INNOVIZ™
TECHNOLOGIES

DESIGNING A LEVEL 3 LIDAR FOR HIGHWAY DRIVING

AN INNOVIZ WHITE PAPER **By Omer David Keilaf**

TABLE OF CONTENTS

Preface

Readers are invited to contact Innoviz at WhitePaper@Innoviz-tech.com to comment on, correct, engage in a discussion, and add ideas about this publication. Using the content of this white paper is subject to the restrictions listed in the agreement accepted by the reader in order to download this document from the Innoviz web site.

Executive Summary

A LiDAR must be robustly designed to ensure safety and system performance at all speeds. This is particularly critical for highway driving. In this white paper we discuss the importance of proper LiDAR design for safe driving of a Level 3 autonomous vehicle at highway speed. We analyze the requirements for such parameters as horizontal and vertical Field of View, horizontal and vertical resolution, detection range, frame rate, and reflectivity response.

One of the greatest challenges for an autonomous vehicle is to detect a small, extremely low reflectivity object, such as a tire, at long distance while driving at highway speeds. Many trade-offs are required in order to ensure safe braking and avoid a crash. This safety event use case illustrates the requirement for long-range detection and the highest possible resolution and frame rate to ensure detection of objects and obstacles at high speeds.

In addition to these parameters, the autonomous driving system must consider the influence of varied weather, lighting, and road conditions, including slopes and curvature. It must also consider the braking time required by driving at highway speed. All these factors impose increased demands on the autonomy stack to meet key performance indicators under any driving condition.

The mounting position of the LiDAR in the vehicle has a significant impact on both the LiDAR design and the vehicle design because a Level 3 LiDAR should be unobtrusive. The design of the car can easily determine the design of the LiDAR. Engineers must accommodate themselves to the engineering demands dictated by the car design.

Having taken these requirements into consideration and based on our experience with InnovizOne, Innoviz has developed InnovizTwo, a next-generation, high-performance LiDAR. In this white paper we introduce InnovizTwo, which is designed to ensure safe and comfortable driving in all situations, especially at highway speeds.

Introduction

Every year over 1.3 million people across the globe die in traffic accidents[1] and up to 50 million people suffer from road-related injuries [2]. Human error accounts for 80-90% of the probable cause of all accidents and 94% of all fatal ones [3]. Since transportation is part of our lives, best engineering practices should be applied to minimize the potential for human error.

Autonomous driving is one of the most challenging tasks in our time. In the quest to design the 'perfect' autonomous car, thousands of engineers in different car companies are engaged in describing different driving scenarios and use cases in which a car must perform safe autonomous driving.

A pareto analysis of car accidents shows that most fatalities on the road are caused on highways and straight urban roads, primarily due to human distraction [4]. Although highway driving simplifies driving decision software, it poses more requirements on the LiDAR sensor, which makes LiDAR selection more difficult for carmakers. Highway driving occurs at higher speeds, requiring longer braking times and imposing increased demands on the autonomy stack. Increased detection range at the sensor level affords the software stack longer response time.

Defining LiDAR Requirements

When designing an autonomous driving feature, many parameters must be considered, including the velocity at which the car is driving automatically, vehicle location, weather conditions and more. These parameters are used to calculate the abstract requirements for the LiDAR, camera, and radar.

The perception parameters (such as determining the distance to detect or classify a person, the allowed latency, false positive rate or misdetections) are split among the various sensors in order to compensate for a different sensor's limitations. The key performance indicators (KPIs) provided by the LiDAR, camera, and radar complement each other and provide redundancy, since the three KPIs will be compared by the domain controller used for sensor fusion.

The LiDAR's perception outputs must meet all the KPIs defined by the car manufacturer for making driving decisions. Meeting those KPIs within the limitations of the available processing power and allowed latency will determine the design of the LiDAR and the perception software that converts the LiDAR's raw point cloud data into the processed data used for sensor fusion.

Car speed
Driving location
Weather
Ambient light

Driving Decision Requirements

Attributes:
Classification, detection, speed, shape, etc.

KPIs:
Accuracy, distance, location, latency, false positive rate, minimal object size

LiDAR/Camera/ Radar Perception Requirements

Attributes:
Range, resolution, field of view, frame rate, reflectivity

KPIs:
Min/max range, position, and accuracy, confidence level, false positive rate, etc.

LiDAR Sensor Requirements

Figure 1: Partial requirements flow

1   World Health Organization (WHO). Global Status Report on Road Safety 2018. December 2018
2   ibid
3   National Highway Traffic Safety Administration (NHTSA). 2016 Fatal Motor Vehicle Crashes. October 2018
4   National Highway Traffic Safety Administration (NHTSA). 2019 Fatal Analysis Reporting System (FARS). March 2021

LiDAR design also is greatly influenced by the constraints dictated by the platform design.

Figure 2: Platform design considerations

SAE Autonomy Levels

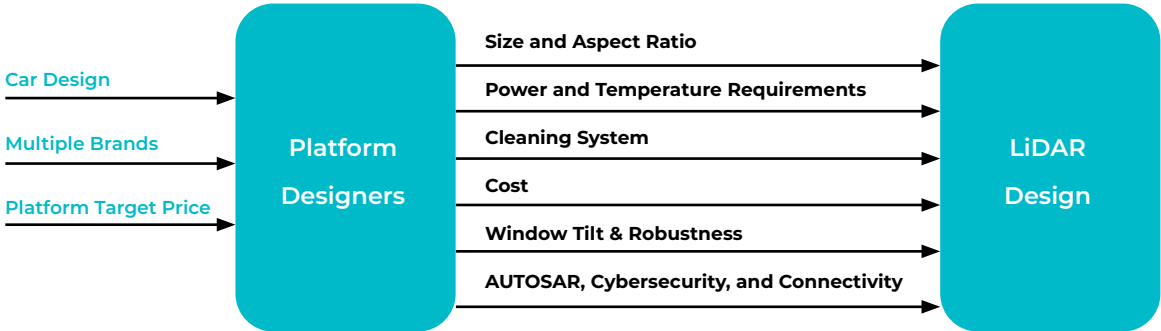L3 autonomous driving, also known as "Highly Automated Driving", is described by the SAE J3016 standard as a vehicle that can perform self-driving within limited conditions decided by the automotive Original Equipment Manufacturer (OEM) without the need for human supervision[5]. The 'driver' is required to respond to a system (vehicle) request for him to regain control within a sufficient period of time.

The requirement for total redundancy in L3 systems is a quantum leap from today's L2+ systems (an unofficial hybrid level between Level 2 and Level 3), which still require a human to supervise the system and not remove his eyes from the road or his hands from the wheel. While L3 driving is likely to occur primarily on a highway, human attention is required to provide the supervision necessary for L2+ driving when exiting the highway and on local roads.

