Article

Not peer-reviewed version

Reading the Fog: Real-Time LiDAR-Based Detection Range Estimation for Autonomous Vehicles in Adverse Weather

Jongmin Lee *

Posted Date: 15 May 2026

doi: 10.20944/preprints202605.1004.v1

Keywords: LiDAR; autonomous driving; adverse weather; fog; detection range; visibility estimation; real-time perception; fail-operational

 doi:10.20944/preprints202605.1004.v1

Article

Reading the Fog: Real-Time LiDAR-Based Detection Range Estimation for Autonomous Vehicles in Adverse Weather

Jongmin Lee [1,2,*]

[1] Department of Mechanical Engineering, Seoul National University, Seoul 08826, South Korea
[2] Autonomous Driving SW Division, Hyundai Mobis Company Ltd., Seoul, South Korea
* Correspondence: jesus911021@snu.ac.kr

Abstract

Detecting objects with LiDAR in fog has long been treated as a noise-removal problem: identify the fog returns and throw them away. This work takes the opposite view. The same fog returns that other methods discard are read as a direct measurement of how far the sensor can still see. After an off-the-shelf fog-segmentation step, the near-range fog points are used to estimate the atmosphere's optical thickness, which is then converted into a per-frame Maximum Detection Range (MDR) — a single, human-readable number, in meters, that tells planning and safety modules how far the LiDAR is currently usable. The estimator has no learnable parameter(The physics core — a per-frame Beer–Lambert slope fit followed by the Koschmieder mapping) and runs in real time during driving. We validate it with real-vehicle experiments in controlled heavy-fog scenarios. Across six fog runs the estimated MDR recovers the true first-detection distance with a mean absolute error of 1.99 m and R^2 = 0.741 — a 58 % error reduction over a no-feature baseline. To rule out that the result is an artefact of our own recording, the same estimate is cross-checked against two independent references used by the wider community: the STF transmissometer dataset and the Hahner et al. automotive fog simulator. Our estimate agrees with both, which is the key differentiator from prior LiDAR-fog work that reports performance only on its own data. The result is a real-time, parameter-free, externally-corroborated visibility readout that fits inside an existing fog-denoising pipeline and gives downstream autonomy a quantitative answer to the question "how far can I see right now?"

Keywords: LiDAR; autonomous driving; adverse weather; fog; detection range; visibility estimation; real-time perception; fail-operational

Highlights

What are the main findings?

• LiDAR fog returns, conventionally treated as noise to be denoised, can be read directly as an optical measurement of how far the sensor can still see — yielding a per-frame Maximum Detection Range $\widehat{MDR}$ (in meters) without any additional inference cost on top of an existing fog-denoising network.

• On two independent real-vehicle heavy-fog recordings, the recovered fog density $\hat{\beta}_{ext}$ agrees to 0.21–0.22 m^{-1} across recordings ($\widehat{MDR} \approx$ 17–19 m), and on the scenarios corpus a leave-one-out per-frame regression recovers the first-detection distance to within 1.99 m mean absolute error (R^2 = 0.741, a 58 % MAE improvement over the no-feature baseline) at LiDAR scan rate (≈14 Hz) on a desktop-class GPU.

What are the implications of the main findings?

 doi:10.20944/preprints202605.1004.v1

• The proposed Maximum Detection Range scalar provides exactly the per-frame, traceable, bounded visibility input that ISO 26262 hazard analysis, ISO 21448 SOTIF Clause 7 triggering-condition specification, and SAE J3016 Operational Design Domain definitions all require — making the perception-layer visibility signal directly consumable by safety arguments rather than only by downstream object detectors.

• Because the visibility estimator reuses the segmentation output that any fog-denoising stack already produces, the proposed pipeline is a drop-in upgrade for existing LiDAR perception systems: the cleaned cloud continues to feed the object detector while the previously-discarded fog-labelled subset feeds a new MDR̂ head, enabling safety controllers (e.g., speed capping, emergency pull-over) to react to silent detector failures *before* they propagate to vehicle behavior.

1. Introduction

Adverse weather is one of the longest-standing open problems of LiDAR-based autonomous driving. The dominant body of work treats fog, snow, and rain as *noise* — physical returns that contaminate the point cloud and must be removed before downstream object detection. Per-point denoising classifiers have steadily improved in this framing, yet the underlying premise — fog is a contaminant — has remained unchallenged. The present work proposes a different premise. Once fog returns have been segmented from real returns by the same per-point classifier the denoiser would use, the very same segmented set carries a different kind of information: how far the sensor can still see through the current fog. The denoiser removes fog; the proposed pipeline reads it.

The reason this distinction matters beyond a cleaner point cloud lies in what the autonomous-driving safety standards actually require. The maximum detection range (MDR) — defined here as the farthest range at which the LiDAR can still register a real object return given the current weather state — is precisely the quantity these standards demand. ISO 26262 hazard analysis requires that the sensor reach be bounded for any safety case to close. ISO 21448 (SOTIF) treats unbounded perceptual degradation as a primary trigger of *performance limitation* hazards. SAE J3016 defines an Operational Design Domain (ODD) in part by the visibility envelope within which the perception stack is qualified to operate. None of these standards consume a denoised point cloud directly; all of them consume a quantitative, per-frame, scalar answer in meters. Producing that answer directly from the LiDAR point cloud, at sensor frame rate, is the goal of this work.

Concretely, the per-point intensity I(r) of the fog-labelled subset within the ego-proximal volume $r \in [0.5, 3]$ m — a region that, on the dual-LiDAR roof-mount instrumented vehicle used in this work, contains a dense and contamination-free fog return — is collected and log I(r) is fit linearly against r. By the Beer–Lambert law, the negative half-slope of this fit is the extinction coefficient $\hat{\beta}_{ext}$, and the Koschmieder relation yields the maximum detection range as $\widehat{MDR} = 3.912 / \hat{\beta}_{ext}$ directly, with no concentric-ring partition, no ring-radius tuning, and no engineered density vector. A ±5-frame median over per-frame slopes suppresses single-frame jitter without introducing any temporal model. A consequence of this design is that the fog-segmentation network — the same network the denoising community already runs at the LiDAR front-end — becomes *dual-purpose*: a single forward pass yields both the cleaned cloud (the conventional product) and the per-frame visibility signal (the product of this work), at no additional inference cost.

Two further design choices distinguish the present work. First, the entire chain — raw dual roof-mounted LiDAR point cloud, fog segmentation, ego-proximal window selection, Beer–Lambert slope fit, and Koschmieder mapping — runs in real time on a single in-vehicle desktop-class GPU (NVIDIA RTX 2070, 8 GB VRAM). On-vehicle measured end-to-end latency is 60–70 ms per frame, comfortably within the LiDAR scan cycle (≈70 ms, ≈14 Hz). Second, the estimator is validated on real-vehicle heavy-fog data acquired on a closed instrumented track (FMTC, SNU) by the SNU instrumented test vehicle, with two independent reference recordings yielding cross-recording-concordant $\hat{\beta}_{ext}$ (Section 4). The dataset is released on Hugging Face [1,2] for community benchmarking — an asset that the adverse-weather LiDAR community has long lacked outside of simulated or geographically narrow public sources.

 doi:10.20944/preprints202605.1004.v1

Contributions of this work. The paper makes the following contributions, organized so that each item points to the section in which it is developed.

