

# Online Pole Segmentation on Range Images for Long-term LiDAR Localization in Urban Environments

Hao Dong\* Xieyuanli Chen\*† Simo Särkkä Cyrill Stachniss

## ARTICLE INFO

### Keywords:

Localization  
Pole  
LiDAR  
Range image  
Mapping  
Autonomous driving  
Deep learning  
Semantic segmentation

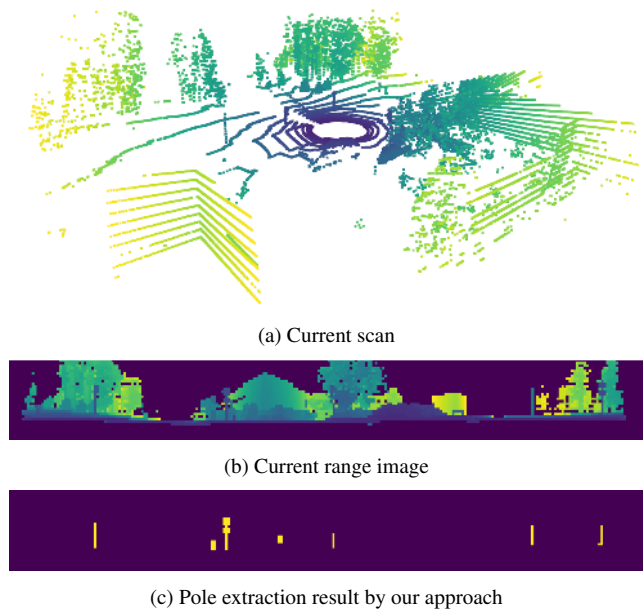
## ABSTRACT

Robust and accurate localization is a basic requirement for mobile autonomous systems. Pole-like objects, such as traffic signs, poles, and lamps are frequently used landmarks for localization in urban environments due to their local distinctiveness and long-term stability. In this paper, we present a novel, accurate, and fast pole extraction approach based on geometric features that runs online and has little computational demands. Our method performs all computations directly on range images generated from 3D LiDAR scans, which avoids processing 3D point clouds explicitly and enables fast pole extraction for each scan. We further use the extracted poles as pseudo labels to train a deep neural network for online range image-based pole segmentation. We test both our geometric and learning-based pole extraction methods for localization on different datasets with different LiDAR scanners, routes, and seasonal changes. The experimental results show that our methods outperform other state-of-the-art approaches. Moreover, boosted with pseudo pole labels extracted from multiple datasets, our learning-based method can run across different datasets and achieve even better localization results compared to our geometry-based method. We released our pole datasets to the public for evaluating the performance of pole extractors, as well as the implementation of our approach.

## 1. Introduction

Robust and precise localization is a crucial capability for an autonomous robot and a commonly performed state estimation task [1]. The accurate estimation of the robot's pose helps to avoid collisions, navigate in a goal-directed manner, follow the traffic lanes, and perform other tasks. Reliability here means that the robot should adapt to changes in the environment, such as different weather conditions [5], day and night [39], or seasonal changes [25].

Global navigation satellite system-based localization systems are robust to appearance changes of the environments. However, in urban areas, they may suffer from low availability due to building and tree occlusions. Additional, map-based approaches are needed for precise and reliable localization for mobile robots. Multiple different types of sensors have been used to build the map of the environments, including Light Detection and Ranging (LiDAR) scanners [18, 3, 41], monocular [28] and stereo cameras [15]. Among them, LiDAR sensors are more robust to the illumination changes, and multiple LiDAR-based effective and efficient mapping approaches have been proposed, for example, by Behley and Stachniss [3] or by Droschel and Behnke [18]. However, these approaches often need substantial amounts of memory due to map representations, thus cannot generalize easily to large-scale scenes. If only



**Figure 1:** Visualization of range image and pole extraction. On the top is the raw LiDAR scan. The corresponding range image generated from this scan is in the middle. The bottom is the pole extraction result based on the range image.

specific features are used to build the map, such as traffic signs, trunks and other pole-like structures, the map size can be reduced significantly [45].

The main contribution of this paper is a novel range image-based pole extractor that can be used for long-term localization of autonomous mobile systems. Instead of using the raw point clouds obtained from 3D LiDAR sensors directly, we investigate the use of range images for pole extraction. Range image is a light and natural representation of the scan from a rotating 3D LiDAR such as a Velodyne or Ouster sensors. Operating on such an image is considerably

\*Authors with equal contribution.

† Corresponding author: xieyuanli.chen@igg.uni-bonn.de

H. Dong and S. Särkkä are with the Aalto University, Finland. X. Chen and C. Stachniss are with the University of Bonn, Germany. C. Stachniss is additionally with the Department of Engineering Science at the University of Oxford, UK, and with the Lamarr Institute for Machine Learning and Artificial Intelligence, Germany. This work has been funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy, EXC-2070 – 390732324 (PhenoRob).

ORCID(s):

faster than on the raw 3D point cloud. Besides, a range image keeps the neighborhood information implicitly in its 2D structure and we can use this information for segmentation. The detected poles in the range image can further be used as pseudo pole labels to train a pole segmentation neural network. After training once with pseudo pole labels generated from different datasets, our learning-based method can detect poles in different environments and achieve even better localization performance than our geometry-based method. To achieve LiDAR localization, in the mapping phase, we first project the raw point cloud into a range image and then extract poles from that image, as shown in Fig. 1. After obtaining the position of poles in the range image, we use the ground-truth poses of the robot to reproject them into the global coordinate system to build a global map. During localization, we utilize Monte Carlo localization (MCL) for updating the importance weights of the particles by matching the poles detected from online sensor data with the poles in the global map.

In sum, we make three key claims that our approach is able to (i) extract more reliable poles in the scene compared to the baseline method, as a result, (ii) achieve better online localization performance in different environments, and (iii) generate pseudo pole labels to train a pole segmentation network achieving better localization results and faster runtime compared to the geometric method. These claims are backed up by the paper and our experimental evaluation. The code of our approach and the pole dataset are released at: <https://github.com/PRBonn/pole-localization>.

## 2. Related Work

For localization given a map, there exists a large amount of research. While many different types of sensors have been used to tackle this problem [37], in this work, we mainly concentrate on LiDAR-based approaches.

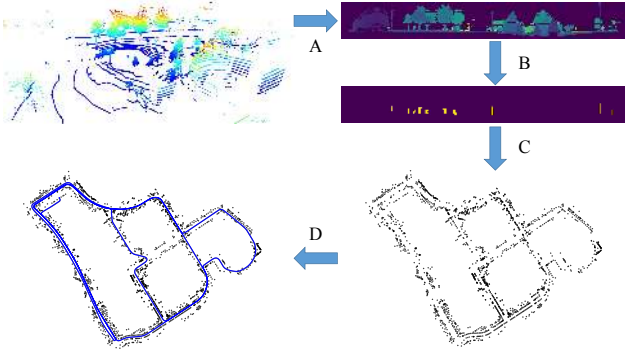
Traditional approaches to robot localization rely on probabilistic state estimation techniques. A popular framework is Monte Carlo localization [16], which uses a particle filter to estimate the pose of the robot and is still widely used in robot localization systems [4, 9, 13, 22, 36, 40, 47].

Besides the traditional geometry-based methods, more and more approaches recently exploit deep neural networks and semantic information for 3D LiDAR localization. For example, Ma et al. [24] combine semantic information such as lanes and traffic signs in a Bayesian filtering framework to achieve accurate and robust localization within sparse HD maps, whereas Tinchev et al. [38] propose a learning-based method to match segments of trees and localize in both urban and natural environments. Sun et al. [36] use a deep-probabilistic model to accelerate the initialization of the Monte Carlo localization and achieve a fast localization in outdoor environments. Shi et al. [33] exploit a graph-based network to register LiDAR point clouds. Wiesmann et al. [44] propose a deep learning-based 3D network to compress the LiDAR point cloud, which can be used for large-scale LiDAR localization. In our previous work [7, 8, 9], we

also exploit CNNs with semantics to predict the overlap between LiDAR scans as well as their yaw angle offset and use this information to build a learning-based observation model for Monte Carlo localization. The learning-based methods perform well in the trained environments, while they usually cannot generalize well in different environments or different LiDAR sensors.

