Ref. Ares(2025)6666906 - 18/08/2025

ROADVIEW

Robust Automated Driving in Extreme Weather

Project No. 101069576

Deliverable 4.2

Evaluation of sensors in-situ

WP4 – Data logging, filtering, annotation, and quality assurance

Authors	Martin Sanfridson (RISE) and Henrik Toss (RISE)
Lead participant	RISE
Delivery date	18 August 2025
Dissemination level	Public
Type	Document, report

Version 1.0

Co-funded by the European Union

Co-funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or European Climate, Infrastructure and Environment Executive Agency (CINEA). Neither the European Union nor the granting authority can be held responsible for them. Project grant No. 101069576.

UK participants in this project are co-funded by Innovate UK under contract No. 10045139.
Swiss participants in this project are co-funded by the Swiss State Secretariat for Education, Research and Innovation (SERI) under contract No. 22.00123.

Revision history

Author(s)	Description	Date
Martin Sanfridson (RISE), Henrik Toss (RISE)	Draft deliverable, version 0.6	06/08/2025
Stefan Parzefall (AVL)	Revision 1	12/08/2025
Martin Sanfridson (RISE), Henrik Toss (RISE)	Version 0.7	13/08/2025
Isabelle Morvan (CRF)	Revision 2	14/08/2025
Martin Sanfridson (RISE)	Final version v1.0	15/08/2025
Johannes Ripperger (accelCH)	Final checks	18/08/2025

Contents

List of Figures

List of Tables

Partner short names

RISE	RISE Research Institutes of Sweden AB
CRF	Canon Research Centre France
accelCH	accelopment Schweiz AG
VTT	VTT Technical Research Centre of Finland Ltd
AVL	AVL Software and Functions GmbH

Abbreviations

D	Deliverable
Euro NCAP	European New Car Assessment Programme
FoV	Field of View
GNSS	Global Network Satellite System
IR	Infrared
RCS	Radar Cross Section
SNR	Signal-to-Noise Ratio
ToF	Time of flight
UNECE	United Nations Economic Commission for Europe
WP	Work Package

Executive summary

Objectives

The objective of this deliverable is to study water-based contamination on radar and lidar. All sensors require shields to protect the optics and electronics. One way to divide weather disturbances is to split them according to location: contamination on the sensor, aerosols in the media (air) and contamination on the target object or environment. In practice, it is often a combination. The focus of this deliverable is contamination deposited on the sensor, on its aperture. This contamination causes degradation of the sensor output and in the extension malfunction of the intended feature and thus has to be removed. Removing it is however costly in terms of installation, fluids and energy. The deliverable is primarily an investigation to build knowledge. The deliverable aims to support the development of more resilient perception systems by providing a comprehensive understanding of how surface-level contaminants affects sensor performance that can develop into a termination of the automated driving, but also to help limit the need of cleaning the shield from contaminants.

Methodology and implementation

The methodology comprises both lab tests and investigation of field operational test data. In the lab tests, the target is known or constant and the contaminant, deliberately causing a disturbance, is added in front of the sensor in way that makes them possible to estimate or control. From the field operational test, targets or objects are typically not known, and the amount of contamination is not controlled and might also be difficult to estimate.

Outcomes

Snow buildup on the sensor surface, especially in the forward-facing direction, can quickly obstruct large parts of the aperture. Wet snow is particularly problematic—it sticks easily, attenuates radar signals, and reflects lidar wavelengths, causing a loss of lidar points in a critical forward sector. A field test shows that this buildup can occur within minutes, depending on weather and vehicle design. Lidar can go from fully functional to having a blocked field of view (FoV) rapidly. Radar performance degrades more gradually, as attenuation increases with the thickness of the snow layer. This leads to reduced signal strength, lower signal-to-noise ratio (SNR), and difficulty detecting low-reflectivity targets, effectively shrinking the radar's usable FoV. A proposed method to assess radar blockage in run-time is to monitor the consistent ground signal just ahead of the vehicle. This can help estimate the reliable FoV, where radar can still detect typical targets despite contamination. For lidar, a camera monitoring the aperture can predict signal degradation in lab environment. Another method that will work better at run-time involves observing fixed surfaces like road segments, walls, or traffic signs. These realistic targets help identify whether signal loss is due to blockage, reflection, or aperture contamination. A useful alternative is to monitor a consistently visible object, such as part of the vehicle's hood.

Next steps

For the lidar, figuring out damping in real-time in order to find signs of beginning of a blockage is a good aim. An approach to try is graph neural networks. Another area of improvement is not only to characterize the impact on the output from the lidar but also the output from a typical perception feature, or at least the mapping of these. Further, it is cumbersome to work with snow, so finding a substitute would be great as an embryo for standardized testing.

The radar signature of part of the road suggested as an indicator surface needs to be further investigated. It has now only been characterized under static conditions, but also needs to be recorded for a moving sensor as this should affect the signal, and thus the reported detections. Furthermore, different types of road surfaces need to be characterized to ensure the method is valid also for driving on road surfaces such as gravel, snow or wet asphalt. Suggestions on determination of the reliable FoV and adaptations of functions could also be part of future work.

1 Background

1.1 Blockage of sensor shields

This report brings up the case of water-based contaminants directly on sensor shields for radar and lidar. The majority of work in the ROADVIEW project concerns disturbances in the air, precipitation such as snowfall or rainfall. When snow with high moisture content falls on a surface, it can quickly build up and prohibit the sensor to detect its surroundings. When sensors are blocked/blinded the availability of automatic driving and of driver assistance features are impaired.

A simple model of weather disturbances is to split them into three domains: contamination of sensors surfaces, aerosols in the air and deposited matter on targets and environment. In a practical situation it isn't always possible to completely separate those, but in this report we concentrate on the first one. Sensor models for simulation of compound disturbances have also been worked on in the project [4].

The report is structured in the following way. In this first section, an introduction to the hardware: the use of shields, and state-of-the art concerning blockage/blinding of radars and lidars. Section 2 brings up the methodology and describes the equipment used and experimental setups. Section 3 and 4 describe, for lidar and radar respectively, results from the lab experiments and from logged data on the road. Finally, Section 5 offers a discussion of the results and future work.

Lidar and radar need protection from dirt and water. In the ROADVIEW project there is no specific design for shields, but the original shields of the electronics. Vehicle manufacturers typically want to embed sensor in to make an aesthetic design which also is aerodynamically sound to avoid contamination of surfaces. Radars are typically placed behind the plastic skirts of the vehicle, and it is possible to put the lidar behind the wind shield. A shield for a radar is commonly called radome.

It might seem straightforward to clean a sensor as soon as contamination occurs. However, this is costly. It requires installing a cleaning system and feeding it with energy or fluids. To reduce cost, activation of the cleaning system should be optimized with constraints. We touch upon such diagnostics in the report.

We use the term 'contamination' or 'soilage' to refer to a contaminant or disturbance, close to the sensor, that obstructs the transmission or reception of light or radio signals. The contaminants in focus here are snow and ice. Blockage is due to contamination of the sensor's aperture or optical window, and we use it here also for partial blocking. The term 'occlusion' is more general than blocking and can also be used for phenomena at a longer distances, where possible counter measures for snow and ice like heating, flushing or similar cannot remove the contaminant. Blockage can of course also be due to other objects that stick to the shield. We define blinding as when the received signal saturates the sensor, which can be due to a close retro reflector, and we will use the term 'blockage' even though blinding arguable might be present at the same time.

1.2 State-of-the-art regarding contamination of lidar

The aim here is to review current research on the characterization and impact of surface contaminants such as snow and ice accumulating on the lidar sensor's optical aperture (i.e., the transparent window through which the laser is emitted and received), and how that degrades sensor functionality and perception reliability. Disturbances in the form of aerosols are left out here.

In UN ECE R157 [3], which regulates general requirements regarding safety of active lane keeping support, it is regulated that non-availability should be notified to the driver within a specific time frame, but it is not mentioned how the testing should be carried out. Euro NCAP is in their vision for 2030 [2] aims to include more elaborate testing conditions including weather conditions. There is no (de facto) standardized method for testing this type of blockage today.

Six different contaminants (no snow or ice, but dew) was experimented with a well-known static target, an Ouster sensor and also a Riegl sensor that allowed inspecting the full waveform, i.e. the time-series return [7]. A typical effect is few points in the cloud and higher average distance to target. Wet droplets from condensation (dew) was the most severe one. With some exception, the amount of contaminant was not estimated.

In an experiment, four different types of contaminants (one was frost) were applied [5] and a neural network based approach for classifying the contaminants was proposed. Data was collected in a lab test outdoors with varying lighting conditions, with contaminants that covered the complete aperture, and with no other objects in the near field apart from respective contaminant. The object and background can thus be considered known. This may not accurately reflect real-world environmental conditions.

Seven different contaminants (no ice or snow) are recorded in two similar test environments and a Graph Neural Networks method is employed [6]. Graph neural networks are mathematical models that can learn functions over graphs and are a contemporary approach for building predictive models on graph-structured data. The test environments contained typical objects. Dynamic aspects were covered by varying the location of the sensor to mimic a sampled trajectory. This is a more realistic stetting compared to [5]. First, the point cloud is voxelized, then voxel points are treated as graph nodes and a graph attention layer is added. All nodes are taken through a global mean pooling to go from node based decision on contamination or not, into a global probability.

Related, but not considered in ROADVIEW is to consider defects of the shield, that is holes, cracks and scratches [8]. An Ouster with a static well-defined target was used. Most holes and cracks prolonged the measured distance slightly outside on standard deviation for the non-contaminated reference. The conclusion is that clearly visible damages to the shield should be avoided.

A related study with homogeneous well-defined dust on a flat plastic sheet and well-defined low reflectivity target and a rotating Velodyne lidar at different distances, concluded the number of points returned from the target, with respect to number of theoretical points, is a relevant measure [9] and the range reduction is readily estimated and compared to the sensor's specification.

Models trained solely on data without contaminants on the surface suffer under contaminated conditions, suggesting that incorporating contaminated data during the training phase is imperative [10]. A dataset containing a collection of uncontaminated and realistically contaminated lidar samples is presented. A measure of accuracy for the perception is proposed rather than proposing the intermediate degradation of sensor data. The approach depends on a well-known open perception to render repeatable figures.

1.3 State-of-the-art regarding contamination of radar

Previous efforts have been made at detecting blockage wherein the issue of false blockage detection in open environments has also been addressed [15]. The method presented has two major drawbacks – it only works in a binary state, i.e. blocked/unblocked, and it is dependent on the raw range-doppler data heat map, a type of data that cannot be said to typically be readily available as a radar sensor output. A more recent patent claims to address the issue of quantifying the blockage but also appears to rely on low level radar data [16], [17].

Another effort at blockage detection has been made instead using machine learning and statistics of all detection over time [18]. This has the obvious drawback of only working in the type of environments it has been trained on, but is more accessible than algorithms relying on lower level data.

The issue of detecting blockage, and testing thereof, has also been addressed in Euro NCAP technical note on blockage of radar [1].

2 Methodology

This section describes the experimental setups used in the report. Corresponding result are found in Section 3 and Section 4. The basis is a mix of lab setups and using data from a test vehicle in the project. The working principles of the radar and lidar are quite different which is reflected in the tools used and contamination manufactured.

The following delimitations apply. Camera is not part of the sensor suite studied in this report. The impact on the sensor shield itself is not studied. The only contaminants aimed at are ice, snow and water. Removing of contamination is not studied. Blinding is not studied. When to discontinue or inhibit a feature that uses lidar or radar is not discussed.

2.1 Lidar in laser test bench

The laser test bench setup is originally used to investigate the stability of a laser. The test bench consists of an image sensor onto which the laser is concentrated using a parabolic mirror. A lab computer samples images and displays the cross section of the laser beam at a long distance. It is not possible to use the reading of the lidar itself in this setup. Some experiments were conducted to dampen the laser by contaminants put into the path, but the approach was abandoned. In this report, the test bench help us understand more aspects of how the this particular lidar works.

2.2 Lidar in ice and glass damping experiment

The purpose of this straight forward setup is to investigate the impact of material in front of the sensor. The distance to the Lambertian target was 3.5 m. For each sample tested, 100 lidar frames were captured and cropped to 110 x 70 points. A camera without infrared (IR) filter was aimed at the target to detect changes in the pattern. The lidar was configured to highest resolution and 10 Hz rotation. The room was dark. A bracket was constructed to hold samples 2 cm from the lidar in the path to the target.

2.3 Lidar in snowfall

In this static setup the aim was to study the relative change on lidar point clouds due to accumulated snow. The setup consisted of a rotating lidar (Ouster OS2), a camera for automated capturing, and a target scenery with a checkered board, a couple of larger Lambertian boards and a couple of retro reflectors for orientation of the FoV, see the sketch in Figure 1. The lidar was placed outdoors 1.8 m away from a roof under which the targets were placed. The camera was 4 dm away from shield and it had to be located at an angle from the main path to the targets to not disturb lidar or snowfall.

