

Source-Only Cross-Weather LiDAR via Geometry-Aware Point Drop

YoungJae Cheong, Jonghyun An*

Abstract—Adverse weather conditions, such as rain, snow, and fog, severely degrade LiDAR semantic segmentation by introducing refraction, scattering, and point dropouts that compromise geometric integrity. While prior approaches ranging from weather simulation and mixing-based augmentation to domain randomization and regularization enhance robustness, they frequently overlook structural vulnerabilities inherent to object boundaries, corners, and highly sparse regions.

To address this limitation, we propose a Light Geometry-Aware Adapter. This module aligns azimuths and applies horizontal circular padding to preserve neighbor continuity across the 0° – 360° wrap-around boundary. Using a local-window K-Nearest Neighbors (KNN) search, it aggregates nearby points and computes lightweight local statistics, compressing them into compact geometry-aware cues. During training, these cues facilitate region-aware regularization, which effectively stabilizes predictions in structurally fragile areas. The proposed adapter is designed to be plug-and-play, complements existing augmentation techniques, and operates exclusively during training, incurring negligible inference overhead.

Operating under a rigorous *source-only* cross-weather paradigm wherein models are trained on SemanticKITTI and evaluated on SemanticSTF without target-domain labels or fine-tuning, our adapter achieves a +3.4 mIoU improvement over strong data-centric augmentation baselines. Furthermore, it demonstrates performance comparable to advanced class-centric regularization methods. These findings highlight that geometry-driven regularization constitutes a critical pathway toward achieving highly robust, all-weather LiDAR segmentation.

I. INTRODUCTION

LiDAR provides accurate ranging for autonomous driving, but adverse weather such as rain, snow, and fog induces refraction, scattering, and point dropouts that distort point clouds and degrade segmentation [1], [2]. *Foreground* classes (vehicles, pedestrians, riders) are especially vulnerable due to low point density, making errors **safety-critical**. Generic regularizers such as Dropout [3] do not directly address *structural* distortions, including boundary breaks, corner loss, and noisy sparse regions. We study *source-only* cross-weather transfer, training on SemanticKITTI [4] and evaluating on SemanticSTF [5] without target labels or fine-tuning.

Prior work largely follows two strategies. The first uses data augmentation and *domain randomization* to mimic weather via dropping, jittering, or simulation [5], [6], [7]. The second encourages *invariance and consistency* under

This work was supported by the Project for Collaboration RD between Industry, University, and Research Institute, funded by the Ministry of SMEs and Startups of Korea in 2025 (RS-2025-02220569) and in part by the Gachon University, South Korea under Grant GCU-202400470001.

YoungJae Cheong and Jonghyun An are with the School of Computing, Gachon University, Seongnam, Republic of Korea (e-mail: bluebull777@gachon.ac.kr; jhonghyun@gachon.ac.kr). Corresponding author: Jonghyun An.




Fig. 1: Failure Cases in Adverse Weather [2]. (a) The baseline misclassifies a bus, merging it into a nearby building. Conversely, our *Light Geometry-Aware Adapter* preserves local neighbor continuity, accurately segmenting the bus (blue). (b) The baseline blurs the semantic boundary between a building and a fence, whereas our method distinctly preserves this structural boundary.

distribution shifts through mixing-based generalization, uncertainty mitigation, or regularization. However, augmentation alone often fails to capture **structural vulnerabilities** at boundaries and in sparse regions, and consistency-based methods frequently omit explicit local geometry, which can increase confusion between boundary-adjacent classes. We adopt the source-only setting throughout.

Fig. 1 elucidates the core motivation of our approach. A standard data-centric augmentation baseline [2] is prone to merging a *bus* into an adjacent *building* or blurring the precise boundary between a *building* and a *fence* when neighbor relations are corrupted by weather noise. In contrast, our proposed methodology actively preserves neighbor continuity and curtails boundary leakage by explicitly modeling local point-to-point structural relationships.

We therefore introduce a **Light Geometry-aware adapter** that injects two geometry signals into training. First, it preserves *neighbor continuity* from the LiDAR scanning pattern by aligning azimuth and applying horizontal circular padding across the 0° – 360° seam. Second, it aggregates a *local-window KNN* neighborhood to compute simple local statistics (offsets and dispersion) that summarize *structural risk* at boundaries, corners, and sparse zones. These geometry-aware cues guide *region-aware regularization* and stabilize learning under weather-induced sparsity. Unlike PointDR [5], which broadens the training distribution via

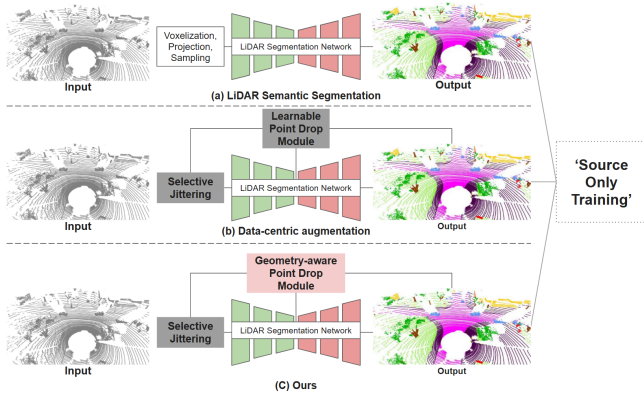


Fig. 2: Comparison of LiDAR segmentation pipelines under the *source-only* setting. (a) Conventional networks lacking structural noise modeling. (b) Data-centric augmentation employing selective jittering and learnable point drop [2]. (c) Our pipeline, featuring the *Light Geometry-Aware Adapter*, which systematically enhances region-level dropping by anchoring it to local geometric continuity.

large-scale domain randomization, our adapter explicitly links local geometry to the learning signal. It is lightweight, plug-and-play, and complementary to augmentation. On SemanticKITTI→SemanticSTF, we obtain consistent gains over strong data-centric baselines and also improve over a class-centric regularization baseline [1]. Fig. 2 illustrates how the adapter enhances region-level dropping for robust segmentation in adverse weather.

