The Fortieth AAAI Conference on Artificial Intelligence (AAAI-26)

Physically-Based LiDAR Smoke Simulation for Robust 3D Object Detection

Shijun Zheng[1,2], **Yu Guo**[1,2], **Weiquan Liu**[3]*,
Yu Zang[1,2], **Siqi Shen**[1,2], **Ming Cheng**[1,2], **Cheng Wang**[1,2]*
[1]Fujian Key Laboratory of Sensing and Computing for Smart Cities, School of Informatics, Xiamen University, China
[2]Key Laboratory of Multimedia Trusted Perception and Efficient Computing, Ministry of Education of China,
School of Informatics, Xiamen University, China
[3]College of Computer Engineering, Jimei University, China
{zhengshijun, gouyu}@stu.xmu.edu.cn, {siqishen, chm99, cwang}@xmu.edu.cn, wqliu@jmu.edu.cn, zangyu7@126.com

Abstract

3D object detection in adverse weather is crucial for autonomous driving, especially in smoke where LiDAR data becomes sparse and noisy. Due to the lack of real smoke data, this paper introduces a physics-based simulation framework to generate realistic LiDAR point clouds of smoke and augment large-scale driving datasets. First, we present a 3D fluid dynamics-based smoke simulation framework in Unity, which models the realistic spatial diffusion and temporal evolution of smoke particles. Coupled with a physically accurate LiDAR perception module, our system captures complex light interactions—such as beam attenuation, scattering, and multi-path effects—to generate high-fidelity, physically consistent smoke point clouds. Second, we propose a range image-based data fusion strategy that seamlessly integrates the simulated smoke point clouds into large-scale real-world LiDAR datasets (e.g., Waymo). This approach accurately emulates LiDAR scanning characteristics and naturally incorporates occlusion effects, enabling realistic smoke integration without compromising spatial consistency. To validate our approach, we collect a real-world LiDAR smoke dataset (LiSmoke) and conduct extensive experiments using state-of-the-art 3D detectors. Results demonstrate that models trained with our augmented synthetic data achieve significant improvements in smoke-affected scenarios, while maintaining competitive performance in clear-weather conditions. Our work provides a cost-effective solution for enhancing perception robustness in safety-critical environments.

Introduction

Autonomous driving technologies have advanced rapidly, driven largely by improvements in 3D object detection using LiDAR point clouds (Zhang et al. 2024; Caesar et al. 2020). However, state-of-the-art detectors often degrade significantly in adverse weather conditions such as fog, smoke, or heavy precipitation. Among these, smoke presents unique challenges due to its transient, unstructured nature and frequent appearance in safety-critical scenarios like wildfires, traffic accidents, and industrial zones. Unlike fog, which can be statistically modeled, smoke involves turbulent motion and diverse particle sizes, resulting in multi-scale scattering fundamentally different from fog or spray.

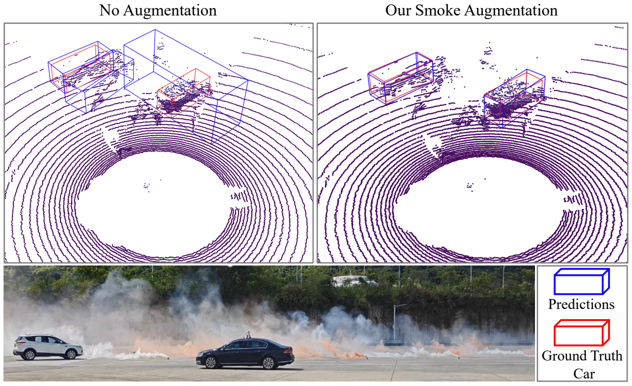

Figure 1: Comparison of 3D object detection performance on the real-world smoke dataset (LiSmoke). The top-left image shows the detection results without any data augmentation, while the top-right image shows the results after training with simulated smoke augmentation. The bottom image illustrates the real-world scene where the smoke point cloud was collected.

LiDAR is a key sensor for autonomous vehicles, providing precise 3D detection via direct depth measurement and high-resolution point clouds. Yet, LiDAR performance deteriorates under adverse weather due to scattering and absorption effects that distort measurements. While prior works have simulated fog, rain, and snow to improve robustness (Hahner et al. 2021, 2022; Shih et al. 2022), smoke-although equally pervasive-remains underexplored. Smoke induces severe signal attenuation, spurious noise, and transient occlusions that impair LiDAR detection (Zheng et al. 2025). Addressing this is critical for safety in scenarios like urban fires and industrial accidents.

Public datasets rarely contain smoke-affected LiDAR scenes due to the difficulty of capturing such data in real environments. Moreover, naive fusion of synthetic smoke with real point clouds often introduces spatial and intensity inconsistencies, causing domain shifts that reduce model generalization. These challenges limit the development and evaluation of robust perception models under smoke interference.

To fill this gap, we propose Physically-Based LiDAR

---

*Corresponding Author.

Smoke Simulation, a novel framework integrating dynamic smoke modeling, LiDAR signal physics, and cross-domain alignment. Our contributions include:

- `Dynamic Smoke Simulation`: A dynamic smoke particle system implemented in Unity3D that captures turbulent motion, particle size variability, and optical effects such as absorption, scattering, and multi-path reflections.
- `LiDAR Perception Modeling`: A LiDAR perception module modeling wavelength-dependent attenuation, beam divergence, and sensor noise, generating realistic smoke-affected point clouds.
- `Cross-Domain Fusion`: A cross-domain fusion strategy that spatially aligns simulated smoke with real datasets (e.g., Waymo (Sun et al. 2020)) while preserving physical plausibility and intensity consistency.
- `LiSmoke Dataset`: A real-world LiDAR smoke dataset that covers various smoke densities and dynamic scenarios, enabling comprehensive benchmarking.
- `Empirical Validation`: Extensive experiments demonstrating smoke's impact on detection and the effectiveness of smoke-augmented training in enhancing model robustness.

