FogTwin: Physically-Grounded Synthetic Adverse-Weather LiDAR in NVIDIA Isaac Sim for Fog-Robust Perception

Jongmin Lee[1,2]
[1]Department of Mechanical Engineering, Seoul National University, Seoul 08826, South Korea
[2]Hyundai Mobis Co., Ltd., South Korea
E-mail: jesus911021@snu.ac.kr

***Abstract*—LiDAR perception degrades sharply in fog, yet labeled foggy data are scarce and, when fog is added in simulation, its physical realism is rarely verified. We present FogTwin, a pipeline that makes synthetic adverse-weather LiDAR both *physically faithful* and *demonstrably useful*. Rather than learning or hand-tuning a fog appearance, FogTwin injects a closed-form, peer-reviewed model of how fog attenuates and scatters light—the Beer–Lambert extinction law, the Koschmieder visibility relation, and near-range backscatter—directly into the ray-traced LiDAR of a modern robotics simulator (NVIDIA Isaac Sim). Combined with Omniverse Replicator, it automatically generates foggy point clouds paired with clean ground-truth labels and tags every frame with a single, physically meaningful measure of fog severity, turning fog density into the continuous, interpretable annotation that most synthetic datasets lack. We argue that trustworthy synthetic weather must be validated on two fronts, and we evaluate both: that the simulator obeys its own physics—simulated returns fade with distance exactly as the attenuation law predicts—and that the synthetic fog resembles real fog, through a reproducible protocol that compares simulated and real foggy LiDAR with distribution-distance statistics on a public adverse-weather benchmark, using no proprietary data. Finally, we show the data are not merely realistic but useful: a perception model trained with our physically grounded fog augmentation stays reliable in dense fog where a fog-naive baseline collapses. We further extend the static factory to dynamic driving sequences with a per-target observability ground truth unavailable from real recordings, and find that a single fog-severity scalar explains only a small fraction of per-target detection—fog-time visibility is class- and range-dependent. We report failure modes honestly—including the regime where the sensor itself, not the model, runs out of usable signal—and all released artifacts are synthetic and reproducible.**

***Index Terms*—synthetic data, LiDAR, fog simulation, adverse weather, sim-to-real, domain randomization, semantic segmentation, NVIDIA Isaac Sim, OpenUSD, robust perception.**

I. Introduction

Light detection and ranging (LiDAR) underpins modern autonomous-driving and mobile-robot perception, but its performance collapses in fog: water droplets attenuate the emitted pulse, randomly drop distant returns, and inject spurious near-range points through backscatter. Collecting and annotating real foggy LiDAR at scale is expensive, weather-dependent, and dangerous, so practitioners increasingly turn to simulation. Two questions then become unavoidable: (i) *is the synthetic fog physically faithful?* and (ii) *does the synthetic data actually improve downstream perception?* Synthetic datasets are easy to generate; datasets whose adverse-weather degradation is *measured against physics* and whose *utility is quantified* are rare.

This work answers both questions within the NVIDIA Omniverse / Isaac Sim ecosystem. Our central idea is deliberately conservative: rather than learn a fog appearance model, we *inject a closed-form, peer-reviewed sensor-physics model* into the simulator's ray-traced LiDAR. The physics (Beer–Lambert extinction, Koschmieder visibility, backscatter) is invariant; only the software interface to the simulator changes between versions. On top of this we build a labeled data factory and a two-stage evaluation—physical self-consistency and downstream robustness.

Contributions.

- A reusable fog-injection module for Isaac Sim RTX LiDAR that applies round-trip Beer–Lambert attenuation, stochastic range-dependent dropout, and near-range backscatter, parameterized by a single physically meaningful knob (target MDR, equivalently $\beta = 3.912/\mathrm{MDR}$).
- A 1,000-frame fog-severity-swept synthetic dataset generated with Omniverse Replicator, where each frame is annotated with semantic segmentation, tight 2-D boxes, depth, *and* quantitative β/MDR labels—a continuous, physically interpretable measure of fog severity that most synthetic sets omit.
- A sim-to-real fidelity *protocol* (distance-binned survival curves; KS and Earth-Mover's-Distance on β) demonstrated on the public STF benchmark, together with a direct check that simulated survival tracks the theoretical $e^{-2\beta r}$ curve.
- A controlled augmentation ablation showing a $\Delta{+}0.54$ foreground mIoU gain in dense fog, with an explicit, honest account of limitations.
- A dynamic extension that drives the ego through fog while logging a *per-target observability ground truth*—a label real recordings cannot provide—and a quantitative

Corresponding author: Jongmin Lee (e-mail: jesus911021@snu.ac.kr).

finding over 1.98×10^4 target-frames from 18 driving scenarios: a single fog scalar (MDR) explains only R^2=0.11 of per-target detection, rising to 0.36 with range and object-class cues, establishing that fog-time observability is class- and geometry-dependent rather than a single-number property.

We emphasize honest reporting throughout: we identify a small systematic offset between simulated and target fog density, we mark out-of-distribution test regimes, and we acknowledge that below a visibility floor neither model can recover information that the sensor never received.

II. Related Work

Physics-based fog simulation for LiDAR. Hahner *et al.* [1] apply a Beer–Lambert/Koschmieder fog model directly to real LiDAR point clouds, adding hard scattering and backscatter to improve 3-D detection in adverse weather; their follow-up extends the idea to snowfall [2]. Our fog model adopts the same physical core but applies it inside a *ray-traced simulator* (Isaac Sim RTX LiDAR), so that fog can be swept jointly with full scene and label randomization rather than post-hoc on a fixed real scan.

Image-domain fog and synthetic adverse weather. For cameras, Sakaridis *et al.* [3] synthesize Foggy Cityscapes via the optical (Koschmieder) scattering model and show segmentation gains. We reuse the same camera scattering model for our 2-D robustness study, but pair it with a LiDAR-consistent extinction parameter so that the camera and range-sensor degradations share one β.

Real adverse-weather benchmarks. The *Seeing Through Fog* (STF) dataset [4] provides real multi-sensor recordings under fog, rain and snow and documents the characteristic range collapse of LiDAR in fog. We use STF as the *public* reference distribution for our sim-to-real fidelity protocol; no proprietary recordings are used in this paper.

