Master of Science in Transportation Systems

Ingenieurfakultät Bau Geo Umwelt

Technische Universität München

Quantitative Cybersecurity Risk Management for Autonomous Vehicle Systems

David Bailey

October 6, 2018 (with comments incorporated on October 25, 2018)

> *Those with access to these resources – students, librarians, scientists – you have been given a privilege. You get to feed at this banquet of knowledge while the rest of the world is locked out. But you need not – indeed, morally, you cannot – keep this privilege for yourselves. You have a duty to share it with the world. And you have: trading passwords with colleagues, filling download requests for friends.*
>
> *Meanwhile, those who have been locked out are not standing idly by. You have been sneaking through holes and climbing over fences, liberating the information locked up by the publishers and sharing them with your friends.*
>
> *But all of this action goes on in the dark, hidden underground. It's called stealing or piracy, as if sharing a wealth of knowledge were the moral equivalent of plundering a ship and murdering its crew. But sharing isn't immoral – it's a moral imperative. Only those blinded by greed would refuse to let a friend make a copy.* – Aaron Swartz (2008)

Contents

List of Figures

List of Tables

Declaration of Authorship

I declare that this Master's thesis is my own work and I have documented all sources and materials used. This thesis has not been previously presented to another examination board and has not been published.

Munich, October 6, 2018

Place, date

Signature

Acknowledgements

> *The most important thing I’ve accomplished, other than building the compiler, is training young people. They come to me, you know, and say, 'Do you think we can do this?' I say, "Try it." And I back 'em up. They need that. I keep track of them as they get older and I stir 'em up at intervals so they don’t forget to take chances.* – Grace Hopper (Lynn and Moore, 2012)

I would like to thank Manos Chaniotakis and Costas Antoniou from the Technical University of Munich for their guidance and support during this thesis. I would also like to thank Stefan, Kartik, Dave, Dan, and everyone else from Starsky Robotics for allowing me to complete this thesis with them.

Abstract

Today, autonomous vehicles are driving themselves in cities from Munich to San Francisco. The failure of the security controls of these vehicles could cause a danger to public safety and shake public confidence in the technology. The ability to assess the security of autonomous vehicle systems and provide assurances that they are safe to operate is critical for their implementation to be successful and accepted. This thesis explores the potential for risk management frameworks to quantitatively measure the security of autonomous vehicle systems and select the appropriate security controls for a system.

The introduction begins with a short history of autonomous vehicles and a high-level overview of modern autonomous vehicles. Next is an overview of successful attacks against vehicles and an introduction to information security, risk management, and risk assessments. The methods section describes an adaptation of the National Institute of Standards and Technology (NIST) risk assessment process using quantitative methods instead of qualitative methods and an optimization algorithm for prioritizing security controls for an autonomous vehicle system. This framework provides the ability to perform both discrete calculations and Monte Carlo simulations. The results and discussion sections demonstrate the ability of quantitative risk management frameworks to model the risks present in an autonomous vehicle system and the cost-effectiveness of various security controls. As others have previously noted, we find that quantitative risk assessment methods can provide more value than qualitative methods, but often require more data (National Institute of Standards and Technology, 2012). While this thesis focuses on security risk management for autonomous vehicle systems, the same techniques can be used in supporting decisions in other areas of transportation including safety, economic/environmental modeling, and traffic management.

1 Introduction

The autonomous vehicle systems driving today are research projects with human safety drivers providing a safeguard against cybersecurity attacks. However, transportation experts predict that cars, trucks, and buses will soon be driving themselves around the world (Littman, 2018; Kockelman et al., 2016). The failure of the security controls of these vehicles could cause a danger to public safety and shake public confidence in the technology. Of people who state they never plan to buy an autonomous car, 30% list ”concerned about the risk of hacking” as the most important reason for not purchasing an autonomous automobile (Ponemon, 2017). However, companies typically spend millions of dollars per year on cybersecurity and still fall victim to breaches (Richards et al., 2017). The ability to assess the security of autonomous vehicle systems and provide assurances that they are safe to operate is critical for their implementation to be successful.

1.1 Benz and Stanley

> *Have you ever noticed, when you're driving, that anyone who's driving slower than you is an idiot, and anyone driving faster than you is a maniac?* – George Carlin (1984)

The first motor vehicles, engineered by Karl Benz in 1885, (Figure 1) were mechanical devices, controlled by a driver (Benz & Co. in Mannheim, 1886). Over the next one-hundred years, electromechanical systems replaced mechanical systems and drivers now control computers which control vehicles. In 2004, the US Department of Defense held the first DARPA Grand Challenge to accelerate the advancement of autonomous vehicle technologies. The 2004 DARPA Grand Challenge was a race for vehicles to autonomously traverse a 150-mile course in the California desert. After no vehicles completed the course, the Defense Advanced Research Projects Agency (DARPA) raised the prize money from $1 million to $2 million for the 2005 DARPA Grand Challenge (Defense Advanced Research Projects Agency, 2004). Five vehicles completed the second race, led by Stanford Racing Team's Stanley (Figure 2) (Davis, 2006; Thrun et al., 2007).

These two races started a revolution in autonomous vehicles that are completely controlled by computers, not drivers. This drive towards autonomous vehicles, along with complemen-

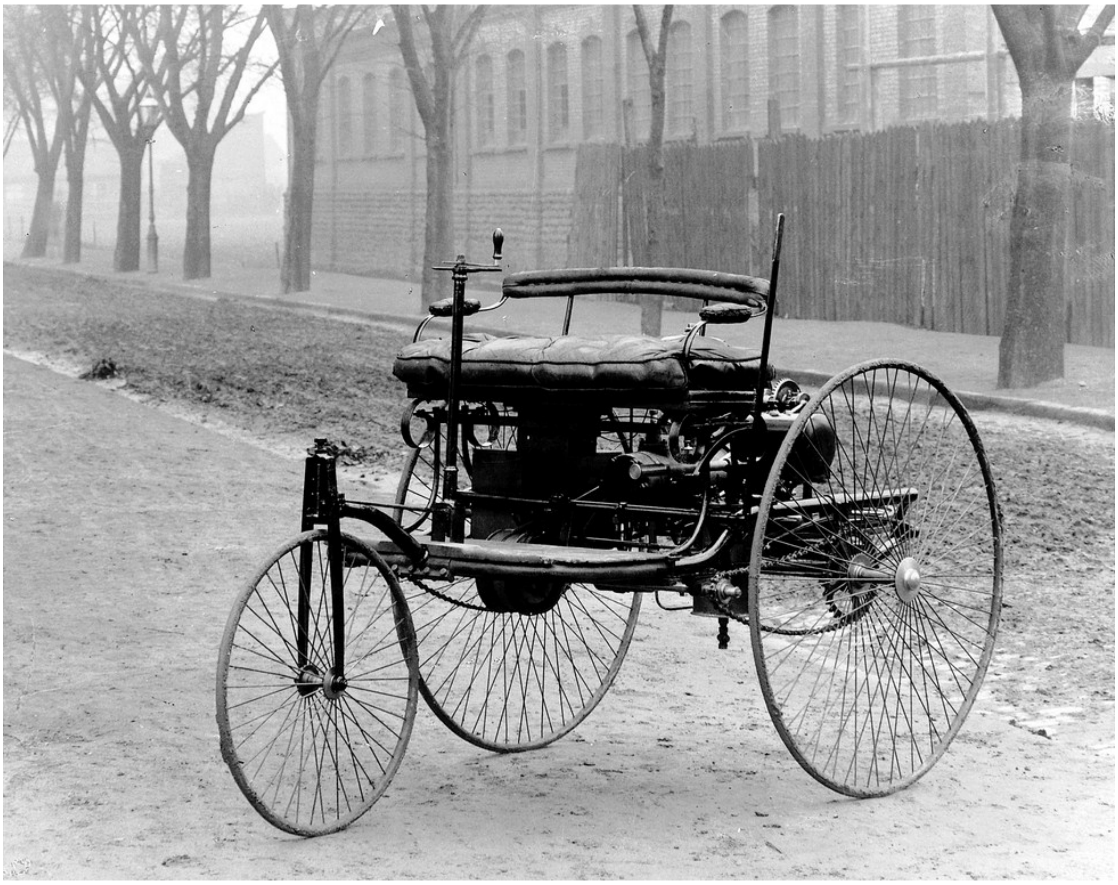

Figure 1: Benz Patent-Motorwagen (public domain)

tary technologies like connected, electric, and shared vehicles, has the potential to transform how people travel. Researchers predict these vehicles will increase transportation safety, convenience, access, and efficiency (Littman, 2018; Kockelman et al., 2016). However, these vehicles also present new security risks not found in human-controlled (level 0) vehicles (Table 1). (It is possible for a level 0 car to be equipped with drive-by-wire controls and therefore vulnerable to security risks, but such a configuration is unlikely because it would increase the cost of a vehicle and provide few benefits over mechanical controls.) Consequently, as the number of autonomous vehicles operating on public roads increases, so does the impact of new security risks. This thesis focuses on level 4 and level 5 vehicles because these vehicles do not expect the driver to respond in any situations. This thesis builds on prior security work inside and outside the automotive industry and presents a security risk assessment framework for autonomous

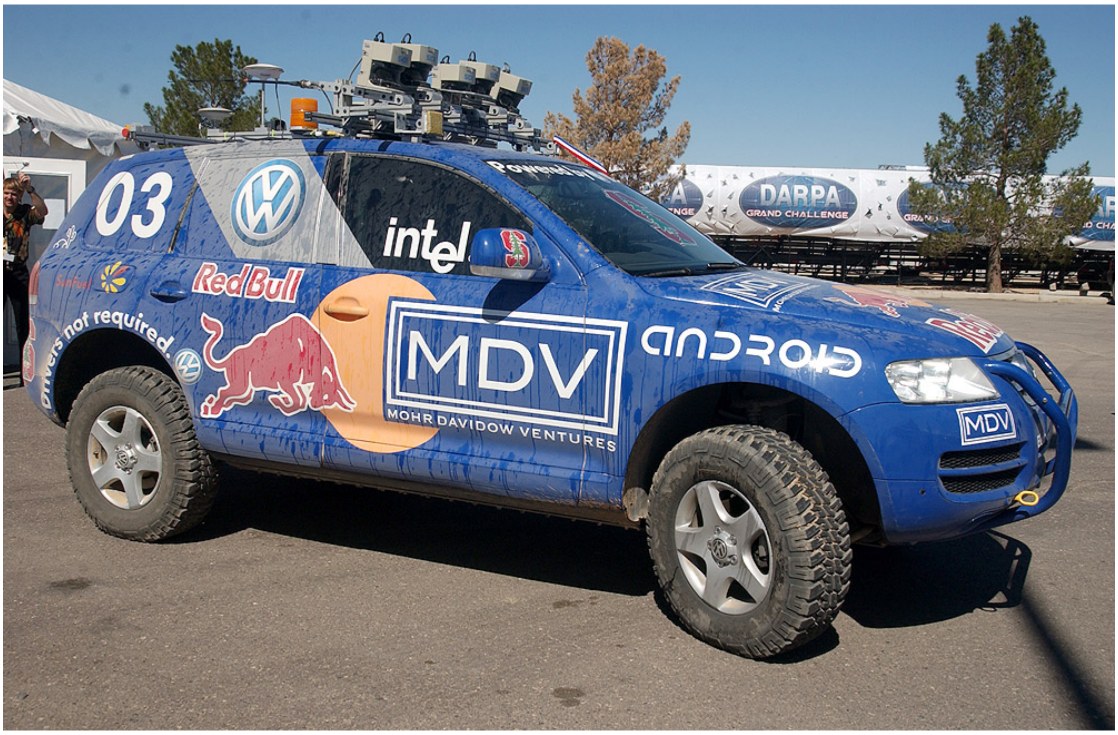

Figure 2: Stanley (public domain)

vehicles.

Level	Name	Dynamic driving task fallback
0	No Driving Automation	Driver
1	Driver Assistance	Driver
2	Partial Driving Automation	Driver
3	Conditional Driving Automation	Driver
4	High Driving Automation	System
5	Full Driving Automation	System

Table 1: Summary of levels of driving automation (SAE International, 2018)

1.2 Autonomous and Teleoperated Vehicles

> *Programs must be written for people to read, and only incidentally for machines to execute.*
>
> – Harold Abelson (Abelson et al., 1996)

Autonomous vehicles are robotic systems based on the sense, plan, act paradigm (Figure 3) (Dickmanns et al., 1994; Christensen et al., 2015). Sensors on the vehicle gather information from the environment. These include lidar, radar, cameras, microphones, and ultrasonic sensors (for slow speeds) (Waymo, 2017). The vehicle's computer systems then combine the sensor inputs to build an understanding about the environment. Next, the vehicle plans its possible actions and compares those possibilities with an overall goal. Planning occurs at several levels: high-level planing including route selection, mid-level planning including lane selection, and low-level planning including vehicle speed and steering angle. The vehicle chooses the best set of actions and actuates electro-mechanical devices to control the vehicle. The entire process continually repeats until the vehicle arrives at its destination.

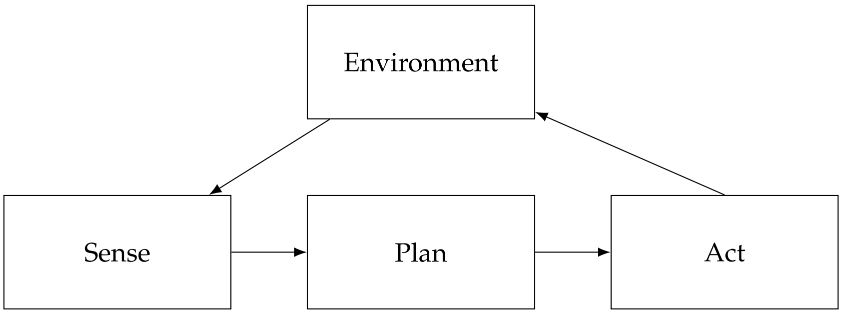

Figure 3: Sense, plan, act (Brooks, 1986)

Teleoperated autonomous vehicles can be controlled by a remote operator that complements or supplements the plan step in an autonomous system. (Teleoperated nonautonomous vehicles also exist but are rare.) In addition to the vehicle, teleoperation requires a command station to control the vehicle from and a communication network for passing information (sense inputs and act outputs) between the vehicle and the command station (Figure 4) (Tiwari and Seltz-Axmacher, 2018). Teleoperation systems can vary from a simple systems for a robotic taxi booking system that simply informs the vehicle of its destination to a complex system where a remote driver views video from the vehicle's cameras and controls acceleration, braking, and steering of the vehicle (Nissan Motor Corporation, 2017; Levinson et al., 2016; Okumura and Prokhorov, 2016; Fairfield and Herbach, 2016; Rust, 2017). California regulations currently require driverless vehicles to maintain a "communication link between the vehicle and remote operator" (Department of Motor Vehicles, 2018). To realize cost savings of autonomous vehicles companies would like to operate a greater number of autonomous vehicles than remote operators. The communication link tends to be multiple bonded cellular connections (Korosec et al., a,b).

