

What Is Artificial Intelligence?

Understand the definition of AI, its history from the Turing Test to modern LLMs, and how AI differs from machine learning and deep learning.

Learning Objectives

- Define artificial intelligence and distinguish it from machine learning and deep learning
- Trace the key milestones in AI history from 1950 to the present
- Identify real-world AI applications across industries

Defining Artificial Intelligence

Artificial intelligence (AI) refers to computer systems that perform tasks that normally require human intelligence — things like understanding language, recognizing images, making decisions, and solving problems.

AI is not one technology. It is a broad field that includes many different approaches, techniques, and subfields. The most important distinction to understand is:

- **Narrow AI (ANI):** Systems designed to do one specific task well. Every AI system in production today is narrow AI — a spam filter, a recommendation engine, a language model.
- **General AI (AGI):** A hypothetical system that can perform any intellectual task a human can. AGI does not exist yet, and researchers disagree on whether it will be achieved in years, decades, or ever.
- **Superintelligence (ASI):** A hypothetical system that surpasses human intelligence across all domains. This is the subject of significant debate among AI researchers and philosophers.

KEY CONCEPT

AI vs. Machine Learning vs. Deep Learning: These terms are often used interchangeably but they have specific meanings. AI is the broadest category. Machine learning is a *method* of achieving AI. Deep learning is

a *technique* within machine learning. Think of nested circles: AI \supset Machine Learning \supset Deep Learning.

A Brief History of AI

Understanding where AI came from helps make sense of where it is today.

The Founding Era (1950–1970)

Alan Turing proposed the "Imitation Game" in 1950 — later known as the Turing Test — asking: *can a machine exhibit intelligent behavior indistinguishable from a human?* This was the first serious philosophical framework for thinking about machine intelligence.

In 1956, John McCarthy, Marvin Minsky, and colleagues organized the **Dartmouth Conference**, widely considered the founding event of AI as a formal field. They were optimistic — they believed human-level AI was a few decades away.

AI Winters (1970s–1990s)

Progress was slower than expected. Twice, funding dried up and interest waned — periods known as **AI winters**. Early systems were brittle: they could solve toy problems but failed on real-world complexity.

The Deep Learning Renaissance (2012–present)

In 2012, a neural network trained on GPUs won the ImageNet image recognition competition by a large margin. This triggered a wave of investment and research. By 2017, the **Transformer architecture** was introduced, enabling the large language models (LLMs) that power ChatGPT, Claude, and Gemini today.

Real-World AI Applications

AI is deployed across nearly every major industry today:

INDUSTRY	EXAMPLE APPLICATIONS
Healthcare	Medical imaging diagnosis, drug discovery (AlphaFold), clinical documentation

INDUSTRY	EXAMPLE APPLICATIONS
Finance	Fraud detection, algorithmic trading, credit scoring
Law	Contract review, legal research, document classification
Education	Personalized tutoring (Khan Academy), language learning (Duolingo)
Manufacturing	Predictive maintenance, quality inspection, supply chain optimization
Daily life	Search engines, recommendation systems, voice assistants, navigation

Key Takeaways

- AI is a broad field covering many techniques; the systems we use today are all narrow AI
- The history of AI includes dramatic swings between optimism and "AI winters"
- The 2012 deep learning breakthrough and the 2017 Transformer paper are the two most important milestones of the modern era
- AI is already embedded in most industries — this is not a future technology