Autonomous Cars & Software Testing – Part 1 of 3

Introduction

This is the first of a three-part article on autonomous cars and software testing.

This part initially introduces reasons for developing autonomous cars and the high-level technological challenges that are involved. It then considers some of the ethical issues, and suggests areas where regulation is needed.

In part 2, the new technologies that are needed for autonomous systems are described in more detail, such as in the areas of sensors, vehicle-to-vehicle communications, and machine learning.

Part 3 provides explanations of how traditional automotive lifecycle practices of requirements specification and architectural design must change for autonomous cars. The article then suggests how autonomous cars will provide new challenges and opportunities for software testing.

The first two parts are aimed at anyone who wants to understand more about autonomous cars, while part 3 is more focused at those involved in managing or performing software testing of the autonomous systems inside them.

Autonomous Cars Save Lives and Money!

According to the World Health Organization, more than 1,250,000 people die each year through road traffic accidents, with at least a further 20 million suffering non-fatal injuries, which often leave them disabled for life. These accidents cost countries approximately 3% of their GDP.

With 93% of these accidents being caused by human error, and the required technology now becoming available, autonomous cars could be an ideal solution. With many car makers (and other new entrants into the car market), promising us that fully autonomous cars will be available within the next 3 or 4 years, they could dramatically reduce accident rates in a relatively short time.

While saving lives (and money) should make autonomous cars particularly attractive to governments, many drivers also see the latest features as highly desirable when they buy a new car, which is another reason that autonomous cars are now a major talking point.

What is an Autonomous Car?

Different people have varying ideas of what is meant by the term ‘autonomous car’. The dictionary tells us that autonomous simply means having the freedom to act independently. In the context of autonomous cars, we normally interpret it to mean that the car can drive from A to B without any human input, while being able to cope with uncertainties in the environment, such as other vehicles and pedestrians, that may not always be following the ‘rules’.

Some people discriminate between ‘automated cars’ and ‘autonomous cars’ by considering a car to only be truly autonomous if it takes no help from outside the car, but in today’s connected world

this is unrealistic. There is no point restricting autonomous cars from accessing useful information that is available from other cars and from transport networks. We typically expect our car of the future to be a connected, autonomous car that requires less and less input from its human passengers.

Levels of Autonomy

There are various schemes for measuring the level of autonomy of cars. By far the most popular scheme is provided by the SAE [1], as shown in the table below.

SAE Level	Description
0 - No driving automation	The same as driving was 20 years ago – you do everything.
1 - Driver assistance	The car provides assistance in just one axis – for instance, accelerating and decelerating with a cruise control, or steering with lane control (but not both). The driver must do the rest of the driving.
2 - Partial driving automation	The car provides assistance in two axes – for instance, both accelerating, decelerating and steering. This means the driver can take their foot off the accelerator and their hands off the steering wheel, such as for self-parking, but they must be ready to take control immediately.
3 - Conditional driving automation	The car does all the driving in certain situations (e.g. on the motorway). For example, from November 2017 the Audi A8 will be the world's first production car with this level of autonomy, through its 'Traffic Jam Pilot', which drives the car in traffic travelling at up to 60 km/hr. This only works on a motorway with a safety barrier – and it checks this condition. With this level of autonomy active, the driver can do other tasks (such as reading email), but should be available to take back control if something goes wrong. However, the driver does not have to pay continuous attention – the car will come to a safe stop if the driver does not take back control.

4 - High driving automation	The car potentially does all the driving for the entire journey, but this is limited to certain types of road and environmental conditions. For instance, it probably cannot go off road (e.g. if it needs to avoid a broken-down car) or may not cope in certain extreme weather conditions. The driver should not need to take control on most journeys, but the car needs a steering wheel, etc., for those 'special' occasions.
5 - Full driving automation	The car does everything – anytime, anyplace, anywhere (nearly). So, you should be able to stick the kids in the car and let it take them to school. It needs no steering wheel or other similar controls.

Other classification schemes are also used, for instance, another, simpler approach uses only 'automated' (SAE levels 1 and 2), 'autonomous' (SAE levels 3 and 4), and 'driverless' (SAE level 5).

