Driverless Vehicle Security: Considering Potential Attacks and Countermeasures for Military Applications

Nicola Bates

Technical Report

RHUL–ISG–2020–1

22 June 2020

Student Number: 100905318
Nicola Bates

Driverless Vehicle Security: Considering Potential Attacks and Countermeasures for Military Applications

Supervisor: Dr Raja Naeem Akram

Submitted as part of the requirements for the award of the
MSc in Information Security
At Royal Holloway, University of London

I declare that this assignment is all my own work and that I have acknowledged all quotations from published or unpublished work of other people. I also declare that I have read the statement on plagiarism in Section 1 of the Regulations Governing Examination and Assessment Offences, and in accordance with these regulations I submit this project report as my own work.

Signature:

Date:

Acknowledgements

A special thank you to Dr Raja Naeem Akram who has provided so much inspiration and has made completing this project an absolute joy. Also, thanks to Peter Davies and Professor Siraj Shaikh who gave up their time to discuss driverless vehicle security, David Burge and Jannette Van Halteren who kindly helped in proof reading and Simon Butler who shared his 2017 MSc project. Last but not least, thank you to my husband Dr Robert Bates who has been there for me throughout my MSc.

Table of Contents   **Page**

Table of Contents (continued) **Page**

List of figures and tables **Page**

Executive Summary

There is much interest from politicians and mainstream press about the driverless vehicle with a host of interested parties fervently researching into this technology. Autonomous Vehicles (AVs) promise mitigation of accidents, reduction in greenhouse gas emissions and more efficient use of infrastructure within the civilian world, as well as opportunities within the military to reduce exposure of troops in warzones. The automotive industry has a lot of work to do to secure these systems, however, if AVs are going to be seen on the road or deployed in a warzone in the not too distant future.

A modern vehicle is an extremely complicated cyber physical system requiring effective operation of between 70 and 100 Electronic Control Units (ECUs) to maintain function and safety governed by around 100 million lines of code [1]. This number is likely to grow by an order of magnitude with AVs with all systems needing to be robust and free from defects. Increased connectivity combined with autonomous functions allow for remote access for malicious hackers posing a considerable counterbalance to the socioeconomic benefits offered.

A major area which needs to be resolved may not be technological, however, but psychological. How civilians will take to this technology once available may be one of the biggest challenges, as negative reactions over the Waymo AVs in Phoenix, Arizona have shown [2]. Back in the 1900's 'self-driving elevators' were invented but widespread adoption was slowed for decades through people refusing to use them [3]. Within a civilian setting there are additional requirements for clear legislation and regulation with insurance liability issues to be resolved.

There are predictions that the army will get AV technology before cities do [4]. A military environment is more flexible in terms of regulations with the arena generally only governed by rules of warfare, which mainly cover humanitarian issues [5]. Added to reduced legislation in the military is the ability to order troops to work with autonomy without having a choice in the matter. In 2016 the US Department of Defence (DoD) stated they will "…exploit all advances in artificial intelligence and autonomy and insert them into DoD's battle networks" with the Pentagon budget released in September 2018 allocating $2 billion over 5 years for artificial intelligence technologies [6].

The purpose of this paper is to describe how AVs work and produce a comprehensive review of both realised and theoretical attacks within the civilian domain. This will give a view as to the types of vulnerabilities which may exist in a military setting, specifically supply line vehicles operating in desert warzone environment. A risk assessment will then be completed which takes as input attacks which would achieve specific enemy objectives. Using the outcomes of this work recommendations will be proposed for how the high and medium level risks identified could be mitigated against.

Analysis highlights the differences between technology requirements for a civilian setting and that which would be needed in a military arena. One of the highest rated attacks simply involved a person walking in front of a military AV to make it stop and allow its capture. In a civilian setting this would be an essential feature and would save many lives. However, in a military scenario this has the potential to cost many lives and allow an enemy to capture the AV – and the associated mission data and autonomous technology.

AVs will need the highest level of security with a 'secure by design' mindset being adopted, rather than adding on features to an existing vehicle, with security being included from the design phase. However, with military vehicles having a lifetime of around 20 years and with model changes approximately every 30 years [7] the ability to update security technology is more restrictive.

A key finding from the report is that by linking all vehicle systems through the Controller Area Network (CAN) bus gives the opportunity for a minor component to enable compromise of safety critical devices. The infotainment system can not only be used to leak troop discussions and vehicle movements but also connect to any ECU which is also connected to the CAN. Supply line AVs may not carry troops making this system redundant, and even simple changes such as troops having an infotainment system isolated from the safety critical devices would increase system security.

Fortunately for military AVs there exist a series of countermeasures and means of mitigating attacks which do not exist within the normal civilian space. In addition to removal of vulnerable systems frequent service schedules permit software updates through physical, not wireless measures which allow some of the most dangerous attack surfaces to be removed. The benefit of military budgets, during active conflict, permits a luxury civilian AVs will not be able to afford in terms of duplication of sensors and systems creating levels of redundancy which can prevent all but the most sophisticated spoofing attacks.

Finally, whilst not recommended in this review, the nature of a military setting would permit an AV to be destroyed to defend against its capture. Should there be clear signs it was operating outside of critical parameters or capture was known, self-destruction would be a viable option, with commanders preferring the AV be destroyed than giving the enemy valuable information.

1. Introduction and background

1.1 Vehicle basics and history

The origins of the modern motor vehicle can be traced to patent number 37435 filed on 29 January 1886 for a "vehicle with gas engine operation" [8]. However, it is only in 1908 with the mass production processes pioneered by Henry Ford did car ownership become within reach of the middle classes with the Model T-Ford [9].

These early cars were fully operated by a human driver. The accelerator pedal was connected to the throttle valve which regulated the air flow into the engine, an increased air flow caused by pressing on the pedal harder gave an increase in car speed. Steering was through direct manual rotation of the steering wheel, which is converted into a swivelling movement of the road wheels. In early cars breaking was done through a wooden block moving against the wheel and then via mechanical drum brakes from 1902 [10]. From 1918 the concept of a hydraulic braking mechanism was proposed using fluid to transfer force from the brake pedal to the discs on the wheels when the brake pedal was pressed [10].

Up until the latter half of the twentieth century little of the basic concepts of automobile mechanics changed. Then came the advent of driver assist technologies which aimed to make driving both easier and safer. This started off with the simplicity of power steering, commercially available from the early 1950's [10], moving onto Antilock Braking Systems (ABS), which were developed throughout the 1970's [10], and then to cruise control, put into vehicle mass production in the mid 1980's [10]. Electronic stability control, which improves vehicle stability by detecting and reducing loss of traction, was also introduced by three automobile manufacturer models in 1995 [10].

Since 1975 the development of integrated circuits and microprocessors made it possible to mass produce driver assist ECUs [10]. These units cover: cruise control; power steering; engine management systems; ABS, stability controllers; and tyre pressure sensors. ECUs need inputs from other vehicle systems, as well as the driver, to produce scheduled outputs in real time. Due to the success of this technology modern cars can have up between 70 and 100 ECUs embedded [1].

ECUs have been merged into one system which in the majority of cases are linked with a Controller Area Network (CAN) bus system, allowing communication between all parts of the vehicle. The need for simplicity of driver control has seen these CAN bus systems combined with the infotainment system along with external means of communications such as Wi-Fi.

The emergence of autonomous driving features has been developing in more recent years with such vehicles already operating on the streets of America.

1.2 Levels of autonomy

With driverless technology, and driver assist becoming more commonplace it has led to the development of a scale which defines more formally the levels of autonomy. This provides a unified approach to levels of driving features, as shown in table 1.

Table 1: SAE levels of autonomy [11]

As can been seen: levels 0-2 require a human driver to monitor the driving environment (driver support features) with levels 3-4 consisting of an automated driving system to monitor the driving environment.

This report will not be considering levels of autonomy 0-2 because these are not considered as autonomous driving. The focus will be on level 5, with the system taking on all the driving modes including execution of steering, acceleration and deceleration, monitoring of driving environment and fall-back performance of dynamic driving tasks. In this report AVs will refer to a fully autonomous system at level 5, unless otherwise specified.

Currently, as at July 2019, fully autonomous cars are not yet available for the public to purchase. However, designs such as the Tesla Model S can be bought with all the hardware needed for fully autonomous driving which can be 'switched on' when available and legally allowable [12].

Waymo, have been working on self-driving technology since 2009 and have a global fleet of around 600 vehicles [13], the majority of which are in Phoenix, Arizona. The streets of Phoenix have been mapped for over two years allowing these cars to be used as a taxi service, however, they are required to have a safety driver behind the wheel [13], thus limiting them from achieving true level 5 autonomy.

1.3 Key similarities and differences in military and civilian autonomy

This report considers potential attacks and countermeasures on AVs in the context of a military environment, I will however be using civilian uses as a base of knowledge to complete this. For effective mapping between civilian and military AV domains there is a need to understand the key similarities and differences between these two use cases.

1.3.1 Key similarities

Interoperability: A civilian owner will expect to be able to travel in the same frictionless way with an AV as they currently do with a conventional vehicle. If an AV owner could not drive freely from Texas to New York, or in Europe, for example this would not be a desirable situation. Interoperability issues from external signalling and regulations in safety standards between areas would need to be in place before this could be a reality. Having system unable to communicate with each other would create a barrier to ownership and use of an AV.

Within the military there is the need to work closely with allied nations in order to coordinate fighting against a common enemy. The US, UK, Canada, France and Germany for example could be working closely on the front line in enemy territory and their equipment would need to be interoperable for effective working and to prevent confusion and friendly fire. Some countries such as the UK would have a 'sovereign' capability in addition to interoperable devices which are deliberately isolated [7]. This is seen most commonly in weapons systems for example.

Attack resilience: In both a civilian and military setting the ability to attack an AV in the same way as a conventional vehicle still exits. Throughout this report level 5 autonomy will be the primary focus of attacks but just because fully autonomous features have been added it does not take away conventional attacks such as physical damage to the car.

From a cyber security point of view both civilian and military vehicles are susceptible to over the air attacks and will have a number of attack surfaces where systems directly interact with the external environment. Both will require extensive communication between different components, some of which will be more vulnerable to attack than others, but all of which require securing to prevent malicious attacks.

Privacy: In order to achieve level 5 autonomy the vehicle needs to know where you are and where you are heading at all times of operation. The wealth of sensors will also be recording information about the external environments, including images of passers-by and potentially the location and identity of other autonomous vehicle users. Voice activation, if used, will require 'always on' microphones which have the potential to record general conversation and information.

With the introduction of the General Data Protection Regulations (GDPR) legislation in Europe in 2018 [14] and similar legal regulation for general privacy being considered around the world the need to protect personal data has never been greater. This has been even more urgent since consultancy firm Cambridge Analytica hit the headlines with its data harvesting tactics from Facebook [15].

In the military it is obviously essential to keep data private in order to prevent the enemy knowing your location history and details of vehicle occupants for example. However, in both civilian and military cases the need to protect data is of high priority and systems need to be secure to prevent unauthorised access. There needs to be efficient ways of erasing all data, whilst not impacting the ability of the system to learn for example.

1.3.2 Key differences

Environment: In both civilian and military settings the environment can obviously be wide open roads. However, if we concentrate on the more extreme environments differences emerge. In a civilian setting manoeuvres through large cities which can have very narrow roads, traffic signals, pedestrians, cyclists, animals, and other road markings to process and interpret are commonplace.

Within the military the terrain maybe unmapped, uneven, with changing routes due to enemy movements or artillery damage. Lots of journeys would be 'off road' with the First Gulf War across desert and in the Balkan war it was more secure to go off road rather than use existing paths [7]. A military setting would be inhospitable and hostile with no recovery services to call upon if help were needed due to technical or mechanical problems.

Specialised: Although civilian vehicles will be specialised with designs such as people carriers, sports edition, vans and trucks; the basic mechanisms and terrains will be of similar nature. However, in a military arena the equipment could be highly specialised and require complete redesign due to the niche operation. Consider for example the differences needed in bomb disposal mechanisms versus reconnaissance vehicles. Standard safety designs such as air bags are also not included in all terrain vehicles whilst roll cages and emergency exits are fitted in addition [7].

Attack Threat: For civilian uses, the biggest risk to life currently on the roads are through human driver error, causing over 90% of all accidents [16], with deliberate attacks a rarity. However, deliberate attacks on military vehicles are often a major focus of enemy combatants. Consequently, a risk to life from not only human errors but deliberate attacks, both physical and through cyber methods, is far higher. The UK is very risk averse so analysis for risk to life will be significant if the vehicles are to carry personnel as opposed to Unmanned Aerial Vehicles (UAV) [7].

Costs: In order to optimise the civilian use of AVs, it has been stated they need to be the main users of the road, replacing conventional vehicles, which currently number 1.2 billion [17]. Getting into figures of tens of millions, gives economies of scale and thus price reductions in the cost per vehicle for developing AV technology.

Within a military setting, even taking into account the use of swarming technologies and smaller vehicles, there is not expected to be such high numbers. The US currently have around 40,000 armoured response vehicles [18] for example and the Defence Advanced Research Projects Agency's (DARPA's) Offensive Swarm Enabled-Tactics (OFFSET) program has the aim of "using swarms comprising upwards of 250 unmanned aircraft systems and/or unmanned ground systems…" [19]. This indicates tens of thousands of units for military settings rather than many millions in the future. In this instance development costs are borne by far fewer units, adding to this the variety of niche uses which further increases per vehicle cost.

In the UK military over the past decade austerity measures have impacted spending. If projects are sponsored under a 'urgent operational requirement' (UOR) this is funded by current conflicts as opposed to 'core' spend. This UOR is likely to be where AV technology advances in leaps with research and development spending expedited. However, after conflict UOR equipment is unfunded which could leave lots of AVs unsupported unless it gets taken into core budget, impacting spending elsewhere [7].

Life Expectancy: A civilian vehicle has an average expected lifetime of approximately 11-12 years with models changing every 5-7 years, which allows incremental improvements and new technologies to be periodically added to a fleet en-mass. In the military these figures are closer to a 20-year lifetime with model changes closer to 30 years [7]. This has implications for the ability to update the mechanics and technology used in the military AV compared to the civilian AV. The increased lifetimes and model changes are likely to result in retrofitting within a military setting as opposed to security by design being baked in from the outset.

When a vehicle is retrofitted it is stripped down and updated but this will not be all at once leaving legacy and interoperability issues which causes problems. Even if equipment is brand new there will be the need for it to be interoperable with existing technology which results in its security being affected to ensure it works alongside other 'onboard kit'. Given the reluctance to take equipment 'offline' for any amount of

time it may not even get a refit leaving it running full time. Conversely, military vehicles can also be parked up and left for months or years because of lack of staff. Equipment management is a constant problem in the military with the evolution cycle less cumbersome for civilian vehicles [7].

1.4 Military uses of autonomy

Core US military doctrine define at least five 'domains of warfare', being: land, sea, air, space and cyberspace [20]. This project will focus on the land domain, but obviously cyberspace will be included as well as space due to the uses of Global Positioning System (GPS) for location data.

The rationale for this choice is due to the amount of activity in this domain from large civilian automotive players, namely Tesla, Google and Apple, as well as university research which will give the information needed for mapping to a military setting. Obviously work in autonomy is also available in avionics, and indeed can be argued to be more mature with the Airforce having more of a reach than any other service too. But the persistence needed to effect change is done by having 'boots on the ground' with the land domain having always been in the hearts and minds of military commanders. It has been stated that "no commander has ever said they have captured territory simply by flying over it" [20], so the land arena of operation can be argued to be the most critical to success.

This project will consider the vulnerabilities and applications of AVs where lethal force will not be utilised. The use of lethal force has been the topic of significant debate with over 5 years of discussion at the United Nations; and many organisations and thousands of scientists having spoken out against lethal autonomy [21]. The current view of the UK military is "…the application of lethal force must be directed by a human, and that a human will always be accountable for the decision." [22], so called 'human in the loop' decision making.

Focussing in on the land arena, ground vehicles in use within the military include fighting platforms, troop carriers, medical platforms, reconnaissance vehicles, command and control vehicles, supply-line vehicles, extraction platforms and explosive device disposal [7]. All of these areas aim to replace soldiers with machines in order to prevent as much loss to life as possible in the field and improve operations.

Project work will focus upon supply line AVs, responsible for getting food, ammunition, troops and other equipment to the front line. It has been said that much of a war is concerned with logistics and many a war lost and won due to supply chain issue [20] so this is a crucial area of focus. Perhaps because of this, and obvious uses in commercial logistical operations, there has also been much work completed in the area of platooning in supply chain scenarios, which will assist with available

information. Supply chain AVs are likely to be in less hostile territory that other areas of use, so more likely to be deployed ahead of more hostile areas of use where a vehicle could be lost to the enemy if problems occurred.

For the land based, supply-line focus, I will also add that my work will relate to supplies in enemy territory rather than in a home country scenario. The risk factors for home supply chain would be more related to a civilian study rather than a military analysis, which would defeat the purpose of the projects aims. It will be assumed supplies will be through desert terrain, so will have uneven ground but fewer obstructions found in built up, mountainous or forested areas which will keep things simpler to review. In addition, recent wars, such as in Iraq, have been fought in these desert terrains so are a useful scenario to be studying in terms of applicability. It will be assumed at the start and end of the supply chain AVs could be built up areas which is a more realistic situation.

1.5 Summary and methodology

Having considered the various areas associated with military AV applications, I will focus my attention on the cyber security of autonomous supply chain land vehicles which will be operating in a warzone with mainly desert terrain but could include inhabited areas such as in the wars in Iraq and Afghanistan. Obviously traditional physical attacks will still be possible, however these will not be the main focus of this report, with level 5 autonomy specifically being analysed, unless stated otherwise.

To consider potential attacks and countermeasures for military applications the area of civilian autonomy will be heavily drawn upon given the confidentiality and lack of published material in the military domain. This information will mainly be from journal papers but also through meeting with experts in the field of AVs, academics and other publications. The attack surfaces of AVs will be identified along with a review of published attacks. This analysis will be completed in a logical ordering through considering a car owners journey to ensure full coverage of AV weaknesses.

The analysis of civilian and military similarities and differences will allow consideration of how the change in reference frame alters needs and security challenges. This will allow a risk assessment within a military setting to be completed which will indicate where connected features are applicable and useful, where these could be eliminated and areas in which extra security may be required. Analysis will also feed into determining countermeasures which could be employed to reduce the risks identified.

When considering the military setting, risk assessment and countermeasures will be conducted by considering the objectives of the enemy, which could range from individual AV capture to AV surveillance or even whole fleet infection with malware. Known civilian attacks will be considered which could be applied in a military setting to meet the enemy's objectives.

2. Devices in an autonomous vehicle

2.1 Introduction

In order for a vehicle to be autonomous its systems must be designed to drive the vehicle under all conditions and not require human intervention. To make this a reality requires a complex mix of sensors and actuators continually recording and responding to the local environment, controlled by many millions of lines of computer code to interpret these and act upon the data received. This creates a fundamental problem given there are over 100 million lines of code in a driverless vehicle [1]. With Carnegie Mellon University research putting the average rate of defect per million line of code at 6,000 [23], this gives 600,000 defects per vehicle, which could create issues from simple inefficiency to dropped commands and safety critical failure.

In this report the coding aspect of AVs is not covered, however, it is recommended that to prevent a single point of failure a variety of algorithms be used rather than a single master design. By having multiple variants of an operating system vehicles on the road, or 'herd', are better able to cope with failure of a single variant if a high percentage remains unaffected [24]. For example, if one version in ten failed, the remaining 90% could likely function around those vehicles. Should there only be two variants and one failed, with 50% displaying potentially erratic behaviour, those impacted would struggle to cope, much like vaccinations providing herd immunity in human populations.

The challenge for AV reliability and security is not just limited to coding, however, in addition the components within the critical systems are crucial. If even a small element were to fail in the wrong place this could cause the whole system to also fail.

In terms of physical AV risks, the supply chain is an obvious source of weakness where a defect in a product from a single supplier could disable the whole fleet of AVs. The UK military for example like to have one supplier with certified security assessed and built in. A way to mitigate this risk is to use of a variety of suppliers for components in vehicle assembly, therefore spreading the risk of component faults across multiple suppliers. Multiple suppliers of chips for the electronic components could prove difficult given there are very few chip manufacturers in the world.

Electronic chips would be used for the ECUs which are embedded systems controlling electrical systems in the AV. There are around 70 to 100 of these ECUs in modern vehicles [1], a variety of which are shown in figure 1.

To review the different functions incorporated within AVs we will follow the flow of user interaction with the vehicle. This starts with getting access and starting the engine, followed by the AV finding its position both locally and globally. An AV would then proceed to move position, velocity and orientation, with the various ECUs sensing and

feeding back on vehicle and driving conditions. The AV would be communicating with its surroundings including other AVs on the road and maybe roadside devices, giving the vehicle occupants a choice of entertainment systems to occupy themselves with when they do not have to focus on driving.
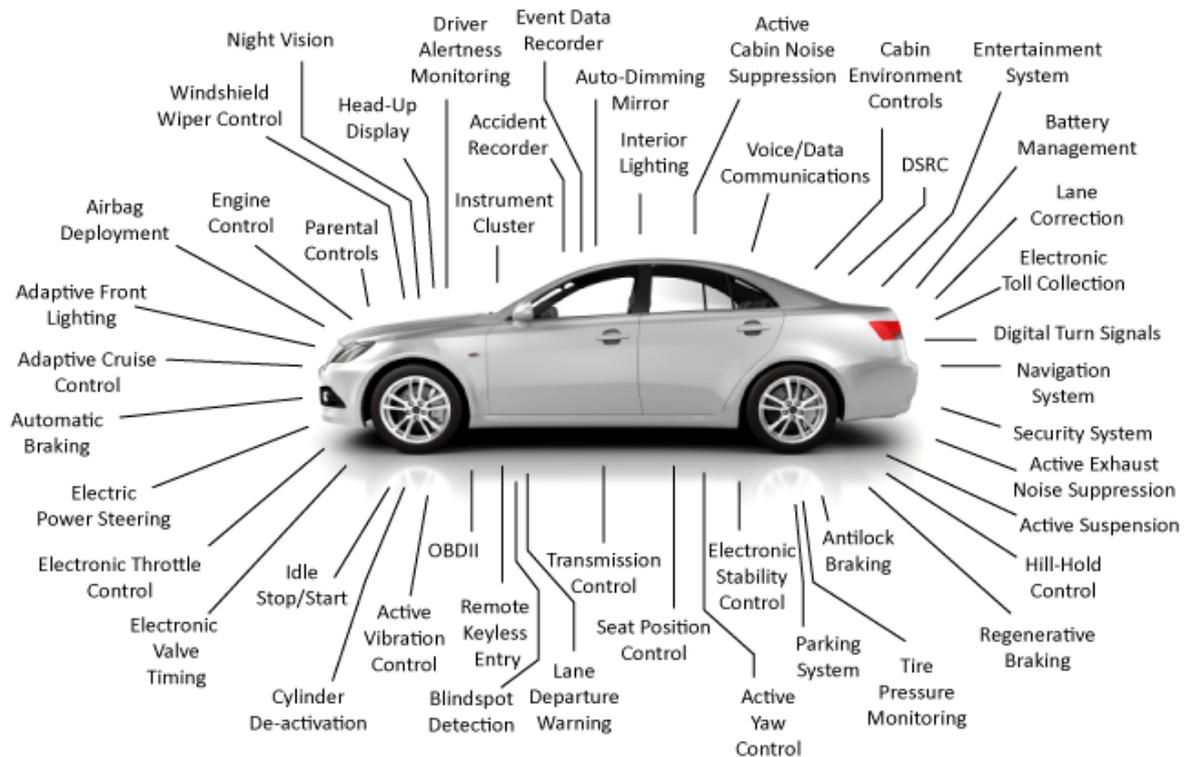
Figure 1: Vehicle ECU's [25]

A summary of the CAN bus which pulls all of these ECU devices together will then be undertaken, followed by maintenance and update mechanisms for the AV which will need to be completed at various points in its lifetime.

2.2 Access and ignition

In early cars, a conventional metal key was used to access the vehicle and to start the engine. An additional level of security, to prevent criminals picking a lock, was to add an 'immobiliser' to the key - essentially a transponder chip which was required to start the engine. This prevented the engine running unless the correct token was present and avoided thefts from a vehicle being 'hot-wired'.

Nowadays mechanical keys and locks have been replaced by transponders and software, allowing users to remotely open and start the vehicle. The key contains a button, which when pushed sends a coded signal via radio waves to a receiver unit in the vehicle which unlocks it. Once inside the key connects to the vehicles computer system by sending a low frequency signal allowing 'push button' start.

2.3 Position

There are two basic things which are required in order for an AV to move around in space: it needs to know location information from both a global and a local perspective.

2.3.1 Global position

GPS: The most common method to find the location of something globally is GPS, which is a network of approximately 30 satellites which are in a geostationary orbit above the Earth [10]. These satellites send out radio signals constantly which a GPS receiver will listen out for and use to calculate its relative distance from multiple satellites, based on the time taken for the message to arrive travelling at the speed of light. Once the GPS receiver has calculated the distance from four or more GPS satellites it can use the data to calculate where on Earth the receiver is.

