

Fast Point Ranking - Robust Cloud Voxelization and Denoising for Lidar Odometry and Mapping in Adverse Weather Conditions

Eugeniu Vezeteu^{1,2}, Heikki Hyyti¹, Ville Kyrki², Juha Hyypä¹

¹ Finnish Geospatial Research Institute (FGI), 02150 Espoo, Finland - eugeniu.vezeteu@maanmittauslaitos.fi

² Electrical Engineering and Automation, Aalto University, 02150 Espoo, Finland

Keywords: Adverse weather, lidar odometry, localization, mapping, denoising

Abstract

Lidar Odometry (LO) is crucial for autonomous navigation, forming the foundation for simultaneous localization and mapping, and providing essential feedback for control systems. Adverse weather conditions, however, introduce false readings, missing echoes, and noise to lidar measurements, severely degrading point cloud quality and compromising LO effectiveness. This study proposes Fast Point Ranking (FPR), a technique that effectively minimizes the impact of adverse weather effects during registration and map denoising via a robust rank-based point cloud voxelization. Experiments on the real-world KITTI-360 and the novel, openly shared Adverse-Weather-KITTI-360 dataset demonstrate that FPR significantly enhances localization accuracy in adverse weather, providing up to 10 m smaller root mean square errors in positioning. Furthermore, FPR shows increased resilience to adverse weather, maintaining consistent localization accuracy despite the weather conditions.

1. Introduction

Lidar Odometry (LO) is an essential technology for any autonomous robot, drone, or vehicle in various outdoor environments, serving as a foundation for Simultaneous Localization and Mapping (SLAM), and providing crucial feedback for control systems. Current LO methods primarily depend on variants of Iterative Closest Point (ICP) (Censi, 2008, Vizzo et al., 2023) or Normal Distributions Transform (NDT) (Magnusson, 2009), with differences mainly in their cost functions. However, the performance of these methods is compromised by adverse weather conditions, characterized by data sparsity and outlier measurements induced by fog, rain, or snowfall. The outlier points can distort the voxel mean in point cloud voxelization or be mistakenly selected for registration, while valid points may be discarded.

The challenge of outliers is commonly addressed by applying robust kernel functions (Kim and Scott, 2012). However, when outliers are abundant, this approach alone may prove insufficient. Traditional methods for scan pre-cleaning, such as Statistical Outlier Removal (SOR), Radius Outlier Removal (ROR), and Dynamic Radius Outlier Removal (DROR) (Prio et al., 2022), filter out incorrect measurements based on local point density. Additionally, learning-based methods (Seppänen et al., 2022, Heinzler et al., 2020, Charron et al., 2018, Seppänen et al., 2023), perform segmentation or multi-echo filtering for outliers. However, these techniques have limitations, as they may remove valid data points or fail to eliminate outliers completely due to their reliance on threshold parameters or binary classification of points.

In this study, we present Fast Point Ranking (FPR), an extension of the KISS-ICP framework (Vizzo et al., 2023), designed to improve LiDAR odometry (LO) robustness in adverse weather conditions. Unlike traditional methods that rely on pre-cleaning scans, FPR directly addresses weather-induced noise during the voxelization stage by ranking points based on neighborhood density and range consistency in a cylindrical range image representation. We hypothesize that higher-ranked




Figure 1. **Fast Point Ranking (FPR):** The foggy point cloud (a) (coloured by the z-axis) is projected on the cylindrical image (b) for the fast nearest neighbour search. Next, the point ranking is computed and displayed on (c) (coloured by the rank). Red pixels on the cylindrical image indicate missing points or points with low ranks.

points are less likely to originate from adverse weather effects (see Fig. 1) and are more likely to have valid data associations with the map. This ranking is then used to perform robust scan voxelization by selecting the best-ranked point in each voxel, leading to improved LO performance. Additionally, we apply the same ranking strategy for map denoising, updating the map only with top-ranked points to reduce the integration of weather-corrupted measurements. We evaluate our method on the KITTI-360 (Liao et al., 2022) and Adverse-Weather-KITTI-360 (Vezeteu, 2024) datasets, where adverse weather conditions (fog, rain, and snowfall) are simulated using point-wise augmentations. Results show that FPR enhances localization accuracy under various weather conditions compared to the state-of-the-art KISS-ICP baseline. In this work, we address the following experimental research questions:

- How does LO perform under adverse weather conditions using Adverse-Weather-KITTI-360 (Vezeteu, 2024) and KITTI-360 (Liao et al., 2022) datasets?
- What are the accuracy and robustness of plain KISS-ICP and KISS-ICP with FPR voxelization in challenging weather scenarios?

