



Robust Automated Driving in Extreme Weather

Project No. 101069576

Deliverable 8.1

ROADVIEW Demonstration 1

WP8 – Integration and demonstration

Authors	Leila Ghasemzadeh (FORD), Eren Erdal Aksoy (HH), Petri Manninen (FGI), Jyri Maanpää (FGI), Isabelle Morvan (CRF), Hervé Ruellan (CRF), Mikko Tarkiainen (VTT), Risto-Matti Hoikka (VTT), Topi Miekkala (VTT), Pasi Pyykönen (VTT)
Lead participant	FORD
Delivery date	31/08/2024
Dissemination level	Public
Type	Demonstrator, pilot, prototype

Version 1.0



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
Leila Ghasemzadeh (FORD)	Draft deliverable	30/07/2024
Eren Erdal Aksoy (HH), Petri Manninen (FGI), Jyri Maanpää (FGI), Isabelle Morvan (CRF), Hervé Ruellan (CRF), Mikko Tarkiainen (VTT), Risto-Matti Hoikka (VTT), Topi Miekkala (VTT), Pasi Pyykönen (VTT)	Input	16/08/2024
Leila Ghasemzadeh (FORD)	Draft revision 1	25/08/2024
Isabelle Morvan (CRF)	Internal review	26/08/2024
Eren Erdal Aksoy (HH)	Internal review	28/08/2024
Leila Ghasemzadeh (FORD)	Second version	29/08/2024
Leila Ghasemzadeh (FORD)	Final version	30/08/2024
Mario Ceccarelli (accelCH)	Formatting and formal check	31/08/2024

Contents

List of Figures	4
List of Tables	4
Partner short names	5
Abbreviations	5
Executive summary	7
Objectives	7
Methodology and implementation	7
Outcomes	7
Next steps	8
1 Introduction	9
2 Demonstration #1	9
2.1 HH	10
2.1.1 Unsupervised Filtering on the FORD Data	10
2.1.2 Supervised Filtering on the THI Data	10
2.2 FGI	12
2.2.1 HD Mapping and Localisation with HD Map	12
2.2.2 Slipperiness Prediction and Drivable Area Segmentation	14
2.3 VTT	15
2.3.1 Visibility Detection	15
2.3.2 Weather Conditional Navigation System	17
2.3.3 Weather Conditional Velocity Controller	19
2.4 CRF	22
2.4.1 V2X Road Weather Information Exchanges	22
2.4.2 Weather-aware V2X Driving Advice	28
2.4.3 Roadside Sensor Fusion	31
3 Conclusions	34
References	35

List of Figures

Figure 1: Point Cloud Filtering on FORD Data Captured under Rainy Conditions.....	11
Figure 2: Point Cloud Filtering on THI Data Captured under Rainy Conditions in Daytime and Nighttime.....	11
Figure 3: Example capture of semantic segments that are used to build EA-NDT HD Map. Each colour represents a different semantic segment.	12
Figure 4: The resulting 3D normal distributions of the built EA-NDT HD map are visualised with the semantic segments.....	13
Figure 5: Example of using EA-NDT HD map for localisation. The points in the current scan and the car trajectory are visualised together with the EA-NDT HD Map.	14
Figure 6: Three example frames on grip prediction and road area segmentation.	15
Figure 7: Pytorch3D render of the generated mesh (above), corresponding depth image (below)	16
Figure 8: Real-time visibility estimation using a TensorRT accelerated estimate	17
Figure 9: Software modules of the Weather-conditional navigation system included in the demonstration	18
Figure 10: MRM execution from automated driving mode (recorded inside the VTT vehicle).	19
Figure 11: MRM execution from automated driving mode (recorded outside the VTT vehicle).	19
Figure 12: Software modules of the Weather-conditional velocity controller included in the demonstration.	20
Figure 13: Full desktop screenshot of the current velocity controller software.	20
Figure 14: Calculated indicators.	22
Figure 15: Demonstrated components for V2X Road Weather Message generation by the roadside unit.	23
Figure 16: Demonstration of the graphical interface for the V2X RWM generation from the roadside unit.	24
Figure 17: Demonstrated components for the reception of V2X RWM by the roadside unit. [1]	25
Figure 18: V2X Road weather message monitoring in snowy or icy conditions.....	26
Figure 19: V2X Road weather message monitoring in foggy conditions.....	27
Figure 20: V2X Road weather message monitoring in rainy conditions.....	27
Figure 21: V2X Road weather message monitoring in clear conditions.....	28
Figure 22: Demonstrated components of weather-aware roadside decision-making system.....	28
Figure 23: Use of collected weather estimation to trigger driving advice using V2X MCM.	30
Figure 24: Triggering of Minimal Risk Manoeuvre (MRM) with Safe Stop indication.	30
Figure 25: Demonstrated components for the roadside sensor data fusion.	31
Figure 26: Roadside sensor fusion graphical user interface layout.....	32
Figure 27: Roadside sensor fusion – clear visibility conditions.	33
Figure 28: Roadside sensor fusion – emulating reduced visibility for the camera sensor (reduced field of view).	33

List of Tables

Table 1: List of Demonstrations.....	9
Table 2: Calculated Indicators and their descriptions.....	21
Table 3: Variation of weather conditions (snow, fog, rain, clear) at the roadside system.	24
Table 4: Variation of weather conditions (snow, fog, rain, clear).....	25
Table 5: Conditions triggering driving advices.....	29
Table 6: Roadside sensor fusion test conditions	32

Partner short names

HH	Hogskolan Halmstad
LUA	Lapin Ammattikorkeakoulu OY
THI	Technische Hochschule Ingolstadt
VTI	Statens Vag- och Transportforskningsinstitut
CE	Centre d'études et d'expertise sur les risques, l'environnement, la mobilité et l'aménagement
RISE	RISE Research Institutes of Sweden AB
FGI	Maanmittauslaitos – Finnish Geospatial Research Institute
Repli5	Repli5
VTT	VTT Technical Research Centre of Finland Ltd
KO	Konrad GmbH
FORD	Ford Otomotiv Sanayi A. S.
CRF	Canon Research Centre France
accelCH	accelopment Schweiz AG
WMG	The University of Warwick
AVL	AVL Software and Functions GmbH

