

Learning an Overlap-based Observation Model for 3D LiDAR Localization

Xieyuanli Chen

Thomas Labe

Lorenzo Nardi

Jens Behley

Cyrill Stachniss

Abstract—Localization is a crucial capability for mobile robots and autonomous cars. In this paper, we address learning an observation model for Monte-Carlo localization using 3D LiDAR data. We propose a novel, neural network-based observation model that computes the expected overlap of two 3D LiDAR scans. The model predicts the overlap and yaw angle offset between the current sensor reading and virtual frames generated from a pre-built map. We integrate this observation model into a Monte-Carlo localization framework and tested it on urban datasets collected with a car in different seasons. The experiments presented in this paper illustrate that our method can reliably localize a vehicle in typical urban environments. We furthermore provide comparisons to a beam-endpoint and a histogram-based method indicating a superior global localization performance of our method with fewer particles.

I. INTRODUCTION

All mobile systems that navigate autonomously in a goal-directed manner need to know their position and orientation in the environment, typically with respect to a map. This task is often referred to as localization and can be challenging especially at the city-scale and in the presence of a lot of dynamic objects, e.g., vehicles and humans. Over the last decades, a wide range of localization systems have been developed relying on different sensors. Frequently used sensors are GPS receivers, inertial measurement units, cameras, and laser range scanners. All sensing modalities have their advantages and disadvantages. For example, GPS does not work indoors, cameras do not work well at night or under strong appearance changes, and LiDARs are active sensors and are still rather expensive.

Most autonomous robots as well as cars have a 3D LiDAR sensor onboard to perceive the scene and directly provide 3D data. In this paper, we consider the problem of vehicle localization *only* based on a 3D LiDAR sensor. For localization, probabilistic state estimation techniques such as extended Kalman filters (EKF) [8], particle filters [12], or incremental factor graph optimization approaches [24] can be found in most localization systems today. Whereas EKFs and most optimization-based approaches track only a single mode, particle filters inherently support multi-modal distributions. Furthermore, PFs do not restrict the motion or observation model to follow a specific distribution such as a Gaussian. The motion model can often be defined

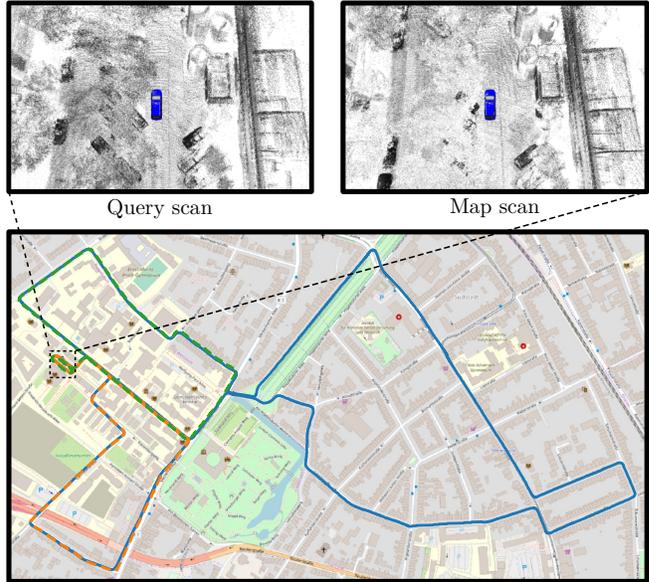


Fig. 1. The image in the lower part shows the trajectories of the dataset used in this paper, overlaid on OpenStreetMap. The blue trajectory represents the sequence used to generate a map for localization. The green and the orange trajectories represent two different test sequences. The 3D point clouds shown in the upper part show LiDAR scans of the same place, once during mapping and once during localization. Since the LiDAR data was collected in different seasons, the appearance of the environment changed quite significantly due to changes in the vegetation but also due to parked vehicles at different places.

quite easily based on the physics of vehicle motion. The observation model, however, is trickier to define. Often, these hand-designed models are used and they strongly impact the performance of the resulting localization system. Frequently used observation models for LiDARs are the beam-end point model, also called the likelihood field [27], the ray-casting model [9], or models based on handcrafted features [26], [35]. Recently, researchers also focused on learning such models completely from data [15], [30].

In this paper, we address the problem of learning observation models for 3D LiDARs and propose a deep neural network-based approach for that. We explore the possibilities of learning an observation model based on OverlapNet [3] that predicts the overlap between two LiDAR scans, see Fig. 1, where the overlap is defined as the ratio of points that can be seen from both LiDAR scans. In this work, we investigate this concept for 3D LiDAR-based observation models.

All authors are with the University of Bonn, Germany.

This work has been funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy, EXC-2070 - 390732324 - PhenoRob as well as under grant number BE 5996/1-1 and by the Chinese Scholarship Committee.

