WEATHERSAFE

SAFETY ENHANCEMENT UNDER ADVERSE WEATHER CONDITIONS THROUGH IMPROVED ADAS/AD ROBUSTNESS

Deliverable D_Simu.2

Results from the comparison between real and synthetic sensor data

Project Name	WeatherSafe	
Dissemination level	Consortium	
Written by	Dr. Philipp Rosenberger	Persival GmbH
Issue date	31/03/2026	
Reviewer	Pierre Duthon, Sébastien Liandrat, Amine Ben-Daoued	
Keywords	Adverse Weather, Simulation, Lidar, Radar, Camera	
Version	1.1	

EXECUTIVE SUMMARY

This deliverable (D_Simu.2) provides a traceable basis for comparing real and synthetic perception sensor data under adverse environmental conditions. It combines (i) a measurement-driven performance evaluation of a camera and four lidar sensors recorded at the Persival Weather Test Bench during winter 2025/2026 with (ii) a methodological "role model" for simulation and validation based on Linnhoff (2023). Together, these elements support the Simulation Work Package by linking measurable weather severity parameters to sensor-level effects and by outlining how such effects can be implemented and validated in simulation.

The Persival Weather Test Bench is a stationary rooftop setup with protected Lambertian reference targets at 21.5 m and 78.5 m and synchronized reference sensing for precipitation, brightness and fog (including meteorological optical range, MOR, and the forward-scatter extinction coefficient). For the camera, mean intensity, variance/SNR and inter-target contrasts were evaluated on fixed regions of interest. For lidar, target detections and intensities/reflectivities were evaluated together with atmospheric "noise" detections, expressed as a detection probability per beam in a defined atmospheric region. Across all analyses, the KPIs are related to quantitative weather descriptors (lux, mm/h, extinction), yielding performance curves and distributions that can be used for model calibration and plausibility checks of synthetic data.

The results show distinct degradation signatures depending on sensor modality and weather type. For the camera, fog acts as a global visibility limiter, producing a consistent contrast reduction with increasing extinction and a simultaneous intensity lift due to veiling luminance (airlight). In contrast, rain, snow and hail mainly increase KPI variability through transient, spatially inhomogeneous disturbances (streaks/occlusions, glare), without a similarly monotonic contrast attenuation within the measured ranges; brightness-dependent camera adaptation was identified as an additional confounder, especially for contrasts involving the darkest reference. For lidar, degradation is strongly configuration-dependent (FoV, range class, wavelength): fog is the dominant object-independent effect, driving range-dependent target detection loss and (for rotating sensors) pronounced atmospheric backscatter detections; patchy fog can additionally generate clustered atmospheric returns for directional sensors. In the observed precipitation regimes, rain and snow show low atmospheric detection probabilities and largely stable target detections, while hail introduces measurable atmospheric detections above roughly 10 mm/h but does not substantially reduce target detectability.

From a simulation and credibility-assessment perspective, the findings motivate a separation between object-independent and object-dependent environmental effects and support performance-driven model parameterization. Fog should be represented with explicit attenuation/veiling (camera) and attenuation plus stochastic atmospheric-return modeling (lidar), parameterized by extinction (or MOR). Precipitation effects are better represented as stochastic disturbance processes that change variability and generate outliers rather than purely deterministic attenuation curves, with precipitation rate treated as a proxy of limited transferability. These conclusions directly support T_Simu.2 by clarifying which environmental parameters must be represented and how they map to sensor-level effects, and they support T_Simu.3 by providing measurable reference curves and validation criteria (trend agreement and variability envelopes) for comparing real and synthetic sensor data.
The deliverable is structured along the main tasks of the Simulation Work Package:

Task	Description
T_Simu.1	State of the art: review of existing European simulation means, studies and works dealing with adverse weather conditions in road safety based on PerCollECT.
T_Simu.2	Adverse weather conditions simulation: based on the work in WP1 and WP2, possibilities for implementing relevant weather effects in sensor simulation are demonstrated.
T_Simu.3	Validation of weather effect modelling: comparison of real and synthetic sensor data with weather influence based on a case study
T_Simu.4	Simulations' assessment guidelines: Constitute and promote a list of recommendations to evaluate different aspects of the simulation of adverse weather, such as physical-numerical correlation and methodologies as addon for WP3.

The work package results are summarized in three deliverables:

Task	Description
D_Simu.1	Report that summarizes the state of the art (T_Simu.1) and the prospect of simulation means able to reproduce adverse weather conditions.
D_Simu.2	Results from the comparison between real and synthetic sensor data with weather influence based on a case study. It reviews discussions, works and outputs of T_Simu.2 and T_Simu.3. Especially the methodologies, parameters and requirements needed to simulate adverse weather conditions. It will include the description of the variables usable to play a scenario and sensor perception levels.
D_Simu.3	Guidelines for sensor model validation including the important points to consider during data collection. It concludes on project guidelines (T_Simu.4) on the assessment of simulation and the parameters to be considered.

REVISION HISTORY

Revision	Date	Description, updates and changes	Status
1.0	31/03/2026	Consolidated version, open for final review	Draft
1.1	26/05/2026	Consolidated version after final review	Approved

ABBREVIATIONS

ASAM OSI	ASAM Open Simulation Interface
eCDF	Empirical cumulated distribution function
FMI	Functional Mock-up Interface
FMU	Functional Mock-up Unit
FoV	Field of View
KPI	Key Performance Indicator
Lidar	Light Detection and Ranging
MOR	Meteorological Optical Range
OSMP	OSI Simulation Model Packaging
Radar	Radio Detection and Ranging
ROI	Region of Interest

TABLE OF CONTENTS

EXECUTIVE SUMMARY 2
REVISION HISTORY 4
ABBREVIATIONS 5
TABLE OF CONTENTS 6
1. INTRODUCTION 7
2. PERSIVAL WEATHER TEST BENCH 9
3. LIDAR SENSOR CHARACTERISTICS 15
4. ACTUAL PERFORMANCE CURVES FROM WINTER 2025/2026 17
5. SIMULATION OF LIDAR IN BAD WEATHER 89
6. CONCLUSION 99
7. ACKNOWLEDGEMENTS 101
8. BIBLIOGRAPHY 102
9. TABLE OF ILLUSTRATIONS 103

1. Introduction

The assessment of perception system robustness under adverse environmental conditions requires a clear understanding of how real sensors behave in the presence of weather influences. While simulation plays a central role in the development and validation of advanced driver assistance systems and automated driving functions, its credibility ultimately depends on how well simulated sensor outputs reflect measured real‑world performance. This deliverable therefore focuses on the evaluation of real perception sensor performance for both camera and Lidar modalities, under controlled and reproducible adverse weather conditions.

Cameras and lidar sensors are affected by environmental influences in fundamentally different ways. Camera performance is impacted by reduced visibility, contrast loss, glare, occlusions, and changes in scene appearance caused by fog, rain, snow, hail, or surface wetting. Lidar sensors, in contrast, are primarily affected by atmospheric attenuation, scattering, backscatter, and signal‑to‑noise degradation, with effects that depend on distance, wavelength, and sensor configuration. As a result, a meaningful evaluation of perception performance as the basis for sensor simulation models must be based on sensor‑specific performance measurements with reference weather sensors.

This deliverable introduces the Persival weather test bench, which has been explicitly designed to characterize perception sensor performance under adverse environmental conditions. The test bench enables the recording of weather phenomena such as rain, fog, snow, and hail, while simultaneously capturing high‑quality weather sensor data as the reference for modeling. By capturing the environmental parameters and observing their impact on sensor outputs, the Persival test bench makes it possible to derive quantitative performance curves for both camera and Lidar sensors. These curves describe how key performance indicators, such as detection range, noise level, uncertainty, and robustness, evolve as a function of weather severity and scene configuration.

For camera sensors, the focus of this deliverable is on the empirical characterization of performance degradation. Rather than attempting to directly model all physical effects in simulation, the measured data provide the possibility to model camera performance on a data-driven, phenomenological level based on performance curves. Additionally, the results establish reference performance envelopes that can be used to assess the realism and relevance of camera simulation outputs.

For lidar sensors, the deliverable goes one step further by demonstrating how sensor models can be derived directly from measured performance data. Using the controlled measurements obtained on the Persival weather test bench, lidar performance curves are extracted and translated into simulation‑ready representations. These models capture the combined effects of atmospheric conditions, sensor geometry, wavelength, and signal processing, without relying on manufacturer specific implementations. This performance driven
approach enables lidar simulation models to reproduce realistic degradation behavior under adverse weather, supporting quantitative comparison between real and synthetic data.

By jointly addressing camera and Lidar performance, this deliverable provides a consistent experimental foundation for subsequent simulation and validation activities. It establishes the link between real sensor behavior, performance oriented sensor modeling, and credible simulation inputs, which is essential for safety oriented evaluation under adverse

environmental conditions. The results presented here directly support the objectives of the Simulation Work Package by enabling informed implementation choices in simulation (T_Simu.2) and by providing data for the comparison of real and synthetic sensor outputs (T_Simu.3).

Please note that these results correspond to an initial, upstream work. Elements in this deliverable are preliminary examples showing how datasets could be constructed to train models, and to present one example model derived from a specific dataset. This is neither exhaustive nor complete, but rather indicative and a starting point for further work. In this context, focus is made on camera an LiDAR performances only.

1.1 SCOPE AND DATA BASIS FOR SENSOR MODELING AND VALIDATION

Within the WeatherSafe project, no measurement data was available that could be directly used for perception sensor modeling or quantitative model validation. This was primarily a consequence of the late inclusion of the simulation-focused work package, as well as the different objectives pursued during the joint measurement workshops conducted at the various partner facilities. These workshops were not designed with the specific requirements of perception sensor modeling and validation in mind, such as controlled variation of environmental parameters, reproducible reference targets, or synchronized ground-truth signals.

To nevertheless demonstrate the general procedure and methodological requirements for environment-aware perception sensor modeling, additional measurement data was collected at the Persival Weather Test Bench. These measurements were explicitly designed to illustrate how controlled environmental conditions, reference targets, and sensor data can be combined to support model development and validation. However, as this data acquisition campaign concluded together with the project, the collected data could not be used to derive or calibrate final sensor models within the scope of WeatherSafe.

As a result, the modeling and model validation aspects presented in this deliverable are primarily based on the methodological report by Linnhoff (Linnhoff, 2023), which provides a detailed and systematic description of the relevant steps for comparing real and synthetic sensor data under adverse environmental conditions. Rather than serving as project-specific validation results, this material is intentionally used as a role model to illustrate best practices for environmental effect modeling, parameter validation, and credibility assessment. The concepts and evaluation patterns described therein are therefore transferable and serve as guidance for future sensor modeling activities within and beyond the WeatherSafe context.

Overall, this deliverable contributes to closing the gap between experimental testing and virtual validation. By grounding simulation in measured sensor performance obtained from dedicated weather testing infrastructure, it supports the development of robust, traceable, and trustworthy assessment methodologies for perception systems operating under adverse weather conditions.

2. Persival Weather Test Bench

The measurement setup consists of a stationary unit and two protected target housings positioned at defined distances on a roof in Darmstadt, Germany. The station includes a waterproof enclosure for the power supply and the measurement PC, as well as modularly mounted perception sensors and reference weather sensors. The targets are placed at 21.5 m and 78.5 m from the station inside weather-protected housings. The front housing contains four 0.5 m × 0.5 m diffuse reflecting targets with reflectance levels of 5%, 50%, 80%, and 95%, while the rear housing holds a single 1 m × 1 m target with 10% reflectance.

Measurements are triggered automatically when defined thresholds of weather parameters are reached, as reported by the reference sensors. These parameters include precipitation rate, ambient temperature, relative humidity, air pressure, wind velocity and direction, brightness, and SYNOP present weather codes, as well as fog-related quantities such as meteorological optical range (MOR) and extinction coefficient. This automated approach ensures that data is collected under representative environmental conditions without manual intervention.



Figure 1: Persival Weather Test Bench setup on a roof top

2.1 Sensor station

The measurement station consists of a waterproof housing that contains the power supply and an OnLogic Karbon 800 industrial PC running Ubuntu and ROS 2 for data acquisition and synchronization. The perception sensors under test are positioned at the front of the station at a height of 0.4–0.6 m above ground level, and their exact mounting positions are recorded in the /tf_static topic within the measurement data for precise spatial referencing.



Figure 2: Weather station with lidars and a camera in the front and weather reference sensors on top

To capture the present weather conditions, the station is equipped with multiple reference sensors. The general weather sensor (Thies Clima 4.9200.00.000) operates at a measurement interval of 0.5 s and provides:

- Precipitation rate (mm/h)
- Ambient temperature (°C)
- Relative humidity (%)
- Air pressure (hPa)
- Wind velocity and direction (m/s)
- Brightness (lx) with light source direction
- SYNOP present weather code for classification

A disdrometer (Thies Laser Precipitation Monitor) with a 60 s interval measures:

- Number of drops per drop size and velocity bin

Additionally, a fog sensor (Biral SWS-050) with a 10 s interval provides:

- MOR (m)
- Extinction coefficient (1/m)

2.2 TARGETS

The measurement setup in winter season 2025/2026 included two target housings that protected the targets from environmental influences. One housing was positioned at a distance of 78.5 m from the station and contains a single target with dimensions of 1 m × 1 m:

- Sphere Optics Zenith Lite Diffuse Reflecting Target, 10% reflectance (PTFE foil) mounted on a honeycomb structure board

The second housing was located at 21.5 m and contained four targets provided by the customer, each measuring 0.5 m × 0.5 m, with different reflectance levels:

- Target 2: Sphere Optics Zenith Lite Diffuse Reflecting Target, 95%
- Target 3: Sphere Optics Zenith Lite Diffuse Reflecting Target, 5%
- Target 4: Sphere Optics Zenith Lite Diffuse Reflecting Target, 50%
- Target 5: Sphere Optics Zenith Lite Diffuse Reflecting Target, 80%



Figure 3: Target housing with 4 different Lambertian targets

These targets provided controlled and reproducible reflection characteristics for evaluating the perception sensor performance under varying environmental conditions.

2.3 MEASUREMENT AUTOMATION

The station records data of the reference sensors as well as the perception sensors listed below in ROS2 MCAP files. Recordings are started automatically depending on the following thresholds of the different weather sensors. The thresholds are derived from commonly used classifications of weather conditions.

- MOR in m: 2000, 1000, 500, 220
- Precipitation intensity in mm/h: 1.9, 8.1, 34, 50
- SYNOP: 65, 69, 70, 71, 72, 73, 74, 87, 88, 89, 90

Recordings are stopped after a duration of 2 min.

Additionally, daily verification measurements are taken automatically at 04:00 and 13:00. These measurements are only taken in clear conditions.

2.4 DATA STRUCTURE

The measurement data of all perception sensors and the weather reference sensors is captured in ROS2 MCAP files. On the first hierarchy level, the data is sorted into folders according to the weather category: fog, hail, rain, and snow. The daily measurements are in a separate folder. Each measurement of the respective category has its own folder with the date and time of the start of the measurement. Within that folder, there is the MCAP file with the corresponding metadata.yaml.

Additionally, each measurement has a JSON file with the same file name as the measurement. This JSON file contains information about the weather condition in the measurement, such as minimum, mean and maximum values for MOR, brightness, precipitation intensity, etc. The folder also contains preprocessed lidar data in a Python Pickle file for each sensor. These Pickle files contain evaluations of the number of detections, intensity, reflectivity and distance on the

5 targets and in the atmosphere for each timestep in the measurement. For each lidar timestep, the matching weather sensor data is also contained.

2.5 EVALUATION PROCEDURE FOR CAMERAS

In the first step, all measurements are classified into four categories: fog, rain, snow, and hail. Rain, snow and hail are distinguished by their respective SYNOP codes, while fog is identified by low visibility combined with a non-precipitation SYNOP code. Brightness is considered as another category. The daily reference measurements are taken for this evaluation.
For each category, environmental parameters (e.g., precipitation rate in mm/h, extinction coefficient in 1/meters or brightness in lux) are paired with perception sensor KPIs across all measurements within that category. Filter regions in the images are defined, to evaluate a square part of the image on each of the 5 targets. The following KPIs are all calculated per target area by always considering 3 consecutive images.

- Mean intensity
- Variance of intensity
- Signal to noise ratio (mean / variance)

Additionally, the contrast between pairs of target regions is calculated as the difference of the means divided by the sum of the means.
The following image shows the 5 target regions in the image.



Figure 4: Different target regions in the camera image

2.6 EVALUATION PROCEDURE FOR LIDARS

In the first step, all measurements are classified into four categories: fog, rain, snow, and hail. Rain, snow and hail are distinguished by their respective SYNOP codes, while fog is identified by low visibility (high extinction coefficient) combined with a non-precipitation SYNOP code. Brightness is considered as another category. The daily reference measurements are taken for this evaluation.

For each category, environmental parameters (precipitation rate in mm/h, extinction coefficient in 1/meters or brightness in lux) are paired with sensor detection data across all measurements within that category. The analysis focuses on the (relative) number of detections and the detection intensities to quantify environmental impact.
Two kinds of regions of interest are defined for filtering sensor detections:

- Atmospheric Region: This region covers azimuth angles from 0° to 10° and a maximum distance of 20 m from the station. Under clear conditions, no detections are expected in this region; detections originate only from precipitation or fog. The filter excludes ground reflections by applying only positive elevation angles.
- Target Regions: A Cartesian filter is applied to isolate detections corresponding to individual targets.

The primary metric to determine the influence of a weather condition on the perception sensor in the atmosphere is the relative number of detections, defined as the ratio of detection count within the filter region to the total number of beams entering that region. This yields a detection probability per beam under given environmental conditions. On the targets, the absolute numbers of detections are considered.
As a second metric, the intensities of the individual detection points are evaluated. They are affected by the attenuation due to environmental influences.

For each evaluation, not only measurements from the respective weather categories are considered, but also the daily reference measurements. This way, a solid baseline for 0 mm/h precipitation intensity and 0 1/m extinction coefficient is established and present in all plots.

2.7 FORWARD SCATTER EXTINCTION COEFFICIENT FOR FOG MEASUREMENTS

In fog, the degradation of optical sensor performance and human visibility is governed by the attenuation of electromagnetic radiation due to scattering by suspended water droplets. This attenuation is quantified by the extinction coefficient, which describes the combined effect of scattering and absorption per unit path length. For fog, absorption is typically negligible in the visible and near-infrared spectral range, such that extinction is dominated by scattering processes. The forward scatter extinction coefficient is commonly measured using visibility sensors that emit light over a short baseline and evaluate the amount of light scattered in the forward direction by fog droplets. This coefficient is directly related to the atmospheric optical density and thus provides a quantitative measure of how strongly the medium attenuates propagating light.

Meteorological visibility is defined as the maximum distance at which a black object of suitable dimensions can be recognized against the horizon sky with a given contrast threshold. Under internationally accepted standards, this contrast threshold is typically set to 5%, which links visibility to the exponential decay of contrast with distance as described by the Beer–Lambert law. As the extinction coefficient increases due to higher droplet concentration or larger droplet sizes in fog, visibility decreases proportionally. Conversely, lower extinction coefficients correspond to clearer atmospheric conditions and greater visibility ranges.

The forward scatter extinction coefficient therefore serves as a physically meaningful proxy for visibility, as both parameters originate from the same underlying scattering mechanisms. Importantly, this relationship is largely independent of the observer or sensor geometry and depends primarily on the microphysical properties of the fog, such as droplet size distribution and liquid water content.

For perception sensor evaluation and simulation, this correlation enables the use of measured extinction coefficients to parameterize environmental conditions in a reproducible and

quantitative manner. By mapping extinction values to corresponding visibility ranges, simulated fog conditions can be aligned with real-world meteorological classifications, allowing consistent comparison between real sensor measurements and synthetic data. This forms a crucial basis for credible modeling of environment-induced performance degradation in optical perception systems.

3. Lidar sensor characteristics

3.1 Sensor characteristics and implications for evaluation

Perception sensors of the same modality can exhibit significantly different behavior depending on their field of view configuration, effective range, and angular resolution, even when they share a comparable number of laser beams. In particular, sensors optimized for wide vertical coverage distribute their measurement channels over a larger angular span, resulting in strong near-field awareness but a reduced effective detection range. Conversely, sensors with a narrower vertical field of view concentrate their sensing elements over a smaller angular region, leading to increased angular resolution and improved mid- to long-range perception.

These design trade-offs have direct implications for sensor performance evaluation, especially under adverse environmental conditions. Wide-field configurations tend to be more sensitive to near-field phenomena such as spray, ground reflections, and clutter close to the vehicle, while narrower-field configurations are more affected by atmospheric attenuation, scattering, and signal loss over distance in conditions such as fog, rain, or snowfall. As a result, performance degradation profiles under adverse weather are not only sensor-technology-dependent, but also strongly influenced by the configuration of the sensor.

For evaluation purposes, this means that sensor performance should not be assessed solely on nominal specifications, but rather in terms of measured performance curves that relate detection probability, effective range, noise characteristics, and uncertainty to environmental parameters. Such curves enable a consistent comparison across different sensor configurations and provide a meaningful basis for performance-driven simulation. They also allow weather intensity to be treated as a continuous or discretized parameter that directly influences sensor output quality.

From a simulation perspective, these differences underline the importance of using configuration-aware sensor models. Applying identical degradation parameters to sensors with different fields of view or range characteristics can lead to unrealistic results and biased safety assessments. Instead, sensor evaluation data should be used to calibrate simulation models in a way that preserves the relative strengths and limitations of each sensor configuration.

This approach aligns with recent standardization efforts that emphasize performance-based evaluation under adverse conditions. It also supports the separation of concerns between sensor characterization, simulation implementation, and scenario definition. Within the WeatherSafe project, this principle provides a clear rationale for evaluating sensors in the second deliverable and using the resulting performance descriptions as inputs to subsequent simulation and validation activities.

3.2 WAVELENGTH CONSIDERATIONS AND IMPLICATIONS FOR PERFORMANCE UNDER ADVERSE WEATHER

In addition to geometric configuration and range optimization, operating wavelength is a key differentiating factor between Lidar sensor designs and has a measurable impact on performance under adverse environmental conditions. Automotive lidar systems commonly operate in the near-infrared (NIR) spectrum, with typical wavelength bands around 850 nm, 905 nm, 940 nm, 1310 nm, or 1550 nm. While all of these wavelengths are eye-safe when properly implemented, they interact differently with the environment and with external interference sources.

Lidar sensors operating at shorter NIR wavelengths generally benefit from high detector sensitivity and mature optoelectronic components. However, these wavelengths tend to exhibit stronger scattering and backscatter effects in the presence of fog, rain, and snow, which can increase noise levels and false detections under adverse weather conditions. They are also more susceptible to interference from solar radiation, particularly under high ambient light conditions.

Sensors operating around 905 nm represent a widely used compromise between component availability, performance, and regulatory constraints. This wavelength range is moderately affected by atmospheric scattering and precipitation, and its behavior under adverse weather is well documented in experimental studies. As a result, many existing performance evaluation methodologies and test infrastructures are implicitly aligned with this wavelength band, making it a common reference point for comparative assessment.

Lidar systems operating at longer NIR wavelengths tend to show reduced sensitivity to solar background radiation and, in some cases, improved robustness to certain atmospheric effects. The longer wavelength can lead to slightly lower scattering efficiency for small particles, which may reduce backscatter noise in fog or snowfall. At the same time, detector technology and signal-to-noise characteristics differ, meaning that performance advantages are strongly dependent on system design rather than wavelength alone.

From an evaluation and simulation perspective, these wavelength differences imply that weather-induced performance degradation cannot be treated uniformly across lidar sensors, even when their geometric configurations are similar. Atmospheric attenuation, backscatter probability, and noise characteristics vary with wavelength and should therefore be reflected in wavelength-aware performance curves derived from experimental testing. This is particularly relevant for adverse weather scenarios involving fog, heavy rain, or snow, where particle size distributions interact differently with optical wavelengths.

For simulation, this reinforces the need to parameterize lidar sensor models not only by field of view, range, and angular resolution, but also by operating wavelength class. Incorporating wavelength-dependent behavior enables more realistic reproduction of detection range loss, noise increase, and false return generation under adverse weather. Such differentiation is consistent with performance-oriented evaluation approaches and supports the development of simulation models that remain credible across heterogeneous sensor configurations without relying on manufacturer-specific implementations.

4. Actual performance curves from winter 2025/2026

4.1 Actual performance curves of a camera

4.1.1 Camera performance in brightness

Around 1859 measurement time steps were taken from the daily measurements during daytime for the brightness evaluation. The weather sensor measures brightness in lux in 4 directions and additionally gives the maximum brightness in lux per measurement time step. The distribution of brightness values in all considered daily measurements is shown in the following figure.

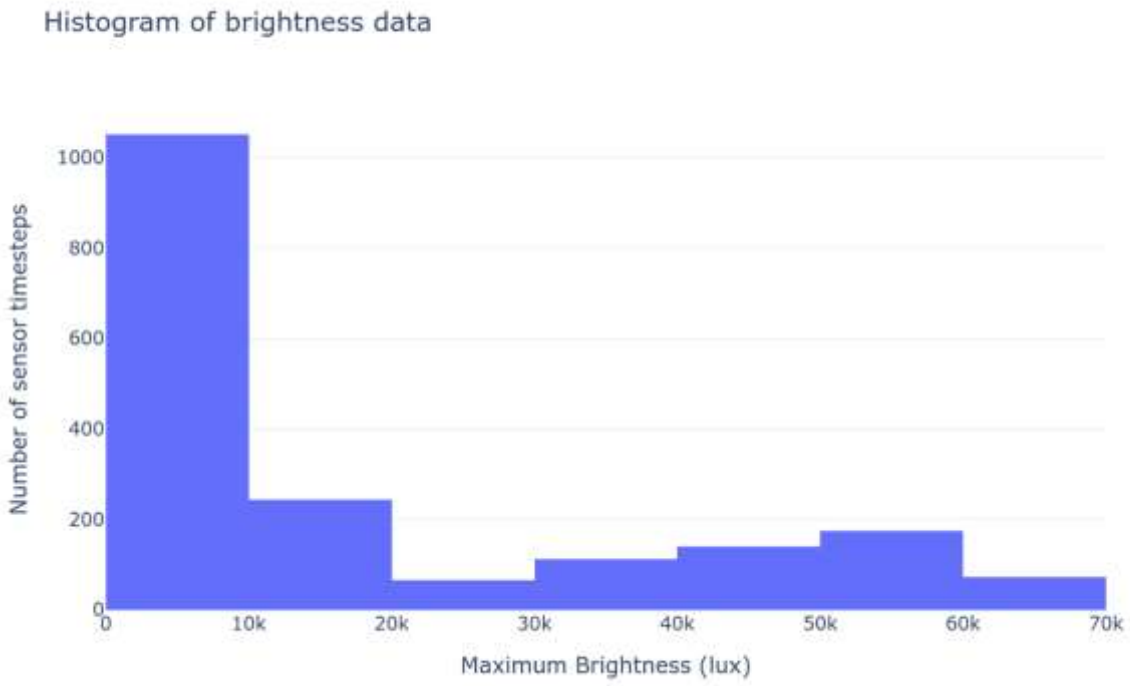

Figure 5: Distribution of brightness values in all considered daily camera measurements

The following images show some impressions of the camera image in different lighting conditions.



Figure 6: Camera image in sunlight



Figure 7: Camera image with bright sunlight and snow on the ground



Figure 8: Camera image with distinct clouds

4.1.1.1 Camera mean intensity in brightness

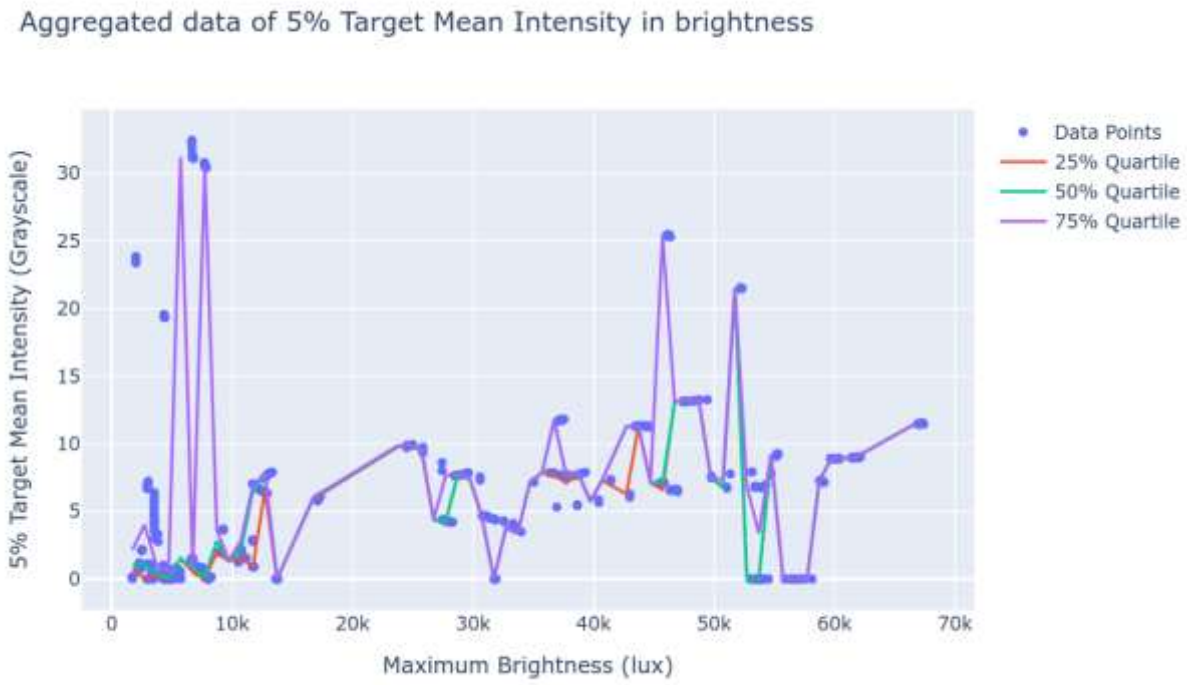

Figure 9: Aggregated data of 5% Target: Mean intensity in brightness

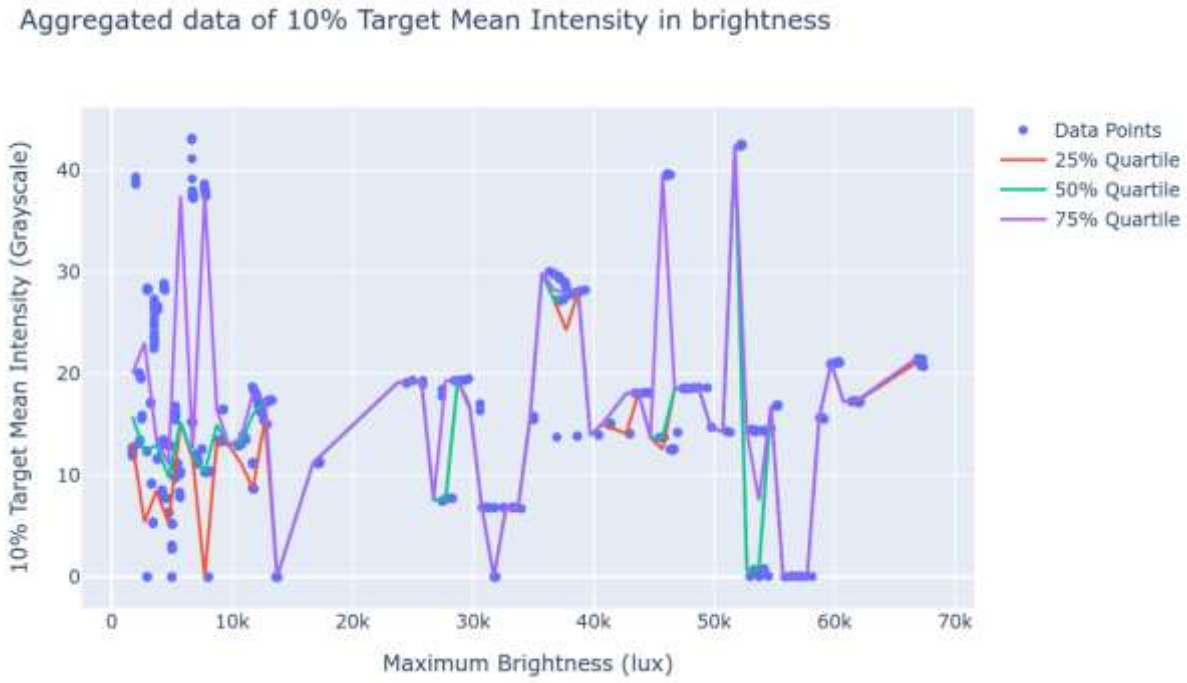

Figure 10: Aggregated data of 10% Target: Mean intensity in brightness



Figure 11: Aggregated data of 50% Target: Mean intensity in brightness



Figure 12: Aggregated data of 80% Target: Mean intensity in brightness



Figure 13: Aggregated data of 95% Target: Mean intensity in brightness

4.1.1.2 Camera contrasts in brightness

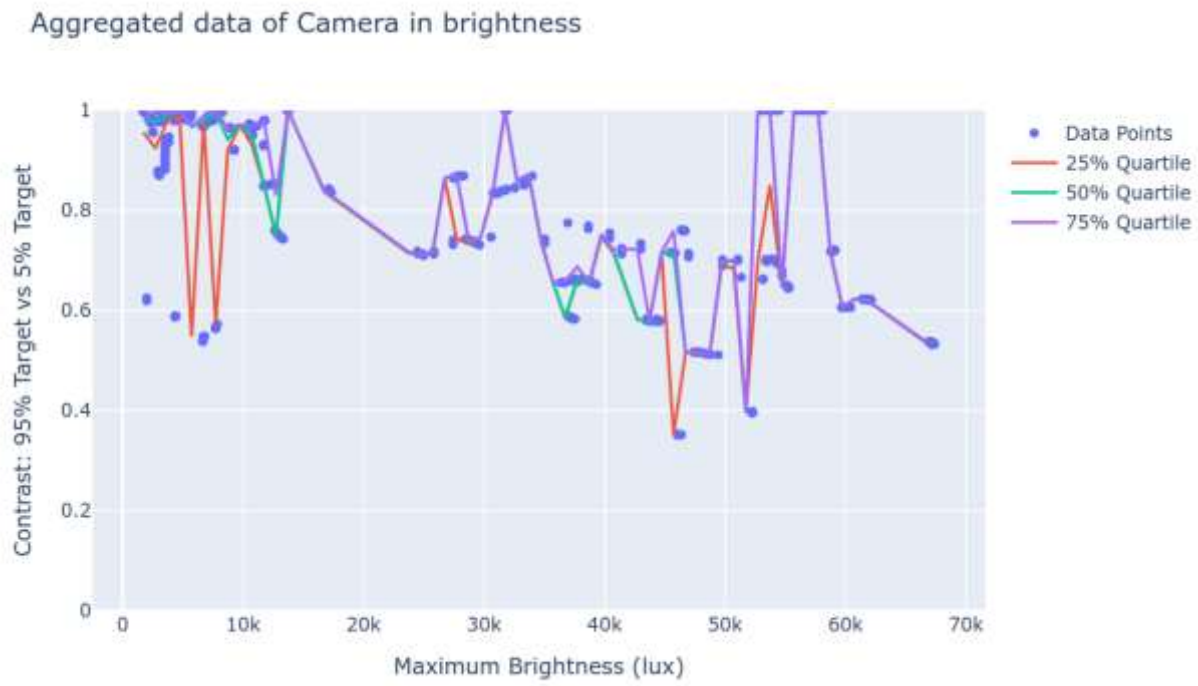

Figure 14: Aggregated contrast of camera in brightness: 95% target vs. 5% target

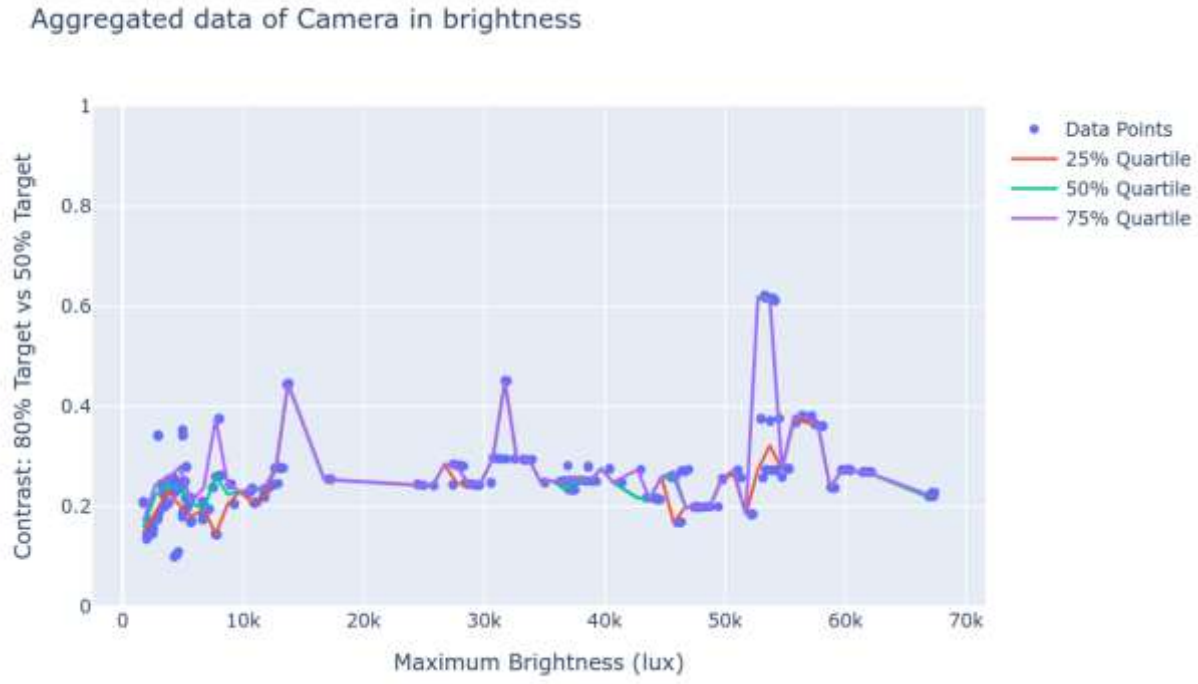

Figure 15: Aggregated contrast of camera in brightness: 80% target vs. 50% target

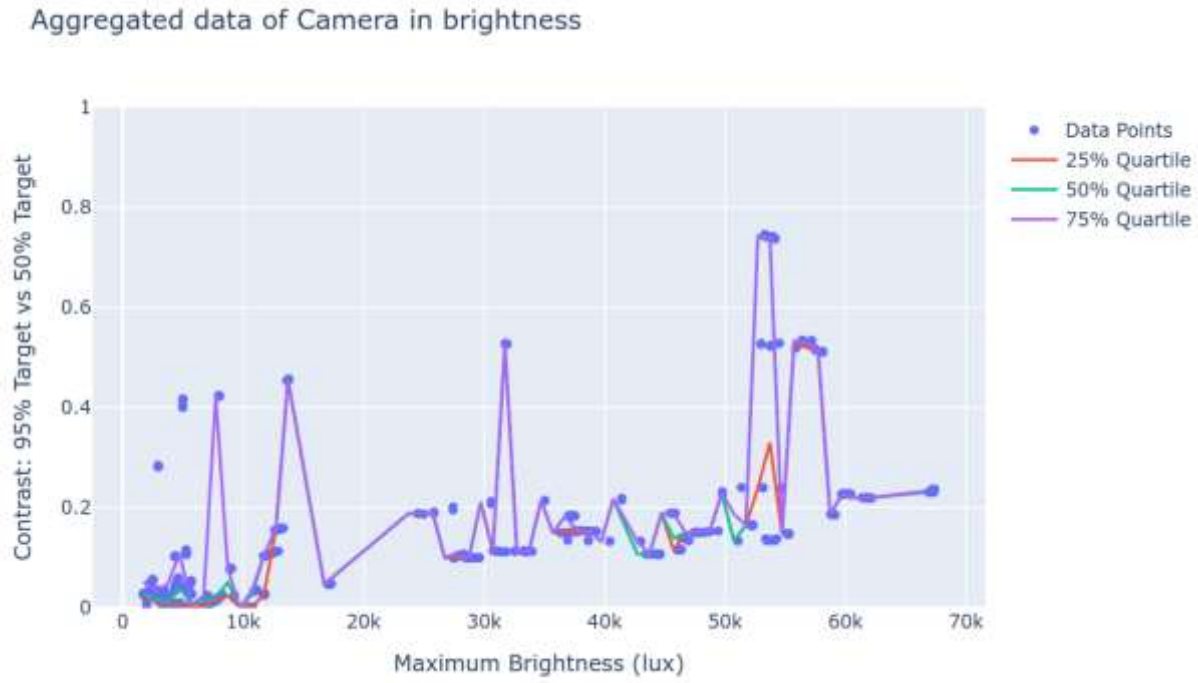

Figure 16: Aggregated contrast of camera in brightness: 95% target vs. 50% target

4.1.1.3 Interim conclusion for camera performance in brightness

Overall, the brightness study indicates that ambient illumination mainly affects contrasts that involve the very low-reflectivity reference (5%), while the contrasts between mid/high-reflectivity targets remain comparatively stable. The plots suggest that with increasing brightness the mean intensity measured on the dark targets (5% and, to a lesser extent, 10%) rises, whereas the mean intensity on the brighter targets (50%, 80%, 95%) appears to decrease. Because the contrast metric is computed as the difference of the means divided by the sum of the means, an increased response on the dark target reduces the numerator and increases the denominator, which naturally lowers the contrast for pairs such as 95% vs. 5% even when the bright target itself does not change significantly.

A plausible explanation is that the camera image pipeline adapts to the global illumination level (e.g., auto-exposure and/or tone mapping). Under high brightness, darker target regions may be lifted more strongly (higher digital counts) to preserve visibility, while bright regions approach saturation or are compressed by the processing chain, reducing their dynamic range. In addition, stray light, lens flare, and increased veiling glare at high illumination can add an approximately constant offset to darker regions, again reducing relative contrast. For the interpretation and later use of these results, this implies that "brightness" should be treated as a confounding factor when deriving performance curves for appearance-based KPIs: contrasts against a very dark reference are not purely material-dependent but also depend on camera adaptation. If these curves are used to parameterize or validate camera simulation, it is therefore recommended to document the camera settings (especially exposure/gain behavior if available) and to avoid extrapolating beyond the measured lux range, as different exposure regimes can change the trend qualitatively.