Instead of using dense semantic information estimated by neural networks [27, 26, 23, 14], a rather lighter solution has been proposed for long-term localization, which extracts only pole landmarks from point clouds. There are usually two parts in pole-based approaches, pole extraction and pose estimation. For pole extraction, Sefati et al. [32] first remove the ground plane from the point cloud and project the rest points on a horizontal grid. After that, they cluster the grid cells and fit a cylinder for each cluster. Finally, a particle filter with nearest-neighbor data association is used for pose estimation. Weng et al. [42] and Schaefer et al. [30] use similar particle filter-based methods to estimate the pose of the robot with different pole extractors. Weng et al. [42] discretize the space and extract poles based on the number of laser reflections in each voxel. Based on that, Schaefer et al. [30] consider both the starting and end points of the scan and thus model the occupied and free space explicitly. Kümmerle et al. [21] use a nonlinear least-squares optimization method to refine the pose estimation. Spangenberg et al. [35] use stereo camera images to extract poles and then feed them into a particle filter with odometry and GPS data. Shi et al. [34] extract pole-like objects from the point cloud by spatial independence analysis and cylindrical or linear feature detection. They also classify the pole-like objects into street lamps, traffic signs and utility poles by 3D shape matching. Weng et al. [43] exploit the reflective intensity information to extract traffic signs which are always painted with highly reflective materials. Chen et al. [6] fuse poles information into a non-linear optimization problem to obtain the vehicle location. Plachetka et al. [29] use a deep neural network for pole extraction by learning encodings of the point cloud input. In contrast to the aforementioned approaches, we use a projection-based method and avoid the comparable costly processing of 3D point cloud data. Thus, our implementation is fast.

This article is an extension of our previous conference paper [17]. In our previous work, we propose a geometry-based pole extractor on LiDAR point clouds, which uses only range information without exploiting neural networks or deep learning. Thus, it generalizes well to different environments and different LiDAR sensors and does not require new training data when moving to different environments. Inspired by an automatically labeling method [11], in this article, we further use the poles extracted by our geometry-based method as pseudo labels to train a pole segmentation network. Trained with a large number of pseudo pole labels automatically generated by our geometry-based pole extractor on different datasets, our learning-based method can generalize well in different environments and outperforms



**Figure 2:** Overview of our pole-based localization system. A, we project the LiDAR point cloud into a range image, and B, extract poles in the image. C, based on the extracted poles, we then build a global pole map of the environment. D, we finally propose a pole-based observation model for MCL to localize the robot in the map.

the geometry-based method. As for the network architecture, we use SalsaNext [14], a state-of-the-art range image-based semantic segmentation network. Instead of segmenting the environment into multiple classes like ground, structure, vehicle and human, in our case, we only distinguish the poles from other objects.

### 3. Methodology

In this paper, we propose a range image-based pole extractor for long-term localization using a 3D LiDAR sensor. As shown in Fig. 2, we first project the LiDAR point cloud into a range image (Sec. 3.1) and extract poles from it using either a geometric (Sec. 3.2) or a learning-based (Sec. 3.3) method. Based on the proposed pole extractor, we then build a global pole map of the environment (Sec. 3.4). In the localization phase, we extract poles online using the same extractor and use a novel pole-based observation model for Monte Carlo localization (Sec. 3.5).

#### 3.1. Range Image Generation

The key idea of our approach is to use range images generated from LiDAR scans for pole extraction. Following the prior work [12, 13], we utilize a spherical projection for range images generation. Each LiDAR point  $\mathbf{p} = (x, y, z)$  is mapped to spherical coordinates via a mapping  $\Pi : \mathbb{R}^3 \mapsto \mathbb{R}^2$  and finally to image coordinates, as defined by

$$\begin{pmatrix} u \\ v \end{pmatrix} = \begin{pmatrix} \frac{1}{2} [1 - \arctan(y, x) \pi^{-1}] w \\ [1 - (\arcsin(z r^{-1}) + |f_{\text{down}}|) f^{-1}] h \end{pmatrix}, \quad (1)$$

where  $(u, v)$  are image coordinates,  $(h, w)$  are the height and width of the desired range image,  $f = |f_{\text{up}}| + |f_{\text{down}}|$  is the vertical field-of-view (FOV) of the sensor ( $f_{\text{up}}$  is the up vertical FOV and  $f_{\text{down}}$  is the down vertical FOV). The  $f_{\text{up}}$  is a positive value in radian, while  $f_{\text{down}}$  is a negative value in radian. The  $r = \|\mathbf{p}\|_2$  is the range value of each point. This procedure results in a list of  $(u, v)$  tuples containing

#### Algorithm 1: Range Image-based Pole Extraction

**Input:** Range Image  $\mathbf{I}_{\text{range}}$

**Output:** PoleParameters  $\mathbf{P}$  with circle centers and radiuses


```

1 Let  $\mathbf{O}$  be the set of all valid pixel coordinates in  $\mathbf{I}_{\text{range}}$ .  $T_d$ 
  is the distance threshold to find neighbors,  $T_n$  is the
  pixel count threshold, and  $T_h$  is the object height
  threshold.  $H_a$  and  $H_b$  are the height lower bound and
  higher bound.  $R_a$  and  $R_b$  are the radius lower bound and
  higher bound.
2 while  $\mathbf{O} \neq \emptyset$  do
3   create a new  $\mathbf{c}$  in  $\mathbf{C}$ ;  $\mathbf{p} \leftarrow \mathbf{O}[0]$ ;  $\mathbf{O} \leftarrow \mathbf{O} - \mathbf{p}$ ;
    $\mathbf{c} \leftarrow \mathbf{c} + \{\mathbf{p}\}$ ; create an empty  $\mathbf{N}$ 
4   foreach  $\mathbf{p}' = \text{Neighbor}(\mathbf{p}) \in \mathbf{O}$  and
      $\text{Distance}(\mathbf{p}', \mathbf{p}) < T_d$  do
5      $\mathbf{N} \leftarrow \mathbf{N} + \mathbf{p}'$ 
6   end
7   while  $\mathbf{N} \neq \emptyset$  do
8      $\mathbf{p} \leftarrow \mathbf{N}[0]$ ;  $\mathbf{N} \leftarrow \mathbf{N} - \mathbf{p}$ ;  $\mathbf{O} \leftarrow \mathbf{O} - \mathbf{p}$ ;  $\mathbf{c} \leftarrow \mathbf{c} + \{\mathbf{p}\}$ 
9     foreach  $\mathbf{p}' = \text{Neighbor}(\mathbf{p}) \in \mathbf{O}$  and
        $\text{Distance}(\mathbf{p}', \mathbf{p}) < T_d$  do
10       $\mathbf{N} \leftarrow \mathbf{N} + \text{Neighbor}(\mathbf{p}')$ 
11    end
12  end
13   $N_{\text{pixel}} \leftarrow$  the number of pixels in  $\mathbf{c}$ 
14  if  $N_{\text{pixel}} < T_n$  then
15     $\mathbf{C} \leftarrow \mathbf{C} - \mathbf{c}$ 
16  end
17 end
18 foreach  $\mathbf{c} \in \mathbf{C}$  do
19    $w, h \leftarrow \text{Width}(\mathbf{c}), \text{Height}(\mathbf{c})$ 
20    $N_{\text{SmallR}} \leftarrow$  the number of pixels in  $\mathbf{c}$  whose range
     value is smaller than its neighbor outside  $\mathbf{c}$ 
21   if  $h/w < 1$  or  $N_{\text{SmallR}} < \delta \cdot \text{Len}(\mathbf{c})$  then
22      $\mathbf{C} \leftarrow \mathbf{C} - \mathbf{c}$ 
23   end
24 end
25 foreach  $\mathbf{c} \in \mathbf{C}$  do
26    $\mathbf{x}, \mathbf{y}, \mathbf{z} \leftarrow$  3D coordinates of pixels in  $\mathbf{c}$ 
27   if  $\max(\mathbf{z}) > H_a$  and  $\min(\mathbf{z}) < H_b$  and
      $(\max(\mathbf{z}) - \min(\mathbf{z})) > T_h$  then
28      $x_c, y_c, r_c \leftarrow \text{FitCircle}(\mathbf{x}, \mathbf{y})$ 
29      $N_{\text{Free}} \leftarrow$  the number of the pixels in a small free
       space outside the radius of the pole
30     if  $r_c < R_a$  and  $r_c > R_b$  and  $N_{\text{Free}} < \mu \cdot \text{Len}(\mathbf{z})$ 
       then
31        $\mathbf{P} \leftarrow \mathbf{P} + \{x_c, y_c, r_c\}$ 
32     end
33   end
34 end
    
```