Related Work

Adverse weather conditions, such as fog, rain, and snow, pose significant challenges to LiDAR-based perception systems by introducing noise and reducing visibility in point cloud data. To address these challenges, researchers have developed various simulation methods (Manivasagam et al. 2020; Bijelic et al. 2020) to augment clear-weather LiDAR datasets with realistic adverse weather effects, facilitating the training of more robust 3D object detection models (Kilic et al. 2025; Piroli et al. 2023; Chae, Kim, and Yoon 2024).

Fog and Snowfall Simulation. Hahner et al. (Hahner et al. 2021) proposed a physically based method to simulate fog on real LiDAR point clouds. Their approach modifies clear-weather measurements by incorporating the optical properties of fog, such as attenuation and backscattering, to generate realistic foggy point clouds. This simulation enables the augmentation of existing datasets, improving the performance of 3D object detection models under foggy conditions. Building upon their fog simulation framework, Hahner et al. (Hahner et al. 2022) developed a snowfall simulation technique for LiDAR data. This method involves sampling snow particles in 2D space for each LiDAR line and adjusting the measurements to reflect the presence of snow. Additionally, the simulation accounts for ground wetness caused by snowfall. Training 3D object detection models with these augmented datasets has been shown to enhance their robustness in snowy environments. Yang et al. (Yang et al. 2024) introduced a realistic rainy weather simulation for LiDAR sensors within the CARLA simulator. Their approach models rain effects, including splashes and sprays caused by vehicle wheels, and incorporates these phenomena into the LiDAR data. By generating synthetic point clouds under various rainfall scenarios, they demonstrated improvements in object detection performance in rainy conditions.

Rain and Spray Simulation. Recent efforts have moved beyond generic weather effects to focus on localized, high-density particulate phenomena such as water spray (Yang et al. 2024; Linnhoff et al. 2024). Scheuble et al. (Scheuble et al. 2024) proposed a physically-informed method for simulating road spray caused by vehicle tires. Their approach combines computational fluid dynamics with probabilistic LiDAR sensor models to synthesize realistic spray-affected returns, with special attention to multipath effects and occlusions. Complementary to simulation, Shih et al. (Shih et al. 2022) presented a reconstruction-based method for modeling spray from real-world LiDAR scans. Their system uses a multi-pass scanning strategy with water sprays in controlled environments to build a volumetric representation of the spray and extract LiDAR response characteristics. They further synthesize new point clouds by sampling from this reconstructed volume, producing realistic and physically plausible spray artifacts. These simulation techniques provide cost-effective means to generate large-scale annotated datasets under adverse weather conditions, thereby enhancing the resilience of LiDAR-based perception systems in autonomous driving applications.

Smoke Simulation. To the best of our knowledge, the only work that explores smoke in the 3D LiDAR domain is Zheng et al. (Zheng et al. 2025), which introduces a generative adversarial network (GAN)-based framework to simulate smoke-affected LiDAR point cloud sequences. However, their goal is to investigate adversarial vulnerabilities of 3D detectors, rather than to enhance perception performance under real-world smoke conditions. Moreover, their approach lacks physical grounding, as it does not model the underlying optical or dynamic properties of smoke. In contrast, our work adopts a physics-based simulation strategy aimed at improving robustness through physically plausible, domain-aligned augmentation.

Despite the growing interest in simulating adverse environments, existing works predominantly focus on fog, rain, and snow, while smoke, a frequent hazard in real-world driving due to wildfires, industrial leaks, and urban accidents, remains significantly underexplored. Moreover, most methods operate by modifying existing clean LiDAR data, rather than synthesizing sensor responses from first principles. Our work addresses this gap by using a particle-based simulation of smoke in Unity3D, coupled with a shader-driven LiDAR scanning pipeline, to generate smoke point clouds directly, enabling precise control over density, motion, and spatial extent.

Method

Unified Smoke Dynamics Model

To characterize the spatiotemporal behavior of smoke, we formulate a unified representation that jointly evolves three fields: smoke concentration $\rho(x,t)$, temperature $T(x,t)$, and fluid velocity $\mathbf{u}(x,t)$. Although these quantities differ in their physical roles, they share a common evolution structure governed by advection, diffusion, and field-specific sourc terms. Formally, for any field $\phi \in \rho, T, \mathbf{u}$, we adopt the generalized transport formulation:

$$\frac{\partial \phi}{\partial t} + (\mathbf{u} \cdot \nabla)\phi = \mathcal{D}(\phi) + \mathcal{S}(\phi), \quad (1$$

where the left-hand side models material advection b the velocity field $\mathbf{u}(x, t)$, while the right-hand side includes field-specific diffusion $\mathcal{D}(\phi)$ and physically motivated source terms $\mathcal{S}(\phi)$. Expanding this framework, th governing equations are instantiated as follows:

Smoke Concentration Dynamics. Let $\rho(x, t) \in [0, 1$ denote the normalized smoke concentration at spatial position x and time t. Its evolution is governed by an advection diffusion-reaction equation:

$$\frac{\partial \rho}{\partial t} + \mathbf{u} \cdot \nabla \rho = D_\rho \nabla^2 \rho - \gamma_\rho \rho + S_\rho(x, t), \quad (2)$$

where D_ρ is the diffusion coefficient, γ_ρ is the decay rate due to dissipation or chemical loss, and $S_\rho(x, t)$ captures time-varying smoke injection from the emitter.

