ADC: LIDAR LOCALIZATION FOR AUTONOMOUS DRIVING IN VARIOUS WEATHER CONDITIONS

Anonymous authors
Paper under double-blind review



Figure 1: **Teaser.** While existing localization methods perform well in clear weather, they experience severe localization degradation under adverse rain and snow conditions due to airborne noise and structural corruption. In contrast, our proposed **ADC** maintains accurate and robust localization across all weather conditions, achieving the top-1 rank on both the standard Oxford benchmark and the adverse-weather Boreas dataset.

ABSTRACT

Accurate and robust 6-DoF localization is essential for autonomous vehicles, where LiDAR has emerged as a mainstream alternative to vision-based sensors due to its direct 3D geometric perception. However, existing LiDAR localization methods are primarily designed for clear weather, overlooking severe signal degradation in adverse environments. In rain and snow, precipitation-induced backscattering, laser attenuation, and snow accumulation lead to airborne clutter, point dropouts, and structural corruption in LiDAR observations, severely degrading localization performance. To break this bottleneck, this paper presents **ADC**, the first map-free localization framework specifically designed for robust LiDAR localization under adverse weather conditions. To combat weather-induced noise and observation inconsistency, we propose a hierarchical sparse encoding architecture coupled with perspective-invariant global aggregation and an adaptive point-wise feature recalibration module to effectively filter out airborne clutter while amplifying stable rigid structures. Furthermore, we introduce a local geometric rigidity constraint that enforces an implicit structural skeleton on degraded point clouds, resolving coordinate drift and preserving topological consistency. Extensive evaluations on the challenging Boreas benchmark and the standard Oxford RobotCar dataset demonstrate that ADC achieves state-of-the-art performance. Specifically, our method reduces mean translation error by 70% on the adverse-weather Boreas dataset and by 18% on the Oxford benchmark relative to the strongest competing map-free baselines, demonstrating robust and weather-resilient localization across diverse driving conditions.

1 INTRODUCTION

Accurate and robust 6-DoF re-localization is the cornerstone of autonomous navigation for unmanned systems, including autonomous vehicles and mobile robots (Yin et al., 2024; Lu et al., 2019). While mainstream localization paradigms heavily rely on visual cameras and Light Detection and Ranging (LiDAR) sensors (Jiang et al., 2025; Wang et al., 2021), vision-based methods (Sattler et al., 2016; Mao et al., 2026) are notoriously vulnerable to illumination variations, low-

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

light environments, and featureless scenes. Consequently, LiDAR has gained widespread attention and adoption owing to its ability to provide direct, high-precision 3D geometric structures.

However, traditional LiDAR localization systems typically operate on a map-based retrieval paradigm Liu et al. (2025); Luo et al. (2025). These methods require pre-building and storing massive global feature databases, which incurs severe storage overhead and prohibitive communication costs during large-scale deployment. To bypass these limitations, map-free regression methods (Wang et al., 2021; Yu et al., 2022) have emerged as a promising alternative. By implicitly encoding environmental landmarks into neural network parameters, map-free approaches maintain a constant memory footprint regardless of scene scale while enabling efficient inference.

Existing map-free localization approaches can be broadly categorized into Absolute Pose Regression (APR) and Scene Coordinate Regression (SCR). APR models directly predict the sensor's 6-DoF pose from raw point clouds (Yu et al., 2023; Wang et al., 2023). In contrast, SCR frameworks (Li et al., 2023) first regress dense 3D scene coordinates for individual points and subsequently resolve the final pose through RANSAC algorithm (Chum & Matas, 2008). Compared with Absolute Pose Regression (APR), Scene Coordinate Regression (SCR) has rapidly become the mainstream paradigm due to its superior geometric fidelity, robustness, and interpretability.

Despite their notable success, nearly all existing LiDAR map-free regression models are designed under the idealized assumption of benign clear weather conditions (Li et al., 2023; Yang et al., 2024; Chen et al., 2025), severely overlooking the physical degradation of LiDAR in adverse environments. In rain and snow, laser beams suffer from severe attenuation, absorption, and refraction caused by suspended airborne particles. This leads to heavy backscatter clutter such as false reflection points from raindrops and snowflakes as well as dynamic point dropouts. Furthermore, wet ground surfaces and snow accumulation physically alter and conceal stable rigid structures including building facades, curbs, and pole-like landmarks. As illustrated in Fig. 1, while existing methods perform well in clear scenarios, they experience catastrophic performance degradation in adverse weather due to coordinate drift and topological collapse. Overcoming LiDAR degradation caused by weather-induced noise and structural distortion remains an open and critical challenge for map-free re-localization.

To break this bottleneck, this paper presents **ADC**, the first map-free re-localization framework specifically tailored for robust LiDAR localization under adverse weather conditions. Instead of forcing the network to blindly predict coordinates from corrupted and noisy reflections, we redesign the regression architecture to be structurally aware and noise-resilient. First, we propose a hierarchical sparse encoding architecture coupled with perspective-invariant global aggregation to capture multi-scale geometric contexts across canonical viewpoints, providing a stable spatial prior that mitigates weather-induced observation inconsistency. Second, an adaptive point-wise feature recalibration module is introduced to dynamically filter out transient airborne clutter and rain or snow noise while amplifying stable and salient rigid geometric cues. Finally, to prevent predicted scene coordinates from local stretching and anisotropic drift caused by missing returns and snow deformation, we design a local geometric rigidity constraint, termed EdgeLoss. By enforcing topological distance consistency across neighboring points, this constraint imposes an implicit structural skeleton on degraded LiDAR points to preserve the underlying rigidity of the scene. In summary, the main contributions of this work are fourfold:

- We propose **ADC**, the first map-free re-localization framework specifically designed to overcome airborne noise, point dropouts, and structural corruption for LiDAR localization under adverse weather.
- We design a hierarchical sparse encoding architecture with perspective-invariant aggregation and an adaptive point-wise feature recalibration module to effectively suppress rain and snow backscatter clutter and extract stable geometric features.
- We develop a geometric rigidity constraint based on an implicit structural skeleton to resolve coordinate distortion and topological collapse in adverse-weather condition coordinate regression.
- Extensive experiments on the challenging adverse-weather Boreas benchmark and the standard Oxford RobotCar dataset demonstrate that ADC achieves state-of-the-art performance, outperforming leading map-free methods by 70% under severe rain and snow conditions.

2 Related Work

2.1 Map-based LiDAR Localization

Map-based localization matches query LiDAR scans against pre-built 3D point clouds or global feature databases, encompassing retrieval-based and registration-based paradigms. Retrieval methods, such as PointNetVLAD Uy & Lee (2018), LPD-Net Liu et al. (2019), and HiTPR Hou et al. (2022), extract global descriptors for coarse place recognition. Registration approaches, including BUFFER Ao et al. (2023), HRegNet Lu et al. (2021), Bevplace++ Luo et al. (2025), and UniLGL Shen et al. (2026), estimate precise 6-DoF poses by aligning scans with dense submaps. Although accurate under clear conditions, these methods suffer from heavy storage and communication costs that scale linearly with the mapped area, limiting large-scale deployment.

2.2 Map-free Regression Localization

Map-free methods implicitly encode environmental landmarks into neural network parameters, maintaining a constant memory footprint. In absolute pose regression (APR), models like HypLiLoc Wang et al. (2023) and DiffLoc Li et al. (2024) directly regress 6-DoF poses via multi-projection fusion or diffusion denoising. In scene coordinate regression (SCR), networks predict dense 3D scene coordinates and solve poses via RANSAC Chum & Matas (2008). Representative SCR works, including SGLoc Li et al. (2023), LiSA Yang et al. (2024), and PELoc Chen et al. (2025), improve accuracy through multi-scale enhancement, semantic distillation, and robust augmentation. However, these methods assume benign weather and suffer catastrophic degradation when exposed to rain or snow.

2.3 Challenges in Adverse-Weather LiDAR Localization

Achieving robust map-free LiDAR localization in adverse weather remains challenging due to physical degradation, including airborne backscatter noise, beam attenuation, and geometric occlusions from snow accumulation. These corruptions introduce outliers and disrupt local geometric rigidity, causing coordinate drift and topological distortion in existing regression pipelines. To address these issues, we propose **ADC**, a structurally aware and noise-resilient map-free framework that suppresses airborne clutter and enforces geometric rigidity for robust all-weather localization.

3 Method

3.1 Hierarchical Sparse Representation and Perspective-Invariant Global Encoding

To mitigate LiDAR point dropouts and observation corruption caused by adverse weather, as shown in Fig. 2, we propose a hierarchical framework integrating sparse convolutions and perspective-invariant aggregation. This architecture incorporates a spatial prior to narrow the solution space and ensure positional stability despite weather-degraded observations.