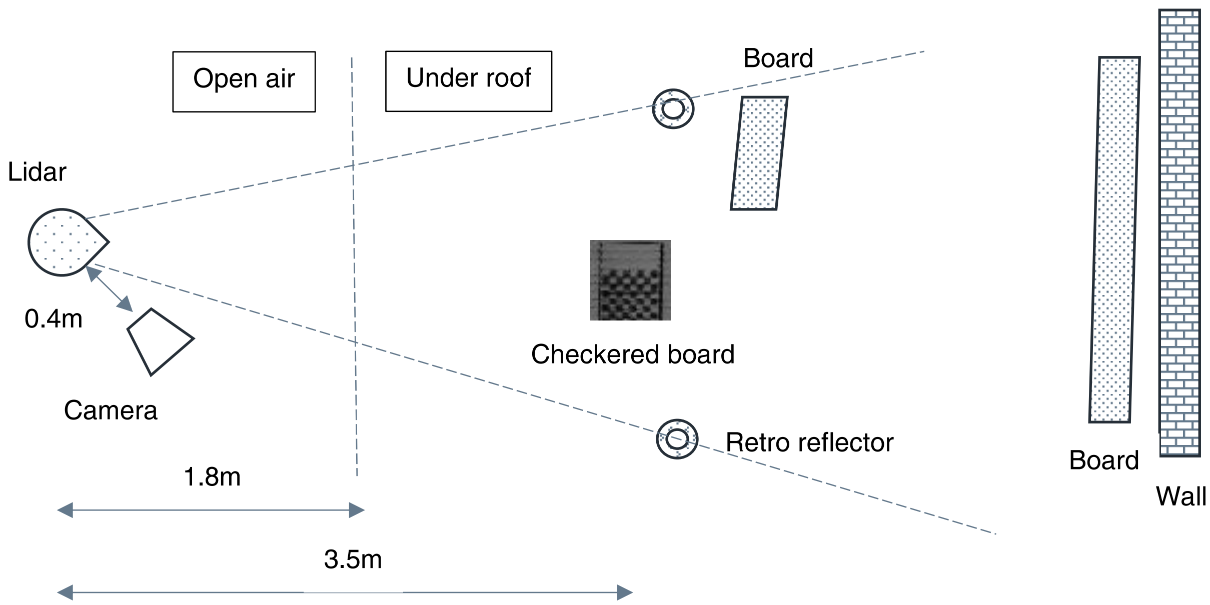

Figure 1. Sketch of setup. The picture of the checkered board is from the lidar, intensity with clean shield.

As the targets facing the lidar were located under a roof they are thus not affected by the snow. The amount of snow in the air volume is low because the distance between the sensor and the overhead is short. The greyness of the camera was not calibrated.

The surface temperature of the lidar will in many cases make the snow melt away. To avoid this, a plastic shield (styrene acrylonitrile) is used as is shown in Figure 2, with a 3-4 cm space between the lidar and the shield in the direction towards the targets. It has also been given a small angle to the gravity gradient to help snow stick to its surface. Throughout the experiment the surface of the lidar is dry. Note that the plastic shield also dampens the signal and another drawback with the design is that a flat shield is not suitable for a rotating lidar. The lidar gets warm during operation and the plastic shield avoids heat transfer to the accumulated snow.

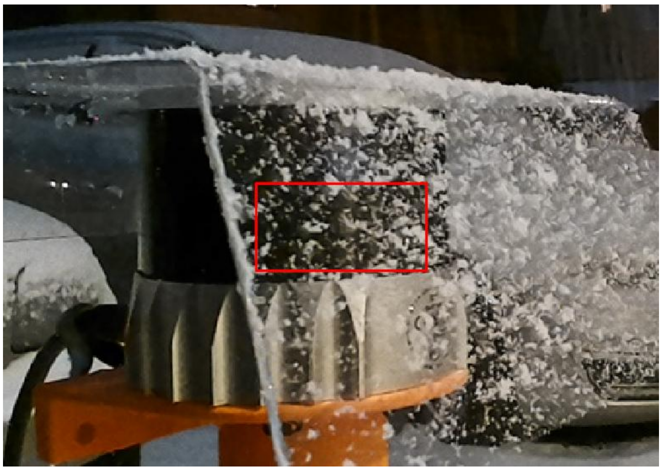

Figure 2. Snow on a flat shield. The red box is the area used to estimate the coverage of snow. The amount of snow in the picture has impact on the lidar.

The primary recording was done in light snow fall, that is a snow fall rate of around 1 mm/h but very fluctuating over time. The outdoor temperature started at -1 °C in the morning and rose to 0 °C during the day. The water content of the snow was estimated and it increased during the day. The wind was between 8 and 12 m/s creating a draft. The rate of the snowfall at the time of measurement is estimated to be 1-2 mm/h.

Depending on the current outdoor temperature the surface of the shield was either snowy as in the picture or full of droplets. The latter condition has not been studied. The area used to estimate the snow coverage, or more precisely its shape, size and location, is not exactly the same area as the laser rays pass, so an assumption is that the shield gets covered in a uniform random way.

2.4 Lidar from AVL's vehicle driving in snowy conditions

For the in-situ test, data from AVL's vehicle has been used. Lidar point clouds from driving in rural areas in southern Germany on a snowy day has been the main data set. There are some differences to lab tests. First, the lidar is in motion which changes the way snow accumulates. Second, there is no guaranteed target object for comparisons. Third, there are other reasons for loss of target for example a specular surface. The aim is to understand the current capability of the lidar, or to be more precise, the current incapability as summarized in the question: is the lidar partially blocked or attenuated?

Figure 3. Overtaking position.

AVL's vehicle, see D4.3 [14], has three lidars but in this work only the one mounted on the top of the vehicle, a Velodyne VLP-32 is used. As the name suggests, it has 32 layers. It has a different definition of intensity, but is otherwise very similar to Ouster OS2 used in the lab experiments in this report. A front looking camera is located 1.5 dm directly below behind a small plastic dome. Both the lidar and the camera can see the hood of the vehicle as seen in Figure 3, which comes from the Sunny dataset, see Table 1. Any contaminant is likely to end up both on the camera and on the lidar.

Table 1. List of datasets used for the present study.

Snow at day	Used mainly the first three minutes. Starts at parked position, then drives through a village. The weather is overcast and it snows moderately, the road is quite wet with patches of snow mainly beside the road. The data was recorded in January.
Sunny data	Used mainly the first three minutes. Starts parked, then drives in suburban setting. The road is dry and the sky is clear, the sun is low. The data was recorded in June.

The data logs are arranged by a file describing the path of the vehicle in global coordinates. To limit size, each file covers one minute of recording. Coordinate transformations to other sensors are given. For each waypoint in the path, pointers to synchronized sensor data is given. The lidar data frames are adjusted for the velocity vectors to account for the slow frame sampling of the lidar (10 Hz). The GNSS data is important in order to translate between local vehicle frame and global road frame.

2.5 Radar

Unfortunately the data set available from a moving automotive radar uses a radar with a very narrow beam in elevation and no elevation information. It is thus very difficult to discern signal from the road surface if at all present. No other constant features that could be used as similar indicators have been identified. A stationary test setup with a wide elevation beam 4D radar was instead used. For a future dynamic test set up it is important to also include some independent measure of the velocity of the radar so that the velocity reported by the radar for the reference surface can be compared to said independent measure. For further information on the measurement set-up and results see section 4.4.

3 Impact on the lidar

In this section, the impact of snow and ice on a lidar sensor will be presented. We are limited to the type of lidar used in the project that is, a rotating multi-layer time-of-flight lidar. First, a short recap of how the sensor works, focusing on parts related to the nature of the contaminants. Second, the two following subsections presents results from lab experiments on ice and snow, concluded by a subsection on the study of lidar when driving in snowfall.

Contamination and damages to the lidar's shield reduce the amount of transmitted light both in the sender and the receiver path. This is caused by scattering and absorption of the laser light.

3.1 Principle time-of-flight lidar point cloud estimation

This section describes briefly how a rotating time-of-flight (ToF) lidar produces a point cloud in order to give an idea of how a soilage on the sensor will interfere. We start in the lab setup described in Section 2.1 which has been used to produce Figure 4. The lidar is orientated such that each row corresponds to layer: Velodyne VLP-32 has 32 layers and Ouster OS2 has a maximum of 128 layers.

The intensity of each point in the cloud tells the return strength of a laser beam after striking an object. Different materials have different reflectivity and an object's distance and orientation can diminish the returned signal intensity. Typically, the longer the distance from the sensor to the object, the lower the object's intensity.

In a ToF lidar, each echoed laser beam is sampled which forms a time series. The azimuth and elevation angles are known. Peaks in the time series correspond to reflections and tell the range, thus completing the spherical annotation, and can result in a point in the resulting cloud. Both the Velodyne VLP-32 and the Ouster OS2 version used here only supports finding one peak per beam. The height of the peak corresponds to the intensity of the returned signal. For the Ouster sensor, the minimum range is around 1.0 m, which means that no point will be reported within that range, but the laser beam can of course still be attenuated or scattered in that interval. The wavelength of the laser is 865 nm.

The configuration of the lidar can matter to the interference of contaminants. For all layers together, the azimuth angle resolution is configurable to 512, 1024, or 2048 points; the maximum number of echoed laser beams is 128 times the configured azimuth angle resolution. More points per revolution means less time and less energy for estimating a point in the point cloud. Another configuration parameter is the spinning rate, 10 or 20 Hz. Higher revolution speed means less time and less energy to estimate a point in the point cloud.

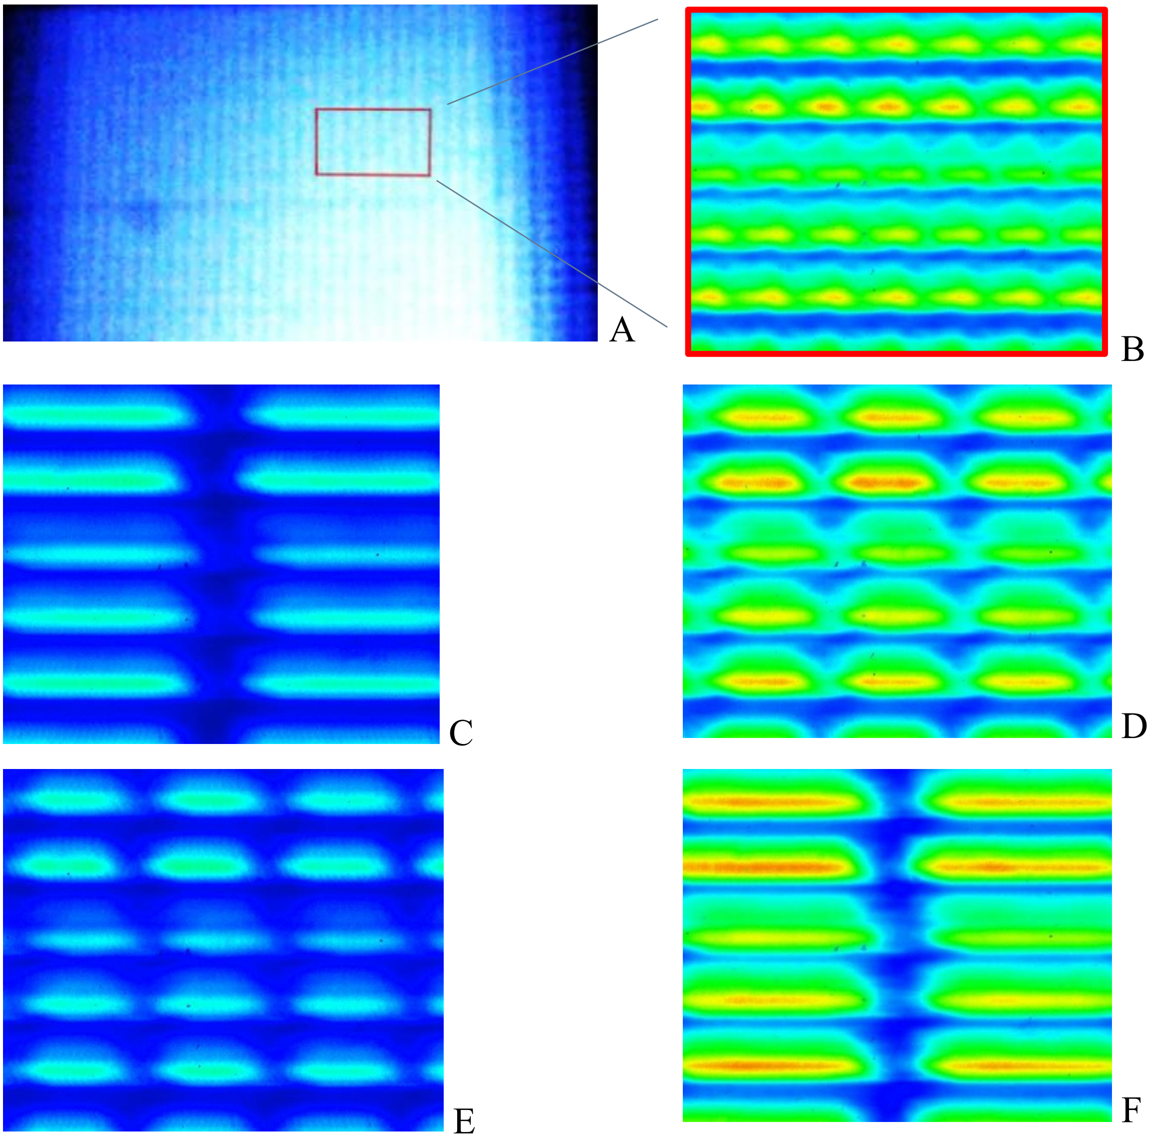

Figure 4. Laser footprint investigation for available operation modes.

The subplots in Figure 4 represent the far field pattern of individual laser beams at some long distance away from the lidar. Laser footprint for configurable operation modes are: 2048p at 10Hz, 512p at 10Hz, 512p at 20Hz, 1024p at 10Hz, 1024p at 20Hz. Subplot A shows the laser dots projected on a wall, which in contrary to the other subplots is photographed by a camera without traditional IR filter. Subplot B "zooms in" on a smaller part such as the red square in A for the mode 2048 points at 10 Hz. Subplot C shows 512 points at 20 Hz, subplot D show 1048 points at 10 Hz, subplot E shows 1048 points at 20 Hz and finally subplot F shows 512 points at 10 Hz. A laser beam is diverging and according to data sheet of the Ouster OS2 sensor the divergence is 1.6 mrad. Examples: the footprint across size is of 2 cm at 1 m, 2.7 cm at 10 m and 9.8 cm at 100 m. The distance between beams in the elevation and azimuth directions of course also increases accordingly, from part of centimeter to several decimeters, depending on the settings.

