Autonomous Cars & Software Testing – Part 2 of 3

Introduction

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

The first part of the article introduced reasons for developing autonomous cars and the high-level technological challenges that are involved. It then considered some of the ethical issues, and suggested areas where regulation is needed.

In this part, 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.

Specialized Technology for Autonomous Cars

Sensors, Communication Technology and Situational Awareness

For an autonomous system to make sensible decisions, it needs as much information as possible about its immediate environment - in all weather conditions and at all times of day and night. The system needs to be able to recognize the road space and road markings, moving objects, such as vehicles and pedestrians, as well as stationary objects close to the road, such as buildings and rivers (this is because to avoid a potential collision with a moving object the car may decide to leave the road space and needs to know how this can be done most safely, e.g. not into a tree). It needs situational awareness, and to achieve this it creates a digital picture of its surroundings.

To be fully autonomous, these cars need suitable information to travel safely at speed on motorways and for manoeuvring into parking spaces. The information requirements for the different driving scenarios can vary quite widely. A fast car can travel about 50 metres per second, and, at this speed will need to be looking ahead for obstacles perhaps 200 metres ahead, whereas for parking the system needs a more localized all-round view. This means that autonomous cars need a comprehensive sensor fit, with some sensors selected to provide information for one specific function, while some other sensors support a range of different functions. As an example, Figure 1 shows a proposed sensor fit for Delphi and Mobileye's Centralized Sensing Localization and Planning (CSLP) automated driving system planned for 2019 [1].

To navigate, autonomous cars need to know exactly where they are and to have a map of the route they are following. However, fully autonomous cars do not just need to map their route in terms of which roads they are going to take, they also need to map their route in far more detail, such as where they position themselves in road lanes (in the middle is not always the safest position) and how they cross multi-lane junctions, such as roundabouts.

Access to detailed mapping information (more than is needed by your satellite navigation system) can be used to supplement the real-time sensor information (e.g. data on road width and nearby obstacles), so several companies are spending a lot of time and effort creating very detailed maps to support autonomous driving. These detailed maps obviously make the tasks performed by autonomous cars easier, however autonomous cars also helping to create the detailed maps. While autonomous cars are driving, their sensors pick up detailed information about their environment, some of which can be used to improve and supplement the detailed maps. Therefore, as more journeys are driven, the detailed maps will get better (a form of crowd sourcing) – and so the autonomous driving will get both safer and more efficient.