For GPS to work as a navigational aid there needs to already be maps in existence of the area. For most of the world there exist constructed maps using either normal cartographic methods or, more recently, from aerial and satellite photography. Such maps need to be kept updated in order to capture new road changes and layouts. Well documented reports of satellite navigation systems sending people the wrong way and into hazards clearly shows the necessity of these updates [26].

Accuracy of GPS defines the range of locations which a user could be in from the data gathered, as shown in figure 2. It depends on many things such as signal blockage, atmospheric conditions, signals reflected off of buildings or walls (as shown in figure 3) and receiver quality. For example, on flat terrain a smartphone GPS is typically accurate to 4.9 metres [27], which is good enough for user guidance where an individual can interpret the location by their surroundings but less suitable if it is the only source of position data for an AV.

In order to achieve centimetre levels of accuracy distortion from the Earth's atmosphere are corrected for by using with high-end dual-frequency receivers where observing two GPS frequencies increases accuracy. The size and cost of this system means it is out of reach of most civilian uses but is used in military scenarios.

Another factor in the accuracy of GPS is radio interference and spoofing or jamming of signals. Within the military there is the option to have encrypted GPS and authentication mechanisms which offers an increased level of security to spoofing. Encryption, however, shortens the signal making it more cumbersome requiring more rebroadcasting stations [7] and it will also not secure against jamming attacks. Attack types on GPS will be discussed further in section 3.3.1.
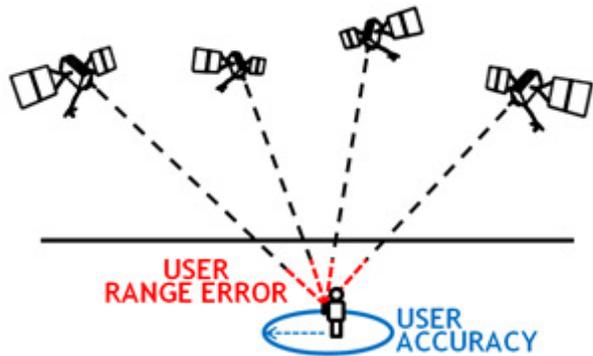
Figure 2: GPS user accuracy [27]   Figure 3: GPS signal reflection [27]
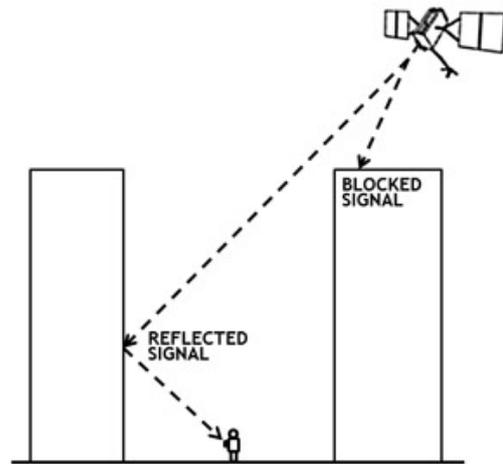
Inertial Navigation System (INS): INS is an alternative means to GPS for determining global location. This system uses gyroscopes and accelerometers to calculate a position relative to an initial inputted starting point. When used with GPS this system gives redundancy and allows location to be determined if GPS communication signals are lost for example. Use of INS for position, velocity and orientation monitoring are covered in more detail in section 2.4.

Mobile mapping: Waymo's approach to getting AVs on the road has been to build up very accurate maps of a town before their vehicles drive through it [13]. To complete detailed digital maps of an area a mobile vehicle has to collect data with a range of sensors which include cameras, Radio Detection And Ranging (RADAR), laser and Light Detection And Ranging (LiDAR).

Mobile mapping technology would be unsuitable for use in an active military domain: the hostile environment and the long timescales needed to build detailed maps make it difficult to do safely. Even in a desert situation the terrain is constantly changing with storms causing reformation of sand structures and the enemy adding road blocks, Improvised Explosive Devices (IED's) or even landmines (if not signatures to the Ottowa Convention) [7], which would block previously safe mapped routes. Further work on mobile maps will therefore not be completed.

2.3.2 Local position

Within the modern vehicle there are multiple sensors allowing the car to map its local surroundings. The configuration of sensors on the Tesla Model S, which is designed to offer level 5 autonomy by activation of a switch [12], are shown in figure 4. This design includes eight cameras and twelve ultrasonic sensors, both providing full 360-degree sight, along with a forward-facing radar which is able to see in poor visibility.
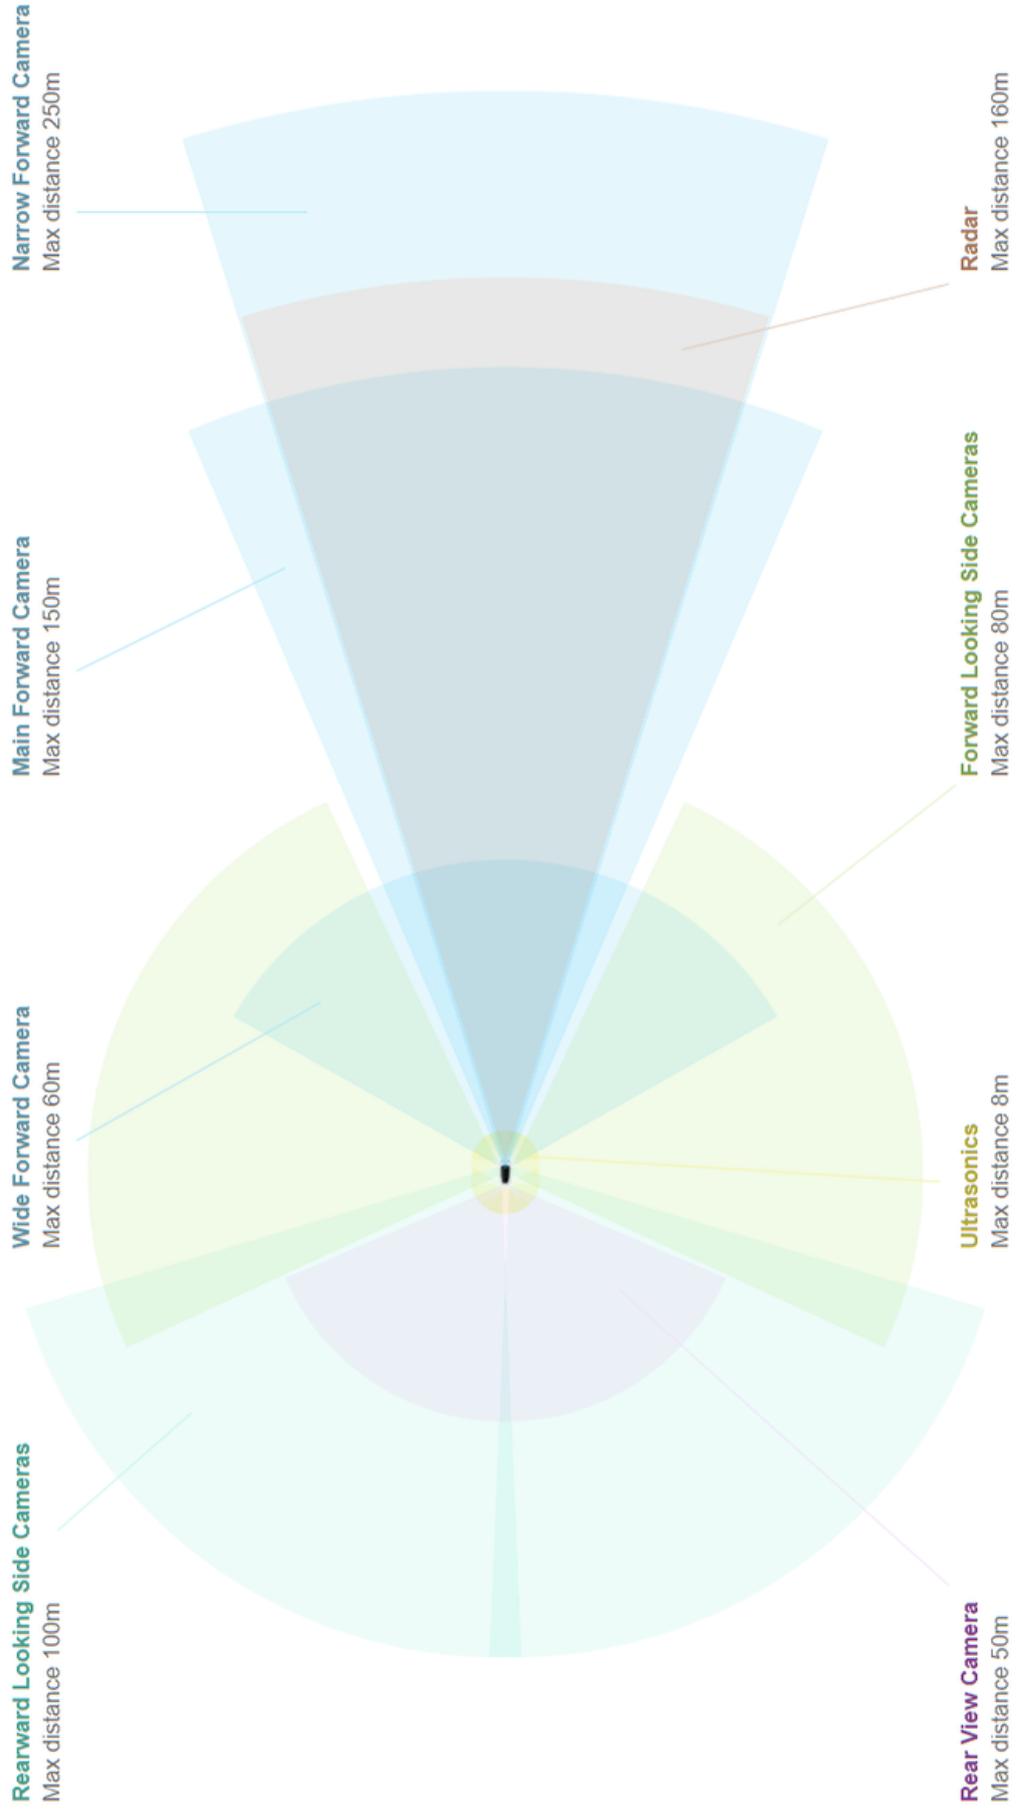
Figure 4: Overview of the Tesla vision system [12]

Narrow Forward Camera
Max distance 250m

Main Forward Camera
Max distance 150m

Wide Forward Camera
Max distance 60m

Rearward Looking Side Cameras
Max distance 100m

Rear View Camera
Max distance 50m

Ultrasonics
Max distance 8m

Forward Looking Side Cameras
Max distance 80m

Radar
Max distance 160m

The Tesla system does not use LiDAR technologies, which is another way of mapping surroundings. However, Waymo has been using LiDAR located on the top of its cars since the early days of its projects [13].

Cameras: The five forward cameras allow appropriate viewing of the road based upon different speeds of the vehicle, road infrastructure and traffic flow. In the designed configuration there is also the added benefit of built in redundancy if faults occur with cameras even able to serve as backup to the radar system. The overlapping of camera fields of vision also gives depth perception capabilities.

The narrow forward camera can see up to 250 metres ahead, which, at the highest speed of 70mph on a UK road, would give an 8 second window of visibility. A main forward camera, with sight up to 150 metres, is used to see the edge of an intersection coming up to a slip road or someone waiting at a junction, for example.

The wide forward camera, with a 60-metre viewing distance, has a wider field of view and is used to observe road signs; obstacles, such as roundabouts; and monitoring traffic flow, allowing the vehicle to react accordingly. A forward-looking side camera has coverage up to 80 metres and gives more resolution than the wide forward camera.

A rear-view camera observes up to 50 metres and is mainly for reversing, so low speed manoeuvrings. The two rearward looking side cameras, with sight up to 100 metres, allows appropriate checking of other vehicles approaching from the rear before changing lanes, to ensure enough space before pulling out. Again, having multiple cameras not only accounts for vehicle speed but also gives redundancy and depth perception.

Radar: Radar works by sending out radio waves and sensing their reflections, which it then uses to map out its surroundings. It is a more desirable sensor for measuring distance than a camera and is not affected by adverse weather conditions (like fog, snowstorms or sandstorms), darkness or glare, which would make visibility difficult for not only cameras but for a human driver.

Cameras are also affected by the material from which an object is made which can give very different reflective properties for light but less so for radar. A person standing on the side of a road in a black matt finish coat for example would not be very visible to light but if they were wearing a reflective high visibility jacket they would be highly visible. With radar both of these situations would be seen equally offering greater safety when light levels are low.

Radar is also the most reliable method to determine if objects are present at further distances, a task sonar would not be able to do. The radar system used on the Tesla can see up to 160 metres ahead and act as a primary sensor to detect the vehicles

surrounding, offering vision capabilities which are above and beyond cameras and the human eye.

Ultrasonics: Ultrasonics work in a similar way to radar but instead of sending out radio waves it sends out ultrasound waves and through their reflections forms a picture of the surroundings. Also, like radar it offers above and beyond human eye and camera ability by being able to see in events such as adverse weather conditions, darkness or glare.

The ultrasonic sensors can detect objects such as humans and animals which again would be seen no matter what they were wearing or the lighting conditions. Sensors can detect objects in blind spots and offer assistance when changing lanes on a motorway to ensure a safe transition. Given its range of coverage it can also give redundancy if information from cameras is not available to the AV.

On the Tesla sonar can detect obstacles within an 8-metre radius over a full 360-degree angle. Whilst it can work at any vehicle speed if travelling at 70mph with just an 8-metre distance they would give just 0.26 seconds of visibility.

LiDAR: Waymo uses a medium range LiDAR system located on the top of its vehicles for its primary mapping sensor and is considered by some a key sensor needed to safely deploy AVs [29]. LiDAR uses differences in laser light return times to build up highly accurate three-dimensional representation of the AVs surroundings.

Figure 5 shows the outgoing laser beam, generated by the infrared transmitter diode, and the reflected echo which is angled into the photo diode receiver by a rotating mirror. On the top of this device is a motor with an angle encoder providing information about the motion of the device, which is processed into information such as object position, speed and distance. The rotation of the LiDAR allows a 360-degree image to be collected.
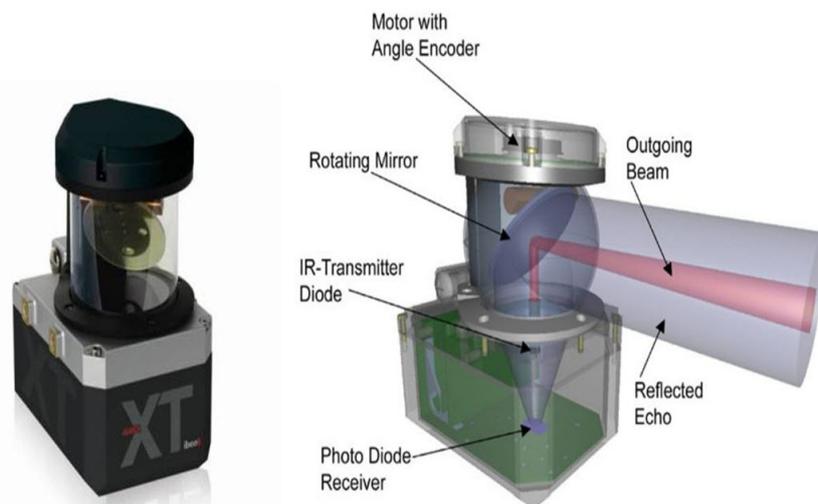
Figure 5: How LIDAR works [30]

Good resolution from a LiDAR system relies on the intensity of light being reflected and scattered from objects it hits. As the reflected and scattered beam will see the intensity reduce with distance from the source, both the resolution and quality of image will diminish with distance. Unfortunately, to overcome this the power of the beam cannot continually be increased to adjust for this, with wavelengths limited to ones which do not damage human retinas for safety reasons. However, given objects further out are less relevant to the immediate safety of the AV, this does not cause significant issues.

LiDAR generates the radiation source for measurement so works well in all light conditions which is a benefit over cameras and the human eye. It can determine velocity of objects using doppler shift, which is an extremely useful measurement in a driving environment. LiDAR is also more accurate at measuring angular changes and keeping sight of target vehicles on curves than radar and sonar systems [31].

The LiDAR system comes with a few disadvantages in that it can only detect objects located in the range of 30 metres to 200 metres [29], with close objects being hard to map. Its performance also deteriorates in fog, rain and dusty weather conditions which radar and sonar can perform in [29]. The cost of a LiDAR system is also more expensive than deployment of say a radar system, which is one of the reasons Tesla does not use LiDAR in its fleet [32].

Like all positional sensors discussed, LiDAR is best used alongside secondary sensors such as cameras and ultrasonics in order to provide an overall mapping functionality. Multiple use of technologies also gives redundancy if anything goes wrong in any of the systems, or conditions are such on the roads which causes some sensors to become ineffective.

2.4 Position, velocity and orientation changes

Wheel rotation: Early speedometers worked by using the rotation of the vehicle wheels. A magnet fitted to the wheel generated a current which was translated into a measure of the number of wheel rotations per second. This measurement could then be used along with the circumference of the tyre to calculate the distance travelled and the speed of travel. This is a basic mechanical calculation and measurements could vary by around 5% [33] depending on the amount the tyres are inflated by for example. Within an AV, these sensors could be used as a check on other devices, such as GPS. Given the sensors do not have to rely on any external signals and only rely on calibration they work even with no communication channels to the vehicle.

GPS: As discussed in section 2.3.1, GPS can be used as a way to map changes in location at a global level. This could also give information on relative position, velocity

and orientation changes although would suffer from accuracy issues at low velocity due to its 4.9 metre resolution [27], as well as jamming and spoofing issues. GPS would work well with other systems, however, and give potential redundancy and validation in systems.

INS: An INS can be used to track movements of an AV and typically contains three gyroscopes and three accelerometers [10]. Gyroscopes are rotation sensors used to measure angular velocity so as to give information on what direction the car is facing. Accelerometers are motion sensor to measure linear acceleration to give information on how the vehicle is accelerating relative to itself, for example either forwards, backward, left, right, up or down.

After an initial input to give the current position, velocity and orientation of the AV it is then possible to track its movements by continually computing the distance from the known starting parameters. This uses both measurements of angular velocity and linear acceleration of the AV measured relative to the moving system in a process of 'dead reckoning', illustrated in figure 6.
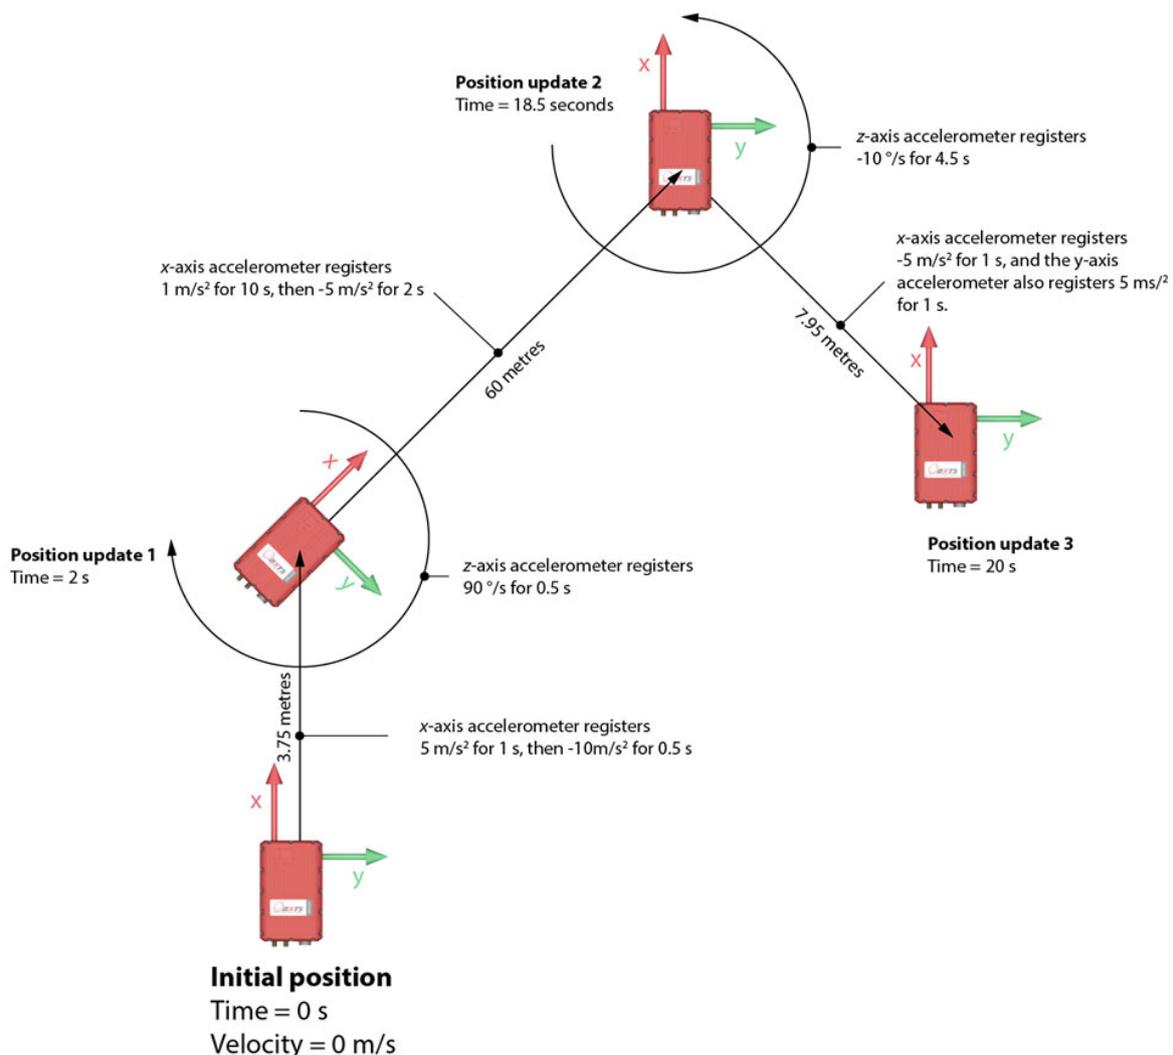
Figure 6: Dead reckoning calculations in an INS [16]

A key benefit of inertial navigation is after feeding in initial data the AV does not need external communication to obtain position, velocity or orientation. A downside of INS is that it suffers from small errors in the measurement of acceleration and angular velocity building up into progressively larger errors in velocity and position [10]. Since the new position is calculated from the previous calculated position and the measured acceleration and angular velocity, the position must be periodically corrected by input from some other type of navigation system.

Due to error build-up inertial navigation is usually used as a part of another system to give higher accuracy than a single system can provide and also to give redundancy if another system fails. Given developments in technology it is now possible to manufacture small and light INSs which makes this more of a realistic proposal.

2.5 Functionality and usability

Within a modern vehicle there are devices which act to inform the driver about conditions which could cause functionality to be impaired - the vehicle's maintenance and warning signals. There are also devices which act as aids to either enhance the driving experience or make it safer; these tend to interact with systems such as engine control, automatic gear box management, Antilock Braking System (ABS) and electronic stability control.

2.5.1 Vehicle maintenance and warning signals

Tyre Pressure Monitoring System (TPMS): A TPMS has a sensor valve in each tyre which monitors the pressure and reports back on low or imbalanced pressures to the dashboard. All new vehicles manufactured since 2014 have TPMS integrated as standard to ensure drivers are aware of the state of their tyres [35].

Other sensors: There are various other sensors on a modern vehicle such as seat belt monitors to make sure an occupant is correctly buckled up. External temperature sensors and moisture sensors inform of potential icy or wet conditions respectively, with sensors for low oil monitoring and engine maintenance to name just a few; there are too many to list all of them.

These sensors perform a similar role of feeding back to the driver, usually via the dashboard or audio warning signals, a condition which requires attention. This could result in the driver performing an action (fill the vehicle with oil, putting on a seatbelt) or adjusting an action (drive slower because the road may be slippery).

As vehicles move closer and eventually reach level 5 autonomy elements of these sensors move from being simply information for a driver to act on to being data an AV would use to adjust an action. For example, it could be if an occupant in the car is

sensed through pressure monitors and the seatbelt has not been engaged the AV will not move or if the temperature drops below zero the AV speed is limited due to the higher risk of ice on the road surface.

2.5.2 Driving aids

Engine control: The engine is what powers the vehicle. An engine management system is controlled by a series of actuators on the engine which feed in more or less energy (fuel) to either increase or decrease the speed of the vehicle. These units constantly monitor the engines performance, adjusting timings of actions and response. Subtle changes here can significantly impact the performance of a vehicle.

Automatic gear box: Automatic gear boxes simplify driving by selecting the most appropriate gear ratio for the speed and present acceleration of the vehicle.

Within conventional vehicles settings are available to the driver allowing them to select different 'modes' in order to balance out fuel efficiency, speed and acceleration. These modes control how quickly the vehicle moves up through the gears and also the amount of fuel injected into the engine when the accelerator is pressed. For example, sports mode would result in more fuel being injected into the engine, hence more power output, and slower progression through the gears, leading to more rapid acceleration. Although level 5 autonomy would not involve a driver the AV would be making the decision as to what 'mode' the car would be best driven in.

ABS: ABS is an automated safety feature on modern vehicles which reduces the braking pressure when the wheel locks, allowing the wheels to retain traction with the road. This improves steering control and reduces stopping distances on most road surfaces with the changes in braking activated at a much faster rate than most drivers could achieve.