Thermal Field Evolution. The temperature field evolves under advection and thermal diffusion, and additionally undergoes radiative cooling modeled by the normalized quartic term $(T/T_{\max})^4$. Here, $T_{\max}$ serves as a reference temperature scale that stabilizes the nonlinear cooling term and reflects the expected upper bound of heat released by the smoke source. This coupling enables buoyant forces in the velocity field and governs the rise and spread of hot smoke plumes.

$$\frac{\partial T}{\partial t} + \mathbf{u} \cdot \nabla T = D_T \nabla^2 T - \sigma_{\text{cool}} \left(\frac{T}{T_{\max}} \right)^4 + S_T(x, t). \quad (3)$$

Incompressible Flow Dynamics. The velocity field $\mathbf{u}(x, t)$ governing smoke motion follows the incompressible Navier–Stokes equations:

$$\frac{\partial \mathbf{u}}{\partial t} + (\mathbf{u} \cdot \nabla)\mathbf{u} = -\frac{1}{\rho_0} \nabla p + \nu \nabla^2 \mathbf{u} + \mathbf{f}, \quad (4)$$

$$\nabla \cdot \mathbf{u} = 0, \quad (5)$$

where ν is the kinematic viscosity, ρ_0 is the reference density, and $\mathbf{f}$ encompasses external forces such as buoyancy and vorticity confinement. These forces couple the temperature and velocity fields, enabling realistic upward smoke motion and turbulent detail.

Having established the governing physics, we next describe the implementation of our simulation pipeline (Figure 2), its integration with LiDAR sensing, and the cross-domain fusion procedure used to convert simulated smoke into training-ready point clouds.

Physically-Based Smoke Simulation

This section describes the physically based smoke simulation, which represents both the fluid velocity field $\mathbf{u}(\mathbf{x}, t)$ and the smoke density field $\phi(\mathbf{x}, t)$ on a uniform three-dimensional Eulerian grid. The simulation advances in time through four primary stages: advection of the velocity and scalar fields, application of external forces, pressure projection to enforce incompressibility, and updating of the smoke density. All computations are performed on a fixed Eulerian grid using efficient numerical methods.

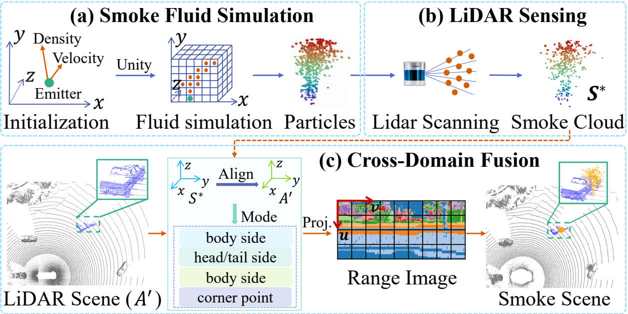

Figure 2: Pipeline of the smoke-scene synthesis framework. Unity-based fluid simulation generates smoke particles, which are converted into LiDAR returns and fused with real-scene point clouds in range-image space to produce geometry-consistent smoke-scene data.

Advection The advection step simulates the transport of fluid quantities along the velocity field. Given the current time step n, the smoke density and velocity at the next step are computed via the semi-Lagrangian method (Stam 1999):

$$\phi^{n+1}(\mathbf{x}) = \phi^n\big(\mathbf{x} - \Delta t\, \mathbf{u}^n(\mathbf{x})\big), \quad (6)$$

$$\mathbf{u}^*(\mathbf{x}) = \mathbf{u}^n\big(\mathbf{x} - \Delta t\, \mathbf{u}^n(\mathbf{x})\big), \quad (7)$$

Here, $\mathbf{x}$ denotes the sampling location. The semi-Lagrangian method traces the departure point of each grid node backward along the velocity field, thus circumventing the strict time-step constraints of explicit Eulerian schemes. Since departure points generally do not lie on grid nodes, trilinear interpolation is employed to estimate the field values.

External Forces The external force term $\mathbf{f}$ in the Navier-Stokes equations models physical effects such as buoyancy and vorticity confinement, which are critical for realistic smoke behavior (Fedkiw, Stam, and Jensen 2001).

Buoyancy force. Buoyancy arises from density and temperature differences within the fluid. It is modeled as:

$$\mathbf{f}_{\text{buoy}} = \alpha \phi \mathbf{g} - \beta (T - \bar{T}) \mathbf{g}, \quad (8)$$

where α is the buoyancy coefficient, β is the temperature coefficient, ϕ is the smoke density, $\mathbf{g}$ is the gravity vector, T is the local temperature, and $\bar{T}$ is the average temperature. In practice, during numerical computation, the temperature term is often substituted with the density field to simplify the simulation.

Vorticity confinement force. To enhance small-scale swirling structures in the smoke, we apply a vorticity confinement force defined by:

$$\mathbf{f}_{\text{vort}} = \epsilon \left(\frac{\nabla |\boldsymbol{\omega}|}{|\nabla |\boldsymbol{\omega}||} \times \boldsymbol{\omega} \right), \quad (9)$$

$$\boldsymbol{\omega} = \nabla \times \mathbf{u}, \tag{10}$$

where ϵ controls the strength of this force (Steinhoff and Underhill 1994).

