

Technology helps self-driving cars learn from their own memories

21 June 2022, by Tom Fleischman



Credit: Pixabay/CC0 Public Domain

An autonomous vehicle is able to navigate city streets and other less-busy environments by recognizing pedestrians, other vehicles and potential obstacles through artificial intelligence. This is achieved with the help of artificial neural networks, which are trained to "see" the car's surroundings, mimicking the human visual perception system.

But unlike humans, cars using [artificial neural networks](#) have no memory of the past and are in a constant state of seeing the world for the first time—no matter how many times they've driven down a particular road before. This is particularly problematic in adverse weather conditions, when the car cannot safely rely on its sensors.

Researchers at the Cornell Ann S. Bowers College of Computing and Information Science and the College of Engineering have produced three concurrent research papers with the goal of overcoming this limitation by providing the car with the ability to create "memories" of previous experiences and use them in future navigation.

Doctoral student Yurong You is lead author of "HINDSIGHT is 20/20: Leveraging Past Traversals to Aid 3D Perception," which You presented virtually in April at ICLR 2022, the International Conference on Learning Representations. "Learning representations" includes deep learning, a kind of machine learning.

"The fundamental question is, can we learn from repeated traversals?" said senior author Kilian Weinberger, professor of computer science in Cornell Bowers CIS. "For example, a car may mistake a weirdly shaped tree for a pedestrian the first time its laser scanner perceives it from a distance, but once it is close enough, the object category will become clear. So the second time you drive past the very same tree, even in fog or snow, you would hope that the car has now learned to recognize it correctly."

"In reality, you rarely drive a route for the very first time," said co-author Katie Luo, a doctoral student in the research group. "Either you yourself or someone else has driven it before recently, so it seems only natural to collect that experience and utilize it."

Spearheaded by doctoral student Carlos Diaz-Ruiz, the group compiled a dataset by driving a car equipped with LiDAR (Light Detection and Ranging) sensors repeatedly along a 15-kilometer loop in and around Ithaca, 40 times over an 18-month period. The traversals capture varying environments (highway, urban, campus), weather conditions (sunny, rainy, snowy) and times of day.

This resulting dataset—which the group refers to as Ithaca365, and which is the subject of one of the other two papers—has more than 600,000 scenes.

"It deliberately exposes one of the key challenges in self-driving cars: poor [weather conditions](#)," said Diaz-Ruiz, a co-author of the Ithaca365 paper. "If the street is covered by snow, humans can rely on

memories, but without memories a neural network is heavily on costly human annotated data) and make heavily disadvantaged."

HINDSIGHT is an approach that uses [neural networks](#) to compute descriptors of objects as the car passes them. It then compresses these descriptions, which the group has dubbed SQuaSH (Spatial-Quantized Sparse History) features, and stores them on a virtual map, similar to a "[memory](#)" stored in a [human brain](#).

The next time the self-driving car traverses the same location, it can query the local SQuaSH database of every LiDAR point along the route and "remember" what it learned last time. The database is continuously updated and shared across vehicles, thus enriching the information available to perform recognition.

"This information can be added as features to any LiDAR-based 3D object detector;" You said. "Both the detector and the SQuaSH representation can be trained jointly without any additional supervision, or human annotation, which is time- and labor-intensive."

While HINDSIGHT still assumes that the artificial neural network is already trained to detect objects and augments it with the capability to create memories, MODEST (Mobile Object Detection with Ephemerality and Self-Training)—the subject of the third publication—goes even further.

Here, the authors let the car learn the entire perception pipeline from scratch. Initially the artificial neural network in the vehicle has never been exposed to any objects or streets at all. Through multiple traversals of the same route, it can learn what parts of the environment are stationary and which are moving objects. Slowly it teaches itself what constitutes other traffic participants and what is safe to ignore.

The algorithm can then detect these objects reliably—even on roads that were not part of the initial repeated traversals.

The researchers hope that both approaches could drastically reduce the development cost of [autonomous vehicles](#) (which currently still relies

heavily on costly human annotated data) and make such vehicles more efficient by learning to navigate the locations in which they are used the most.

Both Ithaca365 and MODEST will be presented at the Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR 2022), to be held June 19-24 in New Orleans.

Other contributors include Mark Campbell, the John A. Mellowes '60 Professor in Mechanical Engineering in the Sibley School of Mechanical and Aerospace Engineering, assistant professors Bharath Hariharan and Wen Sun, from computer science at Bowers CIS; former postdoctoral researcher Wei-Lun Chao, now an assistant professor of computer science and engineering at Ohio State; and doctoral students Cheng Perng Phoo, Xiangyu Chen and Junan Chen.

More information: Conference: cvpr2022.thecvf.com/

Provided by Cornell University

APA citation: Technology helps self-driving cars learn from their own memories (2022, June 21) retrieved 6 July 2022 from <https://techxplore.com/news/2022-06-technology-self-driving-cars-memories.html>

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.