




TIER 1 - FOUNDATIONS * V1.0 -- MAY 2026

MAKE AI KNOW YOU

The single highest-ROI move in the whole Tier 1 catalog. How to set up memory, custom instructions, and a personal-context doc so the model stops being a stranger every time you open it.



BY

Alex Jahn / Agent Logic

v1.0 -- May 2026

Anyone who's used AI for months and still feels like they're starting over from scratch every conversation

15-20 minutes

Free. Forever.

EDITION

AUDIENCE

READ TIME

COST

Prepared by Agent Logic / alexanderjahn79@icloud.com / theaiguwi.com

CONTENTS

What's in here

- 1 Why your AI keeps treating you like a stranger 3**
Open ChatGPT. Open Claude. Open Gemini. Ask any of them: "What do I do for a living?"
- 2 Memory vs custom instructions vs personal context -- three different things 4**
Three terms get muddled. Get them straight up front.
- 3 What memory actually is -- and what it isn't 4**
Memory is the model writing down facts you've mentioned and silently using them in later conversations. It works well for:
- 4 Custom instructions -- the part you write 6**
Custom instructions are the highest-leverage 15 minutes you'll spend on AI all year. The structure that consistently produces sharp output:
- 5 The personal-context doc -- your portable version 7**
Memory is per-tool. Custom instructions are per-tool. If you only ever use one AI, you can stop reading here. But most people end up...
- 6 The privacy red line -- what NEVER goes in 8**
This is the part most articles skip. AI tools have had memory leaks. They've had bugs that surfaced one user's conversations to another...
- 7 Updating, resetting, and the yearly review 9**
Three things you should do periodically:
- 8 When NOT to use memory + context 10**
The framework breaks down in three situations. Honest about them:
- 9 Where to go from here 11**
You've turned the model from a stranger into something that knows you. One module left to go in this set:

SECTION 1

Why your AI keeps treating you like a stranger

The setup most people never turn on

Open ChatGPT. Open Claude. Open Gemini. Ask any of them: "What do I do for a living?"

If you've used these tools for six months and the answer is some version of "I don't know -- could you tell me?" -- you've been leaving the single biggest quality multiplier turned off the whole time.

Every modern chat AI ships with two settings that tell the model who you are, what you do, and how you want it to talk to you. Most people never touch them. So every conversation starts from zero. Every time. The model has no idea you're a working carpenter, that your shop is in Fond du Lac, that your kids are 7 and 4, that you're allergic to corporate marketing language. So it gives you the same generic answer it would give a college sophomore in Connecticut. Then you do the work of customizing the answer in your head, every single time.

That's not how it's supposed to work. The fix takes 15 minutes. You do it once. Every conversation after that is sharper.

What you'll have by page 13

By the end of this primer:

- You'll know what
- You'll have a working
- You'll have a
memory yet, and that you'll refine over the next year.
- You'll know what to keep OUT of memory (the privacy red line).
- You'll know how to update, reset, and clean your context when life changes.

Plain English, fifteen-minute setup, immediate quality jump. The compounding starts the day you turn it on.

memory i
custom in
personal-

Most people use AI like meeting a stranger every five minutes. The fix isn't a better model. It's a better introduction -- once.

SECTION 2

Memory vs custom instructions vs personal context

Three terms get muddled. Get them straight up front.

Memory -- facts the model writes down about you between conversations. ChatGPT calls it Memory. Claude calls it (depending on which interface you're in) Projects or Memory. Gemini calls it Saved Info. The model decides what's worth saving -- a thing you mentioned, a preference you stated, a fact that came up. Some platforms let you edit the list directly. Most update silently in the background.

Custom instructions -- a block of text

ChatGPT has Custom Instructions. Claude has System Prompts at the project level (or a similar setting per-account on some tiers). Gemini has Saved Info that doubles as instructions. This is YOUR voice telling the model how to behave: who you are, what you're working on, how you want answers shaped.

you write,

Personal-context doc -- a one-pager YOU control, stored somewhere you trust (Notes app, Apple Notes, iCloud, a text file on your Mac). It's the canonical version of your context. When a tool doesn't have memory, you paste this in at the start of a conversation. When you switch tools, you carry it with you. When you update it, you update one file -- not five settings across four tools.