4.1.2 CAMERA PERFORMANCE IN FOG

Around 4675 measurement time steps were taken during foggy conditions. Fog is quantified in either the Meteorological Optical Range (MOR) or the Forward Scattering Extinction Coefficient (EXCO). MOR in km is defined as 3/EXCO. In the following plots the EXCO will be used, as it has a finer resolution in the visibility sensor data. The distribution of extinction coefficients in all fog measurements is shown in the following figure.

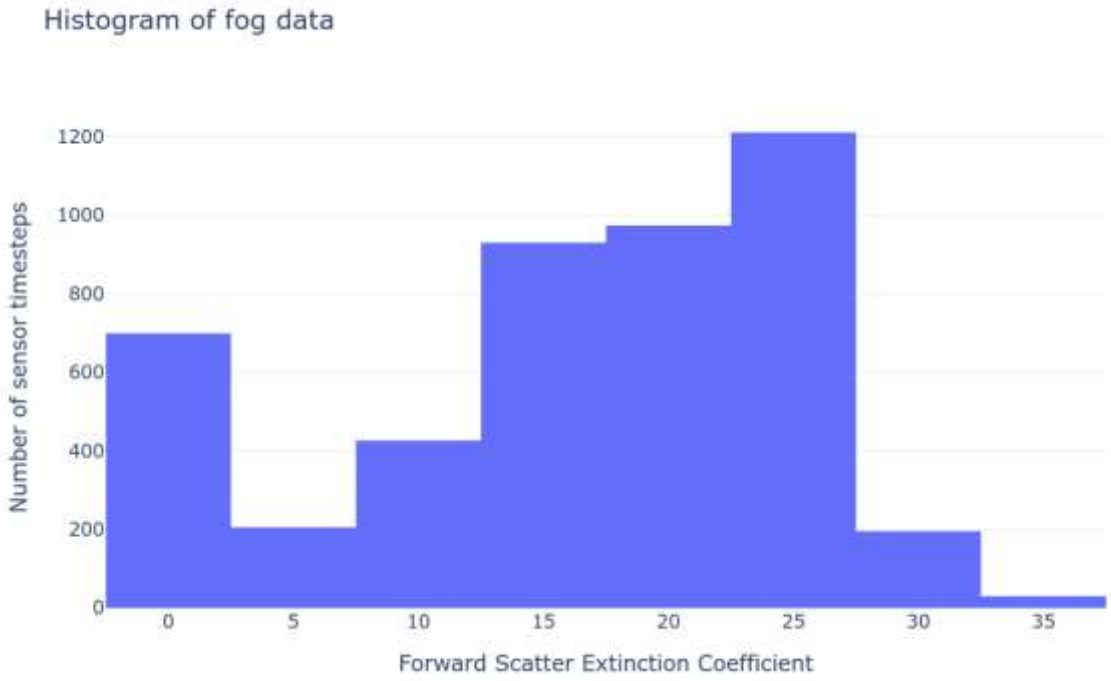

Figure 17: Histogram of fog data

The following images show some impressions of the camera image in light, medium and heavy fog.



Figure 18: Camera image in light fog



Figure 19: Camera image in medium fog

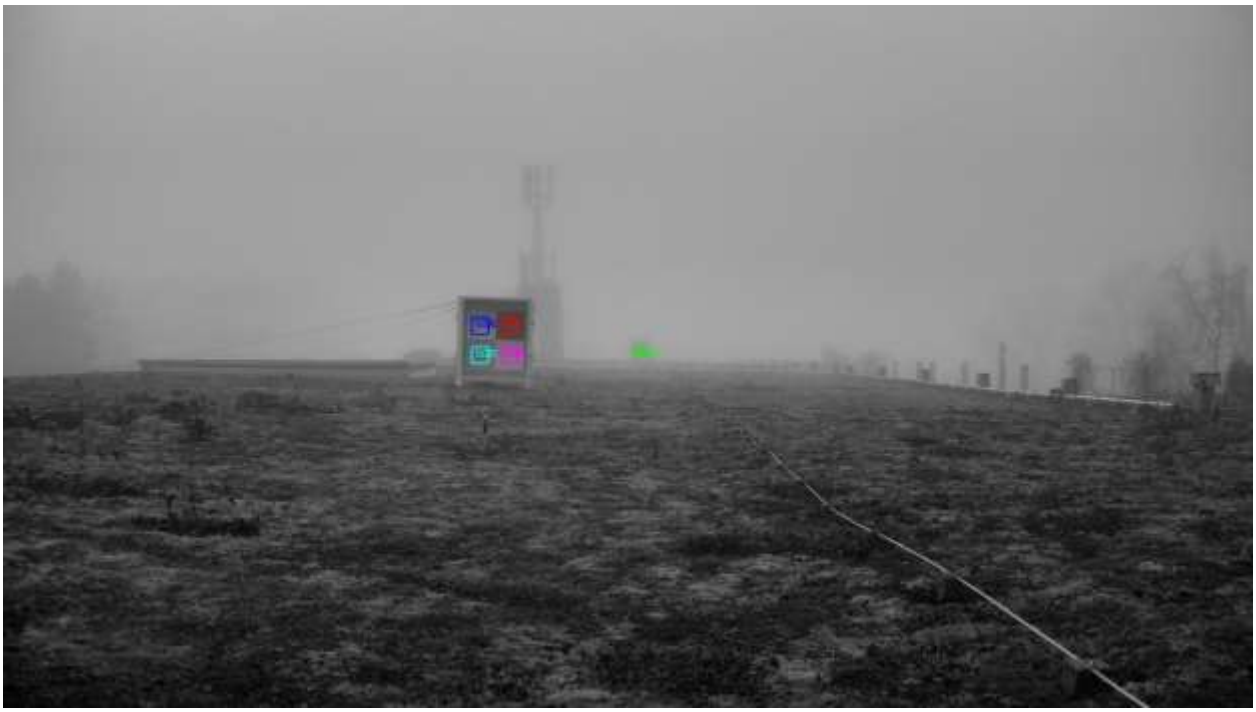

Figure 20: Camera image in heavy fog

4.1.2.1 Camera mean intensity in fog



Figure 21: Aggregated data of 5% Target: Mean intensity in fog



Figure 22: Aggregated data of 10% Target: Mean intensity in fog



Figure 23: Aggregated data of 50% Target: Mean intensity in fog



Figure 24: Aggregated data of 80% Target: Mean intensity in fog



Figure 25: Aggregated data of 95% Target: Mean intensity in fog

4.1.2.2 Camera contrasts in fog



Figure 26: Aggregated contrast of camera in fog: 95% target vs. 5% target

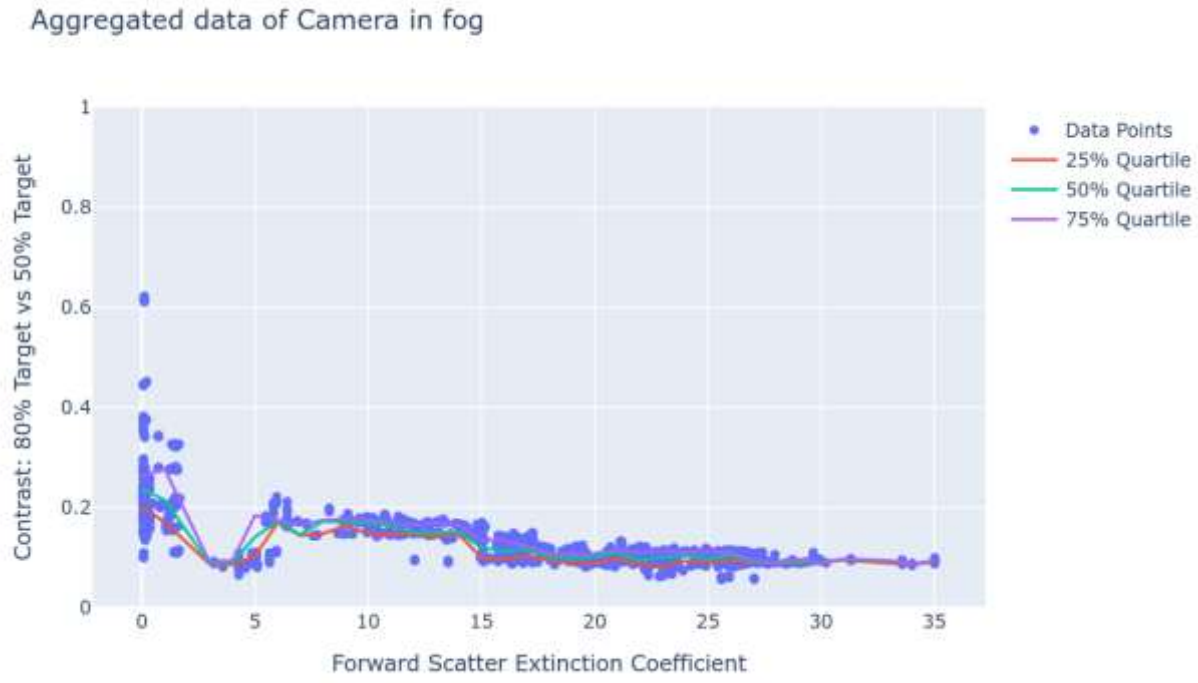

Figure 27: Aggregated contrast of camera in fog: 80% target vs. 50% target



Figure 28: Aggregated contrast of camera in fog: 95% target vs. 50% target

4.1.2.3 Interim conclusion for camera performance in fog

Overall, the fog measurements show a consistent degradation of image contrast with increasing forward scatter extinction coefficient across all target pairs. The contrast drop starts already at low extinction values and becomes progressively stronger as fog density increases, which indicates that fog acts primarily as a global visibility limiter rather than affecting only specific reflectance levels. In parallel, the mean intensity measured on all targets increases with higher extinction. This trend is consistent with the visual impression in the example images: fog introduces a bright veiling component ("airlight") that adds an approximately uniform radiance offset to the scene, making dark targets appear brighter and compressing the dynamic range between dark and bright surfaces.

From a physical perspective, the simultaneous increase in measured intensity and decrease in contrast can be explained by the superposition of (i) attenuation of the target radiance along the line of sight and (ii) additional path radiance generated by forward scattering in the fog layer. As the extinction increases, the contribution of airlight grows and tends to equalize pixel values across the image, which directly reduces relative contrast for all target combinations (including mid-level pairs such as 80% vs. 50%). Depending on the camera's internal processing, auto-exposure and white balance may further amplify the effect by stabilizing overall brightness while effectively lifting low-intensity regions. For later use of these results, this implies that fog should be treated as a dominant confounder for appearance-based KPIs: contrast is expected to decay nonlinearly with extinction and may show different slopes once exposure regimes change. If the curves are used for camera simulation or validation, it is recommended to parameterize fog severity via the measured extinction coefficient (or MOR) and to model the fog effect explicitly as an added veiling luminance term plus attenuation, rather than attributing the behavior to target reflectance changes. Finally, extrapolation beyond the observed extinction range should be avoided, as very dense fog can drive the image into a near-uniform regime where small processing differences (tone mapping, clipping) can dominate the measured KPIs.

4.1.3 CAMERA PERFORMANCE IN HAIL

Around 1049 measurement time steps were taken in hail. Hail intensity is quantified with the precipitation intensity in mm/h. The distribution of precipitation intensity in all hail measurements is shown in the following figure.

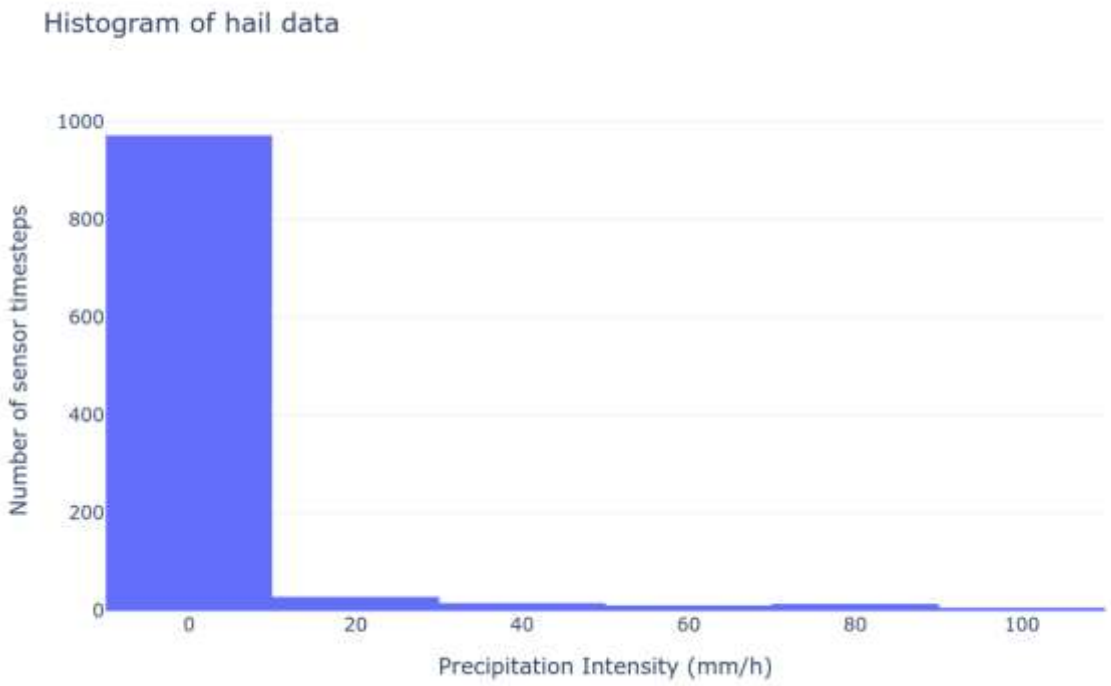

Figure 29: Histogram of hail data for camera evaluation

4.1.3.1 Camera mean intensity in hail

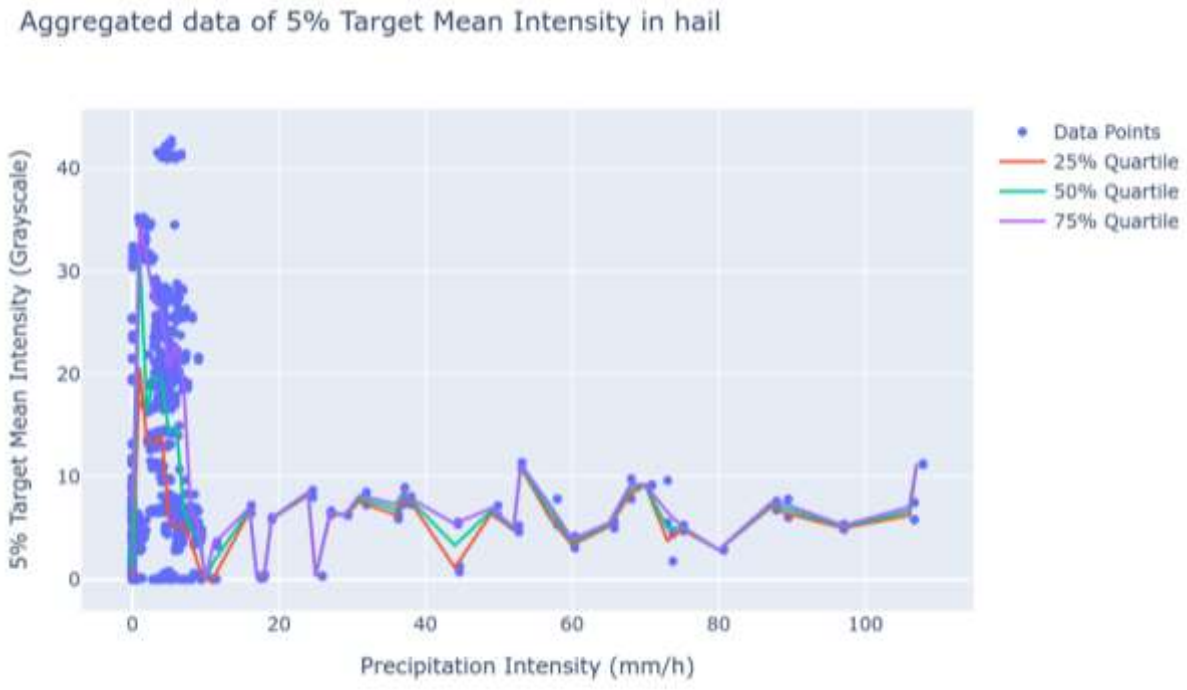

Figure 30: Aggregated data of 5% Target: Mean intensity in hail

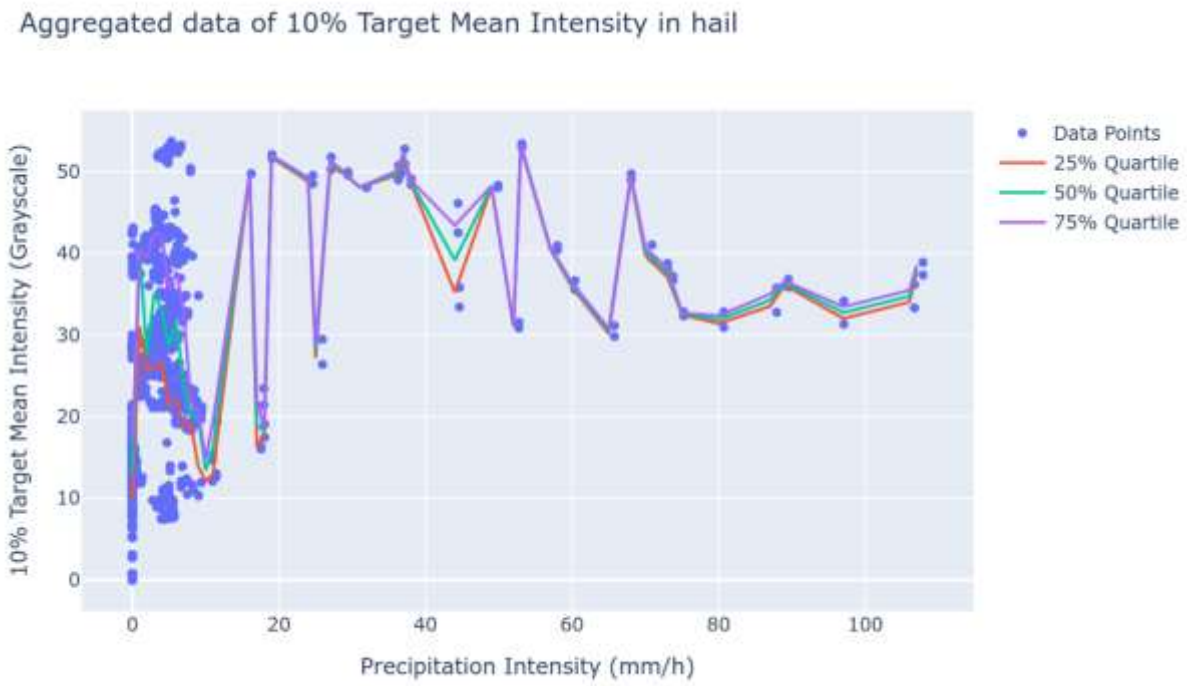

Figure 31: Aggregated data of 10% Target: Mean intensity in hail

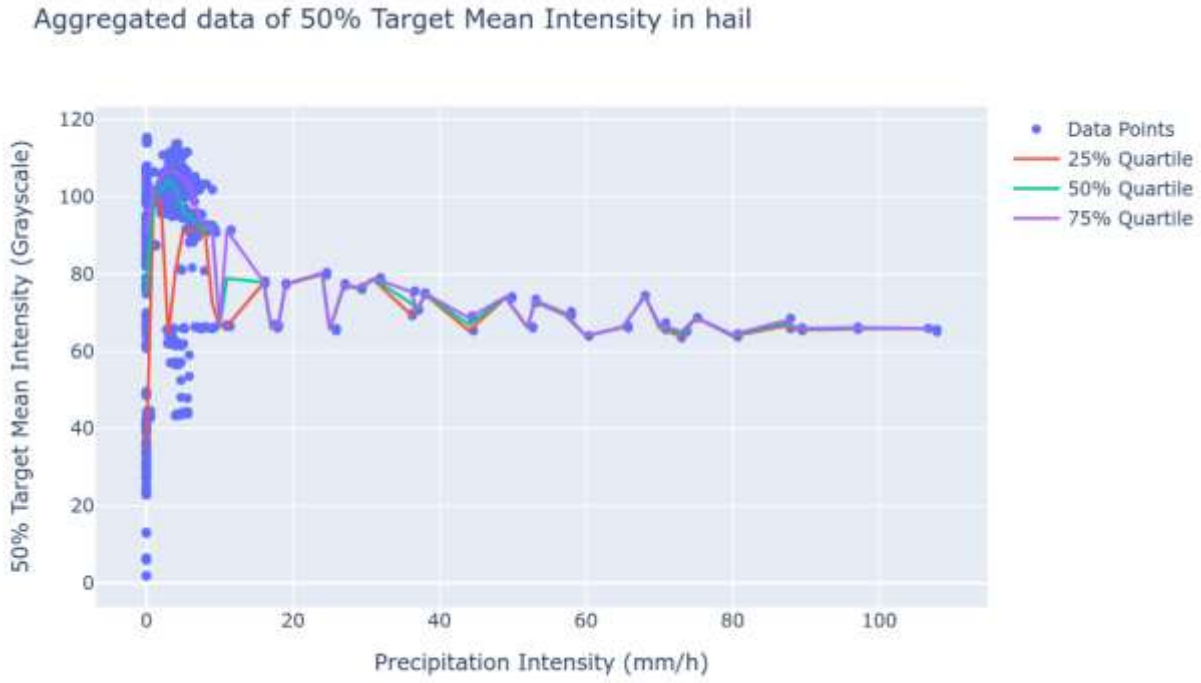

Figure 32: Aggregated data of 50% Target: Mean intensity in hail

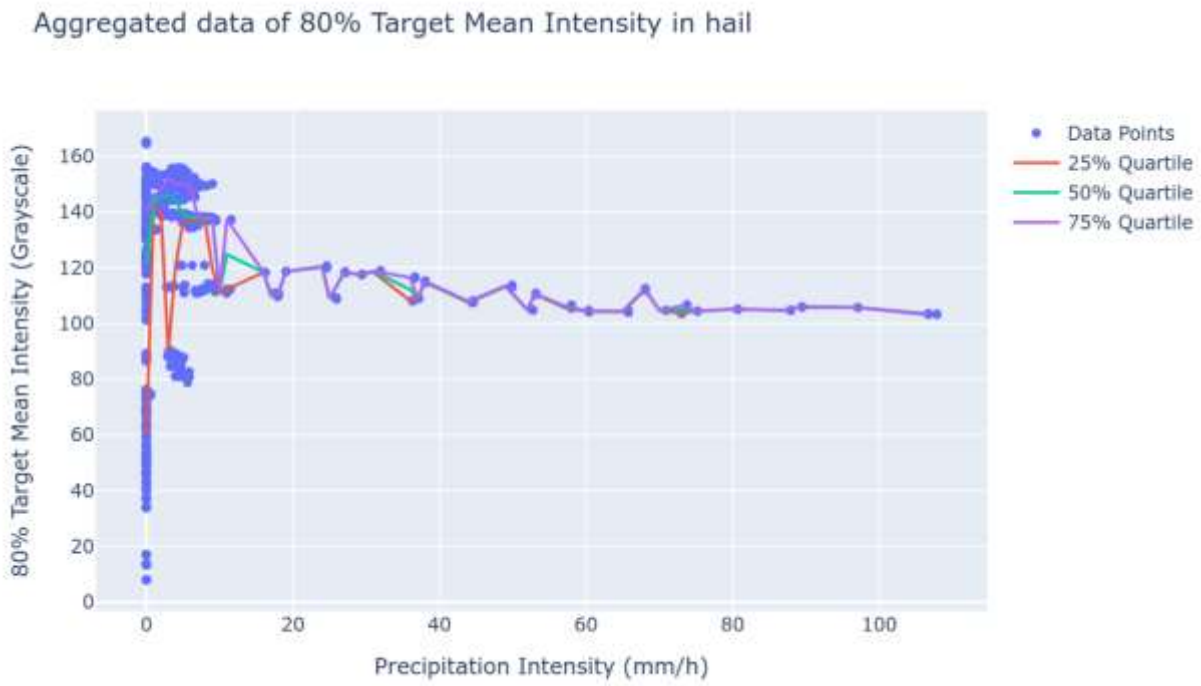

Figure 33: Aggregated data of 80% Target: Mean intensity in hail

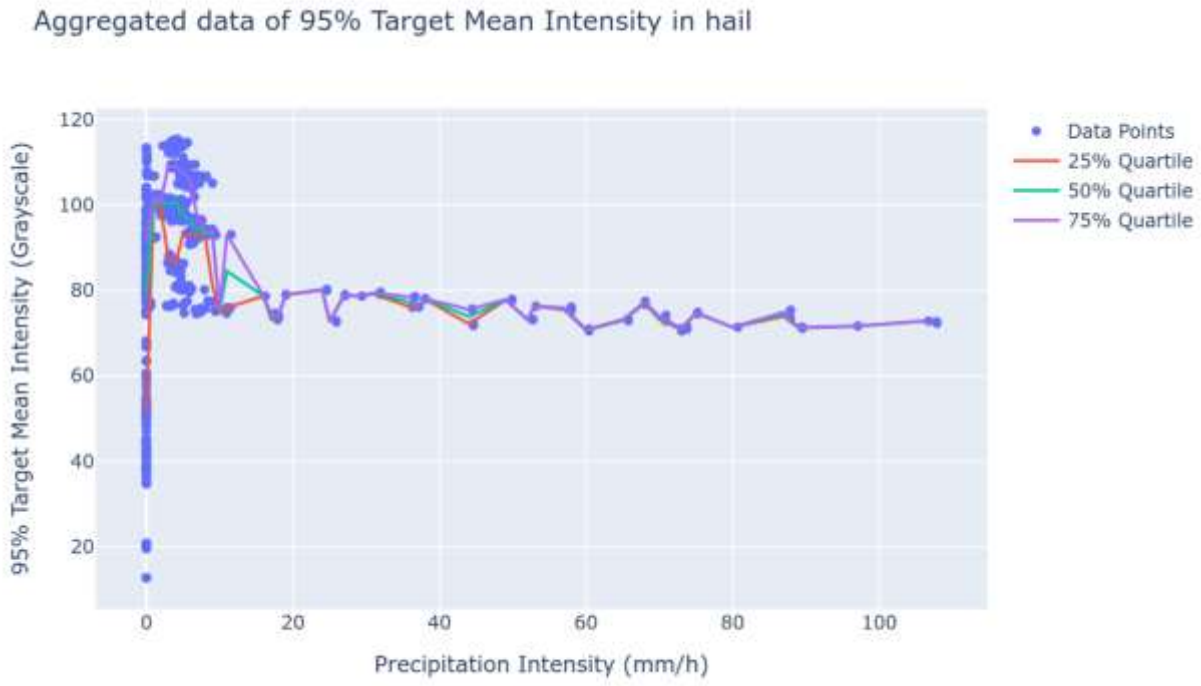

Figure 34: Aggregated data of 95% Target: Mean intensity in hail

4.1.3.2 Camera contrasts in hail

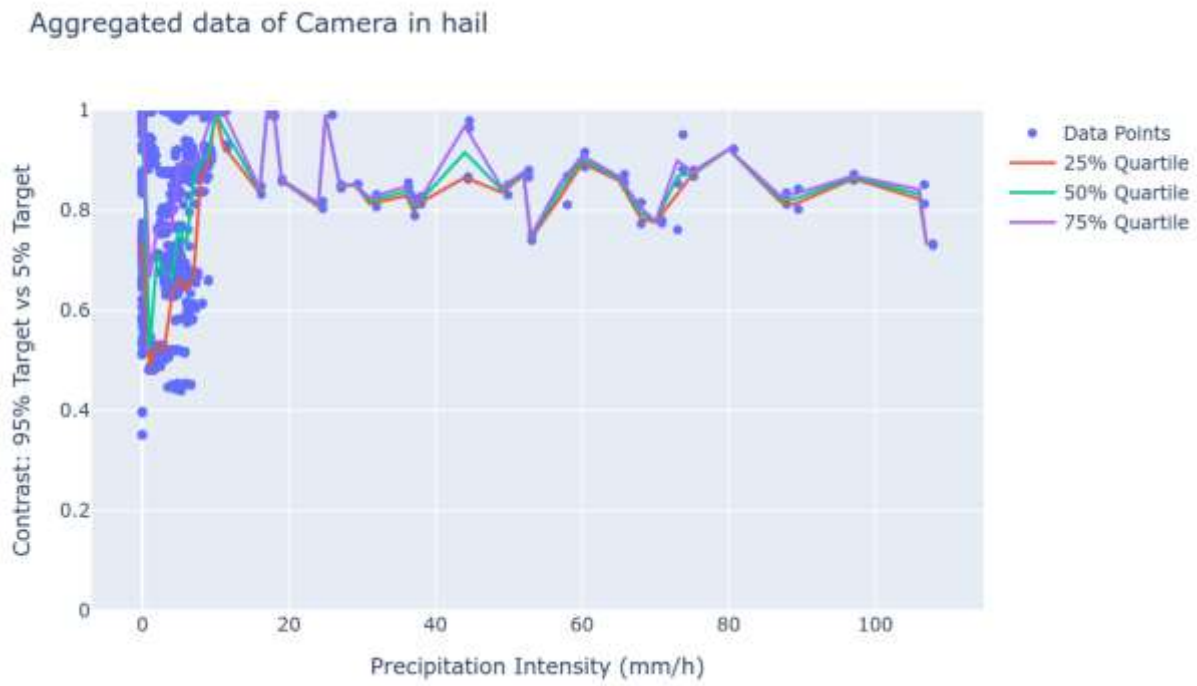

Figure 35: Aggregated contrast of camera in hail: 95% target vs. 5% target

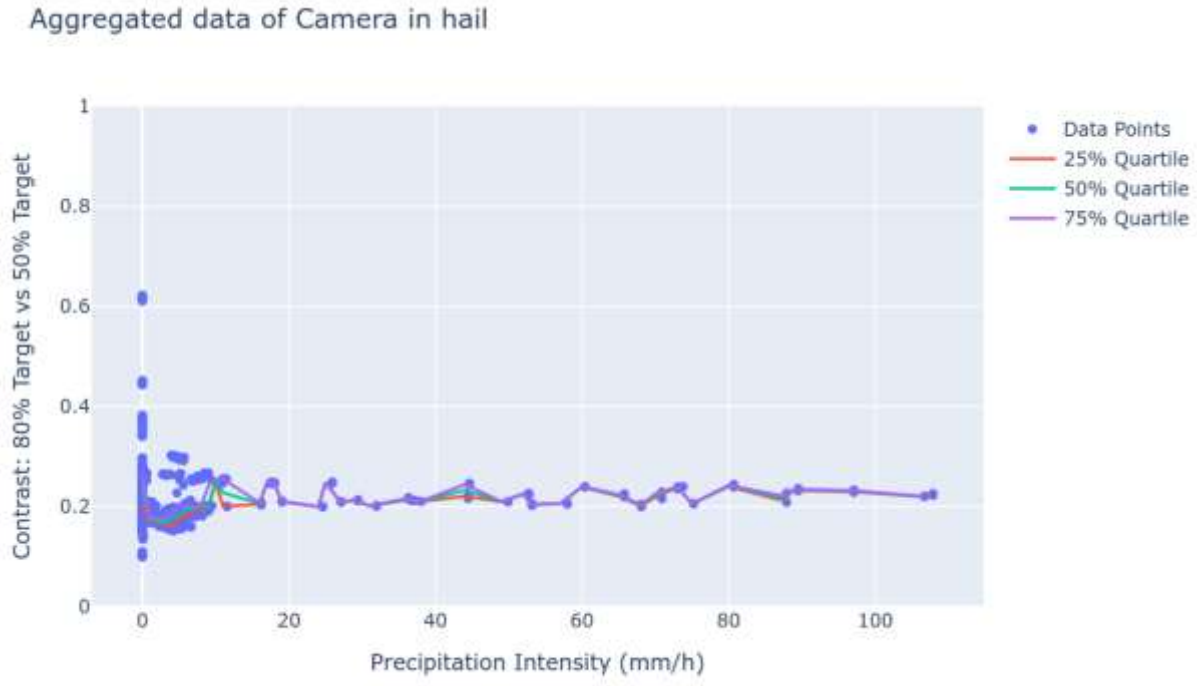

Figure 36: Aggregated contrast of camera in hail: 80% target vs. 50% target

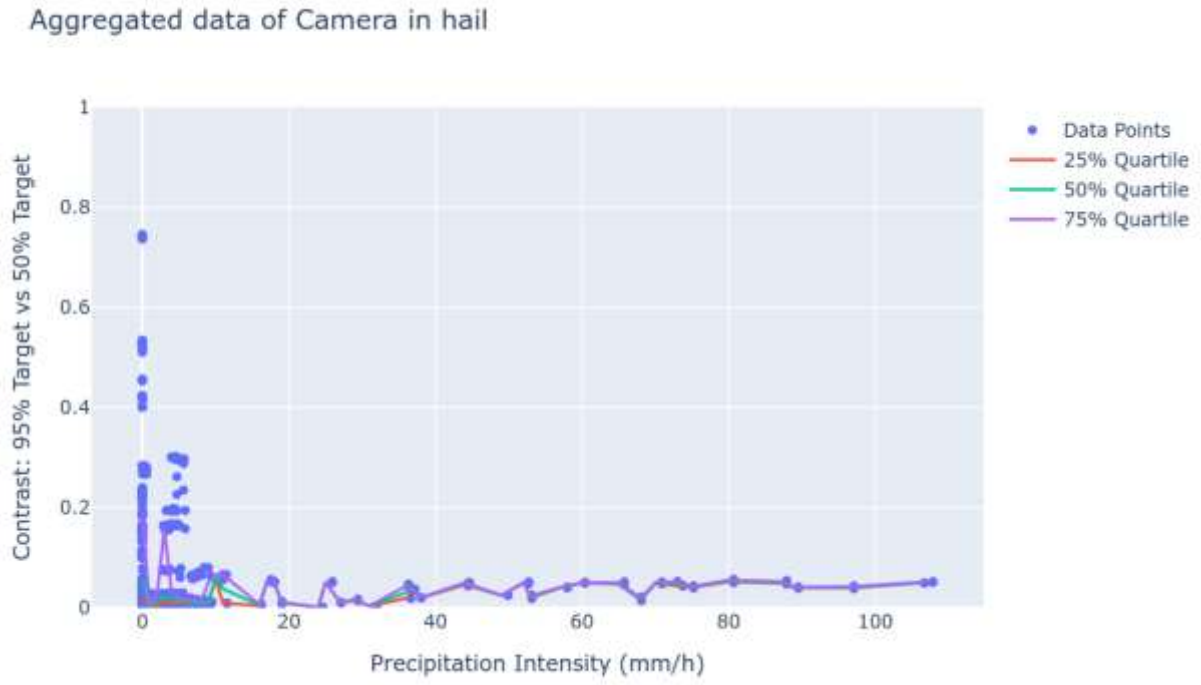

Figure 37: Aggregated contrast of camera in hail: 95% target vs. 50% target

4.1.3.3 Interim conclusion for camera performance in hail

Based on the available hail recordings, no systematic dependence of the evaluated camera KPIs on hail precipitation intensity can be identified. Both the mean intensity measured on the target regions and the derived contrasts between targets remain approximately constant over the covered intensity range. This indicates that, in the current evaluation setup, hail does not act as a persistent visibility limiter in the same way as fog and therefore does not lead to a sustained reduction of target-to-target contrast.

A plausible explanation is that hail primarily introduces short-lived, spatially sparse occlusions (individual hailstones crossing the line of sight) and localized specular highlights rather than a homogeneous veiling component. When statistics are computed on small, fixed target patches and then aggregated temporally, these transient effects are largely averaged out unless hail becomes dense enough to cause frequent full-pixel occlusions or strong motion blur across the target area. In addition, the precipitation intensity used here is an indirect proxy for the actual optical disturbance in the camera's field of view; variations in hailstone size distribution, wind-driven trajectories, and concurrent illumination conditions can dominate the appearance even at similar mm/h values. Importantly, the dataset contains only few samples above approximately 10 mm/h, so the conclusion should be interpreted as valid only for the measured (mostly light-to-moderate) hail regime. For simulation and validation, this suggests that, within the covered range, hail can be treated as a negligible factor for mean-intensity/contrast KPIs on static Lambertian targets, while potential impacts are more likely to appear in other metrics (e.g., temporal artifacts, streaks, or algorithm-level failures). Additional measurements at higher hail intensities and under controlled exposure settings would be required to quantify a possible threshold behavior where hail begins to measurably degrade these aggregated image statistics.

4.1.4 CAMERA PERFORMANCE IN RAIN

Around 3268 measurement time steps were taken during rainy conditions. Rain is quantified in rain intensity (mm/h). The distribution of rain intensities in all rain measurements is shown in the following figure.

