Advancing LiDAR perception in degraded visual environments: a probabilistic approach for degradation analysis and in inference of visibility

Karl Montalban

To cite this version:

Karl Montalban. Advancing LiDAR perception in degraded visual environments: a probabilistic approach for degradation analysis and in inference of visibility. Robotics [cs.RO]. ISAE Université de Toulouse, 2023. English. NNT : 2023ESAE0044 . tel-04260401

HAL Id: tel-04260401
https://laas.hal.science/tel-04260401v1

Submitted on 27 Oct 2023

THÈSE

En vue de l'obtention du

DOCTORAT DE L'UNIVERSITÉ DE TOULOUSE

Délivré par l'Institut Supérieur de l'Aéronautique et de l'Espace

---

Présentée et soutenue par

Karl MONTALBAN

Le 18 septembre 2023

Perception LiDAR dans des Environnements Visuels Dégradés : Une Approche Probabiliste pour l'Analyse de la Dégradation et l'Inférence de la Visibilité

Ecole doctorale : **AA - Aéronautique, Astronautique**

Spécialité : **Photonique et Systèmes Optoélectroniques**

Unité de recherche :
ISAE-ONERA OLIMPES Optronique, Laser, Imagerie Physique et Environnement Spatial

Thèse dirigée par
Simon LACROIX et Nicolas RIVIERE

Jury
M. Andrew Mickael WALLACE, Rapporteur
M. Fawzi NASHASHIBI, Rapporteur
Mme Mariana BATISTA CAMPOS, Examinatrice
M. Thierry PEYNOT, Examinateur
M. Simon LACROIX, Directeur de thèse
M. Nicolas RIVIERE, Co-directeur de thèse
M. Tristan ALLOUIS, CTO YellowScan, Invité
M. Paul-Édouard DUPOUY, Encadrant ONERA, Invité
M. Dinesh ATCHUTHAN, Encadrant EASYMILE, Invité

Advancing LiDAR Perception in Degraded Visual Environments : a Probabilistic Approach for Degradation Analysis and Inference of Visibility

Karl Montalban
karlmontalban@gmail.com

October 15, 2023

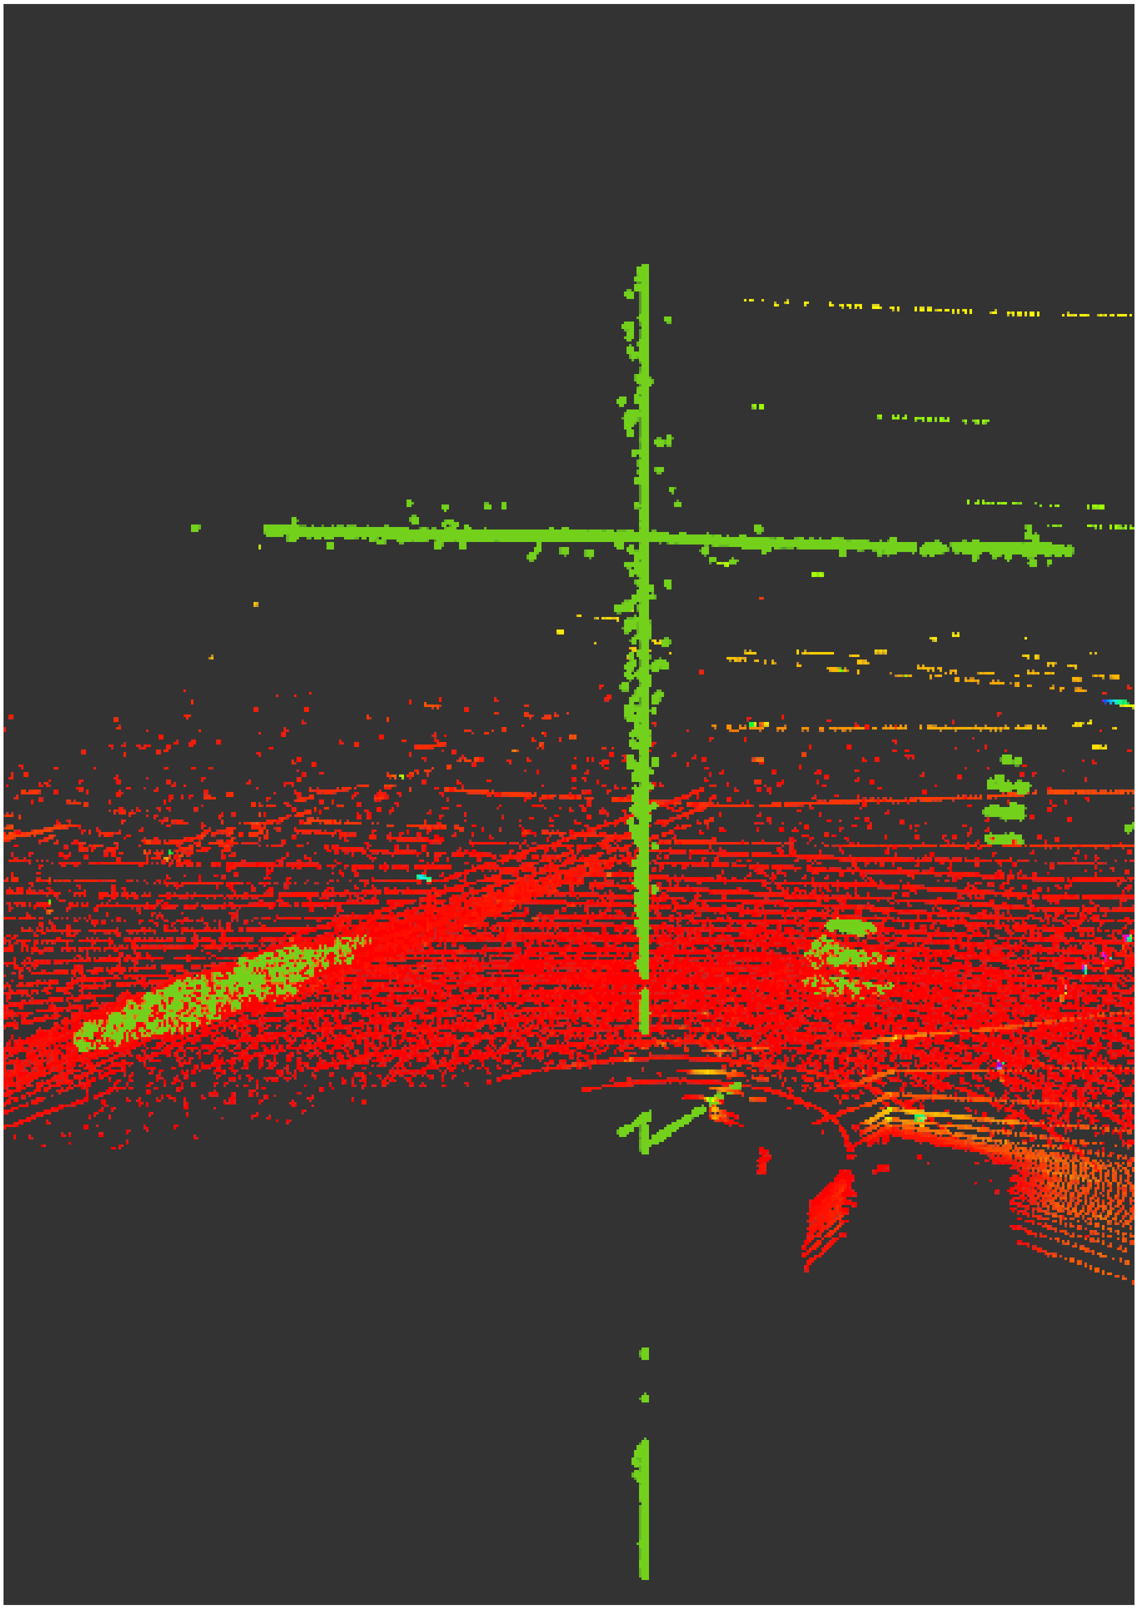

Divine intervention during an imprecise point cloud manipulation

Contents

Chapter 1

Introduction

To ease the reading with (I hope) visual appreciation, every chapter of this thesis is introduced with a painting of my own.

This thesis was part of the CIFRE[1] french program, meaning that I was employed by a Company (EasyMile) and supervised by one or several laboratories (two in my case, ONERA and LAAS-CNRS). This type of configuration is great in the sense that it opens for work at the frontier between industry and research.

1.1 Autonomous driving (AD) ...

Before everyone became really excited about watching a movie while driving at 130 km/h on the highway without looking at the road or even sleeping, autonomous driving (AD) started quietly in 2004, during the DARPA Grand Challenge [1]. Well, at least the autonomy applied to land transport, because the aeronautical, aerospace and maritime sectors have different histories. This challenge showed that a combination of sensors (LiDARs, cameras, GPS, gyroscopes, accelerometers and odometers) allowed the Stanley vehicle to finish the race at the first place. Since, Waymo (company of Google) has developed its own technology and is considered today as the major actor in this field with the commercialization of an autonomous robot-taxi product in San Fransisco. But others joined the race. And whether through technology or market competition, many companies actually try to thrive in the AD business. Along with emerging technologies in software, hardware and also a great amount of funds, AD solutions have greatly improved and can now perform impressive tasks, as shown by videos showcasing the performance of recent Autonomous Vehicles (AVs) [2, 3]. Of course, failures exist as well [4, 5]. The progress of AD over the years is evident, but there are still challenging edge-cases that present unresolved limitations.

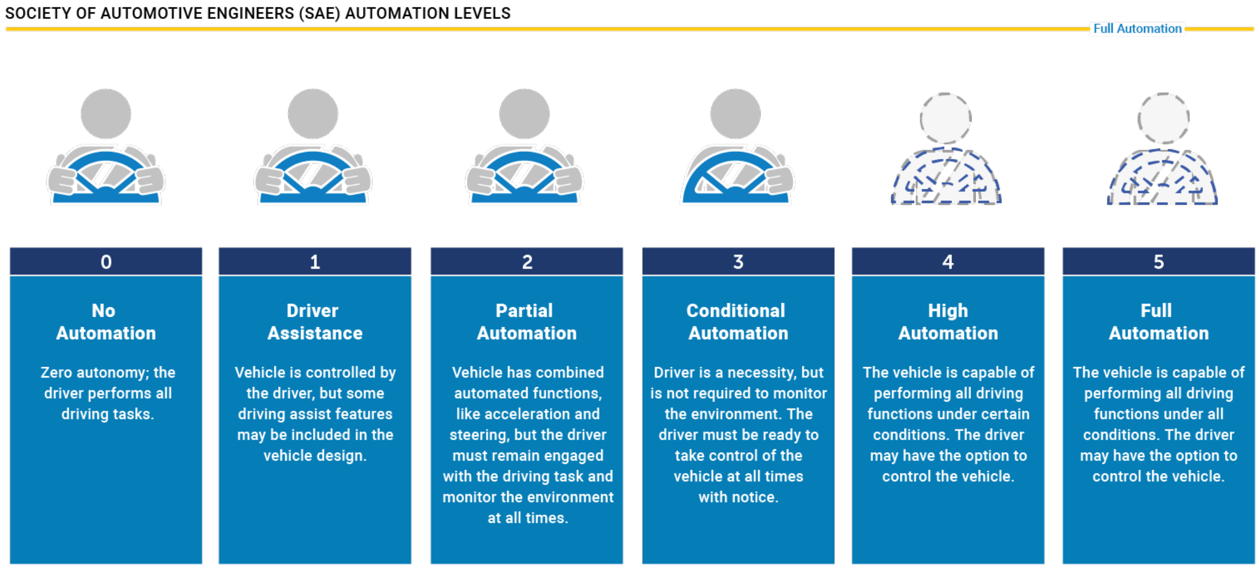
SOCIETY OF AUTOMOTIVE ENGINEERS (SAE) AUTOMATION LEVELS

0	1	2	3	4	5
No Automation	**Driver Assistance**	**Partial Automation**	**Conditional Automation**	**High Automation**	**Full Automation**
Zero autonomy; the driver performs all driving tasks.	Vehicle is controlled by the driver, but some driving assist features may be included in the vehicle design.	Vehicle has combined automated functions, like acceleration and steering, but the driver must remain engaged with the driving task and monitor the environment at all times.	Driver is a necessity, but is not required to monitor the environment. The driver must be ready to take control of the vehicle at all times with notice.	The vehicle is capable of performing all driving functions under certain conditions. The driver may have the option to control the vehicle.	The vehicle is capable of performing all driving functions under all conditions. The driver may have the option to control the vehicle.

Figure 1.1: Levels of autonomy defined by the Society of Automotive Engineers (SAE) [6].

All tasks are not equal and 5 levels of autonomy have been engraved to divide technologies, shown fig.1.1. Advanced Driver Assistance Systems (ADAS) solutions focus on automating features to existing personal cars, with for instance automatic parking or automatic highway driving, corresponding to level 3. Tesla cars are level 3 autonomous cars [7], if the Full-Self-Driving (FSD) option is added, for a supplement of 15k€. Level 4 brings autonomy one step further where an operator is on board but the vehicle should perform all driving functions on its own. The last level, holy grail of AD, means complete autonomy, without any operator and in all conditions.

---

[1] ”Conventions industrielles de formation par la recherche”

This last concern is most suited in this introduction as it relates to the heart of this thesis. The level 5 objective is quite ambitious, perhaps even too much. Basically, it consists on replacing human driving to reduce costs and potentially increase safety and availability on transportation tasks. However, achieving such service that matches or surpasses human efficiency necessitates to address a wide range of limitations. For example, the number of automatic stops should be minimized, and if a vehicle does stop, it shall be capable of autonomously restart. The vehicle shall not halt for obstacles that a human would overlook, such as insects, small vegetation, rain, clouds, high sidewalks, or speed bumps (the list is long). The vehicle shall maintain an appropriate speed while actively avoiding collisions and navigating around obstacles to stay on its designated path. The vehicle shall adapt to various loads of people of goods. The vehicle shall evolve in mix-traffic situations. This list of requirements is also long. While progress in the AD field has been significant, there are still many challenges that need to be overcome, involving profound research topics.

EasyMile company, sponsor of this thesis, develops level 4-5 autonomous solutions since 2014 and today, two main products are in production, the EZ10 passenger shuttle and the TractEasy tow tractor, shown fig.1.2. Designed with dedicated sensor sets and software stack, they are respectively deployed to transport passengers in different kind of sites (recently on public roads) and to transport goods in industrial scenarios or airports. Additionally, EasyMile is involved in several projects to automate different kind of platforms (bus, tram, trucks). One concept that differentiates it from other companies and makes it promising is its approach related to safety. Indeed, a real emphasis is put on reaching safety regulations [8], which is critical for the deployment of autonomous solutions on public roads [9, 10].

(a) EZ10

(b) TractEasy

Figure 1.2: EasyMile vehicles [11].

In order to achieve the needed autonomy tasks, these systems rely on sensors and perception algorithms to understand their surroundings. Then, artificial intelligence (in the large sense) translates the perception data into decisions for the vehicle to navigate autonomously. The perception capabilities of the vehicle are thus one of the very first blocks that needs to be addressed as their output is used extensively in the rest of the software stack in the vehicle. Perception algorithms rely on perception sensors which measure features of the environment.

The biggest limitation of AD (and probably robotics in general) lie in this so-called environment because it can severely impact the sensors data. Sporadic events, artificial or not, take the algorithms out of their nominal behavior and reduce the availability of the vehicle. Weather phenomena is one of the major example where the water particles involved interact with the sensors signals and alter the data. The next section introduces the impacts caused by weather conditions on perception and thus the availability of an AV.

1.2 ... in Degraded Visual Environment (DVE)

Achieving full autonomy requires enormous amount of development (and funds). This is probably why we do not see any AV on our streets except for Waymo in California and Tesla. One major obstacle heavily slows down progress : the problem of the environmental conditions. Going by many names, All weather, Harsh weather, Degraded weather, we choose to state it as Degraded Visual Environment (DVE).

First, it is important to recall that DVE is already an issue in level 0 autonomy, also known as normal driving, and that the principles involved are the same. Our eyes provide us with the necessary perception data to drive in the majority of the conditions. Light travels from the sun (or artificial sources) through the environment to reach our eyes, resulting in accurate visual data in nominal conditions. And, intelligence enables us to take decisions accordingly and actually drive. In DVE scenarios, particles alter the behavior of light and thus our perception of the environment. The resulting data degradation makes it difficult to perform the visual tasks of object recognition and risk assessment correctly. In addition, weather conditions have an impact on road conditions, making it difficult for the vehicle to adhere to the road. As a result, the number of road crashes is increased in DVE conditions and is directly correlated with weather variables [12].

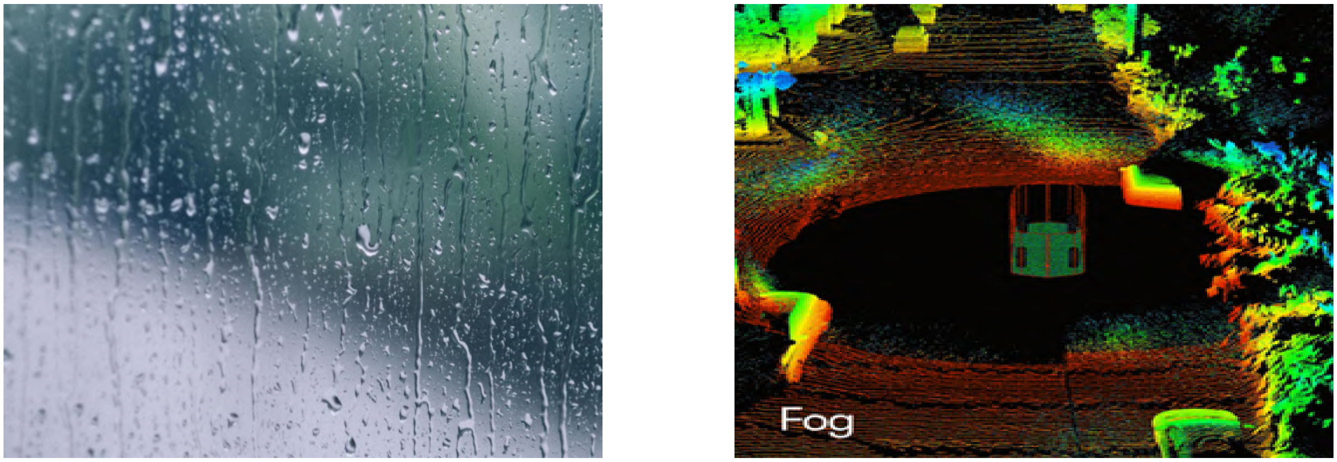

Figure 1.3: Image and point cloud data altered by rain and fog conditions respectively, excerpt from [13].

Just like for us, the perception sensors of an autonomous vehicle are impacted by the degraded environment (as shown on fig.1.3), reducing the capacity to navigate in such conditions [14]. Active or passive perception sensors, much needed for autonomous navigation, use different technologies but all imply the manipulation of electromagnetic waves, namely radio, light and ultrasonic waves. Evolving in the atmosphere, these sensors have to face the changing nature of the environmental conditions because the signals involved have altered behaviors in the presence of particles, just like our eyes. All the sensors technologies are impacted by DVE conditions but their intrinsic parameters cause different results. For example, LiDARs are heavily impacted by small-sized aerosol particles because their wavelength is close to the particle sizes. Radio waves, on the other side, are of larger wavelengths so radar sensors are the solution for autonomous driving in all weather conditions. This kind of statement is quite common in marketing, but relatively dangerous because also wrong. Although the impact difference is true for small size particles, radio signals must still be attenuated and the degradation should be more quantified before being too confident. Additionally, rain consists of larger particles and alter the data of all perception sensors, resulting in reduction of range and lower signal to noise ratio (which is already an issue with radar sensors). Anyway, these degradations of sensors data result in degradation of perception algorithms performances and thus of the vehicle's capacity to navigate autonomously. Although all the sensors technologies are impacted by degraded environmental conditions, the work of this thesis focuses on LiDAR technology.

False positive echoes caused by DVE particles can be deleted with denoising methods. Their goal is to clean the point clouds before they are handled by other autonomy features. But as the changing nature of the events result in different levels of degradation, these methods (as well as other algorithms of the autonomy software stack) should adapt to the changing DVE properties to increase their efficiency. In any case, the vehicles lack of intelligence and information about the environmental conditions around them. Most of the algorithms useful to address these problems have parameters which could be changed according to an evolution in the conditions impacting their input data. Therefore, the classification of DVE properties is key to enhance the performances of AD. This thesis contributes within this particular context.

1.3 Overall approach

This thesis started with the objective of quantifying the degradations induced by DVE on LiDAR point clouds. To do so, the use of climatic chambers arrived quickly. Their capability to generate controlled conditions and assess the evolution of the environmental conditions inside is highly valuable in our objective of quantifying the degradation. In this thesis, two experimental campaigns in climatic chambers are presented.

Following specific test protocols, probabilities are drawn to assess degradations concerning the detection capabilities of LiDARs. Using probabilities enables the comparison of sensors with different technologies and helps to approach safety regulations, essential for the deployment of autonomous driving solutions. These campaigns showed that different LiDAR technologies lead to different degradation results, and that unknown internal sensor factors play a major role. Finally, statistics regarding the point clouds information are studied. Interesting distance distributions of false positive points in fog and smoke conditions with the Ouster OS1-128 LiDAR are revealed, leading to a Bayesian inference algorithm.

Bayesian inference is a major concept in mathematical statistics to estimate information and is used in many scientific fields and also philosophy, medicine, sport and law. As opposed to machine learning, composed of many techniques (the most famous one being deep learning today), it allows the manipulation of probabilities and requires precise modelling of the data. Challenges lie in the explicability of machine learning techniques, making them potentially constraining in this respect [15]. And, the choice of using deep learning or not was guided by the additional components to be implemented in the already existing sensor stack for the deployment of those methods in a production context and the lack of available data that we had to contend with. On the other side, manipulating Bayesian probabilities brings closer to safety standards, which is needed to certify the safety of an AD solution.

Therefore, a Bayesian inference model is developed and represents the main contribution of this work. Its goal is to classify visibility from the identification of noise detections from a 3D-LiDAR in degraded visibility conditions. Our approach tries to reach generalization as much as possible and is interesting because it makes use of a sensor that is already embedded on the company vehicles and used for other perception algorithms.

Finally, note that the physics involved between LiDAR signals and DVE particles is a quite complex topic, which is not covered in depth in this work. Instead, it is rather considered as the context behind our results and is sometimes used to explain some of the phenomena.

1.4 Contributions

To overcome all the issues caused by DVE conditions, a single thesis might not be enough, I hope not to disappoint anyone. Our work is addressed on signal processing techniques using the data produced by the LiDAR sensors (point clouds). Altered by the DVE conditions, point clouds are analyzed to provide a weather classification algorithm. Our contributions are :

- A quantitative analysis of the impacts caused by artificial DVE conditions on LiDAR point clouds.

- A Bayesian inference model aiming to retrieve visibility value from LiDAR point clouds.
- Extensions on how to use our inference model to enhance perception awareness in DVE for AD.

1.5 Outline

This manuscript is divided in two parts. The first one presents the context encompassing this thesis, with the description of the 3D-LiDAR technology, especially the one designed for AD. Then, DVE scenarios are presented, including only scenarios that have been experimented : rain, fog and smoke.

The second part focuses on our contributions. It starts by presenting perception algorithms, with their relation with DVE conditions. Then it describes the State-Of-The-Art (SOTA) methods related to our objective and gives an overview of our approach, with in-depth explanation of the methodology, consistently applied throughout the developments. Next, quantitative results drawn from experiments performed in artificial DVE conditions are presented. These results show interesting data, leading to the development of a visibility inference model, presented in the following chapter. Our final contribution involves potential extensions of our approach and is given in the last chapter before this thesis is concluded.

1.6 Publications and participation to conferences

The works of this thesis have been published in the following scientific publications [16, 17, 18] and presented in ITBMS 2022 [19] and ICRA 2023 [20] international conferences.

Part I

Context

Chapter 2

3D-LiDAR technology

The first lidar (light detection and ranging) system ever created was designed to track satellites from the Earth [21]. The system uses laser (light amplification by stimulated emission of radiation, invented by Gould in the beginning of the '60s) with appropriate optics and electronics to obtain a target distance by measuring the time for the backscattered light to return to a detector, known ad the time-of-flight (ToF) technology. 60 years later, lidars are used in multiple types of applications and situations (terrestrial, maritime, airborne, space) [22, 23, 24, 25, 26, 27, 28]. Surveying, geodesy, geomatics, archaeology, geography, geology, geomorphology, seismology, forestry, atmospheric physics, defense, space and robotics commonly use lidar systems.

Recently, in the world of AD, lidars are used to provide point clouds of the environment. Their capabilities to function as 3D-scanners has garnered significant attention, primarily due to their capacity to generate precise and high-resolution 3D measurements. To become a 3D-liDAR[1], the ranging capability of a 1D-Rangefinder LiDAR is extended with optical and mechanical components to spend laser beams and obtain measurements within diverse Field-of-Views (FoVs) surrounding the sensor. Multiple designs exist for the manipulation of laser beams, detection schemes and overall creation of point clouds.

Fig.2.1 shows a schematic representation of the different elements involved in a 3D-LiDAR system. Produced point clouds are the result of all the different blocks. Lasers, detectors and optical components in a 3D-imaging architecture constitute the core hardware of the technology. The laser beams are emitted from the sensor and propagated in the atmosphere where they can interact with various molecules and particles. Solid targets in the path of the beams reflect the laser light. After back propagation in the atmosphere, the signals are received on the detectors through a detection scheme. Finally, a measurement principle is applied to retrieve information about the impacted targets (range, signal intensity). In this chapter, these different blocks are described. Because the main contribution of this thesis is carried out on a spinning LiDAR, an emphasize is done on its technology.

The work of this thesis focuses on LiDAR sensors designed for AD but tries to give enough details about the overall technology.

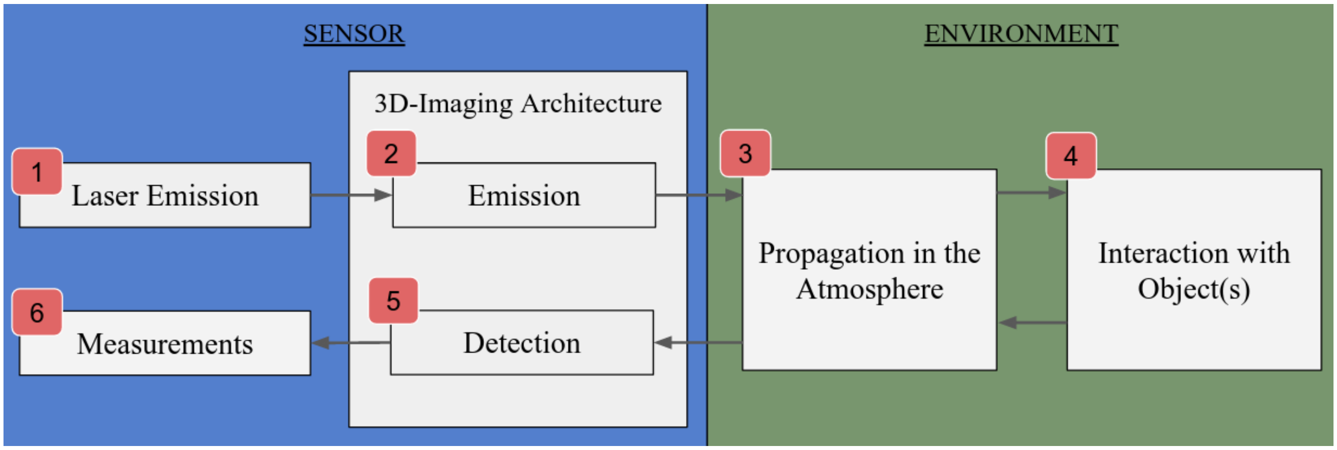

Figure 2.1: Block diagram of a LiDAR principle

3D-LiDAR systems today are relatively expensive due to their complex composition, which can slow down the widespread adoption of this technology. However, historical trends have demonstrated a consistent decrease in costs as technological advancements are made and production processes become more efficient. With the rise of autonomous driving and all the other applications using LiDARs, the issue of cost is expected to diminish significantly. Recently, even smartphones are equipped with LiDARs which confirms the fast development of this technology and its adoption in different domains [29].

---

[1] The acronyms LIDAR, LiDAR or even LADAR (for laser radar) refer to lidar used as a 3D imager or scanner whereas the term lidar is often used for atmospheric lidar and more physical measurements. In the rest of this thesis, the term LiDAR is used to address a 3D-LiDAR imager, as it is the common practice in the literature.

2.1 Laser emission

The laser emission block is the initial component in the overall LiDAR processing chain, responsible for generating the laser light or beam that forms the core of the system. A laser beam is the result of optical amplification of photons inside a gain medium which results in high-energety spatially and temporally coherent light. The spatial coherence of laser allows the beams to be focused to precise spots while temporal coherence results in monochromatic signals. Thanks to these properties, laser emission is used today in a wide range of applications such as information technology (optical disk drives, communication through fiber optics and free space), printing, welding, surgery, military, displays and ranging measurements, the latter being of course our concern.

In general, a laser beam can be defined by its wavelength, amplitude, temporal width and divergence. The choice of wavelength depends on the characteristics of the gain medium and the cavity size. Commonly used wavelengths for LiDAR applications fall within the near-infrared spectrum. For example, 850 nm, 905 nm, 1064 nm and 1550 nm are the most used wavelengths for AD. However, specific applications require other wavelength due to different environmental constraints, such as submarine operations which require the utilization of 534 nm to minimize absorption effects and achieve optimal performance [30]. Laser light is often sent in the form of pulses of specific amplitude (pulse energy) and temporal width, as shown on fig.2.2.

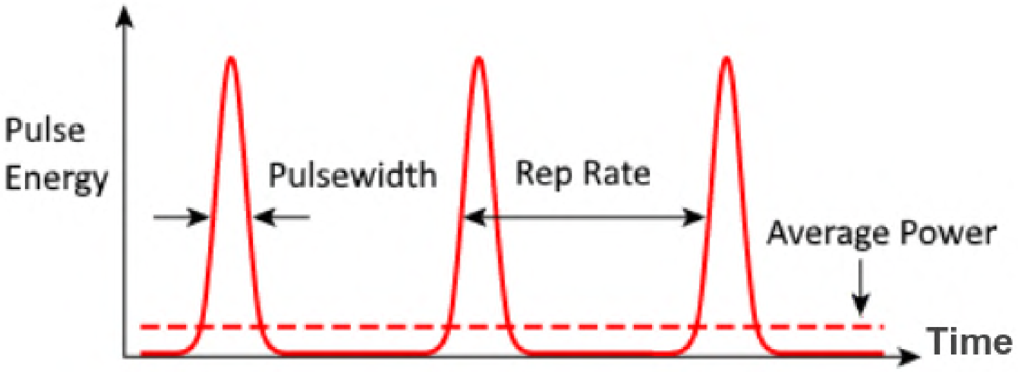

Figure 2.2: Energy of emitted pulsed laser shots in time.

Fig.2.3 shows the geometric properties of an emitted laser beam, where the divergence of the beam is the angle Θ. This divergence defines the footprint of the laser once it has travelled over a certain distance which impacts LiDAR precision and the multi-echo concept (explained in sec.2.7).

Because of their small size and efficiency, two main semiconductor laser sources are commonly used to produce laser beams for LiDAR: laser diodes and Vertical Cavity Surface Emitting Laser (VCSEL) diodes. A laser diode is a semiconductor component, similar to a Light-Emitting-Diode (LED), capable of creating laser light from an electric current. VCSEL diodes, while belonging to the category of laser diodes, exhibit a distinct characteristic in which the emission of laser light occurs perpendicular to the surface, in contrast to conventional semiconductor lasers that emit light from the edge. Moreover, VCSEL diodes are notable for their utilization of more efficient fabrication processes, despite the lengthier and costlier nature of these processes.

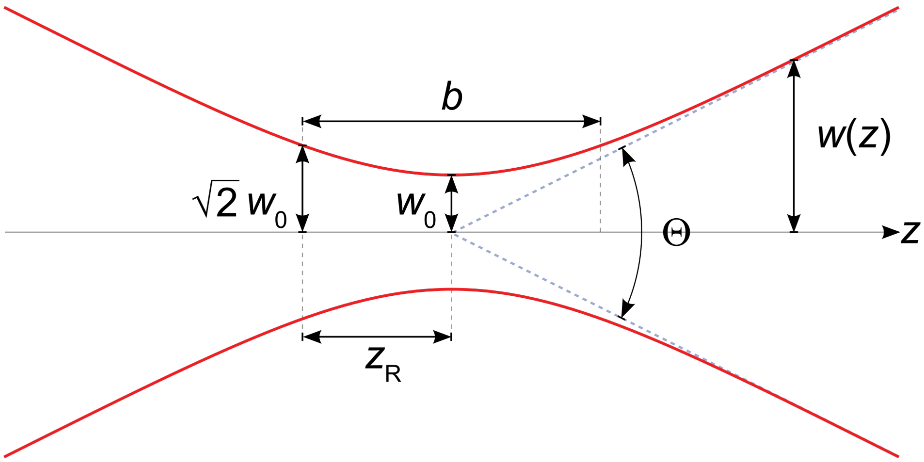

Figure 2.3: Geometric properties of a gaussian beam width $w(z)$ as a function of the distance z, which forms a hyperbola. w_0 is the beam waist, b is the depth of focus, Z_R is the Rayleigh range and Θ is the total angular spread.

Not all laser sources can be used to build LiDAR systems and especially the ones designed for autonomous driving applications. First, the atmosphere poses a significant limitation for LiDAR systems due to scattering effects. Some wavelengths are unsuitable for LiDAR systems because they are prone to complete absorption by the atmosphere (more details in sec.2.4). Furthermore, eye-safety regulations impose limitations on both the amplitude of emitted signals and the choice of wavelengths as the human eye is not sensitive to all wavelengths equally (following safety standard [31]). Fig.2.4 shows the transmittance of light inside elements of the eye and depending on wavelength from 300 nm to 2500 nm. The transmittance is relatively higher at 900 nm than at 1500 nm which means that the eye absorbs more at the latter wavelength and thus higher energy pulses could be used for LiDAR systems operating at 1500 nm without exceeding eye-safety limits.

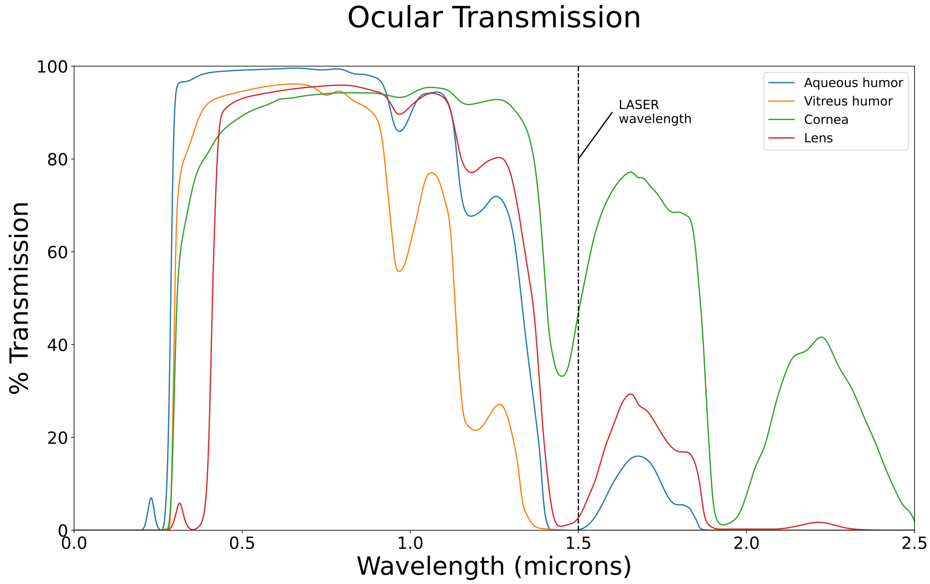

Figure 2.4: Transmittance of light plotted inside various elements of the eye and depending of wavelength [**zuclich˙ocular˙1995**].

2.2 3D-Imaging architecture

Once the laser beam is emitted, it is steered by the LiDAR system to perform scanning. This section depicts the different LiDAR technologies concerning beam steering approaches, which can ultimately define the overall imaging architecture. Two main categories are commonly recognized: point-wise scanning and flash LiDAR where the laser beam is respectively shot as a narrow beam or a highly divergent one. Extensive descriptions of imaging architectures of LiDARs are avaliable in [28, 26, 32]

It is worth noting that the term ”solid-state” LiDAR has emerged as a category that aims to differentiate from mechanical LiDAR technologies (that are potentially prone to shorter lifetime). However, the inclusion of LiDAR systems in the ”solid-state” category is not clearly defined and can sometimes be driven by marketing considerations rather than scientific criteria. The ”solid-state” category is thus not further addressed here.

2.2.1 Point-wise scanning

In point-wise scanning LiDAR systems, narrow laser beams are sent out of the sensor and the 3D-scanning is done by moving them inside the FoV with beam steering techniques.

- Mechanical scanning LiDAR

Mechanical scanning LiDARs leverage the reflection effect of mirrors to manipulate the path of laser beams and precisely steer them in desired directions. The mirror can then be paired with a rotating mechanism to rotate around any axis and extend the sensor’s FoV. Classical static 3D scanning LiDAR use Polygonal or Galvanometric mirrors to steer the laser beamsin the vertical axis while the entire LiDAR system mechanically rotates around the horizontal axis to complete 3D-scanning, see fig.2.5.

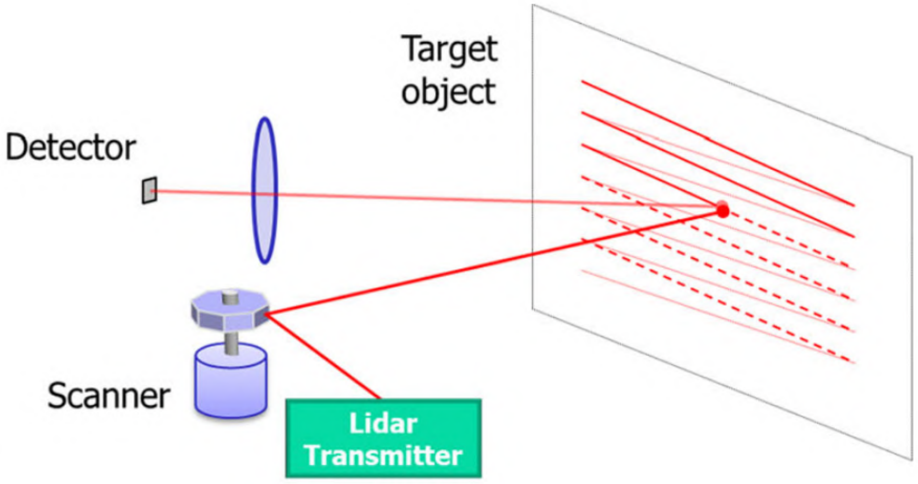

Figure 2.5: Schematic representation of a mechanical beam steering with a mirror, excerpt from [33].

Companies like Riegl, Faro and Leica are big players in this field. These LiDAR systems provide very accurate measurements but require a lot of time to produce 3D point clouds, see fig.2.6.

(a) Riegl VZ (b) Leica RTC (c) Faro Focus

Figure 2.6: Terrestrial scanning LiDARs.

In robotics, the dynamic nature of the environment and platform movement present challenges for LiDAR systems that rely on slow rotation around the vertical axis. Waiting for a complete rotation in such scenarios may be impractical due to the need for real-time data and rapid decision-making. To overcome this limitation, one option is to orient the mirror to primarily scan in the horizontal axis while scanning only a single (or a few) elevation angles. The company Sick proposes such sensors (fig.2.7) and focuses on safety approaches, mandatory for AD or industrial applications.

Figure 2.7: Sick OutdoorScan

This architecture is also commonly used in embedded sensor in mobile mapping activities. To scan the ground surface and the objects on it, the LiDAR performs line-scanning while the plane or drone evolves above the area of interest, see fig.2.8.

(a) Riegl Vux (b) YellowScan Explorer

Figure 2.8: Airborne scanning LiDARs.

Two mirrors can also be used, one for each axis (horizontal and vertical), to enable scanning of a rectangular FoV. This configuration captures the frontal environment of a vehicle, with for example the ADAS application of highway driving, see fig.2.9.

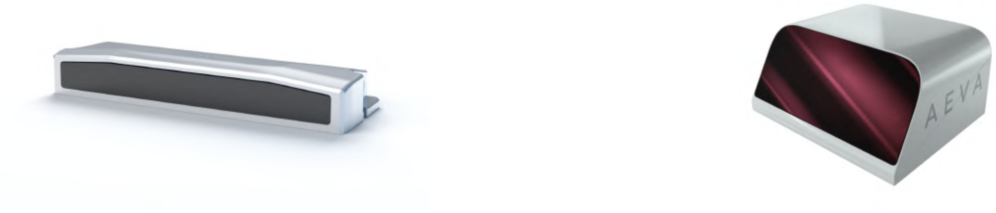

(a) Luminar Iris (b) Aeva Aeries

Figure 2.9: Mechanical scanning LiDARs with a rectangular FoV.

One original architecture uses Risley prisms to perform 3D-scanning [34]. Two Risley prisms rotating in relation to each other can create non repeated scan patterns which periodically explore an entire area with the laser beam, see fig.2.10. The main interest of this technology is that the sensor can capture data from its entire FoV. This effect is particularly useful to perform scanning when it is needed to perceive challenging structures like through obscurant media, vegetation or holes [35]. The Livox company is the most famous company providing such systems for autonomous driving.

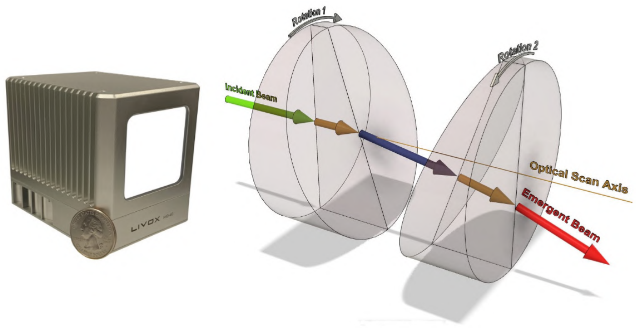

Figure 2.10: Livox LiDAR and schematic representation of the Risley prisms beam steering, excerpt from [34].

- MEMS LiDAR

Micro-Electro-Mechanical-Systems (MEMS) are small mirrors precisely controlled by electromagnetic voltage. In both oscillating and non-oscillating modes of operation, MEMS LiDAR systems have the flexibility to set and spread the angles of reflection, allowing the laser beam to cover a wider area in the scene. Although its monitoring is more complex compared to a classic mirror, a MEMS mirror offers several advantages as it is smaller, cheaper and capable of dual-axis scanning with a single mirror unit, see figures 2.12. Commercially available MEMS LiDARs are presented on fig.2.11.

(a) Innoviw two (b) Aeye 4SightM (c) Blickfeld cube

Figure 2.11: MEMS LiDARs.

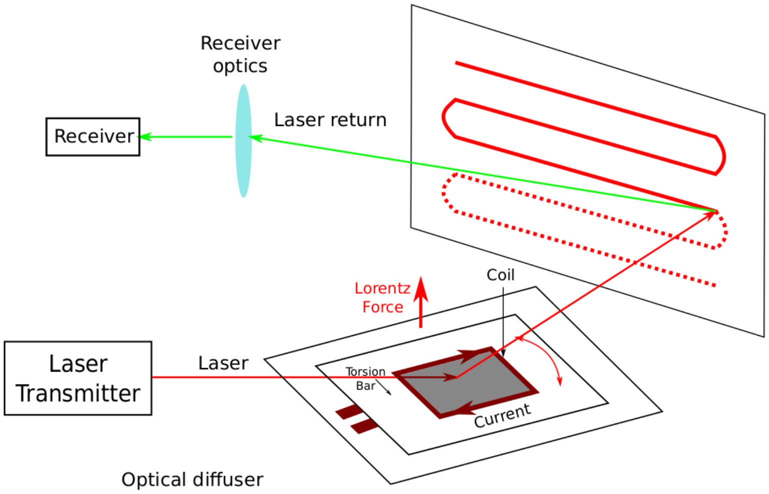

Figure 2.12: Schematic representation of laser beam steering using a dual-axis MEMS mirror, excerpt from [28].

- Spinning LiDAR

A spinning LiDAR is composed of vertical pairs of emitters and detectors which spin at 360° around the vertical axis and perform range measurements for each heading angle. Because the emitter-detector pairs are optically conjugated and their elevation angle remain static, the produced point cloud is said to be composed of layers inside which all points have a common elevation angle value. The addition of pairs increases the vertical resolution of the point cloud while the horizontal resolution is defined by the rotation speed. Spinning LiDARs have become a very prominent 3D-LiDAR technology in AD and also mapping activities due to their ability to provide large scan coverage at high frame rate. Today, common available spinning sensors have up to 128 point layers. By using specific optics, different FoV configurations can be achieved, enabling different range capabilities based in the desired application requirements. Fig.2.13 shows common spinning LiDARs available in the market.

(a) Velodyne VLP32 (b) Ouster OS1-128 (c) Hesai Pandar 64 (d) Robosense RS 32

Figure 2.13: State of the art spinning LiDARs.

Physical considerations of the LiDAR architecture and measurement process (the measurement is performed in the polar space and the points are given in a cartesian coordinate system from the sensor origin) create offsets in the resulting 3D coordinates of the points. The intrinsic calibration of spinning LiDARs concerns each layer (or emitter-detector pair) because the measurements of one layer are performed independently to the other layers. Although some changes in design exist within spinning LiDARs, common intrinsic parameters allow to quantify errors and calibrate the measurements of each layer. For each data point, geometrical corrections can be applied to determine the 3D-location. Common corrections are : vertical correction angle, rotational correction angle, distance correction offset, horizontal parallax correction and vertical parallax correction. To improve the precision of these LiDAR systems and thus the algorithms using point clouds, calibration techniques can be used to measure such parameters [36, 37]. Companies operating the LiDARs on vehicles can develop their own calibration algorithms but it is even more valuable when the supplier itself provides the required information. Some suppliers do provide the calibration parameters. However, those parameters are often not specific to a unit but common to a batch. The intrinsic calibration is a unique characteristic to each single LiDAR unit. Unfortunately, this is usually not the case and the sensors are often considered as black boxes.

Known limitations of spinning 3D-LiDAR sensors first lie in Mean-Time-Between-Failures (MTBF) time. Indeed, mechanical rotating parts have higher chances of failures in time compared to static architectures which makes spinning LiDARs less interesting in the long term. Companies have therefore to invest in the robustification of the spinning and mechanical technology and improve the MTBF.

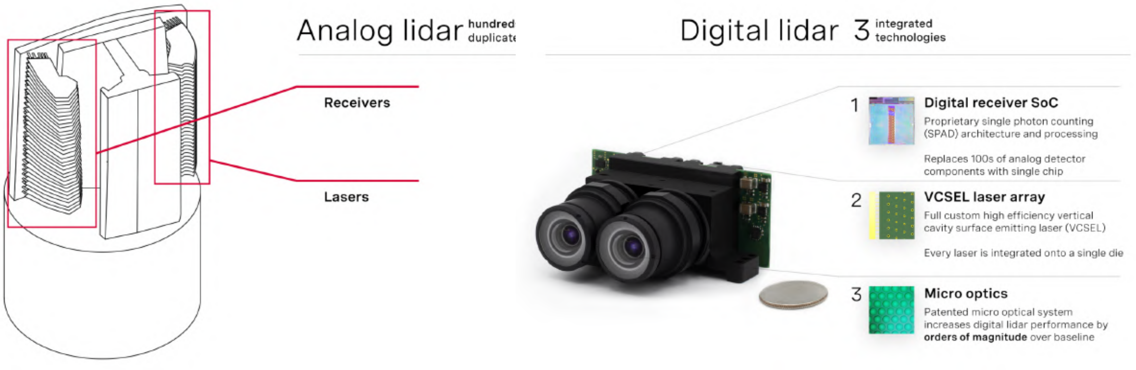

Figure 2.14: Comparison between analog and digital LiDAR architectures, excerpt from [38].

Another limitation is the size caused by the vertical stacking of emitter-detector pairs. The classical architecture of spinning LiDAR systems, which uses laser diodes and Avalanche Photo Diode (APD), is known to be bulky. The accumulation of pairs of components, along with increased power consumption, contribute to a significant extension in size, see fig.2.14. The adoption of smaller components, such as the use of VCSEL instead of traditional laser diodes for emission, and on-chip designs are required to improve the size efficiency.

2.2.2 Flash LiDAR

In a flash LiDAR, the laser beam is spread into the scene using a divergent optic (or diffuser on fig.2.16) so that the emitted light covers the totality of the FoV [39]. Fig.2.15 shows commercially available flash LiDAR units, designed for automotive. Products from Xenomatics, like the XenoLiDAR of fig.2.15c are specifically designed to produce accurate measurements of the road.

(a) Leddartech Pixell (b) Continental HFL (c) Xenomatics XenoLiDAR

Figure 2.15: State of the art flash LiDARs.

Using such design, the detection architecture is a matrix of pixels which resolution defines the point cloud resolution. Each detector of the matrix is synchronized with the emitted pulse time and can perform the ToF measurement from the received signal.

Flash LiDAR technology does face a limitation concerning eye-safety regulations. The laser pulse being concentrated in a divergent beam, its energy distribution spreads with distance. As a result, the compromise of pulse amplitude, maximum range and FoV is even more restrictive than with point-wise scanning.

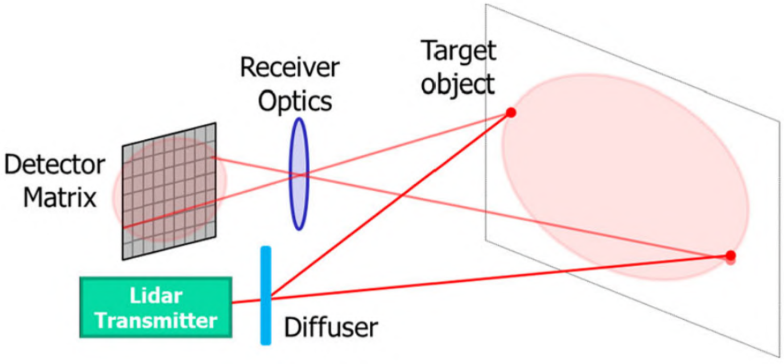

Figure 2.16: Schematic representation of divergent beam emitted in a flash LiDAR architecture, excerpt from [33].

Another technology, known as "laser-gated-camera", works with the same principle as flash LiDAR but the detection matrix is similar to that of a regular camera detector without measuring the ToF [40]. To achieve this property, the detector of the laser-gated-camera is monitored to be electronically open for certain slices of time and therefore space. The detector then integrates all the returned signal received during each slice of time and creates a 2D image. As a result, the contrast within each slice is increased and it seems to improve the performances in scattering media [41]. The technology is not very present on the market but BrightWayVision is the most advanced company to propose such technology for autonomous driving [42].

2.2.3 Mono-static or bi-static architecture

LiDAR architectures also differ by using a mono-static or bi-static architecture. Fig.2.17 shows the overlapping function $O(s)$ for both mono-static and bi-static architectures in a LiDAR system, which represents the proportion of surface shared by the emission and detection FoVs. One can see that the main difference is the blind zone at short range in a bi-static architecture, where $O(s) < 1$.

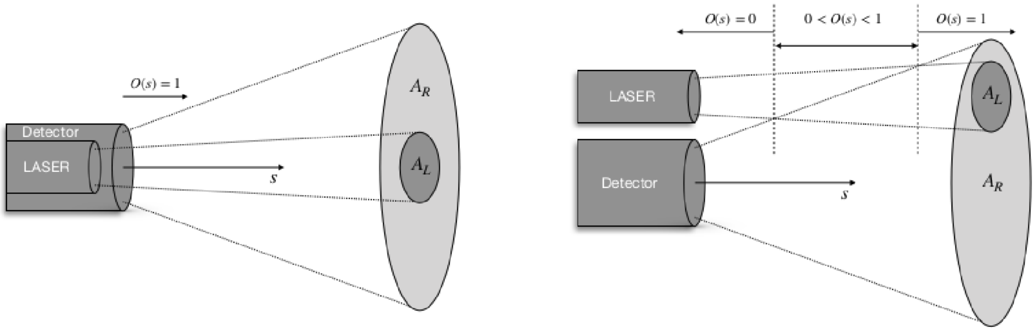

Figure 2.17: Mono-static (left) and bi-static (right) architecture in a LiDAR system, excerpt from [43].

Although a mono-static design is more efficient in terms of weight and size, a bi-static architecture, ideal for long-range applications, presents several advantages. These advantages include being less affected by direct backscattering of light from aerosols because the direct backscattered light within the detector FoV is avoided [26].

2.2.4 Hybrid approaches

Some 3D-LiDAR architectures do not fit in predefined categories due to their unique design characteristics and approaches. Cepton LiDARs use a proprietary micro-motion technology for both emission and detection [44] Their common design is to use a bi-static architecture with arrays of emitters on one side and detectors on the other side. Then, as the whole system vibrates, the laser beams are steered across the FoV. AEye company proposes LiDARs using the combination of a MEMS mirror and a matrix of pixels for detection [45]. This technology is potentially similar to the LiDAR presented in [33] where a 1D MEMS mirror steers a vertical line of laser light in the horizontal axis. Screen shots of point clouds captured by a Cepton and AEye LiDAR are given in chapter 6.

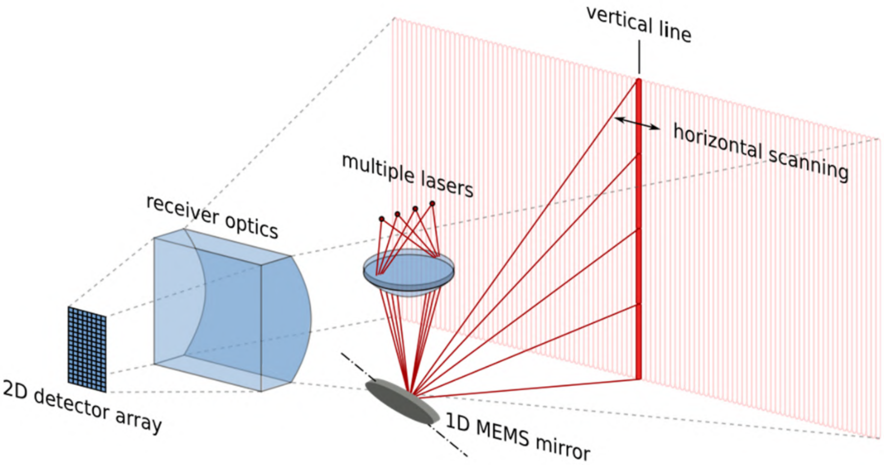

Figure 2.18: Schematic representation of LiDAR using a 1D MEMS mirror and a matrix of pixels, excerpt from [33].

The Australian company Baraja uses refractive prisms to steer laser beams and perform line scanning [46]. However, instead of rotating the prism, they use the fact the refraction angle of light inside a prism depends of the wavelength. Using a laser beam which wavelength slightly shifts, each beam has a different wavelength and is refracted to a different angle. This refraction principle is used for one axis, while the other axis can be scanned using a classic method like a mirror.

All these hybrid approaches show that a large variety of technologies are being developed to cover the 3D scanning needs of numerous and new applications which have specific requirements.

2.3 Detector technology

Depending on the imaging technology, the detection scheme can be mono or multi pixel. The most common pixels found in LiDARs are Avalanche Photo Diode (APD) and Single Photon Avalanche Diode (SPAD). An APD is a semiconductor device which converts light into electrical current through the photoelectric effect. Similarly to semiconductor laser sources, the wavelength detectable by a semiconductor detector depends on the nature of the material. For instance, Silicon detectors are sensible to visible and near infra-red light (usually for LiDARs working around 900 nm) while InGaAs are typically used for Short-Wave-Infra-Red (SWIR) (around 1.5 μm), as shown on fig.2.19).

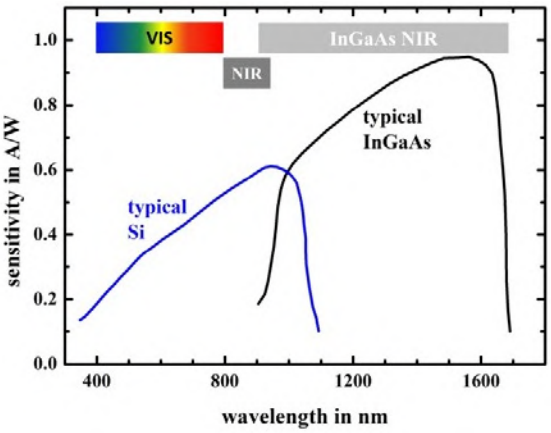

Figure 2.19: Spectral sensitivity of Si and InGaAs detectors, excerpt from [47]

A received backscattered signal measured by an APD is in the form of a Full WaveForm (FWF) signal, fig.2.20. It represents the quantity of light energy received in a time interval after the emission of the laser pulse. From this FWF signal, peak detection algorithms convert the signals in detections and thus points in the point clouds. Other algorithms can also be applied to these signals because they potentially contain information about the targets properties or environmental conditions [48].

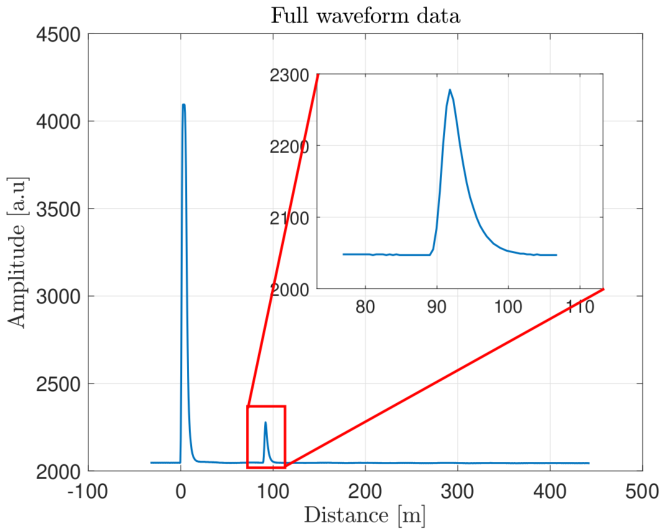

Figure 2.20: FWF signal return from an emitted laser pulse and APD detector, excerpt from [43]. The first peak corresponds to the reflection of the emitted laser on the optics and the second peak is the reflection of a target located at 90 m.

The difference between APD and SPAD lies in the operating voltage. If the detector is biased below a breakdown voltage value, its response to photons is linear according to the amount of photons received, see fig.2.21. When the detector is biased above its breakdown voltage value (and with appropriate structure), it works in Single-Photon (or Geiger) mode. In Single-Photon mode, each received photon triggers a detection and acts as a binary photon detector instead of producing a FWF signal. The detector is also temporally synchronized with the emitted laser so that the binary detection is associated with a ToF and thus a measured distance. Once a photon is detected, the cell is deactivated and reset when the next laser pulse is shot.

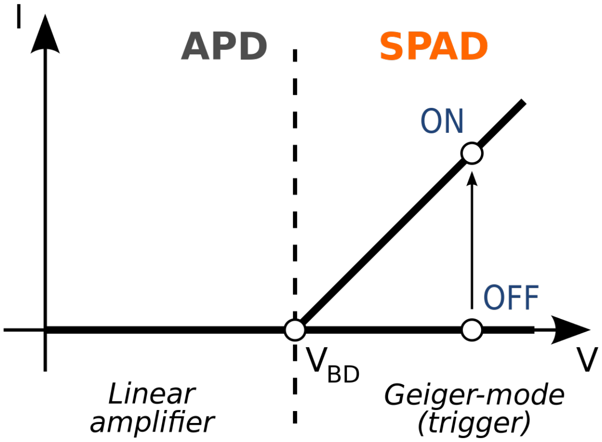

Figure 2.21: Linear and Geiger-mode operating regions of a SPAD cell according to the applied voltage.

This method is incompatible with penetration capabilities (or multi-echo, explained in sec 2.7) as only first echos can be detected. To leverage the advantages of this technology, one should go beyond single binary detection and instead consider multiple acquisitions to gather multiple binary photon detections along with their corresponding ToF information. This is possible by extending the number of pixels associated to each part of the FoV or by repeating the acquisitions (pulse emission and Single-Photon triggering) at high frequency. Compressive sensing is also a promising technology to increase the resolution of 3D-LiDARs using Single-Photon detectors [49]. The result of using Single-Photon detectors can be displayed as a photon count histogram which represents the quantity of triggered detections over some time (fig.2.22).

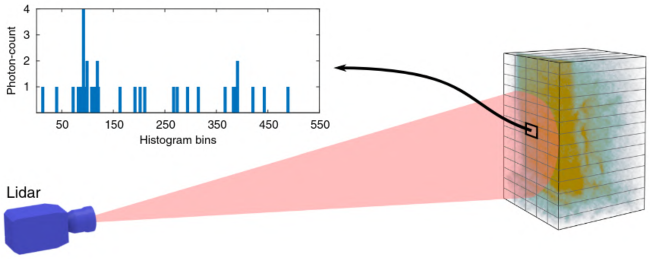

Figure 2.22: Single-Photon histogram return from an emitted laser pulse and SPAD detector, excerpt from [50].

Single Photon or Geiger mode systems are then highly sensitive because they are capable of detecting photons individually. But, high sensitivity also means higher level of noise. Even under nominal conditions, photons can experience backscattering or reflection from aerosol particles. Initial point clouds created in this way often contain a significant amount of false positive points, which are erroneous data points that do not correspond to actual objects or surfaces in the environment. It is necessary to apply pre-processing filtering techniques to obtain a clean and reliable point cloud (fig.2.23). For more details about the Single-Photon technology for LiDARs, see [51].

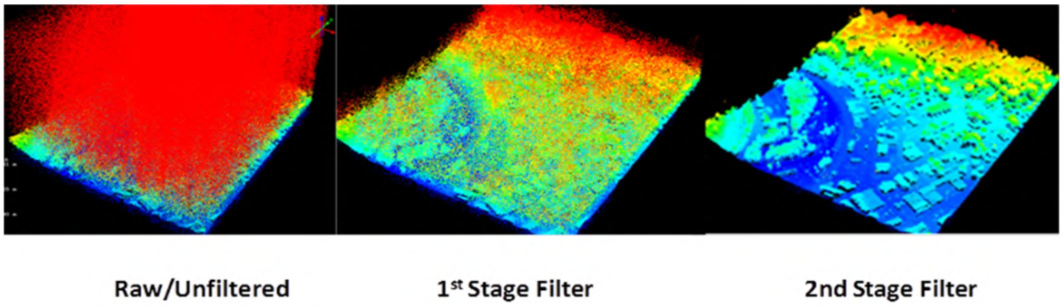

Figure 2.23: Comparison between raw and filtered data from a Single-Photon LiDAR airborne acquisition, excerpt from [52].

2.4 Propagation in the atmosphere

Once a laser beam is emitted by a LiDAR system, it propagates through the atmosphere. Photons interact with the atmospheric particles or aerosols and influence the overall behavior of the beam. This refers to the theory of light scattering which shall be addressed later in chapter 3 as it is the heart of this thesis. In this section, we'll confine ourselves to examining light propagation in a "clear" atmosphere (with no disturbing particles). In this case, two main effects are involved: absorption and turbulence.

The absorption effect is represented on fig.2.24, showing the transmittance of the atmosphere depending on the wavelength, indicating how much the atmosphere does not absorb specific wavelengths. Visible light is minimally absorbed by the atmosphere, which aligns with its perceptible nature as these wavelengths correspond to what we perceive with our eyes. Near infrared signals are also fairly well transmitted through the atmosphere, rendering them suitable for utilization in LiDAR systems.

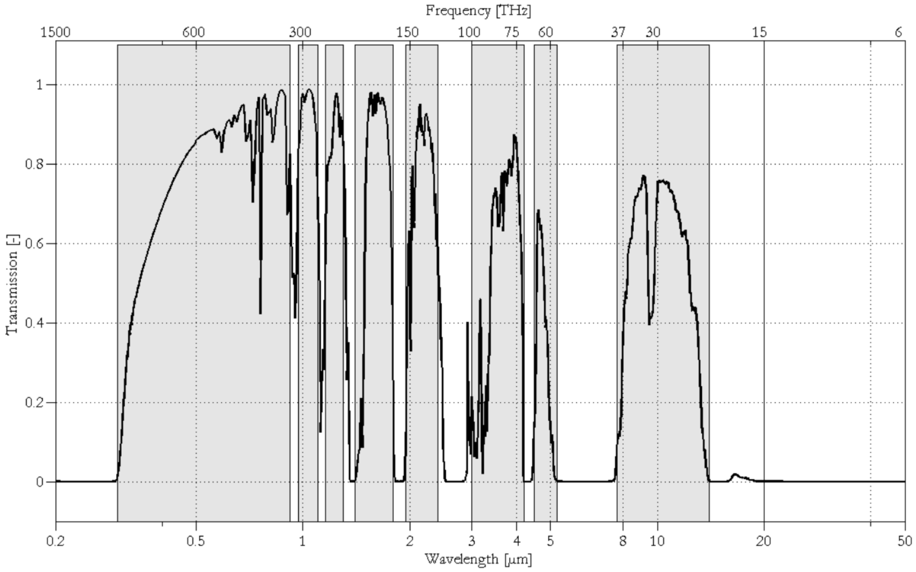

Figure 2.24: Atmospheric transmittance for wavelengths from 0.2 μm to 50 μm, excerpt from [53].