	SAE Level 0	SAE Level 1	SAE Level 2	SAE Level 3	SAE Level 4	SAE Level 5
What does the human in the driver's seat have to do?	You are driving whenever these driver support features are engaged - even if your feet are off the pedals and you are not steering			You **are not** driving when these automated driving features are engaged - even if you are seated in "the drivers seat"		
	You must constantly supervise these support features: you must steer, brake or accelerate as needed to maintain safety			When the feature requests you must drive	These automated driving features will not require you to take over driving	
	These are driver support features			**These are automated driving features**		
What do these features do?	These features are limited to providing warnings and momentary assistance	These features provide steering OR brake/ acceleration support to the driver	These features provide steering AND brake/ acceleration support to the driver	These features can drive the vehicle under limited conditions and will not operate unless all required conditions are met		These features can drive the vehicle under all conditions
Example Features	• automatic emergency braking • blind spot warning • lane departure warning	• lane centering OR • adaptive cruise control	• lane centering AND • adaptive cruise control at the same time	traffic jam chauffeur	• local driverless taxi pedals/ • steering wheel may or may not be installed	same as level 4, but feature can drive everywhere in all conditions

Figure 3: SAE levels of driving automation
Source: *SAE International*

To reach L3 or a higher level of automation, engineers need to prove that the automated driving system meets the auto industry ISO 26262 standard's Functional Safety level ASIL D, which is the highest level of automotive hazard. This level requires analyzing different driving use cases and proving that in any possible event the car will ensure the safety of people within or outside the car, even in the event of a single point of failure or degradation in one of the system's components. The analysis is done through the entire platform for various driving scenes, temperatures, weather conditions, and other internal or external causes.

---

5   Society of Automotive Engineers (SAE). SAE J3016™ Recommended Practice: Taxonomy and Definitions for Terms Related to Driving Automation Systems for On-Road Motor Vehicles. May 2021

Balancing Safety and Comfort

While safety is critical, if the system does not behave in a way that is "comfortable" to the driver, he will lose confidence in the automation system and will not use it. By saying "comfortable", we mean that the car does not delay in decision-making and therefore does not make "last minute" decisions, such as hard braking and sudden maneuvers. The system should ensure smooth driving, which derives from perfect understanding of the scene and fast, correct decision-making about objects that are ahead of the ego vehicle.

The Risk and Potential of L2+

Many drivers confuse autonomous driving with features such as Autopilot. Drivers think that they can keep their eyes off the road and hands off the wheel with Autopilot. The following list presents a snapshot of the current state of Level 2+ vehicles and what LiDARs can offer to L2+ vehicles.

1. L2+ platforms are becoming more and more popular.

2. This unofficial hybrid SAE level (also known as "Autopilot") still requires human supervision.

3. Most accidents happen on straight roads due to human over-confidence regarding road conditions.

4. L2+ without a LiDAR provides a false sense of confidence, resulting in fatal accidents.

5. Research shows drivers become inattentive when autopilot in engaged [6].

6. L2+ without LiDAR gives a bad name to autonomous driving.

7. Accidents caused by L2+ vehicles without LiDARs delay the adoption of AVs and safer mobility.

8. Only a LiDAR that does not incorporate a camera for detection provides redundancy according to ISO 26262.

9. L2+ with LiDARs provides a safer platform for customers.

10. Car makers can leverage on L2+ with LiDARs to crowdsource data in real-life conditions and validate the platform safety.

11. Use of LiDARs on L2+ will save customer lives, will accelerate the development process, will save data collection cost, and will reduce the risk to automakers by launching L3 services at higher confidence.

12. L2+ using a LiDAR and over-the-air (OTA) software upgrades can incrementally improve vehicle safety and reach L3 without changing the LiDAR hardware.

STEP 1

Safe and affordable L2+

STEP 2

Collect data and validate

STEP 3

Software upgrade to L3

Figure 4: Software updates transform L2+ to L3

6   https://techcrunch.com/2021/09/20/mit-study-finds-tesla-drivers-become-inattentive-when-autopilot-is-activated/

Level 3 is the Autonomous Driving MVP

As in any new market, a clear definition of a Minimum Viable Product (MVP) is needed. The definition of such an MVP can be translated to the reduction of significant complexity related to a variety of urban scenarios, intersections, and the high likelihood of Vulnerable Road Users (VRUs), such as pedestrians and bicyclists, present in the middle of the road.

A pareto analysis of car accidents shows that most fatalities on the road occur on highways and straight urban roads, primarily due to human distraction [7]. Restricting L3 vehicles mostly to highways would simplify overall autonomous system requirements and costs, eliminate a huge daily risk from our lives, and reduce the need for human control of a car for many hours every day.

Although highway driving simplifies driving decision software, it poses more requirements on the LiDAR sensor, which makes LiDAR selection more difficult for carmakers. Highway driving occurs at higher speeds, requiring longer braking times and imposing increased demands on the autonomy stack. Increased detection range at the sensor level affords the software stack longer response time.
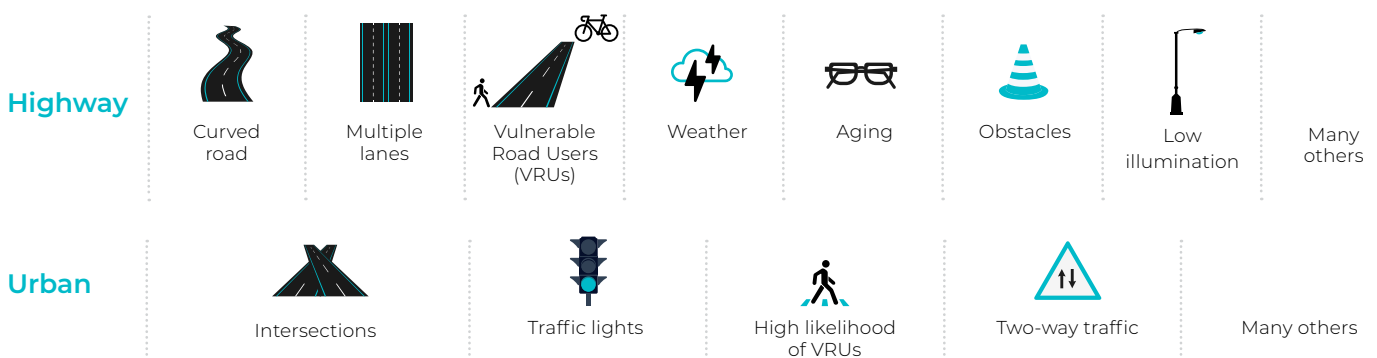
Highway

| Curved road | Multiple lanes | Vulnerable Road Users (VRUs) | Weather | Aging | Obstacles | Low illumination | Many others |

Urban

| Intersections | Traffic lights | High likelihood of VRUs | Two-way traffic | Many others |

Figure 5: Driving conditions and considerations of highway and urban environments

In section 3, we discuss sensor fusion and compare cameras, radars, and LiDARs. In section 4, we discuss LiDAR specifications and posit that LiDAR manufacturers' claims can be misleading and must be carefully evaluated. In section 5, we present a use case of a highway safety event and the required performance in order to ensure safe driving.
In section 6, we present the InnovizTwo LiDAR, which is based on the requirements revealed in thousands of hours of face-to-face discussions with the automotive industry's technology leaders.

---

7   National Highway Traffic Safety Administration (NHTSA). 2019 Fatal Analysis Reporting System (FARS). March 2021

Sensor Fusion

The utilization of various types of sensors in the vehicle's driving decision-making is defined as "sensor fusion". Most L2 autonomous vehicles use cameras and radars as the sensors for making driving decisions (changes in car speed and trajectory). Camera and radars are not able to provide redundancy to each other and fulfill functionality safety requirements. In order to meet ISO 26262 ASIL D, the driver is defined as the redundant sensor by observing the road and providing backup to any wrong detection by the camera or radar. As such, the driver is expected to keep his eyes on the road and hands on the steering wheel so that he can take over in case the car's driving decision-making system makes an error. In L2 and L2+ systems, the driver is liable in case of an accident.

