The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences, Volume XLIX-B2-2026
XXV ISPRS Congress 2026 “From Imagery to Understanding”, Commission II, 4–11 July 2026, Toronto, Canada

GDC: Geometric Diffusion Consistency for Weather-Robust 3D Point Cloud Segmentation

Jing Du[1], John Zelek[2], Michael A. Chapman[3], Jonathan Li[4*]

[1] Department of Systems Design Engineering, University of Waterloo, Waterloo, ON N2L 3G1, Canada - j7du@uwaterloo.ca
[2] Department of Systems Design Engineering, University of Waterloo, Waterloo, ON N2L 3G1, Canada - jzelek@uwaterloo.ca
[3] Department of Civil Engineering, Toronto Metropolitan University, Toronto, ON M5B 2K3, Canada - mchapman@torontomu.ca [4]Department of Geography and Environmental Management, University of Waterloo, Waterloo, ON N2L 3G1, Canada – junli@uwaterloo.ca

Keywords: 3D Point Cloud Segmentation, Domain Generalization, Adverse Weather, Geometric Diffusion Consistency, LiDAR

Abstract

Semantic segmentation of outdoor 3D point clouds degrades significantly under adverse weather, as rain, fog, and snow corrupt the geometric structure of LiDAR returns through backscatter insertion, range-dependent attenuation, and volumetric scattering. Existing domain generalization methods constrain feature values directly, which becomes less effective when weather-induced perturbations alter the local neighborhood topology that underlies feature aggregation. This work proposes Geometric Diffusion Consistency (GDC), a training-time regularizer that enforces consistent feature propagation behavior across geometrically divergent views of the same point cloud. A dual-view augmentation pipeline generates training pairs through weak and strong perturbations, where the strong branch incorporates dual-mode atmospheric extinction modeling, semantic-aware geometric corruption, and weather-coordinated structural perturbation. A lightweight learnable diffusion operator, implemented via sparse convolutions with a gated residual connection, propagates encoder bottleneck features through local voxel neighborhoods. The consistency loss aligns diffused representations at corresponding points across views, preserving topological relationships essential for dense prediction while allowing feature values to adapt to altered geometry. On the SemanticKITTI to SemanticSTF domain generalization benchmark, GDC achieves 38.6% mIoU, exceeding the previous best method by 3.8%, with consistent improvements across dense fog, light fog, rain, and snow conditions.

1. Introduction

Semantic segmentation of 3D point clouds is fundamental to autonomous driving and urban scene understanding (Tang et al., 2020; Zhu et al., 2021; Du et al., 2021; Lu et al., 2024; Du et al., 2025a,b). LiDAR sensors capture precise geometric representations through active illumination, independent of ambient lighting. Deployment under adverse weather, however, reveals a critical challenge: atmospheric conditions alter point cloud geometry at the physical level. Rain generates spurious backscatter returns while occluding genuine targets. Fog attenuates returns following Beer-Lambert extinction, causing progressive density loss at increasing range (Hahner et al., 2021). Snow combines ground-level clutter with volumetric scattering that corrupts measurements across the scan (Hahner et al., 2022). Unlike 2D imagery where weather primarily degrades photometric properties, adverse weather in 3D point clouds directly modifies the local neighborhood topology that segmentation networks rely on for spatial reasoning.

Data augmentation methods address this challenge by simulating weather corruption during training. PolarMix (Xiao et al., 2022) swaps point cloud sectors in polar coordinates across training samples to increase geometric diversity. Weather simulators generate synthetic fog and rain patterns to expose models to realistic corruption (Park et al., 2024). These approaches improve robustness to corruption patterns seen during training, but their effectiveness depends on how closely simulated perturbations match real conditions. A complementary direction is to impose explicit constraints on how the network propagates features through geometrically degraded neighborhoods.

Domain generalization methods impose feature-space constraints to improve transferability in 3D point cloud segmentation. Xiao et al. (2023) proposed PointDR, introducing the SemanticSTF benchmark and a geometry-style randomization technique that alternately randomizes point cloud geometry and aggregates encoded embeddings to learn generalizable representations. Du et al. (2024) introduced WADG, incorporating adaptive feature normalization, dual-attention fusion for cross-domain feature integration, and proxy label generation to improve segmentation robustness across weather conditions. Zhao et al. (2024) proposed UniMix, constructing a bridge domain through weather simulation and applying a universal mixing operator across spatial, intensity, and semantic distributions within a teacher-student framework. These methods constrain feature values directly, which becomes less effective when adverse weather fragments neighborhoods through point deletion, displacement, and insertion, as the geometric context underlying those features has itself changed.

A central observation underlying this work is that the primary effect of adverse weather on point clouds is structural: it alters the composition and connectivity of local neighborhoods used for feature aggregation. Under such conditions, a more robust invariant exists at the level of feature propagation behavior. The pattern by which information flows through spatial neighborhoods can remain consistent even when specific neighborhood compositions differ. If the propagation process is stable, the network can recover semantic information from surviving local connections regardless of which points are present or absent. This motivates a consistency constraint on the diffusion process through which features spread across neighborhoods, preserving topological relationships essential for dense prediction while allowing feature values to adapt to altered geometry.

This work introduces Geometric Diffusion Consistency (GDC),

* Corresponding author

a training-time regularizer that enforces invariant feature diffusion across stochastically augmented views of the same point cloud. Given two views generated through weak and strong augmentation of each input scan, a shared encoder extracts bottleneck features from both. A lightweight learnable diffusion operator propagates these features through local sparse neighborhoods, and a consistency loss aligns the diffused representations at corresponding points across views. The strong augmentation incorporates atmospheric degradation, semantic-aware geometric corruption, and weather-coordinated structural perturbation to expose the encoder to diverse and realistic weather effects. Point-level correspondence is maintained throughout augmentation via unique identifiers, and the diffusion operator is active only during training, introducing no additional cost at inference time.

The contributions of this work are summarized as follows:

- A structural regularization perspective is proposed for weather-robust 3D point cloud segmentation, enforcing consistency in feature propagation topology across geometric perturbations.
- To generate realistic and diverse training views, a dual-view augmentation pipeline is developed, incorporating dual-mode atmospheric extinction estimation, semantic-aware geometric corruption, and cross-stage weather coordination.
- The geometric diffusion operator, implemented as a lightweight sparse convolution network with a gated residual connection, is introduced to regularize spatial information flow through point-wise consistency constraints on encoder bottleneck features.

2. Method

GDC trains a segmentation network on source-domain point clouds by enforcing consistent feature diffusion across two stochastically augmented views of each scan. Algorithm 1 summarizes the training procedure: each input undergoes weak and strong augmentation to produce two views, a shared encoder extracts features from both, a learnable diffusion operator propagates features through sparse neighborhoods, and a consistency loss aligns the diffused representations at corresponding points.