These forces are applied explicitly to the intermediate velocity field after advection:

$$\mathbf{u}^{**} = \mathbf{u}^{*} + \Delta t \left(\mathbf{f}_{\text{buoy}} + \mathbf{f}_{\text{vort}} \right). \tag{11}$$

Viscous Diffusion Viscous effects are modeled by the diffusion term:

$$\frac{\partial \mathbf{u}}{\partial t} = \nu \nabla^2 \mathbf{u}, \tag{12}$$

Discretized implicitly as:

$$\frac{\mathbf{u}^{n+1} - \mathbf{u}^{**}}{\Delta t} = \nu \nabla^2 \mathbf{u}^{n+1}, \tag{13}$$

which results in solving a linear system for $\mathbf{u}^{n+1}$ to maintain numerical stability.

Pressure Projection To enforce incompressibility, the velocity must satisfy:

$$\nabla \cdot \mathbf{u}^{n+1} = 0. \tag{14}$$

The corrected velocity is obtained by subtracting the pressure gradient:

$$\mathbf{u}^{n+1} = \mathbf{u}^{**} - \nabla p. \tag{15}$$

Taking the divergence and applying the incompressibility constraint yields the pressure Poisson equation:

$$\nabla^2 p = \frac{1}{\Delta t} \nabla \cdot \mathbf{u}^{**}. \tag{16}$$

This elliptic equation is discretized on the grid and solved using an iterative method with boundary conditions ensuring no flow through solid boundaries:

$$\mathbf{u}^{n+1} \cdot \mathbf{n} = 0. \tag{17}$$

LiDAR Perception Modeling

The interaction between LiDAR beams and smoke particles involves inherent uncertainties, resulting in discrepancies between the acquired point clouds and ideal measurements in terms of density, spatial accuracy, and temporal consistency. To realistically reflect the statistical characteristics of measured data, the following particle-aware perturbation modeling strategies are adopted (Shih et al. 2022).

Ray-Particle Intersection Each smoke particle is modeled as a sphere with radius r_m centered at $\mathbf{c}$. A LiDAR ray is represented by its origin $\mathbf{o}$ and unit direction $\mathbf{d}$, described by the parametric form:

$$\mathbf{r}(t) = \mathbf{o} + t \cdot \mathbf{d}, \quad t \geq 0 \tag{18}$$

The ray intersects the particle if there exists t such that

$$\|\mathbf{r}(t) - \mathbf{c}\|^2 \leq r_m^2. \tag{19}$$

In practice, an optimized KD-Tree is used to accelerate intersection queries for multiple rays and particles, and the closest intersection point to the ray origin is recorded as the observation.

Point Retention Mechanism Not all particles hit by rays yield effective returns. A random retention mechanism is introduced, where each ray is assigned an independent uniform random variable $u \sim \mathcal{U}[0, 1)$. The point is retained only if $u \leq \tau$, where $\tau \in (0, 1]$ denotes the retention probability:

$$\text{retain}(x) = \mathbf{1}(u \leq \tau). \tag{20}$$

The retention rate τ can be dynamically adjusted based on smoke density, particle distance from the LiDAR, or scene semantics to simulate uneven attenuation and occlusion effects.

Range Measurement Perturbation Range measurement errors are ubiquitous in real LiDAR sensors. To model this, zero-mean Gaussian noise $\epsilon \sim \mathcal{N}(0, \sigma^2)$ is added to the true range r:

$$\hat{r} = r + \epsilon. \tag{21}$$

The perturbed point position is then

$$\hat{\mathbf{x}} = \mathbf{o} + \hat{r} \cdot \mathbf{d}, \tag{22}$$

where the noise standard deviation σ can be set based on smoke concentration and sensor characteristics to reflect environment-dependent accuracy.

Rolling Shutter Effect Considering the sequential scanning nature of multi-beam LiDARs and platform motion, points emitted at different azimuth angles are measured at different times. This temporal offset is modeled by a displacement correction along the lateral scanning direction:

$$\hat{x} = x + \frac{v \cdot \gamma}{2\pi f}, \tag{23}$$

where v is the platform velocity, γ is the azimuth angle of the ray, and f is the LiDAR sampling frequency. This correction is applied prior to ray casting and accounts for spatial drift induced by motion during scanning.

Cross-domain point cloud synthesis

To generate complex perception scenes with non-rigid interferences such as smoke and improve the generalization of point cloud models, we perform geometric fusion of simulated smoke point clouds with real-world scene point clouds in the range image representation space. The range image, a 2D tensor derived from LiDAR scanning sequences, encodes visibility explicitly, enabling occlusion-aware and spatially consistent fusion.