The main contribution of this paper is a novel observation model for 3D LiDAR-based localization. Our model is learned from range data using a deep neural network. It estimates the overlap and yaw angle offset between a query frame and map data. We use this information as the observation model in Monte-Carlo localization (MCL) for updating the importance weights of the particles. Based on our novel observation model, our approach achieves online localization using 3D LiDAR scans over extended periods of time with a comparably small number of particles.

The source code of our approach is available at:

https://github.com/PRBonn/overlap_localization

II. RELATED WORK

Localization is a classical topic in robotics [27]. For localization given a map, one often distinguishes between pose tracking and global localization. In pose tracking, the vehicle starts from a known pose and the pose is updated over time. In global localization, no pose prior is available. In this work, we address global localization using 3D laser scanners without assuming any pose prior from GPS or other sensors. Thus, we focus mainly on LiDAR-based approaches in this section.

Traditional approaches to robot localization rely on probabilistic state estimation techniques. These approaches can still be found today in several localization systems [19], [6]. A popular framework is Monte-Carlo localization [9], [28], [11], which uses a particle filter to estimate the pose of the robot.

In the context of autonomous cars, there are many approaches building and using high-definition (HD) maps for localization, i.e., tackling the simultaneous localization and mapping problem [25] and additionally adding relevant information for the driving domain. Levinson *et al.* [17] utilize GPS, IMU, and LiDAR data to build HD maps for localization. They generate a 2D surface image of ground reflectivity in the infrared spectrum and define an observation model that uses these intensities. The uncertainty in intensity values has been handled by building a prior map [18], [32]. Barsan *et al.* [15] use a fully convolutional neural network (CNN) to perform online-to-map matching for improving the robustness to dynamic objects and eliminating the need for LiDAR intensity calibration. Their approach shows a strong performance but requires a good GPS prior for operation. Based on this approach, Wei *et al.* [30] proposed a learning-based compression method for HD maps. Merfels and Stachniss [20] present an efficient chain graph-like pose-graph for vehicle localization exploiting graph optimization techniques and different sensing modalities. Based on this work, Wilbers *et al.* [31] propose a LiDAR-based localization system performing a combination of local data association between laser scans and HD map features, temporal data association smoothing, and a map matching approach for robustification.

Other approaches aim at performing LiDAR-based place recognition to initialize localization. For example, Kim *et al.* [16] transform point clouds into scan context images

and train a CNN based on such images. They generate scan context images for both the current frame and all grid cells of the map and compare them to estimate the current location as the cell presenting the largest score. Yin *et al.* [34] propose a Siamese network to first generate fingerprints for LiDAR-based place recognition and then use iterative closest points to estimate the metric poses. Cop *et al.* [7] propose a descriptor for LiDAR scans based on intensity information. Using this descriptor, they first perform place recognition to find a coarse location of the robot, eliminate inconsistent matches using RANSAC, and then refine the estimated transformation using iterative closest points. In contrast to approaches that perform place recognition first, our approach integrates a neural network-based observation model into an MCL framework to estimate the robot pose.

Recently, several approaches exploiting semantic information for 3D LiDAR localization have been proposed. In our earlier work [5], we used a camera and a LiDAR to detect victims and localize the robot in an urban search and rescue environment. Ma *et al.* [19] combine semantic information such as lanes and traffic signs in a Bayesian filtering framework to achieve accurate and robust localization within sparse HD maps. Yan *et al.* [33] exploit buildings and intersections information from a LiDAR-based semantic segmentation system [21] to localize on OpenStreetMap. Schaefer *et al.* [23] detect and extract pole landmarks from 3D LiDAR scans for long-term urban vehicle localization. Tinchev *et al.* [29] propose a learning-based method to match segments of trees and localize in both urban and natural environments. Dubé *et al.* [10] propose to perform localization by extracting segments from 3D point clouds and matching them to accumulated data. Whereas, Zhang *et al.* [35] utilize both ground reflectivity features and vertical features for localizing autonomous car in rainy conditions. In our previous work [4], we also exploit semantic information [21], [1] to improve the localization and mapping results by detecting and removing dynamic objects.

Different to the above discussed methods [15], [19], [30], which use GPS as prior for localization, our method only exploits LiDAR information to achieve global localization without using any GPS information. Moreover, our approach uses range scans without explicitly exploiting semantics or extracting landmarks. Instead, we rely on CNNs to predict the overlap between range scans and their yaw angle offset and use this information as an observation model for Monte-Carlo localization. The localization approach proposed in this paper is based on our previous work called OverlapNet [3], which focuses on loop-closing for 3D LiDAR-based SLAM.

III. OUR APPROACH

The key idea of our approach is to exploit the neural network, OverlapNet [3] that can estimate the overlap and yaw angle offset of two scans for building an observation model for localization in a given map. To this end, we first generate virtual scans at 2D locations on a grid rendered from the aggregated point cloud of the map. We train the network completely self-supervised on the map used for

localization (see Sec. III-A). We compute features using our pre-trained network that allows us to compute the overlap and yaw angles between a query and virtual scans (see Sec. III-B). Finally, we integrate an observation model using the overlap (see Sec. III-D) and a separate observation model for the yaw angle estimates (see Sec. III-E) in a particle filter to perform localization (see Sec. III-C).

