THE UBER FATAL CRASH,

WAS IT A TRAGEDY OR ONE MORE ACCIDENT?

Figure 1 Uber's modified Volvo XC90 sport-utility vehicle. SOURCE: UBER

The Uber's redundant suite of sensors include radar, LIDAR (Light Detection and Ranging) and a camera array, all focused on sensing and seeing persons and objects in the roadway, day or night. So, how could they all not have detected Elaine Herzberg walking across an open street? By all accounts, she was a common person in usual clothing, doing what common people do—crossing a street. She was also pushing a bicycle full of bags, which should have created a larger profile that would have made her even more likely to be detected by the car.

The tragedy of the first person killed by an autonomous vehicle points to a potential vulnerability with the nascent technology now being tested on the open roads: while robot-cars, powered by sophisticated sensors and cameras, can reliably see their surroundings, the software doesn't always understand what it detects.

Unfortunately, there are numerous ways autonomous vehicle hardware and software can fail, alone or in combination. It is important to emphasize that these are merely speculative points of discussion.

First, sensors may not sense. Ms. Herzberg was walking at night, and less light makes detection more difficult for the human eye and for camera sensors. However, while darkness can limit the vision of the cameras, radar functions equally well in day or night. LIDAR actually functions better in the dark because the

brilliance of sunshine can sometimes create interference. Therefore, absence of light should be irrelevant to LIDAR and radar performance, but each sensor can separately be obstructed or dirty. Investigators will undoubtedly attempt to verify that all sensors were clean, calibrated and operational. After that — more sensors or different kinds of sensors might be prudent to fill any detection gaps in similar operating domains.

Second, sensors can often see an obstruction, but machine learning algorithms can assess it as a shadow or something else "pre-built" by programmers as objects to ignore. Recording that the investigation with Tesla Autopilot in 2016 focused on whether the semi-tractor trailer was misclassified as an item to be ignored. If an object is sensed but misclassified, programming can select the wrong course of action—including no action—in response.

Third and finally, a car's hardware could sense a figure, but due to conflicting inputs from different sensors, or prioritization of one sensor over the other or due to insufficient learning data, the artificial intelligence 'brain' might find no definitive object classification match and simply not know how to respond. Given a limited library of objects and experiences, and a short list of actions to take in response, software might select proceeding on course.

I believe, being a vehicle technology specialist at Bright Consulting, the Uber car, which had a backup driver sitting in it at the time of the crash, should have detected the woman victim on its built-in sensors. It should have been able to detect her unless Uber engineers have introduced a blind zone around the SUV perimeter that could not detect pedestrians.

The victim did not come out of nowhere. She was walking on a dark road, but it was an open road, so LIDAR and radar should have detected and classified her as a human. The video suggests multiple failures of Uber automated system and its safety driver.

Besides that, in the case that everything goes wrong, these systems, if responsibly implemented, are supposed to prevent accidents – that is the main purpose of the autonomous vehicles.

AUTONOMOUS VEHICLES – THE SAFETY DRIVER

More than 600 autonomous vehicles are being tested on public roads in Arizona. Auto and technology companies need to ride several miles in real roads to improve their self-driving systems. Most of those vehicles need one, or often more, people in the car to make sure everything is running safe.

Safety drivers typically operate robotic vehicles in 6-8 hours work period and provide critical feedback to engineers. However, qualifications, training and practices vary widely. Most companies have two people in the car at all times: one ready to take the wheel and the other monitoring code. Uber seemed to have only one in the vehicle that hit and killed a pedestrian last month.

State regulations also vary. In Arizona, autonomous vehicle safety drivers are only required to have a valid driver's license, just like any other driver on the road. In California, the drivers must have completed a test driver training program by the company they work for, not have any at-fault collisions resulting in injury or death on their driving record and be clean of any citations for driving under the influence of alcohol or drugs in the last decade.

The video indicates that neither the self-driving system nor the human safety driver behind the wheel hit the brakes when she apparently stepped off a median and onto the roadway. The human driver (who appears to

look down at something off-camera in the moments before impact) told police he didn't see the pedestrian coming, and the autonomous system behaved as if it hadn't either.