Our main contributions are:

- 1) **Light Geometry-aware adapter.** We distill a lightweight geometry module that preserves neighbor continuity via circular padding and local-window KNN, providing geometry-aware cues before dropping with low computation.
- 2) **Geometry-conditioned robustness.** We introduce region-aware regularization driven by these cues to reduce confusion at boundaries, corners, and sparse regions under adverse weather.

II. RELATED WORK

A. Semantic Segmentation of LiDAR Point Clouds

LiDAR semantic segmentation methods are typically categorized into *voxel-based*, *projection-based (range-view)*, and *point-based* approaches. Voxel-based models exploit sparse 3D convolutions for large-scale outdoor scenes, e.g., SPVCNN [9] and Cylinder3D [10], and transformer variants further extend this line, e.g., SphereFormer [11]. Projection-based methods render scans to range images and apply 2D CNN backbones, such as SqueezeSeg [12] and RangeNet++ [13]. Point-based models operate directly on points, from PointNet/PointNet++ [14], [15] to kernel-based and lightweight designs, e.g., KPConv [16] and RandLA-Net [17]. While accuracy has improved steadily, adverse-weather robustness remains challenging, motivating geometry-aware components that stabilize local structure.

B. Augmentation and Domain Generalization for LiDAR

Data augmentation and synthesis broaden training distributions through geometric and intensity perturbations, as well as mixing-based strategies. PolarMix [6] and UniMix [7] recombine scans to increase diversity and reduce domain gaps. Domain generalization and adaptation further align representations with distribution-matching objectives, e.g., DGAF-L [18] and PCL [19]. These methods improve coverage or alignment but do not explicitly preserve neighborhood continuity under weather-induced sparsity. Our adapter is complementary, providing local geometry cues that can be used alongside existing augmentations.

C. LiDAR Perception in Adverse Weather

Environmental phenomena such as rain, snow, and fog introduce severe refraction, scattering, and point dropouts, precipitously degrading perception accuracy [1], [2]. The SemanticSTF dataset [5] provides a critical real-world benchmark for evaluating these specific failure modes. Recent data-centric pipelines have integrated selective jittering with Reinforcement Learning (RL)-guided point dropping to systematically target noise-driven vulnerabilities [2]. In a related vein, TripleMixer [8] employs a Geometry Mixer to model local relations. Distinct from these approaches, our *Light Geometry-Aware Adapter* extracts non-parametric geometric cues from local-window neighborhoods utilizing circular wrapping. These specific cues directly inform and support region-aware training objectives, dynamically adapting to structural fragility in adverse weather.

III. METHOD

A. Overview

Our objective is to strengthen a *data-centric augmentation* pipeline based on Selective Jittering (SJ) and an RL-based Learnable Point Drop (LPD) module [2] by integrating a **Light Geometry-Aware Adapter**, so that *structural vulnerabilities* under adverse weather are explicitly reflected in the training signal (Fig. 3). SJ emulates weather-induced corruptions, while LPD selects dropping regions and ratios from region-level statistics. The adapter is inserted **only before LPD**. It preserves neighborhood continuity across the 0° – 360° seam via horizontal *circular wrapping/padding* and extracts compact *geometry-aware cues* using a *local-window KNN*. These cues are injected into the RL agent state to steer region-level dropping toward boundaries, corners, and sparse regions, reducing boundary-adjacent semantic confusion (e.g., *bus*↔*building*; Fig. 1). Unlike PointDR [5], which relies on domain randomization, our adapter deterministically couples local geometry with the dropping policy. We evaluate this architecture under the SemanticKITTI→SemanticSTF *source-only* transfer setting.

B. In Detail

Selective Jittering (SJ) : We adopt SJ as introduced in prior *data-centric augmentation* work [2]. SJ injects *non-uniform*, frame-wise perturbations that mimic adverse weather by selectively applying small offsets or noise to a