Hierarchical Sparse Feature Extraction. Given an input point cloud P with N points, we discretize the raw LiDAR reflections into a voxelized representation. We design a sparse encoder $\Phi(\cdot)$ to capture multi-scale geometric patterns. Initial feature extraction utilizes a sequence of strided sparse convolutions that progressively downsample spatial resolution while increasing channel depth. The transformation of the k-th layer is formulated as follows:

$$\mathbf{f}_k = \text{ReLU}(\text{SparseConv}(\mathbf{f}_{k-1}, \mathbf{W}_k, s_k)), \tag{1}$$

where $\mathbf{W}_k$ and s_k denote the learnable kernels and the stride of the k-th layer respectively. To further refine the local features and ensure stable gradient flow, we stack multiple residual blocks. For a given intermediate feature $\mathbf{f}_{res}$, the refined output is computed through a residual mapping $M(\cdot)$ as follows:

$$\mathbf{f}_{out} = \mathbf{f}_{res} + M(\mathbf{f}_{res}, \mathbf{W}_{res}), \tag{2}$$

where M consists of multiple sparse convolutional layers with bottleneck structures, and $\mathbf{W}_{res}$ represents the learnable weight matrices of the residual branch. By stacking these hierarchical layers,

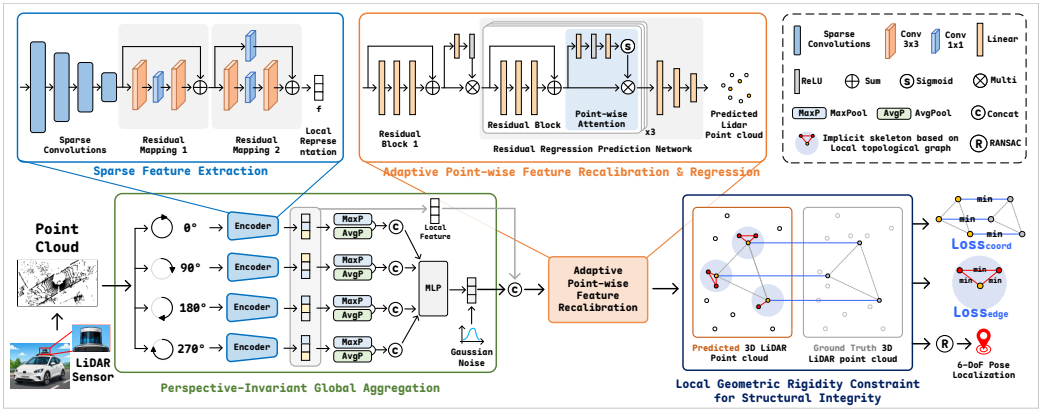

Figure 2: **The framework of ADC.** Hierarchical sparse feature extraction first captures multi-scale geometric patterns via sparse convolutions and residual mappings. Perspective-invariant global aggregation distills a rotation-robust scene context across a discrete rotation group. The adaptive point-wise feature recalibration module purifies features by suppressing rain/snow clutter while amplifying stable rigid structures. A local geometric rigidity constraint preserves structural integrity by imposing an implicit skeleton on weather-corrupted LiDAR points. The final 6-DoF pose is solved via RANSAC based on the predicted 3D scene coordinates.

the encoder yields a discriminative local representation $\mathbf{f} \in \mathbb{R}^{m \times d}$ where m is the number of sparse voxels and d is the feature dimension.

Perspective-Invariant Global Aggregation. To achieve invariance to sensor heading and alleviate viewpoint-dependent weather artifacts, we introduce a perspective-invariant mechanism based on discrete group symmetry. We define a discrete rotation group G representing orthogonal transformations around the vertical axis:

$$G = \{R_\theta \mid \theta = k \cdot 90^\circ, k \in \{0, 1, 2, 3\}\}. \tag{3}$$

The encoder $\Phi(\cdot)$ processes each rotated input $R_\theta P$ to obtain viewpoint-specific feature matrices:

$$\mathbf{F}_\theta = \Phi(R_\theta P). \tag{4}$$

To distill a holistic scene context, we apply a dual-pooling strategy followed by symmetric aggregation across the rotation orbit. The perspective-invariant global descriptor $\mathbf{g} \in \mathbb{R}^{2d}$ as follows:

$$\mathbf{g} = \text{Concat}\left(\max_{\theta \in G} \text{AvgPool}(\mathbf{F}_\theta), \max_{\theta \in G} \text{MaxPool}(\mathbf{F}_\theta)\right). \tag{5}$$

This descriptor encapsulates the global context while remaining robust to the orientation and heading variations of the LiDAR sensor.

Spatial Prior Conditioning. The global descriptor $\mathbf{g}$ estimates a spatial topology prior $\mathbf{g}^*$ through a lightweight estimator $\Omega(\cdot)$ which maps the scene context to a coarse regional response. The resulting prior is L_2-normalized for numerical stability. Finally, we broadcast the spatial prior and concatenate it with the local geometric features $\mathbf{f}$ to obtain the conditioned latent representation $\mathbf{h}_i \in \mathbb{R}^{d_h}$ as follows:

$$\mathbf{h}_i = \text{Concat}(\mathbf{f}, \mathbf{g}^* + \boldsymbol{\epsilon}), \tag{6}$$

where $\boldsymbol{\epsilon}$ is a Gaussian perturbation used during training to model coarse prior uncertainty. The representation $\mathbf{h}_i$ resolves spatial ambiguity caused by weather-induced dropouts and provides a reliable foundation for subsequent tasks.

3.2 Adaptive Point-wise Feature Recalibration

To mitigate noise from rain/snow backscatter reflections in degraded LiDAR returns, we propose an adaptive purification mechanism. This module recalibrates channel importance via point-wise attention to suppress noise-related activations while amplifying stable structural information.

Point-wise Feature Recalibration. Unlike squeeze-and-excitation-based channel recalibration modules commonly adopted in image-based architectures Deng et al. (2025), our module performs channel recalibration on a per-point basis to accommodate the irregular and non-uniform distribution of degraded LiDAR scans. Given the enriched feature $\mathbf{h}_i \in \mathbb{R}^{d_h}$, the module learns a channel-wise gating vector $\mathbf{s}_i$ to adaptively adjust the feature response:

$$\mathbf{s}_i = \text{Sigmoid}(\mathbf{W}_2\text{ReLU}(\mathbf{W}_1\mathbf{h}_i)), \tag{7}$$

where $\mathbf{W}_1$ and $\mathbf{W}_2$ are learnable weight matrices. The purified $\hat{\mathbf{h}}_i$ is then obtained by:

$$\hat{\mathbf{h}}_i = \mathbf{h}_i \odot \mathbf{s}_i. \tag{8}$$

This recalibration process allows the network to distinguish between reliable geometric channels and noisy weather-induced clutter by modeling the inter-dependencies between the local geometric embeddings and the injected spatial priors.

Residual Regression Prediction Network. The coordinate mapping is performed by a deep residual network that iteratively refines the representation $\mathbf{h}_i$ through recalibrated blocks to prevent feature degradation. Let $\mathbf{r}_i^{(k)}$ denote the feature at block k with the initial state $\mathbf{r}_i^{(0)} = \mathbf{h}_i$. The transformation within each block is formulated as follows:

$$\mathbf{r}_i^{(k+1)} = \mathcal{A}^{(k)}(\mathbf{r}_i^{(k)} + \Psi^{(k)}(\mathbf{r}_i^{(k)})), \tag{9}$$

where $\Psi^{(k)}$ represents a non-linear mapping with linear layers and ReLU activations while $\mathcal{A}^{(k)}$ denotes the recalibration module. This iterative structure filters transient particulate clutter and amplifies stable scene structures such as building facades and pole landmarks. The final layer of the network maps the refined high-dimensional features to absolute world coordinates:

$$\hat{\mathbf{X}}_i = \Gamma(\mathbf{r}_i^{(L)}), \quad \hat{\mathbf{X}}_i \in \mathbb{R}^3, \tag{10}$$

where $\Gamma(\cdot)$ is the output projection layer and L is the total number of residual blocks. To supervise the coordinate mapping, we employ a robust regression loss based on the L_1 norm:

$$\mathcal{L}_{\text{coord}} = \frac{1}{N}\sum_{i=1}^{N} \|\hat{\mathbf{X}}_i - \mathbf{X}_i\|_1, \tag{11}$$

where $\mathbf{X}_i$ is the ground-truth coordinate in the world frame.

3.3 Local Geometric Rigidity Constraint for Structural Integrity

To address coordinate drift and geometric distortion caused by unstable LiDAR noise, missing returns, and snow accumulation, we introduce a local geometric rigidity constraint to preserve structural coherence. This mechanism establishes topological relationships among neighboring points and imposes an implicit structural skeleton on weather-corrupted point clouds to maintain geometric consistency during coordinate mapping.

