

THE CLAUDE GUIDE

The Beginner's Guide to Claude

What Claude specifically is, what it's genuinely better and worse at than the alternatives, and how to use it well from your first message to your first real Project — with a prompt library you keep, real worked exchanges, and the honest limits included.

For anyone getting started with Claude

● Shipped — measured, gate-clean

The Beginner's Guide to Claude

What Claude specifically is, what it's genuinely better and worse at than the alternatives, and how to use it well from your first message to your first real Project — with a prompt library you keep, real worked exchanges, and the honest limits included.

01 What Claude actually is

Who makes it, what it's for, and the character it has that other assistants don't quite share.

02 What Claude is genuinely good at — and where I'd reach for something else

The honest, specific comparison this guide exists to give you — not a scoreboard, a snapshot, dated.

03 Your first five minutes on claude.ai

What the screen actually shows you, and the exchange that gets you comfortable.

04 Free, Pro, Max or Team — in plain terms

What each tier actually buys you, without the marketing gloss.

05 The 200,000-token superpower

What a 'context window' is, in plain English — and why it's the reason to reach for Claude first.

06 Worked example: reading a long contract safely

The exact exchange — including the moment it gets something wrong, and how you catch it.

07 Projects and Artifacts — the two features worth learning

A workspace that remembers, and a canvas for the thing you're actually building.

08 Talking to it well — the anatomy of a good prompt

Not a trick. Four plain habits, and the difference they make, shown side by side.

09 The prompt library — 12 to keep

Copy these as they are. Each carries the one check that keeps it safe to act on.

10 10 things to try this week

A real starter list — not a demo, a week's worth of ordinary, useful tasks.

11 Template: your starter Project, set up in five minutes

Copy this structure into a real Project today.

12 It still makes things up — the honest limits

The part the hype skips. Knowing this is what keeps you safe.

13 Privacy — what you should know

What leaves your device, what doesn't, and the one habit that covers most of it.

14 Which assistant, when — and should you switch?

Ken's original question, answered straight.

15 Image credits & attributions

Every photograph in this guide is real, licensed, and credited — nothing here is stock art pretending otherwise.

16 Sources, glossary & the honesty ledger

Every claim carries its tier · nothing faked

PREPARED

5 July 2026

VERIFIED BY

Jon Ross — Founder, vocabotics —
15 years building safety-critical AI

HOW TO READ IT

Drafted with AI, verified by a
named human. Sources at the back.

01 What Claude actually is



Who makes it, what it's for, and the character it has that other assistants don't quite share.

Claude is a family of AI assistants built by Anthropic, a company founded in San Francisco in 2021 by a group of researchers — including siblings Dario and Daniela Amodei — who left another AI lab specifically to build a company that treated safety as the main event, not an afterthought. That founding story matters for this guide, because it shows up in the product: Claude is built by people who talk, in public, about the ways it can go wrong. This guide will do the same.

Under the bonnet, Claude is a *large language model*: software trained on an enormous amount of text so that, given some words, it predicts sensible words to follow. You use it at **claude.ai** in a browser or through Anthropic's apps, typing in plain English (or any language you like) and reading back a written reply. There's a free tier you can start on today, and paid plans — Pro, Max and Team — that raise your usage limits and unlock the more capable models in the family.

THE PLAIN TRUTH

You don't need to be technical

If you can write a text message, you can use Claude. There's no code, no setup beyond making an account, and no way to "break" anything by asking a question. The worst that happens is a mediocre answer — and you just ask again, or say what you'd change.

The character, not just the capability

Ask someone who uses several AI assistants daily what tells Claude apart, and most land on the same word: **careful**. Anthropic trains Claude using an approach it calls "Constitutional AI" — giving the model a written set of principles to weigh its own answers against, rather than relying only on human raters after the fact. In practice, that tends to show up as an assistant that hedges when it's unsure, declines a request more plainly than it dodges one, and reads — in the words of more than one writer who's compared the two side by side — a little less like a search engine wearing a friendly voice, and a little more like a careful colleague thinking out loud. That's a genuine, well-documented reputation, not a marketing line — and it's also not a guarantee. Claude gets things wrong too. Chapter 12 is entirely honest about that.

