

My Tesla Self-Driving Car Crash

The Tesla was driving perfectly—until it wasn't.

By Raffi Krikorian

The smell was strange. Sharp. Chemical. Wrong. The concrete wall was too close. My glasses were gone. One of my kids was standing on the sidewalk next to our car—not crying, just confused.

The seat belt had held. The crumple zone had crumpled. The airbag had fired. Everything designed to protect bodies had done its job. But the car, a Tesla Model X, was totaled.

One Sunday last fall, my kids and I were on a drive we'd done hundreds of times, winding through the residential streets of the Bay Area to drop my son off at his Boy Scouts meeting. The Tesla was in Full Self-Driving mode, driving perfectly—until it wasn't.

What happened next, I've had to piece together. My memory is hazy, and some of it comes from one of my sons, who watched the whole thing unfold from the back seat. The car was making a turn. Something felt off—the steering wheel jerked one way, then the other, and the car decelerated in a way I didn't expect. I turned the wheel to take over. I don't know exactly what the system was doing, or why. I only know that somewhere in those seconds, we ended up colliding with a wall.

You might think I'd have known what to do in this situation. I used to run the self-driving-car division at Uber, trying to build a future in which technology protects us from accidents. I had thought about edge cases, failure modes, the brittleness hiding behind smooth performance. My team trained human drivers on when and how to intervene if a self-driving

car made a mistake. In the two years I ran the division, we had no injuries in our early pilot programs.

With my own Tesla, I started out using Full Self-Driving as the default setting only on highways. That's where it makes sense: You have clear lane markers and predictable traffic patterns. Then, one day, I tried it on a local road, and it worked well enough to become a habit.

Despite the accident, we were lucky. I walked away with a stiff neck, a concussion, a few days of headaches, and some memories I can't shake. The kids climbed out unharmed. Still, you could say I was crushed in what the researcher Madeleine Clare Elish [calls the moral crumple zone](#). Some parts of a car are specifically designed to absorb damage in a crash, to protect the people inside. But when

complex automated systems fail, Elish argues, it's the human users who take the blame. My car's Full Self-Driving mode logged flawless miles for three years, but when the accident happened, it was my name on the insurance report.

And the car has evidence. While you're at the wheel, it logs your hand position, your reaction time, whether you're keeping your eyes on the road—thousands of data points, processed by the vehicle. After crashes, Tesla has used these data to shift blame onto drivers. Following a fatal collision in Mountain View, California, in 2018, [the company released a statement](#) in which it noted that “the vehicle logs show that no action was taken.” (Tesla did not respond to a request for comment.)

While Tesla can access these records, it's not so easy for drivers. They can request their data,

but some say they've received only fragments
—and have had to go to court to get more.

When [plaintiffs in a Florida wrongful-death case sought key evidence](#) of how one of Tesla's driver-assistance systems had failed, the company said it didn't have the data. The plaintiffs had to hire a hacker, who recovered them from a computer chip in the crashed vehicle. Later, Tesla stated that the data had been sitting on its own servers for years, and that the company failed to locate them by mistake. (A judge did not find "sufficient evidence" to conclude that Tesla had sought to hide the data.)

For now, the _____

legal

principle is simple: You're responsible. Though Tesla originally called its technology "Full Self-Driving Capability," the system is officially

classified as “Level 2” partial driver automation, which means the human must remain in control at all times. Last year, a judge in California [found Tesla’s original name “unambiguously false”](#) and misleading to consumers; Tesla now uses “Full Self-Driving (Supervised).” When a Tesla using a version of the technology [killed two people in California in 2019](#), the car’s own logs were used to prosecute the driver for failing to prevent the crash—not the company that designed the system. The company was held accountable in a major verdict for the first time only last year, when a jury found Tesla partly liable in the Florida wrongful-death case and awarded \$243 million to the plaintiffs.

A similar pattern is emerging everywhere algorithms are asked to work alongside humans: in our inboxes, our search results,

our medical charts. These systems are building toward full automation, but they're not there yet. Computers still regularly make mistakes that require human oversight to avoid or fix.

Full Self-Driving works almost all of the time—Tesla's fleet of cars with the technology logs millions of miles between serious incidents, by the company's count. And that's the problem: We are asking humans to supervise systems designed to make supervision feel pointless. A machine that constantly fails keeps you sharp. A machine that works perfectly needs no oversight. But a machine that works almost perfectly? That's where the danger lies. After a few hours of flawless performance, research shows, drivers are prone to start overtrusting self-driving systems. After a month of using adaptive cruise control, drivers were more than six times as likely to look at their

phone, [according to one study](#) from the Insurance Institute for Highway Safety.

