Warning: Objects in driverless car sensors may be closer than they appear

March 14 2022, by Ken Kingery

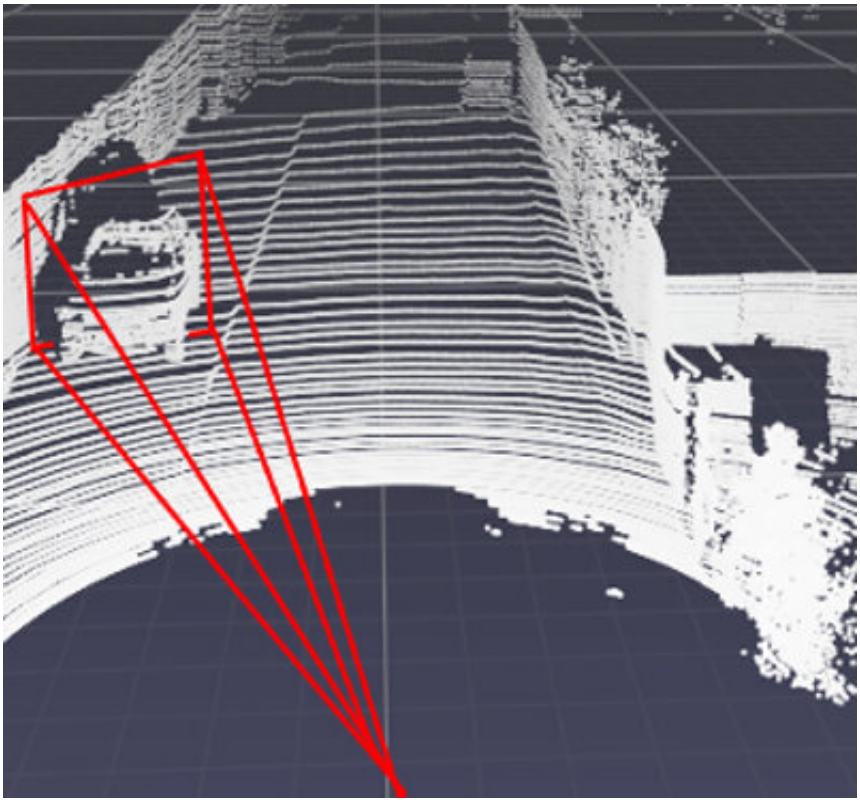

The area shown to be vulnerable to attacks in new research stretches out in front of a camera's lens in the shape of a frustum—a 3D pyramid with its tip sliced off. Credit: Spencer Hallyburton, Duke University

Researchers at Duke University have demonstrated the first attack strategy that can fool industry-standard autonomous vehicle sensors into believing nearby objects are closer (or further) than they appear without

being detected.

The research suggests that adding optical 3D capabilities or the ability to share data with nearby cars may be necessary to fully protect autonomous cars from attacks.

The results will be presented Aug. 10–12 at the 2022 USENIX Security Symposium, a top venue in the field.

One of the biggest challenges researchers developing autonomous driving systems have to worry about is protecting against attacks. A common strategy to secure safety is to check data from separate instruments against one another to make sure their measurements make sense together.

The most common locating technology used by today's autonomous car companies combines 2D data from cameras and 3D data from LiDAR, which is essentially laser-based radar. This combination has proven very robust against a wide range of attacks that attempt to fool the visual system into seeing the world incorrectly.

At least, until now.

"Our goal is to understand the limitations of existing systems so that we can protect against attacks," said Miroslav Pajic, the Dickinson Family Associate Professor of Electrical and Computer Engineering at Duke. "This research shows how adding just a few data points in the 3D point cloud, ahead or behind of where an object actually is, can confuse these systems into making dangerous decisions."

The new attack strategy works by shooting a laser gun into a car's LIDAR sensor to add false data points to its perception. If those data points are wildly out of place with what a car's camera is seeing,

previous research has shown that the system can recognize the attack. But the new research from Pajic and his colleagues shows that 3D LIDAR data points carefully placed within a certain area of a camera's 2D field of view can fool the system.

