

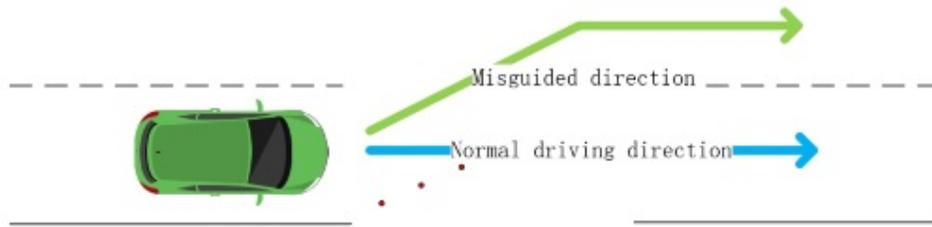
TESLA AUTO-PILOT CAN EASILY KILL-YOU-ON-COMMAND

Researchers have tried a simple trick that might make a Tesla to automatically drive into approaching traffic under certain conditions. The proof-of-concept manipulate works not by hacking into the car's onboard computing system but by working with small, unnoticeable stickers that trick the Enhanced Autopilot of a Model S 75 into detecting and subsequently following a change in the current lane.

Tesla's Enhanced Autopilot assists a range of features, including lane-centering, self-parking, and the ability to systematically change lanes with the driver's confirmation. The feature is now mostly called "Autopilot" after Tesla rearranged the Autopilot price structure. It primarily depends on cameras, ultrasonic sensors, and radar to gather information about its environments, including close-by barriers, landscapes, and route changes. It then submits the data into onboard computers that use machine understanding to make judgments in real time about the most effective way to take action.

A team from Tencent's Keen Security Lab recently reverse-engineered several of Tesla's automated processes to see how they responded to changes in the environmental variables. One of the most impressive breakthroughs was a technique to cause Autopilot to steer into oncoming traffic. The attack worked by carefully attaching 3 stickers to the road surface. The stickers were nearly invisible to drivers, but machine-learning algorithms used by the Autopilot identified them as a line that suggested the lane was switching to the left. As a result, Autopilot steered in that direction.

In a [detailed, 37-page report](#), the researchers wrote:



Fake lane mode in physical world



In-car perspective when testing, the red circle marks, the interference markings are marked with red circles

Tesla autopilot module's lane recognition function has excellent robustness in an ordinary external environment (no intense light, snow, rain, dust and sand interference), but it still doesn't handle the situation correctly in our test scenario. This kind of attack is straightforward to deploy, and the materials are easy to obtain. As we talked in the previous introduction of Tesla's lane recognition function, Tesla uses a pure computer-vision solution for lane recognition, and we found in this attack experiment that the vehicle driving decision is only based on computer-vision lane recognition results. Our tests prove that this architecture has some security risks and reverse lane recognition is one of the necessary functions for autonomous driving in non-closed roads. In the scene we build, if the vehicle knows that the fake lane is pointing to the opposite lane, it should ignore this fake lane and then it could

avoid a traffic accident.

The investigators said autopilot adopts a function called `detect_and_track` to detect lanes and update an internal map that sends the latest information to the controller. The function 1st calls several CUDA kernels for different jobs.



Adding some patches around lane line in physical world, and there is only right lane in the central information display

The researchers noted that Autopilot uses a range of measures to avoid inaccurate detections. The procedures include the position of lane backgrounds, road shoulders, and the distance and size of various objects.

A separate section of the report showed how the analysts- exploiting a now-patched root-privileged access vulnerability in Autopilot ECU (or APE)- were capable of using a gamepad to remotely control a car. That susceptibility was recently fixed in Tesla's 2018.24 firmware release.

However, another section demonstrated how researchers could meddle with Tesla's auto wiper system to activate wipers when rain wasn't falling. Unlike conventional auto wiper systems- which use optical sensors to detect humidity — Tesla's system uses a set of cameras that feeds data into an artificial intelligence network to figure out when wipers should be activated. The researchers found that minor changes in an image can easily deceive artificial intelligence-based image recognition (for example, changes that cause an AI system to mistake a panda for a gibbon) — it

wasn't hard to trick Tesla's auto wiper feature into assuming that rain was falling even when it was not.

Until now, the researchers have only been able to trick auto wiper when they feed images directly into the system. Eventually, they said, it may be possible for attackers to display an "adversarial image" that's shown on road signs or other cars that do the same thing.

In an emailed statement, Tesla officials wrote:

"We developed our bug-bounty program in 2014 in order to engage with the most talented members of the security research community, with the goal of soliciting this exact type of feedback. While we always appreciate this group's work, the primary vulnerability addressed in this report was fixed by Tesla through a robust security update in 2017, followed by another comprehensive security update in 2018, both of which we released before this group reported this research to us. The rest of the findings are all based on scenarios in which the physical environment around the vehicle is artificially altered to make the automatic windshield wipers or Autopilot system behave differently, which is not a realistic concern given that a driver can easily override Autopilot at any time by using the steering wheel or brakes and should always be prepared to do so and can manually operate the windshield wiper settings at all times.

You May Like: [Uber makes confidential filing for long-awaited IPO](#)

Although this report isn't eligible for an award through our bug-bounty program, we know it took an extraordinary amount of time, effort, and skill, and we look forward to reviewing future reports from this group."

The ability to modify self-driving cars by altering the environment isn't new. In late 2017, researchers demonstrated how stickers tagged to road signs could cause similar problems. Currently, changes to physical settings are generally considered beyond the reach of intrusions against self-driving systems. The point of the research is that companies designing such systems potentially should take into account such maneuvers in scope.