The Challenges Facing Autonomous Vehicles and The Progress in Addressing Them

Garrett Johnson

A Senior Thesis submitted in partial fulfillment
of the requirements for graduation
in the Honors Program
Liberty University
Spring 2019

Acceptance of Senior Honors Thesis

This Senior Honors Thesis is accepted in partial fulfillment of the requirements for graduation from the Honors Program of Liberty University.

____
Dr. Kyung Kyoon Bae, Ph.D
Thesis Chair

____
Dr. Young-Man Kim, Ph.D
Committee Member

____
Dr. David Schweitzer, Ph.D
Assistant Honors Director

____
Date

Abstract

Autonomous vehicles are an emerging technology that faces challenges, both technical and socioeconomic. This paper first addresses specific technical challenges, such as parsing visual data, communicating with other entities, and making decisions based on environmental knowledge. The technical challenges are to be addressed by the fields of image processing, Vehicle to Everything Communication (V2X), and decision-making systems. Non-technical challenges such as ethical decision making, social acceptance, and economic pushback are also discussed. Ethical decision making is discussed in the framework of deontology vs utilitarianism, while social acceptance of utilitarian autonomous vehicles is also investigated. Last, the likely economic impact is described.

The Challenges Facing Autonomous Vehicles and the Progress in Addressing Them

Autonomous vehicles are a technology under heavy research and development, for good reasons. Transportation is a major industry worldwide, and this technology has the potential to revolutionize it. Despite all this attention, there have been relatively few functional autonomous vehicles. The technology always appears to be only a few years away and has been like this since the turn of the century. This raises the question of why autonomous vehicles have taken so long to develop. This paper seeks to provide an answer to this, by describing the challenges that face autonomous vehicle development, both technical and socioeconomic. Along with stating the challenges, this paper also describes the ways in which they can be overcome.

The technical challenges facing autonomous vehicles are from a broad array of fields and are based primarily on computational tasks. The first challenge is sensing the environment around the autonomous vehicle, then parsing these data into a useful form for the vehicle's processor. The technical fields that are working to address this challenge are image processing, and Vehicle to Everything (V2X) Communication. Second, the possible additional benefits of V2X communication are also discussed. The third technical challenge to be discussed is developing a decision-making system for the vehicle and is covered in the section on decision making.

Decision making leads into the non-technical challenges such as ethical dilemmas, and socioeconomic factors. This paper describes the relevance of ethical systems for autonomous vehicle applications, with possible principles by which to program the artificial intelligence. Lastly, socioeconomic factors such as social acceptance and potential job losses are also discussed.

Image Processing

The first step to being able to effectively navigate an environment is to have a clear understanding of the surrounding area. For a person, this is not a difficult task: visual information is easily taken in and interpreted by the brain. The process of separating objects from the background, tracking their movement, and classifying them is effortless. Autonomous vehicles need to accomplish the same task, without the help of an innate object recognition system. From this challenge comes the application of image processing technology to autonomous vehicles. Image processing is how optical data, from either LIDAR arrays or cameras, is analyzed and understood. The applications of this technology in the field of autonomous vehicles are numerous. These include recognizing traffic signs, navigating road construction areas, and tracking pedestrians (Teichman and Thrun, 2011). This process can be split into two primary steps: segmentation and classification.

Segmentation

The first step to deciphering optical data comes in the form of segmentation. While a camera can take in a large stretch of road, there will only be a few items of relevance in any given image. To save processing power, it is advantageous to select only items in the image that differ from the expected background, for further analysis. This process is known as segmentation, which marks out unusual areas in an image for the purpose of a more thorough analysis. According to Harshitha & Manikandan (2017), this step is accomplished by creating a Histogram of Oriented Gradients, also known as an HOG. This algorithm takes an image frame, then divides it into a grid of cells. Within these cells, the algorithm determines the change in pixel brightness in the *x* and *y* directions for each pixel, then uses this to generate a gradient vector. These vectors have a specific magnitude and angle, decomposing into vertical and horizontal components, G_x and G_y. They refer to the relative change in pixel brightness between the above pixel and the pixel below. For example, if the pixel above had a value of 155, and the pixel below was 100, G_y would be -55. G_x is calculated in the same manner. These values are used to generate gradient vectors for each pixel.

This collection of gradient vectors is then arranged in a histogram, with bins based upon the angle. Vectors with higher magnitudes will contribute more to each bin, allowing the contents of a cell to be compressed and generalized. The end result will be a set of compressed cell gradients, creating a generalization of the high contrast areas. Instead of storing thousands of pixel gradients, this algorithm combines their overall effect into a set of cell gradients, which are stored in a string of values. This string of values has all the information gathered by breaking down the image into areas of high contrast, the object outlines. These gradients can later be used to classify specific objects, with the help of a classification algorithm.

