How to actually use AI

A plain-English guide for smart, busy people who want useful AI without becoming software engineers.

Luba Muzichenko
San Francisco Realtor · Published 31 August 2026 · Updated 8 September 2026

START HERE / NO TECHNICAL BACKGROUND REQUIRED

If you can send an email, you can use AI.

AI is a helper. Start by giving it one ordinary job you already understand. Then check the work before you use it.

PUBLISHED
31 AUGUST 2026

UPDATED
8 SEPTEMBER 2026

AI changes every day. Model names, prices and features can change without warning. A feature may reach one person before another.[S23] Check the linked source before you pay, publish, connect an account or let AI do a task on its own.

WHY I WROTE THIS

I am a San Francisco native and Realtor with 20 years in the business. I am not a software engineer. My work is full of research, explanations, emails, checklists, marketing and details that cannot be dropped.

I use AI to help with the repeatable parts. I do not hand over my judgment, my voice or my responsibility to a client.

This is the guide I would hand another smart, busy professional who wants a useful result, not a new hobby.

YOUR FIRST USEFUL WIN

1. Pick a small task you already know how to do.
2. Say who it is for and what they need.
3. Ask for a draft, not a final decision.
4. Check every name, date, number and promise.
5. Use it only if it is true and sounds like you.

MAKE ONE FOLDER YOU CAN TRUST

Keep your current bio, brand rules, writing examples, approved facts and active project files in one folder in Google Drive, iCloud Drive or another service you trust. Give AI only the files needed for the job. The folder is reliable memory. AI is not.

Already using Projects? Good on ya. Start there anyway. The next steps only work when your basic information is clean. The five-step structure is adapted, with credit, from Grant Harvey's 2026 guide for The Neuron.[S1]

Start at one.
Stop when it gets useful.

Each step adds one idea. You do not need all five. Most people can get a lot from Projects and clearer instructions.



> Projects and prompts are not beginner mode. They are where useful AI begins.

Project: one place for one kind of work.

Prompt: the instruction you give AI.

Skill: a saved method that worked.

Automation: a known task repeated by itself.

Agent: AI that uses tools and takes several steps. This is the advanced end.

FLOOR ONE / FOUNDATION

01

Give it a desk.

A Project is one place for the files, instructions and conversations that belong to one kind of work.

Start with one Project.

WHAT BELONGS INSIDE

- **A short description:** what this Project is for.
- **Your preferences:** how you want the answer written.
- **The right files:** only the current documents for this work.
- **Related chats:** conversations you want to keep together.

A USEFUL FIRST PROJECT

Create one called **Client Onboarding**. Add your current welcome email, contact list, key-date checklist and one good example written in your voice. Do not add every transaction you have ever handled.

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 and Claude both offer Projects. Access and file limits can change, so check the current help page before a large upload. [S2] [S5] ChatGPT also has account-wide writing preferences called Custom Instructions. [S3]

FLOOR TWO / CONVERSATION

02

Draw the finish line.

A prompt is simply your instruction. Say what you need, who it is for and what the AI must not guess.

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 sources from the company or agency responsible. Put a link right after each claim. Separate what the source says from what you think it means. If sources disagree, show me the disagreement.
```

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 useful instructions. Say what to verify, what to preserve, what could go wrong and what the finished answer must include.

03

Save the recipe.

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

A skill is a saved way of doing a job.

When a prompt works more than once, save the method. It might be a checklist, template, saved prompt or instruction file. Different tools use different names for the same basic idea.

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 an example with missing information. If it only works once, it is not ready to save.

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

Claude currently has a feature called Skills. Access can depend on your plan and settings. [S6]

FLOOR FOUR / REPETITION

04

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.

An automation repeats a known task at a set time or after a clear event. Start with work you would be comfortable doing again tomorrow in the same way.

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.

THE EXACT FEATURE DEPENDS ON THE APP

Claude Cowork is a paid feature on the web, desktop and mobile.[S7] Google and OpenAI offer similar tools, but the names, prices and permissions change. Instructions for one app may not work in another.

05

Give a goal. Draw a fence.

An agent can choose what step comes next. That makes it useful. It also means you must control what it can touch and when it must stop.

An agent takes more than one step.

It can use tools, look at the result and decide what to do next. That is different from a normal chat, where you give each instruction yourself.

1

GOAL

One outcome a new assistant could repeat back correctly.

2

TOOLS

Only the accounts and actions required for that goal.

3

CHECK

How it confirms what happened after each action.

4

LIMITS

What needs your approval and what it may never do.

5

STOPPING POINT

What done means, how many times to retry and what to do if it gets stuck.

ONLY WHEN YOU GET THIS FAR: LET IT WORK WHILE YOU ARE AWAY

An approved browser agent can research, collect, draft and test while you do something else. Give it one narrow job and a clear stopping point. For a local Mac run, the built-in `caffeinate` command can keep the computer awake. [S22] A separate computer or always-on server is advanced. It does not belong in the opening lesson.

Technical tools can connect AI to other apps and files. A connection does not make an agent safe. Test it with harmless information first. [S9] [S10]

Decide what it is allowed to do.

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. [S11] The project also maintains an active list of security advisories. [S12]

This kind of software can reach files, accounts and other tools. That makes a mistake much more serious than a bad chat answer. Leave it alone unless someone qualified can set it up and watch it.

> 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 process. 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 AN OPTIONAL VIDEO
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.
Web research with sources	Perplexity	Searches the live web and puts citations beside its summaries. Open the cited pages before relying on them. [S17]
Ask questions about your files	NotebookLM	Useful when the answer should stay close to a set of sources you chose.
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]
Edit code or a website	Cursor	Can read a site, edit files and run checks. Review the changes before publishing. [S18]
Build a simple app from words	Lovable or Bolt	A fast way to turn a plain-English idea into a working test.
Connect apps	Make or Zapier	Visual automation tools. Review each app permission before connecting it.
Advanced, always-on agent	Hermes Agent	**Later, not now.** It can remember things and run scheduled work. It should live on an isolated computer or server, not your everyday machine. [S19] [S20]

YOU DO NOT NEED ONE MAGIC APP

One tool can research. Your folder can hold approved facts. Another tool can draft or create video. An agent can connect the steps later. Your review is part of the process.

Check the important parts. Ignore the hype.

AI can sound certain when it is wrong. Before you pay, publish or use client information, check the few facts that could hurt someone.

1

OPEN THE COMPANY'S OWN PAGE

Use its current help page, 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 source link lets you check the answer. It does not prove the answer is right.

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 important action? If any answer is unclear, stop before connecting the account.

PUBLISHED 31 AUGUST 2026 · UPDATED 8 SEPTEMBER 2026

Primary sources.

Published 31 August 2026; updated 8 September 2026. Product claims were checked through 6 September. The guide was rewritten for clarity on 8 September. Product pages can change without notice, so 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.

Product claims checked through 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
Memory, saved methods, scheduled work and running the software on your own computer or server.

S20 Hermes Agent: Security
Command approvals, file-write controls, user authorization and container isolation.

S21 Higgsfield: Creative suite
Current image, video, audio and editing tools.

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 native and Realtor with 20 years in the business. She uses AI in real work, including research, writing, marketing and client communication. This guide comes from testing what helps, what wastes time and what still needs a human.

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.