2023 IEEE International Conference on Robotics and Automation (ICRA 2023)
May 29 - June 2, 2023. London, UK

Large-Scale Radar Localization using Online Public Maps

Ziyang Hong[1], Yvan Petillot[1], Kaicheng Zhang[1], Shida Xu[1] and Sen Wang[2]

***Abstract*— In this paper, we propose using online public maps, e.g., OpenStreetMap (OSM), for large-scale radar-based localization without needing a prior sensing map. This can potentially extend the localization system to anywhere worldwide without building, saving, or maintaining a sensing map, as long as an online public map covers the operating area. Existing methods using OSM only use route network or semantics information. These two sources of information are not combined in the previous works, while our proposed system fuses them to improve localization accuracy. Our experiments, on three open datasets collected from three different continents, show that the proposed system outperforms the state-of-the-art localization methods, reducing up to 50% of position errors. We release an open-source implementation for the community.**

I. INTRODUCTION

In mobile robotics, localization using state estimation is one of the most fundamental capabilities toward long-term autonomy. There are usually two steps to building a localization system for autonomous robots: simultaneous localization and mapping (SLAM) and map-based localization. First, in a SLAM system, the robot localizes itself while simultaneously building a consistent map of the area where the robot operates. This pre-built map is then used as prior information to localize robots in the mapped area. However, constructing, updating, and maintaining such SLAM-built prior maps, particularly in large-scale environments for long-term autonomy, are labor-intensive and expensive since a dedicated maintenance team is required to re-visit the mapping areas periodically with tedious data and map post-processing. Meanwhile, it is difficult to guarantee full coverage of the vast regions in the built maps, limiting the robots' operation areas.

To ease this burden, online public maps, e.g., OpenStreetMap (OSM), can be considered for robot localization. As Geographic Information Systems (GIS), these public maps provide rich geographic contexts, e.g., building layouts and road networks, worldwide with crowd-sourcing data contributed by millions in real-time. This ensures broad coverage of global areas and relatively up-to-date geographic information. However, utilizing online public maps for robot localization is challenging since they are mainly designed for geo-services and human navigation, lacking sensory maps for correlating using robot sensors. Moreover, there are likely discrepancies between these maps and the real environments due to recent developments and temporary changes, apart from the fact that metric scales can be distorted from the real physical world.

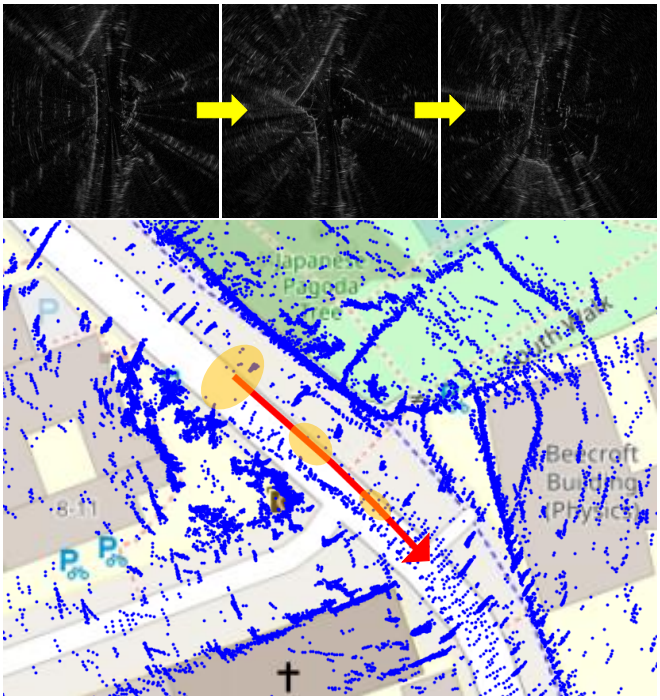

Fig. 1: An example of the localization process using radar data (top) and OpenStreetMap (bottom). The local radar map (blue) is merged using the trajectory (red) estimated by our proposed system by correlating the radar data and OSM map.

This paper explores the potential of using online public maps for large-scale robot localization combined with radar perception, which is capable of observing environmental structures. We show the feasibility of using radar data and OSM to perform accurate localization in a probabilistic manner (see Fig. 1). The main contributions are as follows:

1) To the best of our knowledge, we propose the first radar-only localization system using a public map service. It deals with map errors through a sophisticated and robust probabilistic model.
2) It achieves accurate localization performance compared to the state-of-the-art methods using OSM. We also demonstrate that it can be deployed directly on various continents and cities without needing to build a prior high-definition map using sensors.
3) We release our open-source codes[1] for the community.

