

From: Intro to AI from MIT Technology Review newsletters@technologyreview.com
Subject: What is AI? 🤖
Date: September 17, 2025 at 6:17 AM
To: rkbassett@gmail.com



MIT Technology Review

Intro to AI

Week 1: What is AI?

Welcome to **Intro to AI**—a mini-course from *MIT Technology Review* about **how to get started with artificial intelligence**. Over the next six weeks, we'll take you on a journey that will give you a better understanding of what AI is and how to use it.

We've spoken with experts and curated our best coverage into **six editions to help you navigate the weird world of AI**. **This week's newsletter starts with the basics**. We'll begin by introducing you to key terms and concepts to ground you in the theories and technical advances that underpin modern AI.

Next week, we'll explore ways to use AI and share tips for how to test it out. And in the weeks to come, we'll cover how to talk about AI, ways to hone your media literacy skills, and AI's impact on the workforce. Finally, we'll end with a primer on AI regulations.

Right—let's dive in!





JUN IONED

A (very) brief history

Let's start at the beginning: **What is AI?** **Artificial intelligence** is an umbrella term for **a set of technologies that enable computers to do things that are thought to require intelligence when done by people**. Think of recognizing faces, understanding speech, driving cars, writing sentences, answering questions, or creating pictures.

AI isn't new. People have been working on artificial intelligence for more than 70 years, starting with a conference at Dartmouth College in the 1950s. There, computer scientists and mathematicians set a very ambitious goal: to develop computers that could understand language, translate it, see patterns, and more.

None of the projects that were spun up after that conference achieved those grand ambitions, but the vision stuck. And even though the technologies powering AI have changed over time, many of the field's ultimate goals remain the same.

By the 2010s, there was enough data and computational power to train AI systems to recognize images more accurately than any other computer program. The technique underpinning these systems, known as **deep learning** (more on that below), led to a big boom of investment, with large tech companies entering the space.

From there, deep learning was expanded and adapted for many applications, including recommendation algorithms and voice assistants. Then, within the last few years, we saw the first really powerful **generative AI** models, **GPT-3** and **DALL-E 2**.

Can it → "see"?

The grand ambition driving much AI research these days is to develop something resembling human intelligence, which is often referred to as **artificial general intelligence**, or AGI. Whether or not this is even possible is a matter of debate.

In that sense, AI is also aspirational, and its definition is constantly evolving. Even prominent people within tech have a hard time [agreeing on a definition](#). What would have been considered AI in the past may not be considered AI today. The bounds of AI can get really confusing, and the term often gets stretched to include any kind of algorithm or computer program. And Silicon Valley constantly inflates the capabilities of AI, so take tech companies' claims with a grain of salt.

AI Dictionary

Though there are many different flavors of AI, learning just a few basic terms will help you follow along. Let's start with these:

Neural Networks

Neural networks are the computer programs on which deep learning is built. Very loosely based on the brain, neural networks consist of many artificial neurons wired together through complex mathematical equations called **algorithms**. When a new piece of data (a picture of a cat, for example) is uploaded to a system, it cascades through those algorithms until different data (such as the word "cat") comes out the other end. An untrained neural network will spit out rubbish. A trained neural network, known as a **model**, is one that has learned to reproduce patterns seen in its training data to produce correct answers (most of the time).

Machine Learning

Machine learning uses statistics to find patterns in massive amounts of data. And data, here, can mean a lot of things—numbers, words, images, clicks, what have you. If it can be digitally stored, it can be fed into a **machine-learning algorithm**. Machine learning powers some of the AI that you're most likely to have used. Recommendation algorithms like the ones on Netflix, YouTube, and Spotify use machine learning, as do search engines, social media feeds, and voice assistants like Siri and Alexa.

Deep Learning

Deep learning is machine learning on steroids. It uses massive data sets, large neural networks, and a lot of computing power to give machines an enhanced ability to find—and amplify—even the smallest patterns. And some of these systems can spot patterns without even being told what to look for.

Generative AI

Generative AI is a popular and powerful kind of AI that's emerged in the last few years. **Image-generating AI** developed by Google, OpenAI, and others can now create stunning artworks based on just a few prompts. Type in a short description of pretty much anything, and you get a picture of it in seconds. ChatGPT does the same thing, but for text. There are engines to generate music, videos, and more. These tools have the power to **revolutionize the economy** and transform entire industries. For example, scientists are training models on specialized data about chemical compounds to try and **discover new drugs**. (Spoiler alert: no one really knows where **the rise of generative AI will lead us**.)

Large Language Models

Large language models are a kind of generative AI. They're built on deep-learning algorithms that are trained on **enormous amounts of text**. That means they can figure out which words are most likely to appear together in a sentence or paragraph and generate passages that sound like what a person might write. ChatGPT is a chatbot based on a large language model (GPT-3 and GPT-4).

(A quick note: You may have heard that large language models have aced tests like the bar exam or the US Medical Licensing Exam. The models tend to do well on exams, probably because there are plenty of sample questions and answers in their training data. But **we know remarkably little** about how these models work. By giving them these tests, we're trying to measure their "intelligence" based on their outputs, without fully understanding how they function.)

Multimodal AI

Multimodal AI refers to generative AI that processes multiple types of data, like audio, video, text, and images, to return results in any of these formats. Right now, we have a few multimodal models, like **Gemini** and **GPT-4**.

Keep in mind—AI isn't perfect and will get things wrong. It doesn't actually "know" anything (as far as we know), and there are serious flaws with AI-powered [search engines](#). ChatGPT even reminds you to double check any important information it generates with another source.

Alignment

Alignment is the work being done to ensure that AI systems do what users want them to do and nothing else.


Where to learn more

If you want to dive deeper into these topics, here's a little recommended reading:

- [Large language models can do jaw-dropping things. But nobody knows exactly why](#)
- [AI hype is built on high test scores. Those tests are flawed](#)
- [Why does AI hallucinate?](#)

Now that you know some key AI terms and history, we'll explore popular AI tools in our next edition. We'll cover what they can do and, perhaps more important, what they can't do yet.

Also from the editors of MIT Technology Review



The Algorithm Newsletter

Gain a better understanding of how artificial intelligence technology impacts the world around us.

[SIGN UP](#)

Download the MIT Technology Review app





LET'S STAY CONNECTED



[View in browser](#) | This email was sent to rkbassett@gmail.com.

[Manage your preferences](#) | [Unsubscribe](#) | [Terms of Service](#) | [Privacy Policy](#)

MIT Technology Review · 196 Broadway, 3rd fl, · Cambridge, MA 02139 · USA

Copyright © 2025 MIT Technology Review, All rights reserved.

[Opt out](#) of all promotional emails and newsletters from MIT Technology Review