THE PENNSYLVANIA STATE UNIVERSITY
SCHREYER HONORS COLLEGE

COLLEGE OF INFORMATION SCIENCES & TECHNOLOGY

THE CHALLENGES INVOLVING THE ETHICS AND SECURITY OF AUTONOMOUS MOTOR VEHICLES

CHELSEA MARTIN
Spring 2016

A thesis
submitted in partial fulfillment
of the requirements
for baccalaureate degrees
in Information Sciences & Technology and Management
with honors in Information Sciences & Technology

Reviewed and approved* by the following:

Steven Haynes
Senior Lecturer in Information Sciences & Technology
Thesis Supervisor, Honors Adviser

Linda Trevino
Professor of Management
Thesis Reader

* Signatures are on file in the Schreyer Honors College.

ABSTRACT

The purpose of my thesis is to highlight the controversy concerning self-driving cars. In particular, I am going to focus on the ethics surrounding autonomous cars and whether or not the benefits outweigh the negatives. If self-driving cars are found to be unethical, do we still create them for the benefits? Can autonomous cars be programmed to make the correct moral decisions? Additionally, there are numerous security issues involving certain car technologies. My rationale for conducting this research and writing this thesis revolves around the fact that this is an extremely relevant topic in today's society considering the prevalence self-driving cars are expected to have. In a world where technology is taking over most aspects of life, it is important to ask the question: just because we can, does it mean we should?

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

ACKNOWLEDGEMENTS

I would like to thank my parents for their undying support of all of my endeavors. I would also like to thank Professor Haynes and Professor Trevino for offering invaluable feedback and assistance throughout this process.

Chapter 1

Introduction

The automotive industry is one that is about to experience a rapid revolution. Currently, many car companies are working on and researching the best way to develop a self-driving vehicle. Autonomous technology is paving the way for the new frontier of automotive vehicles. In the near future, we may see cars that can not only parallel park themselves, but drive too. It is possible that we will not only revolutionize private cars, but public transportation as well.

Before self-driving cars become mainstream, numerous actions need to occur in order to ensure that we are safe when using this technology. We need to address the many downfalls to the self-driving technology that present problems. Regulators would need to put laws in place and society itself would need to become comfortable with the new experience. It is possible that people will not feel comfortable making the move of giving up control of their vehicles. Additionally, numerous ethical questions need to be answered regarding the programming of autonomous vehicles. However, if car companies, lawmakers, and insurance companies can all agree on the proper regulations and be successful in their goal of creating autonomous cars, then there will be many benefits obtained by all.

Chapter 2

Autonomous Vehicle Definition

An autonomous car can be referred to as a driverless car, robot car, self-driving car, or autonomous vehicle. Self-driving vehicles are those where operation of the vehicle occurs without direct driver input to control the steering, acceleration, and braking and are designed so that the driver is not expected to constantly monitor the roadway while operating in self-driving mode (Aldana, 2013). Alternatively, an autonomous vehicle is a motor vehicle that uses artificial intelligence, sensors, and global positioning system coordinates to drive itself without the active intervention of a human operator. Essentially, the vehicle is able to sense its environment by updating maps based on sensory input. To qualify as fully autonomous, a vehicle must be able to navigate without human intervention to a predetermined destination over roads that have not been adapted for its use (Rouse, 2011).

Autonomous cars are able to operate without human intervention because they have control systems that can analyze sensory data to distinguish between different cars on the road. This capability allows the car to plan a path to the desired destination. Autonomous cars can use a variety of technologies, such as sensors to avoid collisions and augmented reality to display information to drivers. Like cars that rely on drivers, they can be built with GPS sensing knowledge to help them with navigation ("What is", 2016).

While driverless car technology is relevant today, it did not always have many supporters. In 2004, MIT economists Frank Levy and Richard Murnane believed that robotic cars would not be able to handle the complexity involved with turning left in traffic (Karsten &

West, 2016). That same year, no team finished the Defense Advanced Research Projects Administration Challenge to navigate a desert racecourse with an autonomous car (Karsten & West, 2016). In 2005, the competition yielded a winner (Karsten & West, 2016). In 2009, Google began testing its first driverless cars (Karsten & West, 2016). For years now, automakers have been striving to add more and more automated safety features in cars. Some of these advancements include forward collision avoidance with automatic breaking, lane departure warnings, and adaptive headlights that pivot around curves in the road. It would be easy to assume that if those adaptations make driving safer for humans, then fully automated driving would be even safer. Because car companies are in competition with each other to create the newest and best safety features, it is likely that the first company to develop and mass-produce a driverless car would gain an advantage (Karsten & West, 2016).

Significant research on autonomous vehicles is underway in both the U.S. and Europe. Companies developing and testing driverless cars include Audi, BMW, Ford, Google, General Motors, Volkswagen, and Volvo.

Figure 1 BMW's Driverless Car Interior

Section 1 (Tesla)

Tesla Motors pushed an autopilot software update to its Model S line of electric sports cars in October 2015. Model S owners will have the option for an over-the-air version 7 update, which includes steering and lane changing on highways, automatic parallel parking, and upgraded side collision avoidance (Crothers, 2014). The autopilot hardware includes the ultrasonic sonar, which indicates where everything is within 16 feet of the car and a forward facing camera, which can see where the lanes are, where the cars are ahead of it, and can read speed signs. The hardware also includes a forward radar, which detects fast moving large objects and can see through fog, rain, snow, and dust (Crothers, 2014).

The CEO of Tesla, Elon Musk, stated in a conference call with the media October 14, 2015, “Autopilot is a profound experience for drivers. Although we have been testing it for over a year and gotten used to it, when owners try it out and see the car drive by itself, they're blown away." (Crothers, 2014). However, he has also warned that Autopilot is not synonymous with fully automated driving. Musk emphasized that Tesla's Autopilot is different from rivals (like Mercedes Benz and Audi) because the whole Tesla fleet operates as a network. When one car learns something, the whole fleet learns something. The network uploads data to the central server, where it can be collected, do system analysis, and then feed that back into the cars. Musk claims that other companies are not on that level. Other companies will not be able to have an autonomous driving system in which updates will happen regularly and cars will improve each week (Crothers, 2014).

Another critical element Tesla claims to have over its competitors is its high-precision GPS navigation. The level of detail they provide is an order of magnitude greater than current navigation systems. In terms of data volume, the maps are 100 times more complex than a standard navigation map (Crothers, 2014).

Although the Autopilot system uses twelve long-range ultrasonic sensors and fast processors, it can primarily only handle straight-ahead predictable highway driving. For example, one of the Model S users admitted to ignoring the car's warnings until the vehicle almost swerved into the other lane (Berman, 2015). Tesla has made it very clear that drivers still need to maintain control and responsibility of their vehicle while enjoying the convenience of the Autopilot. There is no correct way to alert drivers to take control back over the vehicle, which is something automakers are trying to define. In contrast with Tesla, Google aims for complete autonomy so that there is no question about whether or not the driver needs to be attentive (Berman, 2015).

Tesla informs each customer of Autopilot's functions through release notes that come with every update, an update to the owner's manual, and emails (Berman, 2015). However, people have a tendency to use technology in a way that was not its intended purpose (Berman, 2015). This can be extremely dangerous and sets a precedent for automakers to be completely clear when describing the purpose of their products.

Section 2 (Google)

Google has already made significant progress testing fleets of self-driving cars. Overall, their self-driving cars have logged over 140,000 miles (Thrun, 2010). They have been testing their self-driving cars since 2009. For testing, Google uses 22 Lexus SUVs and 33 other prototypes. Google tests most of their cars in Mountain View, California, while others are tested in Austin, Texas. While testing, Google always makes sure that a trained safety driver is behind the wheel in case cruise control needs to be disengaged and a trained software operator is in the passenger seat to monitor the software (Metz, 2016).

