Accelerating self-driving car innovation

February 27 2018, by Bradley Allf

On a Monday morning in January of this year a man got into his Tesla sedan, pulled onto a highway outside of Los Angeles, and engaged the vehicle's semi-autonomous "Autopilot" mode. The car juggled inputs from 8 cameras, 12 ultrasonic sensors and one radar to navigate the highway, avoiding other vehicles and staying within the dotted lines. Then, it plowed straight into the back of a stopped firetruck.

Autonomous vehicle technology is often praised for its capacity to eventually make the roads safer—and for good reason. But situations like that described above demonstrate that even today's semi-autonomous vehicles have some serious hitches to work out. Fortunately, researchers like Xin Li are working hard to fix some of the most fundamental problems currently faced by self-driving car technology in order to eventually reduce motor vehicle deaths.

"A few years ago I started to talk to a few of the auto companies in the US about the problems they were facing with self-driving cars," says Li, a professor in the Department of Electrical & Computer Engineering at both Duke and Duke Kunshan University. Because of his dual appointment, Li spends a lot of time in China, where he has continued the self-driving car conversation with Chinese automakers.

Li says there are two main issues currently standing in the way of fully autonomous driving technology: the verification problem and the integration problem.

The verification problem boils down to the following dilemma: for

something to be proven safe, you must first learn what circumstances cause it to fail. "If you are asking your customer to drive a self-driving car on the road, the failure rate must be very, very small," says Li. Figuring out how to calculate that rate—and, by extension, how then to correct those failures—has become a central research question for Li.

So how do you calculate the failure rate? One solution is to simply let the self-driving car navigate an area until it makes a mistake—not stopping at a stop sign, for instance. But most of these failures happen under very rare circumstances called "corner cases" (think, a boulder tumbling onto a mountain road), meaning this approach would take a prohibitive amount of time.

"It's difficult to physically observe all possible scenarios over all corner cases, so we thought, 'why don't we use a computer and ask the computer to synthetically generate these possible cases?'" says Li. Instead of waiting around for a car on the road to uncover every bizarre confluence of weather and erratic drivers that causes the system to misread a stoplight, Li's team's unique approach creates these scenarios using intelligent computer programs. These programs use statistical models and algorithms to generate synthetic situations for the autonomous system to navigate.

In one recent paper published in the November 2017 issue of Computer-Aided Design, his team showed that high temperature and aged circuitry can degrade the images gathered by a self-driving car's cameras. The team used a machine learning technique called Cycle-Consistent Generative Adversarial Networks to generate synthetic image data as if it were gathered at a high temperature or using old circuits. When they fed those degraded images into the autonomous system, it struggled to properly identify stop signs. By creating synthetic failure cases like this, Li's research allows automakers to more accurately verify how often and, more importantly, under which scenarios a system might fail. Li

then works with his collaborators in industry to implement this new technology into hardware and software that makes autonomous vehicles safer.

Unfortunately, the more components added to an autonomous vehicle, the higher the chances of something going wrong. This is the paradigm of the second hurdle for autonomous vehicles: the integration problem. Self-driving cars must seamlessly combine all kinds of complicated systems—from collision avoidance modules to stoplight detection—all while navigating an environment at 60 miles per hour. Figuring out how to ensure that components work together quickly and effectively will form the next phase in Li's research.

"Honestly, it's a very difficult problem and we don't have a good solution yet," says Li, though he's hopeful that developing a more centralized computing system for the car could go a long way towards integrating the vehicle's functionality.

Li says there is also a more fundamental issue facing the field of self-driving car development—education. While conventional vehicle design is the domain of mechanical engineers, building autonomous vehicles will require broad expertise in artificial intelligence, software design and system engineering. To face the problems of tomorrow will take a novel interdisciplinary approach. "I think it's very important for an educational institution like Duke to train students with the right background so they can engage in these emerging areas when they graduate," says Li.

Doing foundational research on self-driving cars has unbelievably high stakes. "When industry adopts my solution, I do have the responsibility to make sure this is going to work…If there's an error then people could get killed," says Li. At the same time, working on some of the biggest questions in self-driving car development is an incredible opportunity. "Right now autonomous driving is an emerging field with lots of open

problems. This makes it a very exciting time for academic people like me," says Li.

Provided by Duke University