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

WHEN NOT TO USE AI

Four red flags every AI user needs to know -- privacy, accuracy, speed, judgment -- and the 30-second decision framework that catches them before you fire up the chat window.

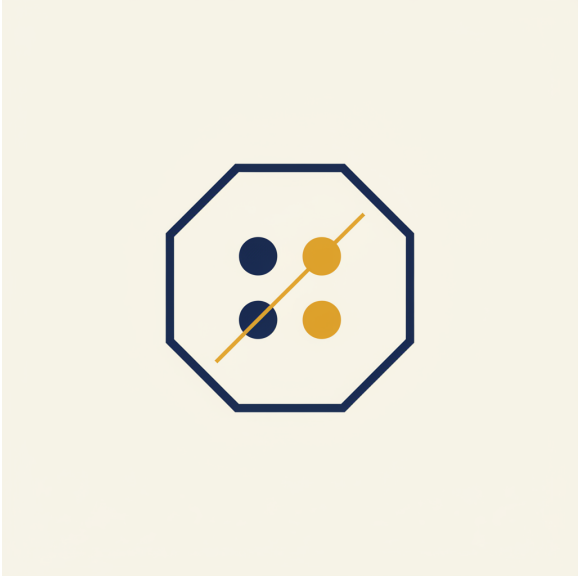

BY Alex Jahn / Agent Logic

EDITION v1.0 -- May 2026

AUDIENCE Anyone who's been told "use AI for everything" and is starting to suspect that's wrong

READ TIME 15-20 minutes

COST Free. Forever.

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

CONTENTS

What's in here

SECTION 1

The cost of misusing the hammer

AI is a hammer. Some of these problems aren't nails.

There's a marketing pressure right now that the answer to every question is "use AI for it." Every newsletter, every productivity guru, every LinkedIn post. *Use AI to your week. Use AI to make your decisions. Use AI for everything.*

It's wrong. Not "wrong as a slogan, right in spirit." Just wrong. There's a real and growing set of jobs where firing up a chat window is the slower, riskier, or dumber choice. People who can't tell which jobs those are end up paying real costs for the lesson.

This module is the list. Four red flags that mean *put the ke* framework you can run before you open the chat window. And the honest case for AI literacy: not just knowing how to use it, but knowing when not to.

Three real losses I've seen from misuse

I'll keep this short because the abstract version of "AI can hurt you" doesn't land. Specific does.

- A small-business owner pasted a customer's full tax return into ChatGPT to "help organize her vendors." Customer SSN, account numbers, the works. He didn't know that consumer tier prompts can be used for model training and were, at minimum, logged. He's now in a months-long disclosure conversation with the customer.
- A contractor asked the model "what permits do I need for X kind of work in my county" and got a confident, plausible answer with specific code references. He acted on it. Three of the code references didn't exist. The county did, however, exist -- and so did the fine.
- A friend in a knowledge-work job started "AI-drafting" every email he sent. The drafts took longer to clean up than the originals would have taken to write from scratch. He spent three months at lower throughput before he noticed.

None of these were stupid people. They were operating on the wrong default -- *use AI for* -- and the failure modes weren't visible until the damage showed up. This module is so you don't repeat their experiments.

> ***AI is a force multiplier. So is a chainsaw. The question isn't whether it makes you stronger -- it's whether the thing you're using it on is the thing you should be using it on.***

SECTION 2

The four red flags

Four categories where AI is the wrong tool. Each one has a clear marker -- once you can spot the marker, the decision becomes automatic.

The four red flags:

1. **Privacy red line** -- There's information in the task that shouldn't enter a third party's logs.
2. **Accuracy red line** -- The cost of being wrong is bigger than the cost of doing it manually.
3. **Speed red line** -- A human reflex would be faster than the AI round-trip.
4. **Judgment red line** -- The decision needs real-world skin in the game the model doesn't have.

Each one gets its own section. Read them once and you'll spot them in your own week.

The pattern across all four: AI shines when *plausible-*
flag is a category where *plausible-*

SECTION 3

The privacy red line

What never enters a prompt window

Some information should not be pasted into a consumer AI tool, full stop. The list is shorter than people think but more important than people realize.

- **Customer or client PII.** Social security numbers, full birthdates, drivers license numbers, account numbers, full credit card numbers. Anything you'd be obligated to disclose if it leaked.
- **Regulated data.** HIPAA-protected health information. Attorney-client material. Anything under an NDA. Anything subject to GDPR or state privacy law.
- **Financial account specifics.** Bank routing numbers, online banking credentials, card numbers, tax-return-level detail tied to a real person.
- **Trade secrets and proprietary IP.** Code, algorithms, formulas, manufacturing specs, anything where leakage = competitor advantage.
- **Anyone else's information you don't have permission to share.** This is the one most people miss. Your employee's salary, your contractor's pricing, your customer's contract terms -- even if it's "for me to organize," it's not yours to paste.