(i) Reframing of LiDAR fog perception (§3). The paper recasts adverse-weather LiDAR perception from *fog removal* to *fog reading*. The fog-labelled returns produced by an existing denoising network — conventionally discarded as noise — are read here as the visibility signal itself, which is the central methodological inversion of the work.

(ii) Direct extinction estimator without ring engineering (§3.2). The fog density, expressed as the Beer–Lambert extinction coefficient β_{ext} (in m^{-1}, larger values meaning denser fog), is recovered as the negative half-slope of log I(r) on the fog-labelled subset within an ego-proximal window $r \in [0.5, 3]$ m. The estimator uses no concentric-ring partition, no per-ring radius schedule, and no engineered density vector; the window bounds are dictated by the geometry of the dual roof-mount, not tuned. A ±5-frame temporal median suppresses single-frame jitter without introducing any sequence model.

(iii) Dual-purpose use of the segmentation network (§3.1). The same fog-segmentation network that any denoising stack already runs is used here for two purposes simultaneously: its output feeds (i) the conventional cleaned cloud consumed by downstream object detectors and (ii) the fog-labelled cloud consumed by the visibility estimator. The visibility estimator therefore adds zero inference cost on top of an existing denoising pipeline — a drop-in upgrade.

(iv) Koschmieder-grounded mapping to a Maximum Detection Range scalar (§3.3). The extinction coefficient β_{ext} is mapped to a per-frame Maximum Detection Range, $\widehat{MDR} = 3.912 / \beta_{ext}$ (in meters), via the Koschmieder relation under the standard 5 % contrast threshold. The mapping introduces no further tunable parameter; the only constant is the textbook Koschmieder coefficient $3.912 = \ln(1/0.05)$.

(v) Standards-aligned auditable output (§3.3, §5). $\widehat{MDR}$ is precisely the per-frame, traceable, reproducible, bounded scalar that ISO 26262 hazard analysis (Exposure-class assignment), ISO 21448 SOTIF Clause 7 (triggering-condition specification), and SAE J3016 ODD definition (visibility envelope) all require. Every step of the chain — segmentation output, linear fit, constant divide, median — is inspectable, which a black-box deep regressor would not satisfy.

(vi) Cross-recording and cross-instrument concordance (§4). On two independent real-vehicle heavy-fog reference recordings the per-frame median β_{ext} agrees to 0.21–0.22 m^{-1} across recordings ($\widehat{MDR} \approx$ 17–19 m), supporting a sensor-universal fog signature rather than a recording-specific artefact. The recovered β_{ext} is additionally cross-validated against (a) the SeeingThroughFog on-vehicle transmissometer [3] for cross-sensor / cross-site agreement and (b) the Hahner et al. [4] automotive fog forward-simulator [4] for fog-regime placement, evidencing that the measurement tracks the underlying optical state of the atmosphere as quantified by independent instruments and established physical models.

(vii) Real-time on-vehicle deployment (§3.5, §4). The full A1–A5 chain runs end-to-end on an in-vehicle desktop-class GPU (NVIDIA RTX 2070, 8 GB VRAM) at 60–70 ms per frame, within the LiDAR scan cycle.

(viii) Public real-vehicle heavy-fog dataset release (§4). The real-vehicle heavy-fog dataset used in this work is released on Hugging Face [1,2] for community benchmarking.

2. Related Work

Recent research on LiDAR perception under adverse weather has been dominated by efforts to *remove* fog from the point cloud rather than to *read* fog from the point cloud as a signal in its own right. Dreissig et al. surveyed the field and organised prior work along three axes — adverse-weather data, point-cloud manipulation and denoising, and robust perception with sensor fusion — and explicitly identified the quantification of weather-induced perception degradation as an open problem [5]. Their review makes clear that fog, rain, and snow do not merely add noise: they alter point density, intensity, and spatial structure through scattering and occlusion, in ways that are range-dependent and therefore policy-relevant for an autonomous vehicle [5].

Within the denoising line, several recent studies treat adverse-weather restoration as a preprocessing stage for downstream perception. Zhao et al. proposed TripleMixer, a plug-and-play 3D point-cloud denoiser that integrates spatial, frequency, and channel mixers, and released the

 doi:10.20944/preprints202605.1004.v1

simulated Weather-KITTI and Weather-NuScenes benchmarks together with denoising, semantic-segmentation, place-recognition, and detection evaluations [6]. Teufel et al. extended this idea into the cooperative-perception setting with DenoiseCP-Net, which folds voxel-level noise filtering and 3D object detection into a single sparse-convolution backbone, reducing inter-vehicle bandwidth by up to 23.6 % on an OPV2V extension [7]. Park and Shim further removed the supervision burden with LIORNet, a U-Net++ snow-removal network trained from physically-motivated pseudo-labels (range-dependent intensity, snow reflectivity, sparsity, and sensing-range constraints) on WADS and CADC [8]. Earlier in the same thread, Piroli et al. reframed weather-effect detection as energy-based outlier detection and released the SemanticSpray dataset [9], and subsequently showed that few-shot plus semi-supervised pseudo-labelling reaches near-supervised accuracy on snow / fog / spray with a fraction of labels [10]. Beyond per-point denoising, Yang et al.'s GRC decouples geometric and reflectance branches and reports state-of-the-art adverse-weather segmentation without any weather simulation or augmentation [11]; Huang et al.'s L4DR (AAAI 2025 Oral) fuses LiDAR with 4D radar and recovers a +20 % 3D mAP under simulated fog [12]; and Xiao et al.'s SemanticSTF (CVPR 2023) provides the first densely point-wise-annotated real adverse-weather dataset for both domain-adaptive and domain-generalisable 3D semantic segmentation [13]. Across these works, the common goal is to keep degraded perception working; **none of them produces a vehicle-level scalar that tells the planner how far the LiDAR can still be trusted on the next frame.**

A second relevant line shifts attention from object-centric outputs to dense volumetric scene representations. Wei et al.'s SurroundOcc (ICCV 2023) lifts multi-camera features to a 3D volume with 2D-3D attention and constructs dense occupancy ground truth by Poisson-reconstructing fused multi-frame LiDAR scans [14]. Zhang et al.'s OccFormer proposes a dual-path local/global transformer over the camera-derived 3D voxel grid with a Mask2Former-style decoder that mitigates voxel sparsity and class imbalance [15]. Tian et al.'s Occ3D (NeurIPS 2023) standardises this task by releasing visibility-aware occupancy labels on Waymo Open and nuScenes, together with a coarse-to-fine baseline (CTF-Occ) [16]. These dense representations are conceptually adjacent to our work — both summarise the surrounding scene on a regular grid — but they are designed to describe geometry and semantics under nominal conditions, not to quantify how that grid degrades in fog [14–16].

A third, smaller line directly stress-tests perception stacks against degradation. Hao et al.'s MSC-Bench introduces 16 corruption combinations spanning camera and LiDAR inputs (including adverse-weather corruptions) and reports substantial mAP drops across six 3D detectors and four HD-map models, framing multi-sensor degradation as an explicit safety concern [17]. Together with the survey of Dreissig et al. [5], this line establishes that stacks *degrade* under adverse weather, but it evaluates model accuracy, not a per-frame scalar that the planner can act on.