Our main contributions are:

- A simple and efficient rank-based scan voxelization that serves both LO and map denoising in adverse weather scenarios. By selecting the most suitable point from each voxel, we demonstrate a reduction of the adverse weather impact.
- The proposed method does not require any scan cleaning for LO to work. It avoids the use of weather-dependent methods, parameters, or thresholds for cleaning the point cloud prior to the registration.
- We extensively test the proposed method with a novel, openly shared Adverse-Weather-KITTI-360 dataset (Vezeteu, 2024) that provides point clouds attenuated by simulated fog, rain, and snowfall.
- We share an open-source implementation on GitHub¹.

The remaining of this paper is organised as follows: we review the relevant literature in Section 2 and present the key components of the system in Section 3. The experimental setup and results are described in Section 4 and the discussions in Section 5. Finally, Section 6 concludes the paper.

2. Related Work

In this section, we review key lidar odometry (LO) and point cloud cleaning methods, focusing on the most relevant approaches.

2.1 Lidar odometry

LO is an extensively explored topic. It aims to compute the trajectory of a mobile sensor by sequentially aligning ego-centric point clouds (Vizzo et al., 2023). Most scan registration methods, like variants of Iterative Closest Point (ICP) (Censi, 2008) (e.g., LOAM (Zhang and Singh, 2014), F-LOAM (Wang et al.,

2021), LegoLOAM (Shan and Englot, 2018), G-ICP (Kuramachi et al., 2015)), use point-to-point, point-to-line, or point-to-plane cost functions to estimate the alignment. Alternatively, methods such as NDT (Magnusson, 2009) adopt a probabilistic approach and minimise a point-to-distribution cost function. To achieve computational efficiency or enable real-time processing of large point clouds, techniques such as voxelization or down-sampling are frequently applied. A common approach is voxel-mean, where an average of points within a voxel is computed. Another method involves keeping only K points per voxel and discarding the rest.

Choosing the best LO method hinges on the particular problem, as these methods are quite similar and highly dependent on their parameters and data quality. Also, these techniques commonly encounter the same challenges with sparse point clouds, especially in adverse weather conditions like fog, rain, or snowfall, which attenuate sensor signals, leading to data loss and the introduction of outliers. The Keep It Small and Simple ICP (KISS-ICP) (Vizzo et al., 2023) developed an adaptive threshold for point correspondences and employed a robust kernel on a point-to-point cost function to mitigate outliers. However, it uses the first point from each voxel for voxelization, which can result in outlier points being selected for registration and map update. We believe that while employing robust kernels on cost functions helps mitigate outliers, additional scan preprocessing or cleaning remains necessary.

2.2 Map denoising

Classical point cloud cleaning methods typically involve filtering low-density points, like Statistical Outlier Removal (SOR) and Radius Outlier Removal (ROR), which use a local point density threshold. However, these methods struggle with varying density, often discarding points from distant areas. To solve this, Dynamic Radius Outlier Removal (DROR) (Prio et al., 2022), adapts the threshold based on point range. Low-Intensity Outlier Removal (LIOR) (Park et al., 2020) considers intensity, preserving points with low intensity and high density. A fusion of DROR and LIOR was proposed in Dynamic Distance-Intensity Outlier Removal (DDIOR) (Xu et al., 2020), to achieve better performance. A common drawback of these methods is their slowdown with large point clouds due to reliance on k-d tree structures for nearest-neighbour (NN) searches.

In contrast, learning-based approaches perform point cloud segmentation to filter weather-induced outliers. Recent advancements like WeatherNet (Heinzler et al., 2020) and 4DenoiseNet (Seppänen et al., 2022) accurately segment fog and rain clutter. (Seppänen et al., 2023) implements a multi-echo approach to denoise point clouds in challenging weather conditions. Rather than selecting points from the strongest echo, their model utilizes points from alternative echoes, which are not present in conventional strongest echo point clouds. While providing promising results, these networks require extensive labelled data, memory, and computational power due to numerous trainable parameters. Additionally, they often rely on graphics processing units (GPUs) for online computation and require re-training.

On the contrary, our method operates efficiently without requiring a GPU or powerful hardware, it uses cylindrical range images for NN searches, eliminating the need for k-d trees. Also, rather than filtering data based on a density threshold, we pick the most suitable points from each voxel according to their rank. This approach prioritizes the selection of points with the

¹ <https://github.com/eugeniu1994/FPR>

highest rank, allowing even those attenuated by adverse weather to be considered for registration when no alternatives are available. Subsequently, we update the map with the highest rank points. Thus, our method serves a dual purpose: selecting optimal points for registration and cleaning the map, as outliers typically exhibit smaller rankings.