Implicit Structural Skeleton Construction. We construct a local topological graph by identifying point pairs i and j within close proximity in the ground-truth frame. The set of local edges $\mathcal{E}$ is defined as follows:

$$\mathcal{E} = \{(i, j) \mid 0 < \|\mathbf{X}_i - \mathbf{X}_j\|_2 < \tau\}, \tag{12}$$

where the distance threshold τ defines the local receptive field. By focusing on the relative geometric layout instead of absolute height or reflectivity confidence, this strategy enables the network to preserve the underlying rigidity of physical scene objects even when their surfaces are partially obscured by snow or rain spray.

Geometric Consistency Loss. To penalize structural deformations such as local stretching, anisotropic drift, or geometric collapsing, we define the geometric consistency loss as follows:

$$\mathcal{L}_{\text{edge}} = \frac{1}{|\mathcal{E}|}\sum_{(i,j)\in\mathcal{E}} \left| \|\hat{\mathbf{X}}_i - \hat{\mathbf{X}}_j\|_2 - \|\mathbf{X}_i - \mathbf{X}_j\|_2 \right|. \tag{13}$$

This constraint matches predicted relative distances to ground truth counterparts to ensure topological stability. The total objective function for the regression task is a multi-task loss as follows:

$$\mathcal{L}_{\text{reg}} = \mathcal{L}_{\text{coord}} + \lambda_e \mathcal{L}_{\text{edge}}, \tag{14}$$

where $\lambda_e = 0.1$ balances point-level accuracy and structural rigidity. Finally, the regressed world coordinates are processed through the RANSAC algorithm to solve for the 6-DoF localization result and ensure robust pose estimation.

4 EXPERIMENT

4.1 EXPERIMENTAL SETUP

Datasets and Metrics. Performance is evaluated on the challenging **Boreas** Burnett et al. (2023) dataset and the standard **Oxford RobotCar** Maddern et al. (2017) LiDAR benchmark. Boreas dataset contains diverse and severe adverse weather conditions, while Oxford serves as a canonical baseline under clear urban conditions. To ensure fair comparison, all methods share identical training and evaluation splits, as well as a consistent pre-processing pipeline. Mean position error (m) and mean rotation error ($^\circ$) are adopted as the primary evaluation metrics.

1. Boreas is a large-scale, multi-weather dataset collected along repeated urban driving routes in Toronto, Canada, spanning over an 8 km trajectory per run. Point clouds are captured at 10 Hz using a 128-beam Velodyne Alpha-Prime LiDAR, while ground-truth 6-DoF poses are provided by a high-precision Applanix POS LV GNSS-INS system post-processed with RTK corrections. We conduct evaluations on the main repetitive Glen Shields route. To rigorously evaluate weather resilience, the training set consists of four clear/sunny sequences: 2021-03-23, 2021-03-30, 2021-04-13, and 2021-06-17. The testing set encompasses diverse adverse weather sequences: 2020-11-26 (Snow), 2020-12-01 (Snow), 2021-02-02 (Snow), 2021-03-02 (Sun), and 2021-04-29 (Rain). The heavy snow and rain noise makes this dataset particularly challenging for regression-based localization.

2. Oxford is a standard benchmark for LiDAR localization tasks featuring multiple urban driving sequences of 10 km. Point clouds are captured at 20Hz using a Velodyne 32-beam LiDAR sensor. The ground truths are refined following the strategy proposed in SGLoc Li et al. (2023). According to established protocols, we utilize sequences 15-13-06-37, 17-13-26-39, 17-14-03-00, and 18-14-14-42 for testing and performance evaluation.

Implementation Details. All experiments are conducted on two NVIDIA GeForce RTX 5090 GPUs. We implement our method using PyTorch and MinkowskiEngine Choy et al. (2020). The raw LiDAR point clouds are voxelized with a voxel size of 0.50m for the Boreas dataset and 0.25m for the Oxford dataset. The network is optimized using the AdamW optimizer with a OneCycleLR scheduler Loshchilov & Hutter (2017); Smith & Topin (2019), where the initial learning rates for the encoder and regression head are set to 5×10^{-4} and 5×10^{-3}, respectively. The batch size is set to 256, and the model is trained for 120 epochs. The MaxPool and AvgPool operations each produce 512-dimensional features, yielding a 1024-dimensional global descriptor $\mathbf{g}$ and a 537-dimensional latent feature $\mathbf{h}_i$ in Eq. (6). For Eq. (12), the local receptive field distance threshold τ is set to 10m.

Baselines and Comparisons. To evaluate the performance of ADC, we compare it against various state of the art localization methods. The map-based retrieval group includes TransLoc4D (Peng et al., 2024) and LighterBEV (Liu et al., 2025) (all augmented with registration). In the category of map-free absolute pose regression, we select PosePN++ (Yu et al., 2022), NIDALoc (Yu et al., 2023), and HypLiLoc (Wang et al., 2023) as baselines. For scene coordinate regression, we compare our approach with LightLoc (Li et al., 2025), RALoc (Yang et al., 2025), GTRLoc (Yu et al., 2025) and TempLoc (Zhu et al., 2026).

4.2 COMPARISON WITH STATE-OF-THE-ART METHODS

Results on Boreas dataset. Tab. 1 reports quantitative results across snow, rain, and sunny conditions. Compared with the leading retrieval method TransLoc4D (1.71m/1.10°), ADC achieves 0.98m/0.89°, representing a 42.7% translation error reduction with zero map storage. While APR baselines fail with errors exceeding 4.3m, ADC consistently achieves the lowest error across all five test sequences within the SCR group, reducing the mean translation error by 70.2% over TempLoc (3.29m to 0.98m). Crucially, despite severe airborne clutter in snow sequences, ADC reliably maintains sub-meter precision (≤ 1.00m), demonstrating exceptional weather resilience.

Fig. 3 presents qualitative comparisons on Boreas, where red points denote localization errors $>$ 10m. While baselines like RALoc, GTRLoc, and TempLoc suffer from heavy coordinate drift and dense red outliers in snow-covered areas, ADC maintains a continuous navy-blue trajectory with minimal outliers along the entire 8 km route. This confirms that our hierarchical encoding, adaptive feature recalibration, and rigidity constraints effectively preserve geometric consistency against rain and snow corruptions.

Table 1: **Quantitative results on the Boreas Dataset.** Performance is evaluated under Snow ❄, Rain 🌧 and Sunny ☀ condition using mean localization error [m/°]↓. Best results in **bold**, second best results underlined.

Category	Method	Venue & Year	Boreas Scenes Localization Error[m,°]↓					Mean Error [m,°]↓
			2020-11-26 Snow ❄	2020-12-01 Snow ❄	2021-02-02 Snow ❄	2021-03-02 Sun ☀	2021-04-29 Rain 🌧	
Retrieval	TransLoc4D	24'CVPR	1.57,1.01	1.72,1.14	1.88,1.26	1.87,1.10	1.52,0.99	1.71,1.10
	LighterBEV	25'RAL	1.34,1.02	3.82,2.99	1.30,1.86	1.44,1.14	1.31,**0.79**	1.84,1.56
APR	PosePN++	22'PR	5.57,0.87	7.83,1.66	5.71,0.97	5.75,0.93	5.91,0.90	6.15,1.07
	NIDALoc	23'TITS	7.93,0.91	10.41,1.25	8.23,1.06	7.98,1.02	8.44,0.93	8.60,1.03
	HypLiLoc	23'CVPR	2.18,**0.50**	10.29,1.86	3.17,**0.75**	2.10,**0.57**	3.80,1.01	4.31,0.94
SCR	LightLoc	25'CVPR	2.90,2.17	5.68,2.78	3.01,2.34	2.78,2.06	2.79,2.13	3.43,2.30
	RALoc	25'ICCV	6.49,3.02	9.50,4.20	6.63,3.71	6.25,4.54	6.26,3.22	7.03,3.74
	GTRLoc	25'NeurIPS	7.58,7.59	8.18,7.71	7.65,7.16	7.44,7.38	7.48,7.16	7.67,7.40
	TempLoc	26'ACM MM	3.17,1.14	3.98,1.34	3.15,1.28	3.19,1.17	2.98,1.15	3.29,1.22
	ADC	Ours	**0.96**,0.86	**1.00**,**0.94**	**1.00**,0.92	**0.98**,0.87	**0.95**,0.85	**0.98**,**0.89**

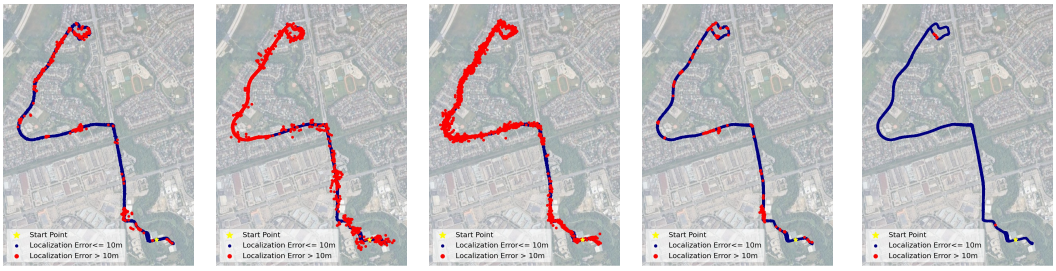

LightLoc(25'CVPR) RALoc(25'ICCV) GTRLoc(25'NeurIPS) TempLoc(26'MM) **ADC(Ours)**

Figure 3: **Visualization Localization Performance on the Boreas Dataset.** Localization estimates with errors >10 m are marked in red. Fewer red points indicate better localization performance.

