



FITNESS MARKETING MACHINE
SCALE WITHOUT LIMITS

The words you need before **the tools** make sense.

Every AI glossary is alphabetical, which is useless: you cannot understand *context* *rot* before you understand what a *token* is. This one is ordered the way the ideas actually stack. Read it top to bottom once and the rest of this world stops sounding like a foreign language.

37 TERMS 8 LEVELS ~20 MIN READ FOR COACHES, NOT ENGINEERS

HOW TO READ IT

Each level uses only words defined in the levels above it. **Nothing is explained using a term you have not met yet.** That constraint is the entire design, and it is why this works when other glossaries don't.

Every entry gives you four things: the plain-English definition, the same idea in coaching terms, and the part that actually matters most: **the one behavior it should change.** If you only skim one line per entry, skim that one.

You do not need to memorize any of this. You need to have *heard* it once, so that when it shows up in a video or a settings screen, it lands as familiar instead of intimidating.

01

What this thing actually is

Six words that explain why AI behaves the way it does. Almost every frustration people have with AI traces back to not knowing one of these.

BUILDS ON **NOTHING** · START HERE

01

LLM

Large Language Model · say it "ell-ell-em"

The AI itself. A program trained on an enormous amount of text that predicts what should come next, one small piece at a time.

IN YOUR WORLD

It is not a database you search. It is closer to an extremely well-read intern who has read most of the internet and has never once seen your client roster.

WHY YOU CARE

It does not look anything up unless you give it a way to. Everything it says about *your* business is a guess until you hand it your actual data.

02

Token

The unit of measurement for all of this

The chunk of text the AI reads and writes in, roughly three quarters of a word. "Coach" is one token. "Periodization" is three or four.

IN YOUR WORLD

Tokens are to AI what reps are to a set. Everything gets counted in them: what you can send, what it can say back, and what you pay.

WHY YOU CARE

Every limit and every bill in AI is measured in tokens, not messages. A 40-page PDF is not "one message". It is thousands of tokens.

03

Context window

The single most important idea on this page

The maximum number of tokens the AI can hold in mind at one time: your instructions, the conversation so far, every file you have given it, and everything it has already said back.

IN YOUR WORLD

It is working memory, not a file cabinet. Like a client who can hold their whole plan in their head during the session and remembers none of it on Thursday.

WHY YOU CARE

This is the ceiling on everything. Nearly every "the AI got dumber" moment is really "we filled up the context window."

04

Model

Opus · Sonnet · Haiku

A specific version of the AI, with a name, a capability level and a price. Claude comes in a few: Opus is the most capable, Sonnet is the balanced default, Haiku is the fastest and cheapest.

IN YOUR WORLD

Same logic as programming a session. You do not send every client to failure on back squats. You match the tool to the job in front of you.

WHY YOU CARE

Reaching for the biggest model on everything is like doing max singles every workout: expensive and slow for no benefit. Use the small fast one for simple, repetitive work.

05

Hallucination

The failure mode you must plan around

When the AI states something false with complete confidence: an invented statistic, a study that does not exist, a feature that was never built.

IN YOUR WORLD

It is not lying to you. It is the client who tells you they hit 10,000 steps every day and genuinely believes it. The confidence is real. The fact is not.

WHY YOU CARE

Confidence is not accuracy. Anything touching money, a medical claim, or a client's personal data gets checked by a human being. Every single time.

06

Non-deterministic

Why it never answers the same way twice

Ask the identical question twice and you get two different answers. Both can be completely correct. This is by design, not a defect.

IN YOUR WORLD

Two excellent coaches write two different programs for the same client. Neither of them is wrong.

WHY YOU CARE

Stop treating it like a calculator. If you need byte-identical output every time, you need software, or you need the AI to *write* you that software, which it is very good at.

02

How you talk to it

The gap between people who get astonishing results and people who get mediocre ones is almost never the model. It is these five words.

BUILDS ON **TOKEN** · **CONTEXT WINDOW**

07

Prompt

What you send it

The question, instruction, or raw material you hand over. Everything you type is a prompt.