Simulation and synthetic-data infrastructure. NVIDIA Isaac Sim [5] provides physically based rendering and ray-traced sensors on top of OpenUSD [6]; Omniverse Replicator [7] automates randomized scene generation with ground-truth annotators. World-foundation models such as Cosmos [8] can further photorealize simulated frames; we leave this as an augmentation extension. For perception we use the standard U-Net [9] architecture and KITTI-style label conventions [10].

Real foggy-LiDAR de-noising and detection range. Our fog-severity parameterization builds directly on prior work that analyzes real foggy LiDAR and performs point-wise segmentation-based de-noising for autonomous-driving perception [11]; from it we adopt the *maximum detection range* (MDR) as the operational, sensor-referred measure of fog severity used throughout this paper (Section III). Our distinct emphasis relative to prior fog-simulation work is on *quantified* sim-to-real fog fidelity and on per-frame, physically interpretable fog-severity metadata.

III. Fog Sensor-Physics Model

A. Extinction and visibility

A collimated beam traversing a scattering–absorbing medium loses power exponentially with path length r. The single-pass transmittance is the Beer–Lambert law

$$T(r) = e^{-\beta r}, \tag{1}$$

where β (units m^{-1}) is the *extinction coefficient* (scattering plus absorption) and $\tau = \beta r$ is the optical depth. The Koschmieder relation [12] converts β into a human-interpretable meteorological visibility

$$V = \frac{-\ln(C_{\mathrm{th}})}{\beta} = \frac{3.912}{\beta}, \qquad C_{\mathrm{th}} = 0.02, \tag{2}$$

the 2% contrast threshold giving $-\ln(0.02) = \ln(50) = 3.912$. As our operational, sensor-referred measure of fog severity we adopt the *maximum detection range* (MDR): the greatest distance at which the LiDAR still returns reliable points through fog. We follow the MDR characterization introduced in prior work on point-wise LiDAR de-noising for driving perception in fog [11]. Whereas the extinction coefficient β is a property of the medium, MDR is the range-domain quantity an autonomy stack actually cares about—"how far can the sensor still see?"—and the two are dual through Eq. (2), so that a target MDR fixes the single severity knob,

$$\beta = \frac{3.912}{\mathrm{MDR}}. \tag{3}$$

A smaller MDR denotes denser fog: e.g. $\mathrm{MDR} = 40\,\mathrm{m}$ corresponds to $\beta \approx 0.10\,\mathrm{m}^{-1}$ and $\mathrm{MDR} = 20\,\mathrm{m}$ to $\beta \approx 0.20\,\mathrm{m}^{-1}$. Reporting fog severity in MDR rather than β makes every figure and table in this paper directly interpretable as a viewing distance.

B. Round-trip attenuation for LiDAR

Unlike a camera (single pass), a LiDAR pulse travels to the target *and* back, traversing the fog twice. The returned intensity is therefore

$$I(r) = I_0\, e^{-2\beta r}, \tag{4}$$

the factor of two being the dominant practical subtlety (omitting it halves the modeled attenuation). As a worked example, for $\mathrm{MDR} = 40\,\mathrm{m}$ ($\beta = 0.0978$), the single-pass transmittance at $40\,\mathrm{m}$ is exactly the 2% visibility threshold, whereas the round-trip return is $e^{-7.824} \approx 4 \times 10^{-4}$—i.e. such points effectively vanish from the cloud, reproducing the empirical loss of distant returns in fog.

C. Three modeled effects

Given a clean simulated scan $\{(\mathbf{p}_i, I_i)\}$ with ranges $r_i = \|\mathbf{p}_i\|$, the fog operator applies:

1) **Intensity attenuation & hard drop**: scale each return by $e^{-2\beta r_i}$ (Eq. 4) and remove returns falling below an intensity floor $I_{\min}$.

2) **Stochastic range dropout**: keep each point with probability $p_{\text{keep}}(r_i) = e^{-\beta r_i}$, modeling the increasing chance that a distant pulse fails to produce a valid return.
3) **Near-range backscatter**: insert spurious points at short range ($r \in [0.5, 3]\,\text{m}$), emulating droplet backscatter, following the treatment of [1].

For the camera modality (Section VI) the same extinction drives the optical (Koschmieder) scattering model

$$\mathbf{J}_{\text{fog}} = \mathbf{J}\, e^{-\beta d} + \mathbf{A}\left(1 - e^{-\beta d}\right), \tag{5}$$

with scene radiance $\mathbf{J}$, per-pixel depth d, and atmospheric light $\mathbf{A}$ [3]. Equations (4) and (5) are two views of one physics, so a single β simultaneously parameterizes range-sensor and camera degradation.

IV. Synthetic Data Generation

A. Ray-traced LiDAR in Isaac Sim

We instantiate a standalone rotary RTX LiDAR in Isaac Sim and read its point-cloud annotator each frame; intensity is taken as unit-valued where the annotator does not populate it. A practical lesson worth recording is that the sensor API differs across simulator releases, so we fix the exact annotator/return contract by first running the vendor's bundled example and then reusing that verified interface throughout—treating the simulator as a moving target while the physics of Eq. (4) stays fixed.

To obtain a dense, full-surround reference scan we elevate the sensor so that the ground ring is returned, accumulate returns over a full rotation, and voxel-deduplicate at $0.15\,\text{m}$. This yields a $\sim$37 k-point raw surround scan ($\approx$14 k after deduplication, matching a realistic single-sweep density) containing both a ground ring and vehicle/pedestrian objects out to $\sim$35 m. Applying the fog operator at MDR $= 40\,\text{m}$ ($\beta = 0.098$) retains roughly 7% of points, and the measured fraction of surviving points per range bin closely follows the theoretical $e^{-2\beta r}$ of Eq. (4) (Fig. 1). This agreement is our first, purely *physical*, self-consistency check.