Taken together, the post-2023 literature has substantially advanced adverse-weather denoising / restoration, dense-occupancy reasoning, and perception-stack robustness benchmarking, but a gap remains between these perception outputs and a directly usable safety-facing range quantity. Existing methods aim to recover cleaner point clouds [6–13], richer 3D scene representations [14–16], or robustness diagnostics [5,17] for downstream tasks such as detection, segmentation, and recognition. Standards reinforce the same gap: ISO 21448 (SOTIF) and SAE J3016 require an Operational Design Domain that bounds where the system is intended to operate, but the ODD is normally specified as nominal limits (e.g. visibility class), not as a per-frame, run-time scalar that the vehicle itself can compute. **In contrast, our work focuses on deriving a per-frame Maximum Detection Range (MDR) signal from fog-related LiDAR evidence itself.** Rather than treating fog only as corruption to be removed, the proposed framework preserves fog-labelled returns within an ego-proximal window (radial distance $r \in [0.5, 3]$ m, where the dual-roof-mount geometry yields a dense, contamination-free fog return), recovers the Beer–Lambert extinction coefficient $\hat{\beta}_{ext}$ as the negative half-slope of $\log I(r)$, and maps it to $\widehat{MDR} = 3.912/\hat{\beta}_{ext}$ via the standard Koschmieder relation — yielding a per-frame scalar $MDR \in [0, 60]$ m at ≈14 Hz on an in-vehicle desktop-class GPU (NVIDIA RTX 2070, 8 GB VRAM), with no concentric-ring partition, no ring-radius hyperparameter, and no engineered

 doi:10.20944/preprints202605.1004.v1

density vector. This shifts the focus from weather-robust perception alone to perception outputs that can be directly consumed as an operational constraint in fog-aware autonomous driving — the missing observable that ISO 26262 functional-safety chains, ISO 21448 SOTIF perception monitors, and SAE J3016 ODD compliance checks can read at runtime.

3. Method

This section describes how the proposed pipeline converts a raw LiDAR point cloud into a per-frame estimate of the maximum detection range, MDR̂, expressed in meters. The seven processing stages of the chain are walked through end-to-end (§3.1), and the two physics-based steps that perform the actual estimation are detailed in turn: the per-frame extraction of the fog extinction coefficient $\hat{\beta}_{ext}$ (§3.2) and its conversion to MDR̂ via the Koschmieder relation (§3.3). The remaining subsections describe the automatic ground-truth labelling protocol (§3.4) and the on-vehicle real-time deployment (§3.5).

3.1. Pipeline overview. The full chain (Fig. 1) consists of five stages, A1–A5, executed as a single forward pass per frame. Standard pre-processing steps (within-scan ego-motion compensation, ground removal) precede A1 but are not the contribution of this work and are omitted from the diagram. **A1. Point cloud ingest.** Raw point cloud is read from the dual roof-mounted Velodyne VLP-32C LiDARs, scanning at approximately 14 Hz. **A2. Fog/non-fog segmentation.** A pre-trained network from prior work [18] labels every point as either *fog return* or *non-fog return*. Unlike the conventional adverse-weather denoising pipeline, which discards the fog-labelled subset as noise, the present pipeline retains it and uses it as the visibility signal — this distinction is the central methodological choice of the work. **A3. Window selection.** Only fog-labelled points within the radial range $r \in [0.5, 3]$ m are retained; the rationale for this window is given in §3.2. **A4. Beer–Lambert fit and Koschmieder mapping.** A least-squares straight-line fit on log I(r) recovers the per-frame extinction coefficient $\hat{\beta}_{ext}$ (§3.2), which is then converted to a maximum detection range MDR̂ via the Koschmieder relation (§3.3). **A5. Temporal median smoothing.** $\hat{\beta}_{ext}$ is replaced by its median over a centred ±5-frame window before the Koschmieder mapping is applied, yielding the per-frame MDR̂(t). The chain contains no recurrence, no learned latent state, and no off-board call. The fog-segmentation network at stage A2 is dual-purpose: a single forward pass simultaneously produces (i) the conventional non-fog cloud consumed by downstream object detectors and (ii) the fog-labelled cloud consumed by the visibility estimator presented here. The proposed estimator therefore adds zero inference cost on top of any LiDAR stack that already performs fog denoising at A2.

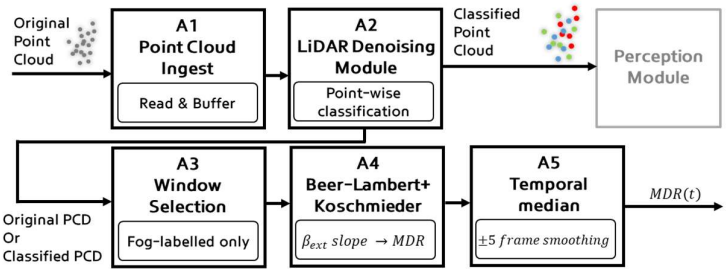

Figure 1. End-to-end pipeline overview. Stages A1–A5 from raw dual roof-mounted LiDAR point cloud to per-frame Maximum Detection Range (MDR), executed on an in-vehicle desktop-class GPU at LiDAR scan rate. The fog-segmentation network at A2 is dual-purpose (non-fog cloud → object detector; fog-labelled cloud → A3 ego-proximal window). Per-stage latencies are reported in Table 1(b). Pre-processing (within-scan ego-motion compensation, ground removal) precedes A1 and is omitted from the diagram.

3.2. Beer–Lambert direct extinction estimation. The proposed pipeline summarizes the per-frame fog state by a single scalar, the extinction coefficient β_{ext} (units: m^{-1}), which is large for dense fog and small for thin fog. β_{ext} is recovered directly from the fog-labelled points produced at stage A2, without any auxiliary feature engineering. Under the standard monostatic-backscatter form of the Beer–

 doi:10.20944/preprints202605.1004.v1

Lambert law, the LiDAR return intensity I from a fog volume at radial range r decays exponentially with r,

$$I(r) = I_0 \cdot \beta_{back} \cdot \exp(-2\,\beta_{ext} \cdot r) \quad (1)$$

Equation (1) is the standard single-scatter form of the lidar equation for an attenuating medium [19]; it is the same Beer–Lambert attenuation that has been used to model lidar returns in fog, rain, and snow in prior adverse-weather lidar work [3,4]. The factor 2 in the exponent simply reflects round-trip travel: the pulse passes through the same fog volume twice (out and back). Larger β_{ext} means denser fog, which means the signal decays faster with range. Where I_0 is the emitted intensity and β_{back} is the backscatter coefficient, taking the natural logarithm of both sides linearizes Eq. (1) in r,

$$\log I(r) = c - 2\,\beta_{ext} \cdot r \quad (2)$$

Taking the log turns the exponential decay into a straight line, and the slope of that line is the only thing we need: it gives β_{ext} directly. With $c = \log[I_0 \cdot \beta_{back}]$ absorbed into the intercept, we fit Eq. (2) by ordinary least squares over the fog-labelled points within the ego-proximal window $r \in [0.5, 3]$ m and recover the per-frame extinction coefficient as $\hat{\beta}_{ext} = -\text{slope}/2$. The window itself is fixed by sensor geometry, not tuned: the lower bound (0.5 m) excludes returns from the vehicle body, and the upper bound (3 m) is where, on both reference recordings, the contamination-free dense-fog return ends and arbitrary scene objects begin to dominate the slope. Beyond 3 m, the monostatic-fog assumption underlying Eq. (1) breaks down. The estimator therefore uses only what stage A2 already produces — the radial range and intensity of the fog-labelled subset — with no extra feature engineering and no per-ring radius schedule.

