

AI FOR REAL PEOPLE · A SHORT COURSE FROM PFI

Getting Started

A short course for people who have never typed into an AI tool, or tried once and stopped. No jargon. No prerequisites. Just plain answers to the questions you'd actually ask.

About 30 minutes Free Five short lessons No videos No quiz you can fail

You will actually try things in this course — open an AI tool in another tab, type one real thing, read what comes back. The people who try get the most out of it.

IN THIS PACKET

1	Welcome — is this course for you?	4 min
2	Try it once	6 min
3	What it actually is	7 min
4	Two habits and one trap	6 min
5	Privacy, signal, and what's next	7 min
+	The most important page, the quick-reference card, and the prompts in one place	

PACIFIC FORWARD INITIATIVE

Pacific Forward Initiative is a 501(c)(3) nonprofit. PFI is a Hawaii-based nonprofit helping individuals, communities, and small businesses build the future-defining skills needed to thrive in a changing economy.

We make artificial intelligence approachable, practical, and human-centered — so people can create new opportunities for themselves, their families, and their communities.

This course is free, and it stays free.

Welcome

If you've never typed into an AI tool, or you tried once and closed the tab — this is for you.

We built this course for the person who feels like everyone got a manual for AI that they missed. There is no such manual. What we're going to do together is small, plain, and honest. Not a lot. Enough.

What you'll leave with

Three things will be true when you finish:

- 1 **You'll have typed into an AI tool.** No more first-time nerves.
- 2 **You'll know one honest sentence about how it works.** No jargon required.
- 3 **You'll have one small habit for catching when it's wrong.**

That's it. Not a lot. Enough.

What this course is *not*

Just so we're clear on the ground:

- **Not a sales pitch** for any specific AI tool. We'll show one, but pick whichever you want.
- **Not a promise** that AI will change your life. Some things will help you. Some won't. You get to decide.
- **Not a test.** Nothing here is graded. Nothing here counts against you.

A SHORT SELF-CHECK BEFORE WE BEGIN

Three quick scenarios. Ask yourself if any of them sound like you.

Scenario 1. Someone at work mentioned they use AI for something ordinary, and you nodded along without asking what they meant.

Scenario 2. You opened one of these tools once, looked at the empty box, could not think of what to type, and closed the tab.

Scenario 3. You are not sure whether the thing is reading your email, keeping your files, or learning about you — and you have not found a plain answer.

If any of those sound familiar, you are in the right course. Keep going.

A note about the live version

Everything in this course also runs as a free live 30-minute class on Zoom. Alex walks through the same demo in real time, answers questions in chat, and the whole thing is over in half an hour. If you would rather be walked through it than read it, sign up at lu.ma/pacforward.

If you would rather just read your way through, that works too. Keep going.

BEFORE LESSON 2 — SET UP AN AI TOOL

This course is tool-agnostic. It works with any popular AI chat tool — Claude, ChatGPT, Copilot, or Gemini. If you do not already use one, take two minutes to set up a free account with whichever you like. The free versions are enough for everything in this course.

Try it once

The whole skill is: type something in plain English, like you'd ask a person, and read what comes back. That is the entire lesson. The rest is you doing it.

TRY IT NOW

Open your AI tool in another tab. You are going to ask it for one small, ordinary thing — a thank-you note. Nothing is riding on this. Nobody sees it but you.

Type this exactly:

COPY AND PASTE

Write a short, warm thank-you note to a colleague who covered a meeting for me last week. Friendly, professional, not overly formal.

Send it. Read what it gives you.

Then change one thing about the request and try again — maybe change "colleague" to a specific role, or add "make it two sentences long." Watch what changes.

Stop here. Actually do this before reading on.

WHAT YOU SHOULD HAVE SEEN

A short note, in reasonable English, that you could send with light editing. Probably a little more formal than you would have written. Probably one phrase you would change.

And when you changed the request, the note changed with it — without you explaining what a thank-you note is or how one is structured.

That's the whole skill. You didn't program it. You didn't tell it a formula for how thank-you notes work. You asked in plain English, and it answered.

