlier language path. Their work did not appear from nowhere; neither did the GPT models built by large teams afterward.

A founder may build a laboratory and a leader may help turn research into a product. Those contributions matter. They are not identical to discovering a learning method or writing the paper that describes it. The story becomes clearer when we honour each kind of work by its real name.

Sources: [2] · [4] · [8] · [9] · [14] · [15] · [19] · [22] · [25] · [27] · [29] · [32] · [40]



CHAPTER 22

Major model families to recognise in 2026

The names below are families or lines, not a ranking. A model family can include several sizes and frequent updates. 'Known for' identifies a useful historical association, not a promise that every version has a given ability.

OpenAI's GPT family (United States) grew from transformer language research into ChatGPT, reasoning models and tools; the GPT-6 series includes Astra, Sol and Luna, alongside GPT-6.1 Sol. Anthropic's Claude family (United States) includes Opus 5.5, Sonnet 5.5, Haiku 5.5 and Fable 5.1 lines as of 8 October 2026 and is associated with assistant design, safety research and tool use. Google's Gemini (United States, with the UK DeepMind lineage) is a multimodal family extending across consumer products, research and agents.

Meta's Llama (United States) became influential for releasing model weights and supporting an ecosystem of adaptations; Meta's Muse Spark is a newer, distinct line. xAI's Grok (United States) developed alongside X integration and large compute investment. France's Mistral family includes open-weight and hosted models such as Mistral Large, Medium and Small, alongside the specialised Magistral and Devstral lines.

In China, Alibaba's Qwen spans language and multimodal tools; DeepSeek's R1 and V4 lines made reasoning and efficiency central to its reputation; Z.ai's GLM serves language, coding and agent work; Moonshot's Kimi developed long-context and agent-oriented uses; and MiniMax's M line sits beside its Hailuo video, speech and music systems. Baidu's ERNIE, Tencent's Hunyuan/Hy and ByteDance's Seed/Doubao also matter to the national and global picture.

A creative map would add DALL-E and OpenAI's image models, FLUX, Stable Diffusion, Midjourney, Runway, Veo, Kling, Hailuo and Seedance for visual media; Suno, Udio and MiniMax Music for songs; and ElevenLabs for speech. These names point to different media, licences and workflows. The relevant question is what a particular version can do in a particular setting - and what evidence supports that answer.

Sources: [41] · [44] · [45] · [47] · [48] · [50] · [53] · [54] · [55] · [56] · [57] · [58] · [59] · [60] · [61] · [62] · [68] · [69] · [70] · [71] · [72] · [73] · [74] · [75] · [76] · [77] · [103] · [104] · [105] · [46]



CHAPTER 23

Where we are in October 2026

By 8 October 2026, AI is an expanding collection of systems: language models, models that spend more computation on difficult reasoning, image and video generators, speech and music tools, scientific models, multimodal assistants, agents that use software, and robots learning to act. Some products save memory outside the model. Some agents call tools or coordinate with other agents. Each addition changes what the whole system can do, and also what its builders must test.

The final days covered by this edition show how quickly the model landscape changes. OpenAI published GPT-6.1 Sol on 29 September and an October update to GPT-6 Sol and GPT-6 Luna on 7 October. Its release record distinguishes the new ChatGPT versions from the September versions then used in Codex and ChatGPT Work. Anthropic released Claude Sonnet 5.5 on 28 September and Haiku 5.5 on 7 October. These are dated releases, rather than evidence that every version is equally capable in every environment.

There is more than one live scientific question. Researchers can measure and alter some learned internal features, but a complete human-readable account of all learned representations is still out of reach. Agents can succeed at long tasks, yet their reliability and controllability across unfamiliar environments remain unsettled. Larger models and more training have helped repeatedly, but no one can state as a fact how far scaling will continue or whether current architectures are sufficient for much more general intelligence.

Nor does the evidence settle whether an AI system has subjective experience, or whether one could have it. Researchers disagree on how even to test that question. We can describe a system's architecture, behaviour and training evidence without claiming access to an inner life. Safety and governance research now runs alongside capability research because systems that persuade, create and take actions affect real people. The field is extraordinary precisely because it is still unfinished.

What will make increasingly capable systems reliable, understandable and accountable in the places where people actually use them?

Sources: [51] · [63] · [78] · [82] · [83] · [87] · [93] · [103] · [104] · [105]

**THE BIG PICTURE**

How all these pieces became AI

Artificial intelligence was not invented in one moment. It began as an old human question about minds and machines. It acquired the tools of logic and computation, a name at Dartmouth, and a series of daring experiments. Some approaches flourished, disappointed their supporters, then returned in forms their first builders could hardly have imagined.