B. Severity sweep and automatic labeling

We extend the scene with semantically tagged objects (vehicle and pedestrian proxies) and drive Omniverse Replicator to randomize object pose and lighting. For every frame, the camera annotators export RGB, semantic segmentation, tight 2-D bounding boxes and depth from the *clean* scene, while the LiDAR cloud receives a fog realization with β sampled from a fog-severity prior. We sample $\beta \sim \mathcal{U}(0.08, 0.26)\,\text{m}^{-1}$ (mean 0.172, p_{10}=0.10, p_{90}=0.24), corresponding to MDR from $\sim$49 down to $\sim$15 m (median MDR $\approx$22 m)—the thick-to-moderate fog regime per Eq. (2). Each frame stores (β, MDR) as a quantitative severity label (Fig. 2, Fig. 3). The result is a 1,000-frame labeled dataset (Table I).

C. The labeling invariant

A single rule governs annotation: *degrade the sensor input, keep the ground truth clean*. Fog is applied only to the sensor stream (the LiDAR cloud here, and the RGB image in Section VI); the semantic, box and depth labels are always taken from the fog-free scene. This prevents a model from "learning the fog" as part of the target and keeps fog severity an *input* nuisance variable rather than a label-corrupting one.

TABLE I
FogTwin synthetic dataset summary.

Property	Value
Frames	1,000
LiDAR	Isaac Sim RTX rotary, full surround
Camera labels (clean GT)	RGB, semantic, 2-D box, depth
LiDAR product	foggy point cloud ($\sim$1.1 k pts/frame)
Severity prior	$\beta \sim \mathcal{U}(0.08, 0.26)\,\text{m}^{-1}$
β statistics	mean 0.172, p_{10} 0.10, p_{90} 0.24
MDR range (median)	$\sim$15–49 m ($\approx$22 m)
Per-frame metadata	(β, MDR)

V. Sim-to-Real Fidelity Protocol

Quantifying *how close* synthetic fog is to real fog is what separates a plausible-looking dataset from a trustworthy one. We define a fidelity protocol with two complementary parts and evaluate it against the public STF benchmark [4]; we deliberately use no proprietary recordings in this public release.

A. Physical self-consistency

First, independent of any real dataset, the simulator must obey its own physics. We bin returns by range and compare the measured survival fraction $N_{\text{fog}}(r)/N_{\text{clean}}(r)$ against the theoretical round-trip transmission $e^{-2\beta r}$ (Eq. 4). As shown in Fig. 1, the two agree closely across the usable range, confirming that the implemented operator reproduces the intended extinction physics rather than an arbitrary thinning.

B. Distributional sim-to-real comparison

Second, we compare distributions of fog severity and range degradation between simulation and real foggy LiDAR. Let $\{\beta_i^{\text{sim}}\}$ be the per-frame severities of our sweep and $\{\beta_j^{\text{real}}\}$ the severities estimated from real foggy scans (e.g. STF). We summarize agreement with two scale-appropriate statistics:

$$\text{KS} = \sup_x \left| F_{\text{sim}}(x) - F_{\text{real}}(x) \right|, \tag{6}$$

$$\text{EMD} = W_1\big(P_{\text{sim}},\, P_{\text{real}}\big) = \int_0^1 \left| F_{\text{sim}}(x) - F_{\text{real}}(x) \right| dx, \tag{7}$$

the Kolmogorov–Smirnov statistic and the first Wasserstein (Earth-Mover's) distance. The same range-binned survival curves are compared between simulated and real foggy sweeps. We also verify that the well-documented STF phenomenon of *LiDAR range collapse in fog* is reproduced: increasing β in simulation compresses the effective maximum range and shifts mass toward short range, mirroring the qualitative behavior reported for real fog in [4].

