



US DOT National
University Transportation Center for Safety

Carnegie Mellon University



Reliable Vehicle Sensing in All Weather Conditions

High-Resolution RF Imaging, Robust 3D Mapping, and Vehicle-Ready Perception

Project #610

Rahul Mangharam, Principal Investigator

<https://orcid.org/0000-0002-3388-8283>

Mingmin Zhao, Co-Principal Investigator

<https://orcid.org/0000-0001-5226-9085>

University of Pennsylvania, School of Engineering and Applied Science
200 South 33rd Street, Moore Building, Philadelphia, PA 19104

Final Report – October 2026

Federal Grant #69A3552344811

DISCLAIMER

The contents of this report reflect the views of the authors, who are responsible for the facts and the accuracy of the information presented herein. This document is disseminated in the interest of information exchange. The report is funded, partially or entirely, under grant number 69A3552344811 from the U.S. Department of Transportation's University Transportation Centers Program. The U.S. Government assumes no liability for the contents or use thereof.

Technical Report Documentation Page

1. Report No. Safety21 Project 610	2. Government Accession No.	3. Recipient's Catalog No.
4. Title and Subtitle Reliable vehicle sensing in all weather conditions: High-resolution RF imaging, robust 3D mapping, and vehicle-ready perception		5. Report Date October 2026
		6. Performing Organization Code
7. Author(s) Rahul Mangharam, https://orcid.org/0000-0002-3388-8283 Mingmin Zhao, https://orcid.org/0000-0001-5226-9085		8. Performing Organization Report No. Safety21 Project 610
9. Performing Organization Name and Address University of Pennsylvania, School of Engineering and Applied Science 200 South 33rd Street, Moore Building, Philadelphia, PA 19104		10. Work Unit No.
		11. Contract or Grant No. Federal Grant #69A3552344811

12. Sponsoring Agency Name and Address U.S. Department of Transportation, Office of the Assistant Secretary for Research and Technology, University Transportation Centers Program (UTC grant, Bipartisan Infrastructure Law) 1200 New Jersey Avenue SE, Washington, DC 20590 Safety21 National University Transportation Center, Carnegie Mellon University (UTC grant recipient) 5000 Forbes Avenue, Pittsburgh, PA 15213		13. Type of Report and Period Covered Final Report (July 1, 2025 – June 30, 2026)
		14. Sponsoring Agency Code USDOT
15. Supplementary Notes Conducted in cooperation with the U.S. Department of Transportation and the Safety21 University Transportation Center. Federal share: \$110,000 in UTC funds. Non-federal match: \$100,000 University of Pennsylvania cost share. Project record: https://ppms.cit.cmu.edu/projects/detail/610 . Publication DOIs: https://doi.org/10.1145/3636534.3649369 , https://doi.org/10.1145/3680207.3723467 , https://doi.org/10.52202/085713-2441 . Datasets: https://upenn.box.com/v/panoradar-dataset , https://upenn.box.com/v/cartoradar-dataset . Full list of research products in Section 7.		
16. Abstract Project 610 develops high-resolution millimeter-wave radar perception for vehicle safety when cameras and LiDAR are degraded by poor visibility, occlusion, or adverse environmental conditions. The work combines a rotating synthetic aperture, motion estimation and compensation, learning-based 3D resolution enhancement, prediction uncertainty, radio-frequency-based simultaneous localization and mapping, and non-line-of-sight scene reconstruction. Published evaluations demonstrate detailed range imaging and semantic perception, centimeter-level 3D mapping, and detection and localization of hidden people around corners. The project also developed and tested a faster vehicle-mounted rotating-radar system, synchronized radar, LiDAR, camera, and INS acquisition, and outdoor data collection. The team is collaborating with SEPTA to extend these methods to public-transportation vehicles for blind-spot and all-weather sensing tests. Together, the research provides a practical path toward robust perception for advanced driver-assistance, automated-driving, and transit systems.		
17. Key Words Millimeter-wave radar; all-weather sensing; autonomous vehicles; advanced driver assistance systems; three-dimensional imaging; simultaneous localization and mapping; uncertainty quantification; object detection; non-line-of-sight sensing		18. Distribution Statement No restrictions. This document is available through the Safety21 project record.
19. Security Classif. (of this report) Unclassified	20. Security Classif. (of this page) Unclassified	21. No. of Pages 39 22. Price

Contents

1	Project Overview	1
1.1	Project Funding and Matching Funds	2
2	PanoRadar: High-Resolution RF Imaging and Recognition	3
2.1	Cylindrical Synthetic Aperture	4
2.2	Radar-Only Motion Estimation and Compensation	5
2.3	Learning-Based Imaging and Recognition	6
2.4	Dataset and Evaluation Protocol	6
2.5	Imaging and Recognition Results	6
3	CartoRadar: Uncertainty-Aware 3D Localization and Mapping	8
3.1	From RF Range Images to 3D Measurements	8
3.2	Training-Free Uncertainty Quantification	9
3.3	Uncertainty-Aware Implicit Occupancy Field	10
3.4	Online Localization and Mapping	11
3.5	Dataset and Comparative Results	11
4	HoloRadar: Non-Line-of-Sight 3D Perception	13
4.1	Two-Stage Reconstruction	13
4.2	Dataset and Test Conditions	14
4.3	Multi-Return Imaging Results	14
4.4	Visible and Hidden Geometry	15
4.5	Hidden-Person Detection	15
4.6	Contribution to Project 610	16
5	PanoRadar++ Vehicle Integration	16
5.1	Faster Rotating Radar Assembly	16
5.2	Vehicle Sensor and Computing Architecture	16
5.3	Common Time Base and Sensor Synchronization	18
5.4	Ego-Motion and Dynamic-Object Compensation	20
5.5	Outdoor and On-Car Data Acquisition	20
5.6	Public-Transportation Testing with SEPTA	21

6	Conclusions, Impact, and Recommendations	22
6.1	Principal Conclusions	22
6.2	Transportation Impact and Deployment	23
6.2.1	SEPTA as a Translation and Co-Design Partner	24
6.2.2	A Broader Transportation Vision	25
6.3	Recommended Next Phase	26
7	Research Products, Data, and Dissemination	27
7.1	Peer-Reviewed Publications	27
7.2	Software	28
7.3	Datasets and Metadata	28
7.4	Awards and Demonstrations	28
7.5	Research Personnel and ORCID Information	29
A	Metrics and Abbreviations	30
A.1	Metric Definitions	30
A.2	Abbreviations	31

List of Figures

1	Shared RF imaging and perception architecture.	3
2	PanoRadar platform and representative geometry and recognition outputs.	4
3	Representative PanoRadar geometry and semantic predictions across diverse buildings. . . .	8
4	CartoRadar converts uncertainty-quantified RF range images into a three-dimensional map and trajectory.	9
5	Perturbation-based uncertainty follows the spatial pattern of RF range errors.	10
6	Representative CartoRadar maps and trajectories reconstructed from RF observations.	12
7	HoloRadar uses multi-bounce RF paths to reconstruct geometry outside direct line of sight. .	13
8	Two-stage HoloRadar architecture for multi-return imaging and reflection-aware reconstruction.	14
9	Intermediate multi-return estimates and final scene reconstruction.	14
10	Representative hidden-scene and hidden-person reconstruction results.	15
11	PanoRadar++ mechanical design and integrated assembly for faster panoramic sensing. . . .	17
12	Project 610 PanoRadar++ rotating sensing unit developed for vehicle integration.	17
13	Roof-rack mounting concept for vehicle deployment and a clear panoramic field of view. . .	18
14	Radar-inclusive synchronization design for the PanoRadar++ vehicle sensing platform. . . .	19
15	LiDAR–camera alignment check used for PanoRadar++ vehicle sensor calibration.	19
16	Ego- and object-motion separation for a rotating radar on a moving vehicle.	20
17	Representative outdoor and on-car data collected with the PanoRadar++ pipeline.	21

List of Tables

1	Principal research and engineering results.	2
2	Project 610 funding sources.	3
3	PanoRadar test-set performance. Lower is better for errors and distances; higher is better for accuracy and average precision.	7
4	Effect of learning-based resolution enhancement on RF point-cloud accuracy.	7
5	CartoRadar uncertainty-quantification results.	10
6	CartoRadar localization and mapping compared with visual SLAM baselines. Lower is better for errors and scan time; higher is better for F-score.	12
7	Effect of uncertainty and joint pose optimization on CartoRadar map quality.	12
8	HoloRadar multi-return range error by reflection order.	14
9	HoloRadar 3D reconstruction results. Lower is better for point-cloud distances; higher is better for F-score.	15
10	Measured PanoRadar++ file sizes and write rates for representative multi-sensor captures. . .	18
11	Proposed SEPTA translation pathway and decision evidence.	24
12	Recommended vehicle and public-transportation validation matrix.	26
13	Peer-reviewed publications associated with the project research line.	27
14	Software developed through the project research line.	28
15	Datasets and descriptive metadata.	28
16	Project investigators, publication authors, and ORCID information.	29
17	Definitions of principal evaluation measures.	30
18	Abbreviations used in this report.	31

1 Project Overview

Safety21 Project 610 addresses a basic limitation in vehicle perception. Cameras provide rich color and texture, and LiDAR provides direct three-dimensional geometry, but both can lose useful information in fog, smoke, dust, precipitation, glare, darkness, and other low-visibility conditions. Millimeter-wave radar remains effective in many of these environments, yet conventional automotive radar usually returns sparse detections with limited angular detail. Led by Principal Investigator Rahul Mangharam and Co-Principal Investigator Mingmin Zhao, this project develops a high-resolution radio-frequency (RF) perception system that combines radar’s environmental resilience with the scene detail needed for vehicle safety.

