

Helping autonomous vehicles and humans share the road

16 November 2016, by Jeffrey C. Peters



Credit: shutterstock.com

A common fantasy for transportation enthusiasts and technology optimists is for self-driving cars and trucks to form the basis of a safe, streamlined, almost choreographed dance. In this dream, every vehicle – and cyclist and pedestrian – proceeds unimpeded on any route, as the rest of the traffic skillfully avoids collisions and even eliminates stop-and-go traffic. It's a lot like the synchronized traffic chaos in "[Rush Hour](#)," a short movie by Black Sheep Films.

Today, autonomous cars are becoming more common, but safety is still a question. More than [30,000 people die](#) on U.S. roads every year – nearly 100 a day. That's despite the best efforts of government regulators, car manufacturers and human drivers alike. [Early statistics from autonomous driving](#) suggest that widespread automation could drive the [death toll](#) down significantly.

There's a key problem, though: Computers like rules – solid, hard-and-fast instructions to follow. How should we program them to handle difficult situations? The hypotheticals are countless: What if the car has to choose between hitting one cyclist or five pedestrians? What if the car must decide to crash into a wall and kill its occupant, or slam

through a group of kindergartners? How do we decide? Who does the deciding?

So far, our transportation system has evolved to be operated by humans, who are good at following guidelines but often interpret them to [properly handle ambiguity](#). We stop midblock and wave a pedestrian across, even though there's no crosswalk. We cross the double yellow line to leave cyclists enough room on the shoulder.

Improving our transportation system to take advantage of the best of machines and humans alike will require melding ambiguity and rigid rules. It will require creating rules that are, in certain ways, even more complex than what we have today. But in other ways it will need to be simpler. The system will not only have to allow automated drivers to function well: It must be easily and clearly understood by the humans at its center.

Human decision-making

[Google cars](#), [Uber self-driving cars](#), [autonomous taxis in Singapore](#), [Tesla's autonomous mode](#) and even [self-driving freight trucks](#) are already on the road. Despite one fatal crash – of a [Tesla on autopilot](#) – [autonomous vehicles](#) are still [safer than a normal human driver](#). Nevertheless, that crash attracted a lot of media attention.

Among the roughly 100 deaths a day on U.S. roads, this one stood out because [people](#) wondered: If the driver was not relying on the [autonomous software](#), what would have happened? What might the human have done differently?


That specific fatal crash was actually fairly straightforward: The [car didn't see a truck](#) in front of it and drove into it. But when people think about accidents, they often worry about having to make moral choices in an instant.

Philosophers call this the "[trolley problem](#)," after a

hypothetical example in which a trolley is hurtling down a track toward some people who cannot get out of the way in time. You have the option to switch the trolley onto a different track, where it will hit some other people.

There are an infinite number of variations on the problem, created by specifying the numbers and types of people, replacing them with animals, sending the trolley into a wall where its passengers die, and more. Would you, for example, save five children and let a senior citizen die? What about saving a dog versus killing a criminal? You can try out many of these variations – and make new ones – [online in a fascinating "Moral Machine" game](#) from which MIT researchers are gathering information on what decisions people make. They hope to find at least some human moral consensus, which can then inform autonomous vehicles and other intelligent machines.

The crux of the problem is whether you choose to switch the trolley or not. In one case, you make an active decision to intervene, deciding to save – and kill – certain groups. In the other, you choose not to act, effectively letting fate take its course. People who use the Moral Machine can see how their results compare to everyone else's. So far the outcomes suggest that people intervene to save younger, fitter people with higher perceived social values (doctors over criminals, for example).



Switch the trolley, or don't? Credit: McGeddon, CC BY-SA

Human – and computer – preferences

To handle these relative preferences, we could equip people with beacons on their cellphones to

signal nearby cars that they are a certain type of person (child, elderly, pedestrian, cyclist). Then programmers could instruct their autonomous systems to make decisions based on priorities from surveys or experiments like the Moral Machine.

But that raises serious problems. For example, would an [autonomous car](#) that noticed a child running in the middle of traffic decide to run over your grandmother on the sidewalk instead?

And what about groups of people? The Moral Machine's creators and other researchers found that society as a whole has a [strong preference for choosing to save more people](#). What if a negligent group of runners steered a car into your path while you walked alone?

The same study also showed that people would be less willing to purchase a vehicle that could include sacrificing the driver (themselves) as an option. If society as a whole is to benefit from the advantages of autonomous vehicles, we need people to buy the cars – so we need to make them more attractive to buyers. That might mean requiring cars to save drivers, as [Mercedes has already decided to do](#).

Breaking the rules

Investigating the trolley problem reveals that "optimizing" for countless specific, but hypothetical, scenarios is not the solution. Further, if we allow autonomous vehicles to break the rules sometimes, under certain circumstances, perhaps malicious humans could game the system. For instance, a pedestrian could walk out in front of traffic without getting hit, but forcing cars to slam on the brakes. That one person might even cause multiple collisions, causing disruption without great risk to the disruptor.

Volvo has already noticed that some [human drivers behave like bullies around autonomous cars](#). For example, a person might cut off an autonomous vehicle because he is confident the other car will avoid a collision itself. As a result, Volvo will not follow the currently common practice of clearly labeling autonomous cars on public roads. At least some of its test vehicles will remain unmarked, in

hopes of measuring differences in human drivers' behavior.

The Mercedes and Volvo developments are the first steps toward trying to clarify human expectations about autonomous cars. By standardizing people's perceptions, it will be easier to predict what humans will do in different scenarios. That will help us engineer ways to keep everyone driving in harmony.

A common set of rules for all autonomous vehicles – whatever those are – will allow people to predict the cars' behavior and adjust our behavior, policy and [transportation infrastructure](#) accordingly.

And if we're going to make [clearer rules](#), perhaps humans should follow them more closely too, as pedestrians, cyclists and drivers. In that world, we probably won't find the perfect chaos of the "Rush Hour" short film. But it will be much more orderly – and safe and efficient – than today.

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

Provided by The Conversation

APA citation: Helping autonomous vehicles and humans share the road (2016, November 16) retrieved 2 July 2022 from <https://techxplore.com/news/2016-11-autonomous-vehicles-humans-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.