


# System trains driverless cars in simulation before they hit the road

24 March 2020, by Rob Matheson



A simulation system invented at MIT to train driverless cars creates a photorealistic world with infinite steering possibilities, helping the cars learn to navigate a host of worse-case scenarios before cruising down real streets. Credit: Massachusetts Institute of Technology

A simulation system invented at MIT to train driverless cars creates a photorealistic world with infinite steering possibilities, helping the cars learn to navigate a host of worse-case scenarios before cruising down real streets.

Control systems, or "controllers," for [autonomous vehicles](#) largely rely on real-world datasets of driving trajectories from human drivers. From these data, they learn how to emulate safe steering controls in a variety of situations. But real-world data from hazardous "edge cases," such as nearly crashing or being forced off the road or into other lanes, are—fortunately—rare.

Some [computer programs](#), called "simulation engines," aim to imitate these situations by rendering detailed virtual roads to help train the controllers to recover. But the learned control from simulation has never been shown to transfer to reality on a full-scale vehicle.

The MIT researchers tackle the problem with their photorealistic simulator, called Virtual Image Synthesis and Transformation for Autonomy (VISTA). It uses only a small dataset, captured by humans driving on a road, to synthesize a practically infinite number of new viewpoints from trajectories that the vehicle could take in the real world. The controller is rewarded for the distance it travels without crashing, so it must learn by itself how to reach a destination safely. In doing so, the vehicle learns to safely navigate any situation it encounters, including regaining control after swerving between lanes or recovering from near-crashes.

In tests, a controller trained within the VISTA simulator safely was able to be safely deployed onto a full-scale driverless car and to navigate through previously unseen streets. In positioning the car at off-road orientations that mimicked various near-crash situations, the controller was also able to successfully recover the car back into a safe driving trajectory within a few seconds. A paper describing the system has been published in *IEEE Robotics and Automation Letters* and will be presented at the upcoming ICRA conference in May.

"It's tough to collect data in these edge cases that humans don't experience on the road," says first author Alexander Amini, a Ph.D. student in the Computer Science and Artificial Intelligence Laboratory (CSAIL). "In our simulation, however, [control systems](#) can experience those situations, learn for themselves to recover from them, and remain robust when deployed onto vehicles in the real world."

The work was done in collaboration with the Toyota Research Institute. Joining Amini on the paper are Igor Gilitschenski, a postdoc in CSAIL; Jacob Phillips, Julia Moseyko, and Rohan Banerjee, all undergraduates in CSAIL and the Department of Electrical Engineering and Computer Science;

Sertac Karaman, an associate professor of aeronautics and astronautics; and Daniela Rus, director of CSAIL and the Andrew and Erna Viterbi Professor of Electrical Engineering and Computer Science.

### Data-driven simulation

Historically, building simulation engines for training and testing autonomous vehicles has been largely a manual task. Companies and universities often employ teams of artists and engineers to sketch virtual environments, with accurate road markings, lanes, and even detailed leaves on trees. Some engines may also incorporate the physics of a car's interaction with its environment, based on complex mathematical models.

But since there are so many different things to consider in complex real-world environments, it's practically impossible to incorporate everything into the simulator. For that reason, there's usually a mismatch between what controllers learn in simulation and how they operate in the real world.

Instead, the MIT researchers created what they call a "data-driven" simulation engine that synthesizes, from real data, new trajectories consistent with road appearance, as well as the distance and motion of all objects in the scene.

They first collect video data from a human driving down a few roads and feed that into the engine. For each frame, the engine projects every pixel into a type of 3-D point cloud. Then, they place a virtual vehicle inside that world. When the vehicle makes a steering command, the engine synthesizes a new trajectory through the point cloud, based on the steering curve and the vehicle's orientation and velocity.

Then, the engine uses that new trajectory to render a photorealistic scene. To do so, it uses a convolutional neural network—commonly used for image-processing tasks—to estimate a depth map, which contains information relating to the distance of objects from the controller's viewpoint. It then combines the depth map with a technique that estimates the camera's orientation within a 3-D scene. That all helps pinpoint the vehicle's location

and relative distance from everything within the virtual simulator.

Based on that information, it reorients the original pixels to recreate a 3-D representation of the world from the vehicle's new viewpoint. It also tracks the motion of the pixels to capture the movement of the cars and people, and other moving objects, in the scene. "This is equivalent to providing the vehicle with an infinite number of possible trajectories," Rus says. "Because when we collect physical data, we get data from the specific trajectory the car will follow. But we can modify that trajectory to cover all possible ways of and environments of driving. That's really powerful."

### Reinforcement learning from scratch

Traditionally, researchers have been training autonomous vehicles by either following human defined rules of driving or by trying to imitate human drivers. But the researchers make their controller learn entirely from scratch under an "end-to-end" framework, meaning it takes as input only raw sensor data—such as visual observations of the road—and, from that data, predicts steering commands at outputs.

"We basically say, 'Here's an environment. You can do whatever you want. Just don't crash into vehicles, and stay inside the lanes,'" Amini says.

This requires "reinforcement learning" (RL), a trial-and-error machine-learning technique that provides feedback signals whenever the car makes an error. In the researchers' simulation engine, the controller begins by knowing nothing about how to drive, what a lane marker is, or even other vehicles look like, so it starts executing random steering angles. It gets a feedback signal only when it crashes. At that point, it gets teleported to a new simulated location and has to execute a better set of steering angles to avoid crashing again. Over 10 to 15 hours of training, it uses these sparse feedback signals to learn to travel greater and greater distances without crashing.

After successfully driving 10,000 kilometers in simulation, the authors apply that learned controller onto their full-scale autonomous vehicle in the real

world. The researchers say this is the first time a controller trained using end-to-end reinforcement learning in simulation has successfully been deployed onto a full-scale autonomous car. "That was surprising to us. Not only has the controller never been on a real car before, but it's also never even seen the roads before and has no prior knowledge on how humans drive," Amini says.

Forcing the [controller](#) to run through all types of driving scenarios enabled it to regain control from disorienting positions—such as being half off the road or into another lane—and steer back into the correct lane within several seconds. "And other state-of-the-art controllers all tragically failed at that, because they never saw any data like this in training," Amini says.

Next, the researchers hope to simulate all types of road conditions from a single driving trajectory, such as night and day, and sunny and rainy weather. They also hope to simulate more complex interactions with other vehicles on the road. "What if other cars start moving and jump in front of the [vehicle](#)?" Rus says. "Those are complex, real-world interactions we want to start testing."

**More information:** Alexander Amini et al. Learning Robust Control Policies for End-to-End Autonomous Driving From Data-Driven Simulation, *IEEE Robotics and Automation Letters* (2020). DOI: [10.1109/LRA.2020.2966414](https://doi.org/10.1109/LRA.2020.2966414)

Provided by Massachusetts Institute of Technology

APA citation: System trains driverless cars in simulation before they hit the road (2020, March 24) retrieved 9 April 2020 from <https://techxplore.com/news/2020-03-driverless-cars-simulation-road.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.*