

Analysis of Sensor Attacks Against Autonomous Vehicles

Søren Bønning Jakobsen, Kenneth Sylvest Knudsen and Birger Andersen^a

DTU Engineering Technology, Technical University of Denmark, Ballerup, Denmark

Keywords: Autonomous Vehicles, LiDAR, Camera, Sensors, Attacks, Countermeasures, Security, Multi-Sensor Fusion.

Abstract: Fully Autonomous Vehicles (AVs) are estimated to reach consumers widely in the near future. The manufacturers need to be completely sure that AVs can outperform human drivers, which first of all requires a solid model of the world surrounding the car. Emerging trends for perception models in the automobile industry are towards combining the data from LiDAR and camera in Multi-Sensor Fusion (MSF). Making the perception model reliable in the event of unforeseen real world circumstances is tricky enough, but the real challenge comes from the security issue that arises when ill-intentioned people try to attack sensors. We analyse possible attacks and countermeasures for LiDAR and camera. We discuss it in context of MSF, and provide a simple framework for further analysis, which we conclude will be required to conceptualise a truly safe AV.

1 INTRODUCTION

Due to advances in technology in recent years, autonomous vehicles are becoming more and more realistic on public roads. Car manufacturers have already dealt with extremely complex challenges such as self-navigation and collision prevention. Some places in the US are even beginning to allow self-driving cars under specific circumstances (Jones, 2022). There are different levels describing how autonomous a car is (Petit, 2022), and the aforementioned cars are examples of category 3 autonomous cars. Today, the race between manufacturers to become the first company to release a fully autonomous category 5 car continues. Considering the end goal means giving the AV complete control with no oversight, it is important during testing, to make sure that the consumer is completely secure. This means that the AVs will have to withstand not only rigorous testing, but also the challenges that arise in the real world, where people might intentionally try to cause crashes.


Researchers, as well as white-hat hacker groups, have been conducting tests on these cars and their flaws. For example Keen Security Labs in China demonstrated flaws in security in a Tesla Model S, that allowed them to remotely hack into the car, and make it change to the reverse lane (Tencent Keen Security Lab, 2019). Another example is found in (Eykholt et al., 2018b), where researchers managed to confuse the perception algorithm to have a stop sign

classified as a speed limit sign using stickers.

In this article, we will analyse attacks, specifically attacks aimed at the AV's ability to perceive the world around it through sensors. Considering all the different technologies involved, the complexity of attacks on AV's sensors vary wildly. We will focus on remote attacks on the physical sensors, but also the underlying algorithms working on the sensor data. We assume that attackers have no access to the vehicle, yet aim to attack the perception model of the AV, since any attacks here will cascade into all other decision-making that an AV does.

2 HOW AN AV MODELS ITS SURROUNDINGS

There are many variables and obstacles on public roads like pedestrians, other cars, turns, traffic lights, bicycles and much more. Before an AV can safely drive on public roads, the car first and foremost needs to see better than humans in order to drive better. This has been a major hurdle for development. By combining different sensing technologies, developers have created detection systems which can "see" better than human eyes (Burke, 2019). The common technologies in AVs used for mapping the environment are LiDAR (Light Detection and Ranging), radar (Radio Detection And Ranging) and cameras. Radar and LiDAR are somewhat interchangeable, as they offer a lot of the same information with different pros and

^a  <https://orcid.org/0000-0003-1402-0355>

cons. This has led to debates in the industry about which is best. Currently, Google, Uber and Toyota all rely heavily on LiDAR, while Tesla are the only real advocates for Radar (Neal, 2018). Because of this, our focus will be on camera and LiDAR.

2.1 Camera

AV perception is achieved by many sensors and sensor systems. One of the first sensors used was the camera. For AV the camera is used to visualize its surroundings. The camera is used for lane detection, horizon/vanishing point detection, object detection and tracking of vehicles and pedestrians, traffic sign recognition and headlight detection, as demonstrated by the colored boxes in Fig. 1.

Cars can have multiple cameras covering a 360 degrees view of their environment. Cameras are very efficient at detecting texture of objects. For implementation, cameras are more affordable than LiDAR and radar sensors. The high pixel quality obtained by the camera comes with a price of computation power. Today cameras can take pictures with millions of pixels in each frame and about 30-60 frames each second. Each of these frames needs to be processed in real time in order for the car to make real-time decisions, this requires lot of computational power (Kocic et al., 2018). Image quality is important for the system to classify the objects. The quality can be affected by lenses, windshield, vibration, environmental conditions like snow, rain, fog and light. All these image disruptions can result in unnoticed objects and increasing image correction processing time. Some cars use a multi-camera setup where some cameras overlap each other (Petit et al., 2015).

