Segmentation-Aware Weighted NDT for Real-Time LiDAR Localization on CPU-Only Platforms

Seokhwan Jeong, Junhyeok Lee, Moonbeom Kim, and Jeongyeup Paek

Department of Computer Science and Engineering, Chung-Ang University, Seoul, Republic of Korea

tjrghks526@cau.ac.kr, ljh102715@cau.ac.kr, mbkim@cau.ac.kr, jpaek@cau.ac.kr

***Abstract*—In mobile robots and autonomous vehicles, Light Detection and Ranging (LiDAR) sensors are widely used because they provide detailed 3D spatial information. The Normal Distributions Transform (NDT) models point distributions within voxel cells as Gaussian distributions for efficient scan matching. However, conventional NDT does not distinguish between ground and non-ground regions, although these regions provide different levels of geometric information for pose estimation. This paper presents a segmentation-aware weighted NDT method that projects LiDAR scans onto 2D range images to classify points as ground or non-ground and reject noise points without relying on learning-based segmentation. The weighting strategy adjusts the contribution of ground and non-ground measurements, while deterministic source ground subsampling reduces the number of processed points. Evaluation on real-world driving sequences from the KITTI odometry dataset shows that Weighted-S50 reduces macro-averaged ATE by 7.6%, yaw MAE by 14.3%, and processing time by 20.1% compared with Classical NDT on a CPU-only platform.**

***Index Terms*—LiDAR localization, Normal Distributions Transform (NDT), weighted NDT, scan matching**

I. Introduction

Light Detection and Ranging (LiDAR) sensors provide accurate 3D information about the surrounding environment and operate reliably under varying illumination conditions, including nighttime [1]. Therefore, LiDAR has been widely used in autonomous driving and robotics [2], [3]. In this work, we focus on map-based LiDAR localization, which estimates the pose of a robot by aligning the current scan data with a pre-built map. LiDAR point clouds typically contain a large number of points. Processing these high-density point clouds in real time imposes a high computational burden. Therefore, LiDAR localization is challenging in resource-constrained hardware environments [4].

Two representative scan matching approaches for map-based LiDAR localization are Iterative Closest Point (ICP) methods and Normal Distributions Transform (NDT) methods. Whereas ICP relies on nearest-neighbor search to find correspondences between two point clouds [5], the NDT represents the point distribution within each voxel cell as a Gaussian distribution and performs scan matching based on these distributions [6]. However, each method has its own limitations. Because ICP relies on nearest-neighbor searches, it is difficult to achieve real-time performance on hardware with limited computational resources when the number of points is large. Although NDT reduces correspondence-search costs by using voxel-based Gaussian representations, real-time NDT-based localization can still be computationally demanding on CPU-only platforms. Moreover, conventional NDT does not explicitly account for differences in the geometric informativeness of individual voxel cells. A prior weighted NDT method accounts for range and geometric characteristics, but does not explicitly distinguish the contributions of ground and non-ground measurements [7].

This study presents and evaluates a segmentation-aware weighted NDT localization method for CPU-only environments. Its design separates weighting from point reduction: segmentation-aware weights adjust the contributions of structurally different measurements, while deterministic ground subsampling reduces the number of source points processed during scan matching. Experiments on real-world driving sequences examine localization accuracy and computational efficiency. Ultimately, this work aims to enable real-time LiDAR localization on a CPU-only platform.

II. Related Work

LiDAR scan matching requires a trade-off between localization accuracy and computational cost, since using all measured points can improve geometric constraints but increases the registration burden. ICP and GICP estimate the relative pose from point correspondences, whereas NDT improves efficiency by representing local point distributions as Gaussian distributions in voxel cells [4]–[6]. NDT-LOAM, an NDT-based odometry and SLAM method, extends the standard NDT formulation by assigning different weights to NDT cells based on range and geometric characteristics for LiDAR motion estimation [7].

Several methods reduce scan matching cost by selecting informative points or regions. Object-based ICP and SP-ICP use object-level or representative points to reduce unnecessary correspondences [8], [9]. LeGO-LOAM, a range-image-based LiDAR odometry and SLAM method, projects 3D scans onto

This research was supported by the MSIT(Ministry of Science, ICT), Korea, under the National Program for Excellence in SW), supervised by the IITP(Institute of Information & Communications Technology Planning & Evaluation) in 2026 (2025-0-00032), by IITP-ITRC (Information Technology Research Center) grant funded by the Korea government (MSIT) (IITP-2026-RS-2022-00156353, 33%), and also by the National Research Foundation of Korea (NRF) grant funded by the Korea government (MSIT) (No. RS-2024-00359450). J. Paek is the corresponding author.