Abbreviations

ARVO	Automated Research Vehicle Observatory
CAV	Connected and Automated Vehicle
CPM	Collective Perception Message
D	Deliverable
DBSCAN	Density-Based Spatial Clustering of Applications with Noise
EA-NDT	Environment-Aware - Normal Distribution Transform
EC	European Commission
EU	European Union
FoV	Field of View
GNSS	Global Navigation Satellite System
GUI	Graphical User Interface
HD	High Definition
IMU	Inertial Measurement Unit
LiDAR	Light Detection and Ranging
M	Month

MCM	Manoeuvre Coordination Message
MQTT	Message Queuing Telemetry Transport
MRM	Minimal Risk Manoeuvre
MS	Milestone
NKSR	Neural Kernel Surface Reconstruction
ONNX	Open Neural Network Exchange
ROS	Robot Operating System
RSU	Roadside Unit
RWM	Road Weather Message
T	Task
TensorRT	Tensor Runtime
V2X	Vehicle-to-everything
WP	Work Package

Executive summary

The ROADVIEW project addresses the significant challenges posed by complex environment and traffic conditions on the safety and operations of Connected and Automated Vehicles (CAVs). Adverse weather conditions not only impact vehicle performance but also affect roadway infrastructure, increasing the risk of collisions and varying traffic scenarios. Most automated vehicles have been primarily trained and tested under optimal weather conditions, but ROADVIEW aims to ensure reliability and accuracy across all weather and road conditions. This deliverable (D) presents the first demonstration of the early results of the ROADVIEW innovations.

Objectives

The primary objective of D8.1 (ROADVIEW Demonstration 1) is to showcase the early implementations of the ROADVIEW system components as part of Task (T) 8.2. This deliverable focuses on demonstrating the progress made in the development of key technologies related to data preprocessing, perception, and decision-making systems. The goal is to verify and validate the initial functionality of these systems by leveraging real-world data collected as part of the ROADVIEW project. This demonstration serves as a critical milestone in advancing the project towards achieving the full integration and validation of the ROADVIEW system across different vehicle platforms.

Specifically, the objectives of this deliverable are to:

- Demonstrate early versions of the ROADVIEW system components, particularly from Work Package (WP) 4, 5, and 6.
- Highlight the performance of perception and decision-making systems on vehicles operated by project partners (FGI and VTT).
- Showcase the initial functionality of data preprocessing algorithms on a FORD truck and THI data, with a focus on offline demonstrations using pre-collected ROADVIEW data in T4.2.

Methodology and implementation

The demonstration outlined in D8.1 follows a structured approach to ensure the reliable presentation of the early-stage components of the ROADVIEW system. The methodology includes:

- **Data collection and preprocessing:** Pre-collected data from the ROADVIEW project serves as the foundation for the offline demonstrations. This data has been gathered under various environmental conditions to simulate the complex scenarios that CAVs are expected to handle.
- **Component integration:** The early-stage ROADVIEW components, including perception, decision-making, and data preprocessing systems, have been integrated into the demonstration vehicles. This integration is a collaborative effort across partners, leveraging the expertise and vehicle platforms of FGI, VTT, and FORD.
- **Offline demonstrations:** The primary focus of this deliverable is on offline demonstrations, which will involve testing the performance of the integrated components using the pre-collected data. This method allows for a controlled environment in which the functionality of the systems can be evaluated without the variability of live testing.
- **Vehicle platforms:** For this demonstration, the perception and decision-making systems have been tested on vehicles provided by FGI and VTT. Meanwhile, the data preprocessing algorithms have been implemented and demonstrated on a FORD truck.
- **Infrastructure-based systems:** Early implementation of the infrastructure-based perception system and decision-making system components that will interact with CAVs through Vehicle-to-everything (V2X) communications were also demonstrated.

Outcomes

The first demonstration of the ROADVIEW system components has successfully achieved the following outcomes:

- **Early validation:** The early-stage implementations of the perception, decision-making, and data preprocessing systems have been validated in an offline environment using the ROADVIEW data. This validation serves as a proof of concept, indicating that the systems are functioning as intended at this stage of development.
- **System performance:** The perception and decision-making systems have been shown to perform effectively on the FGI and VTT vehicles, successfully recognising and reacting to various traffic scenarios based on the

pre-collected data. Additionally, the data preprocessing algorithms have been demonstrated on the FORD truck and THI data, illustrating their capability to handle complex sensor data.

- **Collaborative effort:** This demonstration has underscored the importance of collaboration between project partners, with each partner contributing unique expertise and resources to the successful execution of the demonstration.

Next steps

Following the completion of D8.1, the next steps for the ROADVIEW project include:

- **Refinement of system components:** The insights gained from the first demonstration will be used to refine and optimise the ROADVIEW system components. This includes enhancing the accuracy and reliability of perception and decision-making systems and improving the efficiency of data preprocessing algorithms.
- **Integration and testing:** Future work will focus on the full integration of the ROADVIEW components into the target vehicle platforms and the infrastructure-based systems, moving from offline demonstrations to real-world testing scenarios. This will involve live testing in diverse environmental conditions to further validate system performance.
- **Preparation for final demonstration:** As the project progresses, efforts will shift towards preparing for the final demonstration, which will showcase the fully integrated ROADVIEW system across different vehicle platforms under various use-case scenarios. This will culminate in achieving the project's objective of demonstrating a robust and reliable CAV system capable of operating under harsh weather conditions.

1 Introduction

A demonstration in the context of a project is a structured presentation or test of a system, product, or technology that showcases its functionality, performance, and capabilities. Demonstrations are a critical phase in many projects, especially those involving technological development, as they allow stakeholders to see the practical application of theoretical concepts and designs. By providing a real-world or simulated environment, demonstrations serve as a bridge between the development phase and the deployment or implementation phase.

Demonstrations offer several key benefits to a project by providing a tangible way to validate concepts, identify issues, and gather valuable feedback. They help ensure that the developed system functions as intended and meets project goals before full-scale deployment, reducing the risk of unforeseen problems. Demonstrations also enhance communication among stakeholders by offering a clear visual representation of progress, building confidence and ensuring alignment with project objectives. Additionally, they enable iterative improvement by allowing teams to refine and optimise the system based on real-world performance and feedback, ultimately increasing the chances of a successful project outcome.