[1] Ziyang Hong, Yvan Petillot, Kaicheng Zhang and Shida Xu are with Heriot-Watt University, UK. {zh9, y.r.petillot, kz13, sx2000}@hw.ac.uk

[2]Sen Wang is with the Department of Electrical and Electronic Engineering, Imperial College London, UK sen.wang@imperial.ac.uk

[1]https://github.com/perception-and-robotics-group/radarOSM

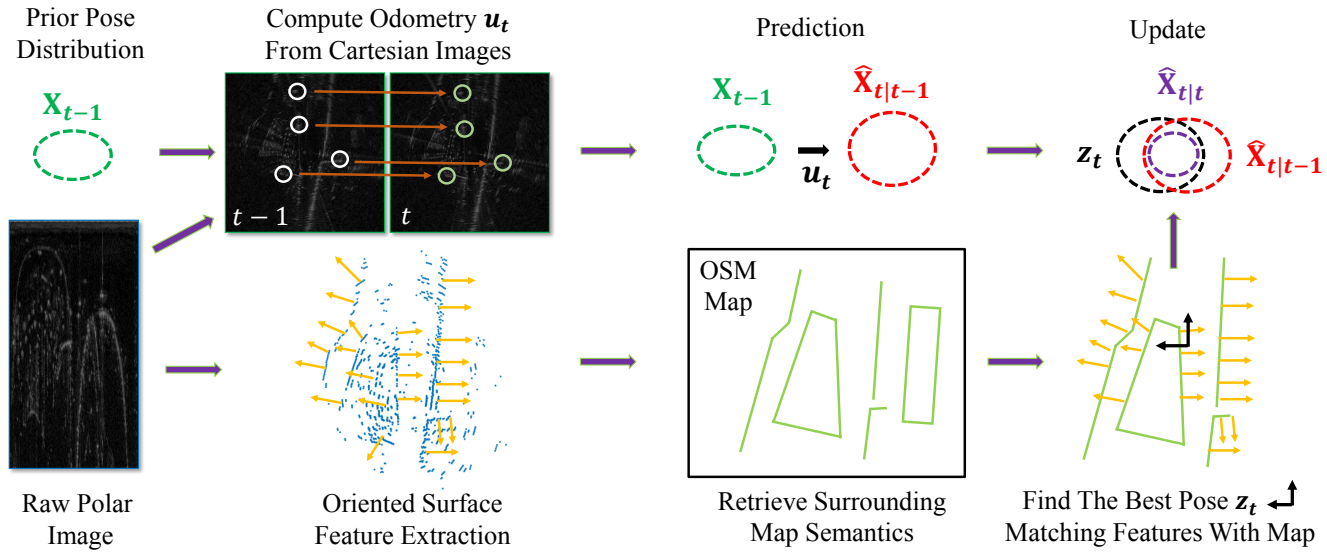

Fig. 2: System overview. The system takes the current polar radar image, and the previous state as input. The radar data is used to compute the radar odometry, then the odometry is used for prediction. Observation is found by matching features with OSM. Finally, in the update step, we calculate the posterior distribution of the pose (in purple).

II. RELATED WORK

The vision and robotics communities have extensively studied the map-based localization problem. Models derived from Bayes filters have been applied in various applications. By integrating vision or laser observation and odometry, many models perform a two-step filtering approach, i.e., prediction and update steps under the Markov assumption. Some models using Monte Carlo method use particle-based representation to approximate the posterior distribution of the robot pose [1], [2], [3] while others use Gaussian representation [4], [5], [6]. The aforementioned map-based localization methods are restricted to small and highly structured environments and require accurate 2D floor plans.

Compared to a 2D map, a high-definition (HD) map provides accurate scene geometry and information to achieve millimeter localization performance. It has attracted a lot of attention in the autonomous driving domain recently. [7], [8] leverage crow-sourced information to improve and produce HD maps. In [9], a sliding window factor graph smoothing method is proposed to fuse odometry, lidar, and map factor to improve the localization reliability. A deep network is learned in [10] to perform place recognition using a LiDAR scan. [11] proposes to build a 3D mesh map and use range-image for localization in an urban environment. However, maintaining HD maps requires high traversal frequencies, which are prohibitively expensive.

The radar sensor is a unique type of active sensor. It is becoming more popular in the mobile robotics communities and has been widely adopted into localization and mapping systems for autonomous vehicles and mobile robots. Existing algorithms using radar sensor are predominantly using it to perform odometry estimation [12], [13], [14], [15], [16], [17] and teach-repeat type of localization and place recognition [18], [19], [20]. For place recognition and teach-repeat localization systems, they must traverse the environment and collect the radar data as a database or reference sequence in advance. The current radar-based localization systems inevitably rely on prior radar maps or similar sensor modality maps like LiDAR maps.

A. Localization using Public Maps

Recently, several works based on OSM to perform localization have been proposed. The self-localization algorithm [21] is one of the earliest works which exploit OSM for large-scale localization. It uses a constant velocity model, route network in the prediction step, and visual odometry information as correction information in a Kalman-style filter. [22] proposes to use the Monte Carlo approach to represent the robot pose in an outdoor environment; laser scan and odometry information are used to compute the robot pose. Similarly, [23] proposes to combine graph-based SLAM and OSM to improve the global positioning accuracy with a laser scan. A compact feature descriptor extracted from semantic range image is proposed for fast inference in [24]. [25], [26] learns a transformation between the panorama images from Google Street View and the tiles data from OSM in the embedded space to perform global localization. By extracting building points from the LiDAR point cloud, [27] performs global localization on the OSM.

B. Localization Against Satellite Imagery

Another type of localization system without using HD maps is to perform urban localization using satellite imagery. [28] proposed to localize in satellite images from a radar mounted on a ground vehicle. They further extend the sensor modalities to both radar and LiDAR in [29], [30] so that the localization system became more universal.