Turbulence or index-of-refraction turbulence (IRT) is responsible for wavefront distortion over distance. When a light beam travels through the atmosphere, the random changes in temperature of the air cause random variations of the refractive index which redistribute the intensity of the beam in a more random pattern. For example and as shown in fig.2.25, a Gaussian shaped intensity profile is changed into a random interference pattern after passing through turbulent atmosphere [53].

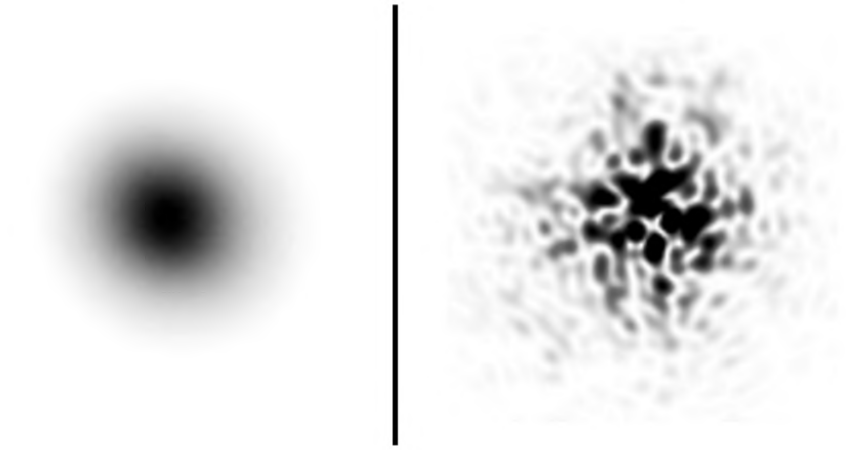

Figure 2.25: Comparison of intensity pattern of an initial light beam (left) and the beam after propagation through turbulent atmosphere (right), excerpt from [53].

Both of these effects have to be taken into account when designing LiDAR systems (or during post-processing) because they influence the signals involved. However, the impact of turbulence is rarely considered for LiDARs designed for autonomous driving systems, with relatively short distance applications (a few hundred meters). Though, it is possible that environments with high temperatures (potentially urban scenarios) would influence 3D-LiDAR measurements from these effects, even at those short distances (mirage effects).

2.5 Interaction with objects

The Bidirectional Reflectance Distribution Function (BRDF) can be used to define how light is reflected when it hits the surface of an object [54, 55]. The BRDF is a fundamental radiometric concept used in various fields including optics design, computer graphics, computer vision or photovoltaic. In LiDAR systems, the BRDF is crucial to assess the behavior of laser beams on objects. It can be measured using specialized optical systems for specific objects and analytical models have been developed to approximate the BRDF of those objects. Specific to objects, the BRDF can be measured using optical systems to approach it. For example, the Lambertian model assumes that the reflected light is perfectly diffuse (fig.2.26a).

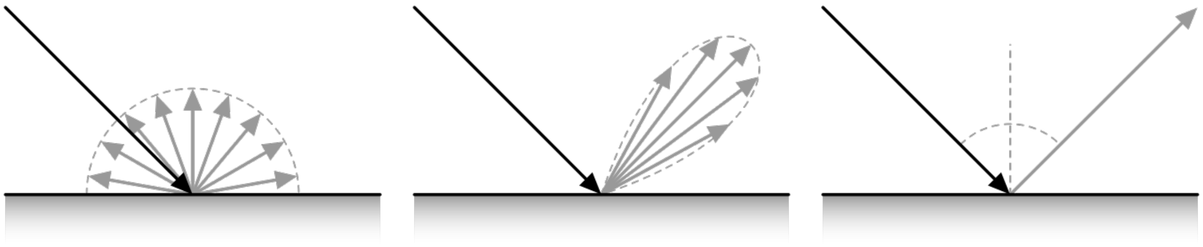

(a) Diffuse (pure lambertian)

(b) Glossy

(c) Mirror (pure specular)

Figure 2.26: Examples of light-surface interactions.

The shape of the objects also influences the resulting signals returning to the detector. Fig.2.27 shows different forms of FWF returns that can be expected with different targets nature and forms. The effect is supposed to be applicable to linear detection (FWF signal from APD cell) as well as Single-Photon (photon counts histograms from SPAD cell).

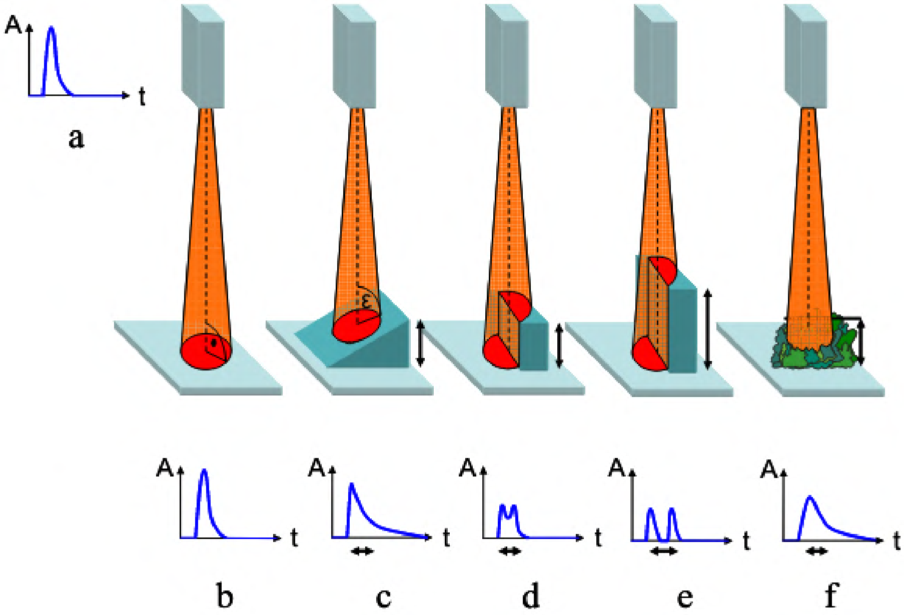

Figure 2.27: Different FWF returns according to objects nature and shape, excerpt from [56].

2.6 Measurement principles

Laser signals can be emitted in continuous or pulsed modes. Depending on the mode, the LiDAR measurement principle differs: pulsed laser signals are emitted in LiDARs working with direct ToF measurements while Continuous-Wave (CW) emission is used in Frequency-Modulated-Continuous-Wave (FMCW) and Amplitude-Modulated-Continuous-Wave (AMCW) technologies, which lead to indirect ToF techniques.

2.6.1 Indirect ToF measurements

Indirect ToF techniques require continuous wave emission. Instead of emitting very short laser signals in the form of pulses, frequency or amplitude modulations are applied to continuous laser emissions. Information about the object are then retrieved by comparing the properties of the emitted and backscattered waves.

Fig.2.28 shows the schematic principle of indirect ToF using the FMCW technique, where the modulation is applied on the frequency of the signal. This method allows to measure the range of objects as well as their velocity. A continuous wave is sent toward the target and the backscattered signal is received after a certain ToF. This backscattered signal is then mixed with a reference. The temporal delay between the received and reference lights causes a constant frequency difference, commonly referred to as the "beat frequency" (f_b on fig.2.28c), when the two signals are mixed together. When the target is stationary, the beat frequency resulting from the mixing of the received backscattered signal and the reference signal remains constant over time. This beat frequency is directly proportional to the range or distance between the LiDAR sensor and the stationary object. However, when the target is in motion, the beat frequency is not only influenced by the range but also by the velocity of the object relative to the LiDAR sensor. This phenomenon is known as the Doppler effect. In the field of autonomous driving, there has been a recent emergence of FMCW LiDAR systems, and among the companies pioneering this technology, Aeva stands out as one of the most advanced [57]. The simultaneous measurement of both range and velocity in LiDAR systems holds great promise for the advancement of autonomous navigation. Traditional perception algorithms often rely on tracking and Bayesian filtering methods to estimate the velocity of surrounding objects. However, obtaining accurate and reliable velocity information can be challenging and prone to errors. By directly measuring the velocity of objects using FMCW LiDAR systems, the results of perception algorithms can be significantly improved.

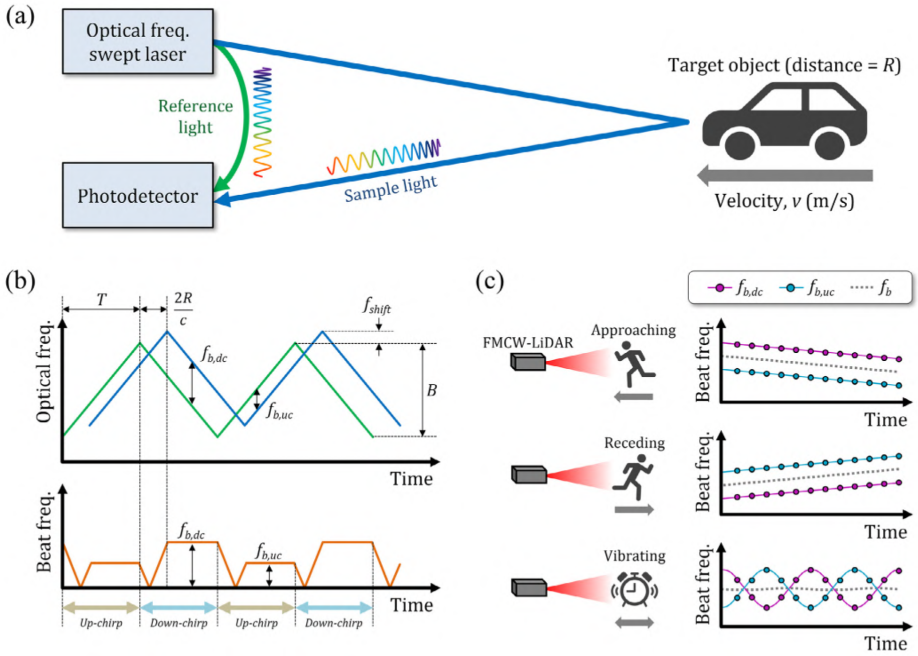

Figure 2.28: FMCW LiDAR range measurement principle. The laser light is emitted continuously and shifts of frequency are applied to perform coherent detection and Doppler velocity measurement in addition to range, excerpt from [58].

The FMCW technology does, however, have its limitations [25]. Concerning the velocity measurements, it is only able to measure the radial velocity and not the lateral one, which considerably reduces its advantage in many scenarios. The other major limitation is time. Since it necessitates the analysis of the waves, more time is needed to perform the measurements, compared to direct ToF techniques. Finally, although promising, this technology has yet to prove itself to be considered major in the realm of LiDAR imaging.

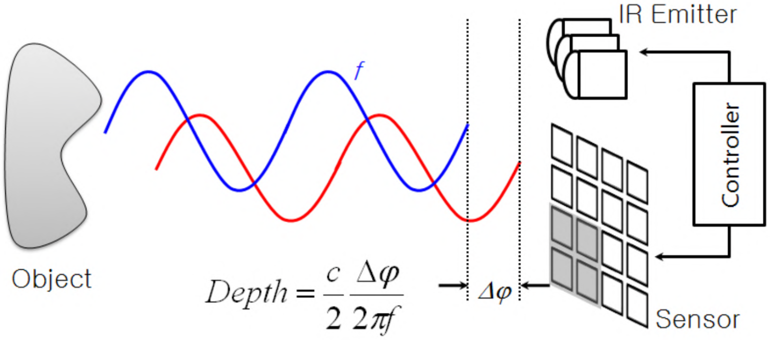

Figure 2.29: Schematic representation of the AMCW range measurement principle, excerpt from [59].

AMCW technology also works with continuous and modulated wave emissions, but instead of frequency modulation, those systems employ amplitude modulation to achieve range measurements by comparing the phases of the emitted and backscattered waves (fig.2.29). This technology obtains good results in terms of range measurement but is limited to moderate ranges due to an ambiguity caused by the frequency of modulation. Often referred as ToF cameras, these sensors are primarily used in indoor environments and have limited applicability in outdoor autonomous navigation scenarios [25]. Commercialized since the 90s, AMCW cameras have been operated in various fields such as video-gaming (Kinect by Microsoft), human machine interaction, robot navigation and object recognition [59].

2.6.2 Direct ToF

Direct ToF, illustrated on fig.2.30 is the most common technology in LiDAR systems. In pulsed laser mode, very short pulses of light are emitted and the ToF principle can be applied to each individual pulse. The overall backscattered signal (FWF or photon count histogram) is received as shown in section 2.3 and detection algorithms are used to convert high energy returns into point detections [60, 61]. The time delay between laser emission and the reception of these peaks is the ToF and thus the range of the objects.

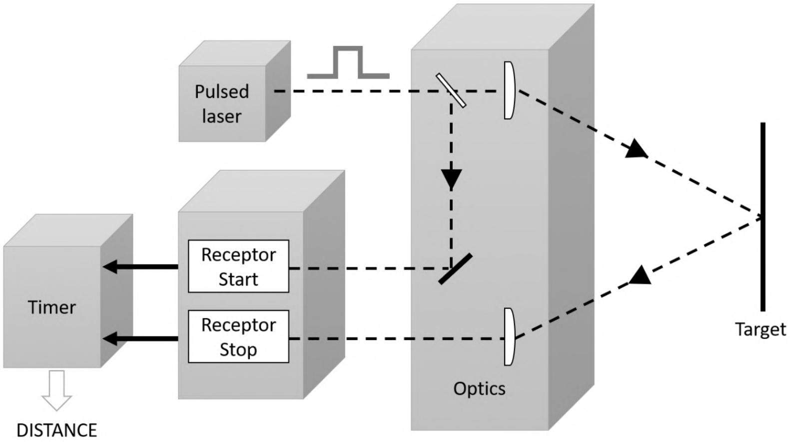

Figure 2.30: Schematic representation of a direct ToF processing chain, excerpt from [62].

Using this technique, the range of objects is measured but additional information concerning the returned pulse shapes can be retrieved depending on the sensor technology. These additional information are valuable as they indicate properties about the objects in the laser's line-of-sight. The data fields provided by the sensor with along with the point clouds can vary depending on the specific sensor and its capabilities. For instance, most sensors only give intensity information but, Riegl LiDARs give the pulse deviation as well, the Valeo Scala gives the pulse width and Ouster sensors give calibrated and non calibrated intensities as well as passive background illumination (quantity of infrared light received without laser emission). Intensity is the first and most common field given along with the points range, it represents the quantity of energy contained in the returned signal. The amplitude and width of the received signal at the peak position provide the information needed to give intensity information to the points. To have the most accurate and reliable intensity data, calibration must be applied to compensate the reduction of amplitude signal caused by distance (inverse square law). Also, to improve the accuracy and reliability of the measurements, characterization must be done on the received signals using targets of various nature (distance, reflectivity, forms) as shown fig.2.27.

2.7 Multi-echo

In direct ToF configuration, a detector is synchronized with the pulse laser emission and captures signal during a defined period of time. During this time, it receives signal from all possible contributions, especially the ones from the backscattering of the emitted laser beam. The resulting FWF signal can vary depending on the scene and spatial properties of objects within the laser beam, fig.2.27. The divergence of the beam increases the laser beam footprint over distance, fig.2.31.

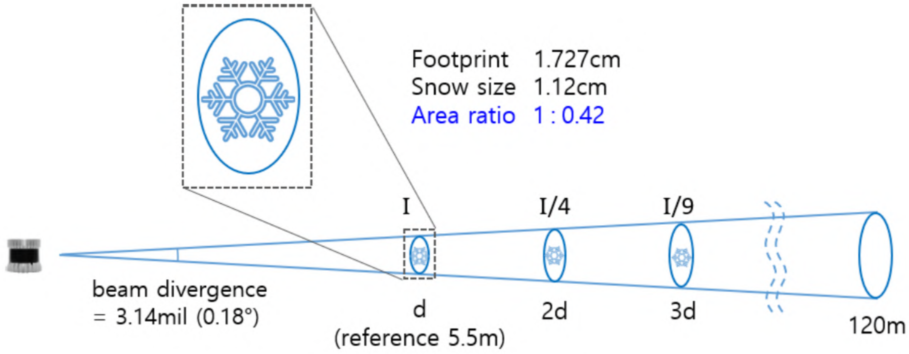

Figure 2.31: Example of laser beam foot-print on a snow flake at various distances [63].

Thus, the energy repartition of a laser beam in the line of sight is contained in this increasing foot print and the potential backscatter energy from an object follows an inverse-square law [64]. And, the beam can impact multiple objects in its path. Each impact backscatters to the detector and contributes to the overall FWF signal. Each peak above the detection threshold can become a detection so that the initial laser beam can produce multiple echoes at different distances but detected within the same line-of-sight. This effect is referred as multi-echo and is particularly noticeable when LiDAR sensors are used in airborne applications, as shown on fig.2.32. Vegetation is often composed of complex structures including multiple small objects (branches, grass) that can each interact with the laser beam.

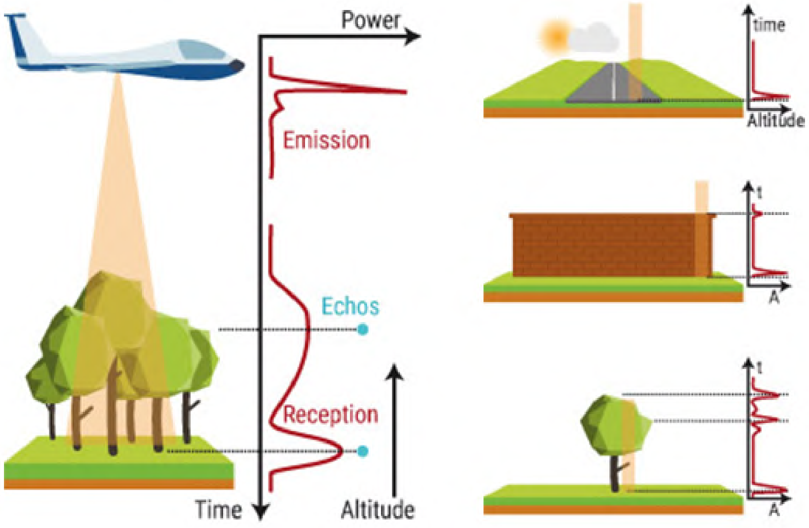

Figure 2.32: Schematic FWF LiDAR returns in airborne scenarios. Different shapes of objects cause different returned signals and potential multiple detected echoes from a single emitted laser beam, excerpt from [65].

In the presence of atmospheric particles or aerosols, it is also very likely that light is backscattered from these particles to the detector. The resulting FWF signal is then influenced with increased noise and potentially strong returns in some cases. These strong returns can be considered as true detections and become part of the constructed point clouds.

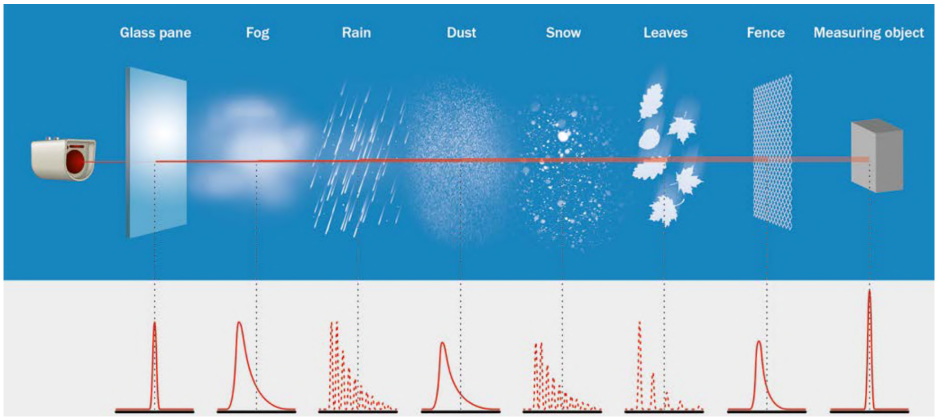

Figure 2.33: Schematic FWF LiDAR return in the presence of disturbing objects including water particles of degraded weather conditions, excerpt from [66].

All LiDARs do not have multi-echo capacity. It requires specific development (and electronic) and care should be taken because increasing the number of echoes detectable from a single laser beam can bring pros and cons. First, a higher amount of echoes naively means more information about the spatial structure of the scene, which is always valuable when performing 3D scanning. On the other side, more echoes also might create higher levels of noise in the point clouds.

LiDARs with this feature offer multiple ways to label echoes. Echoes can be labeled based on their chronological order of appearance (first, second, third...). Alternatively, an echo can be tagged with its amplitude rank. Depending on the application, a certain configuration can be required. For example, the *Last* echo is needed in mapping applications to construct digital elevation model of the ground under the trees.

2.8 Discussions

This last section gives discussions about topics related to the 3D-LiDAR technology. In particular, requirements needed for the applicability in the AD world influence the design of the sensors. The section also provides information about the point clouds generated by LiDARs. These point clouds exist in multiple forms and must meet the diverse needs of autonomous vehicles (AVs). The design of sensor sets is then closely linked to achieving the FoV coverage with these point clouds. Lastly, the section offers insights into the future of 3D-LiDAR technology, giving a glimpse of what lies ahead for this evolving technology.

2.8.1 Automotive LiDAR

A few critical requirements are required for LiDARs to be utilized for AD. First, as other agents in the environment of the vehicle are moving, a high scan rate is critical to capture the moving objects. Typical frequency of point cloud generation on commercial available LiDARs today is 10 Hz but a higher value could be needed for high speed applications like ADAS and highway-driving. As such, LiDARs designed for the mapping industry, which provide high resolution point clouds, are not suitable for AVs. Then, the eye-safety regulation is very present because the vehicles are likely to evolve in scenarios including pedestrians. Therefore, the performances of the sensors is often limited by this aspect. Finally, automotive grade requirements are mandatory for sensors be integrated on cars (ADAS). Overall, a safety approach, much needed to deploy AVs, is also very valuable if applied to the sensors themselves. In addition to the eternal need for better quality in general, these requirements have shaped the evolution of the 3D-LiDAR technology designed for AD.

The first prominent LiDAR sensors for autonomous driving or robotics came from the Velodyne company (with the HDL-64E in 2008 and the VLPs in 2015). Their spinning sensors have led the market for several years but they have not been able to significantly evolve and have recently merged with Ouster. Today, other companies have emerged and are competitive in the LiDAR market: Ouster, Hesai and Robosense are well established companies which propose different models of LiDARs but especially spinning sensors. Ouster only proposes spinning LiDARs today (although their recent acquisition of Sense Photonics will lead to the development of flash LiDARs) but differentiates itself with the combination of VCSEL emitters and SPAD detectors. Hesai and Robosense Chinese companies provide a large variety of LiDAR technologies (spinning, mechanical, maybe MEMS LiDARs) but not a lot of information is available concerning the core technologies inside the sensors. The limitations of the spinning technology (inter-layer effect, MTBF) have led to the development of other architectures, aiming to fill the gaps or to address other use cases.

Luminar company is one of the major actor of mechanical 3D-LiDAR for autonomous driving. Their 1550 nm fiber laser with on-chip designs and especially maturity in terms of integration and safety has convinced many automotive Original Equipment Manufacturers (OEMs). Aeva, is also becoming a major mechanical 3D-LiDAR provider. Their products use the FMCW technology (presented in sec.2.6.1) which is supposed to give valuable velocity information (at the cost of some complications). These mechanical 3D-LiDAR scanners are more designed to suit ADAS applications than robotics because their rectangular FoV focus on frontal detection, e.g. for highway driving.

Risley prisms beam steering technology, leading to non-repeating scan patterns, is attractive to cover the entire FoV but it is difficult to develop algorithms based on these point clouds because of the irregular point density. Flash LiDARs have known their peak of popularity because of well made marketing videos but are now less considered as game changers, rightly so. The technology is promising because it could provide complete coverage of large FoV with no moving parts (true solid state). But, the compromise between the laser emission amplitude and eye-safety regulation is more restrictive than for point-wise scanning approaches and slows the development of the technology. Today, Leddartech company has stopped the development of their Pixell LiDAR. Continental proposes the HFL (for High-resolution Flash LiDAR) sensor but it is not at a very mature state although some recent works have been published using this sensor [67]. The flash technology has more future in applications where the eye-safety is less constraining like defense and space [68, 69].

Finally, embedded signal processing techniques (peak detection, intensity reading) and technical specificities are highly confidential for LiDAR suppliers. As a result, the point-fields information, which is not always reliable, are sometimes difficult to understand and more information would be needed from the constructor. Deeper discussions with them could help the development of algorithms for us (the users) and increase the chances of choosing their sensors but most of them are not open to give these information so the sensors remain black-boxes. Though, LiDAR companies are often ready to loan their sensors for free, at least to a company like EasyMile. In exchange for the generated data and some analysis reports, sensors can be acquired for a period of time to conduct research activities, which is very convenient to benchmark the different technologies.

2.8.2 Point clouds

The design of a LiDAR system originates from requirements inherent of the application and the characteristics of the needed point clouds. In addition to the requirements stated in the previous section, the properties of the produced point clouds are also critical in the applicability of the LiDAR.

Multiple variables are of interest on the point clouds themselves such as density, precision, accuracy, coverage, possibly RGB color and/or penetration capability and additional point-fields information. In addition, characteristics of the acquisition process is also important : pulse repetition rate, scan rate, range and real-time processing capabilities. Dense point clouds captured at low frequency can be needed to provide very precise measurements of a scene while sparser measurements at high scan rate are more useful to capture objects in dynamic scenarios.

Spinning LiDARs have good coverage and frame rate but their architecture and thus their point clouds suffer from inter-layer effect: from a fixed position, the space between two successive pairs of laser-detector is never scanned. Adding vertical resolution solves this problem but it means adding more laser and detectors (see fig.2.34).

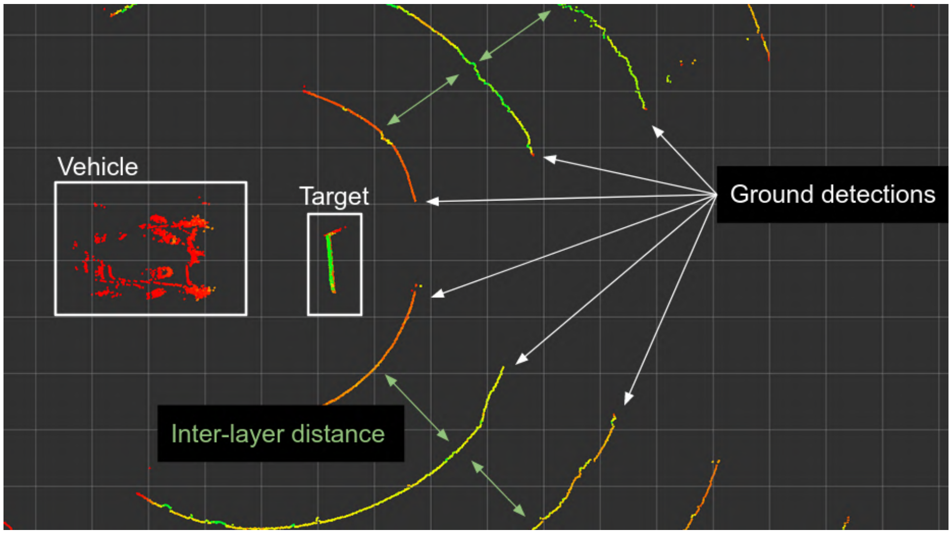

(a) Velodyne VLP32 : 32 vertical layers with a 40° FoV.

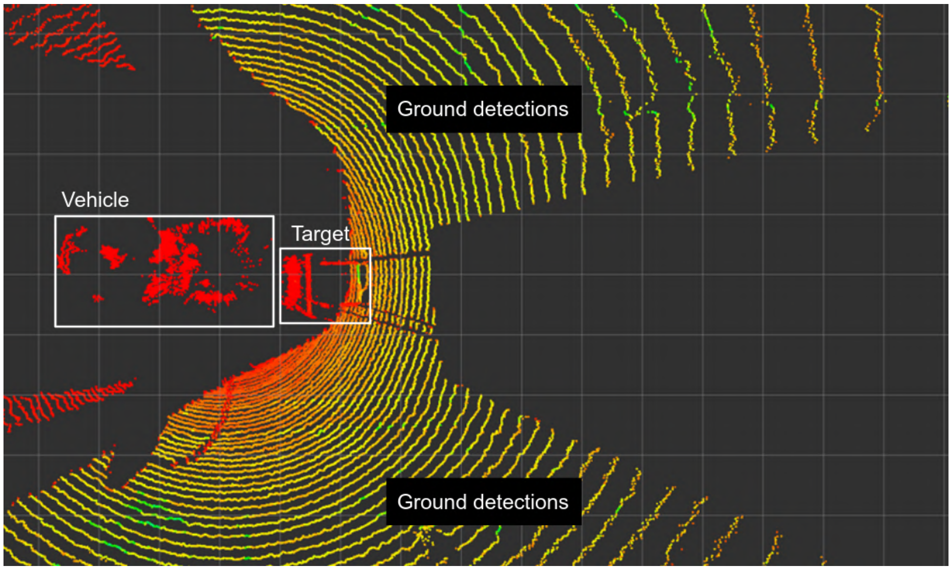

(b) Ouster OS0-128 : 128 vertical layers with a 90° FoV.

Figure 2.34: Top-down views of point clouds produced by two spinning LiDARs with different vertical resolution

Matrix of pixels (potentially present in flash and MEMS LiDARs) provide image-like format of point clouds to theoretically involve less blind spots with large FoV but rounded shapes from straight objects have been witnessed in some point clouds. Single layer scanners do not provide dense point clouds and thus are not used to give much density of information but are rather needed for safety purpose in industrial or autonomy applications. Point clouds obtained from non-repeating scan patterns are useful for their penetration and coverage capacity, to for example reach complex structures (vegetation, clouds). Pictures of point clouds of various LiDAR technologies are given in chapter 6.

Finally, the progress of LiDAR technology (smaller size of laser sources and detectors, better electronics, sensitivity of the detectors and signal processing algorithms) allows the production of denser and more precise point clouds. These advancements are good for the applications but the problem of computational power comes at some point. In static or mapping applications, the issue is less constraining because most of the time the processing is done post acquisitions. For autonomous driving applications and any application which requires online processing, it implies dealing with great amount of data fast enough to be able to adapt the vehicle behavior accordingly. To handle such amount of data, development is needed with the creation of potential innovative data structures and compression techniques [70, 71].

2.8.3 Sensor sets

To enable autonomous navigation, autonomous vehicles (AVs) heavily rely on a variety of sensors that form a sensor set. These sensors are strategically placed at specific locations on the vehicle to ensure comprehensive coverage of all the necessary aspects for the application. Constructing an appropriate sensor set requires a deep understanding of the sensors themselves, including their intrinsic characteristics and the point clouds they produce. In autonomous driving (AD), various functions must be taken into account, such as short-range perception for obstacle detection, long-range perception for localization, obstacle avoidance, and potential target tracking. To address these different scenarios, a combination of point clouds from various regions around the vehicle is required. Ultimately, a multitude of different sensors is used because no single sensor can fulfill all the requirements. For example, it is challenging to simultaneously achieve short-range and long-range detection optically. Fig.2.35 illustrates the sensor set from a recent Waymo AV. It shows that a combination of different sensors is used (LiDARs, RADARs and cameras), located at different positions.

Today, it seems that no ideal solution has been found regarding the design of sensor sets, as no company has really demonstrated any supremacy with a specific architecture. Some manufacturers believe that LiDARs are too expensive and that cameras can be sufficient. Others rely solely on LiDARs because the fusion of multiple modalities is probably too complex to develop. But the combination of several modalities is promising, as these different sensors have different advantages and disadvantages, which should complement each other.

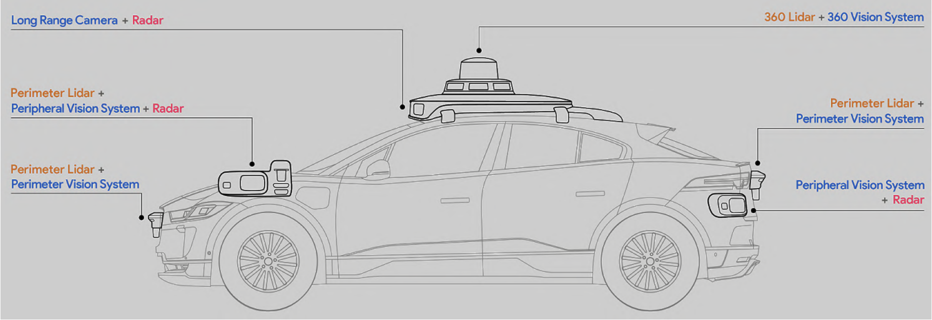

Figure 2.35: Illustration of the sensor set of a Waymo autonomous car.

2.8.4 The future of LiDARs

Just like for all technological systems nowadays, the global performances and cost of 3D-LiDARs are inevitably destined to improve with time. Considered today as relatively expensive, the LiDAR technology should be more available in the future. Ouster, for example, says its sensors are evolving according to Moore's Law because of their components, but more generally, ongoing vendor development will strengthen and improve existing technologies and costs will go down as more sensors come to the market and are mass produced. In addition to the cost, the inner technology is constantly improving. This section discusses on the promising technologies for the LiDAR technology and their future use.

For example, Innoviz company, known for its MEMS LiDAR products has recently announced their Innoviz360 LiDAR sensor which potentially uses a combination of spinning and MEMS technology to perform a 1280 lines scanning. Increasing the resolution of the LiDARs is major concern to increase the performances of AVs which rely on them. As a result, one should expect a continuous increase in the resolution of these sensors in the future.

Some LiDAR manufacturers (mostly classical mechanical scanning LiDARs) propose the capacity to obtain FWF signals in addition to the point clouds but it is not common at all in the world of autonomous driving [72]. The Leddartech company has published a dataset of FWF data from their flash LiDAR along with other sensors data from an equipped vehicle [73] in order to increase the development of FWF algorithms. Processing FWF signals requires higher computation and storage capacity, which is already high when you have to deal with multiple 3D LiDARs on a vehicle. But it is, as an extension of point cloud processing, a way to improve scene understanding because the signals contain more information than the point clouds [74].

The phased array technology is commonly used in radar sensors to electronically steer radio waves in specific directions. When using light beams, the technology is called Optical Phased Array. LiDAR sensors could benefit from this to steer the laser beams without any moving parts (true solid state) and also to scan specific areas instantaneously (active perception) [75]. Some LiDAR manufacturers already propose to increase the resolution of the point clouds in specific parts of the FoV, for example with more laser beams or slower scanning to increase the scanning resolution of road regions. Active perception in LiDAR sensors is proposed by the company Lumotive [76]. Their sensors use Light Control Metasurface (LCM) chips to control the angle of reflection of light and steer laser beams.

Many LiDAR manufacturers are also working on perception software suites along with their LiDAR products. For example, the Aeva company which uses the FMCW for its LiDARs can provide a super resolution algorithm which increases the resolution of the point cloud using the velocity information of the objects. Basically, a point with null velocity is kept during several consecutive scans so the point clouds have additional density in static areas if the sensor is moving [77]. Ouster proposes software for crowd monitoring and retail analytics. Luminar has obstacle and line detection algorithms. These software suites offered by LiDAR vendors are interesting because the developers clearly have privileged information that is not always available to users.

Another research topic concerns the use of a supercontinuum laser source for LiDAR. Indeed, such laser source, produced by non-linear effects inside optic fibers, is composed of a large spectrum of laser-light unlike the very monochromatic ones used so far. With an appropriate system of wavelength splitting, each wavelength of the emitted supercontinuum light could be used uniquely to make the LiDAR measurements. LiDARs working with multiple wavelengths could then be developed and this spectral information could bring valuable information to the point clouds and improve detection and classification possibilities [78].

LiDARs are being used in more and more applications. Recently, Apple decided to include a LiDAR sensor in their iPhone and iPad products. This will accelerate the reduction of size and cost and the LiDAR components as well as the development of algorithms for point cloud data [79, 80].

Finally, some of the limitations of the LiDAR sensors will probably never be solved simply because of the laws of physics. For example, although it is impossible that some degraded scenarios are too difficult to find solutions only with LiDARs, photons will, for example, always interact with fog particles. On the contrary, the wavelengths used by radar sensors will always be less sensitive to these particles. Consequently, the future of LiDARs also lie in its association with other sensors which do not suffer from the same limitations [81, 82].

2.9 Conclusion

3D-LiDAR technology has emerged as a crucial component in the ecosystem of autonomous driving. Its development has been driven by the increasing demand for better perception capabilities and, the diverse range of architectures available today reflects the ongoing efforts to provide accurate and detailed 3D point clouds of the environment.

Autonomous vehicles heavily rely on LiDAR sensors to enhance their performance and ensure safety. By integrating higher quality LiDARs into their sensor sets, vehicles can improve their perception capabilities, including range, field of view, and point density.

For example, the transition from 32 layers spinning LiDARs to higher layer alternatives, such as 128 layers, has addressed previous challenges related to low point density. This improvement enables vehicles to detect and identify objects more effectively, contributing to a safer and more reliable autonomous driving experience.

As the technology continues to advance, it is expected that 3D-LiDARs will play an even more significant role in autonomous driving systems. Further innovations in hardware and software will continue to enhance the accuracy, speed, and reliability of these sensors, enabling autonomous vehicles to navigate complex environments with greater precision and efficiency. The continuous evolution of 3D-LiDAR technology highlights its importance as a fundamental building block for the future of autonomous driving, as well as for other applications.

Chapter 3

Degraded Visual Environment (DVE)

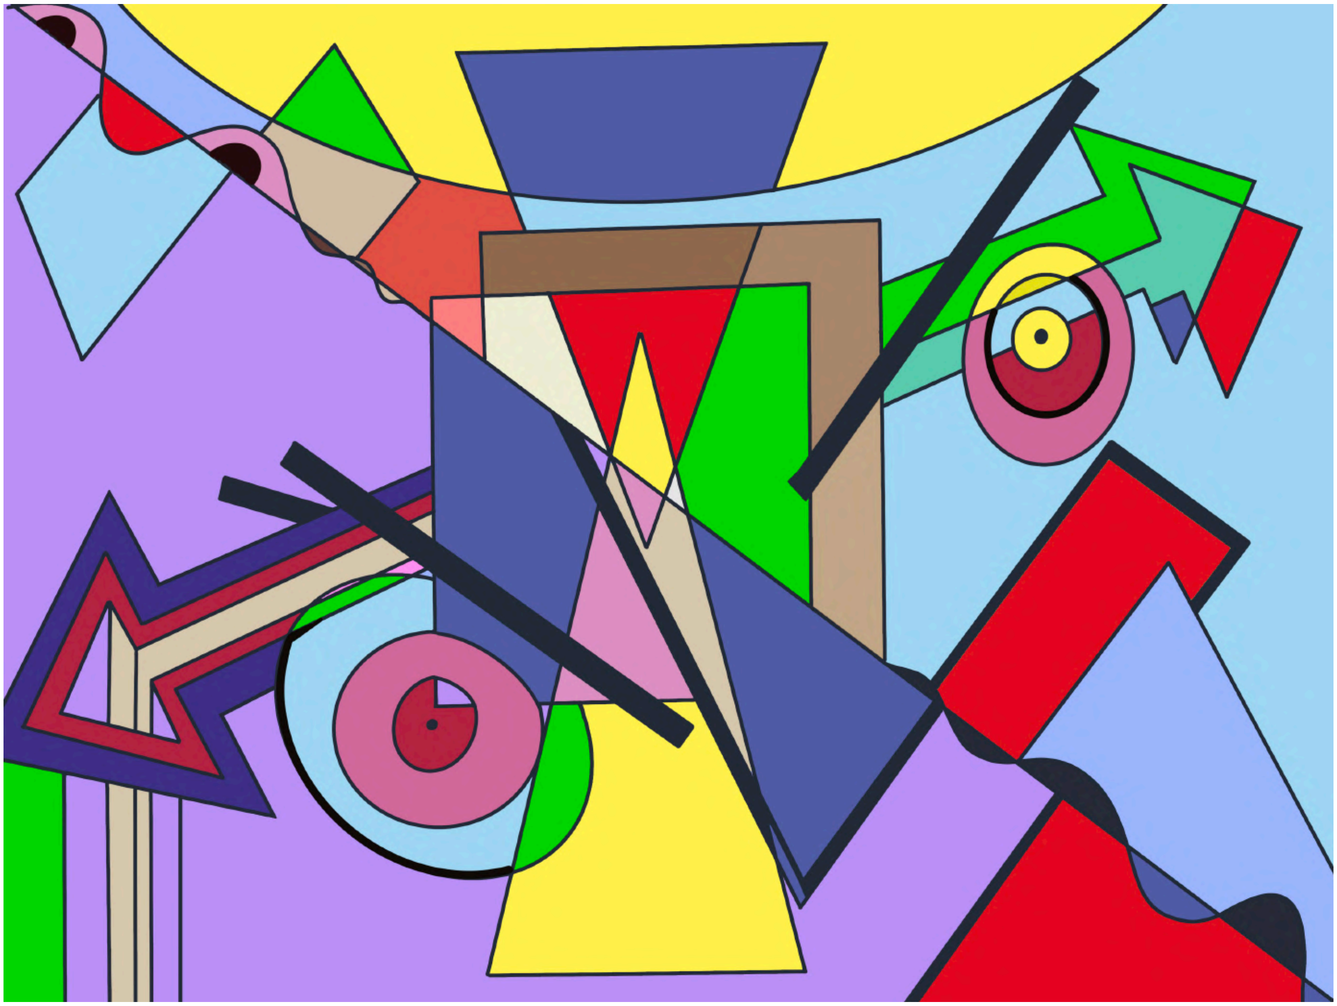

The greatest difficulty with autonomous navigation is the environment in which the vehicle must operate. Indoor applications face weak Global Navigation Satellite System (GNSS) signals and location algorithms must rely solely on perception sensors. On the contrary, outdoor applications offer most of the time the possibility to perform GNSS localization but perception sensors can be impacted by various phenomena. These scenarios are however not exclusive as outdoor applications can suffer from bad GNSS signals (e.g. urban canyons [83]) and indoor ones might as well have degraded visual conditions (e.g. fire smoke [84]). Another challenge appearing in indoor conditions is the lack of detectable features due to the homogeneity of the environments.

The term DVE was first introduced in the military field to address the situation where the landing of helicopters is altered by the environmental conditions. Projects have been carried to enhance visual and situational awareness and mitigate the safety hazards associated with landing in DVE, for example with a specific LiDAR sensor [85]. The same term is used to autonomous navigation when robots or autonomous vehicles struggle to evolve in scenarios where the data acquired by perception sensors is altered by the properties of the environment. This issue is often coined as the “All-Weather”, “Adverse Weather” or “Harsh Weather” problem of autonomous driving but we choose to refer to it as DVE. Experiences from field operations (with EasyMile vehicles) have shown that different scenarios are part of the problem, such as classical weather phenomena (e.g. rain, fog or snow) composed of hydrometeors in the atmosphere. Smoke particles are also found at industrial sites or during fires, sand winds in desert areas, leaf fall in autumn and, although we have killed a fair amount of insects, swarms can sometimes fly through the atmosphere. Overall, DVE denotes any environment in which the data produced by perception sensors is altered due to interactions between the considered signals and particles in the atmosphere [86]. In this thesis, only rain, fog and smoke conditions are considered. Global illumination caused by the sun or artificial light sources can also be part of the problem but are not addressed here. The nature and intensity of DVE events are highly variable phenomena and, the impacts caused on the sensors data vary similarly. To clearly understand the behavior of a perception sensor under particular condition, one should understand the multiple interactions involved in the sensor processing chain (fig.2.1).

Different metrics and models exist to describe the nature of DVE events. Metrics can be specific to the type of event or common and sensors are employed to measure them. Scattering theory explains the interactions involved between electromagnetic signals and particles where several effects occur (single or multiple scattering, backscattering, refraction, absorption, [87]). To enable the quantification of the impacts of DVE on LiDAR and machine learning algorithms, it is necessary to build consistent datasets of LiDAR point clouds labelled by DVE properties. The construction of such datasets with natural conditions is a hard task because it requires a lot of time and infrastructure.

The production of artificial conditions allows to collect data under controlled and reproducible conditions, thus overcoming the impracticality of natural conditions. Most of the results presented in this thesis come from experiments made with artificial conditions. Ideally, the produced conditions have very close properties compared to natural conditions so that the analysis and developed algorithms match real world scenarios as much as possible. As a result, the comprehension of the mechanisms used to generate DVE conditions and properties produced is increased in our case. In practice, the generation process makes it very difficult to approach real world conditions and precisely controlling these events concurrently with specific properties is hard as a multitude of factors influence the behavior of the particles (temperature, pressure, wind). Finally, one should not have absolute confidence in the generation methods of artificial conditions and thus in the data produced, because the produced conditions can sometimes be very different from their natural form. An emphasis should be put on the physical characteristics of the particles to understand and control as much as possible the interactions at stake. Extensive work on artificial generation of DVE conditions for perception sensors is proposed in [88].

This chapter presents each of the DVE conditions considered (rain, fog and smoke), from their natural origin to artificial generation and a look at their characteristics is given along with the specific sensors which aim to measure them. Finally, the different scattering interactions possible between the DVE conditions and LiDAR signals are detailed.

3.1 Rain

Rainfall comes from the condensation of water vapour inside atmospheric clouds, which creates falling drops under the influence of gravity. These drops fall locally under the clouds and are of various numbers with various shapes and speeds. Although they are influenced by wind conditions, the spatial repartition of raindrops on the watered area is generally considered homogeneous. When raindrops fall from the clouds, their shapes evolve as they can interact together and for example fuse to create bigger droplets, or separate into smaller ones. The diameter of a raindrop generally ranges from 0.1 mm to 9 mm but large drops will likely be fragmented into smaller ones [89].

The diameter distribution gives a good insight on the rain nature and is modelled by the Marshall-Palmer law [90]. Droplets of natural rain have time to reach a terminal velocity during their fall and this speed can vary. The two metrics of diameter and speed are often displayed together on a 2D histogram to assess the distribution of both properties, as shown fig.3.1. The relation between the diameter and velocity of natural raindrops is studied in [91, 92] and the models are displayed on fig.3.1.

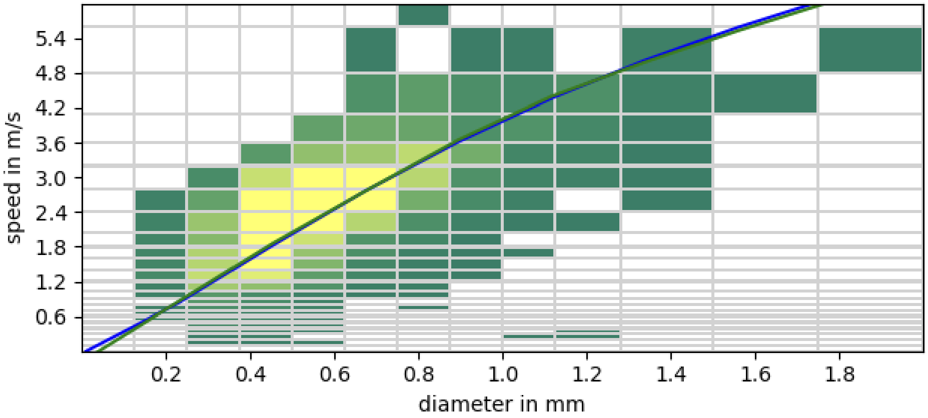

Figure 3.1: Example of a 2D histogram of rain drops diameter and speed measured by the Parsivel optical disdrometer during one of our experiments. The histogram is relatively consistent with the model represented by the blue and green curves [91, 92].

Rainfall rate (or precipitation intensity) is another common metric to quantify rain. Derived from the values of the 2D histogram or measured by a dedicated sensor (pluviometer or rain gauge), it represents the height of water that has fallen during an amount of time and is usually given in mm/h. Rain events can happen very locally in time and space, with short rain showers or long drizzles, which could both result in the same intensity in mm/h but the resulting impacts on perception sensors would greatly differ. This measure suffers from the fact that it is averaged over time and it ignores raindrops diameter and speed. Rain intensity is most likely useful for meteorological and agricultural activities but is less applicable for physics because it does not indicate any information regarding the spatial repartition of the rain nor diameter and speed characteristics of the drops, which are essential for understanding the impacts on LiDAR. We will see later, in chapter 6 that it can hardly be used as such to give context on LiDAR degradation.

Many sensors are available to measure the properties of rain [93]. The most simple one is the mechanical rain gauge (or pluviometer), it measures the amount of liquid in a predefined area and over time. Then, advanced devices give precise measurements of the properties of rain. A disdrometer is a sensor capable of counting rain droplets as well as measuring their diameter and speed. Two disdrometers have been used in this thesis and are presented here : Parsivel optical disdrometer [94] and Lufft WS100 radar disdrometer [95]. Optical disdrometers (illustrated fig.3.2) use an emitter and a detector facing each other which share a flat and rectangular light beam in a small space where drops pass through. When a droplet falls through the beam, the amount of energy received by the detector is reduced during a certain amount of time. This time and amplitude of signal loss respectively give the velocity and diameter of the droplet. The result can be displayed as a 2D histogram, as shown on fig.3.1. As part of the experimental results, the Parsivel disdrometer is described in detail in chapter 6.

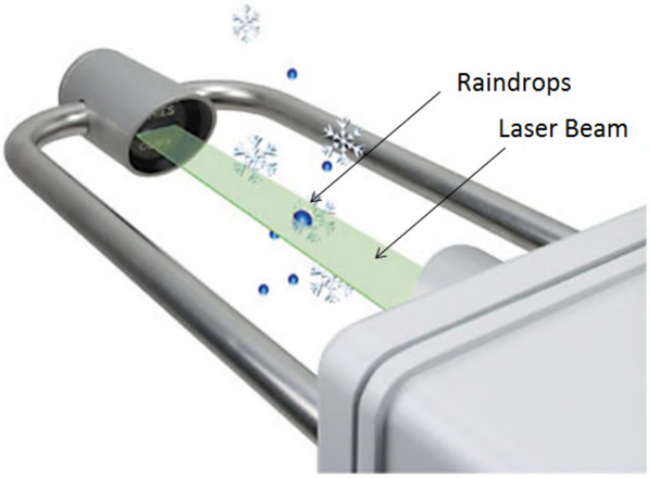

Figure 3.2: Illustration of disdrometer measurement principle, excerpt from [93].

The Lufft WS100 radar disdrometer uses 24 GHz radar emissions and the Doppler effect to count the droplets falling on it while measuring their size and speed. The sensor is also capable to differentiate the types of precipitation between rain, snow, sleet, freezing rain and hail. Some results of this sensor are shown on fig.3.3. Fig.3.3a shows three cumulative rain intensities within 1, 6 and 60 minutes. This shows how an intensity measured in mm/h is averaged over time. Then, figs.3.3b and 3.3c show the evolution of the number of drops according to different size classes and the size histogram during the whole acquisition time. This sensor is not exploited in this thesis.

These sensors are useful to give precise measurements but the obtained data is very local and is sometimes not adapted to changing weather events. Alternatively, external information available on internet could be a solution. Online data produced by specialized company (like MétéoFrance in France [96]) or open Application Programming Interface (API) can be widely spread on territories because they monitor multiple weather stations or even satelite sensors. However, such data can be relatively expensive and does not always contain the appropriate data every time. Moreover, their measurements tend to have a global perspective and may not be directly applicable to our specific local conditions, where the sensors are directly influenced by the DVE particles.

Artificial generation

The production of artificial rain has long existed in active fire protection systems, agriculture and also movies. Nozzles (or sprinklers) of different types allow to water areas with various spray-patterns and rain properties. The interest of such systems is mostly to produce specific quantities of water on predefined areas without being too interested in the properties of the produced rain (diameter and speed). While it is fairly easy to produce precise rain intensities, it is much more difficult to approach natural rain in terms of spatial distribution of raindrops in a whole watered area.

To reach natural rain with an artificial generation system, it is necessary to place the nozzles quite high to give enough time to the drops to reach their terminal velocities. In addition, the spray patterns should be very elaborated in order to obtain a homogeneous repartition of the droplets on the watered area. Overall, achieving this is quite difficult and potentially very expensive.

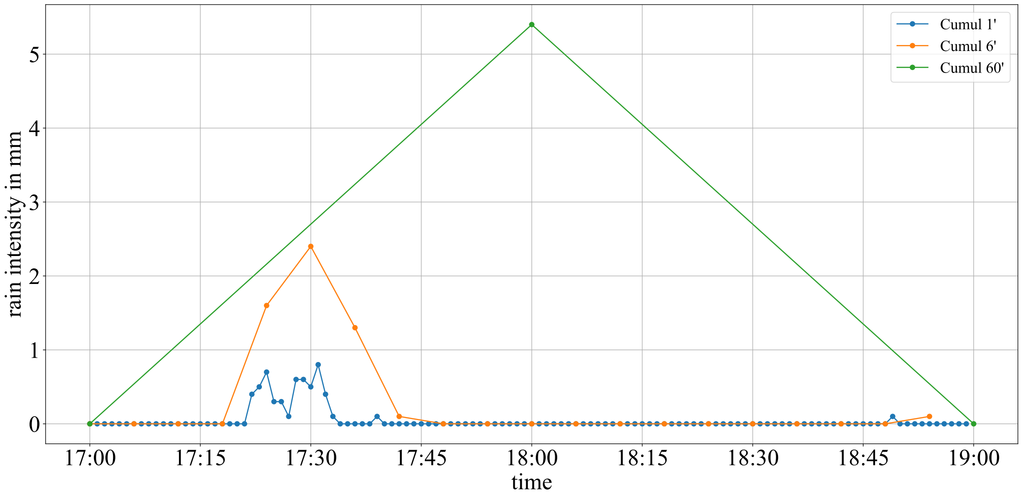

(a) cumulative rain intensities.

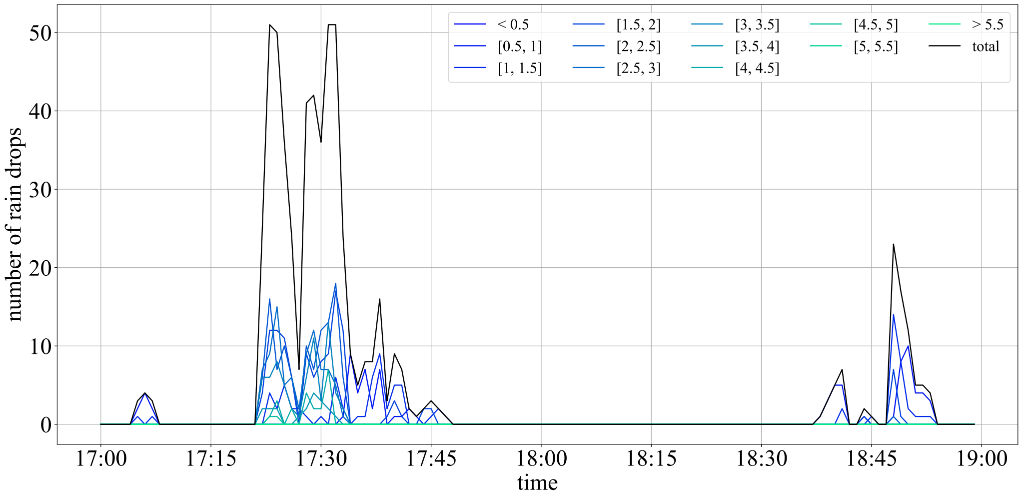

(b) Number of rain drops according to size classes (in mm).

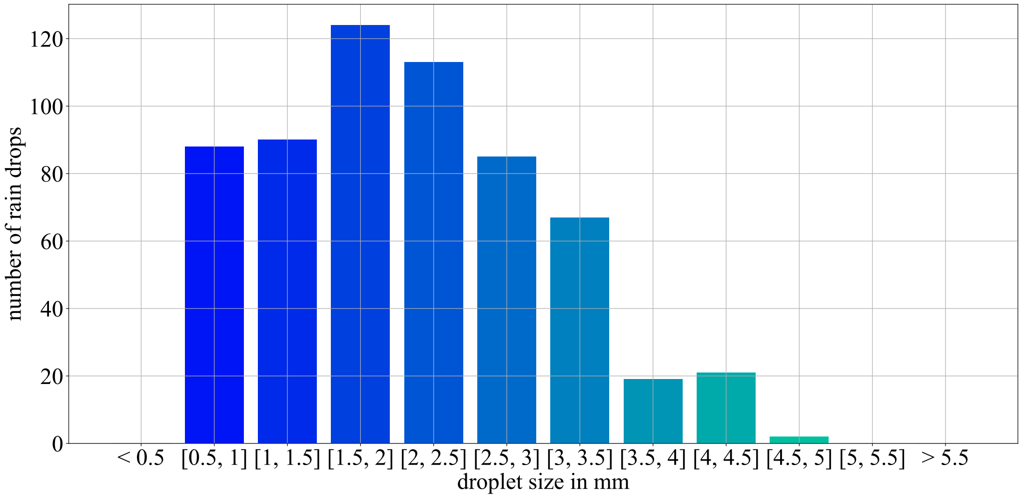

(c) Size distribution of the rain drops.

Figure 3.3: Measurements of the WS100 Lufft radar disdrometer during a 2 hours acquisition.

3.2 Fog

Fog clouds consist of tiny suspended water particles in the atmosphere. Similar to clouds present in the Earth's troposphere. Their creation originates from the condensation of water molecules in gazeous state, similarly to rain but fog droplets differ by remaining light-weighted and only float in the air. Fog is highly dependent on temperature, humidity, pressure and wind. It usually appears when the relative humidity (percentage between the present state of absolute humidity and a maximum humidity given the same temperature) is close to 100% and when the difference between air temperature and the dew point (temperature at which water vapor condenses) is below 2.5°C. In general, fog is caused by the addition of moisture to the air or by a drop in the temperature of the ambient air, causing condensation.

Different types of fog are possible based on how the condensation occurs, the two main ones being radiation and advection fogs. Radiation fog forms when the ground cools down at night, usually under clear skies and light winds. As the ground cools, it cools the air in contact with it, causing the moisture in the air to condense into fog. This type of fog is common in valleys and other low-altitude areas. Radiation fog has the potential to be extremely dense, resulting in strong visual degradation. It is primarily composed of small water droplets, typically measuring several microns in size. Advection fog, on the other hand, forms when warm, moist air moves over a cooler surface, such as a body of water or a cold landmass. As the warm air cools down, the moisture in it condenses into fog. It can occur even during the day and can be persistent and widespread and is composed of bigger droplets. In this thesis, only radiation fog was studied due to its availability in climatic chambers.

Fog conditions can be quantified by assessing visibility or particle size distribution (PSD), both metrics are presented in section 3.4. Extensive research on fog is available in [97].

Artificial generation

Fog particles can be produced using the same kind of system as rain with nozzles or sprinklers [98]. To obtain small particles of fog, the water is forced under high pressure which atomizes the droplets. This results in brief pulses of fog that are repeated to spread throughout a room. To change the nature of the fog conditions, different composition of water can be used, with for example normal tap or demineralized water. The resulting fog is influenced by temperature, which changes the condensation state of the particles. To stabilize the conditions and thus the visibility value, sensors monitor the visibility value and it is adjusted by injecting (more or less) fog pulses.

3.3 Smoke

Smoke consists of airborne particles resulting from the combustion of material, its composition and properties depend on the material nature and combustion process [99]. The most common type of smoke (coarse mode) is composed of particles between 2.5 μm and 10 μm. Smoke conditions can appear in specific scenarios for autonomous navigation. The operation of robots inside fire scenes imposes to handle smoke conditions, both for indoor and outdoor use cases where perceptions sensors suffer [100]. More generally, LiDAR can be used for fire detection to prevent, for example, forest fires [101]. In addition, smoke aerosols can also be present on industrial sites dealing with the combustion of heavy material, which is one of the targeted application for autonomous navigation.

The quantification of smoke conditions can be similar to the one of fog scenarios, with the assessment of visibility or the Particle Size Distribution (PSD) because the interactions are similar (both metrics are presented in section 3.4. The differences between fog and smoke are mostly based on the composition, with water and carbon particles respectively. This difference impacts visibility by altering the scattering effects.

Artificial generation

Producing artificial smoke is common in the world of entertainment to create visual effects associated with light, where people are far from suspecting the multiple scattering interactions at stake. The production of smoke is now made easily possible by the vaporization of glycol based fluid and condensation of the particles outside of the machine at colder temperature. Smoke generation is also interesting for simulating degradation problems in Free Space Optical (FSO) communications [102, 103], where smoke or fog degrades the quality of optical signals, as in our application with LiDAR (although attention must be given to the wavelength difference)

3.4 Visibility

A common metric to quantify the density of fog or smoke conditions is visibility. Visibility is first a very human-eye visual metric commonly used in the aeronautic world to assess the maximum distance observable for pilots around landing strips, also known in this case as the runway visual range [104]. Then, its definition has been precised by the International Civil Aviation Organization as “the greatest distance at which a black object of suitable dimensions, situated near the ground, can be seen and recognized when observed against a bright background”X or quantitatively with a contrast ratio of 0.02 [105]. According to the Office of the Federal Coordinator for Meteorology [106], fog reduces visibility to less than 1 km, whereas mist causes lesser impairment of visibility. Even though visibility is mainly used to asses fog severeness, it also quantifies other DVE conditions as visual information is altered regardless of the condition type. Transmissometers or scatterometers (fig.3.4) are two types of optical sensors designed to measure visibility. The first measures the extinction of a light beam over a certain distance (tens of meters) and converts it into visibility whereas the second one obtains visibility by observing lateral scattering effects with a laser and a receiver close to each other (tens of centimeters). In any case, these measurements are made locally and at specific wavelengths, which are often not the same as those of LiDARs.

(a) Biral scatterometer.

(b) Tranmissiometer system on an airport.

(c) Lighthouse PSD sensor Handled 3016.

Figure 3.4: Pictures of visibility and PSD sensors.

Visibility is related to the atmospheric extinction coefficient during scattering interactions with spherical particles [107], using the Koschmieder Formula

$$V = \frac{3.912}{\sigma} \tag{3.1}$$

where V is visibility [m] and σ [m^{-1}] the extinction coefficient, mostly determined by atmospheric particles for small visibilities (less than a few kilometers).

To quantify fog conditions, the Droplet or Particle Size Distribution (DSD or PSD) is also valuable because it varies with the fog type and thus gives more details on the nature of the fog conditions. The PSD can be measured using a particle counter such as the Lighthouse Handled 3016 used in this thesis (fig.3.4c). This sensor is able to count the number of particles within 6 classes from 0.3 μm to 10 μm. The PSD counts can be converted to visibility using the Mie theory [107]. The extinction coefficient of eq.3.1 is expressed with the particle cross section area πr^2 [m^2] and the extinction efficiency factor Q_{ext} [$\emptyset$],

$$\sigma = \pi r^2 Qext \tag{3.2}$$

In a polydispere distribution of particles, the total extinction is the sum over all particles sizes,

$$\sigma = \sum \pi r^2 Q_{ext} n(r) \tag{3.3}$$

where $n(r)$ is the number of particles for each radius r. The Q_{ext} efficiency factor is computed using Mie theory and the complex refactive index of the particles, the particle size and the wavelength at which the visibility is desired. Finally, the PSD values is converted into visibility values, as shown on fig.3.5, which shows the evolution of particle concentrations measured by the Handled 3016 sensor during smoke conditions and for each size class. As smoke conditions dissipate over time, the concentrations of particles of all sizes decrease. The black dots show the conversion of the PSD counts into visibility.

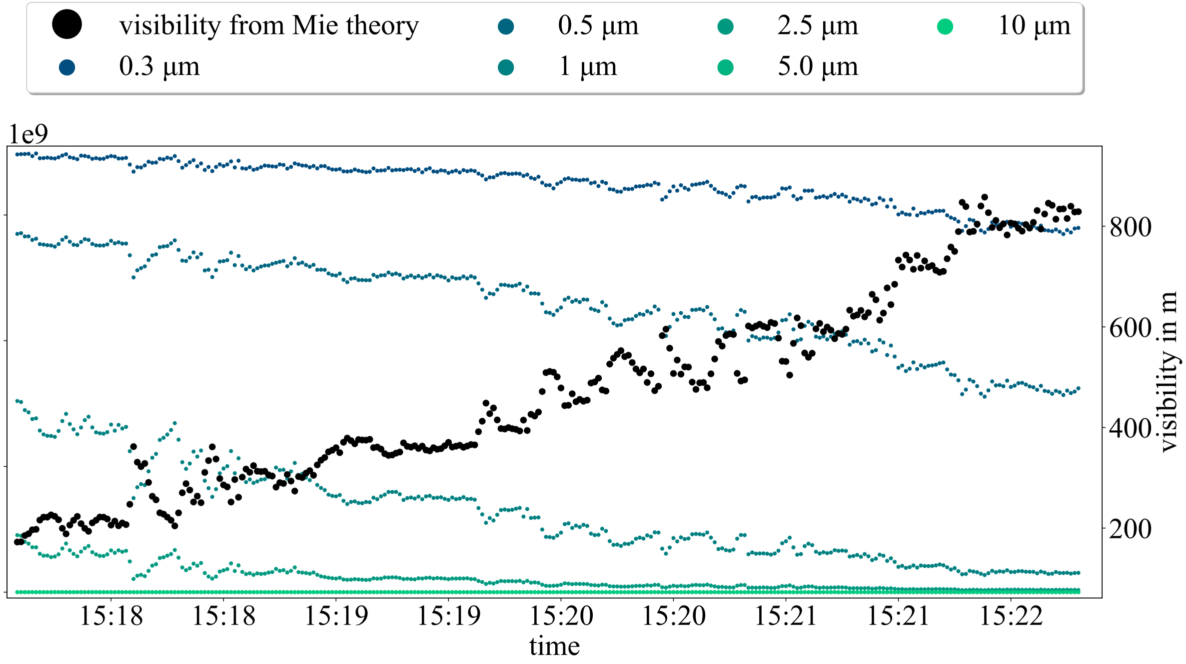

Figure 3.5: Concentration of particles for each size class in smoke conditions and conversion into visibility.

