

CORTEXA ACADEMY

Getting Started with **Claude**

A complete, step-by-step guide to setting up and mastering your Claude account — from your first chat to autonomous, agentic work.

1 Claude Pro · Web

2 Claude Cowork

3 Claude Code

How to use this guide

This guide takes you from zero to confident across the three ways people actually use Claude. You do not need any technical background. Work through it in order — each part builds on the one before, and the skills you learn in Part 1 carry all the way through.

THE SYMBOLS YOU'LL SEE



Numbered steps

Do these in order, exactly as written.



Prompt boxes

Dark boxes are word-for-word text you can type or paste into Claude.



Callouts

Tips, warnings, and things worth knowing before you click.



Screen sketches

Simplified drawings of what you'll see on screen — not live screenshots.



A note on the pictures

Claude's interface changes often — button names and colors may look a little different, but the steps stay the same. When in doubt, look for the label we name in the step.

THE THREE-PART JOURNEY

- **Claude Pro on the web** — your daily driver. Chat, files, Projects, memory, and personalization in the browser.
- **Claude Cowork** — Claude works with real files and completes multi-step tasks for you — now on desktop, web, and mobile.
- **Claude Code** — the most powerful surface. Automate work and build things — approachable even if you've never touched a terminal.

WOVEN THROUGHOUT

A running thread on **personal context files** and **prompt engineering** — the two habits that separate good results from great ones. Both get their own sections near the end.

What's inside

01 · Getting oriented — plans, the three surfaces, choosing a model	p. 4
02 · Part 1 — Claude Pro on the web — sign up, security, the interface, your first chat	p. 8
Projects, files & Artifacts · memory, connectors & Skills	p. 16
03 · Part 2 — Claude Cowork — the agentic workspace for real work	p. 22
04 · Part 3 — Claude Code — automation & building, made approachable	p. 29
05 · Your context files — AboutMe, style guide & CLAUDE.md	p. 36
06 · Prompt engineering — the habits that transform your results	p. 40
07 · The prompt library — copy-paste prompts for every job	p. 44
08 · Top Skills to try — the most-installed Skills & how to add them	p. 50
09 · Pitfalls & ethical use — hallucinations, privacy, and honest AI	p. 54
Quick reference & next steps — cheat sheet and the 4-week rollout	p. 57

THE ONE-MINUTE VERSION

- Subscribe to **Claude Pro (\$20/mo)**. It unlocks all three surfaces — no extra cost.
- Learn the **web app** first. Then add **Cowork** for file work, then **Claude Code** for automation.
- Write your **context files** once. Reuse them everywhere. This is the biggest single upgrade to your results.

Who this is for

Anyone with a Claude account who wants to go from “I typed a question once” to using Claude as a genuine work partner — marketers, researchers, analysts, and teams.

01

Getting oriented

Plans, the three surfaces, and choosing a model — the quick lay of the land before you dive in.

IN THIS SECTION

- The three ways to use Claude
- Which plan you need
- Which model to pick

The three ways to use Claude

Claude shows up in three places. They share one brain and one subscription, but each is built for a different kind of work. You'll learn them in this order — and everything transfers forward.

1 Claude Pro · Web

Type a question, get an answer — plus files, Projects, and memory, all in your browser.

BEST FOR

Everyday thinking, writing, analysis, research, and learning. The place you'll spend most of your time.

WHAT IT FEELS LIKE

Like texting the sharpest colleague you know.
Nothing to install.

2 Claude Cowork

Claude opens, reads, edits, and creates real files — and completes multi-step tasks for you.

BEST FOR

Multi-step jobs on real documents: turn a folder of notes into a report; build a spreadsheet and a deck.

WHAT IT FEELS LIKE

Like handing a project to an assistant who shows you the result — at your desk or from your phone.

3 Claude Code

A tool that runs on your computer to automate tasks and build things — no coding background required.

BEST FOR

Repeatable automations, working with lots of files at once, and anything that touches code.

WHAT IT FEELS LIKE

Like having an engineer on call. More power, a small learning curve.

★ The golden rule

Start in the web app and get comfortable. Move to Cowork when you want Claude to touch your real files. Reach for Claude Code when you want to automate something you do over and over. You do not need all three on day one.

Which plan do you need?

Good news: a single Claude Pro subscription unlocks all three surfaces in this guide. You only pay more for higher usage limits — not for extra features.

Plan	Price	What you get	Who it's for
Free	\$0	Claude Sonnet & Haiku, web search, memory, limited usage, up to 5 Projects	Trying Claude out for light, occasional questions
Pro RECOMMENDED	\$20/mo (\$17/mo billed yearly)	Everything in Free, plus Opus 5 , ~5x more usage, unlimited Projects, Research, Cowork , Claude Code , and the Microsoft 365 add-ins	Almost everyone. The sweet spot for individuals and knowledge workers
Max 5x	\$100/mo	All Pro features with roughly 5x the usage headroom	Heavy daily users who hit limits often
Max 20x	\$200/mo	All Pro features with the most usage headroom	Power users running Cowork or Claude Code all day
Team	\$25/seat (Standard) · \$125/seat (Premium)	Shared billing, central admin, collaboration; Premium seats carry Cowork/Code-scale usage. Annual: \$20 / \$100	Teams that want one bill and shared settings
Enterprise	Custom	SSO, compliance & audit controls, admin governance at scale	Large orgs with IT & compliance needs

📘 Start with Pro

Every step in this guide works on Pro. Only upgrade to Max if you repeatedly hit usage limits — which happens fastest when you lean on Cowork and Claude Code.

★ About “usage limits”

Claude limits how much you can use in a rolling window (plus a weekly cap), not a fixed number of messages. Pro gives you roughly five times the free allowance. Watch your usage in Settings.

Which model should you pick?

Inside Claude you can choose a “model” — think of it as choosing how much horsepower you want. All three below are included in Pro. Here's the simple version.

MOST CAPABLE

Claude Opus 5

The deepest thinker. Best for hard reasoning, nuanced writing, and the deliverables that really matter.

Reach for it when the work is important, complex, or you want the very best output.

BALANCED · DEFAULT

Claude Sonnet 5

Fast and highly capable — the everyday workhorse (and the default since June 2026) that handles the large majority of tasks well.

Reach for it almost always. Start here and only switch if you need to.

FASTEST

Claude Haiku 4.5

The quickest and lightest. Great for simple, high-volume, or throwaway tasks where speed wins.

Reach for it when you want a fast answer to something straightforward.

THE FRONTIER TIER: FABLE 5

Above Opus sits Anthropic's frontier model, **Fable 5**. As of July 20, 2026 it's no longer bundled into plan limits — paid plans buy it as **metered usage credits**, priced per use. Treat it as a deliberate splurge for the hardest work, and check your own account for current terms.

 Model lineups shift fast — what you see in your own account is always the source of truth.

★ The beginner rule of thumb

Default to **Sonnet 5**. Switch up to **Opus 5** for the work that matters most. Drop to **Haiku 4.5** for quick, simple tasks. You can change the model on any message — look for the model name near the message box.

02

PART 1

Claude Pro

Your daily driver: chat, files, Projects, memory, and personalization — all in the browser, nothing to install.

IN THIS PART

- Create your account & upgrade
- Lock it down & set privacy
- A tour of the screen & your first chat
- Projects, files & Artifacts
- Memory, connectors & Skills

Step 1 — Create your account

1 Go to the website

Open your web browser and visit **claude.ai**. You'll land on the sign-in screen shown here.

2 Choose how to sign in

Click **Continue with Google** to use your existing Google account (fastest — no new password to remember). Or click **Continue with email** and Claude will email you a secure code to enter.

3 Verify and set your name


Follow the prompt. If you used email, check your inbox for the code, type it in, and add your first name when asked.

