

Before Hitting the Road, Self-Driving Cars Should Have to Pass a Driving Test

Researchers can't always tell exactly why something works but they can evaluate the outcome

Srikanth Saripalli, *The Conversation US* February 22, 2018



Credit: [Flickr](#) (CC BY-SA 2.0)

The following essay is reprinted with permission from [The Conversation](#), an online publication covering the latest research.

What should a self-driving car do when a nearby vehicle is swerving unpredictably back and forth on the road, as if its driver were drunk? What about encountering a vehicle driving the wrong way? Before autonomous cars are on the road, everyone should know how they'll respond in unexpected situations.

I [develop, test and deploy autonomous shuttles](#), identifying methods to ensure self-driving vehicles are safe and reliable. But there's no testing track like the country's actual roads, and no way to test these new machines as thoroughly as modern human-driven cars have been, with [trillions of miles driven every](#)

year for decades. When self-driving cars do [hit the road, they crash](#) in ways both [serious](#) and [minor](#). Yet all their decisions are made electronically, so how can people be confident they're driving safely?

Fortunately, there's a common, popular and well-studied method to ensure new technologies are safe and effective for public use: The testing system for new medications. The basic approach involves ensuring these systems do what they're intended to, without any serious negative side effects—even if researchers don't fully understand how they work.

Wide-ranging effects

Self-driving cars [are expected](#) to [improve road safety](#), [freeing up drivers' time and attention](#) and [transforming cities](#) and even societies.

The regulations that are created for self-driving cars will have massive effects that [ripple throughout the economy and society](#). The rules are likely to come from some combination of the two current automotive regulators, the federal [National Highway Traffic Safety Administration](#) and [state departments of transportation](#).

Federal rules focus primarily on safety standards for structural, mechanical and electrical components of the vehicles, like airbags and seat belts. States can enforce their own safety rules—for example, regulating emissions and handling driver licensing and vehicle registration, which often also includes requiring insurance coverage.

Current regulations

Today's state and federal rules treat drivers and cars as separate entities. But self-driving cars, by definition, combine the two. Without consistency between those regulations, confusion will reign.

The Obama administration came up with [116 pages of regulations](#) with lots of details, but little understanding of how self-driving cars worked. For example, they called for each car to have human-readable permanent labels listing its specific self-driving capabilities, including limits on speeds, specific highways and weather conditions, all of which would be extremely confusing for users. The regulations also called for ethical decisions to be made “consciously and intentionally”—which is [questionable, if not impossible](#), for a machine.

The Trump administration pared down the rules to [26 pages](#), but have not yet addressed the important issue of testing self-driving cars.

Examining algorithms

Testing algorithms is very like testing medications. In both cases, researchers can't always tell exactly why something works (especially in the case of [machine learning](#) algorithms), but it is nevertheless possible to evaluate the outcome: Does a sick person get well after taking a medication?

The U.S. Food and Drug Administration [requires medicines be tested](#) not for their mechanisms of treatment, but for the results. The two main criteria are effectiveness—how well the medicine treats the condition it's intended to—and safety—how severe any side effects or other problems are. With this method, it's possible to prove a medication is safe and effective without knowing how it works.

Similarly, federal regulations could—and should—require testing for self-driving cars' algorithms. To date, governments have tested cars as machines, ensuring steering, brakes and other functions work properly. Of course, there are also government tests for human drivers.

A machine that does both should have to pass both types of tests—particularly for [vehicles that don't allow for human drivers](#).

Evaluating judgment

In my view, before allowing any specific self-driving car on the road, NHTSA should require test results from the car and its driving algorithms to demonstrate they are safe and reliable. The closest standard at the moment is [California's requirement](#) that all manufacturers of self-driving cars submit annual reports of how many times a human driver had to take control of its vehicles when the algorithms failed to function properly.

That's a good first step, but it doesn't tell regulators or the public anything about what the vehicles were doing or what was happening around them when the humans took over. Tests should examine what the algorithms direct the car to do on freeways with trucks, and in neighborhoods with animals, kids, pedestrians and cyclists. Testing should also look at what the algorithms do when both vehicle performances and sensors' input is compromised by rain, snow or other weather conditions. Cars should run through scenarios with temporary construction zones, four-way intersections, wrong-way vehicles and police officers giving directions that contradict traffic lights and other situations. Human driving tests include some evaluations of a driver's judgment and decision-making, but tests for self-driving cars should be more rigorous because there's no way to rely on human-centered concepts like instinct, reflex or self-preservation. Any action a machine takes is a choice, and the public should be clear on how likely it is that those choices will be safe ones.

Comparing with humans

Self-driving cars' algorithms constantly calculate probabilities. How likely is it that a particular shape is a person? How likely is it that the sensor data means the person is walking toward the road? How likely is it that the person will step into the street? How likely is it that the car can stop before hitting her? This is in fact [similar to how the human brain](#) works.

That presents a straightforward opportunity for testing autonomous cars and any software updates a manufacturer might distribute to vehicles already on the road: They could present human test drivers and self-driving algorithms with the same scenarios and monitor their performance over many trials. Any self-driving car that does as well as, or better than, people, can be certified as safe for the road. This is very much like the method used in drug testing, in which a new medication's performance is rated against existing therapies and methods known to be ineffective, like the typical [placebo sugar pill](#).

Companies should be free to test any innovations they want on their closed tracks, and even on public roads with human safety drivers ready to take the wheel. But before self-driving cars become regular products available for anyone to purchase, the public should be shown clear proof of their safety, reliability and effectiveness.

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