The story passes through the persistence of neural-network researchers, the patient collection of data, the invention of new chips and methods, the transformer team's work, and the effort to make models useful in conversation. It also passes through laboratories on many continents, through people who challenged exaggerated claims, and through users who discovered what these systems could and could not do in daily life.

Today an AI can help write, draw, speak, search, solve scientific problems or guide a robot through a small part of the world. Each ability has its own history and its own limits. The remarkable fact is the continuity: people kept asking whether a machine could perform something that seemed to belong to minds, kept finding partial answers, and kept discovering that the next question was deeper than the last.



A JOURNEY THROUGH TIME

The AI history timeline

A selective route through ideas, papers and widely recognised releases. The field never moved in a perfectly straight line.

Ancient to medieval	Artificial beings in stories; skilled makers create mechanical automata, including al-Jazari's designs.
1830s-1843	Babbage designs the Analytical Engine; Lovelace publishes her expansive notes and algorithm.
1854	Boole's Laws of Thought advances algebraic logic.
1936	Turing clarifies what an abstract computing procedure can do.
1943	McCulloch and Pitts model networks of simplified artificial neurons.
1948-1950	Shannon's information theory; Wiener's cybernetics; Turing's imitation-game paper.
1955-1956	McCarthy names AI in the Dartmouth proposal; the summer workshop helps establish a field. Logic Theorist appears.
1957-1958	Rosenblatt's perceptron gives an early learning system broad visibility.
1960s-early 1970s	ELIZA and SHRDLU show narrow but memorable language interactions; expert systems emerge.
1970s	Overpromises and technical limits contribute to the period commonly called the first AI winter.
1980s	Expert systems spread in business; researchers continue neural-network and statistical work.
1986	Rumelhart, Hinton and Williams make multilayer backpropagation newly influential.



The AI history timeline · continued

Late 1980s-early 1990s	A second downturn follows expert-system and specialised-hardware disappointment.
1997	Deep Blue defeats Kasparov; Hochreiter and Schmidhuber publish LSTM.
2006-2009	CUDA expands GPU computing; ImageNet supplies a large labelled image resource.
2010-2012	DeepMind is founded; AlexNet's 2012 ImageNet result galvanises deep learning.
2014-2016	GANs and sequence-to-sequence language work advance; AlphaGo defeats Lee Sedol.
2017	The Google transformer paper offers a new scalable architecture for sequences.
2018-2020	GPT, GPT-2 and GPT-3 reveal the promise and limits of ever larger language models.
2021-2022	DALL-E and diffusion images expand creation; InstructGPT precedes ChatGPT's public launch.
2023-2024	Llama and Gemini arrive; image, video and audio tools spread; multimodal assistants and agents advance.
2025-Sept 2026	Reasoning models, software agents, scientific systems and learned robotics grow; global labs release new families.
28 September-7 October 2026	Claude Sonnet 5.5, GPT-6.1 Sol, Claude Haiku 5.5 and the October GPT-6 Sol/Luna update extend the model landscape.



WORDS BEHIND THE STORY

A short glossary

These terms help the historical story fall into place. Their meaning in practice depends on the system being discussed.

Artificial intelligence (AI)	The broad field of making systems perform tasks associated with intelligence, such as recognising patterns, reasoning or using language.
Algorithm	A defined procedure or method for working through a task.
Symbolic AI	AI that reasons with explicitly represented things, facts and rules.
Expert system	A program using stored specialist knowledge and rules in a narrow field.
Machine learning	Methods that improve a system's performance by learning patterns from data or experience.
Neural network	A connected arrangement of mathematical units with adjustable numbers learned during training.
Deep learning	Machine learning with multilayer networks that can build increasingly useful features.
Parameter / weight	An adjustable number in a learned model; many act together to shape its behaviour.
Dataset	A collection of examples used for training or evaluation.
Benchmark	A common, specified test used to compare systems; it is narrower than real-world ability.



A short glossary · continued

GPU	A processor good at many similar calculations in parallel, useful for training large networks.
Reinforcement learning	Learning by trying actions and using feedback about their consequences.
Generative AI	Systems that produce new text, images, sound, video or other outputs from learned patterns.
Large language model (LLM)	A large trained model that works with sequences of language or related tokens.
Attention / transformer	Attention relates relevant parts of input; a transformer is an architecture that uses it extensively.
RLHF	Reinforcement learning from human feedback: training that uses human preferences to shape responses.
Multimodal	Able to work across more than one type of input or output, such as text, image and audio.
Agent	A system arranged to pursue a task over steps, often using a model, tools, memory and permissions.
Vision-language-action	A robotics approach linking what a machine sees and is told with actions it can take.
Open-weight	The model's learned numbers are released under a licence; other training details may remain private.



