AI BASICS · START HERE · 10 MIN READ

What is AI, and what is an LLM?

The words behind every AI tool, in plain language: AI, models, LLMs, tokens, context and hallucinations.

Pass the check (4 of 5) online to earn the “AI Explorer” badge:
https://digitalbuilders757.com/university/study/what-is-ai-and-an-llm

AI is a family of tools, not one robot

Artificial intelligence (AI) is software that does things we used to think needed a person: recognising a face in a photo, turning speech into text, recommending a song, or writing a paragraph.

Most modern AI is built with machine learning. Instead of a programmer writing every rule by hand, the software is shown a huge number of examples and learns the patterns in them. The result of that learning is called a model.

- Machine learning: software that learns patterns from examples.
- Model: the learned result you actually use. You send it something; it sends something back.
- Training: the expensive, one-time process of learning from data. Using the model afterwards is called inference.

An LLM is a model that works with language

A large language model (LLM) is a model trained on an enormous amount of text so it can read and write language. ChatGPT, Claude and Gemini are products built on top of LLMs.

At its core an LLM does one thing: given the text so far, it predicts what text should come next, one small piece at a time. Done billions of times with a very large model, that simple skill turns into answering questions, summarising, translating and writing code.

Because it predicts likely text rather than looking facts up, an LLM can sound confident and still be wrong. Knowing this is the most important habit in this whole library.

Tokens and the context window

LLMs don't read words the way we do. They read tokens: small chunks of text that can be a whole word, part of a word, or punctuation. Usage limits and prices for AI tools are usually counted in tokens.

The context window is how much text the model can look at in one go: your instructions, the conversation so far, any files you attached and its own replies. Think of it as the

model's desk, not its memory. When a chat gets very long, older parts can fall off the desk, and quality drops.

- Long chat going off track? Start a fresh chat with a short summary of what matters.
- Paste only the part of a document that matters for the question.

Hallucinations: when AI makes things up

A hallucination is an answer that sounds right but is false: a made-up statistic, a book that doesn't exist, a software setting that was never real. It happens because the model is producing likely-sounding text, not checking a database of facts.

You can reduce hallucinations by giving the model the real source to work from (paste the document, connect the tool, or turn on web search) and by asking it to say when it isn't sure. You can't remove them completely, so anything important gets checked against an original source.

Agents: AI that takes steps

A plain chatbot answers. An AI agent is an LLM that is allowed to use tools and take several steps toward a goal: search the web, read a file, run code, open a pull request, send a draft for your approval.

Agents are powerful because they act, which is also why they need limits: what tools they can use, what needs your approval, and a person who reviews the result.

Key terms

Artificial intelligence (AI)
Software that does tasks we used to think needed a person, like understanding speech or writing text.

Machine learning
Building software that learns patterns from examples instead of following only hand-written rules.

Model
The learned result of training; you send it input and it sends back output.

Training
The one-time, expensive process of a model learning from large amounts of data.

LLM (large language model)
A model trained on huge amounts of text that reads and writes language by predicting what comes next.

Token
A small chunk of text (a word, part of a word or punctuation) that models read and count usage in.

Context window
How much text a model can look at in one go: instructions, chat history, files and its own replies.

Hallucination
A confident AI answer that is false or made up.

AI agent
An AI that can use tools and take several steps toward a goal, not just answer.

Inference
Using a trained model to get an answer. Every chat message is an inference.

Try it: Catch a hallucination on purpose

1. Open any AI assistant (ChatGPT, Claude or Gemini).
2. Ask it for three facts about your town, each with a source link.
3. Open every link. Does the page exist? Does it say what the AI claimed?
4. Write one sentence about what you found. That checking habit is your superpower.

KNOWLEDGE CHECK

What is AI, and what is an LLM?

5 questions. Circle one answer for each; 4 or more correct passes. Answers are at the back.

1. What does an LLM fundamentally do?

A. Predicts what text should come next, based on patterns it learned
B. Looks up each answer in a large database of verified facts, then rewords it for you
C. Searches the web for every question and copies the best matching page
D. Runs a human typist in the background

2. What is a hallucination?

A. When the AI refuses to answer because a question breaks its safety rules
B. A confident answer that is false or made up
C. When the app crashes
D. A picture generated by AI

3. Your chat with an AI has become very long and it keeps forgetting your early instructions. What is the most likely reason?

A. Your internet is slow
B. The AI was updated since you started the chat
C. The early parts no longer fit well in its context window
D. Your device ran out of storage space

4. Which of these is a token?

A. A setting that remembers your favourite answer style
B. A digital coin, like a cryptocurrency, used to pay for AI
C. A saved chat that the AI can look back at
D. A chunk of text, often part of a word

5. What makes an AI agent different from a plain chatbot?

A. It can use tools and take several steps toward a goal
B. It only works on phones
C. It is a bigger model that has memorised far more facts than a chatbot
D. It has no model inside

Official sites

OpenAI ChatGPT: https://chatgpt.com
Anthropic Claude: https://claude.ai
Google Gemini: https://gemini.google.com

Last reviewed 2026-11-25. AI tools change fast; check each tool's own site for current features, plans and prices.

ANSWER KEY

What is AI, and what is an LLM?

1. A: Predicts what text should come next, based on patterns it learned
An LLM is trained to predict the next piece of text. That is why it is fluent, and also why it can be confidently wrong.

2. B: A confident answer that is false or made up
Hallucinations are plausible-sounding but false outputs. Check important facts against original sources.

3. C: The early parts no longer fit well in its context window
The context window is limited. Start a fresh chat with a short summary of the important instructions.

4. D: A chunk of text, often part of a word
Models read and count text in tokens. Limits and prices are usually measured in tokens.

5. A: It can use tools and take several steps toward a goal
Agents are LLMs allowed to use tools and act in steps, which is why they need limits and human review.