The central approach is to rotate a compact 77–81 GHz radar and coherently combine its measurements as a cylindrical synthetic aperture. This creates a dense panoramic RF image with fine range and azimuth resolution. Radar-based motion estimation preserves phase coherence while the platform moves, and efficient learning models recover vertical structure and interpret the RF image. The resulting sensing stack supports dense range, surface normals, semantic segmentation, object detection, human localization, uncertainty-aware three-dimensional mapping, and non-line-of-sight perception around corners.

The research developed through four connected workstreams. PanoRadar established the high-resolution RF imaging and recognition foundation. CartoRadar extended individual scans into uncertainty-aware localization and persistent three-dimensional maps. HoloRadar used multi-bounce reflections to reconstruct hidden spaces and locate people outside direct line of sight. PanoRadar++ then adapted this research line to a faster, synchronized, vehicle-mounted sensing system for outdoor data collection and automotive testing. These are not isolated projects. Each stage adds a capability required for dependable transportation perception.

Overall project outcome

Project 610 established a coherent RF perception stack spanning panoramic three-dimensional imaging, semantic scene understanding, confidence-aware mapping, around-corner sensing, and vehicle integration. Peer-reviewed evaluations demonstrate centimeter-scale geometry and localization, while the vehicle effort demonstrates faster rotation, synchronized multi-sensor acquisition, outdoor and on-car operation, and collection of initial radar, LiDAR, camera, and navigation data.

The published results cover different parts of the sensing problem. PanoRadar demonstrates that physical motion and coherent processing can turn a low-cost commercial radar into a dense panoramic imager. CartoRadar addresses the long-tailed errors that arise when learned RF predictions are accumulated over time, and it uses uncertainty to form consistent maps and trajectories. HoloRadar shows that later RF returns, often discarded as clutter, can instead reveal hidden scene structure. Together, these capabilities provide independent measurements of geometry, motion, confidence, and occluded hazards that can complement cameras and LiDAR in advanced driver-assistance and automated-driving systems [1, 2, 4].

Vehicle translation was a major part of the project. The PanoRadar++ platform increased the panoramic scan rate, integrated motorized rotation, combined radar with LiDAR, cameras, and INS/GNSS, and established common timing, calibration, power, computing, and storage paths. The system was installed and operated on a car for initial outdoor testing, and the team collected synchronized data in the University of Pennsylvania area, including scenes near Levine Hall and Walnut Street. These tests exercised the sensing hardware and data pipeline in real street environments and provide the basis for larger controlled weather and traffic evaluations [3, 5].

The project also has a direct public-transportation path. SEPTA is a deployment partner, and the research team is collaborating with SEPTA to bring the sensing system to trolleys and buses for testing, with an initial

Table 1: Principal research and engineering results.

Workstream	Selected evidence	Transportation contribution
PanoRadar	Mean range error 15.76 cm, median range error 3.39 cm, surface-normal error 8.83°, semantic segmentation 48.00 mIoU, and human range error 12.24 cm.	Converts compact radar measurements into detailed geometry, scene categories, objects, and people rather than isolated detections.
CartoRadar	Offline localization error 14.12 cm, mapping accuracy 7.40 cm, mapping completion 8.09 cm, and 76.48-percent reconstruction F-score.	Provides persistent localization and mapping while estimating which RF predictions are reliable.
HoloRadar	Hidden-person precision 89.96 percent, recall 97.29 percent, F-score 93.48 percent, and mean localization error 13.77 cm.	Uses reflected RF paths to provide information about people and geometry hidden around corners or behind roadside occlusions.
PanoRadar++	Rotation at five revolutions per second, synchronized LiDAR, camera, and INS acquisition, roof-mounted vehicle design, outdoor captures, and initial on-car data collection.	Translates the research platform toward vehicle-mounted, real-time, all-weather evaluation.

focus on blind-spot awareness and robust perception around large transit vehicles. The Autware Foundation provides a complementary path for integrating timestamped RF geometry, objects, and uncertainty into an autonomous-driving software stack [6].

The longer-term opportunity is broader than adding another sensor to one vehicle. The same RF outputs, such as dense geometry, object locations, confidence, persistent maps, and evidence from outside direct line of sight, could form a common safety layer across transit vehicles, municipal fleets, and selected roadside locations. A bus-mounted system could observe the changing space beside a large vehicle. A fixed unit could monitor an occluded intersection. Repeated transit runs could identify locations where visibility is consistently poor. Because the project represents these functions in one sensing framework, observations collected for an immediate driver warning could also support route-level safety analysis, infrastructure planning, and development of more robust automated-driving systems. These are future applications rather than completed project results, but each follows directly from a capability demonstrated in the technical workstreams.

PanoRadar predates the formal 2025–2026 reporting period and serves as the validated technical foundation for the funded research line. CartoRadar and HoloRadar are the principal peer-reviewed extensions, while PanoRadar++ documents the vehicle-engineering work performed to translate these capabilities to transportation use. The following four sections present these technical workstreams in detail, beginning with the shared sensing foundation in PanoRadar. The report then closes with the combined conclusions, deployment recommendations, and research products.

1.1 Project Funding and Matching Funds

Project 610 was funded through the Safety21 National University Transportation Center under U.S. Department of Transportation grant 69A3552344811 for the project period July 1, 2025 through June 30, 2026. Table 2 states the source of all project funds. The University of Pennsylvania provided the full non-federal match as institutional cost share. The cost share requirement has been met and is documented in Penn’s sponsored-projects records. No other matching sources were used. The federal share was fully expended by the end of the project period.

Table 2: Project 610 funding sources.


Funding source	Type	Amount
U.S. Department of Transportation, UTC Program, through Safety21 (grant 69A3552344811)	Federal	\$110,000
University of Pennsylvania institutional cost share	Non-federal match	\$100,000
Total project budget		\$210,000

2 PanoRadar: High-Resolution RF Imaging and Recognition

All four workstreams use the same sensing idea: controlled motion turns a compact radar into a much larger synthetic antenna. The radar sweeps a 4 GHz bandwidth from 77 GHz to 81 GHz, giving an ideal range resolution

$$\Delta R = \frac{c}{2B} = 3.75 \text{ cm}, \quad (1)$$

where c is the speed of light and B is the swept bandwidth. Rotating the radar samples the scene from many positions around a circle. Coherent beamforming combines these measurements into a panoramic range–azimuth–elevation heatmap, while motion estimation corrects the phase changes created by platform translation. Figure 1 summarizes the shared sensing and learning pipeline.

**Figure 1:** Shared RF imaging and perception architecture.

The processing chain has four stages. First, range processing separates reflections by propagation delay. Second, cylindrical beamforming resolves their azimuth and elevation. Third, motion estimation and compensation preserve the synthetic aperture while the platform moves. Fourth, learning models convert the RF heatmap into dense geometry and recognition outputs. PanoRadar develops these stages for a single panoramic scan. CartoRadar adds uncertainty and combines many scans into a consistent map. HoloRadar traces multi-bounce paths into hidden space. PanoRadar++ adapts the stack to faster rotation, dynamic road scenes, synchronized vehicle sensors, and outdoor collection.

The published studies use synchronized LiDAR as geometric reference data and optical labels where semantic supervision is needed. Performance is reported with direct physical measures, including range and surface-normal error, trajectory error, map accuracy and completion, object average precision, semantic intersection over union, and hidden-person localization. The vehicle work adds system measures such as rotation rate, timing alignment, calibration, data throughput, and successful acquisition under outdoor and on-car operation. This section presents the PanoRadar foundation, and Sections 3–5 present the three subsequent technical parts.

PanoRadar is the sensing foundation of Project 610. It demonstrates that a compact commercial mmWave radar can form a dense panoramic representation of an environment and support recognition tasks that are normally associated with cameras or LiDAR. Its main contributions are a rotating cylindrical synthetic aperture, radar-only motion estimation, learning-based elevation enhancement, and a shared RF representation for geometry and semantic perception [2].




Figure 2: PanoRadar platform and representative geometry and recognition outputs.

2.1 Cylindrical Synthetic Aperture

The prototype rotates a Texas Instruments AWR1843 radar on a circle of radius $r = 8$ cm. If receive antenna a is located at height h_a and the motor rotates with angular rate ω , the antenna position at time t is

$$\mathbf{p}_t^a = (r \cos \omega t, r \sin \omega t, h_a). \quad (2)$$

A direction with azimuth θ and elevation ϕ is represented by the unit vector

$$\mathbf{d}(\theta, \phi) = (\cos \phi \cos \theta, \cos \phi \sin \theta, \sin \phi). \quad (3)$$

For complex radar signal S_t^a , coherent beamforming applies the round-trip phase correction for every antenna position and sums the measurements:

$$B(\mathbf{d}) = \sum_{a,t} S_t^a \exp\left(i \frac{4\pi}{\lambda} \mathbf{d}^\top \mathbf{p}_t^a\right), \quad (4)$$

where λ is the radar wavelength. A range Fourier transform applied to $B(\mathbf{d})$ separates reflectors by distance. Querying a grid of directions produces a $512 \times 64 \times 256$ azimuth–elevation–range heatmap for each panoramic scan.