Google's cars work by first requiring that the driver set a destination. The car's software then calculates a route and the car starts moving toward its destination. A rotating, roof-mounted LIDAR (Light Detection and Ranging) sensor monitors a 60-meter range around the car and creates a dynamic 3-D map of the car's current environment. A sensor on the left rear wheel monitors sideways movement to detect the car's position relative to the 3-D map. Radar systems in the front and rear bumpers calculate distances to obstacles. The artificial intelligence software in the car is connected to all the sensors and has input from Google Street View and video cameras inside the car. Additionally, the artificial intelligence simulates human perceptual and decision-making processes and controls actions in driver-control systems such as steering and brakes. The car's software can also consult Google maps in order to gain an advance notice of landmarks, traffic signs, and lights. If the human "driver" would need to take control of the car for any reason, they would be able to by utilizing the override function (Rouse, 2011).

Figure 2 Google's Self-Driving Car Compnents

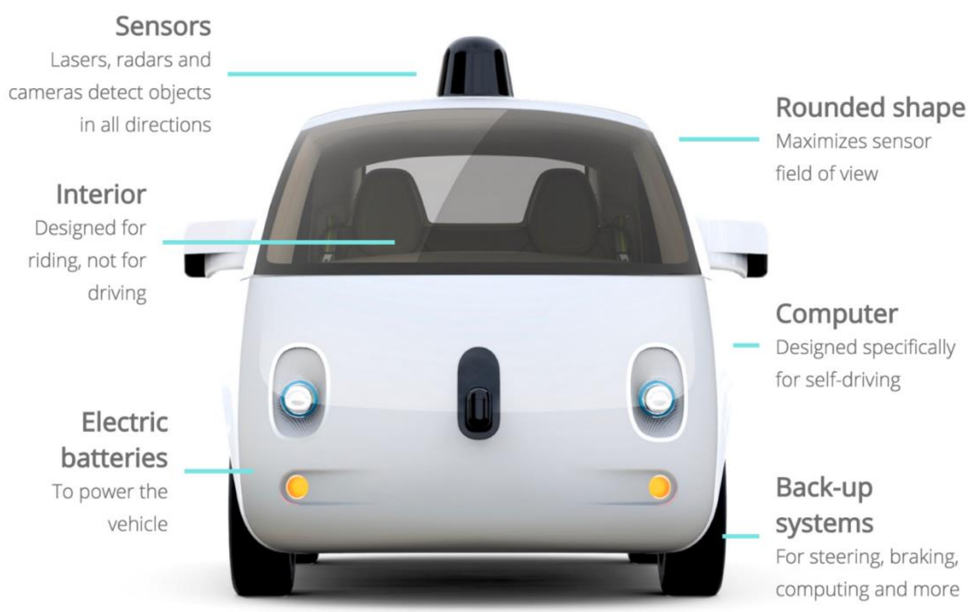

Figure 3 Google's Self-Driving Car

Chapter 3

Benefits of Producing Autonomous Vehicles

The mass-production and widespread use of driverless cars would offer many benefits to society. The most important benefit of all is that they would ultimately save countless lives. Proponents of systems based on driverless cars say they would eliminate accidents caused by driver error, which is currently the cause of almost all traffic accidents (Thrun, 2010). Advocates of the technology argue that robot drivers react faster than humans react, have 360-degree perception and do not get distracted, sleepy, or intoxicated (Markoff, 2010). Humans being distracted are one of the leading causes of accidents in the United States (“Top 20 Pros”, 2016).

According to the World Health Organization, more than 1.2 million lives are lost every year in road traffic accidents (Thrun, 2010). The New York Times echoes that sentiment by stating, “More than 37,000 people died in car accidents in the United States in 2008” (Markoff, 2010). Additionally, 81 percent of car crashes are the result of human error, proving that computers would take a lot of the danger out of the equation (“Top 20 Pros”, 2016). Google believes their technology has the potential to cut that number in half. Researchers estimate that driverless cars, by 2050, could reduce traffic fatalities by up to 90%. Using the number of fatalities in 2013 as a baseline, self-driving cars could save 29,447 lives a year (Lafrance, 2015). In the United States alone, that’s nearly 300,000 fatalities prevented over the course of a decade. According to the Centers for Disease Control, the life-saving estimates for autonomous cars are on par with the efficiency of modern vaccines, which save 42,000 lives for each U.S. birth cohort. On a global scale, World Health Organization reports 1.2 million traffic fatalities per year, meaning that driverless cars could save around 10 million lives per decade (Lafrance, 2015). All of these important statistics rely on the widespread acceptance of self-driving cars.

Other cost savings involving self-driving vehicles includes insurance costs and healthcare costs associated with accident recovery. The U.S. Department of Transportation assigns the value of $9.2 million to each human life. To lose a life due to a driving accident would essentially be losing $9.2 million ("Top 20 Pros", 2016).

In terms of time efficiency, the U.S. Department of Transportation estimates that people spend on average 52 minutes each working day commuting. Google wants people to be able to spend that time more productively. These systems would allow the human to engage in other activities while the car drives. This would result in a cost savings because according to Forbes, there is a cost savings associated with time ("Top 20 Pros", 2016).

Furthermore, the greater precision of an automatic system could improve traffic flow, dramatically increase highway capacity, and reduce or eliminate traffic jams. It is possible that certain manufacturers will keep all of their motor vehicles on the same network. This would allow the cars to operate more efficiently because they can then communicate with one another, allowing each car to identify traffic problems or road risks early on and then pass the message on to other vehicles. This is all due to the complicated algorithms that would determine the appropriate stopping distance, distance from another vehicle, and other data that decreases the chances of car accidents dramatically. This means that cars would be able to drive more safely while closer together. It could also mean that regulators could establish higher speed limits, further reducing driving time (Thrun, 2010).

Proponents are also confident that self-driving cars will transform sharing and reduce car usage, which should cut energy consumption and increase the number of people that can be transported. Because the robot cars would eventually be less likely to crash, manufacturers could build them lighter than regular cars, which would reduce fuel consumption. Another potential benefit would be the reduction of demands on city parking infrastructure (Keynes, 2015).

Many cities struggle with providing sufficient public transportation. Especially large cities oftentimes lack the appropriate infrastructure to support the needs of their residents. Driverless cars could assist with this problem. Moreover, disabled individuals who have to rely on public transportation or assistance from others to get around could utilize self-driving cars- which could allow them more freedom and enhanced mobility. The elderly would also be able to take advantage of those benefits (Keynes, 2015).

If driverless cars became the norm, then police officers could shift their focus from writing traffic tickets and handling accidents to managing other, more serious crimes. Other benefits of driverless cars include the fact that the car can drop a person off and look for a parking space without holding the person up or causing them to be late. There might also be shorter lines at the DMV, which is something everyone would appreciate. Capturing the aforementioned benefits provide a strong incentive for policymakers to anticipate and resolve the legal issues surrounding driverless cars (“Top 20 Pros”, 2016).

In summary…

Table 1 Key Benefits

• Could save around 10 million lives • Eliminate accidents caused by human error • Reduce or eliminate traffic jams • Reduce fuel consumption and carbon emissions • Increase people's productivity • Enhanced mobility for the disabled and elderly • Police could be free to handle more serious crimes

Chapter 4

Regulation Problems

Currently, there are not many regulations for autonomous vehicles (Lin, 2013). Bryant Walker Smith, of Stanford law, suggested that because self-driving cars are not illegal or prohibited, then they are legal in the United States (Lin, 2013). This follows the principle that an act is allowed until it is explicitly banned because we presume that individuals should have as much liberty as possible (Lin, 2013).