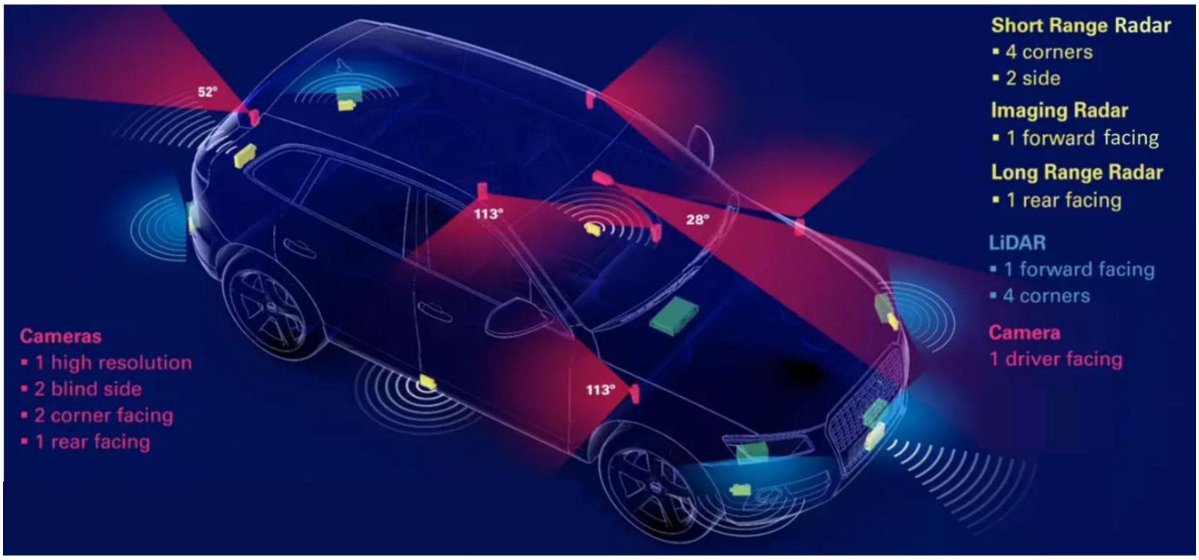

Figure 1: Example Sensor Fit - Delphi and Mobileye's CSLP

One opinion about the sensors required by an autonomous car is that they need to mimic the sensory inputs used by a human driver, which would limit the sensors to stereoscopic cameras able to see in most directions (humans have blind spots). If we also assume that most drivers have some spatial awareness and access to a map, we could also include a GPS (and map data) to the required sensor fit. Many current driver aids at SAE levels 1 and 2 (e.g. Tesla Autopilot) follow this approach and rely heavily on cameras and GPS. However, the use of other sensors allows the autonomous car to potentially have better perception than human drivers (e.g. radar can see through fog), and having more sensors also provides a level of redundancy in the event of a sensor failing.

The following list describes the most popular sensors currently used or planned for use in autonomous cars:

- **Cameras** – Mono and stereo cameras are available, with stereo cameras providing the ability to better judge distance. Colour cameras are also used, for instance, to determine the colour of traffic lights. Note that some countries require rear view cameras to be fitted (e.g. mandatory for all new cars in the US from 2018).

- **Radar** – Uses microwaves to detect and track objects in all weather conditions. Typically, not as accurate, but relatively cheap, compared to LIDAR. Two types of radars are normally used; long range radar works at up to 150 metres, while short range radar works from 20 to 60 metres.
- **LIDAR** – An acronym for light detection and ranging, it works in a similar way to sonar (see ultrasonic, below), but instead of using sound waves, it sends out beams of laser light (many thousands per second), and uses the time they take to reflect back to the sensor to determine the distance of the object they hit (speed is determined by looking at the change in distance over multiple frames). Until very recently they were extremely expensive and large (e.g. the large spinning device shown on the top of the car in Figure 2), but their prices have dropped by about 90% since 2009 and they are now small enough to fit in the car front grille. Their accuracy has also increased, with Waymo reporting the LIDAR on their cars can detect an object the size of a football helmet at 200 m. More accurate than radar, but more expensive, and susceptible to dirty lenses and bad weather. Figure 2 shows an example 3-D image from LIDAR.