Purpose of the demonstrations:

- **Validation of concepts:** Demonstrations help verify that the developed system or components function as intended. This is a crucial step in confirming that the design meets the project's goals and requirements.
- **Identification of issues:** Early demonstrations allow teams to detect and address potential problems before full-scale deployment. By observing the system in action, developers can identify performance gaps, integration challenges, and other issues that may not have been apparent during the design or simulation stages.
- **Stakeholder engagement:** Demonstrations provide a tangible way for stakeholders, including customers, partners, and investors, to assess the progress and viability of the project. Seeing the system in operation builds confidence and ensures alignment between the development team and stakeholders.
- **Iterative improvement:** By conducting demonstrations throughout a project, teams can gather feedback that guides iterative improvements. This continuous refinement helps to optimise the system's performance, ensuring it meets all necessary criteria before final deployment.

In this respect, this deliverable reports the first demonstration of the ROADVIEW project, highlighting the early results of the project innovations.

2 Demonstration #1

The first demonstration of the ROADVIEW project showcases the initial integration of the system's perception, decision-making, and data preprocessing elements. This demonstration uses pre-collected ROADVIEW data from WP3 and WP4 to evaluate the early versions of these components in an offline setting. On the one hand, key project partners, FGI and VTT, use their vehicles to demonstrate the performance of these systems, while FORD highlights the early version of the data preprocessing algorithm on a FORD truck. On the other hand, CRF demonstrates the early version of the infrastructure-based system intended to exchange data with vehicles through V2X communication. This demonstration serves as a crucial step in validating the initial implementations before advancing to more complex real-world scenarios.

Table 1 shows the conducted demonstrations and the link to the demo videos. The Demo Video IDs are encoded as <innovation owner>-<innovation number, innovation subcategory>-<Demo#>.

Table 1: List of Demonstrations

Innovation	Innovation Owner	Vehicle/Data Owner	Demo Video
Fast and Memory Efficient Outlier Detector for 3D LiDAR Point Clouds (T4.4)	HH	FORD/THI	HH-I1-D1 , HH-I2D-D1
HD Mapping and Localisation with HD Map (T5.4)	FGI	FGI	FGI-I1-D1

Slipperiness Prediction and Drivable Area Segmentation (T5.3)	FGI	FGI	FGI-I2-D1
Harsh weather visibility estimator (T5.3)	VTT	VTT	VTT-I1-D1
Weather Conditional Navigation System (T6.2)	VTT	VTT	VTT-I2-D1
Weather Conditional Velocity Controller (T6.4)	VTT	VTT	VTT-I3-D1
V2X Road Weather Information Exchanges (T5.2 and T5.3)	CRF	CRF	CRF-I1A-D1 , CRF-I1S-D1 , CRF-I1F-D1 , CRF-I1R-D1 , CRF-I1C-D1
Weather-aware V2X Driving Advice (T6.3)	CRF	CRF	CRF-I2C-D1 , CRF-I2R-D1
Roadside Sensor Fusion (T5.2)	CRF	CRF	CRF-I3C-D1 , CRF-I3N-D1

2.1 HH

2.1.1 Unsupervised Filtering on the FORD Data

As part of T4.4, HH developed a supervised outlier detector, called 3D-OutDet [1], to detect and filter out outliers in 3D LiDAR point clouds corrupted by raindrops or snowflakes. We demonstrate how the corrupted LiDAR point clouds are removed in data logged by the FORD vehicle in rainy weather conditions in T4.2. Due to the lack of annotation in the FORD dataset, 3D-OutDet underperforms on this dataset. Instead, we demonstrate how a conventional unsupervised statistical method performs on the FORD data. Figure 1 shows a sample image from the demonstration in a rainy environment. Blue points indicate the raw LiDAR point cloud, yellow bounding boxes highlight detected objects, and red points are the detected raindrops. The camera image on the bottom right is shown for visualisation purposes only as the filtering works on the point cloud only. We refer the reader to the supplementary video HH-I1-D1 (Table 1) to see the performance of the filtering and object detection algorithms on the FORD data.

2.1.2 Supervised Filtering on the THI Data

We further demonstrate how 3D-OutDet performs on emulated rain data collected at the THI [2] outdoor proving ground in Germany. In D3.2, THI released the REHEARSE dataset [2], which was collected in three different weather conditions (rain, clear weather, and fog) using an RGB camera (LUCID), a FLIR Camera, a RADAR (ZF ProWave), and LiDAR's (Innoviz One and Ouster OS1). HH developed an automatic annotation framework to label 3D LiDAR point clouds in the REHEARSE dataset. The data were logged in rainy weather conditions at different intensities as a part of T3.2.

Figure 2 shows sample images of obtained filtered point clouds captured during daytime and nighttime. Each colour in the original point cloud (shown on the left) refers to a unique semantic class (e.g., drivable area, rain, car, etc.) in the annotated data. Note that the red points here represent the raindrops (i.e., outliers in the point cloud). Denoised point clouds in the same figure (shown in the middle) highlight the performance of the 3D-OutDet model. To a large extent, raindrops (i.e., red points) are successfully removed in both daytime and nighttime. The RGB images on the right are depicted only for the sake of visualisation. Please refer to the supplementary video HH-I2D-D1 (Table 1) to see the performance of 3D-OutDet on the THI data.




Figure 1: Point Cloud Filtering on FORD Data Captured under Rainy Conditions




Figure 2: Point Cloud Filtering on THI Data Captured under Rainy Conditions in Daytime and Nighttime.

2.2 FGI

2.2.1 HD Mapping and Localisation with HD Map

Based on the work on T5.4 and D5.6, we demonstrate how the High Definition (HD) Map is built and how the vehicle is then localised in this HD Map. The video FGI-I1-D1 (Table 1) demonstrates first how an Environment-Aware - Normal Distributions Transform (EA-NDT) HD map is built and then demonstrates the localisation of FGI's Autonomous Research Vehicle Observatory (ARVO) in the map using a high-resolution Velodyne VLS-128 lidar sensor. EA-NDT HD map is built from a point cloud by performing semantic segmentation, instance clustering, primitive extraction and cell clustering as explained in [2]. Visualisation of semantic segments is shown in Figure 3.