2.1 Problem Formulation

Let $\mathcal{P} = \{p_i\}_{i=1}^{N}$ denote an input point cloud where each point $p_i = (\mathbf{x}_i, \iota_i) \in \mathbb{R}^4$ consists of spatial coordinates $\mathbf{x}_i \in \mathbb{R}^3$ and intensity $\iota_i \in [0, 1]$. A segmentation network f_θ maps $\mathcal{P}$ to per-point predictions over K classes, producing intermediate features $\mathbf{F} \in \mathbb{R}^{N\times C}$ with channel dimension C. We consider the domain generalization setting: f_θ is trained on a source dataset $\mathcal{D}_s$ captured under clear weather and evaluated on a target dataset $\mathcal{D}_t$ collected under adverse conditions (rain, snow, fog), with no access to $\mathcal{D}_t$ during training. The supervised objective minimizes the expected cross-entropy:

$$\mathcal{L}_{\text{seg}} = \mathbb{E}_{(\mathcal{P},\mathbf{Y})\sim\mathcal{D}_s}\left[-\frac{1}{|\mathcal{V}|}\sum_{i\in\mathcal{V}}\sum_{k=1}^{K} y_{i,k}\log\hat{y}_{i,k}\right] \tag{1}$$

where $\mathbf{Y} \in \{1, \ldots, K\}^N$ denotes per-point labels, $y_{i,k}$ is the one-hot encoding of the ground-truth label for point i and class k, $\hat{y}_{i,k}$ is the predicted probability for class k at point i, and $\mathcal{V} = \{i : y_i \neq y_{\text{ign}}\}$ excludes the ignore label.

Algorithm 1 GDC Training Procedure

Require: Source dataset $\mathcal{D}_s$, encoder f_θ, diffusion operator Φ_ω, prototype bank $\{\mathbf{q}_k\}_{k=1}^{K}$
Ensure: Trained encoder f_θ
for each mini-batch $\{(\mathcal{P}, \mathbf{Y})\} \subset \mathcal{D}_s$ **do**
  // Dual-view generation with point-ID tracking
  Assign unique identifiers $\{\text{id}_i\}$ to all points in $\mathcal{P}$
  $\mathcal{P}^{(1)} \leftarrow$ weak augmentation (Eq. 2)
  $\mathcal{P}^{(2)} \leftarrow (\mathcal{T}_{\text{clutter}} \circ \mathcal{T}_{\text{refl}} \circ \mathcal{T}_{\text{atm}} \circ \mathcal{T}_{\text{geom}})(\mathcal{P})$ (Eq. 4)
  Voxelize and subsample both views; propagate id through indexing
  // Forward pass through shared encoder
  for $v \in \{1, 2\}$ **do**
    $\hat{\mathbf{Y}}^{(v)} \leftarrow f_\theta(\mathcal{P}^{(v)})$ per-point logits
    $\mathbf{F}_{\text{bot}}^{(v)} \leftarrow$ encoder bottleneck features (stage 4)
  end for
  // Segmentation loss (View 1)
  $\mathcal{L}_{\text{seg}}^{(1)} \leftarrow \text{CrossEntropy}(\hat{\mathbf{Y}}^{(1)}, \mathbf{Y}^{(1)})$ (Eq. 21)
  // Prototype contrastive loss (View 2)
  $\mathcal{L}_{\text{con}}^{(2)} \leftarrow$ prototype contrastive loss on View 2
  // Geometric Diffusion Consistency
  $\hat{\mathbf{F}}^{(v)} \leftarrow \Phi_\omega(\mathbf{F}_{\text{bot}}^{(v)})$ for $v \in \{1, 2\}$ (Eqs. 15–18)
  $\Omega \leftarrow \{(i, j) : \text{id}_i^{(1)} = \text{id}_j^{(2)},\ \text{id} \geq 0\}$ (Eq. 19)
  $\mathcal{L}_{\text{GDC}} \leftarrow \frac{1}{|\Omega|}\sum_{(i,j)\in\Omega}\|\hat{\mathbf{F}}_i^{(1)} - \hat{\mathbf{F}}_j^{(2)}\|_2^2$ (Eq. 20)
  // Total loss and parameter update
  $\mathcal{L} \leftarrow \mathcal{L}_{\text{seg}}^{(1)} + \lambda_{\text{con}}\,\mathcal{L}_{\text{con}}^{(2)} + \lambda_{\text{GDC}}\,\mathcal{L}_{\text{GDC}}$ (Eq. 22)
  Update θ and ω via $\nabla_{\theta,\omega}\,\mathcal{L}$
end for

Adverse weather corrupts point clouds through deletion, displacement, and insertion of returns. We model this as a stochastic operator $\mathcal{A} : \mathbb{R}^{N\times 4} \to \mathbb{R}^{M\times 4}$ where $M \neq N$ in general. For a standard neighborhood aggregation, the operator $\mathcal{A}$ alters both the neighborhood structure $\mathcal{N}(i)$ and the aggregation weights. GDC regularizes feature propagation to remain consistent under such structural perturbations.

2.2 Dual-View Generation

We assign a unique identifier $\text{id}_i \in \mathbb{Z}$ to each point before augmentation, enabling correspondence tracking. The weak view $\mathcal{P}^{(1)}$ undergoes a random rigid transformation with mild scaling and translational noise:

$$\mathbf{x}_i^{(1)} = s^{(1)} \cdot \mathbf{R}(\theta^{(1)})\,\mathbf{x}_i + \boldsymbol{\epsilon}_i \tag{2}$$

where $\mathbf{R}(\theta)$ is the yaw rotation matrix

$$\mathbf{R}(\theta) = \begin{bmatrix} \cos\theta & \sin\theta & 0 \\ -\sin\theta & \cos\theta & 0 \\ 0 & 0 & 1 \end{bmatrix}, \quad \theta^{(1)} \sim \mathcal{U}(0, 2\pi) \tag{3}$$

with uniformly sampled scale $s^{(1)}$ and Gaussian translational noise $\boldsymbol{\epsilon}_i$. An optional intensity style transfer models sensor calibration variation by applying random gain, gamma correction, additive bias, and per-point noise, all clamped to $[0, 1]$.

The strong view $\mathcal{P}^{(2)}$ is constructed through a composition of four operators:

$$\mathcal{P}^{(2)} = \big(\mathcal{T}_{\text{clutter}} \circ \mathcal{T}_{\text{refl}} \circ \mathcal{T}_{\text{atm}} \circ \mathcal{T}_{\text{geom}}\big)(\mathcal{P}) \tag{4}$$

where $\circ$ denotes function composition and $\mathcal{T}_{\text{geom}}$, $\mathcal{T}_{\text{atm}}$, $\mathcal{T}_{\text{refl}}$, and $\mathcal{T}_{\text{clutter}}$ denote geometric perturbation, atmospheric degradation, mirror-plane reflection with occlusion, and sensor clutter injection, respectively. Each operator propagates the point identifiers, assigning id $= -1$ to all synthetically generated points so that they are excluded from both the supervised loss and the consistency loss.

2.3 Semantic-Aware Geometric Perturbation

The geometric operator $\mathcal{T}_{\text{geom}}$ applies a rigid transform with wider parameter ranges than the weak view, random axis flipping, and coordinate jittering. Four structural corruption stages then modify the point set: class-aware point dropout, azimuthal sector occlusion, elevation-bin dropout, and range-adaptive thinning.

All four stages share a unified probabilistic dropout framework. Given per-point dropout probabilities $\{p_i^{\text{drop}}\}$ from any corruption source, the framework applies: (i) semantic-aware rescaling, where points from rare categories $\mathcal{C}_{\text{rare}}$ have their dropout probability reduced by a stage-specific protection factor γ_k; (ii) Bernoulli sampling; and (iii) a minimum-count guarantee that restores randomly selected dropped points if too few survive. The combined dropout probability is:

$$\tilde{p}_i^{\text{drop}} = \begin{cases} \gamma_k \cdot p_i^{\text{drop}} & \text{if } y_i \in \mathcal{C}_{\text{rare}} \\ \text{clamp}(p_i^{\text{drop}},\ 0,\ 1) & \text{otherwise} \end{cases} \tag{5}$$

where $\mathcal{C}_{\text{rare}}$ includes underrepresented classes such as motorcycle, person, traffic-sign, pole, and bicyclist.