4 You're in — on the Free plan

You now have a working Claude account. Next we'll upgrade to Pro so every feature in this guide is available.

No credit card needed to start

The Free plan doesn't ask for payment. You'll only enter a card when you choose to upgrade to Pro on the next page.



Continue with Google = fastest option

★ Setting this up for work? Sign in with your **work Google account** so the account can join a Team workspace cleanly later.

Step 2 — Upgrade to Pro

1 Open your account menu

Click your initials in the bottom-left corner of the screen. A small menu pops up.

2 Click “Upgrade plan”

Choose **Upgrade plan** from the menu. (You can also go straight to claude.ai/upgrade.)

3 Pick Pro


Select **Pro**. Choose monthly (\$20) or annual (\$17/mo, paid once a year) — annual saves about 15%.

4 Enter payment & confirm

Add your card and confirm. Pro activates immediately — you'll see “Pro” next to your name.

★ What just unlocked

Opus 5 access, far more usage, unlimited Projects, Research mode, and — importantly — **Claude Cowork** and **Claude Code**, which we cover in Parts 2 and 3.



Two clicks from the account menu

❗ Prefer to stay on Free for now?

You can — Part 1's web basics mostly work on Free. But Cowork, Claude Code, Opus, and Research need Pro, so upgrade before Parts 2 and 3.

Step 3 — Lock the account down

Claude's sign-in is passwordless — an emailed code or your Google account. That means your Claude security IS the security of the account behind it.

1 Secure the account behind Claude

Turn on **MFA (an authenticator app) or a passkey** on the Google account or email inbox you sign in with. Whoever controls that inbox controls your Claude.

2 Review active sessions


Settings → check where you're signed in; log out of devices you don't recognize.

3 On a Team/Enterprise plan?

Ask your admin to enforce **SSO** so access follows your company identity — and dies with it at offboarding.

⚠ Why bother?

Your Claude history becomes a diary of your work: drafts, plans, uploaded files, memory. An attacker in your inbox is an attacker in all of it. Treat the sign-in account like the vault it now is.



Passwordless sign-in means **THIS page is your Claude security page**

📘 This slide is the one piece of setup that happens *outside* claude.ai — and it's the highest-leverage two minutes in the guide.

Step 4 — Decide what Anthropic may do with your data

1 Open Privacy settings

Your initials → **Settings** → **Privacy**.

2 Decide on “Help improve Claude”


ON = your chats may be used to train future models, with multi-year retention. OFF = they aren't, and standard deletion windows apply. Turning it off removes **no** features.

3 Know the boundaries

Team, Enterprise, and API traffic are excluded from training **contractually**. Incognito chats are never saved or trained on either way.

★ The plain rule

Decide **once, now**, before the account fills with work. If your chats will ever contain client or confidential material on a personal plan, turn it off today — and put truly sensitive one-offs in Incognito.




The toggle to decide deliberately — not by default

Account type	Training
Free / Pro / Max	Your choice — the toggle above
Team / Enterprise / API	Excluded — contractually

A tour of the screen

Before your first chat, let's get familiar with what you're looking at. Here's the whole screen with the five things you'll use most, labeled.

- 1 New chat**
Starts a fresh conversation. Each chat is separate — a clean slate.
- 2 History**
Every past conversation, newest first. Click any to reopen; hover for a “...” menu to rename, star, or delete.
- 3 Projects**
Reusable workspaces that remember files and instructions. Covered in depth shortly.
- 4 Add / tools**
The + attaches files. Nearby toggles turn on web search and Research.
- 5 Model**
Switch between Opus, Sonnet, and Haiku right here, per message.



The five things you'll use most, labeled

★ Just getting started?

You only truly need two things to begin: the message box in the middle and the Send button. Everything else you'll add as you go.

Your first chat

Time to actually talk to Claude. Click in the message box, type, and press Enter. Here are three starter prompts — type any of them word-for-word to see what Claude can do.

>_ TRY THIS FIRST

[COPY / PASTE](#)

I'm new to Claude. In plain language, explain three ways you could help me with my work this week. Ask me one question about my job first so your answer is specific to me.

>_ LEARN SOMETHING FAST

[COPY / PASTE](#)

Explain **[topic you're curious about]** to me as if I'm smart but completely new to it. Use a simple analogy, then give me three things most people get wrong.

>_ TURN A MESS INTO A PLAN

[COPY / PASTE](#)

Here's a rough brain-dump of everything on my plate: **[paste your notes]**. Turn it into a short summary, a prioritized to-do list, and three risks I might be missing.

THREE TINY HABITS THAT HELP IMMEDIATELY

- **Give context.** Tell Claude who you are and who the work is for.
- **Say the format.** “In five bullets,” “as a table,” “one paragraph.”
- **Keep going.** Reply “shorter,” “more formal,” “add an example” — don't restart.

Enter sends, Shift+Enter makes a new line

Hold Shift while pressing Enter to add line breaks without sending.

The anatomy of a great prompt

You'll get dramatically better results if your request has a few key ingredients. You don't need all of them every time — but the more you include, the sharper the output. We go deeper in the prompt-engineering section.

- R** **Role — tell Claude who to be.**
“You are a senior market researcher...” Sets the tone and expertise.
- C** **Context — the situation, audience, and goal.**
“This is for a client pitch to a tech brand...”
- T** **Task — exactly what you want done.**
“Draft a one-page competitive summary...” Specific and direct.
- E** **Examples — show a sample of the style.**
“Match the tone of this: [paste]” — the single biggest lever.
- F** **Format — name the shape of the output.**
“Five bullets, then a one-line recommendation.”

>_ ALL FIVE, COMBINED

[COPY / PASTE](#)

You are a senior market researcher. I'm preparing a client pitch to a large PC brand and need to look sharp. Draft a one-page competitive summary of the three leading laptop makers. Match the crisp, factual tone of this example: [paste]. Format: a short intro, a comparison table, then three bullets on what it means for us.

An easy shortcut

Can't be bothered to structure it? Just add: “**Ask me any questions you need before you start.**” Claude will fill the gaps itself.

Projects: your reusable workspaces

A Project is a folder that remembers things. Instead of re-explaining your context every time, you set it up once — the files, the rules, the tone — and every chat inside the Project inherits it. This is the single biggest time-saver in the web app.

1 Open Projects

In the left sidebar, click **Projects** (or go to claude.ai/projects), then **+ New Project**.

2 Name and describe it


A clear name like “Market Research” and a one-line description of what it's for.

3 Add knowledge

Upload the files Claude should always know about — brand guides, past reports, personas. Claude reads them in every chat, no re-uploading.

4 Set instructions

Standing rules: who Claude is, your audience, tone, do's and don'ts. These apply to every conversation here.



Set up once — every chat inside inherits it

★ When to use a Project vs. a plain chat

Use a Project for any recurring stream of work with the same context (a client, a research area, a product). Use a plain chat for quick one-offs. Pro gives you unlimited Projects — and Project instructions are the perfect home for the style guide and AboutMe files you'll build in Section 05.

Files in, Artifacts out

PUTTING FILES IN

Click the **+** in the message box to attach PDFs, Word docs, spreadsheets, images, and slides. Claude reads the contents — summarize, analyze, rewrite, or pull data out.

> TRY WITH A FILE ATTACHED

[COPY / PASTE](#)

Summarize this report in five bullets for a busy executive, then give one clear recommendation. **[attach the file first]**

★ Edit by chatting

“Make it shorter,” “add an FAQ,” “change the heading” — Claude revises the Artifact in place, not the whole chat, keeping each version.

GETTING POLISHED WORK OUT: ARTIFACTS

When Claude makes something substantial — a document, a table, a web page, a diagram — it opens in a side panel called an **Artifact**. It's a living document: tell Claude what to change and it updates in place.