The three layered together is the goal. Memory does the easy automatic stuff. Custom instructions handle the structural "how I want to be talked to" stuff. The personal-context doc is your portable version that you own and that survives every tool switch and every time the platform changes its memory features.

SECTION 3

What memory actually is -- and what it isn't

What it does well

Memory is the model writing down facts you've mentioned and silently using them in later conversations. It works well for:

- **Persistent preferences.** "I prefer short answers without bullet lists." "Use metric units." "Never start with 'Great question!'" -- once stated and saved, it sticks.
- **Project context.** "I'm building a small carpentry business in Fond du Lac." "I'm studying for the Wisconsin contractor license exam." Surface-level facts about your situation that the model can reference.
- **Past conversation threads.** "Last time we talked about that proposal for the kitchen remodel." Most platforms now thread some recall across sessions.

What it doesn't do

Memory is not a database. Don't trust it for anything load-bearing. Specifically:

- **Memory hallucinates.** The model will sometimes "remember" things you never said, or get details wrong on things you did say. Verify if it matters.
- **Memory leaks.** Information you put in one conversation can show up in unrelated ones. If you mentioned a client's name in passing, the model may surface it weeks later in a totally different thread. This is a privacy issue, not a feature.
- **Memory resets.** Platforms change their memory features without warning. Your saved facts can get wiped on a tier change, on a privacy-setting reset, or because the company decided to redesign the system. Don't put anything in memory that you'd be upset to lose.
- **Memory is per-platform.** ChatGPT memory does not transfer to Claude. Claude doesn't transfer to Gemini. Each tool's memory is its own island.

The memory test: if a platform deleted everything it has memorized about you tonight, would you mind?

- **No** -> you're using memory right (low-stakes preferences only).
- **Yes** -> you've put something load-bearing in there. Move it to your personal-context doc instead.

SECTION 4

Custom instructions -- the part you write

The structure that works

Custom instructions are the highest-leverage 15 minutes you'll spend on AI all year. The structure that consistently produces sharp output:

The four-part custom instructions block:

- 1. Who I am** -- one sentence. Role, location, situation. *"I'm a working contractor in Fond du Lac, WI. I run a one-person shop with occasional subs."*
- 2. What I'm trying to do** -- one or two sentences. Current focus or recurring tasks. *"I write proposals and invoices, manage scheduling, and build small business systems. I read and write a lot of practical content for trade-business owners."*
- 3. How I want answers** -- three or four bullets. Concrete and concrete. *"Plain English, no marketing fluff. Short answers unless I ask for depth. Skip the 'as a large language model' disclaimers. If I'm wrong about something, tell me directly -- don't soften it."*
- 4. What to skip** -- one or two bullets. Specific irritants. *"No 'Greatest hits' summaries at the end of every reply. Don't add 'I hope this helps!' at the close."*

That's it. Fifteen minutes to write, paste into the settings page, save. From the next conversation onward, every reply starts from this baseline.

Worked example -- what good looks like

Here's a complete custom instructions block, written for a working carpenter:

I'm a working carpenter and small contractor in Fond du Lac, WI, mid-30s, married with two young kids. I run a one-person shop with occasional subs. I'm currently building out the AI side of my business in addition to the trade work, so I move between practical jobsite questions and software/business questions in the same week.

No 'leverage' or 'synergize' or 'unlock value.' Short answers unless I ask for depth. If I ask a yes/no question, lead with yes or no. If I'm wrong about something, tell me directly -- don't soften it. I'd rather be embarrassed for two seconds than walk around wrong all week.

hope this helps!' closers, bullet-pointed summaries at the end of every reply, the 'as an AI language

Make AI Know You
model' disclaimer. If I want a list, I'll ask for one.

That's roughly 130 words. It takes you ten minutes to write, ten seconds to paste, and it changes every reply you ever get from that account afterward.