Class-aware point dropout sets a spatially uniform dropout rate. Azimuthal sector occlusion removes a continuous angular wedge centered at a random azimuth. Elevation-bin dropout discretizes elevations into pseudo-ring bins and removes a random subset. Range-adaptive thinning applies a continuous power-law model:

$$p_{\text{thin}}(r_i) = p_{\max} \cdot \left(\frac{\|\mathbf{x}_i\|_2}{r_{\text{sensor}}} \right)^{\gamma_d} \tag{6}$$

where $p_{\max}$ is the maximum thinning probability, r_{sensor} is the sensor range, and γ_d controls distance sensitivity. The four stages are applied sequentially with independent Bernoulli samples.

2.4 Atmospheric Degradation with Dual-Mode Extinction

The atmospheric operator $\mathcal{T}_{\text{atm}}$ is activated stochastically. A weather type $w \in \mathcal{W} = \{\text{rain, snow, light fog, dense fog}\}$ and an intensity level $\ell \in \{\text{light, medium, heavy}\}$ are sampled from categorical distributions. The operator applies intensity attenuation, range-dependent point loss, and backscatter return replacement in sequence.

The extinction coefficient β_{ext} is estimated through two independent paths and blended probabilistically. The first path derives β_{ext} from a sampled visibility distance V_w via the Koschmieder relation:

$$\beta_{\text{vis}} = \frac{c_k}{V_w} \cdot \xi, \quad V_w \sim \mathcal{U}(V_{\min}^{(w,\ell)}, V_{\max}^{(w,\ell)}) \tag{7}$$

where c_k is the Koschmieder constant and ξ is a multiplicative jitter sampled from a narrow uniform range. Visibility ranges are weather- and level-specific. The second path uses a precomputed Mie-scattering lookup table $\beta_{\text{lut}}^{(w,\ell)}$, computed offline from Gamma-distributed particle size distributions for each weather-level combination. At training time, one of two modes is selected: in the visibility-driven mode, $\beta_{\text{ext}} = \beta_{\text{vis}}$; in the lookup-table mode, the two estimates are blended:

$$\beta_{\text{ext}} = (1-\mu)\, \beta_{\text{vis}} + \mu\, \beta_{\text{lut}}^{(w,\ell)} \cdot s_{\text{lut}} \tag{8}$$

where μ is the blend weight, s_{lut} is a global scaling factor, and β_{ext} is clamped to a valid extinction range. This dual-mode design covers a wider range of extinction profiles: the visibility path produces smooth monotonic attenuation, while the lookup-table path captures particle-size-dependent spectral structure. In the lookup-table mode, the anisotropy parameter g of the Henyey-Greenstein phase function is additionally perturbed, and an azimuthal inhomogeneity term models spatial non-uniformity of particle concentration.

The per-point transmittance and attenuated intensity are computed as:

$$\tau_i = \exp\!\big(-2\, \beta_{\text{ext}}\, r_i\big) \cdot \big(1 + \eta\, \sin(2\phi_i + \psi)\big) \tag{9}$$

$$\iota_i' = \text{clamp}\big(\iota_i \cdot \tau_i,\ 0,\ 1\big) \tag{10}$$

where $r_i = \|\mathbf{x}_i\|_2$, $\phi_i = \text{atan2}(x_{i,2}, x_{i,1})$, η is the inhomogeneity amplitude, and ψ is a random phase offset.

Each point survives with a range-dependent probability that incorporates a second, distinct set of semantic protections. While the geometric perturbation stage (Section 2.3) protects *rare* classes $\mathcal{C}_{\text{rare}}$ from being over-dropped, the atmospheric stage protects *structurally dominant* classes $\mathcal{C}_{\text{prot}}$ (building, road, vegetation, car) from being erased by heavy attenuation. The survival probability is:

$$p_{\text{keep}}(r_i, y_i) = \begin{cases} \min\!\big(p_{\text{base}} \cdot p_{\text{cap}} + \rho_p,\ 1\big) & \text{if } y_i \in \mathcal{C}_{\text{prot}} \\ \text{clamp}\!\big(p_{\text{base}} \cdot p_{\text{cap}},\ 0.05,\ 1\big) & \text{otherwise} \end{cases} \tag{11}$$

where ρ_p is the rescue boost for protected classes. The base survival follows $p_{\text{base}}(r_i) = \exp(-d_w \cdot (r_i/r_{\text{sensor}})^{\gamma_r})$ with weather-specific drop strength d_w and distance exponent γ_r. For fog conditions, a logistic visibility ceiling is applied as $p_{\text{cap}}(r_i) = \sigma((r_{\text{cap}} - r_i)/\delta)$, where $\sigma(\cdot)$ denotes the sigmoid function, r_{cap} is a sampled visibility limit, and δ controls the transition smoothness; for rain and snow, $p_{\text{cap}} = 1$.

A fraction of surviving returns is replaced by backscattered soft targets. For each selected point at range r_{obj}, a scatter distance is sampled via inverse-CDF from a truncated exponential:

$$r_s = -\frac{1}{\beta_{\text{back}}} \ln\!\Big(1 - u \cdot \big(1 - e^{-\beta_{\text{back}}\, r_{\text{obj}}}\big)\Big), \quad u \sim \mathcal{U}(0,1) \tag{12}$$

where $\beta_{\text{back}} = \beta_{\text{ext}} \cdot \rho_b$ with weather-specific backscatter ratio ρ_b. The coordinate is rescaled along its original scan direction as $\mathbf{x}_s = \mathbf{x}_i \cdot r_s / r_{\text{obj}}$, and the scattered intensity follows the Henyey-Greenstein phase function evaluated at the backscatter angle:

$$P_{\text{HG}}(g) = \frac{1 - g^2}{4\pi\, (1 + g^2 + 2g)^{3/2}}$$

$$\iota_s = G_b \cdot \beta_{\text{back}} \cdot P_{\text{HG}}(g) \cdot e^{-2\, \beta_{\text{ext}}\, r_s} \tag{13}$$

where g is the anisotropy parameter, G_b is a backscatter gain, and ι_s is clamped to a valid intensity range. Sparse volumetric particle returns are injected along existing scan directions with near-field-concentrated range sampling ($\text{Beta}(1.5, 4.0)$ distribution) and extinction-consistent intensity attenuation.

2.5 Weather-Coordinated Structural Corruption

The remaining two operators, $\mathcal{T}_{\text{refl}}$ and $\mathcal{T}_{\text{clutter}}$, are conditioned on the weather type w selected by $\mathcal{T}_{\text{atm}}$ (or $w = \text{none}$ if atmospheric degradation was not activated). This cross-stage coordination prevents redundant corruption: when heavy attenuation already introduces substantial geometric noise, structural operators reduce their intensity.