Once the HOG data is extracted, it can be used to represent possible obstacles. Since this process is computationally demanding, it is optimal to identify the object only once, then use other algorithms to track it. The most common tracking algorithm is a particle filter (Nkosi, Hancke, & Santos, 2015). This filter is a nonlinear tracking algorithm, based on probabilistic target trajectories. It assigns a target a set of possible motion vectors, based upon a probability distribution. Then at a predetermined time, it measures where the target is, and uses this motion data to update the vector probability distribution. From this, it can develop an accurate model of where the object is most likely to be, to track the target, regardless of occlusion or nonlinear movement. To summarize, image processing starts with using a HOG algorithm on an image frame to identify areas of interest, and these areas are tracked using a particle filter. The next step is to identify useful information from these areas, which is the goal of classification.

Classification

The gradient data is stored as a string of vectors, which can be generalized as a set of points. The goal of the classification process is to take an unlabeled set of gradient data, then determine what it is most likely to represent. While there are multiple ways to accomplish this, according to Teichman and Thrun (2011), the most common algorithm is to use a Support Vector Machine (SVM). Imagine a graph with points. Some of these points will be a member of a certain class, and other points will be a member of other classes. While each class will vary, it will be possible to draw a line to separate the two, where objects on either side are most likely to be one specific class. A visualization is shown in Figure 1.

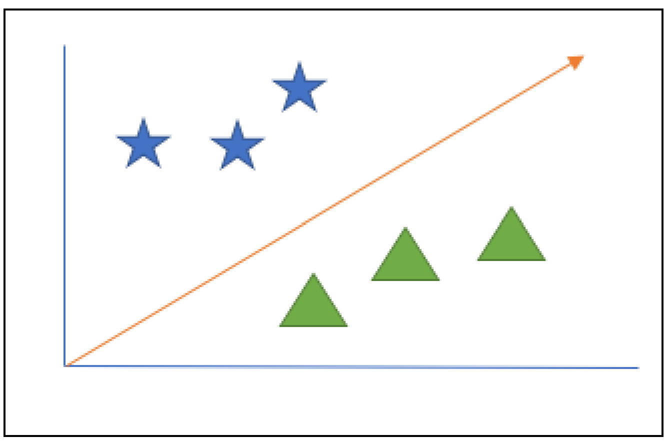

Figure 1: Linear Classification Example

SVMs take this basic principle, then expand on it (Yuan, Hao, Chen, & Wei, 2014). Instead of using only two dimensions of data per each point, they use all of the data points available from the HOG gradient. This creates a multidimensional space in which to place these data points. Then a kernel function is chosen, to separate two distinct classes based upon this data. Kernel functions can be linear, polynomial, sigmoid, or a radical basis function, but they all draw a dividing boundary between classes. These functions are formed based upon labeled training sets, and then are used to categorize unlabeled data. Depending on the data set, different kernel functions have varying performance: for example, nonlinear kernels are better when the class gradients are not related in a linear fashion. Larger training sets reduce the chances of error when classifying objects, but that requires a large database of pre-labeled images. Moreover, it is possible to discover categorical features which allow for easier classification. For example, it could be discovered that signs tend to be vaguely circular, and a few feet above the road. This means that the algorithm does not need to spend time determining if it is a pedestrian, and immediately works to determine what kind of sign it is. The result allows for a set of comparisons to determine which class a data set belongs to. The limiting factor of an SVM is that it only determines which of two possible classes an object is most similar to. This means that many different comparisons need to be done to identify the object, because it only answers basic yes/no questions of object class.

Possible Improvements

Image processing is a rapidly advancing field, due to the myriad of applications. As such, research is being done to increase the speed and accuracy of object recognition and classification. One way that it could improve is by the advance of computer processors: faster computation will lead to faster image processing. Another improvement would be based on kernel functions with larger training sets, allowing for more accurate classification. The last improvement discussed here is improving algorithms by understanding classification data clustering. Some object classes share the same features, such as shape or color, leading them to be clustered around a specific area in the multidimensional space. Algorithms can use this fact to their advantages to rule out unlikely classifications, reducing the number of binary comparisons required to identify the object. While converting optical data into a usable form is a significant challenge, new technologies are being developed and tested to surmount it. The use of image processing allows for a better understanding of the world around an autonomous vehicle. Another way to facilitate this understanding is via active communication, a form of which is another key challenge facing the development of autonomous vehicles.

Vehicle to Everything (V2X) Communication

As stated previously, giving autonomous vehicles an understanding of the world around them is one of the key challenges that face the field. Sensors are helpful to provide an accurate view of the immediate surroundings to an autonomous vehicle, but they are limited in range and application. No sensor can tell that the road up ahead will be closed, nor can it coordinate with surrounding cars to manage traffic. However, a communication system from the autonomous vehicle to other vehicles, infrastructure, and pedestrians can. This communication system can be used to provide autonomous vehicles with accurate and collective information, which it can use to make good driving decisions. Collectively, these communication systems are described as Vehicle to Everything, or V2X communication. This technology has several useful applications in the field of autonomous vehicles, such as object recognition, intersection protocol, and advanced cooperative driving. Current technical limitations with 4G cellular networks limit the use of V2X communication, but upcoming 5G networks could be used to resolve these problems.