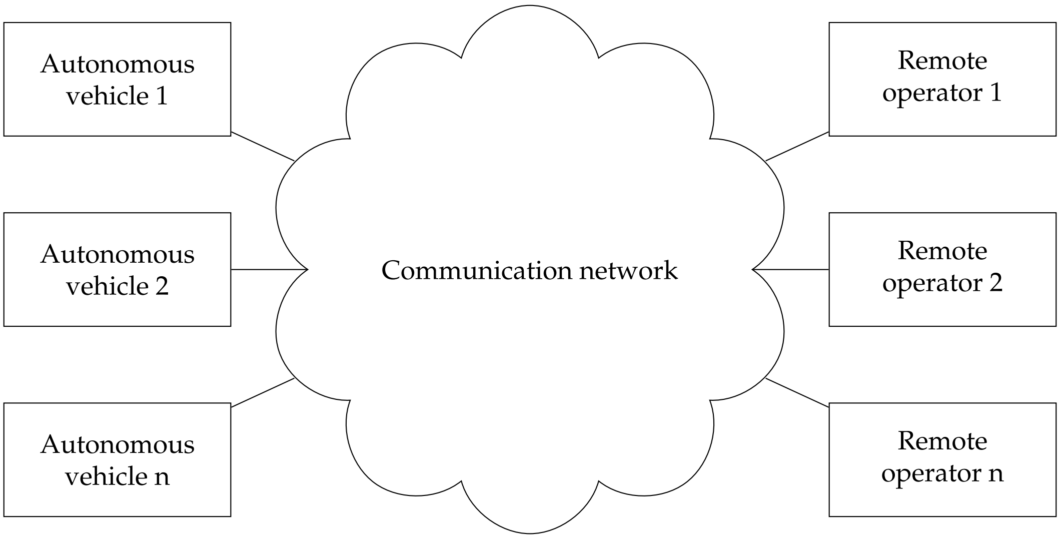

Figure 4: Autonomous vehicles and remote operators

1.3 Information Security and Vehicle Security

I would give all my fame for a pot of ale, and safety. – Boy in Henry V (Shakespeare, 1599)

Historically, computer systems processed information and did not interact with the physical world. Information security risks are classified as threats to confidentiality (unauthorized access), integrity (unauthorized modifications), and availability (uptime) of information (Figure 5) (McCumber, 2004; Gordon, 2015). While confidentiality is often the focus of traditional information security (protecting information from theft), integrity (protecting sensor input data and actuator output commands from tampering) is predicted to be the critical component for systems like autonomous vehicles (Schneier, 2016).

Figure 5: The three pillars of information security: confidentiality, integrity, and availability

In the 1980s, the US Department of Defense created policies and standards for securing computer systems (Department of Defense, 1983). Open standards followed years later (Holbrook and Reynolds, 1991), and as computer systems became prevalent in different industries, laws and standards protecting the information processed by these systems quickly followed (Table 2). However, only the North American Electric Reliability Corporation's (NERC) Critical Infrastructure Protection (CIP) standards which regulate the electric grid mention control systems that interact with the physical world.

Law or Standard	Applicability	Type	Author (Year)
California Senate Bill 1386	California residents	Law	Peace (2002)
CIS Controls	Any	Standard	Center for Internet Security (2018)
Critical Infrastructure Protection standards	North American energy sector	Standard	North American Electric Reliability Corporation
Family Education Rights and Privacy Act of 1974 (FERPA)	US education sector	Law	Buckley (1974)
Federal Information Security Management Act of 2002 (FISMA)	US federal government and contractors	Law	Davis III (2002)
General Data Protection Regulation (GDPR)	Individuals located inside the EU	Law	European Parliament and Council (2016)
Gramm-Leach-Bliley Act (GLBA)	US finance sector	Law	Gramm (1999)
Health Information Technology for Economic and Clinical Health (HITECH) Act	US healthcare sector	Law	McMorris Rodgers (2009)
Health Insurance Portability and Accountability Act of 1996 (HIPAA) Security Rule	US healthcare sector	Law	Archer Jr. (1996)
HITRUST CSF	US healthcare sector	Standard	HITRUST Alliance (2018)
International Standard for Assurance Engagements (ISAE) 3402 Assurance Reports on Controls at a Service Organization	Service organizations	Standard	International Federation of Accountants (2011)
IT Examination Handbook	US finance sector	Standard	Federal Financial Institutions Examination Council (FFIEC)

ISO/IEC 27000 family - Information security management systems	International	Standard	International Organization for Standardization (2013)
NIST Special Publication 800-53	US	Standard	National Institute of Standards and Technology (2013)
Payment Card Industry (PCI) Data Security Standards (DSS)	Credit and debit cards	Standard	PCI Security Standards Council (2016)
Sarbanes-Oxley Act of 2002 (SOX)	US public companies	Law	Oxley (2002)

Table 2: A sample of information security laws and standards

Information security attacks most often target finance and retail organizations. One simple explanation for this is the quote often misattributed to the bank robber Willie Sutton: "because that's where the money is" (Sutton and Linn, 2004). The transportation industry currently makes up less than 5% of security breaches (Verizon, 2017; Trustwave, 2017). Nevertheless, the current generation of vehicles controlled by computer systems may contain vulnerabilities, and if malicious individuals exploit these vulnerabilities, they could compromise the computer systems that control a vehicle (Koscher et al., 2010).

In 2010, researchers from the University of Washington and the University of California, San Diego demonstrated such an attack by controlling the engine, braking, and other systems of a 2009 passenger car (Koscher et al., 2010). In 2015 Charlie Miller and Chris Valasek demonstrated a similar attack by remotely controlling the acceleration, braking, steering, and other systems of an unaltered 2014 Jeep Cherokee (Miller and Valasek, 2015). Their research lead to the recall of 1.4 million vehicles and changes to the Sprint wireless carrier network that connected these vehicles to the internet. In 2016, they demonstrated additional attacks that controlled braking and steering systems (Miller and Valasek, 2016a). In 2016, and again in 2017, Keen Security Lab demonstrated another attack where they remotely controlled the infotainment system, windshield wipers, seat controls, mirror controls, door/trunk controls, and braking system of an unaltered Tesla Model S (Keen Security Lab of Tencent, 2016, 2017).

Other researchers have also demonstrated attacks against autonomous vehicle sensors and including Global Positioning System receivers (Lin and Qing, 2015), Lidar (Shin et al., 2017), and several sensors on the Tesla Model S (Yan et al., 2016). Additionally, researchers have successfully attacked deep learning models commonly used by autonomous vehicles (Evtimov et al., 2017; Goodfellow et al., 2017). However, no non-research vehicle security attacks have been publicly disclosed. The lack of attacks indicates the advanced skills required for these attacks. However, as the prevalence of autonomous vehicles grows, the motivation to conduct attacks will increase, and therefore the likelihood of attacks will increase.

While there are similarities between safety and security (and both words even translate to Sicherheit in German and seguridad in Spanish), this thesis refers to safety as freedom from harm and security as freedom from harm due to actions by external forces. Therefore safety is a broader concept that includes security, but security does not necessarily include accidental harm. Safety systems should include logic and bounds checks (Brooks, 1986). For example, a camera that is blinded by a laser can be detected because pixel values would not match those of a road. Additionally, safety systems should prevent vehicles from driving into pedestrians, cyclists, other vehicles, and fixed objects. Last, if the system commands and immediate turn while traveling at highway speeds, the system should reject this command. If these checks fail, the system should default to a safe state such as stopping the vehicle.

1.4 Risk

The best-laid schemes o' mice an' men / Gang aft agley.

(The best-laid plans of mice and men / Go oft awry.) – Robert Burns (1785)

Federal Information Processing Standard 200 defines risk as ”a measure of the extent to which an entity is threatened by a potential circumstance or event, and typically a function of: (i) the adverse impacts that would arise if the circumstance or event occurs; and (ii) the likelihood of occurrence” (National Institute of Standards and Technology, 2006). Similarly, ISO 31000 defines risk as ”effect of uncertainty on objectives,” which implies that risk is neither positive or negative (International Organization for Standardization, 2018b). Risk is defined in Equation 1.

$$Risk = Impact \times Likelihood \tag{1}$$

NIST Special Publication 800-30 describes risk as a threat source which initiates a threat event which exploits a vulnerability which causes an adverse impact which produces organizational risk (Figure 6) (National Institute of Standards and Technology, 2012). The publication defines four categories for threat sources (adversarial, accidental, structural, and environmental) (National Institute of Standards and Technology, 2012, Appendix D) and two categories for threat events (adversarial and non-adversarial) (National Institute of Standards and Technology, 2012, Appendix E). This process again demonstrates how risk is a combination of an impact and a likelihood of a threat source initiating a threat event which exploits a vulnerability. Adverse impact is typically a loss of confidentiality, integrity, or availability of information.

1.5 Controls

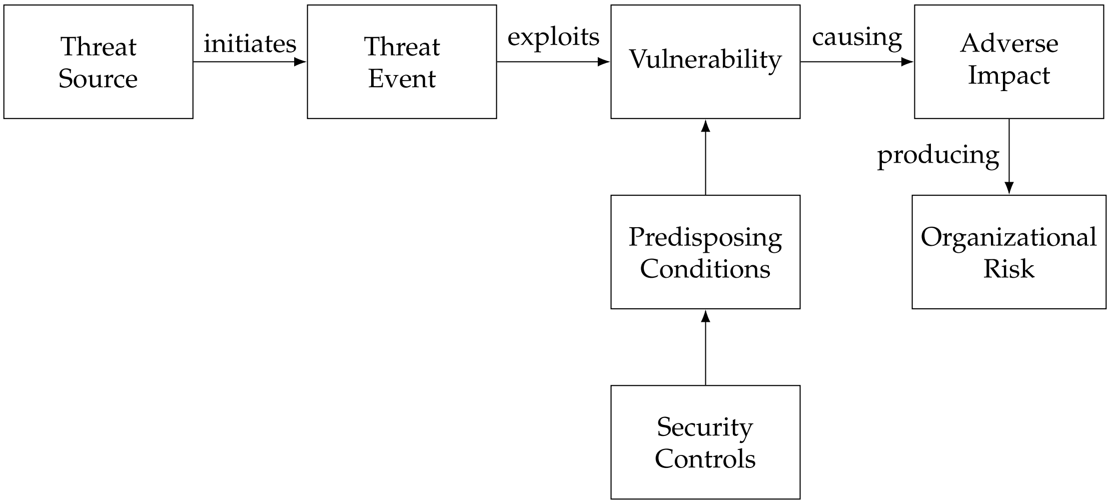

Figure 6: Generic risk model (National Institute of Standards and Technology, 2012)

> *The ultimate cause of our failure was a simple one: despite all statements to the contrary, it was not due to lack of bravery on the part of our men, or to any fault of the Fleet's. We were defeated by one thing only – by the inferior science of our enemies. I repeat – by the inferior science of our enemies.*
> – Arthur C. Clarke (1951)

NIST Special Publication 800-53 defines security controls as ”the safeguards/countermeasures prescribed for information systems or organizations that are designed to: (i) protect the confidentiality, integrity, and availability of information that is processed, stored, and transmitted by those systems/organizations; and (ii) satisfy a set of defined security requirements” (National Institute of Standards and Technology, 2013). Controls can be classified as administrative, technical, or physical. Table 17 lists the Security Control Identifiers from NIST Special Publication 800-53. Charlie McCarthy and Kevin Harentt believe the vehicle sector should develop a ”Security Control Catalog” based on NIST Special Publication 800-53 similar to the guides available for energy, control systems, and public transportation industries (McCarthy and Harnett, 2014). In GM's comments on NHTSA's Guidelines for the Safe Deployment and Operation of Automated Vehicle Safety Technologies, they propose a controls including ”security design reviews and penetration testing, designing and understanding system defensive measures, and the development of monitoring, detection, and response capabilities” (National Highway Traffic Safety Administration, 2016). Other common lists of controls are the CIS Controls from the Center for Internet Security (Figure 18) and the PCI Data Security Standard – High Level Overview from the Payment Card Industry (Table 11). Adding controls modifies the risk equation (Equation 1) by adding a controls factor (Equation 2).

$$Residual\ Risk = Impact \times Likelihood \times Controls \tag{2}$$

1.6 Risk Management

Risk Management: It's not rocket science - it's much more complicated. – John Adams (2005)

Risk Management is an organizational process to measure and manage risk. NIST Special Publication 800-37 defines the Risk Management Framework (Figure 7) which is a common risk management framework. ISO/IEC 27005 - Information security risk management provides an alternative to NIST's framework, but is not freely available (International Organization for Standardization, 2018a). The Federal Information Security Management Act of 2002 (FISMA) requires federal organizations to follow the Risk Management Framework to manage information security risks (Davis III, 2002). McCarthy and Harnett modified the Risk Management Framework into the Risk Management Framework for the Vehicle Sector (Figure 8) (McCarthy and Harnett, 2014). In Google's comments on NHTSA's Guidelines for the Safe Deployment and Operation of Automated Vehicle Safety Technologies, they propose a risk management approach: "self driving vehicles should be designed with a process that can identify known threats (malicious or otherwise) to the vehicle's electronic systems and explain how they have been mitigated by design (e.g. Google's web services use ISO 27001:2013 to validate the design of their security risk management processes)" (National Highway Traffic Safety Administration, 2016).

Risk management is common in many industries as a way to protect an organization from uncertainty. Risks could be financial risk, strategic risk, political risk, or other types of risk. This thesis chooses to focus on security risk because of the unique applicability to autonomous vehicle systems and excludes other types of risk. Organizations have four options for handling risks once identified:

- Avoidance: eliminate the cause of the risk

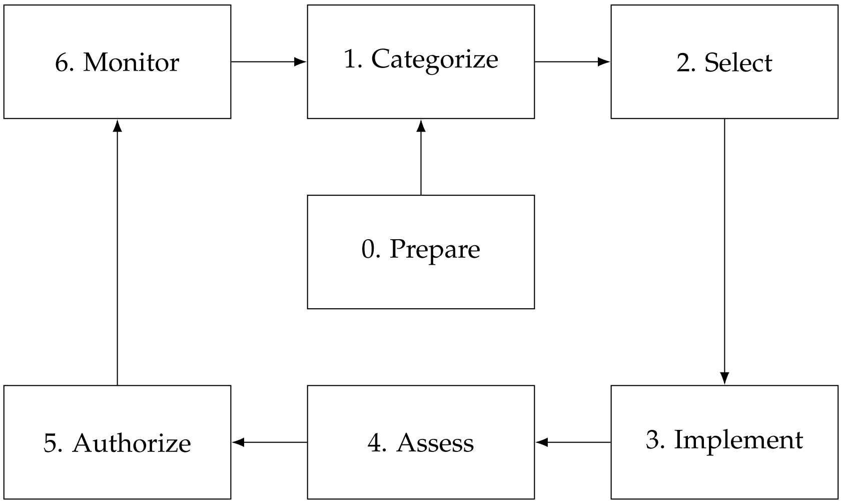

Figure 7: Risk Management Framework (National Institute of Standards and Technology, 2017, Page 8)

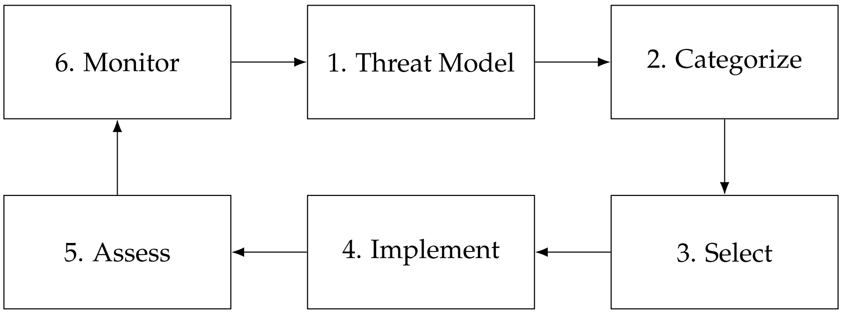

Figure 8: Risk Management Framework for the Vehicle Sector (McCarthy and Harnett, 2014)

- Control: implement controls to reduce the risk
- Transferance: contract with a third party to buy insurance against the risk, hedge against the risk, or outsource the risk
- Acceptance: accept the risk

Scope is an important tool in risk management. Limiting the size of a system or the number of interconnected systems can greatly reduce the impact of risk and therefore the overall risk management process. This is addressed during the categorization step of the risk management framework along with security level classification of the system. For example, a mobile application is public and accessible to everyone, but vehicle's internal control systems need additional controls to protect them. Classifications are common for computer information systems (Markiewicz and Raderman, 2015; Stanford University IT).