3.5 LiDAR facing DVE

A LiDAR system receives the laser signals backscattered and attenuated by constituents of the atmosphere, in addition to passive contributions, as shown on fig.3.6 (more explanations are available in chapter 2). In DVE conditions, the atmospheric contribution is complicated with the addition of large particles. Precipitation, such as rain and snow, is composed of large drops of water and falling snowflakes. Aerosols, like fog and smoke, are made up of numerous "floating" particles. With which the laser beams sent by a LiDAR interact, altering the detected signals.

The principles involved when a beam of light travels through DVE particles relate to the scattering theory [108]. This last section gives a look on the theories which describe the interactions between electromagnetic waves and particles, from single scattering in rain drops to multiple scattering in aerosols. As this subject is highly complex, with still on-going research activities [109, 110], our objective is to give an overview of the concepts involved and give a few references rather than thoroughly describe the underlying physics.

In addition, the conceptual impacts, firt introduced in [111], caused by these scattering phenomena on the LiDAR point clouds are given simultaneously. The reflection effects result in backscattered signals coming back to the detector. Depending on the particles' properties, these backscattered signals have varying amplitude and can potentially become true detections for the LiDAR sensor, namely "false positives". Refraction, absorption and diffraction effects cause loss of laser signal and result in "false negative" detections, degradation of the range measurement and reduction of intensity and maximum range [112].

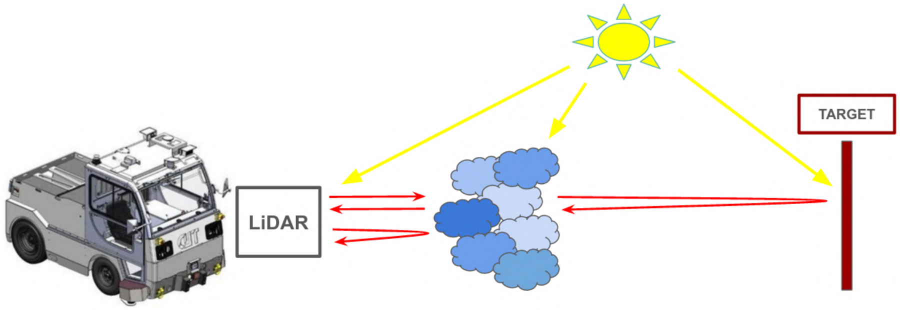

Figure 3.6: Schematic representation of contributions involved in an outdoor LiDAR scenario.

Rain

In rain conditions, single scattering effects happen inside each rain drop with reflection (backscattering), refraction and absorption effects (see fig.3.7). At each interface within a drop, the direction of a light beam undergoes reflection and refraction from the changing index of refraction. Fig.3.7b shows that these effects depend on the wavelength of the light encountering the drop (spoiling the magic of rainbows). Overall, the photometric effect of a raindrop on light is comparable to a wide-angle lens [113]. These effects explain how LiDAR point clouds are potentially impacted in rain conditions. The individual drops can be detected because of reflection at the drop's interface, creating false positives points and, refracted laser beams can either reach objects with biased incidence angle (leading to incorrect range measurements) or never reach any object (false negative). Absorption effects inside water also contribute to the degradation of LiDAR signals. The relation between rain rate, drop size distribution, extinction coefficients and LiDAR signals is studied in [114].

It has been shown in [115] that the number of false positive points generated by a LiDAR in normal rain (a few mm/h) is relatively low and they should be easily cleaned with denoising techniques (see section 4.2). Still, extreme rain conditions will most likely cause more impacts because of more and larger drops. However, the impact on signal intensity from refraction and absorption effects and thus the point cloud quality is significant [116, 117, 118]. Changes in intensity, reduction of range and number of detections are observed when facing outdoor rain conditions. In artificial rain conditions, it is difficult to conclude because the properties of the generated drops differ from natural conditions [119, 120]. Instead of homogeneous falling rain drops, very local rain showers are seen as illustrated on fig.3.8. Our artificial rain experiments conclude similarly (see chapter 6).

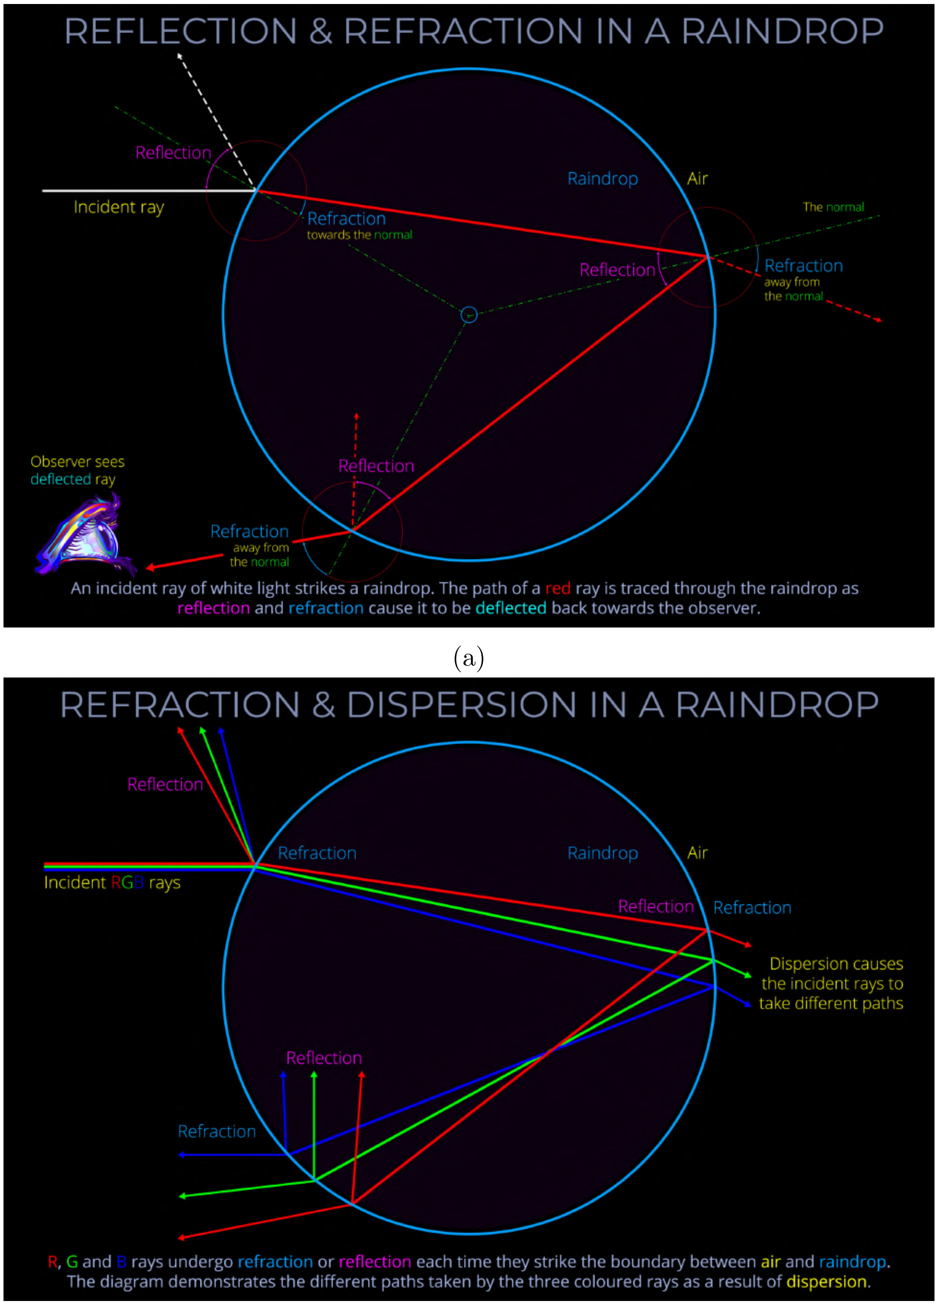

An incident ray of white light strikes a raindrop. The path of a red ray is traced through the raindrop as reflection and refraction cause it to be deflected back towards the observer.

(a)

R, G and B rays undergo refraction or reflection each time they strike the boundary between air and raindrop. The diagram demonstrates the different paths taken by the three coloured rays as a result of dispersion.

(b)

Figure 3.7: Schematic representation of the behavior of a beam of light inside a raindrop [121].

Another difficult problem to solve in rain conditions is the wetting of surfaces and stagnant water. First, water drops can remain on the sensors optical window which can create merging effects [122] and overall degradation of perception quality [123], similar to dirt on optics [124, 125]. This issue is not considered in this thesis and for each experiment with rain conditions, it was ensured that the sensors were free of any direct contamination. This is not the case for aerosol conditions though, where it is difficult to prevent particles from making contact with the sensor window. Then, watered ground strongly impact the quality of ground detections. Specular reflections (mirror effect) are observable in the presence of water puddles, laser beams reflect on them and potentially bounce on further objects to generate false detections under the ground. If no further target is impacted, the laser beam is lost and no point is generated. These effects alter the performance of ground estimation techniques, illustrations showcasing it are shown in section 6.1.

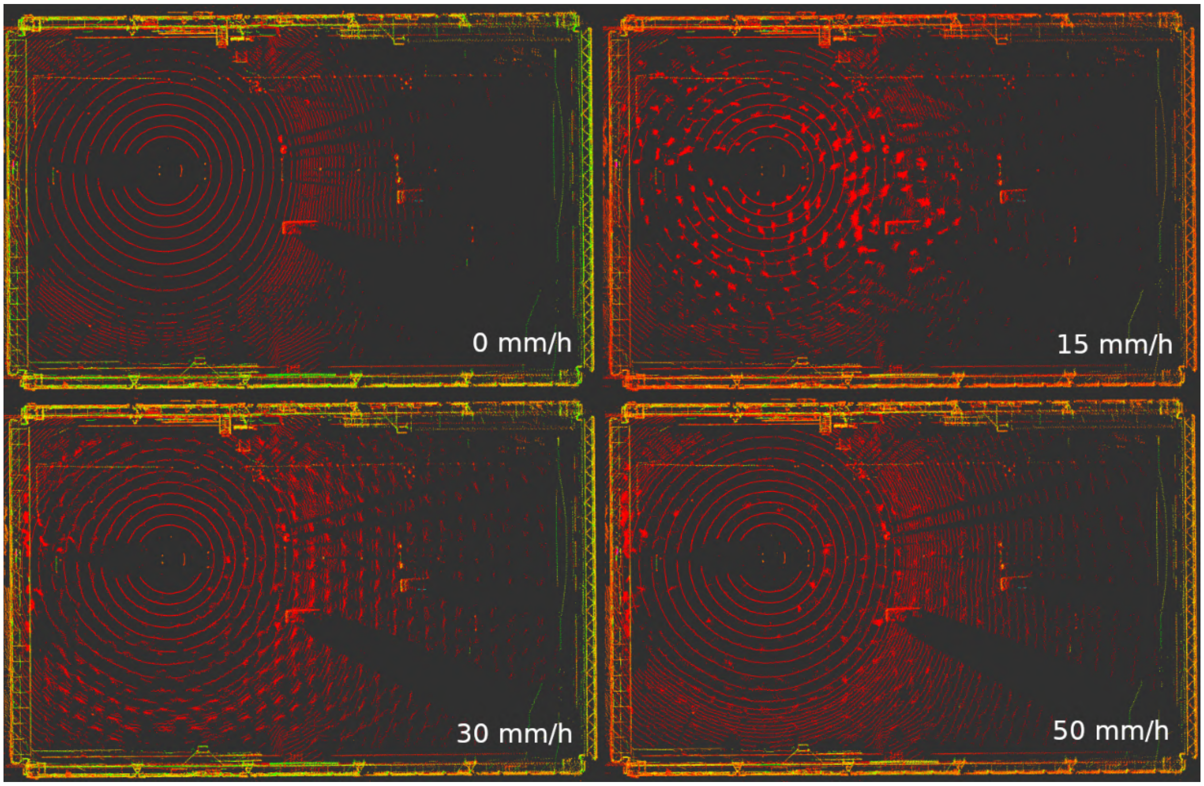

Figure 3.8: Rain showers in LiDAR point clouds in artificial rain conditions, excerpt from [126].

Aerosol

Inside scattering mediums (also called scattering, diffuse or turbid medium), particles are abundant and much smaller than raindrops (a few μm). Multiple scattering effects dictate the behavior of photons as all particles in the path of a light-beam act on the resulting signal [127].

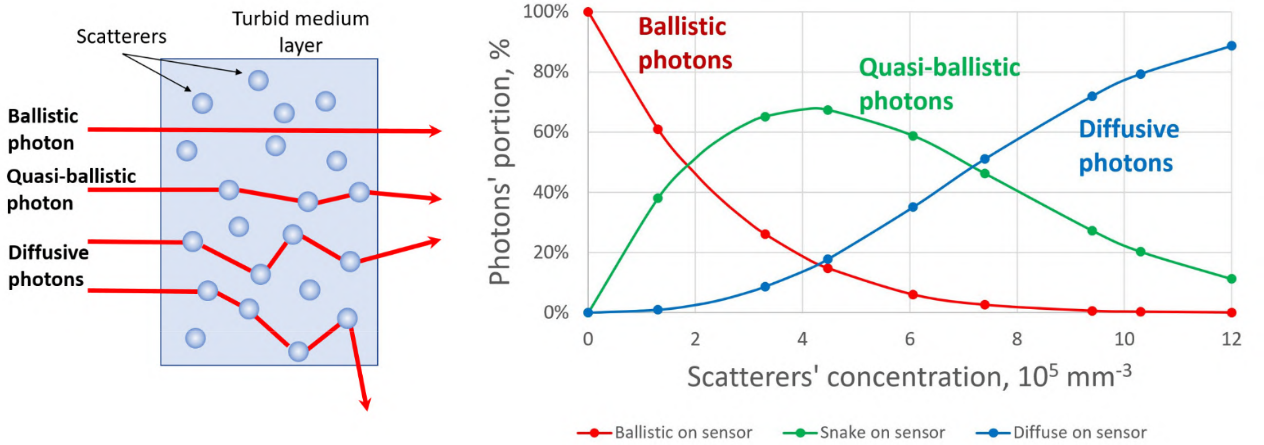

(a) Illustration of the trajectories of photons traveling through a turbid medium.

(b) Monte Carlo simulation of the portions of photons able to cross the medium according to the particles concentration

Figure 3.9: Photons can have multiple behaviors inside scattering medium, excerpt from [128].

Fig.3.9 shows the three types of trajectories of photons passing through a diffuse medium: ballistic, quasi-ballistic and diffuse. These trajectories are the result of multiple individual interactions between the emitted light wave and the particles. Note that the graph could be extended with the backscattered photons which return directly to the source after having bounced on the particles.

The theory of scattering of electromagnetic waves by a particle (considered as a homogeneous sphere) is often divided into Rayleigh and Mie scattering theories, which differ according to the size ratio between the wavelength λ and the particle radius r , defined as

$$x = 2\pi \frac{r}{\lambda} \tag{3.4}$$

As illustrated on fig.3.10, depending on the value of $x \sim r/\lambda$, a different model governs the cross section (the maximum amount of optical flux reflected back to the source).

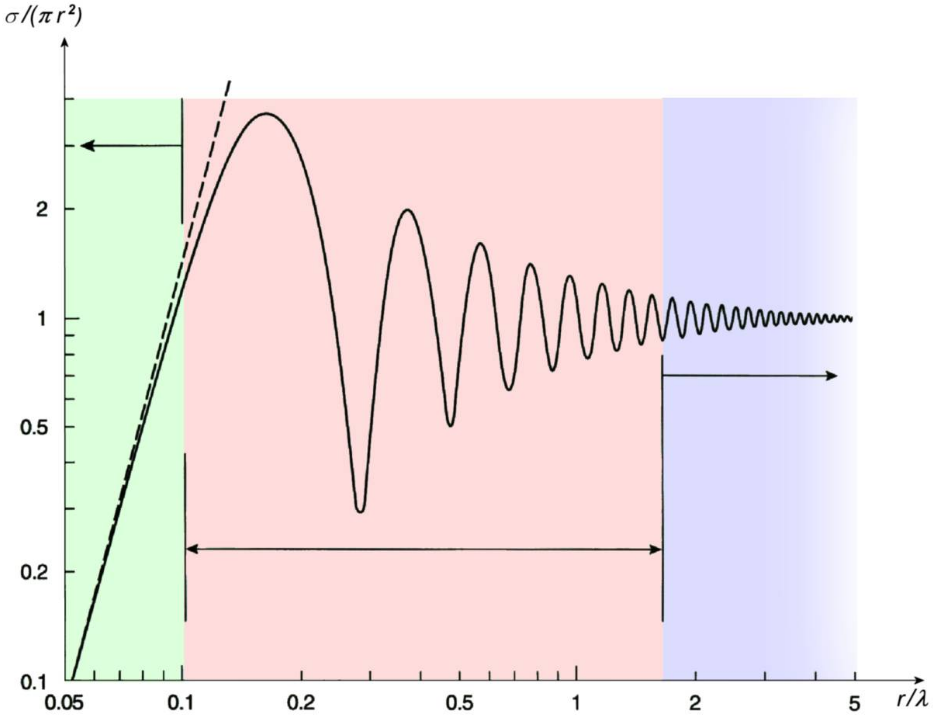

Figure 3.10: Rayleigh (green), Mie (red) and optical (purple) regions for cross sections according to the size ratio.

And, fig.3.11 shows the behavior of scattered light in the case of Rayleigh and Mie scattering. If $x < 0.1$ (green region), the particle size is much smaller than the wavelength and the Rayleigh theory is applied. The graph shows that, in this region, the cross section is highly dependent on the size parameter and thus the wavelength. Additionally, the intensity of the scattered light is identical in the forward and reverse directions (fig.3.11) and inversely proportional to the wavelength to the power of 4 [129]. As a result, the scattering of blue light in air molecules is dominant compared to other wavelengths of the visible spectrum and this is why the sky is blue. When $x \sim 1$ (red region), the particle size is close to the wavelength and the Mie theory is applied. The majority of the scattered light goes in the same direction as the incident light and all wavelengths of the visible spectrum are scattered equally [130], which is why the clouds are white. Otherwise when $x > 1$ (purple region), the particle size is greater than the wavelength, geometry optics laws are applied.

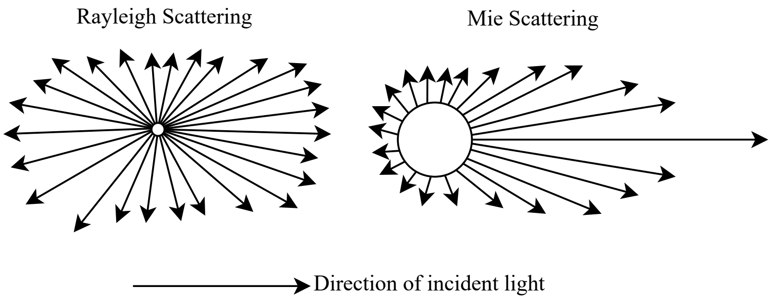

Figure 3.11: Different scattering effects when a laser beam hits a particle.

Fog or smoke conditions are more likely to degrade LiDAR data due to multiple scattering effects and the high density of particles. The conceptual impacts resulting on the point clouds is nicely described in [131], from which is extracted fig.3.12. Facing an aerosol and a solid target, a LiDAR sensor can exhibit different behaviors. Depending on the properties of the cloud and the target (cloud density, composition, reflectivity), the generated detections vary. A very dense aerosol volume can either generate a complete backscattering of the emitted lasers (b) or a total annihilation of the signals inside by multiple scattering effects (d). In comparison, detections solely on the target (a) or on both the target and the cloud (c) can appear if the aerosol is less dense.

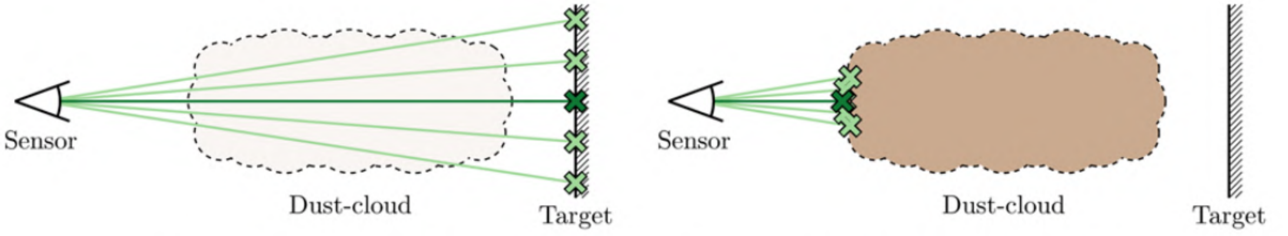

(a) Returns from the target.

(b) Returns from a dust-cloud between the sensor and target.

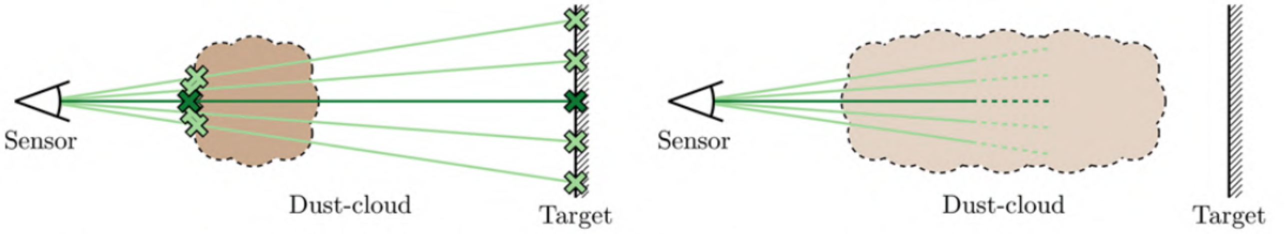

(c) Returns from both the dust-cloud and target.

(d) No returns from either the dust-cloud or the target.

Figure 3.12: Four scenarios where laser light face aerosols, excerpt from [131].

Quantitative analyses have been carried out in various works [132, 85, 119, 133, 134, 135, 136, 120, 137, 138]. In all of these works, strong influence is seen both in terms of false positives and degradation of point cloud quality.

3.6 Conclusion

The presented DVE conditions constitute constitute the physical envelope of this thesis. Present in various forms, they involve falling or floating particles large enough to interact with the laser beams shot by a LiDAR system, thereby causing degradations in the acquired data. Their variety of nature, sometimes unpredictable, makes the problem challenging as the results on LiDAR sensors are highly dependent on this changing behavior. While conceptual models attempt to capture the potential impacts on the point clouds, the underlying physical theories involved show that the reality is highly complex. Quantifying their properties is fundamental and enables the correlation (or labelization) between the DVE conditions and the resulting impacts on the LiDAR point clouds. Although we do not delve into the physics, this thesis tries, as much as possible, to understand the DVE conditions in which we have had the chance to make our experiments.

Chapter 4

LiDAR perception in DVE

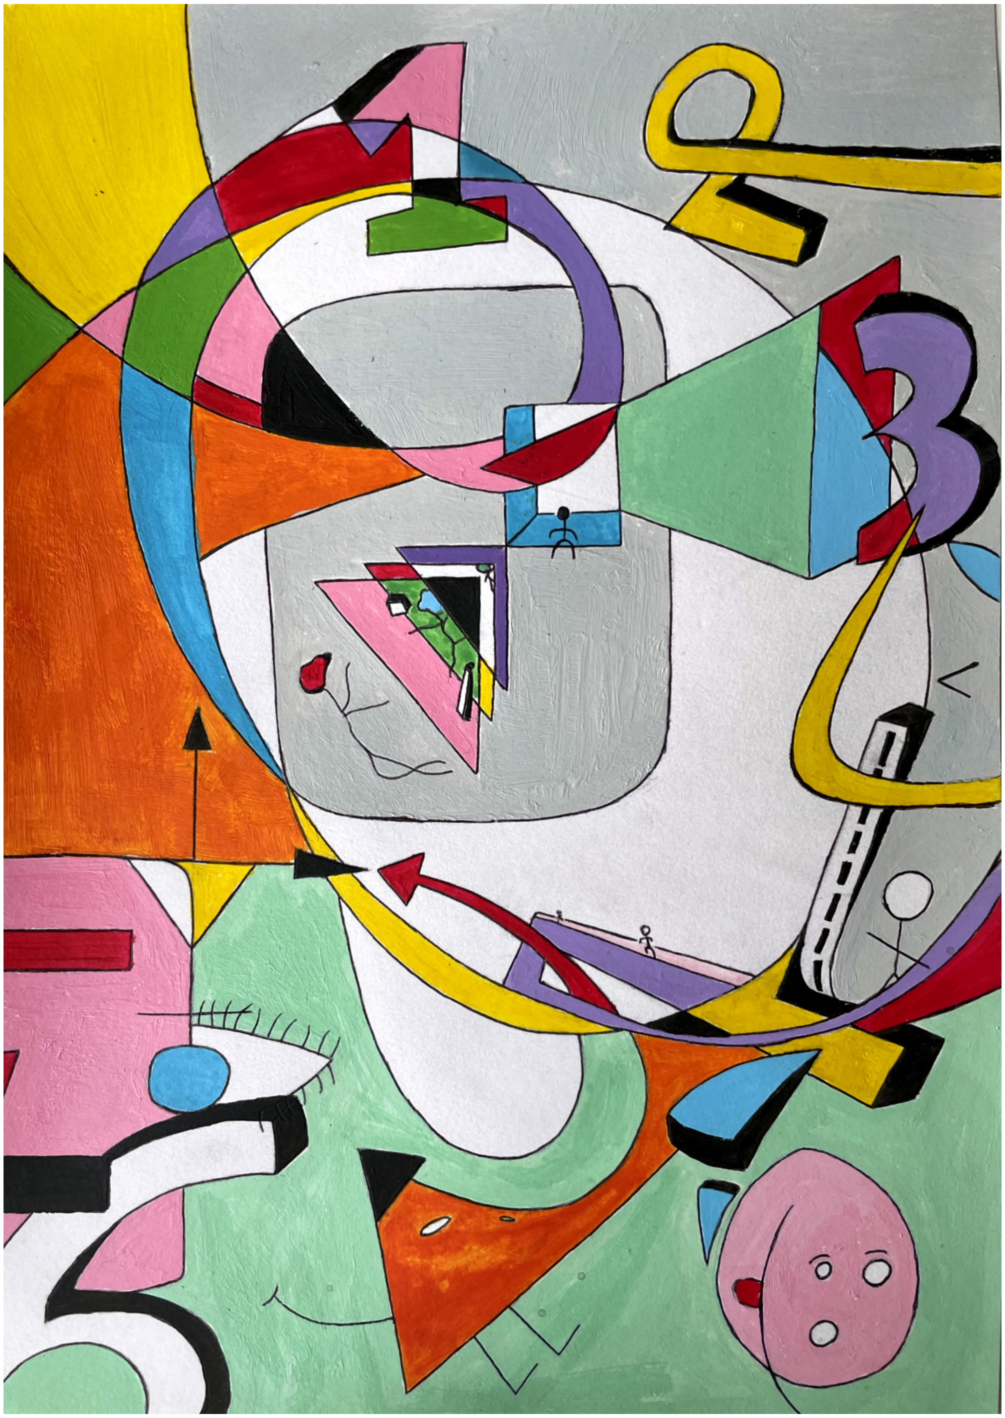

This chapter first introduces the processes involved in an autonomous driving software stack using LiDAR. Then, a look at the state of the art point clouds denoising methods is given, crucial for the enhancement of performances in DVE conditions. We then present approaches focusing on the classification of environmental conditions using LiDAR point clouds, and an overview of the approach we propose.

Prior to the creation of dedicated algorithms, careful attention must be given to optimizing the integration of sensors in order to minimize potential direct degradation. For example, the water flow on a vehicle can lead waterfalls in front of the sensors and cause to serious data degradation or, water splashes on the ground in rain conditions can soil the sensors surfaces. These issues can be tackled by changing the mechanics and integration of the vehicle. Also, the surface of the sensors can globally be often spoiled in DVE conditions, with water or dirt particles. Cleaning systems and weather protection hoods are being designed to mitigate this issue and sensor sets should be thought accordingly. Although this is likely to impact us in all of our experiments, this is not considered in this thesis. Instead, we try to avoid as much as possible any direct contamination and rather focus on the noise produced in the line of sight of the sensors.

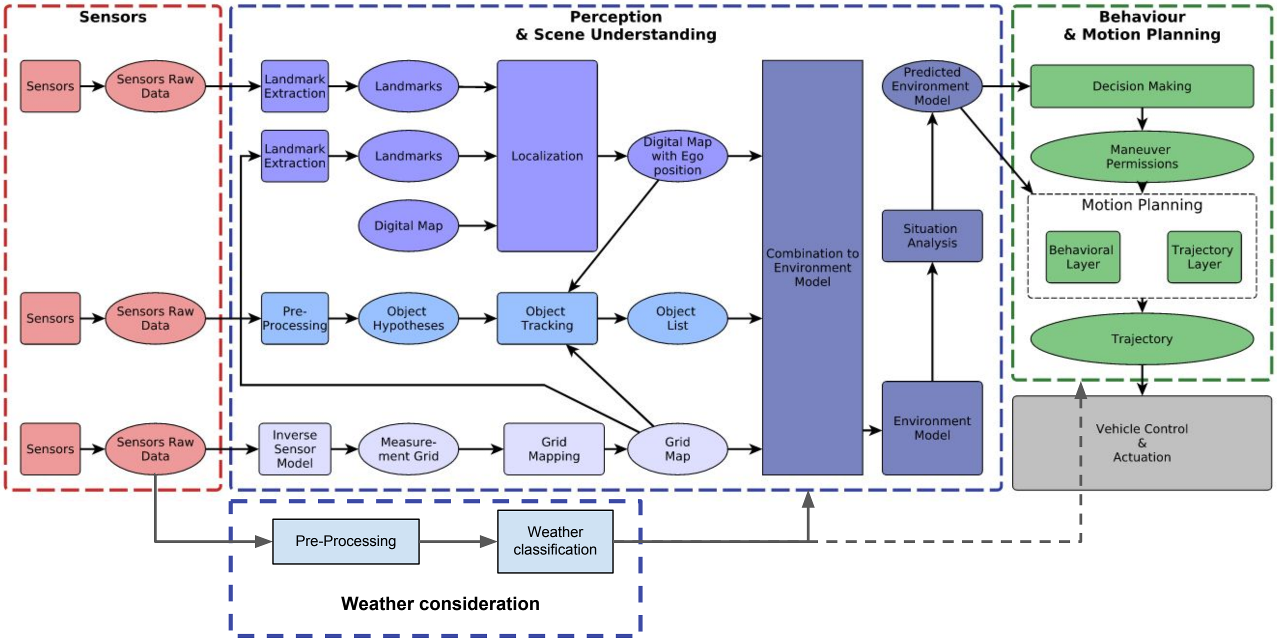

Figure 4.1: Functional system architecture of autonomous driving, excerpt from [139]. Three categories segment the architecture, our work concerns the "Weather consideration" part and belongs to Sensors and Perception and Scene Understanding processes.

4.1 Introduction

Perception tasks, applied prior to navigation, have multiple objectives and focus on the understanding of the scene via the point clouds and the extraction of valuable information [28, 140]. Navigation algorithms (from Behavior & Motion Planning category) do not directly exploit point clouds but rather higher level information such as segmented and labeled objects or target tracks. Localization directly use point clouds to perform Simultaneous Localization And Mapping (SLAM) and place the vehicle within a map [141]. Figures 4.1 and 4.2 show functional system architectures of autonomous navigation with multiple features [142, 139, 143]. Separated into different categories: Sensors provide raw data about the environment, Perception and Scene understanding use the sensors data to extract information and, Behavior and Motion planning defines and controls the dynamic behavior of the vehicle [144, 145].

The semantic information can be extracted to identify the nature of the various objects in the scene (ground, pedestrians, cars), these methods are referred to as segmentation and a multitude of algorithms exist, depending on the application [146, 147, 148, 149]. For instance, in forestry applications, individual trees are extracted for vegetation monitoring [150]. In household robotic tasks or augmented reality (AR) and virtual reality (VR) applications, furniture, stairs, pets, humans, or doors need to be identified [151]. In autonomous driving, the recognition of objects in the urban environment is crucial. A point cloud captured by a LiDAR sensor can contain multiple objects, including the ground, different types of cars, pedestrians, walls, and signs. Understanding the nature of these objects helps to optimize decision-making, as each object may require a specific behavior. For example, the ground and walls are expected to be stationary, unlike moving vehicles and pedestrians. Terrain estimation techniques are used to identify the points belonging to the ground, which then allows for the analysis of drivable areas and estimation of the ground surface [152]. Companies like Xenomatix specialize in flash LiDAR sensors designed to analyze road conditions [153]. Objects above the ground surface can then be extracted and segmented using clustering algorithms, which group points based on their properties [154]. Different classes of objects can be defined based on the properties of these clusters. Once a point cloud is segmented, further processing can be performed using only the points belonging to specific classes. Whether using deep learning or classical approaches, point cloud-based detection tasks rely on point density and the information (coordinates, intensity and others) associated to the points [155]. Finally, the performance of detection algorithms can be affected by weather conditions, which can reduce the number of points detected by LiDAR and alter field values such as 3D coordinates, intensity, and echo pulse width [156, 157]. To address this issue, segmentation algorithms could take into account the current environmental conditions and anticipate the resulting degradation.

The ability to understand the dynamic behavior of objects surrounding a vehicle is crucial for optimizing autonomous navigation decisions. Pedestrians in close proximity to the vehicle can be walking safely on a sidewalk or attempting to cross the road, potentially posing a collision risk. To anticipate appropriate decisions, the vehicle relies on the dynamic state of these objects, typically represented by a 3D velocity vector. FMCW technology for 3D-LiDAR (introduced in section 2.6.1) enables the measurement of object velocity and 3D coordinates through the Doppler effect. This technology provides direct measurements of the required dynamic information, eliminating the need for complex estimation techniques. However, it is not widely available on the market, although some progress has been made with companies like Aeva [57]. Alternatively, significant developments have been made to extract object dynamics from 3D-LiDAR point clouds using direct ToF measurements. While LiDAR sensors and these techniques are relatively recent, target tracking methods, particularly in the military domain, have been developed for a long time [158]. Bayesian filtering methods based on the recursive application of Bayes theorem using measurements and a likelihood model, are commonly employed for target tracking tasks [159]. By incorporating appropriate dynamic models, it becomes possible to estimate the state of multiple variables simultaneously. Similar to other perception tasks, tracking techniques are affected by DVE conditions. False positives present in point clouds can disrupt algorithms and lead to false tracks and associations. Changes in target detection, such as alterations in 3D coordinates and point intensity, also impact the performance of tracking methods designed for nominal conditions. Therefore, addressing these challenges and optimizing tracking algorithms for diverse environmental conditions are essential for robust and reliable autonomous navigation.

But prior to these perception functions, point clouds should be cleared of any false positive points, highly present in DVE conditions. Though, these false positive points can be real assets in the recognition of environmental conditions and the contribution of this thesis and thus should be considered before suppression (see section 4.3).

4.2 Denoising

False positive points can appear when particles are present in the path of the emitted laser beams and their backscattered energy exceed the detection threshold. The condition type (precipitation, aerosol) and its severity (mm/h, visibility, PSD, concentration) influence the impacts and common properties of these parasite points exist and make them predictable and thus removable. Denoising methods are used to remove these points and clean up the point clouds as much as possible before processing them with other algorithms [160, 161, 162, 63, 163, 164, 165]. These work explore various methods and scenarios, from machine learning to pure model-based techniques.

The open source library Point Cloud Library (PCL) offers a series of filtering techniques based on statistical and geometrical information [166]. For example, the Radius Outlier Removal (ROR) suppresses the points depending on their number of neighbors according to a fixed radius. The Statistical Outlier Removal (SOR) filter computes the mean distance of the neighboring points and removes the points considered too far from this mean and a standard deviation. The Voxel Grid (VG) filter downsamples all the points with voxelization. These filters are quite useful to perform quick suppression of noise or downsampling but are not adapted to all situations (precipitation and aerosol of different types and severities) and can have poor computation performances.

An extension of the ROR filter is proposed in [162] by dynamically changing the radius value depending on the distance of the point. This Dynamic ROR (DROR), tested in snow conditions with a Velodyne HDL32 spinning LiDAR, is more efficient than the ROR filter, in terms of computation time and filtering precision, considering the point density repartition of the point clouds with distance. In the Low-Intensity Outlier Removal (LIOR) filter, the intensity point-field information is exploited in snow and rain conditions and an Ouster OS1 spinning LiDAR [63]. It is shown that the intensity returned by the snowflakes is globally lower than the one of real objects so that points with intensity below a threshold value are deleted. To avoid deleting real objects points with low intensity values, a second step similar to the ROR filter is used, that has increased performances compared to the DROR filter. The intensity is indeed valuable information but care should be taken into its characterization because it is very sensor-specific and not always reliable (see 6). By reducing the dimension of the input data (3D point cloud) with Principal Component Analysis (PCA) and thus the complexity of the filtering task, a clustering method to perform noise reduction with the PCA-based adaptive clustering (PCAAC) filter is proposed in [164]. The results are compared with the SOR and ROR methods, using several evaluation indicators. But, the data used for the comparison is taken in clear conditions. An extension of the SOR is proposed in [165] with the Dynamic Statistical Outlier Removal (DSOR) filter and compared to the DROR in snow conditions. The results shown are quite similar than with the DROR but with a better computation time, thanks to a reduced complexity.

Finally, machine learning techniques can also be used to achieve denoising [161, 163]. Support Vector Machine (SVM), K Nearest Neighbors (KNN) and Convolutional Neural Networks (CNN) allow to classify points as false positives with feature vectors of various nature [163]. A voxel-based deep learning technique is presented in [167] to detect airborne particles of smoke and dust.

The denoising task can also be performed directly with FWF signals or Single Photon counts [168, 74]. Such techniques are certainly present inside the sensors but no information is available on this subject for commercial LiDAR sensors.

Ultimately, the challenge of these techniques is as much about the ability to remove false positives as it is about not removing points from real objects (false negative). But these denoising techniques are often based on thresholds and are thus highly tailored to the situation they are designed to solve. An algorithm designed to remove false positive in snow conditions is hardly applicable to fog or smoke conditions because the impacts greatly differ. Therefore, acknowledging the condition type could be needed to know which algorithm to apply. In addition, variations in the intensity of the events change the impacts on the point clouds. Bigger snowflakes will likely return higher energy signals than small ones, and are hence detected at longer ranges, with higher intensity information. Stronger precipitation events contain more falling particles, producing more false positives. Algorithm parameters defined to maximize performance under certain conditions will not work as well under other circumstances. Denoising algorithms should then also be adaptable to the changing properties of the DVE events they are trying to tackle. To that end, the vehicle should be informed of the severity of the conditions (in addition to the type), so as to be able to dynamically set the parameters of the denoising algorithms.

4.3 Classification of environment properties from LiDAR point clouds

DVE conditions, as stated in chapter 3, alter the point clouds produced by 3D-LiDARs. These scenarios exist in multiple forms and happen regularly to produce false positive points, false negatives, reduced range, degraded accuracy and, point-fields measurements. As a result, every point cloud-based algorithm mentioned previously is impacted [156, 157]. Addressing the limitations caused by DVE on autonomous navigation is the subject of ongoing researches and developments to enhance the robustness and capabilities of the perception algorithms. This includes incorporating weather-specific features, optimizing parameter selection strategies, improving training data diversity, considering computational efficiency and, establishing standardized evaluation frameworks to assess performance across different weather conditions. The functional architecture of the autonomous driving system can be extended to consider weather conditions throughout the software stack, as shown in the lower portion of fig.4.1 or fig.4.2.

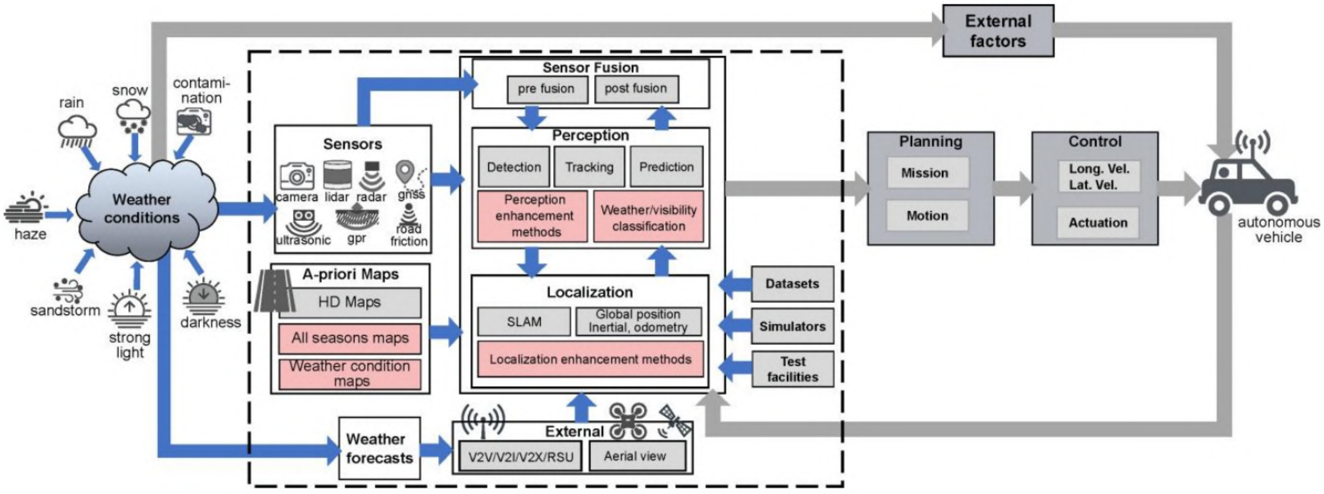

Figure 4.2: Perception stack architecture with weather consideration, excerpt from [156].

As such, the output of a weather classification feature could be sent to further perception processes (or even navigation ones) which benefit from this information. The generalization of filtering capabilities to any weather artifact can be better addressed by studying them and having the ability to better estimate the weather conditions of the autonomous agent. DVE scenarios have highly changeable properties so that the performance of point cloud-based algorithms is highly susceptible to change if the parameters involved are not adapted to the different conditions.

In addition, safety considerations following the IEC TS 62998-1:2019 norm [31]: *Safety of machinery – Safety-related sensors used for the protection of persons*) is crucial to reach high autonomy levels. This norm requires to follow the IEC 60721-3-5 standard [169]: *Classification of environmental conditions – Part 3: Classification of groups of environmental parameters and their severities – Section 5: Ground vehicle installations*. It is stated that “the manufacturer shall specify for all relevant environmental influences the limits for failure to danger condition and for normal operating condition”. Therefore, the consideration of environmental conditions in a vehicle’s intelligence is the key to comply with safety regulations as well as to improve the overall performances of autonomous systems and thus their development

[170, 171].

The classification of environmental conditions by dedicated weather sensors is thoroughly addressed, as described in chapter 3. One could think about installing such sensors on a vehicle, characterize the behavior in dynamic conditions or even design a whole dedicated weather station. This constitutes a research work on its own and has not been considered in this work. These weather sensors have however been used several times during this thesis to study the properties of DVE conditions and provide ground truth metrics and label the LiDAR point clouds accordingly. Also, LiDARs not designed for autonomous driving or robotics but rather environmental physics are relatively well suited to the task with the analysis of FWF backscattered signals from DVE particles [114, 172]. But this type of sensor is not easily integrable into a vehicle. This thesis focuses instead on the exploitation of already embedded sensors and environmental condition classification using 3D-LiDAR generated point clouds, recently treated by a series of authors [173, 136, 174, 126, 175, 176, 177].

These work propose models to estimate either the type of condition (clear, rain, fog) and/or its severity with coarse labels or more precise evaluation of the corresponding metrics (rain fall rate or visibility). Various LiDAR sensors are used but the most common is the Velodyne VLP32 spinning LiDAR sensor, reflecting its dominance in the autonomous vehicle market for the past years. Climatic chambers are used to generate artificial conditions. As mentioned in chapter 3, they enable repeatable and controlled experiments, which are both qualities considering the difficulty to acquire reliable data from outdoor acquisitions.

The impacts produced by the DVE particles on the point clouds are first evaluated, as well as the evolution of meteorological metrics, monitored by weather sensors. Then, the proposed classification methods are either model or pure learning based approaches. The first one represents methods based on the identification of precise variables in the point clouds (coordinates, distance, intensity) and pure learning techniques consist of using specific feature vectors given to learning architectures (SVM, KNN, CNN). Whether model-based or pure learning, these techniques all require large amounts of data. Publicly accessible datasets exist and aim to accelerate the development of autonomous navigation under difficult weather conditions [137, 82, 178, 120].

Part II

Contribution

Chapter 5

Overview of our approach

This chapter describes our approach to facilitate the understanding of the following chapters, which concern quantitative results of LIDAR impacts in DVE (chapter 6), a weather classification model based on the distance distribution of noise echoes (chapter 7) and perspectives for improving our approach (chapter 8). In this thesis, the applied methodology for point cloud processing plays a crucial role, as its significance lies in a probabilistic approach and adaptability to outdoor operational conditions. First, probabilities of detection (POD) of a target and associated noise are assessed. This is useful to compare sensors while assessing their performance in DVE conditions. Also, it enables for considerations regarding safety regulations and contributes to the design of a whole probabilistic solution. Then, these noise detections are analysed and used to build the weather classification model.

5.1 Probabilities of detection (POD)

The proposed approach considers the LiDAR point clouds in relation to the targets and noise characteristics. Following the illustration of fig.5.1, we assess the target and noise POD by extracting both the points detected in a bounding box corresponding to a target location (and dimension) and the points located inside the frustum between the origin of the sensor and the target (noise). In this case (and in single echo configuration), all the lasers fired towards the target can either produce a detection on the target, inside the frustum (from the reflection of DVE particles), or produce no detection at all (the signal is lost).

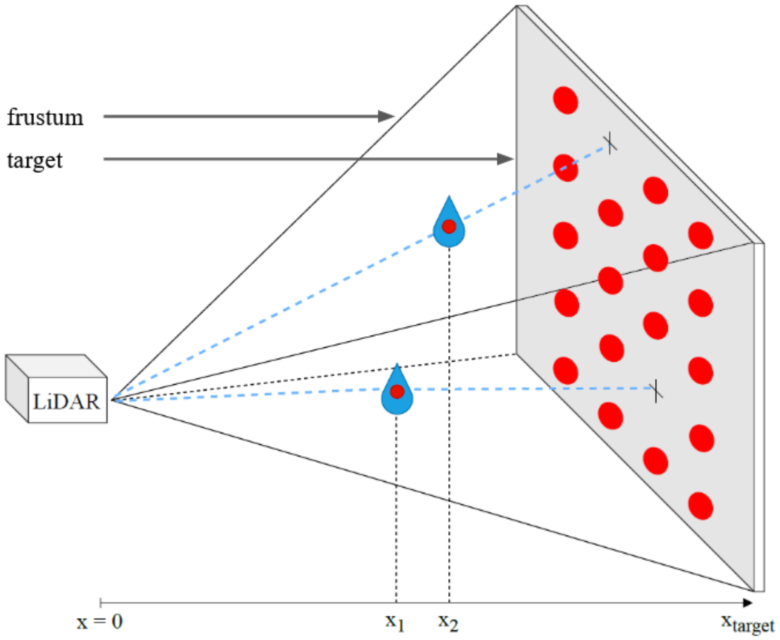

Figure 5.1: Schematic diagram of the point cloud extraction methodology. Points located in a target bounding box and inside the frustum volume are extracted.

Given a fixed number of laser beams shot towards a target, denoted as n_{shots}, the number of echoes detected on the target is n_{target}, while DVE particles can cause additional echoes to be classified as noise, resulting in n_{noise}. Additionally, scattering effects can lead to the loss of certain laser shots, denoted as n_{lost}. We have:

$$n_{shots} = n_{target} + n_{noise} + n_{lost} \tag{5.1}$$

Note that this equation is correct in single echo configuration but a multi-echo capacity would influence it by the consideration of multiple echoes detectable from each laser shot. Then, each number presented can be used to assess the corresponding POD, such that:

$$POD_{target} = n_{target}/n_{shots} \tag{5.2}$$

$$POD_{noise} = n_{noise}/n_{shots} \tag{5.3}$$

This methodology enables the evaluation of POD over time for specific targets and their associated noise under different weather conditions. Normalizing the results based on the number of laser shots facilitates comparisons between multiple sensors operating in the same environment. Additionally, it helps to increase generalization because the drawn probabilities can be attributed to the individual process of emitting a laser shot. As such, the dimensions of the frustum volumes become less important.

In addition to the point counts for assessing POD, other data associated with the points can also offer valuable insights. For instance, analyzing the variations in intensity readings between target detections and those caused by DVE can lead to the identification of potential denoising techniques. Normalization techniques based on the intensity of target points returned under clear conditions can help in displaying and comparing the results as well. But, the intensity values vary significantly between different sensors, and not all sensors provide calibrated measurements. While this information can provide valuable insights into the reflectance properties of objects and potentially contribute to the development of specific perception algorithms (as shown in chapter 6), we found that it is often unreliable. As a result, this information is not extensively studied in this work.

Additionally, the multi-echo labelling has the potential to enhance point cloud density and provide further understanding of DVE conditions, it could also be part of the POD assessments. Unfortunately, this was not considered in our work but our approach is adaptable to it (more details in chapter 8).

5.2 Frustum noise detections

The proposed approach then consists in the analysis of the statistics of the frustum detections. This analysis leads to the weather classification algorithm. Our weather classification technique aims to be agnostic to the environment structure by relying on point clouds extracted from *free-sky* frustums (shown fig.5.2). In outdoor conditions, it is very likely that such space is available. Our approach holds significant value in addressing the problem of DVE using LiDAR due to its focus on improving generalization towards application in operational scenarios. Moreover, it demonstrates compliance with experiments conducted in diverse conditions, encompassing artificial and natural conditions, in both indoor and outdoor environments (if a *free-sky* frustum is available).

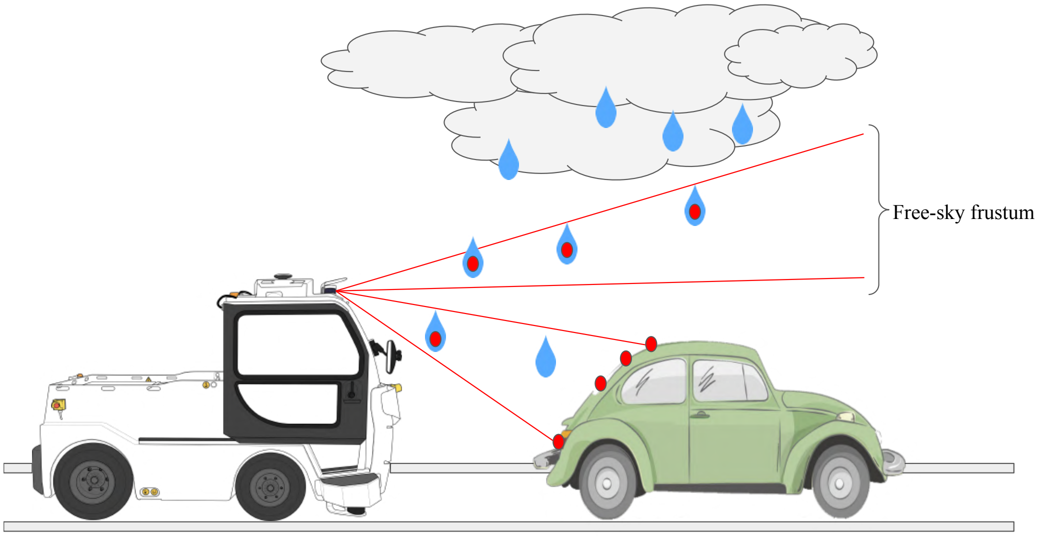

Figure 5.2: Schematic diagram of LiDAR frustums in a DVE situation. A frustum is the geometrical volume between the origin of a sensor and a viewed shape of a target.

The presence of detections within a specific frustum can solely be attributed to the occurrence of DVE conditions. As the severity of these conditions fluctuates, the resulting effects on the point cloud data also change. Consequently, examining the frustum points offers valuable insights into the environmental conditions surrounding the LiDAR system. In addition to depicting the perception capabilities of LiDARs via the POD, this thesis contributes to the analysis of these frustum points. More specifically, the distance repartition of these points is evaluated, leading to an inference model presented later in chapter 7.

One aspect that can be criticized from similar work of the SOTA is their tendency to consider the entire point clouds generated by LiDARs, rather than specific parts of FoV, which poses a challenge in terms of generalization. A 3D-LiDAR (working in direct ToF) detects points from the laser beams it sends out if the signals, received after backscattering in the environment, contain peaks with an amplitude above a threshold [179]. Then, the sensor determines if one or multiple points from these peaks are produced (one point for each peak and depending on the multi-echo capacity). If multiple peaks are found, the order of selection is their magnitude (in most systems).

Solid targets in the path of a laser beam are likely to return relatively high energy signals because of their reflectivity (in very low visibility conditions, the obscurant can return more signal than a solid target at given distance and reflectivity [180]). On the other hand, DVE particles yield scattering interactions and tend to return lower energy signals. This implies that the presence of solid targets can prevent false positive echoes from signals with lower amplitude (e.g. caused by raindrops, fog or smoke particles). This effect is illustrated on fig.5.2 where the presence of the target can prevent from detecting water particles in the frustum defined by the target. Therefore, it is important to note that data obtained in indoor artificial conditions heavily relies on the structure of the climatic chamber. The presence of walls surrounding the area can lead to variations in LiDAR point clouds compared to natural and outdoor conditions, which exhibit diverse structures and a lot of free space. Consequently, training algorithms exclusively on complete point clouds generated under artificial conditions is not adapted for outdoor usage.

During outdoor operations, the sensor is surrounded by a diversity of objects. However, as mentioned earlier, the ability to detect DVE particles is influenced by the presence of a target in the line of sight and their properties. Algorithms trained on noise detections within a specific target frustum become challenging since they require identical targets at the same distance to be applied. To address this, the concept of a *free-sky* frustum becomes relevant, as shown fig.5.2. Such space ensures that no object is located in the line of sight, which is valuable for two reasons. First, it increases the potential for noise detections since no solid targets can reflect the laser beams. Since these detections serve as input for the inference algorithm (introduced hereafter), maximizing their quantity enhances the classification results. Then, it enables generalization as outdoor conditions are very likely to provide parts of the FoV where *free-sky* frustums can be defined. A *free-sky* frustum obtained in artificial conditions is, a priori, similar to one acquired in any outdoor scenario.

To conclude, a *free-sky* frustum volume must be manually defined in the sensor FOV for our methodology to be applied. If such target is not available, it is also possible to consider a solid target (similar to fig.5.1). During this thesis, both these methods are evaluated and the results are presented in the following chapters.

5.3 Inference

Once frustum detections are extracted, we analyse their distance distribution measured by the LiDAR sensor. This statistical analysis leads to the development of a Bayesian inference algorithm designed to estimate visibility and developed in chapter 7. Experiments in rain, fog and smoke artificial conditions, described in chapter 6, are conducted to create datasets including frustum distance distributions (captured by the Ouster OS1-128 LiDAR), paired with ground truth visibility values obtained from the available context sensors (transmissiometer and PSD sensor). In order to model the data appropriately, the Bayesian inference method requires suitable distributions. The parameters of these models are learned through a sampling-based learning process. As the frustum distance distributions vary with changes in visibility, the learned model parameters adapt to capture the characteristics of the data. Consequently, each learning phase yields a set of learned model parameters associated with specific visibility values. During operation, these learned parameters are used to evaluate the likelihood of being in a certain visibility condition. Finally, the algorithm leverages this likelihood assessment to infer the most probable visibility value.

Table 5.1 summarizes the different scientific publications we found that address the estimation of weather properties from LiDAR point clouds.

Reference	LiDAR(s)	Conditions	Method	Input	Output
[173]	VLP16	Rain, fog	Empirical study	Target points intensity	Rain rate
[136, 174]	VLP32	Fog	GPR MDN	Targets range Reflectivity	Disappear visibility
[126]	VLP32 VLS128	Rain	Probabilistic decision tree	Noise points Feature vectors	Rain rate
[175]	Scala	Rain, fog, snow	KNN	Atmosphere & street region Feature vector	DVE type
[176]	VLP32 HDL64	Rain, fog snow	CNN	Full point cloud Feature vector	DVE type
[177]	VLP32 Scala	Rain, fog	KNN, SVM, CNN	Full point cloud Feature vector	DVE type Rain rate Visibility
Thesis contribution	OS1-128	Fog Smoke	Bayesian inference	Frustum extraction Distance	Visibility

Table 5.1: Summary of the SOTA concerning weather classification with 3D-LiDAR

This table shows that our approach is quite unique compared to the SOTA methods. For instance, no similar work is available using the same Ouster sensor, which makes the comparison of classification methods relatively difficult since these sensors produce singular results, depending on the core technology. Although we have been able to use the VLP32 in our experiments, it was not selected for the development of our inference model. The smoke test campaign was added since the results with fog conditions concluded in satisfying results, and we believed that another aerosol condition is likely to produce similar results. Conversely, our rain tests were not pursued into any classification algorithm. And, as described in section 5.2, our approach is designed with a specific frustum extraction and only concerns the distance information of points. Finally, using a Bayesian inference method, our classification objective is primarily to retrieve the visibility value.

Chapter 6

The impacts of weather conditions on automotive LiDAR

In this chapter, we evaluate the performances of LiDARs inside DVE conditions. The first section presents qualitative results obtained in an outdoor test campaign, in order to give insights on the different types of degradations possible with a standard automotive LiDAR.

Then, we present results drawn from two experiments done in climatic chambers to quantitatively assess the degradation of 3D-LiDARs data against DVE conditions, we describe them sequentially. Rain and fog tests, in secion 6.2, first showcase the different behaviors produced by LiDAR technologies, yielding interesting results that influenced the approach of later experiments. Lastly, the smoke experiments, presented in chapter 6.3, focuses on a unique LiDAR sensor and with a different context sensor (PSD sensor) Each section begins with a comprehensive description of the experimental setup, encompassing the selection of sensors and the establishment of test protocols involving targets and generated conditions. Then, the quantitative results are presented, employing an approach consistent with the methodology outlined in section 5. The realization may vary depending on the availability of sensors (context sensors and LiDARs) and the specific test protocol. Note that these results serve the dual purpose of quantifying the impacts of DVE on LiDARs and contributing to a deeper understanding of the data used for the inference model, presented in chapter 7.

6.1 Qualitative analysis in outdoor DVE conditions

This thesis started with an evaluation of the degradation induced by DVE scenarios on a vehicle in outdoor conditions, shown fig.6.1. During one week, data was recorded under various conditions. This section gives some visual insights to better understand how the point clouds can be altered. The proposed images are taken from point clouds measured by the VLP-32 LiDAR sensor.

Figure 6.1: EasyMile TractEasy on an industrial site.

One indirect effect of rain conditions is the creation of water puddles on the ground. Indeed, the laser signals shot by a LiDAR bounce on the puddle surface and potentially reach a solid target in the end, resulting in false detections under the ground. This mirror effect, shown fig.6.2a and illustrated fig.6.2b, results in points located under real targets. This can be a potential issue when estimating the ground height or when performing target detection and tracking.

(a) Point cloud screenshot

(b) Illustration

Figure 6.2: LiDAR mirror effect with wet ground.

But the estimation of the terrain height might not be straightforward because ground point clouds are highly corrupted in rainy and wet conditions due to scattering effects. Fig.6.3 shows the difference between the detections of a wet and a dry ground, where the point cloud quality is heavily degraded in the prior case. Instead, a dry ground is well detected by the LiDAR, as the detected layers are all continuous (fig.6.3b). This difference concerning the ground is very likely to be the same on all objects of the environment.

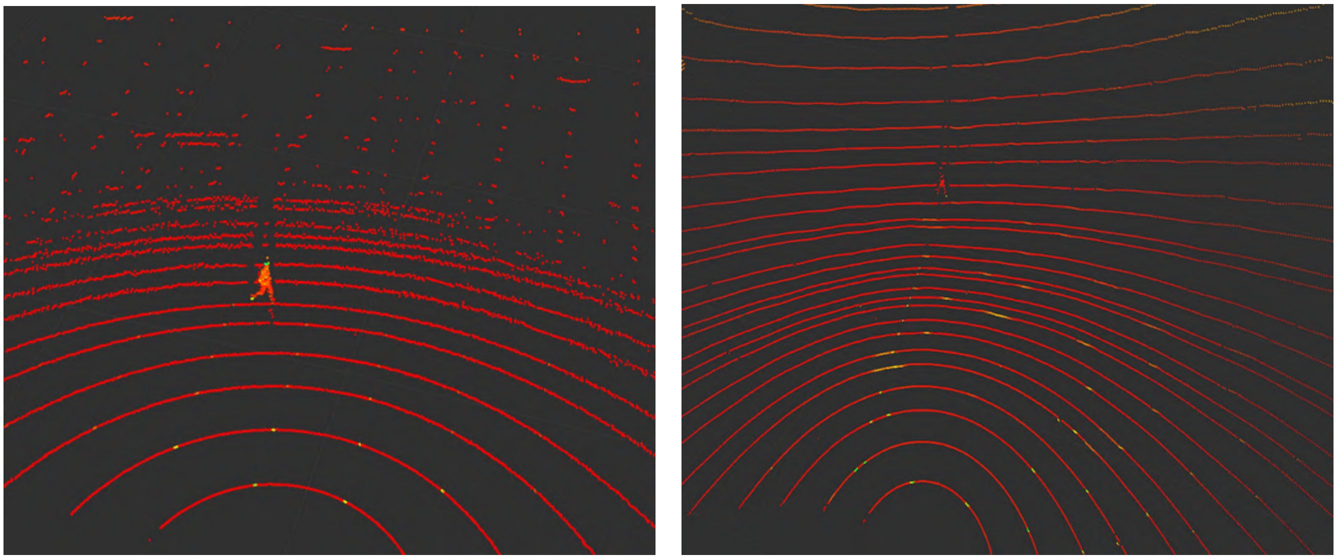

(a) Wet

(b) Dry

Figure 6.3: Comparison of ground detections between wet and dry conditions.

Finally, fig.6.4 shows different DVE scenarios where water particles produce false positive points. The difference between two scans obtained in the same scene but in clear (6.4a) and in heavy hail (6.4b) conditions is highlighted. Frozen rain falling around the sensor interacts with the emitted laser beams and backscatter enough energy to be detected. Considering the homogeneous repartition of the precipitation around the sensor and the quadratic decrease of the signal in the line of sight, more false positive points are expected close to the sensor, as illustrated on fig.2.31. One should note that this scenario corresponds to extreme conditions and this level of noise in the data is relatively rare. Figures 6.4c and 6.4d point out the potential weather-induced difficulty in differentiating targets and artifacts in point clouds. Wet road causes false positive points in the trail of passing cars. These water particles are detected by a LiDAR sensor, which potentially degrades the quality of the car and ground detection. And, high humidity conditions increase the odds of condensation (human breath, fog). Human breath producing fog clouds by condensation is also detectable, altering pedestrian detection.

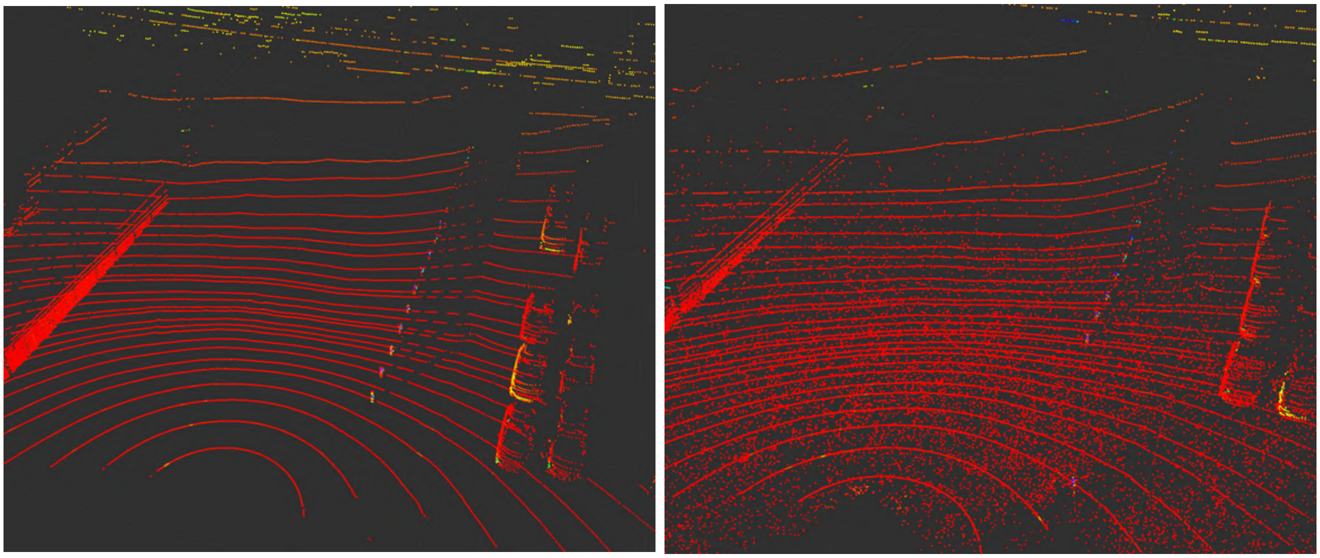

(a) Clear (b) Hail

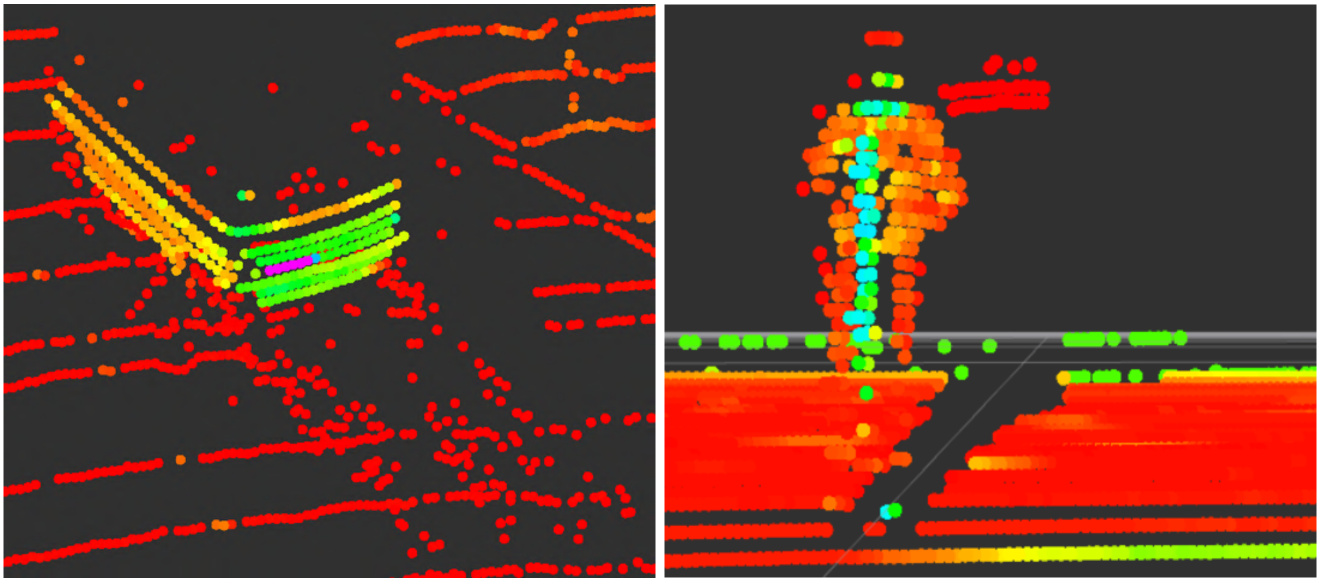

(c) Wet road (d) Condensed human breath

Figure 6.4: DVE scenarios cause false positive points in the point clouds.