Applications of V2X Communication

V2X communication is based upon nearby entities transmitting and receiving data through cellular networks. This communication encompasses Vehicle to Vehicle (V2V), Vehicle to Infrastructure (V2I) and Vehicle to Pedestrian (V2P) communication. For example, a vehicle could send out information in regular pulses regarding its location, heading, speed, which other vehicles could use. Intersections could communicate with vehicles to manage traffic, and pedestrians could have their phones send out information to assist vehicle object recognition software. This technology matters, because it simplifies difficult software problems for autonomous and semi-autonomous vehicles. The first problem that it helps solve is object recognition. Current technologies for recognizing moving images are slow and memory intensive. However, if there was a V2P protocol, a pedestrian could just have their location broadcast to nearby vehicles, circumventing a lengthy object recognition task (Hussein, García, Armingol, & Olaverri-Monreal, 2016). V2X communication would allow a vehicle to classify objects based upon the information they send out, improving their ability to rapidly identify nearby entities. Even for non-autonomous vehicles, this could be used to help prevent collisions, by sending warnings to human drivers, or automatically applying the brakes. V2X communication allows for vehicles to have a better awareness of the world around them, leading to safer road conditions for all (Chen et al., 2017).

Intersection management. The next problem that it helps solve is vehicle interactions at intersections. Intersections pose a problem to autonomous vehicles, because of how many ways they can be set up. They can vary in terms of lanes, lights, turn restrictions, and pedestrian crossings. Along with these numerous setups, safe actions at these intersections is heavily dependent on communication. From drivers waving pedestrians onward, to barely stopping at a stop sign, intersections require a great deal of coordination with other vehicles. Aside from other drivers, one also must be aware of traffic signals, a key feature of city infrastructure. All of these require precise communication, with little delay, and V2X communication systems can facilitate them. This makes information such as signal countdowns, pedestrian activity, and the actions of other vehicles into common knowledge for autonomous vehicles. The demystification of intersections is vital to autonomous vehicle artificial intelligence, and V2X communication is the perfect tool to accomplish this (Le, Festag, Baldessari, & Zhang, 2009).

Collaborative driving. Aside from merely making the processing load on autonomous vehicles easier, V2X communication allows for advanced collaborative driving. Normal drivers need to leave roughly two seconds of following distance to avoid collisions. This is because drivers do not necessarily know what the vehicle ahead of them will do, and maintaining certain distance gives them more time to react to any sudden changes. Autonomous vehicles are not bounded by human reaction times, and V2X communication can address the problem of unpredictable drivers ahead. Due to this, V2X communication could allow for platooning, where autonomous vehicles communicate with each other to drive in tight formations at high speed. This has numerous advantages, such as improving road capacity, fuel efficiency, and greatly reducing traffic. This requires very low latency communications and software but has incredible rewards if executed effectively (Fernandes & Nunes, 2010). This requirement leads to the next subject, the problems facing V2X Communication.

Problems Facing V2X Communication

V2X communications has high baseline requirements. First, there needs to be reliable internet access, to pull information from online sources. This information includes road closures, weather data, and road conditions. Infrastructure data such as this is crucial for determining how an autonomous vehicle chooses a route and requires reliable internet access. This information is to be stored in a curated cloud server system, to be available to all autonomous vehicles that require it (Yao et al., 2018). This is Vehicle to Infrastructure communication, which stands with Vehicle to Vehicle communication and many other communication targets.

A large amount of data used in autonomous driving is stored online, making reliable internet access indispensable. Also, a low latency way of transferring data between devices is critical. The best way to fulfill both requirements is with cellular networks, such as 4G LTE. However, this poses a serious problem: latency. A slight delay between receiving information and acting upon it can have large consequences in a driving scenario, so minimizing latency is a top priority. The acceptable margin of latency is roughly 10 ms for autonomous driving, using V2V and some V2I communication systems (Amjad, Sikora, Hilt, & Lauffenburger, 2018). Current 4G standards can only guarantee a minimum latency of 50 ms, far too high for V2X applications. Aside from network requirements, there are less technical issues at play. For V2X communications to be implemented, there must be a standardized communication protocol. Manufacturers would need to collaborate to ensure that their systems are compatible with other vehicle brands. Another concern is the security risks posed by V2X communications. Incorrect or falsified V2X data could be used to force a vehicle to crash, or otherwise lose control. To combat this fact, there must be strong cybersecurity measures and requirements on these systems. To ensure that security requirements are universal, there needs to be a national standard for these regulations, which would require legislative action (Brecht et al., 2018). Outside of the automotive industry, phone manufacturers would need to implement automatic broadcasting of relevant pedestrian data. To summarize, V2X communication relies on networks with latency problems and faces obstacles to widespread adoption. New developments in 5G networks hold the potential to resolve these technical problems.