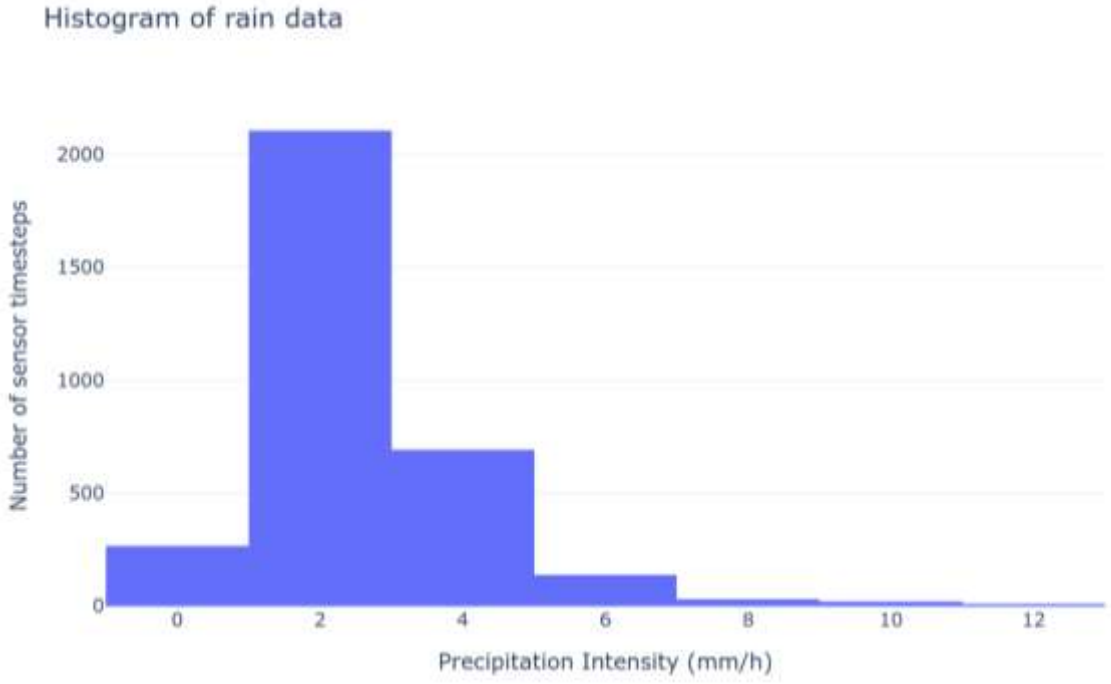

Figure 38: Histogram of rain data for camera evaluation

4.1.4.1 Camera mean intensity in rain

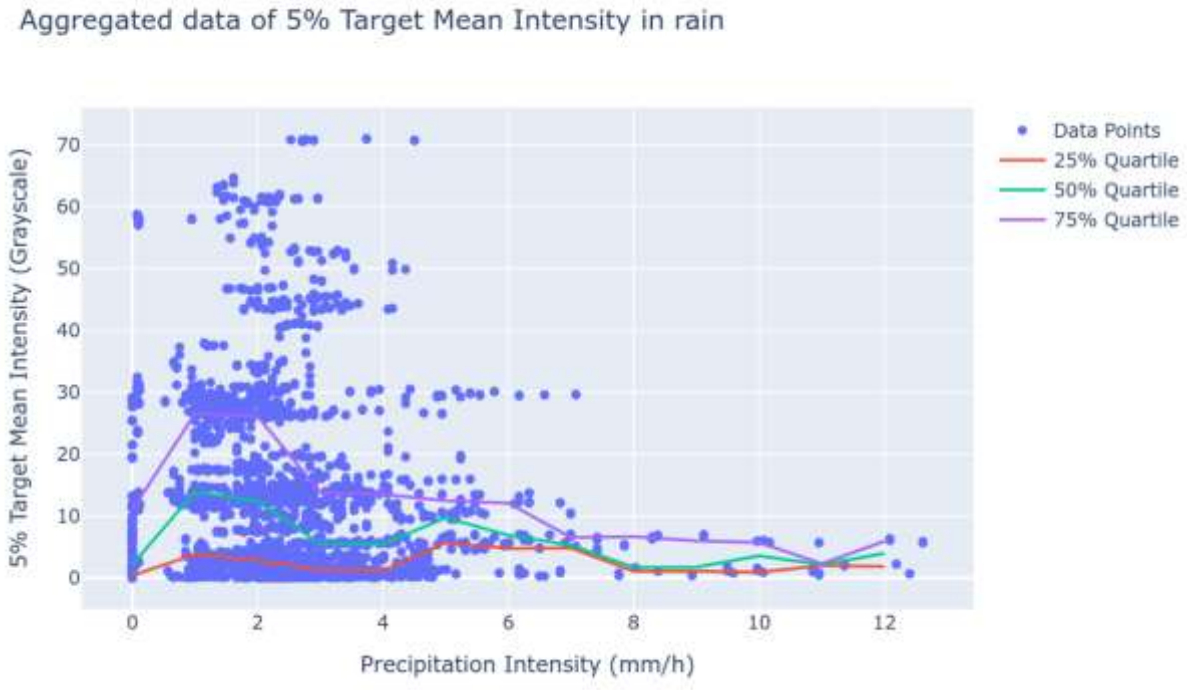

Figure 39: Aggregated data of 5% Target: Mean intensity in rain

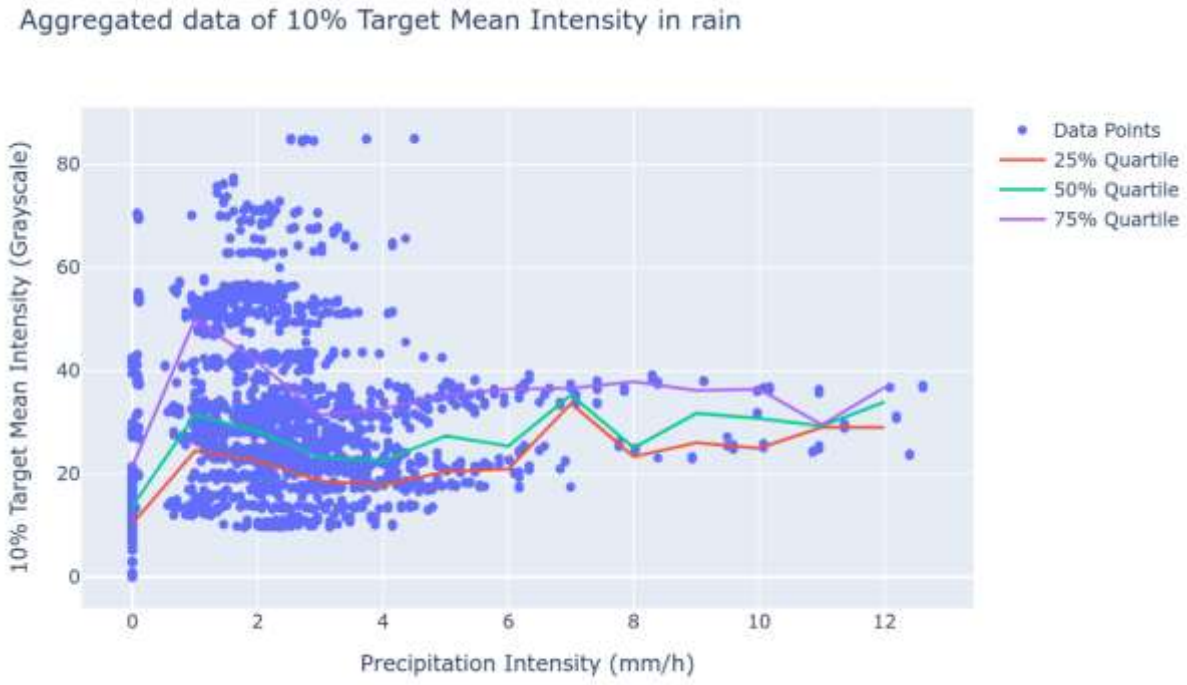

Figure 40: Aggregated data of 10% Target: Mean intensity in rain

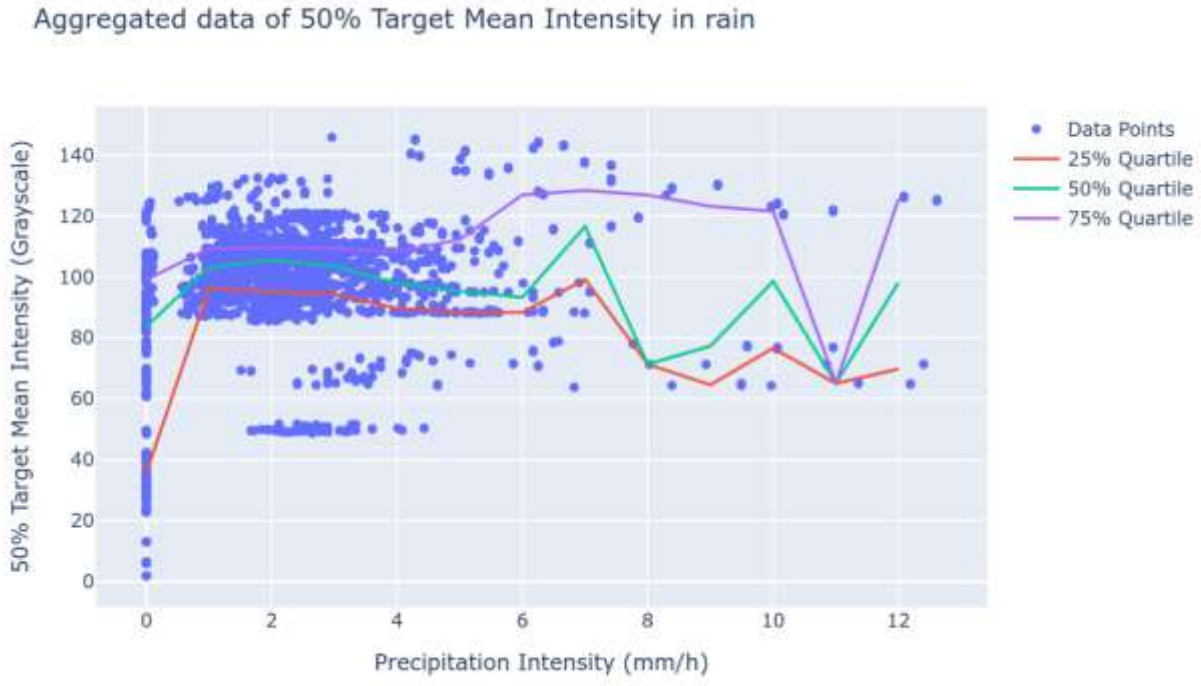

Figure 41: Aggregated data of 50% Target: Mean intensity in rain

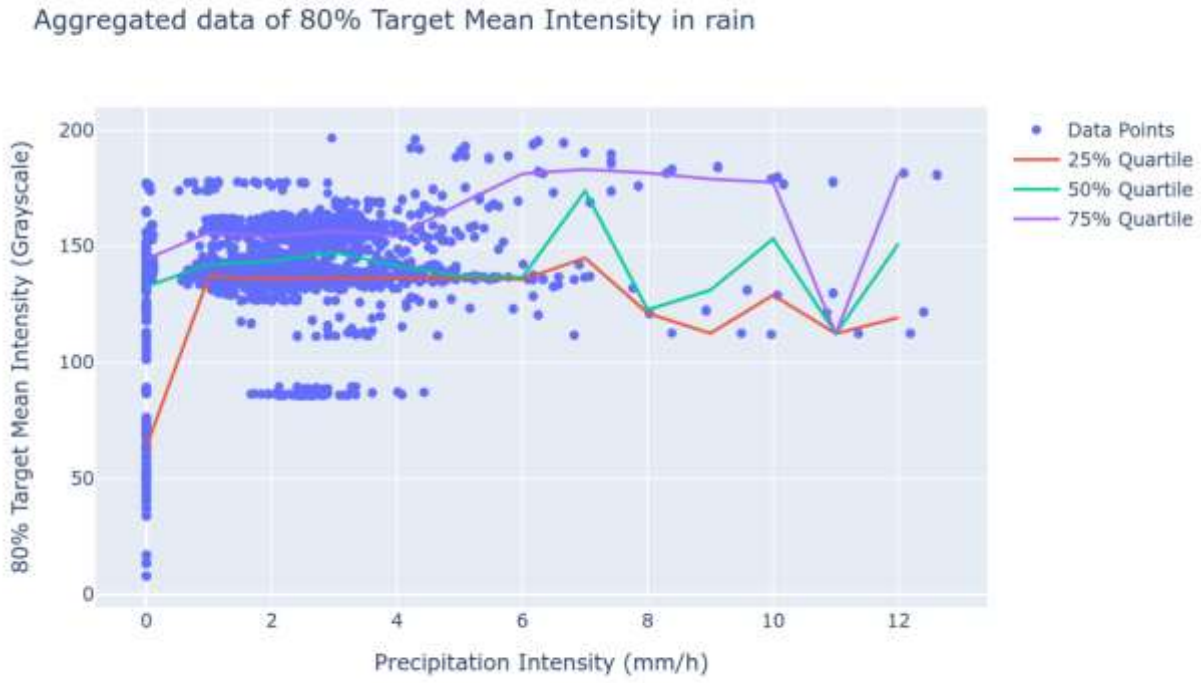

Figure 42: Aggregated data of 80% Target: Mean intensity in rain

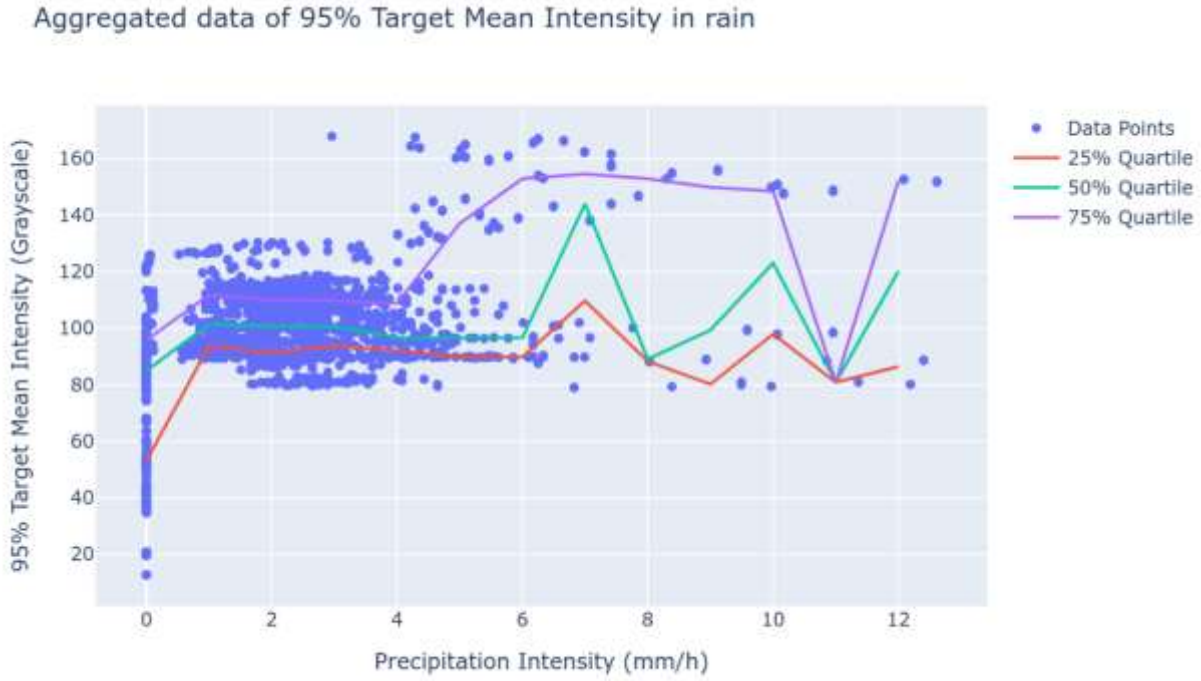

Figure 43: Aggregated data of 95% Target: Mean intensity in rain

4.1.4.2 Camera contrasts in rain

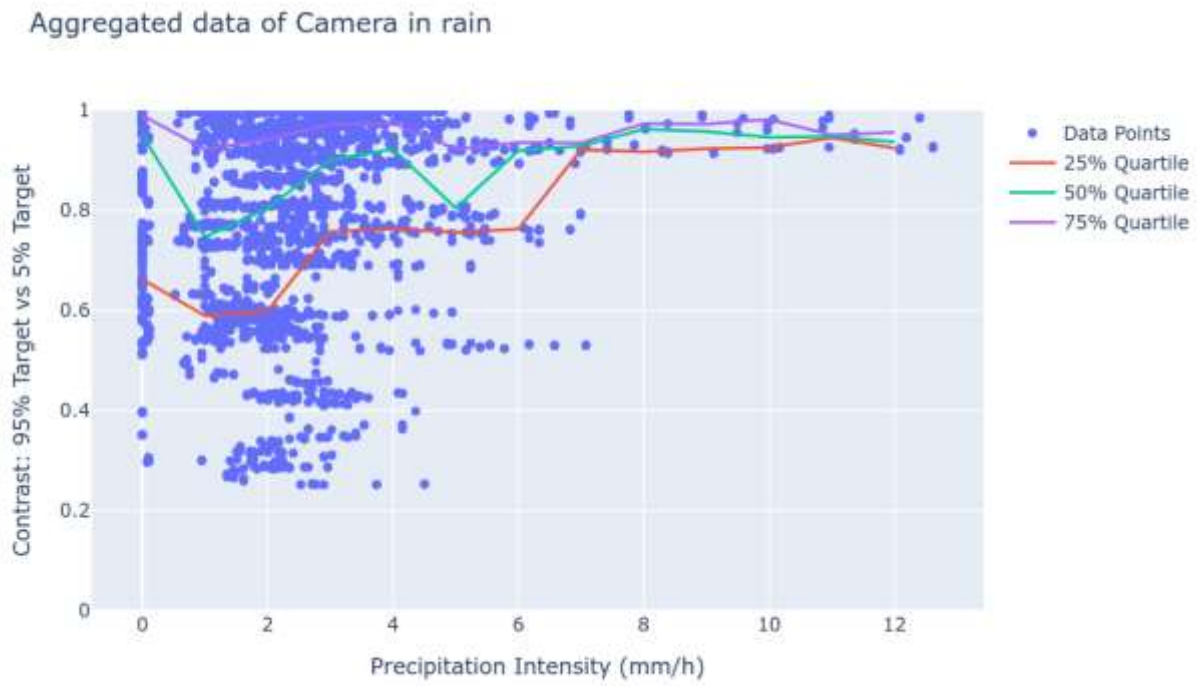

Figure 44: Aggregated contrast of camera in rain: 95% target vs. 5% target

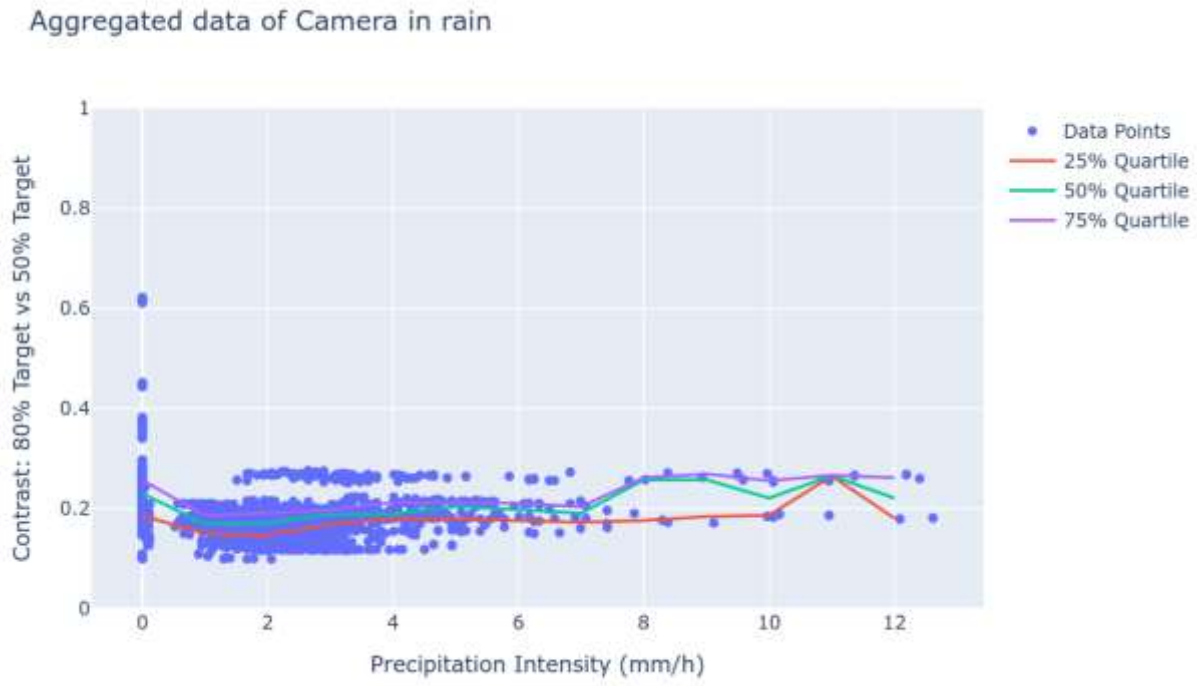

Figure 45: Aggregated contrast of camera in rain: 80% target vs. 50% target

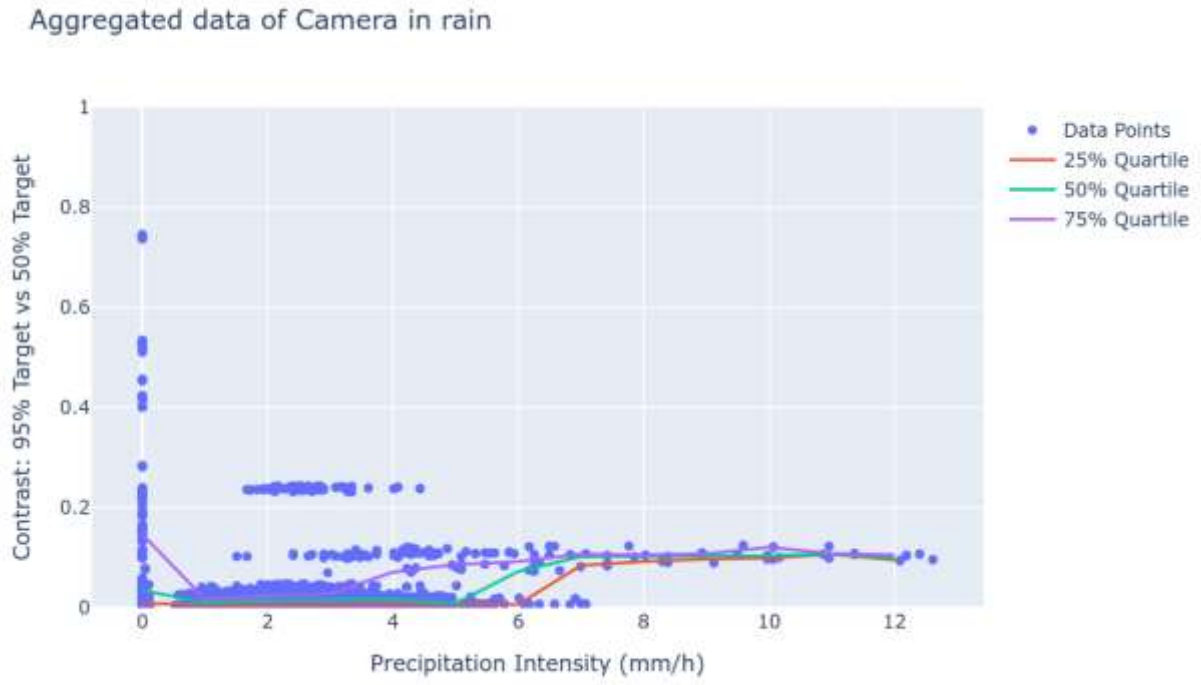

Figure 46: Aggregated contrast of camera in rain: 95% target vs. 50% target

4.1.4.3 Interim conclusion for camera performance in rain

Overall, the rain measurements suggest that rain primarily introduces moderate, highly variable perturbations in the evaluated image KPIs rather than a clear monotonic degradation over rain intensity. Across the covered intensity range, the mean intensities on the target ROIs remain broadly within a comparable band, with an increased scatter at higher rain rates that is consistent with intermittent raindrop occlusions and short-lived brightness fluctuations. The contrast curves show the same qualitative behavior: contrasts against the dark reference (e.g., 95% vs. 5%) tend to be more sensitive and exhibit a larger spread, while mid-level contrasts (e.g., 80% vs. 50%) remain comparatively stable on average. This indicates that, in contrast to fog, rain does not behave like a uniform global visibility reducer in these measurements, but rather affects the image in a temporally and spatially inhomogeneous way.

A plausible explanation is that rainfall impacts the camera image mainly through transient effects: individual drops or streaks can partially cover the target ROIs, produce local specular highlights, and add small veiling components due to scattering of ambient light. In addition, wet surfaces on the target housing window or surrounding structures can increase glare and local reflections, which may bias the intensity of darker patches and thereby reduce contrast intermittently. Because the KPIs are computed from fixed ROIs and aggregated over three consecutive frames, many of these short-lived artifacts are averaged out unless the rain becomes dense enough to cause frequent occlusions within the ROI.

It should also be noted that the precipitation intensity (mm/h) is an indirect proxy for optical disturbance: drop size distribution, wind-driven trajectories, and illumination conditions can lead to different image impacts at the same mm/h. For simulation and validation, this suggests that rain should be modeled less as a deterministic contrast-attenuation curve (as in fog) and

more as a stochastic disturbance process that increases the variance of intensity/contrast and occasionally induces local outliers (drop hits, streaks). Where possible, documenting camera exposure behavior and restricting conclusions to the measured rain-rate range will improve traceability, and additional recordings at higher intensities would be needed to assess whether a threshold exists where rain begins to produce a sustained contrast loss similar to a veiling effect.

4.1.5 CAMERA PERFORMANCE IN SNOW

Around 5137 measurement time steps were taken in Snow. The snowfall is quantified with the precipitation intensity in mm/h. The distribution of precipitation intensity in all snow measurements is shown in the following figure.

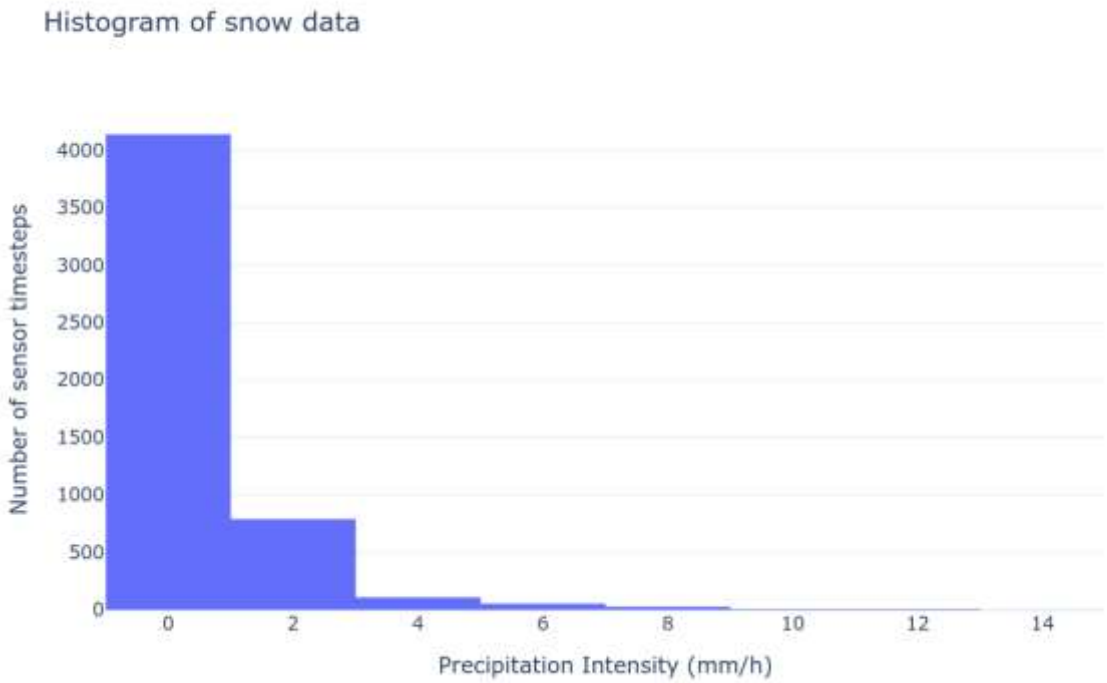

Figure 47: Histogram of snow data for camera evaluation

The following image shows an impression of the camera image during snowfall.



Figure 48: Camera image during snow fall

4.1.5.1 Camera mean intensity in snow

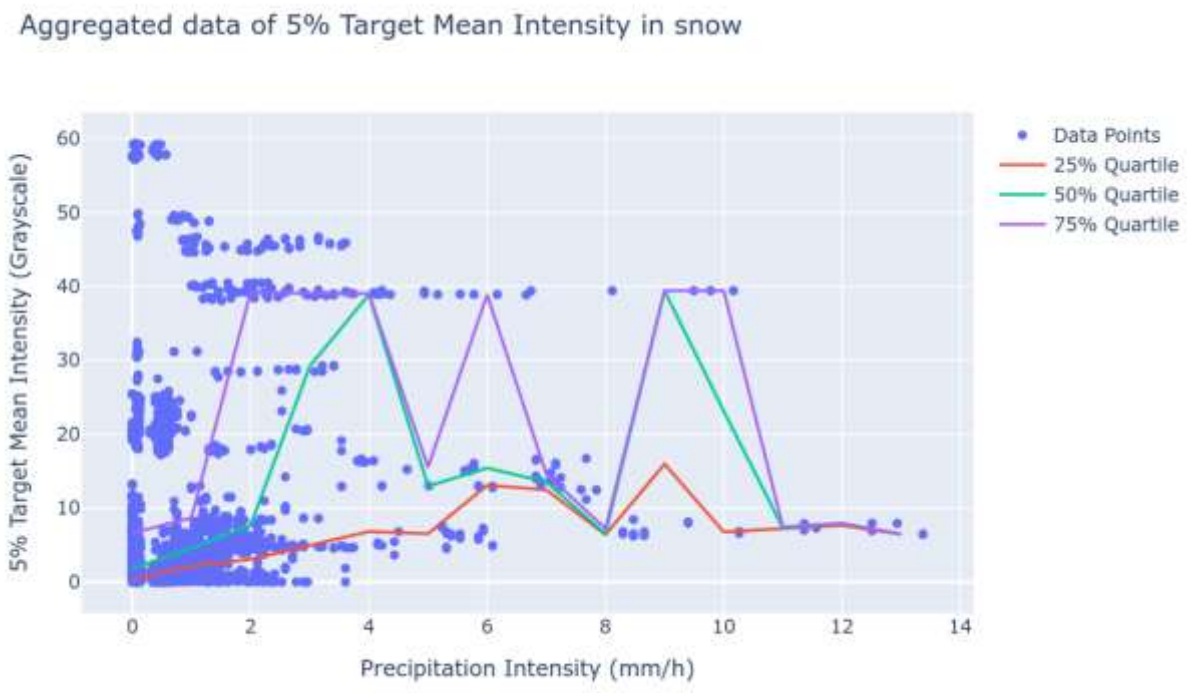

Figure 49: Aggregated data of 5% Target: Mean intensity in snow

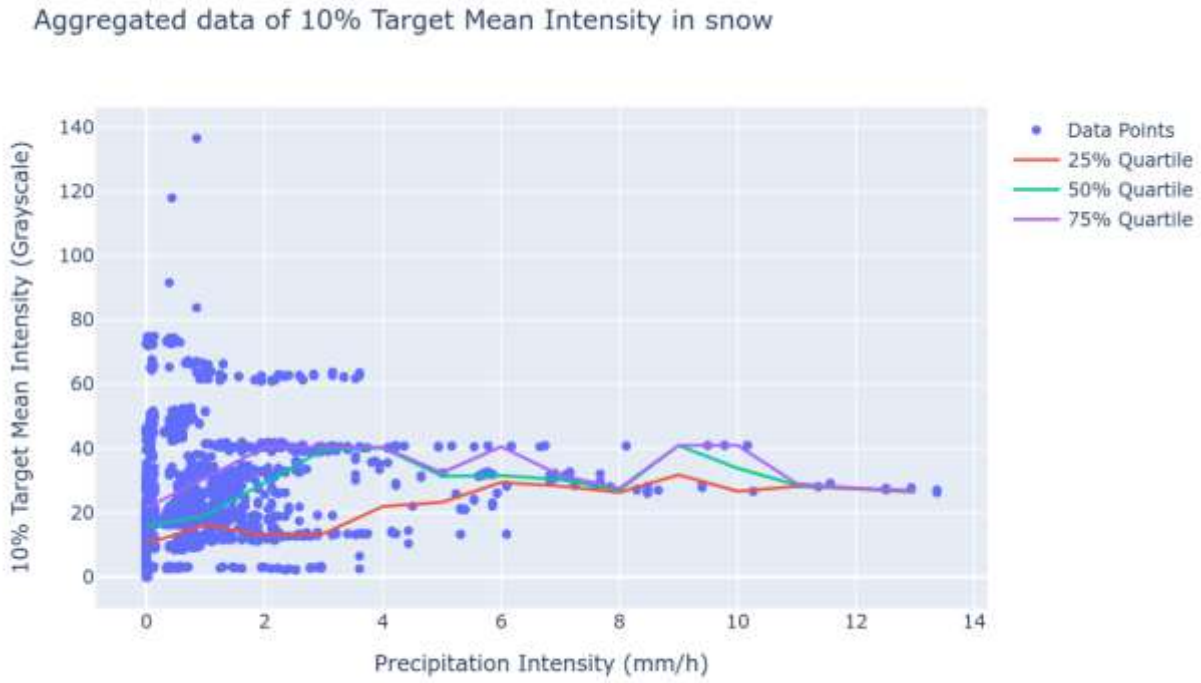

Figure 50: Aggregated data of 10% Target: Mean intensity in snow

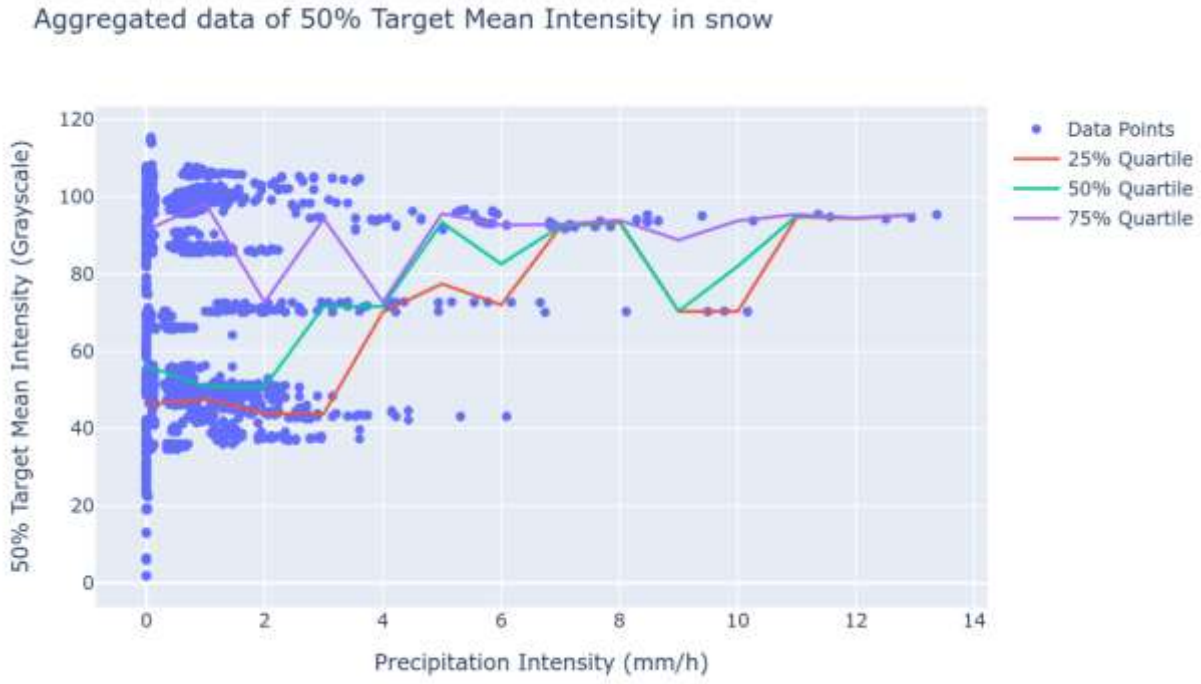

Figure 51: Aggregated data of 50% Target: Mean intensity in snow

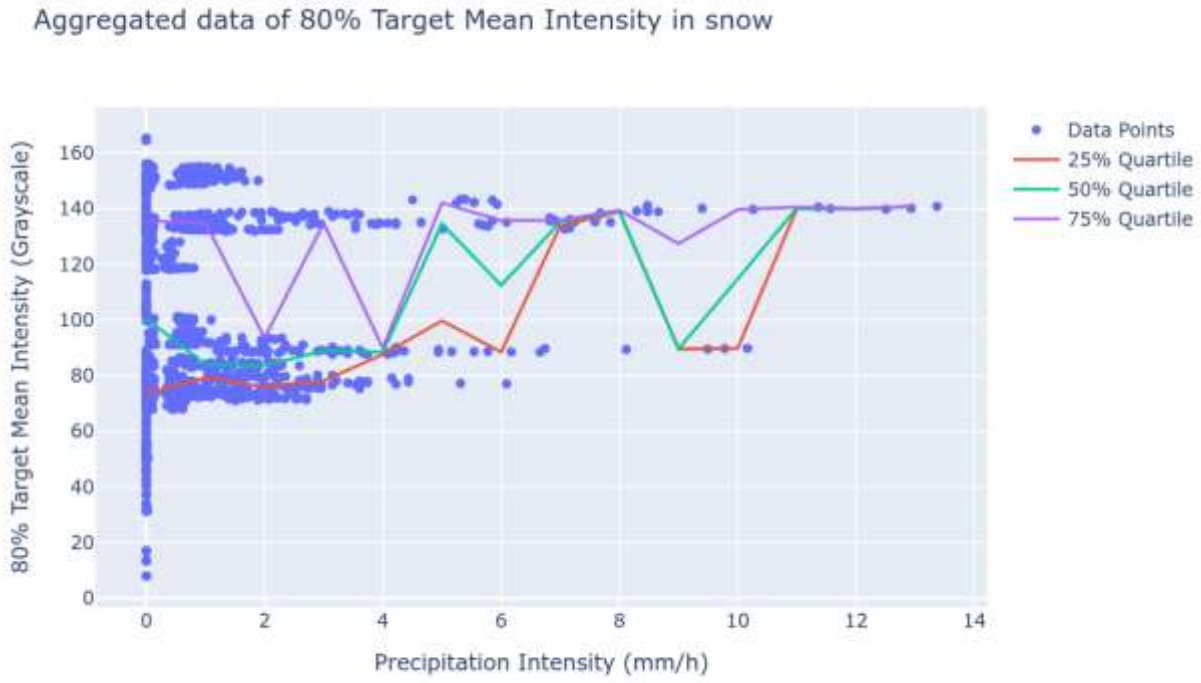

Figure 52: Aggregated data of 80% Target: Mean intensity in snow

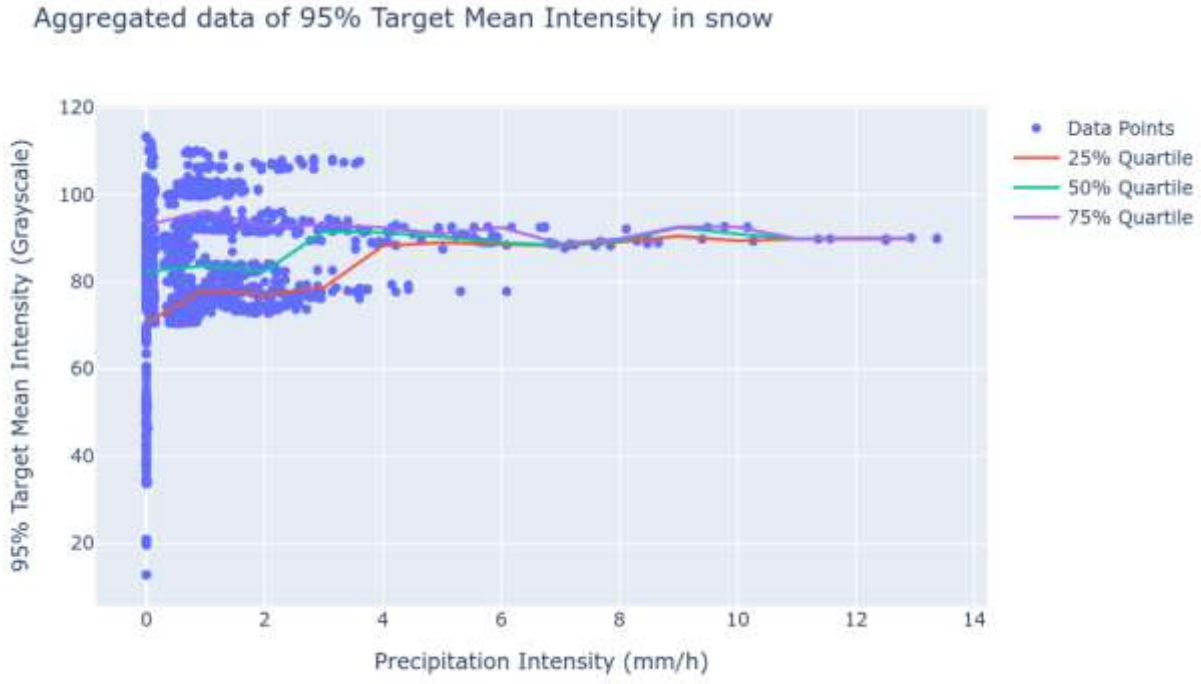

Figure 53: Aggregated data of 95% Target: Mean intensity in snow

4.1.5.2 Camera contrasts in snow

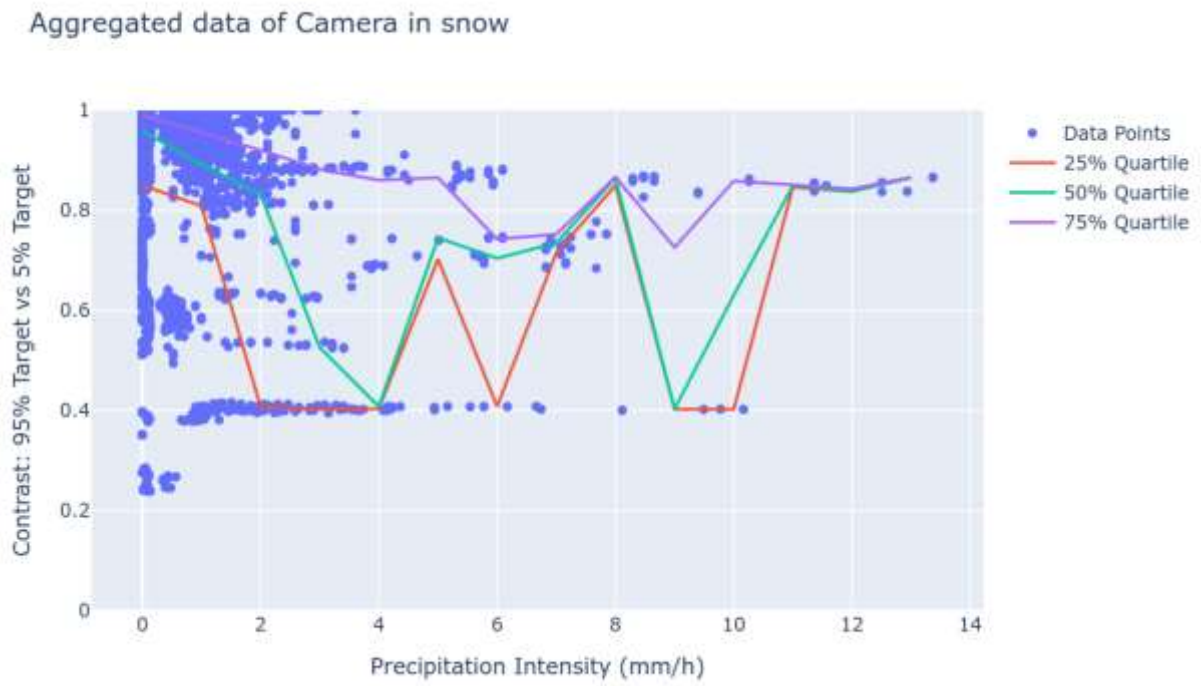

Figure 54: Aggregated contrast of camera in snow: 95% target vs. 5% target

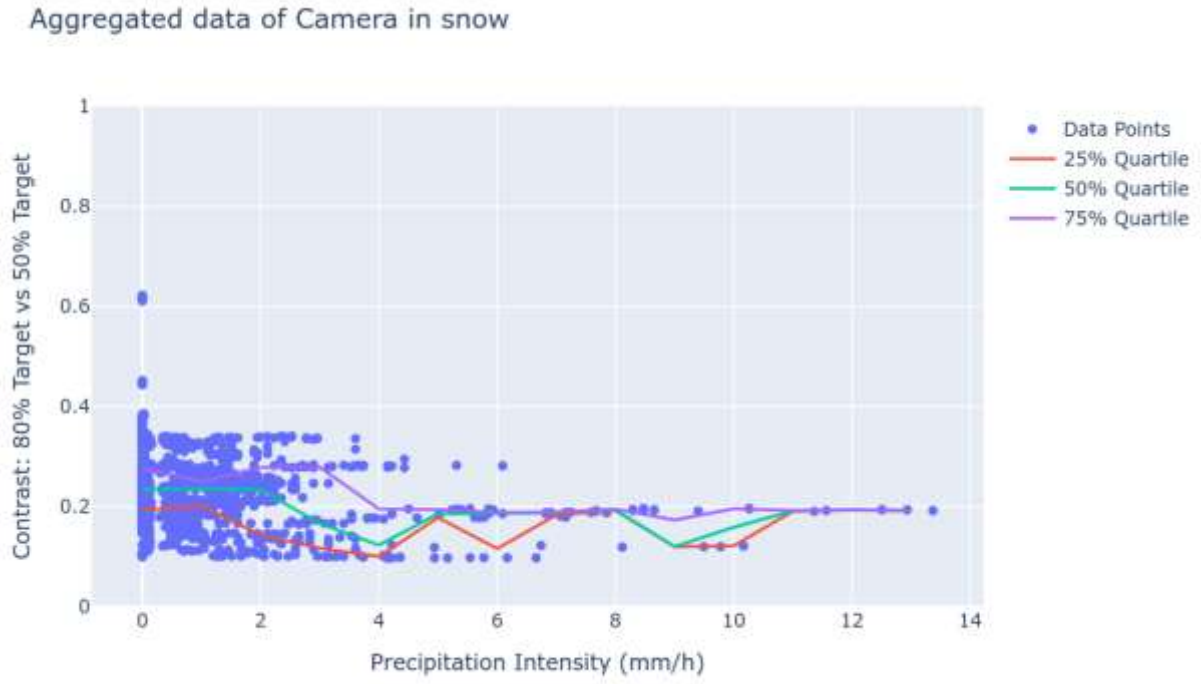

Figure 55: Aggregated contrast of camera in snow: 80% target vs. 50% target

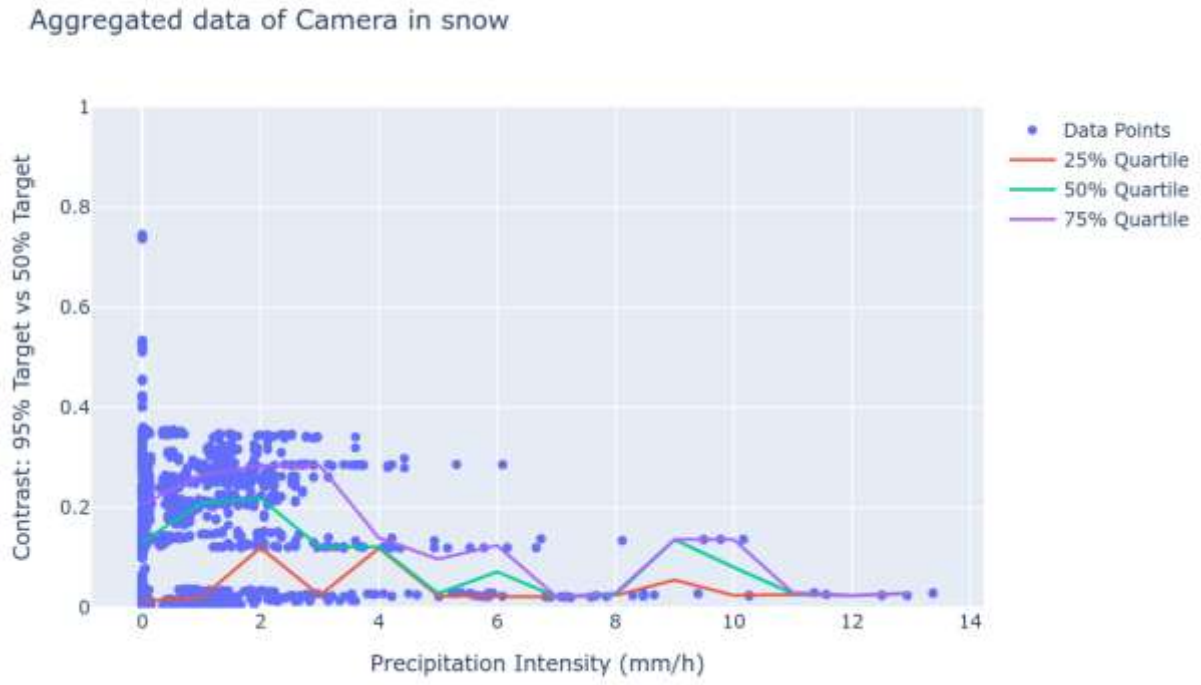

Figure 56: Aggregated contrast of camera in snow: 95% target vs. 50% target

4.1.5.3 Interim conclusion for camera performance in snow

Overall, the snow measurements indicate that snowfall affects the evaluated camera KPIs mainly through increased variability rather than through a strong monotonic trend over the measured precipitation intensities. The mean intensity curves for the individual targets remain broadly within a comparable range, but show a larger scatter than in clear daytime reference recordings, which is consistent with intermittent snowflake occlusions and local brightness fluctuations in the image.