Previously, a human had to be in control of the car at all times. Bernard Lu, senior staff counsel for the California Department of Motor Vehicles, "There are dozens of laws pertaining to the driver of a vehicle, and they all presume to have a human being operating the vehicle" (Markoff, 2010, p. 4). Google researchers feel that their experimental cars should be legal because a human driver can override any error. However, on February 4, 2016, the National Highway Transportation Safety Administration (NHTSA) stated that computers could be considered legal drivers of vehicles. This news came after the U.S. Secretary of Transportation Anthony Foxx visited the Consumer Electronics Show in Las Vegas and the North American Auto Show in Detroit in January. Both industry events highlighted new autonomous features that technology companies and automakers plan to introduce into new prototypes and car models (Karsten & West, 2016).

Google's self-driving car unit, the one that submitted a proposed design for a self-driving car that has no need for a human driver, suggested the change in the definition of driver. The National Highway Transportation Safety Administration will interpret "driver" in the context of

Google's described motor vehicle design as referring to the self-driving system and not any of the vehicle occupants. They have agreed with Google that their car will not have a driver in the traditional sense that vehicles have had drivers since their inception. If the NHTSA names artificial intelligence as a reputable alternative to human-controlled vehicles, it could substantially streamline the process of putting autonomous vehicles on the road. If the car's computer is the driver for legal purposes, then it clears the way for Google or other automakers to design vehicle systems that communicate directly with the vehicle's artificial pilot (Shepardson & Leinert, 2016).

In response to Google's letter, the federal agency offered insight into the many obstacles and regulations that would also need adjusting in order to make commercial driverless cars a reality. Currently, some auto safety equipment requires that braking systems be activated by foot control (Shepardson & Leinert, 2016). If Google builds the cars with steering and breaking mechanisms, then it is possible that humans would try to override the self-driving system's decisions. This is a concern of Google. This means that federal regulations would have to be formally re-written in order to exclude the requirement of equipment like steering wheels and brake pedals. It could take months-or years to rewrite federal regulations governing the design, placement, and operation of vehicle controls. Google could consider applying for exemptions for certain regulations as long as they provide the proper supporting documents to the NHTSA (Shepardson & Leinert, 2016).

Because of the growth of this relevant technology and the efforts of both Silicon Valley and Detroit, the Obama administration announced that they would spend $4 billion on research and development for driverless car technology over the next decade (Karsten & West, 2016). This is significant because it indicates a greater focus at the federal level, which adds to many state and local government's already present focus. Because cars travel between city and state jurisdictions, creating a national strategy for driverless cars is crucial. The surprisingly quick pace of the development of driverless cars has put pressure on federal transportation officials. These officials will have to put rules in place in order to be ready for self-driving cars' mainstream use (Karsten & West, 2016).

Although federal regulators are starting to support the development of self-driving technology, self-driving vehicles are still not legal on most roads. In June 2011, Nevada, U.S. became the first jurisdiction in the world to allow driverless cars on public roadways (Rouse, 2011). As of February 16, 2016, only four states and the District of Columbia have passed laws allowing driverless car testing. California is working on drafting rules requiring the presence of steering wheels in self-driving cars along with a licensed driver (Karsten & West, 2016).

Moreover, laws would have to be put in place to prevent the misuse or abuse of autonomous cars. For example, if developers programmed the cars to drive too conservatively, then they may become a road hazard or cause other drivers to get angry. The drivers with no patience could take advantage of the crash-avoidance system of the self-driving cars by cutting in front of them, knowing that they will slow down or swerve to avoid an accident (Lin, 2013).

A study conducted by the UK's Automobile Association found that 65% of people liked driving too much to want an autonomous car. This implies that if driverless cars were to become the norm then the roads and laws of the roads would have to accommodate both autonomous and manually driven cars (Yeomans, 2014). Therefore, significant autonomous car production could cause problems with existing auto insurance and traffic controls used for human-controlled cars ("What is an Autonomous Car?", 2016).

Another example of an arising issue is the fact that federal rules necessitate alerts on dashboards if tire pressure runs low. NHTSA said they would have to create a test that informs the vehicle's computer of the problem. It is possible that humans in the vehicle should also be informed of car conditions like that (Shepardson & Leinert, 2016).

A culture of alcohol consumption could grow knowing that cars could drive without the assistance of humans. Although drunk driving would decrease, which is a very positive outcome, we do not know whether other alcohol related incidents would increase or decrease. This is something that policymakers and lawmakers would need to consider, along with the fact that to be truly safer, the car must be far more reliable than today's personal computers (for example), which crash occasionally and are frequently infected (Markoff, 2010).

Chapter 5

Ethics Issues

Section 1 (Liability Concerns)

Increasing autonomy of cars provokes theoretical considerations of who should be liable in the event of a crash caused by the car itself. Before self-driving cars can become widespread, carmakers must solve the near impossible ethical dilemma of algorithmic morality. Complex systems inherently have errors and bugs that self-driving cars are not immune to (Ethics, n.d.). There will be situations where a car knows it is about to crash, so it will have to plan how it crashes. There will be inconceivable scrutiny on the engineers who wrote the code to deal with the crash ("Google's Autonomous Vehicle", n.d.). People are unsure of whether or not the car should save the occupant or save someone else. For example, how should the car be programmed to act in the event of an unavoidable accident? Should it minimize the loss of life, even if it means sacrificing the occupants, or should it protect the occupants at all costs? Or should it choose randomly? If the software interprets a worn down sign, does the blame fall on the department of transportation for poorly maintained signage or the company that produced the self-driving software? Answers to these ethical questions are extremely important in order to determine how society will accept the self-driving cars. If the cars are programmed to sacrifice the owners, no one would buy that car ("Google's Autonomous Vehicle", n.d.).

Furthermore, would it be acceptable for an autonomous car to avoid a motorcycle by swerving into a wall if the probability of survival is greater for the passenger of the car than for the rider of the motorcycle? Should manufacturers program the car to make different decisions when children are present in the vehicle? If the manufacturer were to offer different versions of its moral algorithm, then would the buyer be to blame since they would be the one who chose the algorithm?

Presently, in the UK, the user of the car is primarily liable even if their actions did not cause the accident. If defective technology caused the accident, the user has to pursue this legally with the manufacturer (Yeomans, 2014). In addition, drivers are supposed to maintain awareness of their surroundings and car, even if they are not in control of the driving. Currently, the United States is fast to place the blame on car manufacturers.

In the future, it is possible that we will develop full autonomy and car users could be legally permitted to be distracted from driving. However, that would entail tremendous changes to present legislation. If laws are changed such that they no longer require any supervision in cars, a change in liability assignment may also be called for to reflect the fact that users are not expected to have any control over or awareness of the driving (Yeomans, 2014).

Section 2 (Known Accidents)

Since California granted permission to test driverless vehicles on public roads in September 2014, there have been six road traffic accidents. Most of the reported accidents occurred while the vehicles were in self-driving mode and were the fault of other drivers, according to the data released by the U.S. Department of Motor Vehicles (Donnelly, 2015). Five of the incidents involved Google self-driving cars, while the sixth featured a vehicle made with technology from parts provider Delphi. Those Google crashes were attributed to human error in the other cars hitting the Google cars. The car with the parts from Delphi was the only crash to result in a police investigation and none of the incidents resulted in serious injuries (Donnelly, 2015).

On February 14, 2016, Google added another traffic accident to the list. California's Department of Motor Vehicles claimed that one of Google's autonomous Lexus SUVs hit a public bus while trying to avoid sandbags on the road. There was a human present inside the vehicle who had the power of taking control of the wheel if they felt it were necessary. In this case, the person thought the bus would yield. In a statement released by Google, they have taken partial responsibility of the accident. Each vehicle involved in the crash was moving at a slow speed. Google's car was moving at two miles per hour, while the bus was moving at fifteen miles per hour. No one was injured in this accident, which took place at an intersection in Mountain View, the corner of El Camino Real and Castro Street (Metz, 2016).