2D range images to separate ground and non-ground measurements and extract robust features [10]. Unlike NDT-LOAM, our method incorporates ground/non-ground label consistency between source points and target NDT voxels into scan-to-map localization and combines this label-aware weighting with deterministic ground subsampling of the source scan.

III. Method

We propose a segmentation-aware weighted NDT method for CPU-only LiDAR localization. In this work, structural labels denote coarse classes obtained from range-image segmentation, namely ground and non-ground, rather than learning-based object semantics. The method consists of offline weighted NDT map construction and online scan-to-map localization. The proposed method uses ground/non-ground labels for both the target map and the source scan and incorporates structural compatibility and principal component analysis (PCA)-based geometric weights into the NDT optimization.

A. Shared Range-Image Segmentation

Let $\mathcal{P}_t = \{p_i\}_{i=1}^{N_t}$ be the LiDAR scan acquired at time t. Each valid point p_i is projected onto a two-dimensional range image using its vertical scan index and horizontal azimuth angle. The pixel value stores the Euclidean range from the sensor to the corresponding 3D point. Ground points are first identified from the inclination between vertically adjacent range-image measurements. The remaining non-ground pixels are grouped by image-based connected-component segmentation, following the range-image segmentation strategy used in LeGO-LOAM [10]. Components are retained only if they contain a sufficient number of points and span multiple LiDAR rings; otherwise, they are treated as noise. Each retained point is assigned a coarse structural confidence,

$$s_i = \begin{cases} s_g = 0.5, & \text{ground}, \\ s_o = 1.0, & \text{non-ground}. \end{cases}$$

Rejected points are treated as noise and do not enter either the weighted map or the online weighted NDT optimization. The same range-image segmentation and confidence encoding are used during offline map construction and online localization.

B. Offline Weighted NDT Map Construction

Using the pose corresponding to each scan, the labeled ground and non-ground points are transformed into the common map frame and accumulated. The accumulated map is divided into NDT voxels of resolution r as in the standard NDT formulation. For voxel c_j with point set $\mathcal{V}_j$ and cardinality N_j, we compute its mean $\boldsymbol{\mu}_j$, covariance $\boldsymbol{\Sigma}_j$, and average structural confidence

$$\bar{s}_j = \frac{1}{N_j} \sum_{i \in \mathcal{V}_j} s_i.$$

Following the PCA-based geometric characterization used in NDT-LOAM [7], PCA is performed on the sample covariance before the minimum-eigenvalue regularization used by NDT. Let $\sigma_1 \leq \sigma_2 \leq \sigma_3$ denote the square roots of the ordered

Algorithm 1 Online Segmentation-Aware Weighted NDT Localization

Require: Scan S_t, map $\mathcal{M}_w$, previous accepted poses
Ensure: Pose estimate $\hat{T}_t$

Segment S_t and remove rejected components
Deterministically retain 50% of ground points
Keep all non-ground points before applying source voxel filtering
Predict the initial pose from the last two accepted poses
while not converged and below the iteration limit **do**
  Transform source points using the current pose
  Find the target NDT voxels using DIRECT7 neighbor search
  Compute W_{ij} using Eq. (1)
  Accumulate the score, gradient, and Hessian of Eq. (2)
  Update the six-degree-of-freedom pose
end while
if converged, the probability passes the threshold, and the correction is bounded **then**
  Publish and store $\hat{T}_t$
end if
return $\hat{T}_t$

eigenvalues of the sample covariance matrix. The geometric features are

$$f_{\mathrm{L}} = \frac{\sigma_3 - \sigma_2}{\sigma_3}, \quad f_{\mathrm{P}} = \frac{\sigma_2 - \sigma_1}{\sigma_3}, \quad f_{\mathrm{V}} = \frac{\sigma_1}{\sigma_3},$$
$$l_j = \operatorname*{arg\,max}_{k \in \{1,2,3\}} \big(f_{\mathrm{L}}, f_{\mathrm{P}}, f_{\mathrm{V}}\big)_k,$$

where the values 1, 2, and 3 of l_j represent linear, planar, and volumetric structures, respectively. The dimension-dependent factor is

$$\alpha(l_j) = \begin{cases} 0.75, & l_j = 1, \\ 1.25, & l_j = 2, \\ 1.00, & l_j = 3. \end{cases}$$

The target contribution of voxel c_j is

$$w_j^{\text{target}} = \bar{s}_j \alpha(l_j).$$

The voxel index and statistics $\boldsymbol{\mu}_j$, $\boldsymbol{\Sigma}_j$, $\bar{s}_j$, l_j, and N_j are serialized into a precomputed weighted NDT map. The classical baseline instead uses raw map and source points without structural, PCA-based geometric, or compatibility weights.

