

Self-driving cars can't be perfectly safe – what's good enough? 3 questions answered

March 27 2018, by Nicholas G. Evans



Credit: AI-generated image ([disclaimer](#))

On March 19, an Uber self-driving vehicle being tested in Arizona [struck and killed Elaine Herzberg](#), who was walking her bike across the street. This is the first time a self-driving vehicle has killed a pedestrian, and it raises questions about the ethics of developing and testing emerging technologies. Some answers will need to wait until the

full investigation is complete. Even so, [Nicholas Evans](#), a philosophy professor at the University of Massachusetts-Lowell who studies the ethics of autonomous vehicles' decision-making processes, says some questions can be answered now.

1. Could a human driver have avoided this crash?

Probably so. It's easy to think that most [people](#) would have trouble seeing a pedestrian crossing a [road](#) at night. But what's already clear about this particular event is that the road [was not as dark as the local police chief initially claimed](#).

The chief also [originally said](#) Herzberg suddenly stepped out into traffic in front of the car. However, the [disturbing and alarming video footage](#) released by Uber and local authorities shows this isn't true: Rather, Herzberg had already walked across one lane of the two-lane road, and was in the process of continuing the road-crossing when the Uber hit her. (The safety driver also didn't notice the pedestrian, but video suggests [the driver was looking down](#), not through the windshield.)

A normal human driver, someone actively paying attention to the road, would likely have had little problem avoiding Herzberg: With headlights on while [traveling 40 mph](#) on an actually dark road, it's not difficult to avoid obstacles on a straightaway when they're [100 or more feet ahead](#), including people or wildlife trying to get across. This crash was avoidable.

One tragic implication of that fact is clear: A self-driving car killed a person. But there is a public significance too. At least this one Uber car drove itself on populated streets while unable to perform the crucial safety task of detecting a pedestrian, and braking or steering so as not to hit the person.

In the wake of Herzberg's death, the [safety and reliability of Uber's self-driving cars](#) has come into question. It's also worth examining the ethics: Just as Uber has been criticized for [exploiting its drivers for profits](#), the company may arguably be [exploiting the driving, riding and walking public](#) for its own research purposes.

2. Even if this crash was avoidable, are self-driving cars still generally safer than human-driven cars?

Not yet. The death toll on U.S. roads is indeed alarming: approximately [32,000 deaths per year](#). The [federal estimate](#) is that 1.18 people die per 100 million road miles driven by humans. Uber's cars only [drove 3 million miles](#), however, before their first fatality. It's not fair to do statistical analysis from a single point of data, but it's not a great start: Companies should be aiming to make their robots at least as good as humans, if not yet fulfilling the [promise of being significantly better](#).

Even if Uber's autonomous cars were better drivers, the numbers don't tell the whole story. Of the 32,000 people who die on U.S. roads each year, [5,000 to 6,000 are pedestrians](#). When aiming for safety improvements, should the goal be to reduce overall deaths – or to put special emphasis on protecting the most vulnerable victims? It's certainly hypothetically possible to imagine a self-driving car system that cuts overall road deaths in half – to 16,000 – while doubling the pedestrian death rate – to 12,000. Overall, that might seem far better than human drivers – but not from the perspective of people walking along the nation's roads!

My research group has been working to develop ethical decision frameworks for self-driving cars. One potential approach is called "[maximin](#)." Most fundamentally, that way of thinking suggests

people designing [autonomous vehicles](#) – both physically and in terms of software that runs them – should identify the worst possible outcomes of any decision, even if rare, and work to minimize their effects. Anyone who has been unfortunate enough to be hit by a car both as a pedestrian and while in a vehicle knows that being on foot is far worse. Under maximin, people should design and test cars, among other things, to prioritize pedestrian safety.

Maximin probably isn't the best possible – and certainly isn't the only – moral decision theory to use. In some cases, the worst outcome could be avoided if a car never pulls out of its driveway! But maximin provides food for thought about how to integrate self-driving cars into daily life. Even if autonomous cars are always evaluated as safer than humans, what counts as "safer" matters very much.

3. How much better should self-driving cars be than humans before the public accepts them?

Even if people could agree on the ways in which self-driving cars should be safer than humans, it's not clear that people should be okay with self-driving cars when they first become only barely better than humans. If anything, that's when tests on city streets should begin.

Consider a new drug developed by a pharmaceutical company. The company can't market it as soon as it's proven not to kill people who take it. Rather, the drug has to go through a [series of tests proving it is effective](#) at treating the symptom or condition it's intended to. Increasingly, drug tests seek to prove a medication is [significantly better](#) than what's already on the market. People should expect the same with self-driving cars before companies put the

public at risk.

The crash in Arizona wasn't just a tragedy. The failure to see a pedestrian in low light was an avoidable basic error for a [self-driving car](#). Autonomous vehicles should be able to do much more than that before they're allowed to be driven, even in tests, on the open road. Just like pharmaceutical companies, massive technology companies should be required to thoroughly – and ethically – test their systems before their self-driving cars serve or endanger the public.

This article was originally published on [The Conversation](#). Read the [original article](#).

Provided by The Conversation

Citation: Self-driving cars can't be perfectly safe – what's good enough? 3 questions answered (2018, March 27) retrieved 5 July 2024 from <https://phys.org/news/2018-03-self-driving-cars-perfectly-safe-good.html>

<p>This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.</p>
--