

The 7 Levels of Claude — the full ladder

Every level of the Claude app, what each one is actually for, and exactly how to use it — plus a quick self-diagnosis that tells you which level to climb next.

You commented LEVELS, so here's the whole thing.

■ The problem nobody tells you about

Most people open Claude, type a question, read the answer, close the tab. That's it. That's level one, and they never leave it.

Here's the thing: the app you're already paying for (or using free) has seven levels stacked inside it. File reading, saved workspaces, working code, cited research, hooks into your real tools, and tasks that run themselves. Same app, same login. Most people just never scroll past the chat box.

This guide is the full ladder. Read it once, find your level, climb one.

■ The taste — what the top looks like

At level one you ask a question and get a paragraph back.

At level seven you hand Claude a messy multi-step job — "go through these files, pull out every deadline, cross-check them against my calendar, and draft the emails" — and it breaks the job into steps, runs them itself, and checks its own work before showing you. You come back to finished work, not an answer to the first step.

Same tool. The difference is knowing the levels exist. Here they are.

■ Level 1 — Chat

What it is: the basic conversation. The part everyone knows — but most people use it wrong.

Why it matters: Claude keeps the whole thread in context. Every reply you send sharpens what it knows about what you want. So the move isn't to restart when the answer's off — it's to refine. "Shorter." "More concrete." "Now explain it like I'm the client." The fifth reply in one thread beats five fresh chats every time.

How to use it: type your question, then keep replying until it's right. Never open a new chat to fix an answer — push the same thread.

■ Level 2 — Files

What it is: upload a PDF, spreadsheet, or doc straight into the chat. Claude reads every page — not a skim, the whole thing.

Why it matters: you stop retyping and summarizing. Drop in a 60-page contract and ask "what am I agreeing to in section 12?" Drop in a CSV export and ask what's off in the numbers. You question the whole document at once, instead of feeding it to Claude in pasted chunks.

How to use it: drag the file into the chat window, then ask about it. That's the entire move.

■ Level 3 — Projects

What it is: a saved workspace that holds your files AND your custom instructions in one place. Every chat you start inside it already knows your work.

Why it matters: this kills the most annoying habit at the lower levels: re-pasting the same context into every new chat. A Project covers four things at once — your files, your context, your standing instructions ("always write in my voice", "our client is X"), and memory of the work. Set it up once and every chat after starts warm.

How to use it: make a Project, add your files once, write your instructions once, then do all your chats for that piece of work inside it.

■ Level 4 — Artifacts

What it is: Claude writes and runs real code in a side panel next to the chat. Not a code snippet you have to take somewhere else — a working app, chart, or tool, live on the screen.

Why it matters: "describe a dashboard" and "here is the dashboard, click it" are different products. Ask for a calculator for your pricing, a chart of your data, a small interactive tool — and you get the working thing, then refine it by talking.

How to use it: ask for a tool, not an explanation. "Build me a calculator that does X." Then keep chatting to change it.

■ Level 5 — Research

What it is: a mode where Claude searches hundreds of sources, cross-checks them against each other, and comes back with a cited report —

links included.

Why it matters: a normal chat answer is Claude's knowledge. A Research run is the live web, checked. To give you a sense of the scale this runs at: one documented run pulled **507 sources in about 13 minutes**. That's the 15-tabs-open hour of your day, done while you get coffee.

How to use it: turn on Research (the toggle near the message box) *before* you ask. Then ask the big question, not the small one.

■ Level 6 — Connectors

What it is: hooks that let Claude read and act inside the tools you actually work in — Notion, Slack, Google Drive, Gmail, and a growing list.

Why it matters: everything below this level works on copies — things you pasted or uploaded. Connectors put Claude on your real data. "What did we decide in Slack last week?" "Pull the numbers from the Drive sheet." No exporting, no pasting, no stale copy.

How to use it: open Settings, find Connectors, and add one. Start with the tool that holds most of your work.

■ Level 7 — Agents

What it is: the top of the ladder. You hand Claude a big task; it breaks the task into steps, runs the steps itself, and checks its own work along the way.

Why it matters: at every level below, you're the project manager: you feed the next step, it does the step. At level seven, it manages the steps. You define the outcome and it handles the middle. That's the difference between an AI that answers and an AI that works.

How to use it: give it a genuinely multi-step task, in plain words, with the outcome you want — then let it run instead of interrupting after step one.

■ Which level should you climb next?

Don't try to jump to seven. Find your line below, and climb exactly one level this week:

- **You only ever chat** → drag one file into the chat today. Level 2 takes ten seconds to learn.
- **You keep re-pasting the same context** ("here's my business, here's my client...") → make one Project and put it there. Level 3.
- **You ask for code or tools, then rebuild them elsewhere** → ask for an Artifact instead and let it run in the panel. Level 4.
- **You've got 15 research tabs open right now** → turn on Research and hand it the whole question. Level 5.
- **You copy things OUT of Notion/Slack/Drive to show Claude** → add that one tool as a Connector. Level 6.
- **You feed Claude one small step at a time all day** → hand it the whole task once and let it run. Level 7.

One level per week is fast. Most people gain more from climbing one level than from any new AI tool they'll install this month.

■ The bigger picture

The ladder is really one idea, applied harder at each level: **stop carrying context to the AI, and put the AI where the work is.**

- Levels 1–2: you bring pieces to it.
- Levels 3–4: it holds your context and builds real things.

- Levels 5–7: it goes out, works on live data, and runs the steps itself.

Every serious AI setup I build sits at the top of that ladder — the levels below are how you get comfortable enough to trust it there.

■ Your next step

Scroll back to the diagnosis list and pick the one line that described you. Do that one move today — it's a ten-minute climb, not a project.

The cost of skipping it: you keep paying (in money or time) for a seven-level tool and using one. That gap doesn't close on its own.

Built by Axionox. I'm Saad — I build real AI systems for a living, and this ladder is exactly how I move people from "I tried ChatGPT once" to running actual AI-powered workflows. If you want to see what the top levels look like wired into a real business, reply to the DM and I'll walk you through one. Follow for more like this.