IN YOUR WORLD

It is a client intake. "Get me in shape" and "I'm 42, three hours a week, bad left knee, want to pull 315 by June" produce very different programs.

WHY YOU CARE

Most disappointing AI output is a vague brief, not a weak model. Specificity is the entire game.

08

Prompt engineering

Writing the brief on purpose

Deliberately shaping your instructions to get better output: being specific, showing examples, defining what "good" looks like, and stating what to avoid.

IN YOUR WORLD

The difference between telling a client "do some cardio" and prescribing a 40-minute zone-2 session with a heart-rate range and an RPE target.

WHY YOU CARE

The highest-leverage skill on this entire page, and the cheapest to learn. One great example of what you want beats three paragraphs describing it.

09

Context engineering

The step up from prompt engineering

Deliberately managing everything that occupies the context window: which files are loaded, which history is carried forward, which standing instructions apply, so that the AI has exactly what it needs and nothing that distracts it.

IN YOUR WORLD

Prompt engineering is writing a good workout. Context engineering is designing the whole training environment: what equipment is in the room, what is on the whiteboard, and what the client already walked in knowing.

WHY YOU CARE

This is where the people getting remarkable results actually differ. It is not a cleverer prompt. It is a better-loaded context.

Context rot

Why long conversations get worse

The gradual decay of a long conversation. As the context window fills with old detours, abandoned ideas and superseded information, the AI starts confusing what is still true with what was true an hour ago.

IN YOUR WORLD

The client you have coached for three years through six goals, four injuries and two nutrition rewrites. Somewhere in there you lost track of what is actually true today.

WHY YOU CARE

Long conversations are not free. When answers start drifting, do not argue with it. Start a fresh conversation and carry over only what still matters.

Compaction

The automatic summary you should not fully trust

When a conversation runs up against the context window, the AI summarizes what came before and carries on from the summary rather than the full transcript.

IN YOUR WORLD

A handover note between coaches. Genuinely useful, and lossy: it keeps the decisions and loses the texture.

WHY YOU CARE

After a compaction, assume detail was dropped. Restate anything load-bearing instead of trusting that it survived.

How it learns your business

Generic AI gives generic answers. These four terms are how it stops guessing about you and starts working from what is actually true.

BUILDS ON **CONTEXT WINDOW** · **HALLUCINATION**

12

Grounding

The fastest quality upgrade available to you

Giving the AI real source material to work from, instead of letting it answer out of general knowledge.

IN YOUR WORLD

The difference between asking "how is my client doing?" and handing over their check-ins, their weigh-ins and their actual training log.

WHY YOU CARE

Ungrounded AI is an opinion. Grounded AI is analysis. Before you blame the model, ask what you actually gave it.

13

Source of truth

The boring discipline everything else depends on

The one place a fact officially lives. When two systems disagree, the source of truth is the one that wins by definition.

IN YOUR WORLD

Client bodyweight sits in a spreadsheet, in your coaching app, and in a DM from Tuesday. Which one is real? Until you decide, none of them are.

WHY YOU CARE

AI faithfully multiplies whatever you feed it, including your mess. Naming the source of truth for each kind of data is unglamorous work that makes every clever thing downstream possible.

14

Memory

Not the same thing as the context window

What the AI retains between separate conversations. The context window resets every time; memory is written down somewhere and deliberately reloaded.

IN YOUR WORLD

The context window is what is in your head during the session. Memory is what is written in the client's file.

WHY YOU CARE

By default it remembers nothing about you. Anything you want it to still know next Tuesday has to end up in a file it reads on startup.

Project instructions

often a file literally named CLAUDE.md

A file of standing instructions the AI reads at the start of every single session: who you are, how you work, what it must never do, and where things live.

IN YOUR WORLD

The onboarding document you write once for a new assistant coach, instead of explaining your standards on day one for the rest of your life.

WHY YOU CARE

This is the highest-return file you will ever write. Every correction you find yourself repeating belongs in it. Permanently, and that same day.