To conclude, the FPR method does not filter out points during the LO phase. Instead, it focuses on selecting the most suitable points for registration. Points are filtered only in the map denoising phase, as this filtering depends on a threshold that requires tuning. By performing this filtering after registration, we avoid the risk of inlier points being filtered out due to inaccurate or unknown threshold parameters. Therefore, in our experiments, we compare our method with the KISS-ICP algorithm and do not apply any density-based scan cleaning methods discussed in the map denoising section of the literature review.

3. Method

This work introduces a rank-based scan voxelization method to mitigate the impact of adverse weather conditions (fog, rain, and snowfall) on lidar odometry and mapping. The core idea is to assign a reliability rank to each point based on the range deviation of its local neighbors in the cylindrical range image. Points with higher ranks are considered more reliable for odometry estimation.

Within a local neighborhood, the point with the highest rank is more likely to have valid correspondences in the map due to its association with higher local point density. In contrast, noise measurements, often caused by adverse weather, tend to exhibit low point density and yield lower ranks. This behavior is shown in Fig.1.

The process begins by projecting the lidar point cloud into cylindrical range image space, enabling efficient approximation of nearest neighbors. Then, each point is assigned a rank, which is used during voxelization to select the most reliable representative point from each voxel. The following sections provide a detailed explanation of each step.

3.1 Cylindrical range image projection

The range image projection parameters are determined using the sensor's specifications (angular resolution $\Delta\phi$ and the number of rings/channels). The pixel coordinates of a projected point $\mathbf{p}$ are computed as follows

$$\begin{cases} u = \text{ring}, \\ v = \lfloor \frac{\phi}{\Delta\phi} \rfloor \end{cases} \quad (1)$$

where u and v are the row and column values, ϕ is the point's azimuth angle, and $\lfloor \cdot \rfloor$ rounds towards the nearest integer. Since the cylindrical projection is used only for rank computation, no point cloud data used in the later processing is lost nor rounded. Unlike traditional or spherical range image projections, the cylindrical range image is less sparse, as each point's row index is directly taken from the lidar point ring rather than inferred from its elevation angle. See Fig. 1b for a visual representation of a cylindrical range image.

3.2 Fast Point Ranking method

For each 3D point $\mathbf{p}$ mapped to (u, v) range image coordinates, we consider the surrounding neighbourhood within a 5×5 kernel centered at (u, v) with

$$N_{5 \times 5}(u, v) = \left\{ (u', v') \mid \begin{matrix} u-2 \leq u' \leq u+2 \\ v-2 \leq v' \leq v+2 \end{matrix} \right\}. \quad (2)$$

A 5×5 kernel provides enough neighborhood context for rank computation without losing spatial relationships. Larger kernels capture more context but increase the number of neighbors and computational cost, making 5×5 a balanced choice. Next, inspired by the Gaussian function, we compute the rank of the j th point as

$$R_j = \left(1 + \frac{1}{N'} \sum_{i=1}^{N'} e^{-\frac{(r_j - r_i)^2}{2\sigma^2}} \right) \left(1 + \frac{r_j}{r_{\max}} \right), \quad (3)$$

where r_j is the range of the j th point and N' is the number of non-empty neighbours (pixels with at least 1 projected point) with $N' \leq N$, where $N = 25$ is the total number of neighbouring pixels. For weighting, we assume a zero-mean Gaussian distribution with a standard deviation, $\sigma = 1$ m. A higher ranking indicates more non-empty neighbours (N') and smaller differences in their range values. Due to the sparsity of distant points, we weight them using a normalized range factor ($r_j/r_{\max}$) with $r_{\max} = 100$ m. Without this weighting, closer points, having higher density, would naturally receive higher ranks and be prioritized during voxelization. Since both terms are below 1, we have to add 1 to both sides to preserve the intended ranking order and avoid the geometric decay.

In the scan voxelization phase, we utilize the rank values associated with all points. As outlined in Algorithm 1, after computing the FPR for the point cloud, we identify the voxel corresponding to each point and retain only the highest-ranked point per voxel. Unlike selecting the first, middle, or mean point, this approach prioritizes points less affected by adverse conditions. In cases where multiple points share the highest rank, the first occurrence is selected.

