Investigating the security vulnerabilities and solutions for connected and autonomous vehicle technologies

Authors
William Booth, MSc (Royal Holloway, 2021)
Siaw-Lynn Ng, ISG, Royal Holloway

Abstract

Within the last decade alone, the technological developments in the automotive industry have brought forward the emergence of connected and autonomous vehicles (CAVs). Such technologies, however, are subject to both emerging and traditional cybersecurity threats. This article introduces CAVs and their emergence, analysing the underlying technologies and considering their cyber security vulnerabilities and attacks. We identify and assess the existing and emerging countermeasures for such vulnerabilities, and propose high-level recommendations for the vehicle and communication technologies, and the automotive industry, as a whole.[a]

[a] This article is published online by Computer Weekly as part of the 2022 Royal Holloway information security thesis series https://www.computerweekly.com/ehandbook/Royal-Holloway-Securing-connected-and-autonomous-vehicles. It is based on an MSc dissertation written as part of the MSc in Information Security at the ISG, Royal Holloway, University of London. The full thesis is published on the ISG's website at https://www.royalholloway.ac.uk/research-and-teaching/departments-and-schools/information-security/research/explore-our-research/isg-technical-reports/.

Introduction

The development of cheaper and more powerful chipsets, sensors and software components is enabling the production of connected and autonomous vehicles (CAVs) that proposes a future of automotive transport with minimal human input. CAVs are intriguing consumers, industries and governments across global economies, all heavily invested in the potential capabilities. The potential benefits CAVs present to the environment, the economy and public safety are not to be underestimated. The United Kingdom's Centre for Connected and Autonomous Vehicles (CCAV) estimates the market for CAVs to be worth between £52 billion and £62 billion by 2035, capturing around 6% of the £907 billion global market.

However, in order to realise a future of fully connected and autonomous vehicles, the industry must overcome several challenges. Improvements in technological security, regulatory requirements, personal privacy safeguards, industry standardisation and consumer trust are all required to overcome the challenges CAVs face. Additionally, the CAV space has major cybersecurity considerations that must be addressed. Alongside the traditional safety vulnerabilities that concern modern vehicles, CAVs present a vast attack surface for remote attacks on autonomous vehicle hardware, software, user privacy, security and more. While prior research has identified attack on CAV technologies, and applicable legislation has been discussed, there remains a gap in literature for exploring CAV security as a whole.

Connected and autonomous vehicles

CAVs can best be understood as the technologies and capabilities that are inherited from the vehicles being both *connected* and *autonomous*.

- **Autonomous Vehicle (AV)**

  A vehicle which is capable of fulfilling the operational functions of a tradition vehicle, such as the safe and lawful manoeuvring of roads, without human intervention or a back-end control centre. This can be achieved by using a combination of on-board sensors and actuator networks that gathers information on the surrounding environment. Autonomous vehicular decision making is to be supported by computer vision and machine learning capabilities.

- **Connected Vehicle (CV)**

  A vehicle which has the technology that enables it to connect to devices within the vehicle, as well as external networks such as the internet, allowing it to communicate with its surrounding infrastructure and other vehicles. Internally, the communication devices can be connected using a combination of wired or wireless communication technologies, where as externally they are connected using wireless communications.

SAE levels of driving automation

The SAE defines six levels of driving automation, ranging from no driving automation (level 0) to full driving automation (level 5), providing a classification for vehicle driving automation systems that perform part or all of the dynamic driving task (DDT).

SAE Level	Description of Driving Automation	Example Features
Level 0 ***No Driving Automation***	The human driver performs all aspects of the entire Dynamic Driving Task (DDT), even when enhanced by active safety systems.	Automatic emergency braking, Blind spot monitoring, Lane departure warning
Level 1 ***Driver Assistance***	The driver support features perform either the lateral or longitudinal vehicle motion control subtasks of the DDT, i.e. either steering or acceleration/deceleration.	Lane Centering, OR, Adaptive cruise control
Level 2 ***Partial Driving Automation***	The driving automation systems can perform sustained lateral and longitudinal vehicle motion control subtasks of the DDT. The driver is responsible for continually supervising the driving automation system and should remain engaged throughout.	Lane Centering, AND, Adaptive cruise control
Level 3 ***Conditional Driving Automation***	The sustained performance by an autonomous driving system of the entire DDT with the expectation that a user is receptive to DDT fallback requests to intervene from the vehicle.	Traffic jam chauffeur
Level 4 ***High Driving Automation***	The sustained performance by an autonomous driving system of the entire DDT and fallback without any expectation that a user will need to intervene.	Local driverless taxi, Pedals/steering wheel may or may not be installed
Level 5 ***Full Driving Automation***	The sustained and unconditional performance by an autonomous driving system of the entire DDT and DDT fallback without any expectation that a user will need to intervene.	Same as level 4, but features can drive everywhere in all conditions.

