

Who (or what) is behind the wheel? The regulatory challenges of driverless cars

March 7 2016, by Hussein Dia



The Google prototype driverless car faces regulatory hurdles before it can get on the road. Credit: Google

A [legal opinion](#) by the US National Highway Traffic Safety Administration ([NHTSA](#)) [set the internet alight](#) in February.

The US [road safety](#) federal regulator informed Google that the artificial intelligence (AI) software it uses to control its self-driving cars could

effectively be viewed as the "driver" for some (but not all) regulatory purposes.

The NHTSA's letter was in response to a request from Google seeking the NHTSA's interpretations of the US [Federal Motor Vehicle Safety Standards](#).

It was widely viewed in the media as a recognition from the Feds that Google's AI software, the self-driving system (SDS), is [legally the same as a human driver](#). The details of the letter, however, tell a very different story.

First, the letter strictly stated the term "could be" equivalent to a human driver, meaning this definition is yet to be settled.

The NHTSA's letter also suggested that suitable tests would need to be developed to allow the NHTSA to certify the SDS compliance with road safety legislation.

And therein lies the challenge. What procedure can be used to verify compliance? Should the AI self-driving software pass a [benchmark test](#), developed specifically for [autonomous vehicles](#), before it can be recognised as a legal driver? Who should develop such a test and what should it include?



- | | |
|---|--------|
| 1. Q: Is there a person in the blue region? | A: yes |
| 2. Q: Is there a unique person in the blue region?
(Label this person 1) | A: yes |
| 3. Q: Is person 1 carrying something? | A: yes |
| 4. Q: Is person 1 female? | A: yes |
| 5. Q: Is person 1 walking on a sidewalk? | A: yes |
| 6. Q: Is person 1 interacting with any other object? | A: no |
| ⋮ | |

Some of the questions posed in a Visual Turing Test. Credit: Wikimedia Commons, CC BY-SA

Driving the future

Make no mistake, car manufacturers and technology companies are working towards a vision of [fully autonomous vehicles](#), and that vision includes taking the human driver out of the loop. They have already made huge advancements in this space.

The self-driving software that has been developed, based on ["deep neural](#)

[networks](#)", includes millions of virtual neurons that mimic the brain. The [on-board computers](#) have impressive supercomputing power packed inside hardware the size of a lunchbox.

The neural nets do not include any explicit programming to detect objects in the world. Rather, they are [trained](#) to recognise and classify objects using millions of images and examples from data sets representing real-world driving situations.

But the driving task is much more complex than object detection, and detection is not the same as understanding. For example, if a human is driving down a suburban street and sees a soccer ball roll out in front of the car, the driver would probably stop immediately since a child might be close behind.

Even with advanced AI, would a self-driving vehicle know how to react? What about those situations where an accident is unavoidable? Should the car minimise the loss of life, even if it means sacrificing the occupants, or should it protect the occupants at all costs? Should it be given the choice to select between these extremes?

These are not routine instances. Therefore, lacking a large set of examples, they would be relatively resistant to deep learning training. How can such situations be included in a benchmark test?

Turing tests

The question of whether a machine could "think" has been an active area of research since the 1950s, when Alan Turing first proposed his eponymous test.

The basis of the Turing Test is that a human interrogator is asked to distinguish which of two chat-room participants is a computer, and

which is a real human. If the interrogator cannot distinguish computer from human, then the computer is considered to have passed the test.



The Daimler data set used to test driverless cars. Credit: NVIDIA

The Turing Test has [many limitations](#) and is now considered [obsolete](#).


But a group of researchers have come up with a [similar test](#) based on machine vision, which is more suited to today's AI evaluations.

The researchers have proposed a framework for a [Visual Turing Test](#), in which computers would answer increasingly complex questions about a scene.

The [test](#) calls for human test-designers to develop a list of certain attributes that a picture might have. Images would first be hand-scored

by humans on given criteria, and a computer vision system would then be shown the same picture, without the "answers," to determine if it was able to pick out what the humans had spotted.

There are a [few vision benchmark data sets](#) used today to test the performance of neural nets in terms of detection and classification accuracy.



A different data set from Audi. Credit: NVIDIA

The [KITTI data set](#), for example, has been extensively used as a benchmark for self-driving object detection. [Baidu](#), the dominant search

engine company in China, and which is also a leader in self-driving software, is reported to have achieved the [best detection score of 90%](#) on this data set.

At the [Consumer Electronics Show](#) earlier this year, [NVIDIA](#) demonstrated the performance of its self-driving software on new data sets from Daimler and Audi.

The demonstrations showed advanced levels for [single and multi-class detection and segmentation](#), in which the software was able to extract more information from video images.

A modified Visual Turing Test can potentially be used to test the self-driving software if it's tailored to the multi-sensor inputs available to the car's computer, and is made relevant to the challenges of driving.

But putting together such a test would not be easy. This is further complicated by the [ethical questions](#) surrounding self-driving cars. There are also challenges in managing the interface between driver and computer when an acceptable response requires broader knowledge of the world.

Policy remains the last major hurdle to putting driverless cars on the road. Whether the final benchmark bears any resemblance to a Turing-like test, or something else we have not yet imagined, remains to be seen.

As with other fast-moving innovations, policymakers and regulators are struggling to keep pace. Regulators need to engage the public and create a testing and legal framework to verify compliance. They also need to ensure that it is flexible but robust.

Without this, a human will always need to be in the driver's seat and fully autonomous vehicles would go nowhere fast.

This article was originally published on [The Conversation](#). Read the [original article](#).

Source: The Conversation

Citation: Who (or what) is behind the wheel? The regulatory challenges of driverless cars (2016, March 7) retrieved 19 April 2024 from <https://phys.org/news/2016-03-wheel-regulatory-driverless-cars.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>
--