Autonomous Vehicles

Saahith Janapati

January 2021

1 Goal of this Document

This document aims to provide information regarding the development of autonomous vehicles. I haven't included much technical information here, but instead provided high level overviews on areas of the field which seem to be interesting. Where appropriate, I have included related papers and presentations for those who are interested in learning more.

2 Introduction

Developing self driving vehicles has been a longstanding challenge in AI research. There has been much progress in the field over the past several years due to advancements in machine learning, specifically deep learning. Neural networks have yielded impressive results in computer vision tasks which are important when creating self driving systems. Consequently, many car manufacturers and technology companies have started working on self driving cars; the list includes Tesla, Google, Apple, Uber, Audi, Mercedes, Toyota, and more.

3 SAE Levels

In 2016, the National Highway Traffic Safety Association adopted a classification system created by the SAE (the Society of Automotive Engineers) to better categorize the capabilities of vehicles []. Many people in the field of autonomous vehicles also make use of these classification levels, so it is worth learning about them. I've outlined them here:

- Level 0: No autonomy. The driver is fully responsible for operating the vehicle.
- Level 1: Driver Assistance. Examples of level 1 features include adaptive cruise control, which automatically adjusts speed by looking at the vehicle ahead. Another example of a level 1 feature is an automatic steering system, which keeps a car in its lane. In level 1 systems, the human driver must pay full attention to the road.

- Level 2: Partial Automation. The vehicle is capable of both lane centering and adaptive cruise control. However, the human is still responsible for the operation of the vehicle and must constantly be monitoring what the car is doing. The human must be prepared to take control of the vehicle at any time. Tesla's Autopilot system is currently in this category, but rapidly approaching level 3 and beyond.
- Level 3: Conditional Automation. In this category, the human is no longer driving. The vehicle can operate fully autonomously if the necessary conditions are met. However, the human must still pay attention and be able to drive if the vehicle requests the human to do so (which the vehicle can do at any time).
- Level 4: High Automation. This category is similar to Level 3, except that the human is not required to take control. Vehicles of this category may be geofenced, meaning that they are restricted to driving within a set area (e.g. certain portion of a city). However, if the specified conditions are met, these vehicles can drive fully autonomously without any human intervention. Vehicles of this kind do not need to have steering wheels or pedals. Google's Waymo project is currently Level 4, and is operating within a 100-mile radius of Chandler, Arizona [Hug17].

- Level 5: Full Automation: This is the holy grail. Vehicles of this kind can drive on any road that humans can drive on. (It might be a long time before we reach this level due to roads in developing countries. For example, roads in India are very difficult to navigate. Most humans can't even do it!)

4 LiDAR or no LiDAR?

The use of LIDAR sensors has been hotly debated in the research community. Put simply, LiDAR is a method of measuring the distances of objects by emitting lasers and measuring the time it takes for the lasers to be reflected back to a sensor. The lasers can be in the visible spectrum, but can also be in other portions of the light spectrum. LiDAR sensors can be used to create detailed maps of a certain environment (pictured below).

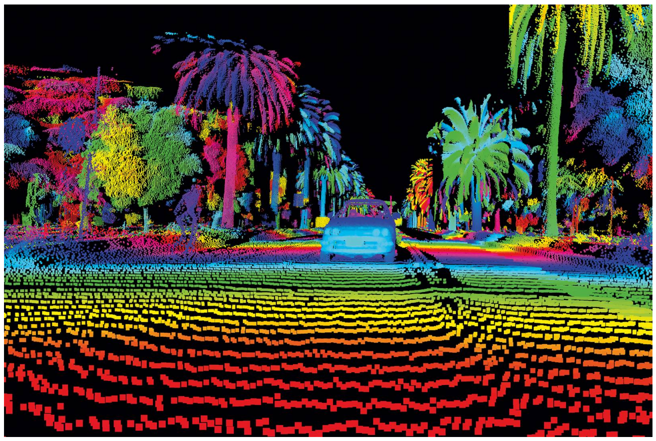

For self-driving vehicles, these maps can be passed to machine learning models to extract relevant information and consequently make a plan for the vehicle to enact.