In L3-L5 autonomous vehicles, SAE J3016 defines the automotive OEM as the legal driver within a designated Operational Design Domain (ODD). Therefore, the OEM is fully responsible for everything that happens when the autonomous driving system completely takes over the full driving task either in defined-use cases (L3-L4) or all use cases (L5), in compliance with ISO 26262 ASIL D. This is intended to guarantee a failure probability so low that automakers, consumers and regulators agree that the vehicle can handle driving without any supervision by a human driver in given areas and at given times.

The sensor suite is defined as a critical part of the autonomous driving stack. Correct driving decisions cannot be made without an accurate understanding of the vehicle's surroundings. A system which complies with ISO 26262 ASIL D requires the "smart" redundancy of critical elements. "Smart" redundancy means using a mix of sensor types, and not simply using several of the same sensor type (i.e., camera, radar, or LiDAR) for backup in case of sensor malfunction.


Figure 6: LiDAR, camera, and radar performance across different dimensions
Source: *Innoviz Technologies Ltd.*

Because all sensor types have limitations, sensor mix improves the probability that there will always be a sensor that performs well in any given situation. Each sensor must work independently without any interaction between them. Thus, correct driving decisions always will be made, regardless of driving conditions that are challenging for a certain sensor type. Each sensor type also must have its own perception software to ensure software redundancy as well as hardware redundancy.

Since each sensor has its advantages and disadvantages, it is necessary to use cameras, radars, and LiDAR, which complement one another and provide the redundancy required by ISO 26262.

Cameras

Cameras are a good source of sensing and are the fundamental sensors used today to achieve scene understanding. Cameras are passive sensors which rely on an external light source (sun, car's headlights, street lighting, etc.). However, they suffer from strong degradation during low-light conditions, strong light saturation, and weather conditions such as water on the lens that cause image blurring. ISO 26262 requires redundancy by a different sensor type. A camera cannot be made redundant by using another camera sensor, since it also will have a high likelihood of a similar failure.

Beside environmental effects, the 2D nature of the camera sensor limits the true positive rate (TPR) of object detection in several scenarios in which the object is obscured or only partially captured. Object detection relies on machine vision which requires prior training from sufficient data. Edge cases of different types of possible objects require real-time detection of objects that were not available in the training set. Such a requirement is difficult to achieve by cameras under the different degradation modes.

Figure 7: Camera vs. LiDAR in low-light conditions

Figure 8: Camera vs. LiDAR in mist/rain

Radars

Radar is excellent for determining the location and velocity of objects moving along the radial direction. Radar sensors can be classified per their operating distance ranges: Short-Range Radar (SRR) 0.2 to 30m range, Medium-Range Radar (MRR) in the 30-80m range and Long-Range Radar (LRR) 80m to over 200m.

Radars are a good type of sensor to reduce sensitivity to environmental events; however, even the most advanced "next-generation' radars under development fail to provide enough scene understanding due to their low resolution. In addition, they have difficulty detecting stationary objects or objects moving perpendicularly across the radial direction. They also cannot accurately determine the shape of objects or the size of small objects. As such, today's radars cannot provide total redundancy to a camera for accurate understanding of the scene.

LiDARs

L2/L2+ platforms that only use cameras and radars require human supervision. Therefore, they do not meet the L3 requirement to provide complete redundancy analysis according to ISO 26262. Currently only LiDAR can guarantee redundancy to cameras and radars.

A LiDAR is an essential component in an autonomous vehicle (AV), enabling sensing that neither radar nor camera can match. LiDARs excel at nighttime operation because they do not rely on ambient light as cameras do. They also offer better object detection than cameras in difficult weather conditions. LiDARs are less susceptible than cameras to blinding by direct sunlight or oncoming headlights. They also provide much higher resolution than radars and enable better object identification.

Figure 9: Camera vs. LiDAR in direct sunlight

LiDAR, which provides a high-resolution, accurate 3D representation of the world, can most easily be used by the perception software layer to create a useful representation of the physical world. LiDAR systems use pulsed laser beams to emit light that reflects (bounces) off objects in the environment and is detected by the LiDAR's sensor when it returns. The differences in return times are used to calculate their exact distance from the LiDAR. This contrasts with cameras, which estimate distance from the sensor.

The points of light that return from the LiDAR scans are collected and processed by the LiDAR's processing unit to create a point cloud, a three-dimensional (3D) representation of the environment, even in low-light conditions. The point cloud displays the environment from multiple points of view (i.e., overhead, driver view, third person, etc.). Each LiDAR point reflects full X, Y, and Z axes reflectivity data over time.

When viewed from multiple points of view, the data points are recognized and classified as objects. Additional information about the object – such as its position, size, velocity, direction, and reflectivity – can be determined from the signal.

While there are many types of LiDARs, most are not able to meet the different requirements for L3 highway driving due to the high bar necessitated by performance, cost, reliability, and size. The following analysis suggests the different attributes an L3 LiDAR must provide based on driving scenarios and events in which cameras and radars are lacking.

Understanding L3 LiDAR Specifications and Performance

Because driving is a dynamic experience, LiDARs must be designed to enable optimal performance under a wide range of conditions and situations. Because performance parameters are interrelated, many trade-offs are required in order to accommodate standard driving conditions and edge cases.

Figure 10: Comfort vs. safety driving requirements
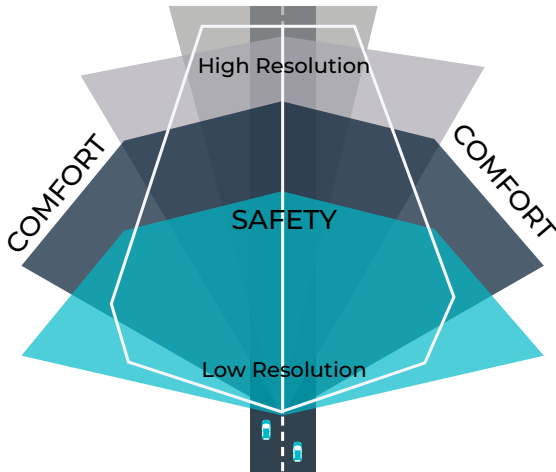
Optimal performance should be considered in the context of safety and comfort requirements. Safety (or emergency) requirements are driven by the need to detect collision-relevant events of moving objects, objects that cannot be driven over, and objects that cannot be driven under. Comfort also factors in for the user experience not to include abrupt stops, turns or any other unexpected action. The combination of comfort and safety defines a nuanced set of requirements when factoring in performance. The faster the autonomous vehicle is designed to drive, the higher the commensurate resolution, range and frame rate that are required for highway driving. However, it should be noted that this higher performance is required only in a small region of interest (ROI), which is the center of the Field of View (FOV) for long-range detection.

The following table presents the parameters that are most important for understanding L3 LiDAR specifications and performance.

SPECIFICATION	ATTRIBUTE
Within ROI (Center)	Field of view
	Minimum range*
	Resolution
	Maximum range*
Outside ROI (Sides)	Field of view
	Minimum range*
	Resolution
	Maximum range*
General	Size
	Weight
	Frame rate
	Power consumption
	Interfaces

All ranges should be indicated at 95% True Positive Rate, 1% False Alarm Rate, 10% reflectivity, 25°C, 20Hz, 20Klux and 100Klux

In this section we present the requirements for highway driving, which is the primary application for Level 3 LiDARs. We explore the requirements and considerations that determine LiDAR design and performance while also weighing comfort and safety. At the end of this section, we present a challenging use-case that is not uncommon, and which has been responsible for several high-visibility accidents involving L2+ vehicles operating under Autopilot.