The camera creates a good representation of the environment, however the depth perception is not nearly as good as that of other sensors, which is why LiDAR technology is used.



Figure 1: An image of an autonomous car using LiDAR as distance measure and camera as object detection (Cameron, 2017).

2.2 LiDAR

The LiDAR sensor fills the existing gap between radar and camera sensors (Kocic et al., 2018). LiDAR works by emitting pulses of infrared light and measuring the time taken to reflect on distant surfaces. These reflections return a point cloud that represents objects from the environment. Most common LiDAR lasers use light in the 900 nm wavelength, longer wavelengths will perform better in poor conditions, such as fog and rain. Because the LiDAR sensor has a more focused laser beam it can create a more dense point cloud, resulting in a high resolution map of the environment (Roriz et al., 2022). Precision is important in LiDAR systems, as lower precision LiDAR sensor originate noisy point clouds. How precise a LiDAR is says how close the estimated point is compared to a point in the real world.

There are different architectural techniques for creating the surroundings of the car. They can be categorized in different groups of spinning and non spinning (solid state). To get a horizontal 360 view, LiDAR sensors can be combined with a mechanical part to spin around while measuring the distance of the surrounding objects, as shown in Fig. 1. This is the most common LiDAR application currently. If LiDAR technology is used but there are no moving parts, this is called solid state sensors. With solid state sensors we get a more narrow angle but usually they are cheaper.

2.3 Multi-Sensor Fusion (MSF)

Multi-sensor fusion is a technique where the input from multiple sensors are combined, in order to leverage the best of both inputs. While it is possible to combine several sensor types, current trends in the automobile industry have gone towards combining camera and LiDAR (Kocic et al., 2018). This minimises hardware complexity, as only two sensor types are involved, and the information from these complements each other nicely. The obvious advantages here is the detailed vision of the camera, allowing for object classification, and LiDAR for accurate object detection and detailed range measurements.

MSF have been a major factor in helping researchers make reliable models of an AVs surroundings. Researchers in (Kocic et al., 2018) cites application of MSF in the detection of objects, grid occupancy mapping for placing these objects in a model around the vehicle and lastly for tracking the objects movements within the model. Their example of an MSF algorithm is the PointFusion network, used for 3D object detection. This algorithm achieves sensor

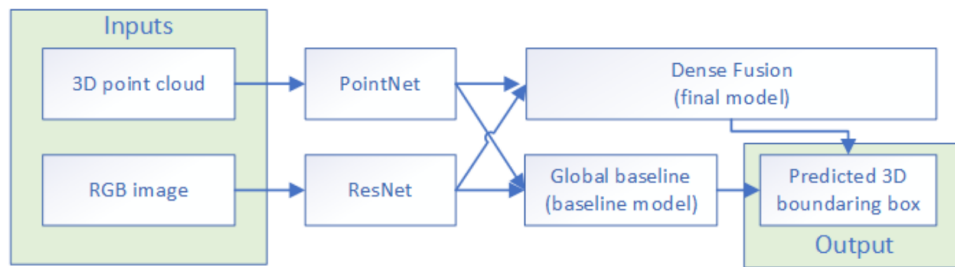


Figure 2: Overview of the neural networks involved in the PointFusion MSF algorithm (Kocic et al., 2018).

fusion by processing each sensor’s data with a different Neural Network (NN), and then feeding the representations into a new neural network, achieving high-level fusion as shown in Fig. 2. As seen, the separate NNs are PointNet and ResNet handling the pointcloud and RGB (Red-green-blue) image respectively. Their results are fed to dense fusion algorithm that for each input point predicts the spatial offset of the corners relative to the input (Xu et al., 2018). The PointNet and ResNet information is also fed to a baseline model that directly regresses the box corner locations. Together, the dense fusion predictions and baseline model results in the predicted 3D boxes, that an AV would have to navigate it’s way around.

While this is just one way of doing this, there are many ways to go about it. It really comes down to balancing the complexity of the mathematics, the computational power in the vehicle and other factors to design the best sensor fusion algorithm.

