

Statistically, self-driving cars are about to kill someone. What happens next?

As autonomous vehicles rack up more and more miles on our roads, the odds of a fatal accident are shortening by the day. How will we react?

Martin Robbins

Tuesday 14 June 2016 10.39 BST

One hundred million. That's the number of miles, on average, that it takes a human driver to kill someone in the United States. It's also the number of miles Tesla's semi-autonomous 'Autopilot' feature had racked up by May this year. Assuming Autopilot is rolled out to Tesla's mass-market Model 3 in 2017, that number will rapidly climb into the billions. Mercedes are deploying similar systems in their new E-class, while Google's fully driverless cars have racked up another 1.6 million miles and counting.

As the miles grow, the odds shrink. At some point, a car driving autonomously or semi-autonomously will cause a fatal accident. If their performance is remotely comparable to a human's, that moment could come within the next 18-24 months. If so, by the law of averages it will probably involve a Tesla Model 3. Self-driving cars may be about to have their Driscoll moment.

In 1896, Bridget Driscoll was attending a summer fete in Crystal Palace, London, when a car travelling at a "tremendous pace" – somewhere under its top speed of eight miles per hour – struck and killed her. She became Britain's first automobile fatality. (It took another three years for the United States to catch up, when an unfortunate pensioner was mown down by a horseless taxicab in New York.)

There wasn't a big reaction. Parliament had just passed an act raising the speed limit to 14 miles per hour, and raised it again to a dizzying 20 miles per hour in 1903. In 1930, despite an annual death toll running into the thousands, limits were removed entirely. Lord Buckmaster explained the reasoning in the House of Lords two years later:

"It is sufficient to say that the reason why the speed limit was abolished was not that anybody thought the abolition would tend to the greater security of foot passengers, but that the existing speed limit was so universally disobeyed that its maintenance brought the law into contempt."

Will modern day lawmakers and regulators be as relaxed? The accidents are starting. On Valentine's Day, Google hit a major milestone in the development of their driverless cars - the first accident:

"On February 14, our vehicle was driving autonomously and had pulled toward the right-hand curb to prepare for a right turn. It then detected sandbags near a storm drain blocking its path, so it needed to come to a stop. After waiting for some other vehicles to pass, our vehicle, still in autonomous mode, began angling back toward the center of the lane at around 2 mph -- and made contact with the side of a passing bus traveling at 15 mph. Our car had detected the approaching bus, but predicted that it would yield to us because we were ahead of it."

"Our test driver, who had been watching the bus in the mirror, also expected the bus to slow or stop. And we can imagine the bus driver assumed we were going to stay put. Unfortunately, all these assumptions led us to the same spot in the lane at the same time."

Not to be outdone, Tesla has been racking up some boo-boos of its own. One driver managed to

‘summon’ his car into a van, while another owner filmed the moment his car collided with a van parked - ridiculously - in the fast lane of a highway.

There’s a pretty big contrast between the Google and Tesla approaches. Google’s cars are trundling slowly around city streets, a strategy that exposes them to more risk and uncertainty, but also means that any accidents are likely to be slow-speed bumps and scrapes. Tesla’s cars are guiding themselves at high speed on freeways, which are simpler environments but with far higher consequences.

Then there’s the human element. Google’s cars are truly self-driving, intended to require no human intervention. Tesla’s vehicles aren’t, but the technology is good enough that drivers can feel as if they are. That leads to complacency, as neatly summed up by the driver in the second incident above, “Yes, I could have reacted sooner, but when the car slows down correctly 1,000 times, you trust it to do it the next time to. My bad.”

We’ve seen this play out in another industry - air travel. In 2009, Air France Flight 447 plunged into the Atlantic Ocean with 228 people on board. William Langewiesche’s 2014 essay on the crash is one of the best things ever written about automation, and serves as a sanity check of the idea that drivers of autonomous road vehicles could simply intervene in the event of a crash.

AF447 was lost after a sequence of events in which sensors failed, causing the autopilot to be disabled, and changing the aircraft’s control system to a mode that left pilot and co-pilot hopelessly confused, as chillingly captured by the flight recorders: “We completely lost control of the airplane, and we don’t understand anything! We tried everything!”

The crash raised a lot of debate about pilots, autopilots, and how the two interact with each other; in particular, whether reliance on technology has resulted in less-experienced - and therefore less-capable - pilots. As Langewiesche put it, “automation has made it more and more unlikely that ordinary airline pilots will ever have to face a raw crisis in flight—but also more and more unlikely that they will be able to cope with such a crisis if one arises.”

The lesson here, for planes, trains or automobiles, is that one plus one does not equal two. Combine an autopilot with a good driver, and you get an autopilot with, if not a bad driver, at least not such a good one.

The question is whether that drop in human performance is matched or surpassed by the improvement that the computer brings to the table. In aircraft that’s proven to be the case, and accidents have generally pushed the industry greater computer control rather than less. Pilots have increasingly become onlookers and supervisors rather than hands-on drivers.

For cars, Tesla and others have launched what amounts to one of the biggest human-computer interaction experiments in the world to find out, trialling novel control modes and algorithms on inexperienced and inexperienced drivers, and streaming data from thousands of vehicles back to the cloud for analysis.

It’s an experiment that’s left regulators like the U.S. National Highway Traffic Safety Association scrambling to catch up, with no clear consensus on questions like, “how do you even test this software?” By the time they do, it’s likely that the technology will already be an accepted fact of life, its safety taken for granted by consumers, its failures written off as the fault of its error-prone human masters.

If that’s the case, then the first death at the hands of a self-driving car could have much the same impact as Bridget Driscoll’s demise, 120 years ago this summer: not much at all.

More blogposts

Topics

Artificial intelligence (AI) Tesla Self-driving cars Computing Transport policy Road safety