Mahmoud, 4:00
Team 9

MACHINE LEARNING IN SELF-DRIVING TESLA AUTOMOBILES

Jared Leuffen jrl139@pitt.edu, Jash Patel jsp69@pitt.edu, Wenjie Xu wex40@pitt.edu

Abstract--*Machine Learning is the next step in the already growing world of artificial intelligence. In recent years, the idea of using machine learning to make our world smarter has become more and more feasible. Although the idea for machine learning is not new, the technology has become more developed and advanced, and it is beginning to be implemented in many industries including Data Security, Financial Trading, and Healthcare. This paper will look at the use of machine learning in self-driving smart cars, specifically at Tesla, and talk about how these self-driving cars are improving upon safety and navigation with the help of machine learning. We first provide background of how the current programing technology in self-driving cars works and the role humans play in this programing. The paper will go into depth about what machine learning is, how it works, and how it plays into the self-driving car industry, specifically at Tesla. We will talk about how machine learning will impact cars through training and show examples of cars using machine learning to become smarter. This paper then breaks down all this information even further, and goes into the applications of machine learning, specifically in the realm of safety and navigation. We will finish the paper with examples of machine learning in action and will include common questions and evaluations of the use and effectiveness of machine learning in self-driving cars at Tesla.*

Key Words—Artificial Intelligence, Machine Learning, Safety, Self-Driving Cars, Tesla

THE BASICS OF SELF-DRIVING CARS AND THEIR PROGRAMMING

The self-driving car industry has really begun to attract attention as the technology has advanced in recent years. Many big-name companies have started to develop self-driving vehicles, including Google, Uber, and Tesla. Self-driving vehicles are vehicles in which human drivers are not required to take control in order to operate the vehicle safely. Self-driving systems create and maintain an internal map of their surroundings by using a wide array of sensors, lasers, and high-powered cameras. An example of this can be seen below.

FIGURE 1 [1]
A map in a self-driving car that was constructed from sensors gathering information

The figure shows a picture of a map that was produced as a result of gathered information. This information is then sent through computer software, which processes it and sends it to the car's hardware, or "actuators." These are the things that control the steering, acceleration, and breaking of the car [1]. Coded rules, predictive modeling, and obstacle avoidance algorithms are all involved, which help the software follow traffic rules and navigate obstacles. Most situations are encoded into the system, so that the cars can recognize the environment and make the associated decisions. Many of these current autonomous systems require constant programming, and reprogramming by humans, in order to operate. This programming is extremely complex, and often requires multiple people to accomplish. This can be very time consuming, and can lead to possible errors, as no human is perfect. There are exceptions to these normal systems, however. Some self-driving systems distinguish themselves as being "connected," or able to

interact and communicate with other cars. Although most of the current systems do not have the capability, many makers of self-driving cars are trying to change that [1].

How Tesla has gotten involved

Tesla is no stranger to various forms of technology. However recently, the company has been in the spotlight for one of their biggest technological advances, self-driving cars. Many have acknowledged the fact that Tesla is ahead of the game when it comes to self-driving car systems [2]. All Tesla vehicles currently being built have the hardware needed for full self-driving capability, known simply as autopilot. Just like any other self-driving car system, Tesla's cars utilize cameras, sensors, and lasers to operate their autopilot mode. The figure below shows an example of the Tesla Autopilot system.

