

# Lesson 01 — What AI actually is

No jargon. No math. Just the plain truth about the thing everyone keeps talking about.

AI for Dummies

Module 1

No prerequisites

**By the end of this lesson you will:** know what AI actually is under the hood, understand the nested-dolls hierarchy (AI → ML → DL → LLM), see how a language model predicts its next word, and know the five terms that unlock every AI conversation.

# The honest definition

**Artificial Intelligence (AI)** is software that does things we normally associate with human thinking.

## What that includes:

- Understanding and generating language
- Recognising what's in a photo
- Answering questions and writing
- Planning, predicting, translating

**The important part:** it's a program. A very sophisticated one — but a program. There's no magic, no brain, no ghost in the machine. The reason this matters: programs can be wrong, and knowing *how* they go wrong lets you use them safely.

# AI is not new — what changed?

The term "artificial intelligence" dates to 1956. The chess program that beat Garry Kasparov is from 1997. So why does AI feel like it arrived last Tuesday?

## Old AI (rules-based)

A human expert sat down and wrote *every single rule*:

```
if "FREE MONEY" in email:
    mark_as_spam()
elif "WINNER" in email:
    mark_as_spam()
# ... 10,000 more rules
```

Fragile. Misses anything the author didn't think of.

## Modern AI (learning from data)

Instead of rules, show it examples:

```
# No rules written — just data
model.train(
    emails=10_000_000,
    labels=["spam", "not spam"]
)
# Model figures out the patterns
```

The software **learns patterns**. It can handle email types the engineer never saw.

The current wave is almost entirely **learning from data** — and the amounts of data are staggering.

---

# AI you already use every day

You're not a newcomer to AI — you've been using it for years without the label:



All of these use machine learning. The LLMs (Claude, ChatGPT) are a newer, more dramatic example of the same idea.  
The pattern: show the system enough examples, and it learns.

The models behind Claude or ChatGPT are bigger, more general versions of the same principle.  
The *scale* is new — the basic idea (learn from examples) has been around for decades.

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

People use "AI," "machine learning," "deep learning," and "LLM" as synonyms. They're not. They nest inside each other:

When someone says "AI" today, they almost always mean the innermost box. But now you know the full picture.

---

# How an LLM works — the next-word predictor

At its core, a language model is a **next-word predictor**. It was trained by reading a huge slice of the internet and books, learning which words follow which. When you type a question, it predicts a likely continuation one word-chunk at a time.

You type: "The capital of France is"

The model now computes probabilities for every possible next word:



It picks "Paris," appends it, then predicts the *next* word given "...France is Paris" — and keeps going until it reaches an end.

**This is literally all that's happening.** One word at a time. Repeated hundreds of times per response.  
The model never "thinks" the whole answer at once.

---

# Try it: next-word predictions

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

Try "once upon a" or "machine learning" or "how are"

Predict next word

Hint: try ending with "the", "a", "is", "once", "deep", "machine", "large", "thank", "how"...

A real LLM does this with 70+ billion numbers (its weights) learned from trillions of words. But the fundamental operation — "given what came before, what word fits best?" — is identical.

# Just bigger autocomplete — really?

Yes. And also: an extraordinary amount of capability emerges from doing that *extremely well at enormous scale*.

## Phone autocomplete

Learned from *your* texts. Predicts the next 1 word.  
Knows "I'll be there" → "soon."

Useful for one sentence. Falls apart immediately after.

Training data: ~thousands of your messages.

## Claude / ChatGPT

Learned from billions of web pages, books, code, papers. Predicts one word at a time but *holds the full context*.

Can write an essay, explain a law, debug your code, and then answer a follow-up — coherently.

Training data: trillions of words.

The gap between "phone autocomplete" and "Claude" is not a different technology. It's a different *scale* of the same technology.


---

# What AI is NOT — #1: Not conscious

When an AI says *"I think..."* or *"I feel excited about this!"* it is not reporting an inner experience. It's producing text that **fits the pattern** of how those phrases are used.

There are no feelings, no goals, no self-awareness, no opinions. There is math producing text — sophisticated math, but math.

**Why it matters:** people over-trust AI that "sounds passionate" or "seems confident." Those signals don't mean anything. A model can sound thrilled about a completely wrong answer.



"Tell me about Paris"

Billions of multiplications  
→ next token probabilities

# What AI is NOT — #2: Not a search engine

## Search engine (Google, Bing)

- **Indexes documents** that actually exist on the web.
- Returns **links to real sources**.
- Knowledge updates **continuously** as the web changes.
- Factual errors usually come from **bad sources**, not the engine itself.

## LLM (Claude, ChatGPT)

- **Generates an answer** from patterns learned during training.
- Does **not look anything up** by default — it predicts.
- Knowledge has a **training cutoff date**.
- Factual errors come from **the model itself** — no source to blame.

Some AI tools *can* search the web as an added capability. But the base model is always a predictor, not a browser. When in doubt, ask: "Did it cite a real URL?"