Another import concept in risk management is defense in depth where multiple redundant layers of security controls provide redundancy in case one or more security controls fails (Gordon, 2015). Defense in depth is also called the onion model where layers of an onion represent layers of security controls. The Swiss cheese model elaborates on this concept by describing each layer of security as a piece of cheese where the holes are vulnerabilities. A system failure requires the holes from all layers of cheese to align.

1.7 Risk Assessment

> *The revolutionary idea that defines the boundary between modern times and the past is the mastery of risk: the notion that the future is more than a whim of the gods and that men and women are not passive before nature. Until human beings discovered a way across that boundary, the future was the mirror of the past or the murky domain of oracles and soothsayers who held a monopoly over knowledge of anticipated events.* – Peter Bernstein (1998)

A risk assessment is a tool for measuring risk in a system. NIST Special Publication 800-30 defines the end result of a risk assessment as ”a determination of risk (i.e., typically a function of the degree of harm and likelihood of harm occurring)” (National Institute of Standards and Technology, 2012, page 1). NIST Special Publication 800-30 publication defines a process for a risk assessment (Figure 9). This process is used as the basis of section 2.

1. Prepare for Assessment
2. Conduct Assessment
   (a) Identify threat sources and events
   (b) Identify vulnerabilities and predisposing conditions
   (c) Determine likelihood of occurrence
   (d) Determine magnitude of impact
   (e) Determine risk
3. Communicate Results
4. Maintain Assessment

Figure 9: Risk Assessment Process (National Institute of Standards and Technology, 2012)

1.8 Risk Quantification

> *I often say that when you can measure what you are speaking about, and express it in numbers, you know something about it; but when you cannot measure it, when you cannot express it in numbers, your knowledge is of a meagre and unsatisfactory kind; it may be the beginning of knowledge, but you have scarcely, in your thoughts, advanced to the stage of science, whatever the matter may be.*
> – Lord Kelvin (Thomson, 1883)

NIST Special Publication 800-30 provides qualitative table (Table 3) for computing risk from likelihood and impact values. Risk values from ”Very Low” to ”Very High” map to a qualitative scale and semi-quantitative (ordinal) scales (Table 4). There are numerous other risk measures including the Common Vulnerability Scoring System (CVSS), Microsoft’s DREAD model, the OWASP Risk Rating Methodology, MIL-STD-882E - System Safety, and the Automotive Safety Integrity Level (ASIL) (Table 5).

Likelihood	Level of Impact				
	Very Low	Low	Moderate	High	Very High
Very High	Very Low	Low	Moderate	High	Very High
High	Very Low	Low	Moderate	High	Very High
Moderate	Very Low	Low	Moderate	Moderate	High
Low	Very Low	Low	Low	Low	Moderate
Very Low	Very Low	Very Low	Very Low	Low	Moderate

Table 3: Assessment Scale - Level of Risk (National Institute of Standards and Technology, 2012, page I-1)

Qualitative Values	Semi-Quantitative Values		Description
Very High	96-100	10	Very high risk means that a threat event could be expected to have multiple severe or catastrophic adverse effects on organizational operations, organizational assets, individuals, other organizations, or the Nation.
High	80-95	8	High risk means that a threat event could be expected to have a severe or catastrophic adverse effect on organizational operations, organizational assets, individuals, other organizations, or the Nation.
Moderate	21-79	5	Moderate risk means that a threat event could be expected to have a serious adverse effect on organizational operations, organizational assets, individuals, other organizations, or the Nation.
Low	5-20	2	Low risk means that a threat event could be expected to have a limited adverse effect on organizational operations, organizational assets, individuals, other organizations, or the Nation.
Very Low	0-4	0	Very low risk means that a threat event could be expected to have a negligible adverse effect on organizational operations, organizational assets, individuals, other organizations, or the Nation.

Table 4: Assessment Scale - Level of Risk (National Institute of Standards and Technology, 2012, page I-2)

Framework	Scale	Mesures	Author
NIST Special Publication 800-30	qualitative and semi-quantitative (ordinal)	likelihood, impact	National Institute of Standards and Technology (2012)
Common Vulnerability Scoring System	semi-quantitative (0-10)	attack vector, attack complexity, privilege required, user interaction, C,I,A impact, exploit code maturity, remediation level, report confidence, and security requirements	FIRST.org, Inc. (2015)
Microsoft's DREAD	semi-quantitative (1-10)	damage, reproducibility, exploitability, affected users, and discoverability	LeBlanc (2007)
OWASP Risk Rating Methodology	qualitative and semi-quantitative (0-9)	likelihood based on threat agent factors and vulnerability factors, impact based on technical impact factors and business impact factors	Open Web Application Security Project (2016)
MIL-STD-882E - System Safety	qualitative	severity, probability	Department of Defense (2012)
Automotive Safety Integrity Level (ASIL)	qualitative	probability, severity, and controllability	International Organization for Standardization (2011)

Table 5: A sample of risk measures

However, these methods are all ”handicapped by a reliance on non-quantitative methodologies” (Soo Hoo, 2000). Qualitative and semi-quantitative risk scales all suffer from problems of interpretability (e.g. Does a highly likely event have a 50%, 60%, 70%, 80%, or 90% probability of occurring?) and computation (e.g. Are three medium risk events worse than one high risk event?) (Hubbard and Seiersen, 2016, chapter 5). Hubbard and Seiersen recommend quantitative (ratio scale) risk assessments based on the probability (likelihood) and impact (dollars) of threat events. Other companies use a quantitative scale based on the value of bug bounty programs (Held and Baghdasaryan, 2017).

1.9 Research Questions and Contributions

This goal of this thesis is to provide insight and answers to several questions regarding the security of autonomous vehicle systems:

1. Can risk management frameworks quantitatively measure the security risks of autonomous vehicle systems?

2. Can risk management frameworks help manufacturers of autonomous vehicle systems prioritize security controls for these systems?

3. Can risk management frameworks provide assurances that autonomous vehicle systems are secure against cyberattacks?

This thesis contributes several original risk management techniques and example applications of those techniques. This thesis presents the first risk assessment of an autonomous vehicle system. This is the first time that a threat assessment has been applied to an autonomous vehicle system to determine the likelihood and impacts of security risks. Additionally, it is the first time that risk reduction through controls has been evaluated as part of such a risk assessment.

This thesis is also one of the first applications of quantitative risk management techniques to the security community. While quantitative techniques are popular in finance, security assessments often rely on qualitative methods which have many drawbacks listed above. This thesis uses quantitative methods to overcome those drawbacks and quantify the potential security impacts of autonomous vehicle systems. Additionally, this is the first time optimization methods have been used to select the optimal security controls to reduce risk in an autonomous vehicle system.

Last, this thesis provides a method for manufacturers of autonomous vehicle systems to provide assurances that the security risks in their systems are as low as reasonably practicable. This is an important step in public acceptance of autonomous vehicle systems.

1.10 Structure of the Thesis

The methods of this thesis are based on an adaptation of NIST Special Publication 800-30 risk assessment process (Figure 9) using quantitative methods instead of qualitative methods. This adapted framework provides the ability to compute risk using either discrete methods based on mean likelihoods and impacts or stochastic methods based on Monte Carlo simulations of likelihood and impact distributions. A controls selection step is added after the assessment which uses simple optimization techniques to balance the costs of controls and the expected loss from the risks they mitigate.

The results section describes the risk assessment and controls selection process for a sample autonomous vehicle system. The results contain many examples of controls that mitigate one or several risks. These examples provide several scenarios for control selection. Some scenarios reduce the expected loss greater than their cost of their controls and some do not. An expected loss curve is also presented for a base scenario with no controls and a optimal controls scenario. As this is an example system there are numerous assumptions behind the risks and controls.

The discussion focuses on areas where the autonomous vehicles industry can improve its security posture and how risk management fits into the overall security picture. Additionally, areas where this thesis is lacking and areas for future research are discussed.

2 Methods

This process generally follows the Risk Assessment Process from NIST Special Publication 800-30 (Figure 9), but uses quantitative values in place of qualitative scales. A risk assessment library was created in the python programming language for all computations (Bailey, 2018).

> *The Three Laws of Robotics*
>
> *1. A robot may not injure a human being or, through inaction, allow a human being to come to harm.*
>
> *2. A robot must obey the orders given it by human beings except where such orders would conflict with the First Law.*
>
> *3. A robot must protect its own existence as long as such protection does not conflict with the First or Second Laws.*
>
> – Isaac Asimov (1950)

2.1 Prepare for Assessment

> *When you fail to prepare, you're preparing to fail.* – John Wooden (Cromwell, 1977)

The tasks in preparation for the assessment are to identify the purpose (desired outputs), scope, assumptions and constraints, information sources, and risk model and analytic approach for the risk assessment (National Institute of Standards and Technology, 2012, pages 24-28). Similar to how the PCI Data Security Standards define their scope as the "people, processes and technologies that store, process, or transmit cardholder data or sensitive authentication data," a vehicle assessment's scope include any people, processes, and technologies that could be used to control a vehicle (PCI Security Standards Council, 2016). Ford considers the scope of their security program to include "not only to the vehicles' electronics, sensors and Virtual Driver System but also to any feature connected to them" (Ford Motor Company, 2018, page 35). Also, Miller and Valasek demonstrated the scope of the Jeep Cherokee they hacked was not limited by proximity to the vehicle when they remotely connected to the vehicle via the internet and were able to control it (Miller and Valasek, 2015). Assessments can be simplified by separating the system into subsystems and with subsystem boundaries and assessing the individual subsystems and the boundaries (National Institute of Standards and Technology, 2017, page 13). This is recommended to reduce scope, cost, difficulty, and risk (PCI Security Standards Council, 2016, page 11).

2.2 Conduct Assessment

The generic risk model (Figure 6) can be represented as figure 10 to match the risk assessment process.

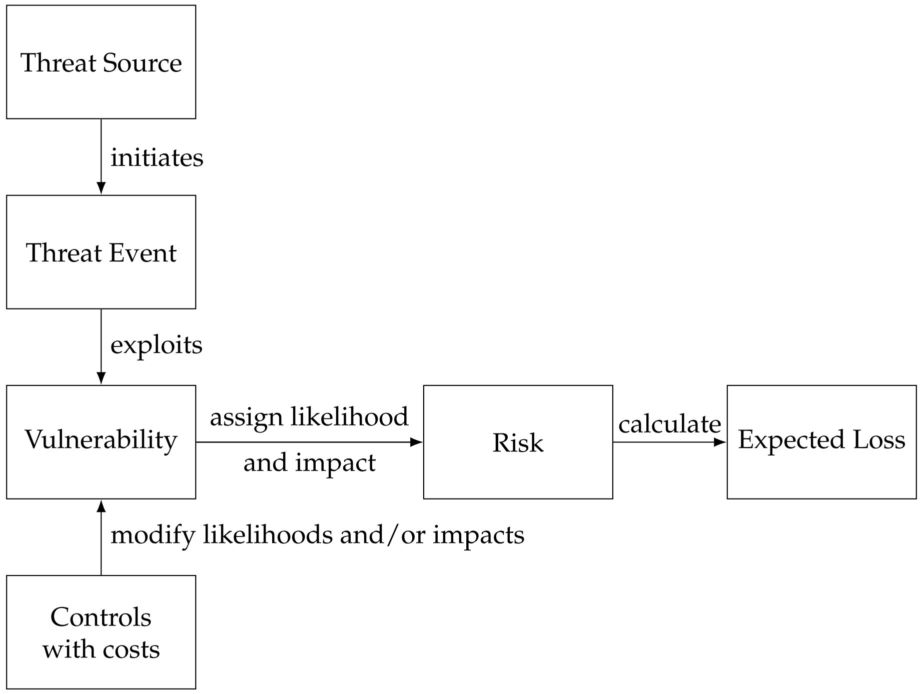

Figure 10: Risk model

2.2.1 Identify threat sources and events

Generally, threat sources can be grouped into two categories: adversarial and non-adversairal (National Institute of Standards and Technology, 2012). Threat events are a function of threat sources (Equation 3).

$$Threat\ Event = f(Threat\ Source) \tag{3}$$

NIST Special Publication 800-30 provides threat event categories and sample threat events (Figure 11). Petit and Shladover created threat models for autonomous and connected vehicles. They determined the attack surface for autonomous vehicles to consist of infrastructure signs, machine vision, GPS, in-vehicle devices, acoustic sensors, radar, lidar, roads, in-vehicle sensors, odometric sensors, electronic devices, and maps (Petit and Shladover, 2015). However they do not include any teleoperation capabilities that are now common in autonomous vehicles or assign a likelihood or impact to any attacks.

Threat Events

- Adversarial threat events
    - Perform reconnaissance and gather information.
    - Craft or create attack tools.
    - Deliver/insert/install malicious capabilities.
    - Exploit and compromise.
    - Conduct an attack (i.e., direct/coordinate attack tools or activities).
    - Achieve results (i.e., cause adverse impacts, obtain information)
    - Maintain a presence or set of capabilities.
    - Coordinate a campaign.
- Non-adversarial threat events

Figure 11: Threat event categories (National Institute of Standards and Technology, 2012, Appendix E)

Other organizations such as Germany's Federal Office for Information Security (Bundesamt für Sicherheit in der Informationstechnik), the Cloud Security Alliance, and the Open Web Application Security Project (OWASP) have created threat catalogs (Bundesamt für Sicherheit in der Informationstechnik, 2008; Cloud Security Alliance Top Threats Working Group, 2016; Open Web Application Security Project, 2017). Additionally, several books list hacking methodologies that can be helpful for brainstorming additional threats (McNab, 2007; Smith, 2016).

Threats to autonomous vehicles can also be classified according to if they attack the sense, plan, or act systems, or the environment (Figure 3). Traditional information security models such as the OSI model or the Internet model (Table 6) are helpful for grouping different threats. Last, data flow diagrams can help trace information through an information system. Similarly, for control systems such as autonomous vehicles, control flow diagrams describe the flow of information within these systems.

OSI model	Internet model
7. Application	Application
6. Presentation	
5. Session	
4. Transport	Transport
3. Network	Internet
2. Data link	
1. Physical	Link

Table 6: OSI model (International Organization for Standardization, 1994) and Internet model (Braden, 1989)

All of the above inputs were combined to form a threat model (Table 7) for an example autonomous vehicle system (Figure 13). Threat sources are a class in the risk assessment library. Threat events are a class which depend on a threat source.

2.2.2 Identify vulnerabilities and predisposing conditions

Vulnerabilities are a function of threat events, security controls, and a system (with predisposing conditions) (Equation 4). Controls each come with a cost and can modify either the impact or likelihood of a risk. Additionally, controls often do not reduce risk to zero, but rather by some factor. For example, anti-malware software will prevent infection from some malware, but not zero-day attacks (Vegge et al., 2009). Also, passwords, another common control, are vulnerable to guessing, observing, viewing when written down, and more (Bryant and Campbell, 2006). Eliminating a vulnerability can reduce the risk of that vulnerability to zero.

The concept of a kill chain can help in identifying vulnerabilities. The phases of a kill chain (reconnaissance, weaponization, delivery, exploitation, installation, command and control, action on objectives) show how attacks must be successful in all phases to create risk while defenses may block an attack at any phase (Pols, 2017, page 19). McNab offers three simpler phases (reconnaissance, enumeration, and exploration) to an attack (McNab, 2007).

$$Vulnerability = f(Threat\ Event,\ Controls,\ System) \tag{4}$$

As described in equation 4, the threat model from section 2.2.1 generates vulnerabilities. In the risk assessment library, vulnerabilities are a class which depend on a threat event, system, and controls. Systems are a tree which describe the autonomous vehicle system. Controls are a class which have a cost and a reduction factor. They also have a boolean to determine if they are implemented or not for use in calculations and optimizations.