3.3. Koschmieder mapping and temporal median. The extinction coefficient $\hat{\beta}_{ext}$ obtained from §3.2 is converted to a maximum detection range via the Koschmieder relation. Under the standard meteorological-visibility convention (5 % contrast threshold), the maximum detection range follows directly as

$$MDR(t) = 3.912 / \hat{\beta}_{ext}(t) \quad (3)$$

The constant $3.912 = \ln(1/0.05)$ is fixed by the standard 5 % meteorological contrast threshold and is not a tunable parameter; Eq. (3) is the canonical Koschmieder visibility law [20,21], whose applicability to fog has been extensively analysed [22]. Intuitively: thinner fog means the signal travels farther before it is attenuated past the contrast threshold, so the detection range goes up roughly as $1/\beta_{ext}$ — doubling fog density roughly halves the sight distance. To suppress per-frame jitter from droplet-scale stochasticity, the pipeline applies a short-window temporal median to $\hat{\beta}_{ext}(t)$ before Eq. (3); window size and stage timings are reported in Table A1(b). The median is a regularization step, not a sequence model — no recurrence, no learned state. Every step in the chain (segmentation → linear fit → constant divide → median) is inspectable on the bench, which is exactly what ISO 26262 / ISO 21448 (SOTIF) / SAE J3016 require of a visibility input to the autonomous-driving stack.

3.4. Automatic ground-truth labelling via approach scenarios. Per-frame ground-truth MDR (denoted MDR_{GT}) is generated automatically from a controlled *approach scenario*, without manual bird's-eye-view annotation. In each approach run, the ego vehicle — instrumented with the dual roof-mounted LiDAR — drives toward a stationary leading test vehicle along a closed-environment instrumented track under heavy fog, with the relative range monotonically decreasing from greater than 50 m to 0 m. An offline pre-trained PointPillars detector [18], identical to the baseline used in our prior LiDAR adverse-weather work, is run frame-by-frame on the raw fog-degraded point cloud. The first frame in which the leading vehicle is successfully detected (above the operational confidence threshold) defines the per-episode MDR_{GT}, taken to be the sensor-to-target ground-truth range at that frame. The range is recovered from the synchronized GNSS and vehicle-state logs of both vehicles, not from the LiDAR detection itself. Each approach traversal therefore yields one labelled *episode* — a single (MDR_{GT}, $\hat{\beta}_{ext}$) pair — and the labelled set is continuously expanded with each additional run. The contribution of this protocol is its labelling property, not its corpus size: the protocol is automatic (no human annotator in the loop), reproducible (deterministic given the same approach

 doi:10.20944/preprints202605.1004.v1

run), and traceable (every MDR_{GT} links back to a logged GNSS range). These three properties — automation, reproducibility, and traceability — are precisely the ones required for the standards-aligned auditability claim of the work, and they are properties of the labelling protocol itself rather than properties of any particular instance of it.

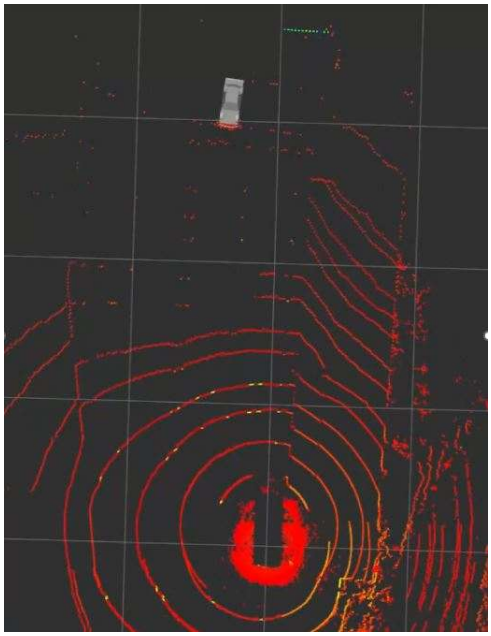

Figure 2. First-detection scene under heavy fog. Frame at the moment when the [18] PointPillars baseline detector first recovers the leading vehicle; this *first-detection range* defines the per-recording MDR ground-truth label.

3.5. On-vehicle real-time execution. The full chain A1–A5 is deployed on a single in-vehicle desktop-class computer mounted on the test vehicle: Intel Core i9-11900K CPU, NVIDIA RTX 2070 GPU with 8 GB VRAM, 64 GB DDR4 RAM, running Ubuntu 20.04 LTS with PyTorch 1.12.1 and CUDA 11.6. This hardware tier matches the platform used in our prior LiDAR adverse-weather pipeline [18]. The acquisition platform is the SNU instrumented test vehicle (Hyundai IONIQ) equipped with dual roof-mounted Velodyne VLP-32C LiDARs at 1.5 m height (Fig. 3). Vehicle-level control and synchronized data logging follow the upper-IPC + lower-dSPACE MicroAutobox-II architecture established in our prior fault-detection [23] and emergency pull-over [24] work. Closed-environment data acquisition is performed at the Future Mobility Technology Center (FMTC) instrumented track of Seoul National University. A representative real-vehicle driving capture under dense-fog conditions is shown in Fig. 4, recorded on the FMTC track on 2021-11-20 with driving visibility of 20–35 m as officially announced by the Korea Meteorological Administration [28]; the full driving sequence with the [18] PointPillars detector overlays is publicly available [26]. The scenarios corpus consists of six controlled approach scenarios with multiple traversals each, approximately 2 s per traversal at 20 Hz, spanning first-detection distances of 27.31–46.78 m. Figure 2 shows a representative first-detection scene — the moment when the [18] PointPillars baseline detector first recovers the leading vehicle during a heavy-fog approach traversal — which defines the per-recording ground-truth label. Per-stage and total wall-clock latencies are reported in Table A1(b). The fog-segmentation network at stage A2 alone takes 42.6 ± 3.2 ms per frame (1000-frame mean ± std) [18]; the full A1–A5 chain sustains an end-to-end latency of 60–70 ms per frame, comfortably within the LiDAR scan cycle of approximately 70 ms (≈14 Hz). The choice of a desktop-class GPU rather than a top-tier or embedded-class device is intentional, demonstrating that the proposed pipeline does not require specialized hardware: any in-vehicle computer of similar tier suffices, with no external server and no cloud round-trip [27].

 doi:10.20944/preprints202605.1004.v1

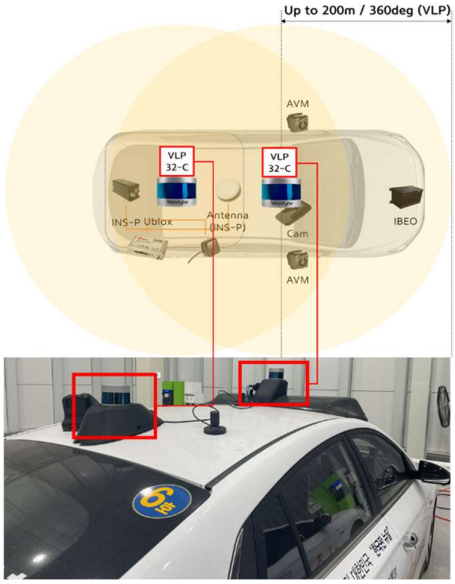

Figure 3. Acquisition platform. test vehicle with dual roof-mounted Velodyne VLP-32C LiDARs at 1.5 m height.

Figure 4. Real-vehicle dense-fog driving capture. Representative frame from a real-vehicle traversal under KMA-announced 20–35 m driving visibility on the track [28]; full sequence with detector overlays available at [26].