LiDAR Mounting Position

It might be assumed that Level 3 platforms might introduce some relaxation on the requirements compared to Level 4 and urban scenarios. But, since Level 3 are targeted for privately owned vehicles, the design of the car and the look of the car in many cases supercedes all other requirements.

Looks sell.

The mounting position of the LiDAR in the vehicle has a significant impact on both the LiDAR design and the vehicle design. The design of the car can easily determine the design of the LiDAR.

Mounting positions such as grille mounting, rooftop mounting, behind the windshield, or colocation with the headlight significantly affect the:

- Field of View
- Size and aspect ratio
- Environmental conditions and air flow
- Window tilt, color and resiliency to stone chipping and different kinds of road dirt
- Cleaning system design

Level 3 cars demands due to the above are more demanding on LiDAR designers than today's Level 4 shuttles or robotaxis because a Level 3 LiDAR should be unobtrusive.

Unfortunately, car designers don't always agree that the LiDAR is the most valuable design element. In some cases the design of the car is made before the technical requirements are all set and they are not considered. Since the design of the car is high up in the car manufacturer's hierarchy, the engineers must accommodate themselves to the engineering demands dictated by the design which in many cases enforces strong design changes on the LiDAR.

LiDAR Configuration for Highway Driving

Forward-Looking LiDAR

L3 highway driving requirements are simplified by the absence of intersections, traffic lights, and complex urban scenarios which include pedestrians and bicyclists. This means that redundancy may be required only for the front view of the traveling vehicle. Hence, only one LiDAR might be needed in the car, assuming it can support the entire design requirements. Multiple LiDARs would be helpful if additional redundancy is required.

This region of interest defines the FOV, resolution, frame rate, and range that will achieve passenger safety, even at the expense of passenger comfort. We refer to this as the "safety FOV". It's important to note that the safety FOV requirement is higher at short-range than at long-range due to the shorter time in which the car can react. Harsh false alarm requirements are defined for the short-range in order to avoid unnecessary braking of the car, which would damage the user experience.

Field of View and Regions of Interest

As discussed before, the requirements of resolution and range are different for the different areas of the FOV. We divide the requirements between two distinct areas that together provide the entire set of requirements. As expected, the requirements for the center of the FOV are higher, due to the assumption that the LiDAR is mounted in the center of the car and faces forward in the direction in which the car is driving.

It is possible to increase the range and resolution of a selected area of the FOV, known as a region of interest. The ROI is the area in which the car should detect hazards at long distances and avoid a safety event in sufficient time.

Figure 11: Vertical and Horizontal Field of View and the ROI subset

Range is increased by dynamically focusing laser energy within the ROI, while reducing energy provided to regions outside the ROI. Regions of interest can be configured independently and can be located anywhere in the FOV.

Figure 12: Horizontal ROI positioning considerations

Key LiDAR Performance Design Considerations

To understand driving requirements, we must first define the specifications under consideration. A LiDAR could be described in a very similar way to a camera, along with some other considerations not relevant to cameras.

1. Horizontal Field of View (HFOV) – The horizontal angle the LiDAR can capture in each frame.
2. Vertical Field of View (VFOV) – The vertical angle the LiDAR can capture in each frame.
3. Horizontal Resolution (HRes) – Angular separation between pixels within same row.
4. Vertical Resolution (VRes) – Angular separation between pixels within same column.
5. Range – The expected maximum range measurement for certain reflectivity under certain light conditions, provided per pixel.
6. Frame rate – How many times a second all pixels are refreshed.

Horizontal Field of View

Horizontal Field of View Considerations

Wide Horizontal FOV (HFOV) is required for safe and comfortable driving. The wider the HFOV is, the better the car can predict collision-relevant objects and track them. The HFOV should enable detection of a cut-in scenario, in which an overtaking vehicle suddenly enters the ego vehicle's lane. The minimum FOV required to detect a cut-in scenario is 120°.


120°

Horizontal FOV

| Road curvature | Traffic awareness | Lane change/keeping | Cut-in scenario |

Figure 13: HFOV considerations

The assumptions of the analysis are:

1. LiDAR is mounted in the grille of the car. In roof-mounting, lower FOV is required since the LiDAR is placed in a location that can detect a cut-in scenario earlier.

2. Ego car width is about 1.9m. Wider vehicles might require a wider FOV.

3. For a highway scenario use case, the overtaking vehicle is 1.6m or longer.

4. The distance between lane markings is 3.7m. The distance between the center of the car and the inside of the lane marking is 1.875m. There is a 90cm safety margin between the side of the ego vehicle and the lane marking.

5. When driving at a very low speed or in high traffic congestion, the car's ultrasonic sensors are capable of providing detection of up to 1m from the side of the adjacent vehicle.

6. The overtaking vehicle drives at a speed similar to or higher than the ego vehicle.

7. At least 40cm length of the front of a vehicle in an adjacent lane must be detected to achieve high confidence of object detection and classification.

8. A fast-driving vehicle can maneuver laterally at about 5°. At 130km/h, the lateral speed is 3.14m per second. Assuming detection in at least three frames at 20FPS, the overtaking vehicle can travel 47cm before the ego vehicle will react.

$$36(\frac{m}{s}) \cdot \sin 5° \cdot 3 \cdot \left(\frac{1}{20}[sec]\right) = 0.47m$$

In order to detect the front of the overtaking vehicle, about 120° HFOV is required at 20FPS. If the frame rate is lower, a bigger safety margin is required and the HFOV should be increased.
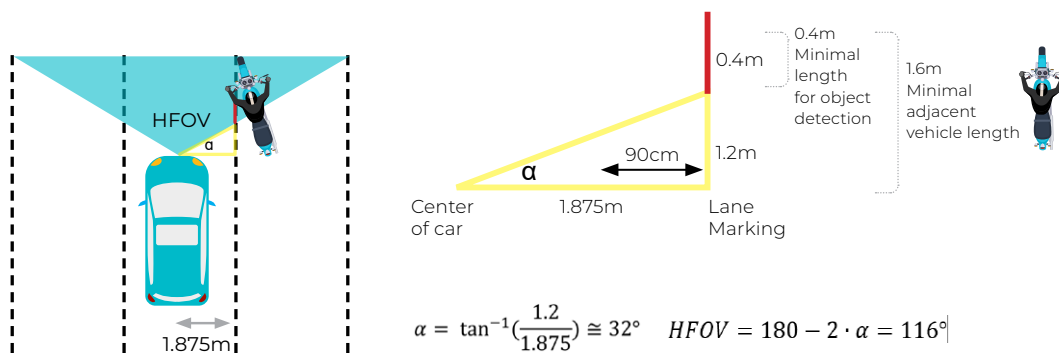

$$\alpha = \tan^{-1}(\frac{1.2}{1.875}) \cong 32° \quad HFOV = 180 - 2 \cdot \alpha = 116°$$

Figure 14: Cut-in scenario with grille mounting

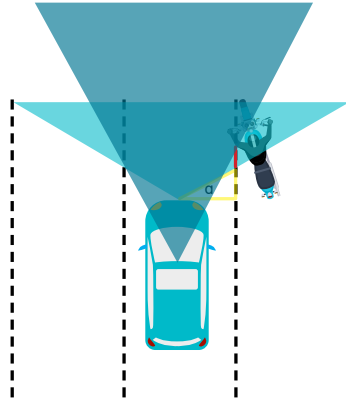A LiDAR mounted in the roof requires a lower HFOV to detect a cut-in, as seen below.