Your chat

WELCOME GUIDE — the Artifact

“Write a one-page welcome guide for new hires.”

Done! I've created it in the panel on the right → edit it by telling me what to change.

The Artifact — edit it live

i Share it

Click **Publish** for a link anyone can open, or download it as a Word, PowerPoint, Excel, or PDF file.

Teach Claude about you

Fill this in once and every future chat gets better. Your profile preferences tell Claude how you like to work — and they carry across chats and into Cowork too.

1 Open Settings

Click your initials (bottom-left) → **Settings** → **Profile**.

2 Add your name & preferences


In “What personal preferences should Claude consider?” describe how you want Claude to respond — as **instructions, not adjectives**.

3 Save

That's it. Claude now applies this to every conversation automatically.

Write rules, not personality traits

“I'm detail-oriented” does nothing. “Double-check every number before presenting it” is an instruction Claude can actually follow. Turn each trait into a behaviour.



This follows you everywhere

THE THREE LAYERS OF PERSONALIZATION

- **Profile preferences** — apply everywhere (this page).
- **Project instructions** — apply inside one Project.
- **The prompt itself** — one message. Most specific wins.

Memory & private chats

Claude can remember useful things about your work across conversations, so you repeat yourself less over time. You're always in control — and you can switch it off whenever you like.

HOW MEMORY WORKS

- Claude quietly builds a summary of what matters from your past chats and uses it to personalize future ones. It tells you when it's drawing on a memory.
- **Manage it:** Settings → Memory — view the summary, edit it, or Reset it (permanent).
- **Ask it directly:** type **“What do you remember about me?”** any time.
- **Per-project memory** keeps each Project's memory separate, so work streams don't bleed together.

Coming from another AI tool?

Claude can import your memories from ChatGPT, Gemini, or Grok. Look under Settings → Memory for **Memory Import** — it takes about a day to process.

INCOGNITO CHATS

- Need a one-off conversation that leaves no trace? Click the **ghost icon** at the top-right of a new chat.
- Incognito chats aren't saved to your history, aren't added to memory, and aren't used for training.
- Use them for sensitive or throwaway questions — the **first** click, not cleanup afterward.

★ Good habit

Keep truly sensitive information (ID numbers, card details, anything confidential) out of memory-enabled chats entirely.

Connect your other tools

Connectors let Claude work with your real accounts — your calendar, your documents, your inbox — from inside a chat. You approve each connection, and Claude asks before it does anything.

1 Open Connectors

Go to **Settings → Connectors**, or click the **+** in the chat box and choose Connectors.

2 Connect an account

Pick a service (say, Google Calendar), click **Connect**, and sign in through the official pop-up. You can revoke access any time.

3 Turn it on per chat

Enable the connector only in the chats where you need it — this keeps things fast and private.

★ A safety note worth knowing

Gmail is **read-and-draft only** — Claude can write a draft but cannot send email for you. Start every connector read-only and watch how it behaves before granting more.

POPULAR CONNECTORS FOR KNOWLEDGE WORK

Google Calendar	read & write
Google Drive	read & save files
Gmail	read & draft only
Slack	read & post
Microsoft 365	via add-ins
HubSpot / Jira	read & update

i Want true “when X happens, do Y” automation?

Connectors act on demand inside a chat. For event-triggered automation across thousands of apps (including Meta & WhatsApp), use a bridge like Zapier — more in the prompt library.

Skills & getting great results

SKILLS: CAPABILITIES CLAUDE PICKS UP AUTOMATICALLY

A **Skill** is a packaged ability Claude uses when it's relevant — without you doing anything special. Claude ships with built-in Skills for creating Excel, PowerPoint, Word, and PDF files. Just ask naturally:

> **THE PPTX SKILL RUNS AUTOMATICALLY**

COPY / PASTE

Create a five-slide PowerPoint summarizing these Q3 results, one clear takeaway per slide. **[attach your data]**

Turn Skills on

Go to **Customize → Skills**. Make sure **Code execution & file creation** is enabled (that's what powers file-making Skills). Much more in Section 08.

ETIQUETTE FOR GREAT RESULTS

- 1 Front-load context.**
The first message sets the whole conversation.
- 2 Let Claude ask.**
For anything big, add “ask me 1–2 questions first.”
- 3 Iterate, don't restart.**
Refine with follow-ups rather than a brand-new chat.
- 4 One chat, one topic.**
Start a new chat when you switch subjects.
- 5 Save Opus for what matters.**
Sonnet handles most things; switch up when the stakes are high.

★ You've finished the web basics

You can now chat, attach files, build Projects, make Artifacts, personalize Claude, and connect your tools. Next: letting Claude do real work with real files.

03

PART 2

Claude Cowork

Move from answers to action. Cowork works directly with real files and completes multi-step tasks — now on desktop, web, and mobile, with cloud sessions that keep working after you close the lid.

IN THIS PART

- From answering to doing
- Set up Cowork
- How it gets work done
- Real jobs to try
- Staying in control

From answering to doing

In the web app, Claude tells you how to do something. In Cowork, Claude does it — opening, reading, editing, and creating real files, then handing you the finished result. It's included in Pro.

REGULAR CHAT

You ask a question

Claude answers with text and advice. To act on it, you do the work — copy things across, open the files, make the edits. Great for thinking & drafting.

CLAUDE COWORK

You describe an outcome

Claude makes a plan, then works directly in the folders you allow — producing finished spreadsheets, decks, and documents. Great for multi-step tasks on real files.

NEW SINCE JULY 2026

- **Everywhere:** desktop (macOS/Windows), the **web** at claude.ai, and **iOS/Android**.
- **Cloud sessions:** hand off a task and close the laptop — it keeps running remotely.
- **Scheduled tasks** fire even with no device online.

How it's different from Claude Code

Cowork and Claude Code share the same powerful engine. Cowork is built for non-technical knowledge work with no terminal — documents, spreadsheets, research, slides. Claude Code (Part 3) is aimed at automation and building.

Set up Cowork

1 Pick your door

Web: open Cowork right on claude.ai — nothing to install. **Desktop:** get the Claude app at claude.com/download (macOS or Windows) for work on local folders.

Mobile: the iOS/Android app can start and steer tasks.

2 Sign in


Same Pro account you use on the web.

3 Switch to Cowork

At the top you'll see the modes: **Chat**, **Cowork**, **Code**. Click Cowork.

4 Grant access

Describe your task, then choose which folder(s) or cloud files Claude may use. It only touches what you allow — nothing else.



Local desktop sessions use your folders; cloud sessions run remotely

★ Local vs. cloud


Working on files that live on your machine? Use the **desktop app** — and keep it open while a local task runs. Working on cloud files, or want it to keep going after you leave? Run the session **in the cloud** from web or mobile.

① Want this step by step?

Our free **Claude Cowork Setup Guide** walks Mac and Windows through the full first hour — install, global instructions, memory files and connectors, with copy-paste templates: cortexa-consulting.ai/cortexa-claude-cowork-setup-guide.html

How Cowork gets work done

When you hand Cowork a task, it doesn't just fire off one response. It works the way a capable assistant would — planning, breaking the job into pieces, doing the work, and checking in with you.



It plugs into everything you set up

Cowork uses your **connectors** (pull from Slack, save to Drive), your **Skills** (build Office files), and even sub-agents that work in parallel. Give it global instructions once, and per-folder instructions for specific projects.

★ Set it and (actually) forget it

Type **/schedule** inside a task to run it on a cadence — a Monday briefing, a weekly cleanup. Since the July 2026 update, scheduled tasks run **in the cloud**, so they fire even when your computer is off.

 Steer from anywhere: a running task can be checked and redirected from the mobile app — hand off at your desk, correct course from the train.

