

FROM PACIFIC FORWARD INITIATIVE · PACFORWARD.ORG

The Field Guide to AI

AI for your real work.

‘A‘ohe pau ka ‘ike i ka hālau ho‘okahi

All knowledge is not learned in one school.

Pacific Forward Initiative is a 501(c)(3) nonprofit. We make free, plain-language materials and run free local classes on practical AI skills. We do not charge for this guide. We do not ask for your email to download it.

BEFORE YOU START

How to read this

This is a guide, not a textbook. Skip around.

It is meant to live on your desk or your phone — open it when you need it, close it when you don't. You do not have to read it in order. You do not have to finish it.

THREE WAYS IN

You have never used AI. Start at Chapter 1, "What AI is and isn't."

You have tried ChatGPT a few times. Start at Chapter 3, "Spotting AI's confident mistakes." That is the most important chapter in this guide.

You attended an AI for Real People live class. Pick up at Chapter 6, "AI for paperwork." Everything before that is review.

Each chapter is short — about 5 minutes of reading.

If a sentence here uses a word you do not know, look at the glossary at the back. If a sentence here makes you feel that we are talking down to you, we did. Email us at hello@pacforward.org. We will fix it.

EVEN IF YOU READ NOTHING ELSE

Four prompts you can copy

These are four things to try, even if you read nothing else in this guide. Copy each one into Claude or ChatGPT. Paste your document, message, or question where it asks. Then check the answer the way the last line tells you.

1 • Find the key points in a document

For pulling the most important parts out of a document when you are short on time.

Read this document and give me the three most important points in plain language. Then quote the exact words from the document that support each point.

Check the answer by: looking at the quoted words. If they are missing or paraphrased, the AI guessed instead of read.

2 • Translate something dense into plain English

For turning a confusing form, letter, or policy into language you can act on.

Read this and rewrite it for someone who has never seen this kind of document before. Use short sentences. After your rewrite, list any places where you were not sure what the original meant.

Check the answer by: reading the "not sure" list. If it is missing or empty, ask the AI to read again — most documents have at least one ambiguous part.

3 • Draft a reply you have been putting off

For getting past the blank page on an email or letter you have been avoiding.

Here is the message I received: [paste it here]. Help me write a professional reply that gets to the point of [your main point in one sentence]. Show me two options — one shorter, one warmer.

Check the answer by: reading both versions out loud before you send. If a sentence does not sound like you, change it before you send.

4 • Ask better questions about a topic you are new to

For getting up to speed without having to start from a question you would rather not ask out loud.

I am trying to learn about [topic]. Ask me three questions to figure out what I already know and what I most need. Then explain the most important thing for me to understand first.

Check the answer by: noticing whether the AI's three questions made sense to you. If they assumed knowledge you do not have, tell it that and try again.

GETTING STARTED — HOW TO USE EITHER TOOL

Using Claude (from Anthropic). Go to claude.ai. Sign in with an email address. Type or paste into the message box at the bottom. The free tier is enough for everything in this guide.

Using ChatGPT (from OpenAI). Go to chatgpt.com. Sign in with an email address. Type or paste into the message box at the bottom. The free tier is enough for everything in this guide.

Either works. The prompts above are written to work in both. If you already use one, stay with what you know.

What AI is and isn't

You have used AI for years.

Every time Google Maps reroutes you around traffic, every time Gmail finishes your sentence, every time Netflix knows you want a third episode of that show — that is AI doing its job behind the scenes.

What changed in the last few years is that you can now *talk to it directly*. What was in the background became something you can talk to. That is the whole change.

So what is it, in plain language?

Today's AI is a tool that recognizes patterns in language. You give it words; it gives you back the words it predicts would most likely come next. There is no thinking behind the curtain. It is not a person. It is not deciding things about you. It is anticipating.

Which means it is very good at the things humans do with language — writing, summarizing, rewriting. It is bad at the things humans do *not* do with language alone — knowing facts it was never shown, remembering what you said three weeks ago, telling the difference between a sentence that is true and a sentence that just sounds true.

A note on a few words you will hear: *algorithm, model, training data*. They are technical names for parts of how the tool works. You do not need to know them to use the tool. They are defined in the glossary at the back if you are curious.

Why this matters for you

You are not behind. You are not late. The tools are new enough that nobody is more than a few hours ahead. What you need is not to "catch up." What you need is a way of working with this that does not waste your time and does not talk down to you. That

is what this guide is for.

Today, Soon, On the Horizon

WHEN	WHAT YOU CAN DO WITH IT
Today	Summarize a document. Draft a reply. Rewrite something in simpler language. Brainstorm a list. Translate. Ask follow-up questions until you understand.
Soon	Sit in on a meeting and write the notes. Pre-fill your forms. Answer common questions from your customers. Cross-check a document against another document.
On the Horizon	Take action on your behalf (book the flight, file the form, schedule the meeting). Most of this works in narrow cases today and will broaden over the next year or two.