a pair of image coordinates for each  $\mathbf{p}_i$ , which we use to generate our proxy representation. Using these indices, we extract for each  $\mathbf{p}_i$ , its range  $r$ , its  $x$ ,  $y$ , and  $z$  coordinates, and store them in the image.

#### 3.2. Geometry-based Pole Extractor

We extract poles based on the range images generated in the previous step. The general intuition behind our pole



**Figure 3:** Visualization of results on each step of our geometric pole extractor. The first image shows the range image. The second image represents the clustering result and the third one shows the pole candidates after applying 2D geometric constraints. The last one is the final pole extraction result.

extraction algorithm is that the range values of the poles are usually significantly smaller than the backgrounds. Based on this idea and as specified in Alg. 1, our first step is to cluster the pixels of the range image into different small regions based on their range values. We first pass through all pixels in the range image, from top to bottom, left to right. We put all pixels with valid range data in an open set  $\mathbf{O}$ . For each valid pixel  $\mathbf{p}$ , we check its neighbors including the left, right and below ones. If there exists a neighbor with a valid value and the range difference between the current pixel and its neighbor is smaller than a threshold  $T_d$ , we add the current pixel to a cluster set  $\mathbf{c}$  and remove it from the open set  $\mathbf{O}$ . We do the same check iteratively with the neighbors until no neighbor pixel meets the above criteria, and we then get a cluster of pixels. After checking all the pixels in  $\mathbf{O}$ , we will get a set  $\mathbf{C}$  with several clusters and each cluster represents one object. If the number of pixels in one cluster is smaller than a threshold  $T_n$ , we regard it as an outlier and ignore it.

The next step is to extract poles from these objects using 2D geometric constraints. To this end, we exploit both the range information and the 3D coordinates  $(x, y, z)$  of each pixel. We first check the aspect ratio of each cluster. Since we are only interested in pole-like objects, whose height is usually larger than its width, we therefore discard the cluster with aspect ratio  $h/w < 1$ . Another heuristic we use is the fact that a pole usually stands alone and has a significant distance from background objects.  $N_{SmallR}$  is the number of points in cluster  $\mathbf{c}$  whose range value is smaller than its neighbor outside  $\mathbf{c}$ , we discard the cluster if  $N_{SmallR}$  is smaller than  $\delta$  times the number of all points in the cluster.

To exploit the 3D coordinates  $(x, y, z)$  of each pixel, we calculate  $\max(z) - \min(z)$  of each cluster and only take a cluster as a pole candidate if  $\max(z) - \min(z) > T_h$ . Besides, we are only interested in poles whose height is higher than  $H_a$ . Based on experience, we also set a threshold  $H_b$  for the lowest position of the pole to filter outliers. For each pole candidate, we then fit a circle using the  $x$  and  $y$  coordinates of all points in the cluster and get the center and

the radius of that pole. We filter out the candidates with too small or too large radiuses and candidates that connect to other objects by checking the free space around them. After the above steps, we finally extract the positions and radiuses of poles. As an example, Fig. 3 visualize the intermediate results on each step of our geometric pole extractor.

### 3.3. Learning-based Pole Segmentation Trained With Pseudo Labels

As shown in [11], geometric information can be used to automatically generate labels for training a LiDAR-based moving object segmentation network and achieve good performance in various environments. Such auto-labeling methods enable network learning in a self-supervised manner, which saves the extensive manual labeling effort and improves the generalization ability of the learning-based method. Inspired by it, we use the poles detected by our geometry-based pole extractor to generate pseudo labels to train an online pole segmentation network.

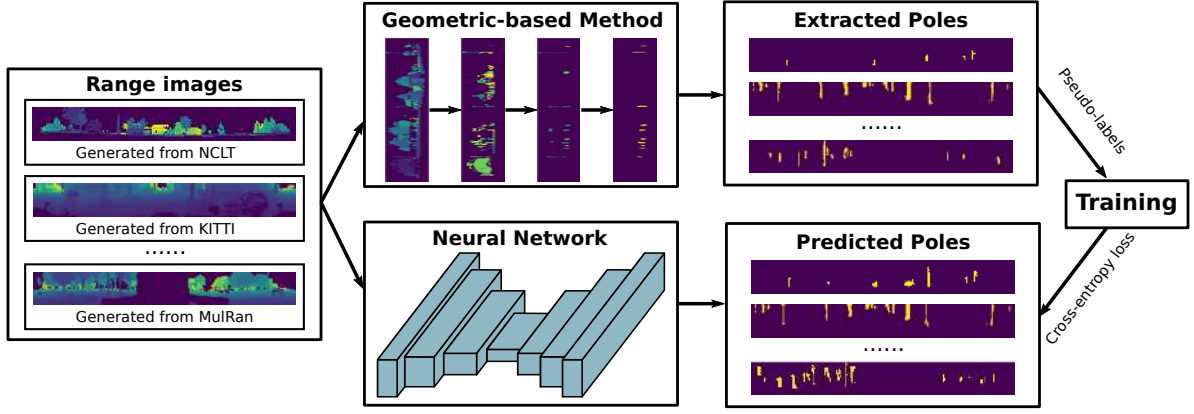
We use our geometry-based method to generate pseudo pole labels from NCLT [5], SemanticKITTI [2], and MulRan [20] datasets. In this work, we do not design a new network architecture but reuse networks that have been successfully applied to LiDAR-based semantic segmentation in the past. We adopt and evaluate SalsaNext [14], an encoder-decoder architecture with solid performances on semantic segmentation tasks. SalsaNext [14] achieves state-of-the-art performances on SemanticKITTI dataset among all range image-based semantic segmentation networks. Therefore, we choose it as the base network architecture for our learning-based method. Instead of segmenting the environment into multiple classes like ground, structure, vehicle and human, in our case, we only distinguish the poles from other objects. After the segmentation, similar filtering steps as used in the geometric method are applied to remove outliers. SalsaNext network is comparably lightweight, and can achieve real-time operation, i.e., run faster than the commonly used frame rate of the employed LiDAR sensor, which is 10 Hz for Ouster and Velodyne scanners. For more detailed information about the network, we refer to the original paper [14].

For training the segmentation network, we directly feed them with the range images plus the pseudo pole labels generated from our geometry-based pole extractor. We use the same loss functions as used in the original segmentation methods, while mapping all classes into two per-point classes, poles and non-poles. We retrain the network and evaluate the pole extracting performance with our pole datasets and also localization tasks. Fig. 4 shows the training pipeline of our proposed learning-based pole segmentation method. Note that, we train the network with pseudo pole labels generated from different datasets, and later use the same model to extract poles in different environments.