The circular aperture has a useful analytic resolution result. For a reflector at the center direction, the array response at a nearby angular offset θ_s is proportional to

$$E(\theta_s) \propto 2\pi J_0\left(\frac{4\pi r}{\lambda} \sin \frac{\theta_s}{2}\right), \quad (5)$$

where J_0 is the Bessel function of the first kind. The corresponding ideal three-decibel azimuth beamwidth is

$$\Delta\theta \approx 0.36 \frac{\lambda}{r} = 0.72 \frac{\lambda}{d}, \quad (6)$$

with aperture diameter $d = 2r$. At $\lambda \approx 3.8$ mm and $r = 8$ cm, the ideal azimuth resolution is approximately 0.96° . The measured resolution is approximately 2.6° after antenna-pattern and calibration effects. The eight vertically spaced receive channels provide an elevation resolution of approximately 14.2° , while the 4 GHz bandwidth provides the 3.75 cm ideal range resolution in Equation (1). This strong range and azimuth resolution, paired with weaker elevation resolution, motivates the learning stage.

2.2 Radar-Only Motion Estimation and Compensation

Synthetic-aperture imaging requires antenna positions to be known within a small fraction of a wavelength. When the sensing platform translates during one rotation, uncompensated phase changes distort walls, people, and other reflectors. PanoRadar estimates this translation directly from the radar measurements.

For a static reflector n at range R_n and azimuth θ_n , a radar rotating at radius r while the platform moves at speed v and heading θ_v has the approximate path length

$$d_n(t) = R_n - r \cos(\omega t - \theta_n) - vt \cos(\theta_v - \theta_n). \quad (7)$$

The first motion term is caused by the controlled radar rotation. The second is caused by platform translation. After compensating for the known rotation and examining short slow-time windows centered at t_c , the strongest frequency associated with a reflector follows

$$f^*(t_c) = -\frac{2v}{\lambda} \cos(\theta_v - \omega t_c). \quad (8)$$

PanoRadar detects responses at multiple ranges and azimuths, then fits this sinusoid with RANSAC. Its amplitude estimates speed and its phase estimates heading. The resulting displacement is inserted into the beamforming phase term:

$$B(\mathbf{d}, \mathbf{v}) = \sum_{a,t} S_t^a \exp\left(i \frac{4\pi}{\lambda} \mathbf{d}^\top (\mathbf{p}_t^a + \mathbf{v}t)\right). \quad (9)$$

Across the 12-building evaluation, the radar-only estimator achieved a mean speed error of 8.48 mm/s and a mean heading error of 1.09° . This step is important beyond the original mobile-robot experiments because a vehicle-mounted synthetic aperture must also compensate for motion during each scan.

2.3 Learning-Based Imaging and Recognition

Signal processing produces fine range and azimuth detail but leaves structures blurred in elevation. PanoRadar uses paired RF and LiDAR observations to learn the missing vertical structure. Instead of applying an expensive three-dimensional convolution to the full heatmap, the network treats the sparse range dimension as feature channels and uses a two-dimensional, seven-stage residual U-Net. The model compresses the range channels early, learns cross-dimensional scene structure, and predicts a dense panoramic range image.

The range objective combines a masked absolute-error term with a perceptual term that preserves fine structure. It can be summarized as

$$\mathcal{L}_{\text{range}} = \|\mathbf{M} \odot (\hat{\mathbf{R}} - \mathbf{R})\|_1 + \alpha \text{LPIPS}(\mathbf{M} \odot \hat{\mathbf{R}}, \mathbf{M} \odot \mathbf{R}), \quad \alpha = 0.1, \quad (10)$$

where $\hat{\mathbf{R}}$ is the RF-derived range image, $\mathbf{R}$ is the LiDAR target, $\mathbf{M}$ removes glass regions from LiDAR supervision, and $\odot$ is elementwise multiplication. Masking is necessary because LiDAR often passes through transparent glass while mmWave radar reflects from it. The perceptual term reduces the over-smoothing produced by an absolute-error objective alone.

The learned RF image supports four additional outputs. A surface-normal head predicts local orientation. A semantic model recognizes 11 categories, including walls, floors, doors, windows, stairs, furniture, and people. A Faster R-CNN head detects human and non-human objects. Human range and direction are then recovered from the detected region and its predicted depth. Circular padding and boundary-aware detection preserve continuity between the left and right edges of the panoramic representation, preventing one object at the zero-degree boundary from being split into two.

This cross-modal learning also reduces common RF artifacts. Training against the first LiDAR return teaches the model to suppress many multipath and sidelobe responses, while the glass mask preserves valid radar reflections that LiDAR does not observe. The result is a single RF representation that supports free-space geometry, object-level reasoning, and semantic scene understanding.

2.4 Dataset and Evaluation Protocol

The prototype uses eight receive channels and approximately 1,200 azimuth samples per rotation at two revolutions per second. An Ouster 64-beam LiDAR supplies synchronized geometry. The complete dataset contains 11,033 paired RF and LiDAR frames from 12 buildings constructed between 1906 and 2013. It covers offices, corridors, classrooms, laboratories, conference rooms, glass surfaces, stairs, and dense furniture layouts. After processing, the dataset occupies approximately 461 GB and contains 11,033 range, surface-normal, and semantic labels and 20,546 object instances [2].

The evaluation uses leave-one-building-out testing. Each building is held out in turn while the remaining 11 are used for training. This creates 12 independent tests and measures whether the model generalizes across materials, floor plans, and building ages instead of memorizing one site.

2.5 Imaging and Recognition Results

Table 3 summarizes the principal test results. The mean range error is 15.76 cm, while the median is only 3.39 cm. This difference shows that most surfaces are reconstructed closely but a smaller set of hard boundaries and weak reflections creates a long error tail. The 90th-percentile range error is 31.98 cm. Surface normals average 8.83° error, semantic segmentation reaches 48.00 mean intersection over union, and object

detection reaches 52.34 AP at a 30-percent overlap threshold. Human localization averages 12.24 cm in range and 1.47° in azimuth.

Table 3: PanoRadar test-set performance. Lower is better for errors and distances; higher is better for accuracy and average precision.

Task	Metric	Result
Dense range	Mean / median / 90th percentile absolute error	15.76 cm / 3.39 cm / 31.98 cm
Surface orientation	Mean / median angular error	8.83° / 2.17°
Semantic segmentation	Mean intersection over union / pixel accuracy	48.00% / 86.33%
Object detection	AP at 30% / 50% overlap	52.34 / 38.30
Human localization	Mean range / angular error	12.24 cm / 1.47°
Ego-motion estimate	Mean speed / heading error	8.48 mm/s / 1.09°
Inference	RTX 3090 / Jetson Orin Nano	95 ms / 726 ms

Machine learning materially improves the geometric point cloud rather than simply adding semantic labels. Table 4 compares constant-false-alarm-rate points extracted from beamforming with the learned reconstruction. The three-dimensional Chamfer distance falls from 26.6 cm to 6.96 cm, and the three-dimensional modified Hausdorff distance falls from 12.0 cm to 3.23 cm. These gains show that the learned model removes multipath and sidelobe artifacts while recovering sharper surfaces.

Table 4: Effect of learning-based resolution enhancement on RF point-cloud accuracy.

Representation	2D Chamfer	3D Chamfer	2D modified Hausdorff	3D modified Hausdorff
Beamforming only	21.4 cm	26.6 cm	12.8 cm	12.0 cm
Beamforming with learning	7.43 cm	6.96 cm	3.12 cm	3.23 cm

The complete learned pipeline requires 95 ms on an NVIDIA RTX 3090, with approximately 51 ms for range estimation and 44 ms for the downstream tasks. It requires 726 ms on a Jetson Orin Nano. These values exclude the physical rotation time, but they establish a practical embedded-computing baseline for the faster PanoRadar++ platform.

PanoRadar answers the first technical question of Project 610. A compact radar can produce a detailed three-dimensional scene representation suitable for recognition. It also supplies the motion model, RF heatmap, processing pipeline, and evaluation data used by the later work. CartoRadar adds confidence and persistent mapping, HoloRadar turns multipath into an around-corner measurement, and PanoRadar++ increases the scan rate and transfers the pipeline to outdoor vehicle testing.




Figure 3: Representative PanoRadar geometry and semantic predictions across diverse buildings.

3 CartoRadar: Uncertainty-Aware 3D Localization and Mapping

PanoRadar produces a detailed panoramic observation at one point in time. CartoRadar extends those observations into a persistent three-dimensional map and an estimated sensor trajectory. Its central challenge is reliability. Learned RF range images contain a small number of large errors, and those errors can accumulate when many frames are fused. CartoRadar therefore estimates uncertainty for every predicted range value and uses that uncertainty when jointly optimizing the map and platform poses [4].

3.1 From RF Range Images to 3D Measurements

Let f be the trained PanoRadar range model and let $\mathbf{X} \in \mathbb{R}^{N_\phi \times N_\theta \times N_r}$ be its RF heatmap. The model predicts a range image $\mathbf{Y} = f(\mathbf{X})$, where Y_{ij} is the first-surface range at elevation ϕ_i and azimuth θ_j . Each range value is converted into a three-dimensional point by

$$\begin{aligned} x_{ij} &= Y_{ij} \cos \phi_i \cos \theta_j, \\ y_{ij} &= Y_{ij} \cos \phi_i \sin \theta_j, \\ z_{ij} &= Y_{ij} \sin \phi_i. \end{aligned} \tag{11}$$

A sequence of these point clouds contains enough structure for mapping, but direct concatenation produces inconsistent surfaces when a few range predictions are wrong. This long-tailed error behavior motivates an




Figure 4: CartoRadar converts uncertainty-quantified RF range images into a three-dimensional map and trajectory.

uncertainty value U_{ij} alongside every point.