Given the number of fatal traffic accidents that involve human error today, it could be considered unethical to introduce self-driving technology too slowly. If technology is available that could save many people, but may still kill due to its imperfection, then how quickly should car companies introduce the technology?

Section 3 (Public Opinion)

During the transition period when conventional and self-driving cars would share the road, we can hypothesize that safety may worsen for the conventional vehicles. Any high-profile crashes that may occur during the transition period could possibly set back the driverless car revolution a couple of years (Lafrance, 2015).

Jean-Francois Bonnefon at the Toulouse School of Economics and his colleagues sought to discover the public's opinion on self-driving cars ("Why Self-Driving Cars", 2015). After believing that the public's view of self-driving cars will influence their acceptance into society, they felt it was important to use experimental ethics to discover the public's opinion. They posed ethical dilemmas to several hundred workers on Amazon's Mechanical Turk and then gathered, analyzed, and reported the results ("Why Self-Driving Cars", 2015). One of the ethical dilemmas posed involved an accident where a self-driving car had a choice to either kill ten people crossing the road or only kill the car occupant by swerving away from the ten people and into a wall. If the goal is to minimize the loss of life, then it is obvious that killing one person is better than killing ten. However, knowing that self-driving cars are programmed to sacrifice their occupants will deter people from buying them. This will then increase the number of regular cars on the roads, which will cause more accidents and in turn kill more people. Some might argue that it is a losing situation ("Why Self-Driving Cars", 2015).

Bonnefon (2015) altered some of the details of the aforementioned dilemma, such as the actual number of pedestrians that could be saved, whether the driver or an on-board computer made the decision to swerve and whether the participants were asked to imagine themselves as the occupant or an anonymous person. Bonnefon found that people are in favor of cars that sacrifice the occupant to save other lives as long as they do not have to drive one themselves. They wished others to drive utilitarian autonomous vehicles more than they wanted to buy utilitarian autonomous vehicles ("Why Self-Driving Cars", 2015).

Section 4 (Theories in Normative Ethics)

The intent and final product of less automobile related deaths would be accepted in both a deontological and consequentialist framework because of the intent to save millions of lives and the result is the elimination of car accidents (“Google’s Autonomous Vehicle”, n.d). However, these philosophical frameworks could diverge in their agreement at a lower level of examination. If a crash is about to occur, humans will usually have a virtuous intention of avoiding the crash even if the crash is not avoided. Consequentialists would still likely favor the autonomous car at this level because the self-driving car will likely outperform the driver in avoiding the crash altogether.

Deontological ethics, or deontology, is an approach to ethics that focuses on the rightness or wrongness of the actions themselves as opposed to the rightness or the wrongness of the consequences of those actions. It judges the morality of an action based on the action’s adherence to the rules. Deontological ethics holds that at least some acts are morally obligatory regardless of their consequences for human welfare (Alexander, 2012).

Consequentialism is the class of ethical theories that states the consequences of one’s actions are the ultimate basis for any judgement about the rightness or the wrongness of those actions (Mizzoni, 2010). Therefore, to be morally correct in the eyes of a consequentialist, a person would have to behave in a way that will produce a good or positive outcome. In this case, it would be important for consequentialists to define “good outcome”. A good outcome for some could mean that the owner of the vehicle would survive in a crash. For others it might mean that the least amount of people are killed in a crash. The consequentialist might stipulate that the lives saved must be at least twice or triple the number of lives lost, even though that is an arbitrary amount with no guiding principle. This makes it difficult to defend with reason. Essentially, there will be a shift in the victim class as it moves from arbitrarily and unpredictable victims to a new set of unlucky victims. With these driverless cars, are we violating people's rights to not be killed? A deontologist would see the shift in the victim class as morally permissible given a non-violation of rights or duties, in addition to the consequentialist reasons based on numbers.

Utilitarianism is a theory in ethics that is a subset of consequentialism and holds that the best moral action is the one that maximizes the well-being of people as a whole (Alexander, 2012). Rule utilitarianism states that we must always choose the option that minimizes the suffering of all involved in any action (Huffstutler, 2015). Therefore, in the case of autonomous cars, someone who identifies with a utilitarian viewpoint would want the cars to be programmed to save as many people as possible. Act utilitarianism says we must consider each individual act as a separate subset action, meaning that each situation is a special case (Huffstutler, 2015). Based off act utilitarianism, it would be extremely difficult to program a car to handle each potential situation. Although it is beneficial to save the largest amount of people possible, the utilitarian does not take into consideration any exceptions. For example, it is more likely that a child would be injured in a car programmed to think with a utilitarianism mindset, because the utilitarian does not differentiate between people. Utilitarianism overall supports the idea of driverless cars over regular cars because of the fact that they are expected to save hundreds of thousands of lives. However, it is extremely likely that if the car were to be programmed with a utilitarian mindset, then the driver would be killed. If the driverless car were about to hit a bus full of people, it would have to swerve and possibly hit something else in order to save the bus full of people.

A deontologist might be against driverless cars because they do not know if the car could act with good will when avoiding a crash. It is unsure of whether or not the car will act with good intention or arbitrarily when it must choose between killing a pedestrian or the driver. If we programmed the car using a deontology viewpoint, then the driver would most likely be safe because it would be seen as wrong to actively kill the driver. A deontologist would expect the car to have a duty or obligation to follow the rules. In this case, rules can mean the laws set forth by regulators that cars must follow. If we program the car to always obey the law, then it is more likely that a deontologist would be happy with the vehicle. Deontology becomes an issue when regulators have to create new laws for the driverless cars to follow. This is because it is expected for there to be many grey areas during the transition from normal cars to autonomous cars. What rules does the car follow: the old ones, new ones, or a combination of both? How do we know that the rules regulators set forth are even moral?

Ethical dilemmas with robot cars are not just theoretical. Many new applied problems could arise, such as emergencies, abuse, theft, equipment failure, and manual overrides. In order to ensure the cars are programmed correctly, it will be important for us to choose one theory of ethics to use as a guiding principle. It is unfortunate that not all theories in normative ethics agree on this issue. Not even people who believe in the same theories will always agree with each other. This makes it extremely difficult when not only programming these vehicles, but creating new regulations for them. What viewpoint do we take when creating autonomous vehicles? Who is more correct in this issue, and which standpoint is better? Some may argue that there is no one who has the authority or high enough moral standing to answer those questions.

Section 5 (Programming Implications)

Chris Gerdes, a professor at Stanford University, leads a research lab that is experimenting with sophisticated hardware and software for automated driving (Knight, 2015). He is also exploring the ethical dilemmas surrounding self-driving vehicles with Patrick Lin, a professor of philosophy at California Polytechnic State University (Knight, 2015). Together they organized a workshop at Stanford that gave engineers and philosophers the chance to discuss and dissect the moral issues surrounding autonomous vehicles. The participants of the workshop implemented different ethical settings in the software that controls automated vehicles and then tested the code in various situations and real vehicles. For example, they “told” the car to prioritize avoiding humans over parked vehicles and not to swerve for squirrels (Knight, 2015).

Asking a car to make a decision is an ethical dilemma. If a child were to run into the road unexpectedly, then the car would have to make a decision. In humans’ eyes, the child has more value than a different object that could be the alternative for the car to hit, such as a van. If the car was programmed to hit the van instead of the child, then the life of the passenger of the self-driving car could be at risk.

Adriano Alessandrini, a researcher at the University de Roma La Sapienza in Italy, leads a project called CityMobil2. CityMobil2 tests automated transit vehicles in various Italian cities. The vehicles in this program are far simpler than the cars Google is developing in that they follow a route and brake if something gets in the way. Alessandrini claims that this eliminates the ethical problems because they do not decide to change lanes if they see something in their paths (Knight 2015).