A QUICK QUESTION FOR YOURSELF · NOT GRADED, NOT SAVED

What did you change on the second try, and did the answer change the way you expected? The gap between what you asked for and what you got back is the thing worth noticing — it is where all of your future editing happens.

SPACE FOR YOUR NOTES ON WHAT HAPPENED

TRY TO REMEMBER · TWO QUESTIONS, UNGRADED

Try answering each one yourself before reading ours. Your answers aren't saved.

1. What did you have to know about AI before typing that prompt?

Our answer: Nothing. That is the point of starting here. The skill is asking in plain English, not knowing the machine.

2. Why start with a thank-you note instead of something that matters?

Our answer: So you learn the pattern — ask, read, adjust, ask again — somewhere nothing bad happens if the answer is wrong.

The free live class runs this same demo live, and you can watch someone else's prompt get answered before you try your own. lu.ma/pacforward

What it actually is

Here is the one sentence to remember from this whole course:

It learned what good answers sound like. Even when it looks something up on the web, it tells you in a voice that sounds sure — so you keep the last word.

Read that once more.

That's the whole thing. It's excellent at how a good answer *sounds*. That's why it's so useful. It's also why it can be so sure — and sometimes so wrong. **Sounding sure and being right are not the same thing.**

The autocomplete analogy

You know when you're texting someone and above the keyboard your phone guesses your next word? You type "I'm running a little—" and it offers "late." It's not reading your mind. It has just seen that word come next, millions of times.

Now picture that same guesser. But instead of just your text history, it read almost everything ever written — books, articles, forum posts, recipes, letters. All of it.

That's this.

The honest fix (the "am I being spied on" question)

Here's the part that keeps you from being scared of the wrong thing:

The AI tool did **not** save all that writing. It did not keep the books. It kept the *patterns* of how we write.

It is not looking things up in a library it keeps in a drawer. **It is not holding your emails.** It's a pattern machine that got very good at sounding like a person who read a lot.

Where it trips up

The tool never says "I don't know."

So sometimes it makes something up — with a completely straight face. This is not the tool being broken. This is how it works.

A confident guess that happens to be wrong sounds exactly like a confident guess that happens to be right.

That's the whole reason we need one small habit — which is the next lesson.

A QUICK QUESTION FOR YOURSELF · NOT GRADED, NOT SAVED

Go back to the thank-you note from Lesson 2. Was there a single hard fact in it — a name, a date, a number? If there had been, would you have caught it being wrong?

TRY TO REMEMBER · TWO QUESTIONS, UNGRADED

1. Why does it sound sure even when it is wrong?

Our answer: Because it learned what good answers sound like. Confidence is part of that sound. It is not a signal about whether the answer is true.

2. Is it holding your emails?

Our answer: No. It kept the patterns of how we write, not the writing itself — which is a different thing from what you type into it today. That part is Lesson 5.

Two habits and one trap

Two habits catch most of what goes wrong. Then there is one trap that catches confident people specifically, and it is worth knowing before it gets you.

Habit 1 — Trust it for shape, check it for facts

It's wonderful at the *shape* of things. A note. A summary. A plan. Trust it there.

The moment it gives you a hard fact — a name, a date, a price, a dose, a phone number — treat that as "**probably right, worth a look,**" not gospel. Look up the specific fact separately.

Habit 2 — For health, money, or the law: starting question, never the answer

Use AI to *understand what you're dealing with*. Use it to *frame the question*.

Then take that framed question to a real doctor, accountant, or lawyer.

THE RULE YOU CANNOT BREAK

For health, money, or the law, this is the one place where the rule is not optional. A starting question, never the answer.

The trap for smart people

This is the move most people don't know about — and it catches the people who feel *most* confident about AI.

You get an answer from the tool. You second-guess it. You type back "**are you sure?**" — and the tool may flip a right answer into a wrong one, just to agree with you.

Every AI tool out there does this, more or less. The companies that build them have tried to reduce it. They haven't fixed it.