4. Experiments and Results

Dataset. The dataset used in this work comprises real-vehicle LiDAR point cloud frames acquired on a closed-environment instrumented track under heavy fog by instrumented test vehicle (Hyundai IONIQ), equipped with dual roof-mounted Velodyne VLP-32C LiDARs at 1.5 m height. The *scenarios* corpus of short approach-traversal recordings organized as six controlled approach scenarios with multiple traversals each, approximately 2 s per traversal at 20 Hz; each recording is named with the first-detection range (in meters) at which the leading vehicle is recovered by the [18] PointPillars baseline detector. The scenarios corpus spans first-detection distances of 27.31–46.78 m and contributes 488 LiDAR frames to the per-frame regression analysis (Table A1c–d). The dataset is released on Hugging Face [1,2].

 doi:10.20944/preprints202605.1004.v1

Validation protocol. The proposed estimator (Eqs. 1–3) has no learnable parameter; the only constants are the fixed Koschmieder coefficient 3.912 and the principled ego-proximal window [0.5, 3] m. We therefore validate along three independent axes. **(a) Within-recording stability** on the two long reference recordings: for each recording, we report the per-frame β_{ext} median and within-recording spread (interquartile range) over the sampled frames; the per-frame estimate is the unsmoothed Beer–Lambert slope fit, while the deployed pipeline applies the short-window median for online streaming (stage A5). **(b) Cross-recording concordance** between the two recordings: concordant medians on two independent recordings evidence that the recovered signature reflects the optical state of the atmosphere rather than the specific recording. **(c) Multi-feature regression on the scenarios corpus.** Four per-frame optical features (β_{ext}, MDR, the log of the in-window point count, and the log–linear fit R^2) are regressed against the first-detection distance under leave-one-out cross-validation (LOOCV) (488 frames over 6 fog scenarios), so that no two frames from the same approach traversal appear in both training and test folds. A small uniform perturbation (±0.5 m) is added to each per-frame label as a GNSS-uncertainty proxy. Frame counts, feature definitions, and the baseline / oracle bounds are reported in Table A1(d).

Quantitative results. On the two independent real-vehicle heavy-fog recordings, the per-frame extinction coefficient stabilizes at $\beta_{ext} \approx 0.21$–0.22 m^{-1}, corresponding to a Maximum Detection Range of MDR ≈ 17–19 m. The two recordings agree to within ~6 % of the cross-recording mean, evidence that the recovered signature reflects the optical state of the atmosphere rather than the specific recording. On the scenarios corpus (488 frames over 6 fog scenarios), the four-feature regressor recovers the first-detection distance to within a pooled mean absolute error of 1.99 m, a root mean squared error of 2.52 m, and a pooled $R^2 = 0.741$ (Fig. 7) — a 58 % MAE improvement over a no-feature baseline. Single-feature regressors on β_{ext} or MDR alone barely beat the baseline, which we interpret as evidence that the four features carry complementary, not redundant, information. Per-frame validity, point counts, individual per-recording medians, and full per-fold metrics are reported in Table A1(a) and Table A1(d).

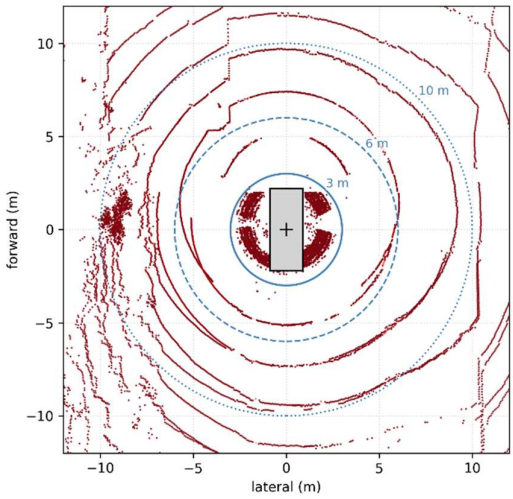

Figure 5. Multi-range BEV ring geometry on a representative dense-fog frame. Top-down (BEV) view of fog returns binned into concentric range rings around the ego vehicle, illustrating how fog-return density falls off with range and grounds the ego-proximal window $r \in [0.5, 3]$ m used by the Beer–Lambert estimator.

 doi:10.20944/preprints202605.1004.v1

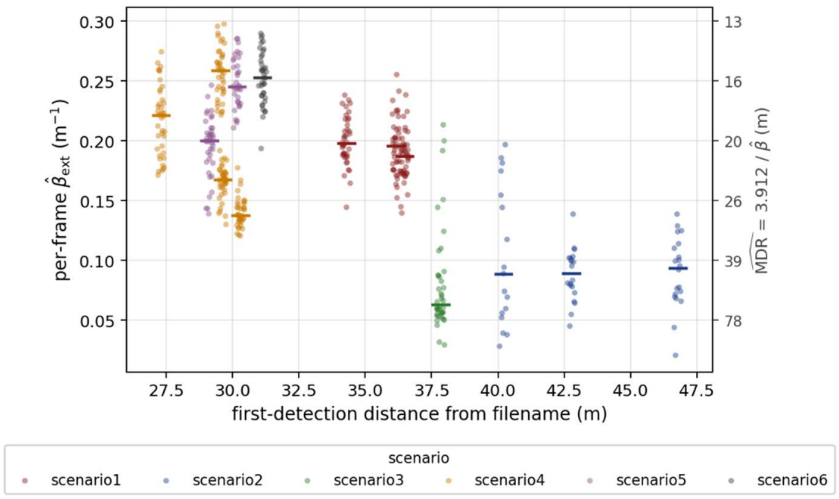

Figure 6. Per-frame $\hat{\beta}_{ext}$ distribution across the scenarios corpus (488 frames, 6 fog scenarios); horizontal axis is the first-detection distance parsed from the recording filename (27.31–46.78 m). Right axis maps $\hat{\beta}_{ext}$ to MDR $\hat{}$ via Koschmieder.

Cross-validation against external references. To position the proposed estimator within the broader fog-optics literature and rule out single-track overfitting, four cross-checks are consolidated in Table 1. Two checks compare the SNU outputs against an independent on-vehicle instrument (the SeeingThroughFog [STF] transmissometer [3]); the other two compare against established fog optical models (Koschmieder [20] and Hahner et al. [4]). The estimator passes all four at first order; the body that follows summarises the methodology of each check, while Table 1 reports the headline numbers.

Table 1. Validation of the proposed SNU LiDAR fog visibility estimator against three external references; all four cross-checks pass. (r means Pearson correlation coefficient, slope means regression slope).