III. PROPOSED SYSTEM

The proposed system uses radar as our perception modality and OSM as the prior online map. We combine both the route network and the semantics information from OSM with a probabilistic model to perform localization. The initial pose, which can be obtained using an initial GPS location or a place recognition algorithm, is assumed to be known in our formulation. The system is overviewed in Fig. 2.

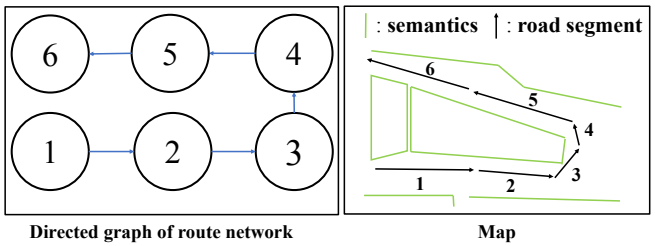

Fig. 3: Map Representation: (left) The route network is converted to a directed graph representation. (right) The map consists of two parts: the directed graph of road segments and the semantics information.

A. Representation of Public Online Maps

Our representation of an online map consists of two main elements, similar to the lane-based map representation in [21]. The first element is the route network which is encoded as a directed graph. Specifically, each road segment is parameterized by its origin, terminus, and orientation. The second element is the map's semantic information, such as building, amenity, sidewalk, man-made facility, and natural. They are represented by a set of segments or polygons, as illustrated in Fig. 3. All the map semantics are stored in an R-tree for fast indexing. However, this map representation is only a coarse approximation to the corresponding real-world environment because some parts of the area can be misaligned or outdated on the map. For instance, newly developed buildings might not be presented, and some types of objects, like trees, are not included in the map. Therefore, to associate the defective map with the radar sensing data, a sophisticated probabilistic model is proposed to perform accurate localization.

B. Radar Data and Oriented Surface Feature

The radar sensor used provides 360-degree surrounding coverage on a 2D plane. Its raw data is streamed as polar coordinated images, which can be transformed into a Cartesian representation. We first extract the point cloud from the raw polar radar power reading using a probabilistic method described in Algorithm 1. Then we detect key points from the point cloud and describe them using their normal orientation, similar to [16]. The extracted oriented surface features (see in Fig. 2) are used later in the update step to match the semantics information in the OSM map.

C. State Model

At time t, we parametrize the robot state as $\mathbf{s}_t = (l, \mathbf{X}_t^l)$ where l is the road segment ID and $\mathbf{X}_t^l = (x_t^l, y_t^l, \theta_t^l)$ denotes the robot pose with respect to a reference frame $\mathcal{F}_l$. The reference frame $\mathcal{F}_l$ is defined by the start pose $\mathbf{X}^L = (x_l, y_l, \theta_l)$ of the road segment l, as shown in Fig. 4. (x_t, y_t) is the position and θ_t is the robot orientation with respect to the frame $\mathcal{F}_l$.

Algorithm 1: Point Cloud Extraction from Radar Images

Input: Radar image $S \in \mathbb{R}^{m \times n}$;
Output: Point Cloud $P \in \mathbb{R}^{z \times 2}$;
Parameters: Minimum peak prominence δ_p and minimum peak distance δ_d;
Initialize empty point cloud set P;
for $i \leftarrow 1$ *to* m **do**
  $Q^{kx1} \leftarrow$ findPeaks($S[i,:]$, δ_p, δ_d);
  $(\mu, \sigma) \leftarrow$ meanAndStandardDeviation($Q^{k \times 1}$);
  for *each peak* q *in* Q **do**
    if $q \geq (\mu + \sigma)$ **then**
      $p \leftarrow$ transformPeakToPoint(q, i);
      Add the point p to P;
    end
  end
end

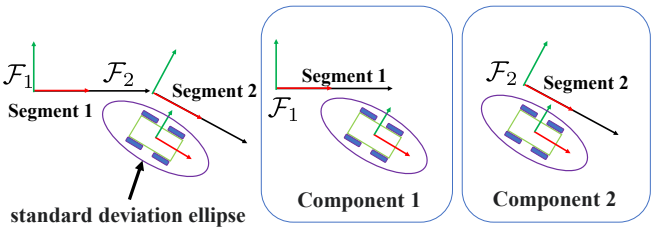

Fig. 4: Road Segments From Map With Robot Pose Components: 2 pose components are shown here.

We hypothesize the robot can be localized on multiple road segments at a time. Here L_t is the set of all the potential road segments that the robot is localizing on at time t and $l \in L_t$. We use L_M to denote a set that contains all the road segments in the map while L_t is a subset of L_M. The cardinality of L_t is far less than that of L_M. The whole probability density function over the state is parameterized by a Gaussian mixture model:

$$p(\mathbf{s}_t) = \sum_{l \in L_t} v_l \mathcal{N}(\mathbf{X}_t^l, \mathbf{P}_t^l) \tag{1}$$

where $\mathbf{P}_t^l \in \mathbb{R}^{3\times3}$ is the covariance matrix while v_l is the weight of the Gaussian component. The global pose $\mathbf{X}_t^G = (x_t^G, y_t^G, \theta_t^G)$ respected to the global UTM coordinate system can be obtained by

$$\begin{aligned} \mathbf{X}_t^G &= g(l, \mathbf{X}_t^l) \\ &= \begin{bmatrix} \cos(\theta_l)x_t^l - \sin(\theta_l)y_t^l + x_l \\ \sin(\theta_l)x_t^l + \cos(\theta_l)y_t^l + y_l \\ \theta_l + \theta_t^l \end{bmatrix} \end{aligned} \tag{2}$$