The contrast curves exhibit the same qualitative pattern: contrasts involving the darkest target (5%) tend to show the highest sensitivity and spread, whereas contrasts between mid/high reflectivity targets (e.g., 80% vs. 50%) remain comparatively stable on average. In contrast to fog, no clear "global" contrast attenuation is visible across all target pairs; instead, the effect appears to be dominated by transient, spatially inhomogeneous disturbances.

A plausible explanation is that snow primarily impacts the image through a combination of short-lived occlusions (flakes crossing the line of sight), motion-streak artifacts, and additional forward scattering that can add a weak veiling component, especially when the scene is illuminated by bright surroundings. Unlike rain, snowflake size, shape, and fall velocity vary strongly, and the precipitation intensity in mm/h is therefore only an indirect proxy for the actual optical disturbance within the camera's field of view. Moreover, snow-covered ground can substantially increase scene brightness and stray light, which may trigger camera adaptation (auto-exposure/tone mapping) and thereby confound a simple mapping between mm/h and image statistics.

For simulation and validation, this suggests treating snowfall as a stochastic disturbance process similar to rain (increased KPI variance and occasional outliers due to flakes/streaks)

potentially coupled with a background/illumination change when snow cover is present. Consequently, the derived curves should be interpreted within the measured intensity range and should not be extrapolated without additional data, because different exposure regimes or denser snowfall could introduce a more pronounced veiling effect and a sustained contrast loss.

4.1.6 Conclusion for camera performance in different weather conditions

Across the evaluated weather conditions, the camera KPIs show two fundamentally different degradation patterns. Fog leads to a clear, physically consistent and largely monotonic loss of contrast with increasing extinction coefficient, accompanied by an increase in mean intensity due to veiling luminance (airlight). This behavior affects all target combinations and therefore represents a true global visibility limitation. In contrast, precipitation conditions (rain, snow, hail) do not exhibit a comparably strong and uniform attenuation trend in the aggregated mean-intensity and contrast metrics. Instead, their main signature is an increased variance and scatter of the KPIs, consistent with spatially and temporally localized disturbances such as transient occlusions, streaks, specular highlights, and window/wet-surface glare that may or may not intersect the fixed ROI at a given time. The brightness evaluation further highlights that camera-internal adaptation (e.g., auto-exposure and tone mapping) can be a confounding factor even in the absence of adverse weather: contrasts involving the darkest reference target are particularly sensitive because lifting dark regions reduces relative contrast by definition and can be reinforced by flare/veiling glare.

Taken together, the results indicate that performance curves for appearance-based KPIs should be interpreted as functions of both environmental severity (e.g., extinction coefficient for fog, precipitation rate as a proxy in rain/snow/hail) and the camera operating regime. For simulation and validation, this supports (i) treating fog with an explicit global veiling/attenuation component parameterized by extinction (or MOR), and (ii) representing rain/snow/hail primarily as stochastic disturbance processes that increase KPI variance and generate outliers rather than as deterministic contrast-attenuation curves.

Finally, the presented curves are specific to the used camera, the fixed target configuration, and the employed aggregation over three consecutive frames. Precipitation intensity in mm/h is only an indirect proxy for optical disturbance and does not capture particle size distribution, wind-driven trajectories, or illumination changes (e.g., snow cover), which can dominate the observed scatter. Several categories also have limited coverage at the high-intensity end (notably hail), so conclusions should be restricted to the measured ranges. To maximize traceability and reusability of the results, it is recommended to document camera settings and any available information on exposure/gain behavior and to avoid extrapolation across illumination regimes or beyond the covered weather-severity intervals.

4.2 Actual performance curves of 4 LiDARs

4.2.1 Four different LiDAR sensors in comparison

All sensors are from the latest generation and reflect 4 different characteristics. Two are 360° rotating sensors and two are directional sensors with ca. 120° horizontal field of view (FoV). We have different maximum ranges and also different wavelengths.

Characteristic	Sensor A	Sensor B	Sensor C	Sensor D
Lidar type	Rotating	Rotating	Directional	Directional
Vertical FoV	~90°	~45°	~20°	~70°
Horizontal FoV	360°	360°	≥120°	≥120°
Typical max range	~45–50 m	~100–120 m	up to ~200 m	~70 m
Angular resolution (V)	Coarse	Medium–fine	Fine (≈0.1°)	Medium (anisotropic)
Vertical scan lines	128	128	96–128	192
Horizontal angular resolution	0.18°	0.18°	0.10°	0.15°
Vertical angular resolution	0.70°	0.35°	0.20°	0.36°
Operating wavelength	**865 nm**	**865 nm**	**940 nm**	**905 nm**
Perception focus	Near-field	Mid-range	Long-range	Side / blind-spot

4.2.2 LIDAR PERFORMANCE IN BRIGHTNESS

Around 42600 measurement time steps were taken of all lidar sensors as daily measurements in clear conditions. The measurements were taken every day if the conditions were clear at 4:00 and at 13:00. The weather sensor measures brightness in lux in 4 directions and additionally gives the maximum brightness in lux per measurement time step. The distribution of brightness measurements in all daily recordings is shown in the following figure.

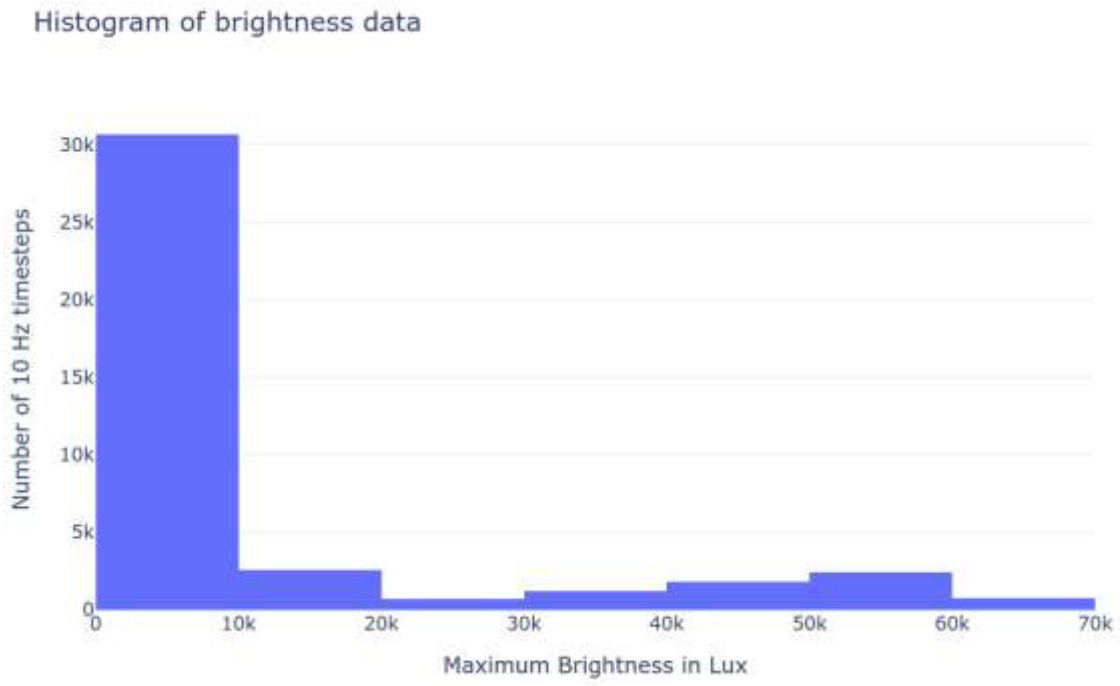

Figure 57: Histogram of brightness data for lidar evaluation

As described in 2.5, the following evaluation will be structured regarding the filter regions: atmosphere and targets 1 - 5. For each filter region, the (relative) number of detections as well as the intensity (or reflectivity) values are plotted over the maximum brightness of the respective time frame. Each plot represents the median value of the respective sensor with the corresponding maximum brightness over all daily measurements.
The plots in this chapter are median values aggregated from all sensors for comparison.

4.2.2.1 Lidar detections in atmosphere in brightness

As expected, the detection probability of all sensors in these clear conditions is at 0 for all brightnesses. There is just one insignificant outlier in the following plot.

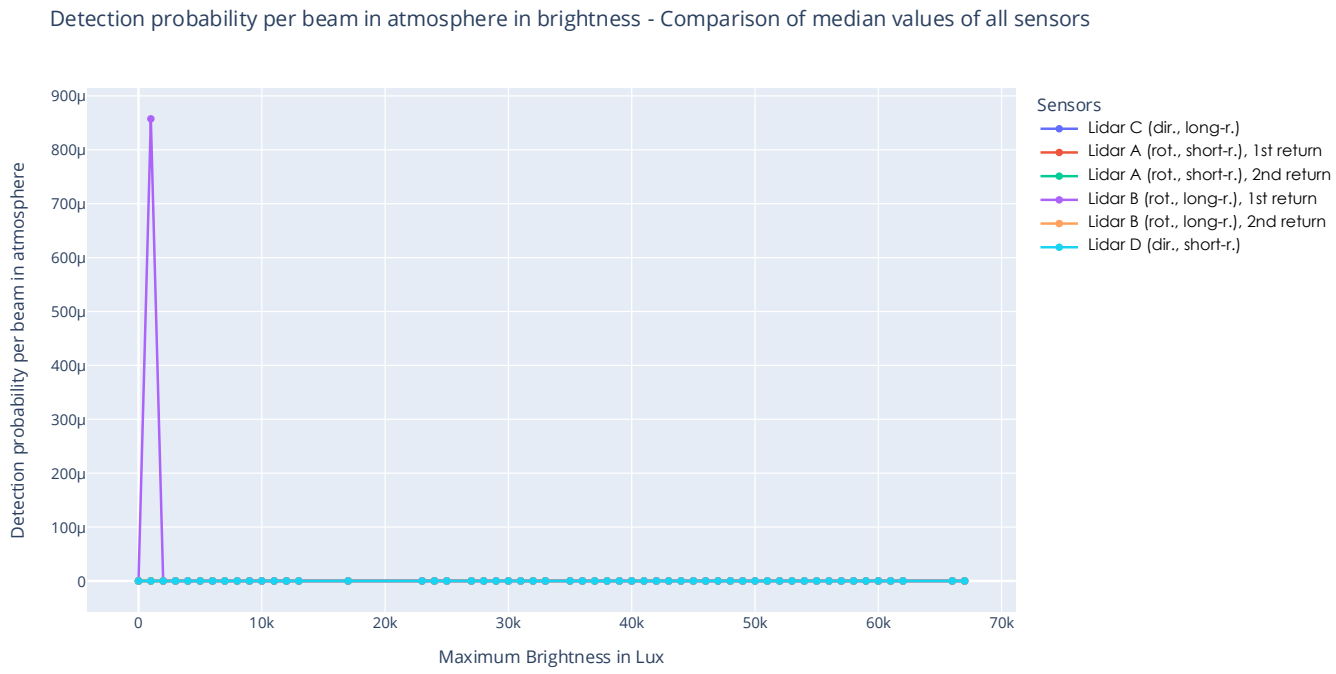

Figure 58: Lidar detection probability per beam in brightness

4.2.2.2 Lidar detections on Target 1 (10 %) in brightness

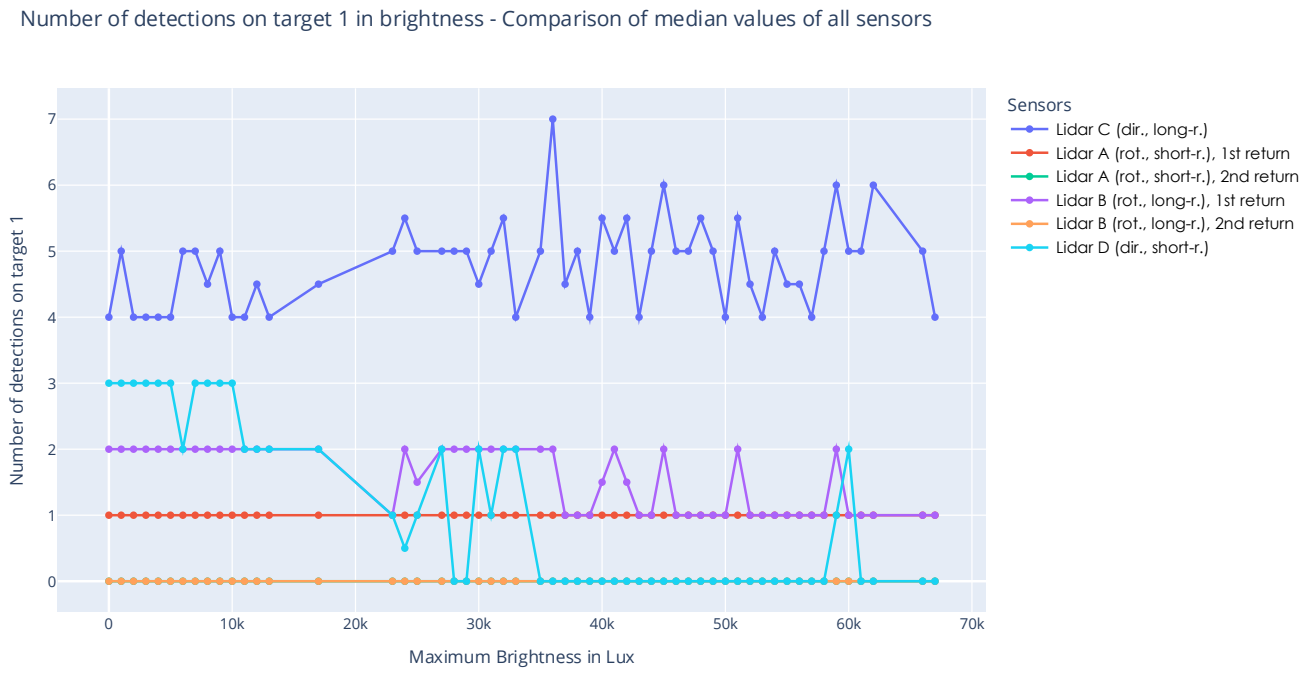

Figure 59: Number of detections on target 1 in brightness

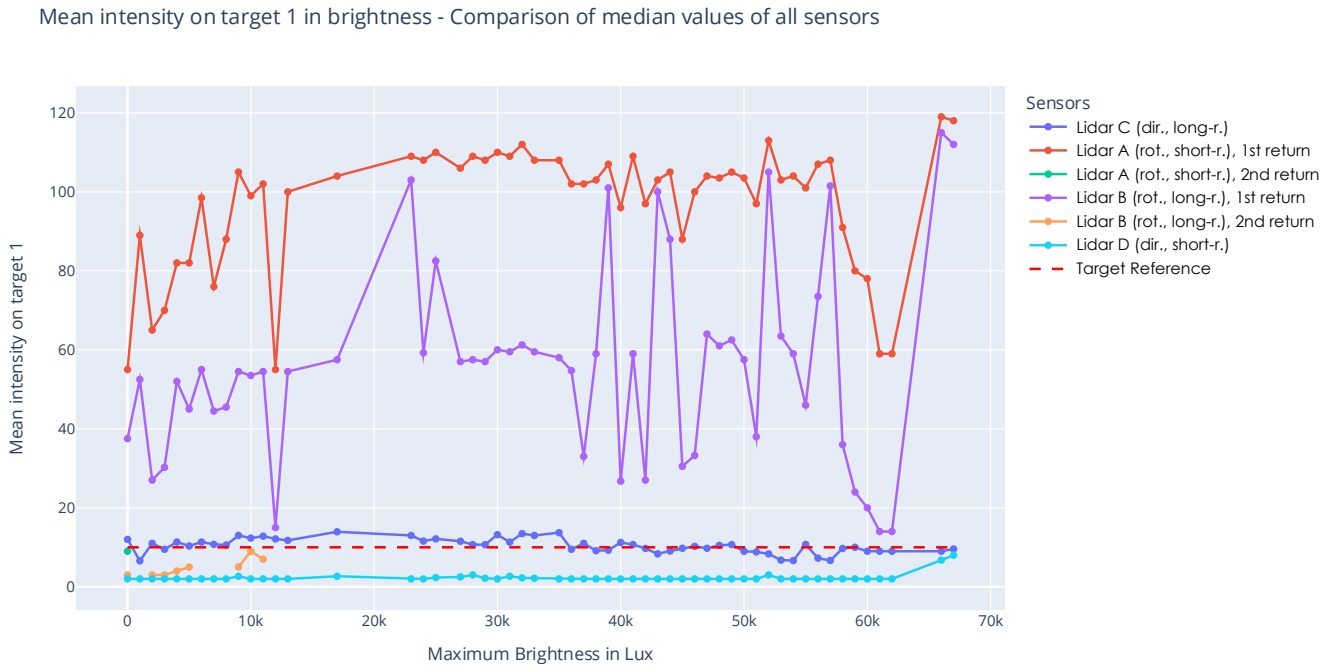

Figure 60: Mean intensity of detections on target 1 in brightness

4.2.2.3 Lidar detections on Target 2 (95 %) in brightness

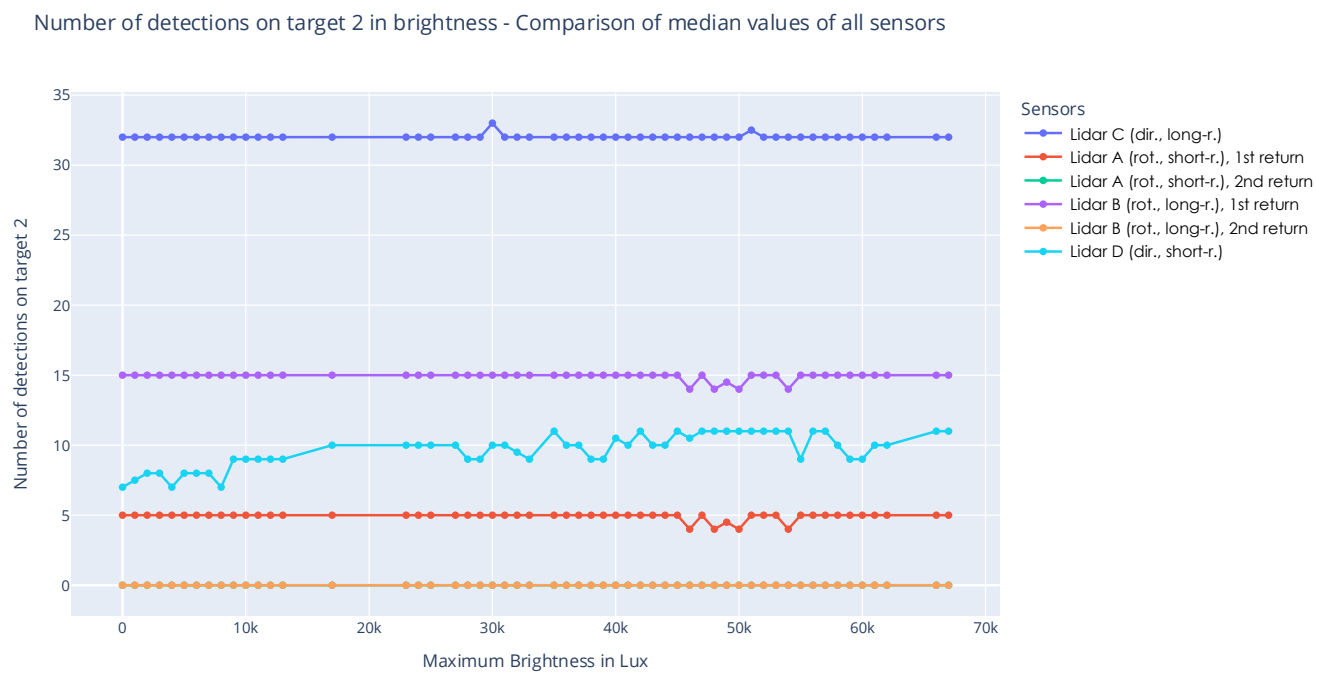

Figure 61: Number of detections on target 2 in brightness



Figure 62: Mean intensity of detections on target 2 in brightness

4.2.2.4 Lidar detections on Target 3 (5 %) in brightness

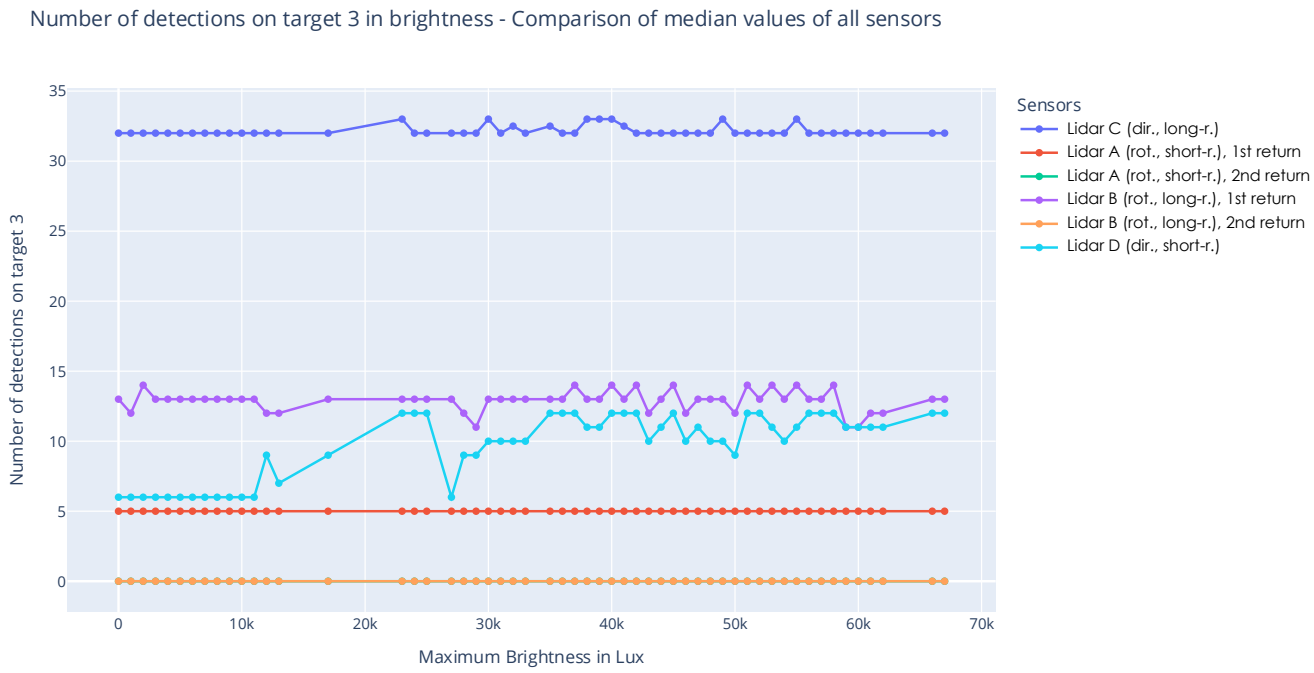

Figure 63: Number of detections on target 3 in brightness



Figure 64: Mean intensity of detections on target 3 in brightness

4.2.2.5 Lidar detections on Target 4 (50 %) in brightness

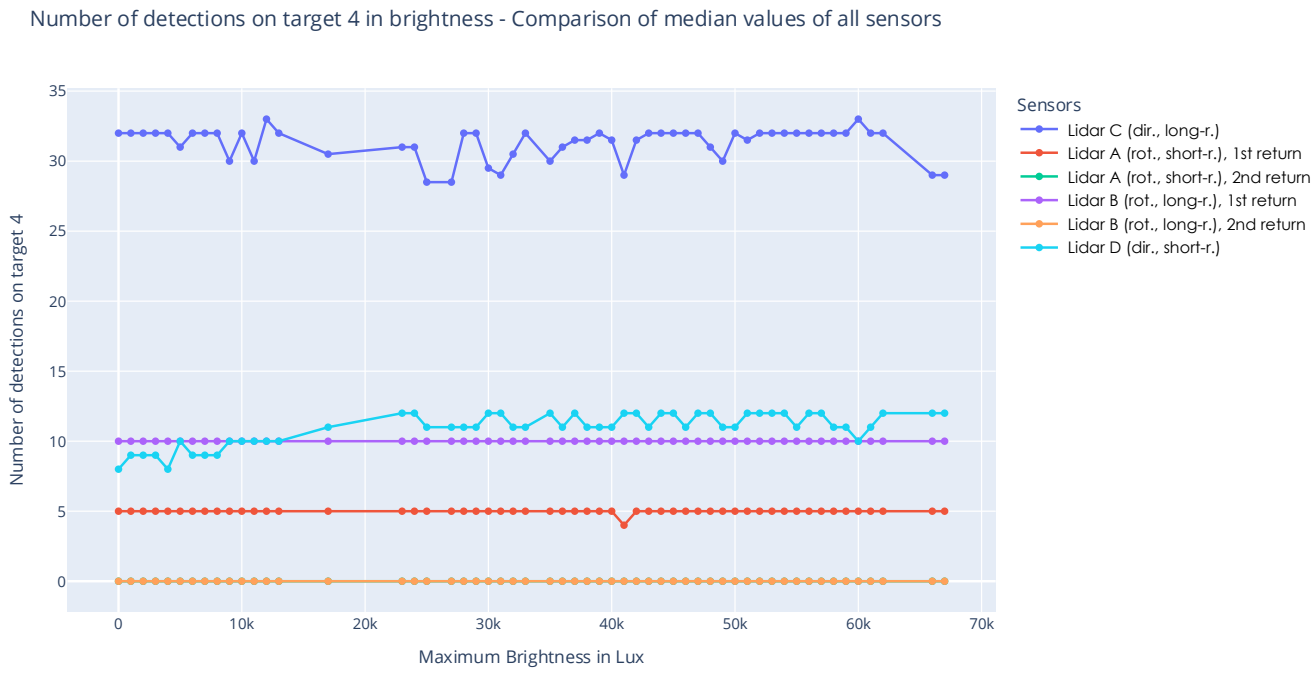

Figure 65: Number of detections on target 4 in brightness



Figure 66: Mean intensity of detections on target 4 in brightness

4.2.2.6 Lidar detections on Target 5 (80 %) in brightness

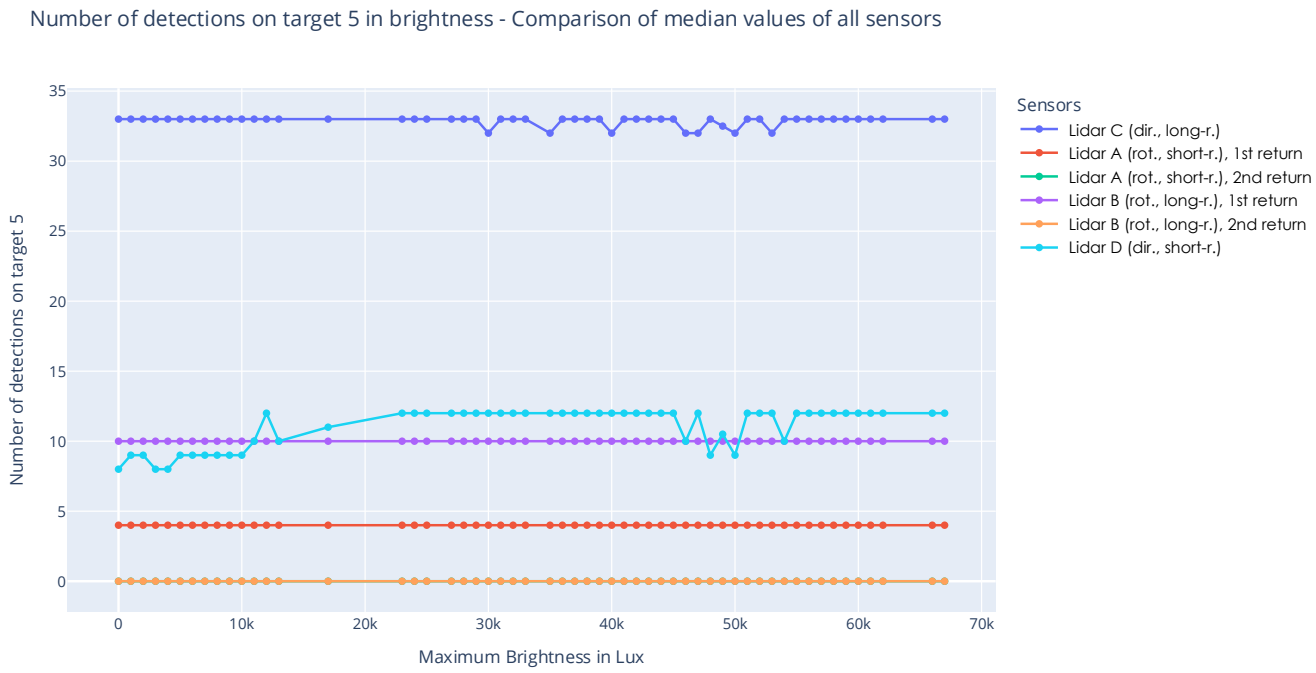

Figure 67: Number of detections on target 5 in brightness

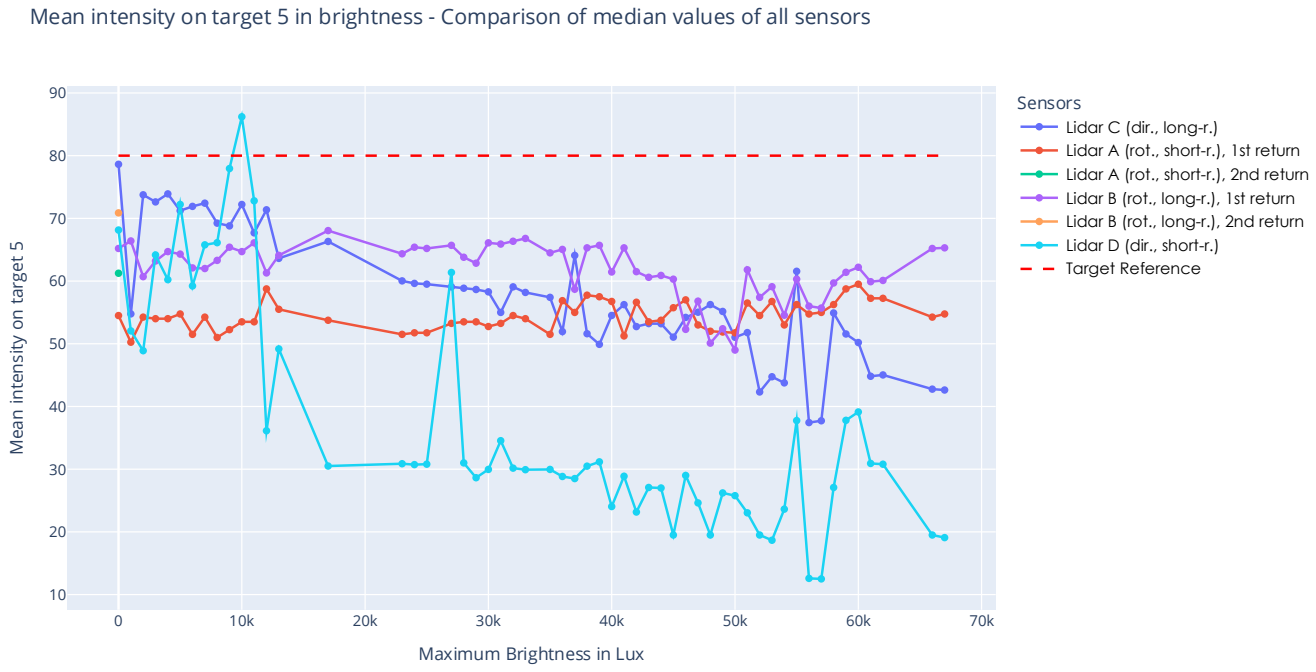

Figure 68: Mean intensity of detections on target 5 in brightness

4.2.2.7 Interim conclusion for lidar performance in brightness

No atmospheric noise detections were observed on any of the sensors across the tested brightness range, indicating that none of the systems produced spurious environmental returns within the tested conditions. The short-range directional lidar (D) showed a reduction in detection performance on the far target (10 % target reflectivity at 78.5 m), with detections beginning to drop at approximately 35,000 lux. The rotating lidar with small vertical FoV (lidar B) exhibited a slight decline in detection probability at around 25,000 lux but continued to register the far target with at least one detection. In contrast, the directional lidar and the rotational lidar with small vertical FoV (lidars B and C) demonstrated robustness at long range, maintaining full detection capability even under elevated brightness.

Reflectivity measurements on the far target varied substantially for both rotating sensors (A and B), with values significantly exceeding expected levels and showing considerable fluctuation under increasing ambient illumination. Both directional lidars (C and D) delivered comparatively stable behavior, with consistent intensity measurements throughout the test conditions.

Across all sensors, no influence of ambient brightness was detected on the detection performance for the four closer targets, confirming strong short□range resilience. However, the intensity output of the short-range directional lidar (D) for these close targets shows a strong dependence on ambient brightness, indicating that its intensity estimation is sensitive to illumination even when detection performance remains stable.

4.2.3 LIDAR PERFORMANCE IN FOG

Around 391000 measurement time steps were taken of all lidar sensors during foggy conditions. Fog is quantified in either the Meteorological Optical Range (MOR) or the Forward Scattering Extinction Coefficient (EXCO). MOR in km is defined as 3/EXCO. In the following plots the EXCO will be used, as it has a finer resolution in the visibility sensor data. The distribution of extinction coefficients in all fog measurements is shown in the following figure.

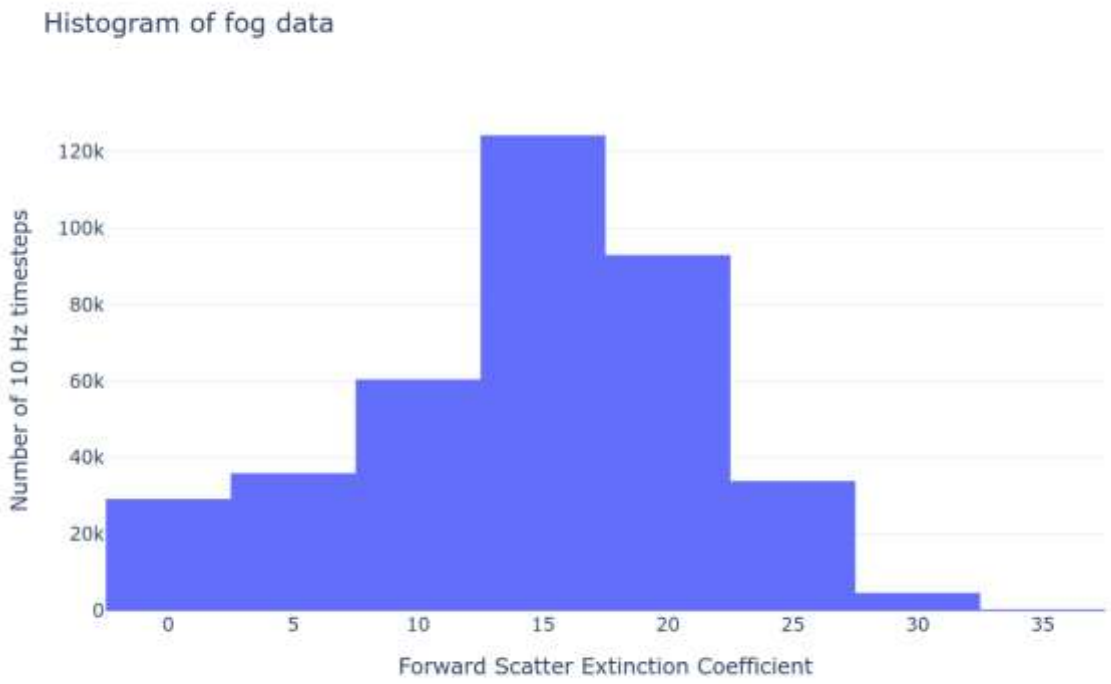

Figure 69: Histogram of extinction coefficients in the fog measurements

As described in 2.5, the following evaluation will be structured regarding the filter regions: atmosphere and targets 1 - 5. For each filter region, the (relative) number of detections as well as the intensity (or reflectivity) values are plotted over the extinction coefficient of the respective time frame. Each plot represents the median value of the respective sensor with the corresponding extinction coefficient over all fog measurements.
The plots in this chapter are median values aggregated from all sensors for comparison.