The above schemes consider the car as a standalone system, but it is also sometimes useful to see autonomous cars as part of a larger system, as this allows us to also look at the level of connectedness of the cars. One such scheme defines three levels of autonomous cars as:

- level 1 - standalone autonomous cars (as SAE scheme above);
- level 2 – convoys, swarms or ad hoc groups of connected autonomous vehicles, which can communicate and share useful information (e.g. the car behind the truck can 'see' the pedestrian by receiving local environment information from the truck);
- level 3 – autonomous cars connected to a transport network, which allows cars to be re-routed around accidents, and passengers to be shifted to other more efficient modes of transport, when necessary.

A further way of looking at levels of autonomy, is to consider the complexity of the situations the autonomous car must deal with. In such a scheme, the environments in which the car can drive autonomously can be used to define the levels, with the urban situation, in which pedestrians cross between parked cars, and dogs run out into the streets defining the highest level of complexity.

Why do we want Autonomous Cars?

As has already been mentioned, there is a strong argument for using autonomous cars if they can reduce the enormous number of road traffic accidents that occur each year. This will save lives and reduce injuries, nearly half of which are not to car drivers and passengers, but to pedestrians, cyclists, and motorcyclists [2]. This will also lead to reduced insurance and healthcare costs.

Other reasons for adoption are that autonomous cars:

- have better perception of their immediate environment (human drivers have blind spots);
- are quicker and more consistent at decision-making, so are more likely to decide on the right manoeuvre, especially in more complex scenarios;

- can execute selected manoeuvres faster and more accurately than humans, by directly controlling the car's acceleration, braking and steering;
- can never get drunk, take drugs, get tired or distracted (these causes account for over 40% of the car crashes in the US);
- allow those who have mobility problems and cannot drive themselves to travel by car;
- reduce stress in drivers in traffic jams and other sub-optimal driving situations;
- connected to a transport network, have reduced journey times and costs, and subsequently cause less congestion and damage to the environment.

There are also several arguments against autonomous cars, such as:

- car owners will no longer have the pleasure of driving (this only applies to owners who opt for cars with full autonomy);
- there will always be new, complex scenarios that autonomous cars cannot cope with (but many human drivers will find these scenarios equally confusing, and once encountered, autonomous cars will be updated so that next time they do cope);
- criminals will perform cyber-attacks on connected autonomous cars (however, these same cars will be far less likely to attract traditional car crime);
- car drivers will become less proficient at driving if they get less practice because they normally travel in a fully autonomous car;
- users may not trust autonomous cars with their lives, when there is so much news coverage of them crashing (we will consider this aspect in more detail later);
- the technology to build autonomous safety-critical systems is too new and too expensive to make them viable in the near future.

The last point raises the question of the technological challenges that stand in the way of autonomous cars. The next section introduces these challenges, as, unless these can be overcome, there is little point in considering other issues, such as ethics, regulation and acceptance by the public.

Technological Challenges

Autonomous cars are extremely complex systems, and, to achieve their full potential, they need to be part of a far larger system of systems, an integrated transport system. For autonomous cars to be acceptable to the public and regulatory bodies, several leading-edge technologies need to work together at the same time.

To understand how these technologies work together, consider the high-level view of the functions in an autonomous car shown in Figure 1.

In this simplistic view, the sensors (e.g. cameras, GPS, RADAR, LIDAR) are used to gather information, and this is used to determine an understanding of the car's environment, such as the positions of nearby cars, pedestrians and information on road signs. Part of this 'sensing' function is also known as localization, which is determining the car's current position in the environment and relating this to detailed offline maps. The 'decision making' decides what the car's next move should be (e.g. braking, turning) depending on the function provided by the autonomous system (e.g. adaptive

cruise control). The 'control' implements the decision by calling on actuators (e.g. to release air, open fuel valve). If we were to add detail to the diagram, then we would also need to show interaction of the system with the passengers (who will tell it where to go) and communication with other systems, such as other cars (i.e. V2V) and other systems (e.g. emergency services).