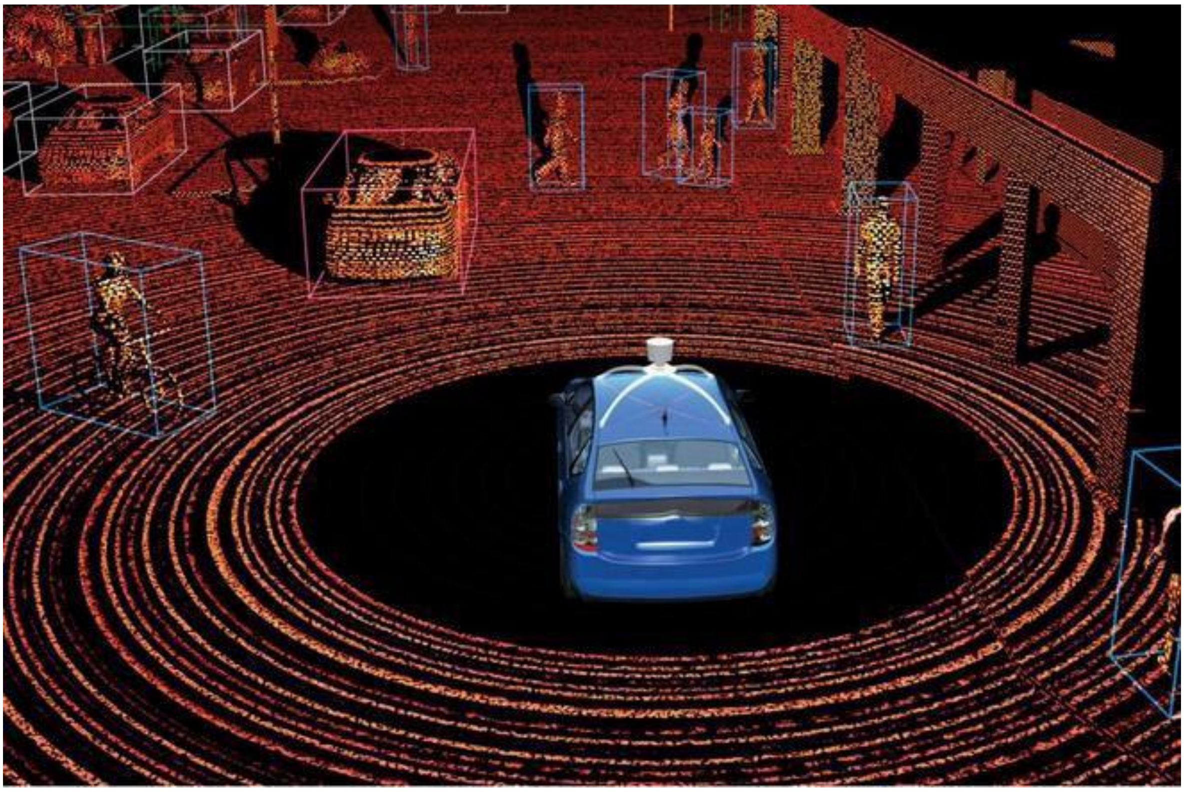

Figure 2: Example LIDAR image

- **GPS** – Provides accurate car position, direction and speed data. When used with detailed (and up-to-date) map data, this provides the system with a detailed awareness of its surroundings.
- **IMU** - Inertial measurement unit, used to supplement localization information provided by the GPS, similar to the old 'dead reckoning' approach of navigators. It can be used as a secondary check on the plausibility of the GPS, but is particularly useful when the GPS signal is lost (e.g. in tunnels).

- **Roadside real-time locating system (RTLS)** – Small, radio beacons tell the car where they are. Useful to provide confidence in the GPS, or act as back-up when GPS signals are not available.
- **Infrared** – These devices provide thermal imaging, which can provide extra information in adverse weather conditions and at night. They are very good at detecting pedestrians and animals at distances up to 50 metres, and the more sensitive devices can also identify changes in the road surface, such as black ice.
- **Ultrasonic** – These sensors provide accurate information on objects close to the car, and are useful when performing low-speed manoeuvring, such as parking.

Given that many of the systems driving autonomous cars are safety-critical, when deciding on the sensor fit, care must be taken that the system is not reliant on a single sensor (a single point of failure). It can also be useful to use information from diverse sensors so that a weakness in one sensor is compensated for by a strength in another.

Connected Cars

Internally, today's cars typically include several types of network (e.g. CAN bus, FlexRay, automotive Ethernet, LIN and MOST), with some networks being duplicated to provide redundancy for safety-related functions. They also receive position information over wireless networks, such as GPS (see above), and traffic information over the Traffic Message Channel, which is integrated with FM, DAB and satellite radio broadcasts.