Bryant Walker-Smith, an assistant professor at the University of South Carolina, studies the legal and social implications of self-driving vehicles. He says that making ethical decisions in automotive engineering is nothing new. He stated, “Ethics, philosophy, law: all of these assumptions underpin so many decisions. If you look at airbags, for example, inherent in that technology is the assumption that you’re going to save a lot of lives, and only kill a few” (Knight, 2015).

Should we trust robotic cars to share our road, just because they are programmed to obey the law and avoid crashes? If a small tree branch pokes out onto a highway and there is no incoming traffic, then as a human driver we would drift a little into the opposite lane and drive around it. In contrast, an automated car might come to a full stop, since it is programmed to obey traffic laws, including the ones that prohibit crossing a double-yellow line. Although the car would avoid hitting the tree branch, it could cause an accident if there are other cars behind it that contain human drivers (Lin, 2013).

Another popular topic of discussion regarding self-driving cars involves the possible tests they would have to pass in order for us to deem them safe for roadways. Is it enough for a robot car to pass a human driving test? Some people believe that it would be unfair to hold manufacturers to a higher standard than humans. However, there are important differences between humans and machines that could warrant a stricter test. For example, human drivers have the capability of exercising judgement in a wide range of dynamic situations, whereas self-driving cars do not (Lin, 2013). We presume that human drivers can act ethically and wisely in situations that do not appear in a standard 40-minute driving test. Moreover, good judgement can often sway people to act illegally. An example of this would be if a driver went faster than the speed limit in an emergency. This beckons the question, should robot cars never break the law in autonomous mode? Should developers program them to give a little leeway? If self-driving cars loyally follow laws and regulations, then they might refuse to drive if a headlight is broken, even during the day.

Section 6 (The Trolley Problem)

When laws cannot offer guidance, people's moral compass and principles are what guide them. The issue is when ethics does not provide the same answer as law. For example, jaywalking and speeding are illegal, but they do not always seem to be unethical when there is no traffic or during an emergency. A policy to arrest jaywalkers would be legal but also seen as harsh and unnecessary. Because the legal framework for self-driving vehicles does not exist yet, the opportunity exists for people to build a legal framework in which ethics contributes. Automakers and lawmakers would be remiss if they did not take advantage of the opportunity to consider ethics now rather than defensively react after a public backlash.

Philosophers Philippa Foot and Judith Jarvis Thomson proposed the Trolley Problem (Lin, 2013). The Trolley Problem describes a runaway trolley that is about to run over and kill five people standing on the tracks. Watching the scene from the outside, someone stands next to a switch that can direct the train to a sidetrack. There is one person standing on the sidetrack. The question is, should the person throw the switch? If so, they would kill the person on the sidetrack who would otherwise live if the person did not throw the switch. Even though one person would be killed, five people would be saved. By quickly analyzing the situation, is it easy to see that it is better that five people should live than only one person. When doing a more in-depth analysis, it is important to consider whether or not there is a moral distinction between "killing" and “letting die”. It seems worse to do something that causes someone to die than to allow someone to die because of events someone did not initiate nor had no responsibility to initiate (Lin, 2013).

Programmers and designers of automated cars do not have the luxury of being forgiven if a problem were to occur with their vehicle. Because they have more time to get the software correct, they bear more responsibility for bad outcomes. By contrast, human drivers have an easier time being forgiven for making an instinctive and split-second decision, even if it is a bad one. Therefore, developers and car manufacturers will have to consider these uncomfortable but possible situations. A discussion about ethics needs to take place to sort out issues such as rights, duties, conflicting values, and other factors.

It is possible that autonomous vehicles would not face situations such as the Trolley Problem very often. In fact, the number of times self-driving cars find themselves in difficult situations might be miniscule. However, programmers still need to instruct the cars on how to act for the entire range of both foreseeable and unforeseen scenarios. It is important to know whether an act was premeditated or committed reflexively without any deliberation.

Section 7 (External Conditions)

Many decisions the cars will have to make depend on environmental conditions, obstacles on and off the road, the size of the obstacles, second-order effects, and lives at risk in and outside the car. For example, if there is a baby present in the vehicle, the car software might give greater weight to protecting its occupants. If the obstacle is small, then it might be okay to hit, whereas a large obstacle could diminish the occupants’ chances of survival.

The owner of the car also influences ethics. For example, police officers have different virtues and duties. Therefore, their cars may also have different responsibilities than other self-driving vehicles, such as a shuttle bus or privately owned car. We would expect that a privately owned robot car would owe allegiance to its owner. Yet, a publicly owned automated vehicle might not have that obligation, which would change moral calculations.

Another obstacle car manufacturers would need to overcome deals with insurance companies. Even if the issues surrounding ethics and regulations are resolved, insurance companies might not be comfortable with the risk. Additionally, self-driving vehicles may be a threat to the insurance industry itself. If the cars avoid most or all accidents, then there would be a decline in the importance and presence of insurance agencies because there would be no risk to insure against. Alternatively, extremely large accidents could occur, as cars are networked together and vulnerable to wireless hacking. Therefore, the insurance industry would need to do what it can to protect itself while not preventing the benefits of technology to be utilized (Lin, 2013).

In summary…

Table 2 Key Ethics Issues

• Liability concerns • Known accidents • Public opinion • Theories in Normative Ethics • Programming implications • The Trolley Problem • External conditions

Chapter 6

Artificial Intelligence

Although the ethics of artificial intelligence applies specifically to robots and other artificially intelligent beings, we can apply it to autonomous vehicles as well. The topic of ethics involving artificial intelligence is divided into two subcategories: roboethics and machine ethics. Machine ethics concerns the moral behavior of artificial moral agents, while roboethics involves the moral behavior of humans as they design and construct the artificially intelligent beings (Veruggio, 2007).

We can apply roboethics to what is happening currently with autonomous cars because it refers to how humans create and use artificially intelligent beings (in this case, the vehicles). It considers how the vehicles may be used to harm humans as well as how they may be used to benefit humans. Autonomous carmakers specifically will have the ethical obligation to be transparent throughout their journey. Some people believe that artificial intelligence developers are representative of future humanity, which is why they should develop open source artificial intelligence if possible (Hibbard, n.d.).

Regarding machine ethics, it is important to develop artificially intelligent machines that behave morally. As previously stated, it is a large dilemma as to whether or not programmers can program self-driving cars to behave ethically and with a moral compass. As proven in Isaac Asimov's “I, Robot”, it is easy for logical rules to break down or create unanticipated behavior. It is arguable that no set of fixed laws can sufficiently foresee all possible circumstances (Asimov, 2008). This makes it difficult for programmers of self-driving cars. What circumstances should they focus on? What rules will they put into place for these new artificially intelligent vehicles? Who is qualified enough to answer these questions?

Another issue involving artificial intelligence is an issue that may not arise until very far into the future. Some technical experts predict that it is hypothetically possible for robots and computers to become self-sufficient and make their own decisions. If this were to happen, then it is possible that they could become smarter than humans. It is unknown the extent to which computers and robots could acquire any level of autonomy. It could be very dangerous for humans if computers were to become smarter than their own creators are (Markoff, 2009).

Therefore, a group of computer scientists is debating whether there should be a limit on research that might lead to a loss of human control over computer-based systems. They are concerned that further advances could create many consequences, such as the transformation of the work force by destroying jobs. If there were driverless cars “employed” by cities to drive around and pick people up, then there would be no need for taxi or Uber drivers anymore. Artificial intelligence researchers and roboticists who met at the Asilomar Conference Grounds on February 25, 2016, agreed that robots that can kill autonomously have either already been created or will be soon. Not only that, but they are worried that cybercriminals will be able to exploit artificial intelligence systems once they are fully developed (Markoff 2009).