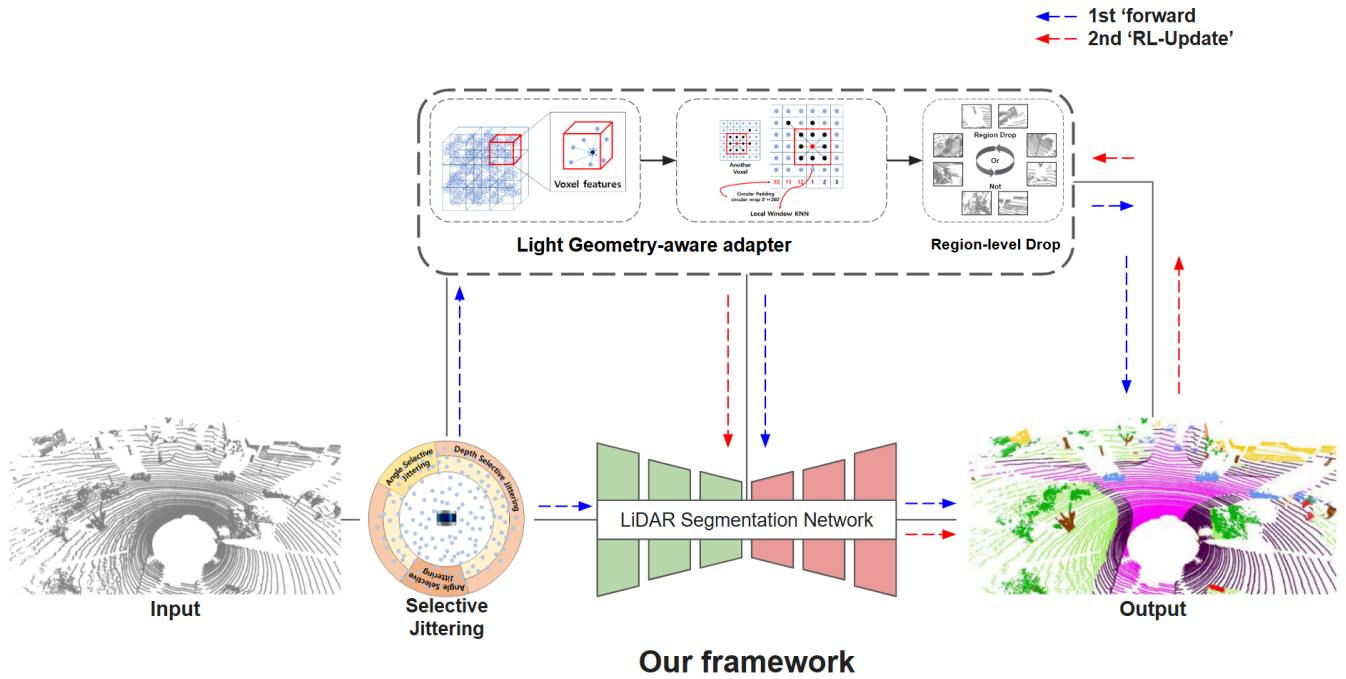


Fig. 3: The Overall Pipeline. Selective jittering emulates weather-induced artifacts [2]. Before the learnable point drop module, the *Light Geometry-Aware Adapter* extracts compact geometry-aware cues via local-window KNN with circular wrapping to preserve $0^\circ\text{--}360^\circ$ continuity. These cues guide **region-level regularization** at boundaries, corners, and sparse structures, improving robustness under adverse weather.




Fig. 4: Global KNN constructs neighborhoods from the entire scan, often matching distant points across the azimuth seam, which increases overhead and weakens locality [8]. In contrast, our *Light Geometry-Aware Adapter* uses horizontal local-window KNN with circular padding to wrap the $0^\circ\text{--}360^\circ$ seam, preserving edge and corner cues with lightweight computation.

subset of points, conditioned on simple factors such as range, bearing, and intensity. Unperturbed points remain unchanged. The perturbed input is fed to both the reference and drop branches. This produces supervision that captures *noise-driven* vulnerabilities alongside *structure-driven* ones.

Light Geometry-aware adapter. For each query point p_i , we retrieve K neighbors $\mathcal{N}_K(i) = \{p_i^1, \dots, p_i^K\}$ using a horizontal *local-window KNN* and apply *circular padding* to preserve neighbor continuity across the $0^\circ\text{--}360^\circ$ azimuth seam. Compared to conventional mixers that use a *global KNN* search over the entire scan [8], our restricted candidate window reduces spurious cross-boundary matches and lowers computation (Fig. 4). We first stabilize local features by voxel-wise mean pooling, yielding a base feature v_i for

p_i within the window, and compute a windowed mean coordinate μ_i . We then summarize local geometry with two lightweight dispersion indicators: $d_i^{(1)} = \|p_i - \mu_i\|$ (off-center geometry near boundaries/corners) and $d_i^{(2)} = \frac{1}{K} \sum_{k=1}^K \|p_i^k - \mu_i\|$ (spread/compactness of local support). Using these quantities, we mix p_i with each neighbor and pool them with attention:

$$\ell_i^k = \phi_p(p_i \oplus p_i^k \oplus (p_i - p_i^k) \oplus \mu_i \oplus [d_i^{(1)}, d_i^{(2)}]), \quad (1)$$

$$\alpha_i^k = \text{FC}(\ell_i^k), \quad c_i^k = \text{Softmax}_k(\alpha_i^k)$$

$$f_i^{\text{pt}} = \sum_{k=1}^K c_i^k \ell_i^k. \quad (2)$$

Finally, we form a geometry-aware cue g_i by fusing the pooled feature f_i^{pt} with the base feature v_i :

$$g_i = \text{MLP}(v_i \oplus f_i^{\text{pt}}) \oplus p_i, \quad (3)$$

where g_i augments only the *pre-LPD* agent state, keeping the inference path unchanged. g_i is the final geometry-aware cue.

RL-Based Region-Level Point Drop: We retain the core LPD architecture established in prior work [2], purposefully *augmenting* the RL agent’s state representation with the newly derived geometry-aware cues. While the standard LPD primarily identifies *noise-induced* vulnerabilities via rudimentary loss and uncertainty metrics, our enhanced state matrix provides an explicit, quantifiable signal for *structural*

TABLE I: Quantitative evaluation on SemanticKITTI→SemanticSTF. Our method improves per-class IoU and overall mIoU over strong baselines, showing that geometry-conditioned, region-aware point dropping enhances structural robustness for cross-weather adaptation.

