SEP 30, 2026 · JOSH BRANDL

It's OK to not trust AI; I don't trust it either

But AI can still help build the thing you can trust

A confession from the technical person

During a recent analysis of summarized collections data[1] for a client, I asked an LLM (Claude, ChatGPT, Copilot, etc.,) for advice on things I might be missing. It confidently asserted the following:

> *Together with the provider's CPT coding being entirely -26 from April onward, this suggests the practice converted billing from global to professional-only around March and April.*

There's one big problem with that assertion: I had only provided a partial view of January, March, and April. It had no clue what the data looked like after April, but it made decisions as if it did.

That's just one example of why I don't trust AI. And I don't think you should trust it either. But that doesn't mean you shouldn't use it. The key is using LLMs in a way that doesn't depend on blind trust.

Why I don't trust LLMs

People get things wrong all the time, so being wrong alone isn't the problem. The problem is that when AI is wrong, it's wrong with total confidence. A colleague who isn't sure will usually hedge, hesitate, or ask for help. AI doesn't do that. In my experience, this shows up in four ways.

- **It's too agreeable.** If you hand it your own bad assumption, it will usually run with it instead of questioning it.
- **It would rather answer than say "I don't know."** When it doesn't have the information, it tends to fill the gap with something that sounds right. That's exactly what happened with my CPT example.
- **It doesn't give the same answer twice.** Ask it the same question on Monday and again on Friday, and you can get two different answers, both delivered with the same confidence. Which one is right? You have no way of telling without something to check them against.
- **Its speed is great, except for when it is bounding in the wrong direction.** AI can produce a week's worth of wrong work in an afternoon, and because it all "looks right," sorting out the mistakes can take longer than starting over.

Working with it anyways

Think of AI as a brilliant new hire: lots of book knowledge, but no experience in your organization. You're happy to have them, but you still wouldn't let them send anything to your board without a review. The comparison has a limit, and it's the part that matters. A new hire learns from your corrections over months, while AI mostly doesn't carry those lessons forward. The guardrails you'd eventually relax for a person stay

 © 2026 Ancore, LLC® | All Rights Reserved.

in place for AI. So the goal isn't to wait for it to earn your trust. It's to set things up so its work can be checked from day one. Each of the four problems above has a habit that does that.

- **If it's too agreeable, make it argue with you.** Before it builds on your assumption, ask it to poke holes in it. Asking "What would have to be true for this to be wrong?" gets you a lot further than "Does this make sense?" Programmers have long explained problems out loud to a rubber duck until they spot their own mistake. AI is a rubber duck 2.0: one that can ask questions back.
- **If it makes things up, make it show its work.** Ask how it got to the answer and what it based that on, the same way you'd walk through a new analyst's first deliverable. If the explanation is vague or goes in circles, slow down.
- **If it's inconsistent, have it build something repeatable.** Don't ask it for the answer again; ask it for the query or formula that produces the answer, so you get the same result every time. More on that below.
- **If speed is the risk, start small and check often.** Start with the tedious stuff where you already know what the right answer looks like: reformatting a report, drafting a first-pass summary, cleaning up a list. You get time back, and you learn how far you can trust it before you hand it anything harder.

Build the thing you can trust

When I use AI for data work, I almost never ask for an answer directly. Instead, I ask it to write the query that gets me the answer. That may sound like a technical distinction, but you probably make the same one every month.

Think about the last time finance and operations showed up to the same meeting with two different numbers. What settled it was opening the spreadsheet and looking at which rows were counted, which formula ran, and which month's data it pulled. **A number typed into a cell can't be traced. A formula can.**

An answer from AI is the typed-in number. It might be right, but you have no way of knowing without redoing the work yourself. If AI writes a query, a script, or a workbook full of formulas instead, your analyst or a trusted partner can open it, follow it, and rerun it next month to get the same result.

Back to my CPT example: AI got the conclusion wrong, but it pointed me toward things I hadn't thought of. Instead of trusting its answer, I asked for its assumptions, conversion tables, and formulas, checked them, and did the analysis myself, a lot faster than starting from scratch. Now when someone asks where a number came from, I can show them.

You don't have to trust it

So it's OK not to trust AI. I don't, and I use it every day. What I do trust is the query I can read, the workbook my team can rerun, and the result I can check against something we already rely on. AI builds most of that for me, and a lot faster than I could on my own.

That's a much easier problem to solve than deciding whether AI is trustworthy. Your organization already solves it every time someone asks where a number came from and your team opens the file to show them.

© 2026 Ancore, LLC® | All Rights Reserved.

1. When using LLMs in a healthcare context, HIPAA compliance and safety are key concerns. Always consider contractual obligations, legal obligations, and your own organization's policies before providing any information to an LLM.

© 2026 Ancore, LLC® | All Rights Reserved.