Each smoke point cloud is first transformed into the coordinate system of the target scene. Let O denote the centroid of the smoke point cloud and Δz its half-height:

$$\Delta z = \frac{1}{2} \left(\max_{(x,y,z) \in P_{\text{smoke}}} z - \min_{(x,y,z) \in P_{\text{smoke}}} z \right), \tag{24}$$

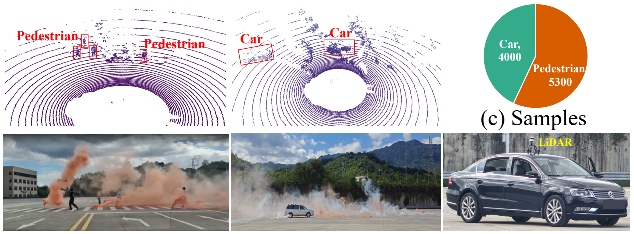

(c) Samples

(a) Smoke & Pedestrians (b) Smoke & Vehicles (d) Data Platform

Figure 3: Illustration of LiSmoke dataset collection and scene composition. From left to right: (a) pedestrian scenes under smoke-LiDAR point cloud (top) and RGB image (bottom); (b) vehicle scenes in smoke-point cloud view (top) and camera image (bottom); (c) dataset composition-frame statistics for each subset (top) and (d) sensor setup on the data collection vehicle (bottom).

where P_{smoke} denotes the smoke point cloud to be fused. A smoke point A is positioned relative to the smoke cloud centroid O with a vertical offset of Δz, i.e.,$A = O - \Delta z$, and is then overlaid onto the corresponding fusion point B on the vehicle point cloud.

We employ four fusion modes to adapt to different vehicle orientations (Zheng et al. 2025):

- **Head/tail Fusion**: When the front or rear of the vehicle faces the LiDAR sensor, the smoke point cloud is attached at the center (point B) of the bottom edge of the corresponding side of the bounding box;
- **Body-Side Fusion**: When one side of the vehicle faces the LiDAR, the fusion point is the center (point B) of the bottom edge of the vehicle's side;
- **Two-Sides Fusion**: When two sides of the vehicle simultaneously face the LiDAR, smoke point clouds are inserted at the bottom midpoints (point B) of both sides;
- **Corner Fusion**: When the LiDAR viewpoint falls on the corner region between two visible faces of the vehicle, the fusion point is set at the bottom vertex (point B) of the intersection edge.

After determining the spatial position and orientation, the smoke point cloud is translated and rotated to complete alignment, and then mapped into the range image space for fusion. This process inherently considers occlusion and geometric consistency, ensuring the inserted points are reasonably integrated into the spatial layout from the original viewpoint. This fusion strategy constructs composite point cloud scenes containing non-rigid interferences, supporting performance evaluation and robustness testing of perception models under complex environmental conditions.

LiSmoke Dataset

To investigate the robustness and modeling capabilities of LiDAR-based perception under smoke interference, we construct a real-world point cloud dataset named **LiSmoke**, specifically designed for smoke-occluded scenarios with non-rigid interference characteristics. The dataset is collected using a vehicle-mounted LiDAR system and covers typical scenes with varying smoke intensities and object categories. LiSmoke serves as a controllable and realistic evaluation benchmark for point cloud perception models in adverse environments. The smoke acquisition setup and scene configuration are illustrated in Figure 3.

Smoke Scene	Light Smoke	Heavy Smoke
Vehicles in Smoke	2000 frames	2000 frames
Pedestrians in Smoke	2900 frames	2400 frames

Table 1: Frame statistics of the LiSmoke dataset for different scene types and smoke levels.

Data Acquisition

LiSmoke is constructed in real outdoor environments using an Ouster OS1-64 LiDAR sensor mounted on the roof of a sedan. To simulate low-visibility conditions typical of emergencies, we release dense smoke using fire safety-grade smoke canisters, each maintaining stable particulate output for about three minutes.

The dataset focuses on two dynamic obstacle types: vehicles and pedestrians. For each, we collect data under two interference levels-light and heavy smoke-by varying the number and duration of smoke releases. Throughout the collection, the LiDAR's pose and parameters remain fixed to ensure consistency across conditions. Frame statistics are summarized in Table 1. LiSmoke contains two subsets:

- **Vehicle Subset**: 2000 frames each under light and heavy smoke.
- **Pedestrian Subset**: 2900 frames (light smoke), 2400 frames (heavy smoke).

Data Characteristics

In the LiSmoke dataset, smoke acts as a non-rigid occluder that causes intensity fluctuations, range perturbations, and localized structural incompleteness in LiDAR point clouds. These effects pose significant challenges to downstream tasks such as detection, segmentation, and tracking, enabling systematic evaluation under degraded conditions.

The dataset features diverse spatial configurations of vehicles and pedestrians across various motion states and occlusion levels, supporting tasks like occlusion-aware modeling and robustness assessment. It also serves as a foundation for exploring cross-modal fusion and uncertainty modeling.

LiSmoke offers the following three key characteristics:

- **Realism**: Acquired using an actual LiDAR sensor under naturally diffused smoke, preserving authentic interference effects from the sensing process.
- **Complexity**: Covers multiple object categories and varying levels of smoke occlusion, forming representative non-rigid disturbance environments.
- **Comparability**: Provides both light and heavy smoke conditions for each scenario, supporting quantitative analysis under different interference levels.

Detection method	Simulation method	heavy smoke ↑		light smoke ↑		clear weather ↑	
		AP	APH	AP	APH	AP	APH
CenterPoint	None	34.06	33.48	46.98	46.50	81.21	80.15
	FogSim	37.88	37.07	51.00	50.45	79.49	78.40
	Zheng et.al	40.73	40.01	51.95	51.32	78.99	77.87
	Ours	43.54	42.81	53.29	52.72	78.63	77.52
PV-RCNN	None	35.34	34.77	49.40	48.87	84.13	82.73
	FogSim	38.76	38.06	55.92	55.30	82.21	80.34
	Zheng et.al	42.03	41.21	57.11	55.95	83.12	81.58
	Ours	44.49	43.32	60.52	59.63	83.05	81.46
Voxel R-CNN	None	40.45	39.86	57.69	57.18	84.19	83.41
	FogSim	42.59	41.91	60.08	59.34	83.03	82.17
	Zheng et.al	45.55	44.87	62.09	61.38	82.99	82.16
	Ours	47.39	46.56	62.50	61.94	83.18	82.37
PointPillars	None	10.58	10.38	4.16	3.92	80.05	78.39
	FogSim	7.34	6.88	8.87	8.09	78.17	76.24
	Zheng et.al	22.19	21.41	33.95	32.77	79.62	77.90
	Ours	26.98	24.48	37.49	34.83	79.81	78.14

Table 2: Car detection performance (AP/APH) under heavy smoke, light smoke from LiSmoke, and clear-weather scenes from the Waymo Open Dataset.