**Figure 15**: Cut-in scenario with rooftop mounting

ROI Horizontal Field of View

The minimal ROI FOV must support rapid changes caused by different road slopes, acceleration and deceleration of the car, wind, and errors in on-line calibration of the sensor direction. All these errors must be covered by ample fixed FOV within the ROI and cannot be compensated by steering the LiDAR scan. Minimal scene understanding also requires the ability to scan multiple lanes and different types of road curvature. Additional FOV is required as a buffer to allow assembly tolerances in the factory.

30°
ROI Horizontal FOV

Road curvature

Multiple lanes on each side

Assembly error

Figure 16: Horizontal FOV ROI considerations

The following analysis is done by the edge case scenario of a small object which must be detected at a minimum of 100m. We compute the ROI to be ~30° (with tolerances) as follows:
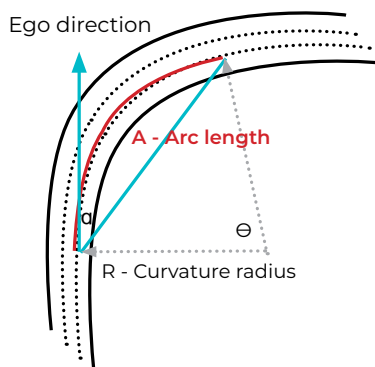
$$ROI_{HFOV[\alpha]} \sim \frac{1}{2} * \left(\frac{A}{R} * \frac{360}{2*pi}\right)°$$
α angle (half HROI) which is extracted from ϴ angle, which is calculated by the arc equation

Road curvature minimum radius (R) = 280m

Maximum safety detection range (A) = 100m

Yaw tolerance 1.5° (mounting tolerance) + 0.1° (yaw calibration error)

Ego lane margins = 1.5 x lanes + 1m = ~6.625m

$$ROI_{HFOV[\alpha]} = \frac{1}{2} \cdot \left(\frac{100}{280}\right) \cdot \left(\frac{360}{2 \cdot pi}\right) + 1.6°yaw + 3.5°(lane\ margins) = ~15°$$
The Horizontal FOV of ROI is 2α (from both sides of ego lane) = ~30°.

Figure 17: HFOV ROI calculations

Now that we have determined the HFOV and knowing the minimum detection distance of 100m, we can solve for the speed at which the vehicle is traveling. We determine the speed to be 130km/h (80mph).

$$Speed(ft)[mph] = 0.5 * \left(-0.03R + \sqrt{(0.03R)^2 + 4R\left(\left(\frac{15E}{100}\right) + 3.6\right)}\right)$$

Where R = Curvature radius; E = Elevation

Note that road curvature usually limits the car's maximum speed. AV architects can adjust the ROI HFOV by taking into account the maximum detection range which is derived from the maximum speed.

Vertical Field of View

Vertical Field of View Considerations

The Vertical FOV (VFOV) is defined by its downward and upward requirements. The required surface detection is defined by minimal detection of ground (small objects) around 3m away.

25-40°

Vertical FOV | Mounting height | Minimum road recognition | Under-drivable | Unsafe driving speed | Dangerous load

Figure 18: VFOV considerations

Figure 19: LiDAR mounting options

The formula for determining the α and β values for the grille mounting vertical FOV is:

$$\alpha = \tan^{-1}\left(\frac{LiDAR\ Grille\ Height}{Ground\ Detection\ from\ LiDAR}\right) = \tan^{-1}\left(\frac{0.5}{3}\right) \cong 9.5°$$

$$\beta = \tan^{-1}\left(\frac{Under-Drivable\ Height\ from\ LiDAR\ Mounting}{Distance\ from\ Cargo}\right)$$
$$= \tan^{-1}\left(\frac{1.6-0.5}{3}\right) \cong 21.8°$$

$$VFOV\,grille = \alpha + \beta \cong 29.5$$

$$\alpha = \tan^{-1}\left(\frac{LiDAR\ Grille\ Height}{Ground\ Detection\ from\ LiDAR}\right) = \tan^{-1}\left(\frac{1.5}{5}\right) \cong 16.7°$$

$$\beta = \tan^{-1}\left(\frac{Under-Drivable\,Height\ from\ LiDAR\ Mounting}{Distance\ from\ Cargo}\right)$$
$$= \tan^{-1}\left(\frac{1.6-1.5}{5}\right) \cong 1.14$$

$$VFOV\,Roof = \alpha + \beta \cong 17.8$$

Adjustable VFOV is necessary in order to shift the LiDAR's scanning pattern to the correct position if a shift from the ideal scanning angle is identified. The shift from the correct position may occur for several reasons:

- Imprecise tolerances in the installation process of the LiDAR in the car body
- Change to the car pose due to unbalanced load or aging
- Change in the car suspension at high speeds

The equations above are provided as an example. Car makers need to consider the mounting position of the LiDAR in all the models in which they plan to install it, since mounting location determines LiDAR design and performance.

For a grille mounting height of 50cm, the total VFOV should be 25-40°. The vertical requirement is defined by the need to detect mounted cargo protruding from a vehicle (~1.5 height) and detection of under-drivable objects, such as low bridges. VFOV requirements also should be calculated based on mounting height (grille, roof, windshield, or headlamps). Roof-mounting should require lower VFOV, since detection of protruding objects is easier and the hood blocks the downward view.

ROI Vertical Field of View

The VFOV of the ROI is determined by a viewing angle which can support different slopes of the road to ensure that any small object on the road can still be detected at long range. The overall VFOV is larger than the ROI, but the range and resolution requirements for the rest of the VFOV are lower. The minimal ROI FOV must support rapid changes in driving conditions – such as different road slopes, acceleration and deceleration of the car, and wind – and on-line errors in sensor calibration. All these errors must be taken into consideration by adequate fixed FOV of the ROI. They cannot be compensated by steering the LiDAR scan because they tend to change rapidly and cannot be predicted; therefore they can only be compensated by a greater VFOV.

A VFOV of 10° is adequate for a scanning LiDAR to eliminate factors that might cause changes in effective VFOV, e.g., assembly tolerances, aging of the chassis, weighted carload, etc. These are elements for which a scanning LiDAR with enough dynamic range could compensate.


10°
ROI Vertical FOV

Slope   Acceleration/Deceleration   Wind   Online calibration error

Figure 20: Vertical FOV ROI

Sufficient VFOV is necessary to eliminate gaps between lines. A dense ROI is needed to allow the car to detect objects that are not over-drivable, since they might damage the car and possibly injure passengers inside the car. Any object above 14cm (which is the standard passenger car suspension length or chassis height) could damage the car. If there are gaps between the lines, it will be impossible for the car to understand that an object exceeds the height that might damage the car if traveling above a certain velocity. Therefore, it is necessary to stop the car at a safe distance if such an object is identified.

Flash LiDAR (which cannot steer its beam) will require an additional 5° ROI (total of 15°).

5°
Steerable Vertical Field of View - every few km

Assembly tolerances   Chassis aging   Car load

Figure 21: Steerable VFOV

Resolution

Resolution Outside the ROI

The capability of perception software to detect a target is detemined by the minimum number of pixels that are required to be detected on the target in order to provide high confidence of identification and classification. The shorter the range is, the lower the required resolution. With resolution of 0.05° at 150m, each pixel is 13cm high. At 50m, resolution of 0.15° is required to achieve the same pixel size.