Real jobs to try first

The best way to learn Cowork is to give it a small, low-stakes job. Six real-world tasks — each with a prompt you can adapt. Point Cowork at the right folder and let it work.

TIDY

Messy folder

“Organize this Downloads folder into subfolders by type and date, and flag any duplicates before deleting anything.”

EXTRACT

Receipts → spreadsheet

“Turn these receipt images into a tidy expense spreadsheet with date, vendor, category, and total, then sum by category.”

SYNTHESIZE

Notes → report

“Read every file in this folder of meeting notes and draft a Q1 product-update report with sections and a summary.”

PRESENT

Data → deck

“Analyze this spreadsheet, find the three biggest trends, and build a clean PowerPoint that tells the story.”

SCHEDULE

Daily briefing

“Each morning, pull my unread Slack highlights and today's calendar into a one-page brief saved to this folder.”

INDEX

Organize documents

“Sort these scattered files into a clear, chronological set and produce an index I can hand to a colleague.”

Start small to build trust

Begin with something reversible — organizing a folder or drafting a doc — and watch how Cowork asks permission and shows its plan. Once you trust its judgment, hand it bigger jobs.

Staying in control

Cowork is powerful because it can act on your files — so it's built to keep you in charge. Here's how the safety model works, and the handful of limits to keep in mind.

YOU'RE ALWAYS IN CHARGE

- ✓ **You choose the folders.** Claude can only see and touch what you explicitly allow.
- ✓ **It asks before big moves** — and always before permanently deleting anything.
- ✓ **It shows its plan before executing**, so there are no surprises.
- ✓ **Work runs in a sandbox** (an isolated environment) for an extra layer of safety.

★ Anthropic's own advice

Always confirm before Claude handles financial, personal, or work-critical tasks. Treat the approval prompts as your seatbelt.

GOOD TO KNOW

- **No memory between sessions.** Each Cowork task starts fresh — put lasting context in folder instructions.
- **Local desktop sessions** still need the app open and the machine awake; **cloud sessions don't** — pick accordingly.
- **It uses your limits faster than chat**, because it does much more per task.



One firm rule

For regulated workloads (health or financial records that must be audited), use Cowork only where your compliance team has signed off — **Enterprise tiers now carry tamper-evident audit logs**; personal plans don't.

04

PART 3

Claude Code

The most powerful surface, made approachable. Automate the repetitive and build real things — in plain English.

IN THIS PART

- Why it's worth it
- Install Claude Code
- Your first session
- Plan mode
- CLAUDE.md & power features

The most powerful surface

Claude Code runs on your computer and can work across many files at once, automate repetitive jobs, and build things — from a data-cleaning script to a small website. It has a reputation as a “developer tool,” but a curious beginner can absolutely use it.

1

You can start without a terminal

The Claude desktop app has a **Code tab** — a friendly window with Claude Code's power, no commands to type. It even has a built-in sandboxed browser now, so Claude can look at web pages while it works. If terminals intimidate you, start here.

2

You talk to it in plain English

You don't write code — you describe what you want. Claude Code figures out the steps, shows you what it's about to do, and asks permission before changing anything. (Asking first is now the **default mode**.)



Two ways in, same power

This part shows the terminal path (the classic way, worth learning) and notes the desktop Code-tab alternative at each step. Pick whichever feels comfortable — you can switch later.

★ Is this for you?

If you have a repetitive task involving lots of files (renaming hundreds of photos, reformatting data, generating reports from a template), or you're curious about building small tools — yes. If you just want to chat and make documents, Parts 1 and 2 already have you covered.

Install Claude Code

1 Install Node.js (one-time)

Claude Code needs the free Node.js runtime (version 18+). Download from nodejs.org and run the installer. Skip this if you use the desktop Code tab.

2 Install Claude Code

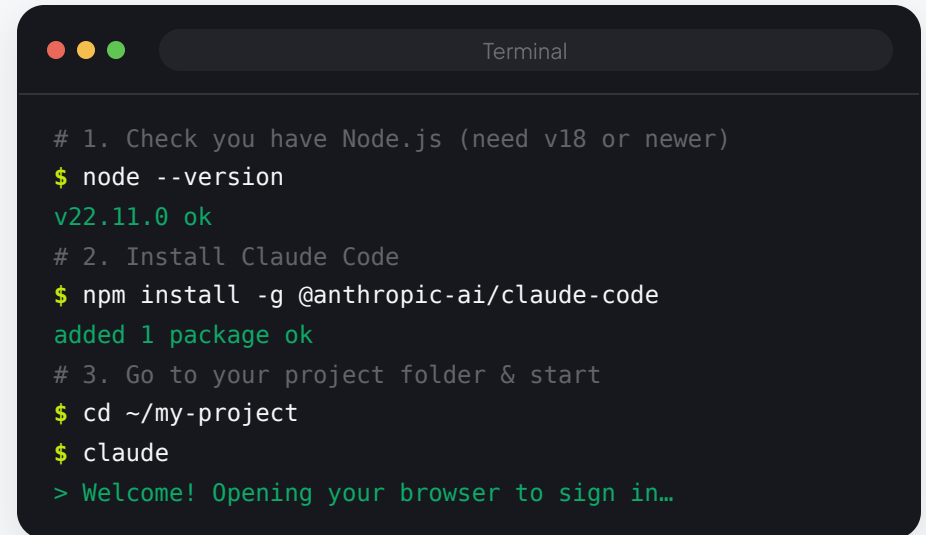
Open your Terminal app and paste the install command shown here, then press Enter. It takes a few seconds.

3 Open a project & run it

Navigate to a folder you want to work in (`cd` means “change directory”), then type `claude` and press Enter.

4 Sign in

Your browser opens — log in with your Pro account. Done. No API key or extra billing needed.

A terminal window with a dark background and light green text. The title bar says "Terminal". The content shows a series of commands and their outputs: 1. Check Node.js version: \$ node --version returns v22.11.0 ok. 2. Install Claude Code: \$ npm install -g @anthropic-ai/claude-code returns added 1 package ok. 3. Go to project folder and start: \$ cd ~/my-project, then \$ claude returns > Welcome! Opening your browser to sign in...

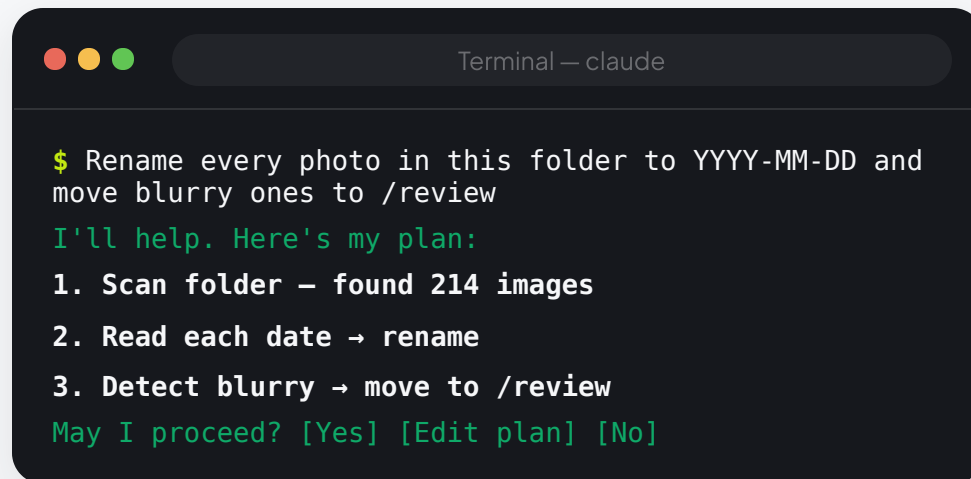
```
# 1. Check you have Node.js (need v18 or newer)
$ node --version
v22.11.0 ok
# 2. Install Claude Code
$ npm install -g @anthropic-ai/claude-code
added 1 package ok
# 3. Go to your project folder & start
$ cd ~/my-project
$ claude
> Welcome! Opening your browser to sign in...
```

Prefer no terminal at all?

