




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

DAY 1 OF LEARNING AI

Skip the studying. Here's what to do in your first hour -- written by someone who uses this stuff every day to run a real business.



BY

Alex Jahn / Agent Logic

v1.0 -- May 2026

Anyone starting from zero today who's tired of reading about AI and ready to actually use it

15-20 minutes

Free. Forever.

EDITION

AUDIENCE

READ TIME

COST

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

CONTENTS

What's in here

- 1 Why Day 1 matters more than Day 100** **3**
If you've already typed "how do I learn AI" into Google or YouTube, you've seen the trap. Hundreds of hours of content. "Mastering..."
- 2 Drop the assumption that you need to learn AI before using it** **4**
Most things adults learn at work, they learn on the job. You learned the company's software by using it. You learned how to run meetings...
- 3 Pick your harness -- what I actually use** **5**
The model is the brain. The harness is the thing you type into. Same brain can sit behind three different harnesses and feel like three...
- 4 Your first 30 minutes -- three concrete prompts to try right now** **6**
These are not impressive demos. They're real tasks you probably have on your plate today. Copy them, fill in the bracketed parts, paste...
- 5 How to read what comes back** **7**
The model produces something. Now what?
- 6 The one thing NOT to do on Day 1** **8**
This is the only landmine on Day 1 that's worth covering up front: assume that anything you type into a free-tier AI chat can be seen by...
- 7 Day 2 and beyond** **9**
You did the prompts. You got something useful. You felt the rhythm of "describe the task -> read the draft -> push back -> ship something..."

SECTION 1

Why Day 1 matters more than Day 100

Most people spend Day 1 reading about AI. That's the trap.

If you've already typed "how do I learn AI" into Google or YouTube, you've seen the trap. Hundreds of hours of content. "Mastering ChatGPT in 30 days." "The complete prompt engineering bootcamp." Nine-part series on transformer architecture. Reading lists. Discord communities. Newsletters with 200 issues in the archive.

Most of it is good. Most of it is also a way to using the thing on real work that matters to you.

put off the

The people I know who got genuinely good at AI in the last two years -- fast -- did not start with study. They started with use. They picked a task they were already doing badly or slowly, opened ChatGPT, and tried. They got a mediocre answer. They tried again. Within two weeks they were doing the task in a third of the time. They went back and read theory

only when

The people I know who are still "going to learn AI someday" -- also fast learners, also smart -- have spent six months reading about it and have not used it on a single piece of real work. They know more about LLMs than the first group does. They produce less with them.

This is the trap. Day 1 is the day you decide which group you join.

What Day 1 actually is

Day 1 is the day you commit to using AI on a real piece of work that's already on your plate. Not a sandbox task. Not a tutorial exercise. A real email you owe someone. A real plan you've been avoiding. A real document you need to draft.

Day 100 is when you've either become someone who reaches for AI as a default -- like a calculator, like a search bar -- or you've become someone who still between those two outcomes is set on Day 1.

thinks abo

This primer is the shortest path from where you are now to actually-using-it by tonight. Twelve pages. No theory until you need it. The next module in this curriculum (*What An LLM Actually Is*) will give you the frame to extend from. But you don't need the frame yet. You need to use it.

The people who got good at AI didn't study it. They used it on real work and read theory only when they hit a wall.

SECTION 2

Drop the assumption that you need to learn AI before you can use it

You don't study a hammer before driving a nail.

Most things adults learn at work, they learn on the job. You learned the company's software by using it. You learned how to run meetings by running meetings. You learned how to manage people by managing people, badly at first.

AI is the same. The tool meets you where you are. You can type a normal English sentence into ChatGPT and get something useful back without knowing the word "transformer." You can use it productively for a year before you need to know what a token is. The model doesn't care that you don't know the theory -- and weirdly, neither should you.

This is unusual for technology. Most software has a learning curve up front. Excel makes you learn formulas. Photoshop makes you learn layers. AI assistants are the first major technology I can think of where the learning curve is *down* front. You can use the interface to describe what you want.

That's it. The rest is practice.

The thing you don't need on Day 1

You don't need a course. You don't need a YouTube binge. You don't need to know the difference between GPT-4 and GPT-4o. You don't need the prompt-engineering tricks people sell on Twitter. You don't need to read about hallucinations. You don't need a vocabulary upgrade.

If you're tempted to read more before starting -- that's exactly the trap from Section 1. Close this PDF, do the prompts in Section 4, come back tomorrow.