Because of this, autonomy needs to be defined before implementing artificial intelligence into our vehicles. The amount of control that an intelligent system has over its actions plays a major factor in describing systems and in distinguishing them from each other. Different levels of autonomy reflect the different abilities of the machines as well as where and how they play a part in our daily lives (Hudson & Reeker, 2007). Will cars be on the same level of autonomy or on different levels of autonomy? If cars are not on the same level of autonomy, then there will need to be regulations for all levels of autonomy. This will be difficult and time-consuming to achieve.

To make sure the artificial intelligence is on the proper level of autonomy, the Machine Intelligence Research Institute suggested that artificial intelligence should be built with a friendly mindset. Friendly artificial intelligence means that the advances which are already occurring with artificial intelligence should also include an effort to make artificial intelligence intrinsically friendly and humane (Muehlhauser, 2012). However, it is difficult to make artificial intelligence's motivations human-friendly because of the complexity of human value systems.

A benefit of all of these questions is that we may become closer to understanding human ethics by motivating humans to address gaps in modern normative theory and by providing a platform for experimental investigation (Muehlhauser, 2012). Because the need for defining what is right and wrong is becoming more and more prevalent, the values of society will become more defined. Both roboethics and machine ethics will play a large part in the creation of the artificial intelligence present in autonomous vehicles.

Artificial intelligence does not only apply to the way in which the cars will be programmed to "think", but the process in which the cars will be programmed, or, "learn". Instead of using more traditional computer algorithms that would take a very long time to write, programmers are using deep learning to transform how cars think. Deep learning is based on a convolutional neural network, which is based on human biology. The process is very similar to how humans learn new skills because instead of trying "now if" statements in code where you tell something to perform an action if this happens, the programmers will show an image of a cat and keep showing that image until the car "understands". This is much like how children try again and keep performing an activity until they learn it ("This is How Cars", 2015).

Another area of artificial intelligence that may be relatable to autonomous vehicles is ambient intelligence. Ambient intelligence refers to electronic environments that are sensitive and responsive to the presence of people. Ambient intelligent devices work together to support people in carrying out their everyday life activities. They are able to do this in a natural way by using information and intelligence that is hidden in the network connecting the devices. As the devices grow smaller, more connected, and more integrated into our environment, they become less noticeable. Essentially, the technology disappears into the surrounding with only the user interface remaining visible. Ambient intelligence uses a human-centric computer interaction design that is context-aware, personalized, adaptive, and anticipatory (Aarts, 2007). It is important that autonomous cars contain all of those features. The more human-centric autonomous vehicles are, the safer we will be.

A new technology that relies on both artificial intelligence and ambient intelligence is Intelligent Transport Systems (ITS). ITS could become mainstream through the potential applications of artificial intelligence to infrastructure, the transport user, or the way in which these interact dynamically to deliver a transport service. The function of ITS is to improve decision-making by transport network controllers and users, which would improve the operation of the entire transport system (Miles & Walker, 2006). Although ITS might be better suited for "smart" public transportation, some of its ideas can help with private driverless vehicles.

Ambient intelligence aids with ITS because it ensures greater user-friendliness, more efficient services support, and user-empowerment. The future for ITS could mean efficient road traffic management, services for smart and driverless cars, co-operative systems, and automated highways (Miles & Walker, 2006). Each of those things would aid in the transition from normal to autonomous cars. The use of ambient intelligence specifically will increase the chances of driverless cars being adopted because it considers the user and guarantees that the car interface will be easy for people to use. This could also help reduce the number of accidents that would occur because people are confused by the driverless technology.

Chapter 7

Security Flaws

Besides the ethical and regulatory issues that need to be solved before autonomous vehicles become mainstream, there are also numerous security risks involved. Whoever owns the vehicle would presumably have the power to remotely take control of the car at any time. This would create an easy path for cyber-attackers to take advantage. Most computing devices created thus far have been hacked. It is important to measure how susceptible robot cars would be to hacking (Abel, 2016). Furthermore, if under attack, it is important to formulate contingency plans regarding what the car should do. Should the car speed away, alert the police, remain at the crime scene to preserve evidence, or defend itself? The issue of cyber-attacking vehicles is already a current problem most car manufacturers are facing because of new cars' use of in-car apps. Cyber criminals have more attack portals than ever before and could remotely access vehicle controls and systems.

It is imperative that car manufacturers do what they can to safeguard the personal information of their customers. If not, then people have to get used to the fact that their privacy rights will be disappearing. Car manufacturers may have good intentions, but most still have trouble making cybersecurity a priority in connected cars. Cost is one of the hurdles to prioritizing security. It would take a holistic, cross-disciplinary approach for the design and implementation of cybersecurity and its interconnection with technology to properly secure connected vehicles (Middleton, 2016).

From a roboethics standpoint, it is possible to argue that if an artificially intelligent program exists that can understand natural languages and speech, then with adequate processing power it could theoretically listen to every phone conversation and read every email sent out by the owner of the vehicle. If every car owner today owns a driverless vehicle featuring artificial intelligence, then that could mean almost every email or phone call in the world. The vehicle could then report to the program's operators exactly what the vehicle occupant said and who is saying it. This could allow governments and other entities to efficiently defeat opposition and attack their enemies ("Ethics of Artificial", 2016).

The Federal Bureau of Investigation and the Department of Transportation released a public service announcement to warn manufacturers and consumers about some of the dangers of connected automobiles. The public service announcement said, "Vulnerabilities may exist within a vehicle's wireless communication functions, within a mobile device-such as a cellular phone or tablet connected to the vehicle via USB, Bluetooth, or Wi-Fi- or within a third-party device connected through a vehicle diagnostic port" (Abel, 2016).

While not all car hacking incidents will result in safety risks, it is important that consumers take the appropriate steps in order to minimize their own risks. For example, consumers should make sure that their vehicle's software is always up to date, as well as be careful when making any modifications to the vehicle software. Consumers should maintain awareness and exercise discretion when connecting third-party devices to their cars. A recent survey of drivers and auto industry representatives found that it could take one to three years before connected car systems are capable of dealing with all the security concerns that are present (Abel, 2016).

Chapter 8

Discussion

Cars may be the most iconic technology in America due to the way they influence culture, the economy, and political landscapes. According to Patrick Lin (2013), an associate professor of philosophy at California Polytechnic State University:

> "They've made new forms of work possible and accelerated the pace of business, but they also waste our time in traffic. They rush to countless patients to hospitals and deliver basic supplies to rural areas, but also continue to kill more than 30,000 people a year in the U.S. alone. They bring families closer together, but also farther away at the same time. They're the reason we have suburbs, shopping malls, and fast-food restaurants, but also new environmental and social problems" (p. 8).

Therefore, whatever ends up happening with driverless vehicles will surely revolutionize everyone's way of life. The incoming prevalence of artificial intelligence itself will change life, as we know it. Personally, I believe there needs to be a limit on how far the advancements in technology go. Although it is highly fictionalized that robots will take over the earth, it would not surprise me if one day the human race were in decline because we became outsmarted by our own technologies. As I write this, I stopped and asked myself, can we really create things that are smarter than we are? Once I thought about it, I realized that we already have. Calculators, smart phones, and computers are all examples of these. Although they do not have the same reasoning, cognitive, and emotional intellect that we do, they are still smarter than we are by being able to perform many tasks that we cannot. If we create autonomous vehicles for mainstream use, then driving can become one more thing that humans are incapable of doing. As time goes on and more and more people stop driving, eventually there will not be anyone left who knows how to drive. This could be very dangerous in the event of a massive breakdown or attack on all of the vehicles on a certain server, for example. Elon Musk is very proud of the fact that Tesla's cars are all updated at the same time and on the same network (Crothers, 2014). I, however, think that is making it easier for someone to hack into the network and cause a massive disruption in all of the Tesla vehicles. Companies nowadays do not seem to do enough to prevent cyberattacks. Nonetheless, they are ready to patch up any holes in their systems and fight back against cyberattacks. What I think needs to happen before driverless cars become manufactured for private use is an increase in the number of security precautions taken.