As the level of autonomy in cars increases, so do their information needs, which has led to a new information source - other vehicles. Vehicle-to-vehicle (V2V) communication and co-operation allows other vehicles to provide useful information to the autonomous system. Vehicle-to-everything (V2X) communication and co-operation allows transport networks to provide traffic and accident data to the vehicle that will allow it to calculate a more efficient route to its destination. At the same time the car provides its own information to the transport network. V2V and V2X communication use the IEEE802.11p (wireless access in vehicular environments - WAVE) protocol that supports Intelligent Transport Systems (ITS) applications, operating in the 5.9 GHz frequency band (ITS-G5) [2]. Vehicles using this protocol will broadcast information 10 times per second, at distances of up to 800 metres.

The Car-2-Car consortium [3] suggest several scenarios that could be supported by V2V communication, such as:

- **Hazardous Location Warning** – One vehicle detects that the road is hazardous (e.g. an oil spill or black ice) and broadcasts a hazard warning to other vehicles. These other vehicles may, in turn, re-broadcast the warning to other vehicles that may be travelling toward the problem area.
- **Traffic Light Optimization** – Traffic lights broadcast information about the layout of the intersection and the sequence of light signals, so allowing approaching vehicles to optimize their approach and transit through the intersection.
- **Motorcycle Warning** – Motorcycles broadcast an awareness message to warn other vehicles of their position and direction. These warning could be especially useful as motorcycles

approach intersections, where the view of the approaching motorcycle may, for instance, be blocked by a truck.

- **Emergency Vehicle Warning** – Emergency vehicles, such as ambulances, broadcast a warning message to other vehicles that they are on an urgent journey. This can be received before drivers see or hear the lights and siren, allowing them to get out of the way more easily.
- **Incident Warning** – Vehicles that have been involved in an accident or have broken down broadcast a message to warn any approaching vehicles. This should help prevent unnecessary collisions with the stationary vehicle and with the people working on it.
- **Roadworks Warning** – Temporary roadside units warn vehicles of the position and size of roadworks and any restrictions (e.g. speed limits) associated with them.
- **Self-Organising Traffic Information System** – Vehicles broadcast information about their journey, such as their average speed on various roads and adverse weather conditions, so allowing other vehicles to optimize their route planning. Surprisingly, only 1 to 2% of vehicles need to support this system for a detailed view of traffic in an area to be available.

Connected cars at higher levels of autonomy will also be able to communicate with each other to optimize manoeuvres, such as overtaking and merging into traffic. Trucks are already using this technology to trial platooning, where several trucks travel close together, typically controlled by a single human driver. 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 (other sensor data may also be shared). Even parking may be made easier as vehicles share information on empty parking spaces that they detect. There is also the potential for vehicles to detect and broadcast information on other vehicles that may be breaking traffic laws (presumably driven by human drivers).

Because V2V and V2X involve radio communication, they all come with extra security implications that do not apply to the sensors that only work across the car's internal wired networks.

Data Fusion

It is obviously beneficial from a redundancy (and, so safety) viewpoint that multiple types of sensors are used in an autonomous car – you never know when a sensor may fail (e.g. a camera may be hit by a stone or not recognise an object). Using multiple sensors also allows the deficiencies in one sensor to be supplemented by other sensors. For instance, cameras are likely to struggle in dense fog, and could be supplemented by radar. However, if we use multiple sensors, we must fuse the data from these multiple sensors and the detailed off-line map.

Typical sensor data that is fused could be the front camera and front radar, or rear camera and ultrasound to assist parking. With these pairs of sensors, the camera identifies the types of objects and the ultrasound and radar determines their distance, but first we must calibrate the sensors to match the data from them.