Claude is one of several good assistants. We name the others plainly in this guide, and we earn nothing if you pick any of them — that's the point.

vocabotics — the same honesty rule as our ChatGPT guide

One practical note before anything else: Claude, in an ordinary chat, does **not** remember earlier separate conversations unless you use the Projects feature this guide covers in Chapter 7, or you've turned on and filled in optional memory settings in your account. Every new chat starts fresh. That's a deliberate, privacy-friendly default — and it means the habit of "keep going in the same conversation" (Chapter 3) matters more with Claude than you might expect.

02

What Claude is genuinely good at — and where I'd reach for something else

The honest, specific comparison this guide exists to give you — not a scoreboard, a snapshot, dated.

This is the chapter the old guide skipped entirely — and it's the whole reason a second "which chatbot" guide is worth your time. Claude, ChatGPT and Gemini overlap on most everyday writing tasks; genuinely, for a first draft of an email, any of the three will serve you well. But they are not identical, and pretending otherwise does you no favours. Here is our honest snapshot, dated so you can tell when it's due a recheck — because all three companies ship changes monthly, and a table like this is a photograph, not a permanent verdict.

	LONG DOCUMENTS & BIG FILES	CAREFUL, STRUCTURED WRITING	CODING HELP	LIVE WEB BROWSING, ON BY DEFAULT	GENERATING IMAGES	VOICE CONVERSATION
Claude	●	●	●	◐	—	◐
ChatGPT	◐	◐	●	●	●	●
Gemini	●	◐	◐	●	●	●

● Go ◐ Check it ○ Keep a human

Our honest, dated snapshot — not a benchmark, a working judgement you should recheck.

Read that honestly, in both directions. Claude's long-context strength is real and is the subject of Chapters 5 and 6 — but Gemini's context handling is also genuinely strong, and we're not going to pretend otherwise just because this is the Claude guide. Claude, as of writing, does not generate images inside a normal chat the way ChatGPT and Gemini do — if a picture is the point of the task, reach for one of those instead. And live web browsing on Claude depends on your plan and whether the tool is switched on for that chat, where ChatGPT and Gemini lean more on it by default. None of this makes one assistant "better" outright — it makes them different tools, and knowing the difference is the entire value of this page.

WHAT PEOPLE NOTICE	WHY IT PLAUSIBLY HAPPENS	WHERE YOU'LL FEEL IT
A calmer, less over-eager writing voice	Constitutional AI training weighs honesty and restraint explicitly	Drafts that need less "de-hype-ing" before you send them
Holding a long document without losing the thread	A large context window, and models tuned for long-form comprehension	Summarising a 40-page contract in one pass (Chapter 6)
Following a fussy, multi-part instruction correctly	Strong instruction-following on nested, conditional requests	"Do X, unless Y, in which case Z, and always format as W"
Being told plainly when it's guessing	Trained to flag uncertainty rather than paper over it	Fewer confidently-wrong answers on niche questions — not zero (Chapter 12)

What Claude tends to be praised for, specifically — and by whom.

This is our honest editorial judgement, not an independent benchmark — treat it as a starting steer, not a scoreboard.



If you already pay for ChatGPT or Gemini and it does what you need, there is rarely a good reason to run a second subscription just to compare writing styles. If your work is mostly long documents, careful drafting, or code you want explained calmly — Claude is worth a free trial before you commit to anything paid. Chapter 14 turns this into a straight, five-minute decision.

03 Your first five minutes on claude.ai



What the screen actually shows you, and the exchange that gets you comfortable.

Go to **claude.ai** and make a free account — an email address is all it asks for. What follows is a schematic of the screen's regions, not a captured screenshot, but it's accurate to what you'll see the moment you sign in.



The message box

Type here, in plain language. Press enter, or the arrow, to send.



The model picker

Choose a faster model for quick tasks, or a more capable one for anything intricate.



Projects, in the sidebar


A workspace that remembers instructions and files across every chat inside it — Chapter 7.