It is noticeable from Figure 4 that the vertical layers have different magnitudes and slightly different shapes. Individual laser footprints within the same layer seems to be consistent. Calibration should be considered for intensity readings. The pattern is stable over time, at least over minutes.

The front lens aperture size affect the impact a contaminant has on the point cloud (or image in case of a camera). The scanning pattern of the lidar means that the momentary aperture is rotating to cover 360 degrees of the sensor window. If the size of an opaque contaminant is larger than the aperture it will completely block. If the contaminant is smaller, light is partially attenuated and a partial impact will be seen. Sensors with large optical apertures are more resilient to contamination when enough of the beam pass around the contaminant.

It is specified [12] that the beam diameter exiting sensors is 19 mm, so a speculation is that the collecting lens is roughly the same size. Blocking one, and only one, specific outgoing beam at a distance that corresponds to the location of the shield is not trivial. Blocking a returning beam would necessitate blocking an area 19 mm across, which definitely affects neighbouring points. The umbra of soilage on the sensor much smaller than 19 mm ends up as damping the echo, lowering the reported returned intensity for the corresponding point. The optical design is robust to small and few marks or scratches on the aperture. A video from the manufacturer gives a nice overview how the rotating sensor is constructed [11].

3.2 Ice, plastic and glass

A lab experiment with ice was carried out, see the setup description in Section 2. To make the experiment repeatable, clear ice was manufactured. The experiment was made in room temperature, and the surface of the ice melted which avoids frost. Ice is seldom clear, but contaminated with air in different ways, or there is frost on the surface which causes diffuse scattering of the light.

The diameter of the sample ice disc is 7 cm, which is large enough to cover the aperture, see Figure 5. Clear discs of different thickness were manufactured. The result is shown in Table 2, where intensity has been scaled by the intensity in the case of no ice. A drawback with the method is that a disc thinner than 3 mm easily breaks apart.

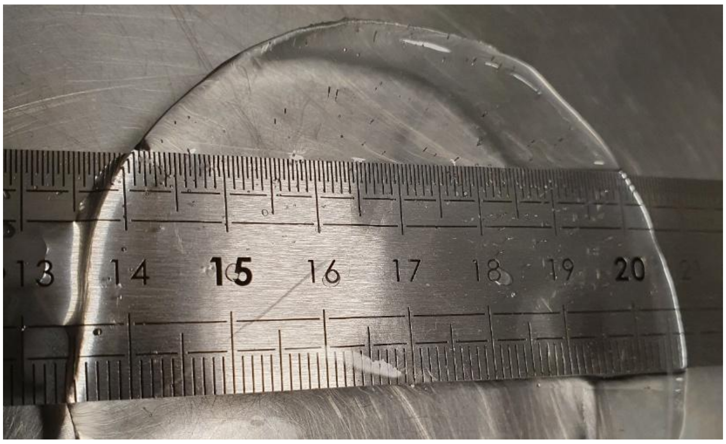

Figure 5. A disc of clear ice on a steel ruler. The surface is wet. A few air bubble impurities are seen.

Table 2. Ice in front of the sensor, from nothing to 9 mm thickness, dampens the returned intensity.

Material	Mean	StdDev
Nothing in the path	100	10
Ice, 2.5 mm	84	15
Ice, 4 mm	64	25
Ice, 9 mm	17	6

For comparison, the same experiment was repeated with two other typical materials. Protecting the sensor by an additional shield by ordinary silicate glass or by plastic glass lowers the mean intensity values over the target by about 10 percent as shown in Table 3. The silicate glass and plastic (styrene acrylonitrile) use here are not specially selected for high end optical applications.

Table 3. Remaining intensity in percentage of having nothing between the lidar and the target (first row).

Material	Mean	StdDev
Nothing in the path	100	5.6
Glass, 1.5mm	89	5.8
Plastic, 2 mm	93	5.8
Plastic, 2+2 = 4 mm	88	5.1

Comparing Table 2 with Table 3, it can be seen that the effect on dampening of clear ice is worse for comparable thickness. A main difference is that the standard deviation is higher. This can be due to impurities in the ice.

The result was expected. Damping can also be calculated based on material properties. Accumulation of ice can happen in very cold weather, possibly in combination with heating. A typical formation is icicles. An additional test was to place an icicle in front of the lidar. This had a small impact on the point cloud, apparently due to damping rather than refraction. The uncertainty of the experiment is high, for example the estimated thickness of the ice. In a real-world driving situation, the properties of the ice are even less certain. Still, the experiment was made to give a rough idea of what to expect from formation of ice.

3.3 Snow

The aim of this experiment is to investigate how the accumulation, following a natural snow fall, affects the lidar sensor, and to correlate the accumulation with the degradation of the point cloud. Clearly, it is possible to completely block the lidar and it is rendered useless unless cleaned, but it is the intermediate behaviour between no blocking and complete blocking which is of interest, where analysis is meaningful.

The setup for this experiment is described in Section 2. In short, a camera is used to estimate the amount of snow covering the aperture of the lidar, recording the change during a snowfall.

As an example from the experiment, the point clouds of Figure 6, which shows the experiment from a viewpoint a few meters above the lidar looking down on the targets and the wall, can be compared. There are two targets in the foreground casting shadows on the floor and on a larger target hanging on the wall in the background. The pictures show the respective amount of snow on the shield. The amount of snow on the sensor shield is also shown. The effect of higher coverage is fewer points in the cloud, a lower average returned intensity, and a range estimation which is less accurate. When the sensor is partly covered by snow, the foreground targets have a considerable variation in the direction of depth. The floor also gets difficult to detect when returned intensity drops. The difference is clearly visible for a human, however it is an open question how much this will affect any perception function. Eventually, when more snow is added, the point cloud will be null.

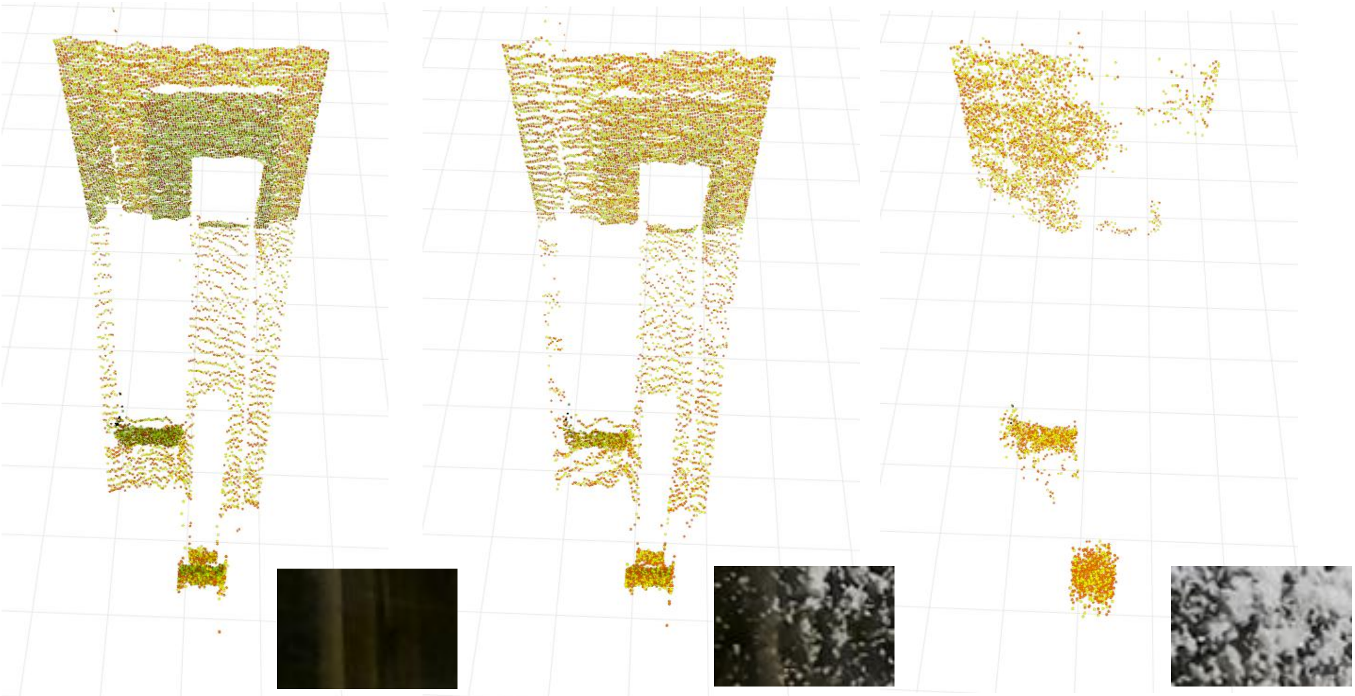

Figure 6. Example of point clouds from single scans, from no snow (left) over 2/5 to about 4/5 covered (right).

In order to make a correlation, the amount of snow needs to be estimated. It was noticed in early experiments that the thickness of the snow layer is not useful. The reason is that once the surface area is completely covered the returned intensity is negligible.

Instead, the coverage of the area was estimated by using a camera using the value of each pixel in the monochrome image. A histogram reveals the distribution for each image. A fully black image suggests no snow and in the other end, a fully white image would suggest a complete coverage of snow. For a black and white image the average greyness of the image is readily calculated. A histogram of the greyness is straightforward to produce and stabilize the value, the greyness is in this report taken to be the median value.

In Figure 7 a sequence starting from top is shown. The left column displays the shield, and the right column the intensity of the point cloud at the target (2D). A bright pixel means higher intensity than a dark pixel. Figure 8 shows a graph from the same sequence with estimates of greyness and lidar intensity as a function of time. This experiment was repeated three times and a typical duration was below 10 min. The actual time from top to bottom would depend on several factors such as the angle of the shield towards aerodynamic and gravitational forces, the adhesive properties of the snow and the surface, the temperature, the ability of snow flakes to entangle, etc. Once the snow flakes start to get stuck and entangle, the pace of accumulation increases.

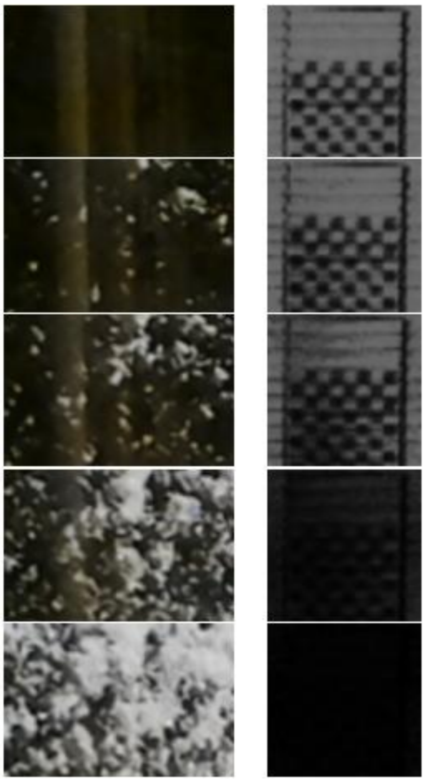

Figure 7. Snow blocks outgoing and returning signal. Left column: pictures of snow deposited on lidar's shield. Right column: 2D lidar point cloud intensity value of the checkered board. Black means low intensity. The more snow on the shield, the more damping of the lidar.

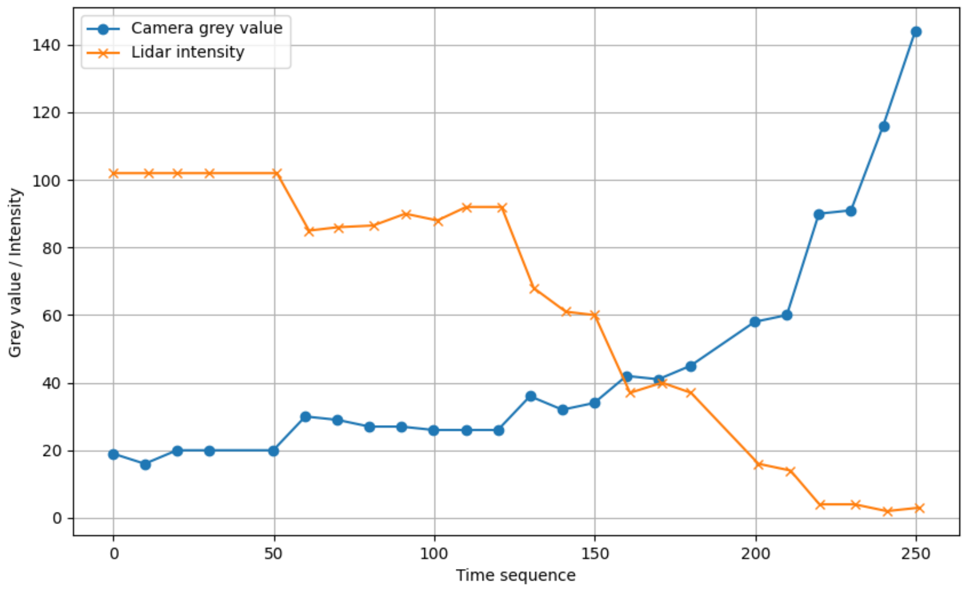

Figure 8. Returned intensity of lidar and greyness of image (in grey scale) taken over.

Figure 9 has another view of the experiment where time is removed, but with pictures of the lidar shield to develop an intuitive idea of the relation.