Results on Oxford dataset. Tab. 2 reports quantitative results on the standard Oxford benchmark under benign weather. While APR frameworks struggle with mean errors around 4m, SCR methods show significantly better accuracy. Our proposed ADC achieves the best performance across all evaluated sequences, yielding an overall mean translation error of 0.59m and a rotation error of 0.90°. Compared with the state-of-the-art baseline TempLoc (0.72m), ADC reduces the mean translation error by 18.1%. This confirms that ADC not only excels in harsh environments but also demonstrates superior generalization and precision in standard clear-weather LiDAR scenarios.

Fig. 4 illustrates the qualitative localization trajectories on the Oxford dataset, where red points indicate localization errors exceeding 3m. Baselines such as LightLoc, RALoc, and GTRLoc produce noticeable red error clusters in turns and complex urban intersections. TempLoc shows improved continuity but still yields visible outliers in open areas. In contrast, ADC delivers a consistently accurate and nearly pure navy-blue trajectory across the entire path. This verifies that our perspective-invariant global encoding and local geometric rigidity constraints provide robust structural representations across diverse urban environments.

4.3 Robustness to Synthetic Rain and Snow

To verify the genuine resilience of ADC to adverse weather rather than potential dataset overfitting, we synthesize rain and snow corruptions following Dong et al. (2023) on the held-out sunny sequence 2021-03-02-13-38. Frozen models are evaluated on identical frames and reference poses across all conditions. Table 3 shows ADC consistently achieves the lowest absolute errors across all conditions. Under severe Snow-L5, ADC maintains $2.61\,\mathrm{m}/1.20°$, compared to $6.01\,\mathrm{m}/3.26°$ for TempLoc and a catastrophic $26.77\,\mathrm{m}/4.55°$ for LightLoc. ADC also registers the smallest relative RE increase (42.9%) at Snow-L5. Despite TempLoc showing a slightly smaller relative TE increase, its absolute error remains significantly higher. These results confirm ADC's robust accuracy advantage under severe perceptual degradation.

Table 2: **Quantitative results on the Oxford Dataset.** We report the mean localization error [m,°]↓. Best results in **bold**, second best results underlined.

Category	Method	Venue & Year	Oxford Scenes Localization Error[m,°]↓				Mean Error [m,°]↓
			15-13-06-37	17-13-26-39	17-14-03-00	18-14-14-42	
APR	PosePN++	22'PR	4.54,1.83	6.44,1.78	4.89,1.55	4.64,1.61	5.13,1.69
	NIDALoc	23'TITS	3.71,1.50	5.40,1.40	3.94,1.30	4.08,1.30	4.28,1.38
	HypLiLoc	23'CVPR	5.03,1.46	4.31,1.43	3.61,1.11	2.61,1.09	3.89,1.27
SCR	LightLoc	25'CVPR	0.82,1.12	0.85,1.07	0.81,1.11	0.82,1.16	0.83,1.12
	RALoc	25'ICCV	1.52,1.28	1.71,1.27	1.37,1.20	1.41,1.23	1.51,1.24
	GTRLoc	25'NeurIPS	0.77,1.02	0.77,1.01	0.67,1.01	0.80,1.07	0.75,1.03
	TempLoc	26'ACM MM	0.74,0.99	0.77,0.99	0.60,0.95	0.77,1.05	0.72,0.99
	ADC	Ours	**0.63,0.91**	**0.61,0.87**	**0.53,0.89**	**0.59,0.95**	**0.59,0.90**

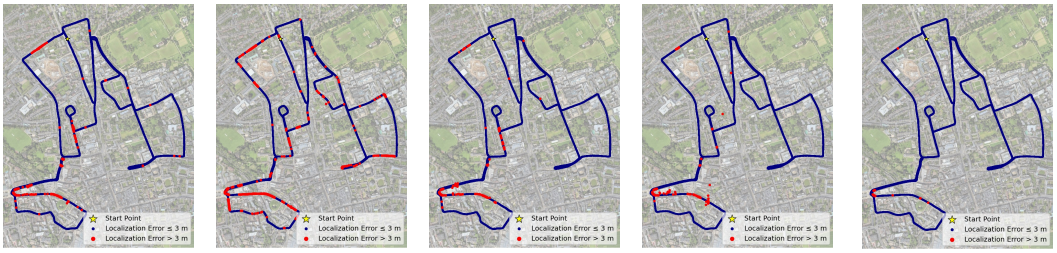

LightLoc(25'CVPR) RALoc(25'ICCV) GTRLoc(25'NeurIPS) TempLoc(26'MM) **ADC(Ours)**

Figure 4: **Visualization Localization Performance on the Oxford Dataset.** Localization estimates with errors >3 m are marked in red. Fewer red points indicate better localization performance.

4.4 Runtime and Storage Consumption

Tab. 4 compares the runtime and storage consumption on the Boreas dataset. Operating at 10 Hz, the system requires an inference latency below 100 ms for real-time localization. ADC achieves an inference latency of 60 ms, satisfying real-time constraints. Furthermore, while map-based methods like Bevplace++ suffer from heavy map storage overhead (348 MB), our map-free ADC requires zero map storage and maintains a compact model footprint of only 72 MB, significantly smaller than RALoc (396 MB) and GTRLoc (412 MB). This confirms the high efficiency and scalability of ADC for practical deployment.

4.5 Ablation Study

Ablation on key modules. Tab. 5 presents the ablation study of the three core components in ADC: Hierarchical Sparse Representation and Perspective-Invariant Global Encoding (**H&P**) in Sec. 3.1, Adaptive Point-wise Feature Recalibration (**APFR**) in Sec. 3.2, and Geometric Consistency Loss (**EdgeLoss**) in Sec. 3.3. Starting from the vanilla baseline without any proposed modules, the localization error reaches 2.50m/2.33° on Boreas and 0.85m/1.14° on Oxford. Integrating H&P provides a drastic performance gain, reducing the mean error to 1.50m/1.96° on Boreas and 0.63m/0.89° on Oxford (a 40.0% and 25.9% translation error reduction, respectively). This verifies that capturing multi-scale sparse geometry while distilling rotation-invariant global priors effectively alleviates viewpoint sensitivity and observation degradation under adverse weather.

When introducing APFR upon H&P, translation error on the standard Oxford dataset further drops to 0.61m, confirming that point-wise recalibration purifies local features and emphasizes stable structural cues. Meanwhile, adding EdgeLoss to H&P regularizes topological consistency to prevent coordinate distortion caused by missing returns. Finally, the full framework integrating all three components achieves the lowest translation error on both datasets, delivering 0.98m/0.89° on the adverse-weather Boreas dataset and 0.59m/0.90° on the Oxford benchmark. Compared with the baseline, the full model achieves overall translation error reductions of 60.8% on Boreas and 30.6% on Oxford, validating the complementary effectiveness of all proposed modules.

Ablation on Perspective-Invariant Global Encoding. Tab. 6 investigates the impact of integrating multiple canonical viewpoints ($k \cdot 90°, k \in \{0, 1, 2, 3\}$) within the perspective-invariant global aggregation module. As the number of viewpoints increases from 1 to 4, localization error on the

Table 3: **Robustness to synthetic rain and snow on Boreas.** We report mean translation / rotation error (TE / RE) and the relative percentage degradation (Δ) from the Sunny condition. L1–L5 indicate corruption severity. Best and second-best values are **boldfaced** and underlined.

Method	Sunny	Rain		Snow			
	Clean	L5	Δ_{L5}	L1	L3	L5	Δ_{L5}
	[m / °] ↓	[m / °] ↓	[% / %] ↓	[m / °] ↓	[m / °] ↓	[m / °] ↓	[% / %] ↓
LightLoc	2.72 / 1.79	4.54 / 2.37	+66.9 / **+32.4**	4.11 / 2.12	4.81 / 2.66	26.77 / 4.55	+884.2 / +154.2
TempLoc	2.36 / 1.14	3.62 / 1.79	**+53.4** / +57.0	3.32 / 1.53	4.20 / 1.98	6.01 / 3.26	**+154.7** / +186.0
ADC (Ours)	**0.96 / 0.84**	**1.57 / 1.15**	+63.5 / +36.9	**1.67 / 1.04**	**1.70 / 1.07**	**2.61 / 1.20**	+171.9 / **+42.9**

Table 4: **Runtime and Storage Consumption.**

Mech.	Method	Boreas Dataset		
		Time (ms↓)	Model (MB↓)	Map (MB↓)
Map-Based	TransLoc4D(24'CVPR)	29	13	42
	Bevplace++(25'TRO)	191	55	348
Map-Free	HypLiLoc(23'CVPR)	29	200	0
	LightLoc(25'CVPR)	45	71	0
	RALoc(25'ICCV)	49	396	0
	GTRLoc(25'NeurIPS)	46	412	0
	ADC(Ours)	60	72	0

Table 5: **Ablation on Key Modules.** Hierarchical Sparse Representation and Perspective-Invariant Global Encoding (**H&P**), Adaptive Point-wise Feature Recalibration (**APFR**), and Geometric Consistency Loss (**EdgeLoss**).