Figure 3: Example capture of semantic segments that are used to build EA-NDT HD Map. Each colour represents a different semantic segment.

The demo starts by showing the model of FGI's autonomous driving vehicle ARVO, after which a point cloud with intensity information is shown. The original point cloud is then semantically segmented, after which the filtering of dynamic objects is performed by utilising semantic segmentation. Equally, the extracted instance clusters, primitives and cell clusters are visualised for the viewer, followed by the visualisation of the 3D normal distributions that model the original point cloud. Figure 4 shows the resulting EA-NDT HD map together with the 3D normal distributions.




Figure 4: The resulting 3D normal distributions of the built EA-NDT HD map are visualised with the semantic segments.

The second part of the demo, HD map localisation, starts directly after building the map. The vehicle starts to move, and EA-NDT HD map is used to localise the vehicle on the map by fitting the current laser scanner view into 3D normal distributions. The registration technique used to localise the vehicle is explained in an earlier D5.6 SW on Improved Localisation Using HD Map Updating in D5.6. In the demo, each complete scan with all semantic segments is visualised on top of the EA-NDT HD map while the vehicle is moving and localising itself. All the scans are visualised in the position and orientation found by the localisation technique. The travelled trajectory is shown with orange colour and the axis visualise the positions and orientations of each localised scan. An example of the localisation is shown in Figure 5.




Figure 5: Example of using EA-NDT HD map for localisation. The points in the current scan and the car trajectory are visualised together with the EA-NDT HD Map.

2.2.2 Slipperiness Prediction and Drivable Area Segmentation

In this demo, we show the performance of our grip prediction and drivable area segmentation methods on example pre-recorded test-drive scenarios. The grip prediction model operates on the fused RGB camera and LiDAR reflectance data whereas the drivable area segmentation model operates only on the RGB camera image. All details on our models are available in D5.5.

The video FGI-I2-D1 (Table 1) presents the performance of our models in three separate data collection recordings, each of which is collected on a different day and in different road conditions. The first recording is collected during dark conditions with snowfall (left side of Figure 6). It starts with a small road which is fully covered with snow, continues to a motorway and ends in an urban area. The second recording is recorded during the daytime and most of it is covered by a snowy road with tyre tracks clear from snow (middle of Figure 6). The last recording is recorded in slushy daytime conditions, and it is mostly recorded within a suburban area featuring slushy, icy and wet roads (right side of Figure 6).

The video consists of three visualisations on top of each other as shown in Figure 6. The uppermost image shows the ground truth grip measurement data overlaid on RGB camera images. Each coloured square corresponds to a measurement with an optical grip sensor, and the colour of the square corresponds to the grip value seen in the colourbar. This data is used for measuring the validation performance of the grip prediction model. The centre image corresponds to the pixel-wise mean grip value predicted by the model, which is evaluated from the grip probability distribution given by the model for each pixel. The colours correspond similarly to the grip values in the colourbar, and the grip predictions are masked by the road area mask predicted by the road area segmentation model. This visualisation can be interpreted as the grip map of the road, which one could use to avoid slippery parts of the road and to adjust the vehicle speed. The bottom image shows the length of the 95% confidence interval for the pixel-wise grip probability distribution given by the grip prediction model. Similarly, the colour of the pixel corresponds to the length value shown in the colourbar. This image is interpreted as the uncertainty of the grip prediction – driving in uncertain conditions requires more caution.



Figure 6: Three example frames on grip prediction and road area segmentation.

In general, the grip prediction model detects well the general conditions of the road surface, such as fully snow-covered roads and slightly wet roads. Also, some local conditions, such as large slushy areas or wider tyre tracks, could be detected. However, narrow tyre tracks are difficult to detect by the model, but they are still considered areas with a lot of uncertainty. This is due to the fact that the ground truth grip data is not perfectly aligned with the corresponding road area in the training data. The road area segmentation model has difficulties in several conditions (as discussed in D5.5), but it could detect most of the road area which is sufficient for grip prediction. A more accurate road area segmentation model is under development related to T5.1.

2.3 VTT

2.3.1 Visibility Detection

The purpose of the visibility estimator developed in T5.3 is to give an estimate of the reliable view distance of the LiDAR sensor at each measurement frame. The LiDAR sensor is used as the primary sensor of the visibility estimation, because of its lower imaging resolution combined with its characteristic of being relatively easily blocked by adverse weather. The challenge of visibility estimation can be divided into two main problems.

First, the definition of the visibility threshold. There should be some reasonable and consistent method of assigning a ground truth visibility distance to the sensor measurements so that we can evaluate the performance of the actual visibility estimation method. One option is manual labelling, but this is very tedious work, especially with 3D point clouds. In manual labelling, there is also the risk of variation in the criterion of labelling due to human error. For this estimator, an automated operation was implemented to obtain ground truth labels for visibility distances. The labelling relies on first converting the 3D point cloud into a 3D mesh surface. This is achieved using the Neural Kernel Surface Reconstruction (NKSr) [3] framework, which is a state-of-the-art method of generating 3D surfaces from sparse point clouds. The output from this is a 3D mesh representing the singular LiDAR scan, and depending on the LiDAR sensor in use, the mesh structure obviously changes due to differences in laser returns and resolution. The 3D mesh is then processed using tools from the Pytorch3D library, which, among other things, has tools for mesh rendering at specific viewpoints. In our case, these tools were used to render the generated mesh from the origin of the LiDAR sensor. This achieves a more 'lifelike' image of the LiDAR frame. From this mesh render, we also obtain a more detailed depth image, than we would from just the sparse point cloud, as illustrated in Figure 7 (a screen capture from the demonstration video).