Detection method	Simulation method	heavy smoke ↑	light smoke ↑	clear weather ↑
		AP	AP	AP
CenterPoint	None	28.39	24.69	66.93
	FogSim	31.90	25.82	64.76
	Zheng et.al	34.93	38.08	63.39
	Ours	35.08	38.24	63.41
PV-RCNN	None	18.61	17.94	64.07
	FogSim	36.12	50.36	65.48
	Zheng et.al	47.31	64.21	62.28
	Ours	47.95	65.02	63.35
Voxel R-CNN	None	27.82	18.90	70.67
	FogSim	37.29	33.81	69.28
	Zheng et.al	43.89	46.08	67.70
	Ours	44.61	46.57	67.85
PointPillars	None	8.82	13.12	52.53
	FogSim	8.13	16.40	49.36
	Zheng et.al	23.61	30.50	53.65
	Ours	24.10	31.06	53.86

Table 3: Pedestrian detection performance (AP) under heavy smoke, light smoke from LiSmoke, and clear-weather scenes from the Waymo Open Dataset.

Experiments

Experimental Setup

Dataset. We conduct experiments on both a large-scale open-source dataset and a real-world smoke dataset. **The Waymo Open Dataset** (Sun et al. 2020) provides high-quality 3D LiDAR point clouds with annotations for vehicles, pedestrians, and other traffic participants; we use its vehicle and pedestrian segments for baseline training and as clean reference data. The **LiSmoke** Dataset is a real-world point cloud dataset captured under controlled smoke interference, containing vehicle and pedestrian subsets recorded at two smoke densities (light and heavy). It enables evaluation of perception robustness under non-rigid occlusions caused by smoke.

Experimental Settings. All 3D object detectors are trained from scratch using OpenPCDet (Team et al. 2020) with default settings. Key parameters for the smoke simulation are set as follows: $\nu = 0.0002$, $\alpha = 0.000625$, $\beta = 0.025$, and $\tau = 0.17$.

Quantitative Results

Performance under Smoke Interference. We evaluate our smoke simulation and synthesis approach on four widely used 3D object detectors-CenterPoint (Yin, Zhou, and Krahenbuhl 2021), PV-RCNN (Shi et al. 2020), Voxel R-CNN (Deng et al. 2021), and PointPillars (Lang et al. 2019)-representing diverse LiDAR architectures. Each model is trained under different simulation settings and tested on heavy smoke, light smoke (LiSmoke), and clear-weather scenes from the Waymo Open Dataset. To validate robustness under smoke conditions, we compare against two state-of-the-art LiDAR simulation baselines: the fog simulation of (Hahner et al. 2021) and the adversarial approach of (Zheng et al. 2025). Specifically, the None setting serves as the baseline simulation method, where models are trained solely on clear-weather data from the Waymo Open Dataset, without any smoke simulation or augmentation. In contrast, FogSim (Hahner et al. 2021), Zheng et al. (Zheng et al. 2025), and Ours denote models trained with a mixture of Waymo data and corresponding simulated smoke-augmented data generated by each method.

The quantitative results are summarized in Table 2 for car detection and Table 3 for pedestrian detection. When evaluated on the LiSmoke, models trained only on clean data (None) show a substantial performance drop, especially under heavy smoke. This indicates the severe domain gap between clean and smoke-degraded environments.

Compared to existing simulation methods, our approach consistently achieves the highest detection performance under both light and heavy smoke settings. For example, under heavy smoke, our simulation boosts the AP of CenterPoint from 34.06 to 43.54 for cars, and from 28.39 to 35.08 for pedestrians. Similar improvements are observed for PV-RCNN, Voxel R-CNN, and PointPillars, demonstrating the generality of our method across diverse architectures.

Furthermore, we observe that all simulation methods slightly reduce performance on the original Waymo validation set due to the domain mixing. However, our method maintains competitive results in clear-weather conditions (e.g., 78.63 AP for CenterPoint), demonstrating that the added robustness under smoke does not come at the cost of significant performance degradation in clean scenarios. These results verify the effectiveness of our physically grounded simulation strategy, which enables the trained models to better generalize to real-world smoke disturbances while preserving performance in normal environments.

Qualitative Analysis

To further demonstrate the robustness of different simulation methods under smoke interference, we provide qualitative comparisons on LiSmoke samples in Figure 4. The visualizations include both vehicle and pedestrian scenarios, each affected by varying degrees of smoke occlusion.