The better move: ask "**what is your reasoning?**" or "**show me the source.**"

Now you're asking for evidence, not for reassurance. The tool has to explain itself instead of caving to you.

Write those two phrases somewhere you'll remember them.

TRY IT NOW

Go back to your AI tab. Ask it something you already know the answer to — a fact about your own town, your own field, your own hobby. Then ask it to show you its reasoning.

COPY AND PASTE

What is your reasoning for that answer? Show me the source.

WHAT YOU SHOULD HAVE SEEN

Either a real explanation you can check, or a vaguer one than the original answer sounded. Both are useful. The second is the tell — when a tool cannot say why, the confidence was doing the work.

TRY TO REMEMBER · THREE QUESTIONS, UNGRADED

1. What is safe to trust it for?

Our answer: The shape of things — a note, a summary, a plan, an outline. Not the hard facts inside them.

2. Why is "are you sure?" the wrong follow-up?

Our answer: Because it asks for reassurance, and the tool will often supply it by changing its answer. Asking for reasoning or a source asks for evidence instead.

3. Where is the rule not optional?

Our answer: Health, money, and the law. A starting question, never the answer.

In the live class we run the "are you sure?" trap on a real question so you can watch an answer flip. lu.ma/pacforward

Privacy, signal, and what's next

One rule, one setting to change. Both together, and you're safer than most people who use AI daily.

The postcard rule

Do not type anything into an AI tool that you would not write on a postcard.

That includes other people's details — a family member's health, a coworker's mistake, a client's real information. Changing the names is not enough.

Why "postcard"? In 2026, the way these tools handle your data changed a lot. That postcard now means one that:

- Gets kept for years
- May be used to train the next version of the tool
- Could be handed over if a court asks

The setting to change (do this today)

Most AI tools have something under **Settings → Data Controls** that says "improve the model" or "allow training on your chats."

Turn it OFF.

Takes about two minutes. Here's roughly where to find it for the big four:

Tool	Where to look
Claude (Anthropic)	Click your name (bottom-left) → Settings → Privacy
ChatGPT (OpenAI)	Click your name (bottom-left) → Settings → Data controls → Improve the model for everyone
Copilot (Microsoft)	Settings → Privacy → Model training on text
Gemini (Google)	Settings → Your activity → Gemini Apps Activity

Menus change. If yours looks different, search the settings for "training" or "improve."

Do it today. Not "someday."

Signal vs. hype

You are going to see a lot of noise about AI in the coming months. Most of it doesn't matter to you. Here is the filter:

WORTH YOUR ATTENTION

- AI quietly being built into the tools you already use — email, search, your work docs
- Data settings that changed this month, especially default settings on the tool you use
- Any specific use that saves you 20 minutes next Tuesday

SAFELY IGNORE

- Predictions about when AI replaces all jobs
- Every new model announcement
- "Prompt hacks" and "secrets"
- Viral demos of one impressive thing

These are content, not signal. They don't affect what you do this week.

Does this affect what I do this week?

If yes, learn just enough. If no, move on.

Your checklist — what you can do now

Tick each one you have actually done.

- ☐ I have typed into an AI tool at least once.
- ☐ I know the one honest sentence about how it works.
- ☐ I know to trust it for shape and check it for facts.
- ☐ I know to ask for reasoning or a source instead of "are you sure?"
- ☐ I have turned off training on my chats.

What to do this week

Open an AI tool. Type one real thing.

A note you've been meaning to write. A paragraph you want re-read for clarity. A question you're curious about. A recipe idea from what's in the fridge. Anything.

Just one. That's what makes it stop being mysterious for good.

WHAT YOU'LL TRY THIS WEEK — WRITE IT DOWN SO YOU ACTUALLY DO IT

Where to go next

The Field Guide to AI

A free downloadable PDF that covers everything in this course plus paperwork, inbox, family conversations, and the honest limits of AI.
pacforward.org/repository

The AI for Real People course series

Self-paced online versions of the same material. Free.