Electronic stability control: Electronic stability control systems detect skidding caused by a loss of traction with the road surface. The system independently monitors each wheel for traction and automatically applies different braking forces or readjusts the power to individual wheels to correct oversteer or understeer scenarios. This helps retain traction and direct the vehicle where the driver is steering. Stability control can also detect the roll of a vehicle and check if the suspension is correctly balanced.

Numerous studies have concluded this technology is very effective in maintaining control of a vehicle and helps in preventing accidents. Since 2012 and 2014 these systems have become mandatory within the US [37] and the European Union (EU) [38] respectively.

2.6 Communication

Telematics is a term which covers all the communication which takes place to and from the AV. Communication within the vehicle by transmitting data collected by sensors have been discussed in sections 2.3 and 2.5. This section will focus externally on so called Vehicle to Everything (V2X) communication, covering Vehicle to Vehicle (V2V), Vehicle to Infrastructure (V2I), Vehicle to Pedestrian (V2P) and Vehicle to Network (V2N) communication. The linkages between these different areas is shown in figure 7.
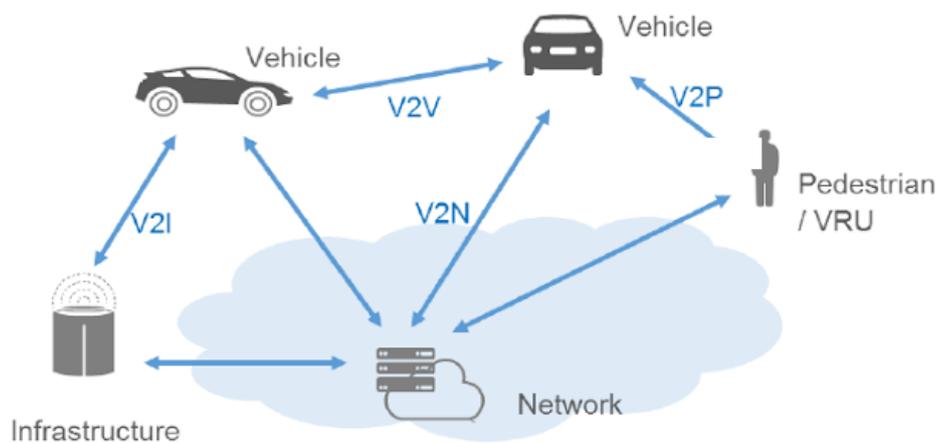
Figure 7: V2X communication types [39]

2.6.1 Communication with other vehicles

The use of V2V communication can be used to form different vehicle formations, namely platooning and swarming, which can offer increased efficiencies over traditional single car journeys.

Platooning: Platooning is a method in which a group of AVs can be driven together using electronic linking. Automated roads could mean AVs being able to organise themselves into platoons with wireless signals between them communicating on their status. Figure 8 gives details of how platooning is achieved.

Platooning would result in vehicles able to be driven closer together with the reaction distance required for the human driver eliminated, increasing space available on the roads. Better traffic flow at faster speeds with fewer collisions along with shorter distances between vehicles achievable through platooning would result in a reduction in road congestion. The ability to platoon with other vehicles would mean AVs not at full autonomy (level 5 in table 1) could be mostly unattended. Another benefit is reduced air resistance giving increased fuel economy and therefore cost savings and environmental benefits.
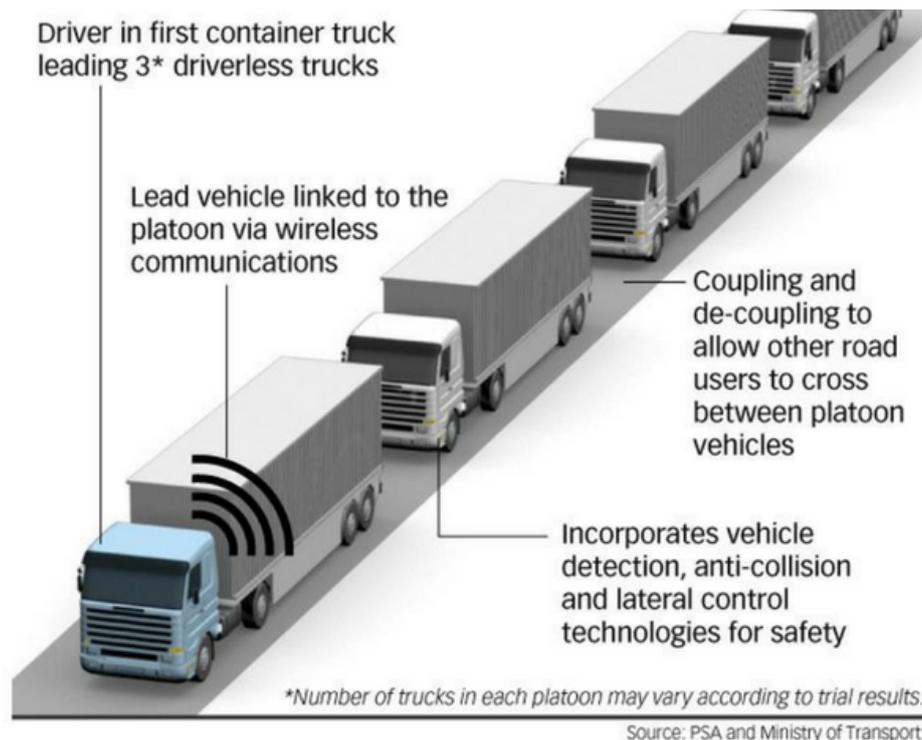
Figure 8: How platooning works [4]

There are downsides to this technology with drivers being less attentive than usual if needed for emergency manoeuvres, but in a level 5 autonomous system, which is the focus of this report, drivers are not expected to be in control. There is also the issue in using wireless communication between platooned vehicles which creates the possibility of attacks, which will be covered in section 3.6.1.

Swarming: Swarming is a technology which tries to mimic the action of social insects to produce a collective action when many individuals are in close proximity. Within an AV environment this swarm intelligence relies on artificial intelligence in which interactions between AVs lead to emergence of 'intelligent' behaviour analogous to schools of fish or ant colonies. Swarming would allow information on driving conditions to be sent electronically to AVs such as traffic congestion, road conditions or upcoming hazards through the vehicle swarm.

In February 2019, the UK Defence Secretary, stated "swarm squadrons will be deployed by the British Armed Forces in the coming years" [41]. Within a military environment swarming has gained interest and have been heavily tested in drone technologies [42] but its adaption to AVs has also attracted large amounts of research [43]. Swarms could be used for covering large areas for reconnaissance or search and rescue missions with their ability to make decisions among themselves highly beneficial. Without the need for communication back to base, interception of messages by the enemy would not be possible for example.

2.6.2 Communication with infrastructure

Wireless V2I communication has been allocated the Dedicated Short-Range Communication (DSRC) spectrum to transfer data since 1999 [44]. Information is transferred through Vehicular-Ad-hoc-Networks (VANETs) which are networks in which the nodes are either vehicles or Road-Side-Units (RSUs). This allows V2I sensors to obtain infrastructure data such as road conditions, congestion, traffic accidents, parking availability and roadwork locations, providing this to drivers in real-time. Traffic management systems can also use the data from VANETs to adjust signals and set variable speed limits to achieve optimum traffic flows.

The ease of deploying various Radio Frequency IDentification (RFID) readers and intelligent transport signage, cameras, street lamps and traffic lights have been questioned however, especially given these infrastructures are funded through public money. Some researchers also believe vehicle sensors on their own would be sufficient to achieve level 5 autonomy making V2I technology redundant [45]. Even if this were the case V2I technology would offer opportunity for redundancy in the system.

This report focusses on military supply chain vehicles operating in a desert environment. This type of setting will have little infrastructure in place or even if present maybe damaged or unusable due to confidentiality issues. I will therefore assume no infrastructure is available for the military scenario so will not discuss V2I technology further.

2.6.3 Communication with pedestrians

There were about 1.35 million deaths in 2016 due to road traffic accidents according to the World Health Organisation [46]. A large majority of these were pedestrians and cyclists with V2P communication being seen as a way to reducing these accidents. Given the focus of this report is supply line AVs in desert warzones, it is unlikely to have need for this technology so it will be considered in section further.

2.6.4 Communication with networks

Waze is a GPS navigation software system which works on smartphones providing navigation data to its users via a mobile network. It operates by gathering data from its large network of users to deliver real-time traffic reports based on their speed and location of travel when on road journeys. This crowdsourcing of traffic data gives useful information about traffic flow and optimal route management.

By using a similar scheme with AVs linked through wireless mesh networks, all AVs can be connected allowing sharing of traffic data, road conditions and other useful information. Development in cloud technology and rollout of the 5[th] generation mobile

network provides necessary structure for this V2N system to be realised, such will be the expected increase in data flow and lower latency required.

2.7 Infotainment: A vehicle infotainment system refers to a combination of entertainment and information conveyance to vehicle occupants. These systems typically include two-way communication with features including radio, CD player functionality and voice command recognition. Newer vehicle systems allow occupants to connect laptops for internet access and smartphones giving hands free use.
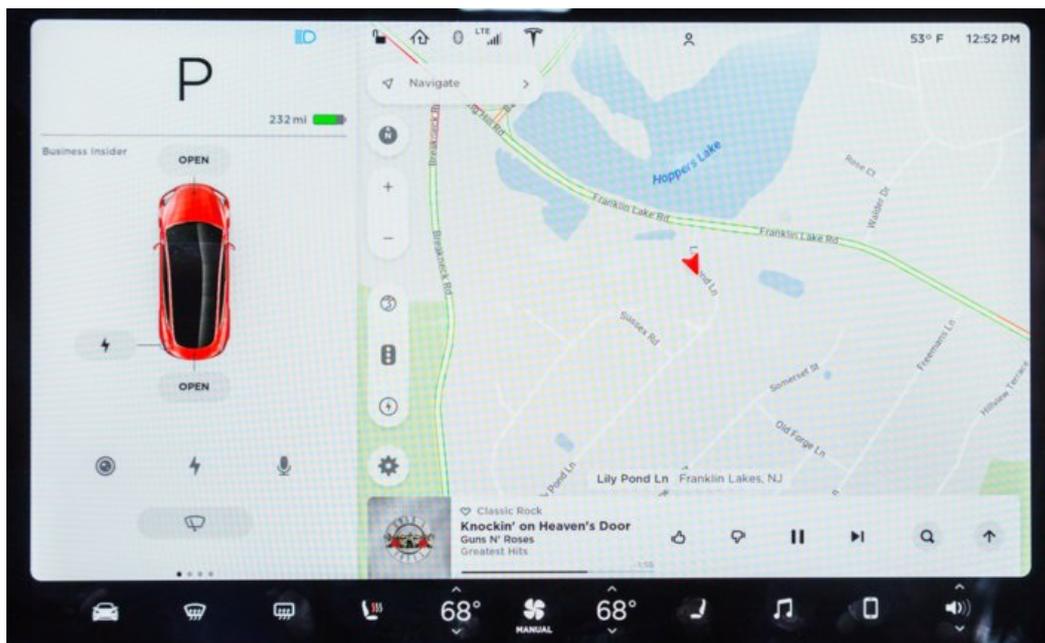
Figure 9: The Tesla Model 3 infotainment system [47]

An example of the Tesla Model 3 infotainment system is shown in figure 9 which has an audio and video interface with control mainly being through a 15inch touchscreen monitor. These infotainment systems are connected to the CAN bus along as with any other ECU devices.

2.8 CAN bus protocol

Network architecture within an AV uses wires to transmit messages in much the same way as has been used for decades in modern vehicles. One of these architectures is the CAN bus system developed in 1986 [48] before which links between vehicle systems were through many interlacing wires as shown in figure [10]. ISO 11898 series gives the specification of the CAN bus protocol stating how messages are sent over the network allowing nodes on the network to communicate [48] as shown in figure 11.
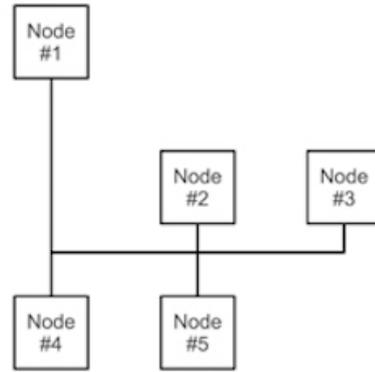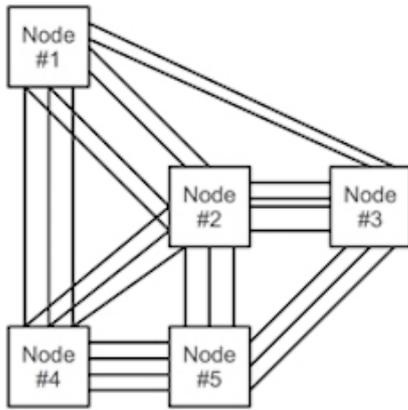
Figure 10: Historic linking of vehicle nodes [49]   Figure 11: A CAN bus system [49]

CAN bus nodes are the ECUs which are typically made up of a microcontroller, a CAN controller and a transceiver which transmit and receive messages in pairs of twisted wires and control all the electronic systems in a vehicle. Figure 1 in section 2.1 shows the large numbers of ECUs which are attached to the CAN bus network. These ECUs work by a sensor, such as for temperature, resulting in information inputs to a processor, which decides on an action, which is outputted as actuator directives for a system, such as the air conditioning unit.

As the number of ECUs have grown, vehicle architecture is now significantly more complex, with an ever increasing cost element of a conventional vehicle comprising the software not hardware components. This is only going to increase with the development in autonomy, which increases the attack surface and results in vulnerabilities comparable to the numbers seen in modern computers.

In terms of network architecture, the AV controllers on the CAN bus are separated from the ethernet part of the network by a gateway, which regulates traffic through the network. The CAN bus delivers a single user interface which has also inadvertently resulted in a single source of failure. How safety critical features are segregated from other part of the network will be a core security feature, discussed further in section 3.8.

There are other communication protocols; such as the more expensive FlexRay, designed to be more reliable and faster than CAN bus and Time Triggered Protocol, used in aviation. However, for simplicity this report focuses on security of the CAN bus, which is the more widely used protocol for vehicles.

2.9 Maintenance and updates

Over-The-Air (OTA): OTA is a way of sending out software updates, adjusting settings or updating encryption keys. This could be done using mobile networks or even home Wi-Fi networks. Given the number of AVs on the road could be in their

millions this is an efficient method for a company to communicate important data to vast numbers of devices. The integrity of messages sent out to ECUs on a network is vital, if an attacker can modify, delete or send out their own messages this can have serious consequences. This is a threat vector the military try to actively reduce [7].

On-Board Diagnostics (OBD): The OBD reports on the status of an AVs systems so when it goes into a garage these can be read by the technician. Early OBDs of the 1980s simply showed a light to indicate a problem but no other information [10]. Modern OBDs have standardised port access to provide real-time data allowing faster diagnostics. Software maintenance and updates can also be performed through the OBD port. Unlike OTA updates, this requires physical presence and is therefore harder for an attacker to gain access in order to modify, delete or inject messages.

2.10 Safety critical devices

Of the devices discussed some are of a safety critical nature which are essential to the functioning of the vehicle whilst others are more useful to have but would not cause the whole vehicle to become unsafe. If an attacker interfered with the ability of the AV to interpret its environment correctly or the processing of the information being relayed there is a risk this could compromise occupant or pedestrian safety.

ECUs which cause the vehicle to brake, steer and accelerate are clearly safety critical with those controlling the windows, air conditioning and infotainment clearly not. External location sensors will be defined as safety critical given they are needed to identify surrounding and so if they are not available a crash could quickly ensue. It could be stated if one failed other location sensors could allow safe operation, however given their essential role it is prudent to have these as safety critical. With some systems such as telematics it would go to the level of individual ECU's to determine if these were safety critical or not.

2.11 Summary

Since the Model T Ford vehicular technology has advanced significantly. What were once simple systems relying primarily upon mechanical interaction to rely driver commands into action have become complex systems of sensors, computers and actuators. This means simple commands pass through many complex systems creating a multitude of opportunities for an attacker to target them.

This project focusses on a military AV environment. The technology described so far for a civilian setting could be unsuitable for use in a military desert situation as it exists currently. Technology may need adjusting to make it more able to withstand very hostile environments. Sand ingress into cameras and sensors would be an issue as

well as optimal operating temperatures. Housing units for sensors and air conditioning may need installing for example to keep technology cool. Presence of troops in vehicles would increase temperatures to maximum operating level so having unmanned AVs would be beneficial [7].

3. Attacks on autonomous vehicles

3.1 Introduction

This report intends to map knowledge of AV attacks from a civilian to a military setting. To complete this a review of reported and theoretical attacks will be completed with a view to understanding which attacks could impact autonomy in a military environment. This will inform performance of a risk assessment in chapter 4 and a review of how these risks could be mitigated against in chapter 5.

The numbers of reported AV attacks and vulnerabilities are now so many that all of them cannot be described here. These do not even include cases of attacks which have not been reported for various reasons. This section will therefore review published and hypothetical attacks across attack surfaces of an AV to illustrate the various areas of functionality which can be affected.

An attack surface of an AV is the sum of the different attack vectors - the points at which an attacker can attempts to inject or extract data in order to compromise security. The number of attack vectors on modern cars have increased in recent years alongside higher complexity and automation which could result in higher numbers of attacks, with attackers having more methods of attack to achieve their objectives.

Attacks will target various systems in the car which were covered in section 2 and are illustrated in Figure 12.
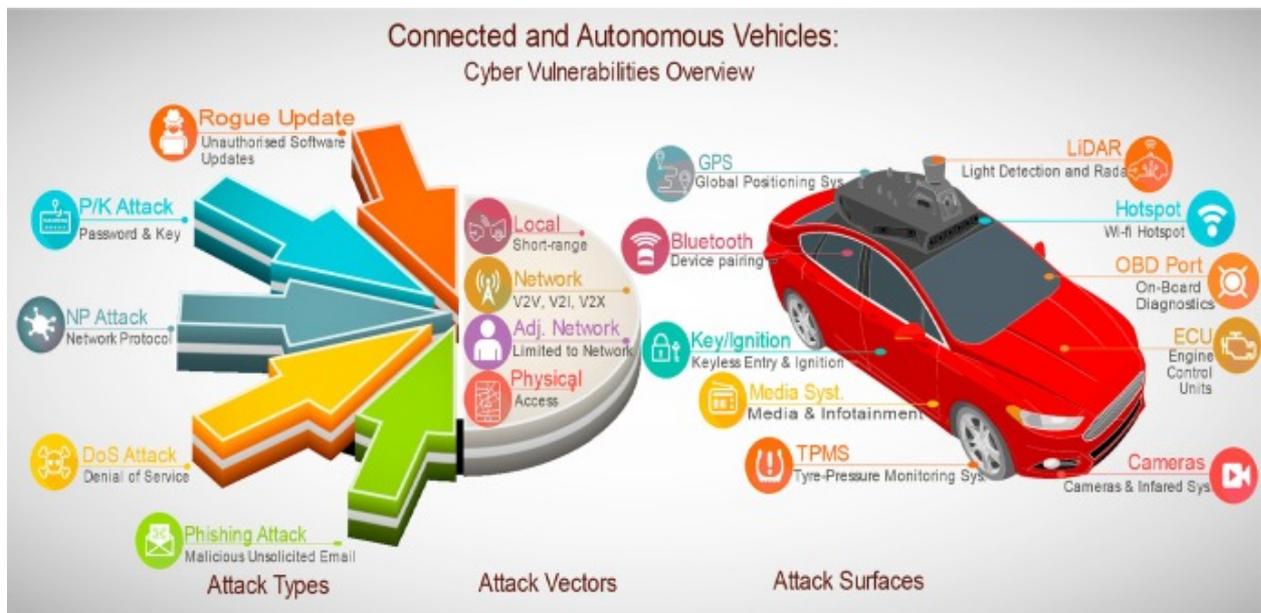
Figure 12: An overview of attack types, vectors and surfaces of an AV [50]

Physical access to a car could be difficult if in use by the driver, parked close to a residence or locked in a garage. Therefore, remote hacks are more desirable with an attacker being far away from the vehicle giving more time to perform the attack and less likely an attacker could be linked to the crime.

The different parts of a vehicle will now be reviewed considering potential and realised attacks.

3.2 Access and ignition

As at July 2019 thieves were continuing to perform jamming and relay attacks on car keys in order to get access to the vehicles [51]. To achieve this a relay transmitter and amplifier are used to pass the key fob signal from the house of the victim to an amplifier close to the target car.

Such a 'key' not only gives access it also allows the vehicle to be started remotely and grants system access. In attacks on the Audi RS4 after gaining entry to the vehicle, thieves simply programmed a new key into the system using the OBD port, which would be valid for the ignition, and drove it away [52]. The ignition could also be started through either brute force attack on the key fob algorithm or by intercepting the cryptographic information sent from the immobiliser during the handshake protocol [53].

Once the attacker has 'legitimate' physical access they can connect to the systems directly. Such attacks could therefore be used on AVs not only to again steal the vehicle but to permit access to software and hardware, specifically the CAN bus or OBD port, in order to commit other crimes which may remain undetected by the owner. By being able to access the car with the 'legitimate' key code the attack is more likely to remain undetected versus physical entry through vehicle damage. If the purpose of the attack was for installing surveillance technology for example, this is a much more desirable position to be in for an attacker.

Additionally, such remote entry systems allow other types of attack. The signal to the key could be jammed temporarily to prevent access, a key could be locked preventing use or even the power of the key fob or battery can be drained by actively probing the immobiliser [53]. Although these attacks would not cause major damage or loss to the vehicle it could be frustrating and inconvenient to the owner and potentially deadly in a military situation. It is noted proximity to the vehicle is needed for this attack so unless a vehicle has already been captured or disabled on the supply-line journey it would be secured behind a fence or moving at speed.

3.3 Position

3.3.1 Global position

Public GPS is an open standard with transparent architecture which relies on maps for navigation. Attacks poisoning on these maps, such as changing the location of places of interest or road layouts, could cause an incorrect location being driven to, wrong manoeuvring or traffic accidents. An example of map poisoning has been demonstrated by researchers showing how traffic data could be controlled [54]. This kind of attack could have been prevented with simple authentication mechanisms.

No physical access is needed for jamming and spoofing of signals to the GPS receiver, which make them attractive targets for attackers. Additionally, where maps are updated OTA, poisoning can also be carried out with no physical access.

GPS jamming: GPS jamming is a simple attack which use radio frequency transmitters at 1575.42 MHz [10] to stop genuine signals being sent to the GPS receiver. This can be performed by devices costing $20 (or about £17) [55].

Jamming could mean there is no accurate positioning information available which could force an AV to stop if GPS is the only navigation capability. Jamming, if used continuously, also enable thieves to disable a vehicle tracking system so it cannot be traced after being stolen.

Mitigation for this attack could be anti-jamming technology or military grade encryption and authentication processes [56, 57] and use of INS. There is also the option of using another secondary source for location measurement, such as India's NAVIC or the EU's Galileo systems [58]. However, as these use similar transmission technologies to GPS, just at different frequencies, they are also susceptible to jamming attacks.

Electronic countermeasures can also be used around vehicles to protect them against effects of electromagnetic attacks at given wavelengths. Each vehicle would generate a sphere of protection around it which would overlap with others in a platoon to create a 'tube' of protection. This would not only interrupt mobile signals allowing vehicles to pass an IED without remote trigger but also prevent an attacker interfering with electronic devices within the platoon [7, 10].

GPS Spoofing: GPS spoofing is the process of a malicious entity broadcasting signals to the GPS receiver in order to provide false location data. An attack would begin by broadcasting signals which are synchronised with genuine location data. The power of the signal would then be gradually increased until it overpowers the genuine signal and the device is then deviated to the location desired of the attacker. Figure 13 shows GPS spoofing - with the attackers signal only needing to be stronger than the

legitimate signals for most receivers to prefer it over the genuine signal This is not difficult given satellites are about 35,800km away orbiting the Earth [10].
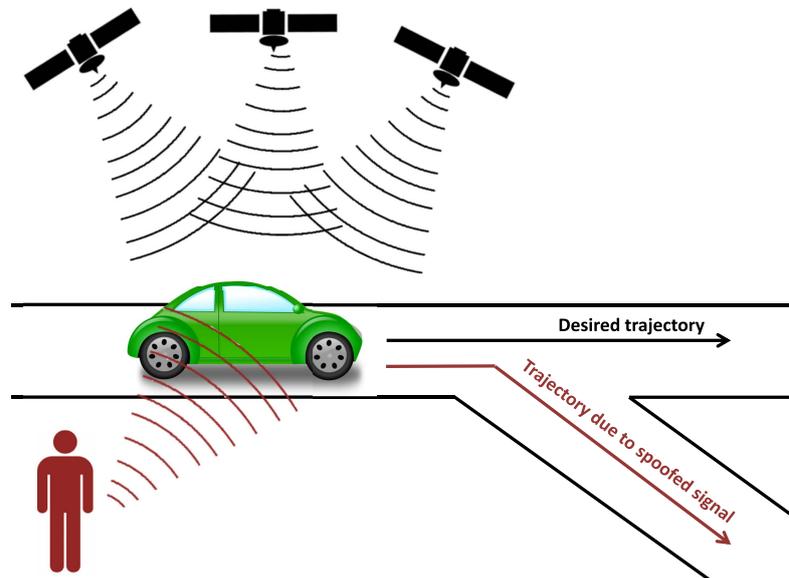
Figure 13: A GPS spoofing attack [59]

Research from the University of Texas showed in 2016 that all commercial off the shelf GPS devices were vulnerable to spoofing attacks [60]. There are also comprehensive details in the public domain as to how to perform successful GPS spoofing [61]. However, when a real attack was conducted in 2013 on a superyacht the system controls reacted to the change in GPS by reporting location variations to the crew who were able to correct this [62].