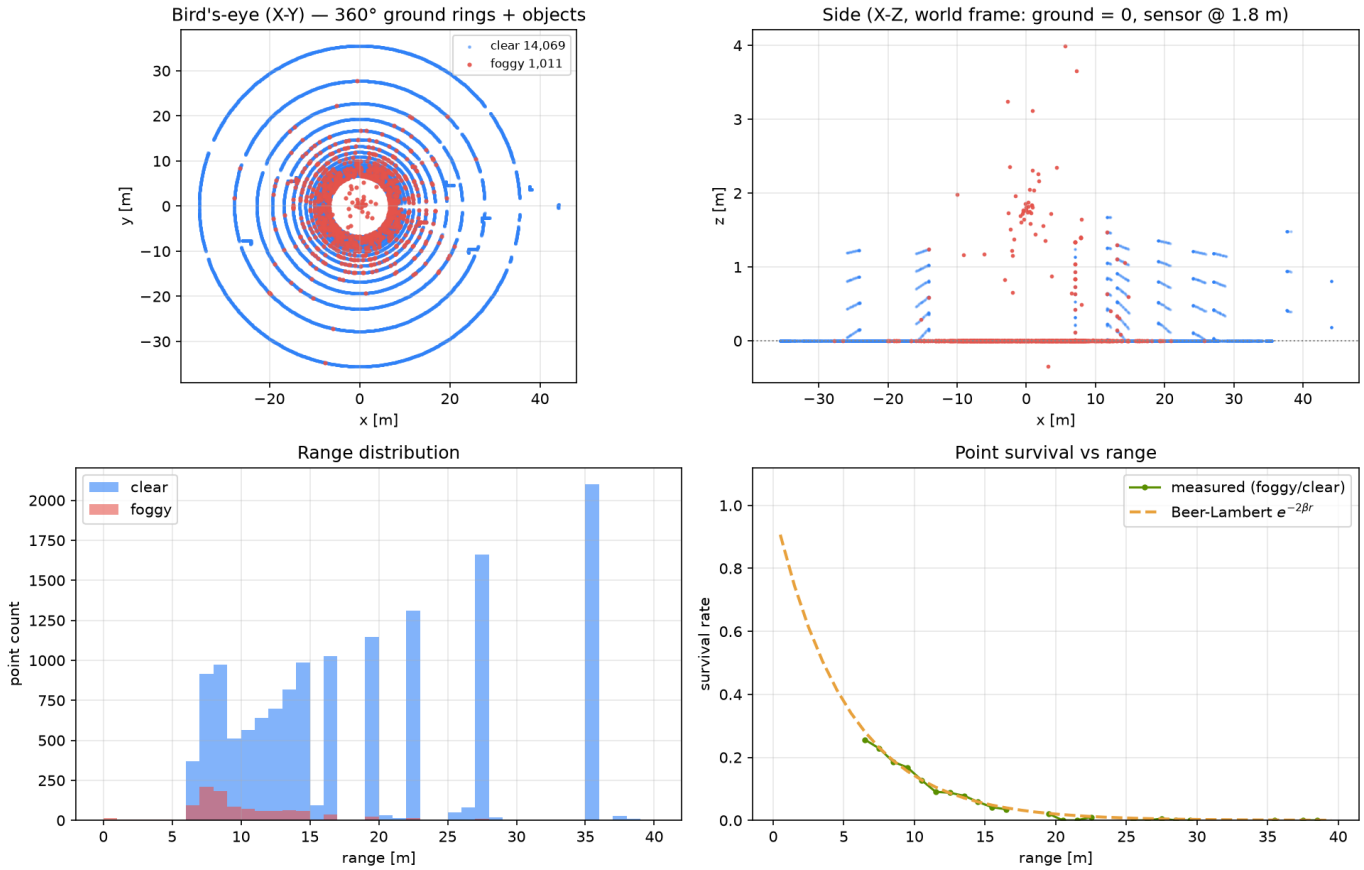

Fig. 1. Full-surround Isaac Sim RTX LiDAR scan. Blue: clean returns (ground ring plus vehicle/pedestrian objects, 8–45 m). Red: after round-trip Beer–Lambert attenuation plus backscatter at MDR = 40 m ($\beta = 0.098$). Two-way attenuation reduces the cloud to ≈7% of returns. Inset: measured per-bin survival tracks the theoretical $e^{-2\beta r}$ curve—physical self-consistency. (Synthetic data; public.)

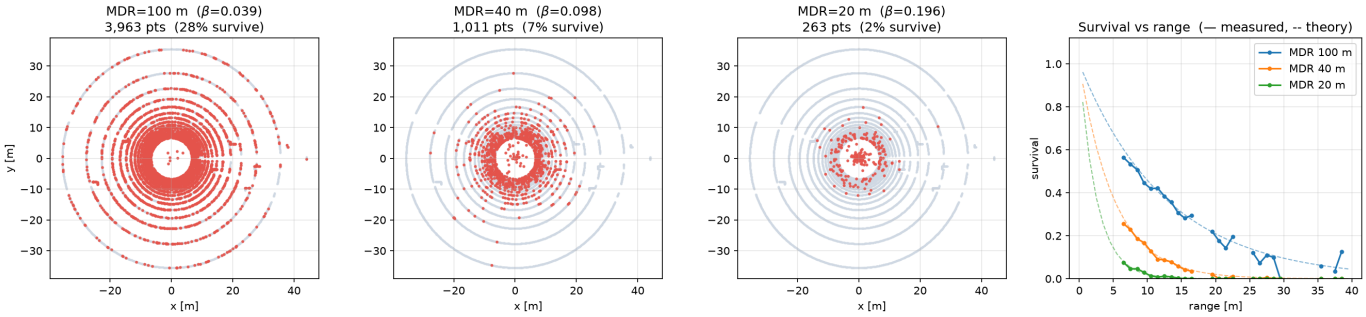

Fig. 2. Fog-severity sweep on one clean scan: as MDR decreases ($100 \to 40 \to 20$ m) the surviving (red) returns fall sharply (≈28% → 7% → 2%). The strong range sensitivity of round-trip attenuation defines the severity range spanned by the dataset. (Synthetic; public.)

C. *Interpretation and honesty*

EMD is the more forgiving, distribution-shape metric and is our headline fidelity number; KS is deliberately sensitive and will flag even a small systematic offset between two large samples. We therefore report both and, crucially, do *not* claim distributional identity: a sensitive test detecting a small, named offset (Section VIII) is more trustworthy than a single optimistic statistic. The point of this section is the *protocol*—a reproducible, physics-anchored way to ask "how real is my synthetic fog?"—not a leaderboard number.

VI. Fog-Robustness Ablation

A. *Setup*

We test whether the physically grounded fog *helps* a downstream model. On the 2-D camera modality we train two otherwise-identical U-Nets [9] (classes: background, vehicle, pedestrian; class weights $0.5/5/2$ to recover the small vehicle class; horizontal-flip augmentation): a **baseline** trained

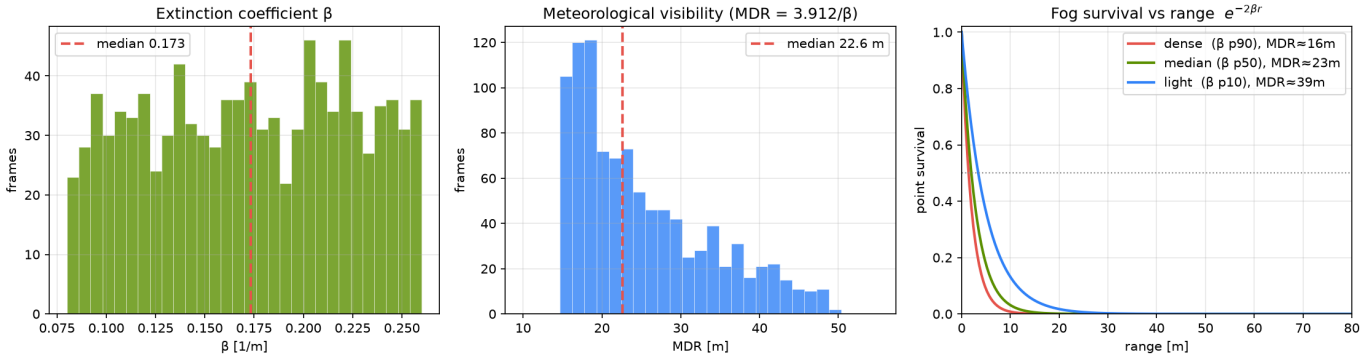

Fig. 3. Characterization of the 1,000-frame sweep. Left: extinction coefficient β distribution. Center: meteorological visibility MDR $= 3.912/\beta$. Right: per-severity range–survival curves. (Synthetic; public.)

TABLE II
FOREGROUND MIOU VS. TEST-TIME FOG SEVERITY. BASELINE (CLEAN-ONLY) VS. ROBUST (CLEAN + PHYSICALLY GROUNDED FOG AUGMENTATION).