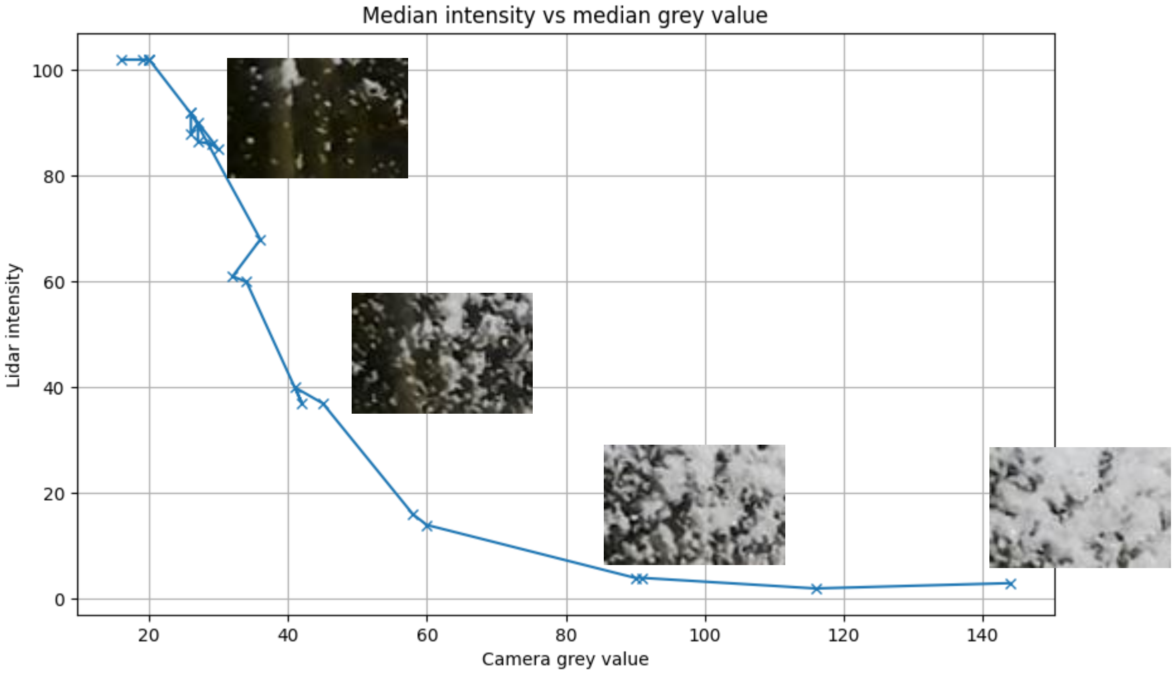

Figure 9. Greyness from camera versus intensity return of lidar.

3.4 In-situ analysis of blocking

This subsection is based on data from field test logging by AVL as described in 2.4. One aim is to determine how much the top mounted sensor is contaminated. A secondary aim is to understand the limitations in the data to enable such a statement. The analysis is based on a dataset where AVL's vehicle is driving in wet snowy conditions. First, a model and a set of assumptions are stated. Second, the algorithm is explained. Third, some results are presented that indicate degradation of the lidar due to buildup of snow. Finally, an alternative method is presented which releases one assumption of the main analysis.

3.4.1 Assumptions and models

Assumptions and simplifications:

1. The sensor is only looking forward. Implication: clogged snow at the front of the senor (that is, in the driving direction) is likely, but clogged snow at the rear of the sensor is less likely. A comparison of forward and rearward direction is a possible extension of the analysis.
2. The road vehicle sensor carrier has some flat area of road in front of it. Its optical properties are unknown, but possibly asphalt with contents of highly reflective lane markings.
3. Everything at either side of the road is opportunistic in nature. Implication: its uncertain presence and possible occlusion makes it not trustworthy to occur at a frequency high enough to be useful as an indicator. Similar reasoning applies to using other vehicles on the road.
4. There is no reference target. Implication: need to assume unknown properties, which in addition are measured from different angles and distances.

To conclude, analysis is based on a small section of the road ahead of the vehicle. For each lidar frame, the complete or partial non-detection of a road area could be due to a combination of (see also Figure 10):

1. Blockage of the sensor, for example snow on its surface
2. Blockage from aerosols, for example snow in the air
3. Occlusion from objects, for example a car in front
4. Reflection from wet asphalt, which depends on angles between laser ray and surface's normal. The asphalt can't be detected, since the laser beam is mirrored away from the sensor.

Non-detection can of course also mean that the lidar is looking into the sky directly and finds no object within maximum range, but it is assumed that the area of interest excludes this.

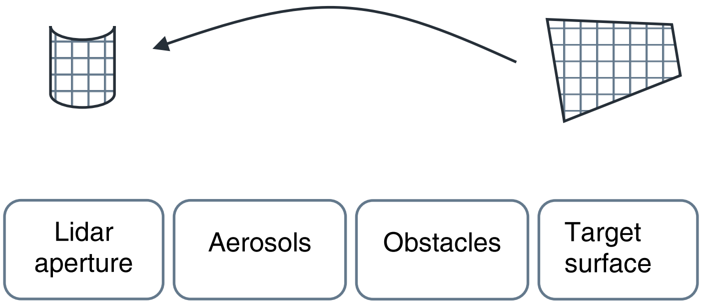

Figure 10. Mapping between sensor surface grid and target surface grid.

There is a chance to detect dynamic obstacles by the lidar and avoid analysis at instances of interference. There is also a chance to detect and possibly clean the lidar data from impact of aerosols like snow. The order of aerosols and obstacles is of less interest. However, it is not possible to avoid the absorption and reflection of the target surface. A typical situation is wet asphalt, where increased water turns the asphalt to become specular from an initial more diffuse reflection, resulting in less returns to the sensor. If the target is beside the road, vertical targets can be chosen instead. The target surface will be seen from different angles and distances.

3.4.2 Analysis data structure

The main analysis is to compare the data from a part on, or along, the road as seen from the sensor at any time. Choose a specific part of the road or large flat object such as traffic sign or wall. The size could be a few meters. The process of selection can be automated. To start with, the region is constructed as a voxel to accumulate data for computation and allow a flat grid to be selected from the set of voxels. Voxel is representation of a value in 3D grid in analogy with pixel representation in 2D grid. Data is collected from the first hit of the voxels to the last, which mean that the vehicle has passed the area. The size of each voxel should be selected such that multiple hits can be captured. Since the beam exit size is quite large it isn't necessary to check for every represented azimuth angle, but the grid in azimuth direction can be a number of degrees, e.g. 5°.

The number of voxels in each direction should be chosen to fill an area of the road or other type of target. The voxel is turned into a grid that represents the road, and a grid that represents the surface of the sensor is also constructed, see the pseudo code in Figure 12. There will be a many-to-many mapping between the grid that represents the road and the grid that represents the surface of the sensor, as illustrated in Figure 11. In the figure, the concentration of scanning lines towards the center of the sensor is visible in the bottom part (azimuth and elevation dimension or s and t coordinates) and there are few hits above the horizon. The road area covers in this case also some part of the structures beside the road. It is also possible to see the road in the upper part (x and y dimension), which reveals that the central part of the road does not return any signals.

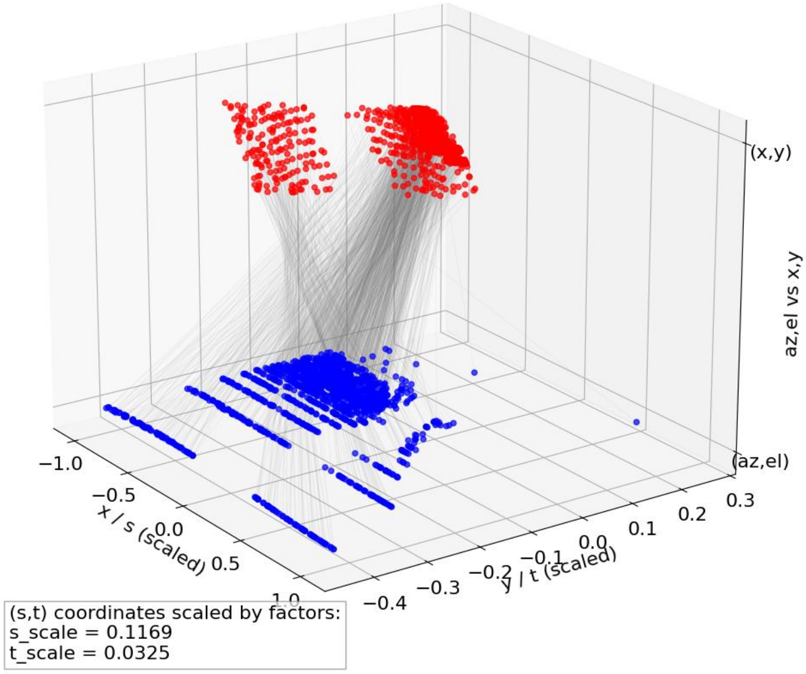

Figure 11. Illustration of many-to-many mapping between a sensor aperture grid (lower, blue) and a road grid (upper, red) with black lines. The scaling is just to allow for plotting in the same diagram.

```
# 1. Define region of interest and voxel grid parameters
set_voxel_grid(volume_of_interest, voxel_size, cell_size)
# 2. Prune and voxelize the point cloud
for point in point_cloud:
    if point in volume_of_interest:
        assign_point_to_voxel_based_on(x, y, z)
# 3. For each voxel, collect intensity, distance, azimuth, elevation
for voxel in voxel_grid:
    store voxel_data = {intensity, distance, azimuth, elevation}
# 4. Create aperture and target grids in (azimuth, elevation) space
aperture_grid = Gridifyer(...)
target_grid = Gridifyer(...)
# 5. Build mappings between target voxels and aperture cells
target_to_aperture_map = find_connecting_keys(aperture_grid)
aperture_to_target_map = reverse_connecting_keys(target_to_aperture_map)
# 6. Predict point cloud at a given position using voxel and mapping data
predicted_pcl = predict_full_pointcloud_at_spot_start(aperture_to_target_map, aperture_grid)
```

Figure 12. Pseudo algorithm for the main method.

An idea is to compare the frame returned by the sensor at the entrance to the spot area (at the time T0) with what could have been seen given all information of the spot area (also from time after T0). The difference between these two grids within the target case as well as the aperture case will reveal information. For the prediction of the target grid there are at least two methods:

- Method A: Collect statistics for the grids, of interest is the azimuth and elevation angles that both sweep the aperture area. Prediction of the frame at T0 is based on random sampling of collected samples.
- Method B: Prediction by sensor modelling in order to find a prediction of the grid at time T0. The sensor model is only geometric. The idea is to use it to project data from the voxel grid to what a single instance at T0 would see of that spot.

Method B is to prefer since it will have higher resemblance with a typical single scan. In the case of a horizontal road, the success depends on how well the azimuth area is swept; the elevation angle will be swept by passing the area and have a high correlation to the distance. This is determined by the motion of the sensor (which is fixed to the vehicle) and this is not controllable. In case of a vertical area, for example a road sign, the angels are swept complementary.

The range detection algorithm from WP5 [13] converts point cloud to mesh in order to connect points into areas and reports on objects furthest away (as visibility range in meters) in azimuth sectors, not necessarily with any temporal behaviour. A main difference to the analysis presented here is the definition of a static spot in the road to collect all data possible, then perform the analysis on what part of the lidar's aperture seems not to be in use.

3.4.3 Results

In many of the examples below, for ease of understanding, the target area has been wide to allow for features along the road side to show on pictures and diagrams. In order to simulate a forward looking lidar, the FoV has been limited to +/- 60 degrees in the azimuth dimension, where zero means straight ahead.

In good weather conditions an example target grid and aperture grid is shown in Figure 13. The axes unit is grid cell. The cell values are capped at 50 to zoom in better. A dark point in the grid represents more hits of that cell. The discontinuity of the lowest layer of the aperture grid is because it has been removed from the data set since it constantly hits the hood of the vehicle. The vehicle took a left turn prior to the target area and this is visible in the aperture view which is "smeared out" in the x-dimension. It is clearly visible in the figure that the rotating lidar is designed to concentrate on the horizon, with sparse coverage of the ground and ever sparser of the sky.

From the data in Figure 13, predicted point clouds are shown in Figure 14. The axes units are meters. Blue on the ground means predicted point cloud (it is actually intensity coded) and red is the point cloud from the sensor at the entrance of the target area, at the location represented by the camera image. The typical rings from the lidar on flat ground are not perfectly recreated.

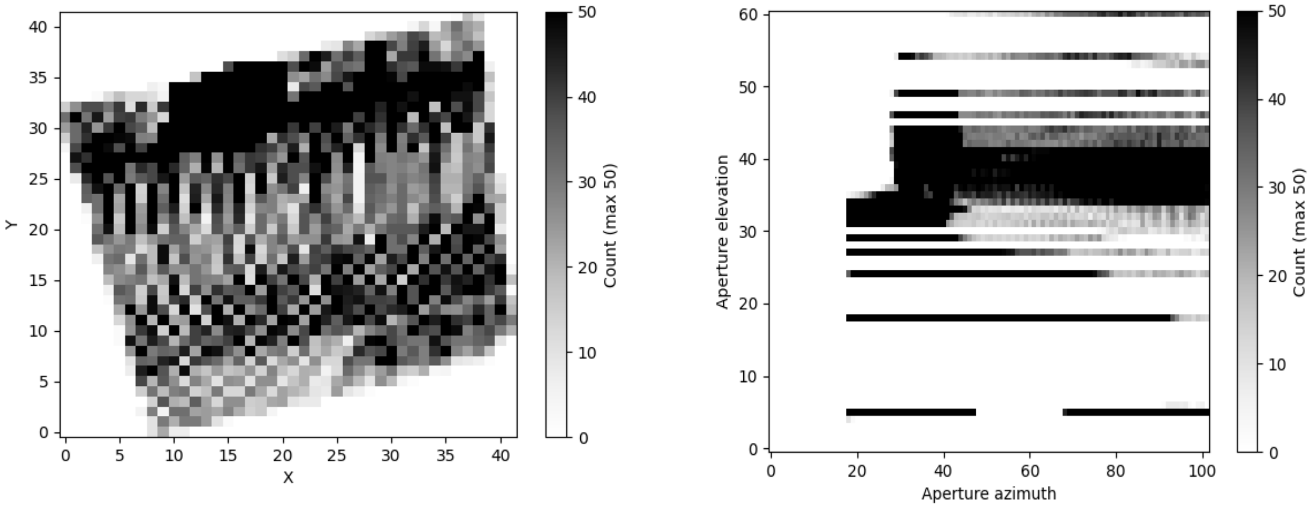

Figure 13. Target (left) and aperture (right) grids. Black means many hits. Clear weather, dry ground, low speed.