3.2 Training-Free Uncertainty Quantification

CartoRadar measures sensitivity rather than modifying the trained PanoRadar network. It adds small independent Gaussian perturbations $\mathbf{E}_n$ to the RF heatmap, evaluates the same model several times, and computes the per-pixel variance:

$$U_{ij} = \text{Var}_{n=1,\dots,N}[f(\mathbf{X} + \mathbf{E}_n)_{ij}], \quad E_{n,ijk} \sim \mathcal{N}(0, \sigma^2). \quad (12)$$

If a predicted surface is strongly supported by the RF measurements, small perturbations produce similar ranges and the variance remains low. Ambiguous patterns caused by weak reflections, specularity, multipath, or object boundaries vary more and receive higher uncertainty. Monte Carlo estimation converges with only $N = 16$ samples in the reported system. The noisy inputs are stacked as one GPU batch, allowing the method to quantify uncertainty without retraining or changing the original model.

CartoRadar also evaluates a hybrid form that combines perturbation sensitivity with a network that predicts the parameters of a range distribution. If $m(f'(\mathbf{X}))$ and $v(f'(\mathbf{X}))$ are the predicted mean and variance, the total uncertainty is

$$U_{ij} = \text{Var}_{\mathbf{E}}[m(f'(\mathbf{X} + \mathbf{E}))_{ij}] + \mathbb{E}_{\mathbf{E}}[v(f'(\mathbf{X} + \mathbf{E}))_{ij}]. \quad (13)$$

The first term captures sensitivity of the predicted mean, while the second captures the distributional uncertainty predicted by the network. The training-free method is attractive for deployment because it can be placed around an existing range model. The hybrid version provides the strongest calibration when retraining is available.

Table 5 reports negative log likelihood (NLL) and area under the sparsification error curve (AUSE). Lower values indicate better uncertainty estimates. The 16-sample training-free method achieves -0.80 NLL and 0.80 AUSE without training another model. Increasing to 32 samples improves both values. The 16-sample




Figure 5: Perturbation-based uncertainty follows the spatial pattern of RF range errors.

hybrid reaches -0.93 NLL and 0.75 AUSE, outperforming the standalone learned Laplace and dropout baselines.

Table 5: CartoRadar uncertainty-quantification results.

Method	NLL	AUSE	Models
Learned Laplace distribution	0.97	0.96	1
MC dropout, 128 samples	1.00	0.87	1
Training-free, 16 samples	-0.80	0.80	0
Training-free, 32 samples	-0.82	0.74	0
Hybrid, 16 samples	-0.93	0.75	1
Hybrid, 32 samples	-0.94	0.70	1

3.3 Uncertainty-Aware Implicit Occupancy Field

CartoRadar represents the scene as a continuous occupancy field rather than a fixed voxel grid. For a sensor location $\mathbf{o}$ and a unit viewing direction $\mathbf{d}$, a ray is $\mathbf{r}(s) = \mathbf{o} + s\mathbf{d}$. A neural field f_O predicts the occupancy $\hat{O}_m = f_O(\mathbf{r}(s_m))$ at sampled distance s_m . Points before the measured surface should be free, the surface should be occupied, and points beyond the first surface are not supervised because they are occluded.

Instead of treating the RF range Y_{ij} as exact, CartoRadar models the first-hit distance D with a Laplace distribution centered at $\mu = Y_{ij}$ and scale

$$b_{ij} = \sqrt{\frac{U_{ij}}{2}}. \quad (14)$$

For a sample at s_m , the probability that the surface lies in a small interval around the sample and the probability that the ray remains free beyond it are

$$P_{\text{hit}}(s_m) = \int_{s_m - \varepsilon}^{s_m + \varepsilon} p(D) dD, \quad P_{\text{free}}(s_m) = \int_{s_m + \varepsilon}^{\infty} p(D) dD. \quad (15)$$

The weighted point-level loss is then

$$\mathcal{L}_{\mathbf{r}(s_m)} = -P_{\text{hit}}(s_m) \log \hat{O}_m - P_{\text{free}}(s_m) \log (1 - \hat{O}_m). \quad (16)$$

Low-uncertainty observations create a narrow distribution and strongly constrain the surface. High-uncertainty observations spread their influence over a wider interval and therefore cannot force the map toward one unreliable range. This is the mechanism that converts uncertainty from a diagnostic image into a direct mapping benefit.

For visualization and pose optimization, CartoRadar renders a range from the occupancy field. With accumulated transmittance $T_m = \prod_{n=1}^{m-1} (1 - \hat{O}_n)$, the rendered range selects the most likely termination point:

$$\hat{D}(\mathbf{r}) = s_j, \quad j = \arg \max_m T_m \hat{O}_m. \quad (17)$$

Because this complete path is differentiable, CartoRadar jointly refines both the occupancy field and the sensor poses. Loop closures are detected by registering low-uncertainty points to previously mapped regions, and a pose graph corrects accumulated odometry drift.

3.4 Online Localization and Mapping

The online system runs four modules in parallel: robust RF sensing, odometry, localization and mapping, and loop-closure detection. New point clouds, initial poses, and per-point uncertainty are stored in a shared-memory queue. The mapping process samples half of each training batch from the newest frame and half from earlier frames, which preserves previous geometry while incorporating new observations. Online mapping samples 64 points along each ray, compared with 512 in the offline version, and one update takes approximately 0.04 s on the reported RTX 3090 system.

This design allows mapping to continue between the two-hertz panoramic measurements. When a loop is detected, the pose graph and stored poses are updated, and joint optimization brings the occupancy field into agreement with the corrected trajectory. The online system therefore produces a map progressively rather than waiting for a complete recording.

3.5 Dataset and Comparative Results

The CartoRadar dataset contains 6,637 synchronized RF and LiDAR observations collected along 1527 m of trajectories on 14 floors in five buildings. It includes narrow corridors, large rooms, turns, loops, repeated structures, glass, stairs, and furniture. The processed dataset occupies approximately 223 GB [4].

Table 6 compares CartoRadar with three visual SLAM systems. For equal scan time, both online and offline CartoRadar provide the lowest trajectory and map errors and the highest F-score. Offline CartoRadar achieves 14.12 cm absolute trajectory error, 7.40 cm map accuracy, 8.09 cm map completion, and a 76.48-percent F-score. The online version remains close, with only 3.96 cm additional trajectory error and less than 2 cm additional mapping error. A manually swept iPad produces the best map but requires 5.5 times the scan time and does not provide an equivalent moving-platform trajectory result.

The component study in Table 7 isolates the contributions of uncertainty and joint pose optimization. Uncertainty alone improves map accuracy by 19.6 percent and completion by 20.3 percent. Joint pose optimization alone improves them by 13.2 and 12.3 percent. Using both produces 4.32 cm accuracy, 4.83 cm completion, and a 91.0-percent F-score on the aligned-segment analysis.

Table 6: CartoRadar localization and mapping compared with visual SLAM baselines. Lower is better for errors and scan time; higher is better for F-score.

Method	ATE (cm)	RPE (deg)	Accuracy (cm)	Completion (cm)	F-score (%)	Rel. scan time
ZED SDK, stereo camera	50.62	5.79	24.55	24.99	39.88	1.0
RTAB-Map, RealSense	72.26	1.29	28.66	28.93	45.00	1.0
RTAB-Map, iPad on cart	53.47	9.17	13.75	28.09	62.76	1.0
CartoRadar online	18.08	1.01	9.24	9.13	72.48	1.0
CartoRadar offline	14.12	0.93	7.40	8.09	76.48	1.0
RTAB-Map, manual iPad sweep	n/a	n/a	5.29	5.92	85.33	5.5

Table 7: Effect of uncertainty and joint pose optimization on CartoRadar map quality.

Configuration	Accuracy (cm)	Improvement	Completion (cm)	Improvement	F-score
Basic occupancy field	6.27	n/a	6.75	n/a	83.1%
With uncertainty	5.04	19.6%	5.38	20.3%	88.5%
With pose optimization	5.44	13.2%	5.92	12.3%	87.0%
With both components	4.32	31.1%	4.83	28.4%	91.0%

Loop closure reduces mean localization error by 0.43 m and the 90th-percentile error by 0.89 m, improvements of at least 76 percent. The completed maps also support later re-localization. Filtering high-uncertainty points before adaptive Monte Carlo localization reduces mean translation error from 35.43 cm to 31.39 cm and orientation error from 0.12 to 0.09 rad.




Figure 6: Representative CartoRadar maps and trajectories reconstructed from RF observations.

CartoRadar turns high-resolution RF imaging into a reliable spatial memory. For transportation, this means that a vehicle can retain an RF-derived estimate of its surroundings and its own motion when optical sensing is weak. Per-point confidence also gives downstream fusion software a principled way to reduce the weight of unstable measurements. PanoRadar++ adds INS/GNSS, LiDAR, and camera timing so these mapping concepts can be evaluated on outdoor routes and transit vehicles.

4 HoloRadar: Non-Line-of-Sight 3D Perception

Line-of-sight sensors cannot directly observe a pedestrian, cyclist, vehicle, or wall that is hidden around a corner. HoloRadar extends the Project 610 sensing stack by using multi-bounce RF returns to reconstruct hidden geometry. The method first estimates the visible surfaces that reflect the signal and then follows the reflected paths into the occluded region [1].




Figure 7: HoloRadar uses multi-bounce RF paths to reconstruct geometry outside direct line of sight.

4.1 Two-Stage Reconstruction

Multi-bounce returns are weaker and less direct than ordinary first-bounce returns. HoloRadar therefore separates imaging into two stages.

1. **Multi-return RF imaging.** A network estimates several range peaks and surface normals for each direction. The first return corresponds mainly to visible surfaces. Later returns contain longer reflection paths.
2. **Reflection-aware reconstruction.** The differentiable reflection model in Equations (18) and (19) converts those paths into 3D interaction points. A refinement model then improves hidden geometry and removes inconsistent points.