4.2.3.1 Lidar detections in atmosphere in fog

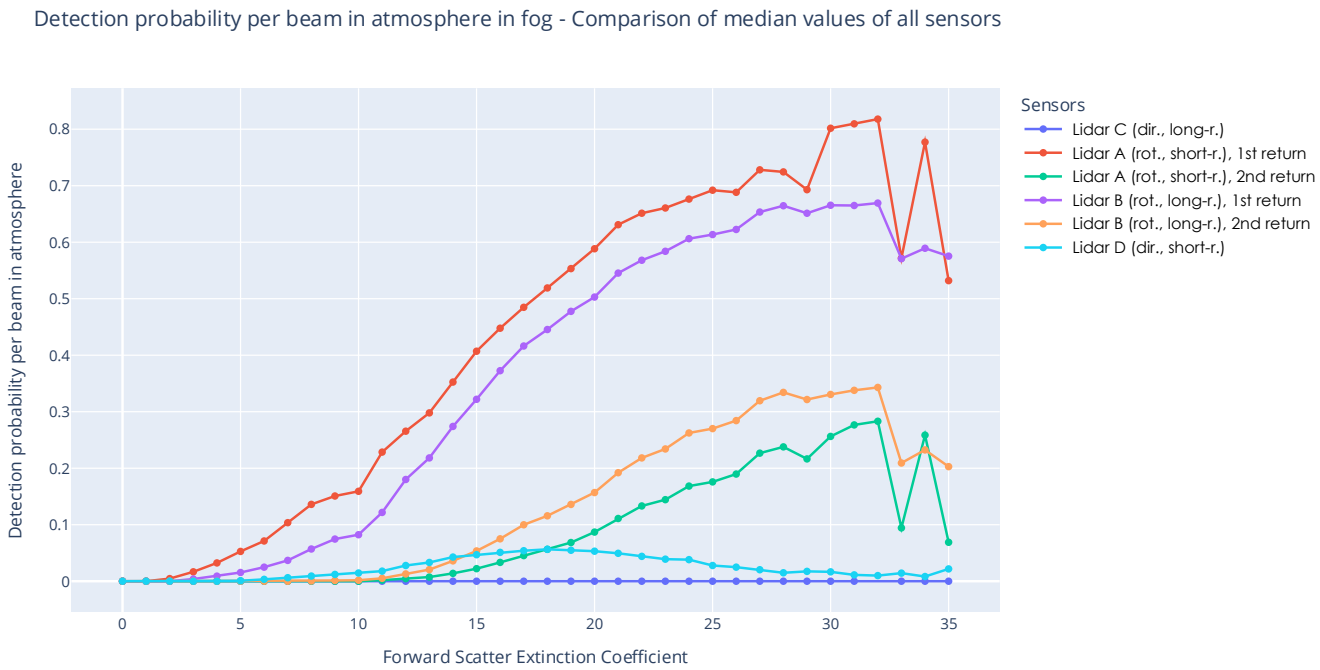

Figure 70: Detection probability per beam in atmosphere in fog

4.2.3.2 Lidar detections on Target 1 (10 %) in fog

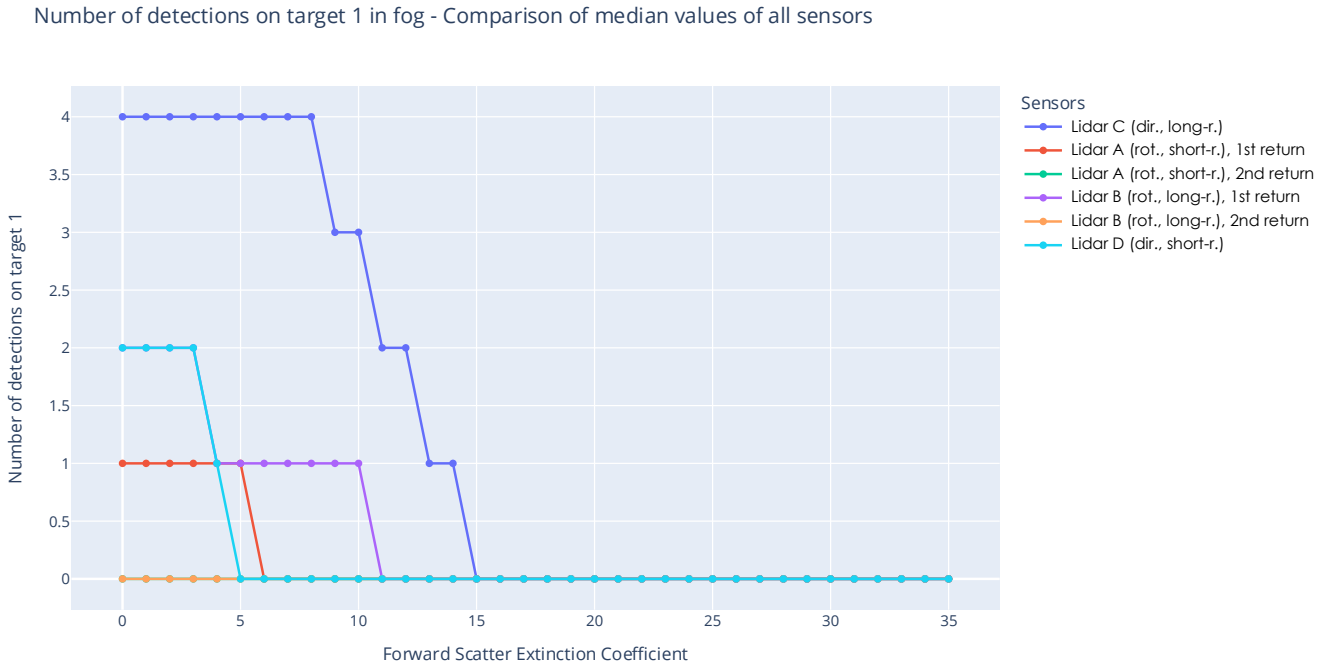

Figure 71: Number of detections on target 1 in fog

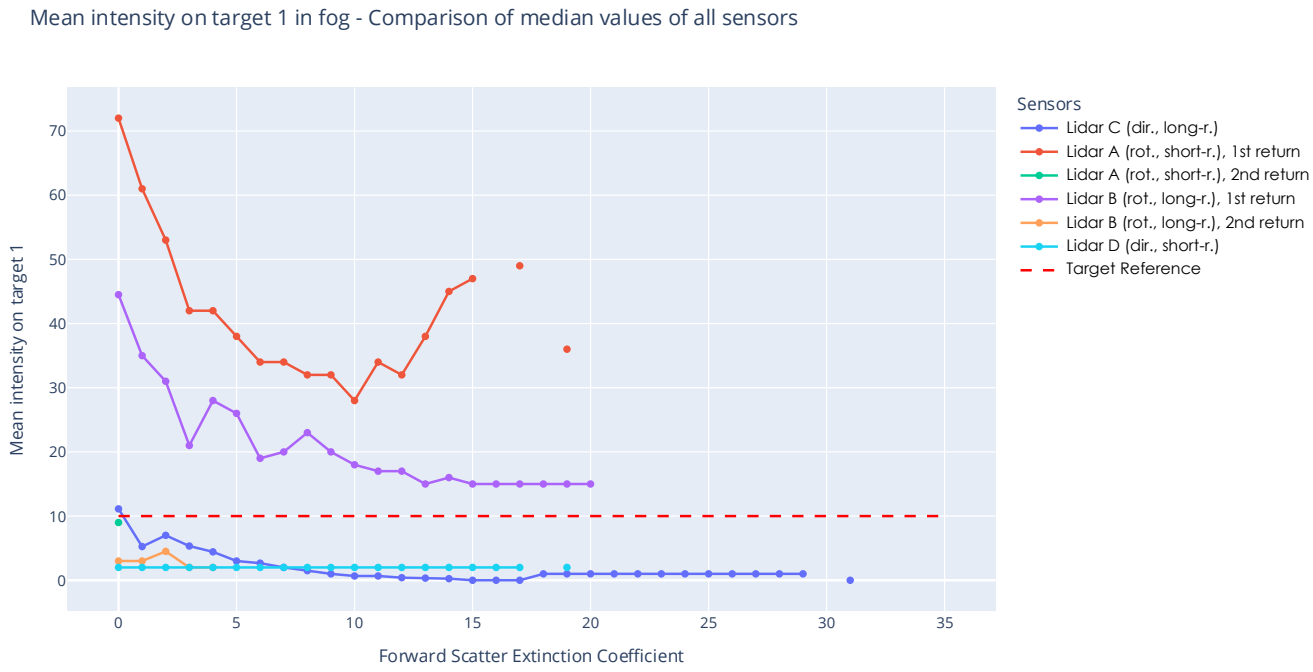

Figure 72: Mean intensity of detections on target 1 in fog

4.2.3.3 Lidar detections on Target 2 (95 %) in fog



Figure 73: Number of detections on target 2 in fog

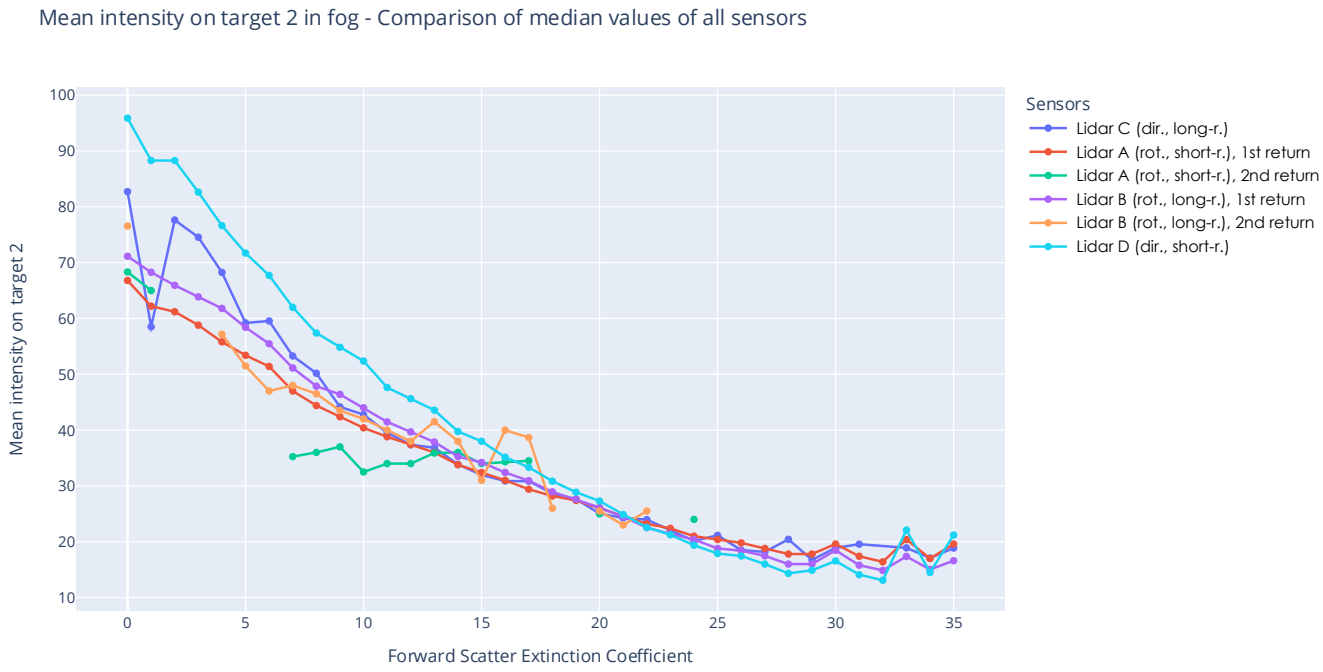

Figure 74: Mean intensity of detections on target 2 in fog

4.2.3.4 Lidar detections on Target 3 (5 %) in fog

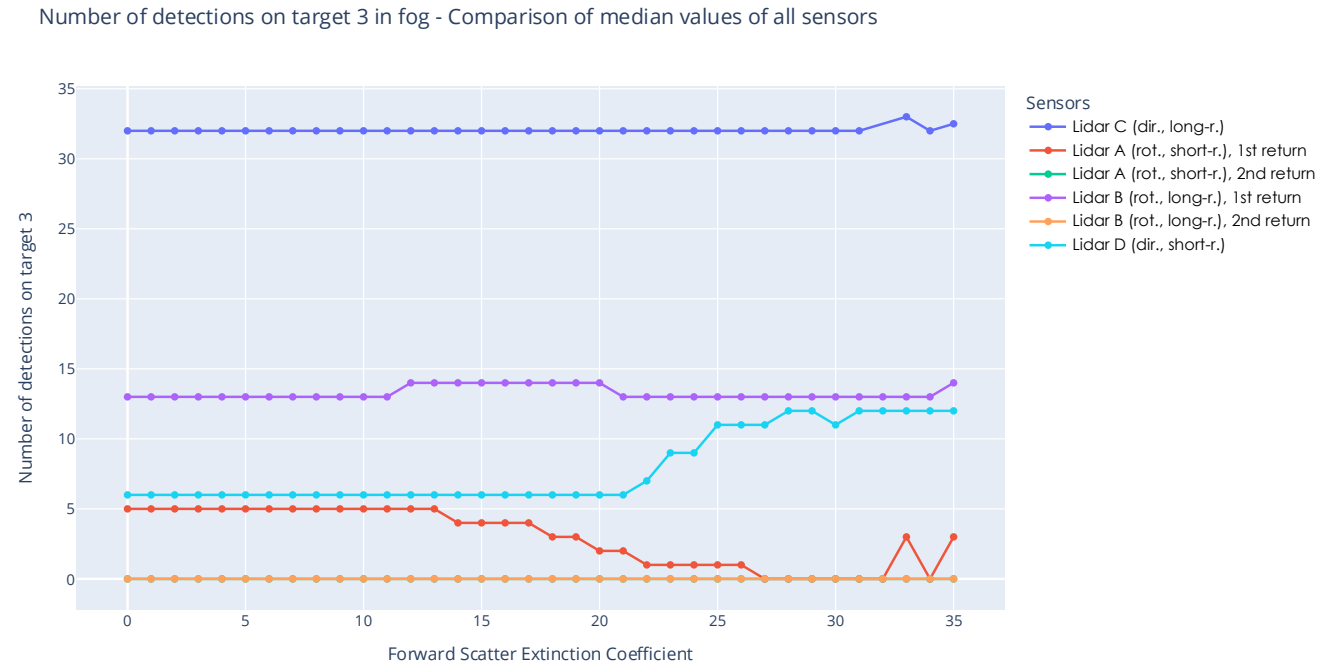

Figure 75: Number of detections on target 3 in fog

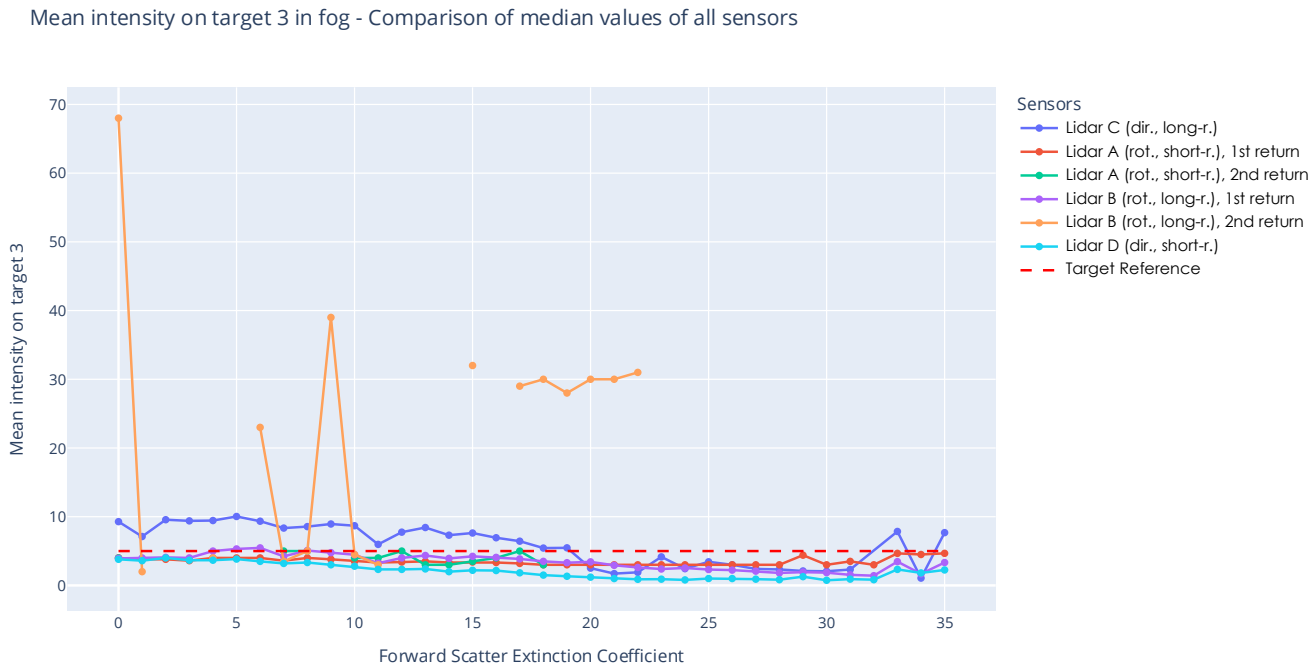

Figure 76: Mean intensity of detections on target 3 in fog

4.2.3.5 Lidar detections on Target 4 (50 %) in fog

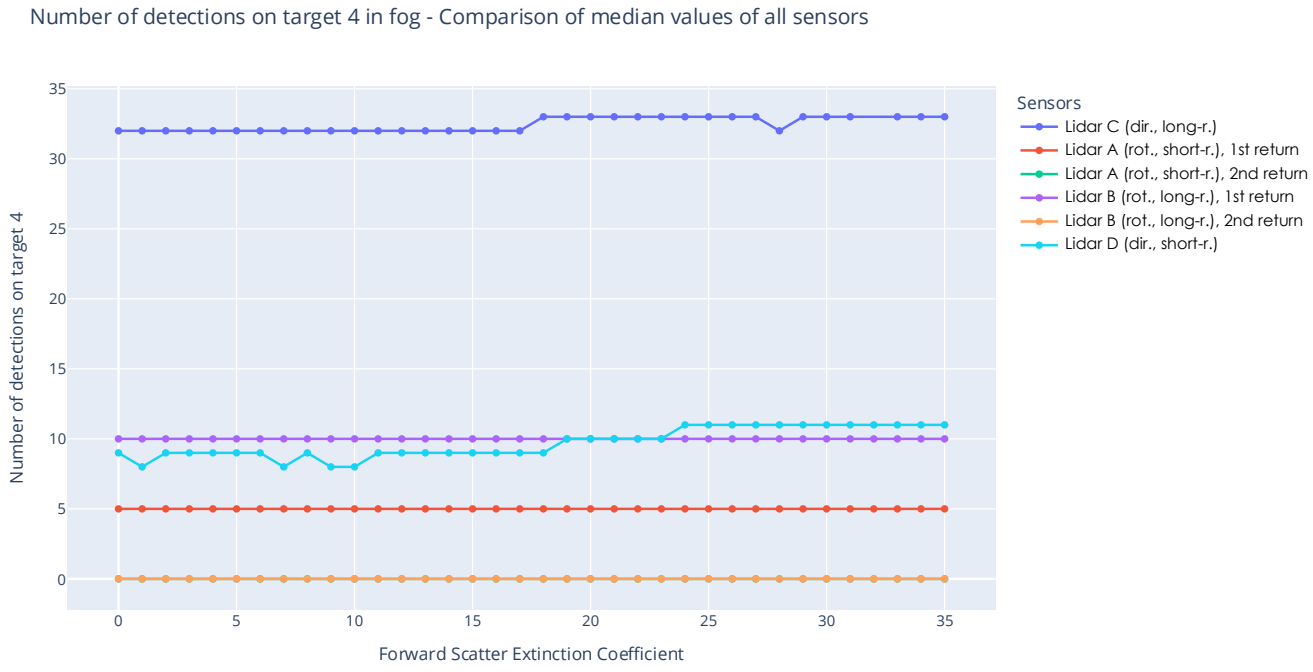

Figure 77: Number of detections on target 4 in fog



Figure 78: Mean intensity of detections on target 4 in fog

4.2.3.6 Lidar detections on Target 5 (80 %) in fog



Figure 79: Number of detections on target 5 in fog

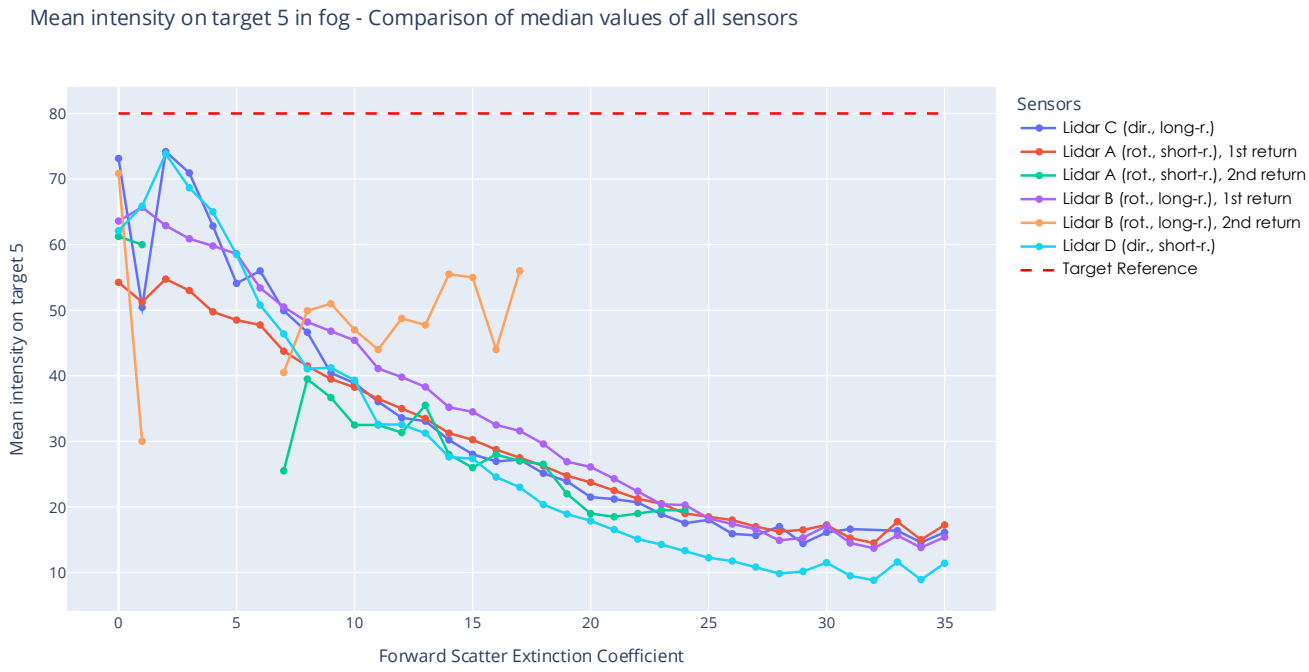

Figure 80: Mean intensity of detections on target 5 in fog

4.2.3.7 Special Consideration: Uneven Fog

The sensors behave differently if the fog is uneven with thicker fog patches in the atmosphere. The images below depict the same time frame of all four lidars in a measurement during patchy fog. The point clouds are colored by z-coordinate. The dark points are from the ground level. Everything in color is above and belongs to the fog patches, except for the target in the center of the images. Especially noteworthy is the long-range directional lidar (C) which is quite robust in homogeneous fog conditions. Here it shows very distinct clusters of points in the thicker fog patches. The behavior is similar in the short-range directional lidar (D) data but not as pronounced. The rotating sensors (A and B) seem not as affected by the unevenness of the fog but show a similar behavior as in other fog measurements.

In Figure 85 the camera image corresponding to the same timestep in the measurement is shown for reference. While different fog patches may be seen in the image upon close inspection, it is not as big of a difference as the lidar point clouds would suggest.
The clustering of points is most critical regarding object detection algorithms which may create false-positive objects due to the fog. An even distribution of atmospheric points would potentially be easier to filter out in post-processing.

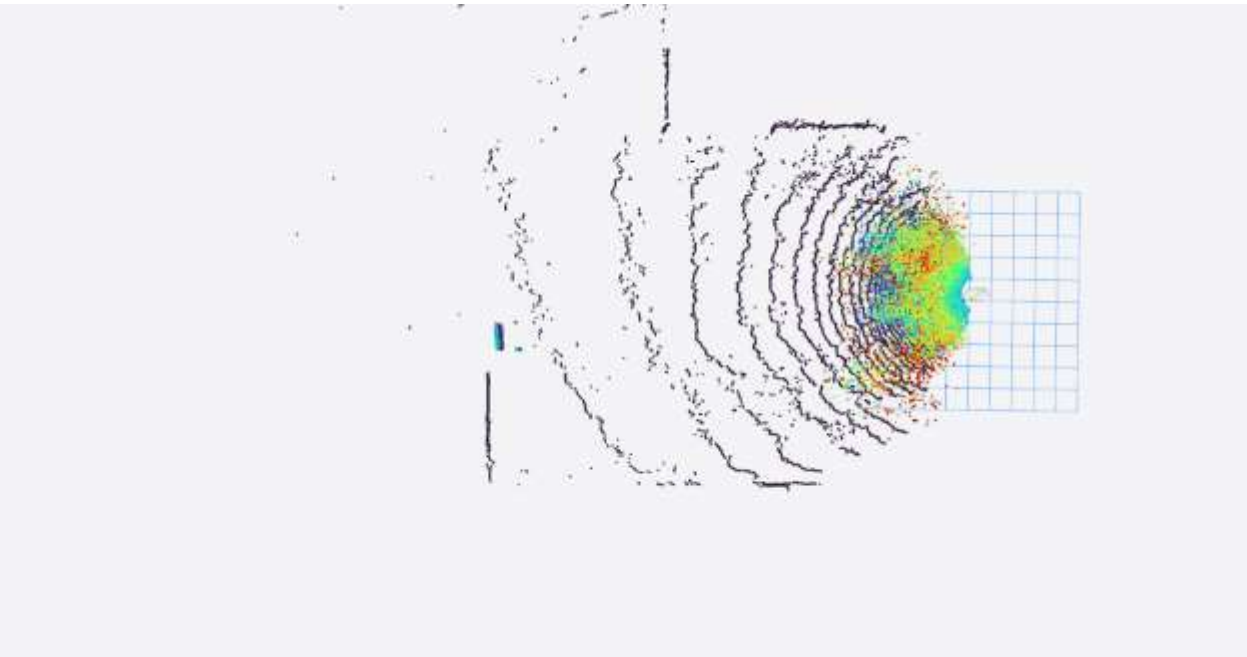

Figure 81: Short-range rotating lidar (lidar A) in uneven fog. The points are colored by z-coordinate. All colored points, except from the target, are due to the fog.

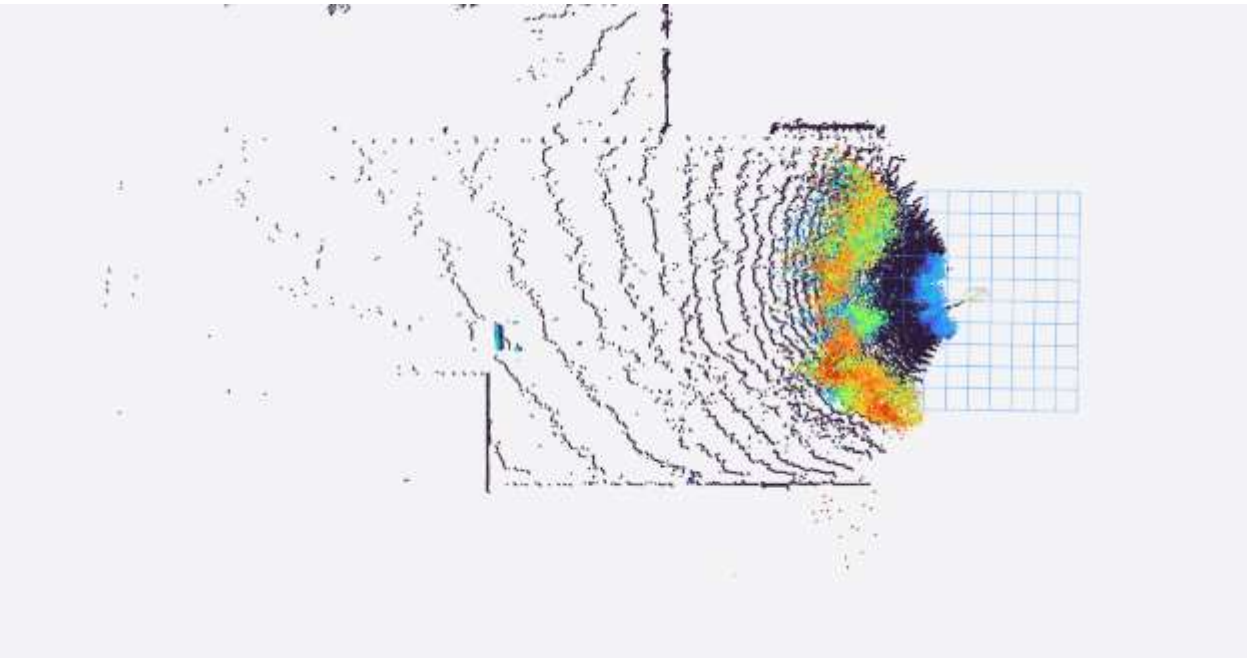

Figure 82: Long-range rotating lidar (lidar B) in uneven fog. The points are colored by z-coordinate. All colored points, except from the target, are due to the fog.

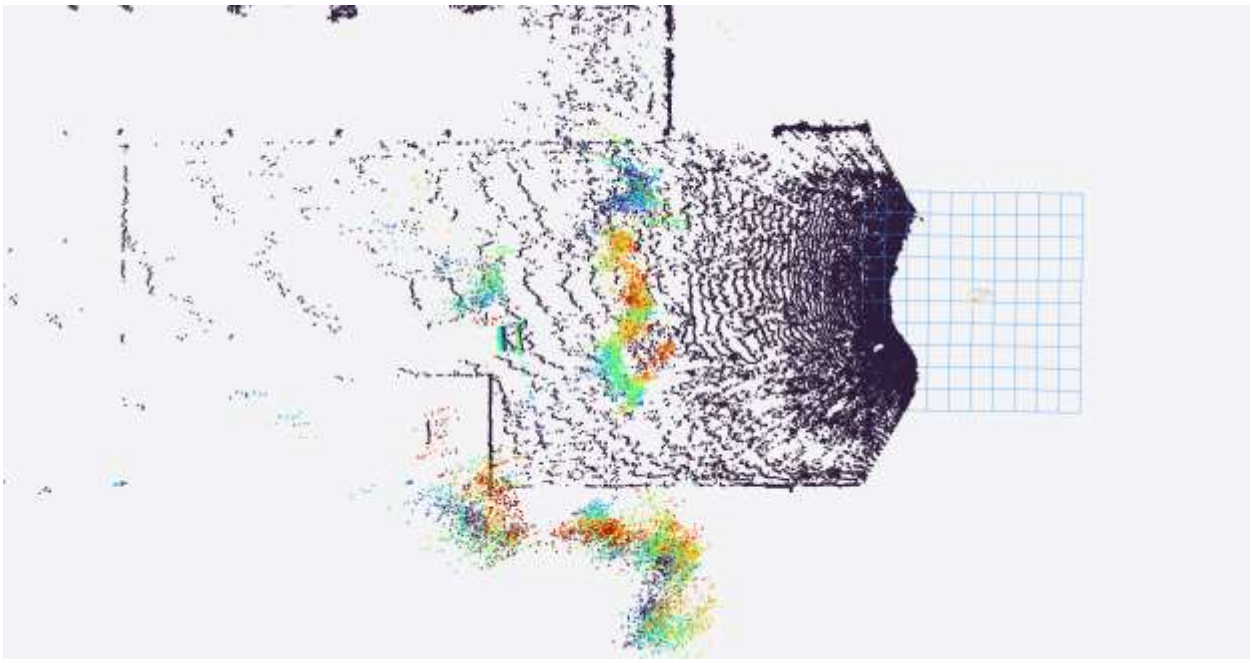

Figure 83: Long-range directional lidar (C in uneven fog. The points are colored by z-coordinate. All colored points, except from the target, are due to the fog.

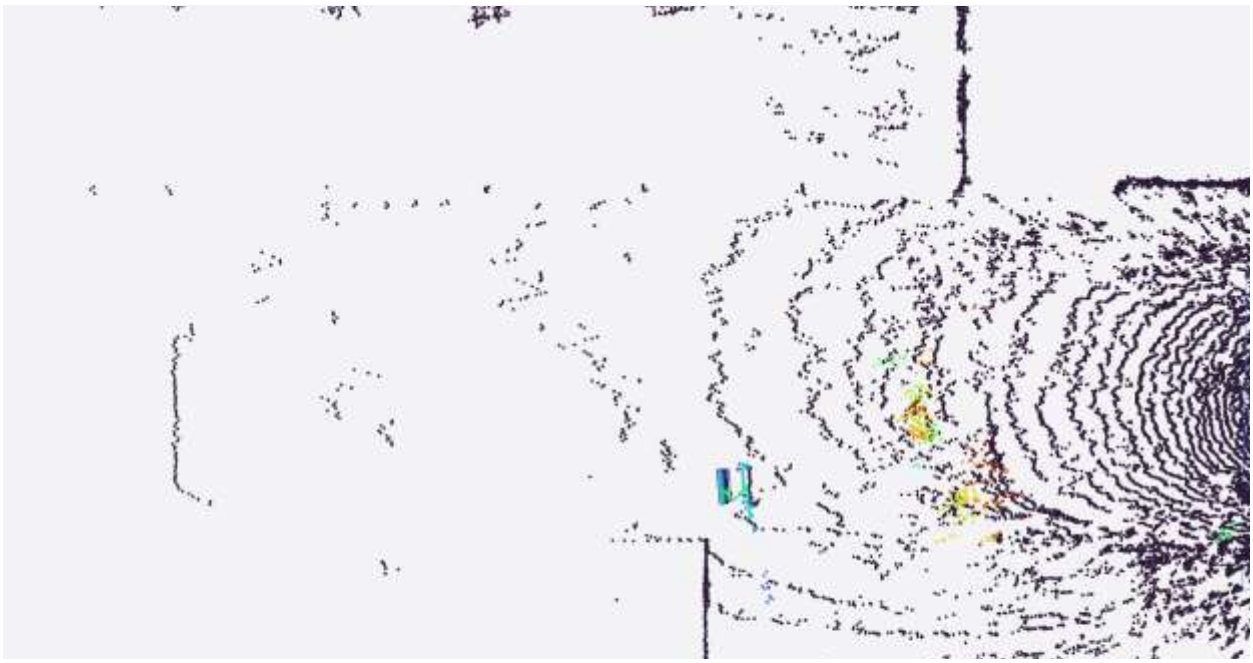

Figure 84: Short-range directional lidar (D) in uneven fog. The points are colored by z-coordinate. All colored points, except from the target, are due to the fog.



Figure 85: Camera image in uneven fog as reference

4.2.3.8 Interim conclusion for lidar performance in fog

In fog, the probability of scatter‑induced atmospheric detections increases markedly in both rotating sensors (A and B) as the extinction coefficient rises. The short-range rotating lidar (A) shows the strongest response, reaching a maximum detection probability of approximately 0.8, meaning that 80 % of all emitted beams generate a first‑return detection within the atmosphere. This maximum occurs at an extinction coefficient of 32, corresponding to a visibility of about 94 m. Alongside the rise in first‑return counts, the number of second returns also increases proportionally, indicating that multiple backscattering events become more prevalent as fog density increases.

The short-range directional lidar (D) shows only a slight increase in atmospheric detections beginning at an extinction coefficient of around 10. Beyond a value of roughly 20, the number of scatter detections decreases again, likely due to the balance between increased backscattering and the strong attenuation that begins to dominate at higher fog densities. The long-range directional lidar (C) remains highly robust throughout all tested fog levels and does not produce notable atmospheric scatter detections.

Detection performance on the far target drops significantly for the short-range directional lidar (D) and the short-range rotating lidar (A) at an extinction coefficient of around 5, after which both sensors can no longer detect the target. The long-range rotational lidar (B) and long-range directional lidar (C) maintain detections up to extinction coefficients of approximately 12. No sensor can detect the far target once the extinction coefficient exceeds 15. Intensity measurements for all sensors decrease steadily with increasing extinction, reflecting the cumulative impact of backscattering and attenuation on received signal strength.

At closer ranges, all targets with reflectivities of 50 % or higher remain stably detected across the full range of tested fog densities. Detection performance of the short-range directional lidar (D) on close targets even shows a slight improvement at higher extinction values, potentially caused by a contrast increase between the illuminated target and the attenuated

background. On the low-reflectivity 5 % target, the short-range rotating lidar (A) exhibits a gradual decline in detection probability beginning at an extinction coefficient of about 13 and loses all detections at a value of approximately 27, corresponding to a visibility near 110 m. Across all close targets and sensors, measured intensity decreases continuously as fog density increases, consistent with the expected attenuation of the optical signal.
The evaluation above only considers fairly homogenous fog. In the special case of patchy fog both directional lidars (C and D) show clusters of points in the atmosphere, presumably in denser fog patches. Especially the long-range directional lidar (C) which is very robust in even fog conditions seems to be very susceptible to spatially distributed fog densities.

4.2.4 LIDAR PERFORMANCE IN HAIL

Around 18000 measurement time steps were taken of all lidar sensors in hail. Hail intensity is quantified with the precipitation intensity in mm/h. The distribution of precipitation intensity in all hail measurements is shown in the following figure.

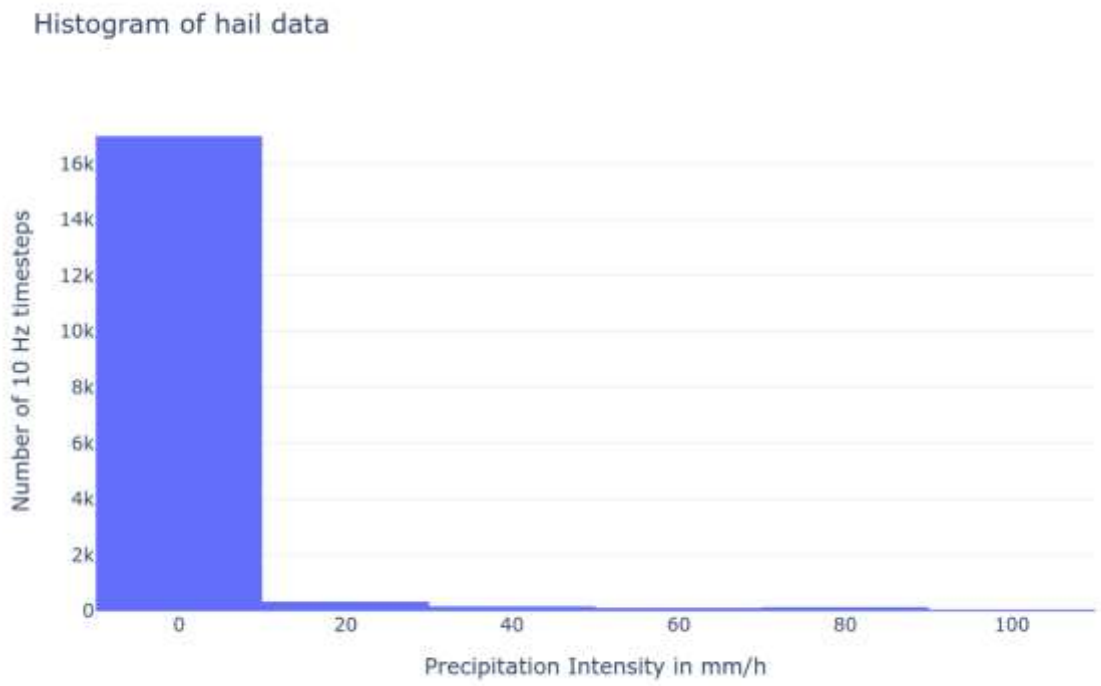

Figure 86: Histogram of hail data for lidar evaluation

As described in 2.5, the following evaluation will be structured regarding the filter regions: atmosphere and targets 1 - 5. For each filter region, the (relative) number of detections as well as the intensity (or reflectivity) values are plotted over the precipitation intensity of the respective time frame. Each plot represents the median value of the respective sensor with the corresponding precipitation intensity over all hail measurements.
The plots in this chapter are median values aggregated from all sensors for comparison.

