NO PREP

A Short History of AI

From Turing to ChatGPT

Includes: Student reading, case study, differentiated questions, teacher key

Name: Period: Date:

A Short History of AI: From Turing to ChatGPT

A Question That Started Everything

In 1950, a British mathematician named Alan Turing published a paper with a deceptively simple question at its center: "Can machines think?" Turing knew the question was philosophically slippery, so he proposed a practical test instead. Imagine you are exchanging written messages with two parties — one human, one machine. If you cannot reliably tell which is which based on the conversation alone, the machine has demonstrated something functionally equivalent to thinking. This thought experiment, now called the Turing Test, became one of the founding ideas of the entire field of artificial intelligence. Turing never used the phrase "artificial intelligence" himself — that term came a few years later. But his question set the agenda for generations of researchers who would spend decades trying to build machines that could learn, reason,

Name: **Period:** **Date:**

A Short History of AI: From Turing to ChatGPT

A Question That Started Everything

In 1950, a British mathematician named Alan Turing published a paper with a deceptively simple question at its center: "Can machines think?" Turing knew the question was philosophically slippery, so he proposed a practical test instead. Imagine you are exchanging written messages with two parties: one human, one machine. If you cannot reliably tell which is which based on the conversation alone, the machine has demonstrated something functionally equivalent to thinking. This thought experiment, now called the Turing Test, became one of the founding ideas of the entire field of artificial intelligence. Turing never used the phrase "artificial intelligence" himself. That term came a few years later. But his question set the agenda for generations of researchers who would spend decades trying to build machines that could learn, reason, and communicate like humans. The path from Turing's question to the AI tools you use today is a story of explosive optimism, crushing setbacks, and finally, in the last decade, a transformation that no one fully predicted. Understanding that history is not just interesting trivia; it explains why AI works the way it does today and why it has the limitations it does.

The Birth of the Field: 1950s-1960s

In the summer of 1956, a small group of scientists gathered at Dartmouth College in New Hampshire for a workshop that became a landmark moment in computing history. The organizer, a mathematician named John McCarthy, coined the term "artificial intelligence" for the event. The participants (who included Marvin Minsky, Claude Shannon, and Nathaniel Rochester) were optimistic to the point of audacity. McCarthy's original proposal suggested that a team working for just two months could make significant progress on teaching machines to use language, form abstractions, and solve problems reserved for humans. They did not solve it in two

months. But the Dartmouth workshop established AI as a serious academic discipline and launched the first wave of serious research. In the late 1950s and 1960s, researchers built early programs that could play checkers, solve algebra problems, and prove mathematical theorems. These systems were impressive given the hardware constraints of the era, and they generated enormous excitement. Government funding poured in, particularly from the U.S. Department of Defense, which saw strategic value in machines that could process intelligence reports and support military decision-making.

The AI Winters: 1970s and 1980s

The excitement of the early decades eventually collided with reality. Researchers discovered that the tasks they had assumed were simple (understanding spoken language, recognizing a face, navigating a room) were in fact extraordinarily difficult for machines. Meanwhile, the tasks that seemed hard (playing chess, proving theorems) turned out to be tractable through brute-force search. This gap between public expectation and actual capability led to two major periods of reduced funding and interest, known as the "AI winters." The first AI winter began in the mid-1970s, when a British government report (the Lighthill Report) concluded that AI had failed to deliver on its promises and recommended cutting funding. The second AI winter hit in the late 1980s after a wave of investment in so-called "expert systems" (programs that encoded human expertise as if-then rules) proved brittle and expensive to maintain. These systems worked well in narrow, controlled environments but collapsed when confronted with situations their designers had not anticipated. The AI field's credibility took serious damage during these years.

The Machine Learning Revolution: 1990s-2000s

The turn back toward optimism began quietly in the 1990s, driven by a shift in approach. Instead of trying to hand-code intelligence — programming every rule explicitly — a new generation of researchers focused on machine learning: giving computers access to massive amounts of data and letting them figure out the patterns themselves. This was a philosophical revolution as much as a technical one. The goal shifted from programming intelligence to training it. A series of landmark achievements in the 1990s and 2000s demonstrated the power of this new approach. IBM's Deep Blue defeated world chess champion Garry Kasparov in 1997, shocking the world. IBM's Watson defeated Jeopardy! champions in 2011, showing that

machines could handle natural language in a structured context. In 2012, a team at the University of Toronto used a type of machine learning called deep learning to achieve a dramatic improvement in image recognition, a result that would reshape the entire field. These milestones attracted new investment, new talent, and new ideas.