Open the Claude desktop app, click the **Code tab**, and pick a folder. You get the same engine in a point-and-click window — skip steps 1–3 entirely.

Your first session

Once you're in, it's a conversation — just like the web app, except Claude can act on the files around it. Type what you want in plain English. Claude proposes a plan and waits for your yes.



```
Terminal — claude

$ Rename every photo in this folder to YYYY-MM-DD and
move blurry ones to /review
I'll help. Here's my plan:
1. Scan folder — found 214 images
2. Read each date → rename
3. Detect blurry → move to /review
May I proceed? [Yes] [Edit plan] [No]
```

Asking first is the default — it waits for your yes

A FEW COMMANDS WORTH KNOWING


Think & plan before touching anything	<code>/plan</code>
Switch between Opus, Sonnet, Haiku	<code>/model</code>
Create a CLAUDE.md memory file	<code>/init</code>
See everything you can do	<code>/help</code>

★ Keep a safety net

Before big changes, make a copy of your folder (or use Git if you know it). Claude asks before editing, but a backup means you can always undo.

Plan mode: look before you leap

The single best habit in Claude Code is Plan mode. It tells Claude to explore and propose a plan first — without changing any files — so you approve the approach before work begins. It prevents almost every “wait, that’s not what I wanted” moment.



i Use it for anything unfamiliar
New folder, big change, or a task you've never done? Plan mode first. For tiny, obvious edits you can skip it.

>_ TURN IT ON

COPY / PASTE

Type `/plan` to enter Plan mode — or press `Shift+Tab` to cycle modes. Then describe your task. Claude will plan without editing until you approve.

CLAUDE.md: your project's memory

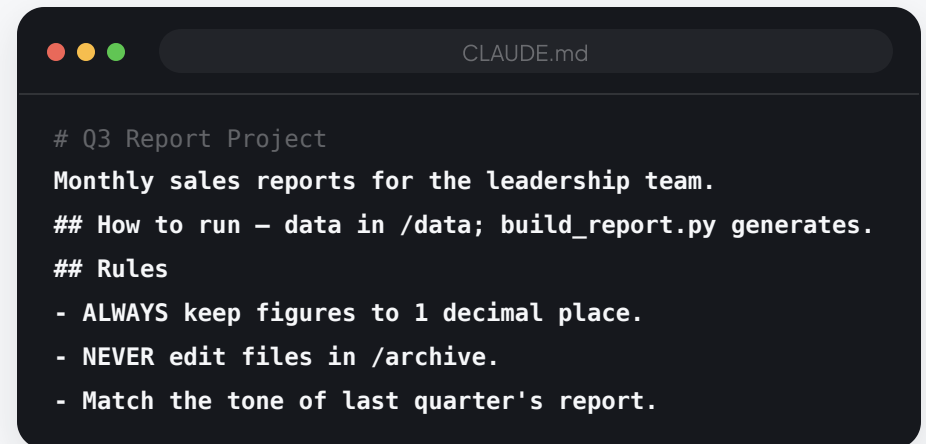
A CLAUDE.md file is a plain text note that lives in your project folder. Claude Code reads it at the start of every session — so it always knows the project's rules without you re-explaining. It's the Claude Code equivalent of the profile preferences you set in Part 1.

CREATE ONE IN TWO SECONDS

Inside a session, just type `/init` — Claude looks at your folder and drafts a starter CLAUDE.md. Then trim it — keep only what matters.

WHAT TO PUT IN IT

- A one-line description of the project
- The commands to run, test, or build it
- Key folders and what they're for
- “Always” and “never” rules Claude can't guess



```
# Q3 Report Project
Monthly sales reports for the leadership team.
## How to run - data in /data; build_report.py generates.
## Rules
- ALWAYS keep figures to 1 decimal place.
- NEVER edit files in /archive.
- Match the tone of last quarter's report.
```

★ Keep it short & high-signal

Aim for under a page. For each line ask: “would removing this cause a mistake?” If not, cut it. Add IMPORTANT or ALWAYS/NEVER for the rules that really matter.

Four ideas to grow into

You don't need these on day one — but as you get comfortable, these four capabilities turn Claude Code from a helper into a genuine force-multiplier.

PACKAGED ABILITIES

Skills

The same packaged abilities from the web app work here too — and you can write your own by dropping a file in a `.claude/skills` folder. Claude uses them automatically.

PARALLEL HELPERS

Subagents

Claude can spin up specialist helpers — one for research, one for review — each focused on its own job, then bring the results together. Great for bigger tasks.

REAL DATA

MCP connectors

Connect Claude Code to outside tools — a database, your calendar, a project tracker — so it can pull and push real data as it works.

EVERY TIME

Hooks

Rules that fire automatically at set moments — for example, “always tidy the formatting after every edit.” Use them when something must happen every time.

The learning path that works

Use Claude Code for a week with just plain requests and Plan mode. Add a `CLAUDE.md` once you repeat yourself. Explore Skills and connectors only when you hit a real need. Let the tool grow with you.

05

Your context files

The highest-leverage habit in the guide: write a few small files once, then reuse them across all three surfaces.

IN THIS SECTION

- Write it once, reuse everywhere
- The three files worth having
- Let Claude interview you
- Anatomy of a style guide

Write it once, reuse it everywhere

Instead of re-explaining who you are and how you work in every chat, write a few small context files once — and feed them into all three surfaces. The result: Claude sounds like it already knows you, from the very first message.

THE THREE FILES WORTH HAVING

- 1 AboutMe**
Who you are, your role, how you like to communicate, and what to always or never do.
- 2 Style guide**
Your voice: words to use, words to avoid, tone, and on-brand vs. off-brand examples.
- 3 CLAUDE.md**
Project-specific rules for Claude Code (and a great model for Cowork folder instructions).

i The golden rule for all three
Write instructions, not adjectives. “I’m detail-oriented” is wasted breath. “Double-check every number before presenting it” is something Claude can actually act on.

WHERE EACH FILE GOES

AboutMe	Profile preferences — applies to every chat & Cowork
Style guide	Project knowledge — or keep as a Skill
CLAUDE.md	Project folder — mirror it in Cowork folder instructions

★ The magic: these **layer**. Your AboutMe sets the baseline everywhere; a Project's instructions add specifics; a single prompt can override both. Most specific always wins.

Let Claude interview you

Here's the trick that makes this effortless: don't write these files yourself. Have Claude interview you one question at a time, then write the file from your answers. It takes about five minutes. Paste one of these to begin.

>_ BUILD YOUR ABOUTME PROFILE

[COPY / PASTE](#)

Interview me one question at a time — wait for each answer before the next — to build a personal profile you'll use to tailor all your future responses to me. Ask about my role, my audience, how I like to communicate, my tone, my formatting preferences, and things you should always or never do. If an answer is vague, push me to be specific. After about 8–10 questions, produce a tight, one-screen “AboutMe” file I can paste into my Claude profile preferences.

>_ BUILD A CLAUDE.MD (IN CLAUDE CODE)

[COPY / PASTE](#)

I want to generate a CLAUDE.md file for this project. Interview me to understand what you'd need to work effectively here — ask about the project, the people involved, and how I want things done. Then draft the file from my answers.

>_ BUILD YOUR BRAND-VOICE STYLE GUIDE

[COPY / PASTE](#)

You are a conversion copywriter and an expert in persuasion and human behaviour. Interview me one question at a time about my brand's voice. Then produce a style guide with: a short intro; 3–4 voice principles (each with what it means, how it shapes writing, an example, and what NOT to do); a vocabulary list (words to use / words to avoid); tone; and cadence.