THE RECORD BEHIND THE HISTORY

Sources and further reading

The chapter source numbers lead to the records below. Titles are clickable. Papers and institutional records document the earlier history; company announcements establish release dates and describe company claims. New October sources were checked on 8 October 2026.

[1] The Met, al-Jazari's mechanical devices.

metmuseum.org

[2] Computer History Museum, Babbage and Ada Lovelace.

computerhistory.org

[3] MacTutor, George Boole and the 1854 Laws of Thought.

mathshistory.st-andrews.ac.uk

[4] Turing, On Computable Numbers (1936).

doi.org

[5] Turing, Computing Machinery and Intelligence (1950).

academic.oup.com

[6] MIT Press, Wiener's Cybernetics (1948).

direct.mit.edu

[7] McCulloch and Pitts, A Logical Calculus (1943).

doi.org

[8] Dartmouth, the 1955 proposal and 1956 workshop.

ai.dartmouth.edu

[9] Computer History Museum, AI & Robotics timeline.

computerhistory.org

[10] MIT CSAIL, films of early AI projects.

projects.csail.mit.edu

[11] Weizenbaum, ELIZA (1966), MIT Press collection.

direct.mit.edu

[12] Minsky and Papert, Perceptrons (1969).

mitpress.mit.edu

[13] Rosenblatt, The Perceptron (1958), MIT Press.

direct.mit.edu

[14] Rumelhart, Hinton and Williams, backpropagation (1986).

nature.com



Sources and further reading · continued

[15] Hochreiter and Schmidhuber, Long Short-Term Memory (1997).

direct.mit.edu

[16] Thomas Haigh, There Was No 'First AI Winter'.

doi.org

[17] Computer History Museum, expert-systems pioneers.

computerhistory.org

[18] John McCarthy, Stanford, on MYCIN and common sense.

www-formal.stanford.edu

[19] ACM, Turing Award to Sutton and Barto.

awards.acm.org

[20] IBM history, Deep Blue.

ibm.com

[21] Stanford Vision Lab, ImageNet paper (2009).

vision.stanford.edu

[22] Krizhevsky, Sutskever and Hinton, AlexNet (2012).

proceedings.neurips.cc

[23] NVIDIA, CUDA's introduction and history.

docs.nvidia.com

[24] ACM, 2018 Turing Award for deep learning.

awards.acm.org

[25] DeepMind, co-founders and early years.

deepmind.google

[26] Mnih and colleagues, Atari deep reinforcement learning.

nature.com

[27] Nobel Prize, AlphaFold and the 2024 Chemistry Prize.

nobelprize.org

[28] Mikolov and colleagues, word representations (2013).

arxiv.org

[29] Sutskever, Vinyals and Le, sequence-to-sequence (2014).

arxiv.org

[30] Goodfellow and colleagues, GANs (2014).

arxiv.org



Sources and further reading · continued

[31] Bahdanau, Cho and Bengio, neural attention.

arxiv.org

[32] Vaswani and colleagues, Attention Is All You Need.

proceedings.neurips.cc

[33] Google Research, the transformer research lineage.

blog.google

[34] OpenAI, first GPT research (2018).

openai.com

[35] OpenAI, GPT-2 (2019).

openai.com

[36] OpenAI, GPT-3 and few-shot learning (2020).

openai.com

[37] OpenAI, InstructGPT and RLHF (2022).

openai.com

[38] OpenAI, introducing ChatGPT (30 November 2022).

openai.com

[39] Kaplan and colleagues, scaling laws (2020).

arxiv.org

[40] OpenAI, founding announcement (2015).

openai.com

[41] OpenAI API, GPT-6 Astra, Sol and Luna guidance.

developers.openai.com

[42] Google, Brain and DeepMind combine (2023).

blog.google

[43] Google, introducing Gemini (2023).

blog.google

[44] Anthropic, 2021 founding-stage research announcement.

anthropic.com

[45] Anthropic, Constitutional AI paper and principles.

anthropic.com



Sources and further reading · continued

[46] Anthropic, current Claude model documentation.

platform.claude.com

[47] Meta, ten years of FAIR.