H&P	APFR	EdgeLoss	Mean Error[m,°]↓	
			Boreas	Oxford
			2.50,2.33	0.85,1.14
✓			1.50,1.96	0.63,0.89
✓	✓		1.48,1.26	0.61,1.00
✓		✓	1.13,1.02	0.63,0.90
✓	✓	✓	0.98,0.89	0.59,0.90

Table 6: **Ablation on Perspective-Invariant Global Encoding.**

k*90°				Dataset Mean Error[m/°]↓	
0	1	2	3	Boreas	Oxford
✓				5.09,4.70	0.62,0.92
✓	✓			1.61,1.82	0.66,0.94
✓	✓	✓		1.23,1.05	0.61,0.92
✓	✓	✓	✓	0.98,0.89	0.59,0.90

Table 7: **Ablation on Geometric Consistency Loss.** Mean translation error [m]↓ under different geometric consistency loss weight λ_e.

Dataset Mean	Parameter λ_e				
Error[m]↓	0.01	0.05	0.1	0.5	1
Boreas	1.22	1.14	**0.98**	2.38	5.17
Oxford	0.65	0.84	**0.59**	0.65	0.89

adverse-weather Boreas dataset drops significantly from 5.09m/4.70° to 0.98m/0.89°, achieving an 80.7% reduction in translation error. On the standard Oxford dataset, the full four-view configuration likewise yields the best precision of 0.59m/0.90°. These results demonstrate that aggregating multi-perspective features across the discrete rotation group effectively eliminates sensor heading ambiguity and compensates for viewpoint-dependent point dropouts in severe weather.

Ablation on Geometric Consistency Loss. Tab. 7 reports the impact of different EdgeLoss weight coefficients λ_e on localization accuracy. Balancing coordinate regression and structural rigidity is crucial for adverse-weather localization. Increasing λ_e from 0.01 to 0.1 consistently improves performance on Boreas, with $\lambda_e = 0.1$ yielding the lowest translation errors on both Boreas (0.98 m) and Oxford (0.59 m). This confirms that enforcing local distance consistency effectively imposes an implicit structural skeleton on degraded LiDAR scans, successfully suppressing coordinate distortion caused by rain or snow noise and point dropouts. However, excessively large weights such as $\lambda_e = 0.5$ and $\lambda_e = 1.0$ severely degrade performance, with translation error increasing to 5.17m on Boreas, as an overly rigid constraint over-regularizes the network and hinders point-level feature optimization. Consequently, $\lambda_e = 0.1$ is adopted as the optimal trade-off in all experiments.

5 Conclusion

This paper presents **ADC**, a robust LiDAR localization framework designed for adverse weather conditions. By integrating hierarchical sparse encoding, adaptive point-wise feature recalibration, and local geometric rigidity constraints, ADC effectively suppresses airborne rain/snow clutter and resolves coordinate drift. Extensive evaluations confirm that ADC achieves state-of-the-art accuracy across both adverse-weather Boreas and standard Oxford benchmarks, delivering a scalable, map-free, and weather-resilient localization solution for autonomous driving.

AI USE STATEMENT

We acknowledge the use of generative AI tools to assist in editing the manuscript language, supporting parts of the experimental workflow, aiding the preliminary interpretation of empirical data, and offering structural recommendations for figure, table, and layout organization to improve clarity. The authors maintained complete scientific oversight at every stage. All AI-generated suggestions were critically reviewed, and all experimental implementations, empirical evaluations, and literature citations were manually inspected and validated by the research team. The authors accept full accountability for the accuracy, integrity, and conclusions of the work presented in this paper.

REFERENCES

Sheng Ao, Qingyong Hu, Hanyun Wang, Kai Xu, and Yulan Guo. Buffer: Balancing accuracy, efficiency, and generalizability in point cloud registration. In *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 1255–1264. IEEE, 2023.

Keenan Burnett, David J Yoon, Yuchen Wu, Andrew Z Li, Haowei Zhang, Shichen Lu, Jingxing Qian, Wei-Kang Tseng, Andrew Lambert, Keith YK Leung, et al. Boreas: A multi-season autonomous driving dataset. *The International Journal of Robotics Research*, 42(1-2):33–42, 2023.

Yidong Chen, Qi Li, Yuyang Yang, Wen Li, Sheng Ao, and Cheng Wang. Unleashing the power of data generation in one-pass outdoor lidar localization. In *ACM Multimedia*, pp. 199–208, 2025.

Christopher Choy, Junha Lee, René Ranftl, Jaesik Park, and Vladlen Koltun. High-dimensional convolutional networks for geometric pattern recognition. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 11224–11233. IEEE, 2020.

Ondřej Chum and Jiří Matas. Optimal randomized ransac. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 30(8):1472–1482, 2008.

Junyuan Deng, Xinyi Wu, Yongxing Yang, Congchao Zhu, Song Wang, and Zhenyao Wu. Acquire and then adapt: Squeezing out text-to-image model for image restoration. In *2025 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 23195–23206. IEEE, 2025.

Yinpeng Dong, Caixin Kang, Jinlai Zhang, Zijian Zhu, Yikai Wang, Xiao Yang, Hang Su, Xingxing Wei, and Jun Zhu. Benchmarking robustness of 3d object detection to common corruptions in autonomous driving. In *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 1022–1032. IEEE, 2023.

Zhixing Hou, Yan Yan, Chengzhong Xu, and Hui Kong. Hitpr: Hierarchical transformer for place recognition in point cloud. In *2022 International Conference on Robotics and Automation (ICRA)*, pp. 2612–2618. IEEE, 2022.

Xudong Jiang, Fangjinhua Wang, Silvano Galliani, Christoph Vogel, and Marc Pollefeys. R-score: Revisiting scene coordinate regression for robust large-scale visual localization. In *2025 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 11536–11546. IEEE, 2025.

Wen Li, Shangshu Yu, Cheng Wang, Guosheng Hu, Siqi Shen, and Chenglu Wen. Sgloc: Scene geometry encoding for outdoor lidar localization. In *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 9286–9295. IEEE, 2023.

Wen Li, Yuyang Yang, Shangshu Yu, Guosheng Hu, Chenglu Wen, Ming Cheng, and Cheng Wang. Diffloc: Diffusion model for outdoor lidar localization. In *2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 15045–15054. IEEE, 2024.

Wen Li, Chen Liu, Shangshu Yu, Dunqiang Liu, Yin Zhou, Siqi Shen, Chenglu Wen, and Cheng Wang. Lightloc: Learning outdoor lidar localization at light speed. In *2025 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 6680–6689. IEEE, 2025.

Binhong Liu, Tao Yang, Haoji Cao, Shuqi Fu, Yangwang Fang, and Zhi Yan. Lighterbev: Lidar global localization meets online learning. *IEEE Robotics and Automation Letters*, 11(2):1170–1177, 2025.

Zhe Liu, Shunbo Zhou, Chuanzhe Suo, Peng Yin, Wen Chen, Hesheng Wang, Haoang Li, and Yun-Hui Liu. Lpd-net: 3d point cloud learning for large-scale place recognition and environment analysis. In *Proceedings of the IEEE/CVF international conference on computer vision*, pp. 2831–2840, 2019.

Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017.

Fan Lu, Guang Chen, Yinlong Liu, Lijun Zhang, Sanqing Qu, Shu Liu, and Rongqi Gu. Hregnet: A hierarchical network for large-scale outdoor lidar point cloud registration. In *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 15994–16003. IEEE, 2021.

Weixin Lu, Yao Zhou, Guowei Wan, Shenhua Hou, and Shiyu Song. L3-net: Towards learning based lidar localization for autonomous driving. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 6382–6391. IEEE, 2019.

Lun Luo, Si-Yuan Cao, Xiaorui Li, Jintao Xu, Rui Ai, Zhu Yu, and Xieyuanli Chen. Bevplace++: Fast, robust, and lightweight lidar global localization for autonomous ground vehicles. *IEEE transactions on Robotics*, 41:4479–4498, 2025.

Will Maddern, Geoffrey Pascoe, Chris Linegar, and Paul Newman. 1 year, 1000 km: The oxford robotcar dataset. *The International Journal of Robotics Research*, 36(1):3–15, 2017.

Xindong Mao, Hang Li, Yuchen Wu, Jiahe Li, Xiao Bai, and Jin Zheng. Color: The devil is in scene coordinate regression for large-scale visual localization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 12206–12216, 2026.