Possible Improvements Through 5G

To understand why 5G is so significant, some background on previous cellular networks is helpful. The first network in this scope is 3G. It offered good data rates, allowing bandwidth intensive applications such as video streaming to be practical. However, the coverage is limited, and the security measures it uses are weak. The next generation, 4G improves on 3G in almost all cases. OFDMA allows for much better spectral efficiency, and MIMO protocols allow for more users to be supported at a time (Ezhilarasan & Dinakaran, 2017). However, both of these technologies have latency specifications that are insufficient for autonomous vehicle applications. 5G technology can change this. The data rate that it can support far outstrips 4G, and its latency is far lower. As stated, V2X communications require less than 10 ms of latency (Amjad et al., 2018). 5G latency is typically under 3 ms, enabling autonomous vehicle communication in a way that both 3G and 4G cannot. This crucial latency reduction is what solves several problems that V2X communication faces. 5G technology also has the benefit of being new enough that standards can be built with autonomous vehicles in mind. 5G can operate at currently unused bandwidths, which means that certain portions of the spectrum can be reserved exclusively for V2X communications. China has done this for vehicle testing purposes, to the great benefit of V2X testing (Chen et al., 2017). Since current devices do not use it, new developments could be subject to IEEE regulations. This means that 5G capable phones could automatically send out useful data for autonomous vehicles. For example, pedestrian data such location and speed would greatly improve the reliability and accuracy of pedestrian recognition by autonomous vehicles (Hussein et al., 2016). The combination of low latency, high data rates, and new regulations makes 5G a promising technology for V2X communication.

Vehicle to Everything communication has the potential to revolutionize autonomous vehicles. By implementing it, autonomous vehicles can work more efficiently, safely, and use complex maneuvers such as platooning. However, this implementation requires strict network specifications. The latency requirements cannot be met by existing networks, limiting the application of V2X communications. Upcoming 5G networks can meet these specifications, making the implementation of V2X communication more feasible. Along with meeting higher specifications, the new network provides the opportunity to standardize communication protocols and security. Overall, 5G technology will be essential for the development of V2X communication systems.

Decision Making

The previous topics discuss the process in which an autonomous vehicle gathers information about the world around it. Furda & Vlacic (2011) move to describe how this can be synthesized into a so-called *World Model*, allowing the autonomous vehicle to understand the current context that it is in. Sensors, V2X communication, and GPS technology all feed into this understanding of the world. This model is a requirement for making driving decisions, because it provides required external context to guide the decisions made by a vehicle. After all, the options available in a highway merging scenario are completely different from those available at a four-way intersection. The concept is to break down the complexity of the world into a set of known events for which the autonomous vehicle has protocols to deal with. Once the vehicle has all of the information it needs, only then can it decide which maneuver it will use. These maneuvers are a sequence of inputs which will achieve an overall driving goal. For example, passing a vehicle on the left requires use of the accelerator, turn signal, and steering system. While there are multiple inputs required to accomplish this task, their combination is referred to as a single maneuver. To accomplish real-time decision making, autonomous vehicles will need a library of preplanned maneuvers to choose from. This will allow them to perform every driving action a human driver can, in a consistent manner. To choose effective maneuvers, information about the world around the vehicle is crucial. The process of gathering information, using it to model the world, and then using that model to decide on a maneuver is the essence of autonomous driving. From V2X communication to sensor processing, each way of perceiving the outside world is crucial for developing a world model, and this model allows informed decisions to be made. The result is a framework that provides the background information for the decision-making system to work as intended.

Decision Making System

To understand this system, one must first understand the prerequisites. The first prerequisite is a set of maneuvers available, such as making a left turn, or merging on a freeway. The second prerequisite is a set of world model events, and the last is a planned route. The goal of this system is to select the most appropriate maneuver in any given situation, given the current events in the world model and overall route. While there are many possible options that can be done in a specific situation, the first stage is to determine which of them are feasible given the current events. Feasible maneuvers are those that are safe and appropriate, and usually adhere to traffic rules. This set of maneuvers describes the overall set of options available to the vehicle. The next stage of decision making is picking an option from these, a much more complicated task. One way to do this is to have an objective hierarchy that maneuvers interact with. Different goals are assigned different priority values, and maneuvers are chosen based upon how well they accomplish these weighted goals. For example, staying in road boundaries is the top priority, with maintaining following distance beneath it, and minimizing waiting time beneath that. Each maneuver would have a utility function based upon a combination of the current events and the objective hierarchy. This method would simplify this decision-making stage to simulating each maneuver and choosing the one with the highest utility value. However, this is only a generalized decision-making model, while driving is filled with highly specific situations. Each of these situations needs to be anticipated, with a decision-making protocol in mind. The following is an example of one such preplanned situation.