Getting Started · Paperwork · Inbox · Job Search · Difficult Conversations

pacforward.org/courses

The free live class

Thirty minutes on Zoom, run monthly. Same material, with someone to ask. ask.luma/pacforward

The Pacific Voyagers Fellowship

A sponsored 45-day program. Pacific Forward Initiative covers a premium AI subscription for you and pairs you with a mentor for the duration. Nominations open on a rolling basis. If this course resonated and you want the deeper version, that's the path. pacforward.org/pacific-voyagers

A QUESTION, IF YOU WANT TO ANSWER IT · OPTIONAL, NOT SAVED

What was the moment in this course where something clicked, or you thought "wait, I could try this on..."?

If you want to tell us, email hello@pacforward.org. A real person writes back. If something here didn't make sense, or you tried one of the tools and got stuck, that's what the email is for. There's no wrong question.

When to stop and ask a real person

This course cannot help with:

- **A medical question that matters.** Symptoms, medications, doses, whether to see someone. Use AI to understand a word your doctor used. Then talk to your doctor.
- **Money you cannot afford to lose.** Taxes, retirement, a loan, a contract you are about to sign. A starting question, never the answer.
- **Anything legal or immigration-related.** The stakes are too high and the details are too local. Use it to learn the vocabulary. Then find a lawyer or an accredited representative.
- **Anything you would not write on a postcard.** Other people's health, a coworker's mistake, a client's real information, your own identifying numbers. Changing the names is not enough.
- **A fact you are about to act on.** A name, a date, a price, a phone number, a deadline. Look it up separately before you rely on it.

You keep the last word. That is not a limitation of the tool — it is the whole arrangement.

Quick-Reference Card

The whole course on one page — for when you sit down to actually try a tool.

ONE SENTENCE TO REMEMBER

It learned what good answers sound like. It's excellent at how a good answer *sounds* — so you keep the last word.

THE DEMO PROMPT TO TRY

Write a short, warm thank-you note to a colleague who covered a meeting for me last week. Friendly, professional, not overly formal.

TWO HABITS

1. Trust it for shape. Check it for facts. (Especially names, dates, prices, doses, phone numbers.)
2. Health, money, or law: starting question, never the answer.

WHEN IT FLIPS ITS ANSWER TO AGREE WITH YOU (DON'T FALL FOR THIS)

Ask "what is your reasoning?" · Ask "show me the source."

THE POSTCARD RULE

If you wouldn't write it on a postcard, don't type it into an AI tool.

THE SETTING TO CHANGE (TODAY)

Settings → look for "improve the model" or "training on your chats" → **OFF**.

THE SIGNAL VS. HYPE FILTER

Does this affect what I do this week? If yes, learn just enough. If no, move on.

THIS WEEK

Open an AI tool. Type one real thing.

WHEN YOU HAVE QUESTIONS

hello@pacforward.org — a real person writes back.

The prompts, in one place

Six prompts you can copy when you need them. Replace the bracketed parts with your own situation.

1 • The first one — from Lesson 2

Write a short, warm thank-you note to a colleague who covered a meeting for me last week. Friendly, professional, not overly formal.

2 • Ask for evidence, not reassurance — from Lesson 4

What is your reasoning for that answer? Show me the source.

3 • Make it sound like you

Here is something I wrote: "[paste your draft]". Keep my meaning and my voice. Just make it clearer. Do not add anything I did not say.

4 • Understand a letter you were sent

Here is a letter I received: "[paste]". Explain in plain English what it is asking me to do, by when, and what happens if I do nothing. List anything I should verify myself.

5 • Frame a question for a professional

I have an appointment with [a doctor / an accountant / a lawyer] about [situation]. Help me write down the three questions I should ask. Do not answer them for me.

6 • Check yourself before you send

List every hard fact in this text — names, dates, numbers, prices, deadlines — so I can verify them one by one. Do not verify them yourself.

Keep this page near your desk. Nothing on it expires.

‘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. Our courses are free for everyone — if this one gave you time back, you can help keep it that way at pacforward.org/donate.

PACIFIC
FORWARD
INITIATIVE

FORWARD TOGETHER