People are very good at detecting things, even though we don't have lasers coming out of our eyes. Our reaction times aren't the best, but if it's clear that the car isn't going to respond, or has responded wrongly, a driver would react.

Possibly a common driver driving a typical vehicle on a dry asphalt road would have perceived and reacted steering or activating brakes before hitting Ms. Herzberg.

REDUNDANT, EXPENSIVE AND SOPHISTICATED SYSTEMS – HOW THE UBER VEHICLES "WORKED" AND WORKS.

The pedestrian killed by the Uber SUV had crossed at least one open lane of road before being hit, according to a video of the crash that raises new questions about autonomous-vehicle technology.

Something unexpectedly entering the vehicle's path is pretty much the first emergency event that autonomous car engineers look at. The situation could be many things — a stopped car, a deer, a pedestrian — and the systems are one and all designed to detect them as early as possible, identify them and take appropriate action. That could be slowing, stopping, swerving, anything.

Uber's vehicles are equipped with several different imaging and detection systems. No less than three different ones should have picked up the victim in this case.

ONE LIDAR or light detection and ranging, a system that produces a 3D image of the car's surroundings multiple times per second. Using infrared laser pulses that bounce off objects and return to the sensor, it can detect static and moving objects in considerable detail, day or night. Heavy snow and fog can obscure a LIDAR's lasers, and its accuracy decreases with range, but for anything from a few meters to a hundred meters, it's an "indispensable" and expensive imaging tool and one that is found on practically every self-driving car.

The LIDAR unit, if operating correctly, should have been able to make out the person in question.

10 Radars, like LIDAR, sends out a signal and waits for it to bounce back, but it uses radio waves instead of light. This makes it more resistant to interference, since radio can pass through snow and fog, but also lowers its resolution and changes its range profile.

The radar array employed in the Uber SUV provides 360 degrees of coverage — or it was supposed to.

The radar signature of a person is not really clear, but it's very likely they would have at least shown up, confirming what the LIDAR detected.

LIDAR and radar are great for locating shapes, but they're no good for reading signs, figuring out what color something is and so on. That's a job for the visible-light cameras with sophisticated computer vision algorithms running in real time.

The 7 cameras on the Uber vehicle watch for patterns that indicate braking vehicles (sudden red lights), traffic lights, crossing pedestrians and so on. Front, rear and side facing cameras were used to get a complete picture of the scene into which the car is driving.

Detecting people is probably the most important issue and the system does that splitting an image - the algorithms that do it have gotten quite good. Splitting an image, generally also involves identifying things like signs, trees, sidewalks and more.

That said, it can be hard at night. But that's an obvious problem, that's why the vehicle has the two previous systems, which work day and night. Even in the darkness, a person wearing all black would show up on LIDAR and radar, warning the car that it should perhaps slow and be ready to see that person in the headlights.

There is also a central computing unit that takes the input from these sources and creates its own more complete representation of the world around the car. A person may disappear behind a car in front of the system's sensors, for instance, and no longer be visible for a second or two, but that doesn't mean they stop existing. This goes beyond simple object recognition and begins to bring in broader concepts of intelligence such as object permanence, predicting actions and the like.

In my opinion, the sophisticated sensors on the autonomous vehicle almost certainly detected the woman pushing her bicycle along the road - unless Uber engineers have introduced a blind zone around the SUV perimeter that could not detect pedestrians. However, it is possible that the car LIDAR and Radar sensors may not have realized it was detecting a person. The real challenge is that you need to distinguish the difference between people, cars, dirt, newspaper and anything else that could be out in the road environment. The detection algorithms may have failed to detect the person or distinguish her.

Driverless cars "see" the world around them using data from cameras as well as radar and LIDAR sensors. That's supposed to enable the vehicle to know, in real time, where to go and when to stop. Nevertheless, pedestrian identification remains the major challenge for self-driving systems.

It isn't clear what the circumstances were under this tragic accident, but the vehicle should have detected the person and caused the car to react appropriately. Furthermore, if one system didn't work, another should have worked.

Author

Murilo Cesar Perin Briganti

Vehicle Technology Specialist

Bright Consulting

murilo.briganti@brightisd.com

www.brightisd.com