2.2.3 Determine likelihood of occurrence

Likelihoods are defined as the number of expected occurrences per year. We represent likelihoods by a Poisson distribution because they are discrete and positive. Likelihoods are a function of vulnerabilities which themselves are a function of threat events, controls, and systems (Equation 5). Likelihoods can be determined by analysis of historical data or expert surveys (Joh and Malaiya, 2017). Two example likelihood histograms are show in Figure 12 representing .5 events/year and 100 events/year.

$$Likelihood = f(Vulnerability) \tag{5}$$

In the risk assessment library, likelihoods are a class defined by the lambda of the Poisson distribution.

2.2.4 Determine magnitude of impact

Impacts are losses defined by monetary units such as United States dollars or euros. These losses are represented by log normal distributions because they are continuous, positive, and have a long tail (Hubbard and Seiersen, 2016). Impacts, like likelihoods, are a function of vulnerabilities (Equation 6). Impacts can be determined by analysis of historical data or expert surveys (Joh and Malaiya, 2017).

Quantitative risk modeling is common in the financial sector and there is quite a large amount of research on financial impacts. One main difference between finance and security is the type of impacts. Financial models define risk as both profit and loss: an investment has the potential to increase or decrease in value (Embrechts et al., 2005; McNeil et al., 2015). However, for security risk, we are only concerned with losses. The same is true for transportation safety

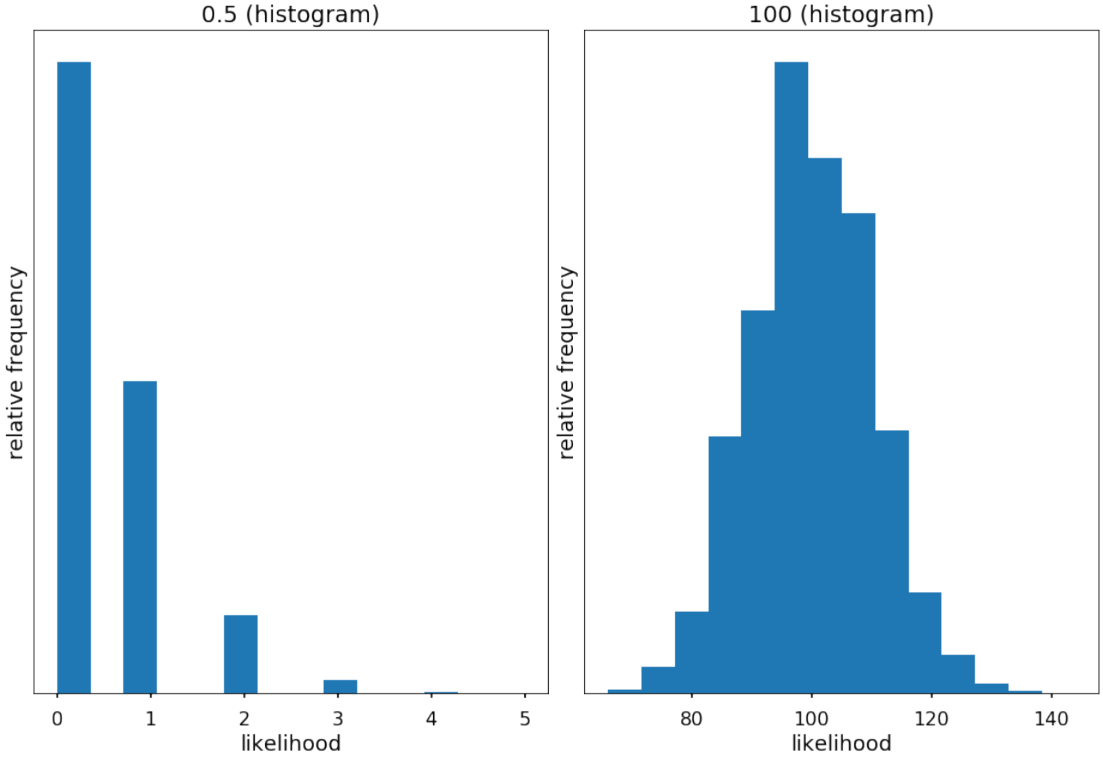

Figure 12: Example likelihood distributions

models such as a road safety models where impacts may be defined in fatalities or crashes.

$$Impact = f(Vulnerability) \tag{6}$$

Impacts to autonomous vehicles can be classified into a few broad categories: attacks where an attacker compromises the control systems of a vehicle to cause an intentional crash, attacks where an attacker causes these systems to fail uncontrollably and the vehicle to crash, and attacks which disable these systems and the vehicle. Additionally, these attacks may involve one vehicle or an entire fleet of vehicles. Other attacks may have impacts similar to traditional information security attacks such as information theft.

In the risk assessment library, impacts are a class defined by the mu and sigma terms of the log normal distribution. There is and additional method for defining an impact based on the lower 90% and upper 90% confidence interval.

2.2.5 Determine risk

Risk is calculated by both deterministic and stochastic methods. Deterministic values are calculated from the mean likelihood and mean impact of each risk using the residual risk equation (Equation 2) from section 1.5. Stochastic risk values are calculated from random sampling of likelihood and impact distributions, again using the residual risk equation. Using the stochastic risk values, a Monte Carlo simulation, over thousands of iterations, samples values of the residual risk at each iteration (Carey et al., 2006, page 477). The result of a deterministic calculation is one value, the annualized expected loss mean, and the result of a stochastic calculation is a probability density function of annualized expected loss. The annualized expected loss is the summation of all residual risks to the autonomous vehicle system (Equation 7).

$$Annual\ Expected\ Loss = \sum Residual\ Risk \tag{7}$$

In the risk assessment library, risks are a class defined by a vulnerability, likelihood, and impact. There are methods to calculate deterministic and stochastic risk values as defined above. There are also methods to perform Monte Carlo simulations of stochastic risk values. For all methods, controls may be implemented or not.

2.3 Communicate Results

The result of the risk assessment, as mentioned in section 1.7, is a determination of risk. In a deterministic calculation, this is the annualized expected loss of the system. The Monte Carlo simulation results in a distribution of annualized expected loss. Additional controls can be evaluated based upon their cost effectiveness. Additionally, these results can be compared with an organization's risk tolerance to determine if additional controls, perhaps cost ineffective, should be implemented.

In the risk assessment library, a risks class is defined by all the risks of a system. This class has a method to calculate the mean annualized expected loss of a system and a method to calculate and plot the annualized expected loss distribution.

2.4 Maintain Assessment

The risk assessment provides the ability to be continuously updated to maintain its effectiveness. When new threat events are discovered, they should be added to the assessment. When additional information about likelihoods and impacts is available, it should be added to the assessment.

2.5 Optimization of Controls

Happiness equals reality minus expectations. – Tom Magliozzi

While selecting controls is step 2 of the Risk Management Framework (Figure 7) or part 3 of the Risk Management Framework for the Vehicle Sector (Figure 8), a quantitative risk assessment provides the ability to select the optimal controls for a system. By balancing the cost of controls with the annual expected loss, the we can compute the total cost of a system (Equation 8).

$$Total\ Cost = Annual\ Expected\ Loss + Control\ Cost \tag{8}$$

We can then minimize this function using optimization techniques. However, because the number of risks and controls is limited, trying all combinations of controls is often the simplest method for finding the optimum. Alternatively, a weight factor can be added to the annual expected loss to add or subtract the impact of this term (Equation 9).

$$Total\ Cost = Weight\ Factor \times Annual\ Expected\ Loss + Control\ Cost \tag{9}$$

In the risk assessment library, the risks class has methods to determine the optimum controls based on the above cost function (equation 8). There is also a method to plot all of the different control combinations.

3 Results

This is our world now... the world of the electron and the switch, the beauty of the baud. We make use of a service already existing without paying for what could be dirt-cheap if it wasn't run by profiteering gluttons, and you call us criminals. We explore... and you call us criminals. We seek after knowledge... and you call us criminals. We exist without skin color, without nationality, without religious bias... and you call us criminals. You build atomic bombs, you wage wars, you murder, cheat, and lie to us and try to make us believe it's for our own good, yet we're the criminals.

Yes, I am a criminal. My crime is that of curiosity. My crime is that of judging people by what they say and think, not what they look like. My crime is that of outsmarting you, something that you will never forgive me for.

I am a hacker, and this is my manifesto. You may stop this individual, but you can't stop us all... after all, we're all alike.

– The Mentor (1986)

3.1 Risk Assessment

The scope of the risk assessment is an example autonomous vehicle system (Figure 13) including all people, processes, and technologies that control a vehicle. As there is a lack of publicly available data about risks and controls in current autonomous vehicle systems, these are assumed and explained below. This assessment assumes a fleet size of ten vehicles. This size demonstrates that an autonomous vehicle system can be more cost effective than employing human drivers. Several security controls have a per vehicle cost and therefore depend on this number. Also, although this thesis does not consider any, some risks could impact an entire vehicle fleet. The cost of security controls are assumed and show in table 8.

The likelihood of occurrence for each threat event is based on historical data. The Common Vulnerabilities and Exposures (CVE) (The MITRE Corporation, 2018a) system tracks public software vulnerabilities, assigns them a CVSS score, and stores them in the National Vulnerability Database (NVD) (National Institute of Standards and Technology). A similar system, Common Weakness Enumeration (CWE), exists for software weaknesses (The MITRE Corporation, 2018b). These databases are then used to model the likelihood of a risk based on historical

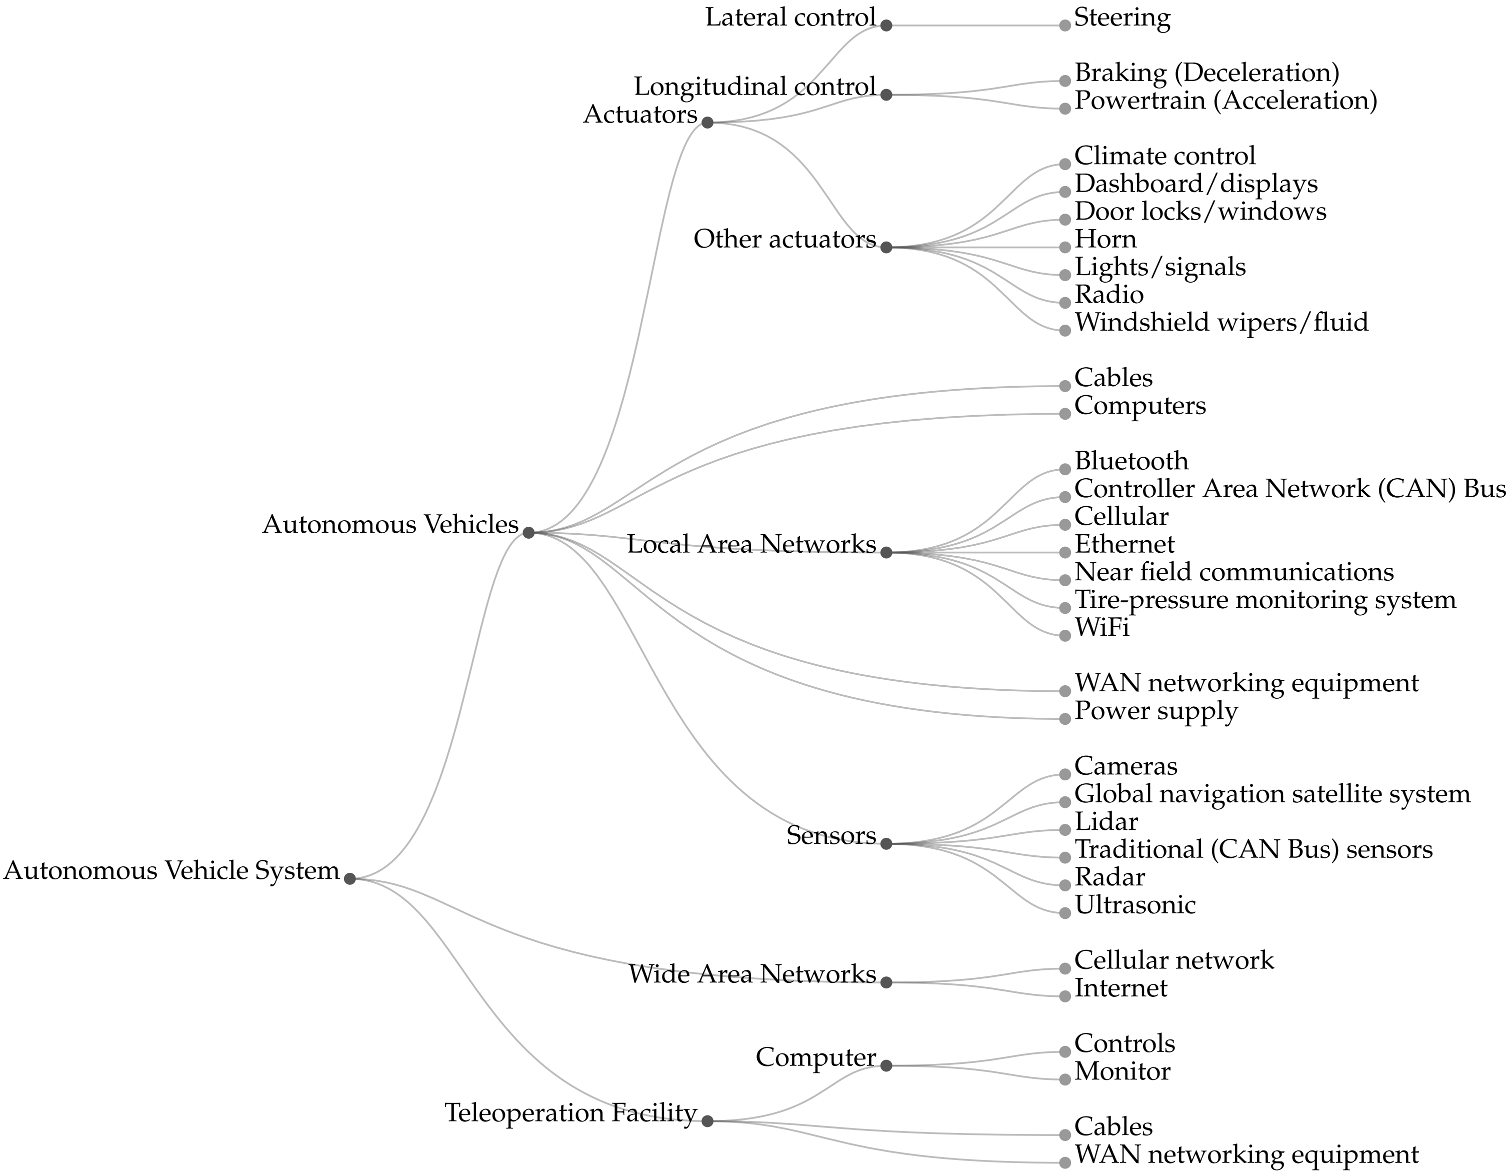

Figure 13: An autonomous vehicle system

occurrence rates. For vulnerabilities that are not in the NVD, likelihoods are assumed. Unfortunately, this approach assumes that future vulnerability rates will be similar to past rates which is often not the case. Also, this requires manual assessment of each vulnerability's inner workings to understand if the vulnerability would or would not apply to the autonomous vehicle system in question. Last, this technique does not take into account the adversary's likelihood of exploiting the vulnerability, only the presence of the vulnerability. It could be assumed that if the vulnerability exists, an adversary will find it and exploit it. However, a manufacturer may fix a vulnerability before it is exploited by an attacker. The likelihood in the model can reflect both cases by implementing a vulnerability management as a control. Penetration tests, code reviews, and other assessments can also discover vulnerabilities.