Algorithm 1 Fast Point Ranking Voxelization

Input: 3D point cloud P , voxel leaf size s
Output: Voxelized point cloud
 Compute the range image for P
 Compute FPR for each point using (3)
for point $\mathbf{p} \in P$ **do**
 Find the corresponding voxel V
 if point in V has a smaller rank **then**
 Set voxel centroid $V \leftarrow \mathbf{p}$
 end if
end for

4. Experiments

4.1 Experimental Setup

We utilized the KITTI-360 dataset (Liao et al., 2022) as a good weather reference, leveraging its extensive driving distance of over 70 km (see Fig. 2), which represents a significant advancement over its predecessor, KITTI (Geiger et al., 2012). The dataset includes 10 Hz point cloud scans with (x, y, z) coordinates and intensity acquired with the Velodyne HDL-64 sensor, which has 64 channels and an angular resolution,




Figure 2. Estimated trajectories on the Adverse-Weather-KITTI-360 dataset, using good-weather sequences as a reference. This visualization highlights the accumulated drift over long-term driving. Note: KITTI-360 does not provide sequences 01 and 08.

$\Delta\phi = 0.2^\circ$. The KITTI-360 dataset includes high-accuracy reference poses, which are measured using position sensors combined with visual features (Liao et al., 2022). The Velodyne data comes in an ordered format, allowing us to infer the ring/channel field (see the shared example code in (Vezeteu, 2024) for details).

The KITTI-360 dataset provides good weather measurements, while some of the existing datasets like the Seeing Through Fog (STF) (Bijelic et al., 2020) include fog, snowfall, and rain but lack ground truth positioning for evaluating LO. The well-known Canadian Adverse Driving Conditions (CADC) (Pitropov et al., 2021) and Winter Adverse Driving dataset (WADS) (Kurup and Bos, 2021) datasets primarily concentrate on snowfall for 3D object detection and contain sequences that are too short, making them unsuitable for LO evaluation. The novel SnowyKITTI dataset (Seppänen et al., 2022) augments the KITTI scans (Geiger et al., 2012) with simulated snowfall. However, it lacks the point ring information, which cannot be inferred after the augmentation process. To evaluate our method, we applied point-wise augmentation techniques to simulate fog (Hahner et al., 2021), snowfall (Hahner et al., 2022), and rain (Kilic et al., 2021) on the KITTI-360 dataset. We generated long-term LiDAR driving sequences under adverse weather conditions, incorporating (x, y, z, intensity, ring/channel) data, following the structure of the Adverse-Weather-KITTI-360 dataset introduced in (Vezeteu, 2024).

Fog (Hahner et al., 2021) models lidar pulse transmission by modifying the impulse response of the optical channel, which converts clear-weather point clouds to foggy counterparts. Rain (Kilic et al., 2021) has a physics-based approach that uses a hybrid Monte Carlo method to simulate rain effects on lidar point clouds by placing large particles randomly and calculating scattering efficiencies from Mie scattering theory (Wriedt, 2012, Kilic et al., 2021) given the rain rate and lidar parameters. The snowfall (Hahner et al., 2022) affects lidar data by sampling snow particles and adding the ground wetness effect. Fig. 3 provides illustrative examples of each weather scenario. The augmented adverse weather on point clouds appears realistic, even though it's simulated. Since we have identical scans under different weather conditions, the results are directly comparable.

We observe that adverse weather conditions cause signal occlusion and introduce false measurements, especially near the

sensor and in the direction of movement, leading to the loss of long-distance points. While most near-sensor noise can be filtered using a minimum range threshold (e.g., in rain), determining an optimal threshold is challenging. Additionally, as shown in Fig. 3d, snowfall attenuates most measurements, making it difficult to distinguish valid points from those that should be removed, limiting the effectiveness of traditional cleaning methods.

Most traditional lidar odometry and mapping techniques are model-based algorithms that require parameter adjustment for




Figure 3. A demonstration of the KITTI-360 scan, showcasing various weather conditions: a) Original scan plotted by ring/channel field, b) Foggy weather, c) Rainy conditions, and d) Snowfall. The scans were not subjected to min-max range filtering. The red points indicate data affected by adverse weather conditions or attenuated scan points. The grey points are original measurements that were not affected by adverse weather. The coloring of gray/red is based on metadata produced by the noise simulation methods and included in the dataset as labels. This visualization emphasizes how points closer to the sensor are more attenuated by adverse weather, while points at longer distances are often lost.

