MILLS & REEVE

briefing

November 2015

Why we should get used to the idea that self-driving cars will sometimes crash

The world is waking up to the prospect of driverless vehicles. Google, Audi, Mercedes, Tesla, Volvo, Apple, among others, are in a race to be first onto the road with a driverless car. Volvo's Drive Me project in Gothenburg is expected to provide vehicles, which will operate fully autonomously on certain roads, to normal families in 2017. Toyota plans to market a vehicle with highway autonomy by 2020. Countries are competing with one another to provide testing grounds to entice manufacturers and developers.

Manufacturers and commentators alike discuss the benefits that these vehicles will bring. For example, reduced emissions, reduced congestion, free time in which to do other things, increased mobility for the young, old and disabled, safety. It is with the last of these factors – safety – that we hit something of a roadblock.

A roadblock?

On the one hand everyone appears to agree that autonomous cars are likely to be safer than human drivers. There have been many studies into the extent to which road traffic accidents, deaths and injuries are caused by driver error. Some studies suggest that in over 90 per cent of accidents, driver error is the major cause. Other (more driver-friendly?), studies suggest that it is upwards of 75 per cent. And indeed these may be conservative estimates as it is likely that humans seek to avoid and underreport their responsibility for accidents. Logic drives us to conclude therefore that autonomous vehicles, without the same human error prone propensity, will lead to a dramatic reduction in road traffic accidents, deaths and injuries. This increased safety, the lives that will be saved and the life changing injuries that will be avoided, is often used as a justification for introducing autonomous cars.

However, just at the point where we conclude that driverless means safer, some difficult questions are asked. How will system designers be able to programme the vehicles to behave appropriately in any given set of circumstances? How do we teach the system to make ethical choices when faced with competing decisions? Should the system choose to hit a child or swerve into an oncoming truck? Other theoretical, and often increasingly unlikely, scenarios are often posed. It is at this point that most people throw the question at the law-makers. It is usually viewed as a legal liability problem which regulators will have to resolve.

Clearing the (legal?) roadblock

The starting point when considering driverless cars is generally assumed to be that our roads are complex environments with many variables and so before we can allow autonomous vehicles on the road we need to be sure that they will be safe in every environment, at all times.

But aren't we missing the point?

Let's bring the question back to the way we currently regulate our roads. How do we educate drivers to behave appropriately in any given set of circumstances? How do we test drivers to ensure that they react safely at all times? How do we teach drivers to make ethical choices when faced with competing decisions? Is a driver instructed on how to choose between hitting a child or swerving into an oncoming truck? Before we allow human

drivers on the road do we make sure that they will be safe in every environment, at all times? Is this how we regulate (and teach) our current drivers?

People are generally taught how to drive from A to B. They are taught to steer a vehicle and to brake and accelerate so as to navigate the road network. Humans are rarely given any kind of emergency scenario training. Some nations require a degree of skid training but this is the exception to the rule. Humans tend to need to: be over the age of about 16/17/18, pass a theoretical test on the rules of the road and pass a short practical test (at best involving an emergency stop). The human is then free to drive a metal box at speeds of up to 70 miles per hour (or more) towards other people (either in other metal boxes or just walking by the roadside). The words of a driving instructor to a newly qualified driver could be as cynical as: "Just make sure you do your best not to hit other cars or people, although if you do hit them, so long as you weren't drunk, drugged, speeding or being unduly reckless, we're likely to let you off any personal liability for any injury or death you cause."

Think about it. Do you sometimes query the ability of other drivers? A driver might pull out in front of you at a junction, or drive too close to your rear bumper, or make an unsafe or erratic last minute manoeuvre. Do you think other drivers often lack the necessary skills and awareness to drive? I think most of us have had that feeling. But still, we are satisfied to trust the decision-making of the user of another vehicle when faced with an emergency. Every time we venture onto the roads, we place our lives in the hands of these other drivers. Is that logical?