This model analyzes three impacts: attacks where an attacker compromises the control systems of a vehicle to cause an intentional malicious crash, attacks where an attacker causes these systems to fail uncontrollably and the vehicle to crash, and attacks which disable these systems and the vehicle (Figure 14). Disabled vehicles are assumed to have a mean cost of $10000 per event. These values will vary depending on the type of vehicle (personal, taxi, truck, bus) and the duration of the outage. Mean values for vehicle crashes in the United States are $70,830 for all vehicle types (Blincoe et al., 2015) and $86,600 for truck and bus crashes (Zaloshnja and Miller, 2004). These values were converted to 2018 dollars based on differences in past year consumer price index values (Bureau of Labor Statistics, 2018). Intentional malicious crashes are rare and their impacts vary greatly. Recent events include the 2016 Nice truck attack that killed 86 and injured 434 and the 2017 Barcelona attacks that killed 13 and injured more than 100 (Le Monde.fr with AFP, 2016; Smith-Spark, 2017). While there are several studies that evaluated the impacts of the September 11 attacks, few studies have assessed the impacts of smaller terrorist attacks (Jackson, 2008; Blalock et al., 2005). For this thesis, impact values were assumed to have a lower 90% confidence interval of $10,000,000 and a higher 90% confidence interval of $100,000,000 (Hubbard and Seiersen, 2016). Figure 14 shows these three impacts as log normal distributions which have a mean determined as above and s equal to 1. Ideally, these distributions would be based on actual historical distributions of the costs associated with autonomous vehicle systems. However, this data is not yet available, and in some cases, such as intentional crashes, may never be available. Therefore, some judgement was used in building these distributions and choosing the above parameters. These distributions are used for calculating stochastic risk values below.

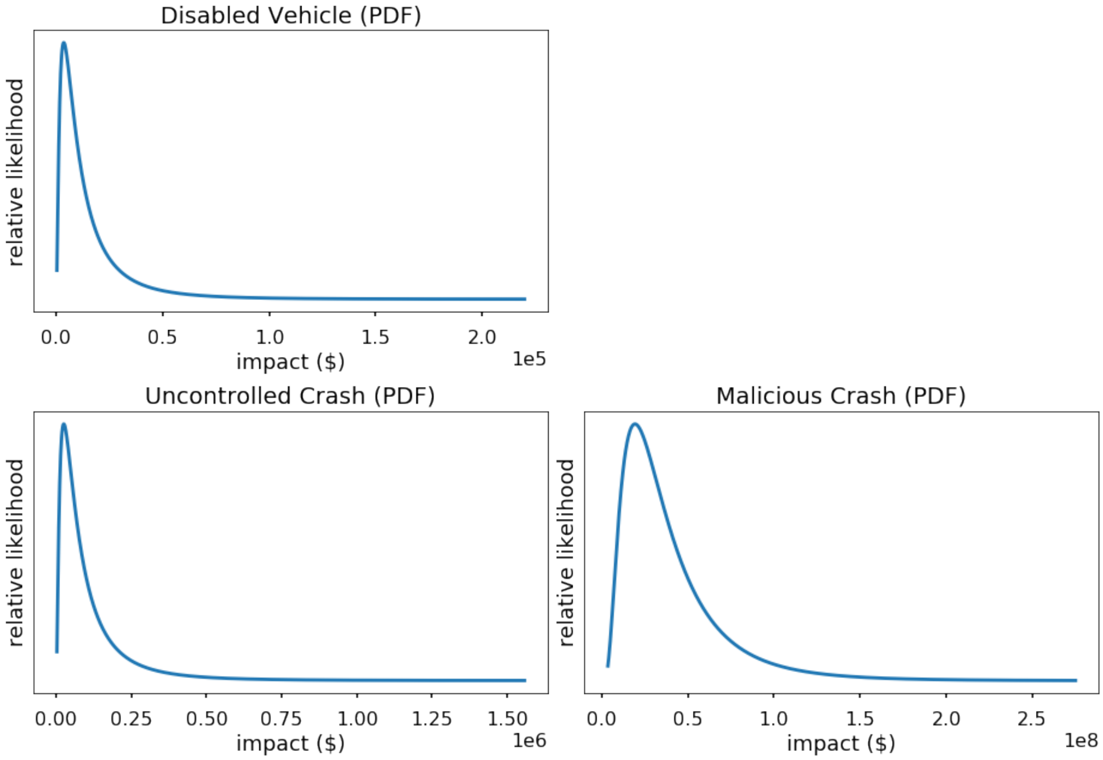

Figure 14: Impact probability density functions

Table 7 presents a summary of threat sources, threat events, systems, controls, impacts, likelihoods, and risks. None of the threat sources, threat events, systems, controls, impacts, likelihoods, or risks are meant to be exhaustive. They merely provide a sample of different situations and how these differences respond to the risk assessment framework. This assessment considers twenty-four threat events:

1. Attacker gains access to teleoperation computer via outdated software on teleoperation computer: In this attack, an attacker is able to run malicious code on a teleoperation computer because of outdated software on the computer. Examples of outdated software are missing operating system patches or missing application patches. The malicious code could be downloaded by a user of the computer, or executed via the network. Controls against this attack include vulnerability management programs and remote operation

protections. The likelihood of this attack is assumed to be once every ten years based on historical data and the impact is an intentional crash of a teleoperated vehicle.

2. Attacker gains access to teleoperation network via VPN: In this attack, an attacker gains remote access to a teleoperation facility via a VPN connection. Examples include weak VPN passwords and compromised accounts. This event is similar to the third event above but targets the teleoperation network instead of the vehicle network. Controls agains this attack include two-factor authentication and remote operation protections. The likelihood of this attack is assumed to be once every ten years based on historical data regarding password breaches and password reuse and the impact is an intentional crash of a teleoperated vehicle.

3. Attacker gains access to vehicle computer via outdated software: In this attack, an attacker is able to exploit a outdated software on a computer inside the vehicle. The attacker is then able to control the vehicle. Controls agains this attack include a vulnerability management program and remote operation protections. The likelihood of this attack is assumed to be once every ten years based on historical data of critical software vulnerabilities and the impact is an intentional crash of a vehicle.

4. Attacker gains access to vehicle via malware on teloperation computer: In this attack, an attacker is able to install malicious software on the teleoperation computer and then use this software to control the vehicle. Examples include a user who is tricked into installing malware or a malware that spreads across the network. Controls against this attack include anti-malware software and remote operation protections. The likelihood of this attack is assumed to be once every one hundred years and the impact is an intentional crash of a teleoperated vehicle.

5. Attacker gains access to vehicle via malware on vehicle computer: In this attack, an attacker is able to install malicious software on the a computer inside the vehicle and then use this software to control the vehicle. Examples include an attacker with physical access to the vehicle or a malware that spreads across the network. This event is similar to the one above except the malware is on a computer in the vehicle instead of a teleoperation

computer. Controls against this attack include anti-malware software and remote operation protections. The likelihood of this attack is assumed to be once every one hundred years and the impact is an intentional crash of a vehicle.

6. Attacker gains access to vehicle wireless network: In this attack, an attacker is able to access a wireless network on the vehicle that is connected to the control system of the vehicle. This assessment does not consider any controls against this event but considers disabling the wireless network which eliminates the risk. The likelihood of this attack is assumed to be once every one hundred years and the impact is an intentional crash of a vehicle.

7. Attacker gains physical access to teleoperation computer: In this attack, an attacker gains physical access to a teleoperation computer and is able to use the computer to control a vehicle. Examples include breaking and entering or social engineering attacks. Controls against this attack include physical security controls at teleoperation facilities and remote operation protections. The likelihood of this attack is assumed to be once every one hundred years and the impact is an intentional crash of a teleoperated vehicle.

8. Attacker gains physical access to vehicle network: In this event, an attacker is able to physically connect to the vehicle network and control the vehicle. Controls against this attack include improving the physical security of the vehicles. The likelihood of this event is assumed to be once every one hundred years and the impact is an intentional crash of a vehicle.

9. Attacker gains remote access to vehicle network by exploiting software bug on WAN networking equipment: In this attack, an attacker is able to exploit a software bug on the network equipment inside an autonomous vehicle. The attacker is then able to access the vehicle's network and control the vehicle. Controls against this attack include a vulnerability management program and remote operation protections. The likelihood of this attack is assumed to be once every one hundred years based on historical data and the impact is an intentional crash of a vehicle.

10. Attacker gains remote access to vehicle network via VPN: In this attack, an attacker gains

remote access to an autonomous vehicle's network via a diagnostic VPN connection. Examples include weak VPN passwords and compromised accounts. Controls against this attack include two-factor authentication and remote operation protections. The likelihood of this attack is assumed to be once every ten years based on historical data regarding password breaches and password reuse and the impact is an intentional crash of a vehicle.

11. Attacker is hired as teleoperation driver: In this attack, an attacker is able to gain employment as a teleoperation driver. Controls against this attack include background checks for teleoperation drivers. The likelihood of this attack is assumed to be once every one hundred years and the impact is an intentional crash of a vehicle.

12. Attacker jams radar signals: In this attack, an attacker is able to disrupt the radar signals from the autonomous vehicle and cause it to enter a fail-safe mode. This assessment does not consider any controls against this event. The likelihood of this attack is assumed to be once every one hundred years and the impact is a disabled vehicle.

13. Attacker spoofs radar signals: In this attack, an attacker manipulates the radar sensors of the autonomous such that surrounding cars are no longer detected. This assessment does not consider any controls against this event. The likelihood of this attack is assumed to be once every ten years and the impact is an accidental crash of a vehicle.

14. Denial of service attack against cellular network: In this event, an attacker disrupts the cellular network connection used for teleoperation of the vehicle. Examples of this attack include cellular blocking devices. This assessment does not consider any controls against this event. The likelihood of this event is assumed to be once every ten years and the impact is a disabled vehicle.

15. Theft of a teleoperation computer: In this event, a thief steals a computer from a teleoperation facility. Controls against this event include improving the physical security of the teleoperation facilities. The likelihood of this event is assumed to be once every ten years and the impact is a disabled vehicle.

16. Theft of a vehicle component: In this event, a thief steals a component from the vehicle. Examples include stealing a sensor or a computer. Controls against this event include improving the physical security of the vehicles. The likelihood of this event is assumed to be once every ten years and the impact is a disabled vehicle.

17. Bumps dislodge vehicle computers: In this event, a substantial bump in the road dislodges a computer in the vehicle. Examples include a cable coming lose, a computer resetting, or a computer failure. This assessment does not consider any controls against this event. The likelihood of this event is assumed to be once every year and the impact is an accidental crash of a vehicle.

18. Fire in the teleoperation facility: In this event, a fire starts in a teleoperation facility while a vehicle is being teleoperated. Examples include office fires. Controls against this event include a redundant teleoperation facility that can take over during a fire. The likelihood of this event is assumed to be once every one hundred years and the impact is an accidental crash of a vehicle.

19. Fire in the vehicle: In this event, a fire starts in the vehicle that destroys the control system. Examples include a passenger accidentally starting a fire or a vehicle component overheating and starting a fire. This assessment does not consider any controls against this event. The likelihood of this event is assumed to be once every ten years and the impact is an accidental crash of a vehicle.

20. Hardware failure of a teleoperation system: In this event, a critical hardware system in the teleoperation facility fails. Examples include a computer failure or power failure. Controls against this event include a redundant teleoperation facility that can take over during a failure. The likelihood of this event is assumed to be once every ten years and the impact is an accidental crash of a vehicle.

21. Hardware failure of a vehicle system: In this event, a critical hardware system in the vehicle fails. Examples include a computer failure, motor failure, or cable failure. This assessment does not consider any controls against this event. The likelihood of this event is assumed to be once every ten years and the impact is an accidental crash of a vehicle.

22. Rain/snow/fog disrupts sensors: In this event, weather events such as rain, snow or fog disrupts sensors on the vehicle. Examples include heavy rains obscuring a camera or snow obscuring a lidar sensor. This assessment does not consider any controls against this event. The likelihood of this event is assumed to be once every ten years and the impact is an accidental crash of a vehicle.

23. Vehicle crash dislodges sensors: In this event, an undetected vehicle crash physically moves a sensor from its calibrated position. Examples include a sideswipe crash that dislodges a side-mounted sensor. This assessment does not consider any controls against this event. The likelihood of this event is assumed to be once every year and the impact is an accidental crash of a vehicle.

24. Weather impacts cellular network: In this event, a weather event (rain, snow, etc.) reduces the ability of the cellular network to provide sufficient bandwidth for teleoperation. This assessment does not consider any controls against this event. The likelihood of this event is assumed to be two hundred times per year and the impact is a disabled vehicle.

Threat Source	Threat Event	System	Controls	Impact ($) (mean)	Likelihood (mean)	Risk (mean)
Adversarial	Attacker gains access to teleoperation computer via outdated software on teleoperation computer	Teleoperation Facilities / Computer	[Vulnerability management program, Autonomous system protections for remote operations]	40398367	0.10	403
Adversarial	Attacker gains access to teleoperation network via vpn	Teleoperation Facilities / WAN networking equipment	[Two-factor authentication, Autonomous system protections for remote operations]	40398367	0.10	4039
Adversarial	Attacker gains access to vehicle computer via outdated software	Autonomous Vehicle / Computers	[Vulnerability management program, Autonomous system protections for remote operations]	40398367	0.10	403
Adversarial	Attacker gains access to vehicle via malware on teloperation computer	Teleoperation Facilities / Computer	[Anti-malware software, Autonomous system protections for remote operations]	40398367	0.01	403
Adversarial	Attacker gains access to vehicle via malware on vehicle computer	Autonomous Vehicle / Computers	[Anti-malware software, Autonomous system protections for remote operations]	40398367	0.01	403
Adversarial	Attacker gains access to vehicle wireless network	Autonomous Vehicle / Local Area Networks / WiFi	[Disable WiFi]	40398367	0.01	0

Adversarial	Attacker gains physical access to teleoperation computer	Teleoperation Facilities / Computer	[Physical security at teleoperation facilities, Autonomous system protections for remote operations]	40398367	0.01	403
Adversarial	Attacker gains physical access to vehicle network	Autonomous Vehicle / Local Area Networks	[Physical security of vehicles]	40398367	0.01	40398
Adversarial	Attacker gains remote access to vehicle network by exploiting software bug on WAN networking equiptment	Autonomous Vehicle / WAN networking equipment	[Vulnerability management program, Autonomous system protections for remote operations]	40398367	0.01	40
Adversarial	Attacker gains remote access to vehicle network via vpn	Autonomous Vehicle / WAN networking equipment	[Two-factor authentication, Autonomous system protections for remote operations]	40398367	0.10	4039
Adversarial	Attacker is hired as teleoperation driver.	People	[Background checks for drivers]	40398367	0.01	4039
Adversarial	Attacker jams radar signals	Autonomous Vehicle / Sensors / Radar	[]	16487	0.01	164
Adversarial	Attacker spoofs radar signals	Autonomous Vehicle / Sensors / Radar	[]	116779	0.10	11677
Adversarial	Denial of service attack against cellular network	Wide Area Networks / Cellular network	[]	16487	0.10	1648

Adversarial	Theft of a teleoperation computer	Teleoperation Facilities / Computer	[Physical security at teleoperation facilities]	16487	0.10	164
Adversarial	Theft of a vehicle component	Autonomous Vehicle	[Physical security of vehicles]	16487	0.10	164
Non-adversarial	Bumps dislodge vehicle computers	Autonomous Vehicle / Computers	[]	116779	1.00	116779
Non-adversarial	Fire in the teleoperation facility	Autonomous Vehicle	[Redundant teleoperation facility]	116779	0.01	11
Non-adversarial	Fire in the vehicle	Autonomous Vehicle	[]	116779	0.10	11677
Non-adversarial	Hardware failure of a teleoperation system	Teleoperation Facilities	[Redundant teleoperation facility]	116779	0.10	116
Non-adversarial	Hardware failure of a vehicle system	Autonomous Vehicle	[]	116779	0.10	11677
Non-adversarial	Rain/snow/fog disrupts sensors	Autonomous Vehicle / Sensors	[]	116779	0.10	11677
Non-adversarial	Vehicle crash dislodges sensors	Autonomous Vehicle / Sensors	[]	116779	1.00	116779
Non-adversarial	Weather impacts cellular network	Wide Area Networks / Cellular network	[]	16487	200.00	3297442

Table 7: Results

Figure 15 shows the annual expected loss for this autonomous vehicle system given no controls (upper blue line) and optimal controls (lower green line). The no controls situation is determined by running the Monte Carlo simulation described in section 2.2.5 with no controls implemented. The optimal controls situation is determined by running the Monte Carlo simulation with optimal controls enabled. Determining optimal controls is discussed below.