A. *OverlapNet*

We proposed the so-called *OverlapNet* to detect loop closures candidates for a 3D LiDAR-based SLAM. The idea of overlap has its origin in the photogrammetry and computer vision community [14]. The intuition is that to successfully match two images and calculate their relative pose, the images must overlap. This can be quantified by defining the overlap percentage as the percentage of pixels in the first image, which can successfully be projected back into the second image.

In *OverlapNet*, we use the idea of overlap for range images generated from 3D LiDAR scans exploiting the range information explicitly. First, we generate from the LiDAR scans a range-image like input tensor $\mathbf{I} \in \mathbb{R}^{H \times W \times 4}$, where each pixel (i, j) in the range-image like representation corresponds to the depth and the corresponding normal, i.e., $\mathbf{I}(i, j) = (\mathbf{r}, \mathbf{n}_x, \mathbf{n}_y, \mathbf{n}_z)$, where $\mathbf{r} = \|\mathbf{p}\|_2$. The indices for inserting points and estimated normals are computed using a spherical projection that maps points $\mathbf{p} \in \mathbb{R}^3$ to two-dimensional coordinates (i, j) , also see [3], [4].

OverlapNet uses a siamese network structure and takes the tensors \mathbf{I}_1 and \mathbf{I}_2 of two scans as input for the legs to generate feature volumes, \mathbf{F}_1 and \mathbf{F}_2 . These two feature volumes are then used as inputs for two heads. One head $\mathcal{H}_{\text{overlap}}(\mathbf{F}_1, \mathbf{F}_2)$ estimates the overlap percentage of the scans and the other head $\mathcal{H}_{\text{yaw}}(\mathbf{F}_1, \mathbf{F}_2)$ estimates the relative yaw angle.

For training, we determine ground truth overlap and yaw angle values using known poses estimated by a SLAM approach [2]. Given these poses, we can train *OverlapNet* completely self-supervised. For more details on the network structure and the training procedure, we are referring to our prior work [3].

B. *Map of Virtual Scans*

OverlapNet requires two LiDAR scans as input. One is the current scan and the second one has to be generated from the map. Thus, we build a map of virtual LiDAR scans given an aggregated point cloud by using a grid of locations with grid resolution γ , where we generate virtual LiDAR scans for each location. The grid resolution is a trade-off between the accuracy and storage size of the map. Instead of storing these virtual scans, we just need to use one leg of the *OverlapNet* to obtain a feature volume \mathbf{F} using the input tensor \mathbf{I} of this virtual scan. Storing the feature volume instead of the complete scan has two key advantages: (1) it uses more than an order less space than the original point cloud (ours: 0.1Gb/km, raw scans: 1.7Gb/km) and (2) we do not need to compute the \mathbf{F} during localization on the map. The features volumes of the virtual scans are then used

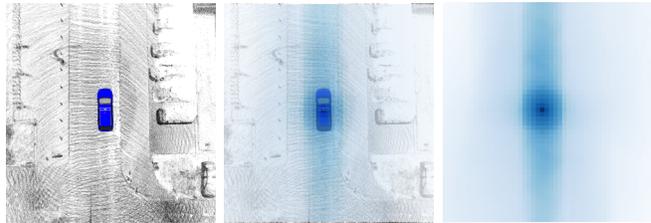


Fig. 2. Overlap observation model. Local heatmap of the scan at the car’s position with respect to the map. Darker blue shades correspond to higher probabilities.

to compute overlap and yaw angle estimates with a query scan that is the currently observed LiDAR point cloud in our localization framework.

C. *Monte-Carlo Localization*

Monte-Carlo localization (MCL) is a localization algorithm based on the particle filter proposed by Dellaert *et al.* [9]. Each particle represents a hypothesis for the robot’s or autonomous vehicle’s 2D pose $\mathbf{x}_t = (x, y, \theta)_t$ at time t . When the robot moves, the pose of each particle is updated with a prediction based on a motion model with the control input \mathbf{u}_t . The expected observation from the predicted pose of each particle is then compared to the actual observation \mathbf{z}_t acquired by the robot to update the particle’s weight based on an observation model. Particles are resampled according to their weight distribution and resampling is triggered whenever the effective number of particles, see for example [13], drops below 50% of the sample size. After several iterations of this procedure, the particles will converge around the true pose.

MCL realizes a recursive Bayesian filtering scheme. The key idea of this approach is to maintain a probability density $p(\mathbf{x}_t | \mathbf{z}_{1:t}, \mathbf{u}_{1:t})$ of the pose \mathbf{x}_t at time t given all observations $\mathbf{z}_{1:t}$ up to time t and motion control inputs $\mathbf{u}_{1:t}$ up to time t . This posterior is updated as follows:

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

where η is a normalization constant, $p(\mathbf{x}_t | \mathbf{u}_t, \mathbf{x}_{t-1})$ is the motion model, and $p(\mathbf{z}_t | \mathbf{x}_t)$ is the observation model.

