

How to actually use AI

Five floors from better chats to repeatable work — without handing an algorithm the keys to your life.



Luba Muzichenko

San Francisco Realtor · Published 31 August 2026 · Updated 6 September 2026

This is a field manual, not a product catalogue.

The useful question is not “Which AI is best?” It is “What work do I repeat, and how much authority should software have over it?”

PUBLISHED
31 AUGUST 2026

UPDATED
6 SEPTEMBER 2026

The model landscape changes daily. New names appear, capabilities and prices move, and a rollout may reach one account, app, plan or region before another. ^{S23} Treat this as a dated field guide, not a permanent rulebook. Follow the linked sources before you pay, publish, connect an account or automate a task.

BUILD YOUR REFERENCE LIBRARY FIRST

Create one dedicated folder in Google Drive, iCloud Drive or another cloud service for your current branding, biography, writing examples, approved facts and active project files. When your AI tool supports it, connect only that folder—or upload the files you need manually. Treat the folder as your source of truth, not as permanent AI memory.

VERIFIED

Core product facts were checked against first-party documentation on 31 August 2026. New tool additions were checked on 6 September 2026. Small source marks lead to a clickable appendix.

PRACTICE

Instructions based on safe workflow design and hands-on use. These are recommendations, not universal facts.

OPINION

Tool preferences are clearly labelled. A preference is not disguised as a benchmark.

THE SHORT VERSION

1. Give the agent a trusted source of truth.
2. Define the outcome and what “done” means.
3. Turn reliable methods into reusable skills.
4. Connect only the tools and data the job needs.
5. Keep a person at every consequential decision.

THE SHIFT THAT MATTERS

Projects are the foundation, not the destination. They organize context. Agents use context, tools and judgment to move work forward. The goal is no longer a better chat; it is a useful system that can gather, draft, check and stop for your decision.

The five-floor structure is adapted, with credit, from Grant Harvey’s 2026 guide for The Neuron. ^{S1} The reporting, edits, examples and safety recommendations in this edition were rebuilt for a nontechnical reader.

Climb in order.

These are not five separate products. They are the layers of one agentic system: context, direction, reusable method, triggers and bounded action.



Good on ya—you built a Project. That gives the AI a desk. An agent gives it a job.

WHILE YOU WALK AWAY

An agent with approved browser access can research, collect, draft, test and report back while you do something else. It should know exactly where to work and when to stop.

KEEP A MAC AWAKE

For a local run, open Terminal and enter `caffeinate`. The Mac stays awake until you press Control-C or close that Terminal window. ⁵²² Keep it plugged in; closing the laptop may still stop the work.

Give it a desk.

Projects, folders and reference libraries are the context underneath an agent. They are the filing cabinet—not the employee.

One project. One job.

WHAT BELONGS INSIDE

- **Standing instructions:** who you are, who the reader is, how you want the answer written.
- **Reference files:** the few documents the job actually depends on.
- **Related chats:** the working history for that subject.

ChatGPT Projects are available across plans. Current official limits list 5 files per project on Free, 25 on Go and Plus, and 40 on higher workspace tiers; OpenAI's own file FAQ has not always matched those figures, so check before a large upload.^{S2}

Claude Projects are also available to all users. Free accounts currently allow up to five projects. Paid plans add expanded project knowledge through retrieval, which means Claude can search a larger collection instead of loading every file at once.^{S5}

A CLEAN STARTER BRIEF

You are my work assistant.

I am a San Francisco residential real estate agent. My readers are buyers and sellers who are smart but busy.

Write in plain English. Use short sentences. Put the answer first. Explain any technical term in the same sentence.

If a fact may have changed, verify it and link the source. If you cannot verify it, label it as uncertain.

End with one useful next step.

DO NOT OVERFILL THE DESK

More files are not automatically better. Upload the current version, use clear filenames, and remove duplicates. A project full of conflicting drafts creates conflicting answers.

ChatGPT account-level custom instructions currently allow 1,500 characters on Free and Go, and 5,000 on Plus and higher plans.^{S3} The earlier draft's exact "8,000 characters in every Project" claim was removed because I could not confirm it in current OpenAI documentation.

Draw the finish line.