The reflection operator $\mathcal{T}_{\text{refl}}$ simulates multipath artifacts through planar reflection and aligned occlusion. Its activation probability and ghost point ratio are both scaled by a weather-dependent factor λ_w, which is highest for rain (where wet-surface reflections are prominent) and lowest for dense fog (where most returns are already attenuated). A reflection plane is sampled as a ground plane, x-wall, or y-wall, estimated from the point distribution statistics. A subset of foreground points from classes $\mathcal{C}_{\text{fg}}$ (car, truck, person, bicyclist, motorcyclist, pole, traffic-sign) is reflected:

$$\tilde{\mathbf{x}}_i = \mathbf{x}_i - 2\left((\mathbf{x}_i - \mathbf{c}) \cdot \hat{\mathbf{n}}\right)\hat{\mathbf{n}} + \boldsymbol{\epsilon}_{\text{ref}} \tag{14}$$

where $\hat{\mathbf{n}}$ is the plane normal, $\mathbf{c}$ lies on the plane, and $\boldsymbol{\epsilon}_{\text{ref}}$ is a translational jitter. A mild normal-axis scaling is applied to avoid perfect symmetry. Reflected points are clipped to the sensor field of view, and an aligned occlusion band removes original points near the reflection plane.

The sensor clutter operator $\mathcal{T}_{\text{clutter}}$ is applied last, independently of weather type, to model sensor-level domain mismatch unrelated to atmospheric conditions. A small fraction of points is injected along existing scan directions at short range with random intensity.

2.6 Geometric Diffusion Operator

The augmentation pipeline produces two views with divergent geometric structures but shared semantic content. GDC regularizes the encoder by requiring a learnable diffusion operator to produce consistent outputs at corresponding points across views. The operator acts on the encoder bottleneck features (stage 4, $C = 256$), where spatial resolution is lowest and semantic abstraction is highest. We define $\Phi_\omega : \mathbb{R}^{N\times C} \to \mathbb{R}^{N\times C}$ as a two-layer sparse convolution network that propagates features through local voxel neighborhoods:

$$\mathbf{h}_i^{(1)} = \sum_{j\in\mathcal{N}_s(i)} \mathbf{W}_{\Delta_{ij}}^{(1)}\,\mathbf{f}_j + \mathbf{b}^{(1)}, \quad \Delta_{ij} = c_j - c_i \tag{15}$$

$$\tilde{\mathbf{h}}_i^{(1)} = \mathrm{ReLU}\big(\mathrm{BN}_1(\mathbf{h}_i^{(1)})\big) \tag{16}$$

$$\mathbf{h}_i^{(2)} = \mathrm{BN}_2\left(\sum_{j\in\mathcal{N}_s(i)} \mathbf{W}_{\Delta_{ij}}^{(2)}\,\tilde{\mathbf{h}}_j^{(1)} + \mathbf{b}^{(2)}\right) \tag{17}$$

where $\mathcal{N}_s(i) = \{j : \|c_j - c_i\|_\infty \le \lfloor\kappa/2\rfloor\}$ is the sparse voxel neighborhood with kernel size $\kappa = 3$, c_i denotes the voxel coordinate of point i, $\mathbf{f}_j \in \mathbb{R}^C$ is the feature vector of point j (the j-th row of $\mathbf{F}$), $\mathbf{W}_{\Delta_{ij}}^{(1)}, \mathbf{W}_{\Delta_{ij}}^{(2)} \in \mathbb{R}^{C\times C}$ are learnable kernel weights for the first and second convolution layers indexed by offset Δ_{ij}, $\mathbf{b}^{(1)}, \mathbf{b}^{(2)}$ are bias vectors, and BN_1, BN_2 denote batch normalization layers. A sigmoid-gated residual connection blends the original and diffused features:

$$\alpha = \sigma(w_\alpha), \quad \Phi_\omega(\mathbf{F})_i = \alpha\,\mathbf{f}_i + (1-\alpha)\,\mathbf{h}_i^{(2)} \tag{18}$$

where $w_\alpha \in \mathbb{R}$ is a learnable scalar, biasing the output toward the original features at the start of training. Φ_ω operates on the encoder bottleneck features and is active only during training.

2.7 Consistency Regularization

We construct a correspondence set Ω by matching point identifiers across views:

$$\Omega = \left\{(i,j) : \mathrm{id}_i^{(1)} = \mathrm{id}_j^{(2)},\ \mathrm{id}_i^{(1)} \ge 0\right\} \tag{19}$$

where id ≥ 0 excludes all synthetically injected points. The GDC loss penalizes inconsistency in diffused bottleneck representations at corresponding locations:

$$\mathcal{L}_{\text{GDC}} = \frac{1}{|\Omega|}\sum_{(i,j)\in\Omega}\left\|\Phi_\omega(\mathbf{F}^{(1)})_i - \Phi_\omega(\mathbf{F}^{(2)})_j\right\|_2^2 \tag{20}$$

where $\mathbf{F}^{(1)}$ and $\mathbf{F}^{(2)}$ are the encoder bottleneck features from View 1 and View 2, respectively. Gradients flow through both branches symmetrically, allowing the diffusion operator to jointly adapt its propagation behavior to both the clean and corrupted geometry.

2.8 Training Objective

The segmentation loss on View 1 uses standard cross-entropy:

$$\mathcal{L}_{\text{seg}}^{(1)} = -\frac{1}{|\mathcal{V}^{(1)}|}\sum_{i\in\mathcal{V}^{(1)}}\sum_{k=1}^{K} y_{i,k}\,\log \hat{y}_{i,k}^{(1)} \tag{21}$$

where $\mathcal{V}^{(v)} = \{i : y_i^{(v)} \neq y_{\text{ign}}\}$ denotes the valid-label point set of View v. A prototype-based contrastive loss $\mathcal{L}_{\text{con}}^{(2)}$ on View 2 reinforces class-discriminative features by matching ℓ_2-normalized projected features against momentum-updated per-class prototypes. The complete training objective is:

$$\mathcal{L} = \mathcal{L}_{\text{seg}}^{(1)} + \lambda_{\text{con}}\,\mathcal{L}_{\text{con}}^{(2)} + \lambda_{\text{GDC}}\,\mathcal{L}_{\text{GDC}} \tag{22}$$

where λ_{con} and λ_{GDC} are loss weights. Φ_ω and the prototype memory bank are used only during training.

3. Experiments

We evaluate GDC on the domain generalization benchmark from SemanticKITTI (Behley et al., 2019) to SemanticSTF (Xiao et al., 2023). SemanticKITTI provides 43,552 outdoor LiDAR scans captured under clear weather with 19-class semantic annotations. SemanticSTF contains 2,076 scans collected under four real-world adverse conditions (dense fog, light fog, rain, and snow) with the same label taxonomy. Following the standard protocol, we train on SemanticKITTI and directly evaluate on SemanticSTF without access to any target-domain data during training. Performance is measured by mean Intersection-over-Union (mIoU). We compare against eight prior methods spanning data augmentation and domain generalization: Dropout (Srivastava et al., 2014), Perturbation, PolarMix (Xiao et al., 2022), MMD (Li et al., 2018), PCL (Yao et al., 2022), PointDR (Xiao et al., 2023), UniMix (Zhao et al., 2024), and WADG (Du et al., 2024).