This paper focuses on the observation model. For the motion model, we follow a standard odometry model for vehicles [27].

We split the observation model into two parts:

$$p(\mathbf{z}_t | \mathbf{x}_t) = p_L(\mathbf{z}_t | \mathbf{x}_t) p_O(\mathbf{z}_t | \mathbf{x}_t), \quad (2)$$

where \mathbf{z}_t is the LiDAR observation at time t , $p_L(\mathbf{z}_t | \mathbf{x}_t)$ is the probability encoding the location (x, y) agreement between the current query LiDAR scan and the virtual scan at the grid where the particle locate and $p_O(\mathbf{z}_t | \mathbf{x}_t)$ is the probability encoding the yaw angle θ agreement between the same pairs of scans.

D. Overlap Observation Model

Given a particle i with the state vector (x_i, y_i, θ_i) , the overlap estimates encode the location agreement between the query LiDAR scan and virtual scans of the grid cells where particles locate. It can be directly used as the probability:

$$p_L(\mathbf{z}_t | \mathbf{x}_t) \propto f(\mathbf{z}_t, \mathbf{z}_i; \mathbf{w}), \quad (3)$$

where f corresponds to the neural network providing the overlap estimation between the input scans $\mathbf{z}_t, \mathbf{z}_i$ and \mathbf{w} is the pre-trained weights of the network. \mathbf{z}_t and \mathbf{z}_i are the current query scan and a virtual scan of one (x, y) location respectively. Note that no additional hyperparameter is needed to formulate our observation model for localization.

For illustration purposes, Fig. 2 shows the probabilities of all grid cells in a local area calculated by the overlap observation model. The blue car in the figure shows the current location of the car. The probabilities calculated by the overlap observation model can represent well the hypotheses of the current location of the car.

Typically, a large number of particles are used, especially when the environment is large. However, a large amount of particles will increase the computation time linearly. When applying the overlap observation model, particles could still obtain relatively large weights as long as they are close to the actual pose, even if not in the exact same position. This allows us to use fewer particles to achieve a high success rate of global localization.

Furthermore, the overlap estimation only encodes the location hypotheses. Therefore, if multiple particles locate in the same grid area, only a single inference against the nearest virtual scan of the map needs to be done, which can further reduce the computation time.

E. Yaw Angle Observation Model

Given a particle i with the state vector (x_i, y_i, θ_i) , the yaw angle estimates encode the orientation agreement between the query LiDAR scan and virtual scans of the corresponding grids where particles locate. We formulate the orientation probability as follows:

$$p_O(\mathbf{z}_t | \mathbf{x}_t) \propto \exp\left(-\frac{1}{2} \frac{(g(\mathbf{z}_t, \mathbf{z}_i; \mathbf{w}) - \theta_i)^2}{\sigma_\theta^2}\right), \quad (4)$$

where g corresponds to the neural network providing the yaw angle estimation between the input scans $\mathbf{z}_t, \mathbf{z}_i$ and \mathbf{w} is the pre-trained weights of the network. \mathbf{z}_t and \mathbf{z}_i are the current query scan and a virtual scan of one particle respectively.

When generating the virtual scans of the grid map, all virtual scans will be set facing the absolute 0° yaw angle direction. By doing this, the estimated relative yaw angle between the query scan and the virtual scan indicates the absolute yaw angle of the current query scan. Eq. (4) assumes a Gaussian measurement error in the heading.

By combining overlap and yaw angle estimation, the proposed observation model will correct the weights of particles considering agreements between the query scan and the map with the full pose (x, y, θ) .



Fig. 3. Sensor setup used for data recording: Ouster OS1-64 LiDAR sensor plus GNSS information from a Emilid Reach RS2.

IV. EXPERIMENTAL EVALUATION

In this paper, we use a grid representation and generate virtual 360° scans for each grid point. The resolution of the grid is $\gamma = 20$ cm. When generating the virtual scans, we set the yaw angle to 0° . Therefore, when estimating the relative yaw angle between the query frame and the virtual frames, the result will indicate the absolute yaw of the query frame. For the yaw angle observation model in Eq. (4), we set $\sigma_\theta = 5^\circ$. To achieve global localization, we train a new model only based on the map scans and the generated virtual scans.

The main focus of this work is a new observation model for LiDAR-based localization. Therefore, when comparing different methods, we only change the observation model $p(\mathbf{z}_t | \mathbf{x}_t)$ of the MCL framework and keep the particle filter-based localization the same. The motion model is the typical odometry model [27].

A. Car Dataset

The dataset used in this paper was collected using a self-developed sensor platform illustrated in Fig. 3. To test LiDAR-based global localization, a large-scale dataset has been collected in different seasons with multiple sequences repeatedly exploring the same crowded urban area. For our car dataset, we performed a 3D LiDAR SLAM [2] combined with GPS information to create near ground truth poses. During localization, we only use LiDAR scans for global localization without using GPS.

