

The Moral Challenges of Driverless Cars


Autonomous vehicles will need to decide on a course of action when presented with multiple less-than-ideal outcomes.

EVERY TIME A car heads out onto the road, drivers are forced to make moral and ethical decisions that impact not only their safety, but also the safety of others. Does the driver go faster than the speed limit to stay with the flow of traffic? Will the driver take her eyes off the road for a split second to adjust the radio? Might the driver choose to speed up as he approaches a yellow light at an intersection, in order to avoid stopping short when the light turns red?

All of these decisions have both a practical and moral component to them, which is why the issue of allowing driverless cars—which use a combination of sensors and pre-programmed logic to assess and react to various situations—to share the road with other vehicles, pedestrians, and cyclists, has created considerable consternation among technologists and ethicists.

The driverless cars of the future are likely to be able to outperform most humans during routine driving tasks, since they will have greater perceptive abilities, better reaction times, and will not suffer from distractions (from eating or texting, drowsiness, or physical emergencies such as a driver having a heart attack or a stroke).

“So 90% of crashes are caused, at least in part, by human error,” says Bryant Walker Smith, assistant professor in the School of Law and chair of the Emerging Technology Law Committee of the Transportation Research Board of the National Academies. “As dangerous as driving is, the trillions of vehicle miles that we travel every year means that crashes are nonetheless a rare event for most drivers,” Smith notes, listing speeding, driving drunk, driving aggressively for conditions, being drowsy, and being distracted as key contributors to accidents. “The hope—though at this point it is a hope—is that automation



can significantly reduce these kinds of crashes without introducing significant new sources of errors.”

However, should an unavoidable crash situation arise, a driverless car’s method of seeing and identifying potential objects or hazards is different and less precise than the human eye-brain connection, which likely will introduce moral dilemmas with respect to how an autonomous vehicle should react, according to Patrick Lin, director of the Ethics + Emerging Sciences Group at California Polytechnic State University, San Luis Obispo. Lin says the vision technology used in driverless cars still has a long way to go before it will be morally acceptable for use.

“We take our sight and ability to distinguish between objects for granted, but it’s still very difficult for a computer to recognize an object as that object,” Lin says, noting that today’s light-detection and ranging (LIDAR)-based machine-vision systems used on autonomous cars simply “see” numerical values related to the brightness of each pixel of the image being scanned, and then infer what the object might be.

Lin says with specific training, it eventually will be technically feasible to

create a system that can recognize baby strollers, shopping carts, plastic bags, and actual boulders, though today’s vision systems are only able to make very basic distinctions, such as distinguishing pedestrians from bicyclists.

“Many of the challenging scenarios that an autonomous car may confront could depend on these distinctions, but many others are problematic exactly because there’s uncertainty about what an object is or how many people are involved in a possible crash scenario,” Lin says. “As sensors and computing technology improves, we can’t point to a lack of capability as a way to avoid the responsibility of making an informed ethical decision.”

Assuming eventually these technical challenges will be overcome, it will be possible to encode and execute instructions to direct the car how to respond to a sudden or unexpected event. However, the most difficult part is deciding what that response should be, given that in the event of an impending or unavoidable accident, drivers are usually faced with a choice of at least two less-than-ideal outcomes.

For example, in the event of an unavoidable crash, does the car’s pro-

gramming simply choose the outcome that likely will result in the greatest potential for safety of the driver and its occupants, or does it choose an option where the least amount of harm is done to any of those involved in an accident, such as having the car hit a telephone pole with the potential to cause the driver a relatively minor injury, instead of striking a (relatively) defenseless pedestrian, bicyclist, or motorcycle rider, if the driver is less likely to be injured?

The answer is not yet clear, though the moral decisions are unlikely to reside with users, given their natural propensity to protect themselves against even minor injuries, often at the expense of others, Lin says.

“This is a giant task in front of the industry,” Lin says. “It’s not at all clear who gets to decide these rules. In a democracy, it’s not unreasonable to think that society should have input into this design decision, but good luck in arriving at any consensus or even an informed decision.”

One potential solution would be the creation and use of institutional review boards, which would compel autonomous vehicle manufacturers to provide potential crash scenarios, explain what its vehicles’ capabilities or responses to those scenarios would be, and document and explain why programmers made those choices.