Why this matters at the consumer tier

Two things people don't always know:

1. Consumer-tier prompts can be used for training. ChatGPT (consumer), Claude (consumer), Gemini (consumer) -- by default, prompts may be used to improve the model. You can opt out, but it's a setting most users haven't found. Even when training is opt-out, prompts are *logged* for review and abuse detection.

2. Even enterprise tiers have a paper trail. Enterprise contracts typically include "no-training" clauses and stronger data handling, but a prompt is still a network request. Logs exist. Subpoenas reach them. Compliance frameworks treat them as data movement.

> **The 60-second privacy gut-check before you paste:**

- Is this data tied to a real, identifiable person?
- Would this need to be disclosed if it leaked?
- Is this regulated under any law I'm subject to (HIPAA, GDPR, state privacy)?
- Is this my data to share, or someone else's?

Yes to any one of those = don't paste it into a consumer-tier chat. Either redact the load-bearing detail, use an enterprise/business tier with proper terms, or do the task yourself.

What this rules out, and what it doesn't

You can still use AI on most of your work. You just need to redact or substitute before pasting. *"Customer A wants to renegotiate her contract -- here's the situation"* is fine. *"Sarah Jo Main Street wants to renegotiate her contract -- here's her SSN for context"* is not. Same task. Different data exposure.

SECTION 4

The accuracy red line

When "good enough" isn't good enough

AI output is *plausible-*
It's also why prompts fail catastrophically in the 20% where being right matters.

The accuracy red line is simple: **if a crede**
answer, that's the cost of the failure mode. You're not saving money by AI-ing it. You're moving the failure to a worse place.

Categories that fall on the wrong side of this line:

- **Legal questions with money attached.** Contracts. Compliance. Liability. The difference between "a lawyer would charge $400 for that opinion" and "I asked Claude" is not money saved -- it's risk shifted to you.
- **Medical questions with action attached.** Drug interactions, dosages, diagnoses. The model can be 99% right and the 1% kills somebody.
- **Tax and accounting at the filing level.** Brainstorming-phase questions are fine. Filing-phase numbers need verification by someone with credentials.
- **Regulatory specifics.** "What permits do I need" / "what code applies" / "what license is required." See my contractor friend earlier in this module.
- **Anything with fake-citation potential.** Case law. Academic citations. Specific statistics. The model invents these confidently.

The verify-first rule

If you're going to use AI in any of these categories anyway -- and sometimes that's reasonable -- the rule is *verify befo*
(the statute, the manual, the licensed professional). Don't let speed-of-output substitute for correctness.

This is the entire premise of Module 3 (*Reading AI Output Critically*). If you skipped it, the short version is: the model has no internal "I'm not sure" signal, so the user has to be the fact-checker.

1

Bad answer is enough.

The model gets the next 99 right and you're still in trouble. Verify load-bearing facts, every time, no exceptions, in any category where being wrong has a real cost.

SECTION 5

The speed red line

When AI is slower than you

This one's counter-intuitive but it shows up constantly. There are tasks the AI is technically capable of, but the round-trip time -- open the chat, frame the prompt, read the answer, copy, paste, edit -- is slower than just doing the thing.

Examples from a normal week:

- "*Reply to Bob: 'looks good, ship it.'*" -> Just type it. The AI round-trip is 30 seconds. The reply is 5.
- "*Add 187 + 412 + 309.*" -> Open the calculator app. Or do it in your head. Both faster than chat.
- "*What's today's date?*" -> It's on your computer. Already.
- "*Is 2027 a leap year?*" -> Two-second mental check. The AI takes 15 seconds and might be wrong.
- "*Reschedule lunch with Tom to Thursday.*" -> Open the calendar. Twenty seconds. Faster than describing the situation to a model.

The tell: **if a comp**
probably the wrong tool. The model is for tasks that require effort to do right, not tasks that require effort only because you haven't done them yet.

Why people get this wrong

AI feels productive even when it's slowing you down. Typing into a chat window has the *shape* of work -- it looks engaged, it produces output, it's a deliverable. So people use it for tasks where doing nothing would be faster than typing, then feel proud of how productive they were.

The fix is observational, not theoretical. For one week, time how long each AI interaction actually takes -- from "thought of the task" to "task is done." Compare to how long the task would have taken without AI. About a third of the time, the answer is *I lost time* the unlock.

30

Seconds.