3 ATTACKS AND COUNTERMEASURES

In this section we will explain different attack types, documenting examples of successful attacks and suggesting possible countermeasures. Since our focus is remote attacks, our assumption of possible setups to perform these attack will be categorized into one of two types, described in (Petit et al., 2015).

Front/Rear/Side Attacks which involve an attacker who installs hardware in his own vehicle to perform an attack. This allows the attacker to keep the hardware within distances of a target vehicle for longer time.

Roadside Attacks which involve a mounted stationary setup, that allows for greater precision. This type of attack is not limited to one installation point, but can have several if need be.

This section summarises results in Table 1 and Table 2.

3.1 Attacks on Cameras

3.1.1 Blinding Attack

Cameras have built in functions to decide how much light is let into its shutter, in order to take better photos in all light conditions. This attack type abuses this function by shining a strong light source into the camera, in order to either completely blind the camera, or partially blind it, causing it to miss objects. In experiments performed in (Qayyum et al., 2020), researchers managed to completely blind a camera for up to 3 seconds by pointing a laser light directly at the camera, and caused irreversible damage by doing so for several seconds from less than 0.5 meters away. They also managed to cause partial blindness by directing a LED (light emitting diode) light matrix at the camera, inhibiting the object recognition and proving that there are multiple tools to perform this attack. The researchers performed the testing in stationary setups, but they also simulated a front/rear/side attack by wobbling the laser at the camera, which was still successful, though the blinding was less effective.

3.1.2 Adversarial Example

While not physically attacking the camera, this attack type attacks the way the information from the camera is processed. The perception models based on machine learning (ML) and deep learning (DL) have proven vulnerable to carefully crafted adversarial perturbations. Generally, Adversarial Examples (AEs) are classified as ‘Appearing Attacks’ and ‘Hiding Attacks’. There have been published several attacks of this type, where most are evaluated on stop signs as they are a critical part of decision making in driving. Researchers in (Eykholt et al., 2018b) executed such a hiding attack, where they designed specific black and white stickers that caused misclassifications of stop signs. In order to fool any human onlookers, they designed the stickers to look like graffiti and still managed to make the sign be classified as a speed limit sign in 87,5% of the tests. In the opposing category, (Eykholt et al., 2018a) managed to make innocent

Table 1: Overview of attacks.

Attack type	Target sensor	Method	Impact	Feasibility
Blinding attack	Camera	Blinding the camera with some kind of light source, making the camera unable to guide the Vehicle	Low	Easy
Adversarial Examples	Camera	Introduce objects with adversarial perturbations, to confuse perception model	High	Hard
Spoofing attack	LiDAR	Relaying light pulses in a different position, creating fake obstacles.	High	Medium
Saturation attack	LiDAR	Jamming or blinding LiDAR sensors by emitting strong light in same wavelength as the LiDAR sensor.	Low	Easy

looking stickers be classified as stop signs. There are other examples of real world applications, such as (Cao et al., 2021), where researchers 3D printed a traffic cone that was ignored by cameras, and (Zhou et al., 2020) who designed a billboard that causes malfunction in the steering angle of AVs. This attack type appears mostly as a roadside attack.

3.2 Attacks on LiDAR

3.2.1 Spoofing by Relaying Attack

LiDAR sensors are what is called active sensors (Shin et al., 2017). This means that the LiDAR sensor to detect an object emits light intentionally from its own sensor and then listen to the echo. Because the speed of light is constant the LiDAR sensor can calculate the distance by measuring the ping time of the signal. In spoofing attacks the attacker uses this signal created by the victim but relays it back from a different position. The goal of spoofing is to deceive the victims LiDAR sensor and to create fake point clouds. The creation of fake points could potentially cause the AV to make sudden erroneous decisions.

Active LiDAR sensors use particular waveform to differentiate echoes from the other inbound signals. This means before the attacker can perform the attack, he needs to obtain the ping waveform. When the waveform of the victims LiDAR sensor is obtained, the attacker can now perform the attack by relaying the signal back to the LiDAR sensor. This attack is effective because the victim car has a hard time to distinguish what signals are real or fake and is unaware of the attack, potentially providing seemingly legitimate but actually erroneous data. There are two ways to perform this attack. One way is to place the attacking device on the roadside and then aiming towards the victim lane (Cao et al., 2019), however it could also be performed as a front/rear/side attack, using computer vision to keep track of the vehicle and to aim precisely. This would however be significantly more difficult to execute. First some of LiDAR sen-

sors used in cars are spinning LiDARs. This means that the victims LiDAR sensor has to be facing the attacking direction. In order for the attack to work the echos has to hit victims sensor at the right receiving angle. The second problem is that LiDAR will only accept echos within a certain delay time. Because of the delay threshold the distance between attacker and victim has a big influence on the attack window. Article (Petit et al., 2015) performs a spoofing attack, and relates timing to the success of spoofing.