Fog $\beta\,[\mathrm{m}^{-1}]$	0.10	0.15	0.20	0.30	0.40	0.50
MDR [m]	39	26	20	13	10	8
Baseline mIoU	0.780	0.866	0.741	0.070	0.000	0.000
Robust mIoU	**0.963**	**0.962**	**0.951**	**0.612**	0.080	0.000
Δ mIoU	+0.184	+0.096	+0.210	**+0.542**	+0.080	0.000

on clean RGB only, and a **robust** model trained on clean RGB *plus* fog-augmented RGB obtained from the depth-based Koschmieder model of Eq. (5), with labels kept clean per the labeling invariant. The only difference between the two runs is the presence of fog augmentation. From 1,001 frames we use a leakage-free split of ∼700 train / ∼150 validation / ∼150 test. Because the foreground classes are small (test ground-truth pixels: vehicle 2.3×10^4, pedestrian 6.3×10^5, background 9.2×10^6), we report *foreground* mIoU (mean of vehicle and pedestrian).

B. Robustness curve

Rather than a single operating point, we sweep test-time fog $\beta \in [0.10, 0.50]$ and measure mIoU at each severity, exposing the full visibility–robustness trade-off (Table II, Fig. 4). In light fog both models perform well; as fog thickens the clean-only baseline collapses past a threshold near $\beta \approx 0.30$ (MDR $\approx 13\,\mathrm{m}$), dropping to mIoU 0.07, whereas the fog-augmented model degrades gracefully to 0.61—a Δ+0.54 (≈8.7×) gain at the decisive severity. The area between the two curves is, in effect, the robustness purchased by the physically grounded augmentation.

C. Why the collapse is physical

The baseline cliff is explained by the physics of Section III: by Eq. (5) the transmittance $e^{-\beta d}$ decays exponentially with depth, so beyond a β-dependent range the scene is dominated by atmospheric light and informative contrast vanishes. The mIoU therefore does not fade linearly but falls off a cliff once the effective visibility crosses the scene's characteristic depth—the same extinction that thins the LiDAR cloud also blinds the camera. Perception degradation is thus *predicted by visibility*, tying the downstream result back to the sensor physics that generated the data.

VII. Dynamic Extension: Temporal Sequences and Per-Class Observability

Sections IV–VI treat fog on *static* scans. We now extend FogTwin to *dynamic* driving sequences, which expose a question static frames cannot ask: *as the ego closes on a target through fog, when does that target first become observable, and how does this depend on what the target is?* Crucially, the simulator can answer this with a label that real recordings cannot provide—a *per-target observability ground truth.*

A. Dynamic scenario engine

An ego vehicle follows a smooth trajectory (forward speed 10–18 m/s, mild curvature) while 4–8 moving actors of four classes—truck, car, motorcycle, pedestrian—advance along lane-offset paths with class-appropriate speeds. Each actor is a class-sized geometric proxy whose extent sets its LiDAR point yield (Stage A; photorealistic USD assets are a drop-in replacement, Section VIII). Every frame we (i) update ego and actor poses, (ii) accumulate a full 360° surround scan as in Section IV, and (iii) apply the fog operator of Eq. (4) with a per-scenario severity that ramps from light to dense fog (MDR $\approx 50 \to 16\,\mathrm{m}$), so each sequence captures a full visibility transition. We generate 18 scenarios $\times$ 200 frames (3,600 frames; 1.98×10^4 target-frames with non-empty clean return). Figure 5 shows one sequence from above as fog thickens.

B. Per-target observability ground truth

For target k at frame t we record, from its *known* bounding box: the clean point yield $y_{k,t}^{\mathrm{clean}}$ (returns before fog), the foggy yield $y_{k,t}^{\mathrm{fog}}$ (returns surviving fog), a visibility flag $v_{k,t} = \mathbf{1}\left[y_{k,t}^{\mathrm{fog}} \geq n_{\min}\right]$ ($n_{\min}$=3), and the per-sequence *first-detection range*—the farthest distance at which the target first yields $\geq n_{\min}$ surviving points. This label exists only in simulation: a real foggy recording cannot distinguish "a

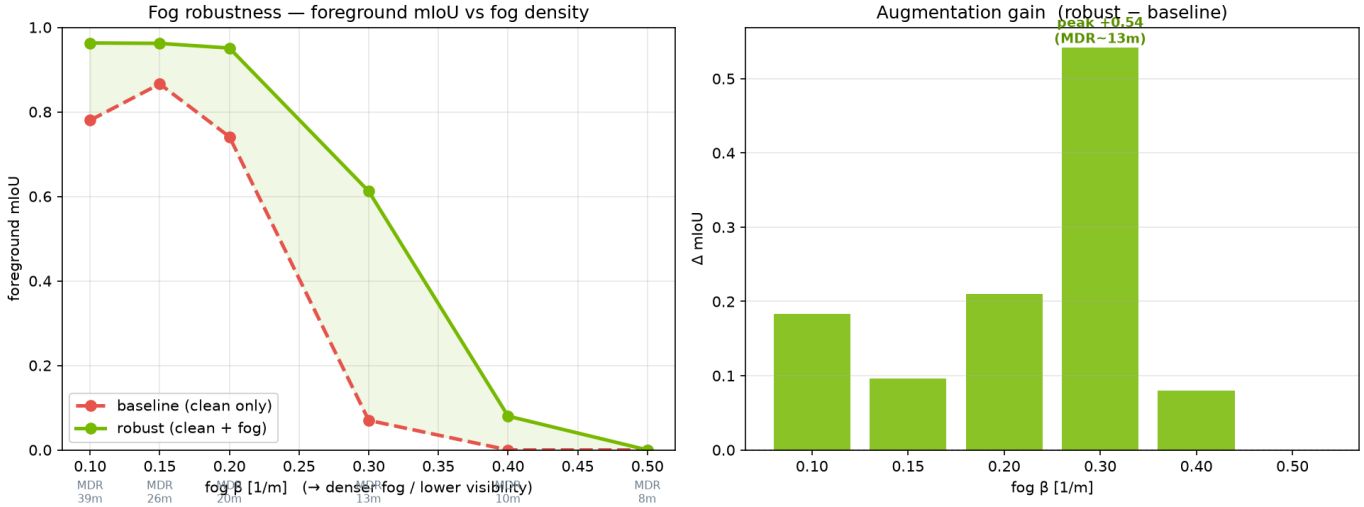

Fig. 4. Fog-robustness curves. As fog thickens the clean-only baseline (red) collapses near $\beta \approx 0.30$ while the fog-augmented model (green) degrades gracefully; the curves diverge at the threshold severity. Right: the Δ mIoU gain peaks at $+0.54$ at MDR $\approx 13\,\mathrm{m}$. (Synthetic evaluation; public.)

