A FYTAHQ COMPANION RESOURCE

AI From Scratch

A Plain-Language Course
for people who did not grow up writing software

LESSON 2 OF 8

"Eighteen plain-language definitions. No code. No math. One example each."

The Vocabulary You'll Hear At The Stammtisch

ISSUE · JUNE 2026 · FACELESSYOUTUBEAUTOMATIONHQ.COM

Lesson 2 — The Vocabulary You'll Hear

WHAT YOU WILL KNOW AFTER THIS LESSON

- The eighteen words you will actually hear — in meetings, at dinner, and in YouTube comments about AI
- One plain sentence per word, plus one short example for each
- The gap between what a word literally means and how people use it in 2026
- Which words matter most for following any AI discussion
- Which words are just decoration that you can safely ignore

You do not need to memorise all eighteen. Three or four of them will carry you through most of the conversations you run into day to day. Read it once. The words settle in by themselves once you have heard them a couple of times.

And there is no test at the end. This is not school. The point is simply that the next time one of these words floats past you, it lands as a thing you already half-know rather than a small wall you bounce off. That is the whole goal — to take the edge off the jargon, so the ideas underneath can actually reach you.

HOW THIS GLOSSARY IS ORGANISED

Three tiers, eighteen words, one example each

The rest of this lesson is a plain-language dictionary. Eighteen words, each with one sentence of explanation and one short example you can picture. To make it easier to find your footing, the words are sorted into three tiers by how often you will actually run into them.

Start with Tier 1. Those six words turn up in almost every conversation about AI, so they earn most of your attention. Tier 2 is worth a calm read so the words do not catch you off guard later. Tier 3 is the decoration — pleasant to recognise, never essential.

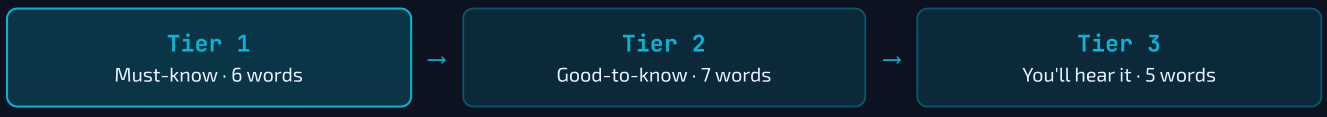

Whenever one word leans on another word you have not met yet, the meaning is spelled out in brackets right where it appears. That way you can read straight through from the top to the bottom without ever getting stuck.

WHY BOTHER LEARNING THE WORDS AT ALL

You are not learning these words to sound clever at the table. You are learning them so that the next time AI comes up — in a meeting, a sales pitch, or a worried headline — you can tell the difference between something that genuinely matters and something that is just ordinary noise dressed up in technical language. The vocabulary is the filter. Once you have it, a good deal of the hype quietly loses its grip on you, and you get to decide for yourself what is worth your attention.

TIER 1 OF 3 — MUST-KNOW

The six words that come up in almost every conversation

If you only ever remember six of these, remember these six. They are the load-bearing words: once you know them, most AI talk stops sounding like a foreign language. None of them is technical in the way an engineer would mean it — each one is just a plain label for a thing you have probably already bumped into without having a name for it.

TIER 1 · MUST-KNOW VOCABULARY

1. Model
The actual AI program itself — the trained thing that does the work when you press enter. Example: Claude is a model made by Anthropic, and ChatGPT runs on a model called GPT. Different companies, same basic idea.

2. Prompt
Whatever you type or say to the AI to set it going. A good prompt is simply a clear request in plain words. Example: "Write me a polite email about a delayed delivery" is a prompt.

3. Response (also: Output)
Whatever the AI sends back after you prompt it. The email, the summary, the answer, the list — all of it is the response. You will hear both words used to mean the same thing.

4. Token
The little chunks of text the AI reads and counts in, each one a bit smaller than a word. As a rough guide, 100 tokens is about 75 words. Example: this sentence is roughly a dozen tokens.