FIGURE 2 [2]
Tesla's Autopilot detection system

As you can see in Figure 2, the autopilot system is using various sensors, and is using information gathered to recognize the types and positions of the vehicles surrounding it. Ever since Tesla came out with the first version of their self-driving car, they have been gathering and analyzing thousands of gigabytes of data. From all this data they've gathered, Tesla has used it to provide various updates to their different self-driving cars, with the goal to improve the way they operate [2]. Although Tesla's self-driving cars are already advanced, just as any good technology company would, Tesla is looking for a way to improve their current technology they utilize in their systems. They want their self-driving cars to be safer, have better navigation, and become overall more efficient. Tesla wants to distinguish itself from other companies creating self-driving cars, such as Uber and Google. Their main way of doing this is to design and create self-driving systems that are "connected" to the world and other autonomous vehicles. In order to accomplish this, Tesla has employed the use of something known as Machine Learning, where the system reprograms itself to improve functionality [2]. This avoids the necessity for constant human interaction and reduces the chances for human error, with the goal of a safer end product. Tesla believes that implementing machine learning into their systems will not only help separate them from their competitors, by will overall make their self-driving cars better [2]. Tesla believes that machine learning is the next step in the advancement of the self-driving care industry.

WHAT IS MACHINE LEARNING

Machine learning is the scientific study of algorithms and models that computers use to solve problems. It is an application of Artificial intelligence (AI) that provides systems with the ability to automatically learn and improve the effectiveness of their programming without being explicitly reprogrammed by humans. Machine learning focuses on the development of computer programs that can access data and use it to learn for themselves [3]. The general goal of machine learning is to understand how data is structured, and how this data fits into models that can be utilized by people [4]. When applied in the field of automobiles, machine learning gives automobile systems the ability to analyze data and learn from it without human programing. Although machine learning falls into the field of computer science, it differs from the traditional computer approach. In traditional computer programming, algorithms are explicitly programmed instructions that computers use to calculate something or solve problems. Machine learning algorithms instead allow for computers to train and evolve from data input and use statistical analysis to produce output values [4].

There are many types of machine learning algorithms however in this paper we will focuses on two: Supervised machine learning algorithms, and Reinforcements machine learning algorithms. The first one we look at; supervised machine learning algorithms is the main algorithm implemented in the "connected' world of self-driving cars. These algortihims work by applying what has been learned in the past to new, current data using labeled examples to be able to predict future events. In supervised learning, the computer is provided with examples that are labeled with their desired outputs. These examples are absorbed by the computer, and the algorithms use patterns to analyze the data [4]. Starting with the analysis of this previously known dataset, the learning algorithm produces functions to make predictions about the output values. It can provide "targets," or goals for new inputs after specific training. It can also compare its output with the correct, intended output, and find errors in order to modify itself accordingly [3].

The other type of machine learning algorithm we are interested in is the reinforcement machine learning algorithm. These are learning methods that interact with the environment by producing actions and discovering errors or rewards. This type of algorithm can be thought of as a trial and error, or goal oriented algorithm [5]. These algorithms allow machines and software to determine the ideal behavior

in a specific situation in order to maximize its performance. Simple reward feedback is required for these algorithms to learn. This is called the reinforcement system [3]. Reinforcement learning can be broken down in different definitions that describe the various stages involved. The first definition is the agent. The agent is responsible for taking actions. Followed by the agent, is the action. The action is the set of all possible moves the agent can make. The agent often chooses to from a list of possible actions and executes them [5]. To put this in terms of a car scenario, an agent would be the car being assigned to make a turn, and the action would be the car either making a right turn or a left turn. The final major component of reinforcement machine learning is the reward factor. The reward is the feedback in which the success of the action is measured [5]. In terms of the previous example, if the car turned left or right as expected, it would be "rewarded," but if the car didn't make a turn and went straight instead, the car did not achieve the desired outcome. This would result in no reward, which the car would process and recognize as a failed decision and use it to improve its performance the next time. Both types of learning algorithms are very powerful and play a big part in machine learning. Using these stages of machine learning, Tesla has developed its way of self-driving and is creating these systems in a more efficient and precise way.

