

How would *we know?*

For AI already operating in the physical world,
safety needs proof somebody else can check.

The answer is not necessarily slower AI. It is verifiable AI.

The question is legitimate.


The answer needs proof.

The debate over frontier AI safety matters. Physical AI is already helping vehicles brake, warning drivers and making calls where mistakes have a physical cost.

Proof that somebody outside the company can check.

Whatever standard gets set at the frontier is the standard the rest of us will inherit.

We already expect this elsewhere.



Drug trials

Evidence of efficacy and safety.



Food inspections

Checks beyond the producer's assurances.



Audited accounts

Independent examination of the record.

The question is not whether these systems are out there. It is how anyone knows they work.

We measured *the road.*

From September 2025 through August 2026, we observed driving across the United States. A record to measure against, before judging a system.

1.105B

miles of driving

74.3M

rides



95,805

collisions



1,815,827

events crossing the
near-miss threshold

A particular set of drivers.


Mostly professional, mostly in dense cities. More than half the events occurred in three states. This is not the whole country.

These numbers describe the road. They do not, by themselves, establish that an automated system is safe.

One collision.

Almost nineteen warnings.

The collision gets a police report, a claim, a file. The other events tell us where the risk sits before anybody gets hurt.



Shares of the two reported event counts combined. Rounded to two decimals. Bars use the same 0–100% scale.


A collision-only baseline misses most of the record.

Nothing in these counts explains why most near misses happened. That is the next problem to solve.



A safety claim needs *a baseline.*

You cannot claim a system made a road safer unless you know what that road does without it. The same standard applies to our own claims.




67%

Fewer most
serious collisions

Reported result in one
national fleet using safety
models built on this record.

Indexed illustration of the reported reduction: 100 → 33. These are not raw collision counts.



One national fleet. A specific outcome.

This is the result reported in the essay. It does not establish the safety of every automated system or every road.

*“Safety is the study of things
that did not happen.”*

Count the events that used to end in a collision and no longer do. A safety claim must be a comparison that can be checked.

Build for verification.

Protect the person.

Privacy has to decide how the record gets built in the first place. It cannot be a promise added after the dataset exists.



What happened on the road.

A record that describes an event without needing to know who was there.



Let somebody else grade it.

Testing your own work on your own data and publishing the result leaves out the point of an audit: third-party validation.

A shared record, across different interests.



Carmaker



Fleet



City



Insurer

We make no vehicles and compete with none of our buyers. Independence should be visible in how the company is built and how results are validated.

The answer is *verifiable AI.*

If you sell a system that acts in the physical world, publish the baseline you measured it against, and let somebody outside your company check it.

01

Publish the baseline

Show what the road does without the system.

02

State the scope

Attach the population, period and outcome to the claim.

03

Open it to an outside check

Let an independent party examine the evidence.

“Not reassurance. A method.”

I do not know what the frontier does in ten years. I know what happened across 1.105 billion miles of road last year, because we counted it.