Trust, but verify. The single most important habit in this guide is checking the answer the AI gives you — Chapter 3 is about that habit specifically.

Your first conversation with AI

There is a single right way to start.

Step 1. Open claude.ai or chatgpt.com in a browser. Sign in with an email.

Step 2. Type something you actually want help with. Not a test question. A real one.

A starting prompt that almost always works:

COPY AND PASTE

I am working on [a one-sentence description of what you are working on]. Ask me three questions to figure out what I most need help with right now. Then help me with the most important one.

That is the whole first conversation. The AI will ask three questions. Answer them. It will then give you something — a draft, a list, an explanation. Read it.

What to do if it gives you something wrong

Tell it. Say: *"That is not what I needed. What I actually needed was X. Try again."*

The AI will not be offended. It will not remember the mistake an hour from now. It is not a person. Correcting it is part of using it well — possibly the most important part. People who get good results from AI are not better at typing. They are better at telling it when it is wrong.

What to do if it gives you something right

Notice what worked. Try the same kind of prompt again next time you need help with that kind of task.

One thing to know

The AI will not remember this conversation tomorrow. Every conversation starts from zero. If something worked, save the prompt somewhere you can find it again. The back of this guide has a prompt notebook for exactly that.

Spotting AI's confident mistakes

A contract came back from the AI with four crisp bullet points. Three were right. The fourth — a renewal clause that would have cost \$4,800 over two years — was made up.

The AI said it confidently. It did not say "I am not sure." It did not flag the part it invented. It wrote the invented part the same way it wrote the three correct ones.

Whoever was reading caught it because they read the contract themselves before signing. That is the habit this chapter is about.

The single most important thing in this whole guide

Today's AI tools are very good at sounding right. They are sometimes wrong. They will almost never tell you which is which. *Your job is to check.*

This is not because AI is broken. This is how the tool works. It does not know the difference between a sentence it is sure about and a sentence it invented to fill in a gap. If it did, it would not invent things.

Check the same way you would check a draft from a brand-new colleague — one who is smart and willing but has been at the job for one day. You like them. You give them work. You read what they hand back before you send it on.

THREE HABITS THAT CATCH MOST MISTAKES

1. Ask for quotes. When the AI summarizes a document for you, ask it to quote the exact words that support each point. If the quotes are missing or paraphrased, it guessed.

2. Ask the same question twice. Open a new conversation. Ask the same question again. If you get a very different answer, neither one is reliable. Ask again until the two answers agree on the parts that matter.

3. Ask what it does not know. Add to the end of any prompt: *"What parts of this are you not certain about? Mark them."* This does not catch everything, but it catches more than you expect.

What "made up" looks like

Names of books that do not exist. Court cases that were never filed. Quotes from people who never said the thing. Statistics that sound plausible. URLs that go to 404 pages. The pattern is almost always the same: very specific, very confident, very wrong.

The thing to remember: the more specific and confident the AI sounds, the more carefully you check.

A risk + what you can do

RISK

WHAT YOU CAN DO

AI invents a fact in a document you send to someone important

Read the AI's output side-by-side with the source.
Use the "ask for quotes" trick.

AI summarizes a long policy and misses the one clause that matters to you

Tell the AI specifically what you are looking for first ("Find any clause about renewal terms"), then check the result against the document.

AI drafts an email that sounds like you but says something you would never say

Read it out loud before sending. If a sentence does not sound like you, change it.

AI gives you a wrong answer with such confidence you don't think to check

This is the most dangerous one. Build the habit of always checking, even when you are sure.

Trust, but verify. Use the tool. Check the work. The check takes 30 seconds and saves you from the \$4,800 mistake.

AI for your inbox

Here is a 20-minute routine that you can do every Monday morning. It will not solve your inbox. Nothing solves anyone's inbox. But it will give you the top of it back.

What you need: an AI chat tool open in one tab, your inbox open in another.

The routine, step by step

Step 1 (5 minutes). Triage what is there.

Copy the subject lines of your last 20 unread emails into the AI chat. Type:

COPY AND PASTE

Group these into three categories — needs a reply this week, can wait, does not need a reply. Show me which is which.

Read the groups. If the AI got one wrong, move it. The point is not for the AI to be perfect. The point is to have a sorted list in front of you in 5 minutes instead of 25.

Step 2 (5 minutes). Draft replies to the "needs a reply this week" pile.

For each one that needs a real reply, paste the email into the AI and use Prompt 3 from the front of this guide ("Draft a reply you have been putting off"). Read the draft. Edit it. Send it. Or save it as a draft to send later.

Step 3 (5 minutes). Acknowledge the "can wait" pile.