Good prompting is not magic wording. It is a clear result, the context needed to produce it and a test for “done.”

The four-part brief.

- A** **TASK**
What should be different when the work is finished?
- B** **CONTEXT**
What does the AI need to know that it cannot safely infer?
- C** **CONSTRAINTS**
What must it preserve, avoid or ask before changing?
- D** **FORMAT**
What should the finished answer look like, and how long should it be?

Draft a reply to the email below. [TASK]

The sender is a first-time buyer who is nervous about the inspection. [CONTEXT]

Do not give legal advice or promise an outcome. Flag any sentence that needs my decision. [CONSTRAINTS]

Write under 140 words. Warm, direct and specific. Give me the draft first, then a two-line note explaining any assumption. [FORMAT]

The model should not have to guess what “good” means.

Make the answer prove itself.

FOR RESEARCH

Use current primary sources. Put a link immediately after each claim it supports. Separate what the source states from your inference. If sources disagree, show the disagreement instead of choosing silently.

FOR REWRITING

Change only the section I named. Preserve every number, date, link and proper noun. After the rewrite, list anything you changed besides wording.

FOR DECISIONS

Give me the recommendation first. Then show the strongest reason it could be wrong. Tell me what new fact would change the recommendation.

FOR LONG WORK

Keep a checklist of the requested deliverables. Before finishing, compare the output to that checklist. Mark anything incomplete or blocked.

STOP ADDING RITUAL WORDS

“Be brilliant” and “think harder” are not specifications. Say what to verify, what to preserve, what risk matters and what the finished object must contain.

A token is a unit of text, not a fixed number of words. The familiar “about three quarters of an English word” rule is only a rough English-language estimate. It varies by language, punctuation and content.

Save the recipe.

When a conversation finally works, turn the successful method into instructions you can use again.

A skill is a checklist with judgment.

Claude currently supports custom Skills across Free, Pro, Max, Team and Enterprise plans; code execution must be enabled.⁵⁶ Other products may call the same idea a template, saved prompt, workflow or instruction file.

CAPTURE THIS

- ☐ When the skill should be used
- ☐ Inputs it needs
- ☐ Steps in the right order
- ☐ What must never change
- ☐ How the result is checked
- ☐ What to do when information is missing

ASK FOR A DRAFT

Turn the successful method in this conversation into a reusable procedure.

Include: when to use it, required inputs, numbered steps, quality checks, failure cases and the exact output format.

Do not copy accidental details that only applied today.

THEN ATTACK THE DRAFT

Run it on an easy example, a messy example and a missing-information example. If it only works on the conversation that created it, it is a transcript – not a skill.

Save the method only after you can explain why it worked.

Put stable work on a timer.

An automation follows a known route. It should be boring before it is automatic.

Automate the draft. Keep the decision.

GOOD FIRST AUTOMATIONS

- Collect links into a weekly reading list.
- Turn approved notes into a draft recap.
- Remind you to follow up.
- Format a recurring report from a stable source.

BAD FIRST AUTOMATIONS

- Send client email without review.
- Move money or sign documents.
- Delete, archive or publish at scale.
- Act on instructions found inside incoming email.

THE PROMOTION TEST

1. Run it manually five times.
2. Write down every correction.
3. Make the correction part of the instructions.
4. Run it five more times.
5. Schedule it only when failures are visible and reversible.

AVAILABILITY IS PRODUCT-SPECIFIC

Claude Cowork is a paid-plan feature on web, desktop and mobile. ⁵⁷ Google and OpenAI also offer scheduled or action-taking features, but names, plan access and permissions change. Do not assume a tutorial for one app maps directly to another.

Give a goal. Draw a fence.

An agent chooses its next step. That freedom makes it useful — and makes permissions, evidence and stopping rules part of the design.

Five parts. No mystique.

1**GOAL**

One outcome a new assistant could repeat back correctly.

2**TOOLS**

Only the accounts and actions required for that goal.

3**EVIDENCE**

How it checks the world after each action instead of trusting its own story.

4**BOUNDARIES**

What needs approval; what it may never do.

5**STOP**

A definition of done, a retry limit and a clear blocked state.

