

# What AI actually is

Reading time: ~12 min · Prereqs: none · 3 interactive demos inside

## The plain definition

---

**Artificial Intelligence (AI)** is software that does things we normally associate with human thinking: understanding language, recognising images, answering questions, writing, planning, or making predictions. That's it. It's a program. A very sophisticated one — but a program.


The kind of AI everyone is talking about right now is mostly built on **machine learning** — instead of a human writing step-by-step rules ("if X then Y"), the software *learns patterns from huge amounts of examples*. Show a system millions of photos labelled "cat" or "dog," and it learns to tell them apart without anyone writing a rule about whiskers.

Why does that matter to you, right now? Because it means the AI's behaviour comes from the patterns in its training data — not from a rulebook anyone wrote down. That's why it can be brilliant about some things and completely wrong about others, sometimes in the same sentence.

## The nested dolls — 90% of the confusion, cleared up

---

People throw the words "AI," "machine learning," "deep learning," and "LLM" around as though they mean the same thing. They don't. They fit inside each other like Russian nesting dolls. Here's the interactive version:



Keep this picture in mind whenever someone uses these terms. They're not interchangeable — they're a hierarchy. ChatGPT, Claude, and Gemini live in the small orange box at the centre. The rest of AI is the bigger territory around it.

## How an LLM works, without the math

An LLM is, at its core, a **next-word predictor**. It was trained by reading a huge slice of the internet and books, over and over, learning which words tend to follow which. When you type a question, it predicts a likely, helpful continuation — one word-chunk at a time — based on everything it learned.

That sounds too simple to be useful. And yet, because it absorbed so much text, "predict the next word really well" turns into "write an email," "explain a tax form," or "debug my code." The analogy that helps: it's like the most powerful version imaginable of your phone's autocomplete — one that has read almost everything and can keep a thought going for pages.

Want to see the idea in action? Try the demo below.

### Next-word predictor

Type the beginning of a sentence and click Predict. The widget uses a small hand-coded probability table — the same idea an LLM uses, just with 30 patterns instead of hundreds of billions.

e.g. "Once upon a" or "How are"

Predict next word

*This is basically what an LLM does — just with billions of examples instead of 30.*

## What AI is NOT

This section is arguably more important than everything above. Knowing what AI *isn't* protects you from the most common mistakes.

**It is not conscious or alive.** It has no feelings, no goals, no awareness of the world. When an AI says "I think" or "I feel," it's producing a pattern that fits the context — not reporting an inner experience. It's math producing text. Sophisticated math, but math.

**It is not a search engine.** A search engine looks up documents that exist. An AI *generates* a likely answer from the patterns it learned. It doesn't "look things up" by default — it predicts what a good answer should look like. Some tools can search the web as an added feature, but the base model is a predictor, not a browser.

**It is not a person.** It can sound warm, empathetic, and thoughtful. That's a product of training on human text. Don't share sensitive personal information, passwords, or private data with it that you wouldn't want on a public forum.

**It is not always right.** This is the one that trips people up the most, so it gets its own section and its own demo.

### The rule that matters most in this course

Treat everything an AI says like advice from a very fast, very well-read intern. Useful — sometimes brilliant — but you check anything that matters against a real source. This habit comes up in every lesson.

## Hallucinations — what they look like up close

A **hallucination** is when an AI states something false as if it were completely true — confidently, fluently, with no hesitation or caveat. The scary part is that the tone is identical whether the AI is right or wrong. There's no "hallucination voice."

Here's a side-by-side demo. Two questions, two answers. One is a real fact. One is entirely made up. Read both, then click the button to reveal what's fabricated in the right card.

### TRUTHFUL ANSWER

Q: "What's the capital of Australia?"

The capital of Australia is **Canberra**. It was purpose-built as the capital city and is located in the Australian Capital Territory, roughly halfway between Sydney and Melbourne. Construction began in 1913.

### HALLUCINATED ANSWER

Q: "Tell me about the 2019 novel 'The Bear's Nightmare' by Maria Chen."

"The Bear's Nightmare" is a Maria Chen novel published in 2019 by Meridian House Press. The story follows a wildlife biologist who discovers a series of bear attacks near a remote Alaskan research station are linked to an experimental pharmaceutical trial gone wrong. Chen's Booker Prize longlisting for this novel marked her as a major new voice in ecological thriller writing.

Show fabricated parts

Both answers above are written in exactly the same confident, fluent tone. That's the danger.

That's the hallucination problem in a nutshell. The AI's tone is identical whether it's reciting a real capital city or inventing an entire bibliography. You cannot tell from the writing alone. You can only tell by checking.

# The 5-term vocabulary you need

Learn these five terms and you'll be able to follow most AI conversations without getting lost.

## Quick-reference vocab

| Term           | Plain meaning                                                                                 |
|----------------|-----------------------------------------------------------------------------------------------|
| Model          | The trained AI itself — e.g., "Claude Opus," "Llama 3." Like a swappable engine.              |
| Prompt         | What you type to the AI — your question or instruction.                                       |
| Token          | A chunk of text (roughly $\frac{3}{4}$ of a word). AI reads and bills in tokens, not words.   |
| Context window | How much text the model can "hold in its head" at once. Run past it and it forgets the start. |
| Hallucination  | When the AI states something false as if it were true.                                        |

## Why this matters to you, practically

You don't need to understand how a transformer neural network works to use AI effectively — just as you don't need to understand combustion to drive a car. But you do need to know that engines can stall.

The three things to take from this lesson into every interaction with AI:

1. **It's a predictor, not a knower.** It generates likely text. "Likely" is not the same as "true." This is why you check anything that matters.
2. **Its confidence is not a signal.** AI sounds equally sure whether it's right or wrong. Never use tone as a proxy for accuracy.
3. **It's enormously useful anyway.** First drafts, explanations, brainstorming, summarising, translating — for all of this, a fast confident predictor is genuinely

powerful. You just keep the verify-anything-important habit running in the background.

### **A one-sentence mental model**

AI is a very fast, very well-read intern: brilliant at drafts and ideas, unreliable for facts and numbers, always worth a sanity-check on anything consequential.

## **Up next**

---

In **Lesson 02** we look at the types of AI that exist — narrow vs general, predictive vs generative, text vs image vs multimodal. Now that you know what AI is, it helps to see the full map of what's out there. Before you go, try the homework for this lesson — it'll put you in front of a real AI chat tool and let you see all three of today's concepts in action.

LAST-REVIEWED: 2026-06-04