ai.meta.com

[48] Meta, Llama 4 and multimodal open weights.

ai.meta.com

[49] xAI, Grok 4.7 release (September 2026).

x.ai

[50] Mistral AI, founding and early open-model goals.

mistral.ai

[51] Stanford HAI, AI Index 2026.

hai.stanford.edu

[52] China State Council, national AI plan (2017).

app.www.gov.cn

[53] Alibaba Cloud, Qwen and September 2026 roadmap.

alibabacloud.com

[54] DeepSeek, R1 original release (2025).

deepseek.com

[55] Z.ai, GLM-5.3 research release.

z.ai

[56] Moonshot AI, Kimi model chronology.

moonshot.ai

[57] MiniMax, multimodal research and models.

minimax.io

[58] ByteDance Seed, 2026 model portfolio.

seed.bytedance.com

[59] Tencent, Hunyuan/Hy3 model update.

tencent.com

[60] Baidu, ERNIE 5.0 technical report.

ernie.baidu.com



Sources and further reading · continued

[61] Huawei, Ascend and Pangu strategy (2026).

huawei.com

[62] Kuaishou, Kling's June 2024 launch.

ir.kuaishou.com

[63] ACM, Hinton/LeCun/Bengio and Canada deep learning.

sigai.acm.org

[64] AI21 Labs, Jamba research (Israel).

ai21.com

[65] Sakana AI, Tokyo research lab.

sakana.ai

[66] NAVER, HyperCLOVA X (South Korea).

navercorp.com

[67] Sarvam AI and India's language-model work.

sarvam.ai

[68] OpenAI, DALL-E images from text (2021).

openai.com

[69] Stability AI, Stable Diffusion public release.

stability.ai

[70] Black Forest Labs, FLUX.1 launch (2024).

bfl.ai

[71] Runway Research, Gen-2 video (2023).

runwayml.com

[72] Google DeepMind, Veo video model family.

deepmind.google

[73] OpenAI, Sora history and product status.

openai.com

[74] Suno, music generation milestone (2024).

about.suno.com

[75] Udio, music-generation tools (2024).

udio.com



Sources and further reading · continued

[76] ElevenLabs, public speech platform (2023).

elevenlabs.io

[77] Kuaishou, Kling 3.0 update (2026).

ir.kuaishou.com

[78] OpenAI, GPT-4o multimodal model (2024).

openai.com

[79] Anthropic, Claude 3 and vision (2024).

anthropic.com

[80] Google, Gemini 1.0 multimodal release.

blog.google

[81] Google, agentic video understanding (2026).

blog.google

[82] OpenAI, ChatGPT agent and tool use.

openai.com

[83] Anthropic, developing computer use.

anthropic.com

[84] Google, Gemini research agent update.

blog.google

[85] Microsoft, persistent Autopilot (September 2026).

blogs.microsoft.com

[86] Meta, Muse Spark multimodal and agent systems.

ai.meta.com

[87] Google DeepMind, RT-2 vision-language-action.

deepmind.google

[88] Google DeepMind, Gemini Robotics.

deepmind.google

[89] Physical Intelligence, pi-zero robot model.

pi.website

[90] Figure, Helix vision-language-action.

figure.ai



Sources and further reading · continued

[91] Unitree Robotics, G1 humanoid overview.

unitree.com

[92] Boston Dynamics, Atlas and DeepMind partnership.

bostondynamics.com

[93] Anthropic, mapping features inside a language model.

anthropic.com

[94] UAE Technology Innovation Institute, Falcon models.

tii.ae

[95] Reuters, Russia's YandexGPT and GigaChat (2026).

reuters.com

[96] Shannon, A Mathematical Theory of Communication (1948).

onlinelibrary.wiley.com

[97] University of Edinburgh, Lighthill report and 1973 debate.

aiai.ed.ac.uk

[98] Google DeepMind, AlphaGo and Lee Sedol.

deepmind.google

[99] Ho, Jain and Abbeel, diffusion models (2020).

arxiv.org

[100] DeepSeek, V4 release and model updates (2026).

deepseek.com

[101] Discord, Midjourney beta widened in July 2022.

discord.com

[102] Tesla, Optimus robotics programme.

tesla.com

[103] OpenAI: GPT-6.1 Sol system-card addendum (29 September 2026)

deploymentsafety.openai.com

[104] OpenAI: GPT-6 Sol and GPT-6 Luna, October update (7 October 2026)

deploymentsafety.openai.com

[105] Anthropic: Claude Platform release notes (checked 8 October 2026)

platform.claude.com