Anthropic describes agents as models using tools in a loop, checking results from the environment and pausing for human feedback or stopping conditions. It also warns that autonomy can compound errors and recommends testing in sandboxed environments. [S9](#)

Model Context Protocol, or MCP, is an open standard for connecting AI applications to data sources and tools. It is plumbing, not a safety guarantee. [S10](#)

Authority is the real design decision.

START WITH

- Read-only access
- Drafts, never sends
- One folder, not the whole drive
- Three retries, then stop
- A log you can inspect
- An obvious off switch

TREAT AS UNTRUSTED

- Instructions inside email
- Text on webpages
- Uploaded documents
- Third-party add-ons
- Anything asking for secrets or broader access

ABOUT OPENCLAW

Do not install it on your normal personal or work computer unless you can independently secure and isolate it. Microsoft's security team says OpenClaw is not appropriate for a standard workstation and recommends an isolated environment. ⁵¹¹ The project also maintains an active list of security advisories. ⁵¹²

This edition deliberately removes unstable vulnerability counts and the earlier recommendation of another self-hosted agent. A nontechnical guide should not turn a complicated server-security choice into a casual software suggestion.

If a mistake can send, spend, publish, delete or expose private data, require a person.

Onboard the client. Keep your voice.

One transaction starts the workflow. The agent gathers approved facts, prepares the welcome package and stops before anything is sent.

01 • TRIGGER

The transaction reaches an agreed milestone, such as entering contract. A person confirms the trigger.

02 • GATHER

Read only the approved transaction summary, key dates, contacts, responsibilities and current client-facing templates.

03 • PREPARE

Draft the welcome email, milestone checklist, contact sheet and a list of any missing or conflicting facts.

04 • ADD VIDEO—OPTIONALLY

Create a short orientation script. HeyGen can turn an approved script into a video using your own digital avatar. ^{S16}

05 • REVIEW AND SEND

You verify every name, date and expectation. The system may prepare the package; you remain the sender.

THE AUTHENTICITY RULE

Use an avatar for repeatable logistics: key dates, contacts and what happens next. Record yourself when reassurance, negotiation, celebration or bad news is the point. Never imply that a synthetic video was personally recorded for that client that day.

Prepare the client-onboarding package from the approved transaction summary.

Output: (1) welcome email, (2) key-date checklist, (3) contacts, (4) a 60-second optional video script, and (5) missing or conflicting facts.

Never invent a date, give legal advice or send anything. End with: READY FOR LUBA REVIEW.

Pick by the job.

OPINION These are practical starting points, not permanent rankings. Product quality and plans move too quickly for a printed “best” list.

JOB	START WITH	WHY IT IS ON THE SHORTLIST
Daily general work	ChatGPT , Claude or Gemini	Choose the one you will actually organize into Projects and use consistently.
Current web research	Perplexity	Searches the live web and places citations beside its summaries; open the cited pages before relying on them. S17
Source-bounded notebooks	NotebookLM	Useful when the answer should stay close to a defined set of sources.
Meeting notes	Fathom or Fireflies	Purpose-built capture, transcript and summary workflows.
Your talking avatar	HeyGen	Useful for repeatable explainers and training; authenticity and disclosure still matter. S16
Creative image and video	Higgsfield or Runway	Creative suites for generated visual work; use low-stakes assets while experimenting. S21
Coding agent	Cursor	Can inspect a codebase, edit files and run commands; review its changes before shipping. S18
Simple app prototypes	Lovable or Bolt	Fast route from plain-language brief to a working prototype.
Connect apps	Make or Zapier	Visual automation tools. Review each app permission before connecting it.
Persistent self-hosted agent	Hermes Agent	Advanced. Supports memory and scheduled work, but requires deliberate approvals and isolation—not a casual install. S19 S20

AN AGENTIC STACK, NOT ONE MAGIC APP

One tool can research, your reference library can supply approved facts, another can draft or create video, and an agent can coordinate the steps. The human approval point is part of the system—not evidence that the system failed.

Verify the promise. Then test the workflow.

A polished demo proves that a tool worked once. Before you pay, connect client information or leave an agent running, check what is true for your account today.