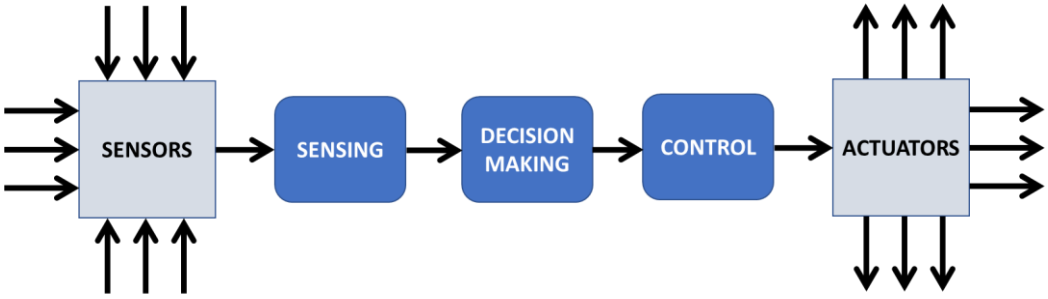

Figure 1: Basic Autonomous Car Functions

Using the view provided in Figure 1, we can identify several potential technological challenges, such as:

- Sensors: Cars can only become fully autonomous if they can completely understand their current situation in all types of weather and at any time of day or night. Some systems are based on cameras that mimic human vision, however adding other sensors can give the system a better understanding (and add needed redundancy in this area).
- Sensing: Collating the enormous amounts of data that are being gathered into a single 'picture' of the environment requires both image processing combined with data fusion. Traditional programming techniques do not appear to be able to do this, especially in real-time, and so machine learning is being used instead.
- Decision making: The complexity of the decision making, which requires an optimal decision to be made based on many different inputs in real-time, typically requires machine learning to be used.
- Control: The low-level control of the actuators should be similar to that used in current automated, but non-autonomous, systems.
- Actuators: The actuators should be similar to those used in current automated, but non-autonomous, systems.

Thus, fully autonomous cars will require new sensors, and to make sense of the data from these sensors, the system will typically use deep learning, a form of machine learning. To perform the necessary decision making, the system will often also use deep learning. See the later sections on 'sensors' and 'machine learning' for more details on these specific areas.

Although these technologies work in theory, when implemented as part of an autonomous car, this all needs to be integrated into a single embedded system that not only works in real-time, but also needs to be safe and secure – and so needs to have fault tolerance built in. And, to make matters worse – we are not building a spacecraft with unlimited funding, instead we want to sell reasonably-priced cars to normal users - the embedded system (including sensors, processors, memory…and software) therefore needs to be relatively cheap.

One thing we already know is that the current approaches to developing E/E systems for cars, compliant with ISO 26262, are not good enough to build fully autonomous systems. Even the new version of ISO 26262, due out in 2018, does not provide sufficient coverage of how the development (and more importantly) testing of systems built using technologies like machine learning. For the sensing and decision-making functions, there will often be inadequate specifications, and they will be implemented using probabilistic approaches, resulting in non-deterministic results (i.e. we cannot predict what the function will do in all circumstances). Suggestions on new approaches for the requirements specification, architectural design and testing of autonomous systems that can meet new regulatory standards are provided in later sections.

When will we get Fully Autonomous Cars?

To answer this question, we need to know how safe autonomous cars need to be before they can be used on our streets.

From a utilitarian viewpoint the answer is simple – "as soon as we can say that autonomous cars are safer than human-driven cars". With this view, we would replace human drivers as soon as the autonomous systems would cause fewer accidents (and deaths) than their human counterparts. Logically, this makes sense – in this case, the introduction of autonomous systems should save lives overall.

Consider the figures from earlier. Let's say we decide to introduce autonomous cars when they are 20% safer than human drivers (assuming we can test well enough to know this and there are enough autonomous cars to replace all our human-driven cars). With the current human drivers, approximately 1,250,000 people die each year from road traffic accidents. In our new 'autonomous' world about 250,000 people who would have died are now saved (and much money is also saved that would have been spent on medical and repair bills). And, because we now have millions of autonomous cars driving many miles each day, the machine learning components of the systems will be sharing their experiences and learning at a great pace, so that in the next year there should be far fewer accidents.

