

HARD ROAD AHEAD

BY DEE SMITH

Our contemporary utopians promise a cleaner, greener world free from accidents via driverless car technology. But like all powerful innovations, this one brings with it serious risks. These can be as mundane as robot drivers having difficulty handling human behavior on the road — or as sinister as the cars being hacked and used as weapons. In this risk analysis, Dee Smith games out what a future filled with automated autos holds and finds the picture far less bright than their advocates make out.



The world is beset with problems: violent extremism, climate change, pandemics, and crises in global governance (to mention just a few). But never fear — the utopians are back. They have reemerged in the form of roboticists, with glowing predictions extolling the golden age soon to be upon us.

Of course, what is driving the robotics revolution is not altruism and the ambition to create a wonderful, limitless future for humanity, but the wish to create a pile of cash. There is nothing wrong with that, of course, unless in the process you generate negative externalities that stop you from keeping and enjoying the large pile of cash.

Robotic or driverless cars are a useful exemplar of many of the general ideas and controversies involved here. According to an IHS Automotive forecast, there may be as many as 54 million autonomous cars on the road by 2035. The utopians predict that so-called autonomous on-demand networks will allow much greater utility for cars, which are at present an asset that sits largely unused most of the time. Cars will pick up passengers and drop them off and run errands, so to speak, to pick up things for consumers. Some estimates claim that with the introduction of this technology, the overall number of cars on the road in developed nations can be reduced by 80 percent, saving resources across the board (although other prognosticators say there will actually be more cars on the road).

How likely is it that driverless cars will take to the roads in any meaningful way in the near future? If you listen to Google's new parent company Alphabet, to Tesla, Uber, GM, and Baidu — all of whom have robust driverless car development programs — that future will be arriving very soon. If and when it does, it will bring problems with it. These fall basically into four categories: technical, regulatory, implementational, and cultural.

It is impressive to see what Google has done in Mountain View, where dozens of prototype driverless pods move

around the town every day with few mishaps. But the seeming beehive of driverless-car activity belies the more sobering truth that from a technical standpoint, it is still very difficult to make driverless cars capable of dealing with unexpected or “fuzzy” situations, particularly when these situations do not follow the rulebook on which the cars’ programs are based. A large part of that rulebook consists of traffic laws, which sounds straightforward until

“Driverless cars could be used for targeted killings.”

—Dee Smith



the reality of what actually happens on the road collides with the rulebook. Consider one simple example: engineers have had an extremely difficult time dealing with driverless cars faced with four-way (or even two-way) stop signs at which human drivers are present. Traffic laws uniformly dictate that a driver should pull up to a stop sign, stop, and then pull forward. The problem is that almost no human drives that way. Human drivers generally slow to a crawl without stopping and then accelerate when they have determined that no traffic is coming. In other words, they violate the letter of the traffic law almost every time they encounter a stop sign. But when faced with such behavior by the humans sharing the road, the robots — who were programmed to expect what the law dictates — refused to move. This is just one example of perhaps the biggest issue with driverless cars: their inability to solve problems when faced with the unexpected. A sudden downpour or hailstorm can stop them in their tracks, let alone more unpredictable elements such as a basketball or rogue trashcan careening down a road.

Think for a moment about what driverless-car technology actually entails. The algorithms that allow it to operate are themselves both new and constantly encountering new situations. They face human drivers prone to mistakes and sometimes unpredictable actions. Perhaps most consternating is the fact that the entire system requires the streaming of all situational data to and from the cloud for analysis on a real-time, microsecond-to-microsecond basis. Not exactly a recipe for stability.

Indeed, systemic instability is already visible. In February of this year, a Google autonomous vehicle detected sandbags by a storm drain that blocked its path, and came to a stop. After it started again and was heading back towards the center of the lane, it hit the side of a bus that was traveling at 15 miles per hour (the driverless car was going about two miles per hour). Google says the car detected the bus, but its algorithmic system predicted that the bus would yield to it. Tesla, which is already deploying its semi-autonomous “Autopilot” on freeways — a very different environment from the sedate Mountain View roads on which Google’s fully automated vehicles are running — now faces a more serious challenge, following the death of a driver, Joshua Brown, when one of their vehicles drove under a truck in May. He had a “need for speed” according to press reports, and may have been playing a DVD when the accident occurred. However, it is notable that Tesla has been trying to explain why its Autopilot (contrast the name of this feature with terms used by Volvo: “semi-autonomous tech” and “park assist pilot”) is not actually liable for Brown’s death. Tesla also has stressed that, in 130 million miles driven by Tesla vehicles using the autopilot feature (which is in “public beta”), this is the first fatal accident (noting that the average in the U.S. is one fatality every 94 million miles).