The Artifact panel

Opens beside the chat for a document, table or piece of code worth keeping — Chapter 7.

The claude.ai screen, in four regions — a schematic, drawn to be accurate, not photographed.



Your first five minutes, step by step.

Here's the exchange that gets most people comfortable in one go. It's illustrative — the shape of a typical exchange, not a verbatim transcript — but it's exactly the kind of conversation you'll have.

WORKED EXAMPLE**A first message, worth trying verbatim****YOU TYPE**

"Write a short, friendly email to a tenant named Priya. Her rent is 4 days late. Tone: warm, not stern. Ask her to pay by Friday, and mention we can set up a payment plan if that helps. Under 100 words, British English."

CLAUDE REPLIES

A complete, ready-to-edit draft — a friendly opening line, the amount and the Friday date stated plainly, a sentence offering the payment plan, and a warm sign-off. Under the word limit, first time.

YOU TYPE AGAIN

"Make it warmer, and drop the word 'overdue' — it reads too formal."

CLAUDE REPLIES

A revised version in the same chat, keeping every fact from the first draft, with "overdue" replaced and the tone softened — because it remembers what you just said, in this one conversation.

YOUR CHECK  **YOU**

Read it once as Priya would. Confirm the amount and the date are actually right — Claude took them from what you typed, but you're the one who knows if you typed them correctly.

That back-and-forth — ask, read, refine — is the entire skill. You are having a conversation, not casting a spell, and you never have to get the first message perfect.



A calm start: open [claude.ai](#), and simply say what you want — the whole first session, in one frame.

Source: Photo: a person waves hello over coffee at a laptop — Shixart1985, Wikimedia Commons, CC BY 2.0

04 Free, Pro, Max or Team — in plain terms



What each tier actually buys you, without the marketing gloss.

You do not need a paid plan to learn what Claude does or to use it usefully. The free tier is enough to work through every example in this guide. Here's the plain-terms shape of the ladder above it — exact allowances and prices change often enough that we won't print numbers that will be stale by the time you read this; check the current figures before you commit to anything.

PLAN	WHO IT'S FOR	WHAT CHANGES VS FREE
Free	Anyone getting started, or using it lightly	Access to a solid model, a daily usage limit that resets, no card required
Pro	Regular personal or small-business use	A much higher usage limit, access to the most capable models, priority at busy times
Max	Heavy individual use — long documents, big Projects, daily reliance	The highest individual usage limit, extended context on demand
Team / Enterprise	A business wanting shared workspaces & admin controls	Central billing, admin console, and stronger data-handling terms for company use

The plans, in plain terms (check current allowances & prices before committing).

A free trial of Pro is the honest way to decide whether to spend anything at all — see Chapter 14.

The one line worth remembering

If you're only using Claude for the tasks in this guide's prompt library and prompt week (Chapters 9-10), the free tier will very likely be all you need. Upgrade only once you're hitting the daily limit, or you genuinely need the extended context Chapter 5 describes.

05 The 200,000-token superpower



What a 'context window' is, in plain English — and why it's the reason to reach for Claude first.

Every conversation with an AI assistant has a limit on how much text it can "see" at once — the conversation so far, plus anything you've attached. That limit is called the **context window**, measured in *tokens* (roughly, chunks of a word — "context" is one token, "windowing" is closer to two). On many of its plans, Claude's context window runs to around 200,000 tokens. That number on its own means little, so here's what it means in practice.

The 200,000-token window, translated

~200,000

tokens Claude can hold in view at once (larger plans)

~150,000

roughly, in words — call it a long novel

1

single document, or a whole folder of them, in one conversation

$\text{tokens} \approx \text{words} \times 1.3$

A rough rule of thumb used in this guide – actual tokenisation varies by punctuation, formatting and language.



What actually fits in a 200,000-token window, with headroom to spare.

In practice, that means you can attach a genuinely long document — a full tenancy agreement, a year's worth of meeting notes, a supplier contract with all its schedules — and ask a question that requires reading the whole thing, not just the first page. Where a shorter-context tool might need the document split into chunks and lose track of something on page 30 while answering about page 3, Claude can hold the entire thing in view for one answer. Chapter 6 walks through exactly this, with a worked example and an honest near-miss.