3.1 Main Results

Table 1 compares GDC against prior methods on the all-weather SemanticSTF benchmark. GDC achieves 38.6% mIoU, exceeding WADG (Du et al., 2024) by 3.8% and PointDR (Xiao et al., 2023) by 10.0%, with consistent gains across all four conditions: 40.7% under dense fog (+1.2% over WADG), 33.8%

Table 1. Domain generalization results (%) from SemanticKITTI to SemanticSTF. Per-class IoU across all weather conditions, per-weather mIoU under four adverse scenarios (dense fog, light fog, rain, snow), and overall all-weather mIoU are reported.

Method	Car	Bicycle	Motorcycle	Truck	Other-Vehicle	Person	Bicyclist	Motorcyclist	Road	Parking	Sidewalk	Other-Ground	Building	Fence	Vegetation	Trunk	Terrain	Pole	Traffic-Sign	Dense-Fog	Light-Fog	Rain	Snow	All-Weather mIoU (%)
	Per-Class IoU (%)																			Per-Weather mIoU (%)				
Baseline	55.9	0.0	0.2	1.9	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**	17.5	29.5	26.0	28.4	21.4	24.4
Dropout	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	29.3	25.6	29.4	24.8	25.7
Perturbation	74.4	0.0	0.0	23.3	0.6	19.7	0.0	0.0	60.3	10.8	33.9	0.7	72.0	45.2	58.7	17.5	42.4	22.1	9.7	26.3	27.8	30.0	24.5	25.9
PolarMix	57.8	1.8	3.8	16.7	3.7	26.5	0.0	2.0	65.7	2.9	32.5	0.3	71.0	48.7	53.8	20.5	45.4	25.9	15.8	29.7	25.0	28.6	25.6	26.0
MMD	63.6	0.0	2.6	0.1	11.4	28.1	0.0	0.0	**67.0**	14.1	37.9	0.3	67.3	41.2	57.1	27.4	47.9	28.2	16.2	30.4	28.1	32.8	25.2	26.9
PCL	65.9	0.0	0.0	17.7	0.4	8.4	0.0	0.0	59.6	12.0	35.0	1.6	74.0	47.5	60.7	15.8	48.9	26.1	27.5	28.9	27.6	30.1	24.6	26.4
PointDR	67.3	0.0	4.5	19.6	9.0	18.8	2.7	0.0	62.6	12.9	**38.1**	0.6	73.3	43.8	56.4	32.2	45.7	28.7	27.4	31.3	29.7	31.9	26.2	28.6
UniMix	82.7	6.6	8.6	4.5	**15.1**	35.5	**15.5**	37.7	55.8	10.2	36.2	1.3	72.8	40.1	49.1	33.4	34.9	23.5	33.5	34.8	30.2	34.9	30.9	31.4
WADG	72.0	0.0	**32.9**	**37.0**	1.9	37.7	6.8	52.9	59.9	10.7	31.8	2.2	**76.0**	48.8	**62.7**	34.0	**49.3**	23.6	20.4	39.5	32.5	31.7	29.4	34.8
Ours	**84.6**	**10.9**	21.6	22.6	11.4	**52.9**	4.0	**67.0**	64.8	**14.5**	36.7	**5.8**	74.2	**48.9**	57.1	**37.1**	48.3	28.7	**41.5**	**40.7**	**33.8**	**37.3**	**34.5**	**38.6**

Table 2. Domain generalization results (%) from SemanticKITTI to SemanticSTF under dense fog. Dash (–) indicates categories with no instances in this scenario.

Method	Car	Bicycle	Motorcycle	Truck	Other-Vehicle	Person	Bicyclist	Motorcyclist	Road	Parking	Sidewalk	Other-Ground	Building	Fence	Vegetation	Trunk	Terrain	Pole	Traffic-Sign	Dense-Fog mIoU (%)
	Per-Class IoU (%)																			
Baseline	74.7	-	-	7.8	0.0	6.4	8.9	0.0	72.2	0.6	**33.8**	0.0	59.6	48.7	56.9	27.4	56.4	27.2	21.1	29.5
Dropout	67.5	-	-	1.9	0.0	**8.9**	2.8	0.0	70.9	5.6	29.0	0.8	64.6	44.0	60.0	31.6	60.6	**28.1**	21.3	29.3
Perturbation	68.6	-	-	8.8	0.0	6.0	0.0	0.0	66.6	**14.8**	24.3	0.1	52.2	43.5	60.1	19.4	54.1	16.3	11.5	26.3
PolarMix	52.3	-	-	17.2	0.0	3.6	0.0	19.3	75.2	0.0	28.7	0.6	62.4	49.5	60.5	29.0	55.4	20.8	30.7	29.7
MMD	75.5	-	-	0.3	0.0	4.2	0.0	0.0	**75.4**	11.2	33.6	0.5	64.8	51.7	64.7	26.1	62.3	23.0	23.0	30.4
PCL	64.3	-	-	11.7	0.0	0.6	0.0	0.0	72.4	3.8	31.3	0.8	63.1	46.5	65.7	19.4	**64.3**	18.5	28.9	28.9
PointDR	69.2	-	-	7.1	0.0	2.4	6.7	0.0	73.5	8.5	33.6	0.2	65.6	47.6	63.6	31.0	60.7	24.4	38.8	31.3
WADG	78.8	-	-	**40.6**	0.0	7.0	**47.7**	67.3	71.6	0.3	29.8	1.4	65.0	48.0	**66.5**	33.0	60.8	21.3	33.0	39.5
Ours	**87.1**	-	-	37.4	0.0	6.7	10.4	**85.5**	74.4	0.6	26.4	**8.6**	**66.1**	**54.0**	58.4	**40.7**	54.4	23.1	**57.3**	**40.7**

Table 3. Domain generalization results (%) from SemanticKITTI to SemanticSTF under light fog.

Method	Car	Bicycle	Motorcycle	Truck	Other-Vehicle	Person	Bicyclist	Motorcyclist	Road	Parking	Sidewalk	Other-Ground	Building	Fence	Vegetation	Trunk	Terrain	Pole	Traffic-Sign	Light-Fog mIoU (%)
	Per-Class IoU (%)																			
Baseline	60.0	0.0	0.0	1.3	10.9	12.3	0.0	0.0	68.6	4.5	36.0	0.0	61.5	53.1	55.6	38.0	44.7	29.2	18.2	26.0
Dropout	63.2	0.0	0.0	3.2	10.2	5.5	0.0	0.0	63.8	4.9	29.4	0.1	62.5	53.1	58.6	42.5	46.6	27.8	14.3	25.6
Perturbation	76.6	0.0	0.0	38.2	0.0	21.9	0.0	0.0	66.6	**8.8**	34.6	0.1	62.4	56.1	63.2	25.3	46.2	22.4	6.5	27.8
PolarMix	42.6	0.2	0.0	29.4	3.3	17.0	0.0	**0.2**	69.8	0.7	33.1	0.1	56.2	56.3	54.9	24.7	44.8	24.1	16.6	25.0
MMD	63.6	0.0	0.0	0.1	**13.3**	25.9	0.0	0.0	**73.9**	5.6	42.8	0.1	64.1	55.3	61.9	36.6	50.7	29.2	9.9	28.1
PCL	66.3	0.0	0.0	26.7	0.2	8.7	0.0	0.0	67.8	5.0	36.7	0.4	64.3	58.0	66.1	21.2	53.1	25.5	24.6	27.6
PointDR	65.9	0.0	0.0	29.7	4.4	11.4	**0.9**	0.0	70.9	**8.8**	**43.3**	0.0	66.5	55.1	61.3	43.0	49.1	29.1	24.3	29.7
WADG	68.5	0.0	0.0	**54.0**	4.8	43.8	0.0	0.0	69.1	4.3	34.3	0.5	**68.4**	58.6	**66.6**	43.8	**55.3**	24.2	20.7	32.5
Ours	**78.2**	**2.8**	0.0	28.2	7.9	**49.4**	0.0	0.0	70.4	8.0	38.8	**5.3**	67.5	**61.3**	61.8	**45.0**	50.0	**29.9**	**38.6**	**33.8**

Table 4. Domain generalization results (%) from SemanticKITTI to SemanticSTF under rain. Dash (–) indicates categories with no instances in this scenario.