For each one, paste the email and type:

COPY AND PASTE

Write a 2-sentence reply that acknowledges this and says I will get back to them by [date].

Send those. Now those people are not waiting on you in silence.

Step 4 (5 minutes). Archive the "does not need a reply" pile.

You do not need AI for this. You just need to stop reading those emails again next Monday.

What the routine actually buys you

The point of the routine is not productivity. The point is that on Wednesday morning, you do not still owe somebody an answer from last Friday. The relationships hold. The work does not pile up on you in the same way. That is the time-back the title of this guide is about.

A NOTE FOR HAWAII INBOXES

A lot of working adults here run side conversations on email that are personal as much as professional — a question about a meeting also includes asking how someone's mother is doing. The AI does not know that context. If a reply needs warmth before it gets to business, write the warmth yourself in one sentence, then ask the AI to draft the business part below it. The warmth has to be yours.

AI for conversations that matter

Some conversations are not about productivity. You can use AI for these too — not to have the conversation for you, but to help you show up to it better prepared.

Using AI to ask better questions

If you sit down with your grandfather and ask "tell me about your life," you will get a short, polite answer.

If you ask: *"Where were you working when you met grandma? Did you like the job? What did you wear to work?"* — you will get an hour.

The AI can help you draft those second kinds of questions. Try this prompt:

COPY AND PASTE

I am going to interview my [relationship] who lived in [Hawaii / a place] from [year] to [year] and worked as a [job]. Give me 8 questions that would help them remember specific moments, not just general impressions. Mix some practical questions (what they ate, what they wore) with some emotional ones (what they were afraid of, what they were proud of).

Read the 8 questions. Pick the 5 you want to actually ask. Bring them on paper, not on your phone — phones change the conversation.

What to do during the conversation

Record it if they say yes. Most phones have a voice memo app. Hold it on the table between you and forget about it. You will remember details from the recording that you do not remember from the conversation, and the recording becomes a thing your family can hear in 30 years.

What to do after

Paste the transcript (or your notes) into the AI. Try this prompt:

COPY AND PASTE

Here are my notes from a conversation with my [relationship]. Pull out three stories that are worth writing down. Quote the parts that have the most specific detail.

Use that as the starting point for a written-down version you can send to your family.

A NOTE ON CONSENT

Ask before you record. Ask before you put anything they said into a chat with an AI. Most people will say yes if you explain why. Some will say no. Respect the no. The conversation is more important than the recording.

AI for paperwork

Most of us have a letter we have been avoiding. A renewal form. A re-verification notice. A medical bill that does not match what we expected. A landlord asking for something. An agency saying our coverage might end.

The AI cannot answer the letter for you. But it can do three things that make answering it possible.

- **It can read it with you.** Even if you have read it three times and it still does not make sense.
- **It can help you draft a response.** Even if you do not know where to start.
- **It can flag what to check.** Before you send anything back.

Move 1 — Read the letter with the AI

Open the letter. Open Claude or ChatGPT. Type:

COPY AND PASTE

I received this letter and I am not sure what it actually wants me to do. Read it for me. Then tell me:

1. In one sentence — what does this letter want from me?
2. By when?
3. What happens if I do nothing?
4. What parts of this letter are unclear or might mean different things?

Paste the letter below. Read what the AI says back. The first three answers should be short and specific. If the AI hedges on any of them — "the letter seems to suggest that you may need to..." — that is a sign that the letter itself is ambiguous, not that the AI is bad at its job. Ambiguity in the letter is the most important thing to know about. That is what the fourth question is for.

Move 2 — Draft a response

Once you understand what the letter wants, use this prompt:

COPY AND PASTE

Help me draft a response to this letter. The main thing I want to say is: [one sentence of what you actually want to say — examples: "I am attaching the documents they asked for," "I am asking for a 30-day extension," "I am disputing the amount they say I owe"].

The response should be calm, polite, and short. Show me two options — one more formal, one more conversational.

The AI will draft two versions. Pick the one that sounds closer to how you would actually write. Edit it until every sentence is a sentence you would actually send.

Move 3 — Check the draft before you send

This is the move that catches problems. Try this prompt:

COPY AND PASTE

Read what I am about to send. Tell me:

1. Am I giving them information they did not ask for?
2. Am I leaving out information they did ask for?
3. Is there anything that could be misread?
4. Anything that sounds defensive or angry when I did not mean it that way?

Read the AI's answers. Fix what needs fixing. Then send.

A REAL EXAMPLE — A RE-VERIFICATION LETTER

Imagine a letter from a Hawaii state agency saying that your enrollment in a benefit program is up for re-verification and you need to send proof of income within 30 days. The letter has three pages and uses the words "ineligibility determination" and "supplemental documentation" and "case action notice."