The dataset has three sequences that were collected at different times of the year, sequence 00 in September 2019, sequence 01 in November 2019, and sequence 02 in February 2020. The whole dataset covers a distance of over 10 km. We use LiDAR scans from sequence 02 to build the virtual scans and use sequence 00 and 01 for localization. As can be seen from Fig. 1, the appearance of the environment changes quite significantly since the dataset was collected in different seasons and in crowded urban environments, including changes in vegetation, but also cars at different locations and moving people.

B. Different Observation Models

In the following experiments, we use the same MCL framework and only exchange the observation models. We compare our observation model with two baseline observation models: the typical beam-end model [27] and a

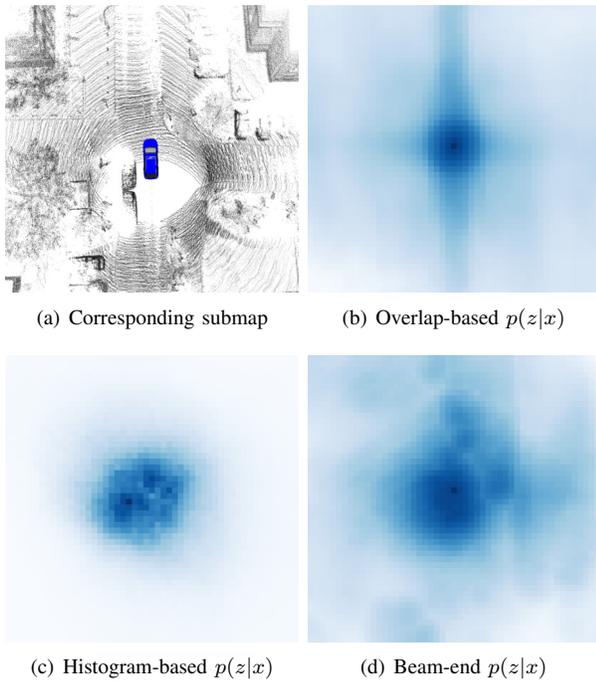


Fig. 4. Heatmaps of different observation models used in the experiments generated for the same query scan and map.

histogram-based model derived from the work of Röhling *et al.* [22].

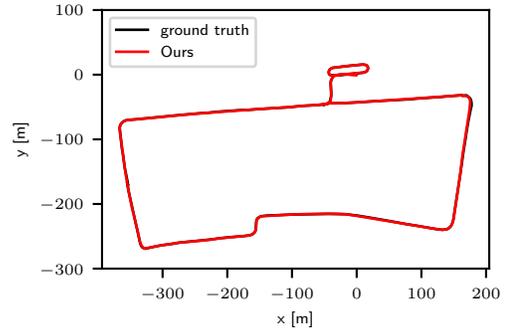
The beam-end observation model is often used for 2D LiDAR data. For 3D LiDAR scans, it needs much more particles to make sure that it converges to the correct pose, which causes the computation time to increase substantially. In this paper, we implement the beam-end model with a down-sampled point cloud map using voxelization with a resolution of 10 cm.

Our second baseline for comparison is inspired by Röhling *et al.* [22], which proposed a fast method to detect loop closures through the use of similarity measures on histograms extracted from 3D LiDAR data. The histogram contains the range information. We use a similar idea, but integrate it into the MCL framework as a baseline observation model. We employ the same grid map and virtual frames as used for our method with the histogram-based observation model. When updating the weights of particles, we will first generate the histograms of the current query scan and the virtual scans of grids where the particles locate. Then, we use the same Wasserstein distance to measure the similarity between them and update the weights of particles as follows:

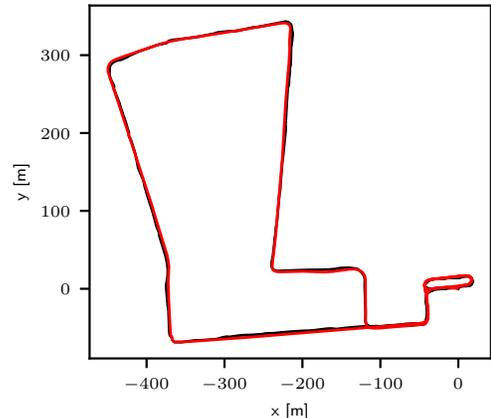
$$p(\mathbf{z}_t|\mathbf{x}_t) \propto d(h(\mathbf{z}_t), h(\mathbf{z}_i)), \quad (5)$$

where d represents the Wasserstein distance between histograms $h(\mathbf{z}_t), h(\mathbf{z}_i)$ of LiDAR scan $\mathbf{z}_t, \mathbf{z}_i$.

Fig. 4 shows the 2D heatmaps in x and y calculated for the different observation models. As can be seen, the proposed observation model tends to give higher weights to the positions along the road, which leads to a higher success rate when the vehicle aims to localize in an urban



(a) Trajectory from sequence 00



(b) Trajectory from sequence 01

Fig. 5. Localization results of our method with 10,000 particles on two sequences recorded with the setup depicted in Fig. 3. Shown are the ground truth trajectory (black) and our estimated trajectory using our observation model (red).

environment. We will show the numerical results which verify that our method can achieve a high success rate with much fewer particles in the next sections.