Jonathan Handel, a computer scientist turned lawyer, explains that rather than try to come up with hard-and-fast rules now, when driverless cars have yet to interact on public roads outside of tightly controlled testing runs, these review boards would provide a process to allow manufacturers, lawyers, ethicists, and government entities to work through these nascent, yet important, ethical decisions.

“I propose ethics review boards, or institutional review boards,” Handel says. “I don’t think that we’re at a place in this technology, nor do I think we will be in the first few years of it [being used], that there would be an obvious, one good answer to all these questions. For the ethics issue, I think we need a procedural answer, not a substantive one.”

He adds, “Eventually, consensus may emerge organically on various issues, which could then be reflected in regulations or legislation.”

Given the near-infinite number of

potential situations that can result in an accident, it would seem resolving these issues before driverless cars hit the road en masse would be the only ethical way to proceed. Not so, say technologists, noting unresolved ethical issues have always been in play with automobiles.

“In some ways, there are ethical issues in today’s products,” Smith says. “If you choose [to drive] an SUV, you are putting pedestrians at greater risk [of injury], even though you would believe yourself to be safer inside, whether or not that’s actually true.”



Further, a high degree of automation is already present in vehicles on the road today. Adaptive cruise control, lane-keeping assist technology, and even self-parking technology is featured on many vehicles, with no specific regulatory or ethical guidelines for use.

In all likelihood, Google, Volkswagen, Mercedes, and the handful of other major auto manufacturers that are pressing ahead with driverless cars are unlikely to wait for ethical issues to be fully resolved. It is likely basic tenets of safe vehicle operation will be programmed, such as directing the car to slow down and take the energy out of a potential crash, avoiding “soft” targets such as pedestrians, cyclists, or other smaller objects, and selecting appropriate trade-offs (choosing the collision path that might result in the least severe injury to all parties involved in an accident) to be employed.

Another option to potentially deal with moral issues would be to cede control back to the driver during periods of congestion or treacherous conditions, so the machine is not required to make moral decisions. However, this approach is flawed: emergency situations can occur at any time; humans are usually unable to respond to a situation fast enough after being disengaged and, in the end, machines are likely able to respond faster and more accurately than humans to emergency situations.

“When a robot car needs its human driver to quickly retake the wheel, we’re going to see new problems in the time it takes that driver to regain enough

situational awareness to operate the car safely,” Lin explains. “Studies have shown a lag-time anywhere from a couple seconds to more than 30 seconds—for instance, if the driver was dozing off—while emergency situations could occur in split-seconds.”

This is why Google and others have been pressing ahead for a fully autonomous vehicle, though it is likely such a vehicle will not be street-ready for at least five years, and probably more. The navigation and control technology has yet to be perfected (today’s driverless cars tooling around the roads of California and other future-minded states still are unable to perform well in inclement weather, such as rain, snow, and sleet, and nearly every inch of the roads used for testing have been mapped.)

Says Lin, “legal and ethical challenges, as well as technology limitations, are all part of the reason why [driverless cars] are not more numerous or advanced yet,” adding that industry predictions for seeing autonomous vehicles on the road vary widely, from this year to 2020 and beyond.

As such, it appears there is time for manufacturers to work through the ethical issues prior to driverless cars hitting the road. Furthermore, assuming the technological solutions can provide enhanced awareness and safety, the number of situations that require a moral decision to be made will become increasingly infrequent.

“If the safety issues are handled properly, the ethics issues will hopefully be rarer,” says Handel. ■

Further Reading

Thierer, A., and Hagemann, R., Removing Roadblocks to Intelligent Vehicles and Driverless Cars, Mercatus Working Paper, September 2014, <http://bit.ly/1CohoV8>

*Levinson, J., Askeland, J., Becker, J., and Dolson, J., Towards fully autonomous driving: Systems and algorithms, *Intelligent Vehicles Symposium IV* (2011), <http://bit.ly/1BKd5A4>*

*Ensar, J., Roadtesting Google’s new driverless car, *The Telegraph*, <http://bit.ly/1x8VgfB>*

Keith Kirkpatrick is principal of 4K Research & Consulting, LLC, based in Lynbrook, NY.

© 2015 ACM 0001-0782/15/08 \$15.00