Guohao Peng, Heshan Li, Yangyang Zhao, Jun Zhang, Zhenyu Wu, Pengyu Zheng, and Danwei Wang. Transloc4d: Transformer-based 4d radar place recognition. In *2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 17595–17605. IEEE, 2024.

Torsten Sattler, Bastian Leibe, and Leif Kobbelt. Efficient & effective prioritized matching for large-scale image-based localization. *IEEE transactions on pattern analysis and machine intelligence*, 39(9):1744–1756, 2016.

Hongming Shen, Xun Chen, Yulin Hui, Zhenyu Wu, Wei Wang, Qiyang Lyu, Tianchen Deng, and Danwei Wang. Unilgl: learning uniform place recognition for fov-limited/panoramic lidar global localization. *IEEE Transactions on Robotics*, 2026.

Leslie N Smith and Nicholay Topin. Super-convergence: Very fast training of neural networks using large learning rates. In *Artificial intelligence and machine learning for multi-domain operations applications*, volume 11006, pp. 369–386. SPIE, 2019.

Mikaela Angelina Uy and Gim Hee Lee. Pointnetvlad: Deep point cloud based retrieval for large-scale place recognition. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 4470–4479. IEEE, 2018.

Sijie Wang, Qiyu Kang, Rui She, Wei Wang, Kai Zhao, Yang Song, and Wee Peng Tay. Hypliloc: Towards effective lidar pose regression with hyperbolic fusion. In *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 5176–5185. IEEE, 2023.

Wei Wang, Bing Wang, Peijun Zhao, Changhao Chen, Ronald Clark, Bo Yang, Andrew Markham, and Niki Trigoni. Pointloc: Deep pose regressor for lidar point cloud localization. *IEEE Sensors Journal*, 22(1):959–968, 2021.

Bochun Yang, Zijun Li, Wen Li, Zhipeng Cai, Chenglu Wen, Yu Zang, Matthias Muller, and Cheng Wang. Lisa: Lidar localization with semantic awareness. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 15271–15280, 2024.

Yuyang Yang, Wen Li, Sheng Ao, Qingshan Xu, Shangshu Yu, Yu Guo, Yin Zhou, Siqi Shen, and Cheng Wang. Raloc: Enhancing outdoor lidar localization via rotation awareness. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 3304–3313, 2025.

Huan Yin, Xuecheng Xu, Sha Lu, Xieyuanli Chen, Rong Xiong, Shaojie Shen, Cyrill Stachniss, and Yue Wang. A survey on global lidar localization: Challenges, advances and open problems. *International Journal of Computer Vision*, 132(8):3139–3171, 2024.

Shangshu Yu, Cheng Wang, Chenglu Wen, Ming Cheng, Minghao Liu, Zhihong Zhang, and Xin Li. Lidar-based localization using universal encoding and memory-aware regression. *Pattern Recognition*, 128:108685, 2022.

Shangshu Yu, Xiaotian Sun, Wen Li, Chenglu Wen, Yunuo Yang, Bailu Si, Guosheng Hu, and Cheng Wang. Nidaloc: Neurobiologically inspired deep lidar localization. *IEEE Transactions on Intelligent Transportation Systems*, 25(5):4278–4289, 2023.

Shangshu Yu, Wen Li, Xiaotian Sun, Zhimin Yuan, Xin Wang, Sijie Wang, Rui She, and Cheng Wang. Gtr-loc: Geospatial text regularization assisted outdoor lidar localization. *Advances in Neural Information Processing Systems*, 38:136468–136490, 2025.

Minghang Zhu, Zhijing Wang, Yuxin Guo, Wen Li, Sheng Ao, and Cheng Wang. From single scan to sequential consistency: A new paradigm for lidar relocalization. *arXiv preprint arXiv:2602.03198*, 2026.

A BOREAS DATASET AND CROSS-WEATHER EVALUATION

Representative environmental conditions. Figure 5 supplements the dataset description in Section 4.1 with representative scenes from the Boreas traversals used in our experiments. The examples illustrate snow accumulation along roads and surrounding surfaces, clear-weather urban scenes, and wet road surfaces during rain. Such conditions can change the geometric evidence available for LiDAR localization: accumulated snow may obscure scene boundaries, while precipitation can introduce transient returns and attenuate surface measurements. The RGB images provide environmental context rather than a direct visualization of these LiDAR effects; ADC uses LiDAR observations for localization.

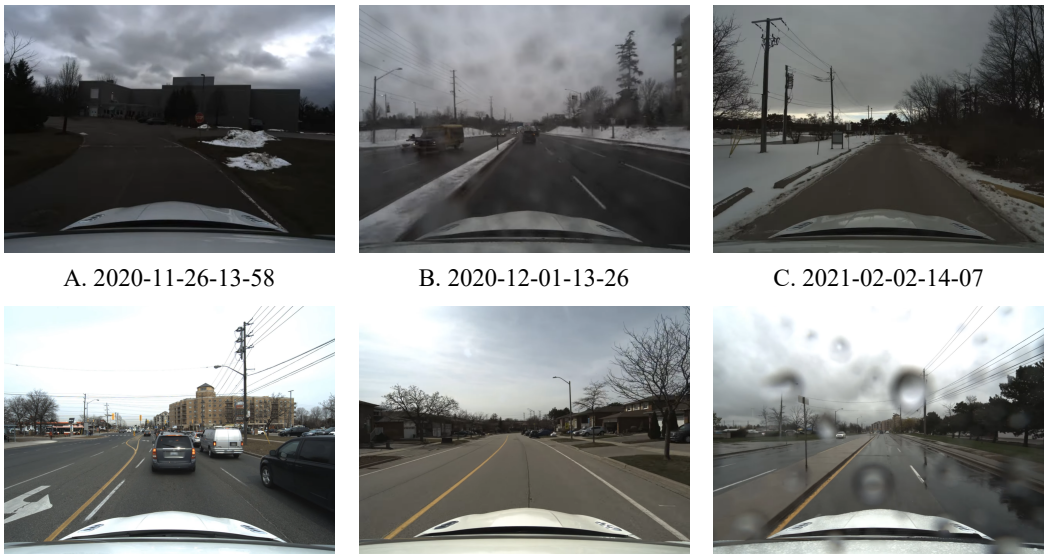

A. 2020-11-26-13-58 B. 2020-12-01-13-26 C. 2021-02-02-14-07

D. 2021-03-23-12-43 E. 2021-03-30-14-23 F. 2021-04-29-15-55

Figure 5: **Representative scenes from the Boreas dataset.** Panels (A)–(C) show examples from snow test sequences, (D)–(E) show clear/sunny training sequences, and (F) shows a rainy test sequence. Sequence identifiers are provided below each panel. RGB images are used only to illustrate environmental conditions.

Relationship to the evaluation protocol. The complete split comprises four clear/sunny training traversals and five test traversals covering snow, sun, and rain, as detailed in Section 4.1. The figure shows three snow test sequences (A–C: 2020-11-26, 2020-12-01, and 2021-02-02), two clear/sunny training sequences (D–E: 2021-03-23 and 2021-03-30), and the rainy test sequence (F: 2021-04-29). The other two training traversals and the sunny test traversal are not displayed. Weather categories follow the sequence-level assignments in the main paper and do not imply visible precipitation in every frame. This split enables evaluation on rain and snow traversals after training on clear-weather observations, with the separate sunny test traversal providing a reference condition.

Interpretation of cross-weather performance. As reported in Table 1, ADC's sequence-level mean translation errors range from 0.95 to $1.00\,\mathrm{m}$, and mean rotation errors range from 0.85° to 0.94°, across the five test traversals. The comparable errors on sunny, rainy, and snowy sequences support stable localization performance under the evaluated weather changes. Since the traversals revisit the same route, this evidence concerns generalization across weather conditions and acquisition sessions within a known environment. The sequences also differ in traffic and other scene conditions, so the results reflect robustness to naturally occurring changes across traversals rather than an isolated measurement of the effect of precipitation.

B ADDITIONAL IMPLEMENTATION DETAILS

We provide architectural and supervision details that complement the main paper. All dimensions below refer to the Boreas configuration with a voxel size of $0.5\,\mathrm{m}$.

B.1 GLOBAL CONTEXT AND SPATIAL CONDITIONING

The encoder produces 512-dimensional local features. The aggregation defined in the main paper yields a 1024-dimensional global descriptor. As shown in Figure 6(a), the regional predictor uses linear layers of widths $1024 \to 512 \to 256 \to 25$, with batch normalization, ReLU, and dropout ($p = 0.5$) after the first two layers. Regional supervision is obtained by clustering training positions.

The output logits are ℓ_2-normalized to form a 25-dimensional response g^* and broadcast to individual points. During regression training, independent per-point Gaussian noise with standard deviation $\sigma = 0.1$ is added to each broadcast response. As illustrated in Figure 6(b), concatenation with local features produces the 537-dimensional point feature h. Noise and dropout are disabled at evaluation. This conditioning supplies regional context when rain or snow makes local geometry ambiguous; prior perturbation encourages tolerance to imperfect regional predictions.