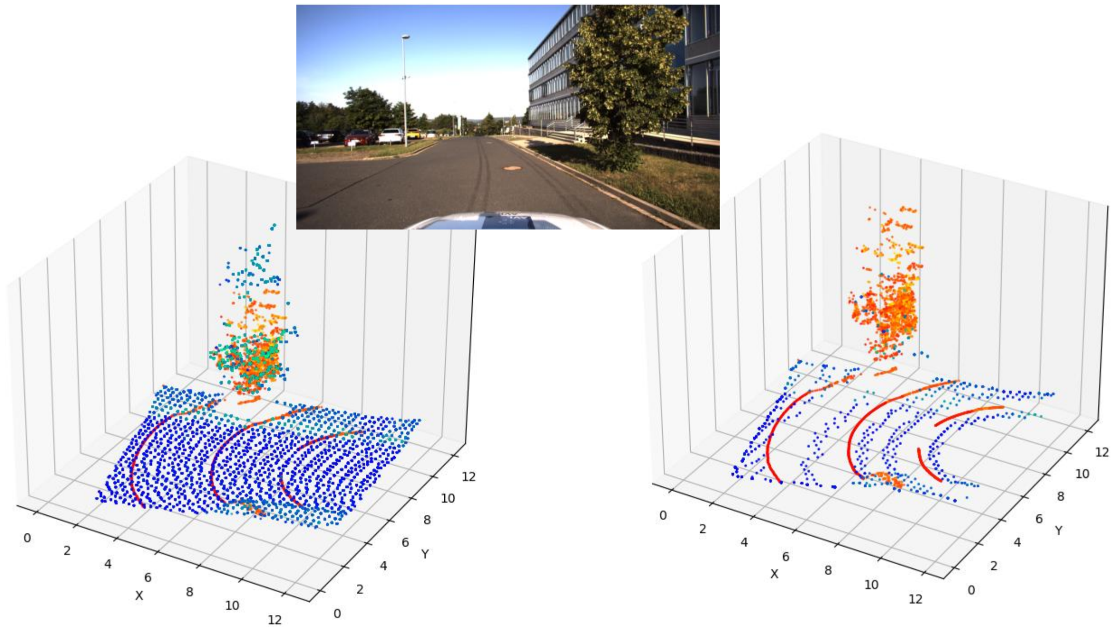

Figure 14. Lumped (left) and single (right) point clouds. Blue/green points are predicted, red points are measured. Clear weather, dry ground, low speed.

From the initial good weather, the contrast to snowy weather is shown in Figure 15 and Figure 16, where the lidar sensor is contaminated but arguably still quite useful. The total number of filled cells in the grids is lower due to the higher speed. It is possible to identify the gap in the aperture grid around its center line at higher elevations (longer distances). The camera image gives a hint of the current contamination. The reconstructed point clouds as well as the point cloud which coincides in time with the image, all have considerably less number of points due to higher speed and the non-cooperative weather. Still, it is possible to see most cells of the target grid and reconstruct (predict) a point cloud at the entrance of the target grid area to compare with the measured one (red colour in Figure 16).

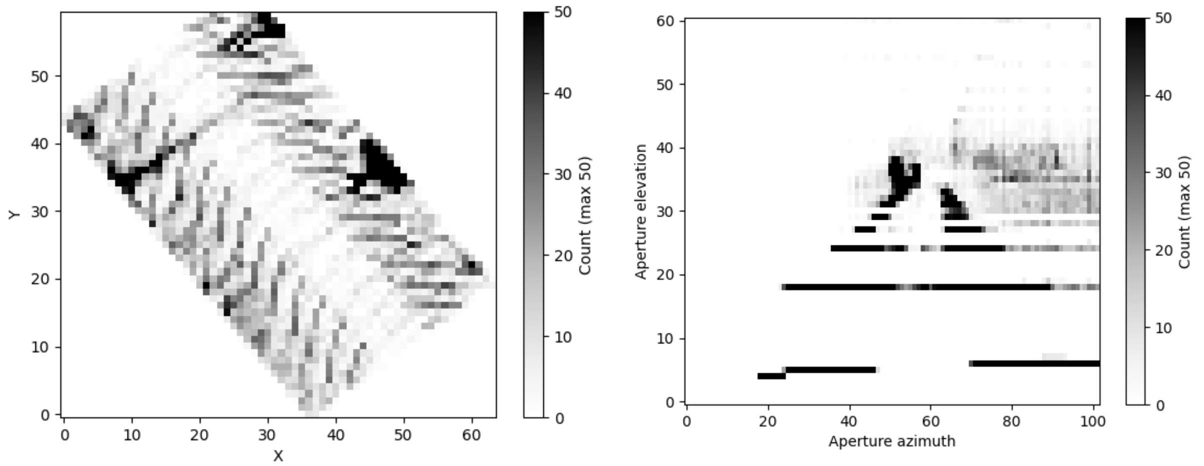

Figure 15. Target (left) and aperture (right) grids. Snowy weather, wet ground, moderate speed, contaminated.

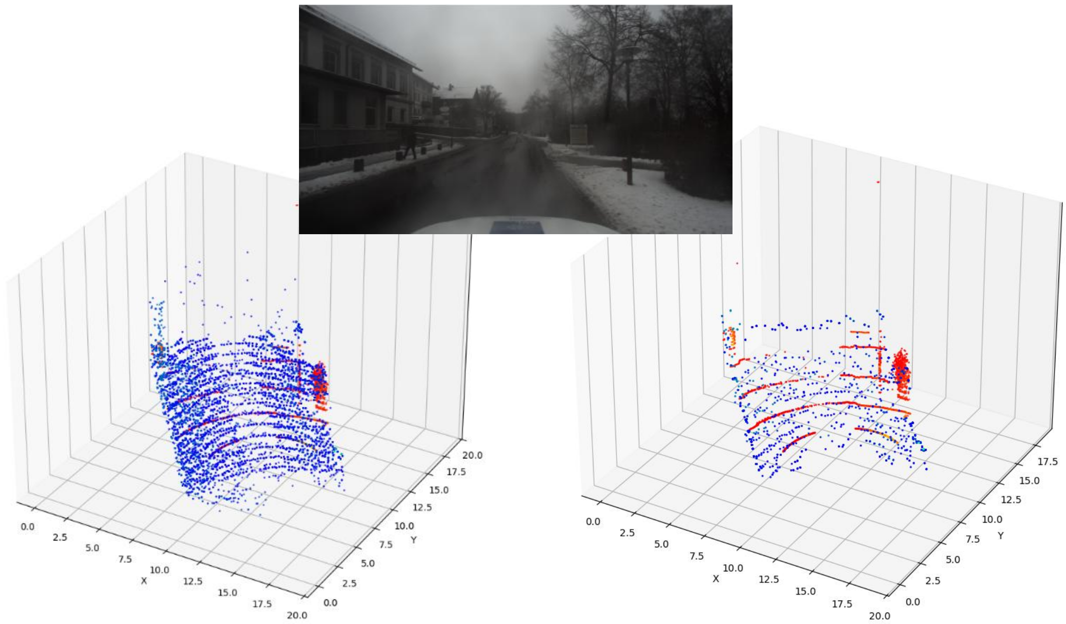

Figure 16. Lumped (left) and single (right) point clouds. Blue/green points are predicted, red points are measured. Snowy weather, wet ground, moderate speed, contaminated.

An even more severe contamination is found in Figure 17 and Figure 18. The lidar does never see a strip of the road in front. There are definitely fewer hits in the grid but the speed is slightly higher. Less than a minute later the output from the lidar in the front direction is completely blocked.

The reason for the lack of response in the middle of the aperture is not simply because that part is contaminated. Another good explanation is the specular surface of the wet asphalt, from which few photons are reflected back to the sensor. The transition between diffuse (Lambertian) and specular depends on the amount of water which is unknown. In practice it is a combination of blocking (absorption) and specular reflection.

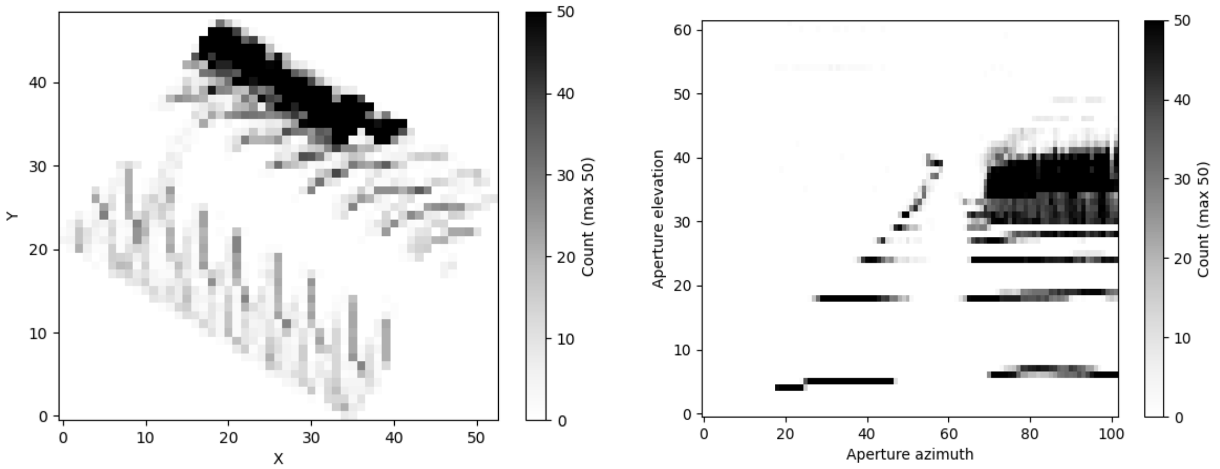

Figure 17. Target (left) and aperture (right) grids. Snowy weather, wet ground, moderate speed, highly contaminated.

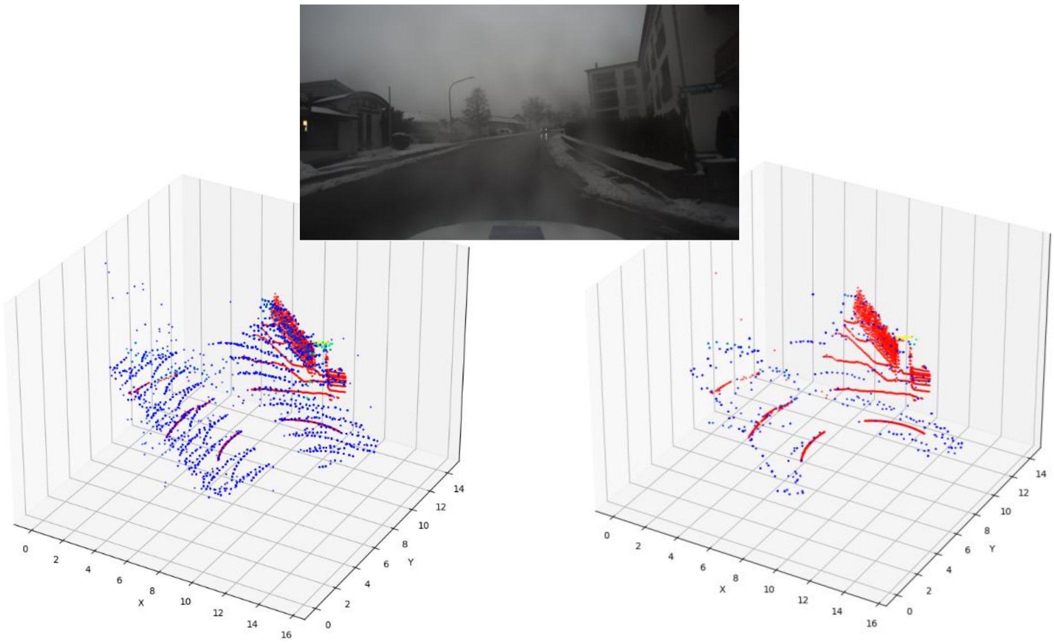

Figure 18. Lumped (left) and single (right) point clouds. Blue/green points are predicted, red points are measured. Snowy weather, wet ground, moderate speed, highly contaminated.

The difference in aperture maps of Figure 13, Figure 15 and Figure 17 is useful to judge the situation. They clearly show the maximum range over azimuth. However, it is not straight forward to discern between the remaining two explanations (blockage and specular reflection). The absence of hits at high elevation angles can suggest that the lidar is contaminated, but it can also be too long a distance to the nearest object. One way to make it possible is to make sure, from the aperture grid's viewpoint, there is a larger variation of where signals come from. This is not very practical since it means moving the sensor or vehicle relative to the ground, or possibly use dynamic objects as targets.