D. Prediction

From the previous time $t-1$ to current time t, the radar odometry $u_{t,t-1}$ is computed using the radar odom-

etry method in [31] which explicitly compensates motion distortion of a scanning radar. The predicted pose $\hat{\mathbf{X}}^l_{t|t-1}$ of each Gaussian component is given by compounding the previous pose $\mathbf{X}^l_{t-1}$ and the odometry which is corrupted by the process noise w_t:

$$\begin{aligned}\hat{\mathbf{X}}^l_{t|t-1} &= f(u_{t,t-1}, \mathbf{X}^l_{t-1}, w_t) \\ \hat{\mathbf{P}}^l_{t|t-1} &= \mathbf{F}_t \mathbf{P}^l_{t-1} \mathbf{F}_t^\top + \mathbf{W}_t \mathbf{Q}_t \mathbf{W}_t^\top\end{aligned} \tag{3}$$

where w_t is assumed to be zero mean Gaussian white noise with covariance $\mathbf{Q}_t$. Here $\mathbf{F}_t$ and $\mathbf{W}_t$ are the Jacobian matrices at $\mathbf{X}^l_{t-1}$.

By exploiting the nonlinear Gaussian transition model of $\hat{\mathbf{X}}^l_{t|t-1}$, we can define the road segment transition probability as

$$p(l|\mathbf{s}_{t|t-1}) = \int_{y_{low}}^{y_{up}} \int_{x_{low}}^{x_{up}} \mathcal{N}(\hat{\mathbf{X}}^l_{t|t-1}, \hat{\mathbf{P}}^l_{t|t-1}) dx dy \tag{4}$$

Here we incorporate the lane information to calculate x_{low}, x_{up}, y_{low} and y_{up}, see Fig. 5. This step is critical since we will use this transition probability to decide which road segment l should be used to draw observation samples from in the update step.

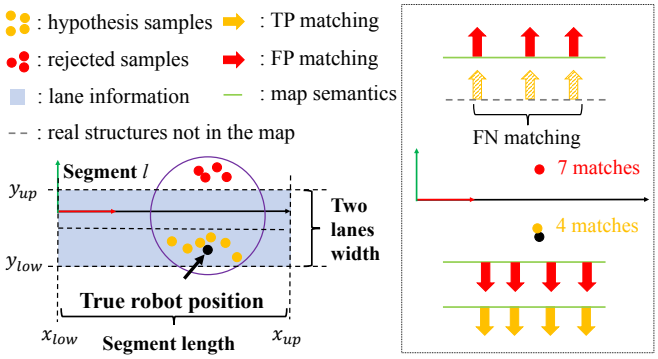

Fig. 5: Lane Information and Observation Sampling: (bottom left) lane information and the samples. (right) the red sample far away from the true position might have more feature matches (false positive matches) compared to the yellow sample that is closer to the true position.

E. Update

In the update step, we aim to develop an association between the extracted radar features and the map representation so that an observation can be obtained to correct the state's drift. Specifically, the pose observation z_t is modeled as

$$p(z_t|\mathbf{s}_t) = \mathcal{N}(z_t|g(l, \mathbf{X}^l_t), \mathbf{R}_t) \tag{5}$$

where $\mathbf{R}_t$ is the covariance matrix describing the observation uncertainty.

However, establishing the association is challenging since the radar features mainly encapsulate low-level geometric information, while the map encodes mostly high-level semantic information. Moreover, the map's semantic information might be outdated or inaccurate, causing large discrepancies between the radar and map data, as an example in Fig. 6.

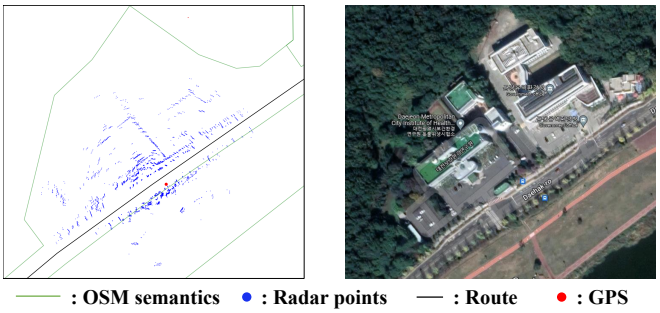

Fig. 6: Discrepancy Between Radar and Map. Left: Radar points superimposed on OSM map using ground truth GPS. Right: Satellite image of the area. We can see most structures detected on the radar are yet present on the OSM map.

To obtain the global pose observation, we propose a sampling approach. From the predicted distribution, we can quickly retrieve the surrounding map semantics information from the map. Then a set of pose samples is drawn from the predicted distribution on the top K road segments according to the road segment transition probability (4). Then, taking each pose sample as a hypothesis, the radar features are matched with the nearby semantics to determine the number of inliers based on two criteria: 1) the normal angle difference between the feature and the segment should be less than a threshold σ_θ; 2) the projection distance between the feature point and the segment should be less than a threshold σ_d. Finally, the pose sample with a maximum number of inliers is selected as a candidate for the global pose observation.