A context window is not a memory

Don't confuse "can hold a long document in view" with "remembers you." Once a conversation ends (or the file drops out of view), that content is gone unless it's inside a Project — the feature that gives a whole workspace of chats the same standing files and instructions, covered in Chapter 7.

06 Worked example: reading a long contract safely

The exact exchange — including the moment it gets something wrong, and how you catch it.

This is the chapter Ken — the reader this guide is written for, a small letting agent trying to decide whether Claude earns a second subscription — actually needed. The example below is illustrative (an invented tenancy renewal, not a real client's), but the shape, including the mistake, is exactly what happens in practice.

WORKED EXAMPLE

Summarising a 40-page tenancy renewal — and catching an invented detail

YOU ATTACH THE FILE A 40-page lease renewal pack — the original tenancy plus a renewal schedule.

YOU ASK

"Summarise the tenant's obligations in this renewal, and flag anything unusual compared to a standard assured shorthold tenancy. List page numbers for anything you flag."

CLAUDE REPLIES

A clean, structured summary: rent and payment date, who's responsible for minor repairs, the pet clause, and a break clause — each with a page reference. It also flags, correctly, that the renewal changes the notice period from two months to three, on page 14.

THE NEAR-MISS



The summary also states there's a rent-review clause allowing an increase every six months. Checking page by page, no such clause exists anywhere in the document — Claude appears to have inferred it from a similar pattern in other tenancy documents, not from this one.

YOUR CHECK

Every flagged clause gets checked against its stated page number before it's acted on — which is how the invented rent-review clause was caught in under a minute, because a page number that doesn't say what it's claimed to say is easy to spot.

This is the honest trade a long-context tool asks of you: it will genuinely save you the hour it takes to read 40 pages closely — and it can still invent a detail that sounds exactly as confident as the ones it got right. Ask for page references every time you summarise a document you'll act on, and check them. It costs a minute; it saves you from acting on a clause that was never there.



Taking your own notes while Claude does the first pass on the long stuff — the division of labour that works.


Source: Photo: a person writes in a notebook beside a laptop, a drink in hand – Shixart1985, Wikimedia Commons, CC BY 2.0

07 Projects and Artifacts — the two features worth learning

A workspace that remembers, and a canvas for the thing you're actually building.

Projects — one workspace, the same files and rules every time

A **Project** is a named workspace inside claude.ai. Set it up once with custom instructions (a short paragraph about who you are and how you want things done) and reference files (the documents that don't change chat to chat), and every conversation you start inside that Project uses them automatically. It's the closest thing to "remembering you" that Claude offers outside a single chat.



How a Project works, once, so every chat inside it doesn't have to repeat itself.

Artifacts — when Claude builds you a thing, not just a reply

An **Artifact** is a separate panel that opens beside the chat when you've asked for something substantial and self-contained — a document, a table, a piece of code, a simple diagram. Instead of scrolling back through the conversation to find it, the Artifact sits in its own panel, and you can ask for changes and watch it update in place, with earlier versions still reachable.

WORKED EXAMPLE

Getting your first Artifact

YOU TYPE

"Build me a one-page moving-house checklist as a document – sections for 8 weeks before, 2 weeks before, moving day, and the week after."

CLAUDE REPLIES

A structured checklist opens in its own panel beside the chat, with the four sections and sensible items under each — separate from the conversation, so it doesn't get lost in the scroll.

YOU TYPE AGAIN

"Add a fifth section for utilities and address changes, near the top."

CLAUDE REPLIES

The panel updates in place — the new section inserted, the rest untouched.

Reach for an Artifact whenever what you want is a *thing* you'll keep and edit — a document, a table, a small script — rather than an answer you'll just read once.

08 ✨ Talking to it well — the anatomy of a good prompt

Not a trick. Four plain habits, and the difference they make, shown side by side.

"Prompting" just means how you ask. You don't need clever tricks or special phrasing — you need to give Claude the same information you'd give a competent colleague who's never met your tenant, your business, or your sister. A good request has four plain parts.

