


Bringing human-like reasoning to driverless car navigation

May 23 2019, by Rob Matheson



To bring more human-like reasoning to autonomous vehicle navigation, MIT researchers have created a system that enables driverless cars to check a simple map and use visual data to follow routes in new, complex environments. Credit: Chelsea Turner

With aims of bringing more human-like reasoning to autonomous

vehicles, MIT researchers have created a system that uses only simple maps and visual data to enable driverless cars to navigate routes in new, complex environments.

Human drivers are exceptionally good at navigating roads they haven't driven on before, using observation and simple tools. We simply match what we see around us to what we see on our GPS devices to determine where we are and where we need to go. Driverless cars, however, struggle with this basic reasoning. In every new area, the cars must first map and analyze all the new roads, which is very time consuming. The systems also rely on complex maps—usually generated by 3-D scans—which are computationally intensive to generate and process on the fly.

In a paper being presented at this week's International Conference on Robotics and Automation, MIT researchers describe an autonomous control [system](#) that "learns" the steering patterns of human drivers as they navigate roads in a small area, using only data from video camera feeds and a simple GPS-like map. Then, the trained system can control a driverless car along a planned route in a brand-new area, by imitating the human driver.

Similarly to human drivers, the system also detects any mismatches between its map and features of the [road](#). This helps the system determine if its position, sensors, or mapping are incorrect, in order to correct the car's course.

To train the system initially, a human operator controlled a driverless Toyota Prius—equipped with several cameras and a basic GPS navigation system—collecting data from local suburban streets including various road structures and obstacles. When deployed autonomously, the system successfully navigated the car along a preplanned path in a different forested area, designated for autonomous vehicle tests.

"With our system, you don't need to train on every road beforehand," says first author Alexander Amini, an MIT graduate student. "You can download a new map for the car to navigate through roads it has never seen before."

"Our objective is to achieve autonomous navigation that is robust for driving in new environments," adds co-author Daniela Rus, director of the Computer Science and Artificial Intelligence Laboratory (CSAIL) and the Andrew and Erna Viterbi Professor of Electrical Engineering and Computer Science. "For example, if we train an autonomous vehicle to drive in an urban setting such as the streets of Cambridge, the system should also be able to drive smoothly in the woods, even if that is an environment it has never seen before."

Joining Rus and Amini on the paper are Guy Rosman, a researcher at the Toyota Research Institute, and Sertac Karaman, an associate professor of aeronautics and astronautics at MIT.

Point-to-point navigation

Traditional navigation systems process data from sensors through multiple modules customized for tasks such as localization, mapping, object detection, motion planning, and steering control. For years, Rus's group has been developing "end-to-end" navigation systems, which process inputted sensory data and output steering commands, without a need for any specialized modules.

Until now, however, these models were strictly designed to safely follow the road, without any real destination in mind. In the new paper, the researchers advanced their end-to-end system to drive from goal to destination, in a previously unseen environment. To do so, the researchers trained their system to predict a full probability distribution over all possible steering commands at any given instant while driving.

The system uses a machine learning model called a convolutional neural network (CNN), commonly used for image recognition. During training, the system watches and learns how to steer from a human driver. The CNN correlates steering wheel rotations to road curvatures it observes through cameras and an inputted map. Eventually, it learns the most likely steering command for various driving situations, such as straight roads, four-way or T-shaped intersections, forks, and rotaries.

"Initially, at a T-shaped intersection, there are many different directions the car could turn," Rus says. "The model starts by thinking about all those directions, but as it sees more and more data about what people do, it will see that some people turn left and some turn right, but nobody goes straight. Straight ahead is ruled out as a possible direction, and the model learns that, at T-shaped intersections, it can only move left or right."

What does the map say?

In testing, the researchers input the system with a map with a randomly chosen route. When driving, the system extracts visual features from the camera, which enables it to predict road structures. For instance, it identifies a distant stop sign or line breaks on the side of the road as signs of an upcoming intersection. At each moment, it uses its predicted probability distribution of steering commands to choose the most likely one to follow its route.

Importantly, the researchers say, the system uses maps that are easy to store and process. Autonomous control systems typically use LIDAR scans to create massive, complex maps that take roughly 4,000 gigabytes (4 terabytes) of data to store just the city of San Francisco. For every new destination, the car must create new maps, which amounts to tons of data processing. Maps used by the researchers' system, however, captures the entire world using just 40 gigabytes of data.

During autonomous driving, the system also continuously matches its [visual data](#) to the map data and notes any mismatches. Doing so helps the autonomous vehicle better determine where it is located on the road. And it ensures the car stays on the safest path if it's being fed contradictory input information: If, say, the car is cruising on a straight road with no turns, and the GPS indicates the car must turn right, the car will know to keep driving straight or to stop.

"In the real world, sensors do fail," Amini says. "We want to make sure that the system is robust to different failures of different sensors by building a system that can accept these noisy inputs and still navigate and localize itself correctly on the road."

More information: Variational End-to-End Navigation and Localization, arXiv:1811.10119 [cs.LG] arxiv.org/abs/1811.10119

Provided by Massachusetts Institute of Technology

Citation: Bringing human-like reasoning to driverless car navigation (2019, May 23) retrieved 23 April 2024 from <https://techxplore.com/news/2019-05-human-like-driverless-car.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>
--