	Absolute error ($t_{abs}[m]$ / $r_{abs}[deg]$)				Relative error ($t_{rel}[\%]$ / $r_{rel}[\%]$)			
	Mean	Median	Std	RMSE	Mean	Median	Std	RMSE
KITTI-360								
KISS-ICP	13.63 / 4.11	10.43 / 3.81	10.27 / 2.10	17.06 / 4.62	1.46 / .70	1.07 / .54	1.64 / .72	2.20 / 1.00
FPR	11.26 / 3.26	8.90 / 2.93	8.67 / 1.78	14.21 / 3.71	1.39 / .69	1.01 / .51	1.62 / .73	2.14 / 1.00
KITTI-360 with simulated fog								
KISS-ICP	13.13 / 3.90	9.90 / 3.45	10.66 / 2.13	16.91 / 4.45	1.56 / .71	1.08 / .55	2.08 / .72	2.60 / 1.02
FPR	10.71 / 3.32	7.92 / 2.96	9.32 / 1.88	14.20 / 3.82	1.44 / .69	1.02 / .51	1.73 / .73	2.25 / 1.00
KITTI-360 with simulated rain								
KISS-ICP	24.45 / 6.26	17.49 / 5.39	20.98 / 3.58	32.22 / 7.21	1.75 / .83	1.32 / .71	1.77 / .71	2.49 / 1.09
FPR	11.11 / 3.61	8.93 / 3.24	8.24 / 1.90	13.83 / 4.08	1.41 / .71	1.05 / .54	1.61 / .72	2.14 / 1.01
KITTI-360 with simulated snowfall								
KISS-ICP	15.07 / 3.93	13.01 / 3.49	10.78 / 2.24	18.54 / 4.58	1.59 / .66	1.09 / .49	2.22 / .71	2.73 / .98
FPR	11.48 / 3.66	8.29 / 3.27	9.31 / 1.96	14.78 / 4.15	1.42 / .67	1.04 / .49	1.63 / .72	2.17 / .98

Table 1. The absolute and relative trajectory error statistics w.r.t. translation [m] and rotation [deg] of the KITTI-360 dataset for each weather condition shared in Adverse-Weather-KITTI-360. **Bold** indicates the best result (smallest error). The table compares the output of KISS-ICP with the proposed Fast Point Ranking (FPR) modification.

various data types and scenarios. In contrast, the KISS-ICP² (Vizzo et al., 2023) method stands out for its ability to perform effectively across different sensor data and scenarios with the same set of parameters. Consequently, in this study, we compared our method against the state-of-the-art KISS-ICP (Vizzo et al., 2023), modifying only the point sampling step while keeping the rest of the pipeline unchanged. Specifically, it applies a double voxelization strategy: a voxel leaf size of 1.5 m is used for selecting points for registration, while a finer leaf size of 0.5 m is used for updating the map. During the voxelization phase, KISS-ICP selects the first point from each voxel and discards the rest. In contrast, we select the highest-ranked point from each voxel, ensuring more robust data association and improved LO performance. We chose KISS-ICP (Vizzo et al., 2023) as our baseline for several reasons. First, it is the current state-of-the-art, demonstrating superior performance compared to other LO methods. Second, its modular and open-source design makes it straightforward to integrate our ranking-

² <https://github.com/PRBonn/kiss-icp>




Figure 4. Two consecutive point clouds from sequence 10 of the Adverse-Weather-KITTI-360 dataset coloured by FPR values. (a) A snow-affected point cloud colored by FPR values. (b) The same scene after 0.1 seconds of motion, from a slightly different viewpoint, showing the stability of FPR values. The enlarged circle in b) highlights FPR values computed for traffic poles and tree trunks.

based voxelization without altering the core algorithm. In the upcoming Results section, "KISS-ICP" denotes the original algorithm, while "FPR" represents our modified version of KISS-ICP. Therefore, we investigate the impact of rank-based voxelization on KISS-ICP.

The scan pre-cleaning methods discussed in the literature review, such as SOR, ROR, LIOR (Park et al., 2020), DROR (Prio et al., 2022), DDIOR (Xu et al., 2020), WeatherNet (Heinzler et al., 2020), and 4DenoiseNet (Seppänen et al., 2022), aim to remove adverse-weather-induced noise but require parameter tuning, which may over-filter valid points or retain outliers depending on weather conditions. In contrast, our approach is not a pre-cleaning method but a ranking-based subsampling strategy designed for LO. While it can be extended for map cleaning, its primary focus is on selecting the most reliable points for LO voxelization without directly filtering the scan. Since our method fundamentally differs from pre-cleaning approaches and does not require removing points before LO, we do not evaluate these methods in our experiments.

