

How to make a driverless car 'see' the road ahead

March 16 2017, by Michael Milford And Jonathan Roberts



Self-driving cars need to 'see' what's going on around them. Credit: Intel/Mobileye

Microchip manufacturer Intel has invested heavily in the driverless car race with the latest US\$15 billion (A\$19.5bn) [purchase of Israeli tech company Mobileye](#).

[Mobileye](#) develops sensors and intelligence technology behind automated driver-assistance systems and many [self-driving cars](#). Its tech

enables a car to "see" and understand the world.

Other recent purchases include the [deep learning tech company Nervana](#), [microchip maker Movidius](#) and [automotive tech company Delphi](#). Intel is also working with the automotive companies [BMW](#) and Volkswagen to [begin trials later this year](#).

Intel is strategically putting together all the critical capabilities required to develop self-driving cars that can "see" and intelligently understand the world around us.

Seeing is safe

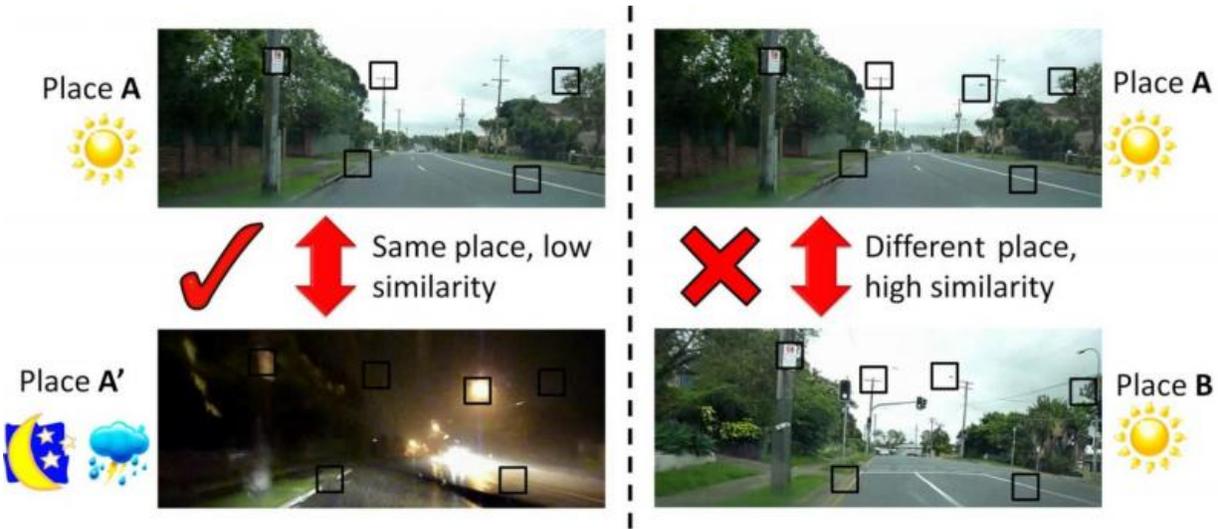
Most self-driving cars use a combination of sensing technologies. These include visual sensors, such as cameras, and range-to-object detecting sensors, such as lasers and radar.

Until the last decade or so, range-based sensors have dominated commercially developed systems in robotics and self-driving cars. These sensors reliably tell the distance to all objects surrounding the platform to ranges of 100 metres or more.

Lasers were typically only used for low-level, simple tasks such as obstacle avoidance, to make sure the system didn't hit anything.

Radar sensors have been used for at least a decade in advanced cruise control systems in some high-end cars, and [have recently transitioned to low-cost cars](#).

But range-only sensors have their limits. A long-range laser or radar scan can give you crude information about the posture of a pedestrian, but won't tell you the expression on that person's face. Range sensors are also poor at reading existing signage, since most signs are visual.



It is hard for a camera-based navigation system to recognise where it is in the world. The photos (left) are the same place, but the photos (right) are from two different places. Credit: Michael Milford, Author provided

In contrast, vision-based sensors like cameras provide a perceptually rich view of the world. They sense colour and fine appearance details, which a laser or radar unit simply does not pick up.