How it touches your other software

This is the level where AI stops being a writing assistant and starts being staff. It is also where most people get stuck for the first time.

BUILDS ON **GROUNDING** · **SOURCE OF TRUTH**

16

API

Application Programming Interface · say it "ay-pee-eye"

The defined way one piece of software talks to another without a human clicking anything. Every serious tool you pay for has one.

IN YOUR WORLD

Your gym's front desk. There is a defined way to ask for things, a defined way answers come back, and you do not get to wander into the back office and help yourself.

WHY YOU CARE

"Does it have an API?" is the single question that decides whether a tool can ever be automated. Ask it *before* you buy the software, not after.

17

Authentication

often shortened to "auth"

Proving to a system that you are who you claim to be, and that you are permitted to do the thing you are asking to do.

IN YOUR WORLD

The key fob at the 24-hour gym. The door does not know you or care about you. It only cares whether your fob is valid, current, and allowed through that particular door.

WHY YOU CARE

This is where nearly everyone gets stuck first. When a connection "just stops working," it is almost always an expired or wrongly-permissioned credential, not a broken tool and not your fault.

18

API key

also: token, secret, credential

The secret string of characters that authenticates you. It looks like `sk-ant-a1b2c3...` Whoever is holding it can act as you.

IN YOUR WORLD

A key to the gym, not a membership card. It does not identify you. It simply opens the door for whoever happens to be holding it.

WHY YOU CARE

Treat it exactly like a password, because that is what it is. Never paste one into a chat, a screenshot, a screen share, or a video. If it is ever exposed, replace it that hour.

Tool use

also called function calling

The AI's ability to *do* things rather than only describe them: read a file, run a search, send a request, update a record. It decides which tool to reach for and when.

IN YOUR WORLD

The jump from a consultant who writes you a beautiful plan to a coach who is on the floor moving the plates.

WHY YOU CARE

This is the exact line between AI that advises and AI that operates. Everything genuinely valuable is on the far side of it.

MCP

Model Context Protocol • say it "em-see-pee"

A shared standard for plugging tools into an AI. Instead of every AI needing custom code for every app, an app publishes one MCP server and any AI can use it.

IN YOUR WORLD

The 45-pound plate. Once the industry agreed on the standard, any plate fits any bar in any gym on earth. MCP is that agreement, for AI tools.

WHY YOU CARE

When you ask "can Claude talk to my CRM?", the real question is "does it have an MCP server?" That is what turns AI from a chat box into something wired into the business you already run.

The Claude Code vocabulary

Eight words that describe the actual machinery. Learn these and you can follow any tutorial, any video, and any settings screen in this world.

BUILDS ON **TOOL USE** · **MCP** · **PROJECT INSTRUCTIONS**

21

Claude Code

the thing this whole page is preparing you for

Claude with hands and a workspace. It runs on your own computer, reads and writes your files, runs commands, and uses connected tools, rather than only answering questions in a browser tab.

IN YOUR WORLD

The difference between texting a contractor photos of your kitchen and handing them the keys to the house.

WHY YOU CARE

Despite the name, it is not only for code. It is the best tool available for running a business made of files, data, and connected software, which is every online coaching business.

22

Terminal

also: command line, CLI, shell

The plain text window where you type commands instead of clicking buttons. CLI stands for command-line interface.

IN YOUR WORLD

Free weights versus machines. More intimidating on day one, fewer guardrails, dramatically more range once you are in.

WHY YOU CARE

This is the one real hurdle on this page, and it is far smaller than it looks. You need about five commands to start, and the AI will hand you every one after that.

23

Agent

the word that changes your job description

An AI given a defined role, a set of tools, and a goal, which then works through the steps on its own rather than answering one question at a time.

IN YOUR WORLD

The difference between an intern who answers exactly what you ask and a staff coach who owns client onboarding end to end.

WHY YOU CARE

You stop writing prompts and start writing job descriptions. That shift is the entire point of everything above.

24

Subagent

the practical cure for context rot

A fresh agent that the main one spins up to handle a chunk of work in its own clean context, reporting back only the conclusion.