The estimated localization was evaluated based on absolute (ATE) and relative trajectory errors (RTE) w.r.t. rotation and translation (Grupp, 2017). For ATE, the trajectories were first aligned. As presented in (Vizzo et al., 2023), relative errors were calculated across all trajectory segments ranging from 100 m to 800 m in size.

4.2 Results

Table 1 presents averaged statistics (mean, median, standard deviation, and root mean square error (RMSE)) on absolute and relative trajectory errors across all scenarios of the KITTI-360 dataset under different weather conditions. Comparisons are made between trajectories generated by KISS-ICP and FPR methods. Overall, FPR demonstrates superior trajectory accuracy. The improvement in translation by using FPR for each dataset scenario is illustrated in Fig. 5, with subplots depicting (a) absolute and (b) relative RMSE differences between KISS-ICP and FPR methods.

Notably, FPR significantly enhances on the absolute scale (Fig. 5a), particularly in rainy conditions, showing an approximate average translation improvement of 10 m. In contrast,

the improvement is around 3 m in other weather conditions. The most significant ATE improvement of approximately 50 m is observed in raining scenario 07. Also, our method exhibits inferior performance in snowfall scenario 07, yet it surpasses the original method in all other cases. On the relative scale (Fig. 5b), the maximum improvement of over 1.5 % is observed in snowfall scenario 10.

The relative errors (Table 1) show that KISS-ICP provides an average 1.46 % translation error, while our proposal yields 1.39 %, almost an insignificant 0.07 % improvement. However, there's a noticeable average relative translation improvement for fog (0.12 %), rain (0.34 %), and snowfall (0.17 %). Additionally, our method consistently yielded similar relative errors of around 1.4 % regardless of the weather conditions. This highlights its improved invariance to adverse weather conditions.

Our method outperformed the original in most error metrics and across all weather conditions, with one exception. Specifically, we observed slightly inferior results on relative rotation in snowy weather, where the original method exhibited a slightly lower mean error by 0.01 degrees (Table 1 mean relative rotation). A possible reason could be that point cloud sparsity induces a lower rank for distant points. Consequently, the algorithm may pick closer points for registration, discarding dis-

tant points that might be better for estimating rotation parameters. Nonetheless, our approach excelled in all other scenarios, highlighting its robustness in challenging weather conditions.

In addition to improving the positioning accuracy by selecting the most suitable point from a voxel, the FPR can be used to clean the collected data after registration. An example of the accumulated point cloud is depicted in Fig. 7. (a) shows all accumulated data colored with the rank R in (3), and (b) shows the result when 10% of the lowest ranked points are removed. We can see that weather-induced measurements create a cloud of outliers around the sensor trajectory, particularly in open areas. Our method effectively gives these points a low rank, and they can be filtered out.

The extra time required to use FPR is minimal. It scans the points once for cylindrical projection $O(n)$ and iterates again to compute ranks $O(n)$. Thus, the overall time complexity is $O(2n)$, indicating efficient, linear scaling with the number of points.

As shown in Fig. 6, the FPR metric effectively mirrors the actual adverse weather-induced noise in the point cloud. The points affected by adverse weather have lower rank values, which have a smaller likelihood of being selected for registration and map updates. This demonstrates the FPR's ability to identify and appropriately down-weight the impact of adverse weather induced points.

5. Discussion

Integrating FPR during scan voxelization improves LO performance, especially under adverse weather conditions (fog, rain, snow), by enhancing both rotational and translational accuracy. While performance gains are evident in good weather, the improvements are more significant in noisy, weather-degraded scans, showcasing the method's robustness. Our analysis shows that weather-affected points tend to receive lower FPR scores, making them less likely to influence scan registration or map updates. Unlike prior work focusing on pre-cleaning, our results highlight the benefit of precise scan voxelization for robust registration.

Given Eq.(3) and Fig.1, vertical structures (e.g., walls) receive higher ranks than the ground due to lidar sampling density. Although small features like poles may have lower ranks, this poses no issue since ranking is done locally within each voxel. Fig. 4 further shows FPR stability across consecutive scans and highlights values for poles and tree trunks.

We validated our method using a newly released Adverse-Weather-KITTI-360 dataset (Vezeteu, 2024), created by simulating fog, rain, and snow on KITTI-360 (Liao et al., 2022). This allows controlled, repeatable comparisons of localization accuracy across weather types using high-quality reference trajectories (see Table 1, Fig. 5). Future work should extend this evaluation to real adverse conditions.