Method	Car	Bicycle	Motorcycle	Truck	Other-Vehicle	Person	Bicyclist	Motorcyclist	Road	Parking	Sidewalk	Other-Ground	Building	Fence	Vegetation	Trunk	Terrain	Pole	Traffic-Sign	Rain mIoU (%)
	Per-Class IoU (%)																			
Baseline	72.4	0.0	-	0.0	16.3	6.9	0.0	-	71.6	12.7	58.1	0.0	70.0	33.0	51.8	9.9	24.2	33.3	22.9	28.4
Dropout	81.3	0.0	-	0.0	**21.2**	5.6	0.0	-	62.2	11.8	44.8	0.6	76.8	44.7	56.0	16.3	23.3	32.8	22.2	29.4
Perturbation	83.9	0.0	-	2.4	0.0	20.9	0.0	-	73.2	12.6	54.7	7.0	71.7	43.2	58.3	5.9	29.4	29.4	16.9	30.0
PolarMix	56.7	**4.0**	-	**9.1**	1.5	29.8	0.0	-	68.2	10.9	50.2	0.5	73.2	**47.2**	48.3	17.8	22.3	32.3	14.1	28.6
MMD	83.9	0.0	-	0.0	8.9	31.6	0.0	-	**77.9**	17.9	**60.2**	0.3	69.6	39.3	58.4	14.1	**32.5**	34.0	30.0	32.8
PCL	84.2	0.0	-	0.0	0.1	4.3	0.0	-	68.1	10.9	55.5	4.6	74.7	43.9	59.6	5.8	27.3	34.2	**38.8**	30.1
PointDR	78.0	0.0	-	0.0	13.8	20.0	0.0	-	72.1	14.7	60.0	1.2	76.1	36.9	58.0	18.3	24.7	36.1	32.5	31.9
WADG	84.3	0.0	-	0.4	0.0	40.6	0.0	-	65.1	**19.2**	51.3	4.2	**77.3**	42.7	**61.3**	15.7	22.8	30.9	23.5	31.7
Ours	**85.8**	2.0	-	0.1	17.4	**63.2**	0.0	-	76.3	15.5	57.7	**9.6**	75.3	45.4	57.4	**28.2**	28.6	**37.7**	34.1	**37.3**

Table 5. Domain generalization results (%) from SemanticKITTI to SemanticSTF under snow. Dash (–) indicates categories with no instances in this scenario.

Method	Car	Bicycle	Motorcycle	Truck	Other-Vehicle	Person	Bicyclist	Motorcyclist	Road	Parking	Sidewalk	Other-Ground	Building	Fence	Vegetation	Trunk	Terrain	Pole	Traffic-Sign	Snow mIoU (%)
	Per-Class IoU (%)																			
Baseline	49.5	0.0	0.3	0.5	11.6	10.8	-	-	42.1	14.9	23.9	0.0	71.5	26.7	29.3	24.0	17.8	**30.8**	10.1	21.4
Dropout	58.5	0.0	30.5	5.4	13.2	5.2	-	-	41.9	18.0	20.4	2.5	76.4	30.5	31.8	**32.7**	19.8	28.2	7.0	24.8
Perturbation	73.6	0.0	0.0	5.5	1.1	19.8	-	-	45.7	10.9	34.4	0.1	**80.6**	32.8	45.2	12.8	20.0	24.4	9.5	24.5
PolarMix	66.5	3.4	9.3	3.5	5.8	32.4	-	-	**55.3**	3.6	30.1	0.1	77.8	36.1	34.2	12.6	25.1	29.8	10.1	25.6
MMD	59.4	0.0	4.7	0.0	14.7	30.5	-	-	50.8	16.9	32.8	0.2	68.4	24.4	36.6	24.1	24.1	30.0	11.4	25.2
PCL	64.0	0.0	0.0	8.2	0.7	9.2	-	-	38.9	15.2	31.6	2.3	79.6	35.1	41.3	11.2	23.1	30.1	**26.8**	24.6
PointDR	66.2	0.0	10.4	0.0	**16.7**	21.3	-	-	43.0	15.2	33.0	1.7	76.8	30.3	36.1	27.6	22.2	30.0	14.1	26.2
WADG	70.6	0.1	**39.8**	**8.7**	0.9	39.4	-	-	39.8	13.8	27.4	3.4	**80.6**	**36.6**	**46.8**	30.6	26.0	23.5	12.0	29.4
Ours	**86.2**	**15.3**	30.2	2.1	13.9	**57.2**	-	-	48.8	**19.6**	**37.0**	**3.6**	78.2	32.7	42.8	31.1	**33.2**	30.0	25.1	**34.5**