However, in that first (hypothetical) year, 1 million people will be killed in autonomous car crashes. Do you remember how many newspaper column inches were devoted to a single fatal crash in a Tesla Model S in 2016? Imagine what would happen if over 27,000 people were dying each day, and autonomous cars were to blame.

If autonomous cars are introduced too early, then it is likely that, when the inevitable crashes occur, users will lose trust in them and then it may take many years to persuade users to try them again. In which time human drivers will continue to cause crashes at a rate far higher than autonomous cars would cause. From a practical perspective, the problem is to decide the tipping point at which it becomes worthwhile introducing autonomous cars. Perhaps it is when they are safer than human-driven taxis or buses. In the US and UK, compared to travelling in a car, taxis appear to be no safer, while buses are about 20 times safer. Ensuring that fully autonomous cars are 20 times safer than their human-driven equivalent, even if this is deemed acceptable, may mean that their introduction is delayed beyond many expectations.

Gradual Introduction by Level of Autonomy – a Difficult Journey?

The alternative to moving straight to full autonomy, that many car manufacturers seem to prefer, is an evolutionary or incremental approach, whereby users are introduced to autonomous cars with gradually increasing levels of autonomy. This is the approach we have taken so far – there are already many cars on the roads at SAE levels 1 and 2, and there is now even one level 3 model on the roads (the Audi A8, as of November 2017).

The problem, however, is that although it's a nice idea to share responsibility for driving between the computer and the driver, it's not clear that this hybrid solution works well in practice. It seems that while the computer drives the car, the human driver is very easily distracted, and, the longer the computer drives without needing help, the more the human driver is tempted to perform secondary tasks, such as reading emails and tweeting.

As an example, Google (now Waymo) were testing their SAE level 3 (human will intervene, if necessary) cars and found that their human passengers were often not in a position to intervene if something had gone wrong. They have video of 'drivers' putting on make-up and even falling asleep while the autonomous car was driving at up to 90 km/h, and their CEO stated that drivers found it "hard to take over because they have lost contextual awareness." They decided to stop development of the SAE Level 3 option, but instead attempt to move directly to the full autonomy provided at SAE levels 4 and 5.

Another example of drivers finding it hard to share responsibility with the autonomous system appears to be the Tesla Model S crash of May 2016. The driver was found to have had their hands on the steering wheel for just 25 seconds of the 37 minutes that the car was on (SAE Level 2) Autopilot, travelling at just below 120 km/h, and the car warned the driver six times about this. Despite this, the driver was distracted enough to not react to a truck crossing his path that he could have seen for at least 7 seconds before he crashed into it.

How do Ethics and Regulations apply to Autonomous Cars?

The final decision on when autonomous cars can drive on our roads should be down to the regulators. A major problem with this is that while governments provide legislation for the regulators, at the same time they are also trying to ensure that their motor industry is one of the leaders in autonomous car development – and so want to encourage the development and use of autonomous cars. At present it seems that most governments have not considered this to be a problem, and are leaving industry to use current (non-autonomous) vehicle legislation and to self-regulate in those specialist areas unique to autonomous vehicles. For instance, the US NHTSA published their 'Federal Automated Vehicles Policy' in September 2016 [3], but have left it as a guidance document rather than implementing it as legislation.

Ethics and Autonomous Cars

We have already considered the question of when is the right time to allow autonomous cars onto our roads, based on the number of lives saved. A common ethical debate around autonomous cars is concerned with the rules that are implemented as part of the decision-making part of these

systems. The debate is often described as an extension of the 'Trolley Problem' (see Figure 2). In the trolley problem, the subject can see a runaway trolley (train) that, if left alone, will run into and kill five people. However, the subject has access to a switch that will divert the trolley onto a second track, the dilemma being that on this track a different person is killed. The subject must decide whether to let matters take their course (in which case, five people are killed) or change the course of the trolley and be instrumental in killing someone.