This test campaign gave valuable insights concerning the impacts induced by natural DVE conditions (mostly rain) on LiDAR point clouds data and also on an AV's behavior. Although it was a good first experience, it did not lead to further developments since the collected data is deprived on any contextual information about the DVE conditions. In the end, it became evident that climatic chambers would be a solution to quantitatively assess the degradation as they provide controlled and repeatable conditions, with precise contextualization.

6.2 Experimental results from artificial rain & fog at CEREMA

This section describes the experimental setup and results of the rain and fog experiments performed in the CEREMA climatic chamber.

6.2.1 Experimental setup

The CEREMA[1] laboratory in Clermont Ferrand proposes a 30 m x 10 m climatic chamber, in which it is possible to produce artificial rain and fog under night or day conditions [181]. This chamber was originally designed to evaluate the impact of rain and fog on visual perception (with humans and cameras). Over time, numerous studies have been conducted using data collected within this chamber, as referenced for instance in [136, 133]. The top cover of the chamber can be removed so as to let the sunlight in, yielding daylight conditions, or set to emulate night conditions and light spots can illuminate the scene. The whole facility is composed of a tunnel where the weather conditions are produced and a control room to monitor the acquisitions. Rain and fog conditions are produced with a combination of nozzles disposed on the ceiling of the chamber at about 2.10 m heigh. Specific nozzles and pressure configurations are set to produce the desired precipitation intensity (in mm/h) or fog visibility (in m). The facility can produce two types of fog conditions using particle diameter modes centered at 1 μm and 10 μm. In the following, we will refer these two types of fog respectively as small fog and big fog.

Figure 6.5: Inside of the climatic chamber during the experiments, with targets denoted by labels.

LiDAR sensors

A total of five LiDAR sensors was gathered for the campaign, namely: Velodyne VLP-32, Ouster OS1-128, Livox Horizon, Cepton 860 and AEye 4SightM. Table 6.1 summarizes the main sensors' characteristics. Finally, fig.6.6 shows point cloud images produced by each sensor in the climatic chamber with clear conditions.

---

[1] ”Centre d’Etudes et d’expertise sur les Risques, l’Environnement, la Mobilité et l’Aménagement“, Clermont-Ferrand, France.

Sensor	**Type**	**Multi-echo**	λ **(in nm)**	**Max points in single scan**	**Intensity (in bit)**
Velodyne VLP-32	Spinning	2	905	35 k	8
Ouster OS1-128	Spinning	1	850	255 k	16
Livox Horizon	Risley prisms	2	905	25 k	8
Cepton 860	Micro motion	1	905	30 k	8
AEye 4SightM	MEMS	4	1550	22 k	16

Table 6.1: List of LiDARs used during the CEREMA experiments.

a. Velodyne VLP-32 [182]

This sensor is a 32 layers 905 nm laser spinning LiDAR with avalanche photodiodes (APD) receivers. It can produce two echoes from one single laser shot, their intensity is coded on 8 bits.

b. Ouster 128 [183]

Also spinning, it uses 850 nm VCSEL lasers and single photon avalanche photodiodes (SPAD) as detectors, also known as Geiger-mode APD. Compared to the VLP-32, this sensor has 128 layers of emitter detectors pairs, offering higher density point clouds. Only one echo can be returned from each laser shot, the intensity of which is coded on 16 bits.

Figures 6.6a and 6.6b show the point clouds of these two spinning LiDARs where the difference in number of vertical layers is observable. The sensors being static and the scanning pattern fixed, some areas are never scanned.

c. Livox Horizon [184]

Working at 905 nm, this LiDAR makes use of Risley prisms to steer its laser beams and generate non repetitive scan patterns. It produces two echoes with an 8 bit intensity. The scan pattern changes over time in a cyclic design so that, after a certain amount of time, the totality of the FoV is scanned [185]. Contrary to spinning LiDARs, over time, Risley prisms LiDARs scan the whole environment when they remain static. It has been observed that LiDARs based on Risley prisms used in military systems are more efficient in DVE like dust clouds for helicopter landings or foliage penetration [85, 35]. Figures 6.6e and 6.6f show the comparison between, respectively, a single scan of a Livox point cloud and the same scene but with a temporal accumulation of several seconds. The accumulation of Livox point cloud increases the coverage of the scene within the FoV. Unfortunately, although the Livox Horizon is capable of recording dual echoes, it was used in single echo mode during the experiments.

d. Cepton 860 [44]

Cepton sensors work with a unique technology based on micro motion technology (MMT), with 905 nm laser emissions. The sensor is composed of 24 channels (each channel is a laser-detector pair) in an optical module and the module as a whole is subject to frictionless oscillation to generate 3D images across a FoV. The model we use produces single returns. Figure 6.6c shows a Cepton 3D-point cloud colored by intensity. The 24 channels overlap on the edges to avoid empty scanned spaces, visible on fig.6.6c. Cepton sensors seem to have stable performances in range variations when it comes to adverse weather conditions, according to [135].

e. AEye 4SightM [76]

The AEye LiDAR uses a MEMS mirror to steer a 1550 nm laser beam and its detector is a Focal Plane Array (FPA). A FPA is an array of pixels performing the ToF measurements for each pixel, as opposed to single pixel architecture. This 1550 nm wavelength is supposed to be more efficient in adverse weather conditions due to the fact that more power can be emitted in the laser pulses while staying eye-safe [119, 186]. The sensor is capable of producing four echoes from a single laser emission. A study of MEMS mirrors for LiDARs is given in [187]. Unfortunately, a mistake was made in the acquisition procedure of the AEye 4SightM sensor and only the number of points is shown for this sensor.

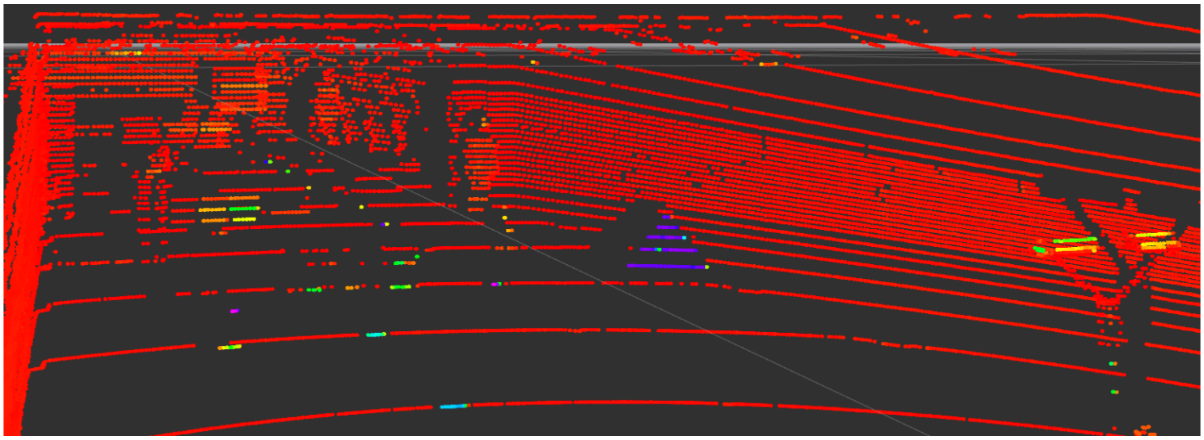

(a) Velodyne VLP-32, colored by intensity.

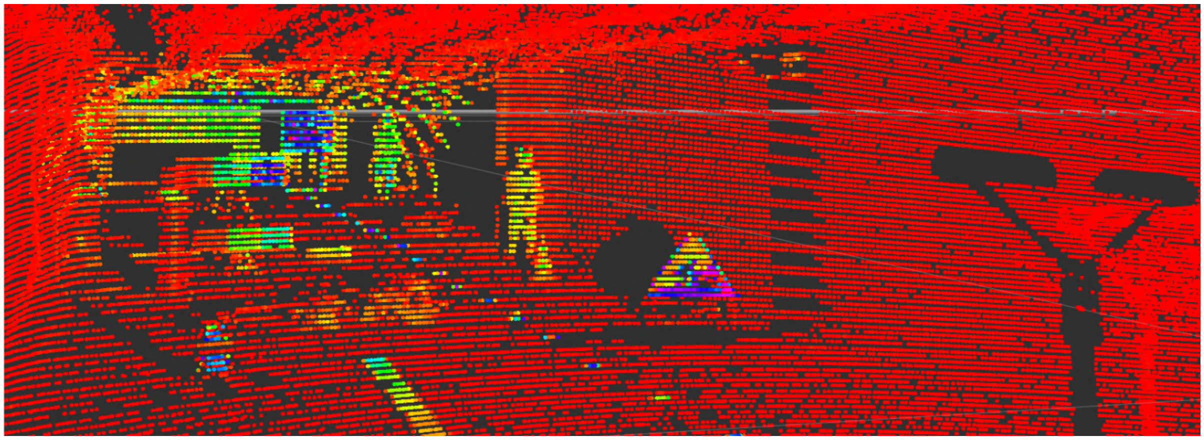

(b) Ouster OS1-128, colored by intensity.

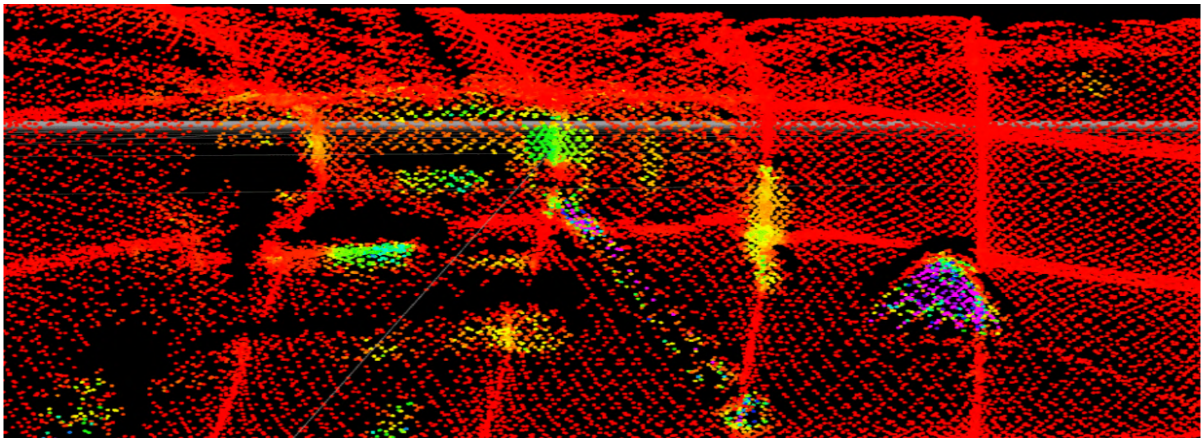

(c) Cepton 860, colored by intensity.

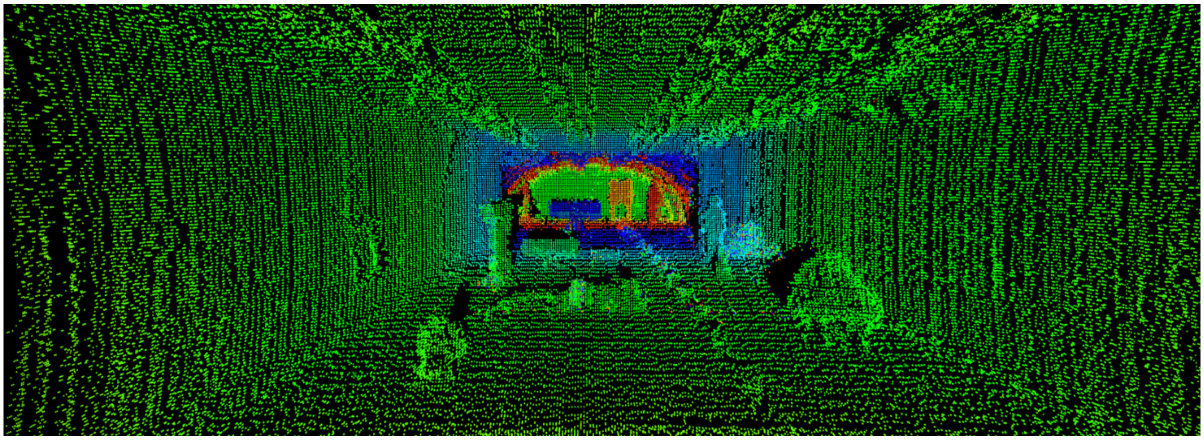

(d) AEye 4SightM, colored by depth.

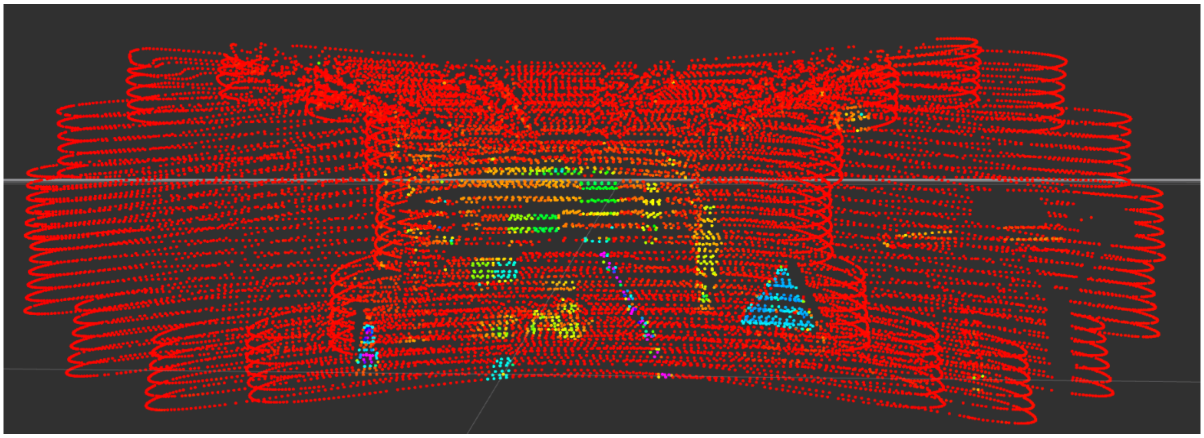

(e) Livox Horizon, single scan, colored by intensity.

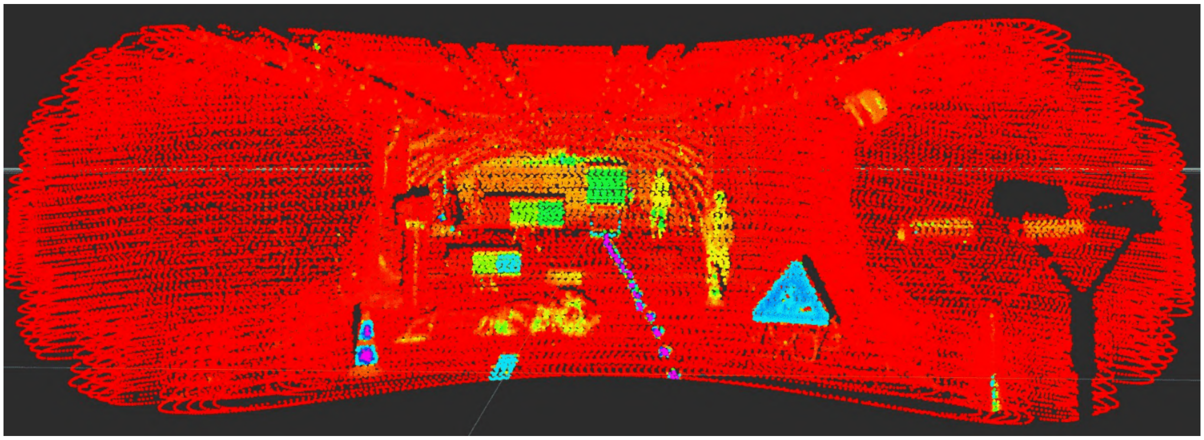

(f) Livox Horizon, accumulated scans, colored by intensity.

Figure 6.6: Captures of point clouds taken by each sensor.

Targets

Figure 6.5 shows the climatic chamber where different targets constitute the static scene used in the experiments. Some targets are typical of urban road environments, others are Lambertian calibrated targets. Table 6.2 describes them with their respective reflectance properties and distance to the sensors.

Type	Objects	Reflectivity	Distance (in m)	Label
Lambertian Surfaces (Flat squares)	1 m × 1 m	80%	23	*a1*
	50 cm × 50 cm	10%	11.3	*b1*
		50%		*b2*
		90%		*b3*
	30 cm × 30 cm	10%	17.3	*c1*
		50%		*c2*
		90%		*c3*
Road objects	Road sign	High	8	*r1*
	Boy dummy	unknown	12.5	*r2*
	Woman dummy	unknown	21	*r3*
	Road cones	High on stripes	6.5 and 10.6	*r4*
	Tire	Low	15.5	*r5*
	Concrete	Low	12.5	*r6*
	Lane	High	0 → 7	*r7*
	Beacons	High	0 → 23	*r8*
	Tree branch	unknown	8	*r9*
To optimize frustum detections	back of the room	Low	30	*back*

Table 6.2: List of targets used during the CEREMA experiments. In the thesis, the assessment of LiDARs performance is given using the a1 target.

Finally, the *back* target is added to the list, as shown on fig.6.7. In this climatic chamber, it is not possible to open the rear doors and reach the methodology described in chapter 5 (with a *free-sky* target). Therefore, we use this target because it is the furthest away in the chamber with a low reflectivity. This allows us to increase the number of frustum detections and thus increase the amount of data available for the weather classification algorithm.

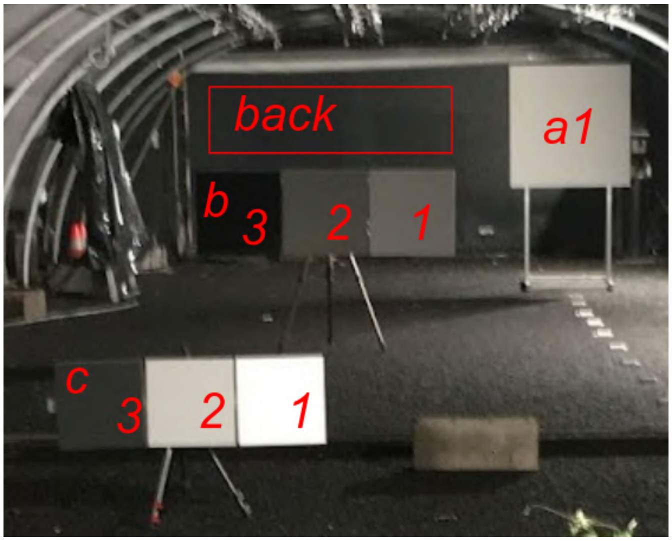

Figure 6.7: The *back* target is added to approach the weather classification methodology.

Weather sensors and DVE control

The evaluation of the conditions differs for fog and rain. For fog conditions, the chamber is equipped with a transmissiometer, a sensor that measures optical visibility at a specific wavelength, as presented in section 3.2 of chapter 3. On the other side, the intensity (or rate) of the rain is not measured but estimated based on the water flow rate inside the tubes and the chamber surface. In addition, rain production is done with a pump system located between the tunnel and the greenhouse, which allows for a more even distribution of water to the nozzles. Fog, on the other hand, is generated by a pump system located around the control room, which leads to a unidirectional distribution of fog in the nozzles from the control room to the end of the greenhouse. We found that the generation of fog leads to an inhomogeneous distribution within the chamber, but this is not considered in this work. Context sensors provide information on the characteristics of DVE. It is essential to have such instruments to gain precise information on the environmental conditions. For this purpose, we used a Parsivel disdrometer and the transmissiometer available in the climatic chamber.

The Parsivel OTT disdrometer [94] is capable of measuring the size, number and speed of water-droplets in rainy conditions. It is an active sensor that measures the extinction coefficients of light at a specific wavelength. Other properties of the precipitation can then be derived such as the intensity of the precipitation in mm/h, the Particle Size Distribution (PSD) or the visibility. This instrument has a minimal sensitivity of 0.2 mm in diameter, which means that it cannot detect particles below this threshold [188], especially fog droplets (around 0.001–0.02 mm). It uses a static emitter-receiver laser system. The only information available about rain in the climatic chamber is the precipitation rate: the disdrometer provides precise knowledge about the droplets diameter and speed. However, it is designed to measure natural rain, and our artificial conditions can mislead the sensor. For example, changing the nozzles' configuration to obtain certain precipitation intensities modifies the experimental conditions for the disdrometer if it remains at the same position. Also, considering the height of the nozzles, rain droplets are not likely to reach a stable state, as opposed to natural conditions [189]. Figure 6.8 shows a diameter and speed histogram of rain drops inside the climatic chamber under 120 mm/h of artificial rain rate. The transmissiometer used these experiments is directly operated by the CEREMA facility and we do not have many information about it. Finally, a passive visible camera is also used in our experiments. The sensor is not tuned for the tests, as its intrinsic parameters of gain and shutter speed are automatically set, but is useful for visually assessing the artificial DVE.

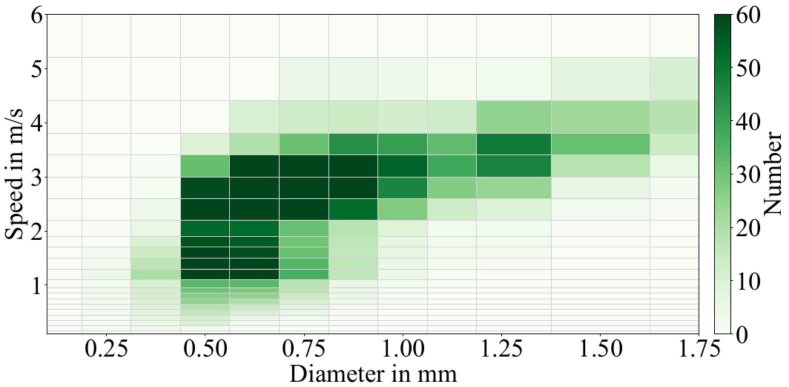

Figure 6.8: Example of a speed-diameter histogram taken with the OTT Parsivel disdrometer under 120 mm/h. Colored by the number of rain drops detected.

6.2.2 Test protocol

In accordance with the capabilities of the climate chamber, our test procedure, and our schedule, acquisitions were made under the following weather conditions:

- Clear conditions: recordings done before any weather condition is generated and dry scene.
- Rain rate (in mm/h): 20, 30, 40, 50, 60, 70, 80, 90 and 120.
- Fog visibility (in m): 10, 20, 30, 40, 50, 60, 70 and 80.
- Fog dissipation : small fog and big fog, from saturation to natural full dissipation.

The dissipation of fog is achieved by saturating the chamber with fog particles until the transmissiometer reaches its minimum value of 5 m. Once this value is reached, the production of particles is stopped. Data recordings start when the conditions within the chamber become stable and homogeneous, and continue until the particles are completely dispersed. This approach allows for a continuous evolution of the fog conditions, while also ensuring a more realistic behavior of the particles than for the stabilized levels of visibility.

Setting up multiple LiDAR systems looking at the same scene can cause direct or indirect cross-talk between the sensors and produce artifacts in the point clouds, thus leading to undesired measurements [120]. To avoid such phenomena and obtain the most consistent results, the sensors are powered one by one, so that they are recorded individually and do not influence each other. One minute acquisitions for each sensor and each weather condition are produced, except for the dissipation experiments. This complicates the acquisition process, but is worth the effort compared to biased and non-reproducible results. The LiDAR sensors are set at 10 Hz point cloud output frequency.

Considering the generated conditions, the smallest rain rates are first considered here, although it is possible to produce rates higher than 120 mm/h. As lower values of rain intensities are more likely to appear in natural conditions, their study is more valuable in the context of this work. The 120 mm/h precipitation value is generated to recreate an extreme rain scenario. We will see that the correlation between the rain rate and the nature of the rain is not straightforward, especially in such artificial conditions. Considering the dissipations, they are performed for each sensor individually.

6.2.3 Experimental results in clear conditions

In this section, the results of the one minute acquisition in clear conditions are shown. Extracted mean and standard deviation (std) values of the number of points on target and their intensities are shown in table 6.3. This serves as a baseline for comparison with the results obtained under DVE conditions. Some of the findings from the clear conditions can be utilized as normalization constants, aiding in the analysis and interpretation of the results obtained in DVE conditions and approaching the probabilistic approach introduced in chapter 5.

Sensor	Mean Number of Points	Std	Mean Intensity	Std
VLP-32	60.55	0.93	11.07	0.50
OS1-128	87.05	0.29	47482.61	1268.17
Livox Horizon	51.82	4.60	104.40	4.13
Cepton 860	109.08	3.31	1.51	0.27
AEye 4SightM	62.85	1.62	∅	∅

Table 6.3: Mean values and std of number of points on the *a1* target and associated intensities in clear conditions.

Compared to the other sensors, the Livox LiDAR presents the largest std and the lowest value of number of points. Taking into account its technology description, this behavior can be explained by its beam steering leading to non repetitive scan patterns on the target, and thus to variability in the number of detected points on the target.

The Cepton 860 also shows an irregular behavior in terms of number of points on target but stands with the highest count. The high point density comes from the fact that the target is placed on an overlap area of the sensor emitter/receiver channels, which increases the number of points, see fig.6.6c.

The Ouster sensor has the most stable number of points. This sensor is known to be functioning with digital technology, compared to an analog one (such as Velodyne LiDARs [38]). Its laser shots are triggered according to precise horizontal angles of the spinning part, thus always aiming at the same direction. On the contrary, the VLP-32 is triggered by a timer. This leads to potential variations in the heading of the lasers shots aiming at the target, thus inducing variations in the number of points. The difference in the number of points between the two spinning LiDARs is due to the number of vertical layers, as seen on figures 6.6a and 6.6b.

The AEye sensor shows a high number of points and also stable values, with its solid-state scanning method and FPA receiving technology.

Intensity values should not be compared directly between the sensors, as differences in resolution (shown table 6.1) and internal processing techniques provide very different results for each sensor. In the analysis of the impacts of DVE, all values are normalized by the values found in clear conditions (shown table 6.3). By doing so, the LiDARs can be compared to each other both in terms of the number of points and intensities.

The mean number of points on target can be considered a good approximation of the number of laser shots n_{shots} aiming at the target, as introduced in chapter 5. So, the mean number of points on target in clear conditions is used as constant for further normalization of the results found in DVE conditions. For the case of the number of points in the frustum, the normalization is also computed using n_{shots}. When the sensor produces multiple echoes (VLP-32, AEye 4SightM), the relation between the number of shots aiming at the target, the number of points on the target and the amount of noise in the frustum is not straightforward. A single shot can produce multiple echoes, explained in the multi-echo section, section 2.7 of chapter 2. As a result, the number of points in the frustum can reach up to $M \times ECHOES$, $ECHOES$ being the maximum number of echoes the sensor can produce. Finally, the intensity of the points measured under clear conditions can also be utilized with the same normalization approach. Normalizing the intensity of the frustum points by the mean intensity of the target found in clear conditions allows differentiation of the properties of solid objects versus noise points for potential intensity-based filters [63, 161].

6.2.4 Experimental results with Rain

Results from the rain experiments are presented here. These results are drawn from 1 minute stabilized levels of rain rate values.

First, a look at visual images captures by the visible camera is given. Next, speed and diameter distributions measured by the disdrometer give a look at the nature of the generated rain conditions. Then, results of target and noise detections are presented. The idea is to compare the output of different sensors facing similar rain conditions. Finally, the distance distributions of frustum points are displayed and analyzed. Although a multiple of targets are available in the CEREMA chamber during the experiments and to simply illustrate the concept here, only two are presented here. Target *a1* is used as to quantify the degradation of LiDAR data against DVE and target *back* is specifically used to assess the distance distributions of frustum points. The data are all displayed and organized based on the rain rate values given by the CEREMA.

Visual information

Images taken by the visible camera at each generated rain rate are presented in fig.6.9. The images allow us to witness the evolution of visual degradation along with increasing precipitation intensity in mm/h, as well as qualitative properties of the artificial rain (size of droplets, spatial distribution, spray profile)

– yet one should carefully interpret these images, as they have been acquired with camera automatic parameters.

As opposed to fig.6.5, taken in clear conditions, the light spots in rainy conditions cause glare effects, caused by the scattering of light by water droplets. Comparing fig.6.9a and 6.9b, one can see that the glare effect is stronger with the 20 mm/h rain rate than with 30 mm/h. Indeed, it is possible to see the end of the chamber at the latter while the glare effect from 20 mm/h makes the identification of objects behind targets *a1*, *a2* and *a3* impossible. This unexpected since an increasing rain rate would intuitively result in higher visual degradation. Different nozzle configurations are used for these two conditions. In the first one, the first row of nozzles out of three is used. It results in droplets of small sizes, suspended in the air. On the other hand, the second row of nozzles is used for the 30 mm/h rain rate. The droplets appear bigger and seem to have a higher falling speed.

Rain rates from 30 mm/h to 80 mm/h keep this last nozzle configuration but an increase in the density of the droplets inside the sprays is observable. Higher pressure and water flow rate must be set to obtain these higher precipitation intensities.

At 90 mm/h, the first and second rows of nozzles are activated. Visual degradation due to the glare effect is similar to results obtained under 20 mm/h rain rate although more droplets are visible.

At 120 mm/h, first and third rows are activated. Visual degradation is lower than at 90 mm/h where parts of the end of the chamber can be seen due to the lower glare effect. The third row of nozzles seems to be alimented with low water pressure, thus producing bigger droplets, similarly to the 30 mm/h rain generation. The first row is activated but its effect is less significant.

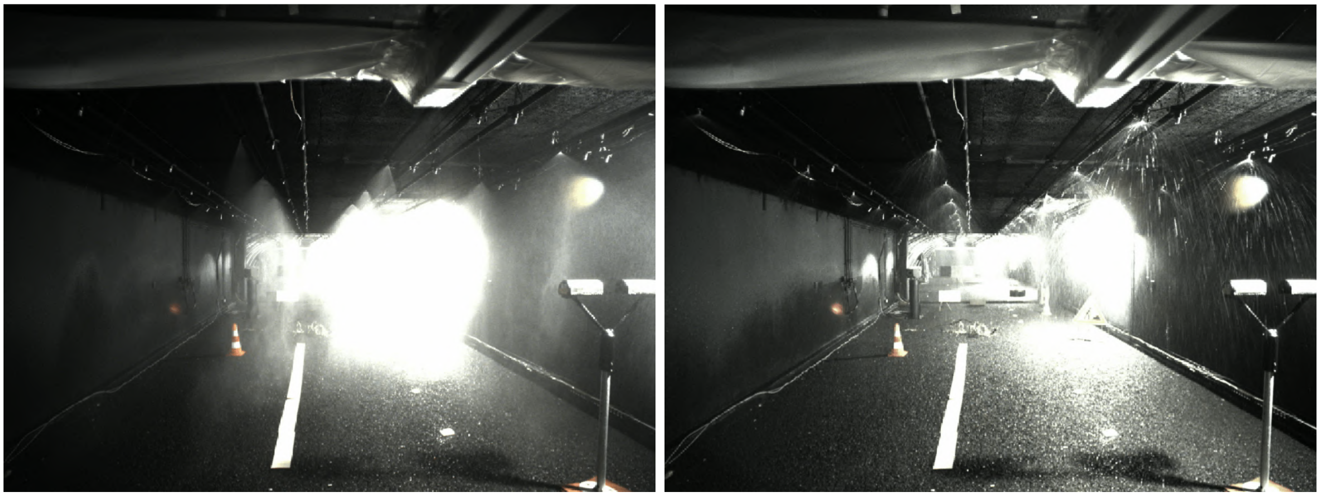

(a) 20 mm/h

(b) 30 mm/h

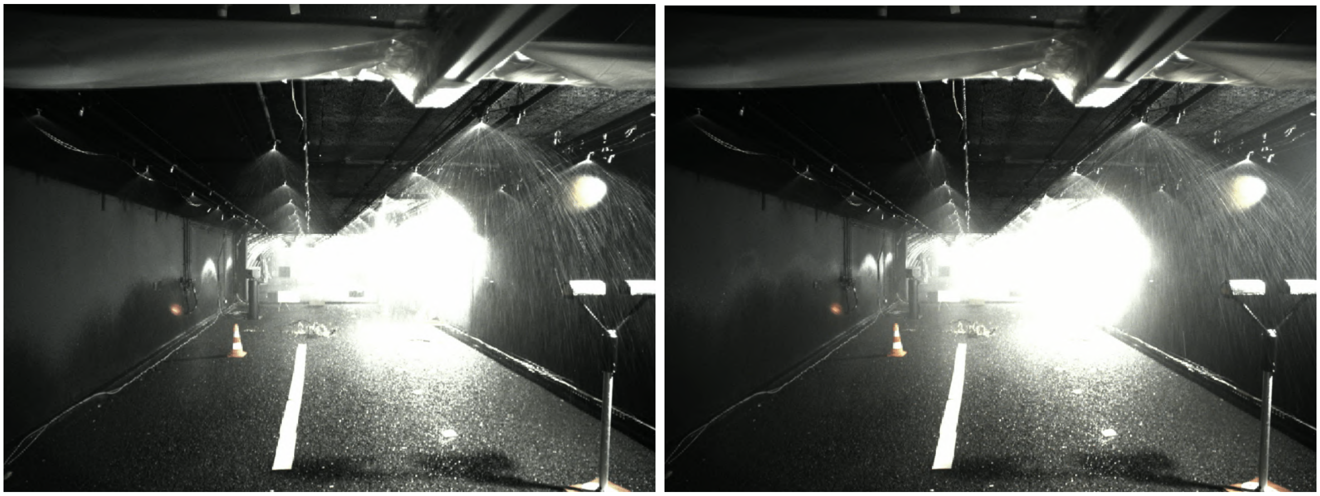

(c) 40 mm/h

(d) 50 mm/h

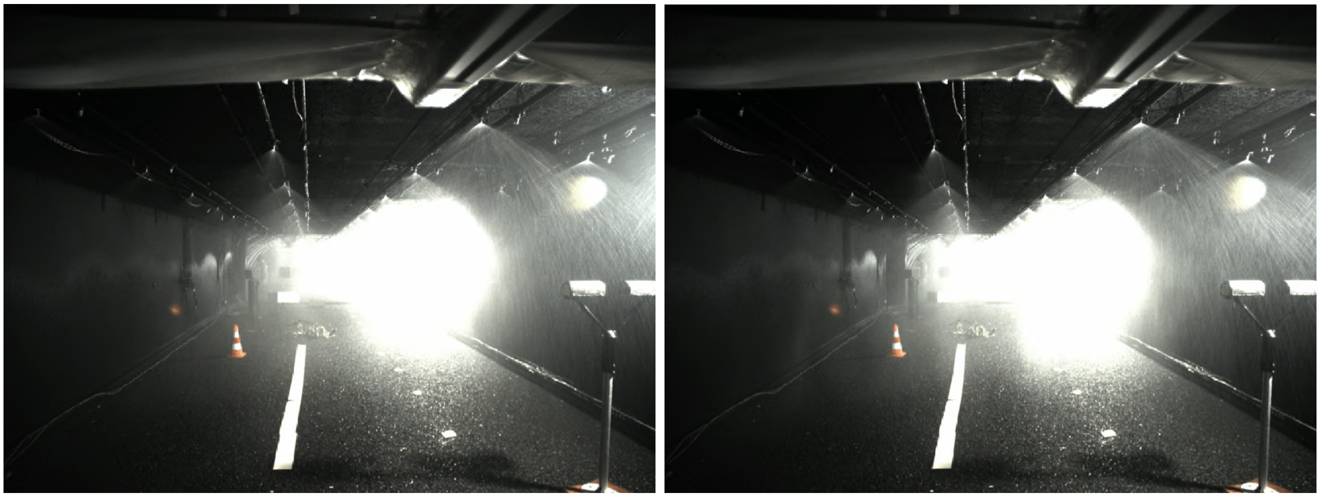

(e) 60 mm/h

(f) 70 mm/h

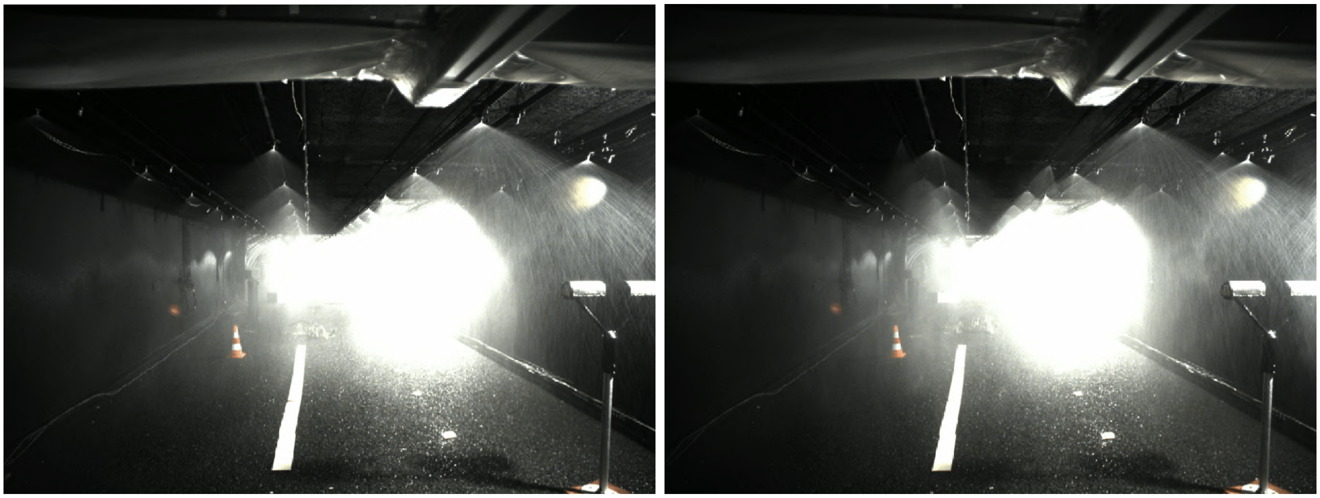

(g) 80 mm/h

(h) 90 mm/h

(i) 120 mm /h

Figure 6.9: Images taken during the rain experiments.

Speed-diameter histograms

The Parsivel OTT disdrometer measures the number and velocity of rain droplets. On fig.6.10, the accumulation of all measurements for each generated rain rate is presented, normalized by the number of measurements. The sensor was designed to give a result every 10 s. Each graph is a 2D histogram of the diameter and speed of droplets, colored by the number of drops.

Similarly to the previous visual results, differences are witnessed according to the precipitation intensity. For example, the histogram from 20 mm/h shows a concentration of small diameter (0.6 mm) and low speed (1 m/s) droplets. A tendency towards lower speeds is also witnessed. Histograms of 30 mm/h and 120 mm/h show singular properties. Indeed, in addition to having closed forms looking similar to the model of Atlas et al. [190], these precipitation values are the only one correctly estimated by the disdrometer (the sensor can output a measure of intensity but this is not deeply investigated in this work). From 30 mm/h to 90 mm/h, histograms tend to shrink to a more vertical form, meaning a more stable diameter value. These observations tend to argue that the rain generation is correct for rain rates of 30 mm/h and 120 mm/h but is affected by imperfections at other precipitation intensities.

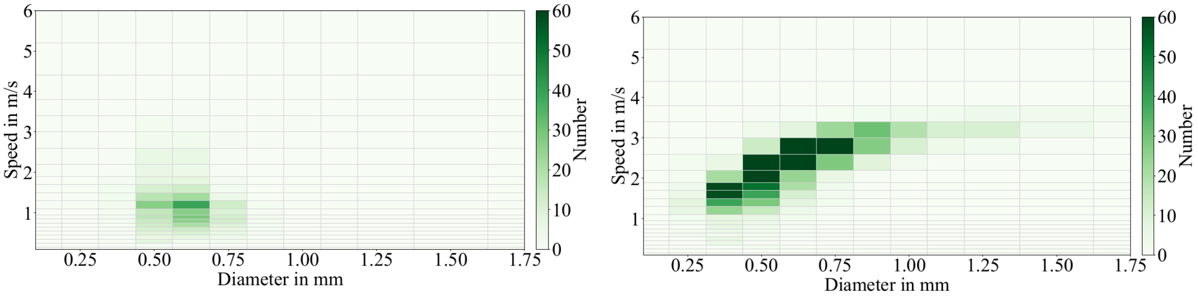

(a) 20 mm/h

(b) 30 mm/h

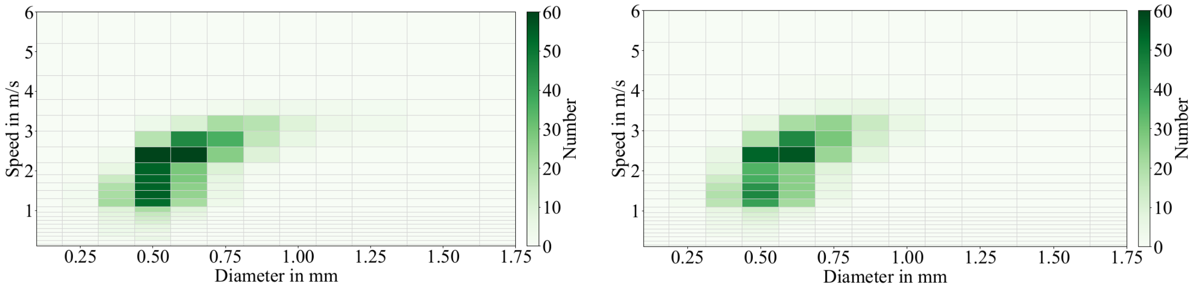

(c) 40 mm/h

(d) 50 mm/h

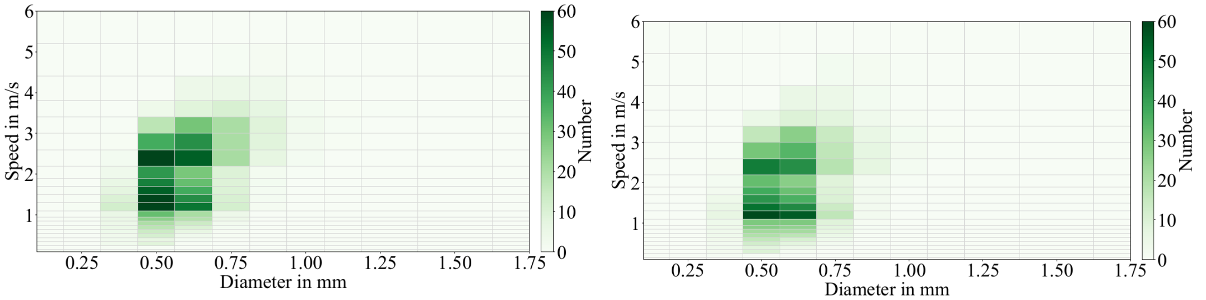

(e) 60 mm/h

(f) 70 mm/h

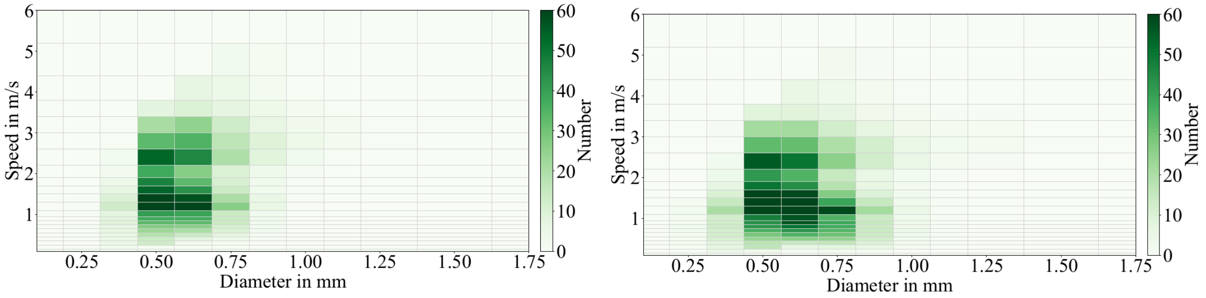

(g) 80 mm/h

(h) 90 mm/h

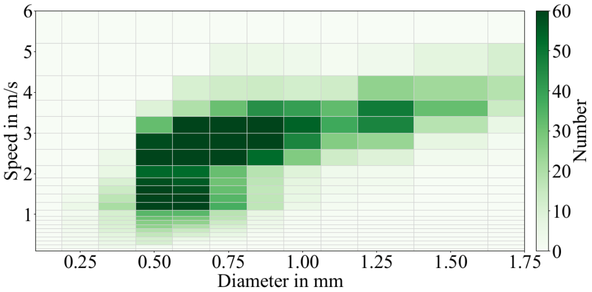

(i) 120 mm /h

Figure 6.10: Speed-diameter histograms from rain experiments.

Target detection and frustum

The impacts produced by rain droplets at different rain rates on the LiDAR sensors are presented in fig.6.11, with POD of both target and points inside the frustum. For each stabilized rain rate and sensor, a 1 min acquisition is acquired. The average values of the number of points and intensity +/− 1 std during these recorded minutes are displayed, normalized by values found in clear conditions. The abscissa is the increasing precipitation rate from 0 mm/h to 120 mm/h, with 0 mm/h being the mean values taken from clear conditions in table 6.3.

An unexpected phenomenon is first observable for the majority of the sensors as performances at 20 mm/h and 120 mm/h present contradictory behaviors with respect to the precipitation values. Indeed, as the precipitation rate increases, one can expect to observe more impacts on the sensors data. On the contrary, at 20 mm/h, almost every sensor (exceptions are described further) shows a reduced number of points on target compared to performances with higher values of precipitation. A comparable effect can also be seen on the intensities of the points which tend to decrease as well. On the other side, at 120 mm/h, while we should have the strongest impact on our data, we observe the opposite. For almost every sensor, performances increase significantly to better values compared to behaviors at lower precipitations. In a general manner, a loss of points on target is linked to a rise of frustum points. Similarly to target detection, the number of points inside the frustum is unexpected for the values of 20 mm/h and 120 mm/h.

These phenomena have to be investigated, they are likely to originate from the rain generation in the climatic chamber. For each sensor, the target and frustum detection behaviors are now presented in more details.

VLP-32 (dual echo): The sensor is capable of producing its nominal number of points on the target for rain precipitations of 30 mm/h to 50 mm/h but the target is almost lost (below 20%) for precipitation values of 70, 80 and 90 mm/h. The intensity of these points are all above 20% of the nominal value and have the tendency to slowly decrease from 80% at 30 mm/h to 30% at 90 mm/h. The 20 mm/h rain rate shows a depletion of the number of points to 80% and 40% of intensity. At 120 mm/h the sensor outputs the nominal amount of points but with lower intensities of 50%. The sensor produces a high number of noise points, with more than 100% of the nominal number of points on target except for 30 mm/h. Finally, points inside the frustum have lower intensity values going up to 20% of the clear weather reference on the target.

OS1-128 (single echo): The Ouster sensor behaves similarly to the VLP-32 but the loss of the target comes later in precipitation intensity, reaching the lowest value of 10% at only 90 mm/h. However it shows greater intensity values variations: a high peak of almost 80% can be observed at 30 mm/h but the intensity decrease is more progressive and reaches almost 0% by 60 mm/h to 90 mm/h. The 20 mm/h precipitation results in a high degradation on the number of points to 40% and intensities to 10%. At 120 mm/h the sensor outputs the nominal quantity of points but with lower intensities of 25% compared to nominal values in clear weather. Ouster's LiDAR outputs less points from the generated water particles. The number of frustum points reaches the maximum values of almost 40% for 20 mm/h and 90 mm/h. An increasing tendency for the number of frustum points can be observed, from 50 mm/h to 90 mm/h, while none of the points are detected at 120 mm/h. The intensity of the frustum points in rainy conditions is nearly 0 for all rain rates, except for 40 mm/h and 50 mm/h but too few points are detected at these conditions.

Livox Horizon (single echo): The Risley Prisms sensor shows intensity values similar to the Ouster sensor. However, its performances in terms of number of points are better. The sensor produces more than 80% of the nominal number for rain rate of 30 mm/h to 80 mm/h and at 90 mm/h, the target is still detected with almost 70% of points. Curiously, it is the only one who presents lower performances at 120 mm/h, with around 30% of points, being its second lowest value after 20mm/h, although these points have higher intensities. The Risley prisms sensor has a peak of frustum points at 20 mm/h of almost 100% of the nominal value. A stable stage is then observed between 30 mm/h to 80 mm/h at around 15%. For the final values of precipitation, the number of points rises to 70%.

Cepton 860 (single echo): Cepton sensor yields good performances at 20 mm/h with almost 100% of points detected on the target, but these points have low intensities of 30%. This high quality of detection is maintained until 60 mm/h with around 80% of points. Then, from 70 mm/h to 90 mm/h, it produces less than 30% of points or even values close to 0%. At 120 mm/h, Cepton's sensor returns its nominal number of points again. This sensor shows similar behavior to the previous sensors in terms of intensity. This sensor has a very low level of noise with less than 5% of its nominal value for all precipitation rates. As a result, the observed intensities for the frustum points are considered as outliers.

AEye 4SightM (4 echoes): The target is detected with 100% of points for precipitation intensities from

20 mm/h to 80 mm/h and 120 mm/h. The performance only decreases to a value of 80% at 90 mm/h. The four echoes capability of this LiDAR brings a large amount of noise compared to the other sensors. The number of noise points vary depending on the precipitation rate, with a peak at 20 mm/h and a value of 250%. We then see a rise from 30 mm/h to 70 mm/h with 50% to 150%. The frustum points finally slowly decrease in number to reach more than 100%.

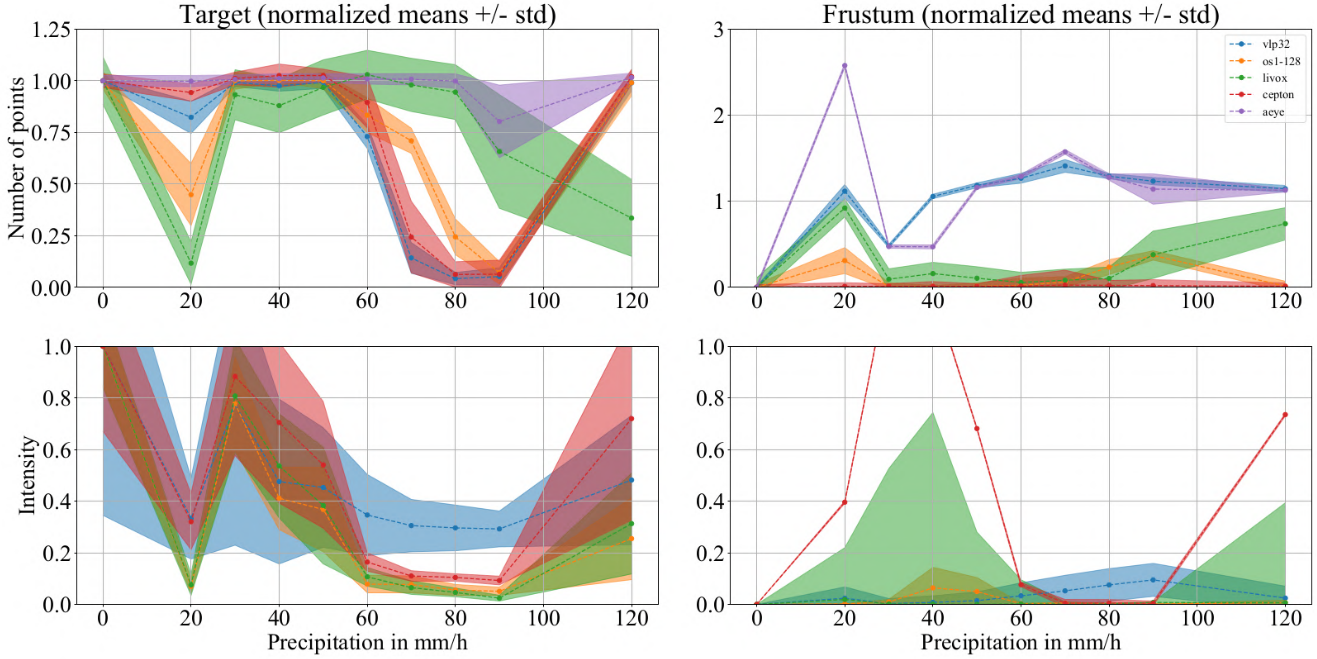

Figure 6.11: Sensors behaviors under rainy conditions.

The studied LiDARs show unexpected behaviors with regards to the edge rain rates of 20 mm/h and 120 mm/h. A few exceptions remain with the Livox, Cepton and AEye sensors. The former is impacted at 20 mm/h but also presents a reduced number of points at 120 mm/h along with higher intensity. On the contrary, Cepton's LiDAR does not show any reduction of number of points on target at the first rain rate. Finally, the AEye sensor is not concerned with these edge cases as its target detection performances are stable across all rain rates, although its frustum detections are affected.

The reason behind these unexpected phenomena comes from the artificial rain generation. To reach a specific rain rate, a configuration between nozzles properties and water flow is used. The value in mm/h is computed accordingly using the flow rate and the room's surface area, which results in various rain characteristics (size of droplets, spatial distribution, spray profile). Impacts of visual passive cameras and near-infrared LiDARs are comparable, observations made from images in the visible spectrum can be applied to impacts on LiDAR data. The images of fig.6.9 show different levels of degradation in the understanding of the scene. Speed-diameter histograms of fig.6.10 add valuable information concerning the properties of the rain because the number of particles and their speed alter the chances of interactions with the laser shots. At 20 mm/h, both the LiDAR and image data are highly degraded. The generation of this rain rate induces a large amount of small droplets in suspension, highly affecting both LiDAR and camera sensors. The disdrometer does not show a high number of droplets but the limitations of the sensor should be considered as water particles may be too small to be detected, being similar to fog particles.

Detections in the frustum at 20 mm/h show high numbers of noise points and sensors with multi-echo capability present the highest counts. This is in agreement with the previous observation of high number of particles at this rain rate, although not detected by the disdrometer. At 30 mm/h and 120 mm/h, while the OTT rain sensor is showing similarities with natural rain conditions and the majority of the studied

LiDARs are not impacted in terms of target detection. Detections in the frustum show a low count of points for the prior rain rate but a relatively high one for the latter. So, it seems that artificial rain close to natural conditions does not degrade LiDAR data as much as less realistic rains and, in this case of near-natural rain, a higher value of mm/h means higher amounts of frustum points. When precipitation grows from 40 mm/h to 90 mm/h, the number of droplets increases but their size converges towards 0.5 mm. The resulting impacts on the LiDARs are globally intensified, although the impacts on the frustum points are not clear for all sensors. To strengthen the analysis and have a better understanding of the results at 120 mm/h, rain rates of 100 mm/h and 110 mm/h should be generated and analyzed.

The LiDAR sensor yielding the best performances in terms of target detection in rainy conditions is the AEye 4SightM. Indeed, almost no reduction of the number of points on target is observable across all rain rates. The higher power of emission available at the 1550 nm wavelength (while staying eye-safe) seems to allow better penetration through raindrops and thus better detection capabilities. Its four echoes capability could also be an asset in this case. Nevertheless, it is not possible to determine whether the MEMS and FPA technology of this LiDAR is more efficient in DVE and it makes it difficult to compare this sensor to the others working at another wavelength. We now compare the other sensors working at the common 905 nm. Both VLP-32 and OS1-128 present a similar behavior in terms of number of points on target but Ouster LiDAR stands with better results. Since the values are normalized, their different number of vertical layers (32 vs. 128) is not sufficient to explain this difference, supposing all laser shots are independent to each others. Cepton sensor also has a similar behavior to the VLP-32 with the exception of the 20 mm/h described above where it stands with better results. Livox sensor shows more stable performances facing rising rain rate. It is able to detect the target with more than 60% of the nominal number of points at rain rates of 80 and 90 mm/h where the other sensors fail.

As expected, the only sensors with multi-echo information (VLP-32 and AEye 4SightM) have higher amounts of noise points. At 20 mm/h, the AEye sensor shows the highest density with more than twice as many points compared to the other sensors. The Velodyne sensor has the second highest count but remains close to the other sensors at this precipitation value. Apart from this rain rate, the two sensors show a similar amount of points coming from the generated raindrops. Although they do not have the same maximum number of echoes (2 for the VLP-32 and 4 for the AEye 4SightM), their behavior shows a convergence to a stable value around 100%. Comparing the two spinning LiDARs with dual or single echo mode, multiple echoes do not seem to offer better target detection but potentially provides additional information about the environment (located in the frustum detections). The AEye LiDAR, with its four echoes capability, provides a high level of points coming from raindrops in the frustum while having the best detection performances on the target.

In terms of intensity, the results given by our LiDARs show, for most rain rates, a clear differentiation between a real target and noise points. When the target is hardly detectable, the intensity of its points becomes very low as well and it could be considered as noise if an intensity-based filter is considered. Sensors seem to follow a common pattern in the intensity of the target points, which varies depending on the rain rate. It allows observation of degradations of the received signals, which is additional information, especially when target detection is optimal.

Frustum distance distributions

Although the analysis of the distance distribution of frustum points under foggy conditions revealed good findings (especially for the ouster sensor, see below), leading to the development of the inference model in this thesis, it is not the case for rain. Fig.6.12 shows the distance distributions of frustum points shooting at the *a1* target, with each sensor and generated rain rate. The interpretation of the results is not straightforward. The amplitude of the distributions (number of noise detections) is in accordance with the ones of Fig.6.11, with for example. Expect for the Cepton sensor, a multitude of peaks can be observed in the distributions, their positions most likely indicate the position of the rain nozzles. However, these positions differ for each sensor. In general, the visual interpretation of the histograms does not indicate a continuous evolution or pattern with regards to the evolving rain rates. Therefore, the erratic nature of the data holds us back in pursuing further development concerning the rain experiments.

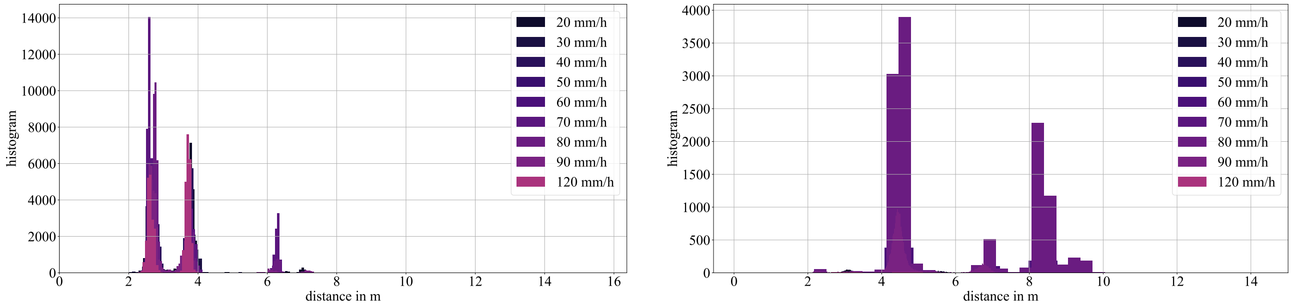

(a) VLP-32

(b) OS1-128

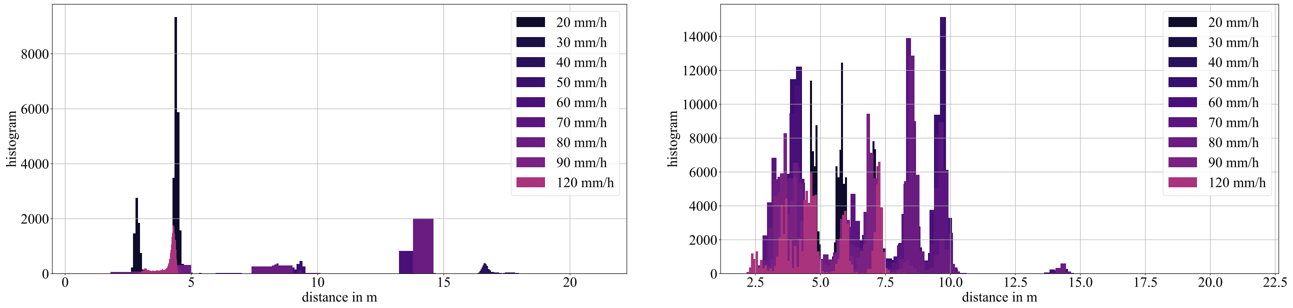

(c) Livox Horizon

(d) AEye 4SightM

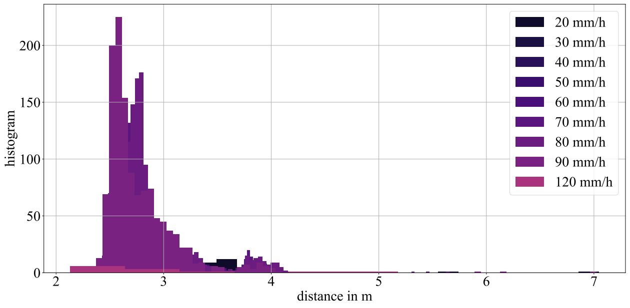

(e) Cepton Vista

Figure 6.12: Distance distribution of frustum points shooting at *a1* target for each sensor and accumulated over all generated rain rates.

To conclude, the rain experiments did not yield positive results. The generated rain conditions deviated significantly from natural states, leading to bias in the correlation between sensor degradation and rain rate. As a result, the rain rate may not the most suitable metric for our approach, as the size of raindrops has a more significant impact. Nozzles positioned too low prevent the drops to reach their terminal velocity, and the spatial distribution of the rain leads to localized showers in the point clouds. The resulting data are altered by this phenomena, particularly affecting the noise distance distributions, rendering them unsuitable for utilization in an inference model.

6.2.5 Experimental results in fog conditions (stabilized visibility levels)

Results from the fog experiments are presented here, where stabilized values of visibility are set. The type of fog used for this experiment is the big fog and the data is drawn from 1 minute recordings for each level of visibility. First, a look at visual images captures by the visible camera is given. Then, results of target and noise detections are presented. The idea is to compare the output of different sensors facing similar fog conditions. Finally, the distance distributions of frustum points are displayed and analyzed. The data are all displayed and organized based on the visibility values measured by the transmissiometer.

Visual information

Fig.6.13 shows images taken during fog experiments at visibilities 10 and 80 m. As opposed to rain conditions, the visual degradation of the generated conditions is intuitive regarding the visibility values : lower visibility values mean increased degraded on the images and inversely.

(a) Visibility of 10 m

(b) Visibility of 80 m

Figure 6.13: Screenshots taken from fog experiments.

Target and frustum detection

Fig.6.14 summarizes the results for each sensor. In contrast to the rain impacts, getting stronger with increasing mm/h (with the exception of the rain generation problems discussed previously), increasing the fog visibility here leads to a decrease in the number of fog particles. As a result, fewer impacts on the sensors data are expected as visibility increases. This can be observed in fig.6.14. Contrary to the results obtained in rainy conditions, there is no unexpected result to be reported in this case.

VLP-32 (dual echo): The Velodyne sensor starts to produce points on target at 50 m visibility, the number of points then rises to reach around 75% at 80 m of visibility. Its intensities have the most surprising values: starting from a peak between 50% to 80%, the values afterwards seem to stabilize around 50% as visibility increases. The Velodyne spinning sensor shows a quasi constant level of noise in fog conditions, at 125% of the nominal number of points on target. A curious peak up to 160% at 60 m of visibility is seen. The intensity of these points starts from 30% and goes to a reduced value of 10%.