Additionally, people (presumably, customers) need to be educated on the various issues that could occur. Most people do not understand the extent of cyberattacking and how easy it is for some people to do so without ever getting caught. If people are not educated on subjects like the growing use of artificial intelligence, then if something horrible were to happen, such as a terrorist hacking into a company's vehicles, then they will not be prepared and not know how to handle it. This would sufficiently ruin the further development of driverless vehicles. Some people already do not trust technology, so for them it will only take one small error for them to stick with regular cars forever. The larger the number of people who would keep using regular cars, the more accidents that will happen as a result of the two different types of cars being on the road. The more accidents there are; the number of people who support driverless cars will continue to diminish. It would become an endless cycle that needs to be stopped immediately in order to see the success of driverless cars proceed.

The success of driverless cars will also only be seen after the proper regulations have been put in place and the appropriate ethics questions have been answered. People will want to make sure that they and their vehicles are still following the law. Additionally, they will want the proper insurance coverage for this new type of vehicle. This means that people would have to place a lot of trust in policymakers and hope that they are keeping our safety and security a priority. One of my concerns is that manufacturers, like many companies today, will not put our safety before their moneymaking. If they were to keep consumers in the dark about the possible negative occurrences and market this new technology really well, then they could possibly make a large profit by risking our safety. Too many people do not know enough about artificial intelligence and its downfalls. Additionally, people would not necessarily be told how their cars are programmed. That would have to be a law that is put into place. Otherwise, people could be buying cars that are programmed to kill their owners in the event of an accident and not even be aware that they have done so. Informed consent should be required to operate or ride in an autonomous vehicle.

In regards to the ethical issues and dilemmas surrounding autonomous vehicles, I believe it will be very difficult to program cars to make ethical decisions. It is already difficult enough to try to figure out the correct ethical actions to take with humans. I am not even sure there is a correct moral authority available with which to raise all the questions that need answering. It is possible to program the cars to obey the laws as they are written; however, not following the laws can be more beneficial in certain scenarios. It would be useful for the car to be able to recognize when such dilemmas occur. That way, they can alert the passenger to make a decision or take control. However, this would not work in certain situations, such as if the occupant were intoxicated. To make decisions regarding the programming of the vehicles, I believe that people who are already working on developing the cars should have a large influence on it because they are the most familiar with the technology. Policymakers may not fully understand the capabilities of these vehicles, and therefore should not decide alone how the cars should be programmed to think and behave. Once the cars are programmed to perform basic functions, it may be necessary to compile a list of the most important objects the car should save in the event of an accident. For example, children could be put as a top priority to be saved no matter what (as long as the car is built with the ability to identify if a child is on the road or in the vehicle). As far as whether or not the car should kill its occupant if it means saving more lives, that depends on the type of person buying the car and the values of the company selling that car.

Concerning questions about liability, I believe the person who bought the car should be responsible if an accident were to occur. This is assuming that the programmer and car manufacturer did everything they could to create a vehicle that puts safety of the occupant first. Once the car is out of the creators' hands, it is not their responsibility to make sure it is used properly. It is the responsibility of lawmakers to enforce the proper use of autonomous subjects, and it is the responsibility of the users to follow those regulations.

The adoption of driverless cars is very much a cultural hurdle as a technological one. I predict the technology for these vehicles will be ready for widespread manufacturing by 2020. I predict the regulations and laws for these vehicles will follow around ten years after that. The widespread social and cultural acceptance of these vehicles could take anywhere between ten and fifty years. The transition itself will take anywhere between ten and fifty years, as well. The transition period alone may scare some individuals, since it will likely be more dangerous. Many people tend to want to be in control of their life and safety. Putting their lives into the hands of technology is something that will take a long time for people to be comfortable with. It is unclear how willing drivers will be to give up their control in favor of safety and convenience No one is going to welcome autonomous driving into the world until there is proof that it is much, much safer than having a human drive. This becomes problematic because the unintended effects of autonomous cars are difficult to predict, just as the great benefits are hard to ensure. Many drivers may take a slightly increased chance of accident in exchange for maintaining their ability to avoid accidents.

Chapter 9

Conclusion

A better future could be on the horizon. If developers, car manufacturers, and lawmakers take proper actions, then countless lives could be saved. That is the true benefit of autonomous vehicles. Saving lives is the main reason why companies such as Google and Tesla are working so hard to create autonomous vehicles. Not only that, but there are smaller benefits too, such as the increase in productivity by people, less traffic, more time for policemen, less fuel consumption, and increased mobility for the disabled or elderly. While all of those goals are quite valiant, there is a chance that many things could go wrong. Several questions have yet to be answered. In the end, it all comes down to efficiency and how ready we are for certain values to be pursued on our behalf. With innovation comes a new form of wastefulness elsewhere, a wastefulness that will have to be addressed.

The widespread adaptation of autonomous vehicles into society would be a colossal change in everyone's lives. It would not only be a technological advancement, but a catalyst for changes in our economy, society, ecology, and possibly even biology. An advancement this impactful needs to be handled with the utmost care, authority, and intelligence. Either way, change is inescapable. Only time will tell what the future holds for autonomous vehicles.

BIBLIOGRAPHY

Aarts, E., Harwig, R., & Schuurmans, M. (2007). *Chapter "Ambient Intelligence" in The Invisible Future: The Seamless Integration Of Technology Into Everyday Life*. McGraw-Hill Companies.

Abel, R. (2016, March 18). FBI, DoT release PSA on dangers of connected cars. Retrieved March 21, 2016, from http://www.scmagazine.com/fbi-dot-psa-motor-vehicles-increasingly-vulnerable-to-remote-exploit/article/484178/?utm_source=feedburner

Aldana, K. (2013, May 30). U.S. Department of Transportation Releases Policy on Automated Vehicle Development. Retrieved January 2, 2016, from http://www.nhtsa.gov/About NHTSA/Press Releases/U.S. Department of Transportation Releases Policy on Automated Vehicle Development

Alexander, L., & Moore, M. (2012, December 12). Deontological Ethics. Retrieved April 11, 2016, from http://plato.stanford.edu/entries/ethics-deontological/

Asimov, Isaac (2008). I, Robot. New York: Bantam.ISBN 0-553-38256-X.

Berman, B. (2015, October 21). Drivers Push Tesla's Autopilot Beyond Its Abilities. Retrieved January 2, 2016, from https://www.technologyreview.com/s/542651/drivers-push-teslas-autopilot-beyond-its-abilities/

Crothers, B. (2014, October 14). "With Autopilot, Tesla Takes Different Road Than Rivals Toward Self Driving Cars." Retrieved February 12, 2016, from http://www.forbes.com/sites/brookecrothers/2015/10/14/tesla-autopilot-version-7-with-fleet-as-network-drives-self-driving-future/#142a20786617

Donnelly, C. (2015, June 19). Californian state officials reveal driverless car crash numbers for first time. Retrieved January 10, 2016, from http://www.computerweekly.com/news/4500248472/Californian-state-officials-reveal-driverless-car-crash-numbers-for-first-time

Ethics of Artificial Intelligence. (n.d.). Retrieved March 23, 2016, from https://en.wikipedia.org/wiki/Ethics_of_artificial_intelligence

“Google’s Autonomous Vehicle”. (n.d.). Retrieved March 21, 2016, from http://googlesautonomousvehicle.weebly.com/ethics.html