What changes when you turn it on

The first conversation after you save custom instructions will feel slightly off -- the model starts dropping the openers, skipping the disclaimers, lining up to your stated preferences. Then by the third conversation it just feels like the model is "getting it." That feeling is the win. You haven't gotten smarter; the model hasn't gotten smarter. The match between what you want and what shows up just got better.

If you want the same starter structure pre-written for your specific role, the Tier 2 *Role-Specific Prompt Templates* primer covers nine common roles in detail, custom-instructions block included. Same training catalog, free.

Role-Specific

SECTION 5

The personal-context doc -- your portable version

Why you own this, not the platform

Memory is per-tool. Custom instructions are per-tool. If you only ever use one AI, you can stop reading here. But most people end up running two or three for different jobs -- Claude for long writing, ChatGPT for image work, Gemini for Google integration. You don't want to write your context block four times. You don't want to update it four times when your situation changes.

The fix: keep one canonical version in a place you control. I keep mine in Apple Notes. Some people use a plain text file. Some use a Google Doc. The format doesn't matter; the ownership does. When a platform memory wipes itself, your doc survives. When you switch tools, you paste the doc in. When something about your life changes -- new job, new kid, new business focus -- you update one place.

What goes in the doc

Mine has six sections. You can copy this layout or build your own; the structure matters less than having one.

Personal-context doc -- the six sections:

1. **Who I am** -- name, age, location, family setup at a high level, role, business name if relevant.
2. **What I'm working on right now** -- current projects, current focus, what's eating my time this month.
3. **What I value in answers** -- voice/style preferences, level of formality, what kinds of disclaimers and hedges to skip.
4. **My recurring tasks** -- proposals, scheduling, learning a new tool, etc. The model uses this to anticipate what you might need.
5. **Background you should know** -- past experience, expertise areas, things I'm specifically NOT an expert in (so the model knows when to teach vs when to assume).
6. **What's off the table** -- privacy red lines, topics not to make assumptions about, things never to volunteer.

The whole doc fits on one page. Mine is about 400 words. It's not a memoir; it's an introduction.

How to use it

For platforms with custom instructions: paste sections 1-4 into the custom instructions field. Keep the full doc handy.

For platforms without custom instructions, or for one-off conversations where you want maximum context: paste the whole doc as your first message.

question." The model reads it, internalizes it, and the rest of the conversation runs on top of that base.

"Quick co

For high-stakes work -- a long writing project, a sensitive personal decision, an important proposal: paste it into a project / starts-of-conversation slot deliberately. Don't trust memory to carry it.

SECTION 6

The privacy red line -- what NEVER goes in

This is the part most articles skip. AI tools have had memory leaks. They've had bugs that surfaced one user's conversations to another user. They've changed their data-retention policies in ways that surprise people. Treat memory and custom instructions as semi-public. Here's what that means:

Never put in memory or custom instructions:

- **Social Security numbers, tax IDs, full date of birth.**
- **Bank account numbers, routing numbers, full credit card numbers.**
- **Passwords, even if you're "just storing them temporarily." Especially not those.**
- **Specific dollar amounts of personal wealth, debt, or income** beyond rough framing ("small business owner, low six figures gross"). The detail isn't load-bearing for the model and you don't want it surfacing in a conversation later.
- **Other people's private information** -- your spouse's medical issues, your kids' specific schools, your client's address. They didn't consent. Don't put their data in your context.
- **Active health diagnoses or therapy notes.** The model will "remember" and reference these in unrelated conversations. Most people don't want that.
- **Anything you'd be upset to see screenshotted.** Treat memory as if it might end up screenshotted somewhere by accident.

If you need the model to know about a sensitive topic for a single conversation, mention it that one time and let it not be saved (most platforms have a "don't save this" toggle, or a private/incognito mode). Don't make it permanent.

The pattern that works: high-level role and preference info goes in custom instructions and memory. Anything specific, sensitive, or load-bearing stays in your personal doc, on your own device, and you paste it in only when you need it for that conversation.

SECTION 7

Updating, resetting, and the yearly review

Memory drifts. Yours will too.