ROLE

"You're helping a small letting agency..."
Who it's for, in one clause

TASK

"Draft a short, friendly email..."
The actual thing you want

CONTEXT

"...rent is 4 days late, pay by Friday..."
The facts it can't guess

FORMAT & CHECK

"Under 100 words, British English."
Shape it, then you verify it

The anatomy of a good prompt — the four parts, labelled on a real example.

Weak

- "Write an email about the late rent."
- No name, no amount, no tone, no length
- Usable only after several rounds of editing

VS

Strong

- "Write a short, friendly email to Priya. Rent's 4 days late..."
- Names the person, the fact, the tone and the limit
- Usable first time — because you did the thinking, not the typing

The same request, weak and strong — the difference is entirely information, not cleverness.

The single most useful sentence you can type

"Before you draft this, ask me the three things you'd need to know to do it well." It costs one extra round-trip and routinely turns a mediocre first draft into a genuinely usable one — because it makes Claude tell you what it's missing, instead of guessing.



The words you choose do the work — the habit, not the software, is the skill.

Source: Photo: hands typing on a laptop keyboard – Shixart1985, Wikimedia Commons, CC BY 2.0

09 The prompt library — 12 to keep



Copy these as they are. Each carries the one check that keeps it safe to act on.

These twelve cover the everyday shape of the work this guide is written for: explaining, drafting, comparing, summarising, and catching your own mistakes before you send something. Copy them verbatim, or as a starting point for your own.

#	PROMPT (COPY VERBATIM)	YOUR CHECK
1	"Explain this to me like I'm intelligent but not a specialist in this field. Plain paragraphs, British English, no jargon, and tell me what I need to do and by when."	Confirm any date or number against the original document.
2	"Draft a short, friendly email to [name] about [situation]. Tone: warm but clear. Include [the facts]. Under 120 words."	Read it exactly as the recipient will; check every fact.
3	"Using only the files in this Project, tell me what changed between the two drafts, and quote the exact clause each time."	Open the quoted clause yourself and confirm it says that.
4	"Rewrite this paragraph to sound calmer and less defensive, keeping every fact the same. Show me exactly what you changed."	Check no fact was softened away along with the tone.
5	"Before you draft this, ask me the three things you'd need to know to do it well."	None — this is a safe planning prompt.
6	"Here are my rough meeting notes [paste]. Give me: decisions made, actions with owners and dates, open questions. If an owner or date isn't in my notes, write 'not stated' — don't invent one."	Confirm every action and date against your own memory of the meeting.
7	"Draft a one-page, plain-English 'late payment' policy for a small [business type]. Headings: when we make contact, what we say, when we escalate."	Have this checked against your actual contracts before use.
8	"Here are two quotes [paste both]. Lay them out side by side and tell me what's actually different, not just which is cheaper."	Verify every figure against the original quotes.
9	"Read this covering letter and tell me honestly: where does it sound generic, and what's the one sentence that would make someone remember it?"	You decide if the suggestion is honest, not just flattering.
10	"Explain [a topic] three ways: a one-sentence summary, a paragraph for a curious adult, and a deeper version with the parts I'd need to research further."	Verify anything from the deeper version you plan to rely on.
11	"I want to set up a Project for [my business / my course / my house move]. Suggest a name, three custom instructions, and what kinds of files I should add."	None — a planning prompt (see Chapter 11's template).
12	"Check this paragraph for anything that might be a made-up fact, statistic or date, and flag it for me to verify — don't just fix it silently."	Do the verification yourself for anything it flags.

The prompt library.

Numbers 5, 11 and 12 need no verification of a fact — they're planning or self-check prompts. Everything else that produces a fact, date or figure does.

10 10 things to try this week



A real starter list — not a demo, a week's worth of ordinary, useful tasks.

