

Teaching cars to drive with foresight

October 28 2019



Single LiDAR scan (left), the superimposed data (right) with descriptions (colors) provided by a human observer and the result of the software (center). Credit: AG Computer Vision der Universität Bonn

Good drivers anticipate dangerous situations and adjust their driving before things get dicey. Researchers at the University of Bonn now also want to teach this skill to self-driving cars. They will present a corresponding algorithm at the International Conference on Computer Vision which is held at Friday, November 1st, in Seoul. They will also present a data set that they used to train and test their approach. It will make it much easier to develop and improve such processes in the future.

An empty street, a row of parked cars at the side: nothing to indicate that you should be careful. But wait: Isn't there a side street up ahead, half covered by the parked cars? Maybe I better take my foot off the gas—who knows if someone's coming from the side. We constantly encounter situations like these when driving. Interpreting them correctly



and drawing the right conclusions requires a lot of experience. In contrast, self-driving cars sometimes behave like a learner driver in his first lesson. "Our goal is to teach them a more anticipatory driving style," explains computer scientist Prof. Dr. Jürgen Gall. "This would then allow them to react much more quickly to dangerous situations."

Gall chairs the "Computer Vision" working group at the University of Bonn, which, in cooperation with his university colleagues from the Institute of Photogrammetry and the "Autonomous Intelligent Systems" working group, is researching a solution to this problem. The scientists now present a first step on the way to this goal at the leading symposium of Gall's discipline, the International Conference on Computer Vision in Seoul. "We have refined an algorithm that completes and interprets socalled LiDAR data," he explains. "This allows the car to anticipate potential hazards at an early stage."

Problem: too little data

LiDAR is a rotating laser that is mounted on the roof of most selfdriving cars. The laser beam is reflected by the surroundings. The LiDAR system measures when the reflected light falls on the sensor and uses this time to calculate the distance. "The system detects the distance to around 120,000 points around the vehicle per revolution," says Gall.

The problem with this: The measuring points become "dilute" as the distance increases—the gap between them widens. This is like painting a face on a balloon: When you inflate it, the eyes move further and further apart. Even for a human being it is therefore almost impossible to obtain a correct understanding of the surroundings from a single LiDAR scan (i.e. the distance measurements of a single revolution). "A few years ago, the University of Karlsruhe (KIT) recorded large amounts of LiDAR data, a total of 43,000 scans," explains Dr. Jens Behley of the Institute of Photogrammetry. "We have now taken sequences from several dozen



scans and superimposed them." The data obtained in this way also contain points that the sensor had only recorded when the car had already driven a few dozen yards further down the road. Put simply, they show not only the present, but also the future.

"These superimposed point clouds contain <u>important information</u> such as the geometry of the scene and the spatial dimensions of the objects it contains, which are not available in a single scan," emphasizes Martin Garbade, who is currently doing his doctorate at the Institute of Computer Science. "Additionally, we have labeled every single point in them, for example: There's a sidewalk, there's a pedestrian and back there's a motorcyclist." The scientists fed their software with a data pair: a single LiDAR scan as input and the associated overlay data including semantic information as desired output. They repeated this process for several thousands of such pairs.

"During this training phase, the algorithm learned to complete and interpret individual scans," explains Prof. Gall. "This meant that it could plausibly add missing measurements and interpret what was seen in the scans." The scene completion already works relatively well: The process can complete about half of the missing data correctly. The semantic interpretation, i.e. deducing which objects are hidden behind the measuring points, does not work quite as well: Here, the computer achieves a maximum accuracy of 18 percent.

However, the scientists consider this branch of research to still be in its infancy. "Until now, there has simply been a lack of extensive data sets with which to train corresponding artificial intelligence methods," stresses Gall. "We are closing a gap here with our work. I am optimistic that we will be able to significantly increase the accuracy rate in semantic interpretation in the coming years." He considers 50 percent to be quite realistic, which could have a huge influence on the quality of autonomous driving.



More information: Project website: semantic-kitti.org

Provided by University of Bonn

Citation: Teaching cars to drive with foresight (2019, October 28) retrieved 10 May 2024 from <u>https://techxplore.com/news/2019-10-cars-foresight.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.