5. Hallucination
When the AI states something false while sounding completely confident. This is the single most important word in the whole list, because it is the one that can quietly catch you out. Example: it invented a book title and an author that simply do not exist.

6. Context Window
How much text the AI can hold in mind during one conversation, before the earliest parts start to slip away. Example: a context window of 200,000 tokens is roughly 150,000 words — about the length of a novel.

HOW THE SIX FIT TOGETHER

Read in order, they tell one small story. You write a **prompt**, the **model** sends back a **response**, all of it is measured in **tokens**, and the conversation can only run as long as the **context window** allows. Every so often the model will **hallucinate** — which is exactly why you always read its answers with one eyebrow raised.

TIER 2 OF 3 — GOOD-TO-KNOW

Seven words worth a calm read

You will not need these every day, but when they turn up you will be glad they do not surprise you. Skim them once and move on.

TIER 2 · GOOD-TO-KNOW VOCABULARY

7. Training
The one-time, expensive process in which the model learned from huge amounts of text, long before you ever arrived. Example: a model trained on text up to early 2025 has not "read" anything published after that date.

8. Inference
What happens each time you actually use the model — fast, cheap, and the part anyone pays for. Example: every single message you send is one inference.

9. Fine-Tuning
A smaller, second round of training that nudges a model towards one specific job. Example: a model fine-tuned on legal contracts will speak fluent legalese.

10. API (A-P-I)
The doorway programmers use to plug their own software into an AI model — you will never touch it directly. Example: an app might quietly call OpenAI's API to translate your messages.

11. RAG (Retrieval-Augmented Generation)
An AI set up to look things up first and answer second. Example: with RAG, the AI checks your own documents before it replies, instead of guessing from memory.

12. Agent
An AI that can take actions, not only answer — though still rough around the edges in 2026. Example: an agent might book the flight for you, not just describe the options.

13. LLM (Large Language Model)
The specific kind of AI that nearly everyone means today when they say "AI". Example: Claude and GPT are both LLMs.

THE ONE PAIR WORTH HOLDING ONTO

If the seven start to blur together, keep just **training** versus **inference**. Training is the slow, costly schooling that happened once, in the past. Inference is the quick, cheap moment that happens every time you actually use it. Almost everything else on this page is a variation on those two ideas.

TIER 3 OF 3 — YOU'LL HEAR IT

Five words it is enough to simply recognise

TIER 3 · YOU'LL-HEAR-IT VOCABULARY

14. Multimodal
An AI that handles more than text — images, sound, and more, not just one of them. Example: you can show a multimodal model a photo and ask what is in it.

15. Open Source vs Closed Source
Who is allowed to run and change the model. Open source: anyone can download and run it. Closed source: only the company that built it can.

16. Local vs Cloud
Where the model actually runs. Local: on your own computer — private, but slower. Cloud: on their servers — faster, but your text leaves your machine.

17. BYOK (Bring-Your-Own-Key)
You connect your own AI subscription to an app instead of paying that app separately. Example: a BYOK note-taking app runs on your own ChatGPT account.

18. Vibe Coding
A half-joking term for building something by chatting with an AI until it somehow works. Example: "I vibe-coded a little website over the weekend."

A CONCRETE EXAMPLE

Picture yourself at a Stammtisch — the regular table at the pub — and somebody at the other end says this:

> *"We tried an LLM through the API to summarise our customer feedback. The prompts were fine, but the model hallucinated a few quotes that were never said. We thought about fine-tuning it on our own data — except the context window is too small for the whole dataset. So we might switch to a RAG setup instead."*

A week ago that sentence was noise. Now every single piece of it has a meaning you can point to. You may not pick the winning side of the debate — but you know what each word means, you can ask a sensible question, and you are genuinely in the conversation.

ONE SMALL THING TO NOTICE

Notice that the person at the other end of the table was not showing off. They were just describing a real problem in the plain working language of 2026 — the same way a builder might mention joists and a load-bearing wall. The words only sound like a private code from the outside. From the inside, they are simply the shortest way to say something true.