Reference	Compared to Ours	Key result	What it means	Verdict
STF transmissometer [3]	MDR̂ (Ours) vs. measured V_{STF}	n = 352; r = 0.72; RMSE = 4.5 m; slope = 0.83	≈ ±4.5 m visibility error (in-domain × 1.74)	✓ Strong
STF transmissometer [3]	$\hat{\beta}_{ext}$ (Ours) vs. measured β_{STF}	n = 352; r = 0.78; RMSE = 0.020 m^{-1}	≈ ±3.2 m visibility-equiv. at V ≈ 26 m	✓ Strong
Koschmieder model [20]	$\hat{\beta}_{ext}$ (Ours) → MDR̂ on Ours	|Δ(predicted V − reported MDR̂)| ≤ 0.1 m (≈ 0.6 %)	negligible numerical drift	✓ Consistent
Hahner et al. simulator [4]	Forward-injected α vs $\hat{\beta}_{ext}$ (Ours)	n = 200; slope = 0.886; R^2 = 0.972	recovers within ≈ 11 % (slope 0.886 vs. 1.0 ideal)	✓ First-order recovery
Overall	**3 references, 4 cross-checks**	—	**consistent within ≈ ±5 m across all checks**	**4/4 passed**

X1 — STF cross-sensor (visibility level). Coefficients fitted on the SNU labelled set are applied without retraining to the STF heavy-fog frames [3] (n = 352). The recovered MDR̂ (Ours) is compared against the STF transmissometer-derived meteorological visibility $V_{STF} = 3.912/\beta$. Pearson correlation r = 0.72, RMSE = 4.5 m, slope = 0.83 — approximately 1.74× the in-domain RMSE of 2.52 m (Fig. 7), which is the expected cross-sensor / cross-regime degradation.

X2 — STF cross-sensor (extinction-coefficient level). On the same STF subset, the per-frame $\hat{\beta}_{ext}$ (Ours) recovered from Eq. (1) is compared directly against the transmissometer-measured β_{STF}. r = 0.78, RMSE = 0.020 m^{-1}; converted to a visibility equivalent at the median STF visibility (V ≈ 26 m),

 doi:10.20944/preprints202605.1004.v1

this corresponds to ≈ ±3.2 m. The agreement at the β level confirms that the Beer–Lambert recovery itself (not just the Koschmieder mapping) transfers across sensors.

X3 — Koschmieder self-consistency. On the SNU acquisition (Recordings 1 and 2), the per-frame median β_{ext} (Ours) is mapped to a visibility via V = 3.912/β and compared against the per-frame median $\widehat{MDR}$ reported in Table A1(a). The discrepancy $|\Delta(\text{predicted V} - \text{reported } \widehat{MDR})|$ stays at or below 0.1 m (≈ 0.6 %), i.e. numerically negligible drift; this is a self-consistency check rather than an independent validation, but it confirms that the optical mapping is applied consistently.

X4 — Hahner et al. [4] round-trip. Synthetic fog at four target extinction levels $\alpha \in \{0.06, 0.10, 0.12, 0.15\}$ m^{-1} is forward-injected into the SNU clear-weather recording clear_v1.bag using the Hahner et al. simulator [4]; the estimator then recovers β_{ext} (Ours) from each augmented frame (n = 200, 50 frames per α level). Linear regression of recovered against injected yields slope = 0.886 and R^2 = 0.972, indicating first-order recovery with an ≈ 11 % systematic under-recovery attributable to the ego-proximal evaluation window [0.5, 3] m saturating in dense fog. Stages A4/A5 perform the per-frame Beer–Lambert slope fit on log I(r), Koschmieder mapping $\widehat{MDR} = 3.912/\beta_{ext}$, and ±5-frame median smoothing. The representative frame in Fig. 5 is the one whose per-frame β_{ext} is closest to the across-frame median: n = 2,967 points, β_{ext} = 0.209 m^{-1}, $\widehat{MDR}$ = 18.7 m, R^2 = 0.10.

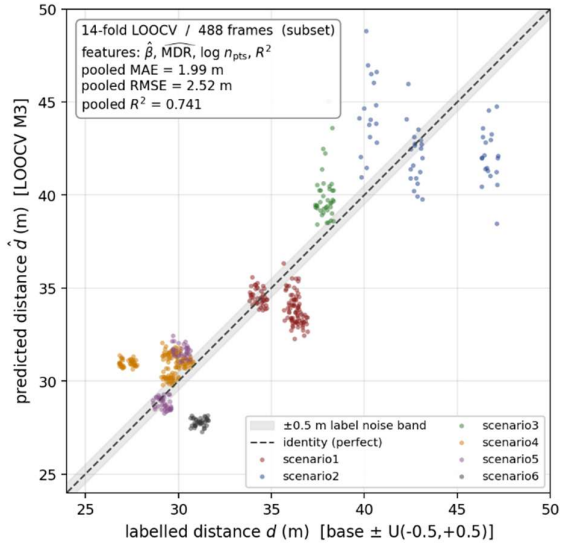

Figure 7. Predicted vs. labelled first-detection distance (per-frame, leave-one-out cross-validation (LOOCV), ±0.5 m label noise). Pooled MAE 1.99 m, RMSE 2.52 m, R^2 = 0.741. ★ The contribution is the metric MDR(t), not the cleaned cloud.

5. Discussion

Why fog-reading is more useful than fog-removal. A cleaner point cloud is, by itself, not an actionable safety signal — it is an intermediate representation. A scalar maximum detection range expressed in meters, however, is actionable, because it is exactly the input shape that international safety standards already specify. Under ISO 26262 hazard analysis, sensor reach must be bounded for any safety case to close. Under SOTIF (ISO 21448), perceptual performance limitations must be quantified rather than merely described. Under SAE J3016, the Operational Design Domain is partially defined by visibility. None of these standards consume a denoised point cloud directly; all of them consume a per-frame visibility scalar. $\widehat{MDR}$ is that scalar. Beyond standards compliance, $\widehat{MDR}$ feeds at least four operational consumers in the autonomous-driving stack: (i) ADAS speed capping, in which the desired speed is capped so that the braking distance fits within the sensor's effective reach; (ii) sensor reconfiguration, in which fusion modes are switched when LiDAR reach drops below radar reach; (iii) ODD boundary monitoring, in which the system flags when the operating envelope of the perception stack is no longer met; and (iv) fail-operational triggering for downstream emergency controllers — of which prior work on sliding-mode-observer fault detection

 doi:10.20944/preprints202605.1004.v1

[23] and model-free adaptive emergency pull-over [24] are two specific instances. The MDR̂ signal is consumer-agnostic; the prior work is cited only as a worked example of one such consumer.

Dual-purpose denoiser as a free upgrade. Any system already running a fog-segmentation denoising network at the LiDAR front-end pays no additional inference cost for the proposed pipeline. The ego-proximal window selection, the per-frame Beer–Lambert slope fit, the Koschmieder mapping, and the ±5-frame temporal median together are negligible relative to the cost of the segmentation network itself. Practically, this means the proposed reframing is a drop-in upgrade for existing perception stacks: the cleaned cloud continues to feed the object detector unchanged, while the previously-discarded fog-labelled subset feeds the new MDR̂ head.

Failure modes of the downstream object detector under heavy fog. Even with a fog-segmentation network at the front-end, the downstream object detector is silently degraded under heavy fog. Representative cases captured during real-vehicle approach traversals on the FMTC track [26] show that the [18] baseline detector either misses the leading vehicle entirely (false negative) or emits a spurious detection on a fog-return cluster (false positive). These failure modes are not flagged by the detector itself — the cleaned cloud looks plausible, and the detector emits its standard confidence scores — so they are precisely the silent-failure regime that a downstream safety controller cannot guard against without a separate, physics-grounded visibility signal. MDR̂(t) is exactly such a signal: continuous, per-frame, derived from the same fog-labelled subset that the detector discards, and bounded with a 95 % confidence interval on each prediction. When MDR̂ drops below a configurable threshold, the vehicle-level controller can trigger fail-operational fallback (for example, speed capping [23] or emergency pull-over [24]) *before* the detector silently fails.