Decision Making Example – Urban Pedestrian Crossings

A 2019 paper by Batkovic, Zanon, Ali, & Falcone provides an example of a specific situation and a protocol to handle it. Their goal is to determine how a vehicle should act when a pedestrian crosses a street in front of it, outside of a crosswalk. They created an algorithm to deal with this event in an urban environment. The algorithm first determines the projected path of the vehicle, which is based upon the preprogrammed route. The algorithm then tracks the motion of the pedestrian, to determine their probable position in the next few seconds. Their probable position is then used to create a zone in which the vehicle will avoid. Based upon this information, the vehicle plots a trajectory that it intends to take, then implements it. This trajectory is based upon required distance tolerances from both the pedestrian, and road boundaries. The algorithm leads to a variety of behaviors, from simply slowing down, changing course, or some combination of the two. Multiple trajectories are planned, then the one to execute is chosen from a utility function. Due to the simplicity of the algorithm, the time required for a decision to be made is very short, with an average decision time of 22.4 milliseconds. This is a good example of how a specific situation that an autonomous vehicle would face can be handled (Batkovic et al., 2019). While this algorithm only works for this one situation, it works exceptionally well. A small protocol such as this one is not close to sufficient to operate a vehicle alone. However, if one combines enough protocols like this together, that changes. A patchwork of protocols can work together to navigate a host of situations as they are recognized by the world model. Along with being modular, algorithms such as this have the advantage of requiring relatively little computation time. The world model can identify any given situation, then hand the decision making off to the assigned protocol. A horde of small algorithms such as this is a possible implementation of an overarching decision-making system. However, preprogrammed protocols are not the only option available. Neural networks present a more flexible option, known as End-To-End Decision-Making.

End-To-End Decision-Making Systems

Recent research conducted by NVIDIA, which is documented in the work by Bojarski indicates another control scheme: End-To-End Decision-Making (Bojarski et al., 2016). This is the concept of training artificial intelligence directly using driver data, instead of programming actions for each situation. The input is controlled on one end, and a neural network is used to provide output commands on the other. The computation in the middle is not directly controlled, and it converts input images to output controls. The study gathered 72 hours of driver data, in the form of video footage. This footage was then divided into single images paired with steering inputs, to make a training set for a convolutional neural network (CNN). The overall goal of this is to allow this network to be sent images, then generate steering commands based upon them. For example, training the network to follow a lane would be done by providing it a set of images where the human driver was following their lane, along with steering commands. The network will then be able to mimic human steering commands in similar situations. Some of the images in the training set were deliberately rotated or set off center to instill error correction into the network.

Empirical tests of the CNN show that it can follow road markings and stay in lane in both simulations and real-world conditions (Bojarski et al., 2016). This process of using training data and neural networks to allow vehicles to make decisions has several useful traits. First, it requires relatively small data sets to function, and does not need to be trained manually. There is no need to create an image classifier or define a set of decision-making protocol. One simply trains the CNN, and it mimics the inputs associated with the training data. This can save countless hours of hard coding responses to specific situations. Alongside this, a CNN is able to work more effectively in novel situations, reacting in a way similar to the training data.

However, this approach is not without drawbacks. It is very difficult to understand how a CNN makes individual decision, and this process is not entirely under the designer's control. For a CNN that is designed to sort handwritten zip codes, this is not an issue. For a CNN that pilots a vehicle, precise knowledge and control of decision making is crucial. A developer needs to know exactly what a vehicle will do at an intersection, versus sending in a trained CNN and hoping for the best. While carefully curated training data can help define how a CNN will act, the developers do not have complete control over the CNN's decision making. End-to-End Decision Making has promise, but this promise should be tempered with caution.

Possible Improvements

There are many projects with the defined goal of developing and testing autonomous vehicle decision making. There are studies in academia, along with research conducted by automotive companies like Tesla and software giants like Google. They are all working toward the same goal, but independently. Instead of having one universal protocol to manage highway merging, you have at least ten. This diversity has benefits, but it also makes it extremely difficult to centralize decision making modules and leads to repeated research on a small pool of problems. A way to fix this would be to make a central database of decision-making software, freely accessible to all researchers. That way, instead of having ten research teams working independently on the same problem, one team could solve a particular problem, and others could incorporate that software into their own research. A centralized database of control software would lead to more efficient research efforts, due to reducing research repetition. Another way to address the challenge of decision-making is by reformulating control systems with field data. Prototype vehicles are gathering test data to improve the algorithms that they follow. From this, developers can experiment with new decision-making systems, and keep those that are most effective. This is of special importance to neural network-based control schemes, because they benefit greatly from higher volumes of test data. Experimentation and feedback are a crucial part of technological development, both of which are provided by road testing. Road testing can improve existing decision-making algorithms and provide the data to train new ones. Decision making is one of the key challenges to creating effective autonomous vehicles, and progress is being made to resolve it. However, decision making is not an entirely technical problem. The choices made by autonomous vehicles have lives at stake. Thus, there is an ethical component to how they should be made. This comprises a substantial portion of the socioeconomic challenges facing autonomous vehicles.