Method	car	bi.cle	mt.cle	truck	bus	pers	bi.clst	mt.clst	road	parki.	sidew.	oth.g.	build.	fence	veget.	trunk	terra.	pole	traf.	mIoU
Oracle	89.4	42.1	0.0	59.9	61.2	69.6	39.0	0.0	82.2	21.5	58.2	45.6	86.1	63.6	80.2	52.0	77.6	50.1	61.7	54.7
Source-only	55.9	0.0	0.2	0.2	10.9	10.3	6.0	0.0	61.2	10.9	32.0	0.0	67.9	41.6	49.8	27.9	40.8	29.6	15.7	24.4
Dropout [3]	62.1	0.0	15.5	3.0	11.5	5.4	2.0	0.0	58.4	12.8	26.7	1.1	72.1	43.6	52.9	34.2	43.5	28.4	15.5	25.7
Perturbation	74.4	0.0	0.0	23.3	0.6	19.7	20.0	0.0	0.0	59.3	10.7	32.0	7.2	70.2	45.2	57.1	47.9	28.2	16.2	25.9
PolarMix [6]	57.8	1.8	3.6	3.7	26.5	3.7	26.5	0.0	65.7	2.9	35.9	48.7	71.0	58.7	53.8	20.5	45.4	29.3	15.8	26.6
MMD [18]	63.6	0.0	2.6	17.4	11.4	28.1	0.0	0.0	67.0	14.1	37.6	41.2	67.1	41.2	57.1	22.4	47.9	28.2	16.2	26.9
PCL [19]	65.9	0.0	0.4	17.3	8.4	8.4	8.4	0.0	59.6	12.0	35.3	63.1	74.0	47.5	60.7	15.8	48.9	26.1	27.5	26.4
PointDR [5]	67.3	0.0	4.5	19.9	18.8	2.7	20.0	0.0	62.6	12.9	36.8	43.8	73.3	43.8	56.4	32.2	45.7	28.7	27.4	26.4
DGLSS [21]	72.6	0.1	11.7	29.4	13.7	48.3	0.5	21.2	65.0	20.2	36.5	3.8	78.9	51.8	57.0	36.4	42.7	26.9	34.9	34.6
UniMix [7]	82.7	6.6	8.6	4.5	19.9	35.5	15.1	15.5	55.8	10.2	36.5	40.1	72.8	40.1	49.1	23.5	39.4	23.5	31.5	31.5
DGUIL [22]	77.9	1.0	19.1	26.0	9.7	46.3	0.6	9.3	69.1	9.8	38.6	9.4	73.3	51.2	59.0	31.8	50.8	31.8	22.3	31.4
LiDARWeather [2]	83.1	1.2	17.2	30.5	18.4	47.5	1.07	18.8	64.0	15.9	38.7	4.6	77.4	50.8	59.7	37.2	47.7	31.1	35.8	35.8
No Thing, Nothing [1]	83.3	3.7	31.3	36.2	18.2	53.3	6.8	45.9	67.2	18.1	37.2	5.4	72.1	41.8	58.0	36.0	46.0	39.8	38.9	38.9
Ours	85.0	9.9	23.9	38.7	23.6	46.3	9.7	28.8	65.4	13.1	37.5	1.8	77.8	50.6	65.2	39.4	53.8	32.7	40.7	39.2

Algorithm 1 LPD with Light Geometry-aware Adapter

Require: input cloud P , backbone f , decision module Q , drop ratios B , selective jittering SJ , Light Geometry-aware Adapter $adapter$

Ensure: perturbed cloud P' , geometry-aware cues g , region features $\{F_k\}$

- 1: $\tilde{P} \leftarrow SJ(P)$; $u \leftarrow \text{Uncertainty}(f(\tilde{P}))$
- 2: $g \leftarrow adapter(\tilde{P})$
- 3: $\{(\mathcal{R}_k, F_k)\} \leftarrow \text{BuildRegions}(\tilde{P}, g)$
- 4: $(k^*, b^*) \leftarrow \varepsilon\text{-greedy } \arg \max_{k,b} Q([F_k, u], b)$
- 5: $P' \leftarrow \text{RegionDrop}(\tilde{P}, \mathcal{R}_{k^*}, B[b^*])$
- 6: $r \leftarrow \text{SegLoss}(f(P')) - \text{SegLoss}(f(\tilde{P})) - \lambda \cdot \text{gt_ratio}$
- 7: update Q with $(s = [F_{k^*}, u], a = (k^*, b^*), r, s')$; update f

vulnerabilities at geometric boundaries, corners, and sparse zones. This geometrically-grounded state guides the DQN policy [20] to selectively target regions exhibiting high structural risk, effectively coupling reinforcement learning decisions with concrete geometric evidence.

Overall Workflow: Algorithm 1 details the integrated procedure. Given an SJ-perturbed input \tilde{P} , the adapter extracts geometry-aware cues $\{g_i\}$ and constructs region descriptors $\{(\mathcal{R}_k, F_k)\}$ prior to the LPD phase. Specifically, BUILDREGIONS groups points into candidate regions via a coarse spatial partitioning (e.g., voxelization with size Δ), where each region \mathcal{R}_k is the set of point indices falling into the same cell. For each region, we compute the region-level descriptor F_k by applying a permutation-invariant aggregation (e.g., mean pooling) over point-level cues $\{g_i \mid i \in \mathcal{R}_k\}$, optionally concatenated with simple region statistics such as point count or density. These descriptors are concatenated with the segmentation backbone’s initial predictions and calculated uncertainty to construct the Q-network state matrix. An ε -greedy policy subsequently dictates the target drop regions and ratios.