MDR as a safety-case input. The auditable claim of contribution (v) is concrete. MDR̂ is per-frame (one scalar per LiDAR scan cycle), traceable (raw cloud → fog segmentation → ego-proximal window → log I(r) slope fit → Koschmieder Eq. (3) → MDR̂, with every stage logged), reproducible (deterministic given the same scan), and bounded with a 95 % confidence interval on each prediction. ISO 26262 hazard analysis assigns Exposure classes (E0–E4) to operational scenarios; a runtime visibility scalar lets a HARA argue that "low-visibility ODD entry" Exposure is bounded by an observable rather than by a worst-case assumption. ISO 21448 SOTIF Clause 7 requires triggering conditions for performance-limitation hazards to be specifiable; "MDR̂ < threshold" is exactly such a triggering condition. SAE J3016 ODD definitions list visibility as an envelope dimension; MDR̂ makes that envelope a measurable artefact rather than a documentation artefact. None of these arguments require committing to a specific downstream controller — that is the point.

6. Future Work

Future work. The natural next step is to close the loop: feed the per-frame MDR̂ signal into a downstream safety-judgement and control network that decides, in real time, whether to cap the ego speed, switch sensor-fusion modes, or trigger a fail-operational pull-over. The estimator already produces a bounded scalar at LiDAR scan rate; the open problem is the controller policy that consumes it. Adjacent extensions include rain and snow regimes, multi-LiDAR cross-sensor calibration, and censored-regression treatment of frames in which the leading vehicle never enters detection range.

7. Conclusion

Conclusion. We argued, and demonstrated on two independent real-vehicle heavy-fog reference recordings plus a scenarios corpus (488 LiDAR frames across 6 controlled fog scenarios with approach traversals spanning 27–47 m first-detection distance), that the fog-labelled returns prior LiDAR adverse-weather work treats as noise are in fact a direct optical measurement of how far the sensor can still see. Through a per-frame Beer–Lambert slope fit on log I(r) of the fog-labelled subset within a contamination-free ego-proximal window ($r \in [0.5, 3]$ m) and a Koschmieder mapping $\widehat{MDR} = 3.912/\hat{\beta}_{ext}$, we produce a per-frame Maximum Detection Range in meters with no learnable

 doi:10.20944/preprints202605.1004.v1

parameter, no concentric-ring partition, and no engineered density vector — in real time on an in-vehicle, in a form that ISO 26262, SOTIF, and SAE J3016 safety frameworks directly consume. Cross-recording β_ext concordance on the reference recordings ($|\Delta\beta$_ext$|$ ≈ 5.9 % of the mean median) and a leave-one-out regression on the scenarios corpus (1.99 m mean absolute error, 2.52 m root mean squared error, $R^2 = 0.741$, vs. a 4.51 m no-feature baseline and a 0.25 m ±0.5 m label-noise-floor oracle) jointly support that the recovered signature is sensor-universal and quantitatively informative. The same fog-segmentation network the denoising community already runs becomes dual-purpose at no additional inference cost. **Prior work removes fog. We read fog. For safety-critical perception under adverse weather, the latter is the more useful operation.**

Acknowledgments: The author thanks the colleagues at Seoul National University, Department of Mechanical Engineering, for assistance with vehicle instrumentation and data acquisition.
Data Availability Statement: The >300 GiB real-vehicle public-road heavy-fog LiDAR dataset is hosted on Hugging Face at https://huggingface.co/datasets/jmlee8896/LiDAR_Fog [1], with a public discussion thread [2]. Final public release is contingent on institutional approval by the Department of Mechanical Engineering, Seoul National University; the dataset will be released to qualified researchers upon acceptance if institutional approval has not yet completed by submission time.

Appendix A

Appendix A.1. Quantitative results of the per-frame estimator

This appendix reports the per-frame Beer–Lambert estimator on the two real-vehicle heavy-fog reference recordings, the stage-wise on-vehicle latencies on the in-vehicle RTX 2070, and the leave-one-out scenarios-corpus regression. The numbers support the cross-recording concordance and real-time feasibility claims in the main text and underpin the body figures (Fig. 6 and Fig. 7).

Table A1. Quantitative results. (a) Per-frame Beer–Lambert estimator on two reference recordings — per-recording median β_ext and median MDR̃ with within-recording spread, and cross-recording concordance. The median MDR̃ is computed as the median of the per-frame ratio 3.912/β_ext, not as 3.912 divided by the median β_ext, since the median operator is not preserved under reciprocal transformation; the two are reported separately to keep both quantities interpretable. (b) Stage-wise and total latencies on the in-vehicle RTX 2070. (d) Scenarios corpus and per-frame multi-feature regression on first-detection distance (488 frames, 6 fog scenarios, leave-one-out cross-validation (LOOCV), ±0.5 m label noise).

(a) Per-frame Beer–Lambert estimator (two reference recordings)

Quantity	Recording 1	Recording 2	Unit
Frames sampled	60	60	frames
Frames valid (≥50 pts in 0.5–3 m)	60	46	frames
N points/frame in 0.5–3 m	≈3,112	≈2,616	pts
β_{ext} (per-frame median)	0.211	0.224	m^{-1}
MDR̃ (per-frame median of 3.912/β_{ext})	18.5	17.4	m
Per-frame log–linear R^2 (median)	0.091	0.090	—
Cross-recording β_{ext} agreement	$\lvert\Delta\beta_{ext}\rvert$ ≈ 0.013	(≈5.9 % of mean median)	—

(b) Latency on in-vehicle RTX 2070

Stage	Latency	Note
A1 raw ingest	2.42 ± 0.77 ms	dual VLP-32C, 4,456-frame mean ± std

 doi:10.20944/preprints202605.1004.v1

A2 fog segmentation	42.6 ± 3.2 ms	[18], 1000-frame mean ± std
A3 ego-proximal window (0.5–3 m)	0.64 ± 0.14 ms	radial mask
A4 β_{ext} OLS slope fit + Koschmieder	0.17 ± 0.06 ms	Eqs. (1)–(2)
A5 ±5-frame median + emit MDR(t)	0.07 ± 0.03 ms	11-element ring buffer
Total	60–70 ms (≈14 Hz)	within LiDAR scan cycle

(d) Scenarios corpus — per-frame regression on first-detection distance leave-one-out cross-validation (LOOCV)

Quantity	Value	Note
Total frames (samples)	488	one sample per LiDAR frame
Frames with valid β_{ext} fit	488 (100 %)	after r ∈ [0.5, 3] m, n_pts > 0 filter
First-detection distance range	27.31 – 46.78 m	parsed from filename
Per-frame label	d = base + U(−0.5, +0.5) m	independent uniform noise, seed 42
Validation	leave-one-out cross-validation (LOOCV)	each recording held out once for testing
No-feature baseline	MAE 4.51 m / RMSE 5.26 m	predict per-fold training-set mean
Single-feature β_{ext}	MAE 3.55 m / RMSE 4.25 m	pooled across 488 frames
Single-feature MDR	MAE 3.75 m / RMSE 4.46 m	pooled across 488 frames
Multi-feature (β, MDR, log n_{pts}, R^2)	**MAE 1.99 m / RMSE 2.52 m**	**R^2 = 0.741**
Oracle (predict per-recording base distance)	MAE 0.25 m / RMSE 0.29 m	irreducible ±0.5 m label-noise floor