IN YOUR WORLD

Delegating properly. Your assistant reads all 40 check-ins and brings you the three that need you. You never load the other 37 into your head.

WHY YOU CARE

Heavy reading happens somewhere else and only the answer comes back. This is how you work on big things without poisoning your own context window.

25

Skill

your SOPs, made executable

A written procedure the AI loads only when it is relevant: a small folder holding the instructions for doing one specific thing your way, every time.

IN YOUR WORLD

A standard operating procedure. Written once, followed identically whether it is 9am Monday or 11pm Friday.

WHY YOU CARE

Every process you currently carry in your head is a skill waiting to be written down. Once written, it is executable: by the AI, by a new hire, or by both at once.

26

Slash command

named for the / you type first

A shortcut you type, like `/debrief` or `/invoice`, that fires a saved instruction instead of you retyping the same request.

IN YOUR WORLD

A saved workout template. You built it carefully once; now it is two taps forever.

WHY YOU CARE

The moment you catch yourself typing roughly the same request for the third time, that is the signal to make it a slash command.

Hook

the difference between a habit and a system

A rule that fires automatically on an event, with no prompting from you at all: before a tool runs, after a file changes, when a session ends.

IN YOUR WORLD

The email that goes out automatically when a client misses two check-ins. Nobody decides to send it. It simply goes.

WHY YOU CARE

Prompts are things you remember to do. Hooks are things that happen whether you remember or not. Anything you must never forget belongs in a hook, not in your head.

Plugin

skills, commands, agents and hooks in one box

A bundle of all of the above, packaged so that someone else can install the whole kit in a single step.

IN YOUR WORLD

Buying a complete program rather than writing every session yourself.

WHY YOU CARE

You do not have to build everything from scratch. And once you have built something genuinely good, this is how you hand it to your team, or sell it.

Letting it work without you

The part everyone wants and almost everyone rushes. These four words are what separate leverage from a mess you have to clean up.

BUILDS ON **AGENT** · **TOOL USE** · **HALLUCINATION**

29

Agentic

the adjective, used constantly, rarely explained

Describes AI that chooses its own next step. You supply the outcome; it works out the sequence.

IN YOUR WORLD

"Get her ready for a half marathon in 16 weeks" instead of "write Tuesday's run."

WHY YOU CARE

The word marks a change in what *you* are responsible for. You stop supervising steps and start defining outcomes and boundaries.

30

Autonomous work

the overnight-run promise

The AI running a long task through to completion without you watching: many steps, sometimes hours, sometimes while you sleep.

IN YOUR WORLD

Handing off a project on Friday afternoon and reviewing the finished thing Monday morning.

WHY YOU CARE

The real constraint is not capability, it is scope. Autonomy works beautifully when the task is clearly bounded and the AI can check its own work. It fails when the goal is vague.

31

Human in the loop

what makes autonomy safe enough to actually use

A deliberate design where the AI prepares the action but a person approves it before it takes effect.

IN YOUR WORLD

Your assistant drafts every client email and you press send. All of the speed, none of the "who on earth sent this?"

WHY YOU CARE

Decide up front which actions are automatic and which need your thumb. Money, client communication, and anything that deletes something are the usual line.

Verification

the habit that separates the pros

Confirming the work is actually finished and correct, rather than accepting the AI's report that it is.

IN YOUR WORLD

The client says they hit their macros all week. The progress photo says otherwise.

WHY YOU CARE

AI reports on itself optimistically. "Did it run?" and "did it produce the right result?" are two different questions, and only the second one has ever mattered.

The files all of this lives in

Three formats. You can learn all three in under ten minutes, and they are what every skill, instruction and note in this world is made of.

BUILDS ON SKILL · PROJECT INSTRUCTIONS

33

Markdown

file extension .md

Plain text with a handful of simple marks for formatting. # makes a heading, ****like this**** makes bold, - starts a bullet. That is genuinely most of it.

IN YOUR WORLD

Writing a program on a napkin instead of in design software. Zero friction, readable by anyone, opens on any device, still works in twenty years.