Society owes a huge debt of gratitude to the pioneers of seat-belts, anti-lock brakes, airbags and the like. Road traffic deaths and injuries have fallen dramatically and there is the potential for them to fall further still. The latest driver assistance feature to be championed by Euro NCAP is autonomous emergency braking (AEB) – a system which brakes for the driver if the driver fails to react in time to an obstruction in front of the vehicle. Just as seat-belts, ABS and air bags have become standard in all new cars, so too AEB is likely to become standard. Recent studies suggest that AEB is having a huge impact by eliminating or reducing the damage caused by certain types of collisions.

One key feature of autonomous emergency braking is the word "autonomous". Whereas safety features that have gone before AEB generally assist the driver in taking action, AEB takes action if the driver fails to do so. It is an autonomous feature that kicks in when "driver error" rears its ugly head. And AEB is arguably the first step towards building the safety system for driverless cars.

The driverless car safety system

Driving a vehicle can be split neatly into two distinct elements. First, the obvious one - navigational control; the activity of getting from A to B. The second? Taking emergency evasive action if an emergency situation arises. We call it critical event control. Currently, the driving test and public testing of driverless cars, focus largely on navigational control. It is only when the examiner slaps the dashboard, usually once, during a driving test, is there a crude attempt to test a driver's ability to exercise critical event control.

In relation to driverless cars, it can be seen that, by splitting control in this way, the safety standards for driverless cars will be dependent on the driverless car's ability to exercise autonomous critical event control.

At any moment in time, when faced with a critical event scenario, a vehicle (whether driven by a human or a system) has three choices: brake, accelerate or steer (in each case in different degrees of magnitude). (A vehicle can alert others to the problem by sounding a horn or flashing lights, but these actions do not positively enable the vehicle to react to the critical event scenario.) AEB deals with braking and manufacturers are developing emergency steering assist (ie, autonomous emergency steering). Any driverless car will need a system that behaves appropriately in any given critical event scenario by exercising autonomous control over the braking, steering and acceleration of the vehicle.

That does however bring us back to what is perceived as the legal and regulatory problem – how do we programme the autonomous critical event control system to behave appropriately in any given set of circumstances? How do we test the system? How do we programme it to make choices – choices that are perceived to be ethical choices?

The ethical dilemma?

Law and regulation generally evolve with changes in public behaviour and changes in public acceptance levels. When the horse was first replaced with an engine (and the horseless carriage was born) a man had to walk in front of the vehicle with a red flag to warn other road users of the approaching car. When the M1 motorway was first opened, there was no speed limit. Road traffic accidents were largely an unknown when the first cars took to our roads. They eventually happened, and the law then had to evolve to deal with them in a way that society would tolerate.

And tolerate we do. We are incredibly tolerant when it comes to forgiving humans for accidents and errors of judgement that take place while a human is behind the wheel. We condemn drivers who are drunk or reckless or behave with similar levels of socially unacceptable conduct, but we are quick to forgive mistakes or other less obviously inconsiderate forms of behaviour. The driver distracted by a passerby, the driver changing the CD or the radio station, the driver blinded by the sun, the driver surprised by the thickness of the snow or the depth of a puddle, the driver who simply did not see the child in the road ahead, the driver who brakes too late or who was driving too close, the driver who slightly misjudged his/her overtaking manoeuvre, the driver who accelerated instead of hitting the brake pedal. All of these “driver error” accidents, in the absence of other wrongful contributory behaviour, are largely forgiven, generally without criticism or penalty.

Society tolerates road traffic accidents that cause death and injury, in the US each year there are 30,000 deaths and 300,000 life changing injuries. Mankind has an illogical tolerance for death and injury that takes place on our roads. Accidents happen, people get killed and unless the driver is driving recklessly or in a manner which society deems morally reprehensible, the driver more often than not walks away, only having to live with the knowledge that their acts (or omissions) have led to an accident (and possibly someone’s death or injury). There is no other area of our daily lives, where society tolerates such a level of death and injury caused by human error.

Like it or not, driverless cars will not be able to eliminate all death and injury on our roads. No matter how many programming hours or how much scenario modelling we do, no matter how much code is written or testing conducted, death and injury involving driverless cars will still happen. Autonomous cars will not be the cure for all road traffic accidents, injuries and deaths.