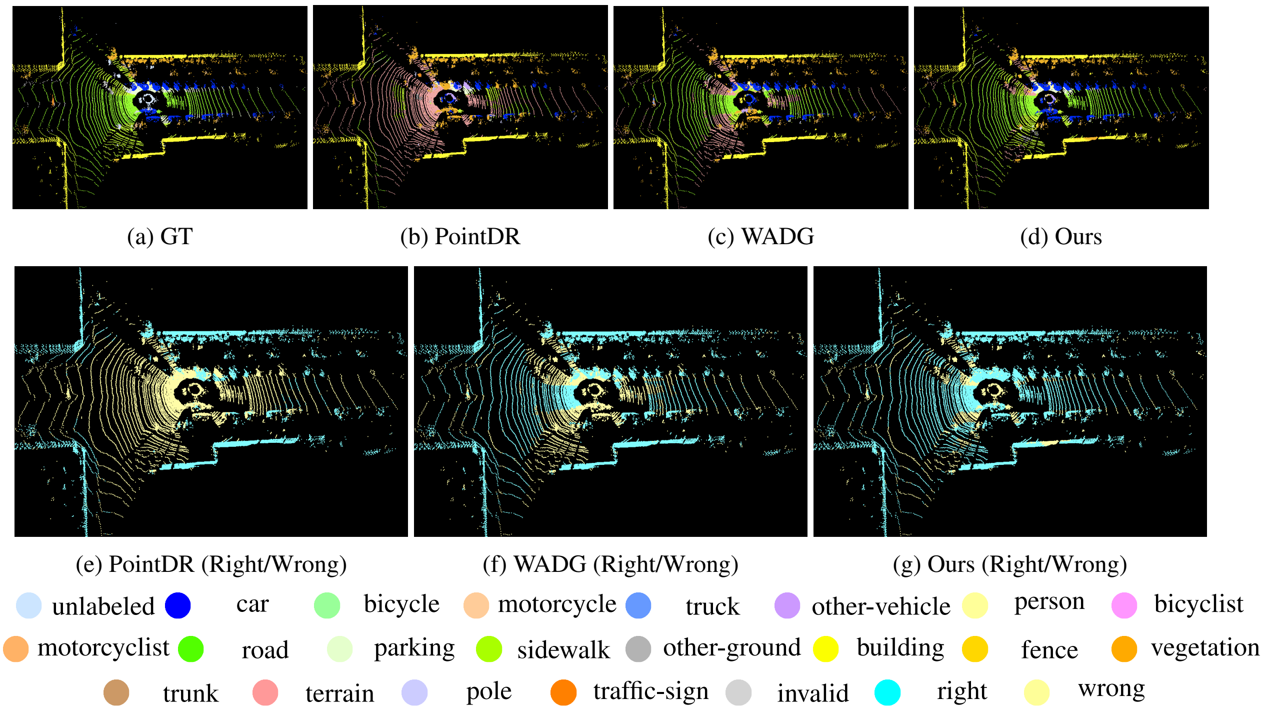

Figure 1. Qualitative results with SemanticKITTI as source and SemanticSTF as target (light fog). Cyan areas (Right) indicate correctly segmented regions, while yellow areas (Wrong) represent erroneously segmented regions.

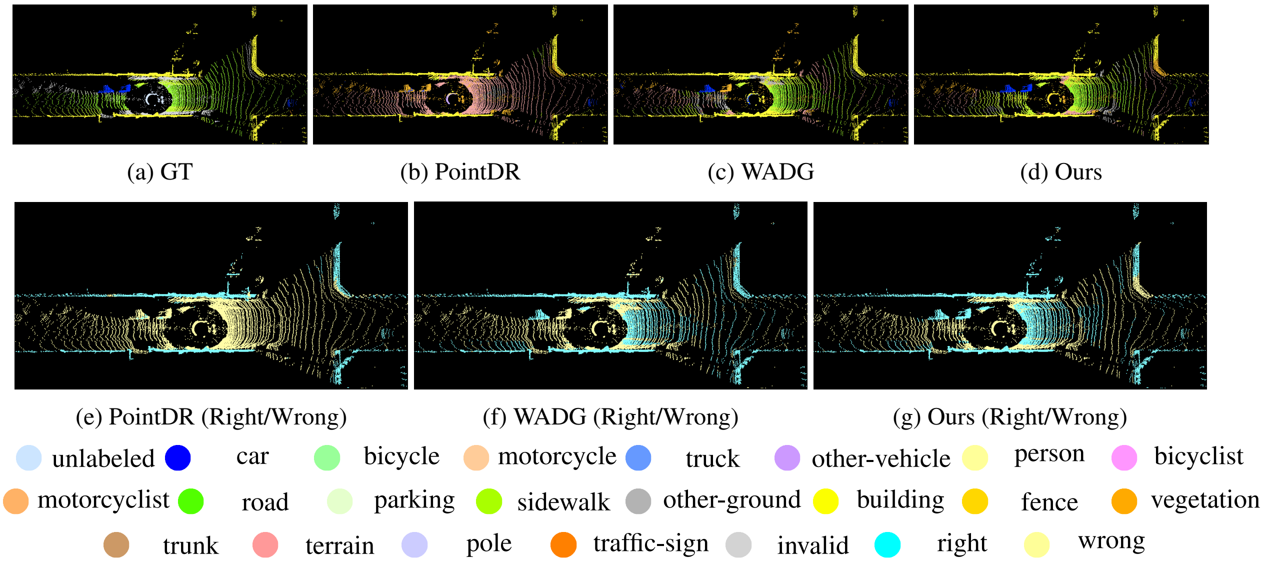

(a) GT (b) PointDR (c) WADG (d) Ours

(e) PointDR (Right/Wrong) (f) WADG (Right/Wrong) (g) Ours (Right/Wrong)

Figure 2. Qualitative results with SemanticKITTI as source and SemanticSTF as target (snow). Cyan areas (Right) indicate correctly segmented regions, while yellow areas (Wrong) represent erroneously segmented regions.

under light fog (+1.3%), 37.3% under rain (+5.6%), and 34.5% under snow (+5.1%). The largest margins appear under rain and snow, where backscatter insertion and volumetric scattering create the most severe geometric distortions. At the per-class level, traffic-sign reaches 41.5% IoU, above WADG (20.4%) by 21.1% and PointDR (27.4%) by 14.1%. Person attains 52.9%, exceeding WADG (37.7%) by 15.2% and PointDR (18.8%) by 34.1%. Motorcyclist reaches 67.0%, above WADG (52.9%) by 14.1%. Car improves to 84.6% from 72.0% (WADG, +12.6%) and 67.3% (PointDR, +17.3%). Bicycle rises to 10.9% from 0.0% under both WADG and PointDR. These categories share sparse point coverage or compact spatial extent, where the diffusion consistency constraint helps preserve class boundaries through degraded neighborhoods. WADG retains higher IoU on truck (37.0% vs. 22.6%) and motorcycle (32.9% vs. 21.6%), both large-volume classes where distribution alignment appears more effective. The gains on person, motorcyclist, and traffic-sign are particularly relevant to autonomous driving safety, as these represent vulnerable road users and navigation-critical infrastructure.

3.2 Per-Weather Analysis

Tables 2–5 report per-class IoU under each weather condition separately. Under dense fog (Table 2), GDC achieves 40.7% mIoU, exceeding WADG (Du et al., 2024) by 1.2% and PointDR (Xiao et al., 2023) by 9.4%. Traffic-sign reaches 57.3% IoU, higher than WADG (33.0%) by 24.3% and PointDR (38.8%) by 18.5%, the largest per-class margin across all methods in this condition. Motorcyclist attains 85.5%, surpassing WADG (67.3%) by 18.2%, while PointDR does not produce valid segmentation for this class. Car reaches 87.1%, the highest absolute IoU in Table 2, above WADG (78.8%) by 8.3% and PointDR (69.2%) by 17.9%. Trunk (40.7% vs. 33.0% from WADG, +7.7%) and fence (54.0% vs. 48.0% from WADG, +6.0%) show further gains on structurally thin objects. Dense fog causes range-dependent attenuation that progressively removes distant returns; the diffusion operator propagates features along surviving local connections, which helps preserve discriminative information for sparse categories such as traffic signs and motorcyclists, whose accurate detection is directly relevant to safe autonomous navigation under reduced visibility.

Light fog (Table 3) represents a milder scenario, and GDC achieves 33.8% mIoU, above WADG by 1.3% and PointDR by 4.1%. The smaller overall margin is consistent with weaker attenuation and less extensive point deletion. Category-level differences, however, remain notable. Traffic-sign reaches 38.6%, nearly doubling WADG (20.7%, +17.9%) and exceeding PointDR (24.3%) by 14.3%. Car improves to 78.2% from 68.5% (WADG, +9.7%) and 65.9% (PointDR, +12.3%). Person attains 49.4%, above WADG (43.8%) by 5.6% and PointDR (11.4%) by 38.0%, while pole reaches 29.9% compared to 24.2% from WADG (+5.7%). These results suggest that the diffusion consistency constraint provides stable feature propagation even when neighborhood perturbations are moderate. Qualitative results under this condition are shown in Figure 1.