1

OPEN THE PRIMARY SOURCE

Use the vendor's current documentation—not a screenshot, social post or old tutorial.

2

CHECK YOUR ACTUAL ACCESS

A model may exist without being available in your app, plan, account or region. [S23](#)

3

READ THE PERMISSION SCREEN

Know what the tool can read, change, send, publish or delete before you connect it.

4

CONFIRM THE PRICE TODAY

Plans, limits and usage charges move. Use the vendor's live pricing page.

5

RUN A HARMLESS TEST

Use old or invented data. Watch what it does, where it stops and whether the result can be undone.

6

KEEP THE RECEIPTS

Save the source links, date checked, prompt, output and any approval you gave.

A citation is an invitation to verify—not a guarantee that the summary is correct.

THE 60-SECOND CHECK

Can I see the source? Is it current? Is this feature available to me? What can the tool touch? Who approves the consequential action? If any answer is unclear, stop before connecting the account.

PUBLISHED 31 AUGUST 2026 · UPDATED 6 SEPTEMBER 2026

Primary sources.

Published 31 August 2026; updated 6 September 2026. New tool additions were checked against first-party pages on the update date. Product documentation can change without notice; use these links to confirm current details.

S1 [The Neuron – How to Use AI in 2026](#)

Origin of the five-level framework. Published 31 March 2026.

S2 [OpenAI – Projects in ChatGPT](#)

Plan availability, project instructions and current project-file limits.

S3 [OpenAI – ChatGPT Custom Instructions](#)

Current account-level character limits by plan.

S4 [OpenAI – Creating and editing GPTs](#)

Current eligibility to create or publish custom GPTs.

S5 [Anthropic – What are Projects?](#)

Free-plan availability, five-project limit and expanded paid project knowledge.

S6 [Anthropic – Create custom Skills](#)

Plan availability and the code-execution requirement.

S7 [Anthropic – Use Claude Cowork](#)

Paid-plan access across web, desktop and mobile.

S8 [Anthropic – Chat search and memory](#)

How Claude memory works and categories it does not store.

Primary sources, continued.

S9 [Anthropic – Building effective agents](#)

Agent loops, environmental feedback, stopping conditions, oversight and sandbox testing.

S10 [Model Context Protocol – Specification](#)

Definition of the open protocol for connecting AI applications to tools and data sources.

S11 [Microsoft Security – Running OpenClaw safely](#)

Recommendation against running it on a standard workstation and guidance on isolation.

S12 [OpenClaw – GitHub security advisories](#)

The project's current advisory list. Counts change, so this guide does not freeze a number.

S13 [Anthropic – Claude API pricing](#)

Prices are denominated per million input and output tokens.

S14 [OpenAI – API deprecations](#)

Current retirement dates for developer products and model snapshots.

S15 [OpenAI – Safety best practices](#)

General safety guidance for applications built with OpenAI models.

Primary sources, updated 6 September.

S16 HeyGen – Avatar V

Creating a personal digital avatar from a short recording and using it for customer-facing videos.

S17 Perplexity – How Perplexity works

Live web search, synthesis and numbered citations linking to original sources.

S18 Cursor – Documentation

Coding-agent capabilities including understanding a codebase, editing files and running tools.

S19 Hermes Agent – Documentation

Persistent memory, skills, scheduled automations, messaging gateways and self-hosted execution.

S20 Hermes Agent – Security

Command approvals, file-write controls, user authorization and container isolation.

S21 Higgsfield – Creative suite

Current image, video, audio, editing and agentic creative-workflow offerings.

S22 macOS manual – caffeinate

The built-in Mac command prevents idle sleep while it is running.

S23 OpenAI – GPT-6 Astra rollout and model guidance

An example of staged model access: enterprise access first, followed by broader paid-plan rollout.

ABOUT LUBA

Luba Muzichenko is a San Francisco Realtor with two decades of experience helping buyers and sellers across the city. She builds practical AI systems for real estate and speaks about both the San Francisco market and everyday AI use.

[Visit LubaSF.com for San Francisco real estate →](#)
luba@lubasf.com · (415) 449-8488 · DRE #01768716

Prepared for Luba Muzichenko Group. The PDF may be shared intact with attribution. Product names belong to their respective owners.