C. Localization Performance

The experiment presented in this section is designed to show the performance of our approach and to support the claim that it is well suited for global localization.

First of all, we show the general localization results tested with two sequences in Fig. 5. The qualitative results show that, after applying our sensor-model, the proposed method can well localize in the map with only LiDAR data collected in highly dynamic environments at different times.

For quantitative results, we calculate the success rate for different methods with different particle numbers comparing our approach to different methods, as shown in Fig. 6. The x-axis represents the number of particles used during localization, while the y-axis is the success rate of different setups. The success rate for a specific setup of one method is calculated using the number of success cases divided by the total numbers of the tests. To decide whether one test is successful or not, we check the location error by every 100 frames after converging. If the location error is smaller than a certain threshold, we count this run as a success case.

TABLE I
LOCATION RESULTS

Sequence	Location error [meter]		
	Beam-end	Histogram-based	Ours
0	0.92 ± 0.27	1.85 ± 0.34	0.81 ± 0.13
1	0.67 ± 0.11	1.86 ± 0.34	0.88 ± 0.07

TABLE II
YAW ANGLE RESULTS

Sequence	Yaw angle error [degree]		
	Beam-end	Histogram-based	Ours
0	1.87 ± 0.47	3.10 ± 3.07	1.74 ± 0.11
1	2.10 ± 0.59	3.11 ± 3.08	1.88 ± 0.09

We test our method together with two baselines using five different numbers of particles $N = \{1,000; 5,000; 10,000; 50,000; 100,000\}$. For each setup, we sample 10 trajectories and perform global localization.

Quantitative results of localization accuracy are shown in Tab. I and Tab. II. Tab. I shows the location error of all methods tested with both sequences. The location error is defined as the root mean square error (RMSE) of each test in terms of (x, y) Euclidean error with respect to the ground truth poses. Tab. I shows the mean and the standard deviation of the error for each observation model. Note that the location error is only calculated for success cases.

Tab. II shows the yaw angle error. It is the RMSE of each test in terms of yaw angle error with respect to the ground truth poses. The table shows the mean and the standard deviation of the error for each observation model. As before, the yaw angle error is also only calculated for cases in which the global localization converged.

As can be seen from the results, our method achieves higher success rates with a smaller number of particles compared to the baseline methods, which also makes the proposed method faster than baseline methods. Furthermore, our method converges already with 100,000 particles in all cases, whereas the other observation models still need more particles to sufficiently cover the state space. Moreover, the proposed method gets similar performance in location error comparing to the baseline methods but it achieves better results in yaw angle estimation. This is because the proposed method decouples the location and yaw angle estimation and, therefore, can exploit more constraint in yaw angle corrections.

To sum up, the proposed method outperforms the baselines method in terms of success rate, while getting similar results in terms of location error. Moreover, our method outperforms baseline methods in yaw angle estimation, because of the proposed de-coupled observation model. Furthermore, our method is faster than the baseline method. The runtime details will be shown in the next experiment.

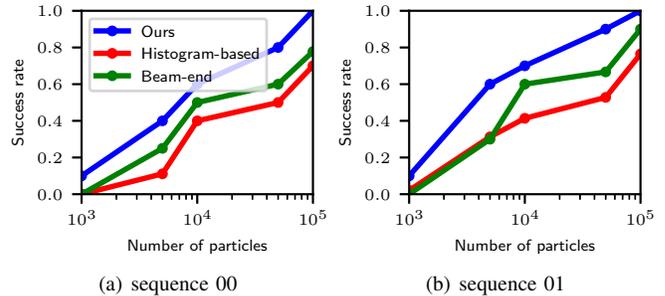


Fig. 6. Success rate of the different observation models for 10 runs with $N = \{1,000; 5,000; 10,000; 50,000; 100,000\}$ particles. Here, we use sequence 00 and sequence 01 to localize in sequence 02. We count runs as success if they converge to the ground truth location within 5 m.

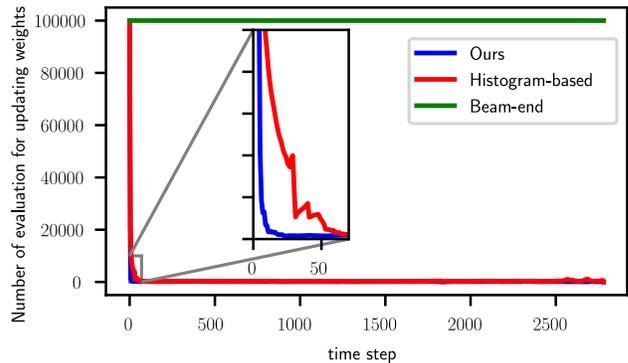


Fig. 7. Number of observation model evaluations for updating the weights at each timestep with 100,000 particles. The beam end model needs to be evaluated for each and every particle individually. The histogram-based method is more computationally efficient, while our proposed method still needs the fewest evaluations.