The thing you do need on Day 1

A real task you're already doing today that you can throw at the model. The task should be:

- **Real.** Something you actually have to do this week.
- **Low stakes.** Day 1 is not the day to draft your divorce papers. Day 1 is the day to draft a routine email or plan a normal week.
- **In writing.** AI is text in, text out. If your task is plumbing a sink, that's not a Day 1 task. If your task is

writing a c

Good Day 1 tasks: drafting an email you've been putting off, planning your week, summarizing a long article you didn't have time to read, brainstorming names for a project, writing a thank-you note, outlining a meeting agenda, generating a packing list, comparing two options you've been mulling. Pick one. Use it tonight.

SECTION 3

Pick your harness -- what I actually use

Three real options.

The model is the brain. The harness is the thing you type into. Same brain can sit behind three different harnesses and feel like three different products. Pick a harness for Day 1 and stick with it for the first week.

ChatGPT in a browser -- chat.openai.com. The most familiar. Most people you know are using this. It's fine. The free tier handles 90% of what beginners need. If you've never used any of these before, this is the lowest-friction starting point.

Claude in a browser -- claude.com. A different model under the hood, made by a different company. In my experience it's noticeably better at thinking through problems, longer-form writing, and not just agreeing with whatever you said last. Free tier works fine.

Claude Code -- claude.com/code. This is the one I actually use to run Agent Logic. It's a terminal-based interface that runs Claude in an

agenti

Which one to pick

If you've never touched any of this before and you want the path of least resistance: **browser, today**. Open it, type, learn the rhythm. There's no wrong answer here.

If you want to see what AI actually looks like when it's -- and you're willing to spend 20 minutes getting set up: the version of this tool that closes the loop between "model produces text" and "thing actually got built." For me it's the difference between an advisor and an apprentice.

I'm telling you what I use, not telling you what to use. ChatGPT is fine. Claude in a browser is fine. Claude Code is what I reach for when the work is real, because it can pull files into context, write code, run scripts, and check its own output -- without me copy-pasting between windows. If that sounds useful, look at it. If it sounds like overhead, don't.

The honest answer: pick one, use it for a week before switching. Switching harnesses every other day is how you stay a beginner.

ChatGPT in a browser is an advisor. Claude Code is an apprentice. They're not the same job.

SECTION 4

Your first 30 minutes -- three concrete prompts

These are not impressive demos. They're real tasks you probably have on your plate today. Copy them, fill in the bracketed parts, paste into your harness of choice. Do all three back-to-back. Twenty to thirty minutes total.

Prompt 1 -- Draft an email you've been avoiding

I need to write an email to [WHO -- name and your relationship to them, e.g. "my landlord, professional but firm"]. The situation is [SITUATION IN 2-3 SENTENCES -- what happened, what you want, why it matters]. Write me three versions: one short and direct, one warmer with more context, one firmer in case the first two don't land. After each, tell me what trade-off it makes.

What this teaches: the model will produce three meaningfully different drafts. Pick the one closest to what you want. Edit. Send. You just compressed a 30-minute task into 8.

ChatGPT

doing wor

Claude C

Prompt 2 -- Plan something you've been putting off

I want to [GOAL -- concrete, e.g. "reorganize my garage" or "launch a side project" or "renegotiate my insurance"]. Here's what I already know: [WHAT YOU KNOW -- 3-5 bullets]. Here's what I don't know: [WHAT YOU DON'T KNOW -- what feels blurry]. What's the smartest sequence to actually get this done in [TIMEFRAME]? Be specific. Push back on the parts of my framing you think are wrong.

What this teaches: the model is a thinking partner who's never tired and won't judge you for not knowing. The "push back" line is critical -- without it, you get a yes-man. With it, you get useful friction.

Prompt 3 -- Find the weakest part of something you wrote

Here's a [EMAIL / PROPOSAL / DOCUMENT] I'm about to send. [PASTE IT.] What's the weakest sentence? What part will the reader skim or misread? If you could change one thing to make this land better, what would you change and why?

What this teaches: the model is a useful editor when you tell it specifically what to look for. Vague prompts ("make this better") get vague results. Specific prompts ("find the weakest sentence") get specific results.

Want the next five Tier 1 modules? theaiguywi.com/training

All free. No email required. New module each week.

SECTION 5

How to read what comes back

The model produces something. Now what?