Socioeconomic Challenges

Ethical Considerations

Driving is an inherently risky activity, with potential impacts that go beyond just the driver. According to a literature review by Fagnant and Kockelman (2015), there are roughly 5.5 million crashes a year in the US, over 90% of which have a human cause as the primary factor. Even with perfect autonomous vehicle artificial intelligence, this indicates that there will be thousands of crashes that cannot be prevented. While autonomous vehicles can reduce the risk of accidents, eventually they will run into crashes, during which they face ethical dilemmas. Most crashes have little warning, leading to no time to decide on how to act during them. There is only a split second to determine what party will take the brunt of the damage. That is not enough time for a human driver to make an informed ethical judgement, but it is enough time for an autonomous vehicle to do so. This leads to the complex topic of what ethical theories an autonomous vehicle should follow, a topic explored in detail by Goodall (2014). The first theory investigated is deontology, where the behavior of a vehicle is governed by strict, unbreakable rules. An example of this type of system is Asimov's three laws of robotics, an ordered list of maxims. Every decision the vehicle could make would be subject to an overarching principle, such as obeying traffic laws, or minimizing injuries. These principles can be arranged in importance, with lesser principles being able to be superseded by more important ones. To demonstrate, a vehicle could swerve into a turn lane to avoid a crash, valuing crash avoidance over following road laws. Deontology poses a problem in implementation, because each correct response for a particular situation would need to be programmed individually. This is because such rules need to be interpreted with common sense, with exceptions and clarifications depending on the situation. This leads to deontology being closer to a code of ethics for the decision-making system than anything else. This entails that a deontological judgement needs to be preprogrammed, hindering its application to unanticipated situations. The next theory of interest is utilitarianism, where the goal is to have the outcome of an action maximize some measure of utility. This means that acts such as saving a life or following rules of the road are given an explicit utility value, used to score the consequences of actions. This can then be used to rank possible maneuvers by their probable outcomes in terms of utility, then pick the highest. While this system is more adaptable than deontological ethics, it runs into issues of how utility should be assigned. This is the stems from the fact that this system would require putting a measured value on people's lives, a large ethical controversy. Beyond that, utilitarianism also tends to maximize overall societal good, without respect to individual rights. For example, imagine a situation where a vehicle would need to swerve to avoid an oncoming obstacle, with a motorcyclist on both sides. One of them is wearing a helmet, and the other is not. To avoid a crash that affects them all, the vehicle would need to swerve into one of the motorcyclists. For this case, a utilitarian system would choose to crash into the helmeted motorcyclist versus a non-helmeted one. This is because the motorcyclist with the helmet is likely to incur less severe brain injuries than their non-helmeted counterpart, minimizing overall injury. Though society is made better off, it does not factor in individual rights or benefits. That said, utilitarianism is a practical basis by which machines can make ethical choices, because it can convert a complex moral choice to a utility function. While this may seem callous, it provides a set of guidelines for an autonomous vehicle to make a moral choice in unexpected situations, unlike deontological ethics. Autonomous vehicles will need a system of ethics, due to the unavoidable nature of certain accidents. However, regardless of the system of ethics chosen, one encounters a social dilemma, based upon the will of the driver.

The Social Dilemma

A study conducted in 2016 showed that survey results indicate that the public at large would prefer autonomous vehicles to operate based upon a utilitarian system of ethics (Bonnefon, Shariff, & Rahwan, 2016). This is not unusual, but the data provides a dilemma along with this finding. Utilitarian ethical systems might determine that it is worth sacrificing the life of the driver to save a greater number of pedestrians. When this fact was posed to survey takers, a moral inconsistency emerged. They were asked two questions about their likelihood of purchasing an autonomous vehicle, which varied in the ethical system the vehicle used. The first option was regulated by utilitarian ethics, with the goal of minimizing overall casualties in the case of an accident. The second (non-regulated) ignored overall casualties and would be willing to sacrifice up to twenty pedestrians to save the driver. Most of the survey respondents indicated that autonomous vehicles should minimize overall road casualties, even at the expense of the driver's safety. The dilemma arises when one realizes that only 38% of respondents said that they would be willing to purchase a car that would sacrifice the driver to save pedestrians. On the other hand, 53% who would purchase a vehicle using an algorithm that prioritizes the protection of the driver. While most respondents favor utilitarian algorithms, most of them would not want to purchase a vehicle that follows one for themselves.

This data shows a classic conflict between self-interest and the common good, one with far-reaching consequences. There are three main groups that have a stake in the ethical systems of autonomous vehicles: consumers, governments, and manufacturers. Consumers have shown that algorithms that favor protecting the driver would be more popular, pushing manufacturers to make vehicles that use said algorithms. Governments have a stake in protecting the common welfare and can impose regulations on manufacturers regarding the ethical systems their vehicles use. These interactions make ethical frameworks a complex issue, with many interconnected parties. While ethical decision making will only be a factor in a tiny portion of crashes, it will bring conflict regardless. Only time will tell whether utilitarian frameworks will be regulated into prominence, or if more selfish algorithms become widespread. Keep in mind, that regardless of the ethical system in place, autonomous vehicles can prevent countless crashes. The projected reduction in automotive deaths provides ample reason to continue developing autonomous vehicles, despite the challenges in choosing an ethical system. While ethical systems are a significant part of the socioeconomic challenges facing autonomous vehicles, they do not stand alone. The economic effects prove to be just as relevant of a factor.