0.1° - 0.25°

Horizontal or Vertical resolution (Out of ROI)

Traffic awareness and late cut-in scenarios

Movement detection

>100m object detection

0.05°

Figure 22: Resolution outside ROI

However, the number of pixels required for detection of people could be the same, regardless of their distance as long as the resolution is high enough. This assumes that the perception layer's ability to classify and identify an object should be at the same level of confidence within the safety zone. At least 2x4 (HxV) pixels are required to detect and classify a person.


Perception level limitation

7 vertical pixels

5 vertical pixels

4 vertical pixels

Figure 23: Perception layer recognition of people at different distances

Without a sufficient number of pixels, a pedestrian would be recognized as a moving object, but not as a pedestrian.


Figure 24: Pedestrian classification requirement

Classification is important in order to have good trajectory prediction.

Horizontal Resolution Inside the ROI

For higher probability of correct identification and classification of VRUs at long range, higher horizontal resolution is needed. Assuming a pedestrian at 150m with width of about 40cm, using 0.07° resolution could enable detection of 2-3 pixels horizontally.

0.07°

Horizontal Resolution

Vulnerable Road Users 150m classification	People classification minimum 2x4 pixels	Reflectivity aids classification	Movement detection

Figure 25: Horizontal resolution considerations

Vertical Resolution Inside the ROI

The resolution that is needed within the ROI is determined by the height of an object that could damage the car or cause it to overturn above a certain velocity. Because the standard passenger car suspension length or chassis height is 14cm, any object above 14cm could damage the car if driven over.

0.05°

Vertical Resolution

Suspension length/ chassis height <14cm	Avoid objects that are not over-drivable	At least two vertical pixels on target	14cm tall object at 100m with 0.05° resolution

Figure 26: Vertical Resolution

A dense ROI is needed in order to allow the car to detect objects that cannot be driven over, since they might damage the car and possibly injure people inside the car. If there are gaps between the scanned lines, the car will not understand that an object which exceeds a certain height might damage the car. Therefore, it is necessary to stop the car at a safe distance in case such an object is identified.

In order to avoid false alarms within the vehicle's drivable area and to achieve sufficient confidence in detection, at least two vertical pixels must be detected across multiple frames. 0.05° resolution is required to guarantee detection of two consecutive pixels on a 14cm height target at around 100m. This limits the L3 speed to 130km/h on a dry road. See the explanation below.

The importance of resolution and range cannot be overstated. A single detected pixel could be an object whose size differs according to the detection range and the angular resolution of the LiDAR. For example, at 200m a pixel detected by a LiDAR with 0.05° resolution could be a very reflective 1cm high object which should be ignored by the car, or it could be a low reflectivity 17cm high object which cannot be ignored. Even if the LiDAR could detect one pixel at 500m, sufficient resolution is required to determine the car's reaction.

200m @ 0.05° Vertical Resolution

Over-drivable Bott's Dot

Black tire not over-drivable

Figure 27: Height ambiguity of an object within a pixel

The vertical size of the ROI is determined by a viewing angle which can support different slopes of the road to ensure that any small object on the road can be detected at long range. While the overall VFOV is larger than the ROI, the range and resolution requirements for the rest of the VFOV are lower.

Range

Range Inside the ROI

The range of the LiDAR is determined by several parameters that affect its capabilities, such as object reflectivity, LiDAR sensitivity, and its resilience to ambient noise and environmental conditions.

Assuming adequate LiDAR resolution, the car will react more smoothly as the detection distance increases. Comfortable driving at 80mph (130km/h) requires the car to identify VRUs at least 200m away. Because of possible LiDAR performance degradation over its lifetime or due to weather conditions, detection beyond 200m is advisable.

200m @ 10%

Range
(95% TPR, 1% FAR, 100Klux)

| Vulnerable Road Users 150m classification | 100 Klux sunlight | Weather | Blockage | Aging | Scratches |

Figure 28: Factors that impact range capability in ROI

Since detected objects consist of several pixels, an object should be detected with a true positive rate (TPR) of 95%-97%. With adequate resolution an object can be detected with 6 or more pixels, resulting in total TPR >99.9999%.

Range Outside the ROI

There are different considerations for long-range and short-range detection. In general, the higher the native resolution of the detector, the lower the amount of light that is collected by each pixel. This degrades the range capability. This tradeoff is mostly painful when it comes to very large, dark objects (such as black cars) that are more easily detected by a LiDAR with low native resolution.

Most LiDAR solutions require trading-off increased resolution for lower sensitivity of the LiDAR per pixel, which yields lower range. For example, black metallic cars can be detected with 10% reflectivity at 150m from the rear due to retro-reflectors (such as the license plate or plastic lenses over the rear lights). However, black metallic cars require 1% reflectivity at 50m when they are traveling in adjacent lanes.

50m @ 1%
150m @ 10%

Range
(95% TPR, 1% FAR, 100Klux)

| Black car | 100 Klux sunlight | Weather | Blockage | Aging | Scratches |

Figure 29: Range outside the ROI

Frame Rate

Frame rate is defined by the number of frames that all pixels are completely refreshed during one second. Twenty frames per second (FPS) means that every 1/20 second the LiDAR provides a new frame, where all the pixels are measured without any reliance on the prior frame. A higher frame rate allows for a greater reaction time and faster movement detection. Frames can be taken in succession to drive performance; a best practice of 4 out of 7 (4oo7) frames can lead to sufficient confidence and true positive rate in accurately detecting objects.

| 10-20Hz | | | | Example: 4oo7 to achieve TPR >99.997% |
| Frame Rate | Reaction time | Braking Distance | Cut-in scenario | 10FPS = 0.7sec., 20FPS = 0.35sec |
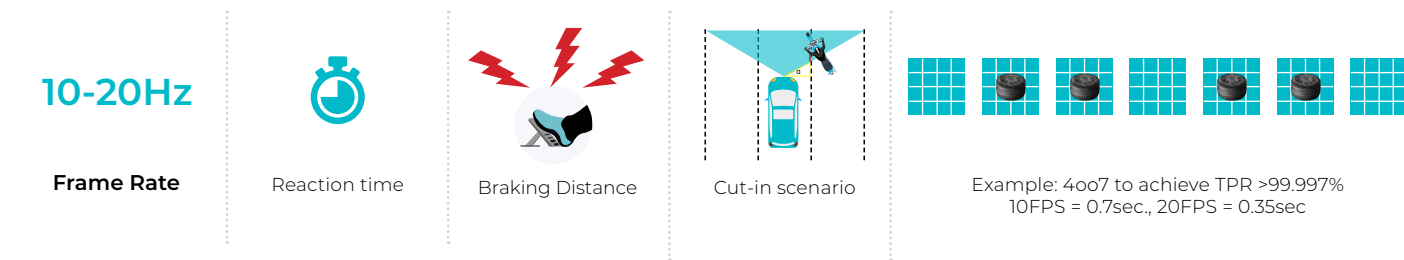
Figure 30: Frame rate considerations

Obtaining High Resolution, Frame Rate and VFOV Simultaneously

Until this point, we have discussed the LiDAR specifications in isolation. While LiDAR performance specifications are defined in datasheets, the upper limit of certain specifications sometimes is measured individually and not in relation to all the other parameters simultaneously. The LiDAR must be able to meet all individual parameters simultaneously. In evaluating a LiDAR, the specifications and parameters should be considered in tandem as there are tradeoffs in improving a single parameter. Since there are dependencies, it is essential not to degrade one by changing another.

The following example illustrates the interrelationship of frame rate, resolution, and Vertical FOV. A LiDAR operating at 10FPS that supports 640 lines per second can provide only 64 lines per frame. If 0.05° resolution is required at 10FPS, it can cover only 3.2° VFOV. Using a higher frame rate will reduce the VFOV even further.