The Modern Era: 2017 to Today

The breakthrough that made today's AI tools possible came in 2017, when researchers at Google published a paper titled "Attention Is All You Need." It introduced the transformer architecture — a new way of building neural networks that turned out to be extraordinarily effective at understanding and generating language. The transformer became the foundation for a new generation of language models, including GPT (Generative Pre-trained Transformer), which OpenAI began developing starting in 2018. The public moment that changed everything came in November 2022, when OpenAI released ChatGPT. The tool reached one million users in five days — faster than any technology product in history at that point. Within months, Google released Bard (later renamed Gemini), Anthropic released Claude, and Microsoft integrated AI into its products. By 2023 and 2024, AI had moved from a specialized research topic to a daily-use tool for hundreds of millions of people. The field had come a long way from Alan Turing's thought experiment in 1950. But in some sense, the field is still grappling with exactly his original question — and still has not fully answered it. One underappreciated dimension of AI history is the role of failure. The AI winters were not simply pauses in progress — they were periods when the community was forced to confront fundamental assumptions about what intelligence requires and how it can be engineered. The failures of expert systems taught researchers that intelligence could not be fully captured in rules a human expert could articulate. The failures of early neural networks taught researchers that architecture matters as much as data. These hard lessons shaped the approaches that eventually succeeded. Understanding this pattern of failure and recovery helps explain why current AI researchers take a longer view of the technology's trajectory than the daily news cycle encourages. It is worth pausing to appreciate how different the trajectory of AI research looks in retrospect versus how it felt to the people living through it. Each generation of researchers genuinely believed they were close to cracking the fundamental problem of machine

intelligence. Each generation was surprised by how much harder the remaining challenges turned out to be. The researchers who built the first chess programs in the 1960s thought language understanding would follow within a decade. It took sixty years. Researchers who celebrated the expert system boom of the 1980s did not foresee that the approach would collapse under its own complexity. Even the founders of deep learning did not predict how dramatically the scaling of transformer models would transform the field between 2017 and 2023.

Key Information: History of AI

Milestone	Year / Details
Turing Test proposed	1950 — Alan Turing asks whether machines can think
Dartmouth Workshop	1956 — John McCarthy coins 'artificial intelligence'
First AI Winter	Mid-1970s — funding cut after overpromised results
Second AI Winter	Late 1980s — expert systems proved brittle and limited
Deep Blue defeats Kasparov	1997 — IBM chess AI beats world chess champion
Deep learning breakthrough	2012 — University of Toronto transforms image recognition
Transformer architecture	2017 — Google paper enables modern language models
ChatGPT launch	November 2022 — reaches 1 million users in 5 days

Case Study: Dr. Patel's Grant Rejection

Dr. Priya Patel is a computer science researcher applying for a federal grant in 1984. Her proposal focuses on neural network approaches to speech recognition — a topic her department chair dismissed as 'fringe work from people who don't understand how real AI is built.' The dominant approach at the time involves expert systems, and her institution's reputation is tied to that school of thought. Her grant application comes back rejected with a note that neural networks had 'failed to demonstrate practical value.' Dr. Patel makes a note in her journal: 'The history of this field is a graveyard of promising ideas dismissed too early.' She decides to keep working on neural networks anyway, taking a reduced salary position at a smaller university. Thirty years later, in 2012, neural networks and deep learning will suddenly become the most important technology in AI. She will have been right all along — but the field had to wait for faster computers and larger datasets to prove it.

1. Dr. Patel's work was dismissed during a period of reduced AI funding and skepticism. Based on the reading, what was happening in the broader AI field in the 1980s, and why did expert systems dominate over approaches like neural networks?

2. Dr. Patel's note says the field is 'a graveyard of promising ideas dismissed too early.' How does the history described in the reading support this observation? What does this tell us about how we should evaluate AI research today?

Multiple Choice

1. What was the primary reason the first AI winter began in the mid-1970s?

a) The U.S. military stopped funding all computer research

b) AI failed to deliver on its promises and funding was cut

c) Computers became too expensive for universities to afford

d) Researchers ran out of mathematical problems to solve

2. What fundamental shift in approach helped restart AI progress in the 1990s?

a) Focusing only on hardware improvements and faster processors

b) Returning to the rule-based expert systems that had worked in the 1980s

c) Hiring more mathematicians to write better algorithms by hand

d) Shifting from hand-coding rules to training systems on large datasets

3. Which 2017 technical development became the foundation for tools like ChatGPT and Claude?

a) The transformer architecture introduced in 'Attention Is All You Need'

b) The development of expert systems that could learn and update dynamically

c) IBM's Watson defeating Jeopardy! champions in a live broadcast

d) The invention of the GPU processor for parallel computing

Matching

Write the letter of the correct definition on the line next to each term.

Term	Definition
1. ___ Turing Test	A) An AI program built from thousands of hand-coded if-then rules by human experts
2. ___ AI Winter	B) The 2017 neural network design that became the foundation for modern language AI
3. ___ Expert System	C) A machine learning approach using brain-inspired neural networks for complex tasks
4. ___ Deep Learning	D) A period of reduced funding and interest after AI fails to meet expectations
5. ___ Transformer Architecture	E) A thought experiment where a machine passes as human in written conversation

Fill-in-the-Blank

1. In 1956, John McCarthy coined the term "artificial intelligence" at a workshop held at _____ College.

2. IBM's Deep Blue made history in 1997 by defeating world chess champion _____ in a six-game match.

3. ChatGPT reached one million users in just five _____, making it the fastest-growing technology product of its time.

Short Answer

Answer each question in 2-4 complete sentences.

1. Explain what an AI winter is and describe what caused the second one in the late 1980s. What does this historical pattern tell us about the relationship between hype and progress in AI?

2. The reading describes the transformer architecture as the technical breakthrough that made ChatGPT and similar tools possible. In your own words, explain what made the transformer paper significant and why the timing — 2017 — matters in the broader story of AI.

3. Alan Turing proposed his test in 1950. Based on what you know from the reading about how modern AI works, do you think today's AI systems like ChatGPT would pass the Turing Test? What does your answer reveal about the limitations of the test itself?