Figure 7: Pytorch3D render of the generated mesh (above), corresponding depth image (below)

The LiDAR Field of View (FoV) is usually rather large. The LiDAR sensors in the VTT vehicle have FoV's of 120° (Luminar and Livox), and 360° (Ouster). It is, of course, possible to find some kind of average visibility throughout the entire FoV, but our approach is to attempt a division of the FoV into smaller sectors, to estimate smaller sections of the entire sensor view in parallel, therefore achieving more specific information about the visibility. In the first test, the 120° area is divided into four 30° segments, each being assigned its own visibility values. We process each sector of the mesh depth image and using a customised Density-Based Spatial Clustering of Applications with Noise (DBSCAN) algorithm [4], find all parts of the rendered depth image which are not considered noise. The thresholds for the DBSCAN are neighbouring pixel distance (set to only allow neighbouring pixels), and the depth difference between the neighbouring pixels. Clusters with too few pixels are discarded. From the output of the clustering, the furthest away cluster is selected as the final visibility distance for the sector. This labelling of the visibility distances is automated, and the sector division can easily be adjusted. The negatives of the method are the heavy computation of NKSR generation and Pytorch3D rendering. Of course, the NKSR generation only gives estimated reconstructions of the point cloud and will never match the absolute real-world environment.

The second main problem is to achieve the visibility estimation in real-time in the vehicle systems. As described above, the process of mesh reconstruction, view rendering and depth image clustering are not something that is achieved in real-time. We can, however, use the generated visibility distances as ground truth labels, to fit a model that takes as an input the LiDAR point cloud, and outputs the estimated visibility for each sector. The current very first implementation is a simple neural network which takes a point cloud depth image (note, in this case, the point cloud depth image, not the mesh depth image), extracts the convolutional features, and from a linear layer, outputs the visibility estimate for each FoV sector. The design of the model requires more work for more stable and accurate estimation but implementing and fitting a separate model allows real-time processing, especially when the trained model is converted to Open Neural Network Exchange (ONNX) [5] and NVIDIA Tensor Runtime (TensorRT) [6] engines. This model is still very much a work in progress, and various methods should be tested, including more complex convolutional architectures, different inputs, adding the radar and camera sensor to support the estimation,

or further improving the method of labelling the visibility distances. A screen capture from the demonstration video illustrating real-time visibility estimation using a TensorRT accelerated estimate is shown in Figure 8. In this example, the LiDAR view has been divided into 4 sectors of 30-degree FoV each. The visibility is estimated for each sector independently based on the LiDAR depth image, and the estimation in meters is visualised as blue lines at each sector. This is preliminary testing for possible methods of real-time visibility estimation. Other options are also to simply give a single visibility estimate for the entire sensor view, for example, through averaging or attempting even more detailed estimation than just dividing the sensor view into segments.




Figure 8: Real-time visibility estimation using a TensorRT accelerated estimate

2.3.2 Weather Conditional Navigation System

The main purpose of this demo is to demonstrate how the current Minimal Risk Manoeuvre (MRM) functionality operates. The idea behind this demo is to take an MRM request as a V2X ROS message and subsequently launch the MRM in the VTT's test vehicle to emulate a request generated by the infrastructure-based decision-making system as described later in section 2.4). The software modules of the Weather-conditional navigation system (developed in T6.2 and defined in D6.2) are illustrated in Figure 9.




Figure 9: Software modules of the Weather-conditional navigation system included in the demonstration

In the demo video, the command window plays the rosbag file, and when the MRM request is activated, the path controller software of the test vehicle initiates the MRM activity and modifies the route depending on the V2X ROS message received (In-lane, road shoulder, or straight-stop as possible MRM request types). The path controller visualises the MRM route using a red colour and displays the message 'MRM active'. A rosbag from CRF containing an MRM request was used for this demo.

At the end of the video, the VTT automated vehicle is shown to execute the MRM functionality on the road. Video was recorded both inside and outside of the vehicle, as shown in Figure 10 and Figure 11. The video was recorded in Finnish Lapland where the MRM implementation was successfully tested. In this video, VTT automated vehicle's MRM fallback function safely pulls over to a road shoulder (bus stop). The MRM was triggered during automated driving at a speed of 50 km/h.



Figure 10: MRM execution from automated driving mode (recorded inside the VTT vehicle).



Figure 11: MRM execution from automated driving mode (recorded outside the VTT vehicle).

2.3.3 Weather Conditional Velocity Controller

The purpose of this demonstration is to showcase the functionality of the "Weather Conditional Velocity Controller for Safe Automated Driving in Harsh Weather Conditions" developed in T6.4. The demonstration primarily aims to display the values used and calculated by the software and to explain its operation. The slipperiness detection (refer to D5.5) provides estimations in 2D map format in front of the vehicle. Friction estimates along the trajectory are calculated, and these are used to estimate the current braking distance, which is used to determine velocity and acceleration limitations. The software modules of the Weather-conditional velocity controller (defined in D6.4) are illustrated in Figure 12.




Figure 12: Software modules of the Weather-conditional velocity controller included in the demonstration.

In the demo video, a recorded grip map and the simulated velocity of the vehicle are used to simulate the controller's implementation. Midway through the video, black rectangles (“Left area” and “Right area”) are moved to the edges of the area to simulate decreased friction values and demonstrate how these changes affect the calculations. A full desktop screenshot of the current software implementation is shown in Figure 13.




Figure 13: Full desktop screenshot of the current velocity controller software.

The main focus of this demonstration is to make estimated predictions based on the vehicle's model, real-time measurements, and the current friction estimate. The slipperiness estimate was provided by FGI (refer to D5.5), as

an example rosbag file. The velocity controller utilises real-time data, such as the Global Navigation Satellite System (GNSS), Inertial Measurement Unit (IMU), and the vehicle's odometry, in its calculations. Table 2: Calculated Indicators and their descriptions. presents the various calculated indicators and their descriptions.

Table 2: Calculated Indicators and their descriptions.

Indicator	Description
Estimated Braking Distance (meters)	The estimated braking distance is calculated using the friction estimate, current velocity, inclination, and deceleration estimate.
Friction Coefficient	Calculated using real-time grip map data.
Brake Position	A limiting parameter is calculated based on the current friction, vehicle model, and velocity. This unitless parameter is specific to the test vehicle's brake actuator.
Maximum Velocity (Limited by braking distance) (m/s)	Calculated using the visibility estimate, friction, inclination, braking distance estimation, and estimated deceleration. Figure 14 illustrates a simulated situation where the current velocity is too high (if the braking distance is intended to match the visibility), necessitating deceleration if the maximum speed limitation is applied.
Current velocity (m/s)	A measurement of the vehicle's velocity.
Max acceleration (m/s ²)	The maximum acceleration is calculated using the friction estimate and inclination.