One limitation is that the simulations assume motion-compensated scans, while noise is added to raw data, compensation is applied later using a constant-velocity model. We believe this absorbs most motion distortion, but the exact impact was not analyzed.

Lastly, while our approach depends on ring/channel info for cylindrical projection, extending it to non-rotating LiDARs would




Figure 5. The absolute (a) and relative (b) RMSE differences calculated for each weather condition on each Adverse-Weather-KITTI-360 (Vezeteu, 2024) sequence. The values represent the RMSE reduction when comparing the proposed FPR method to KISS-ICP.




Figure 6. The illustration of the FPR metric versus the actual adverse weather labels. a) An example of a point cloud affected by fog, sourced from the Adverse-Weather-KITTI-360 (Vezeteu, 2024) dataset trajectory 03. The points are coloured by labels: red represents fog-induced noise, and grey denotes actual measurements. b) The same point cloud coloured by FPR values.

require adaptations. For rotating LiDARs, users may tune the FPR neighborhood size based on sensor specs. Future work could also explore using FPR for weighted voxel means rather than selecting a single point per voxel.

6. Conclusion

In this paper, we proposed a novel approach, **Fast Point Ranking** (FPR), that leverages the cylindrical range image view of a scan and computes a rank for each point based on the number of neighbours and their range discrepancy. Moreover, we introduced a novel rank-based point cloud voxelization method for robust lidar odometry (LO) and map denoising in adverse weather conditions. The method employs a simple yet effective strategy, assigning a rank to each point and selecting the optimal point from each voxel in the scan voxelization phase. Additionally, the ranks can be utilized to remove noisy points




Figure 7. A demonstration of scan accumulation showcasing the initial 100 scans from sequence 02 of the foggy (see Fig. 3b) Adverse-Weather-KITTI-360 (Vezeteu, 2024) dataset. a) Depiction of KISS-ICP output plotted by point rank value, b) The same point cloud with 10% lowest ranked points dropped out. The red line denotes the sensor trajectory. The red and orange points around the sensor trajectory in a) are the fog-induced range measurements.

and to enhance the map's quality.

The experimental results with KITTI-360 (Liao et al., 2022) and Adverse-Weather-KITTI-360 (Vezeteu, 2024) datasets demonstrated significant improvement in LO performance across all tested weather conditions, with notable improvements observed in rainy, foggy, and snowfall scenarios. Furthermore, the approach demonstrates robustness for adverse weather, with a consistent relative translation error of 1.4 % observed regardless of the weather condition. This demonstrates the robustness and reliability of the proposed method under different weather conditions.

Our study is limited to demonstrate the potential of the proposed FPR in positioning accuracy for only KISS-ICP (Vizzo et al., 2023) method. However, the results encourage the future integration of FPR also with other LO and mapping methods, especially, since FPR is simple to implement and computationally efficient. It is important to note that the approach relies on point ring/channel information, which may present limitations with existing data sets. To overcome this, we shared the Adverse-Weather-KITTI-360 dataset (Vezeteu, 2024), which also includes the ring/channel information for each point.

Overall, the proposed method presents a promising solution for addressing the challenges of LO in adverse weather condi-

tions, offering potential benefits for navigation and perception systems across various real-world scenarios. Future research should test the method also with other LO methods and in real challenging adverse weather conditions to observe the limits of the proposed method.

7. Acknowledgements

Co-funded by the European Union. Views and opinions expressed are however, those of the authors only and do not necessarily reflect those of the European Union or European Climate, Infrastructure and Environment Executive Agency (CINEA). Neither the European Union nor the granting authority can be held responsible for them. Project grant no. 101069576.