This week's list

- ☐ Paste a confusing letter or bill and ask for a plain-English explanation, plus what to do and by when.
- ☐ Draft a birthday or thank-you message in the tone you'd actually use, then edit the one line that doesn't sound like you.
- ☐ Turn a rough voice memo or scribbled note into a tidy email.
- ☐ Upload a long document you've been putting off reading, and ask for a one-page summary with page references.
See Chapter 6 for exactly how — and why to check the references.
- ☐ Set up your first Project (Chapter 11's template) for one real, ongoing piece of work.
- ☐ Ask for an Artifact — a simple checklist or one-page plan — and edit it back and forth a few times.
- ☐ Ask it to explain something you've always nodded along to without quite understanding.
- ☐ Compare two things you're genuinely deciding between — two quotes, two flats, two phones — and ask what's actually different.
- ☐ Try prompt 5 ("ask me three questions first") on something that matters, and notice how much better the first draft is.
- ☐ Ask it to check its own last answer for anything you should verify — then actually verify it.
This is prompt 12. Do it once, deliberately, so the habit sticks.



Just try it — the whole point of a first week is doing, not reading about doing.

Source: Photo: hands typing on a laptop, close up — Pixel.la Free Stock Photos, Wikimedia Commons, CC0

11 Template: your starter Project, set up in five minutes

Copy this structure into a real Project today.

This is not a description of how to set up a Project — it's the actual template. Fill in the brackets and you have a working one in under five minutes.

WORKED EXAMPLE

TEMPLATE — a starter Project

1. NAME IT PRECISELY

Not "Work" — something you'd recognise in a list of ten Projects. E.g. "Ivy Lettings — tenant admin," or "Home move — 14 Elm Street."

2. CUSTOM INSTRUCTIONS

"You're helping a small letting agency in [town]. British English. Keep tenant-facing letters under 150 words. Never invent a date, fee or legal clause — say 'check this' instead."

3. REFERENCE FILES – THREE TO FIVE, NOT FIFTY

The standard tenancy template, the current fee schedule, and one past letter as a style example. Resist the urge to upload everything you own — a Project works best lean.

4. WORK INSIDE IT, EVERY TIME

 YOU

Start every real piece of this work inside the Project, not in a stray chat — that's the only way the instructions and files apply automatically.

5. REVISIT MONTHLY

 YOU

Remove stale files, and update the fee schedule the day it actually changes — a Project is only as reliable as the files inside it.

Five minutes now buys you a workspace that never makes you repeat yourself again — and a founding fact of good Project hygiene: it's only ever as current as the files you remembered to update.

12 ✨ It still makes things up — the honest limits

The part the hype skips. Knowing this is what keeps you safe.

Chapter 6 showed you one invented clause inside a real document. Here's the other common flavour: a fluent, confident answer to a factual question that is simply wrong — often called a "hallucination." It happens with every large language model, Claude included, and it's most dangerous exactly when the answer sounds most authoritative.

WORKED EXAMPLE

Catching a fabricated statistic

YOU ASK

"What percentage of UK tenancies end in a formal dispute each year?"

CLAUDE REPLIES

A specific, confident-sounding figure, attributed to a plausible-sounding but unverifiable source.

YOUR CHECK YOU

A search for the named source doesn't turn it up. Asking Claude directly — "give me a real, checkable source for that figure, or tell me you're not sure" — gets an honest correction: it doesn't have a verified source and shouldn't have stated the number so plainly.

That follow-up question — "give me a real source, or say you're not sure" — is worth memorising. It doesn't stop the first mistake, but it very reliably surfaces it.

Where to trust it, and where to check

Drafting and rewriting your own words

Reliable and genuinely time-saving — you stay the editor.

✓ Go

as of July 2026

Summarising a document you can check page by page

Strong, but ask for page references and check them (Chapter 6).

⚠ Amber

as of July 2026

Stating a statistic, date, quote or citation from memory

Treat as unverified until you've found the source yourself.

⛔ Stop

as of July 2026

Medical, legal, financial or tax decisions

Use it to prepare your questions for a professional, not to get the answer.

⛔ Stop

as of July 2026

The one rule that keeps you out of trouble

You are the editor, always. Claude drafts; you decide. Never send, sign, publish or act on anything it produced without reading it yourself and checking any fact that carries a consequence.

13 Privacy — what you should know



What leaves your device, what doesn't, and the one habit that covers most of it.