Connecting through Machine Learning

In addition to the algorithms, one of the key components of machine learning is the ability to "connect," or communicate with other vehicles. It is not possible that a single vehicle can be exposed to every traffic situation. However, with machine learning, each car can greatly expand upon its experience, which will result in better performance. Each self-driving car gathers massive amounts of information whenever it is tested or driven. This data on the various situations encountered by each individual vehicle is shared through a network so that each computer controlling each car can adapt its algorithm to the environment faced by other vehicles [6]. Creating a big database every car can access allows all cars to "learn" from situations other cars have encountered. This type of shared experience and active learning creates a situation where autonomous cars, through Artificial Intelligence algorithms, can improve their ability to react in situations without having experienced those situations [6]. This allows for cars to use one another to learn and help improve their own systems by using other. With all the data that is constantly being gathered by every one of these computer algorithms, these cars that have incorporated machine learning will become smarter, just from being around other cars. The ability for cars to share data and learn from one another is the biggest advantage of machine learning. This proves to be essential for all these systems and can allow for maximum potential when it comes to self-driving cars.

INCORPORATION OF THE TECHNOLOGY

The idea of improving their cars to make them safer, navigate better, and be more efficient was certainly a reason why Tesla began implementing machine learning into their cars. The convenience as well as improved precision has led them to not only use machine learning but design their own algorithms. The original or older way of programming self-driving vehicles is to create a program that has many situations and responses encoded in it so that cars can detect their surrounding and make associated decisions based on the code that has been programmed into them.

Tesla has developed their own new advanced machine learning algorithms that are able to improve their systems and maximize their cars performance. These advanced algorithms that Tesla has created are able to break down thousands of videos and high precision maps and analyzes them to create better navigation and safety by tracking if adjustments made result in positive or negative results. The company effectively utilizes crowdsourcing to gather all of the data from all of their vehicles, as well as the drivers. Tesla's vehicles have both internal and external sensors, which can pick up basic information such as driver hand placement [7]. This data is continuously implemented into the algorithms to produce better results. These results are used to continually reprogram the navigation systems in Tesla's cars and "learn" what works and what doesn't [8]. This will then result in the system being exposed to more situations where it does not know what to do. As more data is processed by their systems, more experience is gained. This experience accumulates, and the system ends up continually evolving and improving without the need for human input [9]. Over time, navigation and safety can both be developed. Through a connection to a network, Tesla vehicles nearby share local information and insights. Machine learning takes care of educating the whole fleet, while at an individual car level, edge computing decides what action the car should be taking now [7]. The figure below shows a simplified example of the Tesla Network.

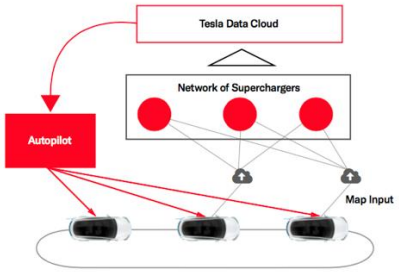

FIGURE 3 [13]
A snapshot of how the Tesla network works

As the figure shows, the cars' autopilot systems gather information, which gets sent though a network, to the Tesla Data cloud. The experiences that one Tesla car goes through is then applied to other autonomous driving cars. This results in multiple cars' computer systems "learning" from each other. What the cars gain from each other can then be applied to their own systems. This makes it much easier to improve all of Tesla's cars that use an autopilot system. As a result, all of Tesla cars will be able to accomplish more, while doing less themselves.

The Applications/Examples

One of Tesla's main concerns, which is common to every car manufacturer, is safety. Thousands of people die in motor vehicle crashes every year. Many of these deaths are the result of human mistakes or errors. A common example of this would be a person looking at their phone and not paying attention to the road in front of them or misjudging the time required to come to a complete stop at an intersection. The idea with self-driving cars is that their systems can prove to be less error-prone than humans, resulting in fewer crashes [1]. Machine learning will be able to help train cars to become safer by allowing them to experience more traffic situations in the real world. Through hundreds of hours of testing on roads, obstacle courses, and parking lots, the computers that these cars run on develop extensive experience in real world driving. By going through training for their own systems, as well as other cars' systems, the cars will begin to accumulate experience for everyday situations, and over time will become accustomed to almost any driving situation the vehicle will encounter. This will allow the car to react even better than a human would. In addition, self-driving cars don't have the chance of becoming distracted, such as a human driving a car would.