Figure 14 illustrates the result values (one sample). In this figure, sliders are used to move the black rectangles shown in Figure 13. The friction estimate is calculated using only the values within these rectangles. Additionally, the user has the option to activate or disable maximum velocity limitations.




Figure 14: Calculated indicators.

The black rectangles are used to limit the grip map area because the provided area is large. These limited areas are placed in the estimated location where the test vehicle's tyres are located. These areas are the most critical for estimating the longitudinal movement of the test vehicle.

2.4 CRF

In this section, we will describe some demonstrations of the early implementations of the ROADVIEW system components from WP5 and WP6 related to the infrastructure-based perception and decision-making systems using V2X communications to interact with the connected and automated vehicle.

Demonstrated features are related to:

- Collaborative perception solutions developed in T5.2 and described in D5.3:
 - Roadside sensor fusion between camera and LiDAR (see section 2.4.3).
 - Dissemination of the weather type classification from the roadside system using V2X Road Weather Message (RWM) (see section 2.4.1.1).
- Onboard weather estimators developed in T5.3 and described in D5.5:
 - Update of local weather map using the received V2X Road Weather Messages from the onboard system with slipperiness and/or visibility estimations (see section 2.4.1.2).
- Infrastructure-based manoeuvre cooperation developed in T6.3:
 - Generation of driving advice using V2X Manoeuvre Coordination Message (MCM) based on the road weather conditions (see section 2.4.2).

2.4.1 V2X Road Weather Information Exchanges

The first set of offline demonstrations is related to the dissemination of road weather conditions using V2X communications.

2.4.1.1 Road Weather Message generation by the Roadside Unit

Figure 15 shows the components used to generate V2X Road Weather Message from the roadside unit to an onboard system.





Figure 15: Demonstrated components for V2X Road Weather Message generation by the roadside unit.

The demonstration is done using the interface defined by CE in charge of roadside weather estimation and CRF in charge of the V2X communication part.


As shown in Figure 16, on the left side of the demonstration video the weather conditions obtained from roadside camera images are displayed. Then the V2X message (Road Weather Message) is generated periodically (1 Hz) and monitored using the Message Queuing Telemetry Transport (MQTT) Explorer. The onboard system is connected to a cellular router to receive the V2X road weather messages through the MQTT broker. Each time there is a change in the weather conditions the console on the onboard system is updated with the latest weather type. One example of a ROS interface message (`/roadview_msgs/v2x/weather_estimation`) is provided on the right side. The measured latency between the sending of the V2X message from the Roadside Unit (RSU) to the onboard system during the test was less than 100 milliseconds (`delta_time` in the ROS topic displayed at the top right in Figure 16).

Roadside Unit (Infrastructure)

Load pre-recorded images



Weather type: clear, rain, fog, snow



V2X Road Weather Message
Generation (1 Hz)



Onboard vehicle (ROS node)




Figure 16: Demonstration of the graphical interface for the V2X RWM generation from the roadside unit.

Table 3 summarises the conditions of the recorded video for this demonstration.

Table 3: Variation of weather conditions (snow, fog, rain, clear) at the roadside system.

Recorded video	Emulated weather conditions	Description	Result
CRF-I1A-D1	Snow, fog, rain and clear	Emulation of weather type classification interface between CE and CRF modules. Generation of V2X Road Weather Message by the RSU at 1 Hz. Reception of the V2X RWM by an onboard ROS node wireless connected using a cellular router.	The onboard system receives the weather classification from the roadside unit every 1 second, with a message latency of less than 100 milliseconds.

2.4.1.2 Road Weather Message Reception by the Roadside Unit

The next offline demonstration illustrates the exchange of weather estimations from the onboard system to the roadside unit. Figure 17 shows the involved components.




Figure 17: Demonstrated components for the reception of V2X RWM by the roadside unit. [1]

At the onboard side, the weather estimations are obtained using the ROS interface (`/roadview_msgs/slipperiness` and `/roadview_msgs/visibility`) emulating possible values that will be provided by FGI for the slipperiness and VTT for the visibility as described in D5.5. The used vehicle trajectory data are obtained by replaying a VTT rosbag (`martti/gnss` ROS topic) recorded while driving in Muonio in Finland or by using a trajectory simulation (using SUMO [7] simulator) to have multiple vehicles situation.

CRF V2X onboard module sends the weather estimations using Road Weather Messages every second.

Next, the roadside unit updates its weather map with received data and displays the results on a local map view for the grip mean value (left side of the map view) and visibility range level value (right side of the map view).

Table 4 summarises the various conditions of the recorded video for this demonstration.

Table 4: Variation of weather conditions (snow, fog, rain, clear).

Recorded video	Emulated weather conditions	Onboard trajectory data	Result
CRF-I1S-D1	Snowy weather type Icy grip mean values Medium visibility	Onboard system: replay of VTT martti vehicle rosbag driving in Muonio.	RSU receives the RWMs every second and updates its local weather map. See Figure 18.
CRF-I1F-D1	Foggy weather type Poor visibility Dry grip mean values	Onboard system: replay of VTT martti vehicle rosbag driving in Muonio.	RSU receives the RWMs every second and updates its local weather map. See Figure 19.
CRF-I1R-D1	Rainy weather type Medium visibility Rainy grip mean values	Onboard system: replay of VTT martti vehicle rosbag driving in Muonio.	RSU receives the RWMs every second and updates its local weather map. See Figure 20.
CRF-I1C-D1	Clear weather Good visibility Dry grip mean values	Simulated trajectories using SUMO [7] of three vehicles driving towards an intersection in Muonio.	RSU receives the RWMs every second from each vehicle and updates its

			local weather map. See Figure 21.
--	--	--	-----------------------------------

Figure 18 is a snapshot from the recorded video where the vehicle generates V2X RWM in snowy or icy conditions.




Figure 18: V2X Road weather message monitoring in snowy or icy conditions.

Figure 19 is a snapshot from the recorded video where the vehicle generates V2X RWM in foggy weather conditions.