This was in a maritime environment, however, and with AVs the margin for error is much smaller than in the superyacht attack with the ability to cause traffic disruption or crashes before being found. This vulnerability has the potential, if an attacker could prevent the system sending an alert, to guide an AV to an undesired location to be stolen or have the produce they are transporting stolen.

There are many simple validations to prevent spoofing, such as monitoring identification codes and time intervals of satellite signals [63]. This would check signals are of an expected strength range (with spoofed signals being many decibels higher), ensure relative changes are as expected and that its strength varies as anticipated (it not being a perfect signal). The use of INS alongside GPS to maintain reliable navigation has also been used in the military for when GPS systems are compromised for short periods [64]. Map reading is still taught as a core skill at phase 1 training in the UK military [7], however this would only be a countermeasure if troops were being transported and they realised the course of the AV was being diverted.

Military-grade cryptographic authentication is also thought to offer strong protection against spoofing and jamming [56, 57]. However, according to Rebert Densmore, an e-war specialist, even modern combat grade GPS is very susceptible to manipulation [65]. It was reported in 2011 that a jamming and spoofing attack forced a US Sentinel reconnaissance drone into auto-pilot and landed it in enemy territory where its technology was subsequently reverse engineered [65].

Even if an attacker could not spoof a signal, the GPS signals could simply be drowned out by overpowering it with a stronger one. Attacks which overwhelm the system default them into a safety mode required by aviation standards which for UAVs would be a program to return them 'home'. If an attacker could spoof a UAV (or by analogy an AV) GPS this 'home' could be manipulated to be an enemy location. This is a big concern for the military with aircraft used to destroy their own equipment rather than it be lost to the enemy [7].

INS: INS technology is self-contained with no external communication, making remote attacks all but impossible. Physical proximity is needed to carry out attacks, which are discussed further in section 3.4.

3.3.2 Local position

Cameras: Devices in cameras can be partially disabled from a 3-metre distance through using an easily obtained low powered laser [66]. A high-powered laser could also cause permanent damage to the cameras giving a remote method of taking out cameras as opposed to physical damage.

Another researched attack used a high-powered torch or vehicle headlights to hide information such as traffic signs, road signs, road edges and obstructions [66]. These attacks have a high probability of success, with Iran shocking western intelligence when in 2011 they were able to 'blind' a CIA spy satellite by "aiming a laser burst quite accurately" [67].

It should be noted, however, that camera blinding does not have to be through a malicious attacker and could be from the sun setting, an oncoming vehicle having full-beams on or even modern glass buildings. This problem was tragically illustrated with a Tesla AV not being able to identify a white truck against the brightly lit sky [68], a task which notably also was not done by the human safety driver.

Further attacks would be to introduce fake pictures into the cameras which could make the AV perform false reactions if there were no other coverage. This attack requires more technical skills if trying to access the camera feeds. However, it could be performed very cheaply by just having a large poster across the road for example.

To mitigate against these attacks more cameras could be added with different angles to add redundancy, give a more complete picture of the environment and make blinding attacks from one direction difficult. Other senses such as radar, ultrasonics and LiDAR could also provide backup sources of the same data in varying wavelengths.

Further mitigation could use filters to remove laser light to prevent blinding by an intense beam which happens frequently in the aviation industry to pilots. Research has also been looking into cameras with the ability to add different filters to improve vision quality [69]. However, the military use 'wavelength-agile' lasers, which have the ability to randomly change colour making filtering of little use if an adversary had this technology [70, 71].

Radar: The simplest method of disruption to radar is through jamming the signal through noise saturation. This would cause the radar to not function correctly causing possible traffic disturbance as the AV would be expected to slow or manoeuvre more cautiously if key systems are impaired. An active electronically scanned array radar deployed on aircraft can be used to jam radar signals, as shown in figure 14.
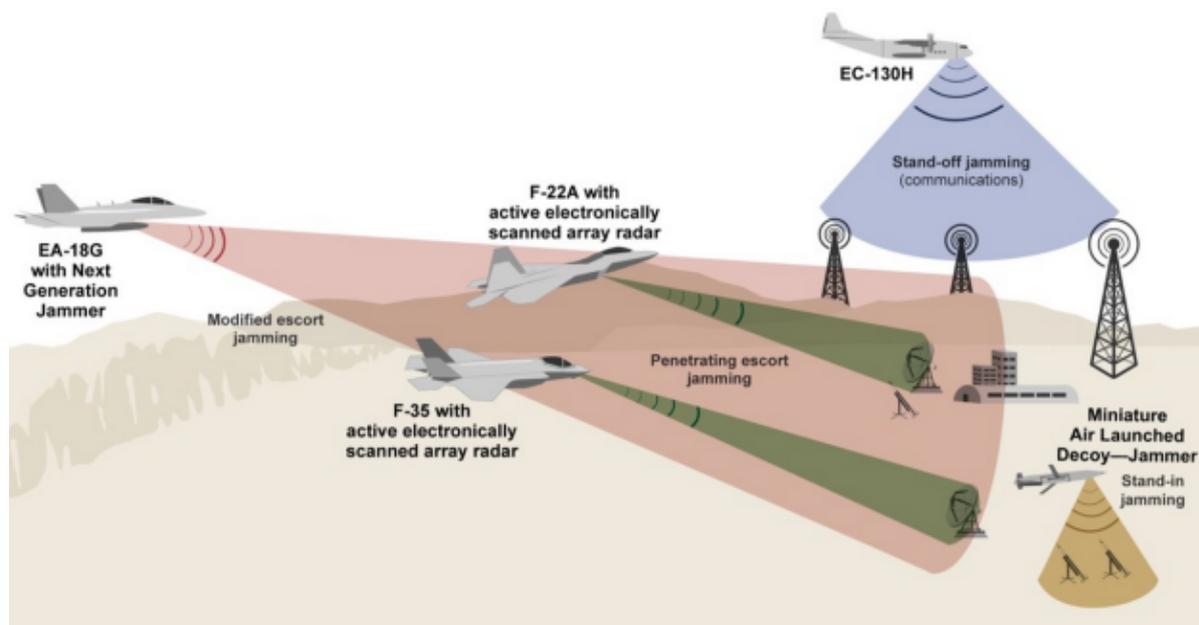
Figure 14: Radar jamming from the air [72]

To remain undetected by radar, smart materials with non-reflective surfaces could be used by an attacker to render objects invisible [73]. If the AV cannot detect objects within its immediate surroundings and there are no other sensors available for that field of view it could lead to the AV crashing.

Another attack would be to create 'ghost' vehicles by using a digital radio frequency memory repeater, which digitizes a received signal and stores it to be retransmitted later [74]. When replayed to an AV the false detection would be unrecognisable from

a legitimate signal and could result in traffic disturbances or cause the AV to become inoperable if it perceives its route is blocked by these 'ghost' vehicles.

To counter these attacks the radio waves attacking the system could be filtered which would cancel the effect of the repeater [74] or there could be other sensors such as cameras, ultrasonics, or LiDAR offering redundancy.

Ultrasonics: Interference caused by inaudible frequencies could cause the sensors to be turned off, resulting in loss of visibility for the AV. This could be mitigated against by applying filters to the returning signals or by use of spectrum analysis [73].

Repeated or faked ultrasonic reflections could allow an attacker to create false positive or false negative obstacles, traffic disturbances or cause a crash. This could be mitigated against by having multiple other sensors such as cameras, radar and LiDAR.

LiDAR: The LiDAR system has no way of checking if the picture constructed of the environment is bona fide or if an attacker has spoofed the data. Spoofing attacks have been demonstrated by researchers and involves emitting light back at the sensor at the same frequency as the laser reflected onto the target [69, 75]. Researchers made the vehicle believe an object was blocking its way, making it stop, and then overwhelming the LiDAR so much that it prevented the vehicle from moving again.

These attacks do not require expensive, sophisticated hardware but were done with a Raspberry Pi and a laser which are widely available and inexpensive to purchase. It has also been noted that smart materials with unusually high absorbency or reflectivity could be used by an attacker to give false detection and delineation [73].

Mitigations to these attacks involve using LiDAR operating at multiple wavelengths which aim to minimise jamming and spoofing with cheap devices [69] requiring an attacker to have more sophisticated hardware and knowledge. Mitigation could also be to use V2V communication for measurement sharing [69], however, this has the potential to scale up an attack to further vehicles.

Another solution is for the device to change the interval between scanning speeds constantly to make it hard to synchronise a laser being used in an attack to the correct frequency [59]. Redundancy could also be used having multiple sensors such as cameras, radar and ultrasonic capability in addition to LiDAR.

3.4 Position, velocity and orientation changes

Wheel rotation: The small magnet and sensor used to count wheel rotation and measure speed is well hidden and behind metal making it hard to attack but for brute

force physical means. If parked an attacker could get physical access but this is time consuming and difficult to do without being detected.

An Electro Magnetic Pulse (EMP) attack would aim to damage the magnetic device, but even though EMPs are easy and cheap to create, to hit and cause damage to the fixed magnet on the wheel would be difficult due to absorption of the EMP by the metal rim. The difficulty in attacking this sensor means it is a useful check to detect if other sensors have been spoofed or would give location data if no other sensor was available.

GPS: GPS map poisoning, jamming and spoofing attacks have been covered in section 3.3.1.

INS: There is not much research on vulnerabilities in these sensors, however, intentionally sending incorrect data to the system would cause issues. For example, spoofing the gyroscopes to show the vehicle was on a steep gradient may force the vehicle to travel very slowly. This could cause the AV to be unusable, allow for the vehicle to be hijacked, cause accidents or delays to other AVs.

This kind of attack would most probably need physical access in order to alter the sensor readings or to intercept and alter the communication on the physical wires or the close proximity wireless links. Research performed using CarShark tools showed the ability to modify a sensor value through changing packets within a CAN network [76]. The sensor reading would be checked by a control unit to ensure it is which is within a boundary of feasibility the system will accept, but if the attacker knows the tolerance they can adjust within this range.

To prevent low level attacks the communication channels could be encrypted which ensures spoofed signals are not easily injected into the network. The signals could also be monitored not just for if they are in range but also for a level of behaviour which could be classified as 'normal'. Additional sensors could also be linked to the INS such as GPS or cameras to determine if the vehicle is actually on a steep gradient.

Attacks using magnets and thermally attacking the gyroscopes in order to give the wrong position have been proposed [73]. These would require skill and access to the vehicle which would make them highly unlikely for anything but a denial of service attack, which could be more easily done by physically disabling them.

3.5 Functionality and usability

3.5.1 Vehicle maintenance and warning signals

TPMS: Researchers have demonstrated that the radio frequencies used to transmit TPMS data to the vehicle control unit can be read from 40 metres away using packet sniffing techniques [77]. The packets sent contain unique identifiers which would enable a vehicle to be tracked, leading to privacy implications. Another attack on TPMS is to consider using the packet identifier as a trigger for some other event, such as setting off an explosion as a vehicle passes by [77].

An attack on the TPMS itself could result in incorrect data being given to the driver which could be a false positive or false negative, resulting in a dangerous reaction, such as a driver continuing to drive fast even though the tyre is flat.

A driver obviously has the chance to see a flat tyre on a car. However, if this system is linked to a vehicle will full level 5 autonomy spoofing this ECU sensor could create false data on which to make decisions with, say, a flat tyre reading instigating the car into entering a 'safe mode' even if there is nothing wrong.

A safe mode is the procedure used when a vehicle goes out of its operational range or comes across a dangerous situation and in this mode the vehicle speed is significantly reduced. Whilst in a Google AV the safety driver would take over to control the car, with level 5 autonomy there would be no safety driver and it would be likely the AV would stop or safely take the vehicle to a default location, such as to a garage or back home. Even outside of Level 5 AVs, research has shown handing control over to the general population after periods of not driving gives poor safety statistics [78] and it may also be the case there are no troops being transported whom to hand control over to.

Further attacks discussed in The Car Hackers Handbook [53], include spoofing the TPMS signal into setting off alarms in the car, tricking the engine control unit into overcorrecting for road conditions which didn't exist and causing a fault in the system which could be exploited.

These examples highlight how serious attacks can be, even in such a simple system as the one used for measuring tyre pressure. The only way to completely remove the risk of attack would be to remove the TPMS, but as this is now required by law in most geographies it is not an option that can be used in a civilian setting.

Asymmetric encryption being used on the TPMS communication channel would be another less drastic mitigation against attack, however, it would take more power and the TPMS would need replacing more frequently [79]. As is the case for many information security situations the use of cryptography would be a balance of what is required from an AV in terms of security and its functionality.

Other sensors: An AV uses many low-level sensors and the type, functionality and use of these is a determining factor in how these could contribute to aiding an attacker.

For example, if seat belt monitors falsely record that an occupant is not buckled up, this could result in the AV refusing to move. External temperature and moisture sensors which have been attacked could give false readings for the driving conditions, for example, and not adjust speed and manoeuvrability for ice or wet creating a dangerous driving situation.

3.5.2 Driving aids

As vehicles become more connected, sensors which have been in a car for many years, such as the engine control unit, become susceptible to attack through not only physical ways but remote probing. An attacker could adjust the settings of these devices so the AV would not be reacting to conditions in the desired way. This could be through attacking the CAN bus, discussed in section 3.8, the OBD port, covered in section 3.9, or even the V2N communications, described in section 3.6.2.

Malware could be injected into these systems which, depending on its type, could have a wide variety of effects [73]. In all of the driver aids described a denial of service attack could be launched which would cause the functionality of the device to stop working. Given these devices are so crucial to an AV, it is likely to be programmed not to move until sensors are fully working.

Mitigation against CAN bus, OBD port and V2N communication attacks are discussed in sections 3.8, 3.9 and 3.6.2 respectively. These include an intrusion detection system, anti-virus software, firewalls, encryption and separation of safety critical ECUs from entertainment systems.

Engine control: An attack on this could stop fuel flow into the engine which would cause the vehicle to stall and stop moving. This would be extremely dangerous in a range of situations, especially if overtaking or travelling at speed.

Automatic gear box: By attacking the transmission system driving modes of the vehicle could be interfered with stopping gear changes. This would prevent the car starting if stuck in high gear or stop higher speeds being attained if stuck in low gear.

ABS: If the ABS is compromised this allows the attacker to control the braking of the vehicle, even allowing the brakes to be turned off altogether, with obvious devastating consequences.

Electronic stability control: The electronic stability control unit monitors numerous parameters through ECUs, with compromise of one of these small sensors having the potential to cause problems with this unit. Researchers found that the ECUs responsible for processing wheel speed could cause incorrect braking resulting in erratic and hazardous driving [80].

3.6 Communication

In 2014 researchers looking into defending AVs against malware attack [81] noted the most severe threats have not as yet been realised, given levels of expected connectivity have not been reached. In this section we will look at the current situation of V2V and V2N communication security.

3.6.1 Communication with other vehicles

It has been quoted that "V2V is the first automotive protocol to consider cybersecurity threats at the design stage, rather than after the fact." [53]. Although V2V devices are still in development, reviews of attacks based on what is currently know about this technology will be completed. V2V communication is likely to use mobile, Wi-Fi and short-range radio technology. Given short radio and Wi-Fi will face similar attacks (being similar technologies operating in similar parts of the electromagnetic spectrum) we will just focus upon Wi-Fi here.

DSRC is also used in some V2V networks, which needs installation of specialised equipment not only in vehicles but also at the side of the road. This project assumes no infrastructure in the warzone supply line region so DSRC technology will not be discussed further.

Connection mechanisms used for V2V communication, also allow attacks over wireless networks, resulting in AVs to be more easily compromised. This also offers a mechanism for an AV to be used as carrier of infection to be passed to other vehicles. A summary of an attacker's objectives is presented in table 2, along with the type of attack which would need to be performed in order to achieve the desired result.

Researchers have recently developed algorithms to implement cryptographic authentication mechanisms to ensure trust can be maintained when data is broadcast between vehicles [59]. This would use asymmetric cryptography which requires the use of central Certificate Authorities (CAs) to validate public keys [53]. This would be similar to its use for securing the internet, however, the identity of who would act as CAs for AVs has still not been determined. For asymmetric cryptography to secure data transmission it would need to have accountability integrity, non-repudiation, privacy and trust.

The vehicles participating in V2V communication use two types of certificates: Long Term Certificate (LTC) which contains vehicle identifiers and can be revoked; and a short-term Pseudonym Certificate (PC) which is used for anonymous transfers of common messages like braking. If an attacker compromised the PC they could only listen into messages for a short time, however, compromise of LTCs would allow

attacks using the correct vehicle identifiers to be used. This would allow an attacker to send updates or malware, which would be accepted by the vehicle system as it would appear to have come from a bona fide source.

			Attacker Objectives						
			O1.1	O1.2	O1.3	O1.4	O1.5	O1.6	O1.7
			Cause an accident	Cause congestion	Cause a driver to change their route	Erode user's faith in the system	Identify a particular driver or track their route	Conceal bad driving behavior	Falsely accuse/report misbehavior
Attacks	A2.1	Cause a false positive to be presented to a driver	X	X	X	X			
	A2.2	Suppress a message that should be presented to the driver (i.e., cause a false negative)	X	X	X	X		X	
	A2.3	Cause the system to be made unreliable, unknown to the driver	X	X	X	X			
	A2.4	Cause the system to be made unreliable, known to the driver	X	X	X	X			
	A2.5	Collect a set of messages from other vehicles and use them to identify a particular vehicle/driver					X		
	A2.6	Prevent the attacker's own vehicle from sending a message						X	
	A2.7	Create messages that will be attributed by the system to a vehicle that did not send them							X
	A2.8	Create messages from "ghost" vehicles to make a target's behavior seem more dangerous than it is, or the attacker's behavior seem safer than it is, from the point of view of an authority reviewing the record						X	X

Table 2: V2V attacks crossed with attacker objectives [53]

3.6.2 Communication with networks

Increased connection makes it possible to compromise vehicle devices which were not intended to be connected to outward facing networks, increasing risk of remote attack. V2N communication must conform to the highest security standards given its use can be life critical and the ability to conduct long range attacks will be most attractive to hackers. This type of attack was seen in 2019 when Iran claim they were

able to intercept live unencrypted video streams from US predator drones giving them crucial knowledge on enemy movements and operations [82].

Enabling communication with different technology will result in the ability to access AVs through the internet or broadcast into public spaces through Wi-Fi, mobile, radio or Bluetooth. As radio and Bluetooth are also used in the infotainment system attacks on these are covered in section 3.7.

Connection of AVs over networks offers the potential for current attacks seen online to be carried out on AVs. Researchers have shown Distributed Denial of Service (DDoS) attacks against VANETs are possible [59] with the ability to cause serious issues given they are used for communicating braking and traffic information.

Also, brute force password attacks could be used to find cryptographic keys in VANETs, a potential discussed in research [83,84]. These attacks could be prevented by having larger keys or a more secure algorithm. Network protocol attacks on the CAN bus are discussed in section 3.8, rogue software updates in section 3.9, with phishing and ransomware attacks covered in section 3.10.

Mobile: A mobile channel is an always on, high bandwidth, two-way, addressable method of communication. These properties could be used by attackers with the ability to conduct long-range attacks on vehicle mobile communications, largely anonymously. Devices in vehicles, such as OnStar, use the mobile network for automatic accident notification as well as permitting remote start, vehicle monitoring and other forms of driver assistance.

In 2011 a remote attack on vehicle mobile communications was completed by academics [85]. A flaw allowed authentication to be bypassed and after exploiting a buffer overflow the researchers could download code which enabled CAN packets to be sent to the vehicle. This gave them remote control over the vehicle in an attack they also showed could be broadcast to multiple cars.

Other attacks on mobile communications include eavesdropping on conversation, tracking vehicle movements and jamming distress calls.

In addition, the security of the mobile network is provided at the wireless level, not at the protocol level. If the connected device is using Internet Protocol (IP) traffic, standard IP security, such as encryption and attack surface reduction need applying.

Wi-Fi: In 2016 an attack by Chinese academics used Wi-Fi to create a fake hotspot using the name used by car dealerships [86]. If a Tesla connected to this hot spot researchers could use a vulnerability to deliver a payload into the vehicle. The gateway controlling access to the CAN was then replaced giving them control of the vehicle. In response to this, Tesla sent out patches in OTA updates and further added

new code signing which cryptographically signs updates with a key only Tesla knows. This mitigation allowed only those with knowledge of the secret key, Tesla, to updated software on the vehicles.

Other attacks on Wi-Fi include intercepting communications or breaking the Wi-Fi password [53], which would allow an attacker to gain information on what is being sent to, and received from the servers such as location, speed and engine status. The vehicle could also be tracked using the unique Wi-Fi host name. Mitigation against this would be to encrypt traffic on the communication channel, so if this is intercepted it would be unreadable. A closed private network could also be used which would only permit access to a whitelisted user with strong authentication mechanisms in place.

It is noted that attacks via Wi-Fi and mobile in V2N communication methods can result in similar outcomes, it is merely the conduit through which access is gained that differs.

3.7 Infotainment

An infotainment system will often accept standard CDs, offer USB connections and support many audio formats. It is also linked directly to the CAN bus in many vehicle architectures, providing an opportunity to inject packets into the network by exploiting vulnerabilities [85]. It has been noted that "The In-Vehicle Infotainment system offers more remote attack surfaces than any other vehicle component" [53].

When connecting other devices to the vehicle infotainment unit, virus and malware can invade into the automotive electronics. This could allow an attacker to spoof data displayed to the user on the infotainment control panel, to eavesdrop into actions or conversations of the vehicle occupants or even gain access to the CAN bus.

Whilst Infotainment systems can communicate through mobile and Wi-Fi channels, attacks on which have been covered in section 3.6.2, here we will focus upon other forms of attack.

CD and USB: Research has been published in 2011 which showed it was possible to send messages on the CAN bus when a corrupt music file on a CD was played [85]. Another paper from 2015, showed remote attacks were possible via the USB in units used for vehicle tracking [87], allowing texting services to be enabled when not needed. Faults noted with the system included inadequate cryptographic key management (all manufactured devices having the same private key) and poor password handling with cracks trivial due to lack of salting.

Through access to CD and USB ports an attacker could install malware which, if linked to the CAN bus, could affect any other ECUs on the network including those for safety critical functioning. Modified software for the vehicle could also be installed to other ECUs through this path. A further attack could include surveillance of occupants through control of the infotainment system and connected microphones.

Bluetooth: Bluetooth is commonly enabled for short range communication below 90 metres [10]. In 2011 researchers found Bluetooth was implemented on the infotainment unit ECU with no boundary checking which led to buffer overflow vulnerabilities. These could be exploited to execute code written by the attacker on other ECU devices attached to the CAN bus [88]. Whilst devices had to be paired for this attack to succeed researchers have detailed how easy it is to brute force the four-digit pin typically used for Bluetooth [89]. This attack required high levels of skill but it is predicted future AV attacks will become financial motivated and lucrative so this will produce lots of attention and motivation into ways of attacking AVs.

Other attacks on Bluetooth include eavesdropping the traffic in order for the attacker to obtain private information and jamming the device to create a denial of service attack [53].

To mitigate against such attack the development of a security layer for smartphone to vehicle communication over Bluetooth is needed. Cryptography can be used not only to make messages sent over the system unreadable but also to introduce asymmetric encryption to allow instructions and updates to be authenticated before executed by the vehicle. Cryptography would prevent packet analysis but it can affect usability and does not stop malicious code being run through memory exploitation [90]. To make things more difficult for attackers, increasing the password length of the Bluetooth system to more than four-digits would be beneficial.

In 2015, researchers looked at the Tesla Model S and found no insecurities which could result in an attack on the Bluetooth system which indicates what can be achieved through using academic work to patch vulnerabilities.

Radio: Radio is a long-range communication medium, typically transmitting up to about 60km [10], used in the infotainment system for digital radio, GPS and traffic messaging channels.

Researchers have shown control of the vehicle braking system and other critical ECUs can be achieved by sending data via the Digital Audio Broadcasting (DAB) radio receiver [91] used to transmit station, artist, and song name to the vehicle. It was noted this could be performed using low-cost, off-the shelf devices with the ability to broadcast the attack to many vehicles simultaneously. Although high technical skill is needed for the attack, researchers noted it could be compiled into an executable file suitable for someone with very little technical knowledge to use.

Mitigation against these attacks would be to patch the infotainment system with robust mechanisms such as ensuring authenticity of data packets at the CAN bus level.

The main cause of concern with attacks on the infotainment system, whichever attack method is used, is that all ECUs on the network have the same broadcast ability. Therefore, attacking the infotainment system, allows you to access the CAN bus, and through this the vehicles safety-critical devices attached to it.