### 3.4. Pole-based Mapping

To build the global map for localization, we follow the same setup as introduced by Schaefer et al. [30], splitting





**Figure 4:** Overview of our learning-based pole extraction approach. We first use the geometric method to generate range images and pseudo pole labels on multiple datasets, including NCLT, KITTI, and MulRan. Then, we train a deep neural network to extract pole-like objects directly on the range images. Our learning-based pole extractor generalizes well to all three datasets with a single model.

the ground-truth trajectory into shorter sections with equal length, extracting poles in these sections separately and finally merging them into a global pole map. Since the provided poses are not very accurate for mapping [30], instead of aggregating a noisy submap, we only use the middle LiDAR scan of each section to extract poles. We merge multiple overlapped pole detections by averaging over their centers and radiuses and apply a counting model to filter out the dynamic objects. Only those candidate poles that appear multiple times in continuous sections are added to the map.

### 3.5. Monte Carlo Localization

Monte Carlo localization (MCL) is commonly implemented using a particle filter [16]. MCL realizes a recursive Bayesian filter estimating a probability density  $p(\mathbf{x}_t | \mathbf{z}_{1:t}, \mathbf{u}_{1:t})$  over the pose  $\mathbf{x}_t$  given all observations  $\mathbf{z}_{1:t}$  and motion controls  $\mathbf{u}_{1:t}$  up to time  $t$ . This posterior is updated as follows:

$$p(\mathbf{x}_t | \mathbf{z}_{1:t}, \mathbf{u}_{1:t}) = \eta p(\mathbf{z}_t | \mathbf{x}_t) \int p(\mathbf{x}_t | \mathbf{u}_t, \mathbf{x}_{t-1}) p(\mathbf{x}_{t-1} | \mathbf{z}_{1:t-1}, \mathbf{u}_{1:t-1}) d\mathbf{x}_{t-1}, \quad (2)$$

where  $\eta$  is a normalization constant,  $p(\mathbf{x}_t | \mathbf{u}_t, \mathbf{x}_{t-1})$  is the motion model,  $p(\mathbf{z}_t | \mathbf{x}_t)$  is the observation model, and  $p(\mathbf{x}_{t-1} | \mathbf{z}_{1:t-1}, \mathbf{u}_{1:t-1})$  is the probability distribution for the prior state  $\mathbf{x}_{t-1}$ .

In our case, each particle represents a hypothesis for the 2D pose  $\mathbf{x}_t = (x, y, \theta)_t$  of the robot at time  $t$ . When the robot moves, the pose of each particle is updated based on a motion model with the control input  $\mathbf{u}_t$  or the odometry measurements. For the observation model, the weights of the particles are updated based on the difference between expected observations and actual observations. The observations are the positions of the poles. We match the online observed poles with the poles in the map via nearest-neighbor search using a k-d tree. The likelihood of the  $j$ -th

particle is then approximated using a Gaussian distribution:

$$p(\mathbf{z}_t | \mathbf{x}_t) \propto \prod_i^N \left( \exp \left( -\frac{1}{2} \frac{d(\mathbf{z}_t^i, \mathbf{z}_j^i)^2}{\sigma_d^2} \right) + \epsilon \right), \quad (3)$$

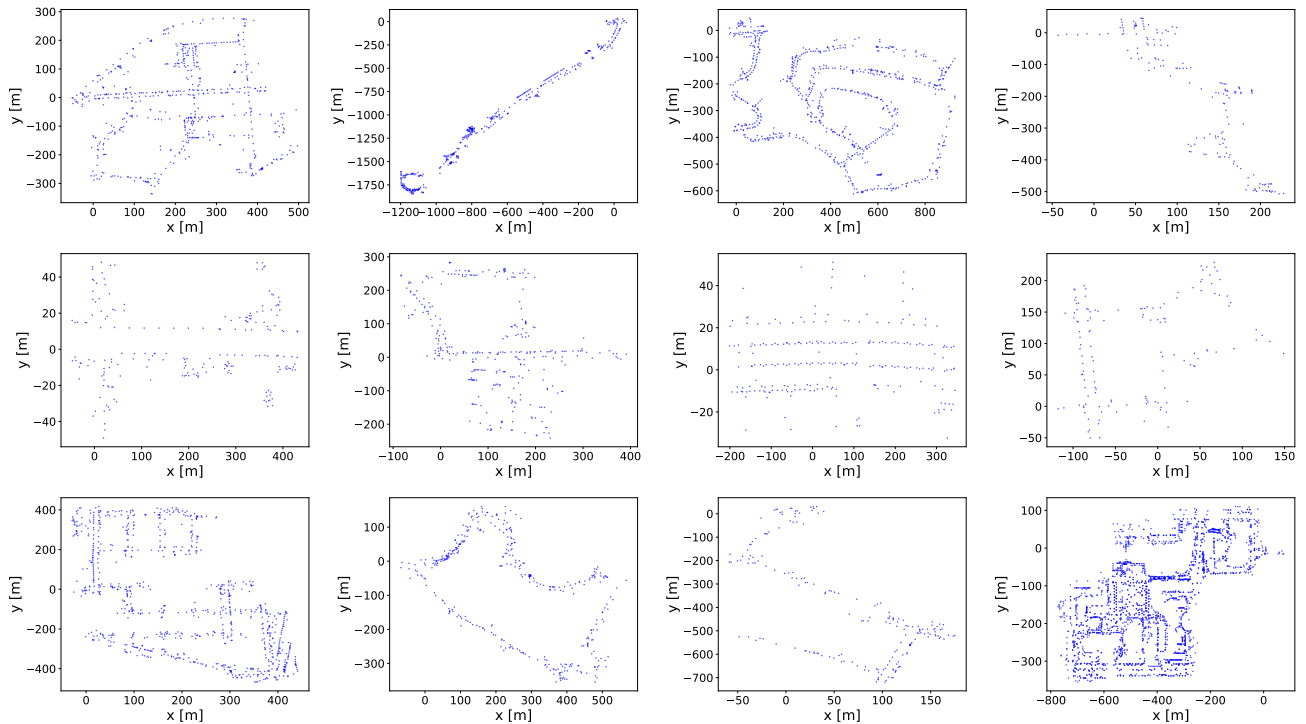
where  $N$  is the number of matched poles in the current scan,  $\sigma_d^2$  is the position uncertainty of the poles,  $d$  corresponds to the difference between the online observed pole  $\mathbf{z}_t^i$  and matched pole in the map  $\mathbf{z}_j^i$  given the position of the particle  $j$ . We use the Euclidean distance between the pole positions to measure this difference. The constant  $\epsilon$  accounts for the probability that a detected pole is not part of the map. This constant is crucial for the robustness of localization when there are many outliers. If the number of effective particles decreases below a threshold [19], the resampling process is triggered and particles are sampled based on their weights.

## 4. Experimental Evaluation

The main focus of this work is an accurate and efficient pole extractor for long-term LiDAR localization. We present our experiments to show the capabilities of our method. The experiments furthermore support our key claims that our method is able to: (i) extract more reliable poles in the environment compared to the baseline method, as a result, (ii) achieve better online localization performance in different environments, and (iii) generate pseudo pole labels to train a pole segmentation network achieving better localization results and faster runtime compared to the geometric method.

### 4.1. Datasets for Pole Extraction and LiDAR Localization

There are few public datasets available to evaluate pole extraction performance. To this end, we label the poles in session 2012-01-08 of NCLT dataset by hand and release



**Figure 5:** Our pole extraction datasets. The first 11 figures show the ground truth pole positions of sequence 00-10 of SemanticKITTI, and the last one (at the right bottom corner) shows pole positions of the session 2012-01-08 of NCLT. Each blue dot represents the position of one pole.

this dataset for public research use. For the reason that the original NCLT ground-truth poses are inaccurate [30], the aggregated point cloud is a little blurry. Therefore, to create the ground-truth pole map of the environment, we partition the ground-truth trajectory into shorter segments of equal length. For each segment, we aggregate the point cloud together and use Open3D [46] to render and label the pole positions. We only label those poles with high certainty and ignore those blurry ones. Besides our own labelled data, we also reorganize the SemanticKITTI [2] dataset sequences 00-10 by extracting the pole-like objects like traffic signs, poles and trunk, and then clustering the point clouds to generate the ground-truth pole instances.

