


New method makes more data available for training self-driving cars

June 17 2020



A technique called scene flow can be used to predict the future position of a cyclist by comparing the current lidar point cloud of a street scene, in green, with the point cloud from the previous time step in the sequence, shown in red. Carnegie Mellon University researchers have developed a method that increases the amount of data available for training such systems. Credit: Carnegie Mellon

For safety's sake, a self-driving car must accurately track the movement of pedestrians, bicycles and other vehicles around it. Training those tracking systems may now be more effective thanks to a new method developed at Carnegie Mellon University.

Generally speaking, the more road and [traffic data](#) available for [training](#) tracking systems, the better the results. And the CMU researchers have found a way to unlock a mountain of autonomous driving data for this purpose.

"Our method is much more robust than previous methods because we can train on much larger datasets," said Himangi Mittal, a research intern working with David Held, assistant professor in CMU's Robotics Institute.

Most autonomous vehicles navigate primarily based on a sensor called a lidar, a laser device that generates 3-D information about the world surrounding the car. This 3-D information isn't images, but a cloud of points. One way the vehicle makes sense of this data is by using a technique known as [scene](#) flow. This involves calculating the speed and trajectory of each 3-D point. Groups of points moving together are interpreted via scene flow as vehicles, pedestrians or other moving objects.

In the past, state-of-the-art methods for training such a system have required the use of labeled datasets—[sensor data](#) that has been annotated to track each 3-D point over time. Manually labeling these datasets is laborious and expensive, so, not surprisingly, little labeled data exists. As a result, scene flow training is instead often performed with simulated

data, which is less effective, and then fine-tuned with the small amount of labeled real-world data that exists.

Mittal, Held and robotics Ph.D. student Brian Okorn took a different approach, using unlabeled data to perform scene flow training. Because unlabeled data is relatively easy to generate by mounting a lidar on a car and driving around, there's no shortage of it.

The key to their approach was to develop a way for the system to detect its own errors in scene flow. At each instant, the system tries to predict where each 3-D point is going and how fast it's moving. In the next instant, it measures the distance between the point's predicted location and the actual location of the point nearest that predicted location. This distance forms one type of error to be minimized.

The system then reverses the process, starting with the predicted point location and working backward to map back to where the point originated. At this point, it measures the distance between the predicted position and the actual origination point, and the resulting distance forms the second type of error.

The system then works to correct those errors.

"It turns out that to eliminate both of those errors, the system actually needs to learn to do the right thing, without ever being told what the right thing is," Held said.

As convoluted as that might sound, Okorn found that it worked well. The researchers calculated that scene flow accuracy using a training set of synthetic data was only 25%. When the synthetic data was fine-tuned with a small amount of real-world labeled data, the accuracy increased to 31%. When they added a large amount of unlabeled data to train the system using their approach, scene flow accuracy jumped to 46%.

The research team presented their method at the Computer Vision and Pattern Recognition (CVPR) conference, which was held virtually June 14-19. The CMU Argo AI Center for Autonomous Vehicle Research supported this research, with additional support from a NASA Space Technology Research Fellowship.

Provided by Carnegie Mellon University

Citation: New method makes more data available for training self-driving cars (2020, June 17)
retrieved 27 April 2024 from

<https://techxplore.com/news/2020-06-method-self-driving-cars.html>

<p>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.</p>
--