3.2.2 Saturation Attack

Typically sensors has a lower and upper bound for input signals (Shin et al., 2017). If signals arrive at the sensor with a low signal power, the sensor will ignore the signal. When the signal increases it eventually exceed the upper threshold, at this point the sensor cannot reflect the input changes well. Therefore the principle of saturation is to expose the target sensor to a signal of high, making the sensor unable to work properly. Because the sensor is unable to receive any new signals while under a saturation attack, this attack is a DoS (Denial of Service) or a blinding attack, since it uses light as a medium.

Saturation attacks are powerful because they are unavoidable, and though it is easy to detect it cannot be prevented. Human drivers and pedestrians will be unaware of the attack because LiDAR operates in eye safety wavelength namely infrared light. In (Shin et al., 2017) they performed saturation attack with both weak and strong light sources. The experiments made was successful and had different outcome for the weak and strong light. When the attack was performed with weak light they observed randomly located fake dots, while with the strong light source, they observed that the sensor became completely blind in a sector of Field of View (FoV).

3.3 Countermeasures on Cameras

3.3.1 Redundancy

In (Petit et al., 2015) argues the simple benefit that more cameras with overlapping view will at least make a blinding attack harder to execute. In (Qayyum et al., 2020) researcher emulated a handheld blinding attack by laser, which would be very hard if several cameras were present. The argument for more cameras could also be supported by the findings in (Liu and Park, 2021), which suggests an algorithm for recognizing perception error attacks in MSF, which requires stereo cameras. This will be further elaborated on in section 3.5. Though this solution seems simple, it is important to remember the extra space and cost associated with it, which are very essential factors in the highly competitive automotive industry. There is also an argument that integrating more cameras introduces more complexity in synchronizing the capturing of frames and maintaining the same exposure (Petit et al., 2015), though this should easily be overcome with today's technology.

3.3.2 Optics and Materials

In (Petit et al., 2015) argues that removable on-demand near-infrared-cut filter, a feature commonly found in security cameras, could serve as a defence against blinding attacks. They argue that such a defence would only be usable during daytime, as the filter would have to be removed during nighttime, in order to make use of infrared light for night vision. Intelligent applications of this countermeasure should be considered, perhaps by implementing software with thresholds for when it considers the camera to be under attack.

Another defensive material is photochromic lenses, which is a type of lenses that change color to filter out specific types of light. Several types of lenses or coating of lenses could be considered, but as an example vanadium-doped zinc telluride will turn more opaque when hit by high-intensity beams, automatically filtering these without effecting image quality in low light conditions (Petit et al., 2015). Once again the hardware and development costs should be considered.

3.3.3 Making AE Robust Perception Models

Researchers in (Qayyum et al., 2020) provide a very thorough examination of how to make adversary robust ML/DL solutions and their efficiency. This is a very wide and highly technical topic, which we consider outside the scope of this article. We will

therefore refer the interested reader to the article, and quickly mention solutions in simple concepts, as well as the conclusions.

Suggestion is re-training a given classifier on images including adversarial attacks, although this approach is easily criticized for being reactive and vulnerable to attacks that simply generate new attack types. Also is suggested training an auxiliary model, whose sole purpose is to detect features commonly found in pictures with AEs and classify a frame as an outlier if it contains it. Conclusion is that there are multiple directions that solutions can go, and that they often can be combined. More research is needed in order to facilitate a solid defense, but it is feasible, and a valuable addition to make safer AVs.

3.4 Countermeasures on LiDAR

3.4.1 Saturation Detection

As described in section 3.2.2 saturation can easily be detected by the sensor system. A victim vehicle could have an inbuilt fail-safe mode, so when the car detects saturation, it slows down the car and pull to the side (Shin et al., 2017). This countermeasure could on a crowded road lead to a dangerous situation. If the car has to pull to the side while having a jammed LiDAR sensor, it is like a person driving with closed eyes.