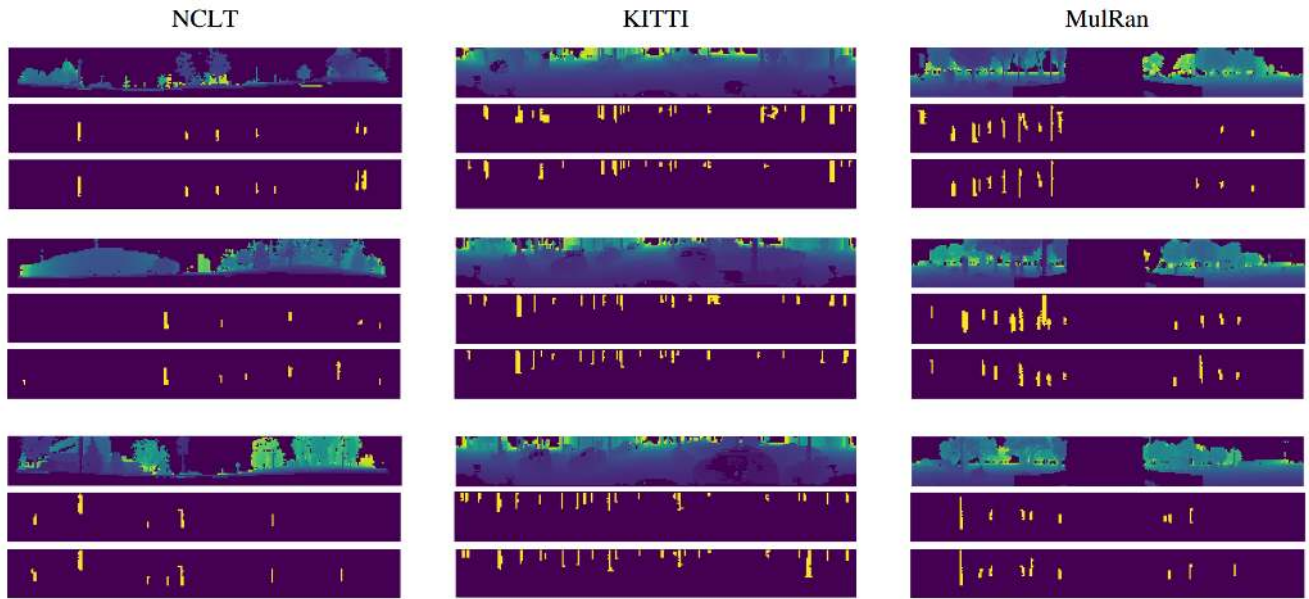
To assess the localization reliability and accuracy of our method, we use NCLT dataset [5] and MulRan dataset [20]. These two datasets are collected in different environments (U.S., Korea) with different LiDAR sensors (Velodyne HDL-32E, Ouster OS1-64). In these two datasets, the robot passes through the same place multiple times with month-level temporal gaps, hence ideal to test the long-term localization performance. We compare our methods to both a pole-based method proposed by Schaefer et al. [30] and the range image-based method proposed by Chen et al. [13]. We reproduce their results using the public available codes. For the SemanticKITTI dataset, there is no overlap area between different sequences for evaluating long-term localization. Therefore, we only used the extracted pole labels from the SemanticKITTI dataset to train our network. Fig. 5 shows examples of our proposed pole datasets.

## 4.2. Pole Extractor Performance

The first experiment evaluates the pole extraction performance of our approach and supports the claim that our range image-based method outperforms the baseline method in pole extraction.

We evaluate both our geometry-based pole extractor, named Ours-G, and our learning-based pole segmentation method, named Ours-L. For training the pole segmentation network, we use data from multiple datasets, including the session 2012-01-08 in NCLT dataset, sequence KAIST 02 in MulRan dataset and sequence 00-02, 05-09 in SemanticKITTI dataset. For validation, we use sequences 03 and 04 in SemanticKITTI dataset and sequence 10 for testing. We train the network for 150 epochs using stochastic gradient descent with an initial learning rate of 0.01 and the learning rate decay is 0.01. The batch size is 12 and the spatial dropout probability is 0.2. The size of the range image is  $32 \times 256$  and the valid range values are normalized between 0 and 1. To prevent overfitting, we augmented the data by applying a random rotation or translation, flipping randomly around the y-axis with a probability of 0.5. During the matching phase, we find the matches via nearest-neighbor search using a k-d tree with 1 m distance bounds.

Tab. 1 summarizes the precision, recall and F1 score of our method and Schaefer et al. [30] with respect to the ground-truth pole map on both NCLT dataset and SemanticKITTI dataset. As can be seen, our methods achieve better performance and extract more poles in both environments compared to the baseline method. Compared to our geometry-based pole extractor, our learning-based method



**Figure 6:** Pole extraction examples using our geometric and learning-based pole extractor. For each sample, the first figure is the range image, the second one is the pole extraction results by our geometric method, and the last one shows the pole extraction results by our learning-based method. Note that, for the learning-based method we get all the results on three datasets with a single trained model.

finds more poles while introducing more false positives, which decreases precision. This can also be seen in Fig. 6, which shows pole extraction examples of our geometric and learning-based pole extractor.

Note that we trained our pole segmentation network only once with pseudo pole labels generated from different datasets and evaluated it on multiple different datasets. As can be seen in Fig. 6, the environments of different datasets vary a lot, while our learning-based method can still extract poles well without fine-tuning, which shows a good generalization ability of our method. The possible reason for that is the range values of the poles are usually significantly different than the backgrounds, which makes poles distinctive and easy to be detected on range images. Compared to multi-class segmentation, it is easier for the neural network to learn a more general model to detect poles based on the range images [10]. Furthermore, the learning-based method has a higher recall than the geometry-based method, but with a lower precision, which means that the learning-based method detects more true positives, but also more false positives. We use the detected poles as landmarks for MCL, which is a very robust probabilistic localization system. Thus, the localization performance will not be influenced by little false positives, but benefits from higher recalls with more landmarks, as shown in the next section.

### 4.3. Localization Performance

The second experiment is presented to support the claim that our approach achieves higher accuracy on localization in different environments. For all the experiments, we use the same setup as used in the baselines and report their results from the original work.

**Table 1**

Pole Extraction Precision, Recall, and F1 Score on NCLT and KITTI datasets.

Dataset	Method	Precision	Recall	F1 Score
NCLT	Schaefer [30]	0.690	0.386	0.495
	Ours-G	<b>0.765</b>	0.657	<b>0.706</b>
	Ours-L	0.675	<b>0.674</b>	0.674
Semantic KITTI	Schaefer [30]	0.621	0.380	0.455
	Ours-G	<b>0.687</b>	0.439	0.515
	Ours-L	0.607	<b>0.582</b>	<b>0.594</b>

#### 4.3.1. Localization on the NCLT Dataset

The NCLT dataset contains 27 sessions with an average length of 5.5 km and an average duration of 1.3 h over the course of 15 months. The data is recorded at different times over a year, different weather and seasons, including both indoor and outdoor environments, and also lots of dynamic objects. The trajectories of different sessions have a large overlap. Therefore, it is an ideal dataset for testing long-term localization in urban environments.