3.4.4 Alternative: approach with a reference target – the hood

In the data set used, the lowest layer of the lidar falls upon the hood of the vehicle. The hood might not be an ideal target or located in an ideal angle to the target, but serves as a well-known ever-present target, and it releases the previous assumption of having to access to a reference target.

The choice of indicator is the number of points returned and the average intensity returned. It would be possible to define other indicators also, for example the sum of returned intensity. Calibration assumes absence of contamination.

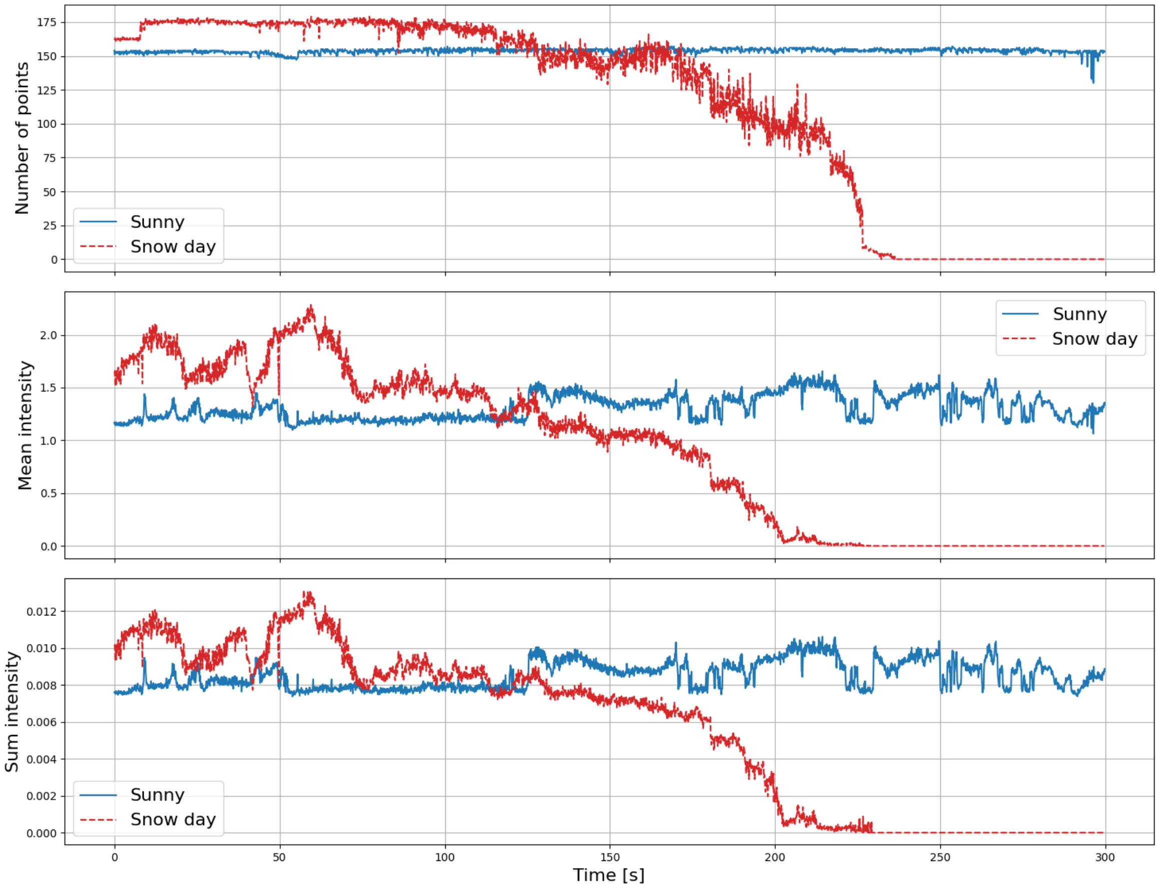

Figure 19. Indicators for sunny day (solid blue) and for snowy day (dashed red).

Figure 19 shows number of points, average intensity, and sum of intensity per frame for both clear and snowy weather. As the buildup of snow increases, the fluctuation also increases for the declining curve describing how many points are echoed from the target. After three minutes the lidar ceases to be useful at this short distance. The time-series from the two occasions (recorded 5 months apart) are not normalized, and for unknown reasons (such as snow on the hood, or slight dislocation of sensor because of remounting it) the number of initial hits is lower in clear weather.

The short distance from target to sensor means less opportunity for aerosols in the air to interfere so the result is likely to depend more on contamination of the sensor or of the hood. In case of snow, as the wet snow melts it typically leaves via the lower part of the sensor's aperture. The hood covers an important part of the FoV, but a drawback is the full FoV isn't covered. Another complication with this approach of a reference target is that the hood might have contaminants such as snow, either building up simultaneously, or from the start. In the main method above, depending on the environment, it cannot be guaranteed that the complete FoV is covered either. In summary, seeing a part of the hood is definitely valuable as an indicator of contamination.

4 Impact on the radar

The impact on radar performance of water, ice, and snow contamination on the radome can roughly be said to depend on the shape, thickness, and electrical permittivity (at frequency of interest) of the contamination. Assuming the thickness of the contamination can be said to be even over the radome, we can simply inspect estimated changes in absorption and reflection of the radar signal based on the permittivity and thickness. The changes we have identified are

- Absorption of radar energy as it passes through the contamination layer
- Reflection of radar energy at the interfaces radome/contamination and contamination/air

If the contamination is homogeneous (compared to the wavelength ~4 mm) an approximation for the expected change in return signal can easily be modelled and calculated through calculating the reflectivity R at the different interfaces

$$R = \left|\frac{n_1 - n_2}{n_1 + n_2}\right|^2,$$

where n_1, n_2 are the complex refractive indices of the two materials of an interface (e.g. radome/snow or snow/air), and Attenuation Loss α

$$\alpha = e^{-\delta \cdot k \cdot D},$$

Where δ is the loss factor (which can be calculated from the more commonly given tanδ), k is the wavenumber, D is material thickness.

4.1 Water

At 77 GHz (20 °C) water has a complex permittivity around $\varepsilon_{water} = 12 - 25\mathrm{i}$ [19]. This means the loss tangent will be fairly high which could be problematic for radar due to the high absorption. It is however unlikely that the buildup of water on its own on the radome will be thick enough for it to cause issues even though there is work suggesting otherwise [20]. A MATLAB script for calculating the reflection and attenuation losses from water/ice/air mixed contaminations of uniform thickness was developed (planned to be published on the ROADVIEW github). Calculating for only a thin water layer shows that it is likely there may indeed be some losses, but if the water layer is thin the main part of the losses will be from increased reflection, while the attenuation losses are small as anticipated.

4.2 Ice

Ice has a complex permittivity around $\varepsilon_{ice} = 3.15 - 0.005\mathrm{i}$ at 77 GHz (-15 °C) [19]. The loss tangent is thus fairly small and absorption might not be a large issue. The permittivity is also quite close to what could be considered normal for the typical plastics found in radomes which would thus mean that reflections in that interface might not cause any greater issue.

4.3 Snow

Snow can roughly be described as a three phase system of water, ice, and air in different proportions. In worst case, this could thus result in a thick contaminating layer with high loss tangent. Snow packing on the radome could thus cause attenuation of the radar signal and obscure objects from detection.

There are methods of estimating the effective permittivity of the contamination based on the permittivities of its constituents and their respective volume ratios (example planned to be published on the ROADVIEW github).

The MATLAB models suggest that wet snow may indeed prove very problematic, should it build up on the radome, causing severe degradation already at a few mm of buildup based on values of “wet snow” given by [21].

The mechanisms and sequence of events leading to build up of snow on the radome has not been investigated in this project, but anecdotal information suggests that after a first seeding layer of snow has attached the buildup increases rapidly, and that a warm radome may even increase the risk of creating such a seeding layer.

Calculations of effective snow permittivity for different types of snow could be valuable input to simulations. In practice it is however possible that it is enough to identify the problematic case and find a way to identify when the radar is affected, regardless of what the contamination is.

4.4 In situ analysis of blocking

4.4.1 Reasoning

Blockage detection as well as rough diagnostic quantification of the reduction in radar visibility could possibly be continuously evaluated using statistic of the detections registered by the radar over time. One identified issue with this strategy is that the environment and detectable objects around the ego-vehicle will be constantly changing and can vary extensively. In all drivable scenarios there should however also be some surface in front of the vehicle. Visibility of the road surface just in front of the vehicle is thus a candidate as an indicator of visibility in general. This would demand the visibility of the road surface itself not to change significantly depending on its constitution, e.g. asphalt, gravel, snow, or ice. Studies have previously been done on using the difference in reflectivity for different polarizations to determine road conditions. From the results of such studies it appears unlikely that the total reflectivity of the road would change more than a few dB due to changes in road surface constitution [22], [23], [24], [25]. Note that the frequencies used in these studies are slightly different than what is typically used for automotive radars and that the backscatter from a snowy surface could be higher than smooth concrete and asphalt, thus somewhat increasing the risk of underestimating the level of blockage.

As mentioned in section 1.3. the idea of radar blockage detection is not new but previous attempts leave room for improvement.

With the introduction of 4D radars with a wider beam in elevation, the radar beam can strike the ground at a much higher angle than would be possible for the older more narrow beam radars. The higher angle should increase the backscattered proportion of the radar signal hitting the ground. The road surface will also be visible at a shorter range, for the same reason, even further increasing the amount of ground scattered radar energy that will make its way back to the sensor. The elevation information and improved azimuth resolution may also provide an opportunity to distinguish the road surface from other clutter, at the very least on a statistical level.

4.4.2 Experimental setup

Blockage was tested with the project 4D radar in a stationary state on a parking lot at RISE in Borås, Sweden(Figure 21). Different blocking states was produced by applying layers of wet tissue paper on the radar surface and freezing them (Figure 20) as an effort to emulate different types and thicknesses of snow, as the tissue paper is able to hold the water content in place even as the emulated snow is thawing. This method is similar to what has previously been done in other work [26].

The wet tissue paper was applied one layer at a time after which the entire radar was placed in the freezer for at least one hour to make sure the water was thoroughly frozen. The first measurement was performed without any contamination to the radar surface. At each measurement time the radar was mounted and adjusted/aligned as quickly as possible to try to avoid thawing before the start of the measurements. This proved quite difficult as the ambient temperature during measurements was well above freezing. Data was collected for at least 2 minutes at each measurement. The ambient temperature nor the contamination temperature was noted during the measurements, which might make it difficult to pinpoint the main reason for the signal loss through the contamination, but as has been mentioned, it may be of less concern than the ability to estimate the loss.

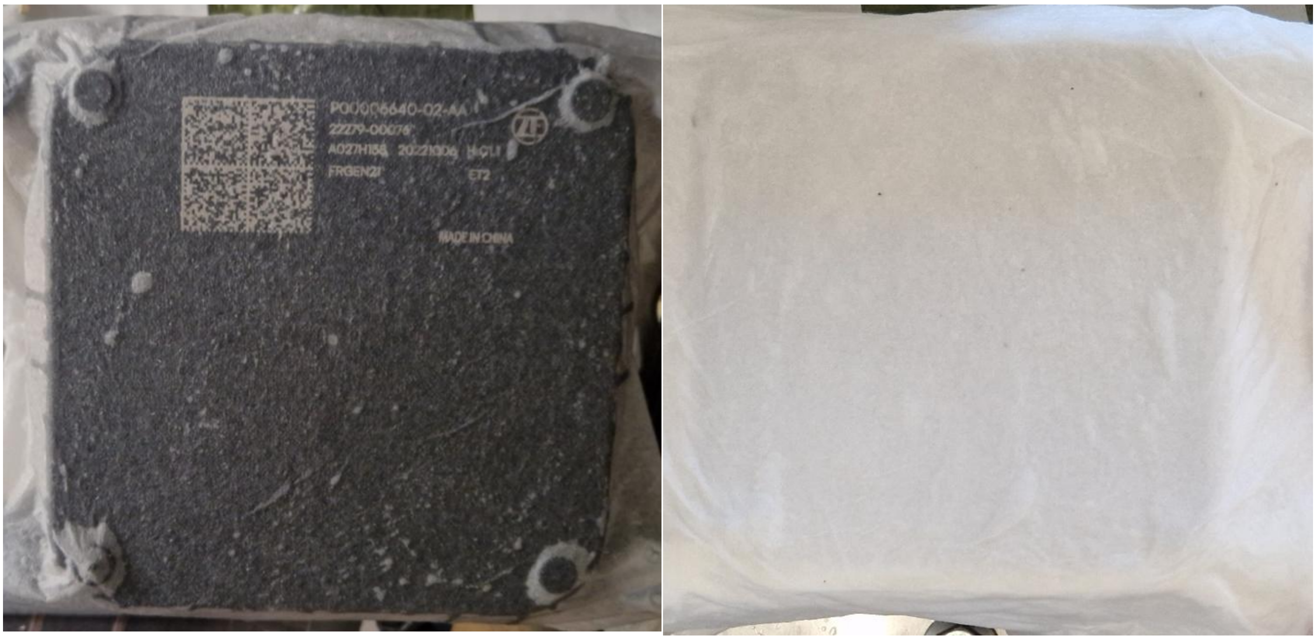

Figure 20. Radar with one layer (left) and several layers (right) of frozen wet tissue.

Figure 21. Radar mounted on stand at parking lot. The two reference corner reflector targets can be seen further away on the flat asphalt space.