So the next time a sentence like that drifts past you, do not let the wrapper fool you. Pick it apart one word at a time, the way you just did here, and you will almost always find a perfectly ordinary thought

Lesson 2 — The Vocabulary You'll Hear

THREE THINGS TO LOOK OUT FOR

A few of these words get bent out of shape in everyday use — stretched, swapped around, or quietly used to mean something grander than they really do. None of it is a trap set on purpose; it is just how language drifts. Here are the three worth keeping a gentle eye on.

1 **"LLM" and "AI" are not the same word, but in 2026 people use them as if they were.** Strictly, AI is the broad category and an LLM is one specific kind of it. In everyday office talk, people usually mean an LLM when they say AI. Knowing this saves you confusion when somebody says "we should use AI for this" — they almost always mean an LLM.

2 **"Tokens" show up the moment money is discussed.** Most cloud AI services charge by the token rather than by the message. One million tokens is roughly 750,000 words of English. Knowing that turns a vague worry about "API costs" into a concrete number you can actually reason about.

3 **"Agent" is the most overloaded word of 2026.** Marketing slaps it on almost everything. Strictly, an agent is an AI that can take real actions — book, send, schedule, change something. If the demo only chats back at you, it is a chatbot (an AI that answers but does not act), not an agent.

WHAT IS SAFE TO SKIP

You do not need to know any of the following to follow a conversation about AI:

- How tokenization works under the hood
- The mathematical definition of "attention"
- Why a context window has a hard limit at all
- The difference between a transformer, a decoder-only transformer, and an encoder-decoder transformer
- The exact per-token pricing of each provider, which changes every quarter anyway

These matter to the engineers who deploy AI for a living. For simply following along — at work, at dinner, or in the comments — none of them is required. You can always come back to them later if your curiosity ever asks.

THE HONEST LINE BETWEEN THE TWO LISTS

Here is the quiet truth that the "safe to skip" list points at: you can take part in almost any conversation about AI without understanding how a single piece of it works on the inside. Knowing **what** a tool does, and roughly **when** to trust it, is a completely different skill from knowing **how** it was built — and only the first one is your job. Leave the wiring to the people who chose that line of work. Your job is to use the thing wisely, ask good questions, and keep your own judgement switched on.

BEFORE YOU CLOSE THIS

Eighteen words is more than most people ever bother to learn, and you just read all of them in one sitting. You do not have to keep them in your head. Keep this page instead, and glance back whenever a word trips you up.

If only one word stays with you, let it be "hallucination" — the habit of double-checking a confident answer is the single most useful thing you can take from this entire lesson.

NEXT TUESDAY

Lesson 3 — What AI Can Actually Do, And What It Cannot

A plain capability map: what AI does genuinely well, where it quietly fails, and the dangerous middle ground where it is confident and wrong at the same time. Free on the resources page — no email, no sign-up.

PREFER A LIVE CONVERSATION?

One-on-one sessions with a real person

If you would rather have a real human walk you through this material in a live one-on-one session than read it on your own, the mrocon side of our world offers exactly that — same philosophy as this course, just guided. One hour, two hours, or three, depending on how deep you want to go.

mrocon.at/pricing →

COMPANION VIDEO

Lesson 2 — Video Version on the channel

The Lesson 2 video is live on YouTube — same eighteen words, in video form. The full "AI From Scratch" playlist updates every Tuesday for the next several weeks. Watch what fits your evening, read what fits your morning.

facelessyoutubeautomationhq.com →

Sources + version note: *Content valid as of June 2026. Reference categories: ChatGPT / OpenAI · Claude / Anthropic · Gemini / Google · open-source model communities · general AI-industry reporting. The vocabulary in this lesson is deliberately drawn from stable, widely used categories rather than fast-moving product names, so it stays useful between update cycles. Specific model versions, token prices, and feature availability shift quickly.*