We first build the map following the setup introduced by Schaefer et al. [30], which uses the laser scans and the ground-truth poses of the first session. Since in later sessions the robot sometimes moves into unseen places for the first session, we therefore also use those scans whose position is 10 m away from all previously visited poses to build the map. During localization, we use 1000 particles and use the same initialization as Schaefer et al. [30] by uniformly sampling positions around the first ground-truth pose within a 2.5 m circle. The orientations are uniformly sampled from  $-5$  to  $5$  degrees. We resample particles if the number of effective particles is less than 50%. To get the


Session Date	$\Delta_{\text{pos}}$ [m]			RMSE <sub>pos</sub> [m]			$\Delta_{\text{ang}}$ [°]			RMSE <sub>ang</sub> [°]		
	Schaefer [30]	Ours-G	Ours-L	Schaefer [30]	Ours-G	Ours-L	Schaefer [30]	Ours-G	Ours-L	Schaefer [30]	Ours-G	Ours-L
2012-01-08	0.130	<b>0.115</b>	0.117	0.178	0.146	<b>0.141</b>	0.663	<b>0.626</b>	0.633	0.857	0.810	<b>0.808</b>
2012-01-15	0.156	<b>0.146</b>	0.147	0.225	<b>0.204</b>	0.205	0.760	<b>0.750</b>	0.753	0.999	<b>0.982</b>	0.989
2012-01-22	0.172	<b>0.149</b>	<b>0.149</b>	0.222	0.190	<b>0.185</b>	0.939	<b>0.911</b>	0.912	1.291	1.238	<b>1.237</b>
2012-02-02	0.155	0.136	<b>0.130</b>	0.205	0.172	<b>0.157</b>	0.720	<b>0.699</b>	0.700	0.975	0.922	<b>0.917</b>
2012-02-04	0.144	0.134	<b>0.130</b>	0.195	0.167	<b>0.158</b>	0.684	0.671	<b>0.670</b>	0.903	0.876	<b>0.866</b>
2012-02-05	0.148	<b>0.138</b>	0.145	0.210	<b>0.206</b>	0.250	<b>0.690</b>	0.694	0.700	0.947	<b>0.937</b>	0.949
2012-02-12	0.269	0.253	<b>0.247</b>	1.005	<b>1.002</b>	1.003	0.802	0.786	<b>0.778</b>	1.040	1.019	<b>0.997</b>
2012-02-18	0.149	0.131	<b>0.129</b>	0.221	0.175	<b>0.161</b>	0.699	<b>0.676</b>	0.681	0.938	<b>0.905</b>	0.909
2012-02-19	0.148	0.137	<b>0.135</b>	0.194	0.183	<b>0.168</b>	0.704	<b>0.692</b>	0.710	0.944	<b>0.923</b>	0.935
2012-03-17	0.149	0.137	<b>0.133</b>	0.191	0.174	<b>0.159</b>	0.830	0.798	<b>0.797</b>	1.062	1.031	<b>1.022</b>
2012-03-25	0.200	0.178	<b>0.170</b>	0.262	0.235	<b>0.221</b>	1.418	1.379	<b>1.365</b>	1.836	1.789	<b>1.767</b>
2012-03-31	0.143	0.135	<b>0.129</b>	0.184	0.176	<b>0.155</b>	0.746	0.729	<b>0.728</b>	0.973	0.936	<b>0.932</b>
2012-04-29	0.170	<b>0.154</b>	0.156	0.251	<b>0.222</b>	0.227	0.829	<b>0.820</b>	0.825	1.079	<b>1.069</b>	1.070
2012-05-11	0.161	0.132	<b>0.131</b>	0.225	0.163	<b>0.155</b>	0.773	<b>0.747</b>	0.766	0.998	<b>0.965</b>	0.987
2012-05-26	0.158	0.142	<b>0.140</b>	0.217	0.183	<b>0.165</b>	0.690	<b>0.672</b>	0.681	0.889	<b>0.871</b>	0.875
2012-06-15	0.180	<b>0.145</b>	0.147	0.238	0.186	<b>0.180</b>	0.659	0.646	<b>0.630</b>	0.879	0.874	<b>0.842</b>
2012-08-04	0.210	0.169	<b>0.159</b>	0.340	0.230	<b>0.192</b>	0.884	<b>0.843</b>	0.847	1.143	1.093	<b>1.091</b>
2012-08-20	0.189	0.156	<b>0.152</b>	0.264	0.207	<b>0.183</b>	0.711	<b>0.688</b>	0.696	0.941	0.906	<b>0.905</b>
2012-09-28	0.206	0.171	<b>0.155</b>	0.311	0.241	<b>0.190</b>	0.731	0.726	<b>0.714</b>	0.952	0.949	<b>0.926</b>
2012-10-28	0.217	0.185	<b>0.168</b>	0.338	0.281	<b>0.230</b>	0.693	0.680	<b>0.678</b>	0.919	0.909	<b>0.900</b>
2012-11-04	0.257	0.208	<b>0.181</b>	0.456	0.317	<b>0.227</b>	0.746	0.718	<b>0.701</b>	0.996	0.973	<b>0.928</b>
2012-11-16	0.403	0.296	<b>0.251</b>	0.722	0.435	<b>0.370</b>	1.467	1.403	<b>1.378</b>	2.031	1.919	<b>1.895</b>
2012-11-17	0.243	0.201	<b>0.172</b>	0.377	0.323	<b>0.219</b>	0.686	0.685	<b>0.677</b>	0.959	0.948	<b>0.914</b>
2012-12-01	0.266	0.226	<b>0.212</b>	0.492	<b>0.429</b>	0.445	0.674	0.665	<b>0.647</b>	0.930	0.887	<b>0.854</b>
2013-01-10	0.217	0.187	<b>0.164</b>	0.278	0.226	<b>0.190</b>	0.689	<b>0.627</b>	0.642	0.911	<b>0.806</b>	0.817
2013-02-23	2.470	0.236	<b>0.207</b>	5.480	0.567	<b>0.492</b>	1.083	<b>0.592</b>	0.593	1.769	0.846	<b>0.845</b>
2013-04-05	0.365	0.295	<b>0.265</b>	0.920	0.869	<b>0.820</b>	0.654	<b>0.642</b>	<b>0.642</b>	1.028	1.036	<b>1.016</b>
Average	0.284	0.174	<b>0.164</b>	0.526	0.293	<b>0.268</b>	0.801	<b>0.761</b>	<b>0.761</b>	1.081	1.016	<b>1.007</b>

**Table 2**

Results of our experiments with the NCLT dataset compared to Schaefer [30], averaged over ten localization runs per session. The variables  $\Delta_{\text{pos}}$  and  $\Delta_{\text{ang}}$  denote the mean absolute errors in position and heading, respectively, RMSE<sub>pos</sub> and RMSE<sub>ang</sub> represent the corresponding root mean squared errors.

pose estimation, we use the average poses of the best 10 % of the particles.

Tab. 2 shows the position and orientation errors for every session. We run the localization 10 times and compute the average means and RMSEs to the ground-truth trajectory. The results show that both our geometric and learning-based methods surpass Schaefer et al. [30] in almost all sessions with an average error of 0.174 m and 0.164 m respectively. Besides, in session 2013-02-23, the baseline method fails to localize resulting in an error of 2.470 m, while our method never loses track of the robot position (Fig. 7). This is because our pole extractor can robustly extract poles even in an environment where there are fewer poles. Schaefer et al. [31] analyze their localization failure in session 2013-02-23 for the reason that the barrels in a construction area are moved a few meters to the right in the later session. As these barrels are detected as poles by their approach, they are built in the map and cause the wrong pole matching to the map in this area. In our pole extraction algorithms, we discard those poles with too large radiuses. Thus, the barrels are not a part of our map and our localization will not be influenced by the movement of these barrels. Interestingly, our learning-based method improves localization results more than our geometry-based method in most sessions. It may be caused by a more general pole segmentation model trained with pseudo labels generated from different environments.