Figure 19: V2X Road weather message monitoring in foggy conditions.

Figure 20 is a snapshot from the recorded video where the vehicle generates V2X RWM in rainy weather conditions.




Figure 20: V2X Road weather message monitoring in rainy conditions.

Figure 21 is a snapshot from the recorded video where the three vehicles driving towards an intersection generate V2X RWM in clear weather conditions.




Figure 21: V2X Road weather message monitoring in clear conditions.

2.4.2 Weather-aware V2X Driving Advice

The next offline demonstration shows how the weather information is used by the roadside decision-making system to trigger driving advice. The onboard system receives driving advice through V2X MCM as described in D6.3. Figure 22 shows the used components for this offline demonstration.




Figure 22: Demonstrated components of weather-aware roadside decision-making system.

The system (see Figure 22) provides weather estimation both from the roadside unit (weather type classification) and the onboard units (slipperiness and visibility estimation) as explained in the previous section 2.4.1.1. The onboard system provides the vehicle's current state (position, speed, heading) and its planned trajectory using V2X MCM Intent messages. The frequency of the V2X MCM Intent message is between 1 Hz and 10 Hz depending on the vehicle dynamics. The roadside unit analyses the vehicle's planned trajectories and when it detects that there are some harsh weather conditions on the trajectory, the RSU triggers driving advice to the vehicle using V2X MCM Request.

Table 5 summarises the recorded videos for this demonstration.

Table 5: Conditions triggering driving advices.

Recorded video	Veh 1 (providing weather estimations)	Veh 2 (receiving driving advice)	Result
CRF-I2C-D1	Replay of VTT martti vehicle rosbag driving in Muonio from North to South. Generation of V2X RWM with icy grip mean values and medium visibility conditions.	Replay of VTT martti vehicle rosbag driving in Muonio after veh1 on the same trajectory from North to South. Generation of V2X MCM Intent containing the vehicle position, speed, heading and planned trajectory.	RSU has detected that veh2 intends to drive on the road with harsh weather conditions, it triggers driving advice (V2X MCM request messages) to veh2. See Figure 23.
CRF-I2R-D1	Use of a trajectory simulator with a vehicle started to drive at a Tornionte junction. Generation of V2X RWM with icy grip mean values and medium visibility conditions.	Replay of VTT martti vehicle rosbag driving in Muonio Tornionte road from South to North in direction to a safe stop location. Generation of V2X MCM Intent containing the vehicle position, speed, heading and planned trajectory.	RSU has detected that veh2 intends to drive on the road with harsh weather conditions, it triggers MCM requests to provide driving advice to veh2 to stop on the road shoulder (Safe Stop indication). See Figure 24.

Figure 23 is a snapshot from the recorded video where the roadside unit detects that a vehicle is driving on the road with harsh weather conditions and then will trigger a driving advice using the V2X MCM protocol.




Figure 23: Use of collected weather estimation to trigger driving advice using V2X MCM.

Figure 24 is an extract of the recorded video where the roadside unit generates a V2X MCM to a vehicle driving on an icy road to request to execute an MRM to stop safely on the road shoulder. It corresponds to the V2X ROS message used as input for the demonstration of the weather condition navigation system already described in section 2.3.2.




Figure 24: Triggering of Minimal Risk Manoeuvre (MRM) with Safe Stop indication.

In this demonstration, we have illustrated that the roadside decision-making system is able to trigger driving advice to the onboard system using V2X MCM, as described in D6.3, when harsh weather conditions are detected. In the next steps of the project, the integration of the V2X advice as additional inputs of the onboard decision-making will be evaluated in T6.5.

2.4.3 Roadside Sensor Fusion

The next offline demonstration is related to the sensor data fusion from roadside cameras and LiDAR sensors. This demonstration is realised with the early implementation of the infrastructure-based perception system described in D5.3. This infrastructure-based perception system is shown in Figure 25 and the involved components are surrounded by the green box.




Figure 25: Demonstrated components for the roadside sensor data fusion.

The demonstration is done using camera and LiDAR data from a pre-recorded dataset containing one car and one pedestrian agent ([SELFY] Multi-Agent Situational Awareness Database, [8]).

The graphical user interface of this demonstration is shown in Figure 26 and contains:

- Left window: map viewer with the display of the perceived object positions from each sensor (camera, LiDAR) and from the fusion. For each sensor (camera in blue, LiDAR in purple and fusion in green) the field of view polygons are represented. For each tracked object, the object current position, its past trajectory and its velocity vector are drawn for each sensor.
- Top right window: display of camera sensor with the detected object bounding boxes
- Bottom right window: display of LiDAR point cloud with the detected object bounding boxes




Figure 26: Roadside sensor fusion graphical user interface layout.

The first video of the demonstration is done assuming clear weather conditions. The second video of the demonstration is done emulating low visibility weather conditions such as fog.

The results of the roadside sensor fusion will be shared with the connected and automated vehicles using V2X communication and Collective Perception Message (CPM) as described in D5.3. CPM contains the description of the perceived object data, such as the position, velocity, classification and confidence level. Additionally, the CPM contains the description of the sensors used to perceive the object with the perception shape area (field of view). In the case where the sensor is affected by the weather conditions the sensor perception area and its confidence level should be reflected in the generated V2X CPM.

Table 6 summarises the recorded video for this demonstration.

Table 6: Roadside sensor fusion test conditions

Recorded video	Weather conditions	Sensors	Result
CRF-I3C-D1	Clear	Roadside camera + LiDAR scene with 1 pedestrian and 1 car from the dataset [8]	Car and pedestrian tracks are displayed on the RSU local map viewer
CRF-I3N-D1	Reduced visibility for the camera sensor (emulated)	Roadside camera + LiDAR scene with 1 pedestrian and 1 car from the dataset [8]	Car and pedestrian tracks are displayed on the RSU local map viewer

Figure 27 is a snapshot from the recorded demonstration with a clear weather situation.




Figure 27: Roadside sensor fusion – clear visibility conditions.

Figure 28 is a snapshot from the recorded demonstration where a reduced visibility situation is emulated for the camera sensor.




Figure 28: Roadside sensor fusion – emulating reduced visibility for the camera sensor (reduced field of view).