C. Loss Formulation

Following the established dual-branch paradigm [2], the reference branch operates solely with selective jittering (SJ),

while the drop branch executes the LPD mechanism informed by our Light Geometry-Aware Adapter. Let $i \in \{1, \dots, N\}$ and $c \in \{1, \dots, C\}$ denote point and class indices, respectively. Ground-truth labels are defined as $Y = \{y_i^{(c)}\}$, with predicted probabilities $\hat{Y} = \{\hat{y}_i^{(c)}\}$ and corresponding class weights w_c . The adapter outputs a normalized vulnerability score $s_i \in [0, 1]$ for each point i . We utilize $\kappa \geq 0$ for geometry-aware loss reweighting, alongside $\alpha \geq 1$ and $\eta \geq 0$ to balance the post-drop and entropy regularization terms.

Base Objective: We compute the standard class-weighted cross-entropy loss:

$$\mathcal{L}_{\text{CE}}(\hat{Y}, Y) = -\frac{1}{N} \sum_{i=1}^N \sum_{c=1}^C w_c y_i^{(c)} \log \hat{y}_i^{(c)}. \quad (4)$$

The reference branch loss is designated as $\mathcal{L}_{\text{before}} = \mathcal{L}_{\text{CE}}(\hat{Y}_{\text{sj}}, Y)$.

Geometry-Aware Reweighting After Drop: To penalize structural degradation, the post-drop loss dynamically reweights the contribution of each point by $(1 + \kappa s_i)$:

$$\mathcal{L}_{\text{after}} = \frac{1}{N} \sum_{i=1}^N (1 + \kappa s_i) \left[-\sum_{c=1}^C w_c y_i^{(c)} \log \hat{y}_{i,\text{drop}}^{(c)} \right]. \quad (5)$$

Entropy Regularization: To mitigate network overconfidence and enforce consistency between the dual branches, we incorporate an entropy term. Defining $H(\mathbf{p}) = -\sum_{c=1}^C p^{(c)} \log p^{(c)}$, the regularization loss is formulated as:

$$\mathcal{L}_{\text{ent}} = \frac{1}{2N} \sum_{i=1}^N \left(H(\hat{\mathbf{y}}_{i,\text{sj}}) + H(\hat{\mathbf{y}}_{i,\text{drop}}) \right). \quad (6)$$

Total Optimization Objective: The final loss optimized during training is:

$$\mathcal{L}_{\text{seg}} = \mathcal{L}_{\text{before}} + \alpha \mathcal{L}_{\text{after}} + \eta \mathcal{L}_{\text{ent}}. \quad (7)$$

IV. EXPERIMENTS

A. Setup

All models are evaluated under the rigorous **SemanticKITTI→SemanticSTF source-only** cross-weather