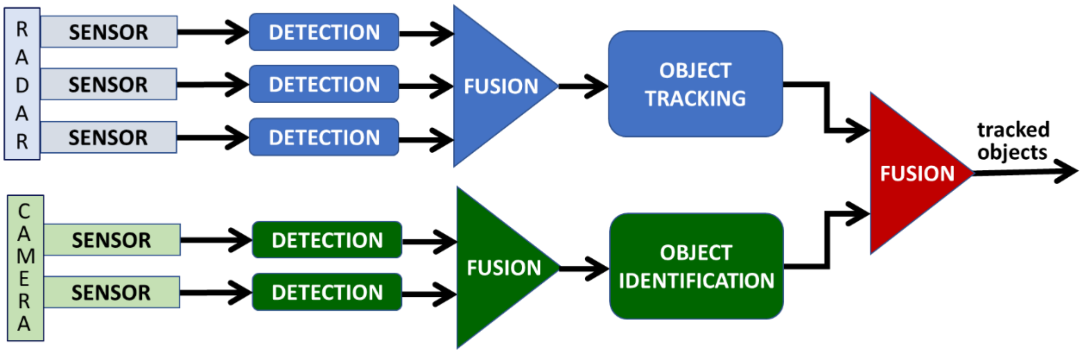

Figure 3: Example Levels of Data Fusion

Figure 3 shows examples of different levels of data fusion within a system. Three radar inputs are fused to track objects (low-level fusion), and two camera inputs (perhaps colour and infrared) are fused to identify objects (low-level fusion). The results are then fused together (high-level fusion) to provide tracked objects that can be used by the decision-making function (as shown in Figure 5).

Artificial Intelligence and Machine Learning

Using traditional software development approaches to the sensing and decision-making required by autonomous cars with their rigid real-time constraints does not work very well. There are several problems with implementing the rules that are needed to identify images and decide on the correct course of action using the traditional 'if-then-else' approach. The first is the size of the task. Very many rules would need to be implemented, and even if we undertook this task, we just cannot predict every situation that the autonomous car will encounter (in terms of the environmental conditions they need to sense, and situations they must react to). Thus, following a traditional approach would lead to our autonomous car occasionally finding that it couldn't recognise something in its environment (e.g. we forgot to program what a swarm of ladybirds or a person in a fancy-dress costume looked like), or it would not know what to do when an unusual situation occurred (e.g. we didn't consider pigs running down the middle of the road). Another reason we don't use a traditional approach is that for many situations there are many different factors that need to be considered, and we know that optimization based on multiple factors is not possible in the timescales needed for a real-time system driving an autonomous car.

Therefore, we need an approach that can implement a practical approach to optimization, and that still works when a situation that has not been pre-considered occurs. The solution is to use a form of artificial intelligence (AI). AI systems can typically sense, decide on a suitable action, perform the action and then adapt based on the results.

As shown in Figure 4, machine learning is a subset of AI, and is widely used, in applications such as internet search engines, voice recognition, dealing on financial markets, language translation and now, in autonomous cars. These applications all rely on the ability of the machine learning system to efficiently identify and use patterns in data. In practice, a specialised form of machine learning, known as deep learning, is typically used in autonomous vehicles. Deep learning uses layers of neural networks to represent the knowledge gathered from the data.

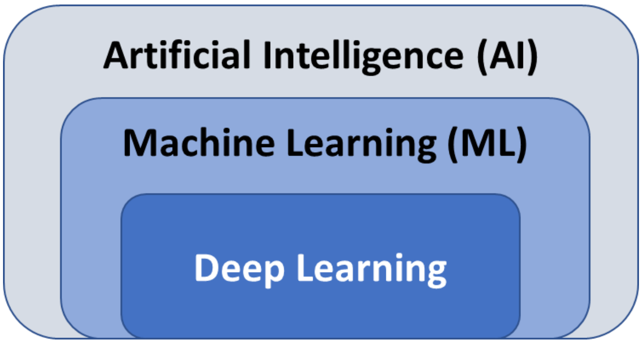

Figure 4: AI, ML and Deep Learning

At a high level, machine learning is used to either *classify* an item or *predict* an outcome. In autonomous cars we use both capabilities. The ability to classify items is used by our sensing function (see Figure 5) to determine what is in the car's environment. The ability to predict outcomes is used by our decision-making function to decide what would be the best action to perform next.