To avoid erroneous observation, we conservatively perform an update which means we might not be able to perform an update on every timestamp. Consequently, the uncertainty of the state might grow over time. To overcome this, we leverage the lane information from the map when the uncertainty is large. This helps us reject erroneous samples, as shown in Fig. 5.

To evaluate whether the candidate observation is reliable, we impose three criteria to evaluate the inlier features. Firstly, a minimum number δ_{in} of inliers is required. Secondly, the inliers must be evenly distributed over the space without being over-concentrated on local regions by inspecting their standard deviation of points' position on the lateral axis using $\delta_{spatial}$. Finally, we observe that if the robot is close to one side, we might have a few matches on the other side of the robot, which we call the "weak side". The number of matches of the weak side should be higher than a threshold δ_{weak}. For the candidate observation that satisfies all these criteria, it is considered as a sound observation z_t to perform the update:

$$\begin{aligned}\hat{\mathbf{P}}^l_{t|t} &= (\mathbf{G}_l^\top (\mathbf{R}_t)^{-1} \mathbf{G}_l + (\hat{\mathbf{P}}^l_{t|t-1})^{-1})^{-1} \\ \hat{\mathbf{X}}^l_{t|t} &= \hat{\mathbf{P}}^l_{t|t} (\mathbf{G}_l^\top (\mathbf{R}_t)^{-1} z_t + (\hat{\mathbf{P}}^l_{t|t-1})^{-1} \hat{\mathbf{X}}^l_{t|t-1})\end{aligned} \tag{6}$$

where $\mathbf{G}_l$ is the Jacobian matrix for the function $g(l, \hat{\mathbf{X}}^l_{t|t-1})$ with respect to $\hat{\mathbf{X}}^l_{t|t-1}$.

TABLE I: Average Position Error (in Meters) on the Three Datasets

	Boreas Sequences				Oxford Sequences						MulRan Sequences						
Method	01	02	03	04	01	02	03	04	05	06	K01	K02	K03	R01	R02	R03	Mean
Lost	6.4	8.1	7.7	7.1	9.2	6.7	6.3	6.8	7.8	7.9	61.8	**5.8**	**5.8**	9.2	9.6	**17.0**	11.5
Ours	**4.7**	**4.9**	**4.9**	**3.8**	**5.4**	**3.4**	**3.1**	**3.9**	**4.1**	**4.1**	**7.0**	6.0	6.2	**6.7**	**5.7**	N/A	**4.9**

Lost has a huge error in K01 since it only considers forward motion in the transition model, and there is a turnaround motion in this sequence. Our method fails to complete R03 due to outdated OSM map information available.
Boreas Sequences: 01: 11-02-11-16, 02: 11-14-09-47, 03: 11-16-14-10, 04: 11-23-14-27.
Oxford Sequences: 01: 10-11-46-21, 02: 10-12-32-52, 03: 16-13-09-37, 04: 17-13-26-39, 05: 18-14-46-59, 06: 18-15-20-12.

IV. EXPERIMENTS

To evaluate the performance, we use the public online map OSM and three public radar datasets. All the parameters of our proposed system are summarized at Table II.

TABLE II: System parameters

Parameter	Value	Note
σ_θ	30	the normal angle threshold between the feature and the segment normal (in degree)
σ_d	2	threshold of projection distance between the point and the map semantic segment (in meters)
δ_{in}	30	minimum number of inliers
δ_{weak}	10	minimum number of inliers at weakside
$\delta_{spatial}$	80	standard deviation of the feature distribution on the lateral axis (in meters)
K	2	Number of segments to sample poses

A. Public Maps and Datasets

1) OpenStreetMap: OpenStreetMap[2] (OSM) is one of the most prominent non-commercial online public map projects, which provides free access to various types of geographic data. OSM fundamentally relies on local knowledge in data collection, and geographical information might be corrected by other public maps, GPS traces, and contributors. In our system, the map consists of semantics, like "sidewalk", "amenity", "buildings", "man_made", "landuse" and "natural", extracted from the OSM.

2) Three Radar Datasets: (a) Boreas Dataset [32]. It was collected by mainly driving a repeated route over one year in Toronto and covers various seasonal and weather conditions. We select four of its sequences for evaluation. (b) Oxford Radar RobotCar Dataset [33]. It provides data from a Navtech CTS350-X radar, traversing the same route multiple times in Oxford, UK. Six sequences are used for evaluation. (c) MulRun Dataset [34]. It provides radar and LiDAR range data, covering a variety of city environments (e.g., bridge, tunnel, and overpass). Both KAIST and Riverside routes are selected for our experiment. Each route contains three traverses on different days.

B. Competing Method

A state-of-the-art method for localization using OSM is *Lost* [21]. It learns a probabilistic transition model with respect to the route network. It also uses a constant velocity model in the prediction step and the odometry information as the observation to correct drifts.

[2]www.openstreetmap.org

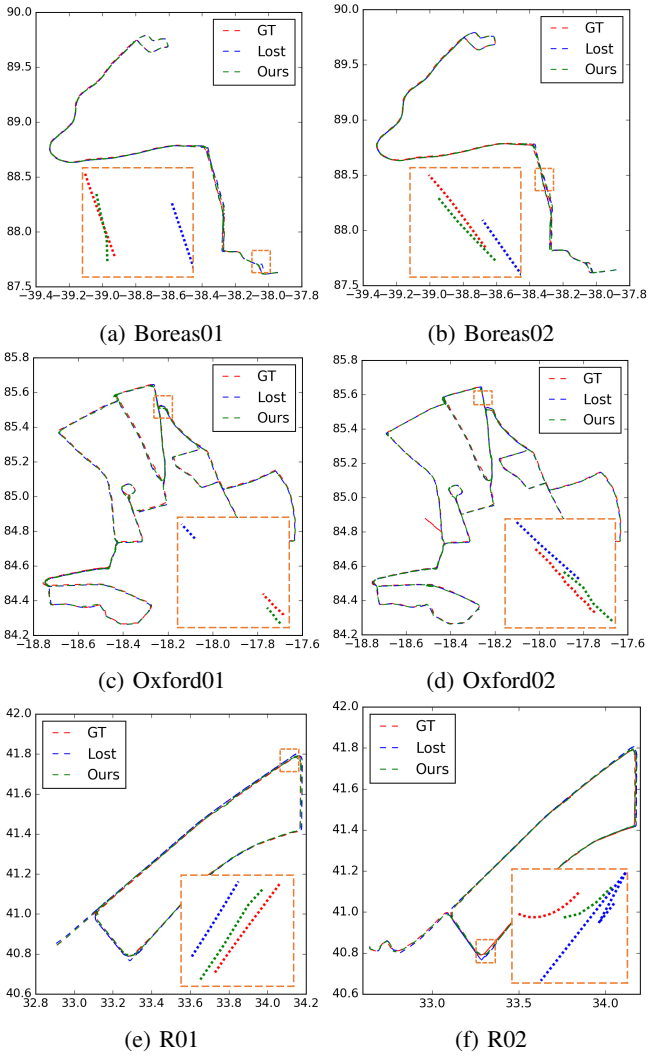

(a) Boreas01 (b) Boreas02

(c) Oxford01 (d) Oxford02

(e) R01 (f) R02

Fig. 7: Estimated Trajectories and Ground Truth.

C. Evaluation Protocol

We initialize all methods with the same initial GPS positions and the headings of the closest road segments. The GPS ground truth positions are then used for benchmarking. The mean of the absolute position errors is computed for each trajectory. For both methods, we use the same radar odometry computed by [31], which has about 2.2% drift over time as reported in [31].

D. Results and Discussion

The average position error for all datasets is shown in Table I. It can be seen that our proposed method outperforms

Lost on most of the sequences. Particularly, our method achieves better position accuracy on the sequences of the Boreas and Oxford datasets. On the MulRan dataset, we achieve degraded yet comparable performance due to the fact the MulRan is collected in Korea, where OSM is restricted. The OSM map information of the areas covered in the MulRan dataset tends to be outdated, without being updated for over ten years. As shown in Fig. 6, a large amount of building information is missing. In contrast, their maps are updated within the last three years for the Boreas and Oxford datasets.

TABLE III: Localization Errors on Boreas Dataset.

	Latitude [m]				Longitude [m]				Heading [degree]			
	Sequence				Sequence				Sequence			
Method	01	02	03	04	01	02	03	04	01	02	03	04
Lost [21]	**5.5**	8.8	7.3	7.1	5.1	5.7	5.7	5.3	1.6	2.2	**1.7**	2.2
Ours	5.8	**5.0**	**5.5**	**3.3**	**2.7**	**3.5**	**3.3**	**2.8**	1.6	2.2	2.2	**2.0**

Since the Boreas dataset provides optimized global position ground truth, we conduct a more in-depth evaluation of it. The latitude, longitude, and heading errors are given in Table III. In terms of latitude and longitude errors, our method is much smaller than Lost since Lost strictly confines its estimated positions on the road segments. Fig. 7 shows the estimated trajectories and their ground truth. In Fig. 8, we can see that the errors of our proposed system do not grow over time and are bound in reasonable ranges, while the error of Lost is sometimes over 30 meters. Surprisingly, Lost and our proposed method both achieve similar heading errors. We assume this is because all the datasets are collected using a car whose heading usually follows the lane heading when driving on the road.

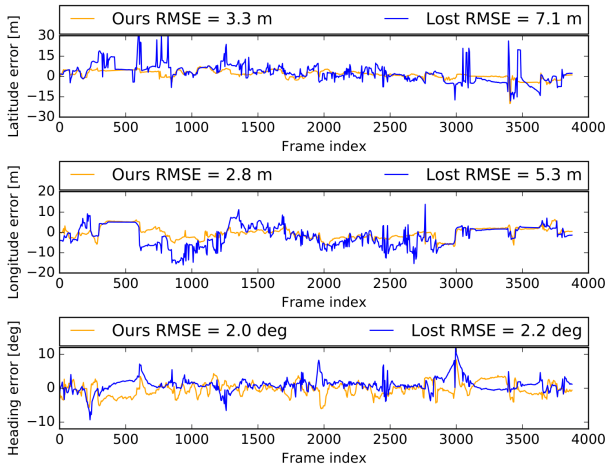

Fig. 8: Localization error on Boreas sequence 04.

Compared to Lost, our method is capable of handling various driving patterns. An example on K01 sequence is given in Fig 9. Our motion model can transit to the correct road segment while Lost estimates a wrong transition using a simple constant velocity motion model, although it uses motion odometry for the update.

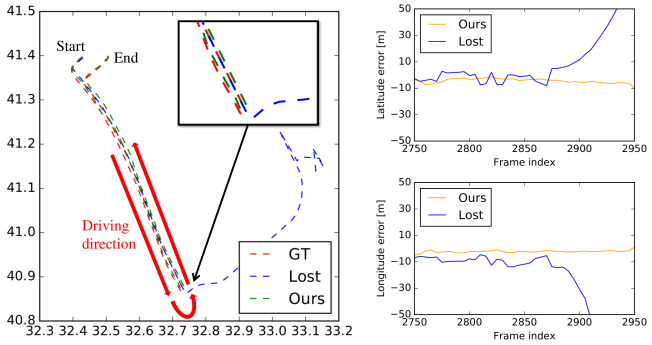

(a) Lost failure (b) Sequence error

Fig. 9: Lost Failure in K01. The motion model of Lost is not able to deal with turnaround situation. The error of Lost becomes unbounded when it transits to the wrong segment.

Failure cases: Our method fails on R03 sequence in the MulRan dataset because the map information is outdated in the OSM. Our system tries to align the radar points with the map; however, there is a large discrepancy between the real environment and the map. Our system performs a sub-optimal update. Also, in the later part of the sequence, the lane information is missing, and our model assumes it is a single-lane road while it is a 4-lane road. Therefore the perception and lane information contradict each other. Without performing update for a long period, our estimation drifts fast (see Fig. 10).

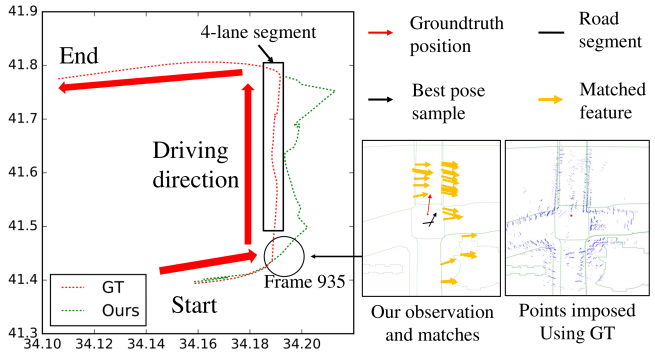

Fig. 10: Failure Case. At Frame 935, due to the large discrepancy between the map and the real radar points, our system performs an incorrect update where the observation is far from the ground truth, so the system starts to drift away.

V. CONCLUSIONS

In this paper, we propose a novel radar-based probabilistic localization system that exploits the information from online public maps. Extensive experiments show the robustness of our proposed system and the feasibility of combining radar data and OSM to perform accurate localization in large-scale environments. We also show that the proposed system outperforms the existing state-of-the-art method using OSM.

REFERENCES

[1] D. Fox, W. Burgard, F. Dellaert, and S. Thrun, "Monte carlo localization: Efficient position estimation for mobile robots," *AAAI/IAAI*, vol. 1999, no. 343-349, pp. 2–2, 1999.

[2] F. Dellaert, W. Burgard, D. Fox, and S. Thrun, "Using the condensation algorithm for robust, vision-based mobile robot localization," in *Proceedings. 1999 IEEE computer society conference on computer vision and pattern recognition (Cat. No PR00149)*, vol. 2. IEEE, 1999, pp. 588–594.

[3] S. Thrun, "Probabilistic robotics," *Communications of the ACM*, vol. 45, no. 3, pp. 52–57, 2002.

[4] J. J. Leonard and H. F. Durrant-Whyte, "Mobile robot localization by tracking geometric beacons," *IEEE Transactions on robotics and Automation*, vol. 7, no. 3, pp. 376–382, 1991.

[5] P. Jensfelt and H. I. Christensen, "Pose tracking using laser scanning and minimalistic environmental models," *IEEE Transactions on Robotics and Automation*, vol. 17, no. 2, pp. 138–147, 2001.

[6] K. O. Arras, J. A. Castellanos, and R. Siegwart, "Feature-based multi-hypothesis localization and tracking for mobile robots using geometric constraints," in *Proceedings 2002 IEEE International Conference on Robotics and Automation (Cat. No. 02CH37292)*, vol. 2. IEEE, 2002, pp. 1371–1377.

[7] G. Máttyus, S. Wang, S. Fidler, and R. Urtasun, "Hd maps: Fine-grained road segmentation by parsing ground and aerial images," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2016, pp. 3611–3619.

[8] D. Pannen, M. Liebner, W. Hempel, and W. Burgard, "How to keep hd maps for automated driving up to date," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2020, pp. 2288–2294.

[9] D. Wilbers, C. Merfels, and C. Stachniss, "Localization with sliding window factor graphs on third-party maps for automated driving," in *2019 International Conference on Robotics and Automation (ICRA)*. IEEE, 2019, pp. 5951–5957.