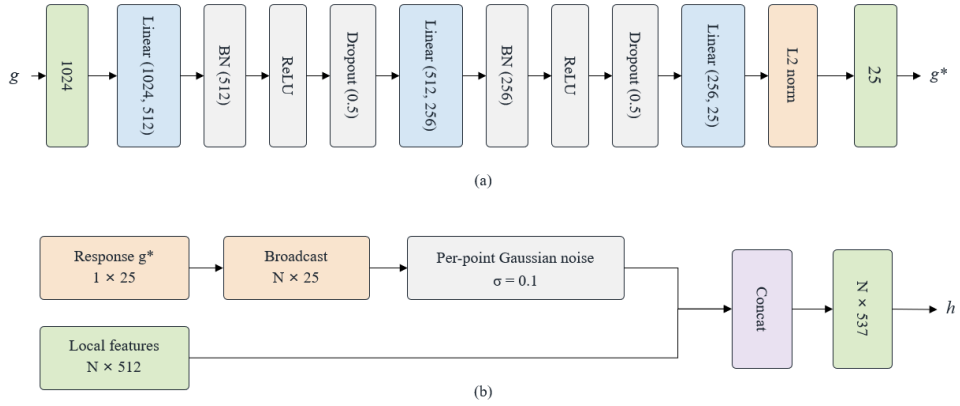

Figure 6: **Spatial prior conditioning.** (a) The regional predictor generates a normalized 25-dimensional regional response g^*. (b) This response is broadcast and concatenated with 512-dimensional local features to output the conditioned feature h. Per-point Gaussian perturbation is used only during training. BN denotes batch normalization.

B.2 POINT-WISE RECALIBRATION AND RESIDUAL REGRESSION

Given an input feature u_i, APFR applies $537 \to 67 \to 537$ linear mappings, with ReLU between them and a final sigmoid activation to produce a gate S_i (Figure 7). The resulting gate multiplies the identity branch element-wise to produce the refined feature $\hat{u}_i$. Parameters are shared across points, while gates are computed per point without cross-point pooling. This permits spatially varying feature weighting within a scan, accommodating the non-uniform influence of weather-related corruption without requiring precipitation-point annotations.

As detailed in Figure 8, the regression network processes the conditioned feature h_i through a residual block architecture. The residual branch follows $537 \to 537 \to 1074 \to 537$, with ReLU after every linear layer. The residual addition is immediately followed by the APFR module. Finally, the output head maps the refined representation to the predicted scene coordinate (x, y, z) using linear layers of $537 \to 537 \to 1074 \to 3$, with ReLU only after the first two layers. No batch normalization or dropout is used in this regression network. Residual paths preserve input features while successive recalibration stages refine the representation used for coordinate prediction.

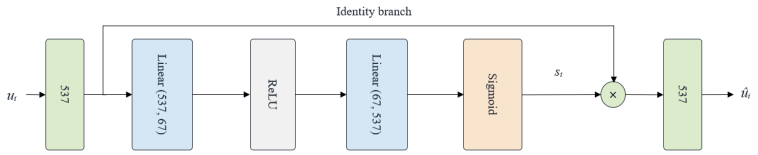

Figure 7: **Adaptive point-wise feature recalibration.** A 67-dimensional bottleneck generates channel gates S_i independently for each point, which scale the identity input u_i via element-wise multiplication to yield $\hat{u}_i$.

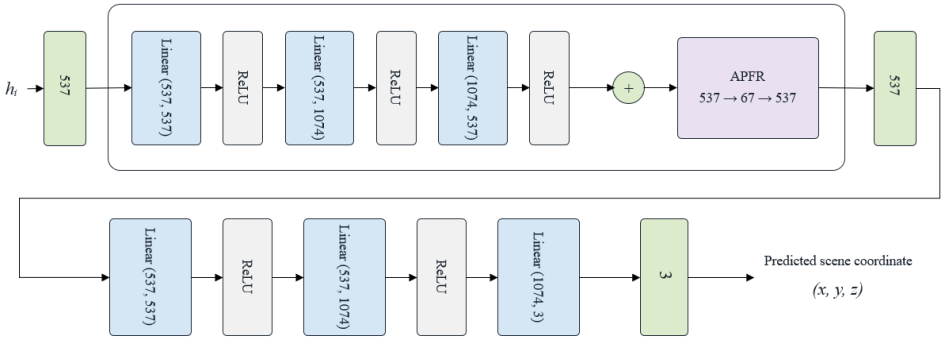

Figure 8: **Residual regression prediction network.** The input feature h_i is processed through a residual block featuring linear refinement and APFR, followed by an output head that yields the predicted three-dimensional scene coordinate (x, y, z).

B.3 Coordinate supervision and geometric consistency

Scene-coordinate targets are centered using the training translation mean and aligned with sparse outputs by average pooling with kernel size and stride eight. The stored offset restores the original world frame.

EdgeLoss uses within-scan point pairs with target distances in $(0, 10\,\mathrm{m})$. Distance discrepancies are averaged over valid pairs in each scan, then over scans containing valid pairs; an empty set contributes zero. Its weight is $\lambda_e = 0.1$. Complementing coordinate supervision, this constraint discourages local stretching and contraction, encouraging structural coherence under incomplete or weather-affected observations.

At inference, sensor-frame points and predicted scene coordinates are used for robust pose estimation with a $2\,\mathrm{m}$ inlier threshold. EdgeLoss is used only during training; test reference poses are used only for evaluation.

C Localized ablation analysis under adverse weather

To complement the aggregate ablations presented in the main manuscript, we examine how the components of ADC affect spatial consistency within a challenging Boreas segment. Figure 9 uses 100 consecutive frames (7875–7974, zero-based indices) from the snow-labeled test sequence 2021-02-02-14-07. This segment is deliberately selected from windows with relatively large vanilla localization errors to visually demonstrate the progressive improvements our modules yield in severe adverse weather. Columns follow the cumulative configurations: vanilla, H&P, H&P+APFR, and H&P+APFR+EdgeLoss.

For each configuration, we transform the same input scans using its predicted 6-DoF poses and accumulate them in a common frame. An accumulation using reference poses is overlaid in grey, allowing pose-induced misalignment to be clearly distinguished from the actual observed scene

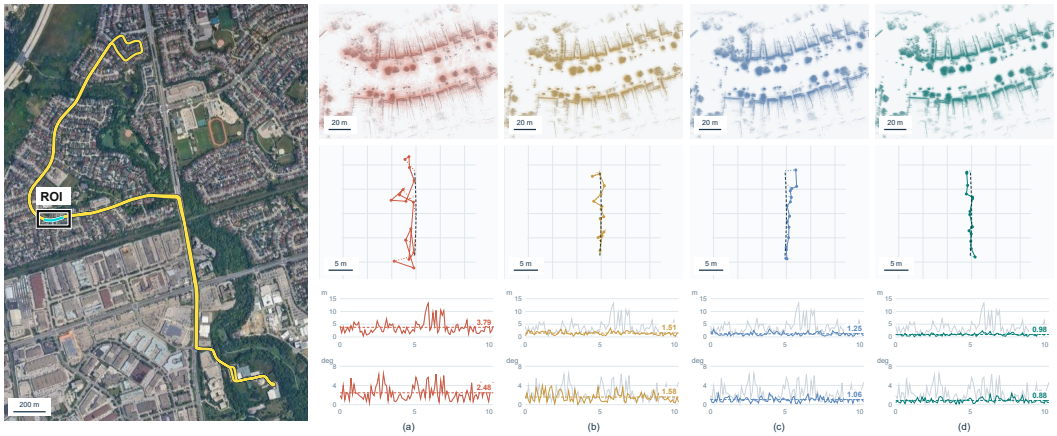

Figure 9: **Spatial effects of progressive module integration on a snow-labeled Boreas segment.** Left: the reference trajectory on an aerial image, with the selected region marked as ROI. Columns (a)–(d): vanilla, H&P, H&P+APFR, and full ADC. First row: 100-frame point-cloud accumulations using reference poses (grey) and predicted poses (colored). Second row: a magnified trajectory over frames 7900–7914, with dashed ground truth and dotted frame-matched displacements. Third and fourth rows: raw translation and rotation errors over the selected segment; horizontal dashed lines and adjacent numbers denote their means. Grey curves repeat vanilla, and time is measured in seconds. All columns use identical inputs, display scales, and rendering settings. The segment is an illustrative improvement case rather than a dataset-wide estimate.

structure. Point selection, spatial bounds, and density normalization are kept identical across all configurations without any additional registration refinement. The local trajectory and error traces provide complementary views of the same segment.