OS1-128 (single echo): This sensor starts to detect the target at 40 m of visibility. The number of points then rises to the maximum value of points on target for a 80 m visibility. Intensities have increasing values from 0% to around 20% at 80 m visibility. This sensor presents a decreasing number of noise points from almost 70% to 0% for increasing visibility. At all visibilities, the frustum points coming from fog stand with an intensity close to 0.

Livox Horizon (single echo): First points appear at 30 m visibility and 100% is reached at 50 m visibility. Intensities have a rising behavior up to 20% as visibility gets to 80 m. The Livox sensor keeps a low number of frustum points at around 10% and shows a slight increase for the last visibility values reaching 30% at 80 m. The intensity of these points is below 5%.

Cepton 860 (single echo): The Cepton sensor produces its first points on target at 50 m of visibility. This number then increases almost linearly to reach 100% at 80 m of visibility. Similarly to rainy conditions, this sensor has the same performances than the OS1-128 and Livox sensors in terms of intensities. The Cepton sensor shows a very reduced number of points in the frustum in fog conditions, almost 0 for all visibility values. The intensity of these points is close to 0 but shows a slight increase for the latter visibilities.

AEye 4SightM (4 echoes): The MEMS sensor starts to detect the target at 30 m visibility. The number of points rises quickly and reaches 60% at 40 m visibility and 100% at 50 m. It also shows an increasing number of frustum points with increasing visibility. Starting from a lowest value of 25% to quickly reach around 200% at 30 m and keeping this high level of frustum points for the latter visibilities.

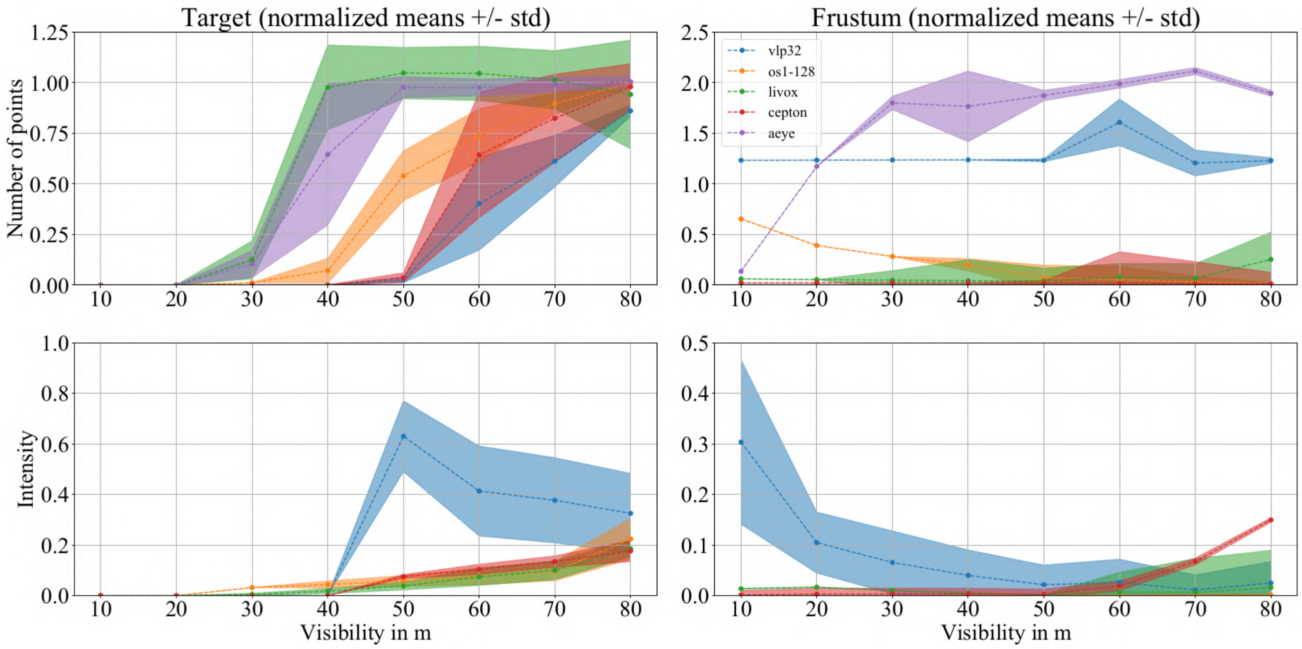

Figure 6.14: Sensors behaviors under foggy conditions.

In a general manner, target detection behaviors under fog conditions are alike for all sensors. When visibility is too low, the target is not detected: the laser signals are highly scattered inside the fog clouds and never reach the detectors (back from the target) or even the target itself. But, they are reflected by the fog particles and create detections in the frustums. Then, as fog dissipates and visibility rises, LiDARs receive more signal leading to better target detection. Finally, all the studied sensors tend to nominal target detection performances but with different results at lower visibility values.

The Livox sensor stands with the best performances in terms of target detection and the AEye LiDAR comes second because at 40 m of visibility, the latter sensor shows 60% of points on target while the Livox is already performing at 100%. For higher visibilities, the two sensors yield nominal target detection performances. Once again, the two spinning sensors have very comparable results as their number of points is rather similar. One could say that the only difference is the starting visibility value. The behavior of Cepton's sensor is very similar to the VLP-32, as both of them show more degradation in fog conditions.

The number of noise points in fog conditions is different than with rain experiments. The Livox sensor presents a low number of points as if it is poorly sensible to fog particles. On the contrary, the sensor can be highly affected by the generated rain detecting up to almost 100% of noise points in the frustum at 20 mm/h. A rise of these frustum points for the latter visibility values is however noticed. This is highly unexpected because the results drawn from the rain experiments showed that the 20 mm/h rain rate was similar to fog conditions. It might be explained by better automatic filtering of the backscattered signals in dense fog than in non homogeneous ones (for the 20 mm/h rain rate). The Cepton sensor does not show any noise points in fog, similarly to rain conditions.

The Velodyne sensor shows an unexpected behavior in terms of intensity for the target points with a decreasing tendency as visibility is increased. However, the intensity of its frustum detections seems to be more logical as denser fog induces higher intensity and inversely. When fog is dense, the high density of particles must act similarly to a solid object and backscatter a lot of laser light, thus increasing the intensity of the frustum points. On the other hand, multiple scattering effects that dissipate the signals must be prevalent in lighter fog situations and reduce the returning intensities. The other sensors show the opposite. As fog dissipates, target detection is improved and the intensity of the target points rises as well for the Livox, Ouster and Cepton LiDARs. However, intensities of the frustum points remain very low for all visibility values, except for the Cepton sensor with the rising effect described previously.

The multi-echo capability here implies again a higher number of points but not necessarily a better analysis of the environment. For example, the Ouster sensor shows a continuous decrease of its frustum points, which could easily potentially to visibility classification. On the contrary, the number of frustum points of the VLP-32 is rather constant over all visibility values. The AEye sensor and its four echoes presents a growing number of noise points. It rises from almost no noise points at first to a stable value of 200% of noise points for visibility values from 30 m to 80 m. Similarly to rain conditions, the two sensors with multi-echo capability show a stable value of frustum points for the majority of visibility values. The evaluation of the frustum points for LiDARs with multi-echo capability does not clearly provide information on the fog conditions. To go deeper, one should look at the labelization of the echoes inside these frustum points. This is not investigated here but rather in the chapter 8 as an opening on future work.

Frustum distance distributions

Fig.6.15 shows, for each sensor, the distance distribution of frustum detections shooting at the *a1* target and for each stabilized visibility value. Results of each sensor greatly differ. The VLP-32 (with dual echo) exhibits distributions with significant amplitudes near a 2 m distance, with slightly understandable change in shape based on visibility values. There are also a few points detected at greater distances (around 6 m and 8 m), but they appear insignificant compared to the others. The Ouster sensor presents a high amplitude peak located very close to the sensor and two other distributions at 3 m and 5 m. The Livox Horizon LiDAR first shows a minimal distance of detection below 1 m which potentially prevents the full data recovery. Then, the results are quite similar than the ones from the VLP-32, but the evolution of the distributions is less consistent. AEye Lidar presents a multitude of peaks up to 12 m. The sensor from Cepton, on the other side, shows characteristic distributions and peaks as well but between 2 m and 4 m.

The multiple peaks found in most of the distributions seem to be related with the position of the fog nozzles, just like with the rain conditions. Indeed, the generation of stabilized visibility values is done by injecting fog clouds with the nozzles to adjust the visibility value measured by the transmissiometer.

Positions of the water nozzles should impact the resulting data because the addition of fog particles under the nozzles increase the odds of interaction with the laser beams locally. Additionally, changes in the distributions caused by the different visibility values are not intuitive.

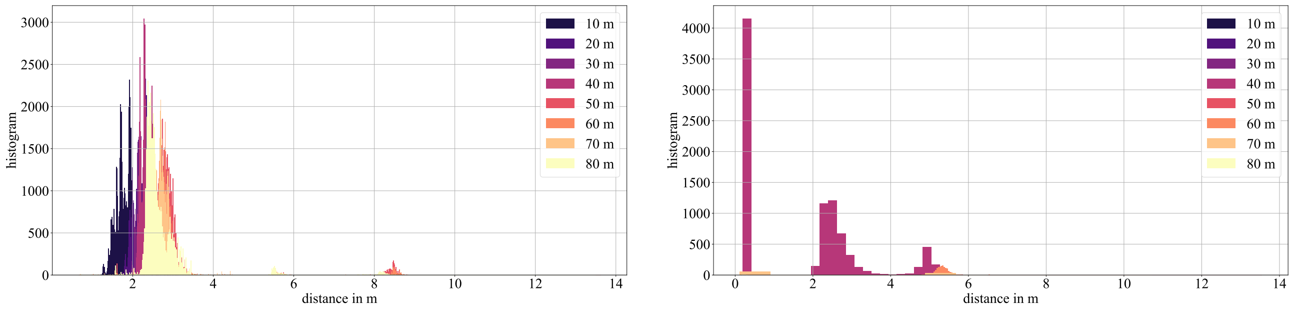

(a) VLP-32

(b) OS1-128

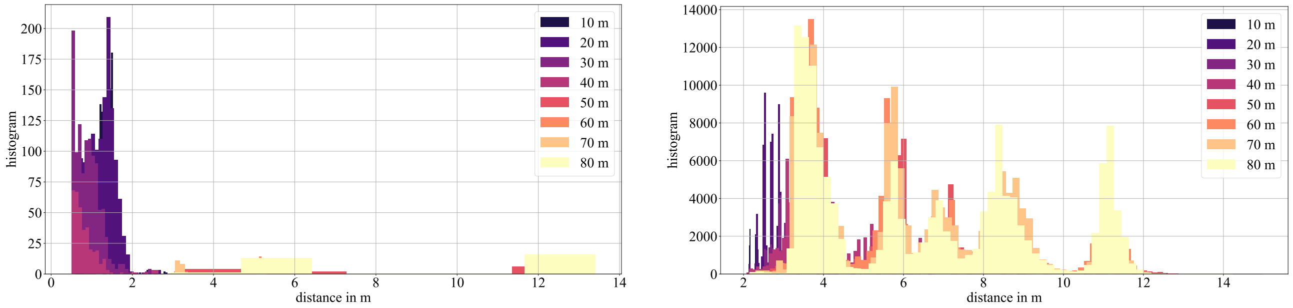

(c) Livox Horizon

(d) AEye 4SightM

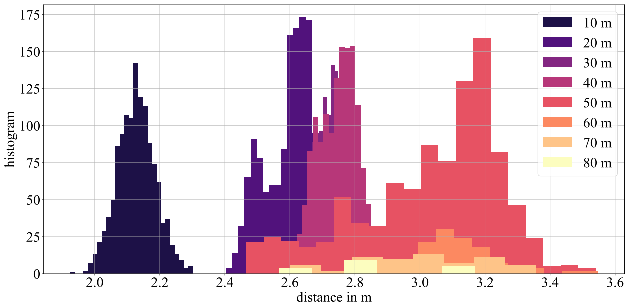

(e) Cepton Vista

Figure 6.15: Distance distribution of frustum points shooting at *a1* target for each sensor and accumulated over all generated visibility values (represented by the color).

Although slightly more promising than with rain, the LiDAR data generated in stabilized fog levels is still not satisfying enough, we choose to not go any further in their analysis. However, the more natural aspect of fog during the dissipation experiments leads to better results.

The histograms displayed on fig.6.15b, which represent the data captured by the Ouster sensor, lack visual consistency. However, upon closer examination of the first peak, a different conclusion can be drawn. Fig.6.16 illustrates the same data, but with the application of a filter that only includes points located below a distance of 1 meter. This time, very distinct distributions are seen and it appears that the varying levels of visibility values have a very gradual impact.

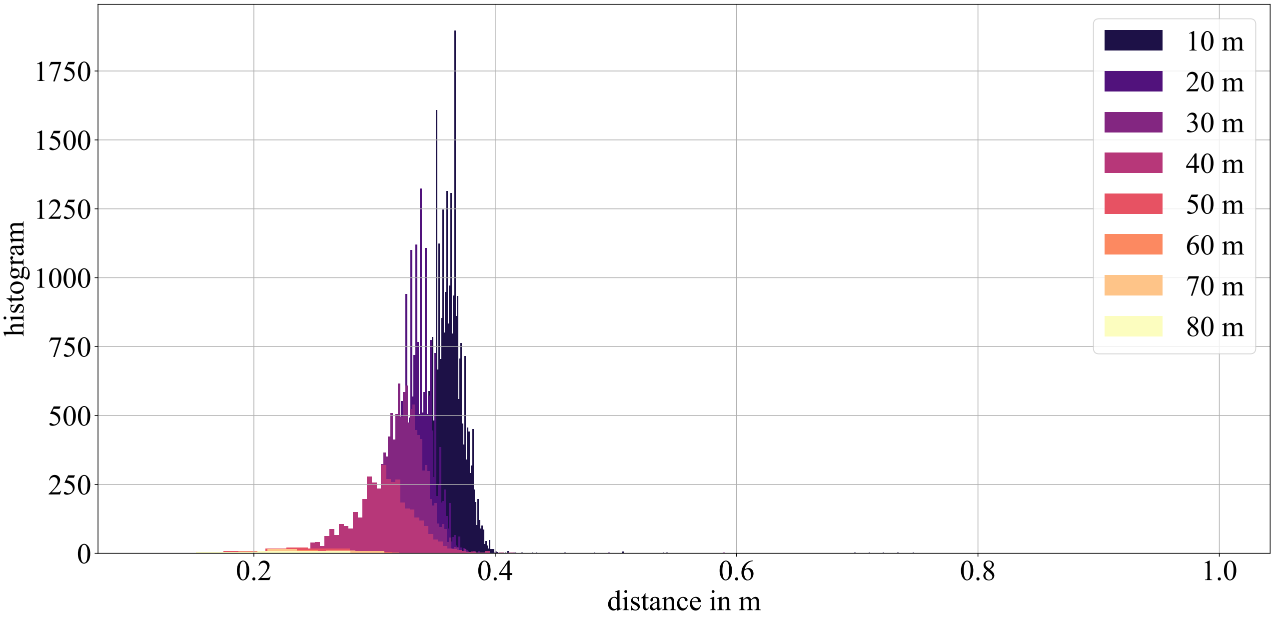

Figure 6.16: Distance distribution of the Ouster frustum points in fog conditions shooting at the *a1* target, filtered with a 1 m maximum distance.

In the end, it became evident that the Ouster sensor is the only one exhibiting data suitable for utilization in an inference method. This discovery was a turning point, motivating us to dig deeper into its analysis because such gradual evolution of statistics can lead to a potential pattern recognition algorithm. These distributions are particularly relevant as they have relatively large amplitudes (hence conducive to calculate probabilities). Furthermore, the frustum points are located very close to the sensor, making the classification easier for potential operational scenarios. Compared to the other sensors (fig.6.15), points located in very close range (at 0.3 m) are more likely to be used for algorithms than at higher distance where real objects can be present. Consequently, it was decided to not go any further with the other sensors and instead focus on the Ouster LiDAR. Hence, for the following of this thesis, only Ouster data and results are presented.

6.2.6 Experimental results with fog dissipations

The fog dissipation experiments consist in the saturation of the chamber with fog particles. Nozzles inject fog clouds until visibility reaches the 5 m lowest value measurable by the transmissiometer and conditions become homogeneous. Then, data is recorded until the fog completely dissipates. As such, this experiment captures a significant and continuous evolution of data over time while potentially proposing more realistic fog conditions (at least in terms of spatial homogeneity). Finally, the dissipation experiments are conducted using both small (1 μ m diameter mode) and big (10 μ m diameter mode) fog conditions and results are therefore presented for each of them. This enables the acquisition of more data and the potential comparison of the two types of fog.

Additionally, maximizing the amount of points in the distributions leads to increasing the amount of data, later used in the inference model (described chapter 7). Consequently, only the *back* target is used in the following. This target is the furthest possible in the chamber and also have low reflectance. As explained in the presentation of the approach in chapter 5, using such target increases the odds of detecting points from DVE particles.

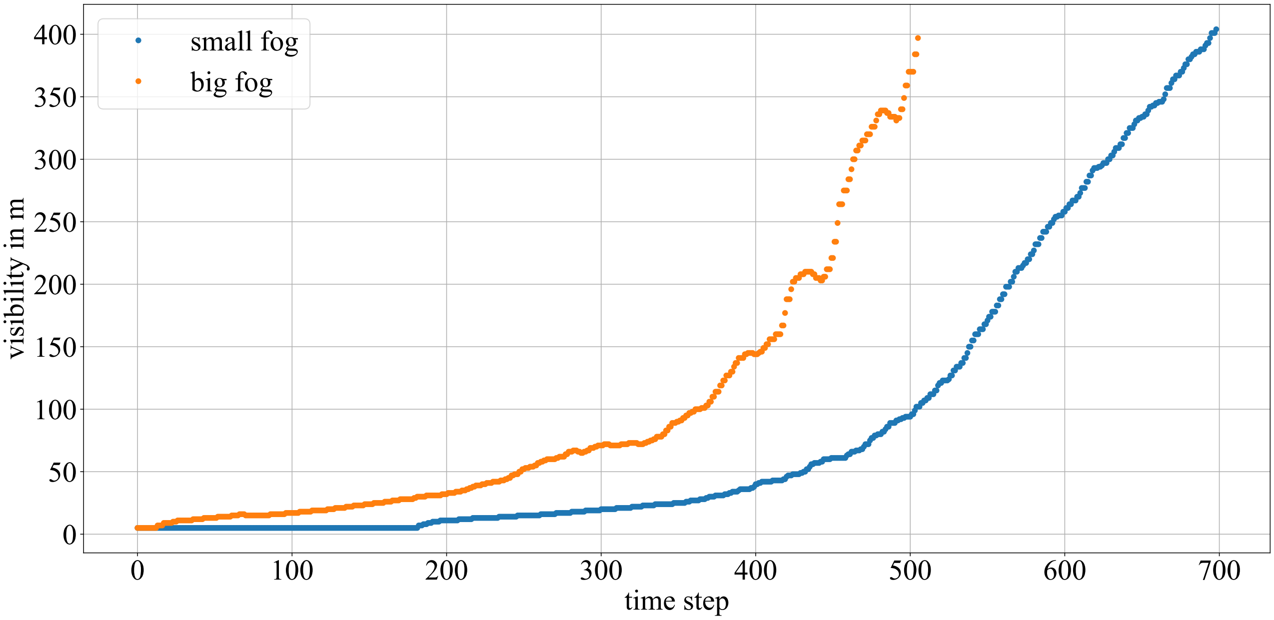

Figure 6.17: Visibility measured by the transmissiometer in small and big conditions at CEREMA.

Visibility

Fig.6.17 shows the evolution of visibility measured by the transmissiometer during both small and big fog dissipations. For both tests, the trend of visibility is not linear but increases as visibility rises and fog dissipates. As a result, the number of point clouds (recorded simultaneously) corresponding to specific visibility values varies and can be low when visibility rises fast. This has a consequence on the inference results because it influences the number of data points available for training and testing the algorithm. In addition, the small fog test shows a constant visibility of 5 m (minimum measurement of the transmissiometer) at the beginning of the dissipation as well as a slower increase compared to the big fog test. The visibility measured in big fog exhibits a faster rise, potentially due to the presence of larger and heavier particles that descend more rapidly.

Target and frustum detection

Similarly to the results drawn from stabilized levels of visibility, the detection capability of the Ouster LiDAR is highly altered during the fog dissipations. Fig.6.18 showcases the evolution of the target and noise POD for the Ouster LiDAR in small and big fog dissipations. As fog dissipates and visibility rises, frustum POD decreases as target POD increases. While the POD of the back target is comparable to a sigmoid function for both tests, the frustum detections do not show a similar behaviour.

Furthermore, the noise levels decrease as visibility rises but the small fog test shows a steeper slope, reaching close to 0 frustum detections at approximately 25 m of visibility. The constant visibility values of 5 m at the beginning of the small fog dissipation (fig.6.17) gives high variation in frustum POD (fig.6.18a) which forces us to consider only the end of the 5 m period for the future labelled dataset. Besides, both tests show a short constant level of frustum POD values (0.9 for big fog and 0.8 for small fog) when visibility is very low. Finally, it is noted that bigger fog particles require higher visibility values for the LiDAR to start detecting the target (around 70 m for big fog and 30 m for small fog). To conclude, it seems that there is no clear trivial relation between the target POD and the frustum POD that generalizes the model for both big and small fog experiments.

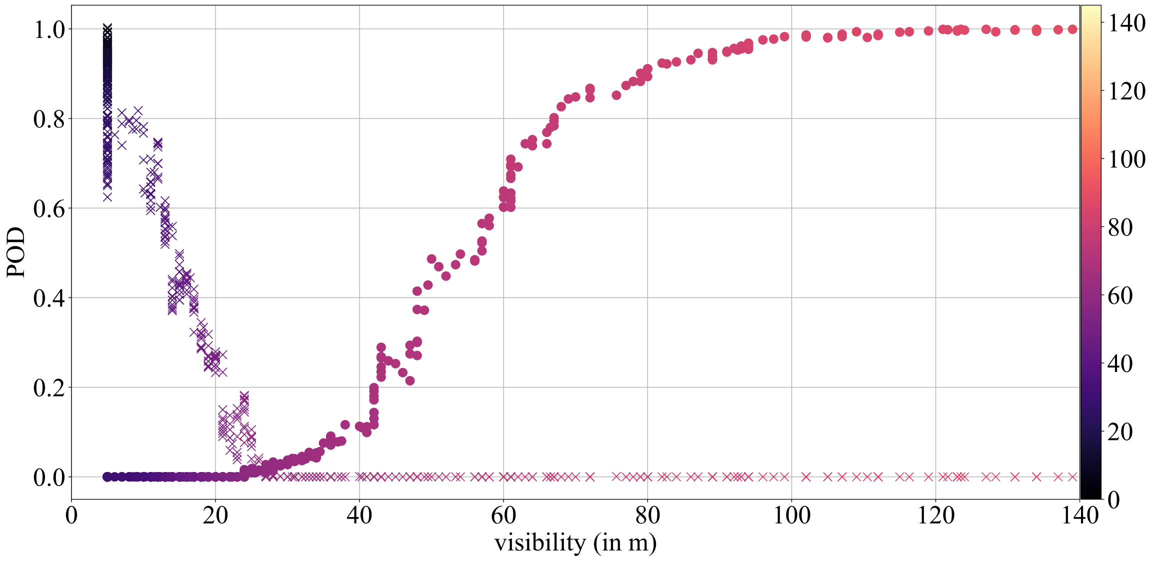

(a) Small fog

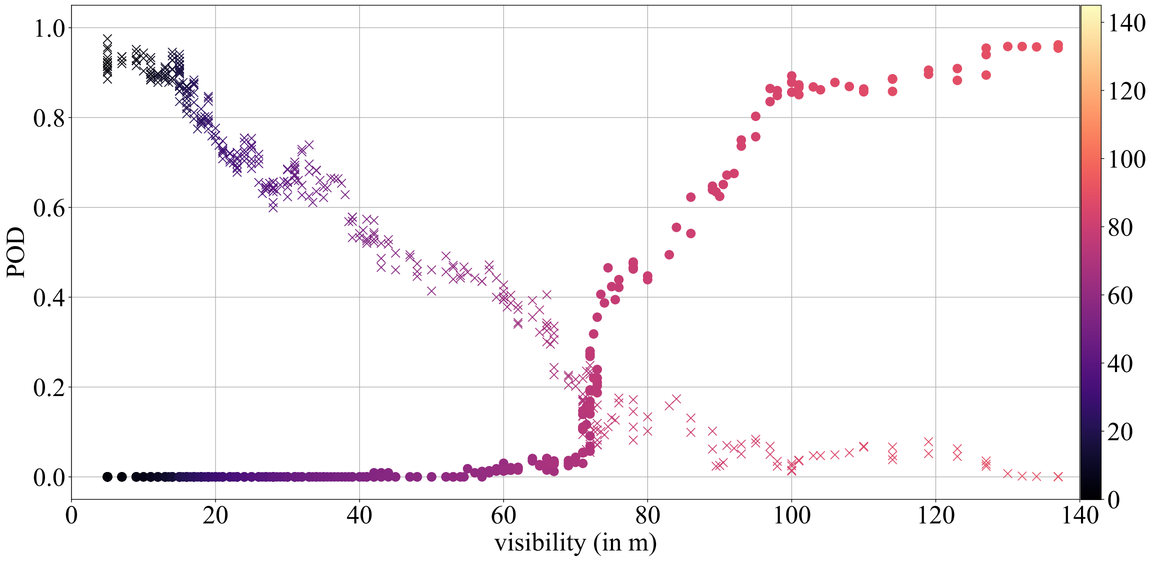

(b) Big fog

Figure 6.18: Target (circles) and frustum (crosses) POD according to visibility during small and big fog dissipations. Color represents visibility.

Frustum distance distributions

An intuitive way to visualize the distance distributions is displayed on the small-multiple graphs of fig.6.19, with respectively the results for small and big fog tests. Each sub-graph represents, in the form of a histogram, the distance distribution of frustum noise detections when the 3D-LiDAR aims at the *back* target according to specific values of visibility (here with classes of 5 m). For both fog tests, classes of visibility start from 5 m, the minimum value of the transmissiometer and the classes step is set to 5 m. The last visibility class is limited by the number of frustum points available at higher visibility values, resulting in 100 m for big fog and 30 m for small fog. These configurations are chosen empirically as they hold the best results and ensure enough data available to later train and test the model on each class, as shown in chapter 7.

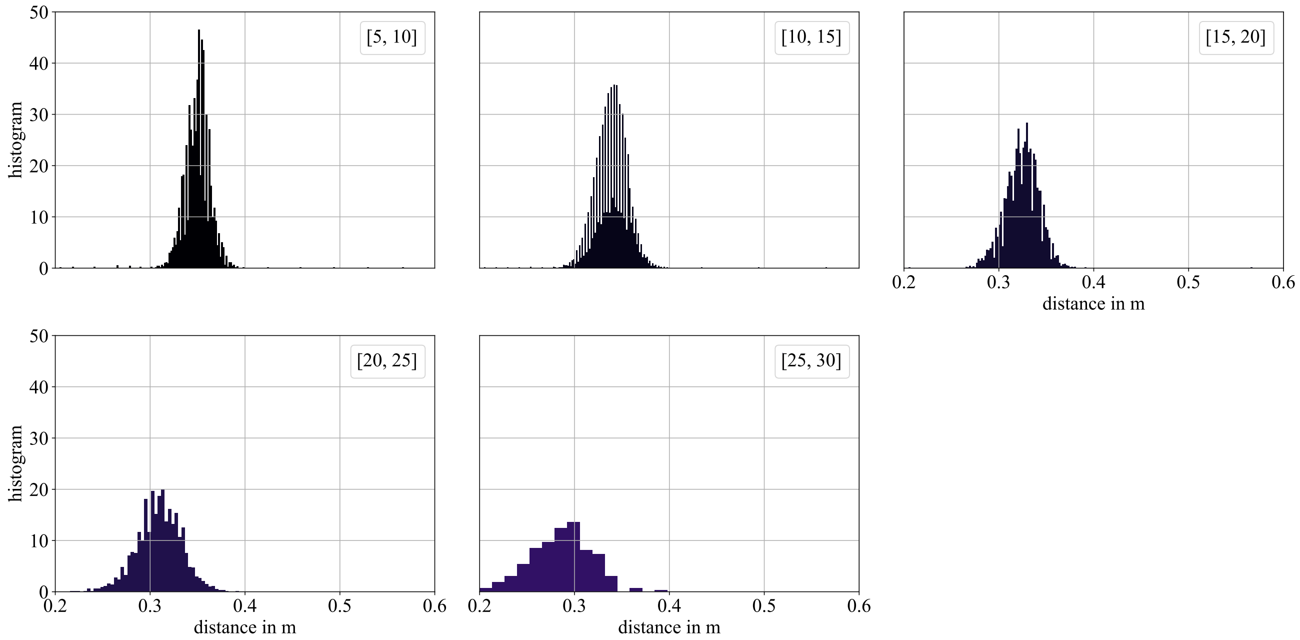

(a) Small fog

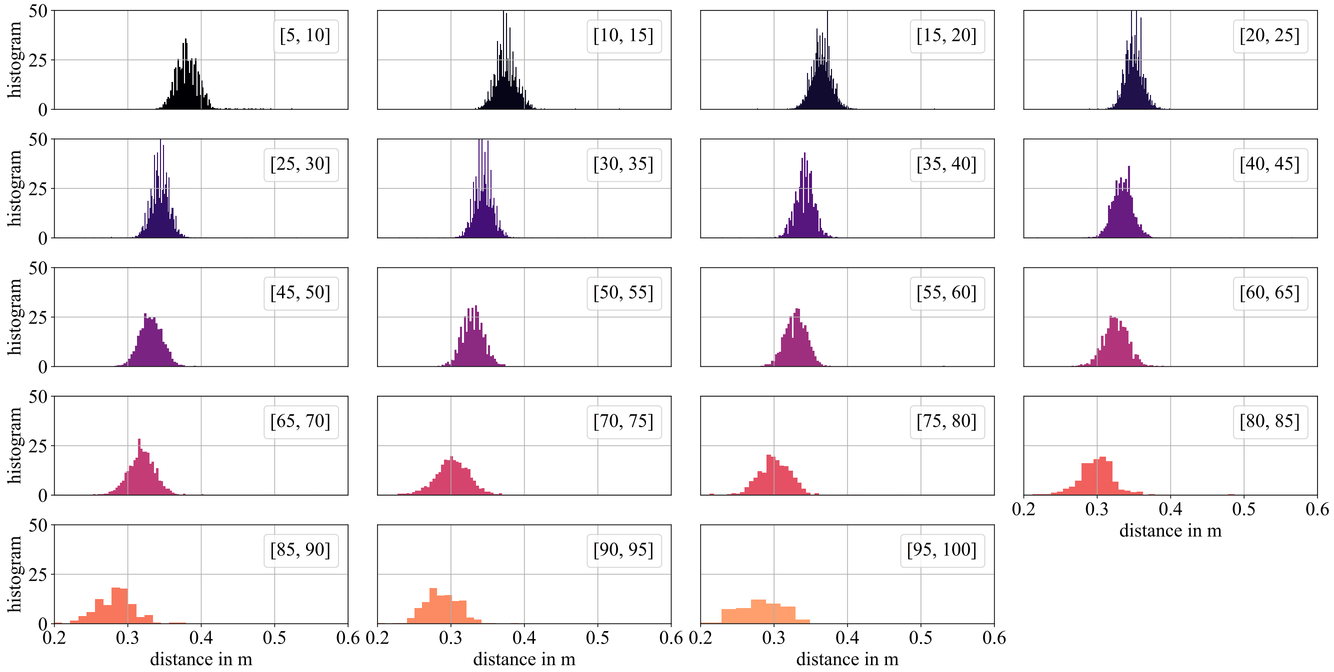

(b) Big fog

Figure 6.19: Distance histograms of frustum fog points detected with the *back* target for each visibility class in both fog conditions. Color refers to visibility in meter.

It first shows that these detections are, again, located at very close-range of the sensor, around 0.3 m. The proximity of these points is advantageous for our inference application, as it is highly improbable for a real object to be situated at such a short distance. In operational conditions, these points could be filtered out and used for inference purposes. The amplitude of the histograms decreases with visibility, aligning with our expectations and as depicted in fig.6.18. Additionally, the mean distance of the echoes tend to shift towards shorter distances as visibility increases, transitioning from approximately 0.4 m to 0.3 m. This phenomenon goes against intuition, as one might expect that lighter fog (higher visibility) would cause noise points to be located at greater distances. However, we can speculate on its underlying cause. In foggy conditions, the signals detected by a single photon LiDAR system can be modeled as Gamma distributions according to [180]. These distributions exhibit a peak and a tail resulting from the attenuated backscatter laser signal. In dense fog, the tail of the Gamma distribution is long due to numerous particles reflecting the photons, causing the LiDAR point to be slightly distanced from the first peak. In contrast, lighter fog produces a shorter tail in the single photon detection, resulting in a smaller mean distance value. Consequently, the LiDAR point generated is positioned at a shorter distance. Again, this is speculation about the sensor's internal signal processing, about which we have no information.

6.3 Experimental results in artificial smoke at Fauga

This section describes the experimental setup and results of the smoke experiments performed in the Fauga climatic chamber.

6.3.1 Experimental setup

The smoke experiments have been conducted in a climatic chamber owned by ONERA and shown on fig.6.20a. The facility is a 15 m × 5 m rectangular room in which it is possible to produce artificial smoke conditions. The smoke is produced by a fog-oil smoke machine with properties described in table 6.4, as measured in [191].

Type of aerosol	oil-fog
Aerosol size distribution	Log-normal
Modal radius	$0.18 \pm 0.01\ \mu m$
Complex refractive index	$1.508 + i10^{-5}$

Table 6.4: Microphysical properties of the artificial smoke particles generated during the smoke experiment.

LiDAR sensor

The Ouster OS1-128 spinning LiDAR, already presented for the CEREMA experiments, is the only LiDAR sensor used in this smoke study.

Targets

The smoke study uses 3 reflectance calibrated targets of 1m^2 each, namely *a1*, *a2* and *a3* for respectively the 80%, 50% and 10% reflectance values, displayed fig.6.20a. A *free-sky* target is introduced, representing the empty space defined within the rear doors at the end of the chamber, shown fig.6.20b. When the rear door of the chamber is open, the laser beams shot by the sensor inside the *free-sky* frustum never reach any solid targets. The rationale behind employing this frustum is explained in the methodology chapter 5.

Weather sensors and DVE control

Two context sensors providing visibility measurements have been used in this campaign. The first sensor is the Biral VPF700 scatterometer, measuring visibility at 0.1 Hz within 10 m to 75 km with a resolution of 10 m. Its characteristics are not suited for measuring short visibility ranges generated during this campaign. The second sensor is the Lighthouse Handheld 3016 particle size diameter (PSD) counters,

(a) Sensors and targets layout in the artificial fog experimental setup. *a1*, *a2*, and *a3* targets in decreasing reflectivity order. At the entrance, from left to right: the PSD sensor, the OS1-128 LiDAR, a Biral VPF700 scatterometer, and the smoke machine.

(b) Illustration of the *free-sky* target during a smoke dissipation with opened rear-doors. Note that the smoke is not uniformly distributed in the chamber during dissipation and evacuates mainly through the top of the chamber.

Figure 6.20: Experimental setup used during the artificial smoke dissipation experiments at ONERA.

already presented in chapter 3. It can simultaneously measure the concentration of particles according to size channels of 0.3, 0.5, 1.0, 2.5, 5.0 and 10 μm at a max frequency of 1 Hz.

6.3.2 Test protocol

Dissipation experiments are performed for the smoke conditions. The chamber is closed and saturated with smoke particles using the smoke machine. Then, the rear doors are opened and the recordings start. As opposed to the CEREMA experiments where the scene remains static for all acquisitions, the targets are, in this case, placed at different positions for each dissipation. Table 6.5 shows, for each dissipation, the target distance from the sensor. As such, the degradation in the point cloud data is assessed for targets at multiple distances. Naturally, the *free-sky* target is not defined with any distance because it is only used for its frustum detections.

Target	***a1***	***a2***	***a3***	***free-sky***
$d_1(m)$	12.4	13.3	13.4	∞
$d_2(m)$	10.9	11.1	11.0	
$d_3(m)$	9.4	9.4	9.3	
$d_4(m)$	7.1	7.2	7.2	
$d_5(m)$	5.5	5.5	5.5	
$d_6(m)$	4.1	3.8	3.7	
$d_7(m)$	1.5	1.4	1.3	

Table 6.5: Targets positions for the smoke experiments. Targets are moved manually before each dissipation. Note that, by definition, the *free-sky* target does not have any defined distance.

6.3.3 Experimental results

The utilization of context sensors, specifically visibility and PSD sensors, in this campaign introduces complexities in creating the labeled dataset compared to the fog experiments described above.

To simplify this part, the correlations between *free-sky* frustum and standard targets detections have not been included here, as previously shown for the CEREMA tests (figs.6.9 and 6.13). Instead, this particular part highlights the evolution of *free-sky* detections in relation to visibility, converted from the PSD data. We believe that emphasizing this data is more relevant, as it is crucial for the subsequent inference model.

From Particle Size Distribution (PSD) to visibility

Measurements from the context sensors of the smoke experiments are available on fig. 6.21. The graph shows simultaneously the concentrations of particles measured by the PSD sensor in color as well as the converted visibility value using the PSD data and Mie theory (computed according to equation 3.1), in black. In addition, timestamps corresponding to the recording times of the LiDAR data are shown with the light purple vertical bands. These time intervals correspond to the dissipation durations and last approximately 5 minutes each. This is significantly shorter than for the fog experiments performed at CEREMA, which results in less LiDAR data to build the labelled dataset. This short duration originates from the test protocol performed and the characteristics of the facility. Indeed, to obtain free-sky frustum LiDAR points, the rear-doors of the chamber are opened at the beginning of the recordings. This leads to faster dissipation of the smoke particles, as opposed to the fog protocol where the chamber remains closed.

Results are separated according to the test day. As the fog dissipates, visibility rises but the concentrations of particles ranging from 0.3 μm to 10 μm decreases. Visibility computed using Mie theory results in values from approximately 70 m to a maximum of 980 m. These results are quite different than the ones from the fog experiments, with visibility values higher by one order of magnitude. Additionally, these high visibility values do not match with either visual observations or impacts on the point clouds.

Errors may exist in our conversion approach, as several factors might alter the results including the non-spherical form of the particles, inaccurate data from the PSD sensor, the use of discrete particle sizes, or wavelength issues. While this discrepancy could be more investigated, it is not a limitation to evaluate our methodology or the performances of the inference model. Consequently, the following shows the results obtained with these Mie-converted visibility values.

Overall, visibility sometimes proves to be inappropriate or challenging to be measured accurately and thus to fit with our problem. During the CEREMA tests, visibility is measured using a transmissiometer that is permanently installed in the chamber and monitored by the staff. The resulting visibility measurement appears to be correct, considering its impact on the sensors. However, the sensor is most certainly calibrated specifically to the chamber's environment. Conversely, smoke experiments utilized a Biral VPF700 scatterometer. The results from this sensor are unsatisfactory, as it appears to be primarily tuned for airport applications rather than short distances (from 0 m to 100 m). Its results are not reported here. And, our PSD to visibility conversion method seems to give inaccurate results. This leads to conclude that the measure of visibility is influenced by multiple factors and should be calibrated specifically for each application. Somewhat similar to the rain rate, this metric might not be the most appropriate to fit with our approach.

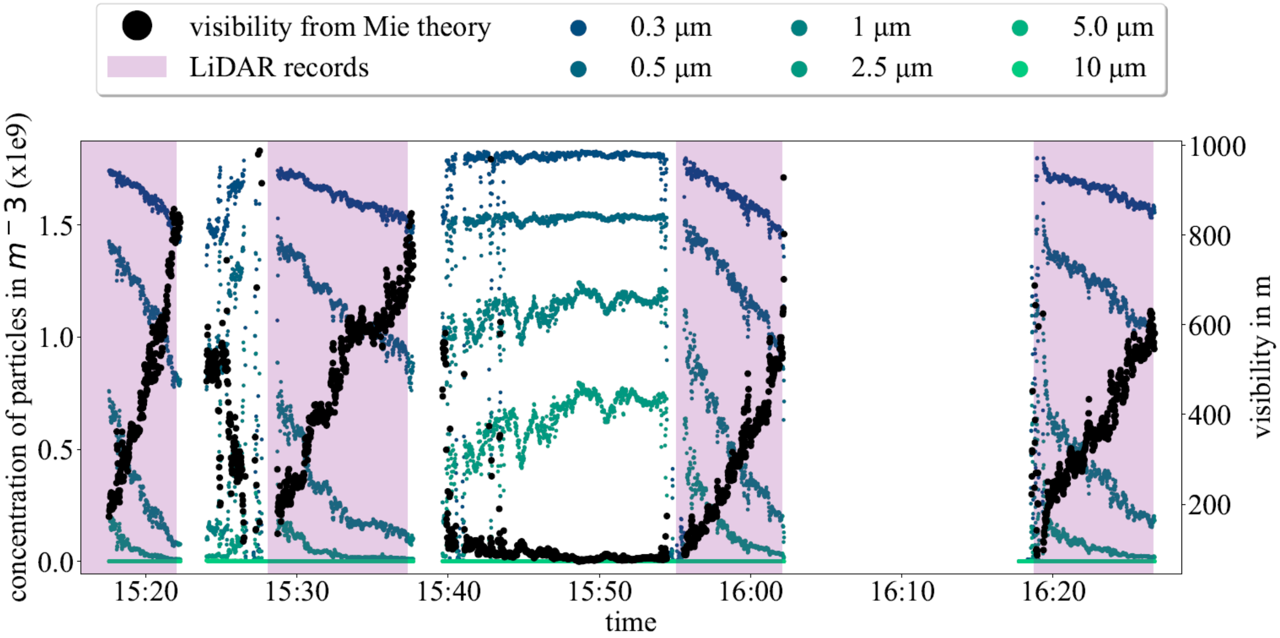

(a) Day 1

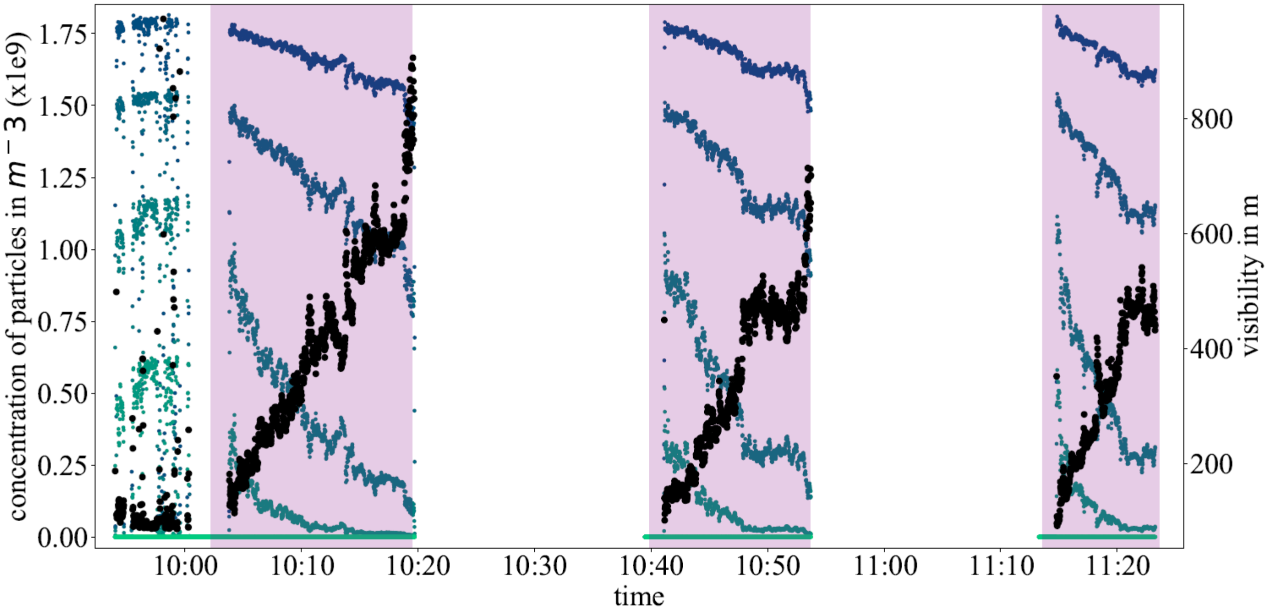

(b) Day 2

Figure 6.21: Measurements from PSD sensor (left y-axis) and conversion into visibility (right y-axis) with synchronized recording times of LiDAR point clouds for all performed dissipations. Although the dissipations are not fully controlled and changes in the dynamics of the visibility evolutions are observed, the particle distributions and computed Mie visibilities show similar behaviors.

Target and frustum detection

The conducted smoke experiments enable the extraction of *free-sky* frustum points, useful to follow the methodology detailed in chapter 5. Consequently, the study of frustum detections now only concerns the ones of the *free-sky*, which can be labeled with DVE properties (visibility, PSD).

The smoke experiments rely on the PSD data. The following presents the correspondence between PSD data and LiDAR *free-sky* frustum points. Fig.6.22 displays, for all recorded dissipations accumulated, the resulting *free-sky* POD according to the concentrations measured by the PSD sensor for each particle size. Across all dissipations, higher concentration of particles leads to higher noise levels in the 3D-LiDAR point clouds and inversely. An exponential tendency is observable, with an increased slope for smaller particle size.

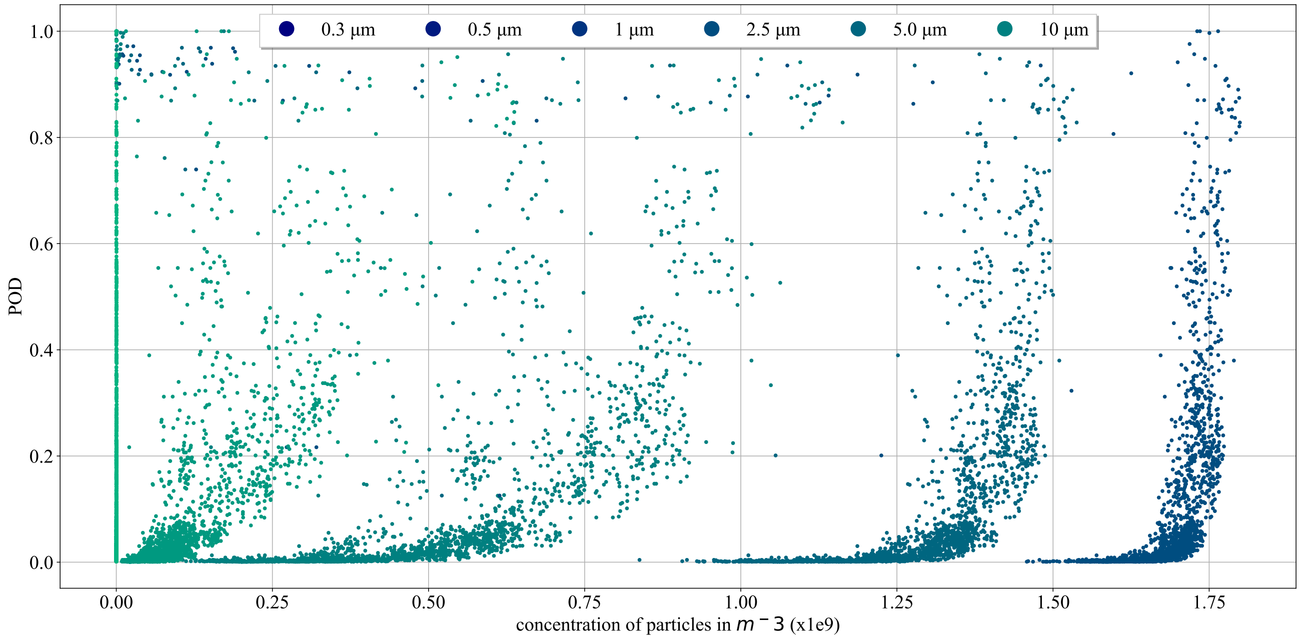

Figure 6.22: *free-sky* frustum POD according to each particle concentration PSD data.

Then, the evolution of the *free-sky* POD according to the converted visibility is displayed on fig.6.23. It provides a breakdown of the results for each dissipation individually to highlight their differences. A common decreasing exponential tendency is observed but disruptions in the evacuation and homogeneity of the smoke particles already discussed cause different results. This has an effect on the inference model performances because the model is trained on data accumulated from every dissipation.

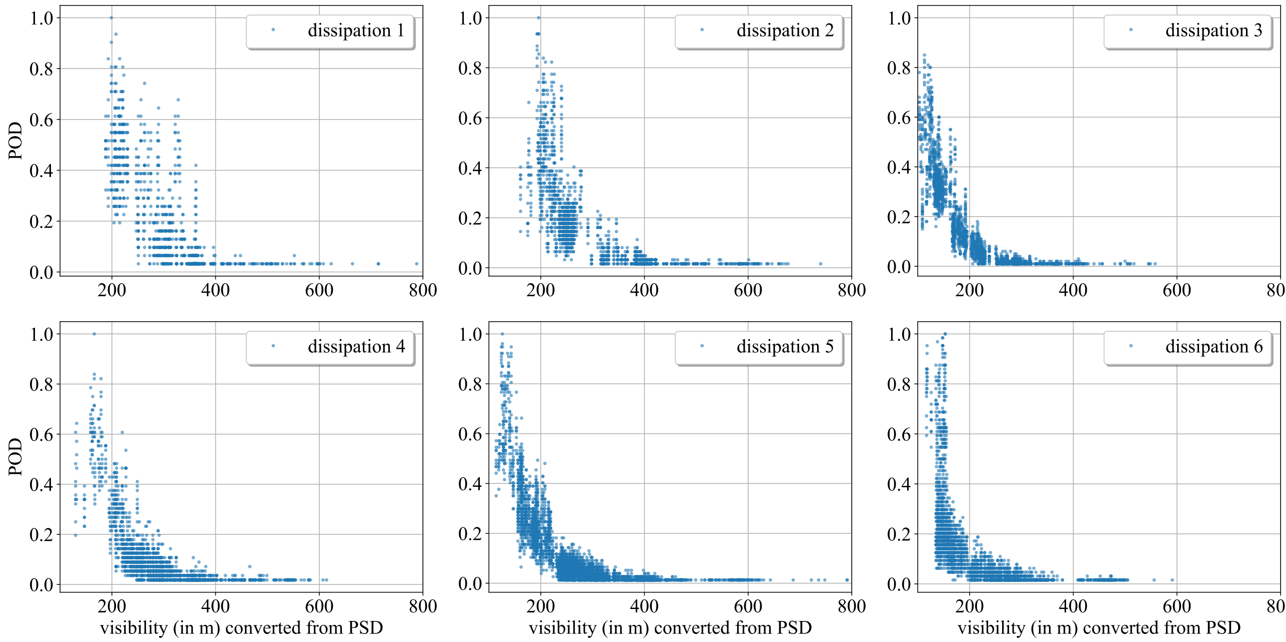

Figure 6.23: *Free-sky* frustum POD according to visibility converted from the PSD data for each dissipation performed.

- Frustum distance distributions

The labelled dataset of frustum detections and visibility classes, built from the smoke experiments, is given on fig.6.24. Here, classes of visibility are set from 75 m (minimum value of the converted visibility) to 425 m (too few points are available above) by steps of 25 m.

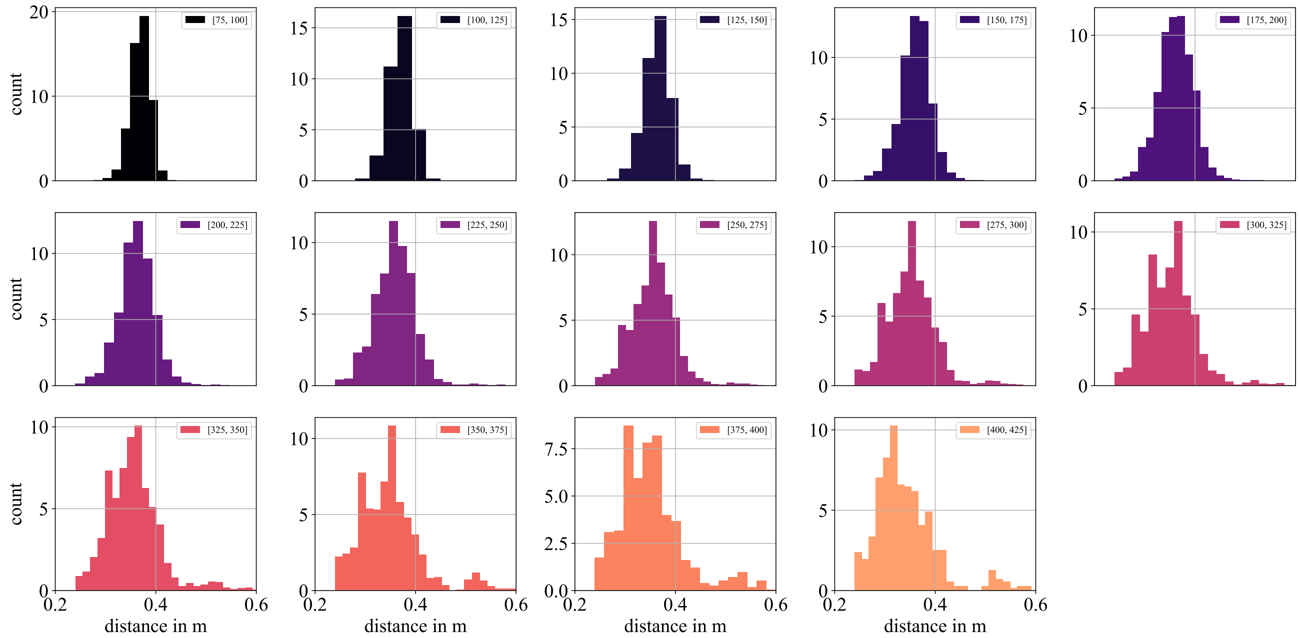

Figure 6.24: Distance histograms of *free-sky* frustum smoke points detected for each visibility class, color refers to visibility in meter.

We classify the visibility into classes of 25 m and label the distance distributions of the *free-sky* frustum points based on these classes. Similar observations regarding the distance distributions of frustum echoes in fog experiments with the OS1-128 LiDAR are observed in the smoke experiments. Detections located around 0.3 m have characteristic shapes with variations in position and amplitudes depending on the evolution of the visibility conditions. Though, the distance distributions of the smoke experiments are more chaotic than the ones of the fog of fig.6.19, which comes from the rather erratic nature of the smoke dissipations as discussed previously.

6.4 Conclusion

Quantitative results from automotive LiDAR sensors in DVE are presented in this study, with artificial rain, fog ad smoke conditions of various severity. A multitude of targets are studied and for each of them, the assessment of target and noise POD is possible by extracting the points respectively on the target and inside the sensor to target frustum. The frustum volume allows for a better understanding of the impacts of DVE on LiDARs while constraining the analysis to a specific space, which can be used in outdoor applications.

Performances differ for all sensors both in number of points and intensity. Sensor specific characteristics can be observed in their behaviors, internal designs of the LiDARs certainly explain the differences, but the information needed for such analysis are hardly accessible. The analysis of the impacts of DVE regarding sensor internal designs is another interesting topic that could be considered during the conception of the sensors. Our study is agnostic with respect to these design parameters because we focused on the use of the point clouds. Though, it seems that the use of 1550 nm laser emission of AEye LiDAR allows better penetration through obscurant media. A non repetitive scan pattern from the Livox LiDAR also shows good performances, especially in foggy conditions. The combined analysis of target and frustum detection aong with multi-echo labeling is valuable for future automatic classification of DVE.

One should be careful in the comparison of artificial and natural weather phenomena where limitations exist. Generated conditions from climatic chambers are useful for their repeatability but weather generation methods show limitations in terms of similarity with natural weather thus limiting further developments. Although care has been taken to optimize the generalization of our tests with real outdoor conditions, future work with outdoor acquisitions is essential in the deployment of our solution.

Additionally, the efficiency of the rain rate and visibility metrics to fit with our approach is mitigated. Their manner of averaging several factors makes them poorly suited for their correlation with the LiDAR impacts. Future work could include the investigation of more precise metrics.

Although not the most impressive sensor, the Ouster OS1-128 (whatever the version during our experiments) shows interesting results in fog and smoke conditions. First, the number of noise detections seems to be closely related to the severity of the DVE conditions (visibility, PSD). And the distance distributions of these noise echoes have characteristic shapes which evolve gradually with changing visibility values. In addition these points are located at very close range which is convenient to be used in algorithms because there is hardly any chance that a real object is located at such small distances. All this indicates a correct physical response, especially compared to other sensors, which is valuable to build an appropriate model. Though, an unexpected behavior is observed concerning the vertical index of the points but this is only presented in appendix A.

The forthcoming chapter describes the inference model that aims to identify the distance distributions of frustum points and make predictions regarding a most probable visibility class.

Chapter 7

Bayesian inference of visibility from 3D-LiDAR point clouds

7.1 Bayesian inference

Contrary to popular belief, AI is not a recent innovation that emerged overnight. While recent advancements in computing power have facilitated the use of previously abandoned techniques, such as deep learning, it is essential to recognize that AI encompasses much more than these methods. IBM defines AI as a field that combines computer science with robust datasets to enable problem-solving [192]. Similarly, Wikipedia characterizes AI as machines demonstrating intelligence by perceiving, synthesizing, and inferring information [193]. Ultimately, AI methods can be applied to virtually any computer task that would typically require human intelligence, and in particular, do not just involve deep learning.

The IBM definition highlights the involvement of datasets, implying that AI relies on the acquisition of data and learning from it. In order to learn from a dataset, it is necessary to label it (except for unsupervised learning techniques), establishing a correspondence between the data and their corresponding labels. This labelization process is considered the lifeblood of training datasets as the AI performance often relies on the accuracy of its dataset labeling. Often accomplished through low-paid and underprivileged jobs and also with the famous online captchas (without anyone noticing that they are working to improve an AI for free), the reliance on such labor-intensive tasks raises ethical concerns, as the individuals involved often face unfavorable working conditions while contributing to the advancement of AI technology. However, AI encompasses various methods that may require less data in exchange for a more precise approach. Developing such methods requires individuals to delve into the realm of data science and statistics to gain a profound understanding of data behavior. This is the direction that was taken in this thesis.

In the domain of robotics, probabilistic pattern recognition methods commonly serve as the foundation for developing features that enable autonomous navigation. Building a perception stack for autonomous driving can somehow consist in building a harmonious collection of algorithms which can jointly work together, as shown chapter 4. One advantage of working with probabilities is that they can be exchanged from one algorithm to another to be able to propagate uncertainties on the required information. In this way, using probabilities helps to construct the aforementioned harmonious system. As a result, the motivation behind developing the inference model in this chapter is to contribute to an existing software stack that already relies on probabilities. Using a probabilistic approach also brings the advantage of facilitating compliance with safety regulations, a critical aspect for autonomous driving. Adhering to standards like the IEC 60721-3-5 (*Classification of environmental conditions – Part 3: Classification of groups of environmental parameters and their severities – Section 5: Ground vehicle installations*) becomes essential to ensure the safe deployment of autonomous vehicles on roads. Probabilistic techniques, which explicitly model uncertainties throughout various processes, are more conducive to meeting these regulatory requirements compared to deep learning methods, which often operate as black boxes.

$$P(A|B) = \frac{P(B|A)P(A)}{P(B)} \tag{7.1}$$

Bayesian probabilities, ruled by the Bayes theorem of eq.7.1, is a major concept for conditional probability and inference. As opposed to the frequentist approach where an event's probability relates to its frequency of happening, a Bayesian probability represents a state of knowledge. Using a prior probability $P(A)$, the state of a random variable $P(A|B)$ is updated given new relevant data B. From this basic concept, a multitude of ideas can be derived, the aim of which is to estimate the state of a variable using data - in other words, to make an inference. Depending on the nature and properties of the random variables and data, different architectures are possible. Extensive descriptions on these methods are available in [194, 195, 196].

P.S.: "Fun fact, the writing of this introduction paragraph on artificial intelligence (and various parts of this thesis) was actually collaborative, with yours truly, ChatGPT, making a substantial contribution. wink". Indeed, this AI is relatively efficient at helping with writing, as it can suggest many alternative ways to write sentences.

7.2 Visibility inference model

7.2.1 Bayesian formulation

Our inference model is a naive Bayes classifier. Its goal is to infer discrete classes of visibility V from independent features – in our case, series of echoes e (fig.7.1). It is a supervised model because a labelled dataset $\mathbf{Y}$ is used[1], and it is generative as it allows to generate samples of distances e from classes of visibility V or model parameters θ, as shown on fig.7.2.

Figure 7.1: Probability Graphical Models (PGMs) of the naive Bayes classifier model for visibility classification using independent LiDAR echoes, related to eq.7.2. (a): developed model, (b): same model using the plate notation.

Fig.7.1a shows the Probabilistic Graphical Model (PGM) representation of a naive Bayes model applied to our problem where i.i.d. features $\{e_1, e_2, ..., e_D\}$ are used to classify a visibility class V. A PGM representation gives visual insight about the causal dependencies between the random variables involved in a probabislic model, the theory behind PGMs is given in [194]. Naive Bayes models are based on the assumption of independence between the features so that the joint model can be expressed as $P(V, e_1, e_2, ..., e_D) = P(V) \prod_{d=1}^{D} P(e_d|V)$. LiDAR echoes are considered independent because the laser beams are fired independently, and each laser emitter is paired optically with a unique detector. For each visibility class, D features (echoes in our case) can be used for inference and D can vary depending on the instance of the class V.

Let $\mathbf{Y} = \{(\mathbf{E_n}, \mathbf{V_n})\}$ be the labelled dataset which consists of distances $\mathbf{E_n} = \{\mathbf{e_1}, \mathbf{e_2}, ..., \mathbf{e_D}\}_\mathbf{n}$ and discrete optical visibilities $\mathbf{V_n}$, for $n \in [1, N]$. At run-time, when operating in DVE conditions, the set of echoes distances $E = \{e_1, e_2, ..., e_D\}$ from a free-sky sensor frustum is acquired (these points are noise artifacts created by the DVE conditions). Using Bayes' theorem, the probability of having a visibility value V knowing the perceived set E and the dataset $\mathbf{Y}$ is $P(V|E, \mathbf{Y})$:

$$P(V|E, \mathbf{Y}) = \frac{f(E|\mathbf{Y}, V) P(V)}{\sum_{i=1}^{N} f(E|\mathbf{Y}, V_n) P(V_n)} \tag{7.2}$$

where $f(E|\mathbf{Y}, V)$ is the likelihood of the echoes distances E for a specific visibility V, and $P(V)$ is the prior probability on the discrete visibility class. Without any knowledge on the status of visibility, $P(V)$ is considered uniform. Information available from a local weather station could be included here and improve the precision of the inference. Additionally, a recursive Bayesian filter would improve inference results over time by using the posterior probability at a certain time step k as the prior probability of the next time step $k + 1$ [197]. A discrete Bayes filter has been developped and is presented in chapter 8. We now introduce model parameters, as the vector θ, used in the model to parametrize the different likelihood distributions. In the following, θ remains a generic hidden variable but more information is

---

[1] **boldface** means that the characters belong to the labelled dataset, following the convention of [195].

given in section 7.2.2. Through marginalization over θ, $f(E|\mathbf{Y},V)$ becomes :

$$f(E|\mathbf{Y},V) = \int_{\theta} f(E|\theta)f(\theta|\mathbf{Y},V)d\theta \tag{7.3}$$

Fig.7.2 adds θ to the PGM representation of the model. It shows that distance samples e can be generated from the random variable θ and $f(E|\theta)$ can be computed as in eq.7.6.

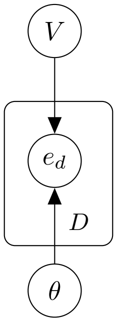

Figure 7.2: PGM of the naive Bayes classifier model for visibility classification using independent LiDAR echoes, with the addition of the model parameters θ.

The rest of the section describes the equations useful to compute the different probabilities and to understand the training step. The Probability Density Function (PDF) $f(E|\theta)$ from eq.7.3 is the likelihood of the echoes given parameters θ. Given that each echo feature is independent, the likelihood $f(E|\theta)$ is the product of probabilities of receiving each echo following the likelihood distribution L,

$$f(E|\theta) = \prod_{e\in E} L(e|\theta) \tag{7.4}$$

The second part of eq.7.3 is the PDF $f(\theta|\mathbf{Y},V)$, it represents the learned model parameters of the likelihood distributions generated by the echoes observed at a given visibility V. For each visibility class $\mathbf{V_n}$ in the training set, each set of echoes $\mathbf{E_n}$ is paired with $\mathbf{V_n}$ in $\mathbf{Y}$, so $f(\theta|\mathbf{Y},V)$ becomes

$$f(\theta_n|\mathbf{Y},V=\mathbf{V_n}) = f(\theta_n|\mathbf{E_n},V=\mathbf{V_n}) \tag{7.5}$$

This means that θ_n only depends on the part of the training dataset $\mathbf{Y}$ acquired at the specific visibility V_n. We use Bayes' theorem again to expose the likelihood of $\mathbf{E_n}$ knowing θ_n,

$$f(\theta_n|\mathbf{E_n},\mathbf{V_n}) = \frac{f(\mathbf{E_n}|\theta_n,\mathbf{V_n})P(\theta_n|\mathbf{V_n})}{\int_{\theta_n} f(\mathbf{E_n}|\theta_n,\mathbf{V_n})P(\theta_n|\mathbf{V_n})d\theta_n} \tag{7.6}$$

Similarly to the PDF $f(E|\theta)$, $f(\mathbf{E_n}|\theta_n,\mathbf{V_n})$ can be modelled by the likelihood distribution with independent events so that :

$$f(\theta|\mathbf{E_n},\mathbf{V_n}) = \frac{\prod_{e\in\mathbf{E_n}} L(e|\theta_n)P(\theta_n,\mathbf{V_n})}{\int_{\theta_n}\prod_{e\in\mathbf{E_n}} L(e|\theta_n)P(\theta_n,\mathbf{V_n})d\theta_n} \tag{7.7}$$

Eq.(7.7) can then be used directly to learn the distribution of θ for each visibility class. Because the integral of the denominator does not have an analytical closed form, we use a sampling technique to produce a numerical approximation on the distribution of the parameters and best approximate $f(\theta|\mathbf{E_n},\mathbf{V_n})$. Namely, a Markov Chain Monte Carlo (MCMC) sampling technique is used. The results consist in a collection of learned parameter samples for each visibility instance, more details concerning the training step and this sampling technique are given section 7.2.3.

