

Phone out. Quick search. Checked!

Bill Smith, Founder, SafetyBuddy, Inc. · October 18, 2026

But did you actually verify anything, or just find another confident answer?

Five votes against Salmonella, and lunch is served!

AI Can Make Mistakes. Check Important Info.

Fine. How often? And who checks the checker?

There it is, that innocent little warning sitting at the bottom of your AI conversation.

“AI can make mistakes. Check important info.”

Well, isn't that comforting?

I'm curious. How many mistakes?

One in a million? One in a thousand? Every tenth answer? Maybe half the time?

So I ask AI.

"What's your error rate?"

Now things get interesting.

It might tell me that error rates depend on the model, the question, the subject, and how you define an error.

Fair enough.

But I didn't ask for a lecture on statistics. I asked how often it gets things wrong.

And apparently there isn't a single number it can give me.

So I try another approach.

"How often are you wrong when you tell me how often you're wrong?"

Now we're having fun.

If AI can make mistakes, then presumably its explanation of its own mistakes can contain mistakes. And if I ask it to check that explanation, the checker can make mistakes too.

I can keep going.

"Check your answer."

"Now check the check."

"And who checked that?"

Eventually I may have a beautifully written explanation, three apologies, four corrections, and absolutely no idea whether the original answer was right.

Welcome to the wonderful world of recursive uncertainty.

Now, imagine walking into your local supermarket.

There's a beautiful head of lettuce sitting in the produce aisle. Nice and green. Looks delicious.

Attached to it is a little label:

"This lettuce may contain salmonella. Check before eating."

Check how?

Should I bring a microscope? A microbiology laboratory? Perhaps ask another head of lettuce?

And while we're at it, how many heads are contaminated?

One in a million? One in ten?

"Sorry, we can't tell you. It depends on the lettuce."

Would you buy it?

How about a bottle of medication?

"These tablets may contain the wrong ingredient. Check before swallowing."

Or an airplane instrument?

"This altimeter may display an incorrect altitude. Please verify before landing."

I'm an engineer. I've worked with instruments, measurements, testing, and systems where a wrong answer can have consequences.

We don't expect perfection. We expect known limitations, testing, calibration, traceability, and some reasonable understanding of the probability and consequences of failure.

If an instrument has an accuracy specification, I can work with it.

If it has no accuracy specification, I can test it.

If it can't tell me its accuracy, can't reliably test itself, and can't tell me when it's wrong, I have a rather different engineering problem.

Now back to AI.

Here is a remarkable technology that can explain quantum mechanics, analyze legal language, interpret complicated medical literature, and write computer programs that actually work.

Sometimes.

And it can do all of those things incorrectly while sounding just as convincing.

The real difficulty isn't that AI makes mistakes. Humans make mistakes. Instruments make mistakes. Software makes mistakes.

The difficulty is that the user may have no reliable way to distinguish the correct answer from the incorrect one.

A physician asking AI about an unfamiliar medical paper may be able to check its sources.

A patient asking about that same paper may not.

An experienced programmer may recognize a dangerous piece of code.

A beginner may copy it straight into a working system.

And the person asking a legal question may be doing so precisely because hiring an attorney is beyond their means.

So we tell all of them:

"Check important info."

With what knowledge? What tools? What independent evidence?

Here's another interesting question.

Who decides what's important?

The AI? The same AI that just warned me it can make mistakes?

Suppose it makes a mistake deciding whether an answer is important enough to check.

Do we need another warning?

"AI can make mistakes, including mistakes about which mistakes you should check."

Perhaps that should be the new label.

Of course, there is a legitimate reason AI companies don't publish one universal error rate. These systems aren't single-purpose measuring instruments. Their accuracy varies enormously across tasks, and a benchmark result doesn't establish reliability for every question a person might ask.

But that makes the engineering challenge more important, not less.

We need to know what has actually been verified, what is inferred, what remains uncertain, and what independent checks have been performed.

Asking the same AI to reconsider its answer can help. It is not the same thing as independent verification.

And asking five AIs doesn't automatically solve the problem either. Five systems can share the same bad information, repeat the same assumption, and confidently agree on the same wrong answer.

Five votes against salmonella don't make the lettuce safe.

So perhaps we're asking the wrong question.

Instead of asking how intelligent the next generation of AI will be, maybe we should ask:

How will it demonstrate that its answers deserve our trust?

Not by saying "Trust me."

Not by asking the user to become an instant expert.

And certainly not by hiding the problem behind a tiny warning at the bottom of the screen.

I use AI extensively. I find it extraordinary. I also challenge it, argue with it, and occasionally discover that its beautifully constructed explanation has wandered completely off the reservation.

That's part of working with a technology still evolving at remarkable speed.

But as AI moves deeper into healthcare, finance, law, engineering, and everyday decision-making, we need something more substantial than a disclaimer.

We need ways to test its answers against independent evidence, expose uncertainty, identify consequential errors, and keep meaningful control in human hands.

Meanwhile, back in the supermarket, I'm still staring at that head of lettuce.

The label says it may contain salmonella.

I ask the lettuce how often.

It replies:

"That's an excellent question."

Oh, brother.

SafetyBuddy, Inc. is a 501(c)(3) nonprofit · safetybuddy.org · Also on LinkedIn with more AI articles by SafetyBuddy's Founder, Bill Smith.