The measured processing time for the roadside sensor object detection and the roadside sensor fusion system is less than 100 milliseconds per frame, which aligns with the CPM requirements to provide the latest up-to-date object data every 100 milliseconds.

The next step in the scope of T5.2 will be to set up a roadside system in a location such as Finland to evaluate the collective perception solution with harsh weather conditions.

3 Conclusions

HH demonstrated the performance of the 3D outlier detector on both the FORD and THI data. First, due to the lack of annotation, an unsupervised statistical method was applied to the 3D LiDAR point cloud data recorded in the FORD truck under rainy weather conditions. Then, the fully annotated THI data was processed by the new supervised detector, called 3DOutDet, to detect and remove noisy raindrops during the day and night. In the same context, HH is currently implementing an unsupervised filtering method to detect falling snowflakes in LiDAR point clouds. As the next step, this novel unsupervised method is planned to be showcased in one of the upcoming demonstrations.

FGI demonstrated the construction of an EA-NDT HD map using data collected with FGI's ARVO. FGI also demonstrated how the vehicle can be positioned on that map using only a vehicle-mounted high-resolution Velodyne VLS-128 LiDAR sensor. In the video demonstration, the map is first built and then the vehicle is positioned on the same map using only a LiDAR sensor. As a next step, FGI is building the HD map of the Demo 4 environment for testing improved localisation on Turkish highways.

FGI also demonstrated road grip prediction at the front of the research vehicle using only camera and LiDAR sensors. The system can estimate the road grip before the vehicle passes the road so the vehicle can slow the speed already before entering a slippery area. The video demonstration presents the performance of road grip prediction in three separate data collection recordings, each of which is collected on a different daytime and in different road and weather conditions. These conditions contain day and nighttime, snowfall, slushy and snowy urban roads as well as highways. As a next step, FGI is implementing the method for real-time use in ROS and test integration of the setup with VTT's weather conditional velocity controller.

VTT demonstrated a Visibility Detector, which gives an estimate of the reliable view distance of the LiDAR sensor. The visibility estimations are provided in real-time from the in-vehicle system to be used by other software modules in ROADVIEW. In addition, VTT demonstrated the Weather Conditional Navigation System showing how the MRM functionality operates on the VTT vehicle and how the MRM in the vehicle can be requested from the infrastructure-based system using V2X. A rosbag from CRF, including V2X MRM request, is used in the demonstration. In this video, VTT automated vehicle's MRM fallback function safely pulls over to a road shoulder (bus stop). Finally, VTT demonstrated the Weather Conditional Velocity Controller for safe automated driving in harsh weather conditions. A recorded grip map, provided as a rosbag by FGI, and the simulated velocity of the vehicle are used to demonstrate the velocity controller's implementation. These demonstrations already showed the functionalities developed by VTT in WP5 and WP6 as well as implementations of ROS interfaces for receiving data from other partners. The further integration and evaluation work for these software modules will now start in T6.5.

CRF's first offline demonstration has enabled the validation of the concept of exchanging road weather information between the infrastructure-based perception system and the CAVs using V2X communication based on the enhanced RWM developed in WP5. This message uses the interfaces defined with other partners (CE, VTT and FGI) that provide the weather estimations. The second CRF demonstration has validated the concept of using road weather information to trigger driving advice to CAV from the infrastructure-based decision-making system developed in WP6. In this demonstration, partners validated the interfaces between partners by using VTT rosbag containing the CAV planned trajectory. The third CRF offline demonstration has shown a test of the early-stage implementation of the roadside sensor fusion between the camera and LiDAR sensor type developed in WP5. In the months to come, this roadside system will be installed in the Finnish Lapland to evaluate its performance under various harsh weather conditions.

As the next step, the ROADVIEW consortium will focus on annotating more data logged by different demonstration vehicle platforms. This will improve the performance of the ROADVIEW innovations in upcoming demonstrations such as Demo 3 and 4 in M38.

References

- [1] A. M. Raisuddin, T. Cortinhal, J. Holmblad and E. E. Aksoy, "3D-OutDet: A Fast and Memory Efficient Outlier Detector for 3D LiDAR Point Clouds in Adverse Weather," in IEEE Intelligent Vehicles Symposium (IV), 2024 .
- [2] M. F. D. L. C. D. P. D. F. B. V. D. P. H. C. Yuri Poledna, "REHEARSE Dataset," ROADVIEW project, 2024. [Online]. Available: <https://s3.ice.ri.se/roadview-WP3-Warwick/T3.2%20-%20Create%20Dataset/rehearse/index.html#>.
- [3] P. Manninen, H. Hyyti, V. Kyrki, J. Maanpää, J. Taher and J. Hyypä, "Towards High-Definition Maps: A Framework Leveraging Semantic Segmentation to Improve NDT Map Compression and Descriptivity," in 2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), IEEE, 2022, pp. 5370--5377.
- [4] J. a. G. Z. a. A. M. a. L. O. a. F. S. a. W. F. Huang, "Neural kernel surface reconstruction," in Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2023, pp. 4369--4379.
- [5] M. a. K. H.-P. a. S. J. a. X. X. a. o. Ester, "A density-based algorithm for discovering clusters in large spatial databases with noise," kdd, vol. 96, no. 34, pp. 226--231, 1996.
- [6] "ONNX: Open Neural Network Exchange," ONNX, [Online]. Available: <https://onnx.ai/>. [Accessed 28 8 2024].
- [7] "TensorRT. NVIDIA Developer," NVIDIA, [Online]. Available: <https://developer.nvidia.com/tensorrt>. [Accessed 28 8 2024].
- [8] "SUMO, Simulation Of Urban Mobility," [Online]. Available: <https://eclipse.dev/sumo/>.
- [9] R. Bellesort, C. Bernier, M. Cancouet, G. Kergourlay, P. Menard, T. Rakotovao and H. Ruellan, "[SELFY] Multi-Agent Situational Awareness Database (1.0.0)," 2024. [Online]. Available: <https://doi.org/10.5281/zenodo.11026094>.
- [10] P. Manninen, "SW on Improved Localization Using High-density Map Updating – First report (Deliverable No. D5.6). Robust Automated Driving in Extreme Weather (ROADVIEW)," 2024.