	Other Suppliers			Innoviz
Lines per Second	640			8,000
Frames per Second	1	10	10	**20**
ROI Vertical Resolution	0.05°	0.4°	0.05°	**0.05°**
VFOV	30°	30°	3.2°	**30°**
Problem	Low Frame Rate	Low Resolution	Low VFOV	**No Problem!**

Figure 31: Configuration trade-offs

Reflectivity Response

Reflectivity data is generated by measuring the amount of light that is collected per-pixel. By normalizing the value by the measured distance, it is possible to extract an object's reflectivity. This information is helpful to identify lane markings and improve the classification of cars and pedestrians.

Reflectivity of an object is defined as the ratio between the amount of light that is collected versus the expected light scattered back from an ideal white Lambertian surface. Reflectivity could vary between 0% to hundreds of percent, where 100% reflectivity is defined by an ideal white object (which does not absorb any light) from which all light scatters back in a Lambertian manner. A Lambertian surface is defined by "an ideal matte" that, regardless of the observer's angle of view, receives the same amount of light.

	White Lambertian surface	Dark Lambertian surface	Mirror (specular reflection)	Retro-Reflector
Reflectivity	100%	10%	0%	500%
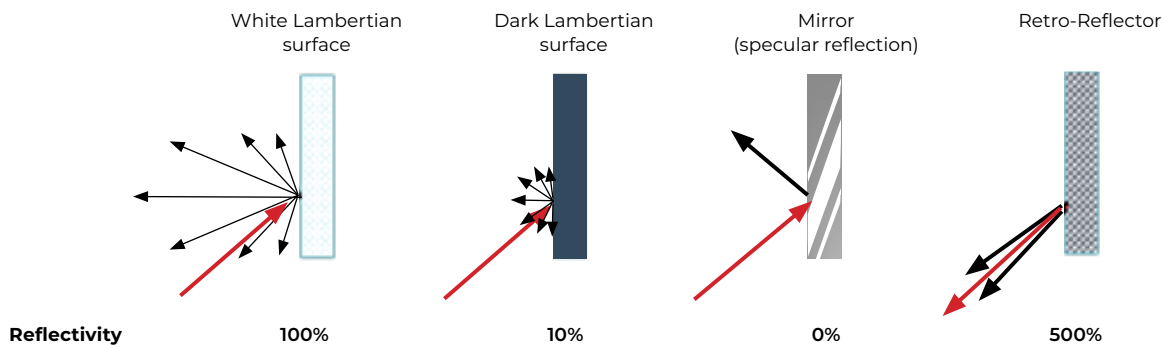
Figure 32: Reflectivity considerations

Whereas diffusely reflecting surfaces scatter light evenly, specular (mirror-like) objects reflect light in a given direction. From this phenomenon, active sensors require an incident angle near normal to detect reflections from specular surfaces. This tendency is leveraged in retro-reflectors: an assembly of specular surfaces are arranged such that light entering from a wide range of angles reflects multiple times before returning from the direction in which it entered. This property allows for reflections which could be many multiples of an ideal matte surface. Finally, highly transmissive surfaces (transparent materials, e.g., glass) reflect very little light. The light they do reflect is generally specular. This results in the phenomenon where, as the incidence angle changes from acute to perpendicular, the surface will suddenly "appear".

A black car, which is both dark and metallic (mirror-like), has very low reflectivity. Following are examples of different objects and the expected reflectivity of each.
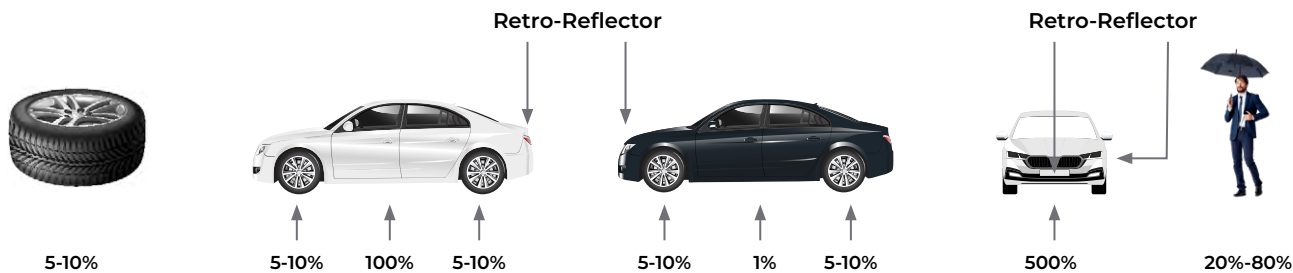

Figure 33: Objects with various reflectivity levels

Figure 34: Point cloud detection of various reflectivities

The AV Challenge: Responding to a Highway Safety Event

The automated driving platform is designed by functional safety engineers who analyze all possible driving scenarios. Each expected risk must be mitigated by the system with the required redundancy. Hundreds of such scenarios are analyzed and then translated to system and LiDAR requirements. Following is one scenario of an edge case that is helpful in defining LiDAR design requirements.

This scenario describes a vehicle (the "follow vehicle") that is driving at 130km/h (80mph) and follows a lead vehicle that is driving at similar velocity. The follow vehicle is driving between two vehicles in adjacent lanes that do not allow it to change lanes. The lead vehicle occludes an object at 100m that is just above 14cm high.

Figure 35: Occluded obstacle on a highway

The lead vehicle changes lanes at the last minute, revealing an object at 120m. It cannot identify whether or not the object is over-drivable.

Safety Incident at 130Kmh (80mph)

Obstacle is NOT detected by two vertical pixels — **> 120m** — VRes 0.05°

Obstacle is detected by two vertical pixels following vehicle — **100m**

Obstacle is detected by at least 4oo7 frames at 99.99% TPR (@20fps = 0.35sec) decision on non-overdrivable obstacle! — **87.5m**

Applying safety brakes
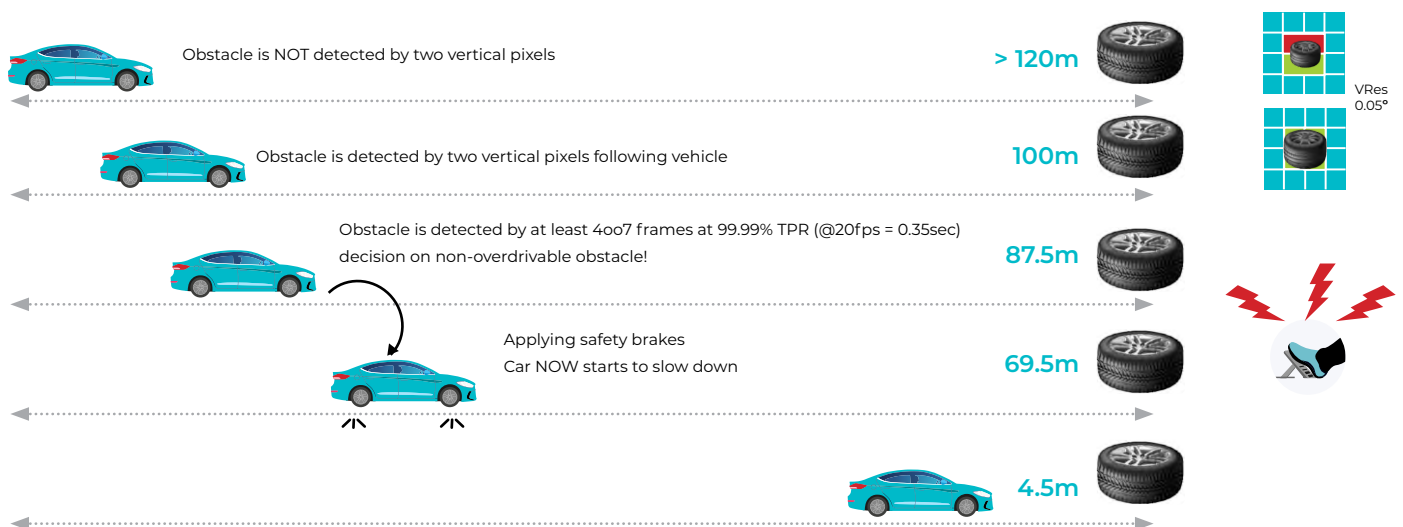Car NOW starts to slow down — **69.5m**