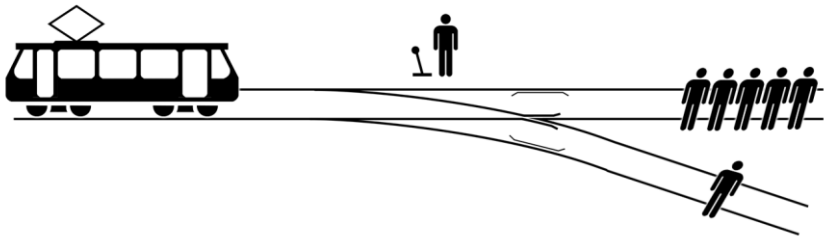

Figure 2: Trolley Problem

The rules that are implemented by an autonomous car will have to cover situations like the trolley problem. For instance, an autonomous car may find itself in a situation where if it drives in one direction it will kill the occupants, while driving in the only alternative direction will kill pedestrians. Some manufacturers have already stated that the answer is simple – they will program their cars to always save the occupants. After all, would you buy a car that was programmed to sacrifice you, if given a choice? Others argue that pedestrians should always be given the benefit, while many believe that the maximum number of people saved would be the right decision.

Since September 2017, Germany has a set of guidelines on autonomous cars that state that they must do the least harm, if crashing into people is unavoidable [4]. So, it seems that in Germany the minimum number of people are expected to be harmed (but how they measure harm is not described). They also provide additional guidance that the cars should not take account of the age, gender, race or disability of the people involved (they must have great expectations for the fidelity of the sensors).

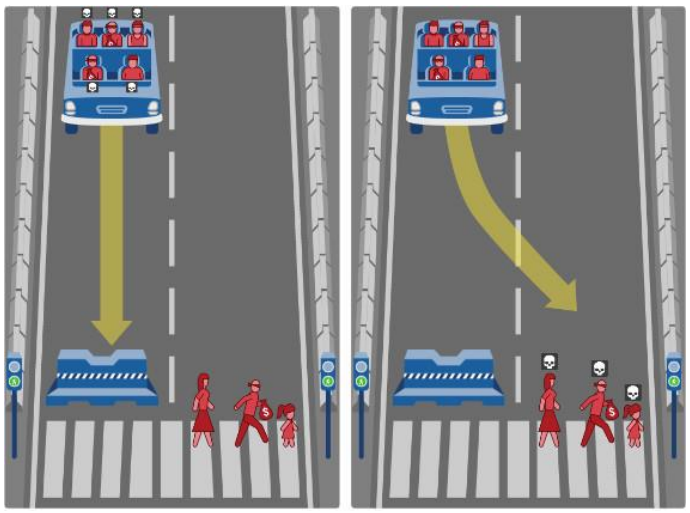

Figure 3: MIT Moral Machine Scenario

Other, less emotive examples of ethical decisions would consider the merits of different collisions between vehicles and stationary objects, and their relative costs. If you are interested in this topic, try looking at (and providing your judgments on) a set of autonomous car scenarios at http://moralmachine.mit.edu/, an example of which is shown in Figure 3.

There is also the ethical question of how easy we make it for autonomous cars to break driving laws. A typical human driver on an empty road will avoid the deep pothole, even if it means temporarily driving on the wrong side of the road, but how do we program our autonomous car to make that sort of judgment?

If the public are going to trust autonomous cars, there is an argument that all the rules that they implement should be transparent and open to debate. Transparent rules could possibly lead to a situation where individual organizations could attempt to optimize their rules if two cars were heading for a collision to benefit their customers. Such 'gaming' of the rules could eventually lead to potentially unsafe behaviours and we believe the way to prevent this is for all autonomous cars to work from an agreed common set of rules that can be included as part of the regulations that autonomous cars must comply with. This is covered in depth in the later section on the specification of requirements for autonomous cars.

At a different level, another ethical decision that needs to be considered is the number of jobs that are likely to be lost by human drivers. Is a government that encourages the use of autonomous cars obligated to provide some financial or employment 'safety net' for out-of-work taxi and truck drivers? At a similar level, what might happen to the car insurance industry if autonomous cars become so safe after a few years of machine learning that accidents are so rare that insurance becomes unnecessary?