Rain (Table 4) introduces backscatter insertion and range-dependent point deletion simultaneously. GDC reaches 37.3% mIoU, above WADG by 5.6% and PointDR by 5.4%. Person reaches 63.2%, compared to 40.6% from WADG (+22.6%) and 20.0% from PointDR (+43.2%), the largest single-class margin in Table 4. Pedestrians occupy compact spatial footprints that are susceptible to backscatter-induced occlusion, making this improvement relevant to real-world deployment where rain is the most frequently encountered adverse condition. Other-vehicle recovers to 17.4% from 0.0% under WADG and 13.8% under PointDR (+3.6%), indicating that GDC can produce valid segmentation for categories on which some prior methods fail entirely. Trunk reaches 28.2%, above WADG (15.7%) by 12.5% and PointDR (18.3%) by 9.9%. Road attains 76.3%, exceeding WADG (65.1%) by 11.2%, and pole reaches 37.7% compared to 30.9% from WADG (+6.8%).

Snow (Table 5) presents the most severe compound corruption through ground-level clutter and volumetric scattering. GDC achieves 34.5% mIoU, exceeding WADG by 5.1% and PointDR by 8.3%. Car reaches 86.2%, above WADG (70.6%) by 15.6% and PointDR (66.2%) by 20.0%, the highest vehicle IoU across all four weather conditions and all methods. Person improves to 57.2% from 39.4% (WADG, +17.8%) and 21.3% (PointDR, +35.9%). Bicycle rises to 15.3% from 0.1% under WADG and 0.0% under PointDR, a category on which nearly all prior methods produce zero IoU; bicycles combine thin frame geometry with low point counts, making them particularly susceptible to snow-induced density loss. Sidewalk reaches 37.0%, above WADG (27.4%) by 9.6%, and terrain improves to 33.2% from 26.0% (WADG, +7.2%) and 22.2% (PointDR, +11.0%), reflecting more reliable ground-surface distinction when snow cover alters surface geometry. Figure 2 provides qualitative comparisons under this condition. Snow remains one of the least-addressed conditions in deployed autonomous perception systems; reliable segmentation of cyclists and pedestrians under winter weather is directly relevant to reducing risk for vulnerable road users.

4. Conclusion

This work presented Geometric Diffusion Consistency (GDC) for 3D point cloud segmentation under adverse weather. The method regularizes feature propagation behavior across augmented views through a lightweight sparse convolution-based diffusion operator, combined with an augmentation pipeline that incorporates dual-mode atmospheric extinction modeling and semantic-aware geometric corruption. On the SemanticKITTI to SemanticSTF domain generalization benchmark, GDC achieves 38.6% mIoU, exceeding WADG by 3.8% and PointDR by 10.0%, with improvements observed across all four weather conditions. The per-class results indicate that categories with sparse point coverage or compact spatial extent, such as person, motorcyclist, and traffic-sign, benefit most from the propagation-level consistency constraint. Several limitations should be noted. The evaluation is conducted on a single source-target pair; additional LiDAR datasets with different sensor configurations and geographic environments would be needed to assess broader generalizability. The current framework does not incorporate temporal information from sequential scans, which could further stabilize segmentation under transient weather effects.

References

Behley, J., Garbade, M., Milioto, A., Quenzel, J., Behnke, S., Stachniss, C., Gall, J., 2019. Semantickitti: A dataset for semantic scene understanding of lidar sequences. *Proc.ICCV*, IEEE, 9296–9306.

Du, J., Cai, G., Wang, Z., Huang, S., Su, J., Marcato Junior, J., Smit, J., Li, J., 2021. ResDLPS-Net: Joint residual-dense optimization for large-scale point cloud semantic segmentation. *ISPRS J. Photogramm. Remote Sens.*, 182, 37-51.

Du, J., Zelek, J., Li, J., 2024. Weather-aware autopilot: Domain generalization for point cloud semantic segmentation in diverse weather scenarios. *ISPRS J. Photogramm. Remote Sens.*, 218, 204-219.

Du, J., Zelek, J., Zhang, D., Li, J., 2025a. 3DLCDM: Hybrid supervision for land cover discovery mapping of emerging urban structures in 3D remote sensing. *Remote Sens. Environ.*, 331, 115018.

Du, J., Zelek, J., Zhang, D., Li, J., 2025b. Gated multi-source fusion with geometric sequence modeling for novel urban structure discovery. *ISPRS J. Photogramm. Remote Sens.*, 230, 495-523.

Hahner, M., Sakaridis, C., Bijelic, M., Heide, F., Yu, F., Dai, D., Gool, L. V., 2022. Lidar snowfall simulation for robust 3d object detection. *Proc. CVPR*, IEEE, 16343–16353.

Hahner, M., Sakaridis, C., Dai, D., Gool, L. V., 2021. Fog simulation on real lidar point clouds for 3d object detection in adverse weather. *Proc. ICCV*, IEEE, 15263–15272.

Li, H., Pan, S. J., Wang, S., Kot, A. C., 2018. Domain generalization with adversarial feature learning. *Proc. CVPR*, IEEE, 5400–5409.

Lu, D., Zhou, J., Gao, K. Y., Du, J., Xu, L., Li, J., 2024. Dynamic clustering transformer network for point cloud segmentation. *Int. J. Appl. Earth Obs. Geoinformation*, 128, 103791.

Park, J., Kim, K., Shim, H., 2024. Rethinking data augmentation for robust lidar semantic segmentation in adverse weather. A. Leonardis, E. Ricci, S. Roth, O. Russakovsky, T. Sattler, G. Varol (eds), *Proc. ECCV*, 15074, 320–336.

Srivastava, N., Hinton, G. E., Krizhevsky, A., Sutskever, I., Salakhutdinov, R., 2014. Dropout: a simple way to prevent neural networks from overfitting. *J. Mach. Learn. Res.*, 15(1), 1929–1958.

Tang, H., Liu, Z., Zhao, S., Lin, Y., Lin, J., Wang, H., Han, S., 2020. Searching efficient 3d architectures with sparse point-voxel convolution. A. Vedaldi, H. Bischof, T. Brox, J. Frahm (eds), *Proc. ECCV*, 12373, Springer, 685–702.

Xiao, A., Huang, J., Guan, D., Cui, K., Lu, S., Shao, L., 2022. Polarmix: A general data augmentation technique for lidar point clouds. *Proc. NeurIPS*.

Xiao, A., Huang, J., Xuan, W., Ren, R., Liu, K., Guan, D., El-Saddik, A., Lu, S., Xing, E. P., 2023. 3d semantic segmentation in the wild: Learning generalized models for adverse-condition point clouds. *Proc. CVPR*, IEEE, 9382–9392.

Yao, X., Bai, Y., Zhang, X., Zhang, Y., Sun, Q., Chen, R., Li, R., Yu, B., 2022. PCL: proxy-based contrastive learning for domain generalization. *Proc. CVPR*, IEEE, 7087–7097.

Zhao, H., Zhang, J., Chen, Z., Zhao, S., Tao, D., 2024. Unimix: Towards domain adaptive and generalizable lidar semantic segmentation in adverse weather. *Proc. CVPR*, IEEE, 14781–14791.

Zhu, X., Zhou, H., Wang, T., Hong, F., Ma, Y., Li, W., Li, H., Lin, D., 2021. Cylindrical and asymmetrical 3d convolution networks for lidar segmentation. *Proc. CVPR*, IEEE, 9939–9948.

https://doi.org/10.5194/isprs-archives-XLIX-B2-2026-161-2026