To use these samples and compute the probability of a visibility class (eq.7.2 and 7.3), we use Monte

Carlo integration [198]. This technique allows to compute an integral using random discrete samples at which the integral is evaluated. In our case, random samples of θ are taken from the sampled representation of $f(\theta|\mathbf{Y}, V)$ (learned with the MCMC method) and injected into $f(E|\theta)$ so that eq.7.3 can be calculated.

The schematic diagram of fig.7.3 shows the overall framework with both learning and testing phases applied to visibility classification. The learning block represents the approximation of the parameters distribution, described later. During the test phase, frustum echoes E are received. Using eq.7.3, the likelihood probability $f(E|\mathbf{Y}, V)$ of each class V is calculated with Monte Carlo integration using the learned parameters θ for this class and the product of likelihood (eq.7.4). Finally, eq.7.2 is used to compute the probability of each visibility class using Bayes' theorem.

Consequently, a probability value is assigned to each visibility class. We decide that the class with the highest probability value is the one inferred.

7.2.2 Gamma or Log-normal likelihood models

The Gamma distribution is introduced in the context of autonomous vehicles sensors in [199] to model the distance repartition of echoes from LiDAR point clouds in snowfall conditions. No physical explanation is given for this choice, but it is showcased that it fits well the data. Also under snow conditions, [200] exposes similar distance distributions of echoes, but suggests to use the Log-normal distribution. They also state that the shapes of the distance distributions arise from the product between an ”optical detection” function and a ”building shielding effect” because the sensors are located at a building window. Satat et al. study the properties of light reflected from fog-particles and objects captured by single-photon avalanche photodiode detectors (SPAD) in a 3D-LiDAR design [180]. They showcase that the time profiles of fog and objects laser returns respectively have Gamma and Gaussian distributions. This allows them to extract real targets inside fog and reconstruct the 3D scene.

Considering the application of laser pulses in fog or smoke conditions, the resulting signal on the LiDAR detector is a combination of the laser pulse shape, considered as Gaussian, an atmospheric extinction function, which can be modelled as a decreasing exponential function, and of the detector response. This yields a product similar to a Gaussian function restricted to the positive domain, and therefore with a longer decreasing tail.

In the absence of a more refined model and according to the literature, both Gamma and Log-normal distributions seem valid candidates to model the distance distribution of fog or smoke echoes returned by a LiDAR. The PDF of a noise echo at a distance x which follows a Gamma distribution with the shape-scale parametrization $\theta = (\gamma, \beta)$ is expressed as :

$$L_G(x) = \frac{x^{\gamma-1} exp(-x/\beta)}{\Gamma(\gamma)\beta^{\gamma}} \tag{7.8}$$

with $\Gamma(\gamma)$ the Gamma function evaluated at γ, and γ is the shape parameter and β is the scale parameter. The PDF of a noise echo at range x which follows a Log-normal distribution is parametrized by $\theta = (\mu, \sigma)$ with mean μ and standard deviation σ and is defined by :

$$L_{LN}(x) = \frac{exp(-(ln(x) - \mu)^2/2\sigma^2)}{x\sigma\sqrt{2\pi}} \tag{7.9}$$

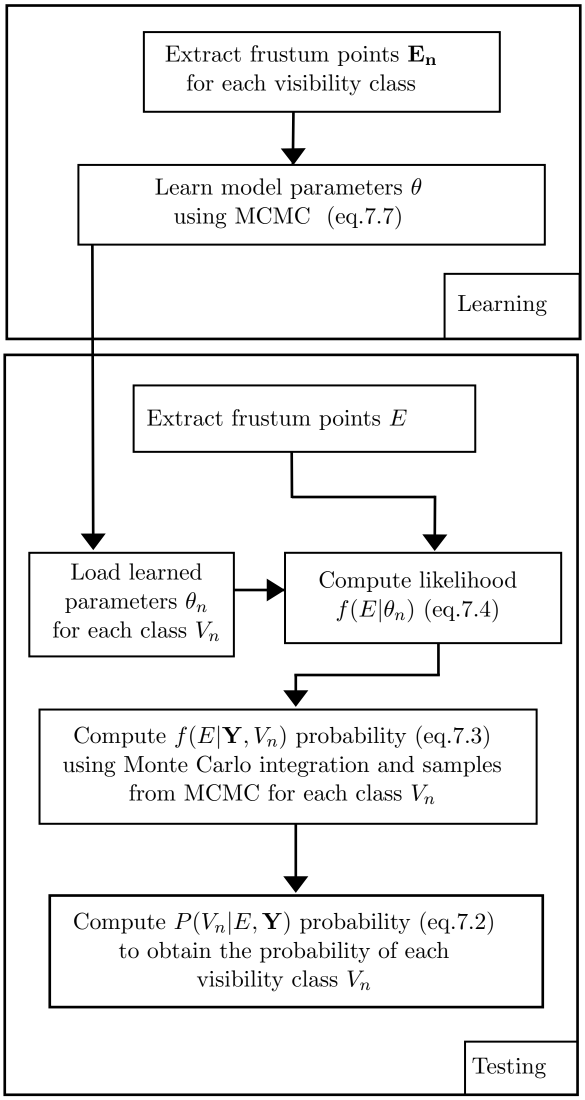

Figure 7.3: Schematic diagram of the Bayesian inference framework.

7.2.3 Training with Markov Chain Monte Carlo (MCMC)

As previously said, a sampling approach is used because no simple analytical solution is available to compute the product of likelihood functions in $f(\theta|\mathbf{E_n}, \mathbf{V_n})$. This PDF represents the statistical distribution of our model parameters. Whether the likelihood model representing the echoes is Gamma or Log-normal (or RFS, introduced in section 7.3), $f(\theta|\mathbf{E_n}, \mathbf{V_n})$ encompasses the value of their parameters for each visibility class v, in the form of samples. MCMC methods are statistical algorithms designed to sample probability distributions. These methods employ random walk methods to create Markov Chains that generate distributions proportional to given functions, such as $\prod_{e \in \mathbf{E_n}} L(e|\theta)$ from eq.7.7 in our case.

The Metropolis Hastings algorithm is used with a Gaussian proposal density and a rejection rule [194]. In the end, our training set consists of samples of Gamma or Log-normal parameters trained on distance distributions and various visibility classes (and different conditions). Fig.7.4 shows the previously presented distance distributions generated from the small fog, big fog and smoke experiments and the addition of Gamma (red) and Log-normal (green) distributions trained using the MCMC method. Using the learned samples for each visibility class, the distributions are plotted with their average value as parameters. Both Gamma and Log-normal fit the distributions relatively well, and neither of them seem to outperform the other.

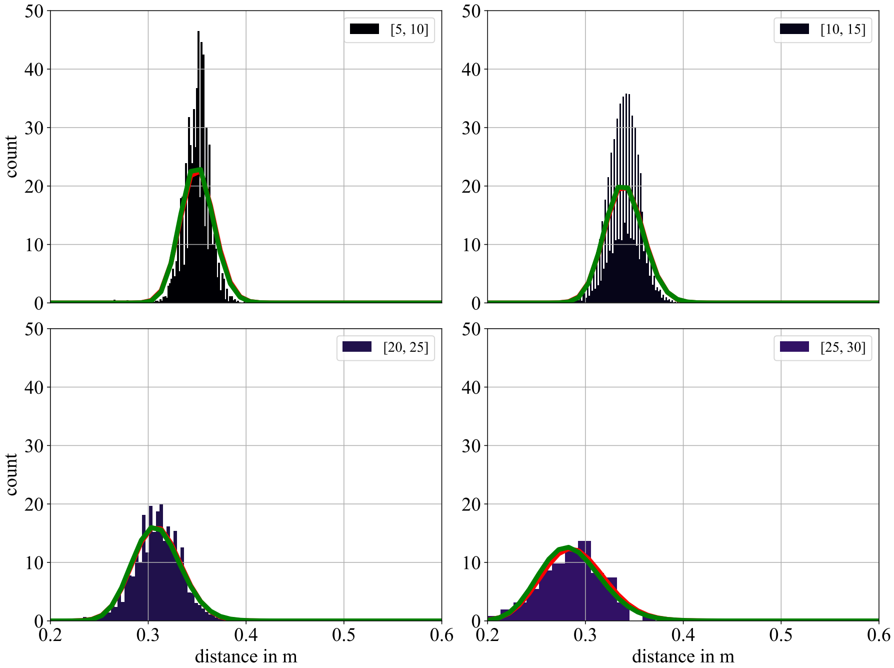

(a) Small fog

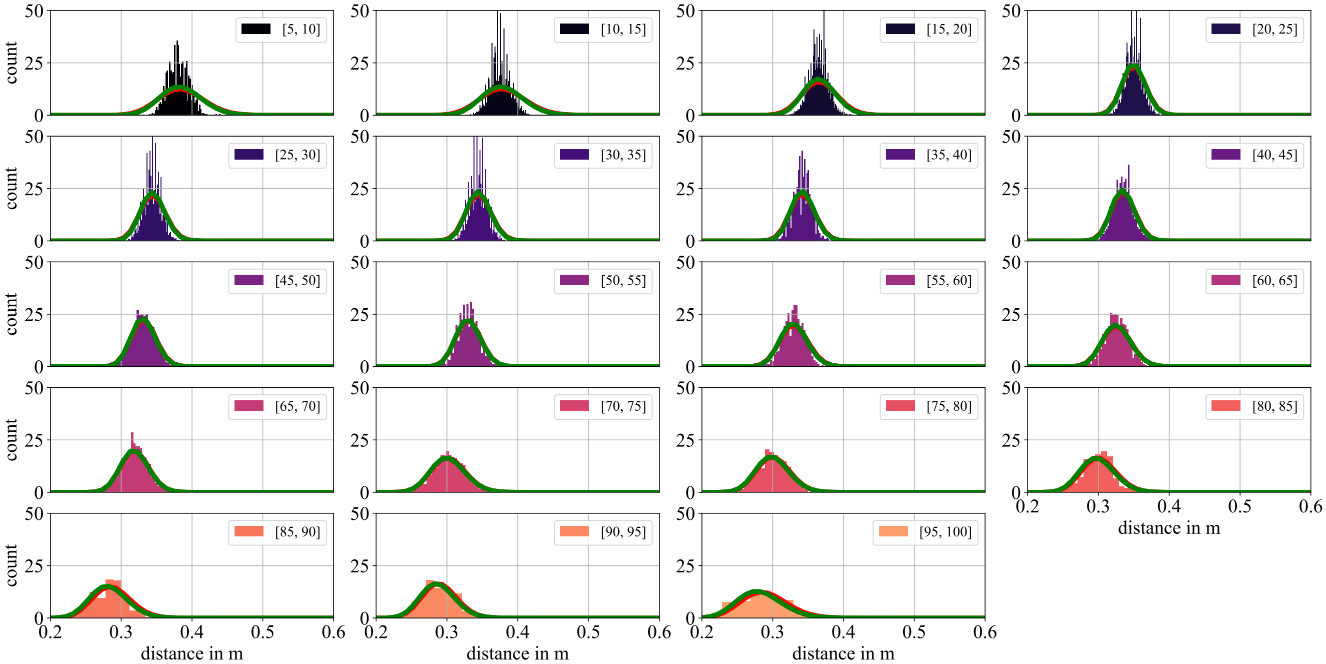

(b) Big fog

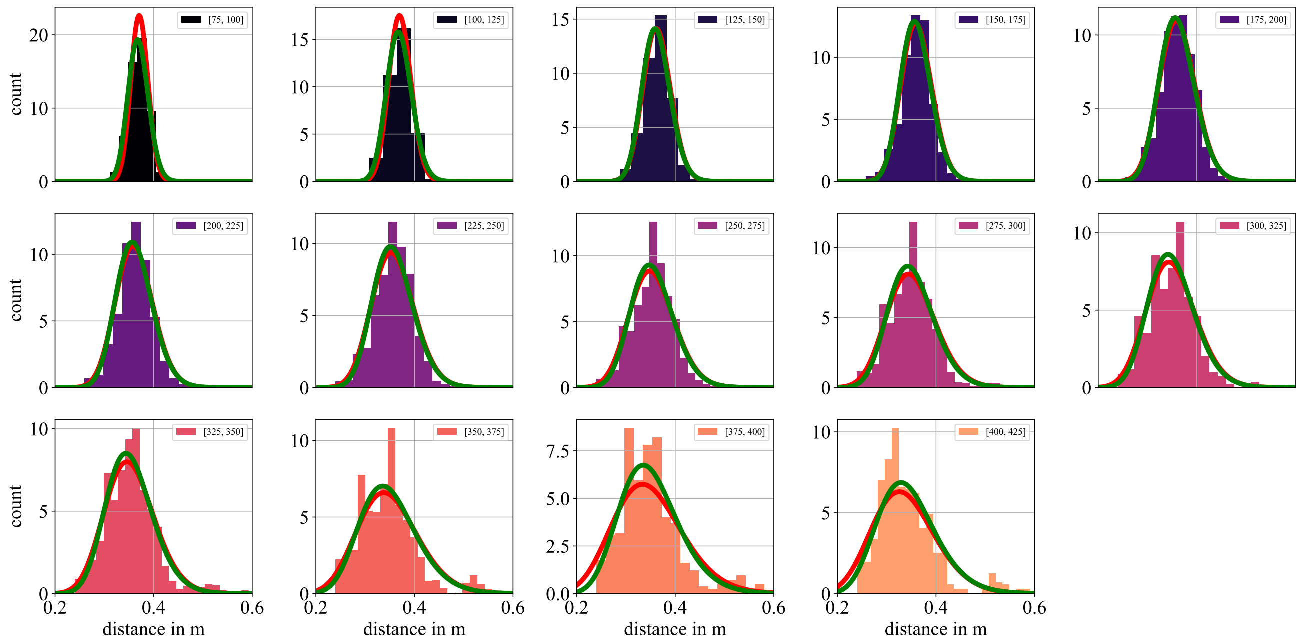

(c) Smoke

Figure 7.4: Distance distributions of the fog and smoke experiments, labeled with visibility classes. Gamma (red) and Log-normal (green) are plotted with parameters learned with the MCMC sampling technique.

7.3 Extension with Random Finite Set (RFS) formalism

The distance distribution of echoes alone does not always provide sufficient information to fully characterize visibility values, as stated in the previous chapter with erratic and counter intuitive behaviors. Analytically, computing the likelihood of a visibility class solely based on the echo distance distribution is a challenging task. Moreover, the results of comparing distance distributions with measured visibility in experimental conditions yield counter-intuitive results, indicating the involvement of sensor-dependent detections. This introduces additional complexity and potential errors in the classification process.

To achieve a more stable inference, two main approaches can be explored. The first approach involves incorporating characteristic information that can improve inference by distinguishing between two visibility conditions that have similar echo distance distributions. Previous sections have demonstrated that the number of frustum point detections is closely related to visibility. As visibility decreases, the number of echoes in the frustum increases, aligning with the intuitive understanding that higher water particle presence leads to more returns in the frustum. The Random Finite Set (RFS) formalism (detailed next section) addresses this concept by statistically capturing the number of frustum detections along with the distance distribution. The second approach involves making assumptions about the dynamics of environmental conditions and the evolution of visibility values. Exploring this idea could lead to more realistic inference, by considering possible dynamics. This approach is lightly elaborated using a discrete Bayes filter and presented chapter 8. However, it requires additional data for analysis and algorithm tuning, which is why it remains presented as an extension.

7.3.1 Random Finite Set

The Random Finite Set (RFS) theory is a mathematical framework originated from the Point Process (PP) theory [201], which is used in information fusion and multi-target detection and tracking [202, 203].

We do not delve into the fundamentals of the PP theory, as it is quite complex. Instead, we rely on the works of Malher et al. [202] and Vo et al. [203], as they offer a more comprehensive understanding of utilizing the PP via the RFS formalism.

An RFS $X = \{x_1, x_2, ..., x_n\}$ is a random variable which takes values as unordered finite sets. The purpose of an RFS is to capture the probability state of elements contained in a set as well as the cardinality of the set, noted $|X| = n$. The cardinality is random and modelled by a discrete distribution $\rho(n) = P\{|X| = n\}$. The elements in the set are modelled by $L_\theta(x_1, x_2, ..., x_n)$, the joint distribution of the elements of X, parametrized by θ. Common probabilistic descriptors apply to it, such as the PDF. Following Mahler's approach [202] and finite set statistics, the PDF of a RFS is described as follows:

$$f(\{x_1, x_2, ..., x_n\}) = n! \cdot \rho(n) \cdot L_\theta(x_1, x_2, ..., x_n) \tag{7.10}$$

7.3.2 RFS models

Our Bayesian inference model can then be extended, using the RFS formalism, by taking into consideration the number of echoes contained in our point clouds, in addition to the distance of the echoes (modelled by Gamma or Log-normal). We now consider our set of echoes E from section 7.2.1 as an RFS :

$$E = \{e_1, ..., e_n\} \tag{7.11}$$

where a sufficient number of echoes $n > n_{lim}$ is received and e is an echo distance. The following describes respectively the Poisson, Bernoulli and Binomial RFS models which are later used for the inference model as likelihood functions $L(e|\theta)$ of eq.7.4. Their discrete cardinality distribution is first given before the overall PDF is deduced from eq.7.10 [203].

- Poisson RFS

The Poisson distribution is a discrete probability distribution which describes the probability of a given number of events occurring in a fixed interval of time or space, parametrized by the Poisson rate l. The Poisson RFS is the most common point process for spatial point patterns, it uses the Poisson distribution as a cardinality distribution.

The number of LiDAR echoes received from DVE particles within a specific time interval can be modelled using a Poisson distribution. The cardinality of a Poisson RFS is Poisson distributed with rate l and likelihood distribution of the model L. In this case, the model parameters θ initially coming from Gamma or Log-normal distributions are now extended with the Poisson rate l. If $E \sim$ Poisson RFS(θ), E is described as :

$$\begin{cases} |E| \sim \text{Poisson}(l) \\ \text{e in E are i.i.d. with likelihood distribution } L(e) \end{cases} \tag{7.12}$$

The Poisson cardinality distribution is described as :

$$\rho(n) = \frac{exp(-l)\ l^n}{n!} \tag{7.13}$$

Given eq.7.10, the PDF of a Poisson RFS E is :

$$P(E|\theta) = exp(-l)\ l^n \prod_{i=1}^{n} L(e_i) \tag{7.14}$$

The Poisson rate l is the only parameter characterizing the Poisson distribution and defines both its mean μ and variance σ. Fig.7.5 shows Poisson distributions plotted with different values of Poisson rate. The Poisson distribution accounts for the random arrival and departure of objects and allow for an unbounded number of objects. In our case, it represents a distribution on the number of echoes received in a certain period of time (scan rate) and space (frustum volume) for a given visibility condition. During the computation of the Poisson RFS likelihood ($f(E|\theta)$ using eq.7.14), we normalize the Poisson rate by the total number of shots emitted in this specific frustum. As a result, the Poisson rate still acts like the μ and σ of the Poisson distribution, but is now comprised between 0 and 1 and captures the probability of having a certain number of returns for each shot. Additionally, it makes it comparable to the success probability of the Binomial RFS described just below. However, this model has its limitations in our case as the number of returns following a laser shot is either 0 or 1 (unless the sensor has multi-echo) and the Poisson distribution (parametrized with a Poisson rate between 0 and 1) can model higher number of returns, as shown on fig.7.5.

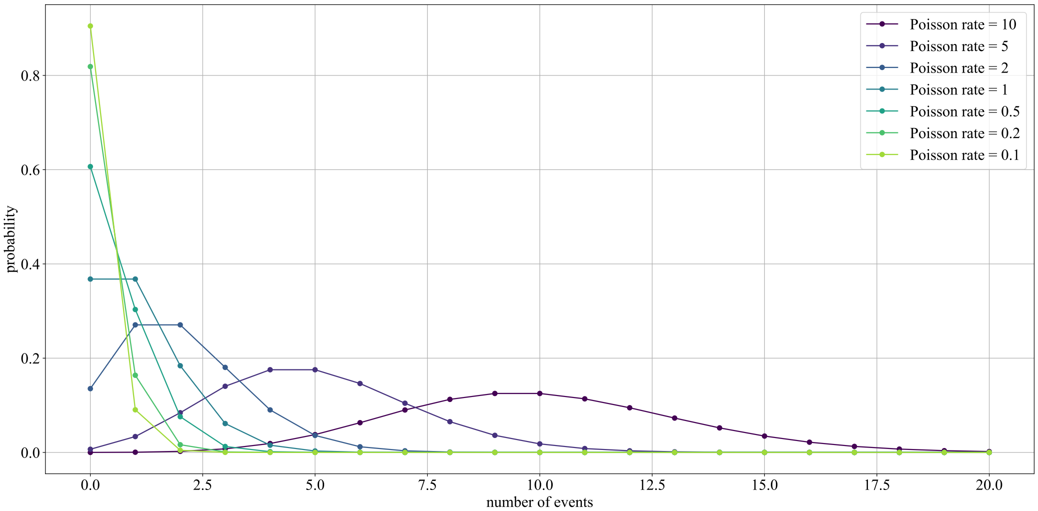

Figure 7.5: Poisson distribution plotted with various Poisson rates

- Bernoulli RFS

The Bernoulli distribution is a discrete probability distribution which describes the outcomes of any single experiment that asks a “yes-no” question. The probability of receiving a LiDAR echo from DVE particles can be modelled using a Bernoulli distribution, where the laser shot is a Bernoulli trial leading to a receiving echo or not. A Bernoulli RFS E is parametrized by the probability of success r and its cardinality distribution follows a Bernoulli distribution. E can be empty with probability $1 - r$ or a singleton with probability r. In the latter case, it contains the distance e of the returned echo following a likelihood distribution L_θ.

$$P(E|r,\psi) = \begin{cases} 1 - r & E = \emptyset \\ r \cdot L_\theta(e) & E = \{e\} \end{cases} \tag{7.15}$$

- Binomial RFS

By repetitively sending independent laser shots and receiving echoes (or not), we repeat Bernoulli trials. The binomial distribution is a discrete probability distribution which describes the number of successes in a sequence of independent Bernoulli trials, which in our case are the independent laser shots. So, the Binomial RFS is an extension of the Bernoulli RFS, where its cardinality distribution is described by the binomial distribution with parameters m (number of binary experiments), n (number of successes) and r (the probability of success of each trial). It is defined by :

$$\rho(n) = \binom{m}{n} r^n \; (1-r)^{m-n} \tag{7.16}$$

Let m be the number of laser shots fired towards a target and $E = \{e_1, e_2, ..., e_n\}$ be a Binomial RFS containing our set of n returned echoes' distances from frustum detections, the PDF of E is :

$$P(E|r,\theta) = \binom{m}{n} r^n (1-r)^{m-n} \prod_{i=1}^{n} L_\theta(e_i) \tag{7.17}$$

In this case, the vector of parameters θ is extended with the success probability r.

Contrary to the Poisson distribution, the binomial distribution does not account for random arrival of objects and rather assumes a fixed number of objects and focuses on the probability of object detection in each trial. Furthermore, the mean and variance properties of the binomial distribution do not have the same constraint as in the Poisson distribution. The mean of a binomial distribution is described by $\mu = mr$, while the variance is given by $\sigma = mr * (1-q)$. Fig.7.6 shows Binomial distributions plotted with different values of success rate and 10 trials.

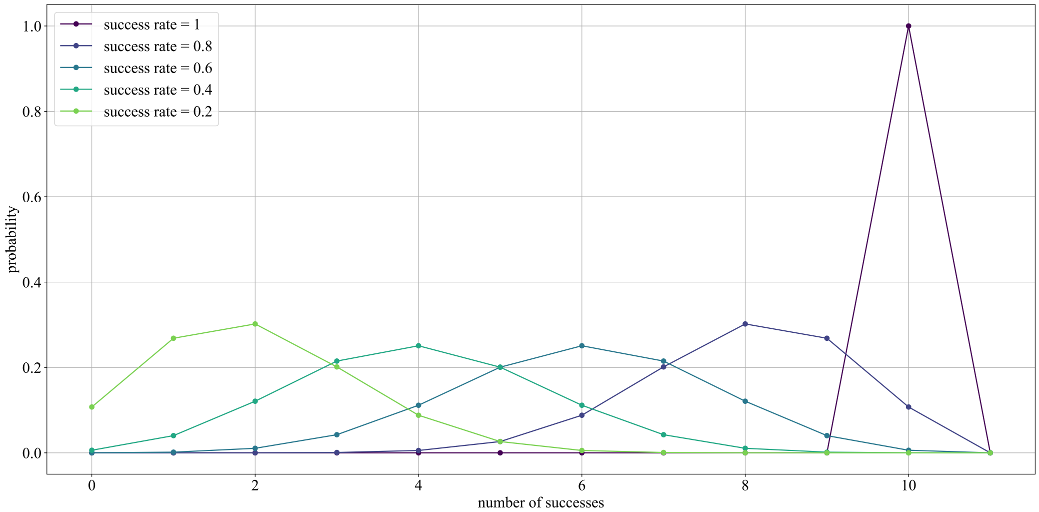

Figure 7.6: Binomial distribution plotted with various success rates and 10 trials.

7.3.3 Training

The training process of the model, incorporating the RFS formalism, follows a similar approach as before, utilizing the MCMC method. However, there is a notable distinction in that an extra parameter is introduced, namely the Poisson rate l or the success probability r for the Poisson and Binomial RFS models, respectively. Consequently, the Gaussian proposal density employed for proposing samples now possesses a 3D dimension. Figures 7.7 and 7.8 show the mean +/- std values of the trained parameters according to the visibility class for each RFS model and DVE type.

Note that these results are drawn using the Gamma distribution as likelihood distance distribution but using the Log-normal globally produces the same results. First of all, one can see that the parameter values are very similar for both RFS models. Although the Poisson and Binomial distributions differ by definition, in the end, they are both suited to capture the probabilities of receiving LiDAR returns. Finally, they share the same behavior as they are similarly influenced by the number of echoes in the frustums. However, the std values observed for the Binomial RFS model are globally higher than the ones of the Poisson model, for all cases, indicating a greater spread in the values of the trained samples. This means that the convergence of the MCMC learning algorithm is relatively less efficient for this model as it struggles to reach the underlying parameter distribution. As a result, the discrimination capacity between visibility classes with via the success parameter r is expected to be reduced.

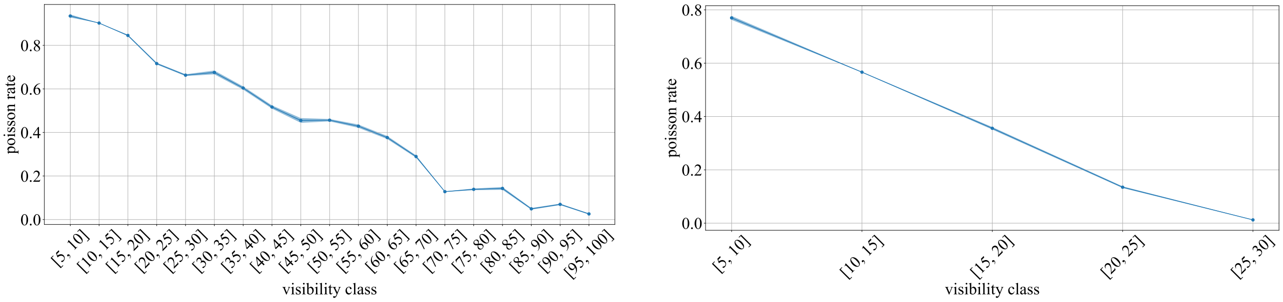

(a) Big fog

(b) Small fog

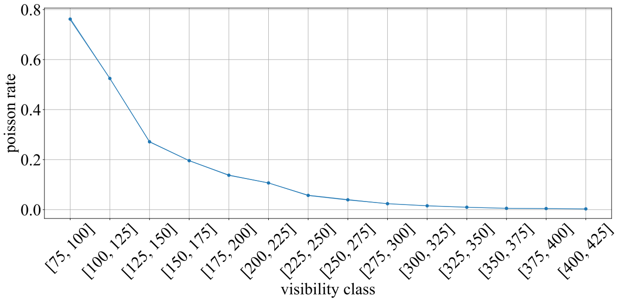

(c) Smoke

Figure 7.7: Mean Poisson rate trained with MCMC +/- std for each DVE type

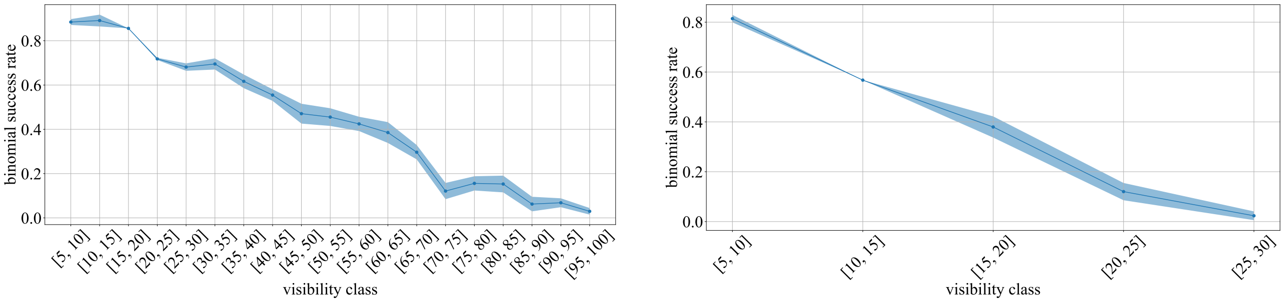

(a) Big fog

(b) Small fog

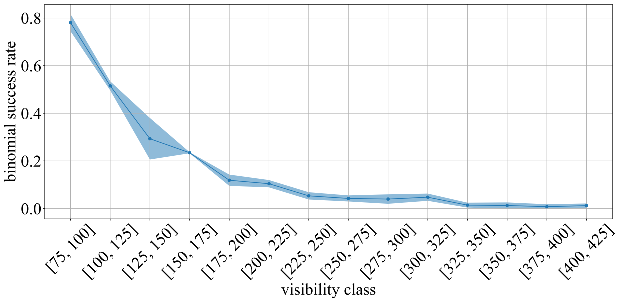

(c) Smoke

Figure 7.8: Mean Binomial rate trained with MCMC +/- std for each DVE type

7.4 Inference results

Our Naive Bayes model is designed to infer visibility classes from point clouds, leading to a classification problem. Note that, to guarantee its performances, our dataset is split into two parts : 2/3 for training and 1/3 for testing. After learning model parameters within the training set, the performance of the model can then be evaluated using the test set. Following the framework displayed on fig.7.3, when echoes are received, each class is tested using the sampled learned model parameters. Formally, $P(V_n|E, \mathbf{Y})$ is computed for each visibility class V_n. The class showcasing the highest probability value is denoted as the inferred class. In order to assess the performance of our solution across the entire dataset, we analyze each visibility class in the test set and aim to accurately infer the correct class based on the corresponding distance distributions.

Employing discrete classes of visibility potentially introduces errors in the differentiation of the distributions. When segmenting with coarse classes, the data can be divided too roughly, resulting in a failure to capture the continuous evolution. However, although using discrete classes of visibility introduces the risk of errors in distribution differentiation, we have decided to maintain this approach given the resolution of our visibility ground truth (1 m) and the relatively small dataset size even in the discrete state space. To enable the inference of continuous visibility, one could repeat the acquisitions and increase the data size.

In order to test the algorithm, it is necessary to define the number of LiDAR scans within each data batch E. Increasing the number of data points is likely to enhance the inference results, as it provides a more representative statistical sample of both number of points and distance. However, this number can also have various effects on the results. Firstly, it influences the time required to collect the data before performing the inference. One may prioritize a faster estimation of visibility, with a short data collection time but in practice, it is likely that the inference can be performed every several seconds without significantly altering any future decisions. Moreover, a higher number of scans increases the chance of collecting sufficient data points for inference. Given that our dataset exhibits some non-homogeneity in terms of available data for each visibility class (explained just below), setting a high number of scans could result in fewer inference tests for a particular class. Conversely, too few LiDAR scans may lead to an insufficient number of data points. To begin with, we have opted for 5 LiDAR scans. This choice is appropriate in terms of the number of points within each batch and the number of batches within each visibility class. Additionally, since the Ouster LiDAR has a scan rate of 10 Hz, this configuration allows for the potential generation of an inference result every 0.5 s. While this time-frame is relatively fast and may yield sub-optimal inference results, we have deliberately chosen it to assess the performance of our solution under a rather constraining configuration.

Confusion matrices and the Root Mean Square Error (RMSE) metric are employed as performance evaluation measures. Confusion matrices are often used to depict the results of supervised learning algorithms. Each row and column represents respectively the instance of a ground truth class and the inferred class and, each cell contains the probability of inferring the corresponding combination. As opposed to the RMSE metric, which gives an average estimate of the classification performance, the confusion matrix gives more insight on the classification of each individual class. As a result, the investigation of classification errors can be approached precisely with the specific classes. Figures 7.10, 7.11 and 7.12 give, in the form of confusion matrices, the classification results for each combination of likelihoods on the whole dataset and each DVE condition. In certain cases, the model may produce a null probability for inferring the correct class. These failures should be investigated closely to identify the underlying causes and constitute potential areas for improvement.

However, given the low amount of data available for this study (explained just below) and the inference of neighboring classes in these worst cases, the proposed methods still remain relevant. Considering the smoke results, fig.7.12, the improvement of classification performance given by the RFS models is clearly observable, but several discontinuities are seen for the classes at high visibility. For instance, it can be seen that the model struggles to classify the last 3 visibility classes (corresponding to visibility values from 85 m to 100 m). This is probably caused by the lack of data points at this stage of visibility.

The RMSE is computed from the confusion matrix using the following equation:

$$RMSE_V = s\sqrt{\sum_{i=1}^{N}(V_i - \hat{V}_i)^2 \;/\; N} \tag{7.18}$$

where V_i is the ground truth visibility class id, $\hat{V}_i$ is the inferred class id, s is the visibility step of each class (5 m for fog tests and 25 m for smoke) and N is the number of classes. For each likelihood model (Gamma or Log-normal) and DVE type, the resulting $RMSE_V$ is given in table 7.1.

DVE conditions Likelihood model		Small fog	Big fog	Smoke
No RFS model	Gamma	1.8 m	6.4 m	85.3 m
	Log-normal	1.8 m	6.3 m	69.4 m
Poisson RFS	Gamma	1.5 m	5.4 m	39.0 m
	Log-normal	1.4 m	4.9 m	39.4 m
Binomial RFS	Gamma	1.7 m	6.0 m	51.6 m
	Log-normal	1.3 m	5.4 m	42.2 m

Table 7.1: $RMSE_V$ results for all likelihood combinations and DVE conditions, with a 0.5 s data collection time-frame. The best combination for each DVE condition is highlighted in green.

Fig.7.9 shows the number of 3D-LiDAR scans within the test dataset of fog and smoke experiments (representing 1/3 of the total dataset) and thus the number of LiDAR points available for the classification differ for each class. The classification results are mitigated by the non-uniformity distribution of data across visibility classes in the dataset. Despite this non uniformity, our model remains robust and demonstrates relatively good performances in fog conditions experiments using the 1 m resolution of the transmissiometer and the 5 m range of our visibility classes. The small fog test exhibits the best results in terms of RMSE but has fewer classes with relatively high number of available LiDAR scans compared to the big fog test. In contrast, the lower performances of the smoke tests can be explained by the properties of the visibility classes and the erratic nature of the distributions inside the dataset, already mentioned in the experimental results in chapter 6.

However, low number of available scans to test the model does not always result in low classification results. For example, classes $[5, 10]$ and $[25, 30]$ for small fog show a low number of scans but are perfectly classified by the model. Conversely, classes disposing of large amount of data files are not necessarily correctly classified. This may be explained by clear differences in the distance histograms from one class to another or by the discretization of the visibility state space.

Finally, the Log-normal distribution seems to hold the best performances for modelling the distance distributions of the frustum echoes, as it presents equal or better $RMSE_V$ results than the Gamma distribution in all cases. There is no clear preference between the Poisson or Binomial RFS models, except maybe a slight advantage for the Poisson RFS model and especially for the smoke experiments. Overall, the RFS extension greatly improves the overall classification results.

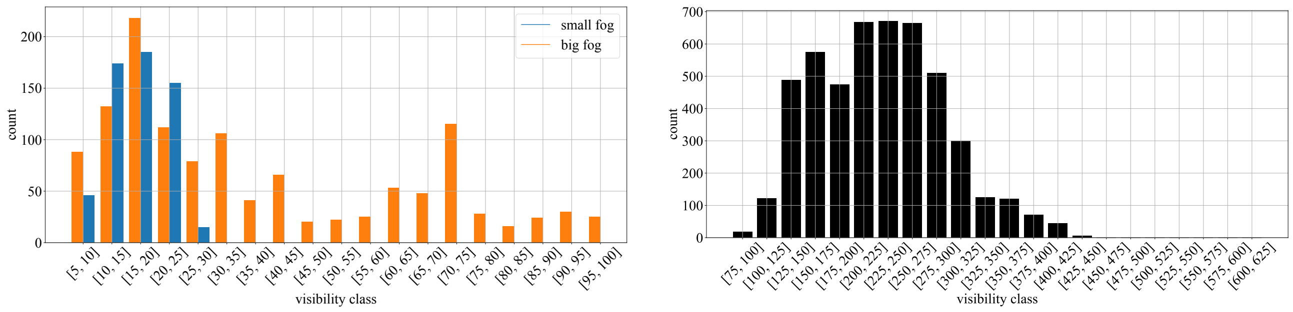

(a) Fog tests

(b) Smoke

Figure 7.9: Number of LiDAR scans used for the test phase for each visibility class, representing 1/3 of the total dataset. The number of scans per visibility class is not homogeneous.

Last, we present the results using a larger rate of data collection of 2 s, presented on table 7.2. Increased classification performance is observed for all model combinations and DVE types, except for the Binomial RFS and Log-normal in smoke. As expected, using more data as input enhances the inference results. In terms of operation, it is likely that having a measurement of visibility every 2 s is enough to take potential decisions while keeping track of the evolving conditions. Table 7.2 also shows that this time, other model combinations exhibit the best results. The Poisson RFS seems to be the most appropriate model, using the Log-normal distance distribution in small fog and the Gamma distribution for big fog and smoke. The results from small fog even show a $RMSE_V$ of 0.9 m, which is below the ground truth sensor resolution.

DVE conditions Likelihood model		Small fog	Big fog	Smoke
No RFS model	Gamma	1.4 m	4.3 m	56.2 m
	Log-normal	1.5 m	3.6 m	59.3 m
Poisson RFS	Gamma	1.2 m	3.2 m	34.4 m
	Log-normal	0.9 m	3.6 m	36.6 m
Binomial RFS	Gamma	1.4 m	3.7 m	49.6 m
	Log-normal	1.2 m	4.0 m	47.3 m

Table 7.2: $RMSE_V$ results for all likelihood combinations and DVE conditions, with a 2 s data collection time-frame. The best combination for each DVE condition is highlighted in green.

Chapter 8 discusses about potential improvements for the inference model. Through these extensions, the given examples will compare the 0.5 s and 2 s times of data collection and use only the Poisson RFS and Gamma model, for simplicity.

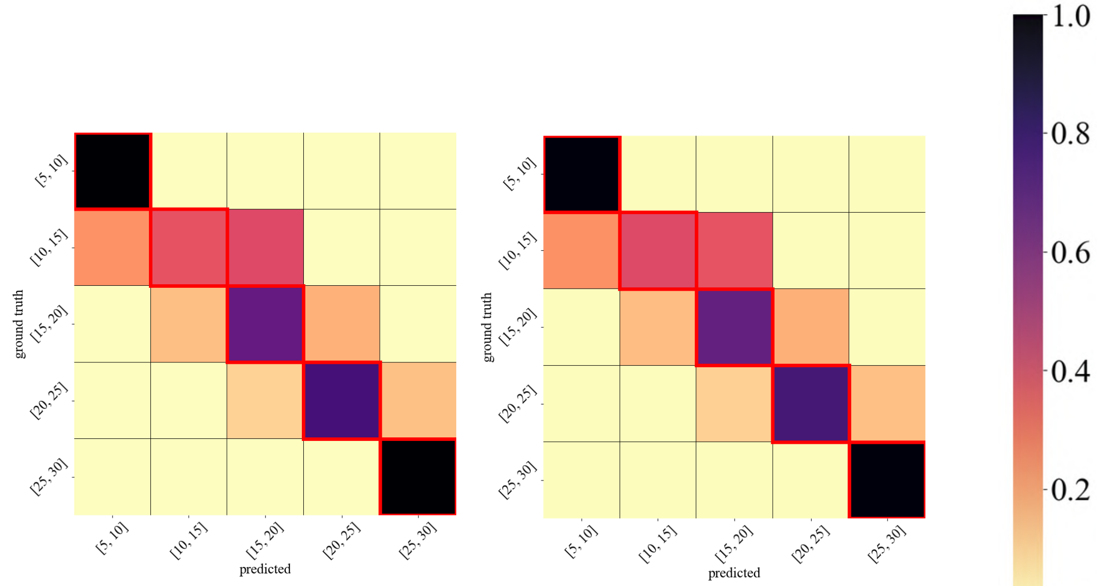

(a) Gamma (b) Log-normal

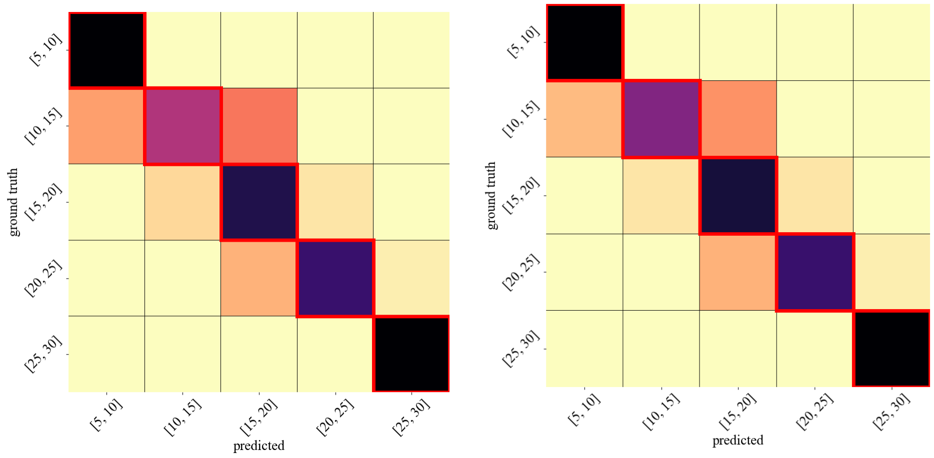

(c) Poisson RFS + Gamma (d) Poisson RFS + Log-normal

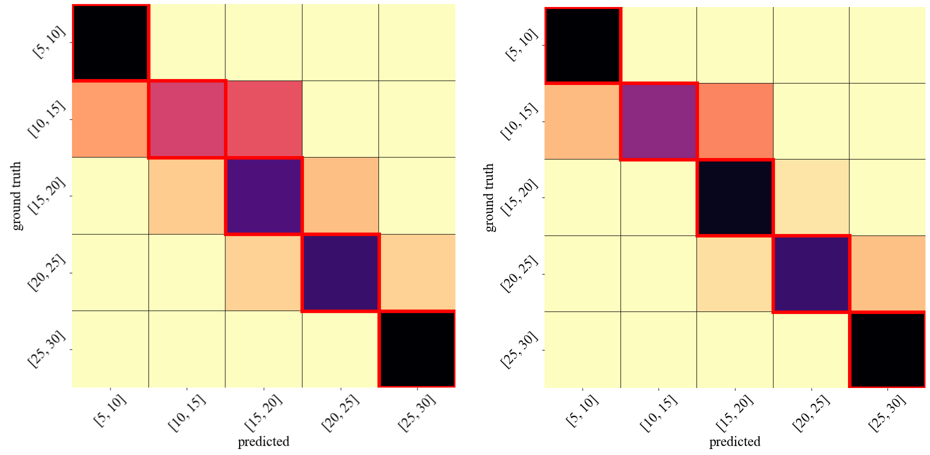

(e) Binomial RFS + Gamma (f) Binomial RFS + Log-normal

Figure 7.10: Confusion matrices for the small fog tests, with a 0.5 s data collection time.

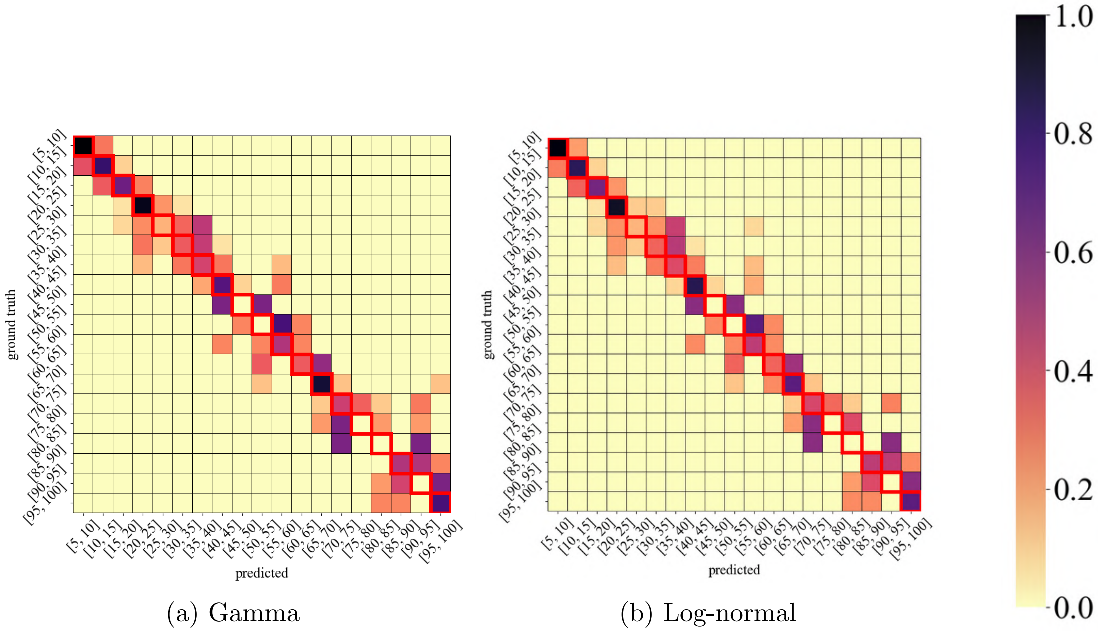

(a) Gamma (b) Log-normal

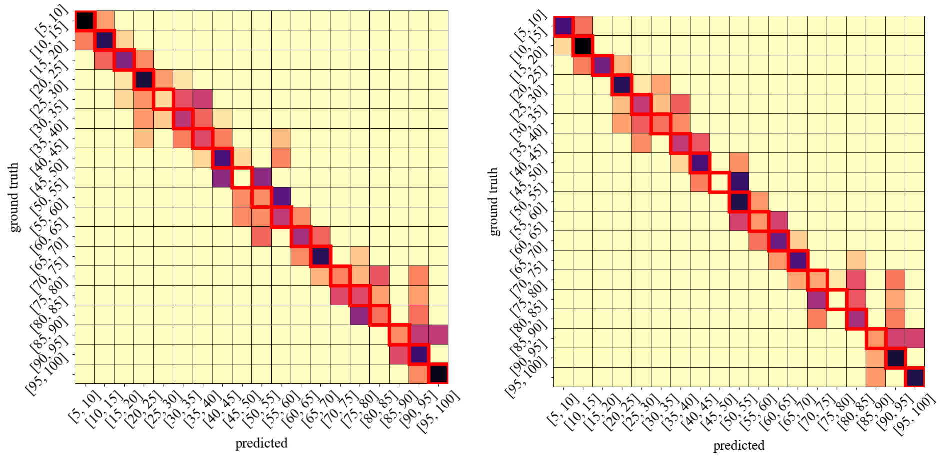

(c) Poisson RFS + Gamma (d) Poisson RFS + Log-normal

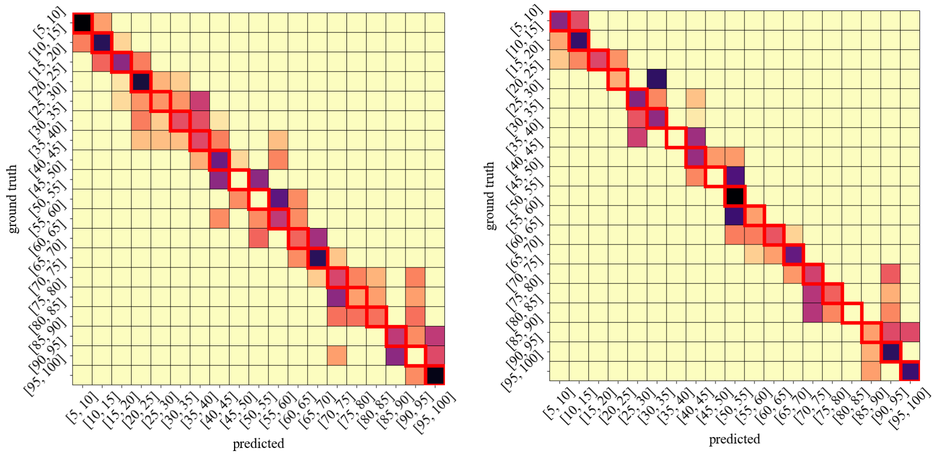

(e) Binomial RFS + Gamma (f) Binomial RFS + Log-normal

Figure 7.11: Confusion matrices for the big fog tests, with a 0.5 s data collection time.

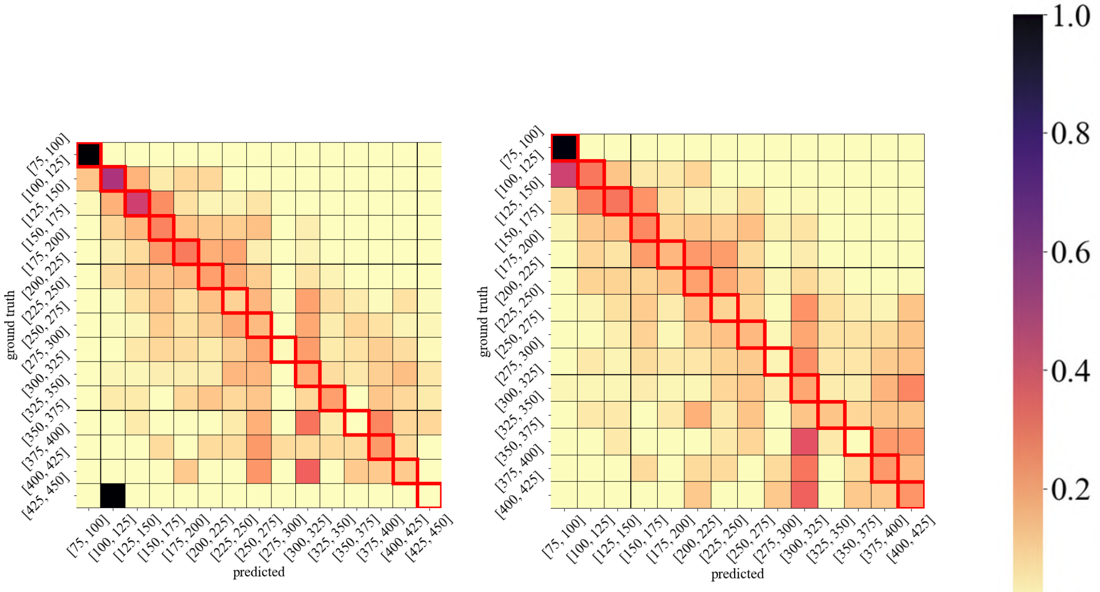

(a) Gamma (b) Log-normal

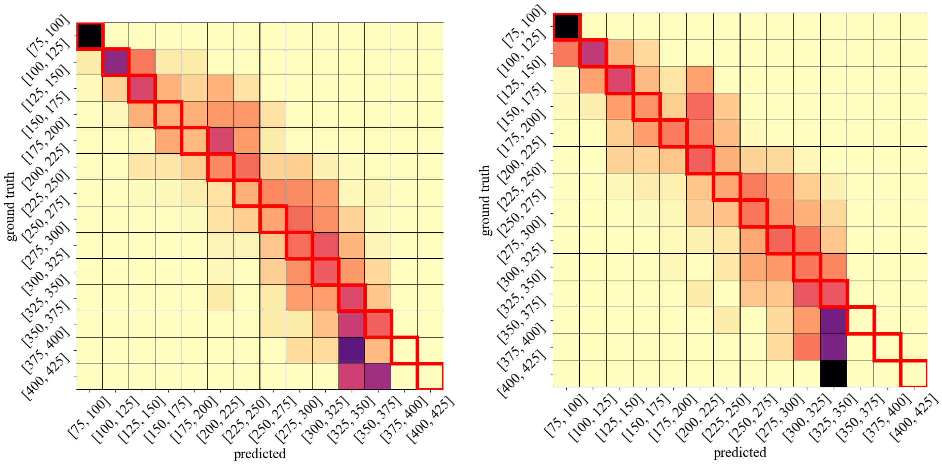

(c) Poisson RFS + Gamma (d) Poisson RFS + Log-normal

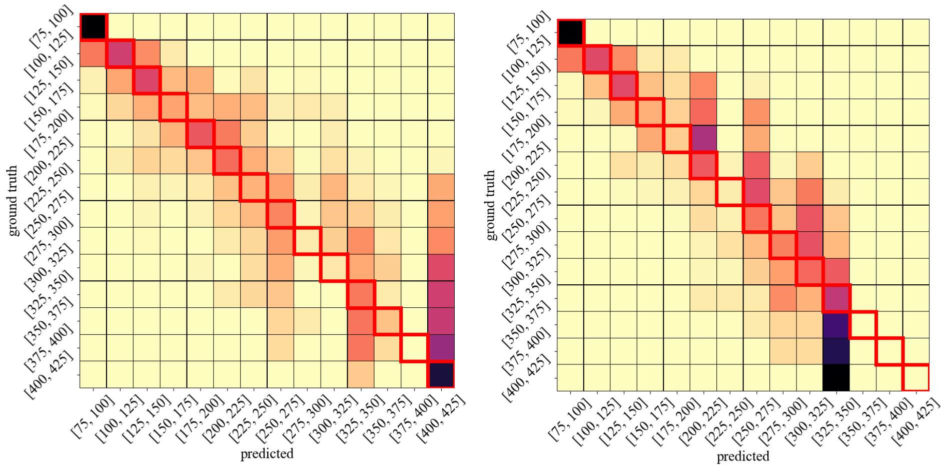

(e) Binomial RFS + Gamma (f) Binomial RFS + Log-normal

Figure 7.12: Confusion matrices for the smoke tests, with a 0.5 s data collection time.

7.5 Conclusion

This chapter presents the inference model developed during this thesis, which aims to classify visibility using point clouds received in fog and smoke conditions. It is a Naive Bayes model built upon the distance distributions of these points and their corresponding visibility values. The results of the classification, considering the datasets used to train the model, are promising, with global RMSE values relatively satisfying. Both Gamma and Log-normal distributions are adequate to model the distance distributions, although the Log-normal one slightly outperforms the other one. The RFS formalism greatly improves the classification results, by statistically capturing the number of detections within the distributions. We consider our inference model to be a significant step forward in addressing the DVE challenge for autonomous driving, thanks to its promising performances and its compliance with operational conditions.

However, in order to enhance the reliability and accuracy of the classification results, it is imperative to conduct additional experiments and gather more data. By repeating the acquisitions with the same test protocol and increasing the dataset size, the training process can be strengthened, leading to improved classification outcomes. This approach will contribute to the robustification of the training phase, ultimately enhancing the overall performance and accuracy of the classification model.

The next chapter develops alternative ways to improve the performances of our approach. Potential applications of the inference model are presented, with the development of a Bayesian filtering technique and the its utilization to differentiate the DVE types. The extension of the model to other variables is also presented, using inference or correlation. Ultimately, long term improvements are proposed, considering multi-echo and the applicability of our method to outdoor conditions.

Chapter 8

Extensions

In progress...

This chapter presents additional information, useful to improve our approach in the future. These concepts are given for demonstrative purpose, as for each of them, we only present examples to highlight its feasibility, showing that they constitute short or long term future work.

First, we present short term extensions of the inference approach. A recursive estimation technique is developped, similar to a classical position algorithm, which prevents the classification of excessive visibility considering prior measurements. And, provided that datasets are similar for each DVE type, the model is applied to differentiate the types of DVE. Additionally, the application to other variables is given. Through inference or correlation, retrieving other variables than visibility might more directly usable in an AD software perspective. Finally, long term extensions are discussed, concerning the multi-echo information and the application of our approach with outdoor conditions.

For all presented examples where we show inference results, the 2 s data collection time is used as well as the Poisson RFS and Gamma model each DVE type, as discussed in section 7.4.

8.1 Direct extension of the inference model

8.1.1 Recursive estimation of visibility

Our inference model can be presented as a batch estimation method. Echoes are received and used as input to infer a most probable visibility class, but the model does not take into account the previous result to improve the new one if one wishes to repeat the process. To achieve this, the development of a recursive estimation method is needed. And following our Bayesian approach, we delve into the concept of Bayesian filtering [197].

Visibility Discrete Bayes filter

This section details the development of a Bayesian recursive filter applied to our visibility inference approach. Our state space is defined by our finite discrete visibility classes, leading to an estimation on a 1D space and the development of a Discrete Bayes filter.

Let the variables x_i and x_k denote individual states. The idea of this algorithm is to predict and update, at each measurement step, the probability value of each state of the state space.

Algorithm 1 Discrete Bayes filter

```
for all $k$ do
    $\bar{p}_{k,t} = \sum_i p(X_t = x_k | X_{t-1} = x_i) p_{i,t-1}$        ▷ Prediction
    $p_{k,t} = \eta p(z_t | X_t = x_k) \bar{p}_{k,t}$        ▷ Update
end for
return $p_{k,t}$
```

In the prediction step, $p(X_t = x_k | X_{t-1} = x_i)$ is the probability of the system transitioning between any two arbitrary states x_i to x_k in one time interval Δt, through the application of a transition model which needs to be defined. Then, the update step updates the states considering a measurement z_t. In our case, the result of the inference model is treated as a measurement and consists of a vector containing a probability value for each class (or bin) of visibility v. So, for each visibility state, the measurement probability $p(z_t | X_t = x_k)$ is multiplied by the predicted probability $\bar{p}_{k,t}$ and the normalization factor η. A measurement model that associates uncertainties to the measurement has to be defined.

To build the transition model, we use similar dynamics than a position estimation problem, with the visibility velocity defined as

$$\dot{v}_{i \to k} = \frac{v_k - v_i}{\Delta t} \tag{8.1}$$

For each transition possible between visibility states, a required velocity value can be calculated. And a probability value is attributed to these velocity values with a Normal distribution centered on the velocity value needed to transition from one state to another and a standard deviation.

$$p(X_t = x_k | X_{t-1}) = \mathcal{N}(\dot{v}_{i \to k} | \mu_{\dot{v}}, \sigma_{\dot{v}}) \tag{8.2}$$

The measurement model is finally defined similarly, using a Normal distribution, centered on the state k.

$$p(z_t | X_t = x_k) = \mathcal{N}(z | x_k, \sigma_z) \tag{8.3}$$

Estimation results

This section showcases the visibility estimation results obtained with this recursive technique. Table 8.1 presents the different $RMSE_V$ results and figures 8.2 and 8.3 enable the visual comparison between the batch and recursive visibility estimations along with the ground truth. Red points correspond to the batch visibility estimation, where the LiDAR's batches of scans are used independently to infer visibility while the temporality is not taken into consideration. These results constitute the measures z of the recursive state estimation described previously. In black, results from the recursive estimation are shown, exhibiting a better alignment with the ground truth evolution of visibility (in blue) Parameters $(\sigma_{\dot{v}}, \sigma_z)$ are defined individually for each DVE type. Their values are set manually as to optimize the global $RMSE_V$ value.

The addition of the presented recursive model enables the enhancement of the visibility estimation, except for the small fog test where the results are the same. Indeed, with a few classes, the batch estimation already showcases precise measurements with a $RMSE_V = 1.5$ m for the 0.5 s batch time and $RMSE_V = 1.2$ m for the 2 s and, the filter does not manage to improve the estimation results. On the contrary, the other DVE types show increased performances and for both batch times. The greatest improvement is observed in the smoke tests, where batch measurements exhibit significant noise, and the recursive estimation facilitates smoother classifications.

DVE type	Small fog ($\sigma_{\dot{v}} = 2, \sigma_z = 1$)	Big fog ($\sigma_{\dot{v}} = 2.5, \sigma_z = 1.4$)	Smoke ($\sigma_{\dot{v}} = 1, \sigma_z = 5$)
Batch estimation (batch of 0.5 s)	1.5 m	5.4 m	39.0 m
Recursive estimation (batch of 0.5 s)	1.5 m	3.9 m	27.6 m
Batch estimation (batch of 2 s)	1.2 m	3.2 m	34.4 m
Recursive estimation (batch of 2 s)	1.1 m	2.9 m	14.9 m

Table 8.1: Comparison of $RMSE_V$ between recursive and batch estimation techniques, using the Poisson RFS and Gamma model and different times of data collection.

These results are promising and constitute the basis for future work. A precise sensitivity analysis can be performed on the $(\sigma_{\dot{v}}, \sigma_z)$ parameters in order to find their optimal values. The recursive model can also be extended, by including the estimation of $\dot{v}$ and/or even $\ddot{v}$, a visibility acceleration. Additionally, the visibility class size could be reduced to lower values than 5 m for the fog tests and 25 m for the smoke ones. Reducing the class size is likely to introduce more noise in the inference results. However, this recursive technique has the potential to be robust and still yield favorable outcomes.

In conclusion, the development of improved datasets continues to be crucial for further advancements. Primarily, it enhances the likelihood of success in the aforementioned concept of shorter data collection times and possibly shorter visibility classes. Conducting experiments under fluctuating visibility conditions would be highly valuable to assess the performance of estimation methods in dynamic environments.

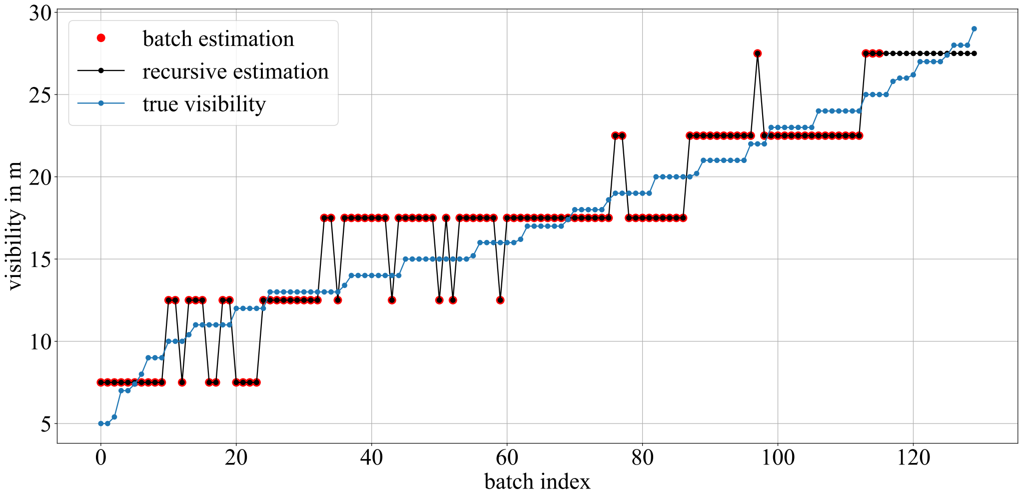

(a) Small fog

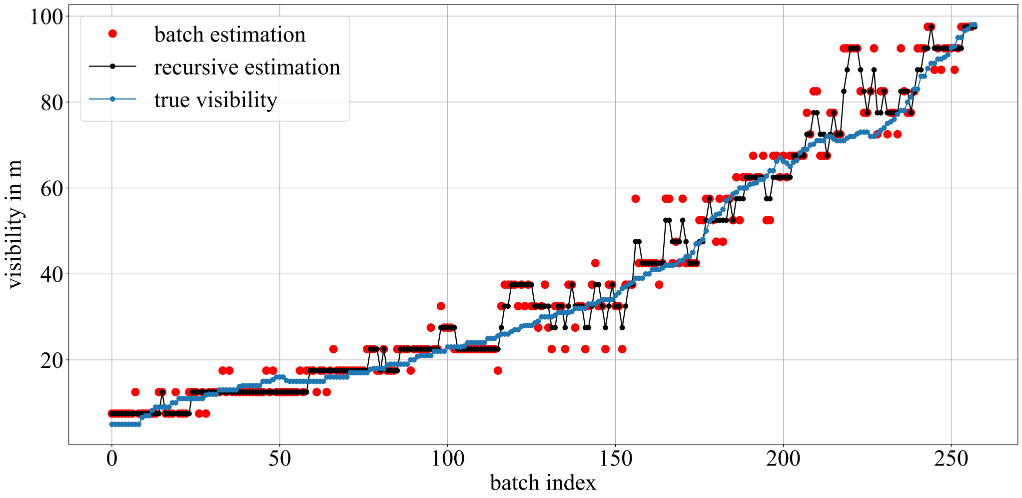

(b) Big fog

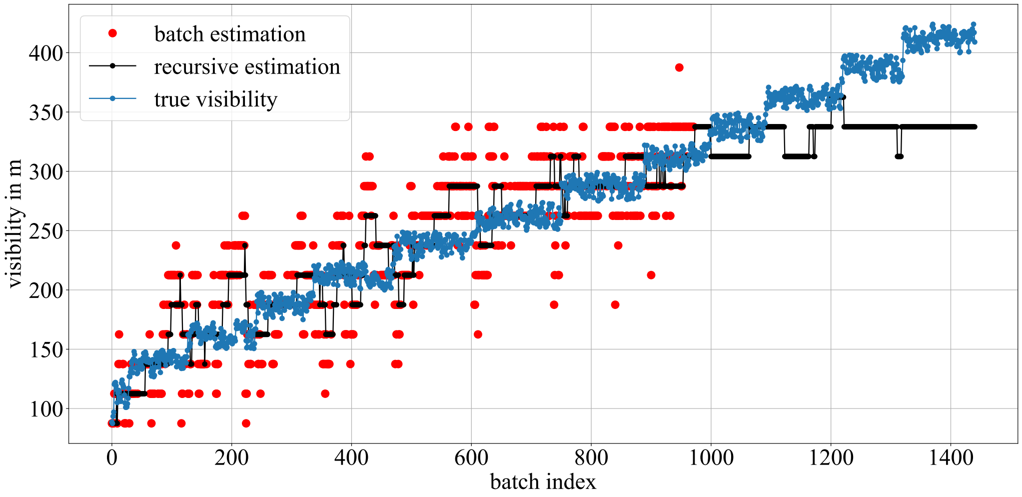

(c) Smoke

Figure 8.2: Recursive visibility estimation results compared to batch estimation and ground truth, using the Poisson RFS + Gamma model and batch data of 0.5 s.