3.4.2 Redundancy and Fusion

One countermeasure could be by having multiple LiDARs overlapping some FoV angles (Shin et al., 2017). With redundant LiDAR sensors, the victim car could under saturation attacks abandon input from the attacked sensor until the attack is over. Though the car knows when it is exposed to saturation, it is significantly harder to detect spoofing. The redundant setup will still work better against spoofing by cross validating the malicious points. If the attacker creates fake points in the non-overlapping zone, redundancy will have no effect. LiDAR sensors are expensive so using multiple sensors will increase the overall cost a lot. This solution is not bullet proof, because the attacker is still able to attack multiple sensors at the same time. Another option proposed in (Petit et al., 2015) is to take advantage of data intercepted by neighboring AVs. Victim vehicle could cross-validate its data with neighboring data to observe inconsistencies. This method only works if there are other vehicles on the road. Vehicle to Vehicle (V2V) solution opens up for more hacking opportunities because one neighboring vehicle could share incorrect data or being tool of the attacker.

Table 2: Overview of countermeasures.

Countermeasure	Sensor	Method	Preventing attack	Feasibility
Redundancy	Camera	Adding more camera with a significant overlap in view	Blinding attack, assists MSF attack detection	High
Optics and materials	Camera	Adding smart materials, that can filter out harmful light	Blinding attack	High
Making AE robust perception models	Camera	Using advanced techniques to make AE attacks more difficult	Adversarial attack	Low
Saturation detection	LiDAR	Builtin Fail-safe mode. Under attack slows down and pull to the side.	Saturation	High
Redundancy and Fusion	LiDAR	Multiple LiDAR setup with overlapping FoV. By comparing input from multiple overlapping sensors it is possible to detect and prevent some attacks	Saturation and spoofing	High
Random probing	LiDAR	By randomizing LiDAR pulse interval makes spoofing very difficult to perform	Spoofing	Medium
Side-channel authentication	LiDAR	Using side-channel information as authentication. Authentication makes it very difficult to spoof a LiDAR, not knowing the secret key	Spoofing	High
Multi-Sensor Fusion	LiDAR and Camera	Use smart MSF to model both LiDAR and camera data, checking for inconsistencies	Camera blinding, LiDAR Spoofing, saturation, relay and rotation	Medium

3.4.3 Random Probing

When making a spoofing attack on LiDAR sensors, a hacker will be interested in the pulse interval. This interval is the timing for when the attacker needs to fire back attacking pulses (Deng et al., 2021). By randomizing the interval it makes it hard for the attacker to synchronize the attack. This method is problematic for spinning LiDAR systems, as they require a constant rotation speed and angle of transmission needs to be known (Petit et al., 2015). Another option here is to skip some pulses, as this will only require some software modification. When the sensor skip a pulse it is still able to listen to incoming pulses, making it possible to detect possible spoofing attacks. If the sensor is skipping some of the pulses, it has to run with a higher rotation speed, to keep the same resolution (Shin et al., 2017). It is important that the skipped pulses are chosen in a secret pseudo-random fashion, so attacker cannot predict the skipped pulse.

3.4.4 Side-Channel Authentication

To understand this countermeasure it is essential to know what is meant by side-channel information. Side-channel information is physically leaked information, which could for example be power consump-

tion or electromagnetic radiation. The suggested side-channel information in (Matsumura et al., 2018) comes from a cryptographic device in the car. The device is making heavy calculations using AES (Advanced Encryption Standard) on a cryptographic key, and the electromagnetic radiation during these calculations are read. This information is then used to modulate and demodulate the amplitude of the laser. It will then only accept returning echoes with exactly this modulation. Though feasible, it becomes very difficult for the attacker to send fake echoes with the correct modulation, and the car can simply change the cryptographic key once in a while to have varying side channel information.

3.5 Countermeasures via MSF

Evaluating countermeasures via MSF can be difficult, as it is an emerging research topic with many approaches mathematically. As such, we will base this section on (Liu and Park, 2021), a newly released paper with meta reflections and criticisms of the current state of MSF algorithms, as well as a suggested new approach.