Hibbard, B. (n.d.). *Open Source AI* (Unpublished master's thesis). University of Wisconsin. Retrieved March 30, 2016, from http://www.ssec.wisc.edu/~billh/g/hibbard_agi_workshop.pdf

Hudson, A. R., & Reeker, L. H. (2007, August 28). *Standardizing measurements of autonomy in the artificially intelligent*. Retrieved April 4, 2016, from http://dl.acm.org.ezaccess.libraries.psu.edu/citation.cfm?id=1660886&CFID=597511227&CFTOKEN=98097737

Huffstutler, J. (2015, June 15). Will your self-driving car be programmed to kill you if it means saving more strangers? Retrieved April 11, 2016, from https://www.sciencedaily.com/releases/2015/06/150615124719.htm

Karsten, J., & West, D. M. (2016, February 16). Driverless cars could arrive sooner than you think. Retrieved February 21, 2016, from http://www.brookings.edu/blogs/techtank/posts/2016/02/16-driverless-cars-closer-than-you-think

Keynes, M. (2015, February 11). How a driverless car will benefit you. Retrieved March 03, 2016, from http://www.telegraph.co.uk/news/uknews/road-and-rail-transport/11403807/How-a-driverless-car-will-benefit-you.html

Knight, W. (2015, July 29). How to Help Self-Driving Cars Make Ethical Decisions. Retrieved March 2, 2016, from https://www.technologyreview.com/s/539731/how-to-help-self-driving-cars-make-ethical-decisions/

Lafrance, A. (2015, September 29). Self-Driving Cars Could Save 300,000 Lives Per Decade in America. Retrieved February 2, 2016, from http://www.theatlantic.com/technology/archive/2015/09/self-driving-cars-could-save-300000-lives-per-decade-in-america/407956/

Lin, P. (2013, October 8). The Ethics of Autonomous Cars. Retrieved February 2, 2016, from http://www.theatlantic.com/technology/archive/2013/10/the-ethics-of-autonomous-cars/280360/

Markoff, J. (2009, July 25). Scientists Worry Machines May Outsmart Man. *The New York Times*. Retrieved April 03, 2016, from http://www.nytimes.com/2009/07/26/science/26robot.html?_r=1

Markoff, J. (2010, October 09). Google Cars Drive Themselves, in Traffic. *The New York Times*. Retrieved January 2, 2016, from http://www.nytimes.com/2010/10/10/science/10google.html?_r=0

Metz, R. (2016, February 29). Google's Self-Driving Car Probably Caused Its First Accident. Retrieved March 2, 2016, from https://www.technologyreview.com/s/600917/googles-self-driving-car-probably-caused-its-first-accident/

Middleton, N. (2016, March 22). Holistic, cross-disciplinary approach to cybersecurity needed, says security specialist. Retrieved March 29, 2016, from http://internationalfleetworld.co.uk/news/2016/Mar/Holistic-cross-disciplinary-approach-to-cybersecurity-needed-says-security-specialist/0444024303

Miles, J. C., & Walker, A. J. (2006). The potential application of artificial intelligence in transport [Abstract]. *IEE Proceedings - Intelligent Transport Systems*, *153*(3), 183-198.

Retrieved April 4, 2016, from http://ieeexplore.ieee.org.ezaccess.libraries.psu.edu/xpl/articleDetails.jsp?arnumber=1706434&queryText=artificial intelligence cars&newsearch=true

Mizzoni, J. (2010). *Ethics: The basics*. Chichester, West Sussex, U.K.: Wiley-Blackwell.

Muehlhauser, L., & Helm, L. (2012). Intelligence Explosion and Machine Ethics. *Machine Intelligence Research Institute*. Retrieved March 28, 2016, from https://intelligence.org/files/IE-ME.pdf.

Rouse, M. (2011, September). What is driverless car? - Definition from WhatIs.com. Retrieved January 13, 2016, from http://whatis.techtarget.com/definition/driverless-car

Shepardson, D., & Leinert, P. (2016, February 10). Exclusive: In boost to self-driving cars, U.S. tells Google computers can qualify as drivers. Retrieved February 15, 2016, from http://www.reuters.com/article/us-alphabet-autos-selfdriving-exclusive-idUSKCN0VJ00H

This Is How Cars Will Start Thinking For Themselves. (2015, December 1). Retrieved April 04, 2016, from http://www.huffingtonpost.co.uk/2016/01/12/self-driving-cars-artificial-intelligence-holds-the-key-to-an-autonomous-world_n_8926596.html

Thrun, S. (2010, October 9). What We're Driving At. Retrieved January 5, 2016, from https://googleblog.blogspot.com/2010/10/what-were-driving-at.html

Top 20 Pros and Cons Associated With Self-Driving Cars. (n.d.). Retrieved March 12, 2016, from http://www.autoinsurancecenter.com/top-20-pros-and-cons-associated-with-self-driving-cars.htm

What is an Autonomous Car? - Definition from Techopedia. (n.d.). Retrieved January 4, 2016, from https://www.techopedia.com/definition/30056/autonomous-car

Why Self-Driving Cars Must Be Programmed to Kill. (2015, October 22). Retrieved February 12, 2016, from https://www.technologyreview.com/s/542626/why-self-driving-cars-must-be-programmed-to-kill/

Veruggio, Gianmarco (2007). The Roboethics Roadmap (PDF). Scuola di Robotica: 2. Retrieved 28 March 2016.

Yeomans, G. (2014). Autonomous Vehicles. Retrieved January 15, 2016, from https://www.lloyds.com/~/media/lloyds/reports/emerging risk reports/autonomous vehicles final.pdf

.

ACADEMIC VITA

Academic Vita of Chelsea Martin
Clm5671@psu.edu

Education
Major(s) and Minor(s): Information Sciences & Technology, Management
Honors: Information Sciences & Technology

Thesis Title: The Challenges Involving the Morality and Security of Autonomous Motor Vehicles
Thesis Supervisor: Steven Haynes

General Electric Power & Water — Van Buren Township, MI — June 2015-August 2015

ERP Change Management Analyst

- Practice communications skills and teamwork in working with broader team members globally to understand and modify documentation
- Utilize new IT tools to better reflect the status of ERP configurations and business process design
- Lead efforts to provide a modern toolkit and capability for capturing and understanding GE PW's ERP business processes and components from a wing-to-wing perspective

White Building — University Park, PA — January 2013-Current

Customer Service

- Provide excellent customer service to professors and students in obtaining fitness memberships, lockers and fitness equipment location
- Facilitate the closing of the facility on weekends
- Report directly to the facilities manager

Rave Patio — Clarks Summit, PA — May 2013-August 2014

Sales & Exterior Design

- Assist and educate customer on high-end patio furniture
- Facilitate delivery and set-up of furniture at customer’s location
- Coordinate incoming and outgoing shipment deliveries
- Design and organize outdoor living spaces

American Eagle Outfitters — Wilkes-Barre, PA — May 2013-August 2013

Sales Associate

- Prioritized and managed daily sales objectives and customer service
- Increased revenue and customer database by educating the customer on the benefits of the credit card rewards program

- Prepared and organized marketing displays and merchandise to attract the customer

SKILLS: Basic java, Basic html, Microsoft Office, Team-building , Leadership
Basic php, Strong communication skills, Fine art & design,
Organizational skills

ACTIVITIES: Undergraduate Teaching Assistant for Dr. Wayne Geerling
2013 Fall Semester
Professor Ron Johnson's Innovation & Quality Team
2013 Fall Semester
Special Events Chair of Bee House, THON Organization
2015-2016
Technology Director-Special Interest Association
2015-2016
THON Committee- OPPerations
2013-2014
THON Committee- Rules & Regulations
2014-2015
Member of the National Society of Collegiate Scholars
2013-2016
Member of Women in IST Club
2014-2016
Member of Red Cell Analytics Lab
2013-2016