3.8 CAN bus protocol

If a malicious attacker has access to a vehicle then they will be able to get access to the physical wires of the CAN bus network, the ultimate goal of an attacker. A major problem with this is it allows messages to be sent to safety critical devices, giving the ability to control over them. Manufacturers defend against this flaw by stating physical access is needed, but this is not impossible to come by when a vehicle is parked overnight or in a parking lot for long periods of time.

Many researchers have looked into attacks in bus systems [76], having little difficulty in getting safety critical ECUs to reply to fake messages without authentication. In 2014 it was found 42% of bus systems tested had no separation between critical ECU's and ones which could be remotely accessed [92].

One attack by Miller and Valasek in 2015 operated a remote attack against a Jeep Cherokee [93]. This was achieved through introducing malicious data into the CAN bus resulting in physical control of critical vehicle controls such as the braking system. This led to the recall of 1.4 million vehicles and cost the manufacturer €761 million (or about £627 million) [94].

Attacks discussed in section 3.7 have shown compromise of a USB port or CD player could lead to the ability to send messages on the CAN bus. In this research it was noted that whilst introducing gateways for separation is encouraged, it is not the answer to all problems [95]. A legitimate command at the wrong time to brake for example or a compromise of the gateway, as seen in an attack against the Tesla [86], would still result in an attacker affecting vehicle control and being able to cause serious disruption.

Research into the Tesla Model S found that even though the network was compromised only legitimate messages could be sent to the gateway [96]. This was enough to enable power to the vehicle but if it was travelling over 5mph it would cause the vehicle to switch to neutral and stop. Although a driver could still steer and brake in this attack, this is still concerning, especially if this occurred in a warzone.

The integrity of messages sent out to ECUs on a network is vital. If an attacker can modify, delete or send out their own messages this can have serious consequences. False messages from sensors could cause accidents as discussed in sections 3.3, 3.4 and 3.5 or cause the vehicle to become unusable.

Other CAN attacks could be to plug a device directly into the bus system in an attempt to start the vehicle without the legitimate key. A malicious device could be installed in to enable remote communication to the CAN bus, making an attack on this system wireless as opposed to physical access being required.

Recent research [97] suggests using asymmetric cryptography using the Advanced Encryption Standard (AES) for ECU authentication and authorisation. This would prevent packet modification on the CAN network by requiring every message be authorised by a central security module with asymmetric key access distributed to the ECUs. This would make it computationally infeasible to modify or inject data packets. If the impact of the extra checks were small enough this would be a useful security mechanism but if delays were excessive this would not be practical in such time-critical, real-time system such as AVs.

Tesla have added a gateway to their vehicles which manages communication between different CAN bus systems [12]. It acts by filtering out messages which are passed from say the infotainment system to the CAN bus. Although this adds a layer of security to the standard linked CAN bus, this still creates a single point of weakness in the gateway.

The aviation industry mitigate against non-safety critical ECUs affecting safety-critical ECUs by having four separate networks. One for safety critical systems, another for air flow, one for cabin doors and another for the passenger entertainment systems [98]. These are completely separate systems with attacks on any individual network ECU not having the ability to influence another. A downside of this scheme would be that it requires four separate networks to be created and maintained. This is not only much more expensive but also not a feasible option for a post design stage retrofit.

3.9 Maintenance and updates

OTA: The use of physical updates will be unmanageable if AVs become as numerous as conventional vehicles are, with researchers proposing technology for distributing firmware updates OTA [99]. The potential for updates to be replaced with malicious code to infect the entire fleet and infrastructure is a severe problem which would cause widespread damage.

Security threats and mitigations associated with mobile networks and Wi-Fi, which would be the main technologies for distributing OTA updates, have been discussed

in section 3.6.2. Specifically, for OTA updates solutions would be needed which have secure protocols using cryptographic techniques to give confidence any update sent to AVs are legitimate.

OBD: The OBD port is mandated by law [100] and is the means by which technicians update the ECUs as well as perform diagnostics. It can access everything on the CAN bus network, so all attacks discussed in section 3.8 for the CAN bus could be achieved with access to the OBD port. Attacks from the OBD port could include changes to ECUs such as those controlling braking and engine functionality [76] as well as extraction of vehicle information or personal details of the owner [101].

A study by Carnegie-Mellon University revealed widespread failures to apply basic security principles to the OBD port [62]. In some vehicles there is no security preventing firmware upload through the OBD port [59] allowing attackers to completely reprogram vehicle behaviour resulting in it being a threat to public safety. Research has even found a way to corrupt the devices used to connect to the OBD port thereby infecting every vehicle the technician connects to [85].

Studies completed in 2013 described security methods for AV protection (which included asymmetric cryptography to ensure firmware uploads came from bona fide source), statistical anomaly detection systems and ECU software integrity [50].

3.10 Human aspects

AVs will operate in an open environment with a population who may have no experience with the technology which runs their vehicle. The AV domain could become rife with social attacks already conducted over the internet, such as phishing and ransomware. An example of ransomware used in the AV domain could be for an attacker to disable an AV and render it immovable unless a bounty is paid.

AVs will generate vast amounts of data - something not only criminals would find appealing, but also commercial companies who could use data to track behaviour [103]. The vast information generated would need to be anonymised, encrypted and protected to prevent attacks on users and also stop the vehicle companies themselves from facing huge fines if this data were compromised. Prosecution, such as under the EU's GDPR carry significant penalties in the event of a data breach, with fines of up to 5% of global turnover [104].

Social attacks and data privacy within the world of AVs could fill many a Masters and PhD project. Although mentioned here only briefly, future use of social engineering by attackers could be highly significant, as could the financial motivation of privacy breaches.

3.11 Summary

The more lines of code and devices added to a system, the more complex it becomes and the more attack surfaces there are for an attacker to exploit. The security and possible safety benefits of implementing a new procedure have to be weighed up against costs and impact on the systems functionality to determine if it will be a workable solution.

There are a multitude of ways in which AV security can be improved and even simple information security techniques used extensively in other domains such as banking have not been employed by the automotive industry. Cryptography is a mature technology, which provides widespread and crucial security services, but advanced vehicles are on the roads which do not utilise this.

Technology is still developing, however, and yet to be subject to significant adversarial pressures with the majority of known vulnerability identification being done by academics, hobbyists and white hat hackers. Bug bounty schemes exist, for example, in which individuals are rewarded for reporting vulnerabilities they have discovered allowing companies to fix these and implement more secure practices.

There are only a few cases recorded of attackers exploiting vulnerabilities requiring little skill [101, 105]. Attack number and sophistication are expected to increase, however, when financial benefits of attacking an AV emerge to motivate black hat hackers.

Within the UK military software maintenance and updates are all done at the Original Equipment Manufacturer (OEM) usually by sending updates over an internet connection for diagnosis in the UK rather than have an OEM on site. This 'at length support' could result in even everyday problems taking days to resolve so if AVs were deployed the military it would be sensible to have representatives from OEMs in the field. This not only increases time to rectify issued but also reduces an attack vector of the internet connection used in sending update information.

With rectifying and finding faults onerous, disabling AVs would cause enormous issues, especially in a supply-line vehicle which would have countless knock-on effects. In a battle scenario soldiers are taught shoot to injure rather than shoot to kill with an injured soldier occupying three individuals as they go to aid an in injured comrade [7]. A similar logic could be used in AV deployment with the enemy disabling AVs which take up more time to fix and occupying more people than complete destruction.

4. Risk analysis in the context of autonomous military vehicles

4.1 Risk analysis methodology

Having completed a review of attacks for civilian AVs, these results can now be mapped and applied to a military AV scenario. In order to optimise the benefit for users in a military setting these have been categorised based upon the different objectives an enemy attack would have. Enemy objectives will then be analysed according to the types of attack needed to achieve these goals.

How the military assesses risk would depend on the deployment scenario. As stated in chapter 1 the environment under consideration will be a war zone requiring supply chain vehicles with level 5 autonomy. It is assumed the capability of the adversary has high technical skills with up to government level funding.

A risk assessment will be completed based on threats and impacts identified using a slightly modified version of the HEAVENS security model classification scheme described in appendix 1. The mapping of threat and impact assessments to a security level score is shown in table 3 [106]. This model is threat centric and is simple to use and understand. Having fewer levels makes classification easier and does not suggest to a reader a higher level of accuracy than is attainable. The modifications made to the HEAVENS model are detailed in appendix 2 and appendix 3. Changes to the model are needed due to the different levels of importance placed on risks between a military and civilian setting. These scales still indicate the level of risk associated with the attack types in order to achieve the final objective, however.

The threat level corresponds to the estimation of how likely it is a threat will be realised. Parameters used to estimate the threat levels in the HEAVENS model are expertise, knowledge about the system, window of opportunity and equipment needed. Added to this list will be the 'cost to perform' the attack. How well-funded the enemy needs to be to carry out the threat will inform the likelihood of an attack occurring, and if it is likely to be a singular threat or something which could be easily repeated.

These five parameters are all rated with a value of 0, 1, 2, or 3 as per the criteria listed in appendix 2. After parameters are rated, totals of all areas are translated into an overall threat level which is used in the risk assessment. The scoring has been slightly modified from the HEAVENS model to account for the addition of another parameter (cost to perform the attack), details of this modification are given in appendix 2 and shown in table 11.

The impact level parameters in the HEAVENS model are safety, financial, operational and privacy and legislation impact. In addition to these parameters, this report also adds political impact which will consider AV issues which could cause political unrest. A decade ago the standard armoured vehicle came under scrutiny due to their

unsuitability and poor design, resulting in these vehicles becoming called 'coffins on wheels' [107] and leading to troop unrest and low morale. Similar impacts on troop engagement with AVs needs to be considered.

In the standard model safety and financial impacts are rated with scores of either 0, 10, 100 or 1,000 with operational and privacy scores 0, 1, 10 or 100 depending on the severity of impact, starting at 0 being no effect. This has been modified within the military environment so the financial and the new political parameters are aligned to the operational and privacy scoring system. The safety parameter will be assigned scores three times these parameters, which will be 0, 3, 30, 300.

The rationale behind modifying the scoring is that a theatre of war is a more hostile and volatile environment where there is already the risk of life-threatening injuries from enemy combatants. Risk in a civilian sector also do not just relate to vehicle occupants, with pedestrians and other road users more likely to suffer life-threatening injuries than those in the vehicle. Within a military scenario there will be fewer people around in general (being a desert) and likely to be enemy combatants. This report will still have safety ranked higher than other levels, but this will be by a factor of three rather than ten in the HEAVENS model.

Within the HEAVENS model a very high financial impact is more is defined in relation to companies where significant financial loss can equate to the inability to continue trading. As the military is underwritten by the government financial risk would not be as catastrophic with issues more being political unrest if military spending was excessive. Therefore, having this risk more aligned to the scoring levels of operational or privacy impact is more appropriate for the focus of this report.

Security Level (SL)	Impact Level (IL)				
		0	1	2	3	4
Threat Level (TL)	0	QM	QM	QM	QM	Low
	1	QM	Low	Low	Low	Medium
	2	QM	Low	Medium	Medium	High
	3	QM	Low	Medium	High	High
	4	Low	Medium	High	High	Critical

Table 3: Security level assessment table [106]

The HEAVENS security model uses the impact level and threat level computed and combines these values to give a final security rating, as shown in table 3 and

completed in section 4.3. The score outputted from the risk assessment can then be used to inform further steps for if a risk needs to be reduced, shared, avoided or indeed retained, as discussed in chapter 5.

4.2 Objectives of an enemy attacker

Threats to an AV in a military setting have been categorised based upon objectives of an enemy attacker. This will help focus attention on outcomes rather than isolated parts of a system to produce a more useable result. Within objectives identified there may be more than one attack which could achieve the same objective, these attacks will be risk assessed separately in the analysis completed.

Objective 1: Capture of an AV, troops or supplies

Capture of a fully working vehicle would allow the autonomy and algorithms used in the AVs design to be analysed. This could allow the enemy not only to use this technology in their own designs (if they could reverse engineer the AV technology) but also to use the captured vehicle to understand all the weaknesses in other vehicles of the same design. Having knowledge of the technology used in the fleet could allow attacks to be performed which would affect enemy systems collectively rather than on a singular basis.

Capture of a vehicle would also be an enabling factor for the attack detailed in objective 2, which involves using a single AV to poison the whole fleet when plugged into the OBD. AV capture would also have the effect of reducing the oppositions military strength but this would be easier to achieve by destroying the vehicle, as discussed in objective 5, or using conventional fire power.

Attacks which would allow a capture scenario, mainly relate to the stopping of a vehicle or slowing it down significantly, specifically attacks 1.1-1.8. These attacks would need to be backed up by physical assistance in order to move the vehicle and cargo to an enemy location once stopped or slowed. In scoring the attack it is assumed this physical attack is successful in getting the AV to enemy territory if the cyber element is completed. Access to the vehicle is also assumed possible using attack techniques described in section 3.2.

It is of note that if the enemy captures a vehicle through very simple attacks, such as walking in front of it to make it stop, this critical weakness would not reflect well politically. The moral of the troops would also be affected if a vehicle could be so easily driven into an enemy base or disabled remotely, with the fear of getting into AVs spreading among troops.

Seven of these attacks involve being able to capture the AV in perfect condition, provided successful extracted once in a 'sitting duck' position. Two of these attacks

involve disabling the vehicle through causing a crash, with the safety impact rating reflecting this. These situations would have to planned not to cause too much damage to the troops, cargo or vehicle technology or they would perform no better than using a physical means for destruction. It is also assumed an adversary would follow the rules of warfare and not hurt troops which may be in the AV after capture.

Financially, whilst the impact of losing a single AV would be high, the much larger impact would be the enemy being able to gain knowledge which would compromise the whole fleet. Operationally, in attacks 1.1-1.3 the vehicles secondary sensors are affected, whilst attacks 1.4-1.9 affect the primary sensors.

1.1 Person walks in front of the AV to enforce a stop situation

This is a very simple attack to perform and requires no skills, cost or expertise. The attacker would need to be in possession of knowledge that the AV would stop rather than run over a human, which if similar algorithms are used as for a civilian setting would be the case. There would be quite a slim window of opportunity to perform this attack with supply chain vehicles travelling fast in the war zone area.

1.2 Flat tyre spoofed to force the AV to stop or slow down

Attacks on the TPMS are described in section 3.5.1, which demonstrates TPMS data sent over radio can be read from 40 metres away [77] giving the attacker an increased window of opportunity. The equipment, skills and costs needed to perform this attack are also relatively low but not something a layperson could perform. There would be the TPMS physical wires connecting to the control panel also which would give redundancy, but the system maybe unsure which reading to trust – if the default is to trust the one showing a flat tyre, for safety reason, this attack would be successful.

With a fully autonomous system the spoofing of a sensor could trigger the vehicle into a safe mode which slows the vehicle down. Given the AV is operating in a warzone there will be systems to automatically re-inflates the tyre [108] because no assistance would be available otherwise. This would, however, create a window of opportunity to capture the vehicle when it slows or stops for tyre re-inflation.

1.3 Force the AV to slow by spoofing motion sensors to show the AV is on a slope

This attack on the INS has been described in section 3.4 and works by spoofing the gyroscopes forcing the AV to travel very slowly. Research indicates this attack would most probably need physical access in order to alter the sensor reading or the communication on the wires. The expertise, knowledge and equipment needed to perform this attack would also be of a high level.

1.4 Force a stop by jamming or spoofing visual sensors to detect an object in front of the vehicle

The ability to complete this attack would depend upon the type and number of sensors on the vehicle as well as the terrain in which it is operating. Such attacks would be less likely to succeed within a barren desert environment (the focus of this project) than in woodland for example.

The local position sensors, discussed in section 3.3.2, would have to agree that in all probability there was an obstruction in front of the vehicle which would cause it to completely stop and not have the ability to change direction to avoid this. As with other attacks discussed, the AV would be moving at speed through the supply chain route leaving the attacker a relatively short window of opportunity in which to perform this attack.

Visual cameras are usually supported by other sensors on the vehicle, such as sonar and/or radar sensors, which would provide contradictory evidence that the path was clear if these had not also been spoofed. To complete the attack at least two sensors would need to be spoofed to give convincing data, which would require more skill than simply spoofing one.

However, if other sensors were unavailable due to jamming this would be easier to complete. The equipment required to do this attack is not sophisticated, examples include injecting pre-recorded sound into the ultrasonic sensors, jamming the cameras with bright light and through putting certain types of material in the road which are more or less absorbent than normal to confuse the LIDAR and radar.

1.5 Jamming primary sensors to force the AV into a safety stop

Primary sensors are required for the AV to function safely and if unavailable could cause it to come to a safety stop. Attacks on these sensors have been covered in section 3.5.2 and describe attacks on the CAN bus, the OBD ports and V2N communication links as ways into these systems. The OBD port would need physical access and would be a more difficult to attack, so wireless methods will be used for the example to perform the risk assessment. Attacking these systems requires a high level of expertise and access is difficult. The knowledge of the system, equipment needed and cost to perform the attack would not be too high, however.

1.6 Disable the engine forcing the AV to stop

Disabling the engine would involve preventing fuel flowing to the engine making the vehicle unable to move. It would involve similar attack techniques to those described in attack 1.5. The attacker would require more knowledge to complete this attack,

however, as access to the engine management systems are required which is more securely protected.

1.7 Alter GPS position so the AV drives into an obstacle

GPS spoofing attacks are covered in section 3.3.1. This shows how comprehensive details of how to perform such an attack are available in the public domain. The equipment needed is both relatively inexpensive and easy to get hold of with someone proficient in technology being able to carry out this attack. The attack can be carried out remotely, but it is noted that the vehicles will be travelling quickly across the supply chain route which would not give an unlimited window of opportunity. It is likely, however, that other sensors, such as the cameras, LiDAR or sonar, would pick up on the obstacle which would prevent this attack from being successful. Other methods of jamming or spoofing would be needed to overcome this issue.

1.8 Alter GPS mapping data so the AV drives into an obstacle

Attack poisoning of GPS maps has been described in section 3.3.1. Changing the location of where the AV is travelling could cause it to drive into an obstacle. This would need for an attacker to gain access to the GPS maps stored on the device within the vehicle, which could be done remotely and which does not require particularly high expertise, knowledge or technical equipment. It is likely, however, that other sensors would identify the obstacle thus preventing the attack from being successful.

1.9 Spoof GPS location so the AV drives into an enemy camp

This attack has been described in section 3.3.1 and a less sophisticated version in attack 1.7. Being able to use the GPS system to control the vehicle over a longer distance and guide it to a precise location is more complicated, however, with expertise of the attacker and knowledge of the system needing to be much higher than in attack 1.7.

Objective 2: Return the captured AV to base and poison other units

The enemy would need to have successfully completed one of the attacks in objective 1 first before being able to complete this attack. The ratings given therefore account for work required to first capturing a vehicle without damage as well as sending it back to base in a poisoning attack on the whole fleet. The enemy would have to install software on the captured vehicle which would need to have the capabilities of poisoning the OBD equipment and to go un-noticed so this malware would spread to other AVs when they are being updated.

Due to the busy nature of a military base vehicles are not strictly accounted for. The vehicle location histories are not routinely checked either for historic location data

before connecting to OBD [109]. This would mean an attacker would not need to construct a history of vehicle movements to account for the its location when being loaded with malware.

This attack would have significant operational impact in terms of the enemy being able to control AVs and a financial cost in terms of lost use, cost of replacements or cost to resolve the problem. Additionally, the safety impact could be material and political fall-out significant should this attack become known.

Objective 3: Cause confusion and break command

This attack would only be effective if the AV is being used to carry troops in its supply line duties. Supplies of food or ammunition would be unaffected by the attacks described here. The ability of the enemy to cause issues such as playing loud noises, or causing erratic movements, would affect the ability to respond to a physical attack by making it hard to hear commands or respond to those commands for example. There is also the ability for the enemy to cause issues by feeding incorrect information, such as video feeds, or just cutting off communication altogether. It is assumed this cyber-attack would be performed in conjunction with a physical attack and would have the most impact on troop safety.

3.1 Mission data made unavailable

This can be completed with a simple denial of service attack on the wireless communication channel, such as radio, Wi-Fi or mobile, described in sections 3.7 and 3.6.2. The knowledge to complete this attack is widely available, with equipment both cheap and easy to get hold of with only a lay persons knowledge needed. The window of opportunity is not unlimited, however, and would rely on an attacker having to access communication links when the vehicle was travelling quickly across the desert.

3.2 Mission data altered

Tampering with the systems which provide information about enemy locations, allied location and mapping for example could cause both mission failure as well as confusion and inappropriate action if the AV was under attack.

This is not a simple attack to perform with high expertise required with the data needing to seem genuine and be fed into the infotainment system within a short window of opportunity. The system knowledge, equipment needed to perform this attack and cost of doing this would not be particularly high, however, as mentioned in section 3.7.

3.3 Activation of in vehicle systems

This attack would cause distraction to troops being transported by, for example, increasing the volume of music, intensity of lighting, temperature levels or causing numerous warning lights to be activated. Attacks on these sensors have been mentioned in section 3.5.1. The attacker expertise, systems knowledge, equipment to carry out and cost to perform the attack would be quite low. With remote attacks possible, the attacker would not have unlimited access, but the system would be easily available.

3.4 Force erratic AV movements through engine control unit or accelerometers

Attacks on the engine control unit have been described in section 3.5.2 and a denial of service against the engine unit covered in objective 1.6. In order to get control over the engine so as to permit changes in speed or direction, rather than just disable it would require more expensive equipment which would be more difficult to source.

3.5 Force erratic AV movements through visual sensors or GPS

Jamming and spoofing of visual sensors has been covered in section 3.3.2, with simple, inexpensive attacks requiring little knowledge and simple equipment able to be performed. What makes this attack more difficult is to continuously create stimuli which will cause movements to disorientate or distract occupants. This would require persistence and multiple lasers, noise saturation devices, smart materials or faked ultrasonic reflections from various locations. This may be easy in a wooded terrain where the vehicle would be travelling slowly, and devices could be hidden in the trees, but would be very difficult in an open desert environment.

Objective 4: Surveillance

Surveillance of a vehicle is possible if the enemy has managed to infiltrate the AV systems and installed malware which will allow monitoring of commands or viewing of camera feeds undetected for example. This could include accessing vehicle movements or turning on microphones to listen in to troop discussions, if being transported. The specific impact of this objective depends on the type of attack completed and the information extracted, but obviously has the ability to feed crucial information to the enemy, which may not just be restricted to supply chain operations.

4.1 In vehicle discussions of troops obtained

This attack can be completed with off the shelf equipment and requires little financing. With remote attack possible on any of the supply line journeys, the enemy has a large window of opportunity. Some knowledge of the system is required with an adversary needing familiarity with attacks and security. As described in section 3.7 the infotainment unit would allow an attacker to eavesdrop through connecting to the vehicles microphones and turning these on.

4.2 A history of AVs recorded movements is obtained

This information is stored in the infotainment system and can be obtained by an attacker with the ability to connect with this system. This attack involves similar threat parameters as described for attack 4.2.

4.3 The enemy is able to see AV movements through camera feeds

The ability to extract and potentially live stream camera feeds from the vehicle would not only allow an enemy to see the current location of a vehicle but also be able to find locations of military bases and other strategically important areas. This attack requires a higher level of expertise and knowledge about the system than the attacks in 4.1 and 4.2. The cost and sophistication of equipment needed to carry it out are also higher, and in the case of live streaming requires continuous access to the AVs communications networks.

Objective 5: Disable or destroy an AV

To disable an AV an attacker would simply need to put the system into a mode which would cause it not to function, such as feeding information to a safety critical sensor that something was not operating as expected.

Alternatively, should an attacker gain operational control, destruction of the AV would be trivial - it could simply be crashed. In practice, this attack would more likely be used for additional mischief by an attacker, as described in objectives 1 and 2.

If the attacker were able to destroy a vehicle, financially a single vehicle would be lost, which although expensive would not be substantial in terms of overall military assets. What would be more impactful is the political impact and the moral of the troops. In terms of safety, attacks 5.6 and 5.7 have the potential to cause at least serious injuries. Given the enemy is destroying the AV there would be no loss of privacy.

The threat levels for attacks 5.1-5.7 have already been described in objective 1, namely attacks 1.2-1.8. These attacks are used again here because they are achieving a separate objective which will give a different impact score and therefore create an altered risk assessment score for this scenario.

4.3 Risk Assessment

A risk assessment on all the attack objectives is shown in table 4, which has been completed using information from section 4.2 and chapter 3. Details of how the Impact Level (IL) and Threat Level (TL) are calculated are shown in Appendix 4, with the Security Level (SL) calculated using table 3 in section 4.1.