Table 1: SAE's levels of driving automation

Exteroceptive and proprioceptive sensing

There are two main types of sensors used within an connected and autonomous vehicles: Exteroceptive sensors, used for sensing and perceiving the surrounding environment, and, proprioceptive sensors, which are used to measure vehicle dynamics. We describe a few here.

LiDAR: Light Detection and Ranging (LiDAR) is a remote sensing technology used to measure distance by processing the time delays (time-of-flight) for emitted optical pulses to be reflected back from an object. It is a popular sensor choice for autonomous vehicles as they are able to generate a detailed three-dimensional view of the surrounding environment. However, the technology is not able to differentiate between objects, meaning a stray plastic bag could be interpreted as a road bump for example, and its efficacy is affected by adverse weather conditions due to the absorption and scattering of light.

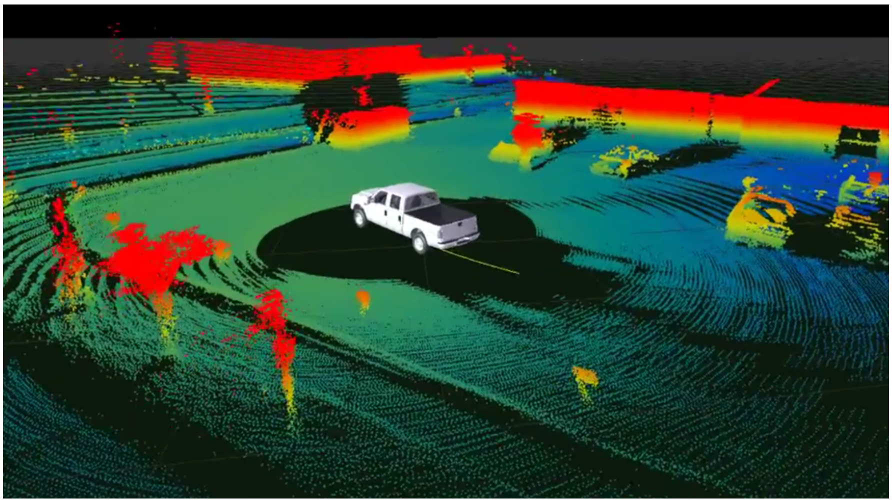

Figure 1: Three-dimensional view of a 360° LiDAR image

Camera: Digital cameras are the most accurate way to create a visual representation of the environment. CAVs incorporate high resolution cameras on each side of the vehicle, producing a three-dimensional view of the surrounding environment. Cameras are often set up so that they overlap, providing depth measuring capabilities in close proximity contexts. Cameras can also distinguish colour, allowing the vehicle to recognise elements in the environment such as traffic lights, road signs, vehicle lights etc. With this, AI systems on the vehicle can identify pedestrians. Although cameras provide considerable performance relative to their cost, image quality diminishes in low light and extreme weather conditions.