An important point raised, is that most of the aforementioned attacks are evaluated on a single sen-

sensor type. They argue that today's AVs do not build their model of their surroundings based on a single sensor type, but rather through the combination of data through the use of MSF. This immediately raises the abstraction level of the discussion, as physical attacks on sensors needs to be evaluated based on the whole perception system. As an example, they criticise the design of some MSF algorithms, as their design seems too focused on working in non-adversarial settings. The algorithm F-PointNet uses a cascade approach to fuse LiDAR and camera data, by generating 2D proposals on the image data, and then projecting these onto 3D space, refined by the LiDAR data. This makes it especially vulnerable to camera attacks, as the detection failures accumulate through to the LiDAR steps of the fusion algorithm. The researchers conclude that any MSF built around the idea of projecting either LiDAR or camera data onto the other, will be significantly more vulnerable to attacks at the sensor considered to be the 'primary'.

To combat the issues, the researchers designs their own sensor fusion algorithm, which uses CV and ML to map features on both the camera and LiDAR data, and analyzing any features that cannot be mapped to both. The results of their design for camera attacks yields a 100% detection rate, and for LiDAR attacks their detection rate for spoofing and saturation are 97% and 96% respectively. Thus proving that MSF can be leveraged for a significant defence against attacks on the sensors providing the data. They do however make a point, that AEs attacking sensory data processing algorithms of camera or LiDAR data would not be detected through their algorithm, as they would not appear as sensor malfunctions.

4 DISCUSSIONS

There is not doubt that AVs will arrive in the near future, but it is clear that one of the major hurdles is not just beating the difficult tasks of modeling the surroundings and navigating them, but also making the perceived information resistant to attacks. As noted in (Jones, 2022) a recent study showed that only 14% of drivers trust an AV to do all the driving, while 54% are too afraid to try it and 32% are unsure. Convincing people will require delivering a product that completely delivers on all safety measures, before public opinion deems AVs to be too dangerous. In Table 1 and Table 2 we have summed up our attacks and countermeasures, but it is still difficult to concretely say that a consumer is sufficiently safe. How do you come up with guarantees in a field that moves so fast in so many directions?

In (Qayyum et al., 2020) they discuss the possibility of threat modeling, as a fundamental approach to safety analysis. Here they mention several important aspects including, but not limited to, *Adversarial Knowledge*, *Adversarial Capabilities* and *Adversarial Specificity*.

Adversarial Knowledge refers to the required knowledge for executing an attack. Typically, adversarial attacks are referred to as either white-box, gray-box, or black-box. White-box assumes that the attacker has full knowledge of the underlying systems, be it hardware or software. Gray-box assumes partial knowledge of the underlying mechanisms and black-box assumes no knowledge, to the point where they might not even know which ML algorithm the perception model they are trying to attack is built on. This is arguably the most important dimension to consider, especially if you factor in who is the perceived target. With the increased tendency towards cyber-warfare that we are seeing internationally, high-priority targets like heads-of-state can expect to be targets of incredibly sophisticated attacks that regular people would have no need to fear.

Adversarial Capabilities defines the assumed capabilities of the attacker, which is important scope to consider. This leans itself towards knowledge as well, but also economic as in (Petit et al., 2015), where they specifically focus on attacks requiring only commodity hardware. They do this under the assumption, that the attacks requiring the least economic and technical knowledge will be most common, an argument that can definitely be extended to our definition as well.

Adversarial Specificity means how specifically targeted an attack is. This could be a consideration of whether or not the laser damage to the cameras tested in (Petit et al., 2015) would translate to the damage on other cameras, or as mention in (Qayyum et al., 2020), that black-box AEs for one ML/DL model are assumed to affect other models trained on datasets with a similar distribution as the original one. It is in the interest of car manufacturers, that attacks do not generalise well across hardware and software.

With this information in mind, one can start reflecting on the future of safety precautions, though nothing is set in stone. We can expect that simple hardware attacks will happen, be it blinding of cameras and attempts at tricking the LiDAR, considering just how low knowledge, capability and specificity requirements are especially for the camera attacks. Whether it is smartest to adopt some of the novel approaches suggested in sections 3.3 and 3.4, or trust in the higher level protection of MSF that checks for physical attacks is hard to tell. As with many IT-security questions, the answer might be somewhere in




Figure 3: Overview of the suggested pipeline for the data (Jahromi et al., 2019).

middle by combining both.

State-of-the-art researches on MSF have proven that it can be a valuable tool to detect attacks on sensors (Liu and Park, 2021), but one should not see this as an excuse to leave them wide open. They note that attacks are still feasible, however, they would need to attack both camera and LiDAR simultaneously and gradually, as sudden shifts would be detected. This raises the knowledge and capability requirements enormously, to the point where one could argue that any attacker with such capabilities should just ram another car into its victim to get the same results with way less effort.