Tesla's description of Full Self-Driving on its website warns, "Do not become complacent," and I didn't think I was. Before my accident, I had my hands on the wheel. But I was driving the way the system had conditioned me to: monitoring instead of steering, trusting the software to make the right call. The familiarity curve bends toward complacency, and the companies building these systems seem to know it. I certainly did. I got lulled anyway.

Psychologists call this the vigilance decrement. Monitoring a nearly perfect system is boring. Boredom leads to mind-wandering. [The research is unforgiving](#): Drivers need five to eight seconds to mentally reengage after an automated driving system gives control back. But emergencies can unfold much faster than

that. The driver's physical reaction might be instantaneous—grabbing the wheel, hitting the brake. But the mental part? Rebuilding context, recognizing what's wrong, deciding what to do? That takes time your brain doesn't have.

The driver in the 2018 Mountain View accident had six seconds before his car steered itself into a concrete median. He never touched the wheel. That same year in Tempe, Arizona, sensors in an Uber test vehicle [detected a pedestrian nearby with 5.6 seconds of warning](#). The safety driver looked up and took the wheel with less than a second left. By then, it was just physics.

In my case, I did take action before my accident. But I was asked to snap from passenger back to pilot in a fraction of a second—to override months of conditioning in

the time it takes to blink. The logs would show that I turned the wheel. They wouldn't show the impossible math.

i don't know enough about what actually happened during my accident to say that Tesla's technology crashed the car. But the problem is bigger than one company's self-driving system. It's about how we're building every AI system, every algorithm, every tool that asks for our trust and trains us to give it. The pattern is everywhere: Condition people to rely on the system. Erode their vigilance. Then, when something breaks, point to the terms of service and blame them for not paying attention.

My car didn't warn me when it was confused. Chatbots don't, either; they deliver their results in the same confident voice, whether they're right or hallucinating. They perform expertise,

even when the sources they cite are dubious or fabricated. They use technical language in an authoritative tone. And we believe them, because why wouldn't we? They've been right so many times before.

Cars train us mile by mile; AI trains us week by week. In week one, you read a chatbot's output carefully. By week three, you're copying and pasting without reading. The errors don't disappear, but your vigilance does. So does your judgment, until one day you realize that you can't remember which ideas in a memo were yours and which were generated by AI. What does it say about us that we've [handed over our thinking](#) so willingly?

When my car failed, it was immediate and palpable. With chatbots, the failure is silent and invisible. You find out about it later, if at all —after the email is sent, the decision made,

the code shipped. By the time you catch the mistake, it's already out there with your name on it. When the system works, you look efficient. When it fails, your judgment is questioned, sometimes with catastrophic consequences. In 2023, a New York lawyer [was sanctioned for citing six cases that didn't exist](#). ChatGPT had invented them, but he'd trusted it, and the court blamed him, not the tool. Because a chatbot never gets fired.

We're experiencing an uncanny valley of autonomy. Computer systems aren't just almost human; they are almost capable of working on their own. When they fail, someone has to absorb the cost. Right now, [that someone is us](#). But when we pay for a self-driving car or an AI tool, we think we're buying a finished product, not signing up to test a work in progress.

This “almost” phase isn’t a brief transition. It’s the product—one that will be with us for years, maybe decades. So it’s important to notice the patterns. When an AI system never admits uncertainty, or when a car’s marketing says “self-driving” but the fine print says “driver responsible,” that’s a warning sign. When you realize that you haven’t really been paying attention for the past 10 miles, or the past 10 auto-composed emails, that’s the trap.

Things don’t have to be this way, but they won’t change unless consumers see the situation clearly and refuse to accept it. We should reject the deal we’ve been handed—the one where the terms of service become a shield for companies and a sword against users. We should demand that companies share the risk they’re enticing us into taking. If

they design for complacency, they should get some of the blame when their product fails.

This isn't a utopian goal. In July 2025, the Chinese carmaker BYD announced that it would pay for the damage caused by crashes involving its self-parking feature, sparing the driver's insurance and record. It's only one company, and only one feature, but it proves that accountability is a choice. Other businesses can be persuaded to opt in, too.

My kids were in the back seat when I had my car accident. One day, they'll have their own cars and use AI in ways that I can't even imagine yet. The systems they inherit will be built either to elevate them or to lull them and blame them when things go wrong. I want them to notice when they're being trained. I want them to ask who absorbs the cost, and the damage.

This article appears in the [April 2026](#) print edition with the headline “My Self-Driving Car Crash.”

ABOUT THE AUTHOR

[Raffi Krikorian](#)

[Raffi Krikorian](#) is the chief technology officer of Mozilla.