Three checks. Run them on every output, every time, until they're automatic.

Check 1 -- Is it factually right?

The model is confident. It's confident whether it's right or wrong. There's no "I'm not sure" tone -- everything sounds equally certain. So

you are th

For Day 1 stakes -- drafting an email, planning your week, writing a thank-you note -- the bar is low. There's not much to fact-check in your own email about your own situation. Fine.

For higher-stakes work -- anything you'll send to a customer, anything that quotes a number, anything that names a specific law or product or person -- verify the load-bearing facts before you act. The model can hallucinate a court case, invent a regulation, or make up a statistic that *true*. It is not lying -- it is generating plausible text. The distinction doesn't matter when you've already sent the email.

sounds

Check 2 -- Is it the right shape for your situation?

The model doesn't know your context. It doesn't know your customers, your tone, your boss, your industry's quirks. It produces a generic-good answer based on what generic-good answers look like in its training data. That's often 70% of what you need. The other 30% is your judgment.

If the email reads like a stranger wrote it about your job -- that's the model not knowing your context. Add context to the prompt. Try again.

Check 3 -- Would I send this if it had my name on it?

This is the gut check. The model's name isn't on the work. Yours is. If you'd be embarrassed by any sentence in the output, edit it before you ship it. If you'd be proud of it -- ship it. The fact that AI helped is irrelevant. People don't care how you wrote the email; they care what it said.

Treat the output like a draft from a smart intern who doesn't know your business yet. Use it. Edit it. Ship work

Redesign

goes deeper into this -- but the three checks above are 90% of what you need on Day 1.

SECTION 6

The one thing NOT to do on Day 1

Don't paste anything you wouldn't email to a stranger.

This is the only landmine on Day 1 that's worth covering up front:

assume t

into a free-tier AI chat can be seen by the people who run that company, used to train future versions of the model, and conceivably surfaced in some form later.

That means:

- Don't paste customer information you wouldn't share publicly.

- Don't paste payroll, employee files, or anything HR-sensitive.
- Don't paste contracts under NDA.
- Don't paste financial statements, account numbers, or credentials.
- Don't paste anything proprietary your company has explicitly told you not to share.

If you wouldn't write it on a public Reddit post, don't paste it into a free AI chat. The threat isn't the model -- it's that the conversation goes through somebody else's servers and somebody else's policies.

This isn't paranoia. Free-tier AI products are explicit in their terms that conversations may be retained and used. Paid plans usually have stronger protections; enterprise plans usually have actual privacy guarantees in writing. On Day 1, you're almost certainly on a free tier. So play safely: redact, generalize, or use a fictional version of the situation if the real version has private info in it.

There's a full Tier 2 module on this called
covers the 95th percentile.

Privacy and

SECTION 7

Day 2 and beyond

The path forward, in order.

You did the prompts. You got something useful. You felt the rhythm of "describe the task -> read the draft -> push back -> ship something better than you would have alone." That's the loop. Everything else is practice and tactics.

Here's the order I'd read these in if I were you:

Tomorrow (Day 2):

it for a day, the mental model lands much harder than reading it cold. This is the foundation everything else sits on.

What An L

Day 3-4:

getting good output the first time, instead of through three rounds of prompting.

The 3-Qu

Day 5-6:

you, and the verification moves that catch it before you ship.

Reading A

Week 2 onward: the rest of Tier 1 (six modules total), then Tier 2 (using AI well at your job), then Tier 3 (being the AI-fluent person on a team). Free curriculum. No drip sequence, no upsells. Read in order or read what's relevant.

The discipline that matters most: going to learn it someday" group. Use it daily for two weeks and you're permanently in the other group. There's no middle ground.

use it dai

Want this rolled out at your business? Text 920-539-8814 or email alexanderjahn79@icloud.com

Real cell, real human. Short call. Honest scope. We figure out together if it's a fit.

Closing -- the lock-in line

One sentence to keep with you. It's the difference between the people who get good at this and the people who don't:

1

Day 1 isn't about learning AI. It's about using it.

Pick one task today. Use it on that task tonight. Read theory tomorrow. The learning happens in the using.

The harness is your choice. The task is your choice. The only thing that's not optional is starting.

Agent Logic --

Fond du Lac, WI. This is the entry-point module for Tier 1 (Personal). After this, read *Actually Is*.

theaiguyn

What An L

© 2026 Agent Logic. Share freely.