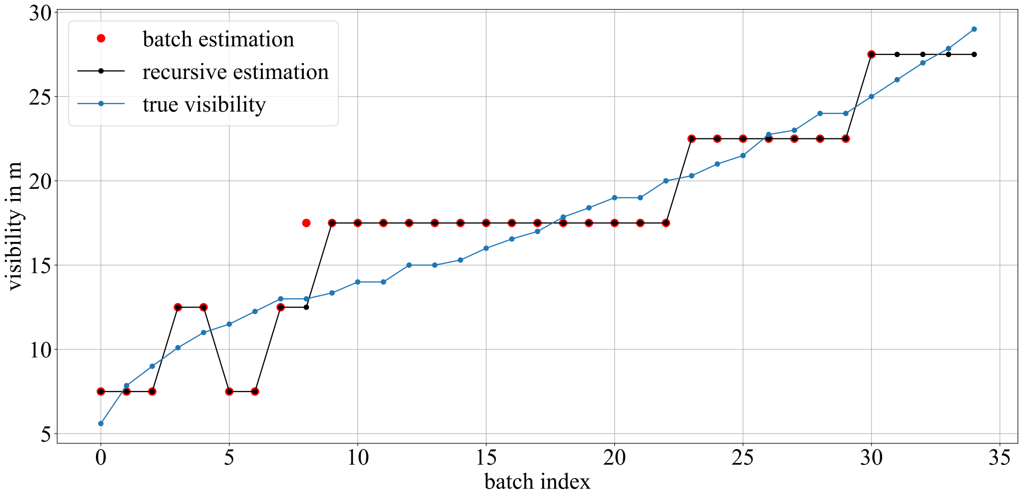

(a) Small fog

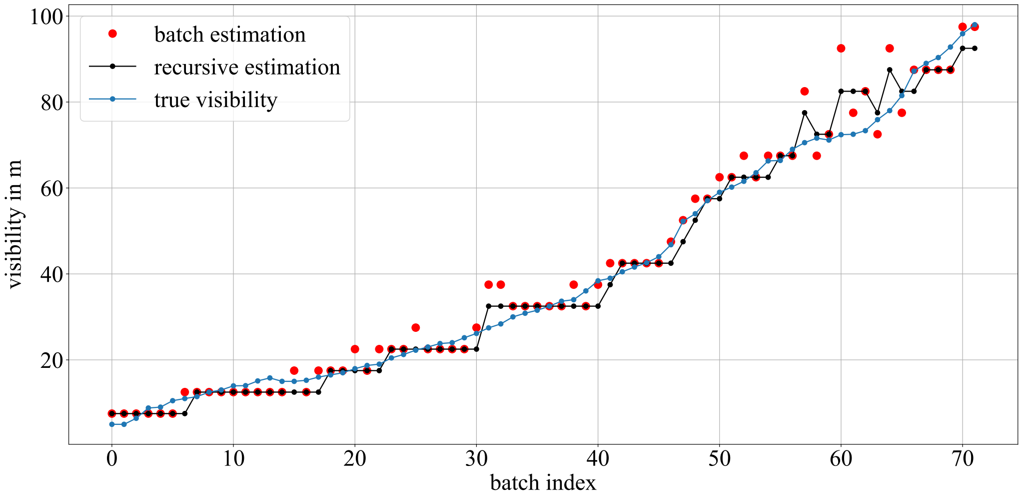

(b) Big fog

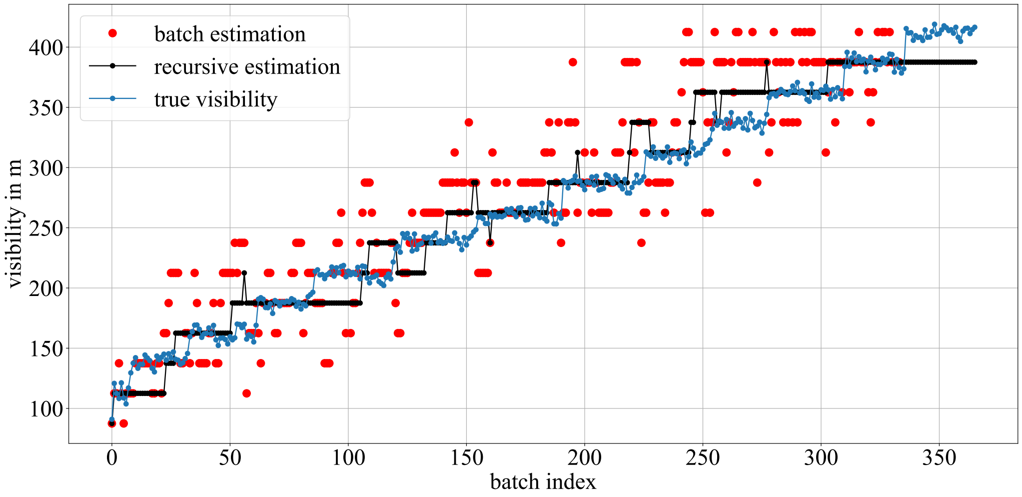

(c) Smoke

Figure 8.3: Recursive visibility estimation results compared to batch estimation and ground truth, using the Poisson RFS + Gamma model and batch data of 2 s.

8.1.2 Differentiating the DVE type

Distinguishing different types of environmental conditions within DVE scenarios is a key capability to enhance the performance of AVs. By accurately identifying the specific DVE type, an AV can potentially tune parameters of algorithms accordingly and potentially improve its behavior. We believe that the inference model presented in chapter 7 can be used to lay the foundations for such a weather differentiation system. In this section, we give an example on how we perform this differentiation, where the idea is to differentiate between the small fog and big fog types. Note that it is not possible to include the smoke conditions here because the data were not acquired similarly (e.g. *back* target frustum for the fog tests and *free-sky* target frustum for smoke).

The classes on which we test the inference model are now composed of the classes of each DVE type we want to differentiate. The previous dataset $\mathbf{Y}$ becomes the union of the datasets of small fog $\mathbf{Y_{small}} = \{(\mathbf{E_n}, \mathbf{V_n})\}_{\mathbf{small}}$ and big fog $\mathbf{Y_{big}} = \{(\mathbf{E_n}, \mathbf{V_n})\}_{\mathbf{big}}$, with $\mathbf{Y} = \mathbf{Y_{small}} \cup \mathbf{Y_{big}}$. To keep the example relatively simple, we use similar classes for each DVE type, so that the classes of each type are similar and $n \in [0, N]$ and $N_{small} = N_{big}$. Following the equations of the inference model detailed section 7.2.1, V, which was previously part of either $\mathbf{Y_{small}}$ or $\mathbf{Y_{big}}$ is now part of both together and, the computation of $f(E|\mathbf{Y}, V_n)$ now concerns all V_n in $\mathbf{Y_{small}} \cup \mathbf{Y_{big}}$. Through this computation, likelihood values are generated for each visibility class and contained in a vector of size $2N$. These likelihood values are then normalized by the total number of classes $2N$ to be converted into probabilities. Ultimately, comparing the cumulative sum of the probabilities associated with each type (small or big) enables the differentiation between the two types, for example, the calculation to retrieve the probability of being in small fog is

$$P(DVE = small) = \frac{\sum_{V}^{V_{small}} P(V|E, \mathbf{Y})}{2N} \tag{8.4}$$

To assess the effectiveness of this approach, each fog type is sequentially selected and tested. This sequential testing allows for the evaluation of the model's ability to differentiate between small and big fogs in both scenarios. Because the small fog has the smallest dataset in terms of frustum points, it is the one limiting the number of classes, defined as $Vclasses_{small} = Vclasses_{big} = [[5, 10], [10, 15], ..., [25, 30]]$

The method is illustrated using only Poisson RFS + Gamma inference model and, using data batches of 0.5 s. In each sub-figure, the bar plot on the left represents the probability of finding each type. And, the box-plot on the right details the occurrences of finding each class using the inference model. The box-plots show that sometimes multiple classes within the same type are chosen as the right visibility class (as shown in the inference model results of section 7.4) but our approach accumulates these findings into their common class.

Fig.8.3 shows the differentiation results testing the small fog dataset. The first class $[5, 10]$ produces the miss-classification of the DVE type because the fog type is classified at 80%. Although the $[5, 10]_{small}$ visibility class is often inferred, it seems that the model often confuses with the $[20, 25]_{big}$ class. The other classes show better results as the small fog type is always the most classified, reaching 100% of classification for classes $[15, 20]$, $[20, 25]$ and $[25, 30]$. Fig.8.3 shows the differentiation results testing the big fog dataset. This time, the approach works even better as the big fog is always the most classified type. Though, the same phenomenon is observed, where there seems to be a confusion between the $[5, 10]_{small}$ and $[20, 25]_{big}$. This confusion should be investigated.

To conclude, our inference model and approach enables the differentiation between different types of DVE. This method requires the dataset of each type to be in a common format. One should build all datasets similarly to enable the differentiation of multiple DVE types. Note the presented method is for illustrative purpose only, and future work is needed to, for example, enable differentiation between aerosols (fog, smoke) and precipitation (rain, snow).

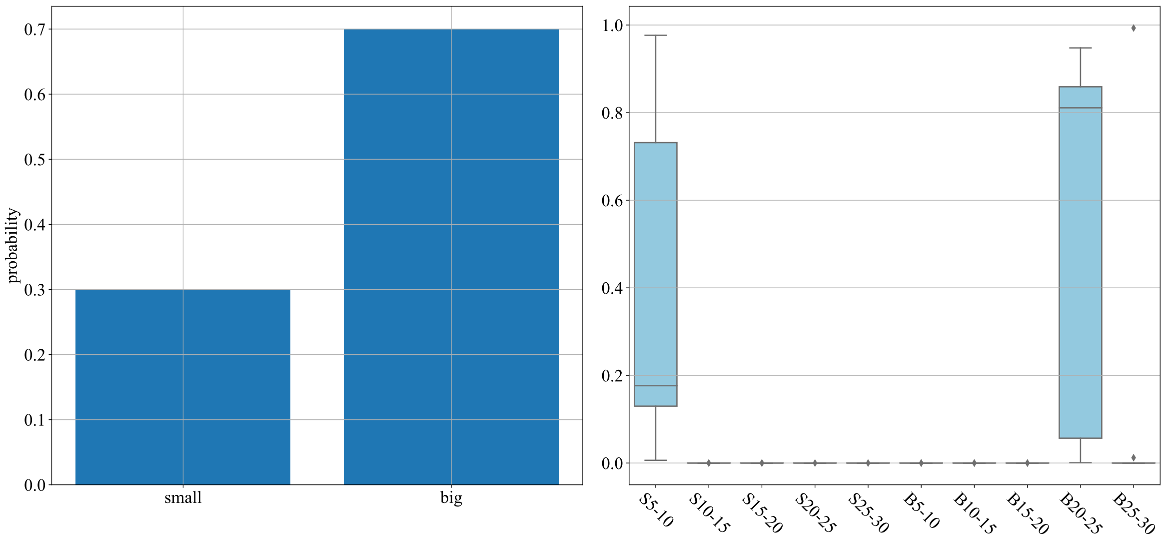

Target class : small [5, 10]

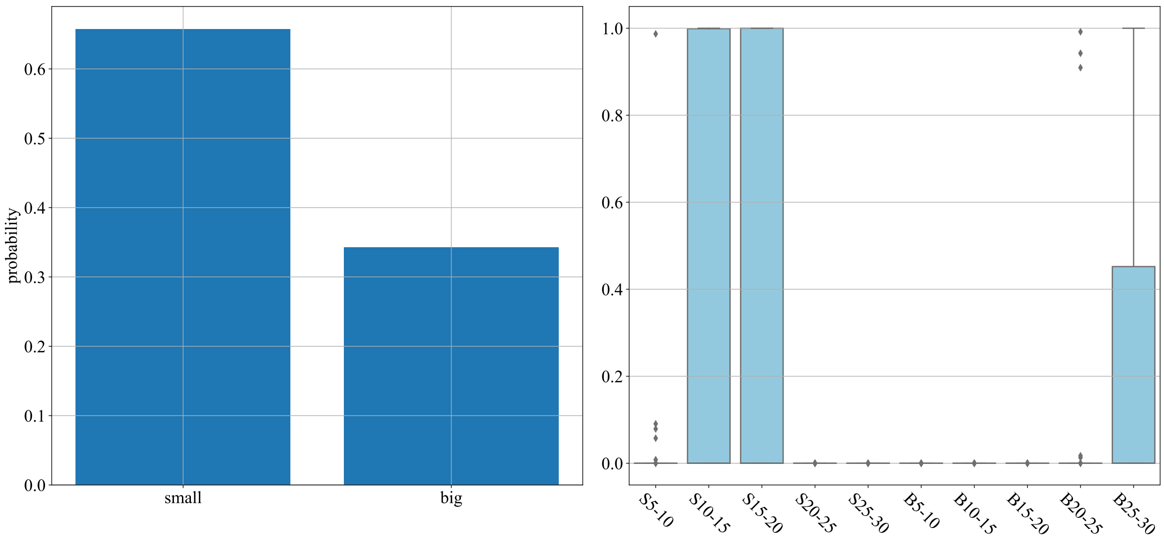

Target class : small [10, 15]

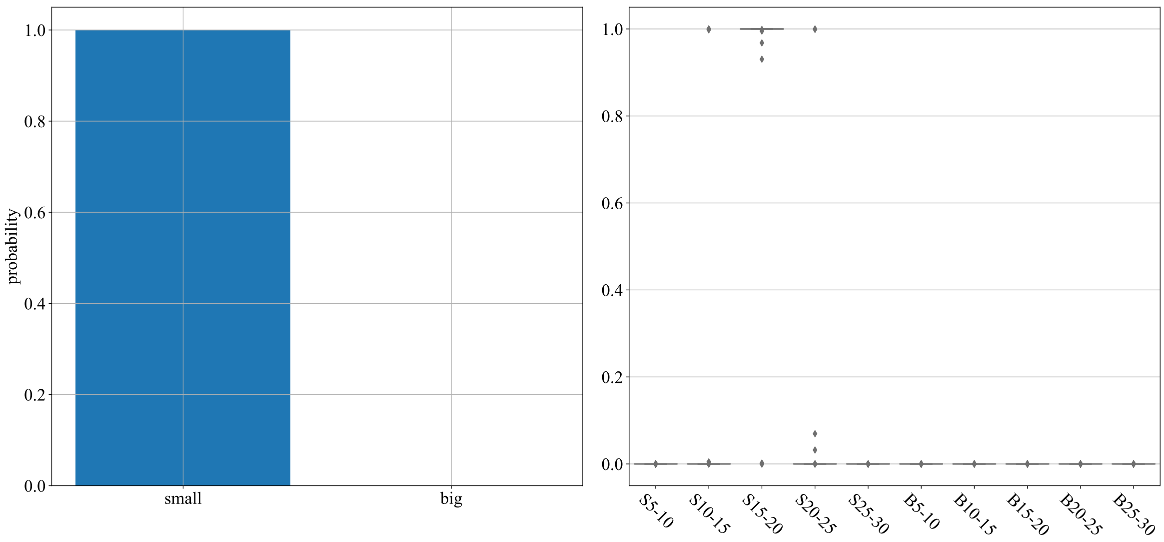

Target class : small [15, 20]

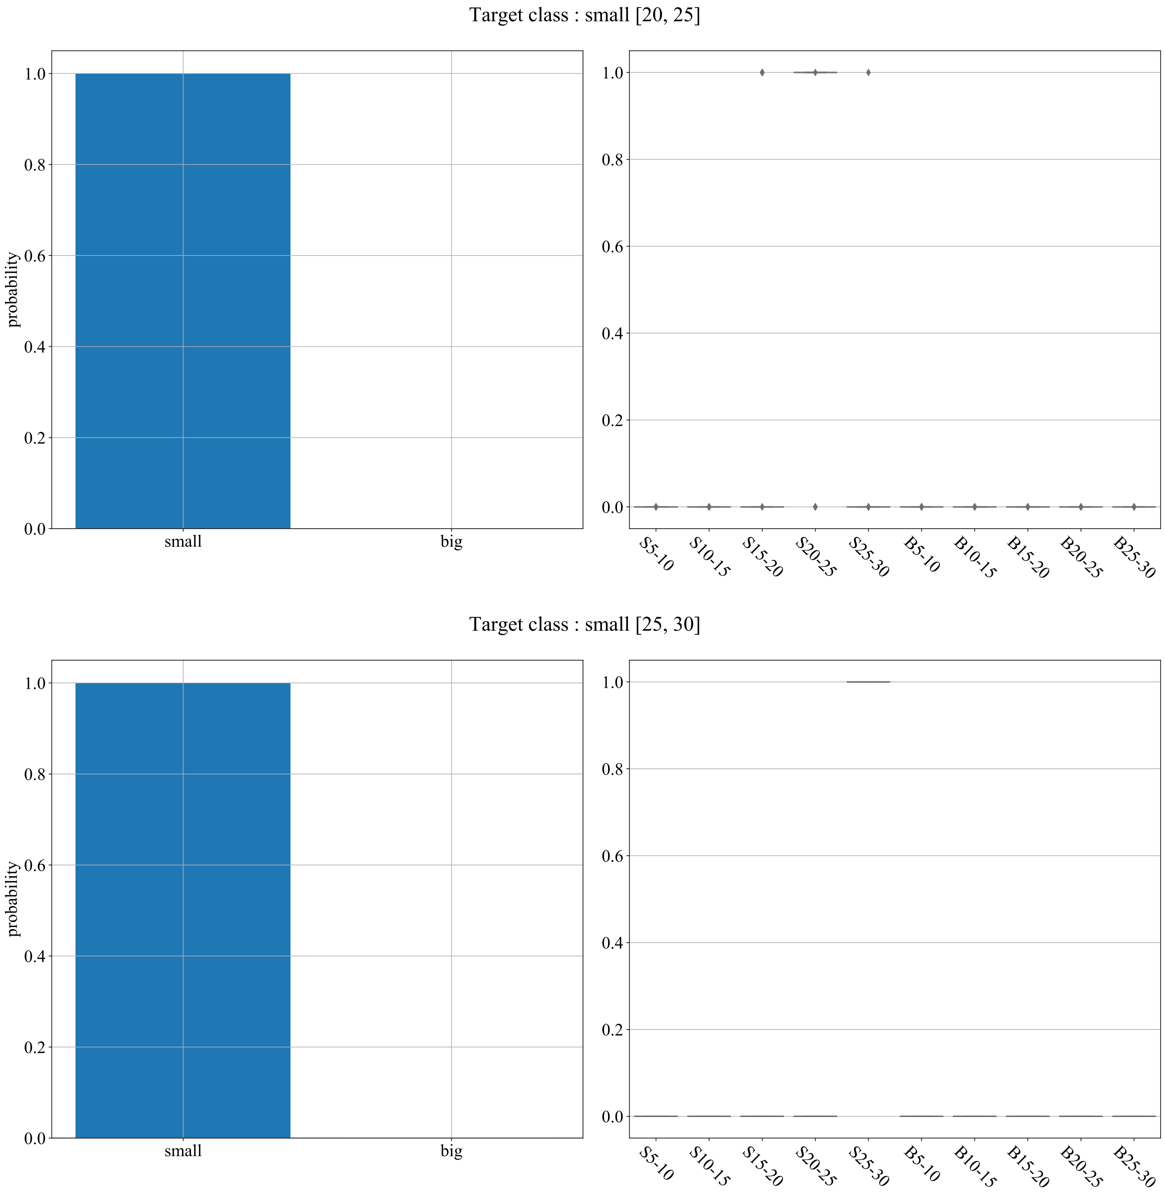

Figure 8.3: Differentiation results between small and big fog types, from the small fog dataset. The inference of visibility classes for the two types determines the probability of being in each of them.

Target class : big [5, 10]

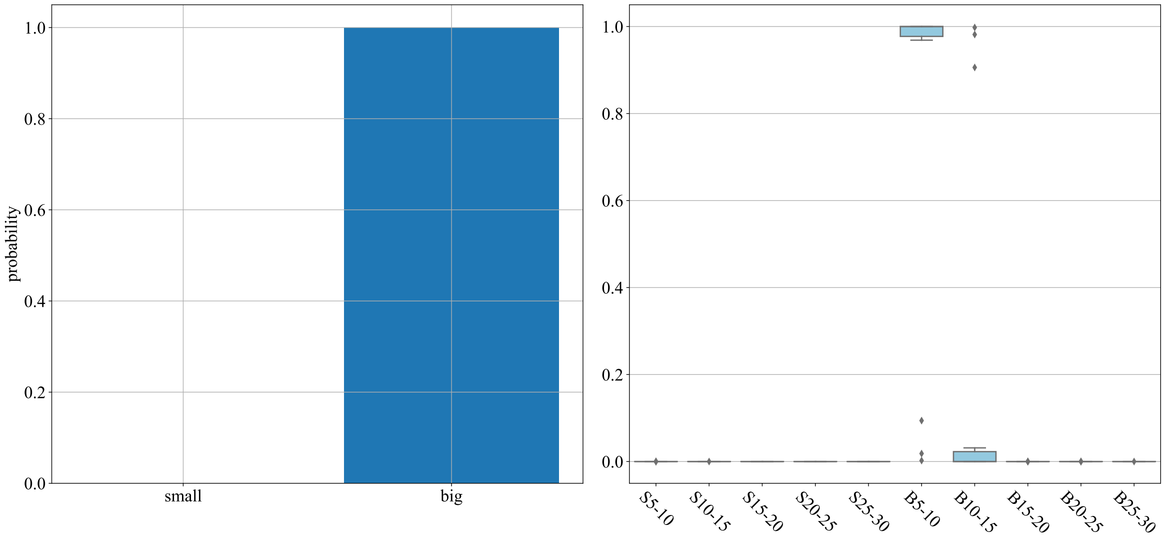

Target class : big [10, 15]

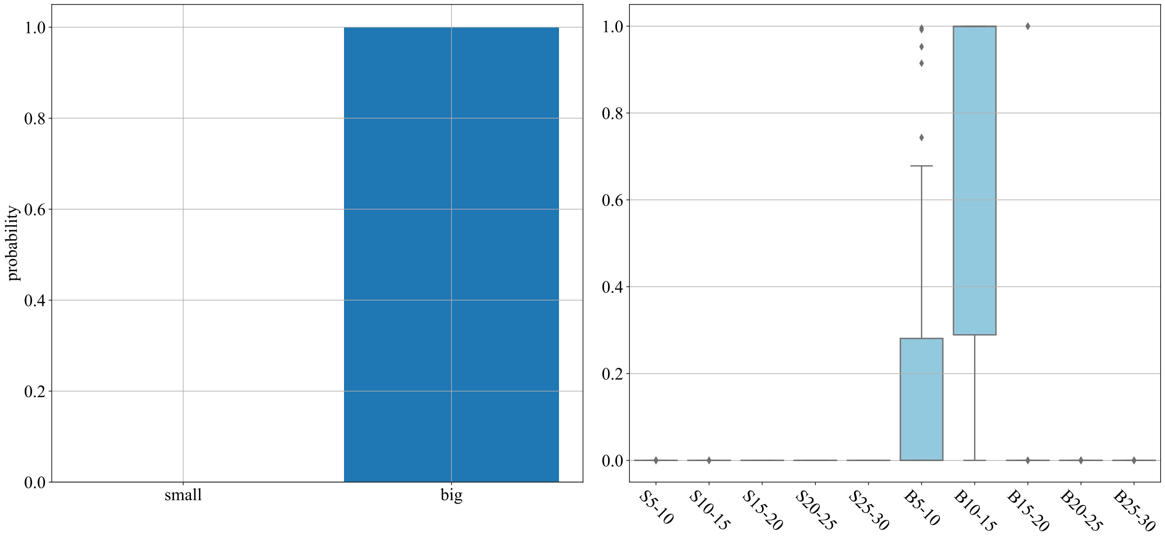

Target class : big [15, 20]

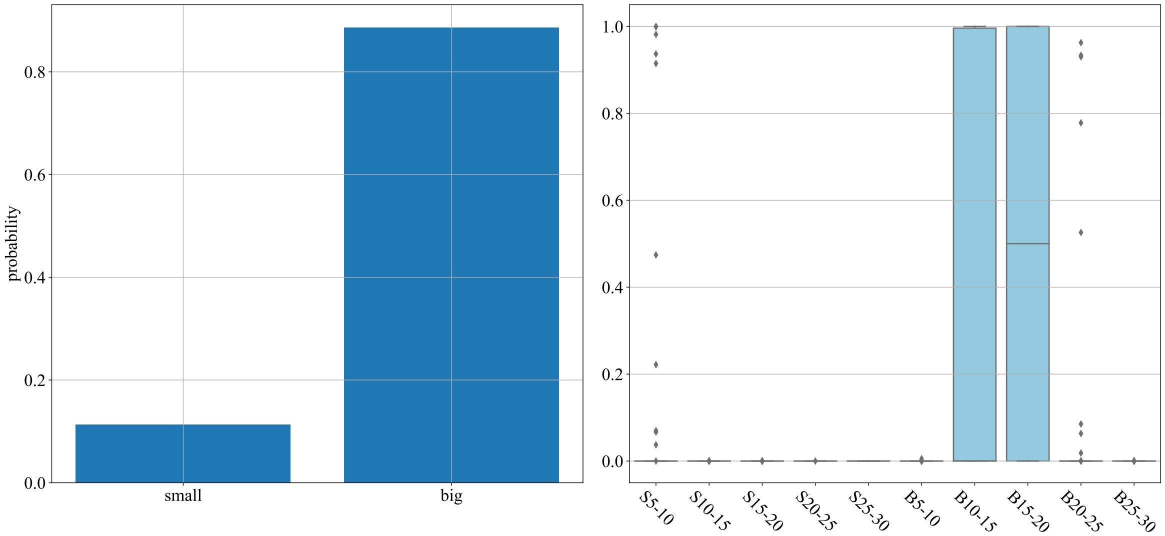

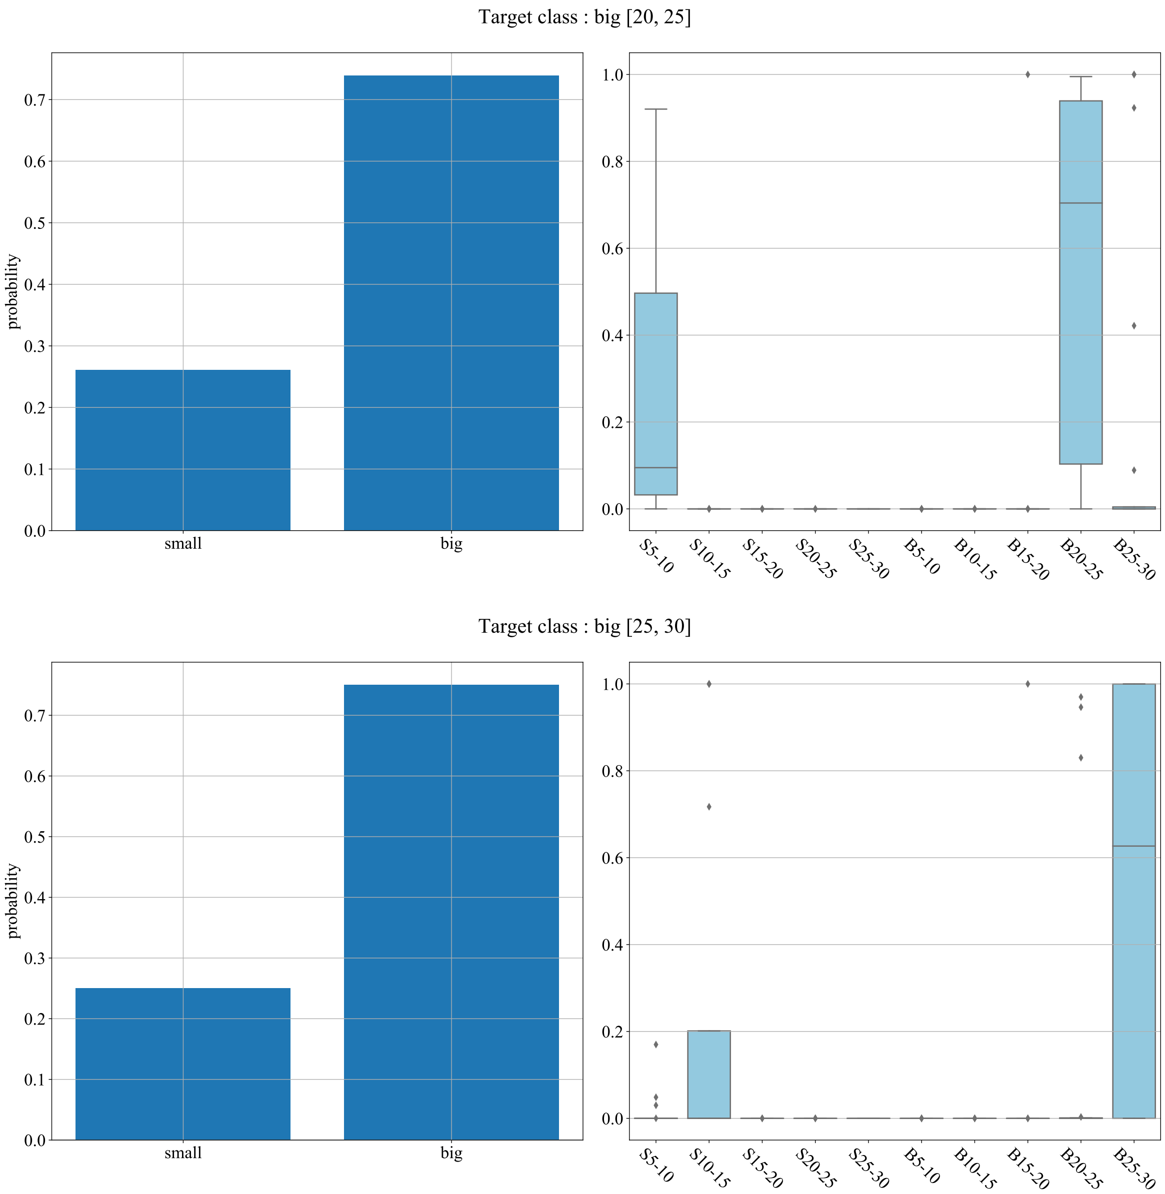

Figure 8.3: Differentiation results between small and big fog types, from the big fog dataset. The inference of visibility classes for the two types determines the probability of being in each of them.

8.2 Application to other variables

The main focus of this work was to emphasize on inferring visibility, which serves as a common metric for assessing the severity of visual degradation. Inferring its values from LiDAR point clouds is valuable as it can be further used by an autonomous vehicle. However, other variables than visibility may be of interest as they are also likely to improve the performance of an AV in DVE. For example, in terms of AV's operability, visibility is less meaningful than the sensors' capacity of detection. The sensor's POD holds the potential to facilitate the adaptation of the speed of an AV or parameters of perception algorithms, focusing on the sensor's performance rather than directly the visibility itself. This section show how the POD of targets can be retreived using within our approach, through inference or correlation.

8.2.1 Inferring a target POD

Replacing visibility for another variable does not lead to significant changes and the inference framework can simply be adapted. One can think of inferring any required variable as long the LiDAR point clouds can be labeled with it. This section depicts an example on the inference of another variable, namely the POD of targets. Note that retrieving targets POD is also possible by simple correlation, as explained later in section 8.2.2.

The frustum point clouds used as input data can also be labelled using the POD information of targets. Indeed, as these noise points are acquired simultaneously to the detections of multiple targets in the chambers, they can all be combined in a dataset. Depending on the application, one can need to estimate the POD of specific targets at a specific distances and this can be achieved in this way. Our fog experiments allow to label *back* frustum points with the POD of *a1, b1, b2, b3, c1, c2, c3* targets. On the other side, the smoke tests enable the labeling of *free-sky* frustum points with targets *a1, a2, a3* and at multiple distances. To illustrate the concept, we choose to infer the POD of the a1 target at CEREMA using the frustum point clouds of the *back* target. The dataset now consists of frustum point clouds, labeled with the POD information of the a1 target. Figure 8.4 illustrates the range covered by the dataset, with each POD interval representing a distinct class. The frustum point clouds falling within the vertical limits depicted in red correspond to each respective POD class. It is worth noting that the selection of these POD intervals was not optimized for maximum performance but rather aimed at visualizing the challenges associated with class selection. To do so, the number of frustum points inside each classes should be increased while keeping the variability of their number at the minimum.

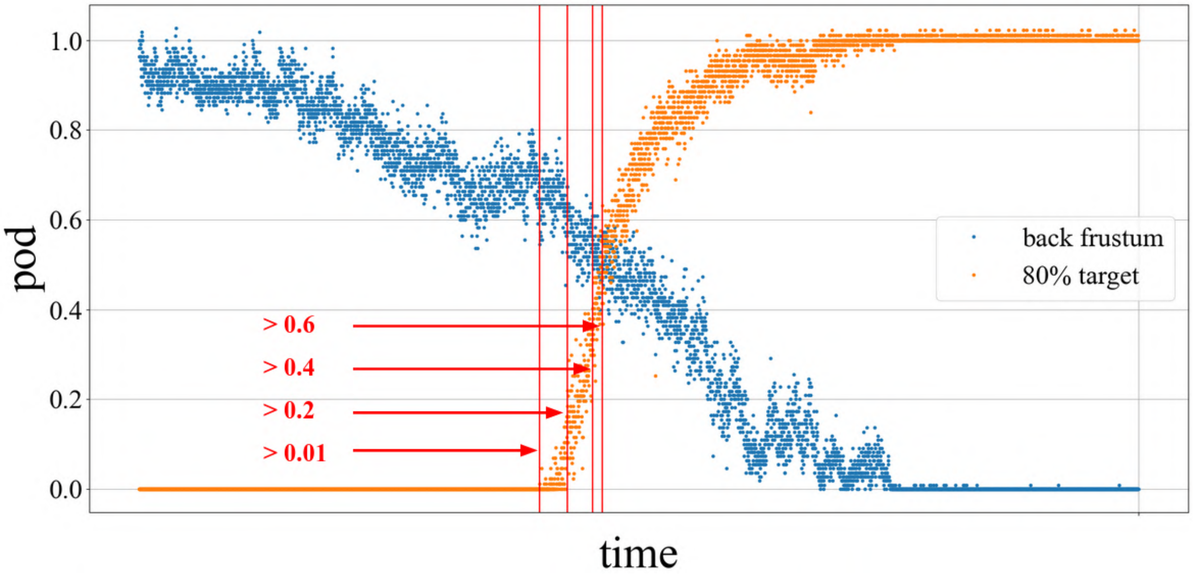

Figure 8.4: POD of both *back* target frustum and *a1* target during the big fog dissipation.

Subsequently, the model can be trained and tested using this dataset, following the same approach as before. Note that in this case, only the Gamma model (without RFS) is utilized, the intention here is to just showcase the concept. Fig.8.5 shows the classification results, with a quite low classification performance although the approach has not been optimized. Given these first results, we believe that the approach is relevant and suited to infer other variables. If other variables are of interest, one should create the appropriate dataset where the variables data are recorded simultaneously than the LiDAR scans used for inference. For example, it could also be interesting to infer the PSD of fog or smoke particles from our experiments.

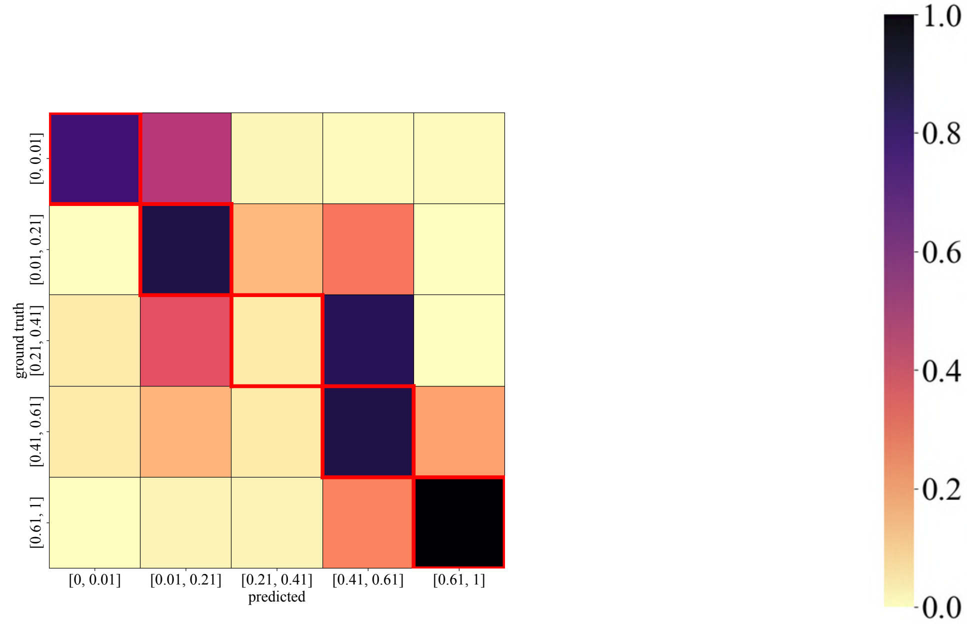

Figure 8.5: Confusion matrix of the POD inference of *a1* target at CEREMA using the frustum point clouds of the *back* target.

8.2.2 Correlation with POD

Without the incorporation of the POD in the inference model as stated previously, it is still possible to retrieve its value thanks to the visibility inference model, because the simultaneous records of targets POD and visibility allows for direct correlation between the two.

Using the influence of visibility over targets POD (shown in chapter 6 of part II) for the back target as an example), the inferred visibility classes can be used to indicate the resulting sensors POD of any target. Fig.8.6 shows two examples of the impacts of the inferred visibility classes on the *a1* target's POD where the targeted visibility classes are respectively $[65, 70]$ and $[45, 50]$. It illustrates that the 5 m resolution of visibility class can lead to high errors in the retrieval of POD for specific visibility classes. Especially for classes containing high variations in POD values, such as the $[40, 45]$ class for which the POD range from 0.2 to 0.6 and the class is inferred with probability p = 0.5 when targeting class $[45, 50]$.

The total classification results and different targets can be used to evaluate RMSE POD errors over the POD of targets at different distances and reflectivity values. It is calculated with the following equation :

$$RMSE_{POD} = \sum_{j=1}^{M} (POD_{gt,i} - min(POD_{inf,j}))^2 p_{inf,j} \tag{8.5}$$

where M is the number of inferred classes and i the targeted class. $POD_{gt,i}$ is the average POD from ground truth, $POD_{inf,j}$ is the POD data points from the inferred classes, and $p_{inf,j}$ is the probability of each inferred class j. Fig.8.7 showcases the mean POD with associated $RMSE_{POD}$ from all visibility classes inferred during the test phase and two different target configurations. The top and bottom plots respectively show the corresponding values for the three available high reflectivity targets $a1$, $b1$ and $c1$ and for the three Lambertian targets $b1$, $b2$ and $b3$ located at 17 m. It should first be noted that errors are of the same order of magnitude for all targets. As expected, the slope of the POD curve strongly influences the resulting $RMSE_{POD}$ of each target. Errors made on the POD are quite low when the POD slope is also low, and inversely higher $RMSE_{POD}$ values are seen when the POD increases. This is partly due to the 5 m step of our visibility classes: smaller classes would lead to higher precision of POD.

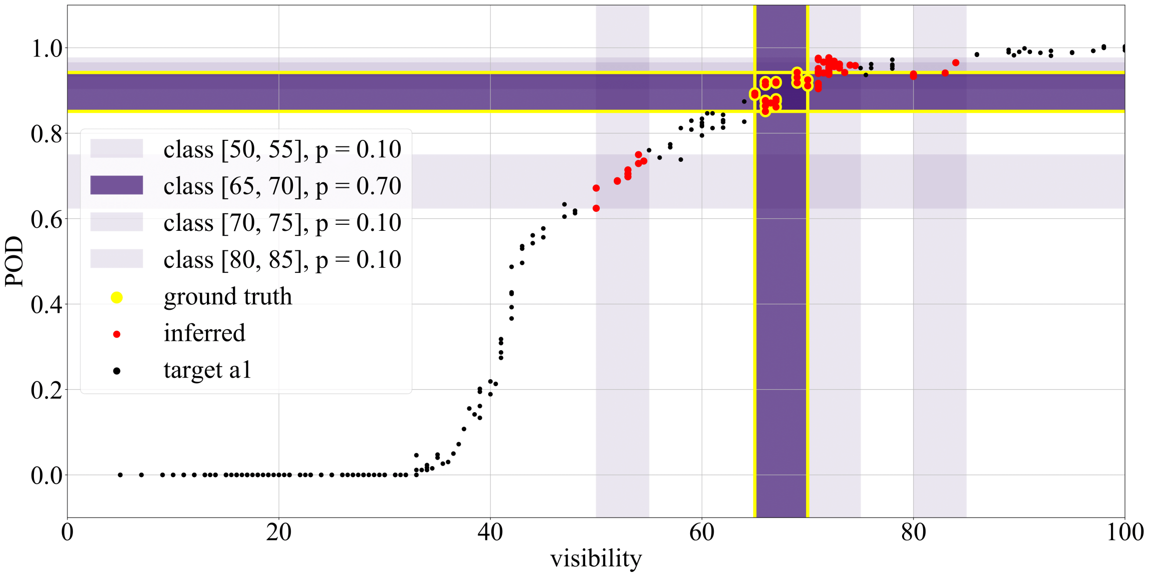

(a) Inferring $[65, 70]$ visibility class

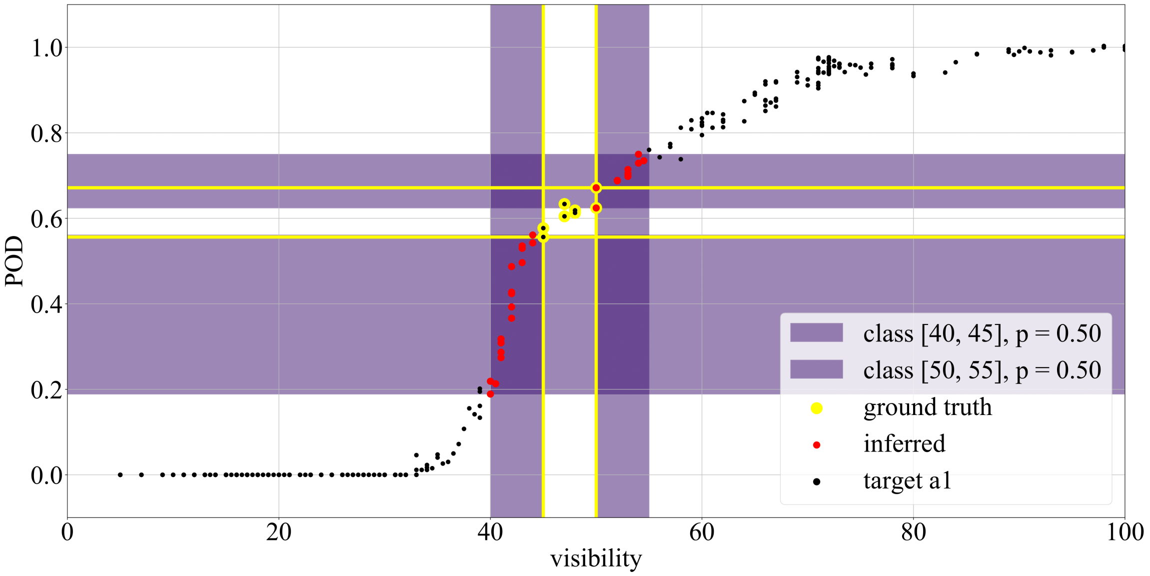

(b) Inferring $[45, 50]$ visibility class

Figure 8.6: Correlation between inferred visibility classes and POD of *a1* target.

This methodology is here used with fog conditions as we are limited by the available dataset. However it can be applied to assess the POD of reference targets in various DVEs. The inference model is thus an enabler to assess the compliance of AVs with safety regulations that require a certain level of detectability of reference targets. Future work consist in applying this approach to the presented experiments, especially the one with smoke where the targets are disposed at multiple positions.

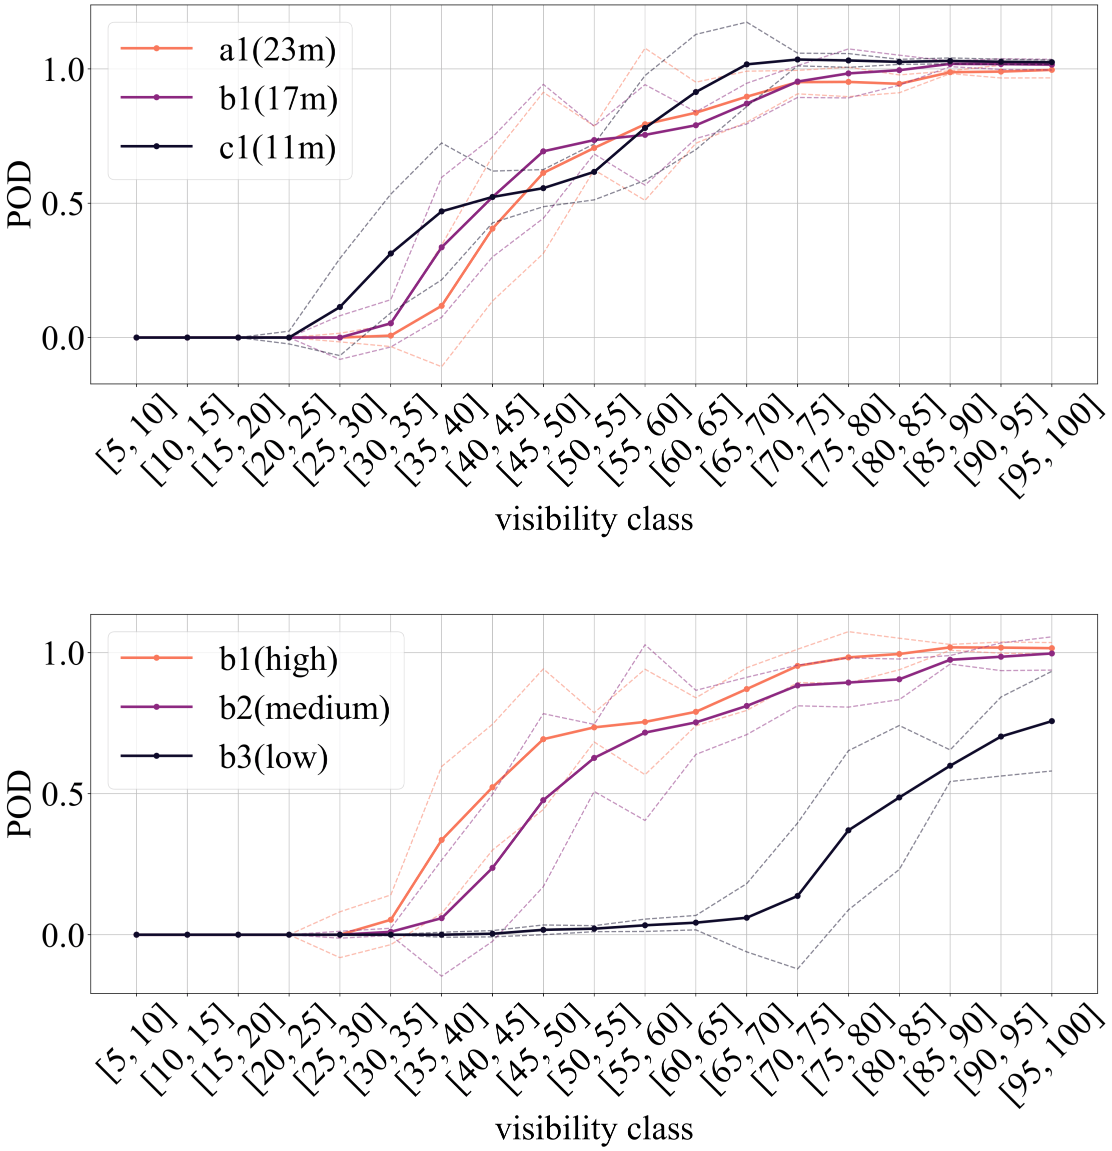

Figure 8.7: Evolution of POD and corresponding $RMSE_{POD}$ for each visibility class and different targets. Top: result for three high reflectivity targets at different ranges. Bottom: results for three targets at the same distance but with different reflectivities.

8.3 Longer term extensions

8.3.1 Multi-echo

Detecting multiple detections from a single laser pulse leads to increased density of point clouds and thus a better 3D representation of the scanned environment. This is particularly true with vegetation and DVE conditions. Though, one should be aware of the increased noise induced by the multiplication of echoes because lowering a detection threshold in the full wave processing can produce undesired detections. LiDARs with accurate multi-echo capacity are mainly used for topographic purposes rather than in the automotive industry. Still, using the multi-echo labeling to enhance the understanding of the phenomena and potentially the results of visibility estimation has always been in mind. This last section proposes some results we found during our experiments.

VLP32

Here, we analyze the labeling of echoes inside the point clouds of the target and the frustum for the VLP-32 LiDAR sensor. This instrument can produce two echoes from a single emitted laser beam. In addition, the labeling of echoes is done according to the intensity of the returned signal. Thus, small intensity echoes (likely to come from small objects or water particles) are labelled as second echoes (echo 2, second *Strongest*). Inversely, solid and bigger obstacles are prone to returning more energy and have points labelled as first echoes (echo 1, *Strongest)*. Fig.8.8 shows the same POD graphs of chapter 6 from rain and fog experiments but with the labeling of echoes 1 and 2.

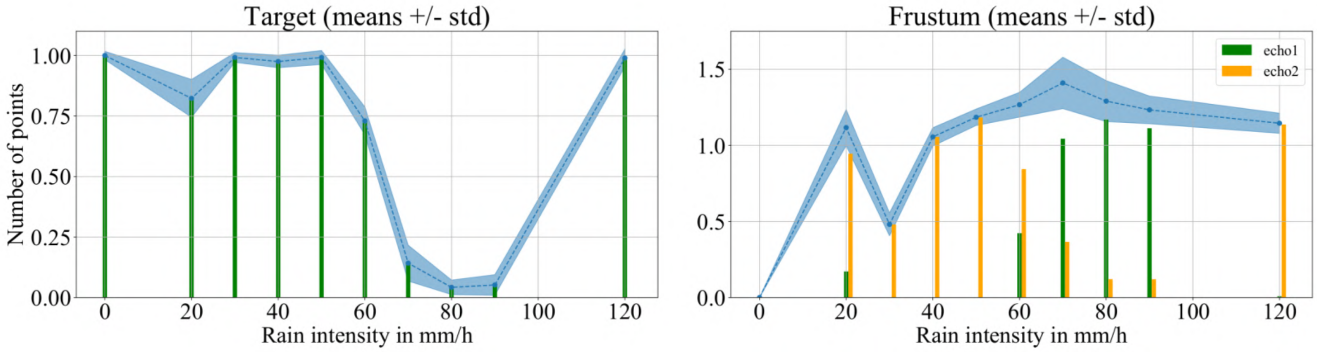

(a) rainy conditions

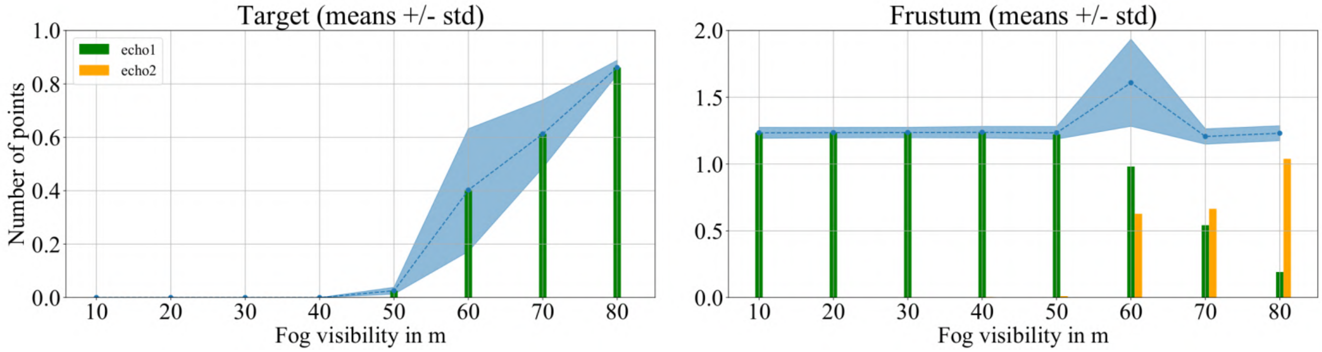

(b) foggy conditions

Figure 8.8: Results from VLP-32 with multi-echo labels.

Detections from the target show that, regardless of the nature of the artificial DVE (rain or fog), the solid target always returns points labelled as first echoes. Even when points are detected both in the frustum and on the target, points from the latter always remain *Strongest* echoes. This is interesting because one could think that when the target is hardly detectable, the sensor could be confused about its *Strongest* echoes. Especially in fog conditions, when we observe that the intensity of the target and frustum points can be close to each other, as shown on Figure 6.14. Then, a mixture of echoes is observable for the frustum points. In rainy conditions, proportions fluctuate depending on the number of points on target, which itself varies with the rain rate. Indeed, when the target is well detected (more than 75% of the nominal number of points), first echoes come from the target and the number of second echoes is the number of points inside the frustum. As the detection performance from the target decreases, first echoes tend to appear in the frustum and we see a combination of first and second echoes. When target detection fails, most echoes in the frustum are labelled as first echoes. Here, the rain droplets become the *Strongest* echoes as the solid target in the line of sight does not return enough energy anymore. Foggy conditions produce similar results. Until the target is detected, all first echoes are located in the frustum while the sensor does not produce any second echo. Then, starting at 60 m of visibility, the number of second echoes rises as visibility increases and the number of first echoes decreases in the frustum. The sensor receives enough energy from the target to detect it and allows for second echoes in-between.

As said in chapter 6, the multi-echo capability of the VLP-32 does not allow it to have better target detection performances with regards to the other spinning LiDAR sensor. Though, the labeling of echoes provides valuable knowledge on the characteristics of the environment when looking at points located in a frustum. A next step is to acquire data from natural conditions, including context weather information and LiDAR point clouds. A deeper analysis on the behavior of each laser shot (which can result on target detection, frustum detection or no detection at all) is expected to add precision.

OS1-128

The CEREMA experiments showed that a mix of single and second echoes are found in the frustum noise detections with rain and fog conditions (using the VLP-32 LiDAR) and also that a majority of strongest echoes are found on the solid targets. These results are quite intuitive and could potentially lead to more analyses and developments, but as the thesis direction turned to the Ouster sensor, it was not investigated. However, the OS1-128 Ouster LiDAR operated during the smoke acquisitions made subsequently was a more recent version of the sensor with an included multi-echo capacity. This section depicts the findings on the analysis of the multi-echo data with the Ouster sensor and how it could possibly be applied to our inference approach. The objective is only to give insights on the data for interested readers but not to delve into its use and quantify any results.

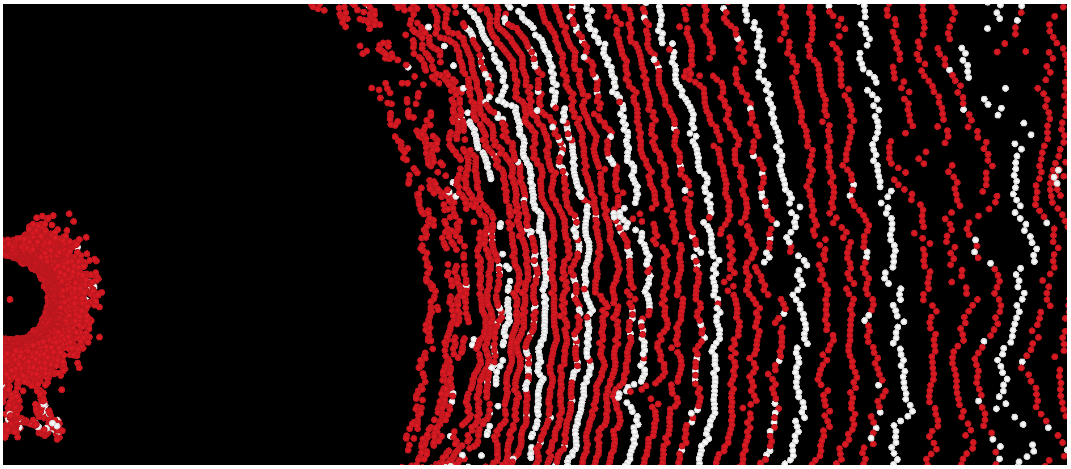

Figure 8.9: Bird-eye view of a point cloud captured by the Ouster sensor in smoke conditions with multi-echo labeling, red for echo 1 and withe for echo 2.

Figures 8.9 and 8.10 show two screenshots of point clouds measured at the beginning of a smoke dissipation. The red points correspond to echoes 1 and white points echoes 2. The first thing to point out is another ring-dependent effect, similar than the one of previous section. The first screen shot shows a bird-eye view of the point cloud where points forming a circle on the left are close-range frustum detections caused by smoke and points on the right correspond to ground detections. It is observed that the majority of echoes labeled as echo 2 lie on specific rings, which are separated by 3 rings every time. Although in accordance with the ring family effect, this result is not intuitive and most likely physically wrong as it cannot be caused by the smoke conditions and ground nature. The second screen shot is taken from the sensor's point of view where the observable points are close-range frustum smoke detections. For information, the low density regions in the middle of the image correspond to parts of the FoV where solid targets are detected at further range, restraining the possible noise detections. Plus, a large amount of echoes 2 are detected in the lower part of the FoV and across a large horizontal region. This result is hardly explainable because it does not seem to be caused by any physical object. One potential interpretation is that the sensor is designed differently in this region than the rest to increase the number of ground detections being highly valuable to AD for terrain estimation techniques. Either the detection threshold is globally lowered for these rings or higher energy lasers are shot, the resulting number of ground detections is increased. If these detections are labeled as echo 1 because the ground returns more energy than the smoke, it could explain this stripe of echo 2 points. This interpretation is quite hypothetical but showcases the fact that, just as with fig.8.9, the multi-echo is influenced by the sensor's characteristics. This also shows how sometimes the black box aspect of these sensors are limitations because these counter-intuitive effects can restrain the analysis of their data.

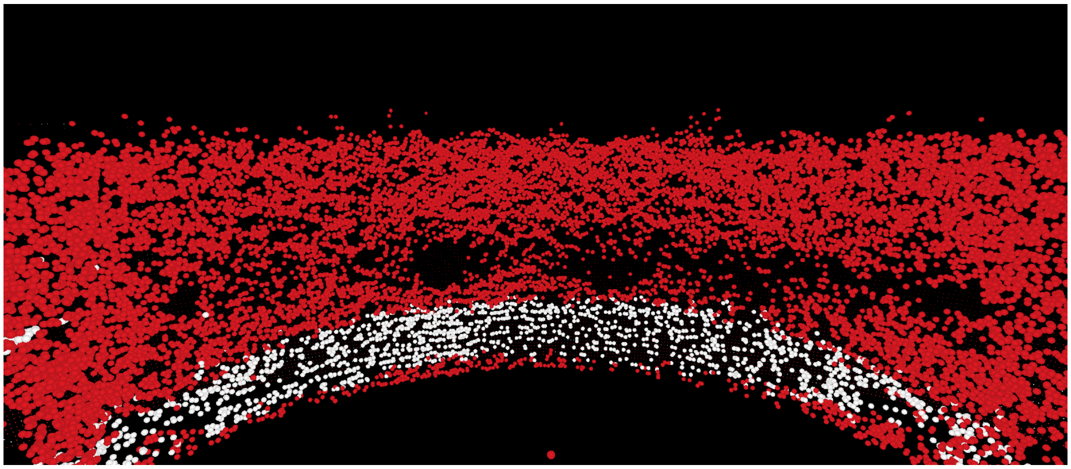

Figure 8.10: Sensor view of a point cloud captured by the Ouster sensor in smoke conditions with multi-echo labeling, red for echo 1 and with for echo 2.

Regardless of these issues, the multi echo capacity still seems applicable to our inference model. Fig.8.11 shows the distance distributions of echo 2 frustum detections shooting at *a1*, *a2* and *a3* targets for each smoke dissipation performed. The d1 to d6 labeling corresponds to the target distance (available on table 6.5). Similar distributions than the ones used in the inference model are observed, at close range and with potentially Gamma or Log-normal shapes. This encourages the use of these echoes for the inference because it increases the amount of data. Another thing observable from these distributions is the evolving number of points. First, closer targets (from d1 to d6) cause an increase in the number of frustum detections. This is caused by the larger frustum spaces available during point cloud extractions. Then, high reflectivity targets show more echo 2 frustum points than low reflectivity ones (a1 > a2 > a3). High reflectivity means more chance of producing echo 1 detections on target and thus, more chance of producing echo 2 detections in the frustum. Note that the echo 2 frustum detections from the *free-sky* are not shown because they are non-existent. Indeed, as their is no solid target in its FoV, the frustum detections are all labeled as echo 1.

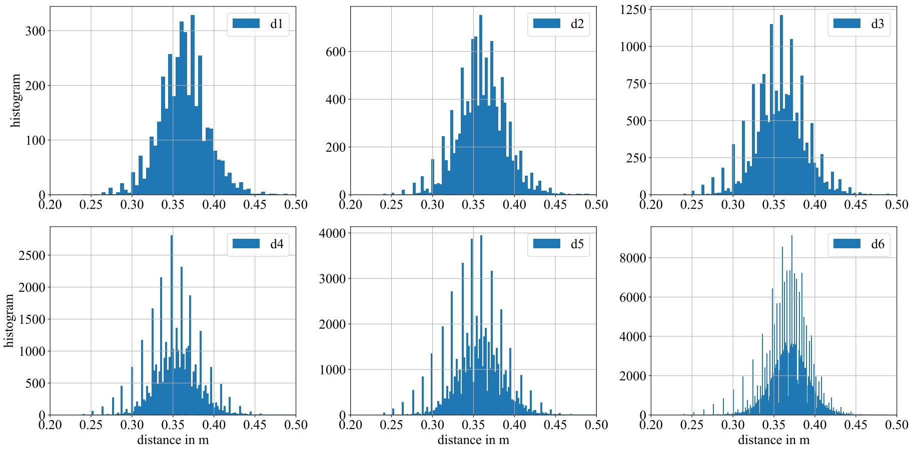

(a) *a1* target

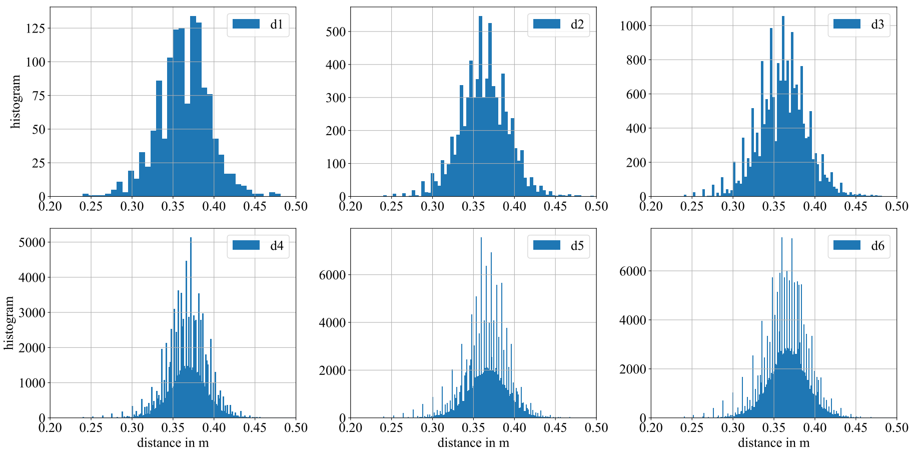

(b) *a2* target

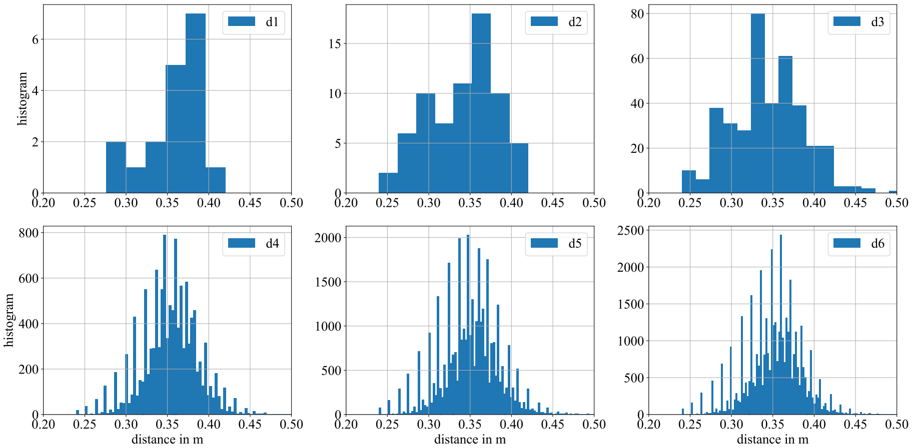

(c) *a3* target

Figure 8.11: Distance distributions of echo 2 frustum points shooting at *a1*, *a2* and *a3* targets for each smoke dissipation.