A common example of a possible safety situation where the car would have to react would be at an intersection or a four-way stop. From personal experience, there have been situations at intersections where cars either run the stop sign or red light, and almost cause an accident. This is a human error that could be prevented. This situation can be programed into the algorithms in self-driving cars, which could face the very same scenario in the future. Now every time the car faces a situation like this, it will react in a way that was different from before. The car is using information from past experiences and is applying it to the situation it currently faces. The information that the car has stored as a result of machine learning will improve its safety. An example of a car incorporating the information it stored would be by doing simple things such as letting others who come to the stop sign first proceed before it does. In this scenario, a self-driving car can recognize this spot where many human drivers would forget to stop. Based on the data of human drivers forgetting about the stop sign, the database of machine learning cars will pay more attention when near or driving through the pathway. Not only for the cars that have experienced dangerous situations at this four-way stop, all cars that can access the machine learning database can detect that spot as a highly dangerous spot, and the performance of every self-driving car is improved. Therefore, a self-driving car will be more careful about its surrounding when passing through the four-way stop. Although this seems simple, many accidents happen as a result of humans not being careful. It is hard to write rules for cars to know what to do in situations like these, but through long periods of self-learning, the system can figure out what to do without much human input. As a result of this "training," self-driving cars have the potential to become safer and more efficient.

In addition to safety, another application of machine learning is improved navigation. The algorithms used at Tesla can take maps that have been programmed into them, and break them down thousands of times, going into the smallest details. Doing this will allow the cars to learn to navigate better and get from one place to another in a quicker and more efficient manner.

Based on past experiences, there have been numerous times where a car that we have driven became lost due to either human navigation error, wrong turns due to outdated maps, or long delays due to accidents. The machine learning aspect of self-driving vehicles will be able to eliminate all of these problems. It will eliminate the outdated maps issue, as well as the accidents issue. Another example would be a road that is constantly jammed. The machine learning algorithms would report to the system that in the case of this road, it is better to avoid it. This is the first step of machine learning. Furthermore, machine learning can detect the time period that the road is heavily jammed, and cars can avoid that road only in certain time periods. This is the further step, compared to just avoiding a particular road due to current traffic. Over time, more cars are informed, and the system can detect the real time traffic flow in that region and cars can decide whether to pass through. This over-time developed analysis of the map, applied with machine learning, can improve driving time and experience. Due to machine learning, the cars' maps and navigation system are constantly updated every minute. By doing this the car will be able to determine when there is an accident ahead and will be able to find a different route that will get to the desired destination faster. The self-driving car aspect will eliminate the human error issues, as there will be not humans controlling the navigation system. The use of machine learning will be able to lead to a safer and more efficient outcome, which is why Tesla has decided to incorporate it into its self-driving cars.

CONCERNS AND EVALUATIONS

The idea of machine learning and the technology behind it is certainly a big step in the new world of artificial

intelligence. However, many people believe that it is not necessarily a good step.

One of the concerns associated with autonomous self-driving cars is still and always will be safety. We have discussed this idea with numerous people. It seems people fear that self-driving cars are programmed very specifically and will not be able react in unusual situations. The figure below shows an example of this.