Regulation of Autonomous Cars

Certification against a Set of Safety Rules (Scenarios)

In the previous section, we suggested that the rules applied by the decision-making function of autonomous cars should be transparent and regulated. In this way all autonomous cars would need to implement and follow the same set of rules (we are not suggesting that the *implementation* of the decision-making function is duplicated, only that the same rules that provide the specification for this function are used by all car makers). This would create an economy of scale, as the derivation of these rules (and their subsequent improvement – as there is no way they will be 'right first time') would only need to be done by one group of experts rather than duplicated by every developer of a decision-making function. If all autonomous cars follow the same set of rules this should ensure a baseline level of inherent safety, while also making the behaviour of autonomous cars predictable to some extent, which would be useful and safer for both other autonomous cars and human drivers.

With a defined set of rules, these can also act as the specification for certifying autonomous cars. Certification would be based on an autonomous car passing a set of tests based on the specification. As is explained in more detail in the later section on testing, it is very important that these tests are completely independent of the development. The execution of the tests should also be

independent, as public trust could be easily eroded if accidents were associated with a particular car model, and its manufacturer had performed the regulatory tests themselves.

One way of specifying these rules would be as a set of scenarios and their expected outcomes. For the developers of the decision-making function, these rules could be used as the basis for training the machine learning (in the safety-related aspects), and the same set of scenarios could also be used by the independent certification body as the basis of their test sets.

Incident Reporting

With such new technology, there will be many opportunities for improvement as new, unexpected scenarios are discovered and better ways of handling existing scenarios are identified. It is expected that most of this activity will be based on incidents with live systems (e.g. crashes and near misses). It should be mandatory for such incidents to be reported to the regulatory body, as this will allow them to identify the reason for the incident and to determine what should be done to prevent future similar incidents. In some cases, this will mean that incident details are shared with other suppliers of autonomous cars, so that they can all improve their systems to avoid similar incidents.

Sharing of Information

Cars at higher levels of autonomy will be connected ('connected cars') and will be able to communicate with each other to optimize manoeuvres, such as overtaking and merging into traffic and they will also share information on road conditions. If one vehicle's sensors can 'see' a pedestrian that may be obscured from another vehicle (e.g. behind a bus or around a corner), then this information can be shared with other vehicles. The sharing of such information may be regulated in as much as it may be mandatory for a connected car to broadcast useful sensor data that may prevent accidents. See part 2 of this article for more on connected cars.

Enhanced ISO 26262 Standards for Autonomous Cars

Development and testing of the E/E systems for autonomous cars should follow the established ISO 26262 set of standards [5]. These should be enhanced for their 3rd edition (it is too late for the 2nd edition, which is due to be published in 2018), to take account of the new technologies (e.g. non-deterministic machine learning systems) that are expected to be the basis of autonomous cars (see earlier section on 'Technological Challenges').

Next Parts

In part 2, the new technologies that are needed for autonomous systems are described in more detail, such as in the areas of sensors, vehicle-to-vehicle communications, and machine learning.

Part 3 provides explanations of how traditional automotive lifecycle practices of requirements specification and architectural design must change for autonomous cars. The article then suggests how autonomous cars will provide new challenges and opportunities for software testing.

References

[1] SAE J3016™: Taxonomy and Definitions for Terms Related to On-Road Motor Vehicle Automated Driving Systems, SAE International, 2016.

[2] Road traffic injuries Fact sheet, World Health Organization, see http://www.who.int/mediacentre/factsheets/fs358/en/, May 2017.

[3] Federal Automated Vehicles Policy, NHTSA, September 2016.

[4] Maßnahmenplan der Bundesregierung zum Bericht der Ethik-Kommission Automatisiertes und Vernetztes Fahren (Ethik-Regeln für Fahrcomputer (www.bundesregierung.de), Sep 2017.

[5] Road vehicles – Functional safety, ISO 26262 Parts 1-10, ISO, 2011.