Economic Pushback

New innovations can have large impacts on old markets, both positive and negative. Old positions become obsolete due to new technologies, through the process of creative destruction. For example, the rise of the automobile was the downfall of countless stables, but also sparked a booming manufacturing industry. The autonomous vehicle is an innovation on a similar scale, with a similar level of impact. A study in 2017 investigated the likely economic impacts of autonomous vehicles, and research suggests that they will be expansive (Clements & Kockelman, 2017). To begin, it makes on-demand vehicle use much more practical, and personal ownership less so. Without the cost of a driver, companies such as Uber can charge lower prices, low enough to where consumers can choose to not own a car. A fleet of on-demand autonomous vehicles is the most efficient way to use the technology, as compared to individual ownership. The next major impact is in Vehicle Miles Travelled (VMT). As the time cost of driving goes down, the average miles travelled will go up. Simulations show that VMT will rise anywhere from 3.6% to 19.6%. This will induce faster vehicle turnover, resulting in greater sales for the automotive industry, a projected increase of anywhere between $29-57 billion USD a year. However, the impact of autonomous vehicles will not be entirely positive. The trucking industry employs over three million drivers in the US. If those drivers were eliminated by autonomous vehicle technology, the trucking industry stands to save $100-500 billion per year. While it is an economically advantageous option, it has the side effect of leaving three million drivers unemployed. The situation with the taxi industry is similar, with the likely result of taxi drivers losing up to 50% of their customer base. Previous statistics show that extremely good autonomous vehicle artificial intelligence can reduce crashes by up to 90%, reducing business for automotive collision repair. While some of this business may be recovered due to the larger number of miles driven, common usage of autonomous vehicles will eliminate much of this market. All told, this would likely cause a $27 billion reduction to revenue for the automotive collision repair industry. These are all examples of industries that will be heavily impacted by the rise of autonomous vehicles, with both positive and negative effects. While society would be better off overall for the adoption of autonomous vehicles, parts of private industry would be devastated. Millions of people would lose their jobs, and companies would lose billions in revenue. Both groups have a large economic incentive to oppose the deployment of autonomous vehicles. The likely result is that there will be heavy lobbying efforts from affected industries to make it more difficult to build and use autonomous vehicles. If mass layoffs in the transportation industry come to pass, there will likely be large scale protests from newly unemployed truckers and taxi drivers. The overall benefits of the technology amount to $3,814 per capita in the US, but said benefits are not distributed evenly. The pushback from that fact becomes a challenge of its own. The economic conflict between the winners and losers in the rise of autonomous vehicles may not be a technical challenge, but it is a major obstacle to the widespread use of autonomous vehicles.

Conclusion

Autonomous vehicles face many challenges that need to be addressed before they achieve widespread adoption. Technologies such as image processing and V2X communication can be used to greatly improve the ability of a vehicle to sense the world around them but will need further improvements. The decision-making process bridges the gap between technical and social challenges, but it will improve with time and testing. Lastly, ethical and economic challenges will lead to social and economic pushback against this technology, but the overall benefits outweigh the costs. These challenges are being addressed, moving ever closer to effective solutions. As processing power grows cheaper and lighter, image processing becomes more practical and effective, as do real-time decision-making systems. 5G technology has been developed and deployed, setting the baseline requirements for V2X communications. While ethical concerns are a difficult matter to address, the overall lives saved by autonomous vehicles outweighs the rare cases in which a predetermined sacrifice will be necessary. Economic concerns will be an issue, but there is a precedent for new technologies displacing obselete industries. To conclude, while autonomous vehicles face several challenges in adoption, they are being addressed. As technical fields develop and research provides new findings, these challenges will fall, one by one. Currently, autonomous vehicles are not ready for widespread use. However, as the challenges are overcome, the viability of autonomous vehicles will grow. There is no longer a question of whether autonomous vehicles will be adopted or not. There is only the question of when.

References

Amjad, Z., Sikora, A., Hilt, B., & Lauffenburger, J.-P. (2018). *Low latency V2X applications and network requirements: performance evaluation.* Paper presented at the 2018 IEEE Intelligent Vehicles Symposium (IV), Changshu, China. doi: 10.1109/ivs.2018.8500531

Batkovic, I., Zanon, M., Ali, M., & Falcone, P. (2019*). Real-time constrained trajectory planning and vehicle control for proactive autonomous driving with road users.* Retrieved from http://arxiv.org/abs/1903.07743

Bojarski, M., Del Testa, D., Dworakowski, D., Firner, B., Flepp, B., Goyal, P., Zieba, K. (2016). *End to end learning for self-driving cars.* Retrieved from https://arxiv.org/abs/1604.07316v1