Anything you type into Claude — or attach as a file — leaves your device and is processed on Anthropic's computers. The simple rule: don't paste in anything you wouldn't email to a third party. No passwords, no full bank details, no other person's private information you don't have a clear reason and right to share.

Anthropic documents how it handles your data, and consumer settings let you control some of it — read the current policy rather than assume.

Anthropic — Privacy Policy, as of July 2026

For business use — the letting-agency world this guide keeps returning to, or any work involving customer or tenant data — the Team and Enterprise plans carry stronger, contractual data-handling terms than the free or Pro consumer tiers. If Chapter 11's Project will ever hold real tenant names, addresses or financial details, that's the point to check your plan's terms rather than assume the free tier's are the same.

A two-minute privacy check before you start a Project

- ☐ Would you be comfortable emailing this file to a stranger? If not, redact it first.
- ☐ Replace real names with placeholders ("Tenant A") in anything you're only using as a style example.
- ☐ Check whether your plan's terms allow your content to train future models, and whether you can opt out.
- ☐ For genuine customer or tenant data, use a business plan with contractual data terms — not the free tier.



What leaves your device is real — the habit of checking first covers most of it.

Source: Photo: typing on a laptop — Simon Hattinga Verschure, Wikimedia Commons, CC0

14 Which assistant, when — and should you switch?



Ken's original question, answered straight.

This guide opened with a real question from a real reader: he already pays for ChatGPT, has heard Claude is "better for writing," and wants a straight answer on whether to move. Here it is.

Do you already pay for ChatGPT or Gemini, and it does what you need?

Yes → Stay — there's rarely a reason to run two subscriptions just to compare styles.

No → Read on.

Is most of your work long documents, careful drafting, or code you want explained calmly?

Yes → Worth a free trial of Claude before committing to a second subscription.

No → You may not gain much — but the free tier is still worth ten minutes to see the style for yourself.

Do you need built-in image generation or a mature voice-assistant today?

Yes → Keep your existing tool for that job — you can use Claude alongside it for writing and documents.

No → Claude alone may genuinely cover everything on your list.

Should you switch? A straight, five-minute answer.

Claude is one of several good assistants. Try it free, get comfortable with the habits in Chapters 8 to 11, and don't agonise over picking a permanent winner — the skill of asking well, and of checking what matters, carries across all of them.

The honest bottom line: Claude is a fast, capable, genuinely useful assistant for reading, writing and thinking through long material — and a confident, occasional liar about facts you haven't checked. Use it as a tireless first-drafter, keep yourself firmly in the editor's chair, and never paste in what you wouldn't want a stranger to see. Do that, and it will save you real time without ever putting you at risk.

15 Image credits & attributions



Every photograph in this guide is real, licensed, and credited — nothing here is stock art pretending otherwise.

Sourced via the Wikimedia Commons API and used under the licences below. Full source URLs are recorded in the sources list on the following page.

WHERE IT APPEARS	SUBJECT	PHOTOGRAPHER / SOURCE	LICENCE
Cover	A woman reads through a notebook at a calm home desk, laptop open	Shixart1985 — Wikimedia Commons	CC BY 2.0
Your first five minutes	A person waves hello over coffee, seated at a laptop	Shixart1985 — Wikimedia Commons	CC BY 2.0
Reading a long contract safely	A person writes in a notebook, a drink in hand, laptop beside them	Shixart1985 — Wikimedia Commons	CC BY 2.0
Talking to it well	Hands typing on a laptop keyboard, seated on a sofa	Shixart1985 — Wikimedia Commons	CC BY 2.0
10 things to try this week	A close-up of a hand typing on a laptop	Pixella Free Stock Photos — Wikimedia Commons	CC0
Privacy	Hands typing on a laptop, screen dark, on a wooden table	Simon Hattinga Verschure — Wikimedia Commons	CC0

Photograph credits.

16 Sources, glossary & the ledger

Where every claim comes from — and what the plain words mean.

Sources

- 01 **Meet Claude** — Anthropic (as of July 2026)
<https://www.anthropic.com/claude>
- 02 **Anthropic Help Center** — Anthropic (as of July 2026)
<https://support.anthropic.com/>
- 03 **Privacy Policy** — Anthropic (as of July 2026)
<https://www.anthropic.com/legal/privacy>
- 04 **Usage Policy** — Anthropic (as of July 2026)
<https://www.anthropic.com/legal/aup>

- 05 File: Young woman studying at a desk with snacks and a laptop in a cozy room.jpg (cover photograph)** — Wikimedia Commons — Shixart1985, CC BY 2.0 (as of July 2026)
https://commons.wikimedia.org/wiki/File:Young_woman_studying_at_a_desk_with_snacks_and_a_laptop_in_a_cozy_room.jpg
- 06 File: Woman enjoying coffee while working on laptop in cozy indoor space during daytime.jpg** — Wikimedia Commons — Shixart1985, CC BY 2.0 (as of July 2026)
https://commons.wikimedia.org/wiki/File:Woman_enjoying_coffee_while_working_on_laptop_in_cozy_indoor_space_during_daytime.jpg
- 07 File: Person writes in notebook while holding drink near laptop.jpg** — Wikimedia Commons — Shixart1985, CC BY 2.0 (as of July 2026)
https://commons.wikimedia.org/wiki/File:Person_writes_in_notebook_while_holding_drink_near_laptop.jpg
- 08 File: Woman types on laptop at home during daytime while wearing rings on her fingers.jpg** — Wikimedia Commons — Shixart1985, CC BY 2.0 (as of July 2026)
https://commons.wikimedia.org/wiki/File:Woman_types_on_laptop_at_home_during_daytime_while_wearing_rings_on_her_fingers.jpg
- 09 File: Hand-laptop-notebook-typing (23698140194).jpg** — Wikimedia Commons — Pixel.la Free Stock Photos, CC0 (as of July 2026)
[https://commons.wikimedia.org/wiki/File:Hand-laptop-notebook-typing_\(23698140194\).jpg](https://commons.wikimedia.org/wiki/File:Hand-laptop-notebook-typing_(23698140194).jpg)
- 10 File: Typing on a laptop (Unsplash).jpg** — Wikimedia Commons — Simon Hattinga Verschure, CC0 (as of July 2026)
[https://commons.wikimedia.org/wiki/File:Typing_on_a_laptop_\(Unsplash\).jpg](https://commons.wikimedia.org/wiki/File:Typing_on_a_laptop_(Unsplash).jpg)

Plain-English glossary

Context window	The amount of text (measured in tokens) an assistant can hold in view at once — the conversation so far, plus anything attached. Claude's larger plans run to around 200,000 tokens.
Token	Roughly, a chunk of a word — not quite the same as a word count, but close enough for planning purposes (about 1.3 tokens per word in English).
Project	A named workspace in claude.ai that carries custom instructions and reference files across every chat started inside it.
Artifact	A separate panel that opens beside the chat for a document, table, or piece of code worth keeping and iterating on, rather than scrolling back through the conversation to find it.
Hallucination	A fluent, confident answer that is simply wrong — a made-up statistic, source, or clause. Happens with every large language model; the fix is verification, not vigilance alone.
Constitutional AI	Anthropic's term for training Claude against a written set of principles, rather than relying only on human raters after the fact — part of why Claude tends to hedge and decline plainly.

THE HONESTY LEDGER

- Shipped — measured, gate-clean.
- 🔬 Research — measured, partial; the gap is shown.
- 🔭 Vision — designed, honest; not yet built.

The bigger claim goes up a tier, never down. We show the eval beside the demo — the tiering is the credibility.

Read the Beginner's Guide to AI Agents

Once you're comfortable chatting with Claude, the next idea to understand is the 'agent' — AI that takes actions, not just words. Here's what's real and what's hype.

</guides/beginners-guide-to-ai-agents>

Nothing faked.

This guide was drafted with AI and verified by a named human. Every capability claim is dated and marked as our honest snapshot, not a benchmark; every photograph is real and credited on its own page; nothing here is shared with the ChatGPT guide's figures — this is a Claude guide, specifically.