This figure shows the probability of a given loss in a given year and the impact security controls can have on reducing losses. The annual expected loss does not include the cost of controls. This figure aids the control selection process by showing the best and worst case scenarios. For example, given a set of controls, it is possible to determine the probability of a loss over a certain amount. If this loss is acceptable, the controls can be implemented and the system can be put in service. If this loss is unacceptable, additional controls (or other methods of managing risk) are necessary. Even in the situation of optimal controls, a considerable loss is expected each year. This is because there are several risks such as "Weather impacts cellular network" that remain unmitigated.

3.2 Optimization of Controls

The controls used in this assessment (Table 8) represent a sample of possible controls that could protect an autonomous vehicle system. A real such system would likely have hundreds, if not thousands, of controls. The costs and reduction factors are samples used to study how the risk assessment framework responds to differences.

The results of optimizing security controls are shown in Table 8 and Figure 16. The table show the cost and reduction factor as well as if the control is implemented in the optimal system. More effective controls that mitigate multiple risks (such as "Autonomous system protections for remote operations") are much more valuable than targeted controls that only protect against specific threat events (such as "Anti-malware software"). Additionally, adding a safety driver control at the cost of an average professional proves far more expensive than equivalent security controls for reducing risk. Interestingly, a sensitivity analysis performed by changing the costs of controls according to a log normal distribution while continually optimizing these controls always resulted in the same set of optimal controls. This is probably because the differences in

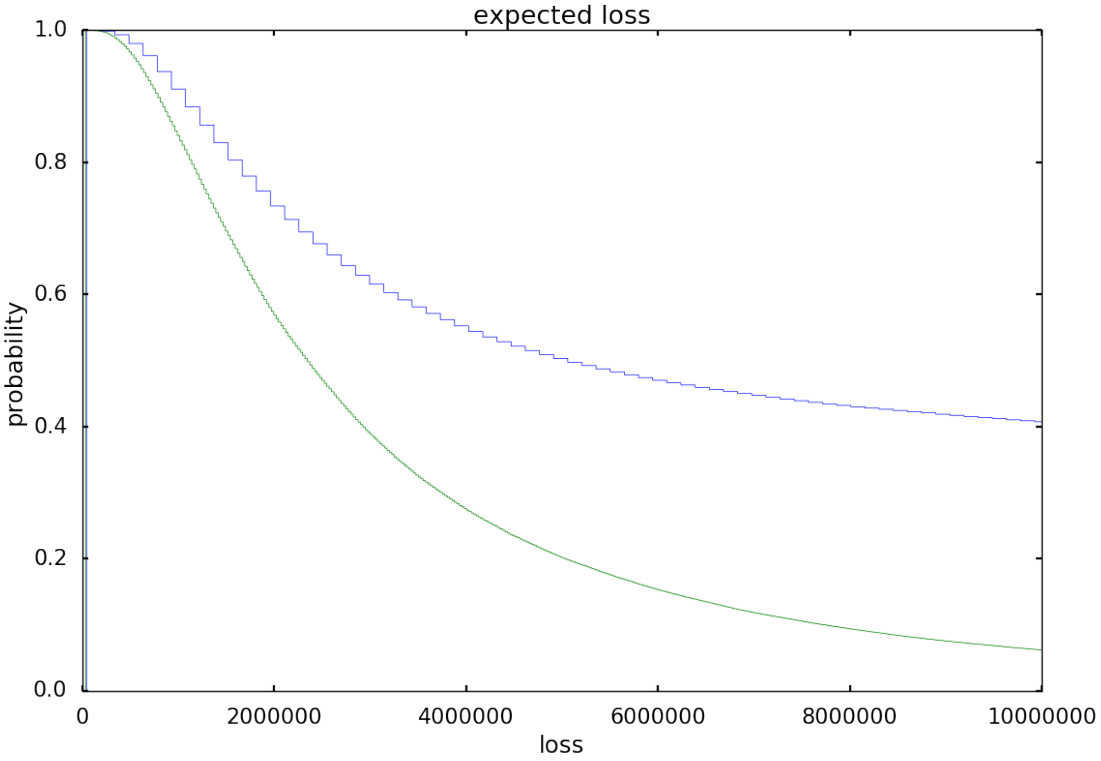

Figure 15: Annual expected loss ($/year); upper blue line is with no controls, lower green line is with optimal controls

costs and reduction factors are relatively large between different controls.

Figure 16 shows one point representing the residual risk and cost for each control combinations. For example, when few controls are implemented as on the left side of the graph, the control cost is low, but the residual risk is high. As the control cost increases, the residual risk tends to increase. The point of optimal controls show in Table 8, defined as the point closest to the origin, is colored red. It is also interesting to see how control combinations cluster and follow patterns around the graph because enabling or disabling one control shifts the group.

	Cost ($)	Reduction factor	Implemented
Vulnerability management program	100000	0.01	False
Anti-malware software	10000	0.10	False
Two-factor authentication	50000	0.10	True
Disable WiFi	1000	0.00	True
Physical security at teleoperation facilities	10000	0.10	False
Physical security of vehicles	100000	0.10	True
Redundant teleoperation facility	100000	0.01	False
Autonomous system protections for remote operations	500000	0.01	True
Background checks for drivers	10000	0.01	True

Table 8: Optimal controls

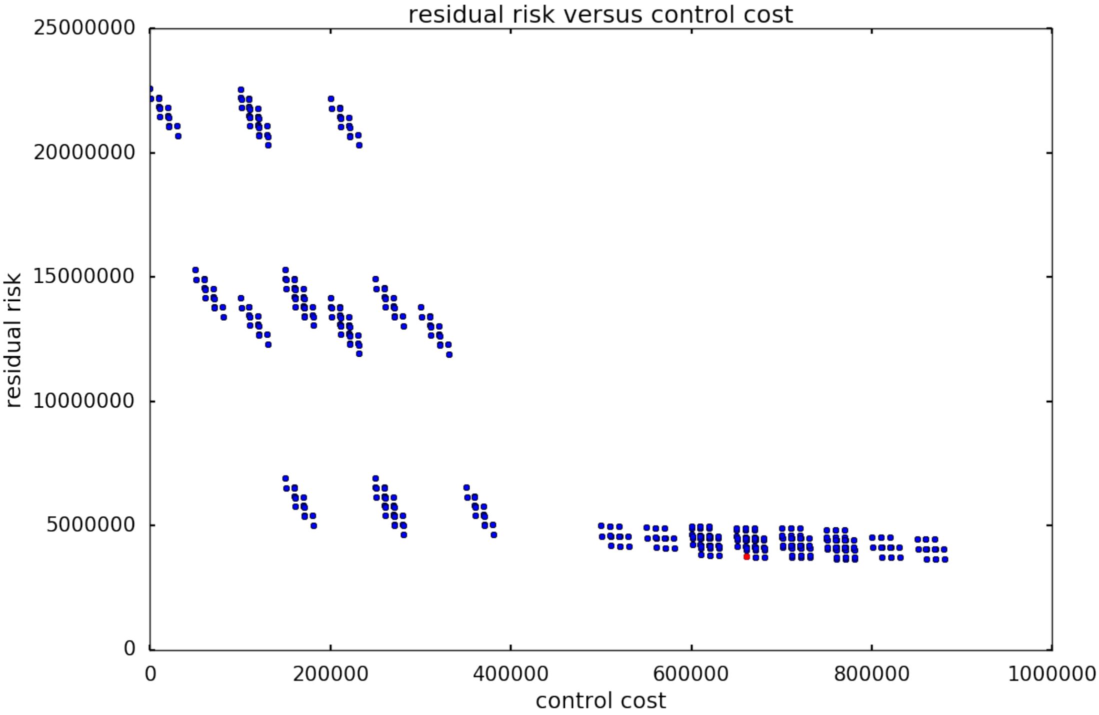

Figure 16: Residual risk versus control cost ($); optimal controls are red

4 Discussion

What they had in common was mainly love of excellence and programming. They wanted to make their programs that they used be as good as they could. They also wanted to make them do neat things. They wanted to be able to do something in a more exciting way than anyone believed possible and show "Look how wonderful this is. I bet you didn't believe this could be done."
– Richard Stallman (1985)

Most of the attacks mentioned in the introduction target the lateral (steering) and longitudinal (acceleration and braking) control systems that can obviously cause a vehicle to crash. However, other attacks against headlights could compromise computer vision systems and also cause vehicles to crash or attacks against entertainment systems could damage passengers' hearing. Specialized vehicles including buses, taxis, and trucks also have unique attack vectors that must be considered for their threat models. For examples, taxis often have a tablet for passengers to interact with, and trucks do not have passengers.

While this thesis focused on autonomous road vehicles, several other autonomous vehicle types are currently under development or in production including unmanned aerial vehicles, autonomous ships, and small autonomous delivery robots. These vehicles bear similar risks and consequently deserve study.

Similar to when other industries transitioned to electronic records, the legal structure around autonomous vehicles is not well defined. For example, California Penal Code 502 defines an injury as "any alteration, deletion, damage, or destruction of a computer system, computer network, computer program, or data caused by the access, or the denial of access to legitimate users of a computer system, network, or program" (Waldron, 2015). However, this ignores cases where an attacker uses a computer to cause physical injury to property or persons. Likewise, the issue who is liable for the damage caused by a hacked autonomous vehicle has not been determined. The United States Congress appears to be aware of autonomous vehicle security, but has failed to pass legislation addressing the issue (Latta, 2017; Thune, 2017).

In addition to legal regulations, autonomous vehicle industry groups are still in their infancy in regards to security. In 2014, the Alliance of Automobile Manufacturers, Inc. and the Association of Global Automakers, Inc. formed the Automotive Information Sharing & Analysis

Center (Auto-ISAC) upon recommendation of the National Highway Traffic Safety Administration (Bainwol and Bozzella, 2014; Beuse, 2014). The purpose of this organization, similar to other ISAC organizations in other industries, is to develop and share security information such as their Automotive Cybersecurity Best Practices (Automotive Information Sharing and Analysis Center (Auto-ISAC)). To date they have released best practices guides for Incident Response and Third Party Collaboration and Engagement.

There are other automotive security industry groups predating modern autonomous vehicles. The Motor Industry Software Reliability Association (MISRA) released versions of C and C++ designed for automotive use. Also, formal methods provide a way to validate safety in safety-critical software. However, it appears many vehicle manufacturers neglect to use these methods and often blame drivers for what may be software faults (Koopman, 2018). It will be interesting to see how this behavior changes when drivers are not be there to blame. While this thesis limited the focus to security events, vehicle safety and security are fundamentally connected. Automotive safety practices such as the failure mode and effects analysis (FMEA) are similar to security risk management. Autonomous vehicle systems should explore combining safety and security efforts when practicable.

Another issue is support for security updates when new vulnerabilities are discovered. GM took 5 years to fix a security flaw in their 2009 Chevrolet Impala (Greenberg, 2015). Some manufacturers, including GM, have technology to provide over-the-air updates (Shavit et al., 2007; Dakroub and Cadena, 2014; Lewis, 2016; Alrabady et al., 2011), but these mechanisms also present a security risk themselves. Also, for how long do manufacturers have to provide security updates? What about software written by third parties? And, do owners have any rights to write their own security software?

Vehicle systems have unique characteristics that must be considered during the risk management process. The Controller Area Network (CAN bus) is an in-vehicle communication network created by BOSCH and first implemented in the 1988 BMW 8-Series (International Organization for Standardization, 2015). It allows open communication between electronic control units (ECUs) and can be exploited via denial of service attacks, spoofed messages, and sniffed messages (Kleberger et al., 2011; Miller and Valasek, 2014, 2016a,b; Smith, 2016). Several researchers have proposed intrusion detections systems (IDS) (Larson et al., 2008; Müter and Asaj, 2011; Boudguiga et al., 2016; Song et al., 2016; Kang and Kang, 2016; Cho and Shin, 2016; Daxin Tian et al., 2018; Spicer, 2018), or encryption systems (Bruton, 2014; Wei et al., 2016; Mukund and Shamrao, 2015) for the CAN bus. Similarly, some autonomous vehicles use the Robot Operating System (ROS) which bears many similarities to CAN: ROS is an open network and any node (analogous to an ECU) can openly communicate with other nodes. Several projects aim to add security features to ROS such as authentication and encryption (Dieber et al., 2017; Breiling et al., 2017).

NIST Special Publication 800-30 notes that quantitative assessments often require more time and effort than qualitative assessments (National Institute of Standards and Technology, 2012). However, the results shown in this thesis would not be possible within a qualitative framework. One area of quantitative assessments that could be greatly improved is the determination of likelihoods and impacts. Further research into methods for robustly gathering this data is welcome. For now, most of this data is based on educated guesses of experienced professionals. Also, because autonomous vehicles are relatively new and large-scale deployments do not yet exist, much of this data is simply unavailable. In the era of big data, rare events prove problematic for quantitive risk management frameworks.

5 Conclusion

No fim, tudo dá certo. Se não deu, ainda não chegou ao fim.

(In the end, everything will be ok. If it's not ok, it's not yet the end.) – Fernando Sabino

Autonomous vehicles have arrived. Qualitative risk management frameworks allow the companies architecting these autonomous vehicle systems to measure their security risks, choose the most appropriate security controls, and assure a level of safety. Existing information security approaches are well suited to securing the robotic systems that control autonomous vehicles. The transportation industry can learn from past security mistakes in other industries.

Traditional, qualitative frameworks have several shortcomings including an inability to mathematically sum, compare, and interpret risks. Quantitative frameworks, while requiring more data, help risk management programs overcome these limitations. Annual expected loss curves to show the probability of different loss levels given different security controls. Optimization techniques to select the best security controls for a system. Quantitative frameworks communicate risk clearly in monetary units instead of "high, medium, and low" risk levels.

Similar to the safety ratings of passenger cars, qualitative risk management provides a method for impartial evaluation of the security risks present in an autonomous vehicle system. Quantitative risk management methods can also be applied to other areas of transportation such as safety modeling, environmental impact modeling, traffic management, and cost-benefit analysis. Also, these techniques can be expanded to cover other modes including active transportation, freight, and private public partnerships.

As Heidi King, Deputy Administrator of the National Highway Traffic Safety Administration said "we should together build and support a cyber security risk management culture for the automotive industry that can serve as a role model to others. Our national automotive safety depends on it" (King, 2018).

6 Appendix I: MIL-STD-882E Risk Tables

Description	Severity	Mishap Result Criteria
Catastrophic	1	Could result in one or more of the following: death, permanent total disability, irreversible significant environmental impact, or monetary loss equal to or exceeding $10M.
Critical	2	Could result in one or more of the following: permanent partial disability, injuries or occupational illness that may result in hospitalization of at least three personnel, reversible significant environmental impact, or monetary loss equal to or exceeding $1M but less than $10M.
Marginal	3	Could result in one or more of the following: injury or occupational illness resulting in one or more lost work day(s), reversible moderate environmental impact, or monetary loss equal to or exceeding $100K but less than $1M.
Negligible	4	Could result in one or more of the following: injury or occupational illness not resulting in a lost work day, minimal environmental impact, or monetary loss less than $100K.