Two separate corner reflectors (henceforth referred to as CR1 and CR2) with well known radar cross sections were placed as reference targets working as ground truth indicators of the radar signal loss due to the emulated contamination. As there are no other distinct target in the vicinity of the reference targets their corresponding detections could be easily identified. The radar was mounted at 0.5 m above ground and adjusted to look straight using a level. Detections with range 0.5 - 9 m from the radar were chosen as being of interest. They were further filtered making sure that the detections used were indeed reported as close to ground (within 1 dm from known ground level). This thus gives us an indicator surface. Four different measures were suggested:

1. Number of detections reported from the indicator surface (ground_N)
2. Median of the RCS of the detections in the indicator surface (rcs_ground)
3. Maximum of the RCS of the detections in the indicator surface (which would come closer and closer to the median as the number of detections deteriorates) (rcs_ground_max)
4. Total sum of RCS of the indicator surface (summation must be performed linearly) (rcs_ground_tot)

4.4.3 Results

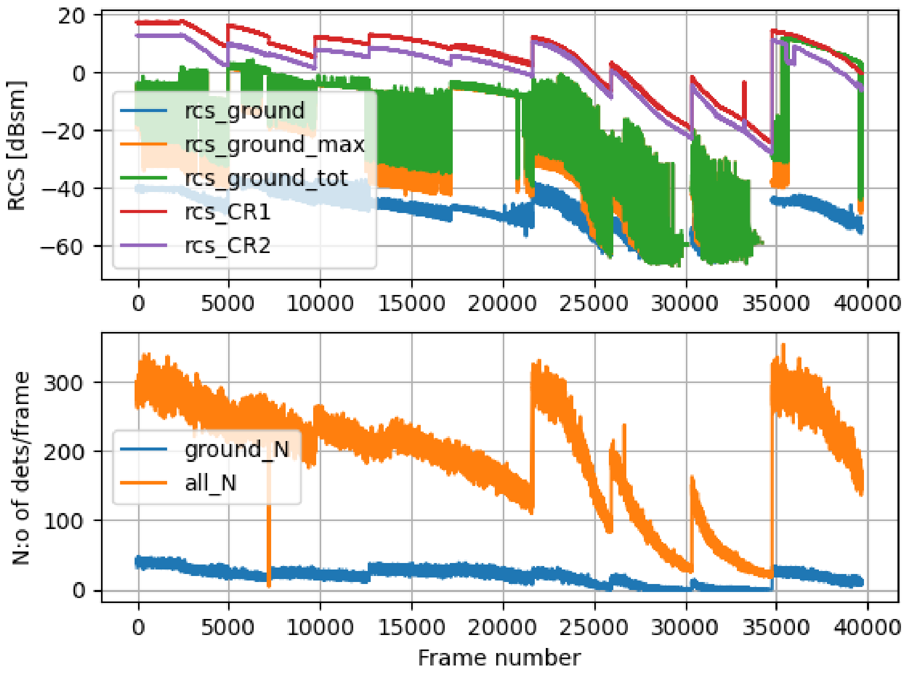

Figure 22. Data for blockage estimation. Top shows the different RCS values and bottom shows number of detections in in entire frame as well as specifically in ground indicator surface. Horizontal axis is number of logged radar frames from the start of the first measurement (frame update frequency is stated by the sensor manufacturer as 20 Hz). The reported RCS for the corner reflectors are rcs_CR1 and rcs_CR2.

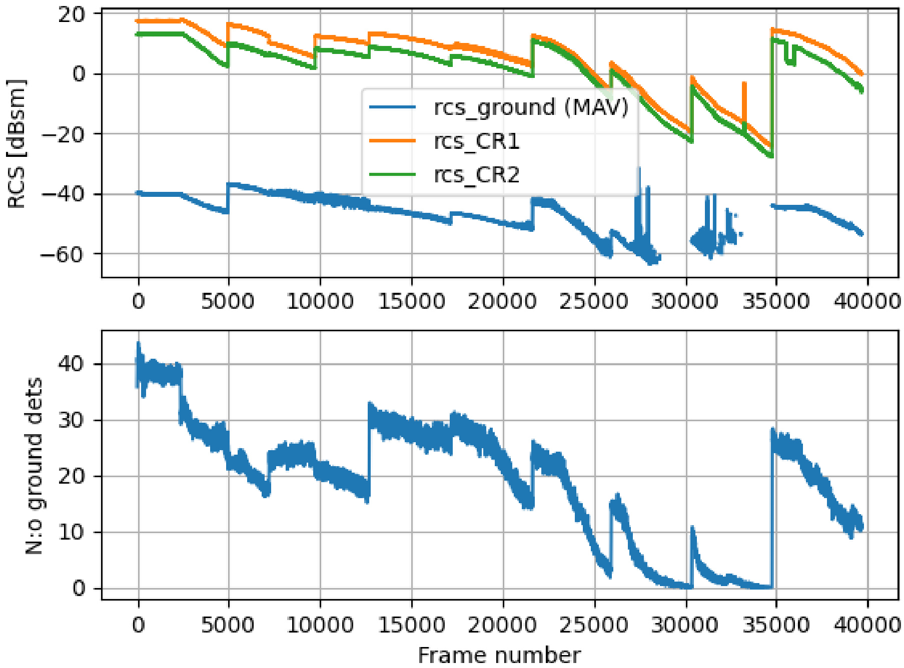

Figure 23. Same data as in Figure 22, but with only the two most promising indicator values together with the corner reflector reported RCS values. Here the indicator value curves have also been smoothed using a moving average window of 11 points (corresponding to 0.5 s).

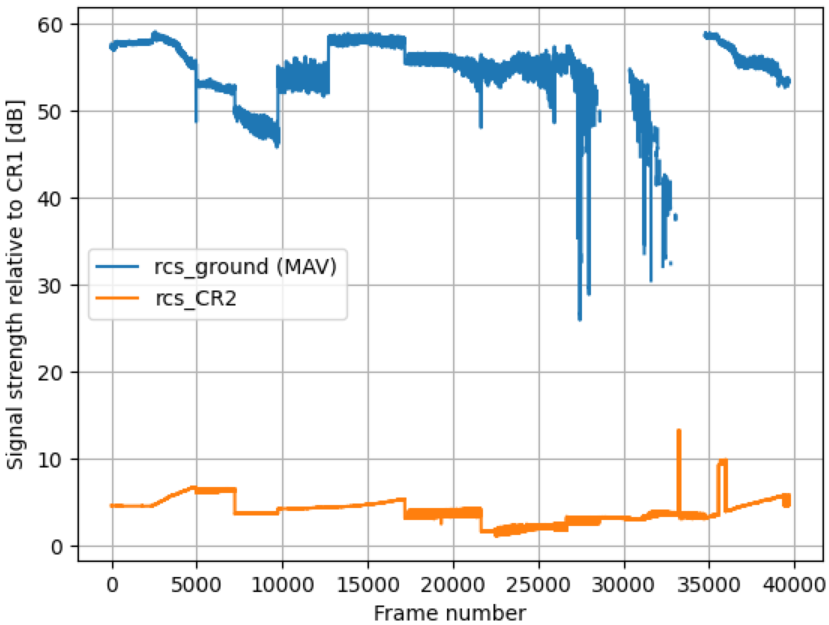

Figure 24. Reported signal strength relative to that of CR1, i.e. difference between the CR1 reference reflector and the ground median indicator value (blue) as well as the second reference reflector CR2 (orange).

Jumps in the reported RCS can be seen in the time series. They correspond to the timepoints where a paper tissue layer has been added and a new measurement has been started. It is clear from Figure 23 that the progressive thawing of the contamination makes the blocking of the radar worse and that re-freezing gets the radar back to a less signal impeded state, even if more material/frozen water has been added. The jumps in the curves could possibly partially also be attributed to discrepancies in mounting after the radar has been moved for the addition of a layer, but as can be seen in Figure 24 the corresponding jumps in RCS of the well known reference targets relative to the suggested ground signal are much less pronounced, especially in the cases where there is a clear ground signal.

The first thing to notice from the results is that the reported RCS of the corner reflectors is not perfect but both of the reflectors follow the same trend and their relative difference appears correct and more or less constant. From the results it can be seen that even though all measures appear to have some correlation with the degradation of the reported RCS of the corner reflectors the total and max values appear less reliable. It is possible that a temporal moving average filter could be applied to aid in getting a more stable measure from these figures. The median values (rcs_ground) and number of detectors (ground_N) appear to be better candidates as indicators up until there starts to appear radar frames in which 0 detections are reported. This happens at a point where the reported RCS from the corner reflectors has been reduced by ~23 dB which corresponds to a reduction of the radar visibility (maximum range where an object of a specific RCS could theoretically be detected) to ~26% of the clean sensor. The sensor is thus not completely blocked but its performance severely impeded.

Excluding the noisiest points, the variation in reported RCS for one corner reflector relative to the other appear to be about 5 dB. For the median RCS of the ground detections the variation is unsurprisingly worse, but still appears to only be about 10 dB (see Figure 24). Deviations beyond this would thus likely need to be taken seriously even if the indicator value has a large uncertainty. A longer moving average time window may be useful for further reduction of the uncertainty. The largest uncertainty is of course when there are almost no detections from the ground reported (close to 0 ground detection counts). Caution should of course be taken long before this happens as the lack of enough indicator detections is an indicator in itself and will, for obvious reasons, not provide enough data for a reliable statistics.

It should be noted that it is obvious from e.g. [15] that for a forward looking radar the ground reflection will appear to have a velocity dependent on the ego-vehicle velocity. This aspect is missed when only inspecting data from a stationary radar. The typical velocity of the ground reflection can be determined as a function of ego-vehicle forward velocity (v_{EGO}) as well as elevation (φ) and azimuth (θ) angles. The expected elevation angle can in turn easily be calculated from the range (r) if the radar mounting height above the road (h) is known.

$$\varphi_{road} \cong \tan^{-1}(-h/r)$$

And the expected detectable velocity from the road would then be $v_{road} = v_{EGO} \cos\theta \cos\varphi_{road}$, which could possibly be simplified to $v_{road} = v_{EGO} \cos\varphi_{road}$ if assuming we are mostly interested in inspecting the signal from the road close to the center of field of view in azimuth.

A limitation to the suggested approach lies in the difficulty to know the exact data processing that leads to the detections reported by the radar. If there is additional filtering, data transformations, or frame-to-frame processing of the detection points this may inadvertently corrupt the detection data in a way that makes the suggested strategy for blockage detection difficult. Some speculative examples of such corruptions are:;

- If the detections from the ground are suspected to induce false alarms or make the data appear messy they may be actively filtered out.
- The detections may be transformed into a ego-vehicle coordinate system which could make some predictions a little bit more difficult.
- The velocity of the detections could utilize frame-to-frame calculations using tracking to try to translate the detectable radial component to more generally familiar cartesian components.

5 Discussion and conclusions

5.1 Methodology

The aim of this deliverable has been to look at contamination on sensor shields. Two sensors have been addressed, lidar and radar. These have similarities such as that they both fundamentally transmit and measure the return of electromagnetic waves and that the results often can be represented as point clouds. The way they transmit does however have dissimilarities. The lidar transmits laser beams with well known directions and reports the return from each beam, while automotive radars typically transmit a much wider beam and is, due to the longer wavelength, instead able to determine the direction of the return through e.g. digital beam forming. The radar can then produce heatmaps in its detection space of returned energy from which peaks can be identified and reported as detection points. The ability to determine phase differences in the radar is also used to include radial velocity as a distinct feature/dimension in its detection space. This difference in working principle has lead us to slightly separate approaches for characterizing the effect of contamination on the surface of the two sensor types.

The spinning mechanical lidar we have used has a rather large aperture which means that a scratch or a small mark from for example one very small bug on the aperture lowers the return intensity rather than blocking parts of the resulting point cloud. For the radar, the main focus in this work has been on the forward looking radar, which is assumed to collect contaminants (to a degree where it could have a significant effect) fairly evenly over the sensor surface, as there is no real reason to suspect otherwise. For side looking radars this assumption should perhaps be re-evaluated as there is reason to believe that the contamination build up may differ more along a side that is not orthogonal to the direction of travel.

The approach has been to try to isolate effects of contaminants so that they can be understood separate from effects of aerosols in the media, and effects of precipitation deposited on targets and environment. However, it is difficult to study ice and snow on the sensor as a completely isolated phenomena.

In lab tests the reference target is typically constant throughout the test and possibly also well characterized and calibrated. This is never a realistic tool in a real-world driving scenario: what is the potential to detect contamination when there is no reference target to compare with? It is advantageous to have a known object for comparison over time, for example a part of the vehicle itself. Knowing the current capability of the sensor is vital for the perception.

For experiments to be repeatable, characterization of the contaminant must be performed by other means than by the lidar under test. The chain has two steps of information: contamination of the aperture decrease the number of points and intensity from the sensor which in turn affects the task and output of the perception. Since measuring the contaminant or aerosol is cumbersome, a totally different approach, as seen in the literature, is to verify the output of the perception, which might be trained on a dataset including adverse weather.

5.2 Observations for design of sensor diagnostics

A build-up of snow on the side facing the driving direction is characteristic. The snow covers a large contiguous portion of the aperture. Wet snow is particularly prone to stick, and its water content attenuate the radar. The snow reflects the wavelengths the lidar is using. For a lidar, this results in loss of measured points possible in a contiguous sector, basically having no vision in the important forward direction. Data from field test shows that the build-up, that obviously depends on the local weather situation and the design of the vehicle, can happen in the course of a few minutes. The lidar can go from functional to having a blocked field of view (FoV) in a short time. The radar is also degraded, but more gradually, as the signal attenuation is largely related to the thickness of the contamination layer. This results in a decreased return power leading to poor signal to noise ratio (SNR), the radar underestimating the radar cross section (RCS) of objects in its surrounding, and makes it more and more difficult to detect low reflectivity targets, effectively limiting the functional FoV of the radar. A method of using the ever-present signal of the ground just in front of the vehicle/radar as an indicator for quantification of the degree of blockage is suggested. This information can be useful to determine the reliable FoV of the radar, i.e., the part of the radar FoV where it can be trusted that the typical radar targets can be detected despite the contamination. Final determination of reliable FoV is not part of this work.