D. Runtime

In this experiment, we show the *number* of observation model evaluations necessary for updating the weights at each time step in Fig. 7. This is a fairer way to compare the computational cost of different methods, since our neural network-based method uses GPU to concurrently updating the weights of particles, while the other methods only use CPU. As can be seen, our method needs a smaller number of observation model evaluations to update the weights for all particles. This is because we only need to perform the network inference for all particles which are localized in the same grid cell once. For one incoming frame and the virtual frame of that grid cell, the inputs of the network and thus the outputs are the same for all particles in that cell.

We tested our method on a system equipped with an Intel i7-8700 with 3.2 GHz and an Nvidia GeForce GTX 1080 Ti with 11 GB of memory. For initializing in the large-scale map, the worst case will take around 43s to process one frame. However, after converging, the proposed method takes only 1 s on average to process one frame with 10,000 particles.

V. CONCLUSION

In this paper, we presented a novel observation model and integrated it into an MCL framework to solve the global localization problem. Our method exploits OverlapNet to estimate the overlap and yaw angle between the current frame and the virtual frames generated at each particle using the pre-built map. This allows us to successfully update the weights of particles with agreements of both location and orientation. We implemented and evaluated our approach on an urban dataset collected with a car in different seasons and provided comparisons to other existing techniques. The experiments suggest that our approach can achieve a similar performance as other approaches in global localization while obtaining a higher success rate and lower computational time.

In future work, we will test our method on more datasets with different types of LiDAR sensors. We also plan to test our method with high definition maps and want to exploit semantic information.