C. Online Weighted NDT Localization

The current scan is segmented in the LiDAR frame using the shared range-image procedure. To reduce redundant ground measurements, the source-subsampling configuration deterministically retains every other ground point, corresponding to a ground retention ratio of $\rho = 0.5$. All non-ground points are retained before source voxel filtering. After segmentation and ground subsampling, points are transformed from the LiDAR frame into the vehicle base frame and voxel-downsampled.

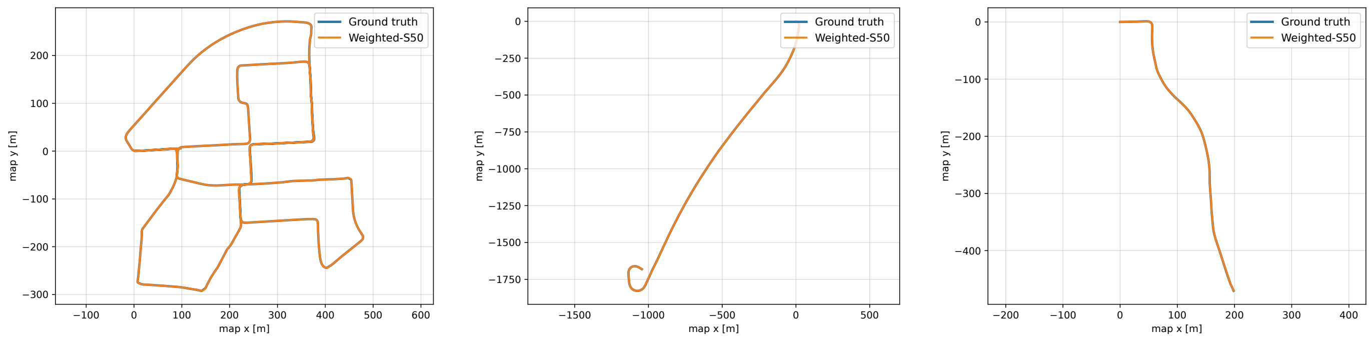

(a) Seq. 00: residential **(b)** Seq. 01: highway **(c)** Seq. 03: rural area

Fig. 1: Representative KITTI odometry trajectories used for evaluation.

The last two accepted poses provide a bounded constant-velocity initial estimate. The prediction initializes NDT but does not replace scan matching. For source point p_i and target voxel c_j, structural compatibility is

$$\eta(s_i, \bar{s}_j) = 1 - (1 - \lambda_m) \min\left(1, \frac{|s_i - \bar{s}_j|}{s_o - s_g}\right),$$

where $\lambda_m = 0.1$ is the mismatch penalty. Matching pure classes obtain $\eta = 1$, opposite pure classes obtain $\eta = \lambda_m$, and mixed target voxels are interpolated continuously. The complete source–target weight is

$$W_{ij} = s_i\, w_j^{\text{target}}\, \eta(s_i, \bar{s}_j). \tag{1}$$

Thus, target structural confidence is included exactly once. Let $\mathbf{q}_{ij} = Tp_i - \boldsymbol{\mu}_j$. The weighted NDT matching score is

$$\mathcal{S}(T) = \sum_i \sum_{j \in \mathcal{N}(Tp_i)} W_{ij} \exp\left(-\frac{1}{2}\mathbf{q}_{ij}^{\mathsf{T}} \boldsymbol{\Sigma}_j^{-1} \mathbf{q}_{ij}\right). \tag{2}$$

The weight W_{ij} is applied during NDT optimization. Therefore, the weights modify the relative contributions of source–voxel pairs and can change the optimization direction rather than merely rescaling the final score. A result is published only when optimization converges, its transformation probability exceeds the acceptance threshold, and the correction from the prediction remains bounded.

IV. Evaluation

This section describes the experimental setup and evaluates how structural weighting and source ground subsampling affect localization accuracy and runtime on real-world driving sequences. Classical NDT is used as the primary baseline because the proposed method targets fixed-map scan-to-map localization. NDT-LOAM is not included as a direct baseline because it is a complete odometry and SLAM framework that applies weighted NDT within a local-map motion-estimation pipeline.