Objective	Target	Attack Type	TL	IL	SL	
1.1 Person walks in front of the AV to enforce a stop situation	Sensors (physical attack)	n/a	4	2	High	
1.2 Flat tyre spoofed to force the AV to stop or slow down	TPMS	Spoofing	2	2	Medium	
1.3 Force an AV to slow by spoofing motion sensors to show the AV is on a slope	Motion sensors	Spoofing	1	2	Low	
1.4 Force a stop by jamming or spoofing visual sensors to detect an object in front of the vehicle	Visual sensors	Spoofing Jamming	2	2	Medium	
1.5 Jamming primary sensors to force the AV into a safety stop	Primary sensors	Jamming DoS	2	2	Medium	
1.6 Disable the engine forcing the AV to stop	Engine control unit	DoS	1	2	Low	
1.7 Alter GPS position so the AV drives into an obstacle	GPS	Tampering	3	2	Medium	
1.8 Alter GPS mapping data so the AV drives into an obstacle	Mapping	Tampering	2	2	Medium	
1.9 Spoof GPS location so the AV drives into an enemy camp	GPS	Spoofing	1	3	Low	
2. Return the captured AV to base and poison other units	CAN, OBD, GPS (physical attack)	Spoofing Tampering	1	4	Medium	
3.1 Mission data made unavailable	Infotainment	DoS	4	1	Medium	
3.2 Mission data altered	Infotainment	Tampering	2	3	Medium	
3.3 Activation of in vehicle systems	Infotainment Sensors	Tampering	3	1	Low	
3.4 Force erratic AV movements through engine control unit or accelerometers	Engine control Accelerometers	Spoofing	1	1	Low	
3.5 Force erratic AV movements through visual sensors or GPS	Visual sensors GPS	Spoofing	2	1	Low	
4.1 In vehicle discussions of troops obtained	Infotainments (audio system)	Information disclosure	4	2	High	
4.2 A history of AVs recorded movements obtained	GPS	Information disclosure	4	2	High	
4.3 The enemy is able to see AV movements through camera feeds	Cameras	Information disclosure	1	3	Low	
5.1 Flat tyre spoofed to force the AV to stop or slow down	TPMS	Spoofing	2	1	Low	
5.2 Force the AV to slow by spoofing motion sensors to show the AV is on a slope	Motion sensors	Spoofing	1	1	Low	
5.3 Force a stop by jamming or spoofing visual sensors to detect an object in front of the vehicle	Visual sensors	Spoofing	2	2	Medium	
5.4 Jamming primary sensors to force the AV into a safety stop	Primary sensors	Jamming DoS	2	2	Medium	
5.5 Disable the engine so AV stops	Engine control unit	DoS	1	2	Low	

| 5.6 Alter GPS position so the AV drives into an obstacle | GPS | Tampering | 3 | 2 | Medium |
| 5.7 Alter GPS mapping data so the AV drives into an obstacle | Mapping | Tampering | 2 | 2 | Medium |

Table 4: A risk assessment based on the attack objectives of an enemy

Note: Attacks are all remote unless specified otherwise in the 'target' column

4.4 Summary

It is easy to be hung up on complexity and the more technical aspects of AV systems and security, but it should not be forgotten that sometimes the simplest of enemy actions can be used to turn the system against itself. The highest risks identified in the risk assessment in table 4 are for relatively simple attacks.

More complicated attacks require sophisticated skill, tough to source equipment and higher budgets with a short window of attack. This makes the likelihood of the attacks being realised low, so even if the impact score is high, the final security level rating will not achieving above a medium.

Conversely if the knowledge, expertise, cost and equipment needed are low with a large window of opportunity then there is a high likelihood of success, however, if the attack is only successful in switching on a dashboard light to a minor sensor then the impact of this is very low, with again no 'high' final impact score.

It is only when an attack with low barriers to entry combines with consequence that are impactful that the risk assessment rating start to get noticeable. This combination of factors is seen in bringing the vehicle to a standstill just by walking in front of it, turning on the microphones in the vehicle to get troop discussions and in extracting AV movement history which are all 'high' rated in the risk assessment.

5. Countermeasures

Common and recommended ways in which risks can be reduced in general for cyber security applications include removing the attack surface altogether, adding a level of redundancy, having duplication, using encryption, requiring authentication and adding authorisation levels. Countermeasures for the various attacks identified using the ratings given in the risk assessment will now be discussed, starting with the highest rated risks.

5.1 Countermeasures for high risks

Attack 1.1: This describes a human essentially walking in front of the vehicle which, if using the same algorithms as civilian AVs, would force a stop situation. This is such a simple attack with no equipment or skills needed but for knowing the vehicle will stop in the presence of a person. Countermeasures against this would need to focus on the algorithm used in the AV and tailoring this to a warfare situation.

An initial countermeasure could be the ability of the AV, through training, to recognise the difference between civilians (protect) and enemy combatants (do not protect) and those surrendering (do not harm). As discussed in an interview with Paul Scharre [110], however, even if signs of surrender were added to the AVs algorithms (such as raised hands or white flag) these could be open to abuse from the enemy if not in the rules of war. Even if they were in the rules of war they may not be adhered to with this rule hard to enforce.

Another resolution could be for the AV not to stop for people when outside of a designated 'safe zone'. Human drivers would have a similar remit if coming across hostile forces within a warzone. Issues arise, however, if a person were not hostile such as playing children or signs of surrender have been given. These scenarios would be picked up quickly by a human driver but would be less obvious for artificial intelligence.

This is a non-simple problem to solve and further work will be required, not only into the technical aspects of writing an algorithm but the human aspects involved in this ethical and philosophical dilemma. Countermeasures to reduce the impact of a vehicle being taken are discussed in section 5.2.

Attack 4.1: This attack uses the infotainment unit to listen to troop discussions. It has been noted in section 3.7 already that the infotainment system contains many attack surfaces.

The countermeasure proposed for this attack is to remove the infotainment system completely from the AV, which will remove microphones from inside the vehicle with which to carry out surveillance. Removing this unit also has the benefit of reducing

the threat of other attacks on this system and attacks which use the infotainment system to gain access to the CAN bus.

With level 5 autonomy there is no need for humans to be in the vehicle, with cargo indifferent over being entertained or given information. If vehicles were to transport troops then an isolated system could be used to play CDs. To prevent attacks from this unit CDs used by troops would need screening or be provided by the military.

It would also been prudent for the isolated unit to be read only with no microphones or recording systems. CDs are an inherently transportable and ubiquitous means of storing data, so losses easily go undetected. Chelsea Manning, for example, used CD's to download 400,000 documents from US military operations [111], demonstrating even with an isolated unit there is the ability for security leaks to occur.

For information updates, again an isolated system separated from the CAN bus can be used when troops are being moved and this system is required.

Attack 4.2: Having the vehicles movement history available to the enemy would risk exposing critical allied locations. With the infotainment system removed from the vehicle, as discussed in the countermeasure to attack 4.1, this would give one less device with the potential to hold historic location information. A GPS receiver would also keep tracking data, however.

A countermeasure for this situation would be to have history of vehicle movements wiped form the GPS system after every mission when connected to the OBD port for updates or servicing. Removing this history would solve the problem of the enemy having multiple location data points but there would still be locations between memory being wiped and an AV being captured. To counter this map history on the device could be loaded with permanent random data which would act as noise to hide actual locations the vehicle had driven to. The noise on vehicles would obviously have to vary between AVs or an enemy capturing more than one vehicle could obtain genuine data just but taking the noise on the vehicles from each other.

There is still an issue if the vehicle has a 'return to base' function such as described in section 3.5.1 for a civilian setting - used if the vehicle has issues. In order for this function to be used the vehicle would need to know the co-ordinates of its base, which the enemy would also know if they captured the vehicle. This could be resolved by having these coordinates being sent OTA if needed from the base location.

This still would not prevent an attack in which the enemy placed a tracking device on the vehicle which would follow its journey to base. This is not a cyberattack so will not be considered further here.

5.2 Countermeasures for medium risks

Objective 1 – reducing impact parameters: Of note in these attacks is the initial condition stipulated that once a vehicle has been stopped or slowed down it is assumed it can be extracted by the enemy using a separate physical mechanism. This is not an easy operation, especially if armed troops are being carried by the AV.

A suitable countermeasure against this would be to fit sensors which could alert the AV to tampering situations such as tilt sensors warning if the vehicle is being moved into a towed position. Alternatively, an alert could be relayed to headquarters in the event of a dramatic reduction in speed or a stop, who could remotely override the vehicle sensors for example.

If the vehicle cannot be saved from being taken by the enemy, then it would be advantageous to stop the enemy being able to study or reverse engineer the technology or extract vital information from the vehicle. One way in which to reduce the impact of capture would be to install a self-destruction mechanism. This could come in different forms which are now discussed.

(1) Certain conditions could cause the vehicle to automatically self-destruct. This would only work under specific conditions which have been programmed into the AV. These conditions would have to be very carefully chosen as you would not want your fleet destroying themselves for no reason when on deployment. The conditions for self-destruct would also have to be kept top secret, this trigger would be a single point of weakness and if these were known by the enemy would cause devastation to allied supply line abilities.

(2) A remote dead switch could be triggered by command headquarters if the vehicle had been captured. This would stop devices self-destructing when not required but it would still create a single point of weakness, giving the enemy the ability to completely destroy all allied supply line vehicles if hacked. If used this would need to have robust encrypted channels with strong authorisation and authentication mechanisms associated with its use.

(3) There is also the option of having a manual self-destruct operated by a human occupant of the vehicle. This would work by entering a secure code into the system before the enemy could take control of a working vehicle. Another method would be to have a particular point which a bullet could be shot to wipe vehicle records and autonomous algorithms, as used in military computer systems [7]. As we are dealing with level 5 autonomy, however, then unless being transported to or from the front line there may not be any one in the vehicle to activate this 'kill switch' so it would not work in all scenarios. This mechanism is less open to abuse to enemy attackers, however.

(4) Multi-factor authentication could be required for accessing systems at rest. Should the authentication fail, the system could lock out an individual for an increasing period of time (such as on an iPhone) and ultimately wipe all data from the AV. Whilst the physical vehicle remains in-tact, any systems and code vulnerabilities would remain protected.

Given all the moving parts and complete destruction which could ensue with the installations of such a self-destruct function, it is decided more research on this is needed before using it as a countermeasure. Aviation have discussed using a pilot override scheme in commercial airlines [112] (not for destruction purposes it is emphasised) which is similar to remote access mechanisms used in point (2). This technology exists however pilots and companies refuse to use it. This technology would suffer from the same single point of failure as discussed for military vehicles risking complete devastation on a massive scale if the override mechanism were to fall into the hands of a malicious actor.

Objective 1 – reducing threat parameters: Reducing the impact of an attack would reduce the risk assessment rating, but to further reduce the risk the threat parameters could also be mitigated against.

Attack 1.2 is an attack on the TPMS. In a military environment the law requiring TPMS to be fitted would not apply, however, if the vehicle were to get a flat tyre in the supply chain route the AV would need to be able to sense this and deploy automatic inflation devices [108].

A countermeasure is to use Bluetooth to transmit TPMS data to the control panel instead of radio waves. Bluetooth has a shorter range than radio waves and would make remote enemy attacks more difficult to perform. This could be used alongside mechanical devices wired into the CAN bus which adds redundancy.

As discussed in section 4.2, the vehicle may be programmed to rank the sensor detecting a puncture over a sensor not showing an issue. This methodology may need reassessing depending on the outcome of a new risk assessment taking into account the countermeasures discussed so far.

Attack 1.3 has the attacker jamming and spoofing visual sensors to trick the AV into stopping, it believing an obstacle is blocking its path. The mitigations against these types of attack have been covered extensively in section 3.3.2 and will now be discussed.

For all physical sensors (cameras, radar, sonar and LiDAR) there is the ability to add additional sensors of the same type, giving the AV more data points about the surroundings and therefore more points an attacker would need to jam or spoof to make an attack realistic. Another method is to add redundancy in the type of sensor

so having a range of these sensors would again cause more work for an attacker - now having to overcome a range of devices in order to complete the attack successfully.

Platooning and swarming offer the opportunity to add redundancy to the system by preventing single units suffering failure by being jammed or spoofed – the collective sensors of the swarm providing the knowledge required. Instead of succeeding by compromising the sensors on a singular vehicle the attack would have to be scaled up in direct proportion to the number of linked vehicles. Aerial drones could also be used within the platooned or swarmed vehicle structure which would offer further visual sensors to give a check on the surrounding.

Countermeasures for attacks against vision sensors depend on the sensor being targeted, we will consider these in turn.

(1) Cameras could be installed with the ability to add different filters to improve vision quality and protect against laser attack. However, within the military technology is available which uses 'wavelength-agile' lasers with the ability to change colour making filtering of little use [70, 71].

(2) Radio attack involves the creation of 'ghost' vehicles using digital radio frequency memory repeaters. These could be countered by using filters and multiple wavelength scans which would cancel the effect of the repeater [74].

(3) For ultrasonic attacks, again an attacker could cause interference through sending inaudible frequencies to the sensor, causing them to be switched off. This could be countered by applying filters to the signals returned to the car or by spectral analysis [74].

(4) Specific mitigation against LiDAR attacks are to use LiDAR systems operating at multiple wavelengths which minimise jamming and spoofing with cheap devices [69]. Another solution is for the LiDAR to constantly change the interval between scanning speeds making it hard to synchronise the laser being used in an attack to the right frequency [59].

Attack 1.5 relies on jamming a primary sensor to force the vehicle into a safety stop. A mitigation against this has already been implemented in section 4.4.1 by taking out the infotainment system. This would reduce the attack surface on which an attacker could use to get onto the CAN bus and access safety critical devices.

Another way to achieve isolation of safety critical systems would be to have separate CAN bus networks for different devices of an AV. This is seen in aviation with four separate networks being used for: safety critical devices, environmental controls, door controls and infotainment [98]. Given the infotainment system has already been

isolated through mitigations in section 4.4.1 we could have separate networks within the vehicle relating to: safety critical devices, environmental sensors and vehicle access. This would allow access to safety critical ECUs only from other safety critical ECUs.

The use of separate CAN buses would need to be assessed in terms of benefits after a further risk assessment has being completed, including the countermeasures discussed so far, to see if this expensive process would be worth implementing. Having sperate networks to the system has to be completed during the design phase and cannot be retrofitted to the system. This would not only make having separate networks expensive but would also take a long time to reach active service.

Countermeasures for the network communications have been discussed in section 3.6.2 and include having authentication and authorisation mechanisms such as those used in network protocol architecture. Brute force password attacks could be mitigated against by having larger cryptographic keys or a more secure algorithm.

Attacks 1.6 and 1.7 rely on altering the GPS and the GPS mapping data respectively. Mitigations against this attack have been discussed in section 3.3.1. A simple countermeasure to prevent against map poisoning attacks would be the additional of authentication mechanisms for any map updates.

There are many countermeasures against GPS spoofing, such as using inertial navigation alongside GPS to validate changes in direction and location. Also including simple validation checks to monitor identification codes and timing intervals of satellite signals would permit the AV to check signals are behaving as expected [63]. The use of military-grade cryptographic authentication adds an extra layer of protection however some research suggests this can also be manipulated [56, 57].

Objective 2 attack: This attack relies on one of the attacks considered in objective 1 being successful. Mitigations discussed for these will therefore reduce the possibility of this attack being achievable given it cannot be completed without first capturing a vehicle.

Further countermeasures for this attack would be to have the OBD port check for malware which could be trying to infect the system. Technicians could also be required to check the movement history for the vehicle before plugging it into any central systems. This would show up anomalous locations and would require the attacker to complete further obfuscation to a vehicles data logs if they were to try and send an AV poisoned with malware back to an allied base.

Garages could also have separation of fleets, so if one of the OBD devices were compromised it would not spread to the whole of the military vehicles. This would

mean poisoning would be contained within a manageable percentage of vehicles so operations could still continue, although with restrictions.

Objective 3: Attacks 3.1 and 3.2 relate to mission data being unavailable and altered respectively with both having a medium level of risk associated with them. In mitigating against higher risk attacks, some mitigations against these will already have been done, including removal of the infotainment system so mission data will now not be linked to the CAN bus system.

However, for any supply line vehicles which would transport troops the requirement of isolated information devices would remain. Whilst isolated from the CAN bus, these would still be vulnerable to remote attack and the nature and type of attack would depend on the type of communication technology being used.

To prevent against signal jamming attacks limiting the flow of information and commands, redundancy could exist in the system. For example, there could be many different ways to communicate commands such as Wi-Fi, mobile and radio. This redundancy would also make attacks from spoofing harder to achieve if there are multiple, independent, sources of data providing information to the troops.

In addition, there would also be other, visual, information which could be used to supplement the AV systems for troop awareness. If a device was relaying that the roadway was clear, for example, the visual sensors on the AV would be able to confirm this.

Further countermeasures against attacks on Wi-Fi are discussed in section 3.6.1 and 3.6.2 which include algorithms to implement asymmetric cryptographic authentication mechanisms to ensure trust in communications. If communications were intercepted an attacker would have to decrypt the message for it to be readable, which would be a virtually impossible problem with for example AES correctly implemented.

Mobile and long-range radio communication countermeasures would also include authentication mechanisms, with encryption used to protect data when travelling over the air. This does have the effect of slowing down communication. The military would also have to be vigilant in protecting cryptographic keys and using best practices throughout the key management life cycle.

Objective 5: Attacks with a medium level risk include: 5.3 spoofing or jamming of visual sensors; 5.4 jamming of primary sensors; 5.6 and 5.7 altering GPS data and mapping data respectively. The countermeasures to these have already been considered in the attacks completed to achieve objective 1, in section 5.2.

5.3 Summary

Countermeasures for risks identified as low will not be considered individually. However, there are other countermeasures which have been discussed in chapter 3 which could be used in these situations in order to reduce risks to a lower level if the threat landscape were to change. This would be through, for example, technology becoming more widely available to the general population, devices used to perform attacks becoming cheaper or knowledge needed to perform an attacking being posted on the dark net or another publicly available channel.

It is noted that risk assessments should not be considered static and the threats and impacts need to be constantly assessed to ensure controls are in place to mitigate the risks identified remain appropriate.

Before countermeasures are employed into a system the knock-on effects to other parts of the system also need to be considered. These may not always have the effect of reducing risks everywhere and could even cause risks to increase by taking away a level of redundancy for example.

The impact on the usability of the area being hardened also needs to be considered. Use of cryptography, for example, is sometimes not used due to engineers reluctant to slow down communications, reduce access to a system or require the need for additional relay stations. There is also the cost of security to consider, as noted in section 4.2 for the separation of CAN bus networks. Sometimes countermeasures would improve security but the amount of cost and time this would require in order to be baked into the design from the start would be restrictively high.

In some situations, it would seem sensible to take out a component which acts as an attack vector, such as OTA updates, and just use the OBD tools when a vehicle has a service or when required. However, by doing this it removes a source of redundancy and if a vehicle is in the field and inaccessible when a vital update or communication needs to be passed on then this could be more of a risk than not having OTA technology.

Even though other sources of data increase the cost, it improves decision making and therefore safety. A challenge of this though as discussed with the TPMS is the fusion of these data sources and which sensor to use if there is a conflicting situation.

Attacks considered would also have a higher risk profile if it were reproducible. An enemy could invest significant capital expenditure to develop devices which could continuously perform the same attack repeatedly without an expert input for example. This would have to be compared against an attack which was cheaper to perform once but with ongoing costs, time and expertise required on each and every subsequent attack.

6. Conclusions

6.1 Conclusion

When reviewing published attacks in chapter 3 it is noted that they are very similar in sophistication to those which have been performed on computer systems before. That manufacturers have not thought of securing against such attacks is of concern. In 2016 research conducted into the state of cybersecurity in the automotive industry only 15% of those surveyed felt security was an integral part of the design process [113]. Whilst in some ways this is not surprising, with AVs being an emerging technology, add-ons have been made to existing vehicle designs which themselves are not particularly secure. It is fortunate that AVs have yet to see notable attacks from black hat hackers with attacks considered nearly all from researchers and white hat hackers. As AVs increase in number and connectivity, however, it is predicted that financial incentives will increase the levels of criminal activity in this space [59].

The security situation seen in AVs mirrors that observed when computers were first being connected to other computers and devices, with vulnerabilities only being found after the fact and design improved accordingly. AVs and vehicles in general are now experiencing a similar revolution in connectivity and the lessons learnt from computers can be applied by having security integral to the design.

This is difficult with initial design of a vehicle to scrappage being many years, and even more in a military setting. With AV capability being added to existing vehicle designs, this does not give the time window to rapidly respond to attacks and add more secure architectures. This type of response is observed in the mobile phone market with this product having turn-around times closer to two years than decades. For the military a way of upgrading encryption and other systems mid-way through a lifecycle would therefore be highly beneficial.

The security architecture in AVs will need to be more secure than any computer sat on an office desk. These mobile 'computers' will be let loose on the street or in a warzone, some weighing several tonnes with real possibilities of fatalities and significant physical damage if things go wrong. The airline industry will be an obvious point of reference in terms of safety and security standards for autonomy, with commercial airlines having auto pilot features fitted for many decades. Even here, however, the challenge of incorporating increasingly sophisticated flight systems remains. These challenges have recently hit the headlines through two fatal crashes involving the Boing 737 MAX 8 aircraft. It is suspected that faulty sensors trigged the flight control system to push the nose down to avoid stalling which could not easily be physically over-ridden by the pilots [114].

With military AVs operating within a supply chain desert warzone, the risk assessment was completed using enemy objectives in terms of what they would want to achieve

through an attack. This led to methods of attack being formulated which would cause a compromise within the AV to realise these objectives. The impact of this success along with a realistic assessment of the threat this posed was conducted and rated using a modified version of the HEAVENS model [106]. For the military political impact was included in the assessment as well as how much financing is required to cause the threat, both of which are important parameters.

Despite all of the potential attacks and levels of resource and complexity available to an enemy attacker it was noted that the most feasible and impactful attacks were surprisingly simple and low tech. These focussed upon either utilising the infotainment system, or relying upon the priority given to civilian safety by the AV.

For the majority of vehicle outings only cargo would be transported through the supply chain, which would not require the infotainment system. By isolating this system there is less complexity and attack vectors are removed making the AV more secure.

In terms of response to people within a battlefield, the algorithms required for a military setting will need to be adjusted from those used in a civilian environment. As was seen in the risk assessment, all that is needed to stop a military AV would be for a person to step in front of it to cause it to stop, resulting in it becoming a 'sitting duck' ready for capture. How the artificial intelligence will deal with this situation is crucial in military AV deployment. This has the potential to cause 'friendly' casualties as well as undermining faith in the technology by assessing a situation differently than a human driver would. Even if signs of surrender were programmed in these have the potential to be abused by the enemy [110].

Additionally, military AVs will need extra means of protecting the technology in case they are captured by the enemy. The possibility of an enemy reverse engineering technology and using this against allies later is obviously not a desirable situation. This is not a hypothetical attack and the consequences of such actions have been demonstrated with Iran able to capture a US Sentinel drone in December 2011 and reverse engineering this [65]. Information on the devices would also need to be carefully protected so allied locations or mission critical data could not be determined from the vehicle.

It was also noted that safety critical features should be isolated from other parts of the network with separate CANs for environmental sensors, vehicle access and safety critical devices. It is more secure to have these as completely separate network systems such as used in the airline industry [98]. This would require extra expense in design and manufacture as well as for updates and maintenance, but it would stop minor systems, such as the TPMS, being able to compromise a critical component, such as the engine control.

With the number of vehicles in a civilian environment physical updates for AVs would be unmanageable and would require OTA or home Wi-Fi updates. In the military there are not only fewer vehicles but these need to be services regularly, especially in a harsh desert environment. This would allow for most updates to be done by a technician using the OBD port. That is not to say the OTA update facility could be disabled though, with redundancy desirable in situations where vehicles could be away for long periods unable to connect to a physical update source. The OBD port would be used in most situations, however, giving a more secure update facility.

The military environment offered an opportunity to remove attack surfaces within the AV such as for infotainment. In other areas such as OTA updates it was determined more prudent to leave in the technology to give redundancy to the system. The subject of redundancy was highly evident in countermeasures to sensor attack from jamming and spoofing. Cameras, radar, sonar, ultrasonic and LiDAR all benefit from having redundancy not only within their own technology but also by using overlapping technologies from different wavelengths. Using other sources of data increases costs but is worthwhile as it significantly improves the decision making and thus safety of an AV. One challenge is how the fusion of all these data sources can be done in order to converge to the most appropriate action however, as noted with the TPMS.

The dilemma between removing and adding systems and sensors is a constant juggling act. So too is the decision to add extra security, such as cryptographic authorisation and authentication mechanisms at the expense of technology functioning and speed. Mitigation techniques such as intrusion detection systems or antijamming GPS technology need either a software update or equipment to be changed. A downside of these countermeasures, however, is the increase in computation overhead in time and power.

6.2 Further research

On completing the mapping from civilian to military environments areas which would need to be further investigated have been discovered. These will now be discussed.

Single-point of failure: Some militaries have multiple suppliers for various different types of vehicles, which ensures a level of redundancy if technical faults were found with any of the products. If all trucks needed to be taken out of service due to a single supplier fault this would hit operations. Having three separate suppliers for example would allow continuing functioning of operations so there is not a single point of failure. This separation technique has also been mentioned with the example of updating vehicles using separate OBD mechanics in-case one has been compromised.

This would also be true in the suppliers of the technology for AVs. So, if the enemy captures a vehicle and engineers an exploit to that particular algorithm, variety in

systems would ensure not all vehicles have to be removed whilst this vulnerability is fixed. Even though this redundancy is desirable it adds to supply chain issues and also the ability to keep the technology secret across additional organisations.

Supply chains: The modern vehicle has an extremely complicated supply chain with part manufactured by many different suppliers. How the security of every component within the supply chain can be assures poses difficult challenges.

Self-destruct mechanism: A self-destruct feature was discussed at length in section 5.2, with four designs suggested. It was decided, however, that more research was needed into this before such a feature was installed onto an AV.