**Figure 7:** Comparison of localization results of Schaefer et al. [30] and our method in session 2013-02-23 on NCLT dataset. The black dots are the poles on the map. The grey line is the ground-truth trajectory. The blue line is our result and the red one is of the baseline method. As can be seen, Schaefer’s method loses track of the robot in some places, while our method always tracks the correct robot poses with respect to the ground truth.

#### 4.3.2. Localization on the MulRan Dataset

To further show the generalization ability of our method, we test both our geometric and learning-based methods on



**Table 3**  
Localization Results on MulRan Dataset.

	Schaefer [30]	Chen [13]	Ours-G	Ours-L
RMSE <sub>pos</sub> [m]	1.82	0.83	<b>0.48</b>	0.49
RMSE <sub>ang</sub> [°]	0.56	3.14	<b>0.27</b>	0.28

**Table 4**  
Pole Extraction Runtime Results.

	NCLT	KITTI	MulRan
Ours-G	12Hz	2Hz	4Hz
Ours-L	17Hz	16Hz	16Hz

the MulRan dataset, which was collected from a different type of LiDAR sensor in a different environment. We use the MulRan dataset KAIST 02 sequence (collected on 2019-08-23) to build the global map and use KAIST 01 sequence (collected on 2019-06-20) for localization. Tab. 3 shows the location and yaw angle RMSE errors on MulRan Dataset. As can be seen, our geometric and learning-based methods consistently achieve a better performance than both baseline methods [30, 13]. Note that, we train our pole segmentation only once, and there is no fine-tuning when applying to a new environment.

#### 4.4. Runtime

This experiment has been conducted to support the claim that our approach runs online at the sensor frame rate. As shown in Tab. 4, we compare our method to the baseline method proposed by Schaefer et al. [30] on three different datasets, including NCLT (session 2012-01-08), KITTI (sequence 09) and MulRan (KAIST 02) datasets. As reported in their paper, on NCLT dataset the baseline method takes an average of 1.33 s for pole extraction on a PC using a GPU. We tested our geometric method without using a GPU and our method only needs 0.09 s for pole extraction and all MCL steps take less than 0.1 s yielding a run time faster than the commonly used LiDAR frame rate of 10 Hz.

The performance of geometry-based pole extractors, both Schaefer's and ours, is influenced by the size of the input data, and it is a trade-off between localization accuracy and speed. To achieve good localization results for the geometric method, we use the range image size of  $32 \times 256$  for NCLT and  $64 \times 500$  for KITTI and MulRan, which leads to a decrease in the runtime performance. However, our learning-based method is not influenced by the size of input data. In our case, we fix the size of network input as  $32 \times 256$ , and our network always works online with good localization performance with a single GPU, which shows a significant advantage of our learning-based method.

## 5. Conclusion

In this paper, we presented a novel range image-based pole extraction approach based on geometric features for online long-term LiDAR localization. Our method exploits

range images generated from LiDAR scans. This allows our method to process point cloud data rapidly and run online. We further use the detected poles by our geometric pole extractor as pseudo labels to train a deep neural network for online pole segmentation. Our learning-based pole extractor can generalize to different types of datasets without fine-tuning, despite the environments of different datasets varying a lot. We implemented and evaluated our approach on multiple different datasets and provided comparisons to other existing techniques and supported all claims made in this paper. The experiments suggest that both our geometric and learning-based methods can accurately extract more poles in the environments and achieve better performance in long-term localization tasks than the baseline methods. Moreover, we release our implementation and pole dataset for other researchers to evaluate their algorithms. In the future, we plan to explore the usage of other features such as road markings, curb, and intersection features, to improve the robustness of our method.

## References

- [1] Barfoot, T.D., 2017. State Estimation for Robotics. 1st ed., Cambridge University Press, USA.
- [2] Behley, J., Garbade, M., Milioto, A., Quenzel, J., Behnke, S., Stachniss, C., Gall, J., 2019. SemanticKITTI: A Dataset for Semantic Scene Understanding of LiDAR Sequences, in: Proc. of the IEEE/CVF Intl. Conf. on Computer Vision (ICCV).
- [3] Behley, J., Stachniss, C., 2018. Efficient Surfel-Based SLAM using 3D Laser Range Data in Urban Environments, in: Proc. of Robotics: Science and Systems (RSS).
- [4] Bennewitz, M., Stachniss, C., Burgard, W., Behnke, S., 2006. Metric Localization with Scale-Invariant Visual Features using a Single Perspective Camera, in: Christensen, H. (Ed.), European Robotics Symposium 2006, Springer Verlag. pp. 143–157.
- [5] Carlevaris-Bianco, N., Ushani, A., Eustice, R., 2016. University of Michigan North Campus long-term vision and lidar dataset. Intl. Journal of Robotics Research (IJRR) 35, 1023–1035. doi:10.1177/0278364915614638.
- [6] Chen, G., Lu, F., Li, Z., Liu, Y., Dong, J., Zhao, J., Yu, J., Knoll, A., 2021a. Pole-curb fusion based robust and efficient autonomous vehicle localization system with branch-and-bound global optimization and local grid map method. IEEE Transactions on Vehicular Technology 70, 11283–11294. doi:10.1109/TVT.2021.3114825.
- [7] Chen, X., Läbe, T., Milioto, A., Röhling, T., Behley, J., Stachniss, C., 2021b. OverlapNet: A Siamese Network for Computing LiDAR Scan Similarity with Applications to Loop Closing and Localization. Autonomous Robots doi:10.1007/s10514-021-09999-0.
- [8] Chen, X., Läbe, T., Milioto, A., Röhling, T., Vysotska, O., Haag, A., Behley, J., Stachniss, C., 2020a. OverlapNet: Loop Closing for LiDAR-based SLAM, in: Proc. of Robotics: Science and Systems (RSS).
- [9] Chen, X., Läbe, T., Nardi, L., Behley, J., Stachniss, C., 2020b. Learning an Overlap-based Observation Model for 3D LiDAR Localization, in: Proc. of the IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS).
- [10] Chen, X., Li, S., Mersch, B., Wiesmann, L., Gall, J., Behley, J., Stachniss, C., 2021c. Moving Object Segmentation in 3D LiDAR Data: A Learning-based Approach Exploiting Sequential Data. IEEE Robotics and Automation Letters (RA-L) 6, 6529–6536. doi:10.1109/LRA.2021.3093567.
- [11] Chen, X., Mersch, B., Nunes, L., Marcuzzi, R., Vizzo, I., Behley, J., Stachniss, C., 2022. Automatic Labeling to Generate Training Data for Online LiDAR-Based Moving Object Segmentation. IEEE