References

- Bijelic, M., Gruber, T., Mannan, F., Kraus, F., Ritter, W., Dietmayer, K., Heide, F., 2020. Seeing through fog without seeing fog: Deep multimodal sensor fusion in unseen adverse weather. *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Censi, A., 2008. An ICP variant using a point-to-line metric. *2008 IEEE international conference on robotics and automation*, IEEE, 19–25.
- Charron, N., Phillips, S., Waslander, S. L., 2018. De-noising of lidar point clouds corrupted by snowfall. *2018 15th Conference on Computer and Robot Vision (CRV)*, IEEE, 254–261.
- Geiger, A., Lenz, P., Urtasun, R., 2012. Are we ready for autonomous driving? the KITTI vision benchmark suite. *Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Grupp, M., 2017. evo: Python package for the evaluation of odometry and slam. <https://github.com/MichaelGrupp/evo>.
- Hahner, M., Sakaridis, C., Bijelic, M., Heide, F., Yu, F., Dai, D., Van Gool, L., 2022. LiDAR snowfall simulation for robust 3D object detection. *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Hahner, M., Sakaridis, C., Dai, D., Van Gool, L., 2021. Fog Simulation on Real LiDAR Point Clouds for 3D Object Detection in Adverse Weather. *IEEE International Conference on Computer Vision (ICCV)*.
- Heinzler, R., Piewak, F., Schindler, P., Stork, W., 2020. CNN-Based Lidar Point Cloud De-Noising in Adverse Weather. *IEEE Robotics and Automation Letters*, 5(2), 2514–2521.
- Kilic, V., Hegde, D., Sindagi, V., Cooper, A. B., Foster, M. A., Patel, V. M., 2021. Lidar light scattering augmentation (LISA): Physics-based simulation of adverse weather conditions for 3D object detection. *arXiv preprint arXiv:2107.07004*.
- Kim, J., Scott, C. D., 2012. Robust kernel density estimation. *The Journal of Machine Learning Research*, 13(1), 2529–2565.
- Kuramachi, R., Ohsato, A., Sasaki, Y., Mizoguchi, H., 2015. G-ICP SLAM: An odometry-free 3D mapping system with robust 6DoF pose estimation. *2015 IEEE International Conference on Robotics and Biomimetics (ROBIO)*, IEEE, 176–181.
- Kurup, A., Bos, J., 2021. Dsor: A scalable statistical filter for removing falling snow from lidar point clouds in severe winter weather. *arXiv preprint arXiv:2109.07078*.
- Liao, Y., Xie, J., Geiger, A., 2022. KITTI-360: A Novel Dataset and Benchmarks for Urban Scene Understanding in 2D and 3D. *Pattern Analysis and Machine Intelligence (PAMI)*.
- Magnusson, M., 2009. The three-dimensional normal-distributions transform: an efficient representation for registration, surface analysis, and loop detection. PhD thesis, Örebro universitet.
- Park, J.-I., Park, J., Kim, K.-S., 2020. Fast and accurate denoising algorithm for LiDAR point clouds. *IEEE Access*, 8, 160202–160212.
- Pitropov, M., Garcia, D. E., Rebello, J., Smart, M., Wang, C., Czarnecki, K., Waslander, S., 2021. Canadian adverse driving conditions dataset. *The International Journal of Robotics Research*, 40(4–5), 681–690.
- Prio, M. H., Patel, S., Koley, G., 2022. Implementation of dynamic radius outlier removal (DROR) algorithm on LiDAR point cloud data with arbitrary white noise addition. *2022 IEEE 95th Vehicular Technology Conference (VTC2022)*, IEEE, 1–7.
- Seppänen, A., Ojala, R., Tammi, K., 2022. 4DenoiseNet: Adverse Weather Denoising From Adjacent Point Clouds. *IEEE Robotics and Automation Letters*, 8(1), 456–463.
- Seppänen, A., Ojala, R., Tammi, K., 2023. Multi-Echo Denoising in Adverse Weather. *arXiv preprint arXiv:2305.14008*.
- Shan, T., Englot, B., 2018. LeGO-LOAM: Lightweight and ground-optimized lidar odometry and mapping on variable terrain. *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, IEEE, 4758–4765.
- Vezeteu, E., 2024. Adverse-Weather-KITTI-360 dataset. <https://doi.org/10.23729/b73a84dc-89cf-4750-a93d-343046dd2ca6>. National Land Survey of Finland, FGI Dept. of Remote sensing and photogrammetry.
- Vizzo, I., Guadagnino, T., Mersch, B., Wiesmann, L., Behley, J., Stachniss, C., 2023. KISS-ICP: In Defense of Point-to-Point ICP—Simple, Accurate, and Robust Registration If Done the Right Way. *IEEE Robotics and Automation Letters*, 8(2), 1029–1036.
- Wang, H., Wang, C., Chen, C.-L., Xie, L., 2021. F-loam: Fast lidar odometry and mapping. *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, IEEE, 4390–4396.
- Wriedt, T., 2012. Mie theory: a review. *The Mie theory: Basics and applications*, 53–71.
- Xu, C., Wu, B., Wang, Z., Zhan, W., Vajda, P., Keutzer, K., Tomizuka, M., 2020. Squeezesegv3: Spatially-adaptive convolution for efficient point-cloud segmentation. *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXVIII*, Springer, 1–19.
- Zhang, J., Singh, S., 2014. LOAM: Lidar odometry and mapping in real-time. *Proceedings of Robotics: Science and Systems*, Berkeley, USA.