

Researchers develop a new way to test machine learning algorithms that control selfdriving cars

March 28 2019



In this example, a perception algorithm misclassifies the cyclist as a pedestrian Credit: Anand Balakrishnan



It's a big question for many people in traffic-dense cities like Los Angeles: When will self-driving cars arrive? But following a series of high-profile accidents in the United States, safety issues could bring the autonomous dream to a screeching halt.

At USC, researchers have published a new study that tackles a longstanding problem for autonomous vehicle developers: testing the system's perception algorithms, which allow the car to "understand" what it "sees."

Working with researchers from Arizona State University, the team's new mathematical method is able to identify anomalies or bugs in the system before the car hits the road.

Perception algorithms are based on convolutional neural networks, powered by machine learning, a type of deep learning. These algorithms are notoriously difficult to test, as we don't fully understand how they make their predictions. This can lead to devastating consequences in safety-critical systems like autonomous vehicles.

"Making perception algorithms robust is one of the foremost challenges for autonomous systems," said the study's lead author Anand Balakrishnan, a USC computer science Ph.D. student.

"Using this method, developers can narrow in on errors in the perception algorithms much faster and use this information to further train the system. The same way cars have to go through crash tests to ensure safety, this method offers a pre-emptive test to catch errors in autonomous systems."

The paper, titled Specifying and Evaluating Quality Metrics for Visionbased Perception Systems, was presented at the Design, Automation and Test in Europe conference in Italy, Mar. 28.



Learning about the world

Typically autonomous vehicles "learn" about the world via machine learning systems, which are fed huge datasets of road images before they can identify objects on their own.

But the system can go wrong. In the case of a fatal accident between a self-driving car and a pedestrian in Arizona last March, the software classified the pedestrian as a "false positive" and decided it didn't need to stop.

"We thought, clearly there is some issue with the way this perception <u>algorithm</u> has been trained," said study co-author Jyo Deshmukh, a USC computer science professor and former research and development engineer for Toyota, specializing in autonomous vehicle safety.

"When a human being perceives a video, there are certain assumptions about persistence that we implicitly use: if we see a car within a video frame, we expect to see a car at a nearby location in the next video frame. This is one of several 'sanity conditions' that we want the perception algorithm to satisfy before deployment."

For example, an object cannot appear and disappear from one frame to the next. If it does, it violates a "sanity condition," or basic law of physics, which suggests there is a bug in the perception system.

Deshmukh and his Ph.D. student Balakrishnan, along with USC Ph.D. student Xin Qin and master's student Aniruddh Puranic, teamed up with three Arizona State University researchers to investigate the problem.

No room for error



The team formulated a new mathematical logic, called Timed Quality Temporal Logic, and used it to test two popular machine-learning tools—Squeeze Det and YOLO—using raw video datasets of driving scenes.

The logic successfully honed in on instances of the machine learning tools violating "sanity conditions" across multiple frames in the video. Most commonly, the <u>machine learning</u> systems failed to detect an object or misclassified an object.

For instance, in one example, the system failed to recognize a cyclist from the back, when the bike's tire looked like a thin vertical line. Instead, it misclassified the cyclist as a pedestrian. In this case, the system might fail to correctly anticipate the cyclist's next move, which could lead to an accident.

Phantom objects—where the system perceives an object when there is none—were also common. This could cause the car to mistakenly slam on the breaks—another potentially dangerous move.

The team's method could be used to identify anomalies or bugs in the perception algorithm before deployment on the road and allows developer to pinpoint specific problems.

The idea is to catch issues with perception algorithm in virtual testing, making the algorithms safer and more reliable. Crucially, because the method relies on a library of "sanity conditions," there is no need for humans to label objects in the test dataset—a time-consuming and often-flawed process.

In the future, the team hopes to incorporate the logic to retrain the <u>perception</u> algorithms when it finds an error. It could also be extended to real-time use, while the car is driving, as a real-time safety monitor.



Provided by University of Southern California

Citation: Researchers develop a new way to test machine learning algorithms that control selfdriving cars (2019, March 28) retrieved 2 May 2024 from <u>https://techxplore.com/news/2019-03-machine-algorithms-self-driving-cars.html</u>

This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.