- Robotics and Automation Letters (RA-L) 7, 6107–6114. doi:10.1109/LRA.2022.3166544.
- [12] Chen, X., Milioto, A., Palazzolo, E., Giguère, P., Behley, J., Stachniss, C., 2019. SuMa++: Efficient LiDAR-based Semantic SLAM, in: Proc. of the IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS).
- [13] Chen, X., Vizzo, I., Läbe, T., Behley, J., Stachniss, C., 2021d. Range Image-based LiDAR Localization for Autonomous Vehicles, in: Proc. of the IEEE Intl. Conf. on Robotics & Automation (ICRA).
- [14] Cortinhal, T., Tzelepis, G., Aksoy, E.E., 2020. SalsaNext: Fast, Uncertainty-Aware Semantic Segmentation of LiDAR Point Clouds, in: Proc. of the IEEE Vehicles Symposium (IV).
- [15] Cvisic, I., Cestic, J., Markovic, I., Petrovic, I., 2017. SOFT-SLAM: Computationally Efficient Stereo Visual SLAM for Autonomous UAVs. Journal of Field Robotics (JFR) 35, 578–595.
- [16] Dellaert, F., Fox, D., Burgard, W., Thrun, S., 1999. Monte carlo localization for mobile robots, in: IEEE International Conference on Robotics and Automation (ICRA).
- [17] Dong, H., Chen, X., Stachniss, C., 2021. Online Range Image-based Pole Extractor for Long-term LiDAR Localization in Urban Environments, in: Proc. of the Europ. Conf. on Mobile Robotics (ECMR).
- [18] Droschel, D., Behnke, S., 2018. Efficient continuous-time slam for 3d lidar-based online mapping, in: Proc. of the IEEE Intl. Conf. on Robotics & Automation (ICRA).
- [19] Grisetti, G., Stachniss, C., Burgard, W., 2007. Improved Techniques for Grid Mapping with Rao-Blackwellized Particle Filters. IEEE Trans. on Robotics (TRO) 23, 34–46.
- [20] Kim, G., Park, Y., Cho, Y., Jeong, J., Kim, A., 2020. Mulran: Multimodal range dataset for urban place recognition, in: Proc. of the IEEE Intl. Conf. on Robotics & Automation (ICRA).
- [21] Kümmerle, J., Sons, M., Poggenhans, F., Kühner, T., Lauer, M., Stiller, C., 2019. Accurate and efficient self-localization on roads using basic geometric primitives, in: Proc. of the IEEE Intl. Conf. on Robotics & Automation (ICRA).
- [22] Kümmerle, R., Ruhnke, M., Steder, B., Stachniss, C., Burgard, W., 2014. Autonomous robot navigation in highly populated pedestrian zones. Journal of Field Robotics (JFR) doi:10.1002/rob.21534.
- [23] Li, S., Chen, X., Liu, Y., Dai, D., Stachniss, C., Gall, J., 2022. Multi-scale Interaction for Real-time LiDAR Data Segmentation on an Embedded Platform. IEEE Robotics and Automation Letters (RA-L) 7, 738–745. doi:10.1109/LRA.2021.3132059.
- [24] Ma, W., Tartavull, I., Bársan, I.A., Wang, S., Bai, M., Mattyus, G., Homayounfar, N., Lakshminanth, S.K., Pokrovsky, A., Urtasun, R., 2019. Exploiting Sparse Semantic HD Maps for Self-Driving Vehicle Localization, in: Proc. of the IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS).
- [25] Maddern, W., Pascoe, G., Linegar, C., Newman, P., 2017. 1 year, 1000 km: The oxford robotcar dataset. The International Journal of Robotics Research 36, 3–15. doi:10.1177/0278364916679498.
- [26] Milioto, A., Stachniss, C., 2019. Bonnet: An Open-Source Training and Deployment Framework for Semantic Segmentation in Robotics using CNNs, in: Proc. of the IEEE Intl. Conf. on Robotics & Automation (ICRA).
- [27] Milioto, A., Vizzo, I., Behley, J., Stachniss, C., 2019. RangeNet++: Fast and Accurate LiDAR Semantic Segmentation, in: Proceedings of the IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS).
- [28] Mur-Artal, R., Montiel, J., Tardos, J.D., 2015. ORB-SLAM: a versatile and accurate monocular SLAM system. IEEE Trans. on Robotics (TRO) 31, 1147–1163.
- [29] Plachetka, C., Fricke, J., Klingner, M., Fingscheidt, T., 2021. Dnn-based recognition of pole-like objects in lidar point clouds, in: 2021 IEEE International Intelligent Transportation Systems Conference (ITSC), pp. 2889–2896. doi:10.1109/ITSC48978.2021.9564759.
- [30] Schaefer, A., Büscher, D., Vertens, J., Luft, L., Burgard, W., 2019. Long-term urban vehicle localization using pole landmarks extracted from 3-D lidar scans, in: Proc. of the Europ. Conf. on Mobile Robotics (ECMR), pp. 1–7.
- [31] Schaefer, A., Büscher, D., Vertens, J., Luft, L., Burgard, W., 2021. Long-term vehicle localization in urban environments based on pole landmarks extracted from 3-d lidar scans. Robotics and Autonomous Systems 136. doi:https://doi.org/10.1016/j.robot.2020.103709.
- [32] Sefati, M., Daum, M., Sondermann, B., Kreisköther, K.D., Kampker, A., 2017. Improving vehicle localization using semantic and pole-like landmarks, in: Proc. of the IEEE Vehicles Symposium (IV).
- [33] Shi, C., Chen, X., Huang, K., Xiao, J., Lu, H., Stachniss, C., 2021. Keypoint Matching for Point Cloud Registration using Multiplex Dynamic Graph Attention Networks. IEEE Robotics and Automation Letters (RA-L) 6, 8221–8228. doi:10.1109/LRA.2021.3097275.
- [34] Shi, Z., Kang, Z., Lin, Y., Liu, Y., Chen, W., 2018. Automatic recognition of pole-like objects from mobile laser scanning point clouds. Remote Sensing 10. doi:10.3390/rs10121891.
- [35] Spangenberg, R., Goehring, D., Rojas, R., 2016. Pole-Based Localization for Autonomous Vehicles in Urban Scenarios, in: Proc. of the IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS).
- [36] Sun, L., Adolfsson, D., Magnusson, M., Andreasson, H., Posner, I., Duckett, T., 2020. Localising Faster: Efficient and precise lidar-based robot localisation in large-scale environments, in: Proc. of the IEEE Intl. Conf. on Robotics & Automation (ICRA).
- [37] Thrun, S., Burgard, W., Fox, D., 2005. Probabilistic Robotics. MIT Press.
- [38] Tinchev, G., Penate-Sanchez, A., Fallon, M., 2019. Learning to see the wood for the trees: Deep laser localization in urban and natural environments on a CPU. IEEE Robotics and Automation Letters (RA-L) 4, 1327–1334.
- [39] Toft, C., Maddern, W., Torii, A., Hammarstrand, L., Stenborg, E., Safari, D., Okutomi, M., Pollefeys, M., Sivic, J., Pajdla, T., Kahl, F., Sattler, T., 2020. Long-term visual localization revisited. IEEE Trans. on Pattern Analysis and Machine Intelligence (TPAMI) doi:10.1109/TPAMI.2020.3032010.
- [40] Trahanias, P., Burgard, W., Argyros, A., Hähnel, D., Baltzakis, H., Pfaff, P., Stachniss, C., 2005. TOURBOT and WebFAIR: Web-Operated Mobile Robots for Tele-Presence in Populated Exhibitions. IEEE Robotics and Automation Magazine (RAM) 12, 77–89.
- [41] Vizzo, I., Chen, X., Chebroly, N., Behley, J., Stachniss, C., 2021. Poisson Surface Reconstruction for LiDAR Odometry and Mapping, in: Proc. of the IEEE Intl. Conf. on Robotics & Automation (ICRA).
- [42] Weng, L., Yang, M., Guo, L., Wang, B., Wang, C., 2018. Pole-based real-time localization for autonomous driving in congested urban scenarios, in: Proc. of the Intl. Conf. on Real-time Computing and Robotics (RCAR).
- [43] Weng, S., Li, J., Chen, Y., Wang, C., 2016. Road traffic sign detection and classification from mobile LiDAR point clouds, in: Proc. of the Intl. Conf. on Computer Vision in Remote Sensing.
- [44] Wiesmann, L., Milioto, A., Chen, X., Stachniss, C., Behley, J., 2021. Deep Compression for Dense Point Cloud Maps. IEEE Robotics and Automation Letters (RA-L) 6, 2060–2067. doi:10.1109/LRA.2021.3059633.
- [45] Wilbers, D., 2021. Graph-based Sliding Window Localization and Map Refinement for Automated Vehicles. Ph.D. thesis. Rheinische Friedrich-Wilhelms-Universität Bonn.
- [46] Zhou, Q., Park, J., Koltun, V., 2018. Open3D: A modern library for 3D data processing. arXiv:1801.09847.
- [47] Zimmerman, N., Wiesmann, L., Guadagnino, T., Läbe, T., Behley, J., Stachniss, C., 2022. Robust Onboard Localization in Changing Environments Exploiting Text Spotting, in: Proc. of the IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS).