DAY ONE

01/10

Meet Claude — *what it is and how it thinks*

Before anyone can use Claude well, they need a clear mental model of what Claude actually is: a family of AI assistants, the differences between its models, what it does brilliantly, where it needs supervision, and the unusual way it "remembers" and reasons.

BY THE END OF DAY 1, YOU WILL BE ABLE TO TEACH OTHERS TO —

- Explain what Claude is in one clear sentence
- Describe Claude's core strengths and limits
- Name the products that give access to Claude
- Choose between Opus, Sonnet & Haiku
- Explain tokens, context & the knowledge cutoff
- Run a confident first session with a newcomer

01 Why today matters

ORIENTATION

Everyone you teach will arrive with a guess about what Claude is — "a chatbot," "Google with sentences," "the thing that writes essays." Those guesses quietly shape how they use it, and usually cap how much value they get. Day 1 replaces the guess with an accurate mental model. Every later day — prompting, Claude Code, the API, design — builds on the vocabulary you establish here.

Keep today concrete and unintimidating. The goal is not technical depth; it is confidence and accuracy. A learner who finishes Day 1 should be able to explain Claude to a colleague without hedging, and should feel that the remaining nine days are an unlock rather than a mountain.

PREREQUISITES

None. Learners need only a device with a browser and a free or paid Claude account (`claude.ai`). No coding, no setup. If you are teaching a group, have everyone sign in before you begin so Day 1 ends with a live first session.

02 What Claude actually is

CORE CONCEPT

Claude is a family of large language models built by Anthropic, an AI safety company. A large language model is a system trained on an enormous body of text to predict and generate language; in practice that training produces something that can read, reason, write, summarise, translate, analyse, plan, and code. "Claude" is both the underlying model and the assistant persona you talk to.

> THE ONE-SENTENCE DEFINITION
>
> **Claude is an AI assistant you collaborate with in natural language.**
>
> You describe what you want — in plain words, the same way you'd brief a capable colleague — and Claude responds with reasoning, drafts, code, analysis, or questions. It is a collaborator, not a search engine and not a database. That distinction is the single most useful idea to install on Day 1.

2.1 Search engine vs. collaborator

A search engine returns links that might contain an answer. Claude returns the work itself — a draft, an explanation, a plan — and can revise it with you across a conversation. It can also search the web when a question needs current information, but the interaction model is dialogue, not query-and-click.

2.2 The values built into Claude

Anthropic trains Claude to be helpful, honest, and harmless, using an approach the company calls Constitutional AI — Claude is trained against an explicit set of principles rather than left to absorb whatever is implicit in raw text. For your learners the practical upshot is simple:

- ✓ It will tell you when it is unsure rather than bluff — though it can still be wrong, so verification still matters.
- ✓ It will push back on requests that are harmful or based on a false premise, and explain why.
- ✓ It aims to be genuinely useful, not just agreeable — good feedback over flattery.

KEY IDEA TO INSTALL

Claude is a reasoning collaborator you direct in plain language. Almost every "Claude isn't working for me" story traces back to treating it like a search box or a vending machine instead of a colleague you brief, correct, and iterate with.

03 The model family — Opus, Sonnet & Haiku

CORE CONCEPT

Claude is not one model but a family, named in ascending order of capability after forms of writing. Every generation ships these tiers so you can trade intelligence against speed and cost. As of this guide, the current flagship is Claude Opus 4.7, the most capable model publicly available; it follows the 4.6 family (Opus 4.6 and Sonnet 4.6), with Haiku 4.5 as the lightweight tier.

TIER	CHARACTER	BEST FOR	TRADE-OFF
Opus	The most capable. Deepest reasoning, hardest problems, long multi-step tasks.	Complex coding, research, agentic work, nuanced writing & analysis.	Highest cost, not the fastest.
Sonnet	The balanced workhorse. Near-flagship quality at lower latency and cost.	The majority of everyday work — drafting, analysis, most coding.	Slightly less headroom on the very hardest tasks.
Haiku	The fast, economical tier. Near-frontier intelligence, lowest price.	High-volume, latency-sensitive jobs — classification, extraction, routing.	Less suited to long, intricate reasoning.

TEACHING SHORTCUT

Names rise alphabetically with capability: Haiku → Sonnet → Opus is small → medium → large. Version numbers (4.5, 4.6, 4.7) are generations — higher is newer. Tell learners: "Default to Sonnet; reach for Opus when the problem is genuinely hard; use Haiku when you need speed at scale."

3.1 Where you actually meet these models

The same model family powers several products. Knowing the surface helps learners pick the right tool later in the series:

- **Claude apps** — the web, desktop and mobile chat interface at `claude.ai`. Where most people start. Days 1–4
- **Claude Code** — a command-line tool that lets Claude work agentically inside a codebase. Days 5–6
- **The Claude Developer Platform & API** — programmatic access for building Claude into your own software. Days 7–8
- **Claude in Chrome, Claude in Excel, and Cowork** — newer surfaces that bring Claude into the browser, the spreadsheet, and general desktop task automation.

04 Strengths, and where to stay in the loop

CORE CONCEPT

An honest map of capability is what turns a nervous beginner into a confident user. Teach both columns with equal energy — overselling creates disappointment, underselling wastes the tool.

CLAUDE IS EXCELLENT AT

- ✓ Drafting, rewriting & adjusting tone
- ✓ Explaining hard things at any level
- ✓ Summarising and restructuring long material
- ✓ Writing, reviewing & debugging code
- ✓ Reasoning through multi-step problems
- ✓ Brainstorming, planning & critique
- ✓ Working with documents and images you provide

STAY IN THE LOOP FOR

- **Facts that must be exact** — names, numbers, citations. Claude can be confidently wrong; verify.
- **Very recent events** — beyond its knowledge cutoff unless it searches the web.
- **High-stakes advice** — legal, medical, financial. Useful for understanding; not a substitute for a professional.
- **Anything you can't check** — if you couldn't tell whether the output is right, don't ship it unreviewed.

> *"Treat Claude's output the way a good editor treats a talented writer's draft — expect quality, but read every line."*

THE WORD FOR THE WRONG ANSWERS

When a model states something false with full confidence, that's called a hallucination. It is not lying — there is no intent — it is the model generating plausible language that happens to be untrue. The fix is not fear; it is verification habits: ask for sources, cross-check facts, and prefer web search for anything current.

05 How Claude "thinks" — context, tokens & memory

CORE CONCEPT

Three mechanical facts explain most of Claude's behaviour. Learners who understand these stop being surprised by it.

5.1 Tokens & the context window

Claude reads and writes in tokens — chunks of text roughly three-quarters of a word. Everything in a conversation — your messages, uploaded files, and Claude's replies — sits inside a context window, the model's working memory for that conversation. Modern Claude models have very large context windows (up to around a million tokens on some models — enough for entire codebases or long documents), but it is still finite. When a conversation grows enormous, the earliest material can fall out of focus.

5.2 Each conversation starts fresh

By default Claude does not carry knowledge from one conversation into the next — each chat begins blank. Newer features change this deliberately: Projects keep shared context together, and a memory capability can carry forward details across chats. But the safe default to teach is: if it matters for this task, include it in this conversation.

5.3 The knowledge cutoff

Claude's training data has a cutoff date. It knows the world up to roughly that point and not beyond — so for genuinely current information it needs to use web search. Teach learners to notice the difference between "explain how mortgages work" (timeless — answer directly) and "what is today's interest rate" (current — needs a search).

THREE FACTS, ONE HABIT

Finite context + fresh-start conversations + a knowledge cutoff all point to the same working habit: put what matters into the conversation, and reach for web search when the question is about now.

LAB 01 ~20 MIN

First Contact — a five-move opening session

Run this live. It proves the "collaborator, not search box" idea by experience rather than assertion. Each learner does all five moves in a single conversation at claude.ai.

1. **Brief, don't query.** Ask Claude to do something real from the learner's own life — "Help me write a clear two-paragraph update to my team about a project running a week late." Notice it asks or assumes context, and produces a finished draft.
2. **Iterate.** Reply in plain language: "Warmer, and cut it to four sentences." Watch it revise rather than restart. This is the collaboration loop.
3. **Push its honesty.** Ask something it cannot know — "What did my city council decide at last night's meeting?" — and watch it flag uncertainty or offer to search instead of inventing.
4. **Switch register.** Ask it to explain a concept the learner finds intimidating "to a complete beginner," then "now in one sentence." Same model, adjustable depth.
5. **Reflect.** Each learner writes one sentence: "The thing that surprised me was ___." Collect these — they are your discussion fuel.

every learner has had Claude revise something based on feedback, and has seen it decline to bluff at least once.

TEACHING NOTES

How to teach Day 1 well

OPEN WITH THIS

Ask the room: "In one sentence — what is Claude?" Write the guesses up. Return to them at the end and let the group correct their own earlier answers. The visible before/after is the lesson.

PACE & EMPHASIS

Spend most of your time on Section 02 (collaborator vs. search box) and the Lab. Sections 03–05 are reference — keep them brisk. Do not let Day 1 drift into prompting tips; that's Day 3, and previewing it dilutes both.

DISCUSSION PROMPTS

· Where in your own work is a "collaborator" more useful than an "answer"? · What's one task you'd never trust to it unreviewed — and why is that the right instinct? · Which model tier fits the work you do most days?

COMMON MISCONCEPTIONS TO PRE-EMPT

"It's basically Google."
Reframe: Google finds pages; Claude does the work and revises it with you. Different tool, different verb.

"It knows everything / it's always right."
Introduce hallucination plainly and make verification feel like normal craft, not distrust.

"It remembers me."
Default is fresh-start. Memory and Projects are deliberate features — that's Day 2.

IF YOU ONLY HAVE 30 MINUTES Teach the one-sentence definition, the collaborator/search-box contrast, and run Lab moves 1–3. Everything else can be reading homework.

Day 1 Cheat Sheet

Claude	A family of large language models from Anthropic — an AI assistant you collaborate with in natural language.
Opus	Most capable tier — hardest reasoning, complex and agentic work.
Sonnet	Balanced workhorse — the sensible default for most tasks.
Haiku	Fastest, most economical tier — high-volume, latency-sensitive jobs.
Token	The chunk of text models read/write in — roughly ¾ of a word.
Context window	The conversation's working memory; large but finite.
Knowledge cutoff	The date Claude's training knowledge ends — use web search for newer facts.
Hallucination	A confident but false output. The fix is verification habits.
Constitutional AI	Anthropic's method of training Claude against explicit principles — helpful, honest, harmless.

Check for understanding

Five questions. Learners should be able to answer all five before Day 2.

1. Define Claude in one sentence, then name the single biggest difference between Claude and a search engine.
2. A teammate needs to classify 50,000 support tickets quickly and cheaply. Which model tier, and why?
3. What is a context window, and what practical habit does its finiteness imply?
4. What is a hallucination, and what is the correct response to the risk of one — fear, or something else?
5. Give one example of a question Claude can answer from training, and one that needs web search.

> **Answer notes —** 1) “An AI assistant from Anthropic you collaborate with in natural language”; it returns the work itself and iterates, vs. returning links. 2) Haiku — built for high-volume, latency-sensitive, low-cost tasks. 3) The conversation’s working memory; finite, so put what matters into the conversation. 4) A confident false output; the response is verification habits, not fear. 5) Training: “explain how compound interest works.” Search: “what is the current cash rate.”

Day 1 in five lines

- Claude is a family of large language models from Anthropic — a collaborator you direct in plain language.
- Opus / Sonnet / Haiku trade capability against speed and cost; default to Sonnet, escalate to Opus, scale with Haiku.
- It is excellent at drafting, explaining, analysing and coding — and needs supervision on exact facts and current events.
- Tokens, a finite context window, fresh-start conversations and a knowledge cutoff explain most of its behaviour.
- The master habit: put what matters into the conversation, verify what you can't check, and search for now.

TOMORROW — DAY 2 → Working in Claude.ai — Projects, Artifacts, memory & files

NOTES & DISCLAIMER

Independent resource. The Field Guide for Humans is an independent, unofficial educational resource produced and published by Kirevra Press, an imprint of Kyvara Pty Ltd (ACN 697 072 049). It is not affiliated with, endorsed by, sponsored by, or officially connected to Anthropic, PBC ("Anthropic"), the maker of Claude.

Trademarks. "Anthropic," "Claude," "Claude Code," and related names and marks are trademarks of Anthropic, PBC, used here for identification and descriptive purposes only. Their use does not imply any endorsement by or affiliation with Anthropic.

Accuracy & currency. The Guide describes third-party products that change frequently. All product details — features, model names, capabilities, commands, interfaces, and pricing — are believed accurate as of the 2026 edition but are provided "as is," without warranty of any kind, and are subject to change without notice. Always confirm current information against Anthropic's official documentation before relying on it.

No professional advice; no guaranteed results. The Guide is provided for general educational and informational purposes only. It does not constitute professional, legal, financial, or technical advice. No particular outcome, result, or level of proficiency is promised or guaranteed; results depend on the individual.