[10] X. Chen, T. Läbe, L. Nardi, J. Behley, and C. Stachniss, "Learning an Overlap-based Observation Model for 3D LiDAR Localization," in *Proceedings of the IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, 2020.

[11] X. Chen, I. Vizzo, T. Läbe, J. Behley, and C. Stachniss, "Range image-based lidar localization for autonomous vehicles," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 5802–5808.

[12] S. H. Cen and P. Newman, "Radar-only ego-motion estimation in difficult settings via graph matching," in *2019 International Conference on Robotics and Automation (ICRA)*. IEEE, 2019, pp. 298–304.

[13] D. Barnes, R. Weston, and I. Posner, "Masking by moving: Learning distraction-free radar odometry from pose information," in *Conference on Robot Learning (CoRL)*, 2019.

[14] Y. S. Park, Y.-S. Shin, and A. Kim, "Pharao: Direct radar odometry using phase correlation," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2020, pp. 2617–2623.

[15] P.-C. Kung, C.-C. Wang, and W.-C. Lin, "A normal distribution transform-based radar odometry designed for scanning and automotive radars," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 14 417–14 423.

[16] D. Adolfsson, M. Magnusson, A. Alhashimi, A. J. Lilienthal, and H. Andreasson, "Cfear radarodometry-conservative filtering for efficient and accurate radar odometry," in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, pp. 5462–5469.

[17] K. Burnett, A. P. Schoellig, and T. D. Barfoot, "Do we need to compensate for motion distortion and doppler effects in spinning radar navigation?" *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 771–778, 2021.

[18] Ş. Săftescu, M. Gadd, D. De Martini, D. Barnes, and P. Newman, "Kidnapped radar: Topological radar localisation using rotationally-invariant metric learning," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2020, pp. 4358–4364.

[19] H. Yin, R. Chen, Y. Wang, and R. Xiong, "Rall: end-to-end radar localization on lidar map using differentiable measurement model," *IEEE Transactions on Intelligent Transportation Systems*, 2021.

[20] K. Burnett, Y. Wu, D. J. Yoon, A. P. Schoellig, and T. D. Barfoot, "Are we ready for radar to replace lidar in all-weather mapping and localization?" *IEEE Robotics and Automation Letters*, vol. 7, no. 4, pp. 10 328–10 335, 2022.

[21] M. A. Brubaker, A. Geiger, and R. Urtasun, "Map-based Probabilistic Visual Self-Localization," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 38, no. 4, pp. 652 – 665, 2016.

[22] P. Ruchti, B. Steder, M. Ruhnke, and W. Burgard, "Localization on openstreetmap data using a 3d laser scanner," in *2015 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2015, pp. 5260–5265.

[23] O. Vysotska and C. Stachniss, "Exploiting building information from publicly available maps in graph-based slam," in *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2016, pp. 4511–4516.

[24] F. Yan, O. Vysotska, and C. Stachniss, "Global localization on openstreetmap using 4-bit semantic descriptors," in *2019 European Conference on Mobile Robots (ECMR)*. IEEE, 2019, pp. 1–7.

[25] N. Samano, M. Zhou, and A. Calway, "You are here: Geolocation by embedding maps and images," in *European Conference on Computer Vision*. Springer, 2020, pp. 502–518.

[26] M. Zhou, X. Chen, N. Samano, C. Stachniss, and A. Calway, "Efficient localisation using images and openstreetmaps," in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2021, pp. 5507–5513.

[27] Y. Cho, G. Kim, S. Lee, and J.-H. Ryu, "Openstreetmap-based lidar global localization in urban environment without a prior lidar map," *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 4999–5006, 2022.

[28] T. Y. Tang, D. De Martini, D. Barnes, and P. Newman, "Rsl-net: Localising in satellite images from a radar on the ground," *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 1087–1094, 2020.

[29] T. Y. Tang, D. De Martini, S. Wu, and P. Newman, "Self-supervised localisation between range sensors and overhead imagery," *arXiv preprint arXiv:2006.02108*, 2020.

[30] ——, "Self-supervised learning for using overhead imagery as maps in outdoor range sensor localization," *The International Journal of Robotics Research*, vol. 40, no. 12-14, pp. 1488–1509, 2021.

[31] Z. Hong, Y. Petillot, A. Wallace, and S. Wang, "Radarslam: A robust simultaneous localization and mapping system for all weather conditions," *The International Journal of Robotics Research*, vol. 41, no. 5, pp. 519–542, 2022.

[32] K. Burnett, D. J. Yoon, Y. Wu, A. Z. Li, H. Zhang, S. Lu, J. Qian, W.-K. Tseng, A. Lambert, K. Y. Leung, A. P. Schoellig, and T. D. Barfoot, "Boreas: A multi-season autonomous driving dataset," *arXiv preprint arXiv:2203.10168*, 2022.

[33] D. Barnes, M. Gadd, P. Murcutt, P. Newman, and I. Posner, "The oxford radar robotcar dataset: A radar extension to the oxford robotcar dataset," in *IEEE International Conference on Robotics and Automation (ICRA)*. Paris: IEEE, 2020, pp. 6433–6438.

[34] G. Kim, Y. S. Park, Y. Cho, J. Jeong, and A. Kim, "Mulran: Multimodal range dataset for urban place recognition," in *IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2020, pp. 6246–6253.