Bonnefon, J.-F., Shariff, A., & Rahwan, I. (2016). The social dilemma of autonomous vehicles. *Science, 352*(6293), 1573–1576. doi: 10.1126/science.aaf2654

Brecht, B., Therriault, D., Weimerskirch, A., Whyte, W., Kumar, V., Hehn, T., & Goudy, R. (2018). A security credential management system for V2X communications. *IEEE Transactions on Intelligent Transportation Systems*, *19*(12), 3850–3871. doi: 10.1109/TITS.2018.2797529

Chen, S., Hu, J., Shi, Y., Peng, Y., Fang, J., Zhao, R., & Zhao, L. (2017). Vehicle-to-everything (v2x) services supported by LTE-based systems and 5G. *IEEE Communications Standards Magazine, 1*(2), 70–76. doi: 10.1109/MCOMSTD.2017.1700015

Clements, L. M., & Kockelman, K. M. (2017). Economic effects of automated vehicles. *Transportation Research Record*, *2606*(1), 106–114. doi: 10.3141/2606-14

Ezhilarasan, E., & Dinakaran, M. (2017). *A review on mobile technologies: 3G, 4G and 5G.* Paper presented at 2017 Second International Conference on Recent Trends and Challenges in Computational Models (ICRTCCM), Tindivanam, India. doi: 10.1109/icrtccm.2017.90

Fagnant, D. J., & Kockelman, K. (2015). Preparing a nation for autonomous vehicles: opportunities, barriers and policy recommendations. *Transportation Research Part A: Policy and Practice*, *77*, 167–181. doi: 10.1016/j.tra.2015.04.003

Fernandes, P., & Nunes, U. (2010). *Platooning of autonomous vehicles with intervehicle communications in SUMO traffic simulator.* Paper Presented at the 13th International IEEE Conference on Intelligent Transportation Systems, Funchal, Portugal. doi: 10.1109/ITSC.2010.5625277

Furda, A., & Vlacic, L. (2011). Enabling safe autonomous driving in real-world city traffic using multiple criteria decision making. *IEEE Intelligent Transportation Systems Magazine*, *3*(1), 4–17. doi.org: 10.1109/MITS.2011.940472

Goodall, N. J. (2014). Machine ethics and automated vehicles. *Road Vehicle Automation Lecture Notes in Mobility*, 93–102. doi: 10.1007/978-3-319-05990-7_9

Harshitha, R., & Manikandan, J. (2017). *Design of a real-time pedestrian detection system for autonomous vehicles.* Paper presented at Region 10 Symposium, Cochin, India. doi: 10.1109/tenconspring.2017.8069981

Hussein, A., García, F., Armingol, J. M., & Olaverri-Monreal, C. (2016). *P2V and V2P communication for pedestrian warning on the basis of autonomous vehicles*. Paper presented at the 2016 IEEE 19th International Conference on Intelligent Transportation Systems (ITSC), Rio de Janeiro, Brazil. doi: 10.1109/ITSC.2016.7795885

Le, L., Festag, A., Baldessari, R., & Zhang, W. (2009). Vehicular wireless short-range communication for improving intersection safety. *IEEE Communications Magazine*, *47*(11), 104–110. doi: 10.1109/MCOM.2009.5307473

Nkosi, M. P., Hancke, G. P., & Santos, R. M. A. D. (2015). *Autonomous pedestrian detection.* Paper presented at AFRICON 2015, Addis Ababa, Ethiopia. doi: 10.1109/afrcon.2015.7332014

Patel, S. (2017). *Chapter 2 : SVM (support vector machine) - theory*. Retrieved from https://medium.com/machine-learning-101/chapter-2-svm-support-vector-machine-theory-f0812effc72.

Siegel, J. E., Erb, D. C., & Sarma, S. E. (2018). A survey of the connected vehicle landscape—architectures, enabling technologies, applications, and development areas. *IEEE Transactions on Intelligent Transportation Systems*, *19*(8), 2391–2406. doi: 10.1109/tits.2017.2749459

Teichman, A., & Thrun, S. (2011). *Practical object recognition in autonomous driving and beyond.* Paper presented at Advanced Robotics and its Social Impacts (ARSO) Conference, Half-Moon Bay, CA, USA. doi: 10.1109/arso.2011.6301978

Yao, Z., Yang, Z., Wu, T., Chen, L., Zhu, K., Zhang, L., & Su, S. (2018). *Implementing ITS applications by LTE-V2X equipment-challenges and opportunities.* Paper presented at the 2018 International Conference on Network Infrastructure and Digital Content (IC-NIDC), Guiyang, China. doi: 10.1109/icnidc.2018.8525768

Yuan, X., Hao, X., Chen, H., & Wei, X. (2014). Robust traffic sign recognition based on color global and local oriented edge magnitude patterns. *IEEE Transactions on Intelligent Transportation Systems*, *15*(4), 1466–1477. doi: 10.1109/TITS.2014.2298912