4.5m

Figure 36: Safety incident at 130km/h

Due to the large horizontal and vertical ROI Field of View, the object is detected immediately as the follow vehicle starts to change lanes, regardless of the slope of the road or its curve. Due to the high horizontal resolution, the object could be detected even when the object is partially obscured.

The system needs to collect 4 out of 7 frames in which the hazard is detected in order to have sufficient confidence when applying the emergency brakes. With 0.05° vertical resolution, at 120m the follow vehicle cannot determine whether the 14cm high object is over-drivable. Only at 100m can the follow vehicle detect 4 out of 7 frames with high confidence.

Figure 37: Detecting 4oo7 frames with high TPR

Figure 38: Detecting 4oo7 frames with high TPR

At 20FPS, the follow vehicle would travel 12.5m at 130km/h for 0.35 seconds. At that range, it will apply emergency braking whose full mechanical delay is about 0.5 seconds, traveling another 18m. Only at 70m will the car start to decelerate. With very strong and uncomfortable braking, it will stop 4.5m away from the object.

The following sensitivity analysis of the parameters that can improve reaction time and increase braking distance indicates that the maximum benefit is derived from a 100% improvement in vertical resolution, as opposed to range or frame rate.

CURRENT PARAMETER		@ 100% IMPROVEMENT		RESULT	CONTRIBUTION
Range 200m @10%	X2	**Range** 400m@10%		No Improvement in braking distance or reaction time	Limited by resolution
Frame Rate 20 Hz	X2	**Frame Rate** 40 Hz		Braking distance improves by 6.25m; reaction time improves by 0.17 seconds	Minor
Vertical Resolution 0.05°	X2	**Vertical Resolution** 0.025°		Braking distance improves by 100m; reaction time improves by 2.7 seconds	Major

While the above example is an extreme scenario, there are several take-aways from this use case:

1. If the follow vehicle had detected the object at less than 95m, it would not have avoided a collision. A clear line of sight of 100m is needed at 130km/h assuming worst case scenario.

2. If the road surface was wet, the required braking period would have been longer, since the car should have driven slower. It is essential to always consider varying road conditions.

3. The above example assumes 20FPS. At a lower frame rate a longer clear line of sight is needed.

4. Even if the system is capable of detection at 500m range, the object will be detected only with sufficient resolution and frame rate.

5. Some driving decision measures could be helpful in such cases, such as automatically slowing in reaction to a follow vehicle changing lanes.

6. The suspension height and braking capability of every vehicle must be taken into consideration when defining the LiDAR performance requirements in an emergency braking scenario.

InnovizTwo: Optimized for High-Speed Highway Driving

InnovizTwo has been designed to ensure safe and comfortable driving. It features unprecedented resolution, frame rate and detection range. The high vertical resolution allows detection of small objects at long range.

Its design is based on valuable experience gained from InnovizOne, which will be included in BMW's first-generation autonomous cars. Its design also benefits from thousands of hours of discussing the requirements of leading automotive OEMs.

InnovizTwo uses the same cost-effective 905nm laser technology and proprietary ASIC as InnovizOne, plus a detector with a higher photon detection efficiency and a next-generation scanner. The hardware design is simplified with fewer components. It provides higher performance than InnovizOne and is expected to cost 70% less.

InnovizTwo supports the -40° to +85°C temperature range, includes a Gigabit Ethernet hardware interface, and a direct software interface to the vehicle's domain controller. It also includes over-the-air (OTA) upgradable firmware which enables remote functionality updates and upgrading a vehicle from L2+ to L3 without changing any hardware. The system software is designed from the beginning according to ASPICE, ISO 26262, and MISRA automotive software development standards. The LiDAR will include complete cybersecurity protection. See https://innoviz.tech/innoviztwo for more information.

The InnovizTwo point cloud will be a 30x improvement over the following images from InnovizOne.

Conclusion

In this paper we presented the need to design an autonomous vehicle for safe, comfortable highway driving. We presented and explained the main considerations that affect LiDAR performance and how they are interrelated.

The design of the car can easily determine the design of the LiDAR. Because a Level 3 LiDAR should be unobtrusive, the mounting position of the LiDAR has a significant impact on both the LiDAR design and the vehicle design. We suggest that OEMs consider all possible mounting locations in all models in which the LiDAR will be installed.

Highway driving at high speed necessitates extended detection range; however, this cannot come at the cost of insufficient resolution or a reduced frame rate. Environmental conditions, such as weather, lighting, road conditions, and blockage of the sensor also can reduce system performance, resulting in a longer braking time. LiDARs must be designed to take into account dynamic conditions while ensuring optimal performance.

We suggest that autonomous vehicle manufacturers carefully evaluate the claims of LiDAR vendors since safe and reliable autonomous system performance requires many trade-offs.

About Innoviz Technologies Ltd.

Innoviz Technologies Ltd. was founded in 2016 and has raised more than US$500 million to date. The company went public in April 2021 and is traded on the Nasdaq stock exchange under the INVZ ticker symbol. As of this writing, Innoviz employs more than 350 full-time employees and continues to hire aggressively.

Innoviz was founded by members of Unit 81, the elite technology unit of the Israel Defense Forces. Their specialty is solving complex technological challenges by employing a multi-disciplinary approach. The unit's motto is "making the impossible possible". Members are obsessed by solving big problems and do not accept "no" for an answer.

About one-quarter of Innoviz's R&D staff are former 81 members. Many of them worked together during their military service. Many are recruited immediately after finishing their service; others come after working in startups, established high-tech companies, or after receiving their doctorate.

The 81 ethos is what makes Innoviz different from the dozens of other LiDAR companies (including a handful that recently went public). Unit 81 graduates know how to work as a team as they think out of the box. They are bored unless they are pushing the limits. This is what drives Innoviz.

About the Author

Omer David Keilaf has spent 21 years driving cutting edge technologies from inception to commercialization. As co-founder and CEO of Innoviz Technologies Ltd, Mr. Keilaf is helping enable the mass-commercialization of autonomous vehicles by providing high-performance LiDAR at mass-market prices. Mr. Keilaf was named CEO of the Year by the 2018 Image Sensors Europe Awards.

Prior to co-founding Innoviz, Omer served as an officer in Unit 81, the elite technological unit of the Intelligence Corps of the Israel Defense Forces (IDF), where he served as Project Manager and System Architecture Manager. He then went on to a distinguished career in opto-mechanics, electrical engineering, MEMS and more. He held senior leadership roles, including System and Product Team Manager at Consumer Physics and R&D manager at STMicroelectronics. Omer holds an MBA, M.Sc., and B.Sc. in Electrical Engineering from Tel Aviv University, where he has also served as a lecturer.