This vulnerable area stretches out in front of a camera's lens in the shape of a frustum—a 3D pyramid with its tip sliced off. In the case of a forward-facing camera mounted on a car, this means that a few data points placed in front of or behind another nearby car can shift the system's perception of it by several meters.

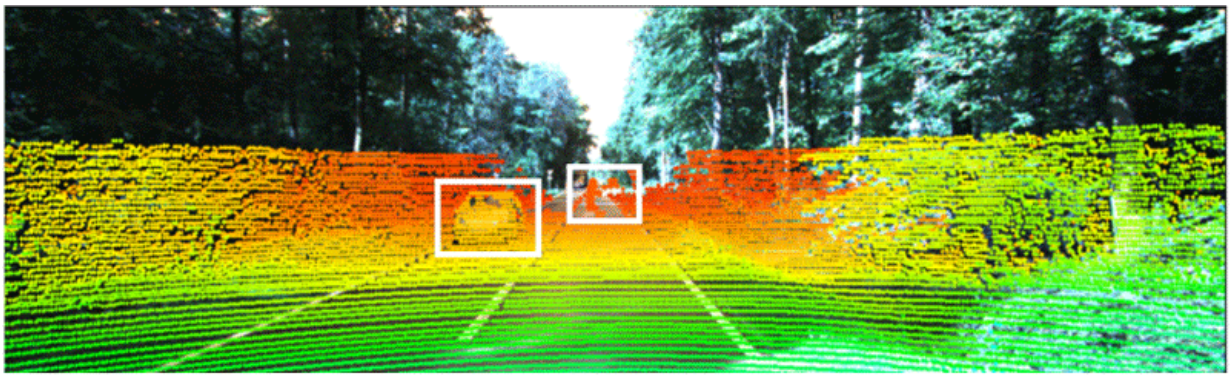

Researchers have shown that a popular method to secure LiDAR sensors against "naive attacks" is still vulnerable at longer distances and only works at short distances. Here, a LiDAR system is fooled into thinking a car is somewhere else until it becomes too late to avoid a sudden and drastic course correction. Credit: Spencer Hallyburton, Duke University

"This so-called frustum attack can fool adaptive cruise control into thinking a vehicle is slowing down or speeding up," Pajic said. "And by the time the system can figure out there's an issue, there will be no way to avoid hitting the car without aggressive maneuvers that could create

even more problems."

According to Pajic, there is not much risk of somebody taking the time to set up lasers on a car or roadside object to trick individual vehicles passing by on the highway. That risk increases tremendously, however, in military situations where single vehicles can be very high-value targets. And if hackers could find a way of creating these false data points virtually instead of requiring physical lasers, many vehicles could be attacked at once.

The path to protecting against these attacks, Pajic says, is added redundancy. For example, if cars had "stereo cameras" with overlapping fields of view, they could better estimate distances and notice LIDAR data that does not match their perception.

"Stereo cameras are more likely to be a reliable consistency check, though no software has been sufficiently validated for how to determine if the LIDAR/stereo camera data are consistent or what to do if it is found they are inconsistent," said Spencer Hallyburton, a Ph.D. candidate in Pajic's Cyber-Physical Systems Lab and the lead author of the study. "Also, perfectly securing the entire vehicle would require multiple sets of stereo cameras around its entire body to provide 100% coverage."

Another option, Pajic suggests, is to develop systems in which cars within close proximity to one another share some of their data. Physical attacks are not likely to be able to affect many cars at once, and because different brands of cars may have different operating systems, a cyberattack is not likely to be able to hit all cars with a single blow.

"With all of the work that is going on in this field, we will be able to build systems that you can trust your life with," Pajic said. "It might take 10+ years, but I'm confident that we will get there."

More information: R. Spencer Hallyburton et al, "Security Analysis of Camera-LiDAR Fusion Against Black-Box Attacks on Autonomous Vehicles," 31st USENIX Security Symposium, Aug. 10-12, 2022

Provided by Duke University