Three things you should do periodically:

Once a quarter -- review what's in memory. Most platforms have a settings page that shows the saved facts. Open it. Read what's there. Delete anything that's wrong, outdated, or that shouldn't have been saved. (You'll find at least one weird thing -- something the model latched onto that you never meant to save.)

Once a year -- refresh the personal-context doc. Your situation changes. Maybe last year's "studying for the contractor license" is this year's "running a four-person crew." Maybe the recurring

task you wrote down at the top has shifted. Reread the doc. Update what's stale. Save the new version.

On any major life change -- full reset, new context. New job, new business, new home, new kid, divorce, big shift in focus. The old context will pull the model toward old assumptions. Wipe the memory, paste the new doc, start fresh. Five minutes. Worth it.

When the model gets a fact wrong

You'll hit conversations where the model "remembers" something incorrectly. Maybe it thinks you live somewhere you don't, or has a stale role. Two responses:

- **For one-off corrections in the conversation:** "Quick correction -- I'm in Fond du Lac, not Madison. Continuing." The model adjusts in the conversation and most platforms write the correction to memory.
- **For persistent errors:** open the memory settings, find the wrong entry, delete it directly. Then re-prime in a new conversation.

Don't argue with the model about what it remembers. The fastest path is always: edit the source (memory page, custom instructions, personal doc), then start fresh.

15

Fifteen minutes once.

That's the setup time. Ten minutes to write a personal-context doc, five minutes to paste it into custom instructions on each tool you use. Every conversation after that runs on top of context the model now actually has. The compounding starts immediately.

SECTION 8

When NOT to use memory + context

The framework breaks down in three situations. Honest about them:

- **Sensitive single-conversation work.** Therapy-style journaling, exploring a hard personal question, working through a conflict with a specific person. Don't save the whole thread to memory. Use a private/incognito mode if your platform has one, or explicitly tell the model "don't save anything from this conversation." Some things should die when the conversation ends.
- **Work-account vs personal-account separation.** If you have an employer-issued AI account, don't pile your full personal context into it. Your employer's IT may have access to that data. Keep work tools work-shaped, personal tools personal-shaped, and don't cross-contaminate.
- **High-stakes professional advice.** Legal questions, medical questions, tax questions, contractor liability questions. The model will use your stored context to give you answers -- that's the whole point -- but stored context doesn't make the answers more accurate. Your context tells the model who's asking, not what's true. Verify with a real professional for anything that has real consequences.

Within those limits, the memory + custom instructions + personal doc combination is the highest-ROI move in the catalog. Outside them, the boundaries from Module 4 (*When NOT to Use AI*) apply.

SECTION 9

Where to go from here

You've turned the model from a stranger into something that knows you. One module left to go in this set:

- **Reading walls of text** -- owner's manuals, contracts, terms of service, court documents, employee handbooks. Drag-and-drop summarization done right, and how to catch what the model misses.

After that, Tier 1 goes into the second six modules: voice mode, free-tools comparison, big decisions, AI alongside your kids, the rest of the Tier 1 expansion.

Get the next module the day it drops: theaiguywi.com/training

One email per release. No drip. No spam. Opt out anytime.

If you want this same setup pattern installed across an entire team -- custom instructions written for each role, the personal-context doc convention adopted shop-wide, the privacy red line trained in once so it sticks -- that's the consulting offer. We do it the way I run it in my own carpentry business.

Reach out: alexanderjahn79@icloud.com

A short call. Honest scope. We figure out together if it's a fit.

Closing -- the lock-in line

The hardest part of this whole thing is going from "I should set this up sometime" to opening the settings page tonight and writing the four-part block. Once it's done, it's done. Every conversation after stands on top of it.

The compounding is invisible -- you don't see "the conversation I would have had without context." You just feel the model getting sharper, week by week. That's what's happening. You introduced yourself once. The model now knows who's asking.

You have the setup. One module left in this batch.

Agent Logic --

Fond du Lac, WI. This is module 7 of 12 in Tier 1 (Personal).

© 2026 Agent Logic. Share freely.

theaiguynw