

Should your driverless car kill you to save a child's life?

1 August 2014, by Jason Millar



Big decision for a small car. Credit: laughingsquid, CC BY-NC-ND

Robots have already taken over the world. It may not seem so because it hasn't happened in the way science fiction author Isaac Asimov imagined it in his book *I, Robot*. City streets are not crowded by humanoid robots walking around just yet, but robots have been doing a lot of mundane work behind closed doors, which humans would rather avoid.

Their visibility is going to change swiftly though. Driverless cars are [projected](#) to appear on roads, and make moving from one point to another less cumbersome. Even though they won't be controlled by [humanoid robots](#), the software that will run them raises many ethical challenges.

For instance, should your robot car kill you to save the life of another in an unavoidable crash?

License to kill?

Consider this [thought experiment](#): you are travelling along a single-lane mountain road in an [autonomous car](#) that is fast approaching a narrow

tunnel. Just before entering the tunnel a child attempts to run across the road but trips in the centre of the lane, effectively blocking the entrance to the tunnel. The car has but two options: hit and kill the child, or swerve into the wall on either side of the tunnel, thus killing you.

Both outcomes will certainly result in harm, and from an ethical perspective there is no "correct" answer to this dilemma. The tunnel problem serves as a good [thought experiment](#) precisely because it is difficult to answer.

The tunnel problem also points to imminent design challenges that must be addressed, in that it raises the following question: how should we program autonomous cars to react in difficult ethical situations? However, a more interesting question is: who should decide how the car reacts in difficult ethical situations?

This second question asks us to turn our attention to the users, designers, and law makers surrounding autonomous cars, and ask who has the legitimate moral authority to make such decisions. We need to consider these questions together if our goal is to produce legitimate answers.

At first glance this second question – the who question – seems odd. Surely it is the designers' job to program the car to react this way or that? I am not so sure.

From a driver's perspective, the tunnel problem is much more than a complex design issue. It is effectively an end-of-life decision. The tunnel problem poses deeply moral questions that implicate the driver directly.

Allowing designers to pick the outcome of tunnel-like problems treats those dilemmas as if they must have a "right" answer that can be selected and applied in all similar situations. In reality they do

not. Is it best for the car to always hit the child? Is it best for the car always to sacrifice the driver? If we strive for a one-size-fits-all solution, it can only be offered arbitrarily.

The better solution is to look for other examples of complex moral decision-making to get some traction on the who question.

Ask the ethicist

Healthcare professionals deal with end-of-life decisions frequently. According to medical ethics, it is generally left up to the individual for whom the question has direct moral implications to decide which outcome is preferable. When faced with a diagnosis of cancer, for example, it is up to the patient to decide whether or not to undergo chemotherapy. Doctors and nurses are trained to respect patients' autonomy, and to accommodate it within reason.

An appeal to personal autonomy is intuitive. Why would one agree to let someone else decide on deeply personal moral questions, such as end-of-life decisions in a driving situation, that one feels capable of deciding on their own?

From an ethical perspective, if we allow designers to choose how a car should react to a tunnel problem, we risk subjecting drivers to paternalism by design: cars will not respect drivers' autonomous preferences in those deeply personal moral situations.

Seen from this angle it becomes clear that there are certain deeply personal moral questions that will arise with autonomous cars that ought to be answered by drivers. A [recent poll](#) suggests that if designers assume moral authority they run the risk of making technology that is less ethical and, if not that, certainly less trustworthy.

As in healthcare, designers and engineers need to recognise the limits of their moral authority and find ways of accommodating user autonomy in difficult [moral](#) situations. Users must be allowed to make some tough decisions for themselves.

None of this simplifies the design of autonomous

cars. But making technology work well requires that we move beyond technical considerations in design to make it both trustworthy and ethically sound. We should work toward enabling users to exercise their autonomy where appropriate when using technology. When robot cars must kill, there are good reasons why designers should not be the ones picking victims.

This story is published courtesy of [The Conversation](#) (under Creative Commons-Attribution/No derivatives).

Source: The Conversation

APA citation: Should your driverless car kill you to save a child's life? (2014, August 1) retrieved 16 May 2021 from <https://phys.org/news/2014-08-driverless-car-child-life.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.