For an incoming unit direction $\mathbf{d}_{i-1}$ and a surface normal $\mathbf{n}_{i-1}$, specular reflection gives the next ray direction

$$\mathbf{d}_i = \mathbf{d}_{i-1} - 2 \left(\mathbf{d}_{i-1}^\top \mathbf{n}_{i-1} \right) \mathbf{n}_{i-1}. \quad (18)$$

If R_i is the accumulated propagation range at bounce i , the next interaction point is

$$\mathbf{p}_i = \mathbf{p}_{i-1} + (R_i - R_{i-1}) \mathbf{d}_i. \quad (19)$$

This division embeds known reflection physics in the learning system. The model does not have to infer every path from examples alone, and the intermediate ranges and normals can be inspected independently.

4.2 Dataset and Test Conditions

The HoloRadar evaluation covers 32 corners in five buildings, with 24 corners for training and eight held-out corners for testing. The dataset contains approximately 28,000 scans. It includes 21 T-shaped intersections, five L-shaped intersections, five cross intersections, and one oblique intersection, with corridor widths from 1.33 m to 4.63 m. People are placed at different hidden locations to test detection and localization. LiDAR observations collected from visible and hidden viewpoints provide reference geometry [1].




Figure 8: Two-stage HoloRadar architecture for multi-return imaging and reflection-aware reconstruction.


This dataset is intentionally structured around occlusion geometry. It allows the study to isolate how well RF reflections reveal hidden space. It does not reproduce every roadside material or moving-traffic condition, so vehicle-oriented validation remains a follow-on task.

4.3 Multi-Return Imaging Results

Table 8 reports range error for successive returns. The first-bounce mean error is 7.03 cm. Error increases for the second and third returns because each additional bounce is weaker and depends on the accuracy of earlier surface estimates. Even so, the later returns provide enough structure for useful hidden-scene reconstruction.

Table 8: HoloRadar multi-return range error by reflection order.

Return	Mean range error	Interpretation
First bounce	7.03 cm	Directly visible geometry
Second bounce	19.03 cm	First hidden interaction or longer reflected path
Third bounce	31.36 cm	Deeper and weaker multi-bounce structure



(a) Estimated return layers.

(b) Reconstructed 3D geometry.

Figure 9: Intermediate multi-return estimates and final scene reconstruction.

4.4 Visible and Hidden Geometry

HoloRadar reaches a 12.2 cm Chamfer distance, 7.1 cm modified Hausdorff distance, and 85.7-percent F-score for visible geometry. For hidden geometry, it reaches a 40.0 cm Chamfer distance, 28.3 cm modified Hausdorff distance, and 54.6-percent F-score. The lower hidden-region accuracy is expected because the signal travels through multiple reflections and hidden labels cover a broader and more difficult volume. The important result is that the hidden region is reconstructed with useful structure rather than remaining entirely unobserved.

Table 9: HoloRadar 3D reconstruction results. Lower is better for point-cloud distances; higher is better for F-score.

Region	Chamfer distance	Modified Hausdorff distance	F-score
Line of sight	12.2 cm	7.1 cm	85.7%
Non-line of sight	40.0 cm	28.3 cm	54.6%

End-to-end inference requires approximately 298 ms, within the 0.5 s rotation time of the reported platform. This supports near-real-time updates for a slow-moving or fixed intersection system and provides a baseline for faster embedded implementation.

4.5 Hidden-Person Detection

The hidden-person evaluation is the most direct transportation-safety result. HoloRadar reports 89.96-percent precision, 97.29-percent recall, and a 93.48-percent F-score. Mean localization error is 13.77 cm, with a 12.55 cm median. High recall is particularly valuable for a warning system because missed hidden road users are costly. Deployment must also manage false alarms and test more varied bodies, clothing, motion, and roadside materials.




Figure 10: Representative hidden-scene and hidden-person reconstruction results.

4.6 Contribution to Project 610

HoloRadar expands the meaning of reliable sensing. Reliability is not only the ability to see through fog or darkness. It is also the ability to provide useful evidence when direct view is blocked. Around intersections, driveways, buildings, stopped buses, construction barriers, and parked vehicles, a reflected RF path may reveal a hazard before a camera or LiDAR has a direct line of sight.

The project result is a strong proof of concept for this capability. A transportation deployment should next test outdoor wall materials, wet surfaces, vehicle body panels, moving hidden targets, wider intersections, interference from other radars, and sensor placement on both vehicles and roadside infrastructure. It should report time-to-warning as well as geometric accuracy, since the safety value depends on how much earlier the hidden road user becomes detectable.

5 PanoRadar++ Vehicle Integration

The PanoRadar, CartoRadar, and HoloRadar publications establish the RF sensing methods. PanoRadar++ brings these methods to a moving vehicle, where scan rate, vibration, ego-motion, dynamic objects, power, timing, calibration, data rate, and weather protection become first-order design constraints. The project produced a faster rotating module, a roof-mounted multi-sensor platform, a common-clock computing architecture, motion-aware fusion methods, and an outdoor acquisition pipeline. The team mounted and operated the system on test cars, conducted initial outdoor runs, and collected synchronized radar, LiDAR, camera, and INS/GNSS data for continued algorithm development [3, 5].

5.1 Faster Rotating Radar Assembly

The original PanoRadar rotates at approximately two revolutions per second. PanoRadar++ increased the demonstrated mechanical rate to five revolutions per second, reducing the time between panoramic updates from approximately 0.5 s to 0.2 s. The faster assembly combines the rotating radar, motor, support frame, electronics, and protective enclosure in a compact structure designed for vehicle installation.


The documented high-channel radar configuration uses 16 receive channels and nine transmit channels, for 144 virtual channel combinations, with 86 non-repeated virtual positions in the selected layout. At the five-hertz operating point, the acquisition design provides approximately 1,850 azimuth observations per second, or about 370 per rotation. These parameters balance panoramic update rate, synthetic-aperture density, and the data volume that the embedded computer must process.

Faster rotation has a direct safety benefit because it reduces scene age. It also creates stricter requirements. Angular position must be known accurately, the structure must remain balanced, radar packets must be timestamped consistently, and the motion compensation must handle a larger platform displacement during a scan. These issues drove the broader vehicle architecture.

5.2 Vehicle Sensor and Computing Architecture

The roof-platform design combines two Texas Instruments cascaded radar assemblies, a Hesai OT128 LiDAR, four Stereolabs ZED X cameras, an SBG Ellipse-D inertial navigation and GNSS unit, and an NVIDIA Jetson Orin computer. A router provides a local data network, and dedicated batteries and regulated supplies

support mobile operation. The exact sensor mix allows RF results to be compared with direct 3D geometry, visual context, and an independent motion estimate.



(a) Integrated assembly CAD.


Rotation Solution

Rotate radars with battery and wireless router.

Printing material: **ASA**, excellent UV and heat ($\sim 95^{\circ}\text{C}$) resistance, excellent water and chemical resistance, dedicated for outdoor applications.

Rated speed: 300 RPM

Rated output frame rate: 10 Hz



(b) Vehicle-oriented rotation module.

Figure 11: PanoRadar++ mechanical design and integrated assembly for faster panoramic sensing.




Figure 12: Project 610 PanoRadar++ rotating sensing unit developed for vehicle integration.

The integrated assembly packages the rotating sensing head and supporting electronics within a protective

structure. The vehicle mounting concept positions this assembly above the roof line, reserves space for cameras and the INS/GNSS antenna, and attaches to roof-rack crossbars while preserving a clear panoramic field of view. This modular design allowed the team to move from laboratory assembly to outdoor car testing without redesigning the complete sensing head.

Car View




Figure 13: Roof-rack mounting concept for vehicle deployment and a clear panoramic field of view.

5.3 Common Time Base and Sensor Synchronization

Multi-sensor comparisons are only meaningful when every observation refers to the same time. The PanoRadar++ architecture uses the Jetson host clock as the common software time base. The radar's digital signal-processing board exchanges time with the onboard computer through Chrony. LiDAR packets arrive through Ethernet and are stored in packet-capture format, ZED X cameras connect through GMSL2, and the INS/GNSS unit sends raw navigation data through USB. Each stream is converted into timestamped records for alignment and calibration.

The timing requirement is easy to see in a driving example. If a five-hertz radar scan is paired with a ten-hertz LiDAR frame by nearest timestamp, the maximum nominal mismatch is 50 ms. At 10 m/s, the vehicle moves 0.5 m in that interval. Without compensation, a stationary roadside object can appear half a meter apart in the two sensors. The project therefore records source timestamps, estimates sensor extrinsics, and transforms observations to a common reference time using the measured vehicle motion.

Table 10: Measured PanoRadar++ file sizes and write rates for representative multi-sensor captures.

Sensor	Captured file or data	Size	Duration	Rate
LiDAR	*.pcap	261.3 MB	31.57 s	8.28 MB/s
SBG INS/GNSS	ellipse_*.sbf (raw)	272.8 KB	24.30 s	11.2 KB/s
SBG INS/GNSS	sbg_arrival_*.csv	175.9 KB	24.30 s	7.2 KB/s
ZED camera	1,036 JPEG frames	453.7 MB	35.42 s	12.8 MB/s
ZED camera	zed_timestamps.csv	42.4 KB	35.42 s	1.2 KB/s

Extrapolating these measured rates gives approximately 30 GB of LiDAR packet data, 66 MB of SBG navigation data, and 46.1 GB of ZED JPEG data per hour, or approximately 76 GB per hour for the combined




Figure 14: Radar-inclusive synchronization design for the PanoRadar++ vehicle sensing platform.