4.2.4.1 Lidar detections in atmosphere in hail

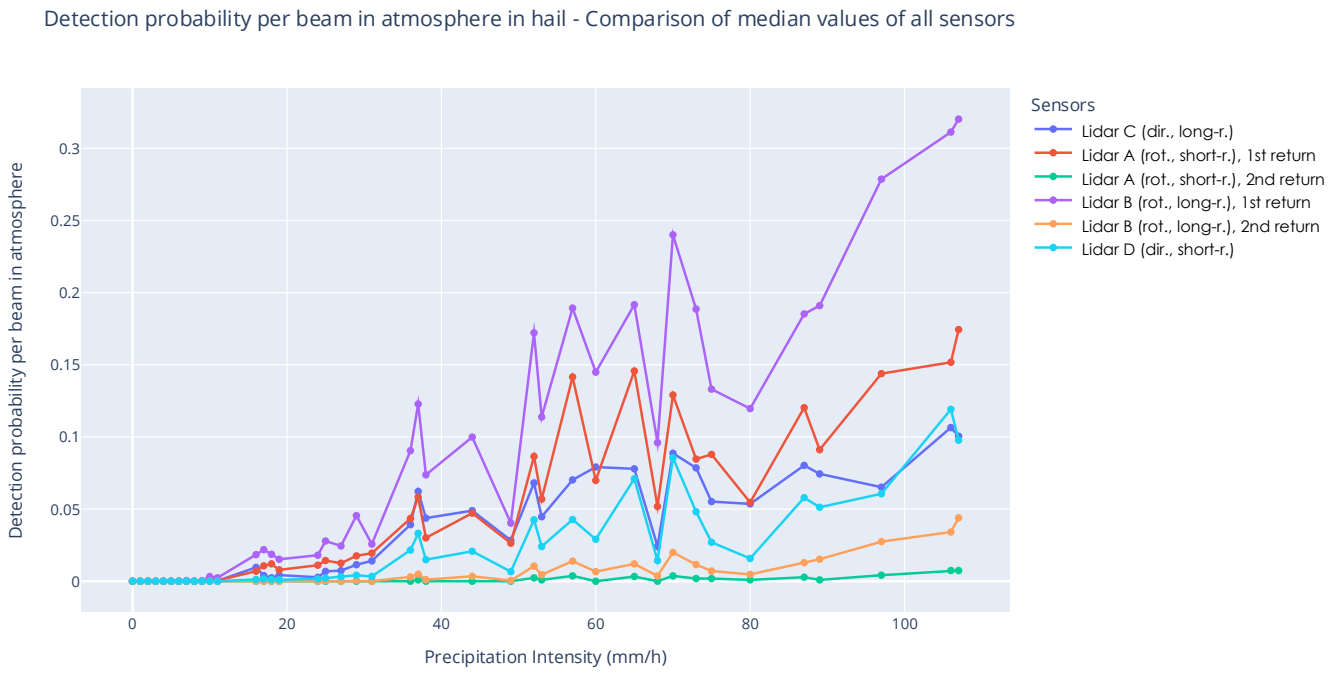

Figure 87: Detection probability per beam in atmosphere in hail

4.2.4.2 Lidar detections on Target 1 (10 %) in hail

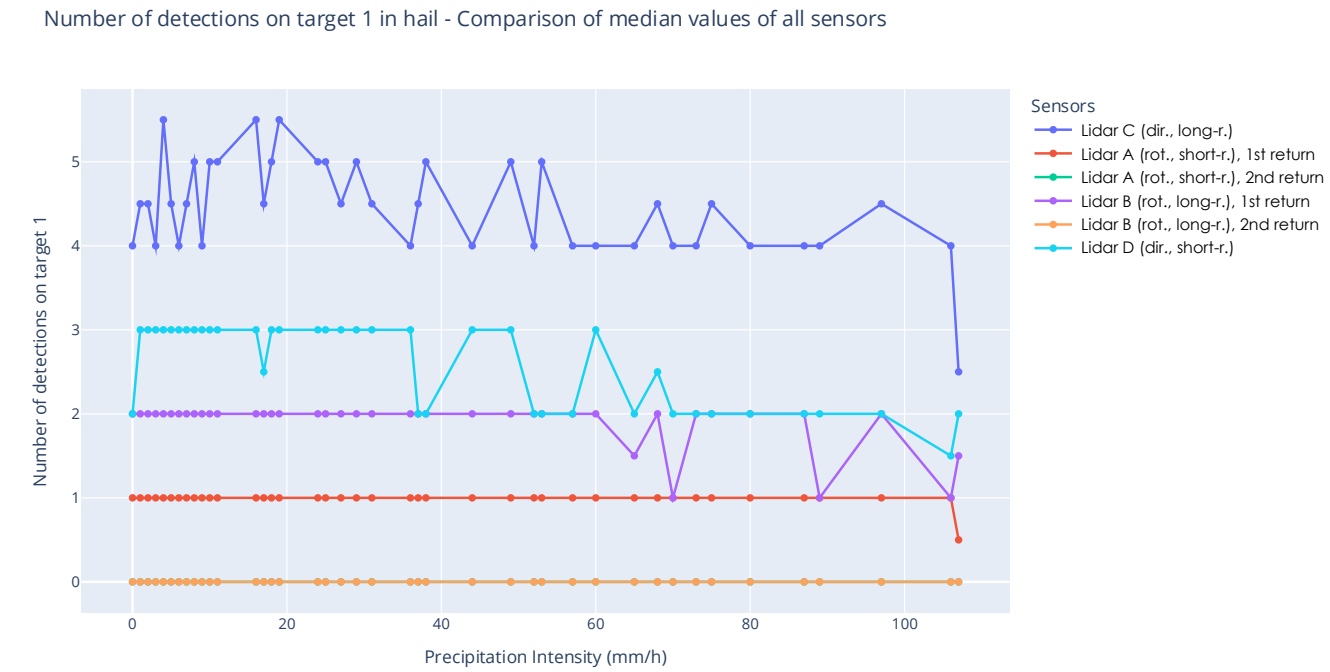

Figure 88: Number of detections on target 1 in hail

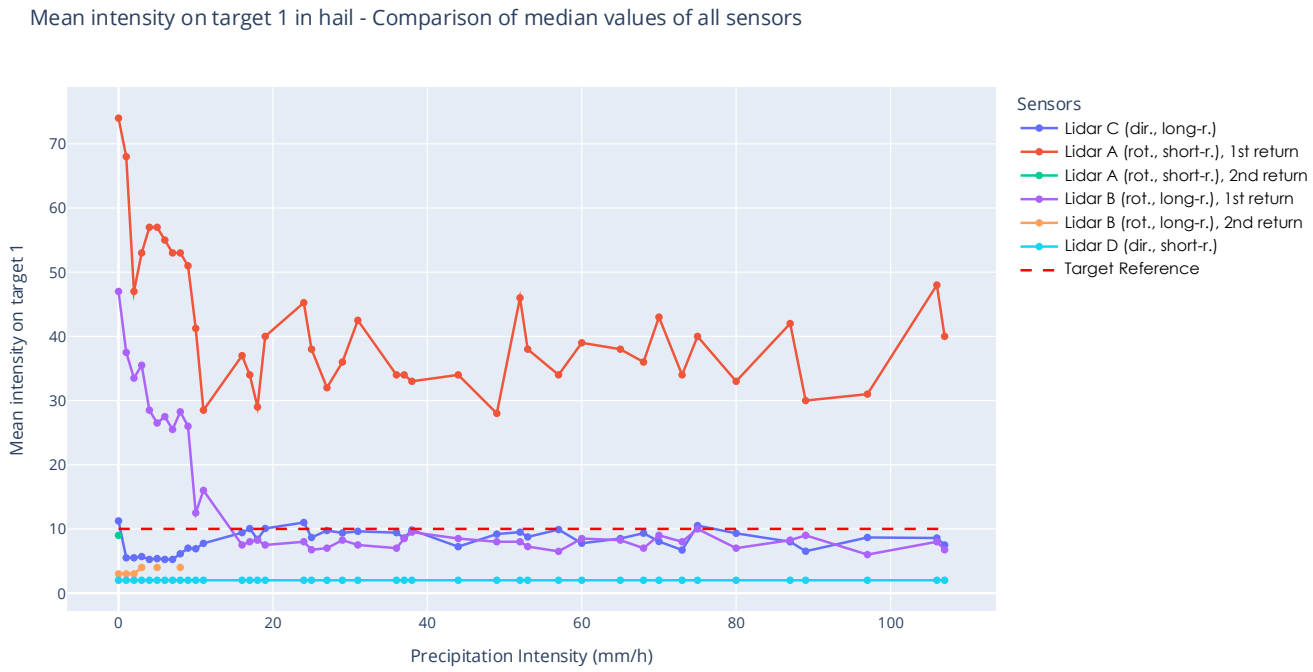

Figure 89: Mean intensity of detections on target 1 in hail

4.2.4.3 Lidar detections on Target 2 (95 %) in hail

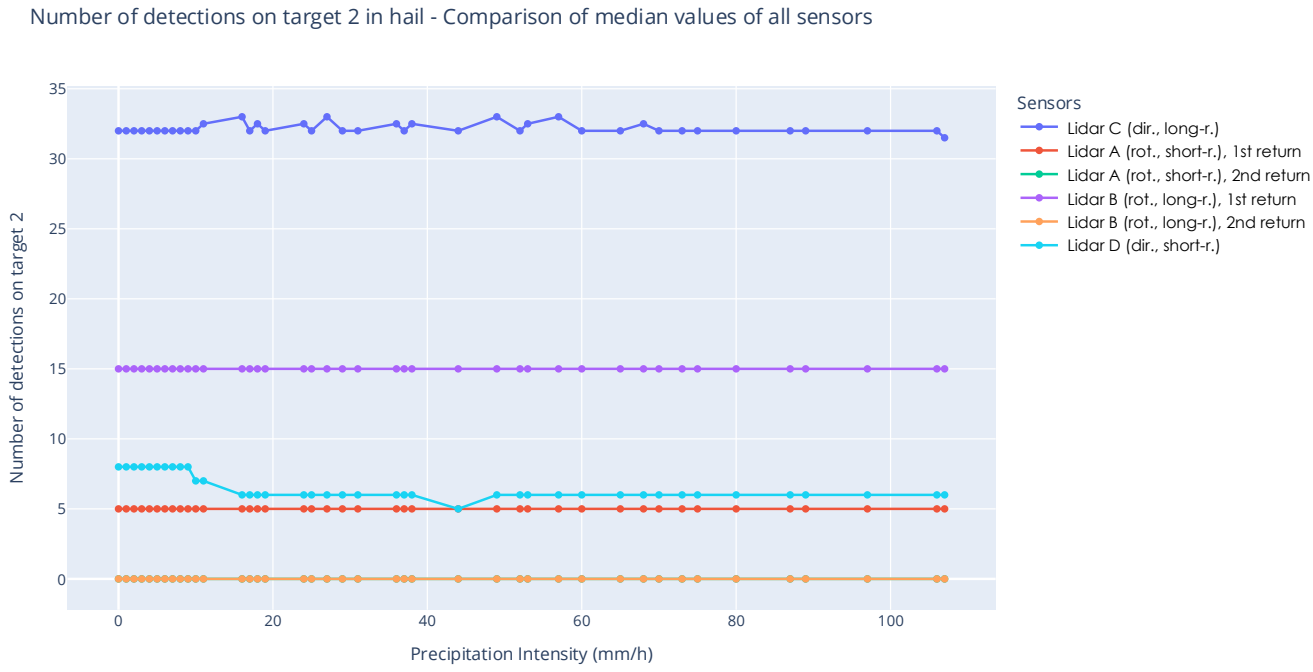

Figure 90: Number of detections on target 2 in hail

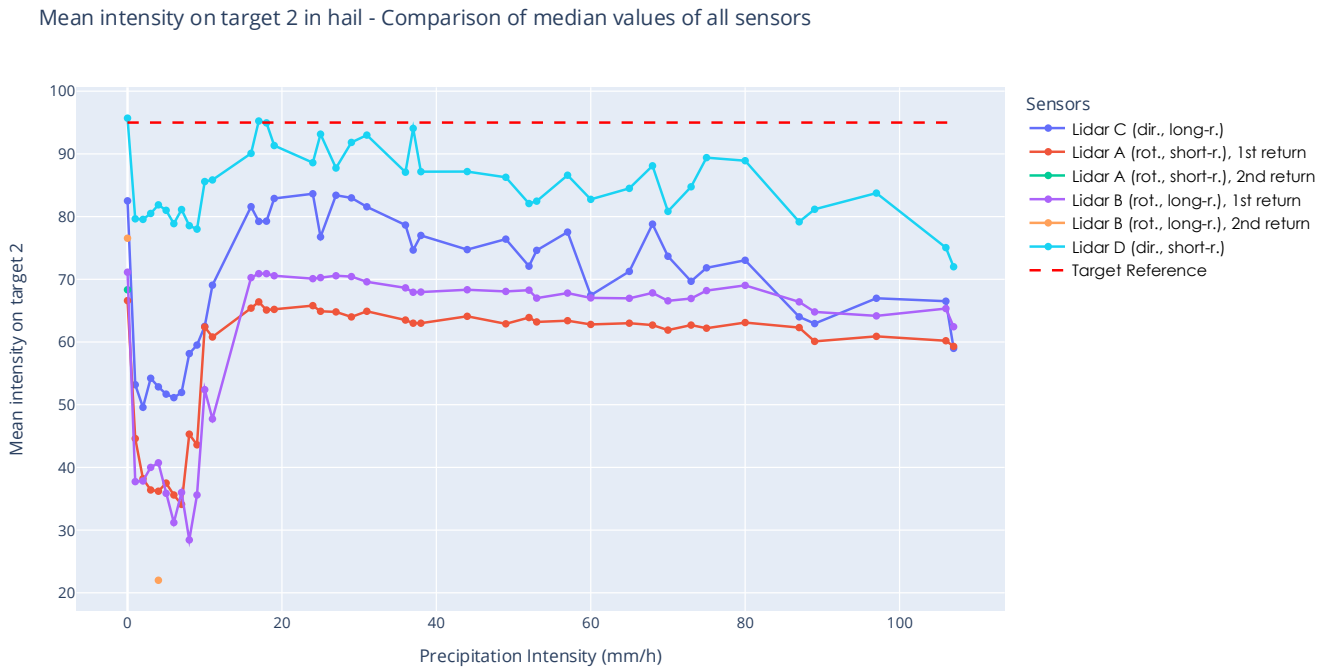

Figure 91: Mean intensity of detections on target 2 in hail

4.2.4.4 Lidar detections on Target 3 (5 %) in hail

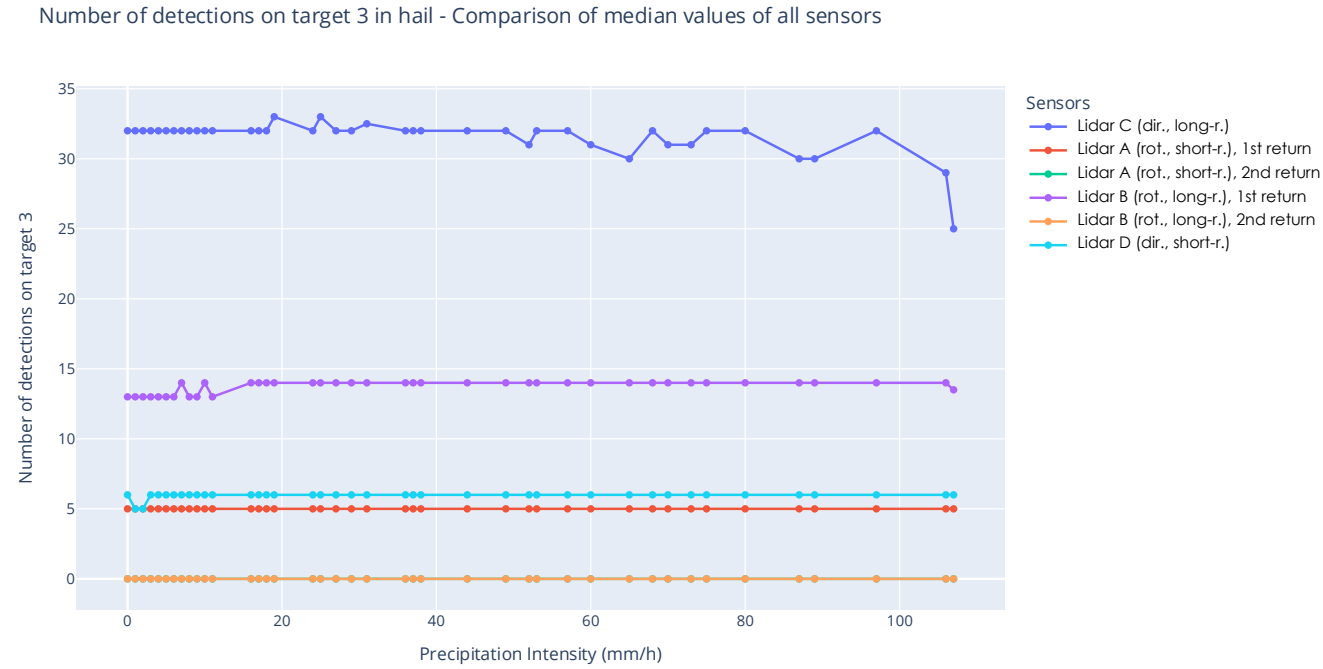

Figure 92: Number of detections on target 3 in hail

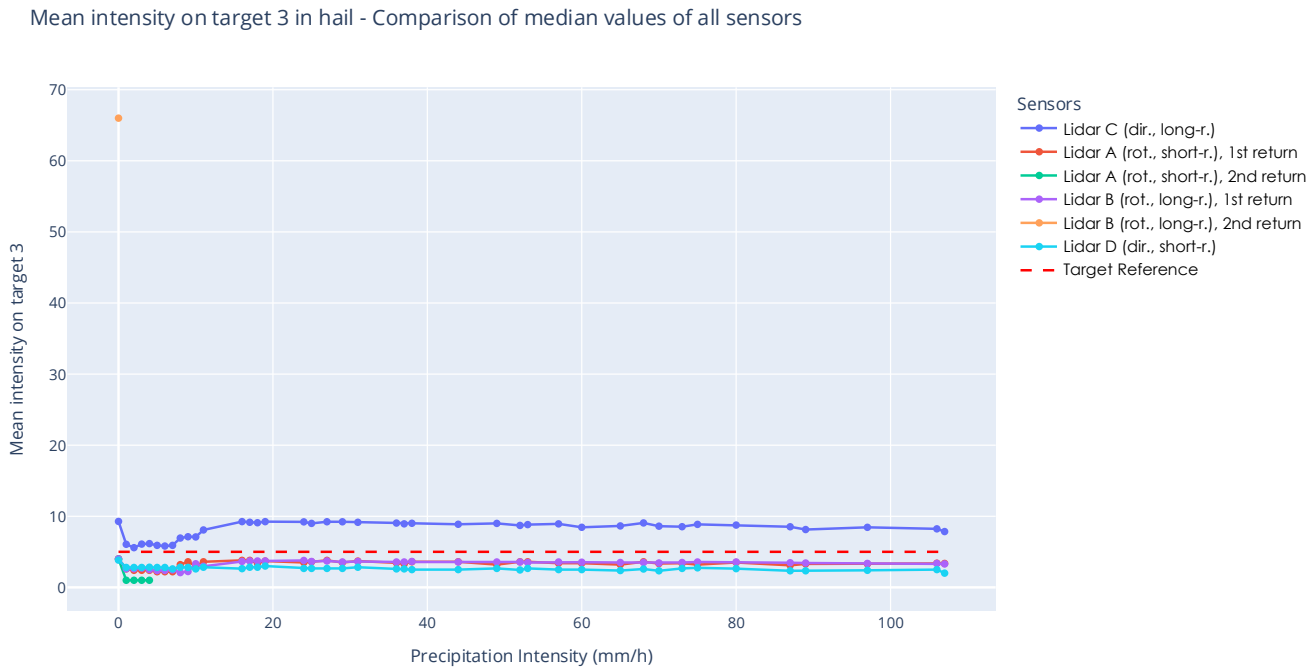

Figure 93: Mean intensity of detections on target 3 in hail

4.2.4.5 Lidar detections on Target 4 (50 %) in hail

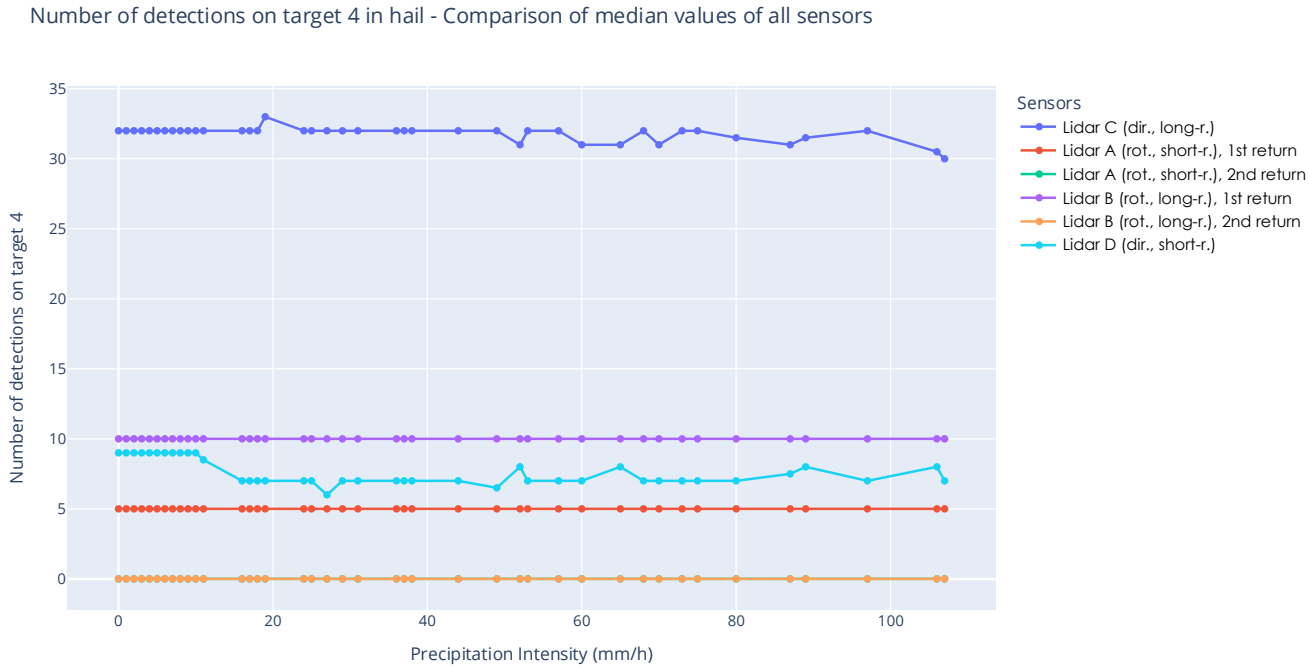

Figure 94: Number of detections on target 4 in hail

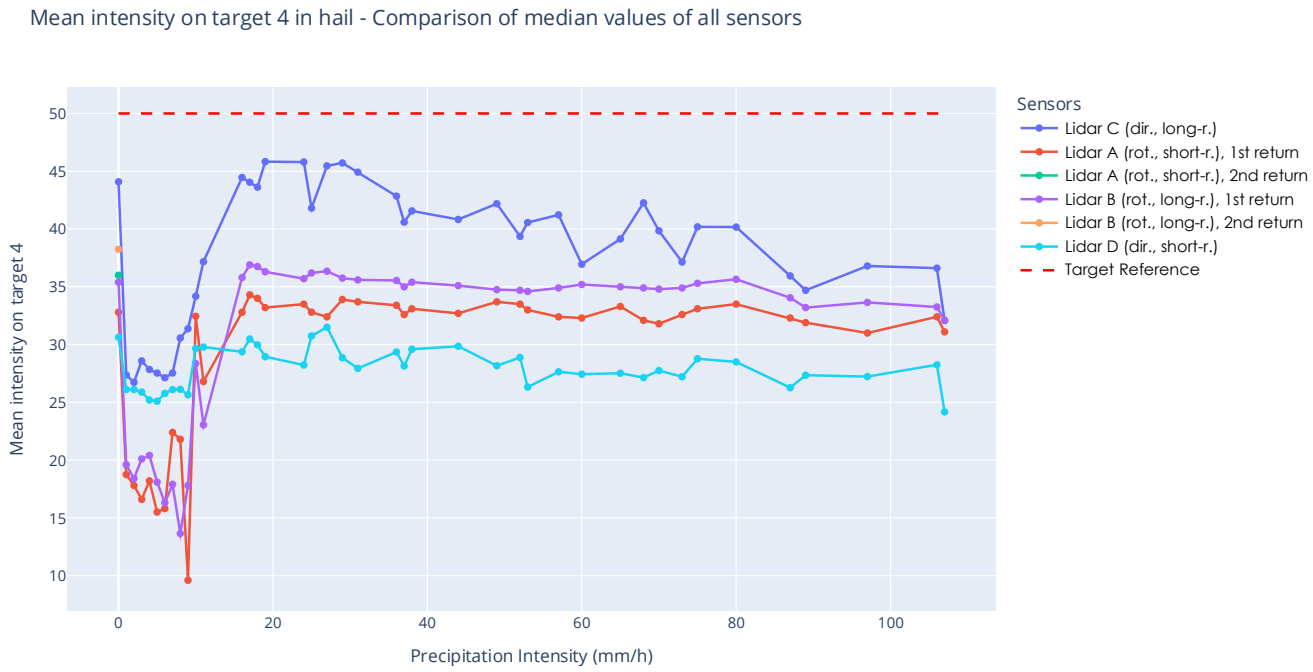

Figure 95: Mean intensity of detections on target 4 in hail

4.2.4.6 Lidar detections on Target 5 (80 %) in hail

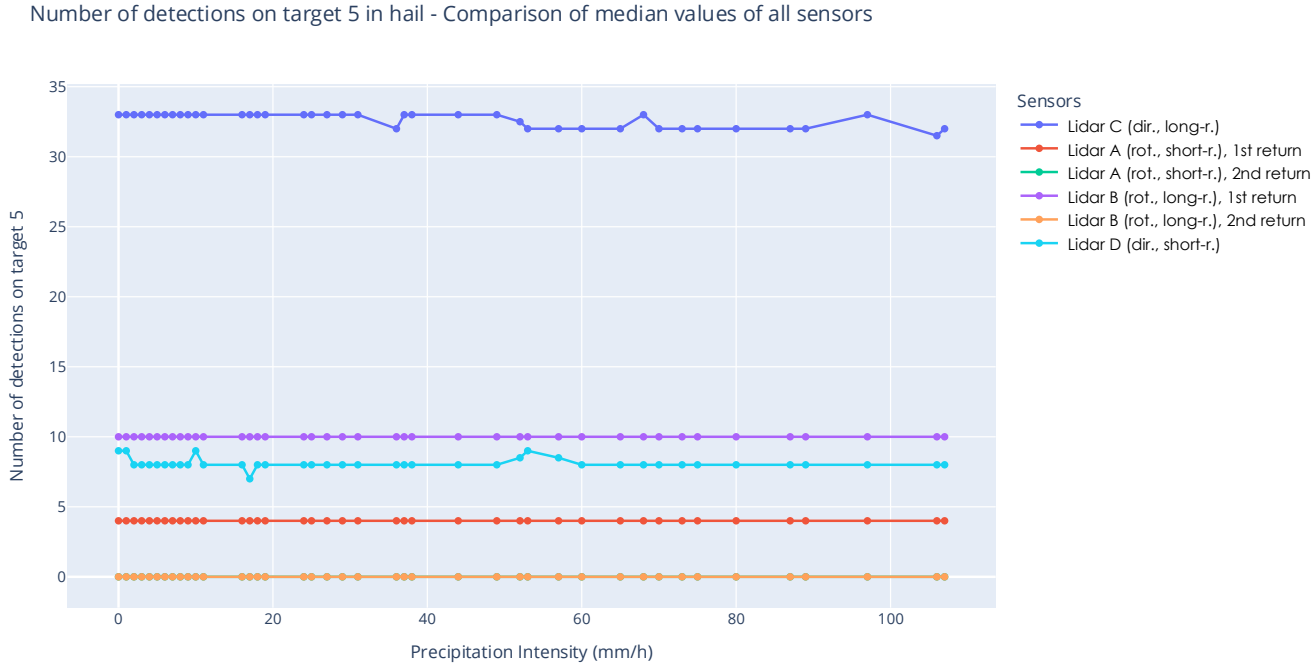

Figure 96: Number of detections on target 5 in hail

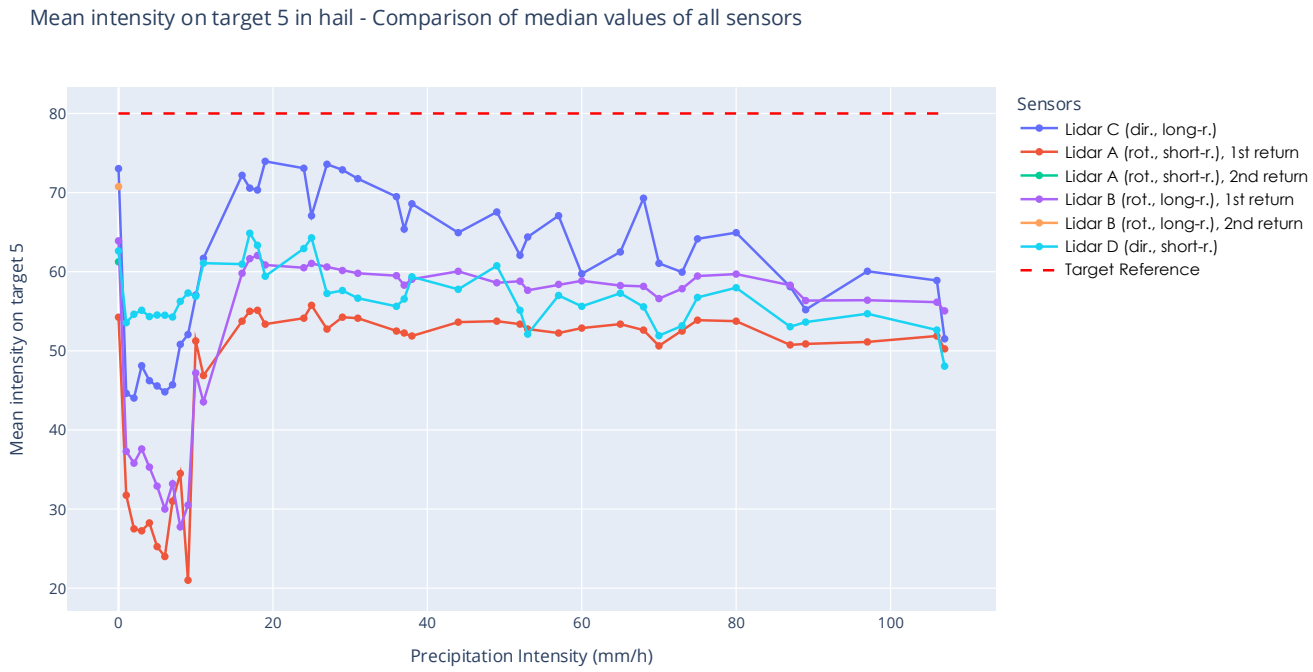

Figure 97: Mean intensity of detections on target 5 in hail

4.2.4.7 Interim conclusion for lidar performance in hail

In hail conditions, the probability of atmospheric detections increases across all tested sensors once precipitation intensity exceeds roughly 10 mm/h. This effect becomes most pronounced in the long-range rotating sensor (lidar B), which reaches a maximum atmospheric detection probability of approximately 0.3 during hailstorms with intensities above 100 mm/h. The short-range rotating lidar (lidar a) shows a lower peak at around 0.15, while both directional lidars (C and D) exhibit a slightly smaller impact, with a maximum around 0.1. These values indicate that larger or more numerous hail particles increase the likelihood of first-return detections within the atmosphere, though the magnitude of the effect remains more moderate compared to fog.

Despite these atmospheric returns, detections on the far target remain largely stable for all sensors throughout the hail scenarios. The only significant effect observed at long range is a reduction in measured intensity for the rotating sensors (A and B), suggesting that hail introduces additional attenuation or scattering that weakens the return signal without disrupting range-level detection.

Detection performance on all near-range targets also remains stable across the full precipitation intensity spectrum. The measured intensity on these close targets initially decreases between 0 and 20 mm/h, likely due to scattering effects from smaller hail particles present at lower precipitation intensities. As precipitation intensity increases further and the hailstone size distribution shifts, intensity values rise again and return to levels comparable to the original baseline.
Overall, hail produces atmospheric noise but does not meaningfully impair detection range capabilities at either near or far distances. The primary measurable impact is a temporary fluctuation in intensity values, particularly in the rotating sensors (A and B), while the directional lidars (C and D) show a smaller sensitivity to hail-related backscatter.

4.2.5 LIDAR PERFORMANCE IN RAIN

Around 99600 measurement time steps were taken during rainy conditions. Rain is quantified in rain intensity (mm/h). The distribution of rain intensities in all rain measurements is shown in the following figure.

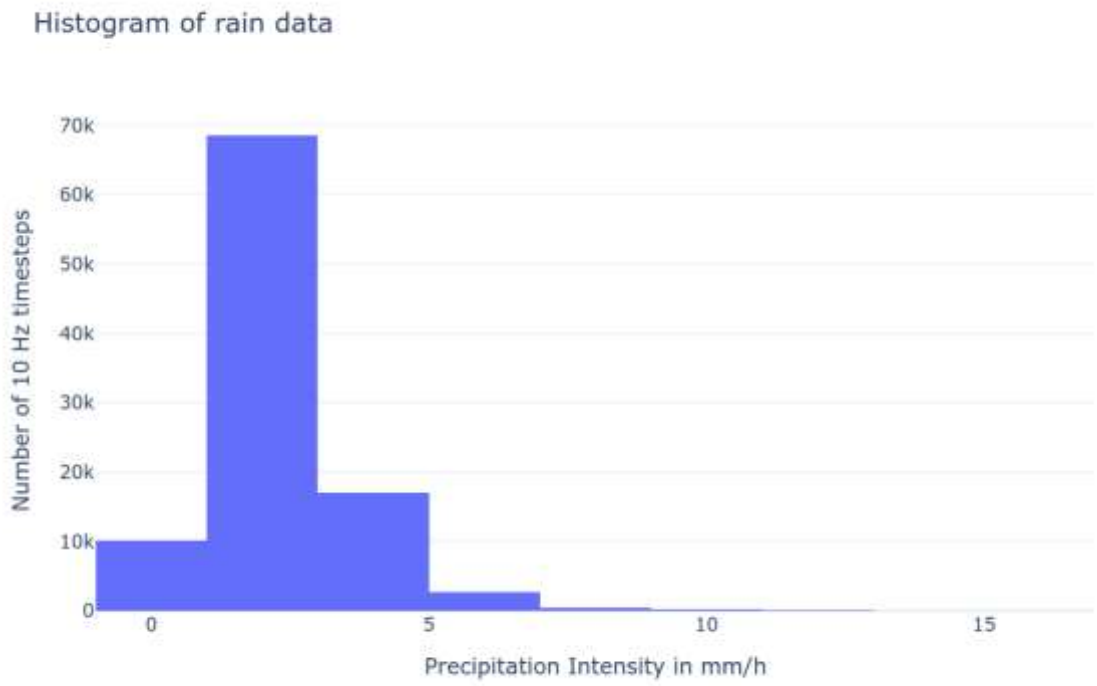

Figure 98: Histogram of rain data for lidar evaluation

As described in 2.5, the following evaluation will be structured regarding the filter regions: atmosphere and targets 1 - 5. For each filter region, the (relative) number of detections as well as the intensity (or reflectivity) values are plotted over the precipitation intensity of the respective time frame. Each plot represents the median value of the respective sensor with the corresponding precipitation intensity over all rain measurements.
The plots in this chapter are median values aggregated from all sensors for comparison.