Swarming: Swarming was discussed as a way of mitigating against attacks on sensors in order to give redundancy. Its use in warfare has received much press with the ability to operate cooperatively and more effectively than a single device [43] and with no long-range communication to base needed, reducing remote attack threats. This technology would rely on V2V communication channels and algorithms for the 'intelligent' behaviour.

Interoperability: There are increasing links between the different domains of war which are land, sea, air, space and cyberspace. As the cyberspace dependence increases so does the coordination this allows within the other domains. The obvious interconnection of these domains would be a logical next step, however, these would need to use technologies which would work across all the domains, not only for one country but for coalitions of forces working together in a warzone. The security architecture for such a system would also need to be risk assessed using similar methodology used in this report.

Lethal uses of AVs: Use of lethal ground vehicles was not covered in this project, due to the UK military noting they would always want human authorisation before any weapons are released. However, on 16th July 2019, Military Aerospace ran an article headed "U.S. military shifting research and technology development toward armed robotic ground vehicles" [115], with field tests planned within a year.

With technology becoming more sophisticated attitudes are changing towards their use and this development heightens the urgency with which cybersecurity needs to be considered and baked into the design. A risk assessment completed on a vehicle fitted with high explosives or hundreds of rounds of ammunition would vary significantly from a benign supply vehicle. Risks highlighted with this use case would be essential before these 'killer robots' become a reality so the technology can be secure.

Failure modes: If an AV is hit by a cyberattack in a civilian setting it will usually be in a position where assistance can be easily called, in a benign environment. However,

with a military supply chain this is not the case and the ability to recover, or not fail catastrophically, is an essential feature.

If the AV 'safe setting' is to return to base for example this could be exploited by the attacker to find allied locations or using the vehicle in a 'Trojan Horse' type attack. If a vehicles failure is to not move until a system is fully functional then there may be many inactive AVs in the desert awaiting repair. Research is therefore needed to identify ways in which military AVs could be programmed to 'fail safely' within a hostile environment.

6.3 Closing thoughts

Is attacking an AV easier in a cyber environment? There is definitely more complexity, which gives rise to more attack surfaces which could cause issues. However, does this scale up to more attacks?

In a civilian environment when AVs become available the main motivation for vehicular attacks will be financial. Most people are benign and not out to cause devastation, random physical attacks on conventional vehicles not being a major problem in society. However, in a military setting there is a motivation to attack an enemy using cyber technologies with cyber skills being increased in countries such as the US, China and Iran. Many articles and books point to the importance of cyberspace in future wars, with the UK Defence Secretary Michael Fallon quoted as saying "Cyberattack is one of the greatest challenges to our security" [116].

The benefits of AVs could be revolutionary, but they need to be designed with security in mind from the outset. The current situation is summarised by Peter Davies of Thales which recognises that with all the complexities of AVs there is never going to be complete safety and we need to make sure when AVs do fail, they will be safe, and the system can recover.

"[As] It is expected that AV will be compromised it is ensuring the failures aren't catastrophic and knowing how to recover from this when it occurs. The AV will only be safe if we have justifiable and enduring confidence they will do what is expected and when we want this." [24]

Appendices

Information within this appendix is derived from pages 77-86 of the Surface Vehicle Recommended Practice by SAE International, Cybersecurity Guidebook for Cyber-Physical Vehicle Systems [106].

Appendix 1: The HEAVENS security model

The HEAVENS security model is used for cybersecurity of vehicle electrical or electronic systems and focuses on threat analysis and risk assessment. The model applies the Microsoft STRIDE approach, shown in table 5, in the context of vehicle systems and establishes a direct mapping between security attributes and threats. Security objectives (such as financial, safety and privacy) are mapped with an estimate of impact levels.

STRIDE threats	Explanation	Security attributes
Spoofing	Attackers pretend to be someone or something else	Authentication, freshness
Tampering	Attackers change data in transit or in a data store, attackers may change functions as well – implemented in software, firmware or hardware	Integrity
Reputation	Attackers perform actions that cannot be traced back to them	Non-repudiation, freshness
Information disclosure	Attackers get access to data in transit or in a data centre	Confidentiality, privacy
Denial of Service	Attackers interrupt a system's legitimate operation	Availability
Elevation of privilege	Attackers perform actions they are not authorized to perform	Authorisation

Table 5: Mapping between STRIDE threats and security attributes

The HEAVENS model derives scores for threat level and impact level parameters as described in appendix 2 and appendix 3 respectively and uses these to derive a security level for each using table 3 in section 4.1. The security level calculation is a measure of the strength of security needed to meet the risk identified.

Appendix 2: Threat level parameters and scoring

Determining the threat level is the first step in completing the HEAVENS model, the threat level corresponds to the estimation of the 'likelihood' of a risk being realised. Parameters used to determine threat along with their scorings are now described.

Expertise parameter: The expertise indicates the knowledge level needed to carry out an attack on a system, including underlying principles, product type or attack methods.

Parameter	Value	Explanation
Layman	0	No particular expertise or knowledge about a system needed. Able to follow simple instructions included in available tools needed to conduct simple attacks. If tools or instructions not performing as expected would not be able to correct themselves.
Proficient	1	Would know about simple and popular attacks and general knowledge about security. Capable of performing attacks using tools available and can improvise in order to achieve result required.
Expert	2	Have familiarity with underlying algorithms, cryptography, protocols, hardware, structures, security behaviours, principles and concepts of security employed. They would know tools available and techniques which could be used for new attacks as well as know classical attack methods.
Multiple experts	3	This defines the situations where more than one expert would be needed in order to complete different distinct steps of an attack.

Table 6: Expertise threat level parameter rating

Knowledge about the system: How easily and widely available information about the system is in terms of sources where an attacker can find information.

Parameter	Value	Explanation
Public	0	Information from the internet, generally available books or shared without needing a non-disclosure agreement.
Restricted	1	Knowledge, such as for design specifications, which is controlled within the developer organisation but shared with other organisations, such as suppliers under a non-disclosure agreement.
Sensitive	2	Information shared only between specific teams or people in the developer organisation such as parameters to enable or disable vehicle features, software source codes or configuration databases.
Critical	3	Knowledge about the system which is only known by a few people such as secret root signing keys. Information would be tightly controlled and only need to be given out on a strict need to know basis depending on individual assignments.

Table 7: Knowledge about the system threat level parameter rating

Window of opportunity: This combines the type of access required, such as physical or logical, and assesses the amount of time typically available for the attack to take place.

Parameter	Value	Explanation
Critical	0	Unlimited physical access is available, as well as logical or remote access, with the system always accessible without any time limitation, for example mobile connections.
High	1	No physical access is needed with an attack possible with logical or remote access. Time is limited but the system is easily available.
Medium	2	Limited physical access to the interior or exterior of the vehicle requiring but no special tools, for example accessing wires under the bonnet. This could also include limited logical access to the system.
Low	3	Physical access is needed to complete an attack which requires complex vehicle part disassembly in order to get access to internal systems.

Table 8: Window of opportunity threat level parameter rating

Equipment: The type of equipment required to identify or exploit a vulnerability or mount an attack is considered in terms of how specialist this is.

Parameter	Value	Explanation
Standard	0	Readily available as part of the system itself or obtained through internet download or attack scripts, for example simple OBD devices.
Specialised	1	Not readily available but can be obtained without much effort such as buying various equipment in moderate numbers or developing more extensive programs or attack scripts.
Bespoke	2	Not available to the public and may be very expensive. Would need to be specially produced or even be so specialised that distribution is controlled or restricted.
Multiple bespoke	3	Different types of bespoke equipment are needed in order to perform distinct steps in an attack.

Table 9: Equipment threat level parameter rating

Cost to perform: This indicates how much funding an attacker needs to have in order to carry out a successful attack on the system.

Parameter	Value	Explanation
Low	0	Very low cost, within reach of an individual.

Medium	1	Costs are substantial to fund for an individual but within easy reach of an organisation.
High	2	Substantial funding needed but could be completed within the department budget but drawing a large amount of funds.
Very high	3	Funding of the attack puts financial strain on the organisation and probably causes funding to other projects to be reduced to pay for this.

Table 10: Cost to perform threat level parameter rating

Threat level value: Modification of the HEAVENS model scoring for threat level is shown in table 11. To adjust for adding a parameter the scoring levels were shifted by 1 point to reflect this so 'critical' now reads 0-2 as opposed to 0-1 and 'low' being now 8-10 rather than 7-9 for example.

Summation of threat level parameters	Threat level	Threat level value
>10	None	0
8-10	Low	1
5-7	Medium	2
3-4	High	3
0-2	Critical	4

Table 11: Mapping of threat level parameter totals to threat level value

Appendix 3: Impact level parameters and scoring

Determining the impact level is the second step in completing the HEAVENS model, the impact level corresponds to the estimation of the 'impact' if a risk is realised. Parameters used to determine impact along with their scorings are now described.

Safety: A fundamental requirement is to ensure safety of vehicle occupants and those in the vicinity of the AV.

Parameter	Value	Explanation
No impact	0	No injury.
Low	3	Light and moderate injuries.
Medium	30	Severe injuries. Life threatening injuries but where survival is likely.
High	300	Fatal injuries. Life threatening injuries but where survival is uncertain.

Table 12: Safety impact level parameter rating

Financial: Considers total financial losses from direct damages such as recalls, loss of business and fines under legislation penalties and indirect losses from reputational damage, loss of market share and intellectual property infringements.

Parameter	Value	Explanation
No impact	0	No effects or consequences noticeable.
Low	1	Financial damage but low level.
Medium	10	Damage resulting from an attack leads to substantial losses but can be financed within the department budget.
High	100	Financial damage puts strain on the department and funding from other areas may be needed to cover losses.

Table 13: Financial impact level parameter rating

Political: Considers political damage such as through troop and civilian unrest as well as the impact on political decision making.

Parameter	Value	Explanation
No impact	0	No appreciable political impact.
Low	1	Troops voice concerns due to AV workings but only minor, few if any news stories report on these and politicians not involved.
Medium	10	Troops concerns grow with numerous news stories covering issues. Civilians discussing the issues as are politicians.
High	100	Protesting from civilians on home territory, troop moral suffering over safety fears. MPs hold emergency meetings to try and resolve issues.

Table 14: Political impact level parameter rating

Operational: Damages caused by loss of vehicle functions such as cruise control, air conditioning and CD-player. Could also be more serious such as loss of primary vehicle functionality which may affect the vehicle functioning safely.

Parameter	Value	Explanation
No impact	0	No noticeable effects
Low	1	Vehicle operates but does not conform with warning light or audible noise affecting 25-75% of AVs.
Medium	10	Vehicle still operable but there is degradation or loss of secondary functions with comfort or convenience functions degraded or not working. Or vehicle inoperable with loss of primary function but with vehicles safety not affected.
High	100	Loss of primary function causing AV to become inoperable. Safety mode functioning affects AVs safe operation (with or without warnings) or noncompliance with regulations.

Table 15: Operational impact level parameter rating

Privacy and legislation: Damages caused by privacy, legislation or regulation violation. This parameter can have a financial impact from fines and operational damage associated with it, but usually no direct injury is caused.

Parameter	Value	Explanation
No impact	0	No noticeable effects in terms of violations.
Low	1	A particular individual has their privacy violated but it may not be used in criminal acts. Legislation violated but without having noticeable operational or financial impact on the business or stakeholders.
Medium	10	A particular individual has their privacy violated and used in criminal acts leading to media coverage. Legislation violated with potential consequences for operations and finances such as penalties and loss of market share.
High	100	Multiple individuals have their privacy violated and used in criminal acts such as identity theft. May lead loss of market share, loss of trust and reputational damage. Legislation violation causes significant consequences for operations and finance as well as extensive media coverage.

Table 16: Privacy and legislation impact level parameter rating

Summation of impact level parameters: The scoring system of the HEAVENS model has been adjusted for adding of another parameter for political impact and also reducing the influence of safety and finance, however, the 'no impact' level will remain the same.

The start of the medium category will start at 16 which would be one parameter, other than safety, being at a medium level with the others being of a low level (10+3+1+1+1=16). The high impact level will begin at a total of 44 being the score of four parameters being at a medium level with safety and another parameter at a low level (10+10+10+10+3+1=44) or the safety parameter and another parameter being at medium impact and the other four at a low level (30+10+1+1+1+1=44). The critical level will begin at over 250 which relates to two parameters, other than safety, being at a high level and the remaining parameters being medium impact (100+100+30+10+10+10=250). Scoring for impact level is shown in table 17.

Summation of impact level parameters	Impact level	Impact level values
0	No Impact	0
1-15	Low	1
16-33	Medium	2

| 34-249 | High | 3 |
| >249 | Critical | 4 |

Table 17: Mapping of impact level parameter totals to impact level values

Appendix 4: Calculations of threat level and impact level

The threat level and impact level scores for a military environment will now be completed. Each objective of the enemy will be considered with attacks identified which would allow the successful completion of the objective. Threat and impact scores will then feed into the security level calculation, shown in table 3, which is the third and final step in the process of completing the HEAVENS security model.

Objective 1: Capture of an AV, troops or supplies

1.1 Person walks in front of the AV to enforce a stop situation
1.2 Flat tyre spoofed to force the AV to stop or slow down
1.3 Force the AV to slow by spoofing motion sensors to show the AV is on a slope
1.4 Force a stop by jamming or spoofing visual sensors to detect an object in front of the vehicle
1.5 Jamming primary sensors to force the AV into a safety stop
1.6 Disable the engine so the AV stops
1.7 Alter GPS position so the AV drives into an obstacle
1.8 Alter GPS mapping data so the AV drives into an obstacle
1.9 Spoof GPS location so the AV drives into an enemy camp

Attack	Expertise	Knowledge	Window	Equipment	Finance	Total	TL	TL value
1.1	0	0	2	0	0	2	Critical	4
1.2	1	1	2	1	0	5	Medium	2
1.3	2	2	3	2	1	10	Low	1
1.4	1	1	2	1	0	5	Medium	2
1.5	2	1	2	1	1	7	Medium	2
1.6	2	2	2	1	1	8	Low	1
1.7	1	0	1	1	1	4	High	3
1.8	1	1	1	1	1	5	Medium	2
1.9	2	2	2	1	1	8	Low	1

Table 18: Threat level ratings for capture scenarios

Attack	Safety	Financial	Political	Operational	Privacy/ legislation	Total	IL	IL value
1.1	0	1	10	1	10	22	Medium	2
1.2	0	1	10	1	10	22	Medium	2
1.3	0	1	10	1	10	22	Medium	2
1.4	0	1	10	10	10	31	Medium	2
1.5	0	1	10	10	10	31	Medium	2
1.6	0	1	10	10	10	31	Medium	2

1.7	3	1	10	10	10	33	Medium	2
1.8	3	1	10	10	10	33	Medium	2
1.9	0	1	100	10	10	121	High	3

Table 19: Impact level ratings for capture scenarios

Objective 2: Return the captured AV to base and poison other units

Attack	Expertise	Knowledge	Window	Equipment	Finance	Total	TL	TL value
2	2	2	2	1	1	8	Low	1

Table 20: Threat level ratings for a fleet poisoning scenario

Attack	Safety	Financial	Political	Operational	Privacy/ legislation	Total	IL	IL value
2	300	100	100	100	10	610	Critical	4

Table 21: Impact level ratings for a fleet poisoning scenario

Objective 3: Cause confusion and break command

3.1 Mission data made unavailable

3.2 Mission data altered

3.3 Activation of in vehicle systems

3.4 Force erratic AV movements through engine control unit or accelerometers

3.5 Force erratic AV movements through visual sensors or GPS

Attack	Expertise	Knowledge	Window	Equipment	Finance	Total	TL	TL value
3.1	0	0	1	0	0	1	Critical	4
3.2	2	1	2	1	1	7	Medium	2
3.3	1	1	1	1	0	4	High	3
3.4	2	2	2	2	2	10	Low	1
3.5	1	1	1	1	1	5	Medium	2

Table 22: Threat level ratings for a scenario to cause confusion

Attack	Safety	Financial	Political	Operational	Privacy/ legislation	Total	IL	IL value
3.1	0	0	1	0	0	1	Low	1
3.2	30	1	1	1	1	34	High	3
3.3	0	0	1	1	0	2	Low	1
3.4	1	0	1	1	0	5	Low	1
3.5	1	0	1	1	0	5	Low	1

Table 23: Impact level ratings for a scenario to cause confusion

Objective 4: Surveillance

4.1 In vehicle discussions of troops obtained

4.2 A history of AVs recorded movements is obtained

4.3 The enemy is able to see AV movements through camera feeds

Attack	Expertise	Knowledge	Window	Equipment	Finance	Total	TL	TL value
4.1	1	1	0	0	0	2	Critical	4
4.2	1	1	0	0	0	2	Critical	4
4.3	2	2	0	1	1	8	Low	1

Table 24: Threat level ratings for a surveillance scenario

Attack	Safety	Financial	Political	Operational	Privacy/ legislation	Total	IL	IL value
4.1	3	0	10	0	10	23	Medium	2
4.2	3	0	10	0	10	23	Medium	2
4.3	30	0	10	0	10	50	High	3

Table 25: Impact level ratings for a surveillance scenario

Objective 5: Disable or destroy an AV:

5.1 Flat tyre spoofed to force the AV to stop or slow down

5.2 Force the AV to slow by spoofing motion sensors to show the AV is on a slope

5.3 Force a stop by spoofing visual sensors to detect an object in front of the vehicle

5.4 Jamming primary sensors to force the AV into a safety stop

5.5 Disable the engine so AV stops

5.6 Alter GPS position so the AV drives into an obstacle

5.7 Alter GPS mapping data so the AV drives into an obstacle

Attack	Expertise	Knowledge	Window	Equipment	Finance	Total	TL	TL value
5.1	1	1	2	1	0	5	Medium	2
5.2	2	2	3	2	1	10	Low	1
5.3	1	1	2	1	0	5	Medium	2
5.4	2	1	2	1	1	7	Medium	2
5.5	2	2	2	1	1	8	Low	1
5.6	1	0	1	1	1	4	High	3
5.7	1	1	1	1	1	5	Medium	2

Table 26: Threat level rating for a scenario to destroy or disable an AV

Attack	Safety	Financial	Political	Operational	Privacy/ legislation	Total	IL	IL value
5.1	0	1	10	1	0	12	Low	1
5.2	0	1	10	1	0	12	Low	1
5.3	0	1	10	10	0	21	Medium	2
5.4	0	1	10	10	0	21	Medium	2
5.5	0	1	10	10	0	21	Medium	2
5.6	3	1	10	10	0	24	Medium	2
5.7	3	1	10	10	0	24	Medium	2

Table 27: Impact level rating for a scenario to destroy or disable an AV

List of definitions and acronyms

ABS	Antilock Braking System
AES	Advanced Encryption Standard
AV	Autonomous Vehicle
Bluetooth	Standard for short-range wireless interconnection of electronic devices
CA	Certificate Authorities A trusted source that can hand out and revoke public keys
CAN	Controller Area Network
DAB	Digital Advanced Broadcasting
DARPA	Defence Advanced Research Projects Agency
DDoS	Distributed Denial of Service
DoD	Department of Defence
DoS	Denial of Service
DSRC	Dedicated Short-Range Communication
ECU	Electronic Control Unit
EMP	Electro Magnetic Pulse
GDPR	General Data Protection Regulation
GPS	Global Positioning System
IED	Improvised Explosive Devices
IL	Impact Level
INS	Inertial Navigation System
IP	Internet Protocol
LASER	Light Amplification by Stimulated Emission of Radiation
LiDAR	Light Detection and Ranging
LTC	Long Term Certificate Contains vehicle identifiers and can be revoked
OBD	On Board Diagnostics

List of definitions and acronyms (continued)

OEM	Original Equipment Manufacturer
OFFSET	Offensive Swarm Enabled-Tactics
OTA	Over The Air
PC	Pseudonym Certificate Used for anonymous transfers of common messages like braking
RADAR	Radio Detection And Ranging
RFID	Radio Frequency Identification
RSU	Road Side Unit
SL	Security Level
TL	Threat Level
TPMS	Tyre Pressure Monitoring System
UAV	Unmanned Aerial Vehicles
V2I	Vehicle to Infrastructure
V2N	Vehicle to Networks
V2P	Vehicle to Pedestrians
V2V	Vehicle to Vehicle
V2X	Vehicle to Everything
VANET	Vehicle Ad-hoc NETwork
Wi-Fi	Allows devices to communicate with one another wirelessly

References

[1] R. Charette, (2009, February 1), "This Car Runs of Code,", IEEE Spectrum. [Online]. Available: https://spectrum.ieee.org/transportation/systems/this-car-runs-on-code [Accessed 12 August 2019].

[2] S. Romero, (2018, December 31), "Wielding Rocks and Knives, Arizonans Attack Self-Driving Cars," The New York Times. [Online]. Available: https://www.nytimes.com/2018/12/31/us/waymo-self-driving-cars-arizona-attacks.html [Accessed 12 August 2019].

[3] D. Araya, (2019, January 29), "The Big Challenges In Regulating Self-Driving Cars". [Online]. Available: https://www.forbes.com/sites/danielaraya/2019/01/29/the-challenges-with-regulating-self-driving-cars/#72010ed4b260 [Accesses 12 August 2019].

[4] G. Barkho, (2019, April 13), "The Army Is Getting Driverless Vehicles Before the Public," Observer. [Online]. Available: https://observer.com/2019/04/army-driverless-vehicles-public-debut/ [Accessed 12 August 2019].

[5] G. Solis, (2016), "The Law of Armed Conflict: International Humanitarian Law in War", 2nd ed. Cambridge University Press.

[6] K. Gronlund, (2019, May 9), "State of AI: Artificial Intelligence, the Military and Increasingly Autonomous Weapons," Future of Life. [Online]. Available: https://futureoflife.org/2019/05/09/state-of-ai/ [Accessed 12 August 2019].

[7] VanHalteren, J. (2019, August 12), "Compiled Report." [email].

[8] UNESCO Mediabank, "Patent DRP 37435 "Vehicle with gas engine operation" submitted by Carl Nenz, Mannheim. Dated 29 January 1886," UNESCO Mediabank. [Online]. Available: https://en.unesco.org/mediabank/25010/ [Accessed 12 August 2019].

[9] B. McCalley, (1994), "Model T Ford: The Car That Changed The World," Motorbooks Intl.

[10] Encyclopaedia Britannica, (2018). "Encyclopaedia Britannica," Encyclopaedia Britannica, Inc..

[11] J. Shuttleworth, (2019, January 7), "SAE Standard News: J3016 automated-driving graphic update," SAE.org. [Online]. Available: https://www.sae.org/news/2019/01/sae-updates-j3016-automated-driving-graphic [Accessed 12 August 2019].

[12] Tesla Motors, (2016, October 19), "All Tesla Cars Being Produced Now Have Full Self-Driving Hardware, " Tesla Motors. [Online]. Available: https://www.tesla.com/en_GB/BLOG/ALL-TESLA-CARS-BEING-PRODUCED-NOW-HAVE-FULL-SELF-DRIVING-HARDWARE [Accessed 12 August 2019].

[13] Waymo, "Our Journey", Waymo. [Online]. Available: https://waymo.com/mission/ [Accessed 12 August 2019].

[14] Gov.uk, (2018, May 25), "Guide to the General Data Protection Regulations," Gov.uk. [Online]. Available: https://www.gov.uk/government/publications/guide-to-the-general-data-protection-regulation [Accessed 12 August 2019].

[15] C. Cadwalladr and E. Graham-Harrison, (2018, March 17), "Revealed: 50 million Facebook profiles harvested for Cambridge Analytica in major data breach," The Guardian. [Online]. Available: https://www.theguardian.com/news/2018/mar/17/cambridge-analytica-facebook-influence-us-election [Accessed 12 August 2019].

[16] H. Lipson and M. Kurman, (2016), "Driverless: Intelligent Cars and the Road Ahead," The MIT Press.

[17] Statistica, "Number of passenger cars and commercial vehicles in use worldwide from 2006 to 2015 in (1,000 units)," Statista Inc.. Available: https://www.statista.com/statistics/281134/number-of-vehicles-in-use-worldwide/ [Accessed: August 12 2019].

[18] Global Fire Power, (2018), "United States Military Strength," Global Fire Power. [Online]. Available: https://www.globalfirepower.com/country-military-strength-detail.asp?country_id=united-states-of-america [Accessed 12 August 2019].

[19] T. Chung, "OFFensive Swarm-Enabled Tactics (OFFSET)", DARPA. [Online]. Available: using swarms comprising 250 unmanned aircraft [Accessed 12 August 2019].

[20] D. Jordan, J. Kirasm D. Lonsdale, I. Speller, C. Tuck and C. Walton (2016), "Understanding modern warfare", 2nd ed. Cambridge University Press.

[21] S. Motoyama, (2018, August 27), "Inside the United Nations effot to regulate autonomous killer robots," The Verge. [Online]. Available: https://www.theverge.com/2018/8/27/17786080/united-nations-un-autonomous-killer-robots-regulation-conference [Accessed 12 August 2019].

[22] K. Gronlund, (2019, May 9), "State of AI: Artificial Intelligence, the Military and Increasingly Autonomous Weapons," Future of Life Institute. [Online]. Available: https://futureoflife.org/2019/05/09/state-of-ai/ [Accessed 12 August 2019].