In the vehicle case (top row), the model trained with FogSim (Hahner et al. 2021) fails to detect the vehicle heavily occluded by smoke, while only the unobstructed vehicle is correctly identified. The model trained with Zheng

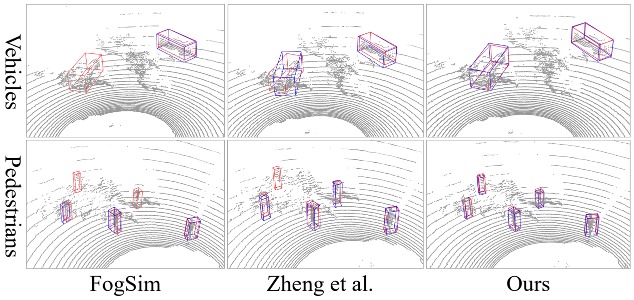

Figure 4: Qualitative comparison of detection results under smoke interference. Blue boxes indicate model predictions, and red boxes denote ground truth.

et al. (Zheng et al. 2025) can detect the occluded vehicle, but the predicted bounding box shows large misalignment with the ground truth. In contrast, the model trained using our method correctly detects both vehicles with more accurate box localization. Similar patterns are observed in the pedestrian case (bottom row). Only our method successfully detects both occluded and unobstructed pedestrians, while other methods either miss detections or exhibit severe box deviations.

Ablation Study

Fusion modes. Employing CenterPoint as the baseline detector, we evaluate four fusion strategies to inject synthetic smoke into clean LiDAR frames: body side, two sides, head-/tail side, and corner point. As shown in Table 4, the head/tail side mode achieves the best overall performance across all weather conditions. It enables more realistic and diverse occlusion patterns, especially for frontal and rear-view detection, leading to improved robustness. The corner point mode also performs well, while the body side fusion gives the lowest performance under smoke due to its limited occlusion variation.

Fusion mode	**heavy smoke ↑**		**light smoke ↑**		**clear weather ↑**	
	AP	APH	AP	APH	AP	APH
body side	43.54	42.81	53.29	52.72	78.63	77.52
two sides	45.38	44.69	52.52	51.93	78.05	76.95
head/tail side	50.65	49.85	56.53	55.90	79.07	77.94
corner point	48.17	47.44	53.45	52.81	79.51	78.43

Table 4: Comparison of different smoke fusion modes on CenterPoint.

Smoke Density. We evaluate the effect of smoke density on CenterPoint's performance by randomly sampling a percentage of points from the original smoke point cloud before fusion. As shown in Table 5, full-density fusion (100%) achieves the best detection under smoke conditions. Reducing the density to 80% leads to a clear performance drop, while lower densities (60%, 40%) slightly recover generalization in clear weather but remain less effective under smoke. These results suggest that denser synthetic smoke helps models better handle real-world occlusions.

Smoke Density	**heavy smoke ↑**		**light smoke ↑**		**clear weather ↑**	
	AP	APH	AP	APH	AP	APH
100%	43.54	42.81	53.29	52.72	78.63	77.52
80%	39.49	38.77	47.42	46.85	58.34	57.71
60%	43.91	43.20	45.06	44.46	58.16	57.49
40%	43.87	43.15	50.24	49.71	58.76	58.11

Table 5: Impact of smoke point sampling density on CenterPoint detection performance.

Spraying Angle. We investigate the impact of different spraying angle ranges on CenterPoint's performance during smoke simulation. Here, the spraying angle denotes the maximum zenith angle within which the smoke is randomly rotated around its base center, forming a cone aligned with the z-axis. As shown in Table 6, moderate ranges such as [0°-20°] lead to the best performance under heavy and light smoke. Smaller angles may result in limited occlusion patterns, while overly large angles (e.g., [0°-40°]) introduce unrealistic artifacts that degrade detection. These findings highlight the importance of simulating physically plausible occlusion distributions.

Spraying Angle	**heavy smoke ↑**		**light smoke ↑**		**clear weather ↑**	
	AP	APH	AP	APH	AP	APH
None	43.54	42.81	53.29	52.72	78.63	77.52
$[0° - 10°]$	49.31	48.54	52.68	52.00	58.79	58.16
$[0° - 20°]$	51.77	50.77	54.01	53.40	59.36	58.70
$[0° - 30°]$	48.72	48.03	51.75	51.09	78.95	77.85
$[0° - 40°]$	47.03	46.05	52.12	51.44	79.19	78.09

Table 6: Effect of simulated smoke spraying angle on CenterPoint detection performance.

Conclusion

This work tackles the difficulty of LiDAR perception in smoke-occluded environments by introducing a physically grounded smoke simulation framework that couples 3D fluid dynamics with an accurate LiDAR scanning model. The proposed approach generates realistic non-rigid occlusion patterns and enables controllable synthesis of smoke-affected point clouds. In addition, we construct LiSmoke, a real-world dataset containing diverse targets and multiple smoke density levels, providing a valuable benchmark for studying degradation effects and robustness in LiDAR-based perception. A range-image–based fusion strategy is further introduced to integrate simulated smoke with large-scale driving datasets while preserving geometric consistency. Extensive experiments demonstrate that training with our synthetic smoke data notably improves 3D object detection performance under smoke interference, while maintaining strong accuracy in clear conditions.

Acknowledgments

This work was supported in part by the National Natural Science Foundation of China (No. 62401225, 62471415), the Natural Science Foundation of Fujian Province of China (No. 2025J0141, 2024J01115, 2023J01004), the Natural Science Foundation of Xiamen of China (No. 3502Z202472018), and the Jimei University Scientific Research Start-up Funding Project (No.ZQ2024034).

References

Bijelic, M.; Gruber, T.; Mannan, F.; Kraus, F.; Ritter, W.; Dietmayer, K.; and Heide, F. 2020. Seeing through fog without seeing fog: Deep multimodal sensor fusion in unseen adverse weather. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 11682–11692.