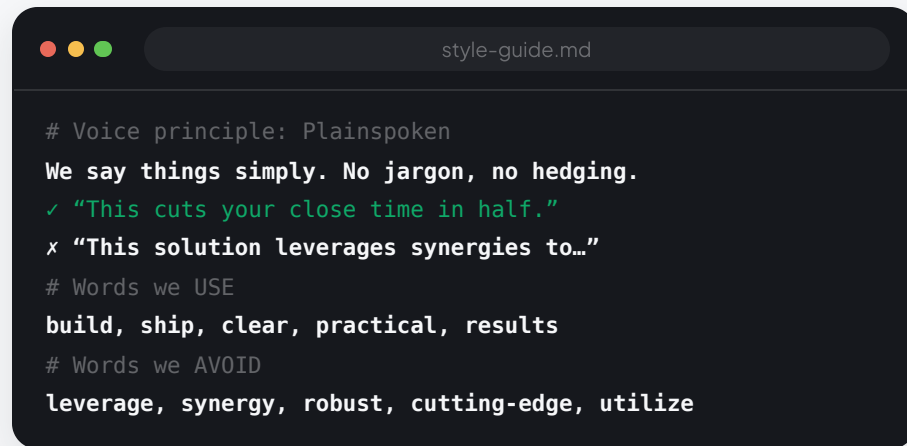
★ The one setting that makes this work

In Claude Code, turn on **Plan mode** first. Without it, Claude tends to rush and write the file after one or two answers — too thin to be useful. Plan mode keeps it asking.

❗ **Keep them alive.** When you correct Claude the same way twice, add that rule to the file. Re-run the interview every few months as your work changes.

Anatomy of a style guide

The style guide is worth extra attention — it's what stops Claude from sounding generic. A strong one is built from specific rules and word lists, not vague adjectives, and it always includes examples of what not to do.



```
# Voice principle: Plainspoken
We say things simply. No jargon, no hedging.
✓ "This cuts your close time in half."
x "This solution leverages synergies to..."

# Words we USE
build, ship, clear, practical, results

# Words we AVOID
leverage, synergy, robust, cutting-edge, utilize
```

THE FIVE SECTIONS THAT MATTER

- 1 Voice principles.**
Three or four, each with a one-line meaning and a “do this / not that” pair.
- 2 Vocabulary.**
Explicit lists: words to use, words to avoid. This does the heaviest lifting.
- 3 Tone.**
How you sound in different situations — confident, warm, direct.
- 4 Cadence.**
Rhythm and length rules — sentence length, paragraph shape, bullets vs. prose.
- 5 Examples.**
A few on-brand and off-brand samples. Claude learns fastest from real pairs.



Put it to work immediately

Once written: **“Using the attached style guide, rewrite this draft in our voice, then list every change and why.”**

06

Prompt engineering

No tricks or magic words — just the handful of habits that make every result noticeably sharper. Treat Claude like a brilliant new hire who has zero context on your project.

IN THIS SECTION

- The habits that transform results
- Structure & the art of iterating
- Prompting reasoning models

The habits that transform results

1

Be clear & direct

- ✗ “Write about our product.”
- ✓ “Write a 100-word product description for first-time buyers, focus on ease of setup.”

2

Give context

- ✗ “Summarize this.”
- ✓ “Summarize this for a busy CEO deciding whether to fund the project.”

3

Show examples

- ✗ “Make it punchy.”
- ✓ “Match the tone of these two headlines: [paste]. Now write five more.” **The single biggest lever.**

4

Ask it to think

- ✗ “Is this a good deal?”
- ✓ “Think through the pros, cons, and risks step by step, then give a verdict.”

5

Assign a role

- ✗ “Give feedback on this.”
- ✓ “You are a veteran editor. Give feedback on this draft...”

6

Name the format

- ✗ “Tell me about options.”
- ✓ “Give three options as a table: name, cost, one-line trade-off.”

❗ You won't use all six every time — even one or two lifts your results noticeably. For anything important, stack several. And remember the universal shortcut: **“Ask me any questions you need before you start.”**

Structure & the art of iterating

GIVE LONG PROMPTS A SKELETON

When a request has several parts, **label them**. Claude pays special attention to clearly marked sections — plain headings or simple tags like `<context>` and `<task>` remove ambiguity about what's background and what's the actual instruction.

>_ A STRUCTURED PROMPT

[COPY / PASTE](#)

`<context>`

We're a B2B agency pitching a retail client. Tone is confident but not salesy.

`</context>`

`<task>`

Write a 3-paragraph intro email requesting a 30-minute call.

`</task>`

`<format>`

Subject line, then body. Under 150 words.

`</format>`

ITERATE — DON'T START OVER

- “Shorter — half the length.”
- “More formal, and drop the exclamation marks.”
- “Add a concrete example to point 2.”
- “Give me three more options in the same style.”
- “What would make this stronger? Then apply it.”

★ One habit to drop

Skip “ignore previous instructions” and other jailbreak-style tricks — they just muddy the request. If something legitimate gets refused, add clear, honest context about why you need it instead.

Reasoning models are briefed, not micromanaged

Opus 5 and Sonnet 5 are concise and direct by default, and plan their own approach on hard problems. Your job shifts from steps to brief: goal, constraints, success criteria — then let them work.

	Quick tasks	Hard problems (extended thinking)
Specify	Detailed guidance helps	Goal, constraints, success criteria
Skip	—	Over-scripted procedures — they cap a better plan than yours
Add	Format & examples	“Flag assumptions rather than guessing” · what a *good* answer contains

★ Want to see the thinking? Just ask: “**show your reasoning.**” For hard problems, turn on **extended thinking** — and audit the logic, not just the verdict.

>_ A GOOD THINKING BRIEF

[COPY / PASTE](#)

Evaluate whether we should switch our reporting stack to **[option]**. Constraints: budget under \$X, no downtime during Q4, team of 3. A good answer weighs cost, migration risk, and lock-in, and ends with a recommendation plus the two strongest counterarguments. Flag any assumption you had to make.

i This is exactly how you brief Cowork on a big task, and exactly what belongs in a CLAUDE.md. One skill, every surface.

07

The prompt library

Copy-paste prompts for the jobs people hand to Claude most — adapt the bracketed bits and go.

IN THIS SECTION

- Everyday knowledge work
- Marketing & market research
- Competitive & technology intelligence
- Automation for your everyday apps
- What's realistic today

Everyday knowledge work

These are the workhorses — useful in almost any role, on almost any day. Text in [brackets] is yours to fill in.

>_ CHIEF-OF-STAFF BRAINDUMP

[COPY / PASTE](#)

You are my chief of staff. Here's a messy braindump of everything on my plate: **[paste]**. Turn it into (1) a one-paragraph summary, (2) a prioritized action list with owners and rough due dates, and (3) three risks I'm not seeing. Ask up to two clarifying questions first if you need to.

>_ EXECUTIVE SUMMARY

[COPY / PASTE](#)

Summarize this document for a busy executive in five bullets, then add one "so what" recommendation line. Document: **[paste or attach]**.

>_ DECISION HELPER

[COPY / PASTE](#)

Help me decide between **[option A]** and **[option B]**. Lay out the trade-offs step by step against these criteria: **[list]**. Then give a clear recommendation and say what would change your mind.

>_ LEARN ANYTHING FAST

[COPY / PASTE](#)

Explain **[topic]** as if I'm smart but completely new to it. Start with a one-sentence analogy, then the three things most people misunderstand, then one thing that will make me look like I know it well.