[23] C. Woody and N. Mead, "Using Quality Metrics and Security Methods to Predict Software Assurance," Software Engineering Institute, Carneigie Mellon University, [Online]. Available: https://insights.sei.cmu.edu/sei_blog/2016/06/using-quality-metrics-and-security-methods-to-predict-software-assurance.html [Accessed 12 August 2019].

[24] P. Davies, (2018, December), "The connected vehicle as in IoT infrastructure - cyber security in a complex socio-Technical system", IEEE Conference, London.

[25] V. Kalyani, M. Bhatnagar, L. Shivnani and E. Verma, (2016, October), "Significance of Electronics and Luxury Car Sophistication through ECU: A Progressive Study of Evolution in Automobile Industry," Management Engineering and Information Technology, vol. 3, issue 5, p2.

[26] BBC.co.uk, (2019, August 8), "Lorry stuck on notoriously steep Tower Hill," BBC.co.uk. [Online]. Available: https://www.bbc.co.uk/news/av/uk-wales-49278763/lorry-stuck-on-notoriously-steep-tower-hill [Accessed 12 August 2019].

[27] GPS.gov, "GPS Accuracy," GPS.gov. [Online]. Available: https://www.gps.gov/systems/gps/performance/accuracy/ [Accessed 12 August 2019].

[28] Wired, (2018, June 2), "What is LiDAR, why do self-driving cars need it, and can it see Nerf bullets?," Wired. [Online]. Available: https://www.wired.com/story/lidar-self-driving-cars-luminar-video/ [Accessed 12 August 2019].

[29] M. Jokela, M. Kutila and P. Pyykonen, (2019, June 7), "Testing and Validation of Automotive Point-Cloud Sensors in Adverse Weather Conditions," Applied sciences, MDPI. [Online]. Available: https://www.mdpi.com/2076-3417/9/11/2341 [Accessed 12 August 2019].

[30] Electronics, Project, Focus, "All You Know About LiDAR Systems and Applications," Electronics, Project, Focus. [Online]. Available: https://www.elprocus.com/lidar-light-detection-and-ranging-working-application/ [Accessed 12 August 2019].

[31] M. Kamelska, (2017, August 3), "RADAR vs. LIDAR sensors in automotive industry," Stanford management science and engineering. [Online]. Available: https://mse238blog.stanford.edu/2017/08/mj2017/radar-vs-lidar-sensors-in-automotive-industry/ [Accessed 12 August 2019].

[32] A. Hawkins, (2019, April 24), "It's Elon Musk vs. everyone else in the race for fully driverless cars," The Verge. [Online]. Available: https://www.theverge.com/2019/4/24/18512580/elon-musk-tesla-driverless-cars-lidar-simulation-waymo [Accessed 12 August 2019].

[33] NewsCore, (2012, May 11), "How fast are you really going? The accuracy of speedometers," Fox News. [Online]. Available: https://www.foxnews.com/auto/how-fast-are-you-really-going-the-accuracy-of-speedometers [Accessed 12 August 2019].

[34] OXTS, "What is an inertial navigation system," OXTS. [Online]. Available: https://www.oxts.com/what-is-inertial-navigation-guide/ [Accessed 12 August 2019].

[35] Tyresafe, "The Law," Tyresafe. [Online]. Available: https://www.tyresafe.org/tyre-safety/tpms/tpms-and-the-law/ [Accessed 12 August 2019].

[36] (2011, June), "Crash prevention effectiveness of light-vehicle electronic stability control: An update of the 2007 NHTSA evaluation", US Department of Transportation. [Online]. Available:

https://crashstats.nhtsa.dot.gov/Api/Public/ViewPublication/811486 [Accessed 12 August 2019].

[37] "Mobility and Transport: Electronic stability control," European Commission. [Online]. Available: https://ec.europa.eu/transport/road_safety/specialist/knowledge/esave/esafety_meas ures_known_safety_effects/electronic_stability_control_en [Accessed 12 August 20190.

[38] (2014, November 27), "Road safety wins as all new vehicles are now equipped with electronic stability control," Ertico. [Online]. Available: https://erticonetwork.com/road-safety-wins-as-all-new-vehicles-are-now-equipped-with-electronic-stability-control-1/ [Accessed 12 August 2019].

[39] M. Eder, M. Wolf. "V2X communication overview and V2I traffic light demonstrator," Munich University of Applied Sciences. [Online]. Available: https://pdfs.semanticscholar.org/bd8d/1639cbc599afe9fca8ef8d83afdc8515a3c1.pdf [Accessed 12 August 2019].

[40] Y. Duan, (2017, November 2017), "Truck Platooning: The Band of Semi-Trailers," Labroots. [Online]. Available https://www.labroots.com/trending/chemistry-and-physics/7405/band-semi-trailers-truck-platooning [Accessed 12 August 2019].

[41] (2019, February 11), "Gavin Williamson: Drone 'swarm squadrons' to be deployed by military," BBC News. [Online]. Available: https://www.bbc.co.uk/news/uk-politics-47192232 [Accessed 12 August 2019].

[42] "The Rise of the Drone Swarm, (2019, February 15), " UK defence journal." [Online]. Available: https://ukdefencejournal.org.uk/the-rise-of-the-drone-swarm/ [Accessed 12 August 2019].

[43] "Raytheon Develops Unmanned Vehicle Swarm Technology, (2018, March 27), " Unmanned Systems Technology. [Online]. Available: https://www.unmannedsystemstechnology.com/2018/03/raytheon-develops-unmanned-vehicle-swarm-technology/ {Accessed 12 August 2019].

[44] "Dedicated Short Range Communications (DSRC) Service, (2019, April 22), " Federal Communications Commission. [Online]. Available: https://www.fcc.gov/wireless/bureau-divisions/mobility-division/dedicated-short-range-communications-dsrc-service [Accessed 12 August 2019].

[45] (2017, November), "Autonomous-driving disruption: Technology, use cases, and opportunities," Mckinsey. [Online]. Available: https://www.mckinsey.com/industries/automotive-and-assembly/our-insights/autonomous-driving-disruption-technology-use-cases-and-opportunities [Accessed 14 August 2019].

[46] "Global status report on road safety 2018, (2018), " World Health Organisation. [Online]. Available:

https://www.who.int/violence_injury_prevention/road_safety_status/2018/en/ [Accessed 14 August 2019].

[47] M. Matousek, (2019, May 2), "Tesla's infotainment system is better than any other auto brand's, according to Consumer Reports," Business Insider. [Online]. Available: https://www.tesla.com/en_GB/BLOG/ALL-TESLA-CARS-BEING-PRODUCED-NOW-HAVE-FULL-SELF-DRIVING-HARDWARE [Accessed 12 August 2019].

[48] "History of CAN technology," CiA. [Online]. Available: https://www.can-cia.org/can-knowledge/can/can-history/ [Accessed 12 August 2019].

[49] "A Brief Introduction to Controller Area Network," Cooperhill technologies. [Online]. Available: https://copperhilltech.com/a-brief-introduction-to-controller-area-network/ [Accessed 12 August 2019].

[50] B. Sheehan, F. Murphy, M Mullins and C. Ryan, (2019, June), "Connected and autonomous vehicles: A cyber-risk classification framework," Science Direct, vol. 124, pp523-536.

[51] (2019, July 7), "Security: New cars that can be stolen in less than 30 seconds," Whatcar. [Online]. Available: https://www.whatcar.com/news/security-new-cars-that-can-be-stolen-in-under-30-seconds/n19898#2 [Accessed 14 August 2019].

[52] D. Tobin, (2014, March 28), "Gone in 90 seconds: How thieves hack into and steal keyless-entry cars," The Sunday Times Driving. [Online]. Available: https://www.driving.co.uk/car-clinic/how-thieves-hack-into-and-steal-keyless-entry-cars/ [Accessed 14 August 2019].

[53] C. Smith, "The car hacker's handbook: a guide for the penetration tester," San Fanccisco, CA: No Starch Press, 2016.

[54] T. Brewster, (2018, July 12), "This GPS Spoofing Hack Can Really Mess With Your Google Maps Trips," Forbes. [Online]. Available: https://www.forbes.com/sites/thomasbrewster/2018/07/12/google-maps-gps-hack-takes-victims-to-ghost-locations/#257b0926335f [Accessed 14 August 2019].

[55] J. Petit and S. Shladover, (2014, September). "Potential Cyberattacks on Autonomated Vehicles." IEEE Transactions on Intelligent Transportation Systems. Available: https://ieeexplore.ieee.org/document/6899663 [Accessed 14 August 2019].

[56] B. O'Hanlon, M. Psiaki, J. Bhatti and D. Shepard. (2013, December). "Real-Time GPS Spoofing Detection via Correlation of Encrypted Signals." Navigation - Journal of The Institute of Navigation. Available: https://www.researchgate.net/publication/259540669_Real-Time_GPS_Spoofing_Detection_via_Correlation_of_Encrypted_Signals [Accessed: 14 August 2019].

[57] T. Humphreys, B. Ledvina, M Psiaki, B. Hanlon and P. Kintner. (2008). "Assessing the Spoofing Threat: Development of a Portable GPS Civilian Spoofer." Available: https://gps.mae.cornell.edu/humphreys_etal_iongnss2008.pdf [Accessed 14 August 2019].

[58] S.-S. Jan and A.-L. Tao, (2016, May), "Comprehensive comparisons of satellite data, signals, and measurements between the BeiDou navigation satellite system and the global positioning system," Sensors, vol. 16, no. 5, p. 689. [Online]. Available: https://www.researchgate.net/publication/303093254_Comprehensive_Comparisons _of_Satellite_Data_Signals_and_Measurements_between_the_BeiDou_Navigation_ Satellite_System_and_the_Global_Positioning_System [Accessed 14 August 2019].

[59] S. Parkinson, P Ward, K Wilson and J Miller. (2017), "Cyber Threats Facing Autonomous and Connected Vehicles: Future Challenges." IEEE transactions on intelligent transport systems, vol 18, no. 11, pp. 2898-2914. [Online]. Available: https://scholar.google.co.uk/scholar?q=Cyber+Threats+Facing+Autonomous+and+C onnected+Vehicles:+Future+Challenges.%22&hl=en&as_sdt=0&as_vis=1&oi=schola rt [Accessed 14 August 2019].

[60] M. Psiaki and T. Humphreys, (2016, April). "GNSS spoofing and detection," Proceesings of the IEEE. [Online]. Available: https://www.researchgate.net/publication/299570659_GNSS_spoofing_and_detectio n [Accessed 14 August].

[61] N. O. Tippenhauer, C. Pöpper, K. B. Rasmussen, and S. Capkun, (2011, October), "On the requirements for successful GPS spoofing attacks," 18th ACM Conference Computing Communication Security, pp. 75–86.

[62] A. M. Wyglinski, X. Huang, T. Padir, L. Lai, T. R. Eisenbarth, and K. Venkatasubramanian, (2013, January), "Security of autonomous systems employing embedded computing and sensors," IEEE Micro, vol. 33, no. 1, pp. 80–86. [Online].

[63] B. W. O'Hanlon, M. L. Psiaki, J. A. Bhatti, D. P. Shepard, and T. E. Humphreys, "Real-time GPS spoofing detection via correlation of encrypted signals," Navigation, vol. 60, no. 4, pp. 267–278, 2013.

[64] E. J. Ohlmeyer, "Analysis of an ultra-tightly coupled GPS/INS system in jamming," in Proc. IEEE/ION Position, Location, Navigat. Symp., Apr. 2006, pp. 44– 53.

[65] (2011, December 19), "Iran claims to have hacked US drone, then landed it," Infosec Magasine. [Online]. Available: https://www.infosecurity-magazine.com/news/iran-claims-to-have-hacked-us-drone-then-landed-it/ [Accessed 14 August 2019].

[66] K. Boccuzzi, (2008), "Investigating the causes of and possible remedies for sensor damage in digital cameras used on the omega laser systems," Univ. Rochester, Rochester, NY, USA, Tech Rep..

[67] D. Cohen, (2011, December 17), "Iran 'blinded' CIA spy satellite," Ynetnews. [Online]. Available: https://www.ynetnews.com/articles/0,7340,L-4162770,00.html [Accessed 14 August 2019].

[68] "A Tragic Loss," Tesla, (2016, 30 June), [Online]. Available: https://www.tesla.com/en_GB/blog/tragic-loss?redirect=no [Accessed 14 August 2019].

[69] B. G. Stottelaar. (2015, February), Practical Cyber-Attacks on Autonomous Vehicles. [Online]. Available: http://essay.utwente.nl/66766/ [Accessed 14 August 2019].

[70] M. Naimark, (2002), "How to ZAP a camera: Using lasers to temporarily neutralize camera sensors," [Online]. Available: http://www.naimark.net/projects/zap/howto.html [Accessed 14 August 2019].

[71] K. N. Truong, S. N. Patel, J. W. Summet, and G. D. Abowd, (2005), "Preventing camera recording by designing a capture-resistant environment," in Proc. 7th Int. Conf. UbiComp, pp. 73–86.

[72] (2016, March 29), "Electronic Countermeasure (ECM)," Aircraft 101. [Online]. Available: https://basicsaboutaerodynamicsandavionics.wordpress.com/2016/03/29/electronic-countermeasure-ecm/ [Accessed 14 August 2019].

[73] J. Petit and S. Shladover. (2015, February), "Potential Cyberattacks on Autonomous Vehicles." IEEE Transactions on Intelligent Transportation Systems. Vol. 16, Issue 1, pp. 546-556. Available: https://ieeexplore.ieee.org/document/6910285/authors#authors [Accessed 14 August 2019].

[74] C. Zhou, E. Liu and Q. Liu (2017, October), "An Adaptive Transmitting Scheme for Interrupted Sampling Repeater Jamming Suppression." MDPI. Available: https://pdfs.semanticscholar.org/39a5/11e000569538c5ac57552d2963119beb9342.pdf [Accessed 14 August 2019].

[75] M. Harris, (2015, September 4), "Researcher Hacks Self-Driving Car Sensors," IEEE Spectrum. [Online]. Available: http://spectrum.ieee.org/cars-that-think/transportation/self-driving/researcher-hacks-selfdriving-car-sensors [Accessed 14 August 2019].

[76] K. Koscher, A. Czeskis, F. Roesner, S. Patel, T. Kohno, S. Checkoway and D. McCoy. (2010, May), "Experimental security analysis of a modern automobile," IEEE Symposium on Security and Privacy, pp. 447–462. [Online]. Available: https://ieeexplore.ieee.org/abstract/document/5504804/authors#authors [Accessed 14 August 2019].

[77] I. Rouf, R. Miller, H. Mustafa, T. Taylor, S. Oh, W. Xu, M. Gruteser, W. Trappe and I. Seskar, (2010, February), "Security and privacy vulnerabilities of in-car

wireless networks: A tire pressure monitoring system case study." [Online]. Available: https://www.usenix.org/legacy/event/sec10/tech/full_papers/Rouf.pdf [Accessed 14 August 2014].

[78] N. Sarter and D. Woods, (1995, March), "How in the world did we ever get into that mode? Mode error and awareness in supervisory control," Human Factors the Journal of the Human Factors and Ergonomics Society, vol. 37, no. 1, pp. 5–19. [Online]. Available: https://www.researchgate.net/publication/258138719_How_in_the_World_Did_We_Ever_Get_into_That_Mode_Mode_Error_and_Awareness_in_Supervisory_Control [Accessed 14 August 2019].

[79] S. Singh, K. Kingsley and C. Chen, (2009, April), "Tire pressure maintenance— A statistical investigation," U.S Dept. Transportation. [Online]. Available: https://crashstats.nhtsa.dot.gov/Api/Public/ViewPublication/811086 [Accessed 14 August 2019].

[80] A. Versprille, (2016, August 16), "Researchers Hack Into Driverless Car System, Take Control Of Vehicle," National Defence. [Online]. Available: https://www.nationaldefensemagazine.org/articles/2015/5/1/2015may-researchers-hack-into-driverless-car-system-take-control-of-vehicle [Accessed 14 August 2019].

[81] T. Zhang, H. Antunes and S. Aggarwal, (2014, February), "Defending connected vehicles against malware: Challenges and a solution framework," IEEE Internet Things Journal, vol. 1, no. 1, pp. 10–21. [Online]. Available: https://www.researchgate.net/publication/264582907_Defending_Connected_Vehicles_Against_Malware_Challenges_and_a_Solution_Framework [Accessed 14 August 2019].

[82] (2019, February 21), "Iran claims it hacked and controlled US drone, shows footage from missions as proof (video)," Russia Today. [Online]. Available: https://www.rt.com/news/452116-iran-claims-control-us-drones/ [Accessed 14 August 2019].

[83] J. Isaac, S. Zeadally and J. Camara, (2010, April), "Security attacks and solutions for vehicular ad hoc networks," IET Commununications, vol. 4, no. 7, pp. 894–903. [Online]. Available: https://digital-library.theiet.org/content/journals/10.1049/iet-com.2009.0191 [Accessed 14 August 2019].

[84] I. Sumra, I. Ahmad, H. Hasbullah and J. Manan, (2011, April), "Classes of attacks in VANET," IEEE Saudi International in Electronics, Communication and Photonics Conference (SIECPC), vol. 1753, pp. 1–5. [Online]. Available: https://ieeexplore.ieee.org/document/5876939 [Accessed 14 August 2019].

[85] S. Checkoway, D. McCoy, B. Kantor, A. Anderson, H. Shacham and S. Savage, (2011, August), "Comprehensive experimental analyses of automotive attack surfaces," USENIX Security Symposium, pp. 77–92. [Online]. Available: http://www.autosec.org/pubs/cars-usenixsec2011.pdf [Accessed 14 August 2019].

[86] R. Ferris. (2016, September 20), "Chinese company hacks Tesla car remotely," CNBC. [Online]. Available: https://www.cnbc.com/2016/09/20/chinese-company-hacks-tesla-car-remotely.html [Accessed 14 August 2019].

[87] A. Greenberg, (2015, July 7), "Hackers remotely kill a jeep on the highway - with me in it," Wired. [Online]. Available: https://www.wired.com/2015/07/hackers-remotely-kill-jeep-highway/ [Accessed 14 December 2019].

[88] K. Haataja, (2009), "Security Threats Countermeasures Bluetooth-Enabled System." Saarbrücken, Germany: LAP Lambert Academic Publishing. [Online]. Available: http://epublications.uef.fi/pub/urn_isbn_978-951-27-0111-7/urn_isbn_978-951-27-0111-7.pdf [Accessed 14 August 2019].

[89] J. Lindberg, (2011, January 1), "Security analysis of vehicle diagnostics using DoIP." [Online] Available: https://scinapse.io/papers/116323999 [Accessed 14 August 2019].

[90] D. Spill and A. Bittau, (2007), "Bluesniff: Eve meets Alice and Bluetooth." [Online]. Available: http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.565.6674&rep=rep1&type=pdf [Accessed 14 August 2019].

[91] M. Jenkins and S. M. Mahmud, (2006, April), "Security needs for the future intelligent vehicles," SAE. [Online]. Available: https://www.researchgate.net/publication/228866911_Security_Needs_for_the_Future_Intelligent_Vehicles [Accessed 14 August 2019].

[92] C. Miller and C. Valasek, (2014, August 20), "A Survey of Remote Automotive Attack Surfaces." [Online]. Available: https://www.slideshare.net/LudovicP/miller-valaceksurveyofremoteattacksurfaces [Accessed 19 August 2019].

[93] C. Miller and C. Valasek, (2015, August 10), "Remote Exploitation of an Unaltered Passenger Vehicle." [Online]. Available: http://illmatics.com/Remote%20Car%20Hacking.pdf [Accessed 14 August 2019].

[94] (2015, July 24), "Fiat Chrysler recalls 1.4m vehicles in wake of Jeep hacking revelation," The Guardian. [Online]. Available: https://www.theguardian.com/business/2015/jul/24/fiat-chrysler-recall-jeep-hacking [Accessed 14 August 2019].

[95] S. Checkoway, D. McCoy, B. Kantor, D. Anderson, H. Shacham, S. Savage, K. Koscher, A. Czeskis, F. Roesner and T. Kohno, (2011), "Comprehensive Experimental Analyses of Automotive Attack Surfaces," [Online]. Available: http://www.autosec.org/pubs/cars-usenixsec2011.pdf [Accessed 14 August 2014].

[96] K. Zetter, (2015, June 8), "Researchers hacked a Model S, but Tesla's already released a patch," Wired. [Online]. Available: https://www.wired.com/2015/08/researchers-hacked-model-s-teslas-already/ [Accessed 14 August 2019].

[97] P. Mundhenk, S. Steinhorst, M. Lukasiewycz, S. Fahmy, and S. Chakraborty, (2015), "Lightweight authentication for secure automotive networks.' [Online]. Available: https://warwick.ac.uk/fac/sci/eng/staff/saf/publications/date2015-mundhenk.pdf [Accessed 14 August 2019].

[98] S. Gao, X. Dai, Y. Hang, Y. Guo and Q Ji, (2018), "Airborne Wireless Sensor Network for Airplane Monitoring System," Wireless Communications and Mobile Computing, vol. 2018. [Online]. Available: https://www.hindawi.com/journals/wcmc/2018/6025825/ [Accessed 14 August 2019].

[99] L. Gomes, (2014, August 28), "Hidden obstacles for googles self-driving cars: Impressive progress hides major limitations of googles quest for automated driving." MIT Technological Review. [Online]. Available: https://www.technologyreview.com/s/530276/hidden-obstacles-for-googles-self-driving-cars/ [Accessed 14 August 2019].

[100] R. Nicolas, (2013, March 4), "Introduction to On Board Diagnostics (OBD)," Car Engineer. [Online]. Available: http://www.car-engineer.com/introduction-to-on-board-diagnostic-obd/ [Accessed 14 August 2019].

[101] A. Yadav, G. Bose, R. Bhange, K. Kapoor, N. C. S. Iyengar, and R. D. Caytiles, (2016), "Security, vulnerability and protection of vehicular on-board diagnostics," Int. J. Secur. Appl., vol. 10, no. 4, pp. 405–422.

[102] D. Kiliedinst and C. King, (2016), "On Board Diagnostics: Risk and Vulnerabilities of the Connected Vehicle", CERT Coordination Center, Tech. Rep..

[103] Q. Wu, J. Domingo-Ferrer, and Ú. Gonzalez-Nicolas, (2010, February), "Balanced trustworthiness, safety, and privacy in vehicle-to-vehicle communications," IEEE Transactions on Vehicular Technology, vol. 59, no. 2, pp. 559–573.

[104] (2018, May 25), "Guide to the General Data Protection Regulations (GDPR)," Information Commissioners Office.

[105] L. Kellon, (2016, February 24), "Nissan Leaf Electric Cars Hack Vulnerability Disclosed" BBC. [Online]. Available: https://www.bbc.co.uk/news/technology-35642749 [Accessed 14 August 2019].

[106] (2016, January), "Surface vehicle recommended practice: Cybersecurity Guidebook for Cyber-Physical Vehicle Systems," SAE International.

[107] A. Johnson, (2009, August 9), "Revealed: How Army's new armoured vehicle is a death trap too," Independent. [Online]. Available: https://www.independent.co.uk/news/uk/home-news/revealed-how-armys-new-armoured-vehicle-is-a-death-trap-too-1769692.html [Accessed 14 August 2019].

[108] "SIT - the Self-Inflating Tire for autonomous driving future." [Online]. Available: http://www.selfinflatingtire.com [Accessed 14 August 2019].

[109] R. Akram, (2019, July 22), Project supervision discussion.

[110] P. Scharre, (2019, January 19), "Autonomous weapons and the new laws of war," The Economist.

[111] A. Withnall, (2019, May 10), "Chelsea Manning: Jailed US analyst walks free after refusing to testify to WikiLeaks grand jury," Independent. [Online]. Available: https://www.independent.co.uk/news/world/americas/chelsea-manning-wikileaks-prison-army-us-grand-jury-jail-sentence-a8907401.html [Accessed 14 August 2019].

[112] S. Marsden and R. Massey, (2015, March 2015), "Why can't airlines seize control of doomed jets from the ground? The technology exists but pilots and companies refuse to use it," The Daily Mail. [Online]. Available: https://www.dailymail.co.uk/news/article-3013858/Why-t-airlines-seize-control-doomed-jets-ground-technology-exists-pilots-companies-refuse-use-it.html [Accessed 14 August 2019].

[113] The Ponemon Institute, (2007), "Car Cybersecurity: A gap still exists,'. [Online]. Available: https://www.slideshare.net/OnBoardSecurity/car-cybersecurity-what-do-automakers-really-think-77268227 [Accessed 14 August 2019].

[114] (2019, April 5), "Boing 737 Max: What went wrong?" BBC. [Online]. Available: https://www.bbc.co.uk/news/world-africa-47553174 [Accessed 14 August 2019].

[115] J. Keller, (2019, July 16), "U.S. military shifting research and technology development toward armed robotic ground vehicles," Military & Aerospace Electronics. [Online]. Available: https://www.militaryaerospace.com/unmanned/article/14036321/ground-vehicles-robotic-armed [Accessed 14 August 2019].

[116] (2019, June 29), "Defence Secretary announces major cyber investment," Ministry of Defence. [Online]. Available: https://www.contracts.mod.uk/blog/defence-secretary-announces-major-cyber-investment/ [Accessed 14 August 2019].