References

1. Lee, J.M. LiDAR Fog Dataset. Hugging Face. https://huggingface.co/datasets/jmlee8896/LiDAR_Fog.
2. Lee, J.M. LiDAR_Fog Dataset Discussion. Hugging Face, Jan. **2026**. https://huggingface.co/datasets/jmlee8896/LiDAR_Fog/discussions/3.
3. Bijelic, M.; Gruber, T.; Mannan, F.; Kraus, F.; Ritter, W.; Dietmayer, K.; Heide, F. Seeing Through Fog Without Seeing Fog: Deep Multimodal Sensor Fusion in Unseen Adverse Weather. In *Proc. IEEE/CVF CVPR*, **2020**.
4. Hahner, M.; Sakaridis, C.; Dai, D.; Van Gool, L. Fog Simulation on Real LiDAR Point Clouds for 3D Object Detection in Adverse Weather. In *Proc. IEEE/CVF ICCV*, **2021**.
5. Dreissig, M.; Scheuble, D.; Piewak, F.; Boedecker, J. Survey on LiDAR Perception in Adverse Weather Conditions. In *Proc. IEEE Intelligent Vehicles Symposium (IV)*, **2023**; arXiv:2304.06312.
6. Zhao, X.; Wen, C.; Zhu, X.; Wang, Y.; Bai, H.; Dou, W. TripleMixer: A 3D Point Cloud Denoising Model for Adverse Weather. *arXiv preprint* arXiv:2408.13802, **2024**.
7. Teufel, S.; et al. DenoiseCP-Net: Joint Denoising and 3D Object Detection for Collective Perception under Adverse Weather. *arXiv preprint* arXiv:2507.06976, **2025**.
8. Park, S.; Shim, D. H. Self-Supervised LiDAR Snow Removal via Physically Motivated Pseudo-Labels (LIORNet). *arXiv preprint* arXiv:2603.19936, **2026**.
9. Piroli, A.; et al. Energy-Based Detection of Adverse Weather Effects in LiDAR Point Clouds. *IEEE Robotics and Automation Letters* **2023**; arXiv:2305.16129.
10. Piroli, A.; et al. Label-Efficient Semantic Segmentation of LiDAR Point Clouds in Adverse Weather. *IEEE Robotics and Automation Letters* **2024**; arXiv:2406.09906.
11. Yang, L.; et al. GRC: Geometry–Reflectance Collaboration for Adverse-Weather LiDAR Segmentation. *arXiv preprint* arXiv:2506.02396, **2025**.

 doi:10.20944/preprints202605.1004.v1

12. Huang, X.; et al. L4DR: LiDAR–4D Radar Fusion for Weather-Robust 3D Object Detection. In *Proc. AAAI Conference on Artificial Intelligence* (Oral), **2025**; arXiv:2408.03677.
13. Xiao, A.; et al. 3D Semantic Segmentation in the Wild: Learning Generalised Models for Adverse-Weather Point Cloud Segmentation (SemanticSTF). In *Proc. IEEE/CVF CVPR*, **2023**; arXiv:2304.00690.
14. Wei, Y.; Zhao, L.; Zheng, W.; Zhu, Z.; Zhou, J.; Lu, J. SurroundOcc: Multi-Camera 3D Occupancy Prediction for Autonomous Driving. In *Proc. IEEE/CVF ICCV*, **2023**; arXiv:2303.09551.
15. Zhang, Y.; Zhu, Z.; Du, D. OccFormer: Dual-Path Transformer for Vision-Based 3D Semantic Occupancy Prediction. *arXiv preprint* arXiv:2304.05316, **2023**.
16. Tian, X.; Jiang, T.; Yun, L.; Mao, Y.; Yang, H.; Wang, Y.; Wang, Y.; Zhao, H. Occ3D: A Large-Scale 3D Occupancy Prediction Benchmark for Autonomous Driving. In *Proc. NeurIPS*, **2023**; arXiv:2304.14365.
17. Hao, X.; et al. MSC-Bench: A Multi-Sensor Corruption Benchmark for Robust Autonomous Driving. *arXiv preprint* arXiv:2501.01037, **2025**.
18. Lee, J. LiDAR Point-Wise Segmentation-Based De-Noising for Autonomous Driving Perception System in Fog Weather. *IEEE Access* **2026**, *14*. doi:10.1109/ACCESS.2026.3674455.
19. Collis, R.T.H.; Russell, P.B. Lidar Measurement of Particles and Gases by Elastic Backscattering and Differential Absorption. In *Laser Monitoring of the Atmosphere*; Hinkley, E.D., Ed.; Springer: Berlin, Heidelberg, **1976**; pp. 71–151. doi:10.1007/3-540-07743-X_18.
20. Koschmieder, H. Theorie der horizontalen Sichtweite. *Beiträge zur Physik der freien Atmosphäre* **1924**, *12*, 33–53 and 171–181.
21. Middleton, W.E.K. *Vision Through the Atmosphere*; University of Toronto Press: Toronto, **1952**. doi:10.3138/9781487586140.
22. Horvath, H. On the applicability of the Koschmieder visibility formula. *Atmos. Environ.* **1971**, *5*, 177–184. doi:10.1016/0004-6981(71)90081-3.
23. Lee, J.; Oh, K.; Yoon, Y.; Song, T.; Lee, T.; Yi, K. Adaptive Fault Detection and Emergency Control of Autonomous Vehicles for Fail-Safe Systems Using a Sliding Mode Approach. *IEEE Access* **2022**, *10*, 27863–27880. doi:10.1109/ACCESS.2022.3155738.
24. Lee, J.; Oh, K.; Oh, S.; Yoon, Y.; Kim, S.; Song, T.; Yi, K. Emergency Pull-Over Algorithm for Level 4 Autonomous Vehicles Based on Model-Free Adaptive Feedback Control With Sensitivity and Learning Approaches. *IEEE Access* **2022**, *10*, 27014–27034. doi:10.1109/ACCESS.2022.3156275.
25. Korea Meteorological Administration. Dense fog advisory: driving visibility 20–35 m, Seoul metropolitan area, 2021-11-20. Official announcement, KMA, 2021.
26. Lee, J. Fog Detection Examples. *YouTube*. https://youtu.be/JYp9aj88Yxg.
27. Lee, J. De-Noising Performance with Inference Time. *YouTube*. https://youtube.com/shorts/mntK6S2jTQw.
28. Lee, J. Fog Scenes, Pointcloud Shapes. *YouTube*, Aug. **2025**. https://youtu.be/g3VE6G_Ceug.
29. Lee, J. LiDAR Point Data in Heavy Fog: Density Changes due to LiDAR Physical Property. *YouTube*, Aug. **2025**. https://youtu.be/eWI3DklK6zo.
30. Heinzler, R.; Piewak, F.; Schindler, P.; Stork, W. CNN-based LiDAR Point Cloud De-Noising in Adverse Weather. *IEEE Robotics and Automation Letters* **2020**, *5*, 2514–2521.
31. Lee, J. Grid Map Videos. *YouTube*. https://youtu.be/JYp9aj88Yxg.
32. Lee, J. Fog De-Noising Results. *YouTube*, Aug. **2025**. https://youtu.be/A8b9FoDMkMY.
33. Lee, J. Un/Mis Recognition Cases in Fog Model (PointPillars). *YouTube*, Aug. **2025**. https://youtu.be/aVR0DtuOfRU.
34. Klett, J.D. Stable analytical inversion solution for processing lidar returns. *Appl. Opt.* **1981**, *20*, 211–220. doi:10.1364/AO.20.000211.