TABLE III
PER-CLASS FOG OBSERVABILITY POOLED OVER 18 DYNAMIC SCENARIOS (108 TARGETS). "DETECT RATE" IS THE FRACTION OF TARGETS DETECTED AT ALL IN FOG; RANGES ARE MEAN FIRST-DETECTION DISTANCES.

Class	n	Clear [m]	Fog [m]	Detect rate
Truck	26	140.0	**28.3**	**73%**
Car	38	97.0	22.2	71%
Motorcycle	22	115.8	17.1	50%
Pedestrian	22	69.1	18.4	**18%**

target is present but its returns were lost to fog" from "no target is there," because ground-truth presence is unknown. The simulator knows the box, so it labels the fog-induced *false negative* exactly. (We label only fog-induced misses; real-world false positives from clutter are out of scope.)

C. Per-class observability collapses unevenly

Table III pools first detections over all 18 scenarios (108 targets). Two trends stand out. First, the *rate* at which a target is detected at all through fog falls monotonically with object size: a truck is seen in fog 73% of the time, a car 71%, a motorcycle 50%, and a pedestrian only 18%—pedestrians are lost outright in roughly four sequences out of five. Second, the fog *range* when a target is detected does *not* follow the same order: motorcycles, though detected more often than pedestrians, are detected at shorter range (17.1 vs. 18.4 m). The two observability axes—how *often* and how *far*—rank the small classes differently, which is precisely why a single severity number cannot summarize fog-time observability.

D. A single severity scalar is insufficient

We test the summarizing power of MDR directly. Treating each target-frame (1.98×10^4 samples) as one observation, we fit a linear model of the fog-time detection indicator

TABLE IV
VARIANCE OF PER-TARGET FOG DETECTION EXPLAINED (R^2) BY PROGRESSIVELY RICHER CUE SETS, OVER 1.98×10^4 TARGET-FRAMES. A SINGLE FOG SCALAR IS A WEAK PREDICTOR.

Predictor set	Detection R^2
MDR / β only (single scalar)	0.11
+ geometry (range, aspect)	0.19
+ class / size (area, density)	**0.36**

$v_{k,t}$ and report the variance explained (R^2). Fog severity alone—the single scalar β (equivalently MDR)—explains only R^2=0.11 of per-target detection. Adding viewing geometry (range, aspect) raises this to 0.19, and adding object class/size cues (frontal area, clean point density) to 0.36: a 3.3× increase (Table IV, Fig. 6). The same pattern holds for point survival: a β-only model reaches R^2=0.06, rising to 0.21 once range enters, consistent with the $e^{-2\beta r}$ law of Eq. (4) in which range and extinction are inseparable. These are deliberately conservative linear estimates; the qualitative message is robust to the model form: *fog-time observability is a property of the (target, range, fog) triple, not of fog severity alone*—the dynamic, sim-only counterpart to the static physical checks of Sections IV–V.

VIII. Limitations and Honest Reporting

We state limitations explicitly; in safety-relevant perception, naming failure modes is part of the contribution.

- **Small density offset.** The sampled simulation severity is slightly milder than typical dense-fog references (sweep mean $\beta \approx 0.17$ vs. a dense-fog target nearer 0.19). A sensitive KS test detects this offset even when EMD is small; re-centering the β prior is a direct fix.
- **Out-of-distribution test severities.** Training fog spans $\beta \in [0.08, 0.26]$, but we test up to 0.50. Results for $\beta \geq$