IT-security often becomes an arms race, as researchers have already tested attacks through adversarial examples, that trick both camera and LiDAR (Cao et al., 2021). This should beat any MSF that cross-validates, as none of its sensors will detect the object. Researchers in (Liu and Park, 2021) cite exactly adversarial examples as a weakness, and though this attack appears to have had huge knowledge and capability requirements to create, the finished 3D model could theoretically be sold on a black market for very little. This would quickly lower the availability of the attack to anyone with a 3D printer.

It is clear that an attack like above certainly complicates things, and it will not be the last to do so. Though the complexity of the topic might seem disheartening, we have no doubts that further research will help increasing safety, to a point where it is safe enough. Together, researchers will have to come up with a definition of what "safe enough" even means, and only the future can tell whether or not it will ever be completely safe.

5 FUTURE WORK

Considering the already mentioned importance of safety and security, it will be necessary for the automobile industry to adopt some sort of safety standards for the AVs perception models, just as one would consider standards for seat belts. As the topic we are trying to concretize are quite more complex than seat belts, there is a dire need for a solid test environment,

e.g., it needs to be measurable for how many frames an AE actually tricks a camera, and whether or not angle has any influence.

We are working on creating a framework which can help to define these standards. Instead of testing with an expensive AV, we plan to build a model of an AV by means of following hardware components:

- Leo Rover Developer Kit 2, to have something driveable to mount the sensors on
- A sufficiently high resolution camera
- A LiDAR sensor
- Signal processing and computational hardware

Our suggested pipeline of the data follows that of (Jahromi et al., 2019) and considers the processing tasks as follows: First camera video frames and the depth channel information from the LiDAR are sent to a Deep Neural Network (DNN) to do the object detection and road segmentation, since DNNs such as a Fully Convolutional Network (FCN) have shown better accuracy for computer vision tasks compared to ML (Jahromi et al., 2019).

The second step is to overlaying the point cloud from the LiDAR with the fused data, before feeding the output to the control layer of the AV. The entire pipeline of the data can be seen in Fig. 3. This will set a foundation for testing possible attacks to see if the AV drives as expected or deviates.

6 CONCLUSION

We are moving towards a future, where we will soon see driverless cars available to the common consumer. It is necessary to secure the safety and well-being of the consumers, as well as earn their trust. To do this, the AVs need to outperform human driving, which first and foremost require that the car can build a reliable model of its surrounding, before the higher level algorithms can navigate the AVs. This requires sensor data that are valid and sensors resistant to attacks, so adversaries cannot manipulate them to cause accidents. We have given our resume of the literature concerning attacks on cameras and LiDAR, two major factors in the current perception models, and also the

two sensors most often combined together in MSF. We have discussed these and their countermeasures in depth, before reviewing them in the context of MSF.

We have opened a discussion into attack complexity and suggested a framework to review them in to better grasp the challenging factors of this issue. It is clear that more countermeasures are needed, however advancements in MSF are looking very promising, and some of the state-of-the-art solutions will be a huge leap in making attacks too complex to be feasible. It is still important to keep in mind that since everyone, including heads-of-state, will be driving around in these cars, complexity alone cannot be seen as a sufficient. Solid standards need to be adopted.