Move 1 — reading it with the AI — gives you back: *"They want you to send three things. Pay stubs from the last 30 days. A current utility bill in your name. A signed and dated form that is in section 4 of the letter. If you don't send these by the date on page 1, your benefit ends on that date. The unclear part is whether the form on page 4 needs to be notarized — the letter does not say either way."*

That is now actionable. You know what you have to gather. You know the deadline. You know the one question to call the agency about (whether the form needs to be notarized).

Move 2 — drafting the response — turns the bureaucratic letter into a transmittal note in your voice: *"Enclosed please find the requested documentation: pay stubs (Aug-Sep), August utility bill, and signed form. Please let me know if the form requires notarization."*

Move 3 — checking before sending — flags anything you missed. That is the whole workflow. Read with the AI. Draft with the AI. Check with the AI. Then you send.

Why this matters

Letters from agencies, doctors, landlords, lawyers — these are the documents that scare people the most. They are full of words designed to be precise and not friendly. People put them off. Then the deadline passes. Then a coverage ends, or a fee gets assessed, or a small problem becomes a big one.

This is the chapter the rest of the guide is for.

What AI can't do for you

The AI cannot

Verify what is true about the world right now. If you ask "is this restaurant still open" or "is this person still in their job" or "did this happen this week," the AI does not know. It can guess based on patterns. The guess is sometimes right and sometimes very confidently wrong.

Make decisions for you. It can lay out options. It cannot tell you what you actually want. If you ask "should I take this job," the AI's answer is somebody else's average opinion about job-taking. It is not yours.

Replace expertise that matters. For medical, legal, tax, immigration, or financial decisions that affect your life: use the AI to understand what is going on, then talk to an actual professional. The AI is a tool for thinking out loud. A doctor is a tool for getting healthy. They are different tools.

Remember anything. Every conversation starts from zero. If you told the AI about your project on Monday, it does not know on Tuesday. Save what mattered. The prompt notebook at the back of this guide is for that.

Apologize sincerely. Some things the AI does are mistakes. It will say "you are right, I apologize." It does not mean it. There is nothing there to mean it. Notice the mistake; do not need the apology.

The AI also cannot tell you

When you should not use it. When you are using it for something that would be better handled by a real conversation with a real person. When the convenience of asking the AI is shutting down a chance you would have had to slow down and think.

That part is on you. The tool is good. The judgment about when to use it is yours. You are the one this guide is for.

APPENDIX A

Plain-English glossary

Words you will hear about AI, in plain English. Skim this once. Do not memorize.

AI / Artificial Intelligence. A tool that recognizes patterns in language, images, or data. The kind of AI in this guide is the chat kind — Claude and ChatGPT.

Algorithm. A set of steps a computer follows to do a task. Search engines use algorithms to rank results; map apps use algorithms to pick routes. The "AI algorithms" people worry about are mostly the recommendation kind.

Model. The specific version of an AI tool. ChatGPT runs on different models depending on which plan you pay for. You usually do not need to know which.

Training data. The examples a tool was shown to learn what good answers look like. The AI you use today was trained on a very large number of internet texts.

Chatbot. Another word for the chat-style AI tools. Claude and ChatGPT are chatbots.

Prompt. What you type into the chat. The prompts at the front of this guide are the kind that work best.

Some words you will hear that we do not use in this guide because they are not necessary: *LLM, transformer, neural network, hallucination*. If you read them somewhere else, just substitute "the AI tool" or "the AI being wrong."

Where to learn more

From PFI

- **AI for Real People** — free 30-minute live classes. lu.ma/pacforward
- **AI for Real People: Paperwork** — the online course this guide pairs with. pacforward.org/courses
- **The PFI blog** — short writeups of tools to try. pacforward.org/blog

Free programs from others

- **AARP Senior Planet** — AI for Older Adults guide. seniorplanet.org
- **Khan Academy** — free AI learning resources for all ages. khanacademy.org
- **Hawaii State Public Library System** — check your local branch for free AI workshops.

One book, if you want one

***Co-Intelligence* by Ethan Mollick.** A short, practical, plain-language read. Honest about what AI can and cannot do.

YOURS TO FILL IN

Your prompt notebook

Use this space to write down prompts that worked for you, so you can use them again. The AI will not remember; this notebook will.

WHAT I WAS TRYING TO DO

THE PROMPT THAT WORKED

Pacific Forward Initiative

A Hawaii 501(c)(3) nonprofit. We make free, plain-language materials and run free local classes on practical AI skills.

We do not charge for this guide. We do not ask for your email to download it. If something here helped you, tell a friend. If something here was wrong, tell us at hello@pacforward.org.

Accessibility: this guide is published as a tagged PDF and meets WCAG 2.1 AA standards. If you need an alternative format, write to accessibility@pacforward.org.

FORWARD TOGETHER