REFERENCES

- [1] J. Behley, M. Garbade, A. Milioto, J. Quenzel, S. Behnke, C. Stachniss, and J. Gall. SemanticKITTI: A Dataset for Semantic Scene Understanding of LiDAR Sequences. In *Proc. of the IEEE/CVF International Conf. on Computer Vision (ICCV)*, 2019.
- [2] J. Behley and C. Stachniss. Efficient Surfel-Based SLAM using 3D Laser Range Data in Urban Environments. In *Proc. of Robotics: Science and Systems (RSS)*, 2018.
- [3] X. Chen, T. Läbe, A. Milioto, T. Röhling, O. Vysotska, A. Haag, J. Behley, and C. Stachniss. OverlapNet: Loop Closing for LiDAR-based SLAM. In *Proc. of Robotics: Science and Systems (RSS)*, 2020.
- [4] X. Chen, A. Milioto, E. Palazzolo, P. Gigure, J. Behley, and C. Stachniss. SuMa++: Efficient LiDAR-based Semantic SLAM. In *Proc. of the IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, 2019.
- [5] X. Chen, H. Zhang, H. Lu, J. Xiao, Q. Qiu, and Y. Li. Robust SLAM system based on monocular vision and LiDAR for robotic urban search and rescue. In *Proc. of the IEEE Intl. Sym. on Safety, Security, and Rescue Robotics (SSRR)*, pages 41–47, 2017.
- [6] Z.J. Chong, B. Qin, T. Bandyopadhyay, M. Ang, E. Frazzoli, and D. Rus. Synthetic 2D Lidar for Precise Vehicle Localization in 3D Urban Environment. In *Proc. of the IEEE Intl. Conf. on Robotics & Automation (ICRA)*, 2013.
- [7] K.P. Cop, P.V.K. Borges, and R. Dub. Delight: An efficient descriptor for global localisation using lidar intensities. In *Proc. of the IEEE Intl. Conf. on Robotics & Automation (ICRA)*, 2018.
- [8] I.J. Cox. Blanche—an experiment in guidance and navigation of an autonomous robot vehicle. *IEEE Trans. on Robotics and Automation*, 7(2):193–204, 1991.
- [9] F. Dellaert, D. Fox, W. Burgard, and S. Thrun. Monte carlo localization for mobile robots. In *IEEE International Conference on Robotics and Automation (ICRA)*, May 1999.
- [10] R. Dube, M.G. Gollub, H. Sommer, I. Gilitschenski, R. Siegwart, C.C. Lerma, and J. Nieto. Incremental segment-based localization in 3d point clouds. *IEEE Robotics and Automation Letters (RA-L)*, 2018.
- [11] D. Fox, W. Burgard, F. Dellaert, and S. Thrun. Monte Carlo Localization: Efficient Position Estimation for Mobile Robots. *Proc. of the Conference on Advancements of Artificial Intelligence (AAAI)*, 1999(343–349), 1999.
- [12] D. Fox, S. Thrun, W. Burgard, and F. Dellaert. Particle filters for mobile robot localization. In *Sequential Monte Carlo methods in practice*, pages 401–428. 2001.
- [13] G. Grisetti, C. Stachniss, and W. Burgard. Improved Techniques for Grid Mapping with Rao-Blackwellized Particle Filters. *IEEE Trans. on Robotics (TRO)*, 23(1):34–46, 2007.
- [14] M. Hussain and J. Bethel. Project and mission planning. In C. McGlone, E. Mikhail, J. Bethel, and R. Mullen, editors, *Manual of Photogrammetry*, chapter 15.1.2.6, pages 1109–1111, 2004.
- [15] I. Andrei I.A. Barsan, S. Wang, A. Pokrovsky, and R. Urtasun. Learning to Localize Using a LiDAR Intensity Map. In *Proc. of the Second Conference on Robot Learning (CoRL)*, pages 605–616, 2018.
- [16] G. Kim, B. Park, and A. Kim. 1-day learning, 1-year localization: Long-term LiDAR localization using scan context image. *IEEE Robotics and Automation Letters (RA-L)*, 4(2):1948–1955, 2019.
- [17] J. Levinson, M. Montemerlo, and S. Thrun. Map-Based Precision Vehicle Localization in Urban Environments. In *Proc. of Robotics: Science and Systems (RSS)*, 2007.
- [18] J. Levinson and S. Thrun. Robust Vehicle Localization in Urban Environments Using Probabilistic Maps. In *Proc. of the IEEE Intl. Conf. on Robotics & Automation (ICRA)*, pages 4372–4378, 2010.
- [19] W. Ma, I. Tartavull, I. A. Bãrsan, S. Wang, M. Bai, G. Mattyus, N. Homayounfar, S. K. Lakshminathan, A. Pokrovsky, and R. Urtasun. Exploiting Sparse Semantic HD Maps for Self-Driving Vehicle Localization. In *Proc. of the IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, 2019.
- [20] Ch. Merfels and C. Stachniss. Pose fusion with chain pose graphs for automated driving. In *Proc. of the IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, 2016.
- [21] A. Milioto, I. Vizzo, J. Behley, and C. Stachniss. RangeNet++: Fast and Accurate LiDAR Semantic Segmentation. In *Proceedings of the IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, 2019.
- [22] T. Röhling, J. Mack, and D. Schulz. A Fast Histogram-Based Similarity Measure for Detecting Loop Closures in 3-D LIDAR Data. In *Proc. of the IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, pages 736–741, 2015.
- [23] A. Schaefer, D. Büscher, J. Vertens, L. Luft, and W. Burgard. Long-term urban vehicle localization using pole landmarks extracted from 3-D lidar scans. In *Proc. of the Europ. Conf. on Mobile Robotics (ECMR)*, pages 1–7, 2019.
- [24] R.C. Smith and P. Cheeseman. On the representation and estimation of spatial uncertainty. *Intl. Journal of Robotics Research (IJRR)*, 5(4):56–68, 1986.
- [25] C. Stachniss, J. Leonard, and S. Thrun. *Springer Handbook of Robotics, 2nd edition*, chapter Chapt. 46: Simultaneous Localization and Mapping. Springer Verlag, 2016.
- [26] B. Steder, M. Ruhnke, S. Grzonka, and W. Burgard. Place Recognition in 3D Scans Using a Combination of Bag of Words and Point Feature Based Relative Pose Estimation. In *Proc. of the IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, 2011.
- [27] S. Thrun, W. Burgard, and D. Fox. *Probabilistic Robotics*. MIT Press, 2005.
- [28] S. Thrun, D. Fox, W. Burgard, and F. Dellaert. Robust Monte Carlo Localization for Mobile Robots. *Artificial Intelligence*, 128(1-2), 2001.
- [29] G. Tinchev, A. Penate-Sanchez, and M. Fallon. 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(2):1327–1334, 2019.
- [30] X. Wei, I. A. Bãrsan, S. Wang, J. Martinez, and R. Urtasun. Learning to Localize Through Compressed Binary Maps. In *Proc. of the IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 10316–10324, 2019.
- [31] D. Wilbers, Ch. Merfels, and C. Stachniss. Localization with Sliding Window Factor Graphs on Third-Party Maps for Automated Driving. In *Proc. of the IEEE Intl. Conf. on Robotics & Automation (ICRA)*, 2019.
- [32] R.W. Wolcott and R.M. Eustice. Fast lidar localization using multiresolution gaussian mixture maps. In *Proc. of the IEEE Intl. Conf. on Robotics & Automation (ICRA)*, pages 2814–2821, 2015.
- [33] F. Yan, O. Vysotska, and C. Stachniss. Global Localization on OpenStreetMap Using 4-bit Semantic Descriptors. In *Proc. of the Europ. Conf. on Mobile Robotics (ECMR)*, 2019.
- [34] H. Yin, Y. Wang, X. Ding, L. Tang, S. Huang, and R. Xiong. 3D LiDAR-Based Global Localization Using Siamese Neural Network. *IEEE Trans. on Intelligent Transportation Systems (ITS)*, 2019.
- [35] C. Zhang, M. H. Ang, and D. Rus. Robust lidar localization for autonomous driving in rain. In *Proc. of the IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, pages 3409–3415, 2018.