

Redefining ‘safety’ for self-driving cars

November 28, 2017 6.26am EST



Author



Srikanth Saripalli

Associate Professor in Mechanical Engineering, Texas A&M University

When self-driving cars get in crashes, who's to blame? Santa Clara Valley Transportation Authority via AP

In early November, a self-driving shuttle and a delivery truck collided in Las Vegas. The event, in which no one was injured and no property was seriously damaged, attracted media and public attention in part because one of the vehicles was driving itself – and because that shuttle had been operating for only less than an hour before the crash.

It's not the first collision involving a self-driving vehicle. Other crashes have involved Ubers in Arizona, a Tesla in “autopilot” mode in Florida and several others in California. But in nearly every case, it was human error, not the self-driving car, that caused the problem.

In Las Vegas, the self-driving shuttle noticed a truck up ahead was backing up, and stopped and waited for it to get out of the shuttle's way. But the human truck driver didn't see the shuttle, and kept backing up. As the truck got closer, the shuttle didn't move – forward or back – so the truck grazed the shuttle's front bumper.

As a researcher working on autonomous systems for the past decade, I find that this event raises a number of questions: Why didn't the shuttle honk, or back up to avoid the approaching truck? Was stopping and not moving the safest procedure? If self-driving cars are to make the roads safer, the bigger question is: What should these vehicles do to reduce mishaps? In my lab, we are developing self-driving cars and shuttles. We'd like to solve the underlying safety challenge: Even when autonomous vehicles are doing everything they're supposed to, the drivers of nearby cars and trucks are still flawed, error-prone humans.

The driver who was backing up a truck didn't see a self-driving shuttle in his way. [Kathleen Jacob/KVVU-TV via AP](#)

How crashes happen

There are two main causes for crashes involving autonomous vehicles. The first source of problems is when the sensors don't detect what's happening around the vehicle. Each sensor has its quirks: GPS works only with a clear view of the sky; cameras work with enough light; lidar can't work in fog; and radar is not particularly accurate. There may not be another sensor with different capabilities to take over. It's not clear what the ideal set of sensors is for an autonomous vehicle – and, with both cost and computing power as limiting factors, the solution can't be just adding more and more.

The second major problem happens when the vehicle encounters a situation that the people who wrote its software didn't plan for – like having a truck driver not see the shuttle and back up into it. Just like human drivers, self-driving systems have to make hundreds of decisions every second, adjusting for new information coming in from the environment. When a self-driving car experiences something it's not programmed to handle, it typically stops or pulls over to the roadside and waits for the situation to change. The shuttle in Las Vegas was presumably waiting for the truck to get out of the way before proceeding – but the truck kept getting closer. The shuttle may not have been programmed to honk or back up in situations like that – or may not have had room to back up.

The challenge for designers and programmers is combining the information from all the sensors to create an accurate representation – a computerized model – of the space around the vehicle. Then the software can interpret the representation to help the vehicle navigate and interact with whatever might be happening nearby. If the system's perception isn't good enough, the vehicle can't make a good decision. The main cause of the fatal Tesla crash was that the car's sensors couldn't tell the difference between the bright sky and a large white truck crossing in front of the car.

If autonomous vehicles are to fulfill humans' expectations of reducing crashes, it won't be enough for them to drive safely. They must also be the ultimate defensive driver, ready to react when others nearby drive unsafely. An Uber crash in Tempe, Arizona, in March 2017 is an example of this.

A human-driven car crashed into this Uber self-driving SUV, flipping it on its side. [Tempe Police Department via AP](#)

According to media reports, in that incident, a person in a Honda CRV was driving on a major road near

the center of Tempe. She wanted to turn left, across three lanes of oncoming traffic. She could see two of the three lanes were clogged with traffic and not moving. She could not see the farthest lane from her, in which an Uber was driving autonomously at 38 mph in a 40 mph zone. The Honda driver made the left turn and hit the Uber car as it entered the intersection.

A human driver in the Uber car approaching an intersection might have expected cars to be turning across its lane. A person might have noticed she couldn't see if that was happening and slowed down, perhaps avoiding the crash entirely. An autonomous car that's safer than humans would have done the same – but the Uber wasn't programmed to.

Improving testing

That Tempe crash and the more recent Las Vegas one are both examples of a vehicle not understanding the situation enough to determine the correct action. The vehicles were following the rules they'd been given, but they were not making sure their decisions were the safest ones. This is primarily because of the way most autonomous vehicles are tested.

The basic standard, of course, is whether self-driving cars can follow the rules of the road, obeying traffic lights and signs, knowing local laws about signaling lane changes, and otherwise behaving like a law-abiding driver. But that's only the beginning.

Sensor arrays atop and along the bumpers of a research vehicle at Texas A&M. Swaroopa Saripalli, CC BY-ND

Before autonomous vehicles can really hit the road, they need to be programmed with instructions about how to behave when other vehicles do something out of the ordinary. Testers need to consider other vehicles as adversaries, and develop plans for extreme situations. For instance, what should a car do if a truck is driving in the wrong direction? At the moment, self-driving cars might try to change lanes, but could end up stopping dead and waiting for the situation to improve. Of course, no human driver would do this: A person would take evasive action, even if it meant breaking a rule of the road, like switching lanes without signaling, driving onto the shoulder or even speeding up to avoid a crash.

Self-driving cars must be taught to understand not only what the surroundings are but the context: A car approaching from the front is not a danger if it's in the other lane, but if it's in the car's own lane circumstances are entirely different. Car designers should test vehicles based on how well they perform difficult tasks, like parking in a crowded lot or changing lanes in a work zone. This may sound a lot like giving a human a driving test – and that's exactly what it should be, if self-driving cars and people are to coexist safely on the roads.



Autonomous vehicles

Driverless cars

Autonomous systems

Car safety