When it comes to driving, our environment has been designed and built with the assumption that a human driver will be able to see. So cars that see like us will fit most naturally into existing infrastructure and signage.

Seeing is difficult

However, cameras are very susceptible to changing environmental conditions. The simplest example encountered on roads is the day-night cycle.

Beyond the darkness, artificial lighting such as from oncoming headlights makes life difficult for the software trying to make sense of what is on the road ahead.

Other changes like weather, seasonal changes, fog, smoke and haze cause other problems. Snow banks can build up on streets, completely obscuring line markings and even signs.

Humans are often able to drive cautiously and work out what to do, but self-driving cars, designed to rely on rigid road rules, struggle.

The biggest challenge occurs when multiple changes happen at once, such as a tropical thunderstorm in the middle of the night. People deal with these conditions reasonably well, [although we have higher accident rate](#) in such conditions.

But no driverless car has yet demonstrated it can reliably drive under extreme conditions. Regularly occurring circumstances, such as moderate rain, is about as well as they can currently do.

Seeing can be taught

Many of the biggest players in the self-driving car world are developing deep learning systems that learn how to drive at a scale far beyond what a human driver does during their hundred hours or so of learner training.

This is where Mobileye comes in. These [deep learning systems](#) typically require huge amounts of labelled data.

Gathering the raw data is expensive but quite doable: just put sensors and computers on a large number of cars and drive millions of hours around road networks.

This leaves them with the labelling problem: labelling people, cars, hazards, traffic lights, lane markings and signs in massive amounts of camera footage.

Mobileye solves this problem by employing hundreds of humans to laboriously label these images. It is one of the leaders in this field and its relationships with dozens of companies working in this area show it has been successful.

Intel's acquisition positions it as a direct challenger to other major companies pursuing the same learning-based approach, such as [NVIDIA](#).

In the longer term, we may see companies like [Mobileye](#) and [Xerox](#) switch increasingly to using photo-realistic simulation to generate much of their data.

This approach has the advantage of not requiring any human labelling, as the simulation environment already knows about all the things in the environment.

Seeing is sensitive and subtle

Current self-driving cars are typically much more cautious than humans. This is because humans are able to reliably make more sense of what is happening in the world around them.

In this Tesla video (below, at 0:54), the car slows and almost stops as it drives past joggers on the side of the road.

A human will see the joggers and likely infer that they are very unlikely to suddenly jump out into the road. A machine tends or is explicitly programmed to be more cautious, at least for now.

Vision technologies like those developed by Mobileye can potentially provide much of the more subtle "scene context" to help the car drive more confidently.

Vision technologies can read facial expressions and analyse the body pose and likely intentions of people standing by the side of the road. They can even see into another human-driven car to see whether the driver is looking at the road or [at their phone](#).

Vision-based technology may also potentially integrate more seamlessly with human drivers in driver-assisted systems such as [Toyota's Guardian system](#), which helps humans avoid mistakes.

Seeing and statistics

More than a [dozen companies have demonstrated self-driving cars](#) driving autonomously under varying circumstances. But we still don't have reliable fleets of vehicles we can book at any time, as we can with human-driven ride-sharing services such taxis, Uber and Lyft.

One of the challenges that remains is that the top self-driving cars are reliable most of the time, but not all of the time.

They are also bad at dealing with those one-in-a-million (or billion) events where something completely unexpected happens, such as a [couch being driven down the road](#) or a [man in a giant chicken suit](#) on the side of the road.

In last year's [Tesla fatality](#), a combination of human error and unlikely but possible events contributed to the accident. One of the problems – the autopilot not noticing the [white side of the tractor trailer against a brightly lit sky](#) – is a classic example of the vision problem that everyone is trying to solve.

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

Provided by The Conversation

Citation: How to make a driverless car 'see' the road ahead (2017, March 16) retrieved 25 April 2024 from <https://techxplore.com/news/2017-03-driverless-car-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.