>_ DRAFT & PRESSURE-TEST

[COPY / PASTE](#)

Draft **[the thing]** for **[audience]**. Then switch roles: critique your own draft as a skeptical reader would, and give me the improved version.

>_ MEETING PREP

[COPY / PASTE](#)

I have a meeting with **[who]** about **[topic]**. Give me a one-page prep: their likely goals, three questions they may ask, my strongest talking points, and one risk to avoid.

Marketing & market research

Two of the most common jobs people hand to Claude. Attach your style guide from the context-files section for on-brand output.

MARKETING

>_ TARGETED EMAIL

[COPY / PASTE](#)

Write a marketing email for **[product]**. Audience: **[segment]**. Highlight three benefits: **[list]**. Tone: professional yet approachable. Output a subject line under 50 characters, a 150–200 word body, and one clear call-to-action.

>_ ON-BRAND REWRITE

[COPY / PASTE](#)

Act as my brand-voice editor. Using the attached style guide, rewrite this draft to match our voice, then list every change you made and why: **[paste]**.

>_ CAMPAIGN CONCEPTS

[COPY / PASTE](#)

Give me three distinct campaign concepts for **[product / launch]**. For each: a big idea, a tagline, the core channel, and why it fits **[audience]**. Then recommend one.

MARKET RESEARCH

>_ MARKET SIZING

[COPY / PASTE](#)

Size the market for **[category]** in **[region]**. Show your assumptions, give both a top-down and a bottom-up estimate, present TAM / SAM / SOM in a table, cite sources, and flag your weakest assumption.

>_ THEME & SENTIMENT ANALYSIS

[COPY / PASTE](#)

Analyze these **[reviews / survey responses]**. Extract the top themes, rate sentiment, quantify how often each theme appears, and surface three non-obvious insights. Data: **[attach]**.

>_ PERSONA BUILDER

[COPY / PASTE](#)

From this research, build three customer personas. For each: goals, frustrations, buying triggers, and the one message most likely to land. Research: **[paste/attach]**.

Competitive & technology intelligence

Turn Claude's web search and Research mode into a standing intelligence function. Turn on Research for the deep ones.

>_ COMPETITIVE TEARDOWN

[COPY / PASTE](#)

Build a competitive teardown of **[competitor]** versus us. Cover positioning, pricing, key features, launches in the last 12 months (with sources), and strategic gaps we can exploit. Output a comparison table plus a five-bullet “what this means for us.”

>_ TECHNOLOGY SCAN

[COPY / PASTE](#)

Give me a briefing on **[emerging technology]**: what it is in plain English, who the serious players are, the state of maturity, and the two or three ways it could affect **[our industry]** in the next two years. Cite sources.

>_ SIGNAL TRACKING

[COPY / PASTE](#)

Search the web for **[competitor or technology]** news from the last 90 days. Cluster it into themes and tell me what's accelerating and what's stalling, with citations for each claim.

>_ WIN/LOSS SYNTHESIS

[COPY / PASTE](#)

Here are notes from recent deals we won and lost: **[paste]**. Find the patterns. Why do we win, why do we lose, and what are the three highest-impact changes to our pitch?

Make it repeatable

Drop any of these into a Project with your competitor list as knowledge, or schedule a monthly version in Cowork with /schedule — and you've built yourself a recurring intelligence report.

Automation: your everyday apps

These use connectors and Skills. Enable the relevant connector in the chat first, and Claude will ask before it acts.

MICROSOFT 365

>_ EXCEL → POWERPOINT

[COPY / PASTE](#)

In this open Excel workbook, identify the key revenue drivers and any anomalies, then build a summary deck in the open PowerPoint using our template. Keep each bullet to one line.

GOOGLE WORKSPACE

>_ INBOX TRIAGE (DRAFT-ONLY)

[COPY / PASTE](#)

Read my Gmail from this week and group it into (1) reply today, (2) can wait, (3) FYI. Draft replies as Gmail drafts for group 1 only. Do not send anything.

TEAMS & SLACK

>_ THREAD TO ACTION ITEMS

[COPY / PASTE](#)

Summarize this Slack/Teams thread into three lists: decisions made, action items with owners, and open questions still unresolved.

CRM — HUBSPOT / SALESFORCE

>_ FOLLOW-UP DRAFTING

[COPY / PASTE](#)

List HubSpot contacts in the “Demo Scheduled” stage with no activity in five days, and draft a friendly check-in email for each. Show me the drafts before doing anything else.

★ Always start with approval on

For any tool that can change data (CRM especially), keep “ask before updating” switched on until you trust the pattern.

Automation: what's realistic today

A clear-eyed look at what Claude can and can't automate on its own — and how to bridge the gap when you need true hands-off automation.

WHAT WORKS TODAY

- ✓ Reading and drafting across your connected apps, inside a chat you start
- ✓ Calendar and Drive read/write; Gmail drafts (never auto-send)
- ✓ Multi-step file work in Cowork — now **on a cloud schedule with /schedule**, no device online needed

WHAT NEEDS A BRIDGE

- ✗ Event-triggered “when X happens, do Y” that reacts to outside events on its own
- ✗ Meta's apps — Facebook, Instagram, Ads Manager, WhatsApp Business
- ✗ Anything across the thousands of apps Claude has no direct connector for

📘 The bridge: Zapier or Make

These services connect Claude to thousands of apps and run on triggers. You build a “when this, then that” rule once, and it runs without you. Claude connects to them through MCP.

>_ EVENT-DRIVEN, VIA A BRIDGE

[COPY / PASTE](#)

When **[trigger, e.g. a new lead form is submitted]**, post **[content]** to our Facebook Page and log the details in a Google Sheet.

- ★ Rule of thumb: if you want Claude to act while you watch, use a connector in a chat or Cowork. If something must react to an outside event by itself, use Zapier or Make.

08

Top Skills to try

Skills are ready-made abilities Claude switches on when they're useful. Here are the ones worth knowing right now — and exactly how to add them.

IN THIS SECTION

- What a Skill is & how to add one
- The most-installed Skills right now
- A word on safety

What a Skill is — and how to add one

A Skill is a small, packaged ability — instructions plus optional files — that Claude switches on automatically when it's relevant. Claude ships with built-in Skills for making Office files; everything else you turn on or install.

1 Turn on the engine

Settings → **Capabilities** → enable *Code execution & file creation*. Skills need this to run.

2 Open Skills


Go to **Customize → Skills**. Toggle on any built-in Skill; it activates when relevant — you don't call it by name.

3 Add a new one

Click **Browse skills** to install from the marketplace, or **Create skill** → upload a .zip of the skill folder (with its SKILL.md inside).

Skills you enable are shared

A Skill switched on in your settings shows up across chat, Cowork, and the Microsoft 365 add-ins — set it once.



Toggle on / off — and add your own

★ **In Claude Code:** drop a folder in `.claude/skills/`, or run `/plugin marketplace add anthropics/skills`. **In Cowork:** the same Skills you've enabled carry over — type / in a task to see them.

The most–installed Skills right now (1–5)

No official leaderboard exists, so this is a considered pick — weighing who made it, install counts and GitHub stars, and how often each is praised in recent write-ups. All work on the web, in Cowork, and in Claude Code. Refreshed July 2026.

BUILT IN

1 · Document skills

PowerPoint, Excel, Word, PDF — the four that let Claude make and edit Office files. Most people's first Skills.

Built in on paid plans · github.com/anthropics/skills

ANTHROPIC

2 · skill-creator

The “make me a Skill” Skill. It interviews you, then builds a custom Skill from your answers — no code required. The one to turn on first.

anthropics.com/skills → skill-creator

ANTHROPIC

3 · frontend-design

Makes Claude produce distinctive, polished web pages and Artifacts instead of generic ones. 275k+ installs — the most-installed Skill in the ecosystem.

anthropics.com/skills → frontend-design

COMMUNITY ★

4 · Superpowers

The most-starred community skill framework (90k+ GitHub stars, now in the official marketplace): a development methodology that makes Claude Code plan, test, and verify like a senior engineer.