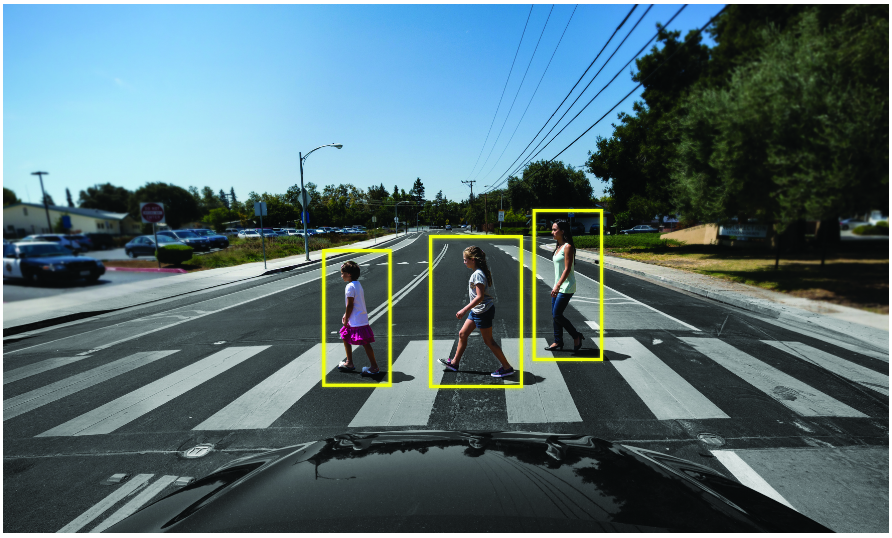

Figure 2: Autonomous vehicle camera using AI to identify pedestrians

GPS: Global Positioning Systems (GPS) provide satellite-based radio-navigation producing longitudinal and latitudinal coordinates. GPS systems are highly accurate and relatively inexpensive. However, radio signals used in GPS do not penetrate buildings, meaning that they are less effective in built-up urban environments. GPS systems are also vulnerable to signal interference, meaning a hacker could deploy spoofing or jamming attacks to the incoming and outgoing signals.

AI: CAV systems rely heavily on machine and deep learning to process the data received form the vehicle's sensors. Machine learning systems train, validate and improve the autonomous driving systems. AI can then be used for autonomous driving applications such as object recognition, vehicle localisation and object tracking. However, AI presents security vulnerabilities that can be exploited in order to disrupt and manipulate the operation of autonomous driving systems.

V2X Communication

The data gathered from each of the aforementioned sensing technologies is not only used by the individual vehicle to support autonomous driving, but is also shared between various vehicles, infrastructure and other nodes on a live vehicular network. This is referred to as vehicle-to-everything communication, or V2X. V2X is an umbrella term for the subset of communication systems of which a CAV possesses. These consist of vehicle to vehicle-to-vehicle (V2V), vehicle-to-infrastructure (V2I), vehicle-to-network (V2N) and vehicle-to-pedestrian (V2P) communications.

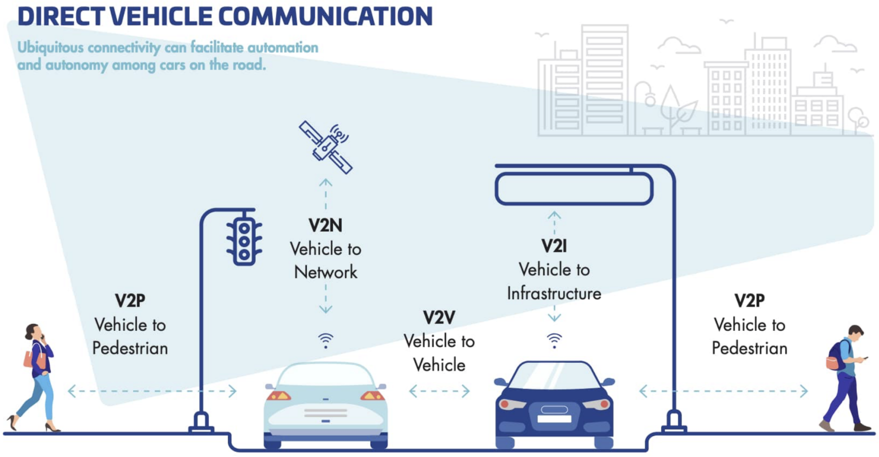
DIRECT VEHICLE COMMUNICATION
Ubiquitous connectivity can facilitate automation and autonomy among cars on the road.
V2N
Vehicle to Network
V2I
Vehicle to Infrastructure
V2P
Vehicle to Pedestrian
V2V
Vehicle to Vehicle
V2P
Vehicle to Pedestrian

Figure 3: The Vehicle-to-Everything (V2X) communication platform

Information collected by the on-board sensors is communicated externally, either to other vehicles, the surrounding infrastructure, pedestrian smartphones, and data centres. V2X communication informs the vehicle with current and predicted environmental conditions, traffic dynamics, and road closures, for example. Each type of vehicular communication is used simultaneously in order to provide reliable and safe mobility. The various communication systems each carry differing requirements as to how the data is transmitted, balancing efficiency, performance and cost.

Attacks and countermeasures

LiDAR & radar:

Attacks: With LiDAR and radar there is no guarantee over the validity of the constructed 3D model. These imaging systems cannot differentiate between objects, and often misinterpret objects in motion. This poses a great security vulnerability, as it allows attackers to manipulate sensor readings through either spoofing, relaying or jamming.

Spoofing attacks use the same physical channels as the sensor unit to emit counterfeit signals, so that what the sensor interprets may not be the same as what the object is in reality. Such attacks do not require complex or expensive hardware. In a relay attack, the attacker captures and delays the original signal from the sensor unit and then relays the signal back to the sensor. This changes the position of an object. For example, an attacker can relay the signal received form the driver's side of the car, and emit it onto the passenger's side. In a jamming attack. the sensor is overwhelmed by signal pulses of the same wavelength and timing as the sensor unit, so there is no usable data for the autonomous driving system. Attacks of this nature do not require expensive equipment, and can be made portable, allowing for remote attacks to be carried out.

Countermeasures: Using different wavelengths can reduce the success of spoofing and jamming attacks. Random probing - periodically and randomly changing the interval between scanning speeds - also makes it harder for the attack to synchronise to the original wavelength speed.

Constantly alternating the radar frequency could limit the effectiveness of jamming attempts, as the attacker will not be able to determine and lock onto a single frequency. However, this may only provide limited success, as, in order for radar to work at each distance, there is a narrow operating window.

Sensor fail safe principles can be applied to radar. Utilising the on board AI and machine learning capabilities, as well as cross-referencing the data with other sensors such as LiDAR, the vehicle could perform anomaly detection on scanned objects to determine if an object has been falsely introduced or its position altered.

Camera:

Attacks: Cameras fitted to AVs typically use either a Charge-Coupled Device (CCD) or Complementary Metal Oxide Semiconductor (CMOS) sensors, which are vulnerable to partial or total blinding, including permanent damage in extreme cases. Such an attack can be achieved by pointing a low cost 'interferer', such as LED spot lasers, directly at the camera.

Automotive cameras are also susceptible to attacks by concealing traffic signs. Research show that it is possible to alter the information and 'hide' traffic signs by surrounding and masking them with other shapes and colours, confusing the AI models. An attacker can further abuse this by placing fake traffic signs in unsuitable locations, or paint additional lane markings on the road, for example.

Camera object tracking capabilities can also be overwhelmed by presenting too many objects to track. Furthermore, the automatic expose controls and auto-focus of the camera can be attacked by pointing a bright light at the camera. When a light is introduced, the camera will reduce its sensitivity and exposure to try and draw out the remaining information from the available image, however, this too can be easily overwhelmed. Thus, a hacker can abuse this by hiding vital information such as traffic signs or pedestrians by introducing light.

Countermeasures: Blinding attacks can be mitigated if the vehicle has a secondary reserve camera. Furthermore, additional cameras could be implemented at different strategic locations on the vehicle, making it difficult for an attacker to blind every camera. Lens filters can also be used to filter out interference. These mitigation methods do introduce additional costs and must be used strategically.

Sensor fail safe principles can be applied to counter both blinding attacks and targeted attacks on camera functionality. The camera's software should have maximum exposure limits, which can shut off the camera unit if a light source causes the exposure to increase to an abnormal level. V2V communication further allows for anomaly detection, as a targeted vehicle will show dramatically different camera data.

GPS

Attacks: Civilian GPS systems used in CAVs are designed without encrypted and authorised transmission and are not intended for safety or security critical operations. Although adopting an open standard for GPS within CAVs may well be robust and inexpensive, the accessible and predictable architecture makes the technology vulnerable to counterfeited or spoofed signals.

Another advanced attack on GPS systems is a black hole attack, where an attacker causes the deliberate loss of GPS information across a V2X network. The attacker can then falsify their GPS data and advertise themselves as having the correct GPS data.

Countermeasures: Navigation Message Authentication (NMA) could be implemented. This method embeds public-key digital signatures into the GPS message which can be validated by the vehicle. This allows for only authorised signals to be accepted by the GPS system on the vehicle. An additional spoofing countermeasure looks to utilise multiple GPS receivers deployed in a static, known formation on the vehicle. This method allows the receivers to exchange their individual locations on the vehicle and can each check if their calculated locations preserve their physical formation.

Attacks on V2X

Whilst attacks on isolated traditional vehicles can have considerable consequences, security breaches do not compromise or propagate over to other vehicles, connected infrastructure or external networks. V2X technologies, however, are susceptible to many more attacks, including the following:

1. *Blackhole attacks:* The attacker receives transmissions from the V2X network but does not route the received data, blocking the spread of information across the network.
2. *Bogus messages:* Compromised nodes spread bogus transmissions across the network, either by generating false messages or modifying existing ones, to misguide other vehicles on the network.
3. *Certificate replication:* Attackers gain access to compromised nodes and exploit replicated certificates to conceal themselves. Certificates that were previously added to a blacklist are recycled by malicious entities on the network
4. *Denial-of-Service:* An attacker injects copious volumes on data into the network, attempting to minimise the packet reception ration (PRR) for nodes on the network, jeopardising the availability and performance of V2X.
5. *Eavesdropping:* Attackers ‘listen in’ to data flowing through a V2X network, aiming to acquire sensitive and confidential information on users, infrastructure and network information.
6. *Masquerading:* The attacker presents himself as a legitimate entity on the V2X network, gaining access to confidential data and abusing authorisation controls. Attackers can then send malicious information to nodes on the network, causing a number of other attacks.
7. *Location tracking:* An attacker listens over a network and analyses the data collected from neighbouring nodes to identify the current and previous locations of the target.
8. *Message modification:* Attackers who have received legitimate messages from nodes on the network modify the message before sending it on.
9. *Replay attacks:* Replay attacks are considered one of the most common attacks in all types of networks. Messages received by adversaries are maliciously replayed repeatedly over the networks, potentially inducing Denial-of-Service across the network.
10. *Sybil attack:* An attacker joins a network using multiple real or fake identities, to generate falsified vehicles on the road, benefiting the attacker.
11. *Unauthorised access:* Network services are accessed by unauthorised users. Confidential data is often targeted as the access controls protecting the given data has been overridden.

Some recommendations

Systematic security validation for AI: The large volume of data that CAVs capture and process provide the foundations for the AI models which enables autonomous driving. However, the models are constantly changing, which can present potential security threats, as model updates can add vulnerabilities that can be exploited.

Therefore, system designers should ensure that the security of the model updates are systematically assessed and validated. so that security vulnerabilities can be identified and rectified before the AI system can be exploited. Furthermore, risk assessments and incident response procedures should be regularly carried out, so that in the event of an attack, the vehicle or system designers can quickly react and neutralise the threat before the security and safety of the vehicle is compromised.

Control access to hardware, firmware and networks: System designers should consider employing strict access control on critical system resources such as firmware, hardware or data. This ensures that only authorised and authenticated personnel can have access to the vehicle's systems. Access control and event logging can provide accountability for actions performed on the vehicle's systems and a certain level of non-repudiation.

System designers should also consider adopting least privilege access control protocols. This ensures that every user of the system should operate using the least set of privileges necessary in order to complete their task. This approach limits the damage that can result from an error, and further limits the interactions between critical CAV systems to the absolute minimum while ensuring autonomous operation and security remains intact.

Network Segmentation: A straightforward solution for protecting in-vehicle networks such as the CAN bus, is to separate the network into multiple sub-networks. Segmentation provides control over which entities or users can access the particular sub-networks, thus reducing the resulting damage from an attack.

Furthermore, network segmentation ensures that errors or attacks do not propagate onto other networks, as an attack on a node in the network will only spread to the specific sub-network in which it is located. This can be used to help protect safety critical systems, with attacks on Bluetooth not propagating to LiDAR sensors for example.

Conclusion

So far there is little in the way of a practical cyber attack history on CAV technologies. However, this is by no means a testament to the security of the vehicle's technologies or systems. Autonomous technologies in automotive vehicles are in their infancy, and while a shift towards more advanced technologies may well provide a multitude of socio-economic benefits, it will also bring forth previously unseen cybersecurity threats and vulnerabilities. Although the automotive industry is capable of dealing with traditional security issues, such as car theft, it will need to adopt a new stance on cybersecuirty and should be aware of the cyber threats it faces.

Biographies

William Booth has previously graduated with a BSc in Computing Science from the University of East Anglia in 2020, and MSc in Information Security with Distinction at Royal Holloway, University of London in 2021. Currently, he is the Information Security Lead for Darktrace, the world leaders in Autonomous Cyber AI, where he is responsible for the management of the vendor information security assurance program, and the oversight of Information Security in project management.

Siaw-Lynn Ng is a Senior Lecturer in the ISG. Her research interests includes combinatorics and finite geometry and their applications in information security.

Series editor: S.- L. Ng