Some of the proposed solutions to the problem of driverless-car accidents (if indeed “solutions” they are) border on the absurd. Google, for example, was recently awarded a patent for a sticky substance to be put on the hood of its driverless cars so that, when a pedestrian is hit by such a car, he will stick to the hood instead of bouncing off and being hit by another vehicle (or a second time by the same driverless car). Lest you wonder how many insects and other detritus might also stick to the

hood, never fear: it will be covered by a hard substance that only shatters when something heavy enough hits it. So apparently when we are hit by a Google driverless car, we will adhere to the vehicle, covered with the shards of the shattered shell and

some substance much stickier than glue. It is hard to determine which will be hurt more: one's body or one's dignity.

Another very serious safety issue relates to the possibility of driverless cars being hacked and used as weapons for murder or assassination. This problem is very real; though it is not discussed much publicly, there is no evident solution for it (as people within the industry will admit if pressed). Systems built by humans can, almost definitionally, be hacked by humans.

There are at least two ways hacked driverless cars could be used for targeted killings: by causing a car carrying the target to crash into a bridge, tree, or another car at high speed; or by commandeering a second car and crashing it into the car carrying the target. On a wider scale, a hacker bent on causing chaos could simply target many cars within a defined geography such as a city and bring traffic to a standstill that could take hours if not days to unwind. Even a mischievous, talented teenager would probably be able to accomplish this. Such an attack would have massive effects on supply chains, stranding trucks carrying just-in-time goods and perishables. Organized criminals could crash several cars into a truck or even an armored car and quickly arrive for the goods. But what terrorists might engineer could be much worse: crashing driverless cars into trucks containing toxic or flammable substances, perhaps instigating dozens of such incidents at the same time in the same city. And when driverless cars arrive, can driverless trucks be far behind? Imagine crashing a truck full of gasoline into a factory conducting a process that utilizes corrosive substances. Or into a museum, or a government building, or a school.

In the regulatory arena, another set of problems arises. One is economic. An 80-percent reduction in cars on the road, if it happens, will reduce income to municipal and state governments as revenue from sales taxes, speeding tickets, and public transportation drops. Another is exemplified by an ongoing debate between the California Department of Motor Vehicles and Google (*inter alia*) about whether driverless cars should have a human component, even if it is confined to assuming control in an emergency. Many autonomous car developers claim this will make their cars less safe. After all, data already shows that many people, when told their car can drive itself most of the time, shift their attention to other things and stop looking at the road. Sometimes they go to sleep. So trying to shift the attention of zonked-out passengers back to driving in an emergency will be very challenging indeed, especially since even proponents of driverless cars admit that it will be a long time before driverless cars are ubiquitous and human drivers disappear from the roads.

This brings up another key regulatory issue. When the inevitable crashes happen, who will be liable: the human driver or the robot? Volvo has recently declared that it will assume liability, because the company realizes that if they are not willing to do so, no one will trust the technology. Many observers feel that in almost all cases, the robot will be blamed.

This brings the discussion around to the cultural issues: will humans actually be comfortable with robotic cars? To put it more precisely, will people be willing to ride in a 4,000 pound robot careening down a road at 60 miles per hour or more — a road filled with other 4,000 pound robots moving at similar speeds? This circles back into the argument for keeping a human element in the equation, an approach championed by Toyota in contrast to Alphabet. Beyond this, to be accepted it may be necessary for driverless cars to drive in a way similar to the way humans drive (both for the comfort of the passengers and for the comfort of other human drivers on the road). Robotic cars conduct the basic business of driving — acceleration, braking, etc. — in ways that humans do not find comfortable. Presumably this can be addressed by programming trial and error, but it is itself an ambitious and subtle undertaking.

Ironically, putting driverless cars on the road that generally do not need human supervision from humans — but sometimes urgently require it — may result in drivers with less experience, and hence less capability, being required to intervene with less

notice in situations with which they are less familiar (if their attention even can be turned to the need at hand quickly enough). The U.S. military, which has far longer and deeper experience than the car industry with trying to make autonomous systems work, does in fact call this cluster of issues the “irony of automation.” Unmanned automated systems end up having poor reliability and needing extensive supervision to operate effectively.

Driverless cars — like the broader issue of autonomous artificial intelligences — are still in their nascent stages. But the questions they raise are stark and troubling. It is all the more imperative that we develop coherent answers and policies sooner rather than later.

Dee Smith is founder and CEO of Strategic Insight Group.