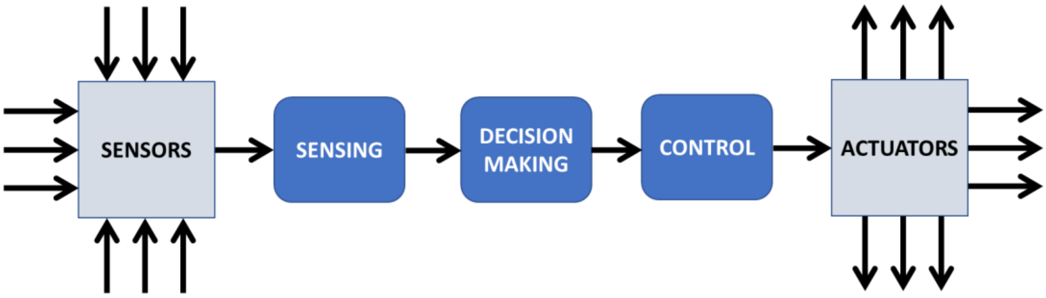

Figure 5: Basic Autonomous Car Functions

Supervised, Unsupervised and Reinforcement Learning

Machine learning uses three basic approaches to decide what the resultant program should do from the data we provide it. These are:

- supervised learning;
- unsupervised learning; and
- reinforcement learning.

An example of supervised machine learning is shown in Figure 6. The system learns by being provided with many images that are tagged with a label that tells the system what they represent (in our example, boys, trees and pumps). When the resultant system is passed a *different* image of a boy, the idea is that the system can recognize that it is a boy, even though it has never seen that particular image before. The system will also often provide a probability that its guess is correct.

The tagged images are the 'training set' for the machine learning system – this training set can be considered as a form of requirements specification. If we want our system to be better in some critical areas, we can use a risk-based approach by providing more training data in these areas. We must also be careful that our training data does include clues (e.g. other data that can be used to determine the classification, known as ‘data leakage’), otherwise the resultant system will only work well when this other data is also available. We test the system using a second (completely independent) set of data, known as the test set.

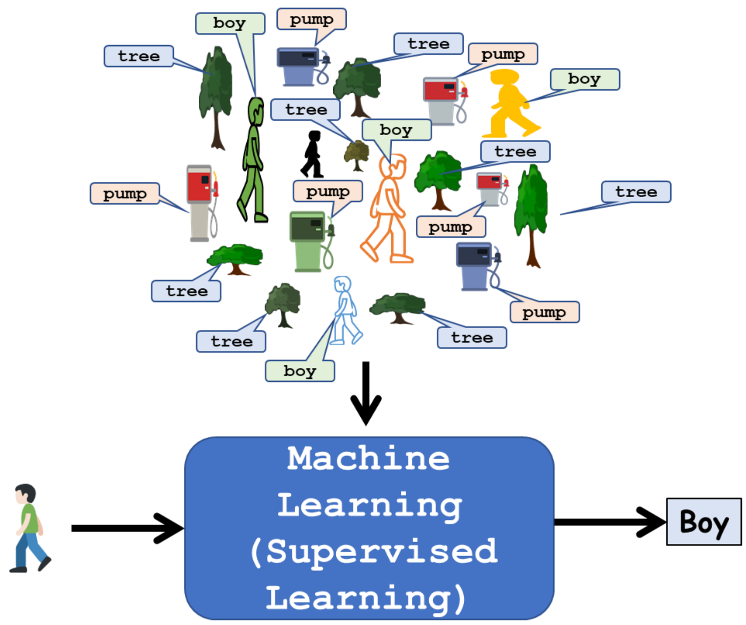

Figure 6: Classification of Pedestrians

Supervised learning is typically used for autonomous cars to teach systems how to interpret their immediate environment (i.e. the ‘sensing’ function in Figure 5). Mobileyes is one of the leaders in the field of sensors for autonomous cars and they predict that they will have a workforce of approximately 1000 people tagging images to support the supervised learning of their machine learning functions by the end of 2017. Supervised learning can also be used to learn good driving by being provided with many scenarios of human drivers (with good driving skills).