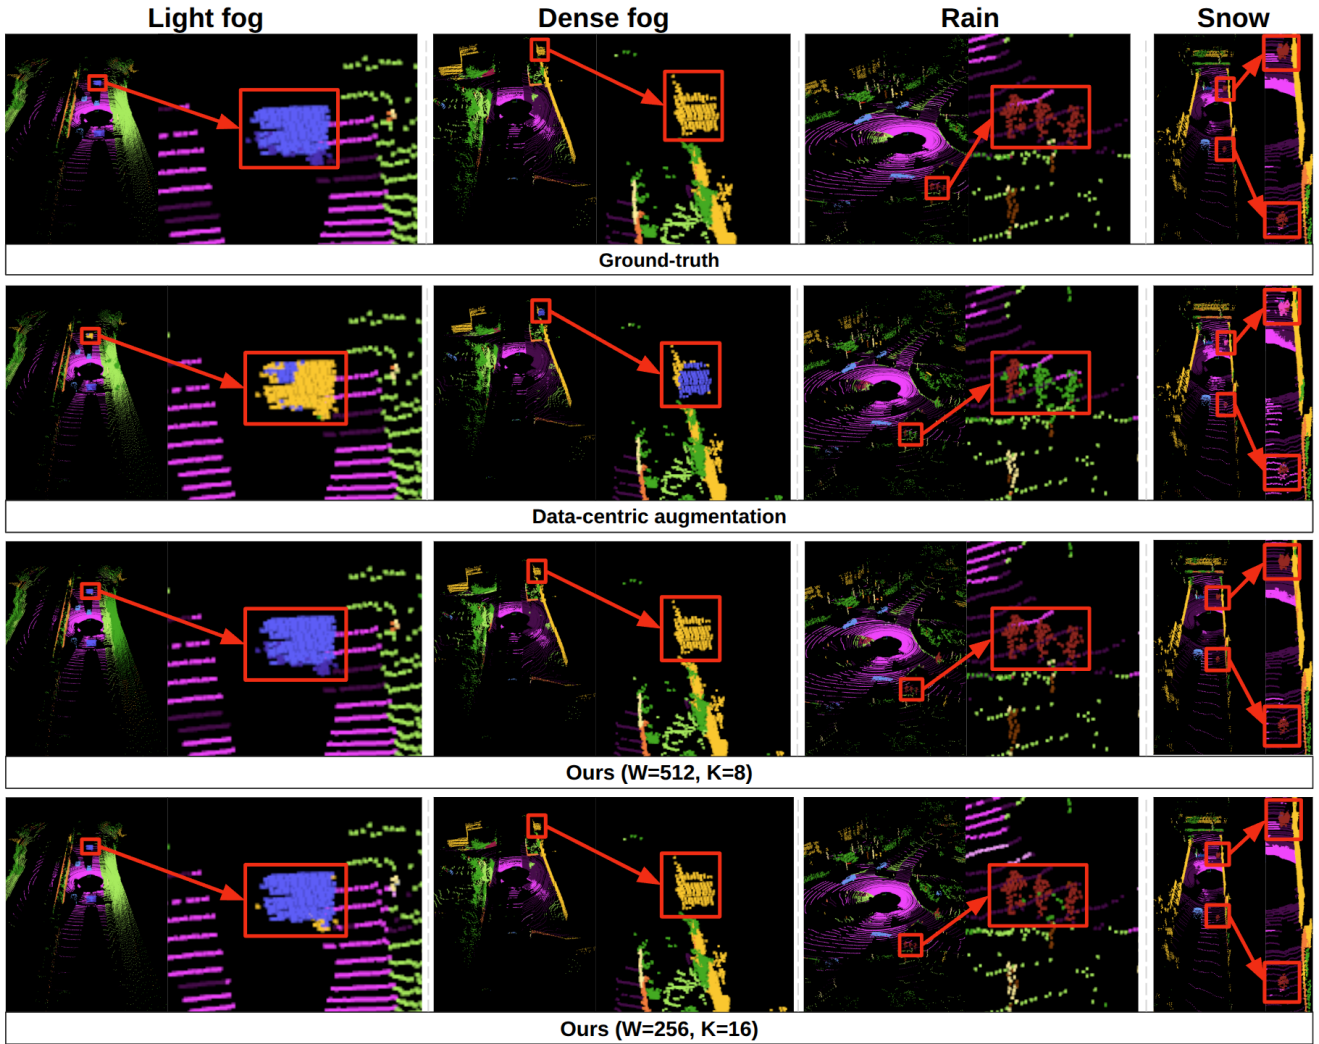


Fig. 5: Qualitative results on SemanticSTF [5] in **light fog**, **dense fog**, **rain**, and **snow**. Top: baseline [2]. Middle/Bottom: ours ($W = 512, K = 8$), ($W = 256, K = 16$). Our adapter preserves continuity and reduces boundary confusion.

TABLE II: Hyperparameter ablation of the Light Geometry-Aware Adapter (window size W , neighbors K) on SemanticKITTI \rightarrow SemanticSTF.

Method	car	bi.cle	mt.cle	truck	bus	pers	mIoU
LiDARWeather [2]	83.1	1.2	17.2	30.5	18.4	47.5	35.8
No Thing, Nothing [1]	<u>83.3</u>	3.7	31.3	<u>36.2</u>	18.2	53.3	38.9
Ours ($W=256, K=8$)	85.6	7.1	28.0	31.5	26.3	51.1	37.8
Ours ($W=256, K=16$)	85.0	9.9	23.9	38.7	23.6	46.3	39.2
Ours ($W=512, K=8$)	84.5	8.4	20.5	41.0	27.3	48.1	38.0
Ours ($W=512, K=16$)	86.0	6.8	19.1	37.7	30.6	49.5	37.1

paradigm. Models are trained entirely on SemanticKITTI [4] and evaluated directly on SemanticSTF [5], strictly forbidding target-domain labels or fine-tuning. We report the mean Intersection over Union (mIoU) and per-class IoU across 19 semantic categories. To ensure statistical reliability given the competitive margins against advanced baselines, we execute our method and the class-centric baseline across three random seeds, observing highly consistent improvements in mIoU and specifically within

TABLE III: Component ablation analysis on SemanticKITTI \rightarrow SemanticSTF. 'C' denotes Circular Padding, 'L' denotes Local-window KNN.

Adapter	Drop	car	bi.cle	mt.cle	truck	bus	pers	mIoU
✓ ×	×	85.3	7.8	17.4	36.4	22.1	51.3	38.0
× ✓	×	85.6	6.5	25.2	37.5	23.7	51.6	37.9
✓ ✓	×	85.8	8.1	27.5	37.2	22.2	49.0	38.3
✓ ×	✓	85.5	9.0	19.7	34.3	27.5	48.2	38.4
× ✓	✓	85.4	10.7	20.6	38.5	23.7	51.6	38.0
× ×	✓	85.2	11.8	18.6	27.4	26.4	56.0	38.0
✓ ✓	✓	85.0	9.9	23.9	38.7	23.6	46.3	39.2

safety-critical classes.

B. Main Results and Comparative Analysis

Table I summarizes the quantitative results. Our full method achieves **39.2** mIoU, outperforming strong data-centric baselines and reaching performance comparable to the advanced class-centric regularization framework [1]. Fig. 5 further shows qualitative results across light fog, dense

TABLE IV: Comparison of methods on the SynLiDAR→SemanticSTF benchmark.