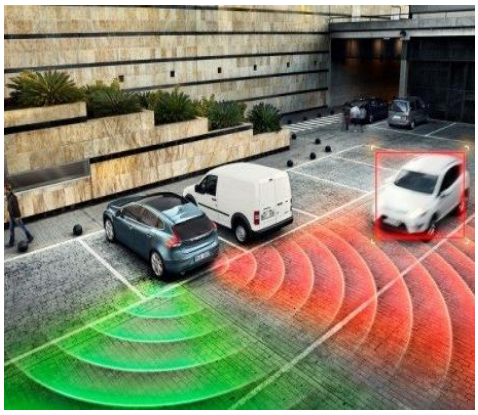

FIGURE 4 [14]
An example of a possible safety situation

In this situation, a self-driving car is getting ready to back out of a parking spot. The concern in this situation is that the car is directed to back out, and therefor will just begin to back up without any consideration of the vehicle behind it. However, the counter-argument to this is exactly what this paper is about: machine learning. Machine learning allows for cars to be reprogrammed all the time. As soon as a situation is faced by any car, the system will be able to process this and teach not only itself but also other cars what to do. In this situation, based on previous experiences, the car will immediately stop when it gets the slightest sense of an object behind it. Throughout the thousands of self-driving cars being developed, massive amounts of data will be able to train each car, so that eventually any car will be able to react to a situation just as well as any human would, if not better. Humans in certain specific scenarios will not be able to predict what they would do, because they have no way of knowing what the exact situation is that they will be facing. The machine learning, self-driving cars on the other hand will remember what has been learned in the past and react accordingly.

Sticking with the safety factor, many people are concerned with how self-driving cars would react in certain weather situations. People ask, "What about when it snows, rains, or the roads become icy?" Prior to the addition of anti-lock braking systems, drivers were trained to pump the brake when trying to slow down on a slick road. However, with the addition of antilock braking systems, this is not the case at all. It is better to just squeeze the brake and hold, and let the system do the work [10]. Machine learning will be able to train cars to do just that. This is where the reinforcement machine learning algorithm comes into play. The cars will be tested in both situations, and because of these algorithms, the system will learn that it is better to not pump the breaks in certain situations. This test that the one car went through will then be applied to the other cars, eventually resulting in many cars using this same information. Although it is understandable that safety is a concern for many people with self-driving cars, we believe that through machine learning, safety can be improved greatly. It is entirely possible that eventually self-driving cars are even safer than human driven cars.

One of the other major concerns when it comes to the idea of self-driving car systems is the issue of data security [11]. With machine learning and self-driving cars, there is no doubt a lot of data is being shared which creates vulnerability. All the data from every car adds up, which in the end translates to massive amount of data. People worry about the system the self-driving cars use being hacked, which could cause chaos and hurt many people. It is true that cyber security is a real concern, and there will always be people trying to steal and incorporate data for bad reasons. However, again we go back to the beauty of machine learning. In the abstract we mentioned the fact that one of the current things machine learning is being used for is just what people are concerned about: data security. Thousands of new malware files are detected every day. However, according to intelligence company Deep Instinct, each piece of new malware tends to have almost the same code as previous versions [12]. According to Deep Instinct only between two and ten percent of the files change from different interactions. With the machine learning algorithms in place, this is barely anything, and the system can predict what files are malware with greater accuracy. The algorithms can also look for various patterns, and can investigate how the data is accessed, and report anomalies that could predict security breaches [10]. As we've mentioned before, machine learning continually causes systems to better themselves, so as even more cyber security issues unfold, the system will become smarter and will be able to overcome most of the possible issues revolving around data security. This means that self-driving cars and their systems will continue to function safely when it comes to possible data security issues.

Our Thoughts

The various types of artificial intelligence make machine learning a viable and practical way to speed the advancement of self-driving cars. Machine learning can gather and assimilate data into these car systems faster than any human programmer could. This will enable the automobile industry to create self-driving cars with extraordinary speed. Machine learning combines large-scale data gathering with its algorithms to make better decisions when encountering complex situations. The ability of

machine learning to program what a car should do when it encounters an unusual situation is more effective than human programming. With machine learning, a self-driving car can learn from a whole database which is gathering massive amounts of data from all self-driving cars in the world to solve a problem one car would encounter at one time. The possibility of getting it wrong and having a car accident is greatly reduced with the existence of such a large database. Despite the advantages of machine learning, many people still fear it will go wrong in some way. In our opinion, those people who are afraid of autopilot are not only fearful of autopilot itself, they are fearful of the potential of having car accidents. Just as we discussed above, the chances of an accident or error are very small, even smaller than a very skilled driver. However, not all accidents can be fully prevented by machine learning. Machine learning can only predict what will happen and make statistical and logical decisions. Sometimes, a driver in another vehicle may act differently than what the computer predicts, or in a way that is inconsistent with what the computer believes based on its massive database. When that happens, an accident may occur.