A. Experimental Setup

The KITTI Odometry dataset [11] was used to evaluate sequences 00, 01, and 03, representing residential, highway, and rural environments, respectively. Their trajectories are shown in Fig. 1.

The sequences were acquired using a 64-channel Velodyne HDL-64E, and each scan was projected onto a 64×1800 range image. Classical and segmented maps were generated from the same timestamped poses and serialized as precomputed 1.0 m NDT voxel maps. The classical map was constructed from unsegmented raw points, whereas the segmented map was constructed from retained ground and non-ground points. All methods used a 0.2 m source voxel size, DIRECT7 neighbor search, identical registration criteria, and four OpenMP threads on an Intel Core i5-12400F CPU. Processing time includes the complete point-cloud callback. ATE is computed as the root-mean-square position error after timestamp matching, excluding the first 3 s to remove initialization transients.

B. Localization Results

We compare Classical NDT, Unweighted-S50, Weighted-S100, and Weighted-S50. Unweighted-S50 uses the same segmentation and 50% source ground subsampling as Weighted-S50 but disables all structural, geometric, and compatibility weights. The comparison between Unweighted-S50 and Weighted-S50 therefore isolates the effect of the proposed weighting. Weighted-S100 retains all segmented source points, whereas Weighted-S50 retains all non-ground source points and 50% of source ground points.

Table I summarizes localization accuracy and runtime after excluding the first 3 s of each sequence. Across the three sequences, Weighted-S50 achieves the lowest macro-averaged ATE, position P95, yaw MAE, yaw P95, and mean runtime among the evaluated methods. Compared with Classical NDT, Weighted-S50 reduces macro-averaged ATE from 0.1621 m to 0.1498 m, macro-averaged yaw MAE from 0.2232° to 0.1912°, and macro-averaged mean runtime from 47.693 ms to 38.088 ms. For Weighted-S50, both the mean and P95 processing times remain below the 100 ms frame period corresponding to 10 Hz operation on all three sequences. These results show that the proposed combination of weighting

TABLE I: KITTI localization results on representative residential, highway, and rural sequences.

Seq.	Method	Time mean / P95 (ms)	ATE/RMSE (m)	Position P95 / Max (m)	Yaw MAE / P95 (deg)
00	Classical NDT	43.508 / **61.868**	**0.3026**	**0.8042** / 1.0890	0.2479 / 0.6661
00	Unweighted-S50	46.931 / 104.699	0.3278	0.9131 / **1.0858**	0.2529 / 0.6203
00	Weighted-S100	44.377 / 86.306	1.4175	1.0109 / 8.7111	0.2837 / 0.6294
00	Weighted-S50	**41.381** / 69.075	0.3104	0.8719 / 1.0944	**0.2388** / **0.5729**
01	Classical NDT	53.542 / 79.229	0.0998	0.2115 / 0.5199	0.2685 / 0.5964
01	Unweighted-S50	34.395 / 46.027	0.0716	0.1464 / **0.2759**	0.2295 / 0.4557
01	Weighted-S100	36.621 / 49.749	**0.0706**	0.1459 / 0.3091	0.1989 / **0.3907**
01	Weighted-S50	**32.580** / **44.568**	0.0708	**0.1431** / 0.3142	**0.1974** / 0.3924
03	Classical NDT	46.029 / 63.067	0.0839	0.1270 / 0.3138	0.1533 / 0.4049
03	Unweighted-S50	43.304 / 76.314	0.0666	**0.0998** / 0.2847	0.1586 / 0.4103
03	Weighted-S100	44.985 / 65.393	**0.0636**	0.1042 / **0.2774**	0.1387 / 0.3499
03	Weighted-S50	**40.302** / **56.880**	0.0682	0.1051 / 0.2875	**0.1374** / **0.3341**
Macro Avg.	Classical NDT	47.693 / 68.055	0.1621	0.3809 / 0.6409	0.2232 / 0.5558
Macro Avg.	Unweighted-S50	41.543 / 75.680	0.1553	0.3864 / **0.5488**	0.2137 / 0.4954
Macro Avg.	Weighted-S100	41.994 / 67.149	0.5172	0.4203 / 3.0992	0.2071 / 0.4567
Macro Avg.	Weighted-S50	**38.088** / **56.841**	**0.1498**	**0.3734** / 0.5654	**0.1912** / **0.4331**

and source ground subsampling reduces mean runtime while improving macro-averaged scan-to-map localization accuracy.