The main reason why companies are using LiDAR is to aid in the task of depth perception. One drawback to using LiDAR is its cost. A couple of years ago, the most widely used LiDAR sensor for self driving vehicles costed 75,000 dollars. Waymo has recently announced that it has decreased this price by nearly an order of magnitude (10x) bringing the cost down to 7,500 dollars [Kor19]. This is very impressive for several years of work, but 7,500 dollars still constitutes a significant fraction of the cost of most cars on the market today.

However, some companies are choosing not to use LiDAR at all. Although LiDAR can yield extremely accurate depth measurements, there are numerous traditional computer vision techniques (e.g. stereo vision) and newer deep learning approaches (see papers in the footnote[1]) to depth perception that only require photos of an environment and can generate depth measurements for each pixel in the photo.

Another argument for not using LiDAR is that humans only need their vision (and occasionally their hearing, but only rarely) when driving. We do not shoot lasers out of our heads to measure depths of nearby obstacles. Furthermore, roads are already designed with the human vision system in mind (this can be seen in light colors, signs, etc.). Thus, we already have proof that self-driving can be accomplished with vision alone.

[1] See 1) Unsupervised Learning of Depth and Ego-Motion from Video and 2) Depth from Videos in the Wild: Unsupervised Monocular Depth Learning from Unknown Cameras

5 Training with Simulated Environments

Some companies are making use of simulated photo-realistic driving environments to train the driving systems in vehicles. Such companies include Waymo, Tesla, and NVIDIA. While this is certainly an interesting way to train machine learning models, it should be noted that simulations become less useful after a certain point is reached. As Elon Musk said in Tesla's Autonomy Day presentation [19], training with simulations is akin to grading your own homework. Simulations can only improve models in scenarios that human programmers have already imagined and incorporated into the simulation. Simulations fail to capture to a full extent the real-life complexities that are present on roads. Creating a simulation that does capture these complexities is probably a harder problem than creating autonomous vehicles themselves). Thus, while simulated environments are useful for training, they cannot be substituted for data acquired from driving on real roads. Pictured below is a visualization of Tesla's simulation environment.

6 Tesla[2,3] Autopilot.

I've chosen to discuss Tesla's Autopilot because it is currently the most advanced semi autonomous system deployed at scale. I think Tesla has a unique position compared to most other players in the self driving industry because of the number of cars they have on the road. Since 2015, Tesla has made impressive strides improving their Autopilot system. It started off as an advanced adaptive cruise control feature but has rapidly improved to also change lanes, navigate on highways, stop at lights, and drive completely autonomously in parking lots towards the driver (Smart Summon). The system is currently classified as Level 2, but is very close to reaching Level 3. Tesla's approach is mainly based on vision and does not make any use of LiDAR. They make extensive use of convolutional neural networks (CNNs) to process input from onboard cameras. They have various neural networks for many different tasks, such as detecting pedestrians,

[2] Interesting Fact 1: Tesla uses PyTorch to train its neural networks.

[3] Interesting Fact 2: Andrej Karpathy, who was the lead lecturer for the first class of Stanford's CS231n, is now the Senior Director of AI at Tesla

finding lane markings, and estimating depth. [4] Furthermore, Tesla has a huge advantage compared to other companies due to the amount of data they collect; all new Tesla vehicles have 8 cameras that transmit data that can be used to improve the aforementioned networks. Furthermore, Tesla has the capability to send over-the-air software updates to its entire fleet, a capability that few other car manufacturers have. Thus, as the neural networks improve, Tesla can make all of its customers safer overnight.

Another interesting aspect of Tesla's approach is that they make their own computing hardware. The Tesla FSD (Full Self Driving) computer is a system dedicated to performing convolution operations to speed on-board forward propagation. Prior to FSD, Tesla used a system based on NVIDIA's (a computer graphics company) Drive PX2 platform for neural network computation. The new FSD computer can perform forward propagation on video frames 21 times faster and is only 80 percent of the cost compared to the Drive PX2 based system. Customers now have the option to purchase Tesla vehicles with this computer (pictured below) installed.