However, even this small potential of car accidents can be prevented with a simple solution, but it may not be easy to execute. The solution is to have every vehicle be self-driving. When there are no human drivers, the uncertainty that may exist can be eliminated, because every move a self-driving car will make is predictable by another self-driving car. When every car is using the same algorithm and the same database, their concluded decision when encountering a specific situation will be the same. Therefore, the possibility of accidents can be eliminated by the universal application of self-driving systems. Self-driving cars can continue to learn over time from the ever evolving database. Therefore, all models of cars will learn from the same situations and learn from them to make the best decisions. This shared knowledge, when consistently applied across all vehicles will make driving much safer.

Aided by machine learning, autopilot has continued to advance to another level. It does not need constant reprograming by many engineers or programmers. These smart vehicles can self-correct themselves and make the best and safest decision in a wide range of scenarios. With the presence of such a large database, Tesla can pick from millions and millions of scenes to resolve one encounter. The ability to rapidly access massive amounts of data and select the response with the best outcome enables these vehicles to make the right choice every time.

Machine learning truly is a breakthrough in artificial intelligence. It has many uses, one of the most exciting ones being its incorporation in self-driving cars. As the idea and realization of self-driving cars evolve, machine learning is truly the key to success. Without machine learning, self-driving cars require constant reprograming. That can lead to the technology advancements taking longer, which is not acceptable in our constantly changing world of technology. With machine learning, self-driving cars can become safer, more efficient, and just overall better in a shorter period.

WRAPPING UP THE IDEA OF MACHINE LEARNING IN TESLA CARS

Considering all the doubts about autopilot, it is a rapidly emerging field of technology that has the possibility of greatly benefiting human society. Many companies have developed their own algorithms for self-driving vehicles. But Tesla is using machine learning in their car models, which has many advantages. When applied in the field of automobiles, machine learning gives automobile systems the ability to analyze data and learn from it without human programing. The rules are used to create statistical designs for self-driving cars. Machine learning combines large-scale data with its algorithms to make better decisions when encountering complex situations, as well as better navigation over time. Tesla's self-driving cars will be able to deal with road situations that automobiles commonly encounter, and through long periods of learning, the cars will be able to figure out what to do when confronted with complex situations. Not only learning from its own experience, cars with machine learning can take data from what other cars have encountered in a mass database and learn from their situations. All the models of vehicles which can access that database will be able to perform at their best with its help. Based on the reward system of machine learning, vehicles can learn from the right or wrong decisions that have been made and can contribute all the decisions in future encounters. Tesla is using machine learning in their cars to upgrade self-driving technology, and the overall experience which will result in fewer accidents and a safer driving experience. Through the massive shared database of the self-driving car industry, all Tesla's self-driving cars will become smarter though each other than they were before with a very safe database. Machine learning truly is the future of self-driving cars, and due to its power, convenience, and efficiency, it is clear why Tesla decided to implement machine learning into its self-driving product line. Based on the success of machine learning in cars so far, the self-driving car industry has a very bright future.

SOURCES

[1] "Self-Driving Cars Explained." Union of Concerned Scientists. Accessed 03.03.19.https://www.ucsusa.org/clean-vehicles/how-self-driving-cars-work

[2] K. Fehrenbacher. "How Tesla is ushering in the age of the learning cars." Fortune. 10.16.15. Accessed 01.16.19. http://fortune.com/2015/10/16/how-tesla-autopilot-learns/

[3] "What is Machine Learning? A Definition." Expert Systems. Accessed 01.16.19. https://www.expertsystem.com/machine-learning-definition/