Figure 15: LiDAR–camera alignment check used for PanoRadar++ vehicle sensor calibration.

raw streams. Encoding the camera stream as video instead of individual JPEG frames reduces its estimate to approximately 15 GB per hour. These values establish practical storage, transfer, and campaign-duration requirements for outdoor and transit testing [5].

Calibration covers radar-to-LiDAR, camera-to-LiDAR, and INS/GNSS-to-platform transforms. The engineering records include a working LiDAR–camera projection, raw INS logging, GNSS configuration, and outdoor position checks. The LiDAR data path converts recorded packets to point-cloud and array formats used by the RF processing tools. These steps are practical but essential. Without stable timestamps and coordinate transforms, an apparent RF error may actually be a fusion error.

5.4 Ego-Motion and Dynamic-Object Compensation

The original PanoRadar motion model assumes that most reflectors are stationary. Road scenes contain moving cars, bicycles, and pedestrians, so PanoRadar++ uses the full range–azimuth–elevation–Doppler cube to separate two effects:

- **Ego-motion** creates a predictable directional Doppler pattern across static background returns. This pattern estimates platform translation during the rotating scan.
- **Object motion** creates returns that disagree with the static-background pattern. Those returns can be grouped and moved to a common time before synthetic-aperture processing or multi-sensor comparison.

Research Method

- **Motivation:** Solve platform and target object motion distortion compensation.
- **Difficulty:** Compensation requires precise imaging first.
- **Solution:** Use 4D cube (Azimuth, Elevation, Range, Doppler)
 - Doppler corresponds to different speed for compensation.
 - Correctly compensated target will show the strongest signal in the 4D cube.





Figure 16: Ego- and object-motion separation for a rotating radar on a moving vehicle.

This design extends the published radar-only ego-motion estimator to the mixed static and dynamic conditions that matter for autonomous driving.

5.5 Outdoor and On-Car Data Acquisition

The team carried out outdoor and on-car data collection in the University of Pennsylvania area, including Walnut Street, the intersection near 34th Street, and the Levine and Skirkanich Hall surroundings. The collection includes panoramic radar measurements, LiDAR packet captures and point clouds, camera streams, raw inertial measurements, GNSS positions, and in-motion trajectories. The data pipeline produced packet-capture files, calibrated point-cloud files, radar heatmaps, and synchronized navigation tables for later training and evaluation. These sessions exercised the complete sensing path on streets and around moving vehicles, pedestrians, building facades, road boundaries, and other urban features that are absent from the indoor publication datasets.

The outdoor and on-car collection provides the first transportation data for the PanoRadar research line. It demonstrates that the rotation, computing, timing, calibration, and storage subsystems can operate together outside the laboratory. The collected data are now available within the project for developing vehicle motion



(a) Outdoor reference point cloud.

(b) Outdoor RF and LiDAR capture.

Figure 17: Representative outdoor and on-car data collected with the PanoRadar++ pipeline.

compensation, multimodal supervision, outdoor perception models, and longer-route mapping. A larger controlled campaign across fog, rain, spray, snow, glare, darkness, speed, and dynamic traffic will build on this operational platform.

5.6 Public-Transportation Testing with SEPTA

SEPTA is the project’s public-transportation deployment partner. The team is collaborating with SEPTA to adapt the PanoRadar++ platform for installation and testing on trolleys and buses. Large transit vehicles are a strong application because they operate in dense pedestrian environments, have substantial blind areas, make frequent low-speed turns, and encounter tunnels, garages, covered stations, poor lighting, precipitation, and street spray.

The initial SEPTA work will transfer the car-tested sensing and acquisition design to a transit-compatible mounting and power configuration. This transfer requires more than scaling the roof rack. Sensor placement must account for the vehicle body, articulation or trolley geometry, operator sight lines, passenger areas, maintenance access, electromagnetic compatibility, wash cycles, vibration, overhead infrastructure, and safe cable and power routing. Working with SEPTA operators, safety staff, vehicle engineers, and maintainers can therefore help the research team select both the highest-value field of view and a configuration that is practical to inspect and service.

A staged deployment can reduce risk while producing useful evidence at each step:

1. **Depot characterization.** Install or temporarily fixture the sensor on a stationary vehicle and measure coverage around the front corner, curb side, rear quarter, doors, and articulation zone. Use walking and cycling targets to identify blind regions and establish repeatable baselines.
2. **Controlled low-speed operation.** Record scripted turns, pull-ins, pull-outs, and stop approaches in a depot or closed area. Compare radar, camera, and LiDAR observations under daylight, darkness, glare, rain or spray, and representative occlusion.
3. **Shadow-mode service recording.** Operate the system without issuing control commands or operator alerts. Record candidate hazards and later compare them with synchronized reference data and operational context. This phase can measure false alarms and missed detections before any human-machine interface is introduced.
4. **Advisory pilot.** If shadow-mode performance is sufficient, test a limited, operator-reviewed warning

interface for one defined maneuver, such as a pedestrian or cyclist entering the curb-side swept path during a turn. Human-factors evaluation should determine whether the alert is timely, understandable, and rare enough to be trusted.

The first pilot should remain deliberately narrow. A curb-side turn or pull-away from a stop offers a clear operational question. Does RF perception identify a vulnerable road user in or approaching the vehicle's swept path, especially when mirrors or optical sensors are affected by vehicle geometry, darkness, glare, spray, or partial occlusion? Relevant measures include coverage, time to first reliable detection, localization error relative to the vehicle, track continuity, warning time before conflict, uncertainty calibration, and false alarms per operating hour. Results should also be separated by maneuver, target type, weather, lighting, and sensor placement so that an overall average does not hide a safety-critical weakness.

SEPTA service also creates a valuable path from vehicle-level sensing to system-level learning. With appropriate privacy and partner controls, repeated shadow-mode runs could reveal locations and maneuvers that repeatedly produce close or occluded interactions. Aggregated confidence and detection patterns could support decisions about mirror placement, stop layout, curb management, operator training, route risk, or the placement of a complementary roadside sensor. This use would not require the radar to identify individuals. The useful unit of analysis could be an anonymous occupancy track, a conflict geometry, or a location-level count.

This collaboration provides a direct route from the project's research results to public-transportation evaluation and possible future operational use [6]. It also supplies requirements that cannot be learned from a passenger-car test alone: the scale and swept path of a transit vehicle, repeated curb interaction, dense boarding activity, specialized maintenance practice, and the need for alerts that support rather than distract a professional operator.

The vehicle effort moves the project beyond an algorithm-only result. It shows how high-resolution RF perception can be mounted, powered, clocked, calibrated, recorded, and compared on a car, and it establishes the system foundation for Autoware integration and SEPTA transit testing.

6 Conclusions, Impact, and Recommendations

6.1 Principal Conclusions

Project 610 demonstrates that a compact millimeter-wave radar can serve as a detailed scene-understanding sensor rather than only a sparse long-range detector. PanoRadar forms a panoramic synthetic aperture, preserves coherence while the platform moves, and produces dense geometry, semantic regions, objects, and human locations. CartoRadar adds measurement confidence and spatial memory, creating consistent three-dimensional maps and trajectories from the same RF sensing foundation. HoloRadar shows that multi-bounce returns can reveal useful information about people and structures outside direct line of sight. PanoRadar++ places these capabilities into a faster vehicle sensing and acquisition system.

Four conclusions emerge from the combined work. First, controlled aperture synthesis and learning can close much of the resolution gap between low-cost radar and optical three-dimensional sensors. Second, uncertainty is part of the measurement. It improves mapping and gives a vehicle a direct way to distinguish strong RF evidence from ambiguous predictions. Third, later RF reflections can be treated as safety information rather than discarded as clutter. Fourth, vehicle deployment depends on the complete measurement system as much as it depends on the perception model. That system includes rotation, timing, calibration, power,

computing, motion compensation, storage, and synchronized reference sensors.

The peer-reviewed experiments provide strong quantitative evidence. PanoRadar achieves centimeter-scale median range and human-localization errors across 12 held-out buildings. CartoRadar produces centimeter-scale trajectories and maps and shows large gains from uncertainty-aware learning and joint pose optimization. HoloRadar detects hidden people with 97.29-percent recall and a 93.48-percent F-score in controlled corner tests. The vehicle program then extends this research line through five-hertz rotation, a roof-mounted sensor system, synchronized acquisition, outdoor operation, and initial data collected on cars.

These results establish a strong base for transportation deployment. The next phase is to broaden the existing outdoor and vehicle evidence across repeatable weather, higher speeds, dynamic traffic, and transit-specific operating conditions. This is a direct continuation of the completed platform rather than a new sensing concept.

6.2 Transportation Impact and Deployment

The central impact of the project is sensor diversity at the point where transportation systems are most vulnerable, which is when a hazard is present but direct optical evidence is weak, incomplete, or unavailable. The project does not imply that radar should replace cameras or LiDAR. Instead, it shows a path to an independent RF representation detailed enough to participate in scene understanding, mapping, and warning. This matters because a safety system built from sensors with different physical failure modes can remain informative when one modality is degraded.

The project can strengthen several vehicle and infrastructure functions:

- **Low-visibility free-space estimation.** Dense RF geometry can preserve road edges, obstacles, walls, and corridor structure when optical contrast or LiDAR return quality is reduced.
- **Vulnerable-road-user detection.** RF range and angle estimates can provide an independent observation of pedestrians and cyclists and can continue through darkness, glare, and partial visual occlusion.
- **Confidence-aware localization.** CartoRadar provides a persistent map and trajectory together with uncertainty that can guide fusion weights and safety margins.
- **Occlusion warning.** HoloRadar can reveal occupancy around corners and behind large objects before a camera or LiDAR obtains direct line of sight.
- **Sensor diversity and health monitoring.** Disagreement among radar, LiDAR, camera, and navigation sensors can expose degraded visibility, timing errors, or calibration drift.