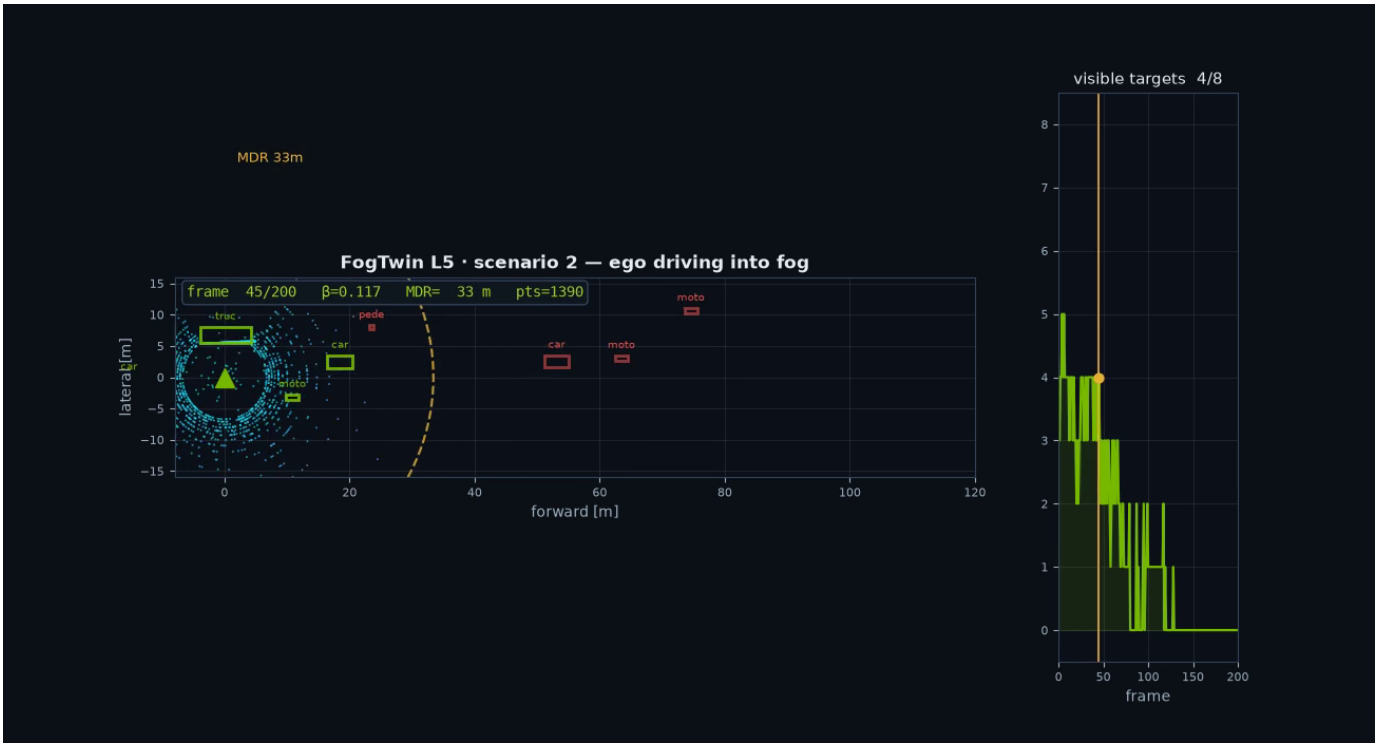

Fig. 5. One frame of a dynamic driving sequence in bird's-eye view: the ego (green triangle) drives through fog (here MDR $\approx 33\,\mathrm{m}$, 8 targets, 4 still observable). Dashed yellow circle: the MDR visibility horizon; gray points are surviving foggy returns; boxes are class-labeled targets, green when observable ($\geq n_{\min}$ returns) and red once fog suppresses them—small targets enter the red (lost) state first. Right panel: number of observable targets across the sequence. (Synthetic; public.)

0.30 therefore probe mild extrapolation; the fairest in-distribution comparison is $\beta \leq 0.26$ (robust ≈0.95 vs. baseline 0.74–0.87), while the $+0.54$ at $\beta = 0.30$ is best described as slightly-OOD generalization.

- **Geometric proxies.** Vehicles and pedestrians are simple geometric proxies that validate the pipeline; photorealistic USD assets (and optionally Cosmos [8] photorealization) would raise visual realism. The vehicle class is small in pixels and is recovered via class weighting; we recommend per-class reporting.
- **Observability-label scope.** The per-target observability ground truth (Section VII) records only fog-induced *false negatives* on known objects; real-world false positives from clutter are not modeled. The reported R^2 values are conservative linear estimates, and the dynamic study uses the same geometric proxies, so absolute detection ranges will shift with photorealistic assets even though the class ordering is expected to persist.
- **Sensor information floor.** Below MDR $< 10\,\mathrm{m}$ both models reach ≈0 mIoU. This is not a model failure but a *sensor* limit: almost no usable signal returns through that much fog. We do not claim recovery where the physics forbids it.
- **Scope of the public release.** This paper validates fidelity against the public STF benchmark and releases only synthetic data and methodology; any in-house real-sensor validation is out of scope here.

IX. Conclusion

FogTwin shows that a small amount of *validated physics*, injected into a modern ray-traced simulator, turns generic synthetic LiDAR into a physically grounded, quantitatively labeled adverse-weather dataset. Simulated point survival follows the theoretical $e^{-2\beta r}$ law; a reproducible fidelity protocol (survival curves plus KS/EMD on β) anchors sim-to-real comparison to public data; and a controlled ablation demonstrates a $\Delta{+}0.54$ foreground-mIoU robustness gain in dense fog. A dynamic extension further shows that fog-time per-target observability is class- and range-dependent—a single fog scalar explains only $R^2{=}0.11$ of detection versus 0.36 with geometry and class—yielding a sim-only observability label that real recordings cannot supply. By coupling per-frame β/MDR metadata with an explicit account of failure modes, FogTwin moves synthetic adverse-weather data from "looks foggy" toward "physically faithful and demonstrably useful." Future work includes re-centering the severity prior to tighten the KS test, replacing geometric proxies with photorealistic USD assets, extending the operator to rain and snow [2], and folding world-foundation-model photorealization into the augmentation loop.

Data, Code, and Video Availability

All released artifacts are fully synthetic; no proprietary or real-world recordings are included. The dynamic dataset—18 driving scenarios (3,600 frames, 1.98×10^4 target-