Caesar, H.; Bankiti, V.; Lang, A. H.; Vora, S.; Liong, V. E.; Xu, Q.; Krishnan, A.; Pan, Y.; Baldan, G.; and Beijbom, O. 2020. nuscenes: A multimodal dataset for autonomous driving. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 11621–11631.

Chae, Y.; Kim, H.; and Yoon, K.-J. 2024. Towards robust 3d object detection with lidar and 4d radar fusion in various weather conditions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 15162–15172.

Deng, J.; Shi, S.; Li, P.; Zhou, W.; Zhang, Y.; and Li, H. 2021. Voxel r-cnn: Towards high performance voxel-based 3d object detection. In *Proceedings of the AAAI conference on artificial intelligence*, volume 35, 1201–1209.

Fedkiw, R.; Stam, J.; and Jensen, H. W. 2001. Visual simulation of smoke. In *Proceedings of the 28th annual conference on Computer graphics and interactive techniques*, 15–22.

Hahner, M.; Sakaridis, C.; Bijelic, M.; Heide, F.; Yu, F.; Dai, D.; and Van Gool, L. 2022. Lidar snowfall simulation for robust 3d object detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 16364–16374.

Hahner, M.; Sakaridis, C.; Dai, D.; and Van Gool, L. 2021. Fog simulation on real LiDAR point clouds for 3D object detection in adverse weather. In *Proceedings of the IEEE/CVF international conference on computer vision*, 15283–15292.

Kilic, V.; Hegde, D.; Cooper, A. B.; Patel, V. M.; and Foster, M. 2025. Lidar light scattering augmentation (lisa): Physics-based simulation of adverse weather conditions for 3d object detection. In *ICASSP 2025-2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 1–5. IEEE.

Lang, A. H.; Vora, S.; Caesar, H.; Zhou, L.; Yang, J.; and Beijbom, O. 2019. Pointpillars: Fast encoders for object detection from point clouds. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 12697–12705.

Linnhoff, C.; Scheuble, D.; Bijelic, M.; Elster, L.; Rosenberger, P.; Ritter, W.; Dai, D.; and Winner, H. 2024. Simulating road spray effects in automotive lidar sensor models. *IEEE Intelligent Vehicles Symposium*.

Manivasagam, S.; Wang, S.; Wong, K.; Zeng, W.; Sazanovich, M.; Tan, S.; Yang, B.; Ma, W.-C.; and Urtasun, R. 2020. Lidarsim: Realistic lidar simulation by leveraging the real world. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 11167–11176.

Piroli, A.; Dallabetta, V.; Kopp, J.; Walessa, M.; Meissner, D.; and Dietmayer, K. 2023. Towards robust 3D object detection in rainy conditions. In *2023 IEEE 26th International Conference on Intelligent Transportation Systems (ITSC)*, 3471–3477. IEEE.

Scheuble, D.; Linnhoff, C.; Bijelic, M.; Elster, L.; Rosenberger, P.; Ritter, W.; and Winner, H. 2024. Simulating Road Spray Effects in Automotive Lidar Sensor Models. In *2024 IEEE Intelligent Vehicles Symposium (IV)*, 659–666. IEEE.

Shi, S.; Guo, C.; Jiang, L.; Wang, Z.; Shi, J.; Wang, X.; and Li, H. 2020. Pv-rcnn: Point-voxel feature set abstraction for 3d object detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 10529–10538.

Shih, Y.-C.; Liao, W.-H.; Lin, W.-C.; Wong, S.-K.; and Wang, C.-C. 2022. Reconstruction and synthesis of lidar point clouds of spray. *IEEE Robotics and Automation Letters*, 7(2): 3765–3772.

Stam, J. 1999. Stable Fluids. In Waggenspack, W. N., ed., *Proceedings of the 26th Annual Conference on Computer Graphics and Interactive Techniques, SIGGRAPH 1999, Los Angeles, CA, USA, August 8-13, 1999*, 121–128. ACM.

Steinhoff, J.; and Underhill, D. 1994. Modification of the Euler equations for “vorticity confinement”: Application to the computation of interacting vortex rings. *Physics of Fluids*, 6(8): 2738–2744.

Sun, P.; Kretzschmar, H.; Dotiwalla, X.; Chouard, A.; Patnaik, V.; Tsui, P.; Guo, J.; Zhou, Y.; Chai, Y.; Caine, B.; et al. 2020. Scalability in perception for autonomous driving: Waymo open dataset. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2446–2454.

Team, O.; et al. 2020. Openpcdet: An open-source toolbox for 3d object detection from point clouds.

Yang, D.; Cai, X.; Liu, Z.; Jiang, W.; Zhang, B.; Yan, G.; Gao, X.; Liu, S.; and Shi, B. 2024. Realistic rainy weather simulation for LiDARs in CARLA simulator. In *2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 951–957. IEEE.

Yin, T.; Zhou, X.; and Krahenbuhl, P. 2021. Center-based 3d object detection and tracking. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 11784–11793.

Zhang, Z.; Chen, M.; Xiao, S.; Peng, L.; Li, H.; Lin, B.; Li, P.; Wang, W.; Wu, B.; and Cai, D. 2024. Pseudo label refinery for unsupervised domain adaptation on cross-dataset 3d object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 15291–15300.

Zheng, S.; Liu, W.; Guo, Y.; Zang, Y.; Shen, S.; and Wang, C. 2025. A New Adversarial Perspective for LiDAR-Based 3D Object Detection. In *Proceedings of the AAAI Conference on Artificial Intelligence*.