With unsupervised learning, the machine learning system is passed lots of data (as with supervised learning), but this time without the tagging information. From this data, the system can identify hidden patterns in the data that are unrecognizable to humans.

Reinforcement learning uses a rewards-based approach to machine learning, which is similar to how some children learn (the autonomous system is also allowed to learn from its mistakes). The system tries different alternatives, and is scored on each choice it makes. From this scoring, the system gradually learns what courses of actions are preferable to others. Figure 7 shows an example reinforcement learning system, that is learning which manoeuvres are most appropriate for different scenarios by being scored against previous choices.

Reinforcement learning can be used to teach a system how to make decisions about what courses of action to take (i.e. the ‘decision-making’ function in Figure 5) in different situations as long as a reasonable scoring system can be created.

Scenario	Manoeuvre	Score
P	Go right	8
Q	Straight on	4
R	Go left	7
:	:	:
:	:	:

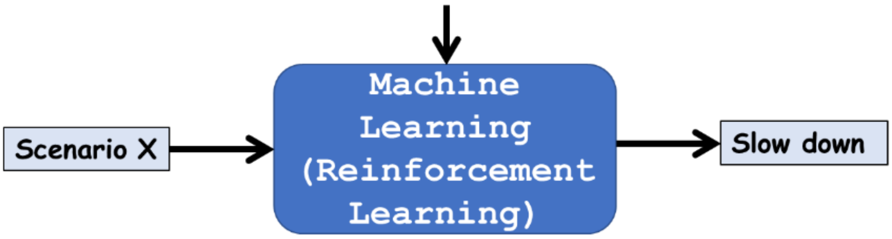

Figure 7: Prediction of Manoeuvres

Online and Offline Learning

We can teach our autonomous systems before they become operational (offline learning), and they can also learn from the real world while they are operational (online learning). Online learning is only useful for real-time systems if the system can learn quickly enough to also make decisions in real-time. A second constraint on online learning is that testing the behaviour of the system in advance is not possible, which could be a problem for safety-critical systems. Thus, for autonomous cars, offline learning that allows us to test the learnt behaviour before it goes live is normally preferable.

If we get our autonomous cars to record the situations they encounter (and we download it to a central database), we can use this information to continuously improve our machine learning systems (the more data we provide a machine learning system, the better it should get). Thus, many companies are gathering as much driving data as possible, so that they can use it to teach their machine learning systems. Data would normally be gathered at the end of each day over the internet when our autonomous car was parked in our garage, and, similarly, over-the-air updates could subsequently be sent to the vehicles with improved software.

Deep Learning

Deep learning looks like it will be the most widely-used form of machine learning used in autonomous cars. It can be used for image processing, data fusion, and decision making. It can even be used across the whole stream of functions shown in **Error! Reference source not found.** (e.g. for avoiding crashing into a vehicle in front), although some argue that providing the training data for teaching a complete function of an autonomous car may be impractical.

One drawback with deep learning is that the resultant data structures and code can be extremely difficult to interpret and reason about. This makes the testing of these systems even more important than usual.

Next Part

Part 3 of this article 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] The Mobileye, Delphi Automated Driving Equation, (see https://www.autonews.com/assets/PDF/CA106710823.PDF), 2017

[2] IEEE802.11p ahead of LTE-V2V for safety applications, (see https://www.auto-talks.com/wp-content/uploads/2017/09/Whitepaper_LTE-V2V_USletter_05-1.pdf), 2017

[3] CAR 2 CAR Communication Consortium (C2C-CC) (see https://www.car-2-car.org)