Table 9: MIL-STD-882E Severity Rankings (Department of Defense, 2012)

Description	Level	Specific Individual Item	Fleet or Inventory
Frequent	A	Likely to occur often in the life of an item.	Continuously experienced
Probable	B	Will occur several times in the life of an item.	Will occur frequently.
Occasional	C	Likely to occur sometime in the life of an item.	Will occur several times.
Remote	D	Unlikely, but possible to occur in the life of an item.	Unlikely, but can reasonably be expected to occur.
Improbable	E	So unlikely, it can be assumed occurrence may not be experienced in the life of an item.	Unlikely to occur, but possible.
Eliminated	F	Incapable of occurrence. This level is used when potential hazards are identified and later elimi nated.	Incapable of occurrence. This level is used when potential hazards are identified and later eliminated

Table 10: MIL-STD-882E Probability Levels (Department of Defense, 2012)

	Catastrophic (1)	Critical (2)	Marginal (3)	Negligible (4)
Frequent (A)	High	High	Serious	Medium
Probable (B)	High	High	Serious	Medium
Occasional (C)	High	Serious	Medium	Low
Remote (D)	Serious	Medium	Medium	Low
Improbable (E)	Medium	Medium	Medium	Low
Eliminated (F)	Eliminated	Eliminated	Eliminated	Eliminated

Table 11: MIL-STD-882E Risk Matrix (Department of Defense, 2012)

7 Appendix II: ASIL Risk Tables

	E1	E2	E3	E4
Duration	-	<1% of operating time	1-10% of operating time	>10% operating time
Frequency	Occur less than once a year	Situation that occurs a few times a year	Situation that occurs once a month	Situations that occur almost every drive
Examples	Driving downhill with engine off	Driving on unsecured steep slope	Slippery roads	Braking

Table 12: ASIL Probability/Exposure (International Organization for Standardization, 2011)

	S1	S2	S3
Description	Light and moderate injuries	Severe injuries, possibly life threatening, survival probable.	Life threatening injuries, survival uncertain, fatal injuries
Example	Collision with tree <20 kpm	Collision with tree 20-40 kpm	Collision with tree >40 kpm

Table 13: ASIL Severity (International Organization for Standardization, 2011)

	C1	C2	C3
Description	Simply controllable	Normally controllable	Difficult to control or uncontrollable
Definition	All drivers will be able to avoid it	90% of all drivers will be able to avoid it	10% of all drivers will be able to avoid it
Example	Starting a vehicle with locked steering	Stopping a vehicle in case of light failure	Loss of breaks

Table 14: ASIL Controllability (International Organization for Standardization, 2011)

8 Appendix III: OWASP Risk Rating Methodology

Likelihood / Impact	Low	Medium	High
High	Medium	High	Critical
Medium	Low	Medium	High
Low	Note	Low	Medium

Table 15: Overall risk severity (Open Web Application Security Project, 2016)

9 Appendix IV: Security Control Identifiers

- Access Control
- Audit and Accountability
- Awareness and Training
- Configuration Management
- Contingency Planning
- Identification and Authentication
- Incident Response
- Maintenance
- Media Protection
- Personnel Security
- Physical and Environmental Protection
- Planning
- Program Management
- Risk Assessment
- Security Assessment and Authorization
- System and Communications Protection
- System and Information Integrity
- System and Services Acquisition

Figure 17: Security Control Identifiers (National Institute of Standards and Technology, 2013)

10 Appendix V: CIS Controls

- Basic CIS Controls
  1. Inventory and Control of Hardware Assets
  2. Inventory and Control of Software Assets
  3. Continuous Vulnerability Assessment and Remediation
  4. Controlled Use of Administrative Privileges
  5. Secure Configuration for Hardware and Software on Mobile Devices, Laptops, Workstations, and Servers
  6. Maintenance, Monitoring, and Analysis of Audit Logs
- Foundational CIS Controls
  7. Email and Web Browser Protections
  8. Malware Defenses
  9. Limitations and Control of Network Ports, Protocols, and Services
  10. Data Recovery Capabilities
  11. Secure Configurations for Network Devices, such as Firewalls, Routers, and Switches
  12. Boundary Defense
  13. Data Protection
  14. Controlled Access Based on the Need to Know
  15. Wireless Access Control
  16. Account Monitoring and Control
- Organizational CIS Controls
  17. Implement a Security Awareness and Training Program
  18. Application Software Security
  19. Incident Response and Management
  20. Penetration Tests and Red Team Exercises

Figure 18: CIS Controls (Center for Internet Security, 2018)

11 Appendix VI: PCI Data Security Standard – High Level Overview

Control Objectives	PCI DSS Requirements
Build and Maintain a Secure Network and Systems	1. Install and maintain a firewall configuration to protect cardholder data
	2. Do not use vendor-supplied defaults for system passwords and other security parameters
Protect Cardholder Data	3. Protect stored cardholder data
	4. Encrypt transmission of cardholder data across open, public networks
Maintain a Vulnerability Management Program	5. Protect all systems against malware and regularly update anti-virus software or programs
	6. Develop and maintain secure systems and applications
Implement Strong Access Control Measures	7. Restrict access to cardholder data by business need to know
	8. Identify and authenticate access to system components
	9. Restrict physical access to cardholder data
Regularly Monitor and Test Networks	10. Track and monitor all access to network resources and cardholder data
	11. Regularly test security systems and processes
Maintain an Information Security Policy	12. Maintain a policy that addresses information security for all personnel

Table 16: PCI Data Security Standard – High Level Overview (PCI Security Standards Council, 2016)

12 Glossary

adverse impact an action that causes harm to an organization. 8

availability the ability of a system to function, uptime, one of the three pillars of information security. 5

confidentiality protecting information from theft, one of the three pillars of information security. 5

confidentiality, integrity, and availability the three pillars of information security. 5

defense in depth the use of multiple layers of security controls to protect a system while allowing for failures of individual controls. 9

integrity protection information from tampering, one of the three pillars of information security. 5

organizational risk the result of a threat source initiating a threat event which exploits a vulnerability which causes an adverse impact. 8

Risk Management an organizational process to measure and manage risk. 8

Risk Management Framework A framework defined in NIST Special Publication 800-37. 8

threat event an event which creates risk. 8

threat source an actor who creates risk. 8

vulnerability a weakness in a system which exposes the system to risk. 8

13 Bibliography

Harold Abelson, Gerald Jay Sussman, and Julie Sussman. *Structure and Interpretation of Computer Programs*. MIT Press, second edition, September 1996. ISBN 0-262-01077-1.

John Adams. Risk Management: It's not rocket science - it's much more complicated, March 2005.

Ansaf I. Alrabady, Howard J. Carver, and Salvatore G. Trupiano. Secure over-the-air modification of automotive vehicular options, April 2011.

William Reynolds Archer Jr. Health Insurance Portability and Accountability Act of 1996, 1996.

Isaac Asimov. Runaround. In *I, Robot*, page 40. Doubleday, New York, 1950. ISBN 0-385-42304-7.

Automotive Information Sharing and Analysis Center (Auto-ISAC). Automotive Cybersecurity Best Practices. https://www.automotiveisac.com/best-practices/.

David Bailey. RAIL: Risk Assessment Library, September 2018.

Mitch Bainwol and John Bozzella. Letter from Mitch Bainwol and John Bozzella to David J. Friedman, July 2014.

Benz & Co. in Mannheim. Fahrzeug mit Gasmotorenbetrieb., January 1886.

Peter L. Bernstein. *Against the Gods: The Remarkable Story of Risk*. Wiley, August 1998. ISBN 978-0-471-29563-1.

Nathaniel Beuse. Docket Submission of One Document Related to an Automotive Cybersecurity Initiative, October 2014.

Garrick Blalock, Vrinda Kadiyali, and Daniel H. Simon. The Impact of 9/11 on Road Fatalities: The Other Lives Lost to Terrorism. Technical Report ID 677549, Social Science Research Network, Rochester, NY, February 2005.

Lawrence Blincoe, Ted R. Miller, Eduard Zaloshnja, and Bruce A. Lawrence. The Economic and Societal Impact of Motor Vehicle Crashes, 2010 (Revised). NHTSA Technical Report DOT

HS 812 013, U.S. Department of Transportation, National Highway Traffic Safety Administration 1200 New Jersey Avenue SE. Washington, DC 20590, National Center for Statistics and Analysis, National Highway Traffic Safety Administration Washington, DC 20590, May 2015.

A. Boudguiga, W. Klaudel, A. Boulanger, and P. Chiron. A simple intrusion detection method for controller area network. In *2016 IEEE International Conference on Communications (ICC)*, pages 1–7, May 2016. doi: 10.1109/ICC.2016.7511098.

R. Braden. RFC 1122: Requirements for Internet Hosts - Communication Layers. https://tools.ietf.org/html/rfc1122, October 1989.

Benjamin Breiling, Bernhard Dieber, and Peter Schartner. Secure Communication for the Robot Operating System. Montreal, April 2017. doi: 10.1109/SYSCON.2017.7934755.

R. Brooks. A robust layered control system for a mobile robot. *IEEE Journal on Robotics and Automation*, 2(1):14–23, March 1986. ISSN 0882-4967. doi: 10.1109/JRA.1986.1087032.

Jennifer Ann Bruton. *Securing CAN Bus Communication: An Analysis of Cryptographic Approaches*. PhD thesis, National University of Ireland, Galway, Galway, August 2014.

Kay Bryant and John Campbell. User Behaviours Associated with Password Security and Management. *Australasian Journal of Information Systems*, 14(1), November 2006. ISSN 1449-8618, 1449-8618. doi: 10.3127/ajis.v14i1.9.

James L. Buckley. Family Educational Rights and Privacy Act of 1974 § 20 U.S.C. § 1232g, 1974.

Bundesamt für Sicherheit in der Informationstechnik. Elementary Threats. https://www.bsi.bund.de/EN/Topics/ITGrundschutz/Download/download_node.html, 2008.

Bureau of Labor Statistics. CPI-All Urban Consumers (Current Series). https://data.bls.gov/timeseries/CUSR0000SA0, 2018.

Robert Burns. To a Mouse. *On Turning up in Her Nest with the Plough*, November 1785.

Mark S. Carey, René M. Stulz, and National Bureau of Economic Research, editors. *The Risks of Financial Institutions*. A National Bureau of Economic Research conference report. University of Chicago Press, Chicago, 2006. ISBN 978-0-226-09285-0. OCLC: ocm65521611.

George Carlin. Carlin on Campus, April 1984.

Center for Internet Security. CIS Controls. https://www.sans.org/critical-security-controls/, 2018.

Kyong-Tak Cho and Kang G. Shin. Fingerprinting Electronic Control Units for Vehicle Intrusion Detection — USENIX. Austin, TX, August 2016. ISBN 978-1-931971-32-4.

Andy Christensen, Andrew Cunningham, Jerry Engelman, Charles Green, Charles Kawashima, Steve Kiger, Danil Prokhorov, Levasseur Tellis, Barbara Wendling, and Frank Barickman. Key Considerations in the Development of Driving Automation Systems. 2015.

Arthur C. Clarke. Superiority. *The Magazine of Fantasy & Science Fiction*, August 1951.

Cloud Security Alliance Top Threats Working Group. The Treacherous 12 Cloud Computing Top Threats in 2016, February 2016.

Carter Cromwell. Wooden Preaches Preparation. *Lubbock Avalanche-Journal*, page D3, January 1977.

Husein Dakroub and Robert Cadena. Analysis of Software Update in Connected Vehicles. *SAE Int. J. Passeng. Cars – Electron. Electr. Syst.*, 7:411–417, April 2014. doi: 10.4271/2014-01-0256.

Joshua Davis. Say Hello to Stanley. *WIRED*, January 2006.

Thomas Milburn Davis III. Federal Information Security Management Act of 2002 § 44 U.S.C. § 3541, 2002.

Daxin Tian, Yuzhou Li, Yunpeng Wang, Xuting Duan, Congyu Wang, Wenyang Wang, Rong Hui, and Peng Guo. An Intrusion Detection System Based on Machine Learning for CAN-Bus — SpringerLink. January 2018. ISBN 978-3-319-74176-5. doi: https://doi.org/10.1007/978-3-319-74176-5_25.

Defense Advanced Research Projects Agency. *DARPA Grand Challenge 2005 Rules*. Defense Advanced Research Projects Agency, 3701 North Fairfax Drive Arlington, VA 22203-1714, August 2004.

Department of Defense. *CSC-STD-001-83: Department of Defense Trusted Computer System Evaluation Criteria*. August 1983.

Department of Defense. MIL-STD-882E - System Safety, May 2012.

Department of Motor Vehicles. Driverless Testing and Public Use Rules for Autonomous Vehicles Approved, February 2018.

E.D. Dickmanns, R. Behringer, D. Dickmanns, T. Hildebrandt, M. Maurer, F. Thomanek, and J. Schiehlen. The seeing passenger car 'VaMoRs-P'. pages 68–73. IEEE, 1994. ISBN 978-0-7803-2135-9. doi: 10.1109/IVS.1994.639472.

Bernhard Dieber, Benjamin Breiling, Sebastian Taurer, Severin Kacianka, Stefan Rass, and Peter Schartner. Security for the Robot Operating System. *Robotics and Autonomous Systems*, 98: 192–203, December 2017. ISSN 0921-8890. doi: 10.1016/j.robot.2017.09.017.

Paul Embrechts, Valerie Chavez-Demoulin, and Johanna Neslehova. Lectures at the Federal Reserve Bank of Boston, September 2005.

European Parliament and Council. Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation) (Text with EEA relevance), May 2016.

Ivan Evtimov, Kevin Eykholt, Earlence Fernandes, Tadayoshi Kohno, Bo Li, Atul Prakash, Amir Rahmati, and Dawn Song. Robust Physical-World Attacks on Deep Learning Models. *arXiv:1707.08945 [cs]*, July 2017.

Nathaniel Fairfield and Joshua Seth Herbach. Remote assistance for an autonomous vehicle in low confidence situations, October 2016.

Federal Financial Institutions Examination Council (FFIEC). IT Examination Handbook.

FIRST.org, Inc. Common Vulnerability Scoring System v3.0. https://www.first.org/cvss/specification-document, 2015.

Ford Motor Company. A Matter of Trust: Ford's Approach to Developing Self-driving Vehicles, 2018.

Ian Goodfellow, Nicolas Papernot, Sandy Huang, Yan Duan, Pieter Abbelle, and Jack Clark. Attacking Machine Learning with Adversarial Examples, February 2017.

Adam Gordon. *Official (ISC)2 Guide to the CISSP CBK*. CRC Press, Boca Raton, FL, fourth edition, February 2015. ISBN 978-1-4822-6275-9.

William Philip Gramm. Gramm–Leach–Bliley Act § 15 U.S.C. Subchapter I, 1999.

Andy Greenberg. GM Took 5 Years to Fix a Full-Takeover Hack in Millions of OnStar Cars. *Wired*, September 2015. ISSN 1059-1028.

Garrett Held and Davit Baghdasaryan. Measuring End-to-End Security Engineering, 2017.

HITRUST Alliance. HITRUST CSF. https://hitrustalliance.net/hitrust-csf/, 2018.

J. P. Holbrook and J. K. Reynolds. RFC 1244: Site Security Handbook. https://tools.ietf.org/html/rfc1244, July 1991.

Douglas W. Hubbard and Richard Seiersen. *How to Measure Anything in Cybersecurity Risk*. July 2016. ISBN 978-1-119-08529-4.

International Federation of Accountants. International Standard for Assurance Engagements (ISAE) 3402 Assurance Reports on Controls at a Service Organization, June 2011.

International Organization for Standardization. ISO/IEC 7498 - Information technology - Open Systems Interconnection - Basic Reference Model. http://standards.iso.org/ittf/PubliclyAvailableStandards/s020269_ISO_IEC_7498-1_1994(E).zip, 1994.

International Organization for Standardization. ISO 26262 - Functional Safety for Road Vehicles. https://www.iso.org/standard/43464.html, 2011.