# What AI is NOT — #3: Not always right

**Hallucination:** when an AI states something false as if it were completely true — confidently, fluently, with no hesitation or caveat.

The terrifying part: **the tone is identical** whether the AI is right or wrong. There is no "hallucination voice." A model sounds exactly as assured when it invents a fake book title as when it correctly names a real capital city.

### **When AI is reliable**

- Writing drafts and summaries
- Explaining concepts and analogies
- Brainstorming and outlining
- Well-documented programming tasks

### **When AI is risky (verify these)**

- Specific facts, dates, statistics
  - Book titles, citations, papers
  - Recent events (past training cutoff)
  - Medical / legal / financial specifics
-

# Hallucination demo — spot the difference

Two AI answers. Both written in the same confident, fluent tone. One is completely true. One is entirely fabricated. Read them, then click to reveal.

## 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 bear attacks near a remote Alaskan station are linked to an experimental pharmaceutical trial. Chen's Booker Prize longlisting for this novel marked her as a major new voice in ecological thriller writing.

Show fabricated parts

Both are written identically. That is the danger.

---

# The mental model that fixes everything

**Treat AI like a very fast, very well-read intern.** Useful — sometimes brilliant — but check anything that matters against a real source.

## Great intern qualities:

- Reads everything you give them in seconds
- Writes excellent first drafts
- Explains complex things simply
- Never bored by repetitive tasks
- Instantly switches topics

## Intern limitations:

- Can sound confident while being wrong
- May invent citations to impress you
- Knowledge has an expiry date
- No common sense about *your* context
- Not responsible for their mistakes — you are

You don't need to understand how an engine works to drive a car. You just need to know that engines can stall. Same here.

# Tokens — what the model actually reads

AI doesn't read *words*. It reads **tokens** — chunks of text roughly  $\frac{3}{4}$  of a word on average. This matters for two practical reasons: billing and memory.

The sentence *"Unhelpfulness is never trivially safe."* splits into 7 tokens:

Un help ful ness is never trivially safe .

Common words = 1 token. Unusual words get split into sub-word pieces. Punctuation is its own token.

**Context window:** the maximum tokens the model can hold in memory at once. Think of it as a desk that fits a certain number of pages. Fill it up, and the model "forgets" the beginning.

**Billing:** cloud AI services charge by the token — both what you send (input) and what comes back (output). ~750 words  $\approx$  1,000 tokens.

# The 5 terms that unlock every AI conversation

Learn these and you'll be able to follow — and participate in — almost any discussion about AI tools, pricing, or limitations.

| Term           | Plain meaning                                            | Real example                        |
|----------------|----------------------------------------------------------|-------------------------------------|
| Model          | The trained AI itself — like a swappable engine.         | Claude Opus 4, GPT-4o, Llama 3.3    |
| Prompt         | What you type to the AI — your question or instruction.  | "Summarise this in 3 bullet points" |
| Token          | A chunk of text (~¾ word). AI reads and bills in tokens. | 1,000 tokens ≈ 750 words            |
| Context window | How much the model can "hold in mind" at once.           | 200 K tokens ≈ ~150,000 words       |
| Hallucination  | When the AI states something false as true.              | Inventing a book title or citation  |

These five terms appear in every AI product page, pricing table, and support ticket. Knowing them fluently makes you a more effective user immediately.

---


# What AI is good at — and what it isn't



**The habit that matters most in this course:** use AI freely for drafts, explanations, and brainstorming. Verify anything with a real-world consequence — facts, numbers, citations, legal or medical specifics — against a trusted source.

**Key insight:** AI is not all-or-nothing. Being a good AI user means knowing where on this spectrum your current task sits.

# Mental model: four sentences



1. **It's a predictor, not a knower.** "Likely" is not the same as "true."
2. **Confidence is not a signal.** AI sounds equally sure whether right or wrong.
3. **It's enormously useful anyway.** First drafts, explanations, brainstorming — a fast predictor is genuinely powerful.
4. **Verify anything that matters.** This habit runs in the background in every lesson of this course.

# Recap & what's coming

## What you now know:

- AI = software that mimics thinking — it's a program
- Modern AI learns from data, not rules
- AI ⊃ ML ⊃ Deep Learning ⊃ LLMs (nested dolls)
- LLMs are next-word predictors at huge scale
- Hallucinations are real, undetectable by tone
- The 5 terms: Model, Prompt, Token, Context window, Hallucination

## Coming up:

- **Lesson 02** — the types of AI that exist (narrow, generative, multimodal)
- **Lesson 03** — the four ways to use AI (Chat, API, CLI, Code tools)
- **Lesson 04** — installing Claude, LM Studio, Ollama
- **Lesson 05** — how to prompt well — the skill that changes everything

**Up next:** Read the Theory page for three interactive demos — the nested-dolls animation, the next-word predictor, and the hallucination toggle. Then do the Homework.

AI for Dummies

Module 1

What AI actually is