The cyclist will still swerve into the path of an autonomous vehicle at the last moment (giving the vehicle no time to swerve or brake) and be hit. The pedestrian will fail to see the autonomous car and will step off the kerb and be injured (again at the last moment giving the car no time to swerve or brake). The child will run after the football, just at the wrong moment and will be struck by the autonomous car. The lorry will shed its load and the driverless car will collide with the wayward cargo. The tree will fall over and onto the driverless car, just as it passes by. These accidents and others like them (accidents which do not involve driver error) will still happen in a world where the autonomous car exists.

Can we be confident that autonomous vehicles will largely eliminate death and injury caused by driver error? That is 3,000 deaths in the UK, 30,000 in the US, roughly 1 million worldwide (and ten times that number of people have life changing injuries). A recent study by the Eno Centre for Transportation in the US, concluded that if 90 per cent of cars on American roads were autonomous, the number of accidents would fall by almost 80 per cent. Google’s autonomous cars have been on the roads of California for six years and about 2 million miles. There have been 16 accidents (only 16) and all but one of these were caused by the other vehicle. The one remaining crash happened when the Google car was in the hands of a human driver.

There'll be an initial period during which driverless cars are present alongside standard vehicles. This may present its own problems. Human drivers may be uncertain as to how to react when faced with driverless cars (and human drivers may make mistakes and cause accidents through a failure to interact appropriately). The smooth convoying of driverless vehicles will be disrupted by the presence of driven vehicles. But once the transition is complete, won't the world be a much safer place?

The question society has to grapple with is whether we are ready to accept (a lower number of) accidents involving robot cars, in a similar way to the way we accept and tolerate (the numerous) accidents that happen when humans are at the wheel. We may be making some progress. An online poll by the Telegraph newspaper asks readers "Would you trust a self-driving car?" So far, 74 per cent say they would.

The road ahead?

There is a place for law and regulation. Legislators can propose parameters for society to debate. One approach would be to view autonomous cars as trains on virtual rails, but with an added benefit that when it is completely safe to do so they can swerve off their virtual rails to avoid hitting the obstruction ahead provided that they can re-join their virtual rails without colliding with another obstruction in the process. This proposal does not need to involve an ethical choice (like the old theoretical dilemma referred to as "the trolley problem"). The only choice for the vehicle is "can it steer around the primary obstruction without colliding with another obstruction?". If yes, it takes evasive steering action, if no, it does its best to brake but then impacts the primary obstruction (in much the same way that it is usually safest for the human driver to do). Clearly irrelevant obstructions that the vehicle would drive over in usual circumstances (crisp packets, leaves, branches etc) are obviously ignored when analysing the potential evasive steering manoeuvre.

Such an approach has many benefits, but it assumes that society will tolerate the fact that those who step or fall or run or drive in front of an autonomous vehicle will, unless the autonomous vehicle has sufficient space to stop or swerve to avoid them, be hit by the autonomous car. This is not a radical approach; those who step or fall or run or drive in front of a driver-controlled vehicle can, unless the driver has sufficient time to react and brake or steer (and it is safe to do so), expect to be hit by the car. And society currently tolerates such accidents. The main difference is that in the world of the autonomous car, the car is likely to react far faster and far better than the human driver.

Perhaps we should all ask ourselves, if I accidently step off the kerb in front of a car tomorrow, would I prefer it to be an autonomous car which has been designed and tested so as to stop as quickly as possible when faced with the prospect of hitting a wayward pedestrian, or would I prefer to take the risk on the driver being sufficiently rational, alert and sober to be able to jump on the brakes?

Autonomous vehicles are an opportunity of a lifetime to make a real difference. The main obstacle is not a legal question, it is one of public acceptance and tolerance.

Stephen Hamilton
Partner
T +44(0)1223 222349
Stephen.Hamilton@mills-reeve.com

Isabel Teare
Technology Lawyer
T +44(0)1223 222402
Isabel.Teare@mills-reeve.com