International Organization for Standardization. ISO/IEC 27000 - Information security management systems. https://www.iso.org/isoiec-27001-information-security.html, 2013.

International Organization for Standardization. ISO 11898 - Road vehicles - Controller area network (CAN). https://www.iso.org/standard/63648.html, 2015.

International Organization for Standardization. ISO/IEC 27005 - Information security risk management. https://www.iso.org/obp/ui/#iso:std:iso-iec:27005:ed-3:v1:en, 2018a.

International Organization for Standardization. ISO 31000 - Risk management. https://www.iso.org/iso-31000-risk-management.html, 2018b.

Olivia A Jackson. The Impact of the 9/11 Terrorist Attacks on the US Economy. March 2008.

HyunChul Joh and Yashwant K Malaiya. Defining and Assessing Quantitative Security Risk Measures Using Vulnerability Lifecycle and CVSS Metrics. 2017.

Min-Joo Kang and Je-Won Kang. Intrusion Detection System Using Deep Neural Network for In-Vehicle Network Security. *PLOS ONE*, 11(6):e0155781, June 2016. ISSN 1932-6203. doi: 10.1371/journal.pone.0155781.

Keen Security Lab of Tencent. Car Hacking Research: Remote Attack Tesla Motors, September 2016.

Keen Security Lab of Tencent. New Car Hacking Research: 2017, Remote Attack Tesla Motors Again, July 2017.

Heidi King. Remarks at SAE/NHTSA Cybersecurity Workshop. January 2018.

P. Kleberger, T. Olovsson, and E. Jonsson. Security aspects of the in-vehicle network in the connected car. In *2011 IEEE Intelligent Vehicles Symposium (IV)*, pages 528–533, June 2011. doi: 10.1109/IVS.2011.5940525.

Kara Kockelman, Paul Avery, Prateek Bansal, Stephen D. Boyles, Pavle Bujanovic, Tejas Choudhary, Lewis Clements, Gleb Domnenko, Dan Fagnant, John Helsel, Rebecca Hutchinson, Michael Levin, Jia Li, Tianxin Li, Lisa Loftus-Otway, Aqshems Nichols, Michele Simoni, and

Duncan Stewart. Implications of Connected and Automated Vehicles on the Safety and Operations of Roadway Networks: A Final Report. Technical Report FHWA/TX-16/0-6849-1, Center for Transportation Research, The University of Texas at Austin 1616 Guadalupe Street, Suite 4.202 Austin, TX 78701, August 2016.

Philip Koopman. Practical Experience Report: Automotive Safety Practices vs. Accepted Principles. page 9, 2018.

Kirsten Korosec, E.W. Niedermeyer, and Alex Roy. #82: Elliot Katz and Jordan Sanders of Phantom Auto, a.

Kirsten Korosec, E.W. Niedermeyer, and Alex Roy. #98: Stefan Seltz-Axmacher of Starsky Robotics, b.

K. Koscher, A. Czeskis, F. Roesner, S. Patel, T. Kohno, S. Checkoway, D. McCoy, B. Kantor, D. Anderson, H. Shacham, and S. Savage. Experimental Security Analysis of a Modern Automobile. In *2010 IEEE Symposium on Security and Privacy*, pages 447–462, May 2010. doi: 10.1109/SP.2010.34.

U. E. Larson, D. K. Nilsson, and E. Jonsson. An approach to specification-based attack detection for in-vehicle networks. In *2008 IEEE Intelligent Vehicles Symposium*, pages 220–225, June 2008. doi: 10.1109/IVS.2008.4621263.

Robert Latta. SELF DRIVE Act, July 2017.

Le Monde.fr with AFP. Le bilan de l'attentat de Nice porté à 86 morts. *Le Monde.fr*, August 2016. ISSN 1950-6244.

David LeBlanc. DREADful, August 2007.

Jesse Sol Levinson, Timothy David Kentley, Gabriel Thurston Sibley, Rachad Youssef Gamara, Ashutosh Gajanan Rege, and Gary Linscott. Teleoperation system and method for trajectory modification of autonomous vehicles, November 2016.

Derek Lane Lewis. Over-the-air vehicle systems updating and associate security protocols, October 2016.

Huang Lin and Yang Qing. GPS Spoofing: Low-cost GPS simulator, August 2015.

Todd Littman. Autonomous Vehicle Implementation Predictions: Implications for Transport Planning, March 2018.

Gilbert Lynn and Gaylen Moore. *Particular Passions: Grace Murray Hopper*. Women of Wisdom. Lynn Gilbert Inc., New York City, 1st edition, December 2012. ISBN 978-1-61979-403-0.

Tom Magliozzi. Car Talk.

Doug Markiewicz and Laura Raderman. Guidelines for Data Classification. http://www.cmu.edu/iso/governance/guidelines/data-classification.html, July 2015.

Charlie McCarthy and Kevin Harnett. National Institute of Standards and Technology (NIST) Cybersecurity Risk Management Framework Applied to Modern Vehicles. Technical Report DOT HS 812 073, Washington DC: National Highway Traffic Safety Administration, October 2014.

John McCumber. *Assessing and Managing Security Risk in IT Systems: A Structured Methodology*. CRC Press, Boca Raton, FL, 2004. ISBN 987-0-203-49042-6.

Cathy Anne McMorris Rodgers. Health Information Technology for Economic and Clinical Health Act, 2009.

Chris McNab. *Network Security Assessment*. O'Reilly Media, Inc., 1005 Gravenstein Highway North, Sebastopol, CA 95472, second edition edition, October 2007. ISBN 978-0-596-00611-2.

Alexander J McNeil, Rüdiger Frey, and Paul Embrechts. *Quantitative Risk Management : Concepts, Techniques and Tools*. Princeton University Press, Princeton, NJ, May 2015. ISBN 978-0-691-16627-8.

Chalie Miller and Chris Valasek. Advanced CAN Injection Techniques for Vehicle Networks, 2016a.

Charlie Miller and Chris Valasek. Adventures in Automotive Networks and Control Units, 2014.

Charlie Miller and Chris Valasek. Remote Exploitation of an Unaltered Passenger Vehicle, August 2015.

Charlie Miller and Chris Valasek. CAN Message Injection: OG Dynamite Edition, June 2016b.

Sutar Tejaswini Mukund and Pawar Sanjay Shamrao. Security for CAN (Controller Area Network) Bus in Vehicle Communication System. *International Journal for Scientific Research & Development*, 3(8), 2015.

M. Müter and N. Asaj. Entropy-based anomaly detection for in-vehicle networks. In *2011 IEEE Intelligent Vehicles Symposium (IV)*, pages 1110–1115, June 2011. doi: 10.1109/IVS.2011.5940552.

National Highway Traffic Safety Administration. Guidelines for the Safe Deployment and Operation of Automated Vehicle Safety Technologies, 2016.

National Institute of Standards and Technology. National Vulnerability Database. https://nvd.nist.gov/.

National Institute of Standards and Technology. Federal Information Processing Standard 200: Minimum Security Requirements for Federal Information and Information Systems. March 2006.

National Institute of Standards and Technology. NIST Special Publication 800-30 Revision 1: Guide for conducting risk assessments. Technical Report NIST SP 800-30r1, National Institute of Standards and Technology, Gaithersburg, MD, September 2012.

National Institute of Standards and Technology. NIST Special Publication 800-53 Revision 4: Security and Privacy Controls for Federal Information Systems and Organizations. Technical report, April 2013.

National Institute of Standards and Technology. NIST Special Publication 800-37: Risk Management Framework for Information Systems and Organizations: A System Life Cycle Approach for Security and Privacy. Technical Report NIST SP 800-37, September 2017.

Issac Newton. Letter from Sir Isaac Newton to Robert Hooke, 1675.

Nissan Motor Corporation. PRESS KIT: Nissan Intelligent Mobility at CES. https://newsroom.nissan-global.com/releases/press-kit-nissan-intelligent-mobility-at-ces?lang=en-US, January 2017.

North American Electric Reliability Corporation. Critical Infrastructure Protection standards. http://www.nerc.com/pa/Stand/Pages/CIPStandards.aspx.

Bunyo Okumura and Danil V. Prokhorov. Remote operation of autonomous vehicle in unexpected environment, November 2016.

Open Web Application Security Project. OWASP Risk Rating Methodology. https://www.owasp.org/index.php/OWASP_Risk_Rating_Methodology, May 2016.

Open Web Application Security Project. OWASP Top 10. https://www.owasp.org/index.php/Category:OWASP_Top_Ten_Project, 2017.

Michael Garver Oxley. Sarbanes–Oxley Act of 2002, 2002.

PCI Security Standards Council. Payment Card Industry (PCI) Data Security Standards. https://www.pcisecuritystandards.org/documents/PCI_DSS_v3-2.pdf, April 2016.

Stephen Peace. California Senate Bill 1386, September 2002.

J. Petit and S. E. Shladover. Potential Cyberattacks on Automated Vehicles. *IEEE Transactions on Intelligent Transportation Systems*, 16(2):546–556, April 2015. ISSN 1524-9050. doi: 10.1109/TITS.2014.2342271.

Paul Pols. *The Unified Kill Chain: Designing a Unified Kill Chain for Analyzing, Comparing and Defending against Cyber Attacks*. PhD thesis, Cyber Security Academy (CSA), December 2017.

Larry Ponemon. Will Privacy & Security Concerns Stall the Adoption of Autonomous Automobiles?, November 2017.

Kevin Richards, Ryan LaSalle, Matt Devost, Floris van den Dool, and Josh Kennedy-White. 2017 Cost of Cyber Crime Study, 2017.

Ian Rust. Expert mode for vehicles, July 2017.

SAE International. Taxonomy and Definitions for Terms Related to Driving Automation Systems for On-Road Motor Vehicles J3016_201806. Technical report, June 2018.

Bruce Schneier. The Internet of Things Will Turn Large-Scale Hacks into Real World Disasters. *Motherboard*, July 2016.

William Shakespeare. Henry V, 1599.

Moshe Shavit, Andy Gryc, and Radovan Miucic. Firmware Update Over The Air (FOTA) for Automotive Industry. In *SAE Technical Paper*. SAE International, August 2007. doi: 10.4271/2007-01-3523.

Hocheol Shin, Dohyun Kim, Yujin Kwon, and Yongdae Kim. Illusion and Dazzle: Adversarial Optical Channel Exploits Against Lidars for Automotive Applications. In Wieland Fischer and Naofumi Homma, editors, *Cryptographic Hardware and Embedded Systems – CHES 2017*, pages 445–467, Cham, 2017. Springer International Publishing. ISBN 978-3-319-66787-4.

Craig Smith. *The Car Hacker's Handbook: A Guide for the Penetration Tester*. No Starch Press, San Francisco, 2016. ISBN 978-1-59327-703-1.

Laura Smith-Spark. Deadly Barcelona attack is worst in a day of violence in Spain. *CNN*, August 2017.

H. M. Song, H. R. Kim, and H. K. Kim. Intrusion detection system based on the analysis of time intervals of CAN messages for in-vehicle network. In *2016 International Conference on Information Networking (ICOIN)*, pages 63–68, January 2016. doi: 10.1109/ICOIN.2016.7427089.

Kevin J Soo Hoo. How Much Is Enough? A Risk-Management Approach to Computer Security, June 2000.

Matthew William Spicer. *Intrusion Detection System for Electronic Communication Buses: A New Approach*. Thesis, Virginia Tech, January 2018.

Richard Stallman. Hackers: Wizards of the Electronic Age, 1985.

Stanford University IT. Risk Classifications. https://uit.stanford.edu/guide/riskclassifications.

Willie Sutton and Edward Linn. *Where the Money Was: The Memoirs of a Bank Robber*. Broadway Books, New York, New York, 2004.

Aaron Swartz. *Guerilla Open Access Manifesto*. Eremo, Italy, July 2008.

The Mentor. The Conscience of a Hacker. *Phrack Magazine*, One(7), January 1986.

The MITRE Corporation. Common Vulnerabilities and Exposures (CVE). http://cve.mitre.org/, 2018a.

The MITRE Corporation. Common Weakness Enumeration (CWE). http://cwe.mitre.org/, 2018b.

William Thomson, 1st Baron Kelvin. Electrical Units of Measurement, May 1883.

Sebastian Thrun, Mike Montemerlo, Hendrik Dahlkamp, David Stavens, Andrei Aron, James Diebel, Philip Fong, John Gale, Morgan Halpenny, Gabriel Hoffmann, Kenny Lau, Celia Oakley, Mark Palatucci, Vaughan Pratt, Pascal Stang, Sven Strohband, Cedric Dupont, Lars-Erik Jendrossek, Christian Koelen, Charles Markey, Carlo Rummel, Joe van Niekerk, Eric Jensen, Philippe Alessandrini, Gary Bradski, Bob Davies, Scott Ettinger, Adrian Kaehler, Ara Nefian, and Pamela Mahoney. Stanley: The Robot That Won the DARPA Grand Challenge. Springer Tracts in Advanced Robotics, pages 1–43. Springer, Berlin, Heidelberg, 2007. ISBN 978-3-540-73428-4 978-3-540-73429-1. doi: 10.1007/978-3-540-73429-1_1.

John Thune. AV START Act, November 2017.

Kartik Tiwari and Stefan Seltz-Axmacher. Vehicle control system and method of use, June 2018.

Trustwave. The 2017 Trustwave Global Security Report, 2017.

H. Vegge, F. M. Halvorsen, R. W. Nergård, M. G. Jaatun, and J. Jensen. Where Only Fools Dare to Tread: An Empirical Study on the Prevalence of Zero-Day Malware. In *2009 Fourth International Conference on Internet Monitoring and Protection*, pages 66–71, May 2009. doi: 10.1109/ICIMP.2009.19.

Verizon. 2017 Data Breach Investigations Report (DBIR), 2017.

Marie Waldron. California Assembly Bill 32 Computer crimes, August 2015.

Waymo. Waymo Safety Report: On The Road to Fully Self-Driving. https://waymo.com/safetyreport/, October 2017.

Zhuo Wei, Yanjiang Yang, and Tieyan Li. Authenticated CAN Communications Using Standardized Cryptographic Techniques. In *Information Security Practice and Experience*, Lecture Notes in Computer Science, pages 330–343. Springer, Cham, November 2016. ISBN 978-3-319-49150-9 978-3-319-49151-6. doi: 10.1007/978-3-319-49151-6_23.

Chen Yan, Wenyuan Xu, and Jianhao Liu. Can you trust autonomous vehicles: Contactless attacks against sensors of self-driving vehicle. *DEF CON*, 24, 2016.

Eduard Zaloshnja and Ted R Miller. Costs of large truck-involved crashes in the United States. *Accident Analysis & Prevention*, 36(5):801–808, September 2004. ISSN 0001-4575. doi: 10.1016/j.aap.2003.07.006.

14 Colophon

If I have seen further it is by standing on the shoulders of Giants. – Sir Isaac Newton (1675)

This thesis was possible because of the dedication and generosoty of open source developers throughout the world. This thesis was written with TeXShop (Richard Koch) using the LaTeX document preparation system (Leslie Lamport) which is based on TeX (Donald E. Knuth). Diagrams were drawn with PGF/Ti*k*Z (Till Tantau). References were managed with Zotero (Center for History and New Media at George Mason University) and exported to BibTeX (Oren Patashnik and Leslie Lamport). The original LaTeXand BibTeX files along with this compiled PDF are available at `https://davidabailey.com/thesis`.

The quantitative risk assessment tool was written with Jupyter Lab (Fernando Pérez) in the Python programming language (Guido van Rossum) and uses numpy (Travis Oliphant) which is based on Numeric (Jim Hugunin) which is based on Jim Fulton's matrix object, pandas (Wes McKinney), matplotlib (John Hunter), and scipy (Travis Oliphant, Pearu Peterson, and Eric Jones). The Python code is available at `https://davidabailey.com/thesis`.