WHY YOU CARE

It is the native language of AI-readable documents. Every skill and instruction file in this world is markdown, and you can learn the whole format in about four minutes.

34

Frontmatter

the block between the --- lines

A small block at the very top of a markdown file holding facts *about* the file: its name, what it is for, and when it should be used.

IN YOUR WORLD

The label on the outside of the folder, as distinct from the paperwork inside it.

WHY YOU CARE

It is how the AI decides whether a file is even relevant before reading it. A skill with a vague description never gets picked up, no matter how good the instructions inside are.

35

Repo & Git

repository · the undo button for your business

A repo is a folder that tracks its own history. Git is the system doing the tracking: every change, who made it, when, and exactly how to undo it.

IN YOUR WORLD

Version history for your entire operation, with an undo button that reaches back years and never expires.

WHY YOU CARE

It is the safety net that makes it reasonable to let AI edit your real files. If something goes wrong, you roll it back with one command and lose nothing.

What it actually costs

Two terms that will change how you use this daily, and stop you being needlessly stingy in exactly the wrong place.

BUILDS ON **TOKEN** · **CONTEXT WINDOW** · **MODEL**

36

Input vs. output tokens

the two halves of every bill

You are charged for what you send *and* what comes back, at different rates. Output typically costs several times more than input.

IN YOUR WORLD

Reading a client's check-in is cheap. Writing their next twelve-week block is not.

WHY YOU CARE

Sending a big document is far cheaper than most people assume. Asking for a big document back is not. Ask for the decision, not the essay, unless you genuinely need the essay.

37

Prompt caching

why being thorough is affordable

When you send the same context repeatedly (your instruction file, your knowledge base), it can be cached and re-read at a steep discount instead of charged in full every time.

IN YOUR WORLD

A membership versus a day pass.

WHY YOU CARE

It is why loading real, substantial context on every session is affordable. Being thorough with your setup is cheaper than being clever about trimming it.

JUST AS IMPORTANT

What you do not need to know

Half of feeling behind is people talking about things that have nothing to do with you. These are real, they are genuinely technical, and you can operate at a very high level without ever touching one of them.

- ✖ **Fine-tuning.** Retraining a model on your own data. Enormously expensive, and good context beats it for essentially everything you would want to do.
- ✖ **RAG.** A specific, formal way of doing grounding. Worth recognising the acronym; not worth building one yourself in 2026.
- ✖ **GPUs, parameters, training runs.** How models get built. Fascinating. Completely irrelevant to using one.
- ✖ **Vector databases & embeddings.** The plumbing behind AI search. Real, useful, and someone else's job. The tools you will use handle it invisibly.
- ✖ **Writing code.** You need to read it well enough to sanity-check, not write it. The AI writes it. You direct it.
- ✖ **Which model is beating which on the leaderboard.** It changes monthly. Pick a good one, learn context engineering, and stop reading the scoreboard.

THE SHORT VERSION

If you remember nothing else, remember these eight.

Every term on this page collapses into one of these. Screenshot this section.

01 **Context is the whole game.**

Not the model, not the prompt. What is loaded into the window when you ask.

02 **Confidence is not accuracy.**

It will be wrong in exactly the same tone it is right. Check anything that matters.

03 **Ground it, or it is guessing.**

Give it your real data. An ungrounded answer about your business is fiction with good grammar.

04 **Write it down once.**

Every correction you repeat is a file you have not written yet. Instructions, skills, memory.

05 **Tools are the line.**

AI that can only talk gives advice. AI with tools does the work. Everything valuable is past that line.

06 **Delegate, do not accumulate.**

Send heavy reading to a subagent. Protect your own context like you protect your calendar.

07 **Automatic beats remembering.**

If it must never be forgotten, it belongs in a hook, not in your head and not in a prompt.

08 **Verify the result, not the report.**

"I have completed the task" is a claim, not evidence. Go look.

None of this requires you to become a developer. It requires you to stop being a tourist. Thirty-seven words is the whole toll.