Method	car	bi.cle	mt.cle	truck	oth-v.	pers.	bi.clst	mt.clst	road	parki.	sidew.	othe.g.	build.	fence	veget.	trunk	terra.	pole	traf.	mIoU
Oracle	89.4	42.1	0.0	59.9	61.2	69.6	39.0	0.0	82.2	21.5	58.2	45.6	86.1	63.6	80.2	52.0	77.6	50.1	61.7	54.7
Source-only	27.1	<u>3.0</u>	0.6	15.8	0.1	25.2	1.8	5.6	23.9	0.3	14.6	0.6	36.3	19.9	37.9	17.9	41.8	9.5	2.3	15.0
Dropout [3]	28.0	<u>3.0</u>	1.4	9.6	0.0	17.1	0.8	0.7	34.2	6.8	30.5	1.1	35.5	19.1	42.3	17.6	36.0	14.0	2.8	15.2
Perturbation	27.1	2.3	2.3	16.0	0.1	23.7	1.2	4.0	27.0	3.6	16.2	0.8	29.2	16.7	35.3	18.3	17.9	5.1	2.4	15.2
PolarMix [6]	39.2	1.1	2.2	8.3	1.5	17.8	0.8	0.7	23.3	1.3	17.5	0.4	45.2	24.8	46.2	20.1	38.7	10.9	0.6	15.7
MMD [18]	25.5	2.3	2.1	13.2	0.7	22.1	1.4	7.5	30.8	0.4	17.6	0.4	30.9	19.7	37.6	19.3	<u>43.5</u>	9.9	2.6	15.1
PCL [19]	30.9	0.8	1.4	10.0	0.4	23.3	4.0	7.9	28.5	1.3	17.7	<u>1.2</u>	39.4	18.5	40.0	18.0	38.6	12.1	2.3	15.5
PointDR [5]	37.8	2.5	2.4	<u>23.6</u>	0.1	26.3	2.2	7.7	27.9	7.7	17.5	0.5	47.6	25.3	45.7	21.0	37.5	17.9	5.5	18.5
DGLSS [21]	47.9	2.9	<u>3.4</u>	17.4	1.1	28.0	2.4	7.3	28.8	<u>10.2</u>	18.1	0.2	48.9	25.3	46.5	21.4	45.2	17.9	4.9	19.8
UniMix [7]	65.4	0.1	3.9	16.9	5.3	32.3	2.0	19.3	52.1	5.0	<u>27.3</u>	3.0	49.4	20.3	58.5	22.7	23.2	26.1	20.9	23.4
DGUIL [22]	43.3	2.8	2.6	23.2	<u>3.2</u>	<u>31.3</u>	2.5	4.4	34.3	9.2	17.9	0.3	57.1	<u>27.6</u>	<u>50.0</u>	<u>24.2</u>	41.5	19.0	6.1	21.1
LiDARWeather [2]	39.0	2.5	2.5	22.3	0.3	27.0	1.8	4.0	36.1	10.3	19.0	1.0	50.6	24.5	45.1	23.2	34.1	21.9	7.2	19.6
No Thing, Nothing [1]	48.4	1.5	2.4	19.4	0.2	29.1	<u>3.2</u>	8.9	43.5	6.7	20.5	0.0	52.2	30.1	49.8	20.0	32.9	<u>24.7</u>	<u>7.5</u>	21.1
Ours	<u>49.1</u>	5.0	2.3	26.2	1.8	29.8	1.9	6.4	<u>48.8</u>	8.5	21.6	0.0	<u>55.2</u>	25.5	48.6	24.3	40.1	20.0	6.2	<u>22.1</u>

fog, rain, and snow, where our approach produces sharper boundaries and fewer label inversions in structurally sparse regions.

Mechanisms of the Adapter and Region-Level Drop: The Light Geometry-Aware Adapter stabilizes learning under weather-induced sparsity by preserving *neighbor continuity* and exposing *structural risks* at boundaries and corners, thereby suppressing boundary fragmentation and local class inversions. Region-level Drop complements this by perturbing *structurally coherent regions* instead of isolated points, reducing supervision damage near object perimeters. Together, geometry-conditioned dropping yields improved robustness, particularly for large rigid and safety-critical categories.

Comparison to Class-Centric Regularization: Class-centric regularization directly enforces class-wise priors, providing notable advantages for rare or structurally thin categories that suffer extreme degradation in adverse weather. Our methodology, by contrast, is intrinsically neighborhood-driven. Thus, its efficacy is maximized when localized structural support remains partially intact. When point clouds become severely decimated, class-centric priors may retain an edge; however, our geometry-driven approach provides a fundamentally orthogonal and highly competitive signal, as evidenced by our superior performance across dominant structural categories.

Limitations and Failure Cases: Failures persist when local evidence is insufficient for stable region formation and cue aggregation. (i) *Extreme long-range sparsity:* Dense fog or snow yields few returns on thin structures, making F_k noisy and degrading geometrically thin classes (e.g., poles, traffic signs). (ii) *Severe occlusion/truncation:* Fragmented visibility near depth discontinuities can merge distinct surfaces, inducing region leakage and boundary label bleeding. (iii) *Missing points:* Large spatial voids bias pooling toward dominant background context, and the RL drop may over-perturb fragile regions. Finally, (W, K) trades sparsity tolerance against boundary sharpness; the full adapter is generally the most reliable. Future work will explore uncertainty-adaptive neighborhoods and occlusion-aware region partitioning to mitigate these cases.

C. Ablation and Generalization

Ablation. Table II studies the hyperparameters (W, K) . The adapter consistently improves over the data-centric baseline, with the best setting reaching **39.2** mIoU. Larger K helps aggregate large rigid objects but can smooth fine boundaries, while larger W increases candidates under sparsity yet may introduce irrelevant neighbors when structure is fragmented. Table III isolates components and their interaction with Region-level Drop. Circular padding (C) stabilizes continuity at the azimuth seam, and local-window KNN (L) suppresses spurious long-range matches. Their combination yields the most reliable cues, and Region-level Drop performs best with the full adapter by forming coherent regions without fragmenting boundaries.

Generalization. To test robustness beyond a single source dataset, we additionally evaluate SynLiDAR→SemanticSTF. As shown in Table IV, our method remains competitive under the synthetic-to-real shift, achieving **22.1** mIoU. Importantly, the qualitative trend is consistent with the main transfer: geometry-aware cues primarily help classes where local continuity and region consistency are informative (e.g., large rigid and safety-critical categories), while extremely sparse or occluded structures remain challenging. This cross-source result supports that the proposed geometry-aware cues and region-level dropping are not tied to a specific source domain and can preserve robustness under substantial distribution changes.

V. CONCLUSIONS

In this paper, we introduced the **Light Geometry-Aware Adapter**, a novel mechanism that drastically enhances region-level point dropping for LiDAR semantic segmentation in adverse weather. Evaluated on the rigorous SemanticKITTI→SemanticSTF benchmark, our approach boosts the mIoU by **+3.4** over robust data-centric baselines, achieving performance on par with advanced class-centric regularization models. The adapter is fundamentally lightweight, highly computationally efficient, and strictly plug-and-play during training. Future research will focus on integrating these geometric abstractions with dedicated, physics-based adverse-weather simulators to further bridge the sim-to-real domain gap.