[4] L. Tagliaferri. An Introduction to Machine Learning." Digital Ocean. 09.28.17. Accessed 03.03.19. https://www.digitalocean.com/community/tutorials/an-introduction-to-machine-learning
[5] "A Beginner's Guide to Deep Reinforcement Learning." Skymind. Accessed 03.03.19. https://skymind.ai/wiki/deep-reinforcement-learning
[6] C. Giarratana." How AI is driving the future of autonomous cars." Readwrtie. 12.20.16. Accessed 02.02.19 https://readwrite.com/2016/12/20/ai-driving-future-autonomous-cars-tl4/
[7] B. Marr. "The Amazing Ways Tesla is Using Artificial Intelligence and Big Data." Forbes. 01.08.18. Accessed 03.07.19 https://www.forbes.com/sites/bernardmarr/2018/01/08/the-amazing-ways-tesla-is-using-artificial-intelligence-and-big-data/#7203e2242704
[8] J. Qu. "Training Self-Driving Cars using Reinforcement Learning." Towards Data Science. 10.22.18. Accessed 01.17.19. https://towardsdatascience.com/reinforcement-learning-towards-general-ai-1bd68256c72d
[9] W. Knight. "Tesla's New AI Guru Could Help Its Cars Teach Themselves" MIT Technology Review. 06.22.17. Accessed 01.17.19 https://www.technologyreview.com/s/608155/teslas-new-ai-guru-could-help-its-cars-teach-themselves/
[10] S. Abuelsamid, "Autopilot Accidents Show Where Tesla Is Lacking: Engineering to Account for Human Element" Forbes. 07.13.16. Accessed 01.17.19. https://www.forbes.com/sites/samabuelsamid/2016/07/13/tesla-autopilot-accidents-demonstrates-need-to-design-for-robustness-in-autonomous-systems/#1ec272675ae8
[11] L. Pratt. S. Thurn "Learning to Learn" Springer. Accessed 02.17.19 https://books.google.com/books?hl=en&lr=&id=X_jpBwAAQBAJ&oi=fnd&pg=PA4&dq=machine+learning+articles&ots=gVR80UI9zr&sig=RP-RnuKEtVUeon7c_60CzvvVgaQ#v=onepage&q=machine%20learning%20articles&f=false
[12] B. Marr. "The top 10 AI and Machine Learning Use Cases Everyone Should Know About" Forbes. 09.30.16. Accessed 02.17.19 https://www.forbes.com/sites/bernardmarr/2016/09/30/what-are-the-top-10-use-cases-for-machine-learning-and-ai/#4b8736dc94c9
[13] C. Strobl, "Tesla Fleet Learning." HackerBay. 05.13.17. Accessed 03.07.19. https://blog.hackerbay.com/update-teslas-fleet-learning-8e34c3cd6ab4
[14] M. Carter. "Volvo Developing Accident-Avoiding Self-Driving Cars for the Year 2020." Inhabitat. Accessed 03.07.19 https://inhabitat.com/volvo-developing-accident-avoiding-self-driving-cars-for-the-year-2020/?variation=d

ADDITIONAL SOURCES

P. Dar. "An Autonomous Car Learned how to Drive itself in 20 minutes using Reinforcement Learning." Analytics Vidhya. 07.09.18. Accessed 01.17.19. https://www.analyticsvidhya.com/blog/2018/07/autonomous-car-learnt-drive-itself-20-minutes-using-reinforcement-learning/
A. Zingade. "Autonomous Driving using Deep Learning and Behavioral Cloning." Medium. 07.01.18 Accessed 01.17.19. https://medium.com/@akarshzingade/autonomous-driving-using-deep-learning-and-behavioural-cloning-97983a57fe10

ACKNOWLEDGEMENTS

We would like to show appreciation to Janine Carlock, for taking time to read over our work and provide feedback on what we did well, and what we can improve upon.

We would also like to thank Bradley Alderama, for meeting with us as a group to look over our paper and provide feedback based on his experience as an upperclassman engineer.