The rough threshold below which AI almost always slows you down. Above 30 seconds of task work, AI starts being faster. Below it, your hands are faster than the chat window.

SECTION 6

The judgment red line

Wisdom vs pattern-match

The hardest red flag because it's the easiest to get fooled by. Some decisions feel like things you can think out loud about -- and the model is great at thinking out loud. But there's a category of decision where the model can produce a beautiful, fluent, well-reasoned answer that's *the wrong* because it lacks something the model fundamentally cannot have: real-world skin in your specific game.

The category:

- **Hiring and firing.** Whether to bring someone on, whether to let someone go.

- **Big strategic pivots.** "Should I expand to a second location?" "Should I take this contract?" "Should I close this product line?"
- **Personal life decisions.** "Should I take the job?" "Should I have the conversation?" "Should I move?"
- **Anything tied to specific human relationships.** What to say to your kid about something hard. How to handle a sensitive customer complaint. Whether a friendship can survive an argument.

The model can pattern-match to similar decisions in the training data. It can list pros and cons. It can play out scenarios. None of that is the same as *judgment*, pattern-recognition with personal cost, real-world experience, and skin in the game.

Why the model's confidence here is dangerous

Module 1 of this series made the case: the model is a prediction engine, and it sounds equally confident whether it's right or wrong. In factual matters that's a manageable problem -- you can verify the fact externally. In judgment matters, there's no external verification. The decision is yours and the model's "confident-sounding answer" can carry false weight.

> ***The model has no skin in your game. Its confidence costs it nothing. Yours costs you everything.***

What AI is good for in judgment-adjacent territory

This isn't "never use AI for hard decisions." Use it as a *thinking p* Concrete uses that work:

- **Steel-manning the option you're leaning against.** "I'm leaning toward firing X. Make the strongest case for keeping them."
- **Surfacing considerations you might have missed.** "What are the three things people in my situation typically forget when deciding Y?"
- **Stress-testing your reasoning.** "Here's my logic for why I should take this contract. Push back."
- **Drafting the conversation, not making the call.** "Help me write the script for the hard talk I need to have with Bob."

The pattern: AI generates the inputs to your judgment. You make the judgment. Don't outsource the part of the decision that's actually yours.

SECTION 7

The 30-second pre-flight check

The whole module compresses to a 30-second mental routine you run *before* ope
window.

> **The 30-second pre-flight check:**

Before you fire up the model, ask:

- **Privacy:** Does this involve data that shouldn't enter a third party's logs? *(If yes -- r*
  tier, or do it yourself.)
- **Accuracy:** Would the cost of a confidently-wrong answer exceed the cost of doing this manually? *(If yes -- a*
- **Speed:** Could a competent person do this in under 30 seconds? *(If yes -- y*
  than the AI round-trip. Just do it.)
- **Judgment:** Does this decision require real-world skin in my specific game? *(If yes -- u*
  surface inputs, not to make the call.)

If any answer is "yes" -- slow down. AI may still be part of the right move, but it's not the *whole* mov
Pause and pick the right tool for the right part of the task.

If all four are "no" -- fire up the model and get the multiplier.

The honest split

For most knowledge workers and small-business owners, my rough estimate is that **70%-80%**
daily AI-eligible tasks are clean -- no red flags, full multiplier available. The remaining 20%-30% has at least one flag in play. That's not "AI is dangerous." That's "AI works on most things, and the user needs to know which things it doesn't."

The whole point of this module is to make the 20%-30% legible. Once you can name the four red flags, you stop blowing up on them by accident.

SECTION 8

Where to go from here

You've now completed four of six Tier 1 modules. The pattern of the foundations:

1. **What an LLM actually is** -- the mental model.
2. **The 3-question prompt framework** -- the operating procedure.
3. **Reading AI output critically** -- the editorial discipline.
4. **When NOT to use AI** *(this one)*

Two modules left in Tier 1 -- *[illegible]*
these foundations once they're internalized. After that, you're into Tier 2 (Professional) and applying the whole stack to your actual job.

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 boundary discipline trained into your team -- the four red flags, the pre-flight check, the patterns that prevent your business from being the next cautionary tale -- that's the consulting offer. I install it the same 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

Internalize this and you'll catch most of the failure modes that bite small-business AI users:

4

Four red flags: privacy, accuracy, speed, judgment.

Each one means put the keyboard down. Run the 30-second pre-flight check before you fire up the chat window. Plausible-on-average is the model's whole game; the four red flags are where plausible-on-average isn't good enough.

You have the boundaries. The next two modules are about pushing the inside of those boundaries hard.

Agent Logic -- *theaiguyw*
Fond du Lac, WI. This is module 4 of 6 in Tier 1 (Personal).