The rise of autonomous vehicles and why ethics matter

By Rudeon Snell
19 Mar 2019

There is little doubt that autonomous vehicles, more commonly known as self-driving cars, are set to be transformational.

Rudeon Snell, Leonardo Lead at SAP Africa

The market alone is set to reach roughly $42bn by 2025 with a compound annual growth rate (CAGR) projected at 21% until 2030. Not only will this disruption be immense, giving birth to new types of businesses, services and business models, but their introduction will trigger similarly immense ethical implications.

The World Health Organisation (WHO) highlights that every year approximately 1.35 million people globally die, in traffic accidents. If there was an effective way to eliminate 90% of those accidents, I don't think there's any doubt that most people would support it. This is the aspirational goal that autonomous vehicles aim to achieve: to eliminate the main source of traffic accidents, namely human error. The benefits of autonomous vehicles are certainly clear - time saved, increased productivity, improved safety, continuous service availability - but the challenges remain and if not addressed, could potentially wreck the promise this technology holds.

Ethical challenges with autonomous vehicles

One of the key challenges for autonomous vehicles is centred around how they value human life. Who decides who lives and who dies in split-second decision-making? Who decides the value of one human life over another? And, more importantly, how is that decision calculated? Being on a road has inherent danger, and this danger means that trade-offs will occur as self-driving cars encounter life or death situations. It is inevitable. And with trade-offs you need ethical guidelines.

If an autonomous vehicle makes a mistake, it could directly lead to loss of life. In these scenarios, the question of who decides who lives and how that decision gets made, becomes very important. Researchers have been grappling with this notion for years. The Trolley problem tests have been around since 1967. First proposed by the philosopher Phillipa Foot, it has subsequently proliferated into many variants. Generally, it is used to assess what actions people would take when asked to take an action that would, for example kill one person, vs killing ten people.

Let's assume an autonomous vehicle experiences a mechanical failure, acceleration increases, and the car is unable to stop. If it continues, it will crash into a large group of pedestrians, but the car may swerve and crash into an obstacle killing the driver inside the car. What should the car do, who decides what the car should do and where would liability rest for that decision? Experienced human drivers have been programmed for years to deal with split-second decisions like these and they still don't always get it right.

Moral utility vs moral duty

Should autonomous vehicles perhaps adopt utilitarian principles when having to choose who lives and who dies? Utilitarian principles, as advocated by the great English philosopher Jeremy Bentham, could offer a framework to help with the ethical decisions AIs need to make. The focus would be on AIs making decisions that result in the greatest good for the greatest amount of people. However, at what cost? How do utilitarian calculations that violate individual rights get reconciled? Why should my life be less important than five strangers I don't know?

Or should autonomous vehicles potentially follow duty-bound principles as advocated by the German philosopher Immanuel Kant? Under this system, a principle such as "thou shalt not kill", would ensure the car is duty-bound to maintain its course even if it harms other people, so long as the driver remains safe. With duty-bound principles, prioritising your individual safety could in effect cause even more harm.

I suspect the key lies in how customers eventually adopt and consume autonomous vehicles and the services they offer. Would we as consumers even be interested in owning self-driving cars that follow utilitarian principles and potentially kill us to save a bunch of strangers? Or would we be interested in only buying self-driving cars that prioritise our own safety, even though it might mean potentially killing a bunch of strangers? Who decides what ethical guidelines AI in autonomous

vehicles will follow? The harsh truth is that ultimately, we as consumers do.

Autonomous vehicles will come to market and won't be perfect. We know autonomous vehicles today struggle to detect small objects like squirrels. It's quite possible that they also cannot detect squirrel sized potholes and rocks that could cause life threatening scenarios with tire blowouts.

As with most technology advancements in the past, we as consumers decide what technology trend gets pushed by technology providers. We do so by voting with our money. If 90% of autonomous vehicle sales are for units that prioritise our lives potentially at the expense of others, and only 10% of autonomous vehicles are sold following utilitarian principles, guess where the focus will be for future autonomous vehicles? Yes, there might be mitigating alternatives. Autonomous vehicle manufacturers could decide to give consumers the choice of how their vehicle operates. But at what cost? And would that capability be feasibly scalable across regions and cultures?

Consumers' power to decide

Autonomous vehicle manufacturers could decide that in order to get their self-driving cars sold, they need to develop them according to consumer preferences. In the past, regulation has been used to address this type of conundrum, but with the world of technology accelerating at an exponential rate, regulation just doesn't seem to be able to keep up. And even if it did, supply and demand, as it more often than not has in the past, will inform the design of future products.

The increased usage of autonomous vehicles is ultimately premised on trust. The people who buy or use them, have to trust the technology and must be comfortable using that technology for true value to be realised. In order to build this trust and acceptance, autonomous vehicle manufacturers must ensure the technology is safe and caters to the needs of their consumers, whatever those needs might be.

#BizTrends2019: The way we love our cars is about to change
Kriben Reddy 7 Jan 2019

In the early 1900s when elevators became autonomous, it made a lot of people very uncomfortable. People were so used to a driver being in the elevator the very idea of an elevator operating autonomously scared them senseless. No one wanted to use them. So, manufacturers invented compromises. Soothing voices were introduced, big red stop buttons were installed. Safety bumpers were introduced, and creative ads were run in the media to help dispel fears. Gradually, people accepted the change. Today we barely give getting into an elevator a second thought.

Waiting until the technology is better is simply not a viable option. There's no guarantee how long it will take to perfect and while we wait, millions will continue to die. Progress should be iterative in nature, because the technology of today can help save millions of lives today. We should not let perfection be the enemy of good enough right now, especially considering the alternative. Is that perhaps the real ethical dilemma we are facing with autonomous vehicles?

An obscured path ahead

Then there's also the notion that perhaps we do not have to decide how to value human life. Would it possibly make sense to let AI determine its own outcomes being guided only by baseline moral principles that are globally accepted? Perhaps a set of principles, similar to Isaac Asimov's three laws of robotics? The challenge is that ethics differ, just as values differ. The way they manifest culturally differs across countries and regions. What is acceptable in Japan might not be in Europe. How we value life in the East, could be different to the West. This diversity has its benefits, for sure, but it also provides massive challenges when trying to design systems that can function universally.

Today, advanced technology is developed under the purview of the leading technology institutions and not under the

purview of governments. An example of this is evidenced in the amount of money spent by the top five US defence contractors on research and development. It comes to less than half of the total research spent of any of the major technology players like Microsoft, Apple, Google, Amazon, and Uber.

A key takeout is that policy makers do not govern advanced technology in the commercial world, ultimately technology providers do. Some consequences can be anticipated and are linked to the promises made on behalf of the technology, while some consequences unfortunately remain unforeseen.

My hope is that we strive not to be surprised by unintended consequences which could derail the promise that breakthrough technology offers, just because we didn't take a moment to anticipate how we could deal with them. Right and wrong is after all, influenced by factors other than just the pros and cons of a situation. If we ask the hard questions now, we can drive our world in the direction we want it, not just for us today, but for our children tomorrow.

ABOUT RUDEON SNELL

Director: Intelligent Solutions at SAP EMEA South Africa
▪ The rise of autonomous vehicles and why ethics matter - 19 Mar 2019

View my profile and articles...