The vanilla configuration produces severely dispersed structural returns and significant departures from the reference trajectory, suffering mean translation and rotation errors of $3.79\,\mathrm{m}$ and $2.48°$. Introducing the H&P module reduces these errors to $1.51\,\mathrm{m}$ and $1.58°$. Visually, this is evidenced by a more compact alignment of repeated observations, proving that the perspective-invariant global encoding successfully provides a stable spatial prior against weather-induced point dropouts. Adding the APFR module further corrects the trajectory, decreasing errors to $1.25\,\mathrm{m}$ and $1.06°$. This reduction visually manifests as cleaner structural aggregations, demonstrating APFR's ability to adaptively filter out transient airborne snow clutter. Finally, integrating EdgeLoss completes the ADC model, achieving the lowest errors of $0.98\,\mathrm{m}$ and $0.88°$. The visualization explicitly proves this success: the reconstructed environmental structures become remarkably crisp, and the predicted local trajectory seamlessly aligns with the ground truth. This progressive visual refinement provides compelling, concrete evidence that each proposed module fundamentally enhances the model's structural awareness and geometric resilience in harsh weather scenarios.

These visual observations strictly align with the complementary roles described in the main paper: H&P supplies robust multi-scale geometry and global context, APFR mitigates noisy local feature responses, and EdgeLoss enforces topological consistency in coordinate prediction. By demonstrating how localization gains translate into reduced spatial dispersion and highly consistent pose estimates during a snowy traversal, this visualization substantiates the theoretical claims behind ADC's weather-resilient design. The broader dataset-wide rain-and-snow evaluation is provided in Table 1.

D Additional Analysis under Synthetic Rain and Snow

To complement the mean-error analysis presented in the main manuscript, we report the median translation and rotation errors under identical synthetic rain and snow conditions. Median errors offer a robust complementary measure of typical localization accuracy, as they are less sensitive to occasional large outliers. We maintain the exact evaluation protocol: synthetic weather corruptions (Dong et al. (2023)) are applied to the same short-frame window from the held-out sunny

Table 8: **Median localization errors under synthetic rain and snow on Boreas.** We report median translation / rotation error (TE / RE) and the relative percentage degradation (Δ) from the Sunny condition. For each metric, relative degradation is computed as $\Delta = 100 \times (\widetilde{E}_{\mathrm{w}} - \widetilde{E}_{\mathrm{s}})/\widetilde{E}_{\mathrm{s}}$, where $\widetilde{E}_{\mathrm{s}}$ and $\widetilde{E}_{\mathrm{w}}$ denote the median errors under the Sunny and the corresponding weather-corrupted conditions, respectively. L1–L5 indicate corruption severity. Lower is better for all metrics; best and second-best values are **boldfaced** and underlined.

Method	**Sunny**	**Rain**		**Snow**			
	Clean	L5	Δ_{L5}	L1	L3	L5	Δ_{L5}
	[m / °] ↓	[m / °] ↓	[% / %] ↓	[m / °] ↓	[m / °] ↓	[m / °] ↓	[% / %] ↓
LightLoc	1.78 / 1.41	3.31 / 2.01	+86.0 / +42.6	2.73 / 1.85	4.09 / 2.09	6.71 / 2.47	+277.0 / +75.2
TempLoc	2.13 / 0.96	3.26 / 1.40	+53.1 / +45.8	2.97 / 1.26	3.82 / 1.60	5.44 / 2.45	+155.4 / +155.2
ADC (Ours)	**0.86 / 0.70**	**0.97 / 0.92**	**+12.7 / +31.4**	**0.90 / 0.90**	**1.00 / 0.89**	**1.06 / 1.02**	**+23.3 / +45.7**

Boreas sequence (`2021-03-02-13-38`). All models operate with frozen weights, and both the selected test frames and reference ground-truth poses remain strictly consistent across all evaluated conditions.

Table 8 demonstrates that ADC consistently achieves the lowest median TE and RE across all reported conditions, alongside the smallest relative performance drops under severe Rain-L5 and Snow-L5. Specifically, under Snow-L5, ADC maintains a highly accurate median estimate of $1.06\,\mathrm{m}/1.02°$, significantly outperforming TempLoc ($5.44\,\mathrm{m}/2.45°$) and LightLoc ($6.71\,\mathrm{m}/2.47°$). Furthermore, ADC's median TE increases by a mere 23.3% relative to the Sunny baseline, in stark contrast to the 155.4% and 277.0% spikes observed in competing methods. While median errors inherently mask the high-error tail and should be interpreted alongside the mean errors detailed in the main text, these findings strongly corroborate ADC's fundamental accuracy advantage and robust generalization in typical adverse-weather scenarios.

To complement the error statistics, Figure 10 visualizes the point cloud accumulations under clear, Rain-L5, and Snow-L5 conditions. For each method and condition, we transform the points from the selected scan sequence into the world coordinate frame using the method's predicted poses, accumulating them into a bird's-eye-view (BEV) density map. For a point $\mathbf{p}_{t,i}$ in the sensor frame, the transformation is defined as $\hat{\mathbf{x}}_{t,i} = \hat{\mathbf{R}}_t \mathbf{p}_{t,i} + \hat{\mathbf{t}}_t$. This visualization intuitively demonstrates how localization errors impact the spatial alignment of repeated observations.

Under synthetic rain and snow, the competing methods exhibit broadened boundaries and structural ghosting, which is particularly pronounced for TempLoc under Snow-L5. This ghosting is a direct result of misalignment caused by pose estimation errors. Although slight dispersion remains under Snow-L5, ADC successfully maintains relatively compact boundaries and a consistent spatial layout across all three conditions. Coupled with the mean and median error statistics, these visual observations further corroborate ADC's ability to deliver highly consistent and robust localization under the evaluated synthetic corruptions.

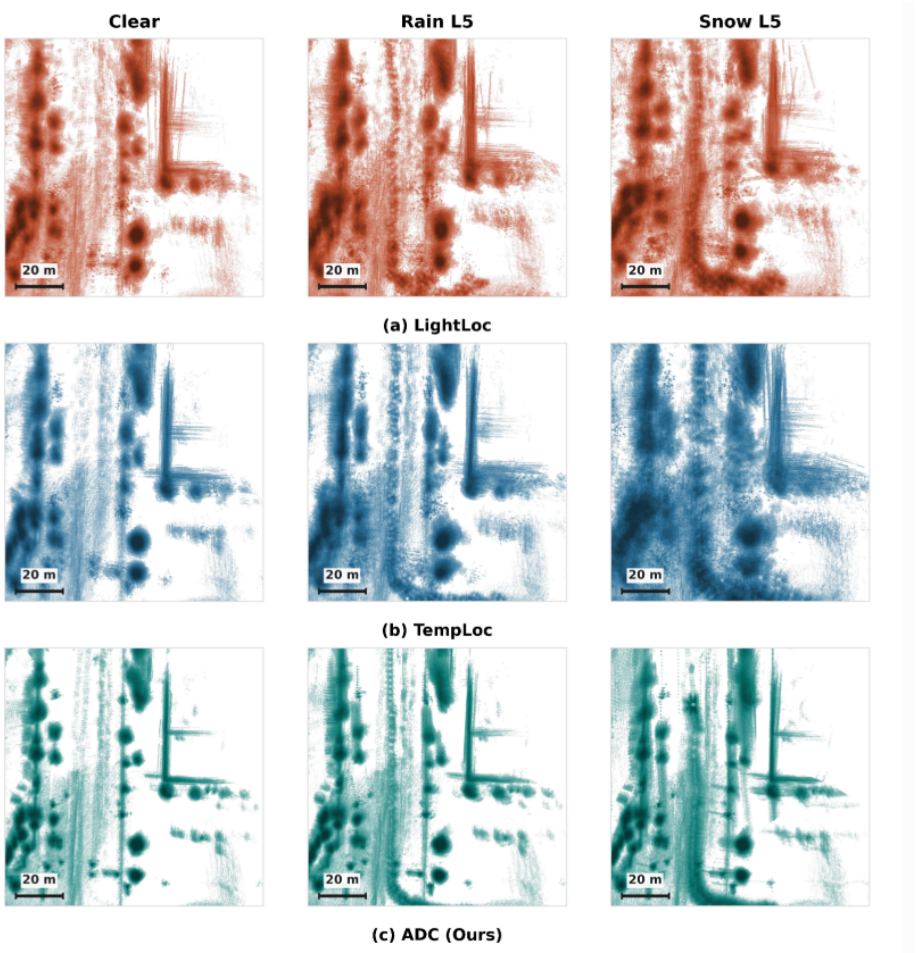

Figure 10: **BEV point cloud accumulations under synthetic rain and snow.** Rows correspond to (a) LightLoc, (b) TempLoc, and (c) ADC; columns represent clear inputs, Rain-L5, and Snow-L5 conditions. Each panel accumulates point clouds into the world frame using the poses predicted by the respective method under that condition. Point density is displayed with a single-hue color scale on a white background; each scale bar represents $20\,\mathrm{m}$. Under weather corruptions, the competing methods exhibit dispersed or ghosted structures, whereas ADC retains highly consistent structural outlines across all three conditions.