A camera to record the coverage of snow on the lidar aperture works well to predict the degradation in returned intensity. For the lidar, a method suggested here is to study piecewise global fixed flat surfaces in the surrounding such as a piece of road preferably on the vehicles path, or a piece of a wall, or traffic sign. The situation is realistic, that is no lab target object with known properties is used. At any location of the lidar, a non-returned transmission of the active sensor can be due to objects obscuring the fixed surface, the fixed surface being out of the field-of-view, reflections such as water on the asphalt, or contamination of the aperture of the sensor. A valuable alternative method is to look at an object that is always present, such as a part of the hood.

The more realistic setting is to figure out contamination without a known object and with a background that is unknown and changes all the time. At the same time, the motion helps activate a larger part of the sensor's surface, which signals that the area is not completely blocked. In general, detecting blocking by means of using returned signal is complicated by the fact that an omitted return can always mean that there is no object in that direction – there is nothing to be seen. The lack of activation of the complete FoV is a generic problem, that might call for a multi-modal approach. Also when leaving the lab, estimating the amount of disturbance on the sensor by other means than the sensor itself should also be a good design.

For the radar the attenuation of the signal by an aerosol such as rain, snow, or fog is less likely to be significant, at least at near range. This leads us to believe that using near range objects as indicators might still be a viable approach if there are such objects or features that are always visible to the radar. The road surface has been identified as a candidate to function as such an indicator feature. Using 4D radars with a wide FoV in elevation further increases both the signal strength from, as well as the ability to correctly identify, the road surface. In a sense it can be seen as irrelevant what constitutes the contamination and that it is more important to quantify the effect of it and understand how that in turn affects the functions which rely on input from the sensor. This notion speaks strongly for the use of an indicator signal, such as the road surface. It can however still be good to continue understanding different processes of sensor contamination to e.g. be able to produce relevant scenarios, in terms of attenuation and contamination progression, for simulation. It should also be pointed out that the placement and elevation FoV of each specific radar will affect the returned radar signal from the ground. This means that even if the principle may be general, this type of blockage detection will need to be part of the design process and some knowledge of the specific sensor used is needed for proper implementation.

For self-driving vehicles it is an absolute must to monitor and understand its capabilities or limitations. Looking at Figure 19 and comparing with results from the main analysis, a diagnostic algorithm for the lidar only should pick up the troubles latest at about 120 seconds. There is room to improve test methodologies for individual sensors and sensor systems. The contamination can build up rather quickly for example in the case of wet snow. The interesting part to work on is when the contaminants affect the measurement but do not block the sensor completely. At this stage it could be possible to pursue operation and decide when to clean the sensor or discontinue normal operation.

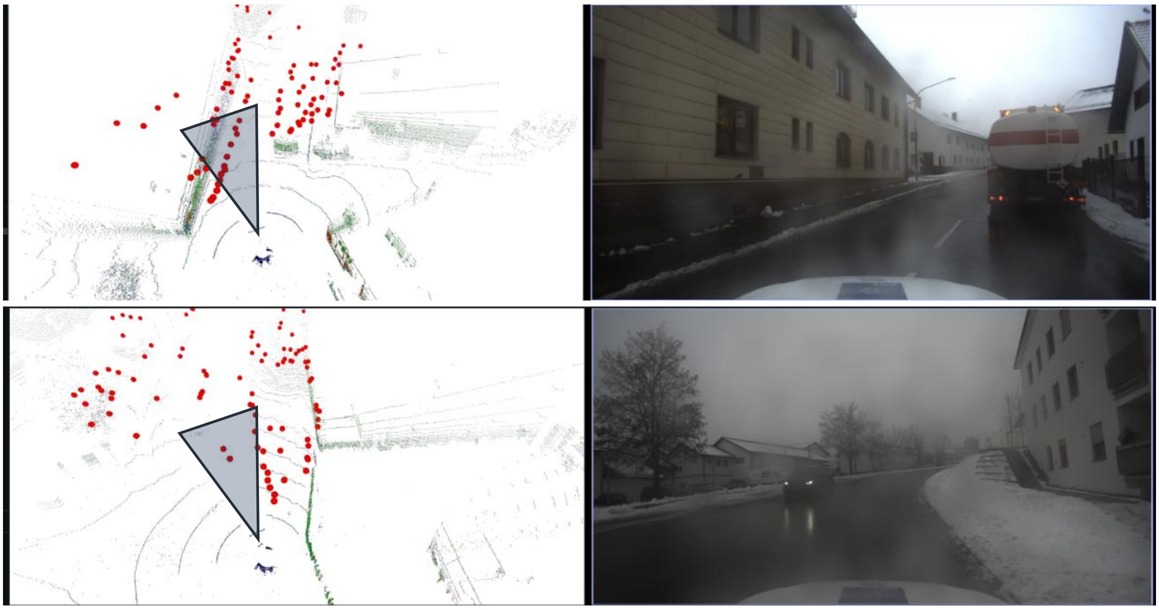

Figure 25. The time difference between the upper and the lower row is 8,5 seconds. The left column shows lidar point cloud (grey) and radar point cloud (red). The difference is also noticeable in the front looking camera.

An observation is that wet snow is worse than dry snow, since it deposits much faster. Figure 25 shows lidar and radar (3D) point clouds from two locations very close to each other. The added cone depicts roughly the near field in the travel direction. The on-coming vehicle gets very close until it is seen by the lidar. The lidar will be completely useless in less than a minute. The radar also starts suffering from snow deposit, but will still work at least five minutes more.

Figure 26 shows where the lidar used in the in-situ study is located on the roof and how it is affected by driving in snowfall in mild winter temperatures. An observation is also that only the part of the lidar that faces the travel direction becomes contaminated, and that the contamination is of a certain size and structure.

Figure 26. Snow on AVL's sensor rack on the roof of their test vehicle. The VLP-32 lidar is covered in snow on the side facing the vehicle's direction of travel. The camera dome below is also covered.

In cold winter conditions as in Figure 27 the snow is dryer and does not stick so easily. In this case, there was no snowfall, but there was an effort to soil the sensors by driving close behind a vehicle that sprayed snow from the ground. There was no build up on the top mounted lidar. However, even in dry conditions there is build-up, as for example on the back of the van due its aerodynamic design with a sharp tail creating a wake.

A build up of dry snow would likely affect the radar to a lesser extent than wet snow due to the lower loss tangent of ice as compared to water. This is reflected in the models mentioned in section 4 (planned to be published on the ROADVIEW github)) nd partially confirmed by the results in the same section where the thawing of the icy contamination clearly increases the attenuation severely.

Figure 27. Test vehicles from VTT in dry snow in cold winter conditions. Snow did not stick to the top mounted lidar of the left vehicle. Aerodynamic effects makes spray snow stick on the back of the right vehicle.

5.3 Future work

Future work should, especially for the radar, include recording during driving on different road surfaces and at different velocities. This would help to verify if and when the suggested method is valid as well as possibly fine tune it by using the detected velocity to e.g. further filter out possible noise and thus ensuring that it is the indicator surface that is detected. Recording during different driving conditions might also yield better indication of what should be considered a suitable indicator surface.

Future improvement for the lidar part, should include detection of damping on the aperture grid and the relation to the grid on the target for run-time detection. Maximum distance detection could be improved using the method in this report. Counting occurrences of points related to a grid should be in relation to the speed of the vehicle. An advantage working with grids is that it allows for subtracting two “images” of the grid to display the difference. A performance index can then be designed. Another future work is to compare consecutive areas along the road, including detection of suitable surfaces automatically. A promising neural network approach is based on graphs, with the idea to work on the relation of unknown targets with some matter on the aperture. Yet a step in the direction from data towards application is improve uncertainty description of the perception output.

6 References

[1] Euro NCAP, Assisted Driving Sensor Blocking Method, Technical Bulletin, June 2023, https://cdn.euroncap.com/media/77310/tb-041-ad-sensor-blocking-v10.pdf, last accessed July 2025.

[2] Euro NCAP, Vision 2030 – A Safer Future for Mobility, https://cdn.euroncap.com/media/74468/euro-ncap-roadmap-vision-2030.pdf, last accessed March 2024.

[3] WP29, "UN/ECE Regulation No 157 – Uniform provisions concerning the approval of vehicles with regards to Automated Lane Keeping Systems [2021/389]", Date of entry into force: 22 January 2021, www.unece.org, last accessed March 2024.

[4] P H Chan, H Haghigi, H Hyyti, V Donzella, "Library of internal sensor noise models and compound noise modelling methodology", ROADVIEW D3.5, February 2025.

[5] Jyothish K. James, Georg Puhlfürst, Vladislav Golyanik, Didier Stricker, Classification of LIDAR Sensor Contaminations with Deep Neural Networks, CSCS 2018, September 13–14, 2018, Munich, Germany.

[6] Grafika Jati et al., AutoGrAN: Autonomous Vehicle LiDAR Contaminant Detectionusing Graph Attention Networks, ICPE '24 Companion, May 7–11, 2024, London, United Kingdom.

[7] Birgit Schlager, Thomas Goelles, Stefan Muckenhuber, Daniel Watzenig, Contaminations on Lidar Sensor Covers: Performance Degradation Including Fault Detection and Modeling as Potential Applications, IEEE Open Journal of Intelligent Transportation Systems, *Digital Object Identifier 10.1109/OJITS.2022.3214094, 2022.*

[8] Schlager, Birgit; Goelles, Thomas; Watzenig, Daniel (2022): Effects of Sensor Cover Damages on Point Clouds of Automotive Lidar. TechRxiv. Preprint. https://doi.org/10.36227/techrxiv.18101279.v1.

[9] M. Trierweiler, P. Caldelas, G. Gröninger, T. Peterseim and C. Neumann, "Influence of sensor blockage on automotive LiDAR systems," *2019 IEEE SENSORS*, Montreal, QC, Canada, 2019, pp. 1-4, doi: 10.1109/SENSORS43011.2019.8956792.

[10] G. Jati, M. Molan, F. Barchi, A. Bartolini, G. Mercurio and A. Acquaviva, "LIDAROC: Realistic LiDAR Cover Contamination Dataset for Enhancing Autonomous Vehicle Perception Reliability," in IEEE Sensors Letters, vol. 8, no. 9, pp. 1-4, Sept. 2024, Art no. 1502404, doi: 10.1109/LSENS.2024.3434624.

[11] Ouster OS0 assembly, How Ouster Digital Lidar works, youtube, 2021, last accessed July 2025.

[12] Ouster OS2 documentation, OS2 - Rev7 — Ouster Sensor Docs documentation, last accessed July 2025, and Datasheets, CAD files, and guides for our lidar sensor | Ouster, also last accessed July 2025.

[13] Kutila, Matti et al., D5.5 Environmental and weather condition state estimates, ROADVIEW, August 2024.

[14] Stefan Parzefall and Heikki Hyyti, D4.3, Logging of Vehicle Data – Part 1, ROADVIEW, August 2024.

[15] M. R. Fetterman and A. Carlsen, "Blockage detection algorithm for automotive radars," *2016 IEEE MTT-S International Conference on Microwaves for Intelligent Mobility (ICMIM)*, San Diego, CA, USA, 2016, pp. 1-4, doi: 10.1109/ICMIM.2016.7533830.

[16] SANSON, Jessica, et al. "Sensor blockage detection based upon raw radar data." U.S. Patent Application No. 18/105,221.

[17] [https://data.epo.org/publication-server/rest/v1.2/patents/EP4407339NWA1/document.pdf].

[18] M. Fetterman, A. Carlsen, J. Ru and Y. Zuo, "Machine Learning Applied to Blockage Classification in Automotive Radar," *2020 IEEE MTT-S International Conference on Microwaves for Intelligent Mobility (ICMIM)*, Linz, Austria, 2020, pp. 1-4, doi: 10.1109/ICMIM48759.2020.9298969.

[19] Jiang, Jonathan H., and Dong L. Wu. "Ice and water permittivities for millimeter and sub-millimeter remote sensing applications." *Atmospheric Science Letters* 5.7 (2004): 146-151.

[20] N. Chen, R. Gourova, O. A. Krasnov and A. Yarovoy, "The influence of the water-covered dielectric radome on 77GHz automotive radar signals," *2017 European Radar Conference (EURAD)*, Nuremberg, Germany, 2017, pp. 139-142, doi: 10.23919/EURAD.2017.8249166.

[21] https://www.hec.usace.army.mil/confluence/hmsdocs/hmstrm/snow-accumulation-and-melt/snow-properties, last accessed July 2025.

[22] J. Häkli *et al.*, "Road surface condition detection using 24 GHz automotive radar technology," *2013 14th International Radar Symposium (IRS)*, Dresden, Germany, 2013, pp. 702-707.

[23] Viikari, Ville V., Timo Varpula, and Mikko Kantanen. "Road-condition recognition using 24-GHz automotive radar." *IEEE transactions on intelligent transportation systems* 10.4 (2009): 639-648.

[24] Sarabandi, Kamal, Eric S. Li, and Abid Nashashibi. "Modeling and measurements of scattering from road surfaces at millimeter-wave frequencies." *IEEE Transactions on Antennas and Propagation* 45.11 (2002): 1679-1688.

[25] Li, Eric S., and Kamal Sarabandi. "Low grazing incidence millimeter-wave scattering models and measurements for various road surfaces." *IEEE Transactions on Antennas and Propagation* 47.5 (2002): 851-861.

[26] Hassen, Alebel Arage. *Indicators for the signal degradation and optimization of automotive radar sensors under adverse weather conditions*. Diss. Darmstadt, Techn. Univ., Diss., 2006, 2007.