Official marketplace → Superpowers

ANTHROPIC

5 · doc-coauthoring

A structured writing workflow: gather context, refine, then “reader-test” the draft with a fresh Claude to catch blind spots. Great for specs and proposals.

anthropics.com/skills → doc-coauthoring

The most–installed Skills right now (6–10)

ANTHROPIC

6 · brand-guidelines

Applies a consistent brand — colours, type, logo rules — to any output. Clone it to encode your own brand, then every deck and doc comes out on-brand.

[anthropics/skills → brand-guidelines](#)

ANTHROPIC

7 · theme-factory

Ready-made colour-and-font themes you can drop onto any deck, doc, or page — or ask it to generate a fresh one. The fast route to a consistent look.

[anthropics/skills → theme-factory](#)

ANTHROPIC

8 · canvas-design

Creates original visual art — posters, covers, infographics — as PNG or PDF. The go-to for marketing visuals without a designer.

[anthropics/skills → canvas-design](#)

PARTNER

9 · Notion Skills

Official partner Skills that let Claude work inside Notion — prep meetings, write sourced research pages, capture knowledge, turn specs into tasks. Needs a Notion connection.

[notion.so → “Notion Skills for Claude”](#)

PARTNER

10 · Remotion

Programmatic **video** — turn data or a script into rendered motion graphics. One of the most-installed skills by weekly usage; a genuine “wow” for social content.

[Official marketplace → Remotion](#)

⚠️ A word on safety

Skills can run code. Stick to Anthropic, the official marketplace, and named partners. Before installing a community Skill, open its SKILL.md and read what it does — audits keep finding third-party skills with hidden risks.

ℹ️ Getting the official Skills

In the web app: Customize → Skills → **Browse skills** → Install. From source: download the folder from github.com/anthropics/skills as a .zip and upload it; in Claude Code, `/plugin marketplace add anthropics/skills`.

09

Pitfalls & ethical use

Capability without judgment is how people get burned. This is the judgment — five pitfalls with cheap, boring, reliable fixes, and the professional standard for honest AI.

IN THIS SECTION

- The five pitfalls — and the fix for each
- Using AI honestly

The five pitfalls — and the fix for each

1 Hallucination (the big one)

Claude states false things fluently — fabricated facts, citations, quotes, numbers. Not lying; pattern-completing. **THE FIX** → verify anything with consequence: click citations, check quotes, recompute key numbers.

2 Stale knowledge

Without web search, it answers from its training cutoff. **THE FIX** → for anything time-sensitive, Search on + citations present.

3 Over-trust on math & details

Plausible ≠ correct — arithmetic in prose, names, dates, legal/medical specifics. **THE FIX** → route real computation to the analysis tool / code execution; treat specifics as unverified until checked.

4 Privacy leaks

Confidential material in the wrong account or tier. **THE FIX** → client & company data only on Team/Enterprise (or with training off); sensitive one-offs in Incognito.

5 The competence mirage

Polished output switches off your critical reading — usually around week three. **THE FIX** → the new-hire rule: brilliant, fast, and reviewed like a junior colleague's work. Every time, forever.

THE ONE-LINE POLICY

AI drafts; humans decide.

If output leaves your hands with your name on it, you've verified it. If your organization adopts a single sentence of AI governance, make it this one.

i It applies double to agentic work — Cowork and Claude Code act on their own conclusions.
Consequence triggers checking — that's the habit.

Using AI honestly — the professional standard

DISCLOSURE

Say so when it's expected

Know when your client, publication, platform, or regulator expects AI use to be declared — many contracts now have AI clauses. When in doubt, disclose. Nobody ever damaged trust by being honest about their tools.

OWNERSHIP

You own the output

The errors, the biases, the plagiarism risk, the consequences. “The AI wrote it” has never once worked as a defense.

FAIRNESS

Check the framing

Models carry biases from their training data. For work touching people — hiring, screening, targeting — ask: * “What perspectives or groups might this framing overlook or disadvantage?” *

PEOPLE FIRST

Augment, not replace

Use AI to make your team faster and sharper — never to quietly hollow out the expertise your clients are actually buying. Clients buy your thinking; AI accelerates it.

Know your policies

✓ your employer's AI policy · ✓ each client's AI clauses (increasingly in MSAs) · ✓ your industry's rules — legal, medical, financial, and academic contexts have real ones.

★ The standard that never goes out of date

Use AI in a way you'd be comfortable explaining — to the client, in plain language, with the “why.” If the explanation would embarrass you, the usage is wrong.

The one–page cheat sheet

Tear this page out (or bookmark it). Everything you'll reach for most, in one place.

PICK A MODEL

Everyday work	Sonnet 5
Important / complex	Opus 5
Quick & simple	Haiku 4.5
The rare splurge (metered)	Fable 5

WHICH SURFACE

Think, write, analyze	Claude Pro · web
Work on my real files	Cowork
Automate / build	Claude Code

WEB APP SHORTCUTS

Send message	Enter
New line	Shift + Enter
Attach a file	+ in the box
Private chat	Ghost icon

CLAUDE CODE BASICS

Start	type <code>claude</code>
Plan first	<code>/plan</code>
Make memory file	<code>/init</code>
Switch model	<code>/model</code>

THE FIVE–PART PROMPT

- **Role** — who Claude should be
- **Context** — the situation & audience
- **Task** — exactly what to do
- **Examples** — show the style
- **Format** — the shape of the output

★ The universal shortcut

Stuck? Add: “Ask me any questions you need before you start.”

i Settings worth memorizing

Privacy (training toggle) · Profile (preferences) · Memory · Customize → Skills

Your 4-week rollout plan

Don't try to do everything at once. This staged plan takes you from sign-up to confident across all three surfaces — one manageable step at a time. Move on when the benchmark feels true.

STAGE 1 Web foundation Day 1	STAGE 2 Connectors & Skills Week 1	STAGE 3 Cowork Week 2	STAGE 4 Claude Code Week 3+
✓ Sign up & upgrade to Pro ✓ Secure the sign-in account; set the Privacy toggle ✓ Fill in profile preferences (interview method) ✓ Turn on Memory; learn Incognito ✓ Build your first Project	✓ Connect Google or Microsoft 365 ✓ Enable document Skills ✓ Start every connector read-only ✓ Write your style guide ✓ Try an on-brand rewrite	✓ Open Cowork (web or desktop app) ✓ Grant one folder, run a small task ✓ Set global & folder instructions ✓ Try a cloud /schedule task ✓ Build trust before bigger jobs	✓ Start in the Code tab (no terminal) ✓ Or install via terminal ✓ Run /init, trim your CLAUDE.md ✓ Always Plan mode first ✓ Keep a backup before big changes

When to upgrade beyond Pro

Only when you repeatedly hit usage limits — which happens fastest with Cowork and Claude Code. Hitting the wall a few times a week points to **Max 5x**; all-day heavy use points to **Max 20x**; a team that needs shared billing points to **Team**.

YOU'RE READY

Start small. Reuse everything. Let it grow with you.

You now have the full picture: the web app for thinking, Cowork for real files, and Claude Code for automation — tied together by the context files that make Claude sound like it already knows you. Pick one task this week and try it.

Keep learning

Anthropic ships new features almost monthly.
Your own account is always the source of truth for what's live.

A few honest caveats

Gmail is draft-only. Regulated data needs the Enterprise audit story. Model names and prices change. Verify specifics before you rely on them.

Need a hand?

Cortexa helps teams put Claude to work — setup, training, and custom workflows. Come say hello.