8.3.2 Application to outdoor conditions

Artificial conditions brought valuable data in this thesis. They helped to test and compare LiDARs in a large variety of conditions with different properties. And, it enabled the development of our inference model with the fog and smoke dissipation experiments. Though, the generated conditions are very likely to be different from the ones of real scenarios, our models are potentially not directly applicable in outdoor conditions. Future work should focus on applying our model and methodology on data taken in natural environments. Unfortunately, we did not have the time to do so. An automatic recording system has been in development during this thesis. It is presented in appendix B but no results are given here.

Recent work published data captured by an outdoor automatic recording system with a few LiDAR systems including a Ouster OS1-32 [137]. They propose data from different weather conditions of fog, rain, snow and low-sun and the environmental conditions are assessed with context sensors (including visibility). In this section, we present the results found which could be of interest for our approach, focusing on the Ouster LiDAR. Fig.8.12 shows different cumulative distance distributions of *free-sky* noise points (which they call ”atmospheric detections”) in fog conditions and various visibility values. This depicts that, similar to our findings, the majority of these echoes are located around 0.3 m from the sensor and that after 0.4 m, no more points are detected.

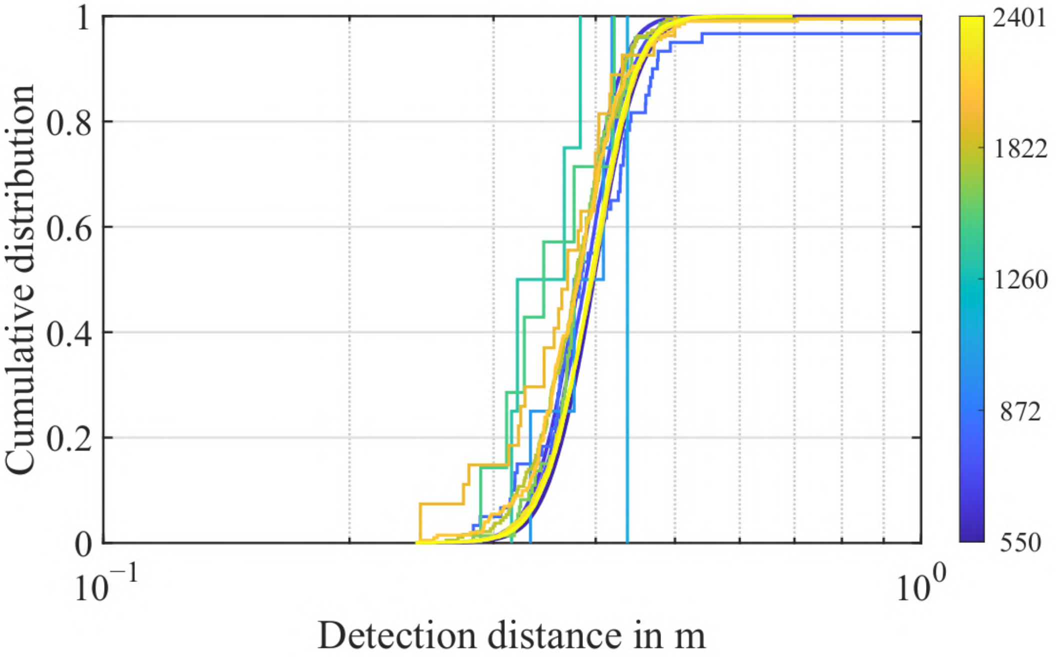

Figure 8.12: Distance cumulative distribution of atmospheric detections in fog. The visibility in m is denoted as color, each line represents a logarithmic visibility bin, excerpt from [137].

This graph encouraged us to download the dataset and dig into the shapes of these distributions and one example is given in the next figure. Fig.8.13 shows the distance distribution of the *free-sky* frustum points accumulated over 5 LiDAR scans. The shape of this distribution seems to be in accordance with our presented results. This is very encouraging for the deployment of our solution in outdoor conditions, as we believe our model would easily fit with this kind of distribution.

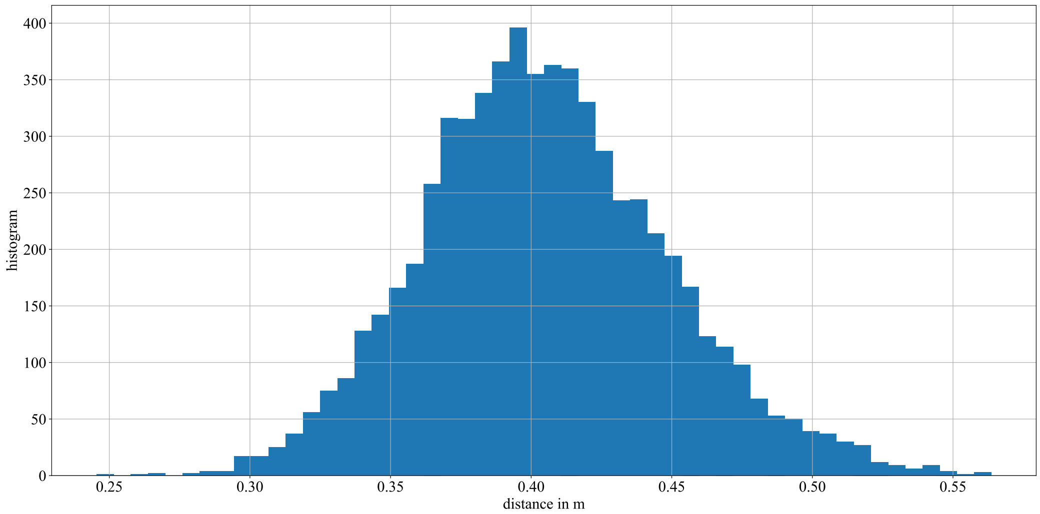

Figure 8.13: Histogram of the distance of atmospheric detections in fog, accumulated on 5 Ouster LiDAR scans.

Chapter 9

Conclusion

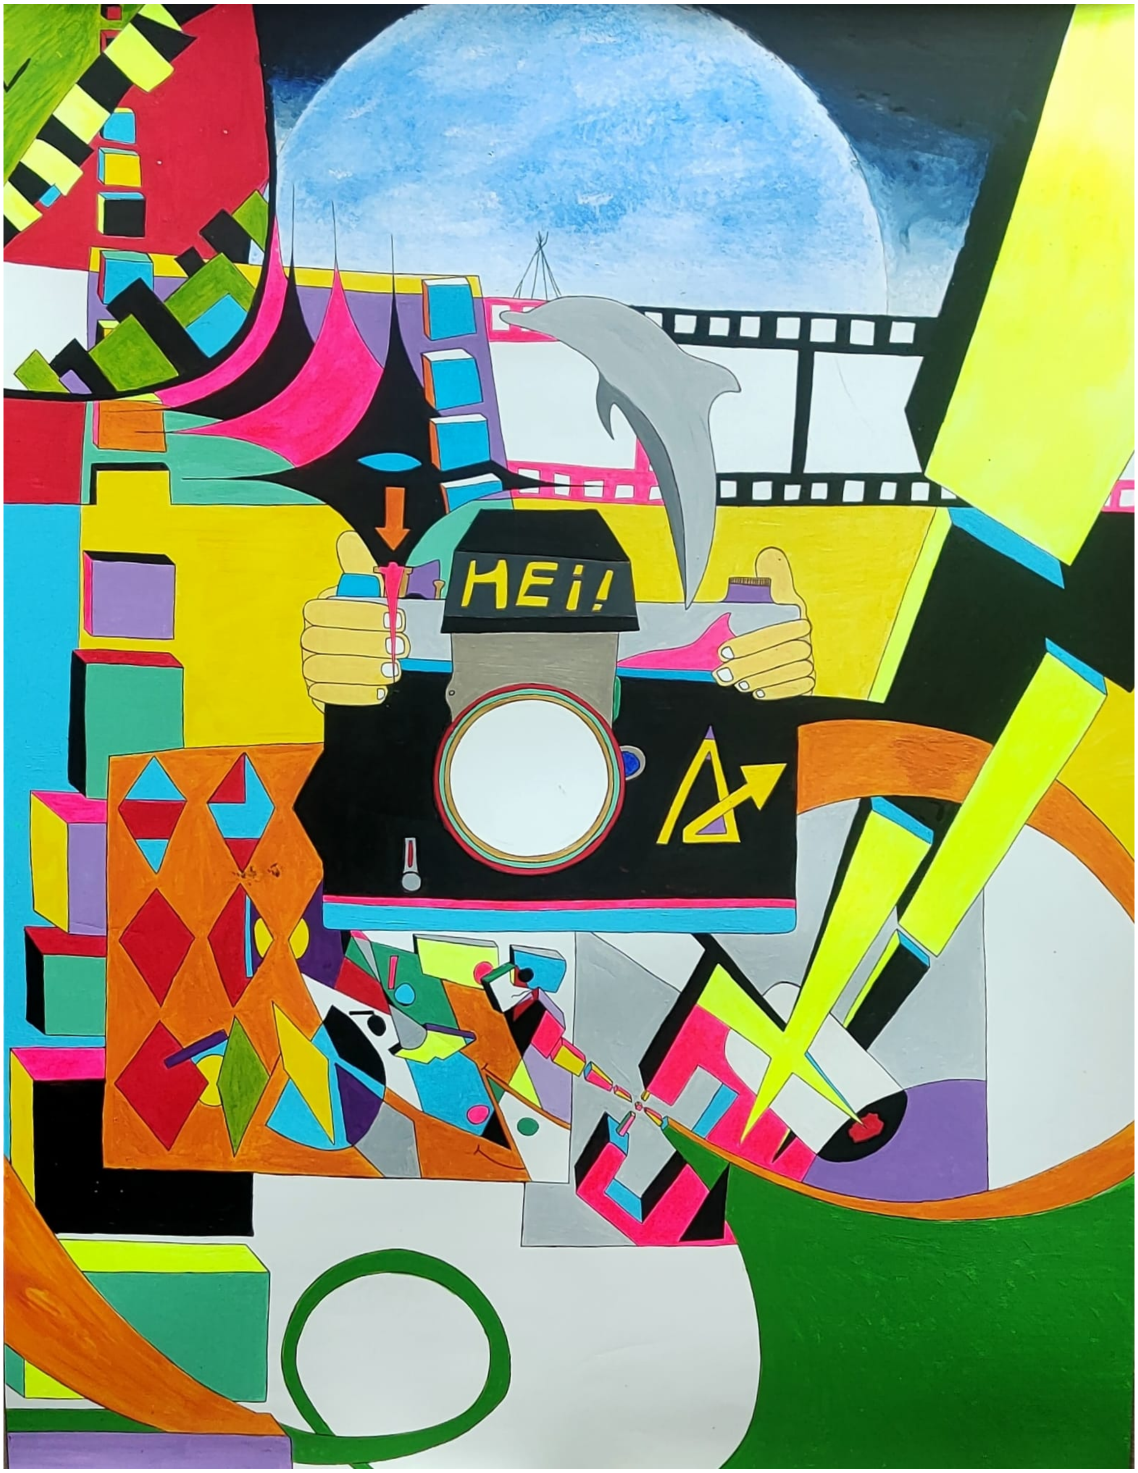

Summary

This thesis proposes advancements in LiDAR perception facing degraded visual environment for autonomous driving. Our work contributes in increasing the awareness of autonomous vehicles in degraded conditions, in particular fog and smoke.

We provided an overview of the 3D-LiDAR technology in **chapter 2** where a wide variety of technologies is presented, which shows progresses, but also uncertainty as to the most suitable design for autonomous driving.

Chapter 3 outlines the various DVE scenarios. Although the physics of the interactions between laser signals and particles is highly complex, models and metrics allow to simply describe these phenomena.

Chapter 4 provides insight on the perception algorithms using LiDAR point clouds for autonomous driving, that are all impacted when the point cloud data is degraded. A focus at denoising and weather classification methods is given before we describe an overivew of our solution: **a Bayesian inference model which aims to retreive the value of visibility from the identification of distance distribution of aerosol echoes**. Our methodology is designed to enable generalization because the point clouds are extracted within specific parts of the LiDAR field-of-view, so that the results can be applied in natural conditions.

Chapter 6 presents a quantitative analysis on the degradation of different LiDAR technologies under artificial rain, fog and smoke conditions. We assess the Probabilities Of Detection (POD) for both targets and associated noise simultaneously, as well as the distance distribution of the noise detections. It enables the evaluation of the degradation of each sensor facing the DVE conditions, the comparison of sensors between each other, and also to get closer to safety standards. Naturally, changes in the severity of environmental conditions are expected to have a related influence on the laser particle interactions and therefore on the properties of the degradation. At least in theory, because the internal processes of commercially available LiDARs are poorly known and can produce counter-intuitive results, indicating unique behaviors for each sensor. Contrary to the other sensors, the Ouster OS1-128 spinning LiDAR showcases identifiable distributions under fog and smoke conditions. The analysis of these distributions leads to the development of the visibility inference model.

In **chapter 7**, we develop a Bayesian inference model which identifies these distributions to classify the visibility conditions. Different models are selected to fit the data, for the shape of the distributions (Gamma and Log-normal), and to exploit the number of elements of the distribution (Poisson and Binomial RFS models). The problem is posed as a classification one, allows to combine different distribution models, and provides satsifying results. The choice of Gamma or Log-normal models does not exhibit any clear difference, whereas the inclusion of RFS models significantly enhances the results, the Poisson RFS model yielding the best performances. These RFS models introduce new possibilities to apply the formalism on point clouds, which, to the best of our knowledge, has not been explored previously.

Chapter 8 constitutes the perspectives of this thesis, which concern direct and longer-term extensions of our approach via the inference model.

Incremental extensions of the inference model

The inference model is extended with a Bayesian filtering technique. Using an appropriate dynamic model of the visibility evolution, the estimation performances are improved by mitigating the influence of inferred classes that deviate significantly from the ground truth. The application of the such method is also likely to help track fluctuating visibility. More data acquisitions with fluctuating visibilities are yet needed to assess the approach.

Our model can also be used to differentiate the type of DVE, as we demonstrate the differentiation between the experimented two types of fog. Using more complete datasets with different types of fog, smoke, and maybe dust, we believe that our approach could be suited for any aerosol conditions. This potentially enables the choice between data processing and denoising algorithms, each of them being designed for a specific DVE type.

Furthermore, we showed that our approach can be adapted to estimate targets' POD, either directly or through correlations. The target POD provides insights into the LiDAR's detection capacity and exhibits high correlation with the performance of denoising or tracking techniques, making it a more interesting variable to estimate.

The evaluation of the DVE severity, through visibility (or targets POD), can also be used to parametrize algorithms to adapt to the changing conditions. The evaluation of how the inferred parameters of our approach can be used to better parametrize the various point cloud processing algorithms (denoising, object detection, localization, ...) is also relevant future work.

Far-fetched extensions

The multi-echo capacity of LiDARs has the potential to be used in our approach because it can increase the number of noise detections in DVE and thus the amount of data available for the inference model. For this, care should be taken into the sensors characterization before any developments, as their outcomes can occasionally be counter-intuitive due to undisclosed internal factors.

The use of artificial conditions is great for the ability to repeat and control the acquisitions because recording data from natural events is rather difficult. But care should be taken into the generalization of features trained on data obtained in such conditions, as the generation mechanism of artificial DVE conditions often hardly approach the properties of natural events. Unfortunately, our approach has not been tested in outdoor conditions, due to the lack of data from natural events. However, we present some results found in a recent publication where similar distributions are found using a similar Ouster LiDAR in natural fog conditions, which confirms that our approach may be suited for outdoor conditions.

Finally, one could think to extend our solution to precipitation conditions such as rain or snow. However, those conditions tend not to produce the same point clouds that we can observe in fog conditions (with points in very close range), but these characteristics also depend on the sensor itself. We have experienced that normal precipitation events (a few mm/h) do not cause high amount of noise points and these points are located homogeneously in the surrounding of the sensor (up to tenth of meters). Whether this method is applicable in other DVE types may depend on the use case and the surrounding environment of the sensor. Furthermore, the approach using RFS may be a good approach for precipitation (the higher the rain rate, higher number of points in the frustums are expected). Future work could include the acquisition of data and potential specific developments.

When addressing the DVE problem as a whole, the selection of appropriate metrics becomes crucial. While commonly used metrics such as precipitation (measured with rain rate in mm/h) and aerosols (measured with visibility in m) have their merits, they may not be the most suitable for correlating with the impacts experienced by LiDAR sensors. Given the complex nature of the involved interactions, which encompass multiple scattering effects, it is likely that detailed information regarding particle properties will provide more valuable insights rather than relying solely on coarse metrics.

Finally, physical limitations will always lie in the interactions between laser signals and disturbing particles. It is very likely that some scenarios will never be resolved using only LiDARs: it is necessary to introduce data gathered by other complementary sensors such as visible and infrared cameras, radar, or even ultrasound sensors.

Appendices

Appendix A

A ring family ?

During the analysis of the point cloud data, an interesting discovery has been made regarding the Ouster OS1-128 LiDAR and a potential effect related to ring indexing. The ring is an index which corresponds to the position of the laser emitter (and associated detector) on the vertical axis. In particular, the observed effect concerns the distance distribution of the DVE frustum points captured during the fog and smoke experiments. Note that this effect has not been witnessed on solid targets. This finding is presented in the appendix part because it is very sensor-specific and it is possible that future versions of this sensor will not behave similarly.

But, this effect can have a direct impact on the classification performance of the presented inference model. In the presence of fog or smoke, it seems that the distance distributions caused by the DVE particles and found at close range exhibit distinct behaviors. These behaviors are closely related to the association with a particular ring indexing, or as we call it, ring family. Empirical evidence suggests that the ring family associated with the specific behaviors in distance distributions includes every ring with an index step of 4. Fig.A.1 shows 2 screenshots where the first 100 LiDAR point clouds from the big fog dissipation are accumulated and displayed.

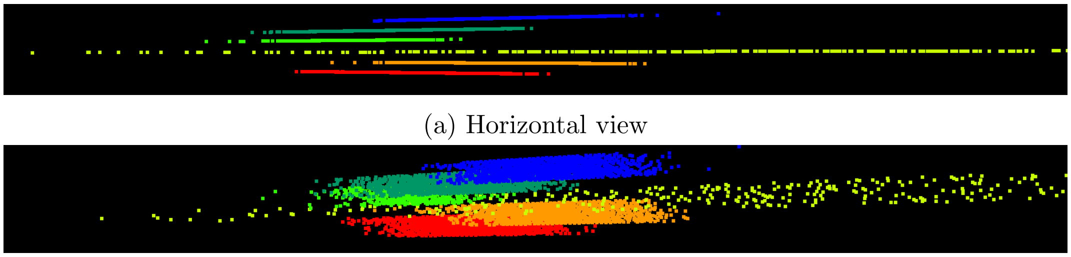

(a) Horizontal view

(b) 3D view

Figure A.1: Screenshots of accumulated frustum (using the *back* target) points in 3D collected in fog conditions. Color represents the vertical ring index.

They consist of frustum points acquired while shooting the *back* target. What can be seen is that depending on the vertical ring indexing of the points, the resulting distance distribution differs. The blue points appear to exhibit the same behavior as the orange ones, as do the dark green and red ones. The light green points could probably be associated with the ones under the red points but this frustum extraction did not include them. Potentially, these 3 associations generate distributions similar to the aforementioned Gamma and Log-normal distributions from the inference model. But having different behaviors, this can potentially influence the classification results as it changes the parameters of the distributions. On the contrary, the yellow points do not show any distributions resembling Gamma or Log-normal characteristics. In this scenario, these data points should be regarded as noise rather than reliable information and should be eliminated to enhance the consistency of the distributions and thus the accuracy of the inference results. Fig.A.2 shows plots of these distributions with the addition of the intensity of the points, for the 4 types of rings. The effect visually observed on the previous screenshots is also seen here. Ring families 1, 2 and 3 show relatively similar distributions but with small changes of position. However, the distributions of the 3rd ring family have less data points. Finally, ring family

4 produces lower number of points and with a different distribution shape. In addition, the intensity of these points seems to be impacted similarly, with common behavior for ring families 1,2 and 3, and a different one for the 4th.

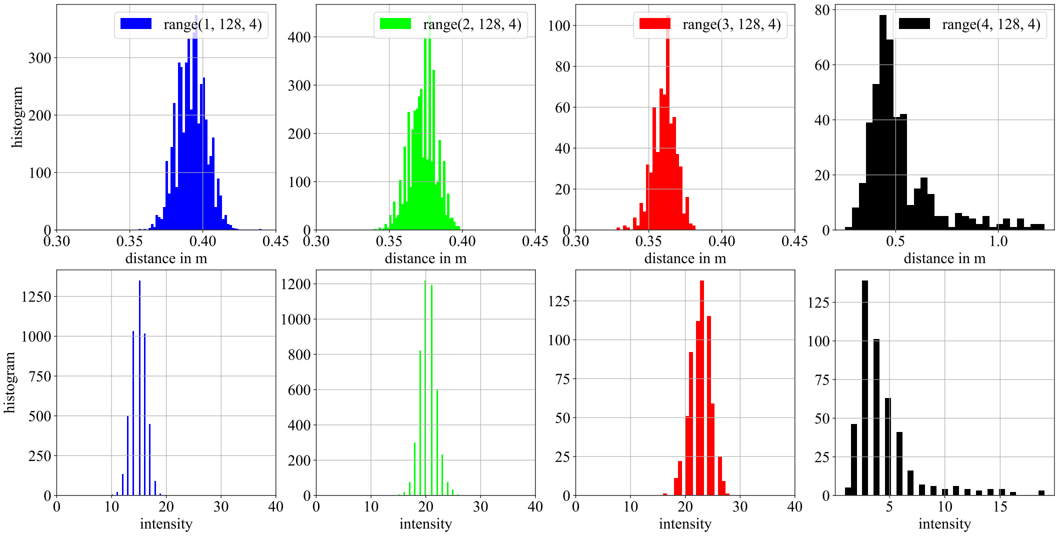

Figure A.2: Distance and intensity distributions of accumulated frustum points in fog conditions, for each of the 4 ring family.

This phenomenon appears to be a limitation of the sensor. However, as mentioned earlier, it is not observed on solid targets at greater distances. Therefore, we can safely disregard this effect for the regular use of the sensor. Ouster was asked regarding this issue but no detailed response has been received. Consequently, we can only speculate about the cause of this phenomenon, which may be related to hardware factors.

It is noteworthy that the colors in this illustration correspond to the distributions shown in fig.A.2. Regardless of its origin, we felt somewhat obliged to check the impact on our model, as it can now be trained and tested on data taken from each ring family individually. Although this phenomenon has been observed within all fog and smoke experiments (showing by the way that the effect is present on different versions of the sensor), to illustrate the results, only the big fog dissipation experiment is used, with the Gamma model (No RFS). The $RMSE_V$ results are shown on table A.1. The ring family 1 holds the best results and even outperforms the ones obtained without any ring selection. On the contrary, ring families 2 and 3 show similar results but they are poorer than the original ones. Finally, the ring family 4 has the worst results.

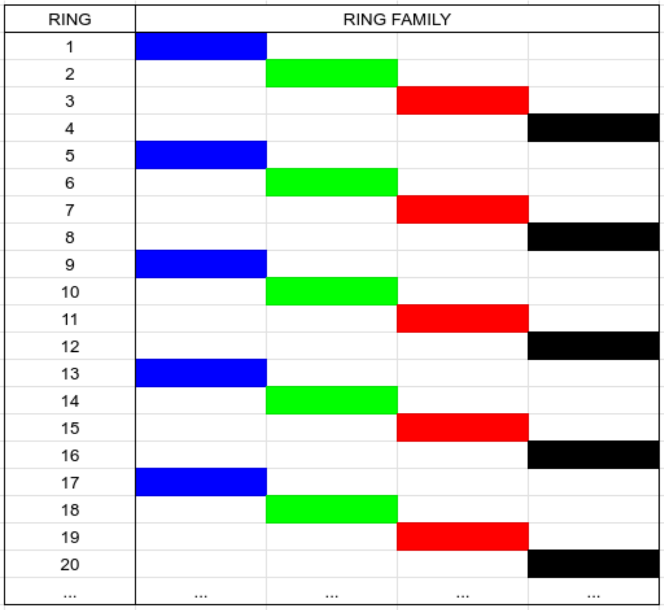

RING	RING FAMILY			
1	🟦			
2		🟩		
3			🟥	
4				⬛
5	🟦			
6		🟩		
7			🟥	
8				⬛
9	🟦			
10		🟩		
11			🟥	
12				⬛
13	🟦			
14		🟩		
15			🟥	
16				⬛
17	🟦			
18		🟩		
19			🟥	
20				⬛
...	...	...	...	...

Figure A.3: Illustration of the ring families

Ring family	1	2	3	4	All rings
$RMSE_V$	5.2 m	7.2 m	7.3 m	13.1 m	6.4 m

Table A.1: Classification errors depending on the ring family for the big fog experiment and Gamma model.

To conclude, although being a very specific flaw (dense aerosol condition, very close frustum detections), the investigation of this sensor specificity improves the inference results. Even though selecting data from only one ring family greatly decreases the data size for both training and testing.

Appendix B

Additional work

- Rain Tunnel

Following the CEREMA test campaign, EasyMile decided to acquire its own system for generating artificial rain, shown fig.B.1. The idea is to produce controlled rain conditions to quantify the vehicles behavior against rain of various properties. As a result, integration issues of the sensors as well as perception degradation can be assessed and the conditions can be repeated to experiment solutions. Since we observed particular behavior during the CEREMA test campaign (chapter 6), the properties of the produced rain have to be qualified. To do so, the Parsivel disdrometer is used, already used in the CEREMA campaign. It is positioned to give measurements at 7 different locations under the tunnel.

Figure B.1: EasyMile TractEasy under the rain tunnel.

The next figure shows a complete plot of the results drawn from this test protocol, where the rain produced is supposed to approach a 15 mm/h rate. The plots titled with number from 1 to 7 correspond to diameter-velocity histograms measured at the different locations. It depicts that regardless of the sensor position, the distributions are relatively similar and seem to follow the empirical model. This shows that the rain has a rather natural form in terms of diameter and speed. But the intensity measured by the disdrometer (bar plot at the center top) shows relatively consistent results which do not align with the supposed 15 mm/h. Again, we can conclude that the intensity metric might be misleading in the evaluation of the rain conditions.

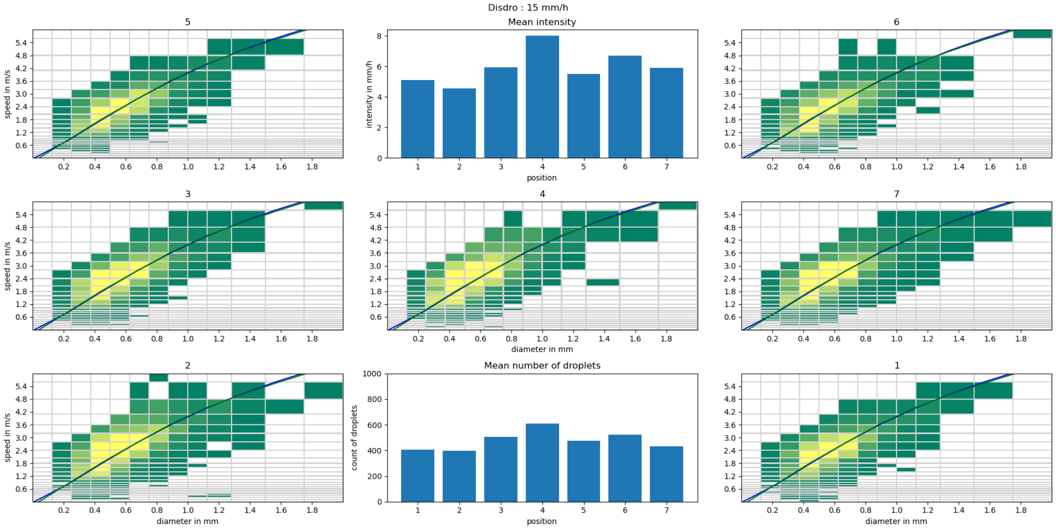

Figure B.2: Rain characteristics measured by the disdrometer according to the position in the tunnel and during 1 minute acquisitions. The plotted line comes from the empirical model of [190]

- Automatic recording system

The need for data of natural conditions and the difficulty of acquisition led to the development of an internship subject, which I had the chance to co-supervise. The idea is to build a system capable of automatically record data when DVE conditions are present, similar to [137, 67]. The system includes LiDARs, weather context sensors (B.3) and targets located at some distance. Using the data continuously measured by the context sensors, the system automatically triggers the powering of the LiDARs and records data. This data is then stored in a specific format to construct a dataset of weather information and LiDAR point clouds. As such, the methodology applied during this thesis and probably other studies can be performed on these records of natural outdoor conditions. Unfortunately, this project has not reached production yet and I am unavailable to present any results. Nonetheless, the system has been set up and the software is under development.

Figure B.3: Automatic recording system.

Bibliography

[1] Sebastian Thrun et al. “Stanley: The robot that won the DARPA Grand Challenge”. In: *Journal of Field Robotics* (2006).

[2] *Full Self-Driving*. URL: `https://www.youtube.com/watch?v=tlThdr3O5Qo` (visited on 06/23/2023).

[3] *My First Ride In An Autonomous Taxi! The Full Waymo Ride Experience From Start To Finish*. 2023. URL: `https://www.youtube.com/watch?v=cSkWNZp9kqA` (visited on 06/23/2023).

[4] *Tesla Autopilot Fails and Disengagements - Compilation - 2019-40-50*. URL: `https://www.youtube.com/watch?v=yi5sVTewmXc` (visited on 06/23/2023).

[5] JJRicks Studios. *Waymo Self Driving Taxi Goes Rogue: Blocks Traffic, Evades Capture — JJRicks Rides With Waymo #54*. 2021. URL: `https://www.youtube.com/watch?v=zdKCQKBvH-A` (visited on 06/23/2023).

[6] On-Road Automated Driving (ORAD) committee. *Taxonomy and Definitions for Terms Related to On-Road Motor Vehicle Automated Driving Systems*. Tech. rep. URL: `https://www.sae.org/content/j3016_201401`.

[7] *Electric Cars, Solar and Clean Energy*. URL: `https://www.tesla.com/` (visited on 06/23/2023).

[8] Mile Easy. *Safety Report*. Tech. rep. URL: `https://easymile.com/sites/default/files/easymile_safety_report.pdf` (visited on 06/22/2023).

[9] Mark Campbell et al. “Autonomous driving in urban environments: approaches, lessons and challenges”. In: *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* (2010).

[10] Nijat Rajabli et al. “Software Verification and Validation of Safe Autonomous Cars: A Systematic Literature Review”. In: *IEEE Access* (2021).

[11] *EasyMile — Autonomous vehicle technology and solutions*. URL: `https://easymile.com` (visited on 06/26/2023).

[12] Ruth Bergel-Hayat et al. “Explaining the road accident risk: Weather effects”. In: *Accident Analysis and Prevention* (2013).

[13] *Waypoint - The official Waymo blog: A fog blog: Understanding a challenge inherent to driving in San Francisco*. URL: `https://waymo.com/blog/2021/11/a-fog-blog.html` (visited on 07/12/2023).

[14] *Korean Competition Shows Weather Still a Challenge for Autonomous Cars - IEEE Spectrum*. URL: `https://spectrum.ieee.org/japan-competition-shows-weather-still-a-challenge-for-autonomous-cars` (visited on 05/10/2023).

[15] Filip Karlo Dosilovic, Mario Brcic, and Nikica Hlupic. “Explainable artificial intelligence: A survey”. In: *International Convention on Information and Communication Technology, Electronics and Microelectronics (MIPRO)*. 2018.

[16] Karl Montalban et al. "A quantitative analysis of point clouds from automotive lidars exposed to artificial rain and fog". In: *Atmosphere* 12.6 (2021), p. 738.

[17] Karl Montalban et al. "Bayesian inference of fog visibility from LiDAR point clouds and correlation with probabilities of detection". In: *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE. 2023, pp. 7076–7082.

[18] Karl Montalban et al. "Bayesian inference of visibility in fog and smoke artificial conditions from 3D-LiDAR point clouds". In: *Journal of Quantitative Spectroscopy and Radiative Transfer* 308 (2023), p. 108666.

[19] *International conference on Target and Background Modeling and Simulation*. URL: `https://itbms.onera.fr/` (visited on 09/26/2023).

[20] *International Conference on Robotics and Automation*. URL: `https://www.icra2023.org/` (visited on 09/26/2023).

[21] Louis D. Smullin and Giorgio Fiocco. "Optical Echoes from the Moon". In: *Nature* (1962).

[22] M. Harris, M. Hand, and A. Wright. *Lidar for Turbine Control: March 1, 2005 - November 30, 2005*. Tech. rep. 2006. URL: `http://www.osti.gov/servlets/purl/881478-40C9A0/`.

[23] John C. Brock and Samuel J. Purkis. "The Emerging Role of Lidar Remote Sensing in Coastal Research and Resource Management". In: *Journal of Coastal Research* (2009).

[24] James Schindling and Cerian Gibbes. "LiDAR as a tool for archaeological research: a case study". In: *Archaeological and Anthropological Sciences* (2014).

[25] Behnam Behroozpour et al. "Lidar System Architectures and Circuits". In: *IEEE Communications Magazine* 10 (2017).

[26] Paul F. McManamon. *LiDAR technologies and systems*. SPIE PM. Washington, 2019.

[27] Martin Beland et al. "On promoting the use of lidar systems in forest ecosystem research". In: *Forest Ecology and Management* (2019).

[28] You Li and Javier Ibanez-Guzman. "Lidar for Autonomous Driving: The principles, challenges, and trends for automotive lidar and perception systems". In: *IEEE Signal Processing Magazine* (2020).

[29] Stefano Tavani et al. "Smartphone assisted fieldwork: Towards the digital transition of geoscience fieldwork using LiDAR-equipped iPhones". In: *Earth-Science Reviews* (2022).

[30] Kévin Walcarius et al. "Impact of blur on 3D laser imaging: Monte-Carlo modelling for underwater applications". In: *Optics Express* (2023).

[31] *IEC 60825-1:2014 — IEC Webstore*. URL: `https://webstore.iec.ch/publication/3587` (visited on 06/17/2023).

[32] Thinal Raj et al. "A Survey on LiDAR Scanning Mechanisms". In: *Electronics* (2020).

[33] Han Woong Yoo et al. "MEMS-based lidar for autonomous driving". In: *Elektrotechnik und Informationstechnik* (2018).

[34] Ryan G. Brazeal, Benjamin E. Wilkinson, and Hartwig H. Hochmair. "A Rigorous Observation Model for the Risley Prism-Based Livox Mid-40 Lidar Sensor". In: *Sensors* (2021).

[35] Richard M Marino and William R Davis. "Jigsaw: A Foliage-Penetrating 3D Imaging Laser Radar System". In: *Lincoln Laboratory Journal* (2005).

[36] N Muhammad and S Lacroix. "Calibration of a rotating multi-beam lidar". In: *IEEE/RSJ International Conference on Intelligent Robots and Systems*. 2010.

[37] Gerardo Atanacio-Jiménez et al. "LIDAR Velodyne HDL-64E Calibration Using Pattern Planes". In: *International Journal of Advanced Robotic Systems* (2011).

[38] *Why Digital Lidar is the Future*. URL: https://ouster.com/blog/why-digital-lidar-is-the-future (visited on 05/05/2023).

[39] Jenkins Godfrey, Vimal Kumar, and Shankar C Subramanian. "Evaluation of Flash LiDAR in Adverse Weather Conditions towards Active Road Vehicle Safety". In: *IEEE Sensors Journal* (2023).

[40] Benjamin Göhler and Peter Lutzmann. "Review on short-wavelength infrared laser gated-viewing at Fraunhofer IOSB". In: *Optical Engineering* (2016).

[41] Frank Christnacher et al. "Influence of gating and of the gate shape on the penetration capacity of range-gated active imaging in scattering environments". In: *Optics Express* (2015).

[42] Yoav Grauer. "Active gated imaging in driver assistance system". In: *Advanced Optical Technologies* (2014).

[43] Marcus Hedlund. *Weather Influence on LiDAR Signals using the Transient Radiative Transfer and LiDAR Equations.pdf*. Tech. rep. 2020. URL: https://www.diva-portal.org/smash/get/diva2:1445891/FULLTEXT01.pdf.

[44] *Cepton Products - Mass-Market Lidar Solutions*. URL: https://www.cepton.com/products/overview (visited on 05/05/2023).

[45] *4Sight™ Intelligent Lidar Platform*. URL: https://www.aeye.ai/4sight-lidar/ (visited on 05/05/2023).

[46] *10 Technical Advantages - Baraja*. en. URL: https://www.baraja.com/en/technology/10-technical-advantages (visited on 07/12/2023).

[47] Michael Vollmer, Klaus-Peter Möllmann, and Joseph A. Shaw. "The optics and physics of near infrared imaging". In: *Education and Training in Optics and Photonics*. 2015.

[48] Clément Mallet. "Analyse des données lidar aéroportées à Retour d'Onde Complète pour la cartographie des milieux urbains". fr. In: ().

[49] Erwan Viala et al. "Increasing spatial resolution of 3D LiDAR by Compressive Sensing". In: *Optro*. 2022.

[50] Julián Tachella et al. "Real-time 3D reconstruction from single-photon lidar data using plug-and-play point cloud denoisers". In: *Nature Communications* (2019).

[51] Mark A. Itzler et al. "Geiger-Mode LiDAR: From Airborne Platforms To Driverless Cars". In: *Imaging and Applied Optics 2017 (3D, AIO, COSI, IS, MATH, pcAOP)*. 2017.

[52] John Degnan. "Scanning, Multibeam, Single Photon Lidars for Rapid, Large Scale, High Resolution, Topographic and Bathymetric Mapping". In: *Remote Sensing* (2016).

[53] Hennes Henniger and Otakar Wilfert. "An Introduction to Free-space Optical Communications". In: *Radioengineering* (2010).

[54] *Bidirectional reflectance distribution function*. URL: https://en.wikipedia.org/w/index.php?title=Bidirectional_reflectance_distribution_function&oldid=1145432670 (visited on 07/12/2023).

[55] Addy Ngan, Frédo Durand, and Wojciech Matusik. "Experimental Analysis of BRDF Models". In: *Eurographics Symposium on Rendering* (2005).

[56] B Jutzi and U Stilla. "Precise range estimation on known surfaces by analysis of fullwaveform laser". In: *Fraunhofer FOM* (2006).

[57] *Home - Aeva*. URL: https://www.aeva.com/ (visited on 07/12/2023).

[58] Hansol Jang et al. "Simultaneous distance and vibration mapping of FMCW-LiDAR with akinetic external cavity diode laser". In: *Optics and Lasers in Engineering* (2023).

[59] Miles Hansard et al. *Time-of-Flight Cameras: Principles, Methods and Applications*. SpringerBriefs in Computer Science. London, 2013.

[60] Wolfgang Wagner et al. “Gaussian decomposition and calibration of a novel small-footprint full-waveform digitising airborne laser scanner”. In: *ISPRS Journal of Photogrammetry and Remote Sensing* (2006).

[61] Gustav Tolt, Christina Grönwall, and Markus Henriksson. “Peak detection approaches for time-correlated single-photon counting three-dimensional lidar systems”. In: *Optical Engineering* (2018).

[62] Santiago Royo and Maria Ballesta-Garcia. “An Overview of Lidar Imaging Systems for Autonomous Vehicles”. In: *Applied Sciences* (2019).

[63] Ji-Il Park, Jihyuk Park, and Kyung-Soo Kim. “Fast and Accurate Desnowing Algorithm for LiDAR Point Clouds”. In: *IEEE Access* (2020).

[64] *Inverse-square law*. URL: https://en.wikipedia.org/w/index.php?title=Inverse-square_law&oldid=1164420364 (visited on 07/12/2023).

[65] *How does LiDAR work ?* URL: https://www.yellowscan.com/knowledge/how-does-lidar-work/ (visited on 05/05/2023).

[66] *Definition for Multi-echo technology — SICK*. URL: https://www.sick.com/fr/en/glossary/multi-echo-technology/g/p555059 (visited on 05/05/2023).

[67] Marcel Kettelgerdes and Gordon Elger. “Modeling Methodology and In-field Measurement Setup to Develop Empiric Weather Models for Solid-State LiDAR Sensors”. In: *IEEE Journal of Radio Frequency Identification* (2023).

[68] Farzin Amzajerdian et al. “Imaging Flash Lidar for Autonomous Safe Landing and Spacecraft Proximity Operation”. In: *AIAA SPACE*. 2016.

[69] Marcus Hammer, Marcus Hebel, and Michael Arens. “Automated object detection and tracking with a flash LiDAR system”. In: *Electro-Optical Remote Sensing*. Ed. by Gary Kamerman and Ove Steinvall. 2016.

[70] Chao Cao, Marius Preda, and Titus Zaharia. “3D Point Cloud Compression: A Survey”. In: *International Conference on 3D Web Technology*. 2019.

[71] Markus Schütz, Stefan Ohrhallinger, and Michael Wimmer. “Fast Out-of-Core Octree Generation for Massive Point Clouds”. In: *Computer Graphics Forum* (2020).

[72] Andrew M. Wallace, Abderrahim Halimi, and Gerald S. Buller. “Full Waveform LiDAR for Adverse Weather Conditions”. In: *IEEE Transactions on Vehicular Technology* (2020).

[73] Jean-Luc Déziel et al. *PixSet : An Opportunity for 3D Computer Vision to Go Beyond Point Clouds With a Full-Waveform LiDAR Dataset*. 2021. URL: http://arxiv.org/abs/2102.12010.

[74] Martin Pfennigbauer et al. “Online waveform processing for demanding target situations”. In: *Defense + Security Symposium*. 2014.

[75] Christopher Vincent Poulton et al. “Long-Range LiDAR and Free-Space Data Communication With High-Performance Optical Phased Arrays”. In: *IEEE Journal of Selected Topics in Quantum Electronics* (2019).

[76] *Technology – Lumotive, Inc.* URL: https://lumotive.com/technology/ (visited on 07/12/2023).

[77] *Aeries II*. URL: https://www.aeva.com/aeries-ii/ (visited on 07/12/2023).

[78] Juha Suomalainen et al. “Demonstration of a virtual active hyperspectral LiDAR in automated point cloud classification”. In: *ISPRS Journal of Photogrammetry and Remote Sensing* (2011).

[79] Gregor Luetzenburg, Aart Kroon, and Anders A. Bjørk. “Evaluation of the Apple iPhone 12 Pro LiDAR for an Application in Geosciences”. In: *Scientific Reports* (2021).

[80] Shinichi Tatsumi, Keiji Yamaguchi, and Naoyuki Furuya. "ForestScanner: A mobile application for measuring and mapping trees with LiDAR-equipped iPhone and iPad". In: *Methods in Ecology and Evolution* (2022).

[81] Thierry Peynot, James Underwood, and Steven Scheding. "Towards reliable perception for Unmanned Ground Vehicles in challenging conditions". In: *IEEE/RSJ International Conference on Intelligent Robots and Systems*. 2009.

[82] Mario Bijelic et al. *Seeing Through Fog Without Seeing Fog: Deep Multimodal Sensor Fusion in Unseen Adverse Weather*. 2020. URL: http://arxiv.org/abs/1902.08913.

[83] Paul D. Groves. "Shadow Matching: A New GNSS Positioning Technique for Urban Canyons". In: *Journal of Navigation* (2011).

[84] João Pedro Machado dos Santos. "SmokeNav – Simultaneous Localization and Mapping in Reduced Visibility Scenarios". PhD thesis. Coimbra, 2013.

[85] Philip Church et al. "Evaluation of a steerable 3D laser scanner using a double Risley prism pair". In: *Degraded Environments: Sensing, Processing, and Display*. Ed. by John (Jack) N. Sanders-Reed and Jarvis (Trey) J. Arthur. 2017.

[86] Stephen J. Fenley and John (Jack) N. Sanders-Reed. "Visibility in degraded visual environments (DVE)". In: *Degraded Environments: Sensing, Processing, and Display*. Ed. by John (Jack) N. Sanders-Reed and Jarvis (Trey) J. Arthur. 2018.

[87] Hendrik Christoffel van de Hulst. *Light scattering by small particles*. Dover classics of science and mathematics. New York, 1981.

[88] Sinan Hasirlioglu. "A Novel Method forSimulation-basedTesting and Validationof Automotive SurroundSensors under AdverseWeather Conditions.pdf". PhD thesis. Linz, 2020.

[89] Emmanuel Villermaux and Benjamin Bossa. "Single-drop fragmentation determines size distribution of raindrops". In: *Nature Physics* (2009).

[90] J. S. Marshall and W. Mc K. Palmer. "The distribution of raindrop with size". In: *Journal of Meteorology* (1948).

[91] K. V. Beard. "Terminal Velocity Adjustment for Cloud and Precipitation Drops Aloft". In: *Journal of the Atmospheric Sciences* (1977).

[92] Ross Gunn and Gilbert D. Kinzer. "The terminal velocity of fall for water droplets in stagnant air". In: *Journal of Meteorology* (1949).

[93] Gopinath Kathiravelu, Terry Lucke, and Peter Nichols. "Rain Drop Measurement Techniques: A Review". en. In: *Water* 8.1 (Jan. 2016), p. 29. ISSN: 2073-4441. DOI: 10.3390/w8010029. URL: http://www.mdpi.com/2073-4441/8/1/29 (visited on 04/17/2023).

[94] *OTT Parsivel² - Laser Weather Sensor*. URL: https://www.ott.com/products/meteorological-sensors-26/ott-parsivel2-laser-weather-sensor-2392/ (visited on 05/05/2023).

[95] *Precipitation Sensors - WS100 Radar Precipitation Sensor / Smart Disdrometer*. URL: https://www.lufft.com/products/precipitation-sensors-287/ws100-radar-precipitation-sensor-smart-disdrometer-2361/ (visited on 07/12/2023).

[96] Pierre Tabary et al. "Le réseau et les produits radars de Météo-France". In: *La Météorologie* (2013).

[97] I. Gultepe et al. "Fog Research: A Review of Past Achievements and Future Perspectives". In: *Pure and Applied Geophysics* (2007).

[98] Pierre Duthon, Michèle Colomb, and Frédéric Bernardin. "Fog Classification by Their Droplet Size Distributions: Application to the Characterization of Cerema's Platform". In: *Atmosphere* (2020).

[99] *Smoke*. URL: https://en.wikipedia.org/wiki/Smoke (visited on 09/27/2023).

[100] Joseph W. Starr and B. Y. Lattimer. “Evaluation of Navigation Sensors in Fire Smoke Environments”. In: *Fire Technology* (2014).

[101] Andrei B. Utkin et al. “Feasibility of forest-fire smoke detection using lidar”. In: *International Journal of Wildland Fire* (2003).

[102] Mohammad Ali Khalighi and Murat Uysal. “Survey on Free Space Optical Communication: A Communication Theory Perspective”. In: *IEEE Communications Surveys and Tutorials* (2014).

[103] M. Ijaz et al. “Modeling of Fog and Smoke Attenuation in Free Space Optical Communications Link Under Controlled Laboratory Conditions”. In: *Journal of Lightwave Technology* (2013).

[104] David Crocker. *Dictionary of aviation*. 2007.

[105] International Civil Aviation Organization. *Meteorological service for international air navigation: international standards and recommended practices : annex 3 to the Convention on International Civil Aviation*. 2010.

[106] Erin McNamara. *Federal Meteorological Handbook Number 1, Surface Weather Observations and Reports*. Tech. rep. 2021. URL: `https://www.icams-portal.gov/resources/icams/fedrep/2021_fedrep.pdf`.

[107] E. M. Patterson, D. A. Gillette, and G. W. Grams. “The Relation Between Visibility and the Size-Number Distribution of Airborne Soil Particles”. In: *Journal of Applied Meteorology* (1976).

[108] Roger G. Newton. *Scattering Theory of Waves and Particles*. Theoretical and Mathematical Physics. Berlin, 1982.

[109] Aki Mayra et al. “Spectral attenuation in low visibility artificial fog: Experimental study and comparison to literature models”. In: *IEEE International Conference on Intelligent Computer Communication and Processing (ICCP)*. 2017.

[110] J. W. Giles et al. “Lidar system model for use with path obscurants and experimental validation”. In: *Applied Optics* (2008).

[111] R. H. Rasshofer, M. Spies, and H. Spies. “Influences of weather phenomena on automotive laser radar systems”. In: *Advances in Radio Science* (2011).

[112] Sven Teufel et al. “Simulating Realistic Rain, Snow, and Fog Variations For Comprehensive Performance Characterization of LiDAR Perception”. In: *Vehicular Technology Conference*. 2022.

[113] Kshitiz Garg Shree K. Nayar. “Photometric Model of a Rain Drop.pdf”. In: *Environmental Science, Physics* (2003).

[114] Piotr A. Lewandowski et al. “Lidar-Based Estimation of Small-Scale Rainfall: Empirical Evidence”. In: *Journal of Atmospheric and Oceanic Technology* (2009).

[115] Thomas Fersch et al. “The influence of rain on small aperture LiDAR sensors”. In: *German Microwave Conference (GeMiC)*. 2016.

[116] A. Filgueira et al. “Quantifying the influence of rain in LiDAR performance”. In: *Measurement* (2017).

[117] Christopher Goodin et al. “Predicting the Influence of Rain on LIDAR in ADAS”. In: *Electronics* (2019).

[118] Chen Zhang et al. “LiDAR Degradation Quantification for Autonomous Driving in Rain”. In: *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. 2021.

[119] M. Kutila et al. “Automotive LiDAR performance verification in fog and rain”. In: *2018 21st International Conference on Intelligent Transportation Systems (ITSC)*. 2018.

[120] Alexander Carballo et al. “LIBRE: The Multiple 3D LiDAR Dataset”. In: (2020). URL: `https://arxiv.org/abs/2003.06129`.

[121] *Explore: Light.* URL: https://lightcolourvision.org/explore-light/ (visited on 07/03/2023).

[122] J. Tuley, N. Vandapel, and M. Hebert. "Analysis and Removal of Artifacts in 3-D LADAR Data". In: *IEEE International Conference on Robotics and Automation*. 2005.

[123] Thomas Goelles, Birgit Schlager, and Stefan Muckenhuber. "Fault Detection, Isolation, Identification and Recovery (FDIIR) Methods for Automotive Perception Sensors Including a Detailed Literature Survey for Lidar". In: *Sensors* (2020).

[124] Jose Roberto Vargas Rivero et al. "Characterization and simulation of the effect of road dirt on the performance of a laser scanner". In: *International Conference on Intelligent Transportation Systems (ITSC)*. 2017.

[125] Manuel Trierweiler et al. "Influence of sensor blockage on automotive LiDAR systems". In: *IEEE SENSORS*. 2019.

[126] Robin Karlsson et al. "Probabilistic Rainfall Estimation from Automotive Lidar". In: (Apr. 2021).

[127] Romain Ceolato and Matthew J. Berg. "Aerosol light extinction and backscattering: A review with a lidar perspective". In: *Journal of Quantitative Spectroscopy and Radiative Transfer* (2021).

[128] Ilya Galaktionov et al. "Focusing of a Laser Beam Passed through a Moderately Scattering Medium Using Phase-Only Spatial Light Modulator". In: *Photonics* (2022).

[129] *Rayleigh scattering.* Apr. 2023. URL: https://en.wikipedia.org/w/index.php?title=Rayleigh_scattering&oldid=1149207455 (visited on 07/12/2023).

[130] *Mie scattering.* Mar. 2023. URL: https://en.wikipedia.org/w/index.php?title=Mie_scattering&oldid=1143880177 (visited on 07/12/2023).

[131] Tyson Govan Phillips, Nicky Guenther, and Peter Ross McAree. "When the Dust Settles: The Four Behaviors of LiDAR in the Presence of Fine Airborne Particulates". In: *Journal of Field Robotics* (2017).

[132] J. Pascoal, L. Marques, and A.T. De Almeida. "Assessment of Laser Range Finders in risky environments". In: *IEEE/RSJ International Conference on Intelligent Robots and Systems*. 2008.

[133] Mario Bijelic, Tobias Gruber, and Werner Ritter. "A Benchmark for Lidar Sensors in Fog: Is Detection Breaking Down?" In: *2018 IEEE Intelligent Vehicles Symposium (IV)*. 2018.

[134] Wenhua Song et al. "The effect of fog on the probability density distribution of the ranging data of imaging laser radar". In: *AIP Advances* (2018).

[135] Maria Jokela, Matti Kutila, and Pasi Pyykönen. "Testing and Validation of Automotive Point-Cloud Sensors in Adverse Weather Conditions". In: *Applied Sciences* (2019).

[136] You Li et al. "What Happens for a ToF LiDAR in Fog?" In: *IEEE Transactions on Intelligent Transportation Systems* (2020).

[137] Clemens Linnhoff et al. "Measuring the Influence of Environmental Conditions on Automotive Lidar Sensors". In: *Sensors* (2022).

[138] Will Northrop. *Can Automated Vehicles "See" in Minnesota? Ambient Particle Effects on LiDAR.* Tech. rep. 2022. URL: https://www.cts.umn.edu/publications/report/can-automated-vehicles-see-in-minnesota-ambient-particle-effects-on-lidar.

[139] Omer Sahin Tas et al. "Functional system architectures towards fully automated driving". In: *IEEE Intelligent Vehicles Symposium (IV)*. 2016.

[140] Abdul Sajeed Mohammed et al. "The Perception System of Intelligent Ground Vehicles in All Weather Conditions: A Systematic Literature Review". In: *Sensors* (2020).

[141] Sampo Kuutti et al. "A Survey of the State-of-the-Art Localization Techniques and Their Potentials for Autonomous Vehicle Applications". In: *IEEE Internet of Things Journal* (2018).

[142] Sagar Behere and Martin Törngren. “A Functional Architecture for Autonomous Driving”. In: *Proceedings of the First International Workshop on Automotive Software Architecture*. 2015.

[143] Rihards Novickis et al. “Functional Architecture for Autonomous Driving and its Implementation”. In: *Biennial Baltic Electronics Conference (BEC)*. 2020.

[144] Brian Paden et al. “A Survey of Motion Planning and Control Techniques for Self-Driving Urban Vehicles”. In: *IEEE Transactions on Intelligent Vehicles* (2016).

[145] Ekim Yurtsever et al. “A Survey of Autonomous Driving: *Common Practices and Emerging Technologies*”. In: *IEEE Access* (2020).

[146] Anh Nguyen and Bac Le. “3D point cloud segmentation: A survey”. In: *IEEE Conference on Robotics, Automation and Mechatronics (RAM)*. 2013.

[147] E. Grilli, F. Menna, and F. Remondino. “A review of point clouds segmentation and classification algorithms”. In: *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences* (2017).

[148] Yuxing Xie, Jiaojiao Tian, and Xiao Xiang Zhu. “Linking Points With Labels in 3D: A Review of Point Cloud Semantic Segmentation”. In: *IEEE Geoscience and Remote Sensing Magazine* (2020).

[149] Thibault Lejemble. “Analyse multi-échelle de nuage de points”. PhD thesis. Université de Toulouse, 2020.

[150] Mariana Batista Campos et al. “A Long-Term Terrestrial Laser Scanning Measurement Station to Continuously Monitor Structural and Phenological Dynamics of Boreal Forest Canopy”. In: *Frontiers in Plant Science* (2021).

[151] A. Kharroubi et al. “Classification and integration of massive 3D point clouds in a virtual reality (VR) environment)”. In: *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences* (2019).

[152] Tiago Gomes et al. “A Survey on Ground Segmentation Methods for Automotive LiDAR Sensors”. In: *Sensors* (2023).

[153] *XenomatiX — True Solid State Lidar*. URL: https://xenomatix.com/ (visited on 05/10/2023).

[154] Amit Saxena et al. “A review of clustering techniques and developments”. In: *Neurocomputing* (2017).

[155] Florent Poux and Roland Billen. “Voxel-based 3D Point Cloud Semantic Segmentation: Unsupervised Geometric and Relationship Featuring vs Deep Learning Methods”. In: *ISPRS International Journal of Geo-Information* (2019).

[156] Yuxiao Zhang et al. “Perception and sensing for autonomous vehicles under adverse weather conditions: A survey”. In: *ISPRS Journal of Photogrammetry and Remote Sensing* (2023).

[157] Mariella Dreissig et al. *Survey on LiDAR Perception in Adverse Weather Conditions*. arXiv:2304.06312 [cs]. 2023. URL: http://arxiv.org/abs/2304.06312.

[158] D.C. Cooper. “Multiple Target Tracking with Radar Applications”. In: *Electronics and Power* (1987).

[159] Simo Särkkä. *Bayesian Filtering and Smoothing*. Institute of Mathematical Statistics Textbooks. 2013.

[160] Xian-Feng Han et al. “A review of algorithms for filtering the 3D point cloud”. In: *Signal Processing: Image Communication* (2017).

[161] Abu Ubaidah Shamsudin et al. “Fog removal using laser beam penetration, laser intensity, and geometrical features for 3D measurements in fog-filled room”. In: *Advanced Robotics* (2016).

[162] Nicholas Charron, Stephen Phillips, and Steven L. Waslander. “De-noising of Lidar Point Clouds Corrupted by Snowfall”. In: *Conference on Computer and Robot Vision (CRV)*. 2018.

[163] Robin Heinzler et al. “CNN-based Lidar Point Cloud De-Noising in Adverse Weather”. In: *IEEE Robotics and Automation Letters* (2020).

[164] Yao Duan and Chuanchuan Yang. “Low-complexity Point Cloud Filtering for LiDAR by PCA-based Dimension Reduction”. en. In: (2021).

[165] Akhil Kurup and Jeremy Bos. *DSOR: A Scalable Statistical Filter for Removing Falling Snow from LiDAR Point Clouds in Severe Winter Weather*. arXiv: 2109.07078. 2021. URL: http://arxiv.org/abs/2109.07078 (visited on 12/07/2021).

[166] *Point Cloud Library*. URL: https://pointclouds.org/ (visited on 04/19/2023).

[167] Leo Stanislas. “Detecting airborne particles in sensor data with deep learning for robust robot perception in adverse environments”. PhD. Queensland University of Technology, 2020.

[168] Tzu-Hsien Sang, SungYou Tsai, and TsungPo Yu. “Mitigating Effects of Uniform Fog on SPAD Lidars”. In: *IEEE Sensors Letters* (2020).

[169] *IEC TS 62998-1:2019 — IEC Webstore*. URL: https://webstore.iec.ch/publication/31009 (visited on 07/03/2023).

[170] Mario Jürgen Berk. “Safety Assessment of Environment Perception in Automated Driving Vehicles”. In: *Environmental Science* (2019).

[171] Wang Hong et al. “Safety Decision of Running Speed Based on Real-time Weather”. In: *IEEE Intelligent Vehicles Symposium (IV)*. 2022.

[172] Florian Gaudfrin. “Lidar supercontinuum pour la caractérisation spectrale des milieux diffusants à haute résolution spatiale : Étude numérique et développement instrumental”. PhD thesis. Toulouse, 2020.

[173] Danick Perrin. “Rainfall estimation using low-cost lidar measurements”. MA thesis. University of Lausanne, 2018.

[174] Tao Yang et al. “LaNoising: A Data-driven Approach for 903nm ToF LiDAR Performance Modeling under Fog”. In: *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. 2020.

[175] Jose Roberto Vargas Rivero et al. “Weather Classification Using an Automotive LIDAR Sensor Based on Detections on Asphalt and Atmosphere”. In: *Sensors* ().

[176] George Sebastian et al. “RangeWeatherNet for LiDAR-only weather and road condition classification”. In: *IEEE Intelligent Vehicles Symposium (IV)*. 2021.

[177] Robin Klaus Heinzler. “LiDAR-based Weather Detection: Automotive LiDAR Sensors in Adverse Weather Conditions”. PhD thesis. Karlsruhe, 2022.

[178] Matthew Pitropov et al. *Canadian Adverse Driving Conditions Dataset*. 2020. URL: https://arxiv.org/abs/2001.10117.

[179] Nicolas Rivière and Paul-Edouard Dupouy. “Perception lidar 3D pour l’aide à la conduite autonome”. In: *Photoniques* (2022).

[180] Guy Satat, Matthew Tancik, and Ramesh Raskar. “Towards photography through realistic fog”. In: *IEEE International Conference on Computational Photography (ICCP)*. 2018.

[181] *Plateforme de simulation de conditions climatiques dégradées*. URL: http://www.cerema.fr/fr/innovation-recherche/innovation/offres-technologie/plateforme-simulation-conditions-climatiques-degradees (visited on 05/05/2023).

[182] *Ultra Puck Surround View Lidar Sensor*. URL: https://velodynelidar.com/products/ultra-puck/ (visited on 05/05/2023).

[183] *High-resolution OS1 lidar sensor: robotics, trucking, mapping — Ouster*. URL: https://ouster.com/products/scanning-lidar/os1-sensor/ (visited on 07/12/2023).

[184] *Livox Horizon user manual v1.0.pdf*. URL: https://www.livoxtech.com/downloads (visited on 05/05/2023).

[185] Zheng Liu, Fu Zhang, and Xiaoping Hong. "Low-Cost Retina-Like Robotic Lidars Based on Incommensurable Scanning". In: *IEEE/ASME Transactions on Mechatronics* (2022).

[186] J. Wojtanowski et al. "Comparison of 905 nm and 1550 nm semiconductor laser rangefinders' performance deterioration due to adverse environmental conditions". In: *Opto-Electronics Review* (2014).

[187] Dingkang Wang, Connor Watkins, and Huikai Xie. "MEMS Mirrors for LiDAR: A Review". In: *Micromachines* (2020).

[188] Adrien Guyot et al. "Effect of disdrometer type on rain drop size distribution characterisation: a new dataset for south-eastern Australia". In: *Hydrol. Earth Syst. Sci.* (2019).

[189] P. K. Wang and H. R. Pruppacher. "Acceleration to Terminal Velocity of Cloud and Raindrops". In: *Journal of Applied Meteorology* (1977).

[190] D. Atlas, R. C. Srivastava, and R. S. Sekhon. "Doppler radar characteristics of precipitation at vertical incidence". In: *Reviews of Geophysics* (1973).

[191] Romain Ceolato, Andres E. Bedoya-Velásquez, and Vincent Mouysset. "Short-Range Elastic Backscatter Micro-Lidar for Quantitative Aerosol Profiling with High Range and Temporal Resolution". In: *Remote Sensing* (2020).

[192] *What is Artificial Intelligence (AI) ?* URL: https://www.ibm.com/topics/artificial-intelligence (visited on 07/10/2023).

[193] *Artificial intelligence*. URL: https://en.wikipedia.org/w/index.php?title=Artificial_intelligence&oldid=1164288564 (visited on 07/10/2023).

[194] Daphne Koller and Nir Friedman. *Probabilistic graphical models: principles and techniques*. Adaptive computation and machine learning. London, 2009.

[195] Christopher M. Bishop. *Pattern recognition and machine learning*. Information science and statistics. New York, 2006.

[196] Kevin P. Murphy. *Machine learning: a probabilistic perspective*. Adaptive computation and machine learning series. London, 2012.

[197] Sebastian Thrun, Wolfram Burgard, and Dieter Fox. *Probabilistic robotics*. Intelligent robotics and autonomous agents. London, 2006.

[198] Stefan Weinzierl. *Introduction to Monte Carlo methods*. 2000. URL: http://arxiv.org/abs/hep-ph/0006269.

[199] Sven Ronnback and Ake Wernersson. "On filtering of laser range data in snowfall". In: *IEEE International Conference "Intelligent Systems" (IS)*. 2008.

[200] Sebastien Michaud, Jean-Francois Lalonde, and Philippe Giguere. "Towards Characterizing the Behavior of LiDARs in Snowy Conditions". In: *IEEE International Conference "Intelligent Robots and Systems" (IROS)*. 2015.

[201] Daryl J. Daley and D. Vere-Jones. *An introduction to the theory of point processes*. Probability and its applications. New York, 2003.

[202] Ronald P. S. Mahler. *Advances in statistical multisource-multitarget information fusion*. 2014.

[203] Branko Ristic et al. “A Tutorial on Bernoulli Filters: Theory, Implementation and Applications”. In: *IEEE Transactions on Signal Processing* (2013).