4.2.5.1 Lidar detections in atmosphere in rain

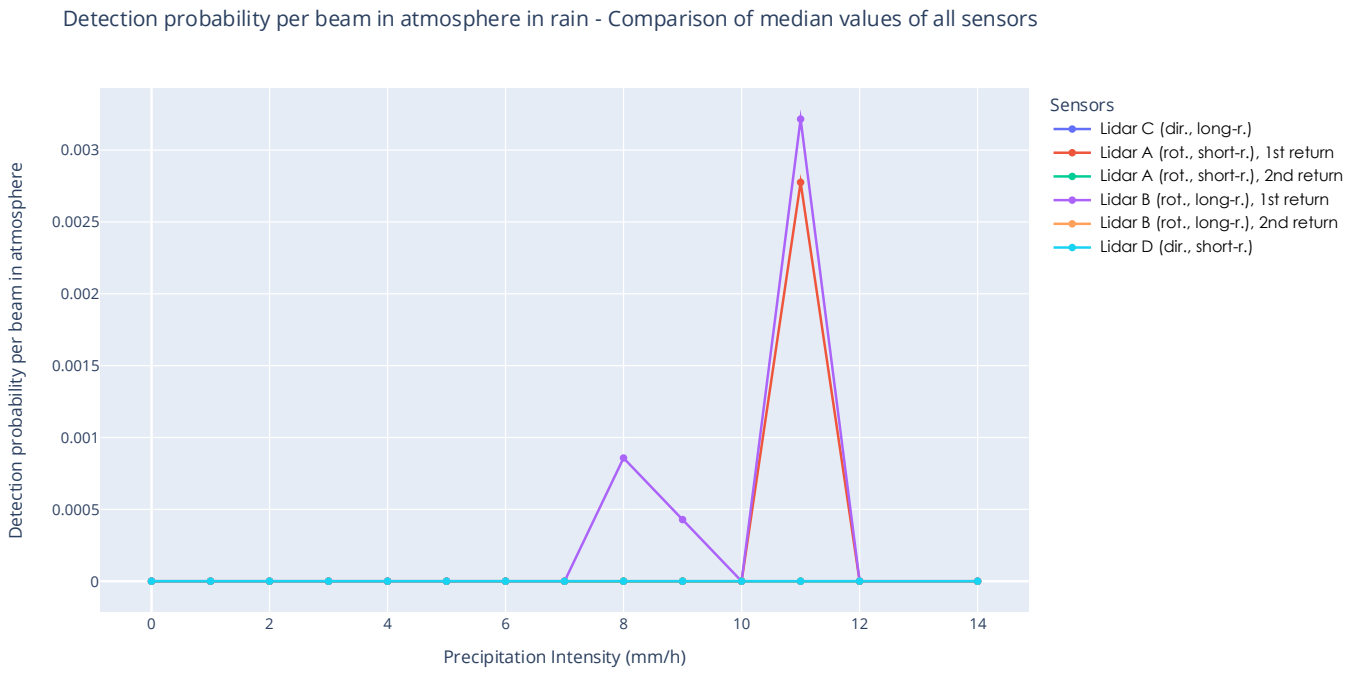

Figure 99: Detection probability per beam in atmosphere in rain

4.2.5.2 Lidar detections on Target 1 (10 %) in rain

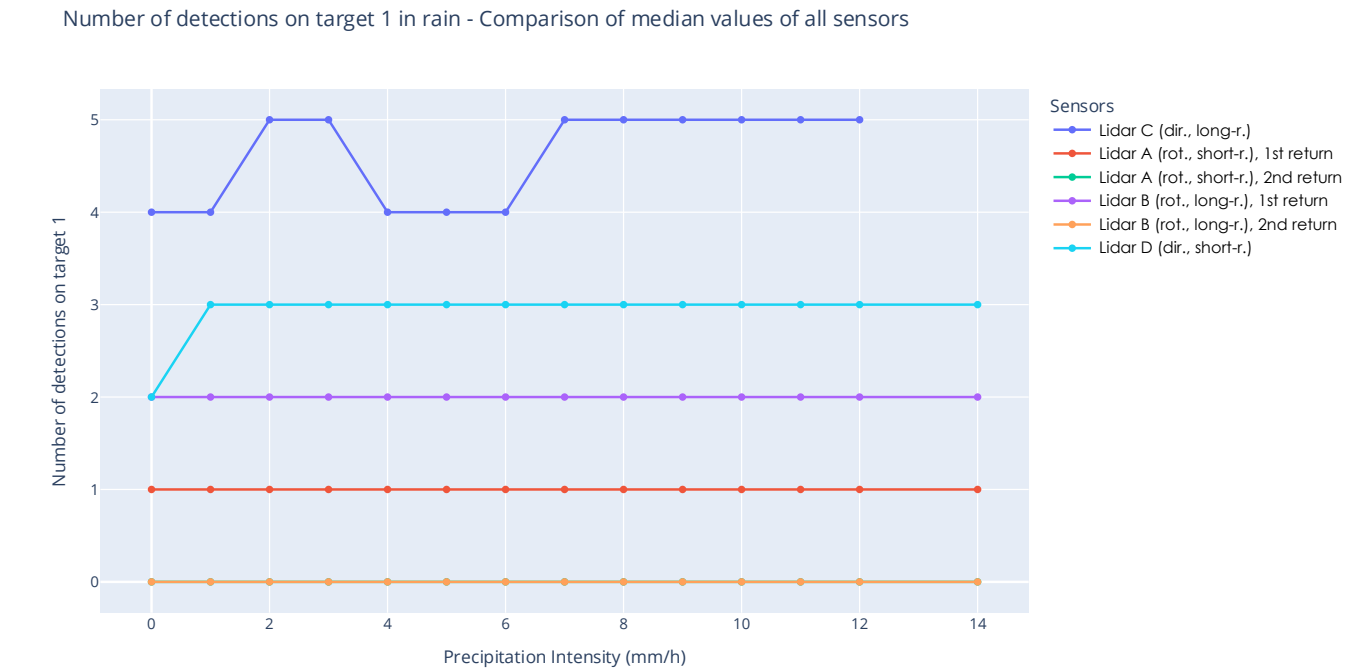

Figure 100: Number of detections on target 1 in rain

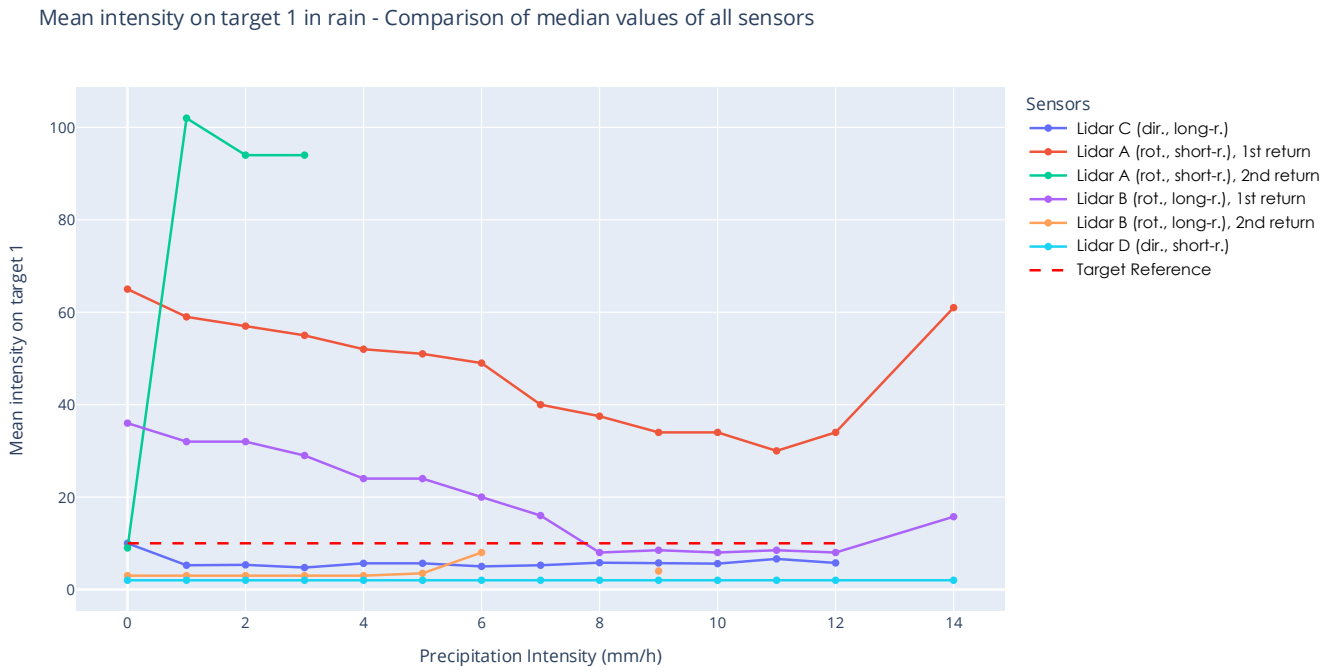

Figure 101: Mean intensity of detections on target 1 in rain

4.2.5.3 Lidar detections on Target 2 (95 %) in rain

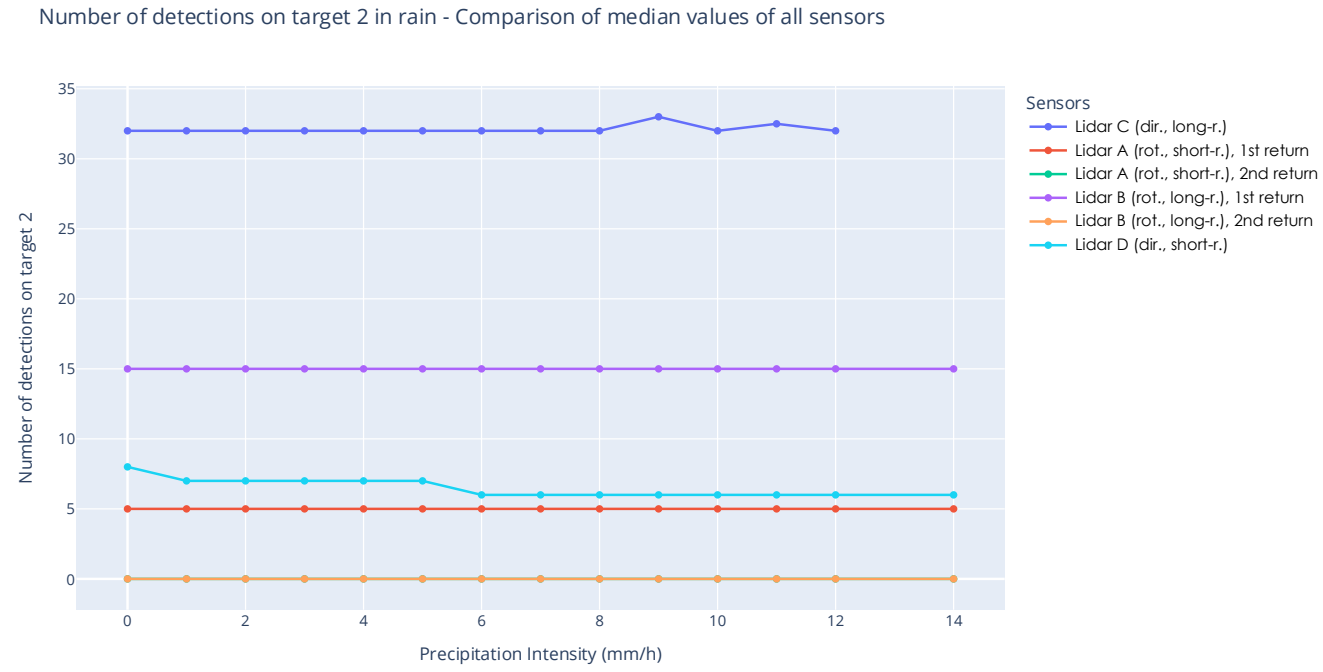

Figure 102: Number of detections on target 2 in rain

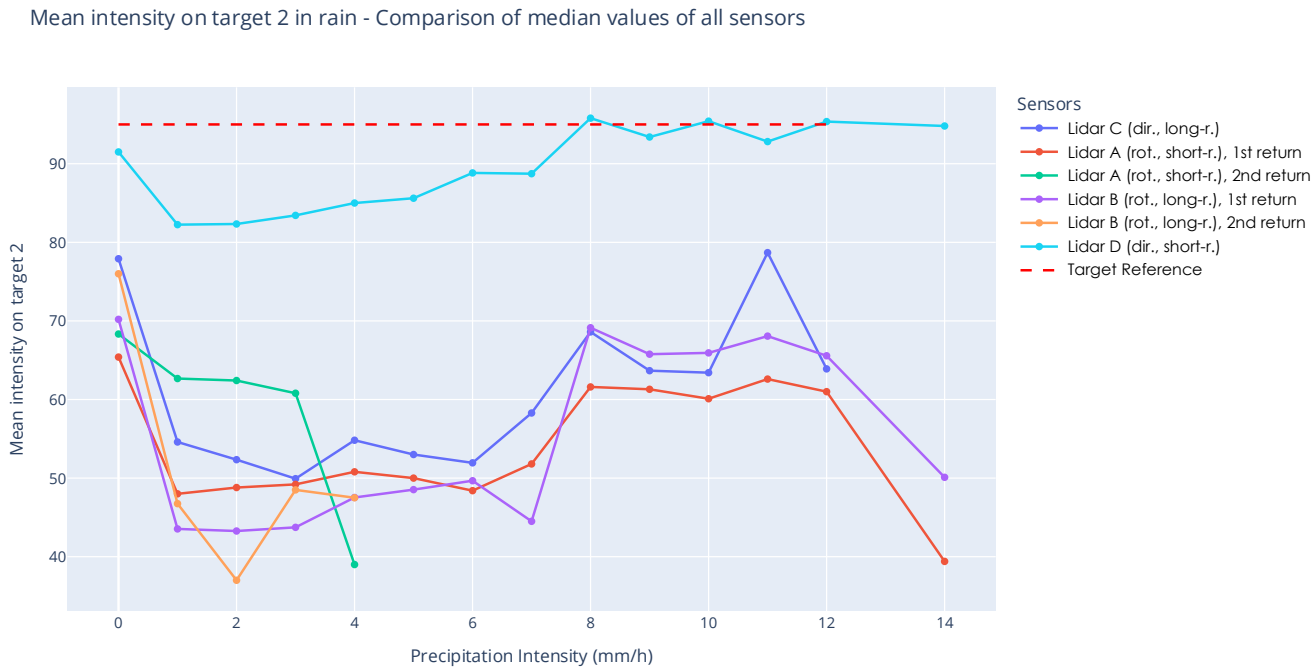

Figure 103: Mean intensity of detections on target 2 in rain

4.2.5.4 Lidar detections on Target 3 (5 %) in rain

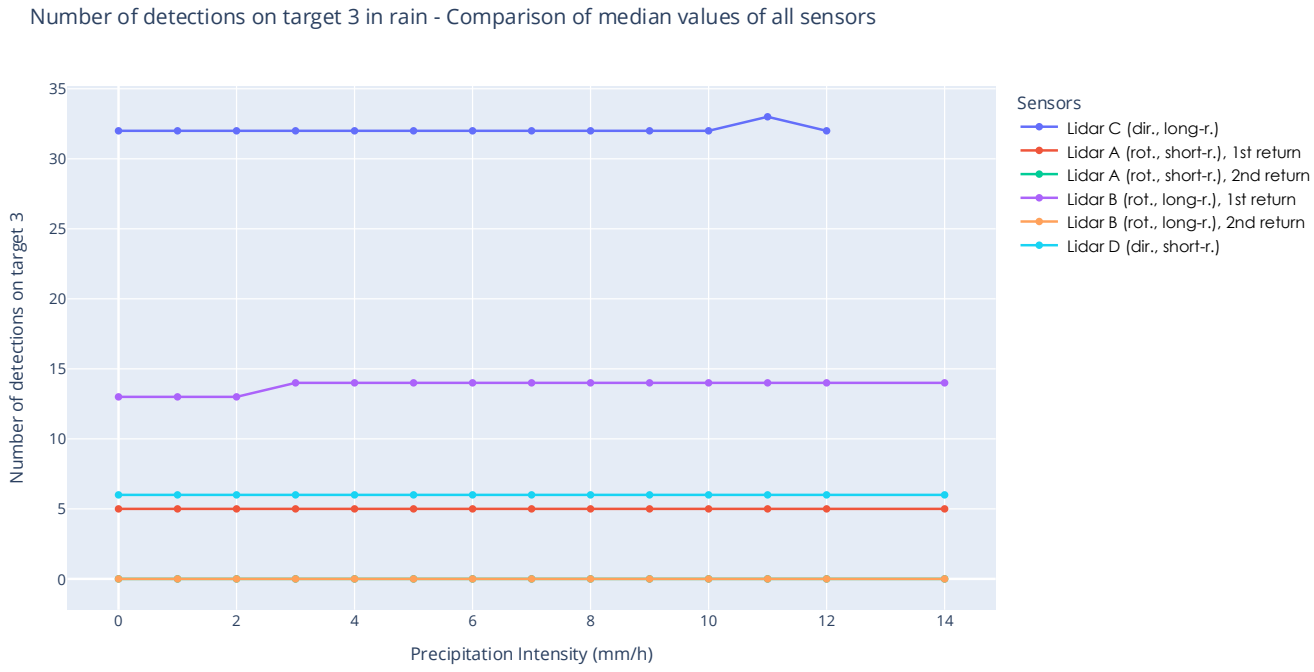

Figure 104: Number of detections on target 3 in rain

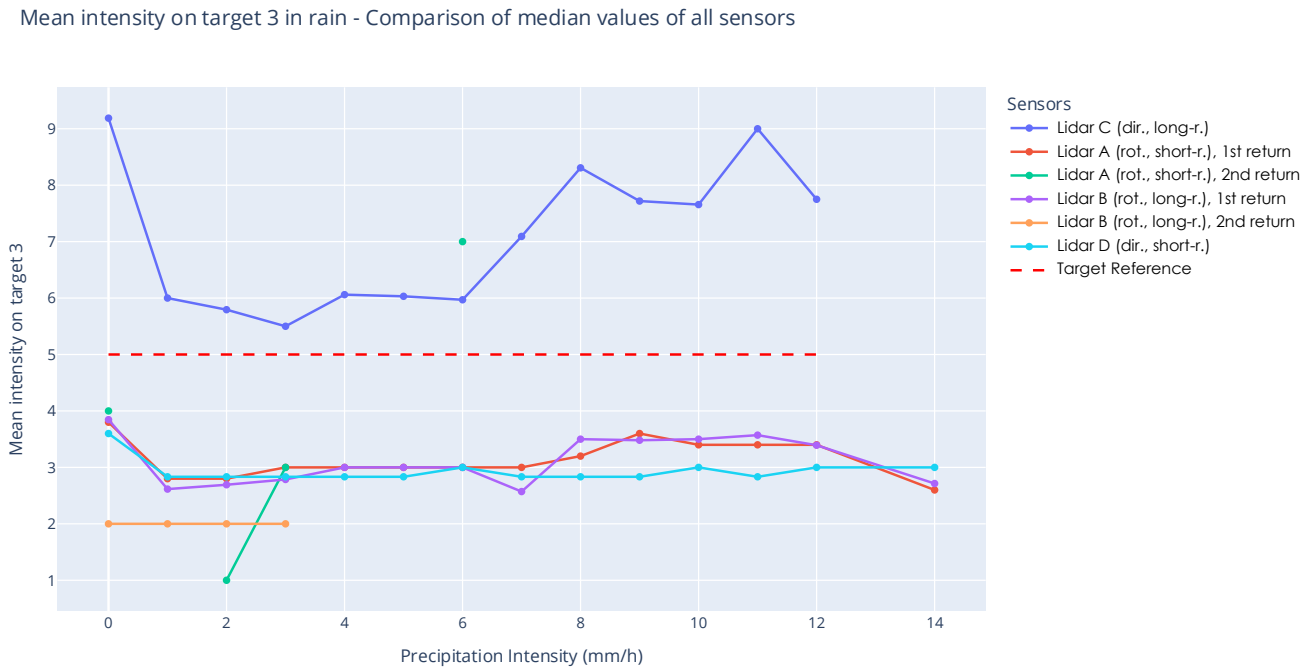

Figure 105: Mean intensity of detections on target 3 in rain

4.2.5.5 Lidar detections on Target 4 (50 %) in rain

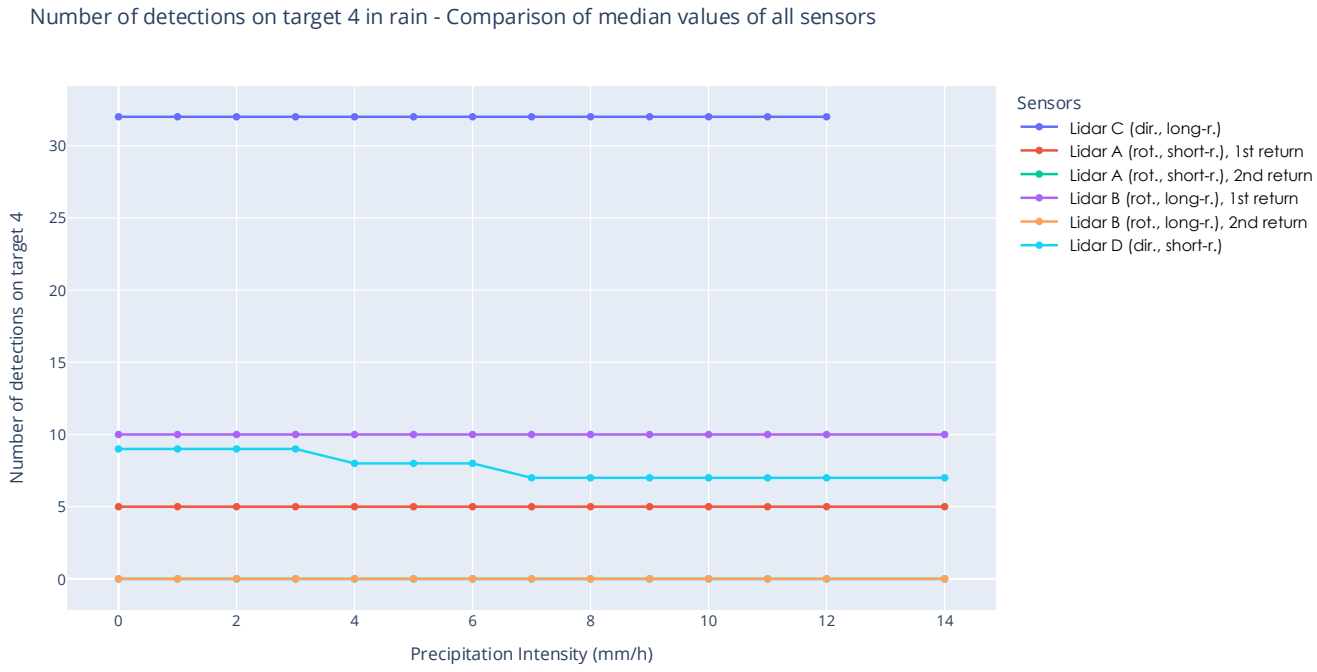

Figure 106: Number of detections on target 4 in rain

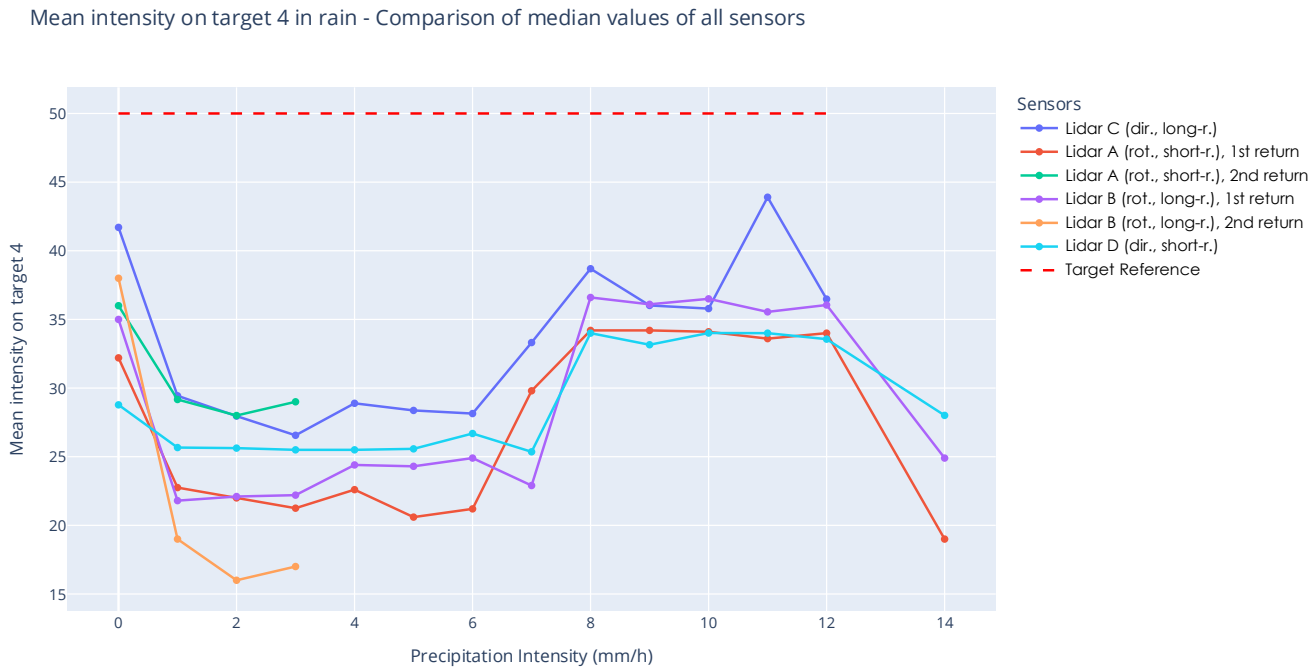

Figure 107: Mean intensity of detections on target 4 in rain

4.2.5.6 Lidar detections on Target 5 (80 %) in rain

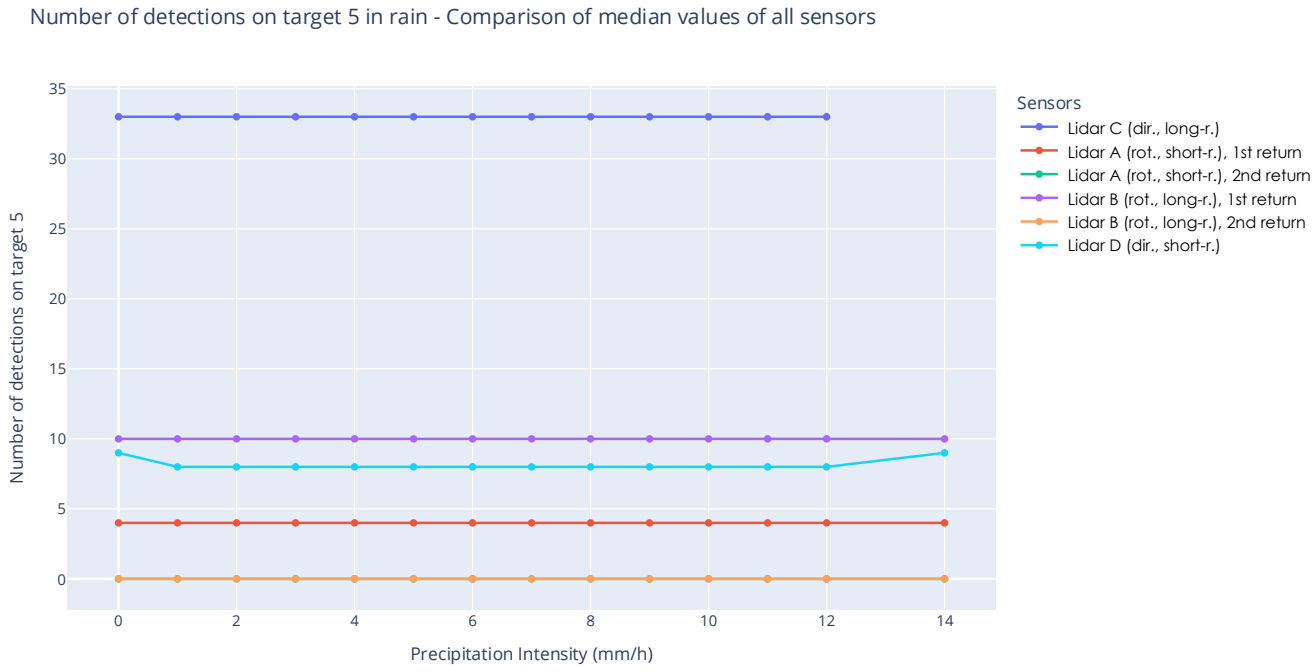

Figure 108: Number of detections on target 5 in rain

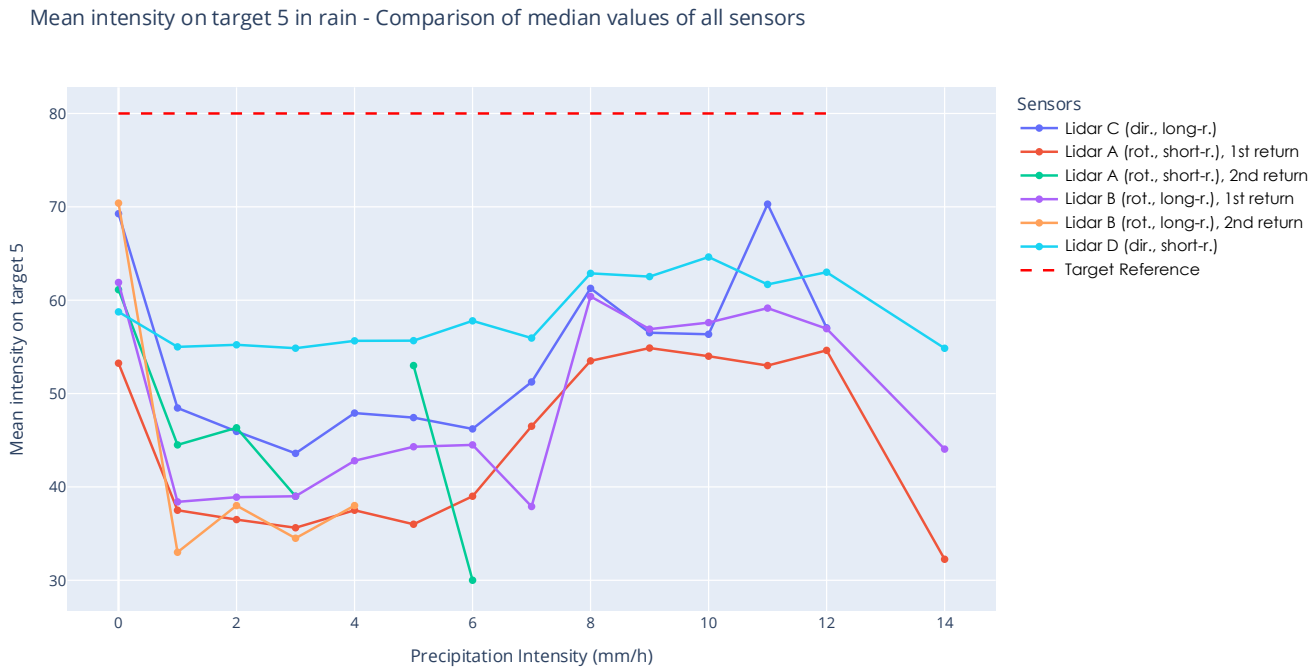

Figure 109: Mean intensity of detections on target 5 in rain

4.2.5.7 Interim conclusion on lidar performance in rain

During the measurement period, rainfall intensities did not exceed 14 mm/h, so the evaluation is limited to this moderate range of conditions. Within this interval, none of the sensors exhibited notable atmospheric detections, indicating that raindrop-induced backscatter remained negligible at these precipitation levels. Detection performance on all targets remained stable across all sensors, with no observed reduction in detection probability at either near or far range.

A slight decrease in reflectivity was observed on the far target for both rotating sensors (A and B), suggesting mild attenuation or scattering effects, although the change remained small and did not influence overall detection reliability. On the near targets, the measured intensity values fluctuated slightly for all sensors, but no consistent trend could be identified. These fluctuations remained within a range that did not affect detection behavior and are likely attributable to the stochastic distribution of raindrops in the measurement volume rather than a systematic sensor response.

Overall, rainfall up to 14 mm/h introduces only minimal optical disturbance, with stable detection performance across all sensors and only minor variations in reflectivity and intensity measurements.

4.2.6 LIDAR PERFORMANCE IN SNOW

Around 145200 measurement time steps were taken in Snow. The snowfall is quantified with the precipitation intensity in mm/h. The distribution of precipitation intensity in all snow measurements is shown in the following figure.

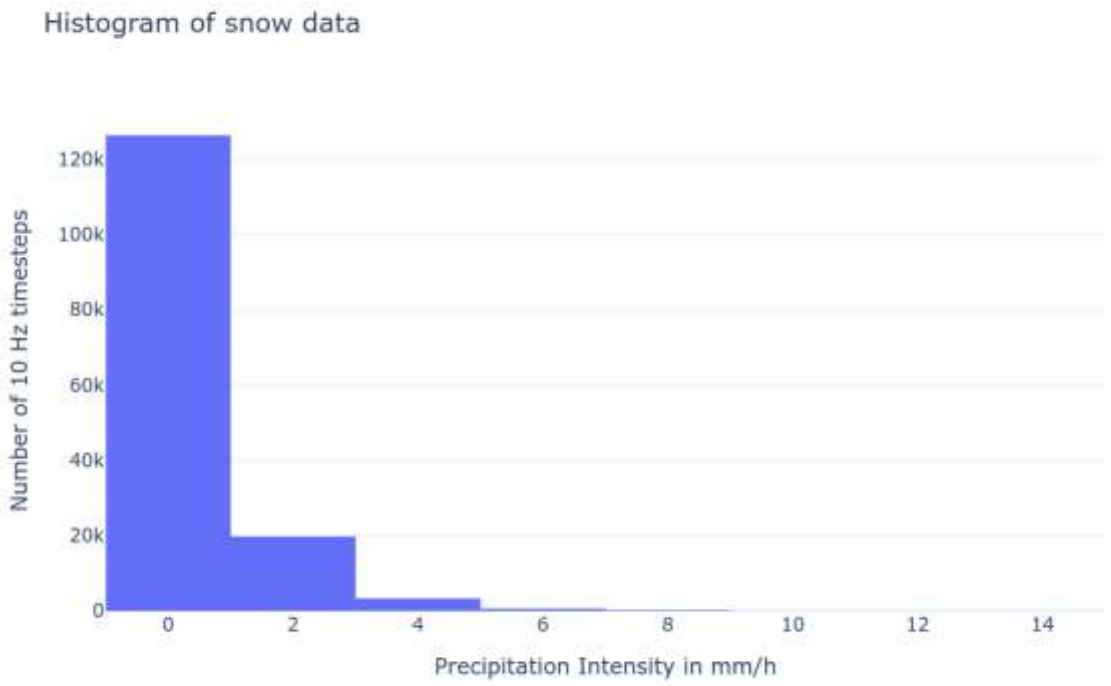

Figure 110: Histogram of snow data for lidar evaluation

As described in 2.5, the following evaluation will be structured regarding the filter regions: atmosphere and targets 1 - 5. For each filter region, the (relative) number of detections as well as the intensity (or reflectivity) values are plotted over the precipitation intensity of the respective time frame. Each plot represents the median value of the respective sensor with the corresponding precipitation intensity over all snow measurements.
The plots in this chapter are median values aggregated from all sensors for comparison.

4.2.6.1 Lidar detections in atmosphere in snow

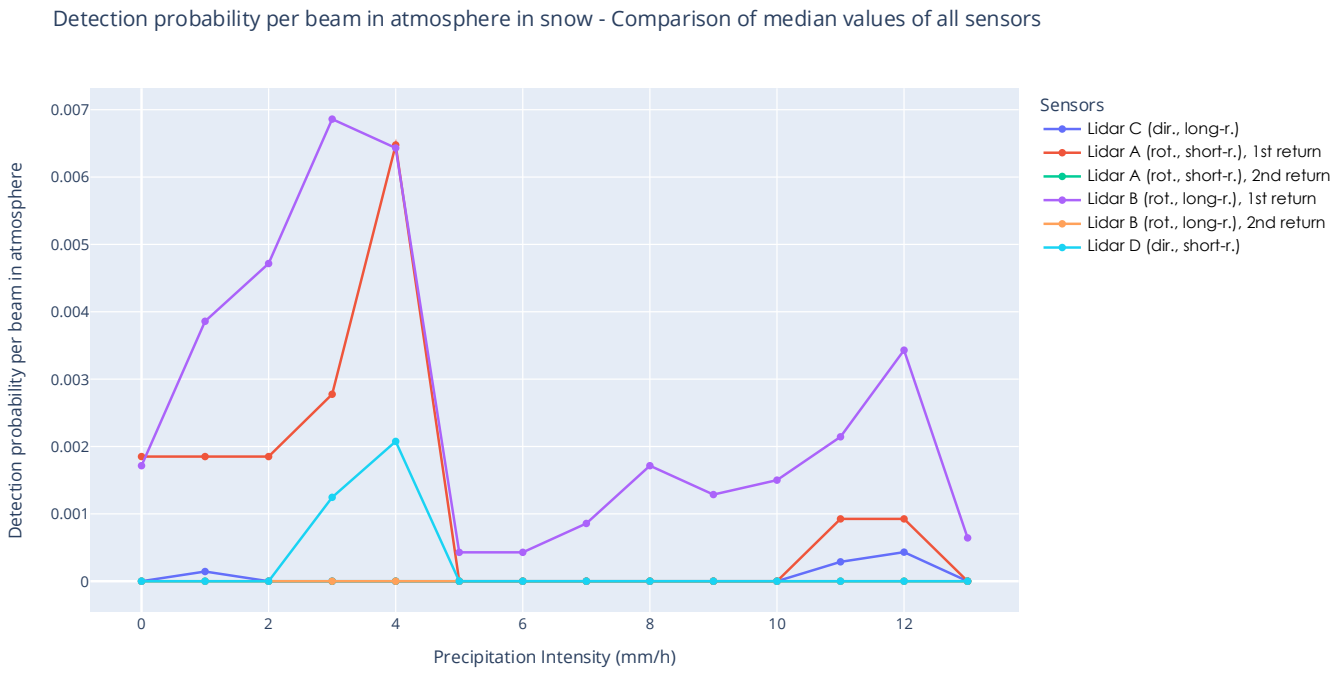

Figure 111: Detection probability per beam in atmosphere in snow

4.2.6.2 Lidar detections on Target 1 (10 %) in snow

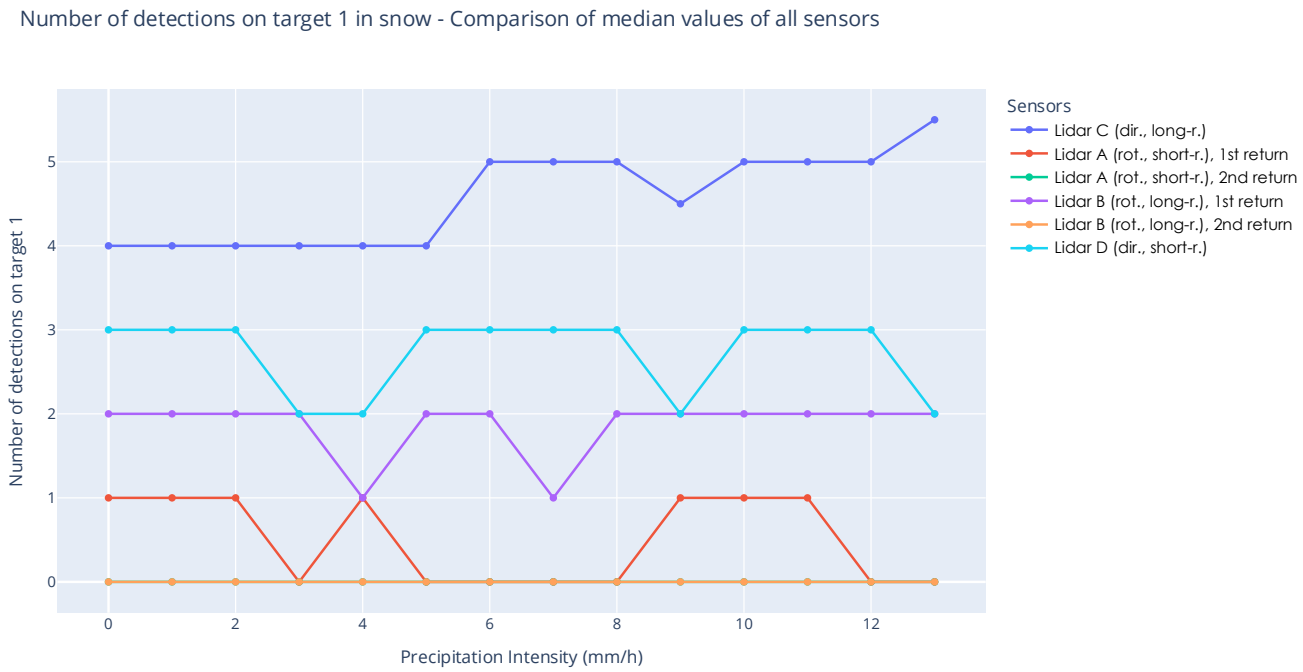

Figure 112: Number of detections on target 1 in snow



Figure 113: Mean intensity of detections on target 1 in snow

4.2.6.3 Lidar detections on Target 2 (95 %) in snow

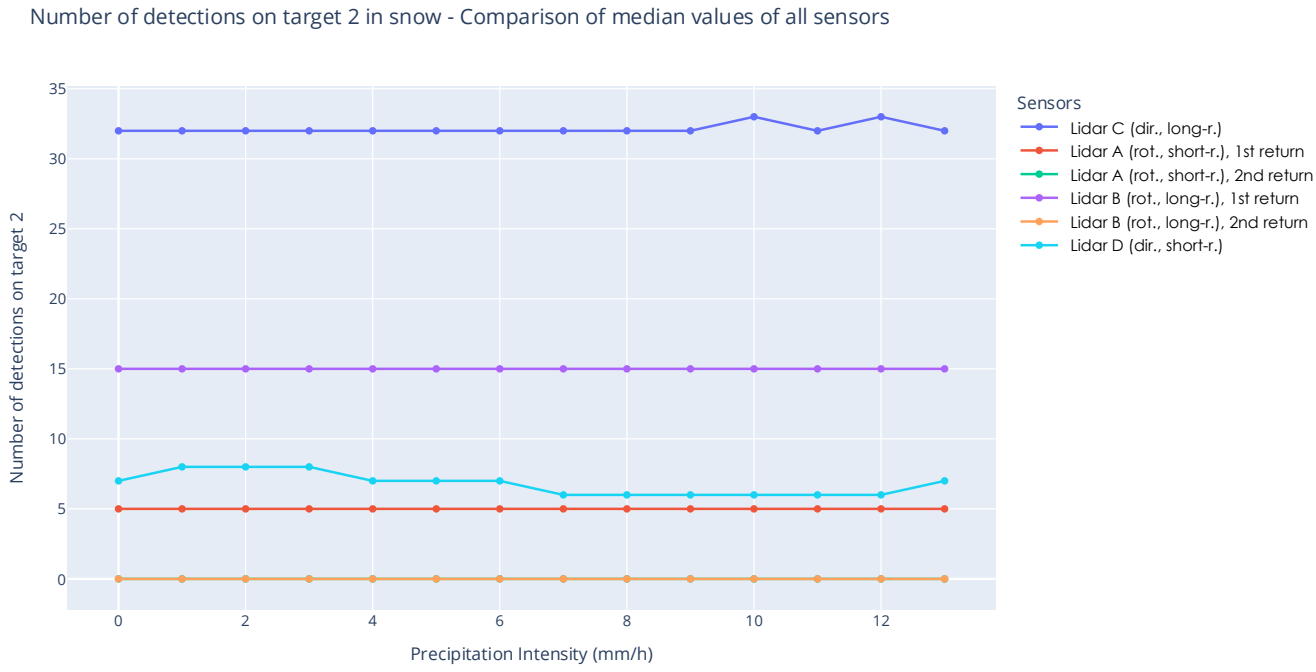

Figure 114: Number of detections on target 2 in snow

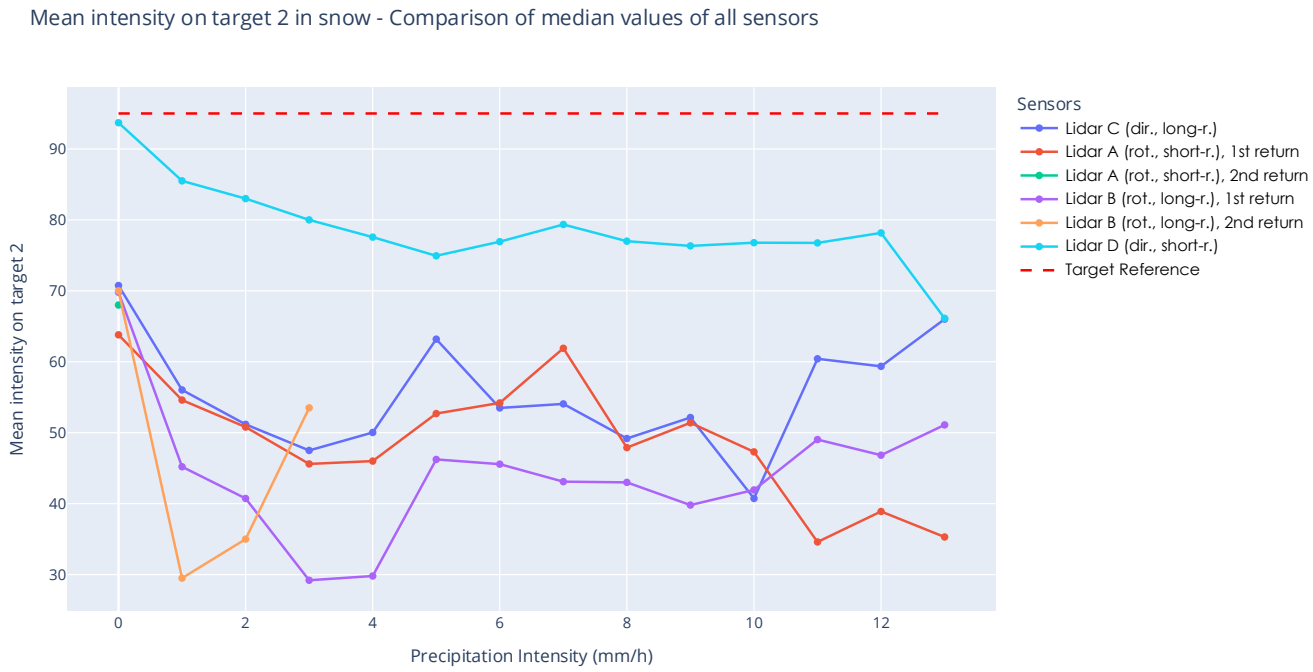

Figure 115: Mean intensity of detections on target 2 in snow

4.2.6.4 Lidar detections on Target 3 (5 %) in snow

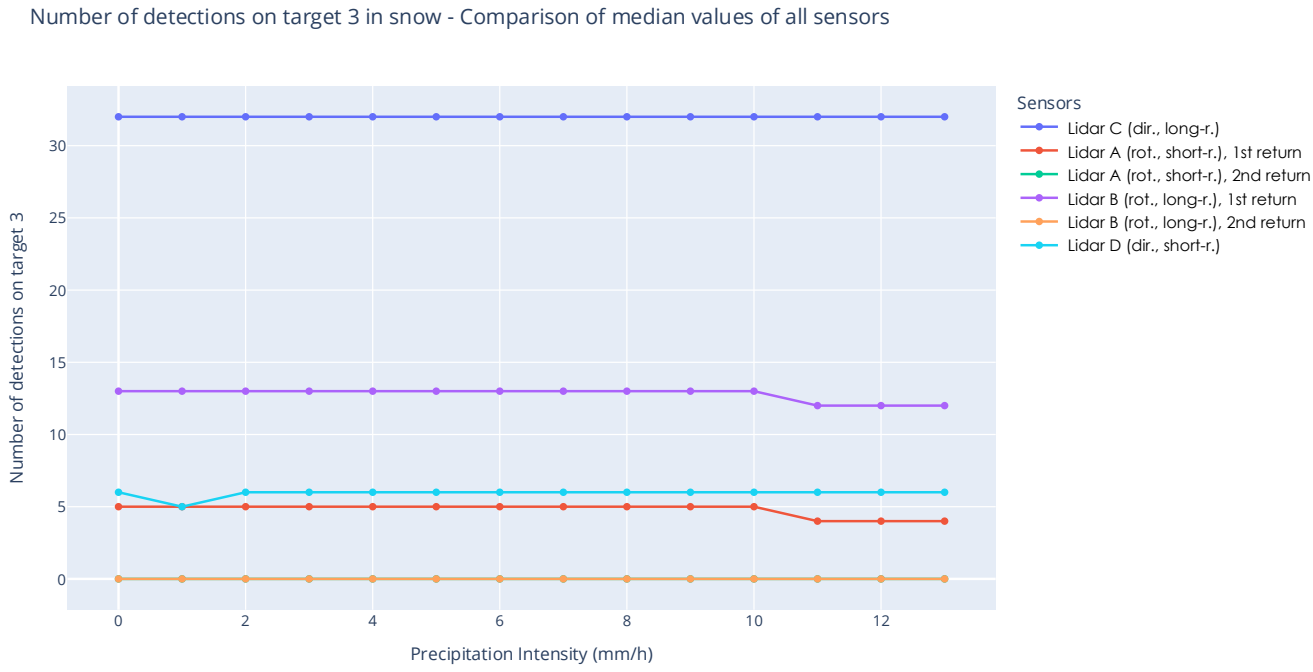

Figure 116: Number of detections on target 3 in snow



Figure 117: Mean intensity of detections on target 3 in snow

4.2.6.5 Lidar detections on Target 4 (50 %) in snow

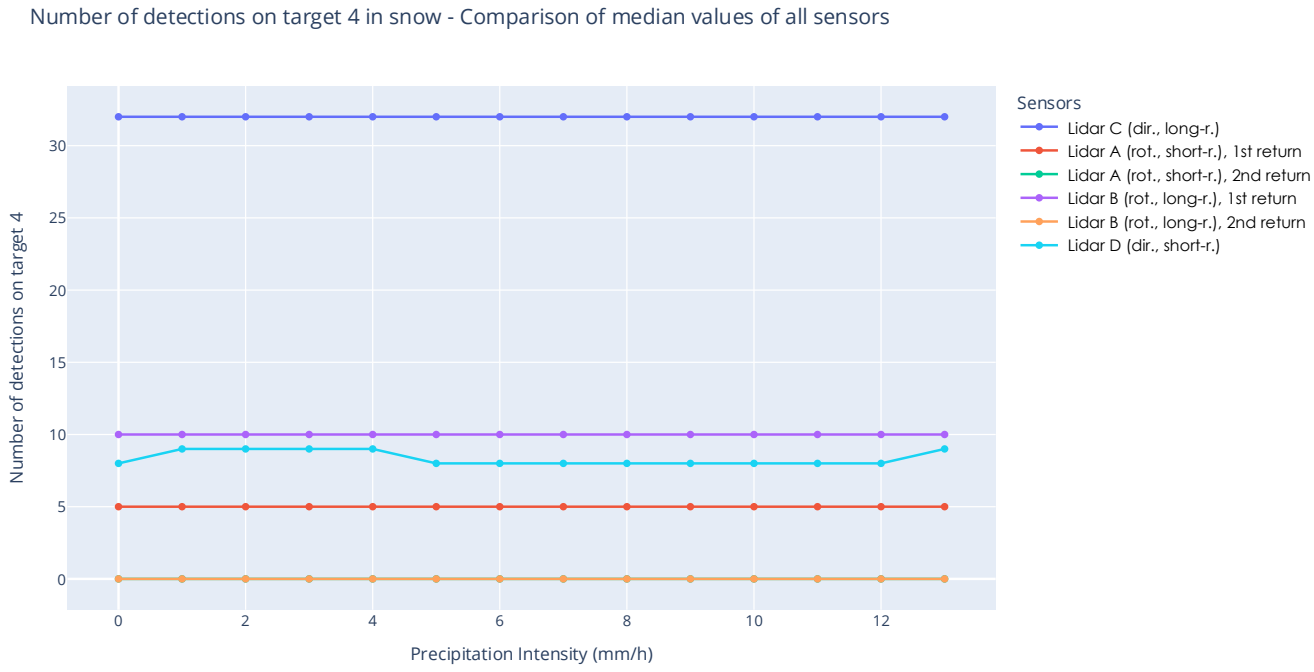

Figure 118: Number of detections on target 4 in snow



Figure 119: Mean intensity of detections on target 4 in snow

4.2.6.6 Lidar detections on Target 5 (80 %) in snow

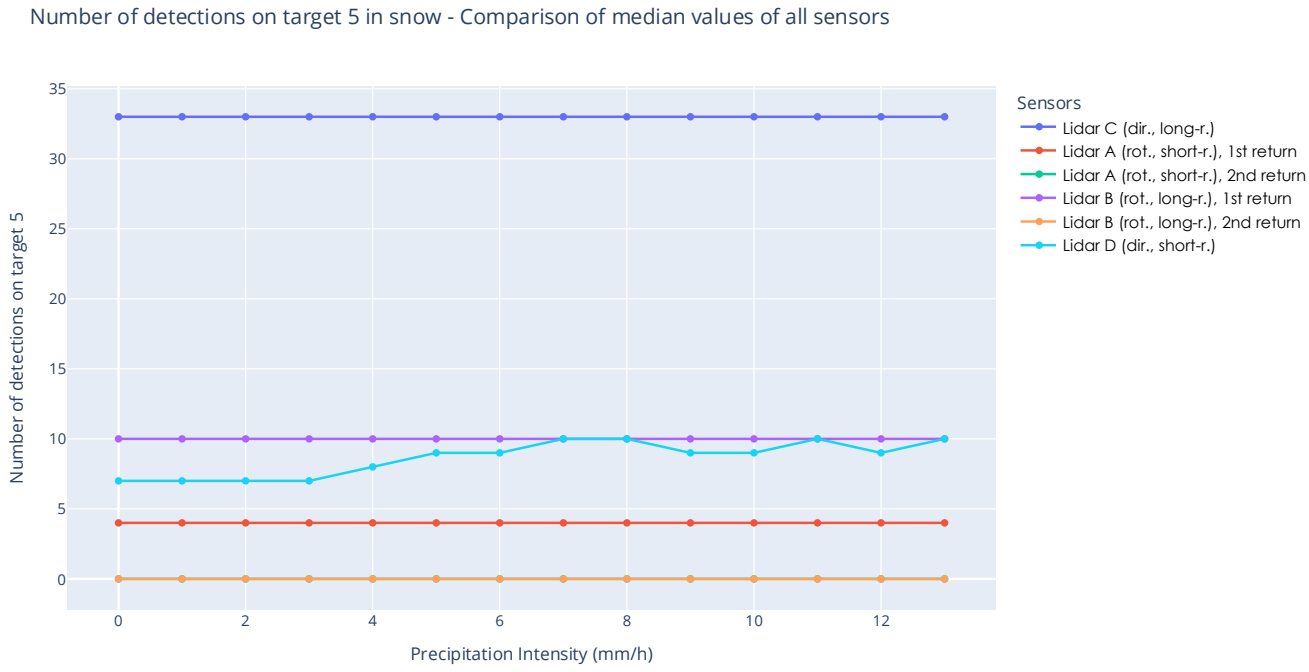

Figure 120: Number of detections on target 5 in snow



Figure 121: Mean intensity of detections on target 5 in snow

4.2.6.7 Interim conclusion for lidar performance in snow

In snowfall conditions with precipitation intensities up to 5 mm/h, the rotating sensors (A and B) exhibit only a very small probability of atmospheric detections, indicating minimal backscatter within this low-intensity snow range. Detection performance on the far target remains largely stable, with the short-range rotating lidar (A) showing occasional drops in detection probability, though these fluctuations do not follow a consistent trend and appear to vary with the irregular nature of the snowfall. As in hail and rain, measured intensities on the far target fluctuate for both rotating sensors (A and B).