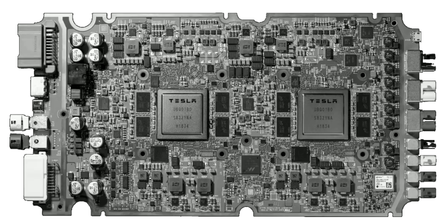

Considering all these facts, Tesla appears to be in a good position with regards to self driving capabilities. Even if there are currently more advanced systems on the road (which is itself debatable), Tesla's existing large fleet with advanced hardware and its ability to send over-the-air software updates will likely allow it to dominate the autonomous vehicle industry.

6.1 Tesla Full Self Driving (FSD)

Over the past several months, many Tesla owners have been able to drive with Tesla FSD Beta installed. I think this is probably the closest system there is to Level 5 in the industry. Here is an early video of FSD beta in action:

`https://www.youtube.com/watch?v=3CR7I1_lOUM`

7 Economic and Societal Impact

The advent of autonomous vehicles will have a profound impact on economies and society at large. The most direct economic impact will be felt by those employed driving such as taxicab drivers and truck drivers. There are currently

[4]You can learn more about the specific technical details of Tesla's autopilot system by watching the following videos (which are on YouTube): 1) Tesla Autonomy Day and 2) PyTorch at Tesla - Andrej Karpathy, Tesla.

3.5 million Americans employed as truck drivers; that's more than 10 percent of the entire U.S. population. As autonomous trucks enter the roads, it will become economically disadvantageous to continue to operate regular trucks. Thus, truck drivers (and similarly taxicab drivers) will have to find work elsewhere.

Additionally, when autonomous vehicles reach SAE levels 4 and 5, robotaxi fleets can operate. If you have an autonomous vehicle, it is not economically optimal to leave it idle in a parking lot while you are at school or work. Instead, you could send it to operate in a robotaxi fleet that would operate similarly to how Uber operates today, except without the drivers. The vehicle could pick customers up and drive them to their desired location completely autonomously, generating positive cash flow for the owner of the car while providing convenience to the customer. When the owner wants the car back, he or she can simply remove it from the fleet. When self-driving vehicles become more popular, the price of such a ride may become cheaper than riding on mass transit systems like subways and buses. While this might seem like science fiction, Tesla already has plans to implement such a system in the near term future, pending regulatory approval.

8 Autonomous Weapons

Although the main focus of this document is self-driving cars, I wanted to include this section because it's pretty important. Many militaries already use partially autonomous weapons. For example, the U.S. military has made extensive use of Predator drones in the Middle East.

While these are usually remotely piloted by humans, the increasing use of these machines is worrisome. It's not hard to conceive of fully autonomous weapons that are capable of identifying and engaging targets without any human input in the not too distant future. Many AI researchers and technologists have voiced their concerns against development of such weapons. While such weapons would save the lives of soldiers, they would also lower the barrier-of-entry to war. If one country is able to develop such technologies, there would inevitably be an arms race. And if terrorists are able to acquire such systems, a lot of harm could be done. Consequently, in 2015, the Future of Life Institute (FLI), a nonprofit organization that focuses on mitigating existential risks to humanity, announced an open letter condemning the development of such weapons. The letter has been signed by many renowned AI researchers and other influential individuals. We should pay close attention to news regarding developments in this area so bad things don't happen.

Now, onto the main focus - autonomous cars!

References

[Hug17] Justin Hughes. *Alphabet's Waymo Is Already Running Level 4 Self-Driving Cars in Arizona*. Nov. 2017. URL: `https://www.thedrive.com/tech/15848/waymo-is-already-running-cars-with-no-one-behind-the-wheel`.

[Kor19] Kirsten Korosec. *Waymo to start selling standalone LiDAR sensors*. Mar. 2019. URL: `https://techcrunch.com/2019/03/06/waymo-to-start-selling-standalone-lidar-sensors/`.

[19] *Tesla Autonomy Day*. 2019. URL: `https://www.youtube.com/watch?v=Ucp0TTmvqOE`.

[] URL: `https://www.sae.org/news/2019/01/sae-updates-j3016-automated-driving-graphic`.