REFERENCES

- [1] J. Park, K. Kim, Y. Kim, J. Lee, and H. Shim, "No Thing, Nothing: Highlighting Safety-Critical Classes for Robust LiDAR Semantic Segmentation in Adverse Weather," in *Proc. CVPR*, 2025.
- [2] J. Park, K. Kim, and H. Shim, "Rethinking Data Augmentation for Robust LiDAR Semantic Segmentation in Adverse Weather," in *Proc. ECCV*, 2024.
- [3] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov, "Dropout: A Simple Way to Prevent Neural Networks from Overfitting," *J. Mach. Learn. Res. (JMLR)*, vol. 15, pp. 1929–1958, 2014.
- [4] J. Behley, M. Garbade, A. Milioto, J. Quenzel, S. Behnke, C. Stachniss, and J. Gall, "SemanticKITTI: A Dataset for Semantic Scene Understanding of LiDAR Sequences," in *Proc. ICCV*, 2019.
- [5] A. Xiao, J. Huang, W. Xuan, R. Ren, K. Liu, D. Guan, A. El Saddik, S. Lu, and E. P. Xing, "PointDR: 3D Semantic Segmentation in the Wild via Domain Randomization," in *Proc. CVPR*, 2023, pp. 9382–9392.
- [6] A. Xiao, J. Gu, Y. Zhang, S. Chen, and Y. Cui, "PolarMix: A General Data Augmentation Technique for LiDAR Point Clouds," in *NeurIPS*, 2022, pp. 11035–11048.
- [7] H. Zhao, H. Qiu, Y. Yao, Y. Wang, S. Wang, and Z. Liu, "UniMix: Towards Domain Adaptive and Generalizable LiDAR Semantic Segmentation in Adverse Weather," in *Proc. CVPR*, 2024.
- [8] X. Zhao, Y. Chen, H. Zhao, Y. Qiao, and D. Lin, "TripleMixer: A 3D Point Clouds Denoising Model for Adverse Weather," *arXiv:2408.13802*, 2024.
- [9] H. Tang, Z. Liu, S. Zhao, Y. Lin, J. Lin, H. Su, and S. Han, "Searching Efficient 3D Architectures with Sparse Point-Voxel Convolution," in *Proc. ECCV*, 2020.
- [10] H. Zhou, Y. Zhu, X. Cheng, D. Lu, X. Yang, and R. Yang, "Cylinder3D: An Effective 3D Framework for Driving-Scene LiDAR Semantic Segmentation," *arXiv:2008.01550*, 2020.
- [11] X. Lai, Y. Chen, S. Jiang, X. Liu, S. Pu, and H. Xiong, "Spherical Transformer for LiDAR-Based 3D Recognition," in *Proc. CVPR*, 2023.
- [12] B. Wu, A. Wan, X. Yue, and K. Keutzer, "SqueezeSeg: Convolutional Neural Nets with Recurrent CRF for Real-Time Road-Object Segmentation from 3D LiDAR Point Clouds," in *Proc. ICRA*, 2018.
- [13] A. Milioto, I. Vizzo, J. Behley, and C. Stachniss, "RangeNet++: Fast and Accurate LiDAR Semantic Segmentation," in *Proc. IROS*, 2019.
- [14] C. R. Qi, H. Su, K. Mo, and L. J. Guibas, "PointNet: Deep Learning on Point Sets for 3D Classification and Segmentation," in *Proc. CVPR*, 2017.
- [15] C. R. Qi, L. Yi, H. Su, and L. J. Guibas, "PointNet++: Deep Hierarchical Feature Learning on Point Sets in a Metric Space," in *NeurIPS*, 2017.
- [16] H. Thomas, C. R. Qi, J.-E. Deschaud, B. Marcotegui, F. Goulette, and L. J. Guibas, "KPConv: Flexible and Deformable Convolution for Point Clouds," in *Proc. ICCV*, 2019.
- [17] Q. Hu, B. Yang, L. Xie, S. Rosa, Y. Guo, Z. Wang, N. Trigoni, and A. Markham, "RandLA-Net: Efficient Semantic Segmentation of Large-Scale Point Clouds," in *Proc. CVPR*, 2020.
- [18] H. Li, S. J. Pan, S. Wang, and A. C. Kot, "Domain Generalization with Adversarial Feature Learning," in *Proc. CVPR*, 2018, pp. 5400–5409.
- [19] X. Yao, Y. Bai, X. Zhang, Y. Zhang, Q. Sun, R. Chen, R. Li, and B. Yu, "PCL: Proxy-based Contrastive Learning for Domain Generalization," in *Proc. CVPR*, 2022, pp. 7097–7107.
- [20] V. Mnih, K. Kavukcuoglu, D. Silver, *et al.*, "Playing Atari with Deep Reinforcement Learning," *arXiv:1312.5602*, 2013.
- [21] H. Kim, Y. Kang, C. Oh, and K.-J. Yoon, "Single-Domain Generalization for LiDAR Semantic Segmentation," in *Proc. CVPR*, 2023, pp. 17587–17598.
- [22] P. He, L. Jiao, L. Li, X. Liu, F. Liu, W. Ma, S. Yang, and R. Shang, "Domain Generalization-aware Uncertainty Introspective Learning for 3D Point Cloud Segmentation," in *Proc. ACM MM*, 2024, pp. 651–660.