All sensors maintain stable detection performance on the close-range targets, with no impact from snowfall up to the tested precipitation level. Intensity measurements on these near targets show fluctuations across all sensors, similar to the behavior observed in rainfall, but without any systematic decrease or increase. Because snowfall produces a wider variation in particle size, density, and fall velocity compared to rain, measuring precipitation intensity in snow is inherently less reliable, and corresponding fluctuations in the evaluation are therefore expected.

Overall, snow at intensities up to 5 mm/h introduces only minor and inconsistent variations in intensity and occasional detection drops for the short-range rotating lidar (A) at long range, while all sensors preserve stable detection capabilities, particularly on close targets.

4.2.7 CONCLUSION FOR LIDAR PERFORMANCE IN DIFFERENT WEATHER CONDITIONS

Across the evaluated environmental conditions, the lidar measurements show that performance degradation is strongly effect- and sensor-configuration-dependent. Fog represents the dominant object-independent visibility limiter: with increasing extinction coefficient, target return intensities decrease and detection performance on the far target drops for all sensors, while near-range high-reflectivity targets remain detectable over a wide range. In addition, fog induces a substantial amount of atmospheric backscatter detections ("noise points") for the rotating sensors, especially the short-range rotating lidar, whereas the long-range directional lidar remains comparatively robust in homogeneous fog. Importantly, the special case of patchy/uneven fog highlights that directional sensors may generate clustered atmospheric detections in denser fog pockets even when they appear stable in homogeneous conditions.

For precipitation conditions within the observed intensity ranges, the primary impact is the appearance of stochastic atmospheric detections and moderate fluctuations in reported intensity/reflectivity rather than a systematic loss of target detectability. Hail produces measurable atmospheric detections across all sensors above roughly 10 mm/h, with the largest effect on the long-range rotating sensor; however, target detections remain largely stable and only the reported intensity shows transient changes. For rain (up to 14 mm/h in the available data) and snow (up to 5 mm/h), atmospheric detection probabilities remain low and detection performance on near and far targets is largely stable, with only minor, inconsistent intensity variations and occasional long-range fluctuations on the short-range rotating sensor. Under varying brightness, no atmospheric detections occur; detection performance is mostly unaffected at close range, while some sensors show reduced robustness on the far target and sensor-specific sensitivity of intensity/reflectivity estimates to ambient illumination.

For simulation and model validation, these results support a two-part modeling strategy for object-independent effects: (i) a physically motivated attenuation component parameterized by the forward scatter extinction coefficient (or MOR) to capture range-dependent detection loss, and (ii) a stochastic atmospheric-return component whose occurrence probability, range distribution, and intensity statistics depend on weather severity and sensor configuration (rotating vs. directional, FoV, wavelength, and processing). For rain, snow, and hail, precipitation rate in mm/h should be treated as a proxy with limited transferability because particle size distributions and spatial inhomogeneity can dominate the observed scatter; consequently, models should emphasize variance/outlier behavior and avoid extrapolation beyond the measured intensity ranges. Finally, the strong differences between sensor types and the sensitivity to non-homogeneous fog underline that credibility assessment should include configuration-specific calibration and, where possible, dedicated tests for spatially varying conditions (e.g., patchy fog) in addition to aggregated performance curves.

5. Simulation of lidar in bad weather

5.1 Object-dependent and -independent environmental effects

Environmental influences on automotive lidar sensors can be systematically classified into object-independent and object-dependent effects. This distinction, introduced and formalized in the context of performance-oriented perception sensor simulation, provides a structured basis for measurement, modeling, and validation of sensor behavior under adverse conditions. (Linnhoff, 2023)

Object-independent effects describe environmental influences that affect the sensor globally, largely independent of specific objects in the scene. Typical examples include fog, rain, snow, and solar radiation, which alter the propagation of optical or electromagnetic signals before they interact with objects. These effects are primarily governed by atmospheric conditions and can be characterized by using external reference sensors such as visibility meters, rain gauges, or irradiance sensors. Due to the available data from the measurements campaign at the Persival Weather Test Bench described earlier, this report mainly focusses on object-independent effects.

For active perception sensors like lidars, object-independent effects manifest as atmospheric attenuation, scattering, and increased background detections. In fog or snowfall, particles in the atmosphere cause signal energy to be scattered back toward the sensor, leading to detections in free space and a reduction of effective detection range. In rain, both attenuation and stochastic backscatter contribute to increased noise and uncertainty. Camera sensors are affected through contrast loss, reduced visibility, glare, and reduced signal-to-noise ratio, impacting downstream perception tasks even when scene geometry remains unchanged.

A key characteristic of object-independent effects is their statistical homogeneity across the sensor's field of view. As shown in long-term stationary measurement campaigns, sensor responses can be aggregated as a function of quantified environmental parameters (e.g. rain rate, visibility range), yielding expectation values such as mean detection density, range distribution shifts, or noise levels. These expectation values form the basis for stochastic and performance-driven sensor models that can be applied consistently across different scenes in simulation.

In contrast, object-dependent effects arise from the interaction between environmental conditions and specific objects or surfaces in the scene. Prominent examples include wet or icy road surfaces, road spray generated by moving vehicles, contamination of sensor windows, and reflections from wet infrastructure. Unlike object-independent effects, these phenomena are highly localized and strongly influenced by object geometry, material properties, relative motion, and vehicle dynamics.

In the case of lidar sensors, object-dependent effects such as road spray lead to characteristic clusters of detections with spatial and temporal structure. Experimental studies show that spray plumes generated by vehicles driving on wet pavement produce dense, turbulent detection clusters with varying detection probability and spatial extent. These effects cannot be described adequately by global noise or attenuation models, as they depend explicitly on object speed, distance, water film thickness, and relative sensor position.
For camera sensors, object-dependent effects include specular reflections on wet roads, partial occlusions by spray or dirt, and rapidly changing local contrast. These effects directly

alter the appearance of specific scene elements and often have a disproportionate impact on perception algorithms, even if global visibility remains acceptable.
Implications for Measurement, Modeling, and Simulation

The distinction between object-independent and object-dependent effects has important implications for simulation. Object-independent effects can be efficiently represented using global, parameter-driven stochastic models derived from aggregated measurements, the already shown performance curves. Object-dependent effects, in contrast, require context-aware and object-centric modeling approaches, where environmental influences are explicitly linked to objects and their dynamics.

The dissertation from Linnhoff (Linnhoff, 2023) demonstrated that both effect classes can be quantified experimentally and translated into performance-oriented sensor models using expectation values derived from real measurements. This approach avoids the need for full physical modeling while preserving realism and traceability. For simulation-based safety validation, combining both effect types is essential to realistically reproduce sensor behavior under adverse weather and to avoid systematic bias in virtual testing.

Within the WeatherSafe project, this classification provides a conceptual backbone for structuring sensor evaluation, deriving simulation models, and defining validation strategies that are grounded in real sensor performance rather than purely phenomenological weather representations.

5.2 LIDAR BASE MODEL

To implement environmental influences in a traceable and modular way, the simulation model builds on a standardized perception-simulation architecture that separates (i) environment and scenario generation, (ii) sensor rendering / reflection generation, and (iii) sensor-internal processing. The underlying principle is that a driving simulator provides the ground truth description of the scene and the scenario with all static and moving objects and their poses for each time step. A sensor model then transforms this input into sensor outputs such as detections (point clouds), logical detections in the vehicle coordinate system, or object-level outputs.

A practical enabler for such modular simulation chains is the use of ASAM Open Simulation Interface (ASAM OSI) for data exchange and OSMP/FMI packaging for tool coupling. In this concept, the sensor model is packaged as a Functional Mock-up Unit (FMU) and receives its inputs as an OSI "Sensor View" container, while producing OSI "Sensor Data" outputs. This makes the sensor model deployable across different simulation tools that support OSI and FMI-based co-simulation.

Within the perception sensor model, a functional decomposition is used to separate sensor behavior into distinct stages. Conceptually, as shown in Figure 122, these stages map to (a) front-end sensing (emission, propagation, interaction, reception), (b) data extraction of detections from raw or semi-processed signals (peak detection, multi-echo logic, intensity calculation), and (c) higher-level perception stages (segmentation, clustering, tracking, object estimation). Such modularization is valuable because it allows targeted extensions without reworking the entire sensor pipeline.

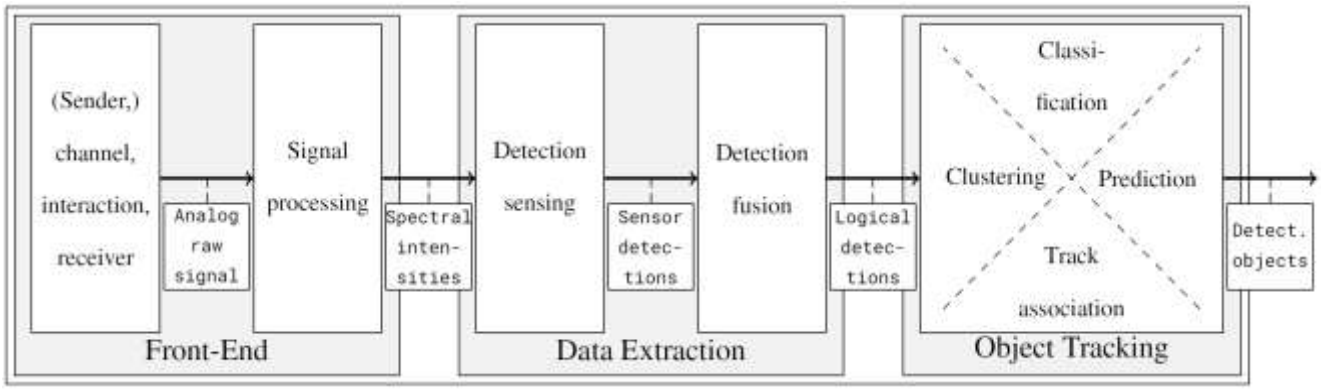

Figure 122: FZD OSMP framework based on a generic functional decomposition of automotive perception sensors (Linnhoff, 2023, p. 95)

In this framework, a reflection-based lidar detection/object model serves as the base model. It consumes reflection/ray-tracing results (oversampled per beam) and performs sensor-specific processing steps such as beam-divergence handling, multi-echo generation, peak detection, and intensity/echo-related signal processing. The resulting detections can be transformed and optionally fused, and detections can be associated with ground-truth objects to emulate segmentation and object estimation.

For integrating environmental effects, there are three generic "injection points":

1. At environment/rendering level (highest physical fidelity): Environmental particles or surfaces are explicitly represented and ray-traced.
2. At reflection level (medium fidelity): Environmental effects modify reflection results before peak detection.
3. At detection level (efficient, pragmatic): Environmental effects are applied after detections have been produced.

The approach used here selects the detection-level injection point, because it offers a good balance between computational efficiency, portability, and practical calibration with real-world measurements. It also supports a consistent strategy for adding "atmospheric" detections and attenuating/removing object detections in a way that can be calibrated to measured sensor behavior, without requiring full particle-level physics or deep access to proprietary simulator internal

5.3 MODELING OF OBJECT INDEPENDENT EFFECTS

Environmental influences are implemented as a dedicated module integrated after the base model's detection generation. The module covers both object-independent effects (fog, rain, snow, direct sunlight) and object-dependent effects (e.g., spray), using the same overarching principle:

1. Add atmospheric detections (false positives) where the environment produces returns without a hard object.
2. Attenuate or remove existing detections (false negatives) where environmental propagation reduces detectability.

To generate atmospheric detections in a controllable way, three ingredients are required:

1. Existence probability per beam: How likely a beam produces an atmospheric detection under a given condition intensity.
2. Distance distribution: Where along the beam those detections appear (range statistics).
3. Intensity distribution: What intensity values those detections carry (sensor-specific).

All three ingredients are calibrated using performance curves derived from referenced measurements (i.e., measurements described in previous chapters). The module can be parameterized for representative lidar configurations used (e.g., a short-/mid-range rotating lidar for object-independent conditions), and the same methodology can be extended to other lidar configurations once calibration data is available.

5.3.1 Implementation of fog

For fog, the key modeled phenomenon is range-dependent detectability loss, i.e., a reduction of detection probability on objects with increasing attenuation. In some lidar configurations, the measurements show that fog may produce few or no atmospheric detections within the tested visibility range; in that case the model focuses on detection drop rather than adding atmospheric returns.

A physically motivated attenuation factor can be expressed via an extinction coefficient derived from visibility. The received signal strength from a target decreases with distance and extinction, and detection occurs only if the received signal exceeds a sensor-internal detection threshold. Because the true thresholding logic is typically proprietary, the model uses a measurement-driven threshold estimation approach: detection probability on reference targets at known distances and reflectivities is used to infer an effective threshold behavior.

To avoid unrealistic "hard cutoffs" (instant drop from full detection to zero), the model introduces stochastic variability (noise) in the effective received signal / threshold comparison. This yields the empirically observed gradual decrease in detection probability over a range of visibility values. The calibration uses characteristic points such as the visibility where detection probability is ~50% (and the spread around it) to tune the stochastic component.

5.3.2 Implementation of rain and snow

For rain and snow, two effects are modeled jointly:

1. Atmospheric detections caused by backscatters from precipitation particles.
2. Attenuation that reduces detection probability on objects and targets.

5.3.2.1 Existence probability

Measurements indicate that the atmospheric detection probability often scales with precipitation intensity. A practical representation is a (piecewise) linear mapping from precipitation rate to per-beam detection probability. Importantly, many rotating multi-layer lidars show layer-dependent sensitivity, i.e., different vertical channels yield different atmospheric detection rates. The implementation therefore uses per-layer compensation factors to distribute detections realistically across elevation layers (a calibration step that is essential for matching measured angular distributions).

5.3.2.2 Distance distribution

Atmospheric detections in precipitation occur primarily in the near field but can extend further depending on intensity and sensor characteristics. The model represents the detection distances with a compact parametric distribution (e.g., lognormal) calibrated to measured distance empirical cumulated distribution functions (eCDFs) across precipitation bins. To keep the model robust under limited data at extreme precipitation rates, the implementation prefers simple parameterizations (often constant parameters across bins) when the measurement uncertainty is high.

5.3.2.3 Intensity distribution

Reported intensity values are highly sensor- and firmware-dependent. For precipitation-induced detections, a discrete non-negative distribution (e.g., Poisson) provides a practical fit to measured intensity histograms and can be calibrated directly from expectation values.

5.3.2.4 Attenuation

Detection loss on reference targets is modeled through an extinction coefficient dependent on precipitation intensity. Rather than using literature coefficients directly, the model fits the attenuation curve to the observed target detection probability decrease. This yields an empirically grounded attenuation parameterization aligned with the sensor's effective signal processing and calibration.

5.4 COMPARISON OF REAL AND SIMULATED LIDAR EFFECTS

The comparison of real and synthetic data is a key element for assessing the credibility and applicability of perception sensor simulation. While synthetic data enables scalable, repeatable, and fully controllable testing, real-world measurements provide the ground truth needed to validate whether simulated sensor responses reflect actual sensor behavior under realistic environmental conditions.

Real sensor data, acquired under controlled yet realistic conditions, captures complex physical effects that arise from the interaction between the sensor, the environment, and the observed objects. These effects include signal attenuation, backscatter, multi-path reflections, noise characteristics, and object-dependent response variations, which are strongly influenced by environmental factors such as precipitation, fog, aerosols, illumination, and surface properties. In contrast, synthetic data is generated based on abstracted sensor models and environmental representations, which inherently rely on assumptions, simplifications, and parameterizations.

A direct comparison between real and synthetic data therefore focuses on identifying systematic deviations and gaps between measured sensor behavior and simulated outputs. This includes, but is not limited to, differences in detection probability, intensity or amplitude distributions, noise levels, false positives, and object visibility as a function of range and environmental severity. Such comparisons reveal whether the simulation reproduces not only idealized sensor performance, but also the degradation effects observed in reality.

Within this report, due to the late availability of the data from the measurement campaign 2025/2026, the evaluation from Linnhoff in his dissertation is described as a role model. Real sensor data obtained from a dedicated measurement setup, similar to the Persival Weather Test Bench, under defined environmental conditions serves as a reference against which synthetic data is evaluated.

A practical limitation identified is that the OSI version used did not provide continuous environmental inputs (e.g., exact rain rate or visibility), only categorical weather descriptors;

therefore, quantitative validation required supplying measured reference signals via an external time series read by the model. A clear-weather baseline is evaluated first to confirm that the synthetic pipeline is sensible without environmental effects and to document any systematic differences caused by measurement-specific constraints (e.g., FoV restrictions in real recordings, off-region vegetation).

5.4.1 VALIDATION PHILOSOPHY

The comparison between measured and simulated sensor data aims to show that the simulation reproduces (i) the correct trends with environmental intensity and (ii) the variability observed in reality. Environmental effects are inherently stochastic and spatially inhomogeneous; therefore, validation does not target a scan-by-scan match, but rather agreement on (a) aggregated behavior and (b) plausible ranges.

Two acceptance criteria are used:

1. Trend criterion: simulated metrics must change in the same direction as measured metrics when environmental intensity changes (e.g., more atmospheric detections with increasing rain rate; reduced target detection probability with increasing attenuation).
2. Variation criterion: simulated values should lie within the spread of measurements (e.g., within mean ± standard deviation bands derived from repeated or aggregated observations).

Validation is limited to the measured validity ranges to avoid extrapolation beyond covered parameter space.

A virtual reference scene replicates the measurement situation (flat test area with reference targets at known distances and defined diffuse reflectivity). Real reference time series (precipitation rate for rain; visibility for fog) are used as ground-truth inputs for parameter sweeps and re-simulations.

5.4.2 SIMULATED AND REAL LIDAR DETECTIONS IN RAIN

Rain validation covers both dominant phenomena:

1. Atmospheric detections (false positives), which increase with rain rate.
2. Attenuation-driven detection loss on reference targets that increases with rain rate.

Additionally, rain validation includes distribution-level checks (e.g., distance distribution of atmospheric detections), because matching only mean detection counts is insufficient for credible synthetic data.

5.4.2.1 Trend-level validation: Atmospheric detection probability vs. rain rate (sweep)

A rain-rate sweep is simulated from near-zero to high precipitation. The key metric is the relative atmospheric detection probability in a defined atmospheric region (normalized by the number of beams entering that region), enabling direct comparison across time and scans.

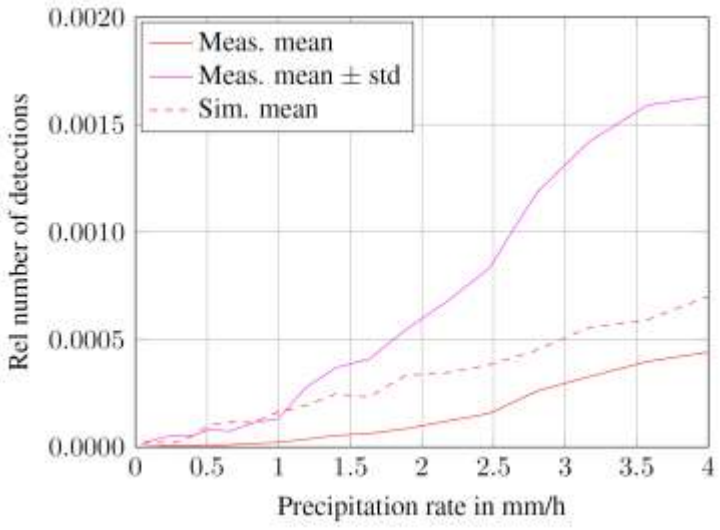

Figure 123: Comparison of atmospheric detection probability in measurement and simulation during light rain. (Linnhoff, 2023, p. 123)

Observed validation behavior:

- Moderate rain: simulation follows the measured trend and stays within the measured spread.
- Light rain regime: simulation may overpredict detections compared to measurements, indicating that a purely linear relationship down to zero can be too permissive and may require an onset threshold or nonlinear mapping near zero.
- Very high rain: deviations can occur near the edges of measurement support, where sample counts are lower and uncertainty increases.

Takeaway: always validate both (i) the full range and (ii) a zoomed light-rain region, because many models otherwise look correct globally while being biased at low intensities.

5.4.2.2 Sample-level validation: re-simulation under dynamic rain (time series)

To test the model's responsiveness under non-stationary rain, individual measurements with rapidly changing rain rates are re-simulated using the recorded precipitation-rate time series. Time traces of detection counts are compared between real and synthetic data.

Key interpretation points:

- In reality, reference precipitation sensors measure at a different location/volume than the lidar FoV and often apply temporal averaging; therefore, perfect temporal alignment between rain-rate peaks and detection peaks is not guaranteed.
- In some sequences, simulation may track the reference more "cleanly" than the measured lidar, reflecting real-world spatial inhomogeneity and reference mismatch.
- In low rain rates, simulation may produce detections when measurements show almost none, again motivating a threshold/offset in the detection-probability model.

Takeaway: time-resolved re-simulation is essential because it exposes reference-to-FoV mismatch and demonstrates why acceptance should be based on bands/trends rather than scan-level equality.

5.4.2.3 Distribution-level validation: distance distribution of atmospheric detections

A second validation dimension is the range distribution of atmospheric detections. Measured distance CDFs are computed in rain-rate bins and compared to the simulated distributions.

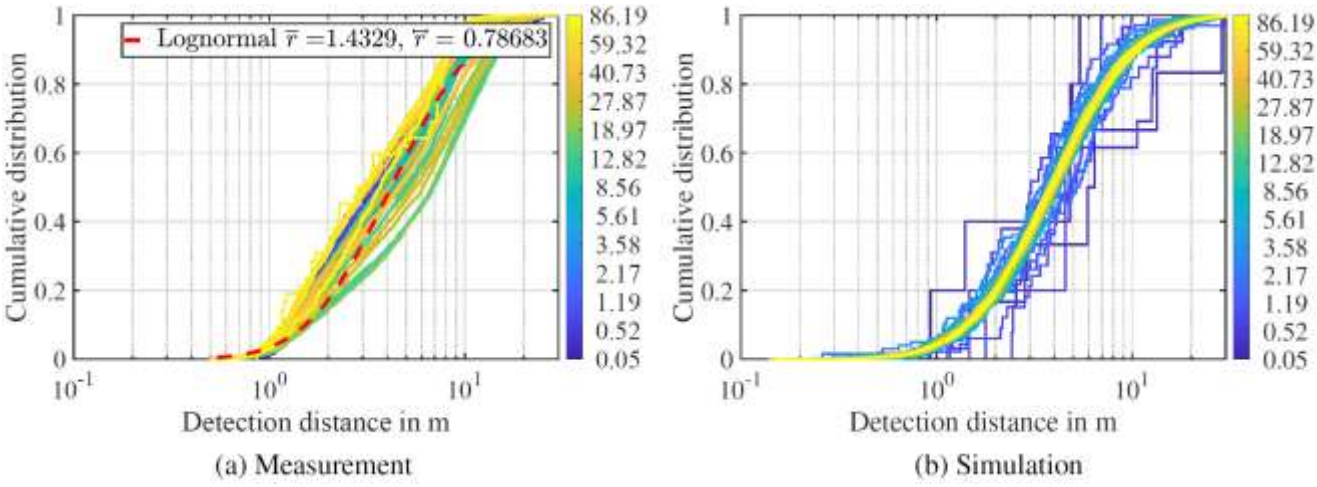

Figure 124: Comparison of detection distributions in atmosphere for rain. (Linnhoff, 2023, p. 124)

Typical results and meaning:

- Simulated distance distributions generally fall within the envelope of measured distributions across rain-rate bins.
- Differences can appear in intermediate rain bins where measured distributions shift toward farther distances; this highlights that a single parametric distribution (e.g., lognormal with constant parameters) may be a pragmatic approximation but not always perfect.

Takeaway: validating distributions (not only means) prevents "false confidence" where a model matches counts but misrepresents where false positives occur spatially—highly relevant for perception pipelines.

5.4.2.4 Attenuation validation: target detection probability vs. rain rate (multi-distance targets)

Attenuation is validated via detection probability on several reference targets at different distances. The expected behavior is a decreasing detection probability with increasing rain rate, with stronger effects at longer range.

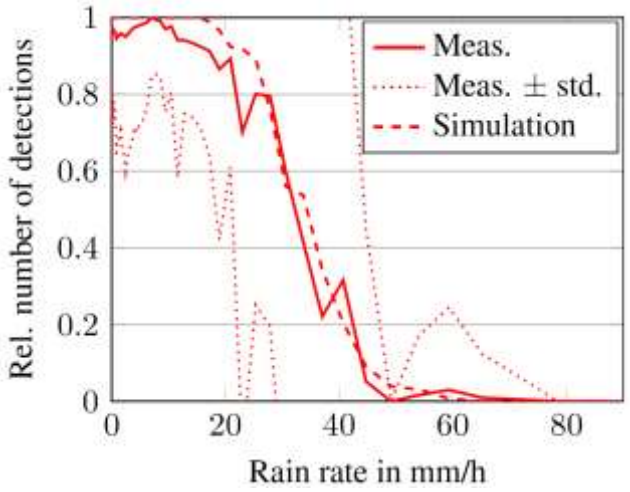

Figure 125: Comparison of target detection probability for measurement and simulation in rain (Linnhoff, 2023, p. 125). The dotted lines denote the measurement with added and subtracted standard deviation as an upper and a lower boundary.

In the evaluation, simulation reproduces the decreasing trends and typically remains within the spread of measured values, though measured spread can be wide due to reference uncertainty and scan-to-scan variability.

Takeaway: multi-target validation is a robust check that attenuation + thresholding behaves plausibly over range, which is essential when using simulation for safety-relevant evaluation.

5.4.3 SIMULATED AND REAL LIDAR DETECTIONS IN FOG

For fog, validation focuses primarily on detection loss on targets (attenuation effect) rather than atmospheric detections, because for the evaluated lidar configuration and the tested visibility range, measurements show no relevant atmospheric detections and the model therefore does not add them.

5.4.3.1 Trend-level validation via visibility sweep (aggregated comparison)

A synthetic visibility sweep is generated from clear to reduced visibility. For each visibility level, the relative number of detections on each reference target is computed and compared to the corresponding measured relationship.

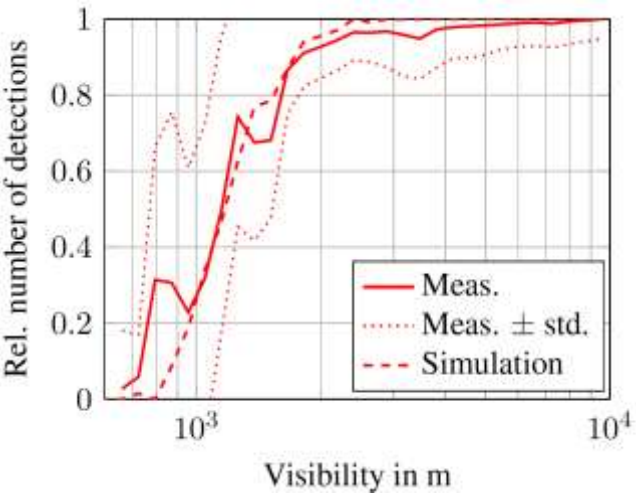

Figure 126: Comparison of target detection probability for measurement and simulation in fog (Linnhoff, 2023, p. 120). The dotted lines denote the measurement with added and subtracted standard deviation as an upper and a lower boundary.

Validation expectations and outcome:

- The simulation must show decreasing target detection probability as visibility decreases.
- The distance ordering must be preserved: farther targets degrade first, nearer targets later.
- Agreement is assessed against measured mean ± standard deviation. Simulated curves can align well for nearer targets while exhibiting a shift for farther targets; such shifts indicate sensitivity to the assumed thresholding behavior (e.g., whether effective thresholding is range-dependent).

Takeaway: fog validation must include multiple target distances to reveal whether attenuation and thresholding interplay is realistic beyond the calibration point.

5.4.3.2 Sample-level validation via re-simulation of an individual fog measurement

Beyond aggregated sweeps, one (or more) real fog measurements are re-simulated using the recorded visibility time series. Target detection counts over time are compared between measurement and simulation.

Typical insights:

- The timing of detection degradation often matches reasonably (degradation begins when visibility drops).
- The simulated degradation may be less steep or may not reach a full "zero detection" state as quickly as the measurement for some targets, highlighting the sensitivity to the stochastic component and effective threshold distribution.

Takeaway: single-measurement re-simulation tests temporal realism and shows where model variance assumptions should be adjusted to reflect the steepness and floor behavior seen in real data.

6. Conclusion

This deliverable addressed the need for measurable and reproducible evidence on how perception sensors behave under adverse environmental conditions, with a particular focus on enabling simulation-oriented modeling and credibility assessment. Using the Persival Weather Test Bench as a long-term measurement setup, synchronized recordings of reference weather parameters and perception sensor outputs were acquired and evaluated. The resulting performance descriptions link sensor KPIs to quantified environmental severity (e.g., brightness in lux, precipitation rate in mm/h, and fog severity via extinction coefficient), providing an experimental foundation for comparing real and synthetic sensor data in a traceable way.

For the evaluated camera, the derived performance curves highlight that different weather types lead to distinct degradation signatures in image-based KPIs. Fog acts as a global visibility limiter: increasing extinction introduces veiling luminance and causes a consistent reduction of contrast across all target pairs. In contrast, precipitation conditions (rain, snow, hail) primarily manifest as increased variability in mean intensity and contrast, consistent with transient, spatially inhomogeneous effects such as streaks/occlusions, specular highlights, and wet-surface glare; within the measured ranges, no uniformly monotonic contrast attenuation comparable to fog was observed. The brightness evaluation further indicates that camera-internal adaptation (e.g., auto-exposure and tone mapping) can strongly influence contrasts involving very dark references and should therefore be treated as a confounding factor when using performance curves for simulation parameterization or validation.

For lidar, the measurements underline that adverse-weather degradation is strongly sensor-configuration dependent (field of view, range class, and wavelength). Fog again emerges as the dominant object-independent effect, driving range-dependent detection loss (especially on the far 10% target) and, for rotating sensors, a pronounced increase in atmospheric backscatter detections. The special case of spatially non-homogeneous (patchy) fog illustrates that directional sensors may show clustered atmospheric detections in denser fog pockets even when they appear robust in homogeneous fog. In the observed precipitation regimes, rain and snow produced only low atmospheric detection probabilities and largely stable target detections, while hail induced measurable atmospheric detections above roughly 10 mm/h but did not significantly reduce target detectability; across precipitation types, the most visible signature is stochastic intensity/reflectivity fluctuation rather than a deterministic range collapse within the covered measurement ranges.

From a simulation perspective, the results support a performance-driven modeling approach that separates object-independent from object-dependent environmental effects. Object-independent effects can be represented efficiently by parameter-driven models calibrated to aggregated measurements: for fog, an explicit attenuation/veiling component parameterized by extinction coefficient (or MOR) is required; for precipitation, the measurements motivate stochastic disturbance models that primarily increase variance and generate atmospheric returns/outliers rather than applying a single deterministic attenuation curve. For lidar, this naturally translates into a two-part representation consisting of range-dependent attenuation (affecting detectability) and a stochastic atmospheric-return process (affecting false positives), with parameters depending on weather severity and sensor configuration.
For camera, the same separation clarifies that global visibility loss (fog) and transient disturbances (rain/snow/hail) should be validated with different acceptance criteria (trend

agreement vs. variability envelopes), and that camera operating regimes (exposure/adaptation) must be documented to make comparisons between real and synthetic data meaningful.

Several limitations should be considered when reusing the presented curves. Precipitation rate in mm/h is an indirect proxy for optical disturbance and does not capture particle size distributions, wind-driven trajectories, or spatial inhomogeneity; this is particularly relevant for snow and hail. Some categories have limited coverage at the high-intensity end (notably hail), so extrapolation beyond the measured intervals is not recommended. Likewise, the reported KPIs depend on the chosen ROIs/filters and the aggregation strategy (three consecutive camera frames, lidar ROI definitions) and are therefore specific to the measurement configuration.

Overall, the deliverable contributes to closing the gap between experimental testing and virtual validation by demonstrating how controlled adverse-weather measurements can be converted into quantitative performance descriptions and simulation-ready model components. The presented results and the referenced methodology (Linnhoff, 2023) directly support the objectives of T_Simu.2 by informing which environmental parameters must be represented and how they can be mapped to sensor-level effects, and they provide a structured basis for T_Simu.3-style comparisons between real and synthetic sensor data. Beyond WeatherSafe, the same workflow can be extended to additional sensors (radar), broader weather-severity ranges, and object-dependent effects (e.g., spray and sensor-window contamination, glare), enabling more comprehensive and credible safety-oriented assessment of perception systems under adverse environmental conditions.

7. Acknowledgements

The WeatherSafe consortium would like to acknowledge for their support and work all the Partners and Third Parties involved:



8. Bibliography

Linnhoff, C. (2023). *Analysis of Environmental Influences for Simulation of Active Perception Sensors.* Darmstadt: Technical University of Darmstadt.

9. Table of illustrations

Figure 1: Persival Weather Test Bench setup on a roof top 9
Figure 2: Weather station with lidars and a camera in the front and weather reference sensors on top 10
Figure 3: Target housing with 4 different Lambertian targets 11
Figure 4: Different target regions in the camera image 12
Figure 5: Distribution of brightness values in all considered daily camera measurements 17
Figure 6: Camera image in sunlight 18
Figure 7: Camera image with bright sunlight and snow on the ground 18
Figure 8: Camera image with distinct clouds 19
Figure 9: Aggregated data of 5% Target: Mean intensity in brightness 19
Figure 10: Aggregated data of 10% Target: Mean intensity in brightness 20
Figure 11: Aggregated data of 50% Target: Mean intensity in brightness 20
Figure 12: Aggregated data of 80% Target: Mean intensity in brightness 21
Figure 13: Aggregated data of 95% Target: Mean intensity in brightness 21
Figure 14: Aggregated contrast of camera in brightness: 95% target vs. 5% target 22
Figure 15: Aggregated contrast of camera in brightness: 80% target vs. 50% target 22
Figure 16: Aggregated contrast of camera in brightness: 95% target vs. 50% target 23
Figure 17: Histogram of fog data 24
Figure 18: Camera image in light fog 25
Figure 19: Camera image in medium fog 25
Figure 20: Camera image in heavy fog 26
Figure 21: Aggregated data of 5% Target: Mean intensity in fog 27
Figure 22: Aggregated data of 10% Target: Mean intensity in fog 27
Figure 23: Aggregated data of 50% Target: Mean intensity in fog 28
Figure 24: Aggregated data of 80% Target: Mean intensity in fog 28
Figure 25: Aggregated data of 95% Target: Mean intensity in fog 29
Figure 26: Aggregated contrast of camera in fog: 95% target vs. 5% target 29
Figure 27: Aggregated contrast of camera in fog: 80% target vs. 50% target 30
Figure 28: Aggregated contrast of camera in fog: 95% target vs. 50% target 30
Figure 29: Histogram of hail data for camera evaluation 32
Figure 30: Aggregated data of 5% Target: Mean intensity in hail 32
Figure 31: Aggregated data of 10% Target: Mean intensity in hail 33
Figure 32: Aggregated data of 50% Target: Mean intensity in hail 33
Figure 33: Aggregated data of 80% Target: Mean intensity in hail 34
Figure 34: Aggregated data of 95% Target: Mean intensity in hail 34
Figure 35: Aggregated contrast of camera in hail: 95% target vs. 5% target 35
Figure 36: Aggregated contrast of camera in hail: 80% target vs. 50% target 35
Figure 37: Aggregated contrast of camera in hail: 95% target vs. 50% target 36
Figure 38: Histogram of rain data for camera evaluation 37
Figure 39: Aggregated data of 5% Target: Mean intensity in rain 37
Figure 40: Aggregated data of 10% Target: Mean intensity in rain 38
Figure 41: Aggregated data of 50% Target: Mean intensity in rain 38
Figure 42: Aggregated data of 80% Target: Mean intensity in rain 39
Figure 43: Aggregated data of 95% Target: Mean intensity in rain 39
Figure 44: Aggregated contrast of camera in rain: 95% target vs. 5% target 40
Figure 45: Aggregated contrast of camera in rain: 80% target vs. 50% target 40
Figure 46: Aggregated contrast of camera in rain: 95% target vs. 50% target 41
Figure 47: Histogram of snow data for camera evaluation 42

Figure 48: Camera image during snow fall 43
Figure 49: Aggregated data of 5% Target: Mean intensity in snow 43
Figure 50: Aggregated data of 10% Target: Mean intensity in snow 44
Figure 51: Aggregated data of 50% Target: Mean intensity in snow 44
Figure 52: Aggregated data of 80% Target: Mean intensity in snow 45
Figure 53: Aggregated data of 95% Target: Mean intensity in snow 45
Figure 54: Aggregated contrast of camera in snow: 95% target vs. 5% target 46
Figure 55: Aggregated contrast of camera in snow: 80% target vs. 50% target 46
Figure 56: Aggregated contrast of camera in snow: 95% target vs. 50% target 47
Figure 57: Histogram of brightness data for lidar evaluation 49
Figure 58: Lidar detection probability per beam in brightness 50
Figure 59: Number of detections on target 1 in brightness 51
Figure 60: Mean intensity of detections on target 1 in brightness 51
Figure 61: Number of detections on target 2 in brightness 52
Figure 62: Mean intensity of detections on target 2 in brightness 52
Figure 63: Number of detections on target 3 in brightness 53
Figure 64: Mean intensity of detections on target 3 in brightness 53
Figure 65: Number of detections on target 4 in brightness 54
Figure 66: Mean intensity of detections on target 4 in brightness 54
Figure 67: Number of detections on target 5 in brightness 55
Figure 68: Mean intensity of detections on target 5 in brightness 55
Figure 69: Histogram of extinction coefficients in the fog measurements 57
Figure 70: Detection probability per beam in atmosphere in fog 58
Figure 71: Number of detections on target 1 in fog 58
Figure 72: Mean intensity of detections on target 1 in fog 59
Figure 73: Number of detections on target 2 in fog 59
Figure 74: Mean intensity of detections on target 2 in fog 60
Figure 75: Number of detections on target 3 in fog 60
Figure 76: Mean intensity of detections on target 3 in fog 61
Figure 77: Number of detections on target 4 in fog 61
Figure 78: Mean intensity of detections on target 4 in fog 62
Figure 79: Number of detections on target 5 in fog 62
Figure 80: Mean intensity of detections on target 5 in fog 63
Figure 81: Short-range rotating lidar (lidar A) in uneven fog. The points are colored by z-coordinate. All colored points, except from the target, are due to the fog. 64
Figure 82: Long-range rotating lidar (lidar B) in uneven fog. The points are colored by z-coordinate. All colored points, except from the target, are due to the fog. 64
Figure 83: Long-range directional lidar (C in uneven fog. The points are colored by z-coordinate. All colored points, except from the target, are due to the fog. 65
Figure 84: Short-range directional lidar (D) in uneven fog. The points are colored by z-coordinate. All colored points, except from the target, are due to the fog. 65
Figure 85: Camera image in uneven fog as reference 66
Figure 86: Histogram of hail data for lidar evaluation 67
Figure 87: Detection probability per beam in atmosphere in hail 68
Figure 88: Number of detections on target 1 in hail 68
Figure 89: Mean intensity of detections on target 1 in hail 69
Figure 90: Number of detections on target 2 in hail 69
Figure 91: Mean intensity of detections on target 2 in hail 70
Figure 92: Number of detections on target 3 in hail 70
Figure 93: Mean intensity of detections on target 3 in hail 71
Figure 94: Number of detections on target 4 in hail 71

Figure 95: Mean intensity of detections on target 4 in hail 72
Figure 96: Number of detections on target 5 in hail 72
Figure 97: Mean intensity of detections on target 5 in hail 73
Figure 98: Histogram of rain data for lidar evaluation 74
Figure 99: Detection probability per beam in atmosphere in rain 75
Figure 100: Number of detections on target 1 in rain 75
Figure 101: Mean intensity of detections on target 1 in rain 76
Figure 102: Number of detections on target 2 in rain 76
Figure 103: Mean intensity of detections on target 2 in rain 77
Figure 104: Number of detections on target 3 in rain 77
Figure 105: Mean intensity of detections on target 3 in rain 78
Figure 106: Number of detections on target 4 in rain 78
Figure 107: Mean intensity of detections on target 4 in rain 79
Figure 108: Number of detections on target 5 in rain 79
Figure 109: Mean intensity of detections on target 5 in rain 80
Figure 110: Histogram of snow data for lidar evaluation 81
Figure 111: Detection probability per beam in atmosphere in snow 82
Figure 112: Number of detections on target 1 in snow 82
Figure 113: Mean intensity of detections on target 1 in snow 83
Figure 114: Number of detections on target 2 in snow 83
Figure 115: Mean intensity of detections on target 2 in snow 84
Figure 116: Number of detections on target 3 in snow 84
Figure 117: Mean intensity of detections on target 3 in snow 85
Figure 118: Number of detections on target 4 in snow 85
Figure 119: Mean intensity of detections on target 4 in snow 86
Figure 120: Number of detections on target 5 in snow 86
Figure 121: Mean intensity of detections on target 5 in snow 87
Figure 122: FZD OSMP framework based on a generic functional decomposition of automotive perception sensors (Linnhoff, 2023, p. 95) 91
Figure 123: Comparison of atmospheric detection probability in measurement and simulation during light rain. (Linnhoff, 2023, p. 123) 95
Figure 124: Comparison of detection distributions in atmosphere for rain. (Linnhoff, 2023, p. 124) 96
Figure 125: Comparison of target detection probability for measurement and simulation in rain (Linnhoff, 2023, p. 125). The dotted lines denote the measurement with added and subtracted standard deviation as an upper and a lower boundary. 97
Figure 126: Comparison of target detection probability for measurement and simulation in fog (Linnhoff, 2023, p. 120). The dotted lines denote the measurement with added and subtracted standard deviation as an upper and a lower boundary. 98