The ablation study separates the effects of point reduction and weighting. Unweighted-S50 improves the macro-averaged ATE and mean runtime compared with Classical NDT, suggesting that range-image segmentation and source ground subsampling contribute to the gain. Weighted-S50 further improves over Unweighted-S50 in macro-averaged ATE, position P95, yaw MAE, yaw P95, and both mean and P95 runtimes, indicating that the proposed weighting provides an additional benefit beyond point reduction alone. The benefit is most pronounced on sequence 01, where Weighted-S50 reduces ATE by 29.0%, position P95 by 32.3%, yaw MAE by 26.4%, and mean runtime by 39.2% compared with Classical NDT. Relative to Classical NDT, Weighted-S50 also improves all reported metrics on sequence 03. On sequence 00, Classical NDT gives slightly lower ATE and position P95, but Weighted-S50 reduces mean runtime and yaw error. Weighted-S100 is less robust on sequence 00, producing a large maximum position error. In contrast, Weighted-S50 lowers the macro-averaged mean and P95 runtimes and the macro-averaged maximum position error relative to Weighted-S100, suggesting that source ground subsampling reduces sensitivity to redundant ground measurements and large tail errors.

V. Conclusion

This paper presented a segmentation-aware weighted NDT method for CPU-only LiDAR localization. The proposed method combines structural and geometric weighting with deterministic source ground subsampling to adjust the contribution of ground and non-ground measurements during scan matching. On representative KITTI sequences, Weighted-S50 reduced the macro-averaged ATE from 0.1621 m to 0.1498 m, the yaw MAE from 0.2232° to 0.1912°, and the processing time from 47.693 ms to 38.088 ms compared with Classical NDT. The results also suggest that retaining all ground points can increase sensitivity to redundant ground measurements, supporting the use of source ground subsampling. Future work will investigate adaptive-resolution NDT to further improve the trade-off among localization accuracy, map size, and computational cost.

References

[1] D. Kim, J. Paek, and S. Bahk, "Lucid: LiDAR-based UAV Classification and Detection for Airspace Surveillance," in *The 46th IEEE International Conference on Distributed Computing Systems (ICDCS'26)*, June 2026.

[2] H. Yin, X. Xu, S. Lu, X. Chen, R. Xiong, S. Shen, C. Stachniss, and Y. Wang, "A Survey on Global LiDAR Localization: Challenges, Advances and Open Problems," *International Journal of Computer Vision*, 2024.

[3] M. Choi, J. Ryu, Y. Son, S. Cho, and J. Paek, "LiDAR-based Localization for Autonomous Vehicles - Survey and Recent Trends," in *The 15th International Conference on Information and Communication Technology Convergence (ICTC'24)*. IEEE, Oct 2024, pp. 456–460.

[4] K. Koide, M. Yokozuka, S. Oishi, and A. Banno, "Voxelized GICP for Fast and Accurate 3D Point Cloud Registration," in *IEEE international conference on robotics and automation (ICRA)*, 2021.

[5] Rusinkiewicz, S. and Levoy, M., "Efficient Variants of the ICP Algorithm," in *Proceedings Third International Conference on 3-D Digital Imaging and Modeling*, 2001, pp. 145–152.

[6] M. Magnusson, "The Three-Dimensional Normal-Distributions Transform: an Efficient Representation for Registration, Surface Analysis, and Loop Detection," Ph.D. dissertation, Örebro universitet, 2009.

[7] S. Chen, H. Ma, C. Jiang, B. Zhou, W. Xue, Z. Xiao, and Q. Li, "NDT-LOAM: A Real-Time LiDAR Odometry and Mapping With Weighted NDT and LFA," *IEEE Sensors Journal*, 2021.

[8] J. Ryu, C. Shin, and J. Paek, "Iterative Closest Point of Object-based Point Clouds for Multi-View Extension," in *The 14th International Conference on Information and Communication Technology Convergence (ICTC'23)*. IEEE, Oct 2023, pp. 543–545.

[9] J. Ryu and J. Paek, "Poster: Fast Field-of-View Expansion for Collaborative Object Detection," in *The 22nd ACM International Conference on Mobile Systems, Applications, and Services (MobiSys'24)*, June 2024.

[10] T. Shan and B. Englot, "LeGO-LOAM: Lightweight and Ground-Optimized LiDAR Odometry and Mapping on Variable Terrain," in *IEEE/RSJ international conference on intelligent robots and systems (IROS)*, 2018.

[11] A. Geiger, P. Lenz, and R. Urtasun, "Are we ready for Autonomous Driving? The KITTI Vision Benchmark Suite," in *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2012.