REFERENCES

- Burke, K. (2019). How Does a Self-Driving Car See? <https://blogs.nvidia.com/blog/2019/04/15/how-does-a-self-driving-car-see/>, Accessed 2023-01-19.
- Cameron, O. (2017). An Introduction to LIDAR: The Key Self-Driving Car Sensor — by Oliver Cameron — Voyage. *Voyage*.
- Cao, Y., Wang, N., Xiao, C., Yang, D., Fang, J., Yang, R., Chen, Q. A., Liu, M., and Li, B. (2021). Invisible for both Camera and LiDAR: Security of multi-sensor fusion based perception in autonomous driving under physical-world attacks. In *Proc. - IEEE Symp. Secur. Priv.*, volume 2021-May, pages 176–194.
- Cao, Y., Zhou, Y., Chen, Q. A., Xiao, C., Park, W., Fu, K., Cyr, B., Rampazzi, S., and Morley Mao, Z. (2019). Adversarial sensor attack on LiDAR-based perception in autonomous driving. In *Proc. ACM Conf. Comput. Commun. Secur.*, pages 2267–2281.
- Deng, Y., Zhang, T., Lou, G., Zheng, X., Jin, J., and Han, Q. L. (2021). Deep Learning-Based Autonomous Driving Systems: A Survey of Attacks and Defenses. *IEEE Trans. Ind. Informatics*.
- Eykholt, K., Evtimov, I., Fernandes, E., Li, B., Rahmati, A., Tramèr, F., Prakash, A., Kohno, T., and Song, D. (2018a). Physical adversarial examples for object detectors. In *12th USENIX Work. Offensive Technol. WOOT 2018, co-located with USENIX Secur. 2018*.
- Eykholt, K., Evtimov, I., Fernandes, E., Li, B., Rahmati, A., Xiao, C., Prakash, A., Kohno, T., and Song, D. (2018b). Robust Physical-World Attacks on Deep Learning Visual Classification. In *Proc. IEEE Comput. Soc. Conf. Comput. Vis. Pattern Recognit.*, pages 1625–1634.
- Jahromi, B. S., Tulabandhula, T., and Cetin, S. (2019). Real-time hybrid multi-sensor fusion framework for perception in autonomous vehicles. *Sensors (Switzerland)*, 19(20).
- Jones, D. (2022). California Approves A Pilot Program For Driverless Rides : NPR. <https://www.npr.org/2021/06/05/1003623528/>
- california-approves-pilot-program-for-driverless-rides, Accessed 2023-01-10.
- Kocic, J., Jovicic, N., and Drndarevic, V. (2018). Sensors and Sensor Fusion in Autonomous Vehicles. In *2018 26th Telecommun. Forum, TELFOR 2018 - Proc.*
- Liu, J. and Park, J. (2021). "Seeing is not Always Believing": Detecting Perception Error Attacks Against Autonomous Vehicles. *IEEE Trans. Dependable Secur. Comput.*
- Matsumura, R., Sugawara, T., and Sakiyama, K. (2018). A secure LiDAR with AES-based side-channel fingerprinting. In *Proc. - 2018 6th Int. Symp. Comput. Netw. Work. CANDARW 2018*, pages 479–482.
- Neal, A. (2018). LiDAR vs. RADAR. <https://www.fierceelectronics.com/components/lidar-vs-radar>, Accessed 2023-01-19.
- Petit, F. (2022). What are the 5 levels of autonomous driving? - Blickfeld. <https://www.blickfeld.com/blog/levels-of-autonomous-driving/>, Accessed 2023-01-14.
- Petit, J., Stottelaar, B., Feiri, M., and Kargl, F. (2015). Remote Attacks on Automated Vehicles Sensors: Experiments on Camera and LiDAR. *Blackhat.com*, pages 1–13.
- Qayyum, A., Usama, M., Qadir, J., and Al-Fuqaha, A. (2020). Securing Connected Autonomous Vehicles: Challenges Posed by Adversarial Machine Learning and the Way Forward. *IEEE Commun. Surv. Tutorials*, 22(2):998–1026.
- Roriz, R., Cabral, J., and Gomes, T. (2022). Automotive LiDAR Technology: A Survey. *IEEE Trans. Intell. Transp. Syst.*, 23(7):6282–6297.
- Shin, H., Kim, D., Kwon, Y., and Kim, Y. (2017). Illusion and dazzle: Adversarial optical channel exploits against lidars for automotive applications. In *Lect. Notes Comput. Sci. (including Subser. Lect. Notes Artif. Intell. Lect. Notes Bioinformatics)*, volume 10529 LNCS, pages 445–467.
- Tencent Keen Security Lab (2019). Experimental Security Research of Tesla Autopilot. page 38. https://keenlab.tencent.com/en/whitepapers/Experimental_Security_Research_of_Tesla_Autopilot.pdf, Accessed 2023-01-13.
- Xu, D., Anguelov, D., and Jain, A. (2018). PointFusion: Deep Sensor Fusion for 3D Bounding Box Estimation. In *Proc. IEEE Comput. Soc. Conf. Comput. Vis. Pattern Recognit.*, pages 244–253.
- Zhou, H., Li, W., Kong, Z., Guo, J., Zhang, Y., Yu, B., Zhang, L., and Liu, C. (2020). Deepbillboard: Systematic physical-world testing of autonomous driving systems. In *Proc. - Int. Conf. Softw. Eng.*, pages 347–358.