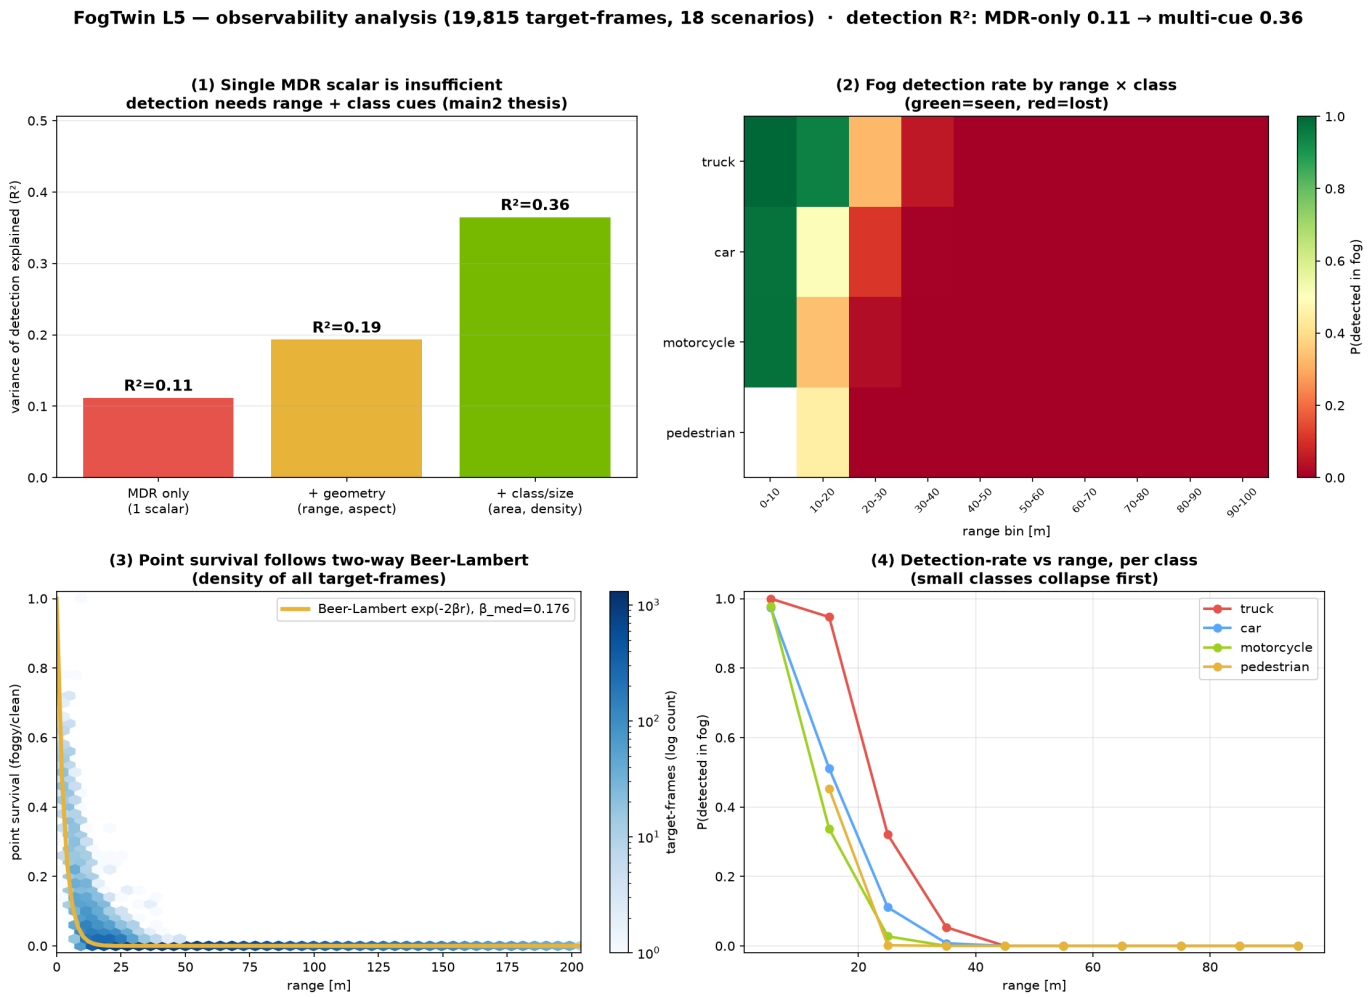

Fig. 6. Dynamic-sequence observability analysis (1.98×10^4 target-frames, 18 scenarios). (1) Variance of fog detection explained: a single MDR scalar reaches only R^2=0.11, rising to 0.36 with range and class cues. (2) Fog detection rate over range × class (green seen, red lost). (3) Per-target point survival follows the two-way Beer–Lambert curve $e^{-2\beta r}$. (4) Detection rate vs. range per class: small classes collapse first. (Synthetic; public.)

frames) with per-target observability ground truth—is released under the CC-BY-4.0 license at https://huggingface.co/datasets/jmlee8896/IssacSimFog. The fog-injection, dynamic sequence-generation, and analysis code is included in the same repository at https://huggingface.co/datasets/jmlee8896/IssacSimFog/tree/main/code. A bird's-eye-view video of all 18 fog scenarios—showing the shrinking MDR visibility horizon and the per-class observability collapse—is available at https://youtu.be/QryhQiNYSEA.

Reproducibility and Ethics

All released data are synthetic and contain no personally identifiable information. The fog operator is closed-form (Eqs. 1–5) and the generation pipeline relies on publicly available simulation tooling [5]–[7]. Sim-to-real fidelity is reported against the public STF benchmark [4]; no proprietary recordings are included.

References

[1] M. Hahner, C. Sakaridis, D. Dai, and L. Van Gool, "Fog simulation on real LiDAR point clouds for 3D object detection in adverse weather," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, 2021, pp. 15 283–15 292.

[2] M. Hahner, C. Sakaridis, M. Bijelic, F. Heide, F. Yu, D. Dai, and L. Van Gool, "LiDAR snowfall simulation for robust 3D object detection," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 2022, pp. 16 364–16 374.

[3] C. Sakaridis, D. Dai, and L. Van Gool, "Semantic foggy scene understanding with synthetic data," *Int. J. Comput. Vis. (IJCV)*, vol. 126, no. 9, pp. 973–992, 2018.

[4] M. Bijelic, T. Gruber, F. Mannan, F. Kraus, W. Ritter, K. Dietmayer, and F. Heide, "Seeing through fog without seeing fog: Deep multimodal sensor fusion in unseen adverse weather," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 2020, pp. 11 682–11 692.

[5] NVIDIA Corporation, "NVIDIA Isaac Sim: Robotics simulation and synthetic data generation," https://developer.nvidia.com/isaac/sim, 2024, accessed: 2026-06-22.

[6] Alliance for OpenUSD (AOUSD), "Universal scene description (OpenUSD)," https://openusd.org, 2024, accessed: 2026-06-22.

[7] NVIDIA Corporation, "Omniverse Replicator: A framework for synthetic data generation," https://developer.nvidia.com/omniverse/replicator, 2024, accessed: 2026-06-22.

[8] ——, "NVIDIA Cosmos: World foundation models for physical AI," https://www.nvidia.com/en-us/ai/cosmos/, 2025, accessed: 2026-06-22.

[9] O. Ronneberger, P. Fischer, and T. Brox, "U-Net: Convolutional networks for biomedical image segmentation," in *Proc. Med. Image Comput. Comput.-Assist. Interv. (MICCAI)*, 2015, pp. 234–241.

[10] A. Geiger, P. Lenz, and R. Urtasun, "Are we ready for autonomous driving? the KITTI vision benchmark suite," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 2012, pp. 3354–3361.

[11] J. Lee, "LiDAR point-wise segmentation-based de-noising for autonomous driving perception system in fog weather," *IEEE Access*, vol. 14, pp. 51 238–51 249, 2026.

[12] H. Koschmieder, "Theorie der horizontalen sichtweite," *Beiträge zur Physik der freien Atmosphäre*, vol. 12, pp. 33–53, 171–181, 1924.