These functions reinforce one another. Dense panoramic geometry defines the vehicle's nearby free space. Semantic and object outputs identify possible road users. Doppler and tracking distinguish motion. CartoRadar preserves context beyond a single scan. Uncertainty indicates when the RF evidence should or should not influence a decision. HoloRadar offers an early cue before direct line of sight. A future system could therefore progress from "something may be present" to a more useful safety statement: a moving road user is likely to occupy a particular region, the observation is supported across several scans, and the vehicle will intersect that region if its current maneuver continues.

6.2.1 SEPTA as a Translation and Co-Design Partner

The on-car PanoRadar++ tests and outdoor collection make these functions ready for a larger transportation study. The SEPTA collaboration is especially important because it can turn a laboratory performance measure into an operational safety question. Trolleys and buses operate close to pedestrians, cyclists, platforms, curbs, parked vehicles, and opposing traffic. Their size creates meaningful blind regions and swept paths, while their routes repeatedly expose a sensing system to tunnels, garages, stops, glare, darkness, precipitation, spray, and dense street activity.

The most credible first use is not unrestricted automation. It is a narrowly scoped driver-assistance study in which PanoRadar++ records in shadow mode and evaluates one maneuver with a clear consequence. Examples include a pedestrian entering the curb-side swept path during a right turn, a cyclist approaching alongside a vehicle pulling away from a stop, or a person moving near the front corner during low-speed depot operation. The system can be evaluated without controlling the vehicle and, initially, without presenting an alert to the operator.

The collaboration should be organized as co-design rather than a one-time equipment installation. Operators can identify maneuvers in which current sight lines and workload create difficulty. Safety personnel can define meaningful conflict and near-conflict events. Vehicle engineers and maintainers can guide sensor placement, power, sealing, inspection, and serviceability. Researchers can then translate those operational needs into measurable sensing requirements. This loop helps prevent a technically successful prototype from solving the wrong problem or creating an alert that is difficult to trust. Table 11 connects the research capabilities to a practical SEPTA pathway.

Table 11: Proposed SEPTA translation pathway and decision evidence.

Stage	Operational focus	Project capability used	Evidence needed to advance
Depot baseline	Coverage around a stationary or slowly moving trolley or bus	Panoramic RF geometry, object localization, synchronized reference sensors	Blind-region coverage, localization error, repeatability, mounting feasibility
Controlled maneuvers	Turns, stop approach and departure, pull-in, pull-out, and close roadside occlusion	Doppler separation, motion compensation, vulnerable-road-user tracking	Time to detection, track continuity, swept-path warning time, performance by condition
Shadow-mode route pilot	Natural traffic without operator alerts or vehicle control	Confidence-aware fusion, persistent mapping, multi-sensor logging	Missed events, false alarms per hour, uncertainty calibration, route and weather coverage
Limited advisory pilot	One operator-reviewed warning use case	Risk logic built from occupancy, motion, confidence, and vehicle trajectory	Human-factors response, alert usefulness, nuisance rate, maintenance burden
Fleet learning	Aggregated, privacy-preserving patterns across repeated service	CartoRadar spatial memory and standardized event records	Repeatability across vehicles, maintainable hardware, governance, measurable safety benefit

The shadow-mode phase is particularly valuable. It permits the team to tune thresholds, characterize rare errors, and compare RF detections with camera and LiDAR references without changing vehicle behavior. It also establishes whether uncertainty behaves as intended. Ambiguous or unfamiliar conditions should reduce confidence rather than produce an unjustifiably strong warning. Only after performance is understood should

the project study an operator-facing alert, and any such study should measure workload, comprehension, timing, and trust in addition to detection accuracy.

If the vehicle-level study succeeds, repeated transit operation offers a second form of impact. Privacy-preserving summaries of occupancy, occlusion, and near-conflict geometry could identify stops, turns, or corridor segments where visibility problems recur. The output could inform stop placement, curb design, mirror or camera configuration, operator training, and selection of sites for a fixed roadside sensor. In this way, the same measurements can support both an immediate warning function and longer-horizon safety planning.

6.2.2 A Broader Transportation Vision

The technical stack suggests several realistic extensions beyond the initial transit pilot:

- **A common RF safety layer for mixed fleets.** Buses, maintenance vehicles, snowplows, emergency vehicles, and passenger cars could provide the same basic outputs even if their sensor hardware and optical suites differ. These outputs are geometry, occupancy, objects, motion, and confidence. Standard interfaces would allow a perception or logging system to use RF information without depending on one research prototype.
- **Vehicle–infrastructure cooperation at occluded locations.** HoloRadar could be deployed on a vehicle, at a roadside pole, or in both locations. A fixed sensor has stable geometry and power, and a vehicle sensor follows the hazard. Their complementary views could provide earlier evidence at blind intersections, driveways, work zones, tunnel portals, or locations blocked by parked vehicles.
- **Weather-resilient continuity rather than weather-specific autonomy.** The near-term goal is not to claim full automation in every storm. A more defensible impact is graceful degradation. The system would detect when optical sensing is weakened, preserve essential geometry and occupancy, increase safety margins, and support a controlled transition to a safer operating state.
- **Living RF maps for change and hazard detection.** CartoRadar’s uncertainty-aware mapping could provide a stable structural reference while repeated observations reveal a new obstruction, altered work zone, displaced barrier, or persistent visibility problem. Confidence can separate a likely environmental change from an isolated weak return.
- **Controlled datasets and benchmarks for transportation evaluation.** Carefully documented transit and adverse-condition records would support consistent comparisons of radar-only and fused methods across project and deployment partners. This can make all-weather and occlusion performance a repeatable evaluation target while respecting privacy, licensing, and partner constraints.

Autoware provides a complementary software translation path. Timestamped RF point clouds, occupancy, objects, motion estimates, and uncertainty can be exposed through standard robotics messages and evaluated beside existing camera- and LiDAR-based modules. The most useful integration would preserve acquisition interval and uncertainty rather than reducing the radar output to an unqualified point cloud. This enables downstream software to account for scene age, motion compensation, and confidence when it fuses observations or sets a safety margin.

These future directions are grounded in the completed research, but they are not yet demonstrated transportation outcomes. The distinction is important. PanoRadar, CartoRadar, and HoloRadar provide technical evidence. PanoRadar++ provides an operational vehicle platform. SEPTA and Autoware provide translation

paths. The next phase must establish whether these elements produce earlier warnings, fewer perception gaps, or more reliable localization under the conditions that matter in service.

The research value extends beyond one prototype. The project produced peer-reviewed systems, supporting project pages, open-source software, and two large public synchronized RF and LiDAR datasets. PanoRadar received a Best Demo Award and first place in the MobiCom Student Research Competition, and CartoRadar received a Best Artifact Award. These outputs support technical validation, student training, and continued development of all-weather transportation sensing [2, 4].

6.3 Recommended Next Phase

The highest-value continuation is a controlled vehicle and transit evaluation using the completed PanoRadar++ platform. Radar-only, camera-only, LiDAR-only, and fused perception should be tested on the same routes and targets. Table 12 gives a minimum evaluation matrix.

Table 12: Recommended vehicle and public-transportation validation matrix.

Dimension	Minimum conditions	Primary measures
Visibility	Clear, darkness, glare, fog, rain or spray, and snow or representative obscurant	Detection precision and recall, range and angle error, free-space error, confidence calibration
Speed	Static, parking-lot speed, urban speed, and the highest safe closed-course speed	Scene age, motion-compensation error, track continuity, end-to-end latency
Geometry	Straight road, bend, intersection, driveway, bus stop, depot, tunnel, garage, and parked-vehicle occlusion	Map error, trajectory drift, NLOS warning time, blind-spot coverage, false-alarm rate
Road users	Stationary and moving pedestrian, cyclist, passenger vehicle, bus, and trolley	Time to first detection, localization error, detection continuity, warning time
Sensor state	Correct calibration, controlled time offset, partial sensor occlusion, and degraded optical sensor	Uncertainty response, fusion robustness, timing sensitivity, fault detection

Five implementation priorities follow:

1. **Complete the dynamic-scene pipeline.** Use the four-dimensional range–azimuth–elevation–Doppler cube to separate static background, ego-motion, and independently moving road users before forming the synthetic aperture. Evaluate crossing, oncoming, overtaking, and stop-and-go motion.
2. **Carry time and uncertainty through every interface.** Each RF output should include its acquisition interval, reference time, compensation pose, calibration state, and uncertainty. Autoware and other fusion software should consume these fields directly.
3. **Advance the SEPTA deployment.** Begin with shadow-mode recording on a trolley or bus at a depot or controlled low-speed route. Focus the first study on a clear operational question, such as pedestrian warning during a turn or blind-spot coverage near a stop.
4. **Strengthen the vehicle hardware.** Continue model acceleration and test the rotating mechanism for balance, bearing life, sealing, temperature, vibration, electromagnetic interference, and safe shutdown. Compare the rotating design with emerging fixed high-channel-count radar arrays using common resolution, cost, power, and latency measures.

5. **Curate a transportation extension dataset.** Maintain a documented subset of the collected vehicle and transit data under project and partner governance. Include radar configuration, coordinate frames, timestamps, calibration files, weather and route metadata, reference labels, access controls, and known limitations.

The recommended near-term decision is to continue the project into this focused field-validation phase. The completed sensing methods, vehicle platform, initial outdoor data, and SEPTA relationship make it possible to produce a quantitative all-weather performance envelope, an Autoware shadow-mode demonstration, and a transit-focused evaluation dataset with direct safety relevance.

7 Research Products, Data, and Dissemination

This section lists the publications, software, datasets, and other outputs associated with the Project 610 research line. Each item is given with its full DOI or URL so the record remains usable in print and with assistive technology.

7.1 Peer-Reviewed Publications

Table 13 lists the peer-reviewed publications. CartoRadar and HoloRadar were published during the project period. PanoRadar was published in 2024, before the project period began, and is listed as the foundational prior work on which the funded research builds. The publications acknowledge Safety21 support under USDOT grant 69A3552344811.

Table 13: Peer-reviewed publications associated with the project research line.

System	Publication and access	Venue and status
PanoRadar	Haowen Lai, Gaoxiang Luo, Yifei Liu, and Mingmin Zhao, “Enabling Visual Recognition at Radio Frequency.” DOI: https://doi.org/10.1145/3636534.3649369 Project page: https://waves.seas.upenn.edu/projects/panoradar/	ACM MobiCom 2024; foundational prior work
CartoRadar	Haowen Lai, Zhiwei Zheng, and Mingmin Zhao, “RF-Based 3D SLAM Rivaling Vision Approaches.” DOI: https://doi.org/10.1145/3680207.3723467 Project page: https://waves.seas.upenn.edu/projects/cartoradar/	ACM MobiCom 2025; project period
HoloRadar	Haowen Lai, Zitong Lan, and Mingmin Zhao, “Non-Line-of-Sight 3D Reconstruction with Radar.” DOI: https://doi.org/10.52202/085713-2441 Paper PDF: https://openreview.net/pdf?id=f4Wd385vHi Project page: https://waves.seas.upenn.edu/projects/holoradar/	NeurIPS 2025; project period

7.2 Software

Table 14: Software developed through the project research line.

Software	Purpose	Access
PanoRadar	Signal processing, learning, evaluation, and dataset-preparation tools for high-resolution panoramic RF imaging.	https://github.com/penn-waves-lab/PanoRadar
CartoRadar	Uncertainty estimation, implicit occupancy mapping, localization, and evaluation tools.	https://github.com/penn-waves-lab/CartoRadar
HoloRadar	Multi-return RF imaging and reflection-aware reconstruction implementation.	Public release planned through the HoloRadar project page.
PanoRadar++ vehicle tools	Sensor acquisition, packet conversion, calibration, synchronization, motor control, and motion-compensation development.	Maintained within the project.

7.3 Datasets and Metadata

Table 15: Datasets and descriptive metadata.

Dataset	Contents and coverage	Access
PanoRadar dataset	Approximately 461 GB; 11,033 synchronized RF and LiDAR frames from 12 buildings; dense range, surface-normal, semantic, object, and human labels; 20,546 object instances.	https://upenn.box.com/v/panoradar-dataset
CartoRadar dataset	Approximately 223 GB; 6,637 paired RF and LiDAR observations; 1527 m of trajectories across 14 floors in five buildings; map and trajectory reference data.	https://upenn.box.com/v/cartoradar-dataset
HoloRadar research dataset	Approximately 28,000 scans at 32 corners in five buildings; T-, L-, cross-, and oblique-intersection layouts; visible and hidden reference geometry; hidden-person placements.	Public release planned through the HoloRadar project page.
PanoRadar++ vehicle records	Initial outdoor and on-car RF, LiDAR, camera, INS/GNSS, hardware, calibration, timing, and trajectory records from the University of Pennsylvania area.	Maintained within the project under the Data Management Plan.

The PanoRadar and CartoRadar datasets are publicly downloadable at the URLs above, and each project page documents the data format and usage. The PanoRadar++ vehicle records do not currently have a public repository URL. They are packaged and archived as described in the project Data Management Plan. Their descriptive metadata covers the sensor streams, radar configuration, file formats and sizes, timing architecture, calibration workflow, and collection locations needed to interpret the records.

7.4 Awards and Demonstrations

- PanoRadar received a Best Demo Award and first place in the MobiCom Student Research Competition.
- CartoRadar received a Best Artifact Award at MobiCom 2025.

- The project websites provide videos and visual demonstrations of high-resolution RF perception, RF-based 3D mapping, and hidden-scene reconstruction.
- PanoRadar++ engineering demonstrations include five-hertz rotation, the integrated multi-sensor platform, calibration and synchronization, outdoor and on-car data collection, and the vehicle roof-mount system.

7.5 Research Personnel and ORCID Information

Rahul Mangharam and Mingmin Zhao are the project investigators. Table 16 gives ORCID IDs for the investigators and for every author of the publications in Table 13.

Table 16: Project investigators, publication authors, and ORCID information.

Name	Project or publication role	ORCID
Rahul Mangharam	Principal investigator	https://orcid.org/0000-0002-3388-8283
Mingmin Zhao	Co-principal investigator; publication author	https://orcid.org/0000-0001-5226-9085
Haowen Lai	Publication author; vehicle research contributor	https://orcid.org/0000-0001-8722-0820
Gaoxiang Luo	Publication author	https://orcid.org/0000-0001-6172-9784
Yifei Liu	Publication author	https://orcid.org/0009-0004-9992-7339
Zhiwei Zheng	Publication author	https://orcid.org/0009-0009-9508-1803
Zitong Lan	Publication author	https://orcid.org/0000-0002-8524-6047

References

- [1] H. Lai, Z. Lan, and M. Zhao, “Non-line-of-sight 3D reconstruction with radar,” in *Advances in Neural Information Processing Systems*, vol. 38, 2025. doi: 10.52202/085713-2441. [Online]. Available: <https://waves.seas.upenn.edu/projects/holoradar/>
- [2] H. Lai, G. Luo, Y. Liu, and M. Zhao, “Enabling visual recognition at radio frequency,” in *Proceedings of the 30th Annual International Conference on Mobile Computing and Networking*, 2024, pp. 388–403. doi: 10.1145/3636534.3649369. [Online]. Available: <https://waves.seas.upenn.edu/projects/panoradar/>
- [3] H. Lai and Project 610 Research Team, “PanoRadar++ vehicle: Project updates and progress,” internal project engineering presentation, January–February 2026.
- [4] H. Lai, Z. Zheng, and M. Zhao, “RF-based 3D SLAM rivaling vision approaches,” in *Proceedings of the 31st Annual International Conference on Mobile Computing and Networking*, 2025, pp. 170–185. doi: 10.1145/3680207.3723467. [Online]. Available: <https://waves.seas.upenn.edu/projects/cartoradar/>
- [5] Project 610 Research Team, “PanoRadar++ vehicle hardware system,” internal project engineering presentation, February–July 2026.
- [6] Safety21 University Transportation Center, “Project 610: Reliable vehicle sensing in all weather conditions.” Accessed: Oct. 6, 2026. [Online]. Available: <https://ppms.cit.cmu.edu/projects/detail/610>

A Metrics and Abbreviations

This appendix defines the principal measures and technical abbreviations used throughout the report.

A.1 Metric Definitions

Table 17: Definitions of principal evaluation measures.

Metric	Meaning
Absolute error	Magnitude of the difference between a predicted and reference value. Range errors in this report are stated in centimeters. Lower is better.
Angular error	Difference between predicted and reference directions or surface normals, in degrees. Lower is better.
mIoU	Mean intersection over union across semantic classes. It compares overlap between predicted and reference regions. Higher is better.
Pixel accuracy	Percentage of labeled image locations assigned the correct class. Higher is better, but it can be influenced by large common classes.
AP30 and AP50	Average precision for object detection at 30-percent or 50-percent overlap thresholds. Higher is better.
ATE	Absolute trajectory error. It measures overall difference between the estimated and reference paths after alignment. Lower is better.
RPE	Relative pose error. It measures local motion error between nearby poses. Lower is better.
Mapping accuracy	Average distance from reconstructed points to the reference surface. Lower is better.
Mapping completion	Average distance from reference-surface points to the reconstruction. Lower is better.
Chamfer distance	A bidirectional average point-cloud distance used to compare reconstructed and reference geometry. Lower is better. Exact sampling and aggregation follow each cited paper.
Modified Hausdorff distance	A robust point-cloud separation measure that emphasizes unmatched geometry. Lower is better. Exact aggregation follows the HoloRadar paper.
Precision	Fraction of reported detections that match a reference target. Higher is better.
Recall	Fraction of reference targets that are detected. Higher is better.
F-score	Harmonic mean of precision and recall, or the paper-defined point-cloud precision and recall. Higher is better.
Negative log likelihood	Calibration measure for a probabilistic prediction. Under the CartoRadar definition, lower is better.
AUSE	Area under the sparsification error curve. It tests whether predicted uncertainty correctly ranks errors. Lower is better.
Inference time	Processing time for the reported model and hardware, excluding physical acquisition unless stated otherwise. Lower is better at equal accuracy.

A.2 Abbreviations

Table 18: Abbreviations used in this report.

Abbreviation	Definition
3D / 4D	Three-dimensional / four-dimensional
ADAS	Advanced driver-assistance system
ADS	Automated driving system
AP	Average precision
ATE	Absolute trajectory error
DMP	Data Management Plan
FMCW	Frequency-modulated continuous wave
GNSS	Global Navigation Satellite System
GMSL2	Second-generation Gigabit Multimedia Serial Link
INS	Inertial navigation system
LiDAR	Light detection and ranging
LOS / NLOS	Line of sight / non-line of sight
mIoU	Mean intersection over union
mmWave	Millimeter wave
PPMS	Safety21 project and performance management system
RF	Radio frequency
RPE	Relative pose error
SEPTA	Southeastern Pennsylvania Transportation Authority
SLAM	Simultaneous localization and mapping
USDOT	U.S. Department of Transportation
UTC	University Transportation Center